Computer Science E-1

Exam 1 Review Session

Administrivia

- Exam is next Monday, right here.
  - 5:30pm - 7:30pm. You will have 2 hours to complete the exam.
- The exam will consist of questions, equally representing lectures 1-4.
  - Including (but not limited to): Fill in the blank, short answer, Multiple Choice.
- No calculators, notes, or other aids.

Study Tips

- Read the recaps
  - Don’t just skim them, pause and make sure you understood what you just read
  - Use the practice questions at the end for, well, practice!
- Go over the problem sets
  - Especially the problems you got wrong!

Computer Science E-1

Review of Lecture 1: Hardware

Motherboard and Power Supply

- Motherboard - circuit board that holds crucial components
- Power Supply - manages voltage sent to individual components

BIOS, CMOS, POSTs

- BIOS (Basic Input Output System) - instructions for computer to start successfully
- CMOS (Complementary Metal-Oxide Semiconductor) - stores information for BIOS
- POSTs (Power On Self Tests) - test that computer's hardware is functioning properly

Graphics Cards and Ports

- PCI Slot - connects graphics cards, sound cards, video cards
- PCI Express - newer, faster PCI standard
- AGP (Accelerated Graphics Port) - higher performance port for graphics cards

RAM

- RAM (Random Access Memory) - computer's short-term memory
  - Typical values - 2GB to 4GB

Ethernet and USB

- Ethernet port - connects computer to router or modem
- USB port - connects computer to cameras, smartphones, printers
  - USB A to B, Micro USB to A

VGA, DVI, HDMI

- Connect computer to external display (monitor)
- VGA - analog, no audio
- DVI - digital, no audio
- HDMI - digital, audio, commonly used for TVs

OS, Kernel

- OS (Operating System) - Windows, OS X, Linux
- Kernel - part of the OS, bridge between hardware and other software

Hard Drive and Connectors

- Hard Drive - computer's long-term storage
  - Typical values - several hundred GBs
- SATA, IDE (PATA) - connects hard drive to motherboard
- HDD vs SSD

CPU and Heat sink

- CPU - Processor, brains of the computer
  - Speed is measured in GHz
- Heat sink - keeps CPU cool

Northbridge and Southbridge

- Northbridge - the half of the motherboard that frequently interfaces with CPU
  - CPU, RAM, graphics card
- Southbridge - interfaces with other components
  - USB, PCI, SATA hard drive

Keyboard

- Key matrix - grid of circuits under keys
  - A circuit is completed each time a key is pressed
- ROM (Read-Only Memory) - used by keyboard to map circuit locations to characters

Mouse

- Ball mouse - tracks motion of ball to determine mouse movements
- Optical mouse - uses LED to track motion

Bits and Binary

- Bits (Binary Digits) - 0s and 1s
- Binary - representation of numbers with bits

Binary to Decimal

10101

Binary to Decimal

1 0 1 0 1

(1 * 16) + (0 * 8) + (1 * 4) + (0 * 2) + (1 * 1)

16 + 4 + 1

21

Decimal to Binary

12

_ _ _ _ _

Decimal to Binary

12

0 1 1 0 0

(0 * 16) + (1 * 8) + (1 * 4) + (0 * 2) + (0 * 1)

ASCII and UTF-8

- ASCII - defines characters as numbers
  - ‘A’ = 65
- UTF-8 - a form of unicode, defines over a million characters

Nibbles and Bytes

- Nibbles - 4 bits
- Bytes - 8 bits
- Kilobytes - ~1000 bytes
- Megabytes - ~1000 kilobytes
- Gigabytes - ~1000 megabytes

Computer Science E-1

Review of Lecture 2: Hardware

Hardware, Continued

- CPU
- RAM
- Hard Drive

CPU

- Processes
- Pipeline
- Parallelism
- Performance

CPU: Processes

- “a single instance of a program being executed”
- each can use time on the CPU, and memory

CPU: Instruction Set

- Different CPUs understand different instructions
- The instructions a CPU can follow is called its “instruction set”
- Three primary types of instructions:
  - **Data -** store or read data from memory
  - **Arithmetic -** perform arithmetic on data
  - **Control Flow -** determine which instruction(s) to execute next

CPU: Pipeline

- Each instruction is broken up into a series of steps, called the pipeline.
- For example, a 4 stage pipeline could look like this:
  - **Fetch -** read instruction into memory
  - **Decode -** determine what needs to be done
  - **Execute -** do it
  - **Writeback -** store result somewhere

CPU: Parallelism

- Superscalar architectures
- Multicore processors

CPU: Superscalar

- Each stage of the CPU pipeline can happen at the same time.
- This allows more instructions to occur in the same number of cycles.

