Lecture Notes on Recursion in the λ-Calculus

15-814: Types and Programming Languages
Jan Hoffmann and Frank Pfenning

Lecture 3
Tuesday, September 1, 2026

1 Introduction

In this lecture, we first conclude our study of computation in the λ-Calculus by discussing how we can express natural numbers, primitive recursion, and general recursion. In the next lecture, we formalize the definitions that we introduced informally using the concept of *rule induction*.

2 Representing Natural Numbers

When we think about the computational power of a programming language or machine model like Turing machines, we generally consider the functions on the natural numbers $0, 1, 2, \ldots$. The first step to defining such functions in the λ-Calculus is to find a representation of the natural numbers. We are looking for an expression $\overline{n}$ of the natural number n, such that $\overline{n}$ is distinct. We obtain this by thinking of the natural numbers as 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.

$$
\begin{array}{lcl}
\overline{0} & \triangleq & \lambda s.\, \lambda z.\, z \\
\overline{1} & \triangleq & \lambda s.\, \lambda z.\, s\, z \\
\overline{2} & \triangleq & \lambda s.\, \lambda z.\, s\, (s\, z) \\
\overline{3} & \triangleq & \lambda s.\, \lambda z.\, s\, (s\, (s\, z)) \\
\ldots & & \\
\overline{n} & \triangleq & \lambda s.\, \lambda z.\, s^n(z)
\end{array}
$$

Here we define for $n \in \mathbb{N}$ and λ-expressions f and e

$$\begin{array}{lcl} f^0(e) & \triangleq & e \\ f^{n+1}(e) & \triangleq & f(f^n(e)) \end{array}$$

Again, there are different ways of encoding natural numbers in the λ-calculus. The encoding $\overline{n}$ that we defined are called the *Church numerals*.

We can prove by induction on n that the representation $\overline{n}$ iterates its first argument n times over its second argument:

$$\overline{n}\, f\, x =_\beta f^n(x)$$

We now define some arithmetic functions. We start with the successor function *succ* that satisfies *succ* $\overline{n} =_\beta \overline{n+1}$. As usual, there is more than one way to define it, here is one (throwing in the definition of *zero* for uniformity):

$$\begin{array}{lcl} \mathit{zero} & \triangleq & \lambda s.\, \lambda z.\, z \\ \mathit{succ} & \triangleq & \lambda n.\, \lambda s.\, \lambda z.\, s\, (n\, s\, z) \end{array}$$

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)

$$\begin{array}{ll} & \mathit{succ}\, \overline{n} \\ =_\beta & \lambda s.\, \lambda z.\, s\, (\overline{n}\, z\, s) \\ =_\beta & \lambda s.\, \lambda z.\, s\, (s^n(z)) \\ =_\beta & \lambda s.\, \lambda z.\, s^{n+1}(z) \\ =_\beta & \overline{n+1} \end{array}$$

A more formal argument uses mathematical induction over n.

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

$$\mathit{plus} \triangleq \lambda n.\, \lambda k.\, n\ \mathit{succ}\ k$$

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

$$\mathit{times} \triangleq \lambda n.\, \lambda k.\, n\ (\mathit{plus}\ k)\ \mathit{zero}$$

3 Primitive Recursion

It is possible to define very fast-growing functions by iteration, such as the exponential function, or the "stack" function iterating the exponential.

$$\begin{array}{lcl} \mathit{exp} & \triangleq & \lambda b.\, \lambda e.\, e\ (\mathit{times}\ b)\ (\mathit{succ}\ \mathit{zero}) \\ \mathit{stack} & \triangleq & \lambda b.\, \lambda n.\, n\ (\mathit{exp}\ b)\ (\mathit{succ}\ \mathit{zero}) \end{array}$$

Everything appears to be going swimmingly until we think of a very simple function, namely the predecessor function. We want the following β-equalities to hold.

$$\begin{array}{lcl} \mathit{pred}\ \overline{0} & =_\beta & \overline{0} \\ \mathit{pred}\ \overline{n+1} & =_\beta & \overline{n} \end{array}$$

