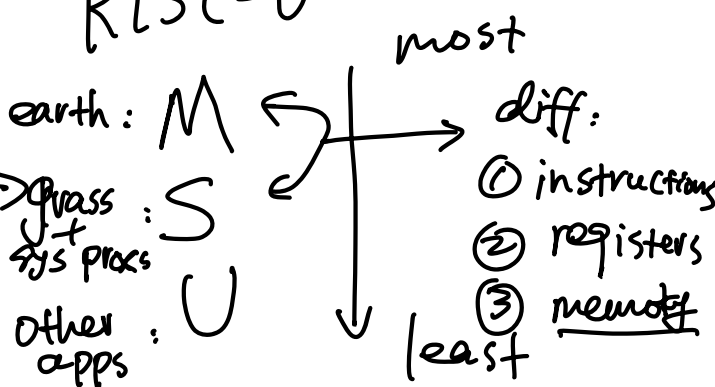


- ✓ 1. last time: egos
- ✓ 2. egos exception handling
- ✓ 3. syscall in egos
- 4. CPU privilege levels
- 5. System virtualization


---  
 • Lab4: OS protection

# RISC-U

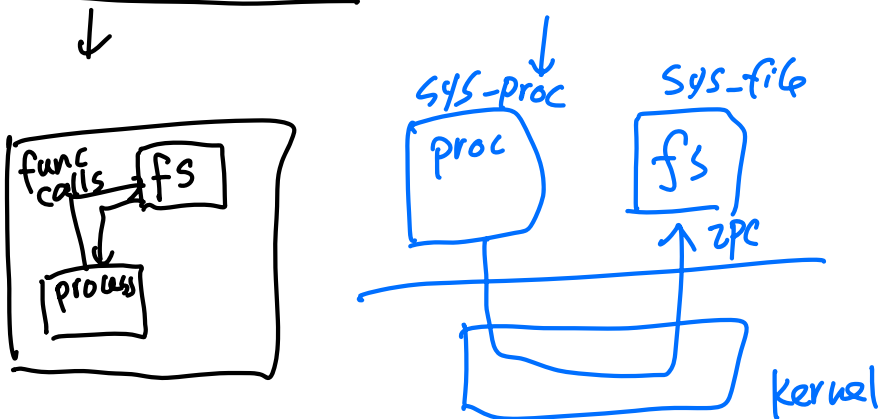


- 1. last time: egos
- ✓ OS organization
- ✓ egos design
- ✓ egos boot

\* kernel ~ three handlers <sup>user</sup>




## monolithic-kernel, micro-kernel



- boot:
  - ↳ CPU 0x2040\_000
  - ↳ earth.S
  - ↳ main (earth.c)
  - ↳ main (grass.c)
  - ↳ sys-proc
    - ↳ sys-file
    - ↳ sys-dir
    - ↳ sys-shell

Q: 3 ways to trap to kernel, ?

- interrupt
- syscall
- exception



(a) Interrupt

→ what code CSR

Q: How does the CPU determine what to execute when an interrupt is triggered?

