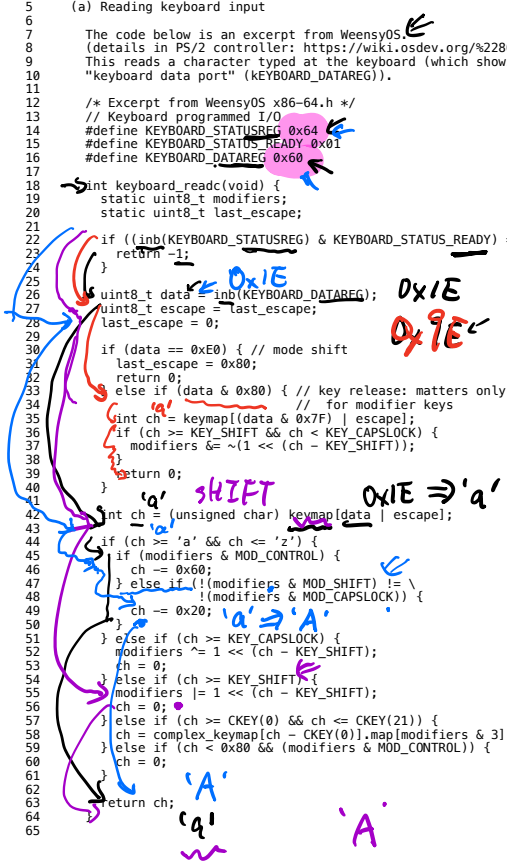


```

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/%28%20PS/2_Controller)
9 This reads a character typed at the keyboard (which shows up on the
10 "keyboard data port" (KEYBOARD_DATAREG)).
11
12 /* Excerpt from WeensyOS x86-64.h */
13 // Keyboard programmed I/O
14 #define KEYBOARD_STATUSREG 0x64
15 #define KEYBOARD_STATUS_READY 0x01
16 #define KEYBOARD_DATAREG 0x60
17
18 int keyboard_readc(void) {
19     static uint8_t modifiers;
20     static uint8_t last_escape;
21
22     if ((inb(KEYBOARD_STATUSREG) & KEYBOARD_STATUS_READY) == 0) {
23         return -1;
24     }
25     uint8_t data = inb(KEYBOARD_DATAREG);
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         uint 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
41     int ch = (unsigned char) keymap[data | escape];
42
43     if (ch >= 'a' && ch <= 'z') {
44         if (modifiers & MOD_CONTROL) {
45             ch -= 0x60;
46         } else if (!(modifiers & MOD_SHIFT) != '\
47             !(modifiers & MOD_CAPSLOCK)) {
48             ch -= 0x20;
49         }
50     }
51     else if (ch >= KEY_CAPSLOCK) {
52         modifiers ^= 1 << (ch - KEY_SHIFT);
53         ch = 0;
54     } else if (ch >= KEY_SHIFT) {
55         modifiers |= 1 << (ch - KEY_SHIFT);
56         ch = 0;
57     } else if (ch >= CKEY(0) && ch <= CKEY(21)) {
58         ch = complex_keymap[ch - CKEY(0)].map[modifiers & 3];
59     } else if (ch < 0x80 && (modifiers & MOD_CONTROL)) {
60         ch = 0;
61     }
62     return ch;
63 }
64
65

```




Press A.
release A
Shift + Press A

```

66
67 (b) Setting the cursor position
68
69 The code below is also excerpted from WeensyOS. It uses I/O
70 instructions to set a blinking cursor. To set the cursor to
71 the upper left of the screen, run: console_show_cursor(0)
72
73 // console_show_cursor(cpos)
74 // Move the console cursor to position 'cpos',
75 // which should be between 0 and 80 * 25.
76
77 void console_show_cursor(int cpos) {
78     if (cpos < 0 || cpos > CONSOLE_ROWS * CONSOLE_COLUMNS)
79         cpos = 0;
80
81     // Command 14 = upper byte of position
82     outb(0x3D4, 14); // Command 14 = upper byte of position
83     // Command 15 = lower byte of position
84     outb(0x3D4, 15); // Command 15 = lower byte of position
85     // Command 15 = lower byte of position
86     // lower byte
87     // = 1
88
89 // if interested, see details: https://wiki.osdev.org/Text_Mode_Cursor

