

15-213
"The Class That Gives CMU Its Zip!"

Introduction to Computer Systems

Seth Goldstein & Andreas Nowatzky
January 13, 2004

Topics:

- n Theme
- n Five great realities of computer systems
- n How this fits within CS curriculum
- n Staff, text, and policies
- n Lecture topics and assignments
- n Lab rationale

class01.ppt

Course Theme

- n Abstraction is good, but don't forget reality!

Courses to date emphasize abstraction

- n Abstract data types
- n Asymptotic analysis

These abstractions have limits

- n Especially in the presence of bugs
- n Need to understand underlying implementations

Useful outcomes

- n Become more effective programmers
 - l Able to find and eliminate bugs efficiently
 - l Able to tune program performance
- n Prepare for later "systems" classes in CS & ECE
 - l Compilers, Operating Systems, Networks, Computer Architecture, Embedded Systems

- 2 -

15-213, S'04

Great Reality #1

Int's are not Integers, Float's are not Reals

Examples

- n Is $x^2 = 0$?
 - l Float's: Yes!
 - l Int's:
 - » $40000 * 40000 \rightarrow 1600000000$
 - » $50000 * 50000 \rightarrow -1794967296$
- n Is $(x + y) + z = x + (y + z)$?
 - l Unsigned & Signed Int's: Yes!
 - l Float's:
 - » $(1e20 + -1e20) + 3.14 \rightarrow 3.14$
 - » $1e20 + (-1e20 + 3.14) \rightarrow 0$

- 3 -

15-213, S'04

Computer Arithmetic

Does not generate random values

- n Arithmetic operations have important mathematical properties

Cannot assume "usual" properties

- n Due to finiteness of representations
- n Integer operations satisfy "ring" properties
 - l Commutativity, associativity, distributivity
- n Floating point operations satisfy "ordering" properties
 - l Monotonicity, values of signs

Observation

- n Need to understand which abstractions apply in which contexts
- n Important issues for compiler writers and serious application programmers

- 4 -

15-213, S'04

Great Reality #2

You've got to know assembly

Chances are, you'll never write program in assembly

- n Compilers are much better & more patient than you are

Understanding assembly key to machine-level execution model

- n Behavior of programs in presence of bugs
 - l High-level language model breaks down
- n Tuning program performance
 - l Understanding sources of program inefficiency
- n Implementing system software
 - l Compiler has machine code as target
 - l Operating systems must manage process state

- 5 -

15-213, S'04

Assembly Code Example

Time Stamp Counter

- n Special 64-bit register in Intel-compatible machines
- n Incremented every clock cycle
- n Read with rdtsc instruction

Application

- n Measure time required by procedure
 - l In units of clock cycles

```
double t;
start_counter();
P();
t = get_counter();
printf("P required %f clock cycles\n", t);
```

- 6 -

15-213, S'04

Code to Read Counter

- n Write small amount of assembly code using GCC's asm facility
- n Inserts assembly code into machine code generated by compiler

```
static unsigned cyc_hi = 0;
static unsigned cyc_lo = 0;

/* Set *hi and *lo to the high and low order bits
of the cycle counter.
*/
void access_counter(unsigned *hi, unsigned *lo)
{
    asm("rdtsc; movl %%edx,%0; movl %%eax,%1"
        : "=r" (*hi), "=r" (*lo)
        : "%edx", "%eax");
}
```

- 7 -

15-213, S'04

Code to Read Counter

```
/* Record the current value of the cycle counter. */
void start_counter()
{
    access_counter(&cyc_hi, &cyc_lo);
}

/* Number of cycles since the last call to start_counter. */
double get_counter()
{
    unsigned ncyc_hi, ncyc_lo;
    unsigned hi, lo, borrow;
    /* Get cycle counter */
    access_counter(&ncyc_hi, &ncyc_lo);
    /* Do double precision subtraction */
    lo = ncyc_lo - cyc_lo;
    borrow = lo > ncyc_lo;
    hi = ncyc_hi - cyc_hi - borrow;
    return (double) hi * (1 << 30) * 4 + lo;
}
```

- 8 -

15-213, S'04

Measuring Time

Trickier than it Might Look

- n Many sources of variation

Example

- n Sum integers from 1 to n

n	Cycles	Cycles/n
100	961	9.61
1,000	8,407	8.41
1,000	8,426	8.43
10,000	82,861	8.29
10,000	82,876	8.29
1,000,000	8,419,907	8.42
1,000,000	8,425,181	8.43
1,000,000,000	8,371,2305,591	8.37

