

### Administrivia

- I say in the syllabus that I try to respond promptly to e-mail. Exceptions are minute essays and homeworks, which I don't always look at right away. If you need a quick reply, make that apparent on the subject line please!

Slide 1

### Minute Essay From Last Lecture

- Difficult: Getting re-acclimated to C. (That was partly the point of the assignment — one more opportunity to practice with the language.) Also figuring out exactly what parameters to pass to the system-call library functions, and doing string parsing in C.
- Interesting: How many system calls are needed for even simple programs. (That also was partly the point of that problem.)
- (There were some other responses, but these stood out.)

Slide 2

Slide 3

### Words of Wisdom?

- A very smart person I know once said the only interesting part of an o/s course was concurrent algorithms (to be covered soon), and the rest is “just details”.

A student a few years ago said “a lot of this just seems like common sense” (once you understand the basic ideas).

Both sort of right . . .

- Goal of this course is to learn/retain basic ideas. Details may help with that — and can be interesting in themselves — but should not be the focus.
- (Both things to keep in mind as you continue reading and we continue discussing . . .)

Slide 4

### Process Abstraction

- We want o/s to manage “things happening at the same time” — applications, hidden tasks such as managing a device, etc.
- Key abstraction for this — “process” — program plus associated data, including program counter.
- True concurrency (“at the same time”) requires more than one CPU (more properly now, “more than one CPU/core”?). Can get apparent concurrency via interleaving — model one virtual CPU per process and have the real processor switch back and forth among them (“context switch”).  
(Aside: In almost all respects, this turns out to be indistinguishable from true concurrency. “Hm!”?)

### Process Abstraction, Continued

Slide 5

- Can also associate with process an “address space” — range of addresses the program can use. Simplifying a little, this is “virtual memory” (like the virtual CPU) that only this process can use. More (lots more) about this later. (Nitpick: Yes, we also want to be able to share memory among processes. More about that later too.)
- How to map this to the real hardware? in this chapter we talk about how to share the real CPU(s) among processes; in the next chapter we talk about how to share the real memory.

### Context Switches

Slide 6

- What is it? switch from one process to another.
- When should this happen?

### Context Switches, Continued

Slide 7

- Should happen
  - when a process's "time slice" is up.
  - when there's an unrecoverable error.
  - when there's something that needs to be done right away (e.g., deal with input/output).
  - maybe other times? (when a process has to wait for something, e.g.).All signalled by some kind of interrupt.
- Goal is to suspend work on a process such that we can later pick up exactly where we left off. How do we make that happen?  
(Think about what the hardware does when an interrupt happens, what's included in that "virtual CPU".)

### Context Switches, Continued

Slide 8

- On interrupt, hardware saves program counter (at least — why?), transfers control to fixed location — which contains o/s code.
- That O/S code has to
  - Save CPU state (program counter, registers, etc.) for the current process.
  - Deal with interrupt (details depend on type — I/O versus timer versus . . .).
  - Restore CPU state for "next" process (previously saved), thereby restarting it.  
(“Next” process? yes, o/s might have to choose — more about that later.)

### Process Creation and Termination

Slide 9

- When are processes created?
  - At system startup.
  - When another process makes a “create process” system call — e.g., to start a new application.
- When are processes destroyed?
  - At program exit.
  - After some kinds of errors.
  - When another process makes a “kill process” system call.

### Process States

Slide 10

- Can think of processes as being in one of three states:
  - “Running” — being executed by a CPU.
  - “Blocked” — waiting for something to happen (I/O to complete, another process to do something, etc.) and unable to do anything useful until it does.
  - “Ready” — not blocked, but waiting because all CPUs are currently executing other processes.
- Possible transitions? Which ones require decision-making?

### Process States, Continued

Slide 11

- Possible transitions (figure in textbook, p. 90):
  - Running to blocked — happens when, e.g., a process makes an I/O request and can't continue until it's complete.
  - Blocked to ready — happens when the event the blocked process is waiting for occurs.
  - Running to ready, ready to running — needed if we want some sort of time-sharing (give all non-blocked processes "a turn" frequently).
- Notice that moving to and from "blocked" state doesn't involve decision-making, but ready/running transitions do.
- The decision-maker — "scheduler" (to be discussed later). Often "running to ready" is triggered by an interrupt (I/O, timer, etc.), and "ready to running" involves this scheduler.

### Implementing Processes

Slide 12

- Think about how you would implement this abstraction . . .
- First, you'd want a data structure to represent each process, to include — what?

### Implementing Processes, Continued

Slide 13

- Data structure to represent each process would include some way to represent such things as:
  - Process ID.
  - Process state (running / ready / blocked).
  - Information needed for context switch — a place to save program counter, registers, etc.
  - Other stuff as needed — e.g., a list of data structures for open files.
- Then you'd collect these into a table (or some similar structure) — “process control table”, with individual data structures being “entries in the process control table” or “process control blocks”.

### Implementing Processes, Example — Linux

Slide 14

- Each process (“task”) is represented by a C `struct` containing information similar to what we described.
- These `structs` are chained as a doubly-linked list; there is also a hash table keyed by PID.
- (This is according to online information about the 2.4 kernel.)

## Processes Versus Threads

Slide 15

- So far I've used "process" in an abstract/general way.
- In typical implementations, though, "process" is more specific — something that has its own address space, list of open files, etc. Often these are called "heavyweight processes".
  - Advantages — such processes don't interfere with each other.
  - Disadvantages — they can't easily share data, switching between them is expensive ("a lot of state" to save/restore).
- For some applications, might be nice to have something that implements the abstract process idea but allows sharing data and faster context switching — "threads".  
(To be continued.)

## Minute Essay

Slide 16

- None — quiz.