


Non-preemptive scheduling




non-presumptive

- FIFD


- SJF

- priority

System calls



this is
still P_1
execution



System call

write to disk

call

handler

wait for disk
to write to disk
to disk

RTI ?

software interrupt

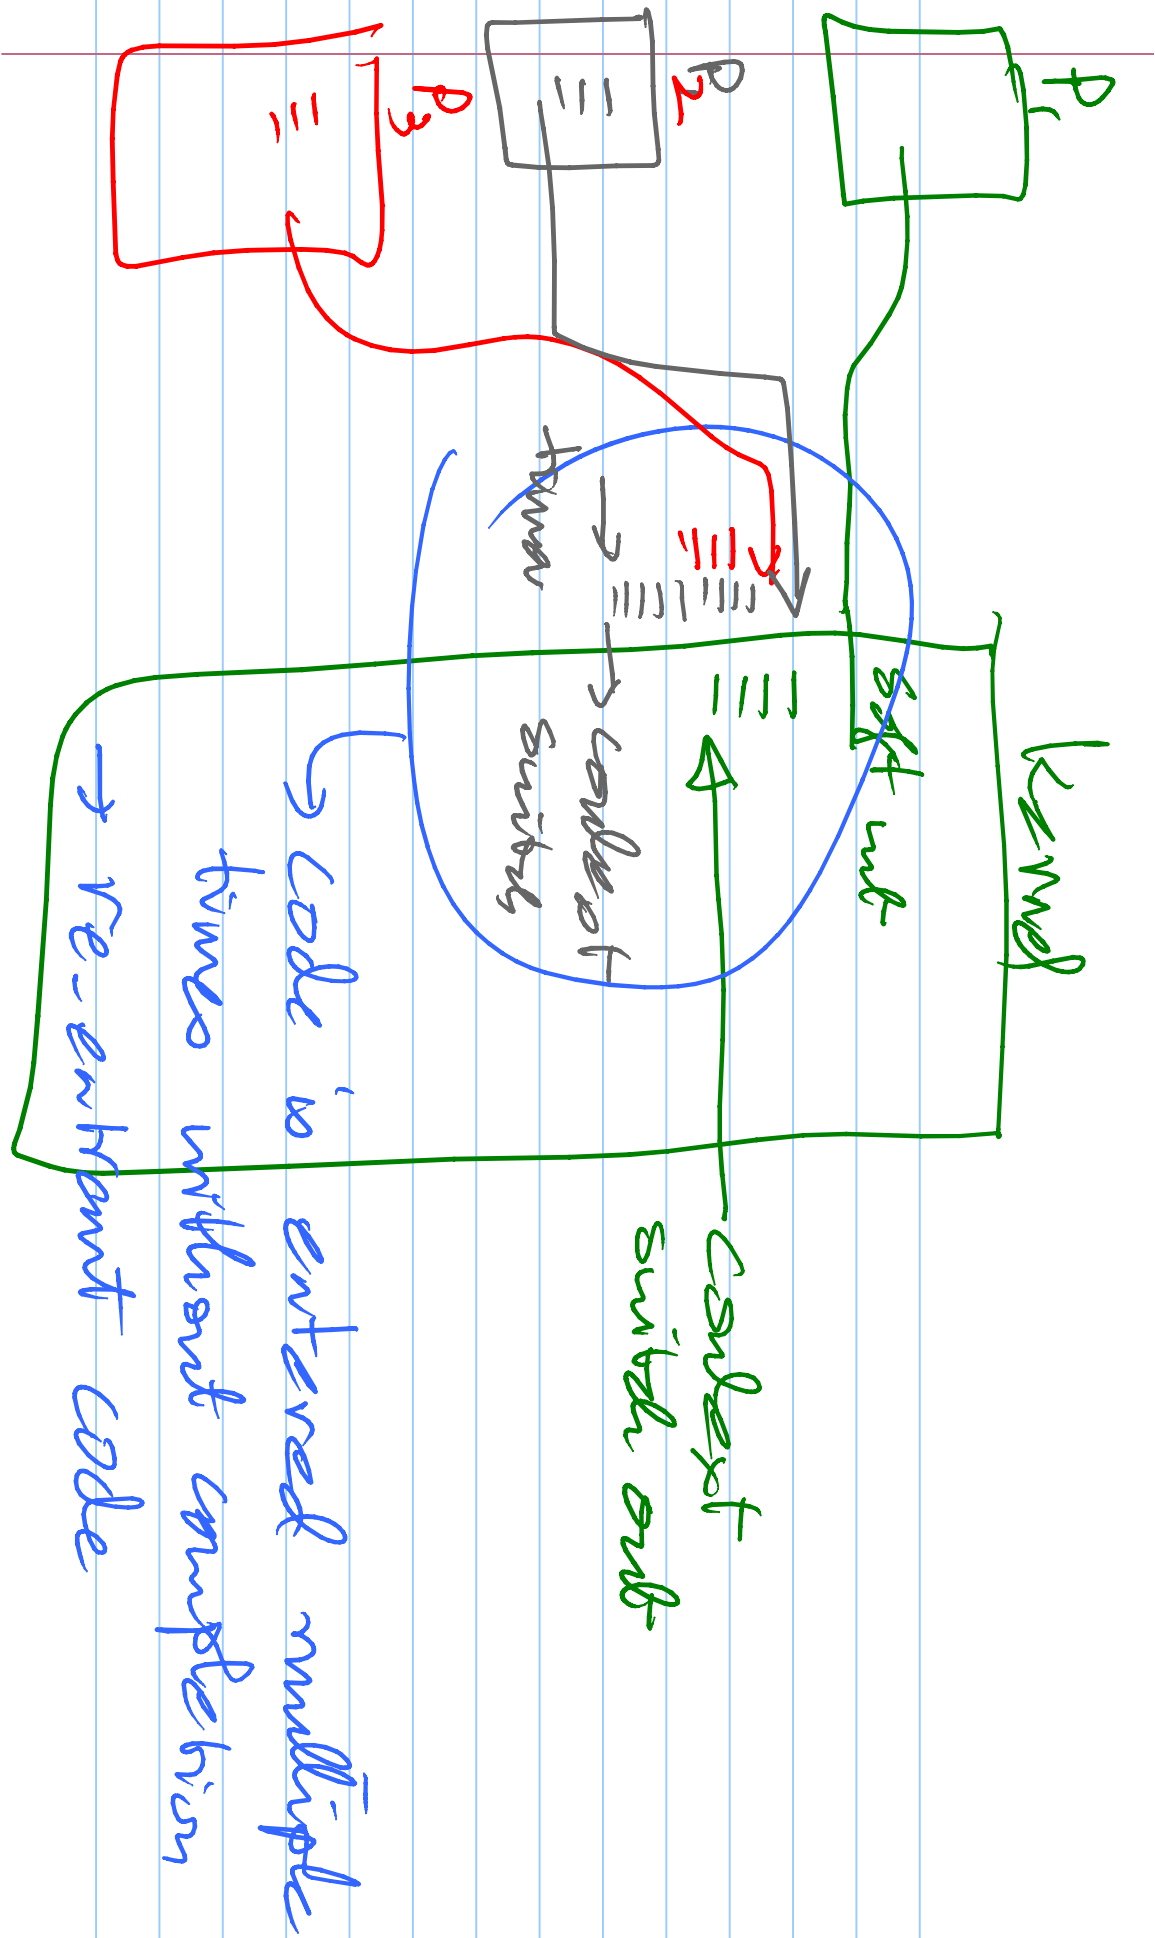
trap

wait
for I/O

call
sched
to context
switch


take the
P/B out of
run &
next process
RT1


→ disk int handler will print P1
back into ready &



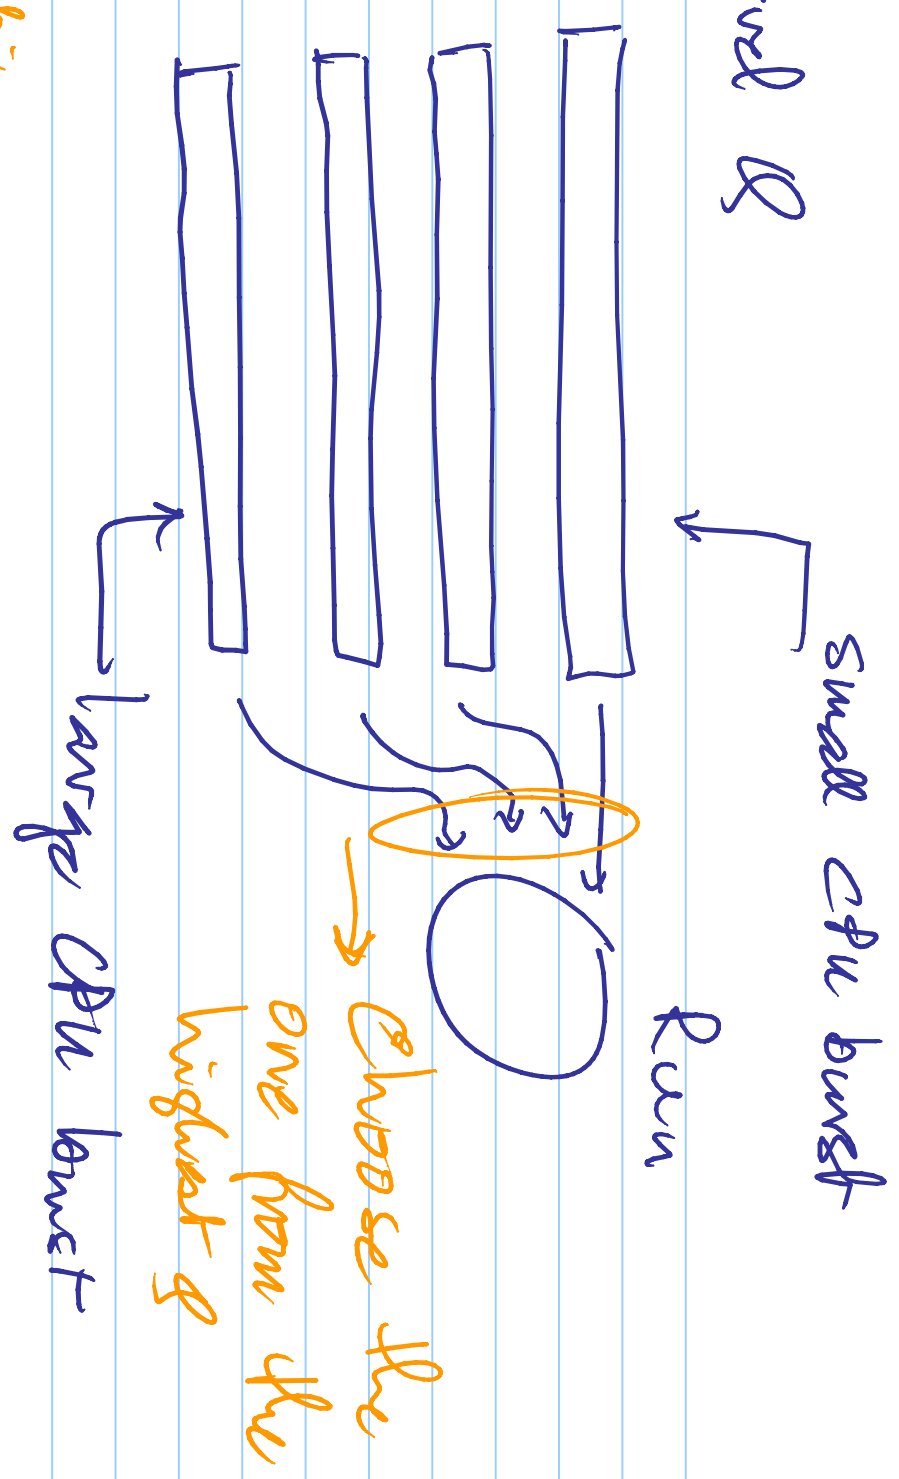
Preemptive Scheduling

① Round Robin





multilevel Q




- Starvation

- Classification of apps

Multitasking Feedback ϕ \rightarrow add copying ϕ \rightarrow more old


New

Run process up



Time Quanta

Small \rightarrow overhead dominates



Response time

→ average time taken
for a process to start
running on CPU after it
gets ready

→ (avg length of Ready) × (quanta)

Windows Scheduler

32 level MCF8 - 1-15 for regular
16-32 for real time


Unix Scheduler ↓ low

NICE → -20 to +19

↑ high
↑ D is normal

Linux 2.4 Epoch scheduler

3 classes interactive, batch, real time



all processes get

1 time $Q \Rightarrow$ also time Q
is different...

Linux 2.6 (original)

- 140 sysns, 100 v.t., 40028
- Active ¹⁸ expired 9
- priority recalculations

Linux 2.6 CFS

↳ Completely Fair Scheduler

→ need black box

O(hgn) to pick / rebalance hze

? key = CPU time used

"

Scheduler Simulations

① process

② queuing models

③ simulation

④ implementation (+ instrument)