You may try for a while to see if you can define the predecessor function, but it is difficult. The problem is that we have to go from $\lambda s.\, \lambda z.\, s\,(\ldots(s\,z))$ to $\lambda s.\, \lambda z.\, s\,(\ldots z)$, that is, we have to *remove* an s rather than add an s as was required for the successor. One possible way out is to change representation and define $\overline{n}$ differently so that predecessor becomes easy (see Exercise 3). We run the risk that other functions then become more difficult to define, or that the representation is larger than the already inefficient unary representation already is. We follow a different path, keeping the representation the same and defining the function directly.

We can start by assessing why the schema of iteration does not immediately apply. The problem is that in

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

the function g only has access to the result of the recursive call of f on n, but not to the number n itself. What we would need is the *schema of primitive recursion*:

$$\begin{array}{lcl} f\ 0 & = & c \\ f\ (n+1) & = & h\ n\ (f\ n) \end{array}$$

where n is passed to h. For example, for the predecessor function we have $c = 0$ and $h = \lambda x.\, \lambda y.\, x$ (we do not need the result of the recursive call, just n which is the first argument to h).

3.1 Defining the Predecessor Function

Instead of trying to solve the general problem of how to implement primitive recursion, let's define the predecessor directly. Mathematically, we write

$n \dot{-} 1$ for the predecessor (that is, $0 \dot{-} 1 = 0$ and $n + 1 \dot{-} 1 = n$). The key idea is to gain access to n in the schema of primitive recursion by *rebuilding it* during the iteration. This requires *pairs*, a representation of which we will construct shortly.

Our specification then is

$$\mathit{pred}_2\, n = \langle n, n \dot{-} 1 \rangle$$

and the key step in its implementation in the λ-calculus is to express the definition by a schema of *iteration* rather than *primitive recursion*. The start is easy:

$$\mathit{pred}_2\, 0 = \langle 0, 0 \rangle$$

For $n + 1$ we need to use the value of $\mathit{pred}_2\, n$. For this purpose we assume we have a function *letpair* where

$$\mathit{letpair}\, \langle e_1, e_2 \rangle\, k = k\, e_1\, e_2$$

In other words, *letpair* passes the elements of the pair to a "continuation" k. Using *letpair* we start as

$$\mathit{pred}_2\, (n + 1) = \mathit{letpair}\, (\mathit{pred}_2\, n)\, (\lambda x.\, \lambda y.\, \ldots)$$

If pred_2 satisfies it specification then reduction will substitute n for x and $n \dot{-} 1$ for y. From these we need to construct the pair $\langle n + 1, n \rangle$ which we can do, for example, with $\langle x + 1, x \rangle$. This gives us

$$\begin{array}{lcl} \mathit{pred}_2\, 0 & = & \langle 0, 0 \rangle \\ \mathit{pred}_2\, (n + 1) & = & \mathit{letpair}\, (\mathit{pred}_2\, n)\, (\lambda x.\, \lambda y.\, \langle x + 1, x \rangle) \\ \\ \mathit{pred}\, n & = & \mathit{letpair}\, (\mathit{pred}_2\, n)\, (\lambda x.\, \lambda y.\, y) \end{array}$$

3.2 Defining Pairs

The next question is how to define pairs and *letpair*. The idea is to just abstract over the continuation itself! Then *letpair* isn't really needed because the functional representation of the pair itself will apply its argument to the two components of the pair, but if want to write it out it would be the identity.

$$\begin{array}{lcl} \langle e_1, e_2 \rangle & \triangleq & \lambda k.\, k\, e_1\, e_2 \\ \mathit{pair} & \triangleq & \lambda x.\, \lambda y.\, \lambda k.\, k\, x\, y \\ \mathit{letpair} & \triangleq & \lambda p.\, p \end{array}$$

Then $\mathit{letpair}\; p\; k =_\beta p\, k$ and $\mathit{pair}\, e_1\, e_2 =_\beta \langle e_1, e_2 \rangle$.

