

ECE351 Course Outline, Winter 2020

Prof Arie Gurfinkel

March 29, 2020

[DRAFT until Jan 6]

arie.gurfinkel@uwaterloo.ca

1 Introduction

Welcome to ECE351! This is an introductory compilers course. The calendar description lists a number of topics:

Programming paradigms, compilation, interpretation, virtual machines. Lexical analysis, regular expressions and finite automata. Parsing, context-free grammars and push-down automata. Semantic analysis, scope and name analysis, type checking. Intermediate representations. Control flow. Data types and storage management. Code generation.

1.1 Course Objective

The objectives of the course are two-fold: First, to teach students the theoretical aspects of compilers and language design, and second, to enable students to build a practical rudimentary compiler of their own. Monday lectures will usually discuss theory, and Friday lectures will usually discuss practice.

This course is very hands on, *i.e.*, you will write a circuit synthesizer and a circuit simulator for circuits described in a subset of VHDL. A circuit synthesizer is essentially a compiler for circuits. The lab notes describe these exercises in detail.

Through the lab you will learn two important skills: First, you will understand how language design can be seen as a general form of problem solving. Second, you will learn how to “implement” a language by constructing an appropriate compiler.

Often language design and problem solving are seen as distinct activities. The point of the lab is to show that language design, in many ways, is the most general way to solve a problem. By most general solution I mean that your solution solves an entire class of problems, not merely one problem or an instance of a problem. For example, suppose that you need to multiply two matrices. The least general, and most direct, thing you could do is to actually multiply the matrices. More generally, you could write a program that multiplies two arbitrary matrices, and then use that program to multiply your two matrices of interest. Most generally, you could design a language that not only allows you to write programs to multiply matrices but also allows you to express programs to do more complex mathematical operations. Your language could be something similar to a suitable subset of Matlab, then write a program in that language that requests your two matrices be multiplied, then execute that program

VHDL is a very complicated hardware description language to specify circuits. The language for which you will develop a synthesizer is a very small subset of VHDL.

Solve an entire class of problems by designing a language to describe those problems.

Good engineers do a cost/benefit analysis before designing and building. Good engineers also have the skills to design a range of solutions.

with the interpreter/compiler you just developed. In order to solve problems at this level of abstraction you need to understand the ideas listed in the course description, and you will need to improve your programming skills.

We will use the term *transformer* or *translator* to describe a program that reads and writes structured text. As you will see, any sufficiently large transformer is actually comprised of many small transformers: the transformation is broken down into steps that typically pass through one or more intermediate forms before the final output is produced. These intermediate forms are typically simpler to process mechanically and less convenient for human use than the original input language.

A compiler is a particular kind of transformer that reads source code and produces assembly or machine code.

For example, the GNU Compiler Collection (GCC) has an intermediate form called RTL (Register Transfer Language). C/C++/Java/Fortran/*etc.* are first translated to RTL, then optimizations are performed on the RTL, and finally assembly code is generated for some particular chip.

2 Coordinates

Instructor:	Arie Gurfinkel	DC-2522	arie.gurfinkel@uwaterloo.ca https://ece.uwaterloo.ca/~agurfink
TAs:	An Qi Zhang	in lab	aq2zhang@uwaterloo.ca
	Reza Babaei	in lab	rbabaeec@uwaterloo.ca
	Samuel Yigzaw	in lab	syigzaw@uwaterloo.ca
Regular Lectures:	MF 10:00AM–11:20AM	E7-4053	Theory Mondays, Lab Fridays (usually)
Extra Lectures:	Fri 12:30PM–1:20PM	E7-4053	See schedule
Theory Tutorials:	W 8:30AM–9:50AM	E7-4053	Even weeks (see schedule), if needed
Regular Labs:	TWTh 3PM–4:20PM	E2-2356A	http://sun5.vlsi.uwaterloo.ca/~eceptc/TimeTables/
Midterm:	Midterm Week	See schedule	https://ece.uwaterloo.ca/electrical-computer-engineering/current-undergraduate-students/academic-planning-and-support/mid-term-examinations/
Quizzes:	bi-weekly	Learn	https://learn.uwaterloo.ca/d2l/home/511703
Discussion:		Campuswire	https://campuswire.com/p/G057DD328 code: 7644
Course Notes:		Git	https://git.uwaterloo.ca/ece351-1201/pdfs
Your Code:	Thursdays Thu 10PM	Git	https://git.uwaterloo.ca/ece351-1201/yourid

Files will be distributed and collected through Git. Each student will have their own private repository. We will have shared repositories for documents and skeleton code. Your Nexus credentials (username and password) should give you appropriate access to the repositories.

