

## Makefile

Page 1/1

```

1: # fiwix/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: TOPDIR := $(shell if [ "$$PWD" != "" ] ; then echo $$PWD ; else pwd ; fi)
8: INCLUDE = $(TOPDIR)/include
9:
10: ARCH = -m32
11: CPU = -march=i386
12:
13: CC = $(CROSS_COMPILE)gcc $(ARCH) $(CPU) -D__KERNEL__ #-D__DEBUG__
14: LD = $(CROSS_COMPILE)ld
15:
16: CFLAGS = -I$(INCLUDE) -O2 -ffreestanding -Wall -Wstrict-prototypes #-Wextra
17: LDFLAGS = -m elf_i386 -nostartfiles -nostdlib -nodefaultlibs -nostdinc
18:
19: DIRS = kernel kernel/syscalls mm fs drivers/block drivers/char lib
20: OBJS = kernel/kernel.o kernel/syscalls/syscalls.o mm/mm.o fs/fs.o \
21:       drivers/block/block.o drivers/char/char.o lib/lib.o
22:
23: export CC LD CFLAGS LDFLAGS INCLUDE
24:
25: all:
26:     @echo "#define UTS_VERSION \" `date` \"" > include/fiwix/version.h
27:     @for n in $(DIRS) ; do (cd $$n ; $(MAKE)) ; done
28:     $(LD) -N -T fiwix.ld $(LDFLAGS) $(OBJS) -o fiwix
29:     nm fiwix | sort | gzip -9c > System.map.gz
30:
31: clean:
32:     @for n in $(DIRS) ; do (cd $$n ; $(MAKE) clean) ; done
33:     rm -f *.o fiwix System.map.gz
34:
35: floppy:
36:     mkfs.ext2 -m 0 /dev/fd0
37:     mount -t ext2 /dev/fd0 /mnt/floppy
38:     @mkdir -p /mnt/floppy/boot/grub
39:     @echo "(fd0) /dev/fd0" > /mnt/floppy/boot/grub/device.map
40:     @grub-install --root-directory=/mnt/floppy /dev/fd0
41:     @tools/MAKEBOOTDISK
42:     @cp -prf tools/etc/* /mnt/floppy/etc
43:     @cp fiwix /mnt/floppy/boot
44:     @cp System.map.gz /mnt/floppy/boot
45:     @cp tools/install.sh /mnt/floppy/sbin
46:     umount /mnt/floppy
47:
48: floppy_update:
49:     mount -t ext2 /dev/fd0 /mnt/floppy
50:     cp fiwix /mnt/floppy/boot
51:     cp System.map.gz /mnt/floppy/boot
52:     umount /mnt/floppy
53:
54:

```

**fiwix.ld**

Page 1/1

```
1: /*
2:  * Linker script for the Fiwix kernel (3GB user / 1GB kernel).
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: OUTPUT_FORMAT("elf32-i386")
9: OUTPUT_ARCH(i386)
10: ENTRY(start)          /* entry point */
11: vaddr = 0xC0000000;    /* virtual base address at 3GB */
12: paddr = 0x100000;     /* physical address at 1MB */
13:
14: /* define output sections */
15: SECTIONS
16: {
17:     . = paddr;
18:
19:     /* kernel setup code */
20:     .setup ALIGN(4096) :
21:     {
22:         *(.setup)
23:     }
24:
25:     . += vaddr;
26:
27:     /* kernel code */
28:     .text : AT(ADDR(.text) - vaddr)
29:     {
30:         *(.text)
31:     }
32:     _etext = .;
33:
34:     /* initialized data */
35:     .data ALIGN(4096) : AT(ADDR(.data) - vaddr)
36:     {
37:         *(.data)
38:         *(.rodata*)
39:     }
40:     _edata = .;
41:
42:     /* uninitialized data */
43:     .bss ALIGN(4096) : AT(ADDR(.bss) - vaddr)
44:     {
45:         *(COMMON*)
46:         *(.bss*)
47:     }
48:     _end = .;
49:
50:     /* remove information not needed */
51:     /DISCARD/ :
52:     {
53:         *(.eh_frame)
54:     }
55: }
```

## kernel/boot.S

```

1: /*
2:  * fiwix/kernel/boot.S
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/const.h>
9:
10: #define ASM      1          /* GRUB stuff */
11: #include <fiwix/multiboot.h>
12:
13: #define CRO_MP   0x00000002 /* CRO bit-01 MP (Monitor Coprocessor) */
14: #define CRO_NE   0x00000020 /* CRO bit-05 NE (Numeric Error) */
15: #define CRO_WP   0x00010000 /* CRO bit-16 WP (Write Protect) */
16: #define CRO_AM   0x00040000 /* CRO bit-18 AM (Alignment Mask) */
17: #define CRO_PG   0x80000000 /* CRO bit-31 PG (Paging) */
18:
19: .section .setup, "a"      /* "a" attribute means Allocatable section */
20:
21: .align 4
22: tmp_gdtr:
23:     .word ((3 * 8) - 1)
24:     .long tmp_gdt
25:
26: .align 4
27: tmp_gdt:
28:     /* NULL DESCRIPTOR */
29:     .word 0x0000
30:     .word 0x0000
31:     .word 0x0000
32:     .word 0x0000
33:
34:     /* KERNEL CODE */
35:     .word 0xFFFF          /* segment limit 15-00 */
36:     .word 0x0000          /* base address 15-00 */
37:     .byte 0x00            /* base address 23-16 */
38:     .byte 0x9A            /* P=1 DPL=00 S=1 TYPE=1010 (exec/read) */
39:     .byte 0xCF            /* G=1 DB=1 0=0 AVL=0 SEGLIM=1111 */
40:     .byte 0x40            /* base address 31-24 */
41:
42:     /* KERNEL DATA */
43:     .word 0xFFFF          /* segment limit 15-00 */
44:     .word 0x0000          /* base address 15-00 */
45:     .byte 0x00            /* base address 23-16 */
46:     .byte 0x92            /* P=1 DPL=00 S=1 TYPE=0010 (read/write) */
47:     .byte 0xCF            /* G=1 DB=1 0=0 AVL=0 SEGLIM=1111 */
48:     .byte 0x40            /* base address 31-24 */
49:
50:
51: .text
52:
53: .globl start; start:
54:     cli
55:     jmp     multiboot_entry
56:
57: .align 4
58: multiboot_header:          /* multiboot header */
59:     .long MULTIBOOT_HEADER_MAGIC /* magic */
60:     .long MULTIBOOT_HEADER_FLAGS /* flags */
61:     /* checksum */
62:     .long -(MULTIBOOT_HEADER_MAGIC + MULTIBOOT_HEADER_FLAGS)
63:
64: #ifndef __ELF__
65:     .long multiboot_header      /* header_addr */
66:     .long _start                /* load_addr */
67:     .long _edata                /* load_end_addr */

```

## kernel/boot.S

Page 2/3

```

68:      .long   _end                /* bss_end_addr */
69:      .long   multiboot_entry     /* entry_addr */
70: #endif /* ! __ELF__ */
71:
72: /*
73:  * We use the CX register in order to keep intact the values in AX and BX
74:  * registers, since they are holding the Multiboot values 'magic' and 'info'
75:  * respectively.
76:  */
77: multiboot_entry:
78:      lgdt    tmp_gdtr            /* load GDTR with the temporary GDT */
79:      movw   $KERNEL_DS, %cx
80:      movw   %cx, %ds
81:      movw   %cx, %es
82:      movw   %cx, %fs
83:      movw   %cx, %gs
84:      movw   %cx, %ss
85:      ljmp   $KERNEL_CS, $1f
86: 1:
87:
88:
89: /*
90:  * WARNING WARNING WARNING WARNING WARNING WARNING WARNING WARNING WARNING !!!
91:  * -----
92:  * The minimal page directory of 4MB only works if the in-memory size of the
93:  * kernel is lesser than 3MB. If you need more space go to the setup_minmem()
94:  * function and set the 'mb4' variable accordingly.
95:  *
96:  * In order to know the current size of the Fiwix kernel, just follow this:
97:  *
98:  * # readelf -l fiwix
99:  * Elf file type is EXEC (Executable file)
100:  * Entry point 0xc0100020
101:  * There are 2 program headers, starting at offset 52
102:  *
103:  * Program Headers:
104:  *   Type           Offset   VirtAddr   PhysAddr   FileSiz MemSiz  Flg Align
105:  *   LOAD           0x000074 0x00100000 0x00100000 0x00020 0x00020 R  0x4
106:  *   LOAD           0x0000a0 0xc0100020 0x00100020 0x33f8c 0x859a0 RWE 0x20
107:  *
108:  *                   check this value --> ^^^^^^^
109:  */
109:      movl   $0xC0010000, %esp    /* default stack address */
110:      pushl  $0                  /* reset EFLAGS */
111:      popf
112:
113:      pushl  %eax                /* save Multiboot magic value */
114:      call   setup_minmem        /* setup a minimal page directory */
115:      movl   %eax, %cr3
116:
117:      movl   %cr0, %eax
118:      andl   $0x00000011, %eax   /* disable all, preserve ET & PE (GRUB)
*/
119:      orl   $CR0_PG, %eax        /* enable PG (Paging) */
120:      orl   $CR0_AM, %eax        /* enable AM (Alignment Mask) */
121:      orl   $CR0_WP, %eax        /* enable WP (Write Protect) */
122:      orl   $CR0_NE, %eax        /* enable NE (Numeric Error) */
123:      orl   $CR0_MP, %eax        /* enable MP (Monitor Coprocessor) */
124:      movl   %eax, %cr0
125:
126:      call   bss_init            /* initialize BSS segment */
127:      call   gdt_init            /* setup and load the definitive GDT */
128:
129:      pushl  %ebx                /* save Multiboot info structure */
130:      call   get_last_boot_addr
131:      add    $4, %esp
132:      popl   %ecx                /* restore Multiboot magic value */
133:      andl   $0xFFFFF000, %eax   /* page aligned */

```

**kernel/boot.S**

Page 3/3

```
134:      addl    $0x3000, %eax      /* 2 whole pages for kernel stack */
135:      subl    $4, %eax
136:      movl    %eax, %esp        /* set kernel stack */
137:
138:      pushl   %esp              /* save kernel stack address */
139:      pushl   %ebx              /* save Multiboot info structure */
140:      pushl   %ecx              /* save Multiboot magic value */
141:      call   start_kernel
142:
143:      .align 4
144:      .globl cpu_idle; cpu_idle:
145:          hlt
146:          jmp     cpu_idle
147:
148:      .align 4
149:      .org 0x1000
150:      .globl _fdc_transfer_area
151:      _fdc_transfer_area:      .fill 512*2*18,1,0
```

## kernel/cmos.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/cmos.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/cmos.h>
10:
11: int cmos_update_in_progress(void)
12: {
13:     return(cmos_read(CMOS_STATA) & CMOS_STATA_UIP);
14: }
15:
16: unsigned char cmos_read_date(unsigned char addr)
17: {
18:     /* make sure an update isn't in progress */
19:     while(cmos_update_in_progress());
20:
21:     if(!(cmos_read(CMOS_STATB) & CMOS_STATB_DM) {
22:         return BCD2BIN(cmos_read(addr));
23:     }
24:     return cmos_read(addr);
25: }
26:
27: void cmos_write_date(unsigned char addr, unsigned char value)
28: {
29:     /* make sure an update isn't in progress */
30:     while(cmos_update_in_progress());
31:
32:     if(!(cmos_read(CMOS_STATB) & CMOS_STATB_DM) {
33:         cmos_write(addr, BIN2BCD(value));
34:     }
35:     cmos_write(addr, value);
36: }
37:
38: unsigned char cmos_read(unsigned char addr)
39: {
40:     unsigned long int flags;
41:
42:     SAVE_FLAGS(flags); CLI();
43:     outport_b(CMOS_INDEX, addr);
44:     RESTORE_FLAGS(flags);
45:
46:     return inport_b(CMOS_DATA);
47: }
48:
49: void cmos_write(unsigned char addr, unsigned char value)
50: {
51:     unsigned long int flags;
52:
53:     SAVE_FLAGS(flags); CLI();
54:     outport_b(CMOS_INDEX, addr);
55:     outport_b(CMOS_DATA, value);
56:     RESTORE_FLAGS(flags);
57: }
```

## kernel/core386.S

Page 1/12

```

1: /*
2:  * fiwix/kernel/core386.S
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/const.h>
9: #include <fiwix/unistd.h>
10:
11: #define CRO_MP    ~ (0x00000002)    /* CR0 bit-01 MP (Monitor Coprocessor) */
12: #define CRO_EM    0x00000004      /* CR0 bit-02 EM (Emulation) */
13:
14: #define SS_RPL3    0x03           /* Request Privilege Level 3 */
15:
16: #define GS         0x00
17: #define FS         0x04
18: #define ES         0x08
19: #define DS         0x0C
20: #define EDI        0x10          /* \
21: #define ESI        0x14          /* |
22: #define EBP        0x18          /* |
23: #define ESP        0x1C          /* | saved by
24: #define EBX        0x20          /* | 'pusha'
25: #define EDX        0x24          /* |
26: #define ECX        0x28          /* |
27: #define EAX        0x2C          /* /
28: #define ERR        0x30          /* error code or padding
29: #define EIP        0x34          /* \
30: #define CS         0x38          /* | saved by processor
31: #define FLAGS      0x3C          /* /
32: #define OLDESP    0x40          /* \ saved by processor on
33: #define OLDSS     0x44          /* / privilege level change
34:
35: #define SAVE_ALL
36:     pushal                      \
37:     pushl   %ds                  ;\
38:     pushl   %es                  ;\
39:     pushl   %fs                  ;\
40:     pushl   %gs                  ;\
41:
42: #define EXCEPTION(exception)
43:     pushl   $exception           \
44:     call    trap_handler         ;\
45:     addl    $4, %esp             ;\
46:
47: #define IRQ(irq)
48:     pushl   $irq                 \
49:     call    irq_handler         ;\
50:     addl    $4, %esp             ;\
51:
52: #define BOTTOM_HALVES
53:     cmpl   $0, intr_count        \
54:     jne    2f                    ;\
55:     incl   intr_count            ;\
56:     call   do_bh                 ;\
57:     decl   intr_count
58:
59: #define CHECK_SIGNALS
60:     call   issig                  \
61:     cmpl   $0, %eax              ;\
62:     je     1f                    ;\
63:     movl   %esp, %eax            ;\
64:     pushl   %eax                  ;\
65:     call   psig                  ;\
66:     addl   $4, %esp              ;\
67: 1:

```

## kernel/core386.S

Page 2/12

```

68:
69: #define SCHEDULE                                     \
70:     cmpl    $0, need_resched                       ;\
71:     je      2f                                     ;\
72:     call    do_sched                               ;\
73: 2:
74:
75: #define RESTORE_ALL                                 \
76:     popl    %gs                                    ;\
77:     popl    %fs                                    ;\
78:     popl    %es                                    ;\
79:     popl    %ds                                    ;\
80:     popal
81:     addl    $4, %esp                               /* suppress error code from stack */
82:
83:
84: .text
85:
86: .align 4
87: .globl except0; except0:                          # DIVIDE ERROR
88:     pushl   $0                                    # save simulated error code to stack
89:     SAVE_ALL
90:     EXCEPTION(0x0)
91:     BOTTOM_HALVES
92:     CHECK_SIGNALS
93:     SCHEDULE
94:     RESTORE_ALL
95:     iret
96:
97: .align 4
98: .globl except1; except1:                          # DEBUG
99:     pushl   $0                                    # save simulated error code to stack
100:    SAVE_ALL
101:    EXCEPTION(0x1)
102:    BOTTOM_HALVES
103:    CHECK_SIGNALS
104:    SCHEDULE
105:    RESTORE_ALL
106:    iret
107:
108: .align 4
109: .globl except2; except2:                          # NMI INTERRUPT
110:    pushl   $0                                    # save simulated error code to stack
111:    SAVE_ALL
112:    EXCEPTION(0x2)
113:    BOTTOM_HALVES
114:    CHECK_SIGNALS
115:    SCHEDULE
116:    RESTORE_ALL
117:    iret
118:
119: .align 4
120: .globl except3; except3:                          # BREAKPOINT INT3
121:    pushl   $0                                    # save simulated error code to stack
122:    SAVE_ALL
123:    EXCEPTION(0x3)
124:    BOTTOM_HALVES
125:    CHECK_SIGNALS
126:    SCHEDULE
127:    RESTORE_ALL
128:    iret
129:
130: .align 4
131: .globl except4; except4:                          # OVERFLOW
132:    pushl   $0                                    # save simulated error code to stack
133:    SAVE_ALL
134:    EXCEPTION(0x4)

```



**kernel/core386.S**

Page 3/12

```

135:         BOTTOM_HALVES
136:         CHECK_SIGNALS
137:         SCHEDULE
138:         RESTORE_ALL
139:         iret
140:
141: .align 4
142: .globl except5; except5:         # BOUND
143:         pushl    $0             # save simulated error code to stack
144:         SAVE_ALL
145:         EXCEPTION(0x5)
146:         BOTTOM_HALVES
147:         CHECK_SIGNALS
148:         SCHEDULE
149:         RESTORE_ALL
150:         iret
151:
152: .align 4
153: .globl except6; except6:         # INVALID OPCODE
154:         pushl    $0             # save simulated error code to stack
155:         SAVE_ALL
156:         EXCEPTION(0x6)
157:         BOTTOM_HALVES
158:         CHECK_SIGNALS
159:         SCHEDULE
160:         RESTORE_ALL
161:         iret
162:
163: .align 4
164: .globl except7; except7:         # NO MATH COPROCESSOR
165:         pushl    $0             # save simulated error code to stack
166:         SAVE_ALL
167:         EXCEPTION(0x7)
168:         clts                    # floating-opcode cached!
169:         BOTTOM_HALVES
170:         CHECK_SIGNALS
171:         SCHEDULE
172:         RESTORE_ALL
173:         iret
174:
175: .align 4
176: .globl except8; except8:         # DOUBLE FAULT
177:         SAVE_ALL
178:         EXCEPTION(0x8)
179:         BOTTOM_HALVES
180:         CHECK_SIGNALS
181:         SCHEDULE
182:         RESTORE_ALL
183:         iret
184:
185: .align 4
186: .globl except9; except9:         # COPROCESSOR SEGMENT OVERRUN
187:         pushl    $0             # save simulated error code to stack
188:         SAVE_ALL
189:         EXCEPTION(0x9)
190:         BOTTOM_HALVES
191:         CHECK_SIGNALS
192:         SCHEDULE
193:         RESTORE_ALL
194:         iret
195:
196: .align 4
197: .globl exceptA; exceptA:         # INVALID TSS
198:         SAVE_ALL
199:         EXCEPTION(0xA)
200:         BOTTOM_HALVES
201:         CHECK_SIGNALS

```

**kernel/core386.S**

Page 4/12

```

202:          SCHEDULE
203:          RESTORE_ALL
204:          iret
205:
206: .align 4
207: .globl exceptB; exceptB:          # SEGMENT NOT PRESENT
208:          SAVE_ALL
209:          EXCEPTION(0xB)
210:          BOTTOM_HALVES
211:          CHECK_SIGNALS
212:          SCHEDULE
213:          RESTORE_ALL
214:          iret
215:
216: .align 4
217: .globl exceptC; exceptC:          # STACK SEGMENT FAULT
218:          SAVE_ALL
219:          EXCEPTION(0xC)
220:          BOTTOM_HALVES
221:          CHECK_SIGNALS
222:          SCHEDULE
223:          RESTORE_ALL
224:          iret
225:
226: .align 4
227: .globl exceptD; exceptD:          # GENERAL PROTECTION FAULT
228:          SAVE_ALL
229:          EXCEPTION(0xD)
230:          BOTTOM_HALVES
231:          CHECK_SIGNALS
232:          SCHEDULE
233:          RESTORE_ALL
234:          iret
235:
236: .align 4
237: .globl exceptE; exceptE:          # PAGE FAULT
238:          SAVE_ALL
239:          EXCEPTION(0xE)
240:          BOTTOM_HALVES
241:          CHECK_SIGNALS
242:          SCHEDULE
243:          RESTORE_ALL
244:          iret
245:
246: .align 4
247: .globl exceptF; exceptF:          # INTEL RESERVED
248:          pushl  $0                # save simulated error code to stack
249:          SAVE_ALL
250:          EXCEPTION(0xF)
251:          BOTTOM_HALVES
252:          CHECK_SIGNALS
253:          SCHEDULE
254:          RESTORE_ALL
255:          iret
256:
257: .globl except10; except10:        # FLOATING POINT ERROR
258:          pushl  $0                # save simulated error code to stack
259:          SAVE_ALL
260:          EXCEPTION(0x10)
261:          BOTTOM_HALVES
262:          CHECK_SIGNALS
263:          SCHEDULE
264:          RESTORE_ALL
265:          iret
266:
267: .globl except11; except11:        # ALIGNMENT CHECK
268:          EXCEPTION(0x11)

```

## kernel/core386.S

Page 5/12

```

269:         SAVE_ALL
270:         BOTTOM_HALVES
271:         CHECK_SIGNALS
272:         SCHEDULE
273:         RESTORE_ALL
274:         iret
275:
276: .globl except12; except12:      # MACHINE CHECK
277:         pushl   $0              # save simulated error code to stack
278:         SAVE_ALL
279:         EXCEPTION(0x12)
280:         BOTTOM_HALVES
281:         CHECK_SIGNALS
282:         SCHEDULE
283:         RESTORE_ALL
284:         iret
285:
286: .globl except13; except13:      # SIMD FLOATING POINT
287:         pushl   $0              # save simulated error code to stack
288:         SAVE_ALL
289:         EXCEPTION(0x13)
290:         BOTTOM_HALVES
291:         CHECK_SIGNALS
292:         SCHEDULE
293:         RESTORE_ALL
294:         iret
295:
296: .globl except14; except14:      # INTEL RESERVED
297: .globl except15; except15:      # INTEL RESERVED
298: .globl except16; except16:      # INTEL RESERVED
299: .globl except17; except17:      # INTEL RESERVED
300: .globl except18; except18:      # INTEL RESERVED
301: .globl except19; except19:      # INTEL RESERVED
302: .globl except1A; except1A:      # INTEL RESERVED
303: .globl except1B; except1B:      # INTEL RESERVED
304: .globl except1C; except1C:      # INTEL RESERVED
305: .globl except1D; except1D:      # INTEL RESERVED
306: .globl except1E; except1E:      # INTEL RESERVED
307: .globl except1F; except1F:      # INTEL RESERVED
308:         pushl   $0              # save simulated error code to stack
309:         SAVE_ALL
310:         EXCEPTION(0x14)
311:         BOTTOM_HALVES
312:         CHECK_SIGNALS
313:         SCHEDULE
314:         RESTORE_ALL
315:         iret
316:
317:
318: #define BUILD_IRQ(num, name)      \
319: .align 4                          ;\
320: .globl name; name:                ;\
321:         pushl   $0              /* save simulated error code to stack */;\
322:         SAVE_ALL                  ;\
323:         IRQ(num)                  ;\
324:         BOTTOM_HALVES             ;\
325:         CHECK_SIGNALS             ;\
326:         SCHEDULE                   ;\
327:         RESTORE_ALL               ;\
328:         iret                      ;\
329:
330: BUILD_IRQ(0, irq0)
331: BUILD_IRQ(1, irq1)
332: BUILD_IRQ(2, irq2)
333: BUILD_IRQ(3, irq3)
334: BUILD_IRQ(4, irq4)
335: BUILD_IRQ(5, irq5)

```

## kernel/core386.S

Page 6/12

```

336: BUILD_IRQ(6, irq6)
337: BUILD_IRQ(7, irq7)
338: BUILD_IRQ(8, irq8)
339: BUILD_IRQ(9, irq9)
340: BUILD_IRQ(10, irq10)
341: BUILD_IRQ(11, irq11)
342: BUILD_IRQ(12, irq12)
343: BUILD_IRQ(13, irq13)
344: BUILD_IRQ(14, irq14)
345: BUILD_IRQ(15, irq15)
346:
347: .align 4
348: .globl unknown_irq; unknown_irq:
349:     pushl    $0                # save simulated error code to stack
350:     SAVE_ALL
351:     IRQ(-1)
352:     RESTORE_ALL
353:     iret
354:
355: .align 4
356: .globl switch_to_user_mode; switch_to_user_mode:
357:     cli
358:     xorl    %eax, %eax          # initialize %eax
359:     movl    %eax, %ebx          # initialize %ebx
360:     movl    %eax, %ecx          # initialize %ecx
361:     movl    %eax, %edx          # initialize %edx
362:     movl    %eax, %esi          # initialize %esi
363:     movl    %eax, %edi          # initialize %edi
364:     movl    %eax, %ebp          # initialize %ebp
365:     movl    $(USER_DS | SS_RPL3), %eax
366:     movw    %ax, %ds
367:     movw    %ax, %es
368:     movw    %ax, %fs
369:     movw    %ax, %gs
370:     pushl   %eax
371:     pushl   $KERNEL_BASE_ADDR - 4 # user stack address
372:     pushl   $0x202                # initialize eflags (Linux 2.2 = 0x292)
373:     popfl
374:     pushfl
375:     movl    $(USER_CS | SS_RPL3), %eax
376:     pushl   %eax
377:     pushl   $KERNEL_BASE_ADDR - 0x1000 # go to init_trampoline() in use
r mode
378:     iret
379:
380: .align 4
381: .globl sighandler_trampoline; sighandler_trampoline:
382:     /*
383:     * This pushes twice the %eax register in order to save the 'signum'
384:     * value from being clobbered by functions like printf(), during the
385:     * signal handler calling.
386:     * FIXME: is this a bug in Newlib or in my side?
387:     */
388:     pushl   %eax                # save 'signum'
389:     pushl   %eax                # 'signum' is the argument of the handle
r
390:     call    *%ecx
391:     addl    $4, %esp            # discard the first push
392:     popl    %ebx                # restore 'signum'
393:
394:     movl    $SYS_sigreturn, %eax
395:     int    $0x80
396:
397:     # never reached, otherwise call sys_exit()
398:     movl    $SYS_exit, %eax
399:     int    $0x80
400:     ret

```

## kernel/core386.S

Page 7/12

```

401: .align 4
402: .globl end_sighandler_trampoline; end_sighandler_trampoline:
403:     nop
404:
405: .align 4
406: .globl syscall; syscall:           # SYSTEM CALL ENTRY
407:     pushl   %eax                   # save the system call number
408:     SAVE_ALL
409:
410:     pushl   %edi                   # \ 5th argument
411:     pushl   %esi                   # | 4th argument
412:     pushl   %edx                   # | 3rd argument
413:     pushl   %ecx                   # | 2nd argument
414:     pushl   %ebx                   # / 1st argument
415:     pushl   %eax                   # system call number
416:     call    do_syscall
417:     addl    $24, %esp              # suppress all 6 pushl from the stack
418:     movl    %eax, EAX(%esp)        # save the return value
419:
420:     BOTTOM_HALVES
421:     CHECK_SIGNALS
422:     SCHEDULE
423: .align 4
424: .globl return_from_syscall; return_from_syscall:
425:     RESTORE_ALL
426:     iret
427:
428: .align 4
429: .globl do_switch; do_switch:
430:     movl    %esp, %ebx
431:     pushal
432:     pushfl
433:     movl    0x4(%ebx), %eax        # save ESP to 'prev->tss.esp'
434:     movl    %esp, (%eax)
435:     movl    0x8(%ebx), %eax        # save EIP to 'prev->tss.eip'
436:     movl    $1f, (%eax)
437:     movl    0xC(%ebx), %esp        # load 'next->tss.esp' into ESP
438:     pushl   0x10(%ebx)             # push 'next->tss.eip' into ESP
439:     movl    0x14(%ebx), %eax      # load 'next->tss.cr3' into CR3
440:     ltr    0x18(%ebx)             # load TSS
441:     movl    %eax, %cr3
442:     ret
443: 1:
444:     popfl
445:     popal
446:
447: .align 4
448: .globl cpuid; cpuid:
449:     pushl   %ebp
450:     movl    %esp, %ebp
451:     pushl   %edi
452:     pushl   %esi
453:     pushl   %ebx
454:
455:     pushf
456:     pop     %eax                   # put original EFLAGS in EAX
457:     mov     %eax, %ecx             # save original EFLAGS in ECX
458:     xor     $0x200000, %eax        # change bit 21 (ID) in EFLAGS
459:     push   %eax                   # save new EFLAGS on stack
460:     popf
461:     pushf
462:     pop     %eax                   # put EFLAGS in EAX
463:     cmp     %ecx, %eax             # compare if both EFLAGS are equal
464:
465:     je     test386                 # can't toggle ID bit, no CPUID
466:     xor     %ebx, %ebx             # CPUID available, will return 0
467:     jmp    end_cpuid

```

## kernel/core386.S

Page 8/12

```

468:
469: test386:
470:     mov     %ecx, %eax           # get original EFLAGS
471:     xor     $0x40000, %eax      # change bit 18 (AC) in EFLAGS
472:     push   %eax                # save new EFLAGS on stack
473:     popf                    # replace current EFLAGS
474:     pushf
475:     pop     %eax
476:     cmp     %ecx, %eax         # compare if both EFLAGS are equal
477:     movb   $3, %bl            # looks like an i386, return 3
478:     je     end_cpuid
479:     movb   $4, %bl            # otherwise is an old i486, return 4
480:
481: end_cpuid:
482:     push   %ecx                # push original EFLAGS
483:     popf                    # restore original EFLAGS
484:     xor     %eax, %eax
485:     movb   %bl, %al           # put return value to AL
486:
487:     popl   %ebx
488:     popl   %esi
489:     popl   %edi
490:     popl   %ebp
491:     ret
492:
493: .align 4
494: .globl getfpu; getfpu:
495:     pushl  %ebp
496:     movl   %esp, %ebp
497:     pushl  %edi
498:     pushl  %esi
499:     pushl  %ebx
500:
501:     fninit
502:     movl   $0x5a5a, _fpstatus
503:     fnstsw _fpstatus
504:     movl   _fpstatus, %eax
505:     cmp    $0, %al
506:     movl   $0, _fpstatus
507:     jne   end_getfpu
508:
509: check_control_word:
510:     fnstcw _fpstatus
511:     movl   _fpstatus, %eax
512:     andl   $0x103f, %eax
513:     cmp    $0x3f, %ax
514:     movl   $0, _fpstatus
515:     jne   end_getfpu
516:     movl   $1, _fpstatus
517:
518: end_getfpu:
519:     movl   _fpstatus, %eax
520:     cmp    $0, %al
521:     jne   1f                  # return if there is a coprocessor
522:     movl   %cr0, %eax         # otherwise (no math processor):
523:     orl    $CR0_EM, %eax      # - set EM (Emulation)
524:     andl   $CR0_MP, %eax     # - clear MP (Monitor Coprocessor)
525:     movl   %eax, %cr0
526:     movl   $0, %eax
527: 1:
528:     popl   %ebx
529:     popl   %esi
530:     popl   %edi
531:     popl   %ebp
532:     ret
533:
534: .align 4

```

## kernel/core386.S

Page 9/12

```

535: .globl vendor_id; vendor_id:
536:     pushl   %ebp
537:     movl    %esp, %ebp
538:     pushl   %ebx
539:     pushl   %edx
540:     pushl   %ecx
541:
542:     mov     $0, %eax
543:     cpuid
544:     movl    %ebx, _vendorid           # save the 12 bytes of vendor ID string
545:     movl    %edx, _vendorid+4
546:     movl    %ecx, _vendorid+8
547:
548:     popl    %ecx
549:     popl    %edx
550:     popl    %ebx
551:     popl    %ebp
552:     ret                                           # EAX returns the highest CPUID value
553:
554: .align 4
555: .globl signature_flags; signature_flags:
556:     pushl   %ebp
557:     movl    %esp, %ebp
558:     pushl   %edi
559:     pushl   %esi
560:     pushl   %ebx
561:     pushl   %edx
562:
563:     mov     $1, %eax
564:     cpuid
565:     movl    %eax, _cpusignature       # signature (model and stepping)
566:     movl    %ebx, _brandid           # misc. information
567:     movl    %edx, _cpuflags         # feature flags
568:     shrl   $8, %eax
569:     andl   $0xF, %eax
570:     movl    %eax, _cputype          # family
571:
572:     popl    %edx
573:     popl    %ebx
574:     popl    %esi
575:     popl    %edi
576:     popl    %ebp
577:     ret
578:
579: .align 4
580: .globl brand_str; brand_str:
581:     pushl   %ebp
582:     movl    %esp, %ebp
583:     pushl   %edi
584:     pushl   %esi
585:     pushl   %ebx
586:
587:     movl    $0x80000000, %eax
588:     cpuid
589:     cmp     $0x80000000, %eax         # check if brand string is supported
590:     jbe    no_brand_str
591:     movl    $0x80000002, %eax         # get first 16 bytes of brand string
592:     cpuid
593:     movl    %eax, _brandstr
594:     movl    %ebx, _brandstr+4
595:     movl    %ecx, _brandstr+8
596:     movl    %edx, _brandstr+12
597:     movl    $0x80000003, %eax         # get more 16 bytes of brand string
598:     cpuid
599:     movl    %eax, _brandstr+16
600:     movl    %ebx, _brandstr+20
601:     movl    %ecx, _brandstr+24

```

## kernel/core386.S

Page 10/12

```

602:        movl    %edx, _brandstr+28
603:        movl    $0x80000004, %eax        # get last 16 bytes of brand string
604:        cpuid
605:        movl    %eax, _brandstr+32
606:        movl    %ebx, _brandstr+36
607:        movl    %ecx, _brandstr+40
608:        movl    %edx, _brandstr+44
609:        jmp     end_brand_str
610:
611: no_brand_str:
612:        movl    $1, %eax
613:
614: end_brand_str:
615:        movl    $0, %eax
616:        popl   %ebx
617:        popl   %esi
618:        popl   %edi
619:        popl   %ebp
620:        ret
621:
622: .align 4
623: .globl tlbinf; tlbinf:
624:        pushl  %edx
625:        pushl  %ecx
626:        mov    $2, %eax
627:        cpuid
628:        movl    %eax, _tlbinf_eax        # store cache information
629:        movl    %ebx, _tlbinf_ebx
630:        movl    %edx, _tlbinf_ecx
631:        movl    %ecx, _tlbinf_edx
632:        popl   %ecx
633:        popl   %edx
634:        ret
635:
636: .align 4
637: .globl inport_b; inport_b:
638:        pushl  %ebp
639:        movl    %esp, %ebp
640:
641:        movw   0x08(%ebp), %dx        # port addr
642:        inb    %dx, %al
643:
644:        jmp    1f                    # recovery time
645: 1:        jmp    1f                    # recovery time
646: 1:        popl   %ebp
647:        ret
648:
649: .align 4
650: .globl inport_w; inport_w:
651:        pushl  %ebp
652:        movl    %esp, %ebp
653:
654:        movw   0x08(%ebp), %dx        # port addr
655:        inw    %dx, %ax
656:
657:        jmp    1f                    # recovery time
658: 1:        jmp    1f                    # recovery time
659: 1:        popl   %ebp
660:        ret
661:
662: .align 4
663: .globl inport_sw; inport_sw:
664:        pushl  %ebp
665:        movl    %esp, %ebp
666:        pushl  %edx
667:        pushl  %edi
668:        pushl  %ecx

```



## kernel/core386.S

Page 11/12

```

669:
670:         cld
671:         mov     0x8(%ebp), %edx        # port addr
672:         mov     0xC(%ebp), %edi       # dest
673:         mov     0x10(%ebp), %ecx      # count
674:         rep
675:         insw
676:
677:         popl   %ecx
678:         popl   %edi
679:         popl   %edx
680:         popl   %ebp
681:         ret
682:
683: .align 4
684: .globl outport_b; outport_b:
685:         pushl  %ebp
686:         movl   %esp, %ebp
687:
688:         movw   0x8(%ebp), %dx        # port addr
689:         movb   0xC(%ebp), %al       # data
690:         outb   %al, %dx
691:
692:         jmp    1f                    # recovery time
693: 1:      jmp    1f                    # recovery time
694: 1:      popl   %ebp
695:         ret
696:
697: .align 4
698: .globl outport_w; outport_w:
699:         pushl  %ebp
700:         movl   %esp, %ebp
701:
702:         movw   0x8(%ebp), %dx        # port addr
703:         movw   0xC(%ebp), %ax       # data
704:         outw   %ax, %dx
705:
706:         jmp    1f                    # recovery time
707: 1:      jmp    1f                    # recovery time
708: 1:      popl   %ebp
709:         ret
710:
711: .align 4
712: .globl outport_sw; outport_sw:
713:         pushl  %ebp
714:         movl   %esp, %ebp
715:         pushl  %edx
716:         pushl  %esi
717:         pushl  %ecx
718:
719:         cld
720:         mov     0x8(%ebp), %edx        # port addr
721:         mov     0xC(%ebp), %esi       # src
722:         mov     0x10(%ebp), %ecx      # count
723:         rep
724:         outsw
725:
726:         popl   %ecx
727:         popl   %esi
728:         popl   %edx
729:         popl   %ebp
730:         ret
731:
732: .align 4
733: .globl load_gdt; load_gdt:
734:         movl   0x4(%esp), %eax
735:         lgdt   (%eax)

```

## kernel/core386.S

Page 12/12

```

736:      movw    $KERNEL_DS, %ax
737:      movw    %ax, %ds
738:      movw    %ax, %es
739:      movw    %ax, %fs
740:      movw    %ax, %gs
741:      movw    %ax, %ss
742:      ljmp   $KERNEL_CS, $1f
743: 1:
744:      ret
745:
746: .align 4
747: .globl load_idt; load_idt:
748:      movl   0x4(%esp), %eax
749:      lidt   (%eax)
750:      ret
751:
752: .align 4
753: .globl activate_kpage_dir; activate_kpage_dir:
754:      movl   kpage_dir, %eax
755:      movl   %eax, %cr3
756:      ret
757:
758: .align 4
759: .globl load_tr; load_tr:
760:      mov    0x4(%esp), %ax
761:      ltr    %ax
762:      ret
763:
764: .align 4
765: .globl get_rdtsc; get_rdtsc:
766:      cpuid
767:      rdtsc
768:      ret
769:
770: .align 4
771: .globl invalidate_tlb; invalidate_tlb:
772:      movl   %cr3, %eax
773:      movl   %eax, %cr3
774:      ret
775:
776:
777: .data
778:
779: .globl _cputype
780: .globl _cpusignature
781: .globl _cpuflags
782: .globl _fpstatus
783: .globl _brandid
784: .globl _vendorid
785: .globl _brandstr
786: .globl _tlbinfo_eax
787: .globl _tlbinfo_ebx
788: .globl _tlbinfo_ecx
789: .globl _tlbinfo_edx
790:
791: _cputype:      .int    0
792: _cpusignature: .int    0
793: _cpuflags:    .int    0
794: _fpstatus:    .int    0
795: _brandid:     .int    0
796: _vendorid:    .fill   13,1,0
797: _brandstr:    .fill   49,1,0
798: _tlbinfo_eax: .int    0
799: _tlbinfo_ebx: .int    0
800: _tlbinfo_ecx: .int    0
801: _tlbinfo_edx: .int    0

```

## kernel/cpu.c

Page 1/5

```

1: /*
2:  * fiwix/kernel/cpu.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/utsname.h>
11: #include <fiwix/pic.h>
12: #include <fiwix/pit.h>
13: #include <fiwix/cpu.h>
14: #include <fiwix/timer.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: char UTS_MACHINE[_UTSNAME_LENGTH];
19:
20: static struct cpu_type intel[] = {
21:     { 4,
22:       { "i486 DX", "i486 DX", "i486 SX", "i486 DX/2",
23:         "i486 SL", "i486 SX/2", NULL, "i486 DX/2 WBE",
24:         "i486 DX/4", NULL, NULL, NULL, NULL, NULL, NULL, NULL }
25:     },
26:     { 5,
27:       { NULL, "Pentium 60/66", "Pentium 75-200", "Pentium ODfor486",
28:         "PentiumMMX", NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL,
29:         NULL, NULL, NULL }
30:     },
31:     { 6,
32:       { NULL, "Pentium Pro", NULL, "Pentium II", NULL, "Pentium II",
33:         "Intel Celeron", "Pentium III", "Pentium III", NULL,
34:         "Pentium III Xeon", "Pentium III", NULL, NULL, NULL, NULL }
35:     }
36: };
37:
38: static const char *cpu_flags[] = {
39:     "FPU", "VME", "DE", "PSE", "TSC", "MSR", "PAE", "MCE", "CX8", "APIC",
40:     "10", "SEP", "MTRR", "PGE", "MCA", "CMOV", "PAT", "PSE-36", "PSN",
41:     "CLFSH", "20", "DS", "ACPI", "MMX", "FXSR", "SSE", "SSE2", "SS",
42:     "HTT", "TM", "30", "PBE"
43: };
44:
45: static unsigned long int detect_cpuspeed(void)
46: {
47:     unsigned long long int tsc1, tsc2;
48:
49:     outport_b(MODEREG, SEL_CHAN2 | LSB_MSB | TERM_COUNT | BINARY_CTR);
50:     outport_b(CHANNEL2, (OSCIL / HZ) & 0xFF);
51:     outport_b(CHANNEL2, (OSCIL / HZ) >> 8);
52:     outport_b(PS2_SYSCtrl_B, inport_b(PS2_SYSCtrl_B) | ENABLE_SDATA | ENABLE
_TMR2G);
53:
54:     tsc1 = 0;
55:     tsc1 = get_rdtsc();
56:
57:     while(!(inport_b(PS2_SYSCtrl_B) & 0x20));
58:
59:     tsc2 = 0;
60:     tsc2 = get_rdtsc();
61:
62:     outport_b(PS2_SYSCtrl_B, inport_b(PS2_SYSCtrl_B) & ~(ENABLE_SDATA | ENAB
LE_TMR2G));
63:
64:     return (tsc2 - tsc1) * HZ;
65: }

```

## kernel/cpu.c

Page 2/5

```
66:
67: /*
68:  * These are the 2nd and 3rd level cache values according to Intel Processor
69:  * Identification and the CPUID Instruction.
70:  * Application Note 485. Document Number: 241618-031. September 2006.
71:  */
72: static void show_cache(int value)
73: {
74:     switch(value) {
75:         /* 2nd level cache */
76:         case 0x39:
77:         case 0x3B:
78:         case 0x41:
79:         case 0x79:
80:             cpu_table.cache = "128KB L2";
81:             break;
82:         case 0x3A:
83:             cpu_table.cache = "192KB L2";
84:             break;
85:         case 0x3C:
86:         case 0x42:
87:         case 0x7A:
88:         case 0x82:
89:             cpu_table.cache = "256KB L2";
90:             break;
91:         case 0x3D:
92:             cpu_table.cache = "384KB L2";
93:             break;
94:         case 0x3E:
95:         case 0x43:
96:         case 0x7B:
97:         case 0x7F:
98:         case 0x83:
99:         case 0x86:
100:            cpu_table.cache = "512KB L2";
101:            break;
102:         case 0x44:
103:         case 0x78:
104:         case 0x7C:
105:         case 0x84:
106:         case 0x87:
107:            cpu_table.cache = "1MB L2";
108:            break;
109:         case 0x45:
110:         case 0x7D:
111:         case 0x85:
112:            cpu_table.cache = "2MB L2";
113:            break;
114:
115:         /* 3rd level cache */
116:         case 0x22:
117:             cpu_table.cache = "512KB L3";
118:             break;
119:         case 0x23:
120:             cpu_table.cache = "1MB L3";
121:             break;
122:         case 0x25:
123:             cpu_table.cache = "2MB L3";
124:             break;
125:         case 0x29:
126:         case 0x46:
127:             cpu_table.cache = "4MB L3";
128:             break;
129:         case 0x49:
130:             cpu_table.cache = "4MB L3 & L2";
131:             break;
132:         case 0x4A:
```

## kernel/cpu.c

Page 3/5

```

133:         cpu_table.cache = "6MB L3";
134:         break;
135:     case 0x47:
136:     case 0x4B:
137:         cpu_table.cache = "8MB L3";
138:         break;
139:     case 0x4C:
140:         cpu_table.cache = "12MB L3";
141:         break;
142:     case 0x4D:
143:         cpu_table.cache = "16MB L3";
144:         break;
145:     default:
146:         break;
147:     }
148: }
149:
150: static void check_cache(int maxcpuid)
151: {
152:     int n, maxcpuids;
153:
154:     maxcpuids = 1;
155:     if(maxcpuid >= 2) {
156:         for(n = 0; n < maxcpuids; n++) {
157:             tlbinfo();
158:             maxcpuids = _tlbinfo_eax & 0xFF;
159:             show_cache((_tlbinfo_eax >> 8) & 0xFF);
160:             show_cache((_tlbinfo_eax >> 16) & 0xFF);
161:             show_cache((_tlbinfo_eax >> 24) & 0xFF);
162:             if(!_tlbinfo_ebx & RESERVED_DESC) {
163:                 show_cache(_tlbinfo_ebx & 0xFF);
164:                 show_cache((_tlbinfo_ebx >> 8) & 0xFF);
165:                 show_cache((_tlbinfo_ebx >> 16) & 0xFF);
166:                 show_cache((_tlbinfo_ebx >> 24) & 0xFF);
167:             }
168:             if(!_tlbinfo_ecx & RESERVED_DESC) {
169:                 show_cache(_tlbinfo_ecx & 0xFF);
170:                 show_cache((_tlbinfo_ecx >> 8) & 0xFF);
171:                 show_cache((_tlbinfo_ecx >> 16) & 0xFF);
172:                 show_cache((_tlbinfo_ecx >> 24) & 0xFF);
173:             }
174:             if(!_tlbinfo_edx & RESERVED_DESC) {
175:                 show_cache(_tlbinfo_edx & 0xFF);
176:                 show_cache((_tlbinfo_edx >> 8) & 0xFF);
177:                 show_cache((_tlbinfo_edx >> 16) & 0xFF);
178:                 show_cache((_tlbinfo_edx >> 24) & 0xFF);
179:             }
180:         }
181:     }
182: }
183:
184: int get_cpu_flags(char *buffer, int offset)
185: {
186:     int n, size;
187:     unsigned int mask;
188:
189:     size = sprintf(buffer + offset, "flags          :");
190:     for(n = 0, mask = 1; n < 32; n++, mask <<= 1) {
191:         if(_cpuflags & mask) {
192:             size += sprintf(buffer + offset + size, " %s", cpu_flags
[n]);
193:         }
194:     }
195:     size += sprintf(buffer + offset + size, "\n");
196:     return size;
197: }
198:

```



**kernel/cpu.c**

Page 5/5

```
259:                printk("80%d86\n", cpu_table.family);
260:                cpu_table.has_cpuid = 0;
261:            }
262:            cpu_table.has_fpu = getfpu();
263: }
```

## kernel/gdt.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/gdt.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/const.h>
10: #include <fiwix/types.h>
11: #include <fiwix/segments.h>
12: #include <fiwix/process.h>
13: #include <fiwix/limits.h>
14: #include <fiwix/string.h>
15:
16: struct seg_desc gdt[NR_GDT_ENTRIES];
17:
18: struct desc_r gdtr = {
19:     sizeof(gdt) - 1,
20:     (unsigned int)&gdt
21: };
22:
23: static void gdt_set_entry(int num, unsigned int base_addr, unsigned int limit, c
har loflags, char hiflags)
24: {
25:     num /= sizeof(struct seg_desc);
26:     gdt[num].sd_lolimit = limit & 0xFFFF;
27:     gdt[num].sd_lobase = base_addr & 0xFFFFFFFF;
28:     gdt[num].sd_loflags = loflags;
29:     gdt[num].sd_hilimit = (limit >> 16) & 0x0F;
30:     gdt[num].sd_hiflags = hiflags;
31:     gdt[num].sd_hibase = (base_addr >> 24) & 0xFF;
32: }
33:
34: void gdt_init(void)
35: {
36:     unsigned char loflags;
37:
38:     gdt_set_entry(0, 0, 0, 0, 0); /* null descriptor */
39:
40:     loflags = SD_CODE | SD_CD | SD_DPL0 | SD_PRESENT;
41:     gdt_set_entry(KERNEL_CS, 0, 0xFFFFFFFF, loflags, SD_OPSIZE32 | SD_PAGE4K
B);
42:     loflags = SD_DATA | SD_CD | SD_DPL0 | SD_PRESENT;
43:     gdt_set_entry(KERNEL_DS, 0, 0xFFFFFFFF, loflags, SD_OPSIZE32 | SD_PAGE4K
B);
44:
45:     loflags = SD_CODE | SD_CD | SD_DPL3 | SD_PRESENT;
46:     gdt_set_entry(USER_CS, 0, 0xFFFFFFFF, loflags, SD_OPSIZE32 | SD_PAGE4KB)
;
47:     loflags = SD_DATA | SD_CD | SD_DPL3 | SD_PRESENT;
48:     gdt_set_entry(USER_DS, 0, 0xFFFFFFFF, loflags, SD_OPSIZE32 | SD_PAGE4KB)
;
49:
50:     loflags = SD_TSSPRESENT;
51:     gdt_set_entry(TSS, 0, sizeof(struct proc) - 1, loflags, SD_OPSIZE32);
52:
53:     load_gdt((unsigned int)&gdtr);
54: }

```



## kernel/idt.c

```

1: /*
2:  * fiwix/kernel/idt.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/const.h>
10: #include <fiwix/types.h>
11: #include <fiwix/segments.h>
12: #include <fiwix/string.h>
13:
14: struct gate_desc idt[NR_IDT_ENTRIES];
15:
16: struct desc_r idtr = {
17:     sizeof(idt) - 1,
18:     (unsigned int)idt
19: };
20:
21: static void set_idt_entry(int num, __off_t handler, unsigned int flags)
22: {
23:     idt[num].gd_looffset = handler & 0x0000FFFF;
24:     idt[num].gd_selector = KERNEL_CS;
25:     idt[num].gd_flags = flags << 8;
26:     idt[num].gd_hioffset = handler >> 16;
27: }
28:
29: void idt_init(void)
30: {
31:     int n;
32:
33:     memset_b(idt, NULL, sizeof(idt));
34:     for(n = 0; n < NR_IDT_ENTRIES; n++) {
35:         set_idt_entry(n, (__off_t)&unknown_irq, SD_32INTRGATE | SD_PRESE
NT);
36:     }
37:
38:     /* FIXME: must be SD_32TRAPGATE for true multitasking */
39:     set_idt_entry(0x00, (__off_t)&except0, SD_32INTRGATE | SD_PRESENT);
40:     set_idt_entry(0x01, (__off_t)&except1, SD_32INTRGATE | SD_PRESENT);
41:     set_idt_entry(0x02, (__off_t)&except2, SD_32INTRGATE | SD_PRESENT);
42:     set_idt_entry(0x03, (__off_t)&except3, SD_32INTRGATE | SD_PRESENT);
43:     set_idt_entry(0x04, (__off_t)&except4, SD_32INTRGATE | SD_PRESENT);
44:     set_idt_entry(0x05, (__off_t)&except5, SD_32INTRGATE | SD_PRESENT);
45:     set_idt_entry(0x06, (__off_t)&except6, SD_32INTRGATE | SD_PRESENT);
46:     set_idt_entry(0x07, (__off_t)&except7, SD_32INTRGATE | SD_PRESENT);
47:     set_idt_entry(0x08, (__off_t)&except8, SD_32INTRGATE | SD_PRESENT);
48:     set_idt_entry(0x09, (__off_t)&except9, SD_32INTRGATE | SD_PRESENT);
49:     set_idt_entry(0x0A, (__off_t)&exceptA, SD_32INTRGATE | SD_PRESENT);
50:     set_idt_entry(0x0B, (__off_t)&exceptB, SD_32INTRGATE | SD_PRESENT);
51:     set_idt_entry(0x0C, (__off_t)&exceptC, SD_32INTRGATE | SD_PRESENT);
52:     set_idt_entry(0x0D, (__off_t)&exceptD, SD_32INTRGATE | SD_PRESENT);
53:     set_idt_entry(0x0E, (__off_t)&exceptE, SD_32INTRGATE | SD_PRESENT);
54:     set_idt_entry(0x0F, (__off_t)&exceptF, SD_32INTRGATE | SD_PRESENT);
55:     set_idt_entry(0x10, (__off_t)&except10, SD_32INTRGATE | SD_PRESENT);
56:     set_idt_entry(0x11, (__off_t)&except11, SD_32INTRGATE | SD_PRESENT);
57:     set_idt_entry(0x12, (__off_t)&except12, SD_32INTRGATE | SD_PRESENT);
58:     set_idt_entry(0x13, (__off_t)&except13, SD_32INTRGATE | SD_PRESENT);
59:     set_idt_entry(0x14, (__off_t)&except14, SD_32INTRGATE | SD_PRESENT);
60:     set_idt_entry(0x15, (__off_t)&except15, SD_32INTRGATE | SD_PRESENT);
61:     set_idt_entry(0x16, (__off_t)&except16, SD_32INTRGATE | SD_PRESENT);
62:     set_idt_entry(0x17, (__off_t)&except17, SD_32INTRGATE | SD_PRESENT);
63:     set_idt_entry(0x18, (__off_t)&except18, SD_32INTRGATE | SD_PRESENT);
64:     set_idt_entry(0x19, (__off_t)&except19, SD_32INTRGATE | SD_PRESENT);
65:     set_idt_entry(0x1A, (__off_t)&except1A, SD_32INTRGATE | SD_PRESENT);
66:     set_idt_entry(0x1B, (__off_t)&except1B, SD_32INTRGATE | SD_PRESENT);

```

## kernel/idt.c

Page 2/2

```

67:     set_idt_entry(0x1C, (__off_t)&except1C, SD_32INTRGATE | SD_PRESENT);
68:     set_idt_entry(0x1D, (__off_t)&except1D, SD_32INTRGATE | SD_PRESENT);
69:     set_idt_entry(0x1E, (__off_t)&except1E, SD_32INTRGATE | SD_PRESENT);
70:     set_idt_entry(0x1F, (__off_t)&except1F, SD_32INTRGATE | SD_PRESENT);
71:
72:     set_idt_entry(0x20, (__off_t)&irq0, SD_32INTRGATE | SD_PRESENT);
73:     set_idt_entry(0x21, (__off_t)&irq1, SD_32INTRGATE | SD_PRESENT);
74:     set_idt_entry(0x22, (__off_t)&irq2, SD_32INTRGATE | SD_PRESENT);
75:     set_idt_entry(0x23, (__off_t)&irq3, SD_32INTRGATE | SD_PRESENT);
76:     set_idt_entry(0x24, (__off_t)&irq4, SD_32INTRGATE | SD_PRESENT);
77:     set_idt_entry(0x25, (__off_t)&irq5, SD_32INTRGATE | SD_PRESENT);
78:     set_idt_entry(0x26, (__off_t)&irq6, SD_32INTRGATE | SD_PRESENT);
79:     set_idt_entry(0x27, (__off_t)&irq7, SD_32INTRGATE | SD_PRESENT);
80:     set_idt_entry(0x28, (__off_t)&irq8, SD_32INTRGATE | SD_PRESENT);
81:     set_idt_entry(0x29, (__off_t)&irq9, SD_32INTRGATE | SD_PRESENT);
82:     set_idt_entry(0x2A, (__off_t)&irq10, SD_32INTRGATE | SD_PRESENT);
83:     set_idt_entry(0x2B, (__off_t)&irq11, SD_32INTRGATE | SD_PRESENT);
84:     set_idt_entry(0x2C, (__off_t)&irq12, SD_32INTRGATE | SD_PRESENT);
85:     set_idt_entry(0x2D, (__off_t)&irq13, SD_32INTRGATE | SD_PRESENT);
86:     set_idt_entry(0x2E, (__off_t)&irq14, SD_32INTRGATE | SD_PRESENT);
87:     set_idt_entry(0x2F, (__off_t)&irq15, SD_32INTRGATE | SD_PRESENT);
88:
89:     /* FIXME: must be SD_32TRAPGATE for true multitasking */
90:     set_idt_entry(0x80, (__off_t)&syscall, SD_32INTRGATE | SD_DPL3 | SD_PRESENT);
91:
92:     load_idt((unsigned int)&idtr);
93: }

```

## kernel/init.c

```

1: /*
2:  * fiwix/kernel/init.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/system.h>
11: #include <fiwix/mm.h>
12: #include <fiwix/timer.h>
13: #include <fiwix/sched.h>
14: #include <fiwix/fcntl.h>
15: #include <fiwix/stat.h>
16: #include <fiwix/process.h>
17: #include <fiwix/syscalls.h>
18: #include <fiwix/unistd.h>
19: #include <fiwix/stdio.h>
20: #include <fiwix/string.h>
21:
22: #define INIT_TRAMPOLINE_SIZE    128    /* max. size of init_trampoline() */
23:
24: char *init_argv[] = { INIT_PROGRAM, NULL, NULL };
25: char *init_envp[] = { "HOME=/", "TERM=linux", NULL };
26:
27: static void init_trampoline(void)
28: {
29:     USER_SYSCALL(SYS_open, "/dev/console", O_RDWR, 0);    /* stdin */
30:     USER_SYSCALL(SYS_dup, 0, NULL, NULL);    /* stdout */
31:     USER_SYSCALL(SYS_dup, 0, NULL, NULL);    /* stderr */
32:     USER_SYSCALL(SYS_execve, INIT_PROGRAM, init_argv, init_envp);
33:
34:     /* only reached in case of error in sys_execve() */
35:     USER_SYSCALL(SYS_exit, NULL, NULL, NULL);
36: }
37:
38: void init_init(void)
39: {
40:     int n;
41:     unsigned int page;
42:     struct inode *i;
43:     unsigned int *pgdir;
44:     struct proc *init;
45:
46:     if(namei(INIT_PROGRAM, &i, NULL, FOLLOW_LINKS)) {
47:         PANIC("can't find %s.\n", INIT_PROGRAM);
48:     }
49:     if(!S_ISREG(i->i_mode)) {
50:         PANIC("%s is not a regular file.\n", INIT_PROGRAM);
51:     }
52:     iput(i);
53:
54:     /* INIT slot was already created in main.c */
55:     init = &proc_table[INIT];
56:
57:     /* INIT process starts with the current (kernel) Page Directory */
58:     if(!(pgdir = (void *)kmalloc())) {
59:         goto init_init__die;
60:     }
61:     init->rss++;
62:     memcpy_b(pgdir, kpage_dir, PAGE_SIZE);
63:     init->tss.cr3 = V2P((unsigned int)pgdir);
64:
65:     memset_b(init->vma, NULL, sizeof(init->vma));
66:     init->ppid = 0;
67:     init->pgid = 0;

```

## kernel/init.c

Page 2/2

```

68:         init->sid = 0;
69:         init->flags = 0;
70:         init->children = 0;
71:         init->priority = DEF_PRIORITY;
72:         init->start_time = CURRENT_TICKS;
73:         init->sleep_address = NULL;
74:         init->uid = init->gid = 0;
75:         init->euid = init->egid = 0;
76:         init->suid = init->sgid = 0;
77:         memset_b(init->fd, NULL, sizeof(init->fd));
78:         memset_b(init->fd_flags, NULL, sizeof(init->fd_flags));
79:         init->root = current->root;
80:         init->pwd = current->pwd;
81:         strcpy(init->argv0, init_argv[0]);
82:         init_argv[1] = init_args;
83:         sprintk(init->pidstr, "%d", init->pid);
84:         init->sigpending = 0;
85:         init->sigblocked = 0;
86:         init->sigexecuting = 0;
87:         memset_b(init->sigaction, NULL, sizeof(init->sigaction));
88:         memset_b(&init->usage, NULL, sizeof(struct rusage));
89:         memset_b(&init->cusage, NULL, sizeof(struct rusage));
90:         init->timeout = 0;
91:         for(n = 0; n < RLIM_NLIMITS; n++) {
92:             init->rlim[n].rlim_cur = init->rlim[n].rlim_max = RLIM_INFINITY;
93:         }
94:         init->rlim[RLIMIT_NOFILE].rlim_cur = OPEN_MAX;
95:         init->rlim[RLIMIT_NOFILE].rlim_max = NR_OPENS;
96:         init->rlim[RLIMIT_NPROC].rlim_cur = CHILD_MAX;
97:         init->rlim[RLIMIT_NPROC].rlim_max = NR_PROCS;
98:         init->umask = 0022;
99:
100:        /* setup the stack */
101:        if(!(init->tss.esp0 = kmalloc())) {
102:            goto init_init__die;
103:        }
104:        init->tss.esp0 += PAGE_SIZE - 4;
105:        init->rss++;
106:        init->tss.ss0 = KERNEL_DS;
107:
108:        /* setup the init_trampoline */
109:        page = map_page(init, KERNEL_BASE_ADDR - PAGE_SIZE, 0, PROT_READ | PROT_
WRITE);
110:        memcpy_b((void *)page, init_trampoline, INIT_TRAMPOLINE_SIZE);
111:
112:        init->tss.eip = (unsigned int)switch_to_user_mode;
113:        init->tss.esp = page + PAGE_SIZE - 4;
114:
115:        init->state = PROC_RUNNING;
116:        nr_processes++;
117:        return;
118:
119:    init_init__die:
120:        PANIC("unable to run init process.\n");
121:    }

```

## kernel/main.c

```

1: /*
2:  * fiwix/kernel/main.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/mm.h>
11: #include <fiwix/system.h>
12: #include <fiwix/timer.h>
13: #include <fiwix/sched.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/cpu.h>
16: #include <fiwix/pic.h>
17: #include <fiwix/fs.h>
18: #include <fiwix/devices.h>
19: #include <fiwix/console.h>
20: #include <fiwix/keyboard.h>
21: #include <fiwix/ramdisk.h>
22: #include <fiwix/version.h>
23: #include <fiwix/limits.h>
24: #include <fiwix/segments.h>
25: #include <fiwix/syscalls.h>
26: #include <fiwix/stdio.h>
27: #include <fiwix/string.h>
28: #include <fiwix/kparms.h>
29: #include <fiwix/i386elf.h>
30: #include <fiwix/bios.h>
31:
32: /*
33:  * check if the bit BIT in Multiboot FLAGS is set
34:  * -----
35:  * bit 11 -> vbe_*
36:  * bit 10 -> apm_table
37:  * bit 9 -> boot_loader_name
38:  * bit 8 -> config_table
39:  * bit 7 -> drives_length and drives_addr
40:  * bit 6 -> mmap_length and mmap_addr
41:  * bit 5 -> ELF symbols
42:  * bit 4 -> a.out symbols
43:  * bit 3 -> mods_count and mods_addr
44:  * bit 2 -> cmdline
45:  * bit 1 -> boot_device
46:  * bit 0 -> mem_lower and mem_upper values
47:  */
48: #define CHECK_FLAG(flags,bit) ((flags) & (1 << (bit)))
49:
50: Elf32_Shdr *symtab, *strtab;
51: unsigned int _last_data_addr;
52: int _memsize;
53: int _extmemsize;
54: int _rootdev;
55: int _noramdisk;
56: int _ramdisksize;
57: char _rootfstype[10];
58: char _rootdevname[DEVNAME_MAX + 1];
59: char _initrd[DEVNAME_MAX + 1];
60: int _syscondev;
61: char *init_args;
62:
63: char cmdline[NAME_MAX + 1];
64:
65: struct new_utsname sys_utsname = {
66:     UTS_SYSNAME,
67:     UTS_NODENAME,

```

## kernel/main.c

Page 2/7

```

68:         UTS_RELEASE,
69:         UTS_VERSION,
70:         "",
71:         UTS_DOMAINNAME,
72: };
73:
74: struct kernel_stat kstat;
75:
76: /*
77:  * This function returns the last address used by kernel symbols or the value
78:  * of 'mod_end' (in the module structure) of the last module loaded by GRUB.
79:  *
80:  * This is intended to setup the kernel stack beyond all these addresses.
81:  */
82: unsigned int get_last_boot_addr(unsigned int info)
83: {
84:     multiboot_info_t *mbi;
85:     Elf32_Shdr *shdr;
86:     elf_section_header_table_t *hdr;
87:     struct module *mod;
88:     unsigned short int n;
89:     unsigned int addr;
90:
91:     symtab = strtab = NULL;
92:     mbi = (multiboot_info_t *)info;
93:     hdr = &(mbi->u.elf_sec);
94:     for(n = 0; n < hdr->num; n++) {
95:         shdr = (Elf32_Shdr *) (hdr->addr + (n * hdr->size));
96:         if(shdr->sh_type == SHT_SYMTAB) {
97:             symtab = shdr;
98:         }
99:         if(shdr->sh_type == SHT_STRTAB) {
100:             strtab = shdr;
101:         }
102:     }
103:
104:     addr = strtab->sh_addr + strtab->sh_size;
105:
106:     /*
107:      * https://www.gnu.org/software/grub/manual/multiboot/multiboot.html
108:      *
109:      * Check if GRUB has loaded some modules and, if so, get the last
110:      * address used by the last one.
111:      */
112:     if(CHECK_FLAG(mbi->flags, 3)) {
113:         mod = (struct module *)mbi->mods_addr;
114:         for(n = 0; n < mbi->mods_count; n++, mod++) {
115:             addr = mod->mod_end;
116:         }
117:     }
118:
119:     return P2V(addr);
120: }
121:
122: /* check the validity of a command line parameter */
123: static int check_parm(struct kparms *parm, const char *value)
124: {
125:     int n;
126:
127:     if(!strcmp(parm->name, "root=")) {
128:         for(n = 0; parm->value[n]; n++) {
129:             if(!strcmp(parm->value[n], value)) {
130:                 _rootdev = parm->sysval[n];
131:                 strncpy(_rootdevname, value, DEVNAME_MAX);
132:                 return 0;
133:             }
134:         }

```

```

135:         return 1;
136:     }
137:     if(!strcmp(parm->name, "noramdisk")) {
138:         _noramdisk = 1;
139:         return 0;
140:     }
141:     if(!strcmp(parm->name, "ramdisksize=")) {
142:         int size = atoi(value);
143:         if(!size || size > RAMDISK_MAXSIZE) {
144:             printk("WARNING: 'ramdisksize' value is out of limits, d
efaulting to 4096KB.\n");
145:             _ramdisksize = 0;
146:         } else {
147:             _ramdisksize = size;
148:         }
149:         return 0;
150:     }
151:     if(!strcmp(parm->name, "initrd=")) {
152:         if(value[0]) {
153:             strncpy(_initrd, value, DEVNAME_MAX);
154:             return 0;
155:         }
156:     }
157:     if(!strcmp(parm->name, "rootfstype=")) {
158:         for(n = 0; parm->value[n]; n++) {
159:             if(!strcmp(parm->value[n], value)) {
160:                 strncpy(_rootfstype, value, sizeof(_rootfstype))
;
161:                 return 0;
162:             }
163:         }
164:         return 1;
165:     }
166:     if(!strcmp(parm->name, "console=")) {
167:         for(n = 0; parm->value[n]; n++) {
168:             if(!strcmp(parm->value[n], value)) {
169:                 if(parm->sysval[n]) {
170:                     _syscondev = parm->sysval[n];
171:                     return 0;
172:                 }
173:                 printk("WARNING: device name for '%s' is not def
ined!\n", parm->name);
174:             }
175:         }
176:         return 1;
177:     }
178:     printk("WARNING: the parameter '%s' looks valid but it's not defined!\n"
, parm->name);
179:     return 0;
180: }
181:
182: static int parse_arg(const char *arg)
183: {
184:     int n;
185:
186:     /* '--' marks the beginning of the init arguments */
187:     if(!strcmp(arg, "--")) {
188:         return 1;
189:     }
190:
191:     for(n = 0; parm_table[n].name; n++) {
192:         if(!strncmp(arg, parm_table[n].name, strlen(parm_table[n].name))
) {
193:             arg += strlen(parm_table[n].name);
194:             if(check_parm(&parm_table[n], arg)) {
195:                 printk("WARNING: invalid value '%s' in the '%s'
parameter.\n", arg, parm_table[n].name);

```

## kernel/main.c

Page 4/7

```

196:         }
197:         return 0;
198:     }
199: }
200: printk("WARNING: invalid cmdline parameter: '%s'.\n", arg);
201: return 0;
202: }
203:
204: static char * parse_cmdline(const char *str)
205: {
206:     char *from, *to;
207:     char arg[CMDL_ARG_LEN];
208:     char c;
209:
210:     from = to = (char *)str;
211:     for(;;) {
212:         c = *(str++);
213:         if(c == ' ' || !c) {
214:             if(to - from < CMDL_ARG_LEN) {
215:                 memcpy_b(arg, from, to - from);
216:                 arg[to - from] = NULL;
217:                 if(arg[0] != NULL) {
218:                     if(parse_arg(arg)) {
219:                         while(*(from++)) {
220:                             if(*from != '-' && *from
!= ' ') {
221:                                 break;
222:                             }
223:                         }
224:                         return from;
225:                     }
226:                 }
227:             } else {
228:                 memcpy_b(arg, from, CMDL_ARG_LEN);
229:                 arg[CMDL_ARG_LEN - 1] = NULL;
230:                 printk("WARNING: invalid length of the cmdline p
arameter '%s'.\n", arg);
231:             }
232:             from = ++to;
233:             if(!c) {
234:                 break;
235:             }
236:             continue;
237:         }
238:         to++;
239:     }
240:
241:     return NULL;
242: }
243:
244: void start_kernel(unsigned long magic, unsigned long info, unsigned int stack)
245: {
246:     struct proc *init, *p_kswapd;
247:     multiboot_info_t mbi;
248:
249:     /* default kernel values */
250:     strcpy(_rootfstype, "ext2"); /* filesystem is ext2 */
251:     _syscondev = MKDEV(VCONSOLES_MAJOR, 0); /* console is /dev/tty0 */
252:
253:     pic_init();
254:     idt_init();
255:     dev_init();
256:     tty_init();
257:
258:     printk("
259:     printk("
Welcome to %s\n", UTS_SYSNAME);
Copyright (c) 2018-2020, Jordi Sanfeliu\n")
;

```



## kernel/main.c

Page 5/7

```

260:         printk("\n");
261:         printk("                kernel v%s for i386 architecture\n", UTS_R
RELEASE);
262:         printk("                (GCC %s, built on %s)\n", __VERSION__, UTS_VERSIO
N);
263:         printk("\n");
264:         printk("DEVICE      ADDRESS      IRQ      COMMENT\n");
265:         printk("-----\n");
-----\n");
266:
267:         if(magic != MULTIBOOT_BOOTLOADER_MAGIC) {
268:             printk("WARNING: invalid multiboot-bootloader magic number: 0x%x
.\n\n", (unsigned long int)magic);
269:             memset_b(&mbi, NULL, sizeof(struct multiboot_info));
270:         } else {
271:             memcpy_b(&mbi, (void *)info, sizeof(struct multiboot_info));
272:         }
273:
274:         memset_b(&kstat, NULL, sizeof(kstat));
275:
276:         cpu_init();
277:
278:         /* check if a command line was supplied */
279:         if(CHECK_FLAG(mbi.flags, 2)) {
280:             int n, len;
281:             char c;
282:             char *p;
283:
284:             p = (char *)mbi.cmdline;
285:             len = strlen(p);
286:             /* suppress 'fiwix' */
287:             for(n = 0; n < len; n++) {
288:                 c = *(p++);
289:                 if(c == ' ') {
290:                     break;
291:                 }
292:             }
293:             strcpy(cmdline, p);
294:             init_args = parse_cmdline(cmdline);
295:         } else {
296:             printk("WARNING: no cmdline detected!\n");
297:         }
298:
299:         printk("kernel      0x%08X      -      cmdline='%s'\n", KERNEL_ENTRY_ADDR,
cmdline);
300:
301:         if(CHECK_FLAG(mbi.flags, 3)) {
302:             int n;
303:             struct module *mod;
304:
305:             mod = (struct module *)mbi.mods_addr;
306:             for(n = 0; n < mbi.mods_count; n++, mod++) {
307:                 if(!strcmp((char *)mod->string, _initrd)) {
308:                     printk("initrd      0x%08X-0x%08X file='%s' size=%
dKB\n", mod->mod_start, mod->mod_end, mod->string, (mod->mod_end - mod->mod_start) / 10
24);
309:                     ramdisk_table[0].addr = (char *)mod->mod_start;
310:                 }
311:             }
312:         }
313:
314:         if(!CHECK_FLAG(mbi.flags, 0)) {
315:             printk("WARNING: values in mem_lower and mem_upper are not valid
!\n");
316:         }
317:         _memsize = (unsigned int)mbi.mem_lower;
318:         _extmemsize = (unsigned int)mbi.mem_upper;

```

## kernel/main.c

Page 6/7

```

319:
320:     if(CHECK_FLAG(mbi.flags, 6)) {
321:         bios_map_init((memory_map_t *)mbi.mmap_addr, mbi.mmap_length);
322:     } else {
323:         bios_map_init(NULL, 0);
324:     }
325:
326:     timer_init();
327:     get_video_parms();
328:
329:     _last_data_addr = stack - KERNEL_BASE_ADDR;
330:     mem_init();
331:
332:     keyboard_init();
333:     vconsole_init();
334:     proc_init();
335:
336:     if(!(CHECK_FLAG(mbi.flags, 5))) {
337:         printk("WARNING: ELF section header table is not valid!\n");
338:     }
339:
340:     /* IDLE is now the current process */
341:     set_tss(current);
342:     load_tr(TSS);
343:     current->tss.cr3 = (unsigned int)kpage_dir;
344:     current->flags |= PF_KPROC;
345:
346:     /* reserve the slot 1 for the INIT process */
347:     init = get_proc_free();
348:     proc_slot_init(init);
349:     init->pid = get_unused_pid();
350:
351:     /* create and setup kswapd process */
352:     p_kswapd = kernel_process(kswapd);
353:
354:     /* kswapd will take over the rest of the kernel initialization */
355:     p_kswapd->state = PROC_RUNNING;
356:     need_resched = 1;
357:
358:     STI();          /* let's rock! */
359:     cpu_idle();
360: }
361:
362: void stop_kernel(void)
363: {
364:     struct proc *p;
365:     int n;
366:
367:     /* stop and disable all interrupts! */
368:     CLI();
369:     for(n = 0; n < NR_IRQS; n++) {
370:         disable_irq(n);
371:     }
372:
373:     printk("\n");
374:     printk("*** Safe to Power Off ***\n");
375:     printk("          -or-\n");
376:     printk("*** Press Any Key to Reboot ***\n");
377:     any_key_to_reboot = 1;
378:
379:     /* put all processes to sleep and reset all pending signals */
380:     FOR_EACH_PROCESS(p) {
381:         p->state = PROC_SLEEPING;
382:         p->sigpending = 0;
383:     }
384:
385:     /* enable keyboard only */

```

**kernel/main.c**

Page 7/7

```
386:         STI ();
387:         enable_irq(KEYBOARD_IRQ);
388:
389:         /* switch to IDLE process */
390:         if(current) {
391:             do_sched();
392:         }
393:
394:         cpu_idle();
395: }
```

**kernel/Makefile**

Page 1/1

```
1: # fiwix/kernel/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = boot.o core386.o main.o init.o gdt.o idt.o syscalls.o pic.o pit.o \
13:     traps.o cpu.o cmos.o timer.o sched.o sleep.o signal.o process.o
14:
15: kernel: $(OBJS)
16:     $(LD) $(LDFLAGS) -r $(OBJS) -o kernel.o
17:
18: clean:
19:     rm -f *.o
20:
```

## kernel/pic.c

Page 1/4

```

1: /*
2:  * fiwix/kernel/pic.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/config.h>
11: #include <fiwix/limits.h>
12: #include <fiwix/errno.h>
13: #include <fiwix/pic.h>
14: #include <fiwix/stdio.h>
15: #include <fiwix/string.h>
16: #include <fiwix/sigcontext.h>
17: #include <fiwix/sleep.h>
18:
19: /* interrupt vector base addresses */
20: #define IRQ0_ADDR      0x20
21: #define IRQ8_ADDR     0x28
22:
23: struct interrupt *irq_table[NR_IRQS];
24: static struct bh *bh_table = NULL;
25: unsigned int intr_count = 0;
26:
27: /*
28:  * This sends the command OCW3 to PIC (master or slave) to obtain the register
29:  * values. Slave is chained and represents IRQs 8-15. Master represents IRQs
30:  * 0-7, with IRQ 2 being the chain.
31:  */
32: static unsigned short int pic_get_irq_reg(int ocw3)
33: {
34:     outport_b(PIC_MASTER, ocw3);
35:     outport_b(PIC_SLAVE, ocw3);
36:     return (inport_b(PIC_SLAVE) << 8) | inport_b(PIC_MASTER);
37: }
38:
39: int register_irq(int num, struct interrupt *new_irq)
40: {
41:     struct interrupt **irq;
42:
43:     if(num < 0 || num >= NR_IRQS) {
44:         return -EINVAL;
45:     }
46:
47:     irq = &irq_table[num];
48:     if(*irq) {
49:         do {
50:             if(*irq == new_irq) {
51:                 printk("WARNING: %s(): interrupt %d already regi
stered!\n", __FUNCTION__, num);
52:                 return -EINVAL;
53:             }
54:             irq = &(*irq)->next;
55:         } while(*irq);
56:     }
57:     *irq = new_irq;
58:     new_irq->ticks = 0;
59:     return 0;
60: }
61:
62: int unregister_irq(int num, struct interrupt *old_irq)
63: {
64:     struct interrupt **irq, *prev_irq;
65:
66:     if(num < 0 || num >= NR_IRQS) {

```

```

67:         return -EINVAL;
68:     }
69:
70:     irq = &irq_table[num];
71:     prev_irq = NULL;
72:
73:     if(*irq) {
74:         do {
75:             if(*irq == old_irq) {
76:                 if((*irq)->next) {
77:                     printk("WARNING: %s(): cannot unregister
interrupt %d.\n", __FUNCTION__, num);
78:                     return -EINVAL;
79:                 }
80:                 *irq = NULL;
81:                 if(prev_irq) {
82:                     prev_irq->next = NULL;
83:                 }
84:                 break;
85:             }
86:             prev_irq = *irq;
87:             irq = &(*irq)->next;
88:         } while(*irq);
89:     }
90:     return 0;
91: }
92:
93: void add_bh(struct bh *new)
94: {
95:     unsigned long int flags;
96:     struct bh **b;
97:
98:     SAVE_FLAGS(flags); CLI();
99:
100:    b = &bh_table;
101:    if(*b) {
102:        do {
103:            b = &(*b)->next;
104:        } while(*b);
105:    }
106:    *b = new;
107:
108:    RESTORE_FLAGS(flags);
109: }
110:
111: void enable_irq(int irq)
112: {
113:     int addr;
114:
115:     addr = (irq > 7) ? PIC_SLAVE + DATA : PIC_MASTER + DATA;
116:     irq &= 0x0007;
117:
118:     outport_b(addr, inport_b(addr) & ~(1 << irq));
119: }
120:
121: void disable_irq(int irq)
122: {
123:     int addr;
124:
125:     addr = (irq > 7) ? PIC_SLAVE + DATA : PIC_MASTER + DATA;
126:     irq &= 0x0007;
127:
128:     outport_b(addr, inport_b(addr) | (1 << irq));
129: }
130:
131: /* each ISR points to this function */
132: void irq_handler(int num, struct sigcontext sc)

```

```

133: {
134:     struct interrupt *irq;
135:     int real;
136:
137:     if(num == -1) {
138:         printk("Unknown IRQ received!\n");
139:         return;
140:     }
141:
142:     irq = irq_table[num];
143:
144:     /* spurious interrupt treatment */
145:     if(!irq) {
146:         real = pic_get_irq_reg(PIC_READ_ISR);
147:         if(!real) {
148:             /*
149:              * If IRQ was not real and came from slave, then send
150:              * an EOI to master because it doesn't know if the IRQ
151:              * was a spurious interrupt from slave.
152:              */
153:             if(num > 7) {
154:                 outport_b(PIC_MASTER, EOI);
155:             }
156:             if(kstat.sirqs < MAX_SPU_NOTICES) {
157:                 printk("WARNING: spurious interrupt detected (un
registered IRQ %d).\n", num);
158:             } else if(kstat.sirqs == MAX_SPU_NOTICES) {
159:                 printk("WARNING: too many spurious interrupts; n
ot logging any more.\n");
160:             }
161:             kstat.sirqs++;
162:             return;
163:         }
164:         if(num > 7) {
165:             outport_b(PIC_SLAVE, EOI);
166:         }
167:         outport_b(PIC_MASTER, EOI);
168:         return;
169:     }
170:
171:     /* all handlers execute with interrupts disabled */
172:     disable_irq(num);
173:
174:     if(num > 7) {
175:         outport_b(PIC_SLAVE, EOI);
176:     }
177:     outport_b(PIC_MASTER, EOI);
178:
179:     kstat.irqs++;
180:     irq->ticks++;
181:     do {
182:         irq->handler(num, &sc);
183:         irq = irq->next;
184:     } while(irq);
185:
186:     enable_irq(num);
187: }
188:
189: /* do bottom halves (interrupts are (FIXME: not yet) enabled) */
190: void do_bh(void)
191: {
192:     struct bh *b;
193:     void (*fn) (void);
194:
195:     if(!lock_area(AREA_BH)) {
196:         //FIXME: STI();
197:         if((b = bh_table)) {

```

## kernel/pic.c

Page 4/4

```
198:             do {
199:                 if(b->flags & BH_ACTIVE) {
200:                     b->flags &= ~BH_ACTIVE;
201:                     fn = b->fn;
202:                     (*fn) ();
203:                 }
204:                 b = b->next;
205:             } while(b);
206:         }
207:         //FIXME: CLI();
208:         unlock_area(AREA_BH);
209:     }
210: }
211:
212: void pic_init(void)
213: {
214:     memset_b(irq_table, NULL, sizeof(irq_table));
215:
216:     /* remap interrupts for PIC1 */
217:     outport_b(PIC_MASTER, ICW1_RESET);
218:     outport_b(PIC_MASTER + DATA, IRQ0_ADDR);          /* ICW2 */
219:     outport_b(PIC_MASTER + DATA, 1 << CASCADE_IRQ); /* ICW3 */
220:     outport_b(PIC_MASTER + DATA, ICW4_8086EOI);
221:
222:     /* remap interrupts for PIC2 */
223:     outport_b(PIC_SLAVE, ICW1_RESET);
224:     outport_b(PIC_SLAVE + DATA, IRQ8_ADDR);          /* ICW2 */
225:     outport_b(PIC_SLAVE + DATA, CASCADE_IRQ);       /* ICW3 */
226:     outport_b(PIC_SLAVE + DATA, ICW4_8086EOI);
227:
228:     /* mask all IRQs except cascade */
229:     outport_b(PIC_MASTER + DATA, ~(1 << CASCADE_IRQ));
230:
231:     /* mask all IRQs */
232:     outport_b(PIC_SLAVE + DATA, OCW1);
233: }
```



**kernel/pit.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/pit.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/pit.h>
10:
11: void pit_beep_on(void)
12: {
13:     outport_b(MODEREG, SEL_CHAN2 | LSB_MSB | SQUARE_WAVE | BINARY_CTR);
14:     outport_b(CHANNEL2, (OSCIL / BEEP_FREQ) & 0xFF);      /* LSB */
15:     outport_b(CHANNEL2, (OSCIL / BEEP_FREQ) >> 8);      /* MSB */
16:     outport_b(PS2_SYSCtrl_B, inport_b(PS2_SYSCtrl_B) | ENABLE_SDATA | ENABLE
_LE_TMR2G);
17: }
18:
19: void pit_beep_off(unsigned int unused)
20: {
21:     outport_b(PS2_SYSCtrl_B, inport_b(PS2_SYSCtrl_B) & ~(ENABLE_SDATA | ENAB
LE_TMR2G));
22: }
23:
24: void pit_init(unsigned short int hertz)
25: {
26:     outport_b(MODEREG, SEL_CHAN0 | LSB_MSB | RATE_GEN | BINARY_CTR);
27:     outport_b(CHANNEL0, (OSCIL / hertz) & 0xFF);      /* LSB */
28:     outport_b(CHANNEL0, (OSCIL / hertz) >> 8);      /* MSB */
29: }
```

## kernel/process.c

Page 1/7

```

1:  /*
2:  *  fiwix/kernel/process.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/kernel.h>
9:  #include <fiwix/mm.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/process.h>
12: #include <fiwix/timer.h>
13: #include <fiwix/sched.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: struct proc *proc_table;
19: struct proc *current;
20:
21: struct proc *proc_pool_head;
22: struct proc *proc_table_head;
23: struct proc *proc_table_tail;
24: unsigned int free_proc_slots = 0;
25:
26: static struct resource slot_resource = { NULL, NULL };
27: static struct resource pid_resource = { NULL, NULL };
28:
29: int nr_processes = 0;
30: __pid_t lastpid = 0;
31:
32: int kill_pid(__pid_t pid, __sigset_t signum)
33: {
34:     struct proc *p;
35:
36:     FOR_EACH_PROCESS(p) {
37:         if(p->pid == pid && (p->state && p->state != PROC_ZOMBIE)) {
38:             return send_sig(p, signum);
39:         }
40:     }
41:     return -ESRCH;
42: }
43:
44: int kill_pgrp(__pid_t pgid, __sigset_t signum)
45: {
46:     struct proc *p;
47:     int found;
48:
49:     found = 0;
50:     FOR_EACH_PROCESS(p) {
51:         if(p->pgid == pgid && (p->state && p->state != PROC_ZOMBIE)) {
52:             send_sig(p, signum);
53:             found = 1;
54:         }
55:     }
56:
57:     if(!found) {
58:         return -ESRCH;
59:     }
60:     return 0;
61: }
62:
63: /* sum up child (and its children) statistics */
64: void add_crusage(struct proc *p, struct rusage *cru)
65: {
66:     cru->ru_utime.tv_sec = p->usage.ru_utime.tv_sec + p->cusage.ru_utime.tv_
sec;

```

## kernel/process.c

Page 2/7

```

67:         cru->ru_utime.tv_usec = p->usage.ru_utime.tv_usec + p->cusage.ru_utime.t
v_usec;
68:         if(cru->ru_utime.tv_usec >= 1000000) {
69:             cru->ru_utime.tv_sec++;
70:             cru->ru_utime.tv_usec -= 1000000;
71:         }
72:         cru->ru_stime.tv_sec = p->usage.ru_stime.tv_sec + p->cusage.ru_stime.tv_
sec;
73:         cru->ru_stime.tv_usec = p->usage.ru_stime.tv_usec + p->cusage.ru_stime.t
v_usec;
74:         if(cru->ru_stime.tv_usec >= 1000000) {
75:             cru->ru_stime.tv_sec++;
76:             cru->ru_stime.tv_usec -= 1000000;
77:         }
78:         cru->ru_maxrss = p->usage.ru_maxrss + p->cusage.ru_maxrss;
79:         cru->ru_ixrss = p->usage.ru_ixrss + p->cusage.ru_ixrss;
80:         cru->ru_idrss = p->usage.ru_idrss + p->cusage.ru_idrss;
81:         cru->ru_isrss = p->usage.ru_isrss + p->cusage.ru_isrss;
82:         cru->ru_minflt = p->usage.ru_minflt + p->cusage.ru_minflt;
83:         cru->ru_majflt = p->usage.ru_majflt + p->cusage.ru_majflt;
84:         cru->ru_nswap = p->usage.ru_nswap + p->cusage.ru_nswap;
85:         cru->ru_inblock = p->usage.ru_inblock + p->cusage.ru_inblock;
86:         cru->ru_oublock = p->usage.ru_oublock + p->cusage.ru_oublock;
87:         cru->ru_msgsnd = p->usage.ru_msgsnd + p->cusage.ru_msgsnd;
88:         cru->ru_msgrcv = p->usage.ru_msgrcv + p->cusage.ru_msgrcv;
89:         cru->ru_nsignals = p->usage.ru_nsignals + p->cusage.ru_nsignals;
90:         cru->ru_nvcsw = p->usage.ru_nvcsw + p->cusage.ru_nvcsw;
91:         cru->ru_nivcsw = p->usage.ru_nivcsw + p->cusage.ru_nivcsw;
92:     }
93:
94: void get_rusage(struct proc *p, struct rusage *ru)
95: {
96:     struct rusage cru;
97:
98:     /*
99:      * Conforms with SUSv3 which specifies that if SIGCHLD is being ignored
100:      * then the child statistics should not be added to the values returned
101:      * by RUSAGE_CHILDREN.
102:      */
103:     if(current->sigaction[SIGCHLD - 1].sa_handler == SIG_IGN) {
104:         return;
105:     }
106:
107:     add_crusage(p, &cru);
108:     memcpy_b(ru, &cru, sizeof(struct rusage));
109: }
110:
111: /* add child statistics to parent */
112: void add_rusage(struct proc *p)
113: {
114:     struct rusage cru;
115:
116:     add_crusage(p, &cru);
117:     current->cusage.ru_utime.tv_sec += cru.ru_utime.tv_sec;
118:     current->cusage.ru_utime.tv_usec += cru.ru_utime.tv_usec;
119:     if(current->cusage.ru_utime.tv_usec >= 1000000) {
120:         current->cusage.ru_utime.tv_sec++;
121:         current->cusage.ru_utime.tv_usec -= 1000000;
122:     }
123:     current->cusage.ru_stime.tv_sec += cru.ru_stime.tv_sec;
124:     current->cusage.ru_stime.tv_usec += cru.ru_stime.tv_usec;
125:     if(current->cusage.ru_stime.tv_usec >= 1000000) {
126:         current->cusage.ru_stime.tv_sec++;
127:         current->cusage.ru_stime.tv_usec -= 1000000;
128:     }
129:     current->cusage.ru_maxrss += cru.ru_maxrss;
130:     current->cusage.ru_ixrss += cru.ru_ixrss;

```

## kernel/process.c

Page 3/7

```

131:     current->cusage.ru_idrss += cru.ru_idrss;
132:     current->cusage.ru_isrss += cru.ru_isrss;
133:     current->cusage.ru_minflt += cru.ru_minflt;
134:     current->cusage.ru_majflt += cru.ru_majflt;
135:     current->cusage.ru_nswap += cru.ru_nswap;
136:     current->cusage.ru_inblock += cru.ru_inblock;
137:     current->cusage.ru_oublock += cru.ru_oublock;
138:     current->cusage.ru_msgsnd += cru.ru_msgsnd;
139:     current->cusage.ru_msgrcv += cru.ru_msgrcv;
140:     current->cusage.ru_nsignals += cru.ru_nsignals;
141:     current->cusage.ru_nvcsw += cru.ru_nvcsw;
142:     current->cusage.ru_nivcsw += cru.ru_nivcsw;
143: }
144:
145: struct proc * get_next_zombie(struct proc *parent)
146: {
147:     struct proc *p;
148:
149:     if(proc_table_head == NULL) {
150:         PANIC("process table is empty!\n");
151:     }
152:
153:     FOR_EACH_PROCESS(p) {
154:         if(p->ppid == parent->pid && p->state == PROC_ZOMBIE) {
155:             return p;
156:         }
157:     }
158:
159:     return NULL;
160: }
161:
162: __pid_t remove_zombie(struct proc *p)
163: {
164:     __pid_t pid;
165:
166:     pid = p->pid;
167:     kfree(p->tss.esp0);           current->rss--;
168:     kfree(P2V(p->tss.cr3));      current->rss--;
169:     release_proc(p);
170:     current->children--;
171:     return pid;
172: }
173:
174: /*
175:  * An orphaned process group is a process group in which the parent of every
176:  * member is either itself a member of the group or is not a member of the
177:  * group's session.
178:  */
179: int is_orphaned_pgrp(__pid_t pgid)
180: {
181:     struct proc *p, *pp;
182:
183:     FOR_EACH_PROCESS(p) {
184:         if(p->pgid != pgid) {
185:             continue;
186:         }
187:         if(p->state && p->state != PROC_ZOMBIE) {
188:             pp = get_proc_by_pid(p->ppid);
189:             /* return if only one is found that breaks the rule */
190:             if(pp->pgid != pgid || pp->sid == p->sid) {
191:                 return 0;
192:             }
193:         }
194:     }
195:     return 1;
196: }
197:

```

## kernel/process.c

Page 4/7

```

198: struct proc * get_proc_free(void)
199: {
200:     struct proc *p = NULL;
201:
202:     if(free_proc_slots <= SAFE_SLOTS && !IS_SUPERUSER) {
203:         printk("WARNING: %s(): the remaining slots are only for root use
r!\n", __FUNCTION__);
204:         return NULL;
205:     }
206:
207:     lock_resource(&slot_resource);
208:
209:     if(proc_pool_head) {
210:
211:         /* get (remove) a process slot from the free list */
212:         p = proc_pool_head;
213:         proc_pool_head = proc_pool_head->next;
214:
215:         free_proc_slots--;
216:     } else {
217:         printk("WARNING: %s(): no more slots free in proc table!\n", __F
UNCTION__);
218:     }
219:
220:     unlock_resource(&slot_resource);
221:     return p;
222: }
223:
224: void release_proc(struct proc *p)
225: {
226:     lock_resource(&slot_resource);
227:
228:     /* remove a process from the proc_table */
229:     if(p == proc_table_tail) {
230:         if(proc_table_head == proc_table_tail) {
231:             proc_table_head = proc_table_tail = NULL;
232:         } else {
233:             proc_table_tail = proc_table_tail->prev;
234:             proc_table_tail->next = NULL;
235:         }
236:     } else {
237:         p->prev->next = p->next;
238:         p->next->prev = p->prev;
239:     }
240:
241:     /* initialize and put a process slot back in the free list */
242:     memset_b(p, NULL, sizeof(struct proc));
243:     p->next = proc_pool_head;
244:     proc_pool_head = p;
245:     free_proc_slots++;
246:
247:     unlock_resource(&slot_resource);
248: }
249:
250: int get_unused_pid(void)
251: {
252:     short int loop;
253:     struct proc *p;
254:
255:     loop = 0;
256:     lock_resource(&pid_resource);
257:
258: loop:
259:     lastpid++;
260:     if(lastpid > MAX_PID_VALUE) {
261:         loop++;
262:         lastpid = INIT;

```

## kernel/process.c

Page 5/7

```

263:     }
264:     if(loop > 1) {
265:         printk("WARNING: %s(): system ran out of PID numbers!\n");
266:         return 0;
267:     }
268:     FOR_EACH_PROCESS(p) {
269:         if(p->state != PROC_UNUSED) {
270:             /*
271:              * Make sure the kernel never reuses active pid, ppid
272:              * or sid values.
273:              */
274:             if(lastpid == p->pid || lastpid == p->ppid || lastpid ==
p->sid) {
275:                 goto loop;
276:             }
277:         }
278:     }
279:     unlock_resource(&pid_resource);
280:     return lastpid;
281: }
282: }
283:
284: struct proc * get_proc_by_pid(__pid_t pid)
285: {
286:     struct proc *p;
287:
288:     FOR_EACH_PROCESS(p) {
289:         if(p->state && p->pid == pid) {
290:             return p;
291:         }
292:     }
293:
294:     PANIC("would return NULL! (current->pid=%d ppid=%d)\n", current->pid, pid
);
295:     return NULL;
296: }
297:
298: int get_new_user_fd(int fd)
299: {
300:     int n;
301:
302:     for(n = fd; n < OPEN_MAX && n < current->rlim[RLIMIT_NOFILE].rlim_cur; n
++) {
303:         if(current->fd[n] == 0) {
304:             current->fd[n] = -1;
305:             current->fd_flags[n] = 0;
306:             return n;
307:         }
308:     }
309:
310:     return -EMFILE;
311: }
312:
313: void release_user_fd(int ufd)
314: {
315:     current->fd[ufd] = 0;
316: }
317:
318: struct proc * kernel_process(int (*fn)(void))
319: {
320:     struct proc *p;
321:
322:     p = get_proc_free();
323:     proc_slot_init(p);
324:     p->pid = get_unused_pid();
325:     p->ppid = 0;
326:     p->flags |= PF_KPROC;

```

## kernel/process.c

Page 6/7

```

327:     p->priority = DEF_PRIORITY;
328:     if(!(p->tss.esp0 = kmalloc())) {
329:         release_proc(p);
330:         return NULL;
331:     }
332:     p->tss.esp0 += PAGE_SIZE - 4;
333:     p->rss++;
334:     p->tss.cr3 = (unsigned int)kpage_dir;
335:     p->tss.eip = (unsigned int)fn;
336:     p->tss.esp = p->tss.esp0;
337:     sprintk(p->pidstr, "%d", p->pid);
338:     p->state = PROC_RUNNING;
339:     return p;
340: }
341:
342: void proc_slot_init(struct proc *p)
343: {
344:     int n;
345:
346:     /* insert process at the end of proc_table */
347:     lock_resource(&slot_resource);
348:     if(proc_table_head == NULL) {
349:         p->prev = NULL;
350:         p->next = NULL;
351:         proc_table_head = proc_table_tail = p;
352:     } else {
353:         p->prev = proc_table_tail;
354:         p->next = NULL;
355:         proc_table_tail->next = p;
356:         proc_table_tail = p;
357:     }
358:     p->sleep_prev = p->sleep_next = NULL;
359:     unlock_resource(&slot_resource);
360:
361:     memset_b(&p->tss, NULL, sizeof(struct i386tss));
362:     p->tss.io_bitmap_addr = (unsigned int)&p->io_bitmap;
363:
364:     /*
365:      * At the moment, all io_bitmap bits are setup to 0, which means full
366:      * access. This must be changed to 1 once we have fixed the ioperm()
367:      * system call.
368:      */
369:     for(n = 0; n < IO_BITMAP_SIZE + 1; n++) {
370:         p->io_bitmap[n] = 0;    /* FIXME: change it to 1 */
371:     }
372:     p->state = PROC_IDLE;
373: }
374:
375: void proc_init(void)
376: {
377:     int n;
378:     struct proc *p;
379:
380:     memset_b(proc_table, NULL, proc_table_size);
381:
382:     /* free list initialization */
383:     proc_pool_head = NULL;
384:     n = (proc_table_size / sizeof(struct proc)) - 1;
385:     do {
386:         p = &proc_table[n];
387:
388:         /* fill the free list */
389:         p->next = proc_pool_head;
390:         proc_pool_head = p;
391:         free_proc_slots++;
392:     } while(n--);
393:     proc_table_head = proc_table_tail = NULL;

```

**kernel/process.c**

Page 7/7

```
394:
395:     /* slot 0 is for the IDLE process */
396:     current = get_proc_free();
397:     proc_slot_init(current);
398: }
```



## kernel/sched.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/sched.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/const.h>
11: #include <fiwix/sched.h>
12: #include <fiwix/process.h>
13: #include <fiwix/segments.h>
14: #include <fiwix/timer.h>
15: #include <fiwix/pic.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: extern struct seg_desc gdt[NR_GDT_ENTRIES];
20: int need_resched = 0;
21:
22: static void context_switch(struct proc *next)
23: {
24:     struct proc *prev;
25:
26:     kstat.ctxt++;
27:     prev = current;
28:     set_tss(next);
29:     current = next;
30:     do_switch(&prev->tss.esp, &prev->tss.eip, next->tss.esp, next->tss.eip,
next->tss.cr3, TSS);
31: }
32:
33: void set_tss(struct proc *p)
34: {
35:     struct seg_desc *g;
36:
37:     g = &gdt[TSS / sizeof(struct seg_desc)];
38:
39:     g->sd_lobase = (unsigned int)p;
40:     g->sd_loflags = SD_TSSPRESENT;
41:     g->sd_hibase = (char)(((unsigned int)p) >> 24);
42: }
43:
44: /* Round Robin algorithm */
45: void do_sched(void)
46: {
47:     int count;
48:     struct proc *p, *selected;
49:
50:     /* let the current running process consume its time slice */
51:     if(!need_resched && current->state == PROC_RUNNING && current->cpu_count
> 0) {
52:         return;
53:     }
54:
55:     need_resched = 0;
56:     for(;;) {
57:         count = -1;
58:         selected = &proc_table[IDLE];
59:
60:         FOR_EACH_PROCESS(p) {
61:             if(p->state == PROC_RUNNING && p->cpu_count > count) {
62:                 count = p->cpu_count;
63:                 selected = p;
64:             }
65:         }

```

## kernel/sched.c

Page 2/2

```
66:             if(count) {
67:                 break;
68:             }
69:
70:             /* reassigns new quantum to all processes */
71:             FOR_EACH_PROCESS(p) {
72:                 if(p->state) {
73:                     p->cpu_count = p->priority;
74:                 }
75:             }
76:         }
77:         if(current != selected) {
78:             context_switch(selected);
79:         }
80:     }
81:
82: void sched_init(void)
83: {
84:     get_system_time();
85:
86:     /* this should be more unpredictable */
87:     kstat.random_seed = CURRENT_TIME;
88: }
```

## kernel/signal.c

Page 1/4

```

1: /*
2:  * fiwix/kernel/signal.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/process.h>
12: #include <fiwix/signal.h>
13: #include <fiwix/sigcontext.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/sched.h>
16: #include <fiwix/syscalls.h>
17: #include <fiwix/mm.h>
18: #include <fiwix/stdio.h>
19: #include <fiwix/string.h>
20:
21: int send_sig(struct proc *p, __sigset_t signum)
22: {
23:     struct proc *z;
24:
25:     if(signum > NSIG || !p) {
26:         return -EINVAL;
27:     }
28:
29:     if(!IS_SUPERUSER && current->euid != p->euid && current->sid != p->sid)
{
30:         return -EPERM;
31:     }
32:
33:     /* kernel processes can't receive signals */
34:     if(p->flags & PF_KPROC) {
35:         return 0;
36:     }
37:
38:     switch(signum) {
39:         case 0:
40:             return 0;
41:         case SIGKILL:
42:         case SIGCONT:
43:             if(p->state == PROC_STOPPED) {
44:                 p->state = PROC_RUNNING;
45:             }
46:             /* discard all pending stop signals */
47:             p->sigpending &= SIG_MASK(SIGSTOP);
48:             p->sigpending &= SIG_MASK(SIGTSTP);
49:             p->sigpending &= SIG_MASK(SIGTTIN);
50:             p->sigpending &= SIG_MASK(SIGTTOU);
51:             break;
52:         case SIGSTOP:
53:         case SIGTSTP:
54:         case SIGTTIN:
55:         case SIGTTOU:
56:             /* discard all pending SIGCONT signals */
57:             p->sigpending &= SIG_MASK(SIGCONT);
58:             break;
59:         default:
60:             break;
61:     }
62:
63:     /*
64:      * Force some signals, that a process cannot ignore, by changing its
65:      * signal disposition to SIG_DFL.
66:      */

```

## kernel/signal.c

Page 2/4

```

67:         switch(signum) {
68:             case SIGFPE:
69:             case SIGSEGV:
70:                 if(p->sigaction[signum - 1].sa_handler == SIG_IGN) {
71:                     p->sigaction[signum - 1].sa_handler = SIG_DFL;
72:                 }
73:                 break;
74:         }
75:
76:         if(p->sigaction[signum - 1].sa_handler == SIG_IGN && signum != SIGCHLD)
{
77:             return 0;
78:         }
79:
80:         /* SIGCHLD is ignored by default */
81:         if(p->sigaction[signum - 1].sa_handler == SIG_DFL) {
82:             /*
83:              * INIT process is special, it only gets signals that have the
84:              * signal handler installed. This avoids to bring down the
85:              * system accidentally.
86:              */
87:             if(p->pid == INIT) {
88:                 return 0;
89:             }
90:             if(signum == SIGCHLD) {
91:                 return 0;
92:             }
93:         }
94:
95:         /* if SIGCHLD is ignored reap its children (prevent zombies) */
96:         if(p->sigaction[signum - 1].sa_handler == SIG_IGN) {
97:             if(signum == SIGCHLD) {
98:                 while((z = get_next_zombie(p))) {
99:                     remove_zombie(z);
100:                 }
101:                 return 0;
102:             }
103:         }
104:
105:         p->sigpending |= 1 << (signum - 1);
106:         p->usage.ru_nsignals++;
107:
108:         /* wake up the process only if that signal is not blocked */
109:         if(!(p->sigblocked & (1 << (signum - 1)))) {
110:             wakeup_proc(p);
111:         }
112:
113:         return 0;
114:     }
115:
116: int issig(void)
117: {
118:     __sigset_t signum;
119:     unsigned int mask;
120:     struct proc *p;
121:
122:     if(!(current->sigpending & ~current->sigblocked)) {
123:         return 0;
124:     }
125:
126:     for(signum = 1, mask = 1; signum < NSIG; signum++, mask <=< 1) {
127:         if(current->sigpending & mask) {
128:             if(signum == SIGCHLD) {
129:                 if(current->sigaction[signum - 1].sa_handler ==
SIG_IGN) {
130:                     /* this process ignores SIGCHLD */
131:                     while((p = get_next_zombie(current))) {

```

## kernel/signal.c

Page 3/4

```

132:                                     remove_zombie(p);
133:                                     }
134:                                     } else {
135:                                     if(current->sigaction[signum - 1].sa_han
dler != SIG_DFL) {
136:                                     return signum;
137:                                     }
138:                                     }
139:                                     } else {
140:                                     if(current->sigaction[signum - 1].sa_handler !=
SIG_IGN) {
141:                                     return signum;
142:                                     }
143:                                     }
144:                                     current->sigpending &= ~mask;
145:                                     }
146:                                     }
147:                                     return 0;
148: }
149:
150: void psig(unsigned int stack)
151: {
152:     int len;
153:     __sigset_t signum;
154:     unsigned int mask;
155:     struct sigcontext *sc;
156:
157:     sc = (struct sigcontext *)stack;
158:     for(signum = 1, mask = 1; signum < NSIG; signum++, mask <= 1) {
159:         if(current->sigpending & mask) {
160:             current->sigpending &= ~mask;
161:
162:             if((unsigned int)current->sigaction[signum - 1].sa_handl
er) {
163:                 current->sigexecuting = mask;
164:                 if(!(current->sigaction[signum - 1].sa_flags & S
A_NODEFER)) {
165:                     current->sigblocked |= mask;
166:                 }
167:
168:                 /* save the current sigcontext */
169:                 memcpy_b(&current->sc[signum - 1], sc, sizeof(st
ruct sigcontext));
170:
171:                 /* setup the jump to the user signal handler */
172:                 len = ((int)end_sighandler_trampoline - (int)sig
handler_trampoline);
173:                 sc->oldesp -= len;
174:                 sc->oldesp -= 4;
175:                 sc->oldesp &= ~3;          /* round up */
176:                 memcpy_b((void *)sc->oldesp, sighandler_trampoline,
len);
177:                 sc->ecx = (unsigned int)current->sigaction[signu
m - 1].sa_handler;
178:                 sc->eax = signum;
179:                 sc->eip = sc->oldesp;
180:
181:                 if(current->sigaction[signum - 1].sa_flags & SA_
RESETHAND) {
182:                     current->sigaction[signum - 1].sa_handle
r = SIG_DFL;
183:                 }
184:                 return;
185:             }
186:             if(current->sigaction[signum - 1].sa_handler == SIG_DFL)
{
187:                 switch(signum) {
188:                     case SIGCONT:

```

## kernel/signal.c

Page 4/4

```

188:                                     current->state = PROC_RUNNING;
189:                                     need_resched = 1;
190:                                     break;
191:     case SIGSTOP:
192:     case SIGTSTP:
193:     case SIGTTIN:
194:     case SIGTTOU:
195:                                     current->exit_code = signum;
196:                                     current->state = PROC_STOPPED;
197:                                     if(!(current->sigaction[signum -
1] .sa_flags & SA_NOCLDSTOP)) {
198:                                     send_sig(get_proc_by_pid
(current->ppid), SIGCHLD);
199:                                     /* needed for job contro
1 */
200:                                     wakeup(&sys_wait4);
201:                                     }
202:                                     need_resched = 1;
203:                                     break;
204:     case SIGCHLD:
205:                                     break;
206:     default:
207:                                     do_exit(signum);
208:                                     }
209:                                     }
210:                                     }
211:     }
212:
213:     /* coming from a system call that needs to be restarted */
214:     if(sc->err > 0) {
215:         if(sc->eax == -ERESTART) {
216:             sc->eax = sc->err;        /* syscall was saved in 'err' */
217:             sc->eip -= 2;            /* point again to 'int 0x80' */
218:         }
219:     }
220: }

```

## kernel/sleep.c

Page 1/3

```

1: /*
2:  * fiwix/kernel/sleep.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/limits.h>
11: #include <fiwix/sleep.h>
12: #include <fiwix/sched.h>
13: #include <fiwix/signal.h>
14: #include <fiwix/process.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: #define NR_BUCKETS          (NR_PROCS * 10) / 100    /* 10% of NR_PROCS */
19: #define SLEEP_HASH(addr)  ((addr) % NR_BUCKETS)
20:
21: struct proc *sleep_hash_table[NR_BUCKETS];
22: static unsigned int area = 0;
23:
24: int sleep(void *address, int state)
25: {
26:     unsigned long int flags;
27:     struct proc **h;
28:     int signum, i;
29:
30:     /* return if it has signals */
31:     if(state == PROC_INTERRUPTIBLE) {
32:         if((signum = issig())) {
33:             return signum;
34:         }
35:     }
36:
37:     if(current->state == PROC_SLEEPING) {
38:         printk("WARNING: %s(): process with pid '%d' is already sleeping
!\n", __FUNCTION__, current->pid);
39:         return 0;
40:     }
41:
42:     SAVE_FLAGS(flags); CLI();
43:     i = SLEEP_HASH((unsigned int)address);
44:     h = &sleep_hash_table[i];
45:
46:     /* insert process in the head */
47:     if(!*h) {
48:         *h = current;
49:         (*h)->sleep_prev = (*h)->sleep_next = NULL;
50:     } else {
51:         current->sleep_prev = NULL;
52:         current->sleep_next = *h;
53:         (*h)->sleep_prev = current;
54:         *h = current;
55:     }
56:     current->sleep_address = address;
57:     current->state = PROC_SLEEPING;
58:
59:     do_sched();
60:
61:     signum = 0;
62:     if(state == PROC_INTERRUPTIBLE) {
63:         signum = issig();
64:     }
65:
66:     RESTORE_FLAGS(flags);

```

## kernel/sleep.c

Page 2/3

```

67:         return signum;
68:     }
69:
70: void wakeup(void *address)
71: {
72:     unsigned long int flags;
73:     struct proc **h;
74:     int i;
75:
76:     SAVE_FLAGS(flags); CLI();
77:     i = SLEEP_HASH((unsigned int)address);
78:     h = &sleep_hash_table[i];
79:
80:     while(*h) {
81:         if((*h)->sleep_address == address) {
82:             (*h)->sleep_address = NULL;
83:             (*h)->state = PROC_RUNNING;
84:             need_resched = 1;
85:             if((*h)->sleep_next) {
86:                 (*h)->sleep_next->sleep_prev = (*h)->sleep_prev;
87:             }
88:             if((*h)->sleep_prev) {
89:                 (*h)->sleep_prev->sleep_next = (*h)->sleep_next;
90:             }
91:             if(h == &sleep_hash_table[i]) { /* if it's the head */
92:                 *h = (*h)->sleep_next;
93:                 continue;
94:             }
95:         }
96:         if(*h) {
97:             h = &(*h)->sleep_next;
98:         }
99:     }
100:    RESTORE_FLAGS(flags);
101: }
102:
103: void wakeup_proc(struct proc *p)
104: {
105:     unsigned long int flags;
106:     struct proc **h;
107:     int i;
108:
109:     if(p->state != PROC_SLEEPING && p->state != PROC_STOPPED) {
110:         return;
111:     }
112:
113:     SAVE_FLAGS(flags); CLI();
114:
115:     /* stopped processes don't have sleep address */
116:     if(p->sleep_address) {
117:         if(p->sleep_next) {
118:             p->sleep_next->sleep_prev = p->sleep_prev;
119:         }
120:         if(p->sleep_prev) {
121:             p->sleep_prev->sleep_next = p->sleep_next;
122:         }
123:
124:         i = SLEEP_HASH((unsigned int)p->sleep_address);
125:         h = &sleep_hash_table[i];
126:         if(*h == p) { /* if it's the head */
127:             *h = (*h)->sleep_next;
128:         }
129:     }
130:     p->sleep_address = NULL;
131:     p->state = PROC_RUNNING;
132:     need_resched = 1;
133:

```



## kernel/sleep.c

Page 3/3

```
134:     RESTORE_FLAGS(flags);
135:     return;
136: }
137:
138: void lock_resource(struct resource *resource)
139: {
140:     unsigned long int flags;
141:
142:     for(;;) {
143:         SAVE_FLAGS(flags); CLI();
144:         if(resource->locked) {
145:             resource->wanted = 1;
146:             RESTORE_FLAGS(flags);
147:             sleep(resource, PROC_UNINTERRUPTIBLE);
148:         } else {
149:             break;
150:         }
151:     }
152:     resource->locked = 1;
153:     RESTORE_FLAGS(flags);
154: }
155:
156: void unlock_resource(struct resource *resource)
157: {
158:     unsigned long int flags;
159:
160:     SAVE_FLAGS(flags); CLI();
161:     resource->locked = 0;
162:     if(resource->wanted) {
163:         resource->wanted = 0;
164:         wakeup(resource);
165:     }
166:     RESTORE_FLAGS(flags);
167: }
168:
169: int lock_area(unsigned int type)
170: {
171:     unsigned long int flags;
172:     int retval;
173:
174:     SAVE_FLAGS(flags); CLI();
175:     retval = area & type;
176:     area |= type;
177:     RESTORE_FLAGS(flags);
178:
179:     return retval;
180: }
181:
182: int unlock_area(unsigned int type)
183: {
184:     unsigned long int flags;
185:     int retval;
186:
187:     SAVE_FLAGS(flags); CLI();
188:     retval = area & type;
189:     area &= ~type;
190:     RESTORE_FLAGS(flags);
191:
192:     return retval;
193: }
194:
195: void sleep_init(void)
196: {
197:     memset_b(sleep_hash_table, NULL, sizeof(sleep_hash_table));
198: }
```

## kernel/syscalls.c

Page 1/7

```

1:  /*
2:  * fiwix/kernel/syscalls.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/types.h>
9:  #include <fiwix/syscalls.h>
10: #include <fiwix/mm.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #endif /* __DEBUG__ */
17:
18: void free_name(const char *name)
19: {
20:     kfree((unsigned int)name);
21: }
22:
23: /*
24:  * This function has two objectives:
25:  * 1. to check the memory address validity of the char pointer supplied by the
26:  *    user, while at the same time limit its length to PAGE_SIZE (4096) bytes.
27:  * 2. to create a copy of 'filename' in the kernel data space before using it.
28:  */
29: int malloc_name(const char *filename, char **name)
30: {
31:     struct vma *vma;
32:     unsigned int start;
33:     short int n, len;
34:     char *b;
35:
36:     /*
37:      * Verifies if the 'vma' array of that process is not empty. It can
38:      * only be empty during the initialization of INIT, when it calls to
39:      * sys_execve and sys_open without having yet a proper setup.
40:      */
41:     if(current->vma[0].s_type != 0) {
42:         if(!filename) {
43:             return -EFAULT;
44:         }
45:         start = (unsigned int)filename;
46:         if(!(vma = find_vma_region(start))) {
47:             return -EFAULT;
48:         }
49:         if(!(vma->prot & PROT_READ)) {
50:             return -EFAULT;
51:         }
52:         len = MIN(vma->end - start, PAGE_SIZE);
53:         if(len < PAGE_SIZE) {
54:             if((vma = find_vma_region(vma->end))) {
55:                 if(vma->prot & PROT_READ) {
56:                     len = PAGE_SIZE;
57:                 }
58:             }
59:         }
60:     } else {
61:         len = PAGE_SIZE;
62:     }
63:     if(!(b = (char *)kmalloc())) {
64:         return -ENOMEM;
65:     }
66:     *name = b;
67:     for(n = 0; n < len; n++) {

```

## kernel/syscalls.c

Page 2/7

```
68:         if(!(*b = *filename)) {
69:             return 0;
70:         }
71:         b++;
72:         filename++;
73:     }
74:
75:     free_name(*name);
76:     return -ENAMETOOLONG;
77: }
78:
79: int check_user_permission(struct inode *i)
80: {
81:     if(!IS_SUPERUSER) {
82:         if(current->euid != i->i_uid) {
83:             return 1;
84:         }
85:     }
86:     return 0;
87: }
88:
89: int check_group(struct inode *i)
90: {
91:     int n;
92:     __gid_t gid;
93:
94:     if(current->flags & PF_USEREAL) {
95:         gid = current->gid;
96:     } else {
97:         gid = current->egid;
98:     }
99:
100:    if(i->i_gid == gid) {
101:        return 0;
102:    }
103:
104:    for(n = 0; n < NGROUPS_MAX; n++) {
105:        if(current->groups[n] == -1) {
106:            break;
107:        }
108:        if(current->groups[n] == i->i_gid) {
109:            return 0;
110:        }
111:    }
112:    return 1;
113: }
114:
115: int check_user_area(int type, const void *addr, unsigned int size)
116: {
117:     struct vma *vma;
118:     unsigned int start;
119:
120:     /*
121:      * Verifies if the 'vma' array of that process is not empty. It can
122:      * only be empty during the initialization of INIT, when it calls to
123:      * sys_execve and sys_open without having yet a proper setup.
124:      */
125:     if(current->vma[0].s_type != 0) {
126:         start = (unsigned int)addr;
127:         if(!(vma = find_vma_region(start))) {
128:             return -EFAULT;
129:         }
130:
131:         for(;;) {
132:             if(type == VERIFY_WRITE) {
133:                 if(!(vma->prot & PROT_WRITE)) {
134:                     return -EFAULT;

```

## kernel/syscalls.c

Page 3/7

```

135:         }
136:     } else {
137:         if(!(vma->prot & PROT_READ)) {
138:             return -EFAULT;
139:         }
140:     }
141:     if(start + size <= vma->end) {
142:         break;
143:     }
144:     if(!(vma = find_vma_region(vma->end))) {
145:         return -EFAULT;
146:     }
147: }
148: }
149:
150: return 0;
151: }
152:
153: int check_permission(int mask, struct inode *i)
154: {
155:     __uid_t uid;
156:
157:     if(current->flags & PF_USEREAL) {
158:         uid = current->uid;
159:     } else {
160:         uid = current->euid;
161:     }
162:
163:     if(uid == 0) {
164:         return 0;
165:     }
166:     if(i->i_uid == uid) {
167:         if((((i->i_mode >> 6) & 7) & mask) == mask) {
168:             return 0;
169:         }
170:     }
171:     if(!check_group(i)) {
172:         if((((i->i_mode >> 3) & 7) & mask) == mask) {
173:             return 0;
174:         }
175:     }
176:     if(((i->i_mode & 7) & mask) == mask) {
177:         return 0;
178:     }
179:
180:     return -EACCES;
181: }
182:
183: /* Linux 2.0.40 ABI system call (some from 2.2.26) */
184: void *syscall_table[] = {
185:     NULL, /* 0 */ /* sys_setup (-ENOSYS) */
186:     sys_exit,
187:     sys_fork,
188:     sys_read,
189:     sys_write,
190:     sys_open, /* 5 */
191:     sys_close,
192:     sys_waitpid,
193:     sys_creat,
194:     sys_link,
195:     sys_unlink, /* 10 */
196:     sys_execve,
197:     sys_chdir,
198:     sys_time,
199:     sys_mknod,
200:     sys_chmod, /* 15 */
201:     sys_chown,

```

## kernel/syscalls.c

Page 4/7

```

202:         NULL,                                /* sys_break (-ENOSYS) */
203:         sys_stat,
204:         sys_lseek,
205:         sys_getpid,                            /* 20 */
206:         sys_mount,
207:         sys_umount,
208:         sys_setuid,
209:         sys_getuid,
210:         sys_stime,                            /* 25 */
211:         NULL, // sys_ptrace
212:         sys_alarm,
213:         sys_fstat,
214:         sys_pause,
215:         sys_utime,                            /* 30 */
216:         NULL,                                /* sys_stty (-ENOSYS) */
217:         NULL,                                /* sys_gtty (-ENOSYS) */
218:         sys_access,
219:         NULL, // sys_nice
220:         sys_ftime,                            /* 35 */
221:         sys_sync,
222:         sys_kill,
223:         sys_rename,
224:         sys_mkdir,
225:         sys_rmdir,                            /* 40 */
226:         sys_dup,
227:         sys_pipe,
228:         sys_times,
229:         NULL, // sys_prof
230:         sys_brk,                              /* 45 */
231:         sys_setgid,
232:         sys_getgid,
233:         sys_signal,
234:         sys_geteuid,
235:         sys_getegid,                          /* 50 */
236:         NULL, // sys_acct
237:         sys_umount2,
238:         NULL,                                /* sys_lock (-ENOSYS) */
239:         sys_ioctl,
240:         sys_fcntl,                            /* 55 */
241:         NULL,                                /* sys_mpx (-ENOSYS) */
242:         sys_setpgid,
243:         NULL,                                /* sys_ulimit (-ENOSYS) */
244:         sys_olduname,
245:         sys_umask,                            /* 60 */
246:         sys_chroot,
247:         sys_ustat,
248:         sys_dup2,
249:         sys_getppid,
250:         sys_getpgrp,                          /* 65 */
251:         sys_setsid,
252:         sys_sigaction,
253:         sys_sgetmask,
254:         sys_ssetmask,
255:         sys_setreuid,                         /* 70 */
256:         sys_setregid,
257:         sys_sigsuspend,
258:         sys_sigpending,
259:         sys_sethostname,
260:         sys_setrlimit,                       /* 75 */
261:         sys_getrlimit,
262:         sys_getrusage,
263:         sys_gettimeofday,
264:         sys_settimeofday,
265:         sys_getgroups,                        /* 80 */
266:         sys_setgroups,
267:         old_select,
268:         sys_symlink,

```

## kernel/syscalls.c

Page 5/7

```

269:      sys_lstat,
270:      sys_readlink,                /* 85 */
271:      NULL, // sys_uselib
272:      NULL, // sys_swapon
273:      sys_reboot,
274:      NULL, // old_readdir
275:      old_mmap,                    /* 90 */
276:      sys_munmap,
277:      sys_truncate,
278:      sys_ftruncate,
279:      sys_fchmod,
280:      sys_fchown,                  /* 95 */
281:      NULL, // sys_getpriority
282:      NULL, // sys_setpriority
283:      NULL,                        /* sys_profil (-ENOSYS) */
284:      sys_statfs,
285:      sys_fstatfs,                 /* 100 */
286:      sys_ioperm,
287:      sys_socketcall, // sys_socketcall XXX
288:      NULL, // sys_syslog
289:      sys_setitimer,
290:      sys_getitimer,              /* 105 */
291:      sys_newstat,
292:      sys_newlstat,
293:      sys_newfstat,
294:      sys_uname,
295:      sys_iopl,                   /* 110 */
296:      NULL, // sys_vhangup
297:      NULL,                        /* sys_idle (-ENOSYS) */
298:      NULL, // sys_vm86old
299:      sys_wait4,
300:      NULL, // sys_swapoff        /* 115 */
301:      sys_sysinfo,
302:      NULL, // sys_ipc
303:      sys_fsync,
304:      sys_sigreturn,
305:      NULL, // sys_clone          /* 120 */
306:      sys_setdomainname,
307:      sys_newuname,
308:      NULL, // sys_modify_ldt
309:      NULL, // sys_adjtimex
310:      sys_mprotect,               /* 125 */
311:      sys_sigprocmask,
312:      NULL, // sys_create_module
313:      NULL, // sys_init_module
314:      NULL, // sys_delete_module
315:      NULL, // sys_get_kernel_syms /* 130 */
316:      NULL, // sys_quotactl
317:      sys_getpgid,
318:      sys_fchdir,
319:      NULL, // sys_bdflush
320:      NULL, // sys_sysfs          /* 135 */
321:      sys_personality,
322:      NULL,                        /* afs_syscall (-ENOSYS) */
323:      sys_setfsuid,
324:      sys_setfsgid,
325:      sys_llseek,                 /* 140 */
326:      sys_getdents,
327:      sys_select,
328:      sys_flock,
329:      NULL, // sys_msync
330:      NULL, // sys_readv          /* 145 */
331:      NULL, // sys_writev
332:      sys_getsid,
333:      sys_fdatasync,
334:      NULL, // sys_sysctl
335:      NULL, // sys_mlock          /* 150 */

```

## kernel/syscalls.c

Page 6/7

```

336:     NULL,    // sys_munlock
337:     NULL,    // sys_mlockall
338:     NULL,    // sys_munlockall
339:     NULL,    // sys_sched_setparam
340:     NULL,    // sys_sched_getparam /* 155 */
341:     NULL,    // sys_sched_setscheduler
342:     NULL,    // sys_sched_getscheduler
343:     NULL,    // sys_sched_yield
344:     NULL,    // sys_sched_get_priority_max
345:     NULL,    // sys_sched_get_priority_min /* 160 */
346:     NULL,    // sys_sched_rr_get_interval
347:     sys_nanosleep,
348:     NULL,    // sys_mremap
349:
350:     NULL,
351:     NULL,
352:     NULL,
353:     NULL,
354:     NULL,
355:     NULL,
356:     NULL,
357:     NULL,
358:     NULL,
359:     NULL,
360:     NULL,
361:     NULL,
362:     NULL,
363:     NULL,
364:     NULL,
365:     NULL,
366:     NULL,
367:     NULL,
368:     sys_chown,          /* 182 */
369:     sys_getcwd,
370:     NULL,
371:     NULL,
372:     NULL,
373:     NULL,
374:     NULL,
375:     NULL,
376:     sys_fork,          /* 190 (sys_vfork) */
377: };
378:
379: static void do_bad_syscall(unsigned int num)
380: {
381: #ifdef __DEBUG__
382:     printk("***** (pid %d) system call %d not supported yet *****\n", current->pid, num);
383: #endif /* __DEBUG__ */
384: }
385:
386: /*
387:  * The argument 'struct sigcontext' is needed because there are some system
388:  * calls (such as sys_iopl and sys_fork) that need to get information from
389:  * certain registers (EFLAGS and ESP). The rest of system calls will ignore
390:  * such extra argument.
391:  */
392: int do_syscall(unsigned int num, int arg1, int arg2, int arg3, int arg4, int arg
5, struct sigcontext sc)
393: {
394:     int (*sys_func)(int, ...);
395:
396:     if(num > NR_SYSCALLS) {
397:         do_bad_syscall(num);
398:         return -ENOSYS;
399:     }
400:     sys_func = syscall_table[num];

```

**kernel/syscalls.c**

Page 7/7

```
401:         if(!sys_func) {
402:             do_bad_syscall(num);
403:             return -ENOSYS;
404:         }
405:         current->sp = (unsigned int)&sc;
406:         return sys_func(arg1, arg2, arg3, arg4, arg5, &sc);
407: }
```



## kernel/timer.c

Page 1/8

```

1: /*
2:  * fiwix/kernel/timer.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/const.h>
11: #include <fiwix/cmos.h>
12: #include <fiwix/pit.h>
13: #include <fiwix/timer.h>
14: #include <fiwix/time.h>
15: #include <fiwix/pic.h>
16: #include <fiwix/sched.h>
17: #include <fiwix/pic.h>
18: #include <fiwix/cmos.h>
19: #include <fiwix/signal.h>
20: #include <fiwix/process.h>
21: #include <fiwix/sleep.h>
22: #include <fiwix/errno.h>
23: #include <fiwix/stdio.h>
24: #include <fiwix/string.h>
25:
26: /*
27:  * timer.c implements a callout table using a singly linked list.
28:  *
29:  * head
30:  * +-----+ +-----+ ... +-----+
31:  * |data|next| |data|next| ... |data|next|
32:  * |-----> |-----> ... |-----> /
33:  * +-----+ +-----+ ... +-----+
34:  * (callout) (callout) ... (callout)
35:  */
36:
37: struct callout callout_pool[NR_CALLOUTS];
38: struct callout *callout_pool_head;
39: struct callout *callout_head;
40:
41: static char month[12] = { 31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 };
42: unsigned int avenrun[3] = { 0, 0, 0 };
43:
44: static struct bh timer_bh = { 0, &irq_timer_bh, NULL };
45: static struct bh callouts_bh = { 0, &do_callouts_bh, NULL };
46: static struct interrupt irq_config_timer = { 0, "timer", &irq_timer, NULL };
47:
48: static unsigned int count_active_procs(void)
49: {
50:     int counter;
51:     struct proc *p;
52:
53:     counter = 0;
54:     FOR_EACH_PROCESS(p) {
55:         if(p->state == PROC_RUNNING) {
56:             counter += FIXED_1;
57:         }
58:     }
59:     return counter;
60: }
61:
62: static void calc_load(void)
63: {
64:     unsigned int active_procs;
65:     static int count = LOAD_FREQ;
66:
67:     if(count-- > 0) {

```

```

68:         return;
69:     }
70:
71:     count = LOAD_FREQ;
72:     active_procs = count_active_procs();
73:     CALC_LOAD(avenrun[0], EXP_1, active_procs);
74:     CALC_LOAD(avenrun[1], EXP_5, active_procs);
75:     CALC_LOAD(avenrun[2], EXP_15, active_procs);
76: }
77:
78: static struct callout *get_free_callout(void)
79: {
80:     struct callout *new;
81:
82:     new = NULL;
83:     if(callout_pool_head) {
84:         new = callout_pool_head;
85:         callout_pool_head = callout_pool_head->next;
86:         new->next = NULL;
87:     }
88:     return new;
89: }
90:
91: static void put_free_callout(struct callout *old)
92: {
93:     old->next = callout_pool_head;
94:     callout_pool_head = old;
95: }
96:
97: static void do_del_callout(struct callout *c)
98: {
99:     struct callout **tmp;
100:
101:     if(callout_head) {
102:         tmp = &callout_head;
103:         while(*tmp) {
104:             if(*tmp == c) {
105:                 if(*tmp->next != NULL) {
106:                     *tmp = (*tmp)->next;
107:                     (*tmp)->expires += c->expires;
108:                 } else {
109:                     *tmp = NULL;
110:                 }
111:                 put_free_callout(c);
112:                 break;
113:             }
114:             tmp = &(*tmp)->next;
115:         }
116:     }
117:     return;
118: }
119:
120: void add_callout(struct callout_req *creq, unsigned int ticks)
121: {
122:     unsigned long int flags;
123:     struct callout *c, **tmp;
124:
125:     del_callout(creq);
126:     SAVE_FLAGS(flags); CLI();
127:
128:     if(!(c = get_free_callout())) {
129:         printk("WARNING: %s(): no more callout slots!\n", __FUNCTION__);
130:         RESTORE_FLAGS(flags);
131:         return;
132:     }
133:
134:     /* setup the new callout */

```

```

135:     memset_b(c, NULL, sizeof(struct callout));
136:     c->expires = ticks;
137:     c->fn = creq->fn;
138:     c->arg = creq->arg;
139:
140:     if(!callout_head) {
141:         callout_head = c;
142:     } else {
143:         tmp = &callout_head;
144:         while(*tmp) {
145:             if((*tmp)->expires > c->expires) {
146:                 (*tmp)->expires -= c->expires;
147:                 c->next = *tmp;
148:                 break;
149:             }
150:             c->expires -= (*tmp)->expires;
151:             tmp = &(*tmp)->next;
152:         }
153:         *tmp = c;
154:     }
155:     RESTORE_FLAGS(flags);
156:     return;
157: }
158:
159: void del_callout(struct callout_req *creq)
160: {
161:     unsigned long int flags;
162:     struct callout *c;
163:
164:     SAVE_FLAGS(flags); CLI();
165:     c = callout_head;
166:     while(c) {
167:         if(c->fn == creq->fn && c->arg == creq->arg) {
168:             do_del_callout(c);
169:             break;
170:         }
171:         c = c->next;
172:     }
173:     RESTORE_FLAGS(flags);
174:     return;
175: }
176:
177: void irq_timer(int num, struct sigcontext *sc)
178: {
179:     if((++kstat.ticks % HZ) == 0) {
180:         CURRENT_TIME++;
181:         kstat.uptime++;
182:     }
183:
184:     timer_bh.flags |= BH_ACTIVE;
185:
186:     /* FIXME: put this in 'timer_bh' */
187:     if(sc->cs == KERNEL_CS) {
188:         current->usage.ru_stime.tv_usec += TICK;
189:         if(current->usage.ru_stime.tv_usec >= 1000000) {
190:             current->usage.ru_stime.tv_sec++;
191:             current->usage.ru_stime.tv_usec -= 1000000;
192:         }
193:         if(current->pid != IDLE) {
194:             kstat.cpu_system++;
195:         }
196:     } else {
197:         current->usage.ru_utime.tv_usec += TICK;
198:         if(current->usage.ru_utime.tv_usec >= 1000000) {
199:             current->usage.ru_utime.tv_sec++;
200:             current->usage.ru_utime.tv_usec -= 1000000;
201:         }

```

## kernel/timer.c

Page 4/8

```

202:         if(current->pid != IDLE) {
203:             kstat.cpu_user++;
204:         }
205:         if(current->it_virt_value > 0) {
206:             current->it_virt_value--;
207:             if(!current->it_virt_value) {
208:                 current->it_virt_value = current->it_virt_interv
al;
209:                 send_sig(current, SIGVTALRM);
210:             }
211:         }
212:     }
213: }
214:
215: unsigned long int tv2ticks(const struct timeval *tv)
216: {
217:     return((tv->tv_sec * HZ) + tv->tv_usec * HZ / 1000000);
218: }
219:
220: void ticks2tv(long int ticks, struct timeval *tv)
221: {
222:     tv->tv_sec = ticks / HZ;
223:     tv->tv_usec = (ticks % HZ) * 1000000 / HZ;
224:     return;
225: }
226:
227: int setitimer(int which, const struct itimerval *new_value, struct itimerval *ol
d_value)
228: {
229:     switch(which) {
230:         case ITIMER_REAL:
231:             if((unsigned int)old_value) {
232:                 ticks2tv(current->it_real_interval, &old_value->
it_interval);
233:                 ticks2tv(current->it_real_value, &old_value->it_
value);
234:             }
235:             current->it_real_interval = tv2ticks(&new_value->it_inte
rval);
236:             current->it_real_value = tv2ticks(&new_value->it_value);
237:             break;
238:         case ITIMER_VIRTUAL:
239:             if((unsigned int)old_value) {
240:                 ticks2tv(current->it_virt_interval, &old_value->
it_interval);
241:                 ticks2tv(current->it_virt_value, &old_value->it_
value);
242:             }
243:             current->it_virt_interval = tv2ticks(&new_value->it_inte
rval);
244:             current->it_virt_value = tv2ticks(&new_value->it_value);
245:             break;
246:         case ITIMER_PROF:
247:             if((unsigned int)old_value) {
248:                 ticks2tv(current->it_prof_interval, &old_value->
it_interval);
249:                 ticks2tv(current->it_prof_value, &old_value->it_
value);
250:             }
251:             current->it_prof_interval = tv2ticks(&new_value->it_inte
rval);
252:             current->it_prof_value = tv2ticks(&new_value->it_value);
253:             break;
254:         default:
255:             return -EINVAL;
256:     }
257:

```

## kernel/timer.c

Page 5/8

```

258:         return 0;
259:     }
260:
261:     unsigned long int mktime(struct mt *mt)
262:     {
263:         int n, total_days;
264:         unsigned long int seconds;
265:
266:         total_days = 0;
267:
268:         for(n = UNIX_EPOCH; n < mt->mt_year; n++) {
269:             total_days += DAYS_PER_YEAR(n);
270:         }
271:         for(n = 0; n < (mt->mt_month - 1); n++) {
272:             total_days += month[n];
273:             if(n == 1) {
274:                 total_days += LEAP_YEAR(mt->mt_year) ? 1 : 0;
275:             }
276:         }
277:
278:         total_days += (mt->mt_day - 1);
279:         seconds = total_days * SECS_PER_DAY;
280:         seconds += mt->mt_hour * SECS_PER_HOUR;
281:         seconds += mt->mt_min * SECS_PER_MIN;
282:         seconds += mt->mt_sec;
283:         return seconds;
284:     }
285:
286:     void irq_timer_bh(void)
287:     {
288:         struct proc *p;
289:
290:         if(current->usage.ru_utime.tv_sec + current->usage.ru_stime.tv_sec > current->rlim[RLIMIT_CPU].rlim_cur) {
291:             send_sig(current, SIGXCPU);
292:         }
293:
294:         if(current->it_prof_value > 0) {
295:             current->it_prof_value--;
296:             if(!current->it_prof_value) {
297:                 current->it_prof_value = current->it_prof_interval;
298:                 send_sig(current, SIGPROF);
299:             }
300:         }
301:
302:         calc_load();
303:         FOR_EACH_PROCESS(p) {
304:             if(!p->state) {
305:                 continue;
306:             }
307:             if(p->timeout > 0 && p->timeout < INFINITE_WAIT) {
308:                 p->timeout--;
309:                 if(!p->timeout) {
310:                     wakeup_proc(p);
311:                 }
312:             }
313:             if(p->it_real_value > 0) {
314:                 p->it_real_value--;
315:                 if(!p->it_real_value) {
316:                     p->it_real_value = p->it_real_interval;
317:                     send_sig(p, SIGALRM);
318:                 }
319:             }
320:         }
321:
322:         /* callouts */
323:         if(callout_head) {

```

## kernel/timer.c

Page 6/8

```

324:         if(callout_head->expires > 0) {
325:             callout_head->expires--;
326:             if(!callout_head->expires) {
327:                 callouts_bh.flags |= BH_ACTIVE;
328:             }
329:         } else {
330:             printk("%s(): callout losing ticks.\n", __FUNCTION__);
331:             callouts_bh.flags |= BH_ACTIVE;
332:         }
333:     }
334:
335:     if(current->pid > IDLE && --current->cpu_count <= 0) {
336:         current->cpu_count = 0;
337:         need_resched = 1;
338:     }
339: }
340:
341: void do_callouts_bh(void)
342: {
343:     struct callout *c;
344:     void (*fn)(unsigned int);
345:     unsigned int arg;
346:
347:     while(callout_head) {
348:         if(callout_head->expires) {
349:             break;
350:         }
351:         if(lock_area(AREA_CALLOUT)) {
352:             continue;
353:         }
354:         fn = callout_head->fn;
355:         arg = callout_head->arg;
356:         c = callout_head;
357:         callout_head = callout_head->next;
358:         put_free_callout(c);
359:         unlock_area(AREA_CALLOUT);
360:         fn(arg);
361:     }
362: }
363:
364: void get_system_time(void)
365: {
366:     short int cmos_century;
367:     struct mt mt;
368:
369:     /* read date and time from CMOS */
370:     mt.mt_sec = cmos_read_date(CMOS_SEC);
371:     mt.mt_min = cmos_read_date(CMOS_MIN);
372:     mt.mt_hour = cmos_read_date(CMOS_HOUR);
373:     mt.mt_day = cmos_read_date(CMOS_DAY);
374:     mt.mt_month = cmos_read_date(CMOS_MONTH);
375:     mt.mt_year = cmos_read_date(CMOS_YEAR);
376:     cmos_century = cmos_read_date(CMOS_CENTURY);
377:     mt.mt_year += cmos_century * 100;
378:
379:     kstat.boot_time = CURRENT_TIME = mktime(&mt);
380: }
381:
382: void set_system_time(__time_t t)
383: {
384:     int sec, spm, min, hour, d, m, y;
385:
386:     sec = t;
387:     y = 1970;
388:     while(sec >= (DAYS_PER_YEAR(y) * SECS_PER_DAY)) {
389:         sec -= (DAYS_PER_YEAR(y) * SECS_PER_DAY);
390:         y++;

```

## kernel/timer.c

Page 7/8

```

391:         }
392:
393:         m = 0;
394:         while(sec > month[m] * SECS_PER_DAY) {
395:             spm = month[m] * SECS_PER_DAY;
396:             if(m == 1) {
397:                 spm = LEAP_YEAR(y) ? spm + SECS_PER_DAY : spm;
398:             }
399:             sec -= spm;
400:             m++;
401:         }
402:         m++;
403:
404:         d = 1;
405:         while(sec >= SECS_PER_DAY) {
406:             sec -= SECS_PER_DAY;
407:             d++;
408:         }
409:
410:         hour = 0;
411:         while(sec >= SECS_PER_HOUR) {
412:             sec -= SECS_PER_HOUR;
413:             hour++;
414:         }
415:
416:         min = 0;
417:         while(sec >= SECS_PER_MIN) {
418:             sec -= SECS_PER_MIN;
419:             min++;
420:         }
421:
422:         /* write date and time to CMOS */
423:         cmos_write_date(CMOS_SEC, sec);
424:         cmos_write_date(CMOS_MIN, min);
425:         cmos_write_date(CMOS_HOUR, hour);
426:         cmos_write_date(CMOS_DAY, d);
427:         cmos_write_date(CMOS_MONTH, m);
428:         cmos_write_date(CMOS_YEAR, y % 100);
429:         cmos_write_date(CMOS_CENTURY, (y - (y % 100)) / 100);
430:
431:         CURRENT_TIME = t;
432:     }
433:
434: void timer_init(void)
435: {
436:     int n;
437:     struct callout *c;
438:
439:     add_bh(&timer_bh);
440:     add_bh(&callouts_bh);
441:
442:     pit_init(HZ);
443:
444:     memset_b(callout_pool, NULL, sizeof(callout_pool));
445:
446:     /* callout free list initialization */
447:     callout_pool_head = NULL;
448:     n = NR_CALLOUTS;
449:     while(n--) {
450:         c = &callout_pool[n];
451:         put_free_callout(c);
452:     }
453:     callout_head = NULL;
454:
455:     printk("clock      -                %d      type=PIT Hz=%d\n", TIMER_IRQ, H
Z);
456:     if(!register_irq(TIMER_IRQ, &irq_config_timer)) {

```

**kernel/timer.c**

Page 8/8

```
457:                enable_irq(TIMER_IRQ);
458:            }
459: }
```



## kernel/traps.c

```

1: /*
2:  * fiwix/kernel/traps.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/traps.h>
11: #include <fiwix/cpu.h>
12: #include <fiwix/mm.h>
13: #include <fiwix/process.h>
14: #include <fiwix/signal.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17: #include <fiwix/sched.h>
18:
19: /*
20:  * PS/2 System Control Port B
21:  * -----
22:  * bit 7 -> system board RAM parity check
23:  * bit 6 -> channel check
24:  * bit 5 -> timer 2 (speaker time) output
25:  * bit 4 -> refresh request (toggle)
26:  * bit 3 -> channel check status
27:  * bit 2 -> parity check status
28:  * bit 1 -> speaker data status
29:  * bit 0 -> timer 2 gate to speaker status
30:  */
31: #define PS2_SYSCCTRL_B 0x61 /* PS/2 system control port B (read) */
32:
33: struct traps traps_table[NR_EXCEPTIONS] = {
34:     { "Divide Error", do_divide_error, 0 },
35:     { "Debug", do_debug, 0 },
36:     { "NMI Interrupt", do_nmi_interrupt, 0 },
37:     { "Breakpoint", do_breakpoint, 0 },
38:     { "Overflow", do_overflow, 0 },
39:     { "BOUND Range Exceeded", do_bound, 0 },
40:     { "Invalid Opcode", do_invalid_opcode, 0 },
41:     { "Device Not Available (No Math Coprocessor)", do_no_math_coprocessor,
0 },
42:     { "Double Fault", do_double_fault, 1 },
43:     { "Coprocessor Segment Overrun", do_coprocessor_segment_overrun, 0 },
44:     { "Invalid TSS", do_invalid_tss, 1 },
45:     { "Segment Not Present", do_segment_not_present, 1 },
46:     { "Stack-Segment Fault", do_stack_segment_fault, 1 },
47:     { "General Protection", do_general_protection, 1 },
48:     { "Page Fault", do_page_fault, 1 },
49:     { "Intel reserved", do_reserved, 0 },
50:     { "x87 FPU Floating-Point Error", do_floating_point_error, 0 },
51:     { "Alignment Check", do_alignment_check, 1 },
52:     { "Machine Check", do_machine_check, 0 },
53:     { "SIMD Floating-Point Exception", do_simd_fault, 0 },
54:     { "Intel reserved", do_reserved, 0 },
55:     { "Intel reserved", do_reserved, 0 },
56:     { "Intel reserved", do_reserved, 0 },
57:     { "Intel reserved", do_reserved, 0 },
58:     { "Intel reserved", do_reserved, 0 },
59:     { "Intel reserved", do_reserved, 0 },
60:     { "Intel reserved", do_reserved, 0 },
61:     { "Intel reserved", do_reserved, 0 },
62:     { "Intel reserved", do_reserved, 0 },
63:     { "Intel reserved", do_reserved, 0 },
64:     { "Intel reserved", do_reserved, 0 },
65:     { "Intel reserved", do_reserved, 0 }
66: };

```

## kernel/traps.c

Page 2/6

```
67:
68: void do_divide_error(unsigned int trap, struct sigcontext *sc)
69: {
70:     if(dump_registers(trap, sc)) {
71:         PANIC("");
72:     }
73:     send_sig(current, SIGFPE);
74:     return;
75: }
76:
77: void do_debug(unsigned int trap, struct sigcontext *sc)
78: {
79:     if(dump_registers(trap, sc)) {
80:         PANIC("");
81:     }
82:     send_sig(current, SIGTRAP);
83:     return;
84: }
85:
86: void do_nmi_interrupt(unsigned int trap, struct sigcontext *sc)
87: {
88:     unsigned char error;
89:
90:     error = inport_b(PS2_SYSCTRL_B);
91:
92:     printk("NMI received: ", error);
93:     switch(error) {
94:         case 0x80:
95:             printk("parity check occurred. Defective RAM chips?\n");
96:             break;
97:         default:
98:             printk("unknown error 0x%x\n", error);
99:             break;
100:    }
101:
102:    if(dump_registers(trap, sc)) {
103:        PANIC("");
104:    }
105:    send_sig(current, SIGSEGV);
106:    return;
107: }
108:
109: void do_breakpoint(unsigned int trap, struct sigcontext *sc)
110: {
111:     if(dump_registers(trap, sc)) {
112:         PANIC("");
113:     }
114:     send_sig(current, SIGTRAP);
115:     return;
116: }
117:
118: void do_overflow(unsigned int trap, struct sigcontext *sc)
119: {
120:     if(dump_registers(trap, sc)) {
121:         PANIC("");
122:     }
123:     send_sig(current, SIGSEGV);
124:     return;
125: }
126:
127: void do_bound(unsigned int trap, struct sigcontext *sc)
128: {
129:     if(dump_registers(trap, sc)) {
130:         PANIC("");
131:     }
132:     send_sig(current, SIGSEGV);
133:     return;
```

```
134: }
135:
136: void do_invalid_opcode(unsigned int trap, struct sigcontext *sc)
137: {
138:     if(dump_registers(trap, sc)) {
139:         PANIC("");
140:     }
141:     send_sig(current, SIGILL);
142:     return;
143: }
144:
145: void do_no_math_coprocessor(unsigned int trap, struct sigcontext *sc)
146: {
147:     /* floating-point emulation would go here */
148:
149:     if(dump_registers(trap, sc)) {
150:         PANIC("No coprocessor/emulation found.\n");
151:     }
152:     send_sig(current, SIGILL);
153:     return;
154: }
155:
156: void do_double_fault(unsigned int trap, struct sigcontext *sc)
157: {
158:     if(dump_registers(trap, sc)) {
159:         PANIC("");
160:     }
161:     send_sig(current, SIGSEGV);
162:     return;
163: }
164:
165: void do_coprocessor_segment_overrun(unsigned int trap, struct sigcontext *sc)
166: {
167:     if(dump_registers(trap, sc)) {
168:         PANIC("");
169:     }
170:     send_sig(current, SIGFPE);
171:     return;
172: }
173:
174: void do_invalid_tss(unsigned int trap, struct sigcontext *sc)
175: {
176:     if(dump_registers(trap, sc)) {
177:         PANIC("");
178:     }
179:     send_sig(current, SIGSEGV);
180:     return;
181: }
182:
183: void do_segment_not_present(unsigned int trap, struct sigcontext *sc)
184: {
185:     if(dump_registers(trap, sc)) {
186:         PANIC("");
187:     }
188:     send_sig(current, SIGBUS);
189:     return;
190: }
191:
192: void do_stack_segment_fault(unsigned int trap, struct sigcontext *sc)
193: {
194:     if(dump_registers(trap, sc)) {
195:         PANIC("");
196:     }
197:     send_sig(current, SIGBUS);
198:     return;
199: }
200:
```

## kernel/traps.c

Page 4/6

```
201: void do_general_protection(unsigned int trap, struct sigcontext *sc)
202: {
203:     if(dump_registers(trap, sc)) {
204:         PANIC("");
205:     }
206:     send_sig(current, SIGSEGV);
207:     return;
208: }
209:
210: /* do_page_fault() resides in mm/fault.c */
211:
212: void do_reserved(unsigned int trap, struct sigcontext *sc)
213: {
214:     if(dump_registers(trap, sc)) {
215:         PANIC("");
216:     }
217:     send_sig(current, SIGSEGV);
218:     return;
219: }
220:
221: void do_floating_point_error(unsigned int trap, struct sigcontext *sc)
222: {
223:     if(dump_registers(trap, sc)) {
224:         PANIC("");
225:     }
226:     send_sig(current, SIGFPE);
227:     return;
228: }
229:
230: void do_alignment_check(unsigned int trap, struct sigcontext *sc)
231: {
232:     if(dump_registers(trap, sc)) {
233:         PANIC("");
234:     }
235:     send_sig(current, SIGSEGV);
236:     return;
237: }
238:
239: void do_machine_check(unsigned int trap, struct sigcontext *sc)
240: {
241:     if(dump_registers(trap, sc)) {
242:         PANIC("");
243:     }
244:     send_sig(current, SIGSEGV);
245:     return;
246: }
247:
248: void do_simd_fault(unsigned int trap, struct sigcontext *sc)
249: {
250:     if(dump_registers(trap, sc)) {
251:         PANIC("");
252:     }
253:     send_sig(current, SIGSEGV);
254:     return;
255: }
256:
257: void trap_handler(unsigned int trap, struct sigcontext sc)
258: {
259:     traps_table[trap].handler(trap, &sc);
260:
261:     /* avoids confusion with -RESTART return value */
262:     sc.err = -sc.err;
263: }
264:
265: static const char * elf_lookup_symbol(unsigned int addr)
266: {
267:     Elf32_Sym *sym;
```

## kernel/traps.c

Page 5/6

```

268:         unsigned int n;
269:
270:         sym = (Elf32_Sym *)symtab->sh_addr;
271:         for(n = 0; n < symtab->sh_size / sizeof(Elf32_Sym); n++, sym++) {
272:             if(ELF32_ST_TYPE(sym->st_info) != STT_FUNC) {
273:                 continue;
274:             }
275:             if(addr >= sym->st_value && addr < (sym->st_value + sym->st_size
)) {
276:                 return (const char *)strtab->sh_addr + sym->st_name;
277:             }
278:         }
279:         return NULL;
280:     }
281:
282: int dump_registers(unsigned int trap, struct sigcontext *sc)
283: {
284:     unsigned int cr2, addr, n;
285:     unsigned int *sp;
286:     const char *str;
287:
288:     printk("\n");
289:     if(trap == 14) { /* Page Fault */
290:         GET_CR2(cr2);
291:         printk("%s at 0x%08x (%s) with error code 0x%08x (0b%b)\n", trap
s_table[trap].name, cr2, sc->err & PFAULT_W ? "writing" : "reading", sc->err, sc->err);
292:     } else {
293:         printk("EXCEPTION: %s", traps_table[trap].name);
294:         if(traps_table[trap].errcode) {
295:             printk(": error code 0x%08x (0b%b)", sc->err, sc->err);
296:         }
297:         printk("\n");
298:     }
299:
300:     printk("Process '%s' with pid %d", current->argv0, current->pid);
301:     if(sc->cs == KERNEL_CS) {
302:         printk(" in '%s()' .", elf_lookup_symbol(sc->eip));
303:     }
304:     printk("\n");
305:
306:     printk(" cs: 0x%08x\teip: 0x%08x\tefl: 0x%08x\t ss: 0x%08x\tesp: 0x%08x\
n", sc->cs, sc->eip, sc->eflags, sc->oldss, sc->oldesp);
307:     printk("eax: 0x%08x\tebx: 0x%08x\tecx: 0x%08x\tedx: 0x%08x\n", sc->eax,
sc->ebx, sc->ecx, sc->edx);
308:     printk("esi: 0x%08x\tedi: 0x%08x\tesp: 0x%08x\tebp: 0x%08x\n", sc->esi,
sc->edi, sc->esp, sc->ebp);
309:     printk(" ds: 0x%08x\t es: 0x%08x\t fs: 0x%08x\t gs: 0x%08x\n", sc->ds, s
c->es, sc->fs, sc->gs);
310:
311:     if(sc->cs == KERNEL_CS) {
312:         printk("Stack:\n");
313:         GET_ESP(sp);
314:         sp = (unsigned int *)sp;
315:         for(n = 1; n <= 32; n++) {
316:             printk(" %08x", *sp);
317:             sp++;
318:             if(!(n % 8)) {
319:                 printk("\n");
320:             }
321:         }
322:         printk("Backtrace:\n");
323:         GET_ESP(sp);
324:         sp = (unsigned int *)sp;
325:         for(n = 0; n < 128; n++) {
326:             addr = *sp;
327:             str = elf_lookup_symbol(addr);
328:             if(str) {

```

## kernel/traps.c

Page 6/6

```
329:                                     printk("<0x%08x> %s()\n", addr, str);
330:                                     }
331:                                     sp++;
332:                                     }
333:                                     }
334:
335:                                     /* panics if the exception has been in kernel mode */
336:                                     if(current->flags & PF_KPROC || sc->cs == KERNEL_CS) {
337:                                     return 1;
338:                                     }
339:
340:                                     return 0;
341: }
```

## kernel/syscalls/access.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/access.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/stat.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /*__DEBUG__*/
18:
19: int sys_access(const char *filename, __mode_t mode)
20: {
21:     struct inode *i;
22:     char *tmp_name;
23:     int errno;
24:
25: #ifdef __DEBUG__
26:     printk("(pid %d) sys_access('%s', %d)", current->pid, filename, mode);
27: #endif /*__DEBUG__*/
28:
29:     if((mode & S_IRWXO) != mode) {
30:         return -EINVAL;
31:     }
32:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
33:         return errno;
34:     }
35:     current->flags |= PF_USERREAL;
36:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS))) {
37:         current->flags &= ~PF_USERREAL;
38:         free_name(tmp_name);
39:         return errno;
40:     }
41:     if(mode & TO_WRITE) {
42:         if(S_ISREG(i->i_mode) || S_ISDIR(i->i_mode) || S_ISLNK(i->i_mode
)) {
43:             if(IS_RDONLY_FS(i)) {
44:                 current->flags &= ~PF_USERREAL;
45:                 iput(i);
46:                 free_name(tmp_name);
47:                 return -EROFS;
48:             }
49:         }
50:     }
51:     errno = check_permission(mode, i);
52:
53: #ifdef __DEBUG__
54:     printk("-> returning %d\n", errno);
55: #endif /*__DEBUG__*/
56:
57:     current->flags &= ~PF_USERREAL;
58:     iput(i);
59:     free_name(tmp_name);
60:     return errno;
61: }

```

## kernel/syscalls/alarm.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/alarm.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/time.h>
10:
11: #ifdef __DEBUG__
12: #include <fiwix/stdio.h>
13: #include <fiwix/process.h>
14: #endif /*__DEBUG__*/
15:
16: int sys_alarm(unsigned int secs)
17: {
18:     struct itimerval value, ovalue;
19:
20: #ifdef __DEBUG__
21:     printk("(pid %d) sys_alarm(%d)\n", current->pid, secs);
22: #endif /*__DEBUG__*/
23:
24:     value.it_interval.tv_sec = 0;
25:     value.it_interval.tv_usec = 0;
26:     value.it_value.tv_sec = secs;
27:     value.it_value.tv_usec = 0;
28:     setitimer(ITIMER_REAL, &value, &ovalue);
29:
30:     /*
31:      * If there are still some usecs left and since the return value has
32:      * not enough granularity for them, then just add 1 second to it.
33:      */
34:     if(ovalue.it_value.tv_usec) {
35:         ovalue.it_value.tv_sec++;
36:     }
37:
38:     return ovalue.it_value.tv_sec;
39: }
```



## kernel/syscalls/brk.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/brk.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9: #include <fiwix/mm.h>
10: #include <fiwix/mman.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_brk(unsigned int brk)
18: {
19:     unsigned int newbrk;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_brk(0x%08x) -> ", current->pid, brk);
23: #endif /*__DEBUG__*/
24:
25:     if(!brk || brk < current->brk_lower) {
26: #ifdef __DEBUG__
27:         printk("0x%08x\n", current->brk);
28: #endif /*__DEBUG__*/
29:         return current->brk;
30:     }
31:
32:     newbrk = PAGE_ALIGN(brk);
33:     if(newbrk == current->brk || newbrk < current->brk_lower) {
34: #ifdef __DEBUG__
35:         printk("0x%08x\n", current->brk);
36: #endif /*__DEBUG__*/
37:         return brk;
38:     }
39:
40:     if(brk < current->brk) {
41:         do_munmap(newbrk, current->brk - newbrk);
42:         current->brk = brk;
43:         return brk;
44:     }
45:     if(!expand_heap(newbrk)) {
46:         current->brk = brk;
47:     } else {
48:         return -ENOMEM;
49:     }
50: #ifdef __DEBUG__
51:     printk("0x%08x\n", current->brk);
52: #endif /*__DEBUG__*/
53:     return current->brk;
54: }
```

## kernel/syscalls/chdir.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/chdir.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/stat.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_chdir(const char *dirname)
19: {
20:     struct inode *i;
21:     char *tmp_name;
22:     int errno;
23:
24: #ifdef __DEBUG__
25:     printk("(pid %d) sys_chdir('%s')\n", current->pid, dirname);
26: #endif /* __DEBUG__ */
27:
28:     if((errno = malloc_name(dirname, &tmp_name)) < 0) {
29:         return errno;
30:     }
31:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS)) {
32:         free_name(tmp_name);
33:         return errno;
34:     }
35:     if(!S_ISDIR(i->i_mode)) {
36:         iput(i);
37:         free_name(tmp_name);
38:         return -ENOTDIR;
39:     }
40:     if((errno = check_permission(TO_EXEC, i)) {
41:         iput(i);
42:         free_name(tmp_name);
43:         return errno;
44:     }
45:     iput(current->pwd);
46:     current->pwd = i;
47:     free_name(tmp_name);
48:     return 0;
49: }
```

## kernel/syscalls/chmod.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/chmod.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/errno.h>
13: #include <fiwix/string.h>
14:
15: #ifdef __DEBUG__
16: #include <fiwix/stdio.h>
17: #include <fiwix/process.h>
18: #endif /*__DEBUG__ */
19:
20: int sys_chmod(const char *filename, __mode_t mode)
21: {
22:     struct inode *i;
23:     char *tmp_name;
24:     int errno;
25:
26: #ifdef __DEBUG__
27:     printk("(pid %d) sys_chmod('%s', %d)\n", current->pid, filename, mode);
28: #endif /*__DEBUG__ */
29:
30:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
31:         return errno;
32:     }
33:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS))) {
34:         free_name(tmp_name);
35:         return errno;
36:     }
37:
38:     if(IS_RDONLY_FS(i)) {
39:         iput(i);
40:         free_name(tmp_name);
41:         return -EROFS;
42:     }
43:     if(check_user_permission(i)) {
44:         iput(i);
45:         free_name(tmp_name);
46:         return -EPERM;
47:     }
48:
49:     i->i_mode &= S_IFMT;
50:     i->i_mode |= mode & ~S_IFMT;
51:     i->i_ctime = CURRENT_TIME;
52:     i->dirty = 1;
53:     iput(i);
54:     free_name(tmp_name);
55:     return 0;
56: }

```

## kernel/syscalls/chown.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/chown.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/errno.h>
13: #include <fiwix/string.h>
14:
15: #ifdef __DEBUG__
16: #include <fiwix/stdio.h>
17: #include <fiwix/process.h>
18: #endif /*__DEBUG__ */
19:
20: int sys_chown(const char *filename, __uid_t owner, __gid_t group)
21: {
22:     struct inode *i;
23:     char *tmp_name;
24:     int errno;
25:
26: #ifdef __DEBUG__
27:     printk("(pid %d) sys_chown('%s', %d, %d)\n", current->pid, filename, own
er, group);
28: #endif /*__DEBUG__ */
29:
30:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
31:         return errno;
32:     }
33:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS))) {
34:         free_name(tmp_name);
35:         return errno;
36:     }
37:
38:     if(IS_RDONLY_FS(i)) {
39:         iput(i);
40:         free_name(tmp_name);
41:         return -EROFS;
42:     }
43:     if(check_user_permission(i)) {
44:         iput(i);
45:         free_name(tmp_name);
46:         return -EPERM;
47:     }
48:
49:     if(owner == (__uid_t)-1) {
50:         owner = i->i_uid;
51:     } else {
52:         i->i_mode &= ~(S_ISUID);
53:         i->i_ctime = CURRENT_TIME;
54:     }
55:     if(group == (__gid_t)-1) {
56:         group = i->i_gid;
57:     } else {
58:         i->i_mode &= ~(S_ISGID);
59:         i->i_ctime = CURRENT_TIME;
60:     }
61:
62:     i->i_uid = owner;
63:     i->i_gid = group;
64:     i->dirty = 1;
65:     iput(i);
66:     free_name(tmp_name);

```

**kernel/syscalls/chown.c**

Page 2/2

```
67:         return 0;  
68:     }
```

## kernel/syscalls/chroot.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/chroot.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/stat.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_chroot(const char *dirname)
19: {
20:     struct inode *i;
21:     char *tmp_name;
22:     int errno;
23:
24: #ifdef __DEBUG__
25:     printk("(pid %d) sys_chroot('%s')\n", current->pid, dirname);
26: #endif /* __DEBUG__ */
27:
28:     if((errno = malloc_name(dirname, &tmp_name)) < 0) {
29:         return errno;
30:     }
31:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS)) {
32:         free_name(tmp_name);
33:         return errno;
34:     }
35:     if(!S_ISDIR(i->i_mode)) {
36:         iput(i);
37:         free_name(tmp_name);
38:         return -ENOTDIR;
39:     }
40:     iput(current->root);
41:     current->root = i;
42:     free_name(tmp_name);
43:     return 0;
44: }
```

## kernel/syscalls/close.c

Page 1/1

```

1:  /*
2:  *  fiwix/kernel/syscalls/close.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/syscalls.h>
9:  #include <fiwix/locks.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/stdio.h>
12:
13: int sys_close(unsigned int ufd)
14: {
15:     unsigned int fd;
16:     struct inode *i;
17:
18: #ifdef __DEBUG__
19:     printk("(pid %d) sys_close(%d)\n", current->pid, ufd);
20: #endif /*__DEBUG__ */
21:
22:     CHECK_UFD(ufd);
23:     fd = current->fd[ufd];
24:     release_user_fd(ufd);
25:
26:     if(--fd_table[fd].count) {
27:         return 0;
28:     }
29:     i = fd_table[fd].inode;
30:     flock_release_inode(i);
31:     if(i->fsop && i->fsop->close) {
32:         i->fsop->close(i, &fd_table[fd]);
33:         release_fd(fd);
34:         iput(i);
35:         return 0;
36:     }
37:     printk("WARNING: %s(): ufd %d without the close() method!\n", __FUNCTION
__, ufd);
38:     return -EINVAL;
39: }

```

**kernel/syscalls/creat.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/creat.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/syscalls.h>
10: #include <fiwix/fcntl.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_creat(const char *filename, __mode_t mode)
18: {
19: #ifdef __DEBUG__
20:     printk("(pid %d) sys_creat('%s', %d)\n", current->pid, filename, mode);
21: #endif /*__DEBUG__*/
22:     return sys_open(filename, O_CREAT | O_WRONLY | O_TRUNC, mode);
23: }
```



## kernel/syscalls/dup2.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/dup2.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/syscalls.h>
9: #include <fiwix/process.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #endif /*__DEBUG__*/
15:
16: int sys_dup2(int old_ufd, int new_ufd)
17: {
18: #ifdef __DEBUG__
19:     printk("(pid %d) sys_dup2(%d, %d)", current->pid, old_ufd, new_ufd);
20: #endif /*__DEBUG__*/
21:
22:     CHECK_UFD(old_ufd);
23:     if(new_ufd < 0 || new_ufd > OPEN_MAX) {
24:         return -EINVAL;
25:     }
26:     if(old_ufd == new_ufd) {
27:         return new_ufd;
28:     }
29:     if(current->fd[new_ufd]) {
30:         sys_close(new_ufd);
31:     }
32:     if((new_ufd = get_new_user_fd(new_ufd)) < 0) {
33:         return new_ufd;
34:     }
35:     current->fd[new_ufd] = current->fd[old_ufd];
36:     fd_table[current->fd[new_ufd]].count++;
37: #ifdef __DEBUG__
38:     printk(" --> returning %d\n", new_ufd);
39: #endif /*__DEBUG__*/
40:     return new_ufd;
41: }
```

## kernel/syscalls/dup.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/dup.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/syscalls.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_dup(unsigned int ufd)
18: {
19:     int new_ufd;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_dup(%d)", current->pid, ufd);
23: #endif /*__DEBUG__*/
24:
25:     CHECK_UFD(ufd);
26:     if((new_ufd = get_new_user_fd(0)) < 0) {
27:         return new_ufd;
28:     }
29:
30: #ifdef __DEBUG__
31:     printk(" -> %d\n", new_ufd);
32: #endif /*__DEBUG__*/
33:
34:     current->fd[new_ufd] = current->fd[ufd];
35:     fd_table[current->fd[new_ufd]].count++;
36:     return new_ufd;
37: }
```

## kernel/syscalls/execve.c

Page 1/6

```

1:  /*
2:  * fiwix/kernel/syscalls/execve.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/syscalls.h>
9:  #include <fiwix/stat.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/mm.h>
12: #include <fiwix/process.h>
13: #include <fiwix/fcntl.h>
14: #include <fiwix/errno.h>
15: #include <fiwix/string.h>
16:
17: #ifdef __DEBUG__
18: #include <fiwix/stdio.h>
19: #endif /*__DEBUG__ */
20:
21: static int initialize_barg(struct binargs *barg, char *argv[], char *envp[])
22: {
23:     int n, errno;
24:
25:     for(n = 0; n < ARG_MAX; n++) {
26:         barg->page[n] = 0;
27:     }
28:     barg->argv_len = barg->envp_len = 0;
29:
30:     for(n = 0; argv[n]; n++) {
31:         if((errno = check_user_area(VERIFY_READ, argv[n], sizeof(char *)))
32: )) {
33:             return errno;
34:         }
35:         barg->argv_len += strlen(argv[n]) + 1;
36:     }
37:     barg->argc = n;
38:
39:     for(n = 0; envp[n]; n++) {
40:         if((errno = check_user_area(VERIFY_READ, envp[n], sizeof(char *)))
41: )) {
42:             return errno;
43:         }
44:         barg->envp_len += strlen(envp[n]) + 1;
45:     }
46:     barg->envc = n;
47:     return 0;
48: }
49:
50: static void free_barg_pages(struct binargs *barg)
51: {
52:     int n;
53:
54:     for(n = 0; n < ARG_MAX; n++) {
55:         if(barg->page[n]) {
56:             kfree(barg->page[n]);
57:         }
58:     }
59: }
60:
61: static int add_strings(struct binargs *barg, char *filename, char *interpreter,
62: char *args)
63: {
64:     int n, p, offset;
65:     unsigned int ae_str_len;
66:     char *page;

```

## kernel/syscalls/execve.c

Page 2/6

```

65:
66:     /*
67:      * For a script we need to substitute the saved argv[0] by the original
68:      * 'filename' supplied in execve(), otherwise the interpreter won't be
69:      * able to find the script file.
70:      */
71:     offset = barg->offset;
72:     p = ARG_MAX - 1;
73:     ae_str_len = barg->argv_len + barg->envp_len + 4;
74:     p -= ae_str_len / PAGE_SIZE;
75:     offset = PAGE_SIZE - (ae_str_len % PAGE_SIZE);
76:     if(offset == PAGE_SIZE) {
77:         offset = 0;
78:         p++;
79:     }
80:     page = (char *)barg->page[p];
81:     while(*(page + offset)) {
82:         offset++;
83:         barg->argv_len--;
84:         if(offset == PAGE_SIZE) {
85:             p++;
86:             offset = 0;
87:             page = (char *)barg->page[p];
88:         }
89:     }
90:     barg->argv_len--;
91:
92:
93:     p = ARG_MAX - 1;
94:     barg->argv_len += strlen(interpreter) + 1;
95:     barg->argv_len += strlen(args) ? strlen(args) + 1 : 0;
96:     barg->argv_len += strlen(filename) + 1;
97:     barg->argc++;
98:     if(*args) {
99:         barg->argc++;
100:    }
101:    ae_str_len = barg->argv_len + barg->envp_len + 4;
102:    p -= ae_str_len / PAGE_SIZE;
103:    offset = PAGE_SIZE - (ae_str_len % PAGE_SIZE);
104:    if(offset == PAGE_SIZE) {
105:        offset = 0;
106:        p++;
107:    }
108:    barg->offset = offset;
109:    for(n = p; n < ARG_MAX; n++) {
110:        if(!barg->page[n]) {
111:            if(!(barg->page[n] = kmalloc())) {
112:                free_barg_pages(barg);
113:                return -ENOMEM;
114:            }
115:        }
116:    }
117:
118:    /* interpreter */
119:    page = (char *)barg->page[p];
120:    while(*interpreter) {
121:        *(page + offset) = *interpreter;
122:        offset++;
123:        interpreter++;
124:        if(offset == PAGE_SIZE) {
125:            p++;
126:            offset = 0;
127:            page = (char *)barg->page[p];
128:        }
129:    }
130:    *(page + offset++) = NULL;
131:    if(offset == PAGE_SIZE) {

```

## kernel/syscalls/execve.c

Page 3/6

```

132:         p++;
133:         offset = 0;
134:     }
135:
136:     /* args */
137:     page = (char *)barg->page[p];
138:     if(*args) {
139:         while(*args) {
140:             *(page + offset) = *args;
141:             offset++;
142:             args++;
143:             if(offset == PAGE_SIZE) {
144:                 p++;
145:                 offset = 0;
146:                 page = (char *)barg->page[p];
147:             }
148:         }
149:         *(page + offset++) = NULL;
150:         if(offset == PAGE_SIZE) {
151:             p++;
152:             offset = 0;
153:         }
154:     }
155:
156:     /* original script ('filename' with path) at argv[0] */
157:     page = (char *)barg->page[p];
158:     while(*filename) {
159:         *(page + offset) = *filename;
160:         offset++;
161:         filename++;
162:         if(offset == PAGE_SIZE) {
163:             p++;
164:             offset = 0;
165:             page = (char *)barg->page[p];
166:         }
167:     }
168:     *(page + offset) = NULL;
169:
170:     return 0;
171: }
172:
173: static int copy_strings(struct binargs *barg, char *argv[], char *envp[])
174: {
175:     int n, p, offset;
176:     unsigned int ae_str_len;
177:     char *page, *str;
178:
179:     p = ARG_MAX - 1;
180:     ae_str_len = barg->argv_len + barg->envp_len + 4;
181:     p -= ae_str_len / PAGE_SIZE;
182:     offset = PAGE_SIZE - (ae_str_len % PAGE_SIZE);
183:     if(offset == PAGE_SIZE) {
184:         offset = 0;
185:         p++;
186:     }
187:     barg->offset = offset;
188:     for(n = p; n < ARG_MAX; n++) {
189:         if(!(barg->page[n] = kmalloc())) {
190:             free_barg_pages(barg);
191:             return -ENOMEM;
192:         }
193:     }
194:     for(n = 0; n < barg->argc; n++) {
195:         str = argv[n];
196:         page = (char *)barg->page[p];
197:         while(*str) {
198:             *(page + offset) = *str;

```

## kernel/syscalls/execve.c

Page 4/6

```

199:         offset++;
200:         str++;
201:         if(offset == PAGE_SIZE) {
202:             p++;
203:             offset = 0;
204:             page = (char *)barg->page[p];
205:         }
206:     }
207:     *(page + offset++) = NULL;
208:     if(offset == PAGE_SIZE) {
209:         p++;
210:         offset = 0;
211:     }
212: }
213: for(n = 0; n < barg->envc; n++) {
214:     str = envp[n];
215:     page = (char *)barg->page[p];
216:     while(*str) {
217:         *(page + offset) = *str;
218:         offset++;
219:         str++;
220:         if(offset == PAGE_SIZE) {
221:             p++;
222:             offset = 0;
223:             page = (char *)barg->page[p];
224:         }
225:     }
226:     *(page + offset++) = NULL;
227:     if(offset == PAGE_SIZE) {
228:         p++;
229:         offset = 0;
230:     }
231: }
232:
233: return 0;
234: }
235:
236: static int do_execve(const char *filename, char *argv[], char *envp[], struct si
gcontext *sc)
237: {
238:     char interpreter[NAME_MAX + 1], args[NAME_MAX + 1], name[NAME_MAX + 1];
239:     __blk_t block;
240:     struct buffer *buf;
241:     struct inode *i;
242:     struct binargs barg;
243:     char *data, *tmp_name;
244:     int errno;
245:
246:     if((errno = initialize_barg(&barg, &(*argv), &(*envp))) < 0) {
247:         return errno;
248:     }
249:
250:     /* save 'argv' and 'envp' into the kernel address space */
251:     if((errno = copy_strings(&barg, &(*argv), &(*envp))) {
252:         return errno;
253:     }
254:
255:     if(!(data = (void *)kmalloc())) {
256:         return -ENOMEM;
257:     }
258:
259:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
260:         kfree((unsigned int)data);
261:         free_barg_pages(&barg);
262:         return errno;
263:     }
264:     strcpy(name, tmp_name);