This is an example how introducing an abstraction (namely pairs) makes it easier to implement a challenging function.

3.3 Proving the Correctness of the Predecessor Function

Summarizing the above and expanding the definition of *letpair* we obtain the following definitions.

$$\begin{array}{lcl} \mathit{pred}_2 & \triangleq & \lambda n.\, n\,(\lambda p.\, p\,(\lambda x.\, \lambda y.\, \mathit{pair}\,(\mathit{succ}\,x)\,x))\,(\mathit{pair}\,\mathit{zero}\,\mathit{zero}) \\ \mathit{pred} & \triangleq & \lambda n.\, \mathit{pred}_2\,n\,(\lambda x.\, \lambda y.\, y) \end{array}$$

Let's do a rigorous proof of correctness of *pred*.[1] For the representation of natural numbers, we use the following equalities.

$$\begin{array}{lcl} \overline{0}\,g\,c & =_\beta & c \\ \overline{n+1}\,g\,c & =_\beta & g\,(\overline{n}\,g\,c) \end{array}$$

Lemma 1 $\mathit{pred}_2\,\overline{n} =_\beta \overline{\langle n, n \dot{-} 1\rangle}$

Proof: By mathematical induction on n.

Base: $n = 0$. Then

$$\begin{array}{lll} \mathit{pred}_2\,\overline{0} & =_\beta \overline{0}\,(\ldots)\,(\mathit{pair}\,\mathit{zero}\,\mathit{zero}) & \\ & =_\beta \mathit{pair}\,\mathit{zero}\,\mathit{zero} & \text{By repn. of } 0 \\ & =_\beta \langle 0,0\rangle = \overline{\langle 0, 0 \dot{-} 1\rangle} & \text{By repn. of 0 and pairs} \end{array}$$

Step: $n = m + 1$. Then

$$\begin{array}{ll} \mathit{pred}_2\,\overline{m+1} =_\beta \overline{m+1}\,(\lambda p.\, p\,(\lambda x.\, \lambda y.\, \mathit{pair}\,(\mathit{succ}\,x)\,x))\,(\mathit{pair}\,\mathit{zero}\,\mathit{zero}) & \\ =_\beta (\lambda p.\, p\,(\lambda x.\, \lambda y.\, \mathit{pair}\,(\mathit{succ}\,x)\,x))\,(\overline{m}\,(\lambda p.\, \ldots)\,(\ldots)) & \text{By repn. of } m+1 \\ =_\beta (\lambda p.\, p\,(\lambda x.\, \lambda y.\, \mathit{pair}\,(\mathit{succ}\,x)\,x))\,(\mathit{pred}_2\,\overline{m}) & \text{By defn. of } \mathit{pred}_2 \\ =_\beta (\lambda p.\, p\,(\lambda x.\, \lambda y.\, \mathit{pair}\,(\mathit{succ}\,x)\,x))\,\overline{\langle m, m \dot{-} 1\rangle} & \text{By ind. hyp. on } m \\ =_\beta \overline{\langle m, m \dot{-} 1\rangle}\,(\lambda x.\, \lambda y.\, \mathit{pair}\,(\mathit{succ}\,x)\,x) & \\ =_\beta \mathit{pair}\,(\mathit{succ}\,\overline{m})\,\overline{m} & \text{By repn. of pairs} \\ =_\beta \overline{\langle m+1, m\rangle} & \text{By repn. of successor and pairs} \\ = \overline{\langle m+1, (m+1) \dot{-} 1\rangle} & \text{By defn. of } \dot{-} \end{array}$$

□

Theorem 2 $\mathit{pred}\,\overline{n} =_\beta \overline{n \dot{-} 1}$

Proof: Direct, from Lemma 1.

