

# Lecture: Scheduling

1

## Questions:

- Difference between alloc / scheduling?
- Pre-emptive vs. non?
- Different policies: how behave?  
(+ 's, - 's)

## Review

Processes + low-level mechanisms  
how to preempt, save/restore ctxt

⇒ Resources: what OS's manage  
things that ~~are~~ process needs/uses  
(CPU, disk, memory)

Today: CPU ~ resource  
scheduling

Resources: Not all alike

### Preemptible

Can take resource away, use for something else  
give back later

e.g. CPU

### Non-preemptible

Once given away, can't really take back  
(until volunteered)

e.g. Disk block, terminal

Key: Given resource/demands ⇒ determine how to manage it

create:  
2 styles:  $V_{fork}/VMS \Rightarrow$   
Spawn: NT  
create process  
(code: fork(), exec())  
e.g. int c=3  
i proc ↓ int rc=fork();  
if (rc < 0) // err  
else if (rc == 0) // child  
exec()  
else // parent  
wait()

⇒ flexible

cat > file  
close(stdout, FILENO)  
fd = open("file", O\_APPEND);  
pipes

⇒ can be restly  
huge copy before  
overlay

# Decisions about Resources

(2)

Allocation: which process gets which resources?

Space Sharing: <sup>[admission control]</sup> Divide up resource, give some to each  
when resources not easily preemptible

e.g. File space

Scheduling: how long process keeps resource

(best order)

Time sharing: more resources can be requested than available

Resource must be pre-emptible

e.g. CPU scheduling

## Role of Dispatcher v. Role of Sched

Dispatcher: low-level mechanism

must ctxt switch

{ save state in PCB for old  
load state of new }

Scheduler: higher-level policy  
which to run

could have both  
space + time  
sharing of CPU  
in parallel  
system,  
memory

# Performance Metrics

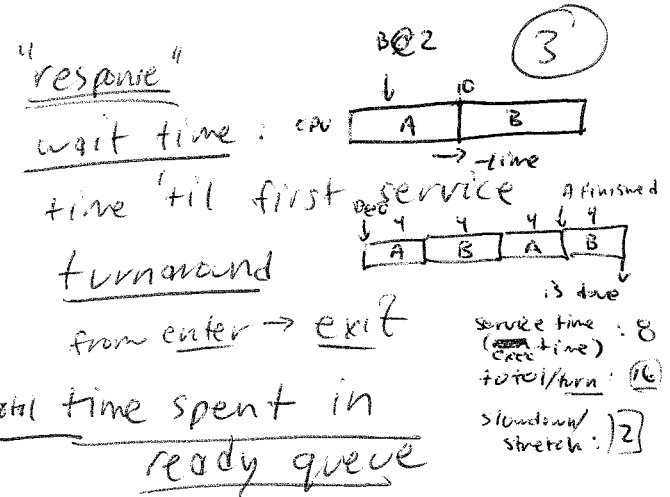
> Minimize <sup>cumulative</sup> waiting time: <sup>ready, not running</sup> (response time)

> Minimize turn-around time

> Minimize "cumulative wait" → total time spent in ready queue

> Minimize overhead  
↓ ctxt switches (costly?)

> <sup>maximize</sup> Fairness  
each process ⇒ same % of CPU



## When does Sched make decisions?

Minimal: Non-preemptive

- 1) Process blocks on I/O
  - 2) Process terminates
- ⇒ (remain scheduled until relinquished)

Also: Pre-emptive

- ⊕ (I/O interrupt)
- 1) event completes: blocked → ready
  - 2) timer interrupts
- ⇒ can run what it wants to

# Algorithms: Overview

① Process: CPU activity, I/O activity interleaved

if CPU bound: long CPU bursts }  
 I/O bound: short CPU bursts } ② might view "job" as □ □ □ □

Best Algorithm depends on workload, environment

specialized: only certain kinds of jobs, much knowledge about them

general purpose: want to do well for all types

Projects                      Book                      Discussion

## FCFS

- > Simplest algorithm:
- > Not preemptive
- > How to implement?  
 FIFO queue  
 => simple
- > why bad?  
 wait/service turnaround time depends on order of arrival  
 => unfair to later jobs

e.g. 3 jobs (A, B, C) arrive @ same time

Gantt chart

A	B	C
---	---	---

10/04 what worked well here:  
 simple 2 job example  
 A: 100  
 B: 20  
 use throughput

