

> Goal: want to


- > run shorter "interactive" jobs first
- > handle I/O bound jobs well
- > also handle CPU bound jobs

=> dynamically try to approx. STCF

> how: priority-based methods

> many out there

> Basic idea



- 1) if $P(A) > P(B)$, A will run
- 2) Among jobs A, B s.t. $P(A) = P(B)$, use round-robin (or similar)

Fall '04
do last

=> [3) Processes priority may vary over time]

4) Time-slice may vary per queue

> Key: How to decide when priority should change?


Goal: short, interactive ↑
long, CPU-bound ↓

Attempt # 7:

- > Start @ top
- > if give up CPU before end of quantum, stay at ~~top~~ same level
- > if use full quantum, move down 1 notch

when good?

- > 1 long-running job has been in system
- > new job comes along
interactive, CPU bound?
 - ↓ will be served first!
 - ↓ ok, will slowly move down, eventually settle



why bad?

- 1) can "game" scheduler
- 2) can lead to starvation

Fall '04

3) job that changes its behavior:
not accounted for

ML #2:

[let them figure these out]


ML-3

> periodically, move ~~it~~^{job} back to top queue

⇒ no more starvation, PRUT still can game

ML #3:

> don't reset quantum counter upon I/O



⇒ no I/O move gaming (no value to it)

Real Implementations

BSD: "Berkeley unix"

⇒ formula for priority

"exponential decay"

SUR4: "AT&T 'unix'"

table-driven approach;

sys admin can tune

NT?!

conclude

? Complex but adaptive scheduler

> works well for many workloads

? How to control what's going on?

Lottery

Problems w/ multi-level feedback

Little control to apps (fixed % of CPU)

Hard to understand parameters

Per job, not per user

Proportional-share scheduling concept

New mechanism: tickets

each user (or process) may get some #

e.g. 2 clients, A, B

A: 100

$$100 / (100 + 200)$$

= $\frac{1}{3}$ of resources

B: 200

$$200 / ()$$

= $\frac{2}{3}$

CPU

Currencies: (can skip)

Users have fixed # of tix in base currency
Can distribute tix in their own user currencies

E.g.

A: 400 \Rightarrow A₁

600 \Rightarrow A₂

[100]

B: 50 \Rightarrow B₁

30 \Rightarrow B₂

20 \Rightarrow B₃

[200]

$$A_1 = \frac{400}{1000} \cdot 100 \Rightarrow \underline{\underline{40}}$$

$$A_2 = \frac{600}{1000} \cdot 100 \Rightarrow \underline{\underline{60}}$$

$$B_1 = \frac{50}{200} \cdot 200 \Rightarrow \underline{\underline{50}}$$

$$B_2 = \frac{30}{200} \cdot 200 \Rightarrow \underline{\underline{30}}$$

$$B_3 = \frac{20}{200} \cdot 200 \Rightarrow \underline{\underline{20}}$$


Implement: PS \Rightarrow Lottery

(L2)

Scheduling decision: hold a lottery!

Winner gets time-slice

How does that work?




\Rightarrow more tickets, more likely to win

random #
between 0 \rightarrow 299 inclusive

How to implement?

> List of jobs



pick $R = \text{random}(1 \rightarrow 300)$

count += fix(ptr)

if (count \geq R) {

winner(ptr);

}

> How to order list for

shortest traversals?

> Could also do deterministically
(think about how)

(L3)

> Positives

Could apply to many resources
user

Proportional share w/ control

Simple to implement

> Negatives

Not general purpose (I/O, interactive)

Projects:

Extra office hours

tomorrow: will send out time,



Monday?

email: cs537-help@cs

at any time for help

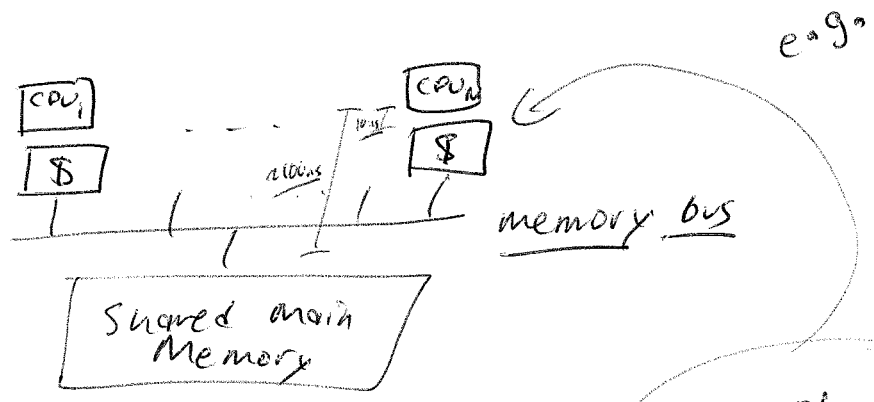
Multiprocessor Issues

M₁



> How to allocate P => C ?

> Hardware



- Usually
- > Small # of CPUs
 - > Uniform access time to memory
 - > Single OS

> caches

Approach 1: Per CPU queue

M₂



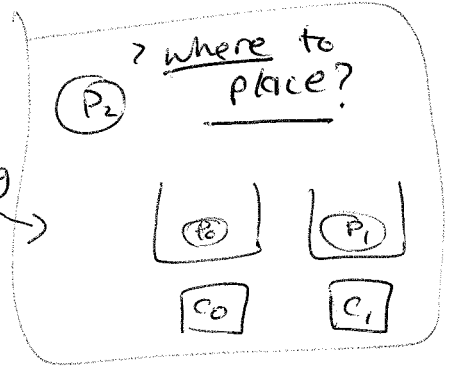
> Process enters, pick queue and stay there forever

Positive:

- > simple
- > Run on same processor => use same cache

Negative:

- > Load imbalance could occur (lose natural advantage of SMP)




Approach 2: [Global Queue] ← do this first (Fall '04)

- > generalize MLFQ
- highest n jobs run

Positive
> allows dynamic load balancing

Disadvantage

- > single queue: contention! might limit scalability
- > might not be cache sensitive




Approach: Hybrid

M₃

> Both local + global queues

> usually have jobs in local queue

> peek in other queues occasionally



> Load balance


> if none in local queue, look elsewhere in ready queues (job stealing)

> use global queue for kernel priority threads

> Processor Affinity

> Try to keep Process on same CPU

Parallel processes



> might want to comm. w/ one another

> Coordinate CPU switches

global ctxt switch => co-scheduling