Lecture topic schedule is subject to change as the term progresses, depending on the class's needs and interests. The deadlines will not change. '+' indicates date of a potential Friday extra lecture

3 Schedule

Starting March 23, 2020 the lectures are live streamed online at <https://www.twitch.tv/profarie>. The lectures will remain accessible on Twitch for some time and will also be archived at YouTube at <https://www.youtube.com/playlist?list=PLUVKdrSz00M4xoUGrYMEahqPy4QQu2kqz>.

Week	Monday 10am	Wednesday 8:30am	Friday 10am
1 Jan 6	Course Intro	TA: Git Labo due Thu 10PM	Lab1 intro (regex + r. descent)
2 Jan 13	Program Understanding	TA: Program Understanding Lab1 due Thu 10PM	Lab2 intro (transformations) 12:30PM open +
3 Jan 20	RegEx to DFA 1	Lab2 due Thu 10PM	Lab3 intro (OOP, equivalence)
4 Jan 27	RegEx to DFA 2	TA: RegEx to DFA Lab3 due Thu 10PM	Lab4 intro (term rewriting) 12:30PM open +
5 Feb 3	LL(1) Proofs 1	<i>Extra time for Lab4</i>	LL(1) Proofs 2
6 Feb 10	Lab5+6 intro (Parboiled)	TA: LL(1) Proofs Lab4 due Thu 10PM	Midterm Review 12:30PM Midterm Review +
— Feb 17	<i>Reading Week</i>		
7 Feb 24	<i>Midterm Week</i> https://ece.uwaterloo.ca/electrical-computer-engineering/current-undergraduate-students/academic-planning-and-support/mid-term-examinations/		
8 March 2	Grammar Design	Labs 5+6 due Thu 10PM	Lab7 intro (visitor, hashing, etc.) 12:30PM open +
9 March 9	TA: Midterm Recap	Lab7 due Thu 10PM	Lab8+11+12 intro
— March 16	<i>COVID19 Week</i>		
10 March 23	Optimization 1 Lab7 due Mon 10PM	Optimization 2 Lab8 due Thu 10PM	Optimization 3
11 March 30	Register Allocation Lab8 bonus due Mon 10PM	Lab11 due Thu 10PM	Garbage Collection 12:30PM open +
12 April 20		Lab12 (bonus) due Thu 10PM	

4 What's New This term?

Every term we do some new things to enhance the course. Many of these are based on student feedback, and are attempts to improve the student experience.

- Everything is new this term!
- New instructor.
- First time teaching this course.

Does every attempt succeed? No. Does every compiler 'optimization' make every program faster/smaller/better? No. Welcome to engineering: the real world is complicated.

- d. Speak up if something does not work or seems missing!
- e. Source L^AT_EX for Course Notes, Lab Manual, Study Questions, *etc.*: Submit a patch for bonus marks!
<https://git.uwaterloo.ca/ece351-notes/ece351-tex>

5 Marking Scheme

There are two marking schemes due to the changes in the teaching and exam schedules. The final mark is a maximum of the two marking schemes. The goal is to be fair to students who expected the final exam to have a significant weight and to the students who find the current circumstances too difficult to do their best on an exam.

5.1 Marking Scheme A

Labo	1%	
Labs (at deadline)	20%	
Labs (at end of term)	15%	
<i>Lab Sub-Total</i>		36%
Quizzes	4%	
Midterm	10%	
Final Exam	50%	
<i>Test Sub-Total</i>		64%
<hr/> <i>Total</i>		<hr/> 100%
Lab 8 bonus (a.k.a. Lab 9)	1%	
Lab 12	3%	
<i>Total (with bonus)</i>		104%

Labs are weighted equally.

Lab marks are based primarily on how your code performs on automated tests.

Both 36 and 64 are perfect squares and representative of a Pythagorean triple.

- You must pass the final exam in order to pass the course. If you do not pass the final exam, then your final mark will be the lower of your earned mark or 48%.
- You must pass the labs in order to pass the course. If you do not pass the labs, then your final mark will be the lower of your earned mark or 48%.
- We will mark the labs at the deadline and again at the end of the term (as indicated above in the marking scheme).
 - The labs are cumulative, so you need to keep up.
 - You may commit+push your work as many times as you want.
 - We mark the code on the *master* branch.
- There are bonus marks available for class participation: creating new tests, submitting patches, answering forum questions, creating new study questions, mentoring other students, scribing lectures *etc.* Any positive contribution outside of the regular marking scheme is open for consideration.

