Alexandria University
Faculty of Engineering
Comp. & Comm. Engineering
CC373: Operating Systems

جامعة الاسكندرية
كلية الهندسة
برنامج هندسة الحاسب والاتصالات
مادة نظم التشغيل

Sheet3
Uniprocessor Scheduling

1) Briefly describe the three types of processor scheduling.

   Long-term scheduling: The decision to add to the pool of processes to be executed. **Medium-term scheduling:** The decision to add to the number of processes that are partially or fully in main memory.

   Short-term scheduling: The decision as to which available process will be executed by the processor

2) What is usually the critical performance requirement in an interactive operating system?

   Response time

3) What is the difference between turnaround time and response time?

   Turnaround time is the total time that a request spends in the system (waiting time plus service time. **Response time** is the elapsed time between the submission of a request until the response begins to appear as output.

4) For process scheduling, does a low-priority value represent a low priority or a high priority?

   In UNIX and many other systems, larger priority values represent lower priority processes. Some systems, such as Windows, use the opposite convention: a higher number means a higher priority

5) What is the difference between preemptive and nonpreemptive scheduling?

   Nonpreemptive: If a process is in the Running state, it continues to execute until (a) it terminates or (b) blocks itself to wait for I/O or to request some operating system service. **Preemptive:** The currently running process may be interrupted and moved to the Ready state by the operating system. The 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.

6) Briefly define FCFS scheduling.

   As each process becomes ready, it joins the ready queue. When the currently-running process ceases to execute, the process that has been in the ready queue the longest is selected for running.

7) Briefly define round-robin scheduling.

   A clock interrupt is generated at periodic intervals. When the interrupt occurs, the currently running process is placed in the ready queue, and the next ready job is selected on a FCFS basis.

8) Briefly define shortest-process-next scheduling.

   This is a nonpreemptive policy in which the process with the shortest expected processing time is selected next

9) Briefly define shortest-remaining-time scheduling.

   This is a preemptive version of SPN. In this case, the scheduler always chooses the process that has the shortest expected remaining processing time. When a new process joins the ready queue, it may in fact have a shorter remaining time than the currently running process. Accordingly, the scheduler may preempt whenever a new process becomes ready

10) Briefly define highest-response-ratio-next scheduling.

When the current process completes or is blocked, choose the ready process with the greatest value of R, where R = (w + s)/s, with w = time spent waiting for the processor and s = expected service time

11) Briefly define 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 (see Figure 9.4). After its first execution, 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. A shorter process will complete quickly, without migrating very far down the hierarchy of ready queues. A longer process will gradually drift downward. Thus, newer, shorter processes are favored over older, longer processes. Within each queue, except the lowest-priority queue, a simple FCFS mechanism is used. Once in the lowest-priority queue, a process cannot go lower, but is returned to this queue repeatedly until it completes execution

12) Consider the following set of processes:

Process	Arrival Time	Processing Time
A	0	1
B	1	9
C	2	1
D	3	9

For each of the following scheduling methods, give (1) a timing chart to illustrate the execution sequence, (2) service time, (3) finish time, (4) turnaround time, (5) normalized turnaround time and (6) waiting time.

a) FCFS
b) RR, q = 1
c) RR, q = 4
d) SPN
e) SRT
f) HRRN

Each square represents one time unit; the Char in the square refers to the currently-running process.

FCFS	A	B	B	B	B	B	B	B	B	B	C	D	D	D	D	D	D	D	D	D
RR, q=1	A	B	C	B	D	B	D	B	D	B	D	B	D	B	D	B	D	B	D	D
RR, q=4	A	B	B	B	B	C	D	D	D	D	B	B	B	B	D	D	D	D	B	D
SPN	A	B	B	B	B	B	B	B	B	B	C	D	D	D	D	D	D	D	D	D
SRT	A	B	C	B	B	B	B	B	B	B	B	D	D	D	D	D	D	D	D	D
HRRN	A	B	B	B	B	B	B	B	B	B	C	D	D	D	D	D	D	D	D	D
Feedback, q=1	A	B	C	D	B	D	B	D	B	D	B	D	B	D	B	D	B	D	B	D
Feedback, $q=2^i$	A	B	C	D	B	B	D	D	B	B	B	B	D	D	D	D	B	B	D	D

a) FCFS

- service time = T_s = 1 + 9 + 1 + 9 = 20
- finish time = 1 + 10 + 11 + 20 = 42 → Mean = 10.5