```

## kernel/syscalls/execve.c

Page 5/6

```

265:         free_name(tmp_name);
266:
267:
268:  loop:
269:         if((errno = namei(name, &i, NULL, FOLLOW_LINKS)) {
270:             free_barg_pages(&barg);
271:             kfree((unsigned int) data);
272:             return errno;
273:         }
274:
275:         if(!S_ISREG(i->i_mode)) {
276:             iput(i);
277:             free_barg_pages(&barg);
278:             kfree((unsigned int) data);
279:             return -EACCES;
280:         }
281:         if(check_permission(TO_EXEC, i) < 0) {
282:             iput(i);
283:             free_barg_pages(&barg);
284:             kfree((unsigned int) data);
285:             return -EACCES;
286:         }
287:
288:         if((block = bmap(i, 0, FOR_READING)) < 0) {
289:             iput(i);
290:             free_barg_pages(&barg);
291:             kfree((unsigned int) data);
292:             return block;
293:         }
294:         if(!(buf = bread(i->dev, block, i->sb->s_blocksize))) {
295:             iput(i);
296:             free_barg_pages(&barg);
297:             kfree((unsigned int) data);
298:             return -EIO;
299:         }
300:
301:         /*
302:          * The contents of the buffer is copied and then freed immediately to
303:          * make sure that it won't conflict while zeroing the BSS fractional
304:          * page, in case that the same block is requested during the page fault.
305:          */
306:         memcpy_b(data, buf->data, i->sb->s_blocksize);
307:         brelse(buf);
308:
309:         errno = elf_load(i, &barg, sc, data);
310:         if(errno == -ENOEXEC) {
311:             /* OK, looks like it was not an ELF binary; let's see if it is a
script */
312:             memset_b(interpreter, 0, NAME_MAX + 1);
313:             memset_b(args, 0, NAME_MAX + 1);
314:             errno = script_load(interpreter, args, data);
315:             if(!errno) {
316:                 /* yes, it is! */
317:                 iput(i);
318:                 if((errno = add_strings(&barg, name, interpreter, args))
) {
319:                     free_barg_pages(&barg);
320:                     kfree((unsigned int) data);
321:                     return errno;
322:                 }
323:                 strcpy(name, interpreter);
324:                 goto loop;
325:             }
326:         }
327:
328:         if(!errno) {
329:             if(i->i_mode & S_ISUID) {

```

## kernel/syscalls/execve.c

Page 6/6

```

330:             current->euid = i->i_uid;
331:         }
332:         if(i->i_mode & S_ISGID) {
333:             current->egid = i->i_gid;
334:         }
335:     }
336:
337:     iput(i);
338:     free_barg_pages(&barg);
339:     kfree((unsigned int)data);
340:     return errno;
341: }
342:
343: int sys_execve(const char *filename, char *argv[], char *envp[], int arg4, int a
rg5, struct sigcontext *sc)
344: {
345:     char argv0[NAME_MAX + 1];
346:     int n, errno;
347:
348: #ifdef __DEBUG__
349:     printk("(pid %d) sys_execve('%s', ...)\n", current->pid, filename);
350: #endif /*__DEBUG__*/
351:
352:     /* copy filename into kernel address space */
353:     strncpy(argv0, argv[0], NAME_MAX);
354:     if((errno = do_execve(filename, &(*argv), &(*envp), sc)) {
355:         return errno;
356:     }
357:
358:     strncpy(current->argv0, argv0, NAME_MAX);
359:     for(n = 0; n < OPEN_MAX; n++) {
360:         if(current->fd[n] && (current->fd_flags[n] & FD_CLOEXEC)) {
361:             sys_close(n);
362:         }
363:     }
364:
365:     current->suid = current->euid;
366:     current->sgid = current->egid;
367:     current->sigpending = 0;
368:     current->sigexecuting = 0;
369:     for(n = 0; n < NSIG; n++) {
370:         current->sigaction[n].sa_mask = 0;
371:         current->sigaction[n].sa_flags = 0;
372:         if(current->sigaction[n].sa_handler != SIG_IGN) {
373:             current->sigaction[n].sa_handler = SIG_DFL;
374:         }
375:     }
376:     current->sleep_address = NULL;
377:     current->flags |= PF_PEXEC;
378:     return 0;
379: }

```



## kernel/syscalls/exit.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/exit.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/syscalls.h>
11: #include <fiwix/process.h>
12: #include <fiwix/sched.h>
13: #include <fiwix/mman.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: void do_exit(int exit_code)
19: {
20:     int n;
21:     struct proc *p, *init;
22:
23: #ifdef __DEBUG__
24:     printk("\n");
25:     printk("sys_exit(pid %d, ppid %d)\n", current->pid, current->ppid);
26:     printk("-----\n");
27: #endif /*__DEBUG__ */
28:
29:     CLI();
30:     current->state = PROC_ZOMBIE;
31:
32:     release_binary();
33:     current->argv = NULL;
34:     current->envp = NULL;
35:     current->sigpending = 0;
36:     current->sigblocked = 0;
37:     current->sigexecuting = 0;
38:     for(n = 0; n < NSIG; n++) {
39:         current->sigaction[n].sa_mask = 0;
40:         current->sigaction[n].sa_flags = 0;
41:         current->sigaction[n].sa_handler = SIG_IGN;
42:     }
43:
44:     init = get_proc_by_pid(INIT);
45:     FOR_EACH_PROCESS(p) {
46:         if(SESS_LEADER(current)) {
47:             if(p->sid == current->sid && p->state != PROC_ZOMBIE) {
48:                 p->pgid = 0;
49:                 p->sid = 0;
50:                 p->ctty = NULL;
51:                 send_sig(p, SIGHUP);
52:                 send_sig(p, SIGCONT);
53:             }
54:         }
55:
56:         /* make INIT inherit the children of this exiting process */
57:         if(p->state && p->ppid == current->pid) {
58:             p->ppid = INIT;
59:             init->children++;
60:             if(p->state == PROC_ZOMBIE) {
61:                 send_sig(init, SIGCHLD);
62:             }
63:         }
64:     }
65:
66:     if(SESS_LEADER(current)) {
67:         disassociate_ctty(current->ctty);

```

## kernel/syscalls/exit.c

Page 2/2

```

68:         }
69:
70:         for(n = 0; n < OPEN_MAX; n++) {
71:             if(current->fd[n]) {
72:                 sys_close(n);
73:             }
74:         }
75:
76:         iput(current->root);
77:         current->root = NULL;
78:         iput(current->pwd);
79:         current->pwd = NULL;
80:         current->exit_code = exit_code;
81:         if(--nr_processes) {
82:             printk("\n");
83:             printk("WARNING: the last user process has exited. The kernel wi
11 stop itself.\n");
84:             stop_kernel();
85:         }
86:
87:         /* notify the parent about the child's death */
88:         if((p = get_proc_by_pid(current->ppid)) {
89:             send_sig(p, SIGCHLD);
90:             if(p->sleep_address == &sys_wait4) {
91:                 wakeup_proc(p);
92:             }
93:         }
94:
95:         need_resched = 1;
96:
97:         /* make sure to recover if the process returns from the death (!?) */
98:         for(;;) {
99:             current->state = PROC_ZOMBIE;
100:            do_sched();
101:        }
102:    }
103:
104: int sys_exit(int exit_code)
105: {
106: #ifdef __DEBUG__
107:     printk("(pid %d) sys_exit()\n", current->pid);
108: #endif /*__DEBUG__ */
109:
110:     /* exit code in the second byte.
111:      * 15          8 7          0
112:      * +-----+-----+-----+
113:      * | exit code (0-255) |          0          |
114:      * +-----+-----+-----+
115:      */
116:     do_exit((exit_code & 0xFF) << 8);
117:     return 0;
118: }

```

## kernel/syscalls/fchdir.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/fchdir.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/process.h>
10: #include <fiwix/stat.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_fchdir(unsigned int ufd)
18: {
19:     struct inode *i;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_fchdir(%d)\n", current->pid, ufd);
23: #endif /*__DEBUG__*/
24:
25:     CHECK_UFD(ufd);
26:     i = fd_table[current->fd[ufd]].inode;
27:     if(!S_ISDIR(i->i_mode)) {
28:         return -ENOTDIR;
29:     }
30:     iput(current->pwd);
31:     current->pwd = i;
32:     current->pwd->count++;
33:     return 0;
34: }
```

## kernel/syscalls/fchmod.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/fchmod.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/errno.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /*__DEBUG__ */
18:
19: int sys_fchmod(int ufd, __mode_t mode)
20: {
21:     struct inode *i;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_fchmod(%d, %d)\n", current->pid, ufd, mode);
25: #endif /*__DEBUG__ */
26:
27:     CHECK_UFD(ufd);
28:     i = fd_table[current->fd[ufd]].inode;
29:
30:     if(IS_RDONLY_FS(i)) {
31:         return -EROFS;
32:     }
33:     if(check_user_permission(i)) {
34:         return -EPERM;
35:     }
36:
37:     i->i_mode &= S_IFMT;
38:     i->i_mode |= mode & ~S_IFMT;
39:     i->i_ctime = CURRENT_TIME;
40:     i->dirty = 1;
41:     return 0;
42: }
```

## kernel/syscalls/fchown.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/fchown.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/process.h>
13: #include <fiwix/errno.h>
14:
15: #ifdef __DEBUG__
16: #include <fiwix/stdio.h>
17: #endif /*__DEBUG__ */
18:
19: int sys_fchown(int ufd, __uid_t owner, __gid_t group)
20: {
21:     struct inode *i;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_fchown(%d, %d, %d)\n", current->pid, ufd, owner, gr
oup);
25: #endif /*__DEBUG__ */
26:
27:     CHECK_UFD(ufd);
28:     i = fd_table[current->fd[ufd]].inode;
29:
30:     if(IS_RDONLY_FS(i)) {
31:         return -EROFS;
32:     }
33:     if(check_user_permission(i)) {
34:         return -EPERM;
35:     }
36:
37:     if(owner == (__uid_t)-1) {
38:         owner = i->i_uid;
39:     } else {
40:         i->i_mode &= ~(S_ISUID);
41:     }
42:     if(group == (__gid_t)-1) {
43:         group = i->i_gid;
44:     } else {
45:         i->i_mode &= ~(S_ISGID);
46:     }
47:
48:     i->i_uid = owner;
49:     i->i_gid = group;
50:     i->i_ctime = CURRENT_TIME;
51:     i->dirty = 1;
52:     return 0;
53: }

```

## kernel/syscalls/fcntl.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/fcntl.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/syscalls.h>
9: #include <fiwix/fcntl.h>
10: #include <fiwix/locks.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_fcntl(int ufd, int cmd, unsigned long int arg)
19: {
20:     int new_ufd, errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_fcntl(%d, %d, 0x%08x)\n", current->pid, ufd, cmd, a
rg);
24: #endif /* __DEBUG__ */
25:
26:     CHECK_UFD(ufd);
27:     switch(cmd) {
28:         case F_DUPFD:
29:             CHECK_UFD(ufd);
30:             if(arg >= OPEN_MAX) {
31:                 return -EINVAL;
32:             }
33:             if((new_ufd = get_new_user_fd(arg)) < 0) {
34:                 return new_ufd;
35:             }
36:             current->fd[new_ufd] = current->fd[ufd];
37:             fd_table[current->fd[new_ufd]].count++;
38: #ifdef __DEBUG__
39:             printk("\t--> returning %d\n", new_ufd);
40: #endif /* __DEBUG__ */
41:             return new_ufd;
42:         case F_GETFD:
43:             return (current->fd_flags[ufd] & FD_CLOEXEC);
44:         case F_SETFD:
45:             current->fd_flags[ufd] = (arg & FD_CLOEXEC);
46:             break;
47:         case F_GETFL:
48:             return fd_table[current->fd[ufd]].flags;
49:         case F_SETFL:
50:             fd_table[current->fd[ufd]].flags &= ~(O_APPEND | O_NONBL
OCK);
51:             fd_table[current->fd[ufd]].flags |= arg & (O_APPEND | O_
NONBLOCK);
52:             break;
53:         case F_GETLK:
54:         case F_SETLK:
55:         case F_SETLKW:
56:             if((errno = check_user_area(VERIFY_READ, (void *)arg, si
zeof(struct flock)))) {
57:                 return errno;
58:             }
59:             return posix_lock(ufd, cmd, (struct flock *)arg);
60:         default:
61:             return -EINVAL;
62:     }
63:     return 0;

```

```
64: }
```

**kernel/syscalls/fdatasync.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/fdatasync.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/syscalls.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #include <fiwix/process.h>
13: #endif /*__DEBUG__ */
14:
15: int sys_fdatasync(int ufd)
16: {
17: #ifdef __DEBUG__
18:     printk("(pid %d) sys_fdatasync(%d)\n", current->pid, ufd);
19: #endif /*__DEBUG__ */
20:
21:     return sys_fsync(ufd);
22: }
```



**kernel/syscalls/flock.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/flock.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/process.h>
10: #include <fiwix/locks.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /* __DEBUG__ */
16:
17: int sys_flock(int ufd, int op)
18: {
19:     struct inode *i;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_flock(%d, %d)\n", current->pid, ufd, op);
23: #endif /* __DEBUG__ */
24:
25:     CHECK_UFD(ufd);
26:     i = fd_table[current->fd[ufd]].inode;
27:     return flock_inode(i, op);
28: }
```

## kernel/syscalls/fork.c

Page 1/3

```

1: /*
2:  * fiwix/kernel/syscalls/fork.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/types.h>
11: #include <fiwix/const.h>
12: #include <fiwix/sigcontext.h>
13: #include <fiwix/process.h>
14: #include <fiwix/sched.h>
15: #include <fiwix/mm.h>
16: #include <fiwix/errno.h>
17: #include <fiwix/stdio.h>
18: #include <fiwix/string.h>
19:
20: int sys_fork(int arg1, int arg2, int arg3, int arg4, int arg5, struct sigcontext
*sc)
21: {
22:     int count, pages;
23:     unsigned int n;
24:     unsigned int *child_pgid;
25:     struct sigcontext *stack;
26:     struct proc *child, *p;
27:     struct vma *vma;
28:     __pid_t pid;
29:
30: #ifdef __DEBUG__
31:     printk("(pid %d) sys_fork()\n", current->pid);
32: #endif /*__DEBUG__ */
33:
34:     /* check the number of processes already allocated by this UID */
35:     count = 0;
36:     FOR_EACH_PROCESS(p) {
37:         if(p->state && p->uid == current->uid) {
38:             count++;
39:         }
40:     }
41:     if(count > current->rlim[RLIMIT_NPROC].rlim_cur) {
42:         printk("WARNING: %s(): RLIMIT_NPROC exceeded.\n", __FUNCTION__);
43:         return -EAGAIN;
44:     }
45:
46:     if(!(pid = get_unused_pid())) {
47:         return -EAGAIN;
48:     }
49:     if(!(child = get_proc_free())) {
50:         return -EAGAIN;
51:     }
52:
53:     /*
54:      * This memcpy() will overwrite the prev and next pointers, so that's
55:      * the reason why proc_slot_init() is separated from get_proc_free().
56:      */
57:     memcpy_b(child, current, sizeof(struct proc));
58:
59:     proc_slot_init(child);
60:     child->pid = pid;
61:     memset_b(&child->tss, NULL, sizeof(struct i386tss));
62:     sprintf(child->pidstr, "%d", child->pid);
63:     child->state = PROC_IDLE;
64:
65:     if(!(child_pgid = (void *)kmalloc())) {
66:         release_proc(child);

```

## kernel/syscalls/fork.c

Page 2/3

```

67:             return -ENOMEM;
68:         }
69:         child->rss++;
70:         memcpy_b(child_pgdirt, kpage_dir, PAGE_SIZE);
71:         child->tss.cr3 = V2P((unsigned int) child_pgdirt);
72:
73:         child->ppid = current->pid;
74:         child->flags = 0;
75:         child->children = 0;
76:         child->cpu_count = child->priority;
77:         child->start_time = CURRENT_TICKS;
78:         child->sleep_address = NULL;
79:
80:         memcpy_b(child->vma, current->vma, sizeof(child->vma));
81:         vma = child->vma;
82:         for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
83:             if(vma->inode) {
84:                 vma->inode->count++;
85:             }
86:         }
87:
88:         child->sigpending = 0;
89:         child->sigexecuting = 0;
90:         memset_b(&child->sc, NULL, sizeof(struct sigcontext));
91:         memset_b(&child->usage, NULL, sizeof(struct rusage));
92:         memset_b(&child->cusage, NULL, sizeof(struct rusage));
93:         child->it_real_interval = 0;
94:         child->it_real_value = 0;
95:         child->it_virt_interval = 0;
96:         child->it_virt_value = 0;
97:         child->it_prof_interval = 0;
98:         child->it_prof_value = 0;
99:
100:        if(!(child->tss.esp0 = kmalloc())) {
101:            kfree((unsigned int) child_pgdirt);
102:            kfree((unsigned int) child->vma);
103:            release_proc(child);
104:            return -ENOMEM;
105:        }
106:
107:        if(!(pages = clone_pages(child))) {
108:            printk("WARNING: %s(): not enough memory, can't clone pages.\n",
__FUNCTION__);
109:            free_page_tables(child);
110:            kfree((unsigned int) child_pgdirt);
111:            kfree((unsigned int) child->vma);
112:            release_proc(child);
113:            return -ENOMEM;
114:        }
115:        child->rss += pages;
116:        invalidate_tlb();
117:
118:        child->tss.esp0 += PAGE_SIZE - 4;
119:        child->rss++;
120:        child->tss.ss0 = KERNEL_DS;
121:
122:        memcpy_b((unsigned int *) (child->tss.esp0 & PAGE_MASK), (void *) ((unsigned
ed int) (sc) & PAGE_MASK), PAGE_SIZE);
123:        stack = (struct sigcontext *) ((child->tss.esp0 & PAGE_MASK) + ((unsigned
int) (sc) & ~PAGE_MASK));
124:
125:        child->tss.eip = (unsigned int) return_from_syscall;
126:        child->tss.esp = (unsigned int) stack;
127:        stack->eax = 0;          /* child returns 0 */
128:
129:        child->state = PROC_RUNNING;
130:

```

## kernel/syscalls/fork.c

Page 3/3

```
131:      /* increase file descriptors usage */
132:      for(n = 0; n < OPEN_MAX; n++) {
133:          if(current->fd[n]) {
134:              fd_table[current->fd[n]].count++;
135:          }
136:      }
137:      if(current->root) {
138:          current->root->count++;
139:      }
140:      if(current->pwd) {
141:          current->pwd->count++;
142:      }
143:
144:      kstat.processes++;
145:      nr_processes++;
146:      current->children++;
147:
148:      return child->pid;      /* parent returns child's PID */
149: }
```

## kernel/syscalls/fstat.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/fstat.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/syscalls.h>
10: #include <fiwix/statbuf.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_fstat(unsigned int ufd, struct old_stat *statbuf)
19: {
20:     struct inode *i;
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_fstat(%d, 0x%08x) -> returning structure\n", current->pid, ufd, (unsigned int) statbuf);
25: #endif /* __DEBUG__ */
26:
27:     CHECK_UFD(ufd);
28:     if((errno = check_user_area(VERIFY_WRITE, statbuf, sizeof(struct old_stat)))) {
29:         return errno;
30:     }
31:     i = fd_table[current->fd[ufd]].inode;
32:     statbuf->st_dev = i->dev;
33:     statbuf->st_ino = i->ino;
34:     statbuf->st_mode = i->i_mode;
35:     statbuf->st_nlink = i->i_nlink;
36:     statbuf->st_uid = i->i_uid;
37:     statbuf->st_gid = i->i_gid;
38:     statbuf->st_rdev = i->rdev;
39:     statbuf->st_size = i->i_size;
40:     statbuf->st_atime = i->i_atime;
41:     statbuf->st_mtime = i->i_mtime;
42:     statbuf->st_ctime = i->i_ctime;
43:     return 0;
44: }

```

## kernel/syscalls/fstatfs.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/fstatfs.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/statfs.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /* __DEBUG__ */
16:
17: int sys_fstatfs(unsigned int ufd, struct statfs *statfsbuf)
18: {
19:     struct inode *i;
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_fstatfs(%d, 0x%08x)\n", current->pid, ufd, (unsigned
d int)statfsbuf);
24: #endif /* __DEBUG__ */
25:
26:     CHECK_UFD(ufd);
27:     if((errno = check_user_area(VERIFY_WRITE, statfsbuf, sizeof(struct statf
s)))) {
28:         return errno;
29:     }
30:     i = fd_table[current->fd[ufd]].inode;
31:     if(i->sb && i->sb->fsop && i->sb->fsop->statfs) {
32:         i->sb->fsop->statfs(i->sb, statfsbuf);
33:         return 0;
34:     }
35:     return -ENOSYS;
36: }
```

## kernel/syscalls/fsync.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/fsync.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/filesystems.h>
10: #include <fiwix/process.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/buffer.h>
13: #include <fiwix/errno.h>
14:
15: #ifdef __DEBUG__
16: #include <fiwix/stdio.h>
17: #endif /*__DEBUG__*/
18:
19: int sys_fsync(int ufd)
20: {
21:     struct inode *i;
22:
23: #ifdef __DEBUG__
24:     printk(" (pid %d) sys_fsync(%d)\n", current->pid, ufd);
25: #endif /*__DEBUG__*/
26:
27:     CHECK_UFD(ufd);
28:     i = fd_table[current->fd[ufd]].inode;
29:     if(!S_ISREG(i->i_mode)) {
30:         return -EINVAL;
31:     }
32:     if(IS_RDONLY_FS(i)) {
33:         return -EROFS;
34:     }
35:     sync_superblocks(i->dev);
36:     sync_inodes(i->dev);
37:     sync_buffers(i->dev);
38:     return 0;
39: }
```

## kernel/syscalls/ftime.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/ftime.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/timeb.h>
11: #include <fiwix/timer.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_ftime(struct timeb *tp)
19: {
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_ftime()\n", current->pid);
24: #endif /* __DEBUG__ */
25:
26:     if((errno = check_user_area(VERIFY_WRITE, tp, sizeof(struct timeb)))) {
27:         return errno;
28:     }
29:     tp->time = CURRENT_TIME;
30:     tp->millitm = ((kstat.ticks % HZ) * 1000000) / HZ;
31:     /* FIXME: 'timezone' and 'dstflag' fields are not used */
32:
33:     return 0;
34: }
```



## kernel/syscalls/ftruncate.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/ftruncate.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/fcntl.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/errno.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /*__DEBUG__*/
18:
19: int sys_ftruncate(int ufd, __off_t length)
20: {
21:     struct inode *i;
22:     int errno;
23:
24: #ifdef __DEBUG__
25:     printk("(pid %d) sys_ftruncate(%d, %d)\n", current->pid, ufd, length);
26: #endif /*__DEBUG__*/
27:
28:     CHECK_UFD(ufd);
29:     i = fd_table[current->fd[ufd]].inode;
30:     if((fd_table[current->fd[ufd]].flags & O_ACCMODE) == O_RDONLY) {
31:         return -EINVAL;
32:     }
33:     if(S_ISDIR(i->i_mode)) {
34:         return -EISDIR;
35:     }
36:     if(IS_RDONLY_FS(i)) {
37:         return -EROFS;
38:     }
39:     if(check_permission(TO_WRITE, i) < 0) {
40:         return -EPERM;
41:     }
42:     if(length == i->i_size) {
43:         return 0;
44:     }
45:
46:     errno = 0;
47:     if(i->fsop && i->fsop->truncate) {
48:         inode_lock(i);
49:         errno = i->fsop->truncate(i, length);
50:         inode_unlock(i);
51:     }
52:     return errno;
53: }

```

**kernel/syscalls/getcwd.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getcwd.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_getcwd(char *buf, __size_t size)
18: {
19:     int errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_getcwd(0x%08x, %d)\n", current->pid, (unsigned int)
buf, size);
23: #endif /*__DEBUG__*/
24:
25:     if((errno = check_user_area(VERIFY_WRITE, buf, size)) {
26:         return errno;
27:     }
28:     return -ENOSYS;
29: }
```

## kernel/syscalls/getdents.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/getdents.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/dirent.h>
10: #include <fiwix/process.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/errno.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_getdents(unsigned int ufd, struct dirent *dirent, unsigned int count)
19: {
20:     struct inode *i;
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_getdents(%d, 0x%08x, %d)", current->pid, ufd, (unsi
gned int)dirent, count);
25: #endif /* __DEBUG__ */
26:
27:     CHECK_UFD(ufd);
28:     if((errno = check_user_area(VERIFY_WRITE, dirent, sizeof(struct dirent))
)) {
29:         return errno;
30:     }
31:     i = fd_table[current->fd[ufd]].inode;
32:
33:     if(!S_ISDIR(i->i_mode)) {
34:         return -ENOTDIR;
35:     }
36:
37:     if(i->fsop && i->fsop->readdir) {
38:         errno = i->fsop->readdir(i, &fd_table[current->fd[ufd]], dirent,
count);
39:     }
40:     #ifdef __DEBUG__
41:     printk(" -> returning %d\n", errno);
42:     #endif /* __DEBUG__ */
43:     return errno;
44: }
45: }

```

**kernel/syscalls/getegid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getegid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #endif /*__DEBUG__*/
13:
14: int sys_getegid(void)
15: {
16: #ifdef __DEBUG__
17:     printk("(pid %d) sys_getegid() -> %d\n", current->pid, current->egid);
18: #endif /*__DEBUG__*/
19:     return current->egid;
20: }
```

**kernel/syscalls/geteuid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/geteuid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #endif /*__DEBUG__*/
13:
14: int sys_geteuid(void)
15: {
16: #ifdef __DEBUG__
17:     printk("(pid %d) sys_geteuid() -> %d\n", current->pid, current->euid);
18: #endif /*__DEBUG__*/
19:     return current->euid;
20: }
```

**kernel/syscalls/getgid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getgid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #endif /*__DEBUG__ */
13:
14: int sys_getgid(void)
15: {
16: #ifdef __DEBUG__
17:     printk("(pid %d) sys_getgid() -> %d\n", current->pid, current->gid);
18: #endif /*__DEBUG__ */
19:     return current->gid;
20: }
```

## kernel/syscalls/getgroups.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getgroups.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__ */
16:
17: int sys_getgroups(__ssize_t size, __gid_t *list)
18: {
19:     int n, errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_getgroups(%d, 0x%08x)\n", current->pid, size, (unsi
gned int)list);
23: #endif /*__DEBUG__ */
24:
25:     /*
26:      * If size is 0, sys_getgroups() shall return the number of group IDs
27:      * that it would otherwise return without modifying the array pointed
28:      * to by list.
29:      */
30:     if(!size) {
31:         for(n = 0; n < NGROUPS_MAX; n++) {
32:             if(current->groups[n] == -1) {
33:                 break;
34:             }
35:         }
36:         return n;
37:     }
38:
39:     if((errno = check_user_area(VERIFY_WRITE, list, sizeof(__gid_t)))) {
40:         return errno;
41:     }
42:     for(n = 0; n < NGROUPS_MAX; n++) {
43:         if(current->groups[n] == -1) {
44:             break;
45:         }
46:         if(size) {
47:             if(n > size) {
48:                 return -EINVAL;
49:             }
50:             list[n] = (__gid_t)current->groups[n];
51:         }
52:     }
53:     return n;
54: }
```

## kernel/syscalls/getitimer.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/getitimer.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/time.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_getitimer(int which, struct itimerval *curr_value)
18: {
19:     int errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_getitimer(%d, 0x%08x) -> \n", current->pid, which,
(unsigned int)curr_value);
23: #endif /*__DEBUG__*/
24:
25:     if((unsigned int)curr_value) {
26:         if((errno = check_user_area(VERIFY_WRITE, curr_value, sizeof(str
uct itimerval)))) {
27:             return errno;
28:         }
29:     }
30:
31:     switch(which) {
32:         case ITIMER_REAL:
33:             ticks2tv(current->it_real_interval, &curr_value->it_inte
rval);
34:             ticks2tv(current->it_real_value, &curr_value->it_value);
35:             break;
36:         case ITIMER_VIRTUAL:
37:             ticks2tv(current->it_virt_interval, &curr_value->it_inte
rval);
38:             ticks2tv(current->it_virt_value, &curr_value->it_value);
39:             break;
40:         case ITIMER_PROF:
41:             ticks2tv(current->it_prof_interval, &curr_value->it_inte
rval);
42:             ticks2tv(current->it_prof_value, &curr_value->it_value);
43:             break;
44:         default:
45:             return -EINVAL;
46:     }
47:     return 0;
48: }

```



## kernel/syscalls/getpgid.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getpgid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/process.h>
10: #include <fiwix/sched.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_getpgid(__pid_t pid)
18: {
19:     struct proc *p;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_getpgid(%d)\n", current->pid, pid);
23: #endif /*__DEBUG__*/
24:
25:     if(pid < 0) {
26:         return -EINVAL;
27:     }
28:     if(!pid) {
29:         return current->pgid;
30:     }
31:     FOR_EACH_PROCESS(p) {
32:         if(p->state != PROC_UNUSED && p->pid == pid) {
33:             return p->pgid;
34:         }
35:     }
36:     return -ESRCH;
37: }
```

**kernel/syscalls/getpgrp.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getpgrp.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #endif /*__DEBUG__*/
13:
14: int sys_getpgrp(void)
15: {
16: #ifdef __DEBUG__
17:     printk("(pid %d) sys_getpgrp() -> %d\n", current->pid, current->pgid);
18: #endif /*__DEBUG__*/
19:     return current->pgid;
20: }
```

**kernel/syscalls/getpid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getpid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #endif /*__DEBUG__*/
13:
14: int sys_getpid(void)
15: {
16: #ifdef __DEBUG__
17:     printk("(pid %d) sys_getpid() -> %d\n", current->pid, current->pid);
18: #endif /*__DEBUG__*/
19:     return current->pid;
20: }
```

**kernel/syscalls/getppid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getppid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #endif /*__DEBUG__*/
13:
14: int sys_getppid(void)
15: {
16: #ifdef __DEBUG__
17:     printk("(pid %d) sys_getppid() -> %d\n", current->pid, current->ppid);
18: #endif /*__DEBUG__*/
19:     return current->ppid;
20: }
```

## kernel/syscalls/getrlimit.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getrlimit.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/resource.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_getrlimit(int resource, struct rlimit *rlim)
18: {
19:     int errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_getrlimit(%d, 0x%08x)\n", current->pid, resource, (
unsigned int)rlim);
23: #endif /*__DEBUG__*/
24:
25:     if((errno = check_user_area(VERIFY_WRITE, rlim, sizeof(struct rlimit))))
{
26:         return errno;
27:     }
28:     if(resource < 0 || resource >= RLIM_NLIMITS) {
29:         return -EINVAL;
30:     }
31:
32:     rlim->rlim_cur = current->rlim[resource].rlim_cur;
33:     rlim->rlim_max = current->rlim[resource].rlim_max;
34:     return 0;
35: }
```

## kernel/syscalls/getrusage.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/getrusage.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/resource.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_getrusage(int who, struct rusage *usage)
19: {
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_getrusage(%d, 0x%08x)\n", current->pid, who, (unsigned int)usage);
24: #endif /* __DEBUG__ */
25:
26:     if((errno = check_user_area(VERIFY_WRITE, usage, sizeof(struct rusage)))
) {
27:         return errno;
28:     }
29:     switch(who) {
30:         case RUSAGE_SELF:
31:             memcpy_b(usage, &current->usage, sizeof(struct rusage));
32:             break;
33:         case RUSAGE_CHILDREN:
34:             memcpy_b(usage, &current->cusage, sizeof(struct rusage))
;
35:             break;
36:         default:
37:             return -EINVAL;
38:     }
39:     return 0;
40: }

```

## kernel/syscalls/getsid.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getsid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/process.h>
10: #include <fiwix/sched.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_getsid(__pid_t pid)
18: {
19:     struct proc *p;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_getsid(%d)\n", current->pid, pid);
23: #endif /*__DEBUG__*/
24:
25:     if(pid < 0) {
26:         return -EINVAL;
27:     }
28:     if(!pid) {
29:         return current->sid;
30:     }
31:
32:     FOR_EACH_PROCESS(p) {
33:         if(p->state != PROC_UNUSED && p->pid == pid) {
34:             return p->sid;
35:         }
36:     }
37:     return -ESRCH;
38: }
```

## kernel/syscalls/gettimeofday.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/gettimeofday.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/process.h>
11: #include <fiwix/time.h>
12: #include <fiwix/timer.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_gettimeofday(struct timeval *tv, struct timezone *tz)
19: {
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_gettimeofday()\n", current->pid);
24: #endif /* __DEBUG__ */
25:
26:     if(tv) {
27:         if((errno = check_user_area(VERIFY_WRITE, tv, sizeof(struct time
val)))) {
28:             return errno;
29:         }
30:         tv->tv_sec = CURRENT_TIME;
31:         tv->tv_usec = ((kstat.ticks % HZ) * 1000000) / HZ;
32:     }
33:     if(tz) {
34:         if((errno = check_user_area(VERIFY_WRITE, tz, sizeof(struct time
zone)))) {
35:             return errno;
36:         }
37:         tz->tz_minuteswest = kstat.tz_minuteswest;
38:         tz->tz_dsttime = kstat.tz_dsttime;
39:     }
40:     return 0;
41: }
```



**kernel/syscalls/getuid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/getuid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #endif /*__DEBUG__*/
13:
14: int sys_getuid(void)
15: {
16: #ifdef __DEBUG__
17:     printk("(pid %d) sys_getuid() -> %d\n", current->pid, current->uid);
18: #endif /*__DEBUG__*/
19:
20:     return current->uid;
21: }
```

## kernel/syscalls/ioctl.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/ioctl.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9: #include <fiwix/errno.h>
10:
11: #ifdef __DEBUG__
12: #include <fiwix/stdio.h>
13: #endif /*__DEBUG__*/
14:
15: int sys_ioctl(unsigned int fd, int cmd, unsigned long int arg)
16: {
17:     int errno;
18:     struct inode *i;
19:
20: #ifdef __DEBUG__
21:     printk("(pid %d) sys_ioctl(%d, 0x%x, 0x%08x) -> ", current->pid, fd, cmd
, arg);
22: #endif /*__DEBUG__*/
23:
24:     CHECK_UFD(fd);
25:     i = fd_table[current->fd[fd]].inode;
26:     if(i->fsop && i->fsop->ioctl) {
27:         errno = i->fsop->ioctl(i, cmd, arg);
28:
29: #ifdef __DEBUG__
30:         printk("%d\n", errno);
31: #endif /*__DEBUG__*/
32:
33:         return errno;
34:     }
35:
36: #ifdef __DEBUG__
37:     printk("%d\n", -ENOTTY);
38: #endif /*__DEBUG__*/
39:
40:     return -ENOTTY;
41: }
```

## kernel/syscalls/ioperm.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/ioperm.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9: #include <fiwix/errno.h>
10:
11: #ifdef __DEBUG__
12: #include <fiwix/stdio.h>
13: #endif /*__DEBUG__ */
14:
15: int sys_ioperm(unsigned long int from, unsigned long int num, int turn_on)
16: {
17: #ifdef __DEBUG__
18:     printk("(pid %d) sys_ioperm(0x%08x, 0x%08x, 0x%08x)\n", current->pid, fr
om, num, turn_on);
19: #endif /*__DEBUG__ */
20:
21:     if(!IS_SUPERUSER) {
22:         return -EPERM;
23:     }
24:
25:     /* FIXME: to be implemented */
26:
27:     return 0;
28: }
```

## kernel/syscalls/iopl.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/iopl.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: /*
9:  * Chapter Input/Output of IA-32 Intel(R) Architecture Software Developer's
10:  * Manual Volume 1 Basic Architecture, says the processor permits applications
11:  * to access I/O ports in either of two ways: by using I/O address space or by
12:  * using memory-mapped I/O. Linux 2.0 and Fiwix uses the first one.
13:  *
14:  * This system call sets the IOPL field in the EFLAGS register to the value of
15:  * 'level' (which is presumably zero), so the current process will have
16:  * privileges to use any port, even if that port is beyond of the default size
17:  * of the I/O bitmap in TSS (which is IO_BITMAP_SIZE = 32). Otherwise the
18:  * processor checks the I/O permission bit map to determine if access to a
19:  * specific I/O port is allowed.
20:  *
21:  * So, we leave it here as in Linux 2.0. That means, leaving to I/O bit map to
22:  * control the ports up to 0x3FF, and the rest of ports will be controlled by
23:  * using this system call.
24:  */
25:
26: #include <fiwix/process.h>
27: #include <fiwix/segments.h>
28: #include <fiwix/sigcontext.h>
29: #include <fiwix/errno.h>
30:
31: #ifdef __DEBUG__
32: #include <fiwix/stdio.h>
33: #endif /*__DEBUG__*/
34:
35: int sys_iopl(int level, int arg2, int arg3, int arg4, int arg5, struct sigcontext
t *sc)
36: {
37: #ifdef __DEBUG__
38:     printk("(pid %d) sys_iopl(%d) -> ", current->pid, level);
39: #endif /*__DEBUG__*/
40:     if(level > USR_PL) {
41: #ifdef __DEBUG__
42:         printk("--EINVAL\n");
43: #endif /*__DEBUG__*/
44:         return -EINVAL;
45:     }
46:     if(!IS_SUPERUSER) {
47: #ifdef __DEBUG__
48:         printk("--EPERM\n");
49: #endif /*__DEBUG__*/
50:         return -EPERM;
51:     }
52:
53:     sc->eflags = (sc->eflags & 0xFFFFCFFF) | (level << EF_IOPL);
54: #ifdef __DEBUG__
55:     printk("0\n");
56: #endif /*__DEBUG__*/
57:     return 0;
58: }

```

## kernel/syscalls/kill.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/kill.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/process.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #endif /*__DEBUG__*/
15:
16: int sys_kill(__pid_t pid, __sigset_t signum)
17: {
18:     int count;
19:     struct proc *p;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_kill(%d, %d)\n", current->pid, pid, signum);
23: #endif /*__DEBUG__*/
24:
25:     if(signum < 1 || signum > NSIG) {
26:         return -EINVAL;
27:     }
28:     if(pid == -1) {
29:         count = 0;
30:         FOR_EACH_PROCESS(p) {
31:             if(p->pid > 1 && p != current) {
32:                 count++;
33:                 send_sig(p, signum);
34:             }
35:         }
36:         return count ? 0 : -ESRCH;
37:     }
38:     if(!pid) {
39:         return kill_pgrp(current->pgrp, signum);
40:     }
41:     if(pid < 1) {
42:         return kill_pgrp(-pid, signum);
43:     }
44:
45:     return kill_pid(pid, signum);
46: }
```

## kernel/syscalls/link.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/link.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/stat.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_link(const char *oldname, const char *newname)
19: {
20:     struct inode *i, *dir, *i_new, *dir_new;
21:     char *tmp_oldname, *tmp_newname, *basename;
22:     int errno;
23:
24: #ifdef __DEBUG__
25:     printk("(pid %d) sys_link('%s', '%s')\n", current->pid, oldname, newname
);
26: #endif /* __DEBUG__ */
27:
28:     if((errno = malloc_name(oldname, &tmp_oldname)) < 0) {
29:         return errno;
30:     }
31:     if((errno = malloc_name(newname, &tmp_newname)) < 0) {
32:         free_name(tmp_oldname);
33:         return errno;
34:     }
35:
36:     if((errno = namei(tmp_oldname, &i, &dir, !FOLLOW_LINKS))) {
37:         if(dir) {
38:             iput(dir);
39:         }
40:         free_name(tmp_oldname);
41:         free_name(tmp_newname);
42:         return errno;
43:     }
44:     if(S_ISDIR(i->i_mode)) {
45:         iput(i);
46:         iput(dir);
47:         free_name(tmp_oldname);
48:         free_name(tmp_newname);
49:         return -EPERM;
50:     }
51:     if(IS_RDONLY_FS(i)) {
52:         iput(i);
53:         iput(dir);
54:         free_name(tmp_oldname);
55:         free_name(tmp_newname);
56:         return -EROFS;
57:     }
58:     if(i->i_nlink == LINK_MAX) {
59:         iput(i);
60:         iput(dir);
61:         free_name(tmp_oldname);
62:         free_name(tmp_newname);
63:         return -EMLINK;
64:     }
65:
66:     basename = get_basename(tmp_newname);

```

## kernel/syscalls/link.c

Page 2/2

```
67:         if((errno = namei(tmp_newname, &i_new, &dir_new, !FOLLOW_LINKS))) {
68:             if(!dir_new) {
69:                 iput(i);
70:                 iput(dir);
71:                 free_name(tmp_oldname);
72:                 free_name(tmp_newname);
73:                 return errno;
74:             }
75:         }
76:         if(!errno) {
77:             iput(i);
78:             iput(dir);
79:             iput(i_new);
80:             iput(dir_new);
81:             free_name(tmp_oldname);
82:             free_name(tmp_newname);
83:             return -EEXIST;
84:         }
85:         if(i->dev != dir_new->dev) {
86:             iput(i);
87:             iput(dir);
88:             iput(dir_new);
89:             free_name(tmp_oldname);
90:             free_name(tmp_newname);
91:             return -EXDEV;
92:         }
93:         if(check_permission(TO_EXEC | TO_WRITE, dir_new) < 0) {
94:             iput(i);
95:             iput(dir);
96:             iput(dir_new);
97:             free_name(tmp_oldname);
98:             free_name(tmp_newname);
99:             return -EACCES;
100:        }
101:
102:        if(dir_new->fsop && dir_new->fsop->link) {
103:            errno = dir_new->fsop->link(i, dir_new, basename);
104:        } else {
105:            errno = -EPERM;
106:        }
107:        iput(i);
108:        iput(dir);
109:        iput(dir_new);
110:        free_name(tmp_oldname);
111:        free_name(tmp_newname);
112:        return errno;
113:    }
```

## kernel/syscalls/llseek.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/llseek.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_llseek(unsigned int ufd, unsigned long int offset_high, unsigned long in
t offset_low, __loff_t *result, unsigned int whence)
19: {
20:     struct inode *i;
21:     __loff_t offset;
22:     __loff_t new_offset;
23:     int errno;
24:
25: #ifdef __DEBUG__
26:     printk(" (pid %d) sys_llseek(%d, %u, %u, %08x, %d)", current->pid, ufd, o
ffset_high, offset_low, result, whence);
27: #endif /* __DEBUG__ */
28:
29:     CHECK_UFD(ufd);
30:     if((errno = check_user_area(VERIFY_WRITE, result, sizeof(__loff_t))) {
31:         return errno;
32:     }
33:     i = fd_table[current->fd[ufd]].inode;
34:     offset = (__loff_t) (((__loff_t)offset_high << 32) | offset_low);
35:     switch(whence) {
36:         case SEEK_SET:
37:             new_offset = offset;
38:             break;
39:         case SEEK_CUR:
40:             new_offset = fd_table[current->fd[ufd]].offset + offset;
41:             break;
42:         case SEEK_END:
43:             new_offset = i->i_size + offset;
44:             break;
45:         default:
46:             return -EINVAL;
47:     }
48:     fd_table[current->fd[ufd]].offset = new_offset;
49:
50:     memcpy_b(result, &new_offset, sizeof(__loff_t));
51:
52: #ifdef __DEBUG__
53:     printk(" -> returning %u\n", *result);
54: #endif /* __DEBUG__ */
55:
56:     return 0;
57: }

```



## kernel/syscalls/lseek.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/lseek.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/syscalls.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_lseek(unsigned int ufd, __off_t offset, unsigned int whence)
19: {
20:     struct inode *i;
21:     __off_t new_offset;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_lseek(%d, %d, %d)", current->pid, ufd, offset, when
ce);
25: #endif /* __DEBUG__ */
26:
27:     CHECK_UFD(ufd);
28:
29:     i = fd_table[current->fd[ufd]].inode;
30:     switch(whence) {
31:         case SEEK_SET:
32:             new_offset = offset;
33:             break;
34:         case SEEK_CUR:
35:             new_offset = fd_table[current->fd[ufd]].offset + offset;
36:             break;
37:         case SEEK_END:
38:             new_offset = i->i_size + offset;
39:             break;
40:         default:
41:             return -EINVAL;
42:     }
43:     if((int)new_offset < 0) {
44:         return -EINVAL;
45:     }
46:     if(i->fsop && i->fsop->lseek) {
47:         fd_table[current->fd[ufd]].offset = new_offset;
48:         new_offset = i->fsop->lseek(i, new_offset);
49:     } else {
50:         return -EPERM;
51:     }
52:
53: #ifdef __DEBUG__
54:     printk(" -> returning %d\n", new_offset);
55: #endif /* __DEBUG__ */
56:
57:     return new_offset;
58: }

```

## kernel/syscalls/lstat.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/lstat.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/stat.h>
10: #include <fiwix/string.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_lstat(const char *filename, struct old_stat *statbuf)
18: {
19:     struct inode *i;
20:     char *tmp_name;
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_lstat('%s', 0x%08x) -> returning structure\n", curr
ent->pid, filename, (unsigned int)statbuf);
25: #endif /*__DEBUG__*/
26:
27:     if((errno = check_user_area(VERIFY_WRITE, statbuf, sizeof(struct old_sta
t)))) {
28:         return errno;
29:     }
30:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
31:         return errno;
32:     }
33:     if((errno = namei(tmp_name, &i, NULL, !FOLLOW_LINKS))) {
34:         free_name(tmp_name);
35:         return errno;
36:     }
37:     statbuf->st_dev = i->dev;
38:     statbuf->st_ino = i->inode;
39:     statbuf->st_mode = i->i_mode;
40:     statbuf->st_nlink = i->i_nlink;
41:     statbuf->st_uid = i->i_uid;
42:     statbuf->st_gid = i->i_gid;
43:     statbuf->st_rdev = i->rdev;
44:     statbuf->st_size = i->i_size;
45:     statbuf->st_atime = i->i_atime;
46:     statbuf->st_mtime = i->i_mtime;
47:     statbuf->st_ctime = i->i_ctime;
48:     iput(i);
49:     free_name(tmp_name);
50:     return 0;
51: }

```

**kernel/syscalls/Makefile**

Page 1/1

```
1: # fiwix/kernel/syscalls/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: SRC = $(wildcard *.c)
13: OBJS = $(patsubst %.c,%.o,$(SRC))
14:
15: lib:     $(OBJS)
16:           $(LD) $(LDFLAGS) -r $(OBJS) -o syscalls.o
17:
18: clean:
19:           rm -f *.o
20:
```

## kernel/syscalls/mkdir.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/mkdir.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/stat.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /*__DEBUG__ */
18:
19: int sys_mkdir(const char *dirname, __mode_t mode)
20: {
21:     struct inode *i, *dir;
22:     char *tmp_dirname, *basename;
23:     int errno;
24:
25: #ifdef __DEBUG__
26:     printk("(pid %d) sys_mkdir('%s', %o)\n", current->pid, dirname, mode);
27: #endif /*__DEBUG__ */
28:
29:     if((errno = malloc_name(dirname, &tmp_dirname)) < 0) {
30:         return errno;
31:     }
32:     basename = remove_trailing_slash(tmp_dirname);
33:     if((errno = namei(basename, &i, &dir, !FOLLOW_LINKS)) {
34:         if(!dir) {
35:             free_name(tmp_dirname);
36:             return errno;
37:         }
38:     }
39:     if(!errno) {
40:         iput(i);
41:         iput(dir);
42:         free_name(tmp_dirname);
43:         return -EEXIST;
44:     }
45:     if(IS_RDONLY_FS(dir)) {
46:         iput(dir);
47:         free_name(tmp_dirname);
48:         return -EROFS;
49:     }
50:
51:     if(check_permission(TO_EXEC | TO_WRITE, dir) < 0) {
52:         iput(dir);
53:         free_name(tmp_dirname);
54:         return -EACCES;
55:     }
56:
57:     basename = get_basename(basename);
58:     if(dir->fsop && dir->fsop->mkdir) {
59:         errno = dir->fsop->mkdir(dir, basename, mode);
60:     } else {
61:         errno = -EPERM;
62:     }
63:     iput(dir);
64:     free_name(tmp_dirname);
65:     return errno;
66: }

```

## kernel/syscalls/mknod.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/mknod.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/stat.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /*__DEBUG__ */
18:
19: int sys_mknod(const char *pathname, __mode_t mode, __dev_t dev)
20: {
21:     struct inode *i, *dir;
22:     char *tmp_name, *basename;
23:     int errno;
24:
25: #ifdef __DEBUG__
26:     printk("(pid %d) sys_mknod('%s', %d, %x)\n", current->pid, pathname, mod
e, dev);
27: #endif /*__DEBUG__ */
28:
29:     if(!S_ISCHR(mode) && !S_ISBLK(mode) && !S_ISFIFO(mode)) {
30:         return -EINVAL;
31:     }
32:     if(!S_ISFIFO(mode) && !IS_SUPERUSER) {
33:         return -EPERM;
34:     }
35:
36:     if((errno = malloc_name(pathname, &tmp_name)) < 0) {
37:         return errno;
38:     }
39:     basename = get_basename(tmp_name);
40:     if((errno = namei(tmp_name, &i, &dir, !FOLLOW_LINKS))) {
41:         if(!dir) {
42:             free_name(tmp_name);
43:             return errno;
44:         }
45:     }
46:     if(!errno) {
47:         iput(i);
48:         iput(dir);
49:         free_name(tmp_name);
50:         return -EEXIST;
51:     }
52:     if(IS_RDONLY_FS(dir)) {
53:         iput(dir);
54:         free_name(tmp_name);
55:         return -EROFS;
56:     }
57:     if(check_permission(TO_EXEC | TO_WRITE, dir) < 0) {
58:         iput(dir);
59:         free_name(tmp_name);
60:         return -EACCES;
61:     }
62:
63:     if(dir->fsop && dir->fsop->mknod) {
64:         errno = dir->fsop->mknod(dir, basename, mode, dev);
65:     } else {
66:         errno = -EPERM;

```

**kernel/syscalls/mknod.c**

Page 2/2

```
67:         }
68:         iput(dir);
69:         free_name(tmp_name);
70:         return errno;
71: }
```

## kernel/syscalls/mount.c

Page 1/4

```

1: /*
2:  * fiwix/kernel/syscalls/mount.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/stat.h>
11: #include <fiwix/buffer.h>
12: #include <fiwix/mman.h>
13: #include <fiwix/filesystems.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/errno.h>
16: #include <fiwix/string.h>
17:
18: #ifdef __DEBUG__
19: #include <fiwix/stdio.h>
20: #include <fiwix/process.h>
21: #endif /* __DEBUG__ */
22:
23: int sys_mount(const char *source, const char *target, const char *fstype, unsigned long int flags, const void *data)
24: {
25:     struct inode *i_source, *i_target;
26:     struct mount *mt;
27:     struct filesystems *fs;
28:     struct vma *vma;
29:     char *tmp_source, *tmp_target, *tmp_fstype;
30:     __dev_t dev;
31:     int len, errno;
32:
33: #ifdef __DEBUG__
34:     printk("(pid %d) sys_mount(%s, %s, %s, 0x%08x, 0x%08x\n", current->pid,
source, target, (int)fstype ? fstype : "<NULL>", flags, data);
35: #endif /* __DEBUG__ */
36:
37:     if(!IS_SUPERUSER) {
38:         return -EPERM;
39:     }
40:
41:     if((errno = malloc_name(target, &tmp_target)) < 0) {
42:         return errno;
43:     }
44:     if((errno = namei(tmp_target, &i_target, NULL, FOLLOW_LINKS))) {
45:         free_name(tmp_target);
46:         return errno;
47:     }
48:     if(!S_ISDIR(i_target->i_mode)) {
49:         iput(i_target);
50:         free_name(tmp_target);
51:         return -ENOTDIR;
52:     }
53:     if((flags & MS_MGC_VAL) == MS_MGC_VAL) {
54:         flags &= ~MS_MGC_MSK;
55:     }
56:
57:     if(flags & MS_REMOUNT) {
58:         if(!(mt = get_mount_point(i_target))) {
59:             iput(i_target);
60:             free_name(tmp_target);
61:             return -EINVAL;
62:         }
63:         fs = mt->fs;
64:         if(fs->fsop && fs->fsop->remount_fs) {
65:             if((errno = fs->fsop->remount_fs(&mt->sb, flags))) {

```

## kernel/syscalls/mount.c

Page 2/4

```

66:             iput(i_target);
67:             free_name(tmp_target);
68:             return errno;
69:         }
70:     } else {
71:         iput(i_target);
72:         free_name(tmp_target);
73:         return -EINVAL;
74:     }
75:
76:     /* switching from RW to RO */
77:     if(flags & MS_RDONLY && !(mt->sb.flags & MS_RDONLY)) {
78:         dev = mt->dev;
79:         /*
80:          * FIXME: if there are files opened in RW mode then
81:          * we can't continue and must return -EBUSY.
82:          */
83:         if(fs->fsop && fs->fsop->release_superblock) {
84:             fs->fsop->release_superblock(&mt->sb);
85:         }
86:         sync_superblocks(dev);
87:         sync_inodes(dev);
88:         sync_buffers(dev);
89:     }
90:
91:     mt->sb.flags &= ~MS_RDONLY;
92:     mt->sb.flags |= (flags & MS_RDONLY);
93:     iput(i_target);
94:     free_name(tmp_target);
95:     return 0;
96: }
97:
98: if(i_target->mount_point) {
99:     iput(i_target);
100:    free_name(tmp_target);
101:    return -EBUSY;
102: }
103:
104: /* check the validity of fstype */
105: if(!(vma = find_vma_region((unsigned int)fstype))) {
106:     iput(i_target);
107:     free_name(tmp_target);
108:     return -EFAULT;
109: }
110: if(!(vma->prot & PROT_READ)) {
111:     iput(i_target);
112:     free_name(tmp_target);
113:     return -EFAULT;
114: }
115: len = MIN(vma->end - (unsigned int)fstype, PAGE_SIZE - 1);
116: if(!(tmp_fstype = (char *)kmalloc())) {
117:     iput(i_target);
118:     free_name(tmp_target);
119:     return -ENOMEM;
120: }
121: memcpy_b(tmp_fstype, fstype, len);
122:
123: if(!(fs = get_filesystem(fstype))) {
124:     iput(i_target);
125:     free_name(tmp_target);
126:     free_name(tmp_fstype);
127:     return -ENODEV;
128: }
129: dev = fs->fsop->fsdev;
130:
131: if((errno = malloc_name(source, &tmp_source)) < 0) {
132:     iput(i_target);

```



## kernel/syscalls/mount.c

Page 3/4

```

133:         free_name(tmp_target);
134:         free_name(tmp_fstype);
135:         return errno;
136:     }
137:     if(fs->fsop->flags == FSOP_REQUIRES_DEV) {
138:         if((errno = namei(tmp_source, &i_source, NULL, FOLLOW_LINKS))) {
139:             iput(i_target);
140:             free_name(tmp_target);
141:             free_name(tmp_fstype);
142:             free_name(tmp_source);
143:             return errno;
144:         }
145:         if(!S_ISBLK(i_source->i_mode)) {
146:             iput(i_target);
147:             iput(i_source);
148:             free_name(tmp_target);
149:             free_name(tmp_fstype);
150:             free_name(tmp_source);
151:             return -ENOTBLK;
152:         }
153:         if(i_source->fsop && i_source->fsop->open) {
154:             if((errno = i_source->fsop->open(i_source, NULL))) {
155:                 iput(i_target);
156:                 iput(i_source);
157:                 free_name(tmp_target);
158:                 free_name(tmp_fstype);
159:                 free_name(tmp_source);
160:                 return errno;
161:             }
162:         } else {
163:             iput(i_target);
164:             iput(i_source);
165:             free_name(tmp_target);
166:             free_name(tmp_fstype);
167:             free_name(tmp_source);
168:             return -EINVAL;
169:         }
170:         dev = i_source->rdev;
171:     }
172:
173:     if(! (mt = get_free_mount_point(dev))) {
174:         if(fs->fsop->flags == FSOP_REQUIRES_DEV) {
175:             i_source->fsop->close(i_source, NULL);
176:             iput(i_source);
177:         }
178:         iput(i_target);
179:         free_name(tmp_target);
180:         free_name(tmp_fstype);
181:         free_name(tmp_source);
182:         return -EBUSY;
183:     }
184:
185:     mt->sb.flags = flags;
186:     if(fs->fsop && fs->fsop->read_superblock) {
187:         if((errno = fs->fsop->read_superblock(dev, &mt->sb))) {
188:             i_source->fsop->close(i_source, NULL);
189:             if(fs->fsop->flags == FSOP_REQUIRES_DEV) {
190:                 iput(i_source);
191:             }
192:             iput(i_target);
193:             release_mount_point(mt);
194:             free_name(tmp_target);
195:             free_name(tmp_fstype);
196:             free_name(tmp_source);
197:             return errno;
198:         }
199:     } else {

```

**kernel/syscalls/mount.c**

Page 4/4

```
200:             if (fs->fsop->flags == FSOP_REQUIRES_DEV) {
201:                 iput (i_source);
202:             }
203:             iput (i_target);
204:             release_mount_point (mt);
205:             free_name (tmp_target);
206:             free_name (tmp_fstype);
207:             free_name (tmp_source);
208:             return -EINVAL;
209:         }
210:
211:         mt->dev = dev;
212:         strncpy (mt->devname, tmp_source, DEVNAME_MAX);
213:         strcpy (mt->dirname, tmp_target);
214:         mt->sb.dir = i_target;
215:         mt->fs = fs;
216:         i_target->mount_point = mt->sb.root;
217:         free_name (tmp_target);
218:         free_name (tmp_fstype);
219:         free_name (tmp_source);
220:         return 0;
221:     }
```

## kernel/syscalls/mprotect.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/mprotect.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/mman.h>
10: #include <fiwix/mm.h>
11: #include <fiwix/fcntl.h>
12: #include <fiwix/errno.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /*__DEBUG__ */
18:
19: int sys_mprotect(unsigned int addr, __size_t length, int prot)
20: {
21:     struct vma *vma;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_mprotect(0x%08x, %d, %d)\n", current->pid, addr, le
ngth, prot);
25: #endif /*__DEBUG__ */
26:
27:     if((addr & ~PAGE_MASK) || length < 0) {
28:         return -EINVAL;
29:     }
30:     if(prot > (PROT_READ | PROT_WRITE | PROT_EXEC)) {
31:         return -EINVAL;
32:     }
33:     if(!(vma = find_vma_region(addr))) {
34:         return -ENOMEM;
35:     }
36:     length = PAGE_ALIGN(length);
37:     if((addr + length) > vma->end) {
38:         return -ENOMEM;
39:     }
40:     if(vma->inode && (vma->flags & MAP_SHARED)) {
41:         if(prot & PROT_WRITE) {
42:             if(!(vma->o_mode & (O_WRONLY | O_RDWR))) {
43:                 return -EACCES;
44:             }
45:         }
46:     }
47:
48:     return do_mprotect(vma, addr, length, prot);
49: }

```

**kernel/syscalls/munmap.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/munmap.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/mman.h>
10:
11: #ifdef __DEBUG__
12: #include <fiwix/stdio.h>
13: #include <fiwix/process.h>
14: #endif /*__DEBUG__*/
15:
16: int sys_munmap(unsigned int addr, __size_t length)
17: {
18: #ifdef __DEBUG__
19:     printk("(pid %d) sys_munmap(0x%08x, %d)\n", current->pid, addr, length);
20: #endif /*__DEBUG__*/
21:     return do_munmap(addr, length);
22: }
```

## kernel/syscalls/nanosleep.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/nanosleep.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/time.h>
11: #include <fiwix/timer.h>
12: #include <fiwix/process.h>
13: #include <fiwix/sched.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/errno.h>
16:
17: #ifdef __DEBUG__
18: #include <fiwix/stdio.h>
19: #endif /*__DEBUG__ */
20:
21: int sys_nanosleep(const struct timespec *req, struct timespec *rem)
22: {
23:     int errno;
24:     long int nsec;
25:
26: #ifdef __DEBUG__
27:     printk("(pid %d) sys_nanosleep(0x%08x, 0x%08x)\n", current->pid, (unsigned
ed int) req, (unsigned int) rem);
28: #endif /*__DEBUG__ */
29:
30:     if((errno = check_user_area(VERIFY_READ, req, sizeof(struct timespec))))
{
31:         return errno;
32:     }
33:     if(req->tv_sec < 0 || req->tv_nsec >= 1000000000L || req->tv_nsec < 0) {
34:         return -EINVAL;
35:     }
36:
37:     /*
38:      * Since the current maximum precision of the kernel is only 10ms, we
39:      * need to convert any lower request to a minimum of 10ms, even knowing
40:      * that this might increase the sleep a bit more than the requested.
41:      */
42:     nsec = req->tv_nsec;
43:     if(nsec < 100000000L) {
44:         nsec *= 10;
45:     }
46:
47:     current->timeout = (req->tv_sec * HZ) + (nsec * HZ / 1000000000L);
48:     if(current->timeout) {
49:         sleep(&sys_nanosleep, PROC_INTERRUPTIBLE);
50:         if(current->timeout) {
51:             if(rem) {
52:                 if((errno = check_user_area(VERIFY_WRITE, rem, s
sizeof(struct timespec)))) {
53:                     return errno;
54:                 }
55:                 rem->tv_sec = current->timeout / HZ;
56:                 rem->tv_nsec = (current->timeout % HZ) * 100000
000L / HZ;
57:             }
58:             return -EINTR;
59:         }
60:     }
61:     return 0;
62: }

```

## kernel/syscalls/newfstat.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/newfstat.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/statbuf.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_newfstat(unsigned int ufd, struct new_stat *statbuf)
18: {
19:     struct inode *i;
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_newfstat(%d, 0x%08x) -> returning structure\n", cur
rent->pid, ufd, (unsigned int )statbuf);
24: #endif /*__DEBUG__*/
25:
26:     CHECK_UFD(ufd);
27:     if((errno = check_user_area(VERIFY_WRITE, statbuf, sizeof(struct new_sta
t)))) {
28:         return errno;
29:     }
30:     i = fd_table[current->fd[ufd]].inode;
31:     statbuf->st_dev = i->dev;
32:     statbuf->__pad1 = 0;
33:     statbuf->st_ino = i->inode;
34:     statbuf->st_mode = i->i_mode;
35:     statbuf->st_nlink = i->i_nlink;
36:     statbuf->st_uid = i->i_uid;
37:     statbuf->st_gid = i->i_gid;
38:     statbuf->st_rdev = i->rdev;
39:     statbuf->__pad2 = 0;
40:     statbuf->st_size = i->i_size;
41:     statbuf->st_blksize = i->sb->s_blocksize;
42:     statbuf->st_blocks = i->i_blocks;
43:     if(!i->i_blocks) {
44:         statbuf->st_blocks = (i->i_size / i->sb->s_blocksize * 2);
45:         statbuf->st_blocks++;
46:     }
47:     statbuf->st_atime = i->i_atime;
48:     statbuf->__unused1 = 0;
49:     statbuf->st_mtime = i->i_mtime;
50:     statbuf->__unused2 = 0;
51:     statbuf->st_ctime = i->i_ctime;
52:     statbuf->__unused3 = 0;
53:     statbuf->__unused4 = 0;
54:     statbuf->__unused5 = 0;
55:     return 0;
56: }

```

## kernel/syscalls/newlstat.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/newlstat.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/statbuf.h>
10: #include <fiwix/string.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_newlstat(const char *filename, struct new_stat *statbuf)
18: {
19:     struct inode *i;
20:     char *tmp_name;
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_newlstat('%s', 0x%08x) -> returning structure\n", c
urrent->pid, filename, (unsigned int)statbuf);
25: #endif /*__DEBUG__*/
26:
27:     if((errno = check_user_area(VERIFY_WRITE, statbuf, sizeof(struct new_sta
t)))) {
28:         return errno;
29:     }
30:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
31:         return errno;
32:     }
33:     if((errno = namei(tmp_name, &i, NULL, !FOLLOW_LINKS))) {
34:         free_name(tmp_name);
35:         return errno;
36:     }
37:     statbuf->st_dev = i->dev;
38:     statbuf->__pad1 = 0;
39:     statbuf->st_ino = i->inode;
40:     statbuf->st_mode = i->i_mode;
41:     statbuf->st_nlink = i->i_nlink;
42:     statbuf->st_uid = i->i_uid;
43:     statbuf->st_gid = i->i_gid;
44:     statbuf->st_rdev = i->rdev;
45:     statbuf->__pad2 = 0;
46:     statbuf->st_size = i->i_size;
47:     statbuf->st_blksize = i->sb->s_blocksize;
48:     statbuf->st_blocks = i->i_blocks;
49:     if(!i->i_blocks) {
50:         statbuf->st_blocks = (i->i_size / i->sb->s_blocksize) * 2;
51:         statbuf->st_blocks++;
52:     }
53:     statbuf->st_atime = i->i_atime;
54:     statbuf->__unused1 = 0;
55:     statbuf->st_mtime = i->i_mtime;
56:     statbuf->__unused2 = 0;
57:     statbuf->st_ctime = i->i_ctime;
58:     statbuf->__unused3 = 0;
59:     statbuf->__unused4 = 0;
60:     statbuf->__unused5 = 0;
61:     iput(i);
62:     free_name(tmp_name);
63:     return 0;
64: }

```

## kernel/syscalls/newstat.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/newstat.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/statbuf.h>
10: #include <fiwix/string.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_newstat(const char *filename, struct new_stat *statbuf)
18: {
19:     struct inode *i;
20:     char *tmp_name;
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_newstat('%s', 0x%08x) -> returning structure\n", cu
rrent->pid, filename, (unsigned int )statbuf);
25: #endif /*__DEBUG__*/
26:
27:     if((errno = check_user_area(VERIFY_WRITE, statbuf, sizeof(struct new_sta
t)))) {
28:         return errno;
29:     }
30:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
31:         return errno;
32:     }
33:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS))) {
34:         free_name(tmp_name);
35:         return errno;
36:     }
37:     statbuf->st_dev = i->dev;
38:     statbuf->__pad1 = 0;
39:     statbuf->st_ino = i->inode;
40:     statbuf->st_mode = i->i_mode;
41:     statbuf->st_nlink = i->i_nlink;
42:     statbuf->st_uid = i->i_uid;
43:     statbuf->st_gid = i->i_gid;
44:     statbuf->st_rdev = i->rdev;
45:     statbuf->__pad2 = 0;
46:     statbuf->st_size = i->i_size;
47:     statbuf->st_blksize = i->sb->s_blocksize;
48:     statbuf->st_blocks = i->i_blocks;
49:     if(!i->i_blocks) {
50:         statbuf->st_blocks = (i->i_size / i->sb->s_blocksize) * 2;
51:         statbuf->st_blocks++;
52:     }
53:     statbuf->st_atime = i->i_atime;
54:     statbuf->__unused1 = 0;
55:     statbuf->st_mtime = i->i_mtime;
56:     statbuf->__unused2 = 0;
57:     statbuf->st_ctime = i->i_ctime;
58:     statbuf->__unused3 = 0;
59:     statbuf->__unused4 = 0;
60:     statbuf->__unused5 = 0;
61:     iput(i);
62:     free_name(tmp_name);
63:     return 0;
64: }

```



## kernel/syscalls/newuname.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/newuname.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/utsname.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_newuname(struct new_utsname *uname)
19: {
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_newuname(0x%08x)\n", current->pid, (int)uname);
24: #endif /* __DEBUG__ */
25:
26:     if((errno = check_user_area(VERIFY_WRITE, uname, sizeof(struct new_utsna
me)))) {
27:         return errno;
28:     }
29:     if(!uname) {
30:         return -EFAULT;
31:     }
32:     memcpy_b(uname, &sys_utsname, sizeof(struct new_utsname));
33:     return 0;
34: }
```

## kernel/syscalls/old\_mmap.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/old_mmap.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/mman.h>
10: #include <fiwix/mm.h>
11: #include <fiwix/fcntl.h>
12: #include <fiwix/errno.h>
13: #include <fiwix/string.h>
14:
15: #ifdef __DEBUG__
16: #include <fiwix/stdio.h>
17: #include <fiwix/process.h>
18: #endif /*__DEBUG__*/
19:
20: int old_mmap(struct mmap *mmap)
21: {
22:     unsigned int page;
23:     struct inode *i;
24:     char flags;
25:     int errno;
26:
27: #ifdef __DEBUG__
28:     printk("(pid %d) old_mmap(0x%08x, %d, 0x%02x, 0x%02x, %d, 0x%08x) -> ",
current->pid, mmap->start, mmap->length, mmap->prot, mmap->flags, mmap->offse
t);
29: #endif /*__DEBUG__*/
30:
31:     if((errno = check_user_area(VERIFY_READ, mmap, sizeof(struct mmap)))) {
32:         return errno;
33:     }
34:     if(!mmap->length) {
35:         return -EINVAL;
36:     }
37:
38:     i = NULL;
39:     flags = 0;
40:     if(!(mmap->flags & MAP_ANONYMOUS)) {
41:         CHECK_UFD(mmap->fd);
42:         if(!(i = fd_table[current->fd[mmap->fd]].inode)) {
43:             return -EBADF;
44:         }
45:         flags = fd_table[current->fd[mmap->fd]].flags & O_ACCMODE;
46:     }
47:     page = do_mmap(i, mmap->start, mmap->length, mmap->prot, mmap->flags, mm
ap->offset, P_MMAP, flags);
48: #ifdef __DEBUG__
49:     printk("0x%08x\n", page);
50: #endif /*__DEBUG__*/
51:     return page;
52: }

```

## kernel/syscalls/old\_select.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/old_select.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/syscalls.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/sleep.h>
12: #include <fiwix/sched.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /*__DEBUG__ */
18:
19: int old_select(unsigned long int *params)
20: {
21:     int nfd;
22:     fd_set *readfds;
23:     fd_set *writefds;
24:     fd_set *exceptfds;
25:     struct timeval *timeout;
26:     int errno;
27:
28: #ifdef __DEBUG__
29:     printk("(pid %d) old_select(0x%08x)\n", current->pid, (int)params);
30: #endif /*__DEBUG__ */
31:
32:     if((errno = check_user_area(VERIFY_READ, (void *)params, sizeof(unsigned
int) * 5))) {
33:         return errno;
34:     }
35:     nfd = *(int *)params;
36:     readfds = *(fd_set **) (params + 1);
37:     writefds = *(fd_set **) (params + 2);
38:     exceptfds = *(fd_set **) (params + 3);
39:     timeout = *(struct timeval **) (params + 4);
40:
41:     return sys_select(nfd, readfds, writefds, exceptfds, timeout);
42: }
```

## kernel/syscalls/olduname.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/olduname.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/utsname.h>
10: #include <fiwix/string.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_olduname(struct oldold_utsname *uname)
18: {
19:     int errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_olduname(0x%0x)", current->pid, uname);
23: #endif /*__DEBUG__*/
24:
25:     if((errno = check_user_area(VERIFY_WRITE, uname, sizeof(struct oldold_utsname)))) {
26:         return errno;
27:     }
28:     memcpy_b(&uname->sysname, &sys_utsname.sysname, _OLD_UTSNAME_LENGTH);
29:     memset_b(&uname->sysname + _OLD_UTSNAME_LENGTH, NULL, 1);
30:     memcpy_b(&uname->nodename, &sys_utsname.nodename, _OLD_UTSNAME_LENGTH);
31:     memset_b(&uname->nodename + _OLD_UTSNAME_LENGTH, NULL, 1);
32:     memcpy_b(&uname->release, &sys_utsname.release, _OLD_UTSNAME_LENGTH);
33:     memset_b(&uname->release + _OLD_UTSNAME_LENGTH, NULL, 1);
34:     memcpy_b(&uname->version, &sys_utsname.version, _OLD_UTSNAME_LENGTH);
35:     memset_b(&uname->version + _OLD_UTSNAME_LENGTH, NULL, 1);
36:     memcpy_b(&uname->machine, &sys_utsname.machine, _OLD_UTSNAME_LENGTH);
37:     memset_b(&uname->machine + _OLD_UTSNAME_LENGTH, NULL, 1);
38:     return 0;
39: }
```

## kernel/syscalls/open.c

Page 1/3

```

1: /*
2:  * fiwix/kernel/syscalls/open.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/syscalls.h>
9: #include <fiwix/stat.h>
10: #include <fiwix/types.h>
11: #include <fiwix/fcntl.h>
12: #include <fiwix/errno.h>
13: #include <fiwix/stdio.h>
14: #include <fiwix/string.h>
15:
16: int sys_open(const char *filename, int flags, __mode_t mode)
17: {
18:     int fd, ufd;
19:     struct inode *i, *dir;
20:     char *tmp_name, *basename;
21:     int errno, follow_links, perms;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_open('%s', %o, %o)\n", current->pid, filename, flag
s, mode);
25: #endif /*__DEBUG__ */
26:
27:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
28:         return errno;
29:     }
30:
31:     basename = get_basename(tmp_name);
32:     follow_links = flags & O_NOFOLLOW ? !FOLLOW_LINKS : FOLLOW_LINKS;
33:     if((errno = namei(tmp_name, &i, &dir, follow_links))) {
34:         if(!dir) {
35:             free_name(tmp_name);
36:             if(flags & O_CREAT) {
37:                 return -ENOENT;
38:             }
39:             return errno;
40:         }
41:     }
42:
43: #ifdef __DEBUG__
44:     printk("\t(inode = %d)\n", i ? i->inode : -1);
45: #endif /*__DEBUG__ */
46:
47:     if(flags & O_CREAT) {
48:         if(!errno && (flags & O_EXCL)) {
49:             iput(i);
50:             iput(dir);
51:             free_name(tmp_name);
52:             return -EEXIST;
53:         }
54:         if(check_permission(TO_EXEC | TO_WRITE, dir) < 0) {
55:             iput(i);
56:             iput(dir);
57:             free_name(tmp_name);
58:             return -EACCES;
59:         }
60:         if(errno) { /* assumes -ENOENT */
61:             if(dir->fsop && dir->fsop->create) {
62:                 errno = dir->fsop->create(dir, basename, mode, &
i);
63:             }
64:             if(errno) {
65:                 iput(dir);
66:                 free_name(tmp_name);

```

## kernel/syscalls/open.c

Page 2/3

```

66:                                     return errno;
67:                                     }
68:                                 } else {
69:                                     iput(dir);
70:                                     free_name(tmp_name);
71:                                     return -EACCES;
72:                                 }
73:                             }
74:     } else {
75:         if(errno) {
76:             iput(dir);
77:             free_name(tmp_name);
78:             return errno;
79:         }
80:         if(S_ISDIR(i->i_mode) && (flags & (O_RDWR | O_WRONLY | O_TRUNC))
) {
81:             iput(i);
82:             iput(dir);
83:             free_name(tmp_name);
84:             return -EISDIR;
85:         }
86:         mode = 0;
87:     }
88:
89:     if((flags & O_ACCMODE) == O_RDONLY) {
90:         perms = TO_READ;
91:     } else if((flags & O_ACCMODE) == O_WRONLY) {
92:         perms = TO_WRITE;
93:     } else {
94:         perms = TO_READ | TO_WRITE;
95:     }
96:     if((errno = check_permission(perms, i))) {
97:         iput(i);
98:         iput(dir);
99:         free_name(tmp_name);
100:        return errno;
101:    }
102:    if((fd = get_new_fd(i)) < 0) {
103:        iput(i);
104:        iput(dir);
105:        free_name(tmp_name);
106:        return fd;
107:    }
108:    if((ufd = get_new_user_fd(0)) < 0) {
109:        release_fd(fd);
110:        iput(i);
111:        iput(dir);
112:        free_name(tmp_name);
113:        return ufd;
114:    }
115:
116: #ifdef __DEBUG__
117:     printk("\t(ufd = %d)\n", ufd);
118: #endif /*__DEBUG__*/
119:
120:     fd_table[fd].flags = flags;
121:     current->fd[ufd] = fd;
122:     if(i->fsop && i->fsop->open) {
123:         if((errno = i->fsop->open(i, &fd_table[fd])) < 0) {
124:             release_fd(fd);
125:             release_user_fd(ufd);
126:             iput(i);
127:             iput(dir);
128:             free_name(tmp_name);
129:             return errno;
130:         }
131:         iput(dir);

```

**kernel/syscalls/open.c**

Page 3/3

```
132:             free_name(tmp_name);
133:             return ufd;
134:         }
135:
136:         printk("WARNING: %s(): file '%s' (inode %d) without the open() method!\n
", __FUNCTION__, tmp_name, i->inode);
137:         release_fd(fd);
138:         release_user_fd(ufd);
139:         iput(i);
140:         iput(dir);
141:         free_name(tmp_name);
142:         return -EINVAL;
143: }
```

## kernel/syscalls/pause.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/pause.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/syscalls.h>
9: #include <fiwix/sched.h>
10: #include <fiwix/sleep.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_pause(void)
19: {
20: #ifdef __DEBUG__
21:     printk("(pid %d) sys_pause()\n", current->pid);
22: #endif /* __DEBUG__ */
23:
24:     for(;;) {
25:         if(sleep(&sys_pause, PROC_INTERRUPTIBLE)) {
26:             break;
27:         }
28:     }
29:     return -EINTR;
30: }
```



**kernel/syscalls/personality.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/personality.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifdef __DEBUG__
9: #include <fiwix/stdio.h>
10: #include <fiwix/process.h>
11: #endif /*__DEBUG__ */
12:
13: int sys_personality(unsigned long int persona)
14: {
15: #ifdef __DEBUG__
16:     printk("(pid %d) sys_personality(%d) -> %d\n", current->pid, persona, pe
rsona);
17: #endif /*__DEBUG__ */
18:     return persona;
19: }
```

## kernel/syscalls/pipe.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/pipe.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/filesystems.h>
10: #include <fiwix/fcntl.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/errno.h>
13: #include <fiwix/stdio.h>
14:
15: int sys_pipe(int pipefd[2])
16: {
17:     int rfd, rufd;
18:     int wfd, wufd;
19:     struct filesystems *fs;
20:     struct inode *i;
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_pipe()", current->pid);
25: #endif /*__DEBUG__ */
26:
27:     if(!(fs = get_filesystem("pipefs"))) {
28:         printk("WARNING: %s(): pipefs filesystem is not registered!\n",
__FUNCTION__);
29:         return -EINVAL;
30:     }
31:     if((errno = check_user_area(VERIFY_WRITE, pipefd, sizeof(int) * 2)) {
32:         return errno;
33:     }
34:     if(!(i = ialloc(&fs->mt->sb, S_IFIFO))) {
35:         return -EINVAL;
36:     }
37:     if((rfd = get_new_fd(i)) < 0) {
38:         iput(i);
39:         return -ENFILE;
40:     }
41:     if((wfd = get_new_fd(i)) < 0) {
42:         release_fd(rfd);
43:         iput(i);
44:         return -ENFILE;
45:     }
46:     if((rufd = get_new_user_fd(0)) < 0) {
47:         release_fd(rfd);
48:         release_fd(wfd);
49:         iput(i);
50:         return -EMFILE;
51:     }
52:     if((wufd = get_new_user_fd(0)) < 0) {
53:         release_fd(rfd);
54:         release_fd(wfd);
55:         release_user_fd(rufd);
56:         iput(i);
57:         return -EMFILE;
58:     }
59:
60:     pipefd[0] = rufd;
61:     pipefd[1] = wufd;
62:     current->fd[rufd] = rfd;
63:     current->fd[wufd] = wfd;
64:     fd_table[rfd].flags = O_RDONLY;
65:     fd_table[wfd].flags = O_WRONLY;
66:

```

**kernel/syscalls/pipe.c**

Page 2/2

```
67: #ifdef __DEBUG__
68:     printk(" -> inode=%d, rufd=%d wufd=%d (rfd=%d wfd=%d)\n", i->inode, rufd
, wufd, rfd, wfd);
69: #endif /*__DEBUG__ */
70:
71:     return 0;
72: }
```

## kernel/syscalls/read.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/read.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/fcntl.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__ */
16:
17: int sys_read(unsigned int ufd, char *buf, int count)
18: {
19:     struct inode *i;
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_read(%d, 0x%08x, %d) -> ", current->pid, ufd, buf,
count);
24: #endif /*__DEBUG__ */
25:
26:     CHECK_UFD(ufd);
27:     if((errno = check_user_area(VERIFY_WRITE, buf, count))) {
28:         return errno;
29:     }
30:     if(fd_table[current->fd[ufd]].flags & O_WRONLY) {
31:         return -EBADF;
32:     }
33:     if(!count) {
34:         return 0;
35:     }
36:     if(count < 0) {
37:         return -EINVAL;
38:     }
39:
40:     i = fd_table[current->fd[ufd]].inode;
41:     if(i->fsop && i->fsop->read) {
42:         errno = i->fsop->read(i, &fd_table[current->fd[ufd]], buf, count
);
43: #ifdef __DEBUG__
44:         printk("%d\n", errno);
45: #endif /*__DEBUG__ */
46:         return errno;
47:     }
48:     return -EINVAL;
49: }

```

## kernel/syscalls/readlink.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/readlink.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/stat.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /* __DEBUG__ */
18:
19: int sys_readlink(const char *filename, char *buffer, __size_t bufsize)
20: {
21:     struct inode *i;
22:     char *tmp_name;
23:     int errno;
24:
25: #ifdef __DEBUG__
26:     printk("(pid %d) sys_readlink(%s, 0x%08x, %d)\n", current->pid, filename
, (unsigned int)buffer, bufsize);
27: #endif /* __DEBUG__ */
28:
29:     if(bufsize <= 0) {
30:         return -EINVAL;
31:     }
32:     if((errno = check_user_area(VERIFY_WRITE, buffer, bufsize)) {
33:         return errno;
34:     }
35:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
36:         return errno;
37:     }
38:     if((errno = namei(tmp_name, &i, NULL, !FOLLOW_LINKS)) {
39:         free_name(tmp_name);
40:         return errno;
41:     }
42:     if(!S_ISLNK(i->i_mode)) {
43:         iput(i);
44:         free_name(tmp_name);
45:         return -EINVAL;
46:     }
47:
48:     if(i->fsop && i->fsop->readlink) {
49:         errno = i->fsop->readlink(i, buffer, bufsize);
50:         iput(i);
51:         free_name(tmp_name);
52:         return errno;
53:     }
54:     iput(i);
55:     free_name(tmp_name);
56:     return -EINVAL;
57: }

```

## kernel/syscalls/reboot.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/reboot.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/syscalls.h>
10: #include <fiwix/reboot.h>
11: #include <fiwix/signal.h>
12: #include <fiwix/process.h>
13: #include <fiwix/errno.h>
14:
15: #ifdef __DEBUG__
16: #include <fiwix/stdio.h>
17: #endif /*__DEBUG__*/
18:
19: int sys_reboot(int magic1, int magic2, int flag)
20: {
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_reboot(0x%08x, %d, 0x%08x)\n", current->pid, magic1
, magic2, flag);
23: #endif /*__DEBUG__*/
24:
25:     if(!IS_SUPERUSER) {
26:         return -EPERM;
27:     }
28:     if((magic1 != BMAGIC_1) || (magic2 != BMAGIC_2)) {
29:         return -EINVAL;
30:     }
31:     switch(flag) {
32:         case BMAGIC_SOFT:
33:             ctrl_alt_del = 0;
34:             break;
35:         case BMAGIC_HARD:
36:             ctrl_alt_del = 1;
37:             break;
38:         case BMAGIC_REBOOT:
39:             reboot();
40:             break;
41:         case BMAGIC_HALT:
42:             sys_kill(-1, SIGKILL);
43:             stop_kernel();
44:             break;
45:         default:
46:             return -EINVAL;
47:     }
48:     return 0;
49: }
```

## kernel/syscalls/rename.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/rename.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/stat.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_rename(const char *oldpath, const char *newpath)
19: {
20:     struct inode *i, *dir, *i_new, *dir_new;
21:     char *tmp_oldpath, *tmp_newpath;
22:     char *oldbasename, *newbasename;
23:     int errno;
24:
25: #ifdef __DEBUG__
26:     printk("(pid %d) sys_rename('%s', '%s')\n", current->pid, oldpath, newpa
th);
27: #endif /* __DEBUG__ */
28:
29:     if((errno = malloc_name(oldpath, &tmp_oldpath)) < 0) {
30:         return errno;
31:     }
32:     if((errno = namei(tmp_oldpath, &i, &dir, !FOLLOW_LINKS)) {
33:         if(dir) {
34:             iput(dir);
35:         }
36:         free_name(tmp_oldpath);
37:         return errno;
38:     }
39:     if(IS_RDONLY_FS(i)) {
40:         iput(i);
41:         iput(dir);
42:         free_name(tmp_oldpath);
43:         return -EROFS;
44:     }
45:
46:     if((errno = malloc_name(newpath, &tmp_newpath)) < 0) {
47:         iput(i);
48:         iput(dir);
49:         free_name(tmp_oldpath);
50:         return errno;
51:     }
52:     newbasename = remove_trailing_slash(tmp_newpath);
53:     if((errno = namei(newbasename, &i_new, &dir_new, !FOLLOW_LINKS)) {
54:         if(!dir_new) {
55:             iput(i);
56:             iput(dir);
57:             free_name(tmp_oldpath);
58:             free_name(tmp_newpath);
59:             return errno;
60:         }
61:     }
62:     if(dir->dev != dir_new->dev) {
63:         errno = -EXDEV;
64:         goto end;
65:     }
66:     newbasename = get_basename(newbasename);

```

## kernel/syscalls/rename.c

Page 2/2

```

67:         if((newbasename[0] == '.' && newbasename[1] == NULL) || (newbasename[0]
== '.' && newbasename[1] == '.' && newbasename[2] == NULL)) {
68:             errno = -EINVAL;
69:             goto end;
70:         }
71:
72:         oldbasename = get_basename(tmp_oldpath);
73:         oldbasename = remove_trailing_slash(oldbasename);
74:         if((oldbasename[0] == '.' && oldbasename[1] == NULL) || (oldbasename[0]
== '.' && oldbasename[1] == '.' && oldbasename[2] == NULL)) {
75:             errno = -EINVAL;
76:             goto end;
77:         }
78:
79:         if(i_new) {
80:             if(S_ISREG(i->i_mode)) {
81:                 if(S_ISDIR(i_new->i_mode)) {
82:                     errno = -EISDIR;
83:                     goto end;
84:                 }
85:             }
86:             if(S_ISDIR(i->i_mode)) {
87:                 if(!S_ISDIR(i_new->i_mode)) {
88:                     errno = -ENOTDIR;
89:                     goto end;
90:                 }
91:             }
92:             if(i->inode == i_new->inode) {
93:                 errno = 0;
94:                 goto end;
95:             }
96:         }
97:
98:         if(check_permission(TO_EXEC | TO_WRITE, dir_new) < 0) {
99:             errno = -EACCES;
100:            goto end;
101:        }
102:
103:        if(dir_new->fsop && dir_new->fsop->rename) {
104:            errno = dir_new->fsop->rename(i, dir, i_new, dir_new, oldbasenam
e, newbasename);
105:        } else {
106:            errno = -EPERM;
107:        }
108:
109:    end:
110:        iput(i);
111:        iput(dir);
112:        iput(i_new);
113:        iput(dir_new);
114:        free_name(tmp_oldpath);
115:        free_name(tmp_newpath);
116:        return errno;
117:    }

```



## kernel/syscalls/rmdir.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/rmdir.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/stat.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__ */
16:
17: int sys_rmdir(const char *dirname)
18: {
19:     struct inode *i, *dir;
20:     char *tmp_dirname;
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_rmdir(%s)\n", current->pid, dirname);
25: #endif /*__DEBUG__ */
26:
27:     if((errno = malloc_name(dirname, &tmp_dirname)) < 0) {
28:         return errno;
29:     }
30:     if((errno = namei(tmp_dirname, &i, &dir, !FOLLOW_LINKS))) {
31:         if(dir) {
32:             iput(dir);
33:         }
34:         free_name(tmp_dirname);
35:         return errno;
36:     }
37:     if(!S_ISDIR(i->i_mode)) {
38:         iput(i);
39:         iput(dir);
40:         free_name(tmp_dirname);
41:         return -ENOTDIR;
42:     }
43:     if(i == current->root || i->mount_point || i->count > 1) {
44:         iput(i);
45:         iput(dir);
46:         free_name(tmp_dirname);
47:         return -EBUSY;
48:     }
49:     if(IS_RDONLY_FS(i)) {
50:         iput(i);
51:         iput(dir);
52:         free_name(tmp_dirname);
53:         return -EROFS;
54:     }
55:     if(i == dir) {
56:         iput(i);
57:         iput(dir);
58:         free_name(tmp_dirname);
59:         return -EPERM;
60:     }
61:     if(check_permission(TO_EXEC | TO_WRITE, dir) < 0) {
62:         iput(i);
63:         iput(dir);
64:         free_name(tmp_dirname);
65:         return -EACCES;
66:     }
67:

```

**kernel/syscalls/rmdir.c**

Page 2/2

```
68:      /* check sticky permission bit */
69:      if(dir->i_mode & S_ISVTX) {
70:          if(check_user_permission(i) {
71:              iput(i);
72:              iput(dir);
73:              free_name(tmp_dirname);
74:              return -EPERM;
75:          }
76:      }
77:
78:      if(i->fsop && i->fsop->rmdir) {
79:          errno = i->fsop->rmdir(dir, i);
80:      } else {
81:          errno = -EPERM;
82:      }
83:      iput(i);
84:      iput(dir);
85:      free_name(tmp_dirname);
86:      return errno;
87: }
```

## kernel/syscalls/select.c

Page 1/3

```

1: /*
2:  * fiwix/kernel/syscalls/select.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/process.h>
11: #include <fiwix/timer.h>
12: #include <fiwix/sched.h>
13: #include <fiwix/sleep.h>
14: #include <fiwix/errno.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: static int check_fds(int nfd, fd_set *rfd, fd_set *wfd, fd_set *efd)
19: {
20:     int n, bit;
21:     unsigned long int set;
22:
23:     n = bit = 0;
24:     while(bit < nfd) {
25:         bit = n * __NFDBITS;
26:         set = rfd->fds_bits[n] | wfd->fds_bits[n] | efd->fds_bits[n];
27:         while(set) {
28:             if(__FD_ISSET(bit, rfd)) {
29:                 CHECK_UFD(bit);
30:             }
31:             set >>= 1;
32:             bit++;
33:         }
34:         n++;
35:     }
36:
37:     return 0;
38: }
39:
40: static int do_check(struct inode *i, int flag)
41: {
42:     if(i->fsop && i->fsop->select) {
43:         if(i->fsop->select(i, flag)) {
44:             return 1;
45:         }
46:     }
47:
48:     return 0;
49: }
50:
51: int do_select(int nfd, fd_set *rfd, fd_set *wfd, fd_set *efd, fd_set *res_rfd, fd_set *res_wfd, fd_set *res_efd)
52: {
53:     int n, count;
54:     struct inode *i;
55:
56:     count = 0;
57:     for(;;) {
58:         for(n = 0; n < nfd; n++) {
59:             if(!current->fd[n]) {
60:                 continue;
61:             }
62:             i = fd_table[current->fd[n]].inode;
63:             if(__FD_ISSET(n, rfd)) {
64:                 if(do_check(i, SEL_R)) {
65:                     __FD_SET(n, res_rfd);
66:                     count++;

```

## kernel/syscalls/select.c

Page 2/3

```

67:         }
68:     }
69:     if(__FD_ISSET(n, wfds)) {
70:         if(do_check(i, SEL_W)) {
71:             __FD_SET(n, res_wfds);
72:             count++;
73:         }
74:     }
75:     if(__FD_ISSET(n, efds)) {
76:         if(do_check(i, SEL_E)) {
77:             __FD_SET(n, res_efds);
78:             count++;
79:         }
80:     }
81: }
82:
83:     if(count || !current->timeout || current->sigpending & ~current->
>sigblocked) {
84:         break;
85:     }
86:     sleep(&do_select, PROC_INTERRUPTIBLE);
87: }
88:
89:     return count;
90: }
91:
92: int sys_select(int nfd, fd_set *readfds, fd_set *writefds, fd_set *exceptfds, s
truct timeval *timeout)
93: {
94:     unsigned long int t;
95:     fd_set rfds, wfds, efds;
96:     fd_set res_rfds, res_wfds, res_efds;
97:     int errno;
98:
99: #ifdef __DEBUG__
100:     printk("(pid %d) sys_select(%d, 0x%08x, 0x%08x, 0x%08x, 0x%08x [%d])\n",
current->pid, nfd, (int)readfds, (int)writefds, (int)exceptfds, (int)timeout, (int)ti
meout ? tv2ticks(timeout): 0);
101: #endif /* __DEBUG__ */
102:
103:     if(nfd < 0) {
104:         return -EINVAL;
105:     }
106:     if(nfd > MIN(__FD_SETSIZE, NR_OPENS)) {
107:         nfd = MIN(__FD_SETSIZE, NR_OPENS);
108:     }
109:
110:     if(readfds) {
111:         if((errno = check_user_area(VERIFY_WRITE, readfds, sizeof(fd_set
))) {
112:             return errno;
113:         }
114:         memcpy_b(&rfds, readfds, sizeof(fd_set));
115:     } else {
116:         __FD_ZERO(&rfds);
117:     }
118:     if(writefds) {
119:         if((errno = check_user_area(VERIFY_WRITE, writefds, sizeof(fd_se
t)))) {
120:             return errno;
121:         }
122:         memcpy_b(&wfds, writefds, sizeof(fd_set));
123:     } else {
124:         __FD_ZERO(&wfds);
125:     }
126:     if(exceptfds) {
127:         if((errno = check_user_area(VERIFY_WRITE, exceptfds, sizeof(fd_s

```

```
et))) {
128:         return errno;
129:     }
130:     memcpy_b(&efds, exceptfds, sizeof(fd_set));
131: } else {
132:     __FD_ZERO(&efds);
133: }
134:
135: /* check the validity of all fds */
136: if((errno = check_fds(nfds, &rfd, &wfd, &efds)) < 0) {
137:     return errno;
138: }
139:
140: if(timeout) {
141:     t = tv2ticks(timeout);
142: } else {
143:     t = INFINITE_WAIT;
144: }
145:
146: __FD_ZERO(&res_rfd);
147: __FD_ZERO(&res_wfd);
148: __FD_ZERO(&res_efds);
149:
150: current->timeout = t;
151: if((errno = do_select(nfds, &rfd, &wfd, &efds, &res_rfd, &res_wfd, &
res_efds)) < 0) {
152:     return errno;
153: }
154: current->timeout = 0;
155:
156: if(readfds) {
157:     memcpy_b(readfds, &res_rfd, sizeof(fd_set));
158: }
159: if(writefds) {
160:     memcpy_b(writefds, &res_wfd, sizeof(fd_set));
161: }
162: if(exceptfds) {
163:     memcpy_b(exceptfds, &res_efds, sizeof(fd_set));
164: }
165: return errno;
166: }
```

## kernel/syscalls/setdomainname.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/setdomainname.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/utsname.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_setdomainname(const char *name, int length)
19: {
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_setdomainname('%s', %d)\n", current->pid, name, len
gth);
24: #endif /* __DEBUG__ */
25:
26:     if((errno = check_user_area(VERIFY_READ, name, length)) {
27:         return errno;
28:     }
29:     if(!IS_SUPERUSER) {
30:         return -EPERM;
31:     }
32:     if(length < 0 || length > _UTSNAME_LENGTH) {
33:         return -EINVAL;
34:     }
35:     memcpy_b(&sys_utsname.domainname, name, length);
36:     sys_utsname.domainname[length] = NULL;
37:     return 0;
38: }

```

**kernel/syscalls/setfsgid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/setfsgid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #include <fiwix/process.h>
13: #endif /*__DEBUG__ */
14:
15: int sys_setfsgid(__gid_t fsgid)
16: {
17: #ifdef __DEBUG__
18:     printk("(pid %d) sys_setfsgid(%d) -> %d\n", current->pid, fsgid);
19: #endif /*__DEBUG__ */
20:     return 0;
21: }
```

**kernel/syscalls/setfsuid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/setfsuid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #include <fiwix/process.h>
13: #endif /*__DEBUG__ */
14:
15: int sys_setfsuid(__uid_t fsuid)
16: {
17: #ifdef __DEBUG__
18:     printk("(pid %d) sys_setfsuid(%d) -> %d\n", current->pid, fsuid);
19: #endif /*__DEBUG__ */
20:     return 0;
21: }
```



**kernel/syscalls/setgid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/setgid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/process.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #endif /*__DEBUG__*/
15:
16: int sys_setgid(__gid_t gid)
17: {
18: #ifdef __DEBUG__
19:     printk("(pid %d) sys_setgid(%d)\n", current->pid, gid);
20: #endif /*__DEBUG__*/
21:
22:     if(IS_SUPERUSER) {
23:         current->gid = current->egid = current->sgid = gid;
24:     } else {
25:         if((current->gid == gid) || (current->sgid == gid)) {
26:             current->egid = gid;
27:         } else {
28:             return -EPERM;
29:         }
30:     }
31:     return 0;
32: }
```

## kernel/syscalls/setgroups.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/setgroups.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_setgroups(__ssize_t size, const __gid_t *list)
18: {
19:     int n, errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_setgroups(%d, 0x%08x)\n", current->pid, size, (unsi
gned int)list);
23: #endif /*__DEBUG__*/
24:
25:     if((errno = check_user_area(VERIFY_READ, list, sizeof(__gid_t)))) {
26:         return errno;
27:     }
28:     if(!IS_SUPERUSER) {
29:         return -EPERM;
30:     }
31:     if(size > NGROUPS_MAX) {
32:         return -EINVAL;
33:     }
34:     for(n = 0; n < NGROUPS_MAX; n++) {
35:         current->groups[n] = -1;
36:         if(n < size) {
37:             current->groups[n] = list[n];
38:         }
39:     }
40:     return 0;
41: }
```

## kernel/syscalls/sethostname.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/sethostname.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/utsname.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_sethostname(const char *name, int length)
19: {
20:     int errno;
21:     char *tmp_name;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_sethostname('%s', %d)\n", current->pid, name, length);
25: #endif /* __DEBUG__ */
26:
27:     if((errno = malloc_name(name, &tmp_name)) < 0) {
28:         return errno;
29:     }
30:     if(!IS_SUPERUSER) {
31:         free_name(tmp_name);
32:         return -EPERM;
33:     }
34:     if(length < 0 || length > _UTSNAME_LENGTH) {
35:         free_name(tmp_name);
36:         return -EINVAL;
37:     }
38:     memcpy_b(&sys_utsname.nodename, tmp_name, length);
39:     sys_utsname.nodename[length] = NULL;
40:     free_name(tmp_name);
41:     return 0;
42: }
```

## kernel/syscalls/setitimer.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/setitimer.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/time.h>
10:
11: #ifdef __DEBUG__
12: #include <fiwix/stdio.h>
13: #include <fiwix/process.h>
14: #endif /*__DEBUG__ */
15:
16: int sys_setitimer(int which, const struct itimerval *new_value, struct itimerval
*old_value)
17: {
18:     int errno;
19:
20: #ifdef __DEBUG__
21:     printk("(pid %d) sys_setitimer(%d, 0x%08x, 0x%08x) -> \n", current->pid,
which, (unsigned int)new_value, (unsigned int)old_value);
22: #endif /*__DEBUG__ */
23:
24:     if((unsigned int)old_value) {
25:         if((errno = check_user_area(VERIFY_WRITE, old_value, sizeof(stru
ct itimerval)))) {
26:             return errno;
27:         }
28:     }
29:
30:     return setitimer(which, new_value, old_value);
31: }
```

## kernel/syscalls/setpgid.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/setpgid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/process.h>
10: #include <fiwix/sched.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_setpgid(__pid_t pid, __pid_t pgid)
18: {
19:     struct proc *p;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_setpgid(%d, %d)", current->pid, pid, pgid);
23: #endif /*__DEBUG__*/
24:
25:     if(pid < 0 || pgid < 0) {
26:         return -EINVAL;
27:     }
28:     if(!pid) {
29:         pid = current->pid;
30:     }
31:
32:     p = get_proc_by_pid(pid);
33:     if(!pgid) {
34:         pgid = p->pid;
35:     }
36:
37:     if(p != current && (p->state == PROC_UNUSED || p->ppid != current->pid))
{
38:         return -ESRCH;
39:     }
40:     if(p->sid == p->pid || p->sid != current->sid) {
41:         return -EPERM;
42:     }
43:
44:     {
45:         struct proc *p;
46:
47:         FOR_EACH_PROCESS(p) {
48:             if(p->state != PROC_UNUSED) {
49:                 if(p->pgid == pgid && p->sid != current->sid) {
50:                     return -EPERM;
51:                 }
52:             }
53:         }
54:     }
55:
56:     if(p != current && p->flags & PF_PEXEC) {
57:         return -EACCES;
58:     }
59:
60:     p->pgid = pgid;
61:
62: #ifdef __DEBUG__
63:     printk(" -> 0\n");
64: #endif /*__DEBUG__*/
65:
66:     return 0;

```

```
67: }
```

## kernel/syscalls/setregid.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/setregid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/process.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #endif /*__DEBUG__*/
15:
16: int sys_setregid(__gid_t gid, __gid_t egid)
17: {
18: #ifdef __DEBUG__
19:     printk("(pid %d) sys_setregid(%d, %d) -> ", current->pid, gid, egid);
20: #endif /*__DEBUG__*/
21:
22:     if(IS_SUPERUSER) {
23:         if(egid != (__uid_t)-1) {
24:             if(gid != (__uid_t)-1 || (current->egid >= 0 && current-
>gid != egid)) {
25:                 current->sgid = egid;
26:             }
27:             current->egid = egid;
28:         }
29:         if(gid != (__uid_t)-1) {
30:             current->gid = gid;
31:         }
32:     } else {
33:         if(egid != (__uid_t)-1 && (current->gid == egid || current->egid
== egid || current->sgid == egid)) {
34:             if(gid != (__uid_t)-1 || (current->egid >= 0 && current-
>gid != egid)) {
35:                 current->sgid = egid;
36:             }
37:             current->egid = egid;
38:         } else {
39:             return -EPERM;
40:         }
41:         if(gid != (__uid_t)-1 && (current->gid == gid || current->egid =
= gid)) {
42:             current->gid = gid;
43:         } else {
44:             return -EPERM;
45:         }
46:     }
47:
48:     return 0;
49: }

```

## kernel/syscalls/setreuid.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/setreuid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/process.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #endif /*__DEBUG__ */
15:
16: int sys_setreuid(__uid_t uid, __uid_t euid)
17: {
18: #ifdef __DEBUG__
19:     printk("(pid %d) sys_setreuid(%d, %d) -> ", current->pid, uid, euid);
20: #endif /*__DEBUG__ */
21:
22:     if(IS_SUPERUSER) {
23:         if(euid != (__uid_t)-1) {
24:             if(uid != (__uid_t)-1 || (current->euid >= 0 && current-
>uid != euid)) {
25:                 current->suid = euid;
26:             }
27:             current->euid = euid;
28:         }
29:         if(uid != (__uid_t)-1) {
30:             current->uid = uid;
31:         }
32:     } else {
33:         if(euid != (__uid_t)-1 && (current->uid == euid || current->euid
== euid || current->suid == euid)) {
34:             if(uid != (__uid_t)-1 || (current->euid >= 0 && current-
>uid != euid)) {
35:                 current->suid = euid;
36:             }
37:             current->euid = euid;
38:         } else {
39:             return -EPERM;
40:         }
41:         if(uid != (__uid_t)-1 && (current->uid == uid || current->euid =
= uid)) {
42:             current->uid = uid;
43:         } else {
44:             return -EPERM;
45:         }
46:     }
47:
48: #ifdef __DEBUG__
49:     printk(" 0\n");
50: #endif /*__DEBUG__ */
51:
52:     return 0;
53: }

```



## kernel/syscalls/setrlimit.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/setrlimit.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/resource.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_setrlimit(int resource, const struct rlimit *rlim)
19: {
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_setrlimit(%d, 0x%08x)\n", current->pid, resource, (
unsigned int)rlim);
24: #endif /* __DEBUG__ */
25:
26:     if((errno = check_user_area(VERIFY_READ, rlim, sizeof(struct rlimit))))
{
27:         return errno;
28:     }
29:     if(resource < 0 || resource >= RLIM_NLIMITS) {
30:         return -EINVAL;
31:     }
32:     if(rlim->rlim_cur > rlim->rlim_max) {
33:         return -EINVAL;
34:     }
35:     if(!IS_SUPERUSER) {
36:         if(rlim->rlim_max > current->rlim[resource].rlim_max) {
37:             return -EPERM;
38:         }
39:     }
40:     memcpy_b(&current->rlim[resource], rlim, sizeof(struct rlimit));
41:     return 0;
42: }

```

## kernel/syscalls/setsid.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/setsid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/string.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #endif /*__DEBUG__*/
15:
16: int sys_setsid(void)
17: {
18:     struct proc *p;
19:
20: #ifdef __DEBUG__
21:     printk("(pid %d) sys_setsid()\n", current->pid);
22: #endif /*__DEBUG__*/
23:
24:     if(PG_LEADER(current)) {
25:         return -EPERM;
26:     }
27:     FOR_EACH_PROCESS(p) { /* POSIX ANSI/IEEE Std 1003.1-1996 4.3.2 */
28:         if(p != current && p->pgid == current->pid) {
29:             return -EPERM;
30:         }
31:     }
32:
33:     current->sid = current->pgid = current->pid;
34:     current->ctty = NULL;
35:     return current->sid;
36: }
```

## kernel/syscalls/settimeofday.c

Page 1/1

```

1:  /*
2:  *  fiwix/kernel/syscalls/settimeofday.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/kernel.h>
9:  #include <fiwix/fs.h>
10: #include <fiwix/time.h>
11: #include <fiwix/timer.h>
12: #include <fiwix/process.h>
13: #include <fiwix/errno.h>
14:
15: #ifdef __DEBUG__
16: #include <fiwix/stdio.h>
17: #endif /*__DEBUG__ */
18:
19: int sys_settimeofday(const struct timeval *tv, const struct timezone *tz)
20: {
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_settimeofday()\n", current->pid);
25: #endif /*__DEBUG__ */
26:
27:     if(!IS_SUPERUSER) {
28:         return -EPERM;
29:     }
30:
31:     if(tv) {
32:         if((errno = check_user_area(VERIFY_READ, tv, sizeof(struct timev
al)))) {
33:             return errno;
34:         }
35:         CURRENT_TIME = tv->tv_sec;
36:         set_system_time(CURRENT_TIME);
37:     }
38:     if(tz) {
39:         if((errno = check_user_area(VERIFY_READ, tz, sizeof(struct timez
one)))) {
40:             return errno;
41:         }
42:         kstat.tz_minuteswest = tz->tz_minuteswest;
43:         kstat.tz_dsttime = tz->tz_dsttime;
44:     }
45:     return 0;
46: }

```

**kernel/syscalls/setuid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/setuid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/process.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #endif /*__DEBUG__*/
15:
16: int sys_setuid(__uid_t uid)
17: {
18: #ifdef __DEBUG__
19:     printk("(pid %d) sys_setuid(%d)\n", current->pid, uid);
20: #endif /*__DEBUG__*/
21:
22:     if(IS_SUPERUSER) {
23:         current->uid = current->suid = uid;
24:     } else {
25:         if((current->uid != uid) && (current->suid != uid)) {
26:             return -EPERM;
27:         }
28:     }
29:     current->euid = uid;
30:     return 0;
31: }
```

**kernel/syscalls/sgetmask.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/sgetmask.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #endif /*__DEBUG__*/
13:
14: int sys_sgetmask(void)
15: {
16: #ifdef __DEBUG__
17:     printk("(pid %d) sys_sgetmask() -> \n", current->pid);
18: #endif /*__DEBUG__*/
19:     return current->sigblocked;
20: }
```

## kernel/syscalls/sigation.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/sigation.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/signal.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__ */
16:
17: int sys_sigation(__sigset_t signum, const struct sigaction *newaction, struct s
igation *oldaction)
18: {
19:     int errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_sigation(%d, 0x%08x, 0x%08x)\n", current->pid, sig
num, (unsigned int)newaction, (unsigned int)oldaction);
23: #endif /*__DEBUG__ */
24:
25:     if(signum < 1 || signum > NSIG) {
26:         return -EINVAL;
27:     }
28:     if(signum == SIGKILL || signum == SIGSTOP) {
29:         return -EINVAL;
30:     }
31:     if(oldaction) {
32:         if((errno = check_user_area(VERIFY_WRITE, oldaction, sizeof(stru
ct sigaction)))) {
33:             return errno;
34:         }
35:         *oldaction = current->sigaction[signum - 1];
36:     }
37:     if(newaction) {
38:         if((errno = check_user_area(VERIFY_READ, newaction, sizeof(struc
t sigaction)))) {
39:             return errno;
40:         }
41:         current->sigaction[signum - 1] = *newaction;
42:         if(current->sigaction[signum - 1].sa_handler == SIG_IGN) {
43:             if(signum != SIGCHLD) {
44:                 current->sigpending &= SIG_MASK(signum);
45:             }
46:         }
47:         if(current->sigaction[signum - 1].sa_handler == SIG_DFL) {
48:             if(signum != SIGCHLD) {
49:                 current->sigpending &= SIG_MASK(signum);
50:             }
51:         }
52:     }
53:     return 0;
54: }

```

## kernel/syscalls/signal.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/signal.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/syscalls.h>
9: #include <fiwix/signal.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #endif /* __DEBUG__ */
17:
18: unsigned int sys_signal(__sigset_t signum, void(* sighandler)(int))
19: {
20:     struct sigaction s;
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_signal()\n", current->pid);
25: #endif /* __DEBUG__ */
26:
27:     if(signum < 1 || signum > NSIG) {
28:         return -EINVAL;
29:     }
30:     if(signum == SIGKILL || signum == SIGSTOP) {
31:         return -EINVAL;
32:     }
33:     if(sighandler != SIG_DFL && sighandler != SIG_IGN) {
34:         if((errno = check_user_area(VERIFY_READ, sighandler, sizeof(void
)))) {
35:             return errno;
36:         }
37:     }
38:
39:     memset_b(&s, 0, sizeof(struct sigaction));
40:     s.sa_handler = sighandler;
41:     s.sa_mask = 0;
42:     s.sa_flags = SA_RESETHAND;
43:     sighandler = current->sigaction[signum - 1].sa_handler;
44:     current->sigaction[signum - 1] = s;
45:     if(current->sigaction[signum - 1].sa_handler == SIG_IGN) {
46:         if(signum != SIGCHLD) {
47:             current->sigpending &= SIG_MASK(signum);
48:         }
49:     }
50:     if(current->sigaction[signum - 1].sa_handler == SIG_DFL) {
51:         if(signum != SIGCHLD) {
52:             current->sigpending &= SIG_MASK(signum);
53:         }
54:     }
55:     return (unsigned int)sighandler;
56: }

```

**kernel/syscalls/sigpending.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/sigpending.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/signal.h>
10: #include <fiwix/process.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_sigpending(__sigset_t *set)
18: {
19:     int errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_sigpending(0x%08x) -> ", current->pid, set);
23: #endif /*__DEBUG__*/
24:
25:     if((errno = check_user_area(VERIFY_WRITE, set, sizeof(__sigset_t)))) {
26:         return errno;
27:     }
28:     memcpy_b(set, &current->sigpending, sizeof(__sigset_t));
29:     return 0;
30: }
```



## kernel/syscalls/sigprocmask.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/sigprocmask.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/signal.h>
10: #include <fiwix/process.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_sigprocmask(int how, const __sigset_t *set, __sigset_t *oldset)
18: {
19:     int errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_sigprocmask(%d, 0x%08x, 0x%08x)\n", current->pid, h
ow, set, oldset);
23: #endif /*__DEBUG__*/
24:
25:     if(oldset) {
26:         if((errno = check_user_area(VERIFY_WRITE, oldset, sizeof(__sigse
t_t)))) {
27:             return errno;
28:         }
29:         *oldset = current->sigblocked;
30:     }
31:
32:     if(set) {
33:         if((errno = check_user_area(VERIFY_READ, set, sizeof(__sigset_t)
))) {
34:             return errno;
35:         }
36:         switch(how) {
37:             case SIG_BLOCK:
38:                 current->sigblocked |= (*set & SIG_BLOCKABLE);
39:                 break;
40:             case SIG_UNBLOCK:
41:                 current->sigblocked &= ~(*set & SIG_BLOCKABLE);
42:                 break;
43:             case SIG_SETMASK:
44:                 current->sigblocked = (*set & SIG_BLOCKABLE);
45:                 break;
46:             default:
47:                 return -EINVAL;
48:         }
49:     }
50:     return 0;
51: }

```

## kernel/syscalls/sigreturn.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/sigreturn.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9: #include <fiwix/sigcontext.h>
10: #include <fiwix/string.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #endif /*__DEBUG__ */
15:
16: int sys_sigreturn(unsigned int signum, int arg2, int arg3, int arg4, int arg5, s
truct sigcontext *sc)
17: {
18: #ifdef __DEBUG__
19:     printk("(pid %d) sys_sigreturn(0x%08x)\n", current->pid, signum);
20: #endif /*__DEBUG__ */
21:
22:     current->sigblocked &= ~current->sigexecuting;
23:     current->sigexecuting = 0;
24:     memcpy_b(sc, &current->sc[signum - 1], sizeof(struct sigcontext));
25:
26:     /*
27:      * We return here the value that the syscall was returning when it was
28:      * interrupted by a signal.
29:      */
30:     return current->sc[signum - 1].eax;
31: }
```

## kernel/syscalls/sigsuspend.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/sigsuspend.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/syscalls.h>
10: #include <fiwix/signal.h>
11: #include <fiwix/process.h>
12: #include <fiwix/errno.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_sigsuspend(__sigset_t *mask)
19: {
20:     __sigset_t old_mask;
21:     int errno;
22:
23: #ifdef __DEBUG__
24:     printk("(pid %d) sys_sigsuspend(0x%08x) -> ", current->pid, mask);
25: #endif /* __DEBUG__ */
26:
27:     old_mask = current->sigblocked;
28:     if(mask) {
29:         if((errno = check_user_area(VERIFY_READ, mask, sizeof(__sigset_t
)))) {
30:             return errno;
31:         }
32:         current->sigblocked = (int)*mask & SIG_BLOCKABLE;
33:     } else {
34:         current->sigblocked = 0 & SIG_BLOCKABLE;
35:     }
36:     sys_pause();
37:     current->sigblocked = old_mask;
38:
39: #ifdef __DEBUG__
40:     printk("-EINTR\n");
41: #endif /* __DEBUG__ */
42:
43:     return -EINTR;
44: }

```

**kernel/syscalls/socketcall.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/socketcall.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/errno.h>
9:
10: #ifdef __DEBUG__
11: #include <fiwix/stdio.h>
12: #include <fiwix/process.h>
13: #endif /*__DEBUG__ */
14:
15: int sys_socketcall(int call, unsigned long int *args)
16: {
17: #ifdef __DEBUG__
18:     printk("(pid %d) sys_socketcall(%d, 0x%08x) -> ENOENT\n", current->pid,
call, args);
19: #endif /*__DEBUG__ */
20:
21:     /* FIXME: to be implemented */
22:
23:     return -ENOENT;
24: }
```

## kernel/syscalls/ssetmask.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/ssetmask.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/process.h>
9: #include <fiwix/signal.h>
10:
11: #ifdef __DEBUG__
12: #include <fiwix/stdio.h>
13: #endif /*__DEBUG__*/
14:
15: int sys_ssetmask(int newmask)
16: {
17:     int oldmask;
18:
19: #ifdef __DEBUG__
20:     printk("(pid %d) sys_ssetmask(0x%08x) -> \n", current->pid, newmask);
21: #endif /*__DEBUG__*/
22:
23:     oldmask = current->sigblocked;
24:     current->sigblocked = newmask & SIG_BLOCKABLE;
25:     return oldmask;
26: }
```

## kernel/syscalls/stat.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/stat.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/statbuf.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_stat(const char *filename, struct old_stat *statbuf)
19: {
20:     struct inode *i;
21:     char *tmp_name;
22:     int errno;
23:
24: #ifdef __DEBUG__
25:     printk("(pid %d) sys_stat(%s, 0x%08x) -> returning structure\n", current
->pid, filename, (unsigned int)statbuf);
26: #endif /* __DEBUG__ */
27:
28:     if((errno = check_user_area(VERIFY_WRITE, statbuf, sizeof(struct old_sta
t)))) {
29:         return errno;
30:     }
31:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
32:         return errno;
33:     }
34:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS))) {
35:         free_name(tmp_name);
36:         return errno;
37:     }
38:     statbuf->st_dev = i->dev;
39:     statbuf->st_ino = i->inode;
40:     statbuf->st_mode = i->i_mode;
41:     statbuf->st_nlink = i->i_nlink;
42:     statbuf->st_uid = i->i_uid;
43:     statbuf->st_gid = i->i_gid;
44:     statbuf->st_rdev = i->rdev;
45:     statbuf->st_size = i->i_size;
46:     statbuf->st_atime = i->i_atime;
47:     statbuf->st_mtime = i->i_mtime;
48:     statbuf->st_ctime = i->i_ctime;
49:     iput(i);
50:     free_name(tmp_name);
51:     return 0;
52: }

```

## kernel/syscalls/statfs.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/statfs.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/statfs.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_statfs(const char *filename, struct statfs *statfsbuf)
19: {
20:     struct inode *i;
21:     char *tmp_name;
22:     int errno;
23:
24: #ifdef __DEBUG__
25:     printk("(pid %d) sys_statfs('%s', 0x%08x)\n", current->pid, filename, (u
nsigned int) statfsbuf);
26: #endif /* __DEBUG__ */
27:
28:     if((errno = check_user_area(VERIFY_WRITE, statfsbuf, sizeof(struct statf
s)))) {
29:         return errno;
30:     }
31:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
32:         return errno;
33:     }
34:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS))) {
35:         free_name(tmp_name);
36:         return errno;
37:     }
38:     if(i->sb && i->sb->fsop && i->sb->fsop->statfs) {
39:         i->sb->fsop->statfs(i->sb, statfsbuf);
40:         iput(i);
41:         free_name(tmp_name);
42:         return 0;
43:     }
44:     iput(i);
45:     free_name(tmp_name);
46:     return -ENOSYS;
47: }

```

## kernel/syscalls/stime.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/stime.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/timer.h>
11: #include <fiwix/errno.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_stime(__time_t *t)
19: {
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_stime(0x%08x)\n", current->pid, (unsigned int)t);
24: #endif /* __DEBUG__ */
25:
26:     if(!IS_SUPERUSER) {
27:         return -EPERM;
28:     }
29:     if((errno = check_user_area(VERIFY_READ, t, sizeof(__time_t)))) {
30:         return errno;
31:     }
32:
33:     set_system_time(*t);
34:     return 0;
35: }
```



## kernel/syscalls/symlink.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/symlink.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/stat.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/string.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_symlink(const char *oldpath, const char *newpath)
19: {
20:     struct inode *i, *dir;
21:     char *tmp_oldpath, *tmp_newpath, *basename;
22:     int errno;
23:
24: #ifdef __DEBUG__
25:     printk("(pid %d) sys_symlink('%s', '%s')\n", current->pid, oldpath, newp
ath);
26: #endif /* __DEBUG__ */
27:
28:     if((errno = malloc_name(oldpath, &tmp_oldpath)) < 0) {
29:         return errno;
30:     }
31:     if((errno = malloc_name(newpath, &tmp_newpath)) < 0) {
32:         free_name(tmp_oldpath);
33:         return errno;
34:     }
35:     basename = get_basename(tmp_newpath);
36:     if((errno = namei(tmp_newpath, &i, &dir, !FOLLOW_LINKS))) {
37:         if(!dir) {
38:             free_name(tmp_oldpath);
39:             free_name(tmp_newpath);
40:             return errno;
41:         }
42:     }
43:     if(!errno) {
44:         iput(i);
45:         iput(dir);
46:         free_name(tmp_oldpath);
47:         free_name(tmp_newpath);
48:         return -EEXIST;
49:     }
50:     if(IS_RDONLY_FS(dir)) {
51:         iput(dir);
52:         free_name(tmp_oldpath);
53:         free_name(tmp_newpath);
54:         return -EROFS;
55:     }
56:
57:     if(check_permission(TO_EXEC | TO_WRITE, dir) < 0) {
58:         iput(dir);
59:         free_name(tmp_oldpath);
60:         free_name(tmp_newpath);
61:         return -EACCES;
62:     }
63:
64:     if(dir->fsop && dir->fsop->symlink) {
65:         errno = dir->fsop->symlink(dir, basename, tmp_oldpath);
66:     } else {

```

**kernel/syscalls/symlink.c**

Page 2/2

```
67:             errno = -EPERM;
68:         }
69:         iput(dir);
70:         free_name(tmp_oldpath);
71:         free_name(tmp_newpath);
72:         return errno;
73: }
```

**kernel/syscalls/sync.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/sync.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/buffer.h>
10: #include <fiwix/filesystems.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /* __DEBUG__ */
16:
17: void sys_sync(void)
18: {
19: #ifdef __DEBUG__
20:     printk("(pid %d) sys_sync()\n", current->pid);
21: #endif /* __DEBUG__ */
22:
23:     sync_superblocks(0);    /* in all devices */
24:     sync_inodes(0);        /* in all devices */
25:     sync_buffers(0);       /* in all devices */
26:     return;
27: }
```

## kernel/syscalls/sysinfo.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/sysinfo.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/system.h>
11: #include <fiwix/sched.h>
12: #include <fiwix/mm.h>
13: #include <fiwix/string.h>
14:
15: #ifdef __DEBUG__
16: #include <fiwix/stdio.h>
17: #include <fiwix/process.h>
18: #endif /*__DEBUG__ */
19:
20: int sys_sysinfo(struct sysinfo *info)
21: {
22:     struct sysinfo tmp_info;
23:     struct proc *p;
24:     int errno;
25:
26: #ifdef __DEBUG__
27:     printk("(pid %d) sys_sysinfo(0x%08x)\n ", current->pid, (unsigned int)in
fo);
28: #endif /*__DEBUG__ */
29:
30:     if((errno = check_user_area(VERIFY_WRITE, info, sizeof(struct sysinfo)))
) {
31:         return errno;
32:     }
33:     memset_b(&tmp_info, NULL, sizeof(struct sysinfo));
34:     tmp_info.loads[0] = avenrun[0] << (SI_LOAD_SHIFT - FSHIFT);
35:     tmp_info.loads[1] = avenrun[1] << (SI_LOAD_SHIFT - FSHIFT);
36:     tmp_info.loads[2] = avenrun[2] << (SI_LOAD_SHIFT - FSHIFT);
37:     tmp_info.uptime = kstat.uptime;
38:     tmp_info.totalram = kstat.total_mem_pages << PAGE_SHIFT;
39:     tmp_info.freeram = kstat.free_pages << PAGE_SHIFT;
40:     tmp_info.sharedram = 0;
41:     tmp_info.bufferram = kstat.buffers * 1024;
42:     tmp_info.totalswap = 0;
43:     tmp_info.freeswap = 0;
44:     FOR_EACH_PROCESS(p) {
45:         if(p->state) {
46:             tmp_info.procs++;
47:         }
48:     }
49:
50:     memcpy_b(info, &tmp_info, sizeof(struct sysinfo));
51:     return 0;
52: }

```

## kernel/syscalls/time.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/time.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/fs.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_time(__time_t *tloc)
18: {
19:     int errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_time() -> ", current->pid);
23: #endif /*__DEBUG__*/
24:
25:     if(tloc) {
26:         if((errno = check_user_area(VERIFY_WRITE, tloc, sizeof(__time_t)
))) {
27:             return errno;
28:         }
29:         *tloc = CURRENT_TIME;
30:     }
31:
32: #ifdef __DEBUG__
33:     printk("%d\n", CURRENT_TIME);
34: #endif /*__DEBUG__*/
35:
36:     return CURRENT_TIME;
37: }
```

## kernel/syscalls/times.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/times.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/syscalls.h>
11: #include <fiwix/times.h>
12:
13: #ifdef __DEBUG__
14: #include <fiwix/stdio.h>
15: #include <fiwix/process.h>
16: #endif /* __DEBUG__ */
17:
18: int sys_times(struct tms *buf)
19: {
20:     int errno;
21:
22: #ifdef __DEBUG__
23:     printk("(pid %d) sys_times(0x%08x) -> ", (unsigned int )buf);
24: #endif /* __DEBUG__ */
25:
26:     if((errno = check_user_area(VERIFY_WRITE, buf, sizeof(struct tms)))) {
27:         return errno;
28:     }
29:     if(buf) {
30:         buf->tms_utime = tv2ticks(&current->usage.ru_utime);
31:         buf->tms_stime = tv2ticks(&current->usage.ru_stime);
32:         buf->tms_cutime = tv2ticks(&current->cusage.ru_utime);
33:         buf->tms_cstime = tv2ticks(&current->cusage.ru_stime);
34:     }
35:
36:     return kstat.ticks;
37: }
```

## kernel/syscalls/truncate.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/truncate.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/stat.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /*__DEBUG__ */
18:
19: int sys_truncate(const char *path, __off_t length)
20: {
21:     struct inode *i;
22:     char *tmp_name;
23:     int errno;
24:
25: #ifdef __DEBUG__
26:     printk("(pid %d) sys_truncate(%s, %d)\n", current->pid, path, length);
27: #endif /*__DEBUG__ */
28:
29:     if((errno = malloc_name(path, &tmp_name)) < 0) {
30:         return errno;
31:     }
32:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS)) {
33:         free_name(tmp_name);
34:         return errno;
35:     }
36:     if(S_ISDIR(i->i_mode)) {
37:         iput(i);
38:         free_name(tmp_name);
39:         return -EISDIR;
40:     }
41:     if(IS_RDONLY_FS(i)) {
42:         iput(i);
43:         free_name(tmp_name);
44:         return -EROFS;
45:     }
46:     if(check_permission(TO_WRITE, i) < 0) {
47:         iput(i);
48:         free_name(tmp_name);
49:         return -EACCES;
50:     }
51:     if(length == i->i_size) {
52:         iput(i);
53:         free_name(tmp_name);
54:         return 0;
55:     }
56:
57:     errno = 0;
58:     if(i->fsop && i->fsop->truncate) {
59:         inode_lock(i);
60:         errno = i->fsop->truncate(i, length);
61:         inode_unlock(i);
62:     }
63:     iput(i);
64:     free_name(tmp_name);
65:     return errno;
66: }

```

**kernel/syscalls/umask.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/umask.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/process.h>
10: #include <fiwix/stat.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #endif /*__DEBUG__*/
15:
16: int sys_umask(__mode_t mask)
17: {
18:     __mode_t old_umask;
19:
20: #ifdef __DEBUG__
21:     printk("(pid %d) sys_umask(%d)\n", current->pid, mask);
22: #endif /*__DEBUG__*/
23:
24:     old_umask = current->umask;
25:     current->umask = mask & (S_IRWXU | S_IRWXG | S_IRWXO);
26:     return old_umask;
27: }
```



## kernel/syscalls/umount2.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/umount2.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/sleep.h>
13: #include <fiwix/devices.h>
14: #include <fiwix/buffer.h>
15: #include <fiwix/errno.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: static struct resource umount_resource = { NULL, NULL };
20:
21: int sys_umount2(const char *target, int flags)
22: {
23:     struct inode *i_target;
24:     struct mount *mt = NULL;
25:     struct filesystems *fs;
26:     struct device *d;
27:     struct inode dummy_i;
28:     struct superblock *sb;
29:     char *tmp_target;
30:     __dev_t dev;
31:     int errno;
32:
33: #ifdef __DEBUG__
34:     printk("(pid %d) sys_umount2(%s, 0x%08x)\n", current->pid, target, flags
);
35: #endif /*__DEBUG__ */
36:
37:     if(!IS_SUPERUSER) {
38:         return -EPERM;
39:     }
40:     if((errno = malloc_name(target, &tmp_target)) < 0) {
41:         return errno;
42:     }
43:     if((errno = namei(tmp_target, &i_target, NULL, FOLLOW_LINKS)) {
44:         free_name(tmp_target);
45:         return errno;
46:     }
47:     if(!S_ISBLK(i_target->i_mode) && !S_ISDIR(i_target->i_mode)) {
48:         iput(i_target);
49:         free_name(tmp_target);
50:         return -EINVAL;
51:     }
52:
53:     if(!(mt = get_mount_point(i_target))) {
54:         iput(i_target);
55:         free_name(tmp_target);
56:         return -EINVAL;
57:     }
58:     if(S_ISBLK(i_target->i_mode)) {
59:         dev = i_target->rdev;
60:     } else {
61:         dev = i_target->sb->dev;
62:     }
63:
64:     if(!(sb = get_superblock(dev))) {
65:         printk("WARNING: %s(): unable to get superblock from device %d,%
d\n", __FUNCTION__, MAJOR(dev), MINOR(dev));

```

## kernel/syscalls/umount2.c

Page 2/2

```

66:         iput(i_target);
67:         free_name(tmp_target);
68:         return -EINVAL;
69:     }
70:
71:     /*
72:      * We must free now the inode in order to avoid having its 'count' to 2
73:      * when calling check_fs_busy(), specially if sys_umount() was called
74:      * using the mount-point instead of the device.
75:      */
76:     iput(i_target);
77:     free_name(tmp_target);
78:
79:     if(check_fs_busy(dev, sb->root)) {
80:         return -EBUSY;
81:     }
82:
83:     lock_resource(&umount_resource);
84:
85:     fs = mt->fs;
86:     if(fs->fsop && fs->fsop->release_superblock) {
87:         fs->fsop->release_superblock(sb);
88:     }
89:     if(sb->fsop->flags & FSOP_REQUIRES_DEV) {
90:         if(!(d = get_device(BLK_DEV, dev))) {
91:             printk("WARNING: %s(): block device %d,%d not registered
!\n", __FUNCTION__, MAJOR(dev), MINOR(dev));
92:             unlock_resource(&umount_resource);
93:             return -EINVAL;
94:         }
95:         memset_b(&dummy_i, 0, sizeof(struct inode));
96:         dummy_i.dev = dummy_i.rdev = dev;
97:         if(d && d->fsop && d->fsop->close) {
98:             d->fsop->close(&dummy_i, NULL);
99:         }
100:     }
101:
102:     sb->dir->mount_point = NULL;
103:     iput(sb->root);
104:     iput(sb->dir);
105:
106:     sync_superblocks(dev);
107:     sync_inodes(dev);
108:     sync_buffers(dev);
109:     invalidate_buffers(dev);
110:     invalidate_inodes(dev);
111:
112:     release_mount_point(mt);
113:     unlock_resource(&umount_resource);
114:     return 0;
115: }

```

**kernel/syscalls/umount.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/umount.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/syscalls.h>
9: #include <fiwix/process.h>
10:
11: #ifdef __DEBUG__
12: #include <fiwix/stdio.h>
13: #endif /*__DEBUG__ */
14:
15: int sys_umount(const char *target)
16: {
17: #ifdef __DEBUG__
18:     printk("(pid %d) sys_umount(%s)\n", current->pid, target);
19: #endif /*__DEBUG__ */
20:
21:     return sys_umount2(target, 0);
22: }
```

## kernel/syscalls/uname.c

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/uname.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/utsname.h>
10: #include <fiwix/string.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /* __DEBUG__ */
16:
17: int sys_uname(struct old_utsname *uname)
18: {
19:     int errno;
20:
21: #ifdef __DEBUG__
22:     printk("(pid %d) sys_uname(0x%08x) -> returning ", current->pid, (unsigned int)uname);
23: #endif /* __DEBUG__ */
24:
25:     if((errno = check_user_area(VERIFY_WRITE, uname, sizeof(struct old_utsname)))) {
26:         return errno;
27:     }
28:     memcpy_b(&uname->sysname, &sys_utsname.sysname, sizeof(sys_utsname.sysname));
29:     memcpy_b(&uname->nodename, &sys_utsname.nodename, sizeof(sys_utsname.nodename));
30:     memcpy_b(&uname->release, &sys_utsname.release, sizeof(sys_utsname.release));
31:     memcpy_b(&uname->version, &sys_utsname.version, sizeof(sys_utsname.version));
32:     memcpy_b(&uname->machine, &sys_utsname.machine, sizeof(sys_utsname.machine));
33:     return 0;
34: }
```

## kernel/syscalls/unlink.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/unlink.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/syscalls.h>
10: #include <fiwix/stat.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: #ifdef __DEBUG__
15: #include <fiwix/stdio.h>
16: #include <fiwix/process.h>
17: #endif /* __DEBUG__ */
18:
19: int sys_unlink(const char *filename)
20: {
21:     struct inode *i, *dir;
22:     char *tmp_name, *basename;
23:     int errno;
24:
25: #ifdef __DEBUG__
26:     printk("(pid %d) sys_unlink('%s')\n", current->pid, filename);
27: #endif /* __DEBUG__ */
28:
29:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
30:         return errno;
31:     }
32:     if((errno = namei(tmp_name, &i, &dir, !FOLLOW_LINKS)) {
33:         if(dir) {
34:             iput(dir);
35:         }
36:         free_name(tmp_name);
37:         return errno;
38:     }
39:     if(S_ISDIR(i->i_mode)) {
40:         iput(i);
41:         iput(dir);
42:         free_name(tmp_name);
43:         return -EPERM; /* Linux returns -EISDIR */
44:     }
45:     if(IS_RDONLY_FS(i)) {
46:         iput(i);
47:         iput(dir);
48:         free_name(tmp_name);
49:         return -EROFS;
50:     }
51:     if(check_permission(TO_EXEC | TO_WRITE, dir) < 0) {
52:         iput(i);
53:         iput(dir);
54:         free_name(tmp_name);
55:         return -EACCES;
56:     }
57:
58:     /* check sticky permission bit */
59:     if(dir->i_mode & S_ISVTX) {
60:         if(check_user_permission(i) {
61:             iput(i);
62:             iput(dir);
63:             free_name(tmp_name);
64:             return -EPERM;
65:         }
66:     }
67:

```

**kernel/syscalls/unlink.c**

Page 2/2

```
68:         basename = get_basename(filename);
69:         if(dir->fsop && dir->fsop->unlink) {
70:             errno = dir->fsop->unlink(dir, i, basename);
71:         } else {
72:             errno = -EPERM;
73:         }
74:         iput(i);
75:         iput(dir);
76:         free_name(tmp_name);
77:         return errno;
78: }
```

## kernel/syscalls/ustat.c

Page 1/1

```

1:  /*
2:  *  fiwix/kernel/syscalls/ustat.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/types.h>
9:  #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/ustat.h>
12: #include <fiwix/statfs.h>
13: #include <fiwix/errno.h>
14: #include <fiwix/string.h>
15:
16: #ifdef __DEBUG__
17: #include <fiwix/stdio.h>
18: #include <fiwix/process.h>
19: #endif /*__DEBUG__ */
20:
21: int sys_ustat(__dev_t dev, struct ustat *ubuf)
22: {
23:     struct superblock *sb;
24:     struct statfs statfsbuf;
25:     int errno;
26:
27: #ifdef __DEBUG__
28:     printk("(pid %d) sys_ustat(%d, 0x%08x)\n", current->pid, dev, (int)ubuf)
;
29: #endif /*__DEBUG__ */
30:     if((errno = check_user_area(VERIFY_WRITE, ubuf, sizeof(struct ustat))))
{
31:         return errno;
32:     }
33:     if(!(sb = get_superblock(dev))) {
34:         return -EINVAL;
35:     }
36:     if(sb->fsop && sb->fsop->statfs) {
37:         sb->fsop->statfs(sb, &statfsbuf);
38:         memset_b(ubuf, NULL, sizeof(struct ustat));
39:         ubuf->f_tfree = statfsbuf.f_bfree;
40:         ubuf->f_tinode = statfsbuf.f_ffree;
41:         return 0;
42:     }
43:     return -ENOSYS;
44: }

```

## kernel/syscalls/utime.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/utime.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/utime.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/errno.h>
13: #include <fiwix/string.h>
14:
15: #ifdef __DEBUG__
16: #include <fiwix/stdio.h>
17: #include <fiwix/process.h>
18: #endif /*__DEBUG__ */
19:
20: int sys_utime(const char *filename, struct utimbuf *times)
21: {
22:     struct inode *i;
23:     char *tmp_name;
24:     int errno;
25:
26: #ifdef __DEBUG__
27:     printk("(pid %d) sys_utime('%s', 0x%08x)\n", current->pid, filename, (in
t)times);
28: #endif /*__DEBUG__ */
29:
30:     if((errno = malloc_name(filename, &tmp_name)) < 0) {
31:         return errno;
32:     }
33:     if((errno = namei(tmp_name, &i, NULL, FOLLOW_LINKS))) {
34:         free_name(tmp_name);
35:         return errno;
36:     }
37:
38:     if(IS_RDONLY_FS(i)) {
39:         iput(i);
40:         free_name(tmp_name);
41:         return -EROFS;
42:     }
43:
44:     if(!times) {
45:         if(check_user_permission(i) || check_permission(TO_WRITE, i)) {
46:             iput(i);
47:             free_name(tmp_name);
48:             return -EACCES;
49:         }
50:         i->i_atime = CURRENT_TIME;
51:         i->i_mtime = CURRENT_TIME;
52:     } else {
53:         if((errno = check_user_area(VERIFY_READ, times, sizeof(struct ut
imbuf)))) {
54:             iput(i);
55:             free_name(tmp_name);
56:             return errno;
57:         }
58:         if(check_user_permission(i)) {
59:             iput(i);
60:             free_name(tmp_name);
61:             return -EPERM;
62:         }
63:         i->i_atime = times->actime;
64:         i->i_mtime = times->modtime;
65:     }

```



**kernel/syscalls/utime.c**

Page 2/2

```
66:
67:         i->i_ctime = CURRENT_TIME;
68:         i->dirty = 1;
69:         iput(i);
70:         free_name(tmp_name);
71:         return 0;
72: }
```

## kernel/syscalls/wait4.c

Page 1/2

```

1: /*
2:  * fiwix/kernel/syscalls/wait4.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/resource.h>
11: #include <fiwix/signal.h>
12: #include <fiwix/sched.h>
13: #include <fiwix/sleep.h>
14: #include <fiwix/errno.h>
15:
16: #ifdef __DEBUG__
17: #include <fiwix/stdio.h>
18: #include <fiwix/process.h>
19: #endif /*__DEBUG__ */
20:
21: int sys_wait4(__pid_t pid, int *status, int options, struct rusage *ru)
22: {
23:     struct proc *p;
24:     int flag, signum, errno;
25:
26: #ifdef __DEBUG__
27:     printk("(pid %d) sys_wait4(%d, status, %d)\n", current->pid, pid, option
s);
28: #endif /*__DEBUG__ */
29:
30:     if(ru) {
31:         if((errno = check_user_area(VERIFY_WRITE, ru, sizeof(struct rusa
ge)))) {
32:             return errno;
33:         }
34:     }
35:     while(current->children) {
36:         flag = 0;
37:         FOR_EACH_PROCESS(p) {
38:             if(p->ppid != current->pid) {
39:                 continue;
40:             }
41:             if(pid > 0) {
42:                 if(p->pid == pid) {
43:                     flag = 1;
44:                 }
45:             }
46:             if(!pid) {
47:                 if(p->pgid == current->pgid) {
48:                     flag = 1;
49:                 }
50:             }
51:             if(pid < -1) {
52:                 if(p->pgid == -pid) {
53:                     flag = 1;
54:                 }
55:             }
56:             if(pid == -1) {
57:                 flag = 1;
58:             }
59:             if(flag) {
60:                 if(p->state == PROC_STOPPED) {
61:                     if(!p->exit_code) {
62:                         continue;
63:                     }
64:                     if(status) {
65:                         *status = (p->exit_code << 8) |

```

## kernel/syscalls/wait4.c

Page 2/2

```
0x7F;
66:                                     }
67:                                     p->exit_code = 0;
68:                                     if(ru) {
69:                                         get_rusage(p, ru);
70:                                     }
71:                                     return p->pid;
72:                                     }
73:                                     if(p->state == PROC_ZOMBIE) {
74:                                         add_rusage(p);
75:                                         if(status) {
76:                                             *status = p->exit_code;
77:                                         }
78:                                         if(ru) {
79:                                             get_rusage(p, ru);
80:                                         }
81:                                         return remove_zombie(p);
82:                                     }
83:                                     }
84:                                     flag = 0;
85:                                     }
86:                                     if(options & WNOHANG) {
87:                                         if(flag) {
88:                                             return 0;
89:                                         }
90:                                         break;
91:                                     }
92:                                     if((signum = sleep(&sys_wait4, PROC_INTERRUPTIBLE)) {
93:                                         return signum;
94:                                     }
95:                                     current->sigpending &= SIG_MASK(SIGCHLD);
96:                                     }
97:                                     return -ECHILD;
98:     }
```

**kernel/syscalls/waitpid.c**

Page 1/1

```
1: /*
2:  * fiwix/kernel/syscalls/waitpid.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/syscalls.h>
10: #include <fiwix/string.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_waitpid(__pid_t pid, int *status, int options)
18: {
19: #ifdef __DEBUG__
20:     printk("(pid %d) sys_waitpid(%d, 0x%08x, %d)\n", current->pid, pid, sta
tus ? *status : 0, options);
21: #endif /*__DEBUG__*/
22:     return sys_wait4(pid, status, options, NULL);
23: }
```

## kernel/syscalls/write.c

Page 1/1

```

1: /*
2:  * fiwix/kernel/syscalls/write.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/fs.h>
9: #include <fiwix/fcntl.h>
10: #include <fiwix/errno.h>
11:
12: #ifdef __DEBUG__
13: #include <fiwix/stdio.h>
14: #include <fiwix/process.h>
15: #endif /*__DEBUG__*/
16:
17: int sys_write(unsigned int ufd, const char *buf, int count)
18: {
19:     struct inode *i;
20:     int errno;
21:
22: #ifdef __DEBUG__
23: /*      printk("(pid %d) sys_write(%d, '%s', %d)\n", current->pid, ufd, buf, cou
nt);*/
24:     printk("(pid %d) sys_write(%d, 0x%08x, %d) -> ", current->pid, ufd, buf,
count);
25: #endif /*__DEBUG__*/
26:
27:     CHECK_UFD(ufd);
28:     if((errno = check_user_area(VERIFY_READ, buf, count))) {
29:         return errno;
30:     }
31:     if(fd_table[current->fd[ufd]].flags & O_RDONLY) {
32:         return -EBADF;
33:     }
34:     if(!count) {
35:         return 0;
36:     }
37:     if(count < 0) {
38:         return -EINVAL;
39:     }
40:     i = fd_table[current->fd[ufd]].inode;
41:     if(i->fsop && i->fsop->write) {
42:         errno = i->fsop->write(i, &fd_table[current->fd[ufd]], buf, coun
t);
43: #ifdef __DEBUG__
44:         printk("%d\n", errno);
45: #endif /*__DEBUG__*/
46:         return errno;
47:     }
48:     return -EINVAL;
49: }

```

## drivers/block/dma.c

Page 1/2

```

1:  /*
2:  *  fiwix/drivers/block/dma.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/asm.h>
9:  #include <fiwix/dma.h>
10: #include <fiwix/string.h>
11:
12: /*
13:  *  DMA Channel  Page  Address  Count
14:  *  -----
15:  *  0 (8 bit)      87h      0h      1h
16:  *  1 (8 bit)      83h      2h      3h
17:  *  2 (8 bit)      81h      4h      5h
18:  *  3 (8 bit)      82h      6h      7h
19:  *  4 (16 bit)     8Fh      C0h     C2h
20:  *  5 (16 bit)     8Bh      C4h     C6h
21:  *  6 (16 bit)     89h      C8h     CAh
22:  *  7 (16 bit)     8Ah      CCh     CEh
23:  */
24:
25: #define LOW_BYTE(addr)  (addr & 0x00FF)
26: #define HIGH_BYTE(addr) ((addr & 0xFF00) >> 8)
27:
28: unsigned char dma_mask[DMA_CHANNELS] =
29:     { 0x0A, 0x0A, 0x0A, 0x0A, 0xD4, 0xD4, 0xD4, 0xD4 };
30: unsigned char dma_mode[DMA_CHANNELS] =
31:     { 0x0B, 0x0B, 0x0B, 0x0B, 0xD6, 0xD6, 0xD6, 0xD6 };
32: unsigned char dma_clear[DMA_CHANNELS] =
33:     { 0x0C, 0x0C, 0x0C, 0x0C, 0xD8, 0xD8, 0xD8, 0xD8 };
34: unsigned char dma_page[DMA_CHANNELS] =
35:     { 0x87, 0x83, 0x81, 0x82, 0x8F, 0x8B, 0x89, 0x8A };
36: unsigned char dma_address[DMA_CHANNELS] =
37:     { 0x00, 0x02, 0x04, 0x06, 0xC0, 0xC4, 0xC8, 0xCC };
38: unsigned char dma_count[DMA_CHANNELS] =
39:     { 0x01, 0x03, 0x05, 0x07, 0xC2, 0xC6, 0xCA, 0xCE };
40:
41:
42: void start_dma(int channel, void *address, unsigned int count, int mode)
43: {
44:     /* setup (mask) the DMA channel */
45:     outport_b(dma_mask[channel], DMA_MASK_CHANNEL | channel);
46:
47:     /* clear any data transfers that are currently executing */
48:     outport_b(dma_clear[channel], 0);
49:
50:     /* set the specified mode */
51:     outport_b(dma_mode[channel], mode | channel);
52:
53:     /* set the offset address */
54:     outport_b(dma_address[channel], LOW_BYTE((unsigned int)address));
55:     outport_b(dma_address[channel], HIGH_BYTE((unsigned int)address));
56:
57:     /* set the physical page */
58:     outport_b(dma_page[channel], (unsigned int)address >> 16);
59:
60:     /* the true (internal) length sent to the DMA is actually length + 1 */
61:     count--;
62:
63:     /* set the length of the data */
64:     outport_b(dma_count[channel], LOW_BYTE(count));
65:     outport_b(dma_count[channel], HIGH_BYTE(count));
66:
67:     /* clear the mask */

```

**drivers/block/dma.c**

Page 2/2

```
68:         outport_b(dma_mask[channel], DMA_UNMASK_CHANNEL | channel);
69:     }
70:
71: int dma_register(int channel, char *dev_name)
72: {
73:     if(dma_resources[channel]) {
74:         return 1;
75:     }
76:     dma_resources[channel] = dev_name;
77:     return 0;
78: }
79:
80: int dma_unregister(int channel)
81: {
82:     if(!dma_resources[channel]) {
83:         return 1;
84:     }
85:
86:     dma_resources[channel] = NULL;
87:     return 0;
88: }
89:
90: void dma_init(void)
91: {
92:     memset_b(dma_resources, NULL, sizeof(dma_resources));
93: }
```

## drivers/block/floppy.c

Page 1/14

```

1: /*
2:  * fiwix/drivers/block/floppy.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/floppy.h>
10: #include <fiwix/ioctl.h>
11: #include <fiwix/devices.h>
12: #include <fiwix/part.h>
13: #include <fiwix/fs.h>
14: #include <fiwix/buffer.h>
15: #include <fiwix/sleep.h>
16: #include <fiwix/timer.h>
17: #include <fiwix/sched.h>
18: #include <fiwix/errno.h>
19: #include <fiwix/pic.h>
20: #include <fiwix/cmos.h>
21: #include <fiwix/dma.h>
22: #include <fiwix/stdio.h>
23: #include <fiwix/string.h>
24:
25: #define WAIT_MOTOR_OFF (3 * HZ)      /* time waiting to turn the motor off */
26: #define WAIT_FDC        WAIT_MOTOR_OFF
27:
28: #define INVALID_TRACK  -1
29:
30: #define DEV_TYPE_SHIFT 2             /* right shift to match with the floppy
31:                                     type when minor > 3 */
32:
33: static int need_reset = 0;
34: static int fdc_wait_interrupt = 0;
35: static int fdc_timeout = 0;
36: static unsigned char fdc_results[MAX_FDC_RESULTS];
37: static struct resource floppy_resource = { NULL, NULL };
38:
39: static struct fddt fdd_type[] = {
40: /*
41:  * R (data rate): 0 = 500Kb/s, 2 = 250Kb/s, 3 = 1Mb/s
42:  * SPEC(IFY) 0xAF: SRT = 6ms, HUT = 240ms (500Kb/s)
43:  * SPEC(IFY) 0xD7: SRT = 6ms, HUT = 240ms (250Kb/s)
44:  * SPEC(IFY) 0xDF: SRT = 3ms, HUT = 240ms (500Kb/s)
45:  * Head Load Time 0x02: HLT = 4ms (500Kb/s), Non-DMA = 0 (DMA enabled)
46:  *
47:  *          SIZE      KB   T   S   H   G_RW  G_FM  R   SPEC  HLT   NAME
48:  *          -----
49:  *          {    0,    0,  0,  0,  0, 0x00, 0x00, 0, 0x00, 0x00, NULL
50:  *          {  720,   360, 40,  9,  2, 0x2A, 0x50, 2, 0xD7, 0x02, "360KB 5.25\""},
51:  *          { 2400, 1200, 80, 15,  2, 0x2A, 0x50, 0, 0xAF, 0x02, "1.2MB 5.25\""},
52:  *          { 1440,   720, 80,  9,  2, 0x1B, 0x54, 2, 0xD7, 0x02, "720KB 3.5\""},
53:  *          { 2880, 1440, 80, 18,  2, 0x1B, 0x54, 0, 0xAF, 0x02, "1.44MB 3.5\""},
54:  *          { 5760, 2880, 80, 36,  2, 0x38, 0x53, 3, 0xDF, 0x02, "2.88MB 3.5\""}, */
55: };
56:
57: /* maximum size of a track for floppy types of 1.44MB */
58: extern char _fdc_transfer_area[BPS * 2 * 18];
59:
60: struct fdd_status {
61:     char type;                /* floppy disk drive type */
62:     char motor;
63:     char recalibrated;
64:     char current_track;
65: };
66:
67: static struct fdd_status fdd_status[] = {

```



## drivers/block/floppy.c

Page 2/14

```

68:         { 0, 0, 0, INVALID_TRACK },
69:         { 0, 0, 0, INVALID_TRACK },
70: };
71:
72: static unsigned char current_fdd = 0;
73: static struct fddt *current_fdd_type;
74: static unsigned int fdd_sizes[256];
75:
76: static struct fs_operations fdc_driver_fsop = {
77:     0,
78:     0,
79:
80:     fdc_open,
81:     fdc_close,
82:     NULL,                /* read */
83:     NULL,                /* write */
84:     fdc_ioctl,
85:     fdc_lseek,
86:     NULL,                /* readdir */
87:     NULL,                /* mmap */
88:     NULL,                /* select */
89:
90:     NULL,                /* readlink */
91:     NULL,                /* followlink */
92:     NULL,                /* bmap */
93:     NULL,                /* lockup */
94:     NULL,                /* rmdir */
95:     NULL,                /* link */
96:     NULL,                /* unlink */
97:     NULL,                /* symlink */
98:     NULL,                /* mkdir */
99:     NULL,                /* mknod */
100:    NULL,                /* truncate */
101:    NULL,                /* create */
102:    NULL,                /* rename */
103:
104:    fdc_read,
105:    fdc_write,
106:
107:    NULL,                /* read_inode */
108:    NULL,                /* write_inode */
109:    NULL,                /* ialloc */
110:    NULL,                /* ifree */
111:    NULL,                /* statfs */
112:    NULL,                /* read_superblock */
113:    NULL,                /* remount_fs */
114:    NULL,                /* write_superblock */
115:    NULL,                /* release_superblock */
116: };
117:
118: static struct device floppy_device = {
119:     "floppy",
120:     FDC_MAJOR,
121:     { 0, 0, 0, 0, 0, 0, 0, 0 },
122:     BLKSIZE_1K,
123:     &fdd_sizes,
124:     &fdc_driver_fsop,
125:     NULL
126: };
127:
128: static struct interrupt irq_config_floppy = { 0, "floppy", &irq_floppy, NULL };
129:
130: static int fdc_in(void)
131: {
132:     int n;
133:     unsigned char status;
134:

```

## drivers/block/floppy.c

Page 3/14

```

135:         if(need_reset) {
136:             return -1;
137:         }
138:
139:         for(n = 0; n < 10000; n++) {
140:             status = inport_b(FDC_MSR) & (FDC_RQM | FDC_DIO);
141:             if(status == FDC_RQM) {
142:                 return 0;
143:             }
144:             if(status == (FDC_RQM | FDC_DIO)) {
145:                 return inport_b(FDC_DATA);
146:             }
147:         }
148:         need_reset = 1;
149:         printk("WARNING: %s(): fd%d: timeout on %s.\n", __FUNCTION__, current_fd
d, floppy_device.name);
150:         return -1;
151:     }
152:
153:     static void fdc_out(unsigned char value)
154:     {
155:         int n;
156:         unsigned char status;
157:
158:         if(need_reset) {
159:             return;
160:         }
161:
162:         for(n = 0; n < 10000; n++) {
163:             status = inport_b(FDC_MSR) & (FDC_RQM | FDC_DIO);
164:             if(status == FDC_RQM) {
165:                 outport_b(FDC_DATA, value);
166:                 return;
167:             }
168:         }
169:
170:         need_reset = 1;
171:         printk("WARNING: %s(): fd%d: unable to send byte 0x%x on %s.\n", __FUNCT
ION__, current_fdd, value, floppy_device.name);
172:     }
173:
174:     static void fdc_get_results(void)
175:     {
176:         int n;
177:
178:         memset_b(fdc_results, 0, sizeof(fdc_results));
179:         for(n = 0; n < MAX_FDC_RESULTS; n++) {
180:             fdc_results[n] = fdc_in();
181:         }
182:         return;
183:     }
184:
185:     static int fdc_motor_on(void)
186:     {
187:         struct callout_req creq;
188:         int errno;
189:
190:         if(fdd_status[current_fdd].motor) {
191:             return 0;
192:         }
193:
194:         /* select floppy disk drive and turn on its motor */
195:         outport_b(FDC_DOR, (FDC_DRIVE0 << current_fdd) | FDC_DMA_ENABLE | FDC_EN
ABLE | current_fdd);
196:         fdd_status[current_fdd].motor = 1;
197:         fdd_status[!current_fdd].motor = 0;
198:

```

## drivers/block/floppy.c

Page 4/14

```

199:      /* fixed spin-up time of 500ms for 3.5" and 5.25" */
200:      creq.fn = fdc_timer;
201:      creq.arg = FDC_TR_MOTOR;
202:      add_callout(&creq, HZ / 2);
203:      sleep(&fdc_motor_on, PROC_UNINTERRUPTIBLE);
204:
205:      errno = 0;
206:
207:      /* check for a disk change */
208:      if(inport_b(FDC_DIR) & 0x80) {
209:          errno = 1;
210:      }
211:
212:      return errno;
213: }
214:
215: static void do_motor_off(unsigned int fdd)
216: {
217:     outport_b(FDC_DOR, FDC_DMA_ENABLE | FDC_ENABLE | fdd);
218:     fdd_status[fdd].motor = 0;
219:     fdd_status[0].motor = fdd_status[1].motor = 0;
220: }
221:
222: static void fdc_motor_off(void)
223: {
224:     struct callout_req creq;
225:
226:     creq.fn = do_motor_off;
227:     creq.arg = current_fdd;
228:     add_callout(&creq, WAIT_FDC);
229: }
230:
231: static void fdc_reset(void)
232: {
233:     int n;
234:     struct callout_req creq;
235:
236:     need_reset = 0;
237:
238:     fdc_wait_interrupt = FDC_RESET;
239:     outport_b(FDC_DOR, 0);          /* enter in reset mode */
240: /* outport_b(FDC_DOR, FDC_DMA_ENABLE); */
241:     for(n = 0; n < 1000; n++) {    /* recovery time */
242:         NOP();
243:     }
244:     outport_b(FDC_DOR, FDC_DMA_ENABLE | FDC_ENABLE);
245:
246:     creq.fn = fdc_timer;
247:     creq.arg = FDC_TR_DEFAULT;
248:     add_callout(&creq, WAIT_FDC);
249:     sleep(&irq_floppy, PROC_UNINTERRUPTIBLE);
250:     if(fdc_timeout) {
251:         need_reset = 1;
252:         printk("WARNING: %s(): fd%d: timeout on %s.\n", __FUNCTION__, cu
rrent_fdd, floppy_device.name);
253:     }
254:     del_callout(&creq);
255:
256:     fdd_status[0].motor = fdd_status[1].motor = 0;
257:     fdd_status[current_fdd].recalibrated = 0;
258:
259:     /* assumes drive polling mode is ON (by default) */
260:     for(n = 0; n < 4; n++) {
261:         fdc_out(FDC_SENSEI);
262:         fdc_get_results();
263:     }
264:

```

## drivers/block/floppy.c

Page 5/14

```

265:      /* keeps controller informed on the drive about to use */
266:      fdc_out(FDC_SPECIFY);
267:      fdc_out(current_fdd_type->spec);
268:      fdc_out(current_fdd_type->hlt);
269:
270:      /* set data rate */
271:      outport_b(FDC_CCR, current_fdd_type->rate);
272: }
273:
274: static int fdc_recalibrate(void)
275: {
276:     struct callout_req creq;
277:
278:     if(need_reset) {
279:         return 1;
280:     }
281:
282:     fdc_wait_interrupt = FDC_RECALIBRATE;
283:     fdc_motor_on();
284:     fdc_out(FDC_RECALIBRATE);
285:     fdc_out(current_fdd);
286:
287:     if(need_reset) {
288:         return 1;
289:     }
290:
291:     creq.fn = fdc_timer;
292:     creq.arg = FDC_TR_DEFAULT;
293:     add_callout(&creq, WAIT_FDC);
294:     sleep(&irq_floppy, PROC_UNINTERRUPTIBLE);
295:     if(fdc_timeout) {
296:         need_reset = 1;
297:         printk("WARNING: %s(): fd%d: timeout on %s.\n", __FUNCTION__, cu
rrent_fdd, floppy_device.name);
298:         return 1;
299:     }
300:
301:     del_callout(&creq);
302:     fdc_out(FDC_SENSEI);
303:     fdc_get_results();
304:
305:     /* PCN must be 0 indicating a successful position to track 0 */
306:     if((fdc_results[ST0] & (ST0_IC | ST0_SE | ST0_UC | ST0_NR)) != ST0_RECAL
IBRATE || fdc_results[ST_PCN]) {
307:         need_reset = 1;
308:         printk("WARNING: %s(): fd%d: unable to recalibrate on %s.\n", __
FUNCTION__, current_fdd, floppy_device.name);
309:         return 1;
310:     }
311:
312:     fdd_status[current_fdd].current_track = INVALID_TRACK;
313:     fdd_status[current_fdd].recalibrated = 1;
314:     fdc_motor_off();
315:     return 0;
316: }
317:
318: static int fdc_seek(int track, int head)
319: {
320:     struct callout_req creq;
321:
322:     if(need_reset) {
323:         return 1;
324:     }
325:
326:     if(!fdd_status[current_fdd].recalibrated) {
327:         if(fdc_recalibrate()) {
328:             return 1;

```

## drivers/block/floppy.c

Page 6/14

```

329:         }
330:     }
331:
332:     if(fdd_status[current_fdd].current_track == track) {
333:         return 0;
334:     }
335:
336:     fdc_wait_interrupt = FDC_SEEK;
337:     fdc_motor_on();
338:     fdc_out(FDC_SEEK);
339:     fdc_out((head << 2) | current_fdd);
340:     fdc_out(track);
341:
342:     if(need_reset) {
343:         return 1;
344:     }
345:
346:     creq.fn = fdc_timer;
347:     creq.arg = FDC_TR_DEFAULT;
348:     add_callout(&creq, WAIT_FDC);
349:     sleep(&irq_floppy, PROC_UNINTERRUPTIBLE);
350:     if(fdc_timeout) {
351:         need_reset = 1;
352:         printk("WARNING: %s(): fd%d: timeout on %s.\n", __FUNCTION__, cu
rrrent_fdd, floppy_device.name);
353:         return 1;
354:     }
355:
356:     del_callout(&creq);
357:     fdc_out(FDC_SENSEI);
358:     fdc_get_results();
359:
360:     if((fdc_results[ST0] & (ST0_IC | ST0_SE | ST0_UC | ST0_NR)) != ST0_SEEK
|| fdc_results[ST_PCN] != track) {
361:         need_reset = 1;
362:         printk("WARNING: %s(): fd%d: unable to seek on %s.\n", __FUNCTIO
N__, current_fdd, floppy_device.name);
363:         return 1;
364:     }
365:
366:     fdc_motor_off();
367:     fdd_status[current_fdd].current_track = track;
368:     return 0;
369: }
370:
371: static int fdc_get_chip(void)
372: {
373:     unsigned char version, fifo, id;
374:
375:     fdc_out(FDC_VERSION);
376:     version = fdc_in();
377:     fdc_out(FDC_LOCK);
378:     fifo = fdc_in();
379:     fdc_out(FDC_PARTID);
380:     id = fdc_in();
381:
382:     if(version == 0x80) {
383:         if(fifo == 0x80) {
384:             printk("(NEC D765/Intel 8272A/compatible)\n");
385:             return 0;
386:         }
387:         if(fifo == 0) {
388:             printk("(Intel 82072)\n");
389:             return 0;
390:         }
391:     }
392:

```

## drivers/block/floppy.c

Page 7/14

```

393:         if(version == 0x81) {
394:             printk("(Very Early Intel 82077/compatible)\n");
395:             return 0;
396:         }
397:
398:         if(version == 0x90) {
399:             if(fifo == 0x80) {
400:                 printk("(Old Intel 82077, no FIFO)\n");
401:                 return 0;
402:             }
403:             if(fifo == 0) {
404:                 if(id == 0x80) {
405:                     printk("(New Intel 82077)\n");
406:                     return 0;
407:                 }
408:                 if(id == 0x41) {
409:                     printk("(Intel 82078)\n");
410:                     return 0;
411:                 }
412:                 if(id == 0x73) {
413:                     printk("(National Semiconductor PC87306)\n");
414:                     return 0;
415:                 }
416:                 printk("(Intel 82078 compatible)\n");
417:                 return 0;
418:             }
419:             printk("(NEC 72065B)\n");
420:             return 0;
421:         }
422:
423:         if(version == 0xA0) {
424:             printk("(SMC FDC37c65C+)\n");
425:             return 0;
426:         }
427:         printk("(unknown controller chip)\n");
428:         return 1;
429:     }
430:
431:     static int fdc_block2chs(__blk_t block, int blksize, int *cyl, int *head, int *s
ector)
432:     {
433:         int spb = blksize / FDC_SECTSIZE;
434:
435:         *cyl = (block * spb) / (current_fdd_type->spt * current_fdd_type->heads)
;
436:         *head = ((block * spb) % (current_fdd_type->spt * current_fdd_type->head
s)) / current_fdd_type->spt;
437:         *sector = (((block * spb) % (current_fdd_type->spt * current_fdd_type->h
eads)) % current_fdd_type->spt) + 1;
438:
439:         if(*cyl >= current_fdd_type->tracks || *head >= current_fdd_type->heads
|| *sector > current_fdd_type->spt) {
440:             return 1;
441:         }
442:
443:         return 0;
444:     }
445:
446:     static void set_current_fdd_type(int minor)
447:     {
448:         current_fdd = minor & 1;
449:
450:         /* minors 0 and 1 are directly assigned */
451:         if(minor < 2) {
452:             current_fdd_type = &fdd_type[(int)fdd_status[current_fdd].type];
453:         } else {
454:             current_fdd_type = &fdd_type[minor >> DEV_TYPE_SHIFT];

```

## drivers/block/floppy.c

Page 8/14

```

455:         }
456:     }
457:
458: void irq_floppy(int num, struct sigcontext *sc)
459: {
460:     if(!fdc_wait_interrupt) {
461:         printk("WARNING: %s(): fd%d: unexpected interrupt on %s.\n", __F
UNCTION__, current_fdd, floppy_device.name);
462:         need_reset = 1;
463:     } else {
464:         fdc_timeout = fdc_wait_interrupt = 0;
465:         wakeup(&irq_floppy);
466:     }
467: }
468:
469: void fdc_timer(unsigned int reason)
470: {
471:     switch(reason) {
472:         case FDC_TR_DEFAULT:
473:             fdc_timeout = 1;
474:             fdc_wait_interrupt = 0;
475:             wakeup(&irq_floppy);
476:             break;
477:         case FDC_TR_MOTOR:
478:             wakeup(&fdc_motor_on);
479:             break;
480:     }
481: }
482:
483: int fdc_open(struct inode *i, struct fd *fd_table)
484: {
485:     unsigned char minor;
486:
487:     minor = MINOR(i->rdev);
488:     if(!TEST_MINOR(floppy_device.minors, minor)) {
489:         return -ENXIO;
490:     }
491:
492:     lock_resource(&floppy_resource);
493:     set_current_fdd_type(minor);
494:     unlock_resource(&floppy_resource);
495:
496:     return 0;
497: }
498:
499: int fdc_close(struct inode *i, struct fd *fd_table)
500: {
501:     unsigned char minor;
502:
503:     minor = MINOR(i->rdev);
504:     if(!TEST_MINOR(floppy_device.minors, minor)) {
505:         return -ENXIO;
506:     }
507:
508:     lock_resource(&floppy_resource);
509:     set_current_fdd_type(minor);
510:     unlock_resource(&floppy_resource);
511:
512:     return 0;
513: }
514:
515: int fdc_read(__dev_t dev, __blk_t block, char *buffer, int blksize)
516: {
517:     unsigned char minor;
518:     unsigned int sectors_read;
519:     int cyl, head, sector;
520:     int retries;

```

## drivers/block/floppy.c

Page 9/14

```

521:         struct callout_req creq;
522:         struct device *d;
523:
524:         minor = MINOR(dev);
525:         if(!TEST_MINOR(floppy_device.minors, minor)) {
526:             return -ENXIO;
527:         }
528:
529:         if(!blksize) {
530:             if(!(d = get_device(BLK_DEV, dev))) {
531:                 return -EINVAL;
532:             }
533:             blksize = d->blksize;
534:         }
535:         blksize = blksize ? blksize : BLKSIZE_1K;
536:
537:         lock_resource(&floppy_resource);
538:         set_current_fdd_type(minor);
539:
540:         if(fdc_block2chs(block, blksize, &cyl, &head, &sector)) {
541:             printk("WARNING: %s(): fd%d: invalid block number %d on %s device
e %d,%d.\n", __FUNCTION__, current_fdd, block, floppy_device.name, MAJOR(dev), MINOR(de
v));
542:             unlock_resource(&floppy_resource);
543:             return -EINVAL;
544:         }
545:
546:         for(retries = 0; retries < MAX_FDC_ERR; retries++) {
547:             if(need_reset) {
548:                 fdc_reset();
549:             }
550:             if(fdc_motor_on()) {
551:                 printk("%s(): %s disk was changed in device %d,%d!\n", _
__FUNCTION__, floppy_device.name, MAJOR(dev), MINOR(dev));
552:                 invalidate_buffers(dev);
553:                 fdd_status[current_fdd].recalibrated = 0;
554:             }
555:
556:             if(fdc_seek(cyl, head)) {
557:                 printk("WARNING: %s(): fd%d: seek error on %s device %d,
%d during read operation.\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(dev)
, MINOR(dev));
558:                 continue;
559:             }
560:
561:             start_dma(FLOPPY_DMA, _fdc_transfer_area, blksize, DMA_MODE_WRIT
E | DMA_MODE_SINGLE);
562:
563:             /* send READ command */
564:             fdc_wait_interrupt = FDC_READ;
565:             fdc_out(FDC_READ);
566:             fdc_out((head << 2) | current_fdd);
567:             fdc_out(cyl);
568:             fdc_out(head);
569:             fdc_out(sector);
570:             fdc_out(2); /* sector size is 512 bytes */
571:             fdc_out(current_fdd_type->spt);
572:             fdc_out(current_fdd_type->gpl1);
573:             fdc_out(0xFF); /* sector size is 512 bytes */
574:
575:             if(need_reset) {
576:                 printk("WARNING: %s(): fd%d: needs reset on %s device %d
,%d during read operation.\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(dev)
, MINOR(dev));
577:                 continue;
578:             }
579:             creq.fn = fdc_timer;

```



## drivers/block/floppy.c

Page 10/14

```

580:         creq.arg = FDC_TR_DEFAULT;
581:         add_callout(&creq, WAIT_FDC);
582:         sleep(&irq_floppy, PROC_UNINTERRUPTIBLE);
583:         if(fdc_timeout) {
584:             need_reset = 1;
585:             printk("WARNING: %s(): fd%d: timeout on %s device %d,%d.
\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(dev), MINOR(dev));
586:             continue;
587:         }
588:         del_callout(&creq);
589:         fdc_get_results();
590:         if(fdc_results[ST0] & (ST0_IC | ST0_UC | ST0_NR)) {
591:             need_reset = 1;
592:             continue;
593:         }
594:         break;
595:     }
596:
597:     if(retries >= MAX_FDC_ERR) {
598:         printk("WARNING: %s(): fd%d: error on %s device %d,%d during rea
d operation,\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(dev), MINOR(dev))
;
599:         printk("\tblock=%d, sector=%d, cylinder/head=%d/%d\n", block, se
ctor, cyl, head);
600:         unlock_resource(&floppy_resource);
601:         fdc_motor_off();
602:         return -EIO;
603:     }
604:
605:     fdc_motor_off();
606:     sectors_read = (fdc_results[ST_CYL] - cyl) * (current_fdd_type->heads *
current_fdd_type->spt);
607:     sectors_read += (fdc_results[ST_HEAD] - head) * current_fdd_type->spt;
608:     sectors_read += fdc_results[ST_SECTOR] - sector;
609:     if(sectors_read * BPS != blksize) {
610:         printk("WARNING: %s(): fd%d: read error on %s device %d,%d (%d s
ectors read).\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(dev), MINOR(dev)
, sectors_read);
611:         printk("\tblock=%d, sector=%d, cylinder/head=%d/%d\n", block, se
ctor, cyl, head);
612:         unlock_resource(&floppy_resource);
613:         fdc_motor_off();
614:         return -EIO;
615:     }
616:
617:     memcpy_b(buffer, (void *)_fdc_transfer_area, blksize);
618:
619:     unlock_resource(&floppy_resource);
620:     return sectors_read * BPS;
621: }
622:
623: int fdc_write(__dev_t dev, __blk_t block, char *buffer, int blksize)
624: {
625:     unsigned char minor;
626:     unsigned int sectors_written;
627:     int cyl, head, sector;
628:     int retries;
629:     struct callout_req creq;
630:     struct device *d;
631:
632:     minor = MINOR(dev);
633:     if(!TEST_MINOR(floppy_device.minors, minor)) {
634:         return -ENXIO;
635:     }
636:
637:     if(!blksize) {
638:         if(!(d = get_device(BLK_DEV, dev))) {

```

## drivers/block/floppy.c

Page 11/14

```

639:                 return -EINVAL;
640:             }
641:             blksize = d->blksize;
642:         }
643:         blksize = blksize ? blksize : BLKSIZE_1K;
644:
645:         lock_resource(&floppy_resource);
646:         set_current_fdd_type(minor);
647:
648:         if(fdc_block2chs(block, blksize, &cyl, &head, &sector)) {
649:             printk("WARNING: %s(): fd%d: invalid block number %d on %s device
e %d,%d.\n", __FUNCTION__, current_fdd, block, floppy_device.name, MAJOR(dev), MINOR(de
v));
650:             unlock_resource(&floppy_resource);
651:             return -EINVAL;
652:         }
653:
654:         for(retries = 0; retries < MAX_FDC_ERR; retries++) {
655:             if(need_reset) {
656:                 fdc_reset();
657:             }
658:             if(fdc_motor_on()) {
659:                 printk("%s(): %s disk was changed in device %d,%d!\n", _
__FUNCTION__, floppy_device.name, MAJOR(dev), MINOR(dev));
660:                 invalidate_buffers(dev);
661:                 fdd_status[current_fdd].recalibrated = 0;
662:             }
663:
664:             if(fdc_seek(cyl, head)) {
665:                 printk("WARNING: %s(): fd%d: seek error on %s device %d,
%d during write operation.\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(dev
), MINOR(dev));
666:                 continue;
667:             }
668:
669:             start_dma(FLOPPY_DMA, _fdc_transfer_area, blksize, DMA_MODE_READ
| DMA_MODE_SINGLE);
670:             memcpy_b((void *)_fdc_transfer_area, buffer, blksize);
671:
672:             /* send WRITE command */
673:             fdc_wait_interrupt = FDC_WRITE;
674:             fdc_out(FDC_WRITE);
675:             fdc_out((head << 2) | current_fdd);
676:             fdc_out(cyl);
677:             fdc_out(head);
678:             fdc_out(sector);
679:             fdc_out(2); /* sector size is 512 bytes */
680:             fdc_out(current_fdd_type->spt);
681:             fdc_out(current_fdd_type->gpl1);
682:             fdc_out(0xFF); /* sector size is 512 bytes */
683:
684:             if(need_reset) {
685:                 printk("WARNING: %s(): fd%d: needs reset on %s device %d
,%d during write operation.\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(de
v), MINOR(dev));
686:                 continue;
687:             }
688:             creq.fn = fdc_timer;
689:             creq.arg = FDC_TR_DEFAULT;
690:             add_callout(&creq, WAIT_FDC);
691:             sleep(&irq_floppy, PROC_UNINTERRUPTIBLE);
692:             if(fdc_timeout) {
693:                 need_reset = 1;
694:                 printk("WARNING: %s(): fd%d: timeout on %s device %d,%d.
\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(dev), MINOR(dev));
695:                 continue;
696:             }

```

## drivers/block/floppy.c

Page 12/14

