

Notes on the λ -calculus

J.R.B. Cockett

Department of Computer Science, University of Calgary
Calgary, T2N 1N4, Alberta, Canada

March 5, 2023

1 Introduction

The λ -calculus was one of the first descriptions produced of computable functions. It grew out of an attempt by Alonso Church to provide a foundation for mathematics using functions as a basic building block. The attempt failed due to a form of Russell's paradox (which gave fixed point combinators). Church decided to extract a part of the calculus which worked perfectly well even if it did not achieve his original objective. This part is what we now know as the λ -calculus. In the course of developing this fragment he (and his students) realized that it was amazingly expressive and, in fact, could express all computable functions.

In this course we follow the realizations of Church and his students Rosser, Kleene, and Turing. We then continue to study the typed λ -calculus and its properties leading to Gödel's system T, Peano arithmetic, and the higher-order primitive recursive functions.

1.1 The λ -calculus

The λ -calculus is an algebra with a binary “application” operation and an “abstraction” operator. The application is sometimes written $f \bullet x$ (pronounced “ f applied to x ”) when we wish to emphasize the presence of a binary operation. However, it is more usually written by juxtaposition with an implicit association to the left (so that $xyz := (x \bullet y) \bullet z$). The abstraction operator binds a variable and is written using a lambda, $\lambda x.M$ here x is being bound: this gives the calculus its name.

More formally the terms in the calculus are formed using the rules in Table 1.


Two important notational conventions are used when writing λ -terms:

$\frac{x \text{ Variable}}{x \text{ Term}} \text{ Var}$
$\frac{M \text{ Term} \quad N \text{ Term}}{MN \text{ Term}} \text{ app} \quad \frac{M \text{ Term} \quad x \text{ Variable}}{\lambda x.M \text{ Term}} \text{ abst}$

Table 1: λ -term construction.

- Application associates to the left: $MNP := (MN)P$.
- Multiple abstractions are flattened: $\lambda xy.M := \lambda x(\lambda y.M)$.

In the term $\lambda x.M$ the variable x becomes bound and all free occurrences of x in the term M now refer to that binding. We may visualize this by drawing a λ -term as a tree with back arrows. For example the term $\lambda x.x(\lambda y.x(yx))$ may be portrayed as:



This suggests one could completely replace the bound variable names with a number which indicates how many λ -bindings one must pass through before one reaches the binding for the variable. These in programming language jargon are called “scope levels”. In this presentation the term $\lambda x.x(\lambda y.x(yx))$ would become $\lambda.0(\lambda.1(01))$. This way of indexing the variables is sometimes known as a “De Bruijn indexing” after the Dutch mathematician, Nicolaas de Bruijn who introduced a method of representing bound variables by numbers.

The free variables of a λ -term are those which occur in the term but are not bound. The free variables of a term are the smallest variable context (i.e. set of variables) in which we build can build the term. Below are the inference rules for building terms in a variable context. A term t can be built from a variable context Γ , written $\Gamma \vdash t$ Term, if and only if it can be built using the rules of the system in Table 2.

Two λ terms are **α -equivalent** if they differ only in the way the bound variables are named. It should be clear that terms are α -equivalent if and only if their de Bruijn indexed terms are the same.

1.2 Substitution

A key operation in the λ -calculus is substitution. By $M[N/x]$ shall be meant, intuitively, the term M with all free occurrences of x replaced by N . In practice this basic idea needs to be adjusted to avoid what is known as “variable capture” which occurs when a variable which is supposed to be free becomes bound as during substitution it is placed within the scope of a bound variable with the same name:

$$(\lambda x.xy)[x/y] \neq (\lambda x.xx)$$

$\frac{x \text{ Variable}}{\{x\} \vdash x \text{ Term}} \text{ Var}$	
$\frac{\Gamma \vdash M \text{ Term} \quad \Gamma \vdash N \text{ Term}}{\Gamma \vdash MN \text{ Term}} \text{ app}$	$\frac{\Gamma, x \vdash M \text{ Term} \quad x \text{ Variable}}{\Gamma \vdash \lambda x.M \text{ Term}} \text{ abst}$
$\frac{\Gamma \vdash M \text{ Term} \quad \Gamma \subseteq \Gamma'}{\Gamma' \vdash M \text{ Term}} \text{ Weaken}$	

Table 2: Variable contexts and λ -term construction.

Here the solution is to rename the bound variable (α -conversion) then do the replacement:

$$(\lambda x.xy)[x/y] =_{\alpha} (\lambda z.zy)[x/y] = \lambda z.zx$$

We may express the process of substitution as follows:

$$\begin{aligned}
x[N/y] &= \begin{cases} N & x == y \\ x & \text{not } x == y \end{cases} \\
(MN)[P/x] &= M[P/x]N[P/x] \\
(\lambda y.M)[N/x] &= \begin{cases} \lambda y.M & x == y \\ \lambda y.(M[N/x]) & \text{not } y \in FV(N) \\ \lambda z.(M[z/y])[N/x] & z \text{ new} \end{cases}
\end{aligned}$$

In the very last step we must do an α -conversion to avoid variable capture.

1.3 β -equality

We have already met α -equality: this, however, is really just a “book-keeping” equality due to having to make choices for variable names. The first significant equality is known as β -equality. It asserts:

$$(\lambda x.M)N = M[N/x].$$

From now on we shall take α -equality to be understood. This means we can define the β -equality relation using Table 3.