[1] We did not carry out this proof in lecture relying on intuition and testing instead.

$$
\begin{aligned}
\mathit{pred}\,\overline{n} &= (\lambda n.\, \mathit{pred}_2\, n\, (\lambda x.\, \lambda y.\, y))\, \overline{n} \\
&=_\beta \mathit{pred}_2 \overline{n}\, (\lambda x.\, \lambda y.\, y) \\
&=_\beta \overline{\langle n, n \dot{-} 1\rangle}\, (\lambda x.\, \lambda y.\, y) && \text{By Lemma 1} \\
&=_\beta (\lambda k.\, k\, \overline{n}, \overline{n \dot{-} 1})\, (\lambda x.\, \lambda y.\, y) && \text{By repn. of pairs} \\
&=_\beta \overline{n \dot{-} 1}
\end{aligned}
$$

□

An interesting consequence of the Church-Rosser Theorem is that if $e =_\beta e'$ where e' is in normal form, then $e \longmapsto^*_\beta e'$.

3.4 Primitive Recursion

The general case of primitive recursion follows by a similar argument. Recall

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

We begin by defining a function f_2 specified with

$$f_2\, n = \langle n, f\, n\rangle$$

We can define f_2 using the schema of iteration.

$$
\begin{aligned}
f_2\, 0 &= \langle 0, c\rangle \\
f_2\, (n+1) &= \mathit{letpair}\, (f_2\, n)\, (\lambda x.\, \lambda y.\, \langle x+1, h\, x\, y\rangle) \\
f\, n &= \mathit{letpair}\, (f_2\, n)\, (\lambda x.\, \lambda y.\, x)
\end{aligned}
$$

To put this all together, we implement a function specified with

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

with the following definition for λ-expressions c and h:

$$
\begin{aligned}
\mathit{pair} &= \lambda x.\, \lambda y.\, \lambda g.\, g\, x\, y \\
f_2 &\triangleq \lambda n.\, n\, (\lambda r.\, r\ (\lambda x.\, \lambda y.\, \mathit{pair}\ (\mathit{succ}\ x)\ (h\ x\ y)))\ (\mathit{pair}\ \mathit{zero}\ c) \\
f &\triangleq \lambda n.\, f_2\, n\ (\lambda x.\, \lambda y.\, y)
\end{aligned}
$$

Recall that for the concrete case of the predecessor function we have $c = 0$ and $h = \lambda x.\, \lambda y.\, x$.

3.5 The Significance of Primitive Recursion

We have used primitive recursion here only as an aid to see how we can define functions in the pure λ-calculus. However, when computing over natural numbers we can restrict the functions that can be formed in schematic ways to obtain a language in which all functions terminate. Primitive recursion plays a central role in this because if c and g are terminating then so is f formed from them by primitive recursion. This can be proved by induction on n.

In this way we obtain a very rich set of functions. However, this set does not, for instance, include a function that simulates general Turing machines. We will see in the next few lectures that for each terminating programming language, there are total and computable functions that cannot be implemented.

4 General Recursion

Recall the schemas of iteration and primitive recursion:

$$\begin{array}{lcl} f\,0 & = & c \\ f\,(n+1) & = & g\,(f\,n) \end{array} \qquad\qquad \begin{array}{lcl} f\,0 & = & c \\ f\,(n+1) & = & g\,n\,(f\,n) \end{array}$$

We have already seen how functions defined by iteration and primitive recursion can be represented in the λ-calculus. We can also see that functions defined in this manner are terminating as long as c and g are.

Since there are computable functions that do not fit such a schema of recursion, there must be a more general way of defining recursive functions in the λ-Calculus.

The most general recursion schema we might think of is

$$f = h\,f$$

which means that in the right-hand side we can make arbitrary recursive calls to f. For the function *plus*, the function h could be defined as follows.

$$\begin{array}{l} h_{\mathsf{plus}} \triangleq \lambda \mathit{plus}.\,\lambda n.\,\lambda k.\ \mathit{ifz}\ n\ \mathit{then}\ k \\ \qquad\qquad\qquad\qquad\quad \mathit{else}\ \mathit{succ}\ (\mathit{plus}\ (\mathit{pred}\ n)\ k) \end{array}$$

Here, we assume that we have a combinator

$$\mathit{ifz}\ e\ \mathit{then}\ e_1\ \mathit{else}\ e_2$$

that branches based on $n = \overline{0}$ (see Exercise 4).

