

# How to read a code repository?

Reading egos-2000 as an example

# Read a repository: 3 passes

## ➔ 1st pass

- read documents and filenames

## • 2nd pass

- track the execution: earth → grass → applications

## • 3rd pass

- read details of specific functionality, such as system call

# Documents of egos-2000

- Explain **why** the project is **important**
  - [README.md](#)
- Explain **how to use** this project
  - [references/USAGES.md](#)
- Explain **the internal design** of the project
  - [references/README.md](#)

Software companies use tools like



# Read filenames: earth

- sd
- ~~bus\_gpio.c~~
- ~~bus\_uart.c~~
- cpu\_intr.c
- cpu\_mmu.c
- dev\_disk.c
- dev\_page.c
- dev\_tty.c
- earth.S
- earth.c
- earth.lids

- `gpio` and `uart` are buses connecting the CPU with I/O devices, just like `usb`
- in the first pass of reading code, knowing what they are on the high-level is enough

# Read filenames: **earth**

sd
<del>bus_gpio.c</del>
<del>bus_uart.c</del>
cpu_intr.c
cpu_mmu.c
<del>dev_disk.c</del>
<del>dev_page.c</del>
<del>dev_tty.c</del>
earth.S
earth.c
earth.lids

- **dev\_disk** controls ROM and SD card
- **dev\_tty** reads keyboard input and print output to the screen
- **dev\_page** does paging from memory to the first 1MB of the SD card

# Read filenames: **earth**

sd
<del>bus_gpio.c</del>
<del>bus_uart.c</del>
<del>cpu_intr.c</del>
<del>cpu_mmu.c</del>
<del>dev_disk.c</del>
<del>dev_page.c</del>
<del>dev_tty.c</del>
earth.S
earth.c
earth.lids

- **cpu\_intr**: interrupt/exception handling
- **cpu\_mmu**: memory management unit (MMU)

# Read filenames: **earth**

sd
<del>bus_gpio.c</del>
<del>bus_uart.c</del>
<del>cpu_intr.c</del>
<del>cpu_mmu.c</del>
<del>dev_disk.c</del>
<del>dev_page.c</del>
<del>dev_tty.c</del>
<del>earth.S</del>
<del>earth.c</del>
<del>earth.lds</del>

- `earth.S` and `earth.c` are for earth layer initialization (e.g., the `main()` of earth)
- `earth.lds` specifies the memory layout

# Read filenames: **earth/sd**

 sd.h

---

 sd\_init.c

---

 sd\_rw.c

---

 sd\_utils.c

- `sd.h` provides basic definitions
- `sd_init.c` initializes the SD card
- `sd_rw.c` provides SD card read and write
- `sd_utils.c` provides helper functions
- There is an optional project P4 on SD driver



# Read filenames: **grass**

grass.S

grass.c

grass.lids

process.c

process.h

scheduler.c

syscall.c

syscall.h

timer.c

- `grass.S` and `grass.c`: initialization
- `grass.lids`: memory layout

## Grass layer (hardware independent)

- `grass/timer`: timer control registers
- `grass/syscall`: system call interfaces to user applications
- `grass/process`: data structures for managing processes (touched by P1)
- `grass/scheduler`: preemptive scheduling and inter-process communication

# Read a repository: 3 passes

- 1st pass

- read documents and filenames

- ➔ 2nd pass

- track the execution: earth → grass → applications

- 3rd pass

- read details of specific functionality, such as system call

The Key:

Find `main()` functions

and track executions from there

# grep is a useful command

```
> cd egos-2000  
> grep "main(" -r *
```

# Main functions in the repository

```
> cd egos-2000
> grep "main(" -r *
earth/earth.S:      /* Call main() of earth.c */
earth/earth.c:int main() {
grass/grass.S:      /* Call main() of grass.c */
grass/grass.c:int main() {
tools/mkrom.c:int main() {
tools/mkfs.c:int main() {

apps/*.c: /* Every application has a main() function */
```

# Main function in **earth**

- Read **earth.s** and **earth.c**
  - Boot loader **disable interrupt** and call `earth main()`
  - `Earth main()` essentially
    - **initialize dev\_tty, dev\_disk, cpu\_intr, cpu\_mmu**
    - load and enter the grass layer

# Main function in **grass**

- Read **grass.s** and **grass.c**
  - Initialize data structures for **processes** (like P1)
  - Initialize and enable **timer interrupt**
  - Load and enter the first application: **GPID\_PROCESS**
- Where is **GPID\_PROCESS** defined?