```

697:         del_callout(&creq);
698:         fdc_get_results();
699:         if(fdc_results[ST1] & ST1_NW) {
700:             unlock_resource(&floppy_resource);
701:             fdc_motor_off();
702:             return -EROFS;
703:         }
704:         if(fdc_results[ST0] & (ST0_IC | ST0_UC | ST0_NR)) {
705:             need_reset = 1;
706:             continue;
707:         }
708:         break;
709:     }
710:
711:     if(retries >= MAX_FDC_ERR) {
712:         printk("WARNING: %s(): fd%d: error on %s device %d,%d during wri
te operation,\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(dev), MINOR(dev)
);
713:         printk("\tblock=%d, sector=%d, cylinder/head=%d/%d\n", block, se
ctor, cyl, head);
714:         unlock_resource(&floppy_resource);
715:         fdc_motor_off();
716:         return -EIO;
717:     }
718:
719:     fdc_motor_off();
720:     sectors_written = (fdc_results[ST_CYL] - cyl) * (current_fdd_type->heads
* current_fdd_type->spt);
721:     sectors_written += (fdc_results[ST_HEAD] - head) * current_fdd_type->spt
;
722:     sectors_written += fdc_results[ST_SECTOR] - sector;
723:     if(sectors_written * BPS != blksize) {
724:         printk("WARNING: %s(): fd%d: write error on %s device %d,%d (%d
sectors written).\n", __FUNCTION__, current_fdd, floppy_device.name, MAJOR(dev), MINOR(
dev), sectors_written);
725:         printk("\tblock=%d, sector=%d, cylinder/head=%d/%d\n", block, se
ctor, cyl, head);
726:         unlock_resource(&floppy_resource);
727:         fdc_motor_off();
728:         return -EIO;
729:     }
730:
731:     unlock_resource(&floppy_resource);
732:     return sectors_written * BPS;
733: }
734:
735: int fdc_ioctl(struct inode *i, int cmd, unsigned long int arg)
736: {
737:     unsigned char minor;
738:     struct hd_geometry *geom;
739:     int errno;
740:
741:     minor = MINOR(i->rdev);
742:     if(!TEST_MINOR(floppy_device.minors, minor)) {
743:         return -ENXIO;
744:     }
745:
746:     lock_resource(&floppy_resource);
747:     set_current_fdd_type(minor);
748:     unlock_resource(&floppy_resource);
749:
750:     switch(cmd) {
751:         case HDIO_GETGEO:
752:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
izeof(struct hd_geometry)))) {
753:                 return errno;
754:             }

```

## drivers/block/floppy.c

Page 13/14

```

755:             geom = (struct hd_geometry *)arg;
756:             geom->heads = current_fdd_type->heads;
757:             geom->sectors = current_fdd_type->spt;
758:             geom->cylinders = current_fdd_type->tracks;
759:             geom->start = 0;
760:             break;
761:         case BLKRRPART:
762:             break;
763:         case BLKGETSIZE:
764:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
izeof(unsigned int)))) {
765:                 return errno;
766:             }
767:             *(int *)arg = fdd_sizes[MINOR(i->rdev)] * 2;
768:             break;
769:         default:
770:             return -EINVAL;
771:     }
772:     return 0;
773: }
774:
775: int fdc_lseek(struct inode *i, __off_t offset)
776: {
777:     unsigned char minor;
778:
779:     minor = MINOR(i->rdev);
780:     if(!TEST_MINOR(floppy_device.minors, minor)) {
781:         return -ENXIO;
782:     }
783:
784:     lock_resource(&floppy_resource);
785:     set_current_fdd_type(minor);
786:     unlock_resource(&floppy_resource);
787:
788:     return offset;
789: }
790:
791: void floppy_init(void)
792: {
793:     short int cmosval, master, slave;
794:
795:     cmosval = cmos_read(CMOS_FDDTYPE);
796:     set_current_fdd_type(0);          /* sets /dev/fd0 by default */
797:
798:     /* the high nibble describes the 'master' floppy drive */
799:     master = cmosval >> 4;
800:
801:     /*
802:      * Some BIOS may return the value 0x05 (for 2.88MB floppy type) which is
803:      * not supported by Fiwix. This prevents from using an unexistent type
804:      * in the fdd_type structure if this happens.
805:      */
806:     if(master > 4) {
807:         master = 4;
808:     }
809:
810:     if(master) {
811:         if(!register_irq(FLOPPY_IRQ, &irq_config_floppy)) {
812:             enable_irq(FLOPPY_IRQ);
813:         }
814:         printk("fd0          0x%04X-0x%04X  %2d    ", FDC_SRA, FDC_CCR, FL
OPPY_IRQ);
815:         printk("%s ", fdd_type[master].name);
816:         fdd_status[0].type = fdd_status[1].type = master;
817:         SET_MINOR(floppy_device.minors, 0);
818:         SET_MINOR(floppy_device.minors, 4);
819:         SET_MINOR(floppy_device.minors, 8);

```

## drivers/block/floppy.c

Page 14/14

```

820:         SET_MINOR(floppy_device.minors, 12);
821:         SET_MINOR(floppy_device.minors, 16);
822:         fdd_sizes[0] = fdd_type[master].sizekb;
823:         fdd_sizes[4] = fdd_type[1].sizekb;
824:         fdd_sizes[8] = fdd_type[2].sizekb;
825:         fdd_sizes[12] = fdd_type[3].sizekb;
826:         fdd_sizes[16] = fdd_type[4].sizekb;
827:         fdc_reset();
828:         fdc_get_chip();
829:     }
830:
831:     /* the low nibble is for the 'slave' floppy drive */
832:     slave = cmosval & 0x0F;
833:     if(slave) {
834:         if(!master) {
835:             if(!register_irq(FLOPPY_IRQ, &irq_config_floppy)) {
836:                 enable_irq(FLOPPY_IRQ);
837:             }
838:         }
839:         printk("fd1      0x%04X-0x%04X  %2d      ", FDC_SRA, FDC_CCR, FL
OPPY_IRQ);
840:         printk("%s ", fdd_type[slave].name);
841:         fdd_status[1].type = slave;
842:         SET_MINOR(floppy_device.minors, 1);
843:         SET_MINOR(floppy_device.minors, 5);
844:         SET_MINOR(floppy_device.minors, 9);
845:         SET_MINOR(floppy_device.minors, 13);
846:         SET_MINOR(floppy_device.minors, 17);
847:         fdd_sizes[1] = fdd_type[slave].sizekb;
848:         fdd_sizes[5] = fdd_type[1].sizekb;
849:         fdd_sizes[9] = fdd_type[2].sizekb;
850:         fdd_sizes[13] = fdd_type[3].sizekb;
851:         fdd_sizes[17] = fdd_type[4].sizekb;
852:         if(!master) {
853:             fdc_get_chip();
854:         } else {
855:             printk("\n");
856:         }
857:     }
858:
859:     if(master || slave) {
860:         need_reset = 1;
861:         dma_init();
862:         if(dma_register(FLOPPY_DMA, floppy_device.name)) {
863:             printk("WARNING: %s(): fd%d: unable to register DMA chan
nel on %s.\n", __FUNCTION__, current_fdd, floppy_device.name);
864:         } else {
865:             if(!register_device(BLK_DEV, &floppy_device)) {
866:                 do_motor_off(current_fdd);
867:             }
868:         }
869:     }
870: }

```

## drivers/block/ide.c

Page 1/14

```

1: /*
2:  * fiwix/drivers/block/ide.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/ide.h>
10: #include <fiwix/ide_hd.h>
11: #include <fiwix/ide_cd.h>
12: #include <fiwix/devices.h>
13: #include <fiwix/sleep.h>
14: #include <fiwix/timer.h>
15: #include <fiwix/sched.h>
16: #include <fiwix/cpu.h>
17: #include <fiwix/pic.h>
18: #include <fiwix/fs.h>
19: #include <fiwix/mm.h>
20: #include <fiwix/errno.h>
21: #include <fiwix/stdio.h>
22: #include <fiwix/string.h>
23:
24: int ide0_need_reset = 0;
25: int ide0_wait_interrupt = 0;
26: int ide0_timeout = 0;
27: int ide1_need_reset = 0;
28: int ide1_wait_interrupt = 0;
29: int ide1_timeout = 0;
30:
31: struct ide ide_table[NR_IDE_CTRL] = {
32:     { IDE_PRIMARY, IDE0_BASE, IDE0_CTRL, IDE0_IRQ,
33:       {
34:         { IDE_MASTER, "hda", IDE0_MAJOR, 0, IDE_MASTER_MSF, NULL
, NULL, NULL, NULL, NULL, { NULL }, {{ NULL }} },
35:         { IDE_SLAVE, "hdb", IDE0_MAJOR, 0, IDE_SLAVE_MSF, NULL,
NULL, NULL, NULL, NULL, { NULL }, {{ NULL }} }
36:       }
37:     },
38:     { IDE_SECONDARY, IDE1_BASE, IDE1_CTRL, IDE1_IRQ,
39:       {
40:         { IDE_MASTER, "hdc", IDE1_MAJOR, 0, IDE_MASTER_MSF, NULL
, NULL, NULL, NULL, NULL, { NULL }, {{ NULL }} },
41:         { IDE_SLAVE, "hdd", IDE1_MAJOR, 0, IDE_SLAVE_MSF, NULL,
NULL, NULL, NULL, NULL, { NULL }, {{ NULL }} }
42:       }
43:     }
44: };
45:
46: static char *ide_ctrl_name[] = { "primary", "secondary" };
47: static char *ide_drv_name[] = { "master", "slave" };
48:
49: static unsigned int ide0_sizes[256];
50: static unsigned int ide1_sizes[256];
51:
52: static struct fs_operations ide_driver_fsop = {
53:     0,
54:     0,
55:
56:     ide_open,
57:     ide_close,
58:     NULL, /* read */
59:     NULL, /* write */
60:     ide_ioctl,
61:     NULL, /* lseek */
62:     NULL, /* readdir */
63:     NULL, /* mmap */

```

## drivers/block/ide.c

Page 2/14

```

64:         NULL,                /* select */
65:
66:         NULL,                /* readlink */
67:         NULL,                /* followlink */
68:         NULL,                /* bmap */
69:         NULL,                /* lockup */
70:         NULL,                /* rmdir */
71:         NULL,                /* link */
72:         NULL,                /* unlink */
73:         NULL,                /* symlink */
74:         NULL,                /* mkdir */
75:         NULL,                /* mknod */
76:         NULL,                /* truncate */
77:         NULL,                /* create */
78:         NULL,                /* rename */
79:
80:         ide_read,
81:         ide_write,
82:
83:         NULL,                /* read_inode */
84:         NULL,                /* write_inode */
85:         NULL,                /* ialloc */
86:         NULL,                /* ifree */
87:         NULL,                /* statfs */
88:         NULL,                /* read_superblock */
89:         NULL,                /* remount_fs */
90:         NULL,                /* write_superblock */
91:         NULL,                /* release_superblock */
92: };
93:
94: static struct device ide0_device = {
95:     "ide0",
96:     IDE0_MAJOR,
97:     { 0, 0, 0, 0, 0, 0, 0, 0 },
98:     0,
99:     &ide0_sizes,
100:    &ide_driver_fsop,
101:    NULL
102: };
103:
104: static struct device idel_device = {
105:     "idel",
106:     IDE1_MAJOR,
107:     { 0, 0, 0, 0, 0, 0, 0, 0 },
108:     0,
109:     &idel_sizes,
110:     &ide_driver_fsop,
111:     NULL
112: };
113:
114: static struct interrupt irq_config_ide0 = { 0, "ide0", &irq_ide0, NULL };
115: static struct interrupt irq_config_idel = { 0, "idel", &irq_idel, NULL };
116:
117: static int ide_identify(struct ide *ide, int drive)
118: {
119:     short int status, *buffer;
120:     struct callout_req creq;
121:
122:     if((status = ide_drvsel(ide, drive, IDE_CHS_MODE, 0))) {
123:         /* some controllers return 0xFF to indicate a non-drive conditio
n */
124:         if(status == 0xFF) {
125:             return status;
126:         }
127:         printk("WARNING: %s(): error on device '%s'.\n", __FUNCTION__, i
de->drive[drive].dev_name);
128:         ide_error(ide, status);

```

## drivers/block/ide.c

Page 3/14

```

129:         return status;
130:     }
131:
132:     outport_b(ide->base + IDE_COMMAND, (ide->drive[drive].flags & DEVICE_IS_
ATAPI) ? ATA_IDENTIFY_PACKET : ATA_IDENTIFY);
133:     if(ide->channel == IDE_PRIMARY) {
134:         ide0_wait_interrupt = ide->base;
135:         creq.fn = ide0_timer;
136:         creq.arg = 0;
137:         add_callout(&creq, WAIT_FOR_IDE);
138:         sleep(&irq_ide0, PROC_UNINTERRUPTIBLE);
139:         if(ide0_timeout) {
140:             status = inport_b(ide->base + IDE_STATUS);
141:             if((status & (IDE_STAT_RDY | IDE_STAT_DRQ)) != (IDE_STAT
_RDY | IDE_STAT_DRQ)) {
142:                 return 1;
143:             }
144:         }
145:         del_callout(&creq);
146:     }
147:     if(ide->channel == IDE_SECONDARY) {
148:         idel_wait_interrupt = ide->base;
149:         creq.fn = idel_timer;
150:         creq.arg = 0;
151:         add_callout(&creq, WAIT_FOR_IDE);
152:         sleep(&irq_idel, PROC_UNINTERRUPTIBLE);
153:         if(idel_timeout) {
154:             status = inport_b(ide->base + IDE_STATUS);
155:             if((status & (IDE_STAT_RDY | IDE_STAT_DRQ)) != (IDE_STAT
_RDY | IDE_STAT_DRQ)) {
156:                 return 1;
157:             }
158:         }
159:         del_callout(&creq);
160:     }
161:
162:     status = inport_b(ide->base + IDE_STATUS);
163:     if((status & (IDE_STAT_RDY | IDE_STAT_DRQ)) != (IDE_STAT_RDY | IDE_STAT
DRQ)) {
164:         return 1;
165:     }
166:
167:     if(!(buffer = (void *)kmalloc())) {
168:         return 1;
169:     }
170:
171:     inport_sw(ide->base + IDE_DATA, (void *)buffer, IDE_HD_SECTSIZE / sizeof
(short int));
172:     memcpy_b(&ide->drive[drive].ident, (void *)buffer, sizeof(struct ide_drv
_ident));
173:     kfree((unsigned int)buffer);
174:
175:     if(ide->drive[drive].ident.gen_config == IDE_SUPPORTS_CFA) {
176:         ide->drive[drive].flags |= DEVICE_IS_CFA;
177:     }
178:
179:     if(ide->drive[drive].flags & DEVICE_IS_ATAPI) {
180:         if((ide->drive[drive].ident.gen_config >> 8) & 0x1F) == ATAPI_I
S_CDROM) {
181:             ide->drive[drive].flags |= DEVICE_IS_CDROM;
182:         }
183:         if(ide->drive[drive].ident.gen_config & 0x3) {
184:             printk("WARNING: %s(): packet size must be 16 bytes!\n")
;
185:         }
186:     }
187:

```



## drivers/block/ide.c

Page 4/14

```

188:      /* only bits 0-7 are relevant */
189:      ide->drive[drive].ident.rw_multiple &= 0xFF;
190:      return 0;
191: }
192:
193: static void get_device_size(struct ide_drv *drive)
194: {
195:     if(drive->ident.capabilities & IDE_HAS_LBA) {
196:         drive->lba_cyls = drive->ident.logic_cyls;
197:         drive->lba_heads = drive->ident.logic_heads;
198:         drive->lba_factor = 0;
199:
200:         while(drive->lba_cyls > 1023) {
201:             if(drive->lba_heads < 255) {
202:                 drive->lba_cyls >>= 1;
203:                 drive->lba_heads <<= 1;
204:             } else {
205:                 break;
206:             }
207:             drive->lba_factor++;
208:         }
209:         drive->nr_sects = drive->ident.tot_sectors | (drive->ident.tot_s
ectors2 << 16);
210:     }
211:
212:     /* some old disk drives (ATA or ATA2) don't specify total sectors */
213:     if(!(drive->ident.capabilities & IDE_HAS_LBA)) {
214:         if(drive->nr_sects == 0) {
215:             drive->nr_sects = drive->ident.logic_cyls * drive->ident
.logic_heads * drive->ident.logic_spt;
216:         }
217:     }
218:
219: }
220:
221: static int get_udma(struct ide *ide, int drive)
222: {
223:     int udma;
224:
225:     if(ide->drive[drive].ident.fields_validity & IDE_HAS_UDMA) {
226:         if((ide->drive[drive].ident.ultradma >> 13) & 1) {
227:             udma = 5;
228:         } else if((ide->drive[drive].ident.ultradma >> 12) & 1) {
229:             udma = 4;
230:         } else if((ide->drive[drive].ident.ultradma >> 11) & 1) {
231:             udma = 3;
232:         } else if((ide->drive[drive].ident.ultradma >> 10) & 1) {
233:             udma = 2;
234:         } else if((ide->drive[drive].ident.ultradma >> 9) & 1) {
235:             udma = 1;
236:         } else {
237:             udma = 0;
238:         }
239:     } else {
240:         udma = -1;
241:     }
242:     return udma;
243: }
244:
245: static void ide_results(struct ide *ide, int drive)
246: {
247:     unsigned int cyl, hds, sect;
248:     __loff_t capacity;
249:     int udma;
250:     int udma_speed[] = { 16, 25, 33, 44, 66, 100 };
251:
252:     cyl = ide->drive[drive].ident.logic_cyls;

```

## drivers/block/ide.c

Page 5/14

```

253:         hds = ide->drive[drive].ident.logic_heads;
254:         sect = ide->drive[drive].ident.logic_spt;
255:
256:         udma = get_udma(ide, drive);
257:         /*
258:          * After knowing if the device is UDMA capable we could choose between
259:          * the PIO transfer mode or the UDMA transfer mode.
260:          * FIXME: Currently only PIO mode is supported.
261:          */
262:
263:         capacity = (__loff_t)ide->drive[drive].nr_sects * BPS;
264:         capacity = capacity / 1024 / 1024;
265:
266:         printk("%s          0x%04X-0x%04X   %2d   ", ide->drive[drive].dev_name, i
de->base, ide->base + IDE_BASE_LEN, ide->irq);
267:         swap_asc_word(ide->drive[drive].ident.model_number, 40);
268:         printk("%s %s ", ide_ctrl_name[ide->channel], ide_drv_name[ide->drive[dr
ive].drive]);
269:
270:         if(!(ide->drive[drive].flags & DEVICE_IS_ATAPI)) {
271:             printk("ATA");
272:         } else {
273:             printk("ATAPI");
274:         }
275:
276:         if(ide->drive[drive].flags & DEVICE_IS_CFA) {
277:             printk(" CFA");
278:         }
279:
280:         if(ide->drive[drive].flags & DEVICE_IS_DISK) {
281:             printk(" DISK drive %dMB\n", (unsigned int)capacity);
282:             printk("                          model=%s\n", ide->drive[
drive].ident.model_number);
283:             if(ide->drive[drive].nr_sects < IDE_MIN_LBA) {
284:                 printk("                          CHS=%d/%d/%d", c
yl, hds, sect);
285:             } else {
286:                 ide->drive[drive].flags |= DEVICE_REQUIRES_LBA;
287:                 printk("                          sectors=%d", ide
->drive[drive].nr_sects);
288:             }
289:             printk(" cache=%dKB", ide->drive[drive].ident.buffer_cache >> 1)
;
290:         }
291:
292:         if(ide->drive[drive].flags & DEVICE_IS_CDROM) {
293:             printk(" CDROM drive\n");
294:             printk("                          model=%s\n", ide->drive[
drive].ident.model_number);
295:             printk("                          cache=%dKB", ide->drive[
drive].ident.buffer_cache >> 1);
296:         }
297:
298:         if(udma >= 0) {
299:             printk(" UDMA%d(%d)", udma, udma_speed[udma]);
300:         }
301:         if(ide->drive[drive].ident.capabilities & IDE_HAS_LBA) {
302:             ide->drive[drive].flags |= DEVICE_REQUIRES_LBA;
303:             printk(" LBA");
304:         }
305:
306:         printk("\n");
307:
308:         if(ide->drive[drive].ident.rw_multiple > 1) {
309:             ide->drive[drive].flags |= DEVICE_HAS_RW_MULTIPLE;
310:         }
311:

```

## drivers/block/ide.c

Page 6/14

```

312:         /*
313:         printk("\n");
314:         printk("%s -> %s\n", ide->drive[drive].dev_name, ide->drive[drive].flags
& DEVICE_IS_ATAPI ? "ATAPI" : "ATA");
315:         printk("general conf = %d (%b) (0x%x)\n", ide->drive[drive].ident.gen_c
onfig, ide->drive[drive].ident.gen_config, ide->drive[drive].ident.gen_config);
316:         printk("logic_cyls = %d (%b)\n", ide->drive[drive].ident.logic_cyls,
ide->drive[drive].ident.logic_cyls);
317:         printk("reserved2 = %d (%b)\n", ide->drive[drive].ident.reserved2, i
de->drive[drive].ident.reserved2);
318:         printk("logic_heads = %d (%b)\n", ide->drive[drive].ident.logic_heads,
ide->drive[drive].ident.logic_heads);
319:         printk("retired4 = %d (%b)\n", ide->drive[drive].ident.retired4, id
e->drive[drive].ident.retired4);
320:         printk("retired5 = %d (%b)\n", ide->drive[drive].ident.retired5, id
e->drive[drive].ident.retired5);
321:         printk("logic_spt = %d (%b)\n", ide->drive[drive].ident.logic_spt, i
de->drive[drive].ident.logic_spt);
322:         printk("retired7 = %d (%b)\n", ide->drive[drive].ident.retired7, id
e->drive[drive].ident.retired7);
323:         printk("retired8 = %d (%b)\n", ide->drive[drive].ident.retired8, id
e->drive[drive].ident.retired8);
324:         printk("retired9 = %d (%b)\n", ide->drive[drive].ident.retired9, id
e->drive[drive].ident.retired9);
325:         printk("serial number = '%s'\n", ide->drive[drive].ident.serial_number);
326:         printk("vendor spec20 = %d (%b)\n", ide->drive[drive].ident.vendor_spec2
0, ide->drive[drive].ident.vendor_spec20);
327:         printk("buffer cache = %d (%b)\n", ide->drive[drive].ident.buffer_cache
, ide->drive[drive].ident.buffer_cache);
328:         printk("vendor spec22 = %d (%b)\n", ide->drive[drive].ident.vendor_spec2
2, ide->drive[drive].ident.vendor_spec22);
329:         printk("firmware rev = '%s'\n", ide->drive[drive].ident.firmware_rev);
330:         printk("model number = '%s'\n", ide->drive[drive].ident.model_number);
331:         printk("rw multiple = %d (%b)\n", ide->drive[drive].ident.rw_multiple,
ide->drive[drive].ident.rw_multiple);
332:         printk("reserved48 = %d (%b)\n", ide->drive[drive].ident.reserved48,
ide->drive[drive].ident.reserved48);
333:         printk("capabilities = %d (%b)\n", ide->drive[drive].ident.capabilities
, ide->drive[drive].ident.capabilities);
334:         printk("reserved50 = %d (%b)\n", ide->drive[drive].ident.reserved50,
ide->drive[drive].ident.reserved50);
335:         printk("pio mode = %d (%b)\n", ide->drive[drive].ident.pio_mode, id
e->drive[drive].ident.pio_mode);
336:         printk("dma mode = %d (%b)\n", ide->drive[drive].ident.dma_mode, id
e->drive[drive].ident.dma_mode);
337:         printk("fields validi = %d (%b)\n", ide->drive[drive].ident.fields_valid
ity, ide->drive[drive].ident.fields_validity);
338:         printk("cur log cyls = %d (%b)\n", ide->drive[drive].ident.cur_log_cyls
, ide->drive[drive].ident.cur_log_cyls);
339:         printk("cur log heads = %d (%b)\n", ide->drive[drive].ident.cur_log_head
s, ide->drive[drive].ident.cur_log_heads);
340:         printk("cur log spt = %d (%b)\n", ide->drive[drive].ident.cur_log_spt,
ide->drive[drive].ident.cur_log_spt);
341:         printk("cur capacity = %d (%b)\n", ide->drive[drive].ident.cur_capacity
| (ide->drive[drive].ident.cur_capacity2 << 16), ide->drive[drive].ident.cur_capacity
| (ide->drive[drive].ident.cur_capacity2 << 16));
342:         printk("mult sect set = %d (%b)\n", ide->drive[drive].ident.mult_sect_se
t, ide->drive[drive].ident.mult_sect_set);
343:         printk("tot sectors = %d (%b)\n", ide->drive[drive].ident.tot_sectors
| (ide->drive[drive].ident.tot_sectors2 << 16), ide->drive[drive].ident.tot_sectors | (
ide->drive[drive].ident.tot_sectors2 << 16));
344:         printk("singleword dma= %d (%b)\n", ide->drive[drive].ident.singleword_d
ma, ide->drive[drive].ident.singleword_dma);
345:         printk("multiword dma = %d (%b)\n", ide->drive[drive].ident.multiword_dm
a, ide->drive[drive].ident.multiword_dma);
346:         printk("adv pio modes = %d (%b)\n", ide->drive[drive].ident.adv_pio_mode
s, ide->drive[drive].ident.adv_pio_modes);

```

## drivers/block/ide.c

Page 7/14

```

347:         printk("min multiword = %d (%b)\n", ide->drive[drive].ident.min_multiwor
d, ide->drive[drive].ident.min_multiword);
348:         printk("rec multiword = %d (%b)\n", ide->drive[drive].ident.rec_multiwor
d, ide->drive[drive].ident.rec_multiword);
349:         printk("min pio wo fc = %d (%b)\n", ide->drive[drive].ident.min_pio_wo_fc
c, ide->drive[drive].ident.min_pio_wo_fc);
350:         printk("min pio w fc = %d (%b)\n", ide->drive[drive].ident.min_pio_w_fc
, ide->drive[drive].ident.min_pio_w_fc);
351:         printk("reserved69 = %d (%b)\n", ide->drive[drive].ident.reserved69,
ide->drive[drive].ident.reserved69);
352:         printk("reserved70 = %d (%b)\n", ide->drive[drive].ident.reserved70,
ide->drive[drive].ident.reserved70);
353:         printk("reserved71 = %d (%b)\n", ide->drive[drive].ident.reserved71,
ide->drive[drive].ident.reserved71);
354:         printk("reserved72 = %d (%b)\n", ide->drive[drive].ident.reserved72,
ide->drive[drive].ident.reserved72);
355:         printk("reserved73 = %d (%b)\n", ide->drive[drive].ident.reserved73,
ide->drive[drive].ident.reserved73);
356:         printk("reserved74 = %d (%b)\n", ide->drive[drive].ident.reserved74,
ide->drive[drive].ident.reserved74);
357:         printk("queue depth = %d (%b)\n", ide->drive[drive].ident.queue_depth,
ide->drive[drive].ident.queue_depth);
358:         printk("reserved76 = %d (%b)\n", ide->drive[drive].ident.reserved76,
ide->drive[drive].ident.reserved76);
359:         printk("reserved77 = %d (%b)\n", ide->drive[drive].ident.reserved77,
ide->drive[drive].ident.reserved77);
360:         printk("reserved78 = %d (%b)\n", ide->drive[drive].ident.reserved78,
ide->drive[drive].ident.reserved78);
361:         printk("reserved79 = %d (%b)\n", ide->drive[drive].ident.reserved79,
ide->drive[drive].ident.reserved79);
362:         printk("major version = %d (%b)\n", ide->drive[drive].ident.majorver, id
e->drive[drive].ident.majorver);
363:         printk("minor version = %d (%b)\n", ide->drive[drive].ident.minorver, id
e->drive[drive].ident.minorver);
364:         printk("cmdset1 = %d (%b)\n", ide->drive[drive].ident.cmdset1, ide
->drive[drive].ident.cmdset1);
365:         printk("cmdset2 = %d (%b)\n", ide->drive[drive].ident.cmdset2, ide
->drive[drive].ident.cmdset2);
366:         printk("cmdsfs_ext = %d (%b)\n", ide->drive[drive].ident.cmdsfs_ext, i
de->drive[drive].ident.cmdsfs_ext);
367:         printk("cmdsfs_enable1 = %d (%b)\n", ide->drive[drive].ident.cmdsfs_enable
1, ide->drive[drive].ident.cmdsfs_enable1);
368:         printk("cmdsfs_enable2 = %d (%b)\n", ide->drive[drive].ident.cmdsfs_enable
2, ide->drive[drive].ident.cmdsfs_enable2);
369:         printk("cmdsfs_default = %d (%b)\n", ide->drive[drive].ident.cmdsfs_defaul
t, ide->drive[drive].ident.cmdsfs_default);
370:         printk("ultra dma = %d (%b)\n", ide->drive[drive].ident.ultradma, id
e->drive[drive].ident.ultradma);
371:         printk("reserved89 = %d (%b)\n", ide->drive[drive].ident.reserved89,
ide->drive[drive].ident.reserved89);
372:         printk("reserved90 = %d (%b)\n", ide->drive[drive].ident.reserved90,
ide->drive[drive].ident.reserved90);
373:         printk("current apm = %d (%b)\n", ide->drive[drive].ident.curapm, ide
->drive[drive].ident.curapm);
374:         */
375:     }
376:
377: void irq_ide0(int num, struct sigcontext *sc)
378: {
379:     if(!ide0_wait_interrupt) {
380:         printk("WARNING: %s(): unexpected interrupt!\n", __FUNCTION__);
381:         ide0_need_reset = 1;
382:     } else {
383:         ide0_timeout = ide0_wait_interrupt = 0;
384:         wakeup(&irq_ide0);
385:     }
386: }

```

## drivers/block/ide.c

Page 8/14

```

387:
388: void irq_idel(int num, struct sigcontext *sc)
389: {
390:     if(!idel_wait_interrupt) {
391:         printk("WARNING: %s(): unexpected interrupt!\n", __FUNCTION__);
392:         idel_need_reset = 1;
393:     } else {
394:         idel_timeout = idel_wait_interrupt = 0;
395:         wakeup(&irq_idel);
396:     }
397: }
398:
399: void ide0_timer(unsigned int arg)
400: {
401:     ide0_timeout = 1;
402:     ide0_wait_interrupt = 0;
403:     wakeup(&irq_ide0);
404: }
405:
406: void idel_timer(unsigned int arg)
407: {
408:     idel_timeout = 1;
409:     idel_wait_interrupt = 0;
410:     wakeup(&irq_idel);
411: }
412:
413: void ide_error(struct ide *ide, int status)
414: {
415:     int error;
416:
417:     if(status & IDE_STAT_ERR) {
418:         error = inport_b(ide->base + IDE_ERROR);
419:         if(error) {
420:             printk("error=0x%x [", error);
421:         }
422:         if(error & IDE_ERR_AMNF) {
423:             printk("address mark not found, ");
424:         }
425:         if(error & IDE_ERR_TKONF) {
426:             printk("track 0 not found (no media) or media error, ");
427:         }
428:         if(error & IDE_ERR_ABRT) {
429:             printk("command aborted, ");
430:         }
431:         if(error & IDE_ERR_MCR) {
432:             printk("media change requested, ");
433:         }
434:         if(error & IDE_ERR_IDNF) {
435:             printk("id mark not found, ");
436:         }
437:         if(error & IDE_ERR_MC) {
438:             printk("media changer, ");
439:         }
440:         if(error & IDE_ERR_UNC) {
441:             printk("uncorrectable data, ");
442:         }
443:         if(error & IDE_ERR_BBK) {
444:             printk("bad block, ");
445:         }
446:         printk("]");
447:     }
448:     if(status & IDE_STAT_DWF) {
449:         printk("device fault, ");
450:     }
451:     if(status & IDE_STAT_BSY) {
452:         printk("device busy, ");
453:     }

```

## drivers/block/ide.c

Page 9/14

```

454:         printk("\n");
455:     }
456:
457: void ide_delay(void)
458: {
459:     int n;
460:
461:     for(n = 0; n < 10000; n++) {
462:         NOP();
463:     }
464: }
465:
466: void ide_wait400ns(struct ide *ide)
467: {
468:     int n;
469:
470:     /* wait 400ns */
471:     for(n = 0; n < 4; n++) {
472:         inport_b(ide->ctrl + IDE_ALT_STATUS);
473:     }
474: }
475:
476: int ide_ready(struct ide *ide)
477: {
478:     int n, retries, status;
479:
480:     SET_IDE_RDY_RETR(retries);
481:     for(n = 0; n < retries; n++) {
482:         status = inport_b(ide->ctrl + IDE_ALT_STATUS);
483:         if(!(status & IDE_STAT_BSY)) {
484:             return 0;
485:         }
486:         ide_delay();
487:     }
488:
489:     inport_b(ide->base + IDE_STATUS);          /* clear any pending interrupt */
490:
491:     return status;
492: }
493: int ide_drvsel(struct ide *ide, int drive, int mode, unsigned char lba24_head)
494: {
495:     int n;
496:     int status;
497:
498:     for(n = 0; n < MAX_IDE_ERR; n++) {
499:         if((status = ide_ready(ide))) {
500:             continue;
501:         }
502:         break;
503:     }
504:     if(status) {
505:         return status;
506:     }
507:
508:     outport_b(ide->base + IDE_DRVHD, (mode + (drive << 4)) | lba24_head);
509:     ide_wait400ns(ide);
510:
511:     for(n = 0; n < MAX_IDE_ERR; n++) {
512:         if((status = ide_ready(ide))) {
513:             continue;
514:         }
515:         break;
516:     }
517:     return status;
518: }
519:

```

## drivers/block/ide.c

Page 10/14

```

520: int ide_softreset(struct ide *ide)
521: {
522:     int error;
523:
524:     error = 0;
525:
526:     outport_b(ide->base + IDE_DRVHD, IDE_CHS_MODE);
527:     ide_delay();
528:
529:     outport_b(ide->ctrl + IDE_DEV_CTRL, IDE_DEVCTR_SRST | IDE_DEVCTR_NIEN);
530:     ide_delay();
531:     outport_b(ide->ctrl + IDE_DEV_CTRL, 0);
532:     ide_delay();
533:
534:     outport_b(ide->base + IDE_DRVHD, IDE_CHS_MODE);
535:     ide_delay();
536:     if(ide_ready(ide)) {
537:         printk("WARNING: %s(): reset error on IDE(%d:0).\n", __FUNCTION_
_, ide->channel);
538:         error = 1;
539:     } else {
540:         /* device is disk by default */
541:         ide->drive[IDE_MASTER].flags |= DEVICE_IS_DISK;
542:
543:         /* check if it's an ATAPI device */
544:         if(inport_b(ide->base + IDE_SECCNT) == 1 && inport_b(ide->base +
IDE_SECCNT) == 1) {
545:             if(inport_b(ide->base + IDE_LCYL) == 0x14 && inport_b(id
e->base + IDE_HCYL) == 0xEB) {
546:                 ide->drive[IDE_MASTER].flags &= ~DEVICE_IS_DISK;
547:                 ide->drive[IDE_MASTER].flags |= DEVICE_IS_ATAPI;
548:             }
549:         }
550:     }
551:
552:     outport_b(ide->base + IDE_DRVHD, IDE_CHS_MODE + (1 << 4));
553:     ide_delay();
554:     if(ide_ready(ide)) {
555:         printk("WARNING: %s(): reset error on IDE(%d:1).\n", __FUNCTION_
_, ide->channel);
556:         outport_b(ide->base + IDE_DRVHD, IDE_CHS_MODE);
557:         ide_delay();
558:         ide_ready(ide);
559:         error |= (1 << 4);
560:     }
561:
562:     outport_b(ide->ctrl + IDE_DEV_CTRL, 0);
563:     ide_delay();
564:     if(error > 1) {
565:         return error;
566:     }
567:
568:     /* device is disk by default */
569:     ide->drive[IDE_SLAVE].flags |= DEVICE_IS_DISK;
570:
571:     /* check if it's an ATAPI device */
572:     if(inport_b(ide->base + IDE_SECCNT) == 1 && inport_b(ide->base + IDE_SEC
NUM) == 1) {
573:         if(inport_b(ide->base + IDE_LCYL) == 0x14 && inport_b(ide->base
+ IDE_HCYL) == 0xEB) {
574:             ide->drive[IDE_SLAVE].flags &= ~DEVICE_IS_DISK;
575:             ide->drive[IDE_SLAVE].flags |= DEVICE_IS_ATAPI;
576:         }
577:     }
578:
579:     return error;
580: }

```

## drivers/block/ide.c

Page 11/14

```

581:
582: struct ide * get_ide_controller(__dev_t dev)
583: {
584:     int controller;
585:
586:     if(MAJOR(dev) == IDE0_MAJOR) {
587:         controller = IDE_PRIMARY;
588:     } else {
589:         if(MAJOR(dev) == IDE1_MAJOR) {
590:             controller = IDE_SECONDARY;
591:         } else {
592:             return NULL;
593:         }
594:     }
595:     return &ide_table[controller];
596: }
597:
598: int get_ide_drive(__dev_t dev)
599: {
600:     int drive;
601:
602:     drive = MINOR(dev);
603:     if(drive) {
604:         if(drive & (1 << IDE_SLAVE_MSF)) {
605:             drive = IDE_SLAVE;
606:         } else {
607:             drive = IDE_MASTER;
608:         }
609:     }
610:     return drive;
611: }
612:
613: int ide_open(struct inode *i, struct fd *fd_table)
614: {
615:     int drive;
616:     struct ide *ide;
617:     struct device *d;
618:
619:     if(!(ide = get_ide_controller(i->rdev))) {
620:         return -EINVAL;
621:     }
622:
623:     if(!(d = get_device(BLK_DEV, i->rdev))) {
624:         return -ENXIO;
625:     }
626:
627:     drive = get_ide_drive(i->rdev);
628:     if(ide->drive[drive].fsop && ide->drive[drive].fsop->open) {
629:         return ide->drive[drive].fsop->open(i, fd_table);
630:     }
631:     return -EINVAL;
632: }
633:
634: int ide_close(struct inode *i, struct fd *fd_table)
635: {
636:     int drive;
637:     struct ide *ide;
638:     struct device *d;
639:
640:     if(!(ide = get_ide_controller(i->rdev))) {
641:         return -EINVAL;
642:     }
643:
644:     if(!(d = get_device(BLK_DEV, i->rdev))) {
645:         return -ENXIO;
646:     }
647:

```



## drivers/block/ide.c

Page 12/14

```

648:         drive = get_ide_drive(i->rdev);
649:         if(ide->drive[drive].fsop && ide->drive[drive].fsop->close) {
650:             return ide->drive[drive].fsop->close(i, fd_table);
651:         }
652:         return -EINVAL;
653:     }
654: }
655: int ide_read(__dev_t dev, __blk_t block, char *buffer, int blksize)
656: {
657:     int drive;
658:     struct ide *ide;
659:     struct device *d;
660:
661:     if(!(ide = get_ide_controller(dev))) {
662:         printk("%s(): no ide controller!\n", __FUNCTION__);
663:         return -EINVAL;
664:     }
665:
666:     if(!(d = get_device(BLK_DEV, dev))) {
667:         return -ENXIO;
668:     }
669:
670:     drive = get_ide_drive(dev);
671:     if(ide->drive[drive].fsop && ide->drive[drive].fsop->read_block) {
672:         return ide->drive[drive].fsop->read_block(dev, block, buffer, bl
ksize);
673:     }
674:     printk("WARNING: %s(): device %d,%d does not have the read_block() metho
d!\n", __FUNCTION__, MAJOR(dev), MINOR(dev));
675:     return -EINVAL;
676: }
677:
678: int ide_write(__dev_t dev, __blk_t block, char *buffer, int blksize)
679: {
680:     int drive;
681:     struct ide *ide;
682:     struct device *d;
683:
684:     if(!(ide = get_ide_controller(dev))) {
685:         printk("%s(): no ide controller!\n", __FUNCTION__);
686:         return -EINVAL;
687:     }
688:
689:     if(!(d = get_device(BLK_DEV, dev))) {
690:         return -ENXIO;
691:     }
692:
693:     drive = get_ide_drive(dev);
694:     if(ide->drive[drive].fsop && ide->drive[drive].fsop->write_block) {
695:         return ide->drive[drive].fsop->write_block(dev, block, buffer, b
lksize);
696:     }
697:     printk("WARNING: %s(): device %d,%d does not have the write_block() meth
od!\n", __FUNCTION__, MAJOR(dev), MINOR(dev));
698:     return -EINVAL;
699: }
700:
701: int ide_ioctl(struct inode *i, int cmd, unsigned long int arg)
702: {
703:     int drive;
704:     struct ide *ide;
705:     struct device *d;
706:
707:     if(!(ide = get_ide_controller(i->rdev))) {
708:         return -EINVAL;
709:     }
710:

```

## drivers/block/ide.c

Page 13/14

```

711:         if(!(d = get_device(BLK_DEV, i->rdev))) {
712:             return -ENXIO;
713:         }
714:
715:         drive = get_ide_drive(i->rdev);
716:         if(ide->drive[drive].fsop && ide->drive[drive].fsop->ioc1) {
717:             return ide->drive[drive].fsop->ioc1(i, cmd, arg);
718:         }
719:         return -EINVAL;
720:     }
721:
722: void ide_init(void)
723: {
724:     int devices, errno;
725:     struct ide *ide;
726:
727:     if(!register_irq(IDE0_IRQ, &irq_config_ide0)) {
728:         enable_irq(IDE0_IRQ);
729:     }
730:     devices = 0;
731:
732:     ide = &ide_table[IDE_PRIMARY];
733:     errno = ide_softreset(ide);
734:     if(!(errno & 1)) {
735:         if(!ide_identify(ide, IDE_MASTER)) {
736:             get_device_size(&ide->drive[IDE_MASTER]);
737:             ide_results(ide, IDE_MASTER);
738:             SET_MINOR(ide0_device.minors, 0);
739:             register_device(BLK_DEV, &ide0_device);
740:             if(ide->drive[IDE_MASTER].flags & DEVICE_IS_DISK) {
741:                 if(!ide_hd_init(ide, IDE_MASTER)) {
742:                     devices++;
743:                 }
744:             }
745:             if(ide->drive[IDE_MASTER].flags & DEVICE_IS_CDROM) {
746:                 if(!ide_cd_init(ide, IDE_MASTER)) {
747:                     devices++;
748:                 }
749:             }
750:         }
751:     }
752:     if(!(errno & 0x10)) {
753:         if(!ide_identify(ide, IDE_SLAVE)) {
754:             get_device_size(&ide->drive[IDE_SLAVE]);
755:             ide_results(ide, IDE_SLAVE);
756:             SET_MINOR(ide0_device.minors, 1 << IDE_SLAVE_MSF);
757:             if(!devices) {
758:                 register_device(BLK_DEV, &ide0_device);
759:             }
760:             if(ide->drive[IDE_SLAVE].flags & DEVICE_IS_DISK) {
761:                 if(!ide_hd_init(ide, IDE_SLAVE)) {
762:                     devices++;
763:                 }
764:             }
765:             if(ide->drive[IDE_SLAVE].flags & DEVICE_IS_CDROM) {
766:                 if(!ide_cd_init(ide, IDE_SLAVE)) {
767:                     devices++;
768:                 }
769:             }
770:         }
771:     }
772:     if(!devices) {
773:         disable_irq(IDE0_IRQ);
774:         unregister_irq(IDE0_IRQ, &irq_config_ide0);
775:     }
776:
777:     if(!register_irq(IDE1_IRQ, &irq_config_ide1)) {

```

## drivers/block/ide.c

Page 14/14

```
778:         enable_irq(IDE1_IRQ);
779:     }
780:     devices = 0;
781:     ide = &ide_table[IDE_SECONDARY];
782:     errno = ide_softreset(ide);
783:     if(!(errno & 1)) {
784:         if(!(ide_identify(ide, IDE_MASTER))) {
785:             get_device_size(&ide->drive[IDE_MASTER]);
786:             ide_results(ide, IDE_MASTER);
787:             SET_MINOR(idel_device.minors, 0);
788:             register_device(BLK_DEV, &idel_device);
789:             if(ide->drive[IDE_MASTER].flags & DEVICE_IS_DISK) {
790:                 if(!ide_hd_init(ide, IDE_MASTER)) {
791:                     devices++;
792:                 }
793:             }
794:             if(ide->drive[IDE_MASTER].flags & DEVICE_IS_CDROM) {
795:                 if(!ide_cd_init(ide, IDE_MASTER)) {
796:                     devices++;
797:                 }
798:             }
799:         }
800:     }
801:     if(!(errno & 0x10)) {
802:         if(!(ide_identify(ide, IDE_SLAVE))) {
803:             get_device_size(&ide->drive[IDE_SLAVE]);
804:             ide_results(ide, IDE_SLAVE);
805:             SET_MINOR(idel_device.minors, 1 << IDE_SLAVE_MSF);
806:             if(!devices) {
807:                 register_device(BLK_DEV, &idel_device);
808:             }
809:             if(ide->drive[IDE_SLAVE].flags & DEVICE_IS_DISK) {
810:                 if(!ide_hd_init(ide, IDE_SLAVE)) {
811:                     devices++;
812:                 }
813:             }
814:             if(ide->drive[IDE_SLAVE].flags & DEVICE_IS_CDROM) {
815:                 if(!ide_cd_init(ide, IDE_SLAVE)) {
816:                     devices++;
817:                 }
818:             }
819:         }
820:     }
821:     if(!devices) {
822:         disable_irq(IDE1_IRQ);
823:         unregister_irq(IDE1_IRQ, &irq_config_idel1);
824:     }
825: }
```

## drivers/block/ide\_cd.c

Page 1/9

```

1: /*
2:  * fiwix/drivers/block/ide_cd.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/buffer.h>
10: #include <fiwix/ide.h>
11: #include <fiwix/ide_cd.h>
12: #include <fiwix/ioctl.h>
13: #include <fiwix/devices.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/timer.h>
16: #include <fiwix/sched.h>
17: #include <fiwix/cpu.h>
18: #include <fiwix/fs.h>
19: #include <fiwix/part.h>
20: #include <fiwix/process.h>
21: #include <fiwix/mm.h>
22: #include <fiwix/errno.h>
23: #include <fiwix/stdio.h>
24: #include <fiwix/string.h>
25:
26: /* default size of 1GB is enough to read a whole CDROM */
27: #define CDROM_DEFAULT_SIZE      (1024 * 1024)    /* in KBs */
28:
29: static struct resource ide_cd_resource = { NULL, NULL };
30:
31: static struct fs_operations ide_cd_driver_fsop = {
32:     0,
33:     0,
34:
35:     ide_cd_open,
36:     ide_cd_close,
37:     NULL,                /* read */
38:     NULL,                /* write */
39:     ide_cd_ioctl,
40:     NULL,                /* lseek */
41:     NULL,                /* readdir */
42:     NULL,                /* mmap */
43:     NULL,                /* select */
44:
45:     NULL,                /* readlink */
46:     NULL,                /* followlink */
47:     NULL,                /* bmap */
48:     NULL,                /* lockup */
49:     NULL,                /* rmdir */
50:     NULL,                /* link */
51:     NULL,                /* unlink */
52:     NULL,                /* symlink */
53:     NULL,                /* mkdir */
54:     NULL,                /* mknod */
55:     NULL,                /* truncate */
56:     NULL,                /* create */
57:     NULL,                /* rename */
58:
59:     ide_cd_read,
60:     NULL,                /* write_block */
61:
62:     NULL,                /* read_inode */
63:     NULL,                /* write_inode */
64:     NULL,                /* ialloc */
65:     NULL,                /* ifree */
66:     NULL,                /* statfs */
67:     NULL,                /* read_superblock */

```

## drivers/block/ide\_cd.c

Page 2/9

```

68:         NULL,                /* remount_fs */
69:         NULL,                /* write_superblock */
70:         NULL,                /* release_superblock */
71: };
72:
73: static char *sense_key_err[] = {
74:     "NO SENSE",
75:     "RECOVERED ERROR",
76:     "NOT READY",
77:     "MEDIUM ERROR",
78:     "HARDWARE ERROR",
79:     "ILLEGAL REQUEST",
80:     "UNIT ATTENTION",
81:     "DATA PROTECT",
82:     "RESERVED",
83:     "RESERVED",
84:     "RESERVED",
85:     "ABORTED COMMAND",
86:     "MISCOMPARE",
87:     "RESERVED"
88: };
89:
90: enum {
91:     RS_NO_SENSE,
92:     RS_RECOVERED_ERROR,
93:     RS_NOT_READY,
94:     RS_MEDIUM_ERROR,
95:     RS_HARDWARE_ERROR,
96:     RS_ILLEGAL_REQUEST,
97:     RS_UNIT_ATTENTION,
98:     RS_DATA_PROTECT,
99:     RS_RESERVED1,
100:    RS_RESERVED2,
101:    RS_RESERVED3,
102:    RS_ABORTED_COMMAND,
103:    RS_MISCOMPARE,
104:    RS_RESERVED4
105: };
106:
107: static int send_packet_command(unsigned char *pkt, struct ide *ide, int drive, i
nt blksize)
108: {
109:     int n, retries, status;
110:
111:     outport_b(ide->ctrl + IDE_DEV_CTRL, 0);
112:     ide_delay();
113:     outport_b(ide->base + IDE_DRVHD, IDE_CHS_MODE);
114:     ide_delay();
115:     if(ide_drvsel(ide, drive, IDE_CHS_MODE, 0)) {
116:         printk("WARNING: %s(): %s: drive not ready to receive PACKET com
mand.\n", __FUNCTION__, ide->drive[drive].dev_name);
117:         return 1;
118:     }
119:
120:     CLI();
121:     outport_b(ide->base + IDE_FEATURES, 0);
122:     outport_b(ide->base + IDE_SECCNT, 0);
123:     outport_b(ide->base + IDE_SECNUM, 0);
124:     outport_b(ide->base + IDE_LCYL, blksize & 0xFF);
125:     outport_b(ide->base + IDE_HCYL, blksize >> 8);
126:     outport_b(ide->base + IDE_DRVHD, drive << 4);
127:     outport_b(ide->base + IDE_COMMAND, ATA_PACKET);
128:     ide_wait400ns(ide);
129:
130: /*
131:  * NOTE: Some devices prior to ATA/ATAPI-4 assert INTRQ if enabled at this
132:  * point. See IDENTIFY PACKET DEVICE, word 0, bits 5-6 to determine if an

```

## drivers/block/ide\_cd.c

Page 3/9

```

133:  * interrupt will occur.
134:  */
135:      SET_IDE_RDY_RETR(retries);
136:
137:      for(n = 0; n < retries; n++) {
138:          status = inport_b(ide->base + IDE_STATUS);
139:          if((status & (IDE_STAT_DRQ | IDE_STAT_BSY)) == IDE_STAT_DRQ) {
140:              break;
141:          }
142:          ide_delay();
143:      }
144:      if(n >= retries) {
145:          printk("WARNING: %s(): %s: drive not ready to receive command pa
cket (retries = %d).\n", __FUNCTION__, ide->drive[drive].dev_name, n);
146:          return 1;
147:      }
148:
149:      outport_sw(ide->base + IDE_DATA, pkt, 12 / sizeof(short int));
150:      return 0;
151: }
152:
153: static int atapi_read_data(__dev_t dev, char *data, struct ide *ide, int blksize
, int offset)
154: {
155:     int errno, status;
156:     char *buffer;
157:     int retries, bytes;
158:     struct callout_req creq;
159:
160:     for(retries = 0; retries < MAX_IDE_ERR; retries++) {
161:         if(ide->channel == IDE_PRIMARY) {
162:             ide0_wait_interrupt = ide->base;
163:             creq.fn = ide0_timer;
164:             creq.arg = 0;
165:             add_callout(&creq, WAIT_FOR_IDE);
166:             sleep(&irq_ide0, PROC_UNINTERRUPTIBLE);
167:             if(ide0_timeout) {
168:                 status = inport_b(ide->base + IDE_STATUS);
169:                 if((status & (IDE_STAT_RDY | IDE_STAT_DRQ)) != (
IDE_STAT_RDY | IDE_STAT_DRQ)) {
170:                     continue;
171:                 }
172:             }
173:             del_callout(&creq);
174:         }
175:         if(ide->channel == IDE_SECONDARY) {
176:             idel_wait_interrupt = ide->base;
177:             creq.fn = idel_timer;
178:             creq.arg = 0;
179:             add_callout(&creq, WAIT_FOR_IDE);
180:             sleep(&irq_idel, PROC_UNINTERRUPTIBLE);
181:             if(idel_timeout) {
182:                 status = inport_b(ide->base + IDE_STATUS);
183:                 if((status & (IDE_STAT_RDY | IDE_STAT_DRQ)) != (
IDE_STAT_RDY | IDE_STAT_DRQ)) {
184:                     continue;
185:                 }
186:             }
187:             del_callout(&creq);
188:         }
189:         status = inport_b(ide->base + IDE_STATUS);
190:         if(status & IDE_STAT_ERR) {
191:             continue;
192:         }
193:
194:         if((status & (IDE_STAT_DRQ | IDE_STAT_BSY)) == 0) {
195:             break;

```

## drivers/block/ide\_cd.c

Page 4/9

```

196:         }
197:
198:         bytes = (inport_b(ide->base + IDE_HCYL) << 8) + inport_b(ide->ba
se + IDE_LCYL);
199:         if(!bytes || bytes > blksize) {
200:             break;
201:         }
202:
203:         bytes = MAX(bytes, IDE_CD_SECTSIZE);    /* read more than 2048 b
ytes is not supported */
204:         buffer = data + offset;
205:         inport_sw(ide->base + IDE_DATA, (void *)buffer, bytes / sizeof(s
hort int));
206:     }
207:
208:     if(status & IDE_STAT_ERR) {
209:         errno = inport_b(ide->base + IDE_ERROR);
210:         printk("WARNING: %s(): error on cdrom device %d,%d, status=0x%x
error=0x%x,\n", __FUNCTION__, MAJOR(dev), MINOR(dev), status, errno);
211:         return 1;
212:     }
213:
214:     if(retries >= MAX_IDE_ERR) {
215:         printk("WARNING: %s(): timeout on cdrom device %d,%d, status=0x%
x.\n", __FUNCTION__, MAJOR(dev), MINOR(dev), status);
216:         /* a reset may be required at this moment */
217:         return 1;
218:     }
219:     return 0;
220: }
221:
222: static int atapi_cmd_testunit(struct ide *ide, int drive)
223: {
224:     unsigned char pkt[12];
225:
226:     pkt[0] = ATAPI_TEST_UNIT;
227:     pkt[1] = NULL;
228:     pkt[2] = NULL;
229:     pkt[3] = NULL;
230:     pkt[4] = NULL;
231:     pkt[5] = NULL;
232:     pkt[6] = NULL;
233:     pkt[7] = NULL;
234:     pkt[8] = NULL;
235:     pkt[9] = NULL;
236:     pkt[10] = NULL;
237:     pkt[11] = NULL;
238:     return send_packet_command(pkt, ide, drive, 0);
239: }
240:
241: static int atapi_cmd_reqsense(struct ide *ide, int drive)
242: {
243:     unsigned char pkt[12];
244:
245:     pkt[0] = ATAPI_REQUEST_SENSE;
246:     pkt[1] = NULL;
247:     pkt[2] = NULL;
248:     pkt[3] = NULL;
249:     pkt[4] = 252;    /* this command can send up to 252 bytes */
250:     pkt[5] = NULL;
251:     pkt[6] = NULL;
252:     pkt[7] = NULL;
253:     pkt[8] = NULL;
254:     pkt[9] = NULL;
255:     pkt[10] = NULL;
256:     pkt[11] = NULL;
257:     return send_packet_command(pkt, ide, drive, 0);

```

## drivers/block/ide\_cd.c

Page 5/9

```

258: }
259:
260: static int atapi_cmd_startstop(int action, struct ide *ide, int drive)
261: {
262:     unsigned char pkt[12];
263:
264:     pkt[0] = ATAPI_START_STOP;
265:     pkt[1] = NULL;
266:     pkt[2] = NULL;
267:     pkt[3] = NULL;
268:     pkt[4] = action;
269:     pkt[5] = NULL;
270:     pkt[6] = NULL;
271:     pkt[7] = NULL;
272:     pkt[8] = NULL;
273:     pkt[9] = NULL;
274:     pkt[10] = NULL;
275:     pkt[11] = NULL;
276:     return send_packet_command(pkt, ide, drive, 0);
277: }
278:
279: static int atapi_cmd_mediumrm(int action, struct ide *ide, int drive)
280: {
281:     unsigned char pkt[12];
282:
283:     pkt[0] = ATAPI_MEDIUM_REMOVAL;
284:     pkt[1] = NULL;
285:     pkt[2] = NULL;
286:     pkt[3] = NULL;
287:     pkt[4] = action;
288:     pkt[5] = NULL;
289:     pkt[6] = NULL;
290:     pkt[7] = NULL;
291:     pkt[8] = NULL;
292:     pkt[9] = NULL;
293:     pkt[10] = NULL;
294:     pkt[11] = NULL;
295:     return send_packet_command(pkt, ide, drive, 0);
296: }
297:
298: static int request_sense(char *buffer, __dev_t dev, struct ide *ide, int drive)
299: {
300:     int errcode;
301:     int sense_key, sense_asc;
302:
303:     errcode = inport_b(ide->base + IDE_ERROR);
304:     sense_key = (errcode & 0xF0) >> 4;
305:     printk("\tSense Key code indicates a '%s' condition.\n", sense_key_err[sense_key & 0xF]);
306:     errcode = atapi_cmd_reqsense(ide, drive);
307:     printk("reqsense() returned %d\n", errcode);
308:     errcode = atapi_read_data(dev, buffer, ide, BLKSIZE_2K, 0);
309:     printk("atapi_read_data() returned %d\n", errcode);
310:     errcode = (int)(buffer[0] & 0x7F);
311:     sense_key = (int)(buffer[2] & 0xF);
312:     sense_asc = (int)(buffer[12] & 0xFF);
313:     printk("errcode = %x\n", errcode);
314:     printk("sense_key = %x\n", sense_key);
315:     printk("sense_asc = %x\n", sense_asc);
316:     return errcode;
317: }
318:
319: void ide_cd_timer(unsigned int arg)
320: {
321:     wakeup(&ide_cd_open);
322: }
323:

```



## drivers/block/ide\_cd.c

Page 6/9

```

324: int ide_cd_open(struct inode *i, struct fd *fd_table)
325: {
326:     int minor;
327:     int drive;
328:     char *buffer;
329:     int errcode;
330:     int sense_key, sense_asc;
331:     int retries;
332:     struct ide *ide;
333:
334:     if(!(ide = get_ide_controller(i->rdev))) {
335:         return -EINVAL;
336:     }
337:
338:     minor = MINOR(i->rdev);
339:     drive = get_ide_drive(i->rdev);
340:     if(drive) {
341:         minor &= ~(1 << IDE_SLAVE_MSF);
342:     }
343:
344:     CLI();
345:     lock_resource(&ide_cd_resource);
346:
347:     if(!(buffer = (void *)kmalloc())) {
348:         unlock_resource(&ide_cd_resource);
349:         return -ENOMEM;
350:     }
351:
352:     if((errcode = atapi_cmd_testunit(ide, drive)) {
353:         printk("WARNING: %s(): cdrom device %d,%d is not ready for TEST_
UNIT, error %d.\n", __FUNCTION__, MAJOR(i->rdev), MINOR(i->rdev), errcode);
354:         request_sense(buffer, i->rdev, ide, drive);
355:     }
356:
357:     for(retries = 0; retries < MAX_CD_ERR; retries++) {
358:         if(!(errcode = atapi_cmd_startstop(CD_LOAD, ide, drive))) {
359:             break;
360:         }
361:         printk("WARNING: %s(): cdrom device %d,%d is not ready for CD_LO
AD, error %d.\n", __FUNCTION__, MAJOR(i->rdev), MINOR(i->rdev), errcode);
362:         atapi_read_data(i->rdev, buffer, ide, BLKSIZE_2K, 0);
363:         errcode = request_sense(buffer, i->rdev, ide, drive);
364:         sense_key = (errcode & 0xF0) >> 4;
365:         /* trying to eject on slim drives may lead to an illegal request
*/
366:         if(!sense_key || sense_key == RS_ILLEGAL_REQUEST) {
367:             break;
368:         }
369:         if(errcode == 0x70 || errcode == 0x71) {
370:             sense_key = (int)(buffer[2] & 0xF);
371:             sense_asc = (int)(buffer[12] & 0xFF);
372:             if(sense_key == RS_NOT_READY && sense_asc == ASC_NO_MEDI
UM) {
373:                 kfree((unsigned int)buffer);
374:                 unlock_resource(&ide_cd_resource);
375:                 return -ENOMEDIUM;
376:             }
377:         }
378:     }
379:
380:     if(retries == MAX_CD_ERR) {
381:         if(sense_key == RS_NOT_READY) {
382:             kfree((unsigned int)buffer);
383:             unlock_resource(&ide_cd_resource);
384:             return -ENOMEDIUM;
385:         }
386:     }

```

## drivers/block/ide\_cd.c

Page 7/9

```

387:
388:     if(atapi_cmd_mediumrm(CD_LOCK_MEDIUM, ide, drive)) {
389:         printk("WARNING: %s(): error on cdrom device %d,%d while trying
to lock medium.\n", __FUNCTION__, MAJOR(i->rdev), MINOR(i->rdev), ATAPI_MEDIUM_REMOVAL)
;
390:         request_sense(buffer, i->rdev, ide, drive);
391:     }
392:
393:     /* this line just to catch interrupt */
394:     atapi_read_data(i->rdev, buffer, ide, BLKSIZE_2K, 0);
395:     kfree((unsigned int)buffer);
396:
397:     unlock_resource(&ide_cd_resource);
398:     return 0;
399: }
400:
401: int ide_cd_close(struct inode *i, struct fd *fd_table)
402: {
403:     int drive;
404:     char *buffer;
405:     struct ide *ide;
406:
407:     if(!(ide = get_ide_controller(i->rdev))) {
408:         return -EINVAL;
409:     }
410:
411:     if(!(buffer = (void *)kmalloc())) {
412:         return -ENOMEM;
413:     }
414:
415:     drive = get_ide_drive(i->rdev);
416:
417:     /* FIXME: only if device usage == 0 */
418:     invalidate_buffers(i->rdev);
419:
420:     if(atapi_cmd_mediumrm(CD_UNLOCK_MEDIUM, ide, drive)) {
421:         printk("WARNING: %s(): error on cdrom device %d,%d during 0x%x c
ommand.\n", __FUNCTION__, MAJOR(i->rdev), MINOR(i->rdev), ATAPI_MEDIUM_REMOVAL);
422:     }
423:
424:     /* this line just to catch interrupt */
425:     atapi_read_data(i->rdev, buffer, ide, BLKSIZE_2K, 0);
426:     kfree((unsigned int)buffer);
427:
428:     return 0;
429: }
430:
431: int ide_cd_read(__dev_t dev, __blk_t block, char *buffer, int blksize)
432: {
433:     int drive;
434:     int sectors_to_read;
435:     int n, retries;
436:     unsigned char pkt[12];
437:     struct ide *ide;
438:
439:     if(!(ide = get_ide_controller(dev))) {
440:         return -EINVAL;
441:     }
442:
443:     drive = get_ide_drive(dev);
444:     blksize = BLKSIZE_2K;
445:     sectors_to_read = blksize / IDE_CD_SECTSIZE;
446:
447:     pkt[0] = ATAPI_READ10;
448:     pkt[1] = NULL;
449:     pkt[2] = (block >> 24) & 0xFF;
450:     pkt[3] = (block >> 16) & 0xFF;

```

## drivers/block/ide\_cd.c

Page 8/9

```

451:     pkt[4] = (block >> 8) & 0xFF;
452:     pkt[5] = block & 0xFF;
453:     pkt[6] = NULL;
454:     pkt[7] = (sectors_to_read >> 8) & 0xFF;
455:     pkt[8] = sectors_to_read & 0xFF;
456:     pkt[9] = NULL;
457:     pkt[10] = NULL;
458:     pkt[11] = NULL;
459:
460:     lock_resource(&ide_cd_resource);
461:     for(n = 0; n < sectors_to_read; n++, block++) {
462:         for(retries = 0; retries < MAX_CD_ERR; retries++) {
463:             if(send_packet_command(pkt, ide, drive, blksize)) {
464:                 printk("\tblock=%d, offset=%d\n", block, block *
blksize);
465:                 unlock_resource(&ide_cd_resource);
466:                 return -EIO;
467:             }
468:             if(atapi_read_data(dev, buffer, ide, blksize, n * IDE_CD
_SECTSIZE)) {
469:                 int errcode;
470:                 int sense_key;
471:                 errcode = inport_b(ide->base + IDE_ERROR);
472:                 sense_key = (errcode & 0xF0) >> 4;
473:                 printk("\tSense Key code indicates a '%s' condit
ion.\n", sense_key_err[sense_key & 0xF]);
474:                 if(sense_key) {
475:                     continue;
476:                 }
477:             }
478:             break;
479:         }
480:         if(retries == MAX_CD_ERR) {
481:             printk("\tblock=%d, offset=%d\n", block, block * blksize
);
482:             unlock_resource(&ide_cd_resource);
483:             return -EIO;
484:         }
485:     }
486:     }
487:     unlock_resource(&ide_cd_resource);
488:     return sectors_to_read * IDE_CD_SECTSIZE;
489: }
490:
491: int ide_cd_ioctl(struct inode *i, int cmd, unsigned long int arg)
492: {
493:     struct ide *ide;
494:
495:     if(!(ide = get_ide_controller(i->rdev))) {
496:         return -EINVAL;
497:     }
498:
499:     switch(cmd) {
500:         default:
501:             return -EINVAL;
502:             break;
503:     }
504:
505:     return 0;
506: }
507:
508: int ide_cd_init(struct ide *ide, int drive)
509: {
510:     struct device *d;
511:     unsigned char minor;
512:
513:     ide->drive[drive].fsop = &ide_cd_driver_fsop;

```

**drivers/block/ide\_cd.c**

Page 9/9

```
514:
515:     minor = !ide->drive[drive].minor_shift ? 0 : 1 << ide->drive[drive].mino
r_shift;
516:
517:     if(!(d = get_device(BLK_DEV, MKDEV(ide->drive[drive].major, minor)))) {
518:         return -EINVAL;
519:     }
520:     SET_MINOR(d->minors, minor);
521:     ((unsigned int *)d->device_data)[minor] = CDROM_DEFAULT_SIZE;
522:
523:     return 0;
524: }
```

## drivers/block/ide\_hd.c

Page 1/9

```

1: /*
2:  * fiwix/drivers/block/ide_hd.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/buffer.h>
10: #include <fiwix/ide.h>
11: #include <fiwix/ide_hd.h>
12: #include <fiwix/ioctl.h>
13: #include <fiwix/devices.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/timer.h>
16: #include <fiwix/sched.h>
17: #include <fiwix/cpu.h>
18: #include <fiwix/fs.h>
19: #include <fiwix/part.h>
20: #include <fiwix/process.h>
21: #include <fiwix/mm.h>
22: #include <fiwix/errno.h>
23: #include <fiwix/stdio.h>
24: #include <fiwix/string.h>
25:
26: static struct resource ide_hd_resource = { NULL, NULL };
27:
28: static struct fs_operations ide_hd_driver_fsop = {
29:     0,
30:     0,
31:
32:     ide_hd_open,
33:     ide_hd_close,
34:     NULL, /* read */
35:     NULL, /* write */
36:     ide_hd_ioctl,
37:     NULL, /* lseek */
38:     NULL, /* readdir */
39:     NULL, /* mmap */
40:     NULL, /* select */
41:
42:     NULL, /* readlink */
43:     NULL, /* followlink */
44:     NULL, /* bmap */
45:     NULL, /* lockup */
46:     NULL, /* rmdir */
47:     NULL, /* link */
48:     NULL, /* unlink */
49:     NULL, /* symlink */
50:     NULL, /* mkdir */
51:     NULL, /* mknod */
52:     NULL, /* truncate */
53:     NULL, /* create */
54:     NULL, /* rename */
55:
56:     ide_hd_read,
57:     ide_hd_write,
58:
59:     NULL, /* read_inode */
60:     NULL, /* write_inode */
61:     NULL, /* ialloc */
62:     NULL, /* ifree */
63:     NULL, /* statfs */
64:     NULL, /* read_superblock */
65:     NULL, /* remount_fs */
66:     NULL, /* write_superblock */
67:     NULL, /* release_superblock */

```

## drivers/block/ide\_hd.c

Page 2/9

```

68: };
69:
70: static void assign_minors(__dev_t rdev, struct ide *ide, struct partition *part)
71: {
72:     int n;
73:     int drive, minor;
74:     struct device *d;
75:
76:     minor = 0;
77:     drive = get_ide_drive(rdev);
78:
79:     if(!(d = get_device(BLK_DEV, rdev))) {
80:         printk("WARNING: %s(): unable to assign minors to device %d,%d.\n
n", __FUNCTION__, MAJOR(rdev), MINOR(rdev));
81:         return;
82:     }
83:
84:     for(n = 0; n < NR_PARTITIONS; n++) {
85:         if(drive == IDE_MASTER) {
86:             minor = (1 << ide->drive[drive].minor_shift) + n;
87:         }
88:         if(drive == IDE_SLAVE) {
89:             minor = (1 << ide->drive[drive].minor_shift) + n + 1;
90:         }
91:         CLEAR_MINOR(d->minors, minor);
92:         if(part[n].type) {
93:             SET_MINOR(d->minors, minor);
94:             ((unsigned int *)d->device_data)[minor] = part[n].nr_sec
ts / 2;
95:         }
96:     }
97: }
98:
99: static __off_t block2sector(__off_t offset, int blksize, struct partition *part,
int minor)
100: {
101:     __off_t sector;
102:
103:     sector = offset * (blksize / IDE_HD_SECTSIZE);
104:     if(minor) {
105:         sector += part[minor - 1].startsect;
106:     }
107:     return sector;
108: }
109:
110: static void sector2chs(__off_t offset, int *cyl, int *head, int *sector, struct
ide_drv_ident *ident)
111: {
112:     *cyl = offset / (ident->logic_spt * ident->logic_heads);
113:     *head = (offset / ident->logic_spt) % ident->logic_heads;
114:     *sector = (offset % ident->logic_spt) + 1;
115: }
116:
117: int ide_hd_open(struct inode *i, struct fd *fd_table)
118: {
119:     return 0;
120: }
121:
122: int ide_hd_close(struct inode *i, struct fd *fd_table)
123: {
124:     sync_buffers(i->rdev);
125:     return 0;
126: }
127:
128: int ide_hd_read(__dev_t dev, __blk_t block, char *buffer, int blksize)
129: {
130:     int minor;

```

## drivers/block/ide\_hd.c

Page 3/9

```

131:         int drive;
132:         int sectors_to_read, cmd;
133:         int n, status, r, retries;
134:         int cyl, head, sector;
135:         __off_t offset;
136:         struct ide *ide;
137:         struct ide_drv_ident *ident;
138:         struct partition *part;
139:         struct callout_req creq;
140:
141:         if(! (ide = get_ide_controller(dev))) {
142:             return -EINVAL;
143:         }
144:
145:         minor = MINOR(dev);
146:         if((drive = get_ide_drive(dev)) {
147:             minor &= ~(1 << IDE_SLAVE_MSF);
148:         }
149:
150:         SET_IDE_RDY_RETR(retries);
151:
152:         blksize = blksize ? blksize : BLKSIZE_1K;
153:         sectors_to_read = MIN(blksize, PAGE_SIZE) / IDE_HD_SECTSIZE;
154:
155:         ident = &ide->drive[drive].ident;
156:         part = ide->drive[drive].part_table;
157:         offset = block2sector(block, blksize, part, minor);
158:
159:         CLI();
160:         lock_resource(&ide_hd_resource);
161:
162:         n = 0;
163:
164:         while(n < sectors_to_read) {
165:             if(ide->drive[drive].flags & DEVICE_HAS_RW_MULTIPLE) {
166:                 output_b(ide->base + IDE_SECCNT, sectors_to_read);
167:                 cmd = ATA_READ_MULTIPLE_PIO;
168:             } else {
169:                 output_b(ide->base + IDE_SECCNT, 1);
170:                 cmd = ATA_READ_PIO;
171:             }
172:
173:             if(ide->drive[drive].flags & DEVICE_REQUIRES_LBA) {
174:                 output_b(ide->base + IDE_SECNUM, offset & 0xFF);
175:                 output_b(ide->base + IDE_LCYL, (offset >> 8) & 0xFF);
176:                 output_b(ide->base + IDE_HCYL, (offset >> 16) & 0xFF);
177:                 if(ide_drvsel(ide, drive, IDE_LBA_MODE, (offset >> 24) &
0x0F)) {
178:                     printk("WARNING: %s(): %s: drive not ready.\n",
__FUNCTION__, ide->drive[drive].dev_name);
179:                     unlock_resource(&ide_hd_resource);
180:                     return -EIO;
181:                 }
182:             } else {
183:                 sector2chs(offset, &cyl, &head, &sector, ident);
184:                 output_b(ide->base + IDE_SECNUM, sector);
185:                 output_b(ide->base + IDE_LCYL, cyl);
186:                 output_b(ide->base + IDE_HCYL, (cyl >> 8));
187:                 if(ide_drvsel(ide, drive, IDE_CHS_MODE, head)) {
188:                     printk("WARNING: %s(): %s: drive not ready.\n",
__FUNCTION__, ide->drive[drive].dev_name);
189:                     unlock_resource(&ide_hd_resource);
190:                     return -EIO;
191:                 }
192:             }
193:             output_b(ide->base + IDE_COMMAND, cmd);
194:             if(ide->channel == IDE_PRIMARY) {

```

## drivers/block/ide\_hd.c

Page 4/9

```

195:         ide0_wait_interrupt = ide->base;
196:         creq.fn = ide0_timer;
197:         creq.arg = 0;
198:         add_callout(&creq, WAIT_FOR_IDE);
199:         sleep(&irq_ide0, PROC_UNINTERRUPTIBLE);
200:         if(!ide0_timeout) {
201:             del_callout(&creq);
202:         }
203:     }
204:     if(ide->channel == IDE_SECONDARY) {
205:         idel_wait_interrupt = ide->base;
206:         creq.fn = idel_timer;
207:         creq.arg = 0;
208:         add_callout(&creq, WAIT_FOR_IDE);
209:         sleep(&irq_idel, PROC_UNINTERRUPTIBLE);
210:         if(!idel_timeout) {
211:             del_callout(&creq);
212:         }
213:     }
214:     for(r = 0; r < retries; r++) {
215:         status = inport_b(ide->base + IDE_STATUS);
216:         if(!(status & IDE_STAT_BSY) && (status & IDE_STAT_DRQ))
{
217:             break;
218:         }
219:         ide_delay();
220:     }
221:     if(status & IDE_STAT_ERR) {
222:         printk("WARNING: %s(): %s: error on hard disk dev %d,%d
during read.\n", __FUNCTION__, ide->drive[drive].dev_name, MAJOR(dev), MINOR(dev));
223:         printk("\tstatus=0x%x ", status);
224:         ide_error(ide, status);
225:         printk("\tblock %d, sector %d.\n", block, offset);
226:         inport_b(ide->base + IDE_STATUS); /* clear any pen
ding interrupt */
227:         unlock_resource(&ide_hd_resource);
228:         return -EIO;
229:     }
230:
231:     if(cmd == ATA_READ_MULTIPLE_PIO) {
232:         inport_sw(ide->base + IDE_DATA, (void *)buffer, (IDE_HD_
SECTSIZE * sectors_to_read) / sizeof(short int));
233:         break;
234:     }
235:     inport_sw(ide->base + IDE_DATA, (void *)buffer, IDE_HD_SECTSIZE
/ sizeof(short int));
236:     inport_b(ide->ctrl + IDE_ALT_STATUS); /* ignore results */
237:     inport_b(ide->base + IDE_STATUS); /* clear any pending int
errupt */
238:     n++;
239:     offset++;
240:     buffer += IDE_HD_SECTSIZE;
241: }
242: inport_b(ide->ctrl + IDE_ALT_STATUS); /* ignore results */
243: inport_b(ide->base + IDE_STATUS); /* clear any pending interrupt */
/
244:     unlock_resource(&ide_hd_resource);
245:     return sectors_to_read * IDE_HD_SECTSIZE;
246: }
247:
248: int ide_hd_write(__dev_t dev, __blk_t block, char *buffer, int blksize)
249: {
250:     int minor;
251:     int drive;
252:     int sectors_to_write, cmd;
253:     int n, status, r, retries;
254:     int cyl, head, sector;

```



## drivers/block/ide\_hd.c

Page 5/9

```

255:     __off_t offset;
256:     struct ide *ide;
257:     struct ide_drv_ident *ident;
258:     struct partition *part;
259:     struct callout_req creq;
260:
261:     if(!(ide = get_ide_controller(dev))) {
262:         return -EINVAL;
263:     }
264:
265:     minor = MINOR(dev);
266:     if((drive = get_ide_drive(dev)) {
267:         minor &= ~(1 << IDE_SLAVE_MSF);
268:     }
269:
270:     SET_IDE_RDY_RETR(retries);
271:
272:     blksize = blksize ? blksize : BLKSIZE_1K;
273:     sectors_to_write = MIN(blksize, PAGE_SIZE) / IDE_HD_SECTSIZE;
274:
275:     ident = &ide->drive[drive].ident;
276:     part = ide->drive[drive].part_table;
277:     offset = block2sector(block, blksize, part, minor);
278:
279:     CLI();
280:     lock_resource(&ide_hd_resource);
281:
282:     n = 0;
283:
284:     while(n < sectors_to_write) {
285:         if(ide->drive[drive].flags & DEVICE_HAS_RW_MULTIPLE) {
286:             outport_b(ide->base + IDE_SECCNT, sectors_to_write);
287:             cmd = ATA_WRITE_MULTIPLE_PIO;
288:         } else {
289:             outport_b(ide->base + IDE_SECCNT, 1);
290:             cmd = ATA_WRITE_PIO;
291:         }
292:
293:         if(ide->drive[drive].flags & DEVICE_REQUIRES_LBA) {
294:             outport_b(ide->base + IDE_SECNUM, offset & 0xFF);
295:             outport_b(ide->base + IDE_LCYL, (offset >> 8) & 0xFF);
296:             outport_b(ide->base + IDE_HCYL, (offset >> 16) & 0xFF);
297:             if(ide_drvsel(ide, drive, IDE_LBA_MODE, (offset >> 24) &
0x0F)) {
298:                 printk("WARNING: %s(): %s: drive not ready.\n",
__FUNCTION__, ide->drive[drive].dev_name);
299:                 unlock_resource(&ide_hd_resource);
300:                 return -EIO;
301:             }
302:         } else {
303:             sector2chs(offset, &cyl, &head, &sector, ident);
304:             outport_b(ide->base + IDE_SECNUM, sector);
305:             outport_b(ide->base + IDE_LCYL, cyl);
306:             outport_b(ide->base + IDE_HCYL, (cyl >> 8));
307:             if(ide_drvsel(ide, drive, IDE_CHS_MODE, head)) {
308:                 printk("WARNING: %s(): %s: drive not ready.\n",
__FUNCTION__, ide->drive[drive].dev_name);
309:                 unlock_resource(&ide_hd_resource);
310:                 return -EIO;
311:             }
312:         }
313:         outport_b(ide->base + IDE_COMMAND, cmd);
314:         for(r = 0; r < retries; r++) {
315:             status = inport_b(ide->base + IDE_STATUS);
316:             if(!(status & IDE_STAT_BSY) && (status & IDE_STAT_DRQ))
{
317:                 break;

```

## drivers/block/ide\_hd.c

Page 6/9

```

318:         }
319:         ide_delay();
320:     }
321:     if(status & IDE_STAT_ERR) {
322:         printk("WARNING: %s(): %s: error on hard disk dev %d,%d
during write.\n", __FUNCTION__, ide->drive[drive].dev_name, MAJOR(dev), MINOR(dev));
323:         printk("\tstatus=0x%x ", status);
324:         ide_error(ide, status);
325:         printk("\tblock %d, sector %d.\n", block, offset);
326:         inport_b(ide->base + IDE_STATUS);          /* clear any pen
ding interrupt */
327:         unlock_resource(&ide_hd_resource);
328:         return -EIO;
329:     }
330:
331:     if(cmd == ATA_WRITE_MULTIPLE_PIO) {
332:         outport_sw(ide->base + IDE_DATA, (void *)buffer, (IDE_HD
_SECTSIZE * sectors_to_write) / sizeof(short int));
333:     } else {
334:         outport_sw(ide->base + IDE_DATA, (void *)buffer, IDE_HD_
SECTSIZE / sizeof(short int));
335:     }
336:     if(ide->channel == IDE_PRIMARY) {
337:         ide0_wait_interrupt = ide->base;
338:         creq.fn = ide0_timer;
339:         creq.arg = 0;
340:         add_callout(&creq, WAIT_FOR_IDE);
341:         sleep(&irq_ide0, PROC_UNINTERRUPTIBLE);
342:         if(!ide0_timeout) {
343:             del_callout(&creq);
344:         }
345:     }
346:     if(ide->channel == IDE_SECONDARY) {
347:         idel_wait_interrupt = ide->base;
348:         creq.fn = idel_timer;
349:         creq.arg = 0;
350:         add_callout(&creq, WAIT_FOR_IDE);
351:         sleep(&irq_idel, PROC_UNINTERRUPTIBLE);
352:         if(!idel_timeout) {
353:             del_callout(&creq);
354:         }
355:     }
356:
357:     inport_b(ide->ctrl + IDE_ALT_STATUS);          /* ignore results */
358:     inport_b(ide->base + IDE_STATUS);              /* clear any pending int
errupt */
359:     if(cmd == ATA_WRITE_MULTIPLE_PIO) {
360:         break;
361:     }
362:     n++;
363:     offset++;
364:     buffer += IDE_HD_SECTSIZE;
365: }
366: inport_b(ide->ctrl + IDE_ALT_STATUS);          /* ignore results */
367: inport_b(ide->base + IDE_STATUS);              /* clear any pending interrupt */
/
368:     unlock_resource(&ide_hd_resource);
369:     return sectors_to_write * IDE_HD_SECTSIZE;
370: }
371:
372: int ide_hd_ioctl(struct inode *i, int cmd, unsigned long int arg)
373: {
374:     int minor;
375:     int drive;
376:     struct ide *ide;
377:     struct ide_drv_ident *ident;
378:     struct partition *part;

```

## drivers/block/ide\_hd.c

Page 7/9

```

379:     struct hd_geometry *geom;
380:     int errno;
381:
382:     if(!(ide = get_ide_controller(i->rdev)) {
383:         return -EINVAL;
384:     }
385:
386:     minor = MINOR(i->rdev);
387:     drive = get_ide_drive(i->rdev);
388:     if(drive) {
389:         minor &= ~(1 << IDE_SLAVE_MSF);
390:     }
391:
392:     ident = &ide->drive[drive].ident;
393:     part = ide->drive[drive].part_table;
394:
395:     switch(cmd) {
396:         case HDIO_GETGEO:
397:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
sizeof(struct hd_geometry)))) {
398:                 return errno;
399:             }
400:             geom = (struct hd_geometry *)arg;
401:             geom->cylinders = ident->logic_cyls;
402:             geom->heads = (char)ident->logic_heads;
403:             geom->sectors = (char)ident->logic_spt;
404:             geom->start = 0;
405:             if(minor) {
406:                 geom->start = part[minor - 1].startsect;
407:             }
408:             break;
409:         case BLKGETSIZE:
410:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
sizeof(unsigned int)))) {
411:                 return errno;
412:             }
413:             if(!minor) {
414:                 *(int *)arg = (unsigned int)ide->drive[drive].nr
_sects;
415:             } else {
416:                 *(int *)arg = (unsigned int)ide->drive[drive].pa
rt_table[minor - 1].nr_sects;
417:             }
418:             break;
419:         case BLKFLSBUF:
420:             sync_buffers(i->rdev);
421:             invalidate_buffers(i->rdev);
422:             break;
423:         case BLKRRPART:
424:             read_msdos_partition(i->rdev, part);
425:             assign_minors(i->rdev, ide, part);
426:             break;
427:         default:
428:             return -EINVAL;
429:             break;
430:     }
431:
432:     return 0;
433: }
434:
435: int ide_hd_init(struct ide *ide, int drive)
436: {
437:     int n;
438:     __dev_t rdev;
439:     struct device *d;
440:     struct partition *part;
441:

```

## drivers/block/ide\_hd.c

Page 8/9

```

442:         rdev = 0;
443:         ide->drive[drive].fsop = &ide_hd_driver_fsop;
444:         part = ide->drive[drive].part_table;
445:
446:         if(ide->channel == IDE_PRIMARY) {
447:             if(drive == IDE_MASTER) {
448:                 rdev = MKDEV(IDE0_MAJOR, drive);
449:                 ide->drive[drive].minor_shift = IDE_MASTER_MSF;
450:                 if(!(d = get_device(BLK_DEV, rdev))) {
451:                     return -EINVAL;
452:                 }
453:                 ((unsigned int *)d->device_data)[0] = ide->drive[drive].
nr_sects / 2;
454:             } else {
455:                 rdev = MKDEV(IDE0_MAJOR, 1 << IDE_SLAVE_MSF);
456:                 ide->drive[drive].minor_shift = IDE_SLAVE_MSF;
457:                 if(!(d = get_device(BLK_DEV, rdev))) {
458:                     return -EINVAL;
459:                 }
460:                 ((unsigned int *)d->device_data)[1 << IDE_SLAVE_MSF] = i
de->drive[drive].nr_sects / 2;
461:             }
462:         } else if(ide->channel == IDE_SECONDARY) {
463:             if(drive == IDE_MASTER) {
464:                 rdev = MKDEV(IDE1_MAJOR, drive);
465:                 ide->drive[drive].minor_shift = IDE_MASTER_MSF;
466:                 if(!(d = get_device(BLK_DEV, rdev))) {
467:                     return -EINVAL;
468:                 }
469:                 ((unsigned int *)d->device_data)[0] = ide->drive[drive].
nr_sects / 2;
470:             } else {
471:                 rdev = MKDEV(IDE1_MAJOR, 1 << IDE_SLAVE_MSF);
472:                 ide->drive[drive].minor_shift = IDE_SLAVE_MSF;
473:                 if(!(d = get_device(BLK_DEV, rdev))) {
474:                     return -EINVAL;
475:                 }
476:                 ((unsigned int *)d->device_data)[1 << IDE_SLAVE_MSF] = i
de->drive[drive].nr_sects / 2;
477:             }
478:         } else {
479:             printk("WARNING: %s(): invalid drive number %d.\n", __FUNCTION__
, drive);
480:             return 1;
481:         }
482:
483:         read_msdos_partition(rdev, part);
484:         assign_minors(rdev, ide, part);
485:         printk("                partition summary: ");
486:         for(n = 0; n < NR_PARTITIONS; n++) {
487:             if(part[n].type) {
488:                 printk("%s%d ", ide->drive[drive].dev_name, n + 1);
489:             }
490:         }
491:         printk("\n");
492:
493:         outport_b(ide->ctrl + IDE_DEV_CTRL, IDE_DEVCTR_NIEN);
494:         if(ide->drive[drive].flags & DEVICE_HAS_RW_MULTIPLE) {
495:             /* read/write in 4KB blocks (8 sectors) by default */
496:             outport_b(ide->base + IDE_SECCNT, PAGE_SIZE / IDE_HD_SECTSIZE);
497:             outport_b(ide->base + IDE_COMMAND, ATA_SET_MULTIPLE_MODE);
498:             ide_wait400ns(ide);
499:             while(inport_b(ide->base + IDE_STATUS) & IDE_STAT_BSY);
500:         }
501:         outport_b(ide->ctrl + IDE_DEV_CTRL, IDE_DEVCTR_DRQ);
502:
503:         return 0;

```

```
504: }
```

**drivers/block/Makefile**

Page 1/1

```
1: # fiwix/drivers/block/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = dma.o floppy.o part.o ide.o ide_hd.o ide_cd.o ramdisk.o
13:
14: block:  $(OBJS)
15:           $(LD) $(LDFLAGS) -r $(OBJS) -o block.o
16:
17: clean:
18:           rm -f *.o
19:
```

## drivers/block/part.c

Page 1/1

```
1: /*
2:  * fiwix/drivers/block/part.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/ide.h>
9: #include <fiwix/ide_hd.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/part.h>
12: #include <fiwix/mm.h>
13: #include <fiwix/errno.h>
14: #include <fiwix/stdio.h>
15: #include <fiwix/string.h>
16:
17: int read_msdos_partition(__dev_t dev, struct partition *part)
18: {
19:     char *buffer;
20:
21:     if(!(buffer = (void *)kmalloc())) {
22:         return -ENOMEM;
23:     }
24:
25:     if(ide_hd_read(dev, PARTITION_BLOCK, buffer, BLKSIZE_1K) <= 0) {
26:         printk("WARNING: %s(): unable to read partition block in device
%d,%d.\n", __FUNCTION__, MAJOR(dev), MINOR(dev));
27:         kfree((unsigned int)buffer);
28:         return -EIO;
29:     }
30:
31:     memcpy_b(part, (void *) (buffer + MBR_CODE_SIZE), sizeof(struct partition
) * NR_PARTITIONS);
32:     kfree((unsigned int)buffer);
33:     return 0;
34: }
```

## drivers/block/ramdisk.c

Page 1/3

```

1:  /*
2:  *  fiwix/drivers/block/ramdisk.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/kernel.h>
9:  #include <fiwix/ramdisk.h>
10: #include <fiwix/ioctl.h>
11: #include <fiwix/devices.h>
12: #include <fiwix/part.h>
13: #include <fiwix/fs.h>
14: #include <fiwix/errno.h>
15: #include <fiwix/mm.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: static unsigned int rd_sizes[256];
20:
21: static struct fs_operations ramdisk_driver_fsop = {
22:     0,
23:     0,
24:
25:     ramdisk_open,
26:     ramdisk_close,
27:     NULL,                /* read */
28:     NULL,                /* write */
29:     ramdisk_ioctl,
30:     ramdisk_lseek,
31:     NULL,                /* readdir */
32:     NULL,                /* mmap */
33:     NULL,                /* select */
34:
35:     NULL,                /* readlink */
36:     NULL,                /* followlink */
37:     NULL,                /* bmap */
38:     NULL,                /* lockup */
39:     NULL,                /* rmdir */
40:     NULL,                /* link */
41:     NULL,                /* unlink */
42:     NULL,                /* symlink */
43:     NULL,                /* mkdir */
44:     NULL,                /* mknod */
45:     NULL,                /* truncate */
46:     NULL,                /* create */
47:     NULL,                /* rename */
48:
49:     ramdisk_read,
50:     ramdisk_write,
51:
52:     NULL,                /* read_inode */
53:     NULL,                /* write_inode */
54:     NULL,                /* ialloc */
55:     NULL,                /* ifree */
56:     NULL,                /* statfs */
57:     NULL,                /* read_superblock */
58:     NULL,                /* remount_fs */
59:     NULL,                /* write_superblock */
60:     NULL,                /* release_superblock */
61: };
62:
63: static struct device ramdisk_device = {
64:     "ramdisk",
65:     RAMDISK_MAJOR,
66:     { 0, 0, 0, 0, 0, 0, 0, 0, 0 },
67:     BLKSIZE_1K,

```



## drivers/block/ramdisk.c

Page 2/3

```

68:         &rd_sizes,
69:         &ramdisk_driver_fsop,
70:         NULL
71: };
72:
73: static struct ramdisk * get_ramdisk(int minor)
74: {
75:     if(TEST_MINOR(ramdisk_device.minors, minor)) {
76:         return &ramdisk_table[minor];
77:     }
78:     return NULL;
79: }
80:
81: int ramdisk_open(struct inode *i, struct fd *fd_table)
82: {
83:     if(!get_ramdisk(MINOR(i->rdev))) {
84:         return -ENXIO;
85:     }
86:     return 0;
87: }
88:
89: int ramdisk_close(struct inode *i, struct fd *fd_table)
90: {
91:     if(!get_ramdisk(MINOR(i->rdev))) {
92:         return -ENXIO;
93:     }
94:     return 0;
95: }
96:
97: int ramdisk_read(__dev_t dev, __blk_t block, char *buffer, int blksize)
98: {
99:     int size;
100:    __off_t offset;
101:    struct ramdisk *ramdisk;
102:
103:    if(!(ramdisk = get_ramdisk(MINOR(dev)))) {
104:        return -ENXIO;
105:    }
106:
107:    size = rd_sizes[MINOR(dev)] * 1024;
108:    offset = block * blksize;
109:    blksize = MIN(blksize, size - offset);
110:    memcpy_b((void *)buffer, ramdisk->addr + offset, blksize);
111:    return blksize;
112: }
113:
114: int ramdisk_write(__dev_t dev, __blk_t block, char *buffer, int blksize)
115: {
116:     int size;
117:     __off_t offset;
118:     struct ramdisk *ramdisk;
119:
120:     if(!(ramdisk = get_ramdisk(MINOR(dev)))) {
121:         return -ENXIO;
122:     }
123:
124:     size = rd_sizes[MINOR(dev)] * 1024;
125:     offset = block * blksize;
126:     blksize = MIN(blksize, size - offset);
127:     memcpy_b((void *)ramdisk->addr + offset, buffer, blksize);
128:     return blksize;
129: }
130:
131: int ramdisk_ioctl(struct inode *i, int cmd, unsigned long int arg)
132: {
133:     struct hd_geometry *geom;
134:     int errno;

```

## drivers/block/ramdisk.c

Page 3/3

```

135:
136:     if(!get_ramdisk(MINOR(i->rdev)) ) {
137:         return -ENXIO;
138:     }
139:
140:     switch(cmd) {
141:         case HDIO_GETGEO:
142:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
sizeof(struct hd_geometry)))) {
143:                 return errno;
144:             }
145:             geom = (struct hd_geometry *)arg;
146:             geom->heads = 63;
147:             geom->sectors = 16;
148:             geom->cylinders = rd_sizes[MINOR(i->rdev)] * 1024 / BPS;
149:             geom->cylinders /= (geom->heads * geom->sectors);
150:             geom->start = 0;
151:             break;
152:         case BLKRRPART:
153:             break;
154:         case BLKGETSIZE:
155:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
sizeof(unsigned int)))) {
156:                 return errno;
157:             }
158:             *(int *)arg = rd_sizes[MINOR(i->rdev)] * 2;
159:             break;
160:         default:
161:             return -EINVAL;
162:     }
163:     return 0;
164: }
165:
166: int ramdisk_lseek(struct inode *i, __off_t offset)
167: {
168:     return offset;
169: }
170:
171: void ramdisk_init(void)
172: {
173:     int n;
174:     struct ramdisk *ramdisk;
175:
176:     if(!_noramdisk) {
177:         for(n = 0; n < RAMDISK_MINORS; n++) {
178:             SET_MINOR(ramdisk_device.minors, n);
179:             rd_sizes[n] = _ramdisksize;
180:             ramdisk = get_ramdisk(n);
181:             printk("ram%d      0x%08x-0x%08x %d RAMdisk(s) of %dKB s
ize, %dKB blocksize\n", n, ramdisk->addr, ramdisk->addr + (_ramdisksize * 1024), RAMDIS
K_MINORS, _ramdisksize, BLKSIZE_1K / 1024);
182:         }
183:         register_device(BLK_DEV, &ramdisk_device);
184:     }
185: }

```



## drivers/char/console.c

Page 2/21

```

68:         "\205\240\203\376\204\206\221\207\212\202\210\211\215\241\214\213"
69:         "\376\244\225\242\223\376\224\366\355\227\243\226\201\376\376\230"
70: ;
71:
72: short int current_cons;
73: unsigned char screen_is_off = 0;
74: int buf_y, buf_top;
75:
76: struct video_parms video;
77: struct vconsole vc[NR_VCONSOLES + 1];
78:
79: /*
80:  * This is the scrollbar history buffer which it's used only in the active
81:  * vconsole. Everytime a vconsole is switched, the screen content of the
82:  * previous vconsole is copied from video memory to kernel memory, and the
83:  * screen content of the new vconsole is copied from kernel memory to video
84:  * memory. Only the visible screen is copied, so switching vconsoles means
85:  * losing the scrollbar history.
86:  */
87: short int vcbuf[80 * 25 * SCREENS_LOG * 2];      /* FIXME: allocate this dynamica
lly */
88:
89: static struct fs_operations tty_driver_fsop = {
90:     0,
91:     0,
92:
93:     tty_open,
94:     tty_close,
95:     tty_read,
96:     tty_write,
97:     tty_ioctl,
98:     tty_lseek,
99:     NULL,          /* readdir */
100:    NULL,          /* mmap */
101:    tty_select,
102:
103:    NULL,          /* readlink */
104:    NULL,          /* followlink */
105:    NULL,          /* bmap */
106:    NULL,          /* lookup */
107:    NULL,          /* rmdir */
108:    NULL,          /* link */
109:    NULL,          /* unlink */
110:    NULL,          /* symlink */
111:    NULL,          /* mkdir */
112:    NULL,          /* mknod */
113:    NULL,          /* truncate */
114:    NULL,          /* create */
115:    NULL,          /* rename */
116:
117:    NULL,          /* read_block */
118:    NULL,          /* write_block */
119:
120:    NULL,          /* read_inode */
121:    NULL,          /* write_inode */
122:    NULL,          /* ialloc */
123:    NULL,          /* ifree */
124:    NULL,          /* statfs */
125:    NULL,          /* read_superblock */
126:    NULL,          /* remount_fs */
127:    NULL,          /* write_superblock */
128:    NULL,          /* release_superblock */
129: };
130:
131: static struct device tty_device = {
132:     "vconsole",
133:     VCONSOLES_MAJOR,

```

## drivers/char/console.c

Page 3/21

```

134:         { 0, 0, 0, 0, 0, 0, 0, 0, 0 },
135:         0,
136:         NULL,
137:         &tty_driver_fsop,
138:         NULL
139: };
140:
141: static struct device console_device = {
142:     "console",
143:     SYSCON_MAJOR,
144:     { 0, 0, 0, 0, 0, 0, 0, 0 },
145:     0,
146:     NULL,
147:     &tty_driver_fsop,
148:     NULL
149: };
150:
151: unsigned short int ansi_color_table[] = {
152:     COLOR_BLACK,
153:     COLOR_RED,
154:     COLOR_GREEN,
155:     COLOR_BROWN,
156:     COLOR_BLUE,
157:     COLOR_MAGENTA,
158:     COLOR_CYAN,
159:     COLOR_WHITE
160: };
161:
162: static int is_vconsole(__dev_t dev)
163: {
164:     if(MAJOR(dev) == VCONSOLES_MAJOR && MINOR(dev) <= NR_VCONSOLES) {
165:         return 1;
166:     }
167:
168:     return 0;
169: }
170:
171: static void update_curpos(struct vconsole *vc)
172: {
173:     unsigned short int curpos;
174:
175:     if(vc->has_focus) {
176:         curpos = (vc->y * vc->columns) + vc->x;
177:         outport_b(VIDEO_PORT + CRT_INDEX, CRT_CURSOR_POS_HI);
178:         outport_b(VIDEO_PORT + CRT_DATA, (curpos >> 8) & 0xFF);
179:         outport_b(VIDEO_PORT + CRT_INDEX, CRT_CURSOR_POS_LO);
180:         outport_b(VIDEO_PORT + CRT_DATA, (curpos & 0xFF));
181:     }
182: }
183:
184: static void show_cursor(int mode)
185: {
186:     int status;
187:
188:     switch(mode) {
189:         case ON:
190:             outport_b(VIDEO_PORT + CRT_INDEX, CRT_CURSOR_STR);
191:             status = inport_b(VIDEO_PORT + CRT_DATA);
192:             outport_b(VIDEO_PORT + CRT_DATA, status & CURSOR_MASK);
193:             break;
194:         case OFF:
195:             outport_b(VIDEO_PORT + CRT_INDEX, CRT_CURSOR_STR);
196:             status = inport_b(VIDEO_PORT + CRT_DATA);
197:             outport_b(VIDEO_PORT + CRT_DATA, status | CURSOR_DISABLE
);
198:             break;
199:     }

```

## drivers/char/console.c

Page 4/21

```
200: }
201:
202: static void get_curpos(struct vconsole *vc)
203: {
204:     unsigned short int curpos;
205:
206:     outport_b(VIDEO_PORT + CRT_INDEX, CRT_CURSOR_POS_HI);
207:     curpos = inport_b(VIDEO_PORT + CRT_DATA) << 8;
208:     outport_b(VIDEO_PORT + CRT_INDEX, CRT_CURSOR_POS_LO);
209:     curpos |= inport_b(VIDEO_PORT + CRT_DATA);
210:
211:     vc->x = curpos % vc->columns;
212:     vc->y = curpos / vc->columns;
213: }
214:
215: static void adjust(struct vconsole *vc, int x, int y)
216: {
217:     if(x < 0) {
218:         x = 0;
219:     }
220:     if(x >= vc->columns) {
221:         x = vc->columns - 1;
222:     }
223:     if(y < 0) {
224:         y = 0;
225:     }
226:     if(y >= vc->bottom) {
227:         y = vc->bottom - 1;
228:     }
229:     vc->x = x;
230:     vc->y = y;
231: }
232:
233: static void delete_char(struct vconsole *vc)
234: {
235:     int offset, count;
236:
237:     offset = (vc->y * vc->columns) + vc->x;
238:     count = vc->columns - vc->x;
239:     memcpy_w(vc->vidmem + offset, vc->vidmem + offset + 1, count);
240:     memset_w(vc->vidmem + offset + count, BLANK_MEM, 1);
241: }
242:
243: static void insert_char(struct vconsole *vc)
244: {
245:     int n, offset;
246:     short int tmp, last_char;
247:
248:     offset = (vc->y * vc->columns) + vc->x;
249:     n = vc->x;
250:     last_char = BLANK_MEM;
251:     while(n++ < vc->columns) {
252:         memcpy_w(&tmp, vc->vidmem + offset, 1);
253:         memset_w(vc->vidmem + offset, last_char, 1);
254:         last_char = tmp;
255:         offset++;
256:     }
257: }
258:
259: static void scroll_screen(struct vconsole *vc, int top, int mode)
260: {
261:     int n, offset, count;
262:
263:     if(!top) {
264:         top = vc->top;
265:     }
266:     switch(mode) {
```

## drivers/char/console.c

Page 5/21

```

267:         case SCROLL_UP:
268:             count = vc->columns * (vc->bottom - top - 1);
269:             offset = top * vc->columns;
270:             top = (top + 1) * vc->columns;
271:             memcpy_w(vc->vidmem + offset, vc->vidmem + top, count);
272:             memset_w(vc->vidmem + offset + count, BLANK_MEM, top);
273:             break;
274:         case SCROLL_DOWN:
275:             count = vc->columns;
276:             for(n = vc->bottom - 1; n >= top; n--) {
277:                 memcpy_w(vc->vidmem + (vc->columns * (n + 1)), v
c->vidmem + (vc->columns * n), count);
278:             }
279:             memset_w(vc->vidmem + (top * vc->columns), BLANK_MEM, co
unt);
280:             break;
281:     }
282:     return;
283: }
284:
285: static void cr(struct vconsole *vc)
286: {
287:     vc->x = 0;
288: }
289:
290: static void lf(struct vconsole *vc)
291: {
292:     if(vc->y == vc->bottom) {
293:         scroll_screen(vc, 0, SCROLL_UP);
294:     } else {
295:         vc->y++;
296:     }
297: }
298:
299: static void ri(struct vconsole *vc)
300: {
301:     if(vc->y == 0) {
302:         scroll_screen(vc, 0, SCROLL_DOWN);
303:     } else {
304:         vc->y--;
305:     }
306: }
307:
308: static void csi_J(struct vconsole *vc, int mode)
309: {
310:     int from, count;
311:
312:     switch(mode) {
313:         case CSI_J_CUR2END: /* Erase Down <ESC>[J */
314:             from = ((vc->y * vc->columns) + vc->x) * 2;
315:             count = (SCREEN_SIZE - from) / sizeof(short int);
316:             break;
317:         case CSI_J_STA2CUR: /* Erase Up <ESC>[1J */
318:             from = 0;
319:             count = (((vc->y * vc->columns) + vc->x) * 2) / sizeof(s
hort int);
320:             break;
321:         case CSI_J_SCREEN: /* Erase Screen <ESC>[2J */
322:             from = 0;
323:             count = SCREEN_SIZE / sizeof(short int);
324:             break;
325:         default:
326:             return;
327:     }
328:     memset_w(vc->vidmem + from, vc->color_attr, count);
329: }
330:

```

## drivers/char/console.c

Page 6/21

```

331: static void csi_K(struct vconsole *vc, int mode)
332: {
333:     int from, count;
334:
335:     switch(mode) {
336:         case CSI_K_CUR2END: /* Erase End of Line <ESC>[K */
337:             from = (vc->y * vc->columns) + vc->x;
338:             count = ((vc->columns - vc->x) * 2) / sizeof(short int);
339:             break;
340:         case CSI_K_STA2CUR: /* Erase Start of Line <ESC>[1K */
341:             from = vc->y * vc->columns;
342:             count = (vc->x * 2) / sizeof(short int);
343:             break;
344:         case CSI_K_LINE: /* Erase Line <ESC>[2K */
345:             from = vc->y * vc->columns;
346:             count = (vc->columns * 2) / sizeof(short int);
347:             break;
348:         default:
349:             return;
350:     }
351:     memset_w(vc->vidmem + from, vc->color_attr, count);
352: }
353:
354: static void csi_X(struct vconsole *vc, int count)
355: {
356:     int from;
357:
358:     from = (vc->y * vc->columns) + vc->x;
359:     count = count > (vc->columns - vc->x) ? vc->columns - vc->x : count;
360:     memset_w(vc->vidmem + from, vc->color_attr, count * 2);
361: }
362:
363: static void csi_L(struct vconsole *vc, int count)
364: {
365:     if(count > (vc->bottom - vc->top)) {
366:         count = vc->bottom - vc->top;
367:     }
368:     while(count--) {
369:         scroll_screen(vc, vc->y, SCROLL_DOWN);
370:     }
371: }
372:
373: static void csi_M(struct vconsole *vc, int count)
374: {
375:     if(count > (vc->bottom - vc->top)) {
376:         count = vc->bottom - vc->top;
377:     }
378:     while(count--) {
379:         scroll_screen(vc, vc->y, SCROLL_UP);
380:     }
381: }
382:
383: static void csi_P(struct vconsole *vc, int count)
384: {
385:     if(count > vc->columns) {
386:         count = vc->columns;
387:     }
388:     while(count--) {
389:         delete_char(vc);
390:     }
391: }
392:
393: static void csi_at(struct vconsole *vc, int count)
394: {
395:     if(count > vc->columns) {
396:         count = vc->columns;
397:     }

```



## drivers/char/console.c

Page 7/21

```

398:         while(count--) {
399:             insert_char(vc);
400:         }
401:     }
402:
403: static void default_color_attr(struct vconsole *vc)
404: {
405:     vc->color_attr = DEF_MODE;
406:     vc->bold = 0;
407:     vc->underline = 0;
408:     vc->blink = 0;
409:     vc->reverse = 0;
410: }
411:
412: /* Select Graphic Rendition */
413: static void csi_m(struct vconsole *vc)
414: {
415:     int n;
416:
417:     if(vc->reverse) {
418:         vc->color_attr = ((vc->color_attr & 0x7000) >> 4) | ((vc->color_
attr & 0x0700) << 4) | (vc->color_attr & 0x8800);
419:     }
420:
421:     for(n = 0; n < vc->nparms; n++) {
422:         switch(vc->parms[n]) {
423:             case SGR_DEFAULT:
424:                 default_color_attr(vc);
425:                 break;
426:             case SGR_BOLD:
427:                 vc->bold = 1;
428:                 break;
429:             case SGR_BLINK:
430:                 vc->blink = 1;
431:                 break;
432:             case SGR_REVERSE:
433:                 vc->reverse = 1;
434:                 break;
435:             /* normal intensity */
436:             case 21:
437:             case 22:
438:                 vc->bold = 0;
439:                 break;
440:             case SGR_BLINK_OFF:
441:                 vc->blink = 0;
442:                 break;
443:             case SGR_REVERSE_OFF:
444:                 vc->reverse = 0;
445:                 break;
446:             case SGR_BLACK_FG:
447:             case SGR_RED_FG:
448:             case SGR_GREEN_FG:
449:             case SGR_BROWN_FG:
450:             case SGR_BLUE_FG:
451:             case SGR_MAGENTA_FG:
452:             case SGR_CYAN_FG:
453:             case SGR_WHITE_FG:
454:                 vc->color_attr = (vc->color_attr & 0xF8FF) | (an
si_color_table[vc->parms[n] - 30]);
455:                 break;
456:             case SGR_DEFAULT_FG_U_ON:
457:             case SGR_DEFAULT_FG_U_OFF:
458:                 /* not supported yet */
459:                 break;
460:             case SGR_BLACK_BG:
461:             case SGR_RED_BG:
462:             case SGR_GREEN_BG:

```

## drivers/char/console.c

Page 8/21

```

463:             case SGR_BROWN_BG:
464:             case SGR_BLUE_BG:
465:             case SGR_MAGENTA_BG:
466:             case SGR_CYAN_BG:
467:             case SGR_WHITE_BG:
468:                 vc->color_attr = (vc->color_attr & 0x8FFF) | ((a
nsi_color_table[vc->parms[n] - 40]) << 4);
469:                 break;
470:             case SGR_DEFAULT_BG:
471:                 /* not supported yet */
472:                 break;
473:         }
474:     }
475:     if(vc->bold) {
476:         vc->color_attr |= 0x0800;
477:     } else {
478:         vc->color_attr &= ~0x0800;
479:     }
480:     if(vc->blink) {
481:         vc->color_attr |= 0x8000;
482:     } else {
483:         vc->color_attr &= ~0x8000;
484:     }
485:     if(vc->reverse) {
486:         vc->color_attr = ((vc->color_attr & 0x7000) >> 4) | ((vc->color_
attr & 0x0700) << 4) | (vc->color_attr & 0x8800);
487:     }
488: }
489:
490: static void init_vt(struct vconsole *vc)
491: {
492:     vc->vt_mode.mode = VT_AUTO;
493:     vc->vt_mode.waitv = 0;
494:     vc->vt_mode.relsig = 0;
495:     vc->vt_mode.acqsig = 0;
496:     vc->vt_mode.frsig = 0;
497:     vc->vc_mode = KD_TEXT;
498:     vc->tty->pid = 0;
499:     vc->switchto_tty = -1;
500: }
501:
502: static void insert_seq(struct tty *tty, char *buf, int count)
503: {
504:     while(count--) {
505:         tty_queue_putchar(tty, &tty->read_q, *(buf++));
506:     }
507:     tty->input(tty);
508: }
509:
510: static void echo_char(struct vconsole *vc, unsigned char *buf, unsigned int coun
t)
511: {
512:     unsigned char ch;
513:
514:     if(vc->has_focus) {
515:         if(buf_top) {
516:             vconsole_restore(vc);
517:             show_cursor(ON);
518:             buf_top = 0;
519:         }
520:     }
521:
522:     while(count--) {
523:         ch = *buf++;
524:         if(ch == NULL) {
525:             continue;
526:

```

## drivers/char/console.c

Page 9/21

```

527:         } else if(ch == '\b') {
528:             if(vc->x) {
529:                 vc->x--;
530:             }
531:
532:         } else if(ch == '\a') {
533:             vconsole_beep();
534:
535:         } else if(ch == '\r') {
536:             cr(vc);
537:
538:         } else if(ch == '\n') {
539:             cr(vc);
540:             vc->y++;
541:             if(vc->has_focus) {
542:                 buf_y++;
543:             }
544:
545:         } else if(ch == '\t') {
546:             while(vc->x < (vc->columns - 1)) {
547:                 if(vc->tty->tab_stop[+vc->x]) {
548:                     break;
549:                 }
550:             }
551:             /*
552:             vc->x += TAB_SIZE - (vc->x % TAB_SIZE); */
553:             vc->check_x = 1;
554:         } else {
555:             if((vc->x == vc->columns - 1) && vc->check_x) {
556:                 vc->x = 0;
557:                 vc->y++;
558:                 if(vc->has_focus) {
559:                     buf_y++;
560:                 }
561:             }
562:             if(vc->y >= vc->bottom) {
563:                 scroll_screen(vc, 0, SCROLL_UP);
564:                 vc->y--;
565:             }
566:             ch = iso8859[ch];
567:             vc->vidmem[(vc->y * vc->columns) + vc->x] = vc->color_at
tr | ch;
568:             if(vc->has_focus) {
569:                 vcbuf[(buf_y * vc->columns) + vc->x] = vc->color
_attr | ch;
570:             }
571:             if(vc->x < vc->columns - 1) {
572:                 vc->check_x = 0;
573:                 vc->x++;
574:             } else {
575:                 vc->check_x = 1;
576:             }
577:         }
578:         if(vc->y >= vc->bottom) {
579:             scroll_screen(vc, 0, SCROLL_UP);
580:             vc->y--;
581:         }
582:         if(vc->has_focus) {
583:             if(buf_y >= VC_BUF_LINES) {
584:                 memcpy_w(vcbuf, vcbuf + SCREEN_COLS, VC_BUF_SIZE
- SCREEN_COLS);
585:                 buf_y--;
586:             }
587:         }
588:     }
589:     update_curpos(vc);
590: }

```

## drivers/char/console.c

Page 10/21

```

591:
592: void vconsole_reset(struct tty *tty)
593: {
594:     int n;
595:     struct vconsole *vc;
596:
597:     vc = (struct vconsole *)tty->driver_data;
598:
599:     vc->top = 0;
600:     vc->bottom = SCREEN_LINES;
601:     vc->columns = SCREEN_COLS;
602:     vc->check_x = 0;
603:     vc->led_status = 0;
604:     set_leds(vc->led_status);
605:     vc->scrlock = vc->numlock = vc->capslock = 0;
606:     vc->esc = vc->sbracket = vc->semicolon = vc->question = 0;
607:     vc->parmv1 = vc->parmv2 = 0;
608:     vc->nparms = 0;
609:     memset_b(vc->parms, 0, sizeof(vc->parms));
610:     default_color_attr(vc);
611:     vc->insert_mode = 0;
612:     vc->saved_x = vc->saved_y = 0;
613:
614:     for(n = 0; n < MAX_TAB_COLS; n++) {
615:         if(!(n % TAB_SIZE)) {
616:             vc->tty->tab_stop[n] = 1;
617:         } else {
618:             vc->tty->tab_stop[n] = 0;
619:         }
620:     }
621:
622:     termios_reset(tty);
623:     vc->tty->winsize.ws_row = vc->bottom - vc->top;
624:     vc->tty->winsize.ws_col = vc->columns;
625:     vc->tty->winsize.ws_xpixel = 0;
626:     vc->tty->winsize.ws_ypixel = 0;
627:     vc->tty->lnext = 0;
628:
629:     init_vt(vc);
630:     vc->blanked = 0;
631:     update_curpos(vc);
632: }
633:
634: void vconsole_write(struct tty *tty)
635: {
636:     int n;
637:     unsigned char ch;
638:     int numeric;
639:     struct vconsole *vc;
640:
641:     vc = (struct vconsole *)tty->driver_data;
642:
643:     if(buf_top) {
644:         vconsole_restore(vc);
645:         buf_top = 0;
646:         show_cursor(ON);
647:         update_curpos(vc);
648:     }
649:
650:     numeric = 0;
651:
652:     while(!vc->scrlock && tty->write_q.count > 0) {
653:         ch = tty_queue_getchar(&tty->write_q);
654:
655:         if(vc->esc) {
656:             if(vc->sbracket) {
657:                 if(IS_NUMERIC(ch)) {

```

## drivers/char/console.c

Page 11/21

```

658:
659:
660:
661:
662:
663:
664:
665:
666:
667:
668:
669:
670:
671:
672:
673:
674:
675:
676:
677:
678:
679:
ESC>[ n @ */
680:
c->parmv1;
681:
682:
683:
684:
*/
685:
c->parmv1;
686:
rmv1);
687:
688:
689:
B */
690:
c->parmv1;
691:
rmv1);
692:
693:
694:
n C */
695:
c->parmv1;
696:
c->y);
697:
698:
699:
[ n D */
700:
c->parmv1;
701:
c->y);
702:
703:
704:
ESC>[ n E */
705:
c->parmv1;
706:
);
707:

```

```

numeric = 1;
if(vc->:semicolon) {
    vc->parmv2 *= 10;
    vc->parmv2 += ch - '0';
} else {
    vc->parmv1 *= 10;
    vc->parmv1 += ch - '0';
}
vc->parms[vc->nparms] *= 10;
vc->parms[vc->nparms] += ch - '0';
continue;
}
switch(ch) {
case ';':
    vc->:semicolon = 1;
    vc->parmv2 = 0;
    vc->nparms++;
    continue;
case '?':
    vc->question = 1;
    continue;
case '@': /* Insert Character(s) <
    vc->parmv1 = !vc->parmv1 ? 1 : v
    csi_at(vc, vc->parmv1);
    CSE;
    continue;
case 'A': /* Cursor Up <ESC>[ n A
    vc->parmv1 = !vc->parmv1 ? 1 : v
    adjust(vc, vc->x, vc->y - vc->pa
    CSE;
    continue;
case 'B': /* Cursor Down <ESC>[ n
    vc->parmv1 = !vc->parmv1 ? 1 : v
    adjust(vc, vc->x, vc->y + vc->pa
    CSE;
    continue;
case 'C': /* Cursor Forward <ESC>[
    vc->parmv1 = !vc->parmv1 ? 1 : v
    adjust(vc, vc->x + vc->parmv1, v
    CSE;
    continue;
case 'D': /* Cursor Backward <ESC>
    vc->parmv1 = !vc->parmv1 ? 1 : v
    adjust(vc, vc->x - vc->parmv1, v
    CSE;
    continue;
case 'E': /* Cursor Next Line(s) <
    vc->parmv1 = !vc->parmv1 ? 1 : v
    adjust(vc, 0, vc->y + vc->parmv1
    CSE;

```

## drivers/char/console.c

Page 12/21

```

708:
709:
s) <ESC>[ n F */
710:
c->parmv1;
711:
);
712:
713:
714:
ition <ESC>[ n G */
715:
716:
rmv1 - 1 : vc->parmv1;
717:
718:
719:
720:
W ; COLUMN H */
721:
osition <ESC>[ ROW ; COLUMN f */
722:
rmv1 - 1 : vc->parmv1;
723:
rmv2 - 1 : vc->parmv2;
724:
1);
725:
726:
727:
tion <ESC>[ n I */
728:
c->parmv1;
729:
730:
ns - 1)) {
731:
stop[++vc->x]) {
732:
733:
734:
735:
736:
737:
738:
739:
) <ESC>[J */
740:
741:
742:
743:
f/) Line <ESC>[K */
744:
745:
746:
747:
n L */
748:
c->parmv1;
749:
750:
751:
752:
n M */
753:
c->parmv1;
754:

```

```

        continue;
case 'F': /* Cursor Previous Line (
        vc->parmv1 = !vc->parmv1 ? 1 : v
        adjust(vc, 0, vc->y - vc->parmv1
        );
        CSE;
        continue;
case 'G': /* Cursor Horizontal Pos
case 'H': /* Cursor Home <ESC>[ RO
case 'f': /* Horizontal Vertical P
        vc->parmv1 = vc->parmv1 ? vc->pa
        vc->parmv2 = vc->parmv2 ? vc->pa
        adjust(vc, vc->parmv2, vc->parmv
        );
        CSE;
        continue;
case 'I': /* Cursor Forward Tabula
        vc->parmv1 = !vc->parmv1 ? 1 : v
        while(vc->parmv1--){
            while(vc->x < (vc->column
                )
                if(vc->tty->tab_
                    )
                    break;
            }
        }
        adjust(vc, vc->x, vc->y);
        CSE;
        continue;
case 'J': /* Erase (Down/Up/Screen
        csi_J(vc, vc->parmv1);
        CSE;
        continue;
case 'K': /* Erase (End of/Start o
        csi_K(vc, vc->parmv1);
        CSE;
        continue;
case 'L': /* Insert Line(s) <ESC>[
        vc->parmv1 = !vc->parmv1 ? 1 : v
        csi_L(vc, vc->parmv1);
        CSE;
        continue;
case 'M': /* Delete Line(s) <ESC>[
        vc->parmv1 = !vc->parmv1 ? 1 : v
        csi_M(vc, vc->parmv1);

```

## drivers/char/console.c

Page 13/21

```

755:                                     CSE;
756:                                     continue;
757:                                     case 'P': /* Delete Character(s) <
ESC>[ n P */
758:                                     vc->parmv1 = !vc->parmv1 ? 1 : v
c->parmv1;
759:                                     csi_P(vc, vc->parmv1);
760:                                     CSE;
761:                                     continue;
762:                                     case 'S': /* Scroll Up <ESC>[ n S
*/
763:                                     vc->parmv1 = !vc->parmv1 ? 1 : v
c->parmv1;
764:                                     while(vc->parmv1--) {
765:                                         scroll_screen(vc, 0, SCR
OLL_UP);
766:                                     }
767:                                     CSE;
768:                                     continue;
769:                                     case 'T': /* Scroll Down <ESC>[ n
T */
770:                                     vc->parmv1 = !vc->parmv1 ? 1 : v
c->parmv1;
771:                                     while(vc->parmv1--) {
772:                                         scroll_screen(vc, 0, SCR
OLL_DOWN);
773:                                     }
774:                                     CSE;
775:                                     continue;
776:                                     case 'X': /* Erase Character(s) <E
SC>[ n X */
777:                                     vc->parmv1 = !vc->parmv1 ? 1 : v
c->parmv1;
778:                                     csi_X(vc, vc->parmv1);
779:                                     CSE;
780:                                     continue;
781:                                     case 'c': /* Query Device Code <ES
C>[c */
782:                                     if(!numeric) {
783:                                         insert_seq(tty, VT100ID,
7);
784:                                     }
785:                                     CSE;
786:                                     continue;
787:                                     case 'd': /* Cursor Vertical Posit
ion <ESC>[ n d */
788:                                     vc->parmv1 = vc->parmv1 ? vc->pa
rmv1 - 1 : vc->parmv1;
789:                                     adjust(vc, vc->x, vc->parmv1);
790:                                     CSE;
791:                                     continue;
792:                                     case 'g':
793:                                         switch(vc->parmv1) {
794:                                             case 0: /* Clear Tab <ES
C>[g */
795:                                                 vc->tty->tab_sto
p[vc->x] = 0;
796:                                                 break;
797:                                             case 3: /* Clear All Tab
<ESC>[3g */
798:                                                 case 5: /* Clear All Tab
<ESC>[5g */
799:                                                     for(n = 0; n < M
AX_TAB_COLS; n++)
800:                                                         vc->tty-
>tab_stop[n] = 0;
801:                                                     break;

```

## drivers/char/console.c

Page 14/21

```

802:                                     }
803:                                     CSE;
804:                                     continue;
805:     case 'h':
806:         if(vc->question) {
807:             switch(vc->parmv1) {
808:                 /* DEC modes */
809:                 case 25: /* Swit
ch Cursor Visible <ESC>[?25h */
810:                                     show_cur
811:                                     break;
812:                 case 4:
813:                                     vc->inse
814:                                     break;
815:                                     }
816:             }
817:             CSE;
818:             continue;
819:     case 'l':
820:         if(vc->question) {
821:             switch(vc->parmv1) {
822:                 /* DEC modes */
823:                 case 25: /* Swit
ch Cursor Invisible <ESC>[?25l */
824:                                     show_cur
825:                                     break;
826:                 case 4:
827:                                     vc->inse
828:                                     break;
829:                                     }
830:             }
831:             CSE;
832:             continue;
833:     case 'm': /* Select Graphic Rendit
ion <ESC> n ... m */
834:                                     vc->nparms++;
835:                                     csi_m(vc);
836:                                     CSE;
837:                                     continue;
838:     case 'n':
839:         if(!vc->question) {
840:             switch(vc->parmv1) {
841:                 case 5: /* Query
Device Status <ESC>[5n */
842:                                     insert_s
843:                                     break;
844:                 case 6: /* Query
Cursor Position <ESC>[6n */
845:                                     {
846:
char curpos[8];
847:
char len;
848:
len = sprintf(curpos, "\033[%d;%dR", vc->y, vc->x);
849:
insert_seq(tty, curpos, len);
850:                                     }
851:                                     break;
852:                                     }
853:             }
854:             CSE;

```



## drivers/char/console.c

Page 15/21

```

855:                                     continue;
856:                                     case 'r': /* Top and Bottom Margin
s <ESC>[r / <ESC>[start];end)r */
857:                                     if(!vc->parmv1) {
858:                                         vc->parmv1++;
859:                                     }
860:                                     if(!vc->parmv2) {
861:                                         vc->parmv2 = SCREEN_LINE
S;
862:                                     }
863:                                     if(vc->parmv1 < vc->parmv2 && vc
->parmv2 <= SCREEN_LINES) {
864:                                         vc->top = vc->parmv1 - 1
;
865:                                         vc->bottom = vc->parmv2;
866:                                         adjust(vc, 0, 0);
867:                                     }
868:                                     CSE;
869:                                     continue;
870:                                     case 's': /* Save Cursor <ESC>[s *
/
871:                                         vc->saved_x = vc->x;
872:                                         vc->saved_y = vc->y;
873:                                         CSE;
874:                                         continue;
875:                                     case 'u': /* Restore Cursor <ESC>[
u */
876:                                         vc->x = vc->saved_x;
877:                                         vc->y = vc->saved_y;
878:                                         CSE;
879:                                         continue;
880:                                     default:
881:                                         CSE;
882:                                         break;
883:                                     }
884:                                     } else {
885:                                         switch(ch) {
886:                                             case '[':
887:                                                 vc->sbracket = 1;
888:                                                 vc->semicolon = 0;
889:                                                 vc->question = 0;
890:                                                 vc->parmv1 = vc->parmv2 = 0;
891:                                                 vc->nparms = 0;
892:                                                 memset_b(vc->parms, 0, sizeof(vc
->parms));
893:                                             continue;
894:                                             case '7': /* Save Cursor & Attrs <
ESC>7 */
895:                                                 vc->saved_x = vc->x;
896:                                                 vc->saved_y = vc->y;
897:                                                 CSE;
898:                                                 continue;
899:                                             case '8': /* Restore Cursor & Attr
s <ESC>8 */
900:                                                 vc->x = vc->saved_x;
901:                                                 vc->y = vc->saved_y;
902:                                                 CSE;
903:                                                 continue;
904:                                             case 'D': /* Scroll Up <ESC>D */
905:                                                 lf(vc);
906:                                                 CSE;
907:                                                 continue;
908:                                             case 'E': /* Move To Next Line <ES
C>E */
909:                                                 cr(vc);
910:                                                 lf(vc);
911:                                                 CSE;

```

## drivers/char/console.c

Page 16/21

```

912:                                     continue;
913:     case 'H': /* Set Tab <ESC>H */
914:         vc->tty->tab_stop[vc->x] = 1;
915:         CSE;
916:         continue;
917:     case 'M': /* Scroll Down <ESC>M */
918:         ri(vc);
919:         CSE;
920:         continue;
921:     case 'Z': /* Identify Terminal <ES
C>Z */
922:         insert_seq(tty, VT100ID, 7);
923:         CSE;
924:         continue;
925:     case 'c': /* Reset Device <ESC>c *
/
926:         vconsole_reset(vc->tty);
927:         vc->x = vc->y = 0;
928:         csi_J(vc, CSI_J_SCREEN);
929:         CSE;
930:         continue;
931:     default:
932:         CSE;
933:         break;
934:     }
935: }
936: }
937: switch(ch) {
938:     case '\033':
939:         vc->esc = 1;
940:         vc->sbracket = 0;
941:         vc->semicolon = 0;
942:         vc->question = 0;
943:         vc->parmv1 = vc->parmv2 = 0;
944:         continue;
945:     default:
946:         echo_char(vc, &ch, 1);
947:         continue;
948: }
949: }
950: if(vc->vc_mode != KD_GRAPHICS) {
951:     update_curpos(vc);
952: }
953: wakeup(&tty_write);
954: }
955:
956: void vconsole_select(int new_cons)
957: {
958:     new_cons++;
959:     if(current_cons != new_cons) {
960:         if(vc[current_cons].vt_mode.mode == VT_PROCESS) {
961:             if(!kill_pid(vc[current_cons].tty->pid, vc[current_cons]
.vt_mode.acqsig)) {
962:                 vc[current_cons].switchto_tty = new_cons;
963:                 return;
964:             }
965:             init_vt(&vc[current_cons]);
966:         }
967:         if(vc[current_cons].vc_mode == KD_GRAPHICS) {
968:             return;
969:         }
970:         vconsole_select_final(new_cons);
971:     }
972: }
973:
974: void vconsole_select_final(int new_cons)
975: {

```

## drivers/char/console.c

Page 17/21

```

976:         if(current_cons != new_cons) {
977:             if(vc[new_cons].vt_mode == VT_PROCESS) {
978:                 if(kill_pid(vc[new_cons].tty->pid, vc[new_cons].vt_mode.
acqsig)) {
979:                     init_vt(&vc[new_cons]);
980:                 }
981:             }
982:             if(buf_top) {
983:                 vconsole_restore(&vc[current_cons]);
984:                 buf_top = 0;
985:                 update_curpos(&vc[current_cons]);
986:             }
987:             if(vc[current_cons].vc_mode != KD_GRAPHICS) {
988:                 vconsole_save(&vc[current_cons]);
989:             }
990:             vc[current_cons].vidmem = vc[current_cons].scrbuf;
991:             vc[current_cons].has_focus = 0;
992:             vc[new_cons].vidmem = VIDEO_ADDR;
993:             vc[new_cons].has_focus = 1;
994:             vconsole_restore(&vc[new_cons]);
995:             current_cons = new_cons;
996:             set_leds(vc[current_cons].led_status);
997:             update_curpos(&vc[current_cons]);
998:
999:             buf_y = vc[current_cons].y;
1000:            buf_top = 0;
1001:            memset_w(vcbuf, BLANK_MEM, VC_BUF_SIZE / sizeof(short int));
1002:            memcpy_w(vcbuf, vc[current_cons].vidmem, SCREEN_SIZE / sizeof(sh
ort int));
1003:                show_cursor(ON);
1004:            }
1005: }
1006:
1007: void vconsole_save(struct vconsole *vc)
1008: {
1009:     memcpy_b(vc->scrbuf, vc->vidmem, SCREEN_SIZE);
1010: }
1011:
1012: void vconsole_restore(struct vconsole *vc)
1013: {
1014:     memcpy_b(vc->vidmem, vc->scrbuf, SCREEN_SIZE);
1015: }
1016:
1017: void vconsole_buffer_srl(int mode)
1018: {
1019:     if(buf_y <= SCREEN_LINES) {
1020:         return;
1021:     }
1022:     if(mode == SCROLL_UP) {
1023:         if(buf_top < 0) {
1024:             return;
1025:         }
1026:         if(!buf_top) {
1027:             vconsole_save(&vc[current_cons]);
1028:             buf_top = (buf_y - SCREEN_LINES + 1) * SCREEN_COLS;
1029:             buf_top -= (SCREEN_LINES / 2) * SCREEN_COLS;
1030:         } else {
1031:             buf_top -= (SCREEN_LINES / 2) * SCREEN_COLS;
1032:         }
1033:         if(buf_top < 0) {
1034:             buf_top = 0;
1035:         }
1036:         memcpy_b(vc[current_cons].vidmem, vcbuf + buf_top, SCREEN_SIZE);
1037:         if(!buf_top) {
1038:             buf_top = -1;
1039:         }
1040:         show_cursor(OFF);

```

## drivers/char/console.c

Page 18/21

```

1041:         return;
1042:     }
1043:     if(mode == SCROLL_DOWN) {
1044:         if(!buf_top) {
1045:             return;
1046:         }
1047:         if(buf_top == buf_y * SCREEN_COLS) {
1048:             return;
1049:         }
1050:         if(buf_top < 0) {
1051:             buf_top = 0;
1052:         }
1053:         buf_top += (SCREEN_LINES / 2) * SCREEN_COLS;
1054:         if(buf_top >= (buf_y - SCREEN_LINES + 1) * SCREEN_COLS) {
1055:             vconsole_restore(&vc[current_cons]);
1056:             buf_top = 0;
1057:             show_cursor(ON);
1058:             update_curpos(&vc[current_cons]);
1059:             return;
1060:         }
1061:         memcpy_b(vc[current_cons].vidmem, vcbuf + buf_top, SCREEN_SIZE);
1062:         return;
1063:     }
1064: }
1065:
1066: void blank_screen(struct vconsole *vc)
1067: {
1068:     if(vc->blanked) {
1069:         return;
1070:     }
1071:     vconsole_save(vc);
1072:     memset_w(vc->vidmem, BLANK_MEM, SCREEN_SIZE / sizeof(short int));
1073:     vc->blanked = 1;
1074:     show_cursor(OFF);
1075: }
1076:
1077: void unblank_screen(struct vconsole *vc)
1078: {
1079:     if(!vc->blanked) {
1080:         return;
1081:     }
1082:     vconsole_restore(vc);
1083:     vc->blanked = 0;
1084:     show_cursor(ON);
1085: }
1086:
1087: void screen_on(void)
1088: {
1089:     unsigned long int flags;
1090:     struct callout_req creq;
1091:
1092:     if(screen_is_off) {
1093:         SAVE_FLAGS(flags); CLI();
1094:         inport_b(INPUT_STAT1);
1095:         inport_b(0x3BA);
1096:         outport_b(ATTR_CONTROLLER, ATTR_CONTROLLER_PAS);
1097:         RESTORE_FLAGS(flags);
1098:     }
1099:     creq.fn = screen_off;
1100:     creq.arg = 0;
1101:     add_callout(&creq, BLANK_INTERVAL);
1102: }
1103:
1104: void screen_off(unsigned int arg)
1105: {
1106:     unsigned long int flags;
1107:

```

## drivers/char/console.c

Page 19/21

```

1108:     screen_is_off = 1;
1109:     SAVE_FLAGS(flags); CLI();
1110:     inport_b(INPUT_STAT1);
1111:     inport_b(0x3BA);
1112:     outport_b(ATTR_CONTROLLER, 0);
1113:     RESTORE_FLAGS(flags);
1114: }
1115:
1116: void vconsole_start(struct tty *tty)
1117: {
1118:     struct vconsole *vc;
1119:
1120:     vc = (struct vconsole *)tty->driver_data;
1121:     if(!vc->scrlock) {
1122:         return;
1123:     }
1124:     vc->led_status &= ~SCLBIT;
1125:     vc->scrlock = 0;
1126:     set_leds(vc->led_status);
1127: }
1128:
1129: void vconsole_stop(struct tty *tty)
1130: {
1131:     struct vconsole *vc;
1132:
1133:     vc = (struct vconsole *)tty->driver_data;
1134:     if(vc->scrlock) {
1135:         return;
1136:     }
1137:     vc->led_status |= SCLBIT;
1138:     vc->scrlock = 1;
1139:     set_leds(vc->led_status);
1140: }
1141:
1142: void vconsole_beep(void)
1143: {
1144:     struct callout_req creq;
1145:
1146:     pit_beep_on();
1147:     creq.fn = pit_beep_off;
1148:     creq.arg = 0;
1149:     add_callout(&creq, HZ / 8);
1150: }
1151:
1152: void vconsole_deltab(struct tty *tty)
1153: {
1154:     int col, n;
1155:     unsigned char count;
1156:     struct vconsole *vc;
1157:     struct cblock *cb;
1158:     unsigned char ch;
1159:
1160:     vc = (struct vconsole *)tty->driver_data;
1161:     cb = tty->cooked_q.head;
1162:     col = count = 0;
1163:
1164:     while(cb) {
1165:         for(n = 0; n < cb->end_off; n++) {
1166:             if(n >= cb->start_off) {
1167:                 ch = cb->data[n];
1168:                 if(ch == '\t') {
1169:                     while(!vc->tty->tab_stop[++col]);
1170:                 } else {
1171:                     col++;
1172:                     if(ISCNTRL(ch) && !ISSPACE(ch) && tty->t
ermios.c_lflag & ECHOCTL) {
1173:                         col++;

```

## drivers/char/console.c

Page 20/21

```

1174:                                     }
1175:                                     }
1176:                                     col %= vc->columns;
1177:                                 }
1178:                             }
1179:                             cb = cb->next;
1180:                         }
1181:                         count = vc->x - col;
1182:
1183:                         while(count--) {
1184:                             tty_queue_putchar(tty, &tty->write_q, '\b');
1185:                         }
1186: }
1187:
1188: void console_flush_log_buf(char *buffer, unsigned int count)
1189: {
1190:     char *b;
1191:     struct tty *tty;
1192:
1193:     if(!(tty = get_tty(_syscondev))) {
1194:         _syscondev = MKDEV(VCONSOLES_MAJOR, 0);
1195:         tty = get_tty(_syscondev);
1196:     }
1197:     b = buffer;
1198:
1199:     while(count) {
1200:         if(tty_queue_putchar(tty, &tty->write_q, *b) < 0) {
1201:             tty->output(tty);
1202:             if(!is_vconsole(tty->dev)) {
1203:                 if(sleep(&tty_write, PROC_INTERRUPTIBLE)) {
1204:                     break;
1205:                 }
1206:             }
1207:             continue;
1208:         }
1209:         count--;
1210:         b++;
1211:     }
1212:     tty->output(tty);
1213: }
1214:
1215: /* get the VGA type from the BIOS equipment information */
1216: void get_video_parms(void)
1217: {
1218:     short int *bios_data;
1219:     memset_b(&video, 0, sizeof(struct video_parms));
1220:
1221:     bios_data = (short int *)0x410;
1222:     if((*bios_data & 0x30) == 0x30) {
1223:         /* monochrome = 0x30 */
1224:         VIDEO_ADDR = (void *)MONO_ADDR;
1225:         VIDEO_PORT = MONO_6845_ADDR;
1226:         VIDEO_TYPE = "VGA monochrome 80x25";
1227:     } else {
1228:         /* color = 0x00 || 0x20 */
1229:         VIDEO_ADDR = (void *)COLOR_ADDR;
1230:         VIDEO_PORT = COLOR_6845_ADDR;
1231:         VIDEO_TYPE = "VGA color 80x25";
1232:     }
1233:
1234:     SCREEN_COLS = 80;
1235:     SCREEN_LINES = 25;
1236: }
1237:
1238: void vconsole_init(void)
1239: {
1240:     int n;

```

## drivers/char/console.c

Page 21/21

```

1241:         struct tty *tty;
1242:
1243:         printk("console 0x%04X-0x%04X - %s", VIDEO_PORT, VIDEO_PORT + 1,
VIDEO_TYPE);
1244:         printk(" (%d virtual consoles)\n", NR_VCONSOLES);
1245:         screen_on();
1246:
1247:         for(n = 1; n <= NR_VCONSOLES; n++) {
1248:             if(!register_tty(MKDEV(VCONSOLES_MAJOR, n)) {
1249:                 tty = get_tty(MKDEV(VCONSOLES_MAJOR, n));
1250:                 tty->driver_data = (void *)&vc[n];
1251:                 tty->stop = vconsole_stop;
1252:                 tty->start = vconsole_start;
1253:                 tty->deltab = vconsole_deltab;
1254:                 tty->reset = vconsole_reset;
1255:                 tty->input = do_cook;
1256:                 tty->output = vconsole_write;
1257:                 vc[n].tty = tty;
1258:                 vc[n].scrbuf = (short int *)kmalloc();
1259:                 vc[n].vidmem = vc[n].scrbuf;
1260:                 memset_w(vc[n].scrbuf, BLANK_MEM, SCREEN_SIZE / sizeof(s
hort int));
1261:                 vconsole_reset(tty);
1262:                 tty_queue_init(tty);
1263:             }
1264:         }
1265:         current_cons = 1;
1266:         vc[current_cons].vidmem = VIDEO_ADDR;
1267:         vc[current_cons].has_focus = 1;
1268:
1269:         memcpy_b(vcbuf, vc[current_cons].vidmem, SCREEN_SIZE);
1270:
1271:         get_curpos(&vc[current_cons]);
1272:         update_curpos(&vc[current_cons]);
1273:         buf_y = vc[current_cons].y;
1274:         buf_top = 0;
1275:
1276:         SET_MINOR(console_device.minors, 0);
1277:         SET_MINOR(console_device.minors, 1);
1278:         for(n = 0; n <= NR_VCONSOLES; n++) {
1279:             SET_MINOR(tty_device.minors, n);
1280:         }
1281:
1282:         register_device(CHR_DEV, &console_device);
1283:         register_device(CHR_DEV, &tty_device);
1284:
1285:         if(is_vconsole(_syscondev)) {
1286:             register_console(console_flush_log_buf);
1287:         }
1288:     }

```

**drivers/char/defkeymap.c**

```

1: /*
2:  * fiwix/drivers/char/defkeymap.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/keyboard.h>
9: #include <fiwix/string.h>
10:
11: #define BS      127      /* backspace */
12:
13: __key_t keymap[NR_MODIFIERS * NR_CODES] = {
14: /*
15:  * Standard US keyboard (default keymap) with 16 modifiers
16:  *
17:  *
18:  *
19:  *
20:  * ===== */
21: /* 00 - NULL */
22: /* 01 - ESC */
23: /* 02 - 1 */
24: /* 03 - 2 */
25: /* 04 - 3 */
26: /* 05 - 4 */
27: /* 06 - 5 */
28: /* 07 - 6 */
29: /* 08 - 7 */
30: /* 09 - 8 */
31: /* 10 - 9 */
32: /* 11 - 0 */
33: /* 12 - _ */
34: /* 13 - += */
35: /* 14 - BS */
36: /* 15 - TAB */
37: /* 16 - q */
38: /* 17 - w */
39: /* 18 - e */
40: /* 19 - r */
41: /* 20 - t */

```

| 0, | 0,       | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
|----|----------|----|----|----|----|----|----|----|----|----|
| 0, | A('['),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('1'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('2'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('3'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('4'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('5'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('6'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('7'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('8'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('9'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('0'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('-'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('='),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A(BS),   | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('\t'), | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('q'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('w'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('e'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('r'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |
| 0, | A('t'),  | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, | 0, |





| drivers/char/defkeymap.c |                |    |        |        |        |        |        |        |    |    | Page 3/5 |
|--------------------------|----------------|----|--------|--------|--------|--------|--------|--------|----|----|----------|
| 0,                       | RSHIFT,        | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 76:                      | /* 55 - *      | */ | ASTSK, | ASTSK, | ASTSK, | ASTSK, | ASTSK, | ASTSK, | 0, | 0, | 0,       |
| 0,                       | ASTSK,         | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 77:                      | /* 56 - ALT    | */ | ALT,   | ALT,   | ALT,   | ALT,   | ALT,   | ALT,   | 0, | 0, | 0,       |
| 0,                       | ALT,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 78:                      | /* 57 - SPC    | */ | ' ' ,  | ' ' ,  | NULL,  | NULL,  | NULL,  | NULL,  | 0, | 0, | 0,       |
| 0,                       | A(' '),        | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 79:                      | /* 58 - CAPS   | */ | CAPS,  | CAPS,  | CAPS,  | CAPS,  | CAPS,  | CAPS,  | 0, | 0, | 0,       |
| 0,                       | CAPS,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 80:                      | /* 59 - F1     | */ | F1,    | SF1,   | NULL,  | NULL,  | F1,    | 0,     | 0, | 0, | 0,       |
| 0,                       | AF1,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 81:                      | /* 60 - F2     | */ | F2,    | SF2,   | NULL,  | NULL,  | F2,    | 0,     | 0, | 0, | 0,       |
| 0,                       | AF2,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 82:                      | /* 61 - F3     | */ | F3,    | SF3,   | NULL,  | NULL,  | F3,    | 0,     | 0, | 0, | 0,       |
| 0,                       | AF3,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 83:                      | /* 62 - F4     | */ | F4,    | SF4,   | NULL,  | NULL,  | F4,    | 0,     | 0, | 0, | 0,       |
| 0,                       | AF4,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 84:                      | /* 63 - F5     | */ | F5,    | SF5,   | NULL,  | NULL,  | F5,    | 0,     | 0, | 0, | 0,       |
| 0,                       | AF5,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 85:                      | /* 64 - F6     | */ | F6,    | SF6,   | NULL,  | NULL,  | F6,    | 0,     | 0, | 0, | 0,       |
| 0,                       | AF6,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 86:                      | /* 65 - F7     | */ | F7,    | SF7,   | NULL,  | NULL,  | F7,    | 0,     | 0, | 0, | 0,       |
| 0,                       | AF7,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 87:                      | /* 66 - F8     | */ | F8,    | SF8,   | NULL,  | NULL,  | F8,    | 0,     | 0, | 0, | 0,       |
| 0,                       | AF8,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 88:                      | /* 67 - F9     | */ | F9,    | SF9,   | NULL,  | NULL,  | F9,    | 0,     | 0, | 0, | 0,       |
| 0,                       | AF9,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 89:                      | /* 68 - F10    | */ | F10,   | SF10,  | NULL,  | NULL,  | F10,   | 0,     | 0, | 0, | 0,       |
| 0,                       | AF10,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 90:                      | /* 69 - NUMS   | */ | NUMS,  | NUMS,  | NUMS,  | NUMS,  | NUMS,  | 0,     | 0, | 0, | 0,       |
| 0,                       | NUMS,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 91:                      | /* 70 - SCRL   | */ | SCRL,  | SCRL3, | SCRL2, | NULL,  | SCRL4, | 0,     | 0, | 0, | 0,       |
| 0,                       | SCRL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 92:                      | /* 71 - HOME/7 | */ | HOME,  | HOME,  | HOME,  | HOME,  | HOME,  | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 93:                      | /* 72 - UP /8  | */ | UP,    | UP,    | UP,    | UP,    | UP,    | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 94:                      | /* 73 - PGUP/9 | */ | PGUP,  | PGUP,  | PGUP,  | PGUP,  | PGUP,  | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 95:                      | /* 74 - MINUS  | */ | MINUS, | MINUS, | MINUS, | MINUS, | MINUS, | 0,     | 0, | 0, | 0,       |
| 0,                       | MINUS,         | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 96:                      | /* 75 - LEFT/4 | */ | LEFT,  | LEFT,  | LEFT,  | LEFT,  | LEFT,  | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 97:                      | /* 76 - MID /5 | */ | MID,   | MID,   | MID,   | MID,   | MID,   | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 98:                      | /* 77 - RIGH/6 | */ | RIGHT, | RIGHT, | RIGHT, | RIGHT, | RIGHT, | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 99:                      | /* 78 - PLUS   | */ | PLUS,  | PLUS,  | PLUS,  | PLUS,  | PLUS,  | 0,     | 0, | 0, | 0,       |
| 0,                       | PLUS,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 100:                     | /* 79 - END /1 | */ | END,   | END,   | END,   | END,   | END,   | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 101:                     | /* 80 - DOWN/2 | */ | DOWN,  | DOWN,  | DOWN,  | DOWN,  | DOWN,  | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 102:                     | /* 81 - PGDN/3 | */ | PGDN,  | PGDN,  | PGDN,  | PGDN,  | PGDN,  | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 103:                     | /* 82 - INS /0 | */ | INS,   | INS,   | INS,   | INS,   | INS,   | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 104:                     | /* 83 - DEL /. | */ | DEL,   | DEL,   | DEL,   | DEL,   | DEL,   | 0,     | 0, | 0, | 0,       |
| 0,                       | DEL,           | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 105:                     | /* 84 -        | */ | NULL,  | NULL,  | NULL,  | NULL,  | NULL,  | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 106:                     | /* 85 -        | */ | NULL,  | NULL,  | NULL,  | NULL,  | NULL,  | 0,     | 0, | 0, | 0,       |
| 0,                       | NULL,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 107:                     | /* 86 - <>     | */ | '<' ,  | '>' ,  | ' ' ,  | NULL,  | NULL,  | 0,     | 0, | 0, | 0,       |
| 0,                       | A('<'),        | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |
| 108:                     | /* 87 - F11    | */ | SF1,   | SF1,   | NULL,  | NULL,  | F11,   | 0,     | 0, | 0, | 0,       |
| 0,                       | AF11,          | 0, | 0,     | 0,     | 0,     | 0,     | 0,     | 0,     | 0, | 0, | 0,       |

| drivers/char/defkeymap.c |          |     |    |           |          |          |          |          |          | Page 4/5 |    |    |
|--------------------------|----------|-----|----|-----------|----------|----------|----------|----------|----------|----------|----|----|
| 109:                     | /*       | 88  | -  | F12       | */       | SF2,     | SF2,     | NULL,    | NULL,    | F12,     | 0, | 0, |
| 0,                       | AF12,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 110:                     | /*       | 89  | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 111:                     | /*       | 90  | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 112:                     | /*       | 91  | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 113:                     | /*       | 92  | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 114:                     | /*       | 93  | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 115:                     | /*       | 94  | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 116:                     | /*       | 95  | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 117:                     | /*       | 96  | -  | EOENTER*/ | EOENTER, | EOENTER, | EOENTER, | EOENTER, | EOENTER, | 0,       | 0, | 0, |
| 0,                       | EOENTER, | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 118:                     | /*       | 97  | -  | RCTRL     | */       | RCTRL,   | RCTRL,   | RCTRL,   | RCTRL,   | RCTRL,   | 0, | 0, |
| 0,                       | RCTRL,   | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 119:                     | /*       | 98  | -  | EOSLASH*/ | EOSLASH, | EOSLASH, | EOSLASH, | EOSLASH, | EOSLASH, | 0,       | 0, | 0, |
| 0,                       | EOSLASH, | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 120:                     | /*       | 99  | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | C('\\"), | 0,       | 0, | 0, |
| 0,                       | C('\\"), | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 121:                     | /*       | 100 | -  | ALTGR     | */       | ALTGR,   | ALTGR,   | ALTGR,   | ALTGR,   | ALTGR,   | 0, | 0, |
| 0,                       | ALTGR,   | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 122:                     | /*       | 101 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 123:                     | /*       | 102 | -  | EOHOME    | */       | EOHOME,  | EOHOME,  | EOHOME,  | EOHOME,  | EOHOME,  | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 124:                     | /*       | 103 | -  | EOUP      | */       | EOUP,    | EOUP,    | EOUP,    | EOUP,    | EOUP,    | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 125:                     | /*       | 104 | -  | EOPGUP    | */       | EOPGUP,  | EOPGUP,  | EOPGUP,  | EOPGUP,  | EOPGUP,  | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 126:                     | /*       | 105 | -  | EOLEFT    | */       | EOLEFT,  | EOLEFT,  | EOLEFT,  | EOLEFT,  | EOLEFT,  | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 127:                     | /*       | 106 | -  | EORIGHT*/ | EORIGHT, | EORIGHT, | EORIGHT, | EORIGHT, | EORIGHT, | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 128:                     | /*       | 107 | -  | EOEND     | */       | EOEND,   | EOEND,   | EOEND,   | EOEND,   | EOEND,   | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 129:                     | /*       | 108 | -  | EODOWN    | */       | EODOWN,  | EODOWN,  | EODOWN,  | EODOWN,  | EODOWN,  | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 130:                     | /*       | 109 | -  | EOPGDN    | */       | EOPGDN,  | EOPGDN,  | EOPGDN,  | EOPGDN,  | EOPGDN,  | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 131:                     | /*       | 110 | -  | EOINS     | */       | EOINS,   | EOINS,   | EOINS,   | EOINS,   | EOINS,   | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 132:                     | /*       | 111 | -  | EODEL     | */       | EODEL,   | EODEL,   | EODEL,   | EODEL,   | EODEL,   | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 133:                     | /*       | 112 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 134:                     | /*       | 113 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 135:                     | /*       | 114 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 136:                     | /*       | 115 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 137:                     | /*       | 116 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 138:                     | /*       | 117 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 139:                     | /*       | 118 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 140:                     | /*       | 119 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 141:                     | /*       | 120 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |
| 0,                       | NULL,    | 0,  | 0, | 0,        | 0,       | 0,       | 0,       | 0,       | 0,       | 0,       | 0, | 0, |
| 142:                     | /*       | 121 | -  | */        | NULL,    | NULL,    | NULL,    | NULL,    | NULL,    | 0,       | 0, | 0, |

**drivers/char/defkeymap.c**

Page 5/5

```

0,      NULL,  0,    0,    0,    0,    0,    0,    0,
143: /* 122 -    */    NULL,  NULL,  NULL,  NULL,  NULL,  0,    0,
0,      NULL,  0,    0,    0,    0,    0,    0,    0,
144: /* 123 -    */    NULL,  NULL,  NULL,  NULL,  NULL,  0,    0,
0,      NULL,  0,    0,    0,    0,    0,    0,    0,
145: /* 124 -    */    NULL,  NULL,  NULL,  NULL,  NULL,  0,    0,
0,      NULL,  0,    0,    0,    0,    0,    0,    0,
146: /* 125 -    */    NULL,  NULL,  NULL,  NULL,  NULL,  0,    0,
0,      NULL,  0,    0,    0,    0,    0,    0,    0,
147: /* 126 -    */    NULL,  NULL,  NULL,  NULL,  NULL,  0,    0,
0,      NULL,  0,    0,    0,    0,    0,    0,    0,
148: /* 127 -    */    NULL,  NULL,  NULL,  NULL,  NULL,  0,    0,
0,      NULL,  0,    0,    0,    0,    0,    0,    0,
149: };

```

## drivers/char/keyboard.c

Page 1/12

```

1: /*
2:  * fiwix/drivers/char/keyboard.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/limits.h>
11: #include <fiwix/keyboard.h>
12: #include <fiwix/console.h>
13: #include <fiwix/pic.h>
14: #include <fiwix/signal.h>
15: #include <fiwix/process.h>
16: #include <fiwix/sleep.h>
17: #include <fiwix/kd.h>
18: #include <fiwix/stdio.h>
19: #include <fiwix/string.h>
20:
21: #define KB_DATA          0x60      /* I/O data port */
22: #define KBC_COMMAND     0x64      /* command/control port */
23: #define KBC_STATUS      0x64      /* status register port */
24:
25: /*
26:  * PS/2 System Control Port A
27:  * -----
28:  * bit 7 -> fixed disk activity led
29:  * bit 6 -> fixed disk activity led
30:  * bit 5 -> reserved
31:  * bit 4 -> watchdog timer status
32:  * bit 3 -> security lock latch
33:  * bit 2 -> reserved
34:  * bit 1 -> alternate gate A20
35:  * bit 0 -> alternate hot reset
36:  */
37: #define PS2_SYSCTRL_A   0x92      /* PS/2 system control port A (write) */
38:
39: #define KB_CMD_RESET    0xFF      /* keyboard reset */
40: #define KB_CMD_ENABLE   0xF4      /* keyboard enable scanning */
41: #define KB_CMD_DISABLE  0xF5      /* keyboard disable scanning */
42: #define KB_CMD_IDENTIFY 0xF2      /* keyboard identify (for PS/2 only) */
43: #define KB_CMD_ECHO     0xEE      /* echo (for diagnostics only) */
44:
45: #define KBC_CMD_RECV_CONFIG 0x20   /* read controller's config byte */
46: #define KBC_CMD_SEND_CONFIG 0x60   /* write controller's config byte */
47: #define KBC_CMD_SELF_TEST  0xAA    /* self-test command */
48: #define KBC_CMD_PS2_1_TEST 0xAB    /* first PS/2 interface test command */
49: #define KBC_CMD_PS2_2_TEST 0xA9    /* second PS/2 interface test command */
50: #define KBC_CMD_DISABLE_PS2_1 0xAD  /* disable first PS/2 port */
51: #define KBC_CMD_ENABLE_PS2_1 0xAE   /* enable first PS/2 port */
52: #define KBC_CMD_DISABLE_PS2_2 0xA7  /* disable second PS/2 port (if any) */
53: #define KBC_CMD_ENABLE_PS2_2 0xA8   /* enable second PS/2 port (if any) */
54: #define KBC_CMD_GET_IFACE  0xCA    /* get interface type (AT or MCA) */
55: #define KBC_CMD_HOTRESET   0xFE    /* Hot Reset */
56:
57: /* flags of the status register */
58: #define KB_STR_OUTBUSY   0x01      /* output buffer full, don't read yet */
59: #define KB_STR_INBUSY   0x02      /* input buffer full, don't write yet */
60: #define KB_STR_TXTMOUT   0x20      /* transmit time-out error */
61: #define KB_STR_RXTMOUT   0x40      /* receive time-out error */
62: #define KB_STR_PARERR    0x80      /* parity error */
63: #define KB_STR_COMMERR   (KB_STR_TXTMOUT | KB_STR_RXTMOUT)
64:
65: #define KB_RESET_OK     0xAA      /* self-test passed */
66: #define KB_ACK          0xFA      /* acknowledge */
67: #define KB_SETTLED      0xED      /* set/reset status indicators (LEDs) */

```

## drivers/char/keyboard.c

Page 2/12

```

68: #define KB_RATE          0xF3      /* set typematic rate/delay */
69: #define DELAY_250        0x00      /* typematic delay at 250ms (default) */
70: #define DELAY_500        0x40      /* typematic delay at 500ms */
71: #define DELAY_750        0x80      /* typematic delay at 750ms */
72: #define DELAY_1000       0xC0      /* typematic delay at 1000ms */
73: #define RATE_30          0x00      /* typematic rate at 30.0 reports/sec (default)
*/
74:
75: #define EXTKEY            0xE0      /* extended key (AltGr, Ctrl-Print, etc.) */
76:
77: __key_t *keymap_line;
78:
79: static unsigned char e0_keys[128] = {
80:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x00-0x07 */
81:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x08-0x0f */
82:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x10-0x17 */
83:     0, 0, 0, 0, E0ENTER, RCTRL, 0, 0, /* 0x18-0x1f */
84:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x20-0x27 */
85:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x28-0x2f */
86:     0, 0, 0, 0, 0, E0SLASH, 0, 0, /* 0x30-0x37 */
87:     ALTGR, 0, 0, 0, 0, 0, 0, 0, /* 0x38-0x3f */
88:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x40-0x47 */
89:     E0UP, E0GUP, 0, E0LEFT, 0, E0RIGHT, 0, E0END, /* 0x48-0x4f */
90:     E0DOWN, E0PGDN, E0INS, E0DEL, 0, 0, 0, 0, /* 0x50-0x57 */
91:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x58-0x5f */
92:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x60-0x67 */
93:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x68-0x6f */
94:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x70-0x77 */
95:     0, 0, 0, 0, 0, 0, 0, 0, /* 0x78-0x7f */
96: };
97:
98: static unsigned char leds = 0;
99: static unsigned char shift = 0;
100: static unsigned char altgr = 0;
101: static unsigned char ctrl = 0;
102: static unsigned char alt = 0;
103: static unsigned char extkey = 0;
104: static unsigned char deadkey = 0;
105:
106: static unsigned char do_sctl_buf = 0;
107: static char do_switch_console = -1;
108:
109: unsigned char kb_identify[2] = {0, 0};
110: char ps2_active_ports = 0;
111: char ps2_supp_ports = 0;
112: char ps2_iface = 0;
113: short int current_cons;
114: char ctrl_alt_del = 1;
115: char any_key_to_reboot = 0;
116:
117: static struct bh keyboard_bh = { 0, &irq_keyboard_bh, NULL };
118: static struct interrupt irq_config_keyboard = { 0, "keyboard", &irq_keyboard, NU
LL };
119:
120: struct diacritic *diacr;
121: static char *diacr_chars = "`'^\"";
122: struct diacritic grave_table[NR_DIACR] = {
123:     { 'A', '\300' },
124:     { 'E', '\310' },
125:     { 'I', '\314' },
126:     { 'O', '\322' },
127:     { 'U', '\331' },
128:     { 'a', '\340' },
129:     { 'e', '\350' },
130:     { 'i', '\354' },
131:     { 'o', '\362' },
132:     { 'u', '\371' },

```

## drivers/char/keyboard.c

Page 3/12

```

133: };
134: struct diacritic acute_table[NR_DIACR] = {
135:     { 'A', '\301' },
136:     { 'E', '\311' },
137:     { 'I', '\315' },
138:     { 'O', '\323' },
139:     { 'U', '\332' },
140:     { 'a', '\341' },
141:     { 'e', '\351' },
142:     { 'i', '\355' },
143:     { 'o', '\363' },
144:     { 'u', '\372' },
145: };
146: struct diacritic circm_table[NR_DIACR] = {
147:     { 'A', '\302' },
148:     { 'E', '\312' },
149:     { 'I', '\316' },
150:     { 'O', '\324' },
151:     { 'U', '\333' },
152:     { 'a', '\342' },
153:     { 'e', '\352' },
154:     { 'i', '\356' },
155:     { 'o', '\364' },
156:     { 'u', '\373' },
157: };
158: struct diacritic diere_table[NR_DIACR] = {
159:     { 'A', '\304' },
160:     { 'E', '\313' },
161:     { 'I', '\317' },
162:     { 'O', '\326' },
163:     { 'U', '\334' },
164:     { 'a', '\344' },
165:     { 'e', '\353' },
166:     { 'i', '\357' },
167:     { 'o', '\366' },
168:     { 'u', '\374' },
169: };
170:
171: static char *pad_chars = "0123456789+~*/\015,.";
172:
173: static char *pad_seq[] = {
174:     "\033[2~",      /* INS */
175:     "\033[4~",      /* END */
176:     "\033[B",       /* DOWN */
177:     "\033[6~",      /* PGDN */
178:     "\033[D",       /* LEFT */
179:     "\033[G",       /* MID */
180:     "\033[C",       /* RIGHT */
181:     "\033[1~",      /* HOME */
182:     "\033[A",       /* UP */
183:     "\033[5~",      /* PGUP */
184:     "+",            /* PLUS */
185:     "-",            /* MINUS */
186:     "*",            /* ASTERISK */
187:     "/",            /* SLASH */
188:     "'\n'",         /* ENTER */
189:     ",",           /* COMMA */
190:     "\033[3~",      /* DEL */
191: };
192:
193: static char *fn_seq[] = {
194:     "\033[A",       /* F1 */
195:     "\033[B",       /* F2 */
196:     "\033[C",       /* F3 */
197:     "\033[D",       /* F4 */
198:     "\033[E",       /* F5 */
199:     "\033[17~",     /* F6 */

```

## drivers/char/keyboard.c

Page 4/12

```

200:         "\033[18~",      /* F7 */
201:         "\033[19~",      /* F8 */
202:         "\033[20~",      /* F9 */
203:         "\033[21~",      /* F10 */
204:         "\033[23~",      /* F11, SF1 */
205:         "\033[24~",      /* F12, SF2 */
206:         "\033[25~",      /* SF3 */
207:         "\033[26~",      /* SF4 */
208:         "\033[28~",      /* SF5 */
209:         "\033[29~",      /* SF6 */
210:         "\033[31~",      /* SF7 */
211:         "\033[32~",      /* SF8 */
212:         "\033[33~",      /* SF9 */
213:         "\033[34~",      /* SF10 */
214:     };
215:
216:     static void keyboard_delay(void)
217:     {
218:         int n;
219:
220:         for(n = 0; n < 1000; n++) {
221:             NOP();
222:         }
223:     }
224:
225:     /* wait controller input buffer to be clear (ready to write) */
226:     static int keyboard_wait_input(void)
227:     {
228:         int n;
229:
230:         for(n = 0; n < 500000; n++) {
231:             if(!(inport_b(KBC_STATUS) & KB_STR_INBUSY)) {
232:                 return 0;
233:             }
234:         }
235:         return 1;
236:     }
237:
238:     static int keyboard_write(const unsigned char port, const unsigned char byte)
239:     {
240:         if(!keyboard_wait_input()) {
241:             outport_b(port, byte);
242:             if(!keyboard_wait_input()) {
243:                 return 0;
244:             }
245:         }
246:
247:         return 1;
248:     }
249:
250:     /* wait controller output buffer to be full (ready to read) */
251:     static int keyboard_wait_output(void)
252:     {
253:         int n, value;
254:
255:         for(n = 0; n < 500000; n++) {
256:             if((value = inport_b(KBC_STATUS)) & KB_STR_OUTBUSY) {
257:                 if(value & (KB_STR_COMMERR | KB_STR_PARERR)) {
258:                     continue;
259:                 }
260:                 return 0;
261:             }
262:         }
263:         return 1;
264:     }
265:
266:     static int keyboard_wait_ack(void)

```



## drivers/char/keyboard.c

Page 5/12

```

267: {
268:     int n;
269:
270:     keyboard_wait_output();
271:     for(n = 0; n < 1000; n++) {
272:         if(inport_b(KB_DATA) == KB_ACK) {
273:             return 0;
274:         }
275:         keyboard_delay();
276:     }
277:     return 1;
278: }
279:
280: static void keyboard_identify(void)
281: {
282:     /* disable */
283:     keyboard_write(KB_DATA, KB_CMD_DISABLE);
284:     if(keyboard_wait_ack()) {
285:         printk("WARNING: %s(): ACK not received on disable command!\n",
__FUNCTION__);
286:     }
287:
288:     /* identify */
289:     keyboard_write(KB_DATA, KB_CMD_IDENTIFY);
290:     if(keyboard_wait_ack()) {
291:         printk("WARNING: %s(): ACK not received on identify command!\n",
__FUNCTION__);
292:     }
293:     if(!keyboard_wait_output()) {
294:         kb_identify[0] = inport_b(KB_DATA);
295:     }
296:     if(!keyboard_wait_output()) {
297:         kb_identify[1] = inport_b(KB_DATA);
298:     }
299:
300:     /* enable */
301:     keyboard_write(KB_DATA, KB_CMD_ENABLE);
302:     if(keyboard_wait_ack()) {
303:         printk("WARNING: %s(): ACK not received on enable command!\n",
__FUNCTION__);
304:     }
305:     keyboard_wait_output();
306:     inport_b(KB_DATA);
307:
308:     /* get the interface type */
309:     keyboard_write(KBC_COMMAND, KBC_CMD_GET_IFACE);
310:     keyboard_wait_output();
311:     ps2_iface = inport_b(KB_DATA);
312: }
313:
314: static void keyboard_reset(void)
315: {
316:     int errno;
317:     unsigned char config;
318:
319:     /* disable device(s) */
320:     keyboard_write(KBC_COMMAND, KBC_CMD_DISABLE_PS2_1);
321:     keyboard_write(KBC_COMMAND, KBC_CMD_DISABLE_PS2_2);
322:
323:     /* flush buffers */
324:     while(!keyboard_wait_output()) {
325:         inport_b(KB_DATA);
326:     }
327:
328:     /* get controller configuration */
329:     keyboard_write(KBC_COMMAND, KBC_CMD_RECV_CONFIG);
330:     keyboard_wait_output();

```

## drivers/char/keyboard.c

Page 6/12

```

331:     config = inport_b(KB_DATA);
332:     ps2_active_ports = config & 0x01 ? 1 : 0;
333:     ps2_active_ports += config & 0x02 ? 1 : 0;
334:     ps2_supp_ports = 1 + (config & 0x20 ? 1 : 0);
335:
336:     /* set controller configuration (disabling IRQs) */
337:     /*
338:     keyboard_write(KBC_COMMAND, KBC_CMD_SEND_CONFIG);
339:     keyboard_write(KB_DATA, config & ~(0x01 | 0x02 | 0x40));
340:     */
341:
342:     /* PS/2 controller self-test */
343:     keyboard_write(KBC_COMMAND, KBC_CMD_SELF_TEST);
344:     keyboard_wait_output();
345:     if((errno = inport_b(KB_DATA)) != 0x55) {
346:         printk("WARNING: %s(): keyboard returned 0x%x in self-test.\n",
__FUNCTION__, errno);
347:     }
348:
349:     /*
350:     * This sets again the controller configuration since the previous
351:     * step may also reset the PS/2 controller to its power-on defaults.
352:     */
353:     keyboard_write(KBC_COMMAND, KBC_CMD_SEND_CONFIG);
354:     keyboard_write(KB_DATA, config);
355:
356:     /* first PS/2 interface test */
357:     keyboard_write(KBC_COMMAND, KBC_CMD_PS2_1_TEST);
358:     keyboard_wait_output();
359:     if((errno = inport_b(KB_DATA)) != 0) {
360:         printk("WARNING: %s(): keyboard returned 0x%x in first PS/2 inte
rface test.\n", __FUNCTION__, errno);
361:     }
362:
363:     if(ps2_supp_ports > 1) {
364:         /* second PS/2 interface test */
365:         keyboard_write(KBC_COMMAND, KBC_CMD_PS2_2_TEST);
366:         keyboard_wait_output();
367:         if((errno = inport_b(KB_DATA)) != 0) {
368:             printk("WARNING: %s(): keyboard returned 0x%x in second
PS/2 interface test.\n", __FUNCTION__, errno);
369:         }
370:     }
371:
372:     /* enable device(s) */
373:     keyboard_write(KBC_COMMAND, KBC_CMD_ENABLE_PS2_1);
374:     keyboard_write(KBC_COMMAND, KBC_CMD_ENABLE_PS2_2);
375:
376:     /* reset device(s) */
377:     keyboard_write(KB_DATA, KB_CMD_RESET);
378:     if(keyboard_wait_ack()) {
379:         printk("WARNING: %s(): ACK not received on reset command!\n", __
FUNCTION__);
380:     }
381:     if(!keyboard_wait_output()) {
382:         if((errno = inport_b(KB_DATA)) != KB_RESET_OK) {
383:             printk("WARNING: %s(): keyboard returned 0x%x in reset.\n
n", __FUNCTION__, errno);
384:         }
385:     }
386:
387:     return;
388: }
389:
390: static void putc(struct tty *tty, unsigned char ch)
391: {
392:     if(tty_queue_putchar(tty, &tty->read_q, ch) < 0) {

```

## drivers/char/keyboard.c

Page 7/12

```

393:             if(tty->termios.c_iflag & IMAXBEL) {
394:                 vconsole_beep();
395:             }
396:         }
397: }
398:
399: static void puts(struct tty *tty, char *seq)
400: {
401:     char ch;
402:
403:     while((ch = *(seq++)) {
404:         putc(tty, ch);
405:     }
406: }
407:
408: void reboot(void)
409: {
410:     CLI();
411:     keyboard_write(PS2_SYSCtrl_A, 0x01);           /* Fast Hot Reset */
412:     keyboard_write(KBC_COMMAND, KBC_CMD_HOTRESET); /* Hot Reset */
413:     HLT();
414: }
415:
416: void set_leds(unsigned char leds)
417: {
418:     keyboard_write(KB_DATA, KB_SETLED);
419:     keyboard_wait_ack();
420:
421:     keyboard_write(KB_DATA, leds);
422:     keyboard_wait_ack();
423: }
424:
425: void irq_keyboard(int num, struct sigcontext *sc)
426: {
427:     __key_t key, type;
428:     unsigned char scode, mod;
429:     struct tty *tty;
430:     struct vconsole *vc;
431:     unsigned char c;
432:     int n;
433:
434:     tty = get_tty(MKDEV(VCONSOLES_MAJOR, current_cons));
435:     vc = (struct vconsole *)tty->driver_data;
436:
437:     scode = inport_b(KB_DATA);
438:
439:     screen_on();
440:     keyboard_bh.flags |= BH_ACTIVE;
441:
442:     /* keyboard said all is OK, perfect */
443:     if(scode == KB_ACK) {
444:         return;
445:     }
446:
447:     /* if in pure raw mode just queue the scan code and return */
448:     if(tty->kbd.mode == K_RAW) {
449:         putc(tty, scode);
450:         return;
451:     }
452:
453:     if(scode == EXTKEY) {
454:         extkey = 1;
455:         return;
456:     }
457:
458:     if(extkey) {
459:         key = e0_keys[scode & 0x7f];

```

## drivers/char/keyboard.c

Page 8/12

