UNIVERSITY OF CALIFORNIA AT BERKELEY
College of Engineering
Department of Electrical Engineering and Computer Sciences
Computer Science Division

CS 162
Fall 1988

John Ousterhout

Topic 9: CPU Scheduling

- Readings for this topic:
  - Silberschatz/Peterson: Chapter 4.

- Until now you have heard about processes. From now on you'll hear about *resources*, the things operated upon by processes. Resources range from cpu time to disk space to channel I/O time.

- Resources fall into two classes:
  - Preemptible: processor or I/O channel. Can take resource away, use it for something else, then give it back later.
  - Non-preemptible: once given, it can't be reused until process gives it back. Examples are file space, terminal, and maybe memory.

  This distinction is a little arbitrary, since anything is preemptible if it can be saved and restored. It generally measures the difficulty of preemption.

- OS makes two related kinds of decisions about resources:
  - Allocation: who gets what. Given a set of requests for resources, which processes should be given which resources in order to make most efficient use of the resources? Implication is that resources aren't easily preemptible.

- Scheduling: how long can they keep it. When more resources are requested than can be granted immediately, in which order should they be serviced? Examples are processor scheduling (one processor, many processes), memory scheduling in virtual memory systems. Implication is that resource is preemptible.

- Resource #1: the processor.

- Processes may be in any one of three general scheduling states: *Draw state diagram, go over which transitions are and are not possible.*
  - Running.
  - Ready. That is, waiting for CPU time. Scheduler and dispatcher determine transitions between this and running state.
  - Blocked. Waiting for some other event: disk I/O, message, semaphore, etc.

- Goals for scheduling disciplines:
  - Efficiency of resource utilization. *(keep CPU and disks busy)*
  - Minimize overhead. *(context swaps)*
  - Minimize response time. *Define response time.*
  - Distribute cycles equitably. What does this mean?

- FIFO (also called FCFS): run until finished.
  - In the simplest case this means uniprogramming.
  - Usually, ‘‘finished’’ means ‘‘blocked’’. One process can use CPU while another waits on a semaphore. Go to back of run queue when ready.
  - Problem: one process can monopolize CPU.

- Solution: limit maximum amount of time that a process can run

without a context switch. This time is called a *time slice*.

- Round Robin: run process for one time slice, then move to back of queue. Each process gets equal share of the CPU. Most systems use some variant of this. What happens if the time slice isn't chosen carefully?
  - *Too long: one process can monopolize the CPU.*
  - *Too small: too much time wasted in context swaps.*

- Originally, Unix had 1 sec. time slices. Too long. Most systems today use time slices of 10,000 - 100,000 instructions.

- Implementation of priorities: run highest-priority processes first, use round-robin among processes of equal priority. Re-insert process in run queue behind all processes of greater or equal priority.

- Even round-robin can produce bad results occasionally. Go through example of ten processes each requiring 100 time slices. *They each take about 1000 time slices to finish, whereas in FIFO they would average 500 time slices to finish. This is like government: it is fair, but uniformly inefficient.*

- What is the best we can do? STCF: shortest time to completion first with preemption. This minimizes the average response time.

- *Lecture 17 ended here.*

- As an example, show two processes, one doing 1 ms computation followed by 10 ms I/O, one doing all computation. Suppose we use 100 ms time slice: I/O process only runs at 1/10th speed, I/O devices are

only utilized 10% of time. Suppose we use 1 ms time slice: then compute-bound process gets interrupted 9 times unnecessarily for each valid interrupt. STCF works quite nicely.

- Unfortunately, STCF requires knowledge of the future. Instead, we can use *past* performance to predict *future* performance. *Joke: man who normally helps little old ladies across street will probably continue to do so; of course, occasionally he may show up with a machete and hack them all to pieces....*

- *If a process has already taken a long time, it will probably take a long time more. We use priorities to disfavor long processes.*

- Exponential Queue (or Multi-Level Feedback Queues): attacks both efficiency and response time problems.
  - Give newly runnable process a high priority and a very short time slice. If process uses up the time slice without blocking then decrease priority by 1 and double time slice for next time.
  - Go through the above example, where the initial values are 1ms and priority 100.
  - Techniques like this one are called *adaptive*. They are common in interactive systems.
  - The CTSS system (MIT, early 1960's) was the first to use exponential queues.

- Fair-share scheduling (similar to what's implemented in Unix):
  - Keep history of recent CPU usage for each process.
  - Give highest priority to process that has used the least CPU time recently. Highly interactive jobs, like vi, will use little CPU time and get high priority. CPU-bound jobs, like compiles, will get

lower priority.

- Can adjust priorities by changing ‘‘billing factors’’ for processes. E.g. to make high-priority process, only use half its recent CPU usage in computing history.

- Summary:
  - In principle, scheduling algorithms can be arbitrary, since the system should produce the same results in any event.
  - However, the algorithms have strong effects on the the system’s overhead, efficiency, and response time.
  - The best schemes are adaptive. To do absolutely best, we’d have to be able to predict the future.

- *Best scheduling algorithms tend to give highest priority to the processes that need the least! Reagonomics?*