Note that lab code is cumulative: there are earlier labs that you will need to have solutions for in order to complete later labs. Dependencies are discussed in the Lab Manual. If you are going to skip a lab, be strategic in your choices. We will cut you some slack for one Git hiccup, subject to a personal interview with course staff.

We recommend that you commit+push frequently — at least at the end of every work session.

You can submit patches to the Course Notes, Lab Manual, and Study Questions in the following repo: `ist-git@git.uwaterloo.ca:ece351-notes/ece351-tex.git`

5.2 Marking Scheme B

Labo	1%	
Labs (at deadline)	20%	
Labs (at end of term)	30%	
<i>Lab Sub-Total</i>		51%
Quizzes	4%	
Midterm	20%	
Final Exam	25%	
<i>Test Sub-Total</i>		49%
<hr/> <i>Total</i>		100%
Lab 8 bonus (a.k.a. Lab 9)	1%	
Lab 12	3%	
<i>Total (with bonus)</i>		104%

- You must pass the labs in order to pass the course. If you do not pass the labs, then your final mark will be the lower of your earned mark or 48%.
- We will mark the labs at the deadline and again at the end of the term (as indicated above in the marking scheme).
 - The labs are cumulative, so you need to keep up.
 - You may commit+push your work as many times as you want.
 - We mark the code on the *master* branch.
- There are bonus marks available for class participation: creating new tests, submitting patches, answering forum questions, creating new study questions, mentoring other students, scribing lectures *etc.* Any positive contribution outside of the regular marking scheme is open for consideration.

Note that lab code is cumulative: there are earlier labs that you will need to have solutions for in order to complete later labs. Dependencies are discussed in the Lab Manual. If you are going to skip a lab, be strategic in your choices. We will cut you some slack for one Git hiccup, subject to a personal interview with course staff. We recommend that you commit+push frequently — at least at the end of every work session.

You can submit patches to the Course Notes, Lab Manual, and Study Questions in the following repo: `ist-git@git.uwaterloo.ca:ece351-notes/ece351-tex.git`

5.3 Course Marking Scheme

The marking scheme of the course is the maximum of the two marking schemes described above.

5.4 Online Final Exam

There will be a take-home final exam on Wednesday, Apr 15, 2020. You will have 48 hours to complete the exam, but the exam must be done in 4 hours. That is, you have 48 hours to start the exam, but once started, you must finish in 6 hours. The 6 hours includes extra time that is required by some students for accessibility reasons, the time that might be required to submit the exam on LEARN, plus any breaks that might be required during a long online exam.

The exam will be broken into an electronic form as a LEARN quiz (multiple choice, yes/no, and short answer questions that should be

answered on a computer), and a pen and paper questions that will be distributed electronically but have to be answered on “paper”. You will scan, photograph your answers and upload them to a Dropbox on LEARN, or complete this portion of the exam using some electronic means such as a tablet.

Randomized questions will be used ensure that answers cannot be shared between exams. The exam is open book and open notes. I will rely on your integrity to conduct yourselves in the best possible way in the given circumstances.

6 Reference Material

Available in <https://git.uwaterloo.ca/ece351-1201/pdfs>

- Course Outline (this document)
- Course Notes
- Lab Manual
- Study Questions

Textbooks available in the library:

- *Crafting a Compiler*¹
- *Modern Compiler Implementation in Java*²
- *Programming Language Pragmatics*³

The tools that we will use are documented largely online.⁴

Source L^AT_EX for these documents is available at <https://git.uwaterloo.ca/ece351-notes/ece351-tex>. Bonus marks for submitting patches.

¹ Charles N. Fischer, Ron K. Cytron, and Richard J. LeBlanc, Jr. *Crafting a Compiler*. Addison-Wesley, 2010

² Andrew W. Appel and Jens Palsberg. *Modern Compiler Implementation in Java*. Cambridge, 2004

³ Michael L. Scott. *Programming Language Pragmatics*. Morgan Kaufmann, 3 edition, 2009

⁴ parboiled.org

7 Collaboration

Interaction is an essential part of learning for most people. We are going to try a novel structure in ECE351 to facilitate learning collaboratively and honestly. We here define different levels of collaboration that correspond to different maximum grades. Your grade for a lab will be the lower of your earned grade or the cap. For each lab you will declare your level of collaboration. The levels are:

Individual: You discuss the labs with fellow students, perhaps making sketches on a blackboard/whiteboard. No written artifacts leave the discussion and get transferred to your code, i.e., your code is completely written by you.