```

460:         } else {
461:             key = scode & 0x7F;
462:         }
463:
464:         if(tty->kbd.mode == K_MEDIUMRAW) {
465:             putc(tty, key | (scode & 0x80));
466:             extkey = 0;
467:             return;
468:         }
469:
470:         key = keymap[NR_MODIFIERS * (scode & 0x7F)];
471:
472:         /* bit 7 enabled means a key has been released */
473:         if(scode & NR_SCODES) {
474:             switch(key) {
475:                 case CTRL:
476:                 case LCTRL:
477:                 case RCTRL:
478:                     ctrl = 0;
479:                     break;
480:                 case ALT:
481:                     if(!extkey) {
482:                         alt = 0;
483:                     } else {
484:                         altgr = 0;
485:                     }
486:                     break;
487:                 case SHIFT:
488:                 case LSHIFT:
489:                 case RSHIFT:
490:                     if(!extkey) {
491:                         shift = 0;
492:                     }
493:                     break;
494:                 case CAPS:
495:                 case NUMS:
496:                 case SCRL:
497:                     leds = 0;
498:                     break;
499:             }
500:             extkey = 0;
501:             return;
502:         }
503:
504:         switch(key) {
505:             case CAPS:
506:                 if(!leds) {
507:                     vc->led_status ^= CAPSBIT;
508:                     vc->capslock = !vc->capslock;
509:                     set_leds(vc->led_status);
510:                 }
511:                 leds = 1;
512:                 return;
513:             case NUMS:
514:                 if(!leds) {
515:                     vc->led_status ^= NUMSBIT;
516:                     vc->numlock = !vc->numlock;
517:                     set_leds(vc->led_status);
518:                 }
519:                 leds = 1;
520:                 return;
521:             case SCRL:
522:                 if(!leds) {
523:                     if(vc->scrlock) {
524:                         tty->start(tty);
525:                     } else {
526:                         tty->stop(tty);

```

## drivers/char/keyboard.c

Page 9/12

```

527:         }
528:     }
529:     leds = 1;
530:     return;
531: case CTRL:
532: case LCTRL:
533: case RCTRL:
534:     ctrl = 1;
535:     return;
536: case ALT:
537:     if(!extkey) {
538:         alt = 1;
539:     } else {
540:         altgr = 1;
541:     }
542:     return;
543: case SHIFT:
544: case LSHIFT:
545: case RSHIFT:
546:     shift = 1;
547:     extkey = 0;
548:     return;
549: }
550:
551: if(ctrl && alt && key == DEL) {
552:     if(ctrl_alt_del) {
553:         reboot();
554:     } else {
555:         send_sig(&proc_table[INIT], SIGINT);
556:     }
557:     return;
558: }
559:
560: keymap_line = &keymap[(scode & 0x7F) * NR_MODIFIERS];
561: mod = 0;
562:
563: if(vc->capslock && (keymap_line[MOD_BASE] & LETTER_KEYS)) {
564:     mod = !vc->capslock ? shift : vc->capslock - shift;
565: } else {
566:     if(shift && !extkey) {
567:         mod = 1;
568:     }
569: }
570: if(altgr) {
571:     mod = 2;
572: }
573: if(ctrl) {
574:     mod = 4;
575: }
576: if(alt) {
577:     mod = 8;
578: }
579:
580: key = keymap_line[mod];
581:
582: if(key >= AF1 && key <= AF12) {
583:     do_switch_console = key - CONS_KEYS;
584:     return;
585: }
586:
587: if(shift && (key == PGUP)) {
588:     do_scr1_buf = VC_BUF_UP;
589:     return;
590: }
591:
592: if(shift && (key == PGDN)) {
593:     do_scr1_buf = VC_BUF_DOWN;

```

## drivers/char/keyboard.c

Page 10/12

```

594:         return;
595:     }
596:
597:     if(extkey && (scode == SLASH_NPAD)) {
598:         key = SLASH;
599:     }
600:
601:     if(any_key_to_reboot) {
602:         reboot();
603:     }
604:
605:     if(tty->count) {
606:         type = key & 0xFF00;
607:         c = key & 0xFF;
608:
609:         switch(type) {
610:             case FN_KEYS:
611:                 puts(tty, fn_seq[c]);
612:                 break;
613:
614:             case SPEC_KEYS:
615:                 switch(key) {
616:                     case CR:
617:                         putc(tty, C('M'));
618:                         break;
619:                 }
620:                 break;
621:
622:             case PAD_KEYS:
623:                 if(!vc->numlock) {
624:                     puts(tty, pad_seq[c]);
625:                 } else {
626:                     putc(tty, pad_chars[c]);
627:                 }
628:                 break;
629:
630:             case DEAD_KEYS:
631:                 if(!deadkey) {
632:                     switch(c) {
633:                         case GRAVE ^ DEAD_KEYS:
634:                             deadkey = 1;
635:                             diacr = grave_table;
636:                             break;
637:                         case ACUTE ^ DEAD_KEYS:
638:                             deadkey = 2;
639:                             diacr = acute_table;
640:                             break;
641:                         case CIRCUM ^ DEAD_KEYS:
642:                             deadkey = 3;
643:                             diacr = circum_table;
644:                             break;
645:                         case DIERE ^ DEAD_KEYS:
646:                             deadkey = 5;
647:                             diacr = diere_table;
648:                             break;
649:                     }
650:                     return;
651:                 }
652:                 c = diacr_chars[c];
653:                 deadkey = 0;
654:                 putc(tty, c);
655:
656:                 break;
657:
658:             case META_KEYS:
659:                 putc(tty, '\033');
660:                 putc(tty, c);

```

## drivers/char/keyboard.c

Page 11/12

```

661:             break;
662:
663:         case LETTER_KEYS:
664:             if (deadkey) {
665:                 for (n = 0; n < NR_DIACR; n++) {
666:                     if (diacr[n].letter == c) {
667:                         c = diacr[n].code;
668:                     }
669:                 }
670:             }
671:             putc(tty, c);
672:             break;
673:
674:         default:
675:             if (deadkey && c == ' ') {
676:                 c = diacr_chars[deadkey - 1];
677:             }
678:             putc(tty, c);
679:             break;
680:     }
681: }
682: deadkey = 0;
683: return;
684: }
685:
686: void irq_keyboard_bh(void)
687: {
688:     int n;
689:     struct tty *tty;
690:
691:     if (do_switch_console >= 0) {
692:         vconsole_select(do_switch_console);
693:         do_switch_console = -1;
694:     }
695:
696:     if (do_scr_l_buf) {
697:         vconsole_buffer_scr_l(do_scr_l_buf);
698:         do_scr_l_buf = 0;
699:     }
700:
701:     tty = &tty_table[0];
702:     for (n = 0; n < NR_VCONSOLES; n++, tty++) {
703:         if (!tty->read_q.count) {
704:             continue;
705:         }
706:         if (tty->kbd.mode == K_RAW || tty->kbd.mode == K_MEDIUMRAW) {
707:             wakeup(&tty_read);
708:             continue;
709:         }
710:         if (lock_area(AREA_TTY_READ)) {
711:             keyboard_bh.flags |= BH_ACTIVE;
712:             continue;
713:         }
714:         tty->input(tty);
715:         unlock_area(AREA_TTY_READ);
716:     }
717: }
718:
719: void keyboard_init(void)
720: {
721:     add_bh(&keyboard_bh);
722:
723:     keyboard_reset();
724:
725:     /* flush buffers */
726:     while (!keyboard_wait_output()) {
727:         inport_b(KB_DATA);

```

**drivers/char/keyboard.c**

Page 12/12

```
728:         }
729:
730:         keyboard_identify();
731:
732:         keyboard_write(KB_DATA, KB_RATE);
733:         keyboard_wait_ack();
734:         keyboard_write(KB_DATA, DELAY_250 | RATE_30);
735:         keyboard_wait_ack();
736:
737:         printk("keyboard 0x%04X-0x%04X   %d   type=%s %s PS/2 devices=%d/%d\n",
", 0x60, 0x64, KEYBOARD_IRQ, kb_identify[0] == 0xAB ? "MF2" : "unknown", ps2_iface & 0x
1 ? "MCA" : "AT", ps2_active_ports, ps2_supp_ports);
738:
739:         if(!register_irq(KEYBOARD_IRQ, &irq_config_keyboard)) {
740:             enable_irq(KEYBOARD_IRQ);
741:         }
742: }
```



## drivers/char/lp.c

Page 1/4

```

1:  /*
2:  *  fiwix/drivers/char/lp.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/asm.h>
9:  #include <fiwix/devices.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/lp.h>
13: #include <fiwix/stdio.h>
14: #include <fiwix/string.h>
15:
16: struct lp lp_table[LP_MINORS];
17:
18: static struct fs_operations lp_driver_fsop = {
19:     0,
20:     0,
21:
22:     lp_open,
23:     lp_close,
24:     NULL, /* read */
25:     lp_write,
26:     NULL, /* ioctl */
27:     NULL, /* lseek */
28:     NULL, /* readdir */
29:     NULL, /* mmap */
30:     NULL, /* select */
31:
32:     NULL, /* readlink */
33:     NULL, /* followlink */
34:     NULL, /* bmap */
35:     NULL, /* lockup */
36:     NULL, /* rmdir */
37:     NULL, /* link */
38:     NULL, /* unlink */
39:     NULL, /* symlink */
40:     NULL, /* mkdir */
41:     NULL, /* mknod */
42:     NULL, /* truncate */
43:     NULL, /* create */
44:     NULL, /* rename */
45:
46:     NULL, /* read_block */
47:     NULL, /* write_block */
48:
49:     NULL, /* read_inode */
50:     NULL, /* write_inode */
51:     NULL, /* ialloc */
52:     NULL, /* ifree */
53:     NULL, /* statfs */
54:     NULL, /* read_superblock */
55:     NULL, /* remount_fs */
56:     NULL, /* write_superblock */
57:     NULL, /* release_superblock */
58: };
59:
60: static struct device lp_device = {
61:     "lp",
62:     LP_MAJOR,
63:     { 0, 0, 0, 0, 0, 0, 0, 0 },
64:     0,
65:     NULL,
66:     &lp_driver_fsop,
67:     NULL

```

## drivers/char/lp.c

Page 2/4

```

68: };
69:
70: struct lp lp_table[LP_MINORS] = {
71:     { LP0_ADDR, LP0_ADDR + 1, LP0_ADDR + 2, 0 }
72: };
73:
74: static void lp_delay(void)
75: {
76:     int n;
77:
78:     for(n = 0; n < 10000; n++) {
79:         NOP();
80:     }
81: }
82:
83: static int lp_ready(int minor)
84: {
85:     int n;
86:
87:     for(n = 0; n < LP_RDY_RETR; n++) {
88:         if(inport_b(lp_table[minor].stat) & LP_STAT_BUS) {
89:             break;
90:         }
91:         lp_delay();
92:     }
93:     if(n == LP_RDY_RETR) {
94:         return 0;
95:     }
96:     return 1;
97: }
98:
99: static int lp_probe(int minor)
100: {
101:     /* first check */
102:     outport_b(lp_table[minor].data, 0x55);
103:     lp_delay();
104:     if(inport_b(lp_table[minor].data) != 0x55) {
105:         return 1;          /* did not retain data */
106:     }
107:
108:     /* second check */
109:     outport_b(lp_table[minor].data, 0xAA);
110:     lp_delay();
111:     if(inport_b(lp_table[minor].data) != 0xAA) {
112:         return 1;          /* did not retain data */
113:     }
114:     return 0;
115: }
116:
117: static int lp_write_data(int minor, unsigned char c)
118: {
119:     unsigned char ctrl;
120:
121:     if(!lp_ready(minor)) {
122:         return -EBUSY;
123:     }
124:     outport_b(lp_table[minor].data, c);
125:     ctrl = inport_b(lp_table[minor].ctrl);
126:     outport_b(lp_table[minor].ctrl, ctrl | LP_CTRL_STR);
127:     lp_delay();
128:     outport_b(lp_table[minor].ctrl, ctrl);
129:     if(!lp_ready(minor)) {
130:         return -EBUSY;
131:     }
132:     return 1;
133: }
134:

```

## drivers/char/lp.c

Page 3/4

```

135: int lp_open(struct inode *i, struct fd *fd_table)
136: {
137:     int minor;
138:
139:     minor = MINOR(i->rdev);
140:     if(!TEST_MINOR(lp_device.minors, minor)) {
141:         return -ENXIO;
142:     }
143:     if(!(lp_table[minor].flags & LP_CTRL_SEL)) {
144:         return -ENXIO;
145:     }
146:     if(lp_table[minor].flags & LP_STAT_BUS) {
147:         return -EBUSY;
148:     }
149:     lp_table[minor].flags |= LP_STAT_BUS;
150:     return 0;
151: }
152:
153: int lp_close(struct inode *i, struct fd *fd_table)
154: {
155:     int minor;
156:
157:     minor = MINOR(i->rdev);
158:     if(!TEST_MINOR(lp_device.minors, minor)) {
159:         return -ENXIO;
160:     }
161:     lp_table[minor].flags &= ~LP_STAT_BUS;
162:     return 0;
163: }
164:
165: int lp_write(struct inode *i, struct fd *fd_table, const char *buffer, __size_t
count)
166: {
167:     unsigned int n;
168:     int bytes_written, total_written;
169:     int minor;
170:
171:     minor = MINOR(i->rdev);
172:     if(!TEST_MINOR(lp_device.minors, minor)) {
173:         return -ENXIO;
174:     }
175:
176:     total_written = 0;
177:     for(n = 0; n < count; n++) {
178:         bytes_written = lp_write_data(minor, buffer[n]);
179:         if(bytes_written != 1) {
180:             break;
181:         }
182:         total_written += bytes_written;
183:     }
184:
185:     return total_written;
186: }
187:
188: void lp_init(void)
189: {
190:     int n;
191:     unsigned char ctrl;
192:
193:     for(n = 0; n < LP_MINORS; n++) {
194:         if(!lp_probe(n)) {
195:             ctrl = inport_b(lp_table[n].ctrl);
196:             ctrl &= ~LP_CTRL_AUT; /* disable auto LF */
197:             ctrl |= LP_CTRL_INI; /* initialize */
198:             ctrl |= LP_CTRL_SEL; /* select in */
199:             ctrl &= ~LP_CTRL_IRQ; /* disable IRQ */
200:             ctrl &= ~LP_CTRL_BID; /* disable bidirectional mode */

```

## drivers/char/lp.c

Page 4/4

```
201:         outport_b(lp_table[n].ctrl, ctrl);
202:         lp_table[n].flags |= LP_CTRL_SEL;
203:         printk("lp%d      0x%04X-0x%04X      -      \n", n,
lp_table[n].data, lp_table[n].data + 2);
204:         SET_MINOR(lp_device.minors, n);
205:     }
206: }
207:
208:     for(n = 0; n < LP_MINORS; n++) {
209:         if(lp_table[n].flags & LP_CTRL_SEL) {
210:             if(register_device(CHR_DEV, &lp_device)) {
211:                 printk("WARNING: %s(): unable to register lp dev
ice.\n", __FUNCTION__);
212:             }
213:             break;
214:         }
215:     }
216: }
```

**drivers/char/Makefile**

Page 1/1

```
1: # fiwix/drivers/char/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = console.o tty.o tty_queue.o vt.o defkeymap.o keyboard.o memdev.o serial.o
lp.o
13:
14: char:     $(OBJS)
15:     $(LD) $(LDFLAGS) -r $(OBJS) -o char.o
16:
17: clean:
18:     rm -f *.o
19:
```

## drivers/char/memdev.c

Page 1/8

```

1: /*
2:  * fiwix/drivers/char/memdev.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/memdev.h>
11: #include <fiwix/devices.h>
12: #include <fiwix/fs.h>
13: #include <fiwix/errno.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/mman.h>
16: #include <fiwix/bios.h>
17: #include <fiwix/stdio.h>
18: #include <fiwix/string.h>
19:
20: static struct fs_operations mem_driver_fsop = {
21:     0,
22:     0,
23:
24:     mem_open,
25:     mem_close,
26:     mem_read,
27:     mem_write,
28:     NULL,                /* ioctl */
29:     mem_lseek,
30:     NULL,                /* readdir */
31:     mem_mmap,
32:     NULL,                /* select */
33:
34:     NULL,                /* readlink */
35:     NULL,                /* followlink */
36:     NULL,                /* bmap */
37:     NULL,                /* lockup */
38:     NULL,                /* rmdir */
39:     NULL,                /* link */
40:     NULL,                /* unlink */
41:     NULL,                /* symlink */
42:     NULL,                /* mkdir */
43:     NULL,                /* mknod */
44:     NULL,                /* truncate */
45:     NULL,                /* create */
46:     NULL,                /* rename */
47:
48:     NULL,                /* read_block */
49:     NULL,                /* write_block */
50:
51:     NULL,                /* read_inode */
52:     NULL,                /* write_inode */
53:     NULL,                /* ialloc */
54:     NULL,                /* ifree */
55:     NULL,                /* statfs */
56:     NULL,                /* read_superblock */
57:     NULL,                /* remount_fs */
58:     NULL,                /* write_superblock */
59:     NULL,                /* release_superblock */
60: };
61:
62: static struct fs_operations kmem_driver_fsop = {
63:     0,
64:     0,
65:
66:     kmem_open,
67:     kmem_close,

```

## drivers/char/memdev.c

Page 2/8

```

68:         kmem_read,
69:         kmem_write,
70:         NULL,                /* ioctl */
71:         kmem_lseek,
72:         NULL,                /* readdir */
73:         mem_mmap,
74:         NULL,                /* select */
75:
76:         NULL,                /* readlink */
77:         NULL,                /* followlink */
78:         NULL,                /* bmap */
79:         NULL,                /* lockup */
80:         NULL,                /* rmdir */
81:         NULL,                /* link */
82:         NULL,                /* unlink */
83:         NULL,                /* symlink */
84:         NULL,                /* mkdir */
85:         NULL,                /* mknod */
86:         NULL,                /* truncate */
87:         NULL,                /* create */
88:         NULL,                /* rename */
89:
90:         NULL,                /* read_block */
91:         NULL,                /* write_block */
92:
93:         NULL,                /* read_inode */
94:         NULL,                /* write_inode */
95:         NULL,                /* ialloc */
96:         NULL,                /* ifree */
97:         NULL,                /* statfs */
98:         NULL,                /* read_superblock */
99:         NULL,                /* remount_fs */
100:        NULL,                /* write_superblock */
101:        NULL,                /* release_superblock */
102:    };
103:
104:    static struct fs_operations null_driver_fsop = {
105:        0,
106:        0,
107:
108:        null_open,
109:        null_close,
110:        null_read,
111:        null_write,
112:        NULL,                /* ioctl */
113:        null_lseek,
114:        NULL,                /* readdir */
115:        NULL,                /* mmap */
116:        NULL,                /* select */
117:
118:        NULL,                /* readlink */
119:        NULL,                /* followlink */
120:        NULL,                /* bmap */
121:        NULL,                /* lockup */
122:        NULL,                /* rmdir */
123:        NULL,                /* link */
124:        NULL,                /* unlink */
125:        NULL,                /* symlink */
126:        NULL,                /* mkdir */
127:        NULL,                /* mknod */
128:        NULL,                /* truncate */
129:        NULL,                /* create */
130:        NULL,                /* rename */
131:
132:        NULL,                /* read_block */
133:        NULL,                /* write_block */
134:

```

## drivers/char/memdev.c

Page 3/8

```

135:     NULL,                /* read_inode */
136:     NULL,                /* write_inode */
137:     NULL,                /* ialloc */
138:     NULL,                /* ifree */
139:     NULL,                /* statfs */
140:     NULL,                /* read_superblock */
141:     NULL,                /* remount_fs */
142:     NULL,                /* write_superblock */
143:     NULL,                /* release_superblock */
144: };
145:
146: static struct fs_operations zero_driver_fsop = {
147:     0,
148:     0,
149:
150:     zero_open,
151:     zero_close,
152:     zero_read,
153:     zero_write,
154:     NULL,                /* ioctl */
155:     zero_lseek,
156:     NULL,                /* readdir */
157:     NULL,                /* mmap */
158:     NULL,                /* select */
159:
160:     NULL,                /* readlink */
161:     NULL,                /* followlink */
162:     NULL,                /* bmap */
163:     NULL,                /* lockup */
164:     NULL,                /* rmdir */
165:     NULL,                /* link */
166:     NULL,                /* unlink */
167:     NULL,                /* symlink */
168:     NULL,                /* mkdir */
169:     NULL,                /* mknod */
170:     NULL,                /* truncate */
171:     NULL,                /* create */
172:     NULL,                /* rename */
173:
174:     NULL,                /* read_block */
175:     NULL,                /* write_block */
176:
177:     NULL,                /* read_inode */
178:     NULL,                /* write_inode */
179:     NULL,                /* ialloc */
180:     NULL,                /* ifree */
181:     NULL,                /* statfs */
182:     NULL,                /* read_superblock */
183:     NULL,                /* remount_fs */
184:     NULL,                /* write_superblock */
185:     NULL,                /* release_superblock */
186: };
187:
188: static struct fs_operations memdev_driver_fsop = {
189:     0,
190:     0,
191:
192:     memdev_open,
193:     NULL,                /* close */
194:     NULL,                /* read */
195:     NULL,                /* write */
196:     NULL,                /* ioctl */
197:     NULL,                /* lseek */
198:     NULL,                /* readdir */
199:     NULL,                /* mmap */
200:     NULL,                /* select */
201:

```



## drivers/char/memdev.c

Page 4/8

```

202:         NULL,                /* readlink */
203:         NULL,                /* followlink */
204:         NULL,                /* bmap */
205:         NULL,                /* lockup */
206:         NULL,                /* rmdir */
207:         NULL,                /* link */
208:         NULL,                /* unlink */
209:         NULL,                /* symlink */
210:         NULL,                /* mkdir */
211:         NULL,                /* mknod */
212:         NULL,                /* truncate */
213:         NULL,                /* create */
214:         NULL,                /* rename */
215:
216:         NULL,                /* read_block */
217:         NULL,                /* write_block */
218:
219:         NULL,                /* read_inode */
220:         NULL,                /* write_inode */
221:         NULL,                /* ialloc */
222:         NULL,                /* ifree */
223:         NULL,                /* statfs */
224:         NULL,                /* read_superblock */
225:         NULL,                /* remount_fs */
226:         NULL,                /* write_superblock */
227:         NULL,                /* release_superblock */
228: };
229:
230: static struct fs_operations urandom_driver_fsop = {
231:     0,
232:     0,
233:
234:     urandom_open,
235:     urandom_close,
236:     urandom_read,
237:     urandom_write,
238:     NULL,                /* ioctl */
239:     urandom_lseek,
240:     NULL,                /* readdir */
241:     NULL,                /* mmap */
242:     NULL,                /* select */
243:
244:     NULL,                /* readlink */
245:     NULL,                /* followlink */
246:     NULL,                /* bmap */
247:     NULL,                /* lockup */
248:     NULL,                /* rmdir */
249:     NULL,                /* link */
250:     NULL,                /* unlink */
251:     NULL,                /* symlink */
252:     NULL,                /* mkdir */
253:     NULL,                /* mknod */
254:     NULL,                /* truncate */
255:     NULL,                /* create */
256:     NULL,                /* rename */
257:
258:     NULL,                /* read_block */
259:     NULL,                /* write_block */
260:
261:     NULL,                /* read_inode */
262:     NULL,                /* write_inode */
263:     NULL,                /* ialloc */
264:     NULL,                /* ifree */
265:     NULL,                /* statfs */
266:     NULL,                /* read_superblock */
267:     NULL,                /* remount_fs */
268:     NULL,                /* write_superblock */

```

## drivers/char/memdev.c

Page 5/8

```

269:         NULL                                     /* release_superblock */
270: };
271:
272: static struct device memdev_device = {
273:     "mem",
274:     MEMDEV_MAJOR,
275:     { 0, 0, 0, 0, 0, 0, 0, 0, 0 },
276:     0,
277:     NULL,
278:     &memdev_driver_fsop,
279:     NULL
280: };
281:
282: int mem_open(struct inode *i, struct fd *fd_table)
283: {
284:     return 0;
285: }
286:
287: int mem_close(struct inode *i, struct fd *fd_table)
288: {
289:     return 0;
290: }
291:
292: int mem_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t count)
293: {
294:     unsigned int physical_memory;
295:
296:     physical_memory = (kstat.physical_pages << PAGE_SHIFT);
297:     if(fd_table->offset >= physical_memory) {
298:         return 0;
299:     }
300:     count = MIN(count, physical_memory - fd_table->offset);
301:     memcpy_b(buffer, (void *)P2V(fd_table->offset), count);
302:     fd_table->offset += count;
303:     return count;
304: }
305:
306: int mem_write(struct inode *i, struct fd *fd_table, const char *buffer, __size_t
count)
307: {
308:     unsigned int physical_memory;
309:
310:     physical_memory = (kstat.physical_pages << PAGE_SHIFT);
311:     if(fd_table->offset >= physical_memory) {
312:         return 0;
313:     }
314:     count = MIN(count, physical_memory - fd_table->offset);
315:     memcpy_b((void *)P2V(fd_table->offset), buffer, count);
316:     fd_table->offset += count;
317:     return count;
318: }
319:
320: int mem_lseek(struct inode *i, __off_t offset)
321: {
322:     return offset;
323: }
324:
325: int kmem_open(struct inode *i, struct fd *fd_table)
326: {
327:     return 0;
328: }
329:
330: int kmem_close(struct inode *i, struct fd *fd_table)
331: {
332:     return 0;
333: }
334:

```

## drivers/char/memdev.c

Page 6/8

```
335: int kmem_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t count
)
336: {
337:     memcpy_b(buffer, (void *)P2V(fd_table->offset), count);
338:     fd_table->offset += count;
339:     return count;
340: }
341:
342: int kmem_write(struct inode *i, struct fd *fd_table, const char *buffer, __size_
t count)
343: {
344:     memcpy_b((void *)P2V(fd_table->offset), buffer, count);
345:     fd_table->offset += count;
346:     return count;
347: }
348:
349: int kmem_lseek(struct inode *i, __off_t offset)
350: {
351:     return offset;
352: }
353:
354: int null_open(struct inode *i, struct fd *fd_table)
355: {
356:     return 0;
357: }
358:
359: int null_close(struct inode *i, struct fd *fd_table)
360: {
361:     return 0;
362: }
363:
364: int null_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t count
)
365: {
366:     return 0;
367: }
368:
369: int null_write(struct inode *i, struct fd *fd_table, const char *buffer, __size_
t count)
370: {
371:     return count;
372: }
373:
374: int null_lseek(struct inode *i, __off_t offset)
375: {
376:     return offset;
377: }
378:
379: int zero_open(struct inode *i, struct fd *fd_table)
380: {
381:     return 0;
382: }
383:
384: int zero_close(struct inode *i, struct fd *fd_table)
385: {
386:     return 0;
387: }
388:
389: int zero_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t count
)
390: {
391:     memset_b(buffer, NULL, count);
392:     return count;
393: }
394:
395: int zero_write(struct inode *i, struct fd *fd_table, const char *buffer, __size_
t count)
```

## drivers/char/memdev.c

Page 7/8

```

396: {
397:     return count;
398: }
399:
400: int zero_lseek(struct inode *i, __off_t offset)
401: {
402:     return offset;
403: }
404:
405: int urandom_open(struct inode *i, struct fd *fd_table)
406: {
407:     return 0;
408: }
409:
410: int urandom_close(struct inode *i, struct fd *fd_table)
411: {
412:     return 0;
413: }
414:
415: int urandom_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t count)
416: {
417:     int n;
418:
419:     for(n = 0; n < count; n++) {
420:         kstat.random_seed = kstat.random_seed * 1103515245 + 12345;
421:         *buffer = (char)(unsigned int)(kstat.random_seed / 65536) % 256;
422:         buffer++;
423:     }
424:     return count;
425: }
426:
427: int urandom_write(struct inode *i, struct fd *fd_table, const char *buffer, __size_t count)
428: {
429:     return count;
430: }
431:
432: int urandom_lseek(struct inode *i, __off_t offset)
433: {
434:     return offset;
435: }
436:
437: int memdev_open(struct inode *i, struct fd *fd_table)
438: {
439:     unsigned char minor;
440:
441:     minor = MINOR(i->rdev);
442:     switch(minor) {
443:         case MEMDEV_MEM:
444:             i->fsop = &mem_driver_fsop;
445:             break;
446:         case MEMDEV_KMEM:
447:             i->fsop = &kmem_driver_fsop;
448:             break;
449:         case MEMDEV_NULL:
450:             i->fsop = &null_driver_fsop;
451:             break;
452:         case MEMDEV_ZERO:
453:             i->fsop = &zero_driver_fsop;
454:             break;
455:         case MEMDEV_RANDOM:
456:             i->fsop = &urandom_driver_fsop;
457:             break;
458:         case MEMDEV_URANDOM:
459:             i->fsop = &urandom_driver_fsop;
460:             break;

```

## drivers/char/memdev.c

Page 8/8

```

461:         default:
462:             return -ENXIO;
463:     }
464:     if(i->fsop->open) {
465:         return i->fsop->open(i, fd_table);
466:     }
467:     return 0;
468: }
469:
470: /*
471:  * This function maps a range of physical addresses marked as not available for
472:  * use in the BIOS memory map, like the video RAM.
473:  */
474: int mem_mmap(struct inode *i, struct vma *vma)
475: {
476:     unsigned int addr, length;
477:
478:     length = (vma->end - vma->start) & PAGE_MASK;
479:
480:     /* this breaks down the range in 4KB chunks */
481:     for(addr = 0; addr < length; addr += PAGE_SIZE) {
482:         /* map the page only if is NOT available in the BIOS map */
483:         if(!addr_in_bios_map(vma->offset + addr)) {
484:             if(!map_page(current, (vma->start + addr) & PAGE_MASK, (
vma->offset + addr) & PAGE_MASK, PROT_READ | PROT_WRITE)) {
485:                 return -ENOMEM;
486:             }
487:         } else {
488:             printk("ERROR: %s(): mapping AVAILABLE pages in BIOS mem
ory map isn't supported.\n", __FUNCTION__);
489:             printk("\tinvalid mapping: 0x%08x -> 0x%08x\n", (vma->st
art + addr) & PAGE_MASK, (vma->offset + addr) & PAGE_MASK);
490:             return -EAGAIN;
491:         }
492:     }
493:     invalidate_tlb();
494:     return 0;
495: }
496:
497: void memdev_init(void)
498: {
499:     SET_MINOR(memdev_device.minors, MEMDEV_MEM);
500:     SET_MINOR(memdev_device.minors, MEMDEV_KMEM);
501:     SET_MINOR(memdev_device.minors, MEMDEV_NULL);
502:     SET_MINOR(memdev_device.minors, MEMDEV_ZERO);
503:     SET_MINOR(memdev_device.minors, MEMDEV_RANDOM);
504:     SET_MINOR(memdev_device.minors, MEMDEV_URANDOM);
505:
506:     if(register_device(CHR_DEV, &memdev_device)) {
507:         printk("ERROR: %s(): unable to register memory devices.\n", __FU
NCTION__);
508:         return;
509:     }
510: }

```

## drivers/char/serial.c

Page 1/9

```

1: /*
2:  * fiwix/drivers/char/serial.c
3:  *
4:  * Copyright 2020, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/devices.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/errno.h>
13: #include <fiwix/pic.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/serial.h>
16: #include <fiwix/tty.h>
17: #include <fiwix/ctype.h>
18: #include <fiwix/stdio.h>
19: #include <fiwix/string.h>
20:
21: static struct fs_operations serial_driver_fsop = {
22:     0,
23:     0,
24:
25:     tty_open,
26:     tty_close,
27:     tty_read,
28:     tty_write,
29:     tty_ioctl,
30:     tty_lseek,
31:     NULL,                /* readdir */
32:     NULL,                /* mmap */
33:     tty_select,
34:
35:     NULL,                /* readlink */
36:     NULL,                /* followlink */
37:     NULL,                /* bmap */
38:     NULL,                /* lockup */
39:     NULL,                /* rmdir */
40:     NULL,                /* link */
41:     NULL,                /* unlink */
42:     NULL,                /* symlink */
43:     NULL,                /* mkdir */
44:     NULL,                /* mknod */
45:     NULL,                /* truncate */
46:     NULL,                /* create */
47:     NULL,                /* rename */
48:
49:     NULL,                /* read_block */
50:     NULL,                /* write_block */
51:
52:     NULL,                /* read_inode */
53:     NULL,                /* write_inode */
54:     NULL,                /* ialloc */
55:     NULL,                /* ifree */
56:     NULL,                /* statfs */
57:     NULL,                /* read_superblock */
58:     NULL,                /* remount_fs */
59:     NULL,                /* write_superblock */
60:     NULL,                /* release_superblock */
61: };
62:
63: static struct device serial_device = {
64:     "ttys",
65:     SERIAL_MAJOR,
66:     { 0, 0, 0, 0, 0, 0, 0, 0, 0 },
67:     0,

```

## drivers/char/serial.c

Page 2/9

```

68:         NULL,
69:         &serial_driver_fsop,
70:         NULL
71: };
72:
73: /* 9600,N,8,1 by default */
74: struct serial serial_table[NR_SERIAL] = {
75:     { 0x3F8, SERIAL4_IRQ, 9600, "ttyS0", UART_LCR_NP | UART_LCR_WL8 | UART_L
CR_1STB, 0, NULL, NULL },
76:     { 0x2F8, SERIAL3_IRQ, 9600, "ttyS1", UART_LCR_NP | UART_LCR_WL8 | UART_L
CR_1STB, 0, NULL, NULL },
77:     { 0x3E8, SERIAL4_IRQ, 9600, "ttyS2", UART_LCR_NP | UART_LCR_WL8 | UART_L
CR_1STB, 0, NULL, NULL },
78:     { 0x2E8, SERIAL3_IRQ, 9600, "ttyS3", UART_LCR_NP | UART_LCR_WL8 | UART_L
CR_1STB, 0, NULL, NULL }
79: };
80:
81: char *serial_chip[] = {
82:     NULL,
83:     "8250",
84:     "16450",
85:     "16550",
86:     "16550A",
87: };
88:
89: static int baud_table[] = {
90:     0,
91:     50,
92:     75,
93:     110,
94:     134,
95:     150,
96:     200,
97:     300,
98:     600,
99:     1200,
100:    1800,
101:    2400,
102:    4800,
103:    9600,
104:   19200,
105:   38400,
106:   57200,
107:  115200,
108:    0
109: };
110:
111: static struct serial *serial_ports = NULL;
112: static struct bh serial_bh = { 0, &irq_serial_bh, NULL };
113: static struct interrupt irq_config_serial0 = { 0, "serial", &irq_serial, NULL };
114: static struct interrupt irq_config_serial1 = { 0, "serial", &irq_serial, NULL };
115:
116: static int is_serial(__dev_t dev)
117: {
118:     if (MAJOR(dev) == SERIAL_MAJOR && ((MINOR(dev) >= (1 << SERIAL_MSF) && MI
NOR(dev) < (1 << SERIAL_MSF) + SERIAL_MINORS))) {
119:         return 1;
120:     }
121:
122:     return 0;
123: }
124:
125: static int serial_identify(struct serial *s)
126: {
127:     int value;
128:
129:     /* set all features in FCR register to test the status of FIFO */

```

## drivers/char/serial.c

Page 3/9

```

130:         outport_b(s->addr + UART_FCR, (UART_FCR_FIFO |
131:                                         UART_FCR_DMA |
132:                                         UART_FCR_FIFO64 |
133:                                         UART_FCR_FIFO14));
134:
135:         value = inport_b(s->addr + UART_IIR);
136:         if(value & UART_IIR_FIFOKO) {
137:             if(value & UART_IIR_FIFO) {
138:                 if(value & UART_IIR_FIFO64) {
139:                     /* 16750 chip is not supported */
140:                 } else {
141:                     s->flags |= UART_IS_16550A | UART_HAS_FIFO;
142:                     return 4;
143:                 }
144:             } else {
145:                 s->flags |= UART_IS_16550;
146:                 return 3;
147:             }
148:         } else {
149:             /*
150:              * At this point we know this device don't has FIFO,
151:              * the Scratch Register will help us to know the final chip.
152:              */
153:             value = inport_b(s->addr + UART_SR); /* save its value */
154:             outport_b(s->addr + UART_SR, 0xAA); /* put a random value */
155:             if(inport_b(s->addr + UART_SR) != 0xAA) {
156:                 s->flags |= UART_IS_8250;
157:                 return 1;
158:             } else {
159:                 outport_b(s->addr + UART_SR, value); /* restore it */
160:                 s->flags |= UART_IS_16450;
161:                 return 2;
162:             }
163:         }
164:         return 0;
165:     }
166:
167:     static void serial_default(struct serial *s)
168:     {
169:         int divisor;
170:
171:         outport_b(s->addr + UART_IER, 0); /* disable all interrupts */
172:
173:         divisor = 115200 / s->baud;
174:         outport_b(s->addr + UART_LCR, UART_LCR_DLAB); /* enable DLAB */
175:         outport_b(s->addr + UART_DLL, divisor & 0xFF); /* LSB of divisor */
176:         outport_b(s->addr + UART_DLH, divisor >> 8); /* MSB of divisor */
177:         outport_b(s->addr + UART_LCR, s->lctrl); /* line control */
178:     }
179:
180:     /* disable transmitter interrupts */
181:     static void serial_stop(struct tty *tty)
182:     {
183:         struct serial *s;
184:         unsigned long int flags;
185:
186:         SAVE_FLAGS(flags); CLI();
187:         s = (struct serial *)tty->driver_data;
188:         outport_b(s->addr + UART_IER, UART_IER_RDAI);
189:         RESTORE_FLAGS(flags);
190:     }
191:
192:     /* enable transmitter interrupts */
193:     static void serial_start(struct tty *tty)
194:     {
195:         struct serial *s;
196:         unsigned long int flags;

```



## drivers/char/serial.c

Page 4/9

```

197:
198:     SAVE_FLAGS(flags); CLI();
199:     s = (struct serial *)tty->driver_data;
200:     outport_b(s->addr + UART_IER, UART_IER_RDAI | UART_IER_THREI);
201:     RESTORE_FLAGS(flags);
202: }
203:
204: static void serial_deltab(struct tty *tty)
205: {
206:     unsigned short int col, n;
207:     unsigned char count;
208:     struct cblock *cb;
209:     unsigned char ch;
210:
211:     cb = tty->cooked_q.head;
212:     col = count = 0;
213:
214:     while(cb) {
215:         for(n = 0; n < cb->end_off; n++) {
216:             if(n >= cb->start_off) {
217:                 ch = cb->data[n];
218:                 if(ch == '\t') {
219:                     while(!tty->tab_stop[++col]);
220:                 } else {
221:                     col++;
222:                     if(ISCNTRL(ch) && !ISSPACE(ch) && tty->t
ermios.c_lflag & ECHOCTL) {
223:                         col++;
224:                     }
225:                 }
226:                 col %= 80;
227:             }
228:         }
229:         cb = cb->next;
230:     }
231:     count = tty->column - col;
232:
233:     while(count--) {
234:         tty_queue_putchar(tty, &tty->write_q, '\b');
235:         tty->column--;
236:     }
237: }
238:
239: static void serial_reset(struct tty *tty)
240: {
241:     termios_reset(tty);
242:     tty->winsize.ws_row = 25;
243:     tty->winsize.ws_col = 80;
244:     tty->winsize.ws_xpixel = 0;
245:     tty->winsize.ws_ypixel = 0;
246:     tty->lnext = 0;
247: }
248:
249: static void serial_errors(struct serial *s, int status)
250: {
251:     struct tty *tty;
252:
253:     tty = s->tty;
254:
255:     if(!(tty->termios.c_iflag & IGNBRK) && tty->termios.c_iflag & BRKINT) {
256:         if(status & UART_LSR_BI) {
257:             printk("WARNING: break interrupt in %s.\n", s->name);
258:         }
259:     }
260:
261:     /* this includes also overrun errors */
262:     if(!(tty->termios.c_iflag & IGNPAR) && tty->termios.c_iflag & PARMRK) {

```

## drivers/char/serial.c

Page 5/9

```
263:         if(status & UART_LSR_OE) {
264:             printk("WARNING: overrun error in %s.\n", s->name);
265:
266:         } else if(status & UART_LSR_PE) {
267:             printk("WARNING: parity error in %s.\n", s->name);
268:
269:         } else if(status & UART_LSR_FE) {
270:             printk("WARNING: framing error in %s.\n", s->name);
271:
272:         } else if(status & UART_LSR_EFIFO) {
273:             printk("WARNING: FIFO error in %s.\n", s->name);
274:         }
275:     }
276: }
277:
278: static void serial_send(struct tty *tty)
279: {
280:     unsigned char ch;
281:     struct serial *s;
282:     int count;
283:
284:     s = (struct serial *)tty->driver_data;
285:
286:     if(!tty->write_q.count) {
287:         outport_b(s->addr + UART_IER, UART_IER_RDAI);
288:         return;
289:     }
290:
291:     count = 0;
292:     while(tty->write_q.count > 0 && count < UART_FIFO_SIZE) {
293:         ch = tty_queue_getchar(&tty->write_q);
294:         outport_b(s->addr + UART_TD, ch);
295:         count++;
296:     }
297:
298:     if(!tty->write_q.count) {
299:         outport_b(s->addr + UART_IER, UART_IER_RDAI);
300:     }
301:     wakeup(&tty_write);
302: }
303:
304: static int serial_receive(struct serial *s)
305: {
306:     int status, count;
307:     unsigned char ch;
308:     struct tty *tty;
309:
310:     count = 0;
311:     tty = s->tty;
312:
313:     do {
314:         ch = inport_b(s->addr + UART_RD);
315:         tty_queue_putchar(tty, &tty->read_q, ch);
316:         count++;
317:         // FIXME: do_sched()?
318:         status = inport_b(s->addr + UART_LSR);
319:     } while(status & UART_LSR_RDA);
320:
321:     return count;
322: }
323:
324: void irq_serial(int num, struct sigcontext *sc)
325: {
326:     struct serial *s;
327:     int status, count;
328:
329:     s = serial_ports;
```

## drivers/char/serial.c

Page 6/9

```

330:         count = 0;
331:
332:         if(s) {
333:             do {
334:                 if(s->irq == num) {
335:                     while(!(inport_b(s->addr + UART_IIR) & UART_IIR_
NOINT)) {
336:                         status = inport_b(s->addr + UART_LSR);
337:                         if(status & UART_LSR_RDA) {
338:                             if(serial_receive(s)) {
339:                                 count++;
340:                             }
341:                         }
342:                         if(status & UART_LSR_THRE) {
343:                             serial_send(s->tty);
344:                         }
345:                         serial_errors(s, status);
346:                     }
347:                 }
348:                 s = s->next;
349:             } while(s);
350:         }
351:         if(count) {
352:             serial_bh.flags |= BH_ACTIVE;
353:         }
354:     }
355:
356: int serial_open(struct tty *tty)
357: {
358:     struct serial *s;
359:     int minor;
360:
361:     minor = MINOR(tty->dev);
362:     if(!TEST_MINOR(serial_device.minors, minor)) {
363:         return -ENXIO;
364:     }
365:
366:     s = (struct serial *)tty->driver_data;
367:
368:     /* enable FIFO */
369:     if(s->flags & UART_HAS_FIFO) {
370:         outport_b(s->addr + UART_FCR, UART_FCR_FIFO | UART_FCR_FIFO14);
371:     }
372:     outport_b(s->addr + UART_MCR, UART_MCR_OUT2 | UART_MCR_RTS | UART_MCR_DT
R);
373:
374:     /* enable interrupts */
375:     outport_b(s->addr + UART_IER, UART_IER_RDAI);
376:
377:     /* clear all input registers */
378:     inport_b(s->addr + UART_RD);
379:     inport_b(s->addr + UART_IIR);
380:     inport_b(s->addr + UART_LSR);
381:     inport_b(s->addr + UART_MSR);
382:
383:     return 0;
384: }
385:
386: int serial_close(struct tty *tty)
387: {
388:     struct serial *s;
389:     int minor;
390:
391:     minor = MINOR(tty->dev);
392:     if(!TEST_MINOR(serial_device.minors, minor)) {
393:         return -ENXIO;
394:     }

```

## drivers/char/serial.c

Page 7/9

```

395:
396:     s = (struct serial *)tty->driver_data;
397:
398:     if(tty->count > 1) {
399:         return 0;
400:     }
401:
402:     /* disable all interrupts */
403:     outport_b(s->addr + UART_IER, 0);
404:
405:     /* disable FIFO */
406:     outport_b(s->addr + UART_FCR, UART_FCR_CRCVCR | UART_FCR_CXMTR);
407:
408:     /* clear all input register */
409:     inport_b(s->addr + UART_RD);
410:
411:     return 0;
412: }
413:
414: void serial_set_termios(struct tty *tty)
415: {
416:     short int divisor;
417:     int baud, size, stop;
418:     int lctrl;
419:     struct serial *s;
420:
421:     s = (struct serial *)tty->driver_data;
422:     lctrl = 0;
423:
424:     if(!(baud = tty->termios.c_cflag & CBAUD)) {
425:         return;
426:     }
427:     divisor = 115200 / baud_table[baud];
428:
429:     outport_b(s->addr + UART_LCR, UART_LCR_DLAB); /* enable DLAB */
430:     outport_b(s->addr + UART_DLL, divisor & 0xFF); /* LSB of divisor */
431:     outport_b(s->addr + UART_DLH, divisor >> 8); /* MSB of divisor */
432:
433:     size = tty->termios.c_cflag & CSIZE;
434:     switch(size) {
435:         case CS5:
436:             lctrl = UART_LCR_WL5;
437:             break;
438:         case CS6:
439:             lctrl = UART_LCR_WL6;
440:             break;
441:         case CS7:
442:             lctrl = UART_LCR_WL7;
443:             break;
444:         case CS8:
445:             lctrl = UART_LCR_WL8;
446:             break;
447:         default:
448:             lctrl = UART_LCR_WL5;
449:             break;
450:     }
451:
452:     stop = tty->termios.c_cflag & CSTOPB;
453:     if(stop) {
454:         lctrl |= UART_LCR_2STB;
455:     } else {
456:         lctrl |= UART_LCR_1STB;
457:     }
458:
459:     if(tty->termios.c_cflag & PARENB) {
460:         lctrl |= UART_LCR_EP;
461:     } else if(tty->termios.c_cflag & PARODD) {

```

## drivers/char/serial.c

Page 8/9

```

462:         lctrl |= UART_LCR_OP;
463:     } else {
464:         lctrl |= UART_LCR_NP;
465:     }
466:
467:     /* FIXME: flow control RTSCTS no supported */
468:
469:     outport_b(s->addr + UART_LCR, lctrl); /* line control */
470: }
471:
472: void serial_write(struct tty *tty)
473: {
474:     struct serial *s;
475:     unsigned long int flags;
476:
477:     SAVE_FLAGS(flags); CLI();
478:     s = (struct serial *)tty->driver_data;
479:     outport_b(s->addr + UART_IER, UART_IER_RDAI | UART_IER_THREI);
480:     RESTORE_FLAGS(flags);
481: }
482:
483: void irq_serial_bh(void)
484: {
485:     struct tty *tty;
486:     struct serial *s;
487:
488:     s = serial_ports;
489:
490:     if(s) {
491:         do {
492:             tty = s->tty;
493:             if(tty->read_q.count) {
494:                 if(!lock_area(AREA_SERIAL_READ)) {
495:                     tty->input(tty);
496:                     unlock_area(AREA_SERIAL_READ);
497:                 } else {
498:                     serial_bh.flags |= BH_ACTIVE;
499:                 }
500:             }
501:             s = s->next;
502:         } while(s);
503:     }
504: }
505:
506: void serial_init(void)
507: {
508:     int n, n2, type, found;
509:     struct serial **sp, *s;
510:     struct tty *tty;
511:
512:     for(n = 0, found = 0; n < SERIAL_MINORS; n++) {
513:         s = &serial_table[n];
514:         if((type = serial_identify(s)) {
515:             printk("ttyS%d 0x%04X-0x%04X %d type=%s%s\n",
n, s->addr, s->addr + 7, s->irq, serial_chip[type], s->flags & UART_HAS_FIFO ? " FIFO=y
es" : "");
516:
517:             SET_MINOR(serial_device.minors, (1 << SERIAL_MSF) + n);
518:             serial_default(s);
519:             sp = &serial_ports;
520:             if(*sp) {
521:                 do {
522:                     sp = &(*sp)->next;
523:                 } while(*sp);
524:             }
525:             if(!register_tty(MKDEV(SERIAL_MAJOR, (1 << SERIAL_MSF) +
n))) {

```

## drivers/char/serial.c

Page 9/9

```

526:         tty = get_tty(MKDEV(SERIAL_MAJOR, (1 << SERIAL_M
SF) + n));
527:         tty->driver_data = (void *)s;
528:         tty->stop = serial_stop;
529:         tty->start = serial_start;
530:         tty->deltab = serial_deltab;
531:         tty->reset = serial_reset;
532:         tty->input = do_cook;
533:         tty->output = serial_write;
534:         tty->open = serial_open;
535:         tty->close = serial_close;
536:         tty->set_termios = serial_set_termios;
537:         serial_reset(tty);
538:         tty_queue_init(tty);
539:         for(n2 = 0; n2 < MAX_TAB_COLS; n2++) {
540:             if(!(n2 % TAB_SIZE)) {
541:                 tty->tab_stop[n2] = 1;
542:             } else {
543:                 tty->tab_stop[n2] = 0;
544:             }
545:         }
546:         tty->count = 0;
547:         s->tty = tty;
548:         *sp = s;
549:         found++;
550:     } else {
551:         printk("WARNING: %s(): unable to register ttyS%d
.\n", __FUNCTION__, n);
552:     }
553: }
554: }
555: if(found) {
556:     add_bh(&serial_bh);
557:     if(register_device(CHR_DEV, &serial_device)) {
558:         printk("WARNING: %s(): unable to register serial device.
\n", __FUNCTION__);
559:     }
560:     if(!register_irq(SERIAL4_IRQ, &irq_config_serial0)) {
561:         enable_irq(SERIAL4_IRQ);
562:     }
563:     if(found > 1) {
564:         if(!register_irq(SERIAL3_IRQ, &irq_config_serial1)) {
565:             enable_irq(SERIAL3_IRQ);
566:         }
567:     }
568: }
569: if(is_serial(_syscondev)) {
570:     register_console(console_flush_log_buf);
571: }
572: }
573: }

```

## drivers/char/tty.c

Page 1/16

```

1: /*
2:  * fiwix/drivers/char/tty.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/ioctl.h>
10: #include <fiwix/tty.h>
11: #include <fiwix/ctype.h>
12: #include <fiwix/console.h>
13: #include <fiwix/pic.h>
14: #include <fiwix/devices.h>
15: #include <fiwix/fs.h>
16: #include <fiwix/errno.h>
17: #include <fiwix/sched.h>
18: #include <fiwix/timer.h>
19: #include <fiwix/sleep.h>
20: #include <fiwix/process.h>
21: #include <fiwix/fcntl.h>
22: #include <fiwix/kd.h>
23: #include <fiwix/stdio.h>
24: #include <fiwix/string.h>
25:
26: struct tty tty_table[NR_TTYS];
27: extern short int current_cons;
28:
29: static void wait_vtime_off(unsigned int arg)
30: {
31:     unsigned int *fn = (unsigned int *)arg;
32:
33:     wakeup(fn);
34: }
35:
36: static void termios2termio(struct termios *termios, struct termio *termio)
37: {
38:     int n;
39:
40:     termio->c_iflag = termios->c_iflag;
41:     termio->c_oflag = termios->c_oflag;
42:     termio->c_cflag = termios->c_cflag;
43:     termio->c_lflag = termios->c_lflag;
44:     termio->c_line = termios->c_line;
45:     for(n = 0; n < NCC; n++) {
46:         termio->c_cc[n] = termios->c_cc[n];
47:     }
48: }
49:
50: static void termio2termios(struct termio *termio, struct termios *termios)
51: {
52:     int n;
53:
54:     termios->c_iflag = termio->c_iflag;
55:     termios->c_oflag = termio->c_oflag;
56:     termios->c_cflag = termio->c_cflag;
57:     termios->c_lflag = termio->c_lflag;
58:     termios->c_line = termio->c_line;
59:     for(n = 0; n < NCC; n++) {
60:         termios->c_cc[n] = termio->c_cc[n];
61:     }
62: }
63:
64: static int opost(struct tty *tty, unsigned char ch)
65: {
66:     int status;
67:

```

## drivers/char/tty.c

Page 2/16

```

68:         status = 0;
69:
70:         if(tty->termios.c_oflag & OPOST) {
71:             switch(ch) {
72:                 case '\n':
73:                     if(tty->termios.c_oflag & ONLCR) {
74:                         if(tty_queue_room(&tty->write_q) >= 2) {
75:                             tty_queue_putchar(tty, &tty->wri
te_q, '\r');
76:                                     tty->column = 0;
77:                                 } else {
78:                                     return -1;
79:                                 }
80:                             }
81:                             break;
82:                 case '\t':
83:                     while(tty->column < (tty->winsize.ws_col - 1)) {
84:                         if(tty->tab_stop[++tty->column]) {
85:                             break;
86:                         }
87:                     }
88:                     break;
89:                 case '\b':
90:                     if(tty->column > 0) {
91:                         tty->column--;
92:                     }
93:                 default:
94:                     if(tty->termios.c_oflag & OLCUC) {
95:                         ch = TOUPPER(ch);
96:                     }
97:                     if(!ISCNTRL(ch)) {
98:                         tty->column++;
99:                     }
100:                    break;
101:            }
102:        }
103:        if(tty_queue_putchar(tty, &tty->write_q, ch) < 0) {
104:            status = -1;
105:        }
106:        return status;
107:    }
108:
109:    static void out_char(struct tty *tty, unsigned char ch)
110:    {
111:        if(ISCNTRL(ch) && !ISSPACE(ch) && (tty->termios.c_lflag & ECHOCTL)) {
112:            if(tty->lnext || (!tty->lnext && ch != tty->termios.c_cc[VEOF]))
113:            {
114:                tty_queue_putchar(tty, &tty->write_q, '^');
115:                tty_queue_putchar(tty, &tty->write_q, ch + 64);
116:                tty->column += 2;
117:            }
118:        } else {
119:            opost(tty, ch);
120:        }
121:    }
122:
123:    static void erase_char(struct tty *tty, unsigned char erasechar)
124:    {
125:        unsigned char ch;
126:
127:        if(erasechar == tty->termios.c_cc[VERASE]) {
128:            if((ch = tty_queue_unputchar(&tty->cooked_q)) &&
129:            if(tty->termios.c_lflag & ECHO) {
130:                tty_queue_putchar(tty, &tty->write_q, '\b');
131:                tty_queue_putchar(tty, &tty->write_q, ' ');
132:                tty_queue_putchar(tty, &tty->write_q, '\b');
133:            }
134:            if(ch == '\t') {

```



## drivers/char/tty.c

Page 3/16

```

133:             tty->deltab(tty);
134:         }
135:         if(ISCNTRL(ch) && !ISSPACE(ch) && tty->termios.c
_lflag & ECHOCTL) {
136:             tty_queue_putchar(tty, &tty->write_q, '\
b');
137:             tty_queue_putchar(tty, &tty->write_q, '
');
138:             tty_queue_putchar(tty, &tty->write_q, '\
b');
139:         }
140:     }
141: }
142: }
143: if(erasechar == tty->termios.c_cc[VWERASE]) {
144:     unsigned char word_seen = 0;
145:
146:     while(tty->cooked_q.count > 0) {
147:         ch = LAST_CHAR(&tty->cooked_q);
148:         if((ch == ' ' || ch == '\t') && word_seen) {
149:             break;
150:         }
151:         if(ch != ' ' && ch != '\t') {
152:             word_seen = 1;
153:         }
154:         erase_char(tty, tty->termios.c_cc[VERASE]);
155:     }
156: }
157: if(erasechar == tty->termios.c_cc[VKILL]) {
158:     while(tty->cooked_q.count > 0) {
159:         erase_char(tty, tty->termios.c_cc[VERASE]);
160:     }
161:     if(tty->termios.c_lflag & ECHOK && !(tty->termios.c_lflag & ECHO
E)) {
162:         tty_queue_putchar(tty, &tty->write_q, '\n');
163:     }
164: }
165: }
166:
167: static void set_termios(struct tty *tty, struct termios *new_termios)
168: {
169:     memcpy_b(&tty->termios, new_termios, sizeof(struct termios));
170:     if(tty->set_termios) {
171:         tty->set_termios(tty);
172:     }
173: }
174:
175: static void set_termio(struct tty *tty, struct termio *new_termio)
176: {
177:     struct termios new_termios;
178:
179:     termio2termios(new_termio, &new_termios);
180: }
181:
182: int register_tty(__dev_t dev)
183: {
184:     int n;
185:
186:     for(n = 0; n < NR_TTYS; n++) {
187:         if(tty_table[n].dev == dev) {
188:             printk("ERROR: %s(): tty device %d,%d already registered
!\n", __FUNCTION__, MAJOR(dev), MINOR(dev));
189:             return 1;
190:         }
191:         if(!tty_table[n].dev) {
192:             tty_table[n].dev = dev;
193:             tty_table[n].count = 0;

```

## drivers/char/tty.c

Page 4/16

```

194:             return 0;
195:         }
196:     }
197:     printk("ERROR: %s(): tty table is full!\n", __FUNCTION__);
198:     return 1;
199: }
200:
201: struct tty * get_tty(__dev_t dev)
202: {
203:     int n;
204:
205:     if(!dev) {
206:         return NULL;
207:     }
208:
209:     /* /dev/console = system console */
210:     if(dev == MKDEV(SYSCON_MAJOR, 1)) {
211:         dev = (__dev_t)_syscondev;
212:     }
213:
214:     /* /dev/tty0 = current virtual console */
215:     if(dev == MKDEV(VCONSOLES_MAJOR, 0)) {
216:         dev = MKDEV(VCONSOLES_MAJOR, current_cons);
217:     }
218:
219:     /* /dev/tty = controlling TTY device */
220:     if(dev == MKDEV(SYSCON_MAJOR, 0)) {
221:         if(!current->ctty) {
222:             return NULL;
223:         }
224:         dev = current->ctty->dev;
225:     }
226:
227:     for(n = 0; n < NR_TTYS; n++) {
228:         if(tty_table[n].dev == dev) {
229:             return &tty_table[n];
230:         }
231:     }
232:     return NULL;
233: }
234:
235: void disassociate_ctty(struct tty *tty)
236: {
237:     struct proc *p;
238:
239:     if(!tty) {
240:         return;
241:     }
242:
243:     /* this tty is no longer the controlling tty of any session */
244:     tty->pgid = tty->sid = 0;
245:
246:     /* clear the controlling tty for all processes in the same SID */
247:     FOR_EACH_PROCESS(p) {
248:         if((p->state != PROC_UNUSED) && (p->sid == current->sid)) {
249:             p->ctty = NULL;
250:         }
251:     }
252:     kill_pgrp(current->pgid, SIGHUP);
253:     kill_pgrp(current->pgid, SIGCONT);
254: }
255:
256: void termios_reset(struct tty *tty)
257: {
258:     tty->kbd.mode = K_XLATE;
259:     tty->termios.c_iflag = ICRNL | IXON | IXOFF;
260:     tty->termios.c_oflag = OPOST | ONLCR;

```

## drivers/char/tty.c

Page 5/16

```

261:         tty->termios.c_cflag = B9600 | CS8 | HUPCL | CREAD | CLOCAL;
262:         tty->termios.c_lflag = ISIG | ICANON | ECHO | ECHOE | ECHOK | ECHOCTL |
ECHOKE | IEXTEN;
263:         tty->termios.c_line = 0;
264:         tty->termios.c_cc[VINTR]     = 3;          /* ^C */
265:         tty->termios.c_cc[VQUIT]    = 28;         /* ^\ */
266:         tty->termios.c_cc[VERASE]   = BS;         /* ^? (127) not '\b' (^H) */
267:         tty->termios.c_cc[VKILL]    = 21;         /* ^U */
268:         tty->termios.c_cc[VEOF]     = 4;          /* ^D */
269:         tty->termios.c_cc[VTIME]    = 0;
270:         tty->termios.c_cc[VMIN]     = 1;
271:         tty->termios.c_cc[VSWTC]    = 0;
272:         tty->termios.c_cc[VSTART]   = 17;         /* ^Q */
273:         tty->termios.c_cc[VSTOP]    = 19;         /* ^S */
274:         tty->termios.c_cc[VSUSP]    = 26;         /* ^Z */
275:         tty->termios.c_cc[VEOL]     = '\n';       /* ^J */
276:         tty->termios.c_cc[VREPRINT] = 18;         /* ^R */
277:         tty->termios.c_cc[VDISCARD] = 15;         /* ^O */
278:         tty->termios.c_cc[VWERASE]  = 23;         /* ^W */
279:         tty->termios.c_cc[VLNEXT]   = 22;         /* ^V */
280:         tty->termios.c_cc[VEOL2]    = 0;
281:     }
282:
283:     void do_cook(struct tty *tty)
284:     {
285:         int n;
286:         unsigned char ch;
287:         struct cblock *cb;
288:
289:         while(tty->read_q.count > 0) {
290:             ch = tty_queue_getchar(&tty->read_q);
291:
292:             if((tty->termios.c_lflag & ISIG) && !tty->lnext) {
293:                 if(ch == tty->termios.c_cc[VINTR]) {
294:                     if(!(tty->termios.c_lflag & NOFLSH)) {
295:                         tty_queue_flush(&tty->read_q);
296:                         tty_queue_flush(&tty->cooked_q);
297:                     }
298:                     if(tty->pgid > 0) {
299:                         kill_pgrp(tty->pgid, SIGINT);
300:                     }
301:                     break;
302:                 }
303:                 if(ch == tty->termios.c_cc[VQUIT]) {
304:                     if(tty->pgid > 0) {
305:                         kill_pgrp(tty->pgid, SIGQUIT);
306:                     }
307:                     break;
308:                 }
309:                 if(ch == tty->termios.c_cc[VSUSP]) {
310:                     if(tty->pgid > 0) {
311:                         kill_pgrp(tty->pgid, SIGTSTP);
312:                     }
313:                     break;
314:                 }
315:             }
316:
317:             if(tty->termios.c_iflag & ISTRIP) {
318:                 ch = TOASCII(ch);
319:             }
320:             if(tty->termios.c_iflag & IUCLC) {
321:                 if(ISUPPER(ch)) {
322:                     ch = TOLOWER(ch);
323:                 }
324:             }
325:
326:             if(!tty->lnext) {

```

## drivers/char/tty.c

Page 6/16

```

327:         if(ch == '\r') {
328:             if(tty->termios.c_iflag & IGNCR) {
329:                 continue;
330:             }
331:             if(tty->termios.c_iflag & ICRNL) {
332:                 ch = '\n';
333:             }
334:         } else {
335:             if(ch == '\n') {
336:                 if(tty->termios.c_iflag & INLCR) {
337:                     ch = '\r';
338:                 }
339:             }
340:         }
341:     }
342:
343:     if(tty->termios.c_lflag & ICANON && !tty->lnext) {
344:         if(ch == tty->termios.c_cc[VERASE] || ch == tty->termios
.c_cc[VWERASE] || ch == tty->termios.c_cc[VKILL]) {
345:             erase_char(tty, ch);
346:             continue;
347:         }
348:
349:         if(ch == tty->termios.c_cc[VREPRINT]) {
350:             out_char(tty, ch);
351:             tty_queue_putchar(tty, &tty->write_q, '\n');
352:             cb = tty->cooked_q.head;
353:             while(cb) {
354:                 for(n = 0; n < cb->end_off; n++) {
355:                     if(n >= cb->start_off) {
356:                         out_char(tty, cb->data[n
]);
357:                     }
358:                 }
359:                 cb = cb->next;
360:             }
361:             continue;
362:         }
363:
364:         if(ch == tty->termios.c_cc[VLNEXT] && tty->termios.c_lfl
ag & IEXTEN) {
365:             tty->lnext = 1;
366:             if(tty->termios.c_lflag & ECHOCTL) {
367:                 tty_queue_putchar(tty, &tty->write_q, '^
');
368:                 tty_queue_putchar(tty, &tty->write_q, '\
b');
369:             }
370:             break;
371:         }
372:
373:         if(tty->termios.c_iflag & IXON) {
374:             if(ch == tty->termios.c_cc[VSTART]) {
375:                 tty->start(tty);
376:                 continue;
377:             }
378:             if(ch == tty->termios.c_cc[VSTOP]) {
379:                 tty->stop(tty);
380:                 continue;
381:             }
382:             if(tty->termios.c_iflag & IXANY) {
383:                 tty->start(tty);
384:             }
385:         }
386:     }
387:
388:     /* using ISSPACE here makes LNEXT working incorrectly, FIXME */

```

## drivers/char/tty.c

Page 7/16

```

389:             if(tty->termios.c_lflag & ICANON) {
390:                 if(ISCNTRL(ch) && !ISSPACE(ch) && (tty->termios.c_lflag
& ECHOCTL)) {
391:                     out_char(tty, ch);
392:                     tty_queue_putchar(tty, &tty->cooked_q, ch);
393:                     tty->lnext = 0;
394:                     continue;
395:                 }
396:                 if(ch == '\n') {
397:                     tty->canon_data = 1;
398:                 }
399:             }
400:
401:             if(tty->termios.c_lflag & ECHO) {
402:                 out_char(tty, ch);
403:             } else {
404:                 if((tty->termios.c_lflag & ECHONL) && (ch == '\n')) {
405:                     out_char(tty, ch);
406:                 }
407:             }
408:             tty_queue_putchar(tty, &tty->cooked_q, ch);
409:             tty->lnext = 0;
410:         }
411:         tty->output(tty);
412:         wakeup(&tty_read);
413:         if(!(tty->termios.c_lflag & ICANON) || ((tty->termios.c_lflag & ICANON)
&& tty->canon_data)) {
414:             wakeup(&do_select);
415:         }
416:     }
417:
418: int tty_open(struct inode *i, struct fd *fd_table)
419: {
420:     int noctty_flag;
421:     struct tty *tty;
422:     int errno;
423:
424:     noctty_flag = fd_table->flags & O_NOCTTY;
425:
426:     if(MAJOR(i->rdev) == SYSCON_MAJOR && MINOR(i->rdev) == 0) {
427:         if(!current->ctty) {
428:             return -ENXIO;
429:         }
430:     }
431:
432:     if(MAJOR(i->rdev) == VCONSOLES_MAJOR && MINOR(i->rdev) == 0) {
433:         noctty_flag = 1;
434:     }
435:
436:     if(!(tty = get_tty(i->rdev))) {
437:         printk("%s(): oops! (%x)\n", __FUNCTION__, i->rdev);
438:         printk("_syscondev = %x\n", _syscondev);
439:         return -ENXIO;
440:     }
441:
442:     if(tty->open) {
443:         if((errno = tty->open(tty) < 0)) {
444:             return errno;
445:         }
446:     }
447:     tty->count++;
448:     tty->column = 0;
449:
450:     if(SESS_LEADER(current) && !current->ctty && !noctty_flag && !tty->sid)
{
451:         current->ctty = tty;
452:         tty->sid = current->sid;

```

## drivers/char/tty.c

Page 8/16

```

453:         tty->pgid = current->pgid;
454:     }
455:     return 0;
456: }
457:
458: int tty_close(struct inode *i, struct fd *fd_table)
459: {
460:     struct proc *p;
461:     struct tty *tty;
462:     int errno;
463:
464:     if(!(tty = get_tty(i->rdev))) {
465:         printk("%s(): oops! (%x)\n", __FUNCTION__, i->rdev);
466:         return -ENXIO;
467:     }
468:
469:     if(tty->close) {
470:         if((errno = tty->close(tty) < 0)) {
471:             return errno;
472:         }
473:     }
474:     tty->count--;
475:     if(!tty->count) {
476:         tty->reset(tty);
477:         termios_reset(tty);
478:         tty->pgid = tty->sid = 0;
479:
480:         /* this tty is no longer the controlling tty of any process */
481:         FOR_EACH_PROCESS(p) {
482:             if((p->state != PROC_UNUSED) && (p->ctty == tty)) {
483:                 p->ctty = NULL;
484:             }
485:         }
486:     }
487:     return 0;
488: }
489:
490: int tty_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t count)
491: {
492:     unsigned int min;
493:     unsigned char ch;
494:     struct tty *tty;
495:     struct callout_req creq;
496:     int n;
497:
498:     if(!(tty = get_tty(i->rdev))) {
499:         printk("%s(): oops! (%x)\n", __FUNCTION__, i->rdev);
500:         return -ENXIO;
501:     }
502:
503:     /* only the foreground process group is allowed to read from the tty */
504:     if(current->ctty == tty && current->pgid != tty->pgid) {
505:         if(current->sigaction[SIGTTIN - 1].sa_handler == SIG_IGN || curr
ent->sigblocked & (1 << (SIGTTIN - 1)) || is_orphaned_pgrp(current->pgid)) {
506:             return -EIO;
507:         }
508:         kill_pgrp(current->pgid, SIGTTIN);
509:         return -ERESTART;
510:     }
511:
512:     n = min = 0;
513:     while(count > 0) {
514:         if(tty->kbd.mode == K_RAW || tty->kbd.mode == K_MEDIUMRAW) {
515:             n = 0;
516:             while(n < count) {
517:                 if((ch = tty_queue_getchar(&tty->read_q))) {
518:                     buffer[n++] = ch;

```

## drivers/char/tty.c

Page 9/16

```

519:                } else {
520:                    break;
521:                }
522:            }
523:            if(n) {
524:                break;
525:            }
526:        }
527:
528:        if(tty->termios.c_lflag & ICANON) {
529:            if((ch = LAST_CHAR(&tty->cooked_q)) {
530:                if(ch == '\n' || ch == tty->termios.c_cc[VEOL] |
| ch == tty->termios.c_cc[VEOF] || (tty->termios.c_lflag & IEXTEN && ch == tty->termios
.c_cc[VEOL2] && tty->termios.c_cc[VEOL2] != 0)) {
531:
532:                tty->canon_data = 0;
533:                /* EOF is not passed to the reading proc
ess */
534:                if(ch == tty->termios.c_cc[VEOF]) {
535:                    tty_queue_unputchar(&tty->cooked
_q);
536:                }
537:
538:                while(n < count) {
539:                    if((ch = tty_queue_getchar(&tty-
>cooked_q)) {
540:                        buffer[n++] = ch;
541:                    } else {
542:                        break;
543:                    }
544:                }
545:                break;
546:            }
547:        }
548:    } else {
549:        if(tty->termios.c_cc[VTIME] > 0) {
550:            unsigned int ini_ticks = kstat.ticks;
551:            unsigned int timeout;
552:
553:            if(!tty->termios.c_cc[VMIN]) {
554:                /* VTIME is measured in tenths of second
*/
555:                timeout = tty->termios.c_cc[VTIME] * (HZ
/ 10);
556:
557:                while(kstat.ticks - ini_ticks < timeout
&& !tty->cooked_q.count) {
558:                    creq.fn = wait_vtime_off;
559:                    creq.arg = (unsigned int)&tty->c
ooked_q;
560:                    add_callout(&creq, timeout);
561:                    if(fd_table->flags & O_NONBLOCK)
{
562:                        return -EAGAIN;
563:                    }
564:                    if(sleep(&tty_read, PROC_INTERRU
PTIBLE)) {
565:                        return -EINTR;
566:                    }
567:                }
568:                while(n < count) {
569:                    if((ch = tty_queue_getchar(&tty-
>cooked_q)) {
570:                        buffer[n++] = ch;
571:                    } else {
572:                        break;
573:                    }

```

## drivers/char/tty.c

Page 10/16

```

574:                                     }
575:                                     break;
576:         } else {
577:             if(tty->cooked_q.count > 0) {
578:                 if(n < MIN(tty->termios.c_cc[VMI
N], count)) {
579:                     ch = tty_queue_getchar(&
tty->cooked_q);
580:                     buffer[n++] = ch;
581:                 }
582:                 if(n >= MIN(tty->termios.c_cc[VM
IN], count)) {
583:                     del_callout(&creq);
584:                     break;
585:                 }
586:                 timeout = tty->termios.c_cc[VTIM
E] * (HZ / 10);
587:                 creq.fn = wait_vtime_off;
588:                 creq.arg = (unsigned int)&tty->c
ooked_q;
589:                 add_callout(&creq, timeout);
590:                 if(fd_table->flags & O_NONBLOCK)
{
591:                     n = -EAGAIN;
592:                     break;
593:                 }
594:                 if(sleep(&tty_read, PROC_INTERRU
PTIBLE)) {
595:                     n = -EINTR;
596:                     break;
597:                 }
598:                 if(!tty->cooked_q.count) {
599:                     break;
600:                 }
601:                 continue;
602:             }
603:         }
604:     } else {
605:         if(tty->cooked_q.count > 0) {
606:             if(min < tty->termios.c_cc[VMIN] || !tty
->termios.c_cc[VMIN]) {
607:                 if(n < count) {
608:                     ch = tty_queue_getchar(&
tty->cooked_q);
609:                     buffer[n++] = ch;
610:                 }
611:                 min++;
612:             }
613:         }
614:         if(min >= tty->termios.c_cc[VMIN]) {
615:             break;
616:         }
617:     }
618: }
619: if(fd_table->flags & O_NONBLOCK) {
620:     n = -EAGAIN;
621:     break;
622: }
623: if(sleep(&tty_read, PROC_INTERRUPTIBLE)) {
624:     n = -EINTR;
625:     break;
626: }
627: }
628:
629: if(n) {
630:     i->i_atime = CURRENT_TIME;
631: }

```



## drivers/char/tty.c

Page 11/16

```

632:         return n;
633:     }
634:
635:     int tty_write(struct inode *i, struct fd *fd_table, const char *buffer, __size_t
count)
636:     {
637:         unsigned char ch;
638:         struct tty *tty;
639:         int n;
640:
641:         if(!(tty = get_tty(i->rdev))) {
642:             printk("%s(): oops! (%x)\n", __FUNCTION__, i->rdev);
643:             return -ENXIO;
644:         }
645:
646:         /* only the foreground process group is allowed to write to the tty */
647:         if(current->ctty == tty && current->pgid != tty->pgid) {
648:             if(tty->termios.c_lflag & TOSTOP) {
649:                 if(current->sigaction[SIGTTIN - 1].sa_handler != SIG_IGN
&& !(current->sigblocked & (1 << (SIGTTIN - 1)))) {
650:                     if(is_orphaned_pgrp(current->pgid)) {
651:                         return -EIO;
652:                     }
653:                     kill_pgrp(current->pgid, SIGTTOU);
654:                     return -ERESTART;
655:                 }
656:             }
657:         }
658:
659:         n = 0;
660:         for(;;) {
661:             if(current->sigpending & ~current->sigblocked) {
662:                 return -ERESTART;
663:             }
664:             while(count && n < count) {
665:                 ch = *(buffer + n);
666:                 /* FIXME: check if *(buffer + n) address is valid */
667:                 if(opost(tty, ch) < 0) {
668:                     break;
669:                 }
670:                 n++;
671:             }
672:             tty->output(tty);
673:             if(n == count) {
674:                 break;
675:             }
676:             if(fd_table->flags & O_NONBLOCK) {
677:                 n = -EAGAIN;
678:                 break;
679:             }
680:             if(tty->write_q.count > 0) {
681:                 if(sleep(&tty_write, PROC_INTERRUPTIBLE)) {
682:                     n = -EINTR;
683:                     break;
684:                 }
685:             }
686:             do_sched();
687:         }
688:
689:         if(n) {
690:             i->i_mtime = CURRENT_TIME;
691:         }
692:         return n;
693:     }
694:
695:     /* FIXME: http://www.lafn.org/~dave/linux/termios.txt (doc/termios.txt) */
696:     int tty_ioctl(struct inode *i, int cmd, unsigned long int arg)

```

## drivers/char/tty.c

Page 12/16

```

697: {
698:     struct proc *p;
699:     struct tty *tty;
700:     int errno;
701:
702:     if(!(tty = get_tty(i->rdev))) {
703:         printk("%s(): oops! (%x)\n", __FUNCTION__, i->rdev);
704:         return -ENXIO;
705:     }
706:
707:     switch(cmd) {
708:         /*
709:          * Fetch and store the current terminal parameters to a termios
710:          * structure pointed to by the argument.
711:          */
712:         case TCGETS:
713:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
izeof(struct termios)))) {
714:                 return errno;
715:             }
716:             memcpy_b((struct termios *)arg, &tty->termios, sizeof(st
ruct termios));
717:             break;
718:
719:         /*
720:          * Set the current terminal parameters according to the
721:          * values in the termios structure pointed to by the argument.
722:          */
723:         case TCSETS:
724:             if((errno = check_user_area(VERIFY_READ, (void *)arg, si
zeof(struct termios)))) {
725:                 return errno;
726:             }
727:             set_termios(tty, (struct termios *)arg);
728:             break;
729:
730:         /*
731:          * Same as TCSETS except it doesn't take effect until all
732:          * the characters queued for output have been transmitted.
733:          */
734:         case TCSETSW:
735:             if((errno = check_user_area(VERIFY_READ, (void *)arg, si
zeof(struct termios)))) {
736:                 return errno;
737:             }
738:             set_termios(tty, (struct termios *)arg);
739:             break;
740:
741:         /*
742:          * Same as TCSETSW except that all characters queued for
743:          * input are discarded.
744:          */
745:         case TCSETSF:
746:             if((errno = check_user_area(VERIFY_READ, (void *)arg, si
zeof(struct termios)))) {
747:                 return errno;
748:             }
749:             set_termios(tty, (struct termios *)arg);
750:             tty_queue_flush(&tty->read_q);
751:             break;
752:
753:         /*
754:          * Fetch and store the current terminal parameters to a termio
755:          * structure pointed to by the argument.
756:          */
757:         case TCGETA:
758:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s

```

## drivers/char/tty.c

Page 13/16

```

759:         return errno;
760:     }
761:     termios2termio(&tty->termios, (struct termio *)arg);
762:     break;
763:
764:     /*
765:     * Set the current terminal parameters according to the
766:     * values in the termio structure pointed to by the argument.
767:     */
768:     case TCSETA:
769:         if((errno = check_user_area(VERIFY_READ, (void *)arg, si
zeof(struct termio)))) {
770:             return errno;
771:         }
772:         set_termio(tty, (struct termio *)arg);
773:         break;
774:
775:     /*
776:     * Same as TCSET except it doesn't take effect until all
777:     * the characters queued for output have been transmitted.
778:     */
779:     case TCSETAW:
780:         if((errno = check_user_area(VERIFY_READ, (void *)arg, si
zeof(struct termio)))) {
781:             return errno;
782:         }
783:         set_termio(tty, (struct termio *)arg);
784:         break;
785:
786:     /*
787:     * Same as TCSETAW except that all characters queued for
788:     * input are discarded.
789:     */
790:     case TCSETAF:
791:         if((errno = check_user_area(VERIFY_READ, (void *)arg, si
zeof(struct termio)))) {
792:             return errno;
793:         }
794:         set_termio(tty, (struct termio *)arg);
795:         tty_queue_flush(&tty->read_q);
796:         break;
797:
798:     /* Perform start/stop control */
799:     case TCXONC:
800:         switch(arg) {
801:             case TCOOFF:
802:                 tty->stop(tty);
803:                 break;
804:             case TCOON:
805:                 tty->start(tty);
806:                 break;
807:             default:
808:                 return -EINVAL;
809:         }
810:         break;
811:     case TCFLSH:
812:         switch(arg) {
813:             case TCIFLUSH:
814:                 tty_queue_flush(&tty->read_q);
815:                 tty_queue_flush(&tty->cooked_q);
816:                 break;
817:             case TCOFLUSH:
818:                 tty_queue_flush(&tty->write_q);
819:                 break;
820:             case TCIOFLUSH:
821:                 tty_queue_flush(&tty->read_q);

```

## drivers/char/tty.c

Page 14/16

```

822:                 tty_queue_flush(&tty->cooked_q);
823:                 tty_queue_flush(&tty->write_q);
824:                 break;
825:             default:
826:                 return -EINVAL;
827:         }
828:         break;
829:     case TIOCSCTTY:
830:         if(SESS_LEADER(current) && (current->sid == tty->sid)) {
831:             return 0;
832:         }
833:         if(!SESS_LEADER(current) || current->ctty) {
834:             return -EPERM;
835:         }
836:         if(tty->sid) {
837:             if((arg == 1) && IS_SUPERUSER) {
838:                 FOR_EACH_PROCESS(p) {
839:                     if((p->state != PROC_UNUSED) &&
(p->ctty == tty)) {
840:                         p->ctty = NULL;
841:                     }
842:                 }
843:             } else {
844:                 return -EPERM;
845:             }
846:         }
847:         current->ctty = tty;
848:         tty->sid = current->sid;
849:         tty->pgid = current->pgid;
850:         break;
851:
852:         /*
853:          * Get the process group ID of the '__pid_t' pointed to by
854:          * the arg to the foreground processes group ID.
855:          */
856:     case TIOCGPRG:
857:         if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
sizeof(__pid_t)))) {
858:             return errno;
859:         }
860:         memcpy_b((void *)arg, &tty->pgid, sizeof(__pid_t));
861:         break;
862:
863:         /*
864:          * Associate the process pointed to by '__pid_t' in the arg to
865:          * the value of the terminal.
866:          */
867:     case TIOCSPGRP:
868:         if(arg < 1) {
869:             return -EINVAL;
870:         }
871:         if((errno = check_user_area(VERIFY_READ, (void *)arg, si
sizeof(__pid_t)))) {
872:             return errno;
873:         }
874:         memcpy_b(&tty->pgid, (void *)arg, sizeof(__pid_t));
875:         break;
876:
877:         /*
878:          * The session ID of the terminal is fetched and stored in
879:          * the '__pid_t' pointed to by the arg.
880:          case TIOCSID:    FIXME
881:          */
882:
883:         /*
884:          * The terminal drivers notion of terminal size is stored in
885:          * the 'winsize' structure pointed to by the arg.

```

## drivers/char/tty.c

Page 15/16

```

886:          */
887:          case TIOCGWINSZ:
888:              if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
sizeof(struct winsize)))) {
889:                  return errno;
890:              }
891:              memcpy_b((void *)arg, &tty->winsize, sizeof(struct winsi
ze));
892:              break;
893:
894:          /*
895:           * The terminal drivers notion of the terminal size is set
896:           * to value in the 'winsize' structure pointed to by the arg.
897:           */
898:          case TIOCSWINSZ:
899:          {
900:              struct winsize *ws = (struct winsize *)arg;
901:              short int changed;
902:
903:              if((errno = check_user_area(VERIFY_READ, (void *)arg, si
zeof(struct winsize)))) {
904:                  return errno;
905:              }
906:              changed = 0;
907:              if(tty->winsize.ws_row != ws->ws_row) {
908:                  changed = 1;
909:              }
910:              if(tty->winsize.ws_col != ws->ws_col) {
911:                  changed = 1;
912:              }
913:              if(tty->winsize.ws_xpixel != ws->ws_xpixel) {
914:                  changed = 1;
915:              }
916:              if(tty->winsize.ws_ypixel != ws->ws_ypixel) {
917:                  changed = 1;
918:              }
919:              tty->winsize.ws_row = ws->ws_row;
920:              tty->winsize.ws_col = ws->ws_col;
921:              tty->winsize.ws_xpixel = ws->ws_xpixel;
922:              tty->winsize.ws_ypixel = ws->ws_ypixel;
923:              if(changed) {
924:                  kill_pgrp(tty->pgid, SIGWINCH);
925:              }
926:          }
927:              break;
928:          case TIOCNOTTY:
929:              if(current->ctty != tty) {
930:                  return -ENOTTY;
931:              }
932:              if(SESS_LEADER(current)) {
933:                  disassociate_ctty(tty);
934:              }
935:              break;
936:          case TIOCLINUX:
937:          {
938:              int val = *(unsigned char *)arg;
939:              if((errno = check_user_area(VERIFY_READ, (void *)arg, si
zeof(unsigned char)))) {
940:                  return errno;
941:              }
942:              switch(val) {
943:                  case 12:          /* get current console */
944:                      return current_cons;
945:                      break;
946:                  default:
947:                      return -EINVAL;
948:                      break;

```

## drivers/char/tty.c

Page 16/16

```
949:             }
950:             break;
951:         }
952:
953:         default:
954:             return vt_ioctl(tty, cmd, arg);
955:     }
956:     return 0;
957: }
958:
959: int tty_lseek(struct inode *i, __off_t offset)
960: {
961:     return -ESPIPE;
962: }
963:
964: int tty_select(struct inode *i, int flag)
965: {
966:     struct tty *tty;
967:
968:     if(!(tty = get_tty(i->rdev))) {
969:         printk("%s(): oops! (%x)\n", __FUNCTION__, i->rdev);
970:         return 0;
971:     }
972:
973:     switch(flag) {
974:         case SEL_R:
975:             if(tty->cooked_q.count > 0) {
976:                 if(!(tty->termios.c_lflag & ICANON) || ((tty->te
rmios.c_lflag & ICANON) && tty->canon_data)) {
977:                     return 1;
978:                 }
979:             }
980:             break;
981:         case SEL_W:
982:             if(!tty->write_q.count) {
983:                 return 1;
984:             }
985:             break;
986:     }
987:     return 0;
988: }
989:
990: void tty_init(void)
991: {
992:     memset_b(tty_table, NULL, sizeof(tty_table));
993: }
```

## drivers/char/tty\_queue.c

Page 1/4

```

1: /*
2:  * fiwix/drivers/char/tty_queue.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/tty.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/string.h>
13:
14: /*
15:  * tty_queue.c implements a queue using a static-sized doubly linked list of a
16:  * central pool of buffers which covers all ttys.
17:  *
18:  * head                                     tail
19:  * +-----+ +-----+ ... +-----+
20:  * |prev|data|next| |prev|data|next| ... |prev|data|next|
21:  * | / | | --> <-- | | --> ... <-- | | / |
22:  * +-----+ +-----+ ... +-----+
23:  * (cblock)      (cblock)      (cblock)
24:  */
25:
26: struct cblock cblock_pool[CB_POOL_SIZE];
27: struct cblock *cblock_pool_head;
28:
29: static struct cblock *get_free_cblock(void)
30: {
31:     struct cblock *new = NULL;
32:
33:     if(cblock_pool_head) {
34:         new = cblock_pool_head;
35:         cblock_pool_head = cblock_pool_head->next;
36:         new->prev = new->next = NULL;
37:     }
38:     return new;
39: }
40:
41: static void put_free_cblock(struct cblock *old)
42: {
43:     old->prev = NULL;
44:     old->next = cblock_pool_head;
45:     cblock_pool_head = old;
46: }
47:
48: static struct cblock *insert_cblock_in_head(struct clist *q)
49: {
50:     struct cblock *cb;
51:
52:     if(q->cb_num >= NR_CB_QUEUE) {
53:         return NULL;
54:     }
55:     if(!(cb = get_free_cblock())) {
56:         return NULL;
57:     }
58:
59:     /* initialize cblock */
60:     cb->start_off = cb->end_off = 0;
61:     memset_b(cb->data, 0, CBSIZE);
62:     cb->prev = cb->next = NULL;
63:     q->cb_num++;
64:
65:     if(!q->head) {
66:         q->head = q->tail = cb;
67:     } else {

```

## drivers/char/tty\_queue.c

Page 2/4

```

68:         cb->prev = NULL;
69:         cb->next = q->head;
70:         q->head->prev = cb;
71:         q->head = cb;
72:     }
73:     return cb;
74: }
75:
76: static struct cblock *insert_cblock_in_tail(struct clist *q)
77: {
78:     struct cblock *cb;
79:
80:     if(q->cb_num >= NR_CB_QUEUE) {
81:         return NULL;
82:     }
83:     if(!(cb = get_free_cblock())) {
84:         return NULL;
85:     }
86:
87:     /* initialize cblock */
88:     cb->start_off = cb->end_off = 0;
89:     memset_b(cb->data, 0, CBSIZE);
90:     cb->prev = cb->next = NULL;
91:     q->cb_num++;
92:
93:     if(!q->tail) {
94:         q->head = q->tail = cb;
95:     } else {
96:         cb->prev = q->tail;
97:         cb->next = NULL;
98:         q->tail->next = cb;
99:         q->tail = cb;
100:    }
101:    return cb;
102: }
103:
104: static void delete_cblock_from_head(struct clist *q)
105: {
106:     struct cblock *tmp;
107:
108:     if(!q->head) {
109:         return;
110:     }
111:
112:     tmp = q->head;
113:     if(q->head == q->tail) {
114:         q->head = q->tail = NULL;
115:     } else {
116:         q->head = q->head->next;
117:         q->head->prev = NULL;
118:     }
119:
120:     q->count -= tmp->end_off - tmp->start_off;
121:     q->cb_num--;
122:     put_free_cblock(tmp);
123: }
124:
125: static void delete_cblock_from_tail(struct clist *q)
126: {
127:     struct cblock *tmp;
128:
129:     if(!q->tail) {
130:         return;
131:     }
132:
133:     tmp = q->tail;
134:     if(q->head == q->tail) {

```



## drivers/char/tty\_queue.c

Page 3/4

```

135:         q->head = q->tail = NULL;
136:     } else {
137:         q->tail = q->tail->prev;
138:         q->tail->next = NULL;
139:     }
140:
141:     q->count -= tmp->end_off - tmp->start_off;
142:     q->cb_num--;
143:     put_free_cblock(tmp);
144: }
145:
146: int tty_queue_putchar(struct tty *tty, struct clist *q, unsigned char ch)
147: {
148:     unsigned long int flags;
149:     struct cblock *cb;
150:     int errno;
151:
152:     SAVE_FLAGS(flags); CLI();
153:
154:     cb = q->tail;
155:     if(!cb) {
156:         cb = insert_cblock_in_tail(q);
157:         if(!cb) {
158:             RESTORE_FLAGS(flags);
159:             return -EAGAIN;
160:         }
161:     }
162:
163:     if(cb->end_off < CBSIZE) {
164:         cb->data[cb->end_off] = ch;
165:         cb->end_off++;
166:         q->count++;
167:         errno = 0;
168:     } else if(insert_cblock_in_tail(q)) {
169:         tty_queue_putchar(tty, q, ch);
170:         errno = 0;
171:     } else {
172:         errno = -EAGAIN;
173:     }
174:
175:     RESTORE_FLAGS(flags);
176:     return errno;
177: }
178:
179: int tty_queue_unputchar(struct clist *q)
180: {
181:     unsigned long int flags;
182:     struct cblock *cb;
183:     unsigned char ch;
184:
185:     SAVE_FLAGS(flags); CLI();
186:
187:     ch = 0;
188:     cb = q->tail;
189:     if(cb) {
190:         if(cb->end_off > cb->start_off) {
191:             ch = cb->data[cb->end_off - 1];
192:             cb->end_off--;
193:             q->count--;
194:         }
195:         if(cb->end_off - cb->start_off == 0) {
196:             delete_cblock_from_tail(q);
197:         }
198:     }
199:
200:     RESTORE_FLAGS(flags);
201:     return ch;

```

## drivers/char/tty\_queue.c

Page 4/4

```

202: }
203:
204: unsigned char tty_queue_getchar(struct clist *q)
205: {
206:     unsigned long int flags;
207:     struct cblock *cb;
208:     unsigned char ch;
209:
210:     SAVE_FLAGS(flags); CLI();
211:
212:     ch = 0;
213:     cb = q->head;
214:     if(cb) {
215:         if(cb->start_off < cb->end_off) {
216:             ch = cb->data[cb->start_off];
217:             cb->start_off++;
218:             q->count--;
219:         }
220:         if(cb->end_off - cb->start_off == 0) {
221:             delete_cblock_from_head(q);
222:         }
223:     }
224:
225:     RESTORE_FLAGS(flags);
226:     return ch;
227: }
228:
229: void tty_queue_flush(struct clist *q)
230: {
231:     unsigned long int flags;
232:
233:     SAVE_FLAGS(flags); CLI();
234:
235:     while(q->head != NULL) {
236:         delete_cblock_from_head(q);
237:     }
238:
239:     RESTORE_FLAGS(flags);
240: }
241:
242: int tty_queue_room(struct clist *q)
243: {
244:     return (NR_CB_QUEUE * CB_POOL_SIZE) - q->count;
245: }
246:
247: void tty_queue_init(struct tty *tty)
248: {
249:     int n;
250:     struct cblock *cb;
251:
252:     memset_b(cblock_pool, NULL, sizeof(cblock_pool));
253:
254:     /* cblock free list initialization */
255:     cblock_pool_head = NULL;
256:     n = CB_POOL_SIZE;
257:     while(n--) {
258:         cb = &cblock_pool[n];
259:         put_free_cblock(cb);
260:     }
261:     tty->read_q.head = tty->read_q.tail = NULL;
262:     tty->cooked_q.head = tty->cooked_q.tail = NULL;
263:     tty->write_q.head = tty->write_q.tail = NULL;
264: }

```

## drivers/char/vt.c

Page 1/4

```

1: /*
2:  * fiwix/drivers/char/vt.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/console.h>
10: #include <fiwix/keyboard.h>
11: #include <fiwix/tty.h>
12: #include <fiwix/vt.h>
13: #include <fiwix/kd.h>
14: #include <fiwix/errno.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: int kbdmode = 0;
19:
20: int vt_ioctl(struct tty *tty, int cmd, unsigned long int arg)
21: {
22:     struct vconsole *vc;
23:     int n, errno;
24:
25:     /* only virtual consoles support the following ioctl commands */
26:     if(MAJOR(tty->dev) != VCONSOLES_MAJOR) {
27:         return -ENXIO;
28:     }
29:
30:     vc = (struct vconsole *)tty->driver_data;
31:
32:     switch(cmd) {
33:         case KDGETLED:
34:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
sizeof(unsigned char)))) {
35:                 return errno;
36:             }
37:             memset_b((void *)arg, vc->led_status, sizeof(char));
38:             break;
39:
40:         case KDSETLED:
41:             if(arg > 7) {
42:                 return -EINVAL;
43:             }
44:             vc->led_status = arg;
45:             set_leds(vc->led_status);
46:             break;
47:
48:             /* FIXME: implement KDGKBLED and KDSKBLED
49:             * it will need to convert 'scrlock, numlock, capslock' into led
50:             */
51:
52:         case KDGKBTYPE:
53:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
sizeof(unsigned char)))) {
54:                 return errno;
55:             }
56:             memset_b((void *)arg, KB_101, sizeof(char));
57:             break;
58:
59:         case KDSETMODE:
60:             if(arg != KD_TEXT && arg != KD_GRAPHICS) {
61:                 return -EINVAL;
62:             }
63:             if(vc->vc_mode != arg) {
64:                 vc->vc_mode = arg;

```

## drivers/char/vt.c

Page 2/4

```

65:                                if(arg == KD_GRAPHICS) {
66:                                    blank_screen(vc);
67:                                } else {
68:                                    unblank_screen(vc);
69:                                }
70:                                }
71:                                break;
72:
73:                                case KDGETMODE:
74:                                    if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
izeof(unsigned char)))) {
75:                                        return errno;
76:                                    }
77:                                    memset_b((void *)arg, vc->vc_mode, sizeof(char));
78:                                    break;
79:
80:                                case KDGKBMODE:
81:                                    if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
izeof(unsigned char)))) {
82:                                        return errno;
83:                                    }
84:                                    memset_b((void *)arg, tty->kbd.mode, sizeof(unsigned cha
r));
85:                                    break;
86:
87:                                case KDSKBMODE:
88:                                    if(arg != K_RAW && arg != K_XLATE && arg != K_MEDIUMRAW)
{
89:                                        arg = K_XLATE;
90:                                    }
91:                                    tty->kbd.mode = arg;
92:                                    tty_queue_flush(&tty->read_q);
93:                                    break;
94:
95:                                case KDSKBENT:
96:                                    {
97:                                        struct kbentry *k = (struct kbentry *)arg;
98:                                        if((errno = check_user_area(VERIFY_WRITE, (void *)k, siz
eof(struct kbentry)))) {
99:                                            return errno;
100:                                        }
101:                                        if(k->kb_table < NR_MODIFIERS) {
102:                                            if(k->kb_index < NR_SCODES) {
103:                                                keymap[(k->kb_index * NR_MODIFIERS) + k-
>kb_table] = k->kb_value;
104:                                            } else {
105:                                                return -EINVAL;
106:                                            }
107:                                        } else {
108:                                            printk("%s(): kb_table value '%d' not supported.
\n", __FUNCTION__, k->kb_table);
109:                                            return -EINVAL;
110:                                        }
111:                                    }
112:                                    break;
113:
114:                                case VT_OPENQRY:
115:                                    {
116:                                        int *val = (int *)arg;
117:                                        if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
izeof(unsigned int)))) {
118:                                            return errno;
119:                                        }
120:                                        for(n = 1; n < NR_VCONSOLES + 1; n++) {
121:                                            tty = get_tty(MKDEV(VCONSOLES_MAJOR, n));
122:                                            if(!tty->count) {
123:                                                break;

```

## drivers/char/vt.c

Page 3/4

```

124:         }
125:     }
126:     *val = (n < NR_VCONSOLES + 1 ? n : -1);
127: }
128:     break;
129:
130:     case VT_GETMODE:
131:     {
132:         struct vt_mode *vt_mode = (struct vt_mode *)arg;
133:         if((errno = check_user_area(VERIFY_WRITE, (void *)vt_mode
e, sizeof(struct vt_mode)))) {
134:             return errno;
135:         }
136:         memcpy_b(vt_mode, &vc->vt_mode, sizeof(struct vt_mode));
137:     }
138:     break;
139:
140:     case VT_SETMODE:
141:     {
142:         struct vt_mode *vt_mode = (struct vt_mode *)arg;
143:         if((errno = check_user_area(VERIFY_READ, (void *)vt_mode
, sizeof(struct vt_mode)))) {
144:             return errno;
145:         }
146:         if(vt_mode->mode != VT_AUTO && vt_mode->mode != VT_PROCE
SS) {
147:             return -EINVAL;
148:         }
149:         memcpy_b(&vc->vt_mode, vt_mode, sizeof(struct vt_mode));
150:         vc->vt_mode.frsig = 0; /* ignored */
151:         tty->pid = current->pid;
152:         vc->switchto_tty = 0;
153:     }
154:     break;
155:
156:     case VT_GETSTATE:
157:     {
158:         struct vt_stat *vt_stat = (struct vt_stat *)arg;
159:         if((errno = check_user_area(VERIFY_WRITE, (void *)vt_sta
t, sizeof(struct vt_stat)))) {
160:             return errno;
161:         }
162:         vt_stat->v_active = current_cons;
163:         vt_stat->v_state = 1; /* /dev/tty0 is always opened */
164:         for(n = 1; n < NR_VCONSOLES + 1; n++) {
165:             tty = get_tty(MKDEV(VCONSOLES_MAJOR, n));
166:             if(tty->count) {
167:                 vt_stat->v_state |= (1 << n);
168:             }
169:         }
170:     }
171:     break;
172:
173:     case VT_RELDISP:
174:         if(vc->vt_mode.mode != VT_PROCESS) {
175:             return -EINVAL;
176:         }
177:         if(vc->switchto_tty < 0) {
178:             if(arg != VT_ACKACQ) {
179:                 return -EINVAL;
180:             }
181:         } else {
182:             if(arg) {
183:                 int switchto_tty;
184:                 switchto_tty = vc->switchto_tty;
185:                 vc->switchto_tty = -1;
186:                 vconsole_select_final(switchto_tty);

```

## drivers/char/vt.c

Page 4/4

```
187:                } else {
188:                    vc->switchto_tty = -1;
189:                }
190:            }
191:            break;
192:
193:            case VT_ACTIVATE:
194:                if(current_cons == MINOR(tty->dev) || IS_SUPERUSER) {
195:                    if(!arg || arg > NR_VCONSOLES) {
196:                        return -ENXIO;
197:                    }
198:                    vconsole_select(--arg);
199:                } else {
200:                    return -EPERM;
201:                }
202:                break;
203:
204:            case VT_WAITACTIVE:
205:                if(current_cons == MINOR(tty->dev)) {
206:                    break;
207:                }
208:                if(!arg || arg > NR_VCONSOLES) {
209:                    return -ENXIO;
210:                }
211:                printk("ACTIVATING another tty!! (cmd = 0x%x)\n", cmd);
212:                break;
213:
214:            default:
215:                return -EINVAL;
216:        }
217:        return 0;
218:    }
```

## fs/buffer.c

Page 1/7

```

1: /*
2:  * fiwix/fs/buffer.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: /*
9:  * buffer.c implements a cache with a free list as a doubly circular linked
10: * list and a chained hash table with doubly linked lists.
11: *
12: * hash table
13: * +-----+ +-----+ +-----+ +-----+
14: * | index | |prev|data|next| |prev|data|next| |prev|data|next|
15: * |   0   --> | / |   |   |   | <--- |   |   |   | <--- |   |   |   |
16: * +-----+ +-----+ +-----+ +-----+
17: * +-----+ +-----+ +-----+ +-----+
18: * | index | |prev|data|next| |prev|data|next| |prev|data|next|
19: * |   1   --> | / |   |   |   | <--- |   |   |   | <--- |   |   |   |
20: * +-----+ +-----+ +-----+ +-----+
21: *                               (buffer)       (buffer)       (buffer)
22: *      ...
23: */
24:
25: #include <fiwix/asm.h>
26: #include <fiwix/kernel.h>
27: #include <fiwix/sleep.h>
28: #include <fiwix/sched.h>
29: #include <fiwix/buffer.h>
30: #include <fiwix/devices.h>
31: #include <fiwix/fs.h>
32: #include <fiwix/mm.h>
33: #include <fiwix/errno.h>
34: #include <fiwix/stdio.h>
35: #include <fiwix/string.h>
36: #include <fiwix/stat.h>
37:
38: #define BUFFER_HASH(dev, block) (((__dev_t)(dev) ^ (__blk_t)(block)) % (NR_BUF_H
ASH))
39: #define NR_BUFFERS          buffer_table_size / sizeof(struct buffer)
40: #define NR_BUF_HASH        buffer_hash_table_size / sizeof(unsigned int)
41:
42: struct buffer *buffer_table;          /* buffer pool */
43: struct buffer *buffer_head;          /* buffer pool head */
44: struct buffer **buffer_hash_table;
45:
46: static struct resource sync_resource = { NULL, NULL };
47:
48: static void insert_to_hash(struct buffer *buf)
49: {
50:     struct buffer **h;
51:     int i;
52:
53:     i = BUFFER_HASH(buf->dev, buf->block);
54:     h = &buffer_hash_table[i];
55:
56:     if(!*h) {
57:         *h = buf;
58:         (*h)->prev_hash = (*h)->next_hash = NULL;
59:     } else {
60:         buf->prev_hash = NULL;
61:         buf->next_hash = *h;
62:         (*h)->prev_hash = buf;
63:         *h = buf;
64:     }
65: }
66:

```

## fs/buffer.c

Page 2/7

```

67: static void remove_from_hash(struct buffer *buf)
68: {
69:     struct buffer **h;
70:     int i;
71:
72:     i = BUFFER_HASH(buf->dev, buf->block);
73:     h = &buffer_hash_table[i];
74:
75:     while(*h) {
76:         if(*h == buf) {
77:             if((*h)->next_hash) {
78:                 (*h)->next_hash->prev_hash = (*h)->prev_hash;
79:             }
80:             if((*h)->prev_hash) {
81:                 (*h)->prev_hash->next_hash = (*h)->next_hash;
82:             }
83:             if(h == &buffer_hash_table[i]) {
84:                 *h = (*h)->next_hash;
85:             }
86:             break;
87:         }
88:         h = &(*h)->next_hash;
89:     }
90: }
91:
92: static void remove_from_free_list(struct buffer *buf)
93: {
94:     buf->prev_free->next_free = buf->next_free;
95:     buf->next_free->prev_free = buf->prev_free;
96:     if(buf == buffer_head) {
97:         buffer_head = buf->next_free;
98:     }
99: }
100:
101: static void buffer_wait(struct buffer *buf)
102: {
103:     unsigned long int flags;
104:
105:     for(;;) {
106:         SAVE_FLAGS(flags); CLI();
107:         if(buf->locked) {
108:             RESTORE_FLAGS(flags);
109:             sleep(&buffer_wait, PROC_UNINTERRUPTIBLE);
110:         } else {
111:             break;
112:         }
113:     }
114:     buf->locked = 1;
115:     RESTORE_FLAGS(flags);
116: }
117:
118: static struct buffer * get_free_buffer(void)
119: {
120:     unsigned long int flags;
121:     struct buffer *buf;
122:
123:     /* no more buffers on free list */
124:     if(buffer_head == buffer_head->next_free) {
125:         return NULL;
126:     }
127:
128:     for(;;) {
129:         SAVE_FLAGS(flags); CLI();
130:         buf = buffer_head;
131:         if(buf->locked) {
132:             RESTORE_FLAGS(flags);
133:             sleep(&buffer_wait, PROC_UNINTERRUPTIBLE);

```



## fs/buffer.c

Page 3/7

```

134:         } else {
135:             break;
136:         }
137:     }
138:
139:     buf = buffer_head;
140:     remove_from_free_list(buf);
141:     buf->locked = 1;
142:
143:     RESTORE_FLAGS(flags);
144:     return buf;
145: }
146:
147: static void sync_one_buffer(struct buffer *buf)
148: {
149:     struct device *d;
150:     int errno;
151:
152:     if(!(d = get_device(BLK_DEV, buf->dev))) {
153:         printk("WARNING: %s(): block device %d,%d not registered!\n", __
FUNCTION__, MAJOR(buf->dev), MINOR(buf->dev));
154:         return;
155:     }
156:
157:     if(d->fsop && d->fsop->write_block) {
158:         errno = d->fsop->write_block(buf->dev, buf->block, buf->data, bu
f->size);
159:         if(errno < 0) {
160:             if(errno == -EROFS) {
161:                 printk("WARNING: %s(): write protection on devic
e %d,%d.\n", __FUNCTION__, MAJOR(buf->dev), MINOR(buf->dev), buf->block);
162:             } else {
163:                 printk("WARNING: %s(): I/O error on device %d,%d
.\n", __FUNCTION__, MAJOR(buf->dev), MINOR(buf->dev), buf->block);
164:             }
165:             return;
166:         }
167:         buf->dirty = 0;
168:     } else {
169:         printk("WARNING: %s(): device %d,%d does not have the write_bloc
k() method!\n", __FUNCTION__, MAJOR(buf->dev), MINOR(buf->dev));
170:     }
171: }
172:
173: static struct buffer * search_buffer_hash(__dev_t dev, __blk_t block, int size)
174: {
175:     struct buffer *buf;
176:     int i;
177:
178:     i = BUFFER_HASH(dev, block);
179:     buf = buffer_hash_table[i];
180:
181:     while(buf) {
182:         if(buf->dev == dev && buf->block == block && buf->size == size)
{
183:             return buf;
184:         }
185:         buf = buf->next_hash;
186:     }
187:
188:     return NULL;
189: }
190:
191: static struct buffer * getblk(__dev_t dev, __blk_t block, int size)
192: {
193:     unsigned long int flags;
194:     struct buffer *buf;

```

## fs/buffer.c

Page 4/7

```

195:
196:     for(;;) {
197:         if((buf = search_buffer_hash(dev, block, size)) {
198:             SAVE_FLAGS(flags); CLI();
199:             if(buf->locked) {
200:                 RESTORE_FLAGS(flags);
201:                 sleep(&buffer_wait, PROC_UNINTERRUPTIBLE);
202:                 continue;
203:             }
204:             buf->locked = 1;
205:             remove_from_free_list(buf);
206:             RESTORE_FLAGS(flags);
207:             return buf;
208:         }
209:
210:         if(!(buf = get_free_buffer())) {
211:             printk("WARNING: %s(): no more buffers on free list!\n",
__FUNCTION__);
212:             sleep(&get_free_buffer, PROC_UNINTERRUPTIBLE);
213:             continue;
214:         }
215:
216:         if(buf->dirty) {
217:             sync_one_buffer(buf);
218:         } else {
219:             if(!buf->data) {
220:                 if(!(buf->data = (char *)kmalloc())) {
221:                     brelse(buf);
222:                     printk("%s(): returning NULL\n", __FUNCT
ION__);
223:                     return NULL;
224:                 }
225:                 kstat.buffers += (PAGE_SIZE / 1024);
226:             }
227:         }
228:
229:         SAVE_FLAGS(flags); CLI();
230:         remove_from_hash(buf); /* remove it from old hash */
231:         buf->dev = dev;
232:         buf->block = block;
233:         buf->size = size;
234:         insert_to_hash(buf);
235:         buf->valid = 0;
236:         RESTORE_FLAGS(flags);
237:         return buf;
238:     }
239: }
240:
241: struct buffer * get_dirty_buffer(__dev_t dev, __blk_t block, int size)
242: {
243:     unsigned long int flags;
244:     struct buffer *buf;
245:
246:     for(;;) {
247:         if((buf = search_buffer_hash(dev, block, size)) {
248:             if(buf->dirty) {
249:                 SAVE_FLAGS(flags); CLI();
250:                 if(buf->locked) {
251:                     RESTORE_FLAGS(flags);
252:                     sleep(&buffer_wait, PROC_UNINTERRUPTIBLE
);
253:                     continue;
254:                 }
255:                 buf->locked = 1;
256:                 remove_from_free_list(buf);
257:                 RESTORE_FLAGS(flags);
258:                 break;

```

## fs/buffer.c

Page 5/7

```

259:         }
260:     }
261:     buf = NULL;
262:     break;
263: }
264:
265:     return buf;
266: }
267:
268: struct buffer * bread(__dev_t dev, __blk_t block, int size)
269: {
270:     struct buffer *buf;
271:     struct device *d;
272:
273:     if(!(d = get_device(BLK_DEV, dev))) {
274:         printk("WARNING: %s(): device major %d not found!\n", __FUNCTION
__, MAJOR(dev));
275:         return NULL;
276:     }
277:
278:     if((buf = getblk(dev, block, size)) {
279:         if(!buf->valid) {
280:             if(d->fsop && d->fsop->read_block) {
281:                 if(d->fsop->read_block(dev, block, buf->data, si
ze) >= 0) {
282:                     buf->valid = 1;
283:                 }
284:             }
285:         }
286:         if(buf->valid) {
287:             return buf;
288:         }
289:         brelse(buf);
290:     }
291:
292:     printk("WARNING: %s(): returning NULL!\n", __FUNCTION__);
293:     return NULL;
294: }
295:
296: void bwrite(struct buffer *buf)
297: {
298:     buf->dirty = 1;
299:     buf->valid = 1;
300:     brelse(buf);
301: }
302:
303: void brelse(struct buffer *buf)
304: {
305:     unsigned long int flags;
306:
307:     SAVE_FLAGS(flags); CLI();
308:
309:     if(!buffer_head) {
310:         buf->prev_free = buf->next_free = buf;
311:         buffer_head = buf;
312:     } else {
313:         buf->next_free = buffer_head;
314:         buf->prev_free = buffer_head->prev_free;
315:         buffer_head->prev_free->next_free = buf;
316:         buffer_head->prev_free = buf;
317:
318:         /* if not valid place the buffer at the head of the free list */
319:         if(!buf->valid) {
320:             buffer_head = buf;
321:         }
322:     }
323:     buf->locked = 0;

```

## fs/buffer.c

Page 6/7

```

324:
325:     RESTORE_FLAGS(flags);
326:
327:     wakeup(&get_free_buffer);
328:     wakeup(&buffer_wait);
329: }
330:
331: void sync_buffers(__dev_t dev)
332: {
333:     struct buffer *buf;
334:     int n;
335:
336:     buf = &buffer_table[0];
337:
338:     lock_resource(&sync_resource);
339:     for(n = 0; n < NR_BUFFERS; n++) {
340:         if(buf->dirty) {
341:             if(!dev || buf->dev == dev) {
342:                 buffer_wait(buf);
343:                 sync_one_buffer(buf);
344:                 buf->locked = 0;
345:                 wakeup(&buffer_wait);
346:             }
347:         }
348:         buf++;
349:     }
350:     unlock_resource(&sync_resource);
351:     return;
352: }
353:
354: void invalidate_buffers(__dev_t dev)
355: {
356:     unsigned long int flags;
357:     unsigned int n;
358:     struct buffer *buf;
359:
360:     buf = &buffer_table[0];
361:     SAVE_FLAGS(flags); CLI();
362:
363:     for(n = 0; n < NR_BUFFERS; n++) {
364:         if(!buf->locked && buf->dev == dev) {
365:             buffer_wait(buf);
366:             remove_from_hash(buf);
367:             buf->valid = 0;
368:             buf->locked = 0;
369:             wakeup(&buffer_wait);
370:         }
371:         buf++;
372:     }
373:
374:     RESTORE_FLAGS(flags);
375:     /* FIXME: invalidate_pages(dev); */
376: }
377:
378: /*
379:  * When kernel runs out of pages, kswapd is awoken and it calls this function
380:  * which goes through the buffer free list, freeing up to NR_BUF_RECLAIM
381:  * buffers.
382:  */
383: int reclaim_buffers(void)
384: {
385:     struct buffer *buf, *first;
386:     int reclaimed;
387:
388:     reclaimed = 0;
389:     first = NULL;
390:

```

## fs/buffer.c

Page 7/7

```

391:         for(;;) {
392:             if(!(buf = get_free_buffer())) {
393:                 printk("WARNING: %s(): no more buffers on free list!\n",
__FUNCTION__);
394:                 sleep(&get_free_buffer, PROC_UNINTERRUPTIBLE);
395:                 continue;
396:             }
397:
398:             remove_from_hash(buf);
399:             if(buf->dirty) {
400:                 sync_one_buffer(buf);
401:             }
402:
403:             /* this ensures the buffer will go to the tail */
404:             buf->valid = 1;
405:
406:             if(first) {
407:                 if(first == buf) {
408:                     brelse(buf);
409:                     break;
410:                 }
411:             } else {
412:                 first = buf;
413:             }
414:             if(buf->data) {
415:                 kfree((unsigned int)buf->data);
416:                 buf->data = NULL;
417:                 kstat.buffers -= (PAGE_SIZE / 1024);
418:                 reclaimed++;
419:                 if(reclaimed == NR_BUF_RECLAIM) {
420:                     brelse(buf);
421:                     break;
422:                 }
423:             }
424:             brelse(buf);
425:             do_sched();
426:         }
427:
428:         wakeup(&buffer_wait);
429:         return reclaimed;
430: }
431:
432: void buffer_init(void)
433: {
434:     struct buffer *buf;
435:     unsigned int n;
436:
437:     memset_b(buffer_table, NULL, buffer_table_size);
438:     memset_b(buffer_hash_table, NULL, buffer_hash_table_size);
439:     for(n = 0; n < NR_BUFFERS; n++) {
440:         buf = &buffer_table[n];
441:         brelse(buf);
442:     }
443: }

```

## fs/devices.c

Page 1/6

```

1: /*
2:  * fiwix/fs/devices.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/devices.h>
12: #include <fiwix/fs.h>
13: #include <fiwix/mm.h>
14: #include <fiwix/process.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: struct device *chr_device_table[NR_CHRDEV];
19: struct device *blk_device_table[NR_BLKDEV];
20:
21: struct fs_operations def_chr_fsop = {
22:     0,
23:     0,
24:
25:     chr_dev_open,
26:     NULL, /* close */
27:     NULL, /* read */
28:     NULL, /* write */
29:     NULL, /* ioctl */
30:     NULL, /* lseek */
31:     NULL, /* readdir */
32:     NULL, /* mmap */
33:     NULL, /* select */
34:
35:     NULL, /* readlink */
36:     NULL, /* followlink */
37:     NULL, /* bmap */
38:     NULL, /* lockup */
39:     NULL, /* rmdir */
40:     NULL, /* link */
41:     NULL, /* unlink */
42:     NULL, /* symlink */
43:     NULL, /* mkdir */
44:     NULL, /* mknod */
45:     NULL, /* truncate */
46:     NULL, /* create */
47:     NULL, /* rename */
48:
49:     NULL, /* read_block */
50:     NULL, /* write_block */
51:
52:     NULL, /* read_inode */
53:     NULL, /* write_inode */
54:     NULL, /* ialloc */
55:     NULL, /* ifree */
56:     NULL, /* stats */
57:     NULL, /* read_superblock */
58:     NULL, /* remount_fs */
59:     NULL, /* write_superblock */
60:     NULL, /* release_superblock */
61: };
62:
63: struct fs_operations def_blk_fsop = {
64:     0,
65:     0,
66:
67:     blk_dev_open,

```

## fs/devices.c

Page 2/6

```

68:         blk_dev_close,
69:         blk_dev_read,
70:         blk_dev_write,
71:         blk_dev_ioctl,
72:         blk_dev_lseek,
73:         NULL,                /* readdir */
74:         NULL,                /* mmap */
75:         NULL,                /* select */
76:
77:         NULL,                /* readlink */
78:         NULL,                /* followlink */
79:         NULL,                /* bmap */
80:         NULL,                /* lockup */
81:         NULL,                /* rmdir */
82:         NULL,                /* link */
83:         NULL,                /* unlink */
84:         NULL,                /* symlink */
85:         NULL,                /* mkdir */
86:         NULL,                /* mknod */
87:         NULL,                /* truncate */
88:         NULL,                /* create */
89:         NULL,                /* rename */
90:
91:         NULL,                /* read_block */
92:         NULL,                /* write_block */
93:
94:         NULL,                /* read_inode */
95:         NULL,                /* write_inode */
96:         NULL,                /* ialloc */
97:         NULL,                /* ifree */
98:         NULL,                /* stats */
99:         NULL,                /* read_superblock */
100:        NULL,                /* remount_fs */
101:        NULL,                /* write_superblock */
102:        NULL,                /* release_superblock */
103: };
104:
105: int register_device(int type, struct device *new_d)
106: {
107:     struct device **d;
108:     int n, minors;
109:
110:     switch(type) {
111:         case CHR_DEV:
112:             if(new_d->major >= NR_CHRDEV) {
113:                 printk("%s(): character device major %d is greater than NR_CHRDEV (%d).\n", __FUNCTION__, new_d->major, NR_CHRDEV);
114:                 return 1;
115:             }
116:             d = &chr_device_table[new_d->major];
117:             break;
118:         case BLK_DEV:
119:             if(new_d->major >= NR_BLKDEV) {
120:                 printk("%s(): block device major %d is greater than NR_BLKDEV (%d).\n", __FUNCTION__, new_d->major, NR_BLKDEV);
121:                 return 1;
122:             }
123:             d = &blk_device_table[new_d->major];
124:             break;
125:         default:
126:             printk("WARNING: %s(): invalid device type %d.\n", __FUNCTION__, type);
127:             return 1;
128:             break;
129:     }
130:
131:     /* make sure there are minors defined */

```

## fs/devices.c

Page 3/6

```

132:         for(n = 0, minors = 0; n < 8; n++) {
133:             minors += new_d->minors[n];
134:         }
135:         if(!minors) {
136:             printk("WARNING: %s(): device major %d with no defined minors.\n",
", __FUNCTION__, new_d->major);
137:             return 1;
138:         }
139:
140:         if(*d) {
141:             if(&(*d)->minors == &new_d->minors || (&(*d)->next && &(*d)->next->minors == &new_d->minors)) {
142:                 printk("WARNING: %s(): duplicated device major %d.\n",
__FUNCTION__, new_d->major);
143:                 return 1;
144:             }
145:             do {
146:                 d = &(*d)->next;
147:             } while(*d);
148:         }
149:         *d = new_d;
150:
151:         return 0;
152:     }
153:
154:     struct device * get_device(int type, __dev_t dev)
155:     {
156:         char *name;
157:         unsigned char major;
158:         struct device *d;
159:
160:         major = MAJOR(dev);
161:
162:         switch(type) {
163:             case CHR_DEV:
164:                 if(major >= NR_CHRDEV) {
165:                     printk("%s(): character device major %d is great
er than NR_CHRDEV (%d).\n", __FUNCTION__, major, NR_CHRDEV);
166:                     return NULL;
167:                 }
168:                 d = chr_device_table[major];
169:                 name = "character";
170:                 break;
171:             case BLK_DEV:
172:                 if(major >= NR_BLKDEV) {
173:                     printk("%s(): block device major %d is greater t
han NR_BLKDEV (%d).\n", __FUNCTION__, major, NR_BLKDEV);
174:                     return NULL;
175:                 }
176:                 d = blk_device_table[major];
177:                 name = "block";
178:                 break;
179:             default:
180:                 printk("WARNING: %s(): invalid device type %d.\n", __FUN
CTION__, type);
181:                 return NULL;
182:         }
183:
184:         while(d) {
185:             if(d->major == major) {
186:                 if(TEST_MINOR(d->minors, MINOR(dev))) {
187:                     return d;
188:                 }
189:                 d = d->next;
190:                 continue;
191:             }
192:             break;

```



## fs/devices.c

Page 4/6

```

193:         }
194:
195:         printk("WARNING: %s(): %s device %d,%d not found.\n", __FUNCTION__, name
, major, MINOR(dev));
196:         return NULL;
197:     }
198:
199: int chr_dev_open(struct inode *i, struct fd *fd_table)
200: {
201:     struct device *d;
202:
203:     if((d = get_device(CHR_DEV, i->rdev)) {
204:         i->fsop = d->fsop;
205:         if(i->fsop && i->fsop->open) {
206:             return i->fsop->open(i, fd_table);
207:         }
208:     }
209:
210:     return -EINVAL;
211: }
212:
213: int blk_dev_open(struct inode *i, struct fd *fd_table)
214: {
215:     struct device *d;
216:
217:     if((d = get_device(BLK_DEV, i->rdev)) {
218:         if(d->fsop && d->fsop->open) {
219:             return d->fsop->open(i, fd_table);
220:         }
221:     }
222:
223:     return -EINVAL;
224: }
225:
226: int blk_dev_close(struct inode *i, struct fd *fd_table)
227: {
228:     struct device *d;
229:
230:     if((d = get_device(BLK_DEV, i->rdev)) {
231:         if(d->fsop && d->fsop->close) {
232:             return d->fsop->close(i, fd_table);
233:         }
234:     }
235:
236:     printk("WARNING: %s(): block device %d,%d does not have the close() meth
od.\n", __FUNCTION__, MAJOR(i->rdev), MINOR(i->rdev));
237:     return -EINVAL;
238: }
239:
240: int blk_dev_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t co
unt)
241: {
242:     __blk_t block;
243:     __off_t total_read;
244:     unsigned long long int device_size;
245:     int blksize;
246:     unsigned int boffset, bytes;
247:     struct buffer *buf;
248:     struct device *d;
249:
250:     if(!(d = get_device(BLK_DEV, i->rdev)) {
251:         return -EINVAL;
252:     }
253:
254:     blksize = d->blksize ? d->blksize : BLKSIZE_1K;
255:     total_read = 0;
256:     if(!d->device_data) {

```

## fs/devices.c

Page 5/6

```

257:             printk("%s(): don't know the size of the block device %d,%d.\n",
__FUNCTION__, MAJOR(i->rdev), MINOR(i->rdev));
258:             return -EIO;
259:         }
260:
261:         device_size = ((unsigned int *)d->device_data)[MINOR(i->rdev)];
262:         device_size *= 1024LLU;
263:
264:         count = (fd_table->offset + count > device_size) ? device_size - fd_table->offset : count;
265:         if(!count || fd_table->offset > device_size) {
266:             return 0;
267:         }
268:         while(count) {
269:             boffset = fd_table->offset % blksize;
270:             block = (fd_table->offset / blksize);
271:             if(!(buf = bread(i->rdev, block, blksize))) {
272:                 return -EIO;
273:             }
274:             bytes = blksize - boffset;
275:             bytes = MIN(bytes, count);
276:             memcpy_b(buffer + total_read, buf->data + boffset, bytes);
277:             total_read += bytes;
278:             count -= bytes;
279:             boffset += bytes;
280:             boffset %= blksize;
281:             fd_table->offset += bytes;
282:             brelse(buf);
283:         }
284:         return total_read;
285:     }
286:
287: int blk_dev_write(struct inode *i, struct fd *fd_table, const char *buffer, __size_t count)
288: {
289:     __blk_t block;
290:     __off_t total_written;
291:     unsigned long long int device_size;
292:     int blksize;
293:     unsigned int boffset, bytes;
294:     struct buffer *buf;
295:     struct device *d;
296:
297:     if(!(d = get_device(BLK_DEV, i->rdev))) {
298:         return -EINVAL;
299:     }
300:
301:     blksize = d->blksize ? d->blksize : BLKSIZE_1K;
302:     total_written = 0;
303:     if(!d->device_data) {
304:         printk("%s(): don't know the size of the block device %d,%d.\n",
__FUNCTION__, MAJOR(i->rdev), MINOR(i->rdev));
305:         return -EIO;
306:     }
307:
308:     device_size = ((unsigned int *)d->device_data)[MINOR(i->rdev)];
309:     device_size *= 1024LLU;
310:
311:     count = (fd_table->offset + count > device_size) ? device_size - fd_table->offset : count;
312:     if(!count || fd_table->offset > device_size) {
313:         printk("%s(): I/O error on device %d,%d, offset %u.\n", __FUNCTION__, MAJOR(i->rdev), MINOR(i->rdev), fd_table->offset);
314:         return -EIO;
315:     }
316:     while(count) {
317:         boffset = fd_table->offset % blksize;

```

## fs/devices.c

Page 6/6

```
318:         block = (fd_table->offset / blksize);
319:         if(!(buf = bread(i->rdev, block, blksize))) {
320:             return -EIO;
321:         }
322:         bytes = blksize - boffset;
323:         bytes = MIN(bytes, count);
324:         memcpy_b(buf->data + boffset, buffer + total_written, bytes);
325:         total_written += bytes;
326:         count -= bytes;
327:         boffset += bytes;
328:         boffset %= blksize;
329:         fd_table->offset += bytes;
330:         bwrite(buf);
331:     }
332:     return total_written;
333: }
334:
335: int blk_dev_ioctl(struct inode *i, int cmd, unsigned long int arg)
336: {
337:     struct device *d;
338:
339:     if((d = get_device(BLK_DEV, i->rdev)) {
340:         if(d->fsop && d->fsop->ioctl) {
341:             return d->fsop->ioctl(i, cmd, arg);
342:         }
343:     }
344:
345:     printk("WARNING: %s(): block device %d,%d does not have the ioctl() meth
od.\n", __FUNCTION__, MAJOR(i->rdev), MINOR(i->rdev));
346:     return -EINVAL;
347: }
348:
349: int blk_dev_lseek(struct inode *i, __off_t offset)
350: {
351:     struct device *d;
352:
353:     if((d = get_device(BLK_DEV, i->rdev)) {
354:         if(d->fsop && d->fsop->lseek) {
355:             return d->fsop->lseek(i, offset);
356:         }
357:     }
358:
359:     return offset;
360: }
361:
362: void dev_init(void)
363: {
364:     memset_b(chr_device_table, NULL, sizeof(chr_device_table));
365:     memset_b(blk_device_table, NULL, sizeof(blk_device_table));
366: }
```

fs/elf.c

```

1: /*
2:  * fiwix/fs/elf.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/asm.h>
10: #include <fiwix/types.h>
11: #include <fiwix/buffer.h>
12: #include <fiwix/fs.h>
13: #include <fiwix/i386elf.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/mman.h>
16: #include <fiwix/fs.h>
17: #include <fiwix/fcntl.h>
18: #include <fiwix/process.h>
19: #include <fiwix/errno.h>
20: #include <fiwix/stdio.h>
21: #include <fiwix/string.h>
22:
23: #define AT_ITEMS          12          /* ELF Auxiliary Vectors */
24:
25: static int check_elf(struct elf32_hdr *elf32_h)
26: {
27:     if(elf32_h->e_ident[EI_MAG0] != ELFMAG0 ||
28:        elf32_h->e_ident[EI_MAG1] != ELFMAG1 ||
29:        elf32_h->e_ident[EI_MAG2] != ELFMAG2 ||
30:        elf32_h->e_ident[EI_MAG3] != ELFMAG3 ||
31:        (elf32_h->e_type != ET_EXEC && elf32_h->e_type != ET_DYN) ||
32:        elf32_h->e_machine != EM_386) {
33:         return -EINVAL;
34:     }
35:     return 0;
36: }
37:
38: /*
39:  * Setup the initial process stack (System V ABI for i386)
40:  * -----
41:  * 0xBFFFFFFF
42:  *   +-----+
43:  *   | envp[] str |
44:  *   +-----+
45:  *   | argv[] str |
46:  *   +-----+
47:  *   | NULL      |
48:  *   +-----+
49:  *   | ELF Aux.Vect. |
50:  *   +-----+
51:  *   | NULL      |
52:  *   +-----+
53:  *   | envp[] ptr |
54:  *   +-----+
55:  *   | NULL      |
56:  *   +-----+
57:  *   | argv[] ptr |
58:  *   +-----+
59:  *   | argc      |
60:  *   +-----+
61:  *   | stack pointer | grows toward lower addresses
62:  *   +-----+ ||
63:  *   | ..... | \
64:  *   | ..... |
65:  *   | ..... |
66:  *   | ..... | /\
67:  *   +-----+ ||

```

elf\_create\_stack() setups this section

## fs/elf.c

Page 2/10

```

68: *      | brk (heap) | grows toward higher addresses
69: *      +-----+
70: *      | .bss section |
71: *      +-----+
72: *      | .data section |
73: *      +-----+
74: *      | .text section |
75: *      +-----+
76: * 0x08048000
77: */
78: static void elf_create_stack(struct binargs *barg, unsigned int *sp, unsigned in
t str_ptr, int at_base, struct elf32_hdr *elf32_h, unsigned int phdr_addr)
79: {
80:     unsigned int n, addr;
81:     char *str;
82:
83:     /* copy strings */
84:     for(n = 0; n < ARG_MAX; n++) {
85:         if(barg->page[n]) {
86:             addr = KERNEL_BASE_ADDR - ((ARG_MAX - n) * PAGE_SIZE);
87:             memcpy_b((void *)addr, (void *)barg->page[n], PAGE_SIZE)
;
88:         }
89:     }
90:
91: #ifdef __DEBUG__
92:     printk("sp = 0x%08x\n", sp);
93: #endif /*__DEBUG__ */
94:
95:     /* copy the value of 'argc' into the stack */
96:     memcpy_l((void *)sp, &barg->argc, 1);
97: #ifdef __DEBUG__
98:     printk("at 0x%08x -> argc\n", sp);
99: #endif /*__DEBUG__ */
100:     sp++;
101:
102:     /* copy as many pointers to strings as 'argc' */
103:     current->argv = (char **)sp;
104:     for(n = 0; n < barg->argc; n++) {
105:         memcpy_l((void *)sp, &str_ptr, 1);
106:         str = (char *)str_ptr;
107: #ifdef __DEBUG__
108:         printk("at 0x%08x -> str_ptr(%d) = 0x%08x (+ %d)\n", sp, n, str_
ptr, strlen(str) + 1);
109: #endif /*__DEBUG__ */
110:         sp++;
111:         str_ptr += strlen(str) + 1;
112:     }
113:
114:     /* the last element of 'argv[]' must be a NULL-pointer */
115:     memset_l((void *)sp, NULL, 1);
116: #ifdef __DEBUG__
117:     printk("at 0x%08x -> ----- = 0x%08x\n", sp, 0);
118: #endif /*__DEBUG__ */
119:     sp++;
120:
121:     /* copy as many pointers to strings as 'envc' */
122:     current->envp = (char **)sp;
123:     for(n = 0; n < barg->envc; n++) {
124:         memcpy_l((void *)sp, &str_ptr, 1);
125:         str = (char *)str_ptr;
126: #ifdef __DEBUG__
127:         printk("at 0x%08x -> str_ptr(%d) = 0x%08x (+ %d)\n", sp, n, str_
ptr, strlen(str) + 1);
128: #endif /*__DEBUG__ */
129:         sp++;
130:         str_ptr += strlen(str) + 1;

```

## fs/elf.c

Page 3/10

```

131:         }
132:
133:         /* the last element of 'envp[]' must be a NULL-pointer */
134:         memset_l((void *)sp, NULL, 1);
135: #ifdef __DEBUG__
136:         printk("at 0x%08x -> ----- = 0x%08x\n", sp, 0);
137: #endif /*__DEBUG__ */
138:         sp++;
139:
140:
141:         /* copy the Auxiliar Table Items (dlinfo_items) */
142:         if(at_base) {
143:             memset_l((void *)sp, AT_PHDR, 1);
144: #ifdef __DEBUG__
145:             printk("at 0x%08x -> AT_PHDR = %d", sp, *sp);
146: #endif /*__DEBUG__ */
147:             sp++;
148:
149:             memcpy_l((void *)sp, &phdr_addr, 1);
150: #ifdef __DEBUG__
151:             printk("\t\tAT_PHDR = 0x%08x\n", *sp);
152: #endif /*__DEBUG__ */
153:             sp++;
154:
155:             memset_l((void *)sp, AT_PHEMT, 1);
156: #ifdef __DEBUG__
157:             printk("at 0x%08x -> AT_PHEMT = %d", sp, *sp);
158: #endif /*__DEBUG__ */
159:             sp++;
160:
161:             memset_l((void *)sp, sizeof(struct elf32_phdr), 1);
162: #ifdef __DEBUG__
163:             printk("\t\tAT_PHEMT = %d\n", *sp);
164: #endif /*__DEBUG__ */
165:             sp++;
166:
167:             memset_l((void *)sp, AT_PHNUM, 1);
168: #ifdef __DEBUG__
169:             printk("at 0x%08x -> AT_PHNUM = %d", sp, *sp);
170: #endif /*__DEBUG__ */
171:             sp++;
172:
173:             memset_l((void *)sp, 0, 1);
174:             memcpy_w((void *)sp, &elf32_h->e_phnum, 1);
175: #ifdef __DEBUG__
176:             printk("\t\tAT_PHNUM = %d\n", *sp);
177: #endif /*__DEBUG__ */
178:             sp++;
179:
180:             memset_l((void *)sp, AT_PAGESZ, 1);
181: #ifdef __DEBUG__
182:             printk("at 0x%08x -> AT_PGSIZE = %d", sp, *sp);
183: #endif /*__DEBUG__ */
184:             sp++;
185:
186:             memset_l((void *)sp, PAGE_SIZE, 1);
187: #ifdef __DEBUG__
188:             printk("\t\tAT_PGSIZE = %d\n", *sp);
189: #endif /*__DEBUG__ */
190:             sp++;
191:
192:             memset_l((void *)sp, AT_BASE, 1);
193: #ifdef __DEBUG__
194:             printk("at 0x%08x -> AT_BASE = %d", sp, *sp);
195: #endif /*__DEBUG__ */
196:             sp++;
197:

```

## fs/elf.c

Page 4/10

```

198:             memset_l((void *)sp, at_base, 1);
199: #ifdef __DEBUG__
200:             printk("\t\tAT_BASE = 0x%08x\n", sp);
201: #endif /*__DEBUG__*/
202:             sp++;
203:
204:             memset_l((void *)sp, AT_FLAGS, 1);
205: #ifdef __DEBUG__
206:             printk("at 0x%08x -> AT_FLAGS = %d", sp, *sp);
207: #endif /*__DEBUG__*/
208:             sp++;
209:
210:             memset_l((void *)sp, NULL, 1);
211: #ifdef __DEBUG__
212:             printk("\t\tAT_FLAGS = %d\n", *sp);
213: #endif /*__DEBUG__*/
214:             sp++;
215:
216:             memset_l((void *)sp, AT_ENTRY, 1);
217: #ifdef __DEBUG__
218:             printk("at 0x%08x -> AT_ENTRY = %d ", sp, *sp);
219: #endif /*__DEBUG__*/
220:             sp++;
221:
222:             memcpy_l((void *)sp, &elf32_h->e_entry, 1);
223: #ifdef __DEBUG__
224:             printk("\t\tAT_ENTRY = 0x%08x\n", *sp);
225: #endif /*__DEBUG__*/
226:             sp++;
227:
228:             memset_l((void *)sp, AT_UID, 1);
229: #ifdef __DEBUG__
230:             printk("at 0x%08x -> AT_UID = %d", sp, *sp);
231: #endif /*__DEBUG__*/
232:             sp++;
233:
234:             memcpy_l((void *)sp, &current->uid, 1);
235: #ifdef __DEBUG__
236:             printk("\t\tAT_UID = %d\n", *sp);
237: #endif /*__DEBUG__*/
238:             sp++;
239:
240:             memset_l((void *)sp, AT_EUID, 1);
241: #ifdef __DEBUG__
242:             printk("at 0x%08x -> AT_EUID = %d", sp, *sp);
243: #endif /*__DEBUG__*/
244:             sp++;
245:
246:             memcpy_l((void *)sp, &current->euid, 1);
247: #ifdef __DEBUG__
248:             printk("\t\tAT_EUID = %d\n", *sp);
249: #endif /*__DEBUG__*/
250:             sp++;
251:
252:             memset_l((void *)sp, AT_GID, 1);
253: #ifdef __DEBUG__
254:             printk("at 0x%08x -> AT_GID = %d", sp, *sp);
255: #endif /*__DEBUG__*/
256:             sp++;
257:
258:             memcpy_l((void *)sp, &current->gid, 1);
259: #ifdef __DEBUG__
260:             printk("\t\tAT_GID = %d\n", *sp);
261: #endif /*__DEBUG__*/
262:             sp++;
263:
264:             memset_l((void *)sp, AT_EGID, 1);

```

## fs/elf.c

Page 5/10

```

265: #ifdef __DEBUG__
266:     printk("at 0x%08x -> AT_EGID = %d", sp, *sp);
267: #endif /*__DEBUG__ */
268:     sp++;
269:
270:     memcpy_l((void *)sp, &current->egid, 1);
271: #ifdef __DEBUG__
272:     printk("\t\tAT_EGID = %d\n", *sp);
273: #endif /*__DEBUG__ */
274:     sp++;
275: }
276:
277:     memset_l((void *)sp, AT_NULL, 1);
278: #ifdef __DEBUG__
279:     printk("at 0x%08x -> AT_NULL = %d", sp, *sp);
280: #endif /*__DEBUG__ */
281:     sp++;
282:
283:     memset_l((void *)sp, NULL, 1);
284: #ifdef __DEBUG__
285:     printk("\t\tAT_NULL = %d\n", *sp);
286: #endif /*__DEBUG__ */
287:     sp++;
288:
289: #ifdef __DEBUG__
290:     for(n = 0; n < barg->argc; n++) {
291:         printk("at 0x%08x -> argv[%d] = '%s'\n", current->argv[n], n, cu
rrent->argv[n]);
292:     }
293:     for(n = 0; n < barg->envc; n++) {
294:         printk("at 0x%08x -> envp[%d] = '%s'\n", current->envp[n], n, cu
rrent->envp[n]);
295:     }
296: #endif /*__DEBUG__ */
297: }
298:
299: static int elf_load_interpreter(struct inode *ii)
300: {
301:     int n, errno;
302:     struct buffer *buf;
303:     struct elf32_hdr *elf32_h;
304:     struct elf32_phdr *elf32_ph, *last_ptload;
305:     __blk_t block;
306:     unsigned int start, end, length;
307:     unsigned int prot;
308:     char *data;
309:     char type;
310:
311:     if((block = bmap(ii, 0, FOR_READING)) < 0) {
312:         return block;
313:     }
314:     if(!(buf = bread(ii->dev, block, ii->sb->s_blocksize))) {
315:         return -EIO;
316:     }
317:
318:     /*
319:      * The contents of the buffer is copied and then freed immediately to
320:      * make sure that it won't conflict while zeroing the BSS fractional
321:      * page, in case that the same block is requested during the page fault.
322:      */
323:     if(!(data = (void *)kmalloc())) {
324:         brelse(buf);
325:         return -ENOMEM;
326:     }
327:     memcpy_b(data, buf->data, ii->sb->s_blocksize);
328:     brelse(buf);
329:

```



```

330:     elf32_h = (struct elf32_hdr *)data;
331:     if(check_elf(elf32_h)) {
332:         kfree((unsigned int)data);
333:         return -ELIBBAD;
334:     }
335:
336:     last_ptload = NULL;
337:     for(n = 0; n < elf32_h->e_phnum; n++) {
338:         elf32_ph = (struct elf32_phdr *) (data + elf32_h->e_phoff + (size
of(struct elf32_phdr) * n));
339:         if(elf32_ph->p_type == PT_LOAD) {
340: #ifdef __DEBUG__
341:             printk("p_offset = 0x%08x\n", elf32_ph->p_offset);
342:             printk("p_vaddr = 0x%08x\n", elf32_ph->p_vaddr);
343:             printk("p_paddr = 0x%08x\n", elf32_ph->p_paddr);
344:             printk("p_filesz = 0x%08x\n", elf32_ph->p_filesz);
345:             printk("p_memsz = 0x%08x\n\n", elf32_ph->p_memsz);
346: #endif /*__DEBUG__ */
347:             start = (elf32_ph->p_vaddr & PAGE_MASK) + MMAP_START;
348:             length = (elf32_ph->p_vaddr & ~PAGE_MASK) + elf32_ph->p_
filesz;
349:             type = P_DATA;
350:             prot = 0;
351:             if(elf32_ph->p_flags & PF_R) {
352:                 prot = PROT_READ;
353:             }
354:             if(elf32_ph->p_flags & PF_W) {
355:                 prot |= PROT_WRITE;
356:             }
357:             if(elf32_ph->p_flags & PF_X) {
358:                 prot |= PROT_EXEC;
359:                 type = P_TEXT;
360:             }
361:             errno = do_mmap(ii, start, length, prot, MAP_PRIVATE | M
AP_FIXED, elf32_ph->p_offset & PAGE_MASK, type, O_RDONLY);
362:             if(errno < 0 && errno > -PAGE_SIZE) {
363:                 kfree((unsigned int)data);
364:                 send_sig(current, SIGSEGV);
365:                 return -ENOEXEC;
366:             }
367:             last_ptload = elf32_ph;
368:         }
369:     }
370:
371:     if(!last_ptload) {
372:         printk("WARNING: 'last_ptload' is NULL!\n");
373:     }
374:     elf32_ph = last_ptload;
375:
376:     /* zero-fill the fractional page of the DATA section */
377:     end = PAGE_ALIGN(elf32_ph->p_vaddr + elf32_ph->p_filesz) + MMAP_START;
378:     start = (elf32_ph->p_vaddr + elf32_ph->p_filesz) + MMAP_START;
379:     length = end - start;
380:
381:     /* this will generate a page fault which will load the page in */
382:     memset_b((void *)start, NULL, length);
383:
384:     /* setup the BSS section */
385:     start = (elf32_ph->p_vaddr + elf32_ph->p_filesz) + MMAP_START;
386:     start = PAGE_ALIGN(start);
387:     end = (elf32_ph->p_vaddr + elf32_ph->p_memsz) + MMAP_START;
388:     end = PAGE_ALIGN(end);
389:     length = end - start;
390:     errno = do_mmap(NULL, start, length, PROT_READ | PROT_WRITE, MAP_PRIVATE
| MAP_FIXED, 0, P_BSS, 0);
391:     if(errno < 0 && errno > -PAGE_SIZE) {
392:         kfree((unsigned int)data);

```

## fs/elf.c

Page 7/10

```

393:         send_sig(current, SIGSEGV);
394:         return -ENOEXEC;
395:     }
396:     kfree((unsigned int)data);
397:     return elf32_h->e_entry + MMAP_START;
398: }
399:
400: int elf_load(struct inode *i, struct binargs *barg, struct sigcontext *sc, char
*data)
401: {
402:     int n, errno;
403:     struct elf32_hdr *elf32_h;
404:     struct elf32_phdr *elf32_ph, *last_ptload;
405:     struct inode *ii;
406:     unsigned int start, end, length;
407:     unsigned int prot;
408:     char *interpreter;
409:     int at_base, phdr_addr;
410:     char type;
411:     unsigned int ae_ptr_len, ae_str_len;
412:     unsigned int sp, str;
413:
414:     elf32_h = (struct elf32_hdr *)data;
415:     if(check_elf(elf32_h)) {
416:         if(current->pid == INIT) {
417:             PANIC("%s has an unrecognized binary format.\n", INIT_PR
OGRAM);
418:         }
419:         return -ENOEXEC;
420:     }
421:
422:     /* check if an interpreter is required */
423:     interpreter = NULL;
424:     ii = NULL;
425:     phdr_addr = at_base = 0;
426:     for(n = 0; n < elf32_h->e_phnum; n++) {
427:         elf32_ph = (struct elf32_phdr *) (data + elf32_h->e_phoff + (size
of(struct elf32_phdr) * n));
428:         if(elf32_ph->p_type == PT_INTERP) {
429:             at_base = MMAP_START;
430:             interpreter = data + elf32_ph->p_offset;
431:             if(namei(interpreter, &ii, NULL, FOLLOW_LINKS)) {
432:                 printk("%s(): can't find interpreter '%s'.\n", _
_FUNCTION__, interpreter);
433:                 send_sig(current, SIGSEGV);
434:                 return -ELIBACC;
435:             }
436: #ifdef __DEBUG__
437:             printk("p_offset = 0x%08x\n", elf32_ph->p_offset);
438:             printk("p_vaddr = 0x%08x\n", elf32_ph->p_vaddr);
439:             printk("p_paddr = 0x%08x\n", elf32_ph->p_paddr);
440:             printk("p_filesz = 0x%08x\n", elf32_ph->p_filesz);
441:             printk("p_memsz = 0x%08x\n", elf32_ph->p_memsz);
442:             printk("using interpreter '%s'\n", interpreter);
443: #endif /*__DEBUG__ */
444:         }
445:     }
446:
447:     /*
448:     * calculate the final size of 'ae_ptr_len' based on:
449:     * - argc = 4 bytes (unsigned int)
450:     * - barg.argc = (num. of pointers to strings + 1 NULL) x 4 bytes (unsi
gned int)
451:     * - barg.envc = (num. of pointers to strings + 1 NULL) x 4 bytes (unsi
gned int)
452:     */
453:     ae_ptr_len = (1 + (barg->argc + 1) + (barg->envc + 1)) * sizeof(unsigned

```

## fs/elf.c

Page 8/10

```

int);
454:         ae_str_len = barg->argv_len + barg->envp_len;
455:         if(ae_ptr_len + ae_str_len > (ARG_MAX * PAGE_SIZE)) {
456:             printk("WARNING: %s(): argument list (%d) exceeds ARG_MAX (%d)!\n"
n", __FUNCTION__, ae_ptr_len + ae_str_len, ARG_MAX * PAGE_SIZE);
457:             return -E2BIG;
458:         }
459:
460: #ifdef __DEBUG__
461:         printk("argc=%d (argv_len=%d) envc=%d (envp_len=%d) ae_ptr_len=%d ae_str
r_len=%d\n", barg->argc, barg->argv_len, barg->envc, barg->envp_len, ae_ptr_len, ae_str
_len);
462: #endif /*__DEBUG__ */
463:
464:
465:         /* point of no return */
466:
467:         release_binary();
468:         current->rss = 0;
469:
470:         current->entry_address = elf32_h->e_entry;
471:         if(interpreter) {
472:             errno = elf_load_interpreter(ii);
473:             if(errno < 0) {
474:                 printk("%s(): unable to load the interpreter '%s'.\n", _
__FUNCTION__, interpreter);
475:                 iput(ii);
476:                 send_sig(current, SIGKILL);
477:                 return errno;
478:             }
479:             current->entry_address = errno;
480:             iput(ii);
481:         }
482:
483:         elf32_ph = last_ptload = NULL;
484:         for(n = 0; n < elf32_h->e_phnum; n++) {
485:             elf32_ph = (struct elf32_phdr *) (data + elf32_h->e_phoff + (size
of(struct elf32_phdr) * n));
486:             if(elf32_ph->p_type == PT_PHDR) {
487:                 phdr_addr = elf32_ph->p_vaddr;
488:             }
489:             if(elf32_ph->p_type == PT_LOAD) {
490:                 start = elf32_ph->p_vaddr & PAGE_MASK;
491:                 length = (elf32_ph->p_vaddr & ~PAGE_MASK) + elf32_ph->p_
filesz;
492:                 type = P_DATA;
493:                 prot = 0;
494:                 if(elf32_ph->p_flags & PF_R) {
495:                     prot = PROT_READ;
496:                 }
497:                 if(elf32_ph->p_flags & PF_W) {
498:                     prot |= PROT_WRITE;
499:                 }
500:                 if(elf32_ph->p_flags & PF_X) {
501:                     prot |= PROT_EXEC;
502:                     type = P_TEXT;
503:                 }
504:                 errno = do_mmap(i, start, length, prot, MAP_PRIVATE | MA
P_FIXED, elf32_ph->p_offset & PAGE_MASK, type, O_RDONLY);
505:                 if(errno < 0 && errno > -PAGE_SIZE) {
506:                     send_sig(current, SIGSEGV);
507:                     return -ENOEXEC;
508:                 }
509:                 last_ptload = elf32_ph;
510:             }
511:         }
512:

```

```

513:         elf32_ph = last_ptload;
514:
515:         /* zero-fill the fractional page of the DATA section */
516:         end = PAGE_ALIGN(elf32_ph->p_vaddr + elf32_ph->p_filesz);
517:         start = elf32_ph->p_vaddr + elf32_ph->p_filesz;
518:         length = end - start;
519:
520:         /* this will generate a page fault which will load the page in */
521:         memset_b((void *)start, NULL, length);
522:
523:         /* setup the BSS section */
524:         start = elf32_ph->p_vaddr + elf32_ph->p_filesz;
525:         start = PAGE_ALIGN(start);
526:         end = elf32_ph->p_vaddr + elf32_ph->p_memsz;
527:         end = PAGE_ALIGN(end);
528:         length = end - start;
529:         errno = do_mmap(NULL, start, length, PROT_READ | PROT_WRITE, MAP_PRIVATE
| MAP_FIXED, 0, P_BSS, 0);
530:         if(errno < 0 && errno > -PAGE_SIZE) {
531:             send_sig(current, SIGSEGV);
532:             return -ENOEXEC;
533:         }
534:         current->brk_lower = start;
535:
536:         /* setup the HEAP section */
537:         start = elf32_ph->p_vaddr + elf32_ph->p_memsz;
538:         start = PAGE_ALIGN(start);
539:         length = PAGE_SIZE;
540:         errno = do_mmap(NULL, start, length, PROT_READ | PROT_WRITE, MAP_PRIVATE
| MAP_FIXED, 0, P_HEAP, 0);
541:         if(errno < 0 && errno > -PAGE_SIZE) {
542:             send_sig(current, SIGSEGV);
543:             return -ENOEXEC;
544:         }
545:         current->brk = start;
546:
547:         /* setup the STACK section */
548:         sp = KERNEL_BASE_ADDR - 4;          /* formerly 0xBFFFFFFC */
549:         sp -= ae_str_len;
550:         str = sp;          /* this is the address of the first string (argv[0]) */
551:         sp &= ~3;
552:         sp -= at_base ? (AT_ITEMS * 2) * sizeof(unsigned int) : 2 * sizeof(unsigned
int);
553:         sp -= ae_ptr_len;
554:         length = KERNEL_BASE_ADDR - (sp & PAGE_MASK);
555:         errno = do_mmap(NULL, sp & PAGE_MASK, length, PROT_READ | PROT_WRITE | P
ROT_EXEC, MAP_PRIVATE | MAP_FIXED, 0, P_STACK, 0);
556:         if(errno < 0 && errno > -PAGE_SIZE) {
557:             send_sig(current, SIGSEGV);
558:             return -ENOEXEC;
559:         }
560:
561:         elf_create_stack(barg, (unsigned int *)sp, str, at_base, elf32_h, phdr_a
ddr);
562:
563:         /* set %esp to point to 'argc' */
564:         sc->oldesp = sp;
565:         sc->eflags = 0x202;          /* FIXME: linux 2.2 = 0x292 */
566:         sc->eip = current->entry_address;
567:         sc->err = 0;
568:         sc->eax = 0;
569:         sc->ecx = 0;
570:         sc->edx = 0;
571:         sc->ebx = 0;
572:         sc->ebp = 0;
573:         sc->esi = 0;
574:         sc->edi = 0;

```

```
575:         return 0;  
576:     }
```

## fs/fd.c

Page 1/1

```
1: /*
2:  * fiwix/fs/fd.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/errno.h>
9: #include <fiwix/types.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/sleep.h>
12: #include <fiwix/stdio.h>
13: #include <fiwix/string.h>
14:
15: struct fd *fd_table;
16:
17: static struct resource fd_resource = { NULL, NULL };
18:
19: int get_new_fd(struct inode *i)
20: {
21:     unsigned int n;
22:
23:     lock_resource(&fd_resource);
24:
25:     for(n = 1; n < NR_OPENS; n++) {
26:         if(fd_table[n].count == 0) {
27:             memset_b(&fd_table[n], NULL, sizeof(struct fd));
28:             fd_table[n].inode = i;
29:             fd_table[n].count = 1;
30:             unlock_resource(&fd_resource);
31:             return n;
32:         }
33:     }
34:
35:     unlock_resource(&fd_resource);
36:
37:     return -ENFILE;
38: }
39:
40: void release_fd(unsigned int fd)
41: {
42:     lock_resource(&fd_resource);
43:     fd_table[fd].count = 0;
44:     unlock_resource(&fd_resource);
45: }
46:
47: void fd_init(void)
48: {
49:     memset_b(fd_table, NULL, fd_table_size);
50: }
```

## fs/filesystems.c

Page 1/2

```

1: /*
2:  * fiwix/fs/filesystems.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_pipe.h>
13: #include <fiwix/fs_iso9660.h>
14: #include <fiwix/fs_proc.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: int register_filesystem(const char *name, struct fs_operations *fsop)
19: {
20:     int n;
21:     __dev_t dev;
22:
23:     for(n = 0; n < NR_FILESYSTEMS; n++) {
24:         if(filesystems_table[n].name) {
25:             if(strcmp(filesystems_table[n].name, name) == 0) {
26:                 printk("WARNING: %s(): filesystem '%s' already r
registered!\n", __FUNCTION__, name);
27:                 return 1;
28:             }
29:         }
30:         if(!filesystems_table[n].name) {
31:             filesystems_table[n].name = name;
32:             filesystems_table[n].fsop = fsop;
33:             if((fsop->flags & FSOP_KERN_MOUNT)) {
34:                 dev = fsop->fsdev;
35:                 return kern_mount(dev, &filesystems_table[n]);
36:             }
37:             return 0;
38:         }
39:     }
40:     printk("WARNING: %s(): filesystems table is full!\n", __FUNCTION__);
41:     return 1;
42: }
43:
44: struct filesystems * get_filesystem(const char *name)
45: {
46:     int n;
47:
48:     if(!name) {
49:         return NULL;
50:     }
51:     for(n = 0; n < NR_FILESYSTEMS; n++) {
52:         if(!filesystems_table[n].name) {
53:             continue;
54:         }
55:         if(strcmp(filesystems_table[n].name, name) == 0) {
56:             return &filesystems_table[n];
57:         }
58:     }
59:     return NULL;
60: }
61:
62: void fs_init(void)
63: {
64:     memset_b(filesystems_table, NULL, sizeof(filesystems_table));
65:
66:     if(minix_init()) {

```

**fs/filesystems.c**

Page 2/2

```
67:                printk("%s(): unable to register 'minix' filesystem.\n", __FUNCT
ION__);
68:                }
69:                if(ext2_init()) {
70:                    printk("%s(): unable to register 'ext2' filesystem.\n", __FUNCTI
ON__);
71:                }
72:                if(pipefs_init()) {
73:                    printk("%s(): unable to register 'pipefs' filesystem.\n", __FUNC
TION__);
74:                }
75:                if(iso9660_init()) {
76:                    printk("%s(): unable to register 'iso9660' filesystem.\n", __FUN
CTION__);
77:                }
78:                if(procfs_init()) {
79:                    printk("%s(): unable to register 'procfs' filesystem.\n", __FUNC
TION__);
80:                }
81: }
```



## fs/inode.c

Page 1/7

```

1: /*
2:  * fiwix/fs/inode.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: /*
9:  * inode.c implements a cache with a free list as a doubly circular linked
10:  * list and a chained hash table with doubly linked lists.
11:  *
12:  * hash table
13:  * +-----+ +-----+ +-----+ +-----+
14:  * | index | |prev|data|next| |prev|data|next| |prev|data|next|
15:  * |   0   --> | / |   |   |   | <--- |   |   |   | <--- |   |   |   |
16:  * +-----+ +-----+ +-----+ +-----+
17:  * +-----+ +-----+ +-----+ +-----+
18:  * | index | |prev|data|next| |prev|data|next| |prev|data|next|
19:  * |   1   --> | / |   |   |   | <--- |   |   |   | <--- |   |   |   |
20:  * +-----+ +-----+ +-----+ +-----+
21:  *                               (inode)       (inode)       (inode)
22:  *      ...
23:  */
24:
25: #include <fiwix/asm.h>
26: #include <fiwix/sleep.h>
27: #include <fiwix/sched.h>
28: #include <fiwix/fs.h>
29: #include <fiwix/filesystems.h>
30: #include <fiwix/stat.h>
31: #include <fiwix/errno.h>
32: #include <fiwix/mm.h>
33: #include <fiwix/stdio.h>
34: #include <fiwix/string.h>
35:
36: #define INODE_HASH(dev, inode) (((__dev_t)(dev) ^ (__ino_t)(inode)) % (NR_INO_H
ASH))
37: #define NR_INODES      (inode_table_size / sizeof(struct inode))
38: #define NR_INO_HASH    (inode_hash_table_size / sizeof(unsigned int))
39:
40: struct inode *inode_table;          /* inode pool */
41: struct inode *inode_head;          /* inode pool head */
42: struct inode **inode_hash_table;
43:
44: int inodes_on_free_list = 0;
45: static struct resource sync_resource = { NULL, NULL };
46:
47: static void insert_to_hash(struct inode *i)
48: {
49:     struct inode **h;
50:     int n;
51:
52:     n = INODE_HASH(i->dev, i->inode);
53:     h = &inode_hash_table[n];
54:
55:     if(!*h) {
56:         *h = i;
57:         (*h)->prev_hash = (*h)->next_hash = NULL;
58:     } else {
59:         i->prev_hash = NULL;
60:         i->next_hash = *h;
61:         (*h)->prev_hash = i;
62:         *h = i;
63:     }
64: }
65:
66: static void remove_from_hash(struct inode *i)

```

```

67: {
68:     struct inode **h;
69:     int n;
70:
71:     n = INODE_HASH(i->dev, i->inode);
72:     h = &inode_hash_table[n];
73:
74:     while(*h) {
75:         if(*h == i) {
76:             if((*h)->next_hash) {
77:                 (*h)->next_hash->prev_hash = (*h)->prev_hash;
78:             }
79:             if((*h)->prev_hash) {
80:                 (*h)->prev_hash->next_hash = (*h)->next_hash;
81:             }
82:             if(h == &inode_hash_table[n]) {
83:                 *h = (*h)->next_hash;
84:             }
85:             break;
86:         }
87:         h = &(*h)->next_hash;
88:     }
89: }
90:
91: static void remove_from_free_list(struct inode *i)
92: {
93:     i->prev_free->next_free = i->next_free;
94:     i->next_free->prev_free = i->prev_free;
95:     inodes_on_free_list--;
96:     if(i == inode_head) {
97:         inode_head = i->next_free;
98:     }
99: }
100:
101: static void inode_wait(struct inode *i)
102: {
103:     unsigned long int flags;
104:
105:     for(;;) {
106:         SAVE_FLAGS(flags); CLI();
107:         if(i->locked) {
108:             RESTORE_FLAGS(flags);
109:             sleep(i, PROC_UNINTERRUPTIBLE);
110:         } else {
111:             break;
112:         }
113:     }
114:     RESTORE_FLAGS(flags);
115: }
116:
117: static struct inode * get_free_inode(void)
118: {
119:     struct inode *i;
120:
121:     /* no more inodes on free list */
122:     if(inode_head == inode_head->next_free) {
123:         return NULL;
124:     }
125:
126:     i = inode_head;
127:     inode_head->next_free->prev_free = inode_head->prev_free;
128:     inode_head->prev_free->next_free = inode_head->next_free;
129:     inode_head = inode_head->next_free;
130:
131:     return i;
132: }
133:

```

## fs/inode.c

Page 3/7

```
134: static int read_inode(struct inode *i)
135: {
136:     int errno;
137:
138:     inode_lock(i);
139:     if(i->sb && i->sb->fsop && i->sb->fsop->read_inode) {
140:         errno = i->sb->fsop->read_inode(i);
141:         inode_unlock(i);
142:         return errno;
143:     }
144:     inode_unlock(i);
145:     return -EINVAL;
146: }
147:
148: static int write_inode(struct inode *i)
149: {
150:     int errno;
151:
152:     errno = 1;
153:
154:     inode_lock(i);
155:     if(i->sb && i->sb->fsop && i->sb->fsop->write_inode) {
156:         errno = i->sb->fsop->write_inode(i);
157:     } else {
158:         /* i.e. PIPE_DEV inodes can't be flushed on disk */
159:         i->dirty = 0;
160:         errno = 0;
161:     }
162:     inode_unlock(i);
163:
164:     return errno;
165: }
166:
167: static struct inode * search_inode_hash(__dev_t dev, __ino_t inode)
168: {
169:     struct inode *i;
170:     int n;
171:
172:     n = INODE_HASH(dev, inode);
173:     i = inode_hash_table[n];
174:
175:     while(i) {
176:         if(i->dev == dev && i->inode == inode) {
177:             return i;
178:         }
179:         i = i->next_hash;
180:     }
181:
182:     return NULL;
183: }
184:
185: static struct inode * get_blank_inode(void)
186: {
187:     unsigned long int flags;
188:     struct inode *i;
189:
190:     SAVE_FLAGS(flags); CLI();
191:
192:     if((i = get_free_inode())) {
193:         remove_from_free_list(i);
194:         remove_from_hash(i);
195:         i->i_mode = 0;
196:         i->i_uid = 0;
197:         i->i_size = 0;
198:         i->i_atime = 0;
199:         i->i_ctime = 0;
200:         i->i_mtime = 0;
```

## fs/inode.c

Page 4/7

```

201:         i->i_gid = 0;
202:         i->i_nlink = 0;
203:         i->i_blocks = 0;
204:         i->i_flags = 0;
205:         i->locked = 0;
206:         i->dirty = 0;
207:         i->mount_point = NULL;
208:         i->dev = 0;
209:         i->inode = 0;
210:         i->count = 0;
211:         i->rdev = 0;
212:         i->fsop = NULL;
213:         i->sb = NULL;
214:         memset_b(&i->u, NULL, sizeof(i->u));
215:     }
216:     RESTORE_FLAGS(flags);
217:     return i;
218: }
219:
220: void inode_lock(struct inode *i)
221: {
222:     unsigned long int flags;
223:
224:     for(;;) {
225:         SAVE_FLAGS(flags); CLI();
226:         if(i->locked) {
227:             RESTORE_FLAGS(flags);
228:             sleep(i, PROC_UNINTERRUPTIBLE);
229:         } else {
230:             break;
231:         }
232:     }
233:     i->locked = 1;
234:     RESTORE_FLAGS(flags);
235: }
236:
237: void inode_unlock(struct inode *i)
238: {
239:     unsigned long int flags;
240:
241:     SAVE_FLAGS(flags); CLI();
242:     i->locked = 0;
243:     wakeup(i);
244:     RESTORE_FLAGS(flags);
245: }
246:
247: struct inode * ialloc(struct superblock *sb, int mode)
248: {
249:     int errno;
250:     struct inode *i;
251:
252:     if((i = get_blank_inode())) {
253:         i->sb = sb;
254:         i->rdev = sb->dev;
255:         if(i->sb && i->sb->fsop && i->sb->fsop->ialloc) {
256:             errno = i->sb->fsop->ialloc(i, mode);
257:         } else {
258:             printk("WARNING: this filesystem does not have the iallo
c() method!\n");
259:             i->count = 1;
260:             i->sb = NULL;
261:             iput(i);
262:             return NULL;
263:         }
264:         if(errno) {
265:             i->count = 1;
266:             i->sb = NULL;

```

## fs/inode.c

Page 5/7

```

267:             iput(i);
268:             return NULL;
269:         }
270:         i->dev = sb->dev;
271:         insert_to_hash(i);
272:         return i;
273:     }
274:     printk("WARNING: %s(): no more inodes on free list!\n", __FUNCTION__);
275:     return NULL;
276: }
277:
278: struct inode * iget(struct superblock *sb, __ino_t inode)
279: {
280:     unsigned long int flags;
281:     struct inode *i;
282:
283:     if(!inode) {
284:         return NULL;
285:     }
286:
287:     for(;;) {
288:         if((i = search_inode_hash(sb->dev, inode)) {
289:             inode_wait(i);
290:             SAVE_FLAGS(flags); CLI();
291:
292:             /* update superblock pointer from mount_table */
293:             i->sb = sb;
294:
295:             if(i->mount_point) {
296:                 i = i->mount_point;
297:             }
298:             /* FIXME: i->locked = 1; ? */
299:             if(++i->count == 1) {
300:                 remove_from_free_list(i);
301:             }
302:             RESTORE_FLAGS(flags);
303:             return i;
304:         }
305:
306:         if(!(i = get_blank_inode())) {
307:             printk("WARNING: %s(): no more inodes on free list! (%d)
.\n", __FUNCTION__, inodes_on_free_list);
308:             return NULL;
309:         }
310:
311:         SAVE_FLAGS(flags); CLI();
312:         i->dev = i->rdev = sb->dev;
313:         i->inode = inode;
314:         i->sb = sb;
315:         i->count = 1;
316:         RESTORE_FLAGS(flags);
317:         if(read_inode(i)) {
318:             iput(i);
319:             return NULL;
320:         }
321:         insert_to_hash(i);
322:         /* FIXME: i->locked = 1; ? */
323:         return i;
324:     }
325: }
326:
327: int bmap(struct inode *i, __off_t offset, int mode)
328: {
329:     if(i->fsop && i->fsop->bmap) {
330:         return i->fsop->bmap(i, offset, mode);
331:     }
332:     return -EPERM;

```

## fs/inode.c

Page 6/7

```

333: }
334:
335: int check_fs_busy(__dev_t dev, struct inode *root)
336: {
337:     struct inode *i;
338:     unsigned int n;
339:
340:     i = &inode_table[0];
341:     for(n = 0; n < NR_INODES; n++, i = &inode_table[n]) {
342:         if(i->dev == dev && i->count) {
343:             if(i == root && i->count == 1) {
344:                 continue;
345:             }
346:             /* FIXME: to be removed */
347:             printk("WARNING: root %d with count %d (on dev %d,%d)\n"
, root->inode, root->count, MAJOR(i->dev), MINOR(i->dev));
348:             printk("WARNING: inode %d with count %d (on dev %d,%d)\n"
, i->inode, i->count, MAJOR(i->dev), MINOR(i->dev));
349:             return 1;
350:         }
351:     }
352:     return 0;
353: }
354:
355: void iput(struct inode *i)
356: {
357:     unsigned long int flags;
358:
359:     /* this solves the problem with rmdir("/") and iput(dir) which is NULL */
360:     if(!i) {
361:         return;
362:     }
363:
364:     if(!i->count) {
365:         printk("WARNING: %s(): trying to free an already freed inode (%d
)! \n", __FUNCTION__, i->inode);
366:         return;
367:     }
368:
369:     SAVE_FLAGS(flags); CLI();
370:
371:     if(--i->count == 0) {
372:         if(!i->i_nlink) {
373:             if(i->sb && i->sb->fsop && i->sb->fsop->ifree) {
374:                 inode_lock(i);
375:                 i->sb->fsop->ifree(i);
376:                 remove_from_hash(i);
377:                 inode_unlock(i);
378:             }
379:         }
380:         if(i->dirty) {
381:             if(write_inode(i)) {
382:                 printk("WARNING: %s(): can't write inode %d (%d,
%d), will remain as dirty.\n", __FUNCTION__, i->inode, MAJOR(i->dev), MINOR(i->dev));
383:                 RESTORE_FLAGS(flags);
384:                 return;
385:             }
386:         }
387:         if(!inode_head) {
388:             i->prev_free = i->next_free = i;
389:             inode_head = i;
390:         } else {
391:             i->next_free = inode_head;
392:             i->prev_free = inode_head->prev_free;
393:             inode_head->prev_free->next_free = i;
394:             inode_head->prev_free = i;

```

```

395:         }
396:         inodes_on_free_list++;
397:     }
398:
399:     RESTORE_FLAGS(flags);
400: }
401:
402: void sync_inodes(__dev_t dev)
403: {
404:     struct inode *i;
405:     int n;
406:
407:     i = &inode_table[0];
408:
409:     lock_resource(&sync_resource);
410:     for(n = 0; n < NR_INODES; n++) {
411:         if(i->dirty) {
412:             if(!dev || i->dev == dev) {
413:                 inode_wait(i);
414:                 if(write_inode(i)) {
415:                     printk("WARNING: %s(): can't write inode
%d (%d,%d), will remain as dirty.\n", __FUNCTION__, i->inode, MAJOR(i->dev), MINOR(i->
dev));
416:                 }
417:             }
418:         }
419:         i++;
420:     }
421:     unlock_resource(&sync_resource);
422:     return;
423: }
424:
425: void invalidate_inodes(__dev_t dev)
426: {
427:     unsigned long int flags;
428:     unsigned int n;
429:     struct inode *i;
430:
431:     i = &inode_table[0];
432:     SAVE_FLAGS(flags); CLI();
433:
434:     for(n = 0; n < NR_INODES; n++) {
435:         if(i->dev == dev) {
436:             inode_wait(i);
437:             remove_from_hash(i);
438:             i->locked = 0;
439:             wakeup(&inode_wait);
440:         }
441:         i++;
442:     }
443:
444:     RESTORE_FLAGS(flags);
445: }
446:
447: void inode_init(void)
448: {
449:     struct inode *i;
450:     unsigned int n;
451:
452:     memset_b(inode_table, NULL, inode_table_size);
453:     memset_b(inode_hash_table, NULL, inode_hash_table_size);
454:     for(n = 0; n < NR_INODES; n++) {
455:         i = &inode_table[n];
456:         i->count = 1;
457:         iput(i);
458:     }
459: }

```

## fs/locks.c

Page 1/4

```

1: /*
2:  * fiwix/fs/locks.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/errno.h>
9: #include <fiwix/types.h>
10: #include <fiwix/locks.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/sleep.h>
13: #include <fiwix/sched.h>
14: #include <fiwix/stdio.h>
15: #include <fiwix/string.h>
16:
17: static struct resource flock_resource = { NULL, NULL };
18:
19: static struct flock_file * get_new_flock(struct inode *i)
20: {
21:     int n;
22:     struct flock_file *ff;
23:
24:     lock_resource(&flock_resource);
25:
26:     for(n = 0; n < NR_FLOCKS; n++) {
27:         ff = &flock_file_table[n];
28:         if(!ff->inode) {
29:             ff->inode = i; /* mark it as busy */
30:             unlock_resource(&flock_resource);
31:             return ff;
32:         }
33:     }
34:
35:     printk("WARNING: %s(): no more free slots in flock file table.\n");
36:     unlock_resource(&flock_resource);
37:     return NULL;
38: }
39:
40: static void release_flock(struct flock_file *ff)
41: {
42:     memset_b(ff, 0, sizeof(struct flock_file));
43: }
44:
45: static struct flock_file * get_flock_file(struct inode *i, int op, struct proc *
p)
46: {
47:     int n;
48:     struct flock_file *ff;
49:
50:     lock_resource(&flock_resource);
51:
52:     ff = NULL;
53:     for(n = 0; n < NR_FLOCKS; n++) {
54:         ff = &flock_file_table[n];
55:         if(ff->inode != i) {
56:             continue;
57:         }
58:         if(p && p != ff->proc) {
59:             continue;
60:         }
61:         break;
62:     }
63:     unlock_resource(&flock_resource);
64:     return ff;
65: }
66:

```



## fs/locks.c

Page 2/4

