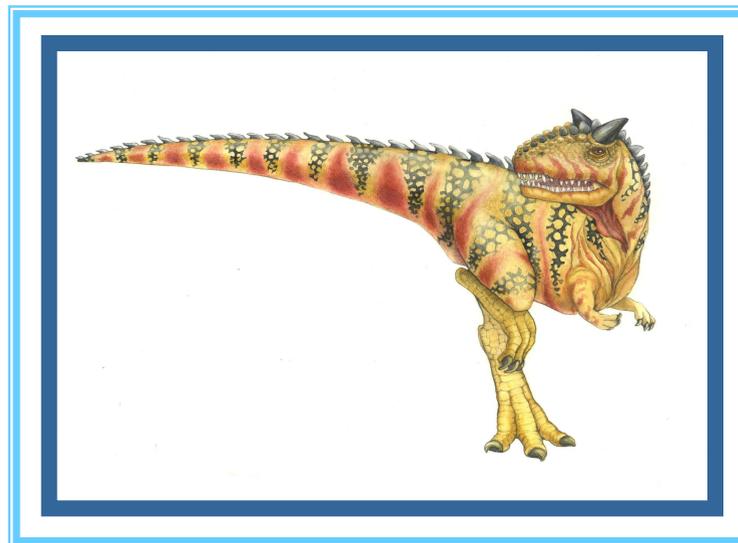
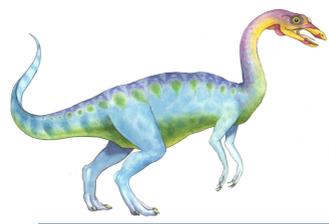


Chapter 5: CPU Scheduling

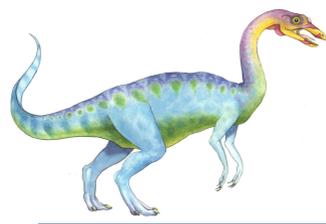




Chapter 5: CPU Scheduling

- Basic Concepts
- Scheduling Criteria
- Scheduling Algorithms
- Examples

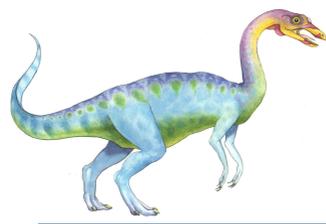




Objectives

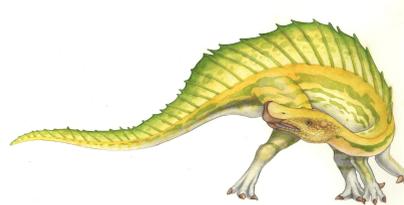
- To introduce CPU scheduling, which is the basis for multiprogrammed operating systems
- To describe various CPU-scheduling algorithms
- To discuss evaluation criteria for selecting a CPU-scheduling algorithm for a particular system

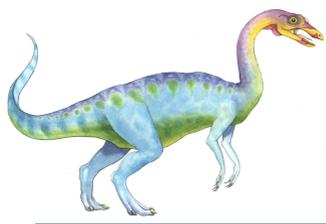




Basic Concepts

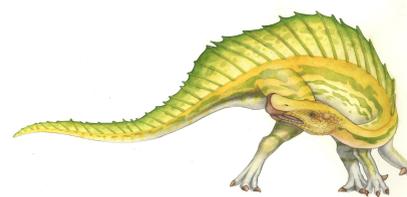
- Maximum CPU utilization obtained with multiprogramming
- **Continuous Cycle :**
 - one process has to wait (I/O)
 - Operating system takes the CPU away
 - Give CPU to another process
 - This pattern continues
- CPU–I/O Burst Cycle – Process execution consists of a *cycle* of CPU execution and I/O wait

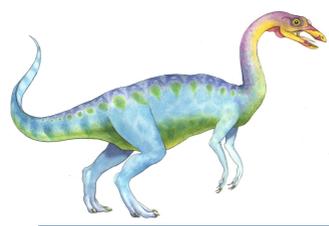




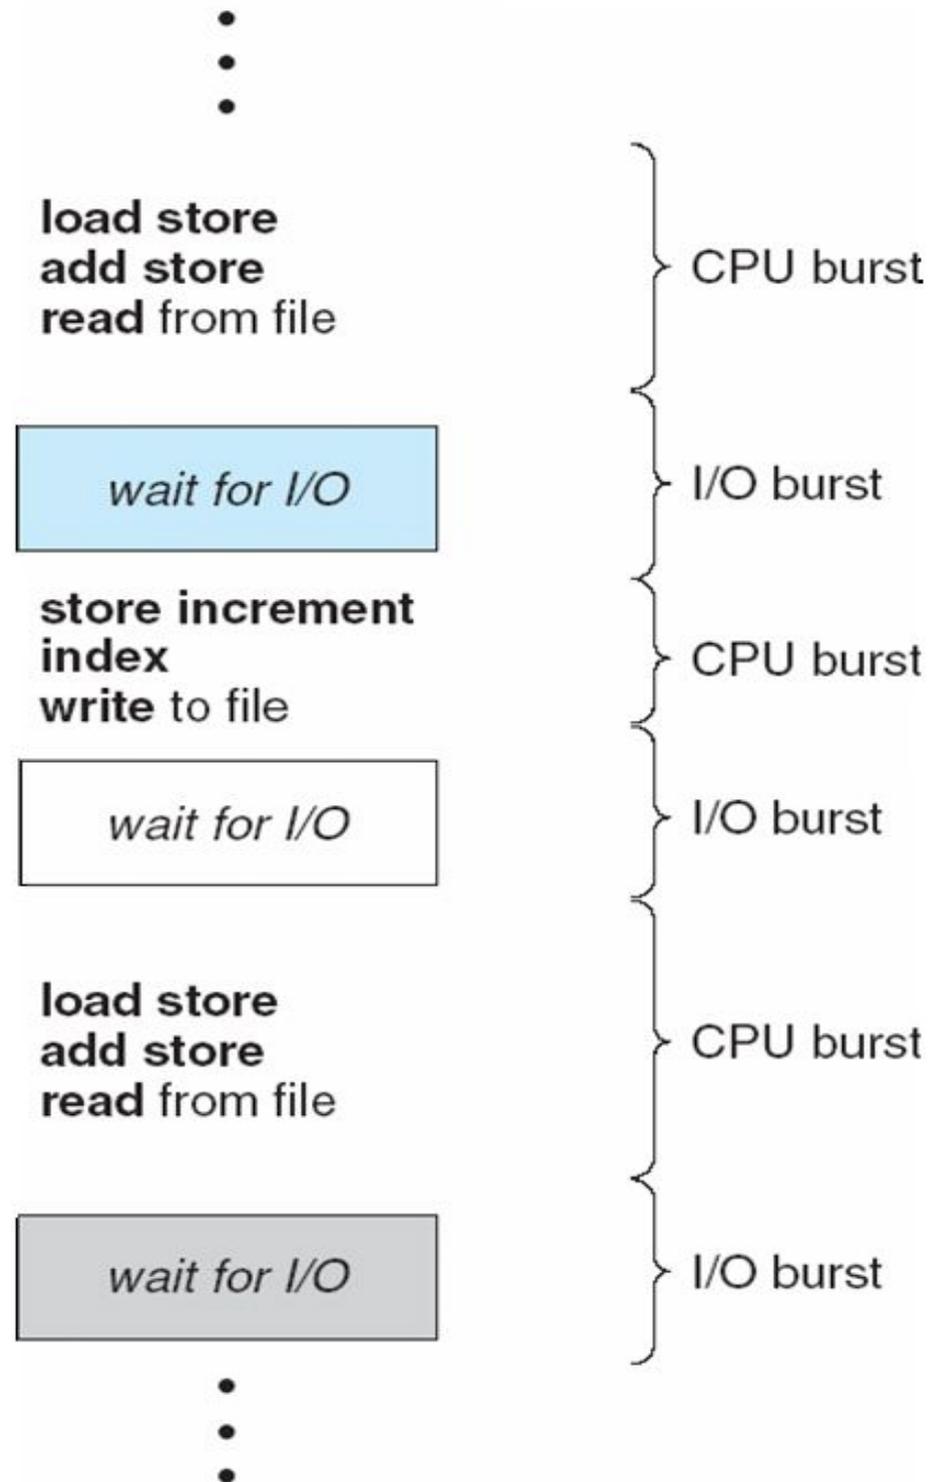
CPU and I/O Burst Cycle

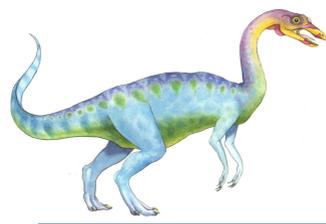
- Almost all processes **alternate between two states** in a continuing **cycle**, as shown in Figure below :
 - **A CPU burst** of performing calculations, and
 - **An I/O burst, waiting for data transfer in or out of the system.**
- Processes alternate back and forth between this two states.





