

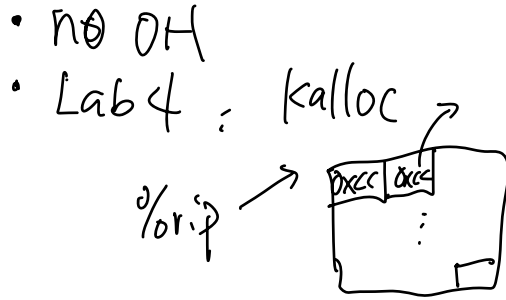
Week 11.b  
 CS 5600  
 03/22 2023

1. Last time
2. mmap
3. I/O architecture

## 1. Last time

- Page fault

• "Paging" (Overcommitting Mem)



INT n/INTO/INT 3—Call to Interrupt Procedure

Opcode	Instruction	Opr En	64-Bit Mode	Compat/ Leg Mode	Description
CC	INT 3	NP	Valid	Valid	Interrupt 3—trap to debugger.
CD/ib	INT imm8	I	Valid	Valid	Interrupt vector specified by immediate byte.
CE	INTO	NP	Invalid	Valid	Interrupt 4—if overflow flag is 1.

Cost:

memory access time ~ 100ns  
 SSD access time ~ 1 ms = 10<sup>6</sup> ns

0.001%, 0.1%, 0.01%

QUESTION: what does page fault probability (p) need to be to ensure that paging hurts performance by less than 10%?

$$\Rightarrow P \cdot t_d + (1-P) \cdot t_m \leq 1.1 \cdot t_m$$

$$\Downarrow P \cdot t_d + t_m - P \cdot t_m \leq 0.1 t_m$$

thrashing

$$P(t_d - t_m) \leq 0.1 t_m$$

$$P \leq \frac{0.1 t_m}{t_d - t_m} \approx \frac{0.1 \cdot 100 \text{ ns}}{10^6 \text{ ns}}$$

$$= \frac{1}{10^5}$$

$$\approx 0.001\%$$



VM:

50 Page

PM:

40 Page

$$20\% \cdot t_d + 80\% \cdot t_m = 0.2 \cdot 10^6 \text{ ns}$$

...

# Mmap

$f_m = 100ns = 2000 \times$   
SLOWER

Some syscalls:

```
fd = open(pathname, mode)
write(fd, buf, sz)
read(fd, buf, sz)
```

A new syscall:

```
void* mmap(void* addr, size_t len, int prot, int flags,
            int fd, off_t offset);
```

• Why mmap is faster?

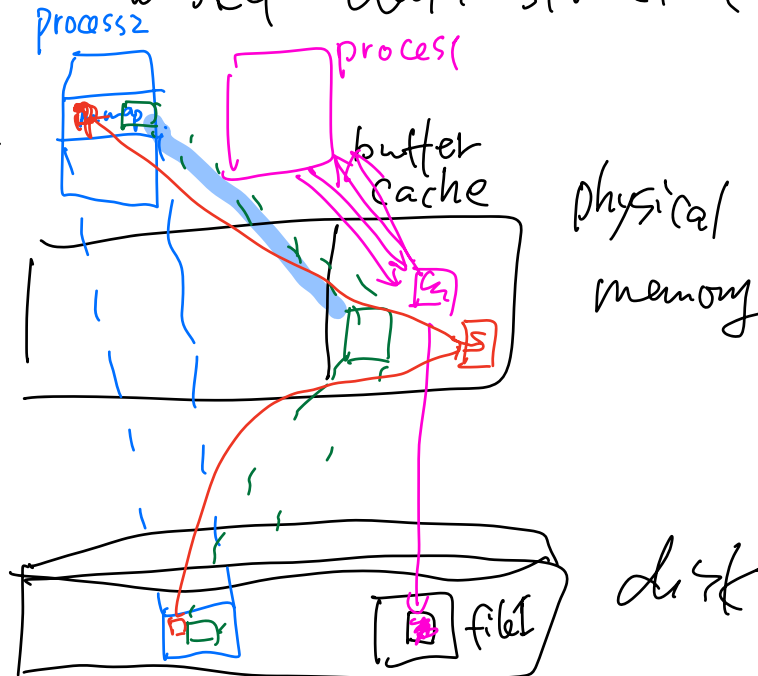
• Other use cases?

