Computer Systems Names and Layers

Mendel Rosenblum
cs110

Today's Agenda

- Complete the discussion of naming
- Examine names and layering in a computer system

Name conflicts

- Names are used to compose modules
  - Each module can be built independently with its own naming context
- Naming conflicts can occur when combining modules together
- Example: Linking large modules together
  - Problem: Linker croaks or wrong routine called
  - Solutions: Renaming? Wrapping?
- In large systems care is needed in naming
  - Careful about what names you export from a module
    - Naming conventions for extern symbols
  - Might not be able to change names once exported
    - Example: Linux device drivers
  - Bad modules – Pollute name space, expose internals, confuse

Name overloading

- Names frequently encode information about the object
  - **Metadata** – data about the object
    - Examples: Type, location, use, time, etc.
  - Useful since you can extract information from name without a lookup or storage
  - **Pure name** – No overloading
    - Only meaningful operations are resolver API
    - Might be hard to find things
- Example of overloaded names:
  CS110-L04-Naming-4-9-09.ppt
  0x6912f439bbaa9f454f309ff42e2b41b44b1e6ef5 - A checksum of a file
  Mr Cook, Mr Farmer, Mr Miller, Mr Smith, …
  Restaurants: 231 Ellsworth, Gary Danko
- Problem: What happens when the object changes?
  - Either name has to change or name is misleading - **Fragility**

Addresses

- Addresses are used as both names and locators, Examples:
  - IP address: 171.64.64.64, 1950s phone number, I/O device address
  - Highly useful but fragile
- Work-around when object moves:
  - Change all references – can be painful
  - Make it work for both new and old
  - Have client search if resolve fail
- Indirection is frequently the solution
  - Update indirection map to handle moves
  - Examples: host names, cell phone numbers, etc.

User-friendly names

- Tension between good machine and good user names
  - Uniqueness vs memorable/pleasing
    - Example of email addresses:
      - » mendel@cs.stanford.edu
      - » Mendel Rosenblum
    - What about Cindy Mendel or Joe Smith?
    - Account number and account names
- Gets complex. File system name example:
  - Case-coercing – Test.txt = test.txt = TEST.TXT
  - Case-sensitive - Test.txt != Test.Txt != Test.txt
  - Case-preserving

Lifetimes and Distribution make naming difficult

- Suppose you want unique identifiers
  - Simple: sequenceID++;
    - What about machine crashes?
    - What if you need million a second from all over the world?
  - Example: timestamps, global partitioning
- Other lifetime problems
  - What happens when objects die and name lives?
    - **Dangling references**
  - What happens when names die and objects live?
    - Lost object, **storage leak**
  - Solution: Reference counting
    Example: i_nlink in the V6 Unix Inode

Computer Systems Layers

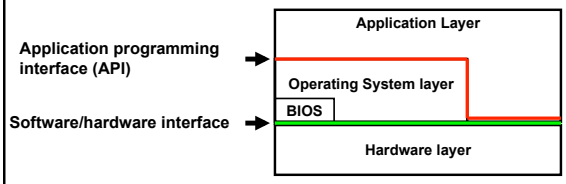

Layers and interfaces
- Tradeoff between what functionality goes into each layer
- Layers sometime **bypass** lower layers

Layers export names

Hardware/software interface 1

Physical Address Name Space

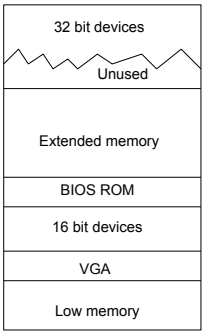

0xFFFFFFFF (4GB)

Memory size

0x00100000 (1MB)

0x000F0000 (960KB)

0x000C0000 (768KB)

0x000A0000 (640KB)