- turnaround time = T_r = 1 + 9 + 9 + 17 = 36 → Mean = 9
- normalized turnaround time = (T_r/T_s) = (1/1)+(9/9) + (9/1) + (17/9) = 12.88889 → Mean = 3.22225
- waiting time = 0 + 0 + 8 + 8 = 16 → Mean = 4

b) RR, q=1
- service time = T_s = 1 + 9 + 1 + 9 = 20
- finish time = 1 + 19 + 3 + 20 = 43 → Mean = 10.75
- turnaround time = T_r = 1 + 18 + 1 + 17 = 37 → Mean = 9.25
- normalized turnaround time = (T_r/T_s) = (1/1)+(18/9) + (1/1) + (17/9) = 5.88889 → Mean = 1.47222
- waiting time = 0 + 0 + 0 + 0 = 0 → Mean = 0

c) RR , q=4
- service time = T_s = 1 + 9 + 1 + 9 = 20
- finish time = 1 + 19 + 6 + 20 = 46 → Mean = 11.5
- turnaround time = T_r = 1 + 18 + 4 + 17 = 40 → Mean = 10
- normalized turnaround time = (T_r/T_s) = (1/1)+(18/9) + (4/1) + (17/9) = 8.88889 → Mean = 2.22225
- waiting time = 0 + 0 + 3 + 3 = 6 → Mean = 1.5

d) SPN
- service time = T_s = 1 + 9 + 1 + 9 = 20
- finish time = 1 + 10 + 11 + 20 = 42 → Mean = 10.5
- turnaround time = T_r = 1 + 9 + 9 + 17 = 36 → Mean = 9
- normalized turnaround time = (T_r/T_s) = (1/1)+(9/9) + (9/1) + (17/9) = 12.88889 → Mean = 3.22225
- waiting time = 0 + 0 + 8 + 8 = 16 → Mean = 4

e) SRT
- service time = T_s = 1 + 9 + 1 + 9 = 20
- finish time = 1 + 11 + 3 + 20 = 35 → Mean = 8.75
- turnaround time = T_r = 1 + 10 + 1 + 17 = 29 → Mean = 7.25
- normalized turnaround time = (T_r/T_s) = (1/1)+(10/9) + (1/1) + (17/9) = 5 → Mean = 1.25
- waiting time = 0 + 0 + 0 + 8 = 8 → Mean = 2

f) HRRN
- service time = T_s = 1 + 9 + 1 + 9 = 20
- finish time = 1 + 10 + 11 + 20 = 42 → Mean = 10.5
- turnaround time = T_r = 1 + 9 + 9 + 17 = 36 → Mean = 9
- normalized turnaround time = (T_r/T_s) = (1/1)+(9/9) + (9/1) + (17/9) = 12.88889 → Mean = 3.22225
- waiting time = 0 + 0 + 8 + 8 = 16 → Mean = 4

13) Most round-robin schedulers use a fixed size quantum. Give an argument in favor of a small quantum. Now give an argument in favor of a large quantum. Compare and contrast the types of systems and jobs to which the arguments apply. Are there any for which both are reasonable?

a) **Small quantum:** favors short processes and more overhead due to frequent context switching.

b) **Large quantum:** favors long processes and poor performance for short interactive requests.

14) Five batch jobs, A through E, arrive at a computer center at essentially the same time. They have an estimated running time of 15, 9, 3, 6, and 12 minutes, respectively. Their (externally defined) priorities are 6, 3, 7, 9, and 4 respectively, with a lower value corresponding to a higher priority. For each of the following scheduling algorithms, determine the turnaround time for each process and the average turnaround for all jobs. Ignore process switching overhead. Explain how you arrived at your answers. In the last three cases, assume that only one job at a time runs until it finishes and that all jobs are completely processor bound.

a) round robin with a time quantum of 1 minute
b) priority scheduling
c) FCFS (run in order 15, 9, 3, 6, and 12)
d) shortest job first

a) **Round robin with time quantum = 1 minute** (Assume they are ordered according to their priorities)

P21, P51, P11, P31, P41
P22, P52, P12, P32, P42
P23, P53, P13, P33, P43
P24, P54, P14, P44
P25, P55, P15, P45
P26, P56, P16, P46
P27, P57, P17
P28, P58, P18
P29, P59, P19
P510, P110
P511, P111
P512, P112
P113
P114
P115