mtvec

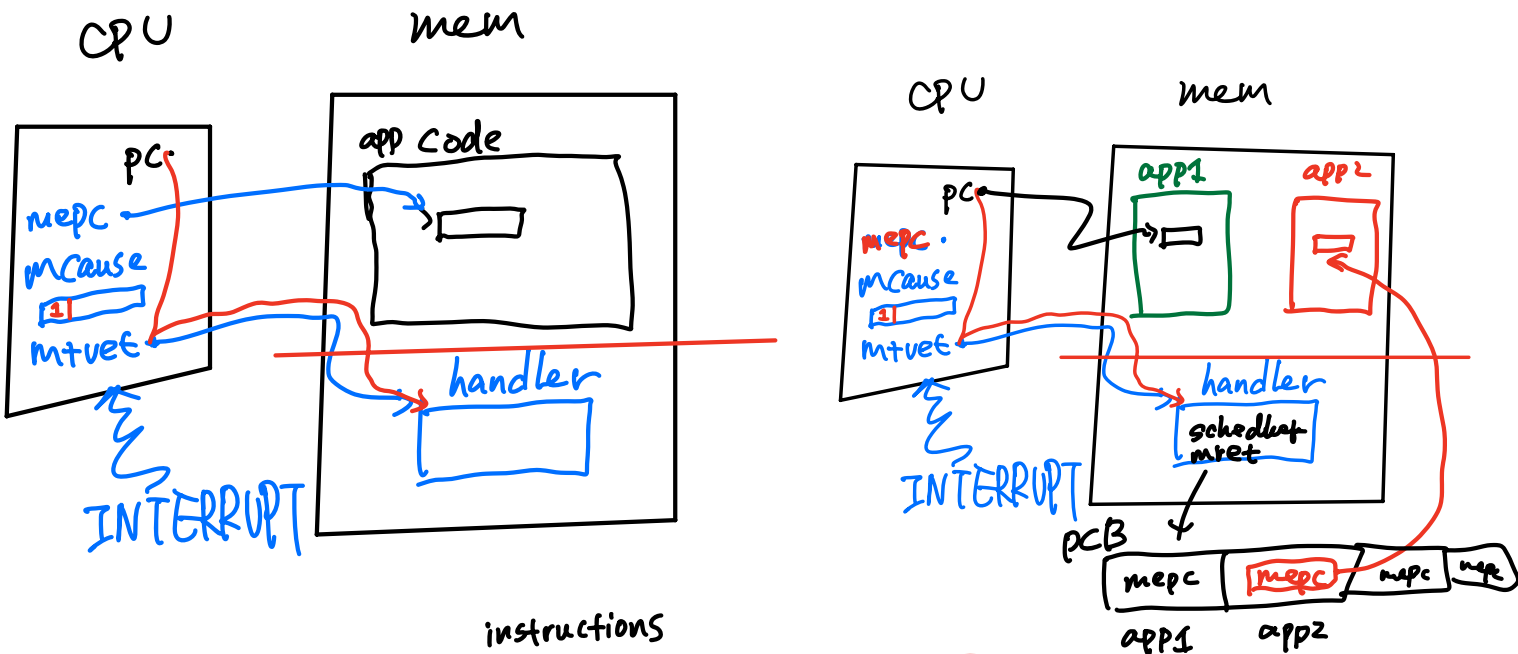
Q: After handling an interrupt, where does the CPU resume execution?