Timing System Performance

```
main(int argc, char** argv)
{
    ...
    for (i=0; i<t; i++) {
        start_counter();
        count(n);
        times[i] = get_counter();
    }
    ...
}
```

```
int count(int n)
{
    int i;
    int sum = 0;

    for (i=0; i<n; i++) {
        sum += i;
    }
    return sum;
}
```

```
int count(int n)
{
    int i;
    int sum = 0;

    for (i=0; i<n; i++) {
        sum += i;
    }
    return sum;
}
```

```
main(int argc, char** argv)
{
    ...
    for (i=0; i<t; i++) {
        start_counter();
        count(n);
        times[i] = get_counter();
    }
    ...
}
```

Timing System Performance

```
main(int argc, char** argv)
{
    ...
}

int count(int n)
{
    ...
}
```

```
int count(int n)
{
    ...
}

main(int argc, char** argv)
{
    ...
}
```

Experiment	n	cycles/n
1	10	1649.2
2	10	17.2
3	1000	24.3
4	1000	6.1

Experiment	n	cycles/n
1	10	1657.6
2	10	26
1a	10	20
2a	10	16.4
3a	1000	1.7
4a	1000	1.6

It's the system, stupid!

Great Reality #3

Memory Matters Random Access Memory is an un-physical abstraction

Memory is not unbounded

- n It must be allocated and managed
- n Many applications are memory dominated

Memory performance is not uniform

- n Cache and virtual memory effects can greatly affect program performance
- n Adapting program to characteristics of memory system can lead to major speed improvements

Memory referencing bugs especially pernicious

- n Effects are distant in both time and space

Memory Performance Example

Implementations of Matrix Multiplication

Multiple ways to nest loops


```

/* ijk */
for (i=0; i<n; i++) {
  for (j=0; j<n; j++) {
    sum = 0.0;
    for (k=0; k<n; k++)
      sum += a[i][k] * b[k][j];
    c[i][j] = sum;
  }
}
    
```


```

/* ikj */
for (i=0; i<n; i++) {
  for (k=0; k<n; k++) {
    sum = 0.0;
    for (j=0; j<n; j++)
      sum += a[i][k] * b[k][j];
    c[i][j] = sum
  }
}
    
```


Matmult Performance (Alpha 21164)




Memory System



Blocked matmult perf (Alpha 21164)



Real Memory Performance



- 17 -

15-213, S'04

Memory Referencing Bug Example

```
main ()
{
    long int a[2];
    double d = 3.14;
    a[2] = 1073741824; /* Out of bounds reference */
    printf("d = %.15g\n", d);
    exit(0);
}
```

	Alpha	MIPS	Linux
-g	5.30498947741318e-315	3.1399998664856	3.14
-O	3.14	3.14	3.14

(Linux version gives correct result, but implementing as separate function gives segmentation fault.)

- 18 -

15-213, S'04

Memory Referencing Errors

C and C++ do not provide any memory protection

- Out of bounds array references
- Invalid pointer values
- Abuses of malloc/free

Can lead to nasty bugs

- Whether or not bug has any effect depends on system and compiler
- Action at a distance
 - Corrupted object logically unrelated to one being accessed
 - Effect of bug may be first observed long after it is generated

How can I deal with this?

- Program in Java, Lisp, or ML
- Understand what possible interactions may occur
- Use or develop tools to detect referencing errors

- 19 -

15-213, S'04

Great Reality #4

There's more to performance than asymptotic complexity

Constant factors matter too!

- Easily see 10:1 performance range depending on how code written
- Must optimize at multiple levels: algorithm, data representations, procedures, and loops

Must understand system to optimize performance

- How programs compiled and executed
- How to measure program performance and identify bottlenecks
- How to improve performance without destroying code modularity and generality