```

67: int posix_lock(int ufd, int cmd, struct flock *fl)
68: {
69:     int n;
70:     struct flock_file *ff;
71:     struct inode *i;
72:     unsigned char type;
73:
74:     lock_resource(&flock_resource);
75:     i = fd_table[current->fd[ufd]].inode;
76:     for(n = 0; n < NR_FLOCKS; n++) {
77:         ff = &flock_file_table[n];
78:         if(ff->inode != i) {
79:             continue;
80:         }
81:         break;
82:     }
83:     unlock_resource(&flock_resource);
84:     if(cmd == F_GETLK) {
85:         if(ff->inode == i) {
86:             fl->l_type = ff->type & LOCK_SH ? F_RDLCK : F_WRLCK;
87:             fl->l_whence = SEEK_SET;
88:             fl->l_start = 0;
89:             fl->l_len = 0;
90:             fl->l_pid = ff->proc->pid;
91:         } else {
92:             fl->l_type = F_UNLCK;
93:         }
94:     }
95:
96:     switch(fl->l_type) {
97:         case F_RDLCK:
98:             type = LOCK_SH;
99:             break;
100:        case F_WRLCK:
101:            type = LOCK_EX;
102:            break;
103:        case F_UNLCK:
104:            type = LOCK_UN;
105:            break;
106:        default:
107:            return -EINVAL;
108:    }
109:    if(cmd == F_SETLK) {
110:        return flock_inode(i, type);
111:    }
112:    if(cmd == F_SETLKW) {
113:        return flock_inode(i, type | LOCK_NB);
114:    }
115:    return 0;
116: }
117:
118: void flock_release_inode(struct inode *i)
119: {
120:     int n;
121:     struct flock_file *ff;
122:
123:     lock_resource(&flock_resource);
124:     for(n = 0; n < NR_FLOCKS; n++) {
125:         ff = &flock_file_table[n];
126:         if(ff->inode != i) {
127:             continue;
128:         }
129:         if(ff->proc != current) {
130:             continue;
131:         }
132:         wakeup(ff);
133:         release_flock(ff);

```

## fs/locks.c

Page 3/4

```

134:         }
135:         unlock_resource(&flock_resource);
136:     }
137:
138: int flock_inode(struct inode *i, int op)
139: {
140:     int n;
141:     struct flock_file *ff, *new;
142:
143:     if(op & LOCK_UN) {
144:         if((ff = get_flock_file(i, op, current))) {
145:             wakeup(ff);
146:             release_flock(ff);
147:         }
148:         return 0;
149:     }
150:
151: loop:
152:     lock_resource(&flock_resource);
153:     new = NULL;
154:     for(n = 0; n < NR_FLOCKS; n++) {
155:         ff = &flock_file_table[n];
156:         if(ff->inode != i) {
157:             continue;
158:         }
159:         if(op & LOCK_SH) {
160:             if(ff->type & LOCK_EX) {
161:                 if(ff->proc == current) {
162:                     new = ff;
163:                     wakeup(ff);
164:                     break;
165:                 }
166:                 unlock_resource(&flock_resource);
167:                 if(op & LOCK_NB) {
168:                     return -EWOULDBLOCK;
169:                 }
170:                 if(sleep(ff, PROC_INTERRUPTIBLE)) {
171:                     return -EINTR;
172:                 }
173:                 goto loop;
174:             }
175:         }
176:         if(op & LOCK_EX) {
177:             if(ff->proc == current) {
178:                 new = ff;
179:                 continue;
180:             }
181:             unlock_resource(&flock_resource);
182:             if(op & LOCK_NB) {
183:                 return -EWOULDBLOCK;
184:             }
185:             if(sleep(ff, PROC_INTERRUPTIBLE)) {
186:                 return -EINTR;
187:             }
188:             goto loop;
189:         }
190:     }
191:     unlock_resource(&flock_resource);
192:
193:     if(!new) {
194:         if(!(new = get_new_flock(i))) {
195:             return -ENOLCK;
196:         }
197:     }
198:     new->inode = i;
199:     new->type = op;
200:     new->proc = current;

```

**fs/locks.c**

Page 4/4

```
201:
202:     return 0;
203: }
204:
205: void flock_init(void)
206: {
207:     memset_b(flock_file_table, NULL, sizeof(flock_file_table));
208: }
```

**fs/Makefile**

Page 1/1

```
1: # fiwix/fs/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: FSDIRS = minix ext2 pipefs iso9660 procfs
13: FILESYSTEMS = minix/minix.o ext2/ext2.o pipefs/pipefs.o iso9660/iso9660.o \
14:     procfs/procfs.o
15: OBJS = filesystems.o devices.o buffer.o fd.o locks.o super.o inode.o \
16:     namei.o elf.o script.o
17:
18: fs:     $(OBJS)
19:     @for n in $(FSDIRS) ; do (cd $$n ; $(MAKE)) ; done
20:     $(LD) $(LDFLAGS) -r $(FILESYSTEMS) $(OBJS) -o fs.o
21:
22: clean:
23:     @for n in $(FSDIRS) ; do (cd $$n ; $(MAKE) clean) ; done
24:     rm -f *.o
25:
```

## fs/namei.c

Page 1/3

```

1: /*
2:  * fiwix/fs/namei.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/sleep.h>
10: #include <fiwix/sched.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/mman.h>
16: #include <fiwix/errno.h>
17: #include <fiwix/stdio.h>
18: #include <fiwix/string.h>
19:
20: static int namei_lookup(char *name, struct inode *dir, struct inode **i_res)
21: {
22:     if(dir->fsop && dir->fsop->lookup) {
23:         return dir->fsop->lookup(name, dir, i_res);
24:     }
25:     return -EACCES;
26: }
27:
28: static int do_namei(char *path, struct inode *dir, struct inode **i_res, struct
inode **d_res, int follow_links)
29: {
30:     char *name, *ptr_name;
31:     struct inode *i;
32:     struct superblock *sb;
33:
34:     int errno;
35:
36:     *i_res = dir;
37:     for(;;) {
38:         while(*path == '/') {
39:             path++;
40:         }
41:         if(*path == NULL) {
42:             return 0;
43:         }
44:
45:         /* extracts the next component of the path */
46:         if(! (name = (char *)kmallocc())) {
47:             return -ENOMEM;
48:         }
49:         ptr_name = name;
50:         while(*path != NULL && *path != '/') {
51:             if(ptr_name > name + NAME_MAX - 1) {
52:                 break;
53:             }
54:             *ptr_name++ = *path++;
55:         }
56:         *ptr_name = NULL;
57:
58:         /*
59:          * If the inode is the root of a file system, then return the
60:          * inode on which the file system was mounted.
61:          */
62:         if(name[0] == '.' && name[1] == '.' && name[2] == NULL) {
63:             if(dir == dir->sb->root) {
64:                 sb = dir->sb;
65:                 iput(dir);
66:                 dir = sb->dir;

```

## fs/namei.c

Page 2/3

```

67:                                dir->count++;
68:                                }
69:                                }
70:
71:                                if((errno = check_permission(TO_EXEC, dir)) {
72:                                    break;
73:                                }
74:
75:                                if((errno = namei_lookup(name, dir, &i)) {
76:                                    break;
77:                                }
78:
79:                                kfree((unsigned int) name);
80:                                if(*path == '/') {
81:                                    if(!S_ISDIR(i->i_mode) && !S_ISLNK(i->i_mode)) {
82:                                        iput(dir);
83:                                        iput(i);
84:                                        return -ENOTDIR;
85:                                    }
86:                                    if(S_ISLNK(i->i_mode)) {
87:                                        if(i->fsop && i->fsop->followlink) {
88:                                            if((errno = i->fsop->followlink(dir, i,
&i)) {
89:                                                iput(dir);
90:                                                return errno;
91:                                            }
92:                                        }
93:                                    }
94:                                } else {
95:                                    if((i->fsop && i->fsop->followlink) && follow_links) {
96:                                        if((errno = i->fsop->followlink(dir, i, &i)) {
97:                                            iput(dir);
98:                                            return errno;
99:                                        }
100:                                    }
101:                                }
102:
103:                                if(d_res) {
104:                                    if(*d_res) {
105:                                        iput(*d_res);
106:                                    }
107:                                    *d_res = dir;
108:                                } else {
109:                                    iput(dir);
110:                                }
111:                                dir = i;
112:                                *i_res = i;
113:                            }
114:
115:                            kfree((unsigned int) name);
116:                            if(d_res) {
117:                                if(*d_res) {
118:                                    iput(*d_res);
119:                                }
120:                                /*
121:                                 * If that was the last component of the path,
122:                                 * then return the directory.
123:                                 */
124:                                if(*path == NULL) {
125:                                    *d_res = dir;
126:                                    dir->count++;
127:                                } else {
128:                                    /* that's an non-existent directory */
129:                                    *d_res = NULL;
130:                                    errno = -ENOTDIR;
131:                                }
132:                                iput(dir);

```

## fs/namei.c

Page 3/3

```
133:         *i_res = NULL;
134:     } else {
135:         iput(dir);
136:     }
137:
138:     return errno;
139: }
140:
141: int parse_namei(char *path, struct inode *base_dir, struct inode **i_res, struct
inode **d_res, int follow_links)
142: {
143:     struct inode *dir;
144:     int errno;
145:
146:     if(!path) {
147:         return -EFAULT;
148:     }
149:     if(*path == NULL) {
150:         return -ENOENT;
151:     }
152:
153:     if(!(dir = base_dir)) {
154:         dir = current->pwd;
155:     }
156:
157:     /* it is definitely an absolute path */
158:     if(path[0] == '/') {
159:         dir = current->root;
160:     }
161:     dir->count++;
162:     errno = do_namei(path, dir, i_res, d_res, follow_links);
163:     return errno;
164: }
165:
166: /*
167:  * namei() returns:
168:  * i_res -> the inode of the last component of the path, or NULL.
169:  * d_res -> the inode of the directory where i_res resides, or NULL.
170:  */
171: int namei(char *path, struct inode **i_res, struct inode **d_res, int follow_lin
ks)
172: {
173:     *i_res = NULL;
174:     if(d_res) {
175:         *d_res = NULL;
176:     }
177:     return parse_namei(path, NULL, i_res, d_res, follow_links);
178: }
```

## fs/script.c

```

1:  /*
2:  *  fiwix/fs/script.c
3:  *
4:  *  Copyright 2019, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/limits.h>
9:  #include <fiwix/process.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/stdio.h>
12: #include <fiwix/string.h>
13:
14: int script_load(char *interpreter, char *args, char *data)
15: {
16:     char *p;
17:     int n, noargs;
18:
19:     /* has shebang? */
20:     if(data[0] != '#' && data[1] != '!') {
21:         return -ENOEXEC;
22:     }
23:
24:     /* discard possible blanks before the interpreter name */
25:     for(n = 2; n < NAME_MAX; n++) {
26:         if(data[n] != ' ' && data[n] != '\t') {
27:             break;
28:         }
29:     }
30:
31:     /* get the interpreter name */
32:     p = interpreter;
33:     noargs = 0;
34:     while(n < NAME_MAX) {
35:         if(data[n] == '\n' || data[n] == NULL) {
36:             noargs = 1;
37:             break;
38:         }
39:         if(data[n] == ' ' || data[n] == '\t') {
40:             break;
41:         }
42:         *p = data[n];
43:         n++;
44:         p++;
45:     }
46:
47:     if(!interpreter) {
48:         return -ENOEXEC;
49:     }
50:
51:
52:     /* get the interpreter arguments */
53:     if(!noargs) {
54:         p = args;
55:         /* discard possible blanks before the arguments */
56:         while(n < NAME_MAX) {
57:             if(data[n] != ' ' && data[n] != '\t') {
58:                 break;
59:             }
60:             n++;
61:         }
62:         while(n < NAME_MAX) {
63:             if(data[n] == '\n' || data[n] == NULL) {
64:                 break;
65:             }
66:             *p = data[n];
67:             n++;

```



**fs/script.c**

Page 2/2

```
68:                                     p++;
69:                                     }
70:     }
71:
72:     return 0;
73: }
```

## fs/super.c

Page 1/4

```

1: /*
2:  * fiwix/fs/super.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/kernel.h>
10: #include <fiwix/types.h>
11: #include <fiwix/errno.h>
12: #include <fiwix/fs.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/filesystems.h>
15: #include <fiwix/sleep.h>
16: #include <fiwix/sched.h>
17: #include <fiwix/stdio.h>
18: #include <fiwix/string.h>
19:
20: struct mount *mount_table;
21: static struct resource sync_resource = { NULL, NULL };
22:
23: void superblock_lock(struct superblock *sb)
24: {
25:     unsigned long int flags;
26:
27:     for(;;) {
28:         SAVE_FLAGS(flags); CLI();
29:         if(sb->locked) {
30:             sb->wanted = 1;
31:             RESTORE_FLAGS(flags);
32:             sleep(&superblock_lock, PROC_UNINTERRUPTIBLE);
33:         } else {
34:             break;
35:         }
36:     }
37:     sb->locked = 1;
38:     RESTORE_FLAGS(flags);
39: }
40:
41: void superblock_unlock(struct superblock *sb)
42: {
43:     unsigned long int flags;
44:
45:     SAVE_FLAGS(flags); CLI();
46:     sb->locked = 0;
47:     if(sb->wanted) {
48:         sb->wanted = 0;
49:         wakeup(&superblock_lock);
50:     }
51:     RESTORE_FLAGS(flags);
52: }
53:
54: struct mount * get_free_mount_point(__dev_t dev)
55: {
56:     unsigned long int flags;
57:     int n;
58:
59:     if(!dev) {
60:         printk("%s(): invalid device %d,%d.\n", __FUNCTION__, MAJOR(dev)
, MINOR(dev));
61:         return NULL;
62:     }
63:
64:     for(n = 0; n < NR_MOUNT_POINTS; n++) {
65:         if(mount_table[n].dev == dev) {
66:             printk("%s(): device %d,%d already mounted.\n", __FUNCTI

```

## fs/super.c

Page 2/4

```

ON__, MAJOR(dev), MINOR(dev));
67:         return NULL;
68:     }
69: }
70:
71:     SAVE_FLAGS(flags); CLI();
72:     for(n = 0; n < NR_MOUNT_POINTS; n++) {
73:         if(!mount_table[n].used) {
74:             /* 'dev' is saved here now for get_superblock() (which
75:              * in turn is called by read_inode(), which in turn is
76:              * called by iget(), which in turn is called by
77:              * read_superblock) to be able to find the device.
78:              */
79:             mount_table[n].dev = dev;
80:             mount_table[n].used = 1;
81:             RESTORE_FLAGS(flags);
82:             return &mount_table[n];
83:         }
84:     }
85:     RESTORE_FLAGS(flags);
86:
87:     printk("WARNING: %s(): mount-point table is full.\n", __FUNCTION__);
88:     return NULL;
89: }
90:
91: void release_mount_point(struct mount *mt)
92: {
93:     memset_b(mt, NULL, sizeof(struct mount));
94: }
95:
96: struct mount * get_mount_point(struct inode *i_target)
97: {
98:     int n;
99:
100:    for(n = 0; n < NR_MOUNT_POINTS; n++) {
101:        if(mount_table[n].used) {
102:            if(S_ISDIR(i_target->i_mode)) {
103:                if(mount_table[n].sb.root == i_target) {
104:                    return &mount_table[n];
105:                }
106:            }
107:            if(S_ISBLK(i_target->i_mode)) {
108:                if(mount_table[n].dev == i_target->rdev) {
109:                    return &mount_table[n];
110:                }
111:            }
112:        }
113:    }
114:    return NULL;
115: }
116:
117: struct superblock * get_superblock(__dev_t dev)
118: {
119:     int n;
120:
121:     for(n = 0; n < NR_MOUNT_POINTS; n++) {
122:         if(mount_table[n].used && mount_table[n].dev == dev) {
123:             return &mount_table[n].sb;
124:         }
125:     }
126:     return NULL;
127: }
128:
129: void sync_superblocks(__dev_t dev)
130: {
131:     struct superblock *sb;
132:     int n, errno;

```

## fs/super.c

Page 3/4

```

133:
134:     lock_resource(&sync_resource);
135:     for(n = 0; n < NR_MOUNT_POINTS; n++) {
136:         if(mount_table[n].used && (!dev || mount_table[n].dev == dev)) {
137:             sb = &mount_table[n].sb;
138:             if(sb->dirty && !(sb->flags & MS_RDONLY)) {
139:                 if(sb->fsop && sb->fsop->write_superblock) {
140:                     errno = sb->fsop->write_superblock(sb);
141:                     if(errno) {
142:                         printk("WARNING: %s(): I/O error
on device %d,%d while syncing superblock.\n", __FUNCTION__, MAJOR(sb->dev), MINOR(sb->
dev));
143:                     }
144:                 }
145:             }
146:         }
147:     }
148:     unlock_resource(&sync_resource);
149: }
150:
151: /* pseudo-file systems are only mountable by the kernel */
152: int kern_mount(__dev_t dev, struct filesystems *fs)
153: {
154:     struct mount *mt;
155:
156:     if(!(mt = get_free_mount_point(dev))) {
157:         return -EBUSY;
158:     }
159:
160:     if(fs->fsop->read_superblock(dev, &mt->sb)) {
161:         release_mount_point(mt);
162:         return -EINVAL;
163:     }
164:
165:     mt->dev = dev;
166:     strcpy(mt->devname, "none");
167:     strcpy(mt->dirname, "none");
168:     mt->sb.dir = NULL;
169:     mt->fs = fs;
170:     fs->mt = mt;
171:     return 0;
172: }
173:
174: int mount_root(void)
175: {
176:     struct filesystems *fs;
177:     struct mount *mt;
178:
179:     /* FIXME: before trying to mount the filesystem, we should first
180:      * check if '_rootdev' is a device successfully registered.
181:      */
182:
183:     if(!(fs = get_filesystem(_rootfstype))) {
184:         printk("WARNING: %s(): '%s' is not a registered filesystem. Defa
ulting to 'minix'.\n", __FUNCTION__, _rootfstype);
185:         if(!(fs = get_filesystem("minix"))) {
186:             PANIC("minix filesystem is not registered!\n");
187:         }
188:     }
189:
190:     if(!(mt = get_free_mount_point(_rootdev))) {
191:         PANIC("unable to get a free mount point.\n");
192:     }
193:
194:     mt->sb.flags = MS_RDONLY;
195:     if(fs->fsop && fs->fsop->read_superblock) {
196:         if(fs->fsop->read_superblock(_rootdev, &mt->sb)) {

```

## fs/super.c

Page 4/4

```
197:                                PANIC("unable to mount root filesystem on %s.\n", _rootd
evname);
198:                                }
199:                                }
200:
201:                                strcpy(mt->devname, "/dev/root");
202:                                strcpy(mt->dirname, "/");
203:                                mt->dev = _rootdev;
204:                                mt->sb.root->mount_point = mt->sb.root;
205:                                mt->sb.root->count++;
206:                                mt->sb.dir = mt->sb.root;
207:                                mt->sb.dir->count++;
208:                                mt->fs = fs;
209:
210:                                current->root = mt->sb.root;
211:                                current->root->count++;
212:                                current->pwd = mt->sb.root;
213:                                current->pwd->count++;
214:                                iput(mt->sb.root);
215:
216:                                printk("mounted root device (%s filesystem) in readonly mode.\n", fs->na
me);
217:                                return 0;
218: }
219:
220: void mount_init(void)
221: {
222:     memset_b(mount_table, NULL, mount_table_size);
223: }
```

## fs/ext2/bitmaps.c

Page 1/5

```

1:  /*
2:  *  fiwix/fs/ext2/bitmaps.c
3:  *
4:  *  Copyright 2019, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/kernel.h>
9:  #include <fiwix/types.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_ext2.h>
13: #include <fiwix/buffer.h>
14: #include <fiwix/errno.h>
15: #include <fiwix/stat.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: static int find_first_zero(struct superblock *sb, __blk_t block)
20: {
21:     unsigned char c;
22:     int blksize;
23:     int n, n2;
24:     struct buffer *buf;
25:
26:     blksize = sb->s_blocksize;
27:
28:     if(!(buf = bread(sb->dev, block, blksize))) {
29:         return -EIO;
30:     }
31:     for(n = 0; n < blksize; n++) {
32:         c = (unsigned char)buf->data[n];
33:         for(n2 = 0; n2 < 8; n2++) {
34:             if(!(c & (1 << n2))) {
35:                 brelse(buf);
36:                 return n2 + (n * 8) + 1;
37:             }
38:         }
39:     }
40:     brelse(buf);
41:     return 0;
42: }
43:
44: static int change_bit(int mode, struct superblock *sb, __blk_t block, int item)
45: {
46:     int byte, bit, mask;
47:     struct buffer *buf;
48:
49:     block += item / (sb->s_blocksize * 8);
50:     byte = (item % (sb->s_blocksize * 8)) / 8;
51:     bit = (item % (sb->s_blocksize * 8)) % 8;
52:     mask = 1 << bit;
53:
54:     if(!(buf = bread(sb->dev, block, sb->s_blocksize))) {
55:         return -EIO;
56:     }
57:
58:     if(mode == CLEAR_BIT) {
59:         if(!(buf->data[byte] & mask)) {
60:             brelse(buf);
61:             return 1;
62:         }
63:         buf->data[byte] &= ~mask;
64:     }
65:     if(mode == SET_BIT) {
66:         if((buf->data[byte] & mask)) {
67:             brelse(buf);

```

## fs/ext2/bitmaps.c

Page 2/5

```

68:             return 1;
69:         }
70:         buf->data[byte] |= mask;
71:     }
72:
73:     bwrite(buf);
74:     return 0;
75: }
76:
77: /*
78:  * Unlike of what Ext2 specifies/suggests, this inode allocation does NOT
79:  * try to assign inodes in the same block group of the directory in which
80:  * they will be created.
81:  */
82: int ext2_ialloc(struct inode *i, int mode)
83: {
84:     __ino_t inode;
85:     __blk_t block;
86:     struct superblock *sb;
87:     struct ext2_group_desc *gd;
88:     struct buffer *buf;
89:     int bg, d, errno;
90:
91:     sb = i->sb;
92:     superblock_lock(sb);
93:
94:     block = SUPERBLOCK + sb->u.ext2.sb.s_first_data_block;
95:     inode = 0;
96:     buf = NULL;
97:
98:     /* read through all group descriptors to find the first unallocated inod
e */
99:     for(bg = 0, d = 0; bg < sb->u.ext2.block_groups; bg++, d++) {
100:         if(!(bg % (sb->s_blocksize / sizeof(struct ext2_group_desc)))) {
101:             if(buf) {
102:                 brelse(buf);
103:                 block++;
104:                 d = 0;
105:             }
106:             if(!(buf = bread(sb->dev, block, sb->s_blocksize))) {
107:                 superblock_unlock(sb);
108:                 return -EIO;
109:             }
110:         }
111:         gd = (struct ext2_group_desc *) (buf->data + (d * sizeof(struct e
xt2_group_desc)));
112:         if(gd->bg_free_inodes_count) {
113:             if((inode = find_first_zero(sb, gd->bg_inode_bitmap)) {
114:                 break;
115:             }
116:         }
117:     }
118:     if(!inode) {
119:         brelse(buf);
120:         superblock_unlock(sb);
121:         return -ENOSPC;
122:     }
123:
124:     errno = change_bit(SET_BIT, sb, gd->bg_inode_bitmap, inode - 1);
125:     if(errno) {
126:         if(errno < 0) {
127:             printk("WARNING: %s(): unable to set inode %d.\n", __FUN
CTION__, inode);
128:             brelse(buf);
129:             superblock_unlock(sb);
130:             return errno;
131:         } else {

```

## fs/ext2/bitmaps.c

Page 3/5

```

132:             printk("WARNING: %s(): inode %d is already marked as use
d!\n", __FUNCTION__, inode);
133:         }
134:     }
135:
136:     inode += bg * EXT2_INODES_PER_GROUP(sb);
137:     gd->bg_free_inodes_count--;
138:     sb->u.ext2.sb.s_free_inodes_count--;
139:     if(S_ISDIR(mode)) {
140:         gd->bg_used_dirs_count++;
141:     }
142:     bwrite(buf);
143:
144:     i->inode = inode;
145:     i->i_atime = CURRENT_TIME;
146:     i->i_mtime = CURRENT_TIME;
147:     i->i_ctime = CURRENT_TIME;
148:
149:     superblock_unlock(sb);
150:     return 0;
151: }
152:
153: void ext2_ifree(struct inode *i)
154: {
155:     struct ext2_group_desc *gd;
156:     struct buffer *buf;
157:     struct superblock *sb;
158:     __blk_t b, bg;
159:     int errno;
160:
161:     if(!i->inode || i->inode > i->sb->u.ext2.sb.s_inodes_count) {
162:         printk("WARNING: %s(): invalid inode %d!\n", __FUNCTION__, i->in
ode);
163:         return;
164:     }
165:
166:     if(i->i_blocks) {
167:         ext2_truncate(i, 0);
168:     }
169:
170:     sb = i->sb;
171:     superblock_lock(sb);
172:
173:     b = SUPERBLOCK + sb->u.ext2.sb.s_first_data_block;
174:     bg = (i->inode - 1) / EXT2_INODES_PER_GROUP(sb);
175:     if(!(buf = bread(sb->dev, b + (bg / EXT2_DESC_PER_BLOCK(sb)), sb->s_bloc
ksize))) {
176:         superblock_unlock(sb);
177:         return;
178:     }
179:     gd = (struct ext2_group_desc *) (buf->data + ((bg % EXT2_DESC_PER_BLOCK(s
b)) * sizeof(struct ext2_group_desc)));
180:     errno = change_bit(CLEAR_BIT, sb, gd->bg_inode_bitmap, (i->inode - 1) %
EXT2_INODES_PER_GROUP(sb));
181:
182:     if(errno) {
183:         if(errno < 0) {
184:             printk("WARNING: %s(): unable to clear inode %d.\n", __F
UNCTION__, i->inode);
185:             brelse(buf);
186:             superblock_unlock(sb);
187:             return;
188:         } else {
189:             printk("WARNING: %s(): inode %d is already marked as fre
e!\n", __FUNCTION__, i->inode);
190:         }
191:     }

```



## fs/ext2/bitmaps.c

Page 4/5

```

192:
193:     gd->bg_free_inodes_count++;
194:     sb->u.ext2.sb.s_free_inodes_count++;
195:     if(S_ISDIR(i->i_mode)) {
196:         gd->bg_used_dirs_count--;
197:     }
198:     bwrite(buf);
199:
200:     i->i_size = 0;
201:     i->i_mtime = CURRENT_TIME;
202:     i->i_ctime = CURRENT_TIME;
203:     i->dirty = 1;
204:
205:     superblock_unlock(sb);
206:     return;
207: }
208:
209: int ext2_balloc(struct superblock *sb)
210: {
211:     __blk_t b, block;
212:     struct ext2_group_desc *gd;
213:     struct buffer *buf;
214:     int bg, d, errno;
215:
216:     superblock_lock(sb);
217:
218:     b = SUPERBLOCK + sb->u.ext2.sb.s_first_data_block;
219:     block = 0;
220:     buf = NULL;
221:
222:     /* read through all group descriptors to find the first unallocated bloc
k */
223:     for(bg = 0, d = 0; bg < sb->u.ext2.block_groups; bg++, d++) {
224:         if(!(bg % (sb->s_blocksize / sizeof(struct ext2_group_desc)))) {
225:             if(buf) {
226:                 brelse(buf);
227:                 b++;
228:                 d = 0;
229:             }
230:             if(!(buf = bread(sb->dev, b, sb->s_blocksize))) {
231:                 superblock_unlock(sb);
232:                 return -EIO;
233:             }
234:         }
235:         gd = (struct ext2_group_desc *) (buf->data + (d * sizeof(struct e
xt2_group_desc)));
236:         if(gd->bg_free_blocks_count) {
237:             if((block = find_first_zero(sb, gd->bg_block_bitmap))) {
238:                 break;
239:             }
240:         }
241:     }
242:     if(!block) {
243:         brelse(buf);
244:         superblock_unlock(sb);
245:         return -ENOSPC;
246:     }
247:
248:     errno = change_bit(SET_BIT, sb, gd->bg_block_bitmap, block - 1);
249:     if(errno) {
250:         if(errno < 0) {
251:             printk("WARNING: %s(): unable to set block %d.\n", __FUN
CTION__, block);
252:             brelse(buf);
253:             superblock_unlock(sb);
254:             return errno;
255:         } else {

```

## fs/ext2/bitmaps.c

Page 5/5

```

256:             printk("WARNING: %s(): block %d is already marked as use
d!\n", __FUNCTION__, block);
257:         }
258:     }
259:
260:     block += bg * EXT2_BLOCKS_PER_GROUP(sb);
261:     gd->bg_free_blocks_count--;
262:     sb->u.ext2.sb.s_free_blocks_count--;
263:     bwrite(buf);
264:
265:     superblock_unlock(sb);
266:     return block;
267: }
268:
269: void ext2_bfree(struct superblock *sb, int block)
270: {
271:     struct ext2_group_desc *gd;
272:     struct buffer *buf;
273:     __blk_t b, bg;
274:     int errno;
275:
276:     if(!block || block > sb->u.ext2.sb.s_blocks_count) {
277:         printk("WARNING: %s(): invalid block %d!\n", __FUNCTION__, block
);
278:         return;
279:     }
280:
281:     superblock_lock(sb);
282:
283:     b = SUPERBLOCK + sb->u.ext2.sb.s_first_data_block;
284:     bg = (block - 1) / EXT2_BLOCKS_PER_GROUP(sb);
285:     if(!(buf = bread(sb->dev, b + (bg / EXT2_DESC_PER_BLOCK(sb)), sb->s_bloc
ksize))) {
286:         superblock_unlock(sb);
287:         return;
288:     }
289:     gd = (struct ext2_group_desc *) (buf->data + ((bg % EXT2_DESC_PER_BLOCK(s
b)) * sizeof(struct ext2_group_desc)));
290:     errno = change_bit(CLEAR_BIT, sb, gd->bg_block_bitmap, (block - 1) % EXT
2_BLOCKS_PER_GROUP(sb));
291:
292:     if(errno) {
293:         if(errno < 0) {
294:             printk("WARNING: %s(): unable to free block %d.\n", __FU
NCTION__, block);
295:             brelse(buf);
296:             superblock_unlock(sb);
297:             return;
298:         } else {
299:             printk("WARNING: %s(): block %d is already marked as fre
e!\n", __FUNCTION__, block);
300:         }
301:     }
302:
303:     gd->bg_free_blocks_count++;
304:     sb->u.ext2.sb.s_free_blocks_count++;
305:     bwrite(buf);
306:
307:     superblock_unlock(sb);
308:     return;
309: }

```

## fs/ext2/dir.c

Page 1/3

```

1: /*
2:  * fiwix/fs/ext2/dir.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/dirent.h>
15: #include <fiwix/mm.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: struct fs_operations ext2_dir_fsop = {
20:     0,
21:     0,
22:
23:     ext2_dir_open,
24:     ext2_dir_close,
25:     ext2_dir_read,
26:     ext2_dir_write,
27:     NULL, /* ioctl */
28:     NULL, /* lseek */
29:     ext2_dir_readdir,
30:     NULL, /* mmap */
31:     NULL, /* select */
32:
33:     NULL, /* readlink */
34:     NULL, /* followlink */
35:     ext2_bmap,
36:     ext2_lookup,
37:     ext2_rmdir,
38:     ext2_link,
39:     ext2_unlink,
40:     ext2_symlink,
41:     ext2_mkdir,
42:     ext2_mknod,
43:     NULL, /* truncate */
44:     ext2_create,
45:     ext2_rename,
46:
47:     NULL, /* read_block */
48:     NULL, /* write_block */
49:
50:     NULL, /* read_inode */
51:     NULL, /* write_inode */
52:     NULL, /* ialloc */
53:     NULL, /* ifree */
54:     NULL, /* statfs */
55:     NULL, /* read_superblock */
56:     NULL, /* remount_fs */
57:     NULL, /* write_superblock */
58:     NULL, /* release_superblock */
59: };
60:
61: int ext2_dir_open(struct inode *i, struct fd *fd_table)
62: {
63:     fd_table->offset = 0;
64:     return 0;
65: }
66:
67: int ext2_dir_close(struct inode *i, struct fd *fd_table)

```

## fs/ext2/dir.c

Page 2/3

```

68: {
69:     return 0;
70: }
71:
72: int ext2_dir_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t c
ount)
73: {
74:     return -EISDIR;
75: }
76:
77: int ext2_dir_write(struct inode *i, struct fd *fd_table, const char *buffer, __s
ize_t count)
78: {
79:     return -EBADF;
80: }
81:
82: int ext2_dir_readdir(struct inode *i, struct fd *fd_table, struct dirent *dirent
, unsigned int count)
83: {
84:     __blk_t block;
85:     unsigned int doffset, offset;
86:     unsigned int size, dirent_len;
87:     struct ext2_dir_entry_2 *d;
88:     int base_dirent_len;
89:     int blksize;
90:     struct buffer *buf;
91:
92:     if(!(S_ISDIR(i->i_mode))) {
93:         return -EBADF;
94:     }
95:
96:     blksize = i->sb->s_blocksize;
97:     if(fd_table->offset > i->i_size) {
98:         fd_table->offset = i->i_size;
99:     }
100:
101:     base_dirent_len = sizeof(dirent->d_ino) + sizeof(dirent->d_off) + sizeof
(dirent->d_reclen);
102:     doffset = offset = size = 0;
103:
104:     while(doffset < count) {
105:         if((block = bmap(i, fd_table->offset, FOR_READING)) < 0) {
106:             return block;
107:         }
108:         if(block) {
109:             if(!(buf = bread(i->dev, block, blksize))) {
110:                 return -EIO;
111:             }
112:
113:             doffset = fd_table->offset;
114:             offset = fd_table->offset % blksize;
115:             while(doffset < i->i_size && offset < blksize) {
116:                 d = (struct ext2_dir_entry_2 *) (buf->data + offs
et);
117:                 if(d->inode) {
118:                     dirent_len = (base_dirent_len + (d->name
_len + 1)) + 3;
119:                     dirent_len &= ~3;          /* round up */
120:                     dirent->d_ino = d->inode;
121:                     if((size + dirent_len) < count) {
122:                         dirent->d_off = doffset;
123:                         dirent->d_reclen = dirent_len;
124:                         memcpy_b(dirent->d_name, d->name
, d->name_len);
125:                         dirent->d_name[d->name_len] = NU
LL;
126:                         dirent = (struct dirent *) ((char

```

```
*)dirent + dirent_len);
127:                                     size += dirent_len;
128:                                     } else {
129:                                         break;
130:                                     }
131:                                     }
132:                                     doffset += d->rec_len;
133:                                     offset += d->rec_len;
134:                                     if(!d->rec_len) {
135:                                         break;
136:                                     }
137:                                     }
138:                                     brelse(buf);
139:                                     }
140:                                     fd_table->offset &= ~(blksize - 1);
141:                                     doffset = fd_table->offset;
142:                                     fd_table->offset += offset;
143:                                     doffset += blksize;
144:                                     }
145:
146:                                     return size;
147: }
```

## fs/ext2/file.c

Page 1/3

```

1: /*
2:  * fiwix/fs/ext2/file.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/buffer.h>
12: #include <fiwix/fs.h>
13: #include <fiwix/filesystems.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/mman.h>
16: #include <fiwix/fcntl.h>
17: #include <fiwix/stdio.h>
18: #include <fiwix/string.h>
19:
20: struct fs_operations ext2_file_fsop = {
21:     0,
22:     0,
23:
24:     ext2_file_open,
25:     ext2_file_close,
26:     file_read,
27:     ext2_file_write,
28:     NULL,                /* ioctl */
29:     ext2_file_lseek,
30:     NULL,                /* readdir */
31:     NULL,                /* mmap */
32:     NULL,                /* select */
33:
34:     NULL,                /* readlink */
35:     NULL,                /* followlink */
36:     ext2_bmap,
37:     NULL,                /* lookup */
38:     NULL,                /* rmdir */
39:     NULL,                /* link */
40:     NULL,                /* unlink */
41:     NULL,                /* symlink */
42:     NULL,                /* mkdir */
43:     NULL,                /* mknod */
44:     ext2_truncate,
45:     NULL,                /* create */
46:     NULL,                /* rename */
47:
48:     NULL,                /* read_block */
49:     NULL,                /* write_block */
50:
51:     NULL,                /* read_inode */
52:     NULL,                /* write_inode */
53:     NULL,                /* ialloc */
54:     NULL,                /* ifree */
55:     NULL,                /* statfs */
56:     NULL,                /* read_superblock */
57:     NULL,                /* remount_fs */
58:     NULL,                /* write_superblock */
59:     NULL,                /* release_superblock */
60: };
61:
62: int ext2_file_open(struct inode *i, struct fd *fd_table)
63: {
64:     if(fd_table->flags & O_APPEND) {
65:         fd_table->offset = i->i_size;
66:     } else {
67:         fd_table->offset = 0;

```

## fs/ext2/file.c

Page 2/3

```

68:         }
69:         if(fd_table->flags & O_TRUNC) {
70:             i->i_size = 0;
71:             ext2_truncate(i, 0);
72:         }
73:         return 0;
74:     }
75:
76: int ext2_file_close(struct inode *i, struct fd *fd_table)
77: {
78:     return 0;
79: }
80:
81: int ext2_file_write(struct inode *i, struct fd *fd_table, const char *buffer, __
size_t count)
82: {
83:     __blk_t block;
84:     __off_t total_written;
85:     unsigned int boffset, bytes;
86:     int blksize;
87:     struct buffer *buf;
88:
89:     inode_lock(i);
90:
91:     blksize = i->sb->s_blocksize;
92:     total_written = 0;
93:
94:     if(fd_table->flags & O_APPEND) {
95:         fd_table->offset = i->i_size;
96:     }
97:
98:     while(total_written < count) {
99:         boffset = fd_table->offset % blksize;
100:        if((block = bmap(i, fd_table->offset, FOR_WRITING)) < 0) {
101:            inode_unlock(i);
102:            return block;
103:        }
104:        bytes = blksize - boffset;
105:        bytes = MIN(bytes, (count - total_written));
106:        if(!(buf = bread(i->dev, block, blksize))) {
107:            inode_unlock(i);
108:            return -EIO;
109:        }
110:        memcpy_b(buf->data + boffset, buffer + total_written, bytes);
111:        update_page_cache(i, fd_table->offset, buffer + total_written, b
ytes);
112:        bwrite(buf);
113:        total_written += bytes;
114:        boffset += bytes;
115:        boffset %= blksize;
116:        fd_table->offset += bytes;
117:    }
118:
119:    if(fd_table->offset > i->i_size) {
120:        i->i_size = fd_table->offset;
121:    }
122:    i->i_ctime = CURRENT_TIME;
123:    i->i_mtime = CURRENT_TIME;
124:    i->dirty = 1;
125:
126:    inode_unlock(i);
127:    return total_written;
128: }
129:
130: int ext2_file_lseek(struct inode *i, __off_t offset)
131: {
132:     return offset;

```

```
133: }
```



## fs/ext2/inode.c

Page 1/8

```

1: /*
2:  * fiwix/fs/ext2/inode.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/fs_ext2.h>
12: #include <fiwix/fs_pipe.h>
13: #include <fiwix/statfs.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/stat.h>
16: #include <fiwix/sched.h>
17: #include <fiwix/buffer.h>
18: #include <fiwix/mm.h>
19: #include <fiwix/process.h>
20: #include <fiwix/errno.h>
21: #include <fiwix/stdio.h>
22: #include <fiwix/string.h>
23:
24: #define BLOCKS_PER_IND_BLOCK(sb)          (EXT2_BLOCK_SIZE(sb) / sizeof(unsigned i
nt))
25: #define BLOCKS_PER_DIND_BLOCK(sb)        (BLOCKS_PER_IND_BLOCK(sb) * BLOCKS_PER_I
ND_BLOCK(sb))
26: #define BLOCKS_PER_TIND_BLOCK(sb)        (BLOCKS_PER_IND_BLOCK(sb) * BLOCKS_PER_I
ND_BLOCK(sb) * BLOCKS_PER_IND_BLOCK(sb))
27:
28: #define EXT2_INODES_PER_BLOCK(sb)        (EXT2_BLOCK_SIZE(sb) / sizeof(struct ext
2_inode))
29:
30: static void free_indblock(struct inode *i, int block, int offset)
31: {
32:     int n;
33:     struct buffer *buf;
34:     __blk_t *indblock;
35:
36:     if(!(buf = bread(i->dev, block, i->sb->s_blocksize))) {
37:         printk("WARNING: %s(): error reading block %d.\n", __FUNCTION__,
block);
38:         return;
39:     }
40:     indblock = (__blk_t *)buf->data;
41:     for(n = offset; n < BLOCKS_PER_IND_BLOCK(i->sb); n++) {
42:         if(indblock[n]) {
43:             ext2_bfree(i->sb, indblock[n]);
44:             indblock[n] = 0;
45:             i->i_blocks -= i->sb->s_blocksize / 512;
46:         }
47:     }
48:     bwrite(buf);
49: }
50:
51: static int get_group_desc(struct superblock *sb, __blk_t block_group, struct ext
2_group_desc *gd)
52: {
53:     __blk_t group_desc_block;
54:     int group_desc;
55:     struct buffer *buf;
56:
57:     group_desc_block = block_group / EXT2_DESC_PER_BLOCK(sb);
58:     group_desc = block_group % EXT2_DESC_PER_BLOCK(sb);
59:     if(!(buf = bread(sb->dev, SUPERBLOCK + sb->u.ext2.sb.s_first_data_block
+ group_desc_block, sb->s_blocksize))) {
60:         return -EIO;

```

## fs/ext2/inode.c

Page 2/8

```

61:         }
62:         memcpy_b(gd, (void *) (buf->data + (group_desc * sizeof(struct ext2_group
_desc))), sizeof(struct ext2_group_desc));
63:         brelse(buf);
64:         return 0;
65:     }
66:
67: int ext2_read_inode(struct inode *i)
68: {
69:     __blk_t block_group, block;
70:     unsigned int offset;
71:     struct superblock *sb;
72:     struct ext2_inode *ii;
73:     struct ext2_group_desc gd;
74:     struct buffer *buf;
75:
76:     if(!(sb = get_superblock(i->dev))) {
77:         printk("WARNING: %s(): get_superblock() has returned NULL.\n");
78:         return -EINVAL;
79:     }
80:     block_group = ((i->inode - 1) / EXT2_INODES_PER_GROUP(sb));
81:     if(get_group_desc(sb, block_group, &gd) {
82:         return -EIO;
83:     }
84:     block = (((i->inode - 1) % EXT2_INODES_PER_GROUP(sb)) / EXT2_INODES_PER_
BLOCK(sb));
85:
86:     if(!(buf = bread(i->dev, gd.bg_inode_table + block, i->sb->s_blocksize))
) {
87:         return -EIO;
88:     }
89:     offset = (((i->inode - 1) % EXT2_INODES_PER_GROUP(sb)) % EXT2_INODES_PE
R_BLOCK(sb)) * sizeof(struct ext2_inode));
90:
91:     ii = (struct ext2_inode *) (buf->data + offset);
92:     memcpy_b(&i->u.ext2.i_data, ii->i_block, sizeof(ii->i_block));
93:
94:     i->i_mode = ii->i_mode;
95:     i->i_uid = ii->i_uid;
96:     i->i_size = ii->i_size;
97:     i->i_atime = ii->i_atime;
98:     i->i_ctime = ii->i_ctime;
99:     i->i_mtime = ii->i_mtime;
100:    i->i_gid = ii->i_gid;
101:    i->i_nlink = ii->i_links_count;
102:    i->i_blocks = ii->i_blocks;
103:    i->i_flags = ii->i_flags;
104:    i->count = 1;
105:    switch(i->i_mode & S_IFMT) {
106:        case S_IFCHR:
107:            i->fsop = &def_chr_fsop;
108:            i->rdev = ii->i_block[0];
109:            break;
110:        case S_IFBLK:
111:            i->fsop = &def_blk_fsop;
112:            i->rdev = ii->i_block[0];
113:            break;
114:        case S_IFIFO:
115:            i->fsop = &pipefs_fsop;
116:            /* it's a union so we need to clear pipefs_i */
117:            memset_b(&i->u.pipefs, NULL, sizeof(struct pipefs_inode)
);
118:            break;
119:        case S_IFDIR:
120:            i->fsop = &ext2_dir_fsop;
121:            break;
122:        case S_IFREG:

```

## fs/ext2/inode.c

Page 3/8

```

123:             i->fsop = &ext2_file_fsop;
124:             break;
125:         case S_IFLNK:
126:             i->fsop = &ext2_symlink_fsop;
127:             break;
128:         case S_IFSOCK:
129:             i->fsop = NULL;
130:             break;
131:         default:
132:             printk("WARNING: %s(): invalid inode (%d) mode %08o.\n",
__FUNCTION__, i->inode, i->i_mode);
133:             brelse(buf);
134:             return -ENOENT;
135:     }
136:     brelse(buf);
137:     return 0;
138: }
139:
140: int ext2_write_inode(struct inode *i)
141: {
142:     __blk_t block_group, block;
143:     short int offset;
144:     struct superblock *sb;
145:     struct ext2_inode *ii;
146:     struct ext2_group_desc gd;
147:     struct buffer *buf;
148:
149:     if(!(sb = get_superblock(i->dev))) {
150:         printk("WARNING: %s(): get_superblock() has returned NULL.\n");
151:         return -EINVAL;
152:     }
153:     block_group = ((i->inode - 1) / EXT2_INODES_PER_GROUP(sb));
154:     if(get_group_desc(sb, block_group, &gd) {
155:         return -EIO;
156:     }
157:     block = (((i->inode - 1) % EXT2_INODES_PER_GROUP(sb)) / EXT2_INODES_PER_
BLOCK(sb));
158:
159:     if(!(buf = bread(i->dev, gd.bg_inode_table + block, i->sb->s_blocksize))
) {
160:         return -EIO;
161:     }
162:     offset = (((i->inode - 1) % EXT2_INODES_PER_GROUP(sb)) % EXT2_INODES_PE
R_BLOCK(sb)) * sizeof(struct ext2_inode);
163:     ii = (struct ext2_inode *) (buf->data + offset);
164:     memset_b(ii, 0, sizeof(struct ext2_inode));
165:
166:     ii->i_mode = i->i_mode;
167:     ii->i_uid = i->i_uid;
168:     ii->i_size = i->i_size;
169:     ii->i_atime = i->i_atime;
170:     ii->i_ctime = i->i_ctime;
171:     ii->i_mtime = i->i_mtime;
172:     ii->i_dtime = i->u.ext2.i_dtime;
173:     ii->i_gid = i->i_gid;
174:     ii->i_links_count = i->i_nlink;
175:     ii->i_blocks = i->i_blocks;
176:     ii->i_flags = i->i_flags;
177:     if(S_ISCHR(i->i_mode) || S_ISBLK(i->i_mode)) {
178:         ii->i_block[0] = i->rdev;
179:     } else {
180:         memcpy_b(ii->i_block, &i->u.ext2.i_data, sizeof(i->u.ext2.i_data
));
181:     }
182:     i->dirty = 0;
183:     bwrite(buf);
184:     return 0;

```

## fs/ext2/inode.c

Page 4/8

```

185: }
186:
187: int ext2_bmap(struct inode *i, __off_t offset, int mode)
188: {
189:     unsigned char level;
190:     __blk_t *indblock, *dindblock, *tindblock;
191:     __blk_t block, iblock, dblock, tblock, newblock;
192:     int blksize;
193:     struct buffer *buf, *buf2, *buf3, *buf4;
194:
195:     blksize = i->sb->s_blocksize;
196:     block = offset / blksize;
197:     level = 0;
198:     buf3 = NULL;    /* makes GCC happy */
199:
200:     if(block < EXT2_NDIR_BLOCKS) {
201:         level = EXT2_NDIR_BLOCKS - 1;
202:     } else {
203:         if(block < (BLOCKS_PER_IND_BLOCK(i->sb) + EXT2_NDIR_BLOCKS)) {
204:             level = EXT2_IND_BLOCK;
205:         } else if(block < ((BLOCKS_PER_IND_BLOCK(i->sb) * BLOCKS_PER_IND
_BLOCK(i->sb)) + BLOCKS_PER_IND_BLOCK(i->sb) + EXT2_NDIR_BLOCKS)) {
206:             level = EXT2_DIND_BLOCK;
207:         } else {
208:             level = EXT2_TIND_BLOCK;
209:         }
210:         block -= EXT2_NDIR_BLOCKS;
211:     }
212:
213:     if(level < EXT2_NDIR_BLOCKS) {
214:         if(!i->u.ext2.i_data[block] && mode == FOR_WRITING) {
215:             if((newblock = ext2_balloc(i->sb)) < 0) {
216:                 return -ENOSPC;
217:             }
218:             /* initialize the new block */
219:             if(!(buf = bread(i->dev, newblock, blksize))) {
220:                 ext2_bfree(i->sb, newblock);
221:                 return -EIO;
222:             }
223:             memset_b(buf->data, 0, blksize);
224:             bwrite(buf);
225:             i->u.ext2.i_data[block] = newblock;
226:             i->i_blocks += blksize / 512;
227:         }
228:         return i->u.ext2.i_data[block];
229:     }
230:
231:     if(!i->u.ext2.i_data[level]) {
232:         if(mode == FOR_WRITING) {
233:             if((newblock = ext2_balloc(i->sb)) < 0) {
234:                 return -ENOSPC;
235:             }
236:             /* initialize the new block */
237:             if(!(buf = bread(i->dev, newblock, blksize))) {
238:                 ext2_bfree(i->sb, newblock);
239:                 return -EIO;
240:             }
241:             memset_b(buf->data, 0, blksize);
242:             bwrite(buf);
243:             i->u.ext2.i_data[level] = newblock;
244:             i->i_blocks += blksize / 512;
245:         } else {
246:             return 0;
247:         }
248:     }
249:     if(!(buf = bread(i->dev, i->u.ext2.i_data[level], blksize))) {
250:         return -EIO;

```

## fs/ext2/inode.c

Page 5/8

```

251:         }
252:         indblock = (__blk_t *)buf->data;
253:         dblock = block - BLOCKS_PER_IND_BLOCK(i->sb);
254:         tblock = block - (BLOCKS_PER_IND_BLOCK(i->sb) * BLOCKS_PER_IND_BLOCK(i->
sb)) - BLOCKS_PER_IND_BLOCK(i->sb);
255:
256:         if(level == EXT2_DIND_BLOCK) {
257:             block = dblock / BLOCKS_PER_IND_BLOCK(i->sb);
258:         }
259:         if(level == EXT2_TIND_BLOCK) {
260:             block = tblock / (BLOCKS_PER_IND_BLOCK(i->sb) * BLOCKS_PER_IND_B
LOCK(i->sb));
261:         }
262:
263:         if(!indblock[block]) {
264:             if(mode == FOR_WRITING) {
265:                 if((newblock = ext2_balloc(i->sb)) < 0) {
266:                     brelse(buf);
267:                     return -ENOSPC;
268:                 }
269:                 /* initialize the new block */
270:                 if(!(buf2 = bread(i->dev, newblock, blksize))) {
271:                     ext2_bfree(i->sb, newblock);
272:                     brelse(buf);
273:                     return -EIO;
274:                 }
275:                 memset_b(buf2->data, 0, blksize);
276:                 bwrite(buf2);
277:                 indblock[block] = newblock;
278:                 i->i_blocks += blksize / 512;
279:                 if(level == EXT2_IND_BLOCK) {
280:                     bwrite(buf);
281:                     return newblock;
282:                 }
283:                 buf->dirty = 1;
284:                 buf->valid = 1;
285:             } else {
286:                 brelse(buf);
287:                 return 0;
288:             }
289:         }
290:         if(level == EXT2_IND_BLOCK) {
291:             newblock = indblock[block];
292:             brelse(buf);
293:             return newblock;
294:         }
295:
296:         if(level == EXT2_TIND_BLOCK) {
297:             if(!(buf3 = bread(i->dev, indblock[block], blksize))) {
298:                 printk("%s(): returning -EIO\n", __FUNCTION__);
299:                 brelse(buf);
300:                 return -EIO;
301:             }
302:             tindblock = (__blk_t *)buf3->data;
303:             block = tindblock[tblock / BLOCKS_PER_IND_BLOCK(i->sb)];
304:             if(!block) {
305:                 if(mode == FOR_WRITING) {
306:                     if((newblock = ext2_balloc(i->sb)) < 0) {
307:                         brelse(buf);
308:                         brelse(buf3);
309:                         return -ENOSPC;
310:                     }
311:                     /* initialize the new block */
312:                     if(!(buf4 = bread(i->dev, newblock, blksize))) {
313:                         ext2_bfree(i->sb, newblock);
314:                         brelse(buf);
315:                         brelse(buf3);

```

## fs/ext2/inode.c

Page 6/8

```

316:                                     return -EIO;
317:                                     }
318:                                     memset_b(buf4->data, 0, blksize);
319:                                     bwrite(buf4);
320:                                     tindblock[tblock / BLOCKS_PER_IND_BLOCK(i->sb)]
= newblock;
321:                                     i->i_blocks += blksize / 512;
322:                                     buf3->dirty = 1;
323:                                     buf3->valid = 1;
324:                                     block = newblock;
325:                                     } else {
326:                                     brelse(buf);
327:                                     brelse(buf3);
328:                                     return 0;
329:                                     }
330:                                     }
331:                                     dblock = tblock;
332:                                     iblock = tblock / BLOCKS_PER_IND_BLOCK(i->sb);
333:                                     if(!(buf2 = bread(i->dev, block, blksize))) {
334:                                     printk("%s(): returning -EIO\n", __FUNCTION__);
335:                                     brelse(buf);
336:                                     brelse(buf3);
337:                                     return -EIO;
338:                                     }
339:                                     } else {
340:                                     iblock = block;
341:                                     if(!(buf2 = bread(i->dev, indblock[iblock], blksize))) {
342:                                     printk("%s(): returning -EIO\n", __FUNCTION__);
343:                                     brelse(buf);
344:                                     return -EIO;
345:                                     }
346:                                     }
347:
348:                                     dindblock = (__blk_t *)buf2->data;
349:                                     block = dindblock[dblock - (iblock * BLOCKS_PER_IND_BLOCK(i->sb))];
350:                                     if(!block && mode == FOR_WRITING) {
351:                                     if((newblock = ext2_balloc(i->sb)) < 0) {
352:                                     brelse(buf);
353:                                     if(level == EXT2_TIND_BLOCK) {
354:                                     brelse(buf3);
355:                                     }
356:                                     brelse(buf2);
357:                                     return -ENOSPC;
358:                                     }
359:                                     /* initialize the new block */
360:                                     if(!(buf4 = bread(i->dev, newblock, blksize))) {
361:                                     ext2_bfree(i->sb, newblock);
362:                                     brelse(buf);
363:                                     if(level == EXT2_TIND_BLOCK) {
364:                                     brelse(buf3);
365:                                     }
366:                                     brelse(buf2);
367:                                     return -EIO;
368:                                     }
369:                                     memset_b(buf4->data, 0, blksize);
370:                                     bwrite(buf4);
371:                                     dindblock[dblock - (iblock * BLOCKS_PER_IND_BLOCK(i->sb))] = new
block;
372:                                     i->i_blocks += blksize / 512;
373:                                     buf2->dirty = 1;
374:                                     buf2->valid = 1;
375:                                     block = newblock;
376:                                     }
377:                                     brelse(buf);
378:                                     if(level == EXT2_TIND_BLOCK) {
379:                                     brelse(buf3);
380:                                     }

```

## fs/ext2/inode.c

Page 7/8

```

381:         brelse(buf2);
382:         return block;
383:     }
384:
385: int ext2_truncate(struct inode *i, __off_t length)
386: {
387:     __blk_t block, dblock, *indblock;
388:     struct buffer *buf;
389:     int blksize, n;
390:
391:     blksize = i->sb->s_blocksize;
392:     block = length / blksize;
393:
394:     if(!S_ISDIR(i->i_mode) && !S_ISREG(i->i_mode) && !S_ISLNK(i->i_mode)) {
395:         return -EINVAL;
396:     }
397:
398:     if(block < EXT2_NDIR_BLOCKS) {
399:         for(n = block; n < EXT2_NDIR_BLOCKS; n++) {
400:             if(i->u.ext2.i_data[n]) {
401:                 ext2_bfree(i->sb, i->u.ext2.i_data[n]);
402:                 i->u.ext2.i_data[n] = 0;
403:                 i->i_blocks -= blksize / 512;
404:             }
405:         }
406:         block = 0;
407:     }
408:
409:     if(!block || block < (BLOCKS_PER_IND_BLOCK(i->sb) + EXT2_NDIR_BLOCKS)) {
410:         if(block) {
411:             block -= EXT2_NDIR_BLOCKS;
412:         }
413:         if(i->u.ext2.i_data[EXT2_IND_BLOCK]) {
414:             free_indblock(i, i->u.ext2.i_data[EXT2_IND_BLOCK], block
);
415:             if(!block) {
416:                 ext2_bfree(i->sb, i->u.ext2.i_data[EXT2_IND_BLOCK
K]);
417:                 i->u.ext2.i_data[EXT2_IND_BLOCK] = 0;
418:                 i->i_blocks -= blksize / 512;
419:             }
420:         }
421:         block = 0;
422:     }
423:
424:     if(block) {
425:         block -= EXT2_NDIR_BLOCKS;
426:         block -= BLOCKS_PER_IND_BLOCK(i->sb);
427:     }
428:     if(i->u.ext2.i_data[EXT2_DIND_BLOCK]) {
429:         if(!(buf = bread(i->dev, i->u.ext2.i_data[EXT2_DIND_BLOCK], blksize))) {
430:             printk("%s(): error reading block %d.\n", __FUNCTION__,
i->u.ext2.i_data[EXT2_DIND_BLOCK]);
431:             return -EIO;
432:         }
433:         indblock = (__blk_t *)buf->data;
434:         dblock = block % BLOCKS_PER_IND_BLOCK(i->sb);
435:         for(n = block / BLOCKS_PER_IND_BLOCK(i->sb); n < BLOCKS_PER_IND_
BLOCK(i->sb); n++) {
436:             if(indblock[n]) {
437:                 free_indblock(i, indblock[n], dblock);
438:                 if(!dblock) {
439:                     ext2_bfree(i->sb, indblock[n]);
440:                     i->i_blocks -= blksize / 512;
441:                 }
442:             }

```

```
443:             dblock = 0;
444:         }
445:         bwrite(buf);
446:         if(!block) {
447:             ext2_bfree(i->sb, i->u.ext2.i_data[EXT2_DIND_BLOCK]);
448:             i->u.ext2.i_data[EXT2_DIND_BLOCK] = 0;
449:             i->i_blocks -= blksize / 512;
450:         }
451:     }
452:
453:     i->i_mtime = CURRENT_TIME;
454:     i->i_ctime = CURRENT_TIME;
455:     i->i_size = length;
456:     i->dirty = 1;
457:
458:     return 0;
459: }
```



**fs/ext2/Makefile**

Page 1/1

```
1: # fiwix/fs/ext2/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6: #
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = inode.o super.o namei.o symlink.o dir.o file.o bitmaps.o
13:
14: ext2:     $(OBJS)
15:             $(LD) $(LDFLAGS) -r $(OBJS) -o ext2.o
16:
17: clean:
18:             rm -f *.o
19:
```

## fs/ext2/namei.c

Page 1/12

```

1: /*
2:  * fiwix/fs/ext2/namei.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_ext2.h>
13: #include <fiwix/buffer.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/errno.h>
16: #include <fiwix/stat.h>
17: #include <fiwix/stdio.h>
18: #include <fiwix/string.h>
19:
20: /* finds the entry 'name' with (optionally) inode 'i' in the directory 'dir' */
21: static struct buffer * find_dir_entry(struct inode *dir, struct inode *i, struct
ext2_dir_entry_2 **d_res, char *name)
22: {
23:     __blk_t block;
24:     unsigned int blksize;
25:     unsigned int offset, doffset;
26:     struct buffer *buf;
27:     int basesize, rlen, nlen;
28:
29:     basesize = sizeof((*d_res)->inode) + sizeof((*d_res)->rec_len) + sizeof(
(*d_res)->name_len) + sizeof((*d_res)->file_type);
30:     blksize = dir->sb->s_blocksize;
31:     offset = 0;
32:
33:     /* nlen is the length of the new entry, used when searching for the first
usable entry */
34:     nlen = basesize + strlen(name) + 3;
35:     nlen &= ~3;
36:
37:     while(offset < dir->i_size) {
38:         if((block = bmap(dir, offset, FOR_READING)) < 0) {
39:             break;
40:         }
41:         if(block) {
42:             if(!(buf = bread(dir->dev, block, blksize))) {
43:                 break;
44:             }
45:             doffset = 0;
46:             do {
47:                 *d_res = (struct ext2_dir_entry_2 *) (buf->data +
doffset);
48:                 if(!i) {
49:                     /* calculates the real length of the current
entry */
50:                     rlen = basesize + strlen((*d_res)->name)
+ 3;
51:                     rlen &= ~3;
52:                     /* returns the first entry where *name can
fit in */
53:                     if(!(*d_res)->inode) {
54:                         if(nlen <= (*d_res)->rec_len) {
55:                             return buf;
56:                         }
57:                     } else {
58:                         if(rlen + nlen <= (*d_res)->rec_
len) {
59:                             return buf;

```

## fs/ext2/namei.c

Page 2/12

```

60:
61:                                nrec_len = (*d_res)->rec
_len - rlen;                                (*d_res)->rec_len = rlen
62:                                ;
63:                                doffset += (*d_res)->rec
_len;                                *d_res = (struct ext2_di
64:                                r_entry_2 *) (buf->data + doffset);
65:                                (*d_res)->rec_len = nrec
_len;                                return buf;
66:                                }
67:                                }
68:                                doffset += (*d_res)->rec_len;
69:                                } else {
70:                                if((*d_res)->inode == i->inode) {
71:                                /* returns the first matching in
72:                                ode */
73:                                if(!name) {
74:                                return buf;
75:                                }
76:                                /* returns the matching inode an
77:                                d name */
78:                                if((*d_res)->name_len == strlen(
name)) {
79:                                if(!strncmp((*d_res)->na
me, name, (*d_res)->name_len)) {
80:                                return buf;
81:                                }
82:                                }
83:                                doffset += (*d_res)->rec_len;
84:                                }
85:                                } while(doffset < blksize);
86:                                brelse(buf);
87:                                offset += blksize;
88:                                } else {
89:                                break;
90:                                }
91:                                }
92:                                }
93:                                *d_res = NULL;
94:                                return NULL;
95:                                }
96:                                }
97: static struct buffer * add_dir_entry(struct inode *dir, struct ext2_dir_entry_2
**d_res, char *name)
98: {
99:     __blk_t block;
100:    struct buffer *buf;
101:
102:    if(!(buf = find_dir_entry(dir, NULL, d_res, name))) {
103:        if((block = bmap(dir, dir->i_size, FOR_WRITING)) < 0) {
104:            return NULL;
105:        }
106:        if(!(buf = bread(dir->dev, block, dir->sb->s_blocksize))) {
107:            return NULL;
108:        }
109:        *d_res = (struct ext2_dir_entry_2 *)buf->data;
110:        dir->i_size += dir->sb->s_blocksize;
111:        (*d_res)->rec_len = dir->sb->s_blocksize;
112:    }
113:
114:    return buf;
115: }
116:

```

## fs/ext2/namei.c

Page 3/12

```

117: static int is_dir_empty(struct inode *dir)
118: {
119:     __blk_t block;
120:     unsigned int blksize;
121:     unsigned int offset, doffset;
122:     struct buffer *buf;
123:     struct ext2_dir_entry_2 *d;
124:
125:     blksize = dir->sb->s_blocksize;
126:     offset = 0;
127:
128:     while(offset < dir->i_size) {
129:         if((block = bmap(dir, offset, FOR_READING)) < 0) {
130:             break;
131:         }
132:         if(block) {
133:             if(!(buf = bread(dir->dev, block, blksize))) {
134:                 break;
135:             }
136:             doffset = 0;
137:             do {
138:                 if(doffset + offset >= dir->i_size) {
139:                     break;
140:                 }
141:                 d = (struct ext2_dir_entry_2 *) (buf->data + doff
set);
142:                 doffset += d->rec_len;
143:                 if(d->inode && d->name_len == 1 && d->name[0] ==
'.') {
144:                     continue;
145:                 }
146:                 if(d->inode && d->name_len == 2 && d->name[0] ==
'.' && d->name[1] == '.') {
147:                     continue;
148:                 }
149:                 if(d->inode) {
150:                     brelse(buf);
151:                     return 0;
152:                 }
153:             } while(doffset < blksize);
154:             brelse(buf);
155:             offset += blksize;
156:         } else {
157:             break;
158:         }
159:     }
160:
161:     return 1;
162: }
163:
164: static int is_prefix(struct inode *dir_new, struct inode *i_old)
165: {
166:     __ino_t inode;
167:     int errno;
168:
169:     errno = 0;
170:     for(;;) {
171:         if(dir_new == i_old) {
172:             errno = 1;
173:             break;
174:         }
175:         inode = dir_new->inode;
176:         if(ext2_lookup("..", dir_new, &dir_new)) {
177:             break;
178:         }
179:         iput(dir_new); /* lookup eats 1 dir_new */
180:         if(dir_new->inode == inode) {

```

## fs/ext2/namei.c

Page 4/12

```

181:             break;
182:         }
183:     }
184:     return errno;
185: }
186:
187: int ext2_lookup(const char *name, struct inode *dir, struct inode **i_res)
188: {
189:     __blk_t block;
190:     unsigned int blksize;
191:     unsigned int offset, doffset;
192:     struct buffer *buf;
193:     struct ext2_dir_entry_2 *d;
194:     __ino_t inode;
195:
196:     blksize = dir->sb->s_blocksize;
197:     inode = offset = 0;
198:     dir->count++;
199:
200:     while(offset < dir->i_size && !inode) {
201:         if((block = bmap(dir, offset, FOR_READING)) < 0) {
202:             return block;
203:         }
204:         if(block) {
205:             if(!(buf = bread(dir->dev, block, blksize))) {
206:                 return -EIO;
207:             }
208:             doffset = 0;
209:             do {
210:                 d = (struct ext2_dir_entry_2 *) (buf->data + doff
set);
211:                 if(d->inode) {
212:                     if(d->name_len == strlen(name)) {
213:                         if(strncmp(d->name, name, d->nam
e_len) == 0) {
214:                             inode = d->inode;
215:                         }
216:                     }
217:                 }
218:                 doffset += d->rec_len;
219:             } while((doffset < blksize) && (!inode));
220:
221:             brelse(buf);
222:             offset += blksize;
223:             if(inode) {
224:                 /*
225:                  * This prevents a deadlock in iget() when
226:                  * trying to lock '.' when 'dir' is the same
227:                  * directory (ls -lai <dir>).
228:                  */
229:                 if(inode == dir->inode) {
230:                     *i_res = dir;
231:                     return 0;
232:                 }
233:
234:                 if(!(*i_res = iget(dir->sb, inode))) {
235:                     return -EACCES;
236:                 }
237:                 iput(dir);
238:                 return 0;
239:             }
240:         } else {
241:             break;
242:         }
243:     }
244:     iput(dir);
245:     return -ENOENT;

```

## fs/ext2/namei.c

Page 5/12

```
246: }
247:
248: int ext2_rmdir(struct inode *dir, struct inode *i)
249: {
250:     struct buffer *buf;
251:     struct ext2_dir_entry_2 *d;
252:
253:     inode_lock(i);
254:
255:     if(!is_dir_empty(i)) {
256:         inode_unlock(i);
257:         return -ENOTEMPTY;
258:     }
259:
260:     inode_lock(dir);
261:
262:     if(!(buf = find_dir_entry(dir, i, &d, NULL))) {
263:         inode_unlock(i);
264:         inode_unlock(dir);
265:         return -ENOENT;
266:     }
267:
268:     d->inode = 0;
269:     i->i_nlink = 0;
270:     dir->i_nlink--;
271:
272:     i->i_ctime = CURRENT_TIME;
273:     i->u.ext2.i_dtime = CURRENT_TIME;
274:     dir->i_mtime = CURRENT_TIME;
275:     dir->i_ctime = CURRENT_TIME;
276:
277:     i->dirty = 1;
278:     dir->dirty = 1;
279:
280:     bwrite(buf);
281:
282:     inode_unlock(i);
283:     inode_unlock(dir);
284:     return 0;
285: }
286:
287: int ext2_link(struct inode *i_old, struct inode *dir_new, char *name)
288: {
289:     struct buffer *buf;
290:     struct ext2_dir_entry_2 *d;
291:     char c;
292:     int n;
293:
294:     inode_lock(i_old);
295:     inode_lock(dir_new);
296:
297:     if(!(buf = add_dir_entry(dir_new, &d, name))) {
298:         inode_unlock(i_old);
299:         inode_unlock(dir_new);
300:         return -ENOSPC;
301:     }
302:
303:     d->inode = i_old->inode;
304:     d->name_len = strlen(name);
305:     /* strcpy() can't be used here because it places a trailing NULL */
306:     for(n = 0; n < NAME_MAX; n++) {
307:         if((c = name[n])) {
308:             d->name[n] = c;
309:             continue;
310:         }
311:         break;
312:     }
```

## fs/ext2/namei.c

Page 6/12

```

313:         d->file_type = 0;          /* not used */
314:
315:         i_old->i_nlink++;
316:         i_old->i_ctime = CURRENT_TIME;
317:         dir_new->i_mtime = CURRENT_TIME;
318:         dir_new->i_ctime = CURRENT_TIME;
319:
320:         i_old->dirty = 1;
321:         dir_new->dirty = 1;
322:
323:         bwrite(buf);
324:
325:         inode_unlock(i_old);
326:         inode_unlock(dir_new);
327:         return 0;
328: }
329:
330: int ext2_unlink(struct inode *dir, struct inode *i, char *name)
331: {
332:     struct buffer *buf;
333:     struct ext2_dir_entry_2 *d;
334:
335:     inode_lock(dir);
336:     inode_lock(i);
337:
338:     if(!(buf = find_dir_entry(dir, i, &d, name))) {
339:         inode_unlock(dir);
340:         inode_unlock(i);
341:         return -ENOENT;
342:     }
343:
344:     /*
345:      * FIXME: in order to avoid low performance when traversing large
346:      * directories plenty of blank entries, it would be interesting
347:      * to merge every removed entry with the previous entry.
348:      */
349:     d->inode = 0;
350:     if(!--i->i_nlink) {
351:         i->u.ext2.i_dtime = CURRENT_TIME;
352:     }
353:
354:     i->i_ctime = CURRENT_TIME;
355:     dir->i_mtime = CURRENT_TIME;
356:     dir->i_ctime = CURRENT_TIME;
357:
358:     i->dirty = 1;
359:     dir->dirty = 1;
360:
361:     bwrite(buf);
362:
363:     inode_unlock(dir);
364:     inode_unlock(i);
365:     return 0;
366: }
367:
368: int ext2_symlink(struct inode *dir, char *name, char *oldname)
369: {
370:     struct buffer *buf, *buf2;
371:     struct inode *i;
372:     struct ext2_dir_entry_2 *d;
373:     __blk_t block;
374:     char c, *data;
375:     int n;
376:
377:     inode_lock(dir);
378:
379:     if(!(i = ialloc(dir->sb, S_IFLNK))) {

```

## fs/ext2/namei.c

Page 7/12

```

380:         inode_unlock(dir);
381:         return -ENOSPC;
382:     }
383:
384:     if(!(buf = add_dir_entry(dir, &d, name))) {
385:         iput(i);
386:         inode_unlock(dir);
387:         return -ENOSPC;
388:     }
389:
390:     i->i_mode = S_IFLNK | (S_IRWXU | S_IRWXG | S_IRWXO);
391:     i->i_uid = current->euid;
392:     i->i_gid = current->egid;
393:     i->dev = dir->dev;
394:     i->count = 1;
395:     i->fsop = &ext2_symlink_fsop;
396:
397:     if(strlen(oldname) >= EXT2_N_BLOCKS * sizeof(__u32)) {
398:         /* this will be a slow symlink */
399:         if((block = ext2_balloc(dir->sb)) < 0) {
400:             iput(i);
401:             brelse(buf);
402:             inode_unlock(dir);
403:             return block;
404:         }
405:         if(!(buf2 = bread(dir->dev, block, dir->sb->s_blocksize))) {
406:             iput(i);
407:             brelse(buf);
408:             ext2_bfree(dir->sb, block);
409:             inode_unlock(dir);
410:             return -EIO;
411:         }
412:         i->u.ext2.i_data[0] = block;
413:         for(n = 0; n < NAME_MAX; n++) {
414:             if((c = oldname[n])) {
415:                 buf2->data[n] = c;
416:                 continue;
417:             }
418:             break;
419:         }
420:         buf2->data[n] = 0;
421:         i->i_blocks = dir->sb->s_blocksize / 512;
422:         bwrite(buf2);
423:     } else {
424:         /* this will be a fast symlink */
425:         data = (char *)i->u.ext2.i_data;
426:         for(n = 0; n < NAME_MAX; n++) {
427:             if((c = oldname[n])) {
428:                 data[n] = c;
429:                 continue;
430:             }
431:             break;
432:         }
433:         data[n] = 0;
434:     }
435:
436:     i->i_size = n;
437:     i->dirty = 1;
438:     i->i_nlink = 1;
439:     d->inode = i->inode;
440:     d->name_len = strlen(name);
441:     /* strcpy() can't be used here because it places a trailing NULL */
442:     for(n = 0; n < NAME_MAX; n++) {
443:         if((c = name[n])) {
444:             d->name[n] = c;
445:             continue;
446:         }

```



## fs/ext2/namei.c

Page 8/12

```

447:             break;
448:         }
449:         d->file_type = 0;          /* EXT2_FT_SYMLINK not used */
450:
451:         dir->i_mtime = CURRENT_TIME;
452:         dir->i_ctime = CURRENT_TIME;
453:         dir->dirty = 1;
454:
455:         bwrite(buf);
456:         iput(i);
457:         inode_unlock(dir);
458:         return 0;
459:     }
460:
461: int ext2_mkdir(struct inode *dir, char *name, __mode_t mode)
462: {
463:     struct buffer *buf, *buf2;
464:     struct inode *i;
465:     struct ext2_dir_entry_2 *d, *d2;
466:     __blk_t block;
467:     char c;
468:     int n;
469:
470:     inode_lock(dir);
471:
472:     if(!(i = ialloc(dir->sb, S_IFDIR))) {
473:         inode_unlock(dir);
474:         return -ENOSPC;
475:     }
476:
477:     i->i_mode = (mode & (S_IRWXU | S_IRWXG | S_IRWXO)) & ~current->umask;
478:     i->i_mode |= S_IFDIR;
479:     i->i_uid = current->euid;
480:     i->i_gid = current->egid;
481:     i->dev = dir->dev;
482:     i->count = 1;
483:     i->fsop = &ext2_dir_fsop;
484:
485:     if((block = bmap(i, 0, FOR_WRITING)) < 0) {
486:         iput(i);
487:         inode_unlock(dir);
488:         return block;
489:     }
490:
491:     if(!(buf2 = bread(i->dev, block, dir->sb->s_blocksize))) {
492:         ext2_bfree(dir->sb, block);
493:         iput(i);
494:         inode_unlock(dir);
495:         return -EIO;
496:     }
497:
498:     if(!(buf = add_dir_entry(dir, &d, name))) {
499:         ext2_bfree(dir->sb, block);
500:         iput(i);
501:         brelse(buf2);
502:         inode_unlock(dir);
503:         return -ENOSPC;
504:     }
505:
506:     d->inode = i->inode;
507:     d->name_len = strlen(name);
508:     /* strcpy() can't be used here because it places a trailing NULL */
509:     for(n = 0; n < NAME_MAX; n++) {
510:         if((c = name[n])) {
511:             if(c != '/') {
512:                 d->name[n] = c;
513:                 continue;

```

## fs/ext2/namei.c

Page 9/12

```

514:         }
515:     }
516:     break;
517: }
518: d->file_type = 0;      /* EXT2_FT_DIR not used */
519:
520: d2 = (struct ext2_dir_entry_2 *)buf2->data;
521: d2->inode = i->inode;
522: d2->name[0] = '.';
523: d2->name[1] = 0;
524: d2->name_len = 1;
525: d2->rec_len = 12;
526: d2->file_type = 0;      /* EXT2_FT_DIR not used */
527: i->i_nlink = 1;
528: d2 = (struct ext2_dir_entry_2 *) (buf2->data + 12);
529: d2->inode = dir->inode;
530: d2->name[0] = '.';
531: d2->name[1] = '.';
532: d2->name[2] = 0;
533: d2->name_len = 2;
534: d2->rec_len = i->sb->s_blocksize - 12;
535: d2->file_type = 0;      /* EXT2_FT_DIR not used */
536: i->i_nlink++;
537: i->i_size = i->sb->s_blocksize;
538: i->i_blocks = dir->sb->s_blocksize / 512;
539: i->dirty = 1;
540:
541: dir->i_mtime = CURRENT_TIME;
542: dir->i_ctime = CURRENT_TIME;
543: dir->i_nlink++;
544: dir->dirty = 1;
545:
546: bwrite(buf);
547: bwrite(buf2);
548: iput(i);
549: inode_unlock(dir);
550: return 0;
551: }
552:
553: int ext2_mknod(struct inode *dir, char *name, __mode_t mode, __dev_t dev)
554: {
555:     struct buffer *buf;
556:     struct inode *i;
557:     struct ext2_dir_entry_2 *d;
558:     char c;
559:     int n;
560:
561:     inode_lock(dir);
562:
563:     if(!(i = ialloc(dir->sb, mode & S_IFMT))) {
564:         inode_unlock(dir);
565:         return -ENOSPC;
566:     }
567:
568:     if(!(buf = add_dir_entry(dir, &d, name))) {
569:         i->i_nlink = 0;
570:         iput(i);
571:         inode_unlock(dir);
572:         return -ENOSPC;
573:     }
574:
575:     d->inode = i->inode;
576:     d->name_len = strlen(name);
577:     /* strcpy() can't be used here because it places a trailing NULL */
578:     for(n = 0; n < NAME_MAX; n++) {
579:         if((c = name[n])) {
580:             d->name[n] = c;

```

## fs/ext2/namei.c

Page 10/12

```

581:             continue;
582:         }
583:         break;
584:     }
585:
586:     i->i_mode = (mode & ~current->umask) & ~S_IFMT;
587:     i->i_uid = current->euid;
588:     i->i_gid = current->egid;
589:     i->i_nlink = 1;
590:     i->dev = dir->dev;
591:     i->count = 1;
592:     i->dirty = 1;
593:
594:     switch(mode & S_IFMT) {
595:     case S_IFCHR:
596:         i->fsop = &def_chr_fsop;
597:         i->rdev = dev;
598:         i->i_mode |= S_IFCHR;
599:         d->file_type = 0;          /* EXT2_FT_CHRDEV not used */
600:         break;
601:     case S_IFBLK:
602:         i->fsop = &def_blk_fsop;
603:         i->rdev = dev;
604:         i->i_mode |= S_IFBLK;
605:         d->file_type = 0;          /* EXT2_FT_BLKDEV not used */
606:         break;
607:     case S_IFIFO:
608:         i->fsop = &pipefs_fsop;
609:         i->i_mode |= S_IFIFO;
610:         /* it's a union so we need to clear pipefs_i */
611:         memset_b(&i->u.pipefs, NULL, sizeof(struct pipefs_inode)
);
612:         d->file_type = 0;          /* EXT2_FT_FIFO not used */
613:         break;
614:     }
615:
616:     dir->i_mtime = CURRENT_TIME;
617:     dir->i_ctime = CURRENT_TIME;
618:     dir->dirty = 1;
619:
620:     bwrite(buf);
621:     iput(i);
622:     inode_unlock(dir);
623:     return 0;
624: }
625:
626: int ext2_create(struct inode *dir, char *name, __mode_t mode, struct inode **i_r
es)
627: {
628:     struct buffer *buf;
629:     struct inode *i;
630:     struct ext2_dir_entry_2 *d;
631:     char c;
632:     int n;
633:
634:     if(IS_RDONLY_FS(dir)) {
635:         return -EROFS;
636:     }
637:
638:     inode_lock(dir);
639:
640:     if(!(i = ialloc(dir->sb, S_IFREG))) {
641:         inode_unlock(dir);
642:         return -ENOSPC;
643:     }
644:
645:     if(!(buf = add_dir_entry(dir, &d, name))) {

```

## fs/ext2/namei.c

Page 11/12

```

646:         i->i_nlink = 0;
647:         iput(i);
648:         inode_unlock(dir);
649:         return -ENOSPC;
650:     }
651:
652:     d->inode = i->inode;
653:     d->name_len = strlen(name);
654:     /* strcpy() can't be used here because it places a trailing NULL */
655:     for(n = 0; n < NAME_MAX; n++) {
656:         if((c = name[n])) {
657:             d->name[n] = c;
658:             continue;
659:         }
660:         break;
661:     }
662:     d->file_type = 0;          /* EXT2_FT_REG_FILE not used */
663:
664:     i->i_mode = (mode & ~current->umask) & ~S_IFMT;
665:     i->i_mode |= S_IFREG;
666:     i->i_uid = current->euid;
667:     i->i_gid = current->egid;
668:     i->i_nlink = 1;
669:     i->i_blocks = 0;
670:     i->dev = dir->dev;
671:     i->fsop = &ext2_file_fsop;
672:     i->count = 1;
673:     i->dirty = 1;
674:
675:     i->u.ext2.i_dtime = 0;
676:
677:     dir->i_mtime = CURRENT_TIME;
678:     dir->i_ctime = CURRENT_TIME;
679:     dir->dirty = 1;
680:
681:     *i_res = i;
682:     bwrite(buf);
683:     inode_unlock(dir);
684:     return 0;
685: }
686:
687: int ext2_rename(struct inode *i_old, struct inode *dir_old, struct inode *i_new,
struct inode *dir_new, char *oldpath, char *newpath)
688: {
689:     struct buffer *buf_old, *buf_new;
690:     struct ext2_dir_entry_2 *d_old, *d_new;
691:     char c;
692:     int n, errno;
693:
694:     errno = 0;
695:     buf_new = NULL;
696:
697:     if(is_prefix(dir_new, i_old)) {
698:         return -EINVAL;
699:     }
700:
701:     inode_lock(i_old);
702:     inode_lock(dir_old);
703:     if(dir_old != dir_new) {
704:         inode_lock(dir_new);
705:     }
706:
707:     if(!(buf_old = find_dir_entry(dir_old, i_old, &d_old, oldpath))) {
708:         errno = -ENOENT;
709:         goto end;
710:     }
711:     if(dir_old == dir_new) {

```

## fs/ext2/namei.c

Page 12/12

```

712:         buf_old->locked = 0;
713:     }
714:
715:     if(i_new) {
716:         if(S_ISDIR(i_old->i_mode)) {
717:             if(!is_dir_empty(i_new)) {
718:                 brelse(buf_old);
719:                 errno = -ENOTEMPTY;
720:                 goto end;
721:             }
722:         }
723:         if(!(buf_new = find_dir_entry(dir_new, i_new, &d_new, newpath)))
{
724:             brelse(buf_old);
725:             errno = -ENOENT;
726:             goto end;
727:         }
728:     } else {
729:         if(!(buf_new = add_dir_entry(dir_new, &d_new, newpath))) {
730:             brelse(buf_old);
731:             errno = -ENOSPC;
732:             goto end;
733:         }
734:     }
735:     if(i_new) {
736:         i_new->i_nlink--;
737:     } else {
738:         i_new = i_old;
739:         d_new->name_len = strlen(newpath);
740:         /* strcpy() can't be used here because it places a trailing NULL
*/
741:         for(n = 0; n < NAME_MAX; n++) {
742:             if((c = newpath[n])) {
743:                 d_new->name[n] = c;
744:                 continue;
745:             }
746:             break;
747:         }
748:     }
749:
750:     d_old->inode = 0;
751:     d_new->inode = i_old->inode;
752:     dir_new->i_mtime = CURRENT_TIME;
753:     dir_new->i_ctime = CURRENT_TIME;
754:     i_new->dirty = 1;
755:     dir_new->dirty = 1;
756:
757:     dir_old->i_mtime = CURRENT_TIME;
758:     dir_old->i_ctime = CURRENT_TIME;
759:     i_old->dirty = 1;
760:     dir_old->dirty = 1;
761:
762:     bwrite(buf_old);
763:     if(buf_new) {
764:         bwrite(buf_new);
765:     }
766:
767: end:
768:     inode_unlock(i_old);
769:     inode_unlock(dir_old);
770:     inode_unlock(dir_new);
771:     return errno;
772: }

```

## fs/ext2/super.c

Page 1/4

```

1: /*
2:  * fiwix/fs/ext2/super.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/fs_ext2.h>
14: #include <fiwix/buffer.h>
15: #include <fiwix/sched.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: struct fs_operations ext2_fsop = {
20:     FSOP_REQUIRES_DEV,
21:     NULL,
22:
23:     NULL,          /* open */
24:     NULL,          /* close */
25:     NULL,          /* read */
26:     NULL,          /* write */
27:     NULL,          /* ioctl */
28:     NULL,          /* lseek */
29:     NULL,          /* readdir */
30:     NULL,          /* mmap */
31:     NULL,          /* select */
32:
33:     NULL,          /* readlink */
34:     NULL,          /* followlink */
35:     NULL,          /* bmap */
36:     NULL,          /* lookup */
37:     NULL,          /* rmdir */
38:     NULL,          /* link */
39:     NULL,          /* unlink */
40:     NULL,          /* symlink */
41:     NULL,          /* mkdir */
42:     NULL,          /* mknod */
43:     NULL,          /* truncate */
44:     NULL,          /* create */
45:     NULL,          /* rename */
46:
47:     NULL,          /* read_block */
48:     NULL,          /* write_block */
49:
50:     ext2_read_inode,
51:     ext2_write_inode,
52:     ext2_ialloc,
53:     ext2_ifree,
54:     ext2_statfs,
55:     ext2_read_superblock,
56:     ext2_remount_fs,
57:     ext2_write_superblock,
58:     ext2_release_superblock
59: };
60:
61: static void check_superblock(struct ext2_super_block *sb)
62: {
63:     if (!(sb->s_state & EXT2_VALID_FS)) {
64:         printk("WARNING: filesystem unchecked, fsck recommended.\n");
65:     } else if ((sb->s_state & EXT2_ERROR_FS)) {
66:         printk("WARNING: filesystem contains errors, fsck recommended.\n");
67:     }
68: };

```

## fs/ext2/super.c

Page 2/4

```

67:         } else if(sb->s_max_mnt_count >= 0 && sb->s_mnt_count >= (unsigned short
int)sb->s_max_mnt_count) {
68:             printk("WARNING: maximal mount count reached, fsck recommended.\n"
n");
69:         } else if(sb->s_checkinterval && (sb->s_lastcheck + sb->s_checkinterval
<= CURRENT_TIME)) {
70:             printk("WARNING: checktime reached, fsck recommended.\n");
71:         }
72:     }
73:
74: void ext2_statfs(struct superblock *sb, struct statfs *statfsbuf)
75: {
76:     statfsbuf->f_type = EXT2_SUPER_MAGIC;
77:     statfsbuf->f_bsize = sb->s_blocksize;
78:     statfsbuf->f_blocks = sb->u.ext2.sb.s_blocks_count;
79:     statfsbuf->f_bfree = sb->u.ext2.sb.s_free_blocks_count;
80:     if(statfsbuf->f_bfree >= sb->u.ext2.sb.s_r_blocks_count) {
81:         statfsbuf->f_bavail = statfsbuf->f_bfree - sb->u.ext2.sb.s_r_blo
cks_count;
82:     } else {
83:         statfsbuf->f_bavail = 0;
84:     }
85:     statfsbuf->f_files = sb->u.ext2.sb.s_inodes_count;
86:     statfsbuf->f_ffree = sb->u.ext2.sb.s_free_inodes_count;
87:     /* statfsbuf->f_fsid = ? */
88:     statfsbuf->f_namelen = EXT2_NAME_LEN;
89: }
90:
91: int ext2_read_superblock(__dev_t dev, struct superblock *sb)
92: {
93:     struct buffer *buf;
94:     struct ext2_super_block *ext2sb;
95:
96:     superblock_lock(sb);
97:     if(!(buf = bread(dev, SUPERBLOCK, BLKSIZE_1K))) {
98:         printk("WARNING: %s(): I/O error on device %d,%d.\n", __FUNCTION
__, MAJOR(dev), MINOR(dev));
99:         superblock_unlock(sb);
100:        return -EIO;
101:    }
102:
103:    ext2sb = (struct ext2_super_block *)buf->data;
104:    if(ext2sb->s_magic != EXT2_SUPER_MAGIC) {
105:        printk("WARNING: %s(): invalid filesystem type or bad superblock
on device %d,%d.\n", __FUNCTION__, MAJOR(dev), MINOR(dev));
106:        superblock_unlock(sb);
107:        brelse(buf);
108:        return -EINVAL;
109:    }
110:
111:    if(ext2sb->s_minor_rev_level || ext2sb->s_rev_level) {
112:        printk("WARNING: %s(): unsupported ext2 filesystem revision.\n",
__FUNCTION__);
113:        printk("Only revision 0 (original without features) is supported
.\n");
114:        superblock_unlock(sb);
115:        brelse(buf);
116:        return -EINVAL;
117:    }
118:
119:    sb->dev = dev;
120:    sb->fsop = &ext2_fsop;
121:    sb->s_blocksize = EXT2_MIN_BLOCK_SIZE << ext2sb->s_log_block_size;
122:    memcpy_b(&sb->u.ext2.sb, ext2sb, sizeof(struct ext2_super_block));
123:    EXT2_DESC_PER_BLOCK(sb) = sb->s_blocksize / sizeof(struct ext2_group_des
c);
124:    sb->u.ext2.block_groups = 1 + (ext2sb->s_blocks_count - 1) / ext2sb->s_b

```

## fs/ext2/super.c

Page 3/4

```

locks_per_group;
125:
126:         if(!(sb->root = iget(sb, EXT2_ROOT_INO))) {
127:             printk("WARNING: %s(): unable to get root inode.\n", __FUNCTION_
_)
);
128:             superblock_unlock(sb);
129:             brelse(buf);
130:             return -EINVAL;
131:         }
132:
133:         check_superblock(ext2sb);
134:         if(!(sb->flags & MS_RDONLY)) {
135:             sb->u.ext2.sb.s_state &= ~EXT2_VALID_FS;
136:             sb->u.ext2.sb.s_mnt_count++;
137:             sb->u.ext2.sb.s_mtime = CURRENT_TIME;
138:             memcpy_b(buf->data, &sb->u.ext2.sb, sizeof(struct ext2_super_blo
ck));
139:             bwrite(buf);
140:         } else {
141:             brelse(buf);
142:         }
143:         superblock_unlock(sb);
144:         return 0;
145:     }
146:
147: int ext2_remount_fs(struct superblock *sb, int flags)
148: {
149:     struct buffer *buf;
150:     struct ext2_super_block *ext2sb;
151:
152:     if((flags & MS_RDONLY) == (sb->flags & MS_RDONLY)) {
153:         return 0;
154:     }
155:
156:     superblock_lock(sb);
157:     if(!(buf = bread(sb->dev, SUPERBLOCK, BLKSIZE_1K))) {
158:         superblock_unlock(sb);
159:         return -EIO;
160:     }
161:     ext2sb = (struct ext2_super_block *)buf->data;
162:
163:     if(flags & MS_RDONLY) {
164:         /* switching from RW to RO */
165:         sb->u.ext2.sb.s_state |= EXT2_VALID_FS;
166:         ext2sb->s_state |= EXT2_VALID_FS;
167:     } else {
168:         /* switching from RO to RW */
169:         check_superblock(ext2sb);
170:         sb->u.ext2.sb.s_state &= ~EXT2_VALID_FS;
171:         sb->u.ext2.sb.s_mnt_count++;
172:         sb->u.ext2.sb.s_mtime = CURRENT_TIME;
173:         ext2sb->s_state &= ~EXT2_VALID_FS;
174:     }
175:
176:     sb->dirty = 1;
177:     superblock_unlock(sb);
178:     bwrite(buf);
179:     return 0;
180: }
181:
182: int ext2_write_superblock(struct superblock *sb)
183: {
184:     struct buffer *buf;
185:
186:     superblock_lock(sb);
187:     if(!(buf = bread(sb->dev, SUPERBLOCK, BLKSIZE_1K))) {
188:         superblock_unlock(sb);

```



```
189:         return -EIO;
190:     }
191:
192:     memcpy_b(buf->data, &sb->u.ext2.sb, sizeof(struct ext2_super_block));
193:     sb->dirty = 0;
194:     superblock_unlock(sb);
195:     bwrite(buf);
196:     return 0;
197: }
198:
199: void ext2_release_superblock(struct superblock *sb)
200: {
201:     if(sb->flags & MS_RDONLY) {
202:         return;
203:     }
204:
205:     superblock_lock(sb);
206:
207:     sb->u.ext2.sb.s_state |= EXT2_VALID_FS;
208:     sb->dirty = 1;
209:
210:     superblock_unlock(sb);
211: }
212:
213: int ext2_init(void)
214: {
215:     return register_filesystem("ext2", &ext2_fsop);
216: }
```

## fs/ext2/symlink.c

Page 1/3

```

1: /*
2:  * fiwix/fs/ext2/symlink.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: struct fs_operations ext2_symlink_fsop = {
19:     0,
20:     0,
21:
22:     NULL,          /* open */
23:     NULL,          /* close */
24:     NULL,          /* read */
25:     NULL,          /* write */
26:     NULL,          /* ioctl */
27:     NULL,          /* lseek */
28:     NULL,          /* readdir */
29:     NULL,          /* mmap */
30:     NULL,          /* select */
31:
32:     ext2_readlink,
33:     ext2_followlink,
34:     NULL,          /* bmap */
35:     NULL,          /* lookup */
36:     NULL,          /* rmdir */
37:     NULL,          /* link */
38:     NULL,          /* unlink */
39:     NULL,          /* symlink */
40:     NULL,          /* mkdir */
41:     NULL,          /* mknod */
42:     NULL,          /* truncate */
43:     NULL,          /* create */
44:     NULL,          /* rename */
45:
46:     NULL,          /* read_block */
47:     NULL,          /* write_block */
48:
49:     NULL,          /* read_inode */
50:     NULL,          /* write_inode */
51:     NULL,          /* ialloc */
52:     NULL,          /* ifree */
53:     NULL,          /* statfs */
54:     NULL,          /* read_superblock */
55:     NULL,          /* remount_fs */
56:     NULL,          /* write_superblock */
57:     NULL,          /* release_superblock */
58: };
59:
60: int ext2_readlink(struct inode *i, char *buffer, __size_t count)
61: {
62:     __u32 blksize;
63:     struct buffer *buf;
64:
65:     if(!S_ISLNK(i->i_mode)) {
66:         printk("%s(): Oops, inode '%d' is not a symlink (!?)\n", __FUNC
TION__, i->inode);

```

## fs/ext2/symlink.c

Page 2/3

```

67:         return 0;
68:     }
69:
70:     inode_lock(i);
71:     blksize = i->sb->s_blocksize;
72:     count = MIN(count, i->i_size);
73:     if(!count) {
74:         inode_unlock(i);
75:         return 0;
76:     }
77:     count = MIN(count, blksize);
78:     if(i->i_blocks) { /* slow symlink */
79:         if(!(buf = bread(i->dev, i->u.ext2.i_data[0], blksize))) {
80:             inode_unlock(i);
81:             return -EIO;
82:         }
83:         memcpy_b(buffer, buf->data, count);
84:         brelse(buf);
85:     } else { /* fast symlink */
86:         memcpy_b(buffer, (char *)i->u.ext2.i_data, count);
87:     }
88:     buffer[count] = NULL;
89:     inode_unlock(i);
90:     return count;
91: }
92:
93: int ext2_followlink(struct inode *dir, struct inode *i, struct inode **i_res)
94: {
95:     struct buffer *buf;
96:     char *name;
97:     __ino_t errno;
98:
99:     if(!i) {
100:         return -ENOENT;
101:     }
102:
103:     if(!S_ISLNK(i->i_mode)) {
104:         printk("%s(): Oops, inode '%d' is not a symlink (!?).\n", __FUNC
TION__, i->inode);
105:         return 0;
106:     }
107:
108:     if(current->loopcnt > MAX_SYMLINKS) {
109:         printk("%s(): too many nested symbolic links!\n", __FUNCTION__);
110:         return -ELOOP;
111:     }
112:
113:     inode_lock(i);
114:     if(i->i_blocks) { /* slow symlink */
115:         if(!(buf = bread(i->dev, i->u.ext2.i_data[0], i->sb->s_blocksize
))) {
116:             inode_unlock(i);
117:             return -EIO;
118:         }
119:         name = buf->data;
120:     } else { /* fast symlink */
121:         buf = NULL;
122:         name = (char *)i->u.ext2.i_data;
123:     }
124:     inode_unlock(i);
125:
126:     current->loopcnt++;
127:     iput(i);
128:     if(buf) {
129:         brelse(buf);
130:     }
131:     errno = parse_namei(name, dir, i_res, NULL, FOLLOW_LINKS);

```

**fs/ext2/symlink.c**

Page 3/3

```
132:         current->loopcnt--;  
133:         return errno;  
134: }
```

## fs/iso9660/dir.c

Page 1/3

```

1: /*
2:  * fiwix/fs/iso9660/dir.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/dirent.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: struct fs_operations iso9660_dir_fsop = {
19:     0,
20:     0,
21:
22:     iso9660_dir_open,
23:     iso9660_dir_close,
24:     iso9660_dir_read,
25:     NULL, /* write */
26:     NULL, /* ioctl */
27:     NULL, /* lseek */
28:     iso9660_dir_readdir,
29:     NULL, /* mmap */
30:     NULL, /* select */
31:
32:     NULL, /* readlink */
33:     NULL, /* followlink */
34:     iso9660_bmap,
35:     iso9660_lookup,
36:     NULL, /* rmdir */
37:     NULL, /* link */
38:     NULL, /* unlink */
39:     NULL, /* symlink */
40:     NULL, /* mkdir */
41:     NULL, /* mknod */
42:     NULL, /* truncate */
43:     NULL, /* create */
44:     NULL, /* rename */
45:
46:     NULL, /* read_block */
47:     NULL, /* write_block */
48:
49:     NULL, /* read_inode */
50:     NULL, /* write_inode */
51:     NULL, /* ialloc */
52:     NULL, /* ifree */
53:     NULL, /* statsfs */
54:     NULL, /* read_superblock */
55:     NULL, /* remount_fs */
56:     NULL, /* write_superblock */
57:     NULL, /* release_superblock */
58: };
59:
60: int iso9660_dir_open(struct inode *i, struct fd *fd_table)
61: {
62:     fd_table->offset = 0;
63:     return 0;
64: }
65:
66: int iso9660_dir_close(struct inode *i, struct fd *fd_table)
67: {

```

## fs/iso9660/dir.c

Page 2/3

```

68:         return 0;
69:     }
70:
71: int iso9660_dir_read(struct inode *i, struct fd *fd_table, char *buffer, __size_
t count)
72: {
73:     return -EISDIR;
74: }
75:
76: int iso9660_dir_readdir(struct inode *i, struct fd *fd_table, struct dirent *dir
ent, unsigned int count)
77: {
78:     __blk_t block;
79:     unsigned int doffset, offset;
80:     unsigned int size, dirent_len;
81:     struct iso9660_directory_record *d;
82:     int base_dirent_len;
83:     int blksize;
84:     struct buffer *buf;
85:     int nm_len;
86:     char nm_name[NAME_MAX + 1];
87:
88:     if(!(S_ISDIR(i->i_mode))) {
89:         return -EBADF;
90:     }
91:
92:     blksize = i->sb->s_blocksize;
93:     if(fd_table->offset > i->i_size) {
94:         fd_table->offset = i->i_size;
95:     }
96:
97:     base_dirent_len = sizeof(dirent->d_ino) + sizeof(dirent->d_off) + sizeof
(dirent->d_reclen);
98:     doffset = offset = size = 0;
99:
100:    while(doffset < count) {
101:        if((block = bmap(i, fd_table->offset, FOR_READING)) < 0) {
102:            return block;
103:        }
104:        if(block) {
105:            if(!(buf = bread(i->dev, block, blksize))) {
106:                return -EIO;
107:            }
108:
109:            doffset = fd_table->offset;
110:            offset = fd_table->offset % blksize;
111:
112:            while(doffset < i->i_size && offset < blksize) {
113:                d = (struct iso9660_directory_record *) (buf->dat
a + offset);
114:                if(isonum_711(d->length)) {
115:                    dirent_len = (base_dirent_len + (isonum_
711(d->name_len) + 1)) + 3;
116:                    dirent_len &= ~3;        /* round up */
117:                    if((size + dirent_len) < count) {
118:                        dirent->d_ino = (block << ISO966
0_INODE_BITS) + (doffset & ISO9660_INODE_MASK);
119:                        dirent->d_off = doffset;
120:                        dirent->d_reclen = dirent_len;
121:                        if(isonum_711(d->name_len) == 1
&& d->name[0] == 0) {
122:                            dirent->d_name[0] = '.';
123:                            dirent->d_name[1] = NULL
;
124:                        } else if(isonum_711(d->name_len
) == 1 && d->name[0] == 1) {
125:                            dirent->d_name[0] = '.';

```

## fs/iso9660/dir.c

Page 3/3

```

126:                                dirent->d_name[1] = '.';
127:                                dirent->d_name[2] = NULL;
;
128:                                dirent_len = 16;
129:                                dirent->d_reclen = 16;
130:                                if(i->u.iso9660.i_parent
) {
131:                                    dirent->d_ino =
i->u.iso9660.i_parent->inode;
132:                                } else {
133:                                    dirent->d_ino =
i->inode;
134:                                }
135:                                } else {
136:                                    nm_len = 0;
137:                                    if(i->sb->u.iso9660.rrip
) {
138:                                        nm_len = get_rri
p_filename(d, i, nm_name);
139:                                }
140:                                if(nm_len) {
141:                                    memcpy_b(dirent->
>d_name, nm_name, nm_len);
142:                                    dirent->d_name[n
m_len] = NULL;
143:                                    dirent->d_reclen
= (base_dirent_len + nm_len + 1) + 3;
144:                                    dirent->d_reclen
&= ~3; /* round up */
145:                                    dirent_len = dir
ent->d_reclen;
146:                                } else {
147:                                    memcpy_b(dirent->
>d_name, d->name, isonum_711(d->name_len));
148:                                    dirent->d_name[i
sonum_711(d->name_len)] = NULL;
149:                                }
150:                                }
151:                                if(!((char)d->flags[0] & ISO9660
_FILE_ISDIR)) {
152:                                    iso9660_cleanfilename(di
rent->d_name, isonum_711(d->name_len));
153:                                }
154:                                dirent = (struct dirent *)((char
*)dirent + dirent_len);
155:                                size += dirent_len;
156:                                } else {
157:                                    break;
158:                                }
159:                                doffset += isonum_711(d->length);
160:                                offset += isonum_711(d->length);
161:                                } else {
162:                                    break;
163:                                }
164:                                }
165:                                brelse(buf);
166:                                }
167:                                fd_table->offset &= ~(blksize - 1);
168:                                doffset = fd_table->offset;
169:                                doffset += blksize;
170:                                fd_table->offset += blksize;
171:                                }
172:                                return size;
173: }

```

## fs/iso9660/file.c

Page 1/2

```

1: /*
2:  * fiwix/fs/iso9660/file.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/mm.h>
14: #include <fiwix/mman.h>
15: #include <fiwix/fcntl.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: struct fs_operations iso9660_file_fsop = {
20:     0,
21:     0,
22:
23:     iso9660_file_open,
24:     iso9660_file_close,
25:     file_read,
26:     NULL, /* write */
27:     NULL, /* ioctl */
28:     iso9660_file_lseek,
29:     NULL, /* readdir */
30:     NULL, /* mmap */
31:     NULL, /* select */
32:
33:     NULL, /* readlink */
34:     NULL, /* followlink */
35:     iso9660_bmap,
36:     NULL, /* lookup */
37:     NULL, /* rmdir */
38:     NULL, /* link */
39:     NULL, /* unlink */
40:     NULL, /* symlink */
41:     NULL, /* mkdir */
42:     NULL, /* mknod */
43:     NULL, /* truncate */
44:     NULL, /* create */
45:     NULL, /* rename */
46:
47:     NULL, /* read_block */
48:     NULL, /* write_block */
49:
50:     NULL, /* read_inode */
51:     NULL, /* write_inode */
52:     NULL, /* ialloc */
53:     NULL, /* ifree */
54:     NULL, /* statfs */
55:     NULL, /* read_superblock */
56:     NULL, /* remount_fs */
57:     NULL, /* write_superblock */
58:     NULL, /* release_superblock */
59: };
60:
61: int iso9660_file_open(struct inode *i, struct fd *fd_table)
62: {
63:     if(fd_table->flags & (O_WRONLY | O_RDWR | O_TRUNC | O_APPEND)) {
64:         return -ENOENT;
65:     }
66:     fd_table->offset = 0;
67:     return 0;

```



```
68: }
69:
70: int iso9660_file_close(struct inode *i, struct fd *fd_table)
71: {
72:     return 0;
73: }
74:
75: int iso9660_file_lseek(struct inode *i, __off_t offset)
76: {
77:     return offset;
78: }
```

## fs/iso9660/inode.c

Page 1/3

```

1: /*
2:  * fiwix/fs/iso9660/inode.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/fs_iso9660.h>
14: #include <fiwix/fs_pipe.h>
15: #include <fiwix/buffer.h>
16: #include <fiwix/stat.h>
17: #include <fiwix/mm.h>
18: #include <fiwix/sched.h>
19: #include <fiwix/stdio.h>
20: #include <fiwix/string.h>
21:
22: static int read_pathtable(struct inode *i)
23: {
24:     int n, offset, pt_len, pt_blk;
25:     struct iso9660_sb_info *sbi;
26:     struct iso9660_pathtable_record *ptr;
27:     struct buffer *buf;
28:
29:     sbi = (struct iso9660_sb_info *)&i->sb->u.iso9660;
30:     pt_len = isonum_733(sbi->sb->path_table_size);
31:     pt_blk = isonum_731(sbi->sb->type_l_path_table);
32:
33:     if(pt_len > PAGE_SIZE) {
34:         printk("WARNING: %s(): path table record size (%d) > 4096, not s
supported yet.\n", __FUNCTION__, pt_len);
35:         return -EINVAL;
36:     }
37:
38:     if(!(sbi->pathtable_raw = (void *)kmalloc())) {
39:         return -ENOMEM;
40:     }
41:     offset = 0;
42:     while(offset < pt_len) {
43:         if(!(buf = bread(i->dev, pt_blk++, BLKSIZE_2K))) {
44:             kfree((unsigned int)sbi->pathtable_raw);
45:             return -EIO;
46:         }
47:         memcpy_b(sbi->pathtable_raw + offset, (void *)buf->data, MIN(pt_
len - offset, BLKSIZE_2K));
48:         offset += MIN(pt_len - offset, BLKSIZE_2K);
49:         brelse(buf);
50:     }
51:
52:     /* allocate and count the number of records in the Path Table */
53:     offset = n = 0;
54:     if(!(sbi->pathtable = (struct iso9660_pathtable_record **)kmalloc())) {
55:         kfree((unsigned int)sbi->pathtable_raw);
56:         return -ENOMEM;
57:     }
58:     sbi->pathtable[n] = NULL;
59:     while(offset < pt_len) {
60:         ptr = (struct iso9660_pathtable_record *) (sbi->pathtable_raw + o
ffset);
61:         sbi->pathtable[++n] = ptr;
62:         offset += sizeof(struct iso9660_pathtable_record) + isonum_711(p
tr->length) + (isonum_711(ptr->length) & 1);
63:     }

```

## fs/iso9660/inode.c

Page 2/3

```

64:         sbi->paths = n;
65:
66:         return 0;
67:     }
68:
69: static int get_parent_dir_size(struct superblock *sb, __blk_t extent)
70: {
71:     int n;
72:     struct iso9660_pathtable_record *ptr;
73:     __blk_t parent;
74:
75:     for(n = 0; n < sb->u.iso9660.paths; n++) {
76:         ptr = (struct iso9660_pathtable_record *)sb->u.iso9660.pathtable
[n];
77:         if(isonum_731(ptr->extent) == extent) {
78:
79:             parent = isonum_723(ptr->parent);
80:             ptr = (struct iso9660_pathtable_record *)sb->u.iso9660.p
athtable[parent];
81:             parent = isonum_731(ptr->extent);
82:             return parent;
83:         }
84:     }
85:     printk("WARNING: %s(): unable to locate extent '%d' in path table.\n", _
_FUNCTION__, extent);
86:     return 0;
87: }
88:
89: int iso9660_read_inode(struct inode *i)
90: {
91:     int errno;
92:     __u32 blksize;
93:     struct superblock *sb;
94:     struct iso9660_directory_record *d;
95:     struct buffer *buf;
96:     __blk_t dblock;
97:     __off_t doffset;
98:
99:     sb = (struct superblock *)i->sb;
100:    if(!sb->u.iso9660.pathtable) {
101:        if((errno = read_pathtable(i))) {
102:            return errno;
103:        }
104:    }
105:
106:    dblock = (i->inode & ~ISO9660_INODE_MASK) >> ISO9660_INODE_BITS;
107:    doffset = i->inode & ISO9660_INODE_MASK;
108:    blksize = i->sb->s_blocksize;
109:
110:    /* FIXME: it only looks in one directory block */
111:    if(!(buf = bread(i->dev, dblock, blksize))) {
112:        return -EIO;
113:    }
114:
115:    if(doffset >= blksize) {
116:        printk("WARNING: %s(): inode %d (dblock=%d, doffset=%d) not foun
d in directory entry.\n", _FUNCTION__, i->inode, dblock, doffset);
117:        brelse(buf);
118:        return -EIO;
119:    }
120:    d = (struct iso9660_directory_record *) (buf->data + doffset);
121:
122:    i->i_mode = S_IFREG;
123:    if((char)d->flags[0] & ISO9660_FILE_ISDIR) {
124:        i->i_mode = S_IFDIR;
125:    }
126:    if(!(char)d->flags[0] & ISO9660_FILE_HASOWNER) {

```

## fs/iso9660/inode.c

Page 3/3

```

127:             i->i_mode |= S_IRUSR | S_IXUSR | S_IRGRP | S_IXGRP | S_IROTH | S
_IXOTH;
128:         }
129:         i->i_uid = 0;
130:         i->i_size = isonum_733(d->size);
131:         i->i_atime = isodate(d->date);
132:         i->i_ctime = isodate(d->date);
133:         i->i_mtime = isodate(d->date);
134:         i->i_gid = 0;
135:         i->i_nlink = 1;
136:         i->i_blocks = 0;          /* FIXME */
137:         i->i_flags = 0;          /* FIXME */
138:         i->locked = 1;
139:         i->dirty = 0;
140:         i->mount_point = NULL;
141:         i->count = 1;
142:
143:         i->u.iso9660.i_extent = isonum_733(d->extent);
144:         check_rrip_inode(d, i);
145:         brelse(buf);
146:
147:         switch(i->i_mode & S_IFMT) {
148:             case S_IFCHR:
149:                 i->fsop = &def_chr_fsop;
150:                 break;
151:             case S_IFBLK:
152:                 i->fsop = &def_blk_fsop;
153:                 break;
154:             case S_IFIFO:
155:                 i->fsop = &pipefs_fsop;
156:                 /* it's a union so we need to clear pipefs_inode */
157:                 memset_b(&i->u.pipefs, NULL, sizeof(struct pipefs_inode)
);
158:                 break;
159:             case S_IFDIR:
160:                 i->fsop = &iso9660_dir_fsop;
161:                 i->i_nlink++;
162:                 break;
163:             case S_IFREG:
164:                 i->fsop = &iso9660_file_fsop;
165:                 break;
166:             case S_IFLNK:
167:                 i->fsop = &iso9660_symlink_fsop;
168:                 break;
169:             case S_IFSOCK:
170:                 i->fsop = NULL;
171:                 break;
172:             default:
173:                 PANIC("invalid inode (%d) mode %08o.\n", i->inode, i->i
mode);
174:         }
175:         return 0;
176: }
177:
178: int iso9660_bmap(struct inode *i, __off_t offset, int mode)
179: {
180:     __blk_t block;
181:
182:     block = i->u.iso9660.i_extent + (offset / i->sb->s_blocksize);
183:     return block;
184: }

```

**fs/iso9660/Makefile**

Page 1/1

```
1: # fiwix/fs/iso9660/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = inode.o super.o namei.o dir.o file.o rrip.o symlink.o
13:
14: iso9660:$ (OBJS)
15:     $(LD) $(LDFLAGS) -r $(OBJS) -o iso9660.o
16:
17: clean:
18:     rm -f *.o
19:
```

## fs/iso9660/namei.c

Page 1/3

```

1: /*
2:  * fiwix/fs/iso9660/namei.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/fs_iso9660.h>
12: #include <fiwix/buffer.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/errno.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: int iso9660_lookup(const char *name, struct inode *dir, struct inode **i_res)
20: {
21:     __blk_t dblock;
22:     __u32 blksize;
23:     int len, dnlen;
24:     unsigned int offset, doffset;
25:     struct buffer *buf;
26:     struct iso9660_directory_record *d;
27:     __ino_t inode;
28:     int nm_len;
29:     char *nm_name;
30:
31:     blksize = dir->sb->s_blocksize;
32:     inode = offset = 0;
33:     len = strlen(name);
34:     dir->count++;
35:
36:     while(offset < dir->i_size && !inode) {
37:         if((dblock = bmap(dir, offset, FOR_READING)) < 0) {
38:             return dblock;
39:         }
40:         if(dblock) {
41:             if(!(buf = bread(dir->dev, dblock, blksize))) {
42:                 return -EIO;
43:             }
44:             doffset = 0;
45:             do {
46:                 d = (struct iso9660_directory_record *) (buf->data
a + doffset);
47:                 if(isonum_711(d->length) == 0) {
48:                     break;
49:                 }
50:                 if(len == 1) {
51:                     if(name[0] == '.' && name[1] == NULL) {
52:                         if(isonum_711(d->name_len) == 1
&& d->name[0] == 0) {
53:                             inode = dir->inode;
54:                         }
55:                     }
56:                 }
57:                 if(len == 2) {
58:                     if(name[0] == '.' && name[1] == '.' && n
ame[2] == NULL) {
59:                         if(isonum_711(d->name_len) == 1
&& d->name[0] == 1) {
60:                             inode = dir->u.iso9660.i
nfo->parent->inode;
61:                         }
62:                     }

```

## fs/iso9660/namei.c

Page 2/3

```

63:         }
64:         if(!(nm_name = (char *)kmalloc())) {
65:             return -ENOMEM;
66:         }
67:         nm_len = 0;
68:         if(dir->sb->u.iso9660.rrip) {
69:             nm_len = get_rrip_filename(d, dir, nm_na
me);
70:         }
71:         if(nm_len) {
72:             dnlen = nm_len;
73:         } else {
74:             dnlen = isonum_711(d->name_len);
75:             if(!((char)d->flags[0] & ISO9660_FILE_IS
DIR)) {
76:                 iso9660_cleanfilename(d->name, d
nlen);
77:                 dnlen = strlen(d->name);
78:             }
79:         }
80:         if(len == dnlen) {
81:             if(nm_len) {
82:                 if(strncmp(nm_name, name, dnlen)
== 0) {
83:                     inode = (dblock << ISO96
60_INODE_BITS) + (doffset & ISO9660_INODE_MASK);
84:                 }
85:             } else {
86:                 if(strncmp(d->name, name, dnlen)
== 0) {
87:                     inode = (dblock << ISO96
60_INODE_BITS) + (doffset & ISO9660_INODE_MASK);
88:                 }
89:             }
90:         }
91:         kfree((unsigned int)nm_name);
92:         doffset += isonum_711(d->length);
93:     } while((doffset < blksize) && (!inode));
94:     brelse(buf);
95:     offset += blksize;
96:     if(inode) {
97:         /*
98:          * This prevents a deadlock in iget() when
99:          * trying to lock '.' when 'dir' is the same
100:          * directory (ls -lai <tmp>).
101:          */
102:         if(inode == dir->inode) {
103:             *i_res = dir;
104:             return 0;
105:         }
106:
107:         if(!(*i_res = iget(dir->sb, inode))) {
108:             return -EACCES;
109:         }
110:         if(S_ISDIR((*i_res)->i_mode)) {
111:             if(!(*i_res)->u.iso9660.i_parent) {
112:                 (*i_res)->u.iso9660.i_parent = d
ir;
113:             }
114:         }
115:         iput(dir);
116:         return 0;
117:     }
118: }
119: }
120: iput(dir);
121: return -ENOENT;

```

```
122: }
```



## fs/iso9660/rrip.c

Page 1/6

```

1: /*
2:  * fiwix/fs/iso9660/rrip.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/stat.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/buffer.h>
14: #include <fiwix/fs_iso9660.h>
15: #include <fiwix/mm.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: void check_rrip_inode(struct iso9660_directory_record *d, struct inode *i)
20: {
21:     unsigned int total_len;
22:     unsigned int len;
23:     unsigned int sig;
24:     int n, nm_len, rootflag;
25:     struct susp_rrip *rrip;
26:     unsigned int dev_h, dev_l;
27:     unsigned int ce_block, ce_offset, ce_size;
28:     struct buffer *buf;
29:     unsigned char *sue;
30:     int sl_len;
31:     struct rrip_sl_component *slc;
32:
33:     ce_block = ce_offset = ce_size = 0;
34:     buf = NULL;
35:     total_len = isonum_711(d->length);
36:     len = isonum_711(d->name_len);
37:     if(!(len % 2)) {
38:         len++;
39:     }
40:     sue = (unsigned char *)d->name;
41:     nm_len = 0;
42:
43: loop:
44:     if(ce_block && ce_size) {
45:         /* FIXME: it only looks in one directory block */
46:         if(!(buf = bread(i->dev, ce_block, i->sb->s_blocksize))) {
47:             return;
48:         }
49:         sue = (unsigned char *)buf->data + ce_offset;
50:         total_len = ce_size;
51:         len = 0;
52:     }
53:
54:     while(len < total_len) {
55:         rrip = (struct susp_rrip *) (sue + len);
56:         if(rrip->len == 0) {
57:             break;
58:         }
59:         sig = GET_SIG(rrip->signature[0], rrip->signature[1]);
60:         switch(sig) {
61:             case GET_SIG('S', 'P'):
62:                 if(rrip->u.sp.magic[0] != SP_MAGIC1 || rrip->u.s
p.magic[1] != SP_MAGIC2) {
63:                     if(ce_block) {
64:                         brelse(buf);
65:                     }
66:                     return;

```

## fs/iso9660/rrip.c

Page 2/6

```

67:         }
68:         break;
69:     case GET_SIG('C', 'E'):
70:         if(ce_block) {
71:             brelse(buf);
72:         }
73:         ce_block = isonum_733(rrip->u.ce.block);
74:         ce_offset = isonum_733(rrip->u.ce.offset);
75:         ce_size = isonum_733(rrip->u.ce.size);
76:         goto loop;
77:         break;
78:     case GET_SIG('E', 'R'):
79:         i->sb->u.iso9660.rrip = 1;
80:         printk("ISO 9660 Extensions: ");
81:         for(n = 0; n < rrip->u.er.len_id; n++) {
82:             printk("%c", rrip->u.er.data[n]);
83:         }
84:         printk("\n");
85:         break;
86:     case GET_SIG('P', 'X'):
87:         i->i_mode = isonum_733(rrip->u.px.mode);
88:         i->i_nlink = isonum_733(rrip->u.px.nlink);
89:         i->i_uid = isonum_733(rrip->u.px.uid);
90:         i->i_gid = isonum_733(rrip->u.px.gid);
91:         break;
92:     case GET_SIG('P', 'N'):
93:         if(S_ISBLK(i->i_mode) || S_ISCHR(i->i_mode)) {
94:             dev_h = isonum_733(rrip->u.pn.dev_h);
95:             dev_l = isonum_733(rrip->u.pn.dev_l);
96:             i->rdev = MKDEV(dev_h, dev_l);
97:         }
98:         break;
99:     case GET_SIG('S', 'L'):
100:         sl_len = rootflag = 0;
101:         slc = (struct rrip_sl_component *)&rrip->u.sl.ar
ea;
102:         while(sl_len < (rrip->len - 5)) {
103:             if(sl_len && !rootflag) {
104:                 nm_len++;
105:             }
106:             rootflag = 0;
107:             switch(slc->flags & 0xE) {
108:                 case 0:
109:                     nm_len += slc->len;
110:                     break;
111:                 case SL_CURRENT:
112:                     nm_len += 1;
113:                     break;
114:                 case SL_PARENT:
115:                     nm_len += 2;
116:                     break;
117:                 case SL_ROOT:
118:                     nm_len += 1;
119:                     rootflag = 1;
120:                     break;
121:                 default:
122:                     printk("WARNING: %s(): u
nsupported RRIP SL flags %d.\n", __FUNCTION__, slc->flags & 0xE);
123:             }
124:             slc = (struct rrip_sl_component *)(((cha
r *)slc) + slc->len + sizeof(struct rrip_sl_component));
125:             sl_len += slc->len + sizeof(struct rrip_
sl_component);
126:         }
127:         i->i_size = nm_len;
128:         break;
129:     case GET_SIG('T', 'F'):

```

## fs/iso9660/rrip.c

Page 3/6

```

130:         n = 0;
131:         if(rrip->u.tf.flags & TF_CREATION) {
132:             i->i_ctime = isodate(rrip->u.tf.times[n+
+].time);
133:         }
134:         if(rrip->u.tf.flags & TF_MODIFY) {
135:             i->i_mtime = isodate(rrip->u.tf.times[n+
+].time);
136:         }
137:         if(rrip->u.tf.flags & TF_ACCESS) {
138:             i->i_atime = isodate(rrip->u.tf.times[n+
+].time);
139:         }
140:         if(rrip->u.tf.flags & TF_ATTRIBUTES) {
141:             i->i_ctime = isodate(rrip->u.tf.times[n+
+].time);
142:         }
143:         break;
144:     }
145:     len += rrip->len;
146: }
147: if(ce_block) {
148:     brelse(buf);
149: }
150: }
151:
152: int get_rrip_filename(struct iso9660_directory_record *d, struct inode *i, char
*name)
153: {
154:     unsigned int total_len;
155:     unsigned int len;
156:     unsigned int sig;
157:     int nm_len;
158:     struct susp_rrip *rrip;
159:     unsigned int ce_block, ce_offset, ce_size;
160:     struct buffer *buf;
161:     unsigned char *sue;
162:
163:     ce_block = ce_offset = ce_size = 0;
164:     buf = NULL;
165:     total_len = isonum_711(d->length);
166:     len = isonum_711(d->name_len);
167:     if(!(len % 2)) {
168:         len++;
169:     }
170:     sue = (unsigned char *)d->name;
171:     nm_len = 0;
172:
173: loop:
174:     if(ce_block && ce_size) {
175:         /* FIXME: it only looks in one directory block */
176:         if(!(buf = bread(i->dev, ce_block, i->sb->s_blocksize))) {
177:             return 0;
178:         }
179:         sue = (unsigned char *)buf->data + ce_offset;
180:         total_len = ce_size;
181:         len = 0;
182:     }
183:
184:     while(len < total_len) {
185:         rrip = (struct susp_rrip *) (sue + len);
186:         if(rrip->len == 0) {
187:             break;
188:         }
189:         sig = GET_SIG(rrip->signature[0], rrip->signature[1]);
190:         switch(sig) {
191:             case GET_SIG('S', 'P'):
```

## fs/iso9660/rrip.c

Page 4/6

```

192:                                if(rrip->u.sp.magic[0] != SP_MAGIC1 || rrip->u.s
p.magic[1] != SP_MAGIC2) {
193:                                    if(ce_block) {
194:                                        brelse(buf);
195:                                    }
196:                                    return 0;
197:                                }
198:                                break;
199:                                case GET_SIG('C', 'E'):
200:                                    if(ce_block) {
201:                                        brelse(buf);
202:                                    }
203:                                    ce_block = isonum_733(rrip->u.ce.block);
204:                                    ce_offset = isonum_733(rrip->u.ce.offset);
205:                                    ce_size = isonum_733(rrip->u.ce.size);
206:                                    goto loop;
207:                                case GET_SIG('N', 'M'):
208:                                    if(rrip->u.nm.flags) { /* FIXME: & ~(NM_CONTINUE
| NM_CURRENT | NM_PARENT)) { */
209:                                        printk("WARNING: %s(): unsupported NM fl
ag settings (%d).\n", __FUNCTION__, rrip->u.nm.flags);
210:                                        if(ce_block) {
211:                                            brelse(buf);
212:                                        }
213:                                        return 0;
214:                                    }
215:                                    nm_len = rrip->len - 5;
216:                                    memcpy_b(name, rrip->u.nm.name, nm_len);
217:                                    name[nm_len] = NULL;
218:                                    break;
219:                                }
220:                                len += rrip->len;
221:                            }
222:                            if(ce_block) {
223:                                brelse(buf);
224:                            }
225:                            return nm_len;
226:    }
227:
228: int get_rrip_symlink(struct inode *i, char *name)
229: {
230:     unsigned int total_len;
231:     unsigned int len;
232:     unsigned int sig;
233:     int nm_len;
234:     struct susp_rrip *rrip;
235:     unsigned int ce_block, ce_offset, ce_size;
236:     struct buffer *buf;
237:     struct buffer *buf2;
238:     unsigned char *sue;
239:     struct iso9660_directory_record *d;
240:     __blk_t dblock;
241:     __off_t doffset;
242:     int sl_len, rootflag;
243:     struct rrip_sl_component *slc;
244:
245:     dblock = (i->inode & ~ISO9660_INODE_MASK) >> ISO9660_INODE_BITS;
246:     doffset = i->inode & ISO9660_INODE_MASK;
247:     /* FIXME: it only looks in one directory block */
248:     if(!(buf = bread(i->dev, dblock, i->sb->s_blocksize))) {
249:         return -EIO;
250:     }
251:     d = (struct iso9660_directory_record *) (buf->data + doffset);
252:
253:     ce_block = ce_offset = ce_size = 0;
254:     buf2 = NULL;
255:     total_len = isonum_711(d->length);

```

## fs/iso9660/rrip.c

Page 5/6

```

256:     len = isonum_711(d->name_len);
257:     if(!(len % 2)) {
258:         len++;
259:     }
260:     sue = (unsigned char *)d->name;
261:     nm_len = 0;
262:
263: loop:
264:     if(ce_block && ce_size) {
265:         /* FIXME: it only looks in one directory block */
266:         if(!(buf2 = bread(i->dev, ce_block, i->sb->s_blocksize))) {
267:             return 0;
268:         }
269:         sue = (unsigned char *)buf2->data + ce_offset;
270:         total_len = ce_size;
271:         len = 0;
272:     }
273:
274:     while(len < total_len) {
275:         rrip = (struct susp_rrip *) (sue + len);
276:         if(rrip->len == 0) {
277:             break;
278:         }
279:         sig = GET_SIG(rrip->signature[0], rrip->signature[1]);
280:         switch(sig) {
281:             case GET_SIG('S', 'P'):
282:                 if(rrip->u.sp.magic[0] != SP_MAGIC1 || rrip->u.s
p.magic[1] != SP_MAGIC2) {
283:                     if(ce_block) {
284:                         brelse(buf2);
285:                     }
286:                     return 0;
287:                 }
288:                 break;
289:             case GET_SIG('C', 'E'):
290:                 if(ce_block) {
291:                     brelse(buf2);
292:                 }
293:                 ce_block = isonum_733(rrip->u.ce.block);
294:                 ce_offset = isonum_733(rrip->u.ce.offset);
295:                 ce_size = isonum_733(rrip->u.ce.size);
296:                 goto loop;
297:             case GET_SIG('S', 'L'):
298:                 sl_len = rootflag = 0;
299:                 slc = (struct rrip_sl_component *)&rrip->u.sl.ar
ea;
300:                 while(sl_len < (rrip->len - 5)) {
301:                     if(sl_len && !rootflag) {
302:                         strcat(name, "/");
303:                         nm_len++;
304:                     }
305:                     rootflag = 0;
306:                     switch(slc->flags & 0xE) {
307:                         case 0:
308:                             nm_len += slc->len;
309:                             strncat(name, slc->name,
slc->len);
310:                             break;
311:                         case SL_CURRENT:
312:                             nm_len += 1;
313:                             strcat(name, ".");
314:                             break;
315:                         case SL_PARENT:
316:                             nm_len += 2;
317:                             strcat(name, "..");
318:                             break;
319:                         case SL_ROOT:

```

## fs/iso9660/rrip.c

Page 6/6

```
320:                                     nm_len += 1;
321:                                     rootflag = 1;
322:                                     strcat(name, "/");
323:                                     break;
324:                                     default:
325:                                     printk("WARNING: %s(): u
nsupported RRIP SL flags %d.\n", __FUNCTION__, slc->flags & 0xE);
326:                                     }
327:                                     slc = (struct rrip_sl_component *)(((cha
r *)slc) + slc->len + sizeof(struct rrip_sl_component));
328:                                     sl_len += slc->len + sizeof(struct rrip_
sl_component);
329:                                     }
330:                                     name[nm_len] = NULL;
331:                                     break;
332:                                     }
333:                                     len += rrip->len;
334:                                     }
335:                                     if(ce_block) {
336:                                     brelse(buf2);
337:                                     }
338:                                     brelse(buf);
339:                                     return nm_len;
340: }
```

## fs/iso9660/super.c

Page 1/4

```

1:  /*
2:  *  fiwix/fs/iso9660/super.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/types.h>
9:  #include <fiwix/errno.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_iso9660.h>
13: #include <fiwix/buffer.h>
14: #include <fiwix/time.h>
15: #include <fiwix/sched.h>
16: #include <fiwix/mm.h>
17: #include <fiwix/stdio.h>
18: #include <fiwix/string.h>
19:
20: struct fs_operations iso9660_fsop = {
21:     FSOP_REQUIRES_DEV,
22:     NULL,
23:
24:     NULL,          /* open */
25:     NULL,          /* close */
26:     NULL,          /* read */
27:     NULL,          /* write */
28:     NULL,          /* ioctl */
29:     NULL,          /* lseek */
30:     NULL,          /* readdir */
31:     NULL,          /* mmap */
32:     NULL,          /* select */
33:
34:     NULL,          /* readlink */
35:     NULL,          /* followlink */
36:     NULL,          /* bmap */
37:     NULL,          /* lookup */
38:     NULL,          /* rmdir */
39:     NULL,          /* link */
40:     NULL,          /* unlink */
41:     NULL,          /* symlink */
42:     NULL,          /* mkdir */
43:     NULL,          /* mknod */
44:     NULL,          /* truncate */
45:     NULL,          /* create */
46:     NULL,          /* rename */
47:
48:     NULL,          /* read_block */
49:     NULL,          /* write_block */
50:
51:     iso9660_read_inode,
52:     NULL,          /* write_inode */
53:     NULL,          /* ialloc */
54:     NULL,          /* ifree */
55:     iso9660_statfs,
56:     iso9660_read_superblock,
57:     NULL,          /* remount_fs */
58:     NULL,          /* write_superblock */
59:     iso9660_release_superblock
60: };
61:
62: int isonum_711(char *str)
63: {
64:     unsigned char *le;
65:
66:     le = (unsigned char *)str;
67:     return le[0];

```

## fs/iso9660/super.c

Page 2/4

```

68: }
69:
70: /* return a 16bit little-endian number */
71: int isonum_723(char *str)
72: {
73:     unsigned char *le;
74:
75:     le = (unsigned char *)str;
76:     return le[0] | (le[1] << 8);
77: }
78:
79: /* return a 32bit little-endian number */
80: int isonum_731(char *str)
81: {
82:     unsigned char *le;
83:
84:     le = (unsigned char *)str;
85:     return le[0] | (le[1] << 8) | (le[2] << 16) | (le[3] << 24);
86: }
87:
88: /* return a 32bit little-endian number */
89: int isonum_733(char *p)
90: {
91:     return isonum_731(p);
92: }
93:
94: /* return a date and time format */
95: unsigned long int isodate(char *p)
96: {
97:     struct mt mt;
98:
99:     if(!p[0]) {
100:         return 0;
101:     }
102:
103:     mt.mt_sec = p[5];
104:     mt.mt_min = p[4];
105:     mt.mt_hour = p[3];
106:     mt.mt_day = p[2];
107:     mt.mt_month = p[1];
108:     mt.mt_year = p[0];
109:     mt.mt_year += 1900;
110:     mt.mt_min += p[6] * 15;
111:
112:     return mktime(&mt);
113: }
114:
115: /* return a clean filename */
116: int iso9660_cleanfilename(char *filename, int len)
117: {
118:     int n;
119:     char *p;
120:
121:     p = filename;
122:     if(len > 2) {
123:         for(n = 0; n < len; n++) {
124:             if((len - n) == 2) {
125:                 if(p[n] == ';' && p[n + 1] == '1') {
126:                     filename[n] = NULL;
127:                     if(p[n - 1] == '.') {
128:                         filename[n - 1] = NULL;
129:                     }
130:                 }
131:             }
132:         }
133:     }
134: }

```



## fs/iso9660/super.c

Page 3/4

```

135:         return 1;
136:     }
137:
138: void iso9660_statfs(struct superblock *sb, struct statfs *statfsbuf)
139: {
140:     statfsbuf->f_type = ISO9660_SUPER_MAGIC;
141:     statfsbuf->f_bsize = sb->s_blocksize;
142:     statfsbuf->f_blocks = isonum_733(sb->u.iso9660.sb->volume_space_size);
143:     statfsbuf->f_bfree = 0;
144:     statfsbuf->f_bavail = 0;
145:     statfsbuf->f_files = 0;          /* FIXME */
146:     statfsbuf->f_ffree = 0;
147:     /* statfsbuf->f_fsid = ? */
148:     statfsbuf->f_namelen = NAME_MAX;
149: }
150:
151: int iso9660_read_superblock(__dev_t dev, struct superblock *sb)
152: {
153:     struct buffer *buf;
154:     struct iso9660_super_block *iso9660sb;
155:     struct iso9660_super_block *pvd;
156:     struct iso9660_directory_record *dr;
157:     __ino_t root_inode;
158:     int n;
159:
160:     superblock_lock(sb);
161:     pvd = NULL;
162:
163:     for(n = 0; n < ISO9660_MAX_VD; n++) {
164:         if(!(buf = bread(dev, ISO9660_SUPERBLOCK + n, BLKSIZE_2K))) {
165:             superblock_unlock(sb);
166:             return -EIO;
167:         }
168:
169:         iso9660sb = (struct iso9660_super_block *)buf->data;
170:         if(strncmp(iso9660sb->id, ISO9660_STANDARD_ID, sizeof(iso9660sb-
>id)) || (isonum_711(iso9660sb->type) == ISO9660_VD_END)) {
171:             break;
172:         }
173:         if(isonum_711(iso9660sb->type) == ISO9660_VD_PRIMARY) {
174:             pvd = (struct iso9660_super_block *)buf->data;
175:             break;
176:         }
177:         brelse(buf);
178:     }
179:     if(!pvd) {
180:         printk("WARNING: %s(): invalid filesystem type or bad superblock
on device %d,%d.\n", __FUNCTION__, MAJOR(dev), MINOR(dev));
181:         superblock_unlock(sb);
182:         brelse(buf);
183:         return -EINVAL;
184:     }
185:
186:     dr = (struct iso9660_directory_record *)pvd->root_directory_record;
187:     root_inode = isonum_711(dr->extent);
188:
189:     sb->dev = dev;
190:     sb->fsop = &iso9660_fsop;
191:     sb->flags = MS_RDONLY;
192:     sb->s_blocksize = isonum_723(pvd->logical_block_size);
193:     sb->u.iso9660.rrip = 0;
194:     if(!(sb->u.iso9660.sb = (void *)kmallocc())) {
195:         superblock_unlock(sb);
196:         brelse(buf);
197:         return -ENOMEM;
198:     }
199:     memcpy_b(sb->u.iso9660.sb, pvd, sizeof(struct iso9660_super_block));

```

**fs/iso9660/super.c**

Page 4/4

```
200:         brelse(buf);
201:
202:         root_inode = (root_inode << ISO9660_INODE_BITS) + (0 & ISO9660_INODE_MAS
K);
203:         if(!(sb->root = iget(sb, root_inode))) {
204:             printk("WARNING: %s(): unable to get root inode.\n", __FUNCTION_
_);
205:             superblock_unlock(sb);
206:             return -EINVAL;
207:         }
208:         sb->u.iso9660.s_root_inode = root_inode;
209:
210:         superblock_unlock(sb);
211:         return 0;
212:     }
213:
214: void iso9660_release_superblock(struct superblock *sb)
215: {
216:     kfree((unsigned int) sb->u.iso9660.sb);
217:     kfree((unsigned int) sb->u.iso9660.pathtable);
218:     kfree((unsigned int) sb->u.iso9660.pathtable_raw);
219: }
220:
221: int iso9660_init(void)
222: {
223:     return register_filesystem("iso9660", &iso9660_fsop);
224: }
```

## fs/iso9660/symlink.c

Page 1/2

```

1: /*
2:  * fiwix/fs/iso9660/symlink.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/fs_iso9660.h>
14: #include <fiwix/stat.h>
15: #include <fiwix/mm.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: struct fs_operations iso9660_symlink_fsop = {
20:     0,
21:     0,
22:
23:     NULL, /* open */
24:     NULL, /* close */
25:     NULL, /* read */
26:     NULL, /* write */
27:     NULL, /* ioctl */
28:     NULL, /* lseek */
29:     NULL, /* readdir */
30:     NULL, /* mmap */
31:     NULL, /* select */
32:
33:     iso9660_readlink,
34:     iso9660_followlink,
35:     NULL, /* bmap */
36:     NULL, /* lookup */
37:     NULL, /* rmdir */
38:     NULL, /* link */
39:     NULL, /* unlink */
40:     NULL, /* symlink */
41:     NULL, /* mkdir */
42:     NULL, /* mknod */
43:     NULL, /* truncate */
44:     NULL, /* create */
45:     NULL, /* rename */
46:
47:     NULL, /* read_block */
48:     NULL, /* write_block */
49:
50:     NULL, /* read_inode */
51:     NULL, /* write_inode */
52:     NULL, /* ialloc */
53:     NULL, /* ifree */
54:     NULL, /* statfs */
55:     NULL, /* read_superblock */
56:     NULL, /* remount_fs */
57:     NULL, /* write_superblock */
58:     NULL, /* release_superblock */
59: };
60:
61: int iso9660_readlink(struct inode *i, char *buffer, __size_t count)
62: {
63:     __off_t size_read;
64:     char *name;
65:
66:     if (!(name = (char *)kmalloc())) {
67:         return -ENOMEM;

```

## fs/iso9660/symlink.c

Page 2/2

```

68:         }
69:
70:         inode_lock(i);
71:         name[0] = NULL;
72:         if((size_read = get_rrip_symlink(i, name))) {
73:             size_read = MIN(size_read, count);
74:             memcpy_b(buffer, name, size_read);
75:         }
76:         kfree((unsigned int)name);
77:         inode_unlock(i);
78:         return size_read;
79:     }
80:
81: int iso9660_followlink(struct inode *dir, struct inode *i, struct inode **i_res)
82: {
83:     char *name;
84:     __off_t size_read;
85:     __ino_t errno;
86:
87:     if(!i) {
88:         return -ENOENT;
89:     }
90:     if(!S_ISLNK(i->i_mode)) {
91:         printk("%s(): Oops, inode '%d' is not a symlink (!?)\n", __FUNC
TION__, i->inode);
92:         return 0;
93:     }
94:
95:     if(!(name = (char *)kmalloc())) {
96:         return -ENOMEM;
97:     }
98:
99:     name[0] = NULL;
100:    if((size_read = get_rrip_symlink(i, name))) {
101:        iput(i);
102:        if((errno = parse_namei(name, dir, i_res, NULL, FOLLOW_LINKS)))
{
103:            kfree((unsigned int)name);
104:            return errno;
105:        }
106:    }
107:    kfree((unsigned int)name);
108:    return 0;
109: }

```

## fs/minix/bitmaps.c

Page 1/3

```

1: /*
2:  * fiwix/fs/minix/bitmaps.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/fs_minix.h>
12: #include <fiwix/buffer.h>
13: #include <fiwix/errno.h>
14: #include <fiwix/stdio.h>
15: #include <fiwix/string.h>
16:
17: #define COUNT          1
18: #define FIRST_ZERO    2
19:
20: static int count_bits(struct superblock *sb, __blk_t offset, int num, int blocks
, int mode)
21: {
22:     unsigned char c;
23:     int blksize;
24:     int n, n2, last, bits, count, mapb;
25:     struct buffer *buf;
26:
27:     count = mapb = 0;
28:     blksize = sb->s_blocksize;
29:
30:     while(offset < blocks) {
31:         if(!(buf = bread(sb->dev, offset, blksize))) {
32:             return -EIO;
33:         }
34:         last = (num / 8) > blksize ? blksize : (num / 8);
35:         for(n = 0; n < last; n++) {
36:             c = (unsigned char)buf->data[n];
37:             bits = n < last ? 8 : num & 8;
38:             for(n2 = 0; n2 < bits; n2++) {
39:                 if(c & (1 << n2)) {
40:                     if(mode == COUNT) {
41:                         count++;
42:                     }
43:                 } else {
44:                     if(mode == FIRST_ZERO) {
45:                         brelse(buf);
46:                         return n2 + ((n * 8) + (mapb * b
blksize * 8));
47:                     }
48:                 }
49:             }
50:         }
51:         offset++;
52:         mapb++;
53:         num -= (blksize * 8);
54:         brelse(buf);
55:     }
56:     return count;
57: }
58:
59: int minix_change_bit(int mode, struct superblock *sb, int map, int item)
60: {
61:     int byte, bit, mask;
62:     struct buffer *buf;
63:
64:     map += item / (sb->s_blocksize * 8);
65:     byte = (item % (sb->s_blocksize * 8)) / 8;

```

## fs/minix/bitmaps.c

Page 2/3

```

66:         bit = (item % (sb->s_blocksize * 8)) % 8;
67:         mask = 1 << bit;
68:
69:         if(!(buf = bread(sb->dev, map, sb->s_blocksize))) {
70:             return -EIO;
71:         }
72:
73:         if(mode == CLEAR_BIT) {
74:             if(!(buf->data[byte] & mask)) {
75:                 brelse(buf);
76:                 return 1;
77:             }
78:             buf->data[byte] &= ~mask;
79:         }
80:         if(mode == SET_BIT) {
81:             if((buf->data[byte] & mask)) {
82:                 brelse(buf);
83:                 return 1;
84:             }
85:             buf->data[byte] |= mask;
86:         }
87:
88:         bwrite(buf);
89:         return 0;
90:     }
91:
92: int minix_balloc(struct superblock *sb)
93: {
94:     int map, block, errno;
95:
96:     superblock_lock(sb);
97:
98:     map = 1 + SUPERBLOCK + sb->u.minix.sb.s_imap_blocks;
99:
100:    if(!(block = minix_find_first_zero(sb, map, sb->u.minix.nzones, map + sb
->u.minix.sb.s_zmap_blocks))) {
101:        superblock_unlock(sb);
102:        return -ENOSPC;
103:    }
104:
105:    errno = minix_change_bit(SET_BIT, sb, map, block);
106:    block += sb->u.minix.sb.s_firstdatazone - 1;
107:
108:    if(errno) {
109:        if(errno < 0) {
110:            printk("WARNING: %s(): unable to set block %d.\n", __FUN
CTION__, block);
111:            superblock_unlock(sb);
112:            return errno;
113:        } else {
114:            printk("WARNING: %s(): block %d is already marked as use
d!\n", __FUNCTION__, block);
115:        }
116:    }
117:
118:    superblock_unlock(sb);
119:    return block;
120: }
121:
122: void minix_bfree(struct superblock *sb, int block)
123: {
124:     int map, errno;
125:
126:     if(block < sb->u.minix.sb.s_firstdatazone || block > sb->u.minix.nzones)
{
127:         printk("WARNING: %s(): block %d is not in datazone.\n", __FUNCTI
ON__, block);

```

## fs/minix/bitmaps.c

Page 3/3

```

128:         return;
129:     }
130:
131:     superblock_lock(sb);
132:
133:     map = 1 + SUPERBLOCK + sb->u.minix.sb.s_imap_blocks;
134:     block -= sb->u.minix.sb.s_firstdatazone - 1;
135:     errno = minix_change_bit(CLEAR_BIT, sb, map, block);
136:
137:     if(errno) {
138:         if(errno < 0) {
139:             printk("WARNING: %s(): unable to free block %d.\n", __FU
NCTION__, block);
140:         } else {
141:             printk("WARNING: %s(): block %d is already marked as fre
e!\n", __FUNCTION__, block);
142:         }
143:     }
144:
145:     superblock_unlock(sb);
146:     return;
147: }
148:
149: int minix_count_free_inodes(struct superblock *sb)
150: {
151:     __blk_t offset;
152:
153:     offset = 1 + SUPERBLOCK;
154:     return count_bits(sb, offset, sb->u.minix.sb.s_ninodes, offset + sb->u.m
inix.sb.s_imap_blocks, COUNT);
155: }
156:
157: int minix_count_free_blocks(struct superblock *sb)
158: {
159:     __blk_t offset;
160:
161:     offset = 1 + SUPERBLOCK + sb->u.minix.sb.s_imap_blocks;
162:     return count_bits(sb, offset, sb->u.minix.nzones, offset + sb->u.minix.s
b.s_zmap_blocks, COUNT);
163: }
164:
165: int minix_find_first_zero(struct superblock *sb, __blk_t offset, int num, int bl
ocks)
166: {
167:     return count_bits(sb, offset, num, blocks, FIRST_ZERO);
168: }

```

## fs/minix/dir.c

Page 1/3

```

1: /*
2:  * fiwix/fs/minix/dir.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/dirent.h>
15: #include <fiwix/mm.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: struct fs_operations minix_dir_fsop = {
20:     0,
21:     0,
22:
23:     minix_dir_open,
24:     minix_dir_close,
25:     minix_dir_read,
26:     minix_dir_write,
27:     NULL, /* ioctl */
28:     NULL, /* lseek */
29:     minix_dir_readdir,
30:     NULL, /* mmap */
31:     NULL, /* select */
32:
33:     NULL, /* readlink */
34:     NULL, /* followlink */
35:     minix_bmap,
36:     minix_lookup,
37:     minix_rmdir,
38:     minix_link,
39:     minix_unlink,
40:     minix_symlink,
41:     minix_mkdir,
42:     minix_mknod,
43:     NULL, /* truncate */
44:     minix_create,
45:     minix_rename,
46:
47:     NULL, /* read_block */
48:     NULL, /* write_block */
49:
50:     NULL, /* read_inode */
51:     NULL, /* write_inode */
52:     NULL, /* ialloc */
53:     NULL, /* ifree */
54:     NULL, /* statfs */
55:     NULL, /* read_superblock */
56:     NULL, /* remount_fs */
57:     NULL, /* write_superblock */
58:     NULL, /* release_superblock */
59: };
60:
61: int minix_dir_open(struct inode *i, struct fd *fd_table)
62: {
63:     fd_table->offset = 0;
64:     return 0;
65: }
66:
67: int minix_dir_close(struct inode *i, struct fd *fd_table)

```



## fs/minix/dir.c

Page 2/3

```

68: {
69:     return 0;
70: }
71:
72: int minix_dir_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t
count)
73: {
74:     return -EISDIR;
75: }
76:
77: int minix_dir_write(struct inode *i, struct fd *fd_table, const char *buffer, __
size_t count)
78: {
79:     return -EBADF;
80: }
81:
82: int minix_dir_readdir(struct inode *i, struct fd *fd_table, struct dirent *diren
t, unsigned int count)
83: {
84:     __blk_t block;
85:     unsigned int doffset, offset;
86:     unsigned int size, dirent_len;
87:     struct minix_dir_entry *d;
88:     int base_dirent_len;
89:     int blksize;
90:     struct buffer *buf;
91:
92:     if(!(S_ISDIR(i->i_mode))) {
93:         return -EBADF;
94:     }
95:
96:     blksize = i->sb->s_blocksize;
97:     if(fd_table->offset > i->i_size) {
98:         fd_table->offset = i->i_size;
99:     }
100:
101:     base_dirent_len = sizeof(dirent->d_ino) + sizeof(dirent->d_off) + sizeof
(dirent->d_reclen);
102:     doffset = offset = size = 0;
103:
104:     while(doffset < count) {
105:         if((block = bmap(i, fd_table->offset, FOR_READING)) < 0) {
106:             return block;
107:         }
108:         if(block) {
109:             if(!(buf = bread(i->dev, block, blksize))) {
110:                 return -EIO;
111:             }
112:
113:             doffset = fd_table->offset;
114:             offset = fd_table->offset % blksize;
115:             while(doffset < i->i_size && offset < blksize) {
116:                 d = (struct minix_dir_entry *) (buf->data + offse
t);
117:                 if(d->inode) {
118:                     dirent_len = (base_dirent_len + (strlen(
d->name) + 1)) + 3;
119:                     dirent_len &= ~3;          /* round up */
120:                     dirent->d_ino = d->inode;
121:                     if((size + dirent_len) < count) {
122:                         dirent->d_off = doffset;
123:                         dirent->d_reclen = dirent_len;
124:                         memcpy_b(dirent->d_name, d->name
, strlen(d->name));
125:                         dirent->d_name[strlen(d->name)]
= NULL;
126:                         dirent = (struct dirent *) ((char

```

## fs/minix/dir.c

Page 3/3

```
*)dirent + dirent_len);
127:                                     size += dirent_len;
128:                                     } else {
129:                                         break;
130:                                     }
131:                                     }
132:                                     doffset += i->sb->u.minix.dirsize;
133:                                     offset += i->sb->u.minix.dirsize;
134:                                     }
135:                                     brelse(buf);
136:                                     }
137:                                     fd_table->offset &= ~(blksize - 1);
138:                                     doffset = fd_table->offset;
139:                                     fd_table->offset += offset;
140:                                     doffset += blksize;
141:                                     }
142:
143:                                     return size;
144: }
```

## fs/minix/file.c

Page 1/3

```

1: /*
2:  * fiwix/fs/minix/file.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/buffer.h>
12: #include <fiwix/fs.h>
13: #include <fiwix/filesystems.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/mman.h>
16: #include <fiwix/fcntl.h>
17: #include <fiwix/stdio.h>
18: #include <fiwix/string.h>
19:
20: struct fs_operations minix_file_fsop = {
21:     0,
22:     0,
23:
24:     minix_file_open,
25:     minix_file_close,
26:     file_read,
27:     minix_file_write,
28:     NULL, /* ioctl */
29:     minix_file_lseek,
30:     NULL, /* readdir */
31:     NULL, /* mmap */
32:     NULL, /* select */
33:
34:     NULL, /* readlink */
35:     NULL, /* followlink */
36:     minix_bmap,
37:     NULL, /* lookup */
38:     NULL, /* rmdir */
39:     NULL, /* link */
40:     NULL, /* unlink */
41:     NULL, /* symlink */
42:     NULL, /* mkdir */
43:     NULL, /* mknod */
44:     minix_truncate,
45:     NULL, /* create */
46:     NULL, /* rename */
47:
48:     NULL, /* read_block */
49:     NULL, /* write_block */
50:
51:     NULL, /* read_inode */
52:     NULL, /* write_inode */
53:     NULL, /* ialloc */
54:     NULL, /* ifree */
55:     NULL, /* statfs */
56:     NULL, /* read_superblock */
57:     NULL, /* remount_fs */
58:     NULL, /* write_superblock */
59:     NULL, /* release_superblock */
60: };
61:
62: int minix_file_open(struct inode *i, struct fd *fd_table)
63: {
64:     if(fd_table->flags & O_APPEND) {
65:         fd_table->offset = i->i_size;
66:     } else {
67:         fd_table->offset = 0;

```

## fs/minix/file.c

Page 2/3

```

68:         }
69:         if(fd_table->flags & O_TRUNC) {
70:             i->i_size = 0;
71:             minix_truncate(i, 0);
72:         }
73:         return 0;
74:     }
75:
76: int minix_file_close(struct inode *i, struct fd *fd_table)
77: {
78:     return 0;
79: }
80:
81: int minix_file_write(struct inode *i, struct fd *fd_table, const char *buffer, _
__size_t count)
82: {
83:     __blk_t block;
84:     __off_t total_written;
85:     unsigned int boffset, bytes;
86:     int blksize;
87:     struct buffer *buf;
88:
89:     inode_lock(i);
90:
91:     blksize = i->sb->s_blocksize;
92:     total_written = 0;
93:
94:     if(fd_table->flags & O_APPEND) {
95:         fd_table->offset = i->i_size;
96:     }
97:
98:     while(total_written < count) {
99:         boffset = fd_table->offset % blksize;
100:        if((block = bmap(i, fd_table->offset, FOR_WRITING)) < 0) {
101:            inode_unlock(i);
102:            return block;
103:        }
104:        bytes = blksize - boffset;
105:        bytes = MIN(bytes, (count - total_written));
106:        if(!(buf = bread(i->dev, block, blksize))) {
107:            inode_unlock(i);
108:            return -EIO;
109:        }
110:        memcpy_b(buf->data + boffset, buffer + total_written, bytes);
111:        update_page_cache(i, fd_table->offset, buffer + total_written, b
ytes);
112:        bwrite(buf);
113:        total_written += bytes;
114:        boffset += bytes;
115:        boffset %= blksize;
116:        fd_table->offset += bytes;
117:    }
118:
119:    if(fd_table->offset > i->i_size) {
120:        i->i_size = fd_table->offset;
121:    }
122:    i->i_ctime = CURRENT_TIME;
123:    i->i_mtime = CURRENT_TIME;
124:    i->dirty = 1;
125:
126:    inode_unlock(i);
127:    return total_written;
128: }
129:
130: int minix_file_lseek(struct inode *i, __off_t offset)
131: {
132:     return offset;

```

```
133: }
```

## fs/minix/inode.c

Page 1/2

```
1: /*
2:  * fiwix/fs/minix/inode.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/fs_minix.h>
12: #include <fiwix/fs_pipe.h>
13: #include <fiwix/statfs.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/stat.h>
16: #include <fiwix/sched.h>
17: #include <fiwix/buffer.h>
18: #include <fiwix/process.h>
19: #include <fiwix/errno.h>
20: #include <fiwix/stdio.h>
21: #include <fiwix/string.h>
22:
23: int minix_read_inode(struct inode *i)
24: {
25:     if(i->sb->u.minix.version == 1) {
26:         return v1_minix_read_inode(i);
27:     }
28:
29:     return v2_minix_read_inode(i);
30: }
31:
32: int minix_write_inode(struct inode *i)
33: {
34:     if(i->sb->u.minix.version == 1) {
35:         return v1_minix_write_inode(i);
36:     }
37:
38:     return v2_minix_write_inode(i);
39: }
40:
41: int minix_ialloc(struct inode *i, int mode)
42: {
43:     if(i->sb->u.minix.version == 1) {
44:         return v1_minix_ialloc(i, mode);
45:     }
46:
47:     return v2_minix_ialloc(i, mode);
48: }
49:
50: void minix_ifree(struct inode *i)
51: {
52:     if(i->sb->u.minix.version == 1) {
53:         return v1_minix_ifree(i);
54:     }
55:
56:     return v2_minix_ifree(i);
57: }
58:
59: int minix_bmap(struct inode *i, __off_t offset, int mode)
60: {
61:     if(i->sb->u.minix.version == 1) {
62:         return v1_minix_bmap(i, offset, mode);
63:     }
64:
65:     return v2_minix_bmap(i, offset, mode);
66: }
67:
```

**fs/minix/inode.c**

Page 2/2

```
68: int minix_truncate(struct inode *i, __off_t length)
69: {
70:     if(i->sb->u.minix.version == 1) {
71:         return v1_minix_truncate(i, length);
72:     }
73:
74:     return v2_minix_truncate(i, length);
75: }
```

**fs/minix/Makefile**

Page 1/1

```
1: # fiwix/fs/minix/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = super.o bitmaps.o inode.o namei.o symlink.o dir.o file.o v1_inode.o v2_in
ode.o
13:
14: minix:    $(OBJS)
15:     $(LD) $(LDFLAGS) -r $(OBJS) -o minix.o
16:
17: clean:
18:     rm -f *.o
19:
```



## fs/minix/namei.c

Page 1/11

```

1: /*
2:  * fiwix/fs/minix/namei.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_minix.h>
13: #include <fiwix/buffer.h>
14: #include <fiwix/errno.h>
15: #include <fiwix/stat.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: static int is_dir_empty(struct inode *dir)
20: {
21:     __blk_t block;
22:     unsigned int blksize;
23:     unsigned int offset, doffset;
24:     struct buffer *buf;
25:     struct minix_dir_entry *d;
26:
27:     blksize = dir->sb->s_blocksize;
28:     doffset = dir->sb->u.minix.dirsize * 2; /* accept only "." and ".." */
29:     offset = 0;
30:
31:     while(offset < dir->i_size) {
32:         if((block = bmap(dir, offset, FOR_READING)) < 0) {
33:             break;
34:         }
35:         if(block) {
36:             if(!(buf = bread(dir->dev, block, blksize))) {
37:                 break;
38:             }
39:             do {
40:                 if(doffset + offset >= dir->i_size) {
41:                     break;
42:                 }
43:                 d = (struct minix_dir_entry *) (buf->data + doffs
et);
44:                 if(d->inode) {
45:                     brelse(buf);
46:                     return 0;
47:                 }
48:                 doffset += dir->sb->u.minix.dirsize;
49:             } while(doffset < blksize);
50:             brelse(buf);
51:             offset += blksize;
52:             doffset = 0;
53:         } else {
54:             break;
55:         }
56:     }
57:
58:     return 1;
59: }
60:
61: /* finds the entry 'name' with inode 'i' in the directory 'dir' */
62: static struct buffer * find_dir_entry(struct inode *dir, struct inode *i, struct
minix_dir_entry **d_res, char *name)
63: {
64:     __blk_t block;
65:     unsigned int blksize;

```

## fs/minix/namei.c

Page 2/11

```

66:         unsigned int offset, doffset;
67:         struct buffer *buf;
68:
69:         blksize = dir->sb->s_blocksize;
70:         offset = 0;
71:
72:         while(offset < dir->i_size) {
73:             if((block = bmap(dir, offset, FOR_READING)) < 0) {
74:                 break;
75:             }
76:             if(block) {
77:                 if(!(buf = bread(dir->dev, block, blksize))) {
78:                     break;
79:                 }
80:                 doffset = 0;
81:                 do {
82:                     *d_res = (struct minix_dir_entry *) (buf->data +
doffset);
83:                     if(!i) {
84:                         /* returns the first empty entry */
85:                         if(!(*d_res)->inode || (doffset + offset
>= dir->i_size)) {
86:                             /* the directory grows by direct
ory entry size */
87:                             if(doffset + offset >= dir->i_si
ze) {
88:                                 dir->i_size += dir->sb->
u.minix.dirsize;
89:                             }
90:                             return buf;
91:                         }
92:                     } else {
93:                         if((*d_res)->inode == i->inode) {
94:                             /* returns the first matching in
ode */
95:                             if(!name) {
96:                                 return buf;
97:                             }
98:                             /* returns the matching inode an
d name */
99:                             if(!strcmp((*d_res)->name, name)
) {
100:                                 return buf;
101:                             }
102:                         }
103:                     }
104:                     doffset += dir->sb->u.minix.dirsize;
105:                 } while(doffset < blksize);
106:                 brelse(buf);
107:                 offset += blksize;
108:             } else {
109:                 break;
110:             }
111:         }
112:
113:         *d_res = NULL;
114:         return NULL;
115:     }
116:
117:     static struct buffer * add_dir_entry(struct inode *dir, struct minix_dir_entry *
*d_res)
118:     {
119:         __blk_t block;
120:         struct buffer *buf;
121:
122:         if(!(buf = find_dir_entry(dir, NULL, d_res, NULL))) {
123:             if((block = bmap(dir, dir->i_size, FOR_WRITING)) < 0) {

```

## fs/minix/namei.c

Page 3/11

```

124:             return NULL;
125:         }
126:         if(!(buf = bread(dir->dev, block, dir->sb->s_blocksize))) {
127:             return NULL;
128:         }
129:         *d_res = (struct minix_dir_entry *)buf->data;
130:         dir->i_size += dir->sb->u.minix.dirsize;
131:     }
132:
133:     return buf;
134: }
135:
136: static int is_prefix(struct inode *dir_new, struct inode *i_old)
137: {
138:     __ino_t inode;
139:     int errno;
140:
141:     errno = 0;
142:     for(;;) {
143:         if(dir_new == i_old) {
144:             errno = 1;
145:             break;
146:         }
147:         inode = dir_new->inode;
148:         if(minix_lookup("../", dir_new, &dir_new)) {
149:             break;
150:         }
151:         iput(dir_new); /* lookup eats 1 dir_new */
152:         if(dir_new->inode == inode) {
153:             break;
154:         }
155:     }
156:     return errno;
157: }
158:
159: int minix_lookup(const char *name, struct inode *dir, struct inode **i_res)
160: {
161:     __blk_t block;
162:     unsigned int blksize;
163:     unsigned int offset, doffset;
164:     struct buffer *buf;
165:     struct minix_dir_entry *d;
166:     __ino_t inode;
167:
168:     blksize = dir->sb->s_blocksize;
169:     inode = offset = 0;
170:     dir->count++;
171:
172:     while(offset < dir->i_size && !inode) {
173:         if((block = bmap(dir, offset, FOR_READING)) < 0) {
174:             iput(dir);
175:             return block;
176:         }
177:         if(block) {
178:             if(!(buf = bread(dir->dev, block, blksize))) {
179:                 iput(dir);
180:                 return -EIO;
181:             }
182:             doffset = 0;
183:             do {
184:                 d = (struct minix_dir_entry *) (buf->data + doffs
et);
185:                 if(d->inode) {
186:                     if(strlen(d->name) == strlen(name)) {
187:                         if(!(strcmp(d->name, name))) {
188:                             inode = d->inode;
189:                         }

```

## fs/minix/namei.c

Page 4/11

```

190:         }
191:     }
192:     doffset += dir->sb->u.minix.dirsize;
193: } while((doffset < blksize) && (!inode));
194:
195: brelse(buf);
196: if(inode) {
197:     if(!(*i_res = iget(dir->sb, inode))) {
198:         iput(dir);
199:         return -EACCES;
200:     }
201:     iput(dir);
202:     return 0;
203: }
204: offset += blksize;
205: } else {
206:     break;
207: }
208: }
209: iput(dir);
210: return -ENOENT;
211: }
212:
213: int minix_rmdir(struct inode *dir, struct inode *i)
214: {
215:     struct buffer *buf;
216:     struct minix_dir_entry *d;
217:
218:     inode_lock(i);
219:
220:     if(!is_dir_empty(i)) {
221:         inode_unlock(i);
222:         return -ENOTEMPTY;
223:     }
224:
225:     inode_lock(dir);
226:
227:     if(!(buf = find_dir_entry(dir, i, &d, NULL))) {
228:         inode_unlock(i);
229:         inode_unlock(dir);
230:         return -ENOENT;
231:     }
232:
233:     d->inode = 0;
234:     i->i_nlink = 0;
235:     dir->i_nlink--;
236:
237:     i->i_ctime = CURRENT_TIME;
238:     dir->i_mtime = CURRENT_TIME;
239:     dir->i_ctime = CURRENT_TIME;
240:
241:     i->dirty = 1;
242:     dir->dirty = 1;
243:
244:     bwrite(buf);
245:
246:     inode_unlock(i);
247:     inode_unlock(dir);
248:     return 0;
249: }
250:
251: int minix_link(struct inode *i_old, struct inode *dir_new, char *name)
252: {
253:     struct buffer *buf;
254:     struct minix_dir_entry *d;
255:     int n;
256:

```

## fs/minix/namei.c

Page 5/11

```

257:     inode_lock(i_old);
258:     inode_lock(dir_new);
259:
260:     if(!(buf = add_dir_entry(dir_new, &d))) {
261:         inode_unlock(i_old);
262:         inode_unlock(dir_new);
263:         return -ENOSPC;
264:     }
265:
266:     d->inode = i_old->inode;
267:     for(n = 0; n < i_old->sb->u.minix.namelen; n++) {
268:         d->name[n] = name[n];
269:         if(!name[n]) {
270:             break;
271:         }
272:     }
273:     for(; n < i_old->sb->u.minix.namelen; n++) {
274:         d->name[n] = 0;
275:     }
276:
277:     i_old->i_nlink++;
278:     i_old->i_ctime = CURRENT_TIME;
279:     dir_new->i_mtime = CURRENT_TIME;
280:     dir_new->i_ctime = CURRENT_TIME;
281:
282:     i_old->dirty = 1;
283:     dir_new->dirty = 1;
284:
285:     bwrite(buf);
286:
287:     inode_unlock(i_old);
288:     inode_unlock(dir_new);
289:     return 0;
290: }
291:
292: int minix_unlink(struct inode *dir, struct inode *i, char *name)
293: {
294:     struct buffer *buf;
295:     struct minix_dir_entry *d;
296:
297:     inode_lock(dir);
298:     inode_lock(i);
299:
300:     if(!(buf = find_dir_entry(dir, i, &d, name))) {
301:         inode_unlock(dir);
302:         inode_unlock(i);
303:         return -ENOENT;
304:     }
305:
306:     d->inode = 0;
307:     i->i_nlink--;
308:
309:     i->i_ctime = CURRENT_TIME;
310:     dir->i_mtime = CURRENT_TIME;
311:     dir->i_ctime = CURRENT_TIME;
312:
313:     i->dirty = 1;
314:     dir->dirty = 1;
315:
316:     bwrite(buf);
317:
318:     inode_unlock(dir);
319:     inode_unlock(i);
320:     return 0;
321: }
322:
323: int minix_symlink(struct inode *dir, char *name, char *oldname)

```

## fs/minix/namei.c

Page 6/11

```

324: {
325:     struct buffer *buf, *buf_new;
326:     struct inode *i;
327:     struct minix_dir_entry *d;
328:     unsigned int blksize;
329:     int n, block;
330:     char c;
331:
332:     inode_lock(dir);
333:
334:     if(!(i = ialloc(dir->sb, S_IFLNK))) {
335:         inode_unlock(dir);
336:         return -ENOSPC;
337:     }
338:
339:     i->i_mode = S_IFLNK | (S_IRWXU | S_IRWXG | S_IRWXO);
340:     i->i_uid = current->euid;
341:     i->i_gid = current->egid;
342:     i->i_nlink = 1;
343:     i->dev = dir->dev;
344:     i->count = 1;
345:     i->fsop = &minix_symlink_fsop;
346:     i->dirty = 1;
347:
348:     block = minix_balloc(dir->sb);
349:     if(block < 0) {
350:         i->i_nlink = 0;
351:         iput(i);
352:         inode_unlock(dir);
353:         return -ENOSPC;
354:     }
355:
356:     if(i->sb->u.minix.version == 1) {
357:         i->u.minix.u.i1_zone[0] = block;
358:     } else {
359:         i->u.minix.u.i2_zone[0] = block;
360:     }
361:     blksize = dir->sb->s_blocksize;
362:     if(!(buf_new = bread(dir->dev, block, blksize))) {
363:         minix_bfree(dir->sb, block);
364:         i->i_nlink = 0;
365:         iput(i);
366:         inode_unlock(dir);
367:         return -EIO;
368:     }
369:
370:     if(!(buf = add_dir_entry(dir, &d))) {
371:         minix_bfree(dir->sb, block);
372:         i->i_nlink = 0;
373:         iput(i);
374:         inode_unlock(dir);
375:         return -ENOSPC;
376:     }
377:
378:     d->inode = i->inode;
379:     for(n = 0; n < i->sb->u.minix.namelen; n++) {
380:         d->name[n] = name[n];
381:         if(!name[n]) {
382:             break;
383:         }
384:     }
385:     for(; n < i->sb->u.minix.namelen; n++) {
386:         d->name[n] = 0;
387:     }
388:
389:     for(n = 0; n < NAME_MAX; n++) {
390:         if((c = oldname[n])) {

```

## fs/minix/namei.c

Page 7/11

```

391:             buf_new->data[n] = c;
392:             continue;
393:         }
394:         break;
395:     }
396:     buf_new->data[n] = 0;
397:     i->i_size = n;
398:
399:     dir->i_mtime = CURRENT_TIME;
400:     dir->i_ctime = CURRENT_TIME;
401:     dir->dirty = 1;
402:
403:     bwrite(buf);
404:     bwrite(buf_new);
405:     iput(i);
406:     inode_unlock(dir);
407:     return 0;
408: }
409:
410: int minix_mkdir(struct inode *dir, char *name, __mode_t mode)
411: {
412:     struct buffer *buf, *buf_new;
413:     struct inode *i;
414:     struct minix_dir_entry *d, *d_new;
415:     unsigned int blksize;
416:     int n, block;
417:
418:     if(strlen(name) > dir->sb->u.minix.namelen) {
419:         return -ENAMETOOLONG;
420:     }
421:
422:     inode_lock(dir);
423:
424:     if(!(i = ialloc(dir->sb, S_IFDIR))) {
425:         inode_unlock(dir);
426:         return -ENOSPC;
427:     }
428:
429:     i->i_mode = ((mode & (S_IRWXU | S_IRWXG | S_IRWXO)) & ~current->umask);
430:     i->i_mode |= S_IFDIR;
431:     i->i_uid = current->euid;
432:     i->i_gid = current->egid;
433:     i->i_nlink = 1;
434:     i->dev = dir->dev;
435:     i->count = 1;
436:     i->fsop = &minix_dir_fsop;
437:     i->dirty = 1;
438:
439:     if((block = bmap(i, 0, FOR_WRITING)) < 0) {
440:         i->i_nlink = 0;
441:         iput(i);
442:         inode_unlock(dir);
443:         return -ENOSPC;
444:     }
445:
446:     blksize = dir->sb->s_blocksize;
447:     if(!(buf_new = bread(i->dev, block, blksize))) {
448:         minix_bfree(dir->sb, block);
449:         i->i_nlink = 0;
450:         iput(i);
451:         inode_unlock(dir);
452:         return -EIO;
453:     }
454:
455:     if(!(buf = add_dir_entry(dir, &d))) {
456:         minix_bfree(dir->sb, block);
457:         i->i_nlink = 0;

```

## fs/minix/namei.c

Page 8/11

```

458:         iput(i);
459:         inode_unlock(dir);
460:         return -ENOSPC;
461:     }
462:
463:     d->inode = i->inode;
464:     for(n = 0; n < i->sb->u.minix.namelen; n++) {
465:         d->name[n] = name[n];
466:         if(!name[n] || name[n] == '/') {
467:             break;
468:         }
469:     }
470:     for(; n < i->sb->u.minix.namelen; n++) {
471:         d->name[n] = 0;
472:     }
473:
474:     d_new = (struct minix_dir_entry *)buf_new->data;
475:     d_new->inode = i->inode;
476:     d_new->name[0] = '.';
477:     d_new->name[1] = 0;
478:     i->i_size += i->sb->u.minix.dirsize;
479:     i->i_nlink++;
480:     d_new = (struct minix_dir_entry *) (buf_new->data + i->sb->u.minix.dirsiz
e);
481:     d_new->inode = dir->inode;
482:     d_new->name[0] = '.';
483:     d_new->name[1] = '.';
484:     d_new->name[2] = 0;
485:     i->i_size += i->sb->u.minix.dirsize;
486:
487:     dir->i_mtime = CURRENT_TIME;
488:     dir->i_ctime = CURRENT_TIME;
489:     dir->i_nlink++;
490:     dir->dirty = 1;
491:
492:     bwrite(buf);
493:     bwrite(buf_new);
494:     iput(i);
495:     inode_unlock(dir);
496:     return 0;
497: }
498:
499: int minix_mknod(struct inode *dir, char *name, __mode_t mode, __dev_t dev)
500: {
501:     struct buffer *buf;
502:     struct inode *i;
503:     struct minix_dir_entry *d;
504:     int n;
505:
506:     inode_lock(dir);
507:
508:     if(!(i = ialloc(dir->sb, mode & S_IFMT))) {
509:         inode_unlock(dir);
510:         return -ENOSPC;
511:     }
512:
513:     if(!(buf = add_dir_entry(dir, &d))) {
514:         i->i_nlink = 0;
515:         iput(i);
516:         inode_unlock(dir);
517:         return -ENOSPC;
518:     }
519:
520:     d->inode = i->inode;
521:     for(n = 0; n < i->sb->u.minix.namelen; n++) {
522:         d->name[n] = name[n];
523:         if(!name[n]) {

```



## fs/minix/namei.c

Page 9/11

```

524:             break;
525:         }
526:     }
527:     for(; n < i->sb->u.minix.namelen; n++) {
528:         d->name[n] = 0;
529:     }
530:
531:     i->i_mode = (mode & ~current->umask) & ~S_IFMT;
532:     i->i_uid = current->euid;
533:     i->i_gid = current->egid;
534:     i->i_nlink = 1;
535:     i->dev = dir->dev;
536:     i->count = 1;
537:     i->dirty = 1;
538:
539:     switch(mode & S_IFMT) {
540:         case S_IFCHR:
541:             i->fsop = &def_chr_fsop;
542:             i->rdev = dev;
543:             i->i_mode |= S_IFCHR;
544:             break;
545:         case S_IFBLK:
546:             i->fsop = &def_blk_fsop;
547:             i->rdev = dev;
548:             i->i_mode |= S_IFBLK;
549:             break;
550:         case S_IFIFO:
551:             i->fsop = &pipefs_fsop;
552:             i->i_mode |= S_IFIFO;
553:             /* it's a union so we need to clear pipefs_i */
554:             memset_b(&i->u.pipefs, NULL, sizeof(struct pipefs_inode)
);
555:             break;
556:     }
557:
558:     dir->i_mtime = CURRENT_TIME;
559:     dir->i_ctime = CURRENT_TIME;
560:     dir->dirty = 1;
561:
562:     bwrite(buf);
563:     iput(i);
564:     inode_unlock(dir);
565:     return 0;
566: }
567:
568: int minix_create(struct inode *dir, char *name, __mode_t mode, struct inode **i_
res)
569: {
570:     struct buffer *buf;
571:     struct inode *i;
572:     struct minix_dir_entry *d;
573:     int n;
574:
575:     if(IS_RDONLY_FS(dir)) {
576:         return -EROFS;
577:     }
578:
579:     inode_lock(dir);
580:
581:     if(!(i = ialloc(dir->sb, S_IFREG))) {
582:         inode_unlock(dir);
583:         return -ENOSPC;
584:     }
585:
586:     if(!(buf = add_dir_entry(dir, &d))) {
587:         i->i_nlink = 0;
588:         iput(i);

```

## fs/minix/namei.c

Page 10/11

