1. **(8%)** Consider the following set of processes, with the length of the CPU-burst time given in milliseconds:

Process	Burst Time	Priority
P_1	10	3
P_2	1	1
P_3	2	3
P_4	1	4
P_5	5	2

The processes are assumed to have arrived in the order P_1, P_2, P_3, P_4, P_5, all at time 0.

   (a) Draw four Gantt charts illustrating the execution of these processes using FCFS, SJF, a nonpreemptive priority (a smaller priority number implies a higher priority), and RR (quantum = 1) scheduling.

   (b) What is the turnaround time of each process for each of the scheduling algorithms in part a?

   (c) What is the waiting time of each process for each of the scheduling algorithms in part a?

   (d) Which of the schedules in part a results in the minimal average waiting time (over all processes)?

2. **(10%)** A semaphore S is an integer variable that, apart from initialization, is accessed only through two standard atomic operations: wait(S) and signal(S).

   (a) Please give the classical definitions of wait(S) and signal(S).

   (b) What are the actions in wait(S) operation that must be executed without interruption?

3. **(16%)** Consider a demand-paging system with the following time-measured utilizations:

CPU utilization	20%
Paging disk	97.7%
Other I/O devices	5%

Which (if any) of the following will (probably) improve CPU utilization? Explain your answer.

(a) Install a faster CPU.

(b) Install a bigger paging disk.

(c) Increase the degree of multiprogramming.

(d) Decrease the degree of multiprogramming.

(e) Install more main memory.

(f) Install a faster hard disk, or multiple controllers with multiple hard disks.

(g) Add prepaging to the page fetch algorithms.

(h) Increase the page size.

4. **(8%)** Consider transferring an enormous file of L bytes from host A to host B. Assume an MSS (maximum segment size) of 1460 bytes.

   (a) What is the maximum value of L such that TCP sequence numbers are not exhausted? Recall that the TCP sequence number field has four bytes.

   (b) For the L you obtain in (a), find how long it takes to transmit the file. Assume that a total of 66 bytes of transport, network, and data-link header are added to each segment before the resulting packet is sent out over a 10 Mbps link. Ignore flow control and congestion control so A can pump out the segments back to back and continuously.

5. **(8%)** In mobile IP, what effect will mobility have on end-to-end delays of datagrams between the source and destination?

6. **(12%)** Explain the following terms.

   (a) Local loops

   (b) Well-known ports

   (c) CIDR

7. **(6%)** What is "hidden terminal" problem in wireless networks? How does IEEE 802.11 prevent this problem?

8. **(12 pionts)** Here is a string of address references given as word addresses: 1, 4, 8, 5, 20, 17, 19, 56, 9, 11, 4, 43, 5, 6, 9, 17. Show the hits and misses and final cache contents for a two-way set associative cache with one-word blocks and a total size of 16 words. Assume LRU replacement.

9. Consider three machines with different cache configurations:

   - Cache 1: Direct mapped with one-word blocks.
   - Cache 2: Direct mapped with four-word blocks.
   - Cache 3: 2-way set associative with four-word blocks.

   The following miss rate measurements have been made:

   - Cache 1: Instruction miss rate is 4%; data miss rate 8%.
   - Cache 2: Instruction miss rate is 2%; data miss rate 5%.
   - Cache 3: Instruction miss rate is 2%; data miss rate 4%.

   For these machines, one-half of the instructions contain a data reference. Assume that the cache miss penalty is 6 + Block size in words. The CPI for this workload was measured on a machine with cache 1 and was found to be 2.0. Answer the following two questions.

   (a) **(10 pionts)** Determine which machine spends the most cycles on cache misses.

   (b) **(10 pionts)** The clock rates for the three machines are 10 ns for the first and second machines and 12 ns for the third machine. Determine which machine is the fastest and which is the slowest.