


CPU Virtualization

Review

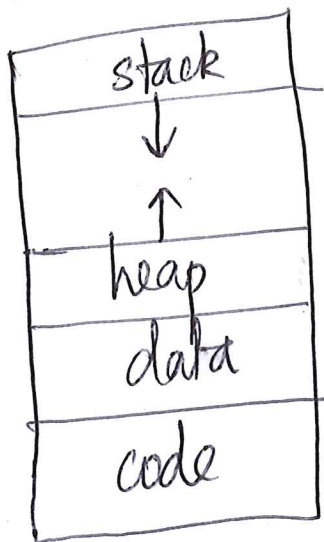


- OS
- Virtual, Concurrency, Persistence




- Process - machine state \Rightarrow PCB.

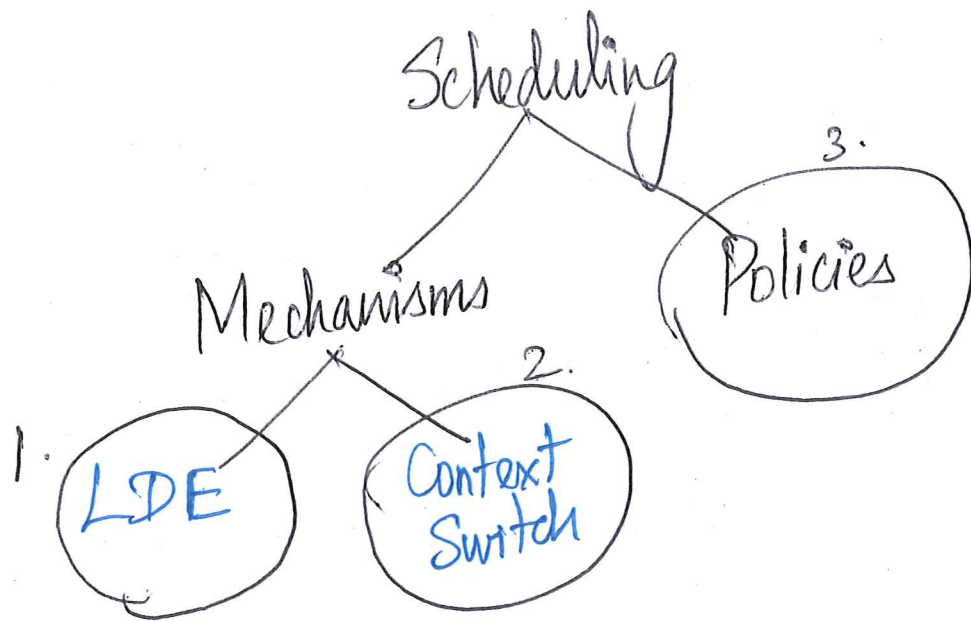


- Program vs Process



CPU: P₁ P₁ P₂ P₂






Scheduler - component of OS.
LDE

Direct Execution.


Q: How can ~~we~~ the OS efficiently virtualize the CPU?



OS P₁ OS P₂ OS P₁ ...

