

UNIVERSITY OF CALIFORNIA
Department of Electrical Engineering
and Computer Sciences
Computer Science Division

Prof. R. Fateman

Fall 2002

General Course Information: CS 164 Programming Languages and Compilers

Personnel

Instructor: Richard J. Fateman, 789 Soda (fateman@CS, cs164@cs)
Teaching Assistants: Ibrahim Merchant (Abe), and Ryan Stejskal
Class Home Page: <http://www-inst.eecs.berkeley.edu/~cs164>.

Purpose of this Course

CS164 is an introduction to the design and implementation of programming languages.

In addition to study of the theory of programming language design and tools, you will learn about the implementation of programming languages through a set of programming exercises. These, taken together comprise most of the parts of a compiler.

We will also discuss a variety of historically influential programming language design ideas as well as current directions in languages.

Each instructor in CS164 provides his or her own stamp on the project sequence. This semester we will ultimately be implementing a language described in the main textbook for this semester. This is a simple declarative language with nested functions, records, arrays, integers, and a few simple control constructs. Although the text uses Java as an implementation language, we will be using COMMON LISP . Why?

1. COMMON LISP will show you what an advanced language of the Lisp family can provide in terms of fast prototyping, on-line debugging, etc.
2. COMMON LISP is already somewhat familiar to you from your acquaintance with Scheme.
3. Language implementation is a fundamental model of program development in COMMON LISP : in CS61a you've seen how Logo can be implemented. Appropriate tools do not have to be re-invented.
4. The model of language implementation by interpretation (the Scheme evaluator) is a powerful idea you've already seen. We will revisit this. Yet generating code, which is the role of a compiler, is surprisingly not much harder, and fits a similar pattern. Since we will generate and execute "virtual machine code" some of the ideas behind Java's machine implementation will emerge as well.

5. Because the tasks you are required to perform can be done with so much less code (and we hope less time and effort), the class will have time to study programming language topics beyond those illustrated by the project. Also, if you allocate your time reasonably well, you will be able to finish your assignments without staying up all night.

Prerequisites

Please take this course *only if you have the prerequisites*: the CS 61ABC sequence. You are assumed to be familiar with Scheme and *skilled* at reading and writing programs in at least one other higher-level language (most likely Java or C++). You are not expected to know COMMON LISP yet, but you will pick up this language easily. It is like Scheme on steroids. We assume you are capable of using a network browser, email, newsgroups, etc. You will use only a small subset of tools in UNIX if you use the class accounts; it is possible to run the same tools on Windows NT or 2000.

While I understand that people can become attached to other (non-emacs) editors, it seems to me the time needed to become familiar with emacs is repaid many times over by its greater level of efficiency in writing and debugging lisp or other programming languages. (Many people unfamiliar with emacs and lisp think that a major issue in lisp programming is “balancing parentheses”. This is not an issue since the editor shows you, by indentation, how expressions are grouped.) Of course there are many interesting additional features of emacs, not the least of which is the possibility of extending its features by writing emacs-lisp code!

Scheduling

There are two lecture meetings each week, meeting in Room 160 Kroeber Hall, Tu/Th 3:30-5:00. The course/lecture slides will be available on-line; sometimes copies may be handed out at lecture. The slides alone are not intended to be a substitute for attendance.

Meeting Times

The meeting times for the 6 discussion sections place them all before the Tuesday lecture. When there is a holiday on Monday (Sept. 2, Nov. 11), the students in the two Monday sections will have to find a Tuesday section. Otherwise you should attend the same discussion section each week. The TAs may consolidate the discussion sections depending on actual enrollment, and will schedule office hours after checking with students. You may change sections *WITH PERMISSION OF THE TAs*.

meeting	days	time	place	person
lecture 1	Tu/Th	3:30–5:00	160 Kroeber Hall	Fateman
disc. 101	M	3:00–4:00PM	405 Davis	TBA
disc. 102	M	4:00–5:00PM	70 Evans	TBA
disc. 103	Tu	9:00–10:00AM	3109 Etcheverry	TBA
disc. 104	Tu	10:00–11:00AM	3111 Etcheverry	TBA
disc. 105	Tu	11:00–12:00PM	405 Davis	TBA
disc. 106	Tu	12:00–1:00PM	310 Hearst Mining	TBA

Grades

Important: Your grade will not depend upon the grades of others in the course. In particular, we will not limit the number of A grades to some percentage of the class. All students who have done excellent work on exams and assignments will get “excellent” grades.

