Note: The homework is due before the class on the due date.

Read Chapters 1, 2, 3.1—3.4, 4, and 5. Answer the following questions:

1. What are the long-term scheduler and the short-term scheduler in process scheduling? Should the two schedulers use the same scheduling algorithm? Explain why.

2. Suppose that a set of processes arrive for execution at the times indicated in the table. The duration of each process execution is also shown in the table.

Process	**Arrival Time** (*msec*)	**Burst Duration** (*msec*)
P_1	0	8
P_2	2	3
P_3	3	1
P_4	4	7

For each of the following cases, compute the average waiting time and average turnaround time (we assume that each process has a single CPU burst, and we ignore context switch cost). Show how you compute your results.

(a) FCFS scheduling.

(b) SJF nonpreemptive scheduling.

(c) SJF preemptive scheduling.

(d) Round-Robin scheduling with quantum size 2.

(e) Round-Robin scheduling with quantum size 4.

(f) Multilevel Feedback Queue scheduling with three levels: quantum size 2, quantum size 4, and FCFS.

3. Explain why the storage system has to be organized in a hierarchy.

4. Explain the advantages and disadvantages of user-level thread support compared to kernel-level thread support.

5. **Optional Bonus Question:** Explain the benefits of the hybrid thread support approach used in Solaris 2 (see Section 4.5.2 of the textbook) with regard to user-level only thread support or kernel-level only thread support.