CSCI315 – Operating Systems Design

Department of Computer Science
Bucknell University

CPU Scheduling Algorithms

Round-Robin, Feedback Queues

Ch 5.3

This set of notes is based on notes from the textbook authors, as well as L. Felipe Perrone, Joshua Stough, and other instructors.
Xiannong Meng, Fall 2021.

CPU Scheduling Algorithms

- In last segment, we discussed some scheduling algorithms, FCFS, Priority, and SJF (the optimal algorithm)
- In this segment, we will discuss Round-Robin (RR) and some other scheduling ideas.

Round Robin (RR)

- Each process gets a small unit of CPU time (time ***quantum***), 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.

RR with Time Quantum = 20

Process	Burst Time
P_1	53
P_2	17
P_3	68
P_4	24

The Gantt chart is:

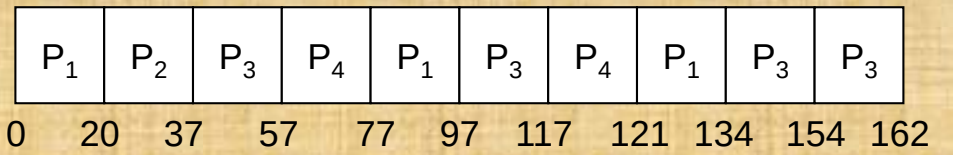

Typically, higher average turnaround than SJF, but better *response.*

Time Quantum and Context Switch Time

quantum	context switches
12	0
6	1
1	9

Question: What influences the choice of value for the quantum?

Turnaround Time Varies with the Time Quantum

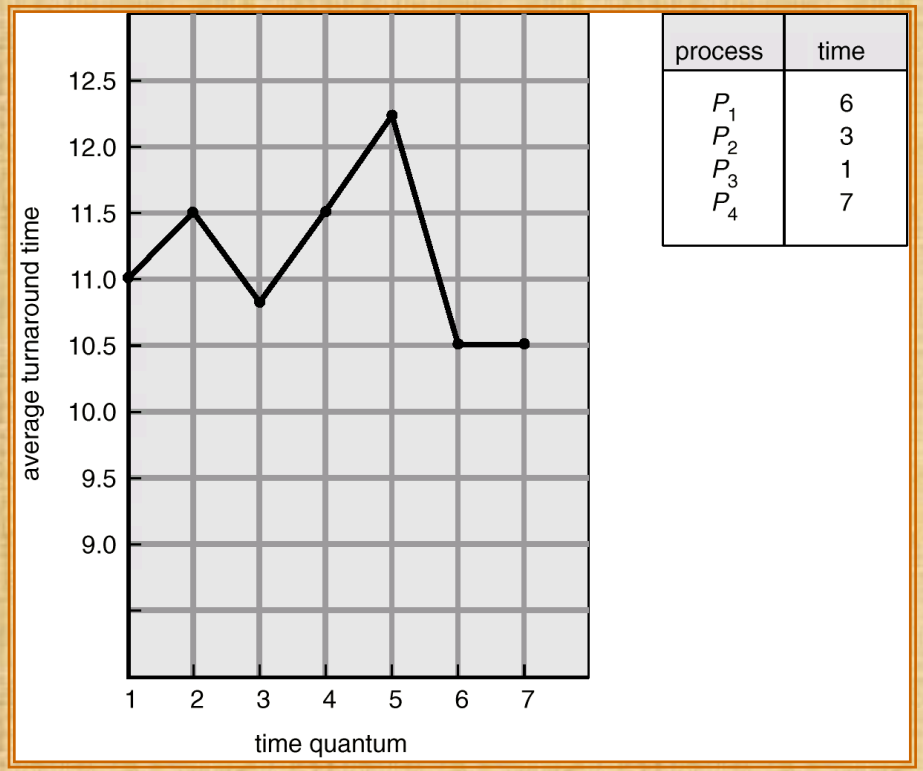

process	time
P_1	6
P_2	3
P_3	1
P_4	7

Performance of RR

- Effects of the quantum length q:
  - q large $\Rightarrow$ FIFO.
  - q small $\Rightarrow$ q must be large with respect to context switch, otherwise overhead is too high.
  - If q is extremely small, and if we ignore the context switch cost, the result is processor sharing.