Alternating Sequence of CPU and I/O Bursts

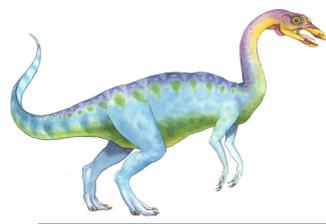




CPU Scheduler

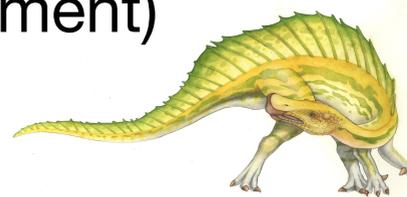
- Selects from the processes in **ready queue**, and allocates the CPU to one of them
 - FIFO queue
 - Priority queue
 - Tree
 - Unordered linked-list
- **CPU scheduling decisions** may take place when a process:
 1. **Switches from running to waiting state (I/O request)**
 2. **Switches from running to ready state (e.g. when interrupt occurs)**
 3. **Switches from waiting to ready (e.g. at completion of I/O)**
 4. **Terminates**
- Scheduling under 1 and 4 is **nonpreemptive**
- All other scheduling is **preemptive**
 - **Consider access to shared data**
 - Consider preemption while in kernel mode
 - Consider interrupts occurring during crucial OS activities

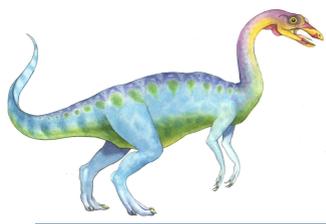




Scheduling Criteria

- **CPU utilization** – keep the CPU as busy as possible
- **Throughput** – # of processes that complete their execution per time unit
- **Turnaround time**
 - amount of time to execute a particular process
 - the interval from the time of submission of a process to the time of the completion.
 - sum of the periods spent waiting to get into memory, waiting in the ready queue, executing on the CPU, doing I/O
- **Waiting time** – amount of time a process has been waiting in the ready queue
- **Response time** – amount of time it takes from when a request was submitted until the first response is produced, not output (for time-sharing environment)

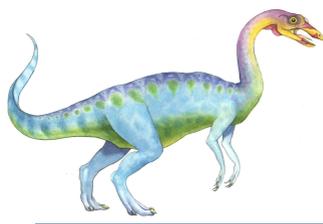




Scheduling Algorithm Optimization Criteria

- Max CPU utilization
- Max throughput
- Min turnaround time
- Min waiting time
- Min response time





First-Come, First-Served (FCFS) Scheduling

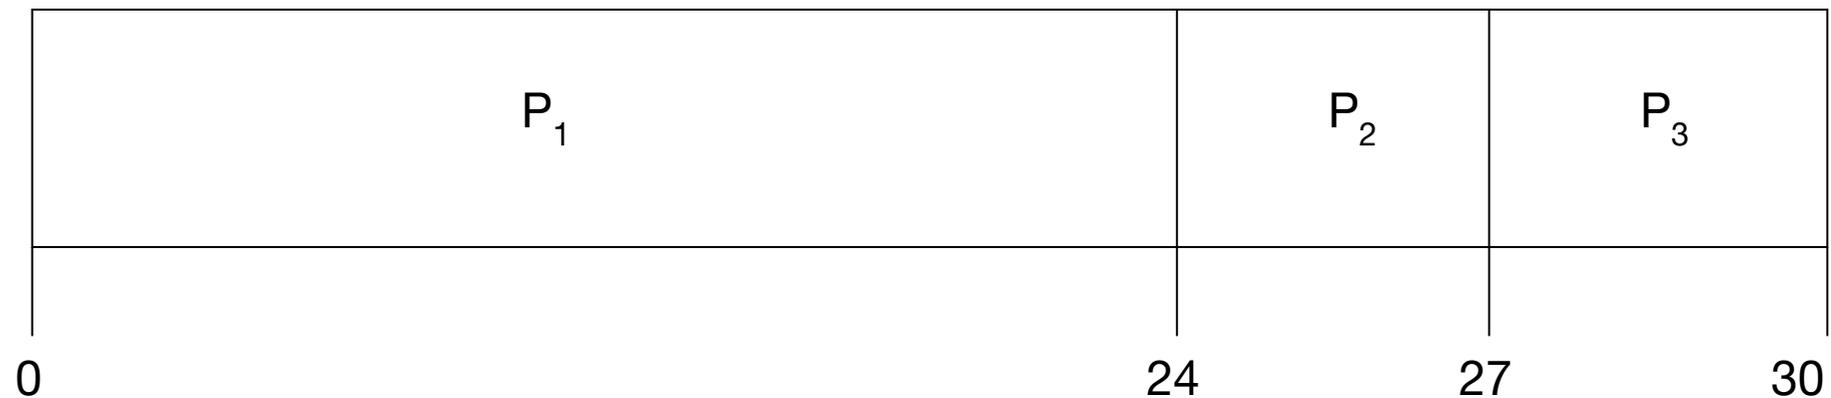
Process Burst Time

P_1 24

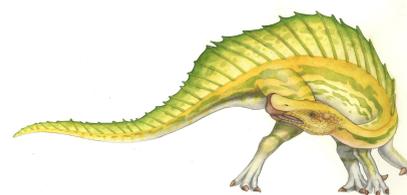
P_2 3

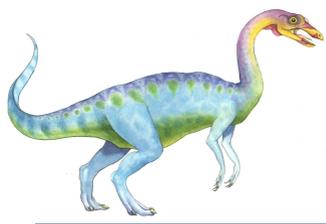
P_3 3

- Suppose that the processes arrive in the order: P_1, P_2, P_3
The Gantt Chart for the schedule is:



- Waiting time for $P_1 = 0$; $P_2 = 24$; $P_3 = 27$
- Average waiting time: $(0 + 24 + 27)/3 = 17$
- Turnaround time $P_1 = 24$; $P_2 = 27$; $P_3 = 30$





FCFS Scheduling (Cont.)

Suppose that the processes arrive in the order:

$$P_2, P_3, P_1$$

- The Gantt chart for the schedule is:



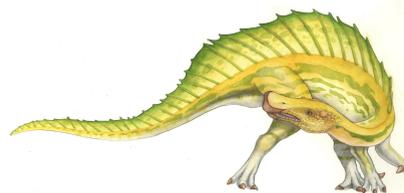
Waiting time for $P_1 = 6$; $P_2 = 0$; $P_3 = 3$

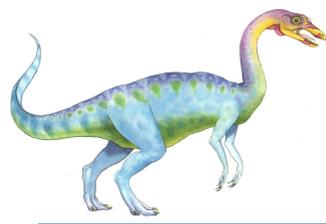
Average waiting time: $(6 + 0 + 3)/3 = 3$

Much better than previous case

Convoy effect - short process behind long process

Consider one CPU-bound and many I/O-bound processes



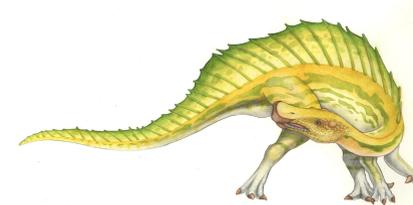


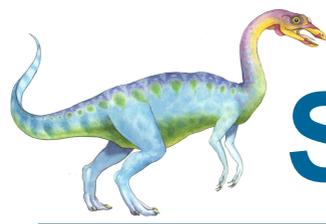
FCFS Scheduling (Cont.)