• large file

• shared data structure (MAP\_SHARED)

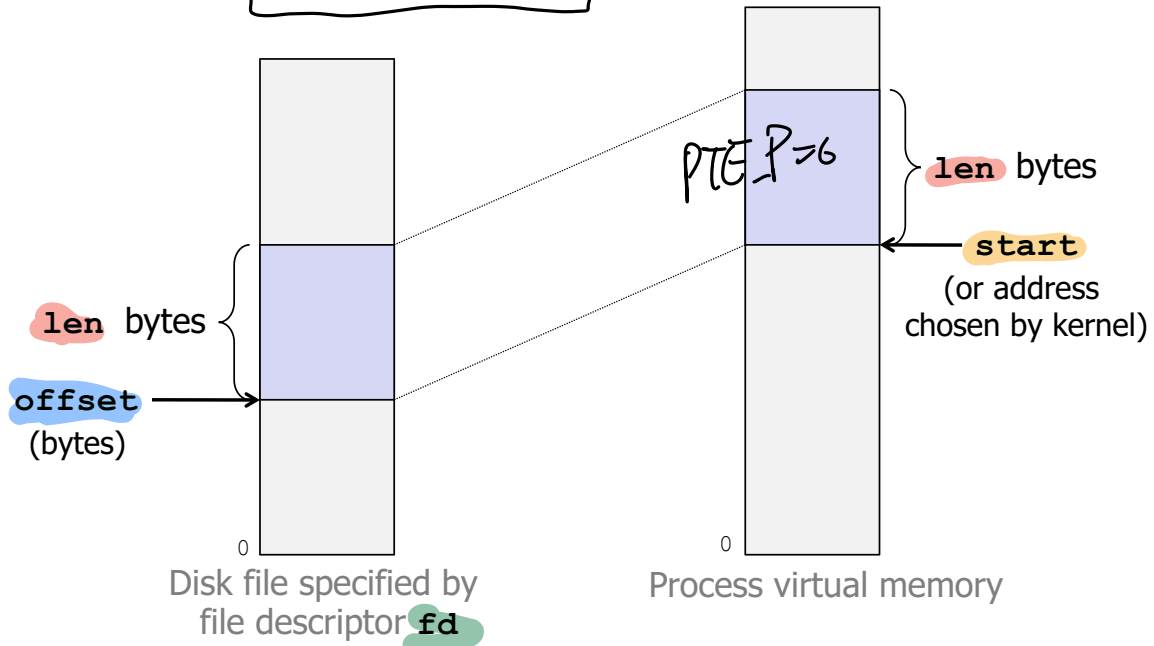
• file-based data structure

• impl?



# User-Level Memory Mapping

```
void *mmap(void *start, int len,  
           int prot, int flags, int fd, int offset)
```



```

1 CS5600, Handout week11.a
2
3 /* file: mmap.c */
4
5 #include <fcntl.h>
6 #include <stdio.h>
7 #include <stdlib.h>
8 #include <sys/mman.h>
9 #include <sys/stat.h>
10 #include <sys/types.h>
11 #include <unistd.h>
12
13 void mmapwrite(int fd, int size);
14 void normalwrite(int fd, int size);
15
16 int main(int argc, char **argv) {
17     struct stat stat;
18     int fd;
19
20     if (argc != 2) { // Check for required cmd line arg
21         printf("usage: %s <filename>\n", argv[0]);
22         exit(0);
23     }
24
25     /* Copy input file to stdout */
26     if ((fd = open(argv[1], O_RDONLY, 0)) < 0)
27         perror("open");
28
29     fstat(fd, &stat);
30
31     // option 1
32     mmapwrite(fd, stat.st_size);
33
34     /* // option 2
35      * normalwrite(fd, stat.st_size);
36      */
37
38     close(fd);
39
40     return 0;
41 }
42
43 void mmapwrite(int fd, int size) {
44     /* Ptr to memory mapped area */
45     char *bufp;
46
47     bufp = mmap(NULL, size, PROT_READ, MAP_PRIVATE, fd, 0);
48
49     write(STDOUT_FILENO, bufp, size);
50     return;
51 }
52
53 }
54
55 void normalwrite(int fd, int size) {
56     char *buf = malloc(size);
57
58     read(fd, buf, size);
59
60     write(STDOUT_FILENO, buf, size);
61
62     return;
63 }
64
65 }

