

# 28 Jan 2025 - Operating Systems - II - Week 04

## CPU Scheduling

CPU scheduler → 4 cases: process

1. switches from running to waiting
2. switches from running to ready
3. switches from waiting to ready
4. terminates

running  
↓  
running in  
the CPU

1 & 4 : non preemptive  
by itself

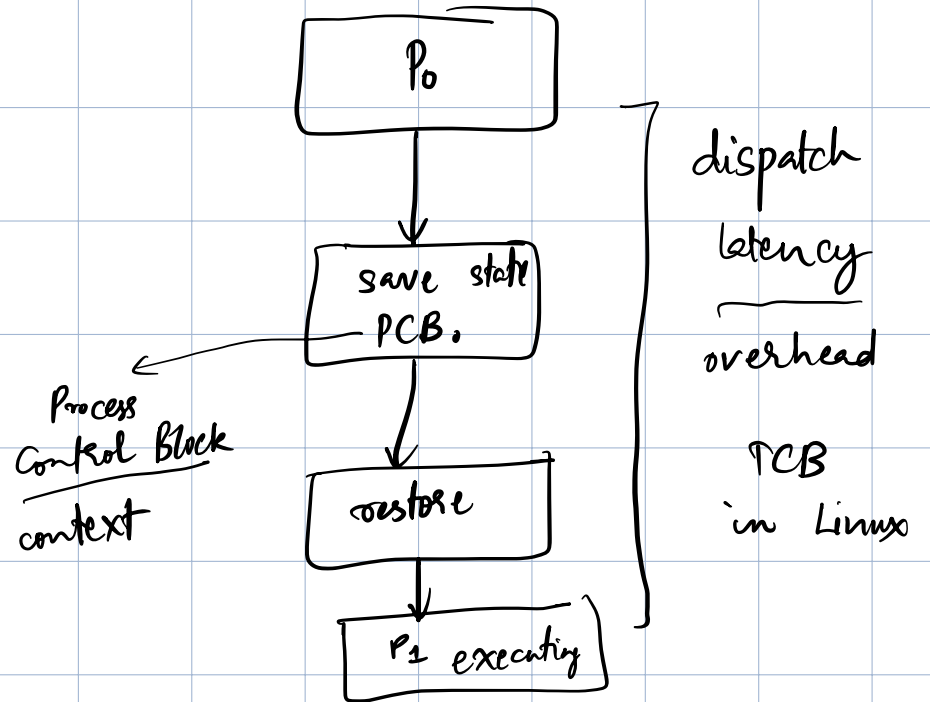
2 & 3 : preemptive

ready state ⇒ in queue

# Process state diagram

why ① non preemptive?  
→ I/O burst for process

## Dispatcher



## Scheduling criteria

→ CPU utilization — keep CPU busy

→ throughput — # of processes

→ Turnaround

→ Waiting

→ Response

→ when is time wasted?  
→ context switching

## FCFS

bad turnaround time

## SJF

→ theoretically best  
→ less predictable

estimation itself  
is an overhead


## Shortest Remaining Time First

→ starvation: smaller jobs keep arriving



Scheduler where?

scheduler is  
also a program  
→ runs on the  
same CPU




CPU

Ready queue



→ timer creates interrupt at regular intervals (time quantum)



ignored  
usually for  
simple analysis } scheduler  
switch

## Context switch time

depends on

→ data: larger PCB ⇒ more time

→ hardware


→ OS

## Scheduling

depends on

→ scheduling algorithm

Linux: BST



quantum

10

5

1

context switches

0

1

9

80% CPU burst  $> q$   $\rightarrow$  time quantum

fork and pthread\_create  $\rightarrow$  create new process

$\downarrow$

system call to invoke scheduler  $\rightarrow$  parent has to wait due to overhead

