

Carnegie Mellon

Introduction to Computer Systems

15-213/18-243, Spring 2009
1st Lecture, Aug. 25th

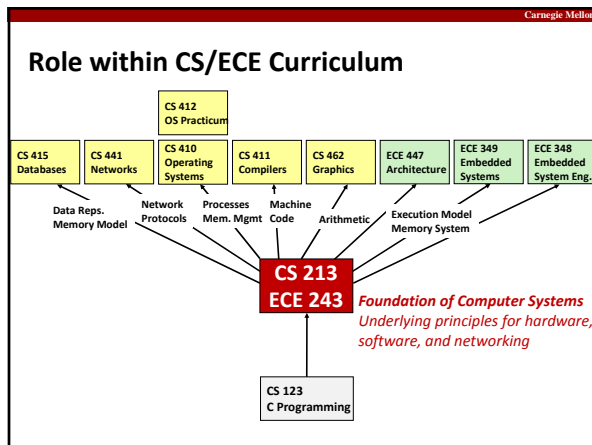
Instructors:
Roger Dannenberg and Greg Ganger

The course that gives CMU its "Zip"!

Carnegie Mellon

Overview

- Course role and theme
- Five realities
- Logistics



Carnegie Mellon

Course Perspective

- Most Systems Courses are Builder-Centric
 - Computer Architecture
 - Design pipelined processor in Verilog
 - Operating Systems
 - Implement large portions of operating system
 - Embedded Systems
 - Implement small-scale embedded systems
 - Networking
 - Implement and simulate network protocols

Carnegie Mellon

Course Perspective (Cont.)

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

Carnegie Mellon

Course Theme:

Abstraction Is Good But Don't Forget Reality

- Most CS courses emphasize abstraction
 - Abstract data types
 - Asymptotic analysis
- These abstractions have limits
 - Especially in the presence of bugs
 - Need to understand details of underlying implementations
- Useful outcomes
 - Become more effective programmers
 - Able to find and eliminate bugs efficiently
 - Able to understand and tune for program performance
 - Prepare for later "systems" classes in CS & ECE
 - Compilers, Operating Systems, Networks, Computer Architecture, Embedded Systems

Carnegie Mellon

Great Reality #1: Int's are not Integers, Float's are not Reals

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

Carnegie Mellon

Computer Arithmetic

- **Does not generate random values**
 - Arithmetic operations have important mathematical properties
- **Cannot assume all "usual" mathematical properties**
 - Due to finiteness of representations
 - Integer operations satisfy "ring" properties
 - Commutativity, associativity, distributivity
 - Floating point operations satisfy "ordering" properties
 - Monotonicity, values of signs
- **Observation**
 - Need to understand which abstractions apply in which contexts
 - Important issues for compiler writers and serious application programmers

Carnegie Mellon

Great Reality #2: You've Got to Know Assembly

- **Chances are, you'll never write program in assembly**
 - Compilers are much better & more patient than you are
- **But: Understanding assembly key to machine-level execution model**
 - Behavior of programs in presence of bugs
 - High-level language model breaks down
 - Tuning program performance
 - Understand optimizations done/not done by the compiler
 - Understanding sources of program inefficiency
 - Implementing system software
 - Compiler has machine code as target
 - Operating systems must manage process state
 - Creating / fighting malware
 - x86 assembly is the language of choice!

Carnegie Mellon

Great Reality #3: Memory Matters Random Access Memory Is an Unphysical Abstraction

- **Memory is not unbounded**
 - It must be allocated and managed
 - Many applications are memory dominated
- **Memory referencing bugs especially pernicious**
 - Effects are distant in both time and space
- **Memory performance is not uniform**
 - Cache and virtual memory effects can greatly affect program performance
 - Adapting program to characteristics of memory system can lead to major speed improvements