# Find GPID\_PROCESS

```
> cd egos-2000
# Find which header file contains GPID_PROCESS
> grep "GPID_PROCESS" -r * | grep "\.h"

library/servers/servers.h:      GPID_PROCESS,
library/servers/servers.h: /* GPID_PROCESS */
```



# Kernel servers (aka. Daemon)

```
enum grass_servers {  
    GPID_UNUSED,  
    GPID_PROCESS,  
    GPID_FILE,  
    GPID_DIR,  
    GPID_SHELL,  
    GPID_USER_START  
};
```

- **GPID\_PROCESS**
  - spawn and kill processes
- **GPID\_FILE & GPID\_DIR**
  - something about file system
- **GPID\_SHELL**
  - shell for entering user commands

# Control flow sketch

- During boot up

- earth main() → grass main() → GPID\_PROCESS

- GPID\_PROCESS → GPID\_FILE

- GPID\_PROCESS → GPID\_DIR

- GPID\_PROCESS → GPID\_SHELL

- After boot up

- GPID\_SHELL → GPID\_PROCESS → user applications

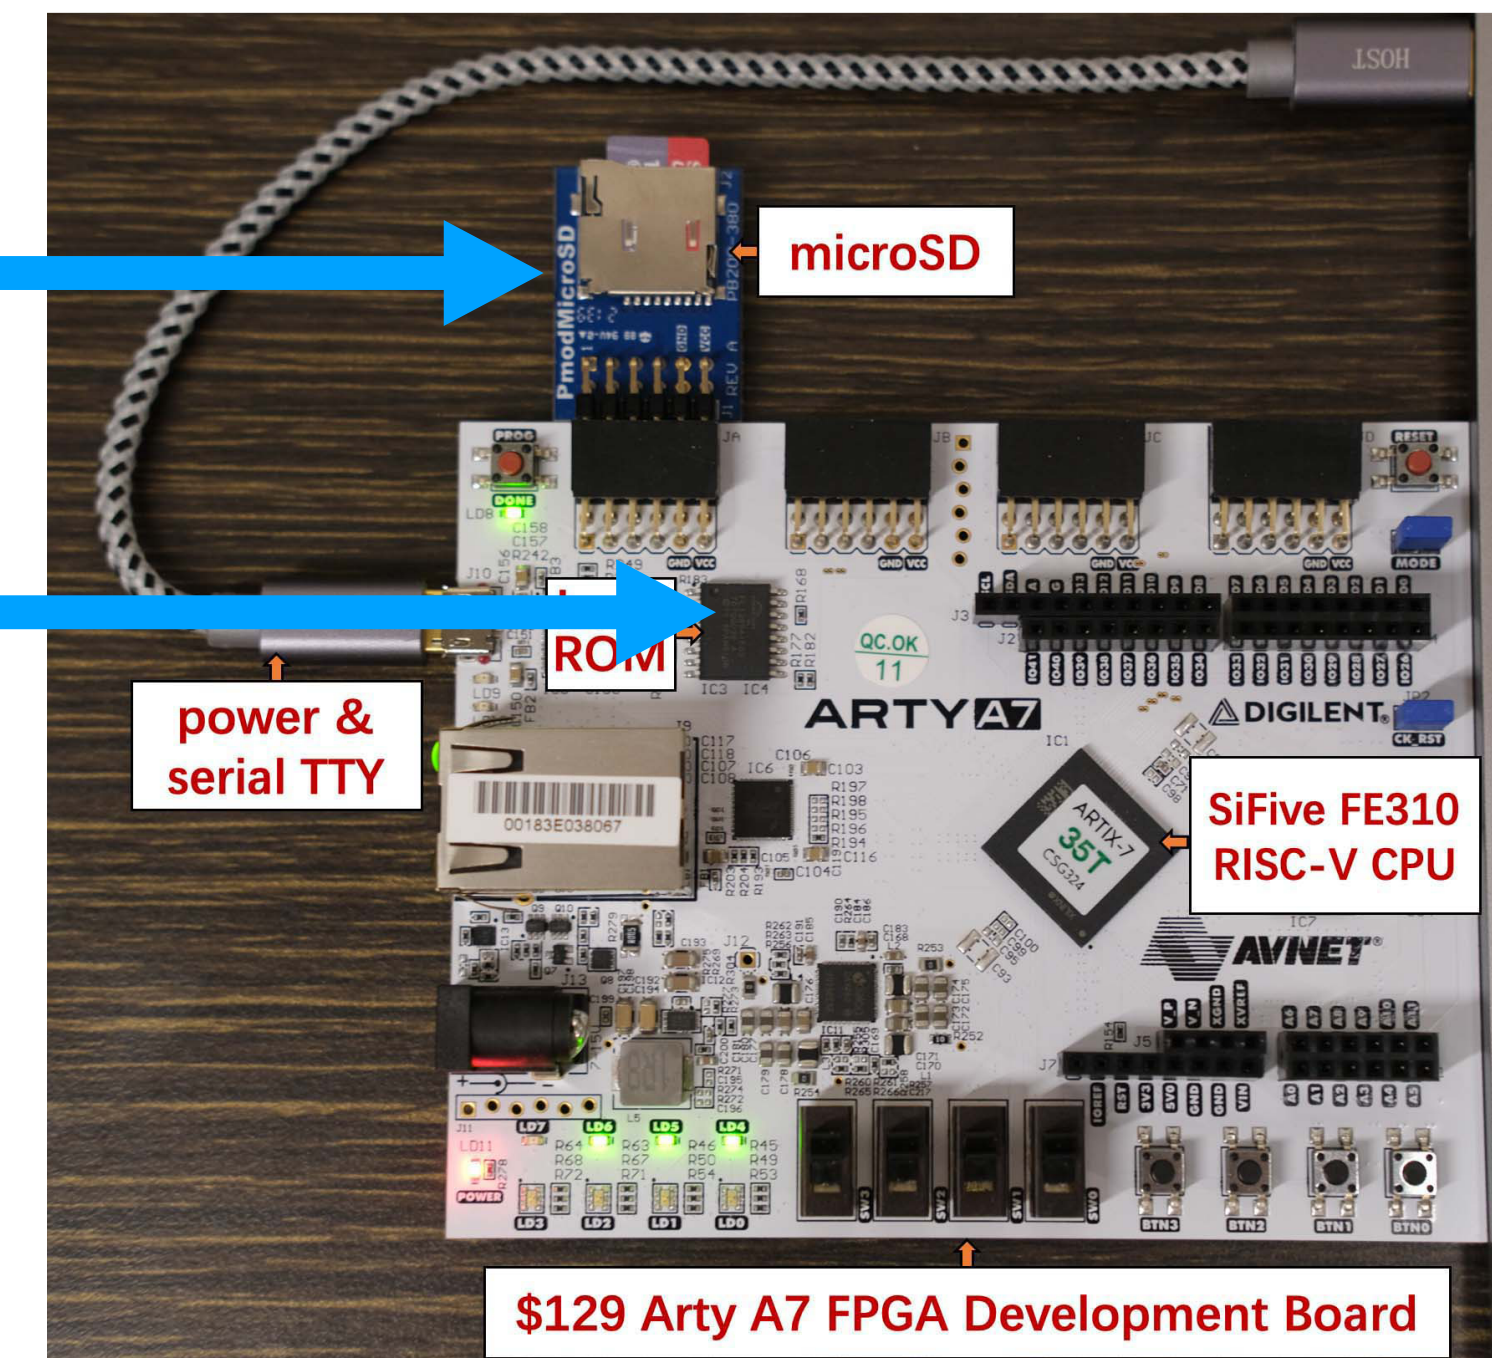
# Two more main functions to read

```
> cd egos-2000
> grep "main(" -r *
earth/earth.S: /* Call main() of earth.c */
earth/earth.c:int main() {
grass/grass.S: /* Call main() of grass.c */
grass/grass.c:int main() {
tools/mkrom.c:int main() {
tools/mkfs.c:int main() {

apps/*.c: /* Every application has a main() function */
```