Past experience suggests that students want to keep track of how they are doing relative to others, and how they should allocate their time among their courses. Most of the activities within CS164 are aimed at helping you learn the material and demonstrate your knowledge on exams. But we will also count other aspects of your performance toward your grade.

Your ranking in this course will depend on four components:

1. Programming Assignments: We are planning on 7 programming assignments. You will be expected to turn in your answers in electronic form by the time given in the assignment. For late assignments we will generally assign a grade of 0 for all questions whose answers have been discussed in class or section (whether or not you attended). To be safe, please turn in your results by 11:59PM on the day assigned (unless otherwise indicated in the assignment). After that, your grade will be decreased by 50% per 24-hour period, or fraction. Not all the programming tasks are expected to be equally challenging. Assignments 1 and 4 should be easier; the hardest is probably assignment 6.

Programming Assignments in total will account for 30% of your course grade.

2. Exams: There will be 2 examinations and a final. These will represent 10, 10, and 30% of your course grade for a total of 50%. We may experiment with in-class brief quizzes factored in to this portion of your grade. Exams will probably be open-book.

This course is scheduled for final exam group 19, which is Wednesday, December 18, 2002 from 12:30-3:30PM. It should be impossible for you to have scheduled another class in the same final exam group.

3. Problem Sets: We may experiment with a few of these (we didn't have any last Fall, and some students wanted to be given some “study questions” for exams). These may be more likely near the beginning of the semester when we are doing more “theory.” The intent is to help you anticipate questions on the exams and final. It will ordinarily be simpler for you to write these out on paper rather than type them in to a computer system, and they will be turned in at class meetings or the homework box opposite 283 Soda Hall. The grade here will represent *no more than* 15% of your course grade.

4. Participation in discussion: Remaining points: your teaching assistant will judge your performance on the basis of in-class participation, any scores on discussion-section quizzes, and partnership cooperation.

Our intention is to provide an absolute grading scale to emphasize that students are *not* competing against each other. The number of A grades is not limited. In the past we have been successful in starting out with 200 possible points, 185–200 is an A+, 175–184 is an A, 165–174 is an A-, with 10 points each for B+, B, B-, C+, C, C-, D+, D, (and below is F). If exams turn out to be harder than expected, we may modify the scale.

Partnerships and Original Work

There is a tradition in this course providing students an opportunity for programming in teams. These partnerships (maximum 2 persons) will receive the same grade on programming assignments. Problem sets and exams must be done individually. Since the exams will contain questions about the program modules, it is extremely unwise to divide up your projects so that each partner does all the work on “half” the modules. In fact, quite a few questions on the exams will be quite unanswerable if you have not done the programming (Or studied your partner’s pieces *verrrry* carefully).

The advantage of partnerships is largely in the ability to practice the social aspects of computing, including having someone else comment on your programs. It is especially helping to find elusive bugs. Partnerships are completely optional, however. You may work alone for one or all of your assignments. Your TA will be monitoring the health of your project partnership.

While you may discuss questions of strategy in your projects with staff, or another student who is not your partner, you *must not take someone else’s code and make it seem as if it is your own*. If you have discussed ideas with others, including the TAs, the professor, and other students, you should give proper credit in comments. For example, “Thanks to xyz@eecs who helped me find the bug in this” or “the idea for storing symbols this way was suggested by abc@eecs”. Clever computer programs have been devised to detect unusual similarities between submitted programs.

In summary: If you visit the staff during office hours and we tell you how to solve a problem, please give us credit. If you use ideas from other students, you must say so: give them credit.

General hints: start your assignments early and ask for help from the staff if you need it. Leave copies of early and intermediate development files on our computers’ disk systems so you will not lose everything on your home computer. This way you also have some evidence of partial efforts on our disk archives. If you work at home, you should nevertheless read and use the newsgroup `ucb.class.cs164`, where many questions and answers will appear. Usually an answer to a well-posed question appears within an hour, at least during the day. Frequently the answer is correct. *Most people benefit greatly from the newsgroup. Provide credit in your submissions!*

Required Texts

There is still no ideal text for CS164 as taught at UCB. In addition to the course notes and slides, we tried the (required!) Appel text for the first time last Fall. It is more up-to-date than the Aho text (optional), but is somewhat brisk in places; it advocates writing a language implementation in Java. Design of the Java code given in the text may be helpful, but detailed examination of the jumble of stuff needed for Java to run might possibly lead you astray. Programming in Lisp should be simpler.

