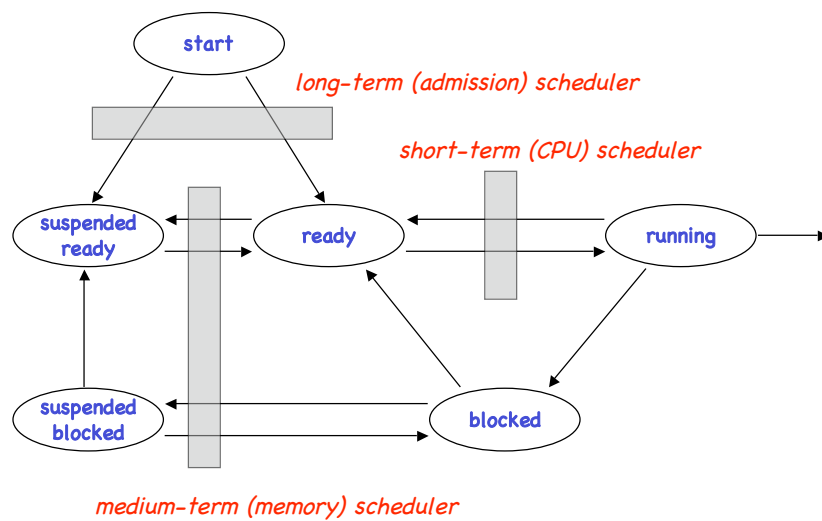


CPU Scheduling


- 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 processes in memory to keep CPU busy.
- Typical execution profile of a process/thread:



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

Scheduling Decisions

"Who is going to use the CPU next?!"

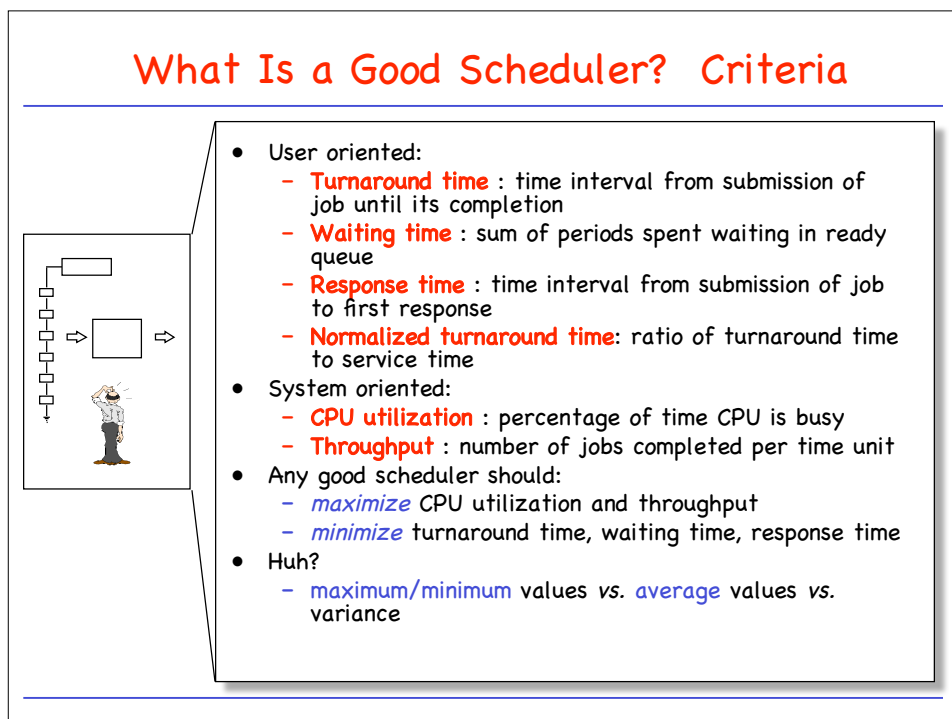
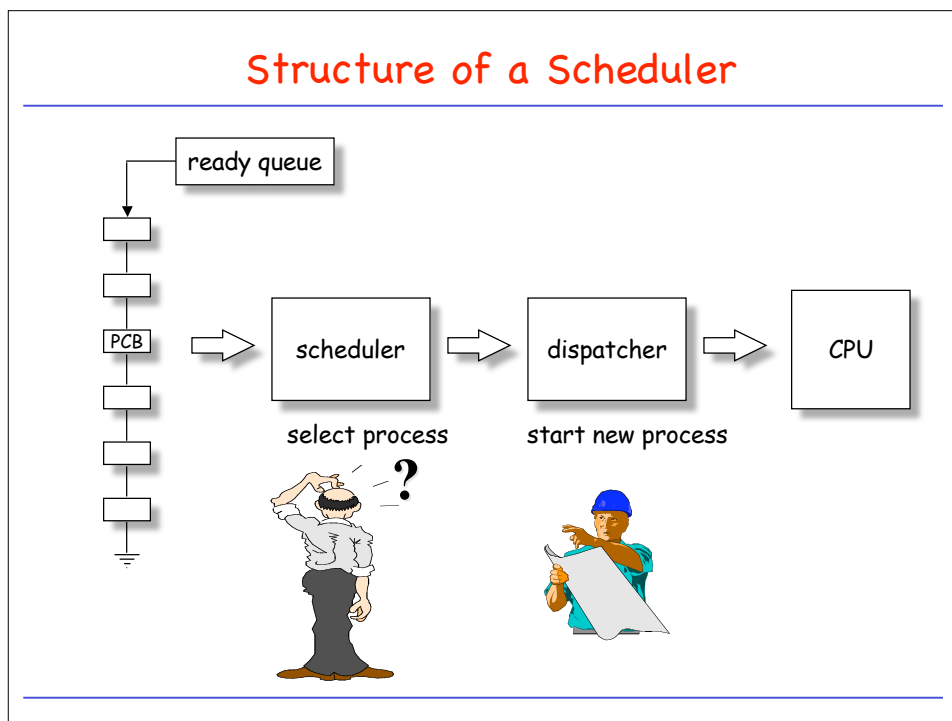