```



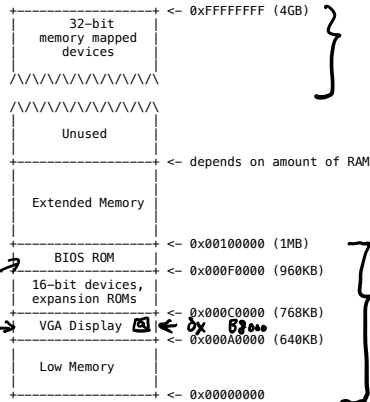
$80 \times 25 = 2000 > 256 = 2^8$

0x1 0x1
1B 1B

89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126

2. Memory-mapped I/O

(a) Here is a 32-bit PC's physical memory map:



} Flag, APIC...

}
0x83...

[Credit to Frans Kaashoek, Robert Morris, and Nikolai Zeldovich for this picture]

127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160

(b) Loads and stores to the device memory "go to hardware".

Here is an excerpt of the console printing code from WeensyOS.

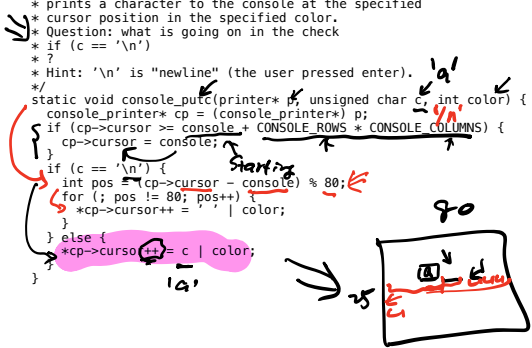
```

/* Compare the address below to the map in panel 2(a). */
PROVIDE(console = 0xB8000);

This is an excerpt about printing; notice how it uses the address
"console":

/*
 * prints a character to the console at the specified
 * cursor position in the specified color.
 * Question: what is going on in the check
 * if (c == '\n')
 * ?
 * Hint: '\n' is "newline" (the user pressed enter).
 */
static void console_putc(printer* p, unsigned char c, int color) {
  console_printer* cp = (console_printer*) p;
  if (cp->cursor >= console + CONSOLE_ROWS * CONSOLE_COLUMNS) {
    cp->cursor = console;
  }
  if (c == '\n') {
    int pos = (cp->cursor - console) % 80;
    for (; pos != 80; pos++) {
      *cp->cursor++ = ' '; color;
    }
  } else {
    *cp->cursor++ = c | color;
  }
}


