

Lecture Notes on The Lambda Calculus

15-814: Types and Programming Languages
Frank Pfenning

Lecture 1
Tuesday, September 4, 2018

1 Introduction

This course is about the principles of programming language design, many of which derive from the notion of *type*. Nevertheless, we will start by studying an exceedingly pure notion of computation based only on the notion of function, that is, Church's λ -calculus [CR36]. There are several reasons to do so.

- We will see a number of important concepts in their simplest possible form, which means we can discuss them in full detail. We will then reuse these notions frequently throughout the course without the same level of detail.
- The λ -calculus is of great historical and foundational significance. The independent and nearly simultaneous development of Turing Machines [Tur36] and the λ -Calculus [CR36] as universal computational mechanisms led to the *Church-Turing Thesis*, which states that the effectively computable (partial) functions are exactly those that can be implemented by Turing Machines or, equivalently, in the λ -calculus.
- The notion of function is the most basic abstraction present in nearly all programming languages. If we are to study programming languages, we therefore must strive to understand the notion of function.
- It's cool!

2 The λ -Calculus

In ordinary mathematical practice, functions are ubiquitous. For example, we might define

$$\begin{aligned} f(x) &= x + 5 \\ g(y) &= 2 * y + 7 \end{aligned}$$

Oddly, we never state what f or g actually are, we only state what happens when we apply them to arbitrary arguments such as x or y . The λ -calculus starts with the simple idea that we should have notation for the function itself, the so-called λ -abstraction.

$$\begin{aligned} f &= \lambda x. x + 5 \\ g &= \lambda y. 2 * y + 7 \end{aligned}$$

In general, $\lambda x. e$ for some arbitrary expression e stands for the function which, when applied to some e' becomes $[e'/x]e$, that is, the result of *substituting* or *plugging in* e' for occurrences of the variable x in e . For now, we will use this notion of substitution informally—in the next lecture we will define it formally.

We can already see that in a pure calculus of functions we will need at least three different kinds of expressions: λ -abstractions $\lambda x. e$ to form function, *application* $e_1 e_2$ to apply a function e_1 to an argument e_2 , and *variables* x, y, z , etc. We summarize this in the following form

$$\begin{array}{ll} \text{Variables} & x \\ \text{Expressions } e & ::= \lambda x. e \mid e_1 e_2 \mid x \end{array}$$

This is not the definition of the *concrete syntax* of a programming language, but a slightly more abstract form called *abstract syntax*. When we write down concrete expressions there are additional conventions and notations such as parentheses to avoid ambiguity.

1. Juxtaposition (which expresses application) is *left-associative* so that $x y z$ is read as $(x y) z$
2. $\lambda x.$ is a prefix whose scope extends as far as possible while remaining consistent with the parentheses that are present. For example, $\lambda x. (\lambda y. x y z) x$ is read as $\lambda x. ((\lambda y. (x y) z) x)$.

We say $\lambda x. e$ *binds* the variable x with scope e . Variables that occur in e but are not bound are called *free variables*, and we say that a variable x may occur free in an expression e . For example, y is free in $\lambda x. x y$ but

not x . Bound variables can be renamed consistently in a term So $\lambda x.x + 5 = \lambda y.y + 5 = \lambda \textit{whatever}. \textit{whatever} + 5$. Generally, we rename variables *silently* because we identify terms that differ only in the names of λ -bound variables. But, if we want to make the step explicit, we call it α -conversion.

$$\lambda x. e =_{\alpha} \lambda y. [y/x]e \quad \text{provided } y \text{ not free in } e$$

The proviso is necessary, for example, because $\lambda x.x y \neq \lambda y.y y$.

We capture the rule for function application with

$$(\lambda x. e_2) e_1 =_{\beta} [e_1/x]e_2$$

and call it β -conversion. Some care has to be taken for the substitution to be carried out correctly—we will return to this point later.