```

589:         inode_unlock(dir);
590:         return -ENOSPC;
591:     }
592:
593:     d->inode = i->inode;
594:     for(n = 0; n < i->sb->u.minix.namelen; n++) {
595:         d->name[n] = name[n];
596:         if(!name[n]) {
597:             break;
598:         }
599:     }
600:     for(; n < i->sb->u.minix.namelen; n++) {
601:         d->name[n] = 0;
602:     }
603:
604:     i->i_mode = (mode & ~current->umask) & ~S_IFMT;
605:     i->i_mode |= S_IFREG;
606:     i->i_uid = current->euid;
607:     i->i_gid = current->egid;
608:     i->i_nlink = 1;
609:     i->dev = dir->dev;
610:     i->fsop = &minix_file_fsop;
611:     i->count = 1;
612:     i->dirty = 1;
613:
614:     dir->i_mtime = CURRENT_TIME;
615:     dir->i_ctime = CURRENT_TIME;
616:     dir->dirty = 1;
617:
618:     *i_res = i;
619:     bwrite(buf);
620:     inode_unlock(dir);
621:     return 0;
622: }
623:
624: int minix_rename(struct inode *i_old, struct inode *dir_old, struct inode *i_new
, struct inode *dir_new, char *oldpath, char *newpath)
625: {
626:     struct buffer *buf_old, *buf_new;
627:     struct minix_dir_entry *d_old, *d_new;
628:     int errno;
629:
630:     errno = 0;
631:     buf_new = NULL;
632:
633:     if(is_prefix(dir_new, i_old)) {
634:         return -EINVAL;
635:     }
636:
637:     inode_lock(i_old);
638:     inode_lock(dir_old);
639:     if(dir_old != dir_new) {
640:         inode_lock(dir_new);
641:     }
642:
643:     if(!(buf_old = find_dir_entry(dir_old, i_old, &d_old, oldpath))) {
644:         errno = -ENOENT;
645:         goto end;
646:     }
647:     if(dir_old == dir_new) {
648:         buf_old->locked = 0;
649:     }
650:
651:     if(i_new) {
652:         if(S_ISDIR(i_old->i_mode)) {
653:             if(!is_dir_empty(i_new)) {
654:                 brelse(buf_old);

```

## fs/minix/namei.c

Page 11/11

```
655:             errno = -ENOTEMPTY;
656:             goto end;
657:         }
658:     }
659:     if(!(buf_new = find_dir_entry(dir_new, i_new, &d_new, newpath)))
{
660:         brelse(buf_old);
661:         errno = -ENOENT;
662:         goto end;
663:     }
664: } else {
665:     if(!(buf_new = add_dir_entry(dir_new, &d_new))) {
666:         brelse(buf_old);
667:         errno = -ENOSPC;
668:         goto end;
669:     }
670: }
671: if(i_new) {
672:     i_new->i_nlink--;
673: } else {
674:     i_new = i_old;
675:     strcpy(d_new->name, newpath);
676: }
677:
678: d_old->inode = 0;
679: d_new->inode = i_old->inode;
680: dir_new->i_mtime = CURRENT_TIME;
681: dir_new->i_ctime = CURRENT_TIME;
682: i_new->dirty = 1;
683: dir_new->dirty = 1;
684:
685: dir_old->i_mtime = CURRENT_TIME;
686: dir_old->i_ctime = CURRENT_TIME;
687: i_old->dirty = 1;
688: dir_old->dirty = 1;
689:
690: bwrite(buf_old);
691: if(buf_new) {
692:     bwrite(buf_new);
693: }
694:
695: end:
696:     inode_unlock(i_old);
697:     inode_unlock(dir_old);
698:     inode_unlock(dir_new);
699:     return errno;
700: }
```

## fs/minix/super.c

Page 1/5

```

1: /*
2:  * fiwix/fs/minix/super.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/types.h>
10: #include <fiwix/errno.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/fs_minix.h>
14: #include <fiwix/buffer.h>
15: #include <fiwix/sched.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: struct fs_operations minix_fsop = {
20:     FSOP_REQUIRES_DEV,
21:     NULL,
22:
23:     NULL, /* open */
24:     NULL, /* close */
25:     NULL, /* read */
26:     NULL, /* write */
27:     NULL, /* ioctl */
28:     NULL, /* lseek */
29:     NULL, /* readdir */
30:     NULL, /* mmap */
31:     NULL, /* select */
32:
33:     NULL, /* readlink */
34:     NULL, /* followlink */
35:     NULL, /* bmap */
36:     NULL, /* lookup */
37:     NULL, /* rmdir */
38:     NULL, /* link */
39:     NULL, /* unlink */
40:     NULL, /* symlink */
41:     NULL, /* mkdir */
42:     NULL, /* mknod */
43:     NULL, /* truncate */
44:     NULL, /* create */
45:     NULL, /* rename */
46:
47:     NULL, /* read_block */
48:     NULL, /* write_block */
49:
50:     minix_read_inode,
51:     minix_write_inode,
52:     minix_ialloc,
53:     minix_ifree,
54:     minix_statfs,
55:     minix_read_superblock,
56:     minix_remount_fs,
57:     minix_write_superblock,
58:     minix_release_superblock
59: };
60:
61: static void check_superblock(struct minix_super_block *sb)
62: {
63:     if(!(sb->s_state & MINIX_VALID_FS)) {
64:         printk("WARNING: filesystem not checked, fsck recommended.\n");
65:     }
66:     if(sb->s_state & MINIX_ERROR_FS) {
67:         printk("WARNING: filesystem contains errors, fsck recommended.\n

```

## fs/minix/super.c

Page 2/5

```

");
68:         }
69:     }
70:
71: void minix_statfs(struct superblock *sb, struct statfs *statfsbuf)
72: {
73:     statfsbuf->f_type = sb->u.minix.sb.s_magic;
74:     statfsbuf->f_bsize = sb->s_blocksize;
75:     statfsbuf->f_blocks = sb->u.minix.nzones << sb->u.minix.sb.s_log_zone_si
ze;
76:     statfsbuf->f_bfree = sb->u.minix.nzones - minix_count_free_blocks(sb);
77:     statfsbuf->f_bavail = statfsbuf->f_bfree;
78:
79:     statfsbuf->f_files = sb->u.minix.sb.s_ninodes;
80:     statfsbuf->f_ffree = sb->u.minix.sb.s_ninodes - minix_count_free_inodes(
sb);
81:     /* statfsbuf->f_fsid = ? */
82:     statfsbuf->f_namelen = sb->u.minix.namelen;
83: }
84:
85: int minix_read_superblock(__dev_t dev, struct superblock *sb)
86: {
87:     struct buffer *buf;
88:     int maps;
89:
90:     superblock_lock(sb);
91:     if(!(buf = bread(dev, SUPERBLOCK, BLKSIZE_1K))) {
92:         printk("WARNING: %s(): I/O error on device %d,%d.\n", __FUNCTION
__, MAJOR(dev), MINOR(dev));
93:         superblock_unlock(sb);
94:         return -EIO;
95:     }
96:     memcpy_b(&sb->u.minix.sb, buf->data, sizeof(struct minix_super_block));
97:
98:     switch(sb->u.minix.sb.s_magic) {
99:         case MINIX_SUPER_MAGIC:
100:             sb->u.minix.namelen = 14;
101:             sb->u.minix.dirsize = sizeof(__u16) + sb->u.minix.namele
n;
102:             sb->u.minix.version = 1;
103:             sb->u.minix.nzones = sb->u.minix.sb.s_nzones;
104:             printk("minix v1 (14 char names) filesystem detected on
device %d,%d.\n", MAJOR(dev), MINOR(dev));
105:             break;
106:         case MINIX_SUPER_MAGIC2:
107:             sb->u.minix.namelen = 30;
108:             sb->u.minix.dirsize = sizeof(__u16) + sb->u.minix.namele
n;
109:             sb->u.minix.version = 1;
110:             sb->u.minix.nzones = sb->u.minix.sb.s_nzones;
111:             printk("minix v1 (30 char names) filesystem detected on
device %d,%d.\n", MAJOR(dev), MINOR(dev));
112:             break;
113:         case MINIX2_SUPER_MAGIC:
114:             sb->u.minix.namelen = 14;
115:             sb->u.minix.dirsize = sizeof(__u16) + sb->u.minix.namele
n;
116:             sb->u.minix.version = 2;
117:             sb->u.minix.nzones = sb->u.minix.sb.s_nzones;
118:             printk("minix v2 (14 char names) filesystem detected on
device %d,%d.\n", MAJOR(dev), MINOR(dev));
119:             break;
120:         case MINIX2_SUPER_MAGIC2:
121:             sb->u.minix.namelen = 30;
122:             sb->u.minix.dirsize = sizeof(__u16) + sb->u.minix.namele
n;
123:             sb->u.minix.version = 2;

```

## fs/minix/super.c

Page 3/5

```

124:             sb->u.minix.nzones = sb->u.minix.sb.s_zones;
125:             printk("minix v2 (30 char names) filesystem detected on
device %d,%d.\n", MAJOR(dev), MINOR(dev));
126:             break;
127:             default:
128:             printk("ERROR: %s(): invalid filesystem type or bad supe
rblock on device %d,%d.\n", __FUNCTION__, MAJOR(dev), MINOR(dev));
129:             superblock_unlock(sb);
130:             brelse(buf);
131:             return -EINVAL;
132:         }
133:
134:         sb->dev = dev;
135:         sb->fsop = &minix_fsop;
136:         sb->s_blocksize = BLKSIZE_1K << sb->u.minix.sb.s_log_zone_size;
137:
138:         if(sb->s_blocksize != BLKSIZE_1K) {
139:             printk("ERROR: %s(): block sizes > %d not supported in this file
system.\n", __FUNCTION__, BLKSIZE_1K);
140:             superblock_unlock(sb);
141:             brelse(buf);
142:             return -EINVAL;
143:         }
144:
145:         /*
146:         printk("s_ninodes          = %d\n", sb->u.minix.sb.s_ninodes);
147:         printk("s_nzones          = %d (nzones = %d)\n", sb->u.minix.sb.s_nzones,
sb->u.minix.nzones);
148:         printk("s_imap_blocks     = %d\n", sb->u.minix.sb.s_imap_blocks);
149:         printk("s_zmap_blocks     = %d\n", sb->u.minix.sb.s_zmap_blocks);
150:         printk("s_firstdatazone = %d\n", sb->u.minix.sb.s_firstdatazone);
151:         printk("s_log_zone_size  = %d\n", sb->u.minix.sb.s_log_zone_size);
152:         printk("s_max_size       = %d\n", sb->u.minix.sb.s_max_size);
153:         printk("s_magic        = %x\n", sb->u.minix.sb.s_magic);
154:         printk("s_state        = %d\n", sb->u.minix.sb.s_state);
155:         printk("s_zones        = %d\n", sb->u.minix.sb.s_zones);
156:         */
157:
158:         /* Minix fs size is limited to: # of bitmaps * 8192 * 1024 */
159:         if(sb->u.minix.version == 1) {
160:             maps = V1_MAX_BITMAP_BLOCKS;    /* 64MB limit */
161:         }
162:         if(sb->u.minix.version == 2) {
163:             maps = V2_MAX_BITMAP_BLOCKS;    /* 1GB limit */
164:         }
165:
166:         if(sb->u.minix.sb.s_imap_blocks > maps) {
167:             printk("ERROR: %s(): number of imap blocks (%d) is greater than
%d.\n", __FUNCTION__, sb->u.minix.sb.s_imap_blocks, maps);
168:             superblock_unlock(sb);
169:             brelse(buf);
170:             return -EINVAL;
171:         }
172:         if(sb->u.minix.sb.s_zmap_blocks > maps) {
173:             printk("ERROR: %s(): number of zmap blocks (%d) is greater than
%d.\n", __FUNCTION__, sb->u.minix.sb.s_zmap_blocks, maps);
174:             superblock_unlock(sb);
175:             brelse(buf);
176:             return -EINVAL;
177:         }
178:
179:         superblock_unlock(sb);
180:
181:         if(!(sb->root = iget(sb, MINIX_ROOT_INO))) {
182:             printk("ERROR: %s(): unable to get root inode.\n", __FUNCTION__
;
183:             brelse(buf);

```

## fs/minix/super.c

Page 4/5

```

184:         return -EINVAL;
185:     }
186:
187:     check_superblock(&sb->u.minix.sb);
188:
189:     if(!(sb->flags & MS_RDONLY)) {
190:         sb->u.minix.sb.s_state &= ~MINIX_VALID_FS;
191:         memcpy_b(buf->data, &sb->u.minix.sb, sizeof(struct minix_super_b
lock));
192:         bwrite(buf);
193:     } else {
194:         brelse(buf);
195:     }
196:
197:     return 0;
198: }
199:
200: int minix_remount_fs(struct superblock *sb, int flags)
201: {
202:     struct buffer *buf;
203:     struct minix_super_block *ms;
204:
205:     if((flags & MS_RDONLY) == (sb->flags & MS_RDONLY)) {
206:         return 0;
207:     }
208:
209:     superblock_lock(sb);
210:     if(!(buf = bread(sb->dev, SUPERBLOCK, BLKSIZE_1K))) {
211:         superblock_unlock(sb);
212:         return -EIO;
213:     }
214:     ms = (struct minix_super_block *)buf->data;
215:
216:     if(flags & MS_RDONLY) {
217:         /* switching from RW to RO */
218:         sb->u.minix.sb.s_state |= MINIX_VALID_FS;
219:         ms->s_state |= MINIX_VALID_FS;
220:     } else {
221:         /* switching from RO to RW */
222:         check_superblock(ms);
223:         sb->u.minix.sb.s_state &= ~MINIX_VALID_FS;
224:         ms->s_state &= ~MINIX_VALID_FS;
225:     }
226:
227:     sb->dirty = 1;
228:     superblock_unlock(sb);
229:     bwrite(buf);
230:     return 0;
231: }
232:
233: int minix_write_superblock(struct superblock *sb)
234: {
235:     struct buffer *buf;
236:
237:     superblock_lock(sb);
238:     if(!(buf = bread(sb->dev, SUPERBLOCK, BLKSIZE_1K))) {
239:         superblock_unlock(sb);
240:         return -EIO;
241:     }
242:
243:     memcpy_b(buf->data, &sb->u.minix.sb, sizeof(struct minix_super_block));
244:     sb->dirty = 0;
245:     superblock_unlock(sb);
246:     bwrite(buf);
247:     return 0;
248: }
249:

```

**fs/minix/super.c**

Page 5/5

```
250: void minix_release_superblock(struct superblock *sb)
251: {
252:     if(sb->flags & MS_RDONLY) {
253:         return;
254:     }
255:
256:     superblock_lock(sb);
257:
258:     sb->u.minix.sb.s_state |= MINIX_VALID_FS;
259:     sb->dirty = 1;
260:
261:     superblock_unlock(sb);
262: }
263:
264: int minix_init(void)
265: {
266:     return register_filesystem("minix", &minix_fsop);
267: }
```



## fs/minix/symlink.c

Page 1/3

```

1:  /*
2:  *  fiwix/fs/minix/symlink.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/types.h>
9:  #include <fiwix/errno.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/stdio.h>
15: #include <fiwix/string.h>
16:
17: struct fs_operations minix_symlink_fsop = {
18:     0,
19:     0,
20:
21:     NULL,                /* open */
22:     NULL,                /* close */
23:     NULL,                /* read */
24:     NULL,                /* write */
25:     NULL,                /* ioctl */
26:     NULL,                /* lseek */
27:     NULL,                /* readdir */
28:     NULL,                /* mmap */
29:     NULL,                /* select */
30:
31:     minix_readlink,
32:     minix_followlink,
33:     NULL,                /* bmap */
34:     NULL,                /* lookup */
35:     NULL,                /* rmdir */
36:     NULL,                /* link */
37:     NULL,                /* unlink */
38:     NULL,                /* symlink */
39:     NULL,                /* mkdir */
40:     NULL,                /* mknod */
41:     NULL,                /* truncate */
42:     NULL,                /* create */
43:     NULL,                /* rename */
44:
45:     NULL,                /* read_block */
46:     NULL,                /* write_block */
47:
48:     NULL,                /* read_inode */
49:     NULL,                /* write_inode */
50:     NULL,                /* ialloc */
51:     NULL,                /* ifree */
52:     NULL,                /* statfs */
53:     NULL,                /* read_superblock */
54:     NULL,                /* remount_fs */
55:     NULL,                /* write_superblock */
56:     NULL,                /* release_superblock */
57: };
58:
59: int minix_readlink(struct inode *i, char *buffer, __size_t count)
60: {
61:     __u32 blksize;
62:     struct buffer *buf;
63:
64:     if(!S_ISLNK(i->i_mode)) {
65:         printk("%s(): Oops, inode '%d' is not a symlink (!?)\n", __FUNC
TION__, i->inode);
66:         return 0;

```

## fs/minix/symlink.c

Page 2/3

```

67:         }
68:
69:         inode_lock(i);
70:         blksize = i->sb->s_blocksize;
71:         count = MIN(count, i->i_size);
72:         if(!count) {
73:             inode_unlock(i);
74:             return 0;
75:         }
76:         count = MIN(count, blksize);
77:         if(i->sb->u.minix.version == 1) {
78:             if(!(buf = bread(i->dev, i->u.minix.u.i1_zone[0], blksize))) {
79:                 inode_unlock(i);
80:                 return -EIO;
81:             }
82:         } else {
83:             if(!(buf = bread(i->dev, i->u.minix.u.i2_zone[0], blksize))) {
84:                 inode_unlock(i);
85:                 return -EIO;
86:             }
87:         }
88:         memcpy_b(buffer, buf->data, count);
89:         brelse(buf);
90:         buffer[count] = NULL;
91:         inode_unlock(i);
92:         return count;
93:     }
94:
95: int minix_followlink(struct inode *dir, struct inode *i, struct inode **i_res)
96: {
97:     struct buffer *buf;
98:     char *name;
99:     __ino_t errno;
100:
101:     if(!i) {
102:         return -ENOENT;
103:     }
104:
105:     if(!S_ISLNK(i->i_mode)) {
106:         printk("%s(): Oops, inode '%d' is not a symlink (!?)\n", __FUNC
TION__, i->inode);
107:         return 0;
108:     }
109:
110:     if(current->loopcnt > MAX_SYMLINKS) {
111:         printk("%s(): too many nested symbolic links!\n", __FUNCTION__);
112:         return -ELOOP;
113:     }
114:
115:     inode_lock(i);
116:     if(i->sb->u.minix.version == 1) {
117:         if(!(buf = bread(i->dev, i->u.minix.u.i1_zone[0], i->sb->s_block
size))) {
118:             inode_unlock(i);
119:             return -EIO;
120:         }
121:     } else {
122:         if(!(buf = bread(i->dev, i->u.minix.u.i2_zone[0], i->sb->s_block
size))) {
123:             inode_unlock(i);
124:             return -EIO;
125:         }
126:     }
127:     name = buf->data;
128:     inode_unlock(i);
129:
130:     current->loopcnt++;

```

**fs/minix/symlink.c**

Page 3/3

```
131:         iput(i);
132:         brelse(buf);
133:         errno = parse_namei(name, dir, i_res, NULL, FOLLOW_LINKS);
134:         current->loopcnt--;
135:         return errno;
136: }
```

## fs/minix/v1\_inode.c

Page 1/7

```

1: /*
2:  * fiwix/fs/minix/v1_inode.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/fs_minix.h>
12: #include <fiwix/fs_pipe.h>
13: #include <fiwix/statfs.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/stat.h>
16: #include <fiwix/sched.h>
17: #include <fiwix/buffer.h>
18: #include <fiwix/process.h>
19: #include <fiwix/errno.h>
20: #include <fiwix/stdio.h>
21: #include <fiwix/string.h>
22:
23: #define BLOCKS_PER_IND_BLOCK(sb)          (sb->s_blocksize / sizeof(__u16))
24: #define MINIX_INODES_PER_BLOCK(sb)       (sb->s_blocksize / sizeof(struct minix_i
node))
25:
26: #define MINIX_NDIR_BLOCKS                  7
27: #define MINIX_IND_BLOCK                   MINIX_NDIR_BLOCKS
28: #define MINIX_DIND_BLOCK                  (MINIX_NDIR_BLOCKS + 1)
29:
30: static void free_zone(struct inode *i, int block, int offset)
31: {
32:     int n;
33:     struct buffer *buf;
34:     __u16 *zone;
35:
36:     if(!(buf = bread(i->dev, block, i->sb->s_blocksize))) {
37:         printk("WARNING: %s(): error reading block %d.\n", __FUNCTION__,
block);
38:         return;
39:     }
40:     zone = (__u16 *)buf->data;
41:     for(n = offset; n < BLOCKS_PER_IND_BLOCK(i->sb); n++) {
42:         if(zone[n]) {
43:             minix_bfree(i->sb, zone[n]);
44:             zone[n] = 0;
45:         }
46:     }
47:     bwrite(buf);
48: }
49:
50: int v1_minix_read_inode(struct inode *i)
51: {
52:     __ino_t block;
53:     short int offset;
54:     struct minix_inode *ii;
55:     struct buffer *buf;
56:     int errno;
57:
58:     block = 1 + SUPERBLOCK + i->sb->u.minix.sb.s_imap_blocks + i->sb->u.mini
x.sb.s_zmap_blocks + (i->inode - 1) / MINIX_INODES_PER_BLOCK(i->sb);
59:
60:     if(!(buf = bread(i->dev, block, i->sb->s_blocksize))) {
61:         return -EIO;
62:     }
63:     offset = (i->inode - 1) % MINIX_INODES_PER_BLOCK(i->sb);
64:     ii = ((struct minix_inode *)buf->data) + offset;

```

## fs/minix/v1\_inode.c

Page 2/7

```

65:
66:     i->i_mode = ii->i_mode;
67:     i->i_uid = ii->i_uid;
68:     i->i_size = ii->i_size;
69:     i->i_atime = ii->i_time;
70:     i->i_ctime = ii->i_time;
71:     i->i_mtime = ii->i_time;
72:     i->i_gid = ii->i_gid;
73:     i->i_nlink = ii->i_nlinks;
74:     memcpy_b(i->u.minix.u.il_zone, ii->i_zone, sizeof(ii->i_zone));
75:     i->count = 1;
76:
77:     errno = 0;
78:     switch(i->i_mode & S_IFMT) {
79:         case S_IFCHR:
80:             i->fsop = &def_chr_fsop;
81:             i->rdev = ii->i_zone[0];
82:             break;
83:         case S_IFBLK:
84:             i->fsop = &def_blk_fsop;
85:             i->rdev = ii->i_zone[0];
86:             break;
87:         case S_IFIFO:
88:             i->fsop = &pipefs_fsop;
89:             /* it's a union so we need to clear pipefs_i */
90:             memset_b(&i->u.pipefs, NULL, sizeof(struct pipefs_inode)
);
91:             break;
92:         case S_IFDIR:
93:             i->fsop = &minix_dir_fsop;
94:             break;
95:         case S_IFREG:
96:             i->fsop = &minix_file_fsop;
97:             break;
98:         case S_IFLNK:
99:             i->fsop = &minix_symlink_fsop;
100:            break;
101:         case S_IFSOCK:
102:             i->fsop = NULL;
103:             break;
104:         default:
105:             printk("WARNING: %s(): invalid inode (%d) mode %o.\n", _
_FUNCTION__, i->inode, i->i_mode);
106:             errno = -ENOENT;
107:             break;
108:     }
109:
110:     brelse(buf);
111:     return errno;
112: }
113:
114: int v1_minix_write_inode(struct inode *i)
115: {
116:     __ino_t block;
117:     short int offset;
118:     struct minix_inode *ii;
119:     struct buffer *buf;
120:
121:     block = 1 + SUPERBLOCK + i->sb->u.minix.sb.s_imap_blocks + i->sb->u.mini
x.sb.s_zmap_blocks + (i->inode - 1) / MINIX_INODES_PER_BLOCK(i->sb);
122:
123:     if(!(buf = bread(i->dev, block, i->sb->s_blocksize))) {
124:         return -EIO;
125:     }
126:     offset = (i->inode - 1) % MINIX_INODES_PER_BLOCK(i->sb);
127:     ii = ((struct minix_inode *)buf->data) + offset;
128:

```

## fs/minix/v1\_inode.c

Page 3/7

```

129:         ii->i_mode = i->i_mode;
130:         ii->i_uid = i->i_uid;
131:         ii->i_size = i->i_size;
132:         ii->i_time = i->i_mtime;
133:         ii->i_gid = i->i_gid;
134:         ii->i_nlinks = i->i_nlink;
135:         if(S_ISCHR(i->i_mode) || S_ISBLK(i->i_mode)) {
136:             ii->i_zone[0] = i->rdev;
137:         } else {
138:             memcpy_b(ii->i_zone, i->u.minix.u.il_zone, sizeof(i->u.minix.u.i
139:         }
140:         i->dirty = 0;
141:         bwrite(buf);
142:         return 0;
143:     }
144:
145: int v1_minix_ialloc(struct inode *i, int mode)
146: {
147:     __blk_t offset;
148:     int inode, errno;
149:     struct superblock *sb;
150:
151:     sb = i->sb;
152:     superblock_lock(sb);
153:
154:     offset = 1 + SUPERBLOCK;
155:
156:     if(!(inode = minix_find_first_zero(sb, offset, sb->u.minix.sb.s_ninodes,
offset + sb->u.minix.sb.s_imap_blocks))) {
157:         superblock_unlock(sb);
158:         return -ENOSPC;
159:     }
160:
161:     errno = minix_change_bit(SET_BIT, sb, offset, inode);
162:
163:     if(errno) {
164:         if(errno < 0) {
165:             printk("WARNING: %s(): unable to set inode %d.\n", __FUN
CTION__, inode);
166:             superblock_unlock(sb);
167:             return errno;
168:         } else {
169:             printk("WARNING: %s(): inode %d is already marked as use
d!\n", __FUNCTION__, inode);
170:         }
171:     }
172:
173:     i->inode = inode;
174:     i->i_atime = CURRENT_TIME;
175:     i->i_mtime = CURRENT_TIME;
176:     i->i_ctime = CURRENT_TIME;
177:     superblock_unlock(sb);
178:     return 0;
179: }
180:
181: void v1_minix_ifree(struct inode *i)
182: {
183:     int errno;
184:     struct superblock *sb;
185:
186:     minix_truncate(i, 0);
187:
188:     sb = i->sb;
189:     superblock_lock(sb);
190:
191:     errno = minix_change_bit(CLEAR_BIT, i->sb, 1 + SUPERBLOCK, i->inode);

```

## fs/minix/v1\_inode.c

Page 4/7

```

192:
193:     if(errno) {
194:         if(errno < 0) {
195:             printk("WARNING: %s(): unable to clear inode %d.\n", __F
UNCTION__, i->inode);
196:         } else {
197:             printk("WARNING: %s(): inode %d is already marked as fre
e!\n", __FUNCTION__, i->inode);
198:         }
199:     }
200:
201:     i->i_size = 0;
202:     i->i_mtime = CURRENT_TIME;
203:     i->i_ctime = CURRENT_TIME;
204:     i->dirty = 1;
205:     superblock_unlock(sb);
206: }
207:
208: int v1_minix_bmap(struct inode *i, __off_t offset, int mode)
209: {
210:     unsigned char level;
211:     __u16 *indblock, *dindblock;
212:     __blk_t block, iblock, dblock, newblock;
213:     int blksize;
214:     struct buffer *buf, *buf2, *buf3;
215:
216:     blksize = i->sb->s_blocksize;
217:     block = offset / blksize;
218:     level = 0;
219:
220:     if(block < MINIX_NDIR_BLOCKS) {
221:         level = MINIX_NDIR_BLOCKS - 1;
222:     } else {
223:         if(block < (BLOCKS_PER_IND_BLOCK(i->sb) + MINIX_NDIR_BLOCKS)) {
224:             level = MINIX_IND_BLOCK;
225:         } else {
226:             level = MINIX_DIND_BLOCK;
227:         }
228:         block -= MINIX_NDIR_BLOCKS;
229:     }
230:
231:     if(level < MINIX_NDIR_BLOCKS) {
232:         if(!i->u.minix.u.il_zone[block] && mode == FOR_WRITING) {
233:             if((newblock = minix_balloc(i->sb)) < 0) {
234:                 return -ENOSPC;
235:             }
236:             /* initialize the new block */
237:             if(!(buf = bread(i->dev, newblock, blksize))) {
238:                 minix_bfree(i->sb, newblock);
239:                 return -EIO;
240:             }
241:             memset_b(buf->data, 0, blksize);
242:             bwrite(buf);
243:             i->u.minix.u.il_zone[block] = newblock;
244:         }
245:         return i->u.minix.u.il_zone[block];
246:     }
247:
248:     if(!i->u.minix.u.il_zone[level]) {
249:         if(mode == FOR_WRITING) {
250:             if((newblock = minix_balloc(i->sb)) < 0) {
251:                 return -ENOSPC;
252:             }
253:             /* initialize the new block */
254:             if(!(buf = bread(i->dev, newblock, blksize))) {
255:                 minix_bfree(i->sb, newblock);
256:                 return -EIO;

```

## fs/minix/v1\_inode.c

Page 5/7

```

257:         }
258:         memset_b(buf->data, 0, blksize);
259:         bwrite(buf);
260:         i->u.minix.u.il_zone[level] = newblock;
261:     } else {
262:         return 0;
263:     }
264: }
265: if(!(buf = bread(i->dev, i->u.minix.u.il_zone[level], blksize))) {
266:     return -EIO;
267: }
268: indblock = (__u16 *)buf->data;
269: dblock = block - BLOCKS_PER_IND_BLOCK(i->sb);
270:
271: if(level == MINIX_DIND_BLOCK) {
272:     block = dblock / BLOCKS_PER_IND_BLOCK(i->sb);
273: }
274:
275: if(!indblock[block]) {
276:     if(mode == FOR_WRITING) {
277:         if((newblock = minix_balloc(i->sb)) < 0) {
278:             brelse(buf);
279:             return -ENOSPC;
280:         }
281:         /* initialize the new block */
282:         if(!(buf2 = bread(i->dev, newblock, blksize))) {
283:             minix_bfree(i->sb, newblock);
284:             brelse(buf);
285:             return -EIO;
286:         }
287:         memset_b(buf2->data, 0, blksize);
288:         bwrite(buf2);
289:         indblock[block] = newblock;
290:         if(level == MINIX_IND_BLOCK) {
291:             bwrite(buf);
292:             return newblock;
293:         }
294:         buf->dirty = 1;
295:         buf->valid = 1;
296:     } else {
297:         brelse(buf);
298:         return 0;
299:     }
300: }
301: if(level == MINIX_IND_BLOCK) {
302:     newblock = indblock[block];
303:     brelse(buf);
304:     return newblock;
305: }
306:
307: iblock = block;
308: if(!(buf2 = bread(i->dev, indblock[iblock], blksize))) {
309:     printk("%s(): returning -EIO\n", __FUNCTION__);
310:     brelse(buf);
311:     return -EIO;
312: }
313: dindblock = (__u16 *)buf2->data;
314: block = dindblock[dblock - (iblock * BLOCKS_PER_IND_BLOCK(i->sb))];
315: if(!block && mode == FOR_WRITING) {
316:     if((newblock = minix_balloc(i->sb)) < 0) {
317:         brelse(buf);
318:         brelse(buf2);
319:         return -ENOSPC;
320:     }
321:     /* initialize the new block */
322:     if(!(buf3 = bread(i->dev, newblock, blksize))) {
323:         minix_bfree(i->sb, newblock);

```



## fs/minix/v1\_inode.c

Page 6/7

```

324:             brelse(buf);
325:             brelse(buf2);
326:             return -EIO;
327:         }
328:         memset_b(buf3->data, 0, blksize);
329:         bwrite(buf3);
330:         dindblock[dblock - (iblock * BLOCKS_PER_IND_BLOCK(i->sb))] = new
block;
331:         buf2->dirty = 1;
332:         buf2->valid = 1;
333:         block = newblock;
334:     }
335:     brelse(buf);
336:     brelse(buf2);
337:     return block;
338: }
339:
340: int v1_minix_truncate(struct inode *i, __off_t length)
341: {
342:     int n;
343:     __blk_t block, dblock;
344:     __u16 *zone;
345:     struct buffer *buf;
346:
347:     block = length / i->sb->s_blocksize;
348:
349:     if(!S_ISDIR(i->i_mode) && !S_ISREG(i->i_mode) && !S_ISLNK(i->i_mode)) {
350:         return -EINVAL;
351:     }
352:
353:     if(block < MINIX_NDIR_BLOCKS) {
354:         for(n = block; n < MINIX_NDIR_BLOCKS; n++) {
355:             if(i->u.minix.u.il_zone[n]) {
356:                 minix_bfree(i->sb, i->u.minix.u.il_zone[n]);
357:                 i->u.minix.u.il_zone[n] = 0;
358:             }
359:         }
360:         block = 0;
361:     }
362:
363:     if(!block || block < (BLOCKS_PER_IND_BLOCK(i->sb) + MINIX_NDIR_BLOCKS))
{
364:         if(block) {
365:             block -= MINIX_NDIR_BLOCKS;
366:         }
367:         if(i->u.minix.u.il_zone[MINIX_IND_BLOCK]) {
368:             free_zone(i, i->u.minix.u.il_zone[MINIX_IND_BLOCK], bloc
k);
369:             if(!block) {
370:                 minix_bfree(i->sb, i->u.minix.u.il_zone[MINIX_IN
D_BLOCK]);
371:                 i->u.minix.u.il_zone[MINIX_IND_BLOCK] = 0;
372:             }
373:         }
374:         block = 0;
375:     }
376:
377:     if(block) {
378:         block -= MINIX_NDIR_BLOCKS;
379:         block -= BLOCKS_PER_IND_BLOCK(i->sb);
380:     }
381:     if(i->u.minix.u.il_zone[MINIX_DIND_BLOCK]) {
382:         if(!(buf = bread(i->dev, i->u.minix.u.il_zone[MINIX_DIND_BLOCK],
i->sb->s_blocksize))) {
383:             printk("%s(): error reading block %d.\n", __FUNCTION__,
i->u.minix.u.il_zone[MINIX_DIND_BLOCK]);
384:         }

```

## fs/minix/v1\_inode.c

Page 7/7

```
385:         zone = (__u16 *)buf->data;
386:         dblock = block % BLOCKS_PER_IND_BLOCK(i->sb);
387:         for(n = block / BLOCKS_PER_IND_BLOCK(i->sb); n < BLOCKS_PER_IND_
BLOCK(i->sb); n++) {
388:             if(zone[n]) {
389:                 free_zone(i, zone[n], dblock);
390:                 if(!dblock) {
391:                     minix_bfree(i->sb, zone[n]);
392:                 }
393:             }
394:             dblock = 0;
395:         }
396:         bwrite(buf);
397:         if(!block) {
398:             minix_bfree(i->sb, i->u.minix.u.i1_zone[MINIX_DIND_BLOCK
]);
399:             i->u.minix.u.i1_zone[MINIX_DIND_BLOCK] = 0;
400:         }
401:     }
402:
403:     i->i_mtime = CURRENT_TIME;
404:     i->i_ctime = CURRENT_TIME;
405:     i->i_size = length;
406:     i->dirty = 1;
407:
408:     return 0;
409: }
```

## fs/minix/v2\_inode.c

Page 1/8

```

1: /*
2:  * fiwix/fs/minix/v2_inode.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/fs_minix.h>
12: #include <fiwix/fs_pipe.h>
13: #include <fiwix/statfs.h>
14: #include <fiwix/sleep.h>
15: #include <fiwix/stat.h>
16: #include <fiwix/sched.h>
17: #include <fiwix/buffer.h>
18: #include <fiwix/process.h>
19: #include <fiwix/errno.h>
20: #include <fiwix/stdio.h>
21: #include <fiwix/string.h>
22:
23: #define BLOCKS_PER_IND_BLOCK(sb)          (sb->s_blocksize / sizeof(__u32))
24: #define MINIX2_INODES_PER_BLOCK(sb)      (sb->s_blocksize / sizeof(struct minix2_
inode))
25:
26: #define MINIX_NDIR_BLOCKS                  7
27: #define MINIX_IND_BLOCK                   MINIX_NDIR_BLOCKS
28: #define MINIX_DIND_BLOCK                   (MINIX_NDIR_BLOCKS + 1)
29: #define MINIX_TIND_BLOCK                   (MINIX_NDIR_BLOCKS + 2)
30:
31: static void free_zone(struct inode *i, int block, int offset)
32: {
33:     int n;
34:     struct buffer *buf;
35:     __u32 *zone;
36:
37:     if(!(buf = bread(i->dev, block, i->sb->s_blocksize))) {
38:         printk("WARNING: %s(): error reading block %d.\n", __FUNCTION__,
block);
39:         return;
40:     }
41:     zone = (__u32 *)buf->data;
42:     for(n = offset; n < BLOCKS_PER_IND_BLOCK(i->sb); n++) {
43:         if(zone[n]) {
44:             minix_bfree(i->sb, zone[n]);
45:             zone[n] = 0;
46:         }
47:     }
48:     bwrite(buf);
49: }
50:
51: int v2_minix_read_inode(struct inode *i)
52: {
53:     __ino_t block;
54:     short int offset;
55:     struct minix2_inode *ii;
56:     struct buffer *buf;
57:     int errno;
58:
59:     block = 1 + SUPERBLOCK + i->sb->u.minix.sb.s_imap_blocks + i->sb->u.mini
x.sb.s_zmap_blocks + (i->inode - 1) / MINIX2_INODES_PER_BLOCK(i->sb);
60:
61:     if(!(buf = bread(i->dev, block, i->sb->s_blocksize))) {
62:         return -EIO;
63:     }
64:     offset = (i->inode - 1) % MINIX2_INODES_PER_BLOCK(i->sb);

```

## fs/minix/v2\_inode.c

Page 2/8

```

65:         ii = ((struct minix2_inode *)buf->data) + offset;
66:
67:         i->i_mode = ii->i_mode;
68:         i->i_nlink = ii->i_nlink;
69:         i->i_uid = ii->i_uid;
70:         i->i_gid = ii->i_gid;
71:         i->i_size = ii->i_size;
72:         i->i_atime = ii->i_atime;
73:         i->i_mtime = ii->i_mtime;
74:         i->i_ctime = ii->i_ctime;
75:         memcpy_b(i->u.minix.u.i2_zone, ii->i_zone, sizeof(ii->i_zone));
76:         i->count = 1;
77:
78:         errno = 0;
79:         switch(i->i_mode & S_IFMT) {
80:             case S_IFCHR:
81:                 i->fsop = &def_chr_fsop;
82:                 i->rdev = ii->i_zone[0];
83:                 break;
84:             case S_IFBLK:
85:                 i->fsop = &def_blk_fsop;
86:                 i->rdev = ii->i_zone[0];
87:                 break;
88:             case S_IFIFO:
89:                 i->fsop = &pipefs_fsop;
90:                 /* it's a union so we need to clear pipefs_i */
91:                 memset_b(&i->u.pipefs, NULL, sizeof(struct pipefs_inode)
);
92:                 break;
93:             case S_IFDIR:
94:                 i->fsop = &minix_dir_fsop;
95:                 break;
96:             case S_IFREG:
97:                 i->fsop = &minix_file_fsop;
98:                 break;
99:             case S_IFLNK:
100:                i->fsop = &minix_symlink_fsop;
101:                break;
102:             case S_IFSOCK:
103:                i->fsop = NULL;
104:                break;
105:             default:
106:                printk("WARNING: %s(): invalid inode (%d) mode %o.\n", _
_FUNCTION__, i->inode, i->i_mode);
107:                errno = -ENOENT;
108:                break;
109:         }
110:
111:         brelse(buf);
112:         return errno;
113: }
114:
115: int v2_minix_write_inode(struct inode *i)
116: {
117:     __ino_t block;
118:     short int offset;
119:     struct minix2_inode *ii;
120:     struct buffer *buf;
121:
122:     block = 1 + SUPERBLOCK + i->sb->u.minix.sb.s_imap_blocks + i->sb->u.mini
x.sb.s_zmap_blocks + (i->inode - 1) / MINIX2_INODES_PER_BLOCK(i->sb);
123:
124:     if (!(buf = bread(i->dev, block, i->sb->s_blocksize))) {
125:         return -EIO;
126:     }
127:     offset = (i->inode - 1) % MINIX2_INODES_PER_BLOCK(i->sb);
128:     ii = ((struct minix2_inode *)buf->data) + offset;

```

## fs/minix/v2\_inode.c

Page 3/8

```

129:
130:     ii->i_mode = i->i_mode;
131:     ii->i_nlink = i->i_nlink;
132:     ii->i_uid = i->i_uid;
133:     ii->i_gid = i->i_gid;
134:     ii->i_size = i->i_size;
135:     ii->i_atime = i->i_atime;
136:     ii->i_mtime = i->i_mtime;
137:     ii->i_ctime = i->i_ctime;
138:     if(S_ISCHR(i->i_mode) || S_ISBLK(i->i_mode)) {
139:         ii->i_zone[0] = i->rdev;
140:     } else {
141:         memcpy_b(ii->i_zone, i->u.minix.u.i2_zone, sizeof(i->u.minix.u.i
2_zone));
142:     }
143:     i->dirty = 0;
144:     bwrite(buf);
145:     return 0;
146: }
147:
148: int v2_minix_ialloc(struct inode *i, int mode)
149: {
150:     __blk_t offset;
151:     int inode, errno;
152:     struct superblock *sb;
153:
154:     sb = i->sb;
155:     superblock_lock(sb);
156:
157:     offset = 1 + SUPERBLOCK;
158:
159:     if(!(inode = minix_find_first_zero(sb, offset, sb->u.minix.sb.s_ninodes,
offset + sb->u.minix.sb.s_imap_blocks))) {
160:         superblock_unlock(sb);
161:         return -ENOSPC;
162:     }
163:
164:     errno = minix_change_bit(SET_BIT, sb, offset, inode);
165:
166:     if(errno) {
167:         if(errno < 0) {
168:             printk("WARNING: %s(): unable to set inode %d.\n", __FUN
CTION__, inode);
169:             superblock_unlock(sb);
170:             return errno;
171:         } else {
172:             printk("WARNING: %s(): inode %d is already marked as use
d!\n", __FUNCTION__, inode);
173:         }
174:     }
175:
176:     i->inode = inode;
177:     i->i_atime = CURRENT_TIME;
178:     i->i_mtime = CURRENT_TIME;
179:     i->i_ctime = CURRENT_TIME;
180:     superblock_unlock(sb);
181:     return 0;
182: }
183:
184: void v2_minix_ifree(struct inode *i)
185: {
186:     int errno;
187:     struct superblock *sb;
188:
189:     minix_truncate(i, 0);
190:
191:     sb = i->sb;

```

## fs/minix/v2\_inode.c

Page 4/8

```

192:         superblock_lock(sb);
193:
194:         errno = minix_change_bit(CLEAR_BIT, i->sb, 1 + SUPERBLOCK, i->inode);
195:
196:         if(errno) {
197:             if(errno < 0) {
198:                 printk("WARNING: %s(): unable to clear inode %d.\n", __F
UNCTION__, i->inode);
199:             } else {
200:                 printk("WARNING: %s(): inode %d is already marked as fre
e!\n", __FUNCTION__, i->inode);
201:             }
202:         }
203:
204:         i->i_size = 0;
205:         i->i_mtime = CURRENT_TIME;
206:         i->i_ctime = CURRENT_TIME;
207:         i->dirty = 1;
208:         superblock_unlock(sb);
209:     }
210:
211: int v2_minix_bmap(struct inode *i, __off_t offset, int mode)
212: {
213:     unsigned char level;
214:     __u32 *indblock, *dindblock, *tindblock;
215:     __blk_t block, iblock, dblock, tblock, newblock;
216:     int blksize;
217:     struct buffer *buf, *buf2, *buf3, *buf4;
218:
219:     blksize = i->sb->s_blocksize;
220:     block = offset / blksize;
221:     level = 0;
222:     buf3 = NULL;    /* makes GCC happy */
223:
224:     if(block < MINIX_NDIR_BLOCKS) {
225:         level = MINIX_NDIR_BLOCKS - 1;
226:     } else {
227:         if(block < (BLOCKS_PER_IND_BLOCK(i->sb) + MINIX_NDIR_BLOCKS)) {
228:             level = MINIX_IND_BLOCK;
229:         } else if(block < ((BLOCKS_PER_IND_BLOCK(i->sb) * BLOCKS_PER_IND
_BLOCK(i->sb)) + BLOCKS_PER_IND_BLOCK(i->sb) + MINIX_NDIR_BLOCKS)) {
230:             level = MINIX_DIND_BLOCK;
231:         } else {
232:             level = MINIX_TIND_BLOCK;
233:         }
234:         block -= MINIX_NDIR_BLOCKS;
235:     }
236:
237:     if(level < MINIX_NDIR_BLOCKS) {
238:         if(!i->u.minix.u.i2_zone[block] && mode == FOR_WRITING) {
239:             if((newblock = minix_balloc(i->sb)) < 0) {
240:                 return -ENOSPC;
241:             }
242:             /* initialize the new block */
243:             if(!(buf = bread(i->dev, newblock, blksize))) {
244:                 minix_bfree(i->sb, newblock);
245:                 return -EIO;
246:             }
247:             memset_b(buf->data, 0, blksize);
248:             bwrite(buf);
249:             i->u.minix.u.i2_zone[block] = newblock;
250:         }
251:         return i->u.minix.u.i2_zone[block];
252:     }
253:
254:     if(!i->u.minix.u.i2_zone[level]) {
255:         if(mode == FOR_WRITING) {

```

## fs/minix/v2\_inode.c

Page 5/8

```

256:         if((newblock = minix_balloc(i->sb)) < 0) {
257:             return -ENOSPC;
258:         }
259:         /* initialize the new block */
260:         if(!(buf = bread(i->dev, newblock, blksize))) {
261:             minix_bfree(i->sb, newblock);
262:             return -EIO;
263:         }
264:         memset_b(buf->data, 0, blksize);
265:         bwrite(buf);
266:         i->u.minix.u.i2_zone[level] = newblock;
267:     } else {
268:         return 0;
269:     }
270: }
271: if(!(buf = bread(i->dev, i->u.minix.u.i2_zone[level], blksize))) {
272:     return -EIO;
273: }
274: indblock = (__u32 *)buf->data;
275: dblock = block - BLOCKS_PER_IND_BLOCK(i->sb);
276: tblock = block - (BLOCKS_PER_IND_BLOCK(i->sb) * BLOCKS_PER_IND_BLOCK(i->
sb)) - BLOCKS_PER_IND_BLOCK(i->sb);
277:
278: if(level == MINIX_DIND_BLOCK) {
279:     block = dblock / BLOCKS_PER_IND_BLOCK(i->sb);
280: }
281: if(level == MINIX_TIND_BLOCK) {
282:     block = tblock / (BLOCKS_PER_IND_BLOCK(i->sb) * BLOCKS_PER_IND_B
LOCK(i->sb));
283: }
284:
285: if(!indblock[block]) {
286:     if(mode == FOR_WRITING) {
287:         if((newblock = minix_balloc(i->sb)) < 0) {
288:             brelse(buf);
289:             return -ENOSPC;
290:         }
291:         /* initialize the new block */
292:         if(!(buf2 = bread(i->dev, newblock, blksize))) {
293:             minix_bfree(i->sb, newblock);
294:             brelse(buf);
295:             return -EIO;
296:         }
297:         memset_b(buf2->data, 0, blksize);
298:         bwrite(buf2);
299:         indblock[block] = newblock;
300:         if(level == MINIX_IND_BLOCK) {
301:             bwrite(buf);
302:             return newblock;
303:         }
304:         buf->dirty = 1;
305:         buf->valid = 1;
306:     } else {
307:         brelse(buf);
308:         return 0;
309:     }
310: }
311: if(level == MINIX_IND_BLOCK) {
312:     newblock = indblock[block];
313:     brelse(buf);
314:     return newblock;
315: }
316:
317: if(level == MINIX_TIND_BLOCK) {
318:     if(!(buf3 = bread(i->dev, indblock[block], blksize))) {
319:         printk("%s(): returning -EIO\n", __FUNCTION__);
320:         brelse(buf);

```

## fs/minix/v2\_inode.c

Page 6/8

```

321:             return -EIO;
322:         }
323:         tindblock = (__u32 *)buf3->data;
324:         block = tindblock[tblock / BLOCKS_PER_IND_BLOCK(i->sb)];
325:         if(!block) {
326:             if(mode == FOR_WRITING) {
327:                 if((newblock = minix_balloc(i->sb)) < 0) {
328:                     brelse(buf);
329:                     brelse(buf3);
330:                     return -ENOSPC;
331:                 }
332:                 /* initialize the new block */
333:                 if(!(buf4 = bread(i->dev, newblock, blksize))) {
334:                     minix_bfree(i->sb, newblock);
335:                     brelse(buf);
336:                     brelse(buf3);
337:                     return -EIO;
338:                 }
339:                 memset_b(buf4->data, 0, blksize);
340:                 bwrite(buf4);
341:                 tindblock[tblock / BLOCKS_PER_IND_BLOCK(i->sb)]
= newblock;
342:                 buf3->dirty = 1;
343:                 buf3->valid = 1;
344:                 block = newblock;
345:             } else {
346:                 brelse(buf);
347:                 brelse(buf3);
348:                 return 0;
349:             }
350:         }
351:         dblock = tblock;
352:         iblock = tblock / BLOCKS_PER_IND_BLOCK(i->sb);
353:         if(!(buf2 = bread(i->dev, block, blksize))) {
354:             printk("%s(): returning -EIO\n", __FUNCTION__);
355:             brelse(buf);
356:             brelse(buf3);
357:             return -EIO;
358:         }
359:     } else {
360:         iblock = block;
361:         if(!(buf2 = bread(i->dev, indblock[iblock], blksize))) {
362:             printk("%s(): returning -EIO\n", __FUNCTION__);
363:             brelse(buf);
364:             return -EIO;
365:         }
366:     }
367:
368:     dindblock = (__u32 *)buf2->data;
369:     block = dindblock[dblock - (iblock * BLOCKS_PER_IND_BLOCK(i->sb))];
370:     if(!block && mode == FOR_WRITING) {
371:         if((newblock = minix_balloc(i->sb)) < 0) {
372:             brelse(buf);
373:             if(level == MINIX_TIND_BLOCK) {
374:                 brelse(buf3);
375:             }
376:             brelse(buf2);
377:             return -ENOSPC;
378:         }
379:         /* initialize the new block */
380:         if(!(buf4 = bread(i->dev, newblock, blksize))) {
381:             minix_bfree(i->sb, newblock);
382:             brelse(buf);
383:             if(level == MINIX_TIND_BLOCK) {
384:                 brelse(buf3);
385:             }
386:             brelse(buf2);

```



## fs/minix/v2\_inode.c

Page 7/8

```

387:             return -EIO;
388:         }
389:         memset_b(buf4->data, 0, blksize);
390:         bwrite(buf4);
391:         dindblock[dblock - (iblock * BLOCKS_PER_IND_BLOCK(i->sb))] = new
block;
392:         buf2->dirty = 1;
393:         buf2->valid = 1;
394:         block = newblock;
395:     }
396:     brelse(buf);
397:     if(level == MINIX_TIND_BLOCK) {
398:         brelse(buf3);
399:     }
400:     brelse(buf2);
401:     return block;
402: }
403:
404: int v2_minix_truncate(struct inode *i, __off_t length)
405: {
406:     int n;
407:     __blk_t block, dblock;
408:     __u32 *zone;
409:     struct buffer *buf;
410:
411:     block = length / i->sb->s_blocksize;
412:
413:     if(!S_ISDIR(i->i_mode) && !S_ISREG(i->i_mode) && !S_ISLNK(i->i_mode)) {
414:         return -EINVAL;
415:     }
416:
417:     if(block < MINIX_NDIR_BLOCKS) {
418:         for(n = block; n < MINIX_NDIR_BLOCKS; n++) {
419:             if(i->u.minix.u.i2_zone[n]) {
420:                 minix_bfree(i->sb, i->u.minix.u.i2_zone[n]);
421:                 i->u.minix.u.i2_zone[n] = 0;
422:             }
423:         }
424:         block = 0;
425:     }
426:
427:     if(!block || block < (BLOCKS_PER_IND_BLOCK(i->sb) + MINIX_NDIR_BLOCKS))
{
428:         if(block) {
429:             block -= MINIX_NDIR_BLOCKS;
430:         }
431:         if(i->u.minix.u.i2_zone[MINIX_IND_BLOCK]) {
432:             free_zone(i, i->u.minix.u.i2_zone[MINIX_IND_BLOCK], bloc
k);
433:             if(!block) {
434:                 minix_bfree(i->sb, i->u.minix.u.i2_zone[MINIX_IN
D_BLOCK]);
435:                 i->u.minix.u.i2_zone[MINIX_IND_BLOCK] = 0;
436:             }
437:         }
438:         block = 0;
439:     }
440:
441:     if(block) {
442:         block -= MINIX_NDIR_BLOCKS;
443:         block -= BLOCKS_PER_IND_BLOCK(i->sb);
444:     }
445:     if(i->u.minix.u.i2_zone[MINIX_DIND_BLOCK]) {
446:         if(!(buf = bread(i->dev, i->u.minix.u.i2_zone[MINIX_DIND_BLOCK],
i->sb->s_blocksize))) {
447:             printk("%s(): error reading block %d.\n", __FUNCTION__,
i->u.minix.u.i2_zone[MINIX_DIND_BLOCK]);

```

## fs/minix/v2\_inode.c

Page 8/8

```
448:         }
449:         zone = (__u32 *)buf->data;
450:         dblock = block % BLOCKS_PER_IND_BLOCK(i->sb);
451:         for(n = block / BLOCKS_PER_IND_BLOCK(i->sb); n < BLOCKS_PER_IND_
BLOCK(i->sb); n++) {
452:             if(zone[n]) {
453:                 free_zone(i, zone[n], dblock);
454:                 if(!dblock) {
455:                     minix_bfree(i->sb, zone[n]);
456:                 }
457:             }
458:             dblock = 0;
459:         }
460:         bwrite(buf);
461:         if(!block) {
462:             minix_bfree(i->sb, i->u.minix.u.i2_zone [MINIX_DIND_BLOCK
]);
463:             i->u.minix.u.i2_zone [MINIX_DIND_BLOCK] = 0;
464:         }
465:     }
466:
467:     i->i_mtime = CURRENT_TIME;
468:     i->i_ctime = CURRENT_TIME;
469:     i->i_size = length;
470:     i->dirty = 1;
471:
472:     return 0;
473: }
```

## fs/pipefs/fifo.c

Page 1/2

```

1:  /*
2:  *  fiwix/fs/pipefs/fifo.c
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #include <fiwix/types.h>
9:  #include <fiwix/errno.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_pipe.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/sleep.h>
16: #include <fiwix/fcntl.h>
17: #include <fiwix/sched.h>
18: #include <fiwix/stdio.h>
19:
20: int fifo_open(struct inode *i, struct fd *fd_table)
21: {
22:     /* first open */
23:     if(i->count == 1) {
24:         if(!(i->u.pipefs.i_data = (void *)kmalloc())) {
25:             return -ENOMEM;
26:         }
27:         i->u.pipefs.i_readoff = 0;
28:         i->u.pipefs.i_writeoff = 0;
29:     }
30:
31:     if((fd_table->flags & O_ACCMODE) == O_RDONLY) {
32:         i->u.pipefs.i_readers++;
33:         wakeup(&pipefs_write);
34:         if(!(fd_table->flags & O_NONBLOCK)) {
35:             while(!i->u.pipefs.i_writers) {
36:                 if(sleep(&pipefs_read, PROC_INTERRUPTIBLE)) {
37:                     if(!--i->u.pipefs.i_readers) {
38:                         wakeup(&pipefs_write);
39:                     }
40:                     return -EINTR;
41:                 }
42:             }
43:         }
44:     }
45:
46:     if((fd_table->flags & O_ACCMODE) == O_WRONLY) {
47:         if((fd_table->flags & O_NONBLOCK) && !i->u.pipefs.i_readers) {
48:             return -ENXIO;
49:         }
50:
51:         i->u.pipefs.i_writers++;
52:         wakeup(&pipefs_read);
53:         if(!(fd_table->flags & O_NONBLOCK)) {
54:             while(!i->u.pipefs.i_readers) {
55:                 if(sleep(&pipefs_write, PROC_INTERRUPTIBLE)) {
56:                     if(!--i->u.pipefs.i_writers) {
57:                         wakeup(&pipefs_read);
58:                     }
59:                     return -EINTR;
60:                 }
61:             }
62:         }
63:     }
64:
65:     if((fd_table->flags & O_ACCMODE) == O_RDWR) {
66:         i->u.pipefs.i_readers++;
67:         i->u.pipefs.i_writers++;

```

**fs/pipefs/fifo.c**

Page 2/2

```
68:                wakeup(&pipefs_write);
69:                wakeup(&pipefs_read);
70:            }
71:
72:            return 0;
73: }
```

**fs/pipefs/Makefile**

Page 1/1

```
1: # fiwix/fs/pipefs/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = super.o fifo.o pipe.o
13:
14: pipefs: $(OBJS)
15:     $(LD) $(LDFLAGS) -r $(OBJS) -o pipefs.o
16:
17: clean:
18:     rm -f *.o
19:
```

## fs/pipefs/pipe.c

Page 1/4

```

1: /*
2:  * fiwix/fs/pipefs/pipe.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_pipe.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/fcntl.h>
15: #include <fiwix/ioctl.h>
16: #include <fiwix/sleep.h>
17: #include <fiwix/sched.h>
18: #include <fiwix/stdio.h>
19: #include <fiwix/string.h>
20:
21: static struct resource pipe_resource = { NULL, NULL };
22:
23: int pipefs_close(struct inode *i, struct fd *fd_table)
24: {
25:     if((fd_table->flags & O_ACCMODE) == O_RDONLY) {
26:         if(!--i->u.pipefs.i_readers) {
27:             wakeup(&pipefs_write);
28:         }
29:     }
30:     if((fd_table->flags & O_ACCMODE) == O_WRONLY) {
31:         if(!--i->u.pipefs.i_writers) {
32:             wakeup(&pipefs_read);
33:         }
34:     }
35:     if((fd_table->flags & O_ACCMODE) == O_RDWR) {
36:         if(!--i->u.pipefs.i_readers) {
37:             wakeup(&pipefs_write);
38:         }
39:         if(!--i->u.pipefs.i_writers) {
40:             wakeup(&pipefs_read);
41:         }
42:     }
43:     return 0;
44: }
45:
46: int pipefs_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t count)
47: {
48:     __off_t bytes_read;
49:     __size_t n, limit;
50:     char *data;
51:
52:     bytes_read = 0;
53:     data = i->u.pipefs.i_data;
54:
55:     while(count) {
56:         if(i->u.pipefs.i_writeoff) {
57:             if(i->u.pipefs.i_readoff >= i->u.pipefs.i_writeoff) {
58:                 limit = PIPE_BUF - i->u.pipefs.i_readoff;
59:             } else {
60:                 limit = i->u.pipefs.i_writeoff - i->u.pipefs.i_readoff;
61:             }
62:         } else {
63:             limit = PIPE_BUF - i->u.pipefs.i_readoff;
64:         }
65:         n = MIN(limit, count);

```

## fs/pipefs/pipe.c

Page 2/4

```

66:         if(i->i_size && n) {
67:             lock_resource(&pipe_resource);
68:             memcpy_b(buffer + bytes_read, data + i->u.pipefs.i_reado
ff, n);
69:             bytes_read += n;
70:             i->u.pipefs.i_readoff += n;
71:             i->i_size -= n;
72:             if(i->u.pipefs.i_writeoff >= PIPE_BUF) {
73:                 i->u.pipefs.i_writeoff = 0;
74:             }
75:             unlock_resource(&pipe_resource);
76:             wakeup(&pipefs_write);
77:             break;
78:         } else {
79:             if(i->u.pipefs.i_writers) {
80:                 if(fd_table->flags & O_NONBLOCK) {
81:                     return -EAGAIN;
82:                 }
83:                 if(sleep(&pipefs_read, PROC_INTERRUPTIBLE)) {
84:                     return -EINTR;
85:                 }
86:             } else {
87:                 if(i->i_size) {
88:                     if(i->u.pipefs.i_readoff >= PIPE_BUF) {
89:                         i->u.pipefs.i_readoff = 0;
90:                         continue;
91:                     }
92:                 }
93:                 break;
94:             }
95:         }
96:     }
97:     if(!i->i_size) {
98:         i->u.pipefs.i_readoff = 0;
99:         i->u.pipefs.i_writeoff = 0;
100:     }
101:     return bytes_read;
102: }
103:
104: int pipefs_write(struct inode *i, struct fd *fd_table, const char *buffer, __siz
e_t count)
105: {
106:     __off_t bytes_written;
107:     __size_t n;
108:     char *data;
109:     int limit;
110:
111:     bytes_written = 0;
112:     data = i->u.pipefs.i_data;
113:
114:     while(bytes_written < count) {
115:         /* if the read end closes then send signal and return */
116:         if(!i->u.pipefs.i_readers) {
117:             send_sig(current, SIGPIPE);
118:             return -EPIPE;
119:         }
120:
121:         if(i->u.pipefs.i_readoff) {
122:             if(i->u.pipefs.i_writeoff <= i->u.pipefs.i_readoff) {
123:                 limit = i->u.pipefs.i_readoff;
124:             } else {
125:                 limit = PIPE_BUF;
126:             }
127:         } else {
128:             limit = PIPE_BUF;
129:         }
130:

```

## fs/pipefs/pipe.c

Page 3/4

```

131:         n = MIN((count - bytes_written), (limit - i->u.pipefs.i_writeoff
));
132:
133:         /*
134:          * POSIX requires that any write operation involving fewer than
135:          * PIPE_BUF bytes must be automatically executed and finished
136:          * without being interleaved with write operations of other
137:          * processes to the same pipe.
138:          */
139:         if(n && n <= PIPE_BUF) {
140:             lock_resource(&pipe_resource);
141:             memcpy_b(data + i->u.pipefs.i_writeoff, buffer + bytes_w
ritten, n);
142:             bytes_written += n;
143:             i->u.pipefs.i_writeoff += n;
144:             i->i_size += n;
145:             if(i->u.pipefs.i_readoff >= PIPE_BUF) {
146:                 i->u.pipefs.i_readoff = 0;
147:             }
148:             unlock_resource(&pipe_resource);
149:             wakeup(&pipefs_read);
150:             continue;
151:         }
152:
153:         wakeup(&pipefs_read);
154:         if(!(fd_table->flags & O_NONBLOCK)) {
155:             if(sleep(&pipefs_write, PROC_INTERRUPTIBLE)) {
156:                 return -EINTR;
157:             }
158:         } else {
159:             return -EAGAIN;
160:         }
161:     }
162:     return bytes_written;
163: }
164:
165: int pipefs_ioctl(struct inode *i, int cmd, unsigned long int arg)
166: {
167:     int errno;
168:
169:     switch(cmd) {
170:         case FIONREAD:
171:             if((errno = check_user_area(VERIFY_WRITE, (void *)arg, s
izeof(unsigned int)))) {
172:                 return errno;
173:             }
174:             memcpy_b((void *)arg, &i->i_size, sizeof(unsigned int));
175:             break;
176:         default:
177:             return -EINVAL;
178:     }
179:     return 0;
180: }
181:
182: int pipefs_lseek(struct inode *i, __off_t offset)
183: {
184:     return -ESPIPE;
185: }
186:
187: int pipefs_select(struct inode *i, int flag)
188: {
189:     switch(flag) {
190:         case SEL_R:
191:             if(i->i_size || !i->u.pipefs.i_writers) {
192:                 return 1;
193:             }
194:             break;

```



**fs/pipefs/pipe.c**

Page 4/4

```
195:                 case SEL_W:
196:                     if(i->i_size < PIPE_BUF || !i->u.pipefs.i_readers) {
197:                         return 1;
198:                     }
199:                     break;
200:                 }
201:                 return 0;
202: }
```

## fs/pipefs/super.c

Page 1/2

```

1: /*
2:  * fiwix/fs/pipefs/super.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_pipe.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/sched.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: static unsigned int i_counter;
20:
21: struct fs_operations pipefs_fsop = {
22:     FSOP_KERN_MOUNT,
23:     PIPE_DEV,
24:
25:     fifo_open,
26:     pipefs_close,
27:     pipefs_read,
28:     pipefs_write,
29:     pipefs_ioctl,
30:     pipefs_lseek,
31:     NULL,                /* readdir */
32:     NULL,                /* mmap */
33:     pipefs_select,
34:
35:     NULL,                /* readlink */
36:     NULL,                /* followlink */
37:     NULL,                /* bmap */
38:     NULL,                /* lookup */
39:     NULL,                /* rmdir */
40:     NULL,                /* link */
41:     NULL,                /* unlink */
42:     NULL,                /* symlink */
43:     NULL,                /* mkdir */
44:     NULL,                /* mknod */
45:     NULL,                /* truncate */
46:     NULL,                /* create */
47:     NULL,                /* rename */
48:
49:     NULL,                /* read_block */
50:     NULL,                /* write_block */
51:
52:     NULL,                /* read_inode */
53:     NULL,                /* write_inode */
54:     pipefs_ialloc,
55:     pipefs_ifree,
56:     NULL,                /* statfs */
57:     pipefs_read_superblock,
58:     NULL,                /* remount_fs */
59:     NULL,                /* write_superblock */
60:     NULL,                /* release_superblock */
61: };
62:
63: int pipefs_read_superblock(__dev_t dev, struct superblock *sb)
64: {
65:     superblock_lock(sb);
66:     sb->dev = dev;
67:     sb->fsop = &pipefs_fsop;

```

## fs/pipefs/super.c

Page 2/2

```
68:         sb->s_blocksize = BLKSIZE_1K;
69:         i_counter = 0;
70:         superblock_unlock(sb);
71:         return 0;
72: }
73:
74: int pipefs_ialloc(struct inode *i, int mode)
75: {
76:     struct superblock *sb = i->sb;
77:
78:     superblock_lock(sb);
79:     i_counter++;
80:     superblock_unlock(sb);
81:
82:     i->i_mode = S_IFIFO;
83:     i->dev = i->rdev = sb->dev;
84:     i->fsop = &pipefs_fsop;
85:     i->inode = i_counter;
86:     i->count = 2;
87:     if(!(i->u.pipefs.i_data = (void *)kmalloc())) {
88:         return -ENOMEM;
89:     }
90:     i->u.pipefs.i_readoff = 0;
91:     i->u.pipefs.i_writeoff = 0;
92:     i->u.pipefs.i_readers = 1;
93:     i->u.pipefs.i_writers = 1;
94:     return 0;
95: }
96:
97: void pipefs_ifree(struct inode *i)
98: {
99:     if(!i->u.pipefs.i_readers && !i->u.pipefs.i_writers) {
100:         /*
101:          * We need to ask before to kfree() because this function is
102:          * also called to free removed (with sys_unlink) fifo files.
103:          */
104:         if(i->u.pipefs.i_data) {
105:             kfree((unsigned int)i->u.pipefs.i_data);
106:         }
107:     }
108: }
109:
110: int pipefs_init(void)
111: {
112:     return register_filesystem("pipefs", &pipefs_fsop);
113: }
```

## fs/procfs/data.c

Page 1/14

```

1: /*
2:  * fiwix/fs/procfs/data.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/system.h>
10: #include <fiwix/types.h>
11: #include <fiwix/process.h>
12: #include <fiwix/cmos.h>
13: #include <fiwix/dma.h>
14: #include <fiwix/ide.h>
15: #include <fiwix/fs.h>
16: #include <fiwix/filesystems.h>
17: #include <fiwix/devices.h>
18: #include <fiwix/locks.h>
19: #include <fiwix/mm.h>
20: #include <fiwix/mman.h>
21: #include <fiwix/fs_proc.h>
22: #include <fiwix/cpu.h>
23: #include <fiwix/pic.h>
24: #include <fiwix/sched.h>
25: #include <fiwix/timer.h>
26: #include <fiwix/utsname.h>
27: #include <fiwix/version.h>
28: #include <fiwix/errno.h>
29: #include <fiwix/stdio.h>
30: #include <fiwix/string.h>
31:
32: #define FSHIFT16          16
33: #define FIXED16_1        (1 << FSHIFT16)
34: #define LOAD_INT(x)      ((x) >> FSHIFT16)
35: #define LOAD_FRAC(x)     LOAD_INT(((x) & (FIXED16_1 - 1)) * 100)
36:
37: static const char *pstate[] = {
38:     "? (unused!)",
39:     "R (running)",
40:     "S (sleeping)",
41:     "Z (zombie)",
42:     "T (stopped)",
43:     "D (idle)",
44: };
45:
46: /*
47:  * procfs root directory related functions
48:  * -----
49:  */
50: int data_proc_self(char *buffer, __pid_t pid)
51: {
52:     return sprintk(buffer, "%s", current->pidstr);
53: }
54:
55: int data_proc_cmdline(char *buffer, __pid_t pid)
56: {
57:     return sprintk(buffer, "%s\n", cmdline);
58: }
59:
60: int data_proc_cpuinfo(char *buffer, __pid_t pid)
61: {
62:     int size;
63:
64:     size = sprintk(buffer, "processor          : 0\n");
65:     size += sprintk(buffer + size, "cpu family          : %d86\n", cpu_table.family
ily <= 6 ? cpu_table.family : 6);
66:     if(cpu_table.model >= 0) {

```

## fs/procfs/data.c

Page 2/14

```

67:             size += sprintf(buffer + size, "model           : %d\n", cpu_tab
le.model);
68:         } else {
69:             size += sprintf(buffer + size, "model           : unknown\n");
70:         }
71:
72:         if(cpu_table.vendor_id) {
73:             size += sprintf(buffer + size, "vendor_id        : %s\n", cpu_tab
le.vendor_id);
74:         }
75:         if(cpu_table.model_name) {
76:             size += sprintf(buffer + size, "model name       : %s\n", cpu_tab
le.model_name);
77:         }
78:         if(cpu_table.stepping >= 0) {
79:             size += sprintf(buffer + size, "stepping         : %d\n", cpu_tab
le.stepping);
80:         } else {
81:             size += sprintf(buffer + size, "stepping         : unknown\n");
82:         }
83:
84:         size += sprintf(buffer + size, "cpu MHz           : ");
85:         if(cpu_table.hz) {
86:             size += sprintf(buffer + size, "%d.%d\n", (cpu_table.hz / 100000
0), ((cpu_table.hz % 1000000) / 100000));
87:         } else {
88:             size += sprintf(buffer + size, "unknown\n");
89:         }
90:         if(cpu_table.cache) {
91:             size += sprintf(buffer + size, "cache size      : %s\n", cpu_tab
le.cache);
92:         }
93:         size += sprintf(buffer + size, "cpuid            : %s\n", cpu_table.has_cp
uid ? "yes" : "no");
94:         size += sprintf(buffer + size, "fpu              : %s\n", cpu_table.has_fp
u ? "yes" : "no");
95:         size += get_cpu_flags(buffer, size);
96:         return size;
97:     }
98:
99: int data_proc_devices(char *buffer, __pid_t pid)
100: {
101:     int n, size;
102:     struct device *d;
103:
104:     size = sprintf(buffer, "Character devices:\n");
105:     for(n = 0; n < NR_CHRDEV; n++) {
106:         d = chr_device_table[n];
107:         while(d) {
108:             size += sprintf(buffer + size, "%3d %s\n", d->major, d->
name);
109:             d = d->next;
110:         }
111:     }
112:
113:     size += sprintf(buffer + size, "\nBlock devices:\n");
114:     for(n = 0; n < NR_BLKDEV; n++) {
115:         d = blk_device_table[n];
116:         while(d) {
117:             size += sprintf(buffer + size, "%3d %s\n", d->major, d->
name);
118:             d = d->next;
119:         }
120:     }
121:     return size;
122: }
123:

```

## fs/procfs/data.c

Page 3/14

```

124: int data_proc_dma(char *buffer, __pid_t pid)
125: {
126:     int n, size;
127:
128:     size = 0;
129:     for(n = 0; n < DMA_CHANNELS; n++) {
130:         if(dma_resources[n]) {
131:             size += sprintf(buffer + size, "%2d: %s\n", n, dma_resou
rces[n]);
132:         }
133:     }
134:     return size;
135: }
136:
137: int data_proc_filesystems(char *buffer, __pid_t pid)
138: {
139:     int n, size;
140:     int nodev;
141:
142:     size = 0;
143:     for(n = 0; n < NR_FILESYSTEMS; n++) {
144:         if(filesystems_table[n].name) {
145:             nodev = 0;
146:             if(filesystems_table[n].fsop->flags != FSOP_REQUIRES_DEV
) {
147:                 nodev = 1;
148:             }
149:             size += sprintf(buffer + size, "%s %s\n", nodev ? "nodev
" : "", filesystems_table[n].name);
150:         }
151:     }
152:     return size;
153: }
154:
155: int data_proc_interrupts(char *buffer, __pid_t pid)
156: {
157:     struct interrupt *irq;
158:     int n, size;
159:
160:     size = 0;
161:     for(n = 0; n < NR_IRQS; n++) {
162:         if((irq = irq_table[n])) {
163:             size += sprintf(buffer + size, "%3d: %9u %s", n, irq->ti
cks, irq->name);
164:             while((irq = irq->next)) {
165:                 size += sprintf(buffer + size, ",%s", irq->name)
;
166:             }
167:             size += sprintf(buffer + size, "\n");
168:         }
169:     }
170:     size += sprintf(buffer + size, "SPU: %9u %s\n", kstat.sirqs, "Spurious i
nterrupts");
171:     return size;
172: }
173:
174: int data_proc_loadavg(char *buffer, __pid_t pid)
175: {
176:     int a, b, c;
177:     int size;
178:     struct proc *p;
179:     int nrun = 0;
180:     int nprocs = 0;
181:
182:     a = avenrun[0] << (SI_LOAD_SHIFT - FSHIFT);
183:     b = avenrun[1] << (SI_LOAD_SHIFT - FSHIFT);
184:     c = avenrun[2] << (SI_LOAD_SHIFT - FSHIFT);

```

## fs/procfs/data.c

Page 4/14

```

185:
186:     FOR_EACH_PROCESS(p) {
187:         if(p->state) {
188:             nprocs++;
189:             if(p->state == PROC_RUNNING) {
190:                 nrun++;
191:             }
192:         }
193:     }
194:
195:     size = sprintf(buffer, "%d.%02d %d.%02d %d.%02d %d/%d %d\n", LOAD_INT(a)
, LOAD_FRAC(a), LOAD_INT(b), LOAD_FRAC(b), LOAD_INT(c), LOAD_FRAC(c), nrun, nprocs, las
tpid);
196:     return size;
197: }
198:
199: int data_proc_locks(char *buffer, __pid_t pid)
200: {
201:     int n, size;
202:     struct flock_file *ff;
203:
204:     size = 0;
205:
206:     for(n = 0; n < NR_FLOCKS; n++) {
207:         ff = &flock_file_table[n];
208:         if(ff->inode) {
209:             size += sprintf(buffer + size, "%d: FLOCK ADVISORY %s
", n + 1, ff->type & LOCK_SH ? "READ " : "WRITE");
210:             size += sprintf(buffer + size, "%d %x:%d:%d 0 EOF\n", ff
->proc->pid, MAJOR(ff->inode->dev), MINOR(ff->inode->dev), ff->inode->inode);
211:         }
212:     }
213:
214:     return size;
215: }
216:
217: int data_proc_meminfo(char *buffer, __pid_t pid)
218: {
219:     struct page *pg;
220:     int n, size;
221:
222:     kstat.shared = 0;
223:     for(n = 0; n < kstat.physical_pages; n++) {
224:         pg = &page_table[n];
225:         if(pg->flags & PAGE_RESERVED) {
226:             continue;
227:         }
228:         if(!pg->count) {
229:             continue;
230:         }
231:         kstat.shared += pg->count - 1;
232:     }
233:
234:     size = 0;
235:     size += sprintf(buffer + size, "          total:    used:    free:    shared
: buffers:  cached:\n");
236:     size += sprintf(buffer + size, "Mem:  %8u %8u %8u %8u %8u %8u\n", kstat.
total_mem_pages << PAGE_SHIFT, (kstat.total_mem_pages << PAGE_SHIFT) - (kstat.free_page
s << PAGE_SHIFT), kstat.free_pages << PAGE_SHIFT, kstat.shared * 1024, kstat.buffers *
1024, kstat.cached * 1024);
237:     size += sprintf(buffer + size, "Swap: %8u %8u %8u\n", 0, 0, 0);
238:     size += sprintf(buffer + size, "MemTotal: %9d kB\n", kstat.total_mem_pag
es << 2);
239:     size += sprintf(buffer + size, "MemFree:  %9d kB\n", kstat.free_pages <<
2);
240:     size += sprintf(buffer + size, "MemShared:%9d kB\n", kstat.shared);
241:     size += sprintf(buffer + size, "Buffers:  %9d kB\n", kstat.buffers);

```

## fs/procfs/data.c

Page 5/14

```

242:         size += sprintf(buffer + size, "Cached:  %9d kB\n", kstat.cached);
243:         size += sprintf(buffer + size, "SwapTotal:%9d kB\n", 0);
244:         size += sprintf(buffer + size, "SwapFree:  %9d kB\n", 0);
245:         return size;
246:     }
247:
248: int data_proc_mounts(char *buffer, __pid_t pid)
249: {
250:     int n, size;
251:     char *flag;
252:
253:     size = 0;
254:     for(n = 0; n < NR_MOUNT_POINTS; n++) {
255:         if(mount_table[n].used) {
256:             if(mount_table[n].fs->fsop->flags != FSOP_KERN_MOUNT) {
257:                 flag = "rw";
258:                 if(mount_table[n].sb.flags & MS_RDONLY) {
259:                     flag = "ro";
260:                 }
261:                 size += sprintf(buffer + size, "%s %s %s %s 0 0\
n", mount_table[n].devname, mount_table[n].dirname, mount_table[n].fs->name, flag);
262:             }
263:         }
264:     }
265:     return size;
266: }
267:
268: int data_proc_partitions(char *buffer, __pid_t pid)
269: {
270:     int n, ctrl, drv, size;
271:     int minor, major;
272:     unsigned int blocks;
273:     struct ide *ide;
274:     struct ide_drv *drive;
275:
276:     size = 0;
277:     size += sprintf(buffer + size, "major minor #blocks name\n\n");
278:
279:     for(ctrl = 0; ctrl < NR_IDE_CTRL; ctrl++) {
280:         ide = &ide_table[ctrl];
281:         for(drv = 0; drv < NR_IDE_DRVS; drv++) {
282:             drive = &ide->drive[drv];
283:             if(!drive->nr_sects) {
284:                 continue;
285:             }
286:             if(drive->flags & DEVICE_IS_DISK) {
287:                 major = (int)drive->major;
288:                 minor = (int)drive->minor_shift;
289:                 blocks = drive->nr_sects / 2;
290:                 size += sprintf(buffer + size, "%4d %4d %9d %s
\n", major, 0, blocks, drive->dev_name);
291:                 for(n = 0; n < NR_PARTITIONS; n++) {
292:                     if(drive->part_table[n].type) {
293:                         blocks = drive->part_table[n].nr
_sects / 2;
294:                         size += sprintf(buffer + size, "
%4d %4d %9u %s%d\n", major, (n + 1) << minor, blocks, drive->dev_name, n + 1);
295:                     }
296:                 }
297:             }
298:         }
299:     }
300:     return size;
301: }
302:
303: int data_proc_rtc(char *buffer, __pid_t pid)
304: {

```



## fs/procfs/data.c

Page 6/14

```

305:         int size;
306:         short int sec, min, hour;
307:         short int day, month, year, century;
308:
309:         sec = cmos_read_date(CMOS_SEC);
310:         min = cmos_read_date(CMOS_MIN);
311:         hour = cmos_read_date(CMOS_HOUR);
312:         day = cmos_read_date(CMOS_DAY);
313:         month = cmos_read_date(CMOS_MONTH);
314:         year = cmos_read_date(CMOS_YEAR);
315:         century = cmos_read_date(CMOS_CENTURY);
316:         year += century * 100;
317:
318:         size = 0;
319:         size += sprintf(buffer + size, "rtc_time\t: %02d:%02d:%02d\n", hour, min
, sec);
320:         size += sprintf(buffer + size, "rtc_date\t: %02d-%02d-%02d\n", year, mon
th, day);
321:         sec = cmos_read_date(CMOS_ASEC);
322:         min = cmos_read_date(CMOS_AMIN);
323:         hour = cmos_read_date(CMOS_AHOUR);
324:         size += sprintf(buffer + size, "alarm\t\t: %02d:%02d:%02d\n", hour, min,
sec);
325:         size += sprintf(buffer + size, "DST_enable\t: %s\n", cmos_read(CMOS_STAT
B) & CMOS_STATB_DSE ? "yes" : "no");
326:         size += sprintf(buffer + size, "BCD\t\t: %s\n", cmos_read(CMOS_STATB) &
CMOS_STATB_DM ? "no" : "yes");
327:         size += sprintf(buffer + size, "24hr\t\t: %s\n", cmos_read(CMOS_STATB) &
CMOS_STATB_24H ? "yes" : "no");
328:         size += sprintf(buffer + size, "square_wave\t: %s\n", cmos_read(CMOS_STA
TB) & CMOS_STATB_SQWE ? "yes" : "no");
329:         size += sprintf(buffer + size, "alarm_IRQ\t: %s\n", cmos_read(CMOS_STATB
) & CMOS_STATB_AIE ? "yes" : "no");
330:         size += sprintf(buffer + size, "update_IRQ\t: %s\n", cmos_read(CMOS_STAT
B) & CMOS_STATB_UIE ? "yes" : "no");
331:         size += sprintf(buffer + size, "periodic_IRQ\t: %s\n", cmos_read(CMOS_ST
ATB) & CMOS_STATB_PIE ? "yes" : "no");
332:         size += sprintf(buffer + size, "periodic_freq\t: %s\n", (cmos_read(CMOS_
STATA) & CMOS_STATATA_IRQF) == 0x6 ? "1024" : "?");
333:         size += sprintf(buffer + size, "batt_status\t: %s\n", cmos_read(CMOS_STA
TD) & CMOS_STATD_VRT ? "okay" : "dead");
334:         return size;
335:     }
336:
337: int data_proc_stat(char *buffer, __pid_t pid)
338: {
339:     int n, size;
340:     unsigned int idle;
341:     struct interrupt *irq;
342:
343:     idle = kstat.ticks - (kstat.cpu_user + kstat.cpu_nice + kstat.cpu_system
);
344:     size = 0;
345:     size += sprintf(buffer + size, "cpu %d %d %d %d\n", kstat.cpu_user, ksta
t.cpu_nice, kstat.cpu_system, idle);
346:     size += sprintf(buffer + size, "disk 0 0 0 0\n");
347:     size += sprintf(buffer + size, "page 0 0\n");
348:     size += sprintf(buffer + size, "swap 0 0\n");
349:     size += sprintf(buffer + size, "intr %u", kstat.irqs);
350:     for(n = 0; n < NR_IRQS; n++) {
351:         irq = irq_table[n];
352:         if(irq) {
353:             size += sprintf(buffer + size, " %u", irq->ticks);
354:         }
355:     }
356:     size += sprintf(buffer + size, "\n");
357:     size += sprintf(buffer + size, "ctxt %u\n", kstat.ctxt);

```

## fs/procfs/data.c

Page 7/14

```

358:         size += sprintf(buffer + size, "btime %d\n", kstat.boot_time);
359:         size += sprintf(buffer + size, "processes %d\n", kstat.processes);
360:         return size;
361:     }
362:
363: int data_proc_uptime(char *buffer, __pid_t pid)
364: {
365:     struct proc *p;
366:     unsigned long int idle;
367:
368:     p = &proc_table[IDLE];
369:     idle = tv2ticks(&p->usage.ru_utime);
370:     idle += tv2ticks(&p->usage.ru_stime);
371:     return sprintf(buffer, "%u.%02u %u.%02u\n", kstat.uptime, kstat.ticks %
HZ, idle / HZ, idle % HZ);
372: }
373:
374: int data_proc_fullversion(char *buffer, __pid_t pid)
375: {
376:     return sprintf(buffer, "Fiwix version %s %s\n", UTS_RELEASE, UTS_VERSION
);
377: }
378:
379: int data_proc_domainname(char *buffer, __pid_t pid)
380: {
381:     return sprintf(buffer, "%s\n", sys_utsname.domainname);
382: }
383:
384: int data_proc_filemax(char *buffer, __pid_t pid)
385: {
386:     return sprintf(buffer, "%d\n", NR_OPENS);
387: }
388:
389: int data_proc_filennr(char *buffer, __pid_t pid)
390: {
391:     int n, nr;
392:
393:     nr = 0;
394:     for(n = 1; n < NR_OPENS; n++) {
395:         if(fd_table[n].count != 0) {
396:             nr++;
397:         }
398:     }
399:     return sprintf(buffer, "%d\n", nr);
400: }
401:
402: int data_proc_hostname(char *buffer, __pid_t pid)
403: {
404:     return sprintf(buffer, "%s\n", sys_utsname.nodename);
405: }
406:
407: int data_proc_inodemax(char *buffer, __pid_t pid)
408: {
409:     return sprintf(buffer, "%d\n", inode_table_size / sizeof(struct inode));
410: }
411:
412: int data_proc_inodenr(char *buffer, __pid_t pid)
413: {
414:     return sprintf(buffer, "%d\n", (inode_table_size / sizeof(struct inode)
- inodes_on_free_list);
415: }
416:
417: int data_proc_osrelease(char *buffer, __pid_t pid)
418: {
419:     return sprintf(buffer, "%s\n", UTS_RELEASE);
420: }
421:

```

## fs/procfs/data.c

Page 8/14

```

422: int data_proc_ostype(char *buffer, __pid_t pid)
423: {
424:     return sprintf(buffer, "%s\n", UTS_SYSNAME);
425: }
426:
427: int data_proc_version(char *buffer, __pid_t pid)
428: {
429:     return sprintf(buffer, "%s\n", UTS_VERSION);
430: }
431:
432:
433: /*
434:  * PID directory related functions
435:  * -----
436:  */
437: int data_proc_pid_cmdline(char *buffer, __pid_t pid)
438: {
439:     int n, size;
440:     char *arg;
441:     char **argv;
442:     unsigned int paddr, offset;
443:     struct proc *p;
444:
445:     size = 0;
446:     if((p = get_proc_by_pid(pid)) {
447:         if(p->argv) {
448:             offset = (int)p->argv & ~PAGE_MASK;
449:             paddr = get_mapped_addr(p, (int)p->argv) & PAGE_MASK;
450:             paddr = P2V(paddr);
451:             argv = (char **) (paddr + offset);
452:             for(n = 0; argv[n]; n++) {
453:                 offset = (int)argv[n] & ~PAGE_MASK;
454:                 paddr = get_mapped_addr(p, (int)argv[n]) & PAGE_
MASK;
455:                 paddr = P2V(paddr);
456:                 arg = (char *) (paddr + offset);
457:                 size += sprintf(buffer + size, "%s", arg);
458:                 buffer[size++] = NULL;
459:             }
460:         }
461:     }
462:     return size;
463: }
464:
465: int data_proc_pid_cwd(char *buffer, __pid_t pid)
466: {
467:     int size;
468:     struct proc *p;
469:     struct inode *i;
470:
471:     size = 0;
472:     if((p = get_proc_by_pid(pid)) {
473:
474:         /* zombie processes don't have current working directory */
475:         if(!p->pwd) {
476:             return -ENOENT;
477:         }
478:
479:         i = p->pwd;
480:         size = sprintf(buffer, "[%02d%02d]:%d", MAJOR(i->rdev), MINOR(i-
>rdev), i->inode);
481:     }
482:     return size;
483: }
484:
485: int data_proc_pid_envron(char *buffer, __pid_t pid)
486: {

```

## fs/procfs/data.c

Page 9/14

```

487:     int n, size;
488:     char *env;
489:     char **envp;
490:     unsigned int paddr, offset;
491:     struct proc *p;
492:
493:     size = 0;
494:     if((p = get_proc_by_pid(pid)) {
495:         if(p->envp) {
496:             offset = (int)p->envp & ~PAGE_MASK;
497:             paddr = get_mapped_addr(p, (int)p->envp) & PAGE_MASK;
498:             paddr = P2V(paddr);
499:             envp = (char **) (paddr + offset);
500:             for(n = 0; envp[n]; n++) {
501:                 offset = (int)envp[n] & ~PAGE_MASK;
502:                 paddr = get_mapped_addr(p, (int)envp[n]) & PAGE_
MASK;
503:                 paddr = P2V(paddr);
504:                 env = (char *) (paddr + offset);
505:                 size += sprintf(buffer + size, "%s", env);
506:                 buffer[size++] = NULL;
507:             }
508:         }
509:     }
510:     return size;
511: }
512:
513: int data_proc_pid_exe(char *buffer, __pid_t pid)
514: {
515:     int size;
516:     struct proc *p;
517:     struct inode *i;
518:
519:     size = 0;
520:     if((p = get_proc_by_pid(pid)) {
521:
522:         /* kernel and zombie processes are programless */
523:         if(!p->vma || !p->vma->inode) {
524:             return -ENOENT;
525:         }
526:
527:         i = p->vma->inode;
528:         size = sprintf(buffer, "[%02d%02d]:%d", MAJOR(i->rdev), MINOR(i-
>rdev), i->inode);
529:     }
530:     return size;
531: }
532:
533: int data_proc_pid_maps(char *buffer, __pid_t pid)
534: {
535:     unsigned int n;
536:     int size, len;
537:     __ino_t inode;
538:     int major, minor;
539:     char *section;
540:     char r, w, x, f;
541:     struct proc *p;
542:     struct vma *vma;
543:
544:     size = 0;
545:     if((p = get_proc_by_pid(pid)) {
546:         if(!p->vma) {
547:             return 0;
548:         }
549:         vma = p->vma;
550:         for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
551:             r = vma->prot & PROT_READ ? 'r' : '-';

```

## fs/procfs/data.c

Page 10/14

```

552:         w = vma->prot & PROT_WRITE ? 'w' : '-';
553:         x = vma->prot & PROT_EXEC ? 'x' : '-';
554:         if(vma->flags & MAP_SHARED) {
555:             f = 's';
556:         } else if(vma->flags & MAP_PRIVATE) {
557:             f = 'p';
558:         } else {
559:             f = '-';
560:         }
561:         switch(vma->s_type) {
562:             case P_TEXT:         section = "text";
563:                                 break;
564:             case P_DATA:         section = "data";
565:                                 break;
566:             case P_BSS:          section = "bss";
567:                                 break;
568:             case P_HEAP:         section = "heap";
569:                                 break;
570:             case P_STACK:       section = "stack";
571:                                 break;
572:             case P_MMAP:         section = "mmap";
573:                                 break;
574:             default:
575:                 section = NULL;
576:                 break;
577:         }
578:         inode = major = minor = 0;
579:         if(vma->inode) {
580:             inode = vma->inode->inode;
581:             major = MAJOR(vma->inode->dev);
582:             minor = MINOR(vma->inode->dev);
583:         }
584:         len = sprintf(buffer + size, "%08x-%08x %c%c%c%c %08x %0
2d:%02d %- 10u [%s]\n", vma->start, vma->end, r, w, x, f, vma->offset, major, minor, in
ode, section);
585:         size += len;
586:     }
587: }
588: return size;
589: }
590:
591: int data_proc_pid_mountinfo(char *buffer, __pid_t pid)
592: {
593:     int n, size;
594:     char *flag, *devname;
595:
596:     size = 0;
597:     for(n = 0; n < NR_MOUNT_POINTS; n++) {
598:         if(mount_table[n].used) {
599:             if(mount_table[n].fs->fsop->flags != FSOP_KERN_MOUNT) {
600:                 flag = "rw";
601:                 if(mount_table[n].sb.flags & MS_RDONLY) {
602:                     flag = "ro";
603:                 }
604:                 devname = mount_table[n].devname;
605:                 if(!strcmp(mount_table[n].devname, "/dev/root"))
{
606:                     devname = _rootdevname;
607:                 }
608:                 size += sprintf(buffer + size, "%d 0 %d:%d %s %s
%s - %s %s %s\n", n, MAJOR(mount_table[n].dev), MINOR(mount_table[n].dev), "/", mount_
table[n].dirname, flag, mount_table[n].fs->name, devname, flag);
609:             }
610:         }
611:     }
612:     return size;
613: }

```

```

614:
615: int data_proc_pid_root(char *buffer, __pid_t pid)
616: {
617:     int size;
618:     struct proc *p;
619:     struct inode *i;
620:
621:     size = 0;
622:     if((p = get_proc_by_pid(pid)) {
623:
624:         /* zombie processes don't have root directory */
625:         if(!p->root) {
626:             return -ENOENT;
627:         }
628:
629:         i = p->root;
630:         size = sprintf(buffer, "[%02d%02d]:%d", MAJOR(i->rdev), MINOR(i-
>rdev), i->inode);
631:     }
632:     return size;
633: }
634:
635: int data_proc_pid_stat(char *buffer, __pid_t pid)
636: {
637:     int n, size, vma_start, vma_end;
638:     unsigned int esp, eip;
639:     int signum, mask;
640:     __sigset_t sigignored, sigcaught;
641:     struct proc *p;
642:     struct sigcontext *sc;
643:     struct vma *vma;
644:     int text, data, stack, mmap;
645:
646:     size = vma_start = vma_end = 0;
647:     text = data = stack = mmap = 0;
648:     if((p = get_proc_by_pid(pid)) {
649:         if(!p->vma) {
650:             return 0;
651:         }
652:         vma_start = p->vma[0].start;
653:         vma_end = p->vma[0].end;
654:
655:         vma = p->vma;
656:         for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
657:             switch(vma->s_type) {
658:                 case P_TEXT:
659:                     text += vma->end - vma->start;
660:                     break;
661:                 case P_HEAP:
662:                     data += vma->end - vma->start;
663:                     break;
664:                 case P_STACK:
665:                     stack += vma->end - vma->start;
666:                     break;
667:                 case P_MMAP:
668:                     mmap += vma->end - vma->start;
669:                     break;
670:             }
671:         }
672:
673:         sigignored = sigcaught = 0;
674:         for(signum = 0, mask = 1; signum < NSIG; signum++, mask <= 1) {
675:             if(p->sigaction[signum].sa_handler == SIG_IGN) {
676:                 sigignored |= mask;
677:             }
678:             if(p->sigaction[signum].sa_handler == SIG_DFL) {
679:                 sigcaught |= mask;

```

```

680:         }
681:     }
682:
683:     esp = eip = 0;
684:     if(p->sp) {
685:         sc = (struct sigcontext *)p->sp;
686:         esp = sc->oldesp;
687:         eip = sc->eip;
688:     }
689:     size = sprintf(buffer, "%d (%s) %c %d %d %d %d %d %d %d %d %d
%u %u %u %u %d %d %d %d %d %d %u %u %u %u %u %u %u %d %d %u %u %u\n",
690:         p->pid,
691:         p->argv0,
692:         pstate[p->state][0],
693:         p->ppid, p->pgid, p->sid,
694:         p->ctty ? p->ctty->dev : 0,
695:         p->ctty ? p->ctty->pgid : - 1,
696:         0, /* flags */
697:         0, 0, 0, 0, /* minflt, cminflt, majflt, cmajflt */
698:         tv2ticks(&p->usage.ru_utime),
699:         tv2ticks(&p->usage.ru_stime),
700:         tv2ticks(&p->cusage.ru_utime),
701:         tv2ticks(&p->cusage.ru_stime),
702:         0, /* counter */
703:         0, /* priority */
704:         0, /* timeout */
705:         0, /* itrealvalue */
706:         p->start_time,
707:         text + data + stack + mmap,
708:         p->rss,
709:         0x7FFFFFFF, /* rlim */
710:         vma_start, /* startcode */
711:         vma_end, /* endcode */
712:         KERNEL_BASE_ADDR - 1, /* startstack */
713:         esp, /* kstkesp */
714:         eip, /* kstkeip */
715:         p->sigpending,
716:         p->sigblocked,
717:         sigignored,
718:         sigcaught,
719:         p->sleep_address
720:     );
721: }
722: return size;
723: }
724:
725: int data_proc_pid_statm(char *buffer, __pid_t pid)
726: {
727:     int n, size;
728:     struct proc *p;
729:     struct vma *vma;
730:     int text, data, stack, mmap;
731:
732:     size = text = data = stack = mmap = 0;
733:     if((p = get_proc_by_pid(pid)) ) {
734:         if(!p->vma) {
735:             return 0;
736:         }
737:         vma = p->vma;
738:         for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
739:             switch(vma->s_type) {
740:                 case P_TEXT:
741:                     text += vma->end - vma->start;
742:                     break;
743:                 case P_HEAP:
744:                     data += vma->end - vma->start;
745:                     break;

```

## fs/procfs/data.c

Page 13/14

```

746:                 case P_STACK:
747:                     stack += vma->end - vma->start;
748:                     break;
749:                 case P_MMAP:
750:                     mmap += vma->end - vma->start;
751:                     break;
752:             }
753:         }
754:
755:         size = sprintf(buffer, "%d", (text + data + stack + mmap) / PAGE
_SIZE);
756:         size += sprintf(buffer + size, " %d", p->rss);
757:         size += sprintf(buffer + size, " 0"); /* shared mappings */
758:         size += sprintf(buffer + size, " %d", text / PAGE_SIZE);
759:         size += sprintf(buffer + size, " 0");
760:         size += sprintf(buffer + size, " %d", (data + stack) / PAGE_SIZE
);
761:         size += sprintf(buffer + size, " 0\n");
762:     }
763:     return size;
764: }
765:
766: int data_proc_pid_status(char *buffer, __pid_t pid)
767: {
768:     int n, size;
769:     int signum, mask;
770:     __sigset_t sigignored, sigcaught;
771:     struct proc *p;
772:     struct vma *vma;
773:     int text, data, stack, mmap;
774:
775:     size = text = data = stack = mmap = 0;
776:     if((p = get_proc_by_pid(pid)) {
777:         if(!p->vma) {
778:             return 0;
779:         }
780:         vma = p->vma;
781:         for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
782:             switch(vma->s_type) {
783:                 case P_TEXT:
784:                     text += vma->end - vma->start;
785:                     break;
786:                 case P_HEAP:
787:                     data += vma->end - vma->start;
788:                     break;
789:                 case P_STACK:
790:                     stack += vma->end - vma->start;
791:                     break;
792:                 case P_MMAP:
793:                     mmap += vma->end - vma->start;
794:                     break;
795:             }
796:         }
797:
798:         size = sprintf(buffer, "Name:\t%s\n", p->argv0);
799:         size += sprintf(buffer + size, "State:\t%s\n", pstate[p->state])
;
800:         size += sprintf(buffer + size, "Pid:\t%d\n", p->pid);
801:         size += sprintf(buffer + size, "PPid:\t%d\n", p->ppid);
802:         size += sprintf(buffer + size, "Uid:\t%d\t%d\t%d\t-\n", p->uid,
p->euid, p->suid);
803:         size += sprintf(buffer + size, "Gid:\t%d\t%d\t%d\t-\n", p->gid,
p->egid, p->sgid);
804:         size += sprintf(buffer + size, "VmSize:\t%8d kB\n", (text + data
+ stack + mmap) / 1024);
805:         size += sprintf(buffer + size, "VmLck:\t%8d kB\n", 0);
806:         size += sprintf(buffer + size, "VmRSS:\t%8d kB\n", p->rss << 2);

```



## fs/procfs/data.c

Page 14/14

```
807:         size += sprintf(buffer + size, "VmData:\t%8d kB\n", data / 1024)
;
808:         size += sprintf(buffer + size, "VmStk:\t%8d kB\n", stack / 1024)
;
809:         size += sprintf(buffer + size, "VmExe:\t%8d kB\n", text / 1024);
810:         size += sprintf(buffer + size, "VmLib:\t%8d kB\n", 0);
811:         size += sprintf(buffer + size, "SigPnd:\t%08x\n", p->sigpending)
;
812:         size += sprintf(buffer + size, "SigBlk:\t%08x\n", p->sigblocked)
;
813:         sigignored = sigcaught = 0;
814:         for(signum = 0, mask = 1; signum < NSIG; signum++, mask <<= 1) {
815:             if(p->sigaction[signum].sa_handler == SIG_IGN) {
816:                 sigignored |= mask;
817:             }
818:             if(p->sigaction[signum].sa_handler == SIG_DFL) {
819:                 sigcaught |= mask;
820:             }
821:         }
822:         size += sprintf(buffer + size, "SigIgn:\t%08x\n", sigignored);
823:         size += sprintf(buffer + size, "SigCgt:\t%08x\n", sigcaught);
824:     }
825:     return size;
826: }
```

## fs/procfs/dir.c

Page 1/4

```

1: /*
2:  * fiwix/fs/procfs/dir.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_proc.h>
13: #include <fiwix/dirent.h>
14: #include <fiwix/stat.h>
15: #include <fiwix/mm.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: struct fs_operations procfs_dir_fsop = {
20:     0,
21:     0,
22:
23:     procfs_dir_open,
24:     procfs_dir_close,
25:     procfs_dir_read,
26:     NULL, /* write */
27:     NULL, /* ioctl */
28:     NULL, /* lseek */
29:     procfs_dir_readdir,
30:     NULL, /* mmap */
31:     NULL, /* select */
32:
33:     NULL, /* readlink */
34:     NULL, /* followlink */
35:     procfs_bmap,
36:     procfs_lookup,
37:     NULL, /* rmdir */
38:     NULL, /* link */
39:     NULL, /* unlink */
40:     NULL, /* symlink */
41:     NULL, /* mkdir */
42:     NULL, /* mknod */
43:     NULL, /* truncate */
44:     NULL, /* create */
45:     NULL, /* rename */
46:
47:     NULL, /* read_block */
48:     NULL, /* write_block */
49:
50:     NULL, /* read_inode */
51:     NULL, /* write_inode */
52:     NULL, /* ialloc */
53:     NULL, /* ifree */
54:     NULL, /* statfs */
55:     NULL, /* read_superblock */
56:     NULL, /* remount_fs */
57:     NULL, /* write_superblock */
58:     NULL, /* release_superblock */
59: };
60:
61: static int proc_listdir(char *buffer)
62: {
63:     int n;
64:     struct proc *p;
65:     struct procfs_dir_entry *pd;
66:     struct procfs_dir_entry d;
67:     int size;

```

## fs/procfs/dir.c

Page 2/4

```

68:
69:     size = 0;
70:     pd = (struct procfs_dir_entry *)buffer;
71:
72:     FOR_EACH_PROCESS(p) {
73:         if(p->state) {
74:             d.inode = PROC_PID_INO + (p->pid << 12);
75:             d.mode = S_IFDIR | S_IRUSR | S_IXUSR | S_IRGRP | S_IXGRP
| S_IROTH | S_IXOTH;
76:             d.nlink = 1;
77:             d.lev = -1;
78:             d.name_len = 1;
79:             n = p->pid;
80:             while(n / 10) {
81:                 n /= 10;
82:                 d.name_len++;
83:             }
84:             d.name = p->pidstr;
85:             d.data_fn = NULL;
86:
87:             if(size + sizeof(struct procfs_dir_entry) >= PAGE_SIZE)
{
88:                 printk("WARNING: kmalloc() is limited to 4096 by
tes.\n");
89:                 break;
90:             }
91:
92:             size += sizeof(struct procfs_dir_entry);
93:             memcpy_b((void *)pd, (void *)&d, sizeof(struct procfs_di
r_entry));
94:             pd++;
95:         }
96:     }
97:     return size;
98: }
99:
100: /*
101:  static int proc_listfd(struct inode *i, char *buffer)
102:  {
103:      int n;
104:      struct proc *p;
105:      struct procfs_dir_entry *pd;
106:      struct procfs_dir_entry d;
107:      int size;
108:
109:      size = 0;
110:      pd = (struct procfs_dir_entry *)buffer;
111:
112:      p = get_proc_by_pid((i->inode >> 12) & 0xFFFF);
113:      for(n = 0; n < OPEN_MAX; n++) {
114:          if(p->fd[n]) {
115:              d.inode = PROC_PID_INO + (p->pid << 12) + n;
116:              d.mode = S_IFREG | S_IRWXU;
117:              d.nlink = 1;
118:              d.lev = -1;
119:              d.name_len = sprintf(d.name, "%d", n);
120:              d.data_fn = NULL;
121:
122:              if(size + sizeof(struct procfs_dir_entry) >= 4096) {
123:                  printk("WARNING: kmalloc() is limited to 4096 by
tes.\n");
124:                  break;
125:              }
126:
127:              size += sizeof(struct procfs_dir_entry);
128:              memcpy_b((void *)pd, (void *)&d, sizeof(struct procfs_di
r_entry));

```

```

129:             pd++;
130:         }
131:     }
132:     memset_b((void *)pd + size, NULL, sizeof(struct procfs_dir_entry));
133:     return size;
134: }
135: */
136:
137: int procfs_dir_open(struct inode *i, struct fd *fd_table)
138: {
139:     fd_table->offset = 0;
140:     return 0;
141: }
142:
143: int procfs_dir_close(struct inode *i, struct fd *fd_table)
144: {
145:     return 0;
146: }
147:
148: int procfs_dir_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t
count)
149: {
150:     __off_t total_read;
151:     unsigned int bytes;
152:     int len, lev;
153:     char *buf;
154:
155:     if(!(buf = (void *)kmalloc())) {
156:         return -ENOMEM;
157:     }
158:
159:     /* create the list of directories for each process */
160:     len = 0;
161:     if(i->inode == PROC_ROOT_INO) {
162:         len = proc_listdir(buf);
163:     }
164:
165:     /* TODO: create the list of fds used for each process
166:     if((i->inode & 0xF0000FFF) == PROC_PID_FD) {
167:         len = proc_listfd(i, buf);
168:     }
169:     */
170:
171:     /* add the rest of static files in the main directory */
172:     lev = i->u.procfs.i_lev;
173:
174:     /* assigns the size of the level without the last entry (NULL) */
175:     bytes = sizeof(procfs_array[lev]) - sizeof(struct procfs_dir_entry);
176:
177:     if((len + bytes) > (PAGE_SIZE - 1)) {
178:         printk("WARNING: %s(): len > 4095 (%d).\n", __FUNCTION__, len);
179:     }
180:     memcpy_b(buf + len, (char *)&procfs_array[lev], bytes);
181:     len += bytes;
182:     total_read = fd_table->offset = len;
183:     memcpy_b(buffer, buf, PAGE_SIZE);
184:     kfree((unsigned int)buf);
185:     return total_read;
186: }
187:
188: int procfs_dir_readdir(struct inode *i, struct fd *fd_table, struct dirent *dire
nt, unsigned int count)
189: {
190:     unsigned int offset, boffset, dirent_offset, doffset;
191:     int dirent_len;
192:     unsigned int total_read;
193:     struct procfs_dir_entry *d;

```

## fs/procfs/dir.c

Page 4/4

```
194:     int base_dirent_len;
195:     char *buffer;
196:     int lev;
197:
198:     if(!i->fsop || !i->fsop->read) {
199:         return -EBADF;
200:     }
201:     if(!(buffer = (void *)kmalloc())) {
202:         return -ENOMEM;
203:     }
204:
205:     lev = i->u.procfs.i_lev;
206:     base_dirent_len = sizeof(dirent->d_ino) + sizeof(dirent->d_off) + sizeof
(dirent->d_reclen);
207:
208:     offset = fd_table->offset;
209:     boffset = dirent_offset = doffset = 0;
210:
211:     boffset = offset % PAGE_SIZE;
212:
213:     total_read = i->fsop->read(i, fd_table, buffer, PAGE_SIZE);
214:     if((count = MIN(total_read, count)) == 0) {
215:         kfree((unsigned int)buffer);
216:         return dirent_offset;
217:     }
218:
219:     while(boffset < fd_table->offset) {
220:         d = (struct procfs_dir_entry *) (buffer + boffset);
221:         if(!d->inode) {
222:             break;
223:         }
224:         dirent_len = (base_dirent_len + (d->name_len + 1)) + 3;
225:         dirent_len &= ~3; /* round up */
226:         if((doffset + dirent_len) <= count) {
227:             boffset += sizeof(struct procfs_dir_entry);
228:             offset += sizeof(struct procfs_dir_entry);
229:             doffset += dirent_len;
230:             dirent->d_ino = d->inode;
231:             dirent->d_off = offset;
232:             dirent->d_reclen = dirent_len;
233:             memcpy_b(dirent->d_name, d->name, d->name_len);
234:             dirent->d_name[d->name_len] = NULL;
235:             dirent = (struct dirent *) ((char *)dirent + dirent_len);
236:             dirent_offset += dirent_len;
237:         } else {
238:             break;
239:         }
240:     }
241:     fd_table->offset = boffset;
242:     kfree((unsigned int)buffer);
243:     return dirent_offset;
244: }
```

## fs/procfs/file.c

Page 1/2

```

1: /*
2:  * fiwix/fs/procfs/file.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_proc.h>
13: #include <fiwix/fcntl.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: struct fs_operations procfs_file_fsop = {
19:     0,
20:     0,
21:
22:     procfs_file_open,
23:     procfs_file_close,
24:     procfs_file_read,
25:     NULL, /* write */
26:     NULL, /* ioctl */
27:     procfs_file_lseek,
28:     NULL, /* readdir */
29:     NULL, /* mmap */
30:     NULL, /* select */
31:
32:     NULL, /* readlink */
33:     NULL, /* followlink */
34:     procfs_bmap,
35:     NULL, /* lookup */
36:     NULL, /* rmdir */
37:     NULL, /* link */
38:     NULL, /* unlink */
39:     NULL, /* symlink */
40:     NULL, /* mkdir */
41:     NULL, /* mknod */
42:     NULL, /* truncate */
43:     NULL, /* create */
44:     NULL, /* rename */
45:
46:     NULL, /* read_block */
47:     NULL, /* write_block */
48:
49:     NULL, /* read_inode */
50:     NULL, /* write_inode */
51:     NULL, /* ialloc */
52:     NULL, /* ifree */
53:     NULL, /* statfs */
54:     NULL, /* read_superblock */
55:     NULL, /* remount_fs */
56:     NULL, /* write_superblock */
57:     NULL, /* release_superblock */
58: };
59:
60: int procfs_file_open(struct inode *i, struct fd *fd_table)
61: {
62:     if(fd_table->flags & (O_WRONLY | O_RDWR | O_TRUNC | O_APPEND)) {
63:         return -EINVAL;
64:     }
65:     fd_table->offset = 0;
66:     return 0;
67: }

```

## fs/procfs/file.c

Page 2/2

```

68:
69: int procfs_file_close(struct inode *i, struct fd *fd_table)
70: {
71:     return 0;
72: }
73:
74: int procfs_file_read(struct inode *i, struct fd *fd_table, char *buffer, __size_
t count)
75: {
76:     __off_t total_read;
77:     unsigned int boffset, bytes, size;
78:     int blksize;
79:     struct procfs_dir_entry *d;
80:     char *buf;
81:
82:     if(!(d = get_procfs_by_inode(i))) {
83:         return -EINVAL;
84:     }
85:     if(!d->data_fn) {
86:         return -EINVAL;
87:     }
88:     if(!(buf = (void *)kmalloc())) {
89:         return -ENOMEM;
90:     }
91:
92:     size = d->data_fn(buf, (i->inode >> 12) & 0xFFFF);
93:     blksize = i->sb->s_blocksize;
94:     if(fd_table->offset > size) {
95:         fd_table->offset = size;
96:     }
97:
98:     total_read = 0;
99:
100:    for(;;) {
101:        count = (fd_table->offset + count > size) ? size - fd_table->off
set : count;
102:        if(!count) {
103:            break;
104:        }
105:
106:        boffset = fd_table->offset % blksize;
107:        bytes = blksize - boffset;
108:        bytes = MIN(bytes, count);
109:        memcpy_b(buffer + total_read, buf + boffset, bytes);
110:        total_read += bytes;
111:        count -= bytes;
112:        boffset += bytes;
113:        boffset %= blksize;
114:        fd_table->offset += bytes;
115:    }
116:
117:    kfree((unsigned int)buf);
118:    return total_read;
119: }
120:
121: int procfs_file_lseek(struct inode *i, __off_t offset)
122: {
123:     return offset;
124: }

```

## fs/procfs/inode.c

Page 1/2

```

1: /*
2:  * fiwix/fs/procfs/inode.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/fs_proc.h>
12: #include <fiwix/statfs.h>
13: #include <fiwix/sleep.h>
14: #include <fiwix/stat.h>
15: #include <fiwix/sched.h>
16: #include <fiwix/mm.h>
17: #include <fiwix/process.h>
18: #include <fiwix/errno.h>
19: #include <fiwix/stdio.h>
20: #include <fiwix/string.h>
21:
22: int procfs_read_inode(struct inode *i)
23: {
24:     int lev;
25:     __mode_t mode;
26:     __nlink_t nlink;
27:     struct procfs_dir_entry *d;
28:
29:     if((i->inode & 0xF0000FFF) == PROC_PID_INO) { /* dynamic PID dir */
30:         mode = S_IFDIR | S_IRUSR | S_IXUSR | S_IRGRP | S_IXGRP | S_IROTH
| S_IXOTH;
31:         nlink = 3;
32:         lev = PROC_PID_LEV;
33:     } else {
34:         if(!(d = get_procfs_by_inode(i))) {
35:             return NULL;
36:         }
37:         mode = d->mode;
38:         nlink = d->nlink;
39:         lev = d->lev;
40:     }
41:
42:     i->i_mode = mode;
43:     i->i_uid = 0;
44:     i->i_size = 0;
45:     i->i_atime = CURRENT_TIME;
46:     i->i_ctime = CURRENT_TIME;
47:     i->i_mtime = CURRENT_TIME;
48:     i->i_gid = 0;
49:     i->i_nlink = nlink;
50:     i->i_blocks = 0;
51:     i->i_flags = 0;
52:     i->locked = 1;
53:     i->dirty = 0;
54:     i->mount_point = NULL;
55:     i->count = 1;
56:     i->u.procfs.i_lev = lev;
57:     switch(i->i_mode & S_IFMT) {
58:         case S_IFDIR:
59:             i->fsop = &procfs_dir_fsop;
60:             break;
61:         case S_IFREG:
62:             i->fsop = &procfs_file_fsop;
63:             break;
64:         case S_IFLNK:
65:             i->fsop = &procfs_symlink_fsop;
66:             break;

```



## fs/procfs/inode.c

Page 2/2

```
67:             default:
68:                 PANIC("invalid inode (%d) mode %08o.\n", i->inode, i->i_
mode);
69:             }
70:             return 0;
71: }
72:
73: int procfs_bmap(struct inode *i, __off_t offset, int mode)
74: {
75:     return i->u.procfs.i_lev;
76: }
77:
78: void procfs_statfs(struct superblock *sb, struct statfs *statfsbuf)
79: {
80:     statfsbuf->f_type = PROC_SUPER_MAGIC;
81:     statfsbuf->f_bsize = sb->s_blocksize;
82:     statfsbuf->f_blocks = 0;
83:     statfsbuf->f_bfree = 0;
84:     statfsbuf->f_bavail = 0;
85:     statfsbuf->f_files = 0;
86:     statfsbuf->f_ffree = 0;
87:     /* statfsbuf->f_fsid = ? */
88:     statfsbuf->f_namelen = NAME_MAX;
89: }
```

**fs/procfs/Makefile**

Page 1/1

```
1: # fiwix/fs/procfs/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = super.o inode.o namei.o dir.o file.o symlink.o tree.o data.o
13:
14: procfs: $(OBJS)
15:     $(LD) $(LDFLAGS) -r $(OBJS) -o procfs.o
16:
17: clean:
18:     rm -f *.o
19:
```

## fs/procfs/namei.c

Page 1/2

```

1: /*
2:  * fiwix/fs/procfs/namei.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/fs.h>
10: #include <fiwix/filesystems.h>
11: #include <fiwix/fs_proc.h>
12: #include <fiwix/process.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/errno.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: int procfs_lookup(const char *name, struct inode *dir, struct inode **i_res)
20: {
21:     int len, lev;
22:     __ino_t inode;
23:     __pid_t pid;
24:     struct proc *p;
25:     struct procfs_dir_entry *pdirent;
26:
27:     pid = inode = 0;
28:     len = strlen(name);
29:     if((dir->inode & 0xF0000000) == PROC_PID_INO) {
30:         pid = (dir->inode >> 12) & 0xFFFF;
31:     }
32:     dir->count++;
33:
34:     lev = bmap(dir, 0, FOR_READING);
35:     pdirent = procfs_array[lev];
36:     while(pdirent->inode && !inode) {
37:         if(len == pdirent->name_len) {
38:             if(!(strcmp(pdirent->name, name))) {
39:                 inode = pdirent->inode;
40:                 if(pid) {
41:                     inode = (PROC_PID_INO + (pid << 12)) + (
inode & 0xFFF);
42:                     if(strcmp(".", name) == 0) {
43:                         inode = dir->inode;
44:                     }
45:                     if(strcmp("../", name) == 0) {
46:                         inode = pdirent->inode;
47:                     }
48:                 }
49:             }
50:         }
51:         if(inode) {
52:             /*
53:              * This prevents a deadlock in iget() when
54:              * trying to lock '.' when 'dir' is the same
55:              * directory (ls -lai <dir>).
56:              */
57:             if(inode == dir->inode) {
58:                 *i_res = dir;
59:                 return 0;
60:             }
61:
62:             if(!(*i_res = iget(dir->sb, inode))) {
63:                 return -EACCES;
64:             }
65:             iput(dir);
66:             return 0;

```

## fs/procfs/namei.c

Page 2/2

```
67:         }
68:         pdirent++;
69:     }
70:
71:     FOR_EACH_PROCESS(p) {
72:         if(p->state) {
73:             if(len == strlen(p->pidstr)) {
74:                 if(!(strcmp(p->pidstr, name))) {
75:                     inode = PROC_PID_INO + (p->pid << 12);
76:                 }
77:             }
78:             if(inode) {
79:                 if(!(*i_res = iget(dir->sb, inode))) {
80:                     return -EACCES;
81:                 }
82:                 iput(dir);
83:                 return 0;
84:             }
85:         }
86:     }
87:     iput(dir);
88:     return -ENOENT;
89: }
```

## fs/procfs/super.c

Page 1/2

```

1: /*
2:  * fiwix/fs/procfs/super.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/filesystems.h>
12: #include <fiwix/fs_proc.h>
13: #include <fiwix/stat.h>
14: #include <fiwix/mm.h>
15: #include <fiwix/sched.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: struct fs_operations procfs_fsop = {
20:     0,
21:     PROC_DEV,
22:
23:     NULL, /* open */
24:     NULL, /* close */
25:     NULL, /* read */
26:     NULL, /* write */
27:     NULL, /* ioctl */
28:     NULL, /* lseek */
29:     NULL, /* readdir */
30:     NULL, /* mmap */
31:     NULL, /* select */
32:
33:     NULL, /* readlink */
34:     NULL, /* followlink */
35:     NULL, /* bmap */
36:     NULL, /* lookup */
37:     NULL, /* rmdir */
38:     NULL, /* link */
39:     NULL, /* unlink */
40:     NULL, /* symlink */
41:     NULL, /* mkdir */
42:     NULL, /* mknod */
43:     NULL, /* truncate */
44:     NULL, /* create */
45:     NULL, /* rename */
46:
47:     NULL, /* read_block */
48:     NULL, /* write_block */
49:
50:     procfs_read_inode,
51:     NULL, /* write_inode */
52:     NULL, /* ialloc */
53:     NULL, /* ifree */
54:     procfs_statfs,
55:     procfs_read_superblock,
56:     NULL, /* remount_fs */
57:     NULL, /* write_superblock */
58:     NULL, /* release_superblock */
59: };
60:
61: int procfs_read_superblock(__dev_t dev, struct superblock *sb)
62: {
63:     superblock_lock(sb);
64:     sb->dev = dev;
65:     sb->fsop = &procfs_fsop;
66:     sb->s_blocksize = PAGE_SIZE;
67:

```

## fs/procfs/super.c

Page 2/2

```
68:         if(!(sb->root = iget(sb, PROC_ROOT_INO))) {
69:             printk("WARNING: %s(): unable to get root inode.\n", __FUNCTION_
_);
70:             superblock_unlock(sb);
71:             return -EINVAL;
72:         }
73:         sb->root->u.procfs.i_lev = 0;
74:
75:         superblock_unlock(sb);
76:         return 0;
77:     }
78:
79: int procfs_init(void)
80: {
81:     return register_filesystem("proc", &procfs_fsop);
82: }
```

## fs/procfs/symlink.c

Page 1/3

```

1: /*
2:  * fiwix/fs/procfs/symlink.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/errno.h>
10: #include <fiwix/buffer.h>
11: #include <fiwix/fs.h>
12: #include <fiwix/filesystems.h>
13: #include <fiwix/fs_proc.h>
14: #include <fiwix/stat.h>
15: #include <fiwix/stdio.h>
16: #include <fiwix/string.h>
17:
18: struct fs_operations procfs_symlink_fsop = {
19:     0,
20:     0,
21:
22:     NULL, /* open */
23:     NULL, /* close */
24:     NULL, /* read */
25:     NULL, /* write */
26:     NULL, /* ioctl */
27:     NULL, /* lseek */
28:     NULL, /* readdir */
29:     NULL, /* mmap */
30:     NULL, /* select */
31:
32:     procfs_readlink,
33:     procfs_followlink,
34:     NULL, /* bmap */
35:     NULL, /* lookup */
36:     NULL, /* rmdir */
37:     NULL, /* link */
38:     NULL, /* unlink */
39:     NULL, /* symlink */
40:     NULL, /* mkdir */
41:     NULL, /* mknod */
42:     NULL, /* truncate */
43:     NULL, /* create */
44:     NULL, /* rename */
45:
46:     NULL, /* read_block */
47:     NULL, /* write_block */
48:
49:     NULL, /* read_inode */
50:     NULL, /* write_inode */
51:     NULL, /* ialloc */
52:     NULL, /* ifree */
53:     NULL, /* statfs */
54:     NULL, /* read_superblock */
55:     NULL, /* remount_fs */
56:     NULL, /* write_superblock */
57:     NULL, /* release_superblock */
58: };
59:
60: int procfs_readlink(struct inode *i, char *buffer, __size_t count)
61: {
62:     __off_t size_read;
63:     struct procfs_dir_entry *d;
64:
65:     if(!(d = get_procfs_by_inode(i))) {
66:         return -EINVAL;
67:     }

```

## fs/procfs/symlink.c

Page 2/3

```

68:
69:     if(!d->data_fn) {
70:         return -EINVAL;
71:     }
72:
73:     size_read = d->data_fn(buffer, (i->inode >> 12) & 0xFFFF);
74:     return size_read;
75: }
76:
77: int procfs_followlink(struct inode *dir, struct inode *i, struct inode **i_res)
78: {
79:     __ino_t errno;
80:     __pid_t pid;
81:     struct proc *p;
82:
83:     if(!i) {
84:         return -ENOENT;
85:     }
86:     if(!(S_ISLNK(i->i_mode))) {
87:         printk("%s(): Oops, inode '%d' is not a symlink (!?)\n", __FUNC
TION_, i->inode);
88:         return 0;
89:     }
90:
91:     p = NULL;
92:     if((pid = (i->inode >> 12) & 0xFFFF)) {
93:         if(!(p = get_proc_by_pid(pid))) {
94:             return -ENOENT;
95:         }
96:     }
97:
98:     /* FIXME!
99:     if(p && p->root) {
100:         printk("(pid %d) p->root->inode = %d (count = %d)\n", p->pid, p-
>root->inode, p->root->count);
101:     }
102:     */
103:
104:     switch(i->inode & 0xF0000FFF) {
105:         case PROC_PID_CWD:
106:             if(!p->pwd) {
107:                 return -ENOENT;
108:             }
109:             *i_res = p->pwd;
110:             p->pwd->count++;
111:             iput(i);
112:             break;
113:         case PROC_PID_EXE:
114:             if(!p->vma || !p->vma->inode) {
115:                 return -ENOENT;
116:             }
117:             *i_res = p->vma->inode;
118:             p->vma->inode->count++;
119:             iput(i);
120:             break;
121:         case PROC_PID_ROOT:
122:             if(!p->root) {
123:                 return -ENOENT;
124:             }
125:             *i_res = p->root;
126:             p->root->count++;
127:             iput(i);
128:             break;
129:         default:
130:             iput(i);
131:             if((errno = parse_namei(current->pidstr, dir, i_res, NUL
L, FOLLOW_LINKS))) {

```



**fs/procfs/symlink.c**

Page 3/3

```
132:                                     return errno;
133:                                     }
134:     }
135:     return 0;
136: }
```

## fs/procfs/tree.c

Page 1/2

```

1: /*
2:  * fiwix/fs/procfs/tree.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/stat.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/fs_proc.h>
12: #include <fiwix/errno.h>
13: #include <fiwix/stdio.h>
14: #include <fiwix/string.h>
15:
16: #define DIR      S_IFDIR | S_IRUSR | S_IXUSR | S_IRGRP | S_IXGRP | S_IROTH | \
17:                S_IXOTH /* dr-xr-xr-x */
18: #define REG      S_IFREG | S_IRUSR | S_IRGRP | S_IROTH /* -r--r--r-- */
19: #define REGUSR   S_IFREG | S_IRUSR /* -r----- */
20: #define LNK      S_IFLNK | S_IRWXU | S_IRWXG | S_IRWXO /* lrwxrwxrwx */
21: #define LNKPID   S_IFLNK | S_IRWXU /* lrwx----- */
22:
23: /*
24:  * WARNING: every time a new entry is added to this array you must also change
25:  * the PROC_ARRAY_ENTRIES value defined in fs_proc.h.
26:  */
27: struct procfs_dir_entry procfs_array[PROC_ARRAY_ENTRIES + 1] = {
28:     { /* [0] / */
29:         { 1, DIR, 2, 0, 1, ".", NULL },
30:         { 2, DIR, 2, 0, 2, "..", NULL },
31:         { 3, DIR, 3, 3, 3, "sys", NULL },
32:         { 4, REG, 1, 0, 7, "cmdline", data_proc_cmdline },
33:         { 5, REG, 1, 0, 7, "cpuinfo", data_proc_cpuinfo },
34:         { 6, REG, 1, 0, 7, "devices", data_proc_devices },
35:         { 7, REG, 1, 0, 3, "dma", data_proc_dma },
36:         { 8, REG, 1, 0, 11, "filesystems", data_proc_filesystems },
37:         { 9, REG, 1, 0, 10, "interrupts", data_proc_interrupts },
38:         { 10, REG, 1, 0, 7, "loadavg", data_proc_loadavg },
39:         { 11, REG, 1, 0, 5, "locks", data_proc_locks },
40:         { 12, REG, 1, 0, 7, "meminfo", data_proc_meminfo },
41:         { 13, REG, 1, 0, 6, "mounts", data_proc_mounts },
42:         { 14, REG, 1, 0, 10, "partitions", data_proc_partitions },
43:         { 15, REG, 1, 0, 3, "rtc", data_proc_rtc },
44:         { 16, LNK, 1, 0, 4, "self", data_proc_self },
45:         { 17, REG, 1, 0, 4, "stat", data_proc_stat },
46:         { 18, REG, 1, 0, 6, "uptime", data_proc_uptime },
47:         { 19, REG, 1, 0, 7, "version", data_proc_fullversion },
48:         { 0, 0, 0, 0, 0, NULL, NULL }
49:     },
50:     { /* [1] /PID/ */
51:         { 1000, DIR, 2, 1, 1, ".", NULL },
52:         { 1, DIR, 2, 0, 2, "..", NULL },
53:         /* { PROC_PID_FD, DIR, 2, 2, 2, "fd", data_proc_pid_fd }, */
54:         { PROC_PID_CMDLINE, REG, 1, 1, 7, "cmdline", data_proc_pid_cmdline }
55:     },
56:     { PROC_PID_CWD, LNKPID, 1, 1, 3, "cwd", data_proc_pid_cwd },
57:     { PROC_PID_ENVIRON, REGUSR, 1, 1, 7, "environ", data_proc_pid_environ }
58: },
59: { PROC_PID_EXE, LNKPID, 1, 1, 3, "exe", data_proc_pid_exe },
60: { PROC_PID_MAPS, REG, 1, 1, 4, "maps", data_proc_pid_maps },
61: { PROC_PID_MOUNTINFO, REG, 1, 1, 9, "mountinfo", data_proc_pid_mountinf
62: },
63: { PROC_PID_ROOT, LNKPID, 1, 1, 4, "root", data_proc_pid_root },
64: { PROC_PID_STAT, REG, 1, 1, 4, "stat", data_proc_pid_stat },
65: { PROC_PID_STATM, REG, 1, 1, 5, "statm", data_proc_pid_statm },
66: { PROC_PID_STATUS, REG, 1, 1, 6, "status", data_proc_pid_status }
67: },

```

## fs/procfs/tree.c

Page 2/2

```

64:         { 0, 0, 0, 0, 0, NULL, NULL }
65:     },
66:
67:     {
68:     },
69:
70:     { /* [3] /sys/ */
71:         { 3,     DIR,  2, 3, 1,  ".",      NULL },
72:         { 1,     DIR,  2, 0, 2,  "..",     NULL },
73:         { 2001,  DIR,  2, 4, 6,  "kernel", NULL },
74:         { 0, 0, 0, 0, 0, NULL, NULL }
75:     },
76:     { /* [4] /sys/kernel/ */
77:         { 2001,  DIR,  2, 4, 1,  ".",      NULL },
78:         { 3,     DIR,  2, 3, 2,  "..",     NULL },
79:         { 3001,  REG,  1, 4, 10, "domainname", data_proc_domainname },
80:         { 3002,  REG,  1, 4, 8,  "file-max",  data_proc_filemax },
81:         { 3003,  REG,  1, 4, 7,  "file-nr",   data_proc_filenr },
82:         { 3004,  REG,  1, 4, 8,  "hostname",  data_proc_hostname },
83:         { 3005,  REG,  1, 4, 9,  "inode-max", data_proc_inodemax },
84:         { 3006,  REG,  1, 4, 8,  "inode-nr",  data_proc_inodenr },
85:         { 3007,  REG,  1, 4, 9,  "osrelease", data_proc_osrelease },
86:         { 3008,  REG,  1, 4, 6,  "ostype",    data_proc_ostype },
87:         { 3009,  REG,  1, 4, 7,  "version",   data_proc_version },
88:         { 0, 0, 0, 0, 0, NULL, NULL }
89:     }
90: };
91:
92: struct procfs_dir_entry * get_procfs_by_inode(struct inode *i)
93: {
94:     __ino_t inode;
95:     int n, lev;
96:     struct procfs_dir_entry *d;
97:
98:     inode = i->inode;
99:     for(lev = 0; procfs_array[lev]; lev++) {
100:         if(lev == PROC_PID_LEV) { /* PID entries */
101:             if((i->inode & 0xF0000000) == PROC_PID_INO) {
102:                 inode = i->inode & 0xF0000FFF;
103:             }
104:         }
105:         d = procfs_array[lev];
106:         for(n = 0; n < PROC_ARRAY_ENTRIES && d->inode; n++) {
107:             if(d->inode == inode) {
108:                 return d;
109:             }
110:             d++;
111:         }
112:     }
113:
114:     return NULL;
115: }

```

## mm/alloc.c

Page 1/1

```
1: /*
2:  * fiwix/mm/alloc.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/mm.h>
9: #include <fiwix/stdio.h>
10: #include <fiwix/string.h>
11:
12: /*
13:  * The implementation of kernel memory allocation is extremely simple, it works
14:  * with a granularity of PAGE_SIZE (4096 bytes). There is indeed a lot of room
15:  * for improvements here.
16:  */
17: unsigned int kmalloc(void)
18: {
19:     struct page *pg;
20:     unsigned int addr;
21:
22:     if((pg = get_free_page())) {
23:         addr = pg->page << PAGE_SHIFT;
24:         return P2V(addr);
25:     }
26:
27:     /* out of memory! */
28:     return 0;
29: }
30:
31: void kfree(unsigned int addr)
32: {
33:     addr = V2P(addr);
34:     release_page(addr >> PAGE_SHIFT);
35: }
```

## mm/bios\_map.c

Page 1/2

```

1: /*
2:  * fiwix/mm/bios_map.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/mm.h>
10: #include <fiwix/bios.h>
11: #include <fiwix/stdio.h>
12: #include <fiwix/string.h>
13:
14: /* check if an specific address is available in the BIOS memory map */
15: int addr_in_bios_map(unsigned int addr)
16: {
17:     int n;
18:     struct bios_mem_map *bmm;
19:
20:     bmm = &bios_mem_map[0];
21:     for(n = 0; n < NR_BIOS_MM_ENT; n++, bmm++) {
22:         if(bmm->to && bmm->type == BIOS_MEM_AVAIL) {
23:             if(addr >= bmm->from && addr < (bmm->to & PAGE_MASK)) {
24:                 return 1;
25:             }
26:         }
27:     }
28:     return 0;        /* not in BIOS map or not available (reserved, ...) */
29: }
30:
31: void bios_map_init(memory_map_t *bmmmap_addr, unsigned long int bmmmap_length)
32: {
33:     unsigned int n;
34:     unsigned int mem_from, mem_to, len_low, len_high;
35:     memory_map_t *bmmmap;
36:     char *bios_mem_type[] = { NULL, "available", "reserved",
37:                             "ACPI Reclaim", "ACPI NVS", "unusable",
38:                             "disabled" };
39:
40:     bmmmap = bmmmap_addr;
41:     if(bmmmap) {
42:         n = 0;
43:
44:         while((unsigned int)bmmmap < (unsigned int)bmmmap_addr + bmmmap_length) {
45:             mem_from = (unsigned int)bmmmap->base_addr_low;
46:             len_low = (unsigned int)bmmmap->length_low;
47:             if(((mem_from >> 16) & 0xFFFF) + ((len_low >> 16) & 0xFFFF) > 0xFFFF) {
48:                 mem_to = ~0;
49:                 len_high = (mem_from >> 16) + (len_low >> 16);
50:                 len_high >>= 16;
51:             } else {
52:                 mem_to = mem_from + len_low;
53:                 len_high = 0;
54:             }
55:             printk("%s      0x%08X%08X-0x%08X%08X %s\n",
56:                  n ? "      " : "memory",
57:                  bmmmap->base_addr_high, bmmmap->base_addr_low,
58:                  len_high, mem_from + len_low,
59:                  bios_mem_type[(int)bmmmap->type]);
60:             /* only memory addresses below 4GB are accepted */
61:             if(!bmmmap->base_addr_high) {
62:                 if(n < NR_BIOS_MM_ENT && len_low) {
63:                     bios_mem_map[n].from = mem_from;
64:                     bios_mem_map[n].to = mem_to;
65:                     bios_mem_map[n].type = (int)bmmmap->type;

```

## mm/bios\_map.c

Page 2/2

```
66:                                     n++;
67:                                     }
68:                                     }
69:                                     bmmap = (memory_map_t *)((unsigned int)bmmap + bmmap->size + sizeof(bmmap->size));
70:                                     }
71:     } else {
72:         printk("WARNING: your BIOS has not provided a memory map.\n");
73:         bios_mem_map[0].from = 0;
74:         bios_mem_map[0].to = _memsize * 1024;
75:         bios_mem_map[0].type = BIOS_MEM_AVAIL;
76:         bios_mem_map[1].from = 0x00100000;
77:         bios_mem_map[1].to = (_extmemsize + 1024) * 1024;
78:         bios_mem_map[1].type = BIOS_MEM_AVAIL;
79:     }
80:     kstat.physical_pages = (_extmemsize + 1024) >> 2;
81:
82:     /*
83:      * This truncates to 1GB since it's the maximum physical memory
84:      * currently supported.
85:      */
86:     if(kstat.physical_pages > (0x40000000 >> PAGE_SHIFT)) {
87:         kstat.physical_pages = (0x40000000 >> PAGE_SHIFT);
88:         printk("WARNING: only up to 1GB of physical memory will be used.
89: \n");
90:     }
```

## mm/fault.c

Page 1/5

```

1: /*
2:  * fiwix/mm/fault.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/sigcontext.h>
10: #include <fiwix/asm.h>
11: #include <fiwix/mm.h>
12: #include <fiwix/process.h>
13: #include <fiwix/traps.h>
14: #include <fiwix/sched.h>
15: #include <fiwix/fs.h>
16: #include <fiwix/mman.h>
17: #include <fiwix/errno.h>
18: #include <fiwix/stdio.h>
19: #include <fiwix/string.h>
20:
21: /* send the SIGSEGV signal to the ofending process */
22: static void send_sigsegv(struct sigcontext *sc)
23: {
24:     dump_registers(14, sc);
25:     printk("Memory map:\n");
26:     show_vma_regions(current);
27:     send_sig(current, SIGSEGV);
28: }
29:
30: static int page_protection_violation(struct vma *vma, unsigned int cr2, struct s
igcontext *sc)
31: {
32:     unsigned int *pgdir;
33:     unsigned int *pgtbl;
34:     unsigned int page, newpage;
35:     unsigned int pde, pte;
36:     struct page *pg;
37:
38:     pde = GET_PGDIR(cr2);
39:     pte = GET_PGTBL(cr2);
40:     pgdir = (unsigned int *)P2V(current->tss.cr3);
41:     pgtbl = (unsigned int *)P2V((pgdir[pde] & PAGE_MASK));
42:     page = (pgtbl[pte] & PAGE_MASK) >> PAGE_SHIFT;
43:
44:     pg = &page_table[page];
45:
46:     /* Copy On Write feature */
47:     if(pg->count > 1) {
48:         /* a page not marked as COW means it's read-only */
49:         if(!(pg->flags & PAGE_COW)) {
50:             printk("Oops!, page %d NOT marked for COW.\n", pg->page)
;
51:             send_sigsegv(sc);
52:             return 0;
53:         }
54:         if(!(newpage = kmalloc())) {
55:             printk("%s(): not enough memory!\n", __FUNCTION__);
56:             return 1;
57:         }
58:         current->rss++;
59:         memcpy_b((void *)newpage, (void *)P2V((page << PAGE_SHIFT)), PAG
E_SIZE);
60:         pgtbl[pte] = V2P(newpage) | PAGE_PRESENT | PAGE_RW | PAGE_USER;
61:         kfree(P2V((page << PAGE_SHIFT)));
62:         current->rss--;
63:         invalidate_tlb();
64:         return 0;

```

## mm/fault.c

Page 2/5

```

65:         } else {
66:             /* last page of Copy On Write procedure */
67:             if(pg->count == 1) {
68:                 /* a page not marked as COW means it's read-only */
69:                 if(!(pg->flags & PAGE_COW)) {
70:                     printk("Oops!, last page %d NOT marked for COW.\n",
n", pg->page);
71:                     send_sigsegv(sc);
72:                     return 0;
73:                 }
74:                 pgtbl[pte] = (page << PAGE_SHIFT) | PAGE_PRESENT | PAGE_
RW | PAGE_USER;
75:                 invalidate_tlb();
76:                 return 0;
77:             }
78:         }
79:         printk("WARNING: %s(): page %d with pg->count = 0!\n", __FUNCTION__, pg-
>page);
80:         return 1;
81:     }
82:
83:     static int page_not_present(struct vma *vma, unsigned int cr2, struct sigcontext
*sc)
84:     {
85:         unsigned int page, file_offset;
86:         struct page *pg;
87:
88:         if(!vma) {
89:             if(cr2 >= (sc->oldesp - 32) && cr2 < KERNEL_BASE_ADDR) {
90:                 if(!(vma = find_vma_region(KERNEL_BASE_ADDR - 1))) {
91:                     printk("WARNING: %s(): process %d doesn't have a
n stack region in vma!\n", __FUNCTION__, current->pid);
92:                     send_sigsegv(sc);
93:                     return 0;
94:                 } else {
95:                     /* assuming stack will never reach heap */
96:                     vma->start = cr2;
97:                     vma->start = vma->start & PAGE_MASK;
98:                 }
99:             }
100:         }
101:
102:         /* if still a non-valid vma is found then kill the process! */
103:         if(!vma || vma->prot == PROT_NONE) {
104:             send_sigsegv(sc);
105:             return 0;
106:         }
107:
108:         /* fill the page with its corresponding file content */
109:         if(vma->inode) {
110:             file_offset = (cr2 & PAGE_MASK) - vma->start + vma->offset;
111:             file_offset &= PAGE_MASK;
112:             pg = NULL;
113:
114:             if(!(vma->prot & PROT_WRITE) || vma->flags & MAP_SHARED) {
115:                 /* check if it's already in cache */
116:                 if((pg = search_page_hash(vma->inode, file_offset))) {
117:                     if(!map_page(current, cr2, (unsigned int)V2P(pg-
>data), vma->prot)) {
118:                         printk("%s(): Oops, map_page() returned
0!\n", __FUNCTION__);
119:                         return 1;
120:                     }
121:                     page = (unsigned int)pg->data;
122:                 }
123:             }
124:             if(!pg) {

```



## mm/fault.c

Page 3/5

```

125:             if(!page = map_page(current, cr2, 0, vma->prot)) {
126:                 printk("%s(): Oops, map_page() returned 0!\n", _
_FUNCTION__);
127:                 return 1;
128:             }
129:             pg = &page_table[V2P(page) >> PAGE_SHIFT];
130:             if(bread_page(pg, vma->inode, file_offset, vma->prot, vm
a->flags)) {
131:                 unmap_page(cr2);
132:                 return 1;
133:             }
134:             current->usage.ru_majflt++;
135:         }
136:     } else {
137:         current->usage.ru_minflt++;
138:         page = 0;
139:     }
140:
141:     if(vma->flags & ZERO_PAGE) {
142:         if(!page) {
143:             if(!page = map_page(current, cr2, 0, vma->prot)) {
144:                 printk("%s(): Oops, map_page() returned 0!\n", _
_FUNCTION__);
145:                 return 1;
146:             }
147:         }
148:         memset_b((void *) (page & PAGE_MASK), NULL, PAGE_SIZE);
149:     }
150:
151:     return 0;
152: }
153:
154: /*
155:  * Exception 0xE: Page Fault
156:  *
157:  *
158:  *
159:  *
160:  *
161:  *
162:  *
163:  *
164:  *
165:  *
166:  *
167:  *
168:  *
169:  * U1 - vma + user + PV + read
170:  *     (vma page in user-mode, page-violation during read)
171:  *     U1.1) if flags match          -> Demand paging
172:  *     U1.2) if flags don't match    -> SIGSEV
173:  *
174:  * U2 - vma + user + PV + write
175:  *     (vma page in user-mode, page-violation during write)
176:  *     U2.1) if flags match          -> Copy-On-Write
177:  *     U2.2) if flags don't match    -> SIGSEGV
178:  *
179:  * U3 - vma + user + PF + (read | write)    -> Demand paging
180:  *     (vma page in user-mode, page-fault during read or write)
181:  *
182:  * K1 - vma + kernel + PV + (read | write)  -> PANIC
183:  *     (vma page in kernel-mode, page-violation during read or write)
184:  * K2 - vma + kernel + PF + (read | write)  -> Demand paging (mmap)
185:  *     (vma page in kernel-mode, page-fault during read or write)
186:  *
187:  * -----
188:  *

```

## mm/fault.c

Page 4/5

```

189:  * U1 - !vma + user + PV + (read | write)      -> SIGSEGV
190:  *      (!vma page in user-mode, page-violation during read or write)
191:  * U2 - !vma + user + PF + (read | write)      -> STACK grows
192:  *      (!vma page in user-mode, page-fault during read or write)
193:  *
194:  * K1 - !vma + kernel + (PV | PF) + (read | write)  -> PANIC
195:  *      (!vma page in kernel-mode, page-fault or page-violation during read
196:  *      or write)
197:  */
198: void do_page_fault(unsigned int trap, struct sigcontext *sc)
199: {
200:     unsigned int cr2;
201:     struct vma *vma;
202:
203:     GET_CR2(cr2);
204:     if((vma = find_vma_region(cr2))) {
205:
206:         /* in user mode */
207:         if(sc->err & PFAULT_U) {
208:             if(sc->err & PFAULT_V) {          /* violation */
209:                 if(sc->err & PFAULT_W) {
210:                     if((page_protection_violation(vma, cr2,
sc))) {
211:                         send_sig(current, SIGKILL);
212:                     }
213:                     return;
214:                 }
215:                 send_sigsegv(sc);
216:             } else {                          /* page not present */
217:                 if((page_not_present(vma, cr2, sc))) {
218:                     send_sig(current, SIGKILL);
219:                 }
220:             }
221:             return;
222:
223:             /* in kernel mode */
224:         } else {
225:             /*
226:              * WP bit marks the order: first check if the page is
227:              * present, then check for protection violation.
228:              */
229:             if(!(sc->err & PFAULT_V)) {      /* page not present */
230:                 if((page_not_present(vma, cr2, sc))) {
231:                     send_sig(current, SIGKILL);
232:                     printk("%s(): kernel was unable to read
a page of process '%s' (pid %d).\n", __FUNCTION__, current->argv0, current->pid);
233:                 }
234:                 return;
235:             }
236:             if(sc->err & PFAULT_W) {        /* copy-on-write? */
237:                 if((page_protection_violation(vma, cr2, sc))) {
238:                     send_sig(current, SIGKILL);
239:                     printk("%s(): kernel was unable to write
a page of process '%s' (pid %d).\n", __FUNCTION__, current->argv0, current->pid);
240:                 }
241:                 return;
242:             }
243:         }
244:     } else {
245:         /* in user mode */
246:         if(sc->err & PFAULT_U) {
247:             if(sc->err & PFAULT_V) {          /* violation */
248:                 send_sigsegv(sc);
249:             } else {                          /* stack? */
250:                 if((page_not_present(vma, cr2, sc))) {
251:                     send_sig(current, SIGKILL);
252:                 }

```

**mm/fault.c**

Page 5/5

```
253:                }
254:                return;
255:            }
256:        }
257:
258:        dump_registers(trap, sc);
259:        show_vma_regions(current);
260:        PANIC("\n");
261: }
```

**mm/Makefile**

Page 1/1

```
1: # fiwix/mm/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = bios_map.o memory.o page.o alloc.o fault.o mmap.o swapper.o
13:
14: mm:     $(OBJS)
15:         $(LD) $(LDFLAGS) -r $(OBJS) -o mm.o
16:
17: clean:
18:         rm -f *.o
19:
```

## mm/memory.c

Page 1/8

```

1: /*
2:  * fiwix/mm/memory.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/asm.h>
10: #include <fiwix/mm.h>
11: #include <fiwix/mman.h>
12: #include <fiwix/bios.h>
13: #include <fiwix/ramdisk.h>
14: #include <fiwix/process.h>
15: #include <fiwix/buffer.h>
16: #include <fiwix/fs.h>
17: #include <fiwix/filesystems.h>
18: #include <fiwix/stdio.h>
19: #include <fiwix/string.h>
20:
21: #define KERNEL_TEXT_SIZE      ((int)_etext - (KERNEL_BASE_ADDR + KERNEL_ENTRY_
ADDR))
22: #define KERNEL_DATA_SIZE      ((int)_edata - (int)_etext)
23: #define KERNEL_BSS_SIZE       ((int)_end - (int)_edata)
24:
25: #define PGDIR_4MB_ADDR        0x90000
26:
27: unsigned int *kpage_dir;
28: unsigned int *kpage_table;
29:
30: unsigned int _last_data_addr;
31:
32: unsigned int proc_table_size = 0;
33: struct proc *proc_table;
34:
35: unsigned int buffer_table_size = 0;
36: unsigned int buffer_hash_table_size = 0;
37: struct buffer *buffer_table;
38: struct buffer **buffer_hash_table;
39:
40: unsigned int inode_table_size = 0;
41: unsigned int inode_hash_table_size = 0;
42: struct inode *inode_table;
43: struct inode **inode_hash_table;
44:
45: unsigned int fd_table_size = 0;
46: struct fd *fd_table;
47:
48: unsigned int mount_table_size = 0;
49: struct mount *mount_table;
50:
51: struct ramdisk ramdisk_table[RAMDISK_MINORS];
52:
53: unsigned int page_table_size = 0;
54: unsigned int page_hash_table_size = 0;
55: struct page *page_table;
56: struct page **page_hash_table;
57:
58: static void map_kaddr(unsigned int from, unsigned int to, int flags)
59: {
60:     unsigned int n;
61:     unsigned int *pgtbl;
62:     unsigned int pde, pte;
63:
64:     for(n = from >> PAGE_SHIFT; n < (to >> PAGE_SHIFT); n++) {
65:         pde = GET_PGDIR(n << PAGE_SHIFT);
66:         pte = GET_PGTBL(n << PAGE_SHIFT);

```

## mm/memory.c

Page 2/8

```

67:             if(!(kpage_dir[pde] & ~PAGE_MASK)) {
68:                 unsigned int addr;
69:                 addr = _last_data_addr;
70:                 _last_data_addr += PAGE_SIZE;
71:                 kpage_dir[pde] = addr | flags;
72:                 memset_b((void *)addr, NULL, PAGE_SIZE);
73:             }
74:             pgtbl = (unsigned int *) (kpage_dir[pde] & PAGE_MASK);
75:             pgtbl[pte] = (kpage_table[(pde * 1024) + pte] & PAGE_MASK) | fla
gs;
76:         }
77:     }
78:
79: void bss_init(void)
80: {
81:     memset_b((void *) ((int)_edata), NULL, KERNEL_BSS_SIZE);
82: }
83:
84: /*
85:  * This function creates a minimal Page Directory covering only the first 4MB
86:  * of physical memory. Just enough to boot the kernel.
87:  * (it returns the address to be used by the CR3 register)
88:  */
89: unsigned int setup_minmem(void)
90: {
91:     int n;
92:     unsigned int addr;
93:     short int pd, mb4;
94:
95:     mb4 = 1;          /* 4MB units */
96:     addr = KERNEL_BASE_ADDR + PGDIR_4MB_ADDR;
97:
98:     kpage_dir = (unsigned int *)addr;
99:     memset_b(kpage_dir, NULL, PAGE_SIZE);
100:
101:     addr += PAGE_SIZE;
102:     kpage_table = (unsigned int *)addr;
103:     memset_b(kpage_table, NULL, PAGE_SIZE * mb4);
104:
105:     for(n = 0; n < (1024 * mb4); n++) {
106:         kpage_table[n] = (n << PAGE_SHIFT) | PAGE_PRESENT | PAGE_RW;
107:         if(!(n % 1024)) {
108:             pd = n / 1024;
109:             kpage_dir[pd] = (unsigned int) (addr + (PAGE_SIZE * pd) +
0x40000000) | PAGE_PRESENT | PAGE_RW;
110:             kpage_dir[GET_PGDIR(KERNEL_BASE_ADDR) + pd] = (unsigned
int) (addr + (PAGE_SIZE * pd) + 0x40000000) | PAGE_PRESENT | PAGE_RW;
111:         }
112:     }
113:     return (unsigned int)kpage_dir + 0x40000000;
114: }
115:
116: /* returns the mapped address of a virtual address */
117: unsigned int get_mapped_addr(struct proc *p, unsigned int addr)
118: {
119:     unsigned int *pgdir, *pgtbl;
120:     unsigned int pde, pte;
121:
122:     pgdir = (unsigned int *)P2V(p->tss.cr3);
123:     pde = GET_PGDIR(addr);
124:     pte = GET_PGTBL(addr);
125:     pgtbl = (unsigned int *)P2V((pgdir[pde] & PAGE_MASK));
126:     return pgtbl[pte];
127: }
128:
129: int clone_pages(struct proc *child)
130: {

```

## mm/memory.c

Page 3/8

```

131:     unsigned int *src_pgdirt, *dst_pgdirt;
132:     unsigned int *src_pgtbl, *dst_pgtbl;
133:     unsigned int pde, pte;
134:     unsigned int p_page, c_page;
135:     unsigned int n, n2, pages;
136:     struct page *pg;
137:     struct vma *vma;
138:
139:     src_pgdirt = (unsigned int *)P2V(current->tss.cr3);
140:     dst_pgdirt = (unsigned int *)P2V(child->tss.cr3);
141:     vma = current->vma;
142:
143:     for(n = 0, pages = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
144:         for(n2 = vma->start; n2 < vma->end; n2 += PAGE_SIZE) {
145:             if(vma->flags & MAP_SHARED) {
146:                 continue;
147:             }
148:             pde = GET_PGDIR(n2);
149:             pte = GET_PGTBL(n2);
150:             if(src_pgdirt[pde] & PAGE_PRESENT) {
151:                 src_pgtbl = (unsigned int *)P2V((src_pgdirt[pde]
& PAGE_MASK));
152:                 if(!(dst_pgdirt[pde] & PAGE_PRESENT)) {
153:                     if(!(c_page = kmalloc())) {
154:                         printk("%s(): returning 0!\n", __
_FUNCTION__);
155:                         return 0;
156:                     }
157:                     current->rss++;
158:                     pages++;
159:                     dst_pgdirt[pde] = V2P(c_page) | PAGE_PRES
ENT | PAGE_RW | PAGE_USER;
160:                     memset_b((void *)c_page, NULL, PAGE_SIZE
);
161:                 }
162:                 dst_pgtbl = (unsigned int *)P2V((dst_pgdirt[pde]
& PAGE_MASK));
163:                 if(src_pgtbl[pte] & PAGE_PRESENT) {
164:                     p_page = src_pgtbl[pte] >> PAGE_SHIFT;
165:                     pg = &page_table[p_page];
166:                     if(pg->flags & PAGE_RESERVED) {
167:                         continue;
168:                     }
169:                     src_pgtbl[pte] &= ~PAGE_RW;
170:                     /* mark write-only pages as COW*/
171:                     if(vma->prot & PROT_WRITE) {
172:                         pg->flags |= PAGE_COW;
173:                     }
174:                     dst_pgtbl[pte] = src_pgtbl[pte];
175:                     if(!valid_page((dst_pgtbl[pte] & PAGE_MA
SK) >> PAGE_SHIFT)) {
176:                         PANIC("%s: missing page %d durin
g copy-on-write process.\n", __FUNCTION__, (dst_pgtbl[pte] & PAGE_MASK) >> PAGE_SHIFT);
177:                     }
178:                     pg = &page_table[(dst_pgtbl[pte] & PAGE_
MASK) >> PAGE_SHIFT];
179:                     pg->count++;
180:                 }
181:             }
182:         }
183:     }
184:     return pages;
185: }
186:
187: int free_page_tables(struct proc *p)
188: {
189:     unsigned int *pgdirt;

```

## mm/memory.c

Page 4/8

```

190:         unsigned int n, count;
191:
192:         pgdir = (unsigned int *)P2V(p->tss.cr3);
193:         for(n = 0, count = 0; n < PD_ENTRIES; n++) {
194:             if((pgdir[n] & (PAGE_PRESENT | PAGE_RW | PAGE_USER)) == (PAGE_PR
ESENT | PAGE_RW | PAGE_USER)) {
195:                 kfree(P2V(pgdir[n]) & PAGE_MASK);
196:                 pgdir[n] = NULL;
197:                 count++;
198:             }
199:         }
200:         return count;
201:     }
202:
203: unsigned int map_page(struct proc *p, unsigned int vaddr, unsigned int addr, uns
igned int prot)
204: {
205:     unsigned int *pgdir, *pgtbl;
206:     unsigned int vpage;
207:     int pde, pte;
208:
209:     pgdir = (unsigned int *)P2V(p->tss.cr3);
210:     pde = GET_PGDIR(vaddr);
211:     pte = GET_PGTBL(vaddr);
212:
213:     if(!(pgdir[pde] & PAGE_PRESENT)) {          /* allocating page table */
214:         if(!(vpage = kmalloc())) {
215:             return 0;
216:         }
217:         p->rss++;
218:         pgdir[pde] = V2P(vpage) | PAGE_PRESENT | PAGE_RW | PAGE_USER;
219:         memset_b((void *)vpage, NULL, PAGE_SIZE);
220:     }
221:     pgtbl = (unsigned int *)P2V((pgdir[pde] & PAGE_MASK));
222:     if(!(pgtbl[pte] & PAGE_PRESENT)) {          /* allocating page */
223:         if(!addr) {
224:             if(!(addr = kmalloc())) {
225:                 return 0;
226:             }
227:             addr = V2P(addr);
228:             p->rss++;
229:         }
230:     }
231:     pgtbl[pte] = addr | PAGE_PRESENT | PAGE_USER;
232:     if(prot & PROT_WRITE) {
233:         pgtbl[pte] |= PAGE_RW;
234:     }
235:     return P2V(addr);
236: }
237:
238: int unmap_page(unsigned int vaddr)
239: {
240:     unsigned int *pgdir, *pgtbl;
241:     unsigned int addr;
242:     int pde, pte;
243:
244:     pgdir = (unsigned int *)P2V(current->tss.cr3);
245:     pde = GET_PGDIR(vaddr);
246:     pte = GET_PGTBL(vaddr);
247:     if(!(pgdir[pde] & PAGE_PRESENT)) {
248:         printk("WARNING: %s(): trying to unmap an unallocated pde '0x%08
x'\n", __FUNCTION__, vaddr);
249:         return 1;
250:     }
251:
252:     pgtbl = (unsigned int *)P2V((pgdir[pde] & PAGE_MASK));
253:     if(!(pgtbl[pte] & PAGE_PRESENT)) {

```



## mm/memory.c

Page 5/8

```

254:             printk("WARNING: %s(): trying to unmap an unallocated page '0x%0
8x'\n", __FUNCTION__, vaddr);
255:             return 1;
256:         }
257:
258:         addr = pgtbl[pte] & PAGE_MASK;
259:         pgtbl[pte] = NULL;
260:         kfree(P2V(addr));
261:         current->rss--;
262:         return 0;
263:     }
264:
265: void mem_init(void)
266: {
267:     unsigned int pages, sizek;
268:     unsigned int n;
269:     unsigned int physical_page_tables;
270:     unsigned int physical_memory;
271:
272:     physical_page_tables = (kstat.physical_pages / 1024) + ((kstat.physical_
pages % 1024) ? 1 : 0);
273:     physical_memory = (kstat.physical_pages << PAGE_SHIFT); /* in bytes */
274:
275:     /* Page Directory will be aligned to the next page */
276:     _last_data_addr = PAGE_ALIGN(_last_data_addr);
277:     kpage_dir = (unsigned int *)_last_data_addr;
278:     memset_b(kpage_dir, NULL, PAGE_SIZE);
279:     _last_data_addr += PAGE_SIZE;
280:
281:     /* Page Tables */
282:     kpage_table = (unsigned int *)_last_data_addr;
283:     memset_b(kpage_table, NULL, physical_page_tables * PAGE_SIZE);
284:     _last_data_addr += physical_page_tables * PAGE_SIZE;
285:
286:     /* Page Directory and Page Tables initialization */
287:     for(n = 0; n < kstat.physical_pages; n++) {
288:         kpage_table[n] = (n << PAGE_SHIFT) | PAGE_PRESENT | PAGE_RW;
289:     }
290:     for(n = 0; n < physical_page_tables; n++) {
291:         kpage_dir[GET_PGDIR(KERNEL_BASE_ADDR) + n] = (unsigned int)&kpag
e_table[n * 1024] | PAGE_PRESENT | PAGE_RW;
292:     }
293:
294:     map_kaddr(0xA0000, KERNEL_ENTRY_ADDR, PAGE_PRESENT | PAGE_RW);
295:     map_kaddr(KERNEL_ENTRY_ADDR, _last_data_addr, PAGE_PRESENT | PAGE_RW);
296:     /* printk("_last_data_addr = 0x%08x-0x%08x (kernel)\n", KERNEL_ENTRY_ADDR,
_last_data_addr); */
297:     activate_kpage_dir();
298:
299:     /* since Page Directory is now activated we can use virtual addresses */
300:     _last_data_addr = P2V(_last_data_addr);
301:
302:
303:     /* reserve memory space for proc_table[NR_PROCS] */
304:     proc_table_size = PAGE_ALIGN(sizeof(struct proc) * NR_PROCS);
305:     if(!addr_in_bios_map(V2P(_last_data_addr) + proc_table_size)) {
306:         PANIC("Not enough memory for proc_table.\n");
307:     }
308:     /* printk("_last_data_addr = 0x%08x-0x%08x (proc_table)\n", _last_data_addr
, _last_data_addr + proc_table_size); */
309:     proc_table = (struct proc *)_last_data_addr;
310:     _last_data_addr += proc_table_size;
311:
312:
313:     /* reserve memory space for buffer_table */
314:     buffer_table_size = (kstat.physical_pages * BUFFER_PERCENTAGE) / 100;
315:     buffer_table_size *= sizeof(struct buffer);

```

## mm/memory.c

Page 6/8

```

316:         pages = buffer_table_size >> PAGE_SHIFT;
317:         buffer_table_size = !pages ? 4096 : pages << PAGE_SHIFT;
318:  /*
319:     printk("_last_data_addr = 0x%08x-0x%08x (buffer_table)\n", _last_data_ad
dr, _last_data_addr + buffer_table_size); */
320:         if(!addr_in_bios_map(V2P(_last_data_addr) + buffer_table_size)) {
321:             PANIC("Not enough memory for buffer_table.\n");
322:         }
323:         buffer_table = (struct buffer *)_last_data_addr;
324:         _last_data_addr += buffer_table_size;
325:
326:         /* reserve memory space for buffer_hash_table */
327:         n = (buffer_table_size / sizeof(struct buffer) * BUFFER_HASH_PERCENTAGE)
/ 100;
328:         n = MAX(n, 10); /* 10 buffer hashes as minimum */
329:         /* buffer_hash_table is an array of pointers */
330:         pages = ((n * sizeof(unsigned int)) / PAGE_SIZE) + 1;
331:         buffer_hash_table_size = pages << PAGE_SHIFT;
332:  /*
333:     printk("_last_data_addr = 0x%08x-0x%08x (buffer_hash_table)\n", _last_da
ta_addr, _last_data_addr + buffer_hash_table_size); */
334:         if(!addr_in_bios_map(V2P(_last_data_addr) + buffer_hash_table_size)) {
335:             PANIC("Not enough memory for buffer_hash_table.\n");
336:         }
337:         buffer_hash_table = (struct buffer **)_last_data_addr;
338:         _last_data_addr += buffer_hash_table_size;
339:
340:         /* reserve memory space for inode_table */
341:         sizek = physical_memory / 1024; /* this helps to avoid overflow */
342:         inode_table_size = (sizek * INODE_PERCENTAGE) / 100;
343:         inode_table_size *= 1024;
344:         pages = inode_table_size >> PAGE_SHIFT;
345:         inode_table_size = pages << PAGE_SHIFT;
346:  /*
347:     printk("_last_data_addr = 0x%08x-0x%08x (inode_table)\n", _last_data_add
r, _last_data_addr + inode_table_size); */
348:         if(!addr_in_bios_map(V2P(_last_data_addr) + inode_table_size)) {
349:             PANIC("Not enough memory for inode_table.\n");
350:         }
351:         inode_table = (struct inode *)_last_data_addr;
352:         _last_data_addr += inode_table_size;
353:
354:         /* reserve memory space for inode_hash_table */
355:         n = ((inode_table_size / sizeof(struct inode)) * INODE_HASH_PERCENTAGE)
/ 100;
356:         n = MAX(n, 10); /* 10 inodes hashes as minimum */
357:         /* inode_hash_table is an array of pointers */
358:         pages = ((n * sizeof(unsigned int)) / PAGE_SIZE) + 1;
359:         inode_hash_table_size = pages << PAGE_SHIFT;
360:  /*
361:     printk("_last_data_addr = 0x%08x-0x%08x (inode_hash_table)\n", _last_dat
a_addr, _last_data_addr + inode_hash_table_size); */
362:         if(!addr_in_bios_map(V2P(_last_data_addr) + inode_hash_table_size)) {
363:             PANIC("Not enough memory for inode_hash_table.\n");
364:         }
365:         inode_hash_table = (struct inode **)_last_data_addr;
366:         _last_data_addr += inode_hash_table_size;
367:
368:         /* reserve memory space for fd_table[NR_OPENS] */
369:         fd_table_size = PAGE_ALIGN(sizeof(struct fd) * NR_OPENS);
370:  /*
371:     printk("_last_data_addr = 0x%08x-0x%08x (fd_table)\n", _last_data_addr,
_last_data_addr + fd_table_size); */
372:         if(!addr_in_bios_map(V2P(_last_data_addr) + fd_table_size)) {
373:             PANIC("Not enough memory for fd_table.\n");
374:         }
375:         fd_table = (struct fd *)_last_data_addr;
376:         _last_data_addr += fd_table_size;

```

## mm/memory.c

Page 7/8

```

376:
377:
378:         /* reserve memory space for mount_table[NR_MOUNT_POINTS] */
379:         mount_table_size = PAGE_ALIGN(sizeof(struct mount) * NR_MOUNT_POINTS);
380:         /*      printk("_last_data_addr = 0x%08x-0x%08x (mount_table)\n", _last_data_addr
r, _last_data_addr + mount_table_size); */
381:         if(!addr_in_bios_map(V2P(_last_data_addr) + mount_table_size)) {
382:             PANIC("Not enough memory for mount_table.\n");
383:         }
384:         mount_table = (struct mount *)_last_data_addr;
385:         _last_data_addr += mount_table_size;
386:
387:
388:         /* reserve memory space for RAMdisk(s) */
389:         if(!_noramdisk) {
390:             if(!_ramdisksize) {
391:                 _ramdisksize = RAMDISK_SIZE;
392:             }
393:             if(!addr_in_bios_map(V2P(_last_data_addr) + (_ramdisksize * 1024
))) {
394:                 printk("WARNING: RAMdisk device disabled (not enough phy
sical memory).\n");
395:                 _noramdisk = 1;
396:             } else {
397:                 /*
398:                  * If the 'initrd=' parameter was supplied, then the fir
st
399:                  * ramdisk was already assigned to the initial ramdisk i
mage.
400:                  */
401:                 if(ramdisk_table[0].addr) {
402:                     n = 1;
403:                     _last_data_addr += _ramdisksize * 1024;
404:                 } else {
405:                     n = 0;
406:                 }
407:                 for(; n < RAMDISK_MINORS; n++) {
408:                     /*      printk("_last_data_addr = 0x%08x-0x%08x (/dev/ra
m%d)\n", _last_data_addr, _last_data_addr + (_ramdisksize * 1024), n); */
409:                     ramdisk_table[n].addr = (char *)_last_data_addr;
410:                     _last_data_addr += _ramdisksize * 1024;
411:                 }
412:             }
413:         }
414:
415:
416:         /* the last one must be the page_table structure */
417:         page_hash_table_size = 1 * PAGE_SIZE; /* only 1 page size */
418:         if(!addr_in_bios_map(V2P(_last_data_addr) + page_hash_table_size)) {
419:             PANIC("Not enough memory for page_hash_table.\n");
420:         }
421:         page_hash_table = (struct page **)_last_data_addr;
422:         /*      printk("_last_data_addr = 0x%08x-0x%08x (page_hash_table)\n", _last_data
_addr, _last_data_addr + page_hash_table_size); */
423:         _last_data_addr += page_hash_table_size;
424:
425:         page_table_size = PAGE_ALIGN(kstat.physical_pages * sizeof(struct page))
;
426:         if(!addr_in_bios_map(V2P(_last_data_addr) + page_table_size)) {
427:             PANIC("Not enough memory for page_table.\n");
428:         }
429:         page_table = (struct page *)_last_data_addr;
430:         /*      printk("page_table_size = %d\n", page_table_size); */
431:         /*      printk("_last_data_addr = 0x%08x-0x%08x (page_table)\n", _last_data_addr
, _last_data_addr + page_table_size); */
432:         _last_data_addr += page_table_size;
433:

```

## mm/memory.c

Page 8/8

```
434:         page_init(kstat.physical_pages);
435:     }
436:
437: void mem_stats(void)
438: {
439:     printk("\n");
440:     printk("memory: total/available=%dKB/%dKB, kernel=%dKB, reserved=%dKB\n"
, kstat.physical_pages << 2, kstat.total_mem_pages << 2, kstat.kernel_reserved, kstat.p
hysical_reserved);
441:     printk("kernel: text=%dKB, data=%dKB, bss=%dKB, i/o buffers=%d (%dKB)\n"
, KERNEL_TEXT_SIZE / 1024, KERNEL_DATA_SIZE / 1024, KERNEL_BSS_SIZE / 1024, buffer_tabl
e_size / sizeof(struct buffer), (buffer_table_size + buffer_hash_table_size) / 1024);
442:     printk("\tinodes=%d (%dKB)\n\n", inode_table_size / sizeof(struct inode)
, (inode_table_size + inode_hash_table_size) / 1024);
443: }
```

## mm/mmap.c

Page 1/8

```

1: /*
2:  * fiwix/mm/mmap.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/asm.h>
9: #include <fiwix/mm.h>
10: #include <fiwix/fs.h>
11: #include <fiwix/fcntl.h>
12: #include <fiwix/stat.h>
13: #include <fiwix/process.h>
14: #include <fiwix/mman.h>
15: #include <fiwix/errno.h>
16: #include <fiwix/stdio.h>
17: #include <fiwix/string.h>
18:
19: void show_vma_regions(struct proc *p)
20: {
21:     __ino_t inode;
22:     int major, minor;
23:     char *section;
24:     char r, w, x, f;
25:     struct vma *vma;
26:     unsigned int n;
27:     int count;
28:
29:     vma = p->vma;
30:     printk("num  address range          flag offset      dev  inode      mod
section cnt\n");
31:     printk("-----\n");
32:     for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
33:         r = vma->prot & PROT_READ ? 'r' : '-';
34:         w = vma->prot & PROT_WRITE ? 'w' : '-';
35:         x = vma->prot & PROT_EXEC ? 'x' : '-';
36:         if(vma->flags & MAP_SHARED) {
37:             f = 's';
38:         } else if(vma->flags & MAP_PRIVATE) {
39:             f = 'p';
40:         } else {
41:             f = '-';
42:         }
43:         switch(vma->s_type) {
44:             case P_TEXT:    section = "text ";
45:                             break;
46:             case P_DATA:    section = "data ";
47:                             break;
48:             case P_BSS:     section = "bss  ";
49:                             break;
50:             case P_HEAP:    section = "heap ";
51:                             break;
52:             case P_STACK:   section = "stack";
53:                             break;
54:             case P_MMAP:    section = "mmap ";
55:                             break;
56:             default:
57:                 section = NULL;
58:                 break;
59:         }
60:         inode = major = minor = count = 0;
61:         if(vma->inode) {
62:             inode = vma->inode->inode;
63:             major = MAJOR(vma->inode->dev);
64:             minor = MINOR(vma->inode->dev);
65:             count = vma->inode->count;

```

## mm/mmap.c

Page 2/8

```

66:         }
67:         printk("[%02d] 0x%08x-0x%08x %c%c%c%c 0x%08x %02d:%02d %- 10u <%=
d> [%s] (%d)\n", n, vma->start, vma->end, r, w, x, f, vma->offset, major, minor, inode
, vma->o_mode, section, count);
68:     }
69:     if(!n) {
70:         printk("[no vma regions]\n");
71:     }
72: }
73:
74: static struct vma * get_new_vma_region(void)
75: {
76:     unsigned int n;
77:     struct vma *vma;
78:
79:     vma = current->vma;
80:
81:     for(n = 0; n < VMA_REGIONS; n++, vma++) {
82:         if(!vma->start && !vma->end) {
83:             return vma;
84:         }
85:     }
86:     return NULL;
87: }
88:
89: /*
90:  * This sorts regions (in ascending order), merging equal regions and keeping
91:  * the unused ones at the end of the array.
92:  */
93: static void sort_vma(void)
94: {
95:     unsigned int n, n2, needs_sort;
96:     struct vma *vma, tmp;
97:
98:     vma = current->vma;
99:
100:    do {
101:        needs_sort = 0;
102:        for(n = 0, n2 = 1; n2 < VMA_REGIONS; n++, n2++) {
103:            if(vma[n].end && vma[n2].start) {
104:                if((vma[n].end == vma[n2].start) &&
105:                    (vma[n].prot == vma[n2].prot) &&
106:                    (vma[n].flags == vma[n2].flags) &&
107:                    (vma[n].offset == vma[n2].offset) &&
108:                    (vma[n].s_type == vma[n2].s_type) &&
109:                    (vma[n].inode == vma[n2].inode)) {
110:                    vma[n].end = vma[n2].end;
111:                    memset_b(&vma[n2], NULL, sizeof(struct v
ma));
112:                    needs_sort++;
113:                }
114:            }
115:            if((vma[n2].start && (vma[n].start > vma[n2].start)) ||
(!vma[n].start && vma[n2].start)) {
116:                memcpy_b(&tmp, &vma[n], sizeof(struct vma));
117:                memcpy_b(&vma[n], &vma[n2], sizeof(struct vma));
118:                memcpy_b(&vma[n2], &tmp, sizeof(struct vma));
119:                needs_sort++;
120:            }
121:        }
122:    } while(needs_sort);
123: }
124:
125: /*
126:  * This function removes all redundant entries.
127:  *
128:  * for example, if for any reason the map looks like this:

```

## mm/mmap.c

Page 3/8