0x00000000

Hardware/software interface 2

IN/OUT Port Name Space

```
0000-000f  Slave DMA controller (8237 chip)
0010-001F System
0020-0021  First Interrupt controller (8259 chip)
0030-0031  Second interrupt controller
0040-0043  Programable Interval Timer 1 (8254 chip)
0048-004B  Programable Interval Timer 2
0070-0071  NMI Enable / Real Time Clock
0080-008B  DMA Page registers
00A0-00A1  Slave interrupt controller
00C0-00DE                Master DMA controller
0168-016F  IDE Interface - Quaternary channel
0200-0207  Games Port (joystick port)
0278-027A  Parallel  Port *
02E8-02EF  Serial Port - COM 4
02F8-02FF  Serial Port - COM 2
0378-037A  Parallel Port *
03E0-03E7  PCIC PCMCIA Port Controller
03E8-03EF  Serial Port - COM 3
03F0-03F6  Floppy Disk Drive Interface
03F8-03FF  Serial Port - COM 1
```

More Hardware Names

- Plug and Play
  - PCI configure space lookup
  - Place devices at physical addresses
- Some devices export names
  - Network Interface Cards:
    - Ethernet: 48bit MAC: Ex 0x0016D3C86305
    - 802.11g 48bit MAC: 0x00197E92127C
  - Also, disks, USB devices, Bluetooth radio,
- Processor Serial Number

BIOS Layer

- Initialize the CPU and chipset
- Interface:
  - INT 10h Video Services
  - INT 13h Low Level Disk Services
  - INT 14h Serial port
  - INT 15h Miscellaneous Services
  - INT 16h communicate with the keyboard
  - INT 17h Print Services
  - INT 19h load the operating system
  - INT 1Ah Real Time Clock Services

Application Programming Interface

- Each process gets its own 32bit address space starting at 0 going to 4GB

Version 6 Unix System calls:

- Process Control
  1 = exit, 2 = fork, 7 = wait, 37 = kill, 34 = nice, 20 = getpid, 23 = setuid, 24 = getuid, 41 = dup, 42 = pipe, 17 = break
- Utility
  25 = stime, 26 = ptrace, 43 = times, 44 = prof, 31 = stty, 32 = gtty, 35 = sleep, 36 = sync,
- File System
  3 = read, 4 = write, 5 = open, 6 = close, 8 = creat, 9 = link, 10 = unlink, 11 = exec, 12 = chdir, 13 = time, 14 = mknod, 15 = chmod, 16 = chown, 18 = stat, 19 = seek,28 = fseek, 21 = mount, 22 = umount,

OS Abstractions – File Systems

- Name mappings
  - Disk: Cylinder Head Sector addressing
  - Unix block driver: Logical block addressing
  - System call: file system
- Files – Hide ugliness of disks
  - File - Linear array of bytes
  - Hierarchical naming system (directories)

File System API

fd = open(pathname, 0)
- lseek(fd, offset, fromWhere)
- read(fd, buffer, len) or write(fd, buffer, len)
- close(fd)

- Adds the "fd" name space
- Supports sequential and non-sequential access
- Why not read(pathname, start, buffer, len)?
  - Resolving pathname can be costly
    - Recursive lookup
    - Size communicated between layers
  - Gives clues when application done with file
    - Atomicity & consistency

Principle of least astonishment

People are part of the system. The design should match the user's experience, expectations, and mental models.
- Be consistent, predictable
- Minimize side effects
- Use names that describe
- Do the obvious thing

- Examples:
  - Deleting a file doesn't normally erase it from disk
  - You have to push submit to have your vote count

Naming Example: URLs

http://cs110.stanford.edu/handouts/L01-Intro.pdf

- Multiple different name context lookups
  - cs110.stanford.edu - DNS name server
  - handouts/L01-Intro.pdf gets set to the web server to be interpreted in anyway it sees fit
- Overloading:

http://www.facebook.com/home.php#/s.php?init=q&q=cs110&ref=ts&sid=129d6bf025962257d7724201174244 93

Reading Assignment

- Read
  - Section 4.1, 4.2