```



inb ~~xxx~~
out ~~xxx~~

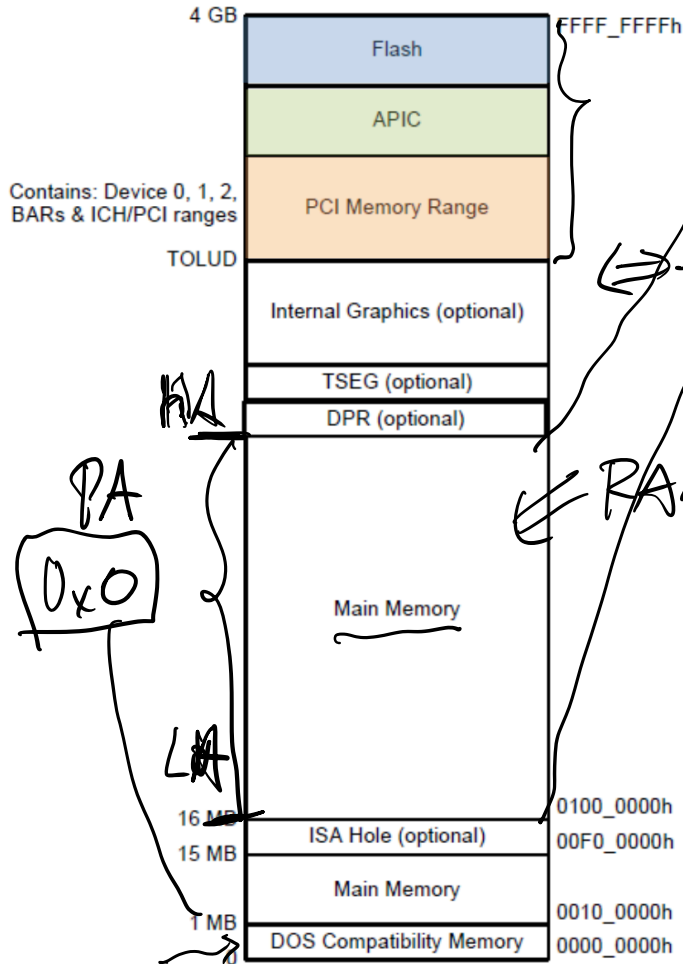
Port I/O Address Space

- Software and hardware architectures of x86 architecture support a separate address space called "I/O Address Space"
 - Separate from memory space ↙
- Access to this separate I/O space is handled through a set of I/O instructions
 - IN, OUT, INS, OUTS
- Access requires Ring0 privileges
 - Access requirement does not apply to all operating modes (like Real-Mode)
- The processor allows 64 KB+3 bytes to be addressed within the I/O space
- Harkens back to a time when memory was not so plentiful
- You may never see port I/O when analyzing high-level applications, but in systems programming (and especially BIOS) you will see lots of port I/O
- One of the biggest impediments to understanding what's going on in a BIOS



