

15-213/15-513/18-213: Introduction to Computer Systems

Gregory Kesden
Carnegie Mellon University
Summer 2014

1 Organization

Class Web Page: <http://www.cs.cmu.edu/~213>

Note: Electronic copies of class handouts and lectures can be found on the class web page.

Note: Multimedia lecture presentations via Panopto are also available to registered students via the link from the lecture Web page.

Contacting Teaching Staff: Send email to 15-213-staff@cs.cmu.edu. This will go to all of the teaching staff.

Note: Please use 15-213-staff@cs.cmu.edu for all questions and concerns, rather than single-person emails, other than scheduling of one-on-one meetings. You will get faster answers, and we will better be able to provide consistency.

Instructors:

Gregory Kesden
gkesden@cs.cmu.edu
GHC 7711, x8-1590
<http://www.cs.cmu.edu/~gkesden/schedule.html>

TAs:

- Shailendra Bingi
- Marjorie Carlson
- Thomas O'Connor
- Vishvesh Selvakumaar
- Shivaraman Shankar
- Aditya Gupta
- Vinay Venkatesh

Lectures:

- First 6 weeks: Tuesday, Wednesday, and Thursday 12:00 - 1:30pm, DH A-302
- Second 6 weeks: Tuesday and Thursday 12:00 - 1:30pm, DH A-302

Recitations:

Thursdays, 4:30-5:50pm, DH A-302

TA Office Hours:

To be selected during the first class

2 Objectives

Our aim in 15-213/15-513/18-213 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 sequence works as follows: 112 is based on a simplified model of program execution. 122 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 122/150/210 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, 2nd Ed.*, Prentice Hall, 2011.

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 is 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 or otherwise utilize the lectures.
2. Preparing for and participating in the recitations.
3. Laboratory assignments.
4. Reading the text.
5. Exams

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 six labs. Some of these are fairly short, requiring just one week, while others are more ambitious, requiring several weeks.

5 Getting Help

For urgent communication with the teaching staff, please send email to 15-213-staff@cs.cmu.edu.

We will use the class website (<http://www.cs.cmu.edu/~213>) and class message board (*Autolab*) as the central repositories for all information about the class. Using these resources, you can:

- Obtain copies of any handouts or assignments. This is especially useful if you miss class or you lose your copy.
- Find links to any electronic data you need for your assignments
- Read clarifications and changes made to any assignments, schedules, or policies.
- Post messages to make queries about the course, specific labs, or exams. (And, of course, see responses to such posts.)

The lab assignments and class message board are offered through a Web service written by Prof. David O'Hallaron called *Autolab*. See the Autolab web page at <http://autolab.cs.cmu.edu> for more information.

If you want to talk to a staff member in person, the posted office hours are the best opportunity, as they represent times when we guarantee that we will be in the location identified. If a meeting is needed outside of the office hours, please use email to arrange a time. (Note that 15-213-staff@cs.cmu.edu reaches all teaching staff, maximizing your odds of a rapid turnaround.)

6 Policies

Working Alone on Assignments

You will work on all assignments by yourself, unless notified otherwise in the assignment description.

Handing in Assignments

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

Handing in Late Assignments

Each student will receive a budget of five *grace days* for the course. These grace days are provided to allow you to cope with most emergencies that prevent completing a lab on time, including computer problems, a cold, getting stuck at the airport, etc. 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. However, a maximum of 2 grace days may be used for any assignment. 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, assignments turned in before the *end date* (see next bullet) will receive a 15% late penalty per day. So, for example, an assignment turned two days late would receive a 0% penalty if you had two or more grace days, a 15% penalty if you had one grace day, and a 30% penalty if you had zero grace days.
- 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 3 days after the due date.

Grace days are a tool to allow you to manage your time in the face of personal issues and to help smooth out burstiness in assignment due dates across classes. They are for when you are sick, when a short-term emergency situation arises, when you have too many deadlines all at once, etc. Except for serious persistent personal issues (see below), you should not anticipate additional deadline leniency. We recommend that you conserve your grace days, saving them for true emergencies and time pressures at the end of the term.

Dealing with Serious Persistent Personal Issues

We hope that everyone in 15-213 will remain happy and healthy. But, if you have a serious persistent personal issue, such as being hospitalized for an extended period or needing to leave the country for a family matter, please talk to your academic advisor as soon as possible. Such issues consistently affect one's ability to succeed in all classes, rather than just 15-213, and the academic advisors are equipped to coordinate plans for dealing with them. We will cooperate with such plans, but we cannot construct them independently of the academic advisors.

Requesting a Re-Grade for an Assignment or an Exam

After each exam and lab assignment is graded, a personalized email will be sent to each of you with your grade (as well as all of your previous grades). We will make the utmost effort to be fair and consistent in our grading. But, we are human. If you believe that you did not receive appropriate credit for an assignment or an exam, you may request a re-grade as follows:

- Submit your request *in writing* within seven calendar days of when the grade notification is sent to you via email, explaining in detail why you think that there was a mistake in the grading. Please note that verbal requests will not be processed; requests must be in writing.