Carnegie Mellon

Memory Referencing Bug Example

```
double fun(int i)
{
    volatile double d[1] = {3.14};
    volatile long int a[2];
    a[i] = 1073741824; /* Possibly out of bounds */
    return d[0];
}
```

```
fun(0) -> 3.14
fun(1) -> 3.14
fun(2) -> 3.1399998664856
fun(3) -> 2.00000061035156
fun(4) -> 3.14, then segmentation fault
```

Carnegie Mellon

Memory Referencing Bug Example

```
double fun(int i)
{
    volatile double d[1] = {3.14};
    volatile long int a[2];
    a[i] = 1073741824; /* Possibly out of bounds */
    return d[0];
}
```

```
fun(0) -> 3.14
fun(1) -> 3.14
fun(2) -> 3.1399998664856
fun(3) -> 2.00000061035156
fun(4) -> 3.14, then segmentation fault
```

Explanation:

Saved State	4
d7 ... d4	3
d3 ... d0	2
a[1]	1
a[0]	0

} Location accessed by fun(i)

Carnegie Mellon

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 or ML
 - Understand what possible interactions may occur
 - Use or develop tools to detect referencing errors

Carnegie Mellon

Great Reality #4: There's more to performance than asymptotic complexity

- **Constant factors matter too!**
- **And even exact op count does not predict performance**
 - 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

Carnegie Mellon

Memory System Performance Example

```
void copyi(int src[2048][2048],
          int dst[2048][2048])
{
    int i,j;
    for (i = 0; i < 2048; i++)
        for (j = 0; j < 2048; j++)
            dst[i][j] = src[i][j];
}
```

```
void copyj(int src[2048][2048],
          int dst[2048][2048])
{
    int i,j;
    for (j = 0; j < 2048; j++)
        for (i = 0; i < 2048; i++)
            dst[i][j] = src[i][j];
}
```

**21 times slower
(Pentium 4)**

- **Hierarchical memory organization (caches)**
- **Performance depends on access patterns**
 - Including how step through multi-dimensional array

Carnegie Mellon

Great Reality #5: Computers do more than execute programs

- **They need to get data in and out**
 - I/O system critical to program reliability and performance
- **They communicate with each other over networks**
 - Many system-level issues arise in presence of network
 - Concurrent operations by autonomous processes
 - Coping with unreliable media
 - Cross platform compatibility
 - Complex performance issues

Carnegie Mellon

Overview

- **Course role and theme**
- **Five realities**
- **Logistics**

Carnegie Mellon

Teaching staff

- **Instructors**
 - Prof. Roger Dannenberg
 - Prof. Greg Ganger
- **TA's**
 - Ben Blum
 - Tessa Eng
 - Jonathan Harbuck
 - Teddy Martin
 - Hunter Pitelka
 - Josh Primera
 - Sean Stangl
 - Tom Tuttle
- **Course Admin**
 - Bara Ammoura ("ECE Course Hub", Hamerschlag Hall, D-level, cube A-10)

We're glad to talk with you, but please send email first or come to office hours.

Textbooks

- **Randal E. Bryant and David R. O'Hallaron,**
 - "Computer Systems: A Programmer's Perspective", Prentice Hall 2003.
 - <http://csapp.cs.cmu.edu>
 - This book really matters for the course!
 - How to solve labs
 - Practice problems typical of exam problems
- **Brian Kernighan and Dennis Ritchie,**
 - "The C Programming Language, Second Edition", Prentice Hall, 1988

Course Components

- **Lectures**
 - Higher level concepts
- **Recitations**
 - Applied concepts, important tools and skills for labs, clarification of lectures, exam coverage
- **Labs (6)**
 - The heart of the course
 - 2 or 3 weeks
 - Provide in-depth understanding of an aspect of systems
 - Programming and measurement
- **Exams (2 + final)**
 - Test your understanding of concepts & mathematical principles

Getting Help

- **Class Web Page**
 - <http://www.cs.cmu.edu/~213>
 - Copies of lectures, assignments, exams, solutions
 - Clarifications to assignments
- **Message Board**
 - <http://autolab.cs.cmu.edu>
 - Clarifications to assignments, general discussion
 - The only board your instructors will be monitoring (No blackboard or Andrew)

