

15-213: Introduction to Computer Systems

Randy Bryant(*) and Andreas G. Nowatzky(**)

(*) School of Computer Science
(**) Robotics Institute and Dept of Electrical and Computer Engineering
Carnegie Mellon University
Fall 2004

1 Organization

Instructors:

Randal E. Bryant	Andreas G. Nowatzky
NSH 4305	NSH 4117
x8-8821	x8-4846
Randy.Bryant@cs.cmu.edu	agn@cs.cmu.edu
Fri., 3–4 pm	Wed., 3–4 pm

TAs:

Benoit Hudson	Naju Mancheril	Chris Rotella
Section E	Section B	Section C
WeH 8019	Wean Cluster	Wean Cluster
x8-5143		
bhudson@cs	ngm@andrew	crotella@andrew
Tues., 7–8 pm	Mon., 5:30–6:30 pm	Tues., 5:30–6:30 pm

Minglong Shao	Craig Soules
Section D	Section A
WeH 1315	HH D-D6
x8-5941	x8-4267
shaoml@cs	soules@ece
Thurs., 5–6 pm	Wed., 4–5 pm

Please see the course Web page for up-to-date office hours.

Course Assistant:

Cindy Chemsak, NSH 4303, x8-7884, cindyc@cs.cmu.edu

Lecture:

Tue and Thu 9:00–10:20am, Wean Hall 7500

Recitations:

A	Mon	10:30–11:20	OSC (Old Student Center) 200	Craig
B	Mon	11:30–12:20	OSC (Old Student Center) 200	Naju
C	Mon	12:30–1:20	OSC (Old Student Center) 200	Chris
D	Mon	1:30–2:20	OSC (Old Student Center) 200	Minglong
E	Mon	2:30–3:20	OSC (Old Student Center) 200	Benoit

Web page: www.cs.cmu.edu/afs/cs/academic/class/15213-f04/www/

Newsgroup: cmu.cs.class.cs213

2 Objectives

Our aim in CS 213 is to help you become a better programmer by teaching you the basic concepts underlying all computer systems. We want you to learn what really happens when your programs run, so that when things go wrong (as they always do) you will have the intellectual tools to solve the problem.

Why do you need to understand computer systems if you do all of your programming in high level languages? In most of computer science, we're pushed to make abstractions and stay within their frameworks. But, any abstraction ignores effects that can become critical. As an analogy, Newtonian mechanics ignores relativistic effects. The Newtonian abstraction is completely appropriate for bodies moving at less than $0.1c$, but higher speeds require working at a greater level of detail.

Oversimplifying matters somewhat, our $21x$ sequence works as follows: 211 is based on a simplified model of program execution. 212 builds further layers of abstraction. 213 introduces greater detail about system behavior and operation. This greater detail is needed for optimizing program performance, for working within the finite memory and word size constraints of computers, and for systems-level programming.

The following “realities” are some of the major areas where the abstractions we teach in 211/212 break down:

1. *Int's are not integers, Float's are not reals.* Our finite representations of numbers have significant limitations, and because of these limitations we sometimes have to think in terms of bit-level representations.
2. *You've got to know assembly language.* Even if you never write programs in assembly, The behavior of a program cannot be understood sometimes purely based on the abstraction of a high-level language. Further, understanding the effects of bugs requires familiarity with the machine-level model.
3. *Memory matters.* Computer memory is not unbounded. It must be allocated and managed. Memory referencing errors are especially pernicious. An erroneous updating of one object can cause a change in some logically unrelated object. Also, the combination of caching and virtual memory provides the functionality of a uniform unbounded address space, but not the performance.

4. *There is more to performance than asymptotic complexity.* Constant factors also matter. There are systematic ways to evaluate and improve program performance
5. *Computers do more than execute instructions.* They also need to get data in and out and they interact with other systems over networks.

By the end of the course you will understand these “realities” in some detail. As a result, you will be prepared to take any of the upper level systems classes at Carnegie Mellon (both CS and ECE). Even more important, you will have learned skills and knowledge that will help you throughout your career.

3 Textbook

The primary textbook for the course is

Randal E. Bryant and David R. O’Hallaron, *Computer Systems: A Programmer’s Perspective*, Prentice Hall, 2003.

In addition, we require you to have the following reference book on the C programming language:

Brian W. Kernighan and Dennis M. Ritchie, *The C Programming Language, Second Edition*, Prentice Hall, 1988.

This the classic *K & R* book, the standard against which all reference manuals are compared. It is an essential part of every computer scientist’s library.

