

Type Systems

Winter Semester 2006

Week 4

November 8

November 15, 2006 - version 1.1

The Lambda Calculus

The lambda-calculus

- ▶ If our previous language of arithmetic expressions was the simplest nontrivial programming language, then the lambda-calculus is the simplest *interesting* programming language...
 - ▶ Turing complete
 - ▶ higher order (functions as data)
- ▶ Indeed, in the lambda-calculus, *all* computation happens by means of function abstraction and application.
- ▶ The *e. coli* of programming language research
- ▶ The foundation of many real-world programming language designs (including ML, Haskell, Scheme, Lisp, ...)

Intuitions

Suppose we want to describe a function that adds three to any number we pass it. We might write

```
plus3 x = succ (succ (succ x))
```

That is, “`plus3 x` is `succ (succ (succ x))`.”

Intuitions

Suppose we want to describe a function that adds three to any number we pass it. We might write

$$\text{plus3 } x \quad = \quad \text{succ } (\text{succ } (\text{succ } x))$$

That is, “`plus3 x` is `succ (succ (succ x))`.”

Q: What is `plus3` itself?

Intuitions

Suppose we want to describe a function that adds three to any number we pass it. We might write

$$\text{plus3 } x \quad = \quad \text{succ } (\text{succ } (\text{succ } x))$$

That is, “`plus3 x` is `succ (succ (succ x))`.”

Q: What is `plus3` itself?

A: `plus3` is the function that, given `x`, yields `succ (succ (succ x))`.

Intuitions

Suppose we want to describe a function that adds three to any number we pass it. We might write

$$\text{plus3 } x = \text{succ } (\text{succ } (\text{succ } x))$$

That is, “`plus3 x` is `succ (succ (succ x))`.”

Q: What is `plus3` itself?

A: `plus3` is the function that, given `x`, yields `succ (succ (succ x))`.

$$\text{plus3} = \lambda x. \text{succ } (\text{succ } (\text{succ } x))$$

This function exists independent of the name `plus3`.

`λx. t` is written “`fun x → t`” in OCaml and “`x ⇒ t`” in Scala.

So `plus3 (succ 0)` is just a convenient shorthand for “the function that, given `x`, yields `succ (succ (succ x))`, applied to `succ 0`.”

$$\begin{aligned} & \text{plus3 (succ 0)} \\ & \quad = \\ & (\lambda x. \text{succ (succ (succ x))}) (\text{succ 0}) \end{aligned}$$

Abstractions Returning Functions

Consider the following variant of `g`:

```
double = λf. λy. f (f y)
```

I.e., `double` is the function that, when applied to a function `f`, yields a *function* that, when applied to an argument `y`, yields `f (f y)`.

Example

```
double plus3 0
=  (λf. λy. f (f y))
   (λx. succ (succ (succ x)))
   0
i.e. (λy. (λx. succ (succ (succ x)))
       ((λx. succ (succ (succ x))) y))
      0
i.e. (λx. succ (succ (succ x)))
       ((λx. succ (succ (succ x))) 0)
i.e. (λx. succ (succ (succ x)))
       (succ (succ (succ 0)))
i.e. succ (succ (succ (succ (succ (succ 0)))))
```

The Pure Lambda-Calculus

As the preceding examples suggest, once we have λ -abstraction and application, we can throw away all the other language primitives and still have left a rich and powerful programming language.

In this language — the “pure lambda-calculus” — *everything* is a function.

- ▶ Variables always denote functions
- ▶ Functions always take other functions as parameters
- ▶ The result of a function is always a function

Formalities

Syntax

$t ::=$

x

$\lambda x. t$

$t t$

terms

variable

abstraction

application

Terminology:

- ▶ terms in the pure λ -calculus are often called *λ -terms*
- ▶ terms of the form $\lambda x. t$ are called *λ -abstractions* or just *abstractions*

Syntactic conventions

Since λ -calculus provides only one-argument functions, all multi-argument functions must be written in curried style.