Partner: You and a partner collaborate. You may look at each other's code. Each student is expected to do all of the typing on his or her computer. Partners are expected to have roughly equivalent skills.

Mentor: A mentor teaches a protégé. The mentor is expected to teach the protégé at the protégé's computer, and only the protégé operates the computer. The protégé does not look at the mentor's computer.

At every level of collaboration, every student is expected to physically key in every character that he or she commits. Students are not permitted to share electronic files with each other, or with students who have taken this course in the past, except as facilitated by the instructors using the course version control system.

7.1 Collaboration & Other Programming Languages

If you choose to implement the labs in another programming language, besides Java, it will be significantly more work, because you will have to build everything from scratch, without being able to use the skeleton code provided by the course staff.

As dispensation for this extra effort, you may collaborate at the *partner* level without the grade cap. You might also be eligible for bonus marks for tackling a more significant technical challenge.

Pre-lab material on the computing environment is excluded from these collaboration restrictions. By all means, please help each other out getting the computing environment working.

Skeleton code provided by course staff is excluded from these collaboration restrictions. You may look at the skeleton code on a computer collaboratively *before* you or your partner start writing your solutions and still declare *verbal* collaboration.

cap 100%

cap 85%

mentor cap 100%

protégé cap 80%

mentor is eligible for bonus points

Course staff are excluded from the collaboration caps, *i.e.*, you may ask questions of the staff, they may look at your code, *etc.*, without that imposing a cap on your mark.

We will share test cases this way, for example.

8 Classroom Conduct Policy

TIGERS! The human visual system has evolved to perceive saber-toothed tigers in the savannah. Fortunately, tigers are rare in Waterloo, Ontario. Unfortunately, your classmates are still human and hence their attention will be drawn to flashing lights (or Facebook, or movies, or video games) in their peripheral vision. We'd like to encourage everyone to be respectful of their classmates and to not distract them.

PRUDENCE. Wise use of computers and the Internet can be helpful for fully engaging in class. You might want to try out some syntax, or you might want to look up an API call, or you might want to verify your instructor's somewhat outrageous-sounding claim.

TEXT MODE. Devices in the first 4 rows of class should be operated in text-mode only: command prompt, text editor, IDE — in full screen mode. Paper is always good, of course.

Notifications should be turned off. Nothing moving on the screen.

Web browsing should be done with a text-mode browser.

TUNING OUT. There are some old-fashioned ways of tuning out that are less distracting for your classmates than the flashing lights on your new-fangled gizmo:

- Doodling.⁵
- Do homework. On paper. Maybe for another class.
- Read a textbook. Or a novel.
- Knit. Crochet. Those of you feeling in an especially masculine mood can explore needlepoint,⁶ macramé,⁷ or whittling.
- Pass notes. Write a joke to your friend. Pass it on.

EXCEPTIONS. You might have a good reason why you want to use GUI programs to support your learning and sit at the front of the class. Just let the instructor know.

ENFORCEMENT. This policy is part of our broader culture of trust and honesty. You are on your honour.

New for 2018! Prof Patrick Lam (ECE459) and I are going to try this out. Let's see how it works. The goal is to balance the benefits of having access to computing devices and the internet against the distractions of same.

tl;dr: paper or text-oriented programs only in first 4 rows of class.

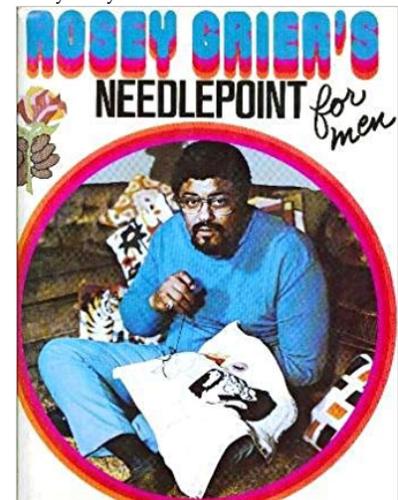
Using the terminal is a good technical skill that complements your learning in this course.

e.g., w3m, lynx, etc.

⁵ Drawing uses the right side of your brain. Language and mathematics are done on the left side of your brain. So doodling gives your brain something to do, but still leaves the left side of it available to tune back in to class if something interesting happens.