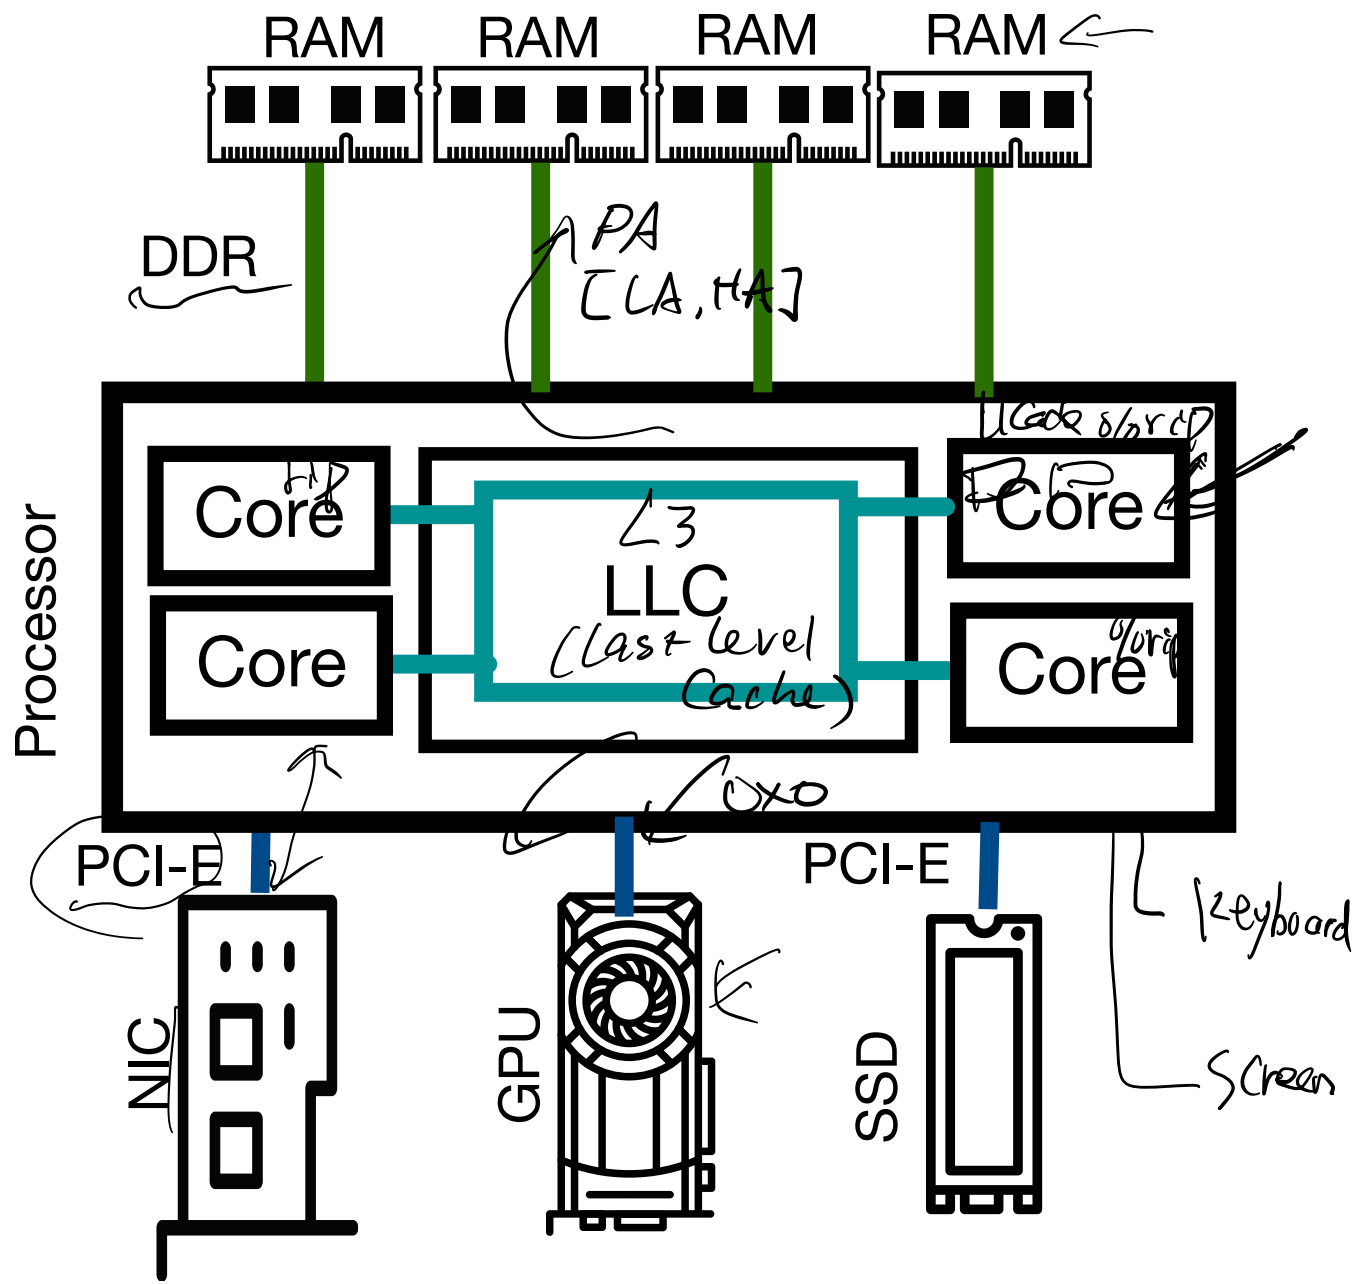
Memory Mapped IO

32bit



- The colored regions are memory mapped devices
- Accesses to these memory ranges are decoded to a device itself
- Flash refers to the BIOS flash
- APIC is the Advanced Programmable Interrupt Controller
- PCI Memory range is programmed by BIOS in the PCIEXBAR