The following conventions make the linear forms of terms easier to read and write:

- ▶ Application associates to the left

E.g., $t u v$ means $(t u) v$, not $t (u v)$

- ▶ Bodies of λ - abstractions extend as far to the right as possible

E.g., $\lambda x. \lambda y. x y$ means $\lambda x. (\lambda y. x y)$, not $\lambda x. (\lambda y. x) y$

Scope

The λ -abstraction term $\lambda x.t$ binds the variable x .

The *scope* of this binding is the *body* t .

Occurrences of x inside t are said to be *bound* by the abstraction.

Occurrences of x that are *not* within the scope of an abstraction binding x are said to be *free*.

Test:

$$\lambda x. \lambda y. x y z$$

Scope

The λ -abstraction term $\lambda x.t$ *binds* the variable x .

The *scope* of this binding is the *body* t .

Occurrences of x inside t are said to be *bound* by the abstraction.

Occurrences of x that are *not* within the scope of an abstraction binding x are said to be *free*.

Test:

$$\begin{array}{l} \lambda x. \lambda y. x y z \\ \lambda x. (\lambda y. z y) y \end{array}$$

Values

$v ::=$
 $\lambda x. t$

values
abstraction value

Operational Semantics

Computation rule:

$$(\lambda x. t_{12}) v_2 \longrightarrow [x \mapsto v_2]t_{12} \quad (\text{E-APPABS})$$

Notation: $[x \mapsto v_2]t_{12}$ is “the term that results from substituting free occurrences of x in t_{12} with v_{12} .”

Operational Semantics

Computation rule:

$$(\lambda x. t_{12}) v_2 \longrightarrow [x \mapsto v_2]t_{12} \quad (\text{E-APPABS})$$

Notation: $[x \mapsto v_2]t_{12}$ is “the term that results from substituting free occurrences of x in t_{12} with v_{12} .”

Congruence rules:

$$\frac{t_1 \longrightarrow t'_1}{t_1 t_2 \longrightarrow t'_1 t_2} \quad (\text{E-APP1})$$

$$\frac{t_2 \longrightarrow t'_2}{v_1 t_2 \longrightarrow v_1 t'_2} \quad (\text{E-APP2})$$

Terminology

A term of the form $(\lambda x. t) v$ — that is, a λ -abstraction applied to a *value* — is called a *redex* (short for “reducible expression”).

Alternative evaluation strategies

Strictly speaking, the language we have defined is called the *pure, call-by-value lambda-calculus*.

The evaluation strategy we have chosen — *call by value* — reflects standard conventions found in most mainstream languages.

Some other common ones:

- ▶ Call by name (cf. Haskell)
- ▶ Normal order (leftmost/outermost)
- ▶ Full (non-deterministic) beta-reduction

Classical Lambda Calculus

Full beta reduction

The classical lambda calculus allows full beta reduction.

- ▶ The argument of a β -reduction to be an arbitrary term, not just a value.
- ▶ Reduction may appear anywhere in a term.

Full beta reduction

The classical lambda calculus allows full beta reduction.

- ▶ The argument of a β -reduction to be an arbitrary term, not just a value.
- ▶ Reduction may appear anywhere in a term.

Computation rule:

$$(\lambda x. t_{12}) t_2 \longrightarrow [x \mapsto t_2]t_{12} \quad (\text{E-APPABS})$$

Full beta reduction

The classical lambda calculus allows full beta reduction.

- ▶ The argument of a β -reduction to be an arbitrary term, not just a value.
- ▶ Reduction may appear anywhere in a term.