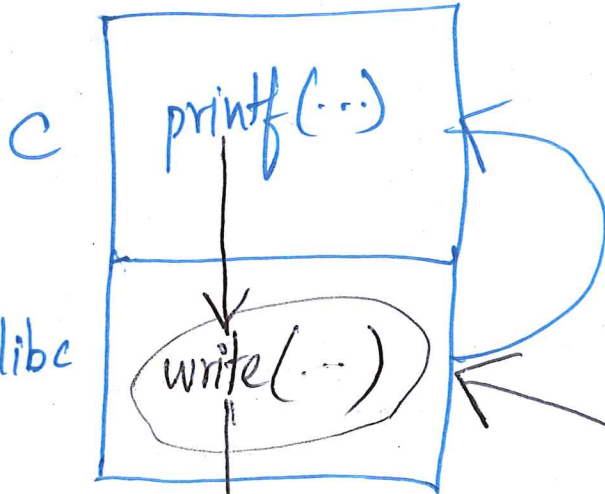
Prob #1: Restricted Ops

soln: kernel mode
vs
user mode



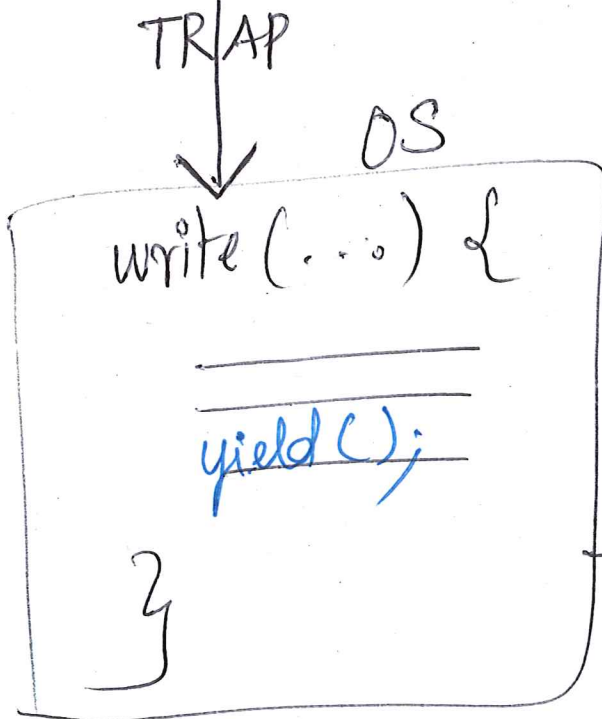
System calls

0 - user
1 - kernel



User level
Prob #2: How to switch from
P_i to OS?

soln: Timer interrupt





Kernel level

System call

write(...)

open 1
read 2
write 3

Assembly:



system calls.

TRAP

OS

Trap Handler
syscall()

write();



Return From Trap

X86 - HW

XV6 - OS


Example: Linux "write" system call


```
C code:  
...  
printf("Hello world\n");  
...
```

```
libc:  
%eax = sys_write;  
int 0x80
```

User mode

Kernel
mode

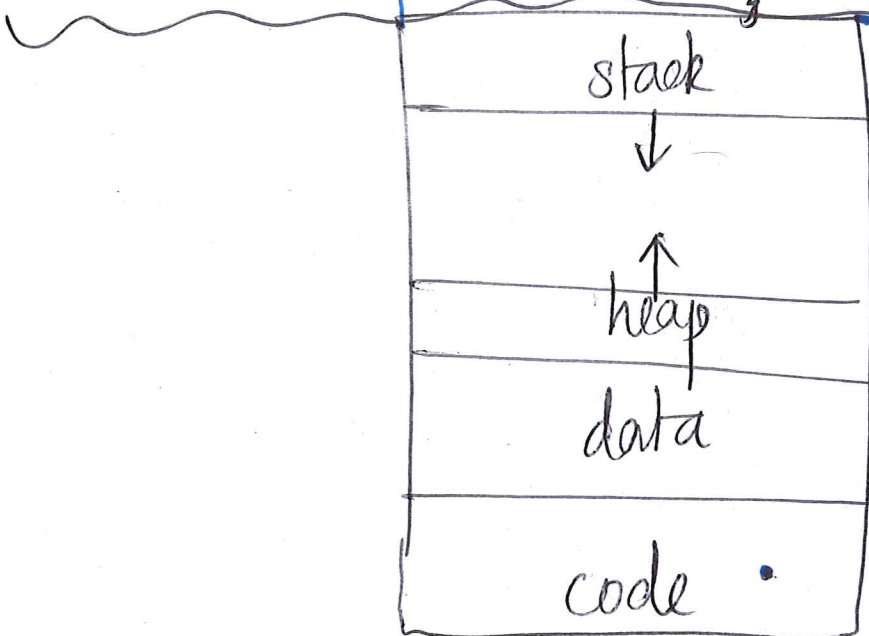
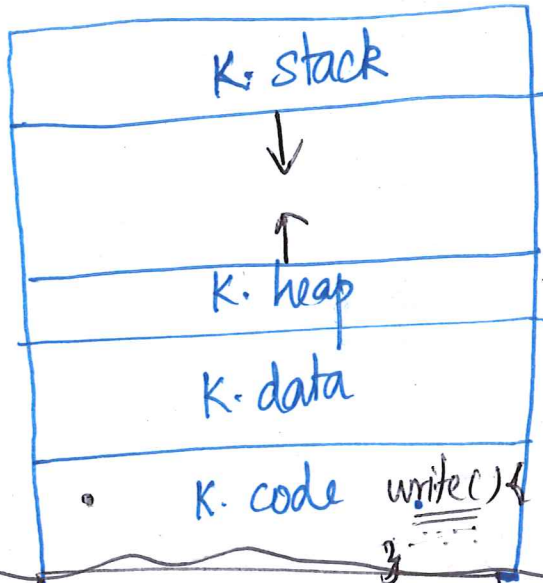