4 Course Organization

Your participation in the course will involve five forms of activity:

1. Attending the lectures.
2. Preparing for and participating in the recitations.
3. Laboratory assignments.
4. Reading the text.
5. Exams

Attendance will not be taken at the lectures or recitation sections. You will be considered responsible for all material presented at the lectures and recitations. Lectures will cover higher-level concepts. Recitations will be more applied, covering important “how-to’s”, especially in using tools that will help you do the labs. In addition, the recitations will help clarify lecture topics and describe exam coverage.

The textbook contains both *practice problems* within the chapter text and *homework problems* at the end of each chapter. The intention is that you work on the practice problems right as you are reading the book. The answers to these problems are at the end of each chapter. Our experience has been that trying out the concepts on simple examples helps make the ideas more concrete. In addition, the schedule (at the end of this document and on the class web page) shows specific homework problems with each lecture topic. The intention is that you try these out and discuss them in the next recitation. You will find that you will get much more out of recitation if you have done some advance preparation.

The only graded assignments in this class will be a set of seven labs. Some of these are fairly short, requiring just one week, while others are more ambitious, requiring several weeks. Most labs will be handed out in class on Thursday and due either one or two weeks later on a Thursday.

5 Getting Help

For urgent communication with the teaching staff, it is best to send electronic mail (preferred) or to phone.

If you want to talk to a staff member in person, remember that our posted office hours are merely nominal times when we guarantee that we will be in our offices. You are always welcome to visit us outside of office hours if you need help or want to talk about the course. However, we ask that you follow a few simple guidelines:

- Prof. Bryant is glad to meet with students, but he's also got a lot of other meetings. The best way to meet with him is to set up an appointment through his assistant, Cindy Chemsak (cindyc@cs.cmu.edu).
- Prof. Nowatzky normally work with his office door open and welcome visits from students whenever his doors are open. However, if his door is closed, he is busy with a meeting or a phone call and should not be disturbed.
- The graduate student TAs share offices with other students. To avoid disturbing these students, please send mail before visiting a TA outside of office hours so he/she can arrange to meet you.

We will use the Web as the central repository for all information about the class. The class home page is at

www.cs.cmu.edu/afs/cs/academic/class/15213-f04/www/

Using the Web, you can:

- Obtain copies of any handouts or assignments. This is especially useful if you miss class or you lose your copy.
- Read clarifications and changes made to any assignments, schedules, or policies.
- Find links to any electronic data you need for your assignments

We have also set up a news group for this class, `cmu.cs.class.cs213`. This group will be used by members of the teaching staff to post announcements and clarifications. You may also post to this group to make queries.

6 Policies

Working Alone on Assignments

You will work on all assignments by yourself.

Handing in Assignments

All assignments are due at 11:59pm (one minute before midnight) on the specified due date. All handins are electronic.

Handing in Late Assignments

Each student will receive a budget of five *grace days* for the course. Here is how grace days work:

- If you hand in an assignment k days late, then you receive full credit for the lab, but you will have spent k of your grace days. For example, if an assignment is due at 11:59pm on Thursday and you hand it in at noon on Saturday, then you will have spent 2 grace days. If you hand it in at 9am on Friday, then you will have spent 1 grace day.
- When you are out of grace days, you can no longer hand in assignments, and must make special arrangements with your professor, as described in the “Making up Exams and Assignments” section.
- Regardless of the number of grace days you have remaining, handins will not be accepted after the *end date* of the lab, which is typically 2 days after the due date.

Grace days are a tool to help you manage your time and to help smooth out burstiness in assignment due dates. We recommend that you conserve your grace days, saving them for the end of the term when things get most hectic.

Making up Exams and Assignments

Missed exams and assignments more than 2 days late can be made up on a case by case basis, but only if you make prior arrangements with Prof. Nowatzky. However you should have a good reason for doing so. It is your responsibility to get your assignments done on time. Be sure to work far enough in advance to avoid unexpected problems, such as illness, unreliable or overloaded computer systems, etc.

Appealing Grades

After each exam, homework, and assignment is graded, Prof. Nowatzky will send each of you a personalized email with your grade (as well as all of your previous grades). You have seven calendar days from the date he sends the email to appeal your grade.

Each TA in 15-213 has the authority to unilaterally change your lab grade, without permission from the professors. So if you have questions about the grade you received on an assignment (homework or lab), please talk first to the person in charge of the assignment, who will be clearly identified in the writeup.

If you are still not satisfied, please come and visit Prof. Nowatzky. If you have questions about an exam grade, please visit Prof. Nowatzky directly.

Final Grade Assignment

Each student will receive a numeric score for the course, based on a weighted average of the following:

- **Assignments:** The assignments will count a combined total of 60% of your score. The exact weighting of the different assignments will be determined near the end of the course based on our perception of the relative effort required. In any case, each lab will count 6–12% of your score. Since small differences in scores can make the difference between two letter grades, you'll want to make a serious effort on each assignment.
- **Exams:** There will be two in-class exams, each counting 10%, plus a final counting 20%.

