

---

# Cyber-Physical Systems

## Scheduling

---



ICEN 553/453 – Fall 2018

Prof. Dola Saha

# Quick Recap

---

1. What characterizes the *memory architecture* of a system?
2. What are the issues with *heaps* in embedded/real-time systems?
3. How do *polling* and *interrupts* compare?
4. What is the difference between *concurrency* and *parallelism*?
5. What are *threads*? What makes them challenging wrt. concurrency?

# Scheduler

---

- A scheduler makes the decision about **what to do next at certain points in time**
- When a processor becomes available, which process will be executed

# Scheduler Policy

---

- Different schedulers will have different goals
  - Maximize throughput
  - Minimize latency
  - Prevent indefinite postponement
  - Complete process by given deadline
  - Maximize processor utilization

# Scheduler Levels

---

## ➤ High-level scheduling

- Determines which jobs can compete for resources
- Controls number of processes in system at one time

## ➤ Intermediate-level scheduling

- Determines which processes can compete for processors
- Responds to fluctuations in system load

## ➤ Low-level scheduling

- Assigns priorities
- Assigns processors to processes

# Priorities

---

## ➤ Static priorities

- Priority assigned to a process does not change
- Easy to implement
- Low overhead
- Not responsive to changes in environment

## ➤ Dynamic priorities

- Responsive to change
- Promote smooth interactivity
- Incur more overhead, justified by increased responsiveness

# How to decide which thread to schedule?

---

## ➤ Considerations:

- Preemptive vs. non-preemptive scheduling
- Periodic vs. aperiodic tasks
- Fixed priority vs. dynamic priority
- Priority inversion anomalies
- Other scheduling anomalies

# Non-Preemptive vs Preemptive

---

## ➤ Non-Preemptive

- Once a process is in the running state, it will continue until it terminates or blocks itself for I/O

## ➤ Preemptive

- Currently running process may be interrupted and moved to ready state by the OS
- Decision to preempt may be performed
  - when a new process arrives,
  - when an interrupt occurs that places a blocked process in the Ready state, or
  - periodically, based on a clock interrupt




# Preemptive Scheduling

---

- Assume all threads have priorities
  - either statically assigned (constant for the duration of the thread)
  - or dynamically assigned (can vary).
- Assume that the kernel keeps track of which threads are “enabled”
- Preemptive scheduling:
  - At any instant, the enabled thread with the highest priority is executing.
  - Whenever any thread changes priority or enabled status, the kernel can dispatch a new thread.

# Periodic scheduling



- Each execution instance of a task is called a **job**.
- For periodic scheduling, the best that we can do is to design an algorithm which will always find a schedule if one exists.
- A scheduler is defined to be **optimal** iff it will find a schedule if one exists.


# Scheduling Policies

---

- First Come First Serve
- Round Robin
- Shortest Process Next
- Shortest Remaining Time Next
- Highest Response Ratio Next
- Feedback Scheduler
- Fair Share Scheduler

# First Come First Serve (FCFS)

- Processes dispatched according to arrival time
- Simplest scheme
- Nonpreemptible
- Rarely used as primary scheduling algorithm
- Implemented using FIFO
- Tends to favor processor-bound processes over I/O-bound processes



