



Chapter 5 Important Points

This is a summary of the parts of Chapter 5 that you should understand and be able to explain.

1. Abstract model of process as sequence of CPU bursts and I/O bursts; properties of processes in these terms.
2. CPU scheduler:
 - (a) definition
 - (b) when it is invoked
 - (c) pre-emptive vs non-pre-emptive
3. Dispatcher; dispatch latency
4. CPU scheduling criteria
5. CPU scheduling algorithms: for each one should be able to determine the **Gantt chart** of a set of processes, the shortcomings and advantages.
 - (a) FCFS
 - i. convoy effect
 - (b) Shortest Job First (SJF)
 - i. exponential averaging
 - (c) Shortest Remaining Time First (SRTF)
 - (d) Round Robin (RR) with various quanta
 - (e) Priority Scheduling
 - i. starvation, aging
6. Multilevel queues/Priority queues
7. Multilevel feedback queues
8. Thread Scheduling
 - (a) Difference between process and system contention scopes
 - (b) Multiprocessor scheduling of threads : common queue versus per-core queues, pros and cons
 - (c) Multithreaded cores - how are threads scheduled
 - (d) Load balancing, push and pull migration and problems with per-core queue scheduling
 - (e) Processor affinity - soft versus hard