mepc

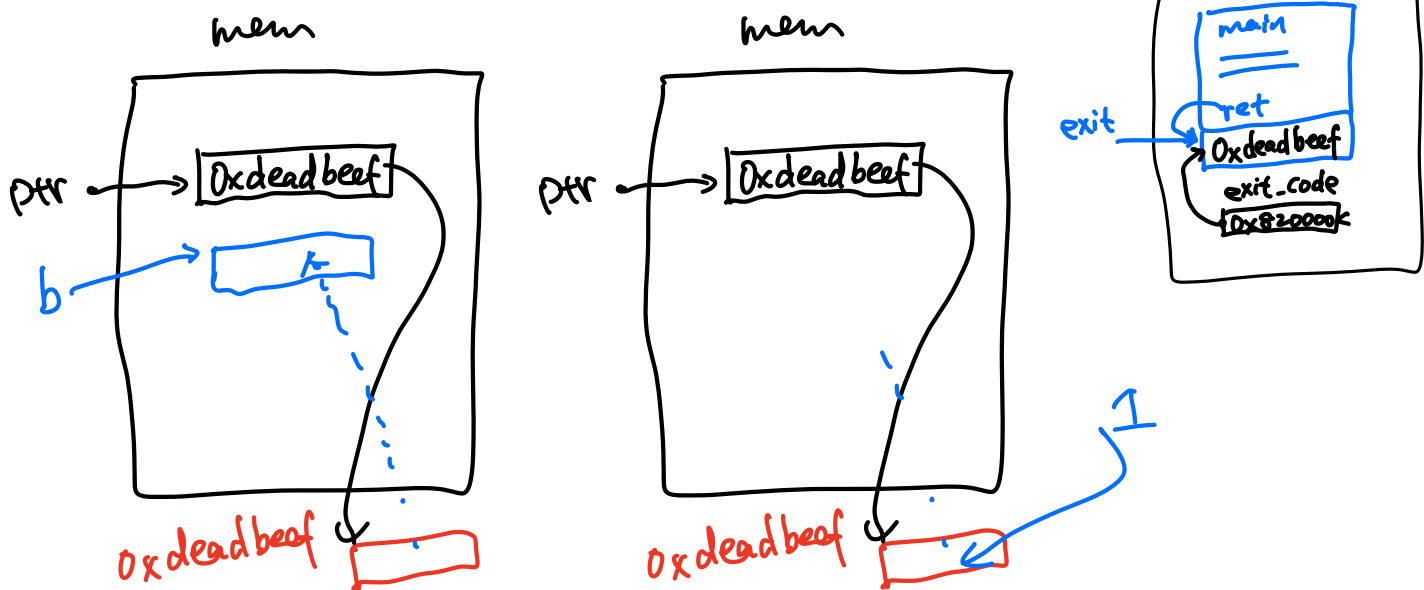

mret

Q: How does the kernel identify which interrupt was triggered?

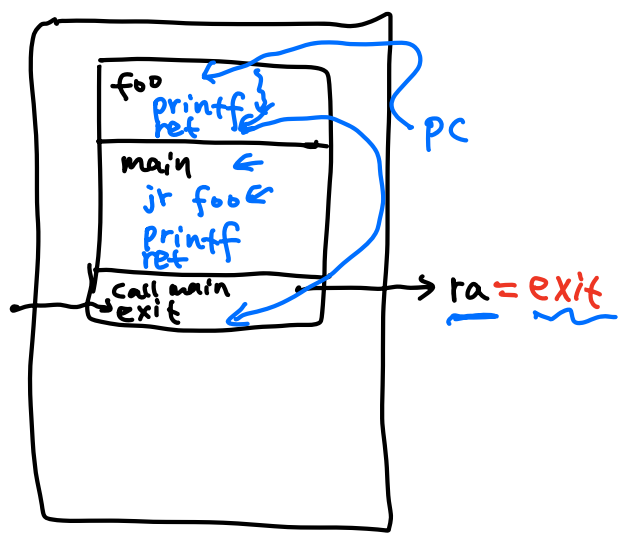
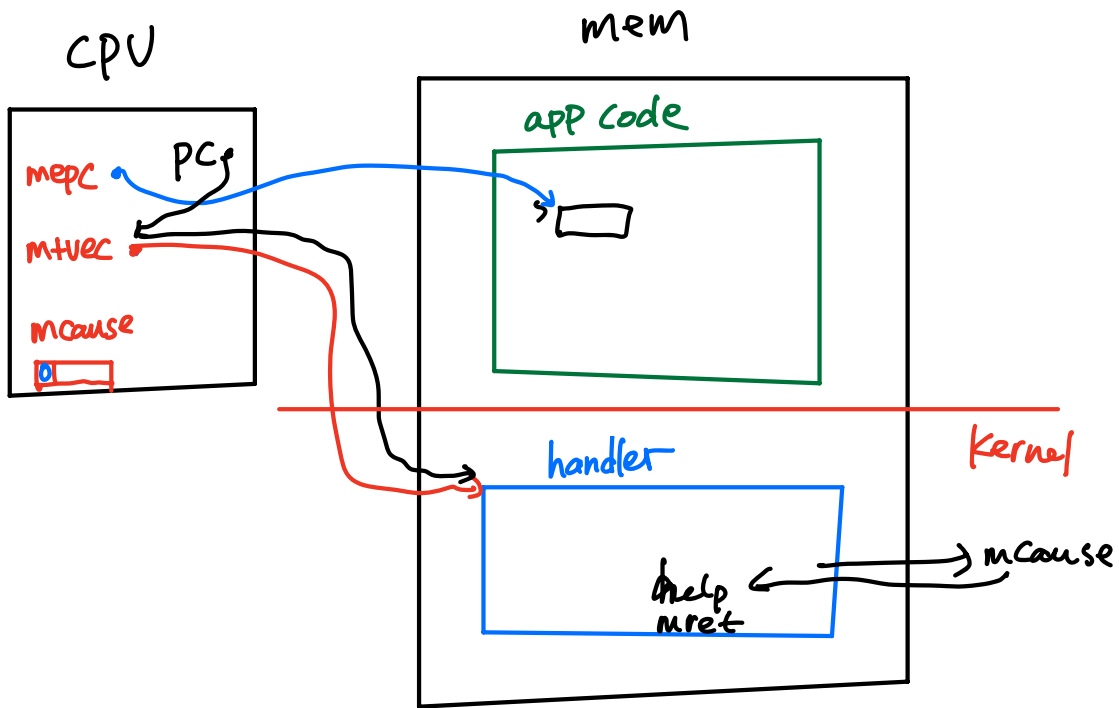
mcause