$\downarrow$

called by a process running in the CPU

In Windows : pthread\_create : lighter than fork

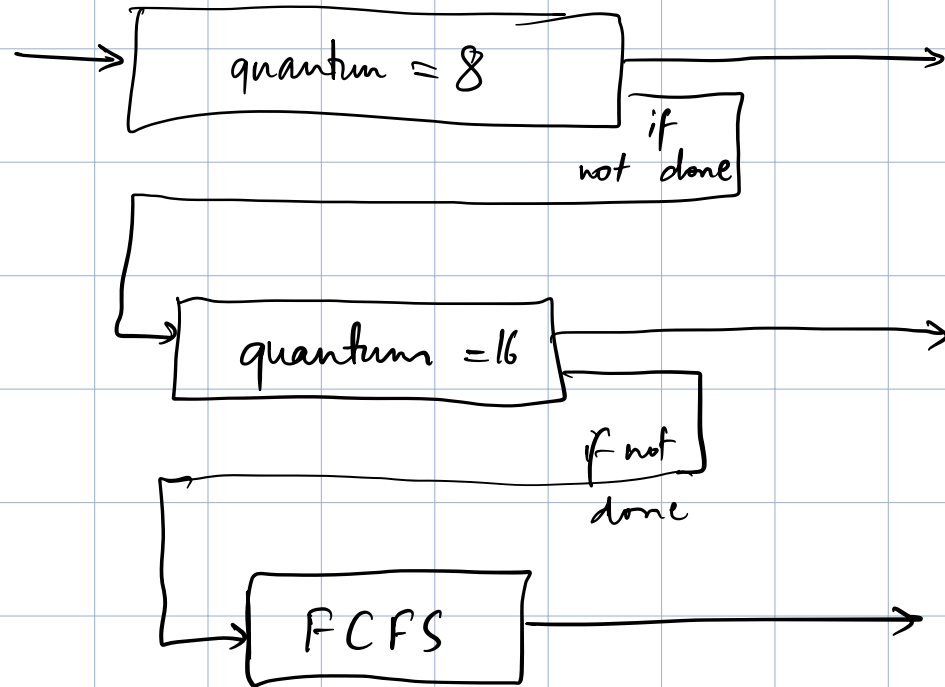
### Priority scheduling

- $\rightarrow$  starvation : low priority processes
- $\rightarrow$  solution : aging
- $\rightarrow$  RR : if same priority

## Multilevel Queue

→ prioritize based on process type.

→ each queue can have different scheduling algorithm

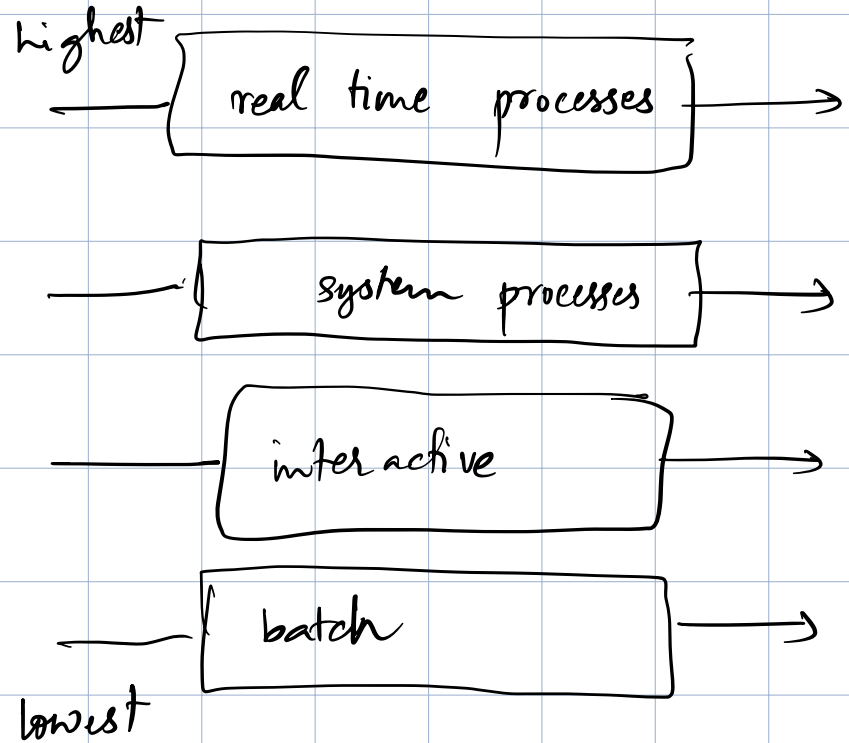


deadline → railway line  
Hard → gates  
→ ABS

Multilevel queue is

defined by

- no. of queues
- scheduling algorithm for each queue





# Single process scheduling


Thread scheduling  $\rightsquigarrow$  distinction b/w user level and kernel

$\rightarrow$  API

level threads

$\rightarrow$  M-to-one & M-to-M models,  
thread library schedules user-level  
threads to run on LWP

$\rightarrow$  process contention scope



→ Pthread Scheduling API

Next class: multiple cores scheduling