Convoy Effect :

many I/O bound process and one CPU bound process

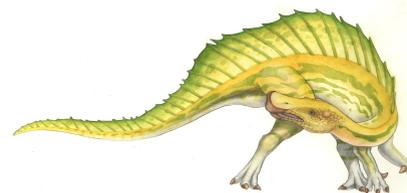
CPU bound process	I/O bound process	Effect
I/O device	I/O queue	CPU site idle
CPU processing	Ready queue	I/O site idle

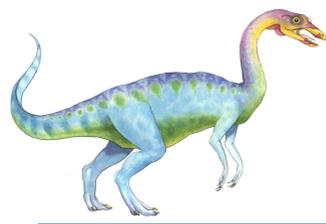




Shortest-Job-First (SJF) Scheduling

- Associate with each process the length of its next CPU burst. Use these lengths to schedule the process with the shortest time
- Two schemes:
 - **Non-preemptive** – once CPU given to the process it cannot be preempted until completes its CPU burst
 - **preemptive** – if a new process arrives with CPU burst length less than remaining time of current executing process, preempt. This scheme is known as the Shortest-Remaining-Time-First (SRTF)
- SJF is optimal – gives minimum average waiting time for a given set of processes

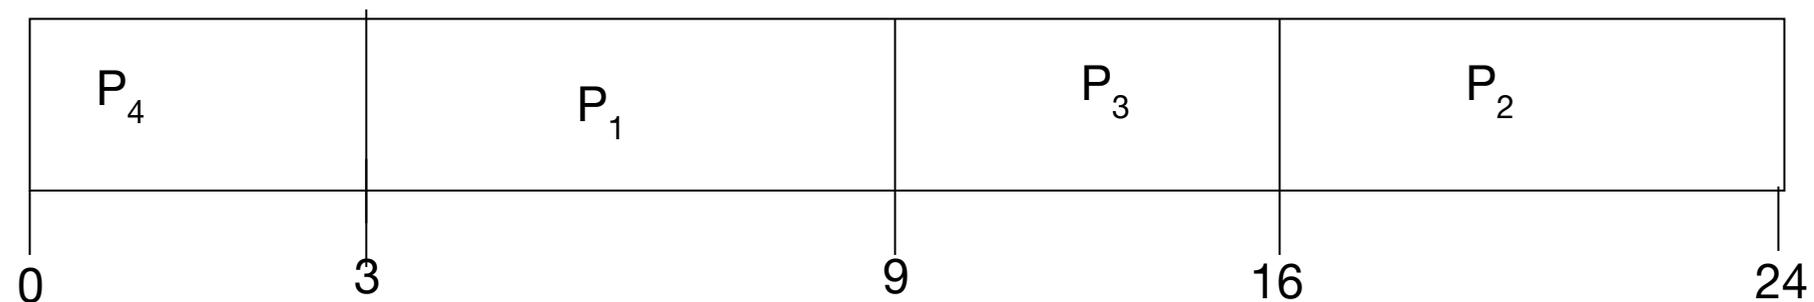




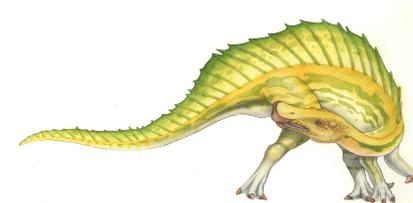
Example of SJF

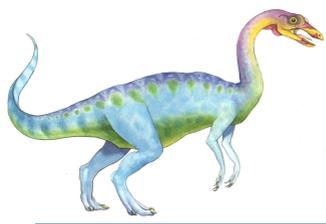
	<u>Process</u>	<u>Burst Time</u>
	P_1	6
	P_2	8
	P_3	7
	P_4	3

- SJF scheduling chart



$$\text{Average waiting time} = (3 + 16 + 9 + 0) / 4 = 7$$

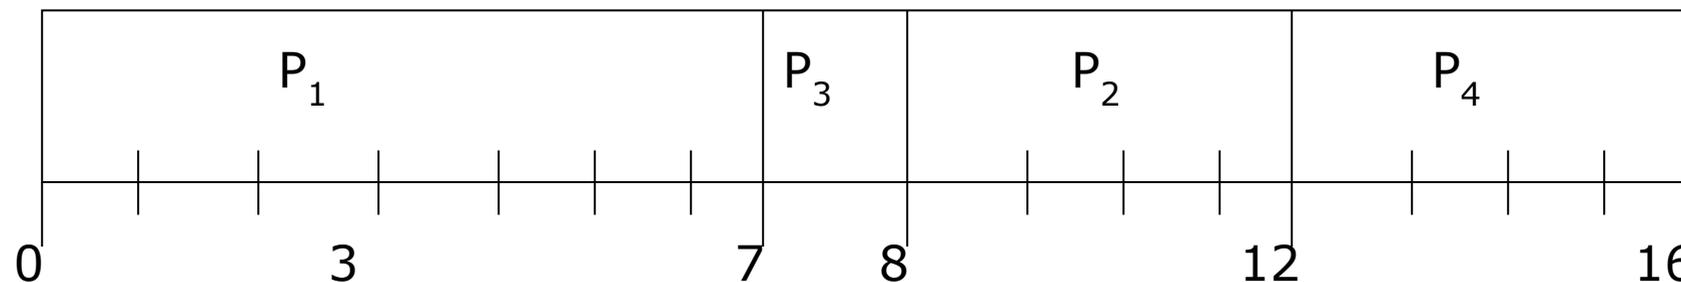




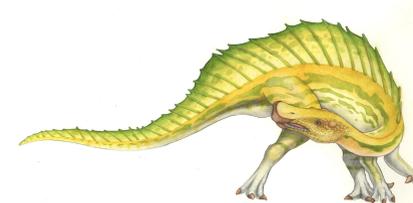
Example of Non-Preemptive SJF

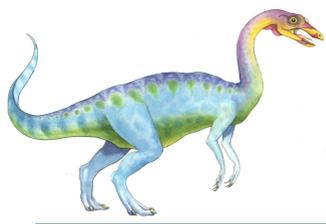
<u>Process</u>	<u>Arrival Time</u>	<u>Burst Time</u>
P_1	0.0	7
P_2	2.0	4
P_3	4.0	1
P_4	5.0	4

- SJF (non-preemptive)



$$\text{Average waiting time} = (0 + 6 + 3 + 7)/4 = 4$$

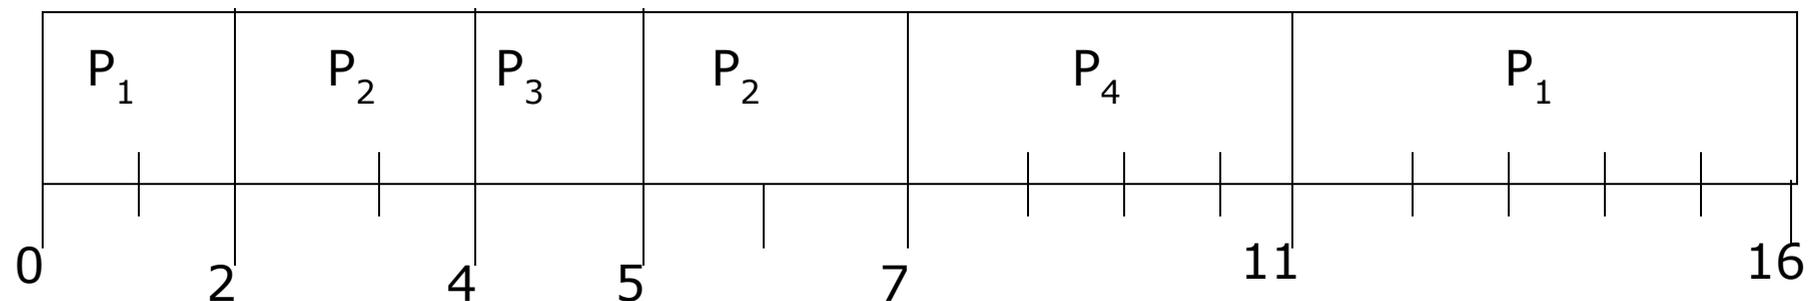




Example of Preemptive SJF

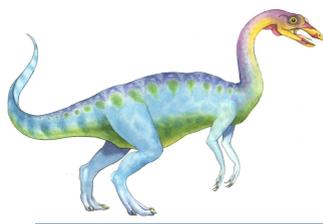
<u>Process</u>	<u>Arrival Time</u>	<u>Burst Time</u>
P_1	0.0	7
P_2	2.0	4
P_3	4.0	1
P_4	5.0	4

- SJF (preemptive)



