

Slide 1

Administrivia

- (None.)

Slide 2

Minute Essay From Last Lecture

- Textbook: Many people using electronic copies, but some paper.
- Most primitive system: MS-DOS, Windows 98, old Apple (Lisa, Mac), custom systems (facial-recognition device, programmable calculator), system for Computer Design, Ubuntu terminal.

Operating System Functionality — Overview

Slide 3

- Provide a “virtual machine”:
 - Filesystem abstraction — files, directories, ownership, access rights, etc.
 - Process abstraction — “process” is a name for one of a collection of “things happening at the same time” (multiple users, multiple applications, background activities such as print spooling, etc.).
- Manage resources (probably on behalf of multiple users/applications):
 - Memory.
 - CPU cycles (one or more CPUs).
 - I/O devices.

Overview of Hardware

Slide 4

- Simplified view of hardware (as it appears to programmers) — processor(s), memory, I/O devices, bus.
- Figure on p. 19 shows simplified view of overall organization — components connected to a single bus. (Actual processors may have more than one bus.)

Processors

Slide 5

- “Instruction set” of primitive operations — load/store, arithmetic/logical operations, control flow.
- Basic CPU cycle — fetch instruction, decode, execute. (Again, this is simplified — pipelined or “superscalar” architectures overlap these steps.)
- Registers — “local memory” for processor; general-purpose registers for arithmetic and other operations, special registers (program counter, stack pointer, program status word (PSW)).
- Not shared among processes in the sense of simultaneous use (but context switches “fake it”).
- Typically also include features useful in writing an operating system . . .

Dual-Mode Operation, Privileged Instructions

Slide 6

- Useful to have mechanism to keep application programs from doing things that should be reserved for o/s.
- Usual approach — in hardware, define two modes for processor (supervisor/kernel and user), privileged instructions.
 - Privileged instructions — things only o/s should do, e.g., enable/disable interrupts.
 - Bit in PSW indicates kernel mode (o/s only, privileged instructions okay) or user mode (application programs, privileged instructions not allowed).
 - When to switch modes? when o/s starts application program, when application program requests o/s services, on error.
 - How to switch? kernel to user seems straightforward, but how about the other way? Usually handled via TRAP or similar instruction, which generates an interrupt (more about interrupts later).

Multithreaded and Multicore Chips

Slide 7

- For many years (at least 20, to my knowledge) advocates of parallel programming have been saying that eventually hardware designers would run out of ways to make single processors faster — and finally it seems to be happening.
- Basic idea — number of transistors one can put on a chip is still increasing, but how to use them to make single processors faster isn't clear. So, instead, hardware designers have chosen to provide (more) hardware support for parallelism. Various approaches, including "hyperthreading" (fast switching among threads), "multicore" (multiple independent CPUs, possibly sharing cache), "GPGPU" (use of graphic card's many processors for computation).

Memory Hierarchy

Slide 8

- In a perfect world — fast, big, cheap, as permanent as desired.
- In this world — hierarchy of types, from fast but expensive to slow but cheap: registers, cache, RAM, magnetic disk, magnetic tape. (See picture, p. 23.)
- Note also — some types volatile, some non-volatile.

Registers and Caches

- Registers — part of processor, fastest to access but most expensive to build. Managed explicitly in software.
- Caches (possibly multiple levels) — less fast, less expensive, bigger. Mostly managed by hardware.
- Aside: Caching is a widely used strategy in computing! virtual memory, disk blocks in memory, etc., etc.

Slide 9

Main Memory (RAM)

- Still less fast, less expensive, bigger.
- Shared among processes — which presents some interesting challenges . . .

Slide 10

Memory Protection

Slide 11

- Very useful to have a way to give each process (including o/s) its own variables that other processes can't alter.
- Usual approach — provide a hardware mechanism such that attempting to access memory out of ranges generates exception/interrupt. Several ways; some simple ones:
 - Limit each process to a range of memory locations; hold starting and ending addresses in special registers.
 - Partition memory into blocks, give each block a numeric key, give each process a key, and only allow processes to access blocks if keys match.

I/O Devices

Slide 12

- What they provide (from the user's perspective):
 - Non-volatile storage (disks, tapes).
 - Connections to outside world (keyboards, microphones, screens, etc., etc.).
- Distance between hardware and "virtual machine" is large here, so usually think in terms of:
 - Layers of s/w abstraction (as with other parts of o/s).
 - Layers of h/w abstraction too: most devices attached via controller, which provides a h/w layer of abstraction (e.g., "IDE controller").

I/O Basics

Slide 13

- CPU communicates with device controller by reading/writing device registers; device controller communicates with device.
- Memory-mapped I/O versus I/O instructions.
- Polling versus interrupts.
- Functionality for a particular device packaged as “device driver”.
- I/O in application programs — make system call to invoke o/s services (more about system calls later).

Minute Essay

Slide 14

- I once had a learning experience about “how DOS is different from a real o/s”.
Summary version: A program using pointers (possibly uninitialized) caused the whole machine to lock up, so thoroughly that the only recovery was to power-cycle.
What do you think went wrong?

Minute Essay Answer

- The program changed memory at the addresses pointed to by the uninitialized pointers — and this memory was being used by the o/s, possibly to store something related to interrupt handling. A “real” o/s wouldn’t allow this!

(Then again, the version of MS-DOS in question was supposedly written to run on hardware that didn’t provide memory protection, so maybe it’s not DOS’s fault.)

Slide 15