⁶ See Roosevelt Grier's book *Rosey Grier's Needlepoint For Men*. Grier was an all-star defensive tackle for the New York Giants and the Los Angeles Rams (both in the National Football League). He is 6'5" tall and weighed 284lbs.

⁷ Macramé is what sailors used to do to pass the time productively at sea and practice their knots. Make something fancy for your sweetheart.



9 University Policies

Reconciliation: We acknowledge that we live and work on the traditional territory of the Neutral, Anishinaabeg and Haudenosaunee peoples. The University of Waterloo is situated on the Haldimand Tract, the land promised to the Six Nations that includes ten kilometres on each side of the Grand River.



New for 2019.

<https://uwaterloo.ca/arts/about-arts/territorial-acknowledgement>

Figure 1: Contemporary map of the original Haldimand Tract and the remaining Six Nations Territory (red).

Intellectual Property: Students should be aware that this course contains the intellectual property of their instructor, TA, and/or the University of Waterloo. Intellectual property includes items such as: *source code*, course notes, lab notes, questions or solution sets, lecture content (and any audio/video recording thereof), *etc.*

New for 2017. <https://uwaterloo.ca/secretariat-general-counsel/faculty-staff-and-students-entering-relationships->
Don't post your code on GitHub!

Academic Integrity: In order to maintain a culture of academic integrity, members of the University of Waterloo community are expected to promote honesty, trust, fairness, respect and responsibility.

<http://uwaterloo.ca/academicintegrity/>

Code Clone Detection Software: ECE351 encourages a culture of trust and honesty, which the pedagogical literature has shown to be the most effective way to encourage learning and reduce cheating. While all labs in this course are to be done individually, the course collaboration policy provides a variety of options for legitimate engagement and learning with your peers.

The instructors might, at their discretion, use code clone detection software on your lab submissions, to cross-check compliance with the course collaboration policy. Please inform the instructors if you object to mechanical analysis of your code.

MOSS is the most common code clone detection tool for academic use, developed by compiler Prof Alex Aiken at Stanford: <https://theory.stanford.edu/~aiken/moss/>

AccessAbility: AccessAbility Services collaborates with all academic departments to arrange appropriate accommodations for students without compromising the academic integrity of the curriculum. If you require academic accommodations, please register with AccessAbility Services at the beginning of each academic term.

Needles Hall, Room 1401.
<https://uwaterloo.ca/accessability-services/>

Grievance: A student who believes that a decision affecting some aspect of his/her university life has been unfair or unreasonable may have grounds for initiating a grievance. When in doubt please be certain to contact the department's administrative assistant who will provide further assistance.

Policy 70, Student Petitions and Grievances, §4, <http://secretariat.uwaterloo.ca/Policies/policy70.htm>

Discipline: A student is expected to know what constitutes academic integrity to avoid committing an academic offence, and to take responsibility for his/her actions. A student who is unsure whether an action constitutes an offence, or who needs help in learning how to avoid offences (*e.g.*, plagiarism, cheating) or about rules for group work/collaboration should seek guidance from the course instructor, academic advisor, or the undergraduate Associate Dean.

<http://uwaterloo.ca/academicintegrity/>
For information on categories of offences and types of penalties, students should refer to Policy 71, Student Discipline, <http://secretariat.uwaterloo.ca/Policies/policy71.htm>
For typical penalties check Guidelines for the Assessment of Penalties, <http://secretariat.uwaterloo.ca/guidelines/penaltyguidelines.htm>

Appeals: A decision made or penalty imposed under Policy 70 (Student Petitions and Grievances) (other than a petition) or Policy 71 (Student Discipline) may be appealed if there is a ground. A student who believes he/she has a ground for an appeal should refer to Policy 72 (Student Appeals).

<http://secretariat.uwaterloo.ca/Policies/policy70.htm>
<http://secretariat.uwaterloo.ca/Policies/policy71.htm>
<http://secretariat.uwaterloo.ca/Policies/policy72.htm>

References

- [1] Andrew W. Appel and Jens Palsberg. *Modern Compiler Implementation in Java*. Cambridge, 2004.
- [2] Charles N. Fischer, Ron K. Cytron, and Richard J. LeBlanc, Jr. *Crafting a Compiler*. Addison-Wesley, 2010.
- [3] Michael L. Scott. *Programming Language Pragmatics*. Morgan Kaufmann, 3 edition, 2009.