- Average waiting time = $(9 + 1 + 0 + 2)/4 = 3$



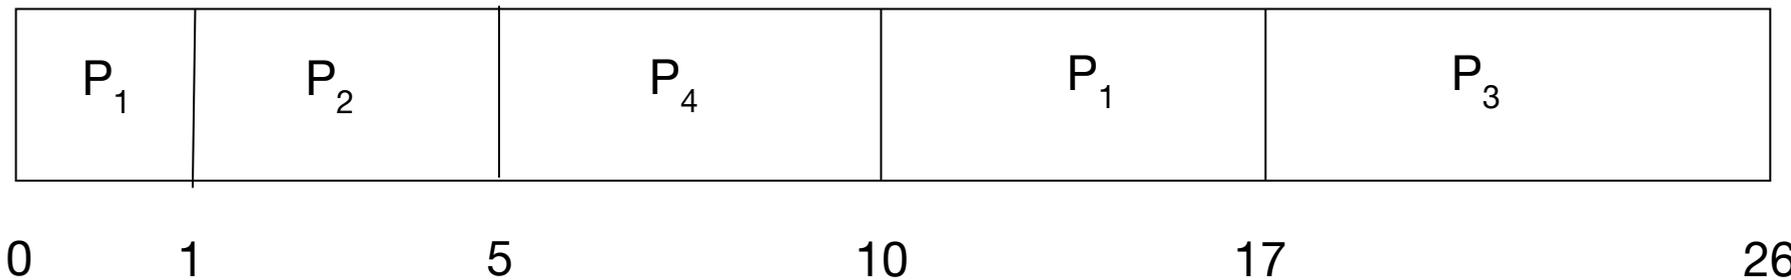


Example of Shortest-remaining-time-first

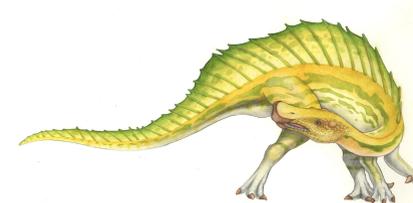
- Now we add the concepts of varying arrival times and preemption to the analysis

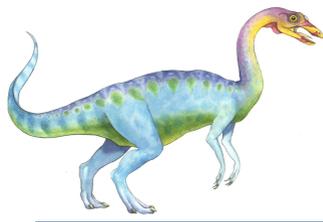
	<u>Process</u>	<u>Arrival Time</u>	<u>Burst Time</u>
	P_1	0	8
	P_2	1	4
	P_3	2	9
	P_4	3	5

- Preemptive SJF Gantt Chart



- Average waiting time = $[(10-1)+(1-1)+(17-2)+5-3]/4 = 26/4 = 6.5$ msec





- Exponential average
Page – 221/946

and pick the process with the shortest predicted CPU burst.

The next CPU burst is generally predicted as an exponential average of the measured lengths of previous CPU bursts. Let t_n be the length of the n th CPU

5.3 Scheduling Algorithms 161

burst, and let τ_{n+1} be our predicted value for the next CPU burst. Then, for α , $0 \leq \alpha \leq 1$, define

$$\tau_{n+1} = \alpha t_n + (1 - \alpha)\tau_n.$$

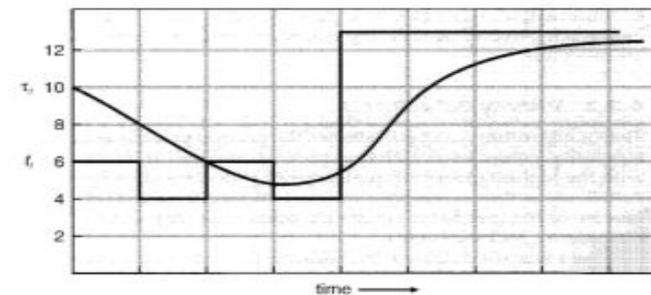
This formula defines an **exponential average**. The value of t_n contains our most recent information; τ_n stores the past history. The parameter α controls the relative weight of recent and past history in our prediction. If $\alpha = 0$, then $\tau_{n+1} = \tau_n$, and recent history has no effect (current conditions are assumed to be transient); if $\alpha = 1$, then $\tau_{n+1} = t_n$, and only the most recent CPU burst matters (history is assumed to be old and irrelevant). More commonly, $\alpha = 1/2$, so recent history and past history are equally weighted. The initial τ_0 can be defined as a constant or as an overall system average. Figure 5.3 shows an exponential average with $\alpha = 1/2$ and $\tau_0 = 10$.

To understand the behavior of the exponential average, we can expand the formula for τ_{n+1} by substituting for τ_n to find

$$\tau_{n+1} = \alpha t_n + (1 - \alpha)\alpha t_{n-1} + \dots + (1 - \alpha)^j \alpha t_{n-j} + \dots + (1 - \alpha)^{n+1} \tau_0.$$

Since both α and $(1 - \alpha)$ are less than or equal to 1, each successive term has less weight than its predecessor.

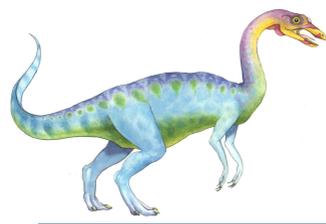
The SJF algorithm can be either preemptive or nonpreemptive. The choice arises when a new process arrives at the ready queue while a previous process is still executing. The next CPU burst of the newly arrived process may be shorter than what is left of the currently executing process. A preemptive SJF algorithm



CPU burst (t_i)	6	4	6	4	13	13	13	...
"guess" (τ_i)	10	8	6	6	5	9	11	12

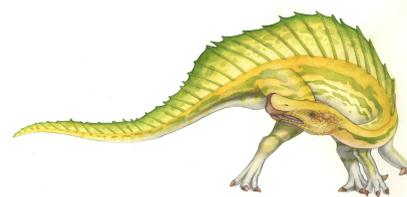
Figure 5.3 Prediction of the length of the next CPU burst.

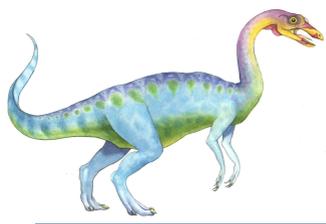




Priority Scheduling

- A priority number (integer) is associated with each process
- The CPU is allocated to the process with the highest priority (smallest integer \equiv highest priority)
 - Preemptive
 - Nonpreemptive
- SJF is priority scheduling where priority is the inverse of predicted next CPU burst time
- Priority can be defined either internally or externally.
 - Factors for internal priority assignment:
 - 4 Time limit, memory requirements, the number or open files etc.
 - Factors for external priority assignment:
 - 4 Importance of the process, the type and amount of funds of funds being paid for computer use, department sponsoring works etc.

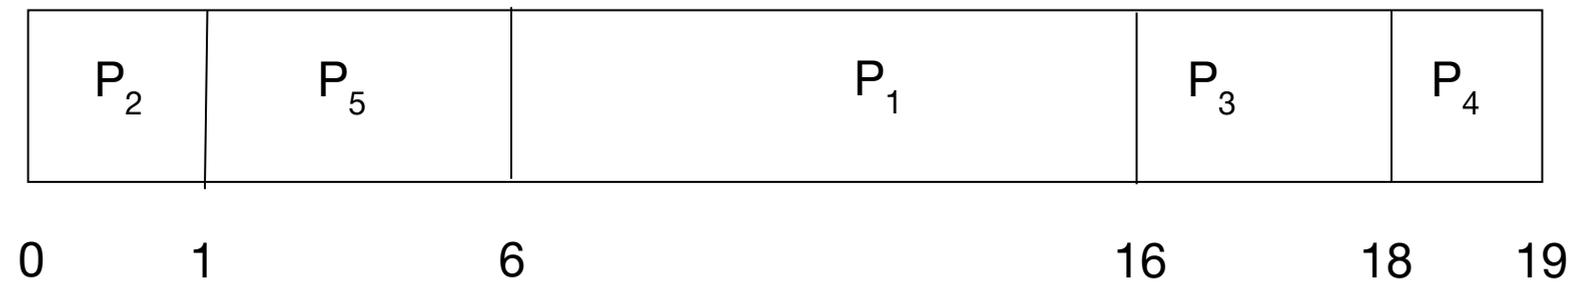




Example of Priority Scheduling

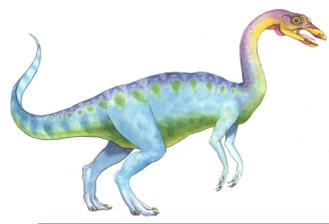
	<u>Process</u>	<u>Burst Time</u>	<u>Priority</u>
P_1	10	3	
P_2	1	1	
P_3	2	4	
P_4	1	5	
P_5	5	2	

- Priority scheduling Gantt Chart



Average waiting time = 8.2 msec

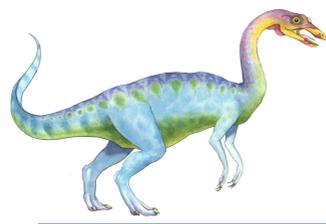




Priority Scheduling

- Problem \equiv **Starvation** – low priority processes may never execute
- Solution \equiv **Aging** – as time progresses increase the priority of the process