If we think beyond mere equality at *computation*, we see that β -conversion has a definitive direction: we apply it from left to right. We call this β -reduction and it is the engine of computation in the λ -calculus.

$$(\lambda x. e_2) e_1 \longrightarrow_{\beta} [e_1/x]e_2$$

3 Function Composition

One of the most fundamental operations on functions in mathematics is to compose them. We might write

$$(f \circ g)(x) = f(g(x))$$

Having λ -notation we can first explicitly denote the result of composition (with some redundant parentheses)

$$f \circ g = \lambda x. f(g(x))$$

As a second step, we realize that \circ itself is a function, taking two functions as arguments and returning another function. Ignoring the fact that it is usually written in infix notation, we define

$$\circ = \lambda f. \lambda g. \lambda x. f(g(x))$$

Now we can calculate, for example, the composition of the two functions we had at the beginning of the lecture. We note the steps where we apply

β -conversion.

$$\begin{aligned}
 & (\circ (\lambda x. x + 5)) (\lambda y. 2 * y + 7) \\
 = & ((\lambda f. \lambda g. \lambda x. f(g(x))) (\lambda x. x + 5)) (\lambda y. 2 * y + 7) \\
 =_{\beta} & (\lambda g. \lambda x. (\lambda x. x + 5)(g(x))) (\lambda y. 2 * y + 7) \\
 =_{\beta} & \lambda x. (\lambda x. x + 5) ((\lambda y. 2 * y + 7)(x)) \\
 =_{\beta} & \lambda x. (\lambda x. x + 5) (2 * x + 7) \\
 =_{\beta} & \lambda x. (2 * x + 7) + 5 \\
 = & \lambda x. 2 * x + 12
 \end{aligned}$$

While this appears to go beyond the pure λ -calculus, we will see in Section 7 that we can actually encode natural numbers, addition, and multiplication. We can see that \circ as an operator is not *commutative* because, in general, $\circ f g \neq \circ g f$. You may test your understanding by calculating $(\circ (\lambda y. 2 * y + 7)) (\lambda x. x + 5)$ and observing that it is different.

4 Identity

The simplest function is the identity function

$$I = \lambda x. x$$

We would expect that in general, $\circ I f = f = \circ f I$. Let's calculate one of these:

$$\begin{aligned}
 & \circ I f \\
 = & ((\lambda f. \lambda g. \lambda x. f(g(x))) (\lambda x. x)) f \\
 =_{\beta} & (\lambda g. \lambda x. (\lambda x. x)(g(x))) f \\
 =_{\beta} & \lambda x. ((\lambda x. x)(f(x))) \\
 =_{\beta} & \lambda x. f(x)
 \end{aligned}$$

We see $\circ I f = \lambda x. f x$ but it does not appear to be equal to f . However, $\lambda x. f x$ and f would seem to be equal in the following sense: if we apply both sides to an arbitrary expression e we get $(\lambda x. f x) e = f e$ on the left and $f e$ on the right. In other words, the two functions appear to be *extensionally* equal. We capture this by adding another rule to the calculus call η .

$$e =_{\eta} \lambda x. e x \quad \text{provided } x \text{ not free in } e$$

The proviso is necessary—can you find a counterexample to the equality if it is violated?

5 Summary of λ -Calculus

λ -Expressions.

Variables x
 Expressions $e ::= \lambda x. e \mid e_1 e_2 \mid x$

$\lambda x. e$ binds x with scope e , which is as large as possible while remaining consistent with the given parentheses. Juxtaposition $e_1 e_2$ is left-associative.

Equality.

Substitution	$[e_1/x]e_2$		<i>(capture-avoiding, see Lecture 2)</i>
α -conversion	$\lambda x. e$	$=_\alpha$	$\lambda y. [y/x]e$ provided y not free in e
β -conversion	$(\lambda x. e_2) e_1$	$=_\beta$	$[e_1/x]e_2$
η -conversion	$\lambda x. e x$	$=_\eta$	e provided x not free in e

We generally apply α -conversion silently, identifying terms that differ only in the names of the bound variables. When we write $e = e'$ we allow $\alpha\beta\eta$ -equality and the usual mathematical operations such as expanding a definition.

