

CS 241 Section Week #5  
(02/26/09)

Topics This Section

- SMP #3
- SMP #4
- 5-state Model
- Review of Scheduling
- Problems

SMP #3

SMP #3

- Multiple Clients / One Server
  - In SMP #3, the general idea was a client/server model where one server interacted with N clients.
  - Major Problems we saw:
    - Clients must run simultaneously
      - Launch all threads in one for loop
      - Join all threads in another for loop
      - DONT do `_create()` and `_join()` in the same loop!
    - Semaphores must be initialized to the correct value

## SMP #4

### SMP4 Forward

In SMP4, you will add code to a simulator for a CPU scheduler.

- ▶ We provide you with the code for the simulator.
  - ▶ You don't need to understand this code to understand this MP.
  - ▶ You should consider the simulator a 'black box'
- ▶ You need to implement these algorithms:  
fcfs, sjf, psjf, pri, ppri, rr#

### SMP4 Forward

- ▶ You need to fill in 3 scheduling functions:
  - ▶ new\_job()
  - ▶ job\_finished()
  - ▶ quantum\_expired()

Note that these are the only times that the scheduler needs to make a decision!
- ▶ A clean\_up() function to clean up any memory your program may've allocated
- ▶ A show\_queue() function to help you debug your program
- ▶ You need to create your own job queue

### SMP4 Forward

- ▶ You also need to fill in 3 statistics functions:
  - ▶ float average\_response\_time()
  - ▶ float average\_wait\_time()
  - ▶ float average\_turnaround\_time()

These are called at the end of the simulation.

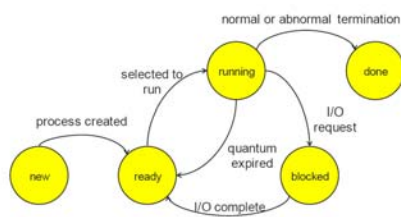
## SMP4 Forward

- ▶ For success on this MP:
  - ▶ Carefully read README.txt for details!
  - ▶ Look at the example runs and compare your results (e.g. using 'diff')!
- ▶ This MP is harder than all previous MPs!!
- ▶ Requires a good understanding of data structures, scheduling, and pointers all in one MP!

Good luck!

## Five State Model

## 5-State Model - Transitions



## Review of Scheduling

## Scheduling

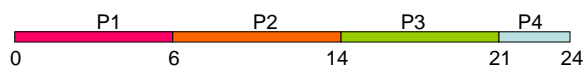
- ▶ The CPU Scheduler decides which thread should be in the running state. It is called when:
  - ▶ A thread is created or finishes
  - ▶ A clock interrupt occurs
  - ▶ An I/O interrupt occurs
  - ▶ A thread yields

## Scheduling

- ▶ The algorithms that we usually talk about are:
  - ▶ First-Come First-Serve (FCFS)
  - ▶ Shortest Job First (SJF)
  - ▶ Priority
  - ▶ Round Robin (RR)

### FCFS Example

Process	Duration	Priority	Arrival Time
P1	6	4	0
P2	8	1	0
P3	7	3	0
P4	3	2	0



### SJF Example

Process	Duration	Priority	Arrival Time
P1	6	4	0
P2	8	1	0
P3	7	3	0
P4	3	2	0



### Priority Example

Process	Duration	Priority	Arrival Time
P1	6	4	0
P2	8	1	0
P3	7	3	0
P4	3	2	0



### RR(1) Example

Process	Duration	Priority	Arrival Time
P1	6	4	0
P2	8	1	0
P3	7	3	0
P4	3	2	0

Quanta = 1 time unit



### Scheduling

- ▶ Scheduling algorithms can be preemptive or non-preemptive
  - ▶ **Non-preemptive**: each thread chooses when to yield to the processor (e.g. when done or system call)
  - ▶ **Preemptive**: scheduler forces the thread to yield (e.g. time quantum expires in RR)

### Scheduling

- ▶ Metrics for a single job
  - ▶ **Response time** = time from job submission until it's running
  - ▶ **Waiting time** = total time that the job is not running but queued
  - ▶ **Turnaround time** = time b/t the job's entry and completion

## Problems

### Problem #1

Job	Duration	Priority #
J1	6	2
J2	4	1
J3	5	1

These three jobs are going to arrive at our scheduler 1 time unit apart from each other (i.e. one job at time 0, one at time 1, and one at time 2), but the order hasn't been decided yet.

### Problem #1

We want to guarantee that the jobs finish in the order  
**J1 then J2 then J3**

### Problem #1

Which arrival order(s) guarantee this if the scheduler uses:

- 1) FCFS?
- 2) non-preemptive SJF?
- 3) preemptive SJF? (use remaining time, and ties are broken by arrival time)
- 4) RR-1? (arriving jobs are placed on ready queue immediately)
- 5) non-preemptive priority?
- 6) preemptive priority?

## Problem #1 Solution

Job	Time	Priority
1	6	2
2	4	1
3	5	1

- FCFS: 1 2 3
- n-p-SJF: 1 2 3 (1 needs to arrive first, then 2 will beat 3)
- p-SJF: 1 3 2 (1 needs to arrive first, and tie breaks will give it control; then 2 beats 3)
- RR-1: 1 3 2 (1 needs to run the longest, then 3, then 2, and it barely works out)
- n-p-Prio: 1 2 3 (1 must come first because of low prio, and then 2 and 3 must follow in order)
- p-Prio: none (2 or 3 will always preempt 1)

## Problem #2

For the SJF and RR examples, calculate:

- 1) Average response time
- 2) Average waiting time
- 3) Average turnaround time

Are either of these clearly better?

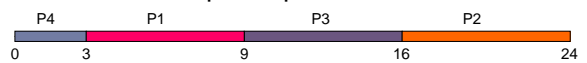
When would you use each?

## SJF Example

Process	Duration	Priority	Arrival Time
P1	6	4	0
P2	8	1	0
P3	7	3	0
P4	3	2	0



## Metrics for Non-preemptive: Shortest Job First



P4 waiting time: 0  
P1 waiting time: 3  
P3 waiting time: 9  
P2 waiting time: 16

Average response time (ART):  
 $(0+3+9+16)/4 = 7$   
 Average waiting time (AWT):  
 $(0+3+9+16)/4 = 7$   
 Average turnaround time (ATT):  
 $(3+9+16+24)/4 = 13$

P4 turnaround time: 3  
P1 turnaround time: 9  
P3 turnaround time: 16  
P2 turnaround time: 24

## RR(1) Example

Process	Duration	Priority	Arrival Time
P1	6	4	0
P2	8	1	0
P3	7	3	0
P4	3	2	0

Quanta = 1 time unit



## Metrics for Round Robin Example



P1 waiting time: 13  
P2 waiting time: 16  
P3 waiting time: 16  
P4 waiting time: 9

Average response time (ART):

$$(0 + 1 + 2 + 3) / 4 = 1.5$$

Average waiting time (AWT):

$$(13 + 16 + 16 + 9) / 4 = 13.5$$

Average turnaround time (ATT):

$$(12 + 24 + 23 + 12) / 4 = 17.75$$

P1 turnaround time: 19  
P2 turnaround time: 24  
P3 turnaround time: 23  
P4 turnaround time: 12

(Stephen Kloder CS241 Spring 2007)