# mkfs and mkrom

- During **make**, the RISC-V compiler compiles egos-2000
  - i.e., create everything under **build/**
- During **make install**,
  - **mkfs** creates **disk.img**
  - **mkrom** creates **bootROM.bin**





# Control flow provides a rough picture

```
> cd egos-2000
> grep "main(" -r *
earth/earth.S: /* Call main() of earth.c */
earth/earth.c:int main() {
grass/grass.S: /* Call main() of grass.c */
grass/grass.c:int main() {
tools/mkrom.c:int main() {
tools/mkfs.c:int main() {

apps/*.c: /* Every application has a main() function */
```

# Control flow provides a rough picture



We now know the **structure** of the work and **some details**.

# Read a repository: 3 passes

- 1st pass
  - read documents and filenames
- 2nd pass
  - track the execution: earth → grass → applications
- ➔ 3rd pass
  - read **details** of specific functionality, such as **system call**



# Consider apps/user/cat.c

```
10 #include "app.h"
11 #include <string.h>
12
13 int main(int argc, char** argv) {
14     if (argc == 1) {
15         INFO("usage: cat [FILE]");
16         return -1;
17     }
18
19     /* Get the inode number of the file */
20     int file_ino = dir_lookup(grass->workdir_ino, argv[1]);
21     if (file_ino < 0) {
22         INFO("cat: file %s not found", argv[1]);
23         return -1;
24     }
25
26     /* Read and print the first block of the file */
27     char buf[BLOCK_SIZE];
28     file_read(file_ino, 0, buf);
29     printf("%s", buf);
30     if (buf[strlen(buf) - 1] != '\n') printf("\r\n");
31
32     return 0;
33 }
```

```
10 #include "app.h"
11 #include <string.h>
12
13 int main(int argc, char** argv) {
14     if (argc == 1) {
15         INFO("usage: cat [FILE]");
16         return -1;
17     }
18
19     /* Get the inode number of the file */
20     int file_ino = dir_lookup(grass->workdir_ino, argv[1]);
21     if (file_ino < 0) {
22         INFO("cat: file %s not found", argv[1]);
23         return -1;
24     }
25
26     /* Read and print the first block of the file */
27     char buf[BLOCK_SIZE];
28     file_read(file_ino, 0, buf);
29     printf("%s", buf);
30     if (buf[strlen(buf) - 1] != '\n') printf("\r\n");
31
32     return 0;
33 }
```



# Send requests to file system

`file_read()`  
`dir_lookup()`

Code of cat → User library

**Application**

→ **Grass kernel**

`sys_send()`

↓  
**Kernel servers** (GPID\_DIR, GPID\_FILE, ...)