Computation rule:

$$(\lambda x. t_{12}) t_2 \longrightarrow [x \mapsto t_2]t_{12} \quad (\text{E-APPABS})$$

Congruence rules:

$$\frac{t_1 \longrightarrow t'_1}{t_1 t_2 \longrightarrow t'_1 t_2} \quad (\text{E-APP1})$$

$$\frac{t_2 \longrightarrow t'_2}{t_1 t_2 \longrightarrow t_1 t'_2} \quad (\text{E-APP2})$$

$$\frac{t \longrightarrow t'}{\lambda x. t \longrightarrow \lambda x. t'} \quad (\text{E-ABS})$$

Substitution revisited

Remember: $[x \mapsto v_2]t_{12}$ is “the term that results from substituting free occurrences of x in t_{12} with v_{12} .”

This is trickier than it looks!

For example:

$$\begin{aligned} & (\lambda x. (\lambda y. x)) y \\ \longrightarrow & [x \mapsto y]\lambda y. x \\ = & ??? \end{aligned}$$

Substitution revisited

Remember: $[x \mapsto v_2]t_{12}$ is “the term that results from substituting free occurrences of x in t_{12} with v_{12} .”

This is trickier than it looks!

For example:

$$\begin{aligned} & (\lambda x. (\lambda y. x)) y \\ \longrightarrow & [x \mapsto y]\lambda y. x \\ = & ??? \end{aligned}$$

Solution:

need to rename bound variables before performing the substitution.

$$\begin{aligned} & (\lambda x. (\lambda y. x)) y \\ = & (\lambda x. (\lambda z. x)) y \\ \longrightarrow & [x \mapsto y]\lambda z. x \\ = & \lambda z. y \end{aligned}$$

Alpha conversion

Renaming bound variables is formalized as α -conversion.

Conversion rule:

$$\frac{y \notin \text{fv}(t)}{\lambda x. t =_{\alpha} \lambda y. [x \mapsto y]t} \quad (\alpha)$$

Equivalence rules:

$$\frac{t_1 =_{\alpha} t_2}{t_2 =_{\alpha} t_1} \quad (\alpha\text{-SYMM})$$

$$\frac{t_1 =_{\alpha} t_2 \quad t_2 =_{\alpha} t_3}{t_1 =_{\alpha} t_3} \quad (\alpha\text{-TRANS})$$

Congruence rules: the usual ones.

Confluence

Full β -reduction makes it possible to have different reduction paths.

Q: Can a term evaluate to more than one normal form?

Confluence

Full β -reduction makes it possible to have different reduction paths.

Q: Can a term evaluate to more than one normal form?

The answer is no; this is a consequence of the following

Theorem [Church-Rosser]

Let t , t_1 , t_2 be terms such that $t \longrightarrow^* t_1$ and $t \longrightarrow^* t_2$. Then there exists a term t_3 such that $t_1 \longrightarrow^* t_3$ and $t_2 \longrightarrow^* t_3$.

Programming in the Lambda-Calculus

Multiple arguments

Consider the function `double`, which returns a function as an argument.

$$\text{double} = \lambda f. \lambda y. f (f y)$$

This idiom — a λ -abstraction that does nothing but immediately yield another abstraction — is very common in the λ -calculus.

In general, $\lambda x. \lambda y. t$ is a function that, given a value v for x , yields a function that, given a value u for y , yields t with v in place of x and u in place of y .

That is, $\lambda x. \lambda y. t$ is a two-argument function.

(Recall the discussion of *currying* in OCaml.)

The “Church Booleans”

`tru` = $\lambda t. \lambda f. t$

`fls` = $\lambda t. \lambda f. f$

`tru` `v` `w`
= $\underline{(\lambda t. \lambda f. t)}$ `v` `w` by definition
→ $\underline{(\lambda f. v)}$ `w` reducing the underlined redex
→ `v` reducing the underlined redex

`fls` `v` `w`
= $\underline{(\lambda t. \lambda f. f)}$ `v` `w` by definition
→ $\underline{(\lambda f. f)}$ `w` reducing the underlined redex
→ `w` reducing the underlined redex

Functions on Booleans

```
not = λb. b fls tru
```

That is, `not` is a function that, given a boolean value `v`, returns `fls` if `v` is `tru` and `tru` if `v` is `fls`.