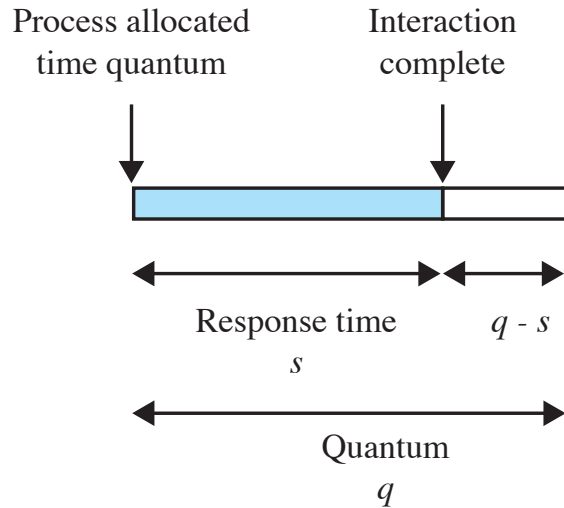
# Round Robin

---

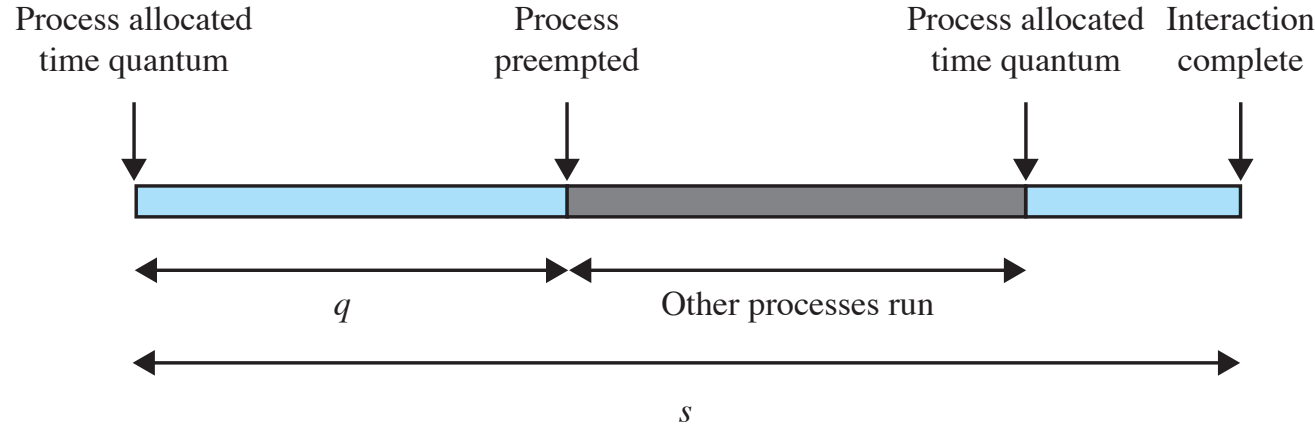
- Based on FIFO
- Processes run only for a limited amount of time called a time slice or a quantum
- Preemptible
- Requires the system to maintain several processes in memory to minimize overhead
- Often used as part of more complex algorithms