# Receive data from file system

`file_read()`  
`dir_lookup()`

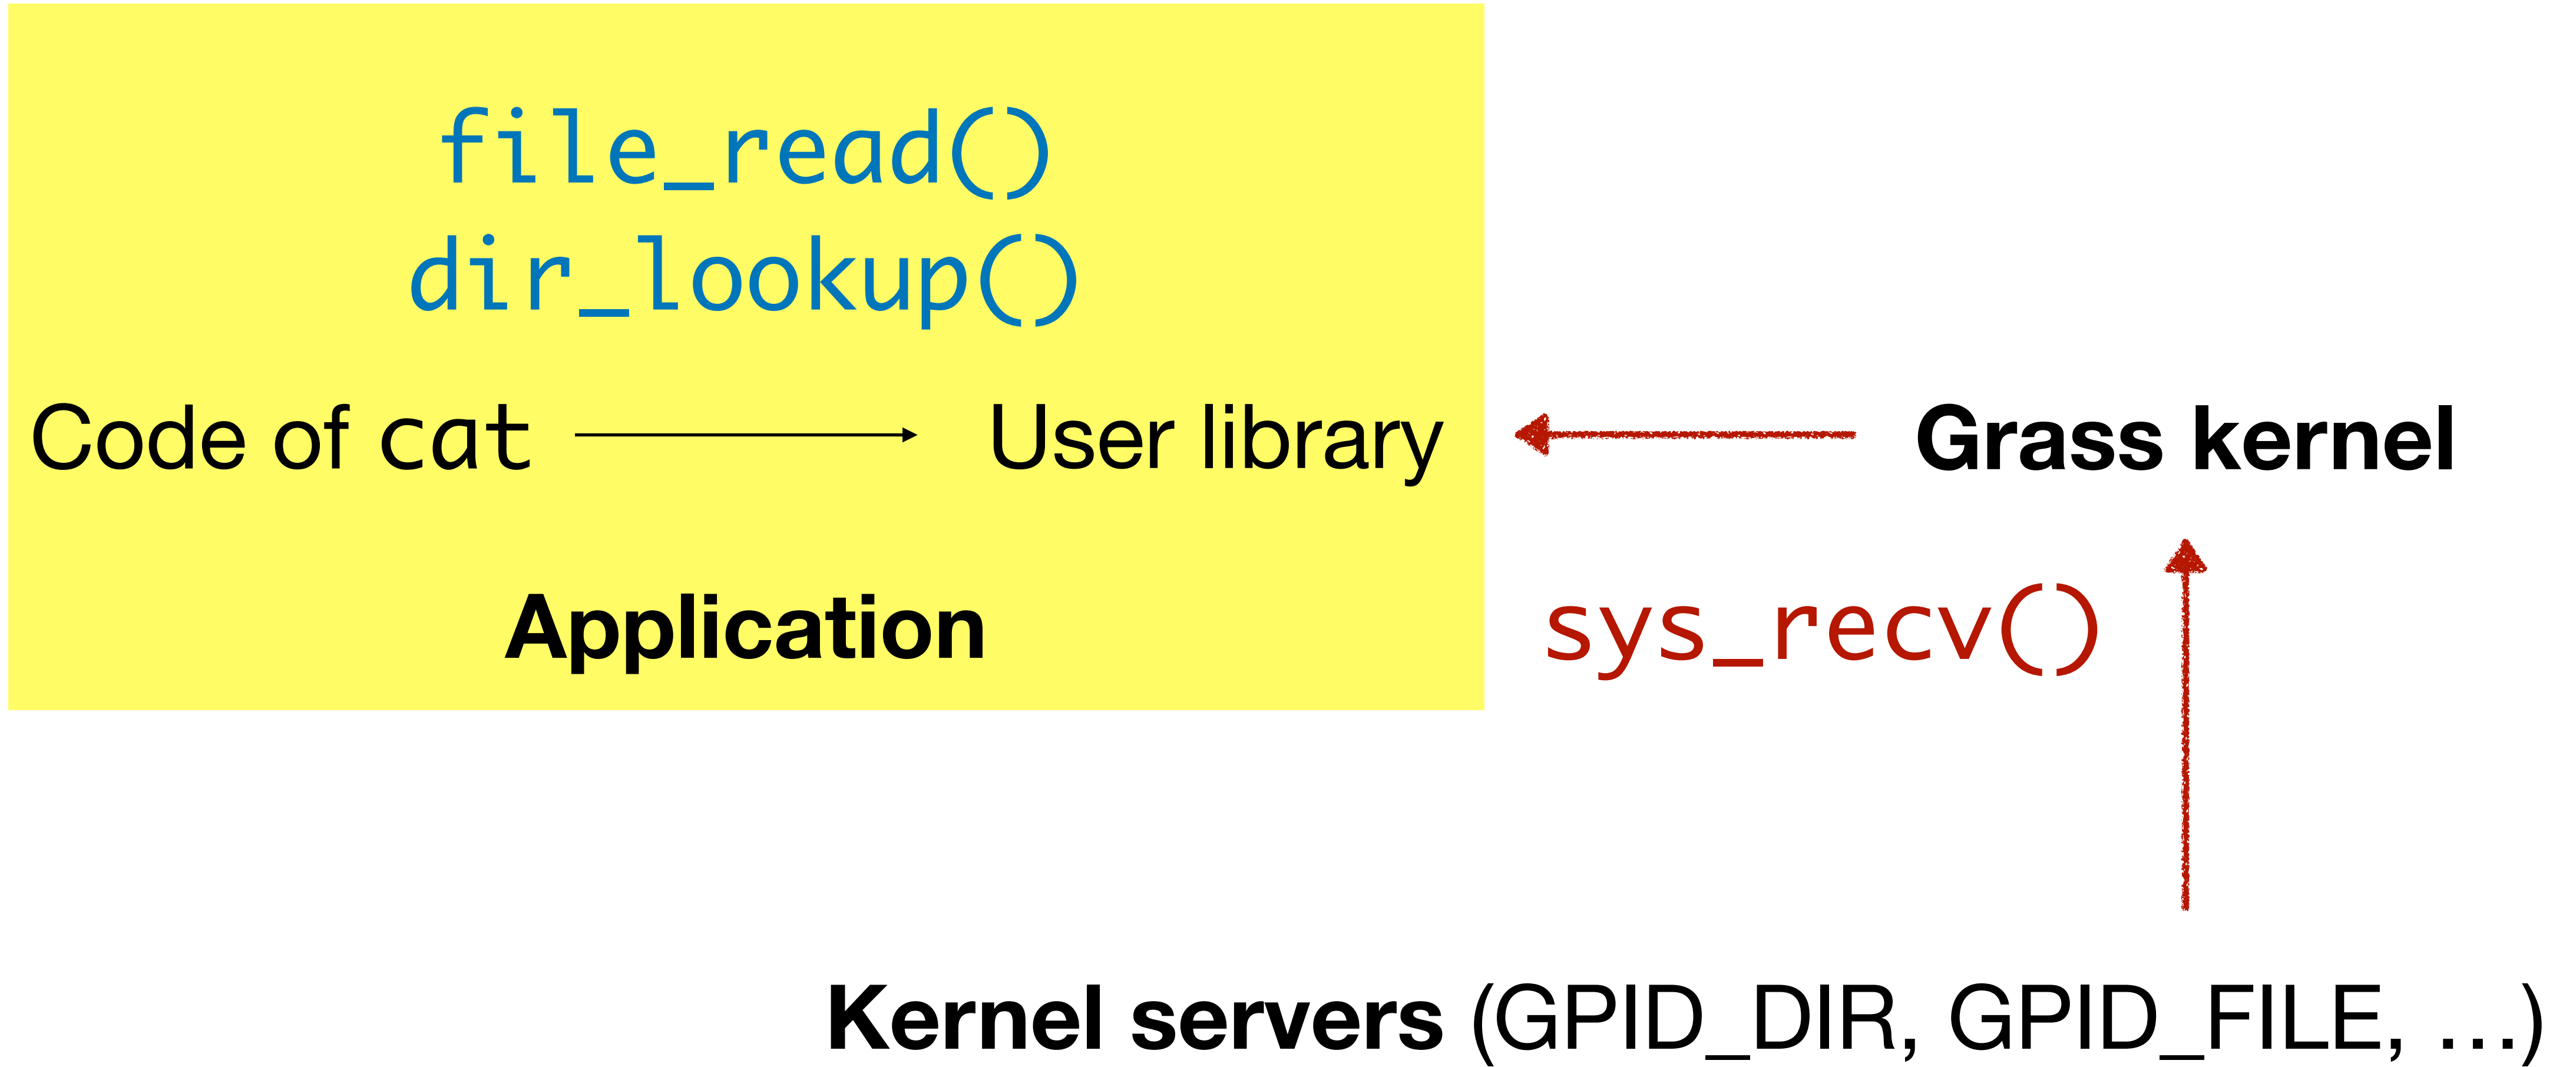
Code of cat → User library

**Application**

← **Grass kernel**

`sys_recv()`

**Kernel servers** (GPID\_DIR, GPID\_FILE, ...)



# Data structures for **system call**

```
struct syscall {  
    enum syscall_type type;  
    struct sys_msg msg;  
    int retval;  
};
```

```
enum syscall_type {  
    SYS_UNUSED,  
    SYS_RECV,  
    SYS_SEND,  
    SYS_NCALLS  
};
```

# App invoking syscall step#1

```
static struct syscall *sc = (struct syscall*)SYSCALL_ARG;
```



a well-known memory address

```
static void sys_invoke() {  
    *((int*)0x2000000) = 1;  
}
```

```
int sys_send(int receiver, char* msg, int size) {  
    if (size > SYSCALL_MSG_LEN) return -1;
```

```
    sc->type = SYS_SEND;  
    sc->msg.receiver = receiver;  
    memcpy(sc->msg.content, msg, size);  
    sys_invoke();  
    return sc->retval;
```

```
}
```

# App invoking syscall step#2

```
