

10/31

Scheduling Wrapup & Address Translation

Scheduling


- Average response time (interactivity)
- FIFO
- SJF
- RR (a single time quantum, impacts long & short jobs differently)
- MLFQ

Use the past to predict the future

- Multiple queues, each with its own time quantum
- start all jobs at the top queue
- if task blocks before the time is up, stay or move up a queue
- if task uses all of its time, need more time on CPU, moves down a queue




MLFQ Example




5ms



10ms




20ms




40ms


① A, B, C arrives,
scheduler picks A to run




5ms



10ms



20ms



40ms

② A ran for 1ms, then blocks
scheduler picks B to run




5ms



10ms




20ms




40ms


③ While B runs, A unblocks,
B runs for 5ms




5ms



10ms



20ms



40ms

④ C gets to run,
blocks after 2ms



Completely Fair Scheduler (CFS)


- ideally, if there are only 2 tasks, each should take 50% of CPU time
- track how long a task had been on the CPU
- red-black tree (self balanced, CPU time-ordered, binary search tree)
efficiently tells you which process has the least time on the CPU.
- * schedule one w/ least amount of CPU share (to make it fair!)

XK scheduler?

A variation of RR (based on priority order)



Address Translation (Memory System)



Resource Allocation Problem

→ How do we allocate?

→ Simple approach: one process at a time

→ give the entire physical memory to the process

→ no translation needed!

Problem?

→ byte addressable

→ ~200 cycles access latency

→ process needs memory to run




Support Multiple Processes

Let process A, B, C run in disjoint sections of physical memory

→ should processes be aware of where it is in physical memory?

→ virtual memory ("infinite" & private memory) vs. physical memory

- virtual address vs. physical address



★ How do we do address translation?

$$PA = VA + \text{base} \quad (VA < \text{bound})$$

base register

0

 for process A
bound register

100

base

100


 bound

200

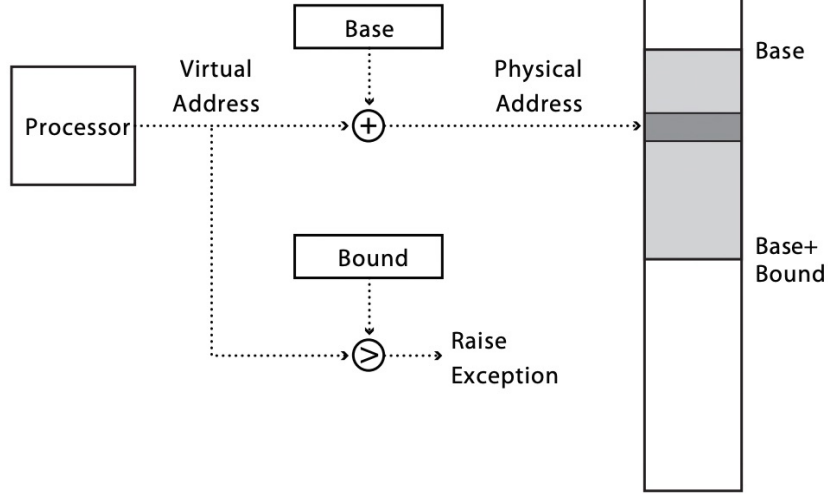
 for process B



Processor's View



Implementation



- Limitations =
- ① processes need to fit into available space in physical memory
 - ② hard to grow the process
 - ③ is it possible to share memory btwn processes?



Another Approach = Paging

- Processes don't need all of its memory at once
- Divide process memory into fixed size chunk & only have what it needs in ^{physical memory}