# Effect of Quantum Size

Time



$q < \text{Typical Interaction Time}$



$q > \text{Typical Interaction Time}$

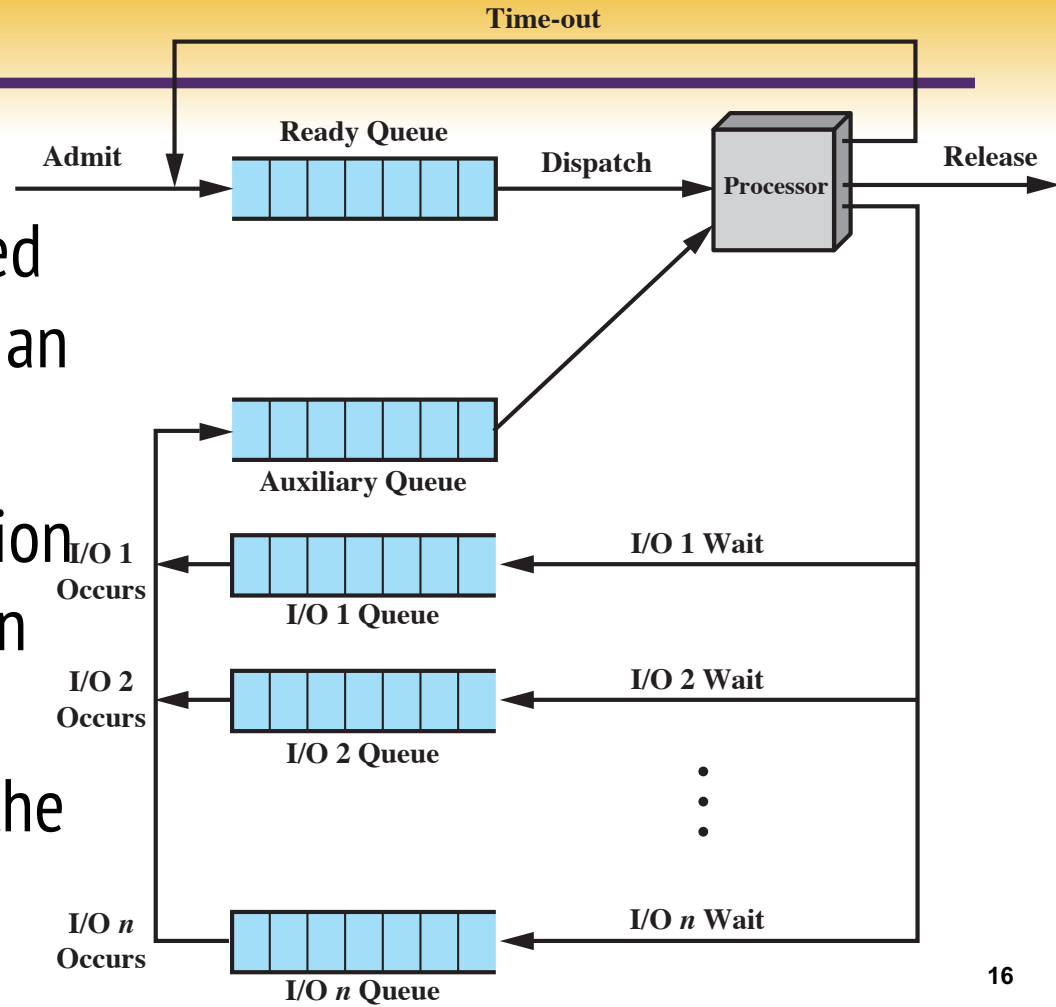
# Quantum Size

---

- Determines response time to interactive requests
- Very large quantum size
  - Processes run for long periods
  - Degenerates to FIFO
- Very small quantum size
  - System spends more time context switching than running processes
- Middle-ground
  - Long enough for interactive processes to issue I/O request
  - Batch processes still get majority of processor time