There is one other equality which is commonly considered called η -equality: $\lambda x.Mx = M$ (here x cannot occur in M). This is linked to “extensionality” because it says that if a function “behaves” just as N it must be N . For this note that $(\lambda x.Mx)N = MN$ for every N . These notes do not discuss this identity any further.

2 β -reduction and the Church-Rosser theorem

Using the above deduction system to determine the β -equality of λ -terms is highly inefficient as it involves a blind search. Church with his student Rosser determined that it was possible to direct the β -rule so it became a “ β -reduction”. A one step β -reduction is defined in Table 4.

$\frac{\Gamma \vdash M \text{ Term}}{\Gamma \vdash M = M} \text{ Reflexive}$	
$\frac{\Gamma \vdash N = M}{\Gamma \vdash M = M} \text{ Symmetry}$	$\frac{\Gamma \vdash M = N \quad \Gamma \vdash N = P}{\Gamma \vdash M = P} \text{ Transitivity}$
$\frac{\Gamma \vdash (\lambda x.M)N \text{ Term}}{\Gamma \vdash (\lambda x.M)N = M[N/x]} \beta\text{-equality}$	$\frac{\forall x, x \vdash M = N}{\Gamma \vdash (\lambda x.M) = (\lambda x.N)} \text{ abst}$
$\frac{\Gamma \vdash M = M' \quad \Gamma \vdash N \text{ Term}}{\Gamma \vdash MN = M'N} \text{ Lapp}$	$\frac{\Gamma \vdash N = N' \quad \Gamma \vdash M \text{ Term}}{\Gamma \vdash MN = MN'} \text{ Rapp}$

Table 3: β -equality between λ -terms.

$\frac{\Gamma \vdash (\lambda x.M)N \text{ Term}}{\Gamma \vdash (\lambda x.M)N \rightarrow M[N/x]} \beta\text{-reduction}$	
$\frac{\Gamma, x \vdash M \rightarrow N}{\Gamma \vdash (\lambda x.M) \rightarrow (\lambda x.N)} \text{ abst}$	
$\frac{\Gamma \vdash M \rightarrow M'}{\Gamma \vdash MN \rightarrow M'N} \text{ Lapp}$	$\frac{\Gamma \vdash N \rightarrow N'}{\Gamma \vdash MN \rightarrow MN'} \text{ Rapp}$

Table 4: One step β -reduction.

$\frac{}{\Gamma \vdash _ = \text{Id}_\Gamma[_] \text{ context}} \text{ Empty}$	
$\frac{\Gamma, x \vdash C_{\Gamma'}[_] \text{ context}}{\Gamma \vdash \lambda x. C_{\Gamma'}[_] \text{ context}} \text{ abst}$	$\frac{\Gamma \vdash C_{\Gamma'}[_] \text{ context} \quad \Gamma'' \subseteq \Gamma'}{\Gamma \vdash C_{\Gamma''}[_] \text{ context}} \text{ weak}$
$\frac{\Gamma \vdash M \text{ term} \quad \Gamma \vdash C_{\Gamma'}[_] \text{ context}}{\Gamma \vdash M C_{\Gamma'}[_] \text{ context}} \text{ Lapp}$	$\frac{\Gamma \vdash C_{\Gamma'}[_] \text{ context} \quad \Gamma \vdash N \text{ term}}{\Gamma \vdash C_{\Gamma'}[_] N \text{ context}} \text{ Rapp}$

Table 5: Construction of contexts

We shall write $N \xrightarrow{*} M$ to indicate that to get from N to M requires 0 or more one step β reductions and $N \xrightarrow{+} M$ to indicate that to get from N to M requires 1 or more one step β reductions. Thus

$$N \xrightarrow{*} M := (N = M_0) \rightarrow M_1 \rightarrow \dots \rightarrow (M_n = M)$$

Another way of thinking of a one step β -reduction is to view it as a top-level β -reduction being performed at exactly one subterm. A *context* is a λ -term with a single hole where one can put a λ expression. This is written $C[_]$ when a term N is put at this position we write $C[N]$. A context, in fact, also makes available some variables to the term which is to fill the hole. We shall indicate the variables which a context makes available in its hole by subscripting with a variable set: $C_\Gamma[_]$. We may then write the rules for building a context, see Table 6.

The first rule creates the hole immediately at the root: thus there is no term surrounding the hole in this case and, of course, only the free variables are available at this position. In the case of the λ -abstraction rule notice that the outside context has strictly fewer variables as the hole has the bound variable x available to it.

To fill the hole of a context one needs a term which can be built using the variables which are available in the hole, one can then fill the hole. For a filled context we shall drop the variable subscription:

$$\frac{\Gamma \vdash C_{\Gamma'}[_] \text{ context} \quad \Gamma' \vdash M \text{ term}}{\Gamma \vdash C[M] \text{ term}}$$


Putting a term N in context is unlike a substitution as a free variable x of N can become bound by the context when the position of the root of the tree is in the scope of a binding for x . The notation of contexts allows a one step rewrite to be viewed as a top-level β -reduction in context:

$$C[(\lambda x.M)N] \xrightarrow[\beta]{} C[M[N/x]].$$

Church with his student Rosser proved the following theorem:

Theorem 2.1 (Church, Rosser) *In the λ -calculus two λ -terms N and M are equal (with respect*

to β -equality) if and only if there is a λ -term L and



2.1 λ -terms in β -normal form

Before proving this theorem we discuss some very simple corollaries of it which – despite their simplicity – are very important and establish the consistency of the λ -calculus. A concern that one might have in dealing with the λ -calculus is that it is not consistent in the sense that one can prove that all terms are equal. The problem with the equality relation is that it really does not tell one when one *cannot* prove that two terms are equal. The Church-Rosser theorem, however, solves this problem rather neatly ...