The function h_{plus} reduces to the addition function *plus* if we apply it to *plus*, that is,

$$\forall m, n \in \mathbb{N} : \ (h_{\mathsf{plus}}\mathsf{plus})\ \bar{m}\ \bar{n} =_\beta \mathsf{plus}\ \bar{m}\ \bar{n}\ .$$

However, as usual, *plus* can also be applied to λ-expressions that are not Church numerals and the behavior of *plus* and h_{plus} *plus* is actually different on some of these arguments, therefore h_{plus} *plus* $\neq_\beta$ *plus*. Can we find a different normal form that is a fixed point of h_{plus}? It's an interesting exercise but the answer seems to be no. As we see below, the λ-expression $Y\,(h_{\mathsf{plus}})$ is a fixed point of h_{plus} but it's not a normal form. In general, not every normal form has a fixed point that is a normal form.

The interesting question now is if we can in fact define an f explicitly when given h so that it satisfies $f = h\,f$. We say that f is a *fixed point* of h, because when we apply h to f we get f back. Since our solution should be in the λ-calculus, it would be $f =_\beta h\,f$. A function f satisfying such an equation may *not* be uniquely determined. For example, the equation $f = f$ (so, $h = \lambda x.x$) is satisfied by every function f. On the other hand, if h is a constant function such as $\lambda x.I$ then $f =_\beta (\lambda x.\,I)\,f =_\beta I$ has a simple unique solution. For the purpose of this lecture, any function that satisfies the given equation is acceptable.

If we believe in the Church-Turing thesis, then any computable function should be representable on Church numerals in the λ-calculus, so there is reason to hope there are explicit representations for such f. The answer is given by the so-called Y combinator.[2] Before we write it out, let's reflect on which laws Y should satisfy? We want that if $f = Y\,h$ and we specified that $f = h\,f$, so we get $Y\,h = h\,(Y\,h)$. We can iterate this reasoning indefinitely:

$$Y\,h = h\,(Y\,h) = h\,(h\,(Y\,h)) = h\,(h\,(h\,(Y\,h))) = \ldots$$

In other words, Y must iterate its argument arbitrarily many times.

The ingenious solution deposits one copy of h and then replicates $Y\,h$.

$$Y = \lambda h.\,(\lambda x.\,h\,(x\,x))\,(\lambda x.\,h\,(x\,x))$$

Here, the application $x\,x$ takes care of replicating $Y\,h$, and the outer application of h in $h\,(x\,x)$ leaves a copy of h behind. Formally, we calculate

$$\begin{array}{lll} Y\,h & =_\beta & (\lambda x.\,h\,(x\,x))\,(\lambda x.\,h\,(x\,x)) \\ & =_\beta & h\,((\lambda x.\,h\,(x\,x))\,(\lambda x.\,h\,(x\,x))) \\ & =_\beta & h\,(Y\,h) \end{array}$$

[2] For our purposes, a *combinator* is simply a λ-expression without any free variables.

In the first step, we just unwrap the definition of Y. In the second step we perform a β-reduction, substituting $[(\lambda x.\, h\,(x\,x))/x]\,h\,(x\,x)$. In the third step we recognize that this substitution recreated a copy of $Y\,h$.