Functions on Booleans

```
and = λb. λc. b c fls
```

That is, `and` is a function that, given two boolean values `v` and `w`, returns `w` if `v` is `tru` and `fls` if `v` is `fls`

Thus `and v w` yields `tru` if both `v` and `w` are `tru` and `fls` if either `v` or `w` is `fls`.

Pairs

```
pair = λf.λs.λb. b f s
fst  = λp. p tru
snd  = λp. p fls
```

That is, `pair v w` is a function that, when applied to a boolean value `b`, applies `b` to `v` and `w`.

By the definition of booleans, this application yields `v` if `b` is `tru` and `w` if `b` is `fls`, so the first and second projection functions `fst` and `snd` can be implemented simply by supplying the appropriate boolean.

Example

$\text{fst } (\text{pair } v \ w)$
 $= \text{fst } ((\lambda f. \lambda s. \lambda b. b \ f \ s) \ v \ w)$ by definition
 $\longrightarrow \text{fst } ((\lambda s. \lambda b. b \ v \ s) \ w)$ reducing
 $\longrightarrow \text{fst } (\lambda b. b \ v \ w)$ reducing
 $= \underline{(\lambda p. p \ \text{tru})} \ (\lambda b. b \ v \ w)$ by definition
 $\longrightarrow \underline{(\lambda b. b \ v \ w) \ \text{tru}}$ reducing
 $\longrightarrow \text{tru } v \ w$ reducing
 $\longrightarrow^* v$ as before.

Church numerals

Idea: represent the number n by a function that “repeats some action n times.”

$$c_0 = \lambda s. \lambda z. z$$

$$c_1 = \lambda s. \lambda z. s z$$

$$c_2 = \lambda s. \lambda z. s (s z)$$

$$c_3 = \lambda s. \lambda z. s (s (s z))$$

That is, each number n is represented by a term c_n that takes two arguments, s and z (for “successor” and “zero”), and applies s , n times, to z .

Functions on Church Numerals

Successor:

Functions on Church Numerals

Successor:

$$\text{succ} = \lambda n. \lambda s. \lambda z. s (n s z)$$

Functions on Church Numerals

Successor:

$$\text{succ} = \lambda n. \lambda s. \lambda z. s (n s z)$$

Addition:

Functions on Church Numerals

Successor:

$$\text{succ} = \lambda n. \lambda s. \lambda z. s (n s z)$$

Addition:

$$\text{plus} = \lambda m. \lambda n. \lambda s. \lambda z. m s (n s z)$$

Functions on Church Numerals

Successor:

$$\text{succ} = \lambda n. \lambda s. \lambda z. s (n s z)$$

Addition:

$$\text{plus} = \lambda m. \lambda n. \lambda s. \lambda z. m s (n s z)$$

Multiplication:

Functions on Church Numerals

Successor:

$$\text{succ} = \lambda n. \lambda s. \lambda z. s (n s z)$$

Addition:

$$\text{plus} = \lambda m. \lambda n. \lambda s. \lambda z. m s (n s z)$$

Multiplication:

$$\text{times} = \lambda m. \lambda n. m (\text{plus } n) c_0$$

Functions on Church Numerals

Successor:

$$\text{succ} = \lambda n. \lambda s. \lambda z. s (n s z)$$

Addition:

$$\text{plus} = \lambda m. \lambda n. \lambda s. \lambda z. m s (n s z)$$

Multiplication:

$$\text{times} = \lambda m. \lambda n. m (\text{plus } n) c_0$$

Zero test:

Functions on Church Numerals

Successor:

$$\text{scc} = \lambda n. \lambda s. \lambda z. s (n s z)$$

Addition:

$$\text{plus} = \lambda m. \lambda n. \lambda s. \lambda z. m s (n s z)$$

Multiplication:

$$\text{times} = \lambda m. \lambda n. m (\text{plus } n) c_0$$

Zero test:

$$\text{iszro} = \lambda m. m (\lambda x. \text{fls}) \text{tru}$$

Functions on Church Numerals

Successor:

$$\text{scc} = \lambda n. \lambda s. \lambda z. s (n s z)$$

Addition:

$$\text{plus} = \lambda m. \lambda n. \lambda s. \lambda z. m s (n s z)$$

Multiplication:

$$\text{times} = \lambda m. \lambda n. m (\text{plus } n) c_0$$

Zero test:

$$\text{iszro} = \lambda m. m (\lambda x. \text{fls}) \text{tru}$$

What about predecessor?