Say that a λ -term is in **normal form** whenever there is no β -reduction leaving it. Here are some examples of terms in normal form:

- (a) Here are two very basic terms **True** := $\lambda xy.x$ and **False** := $\lambda xy.y$;
- (b) Terms which contain no λ -abstraction are always in normal form:

$$x, xy, xx, yy, xyxx, \dots$$

- (c) λ -abstractions of normal form terms are normal form terms:

$$\lambda x.x, \lambda yx.x, \lambda x.xx, \dots$$

- (d) Terms which are an application of a term M , which is in normal form *and* which has no λ -abstractions, applied to any term in normal form is itself in normal form.

$$x(\lambda x.x), y(\lambda x.x(\lambda x.xx)), \dots$$

- (e) Here is an infinite family of closed terms in normal form:

$$\lambda x.x, \lambda x.xx, \lambda x.xxx, \lambda x.xxxx, \lambda x.xxxxx, \dots$$

- (f) When two terms in normal form are applied to each other the result need not be in normal form:

$$(\lambda x.x)(\lambda x.x), \Omega := (\lambda x.xx)(\lambda x.xx), (\lambda x.xxx)(\lambda x.xxx), \dots$$

None of these are in normal form and in fact the second and third have infinite β -reduction sequences.

- $(\lambda x.x)(\lambda x.x) \rightarrow (\lambda x.x)$ is a reduction to normal form.
- $\Omega := (\lambda x.xx)(\lambda x.xx) \rightarrow xx[(\lambda x.xx)/x] = (\lambda x.xx)(\lambda x.xx)$ is a reduction of Ω to itself and thus Ω cannot have a normal form.

- The last term also has an infinite reduction sequence in which the term actually grows in size:

$$\begin{aligned}
& (\lambda x.xxx)(\lambda x.xxx) \rightarrow xxx[(\lambda x.xxx)/x] \\
& = (\lambda x.xxx)(\lambda x.xxx)(\lambda x.xxx) \rightarrow (xxx[(\lambda x.xxx)/x])(\lambda x.xxx) \\
& = (\lambda x.xxx)(\lambda x.xxx)(\lambda x.xxx)(\lambda x.xxx) \rightarrow \dots
\end{aligned}$$

An important corollary of the Church-Rosser theorem is:

Corollary 2.2 *Two terms which are in normal form which are not α -equivalent are distinct in the λ -calculus.*

It is clear that **True** and **False** are closed terms which are not α -equivalent and so:

Corollary 2.3 *The λ -calculus is consistent in the sense that there are at least two unequal closed terms. In fact, there are infinitely many distinct terms.*

Another important result is the “normal form theorem” which says:

Theorem 2.4 *A λ -term is β -equivalent to at most one normal form term.*

PROOF: To prove this requires two observations: first if M is β -equivalent to a normal form term N then there is a reduction sequence $M \xrightarrow{*} N$. We may see this using the Church-Rosser theorem as follows. As M is β -equivalent to N there is an L to which they both reduce. But N cannot be reduced whence $N = L$ so that $M \xrightarrow{*} N$.

Next suppose that M reduces to two normal forms N_1 and N_2 then immediately N_1 and N_2 are β -equivalent so that there is an L to which both reduce. However as both are in normal form $N_1 = L = N_2$ establishing the result. \square

2.2 The proof of the Church-Rosser theorem

At the time it was proven the Church-Rosser theorem was regarded as a fairly major achievement. In fact, it led to a whole field of study called *term rewriting* (which we discuss later). Furthermore, the original proof was quite long. Here we shall present a proof due to Barendregt [1] which uses a marking argument. The proof starts with some simple observations which reduce the problem:

OBSERVATION I:

Two λ -terms are equal if and only if there is a zig-zag of β -reductions:


$$M = N_0 \xleftarrow{*} N_1 \xrightarrow{*} N_2 \xleftarrow{*} \dots \xrightarrow{*} N_p = N$$

(This is immediate!)


OBSERVATION II:

If we can prove:


Proposition 2.5 (Confluence) *Given any divergence of λ -terms and β reductions:*



there is a convergence




Then the Church Rosser theorem holds!
This is because we can fill in any chain as above:




OBSERVATION III:

If we can prove:


Lemma 2.6 (Strip lemma) *Given any divergence of λ -terms and β reductions:*



(note the one-step reduction) there is a convergence



Then we can prove confluence! This is because we can break down confluence into a series of one steps down the left hand reduction:



Unfortunately even the strip lemma is not so easy to prove. This is where we shall employ a marking argument originally due to Barendregt.