# Virtual Round Robin

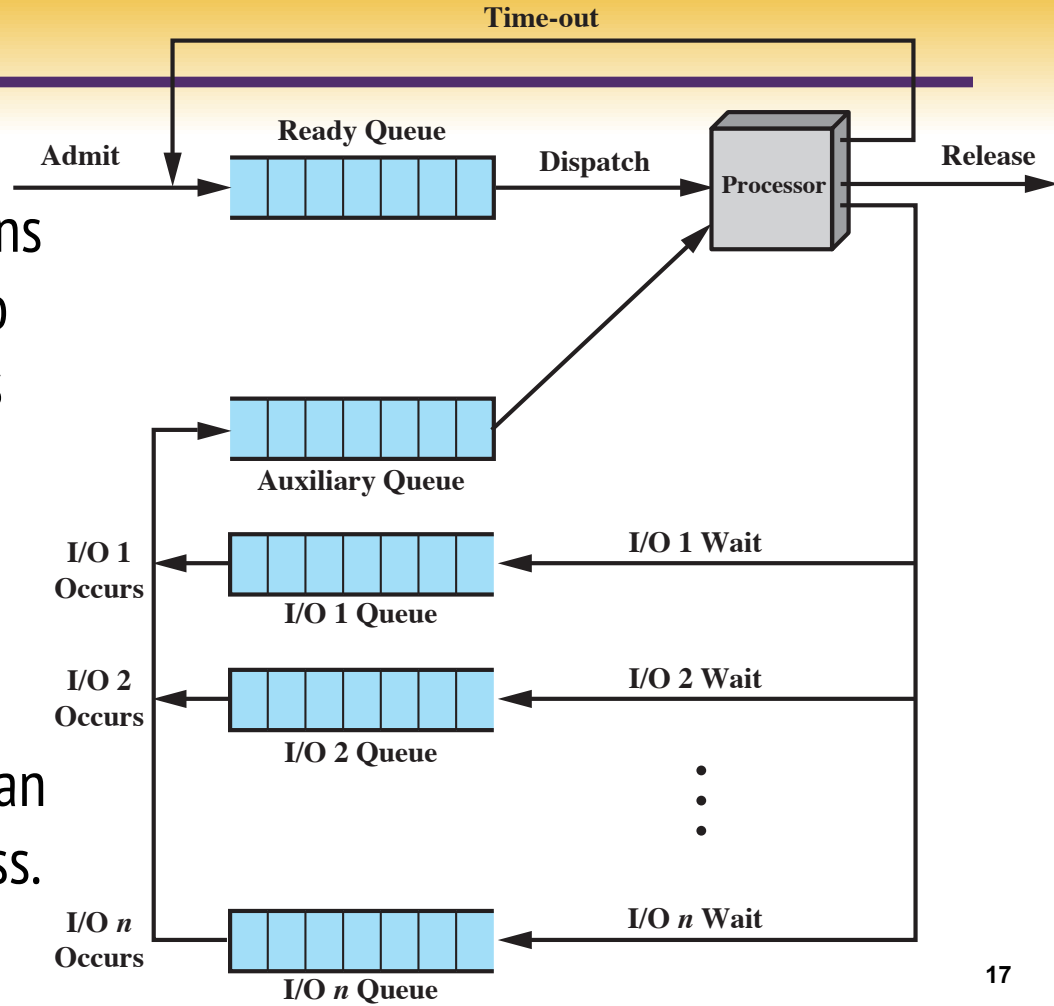
- FCFS auxiliary queue to which processes are moved after being released from an I/O block.
- When a dispatching decision is to be made, processes in the auxiliary queue get preference over those in the main ready queue.





# Virtual Round Robin

- When a process is dispatched from the auxiliary queue, it runs no longer than a time equal to the basic time quantum minus the total time spent running since it was last selected from the main ready queue.
- Performance studies indicate that this approach is better than round robin in terms of fairness.



# Shortest Process Next (SPN) Scheduling

---

- Scheduler selects process with smallest time to finish
  - Lower average wait time than FIFO
    - Reduces the number of waiting processes
  - Potentially large variance in wait times, starvation for longer processes
  - Nonpreemptive
    - Results in slow response times to arriving interactive requests
  - Relies on estimates of time-to-completion
    - Can be inaccurate
  - Unsuitable for use in modern interactive systems

# Shortest Remaining Time (SRT) Scheduling

---

- Preemptive version of SPF
- Shorter arriving processes preempt a running process
- Very large variance of response times: long processes wait even longer than under SPF
- Not always optimal
  - Short incoming process can preempt a running process that is near completion
  - Context-switching overhead can become significant

# Highest Response Ratio Next (HRRN) Scheduling


---

$$\text{Ratio} = \frac{\text{time spent waiting} + \text{expected service time}}{\text{expected service time}}$$

- Chooses next process with the greatest ratio
- Attractive because it accounts for the age of the process
- While shorter jobs are favored, aging without service increases the ratio so that a longer process will eventually get past competing shorter jobs

# Feedback Scheduling

- Scheduling is done on a preemptive (at time quantum) basis, and a dynamic priority mechanism is used.
- When a process first enters the system, it is placed in RQ0.
- After its first preemption, when it returns to the Ready state, it is placed in RQ1.
- Each subsequent time that it is preempted, it is demoted to the next lower-priority queue.