```

129:  * [01] 0x0808e984-0x08092000 rw-p 0x00000000 0
130:  * [02] 0x0808f000-0x0808ffff rw-p 0x000c0000 4066
131:  *
132:  * this function converts it to this:
133:  * [01] 0x0808e984-0x0808f000 rw-p 0x00000000 0
134:  * [02] 0x0808f000-0x0808ffff rw-p 0x000c0000 4066
135:  * [03] 0x08090000-0x08092000 rw-p 0x00000000 0
136:  */
137: static int optimize_vma(void)
138: {
139:     unsigned int n, needs_sort;
140:     struct vma *vma, *prev, *new;
141:
142:     for(;;) {
143:         needs_sort = 0;
144:         prev = new = NULL;
145:         vma = current->vma;
146:         for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
147:             if(!prev) {
148:                 prev = vma;
149:                 continue;
150:             }
151:             if((vma->start < prev->end)) {
152:                 if(!(new = get_new_vma_region())) {
153:                     printk("WARNING: %s(): unable to get a f
ree vma region.\n", __FUNCTION__);
154:                     return -ENOMEM;
155:                 }
156:                 new->start = vma->end;
157:                 new->end = prev->end;
158:                 new->prot = prev->prot;
159:                 new->flags = prev->flags;
160:                 new->offset = prev->offset;
161:                 new->s_type = prev->s_type;
162:                 new->inode = prev->inode;
163:                 new->o_mode = prev->o_mode;
164:                 prev->end = vma->start;
165:                 needs_sort++;
166:                 if(prev->start == prev->end) {
167:                     memset_b(prev, NULL, sizeof(struct vma))
;
168:                 }
169:                 if(new->start == new->end) {
170:                     memset_b(new, NULL, sizeof(struct vma));
171:                 }
172:                 break;
173:             }
174:             prev = vma;
175:         }
176:         if(!needs_sort) {
177:             break;
178:         }
179:         sort_vma();
180:     }
181:
182:     return 0;
183: }
184:
185: /* return the first free address that matches with the size of length */
186: static unsigned int get_unmapped_vma_region(unsigned int length)
187: {
188:     unsigned int n, addr;
189:     struct vma *vma;
190:
191:     if(!length) {
192:         return 0;
193:     }

```

## mm/mmap.c

Page 4/8

```

194:
195:     addr = MMAP_START;
196:     vma = current->vma;
197:
198:     for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
199:         if(vma->start < MMAP_START) {
200:             continue;
201:         }
202:         if(vma->start - addr >= length) {
203:             return PAGE_ALIGN(addr);
204:         }
205:         addr = PAGE_ALIGN(vma->end);
206:     }
207:     return 0;
208: }
209:
210: static void free_vma_pages(unsigned int start, __size_t length, struct vma *vma)
211: {
212:     unsigned int n, addr;
213:     unsigned int *pgdir, *pgtbl;
214:     unsigned int pde, pte, page;
215:     struct page *pg;
216:
217:     pgdir = (unsigned int *)P2V(current->tss.cr3);
218:     pgtbl = NULL;
219:
220:     for(n = 0; n < (length / PAGE_SIZE); n++) {
221:         pde = GET_PGDIR(start + (n * PAGE_SIZE));
222:         pte = GET_PGTBL(start + (n * PAGE_SIZE));
223:         if(pgdir[pde] & PAGE_PRESENT) {
224:             pgtbl = (unsigned int *)P2V((pgdir[pde] & PAGE_MASK));
225:             if(pgtbl[pte] & PAGE_PRESENT) {
226:                 /* make sure to not free reserved pages */
227:                 page = pgtbl[pte] >> PAGE_SHIFT;
228:                 pg = &page_table[page];
229:                 if(pg->flags & PAGE_RESERVED) {
230:                     continue;
231:                 }
232:
233:                 if(vma->prot & PROT_WRITE && vma->flags & MAP_SH
234: ARED) {
235:                     addr = start - vma->start + vma->offset;
236:                     write_page(pg, vma->inode, addr, length)
237: ;
238:                 }
239:
240:                 kfree(P2V(pgtbl[pte]) & PAGE_MASK);
241:                 current->rss--;
242:                 pgtbl[pte] = NULL;
243:
244:                 /* check if a page table can be freed */
245:                 for(pte = 0; pte < PT_ENTRIES; pte++) {
246:                     if(pgtbl[pte] & PAGE_MASK) {
247:                         break;
248:                     }
249:                 }
250:                 if(pte == PT_ENTRIES) {
251:                     kfree((unsigned int)pgtbl & PAGE_MASK);
252:                     current->rss--;
253:                     pgdir[pde] = NULL;
254:                 }
255:             }
256:         }
257:     }
258: static int free_vma_region(struct vma *vma, unsigned int start, __ssize_t length

```



## mm/mmap.c

Page 5/8

```

)
259: {
260:     struct vma *new;
261:
262:     if(!(new = get_new_vma_region())) {
263:         printk("WARNING: %s(): unable to get a free vma region.\n", __FU
NCTION__);
264:         return -ENOMEM;
265:     }
266:
267:     new->start = start + length;
268:     new->end = vma->end;
269:     new->prot = vma->prot;
270:     new->flags = vma->flags;
271:     new->offset = vma->offset;
272:     new->s_type = vma->s_type;
273:     new->inode = vma->inode;
274:     new->o_mode = vma->o_mode;
275:
276:     vma->end = start;
277:
278:     if(vma->start == vma->end) {
279:         if(vma->inode) {
280:             iput(vma->inode);
281:         }
282:         memset_b(vma, NULL, sizeof(struct vma));
283:     }
284:     if(new->start == new->end) {
285:         memset_b(new, NULL, sizeof(struct vma));
286:     }
287:     return 0;
288: }
289:
290: void release_binary(void)
291: {
292:     unsigned int n;
293:     struct vma *vma;
294:
295:     vma = current->vma;
296:
297:     for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
298:         free_vma_pages(vma->start, vma->end - vma->start, vma);
299:         free_vma_region(vma, vma->start, vma->end - vma->start);
300:     }
301:     sort_vma();
302:     optimize_vma();
303:     invalidate_tlb();
304: }
305:
306: struct vma * find_vma_region(unsigned int addr)
307: {
308:     unsigned int n;
309:     struct vma *vma;
310:
311:     if(!addr) {
312:         return NULL;
313:     }
314:
315:     addr &= PAGE_MASK;
316:     vma = current->vma;
317:
318:     for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
319:         if((addr >= vma->start) && (addr < vma->end)) {
320:             return vma;
321:         }
322:     }
323:     return NULL;

```

## mm/mmap.c

Page 6/8

```

324: }
325:
326: int expand_heap(unsigned int new)
327: {
328:     unsigned int n;
329:     struct vma *vma, *heap;
330:
331:     vma = current->vma;
332:     heap = NULL;
333:
334:     for(n = 0; n < VMA_REGIONS && vma->start; n++, vma++) {
335:         /* make sure the new heap won't overlap the next region */
336:         if(heap && new < vma->start) {
337:             heap->end = new;
338:             return 0;
339:         } else {
340:             heap = NULL;    /* was a bad candidate */
341:         }
342:         if(!heap && vma->s_type == P_HEAP) {
343:             heap = vma;    /* possible candidate */
344:             continue;
345:         }
346:     }
347:
348:     /* out of memory! */
349:     return 1;
350: }
351:
352: int do_mmap(struct inode *i, unsigned int start, unsigned int length, unsigned int
nt prot, unsigned int flags, unsigned int offset, char type, char mode)
353: {
354:     struct vma *vma;
355:     int errno;
356:
357:     if(!(length = PAGE_ALIGN(length))) {
358:         return start;
359:     }
360:
361:     if(start > KERNEL_BASE_ADDR || start + length > KERNEL_BASE_ADDR) {
362:         return -EINVAL;
363:     }
364:
365:     /* file mapping */
366:     if(i) {
367:         if(!S_ISREG(i->i_mode) && !S_ISCHR(i->i_mode)) {
368:             return -ENODEV;
369:         }
370:
371:         /*
372:          * The file shall have been opened with read permission,
373:          * regardless of the protection options specified.
374:          * IEEE Std 1003.1, 2004 Edition.
375:          */
376:         if(mode == O_WRONLY) {
377:             return -EACCES;
378:         }
379:         switch(flags & MAP_TYPE) {
380:             case MAP_SHARED:
381:                 if(prot & PROT_WRITE) {
382:                     if(!(mode & (O_WRONLY | O_RDWR))) {
383:                         return -EACCES;
384:                     }
385:                 }
386:                 break;
387:             case MAP_PRIVATE:
388:                 break;
389:             default:

```

## mm/mmap.c

Page 7/8

```

390:                                     return -EINVAL;
391:                                     }
392:                                     i->count++;
393:
394:                                     /* anonymous mapping */
395:                                     } else {
396:                                         if((flags & MAP_TYPE) != MAP_PRIVATE) {
397:                                             return -EINVAL;
398:                                         }
399:
400:                                         /* anonymous objects must be filled with zeros */
401:                                         flags |= ZERO_PAGE;
402:                                     }
403:
404:                                     if(flags & MAP_FIXED) {
405:                                         if(start & ~PAGE_MASK) {
406:                                             return -EINVAL;
407:                                         }
408:                                     } else {
409:                                         start = get_unmapped_vma_region(length);
410:                                         if(!start) {
411:                                             printk("WARNING: %s(): unable to get an unmapped vma reg
ion.\n", __FUNCTION__);
412:                                             return -ENOMEM;
413:                                         }
414:                                     }
415:
416:                                     if(!(vma = get_new_vma_region())) {
417:                                         printk("WARNING: %s(): unable to get a free vma region.\n", __FU
NCTION__);
418:                                         return -ENOMEM;
419:                                     }
420:
421:                                     vma->start = start;
422:                                     vma->end = start + length;
423:                                     vma->prot = prot;
424:                                     vma->flags = flags;
425:                                     vma->offset = offset;
426:                                     vma->s_type = type;
427:                                     vma->inode = i;
428:                                     vma->o_mode = mode;
429:
430:                                     if(i && i->fsop->mmap) {
431:                                         if((errno = i->fsop->mmap(i, vma))) {
432:                                             int errno2;
433:
434:                                             if((errno2 = free_vma_region(vma, start, length))) {
435:                                                 return errno2;
436:                                             }
437:                                             sort_vma();
438:                                             if((errno2 = optimize_vma())) {
439:                                                 return errno2;
440:                                             }
441:                                             return errno;
442:                                         }
443:                                     }
444:
445:                                     sort_vma();
446:                                     if((errno = optimize_vma())) {
447:                                         return errno;
448:                                     }
449:                                     return start;
450:     }
451:
452: int do_munmap(unsigned int addr, __size_t length)
453: {
454:     struct vma *vma;

```

## mm/mmap.c

Page 8/8

```

455:     unsigned int size;
456:     int errno;
457:
458:     if((addr & ~PAGE_MASK) || length < 0) {
459:         return -EINVAL;
460:     }
461:
462:     length = PAGE_ALIGN(length);
463:
464:     while(length) {
465:         if((vma = find_vma_region(addr))) {
466:             if((addr + length) > vma->end) {
467:                 size = vma->end - addr;
468:             } else {
469:                 size = length;
470:             }
471:
472:             free_vma_pages(addr, size, vma);
473:             invalidate_tlb();
474:             if((errno = free_vma_region(vma, addr, size))) {
475:                 return errno;
476:             }
477:             sort_vma();
478:             if((errno = optimize_vma())) {
479:                 return errno;
480:             }
481:             length -= size;
482:             addr += size;
483:         } else {
484:             break;
485:         }
486:     }
487:
488:     return 0;
489: }
490:
491: int do_mprotect(struct vma *vma, unsigned int addr, __size_t length, int prot)
492: {
493:     struct vma *new;
494:     int errno;
495:
496:     if(!(new = get_new_vma_region())) {
497:         printk("WARNING: %s(): unable to get a free vma region.\n", __FU
NCTION__);
498:         return -ENOMEM;
499:     }
500:
501:     new->start = addr;
502:     new->end = addr + length;
503:     new->prot = prot;
504:     new->flags = vma->flags;
505:     new->offset = vma->offset;
506:     new->s_type = vma->s_type;
507:     new->inode = vma->inode;
508:     new->o_mode = vma->o_mode;
509:
510:     sort_vma();
511:     if((errno = optimize_vma())) {
512:         return errno;
513:     }
514:     return 0;
515: }

```

## mm/page.c

Page 1/7

```

1: /*
2:  * fiwix/mm/page.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: /*
9:  * page.c implements a cache with a free list as a doubly circular linked
10:  * list and a chained hash table with doubly linked lists.
11:  *
12:  * hash table
13:  * +-----+ +-----+ +-----+ +-----+
14:  * | index | |prev|data|next| |prev|data|next| |prev|data|next|
15:  * |   0   --> | / |   |   |   | <--- |   |   |   | <--- |   |   |   |
16:  * +-----+ +-----+ +-----+ +-----+
17:  * +-----+ +-----+ +-----+ +-----+
18:  * | index | |prev|data|next| |prev|data|next| |prev|data|next|
19:  * |   1   --> | / |   |   |   | <--- |   |   |   | <--- |   |   |   |
20:  * +-----+ +-----+ +-----+ +-----+
21:  *                               (page)       (page)       (page)
22:  *      ...
23:  */
24:
25: #include <fiwix/asm.h>
26: #include <fiwix/kernel.h>
27: #include <fiwix/mm.h>
28: #include <fiwix/mman.h>
29: #include <fiwix/bios.h>
30: #include <fiwix/sleep.h>
31: #include <fiwix/sched.h>
32: #include <fiwix/devices.h>
33: #include <fiwix/buffer.h>
34: #include <fiwix/errno.h>
35: #include <fiwix/stdio.h>
36: #include <fiwix/string.h>
37:
38: #define PAGE_HASH(inode, offset)      (((__ino_t)(inode) ^ (__off_t)(offset))
% NR_PAGE_HASH)
39: #define NR_PAGES      page_table_size / sizeof(struct page)
40: #define NR_PAGE_HASH  page_hash_table_size / sizeof(unsigned int)
41:
42: struct page *page_table;           /* page pool */
43: struct page *page_head;           /* page pool head */
44: struct page **page_hash_table;
45:
46: static void insert_to_hash(struct page *pg)
47: {
48:     struct page **h;
49:     int i;
50:
51:     i = PAGE_HASH(pg->inode->inode, pg->offset);
52:     h = &page_hash_table[i];
53:
54:     if(!*h) {
55:         *h = pg;
56:         (*h)->prev_hash = (*h)->next_hash = NULL;
57:     } else {
58:         pg->prev_hash = NULL;
59:         pg->next_hash = *h;
60:         (*h)->prev_hash = pg;
61:         *h = pg;
62:     }
63:     kstat.cached += (PAGE_SIZE / 1024);
64: }
65:
66: static void remove_from_hash(struct page *pg)

```

## mm/page.c

Page 2/7

```

67: {
68:     struct page **h;
69:     int i;
70:
71:     if(!pg->inode) {
72:         return;
73:     }
74:
75:     i = PAGE_HASH(pg->inode->inode, pg->offset);
76:     h = &page_hash_table[i];
77:
78:     while(*h) {
79:         if(*h == pg) {
80:             if((*h)->next_hash) {
81:                 (*h)->next_hash->prev_hash = (*h)->prev_hash;
82:             }
83:             if((*h)->prev_hash) {
84:                 (*h)->prev_hash->next_hash = (*h)->next_hash;
85:             }
86:             if(h == &page_hash_table[i]) {
87:                 *h = (*h)->next_hash;
88:             }
89:             kstat.cached -= (PAGE_SIZE / 1024);
90:             break;
91:         }
92:         h = &(*h)->next_hash;
93:     }
94: }
95:
96: static void remove_from_free_list(struct page *pg)
97: {
98:     pg->prev_free->next_free = pg->next_free;
99:     pg->next_free->prev_free = pg->prev_free;
100:    kstat.free_pages--;
101:    if(pg == page_head) {
102:        page_head = pg->next_free;
103:    }
104: }
105:
106: void page_lock(struct page *pg)
107: {
108:     unsigned long int flags;
109:
110:     for(;;) {
111:         SAVE_FLAGS(flags); CLI();
112:         if(pg->locked) {
113:             RESTORE_FLAGS(flags);
114:             sleep(&pg, PROC_UNINTERRUPTIBLE);
115:         } else {
116:             break;
117:         }
118:     }
119:     pg->locked = 1;
120:     RESTORE_FLAGS(flags);
121: }
122:
123: void page_unlock(struct page *pg)
124: {
125:     unsigned long int flags;
126:
127:     SAVE_FLAGS(flags); CLI();
128:     pg->locked = 0;
129:     wakeup(pg);
130:     RESTORE_FLAGS(flags);
131: }
132:
133: struct page * get_free_page(void)

```

## mm/page.c

Page 3/7

```

134: {
135:     unsigned long int flags;
136:     struct page *pg;
137:
138:     /* if no more pages on free list */
139:     while(page_head == page_head->next_free) {
140:         /* reclaim some memory from buffer cache */
141:         wakeup(&kswapd);
142:         sleep(&get_free_page, PROC_UNINTERRUPTIBLE);
143:
144:         if(page_head == page_head->next_free) {
145:             /* definitely out of memory! (no more pages) */
146:             printk("%s(): pid %d ran out of memory. OOM killer neede
d!\n", __FUNCTION__, current->pid);
147:             return NULL;
148:         }
149:     }
150:
151:     SAVE_FLAGS(flags); CLI();
152:
153:     pg = page_head;
154:     remove_from_free_list(pg);
155:     remove_from_hash(pg); /* remove it from its old hash */
156:     pg->count = 1;
157:     pg->inode = NULL;
158:     pg->offset = 0;
159:
160:     RESTORE_FLAGS(flags);
161:     return pg;
162: }
163:
164: struct page * search_page_hash(struct inode *inode, __off_t offset)
165: {
166:     struct page *pg;
167:     int i;
168:
169:     i = PAGE_HASH(inode->inode, offset);
170:     pg = page_hash_table[i];
171:
172:     while(pg) {
173:         if(pg->inode == inode && pg->offset == offset) {
174:             if(!pg->count) {
175:                 remove_from_free_list(pg);
176:             }
177:             pg->count++;
178:             return pg;
179:         }
180:         pg = pg->next_hash;
181:     }
182:
183:     return NULL;
184: }
185:
186: void release_page(unsigned int page)
187: {
188:     unsigned long int flags;
189:     struct page *pg;
190:
191:     if(!valid_page(page)) {
192:         PANIC("Unexpected inconsistency in hash_table. Missing page %d (
0x%x).\n", page, page);
193:     }
194:
195:     pg = &page_table[page];
196:     if(--pg->count > 0) {
197:         return;
198:     }

```

## mm/page.c

Page 4/7

```

199:
200:     SAVE_FLAGS(flags); CLI();
201:
202:     if(!page_head) {
203:         pg->prev_free = pg->next_free = pg;
204:         page_head = pg;
205:     } else {
206:         pg->next_free = page_head;
207:         pg->prev_free = page_head->prev_free;
208:         page_head->prev_free->next_free = pg;
209:         page_head->prev_free = pg;
210:     }
211:
212:     /* if page is not cached then place it at the head of the free list */
213:     if(!pg->inode) {
214:         page_head = pg;
215:     }
216:
217:     kstat.free_pages++;
218:
219:     RESTORE_FLAGS(flags);
220:     wakeup(&get_free_page);
221: }
222:
223: int valid_page(unsigned int page)
224: {
225:     return (page >= 0 && page < NR_PAGES);
226: }
227:
228: void update_page_cache(struct inode *i, __off_t offset, const char *buf, int count)
229: {
230:     __off_t poffset;
231:     struct page *pg;
232:     int bytes;
233:
234:     poffset = offset % PAGE_SIZE;
235:     offset &= PAGE_MASK;
236:     bytes = PAGE_SIZE - poffset;
237:
238:     if(count) {
239:         bytes = MIN(bytes, count);
240:         if((pg = search_page_hash(i, offset))) {
241:             page_lock(pg);
242:             memcpy_b(pg->data + poffset, buf, bytes);
243:             page_unlock(pg);
244:             release_page(pg->page);
245:         }
246:     }
247: }
248:
249: int write_page(struct page *pg, struct inode *i, __off_t offset, unsigned int length)
250: {
251:     struct fd fd_table;
252:     unsigned int size;
253:     int errno;
254:
255:     size = MIN(i->i_size, length);
256:     fd_table.inode = i;
257:     fd_table.flags = 0;
258:     fd_table.count = 0;
259:     fd_table.offset = offset;
260:     if(i->fsop && i->fsop->write) {
261:         errno = i->fsop->write(i, &fd_table, pg->data, size);
262:     } else {
263:         errno = -EINVAL;

```



## mm/page.c

Page 5/7

```

264:         }
265:
266:         return errno;
267:     }
268:
269: int bread_page(struct page *pg, struct inode *i, __off_t offset, char prot, char
flags)
270: {
271:     __blk_t block;
272:     __off_t size_read;
273:     int blksize;
274:     struct device *d;
275:     struct buffer *buf;
276:
277:     blksize = i->sb->s_blocksize;
278:     size_read = 0;
279:
280:     if(!(d = get_device(BLK_DEV, i->dev))) {
281:         printk("WARNING: %s: device major %d not found!\n", __FUNCTION__
, MAJOR(i->dev));
282:         return 1;
283:     }
284:     if(!d->fsop || !d->fsop->read_block) {
285:         printk("WARNING: %s: device %d,%d does not have the read_block()
method!\n", __FUNCTION__, MAJOR(i->dev), MINOR(i->dev));
286:         return 1;
287:     }
288:
289:     pg->inode = i;
290:     pg->offset = offset;
291:     if(!(prot & PROT_WRITE) || flags & MAP_SHARED) {
292:         while(size_read < PAGE_SIZE) {
293:             if((block = bmap(i, offset, FOR_READING)) < 0) {
294:                 return 1;
295:             }
296:             if(block) {
297:                 /* does exist a buffer with recent data? */
298:                 if(!(buf = get_dirty_buffer(i->dev, block, blksize))) {
299:                     if(d->fsop->read_block(i->dev, block, pg
->data + size_read, blksize) < 0) {
300:                         return 1;
301:                     }
302:                 } else {
303:                     memcpy_b(pg->data + size_read, buf->data
, blksize);
304:                     brelse(buf);
305:                 }
306:             } else {
307:                 /* fill the hole with zeros */
308:                 memset_b(pg->data + size_read, 0, blksize);
309:             }
310:             size_read += blksize;
311:             offset += blksize;
312:         }
313:         /* cache all read-only and public (shared) pages */
314:         insert_to_hash(pg);
315:     } else {
316:         while(size_read < PAGE_SIZE) {
317:             if((block = bmap(i, offset, FOR_READING)) < 0) {
318:                 return 1;
319:             }
320:             if(block) {
321:                 /*
322:                  * This feeds the buffer cache by reading only
323:                  * the writable pages which aren't included in
324:                  * the page cache. This will speed up things by

```

## mm/page.c

Page 6/7

```

325:         * keeping in buffer cache the writable pages
326:         * with its original (disk) content (i.e. pages
327:         * of the data section of an ELF).
328:         */
329:         if(!(buf = bread(i->dev, block, blksize))) {
330:             return 1;
331:         }
332:         memcpy_b(pg->data + size_read, buf->data, blksize);
e);
333:         brelse(buf);
334:     }
335:     size_read += blksize;
336:     offset += blksize;
337: }
338: pg->inode = NULL;
339: pg->offset = 0;
340: }
341: return 0;
342: }
343:
344: int file_read(struct inode *i, struct fd *fd_table, char *buffer, __size_t count
)
345: {
346:     __off_t total_read;
347:     unsigned int page, poffset, bytes;
348:     struct page *pg;
349:
350:     inode_lock(i);
351:
352:     if(fd_table->offset > i->i_size) {
353:         fd_table->offset = i->i_size;
354:     }
355:
356:     total_read = 0;
357:
358:     for(;;) {
359:         count = (fd_table->offset + count > i->i_size) ? i->i_size - fd_
table->offset : count;
360:         if(!count) {
361:             break;
362:         }
363:
364:         poffset = fd_table->offset % PAGE_SIZE;
365:         if(!(pg = search_page_hash(i, fd_table->offset & PAGE_MASK))) {
366:             if(!(page = kmalloc())) {
367:                 inode_unlock(i);
368:                 printk("%s(): returning -ENOMEM\n", __FUNCTION__
);
369:                 return -ENOMEM;
370:             }
371:             page = V2P(page);
372:             pg = &page_table[page >> PAGE_SHIFT];
373:             if(bread_page(pg, i, fd_table->offset & PAGE_MASK, 0, MA
P_SHARED)) {
374:                 kfree((unsigned int)pg->data);
375:                 inode_unlock(i);
376:                 printk("%s(): returning -EIO\n", __FUNCTION__);
377:                 return -EIO;
378:             }
379:         }
380:
381:         page_lock(pg);
382:         bytes = PAGE_SIZE - poffset;
383:         bytes = MIN(bytes, count);
384:         memcpy_b(buffer + total_read, pg->data + poffset, bytes);
385:         total_read += bytes;
386:         count -= bytes;

```

## mm/page.c

Page 7/7

```

387:         poffset += bytes;
388:         poffset %= PAGE_SIZE;
389:         fd_table->offset += bytes;
390:         page_unlock(pg);
391:         kfree((unsigned int)pg->data);
392:     }
393:
394:     inode_unlock(i);
395:     return total_read;
396: }
397:
398: void page_init(unsigned int pages)
399: {
400:     struct page *pg;
401:     unsigned int n, addr;
402:
403:     memset_b(page_table, NULL, page_table_size);
404:     memset_b(page_hash_table, NULL, page_hash_table_size);
405:
406:     for(n = 0; n < pages; n++) {
407:         pg = &page_table[n];
408:         pg->page = n;
409:
410:         addr = n << PAGE_SHIFT;
411:         if(addr >= KERNEL_ENTRY_ADDR && addr < V2P(_last_data_addr)) {
412:             pg->flags = PAGE_RESERVED;
413:             kstat.kernel_reserved++;
414:             continue;
415:         }
416:
417:         /*
418:          * Some memory addresses are reserved, like the memory between
419:          * 0xA0000 and 0xFFFFF and other addresses, mostly used by the
420:          * VGA graphics adapter and BIOS.
421:          */
422:         if(!addr_in_bios_map(addr)) {
423:             pg->flags = PAGE_RESERVED;
424:             kstat.physical_reserved++;
425:             continue;
426:         }
427:
428:         pg->data = (char *)P2V(addr);
429:         if(!page_head) {
430:             pg->prev_free = pg->next_free = pg;
431:             page_head = pg;
432:         } else {
433:             pg->next_free = page_head;
434:             pg->prev_free = page_head->prev_free;
435:             page_head->prev_free->next_free = pg;
436:             page_head->prev_free = pg;
437:         }
438:         kstat.free_pages++;
439:     }
440:     kstat.total_mem_pages = kstat.free_pages;
441:     kstat.kernel_reserved <=< 2;
442:     kstat.physical_reserved <=< 2;
443: }

```

## mm/swapper.c

Page 1/1

```

1: /*
2:  * fiwix/mm/swapper.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/limits.h>
10: #include <fiwix/process.h>
11: #include <fiwix/sleep.h>
12: #include <fiwix/sched.h>
13: #include <fiwix/tty.h>
14: #include <fiwix/memdev.h>
15: #include <fiwix/serial.h>
16: #include <fiwix/lp.h>
17: #include <fiwix/ramdisk.h>
18: #include <fiwix/floppy.h>
19: #include <fiwix/ide.h>
20: #include <fiwix/buffer.h>
21: #include <fiwix/mm.h>
22: #include <fiwix/fs.h>
23: #include <fiwix/locks.h>
24: #include <fiwix/filesystems.h>
25: #include <fiwix/stdio.h>
26:
27: /* kswapd continues the kernel initialization */
28: int kswapd(void)
29: {
30:     sprintk(current->argv0, "%s", "kswapd");
31:
32:     /* char devices */
33:     memdev_init();
34:     serial_init();
35:     lp_init();
36:
37:     /* block devices */
38:     ramdisk_init();
39:     floppy_init();
40:     ide_init();
41:
42:     /* data structures */
43:     sleep_init();
44:     buffer_init();
45:     sched_init();
46:     mount_init();
47:     inode_init();
48:     fd_init();
49:     flock_init();
50:
51:     mem_stats();
52:     fs_init();
53:     mount_root();
54:     init_init();
55:
56:     for(;;) {
57:         sleep(&kswapd, PROC_UNINTERRUPTIBLE);
58:         if(reclaim_buffers()) {
59:             continue;
60:         }
61:         printk("WARNING: %s(): out of memory and swapping is not impleme
nted yet, sorry.\n", __FUNCTION__);
62:         wakeup(&get_free_page);
63:     }
64: }

```

## lib/ctype.c

```

1: /*
2:  * fiwix/lib/ctype.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/ctype.h>
9:
10: unsigned char _ctype[] = {
11:     0,
12:     _C,          /* ^@  0x00 (NUL '\0') */
13:     _C,          /* ^A  0x01 (SOH) */
14:     _C,          /* ^B  0x02 (STX) */
15:     _C,          /* ^C  0x03 (ETX) */
16:     _C,          /* ^D  0x04 (EOT) */
17:     _C,          /* ^E  0x05 (ENQ) */
18:     _C,          /* ^F  0x06 (ACK) */
19:     _C,          /* ^G  0x07 (BEL '\a') */
20:     _C,          /* ^H  0x08 (BS '\b') */
21:     _C  _S,     /* ^I  0x09 (HT '\t') */
22:     _C  _S,     /* ^J  0x0A (LF '\n') */
23:     _C  _S,     /* ^K  0x0B (VT '\v') */
24:     _C  _S,     /* ^L  0x0C (FF '\f') */
25:     _C  _S,     /* ^M  0x0D (CR '\r') */
26:     _C,          /* ^N  0x0E (SO) */
27:     _C,          /* ^O  0x0F (SI) */
28:     _C,          /* ^P  0x10 (DLE) */
29:     _C,          /* ^Q  0x11 (DC1) */
30:     _C,          /* ^R  0x12 (DC2) */
31:     _C,          /* ^S  0x13 (DC3) */
32:     _C,          /* ^T  0x14 (DC4) */
33:     _C,          /* ^U  0x15 (NAK) */
34:     _C,          /* ^V  0x16 (SYN) */
35:     _C,          /* ^W  0x17 (ETB) */
36:     _C,          /* ^X  0x18 (CAN) */
37:     _C,          /* ^Y  0x19 (EM) */
38:     _C,          /* ^Z  0x1A (SUB) */
39:     _C,          /* ^[  0x1B (ESC) */
40:     _C,          /* ^\  0x1C (FS) */
41:     _C,          /* ^]  0x1D (GS) */
42:     _C,          /* ^^  0x1E (RS) */
43:     _C,          /* ^_  0x1F (US) */
44:     _S,          /* ' '  0x20 */
45:     _P,          /* '! ' 0x21 */
46:     _P,          /* '" ' 0x22 */
47:     _P,          /* '# ' 0x23 */
48:     _P,          /* '$ ' 0x24 */
49:     _P,          /* '% ' 0x25 */
50:     _P,          /* '& ' 0x26 */
51:     _P,          /* '' ' 0x27 */
52:     _P,          /* '(' ' 0x28 */
53:     _P,          /* ')' ' 0x29 */
54:     _P,          /* '*' ' 0x2A */
55:     _P,          /* '+' ' 0x2B */
56:     _P,          /* ',' ' 0x2C */
57:     _P,          /* '-' ' 0x2D */
58:     _P,          /* '.' ' 0x2E */
59:     _P,          /* '/' ' 0x2F */
60:     _N,          /* '0'  0x30 */
61:     _N,          /* '1'  0x31 */
62:     _N,          /* '2'  0x32 */
63:     _N,          /* '3'  0x33 */
64:     _N,          /* '4'  0x34 */
65:     _N,          /* '5'  0x35 */
66:     _N,          /* '6'  0x36 */
67:     _N,          /* '7'  0x37 */

```

## lib/ctype.c

Page 2/3

```

68:      _N,          /* '8' 0x38 */
69:      _N,          /* '9' 0x39 */
70:      _P,          /* ':' 0x3A */
71:      _P,          /* ';' 0x3B */
72:      _P,          /* '<' 0x3C */
73:      _P,          /* '=' 0x3D */
74:      _P,          /* '>' 0x3E */
75:      _P,          /* '?' 0x3F */
76:      _P,          /* '@' 0x40 */
77:      _U,  _X,     /* 'A' 0x41 */
78:      _U,  _X,     /* 'B' 0x42 */
79:      _U,  _X,     /* 'C' 0x43 */
80:      _U,  _X,     /* 'D' 0x44 */
81:      _U,  _X,     /* 'E' 0x45 */
82:      _U,  _X,     /* 'F' 0x46 */
83:      _U,          /* 'G' 0x47 */
84:      _U,          /* 'H' 0x48 */
85:      _U,          /* 'I' 0x49 */
86:      _U,          /* 'J' 0x4A */
87:      _U,          /* 'K' 0x4B */
88:      _U,          /* 'L' 0x4C */
89:      _U,          /* 'M' 0x4D */
90:      _U,          /* 'N' 0x4E */
91:      _U,          /* 'O' 0x4F */
92:      _U,          /* 'P' 0x50 */
93:      _U,          /* 'Q' 0x51 */
94:      _U,          /* 'R' 0x52 */
95:      _U,          /* 'S' 0x53 */
96:      _U,          /* 'T' 0x54 */
97:      _U,          /* 'U' 0x55 */
98:      _U,          /* 'V' 0x56 */
99:      _U,          /* 'W' 0x57 */
100:     _U,          /* 'X' 0x58 */
101:     _U,          /* 'Y' 0x59 */
102:     _U,          /* 'Z' 0x5A */
103:     _P,          /* '[' 0x5B */
104:     _P,          /* '\' 0x5C */
105:     _P,          /* ']' 0x5D */
106:     _P,          /* '^' 0x5E */
107:     _P,          /* '_' 0x5F */
108:     _P,          /* '`' 0x60 */
109:     _L,  _X,     /* 'a' 0x61 */
110:     _L,  _X,     /* 'b' 0x62 */
111:     _L,  _X,     /* 'c' 0x63 */
112:     _L,  _X,     /* 'd' 0x64 */
113:     _L,  _X,     /* 'e' 0x65 */
114:     _L,  _X,     /* 'f' 0x66 */
115:     _L,          /* 'g' 0x67 */
116:     _L,          /* 'h' 0x68 */
117:     _L,          /* 'i' 0x69 */
118:     _L,          /* 'j' 0x6A */
119:     _L,          /* 'k' 0x6B */
120:     _L,          /* 'l' 0x6C */
121:     _L,          /* 'm' 0x6D */
122:     _L,          /* 'n' 0x6E */
123:     _L,          /* 'o' 0x6F */
124:     _L,          /* 'p' 0x70 */
125:     _L,          /* 'q' 0x71 */
126:     _L,          /* 'r' 0x72 */
127:     _L,          /* 's' 0x73 */
128:     _L,          /* 't' 0x74 */
129:     _L,          /* 'u' 0x75 */
130:     _L,          /* 'v' 0x76 */
131:     _L,          /* 'w' 0x77 */
132:     _L,          /* 'x' 0x78 */
133:     _L,          /* 'y' 0x79 */
134:     _L,          /* 'z' 0x7A */

```

**lib/ctype.c**

Page 3/3

```
135:         _P,           /* '{' 0x7B */
136:         _P,           /* '|' 0x7C */
137:         _P,           /* '}' 0x7D */
138:         _P,           /* '~' 0x7E */
139:         _C,           /* DEL 0x7F */
140:     };
```

**lib/Makefile**

Page 1/1

```
1: # fiwix/lib/Makefile
2: #
3: # Copyright 2018, Jordi Sanfeliu. All rights reserved.
4: # Distributed under the terms of the Fiwix License.
5: #
6:
7: .S.o:
8:     $(CC) -traditional -I$(INCLUDE) -c -o $@ $<
9: .c.o:
10:     $(CC) $(CFLAGS) -c -o $@ $<
11:
12: OBJS = ctype.o strings.o printk.o
13:
14: lib:     $(OBJS)
15:           $(LD) $(LDFLAGS) -r $(OBJS) -o lib.o
16:
17: clean:
18:           rm -f *.o
19:
```



## lib/printk.c

Page 1/6

```

1: /*
2:  * fiwix/lib/printk.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/kernel.h>
9: #include <fiwix/tty.h>
10: #include <fiwix/stdio.h>
11: #include <fiwix/string.h>
12: #include <fiwix/stdarg.h>
13:
14: #define LOG_BUF_LEN      4096
15: #define MAX_BUF          1024    /* printk() and sprintf() size limit */
16:
17: static char log_buf[LOG_BUF_LEN];
18: static unsigned int log_count;
19:
20: static void puts(char *buffer)
21: {
22:     struct tty *tty;
23:     unsigned short int count;
24:     char *b;
25:
26:     /* for special debugging purposes only (X11, SVGAlib, ...)
27:     {
28:         struct inode dummy_i;
29:         memset_b(&dummy_i, 0, sizeof(struct inode));
30:         dummy_i.dev = dummy_i.rdev = 0x0600;           // /dev/lp0
31:         lp_write(&dummy_i, NULL, buffer, strlen(buffer));
32:     }
33:     */
34:
35:     tty = get_tty(_syscondev);
36:     count = strlen(buffer);
37:     b = buffer;
38:
39:     while(count-->0) {
40:         if(!tty) {
41:             if(log_count < LOG_BUF_LEN) {
42:                 log_buf[log_count++] = *(b++);
43:             }
44:         } else {
45:             tty_queue_putchar(tty, &tty->write_q, *(b++));
46:
47:             /* kernel messages must be shown immediately */
48:             tty->output(tty);
49:         }
50:     }
51: }
52:
53: /*
54:  * format identifiers
55:  * -----
56:  *      %d      decimal conversion
57:  *      %u      unsigned decimal conversion
58:  *      %x      hexadecimal conversion (lower case)
59:  *      %X      hexadecimal conversion (upper case)
60:  *      %b      binary conversion
61:  *      %o      octal conversion
62:  *      %c      character
63:  *      %s      string
64:  *
65:  * flags
66:  * -----
67:  *      0      result is padded with zeros (e.g.: '%06d')

```

## lib/printk.c

Page 2/6

```

68:  *          (maximum value is 32)
69:  *      blank result is padded with spaces (e.g.: '% 6d')
70:  *          (maximum value is 32)
71:  *      -    the numeric result is left-justified
72:  *          (default is right-justified)
73:  */
74: static void do_printk(char *buffer, const char *format, va_list args)
75: {
76:     char sw_neg, in_identifier, n_pad, lf;
77:     char ch_pad, basecase, c;
78:     char str[] = {
79:         NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL,
80:         NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL,
81:         NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL,
82:         NULL, NULL, NULL, NULL, NULL, NULL, NULL, NULL,
83:         NULL
84:     };
85:     char nullstr[7] = { '<', 'N', 'U', 'L', 'L', '>', '\0' };
86:     char *ptr_s, *p;
87:     int num, count;
88:     char simplechar;
89:     unsigned int unum, digit;
90:
91:     sw_neg = in_identifier = n_pad = lf = 0;
92:     count = 0;
93:     basecase = 'A';
94:     ch_pad = ' ';
95:     p = NULL;
96:
97:     /* assumes buffer has a maximum size of MAX_BUF */
98:     while((c = *(format++)) && count < MAX_BUF) {
99:         if((c != '%') && !in_identifier) {
100:             *(buffer++) = c;
101:             memset_b(str, NULL, 32);
102:         } else {
103:             in_identifier = 1;
104:             switch(c = *(format)) {
105:                 case 'd':
106:                     num = va_arg(args, int);
107:                     if(num < 0) {
108:                         num *= -1;
109:                         sw_neg = 1;
110:                     }
111:                     ptr_s = str;
112:                     do {
113:                         *(ptr_s++) = '0' + (num % 10);
114:                     } while(num /= 10);
115:                     if(lf) {
116:                         p = ptr_s;
117:                     } else {
118:                         while(*ptr_s) {
119:                             ptr_s++;
120:                         }
121:                     }
122:                     if(sw_neg) {
123:                         sw_neg = 0;
124:                         *(ptr_s++) = '-';
125:                     }
126:                     do {
127:                         *(buffer++) = *(--ptr_s);
128:                         count++;
129:                     } while(ptr_s != str && count < MAX_BUF)
;
130:                     if(lf) {
131:                         while(*p && count < MAX_BUF) {
132:                             *(buffer++) = *(p++);
133:                             count++;

```

## lib/printk.c

Page 3/6

```

134:                                     }
135:                                     }
136:                                     format++;
137:                                     ch_pad = ' ';
138:                                     n_pad = 0;
139:                                     in_identifier = 0;
140:                                     lf = 0;
141:                                     break;
142:
143:                                     case 'u':
144:                                         unum = va_arg(args, unsigned int);
145:                                         ptr_s = str;
146:                                         do {
147:                                             *(ptr_s++) = '0' + (unum % 10);
148:                                         } while(unum /= 10);
149:                                         if(lf) {
150:                                             p = ptr_s;
151:                                         } else {
152:                                             while(*ptr_s) {
153:                                                 ptr_s++;
154:                                             }
155:                                         }
156:                                         do {
157:                                             *(buffer++) = *--ptr_s;
158:                                             count++;
159:                                         } while(ptr_s != str && count < MAX_BUF)
;
160:                                     if(lf) {
161:                                         while(*p && count < MAX_BUF) {
162:                                             *(buffer++) = *(p++);
163:                                             count++;
164:                                         }
165:                                     }
166:                                     format++;
167:                                     ch_pad = ' ';
168:                                     n_pad = 0;
169:                                     in_identifier = 0;
170:                                     lf = 0;
171:                                     break;
172:
173:                                     case 'x':
174:                                         basecase = 'a';
175:                                     case 'X':
176:                                         unum = va_arg(args, unsigned int);
177:                                         ptr_s = str;
178:                                         do {
179:                                             *(ptr_s++) = (digit = (unum & 0x
0F)) > 9 ? basecase + digit - 10 : '0' + digit;
180:                                         } while(unum /= 16);
181:                                         if(lf) {
182:                                             p = ptr_s;
183:                                         } else {
184:                                             while(*ptr_s) {
185:                                                 ptr_s++;
186:                                             }
187:                                         }
188:                                         do {
189:                                             *(buffer++) = *--ptr_s;
190:                                             count++;
191:                                         } while(ptr_s != str && count < MAX_BUF)
;
192:                                     if(lf) {
193:                                         while(*p && count < MAX_BUF) {
194:                                             *(buffer++) = *(p++);
195:                                             count++;
196:                                         }
197:                                     }

```

## lib/printk.c

Page 4/6

```

198:         format++;
199:         ch_pad = ' ';
200:         n_pad = 0;
201:         in_identifier = 0;
202:         lf = 0;
203:         break;
204:
205:     case 'b':
206:         num = va_arg(args, unsigned int);
207:         if(num < 0) {
208:             num *= -1;
209:         }
210:         ptr_s = str;
211:         do {
212:             *(ptr_s++) = '0' + (num % 2);
213:         } while(num /= 2);
214:         if(lf) {
215:             p = ptr_s;
216:         } else {
217:             while(*ptr_s) {
218:                 ptr_s++;
219:             }
220:         }
221:         do {
222:             *(buffer++) = *(--ptr_s);
223:             count++;
224:         } while(ptr_s != str && count < MAX_BUF)
;
225:
226:         if(lf) {
227:             while(*p && count < MAX_BUF) {
228:                 *(buffer++) = *(p++);
229:                 count++;
230:             }
231:         }
232:         format++;
233:         ch_pad = ' ';
234:         n_pad = 0;
235:         in_identifier = 0;
236:         lf = 0;
237:         break;
238:
239:     case 'o':
240:         num = va_arg(args, unsigned int);
241:         if(num < 0) {
242:             num *= -1;
243:         }
244:         ptr_s = str;
245:         do {
246:             *(ptr_s++) = '0' + (num % 8);
247:         } while(num /= 8);
248:         if(lf) {
249:             p = ptr_s;
250:         } else {
251:             while(*ptr_s) {
252:                 ptr_s++;
253:             }
254:         }
255:         do {
256:             *(buffer++) = *(--ptr_s);
257:             count++;
258:         } while(ptr_s != str && count < MAX_BUF)
;
259:
260:         if(lf) {
261:             while(*p && count < MAX_BUF) {
262:                 *(buffer++) = *(p++);

```

## lib/printk.c

Page 5/6

```

263:         }
264:         format++;
265:         ch_pad = ' ';
266:         n_pad = 0;
267:         in_identifier = 0;
268:         lf = 0;
269:         break;
270:
271:     case 'c':
272:         simplechar = va_arg(args, int);
273:         *(buffer++) = simplechar;
274:         format++;
275:         in_identifier = 0;
276:         lf = 0;
277:         break;
278:
279:     case 's':
280:         num = 0;
281:         ptr_s = va_arg(args, char *);
282:         if(n_pad) {
283:             num = n_pad - strlen(ptr_s);
284:             if(num < 0) {
285:                 num *= -1;
286:             }
287:         }
288:         /* if it's a NULL then show "<NULL>" */
289:         if(ptr_s == NULL) {
290:             ptr_s = (char *)nullstr;
291:         }
292:         while((c = *(ptr_s++)) && count < MAX_BU
F) {
293:             *(buffer++) = c;
294:             count++;
295:         }
296:         while(num-- && count < MAX_BUF) {
297:             *(buffer++) = ' ';
298:             count++;
299:         }
300:         format++;
301:         n_pad = 0;
302:         in_identifier = 0;
303:         lf = 0;
304:         break;
305:
306:     case ' ':
307:         ch_pad = ' ';
308:         break;
309:
310:     case '0':
311:         if(!n_pad) {
312:             ch_pad = '0';
313:         }
314:     case '1':
315:     case '2':
316:     case '3':
317:     case '4':
318:     case '5':
319:     case '6':
320:     case '7':
321:     case '8':
322:     case '9':
323:         n_pad = !n_pad ? c - '0': ((n_pad * 10)
+ (c - '0'));
324:         n_pad = n_pad > 32 ? 32 : n_pad;
325:         for(unum = 0; unum < n_pad; unum++) {
326:             str[unum] = ch_pad;
327:         }

```

## lib/printk.c

Page 6/6

```
328:                                     break;
329:
330:                                     case '-':
331:                                         lf = 1;
332:                                         break;
333:                                     case '%':
334:                                         *(buffer++) = c;
335:                                         format++;
336:                                         in_identifier = 0;
337:                                         break;
338:                                     }
339:                                 }
340:                             count++;
341:     }
342:     *buffer = NULL;
343: }
344:
345: void register_console(void (*fn)(char *, unsigned int))
346: {
347:     (*fn)(log_buf, log_count);
348: }
349:
350: void printk(const char *format, ...)
351: {
352:     va_list args;
353:     char buffer[MAX_BUF];
354:
355:     va_start(args, format);
356:     do_printk(buffer, format, args);
357:     puts(buffer);
358:     va_end(args);
359: }
360:
361: int sprintf(char *buffer, const char *format, ...)
362: {
363:     va_list args;
364:
365:     va_start(args, format);
366:     do_printk(buffer, format, args);
367:     va_end(args);
368:     return strlen(buffer);
369: }
```

## lib/strings.c

```

1: /*
2:  * fiwix/lib/strings.c
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #include <fiwix/types.h>
9: #include <fiwix/tty.h>
10: #include <fiwix/mm.h>
11: #include <fiwix/stdio.h>
12: #include <fiwix/string.h>
13:
14: /* convert from big-endian to little-endian (word swap) */
15: void swap_asc_word(char *str, int len)
16: {
17:     int n, n2;
18:     short int *ptr;
19:     char *buf;
20:
21:     if(!(buf = (void *)kmalloc())) {
22:         return;
23:     }
24:
25:     ptr = (short int *)str;
26:
27:     for(n = 0, n2 = 0; n < len; n++) {
28:         buf[n2++] = *ptr >> 8;
29:         buf[n2++] = *ptr & 0xFF;
30:         ptr++;
31:     }
32:     for(n = len - 1; n > 0; n--) {
33:         if(buf[n] == NULL || buf[n] == ' ') {
34:             buf[n] = NULL;
35:         } else {
36:             break;
37:         }
38:     }
39:     memcpy_b(str, buf, len);
40:     kfree((unsigned int)buf);
41: }
42:
43: int strcmp(const char *str1, const char *str2)
44: {
45:     while(*str1) {
46:         if(*str1 != *str2) {
47:             return 1;
48:         }
49:         str1++;
50:         str2++;
51:     }
52:     if(!(*str2)) {
53:         return 0;
54:     }
55:     return 1;
56: }
57:
58: int strncmp(const char *str1, const char *str2, __ssize_t n)
59: {
60:     while(n > 0) {
61:         if(*str1 != *str2) {
62:             return 1;
63:         }
64:         str1++;
65:         str2++;
66:         n--;
67:     }

```

## lib/strings.c

Page 2/5

```
68:         return 0;
69:     }
70:
71:     char *strcpy(char *dest, const char *src)
72:     {
73:         if(!dest || !src) {
74:             return NULL;
75:         }
76:
77:         while(*src) {
78:             *dest = *src;
79:             dest++;
80:             src++;
81:         }
82:         *dest = NULL;          /* NULL-terminated */
83:         return dest;
84:     }
85:
86:     void strncpy(char *dest, const char *src, int len)
87:     {
88:         if(!dest || !src) {
89:             return;
90:         }
91:
92:         while((*src) && len) {
93:             *dest = *src;
94:             dest++;
95:             src++;
96:             len--;
97:         }
98:         *dest = NULL;        /* NULL-terminated */
99:     }
100:
101:     char *strcat(char *dest, const char *src)
102:     {
103:         char *orig;
104:
105:         orig = dest;
106:         while(*dest) {
107:             dest++;
108:         }
109:         while(*src) {
110:             *dest = *src;
111:             dest++;
112:             src++;
113:         }
114:         *dest = NULL;
115:         return orig;
116:     }
117:
118:     char *strncat(char *dest, const char *src, __ssize_t len)
119:     {
120:         char *orig;
121:
122:         orig = dest;
123:         while(*dest) {
124:             dest++;
125:         }
126:         while(*src && len) {
127:             *dest = *src;
128:             dest++;
129:             src++;
130:             len--;
131:         }
132:         *dest = NULL;
133:         return orig;
134:     }
```



## lib/strings.c

Page 3/5

```
135:
136: int strlen(const char *str)
137: {
138:     int n;
139:
140:     n = 0;
141:     while(str && *str) {
142:         n++;
143:         str++;
144:     }
145:     return n;
146: }
147:
148: char * get_basename(const char *path)
149: {
150:     char *basename;
151:     char c;
152:
153:     basename = NULL;
154:
155:     while(path) {
156:         while(*path == '/') {
157:             path++;
158:         }
159:         if(*path != NULL) {
160:             basename = (char *)path;
161:         }
162:         while((c = *(path++)) && (c != '/'));
163:         if(!c) {
164:             break;
165:         }
166:     }
167:     return basename;
168: }
169:
170: char * remove_trailing_slash(char *path)
171: {
172:     char *p;
173:
174:     p = path + (strlen(path) - 1);
175:     while(p > path && *p == '/') {
176:         *p = NULL;
177:         p--;
178:     }
179:     return path;
180: }
181:
182: int is_dir(const char *path)
183: {
184:     while*(path + 1) {
185:         path++;
186:     }
187:     if(*path == '/') {
188:         return 1;
189:     }
190:     return 0;
191: }
192:
193: int atoi(const char *str)
194: {
195:     int n;
196:
197:     n = 0;
198:     while(IS_SPACE(*str)) {
199:         str++;
200:     }
201:     while(IS_NUMERIC(*str)) {
```

## lib/strings.c

Page 4/5

```
202:             n = (n * 10) + (*str++ - '0');
203:         }
204:         return n;
205:     }
206:
207: void memcpy_b(void *dest, const void *src, unsigned int count)
208: {
209:     unsigned char *d;
210:     unsigned char *s;
211:
212:     d = (unsigned char *)dest;
213:     s = (unsigned char *)src;
214:     while(count--) {
215:         *d = *s;
216:         d++;
217:         s++;
218:     }
219: }
220:
221: void memcpy_w(void *dest, const void *src, unsigned int count)
222: {
223:     unsigned short int *d;
224:     unsigned short int *s;
225:
226:     d = (unsigned short int *)dest;
227:     s = (unsigned short int *)src;
228:     while(count--) {
229:         *d = *s;
230:         d++;
231:         s++;
232:     }
233: }
234:
235: void memcpy_l(void *dest, const void *src, unsigned int count)
236: {
237:     unsigned int *d;
238:     unsigned int *s;
239:
240:     d = (unsigned int *)dest;
241:     s = (unsigned int *)src;
242:     while(count--) {
243:         *d = *s;
244:         d++;
245:         s++;
246:     }
247: }
248:
249: void memset_b(void *dest, unsigned char value, unsigned int count)
250: {
251:     unsigned char *d;
252:
253:     d = (unsigned char *)dest;
254:     while(count--) {
255:         *d = value;
256:         d++;
257:     }
258: }
259:
260: void memset_w(void *dest, unsigned short int value, unsigned int count)
261: {
262:     unsigned short int *d;
263:
264:     d = (unsigned short int *)dest;
265:     while(count--) {
266:         *d = value;
267:         d++;
268:     }
}
```

**lib/strings.c**

Page 5/5

```
269: }
270:
271: void memset_l(void *dest, unsigned int value, unsigned int count)
272: {
273:     unsigned int *d;
274:
275:     d = (unsigned int *)dest;
276:     while(count--) {
277:         *d = value;
278:         d++;
279:     }
280: }
```

**include/fiwix/asm.h**

Page 1/3

```

1: /*
2:  * fiwix/include/fiwix/asm.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_ASM_H
9: #define _FIWIX_ASM_H
10:
11: extern void except0(void);
12: extern void except1(void);
13: extern void except2(void);
14: extern void except3(void);
15: extern void except4(void);
16: extern void except5(void);
17: extern void except6(void);
18: extern void except7(void);
19: extern void except8(void);
20: extern void except9(void);
21: extern void exceptA(void);
22: extern void exceptB(void);
23: extern void exceptC(void);
24: extern void exceptD(void);
25: extern void exceptE(void);
26: extern void exceptF(void);
27: extern void except10(void);
28: extern void except11(void);
29: extern void except12(void);
30: extern void except13(void);
31: extern void except14(void);
32: extern void except15(void);
33: extern void except16(void);
34: extern void except17(void);
35: extern void except18(void);
36: extern void except19(void);
37: extern void except1A(void);
38: extern void except1B(void);
39: extern void except1C(void);
40: extern void except1D(void);
41: extern void except1E(void);
42: extern void except1F(void);
43:
44: extern void irq0(void);
45: extern void irq1(void);
46: extern void irq2(void);
47: extern void irq3(void);
48: extern void irq4(void);
49: extern void irq5(void);
50: extern void irq6(void);
51: extern void irq7(void);
52: extern void irq8(void);
53: extern void irq9(void);
54: extern void irq10(void);
55: extern void irq11(void);
56: extern void irq12(void);
57: extern void irq13(void);
58: extern void irq14(void);
59: extern void irq15(void);
60: extern void unknown_irq(void);
61:
62: extern void switch_to_user_mode(void);
63: extern void sighandler_trampoline(void);
64: extern void end_sighandler_trampoline(void);
65: extern void syscall(void);
66: extern void return_from_syscall(void);
67: extern void do_switch(unsigned int *, unsigned int *, unsigned int, unsigned int

```

## include/fiwix/asm.h

Page 2/3

```

, unsigned int, unsigned short int);
68:
69: int cpuid(void);
70: int getfpu(void);
71: int vendor_id(void);
72: int signature_flags(void);
73: int brand_str(void);
74: int tlbinfo(void);
75:
76: unsigned char inport_b(unsigned int);
77: short int inport_w(unsigned int);
78: void inport_sw(unsigned int, void *, unsigned int);
79: void outport_b(unsigned int, unsigned char);
80: void outport_w(unsigned int, unsigned int);
81: void outport_sw(unsigned int, void *, unsigned int);
82:
83: void load_gdt(unsigned int);
84: void load_idt(unsigned int);
85: void activate_kpage_dir(void);
86: void load_tr(unsigned int);
87: unsigned long long int get_rdtsc(void);
88: void invalidate_tlb(void);
89:
90: #define CLI() __asm__ __volatile__ ("cli"::"memory")
91: #define STI() __asm__ __volatile__ ("sti"::"memory")
92: #define NOP() __asm__ __volatile__ ("nop"::"memory")
93: #define HLT() __asm__ __volatile__ ("hlt"::"memory")
94:
95: #define GET_CR2(cr2) __asm__ __volatile__ ("movl %%cr2, %0" : "=r" (cr2));
96: #define GET_ESP(esp) __asm__ __volatile__ ("movl %%esp, %0" : "=r" (esp));
97: #define SET_ESP(esp) __asm__ __volatile__ ("movl %0, %%esp" : "=r" (esp));
98:
99: #define SAVE_FLAGS(flags) \
100:     __asm__ __volatile__( \
101:         "pushfl ; popl %0\n\t" \
102:         : "=r" (flags) \
103:         : /* no input */ \
104:         : "memory" \
105:     );
106:
107: #define RESTORE_FLAGS(x) \
108:     __asm__ __volatile__( \
109:         "pushl %0 ; popfl\n\t" \
110:         : /* no output */ \
111:         : "r" (flags) \
112:         : "memory" \
113:     );
114:
115: #define USER_SYSCALL(num, arg1, arg2, arg3) \
116:     __asm__ __volatile__( \
117:         "movl %0, %%eax\n\t" \
118:         "movl %1, %%ebx\n\t" \
119:         "movl %2, %%ecx\n\t" \
120:         "movl %3, %%edx\n\t" \
121:         "int $0x80\n\t" \
122:         : /* no output */ \
123:         : "eax"((unsigned int)num), "ebx"((unsigned int)arg1), "ecx"((un- \
signed int)arg2), "edx"((unsigned int)arg3) \
124:     );
125:
126: /*
127: static inline unsigned long long int get_rdtsc(void)
128: {
129:     unsigned int eax, edx;
130:
131:     __asm__ __volatile__("rdtsc" : "=a" (eax), "=d" (edx));
132:     return ((unsigned long long int)eax) | (((unsigned long long int)edx) <<

```

**include/fiwix/asm.h**

Page 3/3

```
32);  
133: }  
134: */  
135:  
136: #endif /* _FIWIX_ASM_H */
```

**include/fiwix/bios.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/bios.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_BIOS_H
9: #define _FIWIX_BIOS_H
10:
11: #include <fiwix/multiboot.h>
12:
13: #define BIOS_MEM_AVAIL          1      /* BIOS memory available */
14: #define BIOS_MEM_RES           2      /* BIOS memory reserved */
15: #define BIOS_MEM_ACPI_REC      3      /* BIOS memory ACPI reclaim */
16: #define BIOS_MEM_ACPI_NVS      4      /* BIOS memory ACPI NVS */
17: #define NR_BIOS_MM_ENT         25     /* entries in BIOS memory map */
18:
19: struct bios_mem_map {
20:     unsigned long int from;
21:     unsigned long int to;
22:     unsigned long int type;
23: };
24: struct bios_mem_map bios_mem_map[NR_BIOS_MM_ENT];
25:
26: int addr_in_bios_map(unsigned int);
27: void bios_map_init(memory_map_t *, unsigned long int);
28:
29: #endif /* _FIWIX_BIOS_H */
```

**include/fiwix/buffer.h**

Page 1/1

```

1:  /*
2:  *  fiwix/include/fiwix/buffer.h
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #ifndef _FIWIX_BUFFER_H
9:  #define _FIWIX_BUFFER_H
10:
11: #include <fiwix/types.h>
12: #include <fiwix/fs.h>
13:
14: struct buffer {
15:     __dev_t dev;                /* device number */
16:     __blk_t block;             /* block number */
17:     int size;                  /* block size (in bytes) */
18:     char valid;                /* 1 = valid */
19:     unsigned char locked;      /* 1 = locked */
20:     unsigned char dirty;       /* 1 = delayed write */
21:     char *data;                /* block contents */
22:     struct buffer *prev_hash;
23:     struct buffer *next_hash;
24:     struct buffer *prev_free;
25:     struct buffer *next_free;
26: };
27: extern struct buffer *buffer_table;
28: extern struct buffer **buffer_hash_table;
29:
30: /* values to be determined during system startup */
31: extern unsigned int buffer_table_size;    /* size in bytes */
32: extern unsigned int buffer_hash_table_size; /* size in bytes */
33:
34: struct buffer * get_dirty_buffer(__dev_t, __blk_t, int);
35: struct buffer * bread(__dev_t, __blk_t, int);
36: void bwrite(struct buffer *);
37: void brelse(struct buffer *);
38: void sync_buffers(__dev_t);
39: void invalidate_buffers(__dev_t);
40: int reclaim_buffers(void);
41: void buffer_init(void);
42:
43: #endif /* _FIWIX_BUFFER_H */

```



## include/fiwix/cmos.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/cmos.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_CMOS_H
9: #define _FIWIX_CMOS_H
10:
11: #define CMOS_INDEX      0x70
12: #define CMOS_DATA       0x71
13:
14: #define CMOS_STATB_IRQF 0x0F    /* periodic interrupt frequency */
15: #define CMOS_STATB_UIP  0x80    /* time update in progress */
16:
17: #define CMOS_STATB_DSE  0x01    /* enable daylight savings */
18: #define CMOS_STATB_24H  0x02    /* 24-hour mode (0=12h, 1=24h) */
19: #define CMOS_STATB_DM   0x04    /* time/date in binary mode (0=BCD, 1=binary) */
20: #define CMOS_STATB_SQWE 0x08    /* enable square wave frequency */
21: #define CMOS_STATB_UIE  0x10    /* enable update-ended interrupt */
22: #define CMOS_STATB_AIE  0x20    /* enable alarm interrupt */
23: #define CMOS_STATB_PIE  0x40    /* enable periodic interrupt */
24: #define CMOS_STATB_SET  0x80    /* abort clock update */
25:
26: #define CMOS_STATD_VRT  0x80    /* valid RAM and time */
27:
28: /* CMOS RAM data registers */
29: #define CMOS_SEC        0x00    /* second */
30: #define CMOS_ASEC      0x01    /* alarm second */
31: #define CMOS_MIN       0x02    /* minute */
32: #define CMOS_AMIN     0x03    /* alarm minute */
33: #define CMOS_HOUR      0x04    /* hour */
34: #define CMOS_AHOUR    0x05    /* alarm hour */
35: #define CMOS_DOW       0x06    /* day of week */
36: #define CMOS_DAY       0x07    /* day */
37: #define CMOS_MONTH     0x08    /* month */
38: #define CMOS_YEAR      0x09    /* last two digits of year */
39: #define CMOS_STATA     0x0A    /* status register A */
40: #define CMOS_STATB     0x0B    /* status register B */
41: #define CMOS_STATC     0x0C    /* status register C */
42: #define CMOS_STATD     0x0D    /* status register D */
43: #define CMOS_DIAG     0x0E    /* diagnostics status */
44: #define CMOS_FDDTYPE   0x10    /* floppy disk drive type */
45: #define CMOS_HDDTYPE   0x12    /* hard disk drive type */
46: #define CMOS_CENTURY   0x32    /* century */
47:
48: /* conversions */
49: #define BCD2BIN(bcd)    (((bcd) >> 4) * 10) + ((bcd) & 0x0F)
50: #define BIN2BCD(bin)    ((bin) % 10) | (((bin) / 10) << 4)
51:
52: int cmos_update_in_progress(void);
53: unsigned char cmos_read_date(unsigned char);
54: void cmos_write_date(unsigned char, unsigned char);
55: unsigned char cmos_read(unsigned char);
56: void cmos_write(unsigned char, unsigned char);
57:
58: #endif /* _FIWIX_CMOS_H */

```

**include/fiwix/config.h**

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/config.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_CONFIG_H
9: #define _FIWIX_CONFIG_H
10:
11: /* maximum number of processes */
12: #define NR_PROCS          64
13:
14: /* maximum number of callout functions (timer) */
15: #define NR_CALLOUTS      NR_PROCS
16:
17: /* maximum number of mounted filesystems */
18: #define NR_MOUNT_POINTS  8
19:
20: /* maximum number of opened files in system */
21: #define NR_OPENS         1024
22:
23: /* maximum number of flocks in system */
24: #define NR_FLOCKS        (NR_PROCS * 5)
25:
26:
27:
28: /* percentage of memory that buffer cache will borrow from available memory */
29: #define BUFFER_PERCENTAGE 100
30:
31: /* percentage of hash buckets relative to the size of the buffer table */
32: #define BUFFER_HASH_PERCENTAGE 10
33:
34: /* buffers reclaimed in a single call */
35: #define NR_BUF_RECLAIM   150
36:
37:
38: /* percentage of memory assigned to the inode table and hash table */
39: #define INODE_PERCENTAGE 5
40: #define INODE_HASH_PERCENTAGE 10
41:
42: /* percentage of memory assigned to the page hash table */
43: #define PAGE_HASH_PERCENTAGE 10
44:
45:
46: /* maximum value for PID */
47: #define MAX_PID_VALUE    32767
48:
49: /* number of screens in console' scroll back */
50: #define SCREENS_LOG      6
51:
52: /* maximum number of messages on spurious interrupts */
53: #define MAX_SPU_NOTICES  10
54:
55: /* maximum number of virtual memory mappings (regions) */
56: #define VMA_REGIONS      150
57:
58: #endif /* _FIWIX_CONFIG_H */

```

## include/fiwix/console.h

Page 1/3

```

1: /*
2:  * fiwix/include/fiwix/console.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_CONSOLE_H
9: #define _FIWIX_CONSOLE_H
10:
11: #include <fiwix/config.h>
12: #include <fiwix/termios.h>
13: #include <fiwix/vt.h>
14:
15: #define NR_VCONSOLES          12      /* number of virtual consoles */
16:
17: #define VCONSOLES_MAJOR      4        /* virtual consoles major number */
18: #define SYSCON_MAJOR        5        /* system console major number */
19:
20: #define MONO_ADDR            0xB0000L
21: #define COLOR_ADDR          0xB8000L
22:
23: #define MONO_6845_ADDR       0x3B4    /* i/o address (+1 for data register) */
24: #define COLOR_6845_ADDR     0x3D4    /* i/o address (+1 for data register) */
25:
26: #define ATTR_CONTROLLER      0x3C0    /* attribute controller register */
27: #define ATTR_CONTROLLER_PAS  0x20     /* palette address source */
28: #define INPUT_STAT1         0x3DA    /* input status #1 register */
29: #define BLANK_INTERVAL      (600 * HZ) /* 600 seconds (10 minutes) */
30:
31: #define CRT_INDEX            0
32: #define CRT_DATA             1
33: #define CRT_CURSOR_STR      0xA
34: #define CRT_CURSOR_END      0xB
35: #define CRT_START_ADDR_HI   0xC
36: #define CRT_START_ADDR_LO   0xD
37: #define CRT_CURSOR_POS_HI   0xE
38: #define CRT_CURSOR_POS_LO   0xF
39:
40: #define CURSOR_MASK          0x1F
41: #define CURSOR_DISABLE      0x20
42:
43: /* Graphic Rendition Combination Modes */
44: #define SGR_DEFAULT          0        /* back to the default rendition */
45: #define SGR_BOLD             1        /* set bold */
46: #define SGR_BLINK            5        /* set slowly blinking */
47: #define SGR_REVERSE          7        /* set reverse video */
48: #define SGR_BOLD_OFF        21       /* unset bold */
49: #define SGR_BLINK_OFF       25       /* unset blinking */
50: #define SGR_REVERSE_OFF     27       /* unset reverse video */
51:
52: #define SGR_BLACK_FG         30       /* set black foreground */
53: #define SGR_RED_FG           31       /* set red foreground */
54: #define SGR_GREEN_FG        32       /* set green foreground */
55: #define SGR_BROWN_FG        33       /* set brown foreground */
56: #define SGR_BLUE_FG         34       /* set blue foreground */
57: #define SGR_MAGENTA_FG      35       /* set magenta foreground */
58: #define SGR_CYAN_FG         36       /* set cyan foreground */
59: #define SGR_WHITE_FG        37       /* set white foreground */
60:
61: #define SGR_DEFAULT_FG_U_ON  38       /* set def. fg color (underline on) */
62: #define SGR_DEFAULT_FG_U_OFF 39       /* set def. fg color (underline off) */
63:
64: #define SGR_BLACK_BG         40       /* set black background */
65: #define SGR_RED_BG           41       /* set red background */
66: #define SGR_GREEN_BG        42       /* set green background */
67: #define SGR_BROWN_BG        43       /* set brown background */

```

## include/fiwix/console.h

Page 2/3

```

68: #define SGR_BLUE_BG          44      /* set blue background */
69: #define SGR_MAGENTA_BG      45      /* set magenta background */
70: #define SGR_CYAN_BG         46      /* set cyan background */
71: #define SGR_WHITE_BG        47      /* set white background */
72:
73: #define SGR_DEFAULT_BG      49      /* set default background color */
74:
75: #define NPARMS               16
76:
77: #define COLOR_BLACK          0x0000
78: #define COLOR_BLUE           0x0100
79: #define COLOR_GREEN          0x0200
80: #define COLOR_CYAN           0x0300
81: #define COLOR_RED            0x0400
82: #define COLOR_MAGENTA        0x0500
83: #define COLOR_BROWN         0x0600
84: #define COLOR_WHITE          0x0700
85: #define BG_BLACK             0x0000
86: #define BG_BLUE              0x1000
87: #define BG_GREEN             0x2000
88: #define BG_CYAN              0x3000
89: #define BG_RED                0x4000
90: #define BG_MAGENTA           0x5000
91: #define BG_BROWN             0x6000
92: #define BG_WHITE             0x7000
93:
94: #define DEF_MODE              (COLOR_WHITE | BG_BLACK)
95: #define BLANK_MEM             (DEF_MODE | ' ')
96:
97: #define SCROLL_UP             1
98: #define SCROLL_DOWN          2
99:
100: #define VC_BUF_UP             1
101: #define VC_BUF_DOWN          2
102:
103: #define BS                    127    /* backspace */
104:
105: extern short int current_cons; /* current console (/dev/tty1 ... /dev/tty12) */
106:
107: struct video_parms {
108:     short int *address;
109:     int port;
110:     char *type;
111:     int columns;
112:     int lines;
113: };
114:
115: struct vconsole {
116:     int x; /* current column */
117:     int y; /* current line */
118:     int top, bottom, columns;
119:     short int check_x;
120:     unsigned char led_status;
121:     unsigned char scrlock, numlock, capslock;
122:     unsigned char esc, sbracket, semicolon, question;
123:     int parmv1, parmv2;
124:     int nparms, parms[NPARMS];
125:     unsigned short int color_attr;
126:     unsigned char bold, underline, blink, reverse;
127:     int insert_mode;
128:     short int *vidmem;
129:     short int has_focus;
130:     short int *scrbuf;
131:     int saved_x;
132:     int saved_y;
133:     struct vt_mode vt_mode;
134:     unsigned char vc_mode;

```

**include/fiwix/console.h**

Page 3/3

```
135:         unsigned char blanked;
136:         int switchto_tty;
137:         struct tty *tty;
138: };
139:
140: void vconsole_reset(struct tty *);
141: void vconsole_write(struct tty *);
142: void vconsole_select(int);
143: void vconsole_select_final(int);
144: void vconsole_save(struct vconsole *);
145: void vconsole_restore(struct vconsole *);
146: void vconsole_buffer_scrl(int);
147: void blank_screen(struct vconsole *);
148: void unblank_screen(struct vconsole *);
149: void screen_on(void);
150: void screen_off(unsigned int);
151: void vconsole_start(struct tty *);
152: void vconsole_stop(struct tty *);
153: void vconsole_beep(void);
154: void vconsole_deltab(struct tty *);
155: void console_flush_log_buf(char *, unsigned int);
156: void get_video_parms(void);
157: void vconsole_init(void);
158:
159: #endif /* _FIWIX_CONSOLE_H */
```

**include/fiwix/const.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/const.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_CONST_H
9: #define _FIWIX_CONST_H
10:
11: #define KERNEL_BASE_ADDR      0xC0000000
12: #define KERNEL_ENTRY_ADDR    0x100000
13:
14: #define KERNEL_CS              0x08      /* kernel code segment */
15: #define KERNEL_DS              0x10      /* kernel data segment */
16: #define USER_CS                0x18      /* user code segment */
17: #define USER_DS                0x20      /* user data segment */
18: #define TSS                    0x28      /* TSS segment */
19:
20: #endif /* _FIWIX_CONST_H */
```

## include/fiwix/cpu.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/cpu.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_CPU_H
9: #define _FIWIX_CPU_H
10:
11: #define CPU_FPU          0x00000001    /* Floating-Point Unit on chip */
12: #define CPU_VME          0x00000002    /* Virtual 8086 Mode Enhancements */
13: #define CPU_DE           0x00000004    /* Debugging Extensions */
14: #define CPU_PSE          0x00000008    /* Page Size Extension */
15: #define CPU_TSC          0x00000010    /* Time Stamp Counter */
16: #define CPU_MSR          0x00000020    /* Model Specific Registers */
17: #define CPU_PAE          0x00000040    /* Physical Address Extension */
18: #define CPU_MCE          0x00000080    /* Machine Check Exception */
19: #define CPU_CX8          0x00000100    /* CMPXCHG8B instruction supported */
20: #define CPU_APIC         0x00000200    /* On-chip APIC hardware supported */
21: #define CPU_RES10       0x00000400    /* Reserved */
22: #define CPU_SEP          0x00000800    /* Fast System Call */
23: #define CPU_MTRR         0x00001000    /* Memory Type Range Registers */
24: #define CPU_PGE          0x00002000    /* Page Global Enable */
25: #define CPU_MCA          0x00004000    /* Machine Check Architecture */
26: #define CPU_CMOV         0x00008000    /* Conditional Move Instruction */
27: #define CPU_PAT          0x00010000    /* Page Attribute Table */
28: #define CPU_PSE36        0x00020000    /* 36-bit Page Size Extension */
29: #define CPU_PSN          0x00040000    /* Processor Serial Number */
30: #define CPU_CLFSH        0x00080000    /* CLFLUSH instruction supported */
31: #define CPU_RES20       0x00100000    /* Reserved */
32: #define CPU_DS           0x00200000    /* Debug Store */
33: #define CPU_ACPI         0x00400000    /* Thermal Monitor and others */
34: #define CPU_MMX          0x00800000    /* Intel Architecture MMX Technology */
35: #define CPU_FXSR         0x01000000    /* Fast Floating Point Save and Rest. */
36: #define CPU_SSE          0x02000000    /* Streaming SIMD Extensions */
37: #define CPU_SSE2         0x04000000    /* Streaming SIMD Extensions 2 */
38: #define CPU_SS           0x08000000    /* Self-Snoop */
39: #define CPU_HTT          0x10000000    /* Hyper-Threading Technology */
40: #define CPU_TM           0x20000000    /* Thermal Monitor */
41: #define CPU_RES30       0x40000000    /* Reserved */
42: #define CPU_PBE          0x80000000    /* Pending Break Enable */
43:
44: #define RESERVED_DESC   0x80000000    /* TLB descriptor reserved */
45:
46: struct cpu {
47:     char *vendor_id;
48:     char family;
49:     char model;
50:     char *model_name;
51:     char stepping;
52:     unsigned long int hz;
53:     char *cache;
54:     char has_cpuid;
55:     char has_fpu;
56:     int flags;
57: };
58: struct cpu cpu_table;
59:
60: struct cpu_type {
61:     int cpu;
62:     char *name[20];
63: };
64:
65: int get_cpu_flags(char *, int);
66: void cpu_init(void);
67:

```

**include/fiwix/cpu.h**

Page 2/2

```
68: #endif /* _FIWIX_CPU_H */
```



**include/fiwix/ctype.h**

Page 1/1

```

1:  /*
2:  *  fiwix/include/fiwix/ctype.h
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #ifndef  _FIWIX_CTYPE_H
9:  #define  _FIWIX_CTYPE_H
10:
11: #define  _U      0x01    /* upper case */
12: #define  _L      0x02    /* lower case */
13: #define  _N      0x04    /* numeral (digit) */
14: #define  _S      0x08    /* spacing character */
15: #define  _P      0x10    /* punctuation */
16: #define  _C      0x20    /* control character */
17: #define  _X      0x40    /* hexadecimal */
18: #define  _B      0x80    /* blank */
19:
20: extern unsigned char  _ctype[];
21:
22: #define  ISALPHA(ch)      ((_ctype + 1)[ch] & (_U | _L))
23: #define  ISUPPER(ch)     ((_ctype + 1)[ch] & _U)
24: #define  ISLOWER(ch)     ((_ctype + 1)[ch] & _L)
25: #define  ISDIGIT(ch)     ((_ctype + 1)[ch] & _N)
26: #define  ISALNUM(ch)     ((_ctype + 1)[ch] & (_U | _L | _N))
27: #define  ISSPACE(ch)     ((_ctype + 1)[ch] & _S)
28: #define  ISPUNCT(ch)     ((_ctype + 1)[ch] & _P)
29: #define  ISCNTRL(ch)     ((_ctype + 1)[ch] & _C)
30: #define  ISXDIGIT(ch)    ((_ctype + 1)[ch] & (_N | _X))
31:
32: #define  ISASCII(ch)     ((unsigned) ch <= 0x7F)
33: #define  TOASCII(ch)     ((unsigned) ch & 0x7F)
34:
35: #define  TOUPPER(ch)     ((ch) & ~32)
36: #define  TOLOWER(ch)    ((ch) | 32)
37:
38: #endif  /* _FIWIX_CTYPE_H */

```

**include/fiwix/devices.h**

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/devices.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_DEVICES_H
9: #define _FIWIX_DEVICES_H
10:
11: #include <fiwix/types.h>
12: #include <fiwix/fs.h>
13:
14: #define NR_BLKDEV          128      /* maximum number of block devices */
15: #define NR_CHRDEV          128      /* maximum number of char devices */
16:
17: #define BLK_DEV            1        /* block device */
18: #define CHR_DEV            2        /* character device */
19:
20: #define SET_MINOR(minors, bit)    ((minors[(bit) / 32]) |= (1 << ((bit) % 32)))
21: #define CLEAR_MINOR(minors, bit) ((minors[(bit) / 32]) &= ~(1 << ((bit) % 32)))
22: #define TEST_MINOR(minors, bit)  ((minors[(bit) / 32]) & (1 << ((bit) % 32)))
23:
24: struct device {
25:     char *name;
26:     unsigned char major;
27:     unsigned int minors[8];        /* bitmap of 256 bits */
28:     int blksize;
29:     void *device_data;            /* mostly used for minor sizes, in KB */
30:     struct fs_operations *fsop;
31:     struct device *next;
32: };
33:
34: extern struct device *chr_device_table[NR_CHRDEV];
35: extern struct device *blk_device_table[NR_BLKDEV];
36:
37: int register_device(int, struct device *);
38: struct device * get_device(int, __dev_t);
39: int chr_dev_open(struct inode *, struct fd *);
40: int blk_dev_open(struct inode *, struct fd *);
41: int blk_dev_close(struct inode *, struct fd *);
42: int blk_dev_read(struct inode *, struct fd *, char *, __size_t);
43: int blk_dev_write(struct inode *, struct fd *, const char *, __size_t);
44: int blk_dev_ioctl(struct inode *, int, unsigned long int);
45: int blk_dev_lseek(struct inode *, __off_t);
46:
47: void dev_init(void);
48:
49: #endif /* _FIWIX_DEVICES_H */

```

**include/fiwix/dirent.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/dirent.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_DIRENT_H
9: #define _FIWIX_DIRENT_H
10:
11: #include <fiwix/types.h>
12: #include <fiwix/limits.h>
13:
14: struct dirent {
15:     __ino_t d_ino;           /* inode number */
16:     __off_t d_off;         /* offset to next dirent */
17:     unsigned short int d_reclen; /* length of this dirent */
18:     char d_name[NAME_MAX + 1]; /* file name (null-terminated) */
19: };
20:
21: #endif /* _FIWIX_DIRENT_H */
```

**include/fiwix/dma.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/dma.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_DMA_H
9: #define _FIWIX_DMA_H
10:
11: #define DMA_CHANNELS          8          /* max. number of DMA channels */
12:
13: #define DMA_MASK_CHANNEL     0x04
14: #define DMA_UNMASK_CHANNEL   0x00
15:
16: #define DMA_MODE_VERIFY      0x00
17: #define DMA_MODE_WRITE       0x04      /* read device -> write memory */
18: #define DMA_MODE_READ        0x08      /* read memory -> write device */
19: #define DMA_MODE_AUTOINIT    0x10
20: #define DMA_MODE_ADDRES_DEC  0x20
21: #define DMA_MODE_DEMAND      0x00
22: #define DMA_MODE_SINGLE      0x40
23: #define DMA_MODE_BLOCK       0x80
24: #define DMA_MODE_CASCADE     0xC0
25:
26: char *dma_resources[DMA_CHANNELS];
27:
28: void start_dma(int, void *, unsigned int, int);
29: int dma_register(int, char *);
30: int dma_unregister(int);
31: void dma_init(void);
32:
33: #endif /* _FIWIX_DMA_H */
```

## include/fiwix/errno.h

Page 1/3

```

1: #ifndef _FIWIX_ERRNO_H
2: #define _FIWIX_ERRNO_H
3:
4: #define EPERM          1      /* Operation not permitted - Not owner */
5: #define ENOENT         2      /* No such file or directory */
6: #define ESRCH         3      /* No such process */
7: #define EINTR         4      /* Interrupted system call */
8: #define EIO           5      /* I/O error */
9: #define ENXIO         6      /* No such device or address */
10: #define E2BIG         7      /* Arg list too long */
11: #define ENOEXEC       8      /* Exec format error */
12: #define EBADF         9      /* Bad file number */
13: #define ECHILD        10     /* No child processes */
14: #define EAGAIN        11     /* Try again - No more processes */
15: #define ENOMEM        12     /* Out of memory - No enough space */
16: #define EACCES        13     /* Permission denied */
17: #define EFAULT        14     /* Bad address */
18: #define ENOTBLK       15     /* Block device required */
19: #define EBUSY         16     /* Device or resource busy */
20: #define EEXIST        17     /* File exists */
21: #define EXDEV         18     /* Cross-device link */
22: #define ENODEV        19     /* No such device */
23: #define ENOTDIR       20     /* Not a directory */
24: #define EISDIR        21     /* Is a directory */
25: #define EINVAL        22     /* Invalid argument */
26: #define ENFILE        23     /* File table overflow */
27: #define EMFILE        24     /* Too many open files */
28: #define ENOTTY        25     /* Not a typewriter */
29: #define ETXTBSY       26     /* Text file busy */
30: #define EFBIG         27     /* File too large */
31: #define ENOSPC        28     /* No space left on device */
32: #define ESPIPE        29     /* Illegal seek */
33: #define EROFS         30     /* Read-only file system */
34: #define EMLINK        31     /* Too many links */
35: #define EPIPE        32     /* Broken pipe */
36: #define EDOM          33     /* Math argument out of domain of func */
37: #define ERANGE        34     /* Math result not representable */
38: #define EDEADLK       35     /* Resource deadlock would occur */
39: #define ENAMETOOLONG  36     /* File name too long */
40: #define ENOLCK        37     /* No record locks available */
41: #define ENOSYS        38     /* Function not implemented */
42: #define ENOTEMPTY     39     /* Directory not empty */
43: #define ELOOP         40     /* Too many symbolic links encountered */
44: #define EWOULDBLOCK   EAGAIN /* Operation would block */
45: #define ENOMSG        42     /* No message of desired type */
46: #define EIDRM         43     /* Identifier removed */
47: #define ECHRNG        44     /* Channel number out of range */
48: #define EL2NSYNC      45     /* Level 2 not synchronized */
49: #define EL3HLT        46     /* Level 3 halted */
50: #define EL3RST        47     /* Level 3 reset */
51: #define ELNRNG        48     /* Link number out of range */
52: #define EUNATCH       49     /* Protocol driver not attached */
53: #define ENOCCSI       50     /* No CSI structure available */
54: #define EL2HLT        51     /* Level 2 halted */
55: #define EBADE         52     /* Invalid exchange */
56: #define EBADR         53     /* Invalid request descriptor */
57: #define EXFULL        54     /* Exchange full */
58: #define ENOANO        55     /* No anode */
59: #define EBADRQC       56     /* Invalid request code */
60: #define EBADSLT       57     /* Invalid slot */
61:
62: #define EDEADLOCK     EDEADLK
63:
64: #define EBFONT        59     /* Bad font file format */
65: #define ENOSTR        60     /* Device not a stream */
66: #define ENODATA       61     /* No data available */
67: #define ETIME         62     /* Timer expired */

```

## include/fiwix/errno.h

Page 2/3

```

68: #define ENOSR          63      /* Out of streams resources */
69: #define ENONET         64      /* Machine is not on the network */
70: #define ENOPKG         65      /* Package not installed */
71: #define EREMOTE        66      /* Object is remote */
72: #define ENOLINK        67      /* Link has been severed */
73: #define EADV           68      /* Advertise error */
74: #define ESRMNT         69      /* Srmount error */
75: #define ECOMM          70      /* Communication error on send */
76: #define EPROTO         71      /* Protocol error */
77: #define EMULTIHOP       72      /* Multihop attempted */
78: #define EDOTDOT         73      /* RFS specific error */
79: #define EBADMSG        74      /* Not a data message */
80: #define EOVERFLOW      75      /* Value too large for defined data type */
81: #define ENOTUNIQ       76      /* Name not unique on network */
82: #define EBADF          77      /* File descriptor in bad state */
83: #define EREMCHG        78      /* Remote address changed */
84: #define ELIBACC        79      /* Can not access a needed shared library */
85: #define ELIBBAD        80      /* Accessing a corrupted shared library */
86: #define ELIBSCN        81      /* .lib section in a.out corrupted */
87: #define ELIBMAX        82      /* Attempting to link in too many shared librari
es */
88: #define ELIBEXEC       83      /* Cannot exec a shared library directly */
89: #define EILSEQ         84      /* Illegal byte sequence */
90: #define ERESTART       85      /* Interrupted system call should be restarted */
/
91: #define ESTRPIPE       86      /* Streams pipe error */
92: #define EUSERS         87      /* Too many users */
93: #define ENOTSOCK       88      /* Socket operation on non-socket */
94: #define EDESTADDRREQ   89      /* Destination address required */
95: #define EMSGSIZE       90      /* Message too long */
96: #define EPROTOTYPE     91      /* Protocol wrong type for socket */
97: #define ENOPROTOOPT    92      /* Protocol not available */
98: #define EPROTONOSUPPORT 93      /* Protocol not supported */
99: #define ESOCKTNOSUPPORT 94      /* Socket type not supported */
100: #define EOPNOTSUPP     95      /* Operation not supported on transport endpoint */
*/
101: #define EPFNOSUPPORT   96      /* Protocol family not supported */
102: #define EAFNOSUPPORT   97      /* Address family not supported by protocol */
103: #define EADDRINUSE     98      /* Address already in use */
104: #define EADDRNOTAVAIL  99      /* Cannot assign requested address */
105: #define ENETDOWN       100     /* Network is down */
106: #define ENETUNREACH    101     /* Network is unreachable */
107: #define ENETRESET      102     /* Network dropped connection because of reset */
/
108: #define ECONNABORTED   103     /* Software caused connection abort */
109: #define ECONNRESET     104     /* Connection reset by peer */
110: #define ENOBUFS        105     /* No buffer space available */
111: #define EISCONN        106     /* Transport endpoint is already connected */
112: #define ENOTCONN       107     /* Transport endpoint is not connected */
113: #define ESHUTDOWN      108     /* Cannot send after transport endpoint shutdown */
*/
114: #define ETOOMANYREFS   109     /* Too many references: cannot splice */
115: #define ETIMEDOUT      110     /* Connection timed out */
116: #define ECONNREFUSED   111     /* Connection refused */
117: #define EHOSTDOWN      112     /* Host is down */
118: #define EHOSTUNREACH   113     /* No route to host */
119: #define EALREADY       114     /* Operation already in progress */
120: #define EINPROGRESS    115     /* Operation now in progress */
121: #define ESTALE         116     /* Stale NFS file handle */
122: #define EUCLEAN        117     /* Structure needs cleaning */
123: #define ENOTNAM        118     /* Not a XENIX named type file */
124: #define ENAVAIL        119     /* No XENIX semaphores available */
125: #define EISNAM         120     /* Is a named type file */
126: #define EREMOTEIO     121     /* Remote I/O error */
127: #define EDQUOT         122     /* Quota exceeded */
128:
129: #define ENOMEDIUM     123     /* No medium found */

```

**include/fiwix/errno.h**

Page 3/3

```
130: #define EMEDIUMTYPE      124      /* Wrong medium type          */
131:
132: #endif /* _FIWIX_ERRNO_H */
```

## include/fiwix/fcntl.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/fcntl.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_FCNTL_H
9: #define _FIWIX_FCNTL_H
10:
11: #include <fiwix/types.h>
12:
13: #define O_ACCMODE          0003
14: #define O_RDONLY          00
15: #define O_WRONLY          01
16: #define O_RDWR           02
17:
18: /* for open() only */
19: #define O_CREAT            0100 /* create file if it does not exist */
20: #define O_EXCL             0200 /* exclusive use flag */
21: #define O_NOCTTY          0400 /* do not assign controlling terminal */
22: #define O_TRUNC           01000 /* truncate flag */
23: #define O_NOFOLLOW        0400000 /* do not follow symbolic links */
24:
25: #define O_APPEND           02000
26: #define O_NONBLOCK        04000
27: #define O_NDELAY           O_NONBLOCK
28: #define O_SYNC            010000
29:
30: #define F_DUPFD            0      /* duplicate file descriptor */
31: #define F_GETFD            1      /* get file descriptor flags */
32: #define F_SETFD            2      /* set file descriptor flags */
33: #define F_GETFL            3      /* get status flags and file access modes */
34: #define F_SETFL            4      /* set file status flags */
35: #define F_GETLK            5      /* get record locking information */
36: #define F_SETLK            6      /* set record locking information */
37: #define F_SETLKW           7      /* same as F_SETLK; wait if blocked */
38:
39: /* get/set process or process group ID to receive SIGURG signals */
40: #define F_SETOWN           8      /* for sockets only */
41: #define F_GETOWN           9      /* for sockets only */
42:
43: /* for F_[GET|SET]FL */
44: #define FD_CLOEXEC         1      /* close the file descriptor upon exec() */
45:
46: /* for POSIX fcntl() */
47: #define F_RDLCK            0      /* shared or read lock */
48: #define F_WRLCK            1      /* exclusive or write lock */
49: #define F_UNLCK            2      /* unlock */
50:
51: /* for BSD flock() */
52: #define LOCK_SH             1      /* shared lock */
53: #define LOCK_EX             2      /* exclusive lock */
54: #define LOCK_NB             4      /* or'd with one of the above to prevent
55:                                     blocking */
56: #define LOCK_UN            8      /* unlock */
57:
58: /* IEEE Std 1003.1, 2004 Edition */
59: struct flock {
60:     short int l_type;          /* type of lock: F_RDLCK, F_WRLCK, F_UNLCK */
61:     short int l_whence;       /* flag for 'l_start': SEEK_SET, SEEK_CUR, ... */
62:     __off_t l_start;          /* relative offset in bytes */
63:     __off_t l_len;            /* size; if 0 then until EOF */
64:     __pid_t l_pid;            /* PID holding the lock; returned in F_GETLK */
65: };
66:
67: #endif /* _FIWIX_FCNTL_H */

```



## include/fiwix/filesystems.h

Page 1/3

```

1: /*
2:  * fiwix/include/fiwix/filesystems.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_FILESYSTEMS_H
9: #define _FIWIX_FILESYSTEMS_H
10:
11: #include <fiwix/types.h>
12: #include <fiwix/limits.h>
13:
14: #define NR_FILESYSTEMS          5          /* supported filesystems */
15:
16: /* value to be determined during system startup */
17: extern unsigned int mount_table_size; /* size in bytes */
18:
19: struct filesystems {
20:     const char *name;          /* filesystem name */
21:     struct fs_operations *fsop; /* filesystem operations */
22:     struct mount *mt;          /* mount-table entry (only for nodev) */
23: };
24: struct filesystems filesystems_table[NR_FILESYSTEMS];
25:
26: struct mount {
27:     __dev_t dev;                /* device number */
28:     char devname[DEVNAME_MAX + 1]; /* device name */
29:     char dirname[NAME_MAX + 1]; /* mount point directory name */
30:     unsigned char used;          /* 1=busy, 0=free */
31:     struct superblock sb;        /* superblock */
32:     struct filesystems *fs;      /* pointer to filesystem structure */
33: };
34: extern struct mount *mount_table;
35:
36: int register_filesystem(const char *, struct fs_operations *);
37: struct filesystems * get_filesystem(const char *);
38: void fs_init(void);
39:
40: struct superblock * get_superblock(__dev_t);
41: void sync_superblocks(__dev_t);
42: int kern_mount(__dev_t, struct filesystems *);
43: int mount_root(void);
44: void mount_init(void);
45:
46:
47: /* minix prototypes */
48: int minix_file_open(struct inode *, struct fd *);
49: int minix_file_close(struct inode *, struct fd *);
50: int minix_file_write(struct inode *, struct fd *, const char *, __size_t);
51: int minix_file_lseek(struct inode *, __off_t);
52: int minix_dir_open(struct inode *, struct fd *);
53: int minix_dir_close(struct inode *, struct fd *);
54: int minix_dir_read(struct inode *, struct fd *, char *, __size_t);
55: int minix_dir_write(struct inode *, struct fd *, const char *, __size_t);
56: int minix_dir_readdir(struct inode *, struct fd *, struct dirent *, unsigned int
);
57: int minix_readlink(struct inode *, char *, __size_t);
58: int minix_followlink(struct inode *, struct inode *, struct inode **);
59: int minix_bmap(struct inode *, __off_t, int);
60: int minix_lookup(const char *, struct inode *, struct inode **);
61: int minix_rmdir(struct inode *, struct inode *);
62: int minix_link(struct inode *, struct inode *, char *);
63: int minix_unlink(struct inode *, struct inode *, char *);
64: int minix_symlink(struct inode *, char *, char *);
65: int minix_mkdir(struct inode *, char *, __mode_t);
66: int minix_mknod(struct inode *, char *, __mode_t, __dev_t);

```

## include/fiwix/filesystems.h

Page 2/3

```

67: int minix_truncate(struct inode *, __off_t);
68: int minix_create(struct inode *, char *, __mode_t, struct inode **);
69: int minix_rename(struct inode *, struct inode *, struct inode *, struct inode *,
char *, char *);
70: int minix_read_inode(struct inode *);
71: int minix_write_inode(struct inode *);
72: int minix_ialloc(struct inode *, int);
73: void minix_ifree(struct inode *);
74: void minix_statfs(struct superblock *, struct statfs *);
75: int minix_read_superblock(__dev_t, struct superblock *);
76: int minix_remount_fs(struct superblock *, int);
77: int minix_write_superblock(struct superblock *);
78: void minix_release_superblock(struct superblock *);
79: int minix_init(void);
80:
81:
82: /* ext2 prototypes */
83: int ext2_file_open(struct inode *, struct fd *);
84: int ext2_file_close(struct inode *, struct fd *);
85: int ext2_file_write(struct inode *, struct fd *, const char *, __size_t);
86: int ext2_file_lseek(struct inode *, __off_t);
87: int ext2_dir_open(struct inode *, struct fd *);
88: int ext2_dir_close(struct inode *, struct fd *);
89: int ext2_dir_read(struct inode *, struct fd *, char *, __size_t);
90: int ext2_dir_write(struct inode *, struct fd *, const char *, __size_t);
91: int ext2_dir_readdir(struct inode *, struct fd *, struct dirent *, unsigned int)
;
92: int ext2_readlink(struct inode *, char *, __size_t);
93: int ext2_followlink(struct inode *, struct inode *, struct inode **);
94: int ext2_bmap(struct inode *, __off_t, int);
95: int ext2_lookup(const char *, struct inode *, struct inode **);
96: int ext2_rmdir(struct inode *, struct inode *);
97: int ext2_link(struct inode *, struct inode *, char *);
98: int ext2_unlink(struct inode *, struct inode *, char *);
99: int ext2_symlink(struct inode *, char *, char *);
100: int ext2_mkdir(struct inode *, char *, __mode_t);
101: int ext2_mknod(struct inode *, char *, __mode_t, __dev_t);
102: int ext2_truncate(struct inode *, __off_t);
103: int ext2_create(struct inode *, char *, __mode_t, struct inode **);
104: int ext2_rename(struct inode *, struct inode *, struct inode *, struct inode *,
char *, char *);
105: int ext2_read_inode(struct inode *);
106: int ext2_write_inode(struct inode *);
107: int ext2_ialloc(struct inode *, int);
108: void ext2_ifree(struct inode *);
109: void ext2_statfs(struct superblock *, struct statfs *);
110: int ext2_read_superblock(__dev_t, struct superblock *);
111: int ext2_remount_fs(struct superblock *, int);
112: int ext2_write_superblock(struct superblock *);
113: void ext2_release_superblock(struct superblock *);
114: int ext2_init(void);
115:
116:
117: /* pipefs prototypes */
118: int fifo_open(struct inode *, struct fd *);
119: int pipefs_close(struct inode *, struct fd *);
120: int pipefs_read(struct inode *, struct fd *, char *, __size_t);
121: int pipefs_write(struct inode *, struct fd *, const char *, __size_t);
122: int pipefs_ioctl(struct inode *, int, unsigned long int);
123: int pipefs_lseek(struct inode *, __off_t);
124: int pipefs_select(struct inode *, int);
125: int pipefs_ialloc(struct inode *, int);
126: void pipefs_ifree(struct inode *);
127: int pipefs_read_superblock(__dev_t, struct superblock *);
128: int pipefs_init(void);
129:
130:

```

**include/fiwix/filesystems.h**

Page 3/3

```

131: /* iso9660 prototypes */
132: int iso9660_file_open(struct inode *, struct fd *);
133: int iso9660_file_close(struct inode *, struct fd *);
134: int iso9660_file_lseek(struct inode *, __off_t);
135: int iso9660_dir_open(struct inode *, struct fd *);
136: int iso9660_dir_close(struct inode *, struct fd *);
137: int iso9660_dir_read(struct inode *, struct fd *, char *, __size_t);
138: int iso9660_dir_readdir(struct inode *, struct fd *, struct dirent *, unsigned i
nt);
139: int iso9660_readlink(struct inode *, char *, __size_t);
140: int iso9660_followlink(struct inode *, struct inode *, struct inode **);
141: int iso9660_bmap(struct inode *, __off_t, int);
142: int iso9660_lookup(const char *, struct inode *, struct inode **);
143: int iso9660_read_inode(struct inode *);
144: void iso9660_statfs(struct superblock *, struct statfs *);
145: int iso9660_read_superblock(__dev_t, struct superblock *);
146: void iso9660_release_superblock(struct superblock *);
147: int iso9660_init(void);
148:
149:
150: /* procfs prototypes */
151: int procfs_file_open(struct inode *, struct fd *);
152: int procfs_file_close(struct inode *, struct fd *);
153: int procfs_file_read(struct inode *, struct fd *, char *, __size_t);
154: int procfs_file_lseek(struct inode *, __off_t);
155: int procfs_dir_open(struct inode *, struct fd *);
156: int procfs_dir_close(struct inode *, struct fd *);
157: int procfs_dir_read(struct inode *, struct fd *, char *, __size_t);
158: int procfs_dir_readdir(struct inode *, struct fd *, struct dirent *, unsigned in
t);
159: int procfs_readlink(struct inode *, char *, __size_t);
160: int procfs_followlink(struct inode *, struct inode *, struct inode **);
161: int procfs_bmap(struct inode *, __off_t, int);
162: int procfs_lookup(const char *, struct inode *, struct inode **);
163: int procfs_read_inode(struct inode *);
164: void procfs_statfs(struct superblock *, struct statfs *);
165: int procfs_read_superblock(__dev_t, struct superblock *);
166: int procfs_init(void);
167:
168: #endif /* _FIWIX_FILESYSTEMS_H */

```

## include/fiwix/floppy.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/floppy.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_FLOPPY_H
9: #define _FIWIX_FLOPPY_H
10:
11: #include <fiwix/fs.h>
12: #include <fiwix/sigcontext.h>
13:
14: #define FLOPPY_IRQ        6
15: #define FLOPPY_DMA        2        /* DMA channel */
16:
17: #define FDC_MAJOR        2        /* fdd major number */
18:
19: #define FDC_SECTSIZE      512     /* sector size (in bytes) */
20: #define FDC_TR_DEFAULT    0       /* timer reason is IRQ */
21: #define FDC_TR_MOTOR      1       /* timer reason is motor on */
22:
23: #define FDC_SRA           0x3F0   /* Status Register A */
24: #define FDC_SRB           0x3F1   /* Status Register B */
25: #define FDC_DOR           0x3F2   /* Digital Output Register */
26: #define FDC_MSR           0x3F4   /* Main Status Register */
27: #define FDC_DATA          0x3F5   /* command/data register */
28: #define FDC_DIR           0x3F7   /* Digital Input Register */
29: #define FDC_CCR           0x3F7   /* Configuration Control Register */
30:
31: #define FDC_ENABLE        0x04    /* bit #2 FDC enabled (normal op) */
32: #define FDC_DMA_ENABLE    0x08    /* bit #3 DMA enabled */
33: #define FDC_DRIVE0        0x10    /* motor on for the first drive, the rest will
34:  * be calculated by left-shifting this value
35:  * with 'current_fdd'.
36:  */
37:
38: #define FDC_DIO           0x40    /* bit #6 DIO I/O direction */
39: #define FDC_RQM           0x80    /* bit #7 RQM is ready for I/O */
40:
41: #define MAX_FDC_RESULTS   7
42: #define MAX_FDC_ERR       5
43:
44: #define FDC_RESET         0xFF    /* reset indicator */
45: #define FDC_READ          0xE6
46: #define FDC_WRITE         0xC5
47: #define FDC_VERSION       0x10
48: #define FDC_FORMAT_TRK    0x4D
49: #define FDC_RECALIBRATE   0x07
50: #define FDC_SENSEI        0x08
51: #define FDC_SPECIFY       0x03
52: #define FDC_SEEK          0x0F
53: #define FDC_LOCK          0x14
54: #define FDC_PARTID        0x18
55:
56: #define ST0                0x00    /* Status Register 0 */
57: #define ST1                0x01    /* Status Register 1 */
58: #define ST2                0x02    /* Status Register 2 */
59:
60: #define ST0_IC             0xC0    /* bits #7 and #6 interrupt code */
61: #define ST0_SE             0x20    /* bit #5 successful implied seek */
62: #define ST0_RECALIBRATE    ST0_SE /* bit #5 successful FDC_RECALIBRATE */
63: #define ST0_SEEK          ST0_SE /* bit #5 successful FDC_SEEK */
64: #define ST0_UC             0x10    /* bit #4 unit needs check (fault) */
65: #define ST0_NR             0x8     /* bit #3 drive not ready */
66:
67: #define ST1_NW            0x02    /* bit #1 not writable */

```

**include/fiwix/floppy.h**

Page 2/2

```
68:
69: #define ST_PCN          0x01    /* present cylinder */
70: #define ST_CYL         0x03    /* cylinder returned */
71: #define ST_HEAD        0x04    /* head returned */
72: #define ST_SECTOR      0x05    /* sector returned */
73:
74: /* floppy disk drive type */
75: struct fdtd {
76:     short int size;           /* number of sectors */
77:     short int sizekb;        /* size in KB */
78:     char tracks;             /* number of tracks */
79:     char spt;                /* number of sectors per track */
80:     char heads;             /* number of heads */
81:     char gpl1;              /* GAP in READ/WRITE operations */
82:     char gpl2;              /* GAP in FORMAT TRACK operations */
83:     char rate;              /* data rate value */
84:     char spec;              /* SRT+HUT (StepRate + HeadUnload) Time */
85:     char hlt;               /* HLT (Head Load Time) */
86:     char *name;             /* unit name */
87: };
88:
89: void irq_floppy(int, struct sigcontext *);
90: void fdc_timer(unsigned int);
91:
92: int fdc_open(struct inode *, struct fd *);
93: int fdc_close(struct inode *, struct fd *);
94: int fdc_read(__dev_t, __blk_t, char *, int);
95: int fdc_write(__dev_t, __blk_t, char *, int);
96: int fdc_ioctl(struct inode *, int, unsigned long int);
97: int fdc_lseek(struct inode *, __off_t);
98:
99: void floppy_init(void);
100:
101: #endif /* _FIWIX_FLOPPY_H */
```

## include/fiwix/fs\_ext2.h

Page 1/4

```

1: /*
2:  * fiwix/include/fiwix/fs_ext2.h
3:  *
4:  * This file from: Linux 2.0.40
5:  * Copyright (C) 1992, 1993, 1994, 1995
6:  * Remy Card (card@masi.ibp.fr)
7:  * Laboratoire MASI - Institut Blaise Pascal
8:  * Universite Pierre et Marie Curie (Paris VI)
9:  * Copyright (C) 1991, 1992 Linus Torvalds
10: */
11:
12: #ifndef _FIWIX_FS_EXT2_H
13: #define _FIWIX_FS_EXT2_H
14:
15: #include <fiwix/types.h>
16:
17: #define EXT2_ROOT_INO          2          /* Root inode */
18: #define EXT2_SUPER_MAGIC      0xEF53
19:
20: /*
21:  * Macro-instructions used to manage several block sizes
22:  */
23: #define EXT2_MIN_BLOCK_SIZE    1024
24: #define EXT2_MAX_BLOCK_SIZE    4096
25: #define EXT2_MIN_BLOCK_LOG_SIZE 10
26: # define EXT2_BLOCK_SIZE(s)    ((s)->s_blocksize)
27: # define EXT2_BLOCK_SIZE_BITS(s) ((s)->s_blocksize_bits)
28:
29: /*
30:  * Structure of a blocks group descriptor
31:  */
32: struct ext2_group_desc
33: {
34:     __u32    bg_block_bitmap;          /* Blocks bitmap block */
35:     __u32    bg_inode_bitmap;         /* Inodes bitmap block */
36:     __u32    bg_inode_table;         /* Inodes table block */
37:     __u16    bg_free_blocks_count;    /* Free blocks count */
38:     __u16    bg_free_inodes_count;    /* Free inodes count */
39:     __u16    bg_used_dirs_count;     /* Directories count */
40:     __u16    bg_pad;
41:     __u32    bg_reserved[3];
42: };
43:
44: /*
45:  * Macro-instructions used to manage group descriptors
46:  */
47: #define EXT2_BLOCKS_PER_GROUP(s)    ((s)->u.ext2.sb.s_blocks_per_group)
48: #define EXT2_INODES_PER_GROUP(s)    ((s)->u.ext2.sb.s_inodes_per_group)
49: # define EXT2_DESC_PER_BLOCK_BITS(s) ((s)->u.ext2_sb.s_desc_per_block_bits)
50: #define EXT2_DESC_PER_BLOCK(s)      ((s)->u.ext2.desc_per_block)
51:
52: /*
53:  * Constants relative to the data blocks
54:  */
55: #define EXT2_NDIR_BLOCKS          12
56: #define EXT2_IND_BLOCK            EXT2_NDIR_BLOCKS
57: #define EXT2_DIND_BLOCK          (EXT2_IND_BLOCK + 1)
58: #define EXT2_TIND_BLOCK          (EXT2_DIND_BLOCK + 1)
59: #define EXT2_N_BLOCKS            (EXT2_TIND_BLOCK + 1)
60:
61: /*
62:  * Structure of an inode on the disk
63:  */
64: struct ext2_inode {
65:     __u16    i_mode;                 /* File mode */
66:     __u16    i_uid;                  /* Low 16 bits of Owner Uid */
67:     __u32    i_size;                 /* Size in bytes */

```

## include/fiwix/fs\_ext2.h

Page 2/4

```

68:     __u32    i_atime;        /* Access time */
69:     __u32    i_ctime;        /* Creation time */
70:     __u32    i_mtime;        /* Modification time */
71:     __u32    i_dtime;        /* Deletion Time */
72:     __u16    i_gid;          /* Low 16 bits of Group Id */
73:     __u16    i_links_count;  /* Links count */
74:     __u32    i_blocks;       /* Blocks count */
75:     __u32    i_flags;        /* File flags */
76:     union {
77:         struct {
78:             __u32    l_i_reserved1;
79:         } linux1;
80:         struct {
81:             __u32    h_i_translator;
82:         } hurd1;
83:         struct {
84:             __u32    m_i_reserved1;
85:         } masix1;
86:     } osd1;                    /* OS dependent 1 */
87:     __u32    i_block[EXT2_N_BLOCKS]; /* Pointers to blocks */
88:     __u32    i_generation;    /* File version (for NFS) */
89:     __u32    i_file_acl;      /* File ACL */
90:     __u32    i_dir_acl;       /* Directory ACL */
91:     __u32    i_faddr;         /* Fragment address */
92:     union {
93:         struct {
94:             __u8     l_i_frag;    /* Fragment number */
95:             __u8     l_i_fsize;   /* Fragment size */
96:             __u16    i_pad1;
97:             __u16    l_i_uid_high; /* these 2 fields */
98:             __u16    l_i_gid_high; /* were reserved2[0] */
99:             __u32    l_i_reserved2;
100:        } linux2;
101:        struct {
102:            __u8     h_i_frag;    /* Fragment number */
103:            __u8     h_i_fsize;   /* Fragment size */
104:            __u16    h_i_mode_high;
105:            __u16    h_i_uid_high;
106:            __u16    h_i_gid_high;
107:            __u32    h_i_author;
108:        } hurd2;
109:        struct {
110:            __u8     m_i_frag;    /* Fragment number */
111:            __u8     m_i_fsize;   /* Fragment size */
112:            __u16    m_pad1;
113:            __u32    m_i_reserved2[2];
114:        } masix2;
115:    } osd2;                    /* OS dependent 2 */
116: };
117:
118: /*
119:  * File system states
120:  */
121: #define EXT2_VALID_FS          0x0001 /* Unmounted cleanly */
122: #define EXT2_ERROR_FS         0x0002 /* Errors detected */
123:
124: /*
125:  * Structure of the super block
126:  */
127: struct ext2_super_block {
128:     __u32    s_inodes_count;    /* Inodes count */
129:     __u32    s_blocks_count;    /* Blocks count */
130:     __u32    s_r_blocks_count; /* Reserved blocks count */
131:     __u32    s_free_blocks_count; /* Free blocks count */
132:     __u32    s_free_inodes_count; /* Free inodes count */
133:     __u32    s_first_data_block; /* First Data Block */
134:     __u32    s_log_block_size; /* Block size */

```

## include/fiwix/fs\_ext2.h

Page 3/4

```

135:     __s32    s_log_frag_size;        /* Fragment size */
136:     __u32    s_blocks_per_group;     /* # Blocks per group */
137:     __u32    s_frags_per_group;      /* # Fragments per group */
138:     __u32    s_inodes_per_group;     /* # Inodes per group */
139:     __u32    s_mtime;                /* Mount time */
140:     __u32    s_wtime;                /* Write time */
141:     __u16    s_mnt_count;            /* Mount count */
142:     __s16    s_max_mnt_count;        /* Maximal mount count */
143:     __u16    s_magic;                /* Magic signature */
144:     __u16    s_state;                /* File system state */
145:     __u16    s_errors;               /* Behaviour when detecting errors */
146:     __u16    s_minor_rev_level;      /* minor revision level */
147:     __u32    s_lastcheck;            /* time of last check */
148:     __u32    s_checkinterval;        /* max. time between checks */
149:     __u32    s_creator_os;           /* OS */
150:     __u32    s_rev_level;            /* Revision level */
151:     __u16    s_def_resuid;           /* Default uid for reserved blocks */
152:     __u16    s_def_resgid;           /* Default gid for reserved blocks */
153:     /*
154:      * These fields are for EXT2_DYNAMIC_REV superblocks only.
155:      *
156:      * Note: the difference between the compatible feature set and
157:      * the incompatible feature set is that if there is a bit set
158:      * in the incompatible feature set that the kernel doesn't
159:      * know about, it should refuse to mount the filesystem.
160:      *
161:      * e2fsck's requirements are more strict; if it doesn't know
162:      * about a feature in either the compatible or incompatible
163:      * feature set, it must abort and not try to meddle with
164:      * things it doesn't understand...
165:      */
166:     __u32    s_first_ino;             /* First non-reserved inode */
167:     __u16    s_inode_size;           /* size of inode structure */
168:     __u16    s_block_group_nr;       /* block group # of this superblock */
169:     __u32    s_feature_compat;       /* compatible feature set */
170:     __u32    s_feature_incompat;     /* incompatible feature set */
171:     __u32    s_feature_ro_compat;    /* readonly-compatible feature set */
172:     __u8     s_uuid[16];              /* 128-bit uuid for volume */
173:     char     s_volume_name[16];       /* volume name */
174:     char     s_last_mounted[64];     /* directory where last mounted */
175:     __u32    s_algorithm_usage_bitmap; /* For compression */
176:     /*
177:      * Performance hints. Directory preallocation should only
178:      * happen if the EXT2_COMPAT_PREALLOC flag is on.
179:      */
180:     __u8     s_prealloc_blocks;       /* Nr of blocks to try to preallocate*/
181:     __u8     s_prealloc_dir_blocks;   /* Nr to preallocate for dirs */
182:     __u16    s_padding1;
183:     /*
184:      * Journaling support valid if EXT3_FEATURE_COMPAT_HAS_JOURNAL set.
185:      */
186:     __u8     s_journal_uuid[16];     /* uuid of journal superblock */
187:     __u32    s_journal_inum;         /* inode number of journal file */
188:     __u32    s_journal_dev;          /* device number of journal file */
189:     __u32    s_last_orphan;          /* start of list of inodes to delete */
190:     __u32    s_hash_seed[4];         /* HTREE hash seed */
191:     __u8     s_def_hash_version;     /* Default hash version to use */
192:     __u8     s_reserved_char_pad;
193:     __u16    s_reserved_word_pad;
194:     __u32    s_default_mount_opts;
195:     __u32    s_first_meta_bg;        /* First metablock block group */
196:     __u32    s_reserved[190];        /* Padding to the end of the block */
197: };
198:
199: /*
200:  * Structure of a directory entry
201:  */

```



## include/fiwix/fs\_ext2.h

Page 4/4

```

202: #define EXT2_NAME_LEN 255
203:
204: struct ext2_dir_entry {
205:     __u32    inode;           /* Inode number */
206:     __u16    rec_len;        /* Directory entry length */
207:     __u16    name_len;       /* Name length */
208:     char     name[EXT2_NAME_LEN]; /* File name */
209: };
210:
211: /*
212:  * The new version of the directory entry. Since EXT2 structures are
213:  * stored in intel byte order, and the name_len field could never be
214:  * bigger than 255 chars, it's safe to reclaim the extra byte for the
215:  * file_type field.
216:  */
217: struct ext2_dir_entry_2 {
218:     __u32    inode;           /* Inode number */
219:     __u16    rec_len;        /* Directory entry length */
220:     __u8     name_len;       /* Name length */
221:     __u8     file_type;
222:     char     name[EXT2_NAME_LEN]; /* File name */
223: };
224:
225: /*
226:  * Ext2 directory file types. Only the low 3 bits are used. The
227:  * other bits are reserved for now.
228:  */
229: #define EXT2_FT_UNKNOWN      0
230: #define EXT2_FT_REG_FILE    1
231: #define EXT2_FT_DIR          2
232: #define EXT2_FT_CHRDEV      3
233: #define EXT2_FT_BLKDEV      4
234: #define EXT2_FT_FIFO        5
235: #define EXT2_FT_SOCK        6
236: #define EXT2_FT_SYMLINK     7
237:
238: /* superblock in memory */
239: struct ext2_sb_info {
240:     unsigned int desc_per_block;
241:     unsigned int block_groups;
242:     struct ext2_super_block sb;
243: };
244:
245: /* inode in memory */
246: struct ext2_i_info {
247:     __u32    i_data[EXT2_N_BLOCKS]; /* Pointers to blocks */
248:     __u32    i_dtime;
249: };
250:
251: #endif /* _FIWIX_FS_EXT2_H */

```

## include/fiwix/fs.h

Page 1/4

```

1: /*
2:  * fiwix/include/fiwix/fs.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_FS_H
9: #define _FIWIX_FS_H
10:
11: #include <fiwix/config.h>
12: #include <fiwix/types.h>
13:
14: #define CHECK_UFD(ufd)                                     \
15: {                                                         \
16:     if((ufd) > OPEN_MAX || current->fd[(ufd)] == 0) {   \
17:         return -EBADF;                                   \
18:     }                                                     \
19: }                                                         \
20:
21: struct fd {
22:     struct inode *inode;           /* file inode */
23:     unsigned short int flags;     /* flags */
24:     unsigned short int count;     /* number of opened instances */
25:     __off_t offset;              /* r/w pointer position */
26: };
27:
28: #include <fiwix/statfs.h>
29: #include <fiwix/limits.h>
30: #include <fiwix/process.h>
31: #include <fiwix/dirent.h>
32: #include <fiwix/fs_minix.h>
33: #include <fiwix/fs_ext2.h>
34: #include <fiwix/fs_pipe.h>
35: #include <fiwix/fs_iso9660.h>
36: #include <fiwix/fs_proc.h>
37:
38: #define BPS 512 /* bytes per sector */
39: #define BLKSIZE_1K 1024 /* 1KB block size */
40: #define BLKSIZE_2K 2048 /* 2KB block size */
41: #define SUPERBLOCK 1 /* block 1 is for superblock */
42:
43: #define MAJOR(dev) (((__dev_t) (dev)) >> 8)
44: #define MINOR(dev) (((__dev_t) (dev)) & 0xFF)
45: #define MKDEV(major, minor) (((major) << 8) | (minor))
46:
47: /* filesystem independent mount-flags */
48: #define MS_RDONLY 1 /* mount read-only */
49: #define MS_REMOUNT 32 /* alter flags of a mounted FS */
50:
51: /* old magic mount flag and mask */
52: #define MS_MGC_VAL 0xC0ED0000
53: #define MS_MGC_MSK 0xFFFF0000
54:
55: #define IS_RDONLY_FS(inode) (((inode)->sb) && ((inode)->sb->flags & MS_RDONLY))
56:
57: #define FOLLOW_LINKS 1
58: #define MAX_SYMLINKS 8 /* this prevents infinite loops in symlinks */
59:
60: #define SEEK_SET 0
61: #define SEEK_CUR 1
62: #define SEEK_END 2
63:
64: #define FOR_READING 0
65: #define FOR_WRITING 1
66:
67: #define VERIFY_READ 1

```

## include/fiwix/fs.h

Page 2/4

```

68: #define VERIFY_WRITE      2
69:
70: #define SEL_R              1
71: #define SEL_W              2
72: #define SEL_E              4
73:
74: struct inode {
75:     __mode_t      i_mode;          /* file mode */
76:     __uid_t       i_uid;          /* owner uid */
77:     __size_t      i_size;         /* size in bytes */
78:     __u32         i_atime;        /* access time */
79:     __u32         i_ctime;        /* creation time */
80:     __u32         i_mtime;        /* modification time */
81:     __gid_t       i_gid;          /* group id */
82:     __nlink_t     i_nlink;        /* links count */
83:     __blk_t       i_blocks;       /* blocks count */
84:     __u32         i_flags;        /* file flags */
85:     unsigned char locked;
86:     unsigned char dirty;          /* 1 = delayed write */
87:     struct inode *mount_point;
88:     __dev_t       dev;
89:     __ino_t       inode;
90:     __s16         count;
91:     __dev_t       rdev;
92:     struct fs_operations *fsop;
93:     struct superblock *sb;
94:     struct inode *prev_hash;
95:     struct inode *next_hash;
96:     struct inode *prev_free;
97:     struct inode *next_free;
98:     union {
99:         struct minix_i_info minix;
100:        struct ext2_i_info ext2;
101:        struct pipefs_inode pipefs;
102:        struct iso9660_inode iso9660;
103:        struct procfs_inode procfs;
104:    } u;
105: };
106: extern struct inode *inode_table;
107: extern struct inode **inode_hash_table;
108: extern int inodes_on_free_list;
109:
110: /* values to be determined during system startup */
111: extern unsigned int inode_table_size;          /* size in bytes */
112: extern unsigned int inode_hash_table_size;     /* size in bytes */
113: extern unsigned int fd_table_size;            /* size in bytes */
114:
115: extern struct fd *fd_table;
116:
117: struct superblock {
118:     __dev_t dev;
119:     unsigned char locked;
120:     unsigned char wanted;
121:     struct inode *root;          /* root inode of mounted fs */
122:     struct inode *dir;          /* inode on which the fs was mounted */
123:     unsigned int flags;
124:     unsigned char dirty;        /* 1 = delayed write */
125:     struct fs_operations *fsop;
126:     __u32 s_blocksize;
127:     union {
128:         struct minix_sb_info minix;
129:         struct ext2_sb_info ext2;
130:         struct iso9660_sb_info iso9660;
131:     } u;
132: };
133:
134:

```

## include/fiwix/fs.h

Page 3/4

```

135: #define FSOP_REQUIRES_DEV      1      /* requires a block device */
136: #define FSOP_KERN_MOUNT        2      /* mounted by kernel */
137:
138: struct fs_operations {
139:     int flags;
140:     int fsdev;                  /* internal filesystem (nodev) */
141:
142: /* file operations */
143:     int (*open)(struct inode *, struct fd *);
144:     int (*close)(struct inode *, struct fd *);
145:     int (*read)(struct inode *, struct fd *, char *, __size_t);
146:     int (*write)(struct inode *, struct fd *, const char *, __size_t);
147:     int (*ioctl)(struct inode *, int, unsigned long int);
148:     int (*lseek)(struct inode *, __off_t);
149:     int (*readdir)(struct inode *, struct fd *, struct dirent *, unsigned in
t);
150:     int (*mmap)(struct inode *, struct vma *);
151:     int (*select)(struct inode *, int);
152:
153: /* inode operations */
154:     int (*readlink)(struct inode *, char *, __size_t);
155:     int (*followlink)(struct inode *, struct inode *, struct inode **);
156:     int (*bmap)(struct inode *, __off_t, int);
157:     int (*lookup)(const char *, struct inode *, struct inode **);
158:     int (*rmdir)(struct inode *, struct inode *);
159:     int (*link)(struct inode *, struct inode *, char *);
160:     int (*unlink)(struct inode *, struct inode *, char *);
161:     int (*symlink)(struct inode *, char *, char *);
162:     int (*mkdir)(struct inode *, char *, __mode_t);
163:     int (*mknod)(struct inode *, char *, __mode_t, __dev_t);
164:     int (*truncate)(struct inode *, __off_t);
165:     int (*create)(struct inode *, char *, __mode_t, struct inode **);
166:     int (*rename)(struct inode *, struct inode *, struct inode *, struct ino
de *, char *, char *);
167:
168: /* block device I/O operations */
169:     int (*read_block)(__dev_t, __blk_t, char *, int);
170:     int (*write_block)(__dev_t, __blk_t, char *, int);
171:
172: /* superblock operations */
173:     int (*read_inode)(struct inode *);
174:     int (*write_inode)(struct inode *);
175:     int (*ialloc)(struct inode *, int);
176:     void (*ifree)(struct inode *);
177:     void (*statfs)(struct superblock *, struct statfs *);
178:     int (*read_superblock)(__dev_t, struct superblock *);
179:     int (*remount_fs)(struct superblock *, int);
180:     int (*write_superblock)(struct superblock *);
181:     void (*release_superblock)(struct superblock *);
182: };
183:
184: extern struct fs_operations def_chr_fsop;
185: extern struct fs_operations def_blk_fsop;
186:
187: /* fs_minix.h prototypes */
188: extern struct fs_operations minix_fsop;
189: extern struct fs_operations minix_file_fsop;
190: extern struct fs_operations minix_dir_fsop;
191: extern struct fs_operations minix_symlink_fsop;
192: extern int minix_count_free_inodes(struct superblock *);
193: extern int minix_count_free_blocks(struct superblock *);
194: extern int minix_find_first_zero(struct superblock *, __blk_t, int, int);
195: extern int minix_change_bit(int, struct superblock *, int, int);
196: extern int minix_balloc(struct superblock *);
197: extern void minix_bfree(struct superblock *, int);
198:
199: /* fs_ext2.h prototypes */

```

## include/fiwix/fs.h

Page 4/4

```

200: extern struct fs_operations ext2_fsop;
201: extern struct fs_operations ext2_file_fsop;
202: extern struct fs_operations ext2_dir_fsop;
203: extern struct fs_operations ext2_symlink_fsop;
204: extern int ext2_balloc(struct superblock *);
205: extern void ext2_bfree(struct superblock *, int);
206:
207: /* fs_proc.h prototypes */
208: extern struct fs_operations procfs_fsop;
209: extern struct fs_operations procfs_file_fsop;
210: extern struct fs_operations procfs_dir_fsop;
211: extern struct fs_operations procfs_symlink_fsop;
212: struct procfs_dir_entry * get_procfs_by_inode(struct inode *);
213:
214: /* fs_iso9660.h prototypes */
215: extern int isonum_711(char *);
216: extern int isonum_723(char *);
217: extern int isonum_731(char *);
218: extern int isonum_733(char *);
219: extern unsigned long int isodate(char *);
220: extern int iso9660_cleanfilename(char *, int);
221: extern struct fs_operations iso9660_fsop;
222: extern struct fs_operations iso9660_file_fsop;
223: extern struct fs_operations iso9660_dir_fsop;
224: extern struct fs_operations iso9660_symlink_fsop;
225: void check_rrip_inode(struct iso9660_directory_record *, struct inode *);
226: int get_rrip_filename(struct iso9660_directory_record *, struct inode *, char *)
;
227: int get_rrip_symlink(struct inode *, char *);
228:
229:
230: /* generic VFS function prototypes */
231: void inode_lock(struct inode *);
232: void inode_unlock(struct inode *);
233: struct inode * ialloc(struct superblock *, int);
234: struct inode * iget(struct superblock *, __ino_t);
235: int bmap(struct inode *, __off_t, int);
236: int check_fs_busy(__dev_t, struct inode *);
237: void iput(struct inode *);
238: void sync_inodes(__dev_t);
239: void invalidate_inodes(__dev_t);
240: void inode_init(void);
241:
242: int parse_namei(char *, struct inode *, struct inode **, struct inode **, int);
243: int namei(char *, struct inode **, struct inode **, int);
244:
245: void superblock_lock(struct superblock *);
246: void superblock_unlock(struct superblock *);
247: struct mount * get_free_mount_point(__dev_t);
248: void release_mount_point(struct mount *);
249: struct mount * get_mount_point(struct inode *);
250:
251: int get_new_fd(struct inode *);
252: void release_fd(unsigned int);
253: void fd_init(void);
254:
255: void free_name(const char *);
256: int malloc_name(const char *, char **);
257: int check_user_permission(struct inode *);
258: int check_group(struct inode *);
259: int check_user_area(int, const void *, unsigned int);
260: int check_permission(int, struct inode *);
261:
262: int do_select(int, fd_set *, fd_set *, fd_set *, fd_set *, fd_set *, fd_set *);
263:
264: #endif /* _FIWIX_FS_H */

```

## include/fiwix/fs\_iso9660.h

Page 1/4

```

1: /*
2:  * fiwix/include/fiwix/fs_iso9660.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_FS_ISO9660_H
9: #define _FIWIX_FS_ISO9660_H
10:
11: #include <fiwix/types.h>
12: #include <fiwix/limits.h>
13:
14: #define ISO9660_SUPERBLOCK      16      /* ISO9660 superblock is in block 16 */
15: #define ISO9660_STANDARD_ID    "CD001" /* standard identification */
16: #define ISO9660_SUPER_MAGIC    0x9660
17:
18: #define ISO9660_VD_BOOT        0
19: #define ISO9660_VD_PRIMARY     1
20: #define ISO9660_VD_SUPPLEMENTARY 2
21: #define ISO9660_VD_PARTITION  3
22: #define ISO9660_VD_END        255
23:
24: #define ISODCL(from, to)      ((to - from) + 1)      /* descriptor length */
25:
26: #define ISO9660_MAX_VD       10      /* maximum number of VD per CDROM */
27:
28: /* inodes will have their directory block and their offset packed as follows:
29:  * 7FF7FF
30:  * \-/\-/\
31:  * ^ ^
32:  * | +----- offset value          (11bit entries)
33:  * +----- directory block where to find it (11bit entries)
34:  */
35: #define ISO9660_INODE_BITS    11      /* FIXME: it could be greater (16bit) */
36: #define ISO9660_INODE_MASK    0x7FF
37:
38: #define ISO9660_FILE_NOTEXIST  0x01   /* file shouldn't exists for the user */
39: #define ISO9660_FILE_ISDIR     0x02   /* is a directory */
40: #define ISO9660_FILE_ISASSOC   0x04   /* associated file */
41: #define ISO9660_FILE_HASRECFMT 0x08   /* has a record format */
42: #define ISO9660_FILE_HASOWNER  0x10   /* has owner and group defined */
43: #define ISO9660_FILE_RESERVED5 0x20   /* reserved */
44: #define ISO9660_FILE_RESERVED6 0x40   /* reserved */
45: #define ISO9660_FILE_ISMULTIEXT 0x80  /* has more directory records */
46:
47: #define SP_MAGIC1              0xBE
48: #define SP_MAGIC2              0xEF
49: #define GET_SIG(s1, s2)       ((s1 << 8) | s2)
50:
51: #define SL_CURRENT             0x02
52: #define SL_PARENT              0x04
53: #define SL_ROOT                0x08
54:
55: #define TF_CREATION            0x01
56: #define TF_MODIFY              0x02
57: #define TF_ACCESS              0x04
58: #define TF_ATTRIBUTES         0x08
59: #define TF_BACKUP              0x10
60: #define TF_EXPIRATION          0x20
61: #define TF_EFFECTIVE           0x40
62: #define TF_LONG_FORM           0x80
63:
64: #define NM_CONTINUE            0
65: #define NM_CURRENT             1
66: #define NM_PARENT              2
67:

```

## include/fiwix/fs\_iso9660.h

Page 2/4

```

68: /* formerly Primary Volume Descriptor */
69: struct iso9660_super_block {
70:     char type [ISODCL( 1, 1)]; /* 7.1.1 */
71:     char id [ISODCL( 2, 6)];
72:     char version [ISODCL( 7, 7)]; /* 7.1.1 */
73:     char unused1 [ISODCL( 8, 8)];
74:     char system_id [ISODCL( 9, 40)]; /* a-chars */
75:     char volume_id [ISODCL( 41, 72)]; /* d-chars */
76:     char unused2 [ISODCL( 73, 80)];
77:     char volume_space_size [ISODCL( 81, 88)]; /* 7.3.3 */
78:     char unused3 [ISODCL( 89, 120)];
79:     char volume_set_size [ISODCL(121, 124)]; /* 7.2.3 */
80:     char volume_sequence_number [ISODCL(125, 128)]; /* 7.2.3 */
81:     char logical_block_size [ISODCL(129, 132)]; /* 7.2.3 */
82:     char path_table_size [ISODCL(133, 140)]; /* 7.3.3 */
83:     char type_l_path_table [ISODCL(141, 144)]; /* 7.3.1 */
84:     char opt_type_l_path_table [ISODCL(145, 148)]; /* 7.3.1 */
85:     char type_m_path_table [ISODCL(149, 152)]; /* 7.3.2 */
86:     char opt_type_m_path_table [ISODCL(153, 156)]; /* 7.3.2 */
87:     char root_directory_record [ISODCL(157, 190)]; /* 9.1 */
88:     char volume_set_id [ISODCL(191, 318)]; /* d-chars */
89:     char publisher_id [ISODCL(319, 446)]; /* a-chars */
90:     char preparer_id [ISODCL(447, 574)]; /* a-chars */
91:     char application_id [ISODCL(575, 702)]; /* a-chars */
92:     char copyright_file_id [ISODCL(703, 739)]; /* 7.5 d-chars */
/
93:     char abstract_file_id [ISODCL(740, 776)]; /* 7.5 d-chars */
/
94:     char bibliographic_file_id [ISODCL(777, 813)]; /* 7.5 d-chars */
/
95:     char creation_date [ISODCL(814, 830)]; /* 8.4.26.1 */
96:     char modification_date [ISODCL(831, 847)]; /* 8.4.26.1 */
97:     char expiration_date [ISODCL(848, 864)]; /* 8.4.26.1 */
98:     char effective_date [ISODCL(865, 881)]; /* 8.4.26.1 */
99:     char file_structure_version [ISODCL(882, 882)];
100:    char unused4 [ISODCL(883, 883)];
101:    char application_data [ISODCL(884, 1395)];
102:    char unused5 [ISODCL(1396, 2048)];
103: };
104:
105: struct iso9660_directory_record
106: {
107:     char length [ISODCL( 1, 1)]; /* 7.1.1 */
108:     char ext_attr_length [ISODCL( 2, 2)]; /* 7.1.1 */
109:     char extent [ISODCL( 3, 10)]; /* 7.3.3 */
110:     char size [ISODCL(11, 18)]; /* 7.3.3 */
111:     char date [ISODCL(19, 25)]; /* 7 by 7.1.1 */
112:     char flags [ISODCL(26, 26)];
113:     char file_unit_size [ISODCL(27, 27)]; /* 7.1.1 */
114:     char interleave [ISODCL(28, 28)]; /* 7.1.1 */
115:     char volume_sequence_number [ISODCL(29, 32)]; /* 7.2.3 */
116:     char name_len [ISODCL(33, 33)]; /* 7.1.1 */
117:     char name[0];
118: };
119:
120: struct iso9660_pathtable_record
121: {
122:     char length [ISODCL( 1, 1)]; /* 7.1.1 */
123:     char ext_attr_length [ISODCL( 2, 2)]; /* 7.1.1 */
124:     char extent [ISODCL( 3, 6)]; /* 7.3 */
125:     char parent [ISODCL( 7, 8)]; /* 7.2 */
126:     char name[0];
127: };
128:
129: struct susp_sp {
130:     unsigned char magic[2];
131:     char len_skip;

```

```
132: };
133:
134: struct susp_ce {
135:     char block[8];
136:     char offset[8];
137:     char size[8];
138: };
139:
140: struct susp_er {
141:     char len_id;
142:     char len_des;
143:     char len_src;
144:     char ext_ver;
145:     char data[0];
146: };
147:
148: struct rrip_px {
149:     char mode[8];
150:     char nlink[8];
151:     char uid[8];
152:     char gid[8];
153:     char sn[8];
154: };
155:
156: struct rrip_pn {
157:     char dev_h[8];
158:     char dev_l[8];
159: };
160:
161: struct rrip_sl_component {
162:     unsigned char flags;
163:     unsigned char len;
164:     char name[0];
165: };
166:
167: struct rrip_sl {
168:     unsigned char flags;
169:     struct rrip_sl_component area;
170: };
171:
172: struct rrip_nm {
173:     unsigned char flags;
174:     char name[0];
175: };
176:
177: struct rrip_tf_timestamp {
178:     char time[7];           /* assumes LONG_FORM bit always set to zero */
179: };
180:
181: struct rrip_tf {
182:     char flags;
183:     struct rrip_tf_timestamp times[0];
184: };
185:
186: struct susp_rrip {
187:     char signature[2];
188:     unsigned char len;
189:     unsigned char version;
190:     union {
191:         struct susp_sp sp;
192:         struct susp_ce ce;
193:         struct susp_er er;
194:         struct rrip_px px;
195:         struct rrip_pn pn;
196:         struct rrip_sl sl;
197:         struct rrip_nm nm;
198:         struct rrip_tf tf;
```



**include/fiwix/fs\_iso9660.h**

Page 4/4

```
199:         } u;
200: };
201:
202: struct iso9660_inode {
203:     __blk_t i_extent;
204:     struct inode *i_parent;           /* inode of its parent directory */
205: };
206:
207: struct iso9660_sb_info {
208:     __u32 s_root_inode;
209:     char *pathtable_raw;
210:     struct iso9660_pathtable_record **pathtable;
211:     int paths;
212:     unsigned char rrip;
213:     struct iso9660_super_block *sb;
214: };
215:
216: #endif /* _FIWIX_FS_ISO9660_H */
```

## include/fiwix/fs\_minix.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/fs_minix.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_FS_MINIX_H
9: #define _FIWIX_FS_MINIX_H
10:
11: #include <fiwix/types.h>
12: #include <fiwix/limits.h>
13:
14: #define MINIX_ROOT_INO          1          /* root inode */
15:
16: #define MINIX_SUPER_MAGIC      0x137F    /* Minix v1, 14 char names */
17: #define MINIX_SUPER_MAGIC2    0x138F    /* Minix v1, 30 char names */
18: #define MINIX2_SUPER_MAGIC    0x2468    /* Minix v2, 14 char names */
19: #define MINIX2_SUPER_MAGIC2   0x2478    /* Minix v2, 30 char names */
20:
21: #define MINIX_VALID_FS        1          /* clean filesystem */
22: #define MINIX_ERROR_FS        2          /* needs fsck */
23:
24: #define CLEAR_BIT              0
25: #define SET_BIT                1
26:
27: #define V1_MAX_BITMAP_BLOCKS   8          /* 64MB filesystem size */
28: #define V2_MAX_BITMAP_BLOCKS  128         /* 1GB filesystem size */
29:
30: /*
31:  * Minix (v1 and v2) file system physical layout:
32:  *
33:  *      +-----+
34:  *      |                size in blocks of BLKSIZE_1K (1024 bytes)                |
35:  * +-----+-----+
36:  * | block 0          | 1 |
37:  * +-----+-----+
38:  * | superblock      | 1 |
39:  * +-----+-----+
40:  * | inode map       | number of inodes / (BLKSIZE_1K * 8) |
41:  * +-----+-----+
42:  * | zone map        | number of zones / (BLKSIZE_1K * 8) |
43:  * +-----+-----+
44:  * | inode table    | ((32 or 64) * number of inodes) / BLKSIZE_1K |
45:  * +-----+-----+
46:  * | data zones     | ... |
47:  * +-----+-----+
48:  *
49:  * The implementation of this filesystem in Fiwix might have slow disk writes
50:  * because I don't keep in memory the superblock, nor the blocks of the inode
51:  * map nor the blocks of the zone map. Keeping them in memory would be a waste
52:  * of 137KB per each mounted v2 filesystem (1GB of size).
53:  *
54:  * - superblock    -> 1KB
55:  * - inode map     -> 8KB (1KB (8192 bits) x 8 = 65536 inodes)
56:  * - zone map      -> 128KB (1KB (8192 bits) x 128 = 1048576 1k-blocks)
57:  *
58:  */
59:
60: struct minix_super_block {
61:     __u16 s_ninodes;          /* number of inodes */
62:     __u16 s_nzones;          /* number of data zones */
63:     __u16 s_imap_blocks;     /* blocks used by inode bitmap */
64:     __u16 s_zmap_blocks;     /* blocks used by zone bitmap */
65:     __u16 s_firstdatazone;   /* number of first data zone */
66:     __u16 s_log_zone_size;   /* 1024 << s_log_zone_size */
67:     __u32 s_max_size;        /* maximum file size (in bytes) */

```

## include/fiwix/fs\_minix.h

Page 2/2

```

68:     __u16 s_magic;           /* magic number */
69:     __u16 s_state;          /* filesystem state */
70:     __u32 s_zones;          /* number of data zones (for v2 only) */
71: };
72:
73: struct minix_inode {
74:     __u16 i_mode;
75:     __u16 i_uid;
76:     __u32 i_size;
77:     __u32 i_time;
78:     __u8 i_gid;
79:     __u8 i_nlinks;
80:     __u16 i_zone[9];
81: };
82:
83: struct minix2_inode {
84:     __u16 i_mode;
85:     __u16 i_nlink;
86:     __u16 i_uid;
87:     __u16 i_gid;
88:     __u32 i_size;
89:     __u32 i_atime;
90:     __u32 i_mtime;
91:     __u32 i_ctime;
92:     __u32 i_zone[10];
93: };
94:
95: struct minix_dir_entry {
96:     __u16 inode;
97:     char name[0];
98: };
99:
100: /* super block in memory */
101: struct minix_sb_info {
102:     unsigned char namelen;
103:     unsigned char dirsize;
104:     unsigned short int version;
105:     unsigned int nzones;
106:     struct minix_super_block sb;
107: };
108:
109: /* inode in memory */
110: struct minix_i_info {
111:     union {
112:         __u16 i1_zone[9];
113:         __u32 i2_zone[10];
114:     } u;
115: };
116:
117: int v1_minix_read_inode(struct inode *);
118: int v1_minix_write_inode(struct inode *);
119: int v1_minix_ialloc(struct inode *, int);
120: void v1_minix_ifree(struct inode *);
121: int v1_minix_bmap(struct inode *, __off_t, int);
122: int v1_minix_truncate(struct inode *, __off_t);
123:
124: int v2_minix_read_inode(struct inode *);
125: int v2_minix_write_inode(struct inode *);
126: int v2_minix_ialloc(struct inode *, int);
127: void v2_minix_ifree(struct inode *);
128: int v2_minix_bmap(struct inode *, __off_t, int);
129: int v2_minix_truncate(struct inode *, __off_t);
130:
131: #endif /* _FIWIX_FS_MINIX_H */

```

**include/fiwix/fs\_pipe.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/fs_pipe.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_FS_PIPE_H
9: #define _FIWIX_FS_PIPE_H
10:
11: #define PIPE_DEV          0xFFF0          /* special device number for nodev fs */
12:
13: extern struct fs_operations pipefs_fsop;
14:
15: struct pipefs_inode {
16:     char *i_data;          /* buffer */
17:     unsigned int i_readoff; /* offset for reads */
18:     unsigned int i_writeoff; /* offset for writes */
19:     unsigned int i_readers; /* number of readers */
20:     unsigned int i_writers; /* number of writers */
21: };
22:
23: #endif /* _FIWIX_FS_PIPE_H */
```

## include/fiwix/fs\_proc.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/fs_proc.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_FS_PROC_H
9: #define _FIWIX_FS_PROC_H
10:
11: #include <fiwix/types.h>
12:
13: #define PROC_DEV                0xFFF1 /* special device number for nodev fs */
14: #define PROC_ROOT_INO           1      /* root inode */
15: #define PROC_SUPER_MAGIC        0x9FA0 /* same as in Linux */
16:
17: #define PROC_PID_INO            0x40000000 /* base for PID inodes */
18: #define PROC_PID_LEV            1      /* array level for PID */
19:
20: #define PROC_ARRAY_ENTRIES      20
21:
22: enum pid_dir_inodes {
23:     PROC_PID_FD = PROC_PID_INO + 1001,
24:     PROC_PID_CMDLINE,
25:     PROC_PID_CWD,
26:     PROC_PID_ENVIRON,
27:     PROC_PID_EXE,
28:     PROC_PID_MAPS,
29:     PROC_PID_MOUNTINFO,
30:     PROC_PID_ROOT,
31:     PROC_PID_STAT,
32:     PROC_PID_STATM,
33:     PROC_PID_STATUS
34: };
35:
36: struct procfs_inode {
37:     unsigned int i_lev; /* array level (directory depth) */
38: };
39:
40: struct procfs_dir_entry {
41:     __ino_t inode;
42:     __mode_t mode;
43:     __nlink_t nlink;
44:     int lev; /* array level (directory depth) */
45:     unsigned short int name_len;
46:     char *name;
47:     int (*data_fn)(char *, __pid_t);
48: };
49:
50: extern struct procfs_dir_entry procfs_array[][PROC_ARRAY_ENTRIES + 1];
51:
52: int data_proc_cmdline(char *, __pid_t);
53: int data_proc_cpuinfo(char *, __pid_t);
54: int data_proc_devices(char *, __pid_t);
55: int data_proc_dma(char *, __pid_t);
56: int data_proc_filesystems(char *, __pid_t);
57: int data_proc_interrupts(char *, __pid_t);
58: int data_proc_loadavg(char *, __pid_t);
59: int data_proc_locks(char *, __pid_t);
60: int data_proc_meminfo(char *, __pid_t);
61: int data_proc_mounts(char *, __pid_t);
62: int data_proc_partitions(char *, __pid_t);
63: int data_proc_rtc(char *, __pid_t);
64: int data_proc_self(char *, __pid_t);
65: int data_proc_stat(char *, __pid_t);
66: int data_proc_uptime(char *, __pid_t);
67: int data_proc_fullversion(char *, __pid_t);

```

**include/fiwix/fs\_proc.h**

Page 2/2

```
68: int data_proc_domainname(char *, __pid_t);
69: int data_proc_filemax(char *, __pid_t);
70: int data_proc_filenr(char *, __pid_t);
71: int data_proc_hostname(char *, __pid_t);
72: int data_proc_inodemax(char *, __pid_t);
73: int data_proc_inodenr(char *, __pid_t);
74: int data_proc_osrelease(char *, __pid_t);
75: int data_proc_ostype(char *, __pid_t);
76: int data_proc_version(char *, __pid_t);
77:
78: /* PID related functions */
79: int data_proc_pid_fd(char *, __pid_t);
80: int data_proc_pid_cmdline(char *, __pid_t);
81: int data_proc_pid_cwd(char *, __pid_t);
82: int data_proc_pid_envIRON(char *, __pid_t);
83: int data_proc_pid_exe(char *, __pid_t);
84: int data_proc_pid_maps(char *, __pid_t);
85: int data_proc_pid_mountinfo(char *, __pid_t);
86: int data_proc_pid_root(char *, __pid_t);
87: int data_proc_pid_stat(char *, __pid_t);
88: int data_proc_pid_statm(char *, __pid_t);
89: int data_proc_pid_status(char *, __pid_t);
90:
91: #endif /* _FIWIX_FS_PROC_H */
```

## include/fiwix/i386elf.h

Page 1/5

```

1: /*
2:  * fiwix/include/fiwix/i386elf.h
3:  */
4:
5: #ifndef _FIWIX_ELF_H
6: #define _FIWIX_ELF_H
7:
8: typedef unsigned long   Elf32_Addr;
9: typedef unsigned short Elf32_Half;
10: typedef unsigned long   Elf32_Off;
11: typedef long            Elf32_Sword;
12: typedef unsigned long   Elf32_Word;
13:
14: #define ELFMAG0          0x7f          /* EI_MAG */
15: #define ELFMAG1          'E'
16: #define ELFMAG2          'L'
17: #define ELFMAG3          'F'
18: #define ELFMAG           "\177ELF"
19: #define SELFMAG          4
20:
21: #define EI_NIDENT        16
22:
23: typedef struct elf32_hdr{
24:     unsigned char e_ident[EI_NIDENT]; /* ELF "magic number" */
25:     Elf32_Half    e_type;             /* File type */
26:     Elf32_Half    e_machine;         /* Target machine */
27:     Elf32_Word    e_version;        /* File version */
28:     Elf32_Addr    e_entry;           /* Entry point virtual address */
29:     Elf32_Off     e_phoff;           /* Program header table file offset */
30:     Elf32_Off     e_shoff;           /* Section header table file offset */
31:     Elf32_Word    e_flags;           /* File flags */
32:     Elf32_Half    e_ehsize;          /* Sizeof Ehdr (ELF header) */
33:     Elf32_Half    e_phentsize;       /* Sizeof Phdr (Program header) */
34:     Elf32_Half    e_phnum;           /* Number Phdrs (Program header) */
35:     Elf32_Half    e_shentsize;       /* Sizeof Shdr (Section header) */
36:     Elf32_Half    e_shnum;           /* Number Shdrs (Section header) */
37:     Elf32_Half    e_shstrndx;        /* Shdr string index */
38: } Elf32_Ehdr;
39:
40: #define EI_MAG0          0            /* e_ident[] indexes */
41: #define EI_MAG1          1
42: #define EI_MAG2          2
43: #define EI_MAG3          3
44: #define EI_CLASS         4
45: #define EI_DATA          5
46: #define EI_VERSION       6
47: #define EI_PAD           7
48:
49: #define ELFCLASSNONE     0            /* EI_CLASS */
50: #define ELFCLASS32       1
51: #define ELFCLASS64       2
52: #define ELFCLASSNUM      3
53:
54: #define ELFDATANONE      0            /* e_ident[EI_DATA] */
55: #define ELFDATA2LSB      1
56: #define ELFDATA2MSB      2
57: #define ELFDATANUM       3
58:
59: /* ELF file types */
60: #define ET_NONE          0
61: #define ET_REL           1
62: #define ET_EXEC          2
63: #define ET_DYN           3
64: #define ET_CORE          4
65: #define ET_LOPROC        5
66: #define ET_HIPROC        6
67:

```

## include/fiwix/i386elf.h

Page 2/5

```

68: #define EM_386      3
69:
70: #define EV_NONE      0          /* e_version, EI_VERSION */
71: #define EV_CURRENT   1
72: #define EV_NUM       2
73:
74: typedef struct elf32_phdr{
75:     Elf32_Word      p_type;      /* Entry type */
76:     Elf32_Off       p_offset;    /* File offset */
77:     Elf32_Addr      p_vaddr;    /* Virtual address */
78:     Elf32_Addr      p_paddr;    /* Physical address */
79:     Elf32_Word      p_filesz;   /* File size */
80:     Elf32_Word      p_memsz;    /* Memory size */
81:     Elf32_Word      p_flags;    /* Entry flags */
82:     Elf32_Word      p_align;    /* Memory & file alignment */
83: } Elf32_Phdr;
84:
85: /* segment types stored in the image headers */
86: #define PT_NULL      0
87: #define PT_LOAD      1
88: #define PT_DYNAMIC   2
89: #define PT_INTERP    3
90: #define PT_NOTE      4
91: #define PT_SHLIB     5
92: #define PT_PHDR      6
93: #define PT_NUM       7
94: #define PT_LOPROC    0x70000000
95: #define PT_HIPROC    0x7fffffff
96:
97: /* permission types on sections in the program header, p_flags. */
98: #define PF_R         0x4
99: #define PF_W         0x2
100: #define PF_X         0x1
101:
102: #define PF_MASKPROC 0xf0000000
103:
104: typedef struct {
105:     Elf32_Word      sh_name;     /* Section name, index in string tbl */
106:     Elf32_Word      sh_type;     /* Type of section */
107:     Elf32_Word      sh_flags;    /* Miscellaneous section attributes */
108:     Elf32_Addr      sh_addr;     /* Section virtual addr at execution */
109:     Elf32_Off       sh_offset;   /* Section file offset */
110:     Elf32_Word      sh_size;     /* Size of section in bytes */
111:     Elf32_Word      sh_link;     /* Index of another section */
112:     Elf32_Word      sh_info;     /* Additional section information */
113:     Elf32_Word      sh_addralign; /* Section alignment */
114:     Elf32_Word      sh_entsize;  /* Entry size if section holds table */
115: } Elf32_Shdr;
116:
117: /* sh_type */
118: #define SHT_NULL      0
119: #define SHT_PROGBITS  1
120: #define SHT_SYMTAB    2
121: #define SHT_STRTAB    3
122: #define SHT_RELA      4
123: #define SHT_HASH      5
124: #define SHT_DYNAMIC   6
125: #define SHT_NOTE      7
126: #define SHT_NOBITS    8
127: #define SHT_REL       9
128: #define SHT_SHLIB     10
129: #define SHT_DYNSYM    11
130: #define SHT_NUM       12
131:
132: #define SHT_LOPROC    0x70000000
133: #define SHT_HIPROC    0x7fffffff
134: #define SHT_LOUSER    0x80000000

```



## include/fiwix/i386elf.h

Page 3/5

```

135: #define SHT_HIUSER          0xffffffff
136:
137: /* sh_flags */
138: #define SHF_WRITE           0x1
139: #define SHF_ALLOC           0x2
140: #define SHF_EXECINSTR      0x4
141:
142: /* special section indexes */
143: #define SHN_UNDEF           0
144: #define SHN_LORESERVE      0xff00
145: #define SHN_ABS             0xffff1
146: #define SHN_COMMON         0xffff2
147: #define SHN_HIRESERVE      0xfffff
148: #define SHN_LOPROC         0xff00
149: #define SHN_HIPROC         0xff1f
150:
151: typedef struct elf32_sym{
152:     Elf32_Word    st_name;           /* Symbol name, index in string tbl */
153:     Elf32_Addr    st_value;         /* Value of the symbol */
154:     Elf32_Word    st_size;         /* Associated symbol size */
155:     unsigned char st_info;         /* Type and binding attributes */
156:     unsigned char st_other;        /* No defined meaning, 0 */
157:     Elf32_Half    st_shndx;        /* Associated section index */
158: } Elf32_Sym;
159:
160: #define ELF32_ST_BIND(info)      ((info) >> 4)
161: #define ELF32_ST_TYPE(info)      (((unsigned int) info) & 0xf)
162:
163: /* This info is needed when parsing the symbol table */
164: #define STB_LOCAL              0
165: #define STB_GLOBAL              1
166: #define STB_WEAK                2
167: #define STB_NUM                  3
168:
169: #define STT_NOTYPE              0
170: #define STT_OBJECT              1
171: #define STT_FUNC                2
172: #define STT_SECTION             3
173: #define STT_FILE                4
174: #define STT_NUM                  5
175:
176: typedef struct elf32_rel {
177:     Elf32_Addr    r_offset;         /* Location at which to apply the action */
178:     Elf32_Word    r_info;          /* Index and type of relocation */
179: } Elf32_Rel;
180:
181: typedef struct elf32_rela{
182:     Elf32_Addr    r_offset;         /* Location at which to apply the action */
183:     Elf32_Word    r_info;          /* Index and type of relocation */
184:     Elf32_Sword   r_addend;        /* Constant addend used to compute value */
185: } Elf32_Rela;
186:
187: /* The following are used with relocations */
188: #define ELF32_R_SYM(info)        ((info) >> 8)
189: #define ELF32_R_TYPE(info)       ((info) & 0xff)
190:
191: /* This is the info that is needed to parse the dynamic section of the file */
192: #define DT_NULL                  0
193: #define DT_NEEDED                1
194: #define DT_PLTRELSZ              2
195: #define DT_PLTGOT                3
196: #define DT_HASH                   4
197: #define DT_STRTAB                 5
198: #define DT_SYMTAB                 6
199: #define DT_RELA                   7
200: #define DT_RELASZ                 8
201: #define DT_RELAENT                 9

```

## include/fiwix/i386elf.h

Page 4/5

```

202: #define DT_STRSZ          10
203: #define DT_SYMENT        11
204: #define DT_INIT          12
205: #define DT_FINI          13
206: #define DT_SONAME        14
207: #define DT_RPATH         15
208: #define DT_SYMBOLIC      16
209: #define DT_REL           17
210: #define DT_RELSZ         18
211: #define DT_RELENT        19
212: #define DT_PLTREL        20
213: #define DT_DEBUG         21
214: #define DT_TEXTREL       22
215: #define DT_JMPREL        23
216: #define DT_LOPROC        0x70000000
217: #define DT_HIPROC        0x7fffffff
218:
219: /* Symbolic values for the entries in the auxiliary table
220:    put on the initial stack */
221: #define AT_NULL          0    /* end of vector */
222: #define AT_IGNORE       1    /* entry should be ignored */
223: #define AT_EXECFD       2    /* file descriptor of program */
224: #define AT_PHDR         3    /* program headers for program */
225: #define AT_PHEMT        4    /* size of program header entry */
226: #define AT_PHNUM        5    /* number of program headers */
227: #define AT_PAGESZ       6    /* system page size */
228: #define AT_BASE         7    /* base address of interpreter */
229: #define AT_FLAGS        8    /* flags */
230: #define AT_ENTRY        9    /* entry point of program */
231: #define AT_NOTELF      10    /* program is not ELF */
232: #define AT_UID         11    /* real uid */
233: #define AT_EUID        12    /* effective uid */
234: #define AT_GID         13    /* real gid */
235: #define AT_EGID        14    /* effective gid */
236:
237:
238: typedef struct dynamic{
239:     Elf32_Sword d_tag;          /* entry tabg value */
240:     union{
241:         Elf32_Sword d_val;
242:         Elf32_Addr d_ptr;
243:     } d_un;
244: } Elf32_Dyn;
245:
246: #define R_386_NONE        0
247: #define R_386_32         1
248: #define R_386_PC32       2
249: #define R_386_GOT32      3
250: #define R_386_PLT32      4
251: #define R_386_COPY       5
252: #define R_386_GLOB_DAT   6
253: #define R_386_JMP_SLOT   7
254: #define R_386_RELATIVE   8
255: #define R_386_GOTOFF     9
256: #define R_386_GOTPC     10
257: #define R_386_NUM        11
258:
259: /* Notes used in ET_CORE */
260: #define NT_PRSTATUS      1
261: #define NT_PRFPREG       2
262: #define NT_PRPSINFO      3
263: #define NT_TASKSTRUCT    4
264:
265: /* Note header in a PT_NOTE section */
266: typedef struct elf32_note {
267:     Elf32_Word    n_namesz;    /* Name size */
268:     Elf32_Word    n_descsz;    /* Content size */

```

**include/fiwix/i386elf.h**

Page 5/5

```
269:   Elf32_Word    n_type;           /* Content type */
270: } Elf32_Nhdr;
271:
272: #define ELF_START_MMAP 0x80000000
273:
274: extern Elf32_Dyn __DYNAMIC [];
275: #define elfhdr      elf32_hdr
276: #define elf_phdr    elf32_phdr
277: #define elf_note    elf32_note
278:
279: #endif /* _FIWIX_ELF_H */
```

**include/fiwix/ide\_cd.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/ide_cd.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_IDE_CD_H
9: #define _FIWIX_IDE_CD_H
10:
11: #include <fiwix/fs.h>
12:
13: #define IDE_CD_SECTSIZE          BLKSIZE_2K          /* sector size (in bytes) */
14:
15: void ide_cd_timer(unsigned int);
16:
17: int ide_cd_open(struct inode *, struct fd *);
18: int ide_cd_close(struct inode *, struct fd *);
19: int ide_cd_read(__dev_t, __blk_t, char *, int);
20: int ide_cd_ioctl(struct inode *, int, unsigned long int);
21:
22: int ide_cd_init(struct ide *, int);
23:
24: #endif /* _FIWIX_IDE_CD_H */
```

## include/fiwix/ide.h

Page 1/5

```

1: /*
2:  * fiwix/include/fiwix/ide.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_IDE_H
9: #define _FIWIX_IDE_H
10:
11: #include <fiwix/fs.h>
12: #include <fiwix/part.h>
13: #include <fiwix/sigcontext.h>
14:
15: #define IDE0_IRQ          14      /* primary controller interrupt */
16: #define IDE1_IRQ          15      /* secondary controller interrupt */
17:
18: #define IDE0_MAJOR        3       /* 1st controller major number */
19: #define IDE1_MAJOR        22      /* 2nd controller major number */
20: #define IDE_MINORS        4       /* max. minors/partitions per unit */
21: #define IDE_MASTER_MSF    0       /* IDE master minor shift factor */
22: #define IDE_SLAVE_MSF    6       /* IDE slave minor shift factor */
23:
24: #define IDE_PRIMARY       0
25: #define IDE_SECONDARY    1
26: #define IDE_MASTER       0
27: #define IDE_SLAVE       1
28: #define IDE_ATA          0
29: #define IDE_ATAPI       1
30:
31: #define NR_IDE_CTRL      2        /* IDE controllers */
32: #define NR_IDE_DRVS     2        /* max. drives per IDE controller */
33:
34: /* controller addresses */
35: #define IDE0_BASE        0x1F0    /* primary controller base addr */
36: #define IDE0_CTRL        0x3F4    /* primary controller control port */
37: #define IDE1_BASE        0x170    /* secondary controller base addr */
38: #define IDE1_CTRL        0x374    /* secondary controller control port */
39:
40: #define IDE_BASE_LEN     7        /* controller address length */
41:
42: #define IDE_RDY_RETR_LONG 50000   /* long delay for fast CPUs */
43: #define IDE_RDY_RETR_SHORT 500    /* short delay for slow CPUs */
44: #define MAX_IDE_ERR      10       /* number of retries */
45: #define MAX_CD_ERR       5        /* number of retries in CDRoms */
46:
47: #define SET_IDE_RDY_RETR(retries) \
48:     if((cpu_table.hz / 1000000) <= 100) { \
49:         retries = IDE_RDY_RETR_SHORT; \
50:     } else { \
51:         retries = IDE_RDY_RETR_LONG; \
52:     }
53:
54: #define WAIT_FOR_IDE      (1 * HZ) /* timeout for hard disk */
55: #define WAIT_FOR_CD       (3 * HZ) /* timeout for cdrom */
56:
57: /* controller registers */
58: #define IDE_DATA          0x0     /* Data Port Register (R/W) */
59: #define IDE_ERROR         0x1     /* Error Register (R) */
60: #define IDE_FEATURES      0x1     /* Features Register (W) */
61: #define IDE_SECCNT        0x2     /* Sector Count Register (R/W) */
62: #define IDE_SECNUM        0x3     /* Sector Number Register (R/W) */
63: #define IDE_LCYL          0x4     /* Cylinder Low Register (R/W) */
64: #define IDE_HCYL          0x5     /* Cylinder High Register (R/W) */
65: #define IDE_DRVHD         0x6     /* Drive/Head Register (R/W) */
66: #define IDE_STATUS        0x7     /* Status Register (R) */
67: #define IDE_COMMAND       0x7     /* Command Register (W) */

```

## include/fiwix/ide.h

Page 2/5

```

68:
69: #define IDE_ALT_STATUS          0x2    /* Alternate Register (R) */
70: #define IDE_DEV_CTRL           0x2    /* Device Control Register (W) */
71:
72: /* error register bits */
73: #define IDE_ERR_AMNF           0x01   /* Address Mark Not Found */
74: #define IDE_ERR_TKONF         0x02   /* Track 0 Not Found */
75: #define IDE_ERR_ABRT          0x04   /* Aborted Command */
76: #define IDE_ERR_MCR           0x08   /* Media Change Registered */
77: #define IDE_ERR_IDNF          0x10   /* Sector ID Field Not Found */
78: #define IDE_ERR_MC            0x20   /* Media Changed */
79: #define IDE_ERR_UNC           0x40   /* Uncorrectable Data Error */
80: #define IDE_ERR_BBK           0x80   /* Bad Block */
81:
82: /* status register bits */
83: #define IDE_STAT_ERR           0x01   /* an error occurred */
84: #define IDE_STAT_SENS         0x02   /* sense data available */
85: #define IDE_STAT_CORR         0x04   /* a correctable error occurred */
86: #define IDE_STAT_DRQ          0x08   /* device is ready to transfer */
87: #define IDE_STAT_DSC          0x10   /* device requests service o intr. */
88: #define IDE_STAT_DWF          0x20   /* drive write fault */
89: #define IDE_STAT_RDY          0x40   /* drive is ready */
90: #define IDE_STAT_BSY          0x80   /* drive is busy */
91:
92: #define IDE_CHS_MODE           0xA0   /* select CHS mode */
93: #define IDE_LBA_MODE           0xE0   /* select LBA mode */
94:
95: /* alternate & device control register bits */
96: #define IDE_DEVCTR_DRQ        0x08   /* Data Request */
97: #define IDE_DEVCTR_NIEN       0x02   /* Disable Interrupt */
98: #define IDE_DEVCTR_SRST       0x04   /* Software Reset */
99:
100: /* ATA commands */
101: #define ATA_READ_PIO           0x20   /* read sector(s) with retries */
102: #define ATA_READ_MULTIPLE_PIO 0xC4   /* read multiple sectors */
103: #define ATA_WRITE_PIO          0x30   /* write sector(s) with retries */
104: #define ATA_WRITE_MULTIPLE_PIO 0xC5   /* write multiple sectors */
105: #define ATA_SET_MULTIPLE_MODE  0xC6   /* set multiple mode */
106: #define ATA_PACKET             0xA0   /* packet command */
107: #define ATA_IDENTIFY_PACKET    0xA1   /* identify ATAPI device */
108: #define ATA_IDENTIFY           0xEC   /* identify ATA device */
109:
110: /* ATAPI commands */
111: #define ATAPI_TEST_UNIT        0x00   /* test unit ready */
112: #define ATAPI_REQUEST_SENSE    0x03   /* request sense */
113: #define ATAPI_START_STOP       0x1B   /* start/stop diskette */
114: #define ATAPI_MEDIUM_REMOVAL   0x1E   /* medium removal */
115: #define ATAPI_READ10           0x28   /* read 10 */
116:
117: #define CD_UNLOCK_MEDIUM        0x00   /* allow medium removal */
118: #define CD_LOCK_MEDIUM         0x01   /* prevent medium removal */
119: #define CD_EJECT                0x02   /* eject the CD if possible */
120: #define CD_LOAD                 0x03   /* load the CD (closes tray) */
121:
122: /* ATAPI CD additional sense code */
123: #define ASC_NOT_READY           0x04   /* not ready */
124: #define ASC_NO_MEDIUM          0x3A   /* no medium */
125:
126: /* capabilities */
127: #define IDE_SUPPORTS_CFA        0x848A /* CFA support */
128: #define IDE_HAS_DMA             0x100  /* DMA support */
129: #define IDE_HAS_LBA             0x200  /* LBA support */
130: #define IDE_MIN_LBA             16514064 /* sectors limit for using CHS */
131:
132: /* general configuration bits */
133: #define IDE_HAS_UDMA            0x04   /* device supports UDMA */
134: #define IDE_REMOVABLE           0x80   /* removable media device */

```

## include/fiwix/ide.h

Page 3/5

```

135:
136: /* ATAPI types */
137: #define ATAPI_IS_SEQ_ACCESS      0x01    /* sequential-access device */
138: #define ATAPI_IS_PRINTER        0x02
139: #define ATAPI_IS_PROCESSOR      0x03
140: #define ATAPI_IS_WRITE_ONCE     0x04
141: #define ATAPI_IS_CDROM          0x05
142: #define ATAPI_IS_SCANNER        0x06
143:
144: /* IDE drive flags */
145: #define DEVICE_IS_ATAPI         0x01
146: #define DEVICE_IS_CFA           0x02
147: #define DEVICE_IS_DISK          0x04
148: #define DEVICE_IS_CDROM         0x08
149: #define DEVICE_REQUIRES_LBA     0x10
150: #define DEVICE_HAS_RW_MULTIPLE  0x20
151:
152: /* ATA/ATAPI-4 based */
153: struct ide_drv_ident {
154:     unsigned short int gen_config;      /* general configuration bits */
155:     unsigned short int logic_cyls;      /* logical cylinders */
156:     unsigned short int reserved2;
157:     unsigned short int logic_heads;     /* logical heads */
158:     unsigned short int retired4;
159:     unsigned short int retired5;
160:     unsigned short int logic_spt;       /* logical sectors/track */
161:     unsigned short int retired7;
162:     unsigned short int retired8;
163:     unsigned short int retired9;
164:     char serial_number[20];             /* serial number */
165:     unsigned short int vendor_spec20;
166:     unsigned short int buffer_cache;
167:     unsigned short int vendor_spec22;   /* reserved */
168:     char firmware_rev[8];              /* firmware version */
169:     char model_number[40];             /* model number */
170:     unsigned short int rw_multiple;
171:     unsigned short int reserved48;
172:     unsigned short int capabilities;     /* capabilities */
173:     unsigned short int reserved50;
174:     unsigned short int pio_mode;        /* PIO data transfer mode*/
175:     unsigned short int dma_mode;
176:     unsigned short int fields_validity; /* fields validity */
177:     unsigned short int cur_log_cyls;    /* current logical cylinders */
178:     unsigned short int cur_log_heads;   /* current logical heads */
179:     unsigned short int cur_log_spt;     /* current logical sectors/track
*/
180:     unsigned short int cur_capacity;     /* current capacity in sectors */
/
181:     unsigned short int cur_capacity2;    /* 32bit number */
182:     unsigned short int mult_sect_set;    /* multiple sector settings */
183:     unsigned short int tot_sectors;      /* sectors (LBA only) */
184:     unsigned short int tot_sectors2;    /* 32bit number */
185:     unsigned short int singleword_dma;
186:     unsigned short int multiword_dma;   /* multiword DMA settings */
187:     unsigned short int adv_pio_modes;   /* advanced PIO modes */
188:     unsigned short int min_multiword;   /* min. Multiword DMA transfer */
/
189:     unsigned short int rec_multiword;   /* recommended Multiword DMS tra
nsfer */
190:     unsigned short int min_pio_wo_fc;   /* min. PIO w/o flow control */
191:     unsigned short int min_pio_w_fc;    /* min. PIO with flow control */
192:     unsigned short int reserved69;
193:     unsigned short int reserved70;
194:     unsigned short int reserved71;
195:     unsigned short int reserved72;
196:     unsigned short int reserved73;
197:     unsigned short int reserved74;

```

## include/fiwix/ide.h

Page 4/5

```

198:     unsigned short int queue_depth;           /* queue depth */
199:     unsigned short int reserved76;
200:     unsigned short int reserved77;
201:     unsigned short int reserved78;
202:     unsigned short int reserved79;
203:     unsigned short int majorver;              /* major version number */
204:     unsigned short int minorver;              /* minor version number */
205:     unsigned short int cmdset1;                /* command set supported */
206:     unsigned short int cmdset2;                /* command set supported */
207:     unsigned short int cmdsf_ext;              /* command set/feature sup.ext.
*/
208:     unsigned short int cmdsf_enable1;          /* command s/f enabled */
209:     unsigned short int cmdsf_enable2;          /* command s/f enabled */
210:     unsigned short int cmdsf_default;          /* command s/f default */
211:     unsigned short int ultradma;               /* ultra DMA mode */
212:     unsigned short int reserved89;
213:     unsigned short int reserved90;
214:     unsigned short int curapm;                 /* current APM values */
215:     unsigned short int reserved92_126[35];
216:     unsigned short int r_status_notif;         /* removable media status notif.
*/
217:     unsigned short int security_status;        /* security status */
218:     unsigned short int vendor_spec129_159[31];
219:     unsigned short int reserved160_255[96];
220: };
221:
222: struct ide_drv {
223:     int drive;                                /* master or slave */
224:     char *dev_name;
225:     unsigned char major;                       /* major number */
226:     unsigned int flags;
227:     int minor_shift;                           /* shift factor to get the real minor */
228:     int lba_cyls;
229:     int lba_heads;
230:     short int lba_factor;
231:     unsigned int nr_sects;                       /* total sectors (LBA) */
232:     struct fs_operations *fsop;
233:     struct ide_drv_ident ident;
234:     struct partition part_table[NR_PARTITIONS];
235: };
236:
237: struct ide {
238:     int channel;                               /* primary or secondary */
239:     int base;                                  /* base address */
240:     int ctrl;                                  /* control port address */
241:     short int irq;
242:     struct ide_drv drive[NR_IDE_DRVS];
243: };
244:
245: extern struct ide ide_table[NR_IDE_CTRLIS];
246:
247: extern int ide0_need_reset;
248: extern int ide0_wait_interrupt;
249: extern int ide0_timeout;
250: extern int idel_need_reset;
251: extern int idel_wait_interrupt;
252: extern int idel_timeout;
253:
254: void irq_ide0(int, struct sigcontext *);
255: void ide0_timer(unsigned int);
256: void irq_idel(int, struct sigcontext *);
257: void idel_timer(unsigned int);
258:
259: void ide_error(struct ide *, int);
260: void ide_delay(void);
261: void ide_wait400ns(struct ide *);
262: int ide_ready(struct ide *);

```



**include/fiwix/ide.h**

Page 5/5