2.3 Barendregt's marking argument

We may define marked λ -terms by:

$$M \rightarrow x \mid \lambda x.M \mid MM \mid (\underline{\lambda}x.M)M$$

There is an obvious underlying map to the “unmarked” (ordinary) λ -terms by just dropping the marking:

$$U((\underline{\lambda}x.M)N) = (\lambda x.M)N$$

There is also a by-value evaluation of a marked λ -term:

$$\begin{aligned}
 V(x) &= x \\
 V(\lambda x.M) &= \lambda x.V(M) \\
 V(MN) &= V(M)V(n) \\
 V((\underline{\lambda}x.M)N) &= V(M)[V(N)/x]
 \end{aligned}$$

This can also be written as an inference system where the by-value reduction is presented as $N \rightsquigarrow V(N)$:

$\frac{x \text{ var}}{x \rightsquigarrow x}$	$\frac{M \rightsquigarrow M' \quad N \rightsquigarrow N'}{MN \rightsquigarrow M'N'}$
$\frac{N \rightsquigarrow N'}{\lambda x.N \rightsquigarrow \lambda x.N'}$	$\frac{N \rightsquigarrow N' \quad M \rightsquigarrow M'}{(\underline{\lambda}x.N)M \rightsquigarrow N'[M'/x]}$

Marked by-value reduction

It is easy to see that this evaluation always terminates and $V(M)$ contains no marked terms. Furthermore there is a rewriting sequence mimicking the evaluation;

$$\begin{array}{ccc} & M & \\ & \downarrow & \searrow \\ U(M) & \xrightarrow{*} & V(M) \end{array}$$

Consider again the strip lemma we may view the one step rewrite as a marked reduction:

$$\begin{array}{ccc} N & \xrightarrow{*} & M \\ \downarrow & & \\ M' & & \end{array} \quad \begin{array}{ccc} U(N') = N & \xrightarrow{*} & U(M) \\ \downarrow & & \\ V(N') = M' & & \end{array}$$

we can then prove the strip lemma by proving the following proposition for marked reductions:

Proposition 2.7 *In the marked λ -calculus (with the evident reductions) we have:*

$$\begin{array}{ccc} M & \longrightarrow & M' \\ \downarrow * & & \downarrow * \\ V(M) & \xrightarrow{*} & V(M') \end{array}$$

Given this we may break down the version of the strip lemma above as:

$$\begin{array}{ccccccc} M & \longrightarrow & M_1 & \longrightarrow & M_2 & \longrightarrow & \dots & \longrightarrow & M' \\ \downarrow * & & \downarrow * & & \downarrow * & & & & \downarrow * \\ V(M) & \xrightarrow{\dots} & V(M_1) & \xrightarrow{\dots} & V(M_2) & \xrightarrow{\dots} & \dots & \xrightarrow{\dots} & V(M') \end{array}$$

and thus obtain the desired result.

The proof of the proposition proceeds by considering a series of cases for the single step rewrite:

At the root(unmarked):

$$\begin{array}{ccc} (\lambda x.M)N & \longrightarrow & M[N/x] \\ \downarrow * & & \downarrow * \\ (\lambda x.V(M))V(N) & \longrightarrow & V(M)[V(N)/x] = V(M[N/x]) \end{array}$$

where we must use the substitution lemma 2.8 below for the last step.

At the root (marked):

$$\begin{array}{ccc}
(\underline{\lambda}x.M)N & \longrightarrow & M[N/x] \\
\downarrow * & & \downarrow * \\
(\underline{\lambda}x.V(M))V(N) & & \\
\downarrow & & \\
V(M)[V(N)/x] & \equiv & V(M)[V(N)/x]
\end{array}$$

In context (unmarked): The argument is more complex if the reduction is not at the root:

$$\begin{array}{ccc}
C[(\underline{\lambda}x.M)N] & \longrightarrow & C[M[N/x]] \\
\downarrow * & & \downarrow * \\
C[(\underline{\lambda}x.V(M))V(N)] & \longrightarrow & C[V(M)[V(N)/x]] \\
\downarrow * & & \downarrow * \\
V(C[(\underline{\lambda}x.V(M))V(N)]) & \xrightarrow{*} & V(C[V(M)[V(N)/x]])
\end{array}$$

Here we use the fact that the by value evaluation of the term in (marked) context will evaluate the inner term before it evaluates the context. This allows us to separate the evaluation into evaluating the inner term and then evaluating the result in context. We must the establish that the lower square can be completed: this done in lemma 2.9.

In context (marked):

$$\begin{array}{ccc}
C[(\underline{\lambda}x.M)N] & \longrightarrow & C[M[N/x]] \\
\downarrow * & & \downarrow * \\
C[V(M)[V(N)/x]] & \equiv & C[V(M)[V(N)/x]] \\
\downarrow * & & \downarrow * \\
V(C[V(M)[V(N)/x]]) & \equiv & V(C[V(M)[V(N)/x]])
\end{array}$$

To complete the proof we have two lemmas to prove:

Lemma 2.8

$$V(M[N/x]) = V(M)[V(N)/x].$$

PROOF: We shall prove this by structural induction on M :

M is a variable: If $M = x$ we have $V(x[N/x]) = V(N) = x[V(N)/x] = V(x)[V(N)/x]$. If $M = y$ and $y \neq x$ then we have $V(y[N/x]) = V(y) = y = y[V(N)/x] = V(y)[V(N)/y]$.