Reduction.

β -reduction $(\lambda x. e_2) e_1 \longrightarrow_\beta [e_1/x]e_2$

6 Representing Booleans

Before we can claim the λ -calculus as a universal language for computation, we need to be able to represent *data*. The simplest nontrivial data type are the Booleans, a type with two elements: *true* and *false*. The general technique is to represent the values of a given type by *normal forms*, that is, expressions that cannot be reduced. Furthermore, they should be *closed*, that is, not contain any free variables. We need to be able to distinguish between two values, and in a closed expression that suggest introducing two bound variables. We then define rather arbitrarily one to be *true* and the other to be *false*

$true = \lambda x. \lambda y. x$
 $false = \lambda x. \lambda y. y$

The next step will be to define *functions* on values of the type. Let's start with negation: we are trying to define a λ -expression *not* such that

$$\begin{aligned} \text{not true} &= \text{false} \\ \text{not false} &= \text{true} \end{aligned}$$

We start with the obvious:

$$\text{not} = \lambda b. \dots$$

Now there are two possibilities: we could either try to apply b to some arguments, or we could build some λ -abstractions. In lecture, we followed the first path—you may want try the second as an exercise.

$$\text{not} = \lambda b. b(\dots)(\dots)$$

We suggest two arguments to b , because b stands for a Boolean, and Booleans *true* and *false* both take two arguments. $\text{true} = \lambda x. \lambda y. x$ will pick out the first of these two arguments and discard the second, so since we specified $\text{not true} = \text{false}$, the first argument to b should be *false*!

$$\text{not} = \lambda b. b \text{ false } (\dots)$$

Since $\text{false} = \lambda x. \lambda y. y$ picks out the second argument and $\text{not false} = \text{true}$, the second argument to b should be *true*.

$$\text{not} = \lambda b. b \text{ false true}$$

Now it is a simple matter to calculate that the computation of *not* applied to *true* or *false* completes in three steps and obtain the correct result.

$$\begin{aligned} \text{not true} &\longrightarrow_{\beta}^3 \text{false} \\ \text{not false} &\longrightarrow_{\beta}^3 \text{true} \end{aligned}$$

We write $\longrightarrow_{\beta}^n$ for reduction in n steps, and $\longrightarrow_{\beta}^*$ for reduction in an arbitrary number of steps, including zero steps. In other words, $\longrightarrow_{\beta}^*$ is the reflexive and transitive closure of \longrightarrow_{β} .

As a next exercise we try conjunction. We want to define a λ -expression *and* such that

$$\begin{aligned} \text{and true true} &= \text{true} \\ \text{and true false} &= \text{false} \\ \text{and false true} &= \text{false} \\ \text{and false false} &= \text{false} \end{aligned}$$

Learning from the negation, we start by guessing

$$\text{and} = \lambda b. \lambda c. b (\dots) (\dots)$$

where we arbitrarily put b first. If b is *true*, this will return the first argument. Looking at the equations we see that this should always be equal to the second argument.

$$\text{and} = \lambda b. \lambda c. b c (\dots)$$

If b is *false* the result is always false, no matter what c is, so the second argument to b is just *false*.

$$\text{and} = \lambda b. \lambda c. b c \text{ false}$$

Again, it is now a simple matter to verify the desired equations and that, in fact, the right-hand side of these equations is obtained by reduction.

We know we can represent all functions on Booleans returning Booleans once we have negation and conjunction. But we can also represent the more general conditional *if* with the requirements

$$\begin{aligned} \text{if true } u \ w &= u \\ \text{if false } u \ w &= w \end{aligned}$$

Note here that the variable u and w stand for arbitrary λ -expressions and not just Booleans. From what we have seen before, the conditional is now easy to define:

$$\text{if} = \lambda b. \lambda u. \lambda w. b u w$$

Looking at the innermost abstraction, we have $\lambda w. (b u) w$ which is actually η -convertible to $b u$! Taking another step we arrive at

$$\begin{aligned} \text{if} &= \lambda b. \lambda u. \lambda w. b u w \\ &=_{\eta} \lambda b. \lambda u. b u \\ &=_{\eta} \lambda b. b \\ &= I \end{aligned}$$

In other words, the conditional is just the identity function!