timer() ↓




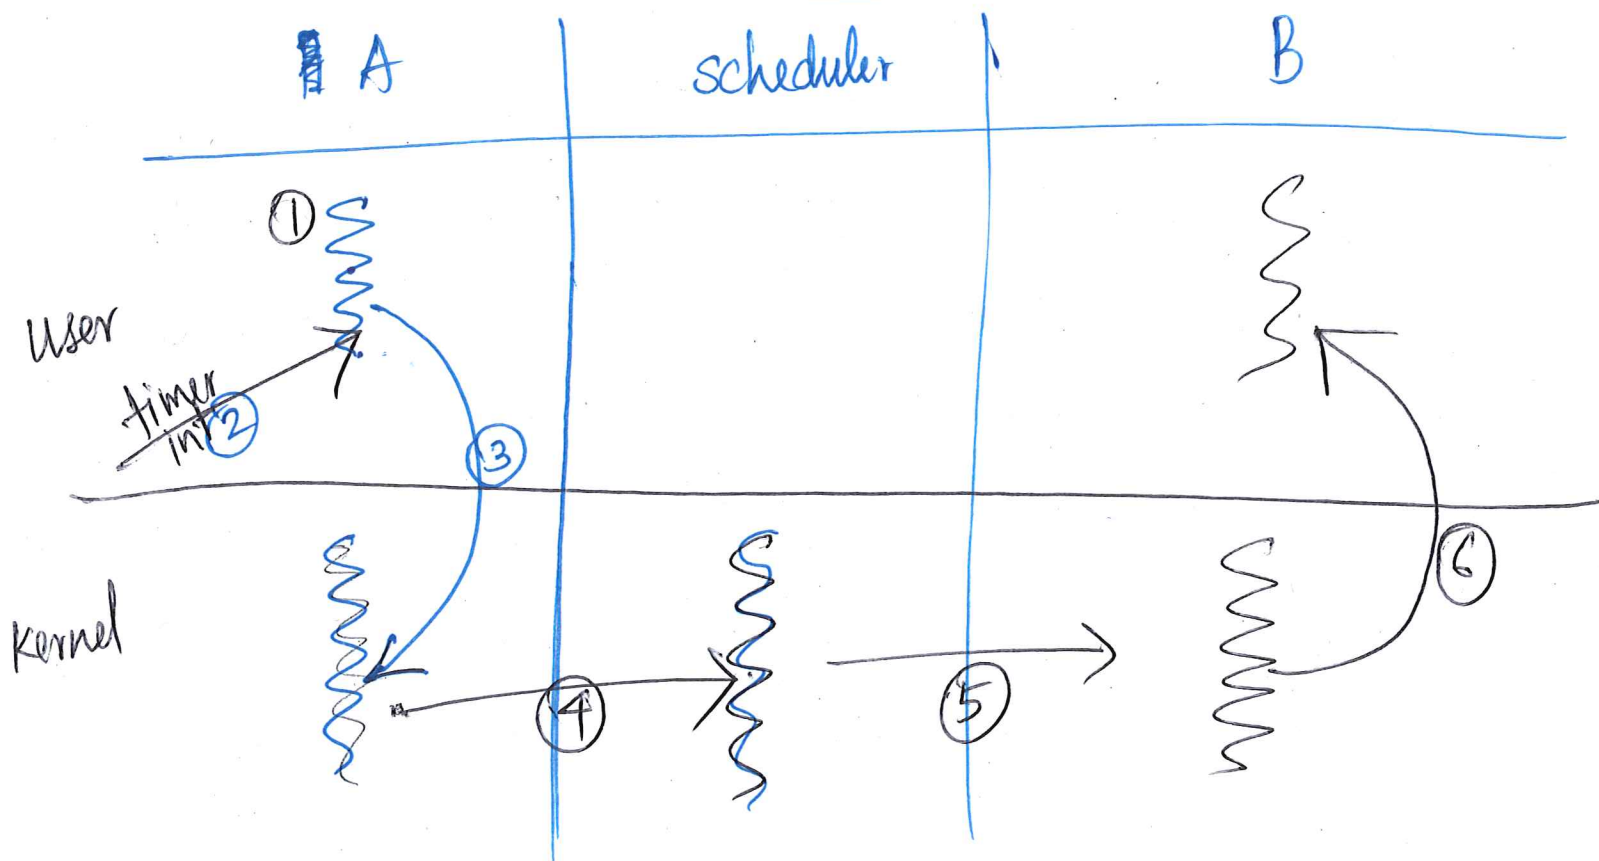
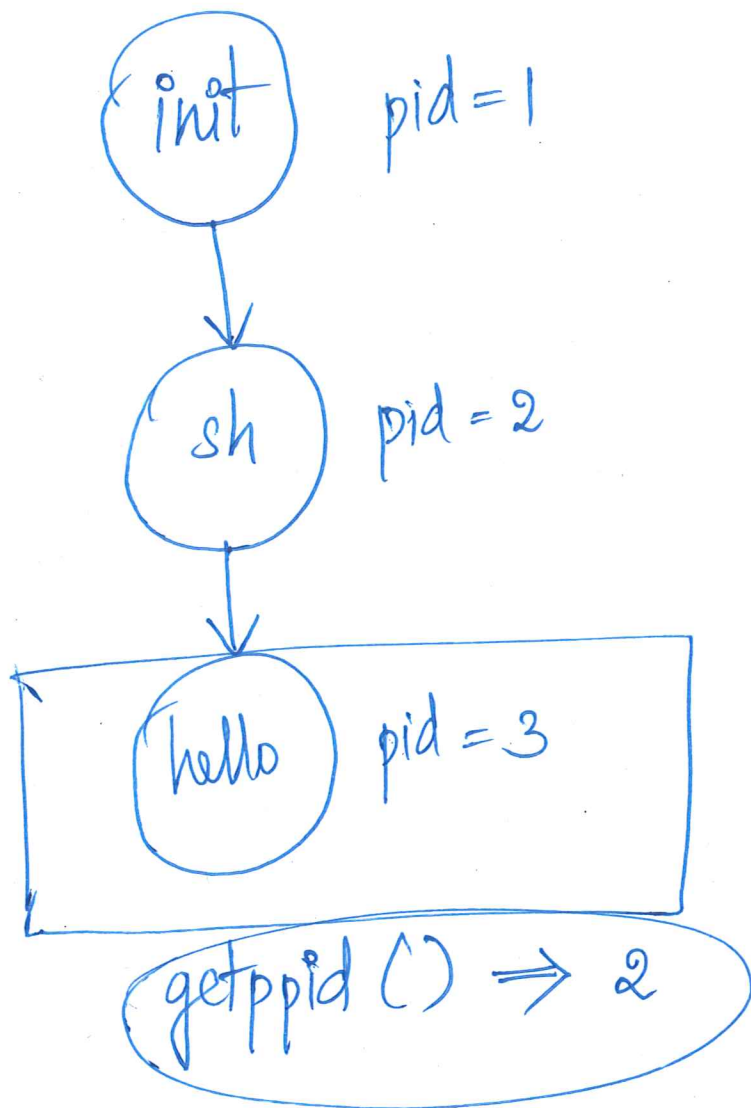
P₁ }