M is an application: If $M = N_1N_2$ then

$$\begin{aligned}
V((N_1N_2)[N/x]) &= V(N_1[N/x]N_2[N/x]) \text{ defn of substitution} \\
&= V(N_1[N/x])V(N_2[N/x]) \text{ defn of marked evaluation} \\
&= V(N_1)[V(N)/x]V(N_2)[V(N)/x] \\
&\quad \text{structural induction } (N_1 \text{ and } N_2 \text{ are smaller terms)} \\
&= (V(N_1)V(N_2))[V(N)/x] \text{ defn of substitution} \\
&= (V(N_1N_2))[V(N)/x] \text{ defn of marked evaluation}
\end{aligned}$$

M is an unmarked abstraction: Suppose $M = \lambda y.N'$ then (assuming without loss of generality $x \neq y$):

$$\begin{aligned}
V((\lambda y.N')[N/x]) &= V(\lambda y.(N'[N/x])) \text{ defn of substitution} \\
&= \lambda y.V(N'[N/x]) \text{ defn of marked evaluation} \\
&= \lambda y.V(N')[V(N)/x] \text{ structural induction } (N' \text{ is a smaller term)} \\
&= V(\lambda y.N')[V(N)/x] \text{ defn of marked evaluation}
\end{aligned}$$

M is a marked λ -abstraction: Suppose $M = (\lambda y.M')N'$ then

$$\begin{aligned}
V(((\lambda y.M')N')[N/x]) &= V(((\lambda y.M')[N/x])(N'[N/x])) \text{ defn of substitution} \\
&= V((\lambda y.(M'[N/x]))(N'[N/x])) \text{ defn of substitution} \\
&= V(M'[N/x])V(N'[N/x]/y) \text{ defn of marked evaluation} \\
&= (V(M')[V(N)/x])[V(N')[V(N)/x]/y] \text{ structural induction} \\
&= (V(M'))[(V(N')/y)][V(N)/x] \text{ property of substitution} \\
&= V((\lambda y.M')N')[V(N)/x] \text{ defn of marked evaluation}
\end{aligned}$$

□

We now have to examine more carefully how to build a marked context, in particular, it is useful to have a structural inductive definition of the construction of context as then we can use proofs by structural induction. Here is the way in which a marked context is built (note the change from table 6):

- The empty context is a context for the term $\Gamma \vdash M$ which uses the variable environment Γ : $C_\Gamma[M] = M$
- Adding application: if $C_\Gamma[M]$ is a context for M then

$$N C_\Gamma[M] \text{ and } C_\Gamma[M] L$$

for any marked terms N and L , are contexts for M .

- Unmarked abstraction: if $C_\Gamma[M]$ is a context for M then $\lambda x.C_\Gamma[M]$ is a context for M . Note here any $x \in M$ becomes bound.

$\frac{}{\Gamma \vdash _ = \text{ld}_{\Gamma}[_] \text{ mcontext}} \text{ Empty}$	
$\frac{\Gamma, x \vdash C_{\Gamma'}[_] \text{ mcontext}}{\Gamma \vdash \lambda x. C_{\Gamma'}[_] \text{ mcontext}} \text{ abst}$	$\frac{\Gamma \vdash C_{\Gamma'}[_] \text{ mcontext} \quad \Gamma'' \subseteq \Gamma'}{\Gamma \vdash C_{\Gamma''}[_] \text{ mcontext}} \text{ weak}$
$\frac{\Gamma \vdash M \text{ mterm} \quad \Gamma \vdash C_{\Gamma'}[_] \text{ mcontext}}{\Gamma \vdash M C_{\Gamma'}[_] \text{ mcontext}} \text{ Lapp}$	$\frac{\Gamma \vdash C_{\Gamma'}[_] \text{ mcontext} \quad \Gamma \vdash N \text{ mterm}}{\Gamma \vdash C_{\Gamma'}[_] N \text{ mcontext}} \text{ Rapp}$
$\frac{\Gamma, x \vdash C_{\Gamma'}[_] \text{ mcontext} \quad \Gamma \vdash N \text{ mterm}}{\Gamma \vdash (\lambda x. C_{\Gamma'}[_]) N \text{ mcontext}} \text{ mAbst}_L$	$\frac{\Gamma, x \vdash M \text{ mterm} \quad \Gamma \vdash C_{\Gamma'}[_] \text{ mcontext}}{\Gamma \vdash (\lambda x. M) C_{\Gamma'}[_] \text{ mcontext}} \text{ mAbst}_R$

Table 6: Construction of marked contexts

- Marked abstraction: if $C_{\Gamma}[M]$ is a context for M then $(\lambda x. C_{\Gamma}[M])$ is a context for M . Note that x becomes bound in M .