You might wonder how a function defined with Y will ever terminate since

$$Y\,h =_\beta h\,(Y\,h) =_\beta h\,(h\,(Y\,h)) =_\beta h\,(h\,(h\,(Y\,h))) = \ldots$$

Well, it sometimes doesn't. Actually, this is important if we are to represent *partial recursive functions* which include functions that are undefined (have no normal form) on some arguments. Reconsider the specification $f = f$ as a recursion schema. Then $h = \lambda g.\, g$ and

$$Y\,h = Y\,(\lambda g.\, g) =_\beta (\lambda x.\, (\lambda g.\, g)\,(x\,x))\,(\lambda x.\, (\lambda g.\, g)\,(x\,x)) =_\beta (\lambda x.\, x\,x)\,(\lambda x.\, x\,x)$$

The expression on the right-hand side here (called Ω) only reduces to itself. It therefore does not have a normal form. In other words, the function $f = Y\,(\lambda g.\, g) = \Omega$ solves the equation $f = f$ by giving us a result which always diverges.

However, some recursive functions always terminate. Consider, for example, a case where f does not call itself recursively at all: $f = \lambda n.\, succ\ n$. Then $h_0 = \lambda g.\, \lambda n.\, succ\ n$. And we calculate further

$$\begin{array}{lcl} Y\,h_0 & = & Y\,(\lambda g.\, \lambda n.\, succ\ n) \\ & =_\beta & (\lambda x.\, (\lambda g.\, \lambda n.\, succ\ n)\,(x\,x))\,(\lambda x.\, (\lambda g.\, \lambda n.\, succ\ n)\,(x\,x)) \\ & =_\beta & (\lambda x.\, (\lambda n.\, succ\ n))\,(\lambda x.\, (\lambda n.\, succ\ n)) \\ & =_\beta & \lambda n.\, succ\ n \end{array}$$

So, fortunately, we obtain just the successor function *if we apply β-reduction from the outside in*. It is however also the case that there is an infinite reduction sequence starting at $Y\,h_0$. By the Church-Rosser Theorem this means that at any point during such an infinite reduction sequence we could still also reduce to $\lambda n.\, succ\ n$. A remarkable and nontrivial theorem about the λ-calculus is that if we always reduce the left-most/outer-most redex (which is the first expression of the form $(\lambda x.\, e_1)\,e_2$ we come to when reading an expression from left to right) then we will definitely arrive at a normal form when one exists. And by the Church-Rosser theorem such a normal form is unique (up to renaming of bound variables, as usual).

If a fixed point is not unique then the result of the Y combinator will return the most undefined result. Consider again the successor function but this time we will use the recursive result. Define

$$h_1 \triangleq \lambda n.\, succ\ n$$

Then we calculate

$$\begin{array}{rcl} Y\,h_1 & = & Y(\lambda n.\,succ\ n) \\ & =_\beta & (\lambda n.\,succ\ n)(Y\,h_1) \\ & =_\beta & succ\ (Y\,h_1) \\ & = & (\lambda n.\,\lambda s.\,\lambda z.\,s\,(n\,s\,z))\ (Y\,h_1) \\ & =_\beta & \lambda s.\,\lambda z.\,s\,((Y\,h_1)\,s\,z) \\ & =_\beta & \lambda s.\,\lambda z.\,s\,((\lambda s.\,\lambda z.\,s\,((Y\,h_1)\,s\,z))\,s\,z) \\ & =_\beta & \lambda s.\,\lambda z.\,s\,(s\,((Y\,h_1)\,s\,z)) \\ & =_\beta & \ldots \end{array}$$

Like Ω, this "infinite" number does not have a normal form but can appear in expressions that have a normal form. For example:

$$(Y\,h_1)(\lambda x.succ\ zero)\ zero =_\beta \overline{1}$$

5 Defining Functions by Recursion

Recall the definition

$$\begin{array}{l} h_{\mathsf{plus}} \triangleq \lambda f.\,\lambda n.\,\lambda k.\ \ ifz\ n\ then\ k \\ \qquad\qquad\qquad\qquad\quad else\ succ\ (f\ (pred\ n)\ k) \end{array}$$

