

Slide 1

Administrivia

- “Useful links” page has links to images of old equipment, and also a video showing a keypunch machine in action.

Slide 2

Minute Essay From Last Lecture

- Oldest / most primitive system?
 - Windows 9x, pre-Windows PC
 - Atari, Super Nintendo
 - IBM mainframe (“Master the Mainframe” competition)
 - first-generation Windows phone
 - programmable calculator
 - Dewey Decimal cards (?)

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 1.6 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-purpose 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 30, 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 Figure 1-9.)
- 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).

Overview of Hardware — Recap

Slide 14

- Idea is to get a sense of what o/s designers/developers have to work with.
- Notice also what features seem intended to make it possible to write an o/s that can defend itself!
- (I won't talk in class about the sections on buses and booting, but do read them.)

Operating System Functionality — Review

Slide 15

- “Operating system as virtual machine” must provide key abstractions (processes, filesystems).
- “Operating system as resource manager” must manage resources (memory, I/O devices, etc.).
- Operating system functionality typically packaged as “system calls” (more next time).
- Details obviously vary among systems, but some ideas are common to most/many (more next time).

Processes — Abstraction

Slide 16

- Basic idea — a program (application or background activity) together with its current state (registers and memory contents).
- In order to have more than one at a time, need some way to share the physical machine among them.
- May be useful to think in terms of each process having its own simulated processor and memory (“address space”), with operating system providing infrastructure to map that onto the hardware. How to do that? (Next slide.)
- Other relevant concepts include process ownership, hierarchical relationships among processes, interprocess communication.

Processes — Implementation

Slide 17

- Managing the “simulated processor” aspect requires some way to timeshare physical processor(s). Typically do that by defining a per-process data structure that can save information about process. Collection of these is a “process table”, and each one is a “process table entry”.
- Managing the “address space” aspect requires some way to partition physical memory among processes. To get a system that can defend itself (and keep applications from stepping on each other), memory protection is needed — probably via hardware assist. Some notion of address translation may also be useful, as may a mechanism for using RAM as a cache for the most active parts of address space, with other parts kept on disk.

Filesystems

Slide 18

- Most common systems are hierarchical, with notions of “files” and “folders”/“directories” forming a tree. “Links”/“shortcuts” give the potential for a more general (non-tree) graph.
- Connecting application programs with files — notions of “opening” a file (yielding a data structure programs can use, usually by way of library functions).
- Many, many associated concepts — ownership, permissions, access methods (simple sequence of bytes, or something more complex?), whether/how to include direct access to I/O devices in the scheme.
- Relevant system calls — create file, create directory, remove file, open, close, etc., etc.
- See text for some UNIXisms — single hierarchy, regular versus special files, pipes, etc.

I/O

- As noted previously — hardware is diverse, and communicating with it may involve a lot of messy details.
- So — typically there is an “I/O subsystem”, often involving multiple layers of abstraction. More later!

Slide 19

Hardware, Software, and History

- Textbook has a section called “Ontogeny Recapitulates Phylogeny”. Many interesting general observations:
- What seems like a good idea in software is strongly influenced by what the hardware can do. (I think it goes the other way too, but that’s speculation.)
- As in other areas of human endeavor, evolution of operating systems is in some ways cyclic: What seems brilliant now may be “ready for the scrap heap” in a few years — and then resurface as brilliant later. (This is why it’s not useless to read about approaches not currently in use?)

Slide 20

The Operating System “Zoo”

- Section of this name in textbook talks briefly about different kinds of operating systems.
- Key point here? a lot of variation in situations (combination of hardware and “use case”) where an o/s is needed, worth thinking about what implications that might have for o/s.

Slide 21

Operating System Structures

- General-purpose operating systems are big — tens of millions of lines of code (probably mostly in something C-like). How to organize all of it? several choices, discussed in textbook:
 - Monolithic systems.
 - Layered systems.
 - Microkernels.
 - Client-server model.
 - Virtual machines.
 - Exokernels.
- A possibly-relevant maxim, origin unknown (to me): “Any programming problem can be solved by adding a layer of abstraction. Any performance problem can be solved by removing a layer of abstraction.” Not always true, but true enough?

Slide 22

Minute Essay

- 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.

Slide 23

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 24