Lemma 2.9

- (i) $V(C[M]) = V(C[V(M)])$;
- (ii) If $M \rightarrow N$ is a rewrite between unmarked terms then:

$$\begin{array}{ccc}
 C[M] & \longrightarrow & C[N] \\
 * \downarrow & & \downarrow * \\
 V(C[M]) & \xrightarrow{*} & V(C[N])
 \end{array}$$

PROOF: We prove both parts by a structural induction on building the context:

- (i) If the context is empty then as $V(V(M)) = V(M)$ the result follows. If the context is built using an application we have (for left application):

$$\begin{aligned}
 V(NC[M]) &= V(N)V(C[M]) \quad \text{defn of marked evaluation} \\
 &= V(N)V(C[V(M)]) \quad \text{structural induction} \\
 &= V(NV(C[V(M)])) \quad \text{defn of marked evaluation}
 \end{aligned}$$

For unmarked abstractions we have:

$$\begin{aligned}
 V(\lambda x. C[M]) &= \lambda x. V(C[M]) \quad \text{defn of marked evaluation} \\
 &= \lambda x. V(C[V(M)]) \quad \text{structural induction} \\
 &= V(\lambda x. C[V(M)]) \quad \text{defn of marked evaluation}
 \end{aligned}$$

Finally for the first case of marked abstraction we have:

$$\begin{aligned}
V((\underline{\lambda}y.C[[M]])N) &= V(C[[M]])[V(N)/y] \text{ defn of marked evaluation} \\
&= V(C[[V(M]])[V(N)/y] \text{ structural induction} \\
&= V((\underline{\lambda}y.C[[V(M]])N) \text{ defn of marked evaluation}
\end{aligned}$$

The second is similar.

(ii) The result is clear for the empty context. For abstraction on the left (on the right is similar) we must produce a rewrite sequence

$$V(LC[[M]]) \rightarrow V(LC[[M]])$$

on the assumption that we have a sequence of rewrites $V(C[[M]]) \rightarrow V(C[[M]])$. But as $V(LC[[M]]) = V(L)V(C[[M]])$ and $V(LC[[M]]) = V(L)V(C[[M]])$ then we may use the original rewrite sequence as applied on the right of the application.

A similar argument works for extending the context by an unmarked lambda.

The difficult case is when the context is extended by a marked λ . There are two cases:

$$\begin{array}{ccc}
(\underline{\lambda}y.C[[M]])M' & \longrightarrow & (\underline{\lambda}y.C[[N]])M' \\
\downarrow * & & \downarrow * \\
V(C[[M]][M'/y]) & & V(C[[N]][M'/y]) \\
\parallel & & \parallel \\
V(C[[M]][V(M')/y]) & \xrightarrow{*} & V(C[[N]][V(M')/y])
\end{array}$$

Here by assumption we have a rewrite sequence $C[[M]] \xrightarrow{*} C[[N]]$ and the last line simply substitutes this (note the step down uses the substitution lemma above).

Finally we have the last case (which is also the hardest case):

$$\begin{array}{ccc}
(\underline{\lambda}y.M')C[[M]] & \longrightarrow & (\underline{\lambda}y.M')C[[N]] \\
\downarrow * & & \downarrow * \\
V(M'[C[[M]]/y]) & & V(M'[C[[N]]/y]) \\
\parallel & & \parallel \\
V(M')[V(C[[M]]/y)] & \xrightarrow{*} & V(M')[V(C[[N]]/y)]
\end{array}$$

In this case wherever y occurs we must do the rewrite sequence $C[[M]] \xrightarrow{*} C[[N]]$. It is then a matter of going through each occurrence of y and doing the rewrites required at that occurrence. This gives a possibly huge number of rewrites as y can occur often in the term $V(M')$... but it is still a finite number of occurrences so the parallel rewriting of subterms can be turned into a sequential rewriting!

□

This completes the proof of the Church-Rosser theorem in some detail.

3 Representing data in the λ -calculus

To represent computations in the λ -calculus one must first agree on how to represent data. The two most basic datatype are the **Booleans** and the ability to form **products**. The element (of the Booleans are **True** and **False**. The elements of a product of two sets are all the pairs $\langle a, b \rangle$ where a is from the first set and b is from the second. However, having data is of no value if one cannot *use* the data. Thus, for Booleans, a key feature is that one can use them as a condition to control a program; for products one must be able to extract the information in the components – that is *project* to obtain the entries.

Next one wants to be able to represent more general (infinite) inductive data such as the natural numbers, lists, and trees. These, it turns out can all be represented very uniformly in the λ -calculus. At the time these ideas were being developed - in the 1930s - representing the natural numbers was of crucial significance as it was for functions between the natural numbers that the first notions of being “computable” or (partial) “recursive” arose. We shall trace these developments showing that, in the λ -calculus, one can represent *all* the computable functions.

Another important way to encode computable functions was to use a Turing machines. Once one has general inductive data (lists in particular) and general recursion it is straightforward to program Turing machines in the λ -calculus. However, this does mean that it is useful to know how these encodings work in the λ -calculus.

3.1 Booleans

In the λ -calculus it is traditional to set:

$$\begin{aligned}\text{True} &:= \lambda xy.x \\ \text{False} &:= \lambda xy.y \\ \text{If } zt f &= zt f \Leftrightarrow \text{If} = \lambda zt f.zt f\end{aligned}$$

In the last of these we are using a the more natural form for the definition of **If**: this uncurries function symbol being defined to remove the outer lambda abstractions. The curried for, which gives **If** as a close expression is given on the right.

