



CIS 500

Software Foundations

Spring 2011

Lecture 1



Course Overview

What is “software foundations” ?

Software foundations (or “theory of programming languages”) is the mathematical study of the **meaning** of programs.

The goal is finding ways to describe program behaviors that are both **precise** and **abstract**.

- **precise** so that we can use mathematical tools to formalize and check interesting properties
- **abstract** so that properties of interest can be discussed clearly, without getting bogged down in low-level details

Why study software foundations?

- To prove specific properties of particular programs (i.e., program verification)
 - Important in some domains (safety-critical systems, hardware design, security protocols, inner loops of key algorithms, ...), but still fairly difficult and expensive
- To develop intuitions for *informal* reasoning about programs
- To prove general facts about all the programs in a given programming language (e.g., safety or isolation properties)
- To deeply understand language features (and their interactions) and develop principles for better language design (*PL is the “materials science” of computer science...*)

Questions This Course Could Help Answer

- “I’m designing a new web scripting language that’s going to take over the world. How can I specify it so that different people implementing compilers will know how it is supposed to behave?”
- I’m writing a compiler and I’d like to add this optimization, but I’m not convinced it’s always correct. How can I be sure?”
- “I want to write a program analysis tool that examines PERL programs and checks for possible command injection attacks. How can I know I haven’t missed any?”

Approaches to Program Meaning

- *Denotational semantics* and *domain theory* view programs as simple mathematical objects, abstracting away their flow of control and concentrating on their input-output behavior.
- *Program logics* such as *Hoare logic* and *dependent type theories* focus on logical rules for reasoning about programs.
- *Operational semantics* describes program behaviors by means of *abstract machines*. This approach is somewhat lower-level than the others, but is extremely flexible.
- *Process calculi* focus on the communication and synchronization behaviors of complex concurrent systems.
- *Type systems* describe *approximations* of program behaviors, concentrating on the shapes of the values passed between different parts of the program.

Overview

We will concentrate on operational techniques and type systems.

1. Part I: Foundations

- 1.1 Functional programming

- 1.2 Inductive definitions and proof techniques

- 1.3 Constructive logic

- 1.4 The Coq proof assistant

2. Basics

- 2.1 Reasoning about simple imperative programs

- 2.2 Operational semantics

3. Type systems

- 3.1 Simply typed lambda-calculus

- 3.2 Type safety

- 3.3 Subtyping

What you'll get out of the course

- A more sophisticated perspective on programs, programming languages, and the activity of programming
 - How to view programs and whole languages as formal, mathematical objects
 - How to make and prove rigorous claims about them
 - Detailed study of a range of basic language features
- Powerful mathematical tools for language (and software) design, description, and analysis
 - Constructive logic
 - Inductive methods of definition and proof
- Expertise using a cutting-edge mechanical proof assistant

Most software designers are language designers, at some point!

What this course is not

- An introduction to programming (see CIT 591)
- A course on functional programming (though we'll be doing some functional programming along the way)
- A course on compilers (you should ideally already have basic concepts such as lexical analysis, parsing, abstract syntax, and scope under your belt)
- A comparative survey of many different programming languages and styles (boring!)



Administrative Staff

Personnel

Instructor: Benjamin Pierce
Levine 562

Teaching Assistants: Anthony Cowley
Brent Yorgey

Administrative Assistant: Marissa Mele
Levine 311

Information

Textbook: Software Foundations
(see web page)

Webpage: <http://www.seas.upenn.edu/~cis500>

Mailing list: cis500@lists.seas.upenn.edu



Exams

- Two in-class midterms
 - Wednesday, Feb 16
 - Wednesday, Mar 30
- Final exam
 - Date not yet announced

Grading

Final course grades will be computed as follows:

- Homework: 20%
- 2 midterms: 20% each
- Final: 40%

(Lack of) extra Credit

1. Grade improvements can *only* be obtained by sitting in on the course next year and turning in all homeworks and exams. (If you are doing this to improve your grade from last year, please speak to me after class so I know who you are.)
2. There will be no extra credit projects, either during the semester or after the course ends. Concentrate your efforts on this course, now.

Collaboration

- Collaboration on homework is *strongly encouraged*
- Studying with other people is the best way to internalize the material
- Form study groups!
 - 2 people is the ideal size.
 - 3 is OK.
 - 4 is too many.

“You never really misunderstand something until you try to teach it...”

— Anon.

Homework

- Small part of your grade, but a large part of your understanding — impossible to perform well on exams without seriously grappling with the homework
- Submit one assignment per study group
- On written parts of homeworks, we will grade a semi-random subset of the problems on each assignment
- Late policy: Late homeworks lose 25% of their value for each day (or partial day) after the announced deadline

First Homework Assignment

- The first homework assignment is **due next Wednesday by 3PM**.
- You will need:
 - An account on a machine where Coq is installed (you can also install Coq on your own machine if you like)
 - Instructions on running Coq (available on the course web page)
 - The files `Preface.v` and `Basics.v` from the course web page



No Class Monday

Next Monday is Martin Luther King day.

Instead of coming to class, go out and do something for your community.



Office Hours

- Office hours will be held in the linux lab so that people can sit and work on their homework with a TA nearby, if they wish.

The WPE-I

- PhD students in CIS must pass a five-section Written Preliminary Exam (WPE-I)
Software Foundations is one of the five areas
- The final for this course is also the software foundations WPE-I exam
- Near the end of the semester, you will be given an opportunity to declare your intention to take the final exam for WPE credit

The WPE-I (continued)

- You do not need to be enrolled in the course to take the exam for WPE credit
- If you are enrolled in the course and also take the exam for WPE credit, you will receive two grades: a letter grade for the course final and a Pass/Fail for the WPE



The Coq Proof Assistant

What is a Proof?

- A proof is an indisputable argument for the truth of some mathematical assertion
- Proofs are ubiquitous in most branches of computer science
- However, unlike proofs in pure mathematics, many CS proofs are very long, shallow, and boring
- Can computers help?

What is a Proof Assistant?

Different ways of proving theorems with a computer:

- **Automatic theorem provers** find complete proofs on their own
 - Huge amount of work, beginning in the AI community in the 50s to 80s
 - In limited domains, extremely successful (e.g., hardware model checking)
 - In general, though, this approach is just too hard
- **Proof checkers** simply *verify* proofs that they are given
 - These proofs must be presented in an extremely detailed, low-level form
- **Proof assistants** are hybrid tools
 - “Hard steps” of proofs (the ones requiring deep insight) are provided by a human
 - “Easy parts” are filled in automatically

Some Proof Assistants

There are now a number of mature, sophisticated, and widely used proof assistants...

- Mizar (Poland)
- PVS (SRI)
- ACL2 (U. Texas)
- Isabelle (Cambridge / Munich)
- Twelf (CMU)
- Coq (INRIA)
- ...

The Coq Proof Assistant

- Developed at the INRIA research lab near Paris over the past 20 years
- Based on an extremely expressive logical foundation
 - The *Calculus of Inductive Constructions*
 - (hence the name!)
- Has been used to check a wide range of significant mathematical results
 - E.g., Gonthier's fully verified proof of the four-color theorem

Why Use Coq in This Course

- Rigor
 - using Coq forces us to be completely precise about the things we define and the claims we make about them
 - Coq's core notions of inductive definition and proof are a good match for the fundamental ways we define and reason about programming languages
- Interactivity
 - Instant feedback on homework
 - Easy to experiment with consequences of changing definitions, different reasoning techniques, etc.
- Useful background
 - Proof assistants are being used more and more widely in industry and academia
- Fun!
 - Coq is pretty addictive...