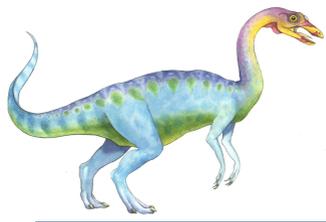




Round Robin (RR)

- Each process gets a small unit of CPU time (**time quantum q**), usually 10-100 milliseconds. After this time has elapsed, the process is preempted and added to the end of the ready queue.
- If there are n processes in the ready queue and the time quantum is q , then each process gets $1/n$ of the CPU time in chunks of at most q time units at once. No process waits more than $(n-1)q$ time units.
- Timer interrupts every quantum to schedule next process
- Performance
 - q large \Rightarrow FCFS
 - q small $\Rightarrow q$ must be large with respect to context switch, otherwise overhead is too high





Example of RR with Time Quantum = 4

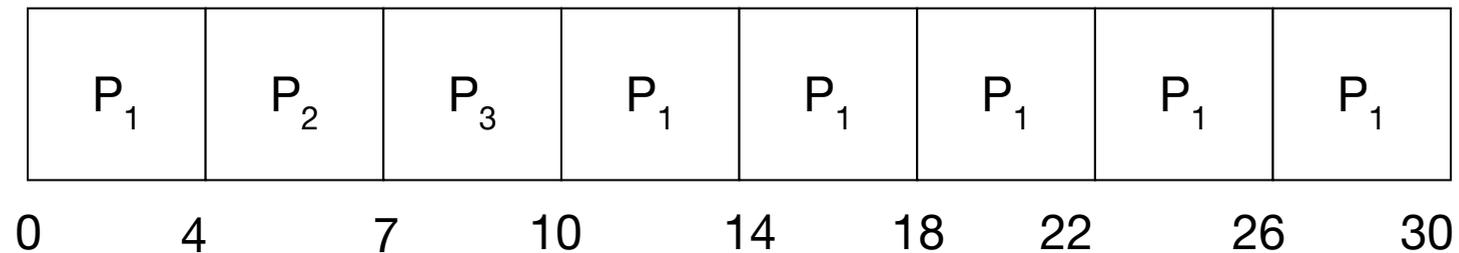
Process Burst Time

P_1 24

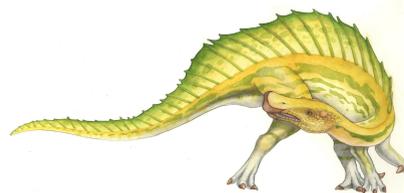
P_2 3

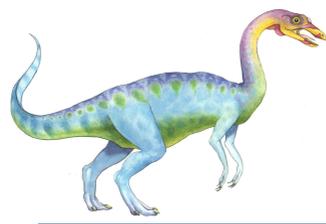
P_3 3

- The Gantt chart is:

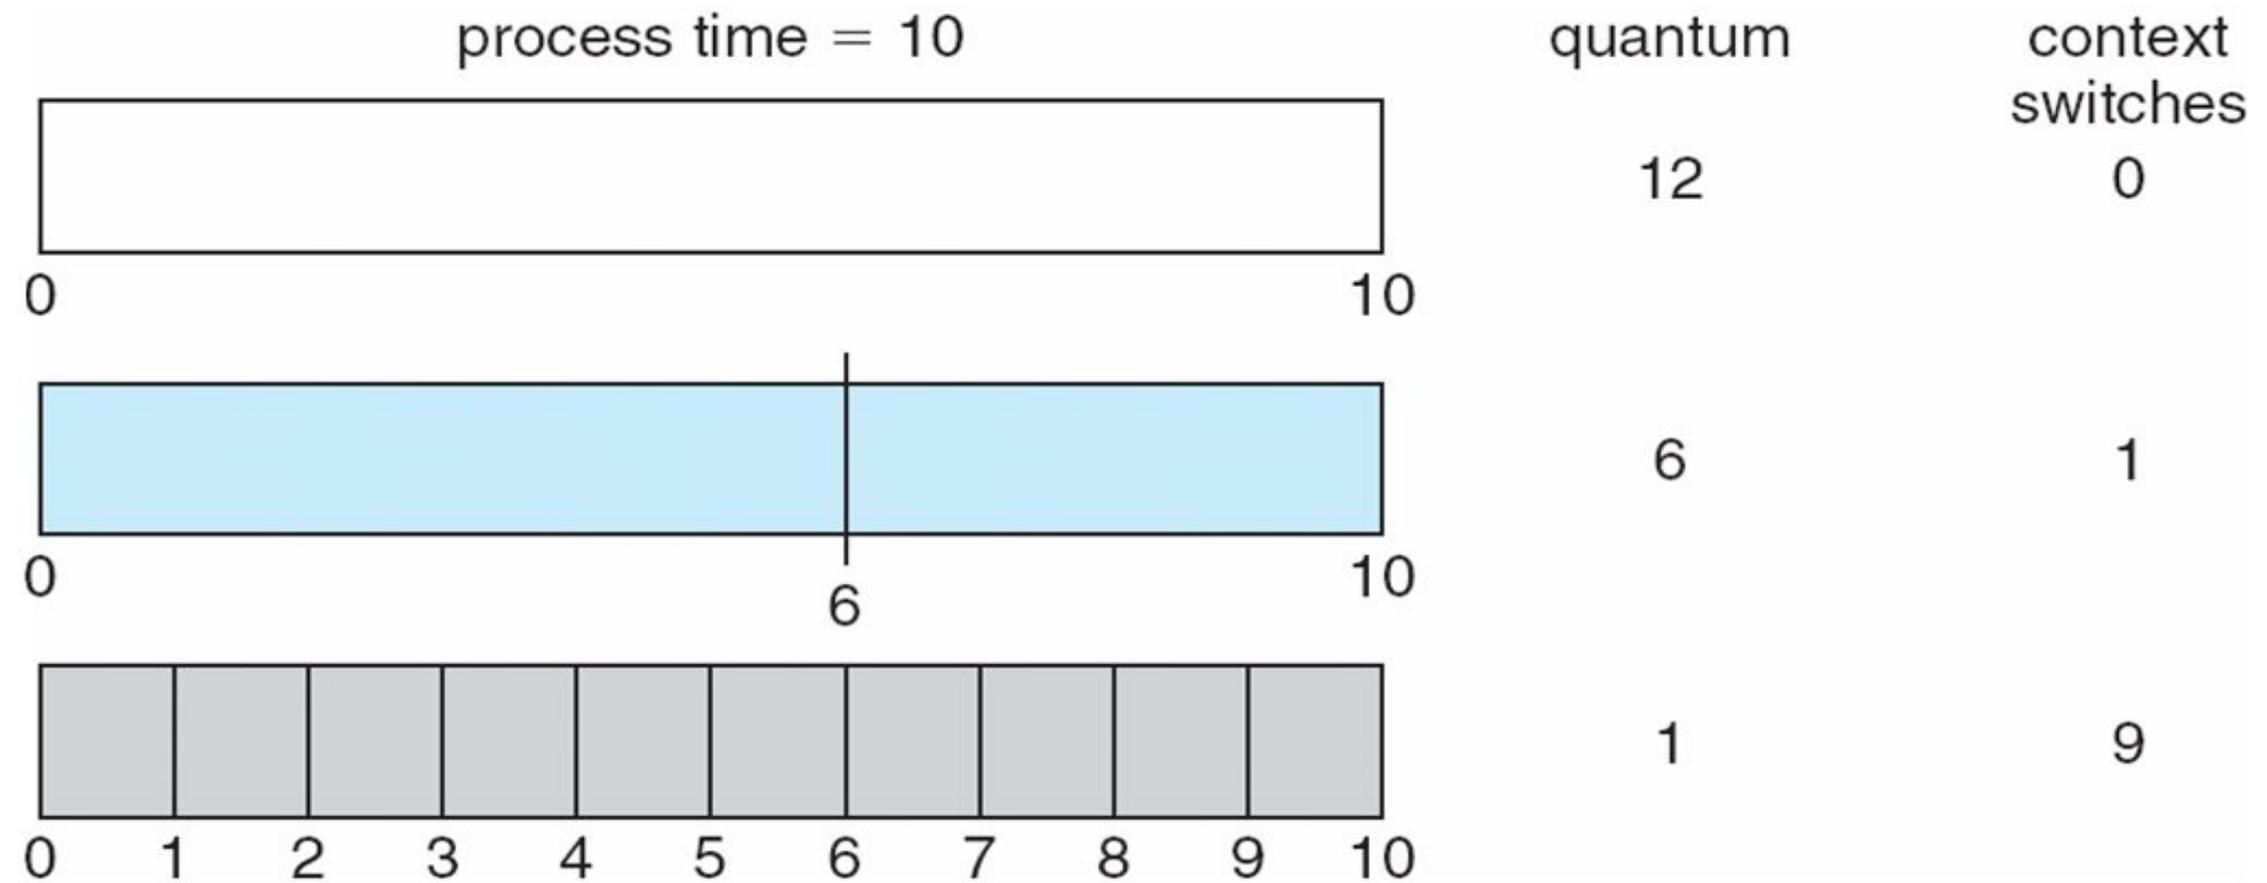


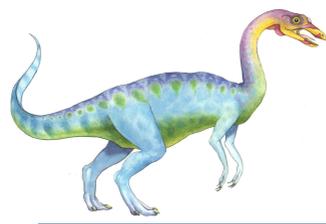
- Average waiting time is $17 / 3 = 5.66$ milisecond
- Typically, higher average turnaround than SJF, but better *response*
- quantum should be large compared to context switch time
- Quantum usually 10ms to 100ms, context switch < 10 usec



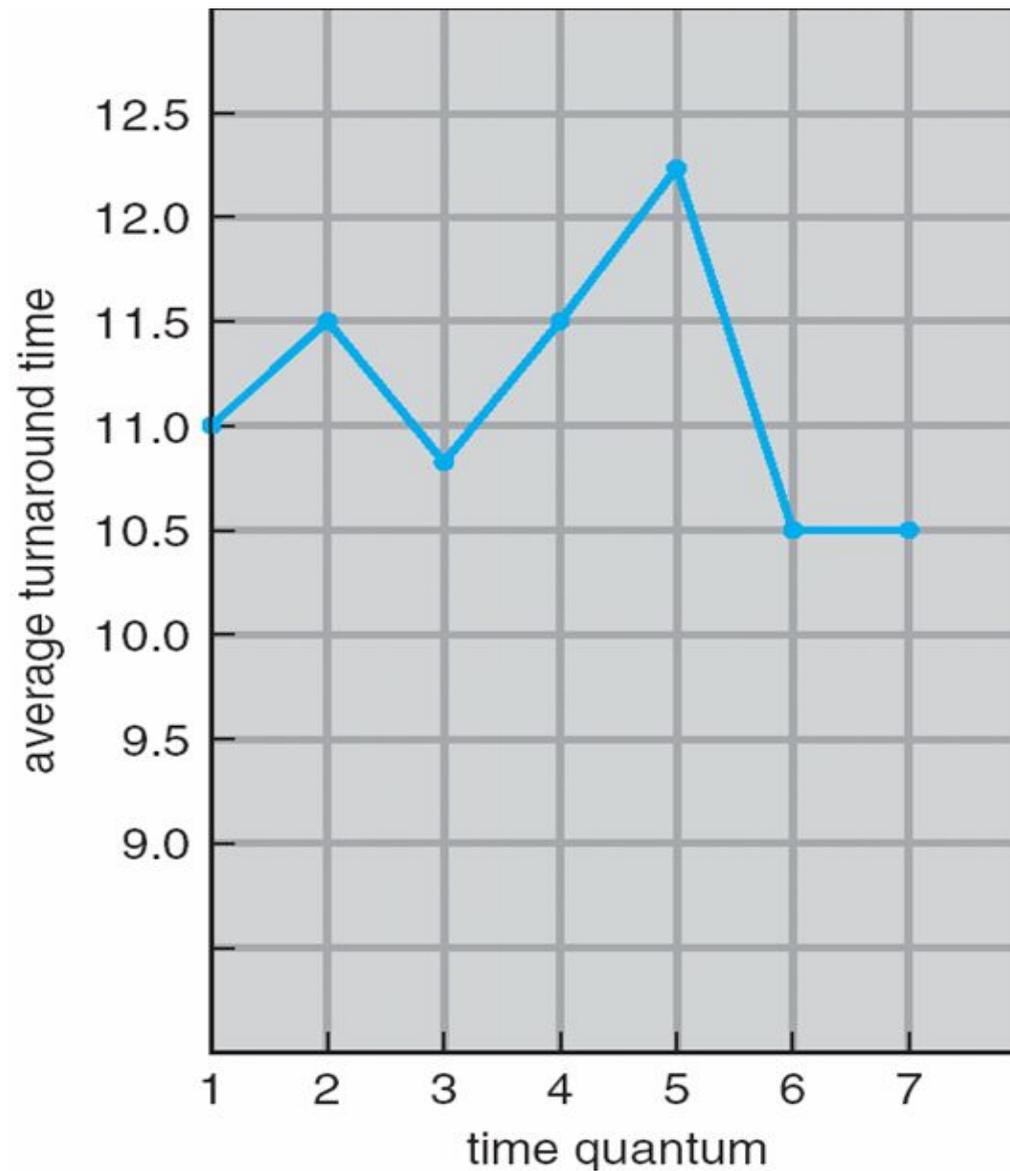


Time Quantum and Context Switch Time





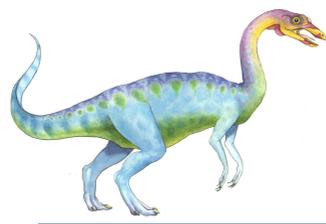
Turnaround Time Varies With The Time Quantum