Andrew Appel, *Modern Compiler Implementation in Java*, Cambridge Univ. Press (1998, reprinted with corrections, 1999) Note: don’t get the version of this book based on ML or C++.

The next (also required) book explains COMMON LISP in a very readable and literate manner. It has especially fine sections on extended examples, and we will use this as a source for programming language features issues.

Paul Graham, *ANSI Common Lisp* Prentice Hall, 1996.

There is also

G. L. Steele, Jr, *Common Lisp: the Language* 2nd edition, Digital Press, 1990.

This is the standard reference, entirely available on-line. After working on this, Guy Steele joined Sun to work on the design of Java.

Optional/Recommended

You may find it useful to have the old CS61A text, as well as your favorite book on Java. If you have never seen Java (but know C++, for example), Flanagan's *Java in a nutshell* may be helpful.

The classic compiler book, used in past CS164 classes, is difficult to read and somewhat out of date. Coverage of early topics in the course is an alternative to Appel. You may hear it referred to as the "Red Dragon" book because of its cover illustration.

A. V. Aho, R. Sethi, J. D. Ullman. *Compilers: Principles, Techniques, and Tools*, Addison Wesley, 1986. Additional notes, programs, test-files will be provided on-line in various forms.

Programming Languages

All your programming assignments will be written in COMMON LISP . We hope that by the end of the semester you will agree that you would have had much more work to do if you had been writing in C++ or Java.

Computer Accounts

If you are a CS or an EECS major, you are expected to have your own "named" computer account. Nevertheless, we still think it is useful for you to pick up a class account at the first lecture. While it is possible for you to use your named account for the course, there may still be an advantage to class accounts, sending email, setting up shared files with a partner, etc.

If you use a named account, be sure to fix up your `.emacs` file so that it includes the class-emacs initialization from the class home page.

If you are intent on using your own home computer, it is especially important to get a copy of the same software we have on our computers. You can download a (free) "Lite" version of Allegro Common Lisp that should provide you with all the facilities you need for this course, as well as documentation, on Windows and Linux. You must also make sure to use emacs. The enormously complex Allegro CL "IDE" is, in my opinion, quite unsuitable for writing simple non-graphical programs such as we will be doing. (On the other hand, if you were writing a menu-based graphical user interface, the IDE would probably help.) You can download from www.franz.com.

Note that if you depend on your home/dorm computer, and something goes wrong, we may not be able to help much. You take full responsibility for getting material to and from the computer systems on campus, making backups, etc. (And once again, make sure you have emacs set up right at home).

Web access

You should bookmark the home page for CS164, and visit it often. It provides links to lecture notes, software, etc.

<http://www-inst.eecs.berkeley.edu/~cs164>

If you need physical access to the workstations in Soda Hall after 6:30PM, you should get a card-key allowing you in to the labs. *These cards will be available to students whose names are on the official enrollment list for CS164.* These will be available in 395 Cory after the first week or so of classes; they will require a deposit.

Office Hours

person	days	time	place	e-mail	telephone
R. Fateman	T/Th	1:00–2:30PM, OBA	789 Soda	fatemann@cs	642-1879
XXX	TBA	TBA	TBA Soda	XXX@eecs	
YYY	TBA	TBA	TBA Soda	YYY@cs	
readers				???	

(OBA = Or By Appointment, TBA = To Be Announced) Other times may be posted from time to time, and other times for meetings can be arranged if you cannot make any of these times. It is a good idea to contact us by e-mail to ask questions or set up appointments. The class newsgroup can be invaluable, as well.

This is a course where 30 seconds of educated help from the staff or other students can save you 3 hours of thrashing.

While you will be doing a substantial amount of programming, we hope it is MUCH LESS than if you were writing in other languages.

Course assignments, homeworks, exam schedule

This is a *tentative* course outline including lecture topics, readings, topics for discussion section, and homeworks. Reading marked AWA are in the text by Andrew W. Appel, *Modern Compiler Implementation*. Reading marked PG are in the text by Paul Graham, *ANSI Common Lisp*.

Note that the lectures are intended to provide perspective on the readings and homework. We hope you are mature enough to read the assigned material in the week given, to help you get full value out of the lectures. **The material in the reading assignments anticipate the lectures by one week.** We might try a pop quiz or two to encourage you to keep up.