We now define the addition function as

$$\mathsf{plus}' \triangleq Y\,h_{\mathsf{plus}}\ .$$

We calculate:

$$\begin{array}{rcl} \mathsf{plus}'\ \overline{1}\ \overline{0} & =_\beta & (h_{\mathsf{plus}}(Y\,h_{\mathsf{plus}}))\ \overline{1}\ \overline{0} \\ & =_\beta & ifz\ \overline{1}\ then\ \overline{0}\ else\ succ\ ((Y\,h_{\mathsf{plus}})\ (pred\ \overline{1})\ \overline{0}) \\ & =_\beta & succ\ ((Y\,h_{\mathsf{plus}})\ (pred\ \overline{1})\ \overline{0}) \\ & =_\beta & succ\ ((Y\,h_{\mathsf{plus}})\ \overline{0}\ \overline{0}) \\ & =_\beta & succ\ ((h_{\mathsf{plus}}(Y\,h_{\mathsf{plus}}))\ \overline{0}\ \overline{0}) \\ & =_\beta & succ\ ((h_{\mathsf{plus}}(Y\,h_{\mathsf{plus}}))\ \overline{0}\ \overline{0}) \\ & =_\beta & succ\ (ifz\ \overline{0}\ then\ \overline{0}\ else\ succ\ ((Y\,h_{\mathsf{plus}})\ (pred\ \overline{0})\ \overline{0})) \\ & =_\beta & succ\ \overline{0} \\ & =_\beta & \overline{1} \end{array}$$

Another example is the factorial function, which is informally given by the following recursive definition of fact

$$\mathsf{fact} = \lambda n.\,ifz\ n\ (succ\ zero)\ (times\ n\ (\mathsf{fact}\ (pred\,n)))$$

where we have already defined *succ*, *zero*, *times*, and *pred*. Of course, this is not directly allowed in the λ-calculus since the right-hand side mentions fact which we are just trying to define. The function h_{fact} which will be the argument to the Y combinator is then

$h_{\mathsf{fact}} = \lambda f.\, \lambda n.\, \mathit{ifz}\ n\ (\mathit{succ}\ \mathit{zero})\ (\mathit{times}\ n\ (f\ (\mathit{pred}\, n)))$

and

$\mathit{fact} = Y\ h_{\mathsf{fact}}$

We can write and execute this now in LAMBDA notation (see file rec.lam)

```
1  defn I = \x. x
2  defn K = \x. \y. x
3  defn Y = \h. (\x. h (x x)) (\x. h (x x))
4
5  defn ifz = \n. \c. \d. n (K d) c
6
7  defn h_fact = \f. \n. ifz n (succ zero) (times n (f (pred n)))
8  defn fact = Y h_fact
9
10 norm _120 = fact _5
11 norm _720 = fact (succ _5)
```

Listing 1: Recursive factorial in LAMBDA

Exercises

Exercise 1 One approach to representing functions defined by the schema of primitive recursion is to change the representation so that $\overline{n}$ is not an iterator but a *primitive recursor*.

$$\begin{array}{lcl} \overline{0} & = & \lambda s.\, \lambda z.\, z \\ \overline{n+1} & = & \lambda s.\, \lambda z.\, s\, \overline{n}\, (\overline{n}\, s\, z) \end{array}$$

1. Define the successor function *succ* (if possible) and show its correctness.
2. Define the predecessor function *pred* (if possible) and show its correctness.
3. Explore if it is possible to directly represent any function f specified by a schema of primitive recursion, ideally without constructing and destructing pairs.

Exercise 2 The unary representation of natural numbers requires tedious and error-prone counting to check whether your functions (such as factorial, Fibonacci, or greatest common divisor in the exercises below) behave correctly on some inputs with large answers. Fortunately, you can exploit that the LAMBDA implementation counts the number or reduction steps for you and prints it in decimal form!

(i) We have

$$\overline{n}\ \mathit{succ}\ \mathit{zero} \longmapsto_\beta^* \overline{n}$$

because $\overline{n}$ iterates the successor function n times on 0. Run some experiments in LAMBDA and conjecture how many leftmost-outermost reduction steps are required as a function of n. Note that only β-reductions are counted, and *not* replacing a definition (for example, *zero* by $\lambda s.\, \lambda z.\, z$). We justify this because we think of the definitions as taking place at the metalevel, in our mathematical domain of discourse.