```
263: int ide_drvsel(struct ide *, int, int, unsigned char);
264: int ide_softreset(struct ide *);
265:
266: struct ide * get_ide_controller(__dev_t);
267: int get_ide_drive(__dev_t);
268:
269: int ide_open(struct inode *, struct fd *);
270: int ide_close(struct inode *, struct fd *);
271: int ide_read(__dev_t, __blk_t, char *, int);
272: int ide_write(__dev_t, __blk_t, char *, int);
273: int ide_ioctl(struct inode *, int, unsigned long int);
274:
275: void ide_init(void);
276:
277: #endif /* _FIWIX_IDE_H */
```

**include/fiwix/ide\_hd.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/ide_hd.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_IDE_HD_H
9: #define _FIWIX_IDE_HD_H
10:
11: #include <fiwix/types.h>
12:
13: #define IDE_HD_SECTSIZE          512      /* sector size (in bytes) */
14:
15: int ide_hd_open(struct inode *, struct fd *);
16: int ide_hd_close(struct inode *, struct fd *);
17: int ide_hd_read(__dev_t, __blk_t, char *, int);
18: int ide_hd_write(__dev_t, __blk_t, char *, int);
19: int ide_hd_ioctl(struct inode *, int, unsigned long int);
20:
21: int ide_hd_init(struct ide *, int);
22:
23: #endif /* _FIWIX_IDE_HD_H */
```

## include/fiwix/ioctl.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/ioctl.h
3:  */
4:
5: #ifndef _FIWIX_IOCTL_H
6: #define _FIWIX_IOCTL_H
7:
8: #define HDIO_GETGEO      0x0301      /* get device geometry */
9:
10: #define BLKRRPART       0x125F      /* re-read partition table */
11: #define BLKGETSIZE      0x1260      /* return device size */
12: #define BLKFLSBUF       0x1261      /* flush buffer cache */
13:
14: /* 0x54 is just a magic number to make these relatively unique ('T') */
15: #define TCGETS          0x5401
16: #define TCSETS         0x5402
17: #define TCSETSW        0x5403
18: #define TCSETSF        0x5404
19: #define TCGETA         0x5405
20: #define TCSETA         0x5406
21: #define TCSETAW        0x5407
22: #define TCSETAF        0x5408
23: #define TCSBRK         0x5409
24: #define TCXONC         0x540A
25: #define TCFLSH         0x540B
26: #define TIOCEXCL       0x540C
27: #define TIOCNXCL       0x540D
28: #define TIOCSCTTY      0x540E
29: #define TIOCGGRP       0x540F
30: #define TIOCSPGRP      0x5410
31: #define TIOCOUTQ       0x5411
32: #define TIOCSTI        0x5412
33: #define TIOCGWINSZ     0x5413
34: #define TIOCSWINSZ     0x5414
35: #define TIOCMGET       0x5415
36: #define TIOCMBIS       0x5416
37: #define TIOCMBIC       0x5417
38: #define TIOCMSET       0x5418
39: #define TIOCGSOFTCAR   0x5419
40: #define TIOCSSOFTCAR   0x541A
41: #define FIONREAD       0x541B
42: #define TIOCINQ        FIONREAD
43: #define TIOCLINUX      0x541C
44: #define TIOCCONS       0x541D
45: #define TIOCGSERIAL    0x541E
46: #define TIOCSSERIAL    0x541F
47: #define TIOCPKT        0x5420
48: #define FIONBIO        0x5421
49: #define TIOCNOTTY      0x5422
50: #define TIOCSETD       0x5423
51: #define TIOCGETD       0x5424
52: #define TCSBRKP        0x5425      /* Needed for POSIX tcsendbreak() */
53: #define TIOCTTYGSTRUCT 0x5426      /* For debugging only */
54: #define TIOCSBRK       0x5427      /* BSD compatibility */
55: #define TIOCCBRK       0x5428      /* BSD compatibility */
56: #define TIOCGSID       0x5429      /* Return the session ID of FD */
57: #define TIOCGPTN       _IOR('T',0x30, unsigned int) /* Get Pty Number (of pty-m
ux device) */
58: #define TIOCSPTLCK     _IOW('T',0x31, int) /* Lock/unlock Pty */
59:
60: #define FIONCLEX       0x5450      /* these numbers need to be adjusted. */
61: #define FIOCLEX        0x5451
62: #define FIOASYNC       0x5452
63: #define TIOCSERCONFIG  0x5453
64: #define TIOCSERGWILD   0x5454
65: #define TIOCSERSWILD   0x5455
66: #define TIOCGLCRMIOS   0x5456

```

**include/fiwix/ioctl.h**

Page 2/2

```
67: #define TIOCSLCKTRMIOS 0x5457
68: #define TIOCSESGSTRUCT 0x5458 /* For debugging only */
69: #define TIOCSEGETLSR 0x5459 /* Get line status register */
70: #define TIOCSEGETMULTI 0x545A /* Get multiport config */
71: #define TIOCSESETMULTI 0x545B /* Set multiport config */
72:
73: #define TIOCMWAIT 0x545C /* wait for a change on serial input line(s) */
74: #define TIOCGICOUNT 0x545D /* read serial port inline interrupt counts */
75: #define TIOCGHAYESESP 0x545E /* Get Hayes ESP configuration */
76: #define TIOCSHAYESESP 0x545F /* Set Hayes ESP configuration */
77:
78: /* Used for packet mode */
79: #define TIOCPKT_DATA 0
80: #define TIOCPKT_FLUSHREAD 1
81: #define TIOCPKT_FLUSHWRITE 2
82: #define TIOCPKT_STOP 4
83: #define TIOCPKT_START 8
84: #define TIOCPKT_NOSTOP 16
85: #define TIOCPKT_DOSTOP 32
86:
87: #define TIOCSER_TEMT 0x01 /* Transmitter physically empty */
88:
89: #endif /* _FIWIX_IOCTL_H */
```

## include/fiwix/kd.h

Page 1/3

```

1: /*
2:  * fiwix/include/fiwix/kd.h
3:  */
4:
5: #ifndef _LINUX_KD_H
6: #define _LINUX_KD_H
7:
8: /* Prefix 0x4B is 'K', to avoid collision with termios and vt */
9:
10: #define GIO_FONT      0x4B60 /* gets font in expanded form */
11: #define PIO_FONT      0x4B61 /* use font in expanded form */
12:
13: #define GIO_FONTX     0x4B6B /* get font using struct consolefontdesc */
14: #define PIO_FONTX     0x4B6C /* set font using struct consolefontdesc */
15: struct consolefontdesc {
16:     unsigned short int charcount; /* characters in font (256 or 512) */
17:     unsigned short int charheight; /* scan lines per character (1-32) */
18:     char *chardata; /* font data in expanded form */
19: };
20:
21: #define PIO_FONTRESET 0x4B6D /* reset to default font */
22:
23: #define GIO_CMAP      0x4B70 /* gets colour palette on VGA+ */
24: #define PIO_CMAP      0x4B71 /* sets colour palette on VGA+ */
25:
26: #define KIOCSOUND     0x4B2F /* start sound generation (0 for off) */
27: #define KDMKTONE      0x4B30 /* generate tone */
28:
29: #define KDGETLED      0x4B31 /* return current led state */
30: #define KDSETLED      0x4B32 /* set led state [lights, not flags] */
31: #define LED_SCR       0x01 /* scroll lock led */
32: #define LED_NUM       0x02 /* num lock led */
33: #define LED_CAP       0x04 /* caps lock led */
34:
35: #define KDGBKTYPE     0x4B33 /* get keyboard type */
36: #define KB_84         0x01
37: #define KB_101        0x02 /* this is what we always answer */
38: #define KB_OTHER      0x03
39:
40: #define KDADDIO       0x4B34 /* add i/o port as valid */
41: #define KDDELIO       0x4B35 /* del i/o port as valid */
42: #define KDENABIO     0x4B36 /* enable i/o to video board */
43: #define KDDISABIO    0x4B37 /* disable i/o to video board */
44:
45: #define KDSETMODE     0x4B3A /* set text/graphics mode */
46: #define KD_TEXT       0x00
47: #define KD_GRAPHICS   0x01
48: #define KD_TEXT0      0x02 /* obsolete */
49: #define KD_TEXT1      0x03 /* obsolete */
50: #define KDGETMODE     0x4B3B /* get current mode */
51:
52: #define KDMAPDISP     0x4B3C /* map display into address space */
53: #define KDUNMAPDISP  0x4B3D /* unmap display from address space */
54:
55: typedef char scrnmap_t;
56: #define E_TABSZ       256
57: #define GIO_SCRNMAP   0x4B40 /* get screen mapping from kernel */
58: #define PIO_SCRNMAP   0x4B41 /* put screen mapping table in kernel */
59: #define GIO_UNISCRNMAP 0x4B69 /* get full Unicode screen mapping */
60: #define PIO_UNISCRNMAP 0x4B6A /* set full Unicode screen mapping */
61:
62: #define GIO_UNIMAP    0x4B66 /* get unicode-to-font mapping from kernel */
63: struct unipair {
64:     unsigned short int unicode;
65:     unsigned short int fontpos;
66: };
67: struct unimapdesc {

```

## include/fiwix/kd.h

Page 2/3

```

68:         unsigned short int entry_ct;
69:         struct unipair *entries;
70: };
71: #define PIO_UNIMAP      0x4B67 /* put unicode-to-font mapping in kernel */
72: #define PIO_UNIMAPCLR   0x4B68 /* clear table, possibly advise hash algorithm */
/
73: struct unimapinit {
74:     unsigned short int advised_hashsize; /* 0 if no opinion */
75:     unsigned short int advised_hashstep; /* 0 if no opinion */
76:     unsigned short int advised_hashlevel; /* 0 if no opinion */
77: };
78:
79: #define UNI_DIRECT_BASE 0xF000 /* start of Direct Font Region */
80: #define UNI_DIRECT_MASK 0x01FF /* Direct Font Region bitmask */
81:
82: #define K_RAW           0x00
83: #define K_XLATE        0x01
84: #define K_MEDIUMRAW    0x02
85: #define K_UNICODE      0x03
86: #define KDGBKMODE     0x4B44 /* gets current keyboard mode */
87: #define KDSKMODE      0x4B45 /* sets current keyboard mode */
88:
89: #define K_METABIT      0x03
90: #define K_ESCPREFIX    0x04
91: #define KDGBKMETA     0x4B62 /* gets meta key handling mode */
92: #define KDSKBMETA     0x4B63 /* sets meta key handling mode */
93:
94: #define K_SCROLLLOCK   0x01
95: #define K_NUMLOCK      0x02
96: #define K_CAPSLOCK     0x04
97: #define KDGBKLED      0x4B64 /* get led flags (not lights) */
98: #define KDSKLED       0x4B65 /* set led flags (not lights) */
99:
100: struct kbentry {
101:     unsigned char kb_table;
102:     unsigned char kb_index;
103:     unsigned short int kb_value;
104: };
105: #define K_NORMTAB      0x00
106: #define K_SHIFTTAB    0x01
107: #define K_ALTTAB      0x02
108: #define K_ALTSHIFTTAB 0x03
109:
110: #define KDGBKENT      0x4B46 /* gets one entry in translation table */
111: #define KDSKENT       0x4B47 /* sets one entry in translation table */
112:
113: struct kbsentry {
114:     unsigned char kb_func;
115:     unsigned char kb_string[512];
116: };
117: #define KDGBKSENT     0x4B48 /* gets one function key string entry */
118: #define KDSKSENT     0x4B49 /* sets one function key string entry */
119:
120: struct kbdiacr {
121:     unsigned char diacr, base, result;
122: };
123: struct kbdiacrs {
124:     unsigned int kb_cnt; /* number of entries in following array */
125:     struct kbdiacr kbdiacr[256]; /* MAX_DIACR from keyboard.h */
126: };
127: #define KDGBKDIACR    0x4B4A /* read kernel accent table */
128: #define KDSKDIACR    0x4B4B /* write kernel accent table */
129:
130: struct kbkeycode {
131:     unsigned int scancode, keycode;
132: };
133: #define KDGETKEYCODE  0x4B4C /* read kernel keycode table entry */

```

## include/fiwix/kd.h

Page 3/3

```

134: #define KDSETKEYCODE    0x4B4D  /* write kernel keycode table entry */
135:
136: #define KDSIGACCEPT     0x4B4E  /* accept kbd generated signals */
137:
138: struct kbd_repeat {
139:     int delay;           /* in msec; <= 0: don't change */
140:     int rate;           /* in msec; <= 0: don't change */
141: };
142:
143: #define KDKBDREP        0x4B52  /* set keyboard delay/repeat rate;
144:                                * actually used values are returned */
145:
146: #define KDFONTOP        0x4B72  /* font operations */
147:
148: struct console_font_op {
149:     unsigned int op;     /* operation code KD_FONT_OP_* */
150:     unsigned int flags;  /* KD_FONT_FLAG_* */
151:     unsigned int width, height; /* font size */
152:     unsigned int charcount;
153:     unsigned char *data; /* font data with height fixed to 32 */
154: };
155:
156: #define KD_FONT_OP_SET    0      /* Set font */
157: #define KD_FONT_OP_GET    1      /* Get font */
158: #define KD_FONT_OP_SET_DEFAULT 2 /* Set font to default, data points to n
ame / NULL */
159: #define KD_FONT_OP_COPY    3      /* Copy from another console */
160:
161: #define KD_FONT_FLAG_DONT_RECALC 1 /* Don't recalculate hw charcell
size [compat] */
162:
163: /* note: 0x4B00-0x4B4E all have had a value at some time;
164:    don't reuse for the time being */
165: /* note: 0x4B60-0x4B6D, 0x4B70-0x4B72 used above */
166:
167: #endif /* _LINUX_KD_H */

```

## include/fiwix/kernel.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/kernel.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_KERNEL_H
9: #define _FIWIX_KERNEL_H
10:
11: #include <fiwix/limits.h>
12: #include <fiwix/i386elf.h>
13:
14: #define PANIC(format, args...)                                \
15: {                                                            \
16:     printk("\nPANIC: in %s()", __FUNCTION__);                \
17:     printk("\n");                                           \
18:     printk(format, ## args);                                \
19:     stop_kernel();                                         \
20: }
21:
22: #define CURRENT_TIME      (kstat.system_time)
23: #define CURRENT_TICKS    (kstat.ticks)
24: #define INIT_PROGRAM      "/sbin/init"
25:
26: extern char *init_argv[];
27: extern char *init_envp[];
28: extern char *init_args;
29:
30: extern Elf32_Shdr *symtab, *strtab;
31: extern unsigned int _last_data_addr;
32:
33: extern int _memsize;
34: extern int _extmemsize;
35: extern int _rootdev;
36: extern int _noramdisk;
37: extern int _ramdisksize;
38: extern char _rootfstype[10];
39: extern char _rootdevname[DEVNAME_MAX + 1];
40: extern char _initrd[DEVNAME_MAX + 1];
41: extern int _syscondev;
42:
43: extern int _cputype;
44: extern int _cpusignature;
45: extern int _cpuflags;
46: extern int _brandid;
47: extern char _vendorid[12];
48: extern char _brandstr[48];
49: extern unsigned int _tlbinfo_eax;
50: extern unsigned int _tlbinfo_ebx;
51: extern unsigned int _tlbinfo_ecx;
52: extern unsigned int _tlbinfo_edx;
53: extern char _etext[], _edata[], _end[];
54:
55: extern char cmdline[NAME_MAX + 1];
56:
57: struct kernel_stat {
58:     unsigned int cpu_user;           /* ticks in user-mode */
59:     unsigned int cpu_nice;          /* ticks in user-mode (with priority) */
60:     unsigned int cpu_system;        /* ticks in kernel-mode */
61:     unsigned int irq;               /* irq counter */
62:     unsigned int sirqs;              /* spurious irq counter */
63:     unsigned int ctxt;              /* context switches */
64:     unsigned int ticks;             /* ticks (1/HZths of sec) since boot */
65:     unsigned int system_time;       /* current system time (since the Epoch) */
66:     unsigned int boot_time;         /* boot time (since the Epoch) */

```



**include/fiwix/kernel.h**

Page 2/2

```
67:      int tz_minuteswest;          /* minutes west of GMT */
68:      int tz_dsttime;             /* type of DST correction */
69:      unsigned int uptime;        /* seconds since boot */
70:      unsigned int processes;     /* number of forks since boot */
71:      unsigned int physical_pages; /* physical memory in pages */
72:      unsigned int kernel_reserved; /* kernel memory reserved in KB */
73:      unsigned int physical_reserved; /* physical memory reserved in KB */
74:      unsigned int total_mem_pages; /* total memory in pages */
75:      unsigned int free_pages;    /* pages on free list (available) */
76:      unsigned int buffers;      /* memory used by buffers in KB */
77:      unsigned int cached;       /* memory used to cache file pages */
78:      unsigned int shared;       /* pages with count > 1 */
79:      unsigned long int random_seed; /* next random seed */
80: };
81: extern struct kernel_stat kstat;
82:
83: unsigned int get_last_boot_addr(unsigned int);
84: void start_kernel(unsigned long, unsigned long, unsigned int);
85: void stop_kernel(void);
86: void init_init(void);
87: void cpu_idle(void);
88:
89: #endif /* _FIWIX_KERNEL_H */
```

**include/fiwix/keyboard.h**

Page 1/3

```

1: /*
2:  * fiwix/include/fiwix/keyboard.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_KEYBOARD_H
9: #define _FIWIX_KEYBOARD_H
10:
11: #define KEYBOARD_IRQ    1
12:
13: #define NR_MODIFIERS    16      /* max. number of modifiers per keymap */
14: #define NR_SCODES      128     /* max. number of scancodes */
15: #define NR_DIACR       10
16:
17: #define SCRLBIT        0x01    /* scroll lock led */
18: #define NUMSBIT       0x02    /* num lock led */
19: #define CAPSBIT       0x04    /* caps lock led */
20:
21: #define C(ch)          ((ch) & 0x1F)
22: #define A(ch)          ((ch) | META_KEYS)
23: #define L(ch)          ((ch) | LETTER_KEYS)
24:
25: #define SLASH_NPAD     53
26:
27: #define EOENTER        96
28: #define RCTRL         97
29: #define EOSLASH       98
30: #define ALTGR         100
31: #define EOHOME        102
32: #define EOUP          103
33: #define EOPGUP        104
34: #define EOLEFT        105
35: #define EORIGHT       106
36: #define EOEND         107
37: #define EODOWN        108
38: #define EOPGDN        109
39: #define EOINS         110
40: #define EODEL         111
41:
42: #define MOD_BASE       0
43: #define MOD_SHIFT      1
44: #define MOD_ALTGR      2
45: #define MOD_CTRL       3
46: #define MOD_ALT        4
47: #define MOD_SHIFTL     5
48: #define MOD_SHIFTR     6
49: #define MOD_CTRLRL     7
50: #define MOD_CTRLR      8
51:
52: #define FN_KEYS        0x100
53: #define SPEC_KEYS      0x200
54: #define PAD_KEYS       0x300
55: #define DEAD_KEYS      0x400
56: #define CONS_KEYS      0x500
57: #define SHIFT_KEYS     0x700
58: #define META_KEYS      0x800
59: #define LETTER_KEYS    0xB00
60:
61: #define CR              (0x01 + SPEC_KEYS)
62: #define SCRL2          (0x02 + SPEC_KEYS)      /* SH_REGS (show registers) */
63: #define SCRL3          (0x03 + SPEC_KEYS)      /* SH_MEM (show memory) */
64: #define SCRL4          (0x04 + SPEC_KEYS)      /* SH_STAT (show status) */
65: #define CAPS           (0x07 + SPEC_KEYS)
66: #define NUMS           (0x08 + SPEC_KEYS)
67: #define SCRL           (0x09 + SPEC_KEYS)

```

## include/fiwix/keyboard.h

Page 2/3

```

68:
69: #define INS                (0x00 + PAD_KEYS)
70: #define END                (0x01 + PAD_KEYS)
71: #define DOWN              (0x02 + PAD_KEYS)
72: #define PGDN              (0x03 + PAD_KEYS)
73: #define LEFT              (0x04 + PAD_KEYS)
74: #define MID                (0x05 + PAD_KEYS)
75: #define RIGHT             (0x06 + PAD_KEYS)
76: #define HOME              (0x07 + PAD_KEYS)
77: #define UP                (0x08 + PAD_KEYS)
78: #define PGUP              (0x09 + PAD_KEYS)
79: #define PLUS              (0x0A + PAD_KEYS)
80: #define MINUS             (0x0B + PAD_KEYS)
81: #define ASTSK             (0x0C + PAD_KEYS)
82: #define SLASH             (0x0D + PAD_KEYS)
83: #define ENTER             (0x0E + PAD_KEYS)
84: #define DEL               (0x10 + PAD_KEYS)
85:
86: #define GRAVE              (0x00 + DEAD_KEYS)
87: #define ACUTE              (0x01 + DEAD_KEYS)
88: #define CIRCUM             (0x02 + DEAD_KEYS)
89: #define DIERE              (0x04 + DEAD_KEYS)
90:
91: #define SHIFT              (0x00 + SHIFT_KEYS)
92: #define CTRL              (0x02 + SHIFT_KEYS)
93: #define ALT                (0x03 + SHIFT_KEYS)
94: #define LSHIFT            (0x04 + SHIFT_KEYS)
95: #define RSHIFT            (0x05 + SHIFT_KEYS)
96: #define LCTRL             (0x06 + SHIFT_KEYS)
97:
98: #define F1                 (0x00 + FN_KEYS)
99: #define F2                 (0x01 + FN_KEYS)
100: #define F3                 (0x02 + FN_KEYS)
101: #define F4                 (0x03 + FN_KEYS)
102: #define F5                 (0x04 + FN_KEYS)
103: #define F6                 (0x05 + FN_KEYS)
104: #define F7                 (0x06 + FN_KEYS)
105: #define F8                 (0x07 + FN_KEYS)
106: #define F9                 (0x08 + FN_KEYS)
107: #define F10                (0x09 + FN_KEYS)
108: #define F11                (0x0A + FN_KEYS)
109: #define F12                (0x0B + FN_KEYS)
110:
111: #define SF1                (0x0A + FN_KEYS)
112: #define SF2                (0x0B + FN_KEYS)
113: #define SF3                (0x0C + FN_KEYS)
114: #define SF4                (0x0D + FN_KEYS)
115: #define SF5                (0x0E + FN_KEYS)
116: #define SF6                (0x0F + FN_KEYS)
117: #define SF7                (0x10 + FN_KEYS)
118: #define SF8                (0x11 + FN_KEYS)
119: #define SF9                (0x12 + FN_KEYS)
120: #define SF10               (0x13 + FN_KEYS)
121: #define SF11               (0x0A + SHIFT)
122: #define SF12               (0x0B + SHIFT)
123:
124: #define AF1                (0x00 + CONS_KEYS)
125: #define AF2                (0x01 + CONS_KEYS)
126: #define AF3                (0x02 + CONS_KEYS)
127: #define AF4                (0x03 + CONS_KEYS)
128: #define AF5                (0x04 + CONS_KEYS)
129: #define AF6                (0x05 + CONS_KEYS)
130: #define AF7                (0x06 + CONS_KEYS)
131: #define AF8                (0x07 + CONS_KEYS)
132: #define AF9                (0x08 + CONS_KEYS)
133: #define AF10               (0x09 + CONS_KEYS)
134: #define AF11               (0x0A + CONS_KEYS)

```

**include/fiwix/keyboard.h**

Page 3/3

```
135: #define AF12                (0x0B + CONS_KEYS)
136:
137: #ifdef __KERNEL__
138:
139: #include <fiwix/types.h>
140: #include <fiwix/sigcontext.h>
141:
142: struct diacritic {
143:     unsigned char letter;
144:     unsigned char code;
145: };
146:
147: extern __key_t keymap[NR_MODIFIERS * NR_SCODES];
148:
149: void set_leds(unsigned char);
150: void irq_keyboard(int num, struct sigcontext *);
151: void irq_keyboard_bh(void);
152: void keyboard_init(void);
153:
154: #endif /* __KERNEL__ */
155:
156: #endif /* _FIWIX_KEYBOARD_H */
```

## include/fiwix/kparms.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/kparms.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_KPARMS_H
9: #define _FIWIX_KPARMS_H
10:
11: #define CMDL_ARG_LEN    100    /* max length of cmdline argument */
12: #define CMDL_NUM_VALUES 30    /* max values of cmdline parameter */
13:
14: struct kparms {
15:     char *name;
16:     char *value[CMDL_NUM_VALUES];
17:     unsigned int sysval[CMDL_NUM_VALUES];
18: };
19:
20: static struct kparms parm_table[] = {
21:     { "root=",
22:       { "/dev/ram0", "/dev/fd0", "/dev/fd1",
23:         "/dev/hda", "/dev/hda1", "/dev/hda2", "/dev/hda3", "/dev/hda4",
24:         "/dev/hdb", "/dev/hdb1", "/dev/hdb2", "/dev/hdb3", "/dev/hdb4",
25:         "/dev/hdc", "/dev/hdc1", "/dev/hdc2", "/dev/hdc3", "/dev/hdc4",
26:         "/dev/hdd", "/dev/hdd1", "/dev/hdd2", "/dev/hdd3", "/dev/hdd4",
27:       },
28:       { 0x100, 0x200, 0x201,
29:         0x300, 0x301, 0x302, 0x303, 0x304,
30:         0x340, 0x341, 0x342, 0x343, 0x344,
31:         0x1600, 0x1601, 0x1602, 0x1603, 0x1604,
32:         0x1640, 0x1641, 0x1642, 0x1643, 0x1644,
33:       }
34:     },
35:     { "noramdisk",
36:       { NULL },
37:       { NULL },
38:     },
39:     { "ramdisksize=",
40:       { NULL },
41:       { NULL },
42:     },
43:     { "initrd=",
44:       { NULL },
45:       { NULL },
46:     },
47:     { "rootfstype=",
48:       { "minix", "ext2", "iso9660" },
49:       { 0, 0 }
50:     },
51:     { "console=",
52:       { "/dev/tty1", "/dev/tty2", "/dev/tty3", "/dev/tty4", "/dev/tty5",
53:         "/dev/tty6", "/dev/tty7", "/dev/tty8", "/dev/tty9", "/dev/tty10",
54:         "/dev/tty11", "/dev/tty12", "/dev/ttyS0", "/dev/ttyS1",
55:         "/dev/ttyS2", "/dev/ttyS3"
56:       },
57:       { 0x401, 0x402, 0x403, 0x404, 0x405,
58:         0x406, 0x407, 0x408, 0x409, 0x410,
59:         0x411, 0x412, 0x440, 0x441, 0x442,
60:         0x443
61:       }
62:     },
63:     { NULL }
64: };
65: };
66:
67: #endif /* _FIWIX_KPARMS_H */

```

**include/fiwix/limits.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/limits.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_LIMITS_H
9: #define _FIWIX_LIMITS_H
10:
11: #define DEVNAME_MAX      50      /* device name length in mount table */
12: #define ARG_MAX          32      /* length (in pages) of argv+env in 'execve' */
13: #define CHILD_MAX        64      /* simultaneous processes per real user ID */
14: #define LINK_MAX         255     /* maximum number of links to a file */
15: #define MAX_CANON        255     /* bytes in a terminal canonical input queue */
16: #define MAX_INPUT        255     /* bytes for which space will be available in a
17:                                     terminal input queue */
18: #define NGROUPS_MAX      32      /* simultaneous supplementary group IDs */
19: #define OPEN_MAX         256     /* files one process can have opened at once */
20: #define FD_SETSIZE       OPEN_MAX /* descriptors that a process may examine with
21:                                     'pselect' or 'select' */
22: #define NAME_MAX         255     /* bytes in a filename */
23: #define PATH_MAX         1024    /* bytes in a pathname */
24: #define PIPE_BUF         4096    /* bytes than can be written atomically to a
25:                                     pipe */
26:
27: #endif /* _FIWIX_LIMITS_H */
```

**include/fiwix/locks.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/locks.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_LOCKS_H
9: #define _FIWIX_LOCKS_H
10:
11: #include <fiwix/config.h>
12: #include <fiwix/fs.h>
13: #include <fiwix/fcntl.h>
14:
15: struct flock_file {
16:     struct inode *inode;    /* file */
17:     unsigned char type;    /* type of lock */
18:     struct proc *proc;     /* owner */
19: };
20:
21: struct flock_file flock_file_table[NR_FLOCKS];
22:
23: int posix_lock(int, int, struct flock *);
24:
25: void flock_release_inode(struct inode *);
26: int flock_inode(struct inode *, int);
27: void flock_init(void);
28:
29: #endif /* _FIWIX_LOCKS_H */
```

## include/fiwix/lp.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/lp.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_LP_H
9: #define _FIWIX_LP_H
10:
11: #include <fiwix/fs.h>
12:
13: #define LP_MAJOR          6          /* major number for /dev/lp[n] */
14: #define LP_MINORS        1
15:
16: /*#define LP0_ADDR        0x3BC */
17: #define LP0_ADDR          0x378
18: /*#define LP2_ADDR        0x278 */
19:
20: #define LP_STAT_ERR      0x08      /* printer error */
21: #define LP_STAT_SEL      0x10      /* select in */
22: #define LP_STAT_PE       0x20      /* paper empty or no paper */
23: #define LP_STAT_ACK      0x40      /* ack */
24: #define LP_STAT_BUS      0x80      /* printer busy */
25:
26: #define LP_CTRL_STR      0x01      /* strobe */
27: #define LP_CTRL_AUT      0x02      /* auto line feed */
28: #define LP_CTRL_INI      0x04      /* initialize printer (reset) */
29: #define LP_CTRL_SEL      0x08      /* select printer */
30: #define LP_CTRL_IRQ      0x10      /* enable IRQ */
31: #define LP_CTRL_BID      0x20      /* bidireccional (on PS/2 ports) */
32:
33: #define LP_RDY_RETR      100       /* retries before timeout */
34:
35: struct lp {
36:     int data;             /* data port address */
37:     int stat;            /* status port address */
38:     int ctrl;           /* control port address */
39:     char flags;         /* flags */
40: };
41:
42: int lp_open(struct inode *, struct fd *);
43: int lp_close(struct inode *, struct fd *);
44: int lp_write(struct inode *, struct fd *, const char *, __size_t);
45:
46: void lp_init(void);
47:
48: #endif /* _FIWIX_LP_H */

```



## include/fiwix/memdev.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/memdev.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_MEMDEV_H
9: #define _FIWIX_MEMDEV_H
10:
11: #include <fiwix/fs.h>
12:
13: #define MEMDEV_MAJOR    1        /* major number */
14: #define MEMDEV_MINORS  5        /* number of supported minors */
15:
16: #define MEMDEV_MEM      1        /* minor for /dev/mem */
17: #define MEMDEV_KMEM     2        /* minor for /dev/kmem */
18: #define MEMDEV_NULL     3        /* minor for /dev/null */
19: #define MEMDEV_ZERO     5        /* minor for /dev/zero */
20: #define MEMDEV_RANDOM   8        /* minor for /dev/random */
21: #define MEMDEV_URANDOM  9        /* minor for /dev/urandom */
22:
23: int mem_open(struct inode *, struct fd *);
24: int mem_close(struct inode *, struct fd *);
25: int mem_read(struct inode *, struct fd *, char *, __size_t);
26: int mem_write(struct inode *, struct fd *, const char *, __size_t);
27: int mem_lseek(struct inode *, __off_t);
28:
29: int kmem_open(struct inode *, struct fd *);
30: int kmem_close(struct inode *, struct fd *);
31: int kmem_read(struct inode *, struct fd *, char *, __size_t);
32: int kmem_write(struct inode *, struct fd *, const char *, __size_t);
33: int kmem_lseek(struct inode *, __off_t);
34:
35: int null_open(struct inode *, struct fd *);
36: int null_close(struct inode *, struct fd *);
37: int null_read(struct inode *, struct fd *, char *, __size_t);
38: int null_write(struct inode *, struct fd *, const char *, __size_t);
39: int null_lseek(struct inode *, __off_t);
40:
41: int zero_open(struct inode *, struct fd *);
42: int zero_close(struct inode *, struct fd *);
43: int zero_read(struct inode *, struct fd *, char *, __size_t);
44: int zero_write(struct inode *, struct fd *, const char *, __size_t);
45: int zero_lseek(struct inode *, __off_t);
46:
47: int urandom_open(struct inode *, struct fd *);
48: int urandom_close(struct inode *, struct fd *);
49: int urandom_read(struct inode *, struct fd *, char *, __size_t);
50: int urandom_write(struct inode *, struct fd *, const char *, __size_t);
51: int urandom_lseek(struct inode *, __off_t);
52:
53: int memdev_open(struct inode *, struct fd *);
54: int mem_mmap(struct inode *, struct vma *);
55: void memdev_init(void);
56:
57: #endif /* _FIWIX_MEMDEV_H */

```

## include/fiwix/mman.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/mman.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_MMAN_H
9: #define _FIWIX_MMAN_H
10:
11: #include <fiwix/fs.h>
12:
13: #define PROT_READ          0x1          /* page can be read */
14: #define PROT_WRITE        0x2          /* page can be written */
15: #define PROT_EXEC         0x4          /* page can be executed */
16: #define PROT_NONE         0x0          /* page cannot be accessed */
17:
18: #define MAP_SHARED        0x01         /* share changes */
19: #define MAP_PRIVATE       0x02         /* changes are private */
20: #define MAP_TYPE          0x0f         /* mask for type of mapping */
21: #define MAP_FIXED         0x10         /* interpret address exactly */
22: #define MAP_ANONYMOUS    0x20         /* don't use the file descriptor */
23:
24: #define MAP_GROWSDOWN     0x0100       /* stack-like segment */
25: #define MAP_DENYWRITE     0x0800       /* -ETXTBSY */
26: #define MAP_EXECUTABLE    0x1000       /* mark it as a executable */
27: #define MAP_LOCKED        0x2000       /* pages are locked */
28:
29: #define ZERO_PAGE         0x80000000    /* this page must be zero-filled */
30:
31: #define MS_ASYNC          1            /* sync memory asynchronously */
32: #define MS_INVALIDATE    2            /* invalidate the caches */
33: #define MS_SYNC           4            /* synchronous memory sync */
34:
35: #define MCL_CURRENT       1            /* lock all current mappings */
36: #define MCL_FUTURE       2            /* lock all future mappings */
37:
38: /* compatibility flags */
39: #define MAP_ANON          MAP_ANONYMOUS
40: #define MAP_FILE          0
41:
42: struct mmap {
43:     unsigned int start;
44:     unsigned int length;
45:     unsigned int prot;
46:     unsigned int flags;
47:     int fd;
48:     unsigned int offset;
49: };
50:
51: void show_vma_regions(struct proc *);
52: void release_binary(void);
53: struct vma * find_vma_region(unsigned int);
54: int expand_heap(unsigned int);
55: int do_mmap(struct inode *, unsigned int, unsigned int, unsigned int, unsigned i
nt, unsigned int, char, char);
56: int do_munmap(unsigned int, __size_t);
57: int do_mprotect(struct vma *, unsigned int, __size_t, int);
58:
59: #endif /* _FIWIX_MMAN_H */

```

## include/fiwix/mm.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/mm.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_MEMORY_H
9: #define _FIWIX_MEMORY_H
10:
11: #include <fiwix/types.h>
12: #include <fiwix/const.h>
13: #include <fiwix/process.h>
14:
15: /*
16:  * Convert only from physical to virtual the addresses below KERNEL_BASE_ADDR
17:  * (formerly 0xC0000000).
18:  */
19: #define P2V(addr)          (addr < KERNEL_BASE_ADDR ? addr + KERNEL_BASE_AD
DR : addr)
20:
21: #define V2P(addr)          (addr - KERNEL_BASE_ADDR)
22:
23: #define PAGE_SIZE          4096
24: #define PAGE_SHIFT         0x0C
25: #define PAGE_MASK          ~(PAGE_SIZE - 1)          /* 0xFFFFF000 */
26: #define PAGE_ALIGN(addr)   ((addr) + (PAGE_SIZE - 1)) & PAGE_MASK)
27: #define PT_ENTRIES         (PAGE_SIZE / sizeof(unsigned int))
28: #define PD_ENTRIES         (PAGE_SIZE / sizeof(unsigned int))
29:
30: #define PAGE_PRESENT       0x001    /* Present */
31: #define PAGE_RW            0x002    /* Read/Write */
32: #define PAGE_USER          0x004    /* User */
33:
34: #define PAGE_RESERVED      0x100    /* kernel, BIOS address, ... */
35: #define PAGE_COW           0x200    /* marked for Copy-On-Write */
36:
37: #define PFAULT_V           0x01     /* protection violation */
38: #define PFAULT_W           0x02     /* during write */
39: #define PFAULT_U           0x04     /* in user mode */
40:
41: #define GET_PGDIR(address) ((unsigned int)((address) >> 22) & 0x3FF)
42: #define GET_PGTBL(address) ((unsigned int)((address) >> 12) & 0x3FF)
43:
44: #define P_TEXT  1          /* text section */
45: #define P_DATA  2          /* data section */
46: #define P_BSS   3          /* BSS section */
47: #define P_HEAP  4          /* heap section (sys_brk()) */
48: #define P_STACK 5          /* stack section */
49: #define P_MMAP  6          /* mmap() section */
50:
51: struct page {
52:     unsigned int page;          /* page number */
53:     unsigned int count;        /* usage counter */
54:     unsigned int flags;
55:     unsigned char locked;      /* 1 = locked */
56:     struct inode *inode;
57:     __off_t offset;           /* file offset */
58:     char *data;               /* page contents */
59:     struct page *prev_hash;
60:     struct page *next_hash;
61:     struct page *prev_free;
62:     struct page *next_free;
63: };
64:
65: extern struct page *page_table;
66: extern struct page **page_hash_table;

```

**include/fiwix/mm.h**

Page 2/2

```
67:
68: /* values to be determined during system startup */
69: extern unsigned int page_table_size; /* size in bytes */
70: extern unsigned int page_hash_table_size; /* size in bytes */
71:
72: extern unsigned int *kpage_dir;
73: extern unsigned int *kpage_table;
74:
75: /* alloc.c */
76: unsigned int kmalloc(void);
77: void kfree(unsigned int);
78:
79: /* page.c */
80: void page_lock(struct page *);
81: void page_unlock(struct page *);
82: struct page * get_free_page(void);
83: struct page * search_page_hash(struct inode *, __off_t);
84: void release_page(unsigned int);
85: int valid_page(unsigned int);
86: void update_page_cache(struct inode *, __off_t, const char *, int);
87: int write_page(struct page *, struct inode *, __off_t, unsigned int);
88: int bread_page(struct page *, struct inode *, __off_t, char, char);
89: int file_read(struct inode *, struct fd *, char *, __size_t);
90: void page_init(unsigned int);
91:
92: /* memory.c */
93: void bss_init(void);
94: unsigned int setup_minmem(void);
95: unsigned int get_mapped_addr(struct proc *, unsigned int);
96: int clone_pages(struct proc *);
97: int free_page_tables(struct proc *);
98: unsigned int map_page(struct proc *, unsigned int, unsigned int, unsigned int);
99: int unmap_page(unsigned int);
100: void mem_init(void);
101: void mem_stats(void);
102:
103: /* swapper.c */
104: int kswapd(void);
105:
106: #endif /* _FIWIX_MEMORY_H */
```

**include/fiwix/multiboot.h**

Page 1/2

```

1:  /* multiboot.h - the header for Multiboot */
2:  /* Copyright (C) 1999, 2001 Free Software Foundation, Inc.
3:
4:  This program is free software; you can redistribute it and/or modify
5:  it under the terms of the GNU General Public License as published by
6:  the Free Software Foundation; either version 2 of the License, or
7:  (at your option) any later version.
8:
9:  This program is distributed in the hope that it will be useful,
10: but WITHOUT ANY WARRANTY; without even the implied warranty of
11: MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
12: GNU General Public License for more details.
13:
14: You should have received a copy of the GNU General Public License
15: along with this program; if not, write to the Free Software
16: Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA. */
17:
18: /* Macros. */
19:
20: #ifndef _FIWIX_MULTIBOOT_H
21: #define _FIWIX_MULTIBOOT_H
22:
23: /* The magic number for the Multiboot header. */
24: #define MULTIBOOT_HEADER_MAGIC      0x1BADB002
25:
26: /* The flags for the Multiboot header. */
27: #ifdef __ELF__
28: # define MULTIBOOT_HEADER_FLAGS    0x00000003
29: #else
30: # define MULTIBOOT_HEADER_FLAGS    0x00010003
31: #endif
32:
33: /* The magic number passed by a Multiboot-compliant boot loader. */
34: #define MULTIBOOT_BOOTLOADER_MAGIC  0x2BADB002
35:
36: /* The size of our stack (16KB). */
37: #define STACK_SIZE                   0x4000
38:
39: /* C symbol format. HAVE_ASM_USCORE is defined by configure. */
40: #ifdef HAVE_ASM_USCORE
41: # define EXT_C(sym)                 _ ## sym
42: #else
43: # define EXT_C(sym)                 sym
44: #endif
45:
46: #ifndef ASM
47: /* Do not include here in boot.S. */
48:
49: /* Types. */
50:
51: /* The Multiboot header. */
52: typedef struct multiboot_header
53: {
54:     unsigned long magic;
55:     unsigned long flags;
56:     unsigned long checksum;
57:     unsigned long header_addr;
58:     unsigned long load_addr;
59:     unsigned long load_end_addr;
60:     unsigned long bss_end_addr;
61:     unsigned long entry_addr;
62: } multiboot_header_t;
63:
64: /* The symbol table for a.out. */
65: typedef struct aout_symbol_table
66: {
67:     unsigned long tabsize;

```

**include/fiwix/multiboot.h**

Page 2/2

```
68:  unsigned long  strsize;
69:  unsigned long  addr;
70:  unsigned long  reserved;
71: } aout_symbol_table_t;
72:
73: /* The section header table for ELF. */
74: typedef struct elf_section_header_table
75: {
76:  unsigned long  num;
77:  unsigned long  size;
78:  unsigned long  addr;
79:  unsigned long  shndx;
80: } elf_section_header_table_t;
81:
82: /* The Multiboot information. */
83: typedef struct multiboot_info
84: {
85:  unsigned long  flags;
86:  unsigned long  mem_lower;
87:  unsigned long  mem_upper;
88:  unsigned long  boot_device;
89:  unsigned long  cmdline;
90:  unsigned long  mods_count;
91:  unsigned long  mods_addr;
92:  union
93:  {
94:    aout_symbol_table_t aout_sym;
95:    elf_section_header_table_t elf_sec;
96:  } u;
97:  unsigned long  mmap_length;
98:  unsigned long  mmap_addr;
99: } multiboot_info_t;
100:
101: /* The module structure. */
102: typedef struct module
103: {
104:  unsigned long  mod_start;
105:  unsigned long  mod_end;
106:  unsigned long  string;
107:  unsigned long  reserved;
108: } module_t;
109:
110: /* The memory map. Be careful that the offset 0 is base_addr_low
111:    but no size. */
112: typedef struct memory_map
113: {
114:  unsigned long  size;
115:  unsigned long  base_addr_low;
116:  unsigned long  base_addr_high;
117:  unsigned long  length_low;
118:  unsigned long  length_high;
119:  unsigned long  type;
120: } memory_map_t;
121:
122: #endif /* ! ASM */
123:
124: #endif /* _FIWIX_MULTIBOOT_H */
```

**include/fiwix/part.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/part.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_PART_H
9: #define _FIWIX_PART_H
10:
11: #define PARTITION_BLOCK          0
12: #define NR_PARTITIONS           4      /* partitions in the MBR */
13: #define MBR_CODE_SIZE           446
14: #define ACTIVE_PART              0x80
15:
16: struct hd_geometry {
17:     unsigned char heads;
18:     unsigned char sectors;
19:     unsigned short int cylinders;
20:     unsigned long int start;
21: };
22:
23: struct partition {
24:     unsigned char status;
25:     unsigned char head;
26:     unsigned char sector;
27:     unsigned char cyl;
28:     unsigned char type;
29:     unsigned char endhead;
30:     unsigned char endsector;
31:     unsigned char endcyl;
32:     unsigned int startsect;
33:     unsigned int nr_sects;
34: };
35:
36: int read_msdos_partition(__dev_t, struct partition *);
37:
38: #endif /* _FIWIX_PART_H */
```

**include/fiwix/pic.h**

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/pic.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_PIC_H
9: #define _FIWIX_PIC_H
10:
11: #include <fiwix/sigcontext.h>
12:
13: #define NR_IRQS          16          /* hardware interrupts */
14: #define PIC_MASTER      0x20        /* I/O base address for master PIC */
15: #define PIC_SLAVE       0xA0        /* I/O base address for slave PIC */
16:
17: #define DATA           0x01        /* offset to data port */
18: #define EOI              0x20        /* End-Of-Interrupt command code */
19:
20: /* Inicialization Command Words */
21: #define ICW1_RESET      0x11        /* ICW1_INIT + ICW1_ICW4 */
22: #define CASCADE_IRQ     0x02
23: #define ICW4_8086EOI    0x01
24:
25: #define PIC_READ_IRR    0x0A        /* OCW3 irq ready */
26: #define PIC_READ_ISR    0x0B        /* OCW3 irq service */
27:
28: /* Operational Command Words */
29: #define OCW1            0xFF        /* mask (disable) all IRQs */
30:
31: struct interrupt {
32:     unsigned int ticks;
33:     char *name;
34:     void (*handler)(int, struct sigcontext *);
35:     struct interrupt *next;
36: };
37: extern struct interrupt *irq_table[NR_IRQS];
38: extern unsigned int intr_count;
39:
40:
41: #define BH_ACTIVE        0x01
42:
43: struct bh {
44:     int flags;
45:     void (*fn)(void);
46:     struct bh *next;
47: };
48:
49: void add_bh(struct bh *);
50: void enable_irq(int);
51: void disable_irq(int);
52: int register_irq(int, struct interrupt *);
53: int unregister_irq(int, struct interrupt *);
54: void irq_handler(int, struct sigcontext);
55: void do_bh(void);
56: void pic_init(void);
57:
58: #endif /* _FIWIX_PIC_H */

```



## include/fiwix/pit.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/pit.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_PIT_H
9: #define _FIWIX_PIT_H
10:
11: /* Intel 8253/82c54 Programmable Interval Timer */
12:
13: #define OSCIL          1193182 /* oscillator frequency */
14:
15: #define MODEREG        0x43 /* mode/command register (w) */
16: #define CHANNEL0      0x40 /* channel 0 data port (rw) */
17: #define CHANNEL1      0x41 /* channel 1 data port (rw) */
18: #define CHANNEL2      0x42 /* channel 2 data port (rw) */
19:
20: #define BINARY_CTR    0x00 /* 16bit binary mode counter */
21: #define TERM_COUNT    0x00 /* mode 0 (Terminal Count) */
22: #define RATE_GEN      0x04 /* mode 2 (Rate Generator) */
23: #define SQUARE_WAVE  0x06 /* mode 3 (Square Wave) */
24: #define LSB_MSB       0x30 /* LSB then MSB */
25: #define SEL_CHAN0     0x00 /* select channel 0 */
26: #define SEL_CHAN2     0x80 /* select channel 2 */
27:
28: /*
29:  * PS/2 System Control Port B
30:  * -----
31:  * bit 7 -> IRQ=1, 0=reset
32:  * bit 6 -> reserved
33:  * bit 5 -> reserved
34:  * bit 4 -> reserved
35:  * bit 3 -> channel check enable
36:  * bit 2 -> parity check enable
37:  * bit 1 -> speaker data enable
38:  * bit 0 -> timer 2 gate to speaker enable
39:  */
40: #define PS2_SYCTRL_B  0x61 /* PS/2 system control port B (write) */
41:
42: #define ENABLE_TMR2G  0x01 /* timer 2 gate to speaker enable */
43: #define ENABLE_SDATA  0x02 /* speaker data enable */
44:
45: #define BEEP_FREQ     900 /* 900Hz */
46:
47: void pit_beep_on(void);
48: void pit_beep_off(unsigned int);
49: void pit_init(unsigned short int);
50:
51: #endif /* _FIWIX_PIT_H */

```

## include/fiwix/process.h

Page 1/3

```

1: /*
2:  * fiwix/include/fiwix/process.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_PROCESS_H
9: #define _FIWIX_PROCESS_H
10:
11: struct vma {
12:     unsigned int start;
13:     unsigned int end;
14:     char prot;                /* PROT_READ, PROT_WRITE, ... */
15:     unsigned int flags;      /* MAP_SHARED, MAP_PRIVATE, ... */
16:     unsigned int offset;
17:     char s_type;             /* section type (P_TEXT, P_DATA, ...) */
18:     struct inode *inode;    /* file inode */
19:     char o_mode;            /* open mode (O_RDONLY, O_RDWR, ...) */
20: };
21:
22: #include <fiwix/config.h>
23: #include <fiwix/types.h>
24: #include <fiwix/signal.h>
25: #include <fiwix/limits.h>
26: #include <fiwix/sigcontext.h>
27: #include <fiwix/time.h>
28: #include <fiwix/resource.h>
29: #include <fiwix/tty.h>
30:
31: #define IDLE          0          /* PID of idle */
32: #define INIT          1          /* PID of /sbin/init */
33: #define SAFE_SLOTS    2          /* process slots reserved for root */
34: #define SLOT(p)       ((p) - (&proc_table[0]))
35:
36: /* bits in flags */
37: #define PF_KPROC       0x00000001 /* kernel internal process */
38: #define PF_PEXEC       0x00000002 /* has performed a sys_execve() */
39: #define PF_USERREAL    0x00000004 /* use real UID in permission checks */
40:
41: #define MMAP_START     0x40000000 /* mmap()'s start at 1GB */
42: #define IS_SUPERUSER   (current->euid == 0)
43:
44: #define IO_BITMAP_SIZE 32        /* 32 * 32bit = 1024 = 0x3FF */
45:                                     /* 2048 * 32bit = 65536 = 0xFFFF */
46:
47: #define PG_LEADER(p)   ((p)->pid == (p)->pgid)
48: #define SESS_LEADER(p) ((p)->pid == (p)->pgid && (p)->pid == (p)->sid)
49:
50: #define FOR_EACH_PROCESS(p)      for(p = &proc_table[INIT]; p; p = p->next)
51:
52: /* value to be determined during system startup */
53: extern unsigned int proc_table_size; /* size in bytes */
54:
55: extern char any_key_to_reboot;
56: extern int nr_processes;
57: extern __pid_t lastpid;
58:
59: struct binargs {
60:     unsigned int page[ARG_MAX];
61:     int argc;
62:     int argv_len;
63:     int envc;
64:     int envp_len;
65:     int offset;
66: };
67:

```

## include/fiwix/process.h

Page 2/3

```

68: /* Intel 386 Task Switch State */
69: struct i386tss {
70:     unsigned int prev_tss;
71:     unsigned int esp0;
72:     unsigned int ss0;
73:     unsigned int esp1;
74:     unsigned int ssl;
75:     unsigned int esp2;
76:     unsigned int ss2;
77:     unsigned int cr3;
78:     unsigned int eip;
79:     unsigned int eflags;
80:     unsigned int eax;
81:     unsigned int ecx;
82:     unsigned int edx;
83:     unsigned int ebx;
84:     unsigned int esp;
85:     unsigned int ebp;
86:     unsigned int esi;
87:     unsigned int edi;
88:     unsigned int es;
89:     unsigned int cs;
90:     unsigned int ss;
91:     unsigned int ds;
92:     unsigned int fs;
93:     unsigned int gs;
94:     unsigned int ldt;
95:     unsigned short int debug_trap;
96:     unsigned short int io_bitmap_addr;
97: };
98:
99: struct proc {
100:     struct i386tss tss;
101:     unsigned int io_bitmap[IO_BITMAP_SIZE + 1];
102:     __pid_t pid; /* process ID */
103:     __pid_t ppid; /* parent process ID */
104:     __pid_t pgid; /* process group ID */
105:     __pid_t sid; /* session ID */
106:     int flags;
107:     int groups[NGROUPS_MAX];
108:     int children; /* number of children */
109:     struct tty *ctty; /* controlling terminal */
110:     int state; /* process state */
111:     int priority;
112:     int cpu_count; /* time of process running */
113:     __time_t start_time;
114:     int exit_code;
115:     void *sleep_address;
116:     unsigned short int uid; /* real user ID */
117:     unsigned short int gid; /* real group ID */
118:     unsigned short int euid; /* effective user ID */
119:     unsigned short int egid; /* effective group ID */
120:     unsigned short int suid; /* saved user ID */
121:     unsigned short int sgid; /* saved group ID */
122:     unsigned short int fd[OPEN_MAX];
123:     unsigned char fd_flags[OPEN_MAX];
124:     struct inode *root;
125:     struct inode *pwd; /* process working directory */
126:     unsigned int entry_address;
127:     char argv0[NAME_MAX + 1];
128:     char **argv;
129:     char **envp;
130:     char pidstr[5]; /* pid number converted to string */
131:     struct vma vma[VMA_REGIONS]; /* virtual memory-map addresses */
132:     unsigned int brk_lower; /* lower limit of the heap section */
133:     unsigned int brk; /* current limit of the heap */
134:     __sigset_t sigpending;

```

**include/fiwix/process.h**

Page 3/3

```

135:     __sigset_t sigblocked;
136:     __sigset_t sigexecuting;
137:     struct sigaction sigaction[NSIG];
138:     struct sigcontext sc[NSIG];      /* each signal has its own context */
139:     unsigned int sp;                 /* current process' stack frame */
140:     struct rusage usage;             /* process resource usage */
141:     struct rusage cusage;           /* children resource usage */
142:     unsigned long int it_real_interval, it_real_value;
143:     unsigned long int it_virt_interval, it_virt_value;
144:     unsigned long int it_prof_interval, it_prof_value;
145:     unsigned long int timeout;
146:     struct rlimit rlim[RLIM_NLIMITS];
147:     unsigned long int rss;
148:     __mode_t umask;
149:     unsigned char loopcnt;          /* nested symlinks counter */
150:     struct proc *prev;
151:     struct proc *next;
152:     struct proc *sleep_prev;
153:     struct proc *sleep_next;
154: };
155:
156: extern struct proc *current;
157: extern struct proc *proc_table;
158:
159: int send_sig(struct proc *, __sigset_t);
160: int kill_pid(__pid_t, __sigset_t);
161: int kill_pgrp(__pid_t, __sigset_t);
162: void add_crusage(struct proc *, struct rusage *);
163: void get_rusage(struct proc *, struct rusage *);
164: void add_rusage(struct proc *);
165: struct proc * get_next_zombie(struct proc *);
166: __pid_t remove_zombie(struct proc *);
167: int is_orphaned_pgrp(__pid_t);
168: struct proc * get_proc_free(void);
169: void release_proc(struct proc *);
170: int get_unused_pid(void);
171: struct proc * get_proc_by_pid(__pid_t);
172:
173: int get_new_user_fd(int);
174: void release_user_fd(int);
175:
176: struct proc * kernel_process(int (*fn) (void));
177: void proc_slot_init(struct proc *);
178: void proc_init(void);
179:
180: int elf_load(struct inode *, struct binargs *, struct sigcontext *, char *);
181: int script_load(char *, char *, char *);
182:
183: #endif /* _FIWIX_PROCESS_H */

```

**include/fiwix/ramdisk.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/ramdisk.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_RAMDISK_H
9: #define _FIWIX_RAMDISK_H
10:
11: #include <fiwix/fs.h>
12:
13: #define RAMDISK_MAJOR 1      /* ramdisk device major number */
14: #define RAMDISK_MINORS 1    /* number of minors */
15: #define RAMDISK_SIZE 4096   /* default ramdisk size in KBs */
16: #define RAMDISK_MAXSIZE 131072 /* maximum ramdisk size in KBs */
17:
18: struct ramdisk {
19:     char *addr;          /* ramdisk memory address */
20: };
21:
22: extern struct ramdisk ramdisk_table[RAMDISK_MINORS];
23:
24: int ramdisk_open(struct inode *, struct fd *);
25: int ramdisk_close(struct inode *, struct fd *);
26: int ramdisk_read(__dev_t, __blk_t, char *, int);
27: int ramdisk_write(__dev_t, __blk_t, char *, int);
28: int ramdisk_ioctl(struct inode *, int, unsigned long int);
29: int ramdisk_lseek(struct inode *, __off_t);
30:
31: void ramdisk_init(void);
32:
33: #endif /* _FIWIX_RAMDISK_H */
```

**include/fiwix/reboot.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/reboot.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_REBOOT_H
9: #define _FIWIX_REBOOT_H
10:
11: #define BMAGIC_HARD      0x89ABCDEF
12: #define BMAGIC_SOFT      0
13: #define BMAGIC_REBOOT    0x01234567
14: #define BMAGIC_HALT      0xCDEF0123
15: #define BMAGIC_POWEROFF  0x4321FEDC
16:
17: #define BMAGIC_1         0xFEE1DEAD
18: #define BMAGIC_2         672274793
19:
20: extern char ctrl_alt_del;
21: void reboot(void);
22:
23: #endif /* _FIWIX_REBOOT_H */
```

## include/fiwix/resource.h

Page 1/2

```

1:  /*
2:  * fiwix/include/fiwix/resource.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_RESOURCE_H
9: #define _FIWIX_RESOURCE_H
10:
11: #include <fiwix/time.h>
12:
13: #define RLIMIT_INFINITY 0x7FFFFFFF /* value to indicate "no limit" */
14: #define RLIM_INFINITY RLIMIT_INFINITY /* traditional name */
15:
16: #define RUSAGE_SELF 0 /* the calling process */
17: #define RUSAGE_CHILDREN (-1) /* all of its termin. child processes */
18:
19: #define RLIMIT_CPU 0 /* per-process CPU limit (secs) */
20: #define RLIMIT_FSIZE 1 /* largest file that can be created
21: (bytes) */
22: #define RLIMIT_DATA 2 /* maximum size of data segment
23: (bytes) */
24: #define RLIMIT_STACK 3 /* maximum size of stack segment
25: (bytes) */
26: #define RLIMIT_CORE 4 /* largest core file that can be created
27: (bytes) */
28: #define RLIMIT_RSS 5 /* largest resident set size (bytes) */
29: #define RLIMIT_NPROC 6 /* number of processes */
30: #define RLIMIT_NOFILE 7 /* number of open files */
31: #define RLIMIT_MEMLOCK 8 /* locked-in-memory address space */
32: #define RLIMIT_AS 9 /* address space limit */
33:
34: #define RLIM_NLIMITS 10
35:
36: struct rusage {
37:     struct timeval ru_utime; /* total amount of user time used */
38:     struct timeval ru_stime; /* total amount of system time used */
39:     long ru_maxrss; /* maximum resident set size (KB) */
40:     long ru_ixrss; /* amount of sharing of text segment
41: memory with other processes
42: (KB-secs) */
43:     long ru_idrss; /* amount of data segment memory used
44: (KB-secs) */
45:     long ru_isrss; /* amount of stack memory used
46: (KB-secs) */
47:     long ru_minflt; /* number of soft page faults (i.e.
48: those serviced by reclaiming a page
49: from the list of pages awaiting
50: relocation) */
51:     long ru_majflt; /* number of hard page faults (i.e.
52: those that required I/O) */
53:     long ru_nswap; /* number of times a process was swapped
54: out of physical memory */
55:     long ru_inblock; /* number of input operations via the
56: file system. Note this and
57: 'ru_outblock' do not include
58: operations with the cache */
59:     long ru_oublock; /* number of output operations via the
60: file system */
61:     long ru_msgsnd; /* number of IPC messages sent */
62:     long ru_msrvcv; /* number of IPC messages received */
63:     long ru_nsignals; /* number of signals delivered */
64:     long ru_nvcsw; /* number of voluntary context switches,
65: i.e. because the process gave up the
66: process before it had to (usually to
67: wait for some resource to be

```

**include/fiwix/resource.h**

Page 2/2

```
68:                                     availabe */
69:         long ru_nivcsw;                /* number of involuntary context
70:                                     switches. i.e. a higher priority
71:                                     process became runnable or the
72:                                     current process used up its time
73:                                     slice */
74: };
75:
76: struct rlimit {
77:     int rlim_cur;                       /* the current (soft) limit */
78:     int rlim_max;                       /* the maximum (hard) limit */
79: };
80:
81: #endif /* _FIWIX_RESOURCE_H */
```



## include/fiwix/sched.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/sched.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_SCHED_H
9: #define _FIWIX_SCHED_H
10:
11: #include <fiwix/process.h>
12:
13: #define PRIO_PROCESS    0
14: #define PRIO_PGRP      1
15: #define PRIO_USER      2
16:
17: #define PROC_UNUSED    0
18: #define PROC_RUNNING   1
19: #define PROC_SLEEPING  2
20: #define PROC_ZOMBIE    3
21: #define PROC_STOPPED   4
22: #define PROC_IDLE      5
23:
24: #define PROC_INTERRUPTIBLE  1
25: #define PROC_UNINTERRUPTIBLE 2
26:
27: #define DEF_PRIORITY      (20 * HZ / 100) /* 200ms of time slice */
28:
29: extern int need_resched;
30:
31: #define SI_LOAD_SHIFT    16
32:
33: /*
34:  * This was brought from Linux 2.0.30 (sched.h).
35:  * Copyright Linus Torvalds et al.
36:  */
37: extern unsigned int avenrun[3];          /* Load averages */
38: #define FSHIFT          11              /* nr of bits of precision */
39: #define FIXED_1         (1<<FSHIFT)    /* 1.0 as fixed-point */
40: #define LOAD_FREQ       (5*HZ)         /* 5 sec intervals */
41: #define EXP_1           1884            /* 1/exp(5sec/1min) as fixed-point */
42: #define EXP_5           2014            /* 1/exp(5sec/5min) */
43: #define EXP_15          2037            /* 1/exp(5sec/15min) */
44:
45: #define CALC_LOAD(load,exp,n) \
46:     load *= exp; \
47:     load += n*(FIXED_1-exp); \
48:     load >>= FSHIFT;
49: /* ----- */
50:
51:
52: void do_sched(void);
53: void set_tss(struct proc *);
54: void sched_init(void);
55:
56: #endif /* _FIWIX_SCHED_H */

```

## include/fiwix/segments.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/segments.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_SEGMENTS_H
9: #define _FIWIX_SEGMENTS_H
10:
11: #include <fiwix/types.h>
12:
13: #define NR_GDT_ENTRIES 6 /* entries in GDT descriptor */
14: #define NR_IDT_ENTRIES 256 /* entries in IDT descriptor */
15:
16: /* low flags of Segment Descriptors */
17: #define SD_DATA 0x02 /* DATA Read/Write */
18: #define SD_CODE 0x0A /* CODE Exec/Read */
19:
20: #define SD_32TSSA 0x09 /* 32-bit TSS (Available) */
21: #define SD_32TSSB 0x0B /* 32-bit TSS (Busy) */
22: #define SD_32CALLGATE 0x0C /* 32-bit Call Gate */
23: #define SD_32INTRGATE 0x0E /* 32-bit Interrupt Gate (0D110) */
24: #define SD_32TRAPGATE 0x0F /* 32-bit Trap Gate (0D111) */
25:
26: #define SD_CD 0x10 /* 0 = system / 1 = code/data */
27: #define SD_DPL0 0x00 /* priority level 0 */
28: #define SD_DPL1 0x20 /* priority level 1 (unused) */
29: #define SD_DPL2 0x40 /* priority level 2 (unused) */
30: #define SD_DPL3 0x60 /* priority level 3 (user) */
31: #define SD_PRESENT 0x80 /* segment present or valid */
32:
33: /* high flags Segment Descriptors */
34: #define SD_OPSIZE32 0x04 /* 32-bit code and data segments */
35: #define SD_PAGE4KB 0x08 /* page granularity (4KB) */
36:
37: /* low flags of the TSS Descriptors */
38: #define SD_TSSPRESENT 0x89 /* TSS present and not busy flag */
39:
40: #define USR_PL 3 /* User Privilege Level */
41:
42: /* EFLAGS */
43: #define EF_IOPL 12 /* IOPL bit */
44:
45: struct desc_r {
46:     __u16 limit;
47:     __u32 base_addr;
48: } __attribute__((packed));
49:
50: struct seg_desc {
51:     unsigned sd_lolimit : 16; /* segment limit 0-15 bits */
52:     unsigned sd_lobase : 24; /* base address 0-23 bits */
53:     unsigned sd_loflags : 8; /* flags (P, DPL, S and TYPE) */
54:     unsigned sd_hilimit : 4; /* segment limit 16-19 bits */
55:     unsigned sd_hiflags : 4; /* flags (G, DB, 0 and AVL) */
56:     unsigned sd_hibase : 8; /* base address 24-31 bits */
57: } __attribute__((packed));
58:
59: struct gate_desc {
60:     unsigned gd_looffset : 16; /* offset 0-15 bits */
61:     unsigned gd_selector : 16; /* segment selector */
62:     unsigned gd_flags : 16; /* flags (P, DPL, TYPE, 0 and NULL) */
63:     unsigned gd_hioffset : 16; /* offset 16-31 bits */
64: } __attribute__((packed));
65:
66: void gdt_init(void);
67: void idt_init(void);

```

**include/fiwix/segments.h**

Page 2/2

```
68:
69: #endif /* _FIWIX_SEGMENTS_H */
```

## include/fiwix/serial.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/serial.h
3:  *
4:  * Copyright 2020, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_SERIAL_H
9: #define _FIWIX_SERIAL_H
10:
11: // #include <fiwix/fs.h>
12:
13: #define SERIAL4_IRQ    4        /* IRQ for serial ports 1 and 3 */
14: #define SERIAL3_IRQ    3        /* IRQ for serial ports 2 and 4 */
15:
16: #define NR_SERIAL      4        /* maximum number of serial ttys */
17: #define SERIAL_MAJOR   4        /* major number for /dev/ttyS[n] */
18: #define SERIAL_MINORS  NR_SERIAL
19: #define SERIAL_MSF     6        /* serial minor shift factor */
20:
21: /* UART registers */
22: #define UART_TD        0        /* W: Transmitter Holding Buffer */
23: #define UART_RD        0        /* R: Receiver Buffer */
24: #define UART_DLL       0        /* RW: Divisor Latch Low Byte */
25: #define UART_DLH       1        /* RW: Divisor Latch High Byte */
26: #define UART_IER       1        /* RW: Interrupt Enable Register */
27: #define UART_IIR       2        /* R: Interrupt Identification Register */
28: #define UART_FCR       2        /* W: FIFO Control Register */
29: #define UART_LCR       3        /* RW: Line Control Register */
30: #define UART_MCR       4        /* RW: Modem Control Register */
31: #define UART_LSR       5        /* R: Line Status Register */
32: #define UART_MSR       6        /* R: Modem Status Register */
33: #define UART_SR       7        /* RW: Scratch Register */
34:
35: /* Interrupt Enable Register */
36: #define UART_IER_RDAI  0x1      /* enable Received Data Available Interrupt */
37: #define UART_IER_THREI 0x2      /* enable Transmitter Holding Register Empty Interrupt */
38: #define UART_IER_RLSI  0x4      /* enable Receiver Line Status Interrupt */
39: #define UART_IER_MSI   0x8      /* enable Modem Status Interrupt */
40:
41: /* Interrupt Identification Register */
42: #define UART_IIR_NOINT 0x01     /* no interrupts pending */
43: #define UART_IIR_MKINT 0x06     /* mask all interrupt flags */
44: #define UART_IIR_MSI   0x00     /* Modem Status Interrupt */
45: #define UART_IIR_THREI 0x02     /* Transmitter Holding Register Empty Interrupt */
46: #define UART_IIR_RDAI  0x04     /* Received Data Available Interrupt */
47: #define UART_IIR_RLSI  0x06     /* Receiver Line Status Interrupt */
48: #define UART_IIR_FIFO0 0xC0     /* FIFO TimeOut interrupt */
49: #define UART_IIR_FIFO64 0x20    /* 64 byte FIFO enabled (16750 only) */
50: #define UART_IIR_FIFO  0x40     /* FIFO is enabled (still needs bit #7 on) */
51: #define UART_IIR_FIFOK 0x80     /* FIFO is enabled, but unusable */
52:
53: /* FIFO Control Register */
54: #define UART_FCR_FIFO  0x07     /* enable FIFO (clear receive and transmit) */
55: #define UART_FCR_CRCVR 0x02     /* clear receiver */
56: #define UART_FCR_CXMTR 0x04     /* clear transmitter */
57: #define UART_FCR_DMA   0x08     /* DMA mode select */
58: #define UART_FCR_FIFO64 0x20    /* enable 64 byte FIFO (16750 only) */
59: #define UART_FCR_FIFO14 0xC0    /* set to 14 bytes 'trigger level' FIFO */
60:
61: /* Line Control Register */
62: #define UART_LCR_WL5   0x00     /* word length 5 bits */
63: #define UART_LCR_WL6   0x01     /* word length 6 bits */
64: #define UART_LCR_WL7   0x02     /* word length 7 bits */
65: #define UART_LCR_WL8   0x03     /* word length 8 bits */

```

## include/fiwix/serial.h

Page 2/2

```

66: #define UART_LCR_2STB 0x04 /* 2 stop bits */
67: #define UART_LCR_1STB 0x00 /* 1 stop bit */
68: #define UART_LCR_NP 0x00 /* no parity */
69: #define UART_LCR_OP 0x08 /* odd parity */
70: #define UART_LCR_EP 0x18 /* even parity */
71: #define UART_LCR_SBRK 0x40 /* Set Break enable */
72: #define UART_LCR_DLAB 0x80 /* Divisor Latch Access Bit */
73:
74: /* Modem Control Register */
75: #define UART_MCR_DTR 0x1 /* Data Terminal Ready */
76: #define UART_MCR_RTS 0x2 /* Request To Send */
77: #define UART_MCR_OUT2 0x8 /* Auxiliary Output 2 */
78:
79: /* Line Status Register */
80: #define UART_LSR_RDA 0x01 /* Received Data Available */
81: #define UART_LSR_OE 0x02 /* Overrun Error */
82: #define UART_LSR_PE 0x04 /* Parity Error */
83: #define UART_LSR_FE 0x08 /* Framing Error */
84: #define UART_LSR_BI 0x10 /* Break Interrupt */
85: #define UART_LSR_THRE 0x20 /* Transmitter Holding Register Empty */
86: #define UART_LSR_EDHR 0x40 /* Empty Data Holding Registers TD and SH */
87: #define UART_LSR_EFIFO 0x80 /* Error in Received FIFO */
88:
89:
90: #define UART_FIFO_SIZE 16 /* 16 bytes */
91: #define UART_HAS_FIFO 0x02 /* has FIFO working */
92: #define UART_IS_8250 0x04 /* is a 8250 chip */
93: #define UART_IS_16450 0x08 /* is a 16450 chip */
94: #define UART_IS_16550 0x10 /* is a 16550 chip */
95: #define UART_IS_16550A 0x20 /* is a 16550A chip */
96:
97: struct serial {
98:     short int addr; /* port I/O address */
99:     char irq;
100:    int baud;
101:    char *name;
102:    short int lctrl; /* line control flags (8N1, 7E2, ...) */
103:    int flags;
104:    struct tty *tty;
105:    struct serial *next;
106: };
107:
108: int serial_open(struct tty *);
109: int serial_close(struct tty *);
110: int serial_ioctl(struct tty *, int, unsigned long int);
111: void serial_write(struct tty *);
112: void irq_serial(int, struct sigcontext *);
113: void irq_serial_bh(void);
114: void serial_init(void);
115:
116: #endif /* _FIWIX_SERIAL_H */

```

**include/fiwix/sigcontext.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/sigcontext.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_SIGCONTEXT_H
9: #define _FIWIX_SIGCONTEXT_H
10:
11: struct sigcontext {
12:     unsigned int gs;
13:     unsigned int fs;
14:     unsigned int es;
15:     unsigned int ds;
16:     unsigned int edi;
17:     unsigned int esi;
18:     unsigned int ebp;
19:     unsigned int esp;
20:     int ebx;
21:     int edx;
22:     int ecx;
23:     int eax;
24:     int err;
25:     unsigned int eip;
26:     unsigned int cs;
27:     unsigned int eflags;
28:     unsigned int oldesp;
29:     unsigned int oldss;
30: };
31:
32: #endif /* _FIWIX_SIGCONTEXT_H */
```

## include/fiwix/signal.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/signal.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_SIGNAL_H
9: #define _FIWIX_SIGNAL_H
10:
11: #define NSIG          32
12:
13: #define SIGHUP        1      /* Hangup or Reset */
14: #define SIGINT        2      /* Interrupt */
15: #define SIGQUIT       3      /* Quit */
16: #define SIGILL        4      /* Illegal Instruction */
17: #define SIGTRAP       5      /* Trace Trap */
18: #define SIGABRT       6      /* Abort Instruction */
19: #define SIGIOT        SIGABRT /* I/O Trap Instruction */
20: #define SIGBUS        7      /* Bus Error */
21: #define SIGFPE        8      /* Floating Point Exception */
22: #define SIGKILL       9      /* Kill */
23: #define SIGUSR1       10     /* User Defined #1 */
24: #define SIGSEGV       11     /* Segmentation Violation */
25: #define SIGUSR2       12     /* User Defined #2 */
26: #define SIGPIPE       13     /* Broken Pipe */
27: #define SIGALRM       14     /* Alarm Clock */
28: #define SIGTERM       15     /* Software Termination */
29: #define SIGSTKFLT     16     /* Stack Fault */
30: #define SIGCHLD       17     /* Child Termination */
31: #define SIGCONT       18     /* Continue */
32: #define SIGSTOP       19     /* Stop */
33: #define SIGTSTP       20     /* Terminal Stop */
34: #define SIGTTIN       21     /* Background Read */
35: #define SIGTTOU       22     /* Background Write */
36: #define SIGURG        23     /* Urgent Data */
37: #define SIGXCPU       24     /* CPU eXceeded */
38: #define SIGXFSZ       25     /* File Size eXceeded */
39: #define SIGVTALRM     26     /* Virtual Time Alarm */
40: #define SIGPROF       27     /* Profile Alarm */
41: #define SIGWINCH      28     /* Window Change */
42: #define SIGIO         29     /* I/O Asynchronous */
43: #define SIGPOLL       SIGIO
44: #define SIGPWR        30     /* Power Fault */
45: #define SIGUNUSED     31
46:
47: typedef unsigned long int __sigset_t;
48: typedef void (*__sighandler_t)(int);
49:
50: struct sigaction {
51:     __sighandler_t sa_handler;
52:     __sigset_t sa_mask;
53:     int sa_flags;
54:     void (*sa_restorer)(void);
55: };
56:
57: #define SIG_DFL        ((__sighandler_t) 0)
58: #define SIG_IGN        ((__sighandler_t) 1)
59: #define SIG_ERR        ((__sighandler_t) -1)
60:
61: /* bits in sa_flags */
62: #define SA_NOCLDSTOP   0x00000001 /* don't send SIGCHLD when children stop
*/
63: #define SA_NOCLDWAIT   0x00000002 /* don't create zombie on child death */
64: #define SA_ONSTACK     0x08000000 /* invoke handler on alternate stack */
65: #define SA_RESTART     0x10000000 /* automatically restart system call */
66: #define SA_INTERRUPT   0x20000000 /* unused */

```

**include/fiwix/signal.h**

Page 2/2

```
67:
68: /* don't automatically block signal when the handler is executing */
69: #define SA_NODEFER      0x40000000
70: #define SA_NOMASK      SA_NODEFER
71:
72: /* reset signal disposition to SIG_DFL before invoking handler */
73: #define SA_RESETHAND    0x80000000
74: #define SA_ONESHOT      SA_RESETHAND
75:
76: /* bits in the third argument to 'waitpid/wait4' */
77: #define WNOHANG         1      /* don't block waiting */
78: #define WUNTRACED      2      /* report status of stopped children */
79:
80: #define SIG_BLOCK       0      /* for blocking signals */
81: #define SIG_UNBLOCK    1      /* for unblocking signals */
82: #define SIG_SETMASK    2      /* for setting the signal mask */
83:
84: /* SIGKILL and SIGSTOP can't ever be set as blockable signals */
85: #define SIG_BLOCKABLE  (~(1 << (SIGKILL - 1)) | (1 << (SIGSTOP - 1)))
86:
87: #define SIG_MASK(sig)  (~(1 << ((sig) - 1)))
88:
89: int issig(void);
90: void psig(unsigned int);
91:
92: #endif /* _FIWIX_SIGNAL_H */
```



**include/fiwix/sleep.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/sleep.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_SLEEP_H
9: #define _FIWIX_SLEEP_H
10:
11: #include <fiwix/process.h>
12:
13: #define AREA_BH 0x00000001
14: #define AREA_CALLOUT 0x00000002
15: #define AREA_TTY_READ 0x00000004
16: #define AREA_SERIAL_READ 0x00000008
17:
18: struct sleep {
19:     unsigned short int next;
20:     void *address;
21:     struct proc *proc;
22: };
23:
24: struct resource {
25:     char locked;
26:     char wanted;
27: };
28:
29: int sleep(void *, int);
30: void wakeup(void *);
31: void wakeup_proc(struct proc *);
32:
33: void lock_resource(struct resource *);
34: void unlock_resource(struct resource *);
35: int lock_area(unsigned int);
36: int unlock_area(unsigned int);
37:
38: void sleep_init(void);
39:
40: #endif /* _FIWIX_SLEEP_H */
```

## include/fiwix/statbuf.h

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/statbuf.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_STATBUF_H
9: #define _FIWIX_STATBUF_H
10:
11: struct old_stat {
12:     __dev_t st_dev;
13:     unsigned short int st_ino;
14:     __mode_t st_mode;
15:     __nlink_t st_nlink;
16:     __uid_t st_uid;
17:     __gid_t st_gid;
18:     __dev_t st_rdev;
19:     unsigned int st_size;
20:     __time_t st_atime;
21:     __time_t st_mtime;
22:     __time_t st_ctime;
23: };
24:
25: struct new_stat {
26:     __dev_t st_dev;
27:     unsigned short int __pad1;
28:     __ino_t st_ino;
29:     __mode_t st_mode;
30:     __nlink_t st_nlink;
31:     __uid_t st_uid;
32:     __gid_t st_gid;
33:     __dev_t st_rdev;
34:     unsigned short int __pad2;
35:     __off_t st_size;
36:     __blk_t st_blksize;
37:     __blk_t st_blocks;
38:     __time_t st_atime;
39:     unsigned int __unused1;
40:     __time_t st_mtime;
41:     unsigned int __unused2;
42:     __time_t st_ctime;
43:     unsigned int __unused3;
44:     unsigned int __unused4;
45:     unsigned int __unused5;
46: };
47:
48: #endif /* _FIWIX_STATBUF_H */
```

**include/fiwix/statfs.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/statfs.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_STATFS_H
9: #define _FIWIX_STATFS_H
10:
11: typedef struct {
12:     long int val[2];
13: } fsid_t;
14:
15: struct statfs {
16:     long int f_type;
17:     long int f_bsize;
18:     long int f_blocks;
19:     long int f_bfree;
20:     long int f_bavail;
21:     long int f_files;
22:     long int f_ffree;
23:     fsid_t f_fsid;
24:     long int f_namelen;
25:     long int f_spare[6];
26: };
27:
28: #endif /* _FIWIX_STATFS_H */
```

## include/fiwix/stat.h

Page 1/1

```

1: #ifndef _FIWIX_STAT_H
2: #define _FIWIX_STAT_H
3:
4: #include <fiwix/statbuf.h>
5:
6: /* Encoding of the file mode.  These are the standard Unix values,
7:    but POSIX.1 does not specify what values should be used.  */
8:
9: #define S_IFMT          0170000          /* Type of file mask */
10:
11: /* File types.  */
12: #define S_IFIFO         0010000          /* Named pipe (fifo) */
13: #define S_IFCHR         0020000          /* Character special */
14: #define S_IFDIR         0040000          /* Directory */
15: #define S_IFBLK         0060000          /* Block special */
16: #define S_IFREG         0100000          /* Regular */
17: #define S_IFLNK         0120000          /* Symbolic link */
18: #define S_IFSOCK        0140000          /* Socket */
19:
20: /* Protection bits.  */
21: #define S_IXUSR         00100           /* USER  --x----- */
22: #define S_IWUSR         00200           /* USER  -w----- */
23: #define S_IRUSR         00400           /* USER  r----- */
24: #define S_IRWXU         00700           /* USER  rwx----- */
25:
26: #define S_IXGRP         00010           /* GROUP  -----x--- */
27: #define S_IWGRP         00020           /* GROUP  -----w--- */
28: #define S_IRGRP         00040           /* GROUP  ----r----- */
29: #define S_IRWXG         00070           /* GROUP  ---rwx--- */
30:
31: #define S_IXOTH         00001           /* OTHERS  -----x */
32: #define S_IWOTH         00002           /* OTHERS  -----w- */
33: #define S_IROTH         00004           /* OTHERS  -----r-- */
34: #define S_IRWXO         00007           /* OTHERS  -----rwx */
35:
36: #define S_ISUID         0004000          /* set user id on execution */
37: #define S_ISGID         0002000          /* set group id on execution */
38: #define S_ISVTX         0001000          /* sticky bit */
39:
40: #define S_IREAD         S_IRUSR          /* Read by owner.  */
41: #define S_IWRITE        S_IWUSR          /* Write by owner.  */
42: #define S_IEXEC         S_IXUSR          /* Execute by owner.  */
43:
44: #define S_ISFIFO(m)     ((m) & S_IFMT) == S_IFIFO)
45: #define S_ISCHR(m)      ((m) & S_IFMT) == S_IFCHR)
46: #define S_ISDIR(m)      ((m) & S_IFMT) == S_IFDIR)
47: #define S_ISBLK(m)      ((m) & S_IFMT) == S_IFBLK)
48: #define S_ISREG(m)      ((m) & S_IFMT) == S_IFREG)
49: #define S_ISLNK(m)      ((m) & S_IFMT) == S_IFLNK)
50: #define S_ISSOCK(m)     ((m) & S_IFMT) == S_IFSOCK)
51:
52: #define TO_READ         4                /* test for read permission */
53: #define TO_WRITE        2                /* test for write permission */
54: #define TO_EXEC         1                /* test for execute permission */
55:
56: #endif /* _FIWIX_STAT_H */

```

**include/fiwix/stdarg.h**

Page 1/1

```
1: /*
2: Copyright (C) 1988 Free Software Foundation
3:
4: This file is part of GNU CC.
5:
6: GNU CC is distributed in the hope that it will be useful,
7: but WITHOUT ANY WARRANTY. No author or distributor
8: accepts responsibility to anyone for the consequences of using it
9: or for whether it serves any particular purpose or works at all,
10: unless he says so in writing. Refer to the GNU CC General Public
11: License for full details.
12:
13: Everyone is granted permission to copy, modify and redistribute
14: GNU CC, but only under the conditions described in the
15: GNU CC General Public License. A copy of this license is
16: supposed to have been given to you along with GNU CC so you
17: can know your rights and responsibilities. It should be in a
18: file named COPYING. Among other things, the copyright notice
19: and this notice must be preserved on all copies.
20: */
21:
22: #ifndef __stdarg_h
23: #define __stdarg_h
24:
25: typedef char *va_list;
26:
27: /* Amount of space required in an argument list for an arg of type TYPE.
28:    TYPE may alternatively be an expression whose type is used. */
29:
30: #define __va_rounded_size(TYPE) \
31:    (((sizeof (TYPE) + sizeof (int) - 1) / sizeof (int)) * sizeof (int))
32:
33: #define va_start(AP, LASTARG) \
34:    (AP = ((char *) &(LASTARG) + __va_rounded_size(LASTARG)))
35:
36: extern void va_end (va_list);
37: #define va_end(AP) /* Nothing */
38:
39: #define va_arg(AP, TYPE) (AP += __va_rounded_size (TYPE), \
40:    *((TYPE *) (AP - __va_rounded_size (TYPE))))
41:
42: #endif /* __stdarg_h */
```

**include/fiwix/stdio.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/stdio.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _INCLUDE_STDIO_H
9: #define _INCLUDE_STDIO_H
10:
11: void register_console(void (*fn)(char *, unsigned int));
12: void printk(const char *, ...);
13: int sprintk(char *, const char *, ...);
14:
15: #endif /* _INCLUDE_STDIO_H */
```

## include/fiwix/string.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/string.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _INCLUDE_STRING_H
9: #define _INCLUDE_STRING_H
10:
11: #include <fiwix/types.h>
12:
13: #ifndef NULL
14: #define NULL    '\0'    /* ((void *)0) */
15: #endif
16:
17: #define MIN(a,b)    ((a) < (b) ? (a) : (b))
18: #define MAX(a,b)    ((a) > (b) ? (a) : (b))
19:
20: #define IS_NUMERIC(c)    ((c) >= '0' && (c) <= '9')
21: #define IS_SPACE(c)    ((c) == ' ')
22:
23: void swap_asc_word(char *, int);
24: int strcmp(const char *, const char *);
25: int strncmp(const char *, const char *, __ssize_t);
26: char * strcpy(char *, const char *);
27: void strncpy(char *, const char *, int);
28: char * strcat(char *, const char *);
29: char * strncat(char *, const char *, __ssize_t);
30: int strlen(const char *);
31: char * get_basename(const char *);
32: char * remove_trailing_slash(char *);
33: int is_dir(const char *);
34: int atoi(const char *);
35: void memcpy_b(void *, const void *, unsigned int);
36: void memcpy_w(void *, const void *, unsigned int);
37: void memcpy_l(void *, const void *, unsigned int);
38: void memset_b(void *, unsigned char, unsigned int);
39: void memset_w(void *, unsigned short int, unsigned int);
40: void memset_l(void *, unsigned int, unsigned int);
41:
42: #endif /* _INCLUDE_STRING_H */

```

## include/fiwix/syscalls.h

Page 1/3

```

1:  /*
2:  * fiwix/include/fiwix/syscalls.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #ifndef _FIWIX_SYSCALLS_H
9:  #define _FIWIX_SYSCALLS_H
10:
11: #include <fiwix/types.h>
12: #include <fiwix/system.h>
13: #include <fiwix/time.h>
14: #include <fiwix/times.h>
15: #include <fiwix/timeb.h>
16: #include <fiwix/utime.h>
17: #include <fiwix/statbuf.h>
18: #include <fiwix/ustat.h>
19: #include <fiwix/signal.h>
20: #include <fiwix/utsname.h>
21: #include <fiwix/resource.h>
22: #include <fiwix/dirent.h>
23: #include <fiwix/statfs.h>
24: #include <fiwix/sigcontext.h>
25: #include <fiwix/mman.h>
26:
27: #define NR_SYSCALLS      (sizeof(syscall_table) / sizeof(unsigned int))
28:
29: int do_syscall(unsigned int, int, int, int, int, int, int, struct sigcontext);
30:
31: int sys_exit(int);
32: void do_exit(int);
33: int sys_fork(int, int, int, int, int, struct sigcontext *);
34: int sys_read(unsigned int, char *, int);
35: int sys_write(unsigned int, const char *, int);
36: int sys_open(const char *, int, __mode_t);
37: int sys_close(unsigned int);
38: int sys_waitpid(__pid_t, int *, int);
39: int sys_creat(const char *, __mode_t);
40: int sys_link(const char *, const char *);
41: int sys_unlink(const char *);
42: int sys_execve(const char *, char **, char **, int, int, struct sigcontext *);
43: int sys_chdir(const char *);
44: int sys_time(__time_t *);
45: int sys_mknod(const char *, __mode_t, __dev_t);
46: int sys_chmod(const char *, __mode_t);
47: int sys_chown(const char *, __uid_t, __gid_t);
48: int sys_stat(const char *, struct old_stat *);
49: int sys_lseek(unsigned int, __off_t, unsigned int);
50: int sys_getpid(void);
51: int sys_mount(const char *, const char *, const char *, unsigned long int, const
void *);
52: int sys_umount(const char *);
53: int sys_setuid(__uid_t);
54: int sys_getuid(void);
55: int sys_stime(__time_t *);
56: int sys_alarm(unsigned int);
57: int sys_fstat(unsigned int, struct old_stat *);
58: int sys_pause(void);
59: int sys_utime(const char *, struct utimbuf *);
60: int sys_access(const char *, __mode_t);
61: int sys_ftime(struct timeb *);
62: void sys_sync(void);
63: int sys_kill(__pid_t, __sigset_t);
64: int sys_rename(const char *, const char *);
65: int sys_mkdir(const char *, __mode_t);
66: int sys_rmdir(const char *);

```



## include/fiwix/syscalls.h

Page 2/3

```

67: int sys_dup(unsigned int);
68: int sys_pipe(int *);
69: int sys_times(struct tms *);
70: int sys_brk(unsigned int);
71: int sys_setgid(__gid_t);
72: int sys_getgid(void);
73: unsigned int sys_signal(__sigset_t, void(*sighandler)(int));
74: int sys_geteuid(void);
75: int sys_getegid(void);
76: int sys_umount2(const char *, int);
77: int sys_ioctl(unsigned int, int, unsigned long int);
78: int sys_fcntl(int, int, unsigned long int);
79: int sys_setpgid(__pid_t, __pid_t);
80: int sys_olduname(struct oldold_utsname *);
81: int sys_umask(__mode_t);
82: int sys_chroot(const char *);
83: int sys_ustat(__dev_t, struct ustat *);
84: int sys_dup2(int, int);
85: int sys_getppid(void);
86: int sys_getpgrp(void);
87: int sys_setsid(void);
88: int sys_sigaction(__sigset_t, const struct sigaction *, struct sigaction *);
89: int sys_sgetmask(void);
90: int sys_ssetmask(int);
91: int sys_setreuid(__uid_t, __uid_t);
92: int sys_setregid(__gid_t, __gid_t);
93: int sys_sigsuspend(__sigset_t *);
94: int sys_sigpending(__sigset_t *);
95: int sys_sethostname(const char *, int);
96: int sys_setrlimit(int, const struct rlimit *);
97: int sys_getrlimit(int, struct rlimit *);
98: int sys_getrusage(int, struct rusage *);
99: int sys_gettimeofday(struct timeval *, struct timezone *);
100: int sys_settimeofday(const struct timeval *, const struct timezone *);
101: int sys_getgroups(__ssize_t, __gid_t *);
102: int sys_setgroups(__ssize_t, const __gid_t *);
103: int old_select(unsigned long int *);
104: int sys_symlink(const char *, const char *);
105: int sys_lstat(const char *, struct old_stat *);
106: int sys_readlink(const char *, char *, __size_t);
107: int sys_reboot(int, int, int);
108: int old_mmap(struct mmap *);
109: int sys_munmap(unsigned int, __size_t);
110: int sys_truncate(const char *, __off_t);
111: int sys_ftruncate(int, __off_t);
112: int sys_fchmod(int, __mode_t);
113: int sys_fchown(int, __uid_t, __gid_t);
114: int sys_statfs(const char *, struct statfs *);
115: int sys_fstatfs(unsigned int, struct statfs *);
116: int sys_ioperm(unsigned long int, unsigned long int, int);
117: int sys_socketcall(int, unsigned long int *);
118: int sys_setitimer(int, const struct itimerval *, struct itimerval *);
119: int sys_getitimer(int, struct itimerval *);
120: int sys_newstat(const char *, struct new_stat *);
121: int sys_newlstat(const char *, struct new_stat *);
122: int sys_newfstat(unsigned int, struct new_stat *);
123: int sys_uname(struct old_utsname *);
124: int sys_iopl(int, int, int, int, int, struct sigcontext *);
125: int sys_wait4(__pid_t, int *, int, struct rusage *);
126: int sys_sysinfo(struct sysinfo *);
127: int sys_fsync(int);
128: int sys_sigreturn(unsigned int, int, int, int, int, int, struct sigcontext *);
129: int sys_setdomainname(const char *, int);
130: int sys_newuname(struct new_utsname *);
131: int sys_mprotect(unsigned int, __size_t, int);
132: int sys_sigprocmask(int, const __sigset_t *, __sigset_t *);
133: int sys_getpgid(__pid_t);

```

**include/fiwix/syscalls.h**

Page 3/3

```
134: int sys_fchdir(unsigned int);
135: int sys_personality(unsigned long int);
136: int sys_setfsuid(__uid_t);
137: int sys_setfsgid(__gid_t);
138: int sys_llseek(unsigned int, unsigned long int, unsigned long int, __loff_t *, u
nsigned int);
139: int sys_getdents(unsigned int, struct dirent *, unsigned int);
140: int sys_select(int, fd_set *, fd_set *, fd_set *, struct timeval *);
141: int sys_flock(int, int);
142: int sys_getsid(__pid_t);
143: int sys_fdatasync(int);
144: int sys_nanosleep(const struct timespec *, struct timespec *);
145: int sys_getcwd(char *, __size_t);
146:
147: #endif /* _FIWIX_SYSCALLS_H */
```

**include/fiwix/system.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/system.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_SYSTEM_H
9: #define _FIWIX_SYSTEM_H
10:
11: #define UTS_SYSNAME      "Fiwix"
12: #define UTS_NODENAME    "(none)"
13: #define UTS_RELEASE     "1.2.0"
14: #define UTS_DOMAINNAME "(none)"
15:
16: struct sysinfo {
17:     long int uptime;           /* seconds since boot */
18:     unsigned long int loads[3]; /* load average (1, 5 and 15 minutes) */
19:     unsigned long int totalram; /* total usable main memory size */
20:     unsigned long int freeram;  /* available memory size */
21:     unsigned long int sharedram; /* amount of shared memory */
22:     unsigned long int bufferram; /* amount of memory used by buffers */
23:     unsigned long int totalswap; /* total swap space size */
24:     unsigned long int freeswap; /* available swap space */
25:     unsigned short int procs;   /* number of current processes */
26:     char _f[22];               /* pads structure to 64 bytes */
27: };
28:
29: #endif /* _FIWIX_SYSTEM_H */
```

## include/fiwix/termbits.h

Page 1/3

```

1: /*
2:  * fiwix/include/fiwix/termbits.h
3:  *
4:  */
5:
6: #ifndef _FIWIX_TERMBITS_H
7: #define _FIWIX_TERMBITS_H
8:
9: /* These definitions match those used by the 4.4 BSD kernel.
10:    If the operating system has termios system calls or ioctls that
11:    correctly implement the POSIX.1 behavior, there should be a
12:    system-dependent version of this file that defines 'struct termios',
13:    'tcflag_t', 'cc_t', 'speed_t' and the 'TC*' constants appropriately. */
14:
15: /* Type of terminal control flag masks. */
16: typedef unsigned long int tcflag_t;
17:
18: /* Type of control characters. */
19: typedef unsigned char cc_t;
20:
21: /* Type of baud rate specifiers. */
22: typedef long int speed_t;
23:
24: /* c_iflag bits */
25: #define IGNBRK 0000001 /* Ignore break condition */
26: #define BRKINT 0000002 /* Signal interrupt on break */
27: #define IGNPAR 0000004 /* Ignore characters with parity errors */
28: #define PARMRK 0000010 /* Mark parity and framing errors */
29: #define INPCK 0000020 /* Enable input parity check */
30: #define ISTRIP 0000040 /* Strip 8th bit off characters */
31: #define INLCR 0000100 /* Map NL to CR on input */
32: #define IGNCR 0000200 /* Ignore CR */
33: #define ICRNL 0000400 /* Map CR to NL on input */
34: #define IUCLC 0001000 /* Convert to lowercase */
35: #define IXON 0002000 /* Enable start/stop output control */
36: #define IXANY 0004000 /* Any character will restart after stop */
37: #define IXOFF 0010000 /* Enable start/stop input control */
38: #define IMAXBEL 0020000 /* Ring bell when input queue is full */
39:
40: /* c_oflag bits */
41: #define OPOST 0000001 /* Perform output processing */
42: #define OLCUC 0000002
43: #define ONLCR 0000004 /* Map NL to CR-NL on output */
44: #define OCRNL 0000010
45: #define ONOCR 0000020
46: #define ONLRET 0000040
47: #define OFILL 0000100
48: #define OFDEL 0000200
49: #define NLDLY 0000400
50: #define NL0 0000000
51: #define NL1 0000400
52: #define CRDLY 0003000
53: #define CR0 0000000
54: #define CR1 0001000
55: #define CR2 0002000
56: #define CR3 0003000
57: #define TABDLY 0014000
58: #define TAB0 0000000
59: #define TAB1 0004000
60: #define TAB2 0010000
61: #define TAB3 0014000
62: #define XTABS 0014000
63: #define BSDLY 0020000
64: #define BS0 0000000
65: #define BS1 0020000
66: #define VTDLY 0040000
67: #define VT0 0000000

```

## include/fiwix/termbits.h

Page 2/3

```

68: #define VT1 0040000
69: #define FFDLY 0100000
70: #define FF0 0000000
71: #define FF1 0100000
72:
73: /* c_cflag bit meaning */
74: #define CBAUD 0010017
75: #define B0 0000000 /* hang up */
76: #define B50 0000001 /* 50 baud */
77: #define B75 0000002 /* 75 baud */
78: #define B110 0000003 /* 110 baud */
79: #define B134 0000004 /* 134 baud */
80: #define B150 0000005 /* 150 baud */
81: #define B200 0000006 /* 200 baud */
82: #define B300 0000007 /* 300 baud */
83: #define B600 0000010 /* 600 baud */
84: #define B1200 0000011 /* 1200 baud */
85: #define B1800 0000012 /* 1800 baud */
86: #define B2400 0000013 /* 2400 baud */
87: #define B4800 0000014 /* 4800 baud */
88: #define B9600 0000015 /* 9600 baud */
89: #define B19200 0000016 /* 19200 baud */
90: #define B38400 0000017 /* 38400 baud */
91: #define EXTA B19200
92: #define EXTB B38400
93: #define CSIZE 0000060 /* Number of bits per byte (mask) */
94: #define CS5 0000000 /* 5 bits per byte */
95: #define CS6 0000020 /* 6 bits per byte */
96: #define CS7 0000040 /* 7 bits per byte */
97: #define CS8 0000060 /* 8 bits per byte */
98: #define CSTOPB 0000100 /* Two stop bits instead of one */
99: #define CREAD 0000200 /* Enable receiver */
100: #define PARENB 0000400 /* Parity enable */
101: #define PARODD 0001000 /* Odd parity instead of even */
102: #define HUPCL 0002000 /* Hang up on last close */
103: #define CLOCAL 0004000 /* Ignore modem status lines */
104: #define CBAUDEX 0010000
105: #define B57600 0010001
106: #define B115200 0010002
107: #define B230400 0010003
108: #define B460800 0010004
109: #define B500000 0010005
110: #define B576000 0010006
111: #define B921600 0010007
112: #define B1000000 0010010
113: #define B1152000 0010011
114: #define B1500000 0010012
115: #define B2000000 0010013
116: #define B2500000 0010014
117: #define B3000000 0010015
118: #define B3500000 0010016
119: #define B4000000 0010017
120: #define CIBAUD 002003600000 /* input baud rate (not used) */
121: #define CMSPAR 010000000000 /* mark or space (stick) parity */
122: #define CRTSCTS 020000000000 /* flow control */
123:
124: /* c_lflag bits */
125: #define ISIG 0000001 /* Enable signals */
126: #define ICANON 0000002 /* Do erase and kill processing */
127: #define XCASE 0000004
128: #define ECHO 0000010 /* Enable echo */
129: #define ECHOE 0000020 /* Visual erase for ERASE */
130: #define ECHOK 0000040 /* Echo NL after KILL */
131: #define ECHONL 0000100 /* Echo NL even if echo is OFF */
132: #define NOFLSH 0000200 /* Disable flush after interrupt */
133: #define TOSTOP 0000400 /* Send SIGTTOU for background output */
134: #define ECHOCTL 0001000 /* Echo control characters as ^X */

```

## include/fiwix/termbits.h

Page 3/3

```

135: #define ECHOPRT 0002000 /* Hardcopy visual erase */
136: #define ECHOKE 0004000 /* Visual erase for KILL */
137: #define FLUSHO 0010000 /* Output being flushed (state) */
138: #define PENDIN 0040000 /* Retype pending input (state) */
139: #define IEXTEN 0100000 /* Enable DISCARD and LNEXT */
140:
141: /* c_cc characters */
142: #define VINTR 0 /* Interrupt character [ISIG] */
143: #define VQUIT 1 /* Quit character [ISIG] */
144: #define VERASE 2 /* Erase character [ICANON] */
145: #define VKILL 3 /* Kill-line character [ICANON] */
146: #define VEOF 4 /* End-of-file character [ICANON] */
147: #define VTIME 5 /* Time-out value (1/10 secs) [!ICANON] */
148: #define VMIN 6 /* Minimum # of bytes read at once [!ICANON] */
149: #define VSWTC 7
150: #define VSTART 8 /* Start (X-ON) character [IXON, IXOFF] */
151: #define VSTOP 9 /* Stop (X-OFF) character [IXON, IXOFF] */
152: #define VSUSP 10 /* Suspend character [ISIG] */
153: #define VEOL 11 /* End-of-line character [ICANON] */
154: #define VREPRINT 12 /* Reprint-line character [ICANON] */
155: #define VDISCARD 13 /* Discard character [IEXTEN] */
156: #define VWERASE 14 /* Word-erase character [ICANON] */
157: #define VLNEXT 15 /* Literal-next character [IEXTEN] */
158: #define VEOL2 16 /* Second EOL character [ICANON] */
159:
160: /* Values for the ACTION argument to 'tcflow'. */
161: #define TCOOFF 0 /* Suspend output */
162: #define TCOON 1 /* Restart suspended output */
163: #define TCIOFF 2 /* Send a STOP character */
164: #define TCION 3 /* Send a START character */
165:
166: /* Values for the QUEUE_SELECTOR argument to 'tcflush'. */
167: #define TCIFLUSH 0 /* Discard data received but not yet read */
168: #define TCOFLUSH 1 /* Discard data written but not yet sent */
169: #define TCIOFLUSH 2 /* Discard all pending data */
170:
171: /* Values for the OPTIONAL_ACTIONS argument to 'tcsetattr'. */
172: #define TCSANOW 0 /* Change immediately */
173: #define TCSADRAIN 1 /* Change when pending output is written */
174: #define TCSAFLUSH 2 /* Flush pending input before changing */
175:
176: #endif /* _FIWIX_TERMBITS_H */

```

**include/fiwix/termios.h**

Page 1/1

```

1:  /*
2:  *  fiwix/include/fiwix/termios.h
3:  *
4:  *  Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  *  Distributed under the terms of the Fiwix License.
6:  */
7:
8:  #ifndef _FIWIX_TERMIOS_H
9:  #define _FIWIX_TERMIOS_H
10:
11:  #include <fiwix/termbits.h>
12:
13:  struct winsize {
14:      unsigned short int ws_row;
15:      unsigned short int ws_col;
16:      unsigned short int ws_xpixel;
17:      unsigned short int ws_ypixel;
18:  };
19:
20:
21:  #define NCC      8
22:
23:  /* old terminal control structure */
24:  struct termio {
25:      unsigned short int c_iflag;      /* input mode flags */
26:      unsigned short int c_oflag;      /* output mode flags */
27:      unsigned short int c_cflag;      /* control mode flags */
28:      unsigned short int c_lflag;      /* local mode flags */
29:      unsigned char c_line;            /* line discipline */
30:      unsigned char c_cc[NCC];         /* control characters */
31:  };
32:
33:
34:  #define NCCS 19
35:
36:  /* new terminal control structure */
37:  struct termios {
38:      tcflag_t c_iflag;      /* input mode flags */
39:      tcflag_t c_oflag;      /* output mode flags */
40:      tcflag_t c_cflag;      /* control mode flags */
41:      tcflag_t c_lflag;      /* local mode flags */
42:      cc_t c_line;           /* line discipline */
43:      cc_t c_cc[NCCS];       /* control characters */
44:  };
45:
46:  #endif /* _FIWIX_TERMIOS_H */

```

**include/fiwix/timeb.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/timeb.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_TIMEB_H
9: #define _FIWIX_TIMEB_H
10:
11: struct timeb {
12:     unsigned int time;           /* in seconds since Epoch */
13:     unsigned short int millitm; /* additional milliseconds */
14:     short int timezone;         /* minutes west of GMT */
15:     short int dstflag;         /* nonzero if DST is used */
16: };
17:
18: #endif /* _FIWIX_TIMEB_H */
```



## include/fiwix/time.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/time.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_TIME_H
9: #define _FIWIX_TIME_H
10:
11: #define ITIMER_REAL    0
12: #define ITIMER_VIRTUAL 1
13: #define ITIMER_PROF   2
14:
15: struct timespec {
16:     long int tv_sec;          /* seconds since 00:00:00, 1 Jan 1970 UTC */
17:     long int tv_nsec;        /* nanoseconds (1000000000ns = 1sec) */
18: };
19:
20: struct timeval {
21:     long int tv_sec;          /* seconds since 00:00:00, 1 Jan 1970 UTC */
22:     long int tv_usec;        /* microseconds (1000000us = 1sec) */
23: };
24:
25: struct timezone {
26:     int tz_minuteswest;      /* minutes west of GMT */
27:     int tz_dsttime;          /* type of DST correction */
28: };
29:
30: struct itimerval {
31:     struct timeval it_interval;
32:     struct timeval it_value;
33: };
34:
35: struct mt {
36:     int mt_sec;
37:     int mt_min;
38:     int mt_hour;
39:     int mt_day;
40:     int mt_month;
41:     int mt_year;
42: };
43:
44: unsigned long int tv2ticks(const struct timeval *);
45: void ticks2tv(long int, struct timeval *);
46: int setitimer(int, const struct itimerval *, struct itimerval *);
47: unsigned long int mktime(struct mt *);
48:
49: #endif /* _FIWIX_TIME_H */

```

**include/fiwix/timer.h**

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/timer.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_TIMER_H
9: #define _FIWIX_TIMER_H
10:
11: #include <fiwix/types.h>
12: #include <fiwix/sigcontext.h>
13:
14: #define TIMER_IRQ      0
15: #define HZ              100      /* kernel's Hertz rate (100 = 10ms) */
16: #define TICK           (1000000 / HZ)
17:
18: #define UNIX_EPOCH     1970
19:
20: #define LEAP_YEAR(y)   ((y % 4) == 0 && ((y % 100) != 0 || (y % 400) == 0))
21: #define DAYS_PER_YEAR(y) ((LEAP_YEAR(y)) ? 366 : 365)
22:
23: #define SECS_PER_MIN   60
24: #define SECS_PER_HOUR (SECS_PER_MIN * 60)
25: #define SECS_PER_DAY   (SECS_PER_HOUR * 24)
26:
27: #define INFINITE_WAIT  0xFFFFFFFF
28:
29: struct callout {
30:     int expires;
31:     void (*fn)(unsigned int);
32:     unsigned int arg;
33:     struct callout *next;
34: };
35:
36: struct callout_req {
37:     void (*fn)(unsigned int);
38:     unsigned int arg;
39: };
40:
41: void add_callout(struct callout_req *, unsigned int);
42: void del_callout(struct callout_req *);
43: void irq_timer(int, struct sigcontext *);
44: void irq_timer_bh(void);
45: void do_callouts_bh(void);
46: void get_system_time(void);
47: void set_system_time(__time_t);
48: void timer_init(void);
49:
50: #endif /* _FIWIX_TIMER_H */

```

**include/fiwix/times.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/times.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_TIMES_H
9: #define _FIWIX_TIMES_H
10:
11: struct tms {
12:     __clock_t tms_utime;    /* CPU time spent in user-mode */
13:     __clock_t tms_stime;    /* CPU time spent in kernel-mode */
14:     __clock_t tms_cutime;   /* (tms_utime + tms_cutime) of children */
15:     __clock_t tms_cstime;   /* (tms_stime + tms_cstime) of children */
16: };
17:
18: #endif /* _FIWIX_TIMES_H */
```

## include/fiwix/traps.h

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/traps.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_TRAPS_H
9: #define _FIWIX_TRAPS_H
10:
11: #include <fiwix/sigcontext.h>
12:
13: #define NR_EXCEPTIONS    32
14:
15: struct traps {
16:     char *name;
17:     void (*handler)(unsigned int, struct sigcontext *);
18:     char errcode;
19: };
20:
21: void do_divide_error(unsigned int, struct sigcontext *);
22: void do_debug(unsigned int, struct sigcontext *);
23: void do_nmi_interrupt(unsigned int, struct sigcontext *);
24: void do_breakpoint(unsigned int, struct sigcontext *);
25: void do_overflow(unsigned int, struct sigcontext *);
26: void do_bound(unsigned int, struct sigcontext *);
27: void do_invalid_opcode(unsigned int, struct sigcontext *);
28: void do_no_math_coprocessor(unsigned int, struct sigcontext *);
29: void do_double_fault(unsigned int, struct sigcontext *);
30: void do_coprocessor_segment_overrun(unsigned int, struct sigcontext *);
31: void do_invalid_tss(unsigned int, struct sigcontext *);
32: void do_segment_not_present(unsigned int, struct sigcontext *);
33: void do_stack_segment_fault(unsigned int, struct sigcontext *);
34: void do_general_protection(unsigned int, struct sigcontext *);
35: void do_page_fault(unsigned int, struct sigcontext *);
36: void do_reserved(unsigned int, struct sigcontext *);
37: void do_floating_point_error(unsigned int, struct sigcontext *);
38: void do_alignment_check(unsigned int, struct sigcontext *);
39: void do_machine_check(unsigned int, struct sigcontext *);
40: void do_simd_fault(unsigned int, struct sigcontext *);
41:
42: void trap_handler(unsigned int, struct sigcontext);
43:
44: int dump_registers(unsigned int, struct sigcontext *);
45:
46: #endif /* _FIWIX_TRAPS_H */
```

## include/fiwix/tty.h

Page 1/2

```

1: /*
2:  * fiwix/include/fiwix/tty.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_TTY_H
9: #define _FIWIX_TTY_H
10:
11: #include <fiwix/termios.h>
12: #include <fiwix/fs.h>
13: #include <fiwix/console.h>
14: #include <fiwix/serial.h>
15:
16: #define NR_TTYS          NR_VCONSOLES + NR_SERIAL
17:
18: #define CBSIZE          32      /* number of characters in cblock */
19: #define NR_CB_QUEUE    10      /* number of cblocks per queue */
20: #define CB_POOL_SIZE   128     /* number of cblocks in the central pool */
21:
22: #define TAB_SIZE       8
23: #define MAX_TAB_COLS   132     /* maximum number of tab stops */
24:
25: #define LAST_CHAR(q)   ((q)->tail ? (q)->tail->data[(q)->tail->end_off - 1] : N
ULL)
26:
27: struct clist {
28:     unsigned short int count;
29:     unsigned short int cb_num;
30:     struct cblock *head;
31:     struct cblock *tail;
32: };
33:
34: struct cblock {
35:     unsigned short int start_off;
36:     unsigned short int end_off;
37:     unsigned char data[CBSIZE];
38:     struct cblock *prev;
39:     struct cblock *next;
40: };
41:
42: struct kbd_state {
43:     char mode;
44: };
45:
46: struct tty {
47:     __dev_t dev;
48:     struct clist read_q;
49:     struct clist cooked_q;
50:     struct clist write_q;
51:     short int count;
52:     struct termios termios;
53:     struct winsize winsize;
54:     struct kbd_state kbd;
55:     __pid_t pid, pgid, sid;
56:     unsigned char lnext;
57:     void *driver_data;
58:     int canon_data;
59:     char tab_stop[132];
60:     int column;
61:
62:     /* formerly tty driver operations */
63:     void (*stop)(struct tty *);
64:     void (*start)(struct tty *);
65:     void (*deltab)(struct tty *);
66:     void (*reset)(struct tty *);

```

**include/fiwix/tty.h**

Page 2/2

```
67:         void (*input)(struct tty *);
68:         void (*output)(struct tty *);
69:         int (*open)(struct tty *);
70:         int (*close)(struct tty *);
71:         void (*set_termios)(struct tty *);
72:     };
73: extern struct tty tty_table[];
74:
75: int register_tty(__dev_t);
76: struct tty * get_tty(__dev_t);
77: void disassociate_ctty(struct tty *);
78: void termios_reset(struct tty *);
79: void do_cook(struct tty *);
80: int tty_putchar(struct tty *, unsigned char);
81: int tty_open(struct inode *, struct fd *);
82: int tty_close(struct inode *, struct fd *);
83: int tty_read(struct inode *, struct fd *, char *, __size_t);
84: int tty_write(struct inode *, struct fd *, const char *, __size_t);
85: int tty_ioctl(struct inode *, int cmd, unsigned long int);
86: int tty_lseek(struct inode *, __off_t);
87: int tty_select(struct inode *, int);
88: void tty_init(void);
89:
90: int tty_queue_putchar(struct tty *, struct clist *, unsigned char);
91: int tty_queue_unputchar(struct clist *);
92: unsigned char tty_queue_getchar(struct clist *);
93: void tty_queue_flush(struct clist *);
94: int tty_queue_room(struct clist *q);
95: void tty_queue_init(struct tty *);
96:
97: int vt_ioctl(struct tty *, int, unsigned long int);
98:
99: #endif /* _FIWIX_TTY_H */
```

## include/fiwix/types.h

Page 1/1

```

1: /*
2:  * fiwix/include/fiwix/types.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_TYPES_H
9: #define _FIWIX_TYPES_H
10:
11: typedef __signed__ char __s8;
12: typedef unsigned char __u8;
13: typedef __signed__ short int __s16;
14: typedef unsigned short int __u16;
15: typedef __signed__ int __s32;
16: typedef unsigned int __u32;
17:
18: typedef __u16 __uid_t;
19: typedef __u16 __gid_t;
20: typedef __u32 __ino_t;
21: typedef __u16 __mode_t;
22: typedef __u16 __nlink_t;
23: typedef __u32 __off_t;
24: typedef __s32 __pid_t;
25: typedef __s32 __ssize_t;
26: typedef __u32 __size_t;
27: typedef unsigned long int __clock_t;
28: typedef __u32 __time_t;
29: typedef __u16 __dev_t;
30: typedef __u16 __key_t;
31: typedef __s32 __blk_t;          /* must be signed in order to return -EIO */
32: typedef __s32 __daddr_t;
33: typedef unsigned long long int __loff_t;
34:
35: /* number of descriptors that can fit in an 'fd_set' */
36: /* WARNING: this value must be the same as in the C Library */
37: #define __FD_SETSIZE      64
38:
39: #define __NFDBITS          (sizeof(unsigned long int) * 8)
40: #define __FDELT(d)        ((d) / __NFDBITS)
41: #define __FDMASK(d)       (1 << ((d) % __NFDBITS))
42:
43: /* define the fd_set structure for select() */
44: typedef struct {
45:     unsigned long int fds_bits[__FD_SETSIZE / __NFDBITS];
46: } fd_set;
47:
48: #define __FD_ZERO(set)      (memset_b((void *) (set), 0, sizeof(fd_set)))
49: #define __FD_SET(d, set)    ((set)->fds_bits[__FDELT(d)] |= __FDMASK(d))
50: #define __FD_CLR(d, set)    ((set)->fds_bits[__FDELT(d)] &= ~__FDMASK(d))
51: #define __FD_ISSET(d, set)  ((set)->fds_bits[__FDELT(d)] & __FDMASK(d))
52:
53: #endif /* _FIWIX_TYPES_H */

```

## include/fiwix/unistd.h

Page 1/4

```

1: /*
2:  * fiwix/include/fiwix/unistd.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_UNISTD_H
9: #define _FIWIX_UNISTD_H
10:
11: /* Linux 2.0.40 ABI (one syscall from Linux 2.2.26) */
12: /* #define SYS_setup */
13: #define SYS_exit 1
14: #define SYS_fork 2
15: #define SYS_read 3
16: #define SYS_write 4
17: #define SYS_open 5
18: #define SYS_close 6
19: #define SYS_waitpid 7
20: #define SYS_creat 8
21: #define SYS_link 9
22: #define SYS_unlink 10
23: #define SYS_execve 11
24: #define SYS_chdir 12
25: #define SYS_time 13
26: #define SYS_mknod 14
27: #define SYS_chmod 15
28: #define SYS_chown 16
29: #define SYS_break 17 /* -ENOSYS */
30: #define SYS_oldstat 18
31: #define SYS_lseek 19
32: #define SYS_getpid 20
33: #define SYS_mount 21
34: #define SYS_umount 22
35: #define SYS_setuid 23
36: #define SYS_getuid 24
37: #define SYS_stime 25
38: /* #define SYS_ptrace */
39: #define SYS_alarm 27
40: #define SYS_oldfstat 28
41: #define SYS_pause 29
42: #define SYS_utime 30
43: #define SYS_stty 31 /* -ENOSYS */
44: #define SYS_gtty 32 /* -ENOSYS */
45: #define SYS_access 33
46: /* #define SYS_nice */
47: #define SYS_ftime 35
48: #define SYS_sync 36
49: #define SYS_kill 37
50: #define SYS_rename 38
51: #define SYS_mkdir 39
52: #define SYS_rmdir 40
53: #define SYS_dup 41
54: #define SYS_pipe 42
55: #define SYS_times 43
56: #define SYS_prof 44 /* -ENOSYS */
57: #define SYS_brk 45
58: #define SYS_setgid 46
59: #define SYS_getgid 47
60: #define SYS_signal 48
61: #define SYS_geteuid 49
62: #define SYS_getegid 50
63: /* #define SYS_acct */
64: #define SYS_umount2 52 /* from 2.2.26, it was sys_phys() */
65: #define SYS_lock 53 /* -ENOSYS */
66: #define SYS_ioctl 54
67: #define SYS_fcntl 55

```



## include/fiwix/unistd.h

Page 2/4

```

68: #define SYS_mpx                56                /* -ENOSYS */
69: #define SYS_setpgid            57
70: #define SYS_ulimit              58                /* -ENOSYS */
71: #define SYS_olduname           59
72: #define SYS_umask               60
73: #define SYS_chroot              61
74: #define SYS_ustat               62
75: #define SYS_dup2                63
76: #define SYS_getppid            64
77: #define SYS_getpgrp            65
78: #define SYS_setsid             66
79: #define SYS_sigaction          67
80: #define SYS_sgetmask           68
81: #define SYS_ssetmask           69
82: #define SYS_setreuid           70
83: #define SYS_setregid           71
84: #define SYS_sigsuspend         72
85: #define SYS_sigpending         73
86: #define SYS_sethostname        74
87: #define SYS_setrlimit          75
88: #define SYS_getrlimit          76
89: #define SYS_getrusage          77
90: #define SYS_gettimeofday       78
91: #define SYS_settimeofday       79
92: #define SYS_getgroups          80
93: #define SYS_setgroups          81
94: #define SYS_oldselect          82
95: #define SYS_symlink            83
96: #define SYS_oldlstat           84
97: #define SYS_readlink           85
98: /* #define SYS_uselib */
99: /* #define SYS_swapon */
100: #define SYS_reboot              88
101: /* #define SYS_oldreaddir */
102: #define SYS_old_mmap           90
103: #define SYS_munmap             91
104: #define SYS_truncate           92
105: #define SYS_ftruncate          93
106: #define SYS_fchmod             94
107: #define SYS_fchown             95
108: /* #define SYS_getpriority */
109: /* #define SYS_setpriority */
110: /* #define SYS_profil */
111: #define SYS_statfs              99
112: #define SYS_fstatfs            100
113: #define SYS_ioperm             101
114: #define SYS_socketcall         102
115: /* #define SYS_syslog */
116: #define SYS_setitimer          104
117: #define SYS_getitimer          105
118: #define SYS_newstat            106
119: #define SYS_newlstat           107
120: #define SYS_newfstat           108
121: #define SYS_uname              109
122: #define SYS_iopl               110
123: /* #define SYS_vhangup */
124: /* #define SYS_idle              112                -ENOSYS */
125: /* #define SYS_vm86old */
126: #define SYS_wait4              114
127: /* #define SYS_swapoff */
128: #define SYS_sysinfo            116
129: /* #define SYS_ipc */
130: #define SYS_fsync              118
131: #define SYS_sigreturn          119
132: /* #define SYS_clone */
133: #define SYS_setdomainname      121
134: #define SYS_newuname           122

```

**include/fiwix/unistd.h**

Page 3/4

```

135: /* #define SYS_modify_ldt */
136: /* #define SYS_adjtimex */
137: #define SYS_mprotect          125
138: #define SYS_sigprocmask      126
139: /* #define SYS_create_module */
140: /* #define SYS_init_module */
141: /* #define SYS_delete_module */
142: /* #define SYS_get_kernel_syms */
143: /* #define SYS_quotactl */
144: #define SYS_getpgid          132
145: #define SYS_fchdir           133
146: /* #define SYS_bdflush */
147: /* #define SYS_sysfs */
148: #define SYS_personality      136
149: /* #define afs_syscall */
150: #define SYS_setfsuid         138
151: #define SYS_setfsgid         139
152: #define SYS_llseek           140
153: #define SYS_getdents         141
154: #define SYS_select           142
155: #define SYS_flock            143
156: /* #define SYS_msync */
157: /* #define SYS_readv */
158: /* #define SYS_writev */
159: #define SYS_getsid           147
160: #define SYS_fdatasync        148
161: /* #define SYS_sysctl */
162: /* #define SYS_mlock */
163: /* #define SYS_munlock */
164: /* #define SYS_mlockall */
165: /* #define SYS_munlockall */
166: /* #define SYS_sched_setparam */
167: /* #define SYS_sched_getparam */
168: /* #define SYS_sched_setscheduler */
169: /* #define SYS_sched_getscheduler */
170: /* #define SYS_sched_yield */
171: /* #define SYS_sched_get_priority_max */
172: /* #define SYS_sched_get_priority_min */
173: /* #define SYS_sched_rr_get_interval */
174: #define SYS_nanosleep        162
175: /* #define SYS_mremap */
176:
177: /* extra system calls from Linux 2.2.26 */
178: /* #define SYS_mremap */
179: /* #define SYS_setresuid */
180: /* #define SYS_getresuid */
181: /* #define SYS_ni_syscall */
182: /* #define SYS_query_module */
183: /* #define SYS_poll */
184: /* #define SYS_nfsservctl */
185: /* #define SYS_setresgid */
186: /* #define SYS_getresgid */
187: /* #define SYS_prctl */
188: /* #define SYS_rt_sigreturn_wrapper */
189: /* #define SYS_rt_sigaction */
190: /* #define SYS_rt_sigprocmask */
191: /* #define SYS_rt_sigpending */
192: /* #define SYS_rt_sigtimedwait */
193: /* #define SYS_rt_sigqueueinfo */
194: /* #define SYS_rt_sigsuspend_wrapper */
195: /* #define SYS_pread */
196: /* #define SYS_pwrite */
197: /* #define SYS_chown */
198: #define SYS_getcwd           183
199: /* #define SYS_capget */
200: /* #define SYS_capset */
201: /* #define SYS_sigaltstack_wrapper */

```

**include/fiwix/unistd.h**

Page 4/4

```
202: /* #define SYS_sendfile */
203: /* #define SYS_ni_syscall */
204: /* #define SYS_ni_syscall */
205: #define SYS_vfork                190
206:
207: #endif /* _FIWIX_UNISTD_H */
```

**include/fiwix/ustat.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/ustat.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_USTAT_H
9: #define _FIWIX_USTAT_H
10:
11: #include <fiwix/types.h>
12:
13: struct ustat {
14:     __daddr_t f_tfree;           /* total free blocks */
15:     __ino_t f_tinode;          /* number of free inodes */
16:     char f_fname;              /* filesystem name */
17:     char f_fpack;             /* filesystem pack name */
18: };
19:
20: #endif /* _FIWIX_USTAT_H */
```

**include/fiwix/utime.h**

Page 1/1

```
1: /*
2:  * fiwix/include/fiwix/utime.h
3:  *
4:  * Copyright 2018, Jordi Sanfeliu. All rights reserved.
5:  * Distributed under the terms of the Fiwix License.
6:  */
7:
8: #ifndef _FIWIX_UTIME_H
9: #define _FIWIX_UTIME_H
10:
11: #include <fiwix/types.h>
12:
13: struct utimbuf {
14:     __time_t actime;           /* access time */
15:     __time_t modtime;        /* modification time */
16: };
17:
18: #endif /* _FIWIX_UTIME_H */
```

## include/fiwix/utsname.h

Page 1/2

```

1:  /* Copyright (C) 1991, 1992, 1994, 1996 Free Software Foundation, Inc.
2:    This file is part of the GNU C Library.
3:
4:    The GNU C Library is free software; you can redistribute it and/or
5:    modify it under the terms of the GNU Library General Public License as
6:    published by the Free Software Foundation; either version 2 of the
7:    License, or (at your option) any later version.
8:
9:    The GNU C Library is distributed in the hope that it will be useful,
10:   but WITHOUT ANY WARRANTY; without even the implied warranty of
11:   MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the GNU
12:   Library General Public License for more details.
13:
14:   You should have received a copy of the GNU Library General Public
15:   License along with the GNU C Library; see the file COPYING.LIB.  If not,
16:   write to the Free Software Foundation, Inc., 59 Temple Place - Suite 330,
17:   Boston, MA 02111-1307, USA.  */
18:
19:  /*
20:   *      POSIX Standard: 4.4 System Identification      <fiwix/utsname.h>
21:   */
22:
23:  #ifndef _FIWIX_UTSNAME_H
24:  #define _FIWIX_UTSNAME_H
25:
26:  #define _OLD_UTSNAME_LENGTH      9
27:  #define _UTSNAME_LENGTH        65
28:
29:  #ifndef _UTSNAME_NODENAME_LENGTH
30:  #define _UTSNAME_NODENAME_LENGTH _UTSNAME_LENGTH
31:  #endif
32:
33:  /* Very OLD structure describing the system and machine.  */
34:  struct oldold_utsname
35:  {
36:      char sysname[_OLD_UTSNAME_LENGTH];
37:      char nodename[_OLD_UTSNAME_LENGTH];
38:      char release[_OLD_UTSNAME_LENGTH];
39:      char version[_OLD_UTSNAME_LENGTH];
40:      char machine[_OLD_UTSNAME_LENGTH];
41:  };
42:
43:  /* OLD structure describing the system and machine.  */
44:  struct old_utsname
45:  {
46:      char sysname[_UTSNAME_LENGTH];
47:      char nodename[_UTSNAME_NODENAME_LENGTH];
48:      char release[_UTSNAME_LENGTH];
49:      char version[_UTSNAME_LENGTH];
50:      char machine[_UTSNAME_LENGTH];
51:  };
52:
53:  /* NEW structure describing the system and machine.  */
54:  struct new_utsname
55:  {
56:      /* Name of the implementation of the operating system.  */
57:      char sysname[_UTSNAME_LENGTH];
58:
59:      /* Name of this node on the network.  */
60:      char nodename[_UTSNAME_NODENAME_LENGTH];
61:
62:      /* Current release level of this implementation.  */
63:      char release[_UTSNAME_LENGTH];
64:      /* Current version level of this release.  */
65:      char version[_UTSNAME_LENGTH];
66:
67:      /* Name of the hardware type the system is running on.  */

```

**include/fiwix/utsname.h**

Page 2/2

```
68:     char machine[_UTSNAME_LENGTH];
69:     char domainname[_UTSNAME_LENGTH];
70: };
71:
72: extern struct new_utsname sys_utsname;
73: extern char UTS_MACHINE[_UTSNAME_LENGTH];
74:
75: #endif /* _FIWIX_UTSNAME_H */
```

**include/fiwix/version.h**

Page 1/1

```
1: #define UTS_VERSION "Sat 12 Dec 2020 09:54:04 AM CET"
```



## include/fiwix/vt.h

Page 1/1

```

1: #ifndef VT_H
2: #define VT_H
3:
4: /* prefix 0x56 is 'V', to avoid collision with terminos and kd */
5:
6: #define VT_OPENQRY      0x5600 /* find available vt */
7:
8: struct vt_mode {
9:     char mode; /* vt mode */
10:    char waitv; /* if set, hang on writes if not active */
11:    short int relsig; /* signal to raise on release req */
12:    short int acqsig; /* signal to raise on acquisition */
13:    short int frsig; /* unused (set to 0) */
14: };
15: #define VT_GETMODE      0x5601 /* get mode of active vt */
16: #define VT_SETMODE      0x5602 /* set mode of active vt */
17: #define VT_AUTO         0x00 /* auto vt switching */
18: #define VT_PROCESS      0x01 /* process controls switching */
19: #define VT_ACKACQ       0x02 /* acknowledge switch */
20:
21: struct vt_stat {
22:    unsigned short int v_active; /* active vt */
23:    unsigned short int v_signal; /* signal to send */
24:    unsigned short int v_state; /* vt bitmask */
25: };
26: #define VT_GETSTATE     0x5603 /* get global vt state info */
27: #define VT_SENDSIG      0x5604 /* signal to send to bitmask of vts */
28:
29: #define VT_RELDISP      0x5605 /* release display */
30:
31: #define VT_ACTIVATE     0x5606 /* make vt active */
32: #define VT_WAITACTIVE   0x5607 /* wait for vt active */
33: #define VT_DISALLOCATE  0x5608 /* free memory associated to vt */
34:
35: struct vt_sizes {
36:    unsigned short int v_rows; /* number of rows */
37:    unsigned short int v_cols; /* number of columns */
38:    unsigned short int v_scrollsize; /* number of lines of scrollbar */
39: };
40: #define VT_RESIZE       0x5609 /* set kernel's idea of screensize */
41:
42: struct vt_consize {
43:    unsigned short int v_rows; /* number of rows */
44:    unsigned short int v_cols; /* number of columns */
45:    unsigned short int v_vlin; /* number of pixel rows on screen */
46:    unsigned short int v_clin; /* number of pixel rows per character */
47:    unsigned short int v_vcol; /* number of pixel columns on screen */
48:    unsigned short int v_ccol; /* number of pixel columns per character
*/
49: };
50: #define VT_RESIZEX      0x560A /* set kernel's idea of screensize + more */
51: #define VT_LOCKSWITCH  0x560B /* disallow vt switching */
52: #define VT_UNLOCKSWITCH 0x560C /* allow vt switching */
53:
54: #endif /* VT_H */

```

**README.md**

Page 1/2

```

1: Fiwix
2: =====
3: Fiwix is an operating system kernel written from scratch, based on the UNIX arch
itecure and fully focused on being POSIX compatible. It is designed and developed main
ly as a hobby OS and, since it serves also for educational purposes, the kernel code is
kept as simple as possible for the benefit of students and OS enthusiasts. It is small
in size (less than 40,000 lines of code), runs on the i386 hardware platform and is co
mpatible with a good base of existing GNU applications.
4:
5: Features
6: -----
7: - Mostly written in C language (Assembly only used in the needed parts).
8: - GRUB Multiboot Specification v1 compliant.
9: - Full 32bit protected mode non-preemptive kernel.
10: - For i386 processors and higher.
11: - Preemptive multitasking.
12: - Protected task environment (independent memory address per process).
13: - POSIX-compliant (mostly).
14: - Process groups, sessions and job control.
15: - Interprocess communication with pipes and signals.
16: - BSD file locking mechanism (POSIX restricted to file and advisory only).
17: - Virtual memory management up to 4GB (1GB physical only and no swapping yet).
18: - Demand paging with Copy-On-Write feature.
19: - Linux 2.0 ABI system calls compatibility (mostly).
20: - ELF-386 executable format support (statically and dynamically linked).
21: - Round Robin based scheduler algorithm (no priorities yet).
22: - VFS abstraction layer.
23: - EXT2 filesystem support with 1KB, 2KB and 4KB block sizes.
24: - Minix v1 and v2 filesystem support.
25: - Linux-like PROC filesystem support (read only).
26: - PIPE pseudo-filesystem support.
27: - ISO9660 filesystem support with Rock Ridge extensions.
28: - RAMdisk device support.
29: - Initial RAMdisk (initrd) image support.
30: - SVGAlib based applications support.
31: - Keyboard driver with Linux keymaps support.
32: - Serial port (RS-232) driver support.
33: - Remote serial console support.
34: - Parallel port printer driver support.
35: - Basic implementation of a Pseudo-Random Number Generator.
36: - Floppy disk device driver and DMA management.
37: - IDE/ATA ATAPI CD-ROM device driver.
38: - IDE/ATA hard disk device driver.
39:
40: Compiling
41: -----
42: The command needed to build the Fiwix kernel is `make clean ; make`. This will
create the files **fiwix** (the kernel itself) and **System.map.gz** (the symbol table)
in the root directory of the source code tree.
43:
44: Before compiling you might want to tweak the kernel configuration by changing th
e default values in `include/fiwix/config.h`.
45:
46: Keep in mind that the kernel doesn't do anything on its own, you need to create
a user-space environment to make use of it. Upon booting, the kernel mounts the root fi
lesystem and tries to run `/sbin/init` on it, so you would need to provide this program
yourself. Fortunately, FiwixOS provides a full user-space UNIX-like environment to te
st the Fiwix kernel.
47:
48: Installing
49: -----
50: You can proceed to install [FiwixOS](https://www.fiwix.org/downloads.html) on a
hard disk either by booting from the CD-ROM or from a floppy. If you chosen the latter,
you will also need the Installation CD-ROM inserted in order to install the packages t
hat form all the system environment.
51:
52: Let the system boot either from the CD-ROM or floppy and, when you are ready, ju

```

**README.md**

Page 2/2

```
st type `install.sh`.
53:
54: The minimal hardware requirements are as follows:
55:
56: - Standard IBM PC-AT architecture.
57: - i386 processor or higher.
58: - 3MB of RAM memory (128MB recommended).
59: - IDE/ATAPI CD-ROM or floppy disk (3.5", 1.44MB).
60: - 1GB ATA hard disk.
61:
62: Please keep in mind that this is a kernel in its very early stages and may well
have serious bugs and broken features which have not yet been identified or resolved.
63:
64: Let me repeat that.
65:
66: Please keep in mind that this is a kernel in its very early stages and may well
have serious bugs and broken features which have not yet been identified or resolved.
67:
68: *****
69: *** USE AT YOUR OWN RISK! ***
70: *****
71:
72: References
73: -----
74: - [Website] (https://www.fiwix.org)
75: - [Mailing List] (https://lists.sourceforge.net/lists/listinfo/fiwix-general)
76: - [IRC] (https://webchat.freenode.net/?channels=fiwix)
77:
78: License
79: -----
80: Fiwix is free software licensed under the terms of the MIT License, see the LICE
NSE file for more details.
81: Copyright (C) 2018-2020, Jordi Sanfeliu.
82:
83: Credits
84: -----
85: Fiwix was created by [Jordi Sanfeliu] (https://www.fibranet.cat).
86: You can contact me at [jordi@fibranet.cat] (mailto:jordi@fibranet.cat).
87:
88:
```

**LICENSE**

Page 1/1

1: MIT License  
2:  
3: **Copyright** (c) 2018 Jordi Sanfeliu  
4:  
5: Permission is hereby granted, free of charge, to any person obtaining a copy  
6: of this software and associated documentation **files** (the "**Software**"), to deal  
7: in the Software without restriction, including without limitation the rights  
8: to use, copy, modify, merge, publish, distribute, sublicense, and/or sell  
9: copies of the Software, and to permit persons to whom the Software is  
10: furnished to **do** so, subject to the following conditions:  
11:  
12: The above copyright notice and this permission notice shall be included in all  
13: copies or substantial portions of the Software.  
14:  
15: THE SOFTWARE IS PROVIDED "**AS IS**", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR  
16: IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,  
17: FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE  
18: AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER  
19: LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM,  
20: OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE  
21: SOFTWARE.  
22:

**CODING**

Page 1/1

```
1: Fiwix kernel coding standards
2: -----
3: It's easier on everyone if all authors working on a shared code base are
4: consistent in the way they write their programs. Fiwix has the following
5: conventions in its code:
6:
7: - Use of snake_case (multi-word names are lower_case_with_underscores) for
8:   everything except for macro and constant names.
9:
10: - No space after the name of a function in a call.
11:   For example, printk("hello") not printk ("hello").
12:
13: - Optional space after keywords "if", "for", "while", "switch".
14:   For example, if(x) or if (x).
15:
16: - Space before braces (except in function declarations).
17:   For example, if(x) { not if(x){.
18:
19: - Space between operands.
20:   For example, for(n = 0; n < 10; n++), not for(n=0;n<10;n++).
21:
22: - Beginning-of-line indentation via tabs, not spaces.
23:
24: - Preprocessor macros are always UPPERCASE.
25:
26: - Pointer types have spaces: (uint16_t *) not (uint16_t*).
27:
28: - Comments in code are always as in C89 /* ... */.
29:
30: - Multiline comments start always with a capital letter and the last sentence
31:   ends with a period.
32:
33: - Functions that take no arguments are declared as f(void) not f().
34:
35: - The open parenthesis is always on the same line as the function name.
36:
37: - There is never a space between the function name and the open parenthesis.
38:
39: - There is never a space between the parentheses and the parameters.
40:
41: - The open curly brace is always at the next line of the function declaration.
42:
43: - Conditionals should always include curly braces, even if there is only a
44:   single statement within the conditional.
45:
46:
47: Recommended for reading
48: -----
49: - http://skull.piratehaven.org/~bapper/298c/cstyle.html
50:
```

**THANKS**

1: THANKS!  
2:  
3: A huge THANKS to all people who created their own hobby, and not so hobby,  
4: operating systems and made them freely available on Internet.  
5:  
6: From those simple "**Hello world!\n**" kernels to a real self-hosting operating  
7: system, their projects were a invaluable source of information and inspiration  
8: to create the Fiwix kernel.  
9:  
10: A special thanks to OSDEV Community **for** their tutorials, documents, wikis, etc.  
11:  
12: Jordi Sanfeliu  
13: **https://www.fiwix.org**  
14:

## Table of Contents

Page -1/5

|    |                                |    |       |     |       |          |          |
|----|--------------------------------|----|-------|-----|-------|----------|----------|
| 1  | Makefile                       | 1  | pages | 53  | lines | 20/12/08 | 22:11:03 |
| 2  | fiwix.ld                       | 1  | pages | 55  | lines | 18/04/21 | 21:19:46 |
| 3  | kernel/boot.S                  | 3  | pages | 151 | lines | 20/05/21 | 16:26:26 |
| 4  | kernel/cmos.c                  | 1  | pages | 57  | lines | 18/04/21 | 21:19:46 |
| 5  | kernel/core386.S               | 12 | pages | 801 | lines | 20/12/12 | 08:41:10 |
| 6  | kernel/cpu.c                   | 5  | pages | 263 | lines | 19/10/23 | 23:20:36 |
| 7  | kernel/gdt.c                   | 1  | pages | 54  | lines | 18/04/21 | 21:19:46 |
| 8  | kernel/idt.c                   | 2  | pages | 93  | lines | 20/06/08 | 07:56:45 |
| 9  | kernel/init.c                  | 2  | pages | 121 | lines | 20/03/07 | 19:06:51 |
| 10 | kernel/main.c                  | 7  | pages | 395 | lines | 20/10/26 | 17:46:01 |
| 11 | kernel/Makefile                | 1  | pages | 20  | lines | 18/04/21 | 21:19:46 |
| 12 | kernel/pic.c                   | 4  | pages | 233 | lines | 20/09/14 | 12:18:05 |
| 13 | kernel/pit.c                   | 1  | pages | 29  | lines | 18/04/21 | 21:19:46 |
| 14 | kernel/process.c               | 7  | pages | 398 | lines | 20/03/15 | 09:11:30 |
| 15 | kernel/sched.c                 | 2  | pages | 88  | lines | 20/03/09 | 16:57:30 |
| 16 | kernel/signal.c                | 4  | pages | 220 | lines | 20/11/27 | 12:30:02 |
| 17 | kernel/sleep.c                 | 3  | pages | 198 | lines | 20/04/21 | 08:28:25 |
| 18 | kernel/syscalls                | 0  | pages | 0   | lines | 20/12/12 | 09:58:04 |
| 19 | kernel/syscalls.c              | 7  | pages | 407 | lines | 20/04/30 | 12:05:40 |
| 20 | kernel/timer.c                 | 8  | pages | 459 | lines | 20/06/07 | 10:17:20 |
| 21 | kernel/traps.c                 | 6  | pages | 341 | lines | 20/12/04 | 19:29:05 |
| 22 | kernel/syscalls/access.c       | 1  | pages | 61  | lines | 18/04/21 | 21:19:46 |
| 23 | kernel/syscalls/alarm.c        | 1  | pages | 39  | lines | 18/04/21 | 21:19:46 |
| 24 | kernel/syscalls/brk.c          | 1  | pages | 54  | lines | 19/12/06 | 23:35:58 |
| 25 | kernel/syscalls/chdir.c        | 1  | pages | 49  | lines | 18/04/21 | 21:19:46 |
| 26 | kernel/syscalls/chmod.c        | 1  | pages | 56  | lines | 19/12/20 | 18:21:44 |
| 27 | kernel/syscalls/chown.c        | 2  | pages | 68  | lines | 18/04/21 | 21:19:46 |
| 28 | kernel/syscalls/chroot.c       | 1  | pages | 44  | lines | 18/04/21 | 21:19:46 |
| 29 | kernel/syscalls/close.c        | 1  | pages | 39  | lines | 18/04/21 | 21:19:46 |
| 30 | kernel/syscalls/creat.c        | 1  | pages | 23  | lines | 18/04/21 | 21:19:46 |
| 31 | kernel/syscalls/dup2.c         | 1  | pages | 41  | lines | 18/04/21 | 21:19:46 |
| 32 | kernel/syscalls/dup.c          | 1  | pages | 37  | lines | 18/04/21 | 21:19:46 |
| 33 | kernel/syscalls/execve.c       | 6  | pages | 379 | lines | 20/09/19 | 11:47:46 |
| 34 | kernel/syscalls/exit.c         | 2  | pages | 118 | lines | 19/10/13 | 11:15:15 |
| 35 | kernel/syscalls/fchdir.c       | 1  | pages | 34  | lines | 18/04/21 | 21:19:46 |
| 36 | kernel/syscalls/fchmod.c       | 1  | pages | 42  | lines | 18/04/21 | 21:19:46 |
| 37 | kernel/syscalls/fchown.c       | 1  | pages | 53  | lines | 18/04/21 | 21:19:46 |
| 38 | kernel/syscalls/fcntl.c        | 2  | pages | 64  | lines | 18/04/21 | 21:19:46 |
| 39 | kernel/syscalls/fdatasync.c    | 1  | pages | 22  | lines | 18/04/21 | 21:19:46 |
| 40 | kernel/syscalls/flock.c        | 1  | pages | 28  | lines | 18/04/21 | 21:19:46 |
| 41 | kernel/syscalls/fork.c         | 3  | pages | 149 | lines | 20/02/20 | 14:47:34 |
| 42 | kernel/syscalls/fstat.c        | 1  | pages | 44  | lines | 18/04/21 | 21:19:46 |
| 43 | kernel/syscalls/fstatfs.c      | 1  | pages | 36  | lines | 18/04/21 | 21:19:46 |
| 44 | kernel/syscalls/fsync.c        | 1  | pages | 39  | lines | 18/04/21 | 21:19:46 |
| 45 | kernel/syscalls/ftime.c        | 1  | pages | 34  | lines | 18/04/21 | 21:19:46 |
| 46 | kernel/syscalls/ftruncate.c    | 1  | pages | 53  | lines | 18/04/21 | 21:19:46 |
| 47 | kernel/syscalls/getcwd.c       | 1  | pages | 29  | lines | 18/04/21 | 21:19:46 |
| 48 | kernel/syscalls/getdents.c     | 1  | pages | 45  | lines | 18/04/21 | 21:19:46 |
| 49 | kernel/syscalls/getegid.c      | 1  | pages | 20  | lines | 18/04/21 | 21:19:46 |
| 50 | kernel/syscalls/geteuid.c      | 1  | pages | 20  | lines | 18/04/21 | 21:19:46 |
| 51 | kernel/syscalls/getgid.c       | 1  | pages | 20  | lines | 18/04/21 | 21:19:46 |
| 52 | kernel/syscalls/getgroups.c    | 1  | pages | 54  | lines | 20/01/11 | 19:01:23 |
| 53 | kernel/syscalls/getitimer.c    | 1  | pages | 48  | lines | 18/04/21 | 21:19:46 |
| 54 | kernel/syscalls/getpgid.c      | 1  | pages | 37  | lines | 18/04/21 | 21:19:46 |
| 55 | kernel/syscalls/getpgrp.c      | 1  | pages | 20  | lines | 18/04/21 | 21:19:46 |
| 56 | kernel/syscalls/getpid.c       | 1  | pages | 20  | lines | 18/04/21 | 21:19:46 |
| 57 | kernel/syscalls/getppid.c      | 1  | pages | 20  | lines | 18/04/21 | 21:19:46 |
| 58 | kernel/syscalls/getrlimit.c    | 1  | pages | 35  | lines | 18/04/21 | 21:19:46 |
| 59 | kernel/syscalls/getrusage.c    | 1  | pages | 40  | lines | 18/04/21 | 21:19:46 |
| 60 | kernel/syscalls/getsid.c       | 1  | pages | 38  | lines | 18/04/21 | 21:19:46 |
| 61 | kernel/syscalls/gettimeofday.c | 1  | pages | 41  | lines | 20/05/24 | 19:04:39 |
| 62 | kernel/syscalls/getuid.c       | 1  | pages | 21  | lines | 18/04/21 | 21:19:46 |
| 63 | kernel/syscalls/ioctl.c        | 1  | pages | 41  | lines | 18/04/21 | 21:19:46 |
| 64 | kernel/syscalls/ioperm.c       | 1  | pages | 28  | lines | 20/06/07 | 10:35:30 |
| 65 | kernel/syscalls/iopl.c         | 1  | pages | 58  | lines | 18/04/21 | 21:19:46 |
| 66 | kernel/syscalls/kill.c         | 1  | pages | 46  | lines | 18/07/08 | 20:28:28 |
| 67 | kernel/syscalls/link.c         | 2  | pages | 113 | lines | 18/04/21 | 21:19:46 |

## Table of Contents

|     |                                 |         |           |          |          |
|-----|---------------------------------|---------|-----------|----------|----------|
| 68  | kernel/syscalls/llseek.c        | 1 pages | 57 lines  | 20/09/15 | 11:56:46 |
| 69  | kernel/syscalls/lseek.c         | 1 pages | 58 lines  | 18/04/21 | 21:19:46 |
| 70  | kernel/syscalls/lstat.c         | 1 pages | 51 lines  | 18/04/21 | 21:19:46 |
| 71  | kernel/syscalls/Makefile        | 1 pages | 20 lines  | 18/04/21 | 21:19:46 |
| 72  | kernel/syscalls/mkdir.c         | 1 pages | 66 lines  | 18/04/21 | 21:19:46 |
| 73  | kernel/syscalls/mknod.c         | 2 pages | 71 lines  | 18/04/21 | 21:19:46 |
| 74  | kernel/syscalls/mount.c         | 4 pages | 221 lines | 18/04/21 | 21:19:46 |
| 75  | kernel/syscalls/mprotect.c      | 1 pages | 49 lines  | 18/04/21 | 21:19:46 |
| 76  | kernel/syscalls/munmap.c        | 1 pages | 22 lines  | 19/12/06 | 23:01:32 |
| 77  | kernel/syscalls/nanosleep.c     | 1 pages | 62 lines  | 20/12/02 | 12:06:48 |
| 78  | kernel/syscalls/newfstat.c      | 1 pages | 56 lines  | 20/09/26 | 17:03:57 |
| 79  | kernel/syscalls/newlstat.c      | 1 pages | 64 lines  | 18/04/21 | 21:19:46 |
| 80  | kernel/syscalls/newstat.c       | 1 pages | 64 lines  | 20/05/24 | 19:04:46 |
| 81  | kernel/syscalls/newuname.c      | 1 pages | 34 lines  | 18/04/21 | 21:19:46 |
| 82  | kernel/syscalls/old_mmap.c      | 1 pages | 52 lines  | 19/12/06 | 23:01:28 |
| 83  | kernel/syscalls/old_select.c    | 1 pages | 42 lines  | 18/04/21 | 21:19:46 |
| 84  | kernel/syscalls/olduname.c      | 1 pages | 39 lines  | 18/04/21 | 21:19:46 |
| 85  | kernel/syscalls/open.c          | 3 pages | 143 lines | 20/04/30 | 12:05:54 |
| 86  | kernel/syscalls/pause.c         | 1 pages | 30 lines  | 18/04/21 | 21:19:46 |
| 87  | kernel/syscalls/personality.c   | 1 pages | 19 lines  | 18/04/21 | 21:19:46 |
| 88  | kernel/syscalls/pipe.c          | 2 pages | 72 lines  | 19/11/09 | 10:13:27 |
| 89  | kernel/syscalls/read.c          | 1 pages | 49 lines  | 20/09/26 | 17:42:42 |
| 90  | kernel/syscalls/readlink.c      | 1 pages | 57 lines  | 18/04/21 | 21:19:46 |
| 91  | kernel/syscalls/reboot.c        | 1 pages | 49 lines  | 18/06/11 | 13:05:20 |
| 92  | kernel/syscalls/rename.c        | 2 pages | 117 lines | 18/04/21 | 21:19:46 |
| 93  | kernel/syscalls/rmdir.c         | 2 pages | 87 lines  | 18/04/21 | 21:19:46 |
| 94  | kernel/syscalls/select.c        | 3 pages | 166 lines | 20/05/24 | 19:04:35 |
| 95  | kernel/syscalls/setdomainname.c | 1 pages | 38 lines  | 18/04/21 | 21:19:46 |
| 96  | kernel/syscalls/setfsuid.c      | 1 pages | 21 lines  | 18/04/21 | 21:19:46 |
| 97  | kernel/syscalls/setfsuid.c      | 1 pages | 21 lines  | 18/04/21 | 21:19:46 |
| 98  | kernel/syscalls/setgid.c        | 1 pages | 32 lines  | 18/04/21 | 21:19:46 |
| 99  | kernel/syscalls/setgroups.c     | 1 pages | 41 lines  | 18/04/21 | 21:19:46 |
| 100 | kernel/syscalls/sethostname.c   | 1 pages | 42 lines  | 18/04/21 | 21:19:46 |
| 101 | kernel/syscalls/setitimer.c     | 1 pages | 31 lines  | 18/04/21 | 21:19:46 |
| 102 | kernel/syscalls/setpgid.c       | 2 pages | 67 lines  | 18/04/21 | 21:19:46 |
| 103 | kernel/syscalls/setregid.c      | 1 pages | 49 lines  | 20/05/05 | 16:52:26 |
| 104 | kernel/syscalls/setreuid.c      | 1 pages | 53 lines  | 20/05/05 | 16:50:01 |
| 105 | kernel/syscalls/setrlimit.c     | 1 pages | 42 lines  | 18/04/21 | 21:19:46 |
| 106 | kernel/syscalls/setsid.c        | 1 pages | 36 lines  | 18/04/21 | 21:19:46 |
| 107 | kernel/syscalls/settimeofday.c  | 1 pages | 46 lines  | 18/04/21 | 21:19:46 |
| 108 | kernel/syscalls/setuid.c        | 1 pages | 31 lines  | 18/04/21 | 21:19:46 |
| 109 | kernel/syscalls/sgetmask.c      | 1 pages | 20 lines  | 18/04/21 | 21:19:46 |
| 110 | kernel/syscalls/sigaction.c     | 1 pages | 54 lines  | 20/01/13 | 11:58:09 |
| 111 | kernel/syscalls/signal.c        | 1 pages | 56 lines  | 20/01/13 | 10:38:50 |
| 112 | kernel/syscalls/sigpending.c    | 1 pages | 30 lines  | 18/04/21 | 21:19:46 |
| 113 | kernel/syscalls/sigprocmask.c   | 1 pages | 51 lines  | 18/04/21 | 21:19:46 |
| 114 | kernel/syscalls/sigreturn.c     | 1 pages | 31 lines  | 20/12/07 | 17:08:39 |
| 115 | kernel/syscalls/sigsuspend.c    | 1 pages | 44 lines  | 18/04/21 | 21:19:46 |
| 116 | kernel/syscalls/socketcall.c    | 1 pages | 24 lines  | 18/10/20 | 11:31:46 |
| 117 | kernel/syscalls/ssetmask.c      | 1 pages | 26 lines  | 20/12/03 | 18:52:06 |
| 118 | kernel/syscalls/stat.c          | 1 pages | 52 lines  | 18/04/21 | 21:19:46 |
| 119 | kernel/syscalls/statfs.c        | 1 pages | 47 lines  | 18/04/21 | 21:19:46 |
| 120 | kernel/syscalls/stime.c         | 1 pages | 35 lines  | 18/04/21 | 21:19:46 |
| 121 | kernel/syscalls/symlink.c       | 2 pages | 73 lines  | 18/04/21 | 21:19:46 |
| 122 | kernel/syscalls/sync.c          | 1 pages | 27 lines  | 18/04/21 | 21:19:46 |
| 123 | kernel/syscalls/sysinfo.c       | 1 pages | 52 lines  | 18/04/21 | 21:19:46 |
| 124 | kernel/syscalls/time.c          | 1 pages | 37 lines  | 18/04/21 | 21:19:46 |
| 125 | kernel/syscalls/times.c         | 1 pages | 37 lines  | 18/04/21 | 21:19:46 |
| 126 | kernel/syscalls/truncate.c      | 1 pages | 66 lines  | 18/04/21 | 21:19:46 |
| 127 | kernel/syscalls/umask.c         | 1 pages | 27 lines  | 18/04/21 | 21:19:46 |
| 128 | kernel/syscalls/umount2.c       | 2 pages | 115 lines | 20/03/30 | 10:43:11 |
| 129 | kernel/syscalls/umount.c        | 1 pages | 22 lines  | 18/04/21 | 21:19:46 |
| 130 | kernel/syscalls/uname.c         | 1 pages | 34 lines  | 18/04/21 | 21:19:46 |
| 131 | kernel/syscalls/unlink.c        | 2 pages | 78 lines  | 18/04/21 | 21:19:46 |
| 132 | kernel/syscalls/ustat.c         | 1 pages | 44 lines  | 18/04/21 | 21:19:46 |
| 133 | kernel/syscalls/utime.c         | 2 pages | 72 lines  | 18/04/21 | 21:19:46 |
| 134 | kernel/syscalls/wait4.c         | 2 pages | 98 lines  | 20/01/13 | 11:50:52 |



## Table of Contents

Page -3/5

|     |                           |    |       |      |       |          |          |
|-----|---------------------------|----|-------|------|-------|----------|----------|
| 135 | kernel/syscalls/waitpid.c | 1  | pages | 23   | lines | 20/08/13 | 09:31:29 |
| 136 | kernel/syscalls/write.c   | 1  | pages | 49   | lines | 20/05/04 | 17:16:31 |
| 137 | drivers/block/dma.c       | 2  | pages | 93   | lines | 18/04/21 | 21:19:46 |
| 138 | drivers/block/floppy.c    | 14 | pages | 870  | lines | 20/04/09 | 19:13:48 |
| 139 | drivers/block/ide.c       | 14 | pages | 825  | lines | 20/04/09 | 19:08:13 |
| 140 | drivers/block/ide_cd.c    | 9  | pages | 524  | lines | 20/04/09 | 19:08:40 |
| 141 | drivers/block/ide_hd.c    | 9  | pages | 504  | lines | 20/09/14 | 22:50:28 |
| 142 | drivers/block/Makefile    | 1  | pages | 19   | lines | 18/04/21 | 21:19:46 |
| 143 | drivers/block/part.c      | 1  | pages | 34   | lines | 18/04/21 | 21:19:46 |
| 144 | drivers/block/ramdisk.c   | 3  | pages | 185  | lines | 20/03/29 | 10:36:52 |
| 145 | drivers/char/console.c    | 21 | pages | 1288 | lines | 20/11/29 | 19:43:32 |
| 146 | drivers/char/defkeymap.c  | 5  | pages | 149  | lines | 20/05/03 | 20:04:48 |
| 147 | drivers/char/keyboard.c   | 12 | pages | 742  | lines | 20/06/07 | 10:31:10 |
| 148 | drivers/char/lp.c         | 4  | pages | 216  | lines | 20/03/29 | 10:35:31 |
| 149 | drivers/char/Makefile     | 1  | pages | 19   | lines | 20/04/21 | 20:40:53 |
| 150 | drivers/char/memdev.c     | 8  | pages | 510  | lines | 20/03/29 | 20:08:05 |
| 151 | drivers/char/serial.c     | 9  | pages | 573  | lines | 20/09/14 | 13:14:20 |
| 152 | drivers/char/tty.c        | 16 | pages | 993  | lines | 20/05/03 | 22:22:23 |
| 153 | drivers/char/tty_queue.c  | 4  | pages | 264  | lines | 20/04/11 | 01:56:57 |
| 154 | drivers/char/vt.c         | 4  | pages | 218  | lines | 20/09/14 | 13:07:54 |
| 155 | fs/buffer.c               | 7  | pages | 443  | lines | 20/03/30 | 10:45:45 |
| 156 | fs/devices.c              | 6  | pages | 366  | lines | 20/09/26 | 17:42:35 |
| 157 | fs/elf.c                  | 10 | pages | 576  | lines | 19/12/25 | 20:02:22 |
| 158 | fs/ext2                   | 0  | pages | 0    | lines | 20/12/12 | 09:58:04 |
| 159 | fs/fd.c                   | 1  | pages | 50   | lines | 18/04/21 | 21:19:46 |
| 160 | fs/filesystems.c          | 2  | pages | 81   | lines | 18/04/21 | 21:19:46 |
| 161 | fs/inode.c                | 7  | pages | 459  | lines | 19/11/09 | 10:06:14 |
| 162 | fs/iso9660                | 0  | pages | 0    | lines | 20/12/12 | 09:58:04 |
| 163 | fs/locks.c                | 4  | pages | 208  | lines | 18/04/21 | 21:19:46 |
| 164 | fs/Makefile               | 1  | pages | 25   | lines | 19/11/23 | 11:39:20 |
| 165 | fs/minix                  | 0  | pages | 0    | lines | 20/12/12 | 09:58:04 |
| 166 | fs/namei.c                | 3  | pages | 178  | lines | 20/04/30 | 12:06:05 |
| 167 | fs/pipefs                 | 0  | pages | 0    | lines | 20/12/12 | 09:58:04 |
| 168 | fs/procfs                 | 0  | pages | 0    | lines | 20/12/12 | 09:58:04 |
| 169 | fs/script.c               | 2  | pages | 73   | lines | 19/12/24 | 18:45:20 |
| 170 | fs/super.c                | 4  | pages | 223  | lines | 18/04/21 | 21:19:46 |
| 171 | fs/ext2/bitmaps.c         | 5  | pages | 309  | lines | 19/11/09 | 10:14:21 |
| 172 | fs/ext2/dir.c             | 3  | pages | 147  | lines | 19/08/25 | 11:00:24 |
| 173 | fs/ext2/file.c            | 3  | pages | 133  | lines | 19/08/19 | 09:57:57 |
| 174 | fs/ext2/inode.c           | 8  | pages | 459  | lines | 19/08/24 | 11:41:18 |
| 175 | fs/ext2/Makefile          | 1  | pages | 19   | lines | 19/06/15 | 17:54:13 |
| 176 | fs/ext2/namei.c           | 12 | pages | 772  | lines | 19/11/09 | 10:06:20 |
| 177 | fs/ext2/super.c           | 4  | pages | 216  | lines | 19/08/25 | 11:09:23 |
| 178 | fs/ext2/symlink.c         | 3  | pages | 134  | lines | 19/06/15 | 18:13:30 |
| 179 | fs/iso9660/dir.c          | 3  | pages | 173  | lines | 18/05/30 | 16:45:16 |
| 180 | fs/iso9660/file.c         | 2  | pages | 78   | lines | 18/04/21 | 21:19:46 |
| 181 | fs/iso9660/inode.c        | 3  | pages | 184  | lines | 18/04/21 | 21:19:46 |
| 182 | fs/iso9660/Makefile       | 1  | pages | 19   | lines | 18/04/21 | 21:19:46 |
| 183 | fs/iso9660/namei.c        | 3  | pages | 122  | lines | 18/04/21 | 21:19:46 |
| 184 | fs/iso9660/rrrip.c        | 6  | pages | 340  | lines | 18/04/21 | 21:19:46 |
| 185 | fs/iso9660/super.c        | 4  | pages | 224  | lines | 18/04/21 | 21:19:46 |
| 186 | fs/iso9660/symlink.c      | 2  | pages | 109  | lines | 18/04/21 | 21:19:46 |
| 187 | fs/minix/bitmaps.c        | 3  | pages | 168  | lines | 19/08/16 | 17:17:09 |
| 188 | fs/minix/dir.c            | 3  | pages | 144  | lines | 18/05/30 | 17:22:07 |
| 189 | fs/minix/file.c           | 3  | pages | 133  | lines | 18/04/21 | 21:19:46 |
| 190 | fs/minix/inode.c          | 2  | pages | 75   | lines | 19/11/09 | 10:08:51 |
| 191 | fs/minix/Makefile         | 1  | pages | 19   | lines | 18/04/21 | 21:19:46 |
| 192 | fs/minix/namei.c          | 11 | pages | 700  | lines | 19/11/09 | 10:09:57 |
| 193 | fs/minix/super.c          | 5  | pages | 267  | lines | 18/04/21 | 21:19:46 |
| 194 | fs/minix/symlink.c        | 3  | pages | 136  | lines | 18/04/21 | 21:19:46 |
| 195 | fs/minix/v1_inode.c       | 7  | pages | 409  | lines | 19/11/09 | 10:10:19 |
| 196 | fs/minix/v2_inode.c       | 8  | pages | 473  | lines | 19/11/09 | 10:10:37 |
| 197 | fs/pipefs/fifo.c          | 2  | pages | 73   | lines | 18/04/21 | 21:19:46 |
| 198 | fs/pipefs/Makefile        | 1  | pages | 19   | lines | 18/04/21 | 21:19:46 |
| 199 | fs/pipefs/pipe.c          | 4  | pages | 202  | lines | 18/04/21 | 21:19:46 |
| 200 | fs/pipefs/super.c         | 2  | pages | 113  | lines | 19/11/09 | 10:10:58 |
| 201 | fs/procfs/data.c          | 14 | pages | 826  | lines | 20/06/21 | 18:13:34 |

**Table of Contents**

Page -4/5

|     |                             |         |           |          |          |
|-----|-----------------------------|---------|-----------|----------|----------|
| 202 | fs/procfs/dir.c             | 4 pages | 244 lines | 20/01/07 | 17:57:09 |
| 203 | fs/procfs/file.c            | 2 pages | 124 lines | 18/04/21 | 21:19:46 |
| 204 | fs/procfs/inode.c           | 2 pages | 89 lines  | 18/04/21 | 21:19:46 |
| 205 | fs/procfs/Makefile          | 1 pages | 19 lines  | 18/04/21 | 21:19:46 |
| 206 | fs/procfs/namei.c           | 2 pages | 89 lines  | 18/04/21 | 21:19:46 |
| 207 | fs/procfs/super.c           | 2 pages | 82 lines  | 18/04/21 | 21:19:46 |
| 208 | fs/procfs/symlink.c         | 3 pages | 136 lines | 18/04/21 | 21:19:46 |
| 209 | fs/procfs/tree.c            | 2 pages | 115 lines | 20/03/14 | 21:46:42 |
| 210 | mm/alloc.c                  | 1 pages | 35 lines  | 20/09/14 | 14:18:41 |
| 211 | mm/bios_map.c               | 2 pages | 90 lines  | 20/05/21 | 14:55:02 |
| 212 | mm/fault.c                  | 5 pages | 261 lines | 20/11/23 | 15:28:03 |
| 213 | mm/Makefile                 | 1 pages | 19 lines  | 18/04/21 | 21:19:46 |
| 214 | mm/memory.c                 | 8 pages | 443 lines | 20/12/07 | 14:47:37 |
| 215 | mm/mmap.c                   | 8 pages | 515 lines | 20/09/17 | 10:52:34 |
| 216 | mm/page.c                   | 7 pages | 443 lines | 20/03/30 | 10:43:20 |
| 217 | mm/swapper.c                | 1 pages | 64 lines  | 20/12/08 | 22:11:06 |
| 218 | lib/ctype.c                 | 3 pages | 140 lines | 20/01/26 | 21:03:16 |
| 219 | lib/Makefile                | 1 pages | 19 lines  | 18/04/21 | 21:19:46 |
| 220 | lib/printk.c                | 6 pages | 369 lines | 20/09/15 | 11:37:49 |
| 221 | lib/strings.c               | 5 pages | 280 lines | 18/04/21 | 21:19:46 |
| 222 | include/fiwix/asm.h         | 3 pages | 136 lines | 20/05/03 | 22:24:20 |
| 223 | include/fiwix/bios.h        | 1 pages | 29 lines  | 20/05/21 | 14:54:55 |
| 224 | include/fiwix/buffer.h      | 1 pages | 43 lines  | 18/04/21 | 21:19:46 |
| 225 | include/fiwix/cmos.h        | 1 pages | 58 lines  | 18/04/21 | 21:19:46 |
| 226 | include/fiwix/config.h      | 1 pages | 58 lines  | 20/12/08 | 21:37:48 |
| 227 | include/fiwix/console.h     | 3 pages | 159 lines | 20/11/28 | 20:09:58 |
| 228 | include/fiwix/const.h       | 1 pages | 20 lines  | 18/04/21 | 21:19:46 |
| 229 | include/fiwix/cpu.h         | 2 pages | 68 lines  | 18/04/21 | 21:19:46 |
| 230 | include/fiwix/ctype.h       | 1 pages | 38 lines  | 18/04/21 | 21:19:46 |
| 231 | include/fiwix/devices.h     | 1 pages | 49 lines  | 20/03/30 | 10:42:21 |
| 232 | include/fiwix/dirent.h      | 1 pages | 21 lines  | 18/04/21 | 21:19:46 |
| 233 | include/fiwix/dma.h         | 1 pages | 33 lines  | 18/04/21 | 21:19:46 |
| 234 | include/fiwix/errno.h       | 3 pages | 132 lines | 18/04/21 | 21:19:46 |
| 235 | include/fiwix/fcntl.h       | 1 pages | 67 lines  | 18/04/21 | 21:19:46 |
| 236 | include/fiwix/filesystems.h | 3 pages | 168 lines | 19/11/09 | 10:04:45 |
| 237 | include/fiwix/floppy.h      | 2 pages | 101 lines | 20/04/09 | 19:06:41 |
| 238 | include/fiwix/fs_ext2.h     | 4 pages | 251 lines | 19/08/18 | 18:47:16 |
| 239 | include/fiwix/fs.h          | 4 pages | 264 lines | 20/09/27 | 07:53:34 |
| 240 | include/fiwix/fs_iso9660.h  | 4 pages | 216 lines | 18/04/21 | 21:19:46 |
| 241 | include/fiwix/fs_minix.h    | 2 pages | 131 lines | 19/11/09 | 10:05:13 |
| 242 | include/fiwix/fs_pipe.h     | 1 pages | 23 lines  | 18/04/21 | 21:19:46 |
| 243 | include/fiwix/fs_proc.h     | 2 pages | 91 lines  | 20/03/14 | 21:46:54 |
| 244 | include/fiwix/i386elf.h     | 5 pages | 279 lines | 18/04/21 | 21:19:46 |
| 245 | include/fiwix/ide_cd.h      | 1 pages | 24 lines  | 18/04/21 | 21:19:46 |
| 246 | include/fiwix/ide.h         | 5 pages | 277 lines | 20/04/09 | 19:08:24 |
| 247 | include/fiwix/ide_hd.h      | 1 pages | 23 lines  | 18/04/21 | 21:19:46 |
| 248 | include/fiwix/ioctl.h       | 2 pages | 89 lines  | 18/04/21 | 21:19:46 |
| 249 | include/fiwix/kd.h          | 3 pages | 167 lines | 18/04/21 | 21:19:46 |
| 250 | include/fiwix/kernel.h      | 2 pages | 89 lines  | 20/03/07 | 18:58:23 |
| 251 | include/fiwix/keyboard.h    | 3 pages | 156 lines | 20/05/07 | 15:04:44 |
| 252 | include/fiwix/kparms.h      | 1 pages | 67 lines  | 20/05/21 | 13:20:00 |
| 253 | include/fiwix/limits.h      | 1 pages | 27 lines  | 20/05/21 | 13:24:50 |
| 254 | include/fiwix/locks.h       | 1 pages | 29 lines  | 18/04/21 | 21:19:46 |
| 255 | include/fiwix/lp.h          | 1 pages | 48 lines  | 18/04/21 | 21:19:46 |
| 256 | include/fiwix/memdev.h      | 1 pages | 57 lines  | 18/04/25 | 22:41:18 |
| 257 | include/fiwix/mman.h        | 1 pages | 59 lines  | 18/04/21 | 21:19:46 |
| 258 | include/fiwix/mm.h          | 2 pages | 106 lines | 20/09/24 | 16:26:57 |
| 259 | include/fiwix/multiboot.h   | 2 pages | 124 lines | 18/04/21 | 21:19:46 |
| 260 | include/fiwix/part.h        | 1 pages | 38 lines  | 18/04/21 | 21:19:46 |
| 261 | include/fiwix/pic.h         | 1 pages | 58 lines  | 20/06/07 | 10:10:26 |
| 262 | include/fiwix/pit.h         | 1 pages | 51 lines  | 18/04/21 | 21:19:46 |
| 263 | include/fiwix/process.h     | 3 pages | 183 lines | 19/11/26 | 15:54:25 |
| 264 | include/fiwix/ramdisk.h     | 1 pages | 33 lines  | 18/04/21 | 21:19:46 |
| 265 | include/fiwix/reboot.h      | 1 pages | 23 lines  | 18/04/21 | 21:19:46 |
| 266 | include/fiwix/resource.h    | 2 pages | 81 lines  | 18/04/21 | 21:19:46 |
| 267 | include/fiwix/sched.h       | 1 pages | 56 lines  | 20/03/09 | 16:52:23 |
| 268 | include/fiwix/segments.h    | 2 pages | 69 lines  | 18/04/21 | 21:19:46 |

**Table of Contents**

Page -5/5

|     |                            |         |           |          |          |
|-----|----------------------------|---------|-----------|----------|----------|
| 269 | include/fiwix/serial.h     | 2 pages | 116 lines | 20/04/19 | 16:47:39 |
| 270 | include/fiwix/sigcontext.h | 1 pages | 32 lines  | 18/04/21 | 21:19:46 |
| 271 | include/fiwix/signal.h     | 2 pages | 92 lines  | 18/04/21 | 21:19:46 |
| 272 | include/fiwix/sleep.h      | 1 pages | 40 lines  | 20/06/01 | 18:53:33 |
| 273 | include/fiwix/statbuf.h    | 1 pages | 48 lines  | 18/04/21 | 21:19:46 |
| 274 | include/fiwix/statfs.h     | 1 pages | 28 lines  | 18/04/21 | 21:19:46 |
| 275 | include/fiwix/stat.h       | 1 pages | 56 lines  | 18/04/21 | 21:19:46 |
| 276 | include/fiwix/stdarg.h     | 1 pages | 42 lines  | 18/04/21 | 21:19:46 |
| 277 | include/fiwix/stdio.h      | 1 pages | 15 lines  | 18/04/21 | 21:19:46 |
| 278 | include/fiwix/string.h     | 1 pages | 42 lines  | 18/04/21 | 21:19:46 |
| 279 | include/fiwix/syscalls.h   | 3 pages | 147 lines | 18/04/21 | 21:19:46 |
| 280 | include/fiwix/system.h     | 1 pages | 29 lines  | 20/12/12 | 09:53:35 |
| 281 | include/fiwix/termbits.h   | 3 pages | 176 lines | 20/04/13 | 11:47:52 |
| 282 | include/fiwix/termios.h    | 1 pages | 46 lines  | 20/04/13 | 11:47:49 |
| 283 | include/fiwix/timeb.h      | 1 pages | 18 lines  | 18/04/21 | 21:19:46 |
| 284 | include/fiwix/time.h       | 1 pages | 49 lines  | 18/04/21 | 21:19:46 |
| 285 | include/fiwix/timer.h      | 1 pages | 50 lines  | 20/04/09 | 19:13:22 |
| 286 | include/fiwix/times.h      | 1 pages | 18 lines  | 18/04/21 | 21:19:46 |
| 287 | include/fiwix/traps.h      | 1 pages | 46 lines  | 18/04/21 | 21:19:46 |
| 288 | include/fiwix/tty.h        | 2 pages | 99 lines  | 20/05/01 | 22:02:05 |
| 289 | include/fiwix/types.h      | 1 pages | 53 lines  | 18/04/21 | 21:19:46 |
| 290 | include/fiwix/unistd.h     | 4 pages | 207 lines | 18/04/21 | 21:19:46 |
| 291 | include/fiwix/ustat.h      | 1 pages | 20 lines  | 18/04/21 | 21:19:46 |
| 292 | include/fiwix/utime.h      | 1 pages | 18 lines  | 18/04/21 | 21:19:46 |
| 293 | include/fiwix/utsname.h    | 2 pages | 75 lines  | 18/04/21 | 21:19:46 |
| 294 | include/fiwix/version.h    | 1 pages | 1 lines   | 20/12/12 | 09:54:04 |
| 295 | include/fiwix/vt.h         | 1 pages | 54 lines  | 18/04/21 | 21:19:46 |
| 296 | README.md                  | 2 pages | 87 lines  | 20/12/12 | 08:49:29 |
| 297 | LICENSE                    | 1 pages | 22 lines  | 20/03/20 | 16:27:21 |
| 298 | CODING                     | 1 pages | 50 lines  | 20/03/23 | 10:47:38 |
| 299 | THANKS                     | 1 pages | 13 lines  | 20/03/26 | 18:47:04 |