# Performance

---

- Any scheduling policy that chooses the next item to be served independent of service time obeys the relationship:

$$\frac{T_r}{T_s} = \frac{1}{1 - \rho}$$

where

$T_r$  = turnaround time or residence time; total time in system, waiting plus execution

$T_s$  = average service time; average time spent in Running state

$\rho$  = processor utilization

# Single Server Queue with Two Priorities

- Assumptions:
1. Poisson arrival rate.
  2. Priority 1 items are serviced before priority 2 items.
  3. First-come-first-served dispatching for items of equal priority.
  4. No item is interrupted while being served.
  5. No items leave the queue (lost calls delayed).

## (a) General formulas

$$\lambda = \lambda_1 + \lambda_2$$

$$\rho_1 = \lambda_1 T_{s1}; \quad \rho_2 = \lambda_2 T_{s2}$$

$$\rho = \rho_1 + \rho_2$$

$$T_s = \frac{\lambda_1}{\lambda} T_{s1} + \frac{\lambda_2}{\lambda} T_{s2}$$

$$T_r = \frac{\lambda_1}{\lambda} T_{r1} + \frac{\lambda_2}{\lambda} T_{r2}$$

# Single Server Queue with Two Priorities

**(b) No interrupts; exponential service times**

$$T_{r1} = T_{s1} + \frac{\rho_1 T_{s1} + \rho_2 T_{s2}}{1 - \rho_1}$$

$$T_{r2} = T_{s2} + \frac{T_{r1} - T_{s1}}{1 - \rho}$$

**(c) Preemptive-resume queuing discipline; exponential service times**

$$T_{r1} = T_{s1} + \frac{\rho_1 T_{s1}}{1 - \rho_1}$$

$$T_{r2} = T_{s2} + \frac{1}{1 - \rho_1} \left( \rho_1 T_{s2} + \frac{\rho T_s}{1 - \rho} \right)$$

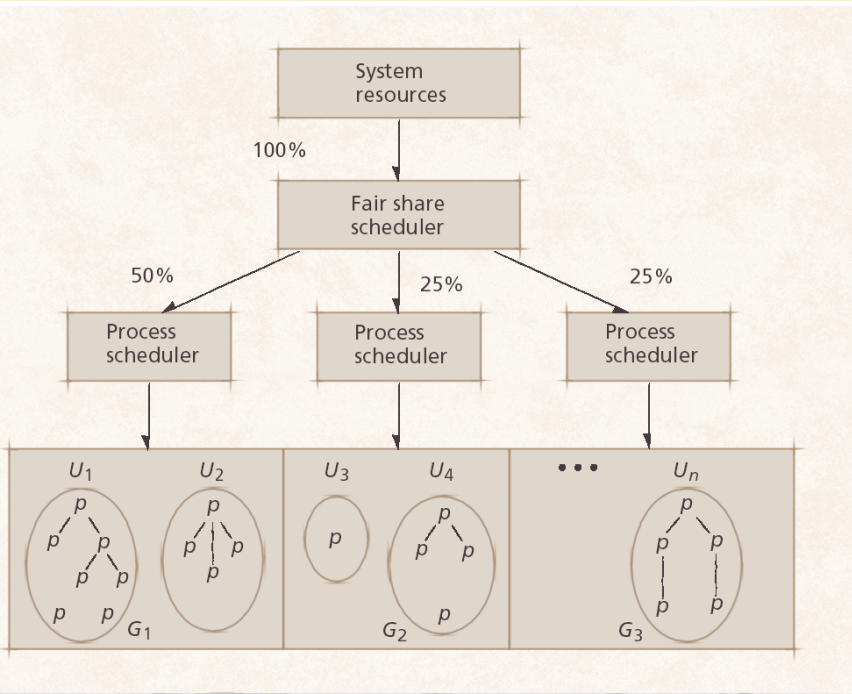


# Fair Share Scheduler

---

- Scheduling decisions based on the process sets
- Each user is assigned a share of the processor
- Objective is to monitor usage to give fewer resources to users who have had more than their fair share and more to those who have had less than their fair share
- Some user groups more important than others
- Ensures that less important groups cannot monopolize resources
- Unused resources distributed according to the proportion of resources each group has been allocated
- Groups not meeting resource-utilization goals get higher priority