Predecessor

```
zz = pair c0 c0
```

```
ss = λp. pair (snd p) (scc (snd p))
```

```
prd = λm. fst (m ss zz)
```

Normal forms

Recall:

- ▶ A *normal form* is a term that cannot take an evaluation step.
- ▶ A *stuck* term is a normal form that is not a value.

Are there any stuck terms in the pure λ -calculus?

Normal forms

Recall:

- ▶ A *normal form* is a term that cannot take an evaluation step.
- ▶ A *stuck* term is a normal form that is not a value.

Are there any stuck terms in the pure λ -calculus?

Does every term evaluate to a normal form?

Divergence

$$\text{omega} = (\lambda x. x x) (\lambda x. x x)$$

Note that `omega` evaluates in one step to itself!

So evaluation of `omega` never reaches a normal form: it *diverges*.

Divergence

$$\text{omega} = (\lambda x. x x) (\lambda x. x x)$$

Note that `omega` evaluates in one step to itself!

So evaluation of `omega` never reaches a normal form: it *diverges*.

Being able to write a divergent computation does not seem very useful in itself. However, there are variants of `omega` that are very useful...

Recursion in the Lambda-Calculus

Iterated Application

Suppose f is some λ -abstraction, and consider the following term:

$$Y_f = (\lambda x. f (x x)) (\lambda x. f (x x))$$

Iterated Application

Suppose f is some λ -abstraction, and consider the following term:

$$Y_f = (\lambda x. f (x x)) (\lambda x. f (x x))$$

Now the “pattern of divergence” becomes more interesting:

$$\begin{aligned} & Y_f \\ & = \\ & \quad \underline{(\lambda x. f (x x)) (\lambda x. f (x x))} \\ & \quad \longrightarrow \\ & \quad f \left(\underline{(\lambda x. f (x x)) (\lambda x. f (x x))} \right) \\ & \quad \longrightarrow \\ & \quad f \left(f \left(\underline{(\lambda x. f (x x)) (\lambda x. f (x x))} \right) \right) \\ & \quad \longrightarrow \\ & \quad f \left(f \left(f \left(\underline{(\lambda x. f (x x)) (\lambda x. f (x x))} \right) \right) \right) \\ & \quad \longrightarrow \\ & \quad \dots \end{aligned}$$

Y_f is still not very useful, since (like ω), all it does is diverge.
Is there any way we could “slow it down”?

Delaying divergence

`poisonpill = λy. omega`

Note that `poisonpill` is a value — it will only diverge when we actually apply it to an argument. This means that we can safely pass it as an argument to other functions, return it as a result from functions, etc.

$$\begin{array}{c} \frac{(\lambda p. \text{fst } (\text{pair } p \text{ fls}) \text{ tru}) \text{ poisonpill}}{\longrightarrow} \\ \text{fst } (\text{pair } \text{poisonpill} \text{ fls}) \text{ tru} \\ \longrightarrow^* \\ \frac{\text{poisonpill } \text{tru}}{\longrightarrow} \\ \text{omega} \\ \longrightarrow \\ \dots \end{array}$$

A delayed variant of omega

Here is a variant of `omega` in which the delay and divergence are a bit more tightly intertwined:

$$\text{omegav} = \lambda y. (\lambda x. (\lambda y. x x y)) (\lambda x. (\lambda y. x x y)) y$$

