

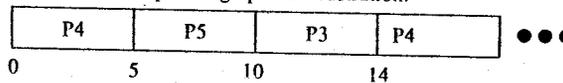
Operating System

1. Consider the following set of processes with their CPU burst times in milliseconds, arrival times in milliseconds and initial priorities. For example, process P1 arrives 2 milliseconds after the start of the system, has priority 10 and needs the CPU for 12 milliseconds. In the table, higher number means higher priority.

Process	Burst Time	Arrival Time	Priority
P1	12	2	10
P2	23	6	20
P3	4	12	0
P4	18	0	30
P5	7	32	15
P6	27	24	10

For each of the scheduling algorithms listed below, graphically show the process execution schedule and calculate the average turnaround time for this set of processes. Show process arrival times, context switch times, and indicate which process is running between context switches. When calculating average turnaround time, you must show your calculation to receive full credit.

Example for graphical illustration.



- a. (10%) Shortest Remaining Time Next
 - b. (10%) Round Robin (with a time quantum of 5 ms)
 - c. (10%) Non-preemptive Priority (priority of a process in the ready queue is increased by 1 each millisecond)
2. (10%) Short Answer Questions (be short!)
- a. What is DMA? What happens during a DMA operation? What will happen if a system does not support DMA?
 - b. What is swap cache? When and why do we need swap cache?
 - c. What is a TLB and where it is used?
 - d. What happens during thrashing (in the memory system) and why it happens?
3. (10%) Short Answer Questions (be short!)
- a. How does file (block) cache replacement/eviction differ from virtual memory replacement?
 - b. Most block caches use a hybrid of exact LRU and MRU policies for victim selection in the file cache. Under what conditions is LRU a better policy than MRU? Under what conditions is MRU a better policy than LRU?
 - c. When LRU is used, block caches typically use exact LRU rather than the approximate LRU replacement used in most virtual memory systems. Why is exact LRU feasible for file block caches but not for virtual memory page caches?

4. (10%) Short Answer Questions about Distributed File Systems (DFS)

- a. What are the functionalities of directory service in DFS?
- b. How does SUN's NFS achieve location transparency?
- c. What are the differences between a hard-mounted and a soft-mounted remote file system?

5. (20%) Suppose a bank system manages the information of customer accounts and passwords. To prevent a customer from obtaining other customers' passwords, you have to prevent information leakage within an application. Please propose an approach to prevent the leakage (Hint: It is a language-based protection).

6. (20%) Formally define a stack algorithm for page replacement. Prove that LRU is a stack algorithm.