

Please read 6.1–6.3, 6.5 and 7.1 of the textbook and then answer the following, trying not to look at your notes or at the textbook. Quiz #7, on Thursday 30th November, will consist exclusively of questions taken from the Part I of this homework.

Part I — Short Questions

Question 1

Explain the difference between preemptive and nonpreemptive scheduling

Question 2

What is the difference between multi-programming and multi-processing?

Question 3

Discuss why a scheduler might decide to give to I/O-bound processes a higher priority than to CPU-bound processes.

Question 4

Which of the following scheduling algorithms could result in starvation?

1. First-come, First-served
2. Shortest job first
3. Round robin
4. Priority

Question 5

What is the difference between response time and turnaround time?

Question 6

Why would a scheduler for an interactive system not necessarily be the best choice for a real-time system?

Question 7

What is the convoy effect?

Question 8

Give an example of internal priority, and an example of external priority for a priority-scheduling algorithm.

Question 9

What is aging? Why is it useful?

Question 10

What is processor affinity?

Question 11

Why is isolation between processes such an important feature?

Question 12

What is the difference between a physical address and a virtual address?



Part II — Problem

As usual, I'll assume that you will have successfully completed those two problems by the time Quiz #7 is taken (Thursday 30th November).

Problem 1

Consider the following workload:

Process	Burst Time	Priority	Arrival Time
P_1	50 ms	4	0 ms
P_2	20 ms	1	20 ms
P_3	100 ms	3	40 ms
P_4	40 ms	2	60 ms

- (a) Ignoring the context-switch overhead, show the schedule using shortest remaining time, nonpreemptive priority (a smaller priority number implies higher priority) and round robin with quantum 30 ms.

Here is the example for first come first served, where 1 unit is 10 ms:

P_1	P_1	P_1	P_1	P_1	P_2	P_2	P_3	P_3	P_3	P_3	P_3	P_3	P_3	P_3	P_3	P_3	P_3	P_4	P_4	P_4	P_4
-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------	-------

- (b) What is the average waiting time of the above scheduling policies?
 (c) What is the mean process turnaround time of the above scheduling policies?

Problem 2

Consider the following list of desirable features for a CPU scheduler:

Name	Description
Meeting deadlines	Avoid losing data
CPU utilization	Keep the CPU busy all the time
Balance	Keeping all parts of the system busy
Fairness	Giving each process a fair share of the CPU
Predictability	Avoid quality degradation
Responsive time	Respond to requests quickly
Throughput	Maximize jobs per hours
Policy enforcement	Seeing that stated policy is carried out
Proportionality	Meet users' expectations
Turnaround time	Minimize time between submission and termination

Some of them are critical for **batch systems**, **interactive systems**, **real-time systems**, or **all systems**, some are clearly dispensable for some systems. Identify the correspondences.