process	time
P_1	6
P_2	3
P_3	1
P_4	7

80% of CPU bursts should be shorter than quantum

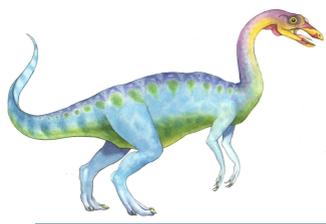




Multilevel Queue

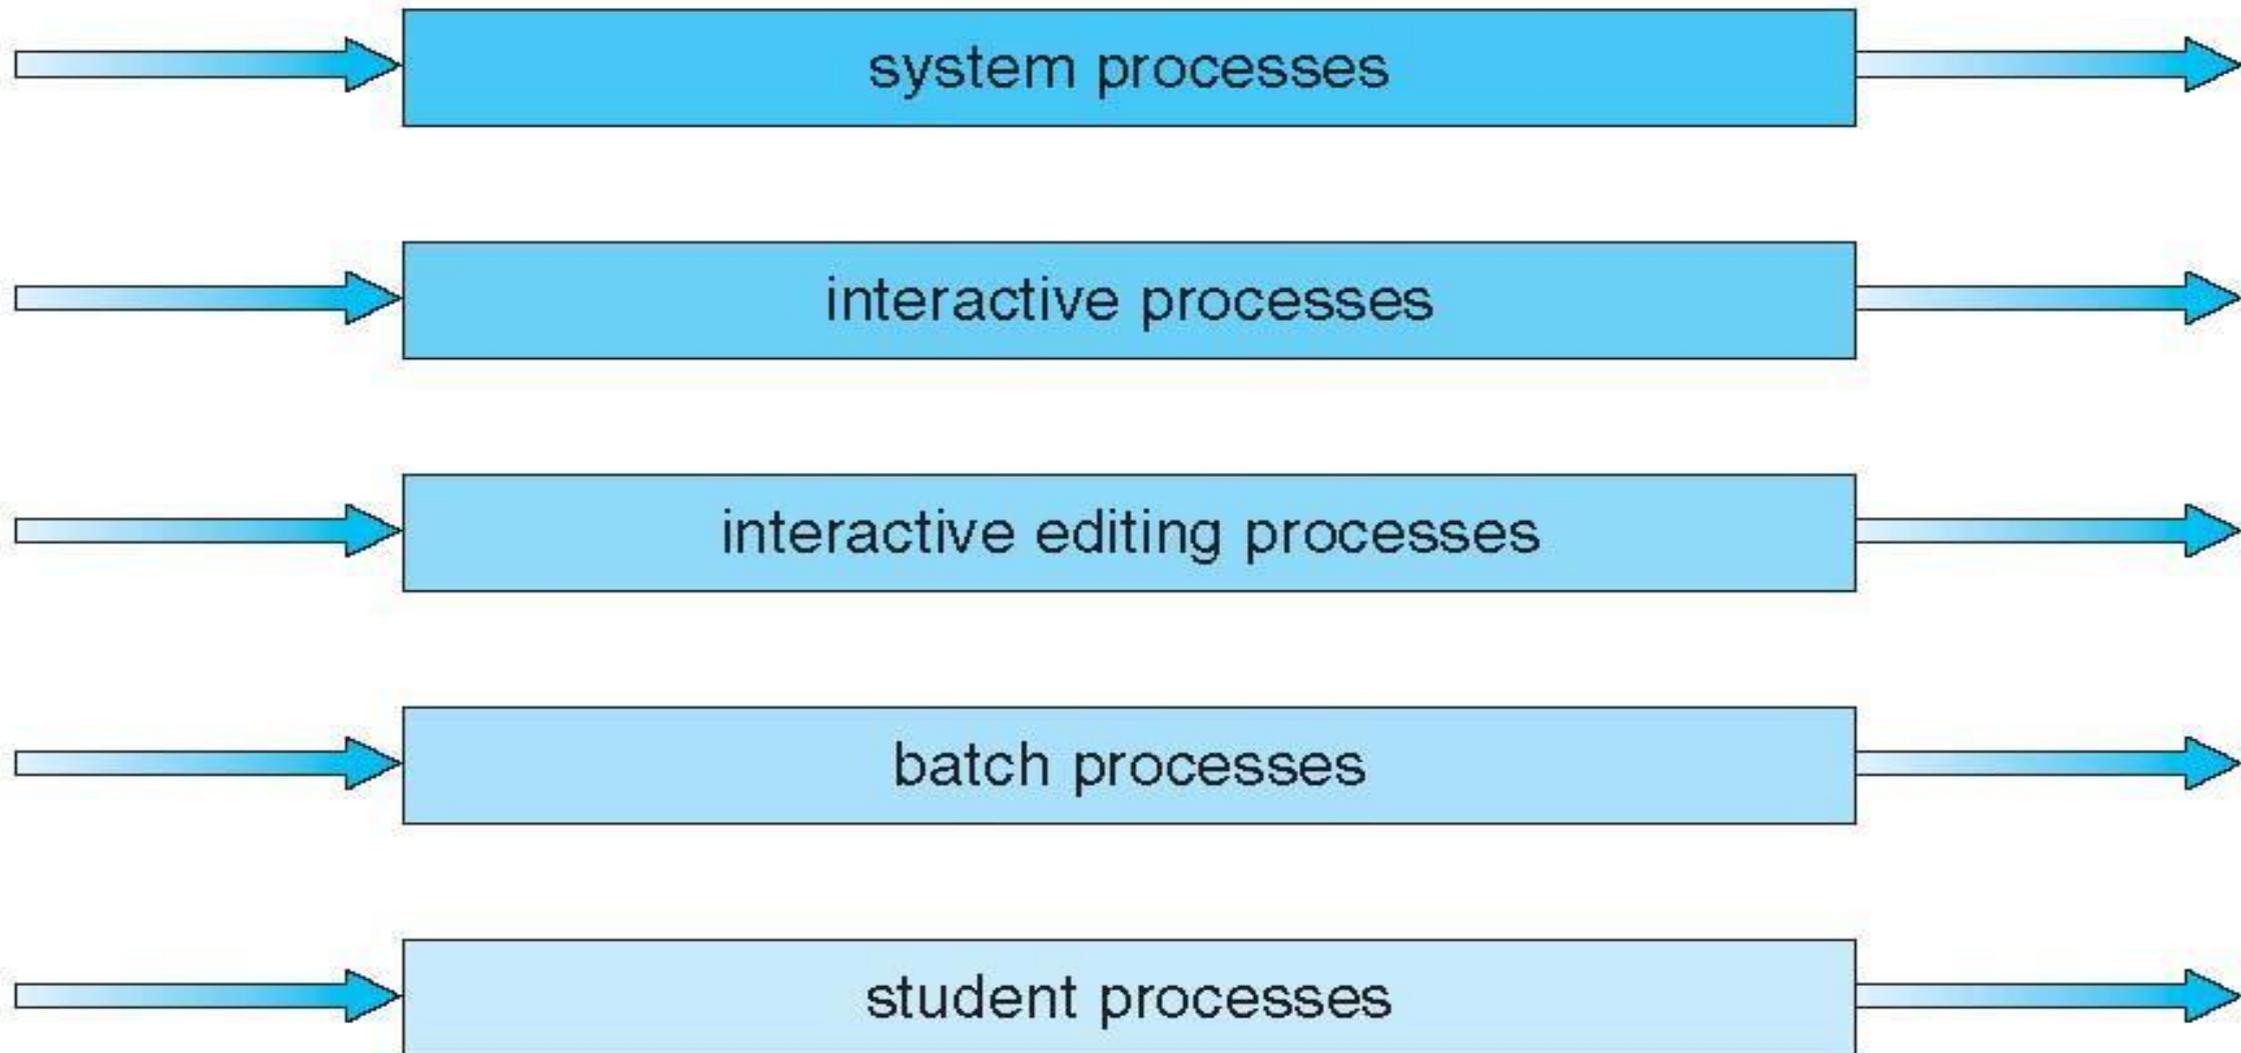
- Another class of scheduling algorithm needs- in which processes are classified into different groups, e.g.:
 - foreground (interactive) processes
 - background (batch) processes
- They have different response time requirements-so different scheduling needs.
- Foreground processes may have priority over background processes.
- A multilevel queue-scheduling algorithm partitions the ready queue into several separate queues-we can see it in the figure of next slide:-
- Each queue has its own scheduling algorithm:
 - Foreground queue scheduled by – RR algorithm
 - Background queue scheduled by – FCFS algorithm
- Scheduling must be done between the queues:
 - Fixed priority preemptive scheduling; (i.e., serve all from foreground then from background). Possibility of starvation.
 - Time slice – each queue gets a certain amount of CPU time which it can schedule amongst its processes; i.e., foreground queue can be given 80% of the CPU time for RR-scheduling among its processes, while 20% to background in FCFS manner.



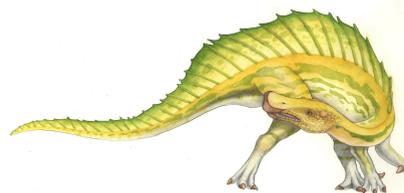


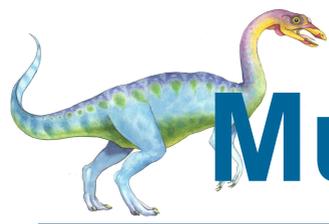
Multilevel Queue Scheduling

highest priority



lowest priority

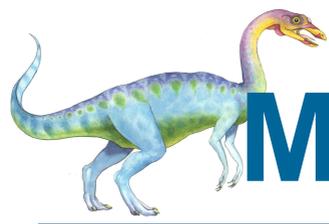




Multilevel Feedback Queue scheduling

- Processes do not move from one queue to the other----But
- Multilevel Feedback Queue scheduling, allows a process to move between queues.
- If a process uses too much CPU time, it will be moved to a lower priority queue.
- Similarly, a process that waits too long in a lower-priority queue may be moved to a higher-priority queue.

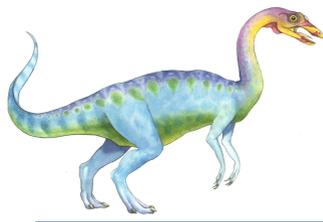




Multilevel Feedback Queue scheduling

- Multilevel-feedback-queue scheduler defined by the following parameters:
 - number of queues
 - scheduling algorithms for each queue
 - method used to determine when to upgrade a process
 - method used to determine when to demote a process
 - method used to determine which queue a process will enter when that process needs service