Getting Help

- **Staff mailing list**
 - 15-213-staff@cs.cmu.edu
 - "The autolab server is down!"
 - "Who should I talk to about ..."
 - "This code {...}, which I don't want to post to the bboard, causes my computer to melt into slag."
- **Teaching assistants**
 - I don't get "associativity"...
 - Office hours, e-mail, by appointment
 - Please send mail to 15-213-staff, *not a randomly-selected TA*
- **Professors**
 - Office hours or appointment
 - "Should I drop the class?" "A TA said ... but ..."

Policies: Assignments (Labs) And Exams

- **Work groups**
 - You must work alone on all but final lab (see Syllabus!)
- **Handins**
 - Assignments due at 11:59pm on Tues or Thurs evening
 - Electronic handins using Autolab (no exceptions!).
- **Conflict exams, other irreducible conflicts**
 - OK, but must make PRIOR arrangements with Prof. Dannenberg/Ganger
- **Appealing grades**
 - Within 7 days of completion of grading.
 - Following procedure described in syllabus

Autolab Web Service

- **Labs are provided by the Autolab system**
 - Autograding handin system developed in 2003 by Dave O'Hallaron
 - Apache Web server + Perl CGI programs
 - Beta tested Fall 2003, very stable by now
- **With Autolab you can use your Web browser to:**
 - Review lab notes, clarifications
 - Download the lab materials
 - Stream autoresults to a *class status Web page* as you work.
 - Handin your code for autograding by the Autolab server.
 - View the complete history of your code handins, autoresult submissions, autograding reports, and instructor evaluations.
 - View the class status page

Carnegie Mellon

Facilities

- **Labs will use the Intel Computer Systems Cluster (aka “the fish machines”)**
 - 15 Pentium Xeon servers donated by Intel for CS 213
 - Dual 3.2 Ghz 64-bit (EM64T) Nocona Xeon processors
 - 2 GB, 400 MHz DDR2 SDRAM memory
 - Rack mounted in the 3rd floor Wean Hall machine room.
 - Your accounts are ready nearing readiness.
- **Getting help with the cluster machines:**
 - See course Web page for login directions
 - Please direct questions to your TA’s first

Carnegie Mellon

Timeliness

- **Grace days**
 - **4 for the course**
 - Covers scheduling crunch, out-of-town trips, illnesses, minor setbacks
 - Save them until late in the term!
- **Lateness penalties**
 - Once grace days used up, get penalized 15%/day
 - Typically shut off all handins 2—3 days after due date
- **Catastrophic events**
 - Major illness, death in family, ...
 - Work with your academic advisor to formulate plan for getting back on track
- **Advice**
 - Once you start running late, it’s really hard to catch up

Carnegie Mellon

Cheating

- **What is cheating? (see Syllabus!)**
 - Sharing code: either by copying, retyping, looking at, or supplying a copy of a file
 - Coaching: helping your friend to write a lab, line by line
 - Copying code from previous course or from elsewhere on WWW
 - Only allowed to use code we supply, or from CS:APP website
- **What is NOT cheating?**
 - Explaining how to use systems or tools
 - Helping others with high-level design issues
- **Penalty for cheating:**
 - Removal from course with failing grade
- **Detection of cheating:**
 - We do check and our tools for doing this are much better than you think!

Carnegie Mellon

Policies: Grading

- **Exams: weighted ¼, ¼, ½ (final)**
- **Labs: weighted according to effort (determined near the end)**
- **The worse of lab score and exam score is weighted 60%, the better 40%:**
 - Lab score: $0 \leq L \leq 100$,
 - Exam score: $0 \leq E \leq 100$
 - **Total score: $0.6 \min(L, E) + 0.4 \max(L, E)$**
- **Guaranteed:**
 - > 90%: A
 - > 80%: B
 - > 70%: C

Carnegie Mellon

Have Fun!