non-preemptive

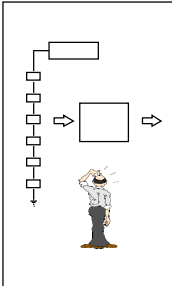
Scheduling decision points:

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

preemptive

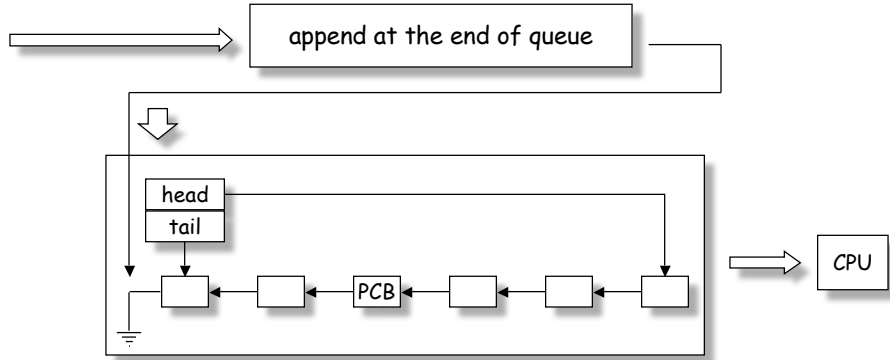


Scheduling Algorithms



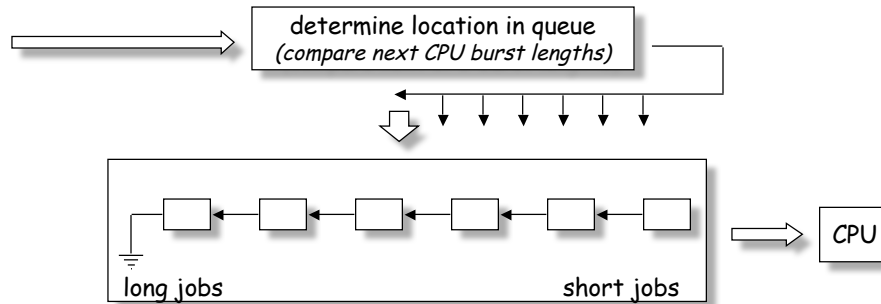
- **FCFS** : First-come-first-served
- **SPN**: Shortest Process Next
- **SRT**: Shortest Remaining Time
- priority scheduling
- **RR** : Round-robin
- **MLFQ**: Multilevel feedback queue scheduling
- Multiprocessor scheduling

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 process with **shortest next CPU burst**.
- Advantages: minimizes average waiting times.
- Problem: How to determine length of **next CPU burst**!?
- Problem: Starvation of jobs with long CPU bursts.

SJF Minimizes Average Waiting Time

- Provably optimal: Proof: swapping of jobs



- 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	$W = 4+10+18 = 32$

