

# workloads

- batch / computational
  - bursty + io
- support both

in these real workloads: do we really know time (#4?)  
do we really know what kind of job?  
(maybe we do! but unix didn't)

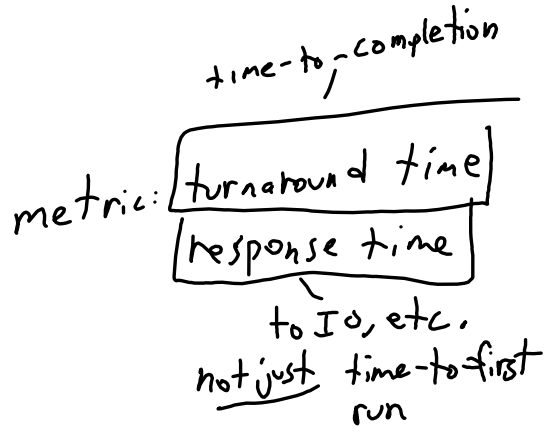
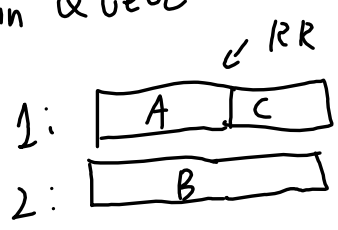
algorithm: MLFQ

Multi-level -  
feedback - queuing

idea: learn type of job over time

Queues with priority  
(DMV, doctor office, etc)  
MVC

RR within Queue



## Side bar / review : RR

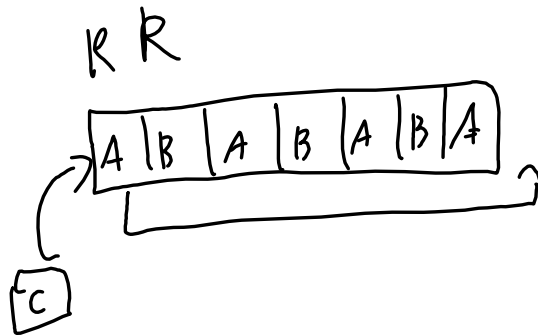
- slice jobs into individual quanta
- (how to set quanta? magic number!  
100 us = 0.1ms on courselab (minimum))  
↳ side bar: /proc/sys/kernel, check it out
- balance responsiveness vs pre-emption

RR: live scheduling?

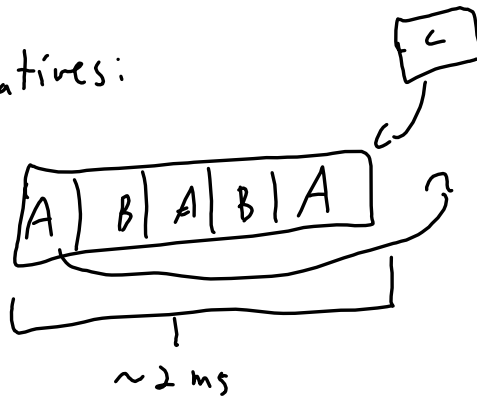
- book assumes fixed schedule

challenge: where to put new jobs?

last time:



alternatives:




sidebar: flexible quantum?

back to MLFQ

1 

Rules:

# 1: run high-priority first

2 

# 2: RR within queue

remember: we have to implement  
what DS? what are we tracking?

Challenge: where to put jobs? How to set priority?

More rules: # 3: enter at top

# 4a: demote on timer

# 4b: keep on yield

(no new info to track!)

problems? (starvation) (game the sched)

benefits? (responsiveness)

---

try again: add: rule # 5: merge into top every N us.

what does this solve?

try again: what's better? what's wrong?  
(gaming)

fix gaming: track more state.

replace rule # 4 a, b!

Full rules:

# 1: run high-priority  
first

# 2: RR within queue

# 3: enter at top

# 4: demote after allotment used

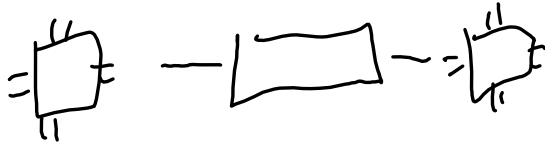
# 5: merge into top every  $N$  us.

- discuss: parameterization!

Nice: user-level control

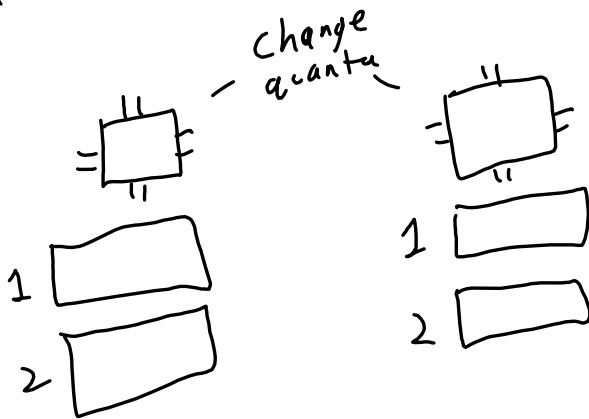
- how long is a  $Q_i$ ? can we enforce this?
- what if we know interactivity?

Modern-day: multicore  
(NOT on the exam)

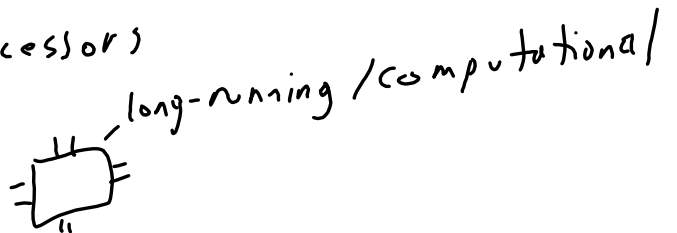
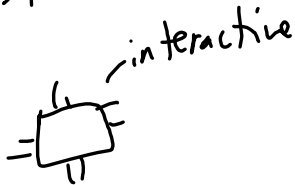


- sched per core?
- cores with specific purpose?

Multi-core MLFQ



"specialize" processors



example: compute clusters

timer: core migration (configurable!)  
look at /proc/sys/kernel

# A fun alternative: lottery

- use randomness!

- process launch: gain  $N$  tickets (random)

- sched: select  $k$  random tickets  
(with replacement)

Benefits:

- ticket transfer
- low state tracking

- subdivide tickets  
↳ give tickets to  
users, e.g.

limitations:

- unpredictable
- source of randomness

---

CFS! (probably won't get here)

- use a min-heap