Administrivia

- Reminder: Homework 1 written problems due today, programming problem due Monday.

- Reminder: Quiz 1 Monday. Short, topic(s) from Chapter 1. "Open book, open notes" — okay to use textbook (in whatever format you have), course Web site, your own notes (again in whatever format you have). No other computer access (so, no using search engines, etc.)

Slide 1

Minute Essay From Last Lecture

- (Review.)

Slide 2

O/S Versus Application Programs — Recap/Review

- Should seem reasonable to make distinction between what O/S can do and what application programs can do.

- But how to enforce that? i.e., how to make it as difficult as possible for buggy or malicious application programs to do what they shouldn't?

Slide 3

Can this problem be solved completely by clever programming? Consider that most current systems can be asked to load and execute machine-level application code . . .

O/S Versus Application Programs, Continued

- If you don't allow that — how do you decide what's okay?

- If you do allow loading and executing arbitrary code, then some sort of hardware mechanism for limiting what it can do seems like the only way. This is the problem "dual-mode operation" is intended to solve.

Slide 4

O/S Versus Application Programs, Continued

Slide 5

- At hardware level, then, need to keep track of which mode we're in and use that information to allow/disallow certain operations (and maybe memory accesses — though that could be a separate problem/solution).

- To do this efficiently — single bit in a register somewhere, probably a special-purpose one, checked by "privileged" instructions.

- What happens if unprivileged program tries . . . ? Hardware version of exception — interrupt.

- How to set this bit? privileged operation, or no?

O/S Versus Application Programs, Continued

Slide 6

- A solution: Include instruction to generate interrupt, and have hardware, on interrupt, transfer control to a fixed location *and* set the "privileged" bit. If what's at the fixed location is O/S code, then it can do more checking (e.g., passwords).

- What if it's not O/S code?

O/S Versus Application Programs, Continued

Slide 7

- So maybe we need memory protection too? but we probably needed that anyway.

- How to make memory protection work? more about that later, but for now — again, seems like the only way to do this reliably and efficiently is with help from hardware.

System Call / Interrupt Processing — Recap/Review

Slide 8

- Recall(?) typical mechanism for regular program calls: Put parameters in agreed-on locations (registers, stack, etc.), issue instruction that saves current program counter (in another register maybe) and transfers control to called program. Called program returns using saved program counter.

- System calls are similar *except* that the "called program" is at a fixed address *and* the transfer of control also puts the processor in supervisor/kernel mode.

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.

Slide 9

- True concurrency ("at the same time") requires more than one CPU/processor/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

- 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.)

Slide 10

- 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 11

- What is it? switch from one process to another.

- When should this happen?

Context Switches, Continued

Slide 12

- 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".)

Slide 13

Context Switches, Continued

- 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.)

Slide 14

Process Creation and Termination

- 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.

Slide 15

Process States

- 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?

Slide 16

Process States, Continued

- 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 17

- 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 18

- 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

- Each process ("task") is represented by a C `struct` containing information similar to what we described.

- These `struct`s are chained as a doubly-linked list; there is also a hash table keyed by PID.

Slide 19

- (This is according to online information about the 2.4 kernel.)

Processes Versus Threads

- 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.

Slide 20

  - 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".

Threads

- So, threads are another way to implement the process abstraction.

- Typically, a thread is "owned" by a (heavyweight) process, and all threads owned by a process share some of its state — address space, list of open files.

Slide 21

- However, each thread has a "virtual CPU" (a distinct copy of registers, including program counter).

- Implementation involves data structures similar to process table.

- Advantages / disadvantages (compared to processes)?

Threads, Continued

- Advantages: threads can share data (same address space), switching from thread to thread is fairly fast.

- Disadvantages: sharing data has its hazards (more about this later).

Slide 22

Implementing Threads

Slide 23

- Two basic approaches — "in user space" and "in kernel space" Various hybrid schemes also possible.

- Basic idea of "in user space" — operating system thinks it's managing single-threaded processes, all the work of managing multiple threads happens via library calls within each process.

- Basic idea of "in kernel space" — operating system is involved in managing threads, the work of managing multiple threads happens via system calls (rather than user-level library calls).

- How do they compare?. . .

Implementing Threads, Continued

Slide 24

- Implementing in user space is likely more efficient — fewer system calls.

- Implementing in kernel space avoids some problems, though:

  - If a thread blocks, it may do so in a way that blocks the whole process.

  - Preemptive multitasking is difficult/impossible without help from the kernel, as is using multiple CPUs.

Adding Multithreading

- If you've written multithreaded applications — moving from single-threaded to multithreaded not trivial:
  - Figure out how to split up computation among threads.
  - Coordinate threads' actions (including dealing properly with shared variables).
- Similar problems in adding multithreading to systems-level programs:
  - Deal properly with shared variables (including ones that may be hidden).
  - Deal properly with signals/interrupts.

Slide 25

Implementing Threads, Example — Linux

- Early versions of Linux provided no support for kernel-space threading, but there were libraries for the user-space version.
- More-recent kernels provide support, but in an interesting way — threads in some ways are just processes with with some different flags allowing them to share memory, etc.

Adding support for threads complicates process creation — the basic mechanism (`fork`) duplicates an existing process, and if that process is multithreaded, things can be interesting. Some details in chapter 10, or read the POSIX standard for `fork`.

Slide 26

Interprocess Communication

- Processes almost always need to interact with other processes:

    - "Ordering constraints" – e.g., process B uses as input some data produced by process A.

    - Use of shared resources — files, shared memory locations, etc.

Slide 27

- Use of shared resources can lead to "race conditions" — output depends on details of interleaving.

- Processes must communicate to avoid race conditions and otherwise synchronize.

- "Classical IPC problems" — simplified versions of things you often want to do.

Minute Essay

- In a system with 8 CPUs and 100 processes, what are the maximum and minimum number of processes that can be running? ready? blocked?

Slide 28

Slide 29

Minute Essay Answer

- Blocked: Maximum of 100 (unless you assume that there's an "idle" operating system process that runs when nothing else does and never blocks, and maybe one of these is needed for every CPU). Minimum of 0.

- Running: Maximum of 8, because there are 8 CPUs. Minimum of 0 (again unless you assume that there's an o/s process that runs when nothing else does).

- Ready: Maximum of 92, since all CPUs will be running processes if there are any that can be run. (Depending on details, you might have to add "except during context switches, when the scheduler is choosing the next process to run on a CPU".) Minimum of 0, since they could all be blocked or running.