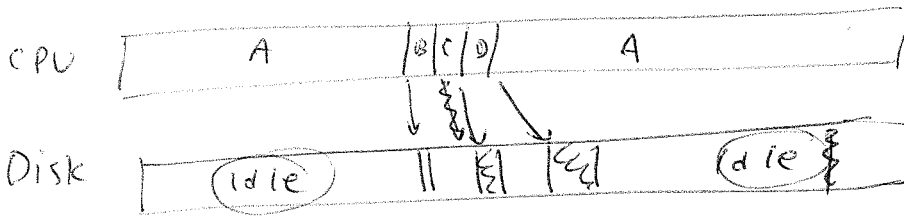
Avg <sup>cum</sup> wait time:  $(0 + 24 + 27) / 3 \Rightarrow 17$

Turnaround:  $\frac{24 + 27 + 30}{3} \Rightarrow 27$

Note: in multiprogrammed world, can put job @ end of Q upon I/O

⇒ Instance of Convoy Effect  
(Slow truck in front of fast cars)

⑤



A: CPU bound

B, C, D: I/O bound

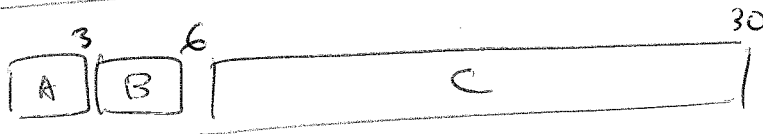
general problems

⇒ can also reduce I/O device utilization

⇒ the short job waiting time

SJF (Shortest Job First)

⇒ Should minimize wait/turnaround time



Avg wait:  $(0 + 3 + 6) / 3 = \boxed{3}$

Avg Turn Around:  $\frac{3 + 6 + 30}{3} = \frac{39}{3} = 13$

⇒ if  $T_A = T_B = T_C \Rightarrow$  FCFS

> Provably optimal (w/o preemption)

10/01

(supermarket example)

insight: moving short jobs up improves their standing A LOT while only hurting the long jobs A LITTLE

> But, not <sup>often</sup> practical

How to know job length?

(maybe can use past ⇒ predict future)

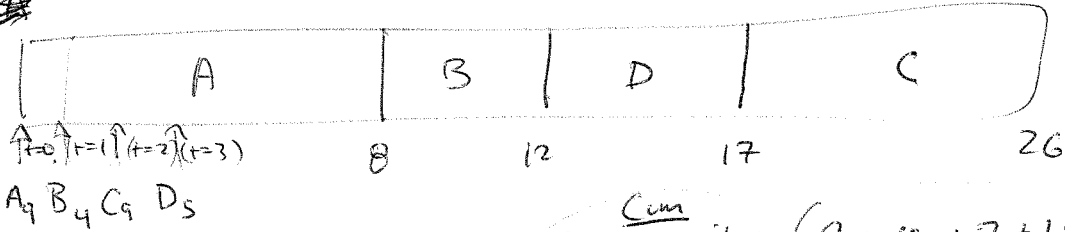
# STCF (Shortest Time to Completion first)

(6)

STCF = SJF w/ preemption

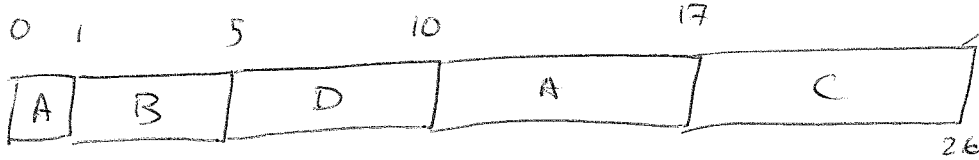
Problem: SJF is fine, IF all jobs are in system  
BUT, some arrive later!

SJF



Avg Cum wait:  $(0 + 7 + 9 + 15) / 4 = 7.75$

w/ preemption!



