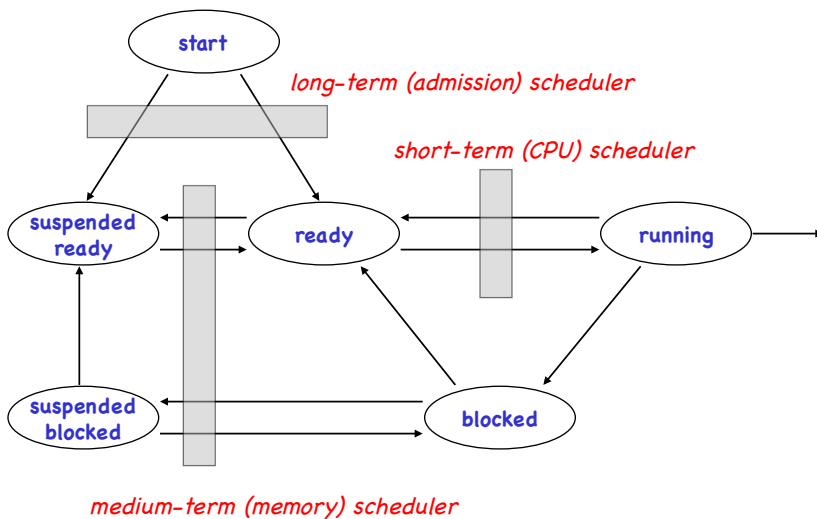


CPU Scheduling


- Schedulers in the OS
- Structure of a CPU Scheduler
 - Scheduling = Selection + Dispatching
- Criteria for scheduling
- Scheduling Algorithms
 - FIFO/FCFS
 - SPF / SRTF
 - Priority - Based

Schedulers



Focus: Short-Term Scheduling


- Recall: Motivation for **multiprogramming** -- have multiple threads in memory to keep CPU busy.
- Typical execution profile of a thread:



- **CPU scheduler** is managing the execution of CPU bursts, represented by threads in ready or running state.

Scheduling Decisions

“Who is going to use the CPU next?!”

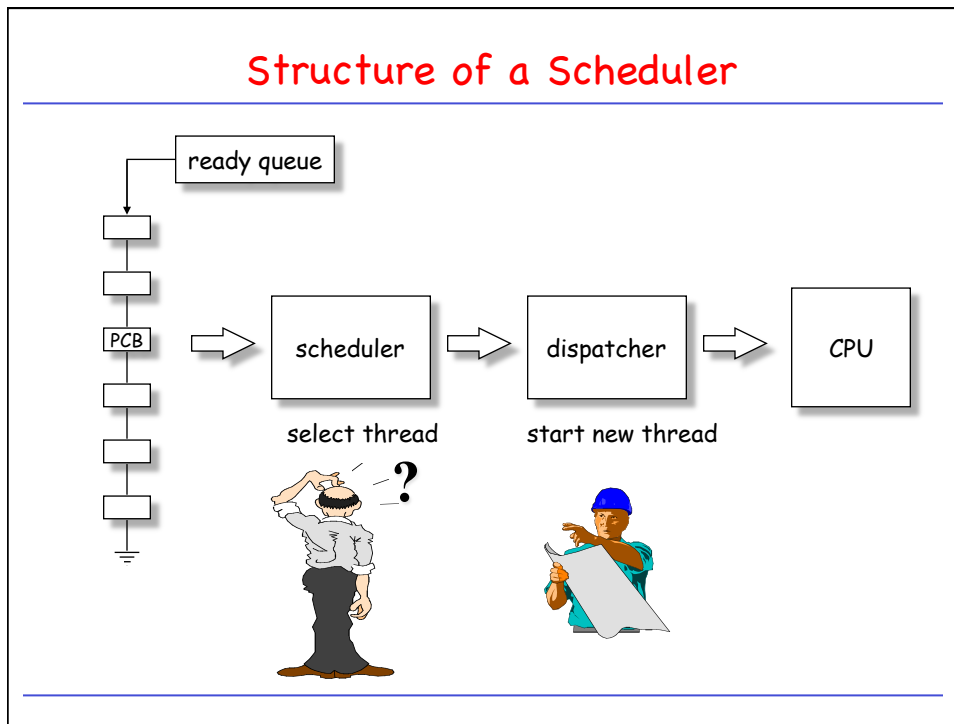


non-preemptive

Scheduling decision points:

- 1. The running thread changes from *running* to *waiting* (current CPU burst of that thread is over).
- 2. The running thread *terminates*.
- 3. A waiting thread becomes *ready* (new CPU burst of that process begins).
- 4. The current thread switches from *running* to *ready*.


preemptive



What Is a Good Scheduler? Criteria

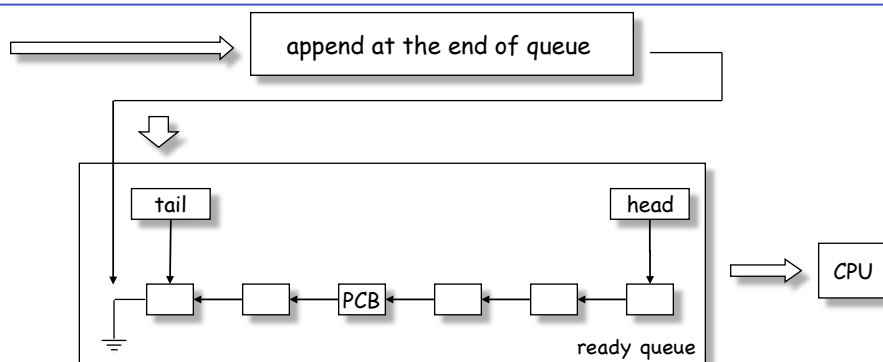
- **User oriented:**
 - **Waiting time** : sum of periods spent waiting in ready queue
 - **Response time** : time interval from submission of job to first response
 - **Normalized response time**: ratio of response time to service time
- **System oriented:**
 - **CPU utilization** : percentage of time CPU is busy
 - **Throughput** : number of jobs completed per time unit
- Any good scheduler should:
 - *maximize* CPU utilization and throughput
 - *minimize* turnaround time, waiting time, response time
- Huh?
 - *maximum/minimum* values vs. *average* values vs. *variance*

Scheduling Algorithms



- **FCFS** : First-come-first-served
- **SPN**: Shortest Process Next
- **SRT**: Shortest Remaining Time
- priority scheduling
- **RR** : Round-robin
- **MLFQ**: Multilevel feedback queue scheduling
- Multiprocessor scheduling

First-Come-First-Served (FCFS/FIFO)



- Advantages:
 - very *simple*
- Disadvantages:
 - long average and worst-case *waiting times*
 - poor dynamic behavior (*convoy effect*)

Shortest Process Next

- Whenever CPU is idle, picks thread with **shortest next CPU burst**.
- Advantages: minimizes average waiting times.
- Problem: Starvation of threads with long CPU bursts.
- Problem: How to determine length of **next CPU burst**!?

SJF Minimizes Average Waiting Time

- Provably optimal: Proof: swapping of jobs

$$dW = t_{short} - t_{long} < 0$$

- Example:

6	12	8	4	W = 6+18+26 = 50
6	8	12	4	W = 6+14+26 = 46
6	8	4	12	W = 6+14+18 = 38
6	4	8	12	W = 6+10+18 = 34
4	6	8	12	W = 4+10+18 = 32

