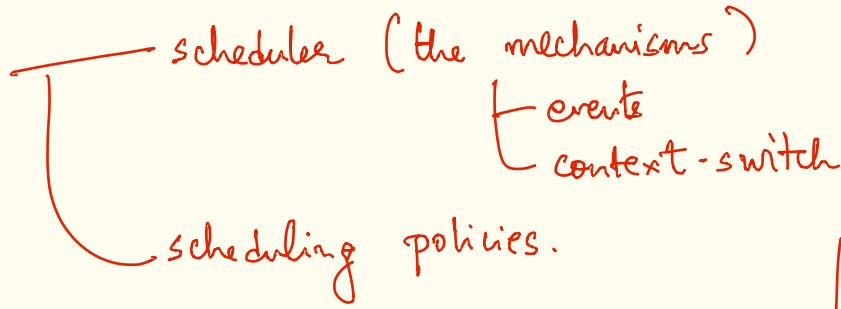


\* scheduling



execution time

\* (i) metrics :

- utilization
- fairness (no starvation)
- finish time
- waiting time
- response time
- throughput
- # deadlines met.

service time  
|  
actual time to do work.

yield()

(ii) categories :

(i) non-preemptive vs.

Schedule only when current process yields, exits or blocks.

preemptive


Schedule a new process even if current process is not blocked/done (basically wants to continue running).

(ii) non-work conserving vs.

if resource is available & work to be done, may not schedule work on resource.

work conserving.

if resource is available & work is to be done, schedule work, no resource idling.



(iii) policy criteria / challenges.

- interactive / bursty vs. CPU-intensive
- real time vs. best effort

## # Examples :

### ① FCFS : first-come-first-serve

- add to end of ready queue
- dequeue from head of queue. (on block (exit/yield))

⊗ work-conserving, no starvation, waiting time could be high!  
(~ non-preemptive)

### ② RR ~ Round Robin

~ preemptive at every time interval.

- dequeue from head of ready queue & execute for a specific max. time interval.
- on time-out or block (move to end of queue).

↳ dequeue next ready process.

⊗ work-conserving, no starvation,  
pre-emptive.

good for: fairness, response time.

time interval  $\rightarrow 0$  fair scheduler  
(context switch overheads increase)  
high.  
time interval  $\gg 0$  (large)

$\Rightarrow$  FCFS.

Execution & finish time can suffer (esp. at high load).

### ③ Priority-based scheduler

- ~ a separate ready queue per priority level.
- ~ each process has a tagged priority.
- ~ FCFS/RR used for each ready queue.

e.g. priority levels  
└─ gold  
└─ silver  
└─ bronze

⊗ ~ always schedule process from higher priority first!

⊕ pre-emptive, priority sensitive, (higher priority processes get priority for CPU usage).

⊖ starvation! - lower priority process - may never get CPU.

## ④ Proportionate Scheduling.

### WRR

weighted round robin

- one ready queue per priority.
- time slice per queue (weight) is proportional to priority.

e.g: 4:2:1 for 3 queues.

↳ 4 tasks of priority 1

2 tasks of " 2

1 task of " 1

and repeat (round robin).

⊕ priority sensitive  
no starvation.

⊖ finish time depends  
on weights & load.

### Lottery Scheduling

- # tickets issued to each process.
- random number generated to select winner / ticket  $\Rightarrow$  process.
- # tickets  $\Rightarrow$  probability of scheduling.
- # tickets  $> 0 \Rightarrow$  no starvation!

⊕ probabilistic & proportionate scheduling.

more policies:

- SJF: shortest / smallest job first.
- STCF: Smallest Time to Completion first.  
Shortest

## ⑤ real world

- arrival of processes (in ready queue) non-deterministic
- process execution time unknown. (and execution times are different for different processes)
- process behaviour (interactive vs cpu intensive) unknown.

⊕ real scheduler example

(i) # multi-level feedback queue based scheduler. (MLFQ).

- part of Linux for a long time
- multiple queues (one per priority)
- priority of each process is dynamic / changing.

## # the MLFQ scheduling policy —

1. if  $\text{priority}(\text{process A}) > \text{priority}(\text{process B})$   
A gets CPU.

2. if  $\text{priority}(A) = \text{priority}(B)$ , A & B run in RR fashion  
using a time slice for the given queue.

3. a process entering the system, Placed at highest priority  
(READY for first time) (topmost queue).

4. Once process uses time slice/allotment at a given level  
its priority is reduced (moved down one queue level).

5. After some time period, move all processes to topmost queue  
or wrap around after last queue

\* note: each queue has a different time slice for RR.

(+) no starvation, preemptive, priority sensitive & fair.

short jobs finish quickly, CPU-intensive make regular progress.  
long jobs

---

(ii) CFS ————— current Linux scheduler  
Completely Fair Scheduler. (default for non real time tasks).

(iii) multi CPU (multi processor) scheduling.

+ multiple queues (one per CPU)

+ which queue a process should be placed in? (load balancing).

+ should process move across queues? (migration).