

Slide 1

### Administrivia

- Reading Quiz 1 assigned. Linked from schedule page. Due in a week. Turn in by putting a PDF or plain-text file in the folder for this assignment in the Google Drive TurnIn folder I set up for you.

Slide 2

### Processes and “Virtualizing the CPU”

- One thing we want an operating system to make possible is having multiple applications “open” at the same time, possibly more than we have processors to run them.  
Useful to think of each “open application” as an instance of “process abstraction”.
- (Aside: Terminology can be confusing. Here, “processor” means anything that can execute a sequence of instructions, and can be a totally independent chip or one “core” on a multicore chip. An older term is “CPU”. Possibly no standard term. A while back my research collaborators and I decided to use “processing element”, but it hasn’t caught on.)
- How? “Virtualize the CPU”.

### Virtualizing the CPU

Slide 3

- Big picture is that we want to have (conceptually) arbitrarily many processes, each with an “address space” (memory it can use — so, yes, we will have to “virtualize memory” as well, which is the next big topic). Need to “time-share” CPU, “space-share” memory.
- The “crux” of the problem is how to implement this, with reasonable efficiency. A lot about this is done should seem kind of like common sense, if you keep this in mind.

### Process — Abstraction

Slide 4

- In CS terms, define an abstraction in terms of possible values, operations.
- Possible values here are a little complicated — some representation of the state of the process.
- Many operations possible; some basic ones include “create”, “destroy”, “wait”, and “signal”.  
Encapsulated as “process API” / set of system calls.

Slide 5

### Process State — Implementation

- Key here is to include everything important-in-context, so includes:
- Machine state (e.g., contents of registers). (Might be a good time to reflect on what you remember from CSCI 2321.)
- Information about address space — e.g., where it is in physical memory. Curiously enough, its *contents* aren't part of this state! (Which is okay, because we might want another process to change something.)
- List of open files.

Slide 6

### Process Creation — Implementation

- Lots involved here. Interesting to reflect on this in the context of what you know about program startup from application perspective? Some things:
- Load program from disk. Details vary, but code has to be in memory to run, not on disk. Includes initializing data defined at startup.
- Initialize stack, so stack-based schemes for function calls work.
- Initialize heap, so we can do `malloc` and equivalents.
- Put any arguments on stack.
- Other stuff — e.g., for UNIX, set up some standard “open files”.
- Start program (C `main()` or the equivalent).

### Process Creation — Implementation, Continued

- Textbook shows semi-real-world example of data structure (figure 4.5).  
Note that this is for an architecture similar to x86, where the registers all have these I-think-ugly semi-symbolic names. Contrast to MIPS's simple r0, r1, etc.!

Slide 7

### Some Specifics — UNIX

- Chapter on processes in UNIX is, I think, interesting. Not all operating systems work exactly like this of course but same basic ideas / principles.
- One noteworthy thing — process creation. Other systems *don't* do it the same way (e.g., Windows has a single "spawn process" system call), and it can be a topic for heated discussion!

Slide 8

### Process Creation in UNIX

Slide 9

- Unusual in needing not one but two system calls:
- `fork()` to create a new process — which is an almost-exact copy of the calling process! including its full address space. Only difference is return code from `fork()` function (zero in “child” process, process ID of child in “parent”, or negative value if operation fails).
- `exec` functions (several options) load new code into process, replacing current code.
- Why oh why? can be useful to retain some things such as list of open files.
- Big potential performance hit; think a minute about what it is before going on.

### Process Creation in UNIX, Continued

Slide 10

- Duplicating address space — probably fine when `fork()` was invented and memory was limited. But now?
- Mitigate by not actually making copy all at once — “copy on write” scheme.

### Minute Essay

- Questions? Is this material making sense?
- Did you learn (some) MIPS assembly in CSCI 2321? (Last year was the first time in a *long* time I didn't teach the course and so don't really know what students were asked to do in the prereq course.)

Slide 11