7 Representing Natural Numbers

Finite types such as Booleans are not particularly interesting. When we think about the computational power of a calculus we generally consider

the *natural numbers* $0, 1, 2, \dots$. We would like a representation \bar{n} such that they are all distinct. We obtain this by thinking of the natural numbers are generated from zero by repeated application of the successor function. Since we want our representations to be closed we start with two abstractions: one (z) that stands for zero, and one (s) that stands for the successor function.

$$\begin{aligned}\bar{0} &= \lambda z. \lambda s. z \\ \bar{1} &= \lambda z. \lambda s. s z \\ \bar{2} &= \lambda z. \lambda s. s (s z) \\ \bar{3} &= \lambda z. \lambda s. s (s (s z)) \\ \dots & \\ \bar{n} &= \lambda z. \lambda s. \underbrace{s (\dots (s z))}_{n \text{ times}}\end{aligned}$$

In other words, the representation \bar{n} iterates its second argument n times over its first argument

$$\bar{n} x f = f^n(x)$$

where $f^n(x) = \underbrace{f(\dots(f(x)))}_{n \text{ times}}$

The first order of business now is to define a successor function that satisfies $\text{succ } \bar{n} = \overline{n+1}$. As usual, there is more than one way to define it, here is one (throwing in the definition of *zero* for uniformity):

$$\begin{aligned}\text{zero} &= \bar{0} &= \lambda z. \lambda s. z \\ \text{succ} &= \lambda n. \overline{n+1} &= \lambda n. \lambda z. \lambda s. s (n z s)\end{aligned}$$

We cannot carry out the correctness proof in closed form as we did for the Booleans since there would be infinitely many cases to consider. Instead we calculate generically (using mathematical notation and properties)

$$\begin{aligned}\text{succ } \bar{n} & \\ &= \lambda z. \lambda s. s (\bar{n} z s) \\ &= \lambda z. \lambda s. s (s^n(z)) \\ &= \lambda z. \lambda s. s^{n+1}(z) \\ &= \overline{n+1}\end{aligned}$$

A more formal argument might use mathematical induction over n .

Using the iteration property we can now define other mathematical functions over the natural numbers. For example, addition of n and k iterates the successor function n times on k .

$$\text{plus} = \lambda n. \lambda k. n k \text{ succ}$$

You are invited to verify the correctness of this definition by calculation. Similarly:

$$\begin{aligned} \text{times} &= \lambda n. \lambda k. n \text{ (plus } k \text{) zero} \\ \text{exp} &= \lambda b. \lambda e. e \text{ (times } b \text{) (succ zero)} \end{aligned}$$

Everything appears to be going swimmingly until we hit the predecessor function defined by

$$\begin{aligned} \text{pred } \bar{0} &= \bar{0} \\ \text{pred } \overline{n+1} &= \bar{n} \end{aligned}$$

You may try for a while to see if you can define the predecessor function, but it is very difficult. The problem seems to be that it is easy to define functions f using the schema of *iteration*

$$\begin{aligned} f \ 0 &= c \\ f \ (n+1) &= g \ (f \ n) \end{aligned}$$

(namely: $f = \lambda n. n \ c \ g$), but not the so-called schema of *primitive recursion*

$$\begin{aligned} f \ 0 &= c \\ f \ (n+1) &= g \ n \ (f \ n) \end{aligned}$$

because it is difficult to get access to n .

More about this and other properties and examples of the λ -calculus in Lecture 2.

References

- [CR36] Alonzo Church and J.B. Rosser. Some properties of conversion. *Transactions of the American Mathematical Society*, 39(3):472–482, May 1936.
- [Tur36] Alan Turing. On computable numbers, with an application to the entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42:230–265, 1936. Published 1937.