(ii) Prove your conjecture from part (i), using induction on n. It may be helpful to use the mathematical notation $f^k c$ to describe a λ-expression generated by $f^0\, c = c$ and $f^{k+1}\, c = f\,(f^k\, c)$ where f and c are λ-expressions. For example, $\overline{n} = \lambda s.\, \lambda z.\, s^n\ z$ or $\mathit{succ}^3\ \mathit{zero} = \mathit{succ}\ (\mathit{succ}\ (\mathit{succ}\ \mathit{zero}))$.

Exercise 3 Once we can define each individual instance of the schemas of iteration and primitive recursion, we can also define them explicitly as combinators.

Define combinators *iter* and *primrec* such that

(i) The function *iter* $g\ c$ satisfies the schema of iteration

(ii) The function *primrec* $h\ c$ satisfies the schema of primitive recursion

You do not need to prove the correctness of your definitions.

Exercise 4 Define the following functions in the λ-calculus using the LAMBDA implementation. Here we take "=" to mean $=_\beta$, that is, β-conversion.

You may use all the functions in nat.lam as helper functions. Your functions should evidently reflect iteration, primitive recursion and pairs. In particular, you should avoid the use of the Y combinator which will be introduced in Lecture 3.

Provide at least 3 test cases for each function.

(i) ifz (definition by cases) satisfying the specification

$$\begin{array}{lcl} \mathsf{ifz}\ \overline{0}\ x\ y & = & x \\ \mathsf{ifz}\ \overline{k+1}\ x\ y & = & y \end{array}$$

(ii) even satisfying the specification

$$\begin{array}{lcl} \mathsf{even}\ \overline{2k} & = & \mathsf{true} \\ \mathsf{even}\ \overline{2k+1} & = & \mathsf{false} \end{array}$$

(iii) half satisfying the specification

$$\begin{array}{lcl} \mathsf{half}\ \overline{2k} & = & k \\ \mathsf{half}\ \overline{2k+1} & = & k \end{array}$$

Exercise 5 We can define binomial coefficients $\mathsf{bin}\ n\ k$ by the following recurrence:

$$\begin{array}{lcl} \mathsf{bin}\ 0\ k & = & 1 \\ \mathsf{bin}\ (n+1)\ 0 & = & 1 \\ \mathsf{bin}\ (n+1)\ (k+1) & = & \mathsf{bin}\ n\ k + \mathsf{bin}\ n\ (k+1) \end{array}$$

Give an implementation of the bin function in the λ-calculus via the LAMBDA implementation.

You may use the functions from nat.lam as helper functions, as well as those from Exercise 4. Your functions should evidently reflect iteration, primitive recursion and pairs. In particular, you should avoid the use of the Y combinator which will be introduced in Lecture 3.

Provide at least 5 test cases.

Exercise 6 Give an implementation of the factorial function in the λ-calculus as it arises from the schema of primitive recursion. How many β-reduction steps are required for factorial of 0, 1, 2, 3, 4, 5 in each of the two implementations?

Exercise 7 The Fibonacci function is defined by

$$\begin{array}{lcl} \mathsf{fib}\ 0 & = & 0 \\ \mathsf{fib}\ 1 & = & 1 \\ \mathsf{fib}\ (n+2) & = & \mathsf{fib}\ n + \mathsf{fib}\ (n+1) \end{array}$$

Give two implementations of the Fibonacci function in the λ-calculus (using the LAMBDA implementation). You may use the functions in (see file rec.lam).

(i) Exploit the idea behind the encoding of primitive recursion using pairs to give a direct implementation of fib without using the Y combinator.

(ii) Give an implementation of fib using the Y combinator.

Test your implementation on inputs 0, 1, 9, and 11, expecting results 0, 1, 34, and 89. Which of the two is more "efficient" (in the sense of number of β-reductions)?