- 20 -

15-213, S'04

Great Reality #5

Computers do more than execute programs

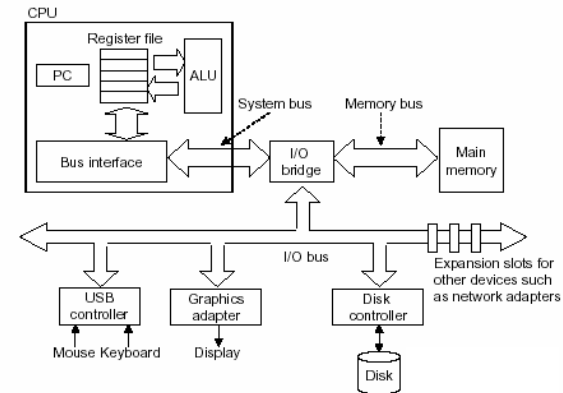
They need to get data in and out

- n I/O system critical to program reliability and performance

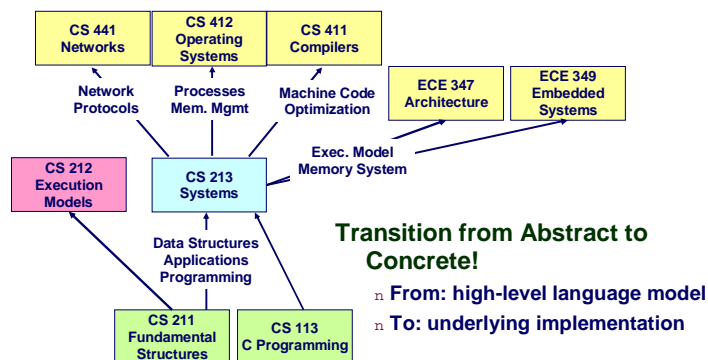
They communicate with each other over networks

- n Many system-level issues arise in presence of network
 - 1 Concurrent operations by autonomous processes
 - 1 Coping with unreliable media
 - 1 Cross platform compatibility
 - 1 Complex performance issues

Hardware Organization (Naïve)



Role within Curriculum



Course Perspective

Most Systems Courses are Builder-Centric

- n Computer Architecture
 - 1 Design pipelined processor in Verilog
- n Operating Systems
 - 1 Implement large portions of operating system
- n Compilers
 - 1 Write compiler for simple language
- n Networking
 - 1 Implement and simulate network protocols

Course Perspective (Cont.)

Our Course is Programmer-Centric

- n Purpose is to show how knowing more about the underlying system, leads one to be a more effective programmer
- n Enable you to
 - l Write programs that are more reliable and efficient
 - l Incorporate features that require hooks into OS
 - » E.g., concurrency, signal handlers
- n Not just a course for dedicated hackers
 - l We bring out the hidden hacker in everyone
- n Cover material in this course that you won't see elsewhere

Teaching staff


- n Instructors
 - l Prof. Seth Goldstein (Wed 1--2pm, WeH 7122)
 - l Prof. Andreas Nowatzky (Tue 3--4pm, NSH 4117)
- n TA's
 - l Ningning Hu (A, Tue 5--6pm, WeH 8205)
 - l Carolyn Au (B, Wed 3--4pm, WeH 3108)
 - l David Charlton (C, Fri 11:30--12:30pm, Weh 3108)
 - l David Fields (D, Wed 12:30am--1:30pm, Weh 3108)
 - l Mike Nollen (E, Thu 3--4pm, Weh 3108)
- n Course Admin
 - l Norene Mears (WeH 7114)

**These are the nominal office hours. Come talk to us anytime!
(Or phone or send email)**

Textbooks


Randal E. Bryant and David R. O'Hallaron,

- n "Computer Systems: A Programmer's Perspective", Prentice Hall 2003.
- n <http://csapp.cs.cmu.edu/>



Samuel P. Harbison III and Guy L. Steele Jr.,

- n "C A Reference Manual 5th Edition", Prentice Hall, 2002
- n <http://careferencemanual.com/>



Course Components

Lectures

- n Higher level concepts

Recitations

- n Applied concepts, important tools and skills for labs, clarification of lectures, exam coverage

Labs

- n The heart of the course
- n 1, 2, or 3 weeks
- n Provide in-depth understanding of an aspect of systems
- n Programming and measurement

Getting Help

Web

- n www.cs.cmu.edu/~213
- n Copies of lectures, assignments, exams, solutions
- n Clarifications to assignments

Newsgroup

- n cmu.cs.class.cs213
- n Clarifications to assignments, general discussion

Personal help

- n Professors: door open means come on in (no appt necessary)
- n TAs: please mail or zephyr first.

