

Mar 05, 15 12:33

handout08.txt

Page 1/4

```

1 CS 202, Spring 2015
2 Handout 8 (Class 10)
3
4 1. Example use of I/O instructions: boot loader
5
6     Below is a WeensyOS boot loader (for WeensyOS Schedos, which is part
7     of lab 4).
8
9     It may be helpful to understand the overall picture (given in the
10    comments; see also our class notes l03, Section 2(b).)
11
12    The main point for today's class is that this code demonstrates I/O,
13    specifically with the disk: the boot loader reads in the kernel from
14    the disk.
15
16    See the functions waitdisk() and readsect(). Compare to Figures 36.4
17    and 36.5 in OSTEP.
18
19    /* boot.c */
20
21    #include "x86.h"
22    #include "elf.h"
23
24    /*****
25     * This a dirt simple boot loader, whose sole job is to boot
26     * an ELF kernel image from the first IDE hard disk.
27     *
28     * DISK LAYOUT
29     * * This program (bootstart.S and boot.c) is the boot loader.
30     *   It should be stored in the disk's sector 0 (the first sector).
31     *
32     * * Sectors 1 through 32 hold the image for the schedos kernel and the
33     *   schedos application images (as binaries).
34     *
35     * * The kernel image must be in ELF executable format.
36     *
37     * BOOT UP STEPS
38     * * When the CPU boots it loads the BIOS into memory and executes it.
39     *
40     * * The BIOS intializes devices, sets up the interrupt routines, and
41     *   reads the first sector of the boot device (e.g., hard-drive)
42     *   into memory and jumps to it.
43     *
44     * * Assuming this boot loader is stored in the first sector of the
45     *   hard-drive, this code takes over.
46     *
47     * * Control starts in bootstart.S, which sets up protected mode,
48     *   and a stack so C code then run, then calls bootmain().
49     *
50     * * bootmain() in this file takes over, reads in the kernel image,
51     *   and jumps to it.
52     *
53     *****/
54
55    #define SECTORSIZE      512
56    #define PAGESIZE       4096
57    #define ELFHDR         ((struct Elf *) 0x10000) // scratch space
58
59    void readsect(void *addr, uint32_t sect);
60    void readseg(uint32_t va, uint32_t filesz, uint32_t memsz, uint32_t sect);
61
62    void
63    bootmain(void)
64    {
65        struct Proghdr *ph, *eph;
66        uint32_t *stackptr;
67
68        // read 1st page off disk
69        readseg((uint32_t) ELFHDR, PAGESIZE, PAGESIZE, 1);
70
71        // is this a valid ELF?
72        if (ELFHDR->e_magic != ELF_MAGIC)
73            return;

```

Mar 05, 15 12:33

handout08.txt

Page 2/4

```

74
75    // load each program segment (ignores ph flags)
76    ph = (struct Proghdr*) ((uint8_t *) ELFHDR + ELFHDR->e_phoff);
77    eph = ph + ELFHDR->e_phnum;
78    for (; ph < eph; ph++)
79        readseg(ph->p_va, ph->p_filesz, ph->p_memsz,
80                1 + ph->p_offset / SECTORSIZE);
81
82    // jump to the kernel, clearing %eax
83    __asm __volatile("movl %0, %%esp; ret" : : "r" (&ELFHDR->e_entry), "a" (0));
84    }
85
86    // Read 'filesz' bytes at 'offset' from kernel into virtual address 'va',
87    // then clear the memory from 'va+filesz' up to 'va+memsz' (set it to 0).
88    void
89    readseg(uint32_t va, uint32_t filesz, uint32_t memsz, uint32_t sect)
90    {
91        uint32_t end_va;
92
93        end_va = va + filesz;
94        memsz += va;
95
96        // round down to sector boundary
97        va &= ~(SECTORSIZE - 1);
98
99        // read sectors
100       while (va < end_va) {
101           readsect((uint8_t*) va, sect);
102           va += SECTORSIZE;
103           sect++;
104       }
105
106       // clear bss segment
107       while (end_va < memsz)
108           *((uint8_t*) end_va++) = 0;
109    }
110
111    void
112    waitdisk(void)
113    {
114        // wait for disk ready
115        while ((inb(0x1F7) & 0xC0) != 0x40)
116            /* do nothing */;
117    }
118
119    void
120    readsect(void *dst, uint32_t sect)
121    {
122        // wait for disk to be ready
123        waitdisk();
124
125        outb(0x1F2, 1);           // count = 1
126        outb(0x1F3, sect);
127        outb(0x1F4, sect >> 8);
128        outb(0x1F5, sect >> 16);
129        outb(0x1F6, (sect >> 24) | 0xE0);
130        outb(0x1F7, 0x20);       // cmd 0x20 - read sectors
131
132        // wait for disk to be ready
133        waitdisk();
134
135        // read a sector
136        insl(0x1F0, dst, SECTORSIZE/4);
137    }
138