HW → user level
regs of P → Trap Frame

OS → kernel level
regs of P → Context




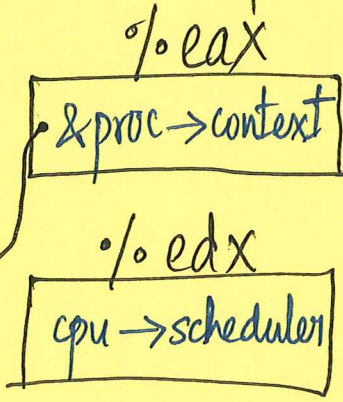
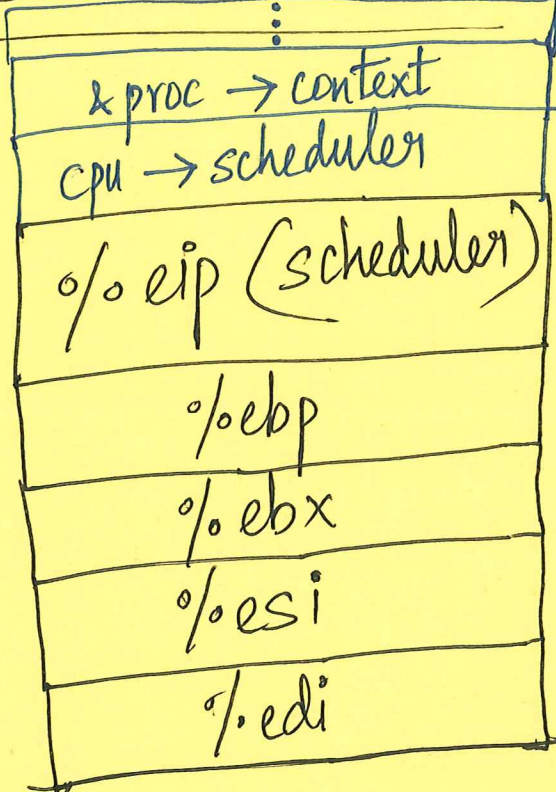