In addition to helping you understand the lectures and answer questions on exams, another benefit of reading the material is that you will be able to impress your fellow students and the

instructional staff with your ability to ask (and answer!) pertinent questions in discussion, and perhaps in lecture too!

Assignment 0

Turn in, in class on Thursday a small photograph or sketch (!) of yourself with your full name, login and preferred form of reference (Ms. Jones, Sneezy, etc.) on the back. This could be a photocopy of your photo on your registration card.

Course Outline TENTATIVE, Fall 2002 ***

Week 1:	(Aug 27, 29)
Lect:	Course overview; Languages, translation, Common Lisp.
Read:	AWA: 1.1–1.4(lightly)
Read:	PG: 1,2,3,4 (you know most of this)
Read:	PG: 7.1-7.3, maybe 7.4
	get comfy with Common Lisp
Sec:	set up computer accounts, Review Lisp vs. Scheme vs. Java
Due:	Programming Assignment 0 due 8/29
Week 2:	(Sep 3, 5)
Lect:	Tiger, Intro to Lexical Analysis
Read:	AWA: 1.4, 2, (light on p 28-30)
Read:	PG: Appendix A (p 287)
Sec:	More LISP. turning in programs; Q & A
Due:	Programming Assignment: 1 due 9/5
Week 3:	(Sep 10, 12)
Lect:	Lexical analysis, Regular expressions, finite automata
	Highlights of Tiger .
Read:	AWA: 3.1-3.2
Sec:	tools for automated LEX
Week 4:	(Sep 17, 19)
Lect:	CF Languages, Parsing, Syntactic analysis.
Read:	AWA: 3.3-3.5
Sec:	Review for exam, recognizers, NDA to DA conversion
Due:	Programming Assignment 2 (lexical analysis) due 9/19
Week 5:	(Sep 24, 26)
Lect:	Review for test; Recursive Descent, Predictive Parsing
Read:	AWA: 4.1-4.2
Sec:	Q & A, LL parser lisp tool;
Exam 1	this week. Sept 26 in class
Week 6:	(Oct 1, 3)
Lect:	Abstract Syntax, LR parsing
Read:	AWA: 5.1-5.4
Read:	PG: 13
Sec:	examples for LALR parser generator, post-exam discussion
Week 7:	(Oct 8, 10)
Lect:	Static analysis
Read:	AWA: 6.1-6.2, 7.1-7.3
Sec:	discuss parsing problems
Due:	Programming Assignment 3 (parser) due 10/11

Week 8:	(Oct 15, 17)	
Lect:	Semantic Analysis, Intro to run-time	
Read	AWA: 13.1-13.7	
Sec:	Discuss type-checking, polymorphism, etc.	
Due:	Programming Assignment 4 (parser, AST)	
Week 9:	(Oct 22, 24)	
Lect:	Language Run-time (continued)	
Read:	notes on interpreters, AWA 15.1-15.7 ***	
Sec:	Review successful parser/checker techniques	
Week 10:	(Oct 29, 31)	
Lect:	Review for Test; Interpretation	
Read:	AWA: 8, 10	
Sec:	properly tail recursive interpreter for Tiger	
Exam 2	Oct 31	
Due:	Programming Assignment 5 (interpreter for AST) 11/1	
Week 11:	(Nov 5, 7)	
Lect:	Virtual Machine, Intermediate Code generation	
Read:	AWA: 14, notes	
Sec:	Review	
	discussion of stack machine architecture, assembler	
Week 12:	(Nov 12, 14)	
Lect:	Code generation and optimization	
Read:	notes	
Sec:	More architecture review	
Due:	Programming Assignment 6 (type checker)	
Week 13:	(Nov 19, 21)	
Lect:	Prog. language features: macro/string processors, other Examples/Critiques	
Read:	PG: 11, 17, Allegro CL on-line notes	
Sec:	Code generation	
Week 14:	(Nov 26)	
Th/Fr Nov 28–29: holiday		
Lect:	Storage (GC), OOP	
Read:	notes	
Sec:	Linking, loading, testing.	
Due:	Programming Assignment 7 (code generator) due 11/26	
Week 15:	(Dec 3, 5)	
Lect:	Other languages and their implementation; review	
Sec:	Review for final, (Dec 18)	