

Slide 1

Administrivia

- (None.)

Slide 2

Minute Essay From Last Lecture

- (Review answer from online notes.)

Implementing Processes

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

Slide 3

Implementing Processes, Continued

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

Slide 4

Processes Versus Threads

Slide 5

- 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 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

Slide 6

- 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.
- 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 7

Implementing Threads

- Two basic approaches — “in user space” and “in kernel space” (next two slides).
- Various hybrid schemes also possible.

Slide 8

Implementing Threads “In User Space”

- Basic idea — operating system thinks it's managing single-threaded processes, all the work of managing multiple threads happens via library calls within each process.
- Advantages / disadvantages?

Slide 9

Implementing Threads “In User Space”, Continued

- Advantages: fewer system calls, hence probably more efficient.
- Disadvantages:
 - If a thread blocks, it may do so in a way that blocks the whole process.
 - Preemptive multitasking is difficult/impossible.
 - Using multiple CPUs is difficult/impossible.

Slide 10

Implementing Threads “In Kernel Space”

- Basic idea — operating system is involved in managing threads, the work of managing multiple threads happens via system calls (rather than user-level library calls).
- Advantages / disadvantages?

Slide 11

Implementing Threads “In Kernel Space”, Continued

- Advantages: avoids the difficulties of implementing in user space.
- Disadvantages: probably less efficient.

Slide 12

Example Implementations

Slide 13

- Unix systems vary as to which they use (see chapter 10 for more info). Until fairly recently Linux did kernel-space threading, but allegedly with some tweaks to make it more efficient. There have been some changes in the latest versions . . .
- Windows NT/2000 apparently is such that *all* processes have at least one thread, and the basic scheme is either kernel-space or a hybrid (see chapter 11 for more info).

Minute Essay

Slide 14

- If you were doing an object-oriented design for an operating system, you might have a `Process` class and a `Thread` class. How might you think of relating them? (class/subclass? something else?)