Given the conditional one may then define all the usual functions on Booleans:

$$\begin{aligned}\text{Not } b &= \text{If } b \text{ False True} \\ \text{And } b_1 b_2 &= \text{If } b_1 (\text{If } b_2 \text{ True False}) \text{ False} \\ \text{Or } b_1 b_2 &= \text{If } b_1 \text{ True } (\text{If } b_2 \text{ True False})\end{aligned}$$

3.2 Products

Here we must encode the ability to form pairs:

$$\begin{aligned}\langle f, g \rangle &:= \lambda h.hfg \\ \pi_0 &:= \lambda z.z(\lambda xy.x) \\ \pi_1 &:= \lambda z.z(\lambda xy.y)\end{aligned}$$

Note that when we project a pair we have

$$\begin{aligned}
 \pi_0 \langle f, g \rangle &:= (\lambda z. z(\lambda xy. x)) \langle f, g \rangle \\
 &\rightarrow \langle f, g \rangle (\lambda xy. x) \\
 &:= (\lambda h. hfg) (\lambda xy. x) \\
 &\rightarrow (\lambda xy. x) fg \\
 &\rightarrow (\lambda y. f) g \\
 &\rightarrow f
 \end{aligned}$$

Notice that with pairing and Boolean operations one can represent all finite stes (as binary numbers of a fixed width) and then one can define any function using Boolean expressions. This does not allow one to define functions on infinite sets, however.

3.3 The natural numbers

The first significant difficulty is to represent inductive data such as the natural numbers in the λ -calculus. The following is the Church encoding for the natural numbers:

$$\begin{aligned}
 0 &:= \lambda xy. x \\
 1 &:= \lambda xy. yx \\
 2 &:= \lambda xy. y(yx) \\
 3 &:= \lambda xy. y(y(yx)) \\
 &\dots
 \end{aligned}$$

They are called the “Church numerals”.

This is by no means the only way to do it but it is the most natural way in a sense we shall soon make clear. But what are **Zero**, **Succ**, and the **fold** or iteration functions over numbers? Here they are:

$$\begin{aligned}
 \text{Zero} &= \lambda xy. x \\
 \text{Succ } n &= \lambda xy. y(nxy) \\
 \text{Fold } zfn &= nzf
 \end{aligned}$$

This allows us to define addition and multiplication:

$$\begin{aligned}
 \text{Add } n m &= \text{Fold } m \text{ Succ } n \\
 \text{Mult } n m &= \text{Fold } \text{Zero} (\text{Add } m) n
 \end{aligned}$$

To illustrate how this works let us add one and one:

$$\begin{aligned}
\text{Add } \mathbf{1} \ \mathbf{1} &:= \text{Fold } \mathbf{1} \ \text{Succ } \mathbf{1} \\
&:= \mathbf{1} \ \mathbf{1} \ \text{Succ} \\
&:= (\lambda xy.yx) \ \mathbf{1} \ \text{Succ} \\
&\rightarrow \text{Succ } \mathbf{1} \\
&:= (\lambda xy.y(\mathbf{1} \ xy)) \\
&:= (\lambda xy.y((\lambda xy.yx)xy)) \\
&\rightarrow (\lambda xy.y(yx)) \\
&=: \mathbf{2}
\end{aligned}$$

A problem which gave Church and his student Kleene pause was the question of how to encode the predecessor function. This seems like an easy problem, however, removing one successor is not so easy when you are trying to do it purely from the primitives the λ -calculus provides. The story goes that Kleene had the inspiration when he was sitting in the dentists chair ... here is what he did:

$$\text{Pred } n = \pi_1(\text{Fold } \langle \mathbf{0}, \mathbf{0} \rangle (\lambda z. \langle \text{Succ}(\pi_1 z), \pi_1 z \rangle) n)$$

Why is the predecessor function so important? A basic ability which is required in many programs is the ability to determine whether two numbers are equal. The basic scheme for this was to use truncated subtraction, or the `monus` function, and to test whether the two ways of doing this subtraction where both zero (using the `iszero` function. However, the easiest way to define truncated subtraction was to use the predecessor function. Here are the programs:

$$\begin{aligned}
\text{iszero } n &= \text{Fold True } (\lambda x.\text{False}) \ n \\
\text{monus } n \ m &= \text{Fold } n \ \text{Pred } m \\
\text{Eqnm} &= \text{And } (\text{iszero } (\text{monus } n \ m)) \ (\text{iszero } (\text{monus } m \ n)).
\end{aligned}$$

The predecessor function, of course, is easily programmed if one has a case function as in Haskell:

$$\text{Pred } x = \text{case } x \ \text{of} \ \left| \begin{array}{l} \text{Zero} \quad \mapsto \ \text{Zero} \\ \text{Succ } x' \quad \mapsto \ x' \end{array} \right.$$

Thus, a more fundamental problem is to program the `case` construct. Not surprisingly this requires a similar trick to the predecessor function. Our first step is to remove the “syntactic sugar” from the case statement to obtain a combinator:

$$\text{case } n \ \text{of} \ \left| \begin{array}{l} \text{Zero} \quad \mapsto \ t_1 \\ \text{Succ } n' \quad \mapsto \ t_2 \end{array} \right. := \text{Case } n \ t_1 \ (\lambda n'. t_2)$$

Next we show how to implement this combinator:

$$\text{Case } n \ t_1 \ f := \pi_1(\text{fold} \langle \text{Nil}, t_1 \rangle (\lambda x. \langle \text{Succ}(\pi_0 x), f(\pi_0 x) \rangle))$$

If we allow pattern matching for products we can present this a little more neatly as:

$$\text{Case } n \ t_1 \ f := \pi_1 (\text{fold} \langle \text{Nil}, t_1 \rangle (\lambda \langle x_0, x_1 \rangle. \langle \text{Succ } x_0, f \ x_0 \rangle)).$$

Note that there is an essential use of products in defining this function too..