Policies: Assignments

Work groups

- n Labs: You must work alone on all labs

Handins

- n Assignments due at 11:59pm on specified due date
- n Typically 11:59pm Wednesday evening
- n Electronic handins only
- n Allowed a **total** of up to 5 late days for the semester

Makeup exams and assignments

- n OK, but must make PRIOR arrangements with either Prof. Goldstein or Nowatzky

Appealing grades

- n Within 7 days of due date or exam date
- n Assignments: Talk to the lead person on the assignment
- n Exams: Talk to either Prof. Goldstein or Nowatzky

Cheating

What is cheating?

- n Sharing code: either by copying, retyping, looking at, or supplying a copy of a file.
- n Using solutions or tools other than those from the course book, lectures, or staff.

What is NOT cheating?

- n Helping others use systems or tools.
- n Helping others with high-level design issues.
- n Helping others debug their code.

Usual penalty for cheating:

- n Removal from course with failing grade.
- n Note in student's permanent record

Policies: Grading

Exams (40%)

- n Two in class exams (10% each)
- n Final (20%)
- n All exams are open book/open notes.

Labs (60%)

- n 7 labs (8-12% each)

Grading Characteristics

- n Lab scores tend to be high
 - l Serious handicap if you don't hand a lab in
- n Tests typically have a wider range of scores

Facilities

Assignments will use Intel Computer Systems Cluster (aka “the fish machines”)

- n 25 Pentium III Xeon servers donated by Intel for CS 213
- n 550 MHz with 256 MB memory.
- n Rack mounted in the 3rd floor Wean machine room.
- n We'll be setting up your accounts this week.

Getting help with the cluster machines:

- n See course Web page for info
- n Please direct questions to your TAs

Programs and Data

Topics

- n Bits operations, arithmetic, assembly language programs, representation of C control and data structures
- n Includes aspects of architecture and compilers
- n Learning the tools

Assignments **L1 Available THUR! (Due 1/25 11:59pm)**

- n L1: Manipulating bits
- n L2: Defusing a binary bomb
- n L3: Hacking a buffer bomb

Performance

Topics

- n High level processor models, code optimization (control and data), measuring time on a computer
- n Includes aspects of architecture, compilers, and OS

Assignments

- n L4: Optimizing Code Performance

The Memory Hierarchy

Topics

- n Memory technology, memory hierarchy, caches, disks, locality
- n Includes aspects of architecture and OS.

Assignments

- n L4: Optimizing Code Performance

Linking and Exceptional Control Flow

Topics

- Object files, static and dynamic linking, libraries, loading
- Hardware exceptions, processes, process control, Unix signals, nonlocal jumps
- Includes aspects of compilers, OS, and architecture

Assignments

- L5: Writing your own shell with job control

Virtual memory

Topics

- Virtual memory, address translation, dynamic storage allocation
- Includes aspects of architecture and OS

Assignments

- L6: Writing your own malloc package

I/O, Networking, and Concurrency

Topics

- High level and low-level I/O, network programming, Internet services, Web servers
- concurrency, concurrent server design, threads, I/O multiplexing with select.
- Includes aspects of networking, OS, and architecture.

Assignments

- L7: Writing your own Web proxy

Lab Rationale

Each lab should have a well-defined goal such as solving a puzzle or winning a contest.

- Defusing a binary bomb.
- Winning a performance contest.

Doing a lab should result in new skills and concepts

- Data Lab: computer arithmetic, digital logic.
- Bomb Labs: assembly language, using a debugger, understanding the stack
- Perf Lab: profiling, measurement, performance debugging.
- Shell Lab: understanding Unix process control and signals
- Malloc Lab: understanding pointers and nasty memory bugs.
- Proxy Lab: network programming, server design

We try to use competition in a fun and healthy way.

- Set a threshold for full credit.
- Post intermediate results (anonymized) on Web page for glory!

Autolab Web Service

Labs are provided by the Autolab system

- n Developed in summer 2003 by Dave O'Hallaron
- n Apache Web server + Perl CGI programs
- n Beta tested in Fall 2003, so of course, bug free now

With Autolab you can use your Web browser to:

- n Review lab notes
- n Download the lab materials
- n Stream autoresults to a class status Web page as you work.
- n Upload (handin) your code for autograding by the Autolab server.
- n View the complete history of your code handins, autoresult submissions, autograding reports, and instructor evaluations.
- n View the class status page

Acknowledgement

15-213 was developed and fine-tuned by
Randal E. Bryant and David O'Hallaron.
They wrote *The Book!*

Have a Great Semester!