CPU: Multicore

- If we have 2 processors, 2 instructions can be run at the same time.
- What problems have we introduced?
  - Managing two processors
  - Making sure instructions don’t interfere with each other
  - Balancing workload

CPU: Performance

- Clockspeed, pipeline length
- Moore's Law - processing power doubles every 18 months
  - Will this keep up?
- "Megahertz Myth"

Memory

- Registers
- RAM
- Endian-ness
- Cache

Memory: Registers

- Small pieces of memory on the CPU
  - VERY tiny (on the order of bytes)
  - used for computation, etc
- Fast, expensive

Memory: RAM

- Short term memory with a larger capacity
- Slower, less expensive than CPU registers
- Used to store things like songs, frames of a video, webpages, etc.
- Broken up into 1-byte “blocks”

Memory: Endian-ness

- Because most data is larger than 1 byte, we need some way to combine blocks.
- If we read bytes left-to-right, this is called **big-endian.**
- If we read bytes right-to-left, this is called **little-endian.**

Endian Example

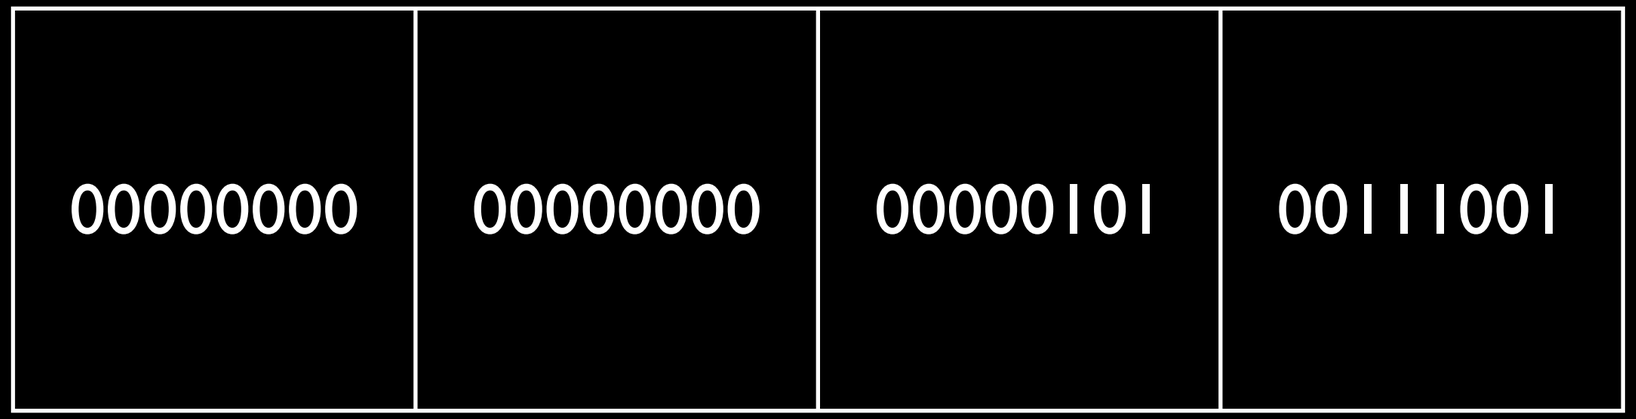

In a big-endian system, these four bytes represent
00000000000000000000101001111001 = 1,337

In a little-endian system, these four bytes represent
00111001000001010000000000000000 = 956,628,992

Memory: Cache

- Small, fast memory located near the processor.
- Typically stores a copy of something on RAM, for quicker access.
- Two types:
  - L1 cache - closer, smaller, and faster
  - L2 cache - slightly farther away, bigger, slower

Hard Drives

- HDDs
- Filesystems
- SSDs
- Virtual Memory

HDD

- consists of magnetic **platters**
- information is encoded on the platters via **read-write heads**
- platter divided into **tracks** and **sectors**
- locations of files are stored in a **file allocation table**, on the drive

Filesystems

- The details of how files are stored and accessed depends on the type of **filesystem** that the operating system employs in organizing files.
- Windows typically uses a filesystem called **NTFS.**
- Thumb drives typically use **FAT.**
- Mac OS X typically uses **HFS+.**

Filesystems, cont.

- What makes filesystems different?
  - The way things are organized on disk (unseen by you)
  - The maximum file size.
  - The maximum disk size.

SSDs

- Solid State Drives (SSDs) have no moving parts, and store their contents in flash memory, much like RAM.
- SSDs are more expensive and store less than HDDs, but are faster and use less power.

Virtual Memory

- I'll be honest, this is pretty confusing.
- Essentially, the operating system allows processes to think that they have access to more memory than is available.
- This is implemented by creating **swap files** on the hard drive, that act as RAM.
- Notice that this is slower than having more RAM.

Shopping!

- display
- weight
- resolution
- CPU cores
- CPU cache size
- RAM size
- storage capacity
- HDD / SSD
- peripherals
- keyboard

Summary

(taken from lecture, so you know it's good!)