Turnaround times:
P1: 45
P2: 36
P3: 15
P4: 27
P5: 42
Average = 33

b) Priority scheduling

P2, P5, P1, P3, P4

Turnaround times:
P1: 36
P2: 9
P3: 39
P4: 45
P5: 21
Average = 30

c) **FCFS**
P1, P2, P3, P4, P5
Turnaround times:
P1: 15
P2: 24
P3: 27
P4: 33
P5: 45
Average = 28.8

d) **Shortest job first**
P3, P4, P2, P5, P1
Turnaround times:
P1: 45
P2: 18
P3: 3
P4: 9
P5: 30
Average = 21

15) Consider the following three processes. Each process makes a CPU burst then an I/O burst, another CPU burst, another I/O burst and terminates with a CPU burst. The lengths of the CPU burst and I/O burst times in milliseconds are given in the following table:

Process	CPU-Burst 1	I/O-Burst 1	CPU-Burst 2	I/O-Burst 2	CPU-Burst	Arrival Time
P1	2	4	2	2	2	0
P2	2	2	3	3	1	1
P3	1	2	1	1	1	1

The processes are assumed to arrive as indicated. Draw three time diagrams that illustrate the execution of these processes using

a) The first come first served (FCFS) algorithm
b) The round robin (RR) scheduling algorithm with quantum = 1 and
c) The round robin (RR) scheduling algorithm with quantum = 2.

If an I/O completion and a CPU timeout of two processes occur at the same time, we treat the I/O completion first. Add three lines to the Gantt chart that represent the ready queue.

d) What is the waiting time of each process for of the above scheduling algorithms? What is the mean value?
e) What is the turn-around time of each process for each of the above scheduling algorithms? What is the mean value?
f) What happens if we assign a priority to each process? The priority of P2 is higher than the priority of P1 and P3. Consider the situations of RR, where quantum=1 and quantum=3. Do not consider FCFS (Why not?).

They are as follows (**Red** means doing CPU work, **green** means doing IO work and **black** means waiting):

a) **FCFS**

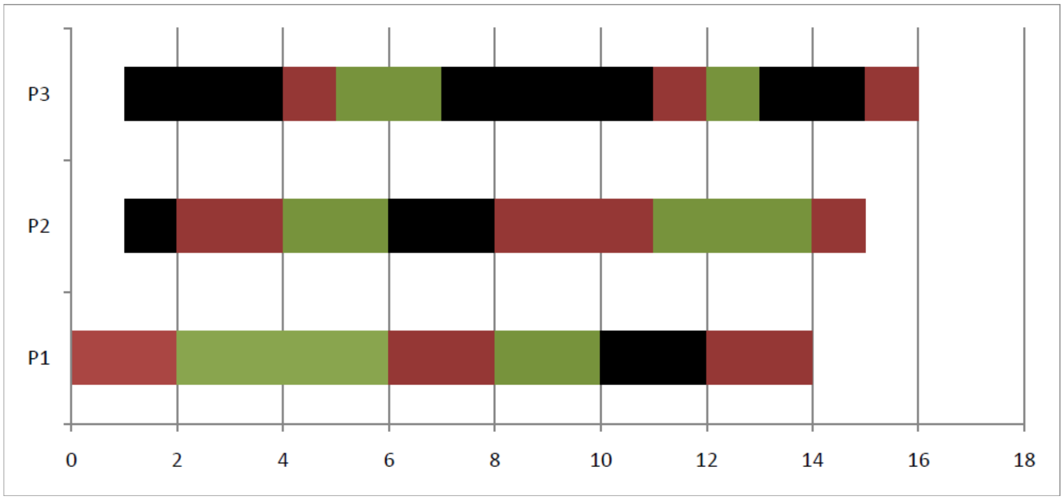

Waiting time (black periods):

P1: 2
P2: 3
P3: 9
Average = 14/3

Response time:

P1: 0
P2: 1
P3: 3
Average = 4/3

Turnaround:

P1: 14
P2: 14
P3: 15
Average = 14.333

b) RR, q=1

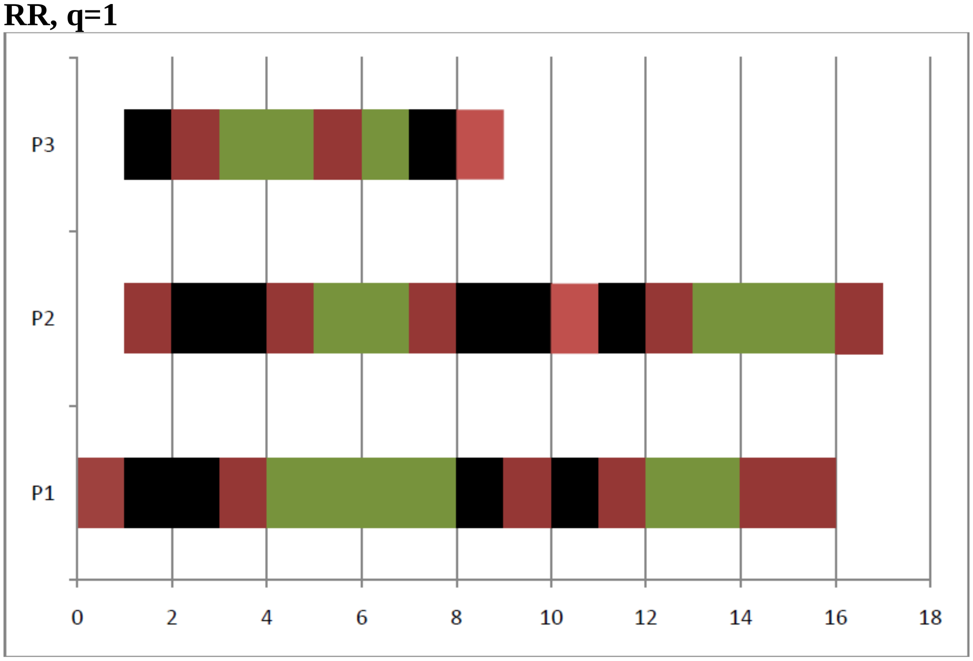

Waiting time (black periods):

P1: 4
P2: 5
P3: 2
Average = 11/3

Response time:

P1: 0
P2: 0
P3: 1
Average = 1/3

Turnaround:

P1: 16
P2: 16
P3: 8
Average = 13.3333

c) RR, q=2

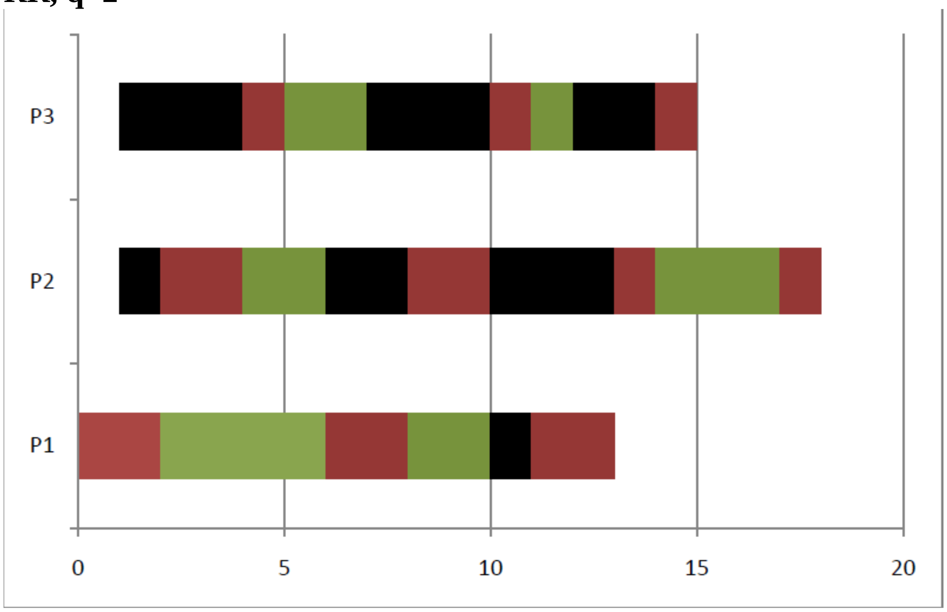

Waiting time (black periods):

P1: 1
P2: 6
P3: 8
Average = 3

Response time:

P1: 0
P2: 1
P3: 3
Average = 4/3

Turnaround:

P1: 13
P2: 17

P3: 14
Average = 14.666