Context Switch (From Proc A to scheduler)

Process A's Kernel stack ①



Scheduler stack

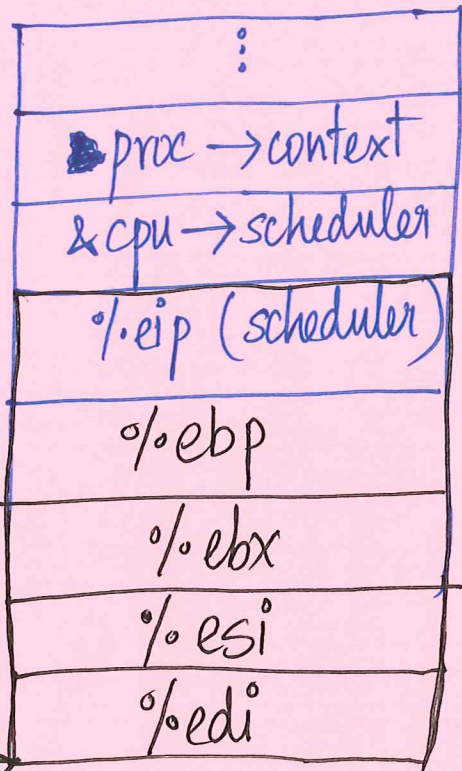


NOTE: Arrows (→) indicate the state after switching from A to scheduler.



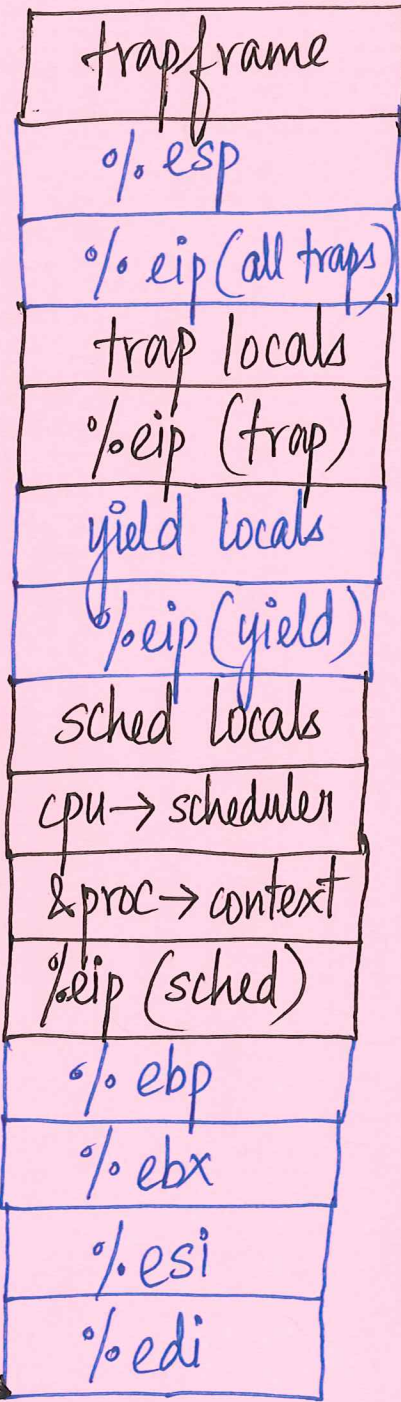
Context Switch

From scheduler to Proc B

Scheduler stack



Process B's kernel stack



NOTE: Arrow indicate the state after switching from scheduler to B.