Grades for the course will be determined by a method that combines both curving and absolute standards. The total score will be plotted as a histogram. Cutoff points are determined by examining the quality of work by students on the borderlines. Individual cases, especially those near the cutoff points may be adjusted upward or downward based on factors such as attendance, class participation, improvement throughout the course, final exam performance, and special circumstances.

Cheating

Each lab assignment must be the sole work of the student turning it in. Assignments will be closely monitored by automatic cheat checkers, and students may be asked to explain any suspicious similarities. The following are guidelines on what collaboration is authorized and what is not:

What is Cheating?

- *Sharing code or other electronic files:* either by copying, retyping, looking at, or supplying a copy of a file.
- *Sharing written assignments:* Looking at, copying, or supplying an assignment.

What is NOT Cheating?

- Clarifying ambiguities or vague points in class handouts or textbooks.
- Helping others use the computer systems, networks, compilers, debuggers, profilers, or other system facilities.
- Helping others with high-level design issues.

- Helping others debug their code.

Be sure to store your work in protected directories.

The usual penalty for cheating is to be removed from the course with a failing grade. We also place a record of the incident in the student's permanent record.

7 Facilities: Intel Computer Systems Cluster

Intel has generously donated a cluster of 25 Linux-based Pentium III servers, specifically for 15-213, that we will use for all labs and assignments. The class Web page has details.

8 Class Schedule

Figure 1 shows the tentative schedule for the class. The reading assignments are all in the new book. You can check on the class web page for suggested readings in K&R. The schedule also indicates suggested homework problems, the lab activities, and the lecturer for each class.

Any changes will be announced on the class news group. An updated schedule will be maintained on the class web page.

Class	Date	Day	Topic	Reading	Problems	Labs	Lecturer
1	08/31	Tue	Overview	1			Both
2	09/02	Thu	Bits and Bytes	2.1	2.44, 2.45,	L1 Out	AGN
3	09/07	Tue	Integers	2.2–2.3	2.49, 2.54		REB
4	09/09	Thu	Floating Point	2.4–2.5	2.59, 2.60, 2.61		REB
5	09/14	Tue	Machine Prog I - Overview	3.1–3.5	3.31		REB
6	09/16	Thu	Machine Prog II - Control	3.6	3.34	L1 Due, L2 Out	REB
7	09/21	Tue	Machine Prog III- Procs	3.7			REB
8	09/23	Thu	Machine Prog IV- Data	3.8–3.11	3.36		REB
9	09/28	Tue	Mach. Prog V - Advanced	3.12–3.13, 3.16	3.24		AGN
10	09/30	Thu	Program Optimization	5		L2 Due, L3 Out	AGN
11	10/05	Tue	Memory Hierarchy	6.1–6.4	6.21, 6.23, 6.24		AGN
12	10/07	Thu	Cache Memories	6.5–6.8	6.25, 6.26, 6.27	L3 Due	AGN
13	10/12	Tue	Exam 1	Time and Location TBA			
14	10/14	Thu	Linking	7	7.8, 7.12	L4 Out	AGN
15	10/19	Tue	Except. Control Flow I	8.1–8.4	8.10, 8.11		AGN
16	10/21	Thu	Except. Control Flow II	8.5–8.8	8.19		AGN
17	10/26	Tue	Virtual Memory	10.1–10.6	10.11, 10.12, 10.13		AGN
18	10/28	Thu	P6/Linux Memory System	10.7–10.8	10.14	L4 Due, L5 Out	REB
19	11/02	Tue	Dynamic Storage Alloc I	10.9	10.15, 10.16		AGN
20	11/04	Thu	Dynamic Storage Alloc II	10.10–10.13	10.18	L5 Due, L6 Out	AGN
21	11/09	Tue	Debugging				AGN
22	11/11	Thu	System-level I/O	11	11.7		REB
23	11/16	Tue	Exam 2	Time and Location TBA			
24	11/18	Thu	Internetworking	12.1–12.3		L6 Due, L7 Out	REB
25	11/23	Tue	Network Programming	12.4			REB
	11/25	Thu	Thanksgiving	12.4			
26	11/30	Tue	Web Services	12.5–12.7	12.10		REB
27	12/02	Thu	Concurrency	13.1–13.3	13.22, 13.23		REB
28	12/07	Tue	Synchronization	13.4–13.8	13.24		REB
29	12/09	Thu	Time Measurement	9		L7 Due	AGN

Figure 1: CS 213 Class Schedule