Avg <sup>Cum</sup> wait:  $(9 + 0 + 2 + 15) / 4 = 6.5$

10/04  
 What worked well:  
 let them figure it out  
 A: 100, B: 20  
 A has run for:  
 { 25 }  
 { 95 }

Note: total time is always same  
 (no magic here)

## RR

Avg Cum wait

- ⇒ more practical time sharing approach
- ⇒ run for time slice, then back of FIFO Q
- ⇒ preempted if still running @ end

## Advantages

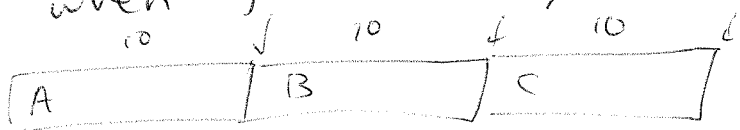
Fair, low average wait if job lengths vary widely

- > good interactive perf  
 (low response time)  
 ⇒ if N jobs time slice T ms  
 longest wait for some?  $(N-1)T$  avg  $\frac{(N-1)T}{2}$

But, RR has problems

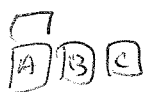
7

Poor when jobs nearly identical



Even FCFS is better!

$$\text{Avg turn around time} = (10 + 20 + 30) / 3 = \underline{\underline{20}}$$



$$= (28 + 29 + 30) / 3 = \underline{\underline{29}}$$

But, avg wait/response time is good!

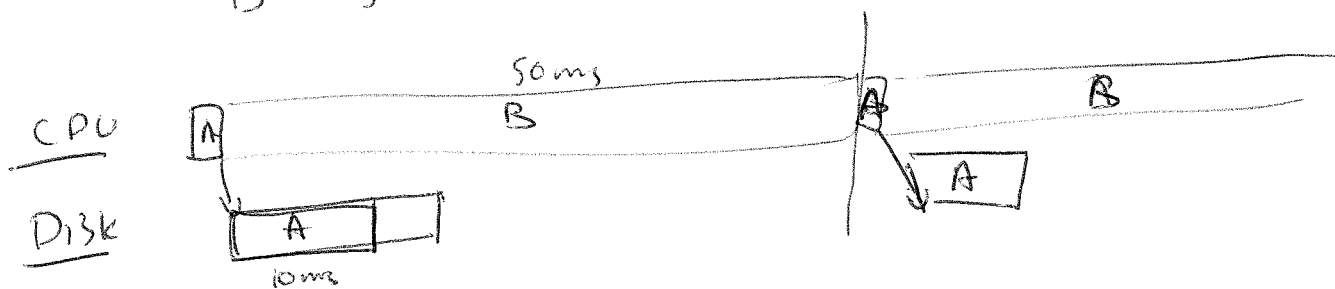
$$(0 + 1 + 2) / 3 \Rightarrow \underline{\underline{1}}$$

more interactive

$$\text{FCFS} : (0 + 10 + 20) / 3 \Rightarrow \underline{\underline{10}}$$

Problem #2: Perf depends on length of time slice  
if too high  $\Rightarrow$  FCFS

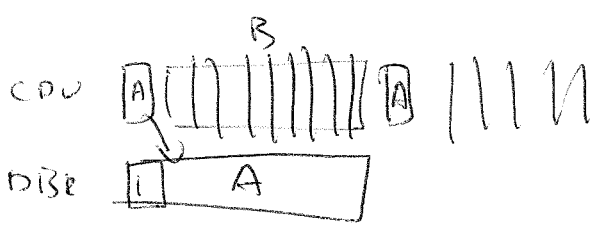
e.g. A 1 ms compute, 10 ms I/O  
B just compute



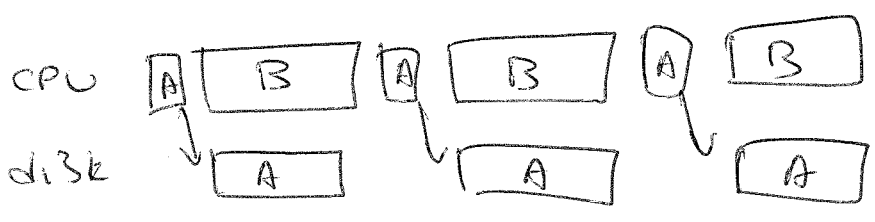
RR cont

> if time-slice too low, ctxt switch overhead

TS = 1ms



But STCF is still better



> How to approximate? => ~~Priorities~~ Dynamic Adaptive w/ priorities

---

> Each process has priority  
 run highest, RR among equal