Machine




```

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
45     /* Ptr to memory mapped area */
46     char *bufp;
47
48     bufp = mmap(NULL, size, PROT_READ, MAP_PRIVATE, fd, 0);
49     write(STDOUT_FILENO, bufp, size);
50
51     return;
52 }
53
54
55
56 void normalwrite(int fd, int size) {
57     char *buf = malloc(size);
58     read(fd, buf, size);
59
60     write(STDOUT_FILENO, buf, size);
61
62     return;
63 }
64
65 }

```

→ Cat "/tmp/file.txt"

→ size of file



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

```

BSX

PTE



⇒ buffer cache.

32 bit \Rightarrow 36 bit \Rightarrow 48 bit \Rightarrow 57 bit

$2^4 = 16$

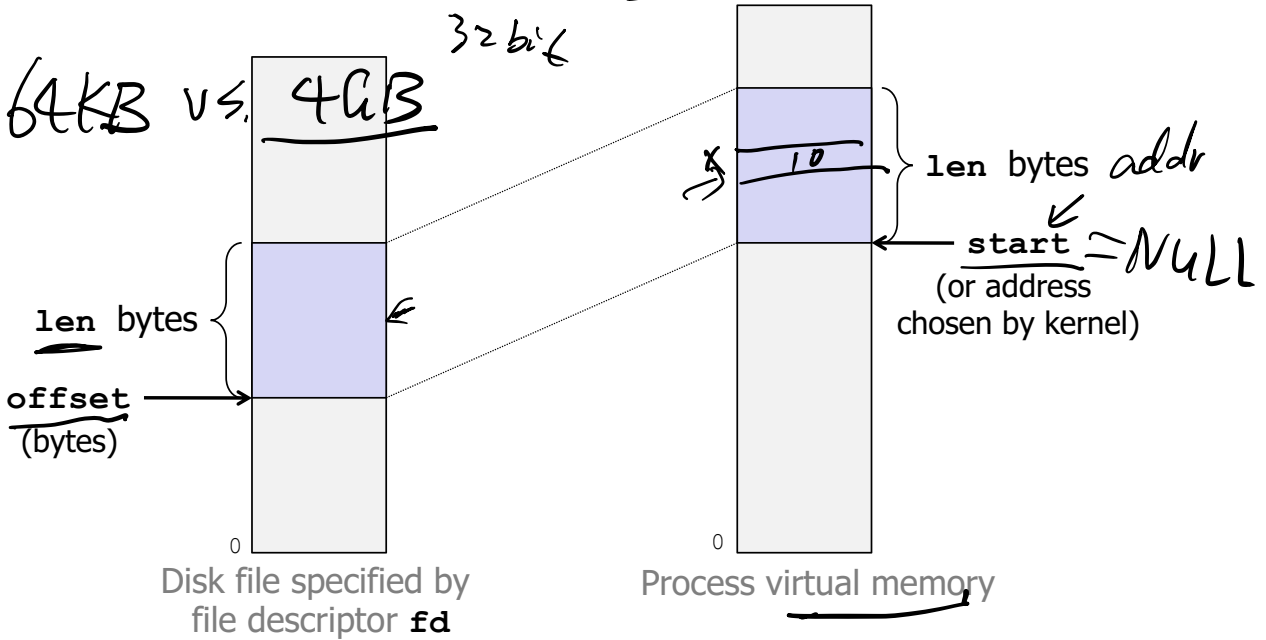
48 bit

$2^{48} = 256 \text{ TB}$

User-Level Memory Mapping

```

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



```

read(fd, buf, size)
write(fd, buf, size)
    
```

$\int a = \text{addr}[x]$

$\text{addr}[x] = 10;$

1. Last time ←
2. mmap
3. I/O architecture
4. device drivers

VM
I/O

• Lab 3.