Multilevel Queue

- Ready queue is partitioned into separate queues:
  - foreground (interactive)
  - background (batch)
- Each queue has its own scheduling algorithm.
  - foreground: RR
  - background: FCFS
- Scheduling must be done between the queues:
  - Fixed priority 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., 80% to foreground in RR.
  - 20% to background in FCFS .

Multilevel Queue Scheduling

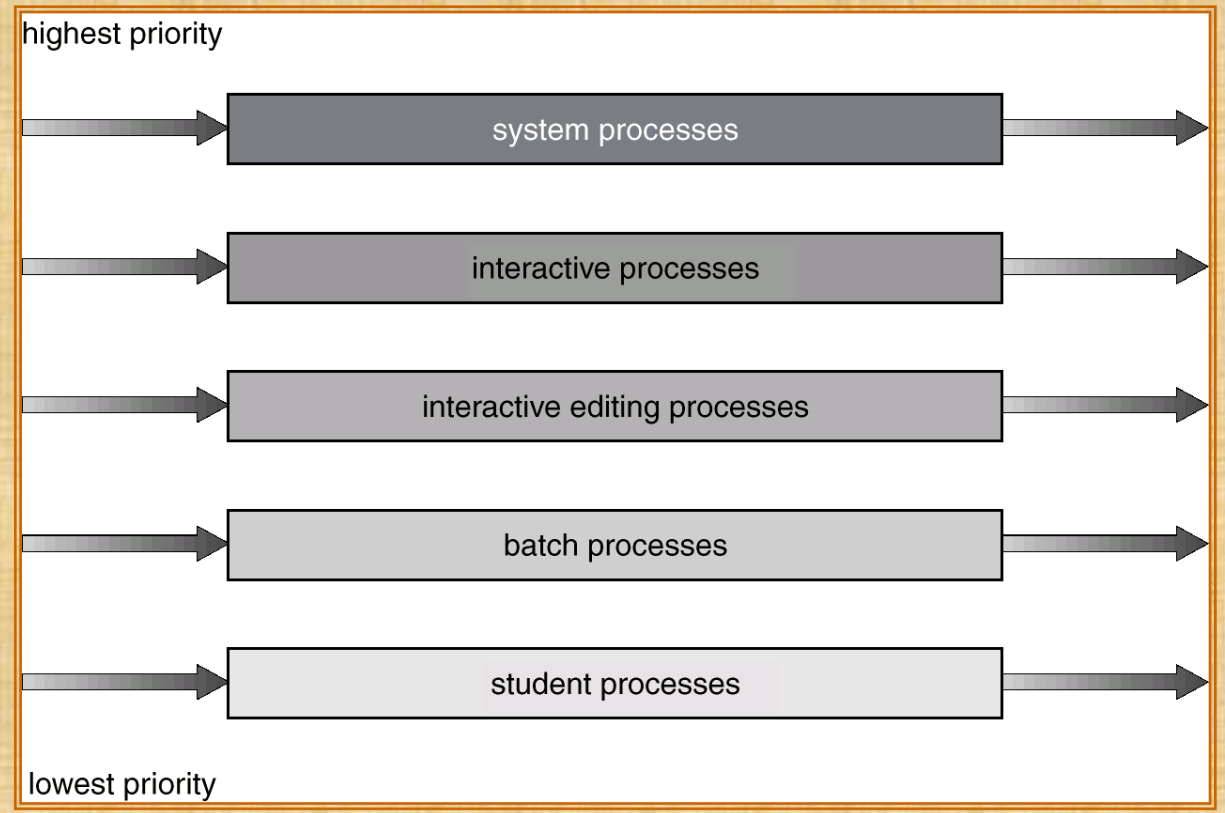

Multilevel Feedback Queue

- A process can move between the various queues; aging can be implemented this way.

- 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:
  - Q_0 – time quantum 8 milliseconds (most favorite queue)
  - Q_1 – time quantum 16 milliseconds
  - Q_2 – FCFS (least favorite queue)

- Scheduling
  - A new job enters queue Q_0 which is served FCFS. When it gains CPU, job receives 8 milliseconds. If it does not finish in 8 milliseconds, job is moved to queue Q_1.
  - At Q_1 job is again served FCFS and receives 16 additional milliseconds. If it still does not complete, it is preempted and moved to queue Q_2.

Multilevel Feedback Queues