

Lecture 9 — xv6 hands-on.

CS 744

Lecture 10 — system calls & memory mgmt.

System call mechanism.

- ✓ + how to invoke?
- ✓ + how to switch mode of execution?
- . + how to handle arguments & return values?
- . + how to identify which system call?
- ✓ + how to handle context of execution?

* Lab 3 available
xv6-based
system calls,
process abstraction
memory abstraction
④ xv6-book |
reference

(i) How assisted mechanism for invocation

is an instruction(s) of an ISA.

x86: $\text{int } \text{0x80}$ ~ explicit S/w interrupt

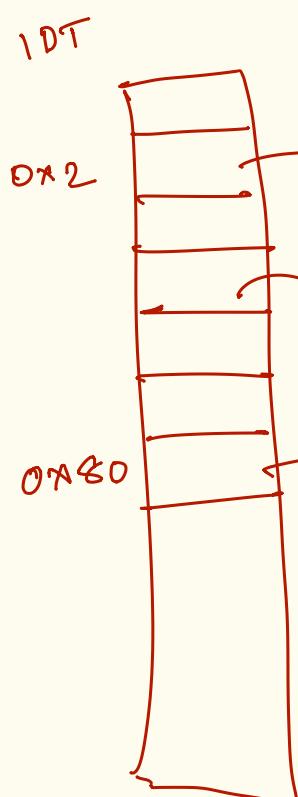
instruction argument - interrupt vector

- * ↗ privilege level to Ring 0
- * save context of user process on kernel stack
- * S/w to using kernel stack
- * jump to interrupt handler. ~ depends on IDTR & interrupt vector

(ii) how to jump to the interrupt handler?

IDT — interrupt descriptor table

— array of function pointers (interrupt handling functions)



⌚ how to get IDT?

IDTR

{ register to
store base
addr of
the IDT.

⌚ OS maintains
IDT & IDTR.

(iii) how to identify the system call and its arguments?

— via syscall number

— before ~~invoking~~ invoking system call setup syscall number & arguments in CPU reg.

eax ← syscall number

ebx

cx

dx

:

int 0x80

} arguments.

user-space
wrapper f'n.

sets ~~this~~ up this
invocation

Open: }

```
    mov eax, SYS_OPEN  
    {  
        mov ebx  
        ecx  
        :  
        int ox 80
```

(iv) system call handler:

```
syscall_handler {
```

```
    fn_ptr = syscall_table [ value of  
                           eax ];
```

^{from}
^{on} kernel stack

```
    call fn_ptr;
```

```
    iret; → in return from interrupt.
```

{ } instruction of the ISA.

|
return
values
are via
register
(eax).

(*) xv6 mechanism details:

traps.h — list of all interrupts

vectors.S — IDT entries for each interrupt vector

trapasm.S ~ alltraps & label / address
entry point for all
interrupts

trap.c ~ generic IDT handler

syscall.c ~ syscall handler