CPU

- instruction set
- pipeline
- parallelism
- superscalar
- multi-core
- clock speed
- Moore's Law
- Megahertz Myth

Memory

- byte, kilobyte, megabyte
- registers
- RAM
- addressing
- big-endian, little-endian
- caching
- L1 cache, L2 cache, L3 cache
- hard disk drive
- platters, tracks, sectors
- read-write heads
- file allocation table
- seek time, data rate
- solid state drive
- flash memory

Computer Science E-1

Review of Lecture 3: Internet

Clients and Servers

- Server - computer that powers a website
- Client - user that makes request to server

ISP, LAN, WAN

- ISP (Internet Service Provider) - connects you to the internet
- LAN (Local Area Network) - connects computers in a limited area, such as a home or office
- WAN (Wide Area Network) - connects computers over a larger area

IP, IPv4, IPv6

- IP Address - unique series of numbers assigned to each device on a network
- IPv4 - 32 bit IP address
- IPv6 - 128 bit IP address

Routers and Routing Table

- Router - routes network requests to correct destinations
- Routing Table - maps ranges of addresses to routers

NAT, Private IPs, Source Port

- NAT (Network Address Translation) - modifies requests from devices with private IPs to give them the router's public IP
- Private IP Address - addresses that aren't publicly accessible
- Source Port - unique number used by the router to identify devices connected to the network

DHCP

- DHCP (Dynamic Host Configuration Protocol) - process through which clients obtain an IP address on a network

Firewalls and VPNs

- Firewall - software that protects a network from malicious connections
- VPN (Virtual Private Network) - allows you to access a LAN you're not connected to

DNS and Domains

- DNS (Domain Name System) - a database that maps domain names to specific IP addresses
- Subdomain - addresses like foo.domain.tld
- TLD (Top Level Domain) - suffix like com, net, or org
- ccTLD - country code TLD like tm, ly, sy

DNS Servers

- Cache DNS server - first server contacted by browser to determine IP address of a domain name
- Root DNS server - contacted after trying several DNS servers, forwards requests to TLD DNS server
- TLD DNS server - handles all request for a certain TLD (.net)
- Authoritative name server - queried by TLD DNS, maintains lists of all addresses in a DNS zone

DNS Records

- NS: name servers
- MX: email
- A: IPv4 address
- AAAA: IPv6 address
- CNAME: domain alias

Registrars and ICANN

- Registrar - Organization that handles reservation of domain names
- ICANN - Non-profit responsible for managing registered domain names

URL

http://foo.example.com:1234/cs/e1/is.html?fun=yes&boring=no#awesome

- Scheme - “http”
- Domain - “foo.example.com”
- Port - “1234”
- Path - “/cs/e1/is.html”
- Query string - “fun=yes&boring=no”
- Fragment - “#awesome”

Key-Value Pairs

query=cse l &page=3

- keys - “query” and “page”
- values - “cse l ” and “3”

URI

- URI (Uniform Resource Identifier) - a unique identifier, more general than URL

APIs

- API (Application Programming Interface) - a way of exposing information in a standard structure
  - Facebook API to allow programmers to get information from profiles

Computer Science E-1

Review of Lecture 4: Internet

Internet, Continued

- Protocols
- HTML
- Email
- TCP/IP

Protocols

- DHCP
- HTTP
- TCP/IP

DHCP

- send DHCPBROADCAST - to locate network's DHCP
- receive DHCPOFFER - contains IP for client
- send DHCPREQUEST - asking to claim the IP
- receive DHCPACK - confirmation

HTTP

- Request
  - method/verb (e.g. GET, POST)
- Response
  - Headers
  - Body

HTTP Example

GET /home.php HTTP/1.1

Host: facebook.com

```
HTTP/1.1 200 OK
Date: Wed, 30 Jan 2013
21:43:11 GMT
Server: Apache/2.2.22 (Fedora)
Content-Length: 2422
Content-Type: text/html;
charset=UTF-8
Connection: close
```

Status Codes

- Comes back with the HTTP Response
- Helps the browser know what to do (render, redirect, show an error, etc.)
- Some good ones to know:
  - 200 (OK), 404 (Not Found), 500 (Internal Server Error)

HTML

- tags - <h1></h1>, for example
- contained in the body of the HTTP response

SMTP

- Simple Mail Transfer Protocol
- SMTP Server managed by mail provider (like Google)
- Queue of messages to send
- Headers/Body - look at the detailed demo in the recap

POP3

- Way of receiving messages from a server
- Commands that your email client (like Gmail) would send to the server (in particular retr and dele)
- Doesn’t store state on the server

IMAP

- Two way communication - can store state on the server
  - Like read vs. unread, etc.

TCP/IP

- How an HTTP request is actually sent/received

TCP

- TCP ensures reliable data transfer
  - all the segments will get there, and they will get there in order
  - How? Sequence numbers and ACKs (confirmation of receipt)