- These requests should be submitted to `15-213-staff@cs.cmu.edu`, whether for an assignment or for an exam. In the case of exams, the exam in question should be hand-delivered to the course assistant together with a printed copy of the email request.
- When you submit a request for a re-grade, the entire assignment or exam may be re-graded (not just the parts that you specify). Your grade may go up or down (or stay the same) as a result of the re-grade request.

Your request will be processed off-line, and we will respond to your request as quickly as possible (typically within a week). This re-grade policy is designed to correct legitimate mistakes in grading, while discouraging frivolous re-grade requests (for the sake of being fair and consistent across the entire class).

Final Grade Assignment

To do well in this course, you must do well on both the lab assignments *and* the exams. These two different forms of evaluation test different aspects of the material—how well you understand the underlying concepts (exams) and how well you can put the concepts into practice (labs).

You will receive an overall score, between 0 and 100 points, for the course. This score is based on both your lab and exam averages. The exam average is 50% of the final grade, the labs weigh 50% of the final grade.

The goal of this grading scheme is to ensure that each student has the necessary knowledge, skill, and experience to succeed in upper-division systems classes.

Local students; The final exam counts 30
Distance students; The final exam counts 35

The weight of each lab is proportional to the effort required for that lab. In order to allow proper tuning of the course, the exact weight of each lab is not determined until very late in the semester. The weight of an individual lab is likely to vary from 10-percent of the lab average for a simple, short lab, to 25-percent of the lab average for a long, complex lab.

Given the overall score grades for the course will be determined by a method that combines both curving and absolute standards. The overall scores will be plotted as a histogram. Cutoff points are determined by examining the quality of work by students on the borderlines.

Cheating

Each exam and lab assignment must be the sole work of the student turning it in. Assignments will be closely monitored by automatic cheat checkers, including comparing turned-in code to the work of students from the same and previous semesters, and students may be asked to explain any suspicious similarities. These cheat checkers are very effective, having been refined over years of research, and they are not fooled by attempts to mask copying of code. Please don't try your luck.

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.

The following are guidelines on what non-exam 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 from this or a previous semester.
- *Sharing written assignments:* Looking at, copying, or supplying an assignment.
- *Using other's code.* Using code from this or previous offerings of 15-213, as well as from courses at other institutions.

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 with high-level (not code-based) debugging..
- Using code from the CS:APP website or from the course web pages.

Be sure to store your work in protected directories.

Electronically Mediated Communication

You are not permitted to engage in any electronically mediated communication during class – the penalty, even for a first offense, is an "R" in the course. No TXTing. No IMing. No email. No surfing the Web. No net games. No talking on your cell phone. You get the idea: If something is being communicated and anything even vaguely electronic is involved, you flunk. If you are not comfortable with this – find a different section.

This might seem extreme to you, but it necessary. The ubiquity and pervasiveness of these technologies is threatening the classroom dynamic. Unless it is controlled, it will become cultural – and the loss to education permanent and tremendous.

Since it is somewhat difficult to prove allegations of this kind from the front of the classroom, the penalty is severe to act as a deterrent. Other instructors, student assistants, staff members, and/or other trusted individuals may be asked to monitor the classroom from various points of observation to aid in the enforcement of this policy and still photographic evidence may be acquired and used to document violations.

Recording

High-fidelity recordings of class, including any audio and/or video recordings, regardless of the media or format, and regardless of the intended or actual use, are not permitted without prior written permission of the instructor. The class will be notified in advance should any such recording be approved. The penalty for violating this policy is an R in the course. If you are not comfortable with this, drop the course now. Students have no right to record classes under any University policy.

This policy is intended to protect the privacy of the students. No student should run the risk of potential employers finding a question or incorrect answer on the Web. The classroom is a learning environment, not an exhibition. Rather than attempt to control the uncontrollable or distinguish between neutral and detrimental uses, all recording is prohibited. Experience has shown that, excluding special cases such as use by students with disabilities or distance learners, undergraduate students do not improve their performance through the use of high fidelity recordings.

Official Panopto presentations of the class are available for your use from our server during this instance of the class only. They will be deactivated after the final exam. You are not permitted to download or otherwise record these videos for your future use, or use of any kind. You are only permitted to view them streamed from our servers.

MIME Attachments

Unless specifically requested, please send only plain-text email to course staff – never send MIME attachments. These bog down the staff's ability to respond to email via cell phones and other light clients.

7 Facilities: Intel Computer Systems Cluster

Intel has generously donated a cluster of 15 Linux-based 64-bit Xeon servers, specifically for 15-213, that we will use for all labs and assignments. The class Web page has details.

8 Class Schedule

Please see the schedule maintained on the class Web page for information about lectures, reading assignments, suggested homework problems, lab start and end dates, and the lecturer for each class. The reading assignments are all from the CS:APP book.