• review session.

• Lab 2.

• page faults.

① PTE's P=0 ←


② permission check

• mmap → syscall

recall:


(3) fd = open("/tmp/file.txt", mode) ←

↳ ? o/i int. { 0, 1, 2 }



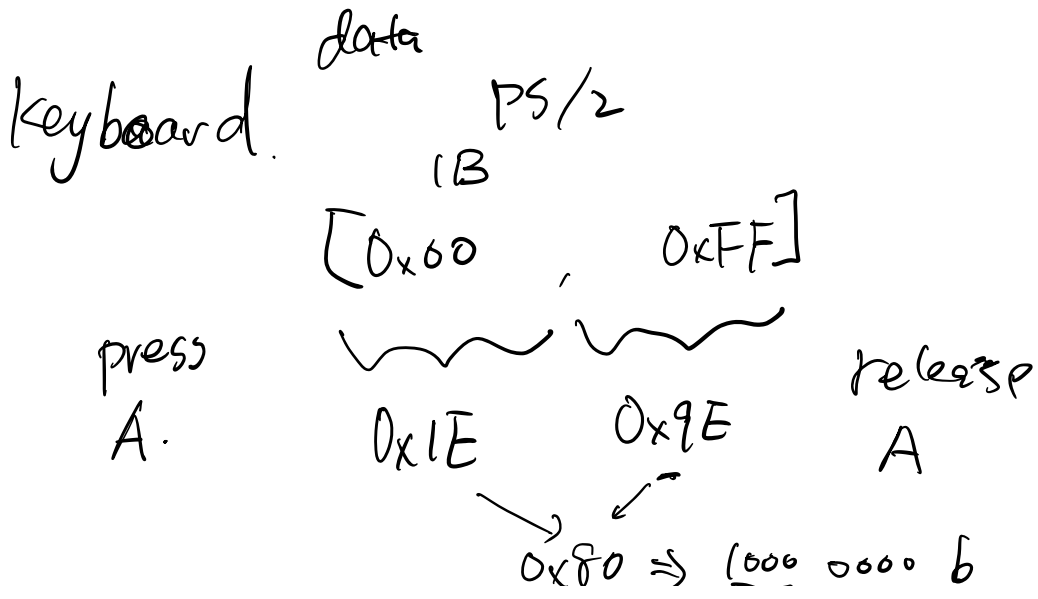
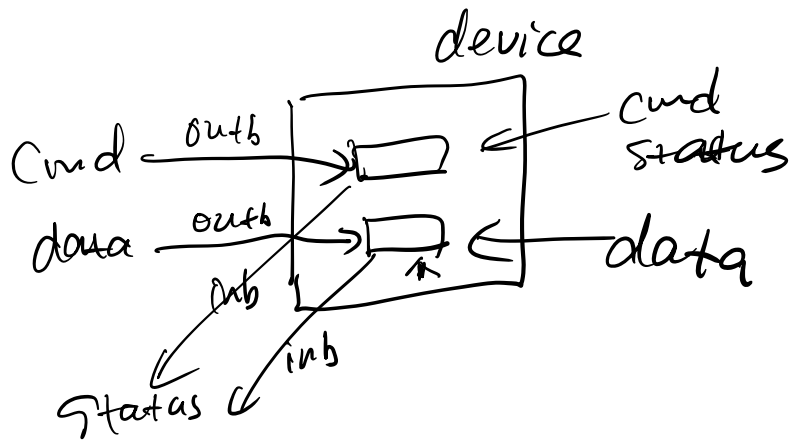
{ read (fd, ...) }
| write (fd, ...)

- file-based data structure.



① Port I/O. \leftarrow

- inb/outb/inw/intw



-
- ① memory-mapped I/O
 - ② interrupt
 - ④ via memory (DMA)