

Today

CPU virtualization \Leftarrow
mechanisms \Rightarrow Limited
 \rightarrow Scheduling Policy [Direct Execution]


Q) How to build a scheduling policy?

Memory Virtualization

Scheduling

\Rightarrow FILO
(FCFS)

bad:



A, B, arrive
C

Metric:
(Turnaround Time:)

$T_{\text{completes}} - T_{\text{arrives}}$

"run to completion"

\Rightarrow SJF "shortest job first"



"generalized" SJF:

STCF (shortest time to completion first)

short long
A, B

preempt B



C arrives,
C very short

⇒ Metric #2: Response Time:

A, B, C arrive

$T_{\text{first runs}} - T_{\text{arrives}}$



TS
context switch

Policy:
Round Robin

TS → Time Slice (Quantum):

multiple of
how often timer
interrupt goes
off

"Real" Scheduler :

- don't know job length
(how to learn?)
- workload:
 - interactive
 - long-running ("batch" or "background")
- jobs ⇒ do I/O
(don't just use CPU)

Key: use history to predict future

Policy: Multi-Level Feedback Queue
(MLFQ)

priority

high	$Q_3 \rightarrow \underline{A}, \underline{B}$	$\overset{10ms}{\text{}} \Rightarrow$	a process can be on only 1 Queue
<u>medium</u>	$Q_2 \rightarrow C$	$\overset{20ms}{\text{}}$	
low	$Q_1 \rightarrow E, L$	$\overset{40ms}{\text{}}$	

Priority: \Rightarrow process at highest queue runs

when to change queues?

\Rightarrow each queue has time slice

(how long job will run at that priority)

\Rightarrow when time slice is used, move down one level

initial? job \Rightarrow highest priority


\Rightarrow if ≥ 1 job @ some priority, use round robin

Example:

1 long running job: L

occasionally, an "interactive" job shows up
 $\leftarrow \leftarrow \leftarrow$

works $\rightarrow 1, \rightarrow 2, \rightarrow 3$



another long job: M



starvation :

"too many" short jobs
arrive + run,
long-running jobs may
never get service

solve?


periodically, reset
jobs \Rightarrow top ("high")

what about I/O?

=> approach #2:

if do I/O, stay @
same queue

(~~reset counter,~~
to run for full
time slice)



Problem: "gaming"
=> retaining high priority
artificially

Memory Virtualization

Each
Process:

Virtual
Memory

C pointers:
virtual addresses

Shared

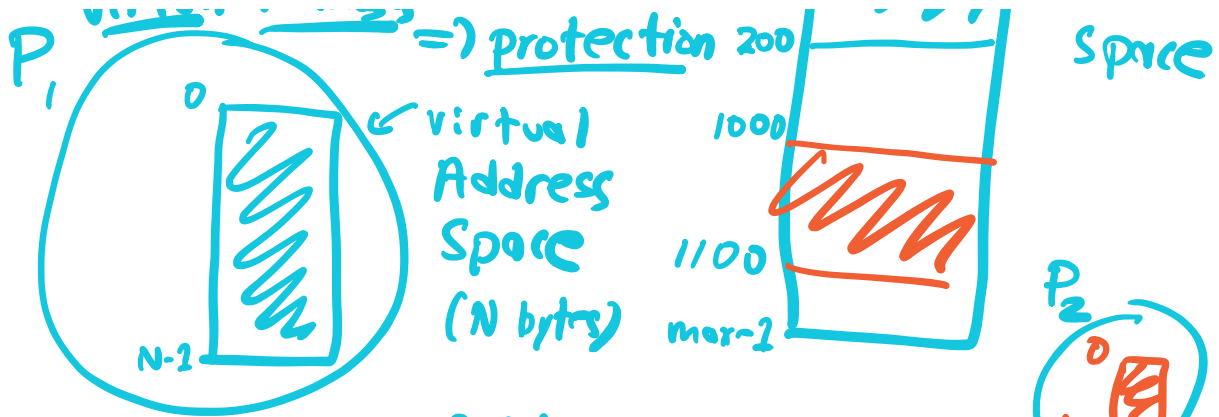
by all
running
programs

=> efficiently

Physical

Memory





Virtual Addr \Rightarrow
Physical Addr.


"Address Translation"

\Rightarrow this happens on every memory access

(virtual)
e.g. load addr, R₁

PC: virtual address
then load: another virt. addr.

upon
memory reference:



(per CPU)

MMU: hardware
(on CPU)

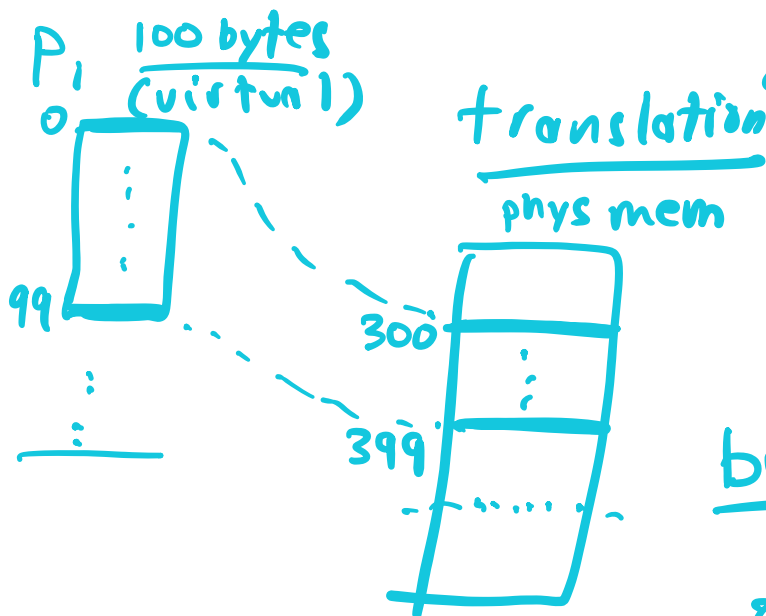


memory management unit
(assist in address translation)

simplest form of MMU:

⇒ base / bounds

MMU
→ 2 registers



base
bounds
protection

base: 300

base: where is first byte of AS in phys memory?

on every mem. access:

virt addr → MMU

base + virt. addr ⇒

phys. addr

bounds: check to make
sure "legal" address

make qemu-nox