

Comp 311
Functional Programming
Introduction

Robert “Corky” Cartwright
Rice University

Robert “Corky” Cartwright

- PhD Stanford 1976[1977]
 - Official Advisor: David Luckham
 - Primary Mentor: John McCarthy
- Research Background
 - PL theory
 - First-order Programming Logic
 - Semantics of types, sequential functional languages
 - Type systems (Soft typing)
 - PL systems (software engineering)
 - Soft type checker for functional languages like Scheme
 - Testing concurrent programs
 - DrJava including Functional Java Subset

Course Overview I

- An Introduction to Functional Programming
- Lectures: TuTh, 9:25am – 10:40am, DH 1075
- Office hours: Corky
 - Duncan Hall 3104
 - Tuesdays and Thursdays 3 - 4 pm
 - By appointment

Course Mechanics

- Course website: Current and former course websites are located at the URL <https://comp311.rice.edu>
 - Syllabus and past lectures posted there
 - Lecture topics are subject to some change, particularly in last third of the class.
- Piazza: <https://piazza.com/rice/fall2024/comp311>
 - Course announcements and Q&A forum
 - Homework assignments and practice exams posted here
- Grading
 - 50% Homework Assignments
 - 25% Mid-term
 - 25% Final
 - Extra credit points on exams, some assignments
 - Late assignment not accepted except for seven 24-hour slip day extensions at your option. Save them for the harder assignments later in the course.

Course Overview II

- No textbook purchase required.
We will draw from a variety of sources including free online textbooks, monographs, and published papers. Some of them are available for purchase in printed form from online bookstores if you choose.
- Coursework consists primarily of weekly homework assignments that are either short programming assignments or written assignments about the underlying theory.
- Make sure you do these assignments! They embody the key ideas and principles covered in the course.

Course Culture

- Basic course on functional programming and “mostly functional” semantics-directed programming.
- With the possible exception of the material on Hoare Logic at the end of the course, the content should be accessible to freshmen with little background in Computer Science who know the rudiments of Java, *e.g.*, have taken an AP Computer Science course that covers Java programming.
- Coursework consists primarily of weekly homework assignments that are either short programming assignments or written assignments about the underlying theory.
- Make sure you do these assignments! They embody the key ideas and principles covered in the course.
- Why functional programming matters (expounded on next slide)

Two basic models of computation

- **Mutate state.** Example: simulating a machine language program for a commodity x86 processor. Too messy to illustrate.
- **Simplify symbolic expressions** (typically using leftmost reductions). Example: reducing an arithmetic expression to a value, *e.g.*,

$$(17 + 5) - 3 = 22 - 3 = 19$$

Read: [whyfp.pdf \(chalmers.se\)](http://www.cse.chalmers.se/~rjmh/Papers/whyfp.pdf)

[<https://www.cse.chalmers.se/~rjmh/Papers/whyfp.pdf>]

Mutating State

- In this model, a computation is a sequence of machine states that begins with a start state and generates a potentially infinite sequence of steps where each step is a new state completely determined by the preceding state and the instruction being performed.
 - The state includes both the program and the data being processed by the program.
 - The state space is infinite if the model is Turing-*complete* (capable of implementing any computable function).
 - The simplest such model is Turing machines, but modern so-called Random Access Machines (RAMs), such as x86 processors that are ubiquitous (billions of microprocessors!), embody the same basic model, except that they impose finite but very large limits on size and number of allowed states. Of course this limit increases if you add more memory (including auxiliary storage) to the machine.
 - Each step in a computation *updates* the state of a machine.
 - A primary advantage of this model is that it makes the costs of computations manifest by counting state-change steps (or equivalently the number of states appearing in a computation that terminates).
 - A disadvantage is that even trivial computation are messy and full of tedious details. Examples are very time-consuming to present so I won't.

Simplifying Expressions

In this model (sometimes called the *reduction* model), a computation is a potentially infinite sequence of “reduction steps” that transform a starting expression to an irreducible expression that is called the “answer”. Each reduction step simply replaces a sub-expression of an expression by an equivalent expression that is closer to being an answer.

- Simplification replaces a *program* by a “simpler” *program*.
- This computational model is actually *very familiar* to most students because we all learn how to do arithmetic in grammar school. Evaluating the arithmetic expression

$$(17 + 5) - 3 = 22 - 3 = 19$$

is a very simple computation in this model.

- In the simplest version of this model, every expression is a tree constructed from a countable collection of primitive operations of fixed arity (number of arguments).
- The linear expression


$$(17 + 5) - 3$$

is really a tree with the $-$ operation at the root. In linear symbolic notation, we typically encode the structure of the tree in text using parentheses as shown above.