(b) Exception







Linux:

fs → (c) Syscall

- open/close/read/write
- socket/bind/connect/listen
- mkdir/chmod/stat
- Pipe/dup2

Q1: how does control transfer to the kernel? (trap to kernel)

exception

interrupt

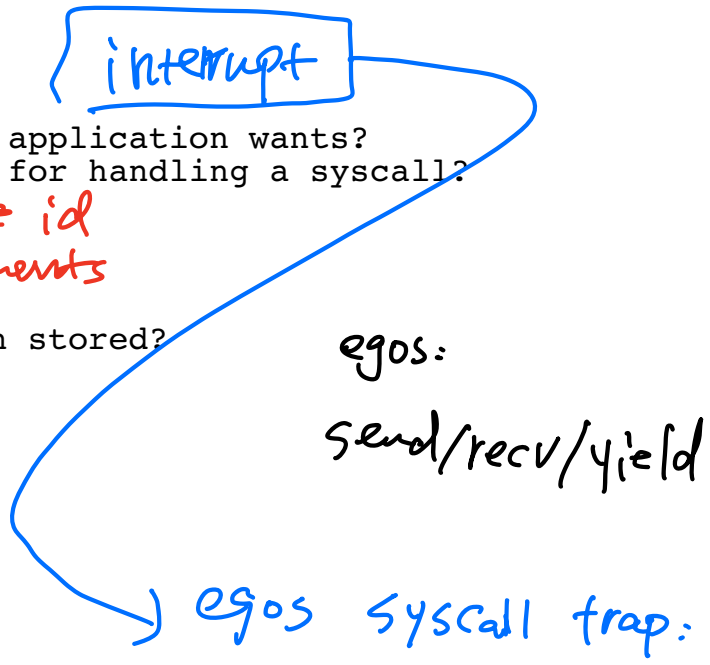
Q2: how does kernel understand what the application wants?  
That is, what information is needed for handling a syscall?

pid (x) ~~addr~~ • ~~type~~ id  
uapc (x) • arguments

Q3: where is syscall-related information stored?

• registers  
OR  
• mem } stack  
          } else

egos:  
send/recv/yield



egos syscall trap:  
0x200\_0000 = 1  
↳ MSIP of \$mip  
↳ pending interrupt  
↳ software int

Q1: sys\_invoke()

Q2: struct syscall

Q3: well-known mem location

[Skipped]

Q: How could I know what privilege level the current CPU is running in?

"""

RISC-V deliberately doesn't make it easy for code to discover what mode it is running in because this is a virtualisation hole. As a general principle, code should be designed for and implicitly know what mode it will run in. Applications code should assume it is in U mode. The operating system should assume it is in S mode (it might in fact be virtualised and running in U mode, with things U mode can't do trapped and emulated by the hypervisor).

"""

[from <https://forums.sifive.com/t/how-to-determine-the-current-execution-privilege-mode/2823>]