Note that `omegav` is a normal form. However, if we apply it to any argument `v`, it diverges:

$$\begin{aligned} & \text{omegav } v \\ & = \\ & \frac{(\lambda y. (\lambda x. (\lambda y. x x y)) (\lambda x. (\lambda y. x x y)) y) v}{\longrightarrow} \\ & \frac{(\lambda x. (\lambda y. x x y)) (\lambda x. (\lambda y. x x y)) v}{\longrightarrow} \\ & (\lambda y. (\lambda x. (\lambda y. x x y)) (\lambda x. (\lambda y. x x y)) y) v \\ & = \\ & \text{omegav } v \end{aligned}$$

Another delayed variant

Suppose f is a function. Define

$$Z_f = \lambda y. (\lambda x. f (\lambda y. x x y)) (\lambda x. f (\lambda y. x x y)) y$$

This term combines the “added f ” from Y_f with the “delayed divergence” of ω_{av} .

If we now apply Z_f to an argument v , something interesting happens:

$$\begin{aligned} & Z_f v \\ &= \\ & \frac{(\lambda y. (\lambda x. f (\lambda y. x x y)) (\lambda x. f (\lambda y. x x y)) y) v}{\longrightarrow} \\ & \frac{(\lambda x. f (\lambda y. x x y)) (\lambda x. f (\lambda y. x x y)) v}{\longrightarrow} \\ & f (\lambda y. (\lambda x. f (\lambda y. x x y)) (\lambda x. f (\lambda y. x x y)) y) v \\ &= \\ & f Z_f v \end{aligned}$$

Since Z_f and v are both values, the next computation step will be the reduction of $f Z_f$ — that is, before we “diverge,” f gets to do some computation.

Now we are getting somewhere.

Recursion

Let

```
f = λfct.  
    λn.  
    if n=0 then 1  
    else n * (fct (pred n))
```

`f` looks just the ordinary factorial function, except that, in place of a recursive call in the last time, it calls the function `fct`, which is passed as a parameter.

N.b.: for brevity, this example uses “real” numbers and booleans, infix syntax, etc. It can easily be translated into the pure lambda-calculus (using Church numerals, etc.).

We can use Z to “tie the knot” in the definition of f and obtain a real recursive factorial function:

$$\begin{aligned}
 & Z_f \ 3 \\
 & \longrightarrow^* \\
 & f \ Z_f \ 3 \\
 & = \\
 & (\lambda fct. \ \lambda n. \ \dots) \ Z_f \ 3 \\
 & \longrightarrow \ \longrightarrow \\
 & \text{if } 3=0 \text{ then } 1 \text{ else } 3 * (Z_f \ (\text{pred } 3)) \\
 & \longrightarrow^* \\
 & 3 * (Z_f \ (\text{pred } 3)) \\
 & \longrightarrow \\
 & 3 * (Z_f \ 2) \\
 & \longrightarrow^* \\
 & 3 * (f \ Z_f \ 2) \\
 & \dots
 \end{aligned}$$

A Generic Z

If we define

$$Z = \lambda f. Z_f$$

i.e.,

$$Z = \lambda f. \lambda y. (\lambda x. f (\lambda y. x x y)) (\lambda x. f (\lambda y. x x y)) y$$

then we can obtain the behavior of Z_f for any f we like, simply by applying Z to f .

$$Z f \longrightarrow Z_f$$

For example:

```
fact      =      Z ( λfct.  
                    λn.  
                      if n=0 then 1  
                      else n * (fct (pred n)) )
```

Technical Note

The term Z here is essentially the same as the `fix` discussed the book.

$$Z = \lambda f. \lambda y. (\lambda x. f (\lambda y. x x y)) (\lambda x. f (\lambda y. x x y)) y$$

$$\text{fix} = \lambda f. (\lambda x. f (\lambda y. x x y)) (\lambda x. f (\lambda y. x x y))$$

Z is hopefully slightly easier to understand, since it has the property that $Z f v \longrightarrow^* f (Z f) v$, which `fix` does not (quite) share.