Multitude of different syntactic definitions of expressions

- You are undoubtedly acquainted with some of these definitions as in Python, Pascal, Java, C, C++, Scala, Swift, Rust, Kotlin, Haskell, JavaScript, Perl, Snobol, APL, COBOL, Fortran, Modula, Cedar, ... , Racket/Scheme. If we ignore tedious differences in syntactic conventions and outright pathologies (*e.g.*, APL), there typically are only a few important differences:
 - Are variables allowed and can they be bound inside expressions? (Necessary for Turing completeness? No.)
 - Can compound expressions appear as operators? (Often stated: are functions “first-class” data values.) The answer regrettably is “No” in most mainstream languages (*e.g.* Pascal, Java, C, C++, JavaScript, ...). Even first-order logic bans such constructions. (Unnecessary for Turing completeness; explicit *apply* operator is a common hack.)
 - Many languages that do not explicitly support “functions as first class values” have tolerable workarounds like anonymous inner classes in Java (now abbreviated by “lambda expressions” [Church is turning in his grave]).

The intuitive key to mastering the “reduction” model of computation is to think of expressions as trees, often called *abstract syntax trees* (ASTs). In such a tree every operator is the root of a subtree where its operands (arguments), represented as ASTs, are its children. Hence, $(17 + 5)$ is represented by the AST



Homework Assignments

Think of the programming assignments in this class as very short essays. Focus as much on style as you would for an essay. In grading homework,

50% is based on clarity and style, and

50% is based on correctness.

Homework Assignments

- Projects are generally due one week after being assigned.
- Each student has 7 “slip days” to address scheduling conflicts and minor sickness. No more than 3 “slip days” can be used on a given assignment unless you get explicit permission from the instructor.
- Hoard your slip days. The assignments will be progressively more challenging. I predict that some students will not use any slip days.
- Expect to spend about 10 hours outside of class per week.
- Block this time off now in your schedule and respect these commitments.

Homework Assignments

- Assignments are published on Monday or Wednesday evening and due exactly one week (11:59 pm) later.
- Start on assignments early so that you have time to ask questions in class, on Piazza, and at office hours.
- A positive attitude and tackling assignments early will help you do your best in the course.

Homework Assignments

- All assignments will be small in scale.
- Most will be given in (the functional subset of) Racket which is a very simple, pure functional language that is easy to simulate in most modern type-safe languages. We will document Racket programs with types (sets of values), which are mandated in the program text of statically typed functional languages like Ocaml and Haskell.
- We will subsequently show how to cleanly express functional programs in Java consistent with good OO design, exposing most of the technology used to implement OO languages like Scala that explicitly support functional programming.
- We will show how to systematically reason about imperative code using Hoare logic; this reasoning typically relies on functional specifications. Programming is really mathematics.

Homework Assignments

- We mandate that you use the DrRacket programming environment to develop and test Racket programs. Racket is a dialect of the Scheme programming language. This standardization of Scheme dialect matters because we will test your code using Racket. The Racket platform runs on Windows, MacOSX, and Linux.
- For Java, you have the option of using DrJava or a professional IDE like IntelliJ or Eclipse. Choose what makes you comfortable. I only use DrJava so I won't be able to answer questions about the professional IDEs.
- We will use GitHub Classroom as the distribution mechanism and the repository for all assignments.
- Instructions on how to format and submit assignments will be posted as part of the assignment instructions.

*What is Functional
Programming?*

Background: Early Models of Computation

- Turing Machines (Turing)
- Type-0 Grammars (Chomsky)
- The Lambda Calculus (Church)
- Post Machines (Post)

The creators of these models were surprised when they all turned out to be equally expressive if computations are confined to functions mapping finite inputs to finite outputs. The notion of computability is an utterly fundamental notion in mathematics and predates all electronic computers

With exception of Lambda Calculus, all of these models are “bottom-up” frameworks for pushing bits or symbols. But even the Lambda Calculus had a grubby syntactic character because there was no underlying model based on defining and applying functions. The functional character of the Lambda Calculus was only an intuitive vision until Scott developed a truly functional model that could handle self-application and support an isomorphism between the domain D of values represented by lambda expressions and the domain $D \rightarrow D$ of (continuous) functions over D .

The Lambda Calculus

- A *calculus* consists of a set of rules for rewriting symbols.
- The *Lambda Calculus* was an attempt to rebuild all of mathematics on the notion of *functions* and *applications*.
- There is no mutation in the lambda calculus; it is a reduction system like rational arithmetic.
- Every program consists solely of applications of functions to arguments (which are also functions in the pure lambda calculus, a misleading restriction IMO)
- Applications of functions return values (which are also functions)
- Encoding numbers as functions does not work out well; in the pure lambda calculus, numbers are actually encoded as syntactic descriptions of functions. The equality of functions is undecidable.
- The Pure Lambda Calculus was a critical step in the right direction but it was NOT a true functional programming language. If you add a few constants (*natural numbers*, the *suc* and *if-zero* (conditional expressions) functions), you get PCF which is a true universal functional programming language. But even PCF is incomplete in fundamental (if practically unimportant) ways. Scheme/Racket is an analog of PCF that supports more interesting forms of ground (non-functional) data values than just the natural numbers.

What is Functional Programming?