```

Mar 05, 15 12:33

handout08.txt

Page 3/4

```

139 2. Two more examples of I/O instructions
140
141     (a) Reading keyboard input
142
143     The code below is an excerpt from WeensyOS's x86.c.
144
145     This reads a character typed at the keyboard (which shows up on the
146     "keyboard data port" (KBDATAP)), and converts it to a digit. This
147     code is not called in lab4; it was called in lab1. (Use grep to
148     convince yourself of this!)
149
150     /* Excerpt from WeensyOS x86.c. Comments from kbd.h in xv6 */
151
152     #define KBSTATP 0x64 // keyboard controller status port (I)
153     #define KBS_DIB 0x01 // keyboard data in buffer
154     #define KBDATAP 0x60 // keyboard data port (I)
155
156     int
157     console_read_digit(void)
158     {
159         uint8_t data;
160
161         if ((inb(KBSTATP) & KBS_DIB) == 0)
162             return -1;
163
164         data = inb(KBDATAP);
165         if (data >= 0x02 && data <= 0x0A)
166             return data - 0x02 + 1;
167         else if (data == 0x0B)
168             return 0;
169         else if (data >= 0x47 && data <= 0x49)
170             return data - 0x47 + 7;
171         else if (data >= 0x4B && data <= 0x4D)
172             return data - 0x4B + 4;
173         else if (data >= 0x4F && data <= 0x51)
174             return data - 0x4F + 1;
175         else if (data == 0x53)
176             return 0;
177         else
178             return -1;
179     }
180
181     (b) Setting the cursor position
182
183     The code below is also excerpted from WeensyOS's x86.c.
184     It clears the console (see next items on the handout) and then uses
185     I/O instructions to set a blinking cursor in the upper left of the
186     screen.
187
188     /*
189     * console_clear
190     *
191     * Clear the console by writing spaces to it, and move the cursor
192     * to the upper left (row 0, column 0).
193     */
194
195     void
196     console_clear(void)
197     {
198         int i;
199
200         /* what's this?? (see next items on handout) */
201         cursorpos = (uint16_t *) 0xB8000;
202
203         for (i = 0; i < 80 * 25; i++)
204             cursorpos[i] = ' ' | 0x0700;
205
206         outb(0x3D4, 14); // Command 14 = upper byte of position
207         outb(0x3D5, 0 / 256); // row 0
208         outb(0x3D4, 15); // Command 15 = lower byte of position
209         outb(0x3D5, 0 % 256); // column 0
210     }
211

```

Thursday March 05, 2015

handout08.txt

Mar 05, 15 12:33

handout08.txt

Page 4/4

```

212 3. Memory-mapped I/O
213
214     a. Here is a 32-bit PC's physical memory map:
215
216     +-----+ <- 0xFFFFFFFF (4GB)
217     | 32-bit |
218     | memory |
219     | mapped |
220     | devices|
221     | \/\|  |
222     | \/\|  |
223     | \/\|  |
224     | \/\|  |
225     | \/\|  |
226     | \/\|  |
227     +-----+ <- depends on amount of RAM
228     | Extended Memory |
229     | \/\|  |
230     | \/\|  |
231     | \/\|  |
232     | \/\|  |
233     +-----+ <- 0x00100000 (1MB)
234     | BIOS ROM |
235     | \/\|  | <- 0x000F0000 (960KB)
236     | \/\|  |
237     | \/\|  |
238     +-----+ <- 0x000C0000 (768KB)
239     | VGA Display |
240     | \/\|  | <- 0x000A0000 (640KB)
241     | \/\|  |
242     | \/\|  |
243     | \/\|  |
244     +-----+ <- 0x00000000
245     | Low Memory |
246
247     [Credit to Frans Kaashoek, Robert Morris, and Nickolai Zeldovich for
248     this picture]
249
250     b. Loads and stores to the device memory "go to hardware".
251
252     An example is in the console printing code from WeensyOS. Here are
253     excerpts from lib.h and lib.c:
254
255     /* Compare the addresses below to the map above. */
256     #define CONSOLE_BEGIN ((uint16_t *) 0x00B8000)
257     #define CONSOLE_END (CONSOLE_BEGIN + 80 * 25)
258
259     /*
260     * prints a character to the console at the specified
261     * cursor position in the specified color.
262     * Question: what is going on in the check
263     * if (c == '\n')
264     * ?
265     * Hint: '\n' is "C" for "newline" (the user pressed enter).
266     */
267
268     static uint16_t *
269     console_putc(uint16_t *cursor, unsigned char c, int color)
270     {
271         if (cursor >= CONSOLE_END)
272             cursor = CONSOLE_BEGIN;
273         if (c == '\n') {
274             int pos = (cursor - CONSOLE_BEGIN) % 80;
275             /* what does this do? */
276             for (; pos != 80; pos++)
277                 *cursor++ = ' ' | color;
278         } else
279             *cursor++ = c | color;
280         return cursor;
281     }
282
283

```

2/2