```

*entire file* (with arrow pointing to `size` in `mmap`)

*offset* (with arrow pointing to `0` in `mmap`)

*don't care where it maps to* (with arrow pointing to `write`)

Question:  
Which runs faster, option 1 or option 2? by how much?

Exercise:  
Try to run both options by yourself:

```

$ cat /dev/urandom | head -c 1000000000 > 1G.file
$ make mmap
$ time ./mmap 1G.file > /dev/null

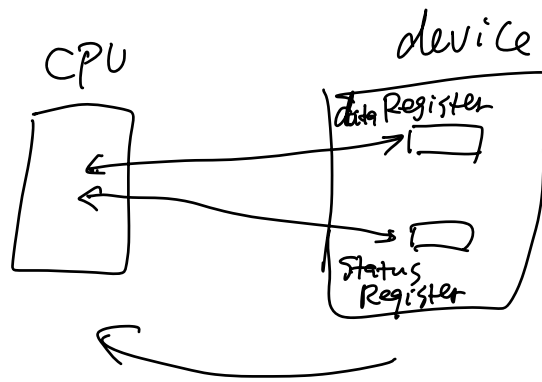
```

```

$ vim mmap.c
// switch to option 2
$ make mmap
$ time ./mmap 1G.file > /dev/null

```

I/O



PS/2 (old keyboard/mouse)

IO Port	Access Type	Purpose
0x60	Read/Write	Data Port
0x64	Read	Status Register
0x64	Write	Command Register

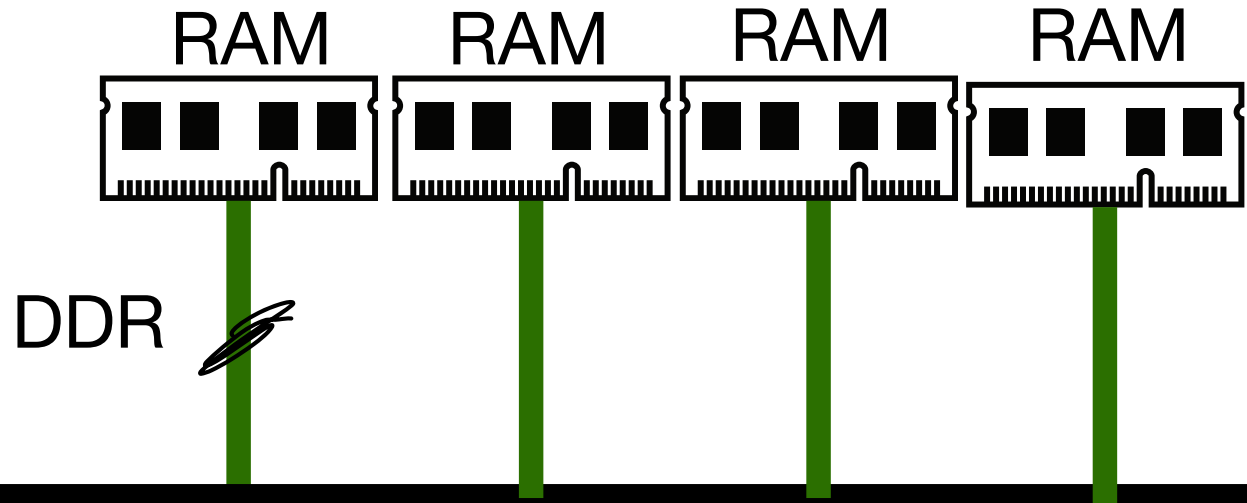
inb / outb / inw / outw.