Every program is a collection of function definitions plus an execution expression. For example, assuming our programming language includes the natural numbers, booleans, and a few simple primitive operations on natural numbers and boolean, we can define a program for computing factorial is as follows:

$$fact(0) = 1$$

$$fact(n + 1) = n \times fact(n - 1)$$

and compute $fact(1000)$ using equational reasoning which can be systematized as a *reduction system*. The result is obviously a large number but you can trivially do this calculation in Racket.

Why Avoid Side Effects?

- **Programs are easier to write:** There are fewer interactions between program components, enabling multiple programmers (or a single programmer on multiple days) to work together more easily. Moreover, essentially all data types have simple inductive definitions which provide a simple framework (recursive definitions of functions) for writing code.
- **Programs are easier to read:** Pieces of a program can be read and understood in isolation.
- **Programs are easier to test:** Less context needs to be built up before calling a function to test it.
- **Programs are easier to debug:** Problems can be isolated more easily, and behavior is inherently deterministic and **local**.
- **Programs are easier to reason about:** The model of computation needed to understand a program without mutation is much simpler; it is ordinary algebra plus induction on the structure of the data.

Why Avoid Side Effects?

- **Programs are easier to execute in parallel:** Because separate pieces of a computation do not interact, it is easy to compute them on separate processors
- This is an increasingly important consideration in the era of multicore chips, big data, and distributing computing.
- *This advantage undermines an often cited efficiency argument for using imperative programming (mutation of variables and data structures) instead of functional programming. Imperative hacks often introduce more synchronization and data-representation sharing.*

General Functional Programming Languages

We already defined the functional model of computation but modern functional languages share two important properties

- The ground (non-functional) domain of data values includes all common primitive types: numbers of various forms, characters, strings, lists, and optional algebraic forms of data that are constructed using program-declared constructors. Functions may appear in such constructions.
- Functions may be dynamically created during computation and returned as values of applications.

Why Emphasize Functions?

- Functions allow us to factor out common code
 - DRY: Don't Repeat Yourself
 - Why is DRY important?
 - Program understanding
 - Program maintenance
 - Passing functions as arguments is often the most straightforward way to abide by DRY
 - Returning functions as values is also important for DRY

Why Emphasize Functions?

- Functions allow us to concisely package computations and move them from one control point to another
- The functional model simplifies implementing and reasoning about parallel and distributed programming (yet again)
- Even reasoning about sequential programs is easier.
 - Equational reasoning + induction is sufficient for proving almost any property of a functional program

A Word on Object-Oriented Programming

- There is no tension between functional and object-oriented programming. In fact, OOP can be cast as an enrichment of FP. See <https://www.cs.rice.edu/~javaplt/papers/OOPEnrichesFP.pdf>
- In many ways, they complement one another.
- Languages like Scala, Swift, Kotlin and even Rust are designed to integrate both styles of programming

Quick Start with Racket

To install Racket on Windows, MacOSX, or Linux,


- Go to <https://racket-lang.org/download/> and download the “regular” version of Racket.
- Execute the downloaded installation file.
- Play with Racket arithmetic and simple functions on numbers. Racket performs rational arithmetic until forced to use inexact approximations.

Generalized Expressions in FP

- Most mainstream languages support a very restricted set of expressions built from operators (which are typically not values that can be passed as arguments to other operators) and constants because this restriction facilitates simple notation and simple implementation. Most functional programming languages support a much more general expression language where there is little or no distinction between functions and other ground (non-function) values nor between built-in operators like `+` and program-defined or library functions.
- In this realm of generalized expressions, we need to use a more explicit syntax (or adopt some slick syntactic conventions as in Ocaml and Haskell) to support the use of operators as function constants. Infix operators are particularly problematic, so the Lisp family of languages (which includes Racket) does not support infix notation. Instead of infix operators there are corresponding primitive functions. This convention is also common in mathematical logic, where $2+2$ is usually written $+(2,2)$. In such frameworks, there can be 0-ary operators where the application of 0-ary operator must be distinguished from the value itself (which is a function of 0 arguments). A putative example is an operator `!` designating a function that aborts a computation and returns a special failure value. Passing such an operator as a value is very different than applying it.
- Languages like Racket with Lisp-like syntax use parenthesized prefix notation where the function appears after the open parenthesis and space is the separator. Hence $2+2$ is written `(+ 2 2)`.

Program Syntax as Trees

In tools that process programs, programs are almost always represented as trees, a representation called *abstract syntax*. Parentheses do not appear in such representations because all “grouping” is explicit in the tree structure. For example, the expression $(17 + 5) - 3$ can be represented:



where the root is the “outermost” operator and the children are trees representing the operand expressions. In such a tree, all operators including constants are tree nodes and every operator node of arity $k > 0$ has k subtrees representing the operands. There is a one-to-one correspondence between fully parenthesized prefix notation (since the operator appears immediately before its operands) and tree notation for simple expressions.

In Racket, the expression above is written `(- (+ 17 5) 3)`