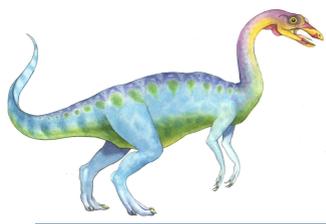




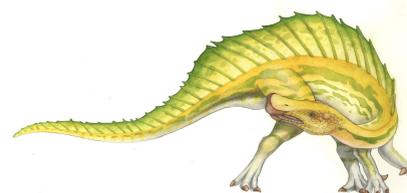
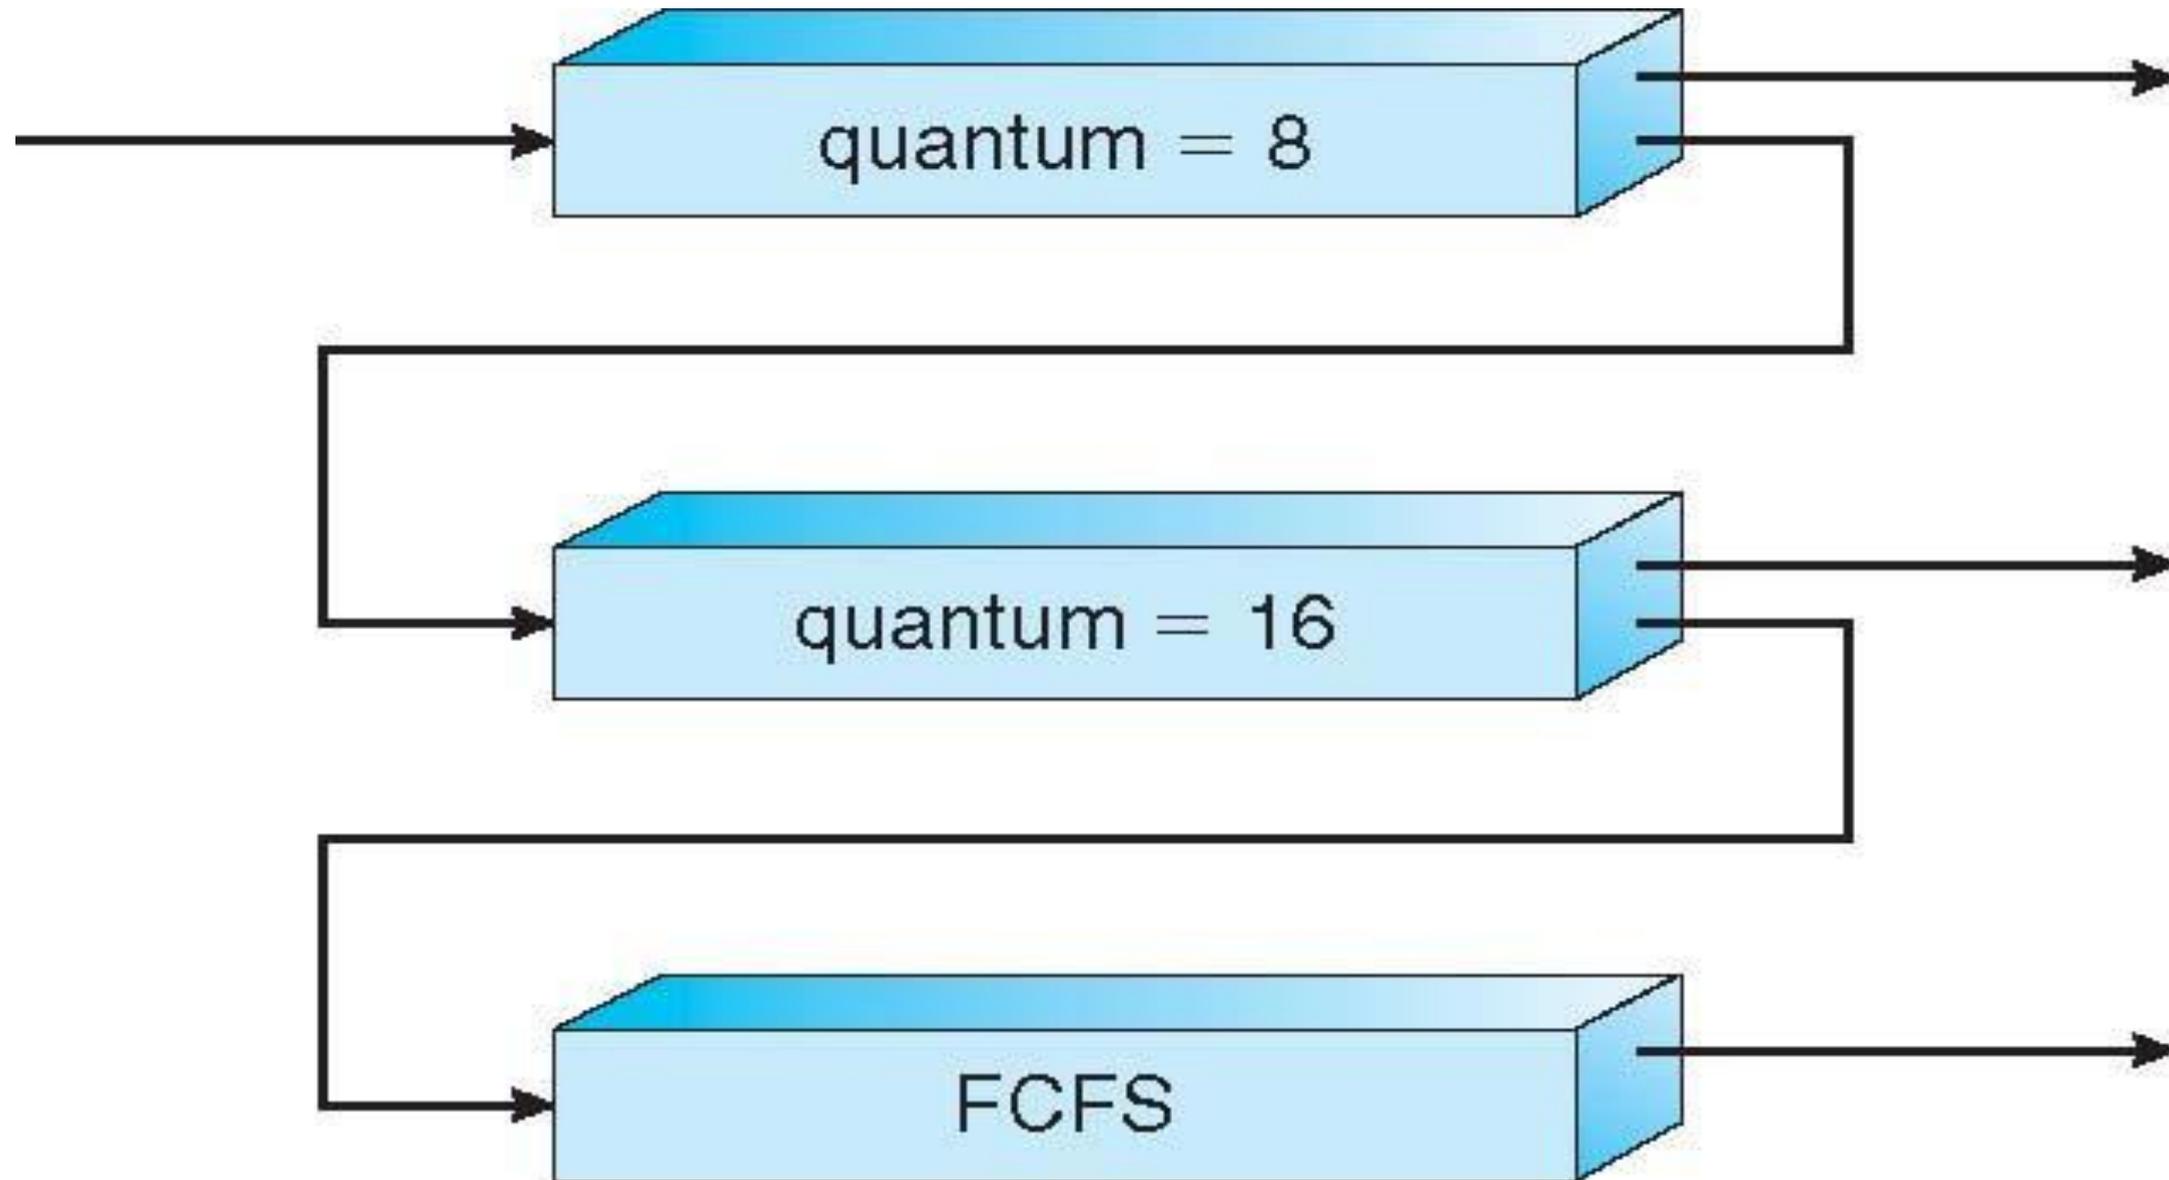
Example of Multilevel Feedback Queue

- Three queues: (can see the figure in next slide)
 - Q_0 – RR with time quantum 8 milliseconds
 - Q_1 – RR time quantum 16 milliseconds
 - Q_2 – FCFS
- Scheduling
 - A new job enters queue Q_0 which is served for RR
 - 4 When it gains CPU, job receives 8 milliseconds
 - 4 If it does not finish in 8 milliseconds, job is moved to queue Q_1
 - At Q_1 job is again served RR and receives 16 additional milliseconds
 - 4 If it still does not complete, it is preempted and moved to queue Q_2





Multilevel Feedback Queues



End of Chapter 5