# Fair Share



$$CPU_j(i) = \frac{CPU_j(i-1)}{2}$$

$$GCPU_k(i) = \frac{GCPU_k(i-1)}{2}$$

$$P_j(i) = Base_j + \frac{CPU_j(i)}{2} + \frac{GCPU_k(i)}{4 \times W_k}$$

where

$CPU_j(i)$  = measure of processor utilization by process  $j$  through interval  $i$ ,

$GCPU_k(i)$  = measure of processor utilization of group  $k$  through interval  $i$ ,

$P_j(i)$  = priority of process  $j$  at beginning of interval  $i$ ; lower values equal higher priorities,

$Base_j$  = base priority of process  $j$ , and

$W_k$  = weighting assigned to group  $k$ , with the constraint that and

$$0 < W_k \leq 1 \text{ and } \sum W_k = 1.$$

# Example

Time	Process A			Process B			Process C		
	Priority	Process CPU count	Group CPU count	Priority	Process CPU count	Group CPU count	Priority	Process CPU count	Group CPU count
	0	60	0	0	60	0	0	60	0
1		1	1						
		2	2						
		•	•						
		•	•						
		60	60						
2	90	30	30	60	0	0	60	0	0
					1	1			1
					2	2			2
					•	•			•
					•	•			•
3					60	60			60
	74	15	15	90	30	30	75	0	30
		16	16						
		17	17						
		•	•						
4		•	•						
		75	75						
	96	37	37	74	15	15	67	0	15
						16		1	16
						17		2	17
5						75		•	•
						•		•	•
						•		•	•
						75		60	75
6	78	18	18	81	7	37	93	30	37
		19	19						
		20	20						
		•	•						
		•	•						
7		78	78						
	98	39	39	70	3	18	76	15	18

Group 1

Group 2

# UNIX Scheduler

---

- Designed to provide good response time for interactive users while ensuring that low-priority background jobs do not starve
- Employs multilevel feedback using round robin within each of the priority queues
- Makes use of one-second preemption
- Priority is based on process type and execution history

# Scheduling Formula

---

$$CPU_j(i) = \frac{CPU_j(i-1)}{2}$$

$$P_j(i) = Base_j + \frac{CPU_j(i)}{2} + nice_j$$

where

$CPU_j(i)$  = measure of processor utilization by process  $j$  through interval  $i$

$P_j(i)$  = priority of process  $j$  at beginning of interval  $i$ ; lower values equal higher priorities

$Base_j$  = base priority of process  $j$

$nice_j$  = user-controllable adjustment factor



# Characteristics of Various Scheduling Policies

	FCFS	Round Robin	SPN	SRT	HRRN	Feedback
<b>Selection Function</b>	$\max[w]$	constant	$\min[s]$	$\min [e - e]$	$\max \left( \frac{w+s}{s} \right)$	(see text)
<b>Decision Mode</b>	Non-preemptive	Preemptive (at time quantum)	Non-preemptive	Preemptive (at arrival)	Non-preemptive	Preemptive (at time quantum)
<b>Throughput</b> ⓘ	Not emphasized	May be low if quantum is too small	High	High	High	Not emphasized
<b>Response Time</b>	May be high, especially if there is a large variance in process execution times	Provides good response time for short processes	Provides good response time for short processes	Provides good response time	Provides good response time	Not emphasized
<b>Overhead</b>	Minimum	Minimum	Can be high	Can be high	Can be high	Can be high
<b>Effect on Processes</b>	Penalizes short processes; penalizes I/O-bound processes	Fair treatment	Penalizes long processes	Penalizes long processes	Good balance	May favor I/O-bound processes
<b>Starvation</b>	No	No	Possible	Possible	No	Possible