mdq

# Machine


CXL



Processor



SATA



1 CS5600 Week11.b

2  
3 1. Two examples of I/O instructions

4  
5 (a) Reading keyboard input ←

6  
7 The code below is an excerpt from WeensyOS.  
8 (details in PS/2 controller: [https://wiki.osdev.org/%228042%22\\_PS/2\\_Controller](https://wiki.osdev.org/%228042%22_PS/2_Controller))  
9 This reads a character typed at the keyboard (which shows up on the  
10 "keyboard data port" (KEYBOARD\_DATAREG)).

```

11 /* Excerpt from WeensyOS x86-64.h and k-hardware.cc */
12 // Keyboard programmed I/O
13 #define KEYBOARD_STATUSREG 0x64 ←
14 #define KEYBOARD_STATUS_READY 0x01 ←
15 #define KEYBOARD_DATAREG 0x60 ←
16
17 int keyboard_read() {
18     static uint8_t modifiers;
19     static uint8_t last_escape;
20
21     if ((inb(KEYBOARD_STATUSREG) & KEYBOARD_STATUS_READY) == 0) {
22         return -1;
23     }
24
25     uint8_t data = inb(KEYBOARD_DATAREG); ← 0x1E
26     uint8_t escape = last_escape;
27     last_escape = 0;
28
29     if (data == 0xE0) { // mode shift
30         last_escape = 0x80;
31         return 0;
32     } else if (data & 0x80) { // key release: matters only
33         // for modifier keys
34         int ch = keymap[(data & 0x7F) | escape];
35         if (ch >= KEY_SHIFT && ch < KEY_CAPSLOCK) {
36             modifiers &= ~(1 << (ch - KEY_SHIFT));
37         }
38         return 0;
39     }
40     int ch = (unsigned char) keymap[data | escape];
41     if (ch >= 'a' && ch <= 'z') {
42         if (modifiers & MOD_CONTROL) {
43             ch -= 0x60;
44         } else if (!(modifiers & MOD_SHIFT) != \
45             !(modifiers & MOD_CAPSLOCK)) { } ← } ←
46         ch -= 0x20; ← 'a' ⇒ 'A'
47     } else if (ch >= KEY_CAPSLOCK) {
48         modifiers ^= 1 << (ch - KEY_SHIFT);
49         ch = 0;
50     } else if (ch >= KEY_SHIFT) {
51         modifiers |= 1 << (ch - KEY_SHIFT);
52         ch = 0;
53     } else if (ch >= CKEY(0) && ch <= CKEY(21)) {
54         ch = complex_keymap[ch - CKEY(0)].map[modifiers & 3];
55     } else if (ch < 0x80 && (modifiers & MOD_CONTROL)) {
56         ch = 0;
57     }
58     return ch; 'a'
59 }
60 }
61 }
62 }
63 }
64 }
65 }

```

A pressed : 0x1E  
 A released : 0x9E  
 0x80  
 ↓  
 1000 000

66 (b) Setting the cursor position

67  
68 The code below is also excerpted from WeensyOS. It uses I/O  
69 instructions to set a blinking cursor. To set the cursor to  
70 the upper left of the screen, run: console\_show\_cursor(0)

```

71 // console_show_cursor(cpos)
72 // Move the console cursor to position 'cpo',
73 // which should be between 0 and 80 * 25.
74
75 void console_show_cursor(int cpos) {
76     if (cpo < 0 || cpo > CONSOLE_ROWS * CONSOLE_COLUMNS)
77         cpo = 0;
78     outb(0x3D4, 14); // Command 14 = upper byte of position
79     outb(0x3D5, cpo / 256); // upper byte (256 = 2^8)
80     outb(0x3D4, 15); // Command 15 = lower byte of position
81     outb(0x3D5, cpo % 256); // lower byte
82 }
83 // if interested, see details: https://wiki.osdev.org/Text\_Mode\_Cursor
84 }
85 }
86 }
87 }
88 }

```