static struct syscall *sc = (struct syscall*)SYSCALL_ARG;
```

```
static void sys_invoke() {  
    *((int*)0x2000000) = 1;  
}
```

```
// The sys_send function takes 3 parameters
```

```
int sys_send(int receiver, char* msg, int size) {  
    if (size > SYSCALL_MSG_LEN) return -1;  
  
    sc->type = SYS_SEND;  
    sc->msg.receiver = receiver;  
    memcpy(sc->msg.content, msg, size);  
    sys_invoke();  
    return sc->retval;  
}
```

# App invoking syscall step#3

```
static struct syscall *sc = (struct syscall*)SYSCALL_ARG;
```

```
static void sys_invoke() {  
    *((int*)0x2000000) = 1;  
}
```

```
int sys_send(int receiver, char* msg, int size) {  
    if (size > SYSCALL_MSG_LEN) return -1;  
    // Prepare the system call data structure  
    sc->type = SYS_SEND;  
    sc->msg.receiver = receiver;  
    memcpy(sc->msg.content, msg, size);  
    sys_invoke();  
    return sc->retval;  
}
```

# App invoking syscall step#4

```
static struct syscall *sc = (struct syscall*)SYSCALL_ARG;
```

```
static void sys_invoke() {  
    *((int*)0x2000000) = 1; // Trigger a software interrupt  
} // which is interrupt #3
```

```
int sys_send(int receiver, char* msg, int size) {  
    if (size > SYSCALL_MSG_LEN) return -1;
```

```
    sc->type = SYS_SEND;  
    sc->msg.receiver = receiver;  
    memcpy(sc->msg.content, msg, size);  
    sys_invoke();  
    return sc->retval;
```

```
}
```

# App invoking syscall step#5

```
void kernel() {
    int mcause;
    __asm__ volatile("csrr %0, mcause" : "=r"(mcause));

    int id = mcause & 0x3ff;
    if (mcause & (1 << 31)) {
        if (id == 3) { syscall_handler(); }
        if (id == 7) { timer_handler(); } // last lecture
    } else {
        fault_handler();
    }
}
```



# Homework

- Handle system call with the `ecall` instruction
- Replace `*((int*)0x2000000) = 1` by `asm("ecall")` which triggers exception#8/#11 instead of interrupt#3
- **P2** will be due on **Mar 24**
- Next lecture: memory exception and protection