3.4 Representing list

One way to prove that one can program all computable functions is to show that one can program a Turing machine. To do this the most difficult aspect is to represent the tape: a standard way of doing this is to view it as a list of cells left of head, the cell under the head, and a list of cells to the right of the head. Moving the head to the left involves, for example, pushing the current cell under the head onto the right list and popping the top element off the left list to become the new cell under the head. If the list to be popped is empty then one inserts a black cell under the head. Thus, once one has lists one can, in principle, program the basic mechanism of any Turing machine.

There are two problems to surmount: first one must be able to represent lists in the λ -calculus and second one must be able to iterate its “step” until the machine reaches a halt state. We shall see how this can be done soon. Modulo this aspect one will then be confident that all computable functions can be programmed.

Here is how lists are represented:

$$\begin{aligned} [] &= \lambda nc.n \\ [a_1] &= \lambda nc.ca_1n \\ [a_1, a_2] &= \lambda nc.ca_1(ca_2n) \\ [a_1, a_2, a_3] &= \lambda nc.ca_1(ca_2(ca_3n)) \\ &\dots \end{aligned}$$

Here are the basic functions associated with lists:

$$\begin{aligned} \text{Nil} &= \lambda nc.n \\ \text{Cons } a \ l &= \lambda nc.ca(lnc) \\ \text{FoldList } v \ g \ l &= lvg \end{aligned}$$

We can than define the basic functions associated with list such as appending lists, mapping over lists, etc.

$$\begin{aligned} \text{append } l_1 \ l_2 &= \text{FoldList } l_2 \ \text{Cons } l_1 \\ \text{MapList } f \ l &= \text{FoldList } \text{Nil} \ (\lambda az.\text{Cons } (f \ a) \ z)l \end{aligned}$$

Of course, as for the natural numbers, there is still the problem of programming the case construct. We start by removing the syntactic sugar:

$$\text{case } l \ \text{of} \ \left| \begin{array}{ll} \text{Nil} & \mapsto \ t_1 \\ \text{Cons } a \ as & \mapsto \ t_2 \end{array} \right. := \text{CaseList } l \ t_1 \ (\lambda a \ as.t_2)$$

Then `CaseList` can be programmed as follows:

$$\text{CaseList } l \ t \ f := \pi_1 (\text{FoldList } \langle \text{Nil}, t \rangle (\lambda a \langle l, - \rangle. \langle \text{Cons } a \ l, f \ a \ l \rangle) \ l)$$

Again notice the essential use of the product and the use of (product) patterns. This allow us, for example, to program the “tail” function for a list as follows:

$$\text{Tail } l := \text{case } l \text{ of } \left\{ \begin{array}{l} \text{Nil} \quad \mapsto \text{Nil} \\ \text{Cons } a \ as \mapsto as \end{array} \right.$$

3.5 Representing trees

How do we represent trees? By now you may be suspecting that there is a definite pattern emerging! Before describing the general pattern let us do one more simple example. Start with the Haskell data declaration:

$$\text{data Tree } a = \left\{ \begin{array}{l} \text{Leaf } a \\ \text{Node } (\text{Tree } a) (\text{Tree } a) \end{array} \right.$$

The first step is to write down some simple trees:

$$\begin{aligned} \text{Leaf } x &= \lambda l \ n. l \ x \\ \text{Node } (\text{Leaf } x_1) (\text{Leaf } x_2) &= \lambda l \ n. n \ (l \ x_1) \ (l \ x_2) \\ &\dots \end{aligned}$$

Now one defines:

$$\begin{aligned} \text{Leaf } x &= \lambda l \ n. l \ x \\ \text{Node } t_1 \ t_2 &= \lambda l \ n. n \ (t_1 \ l \ n) \ (t_2 \ l \ n) \\ \text{FoldTree } f \ g \ t &= t \ f \ g. \end{aligned}$$

The last definition indicates that the objective of the representatin is to make a tree into a curried fold function for the tree. Thus, in a very real sense data in the λ -calculus is potential computation – as it curried computation.

Finally, and somewhat more tricky, is to write down the `case` function for a tree. First, remove the syntactic sugar:

$$\text{case } t \text{ of } \left\{ \begin{array}{l} \text{Leaf } x \quad \mapsto v_1 \\ \text{Node } t_1 \ t_1 \mapsto v_2 \end{array} \right. := \text{CaseTree } t \ (\lambda x. v_1) \ (\lambda t_1 \ t_2. v_2)$$

Next write down the `CaseTree` function using products:

$$\text{CaseTree } t \ f \ g := \pi_1 (\text{FoldTree } (\lambda x. \langle \text{Leaf } x, f \ x \rangle) (\lambda \langle t_1, - \rangle \langle t_1, - \rangle. \langle \text{Node } t_1 \ t_2, g t_1 t_2 \rangle) \ t)$$

3.6 Inductive data in general

Representing data in the λ -calculus is just a translation process. Unfortunately, this does not mean it is particularly easy BUT once you have tried it a few times you will probably be able to appreciate the general process. Here is an explicit description of the process:

1. Start with an arbitrary datatype definition:

$$\text{data } D \ a_1 \ \dots \ a_n = \begin{array}{l} \text{Cons}_1(T_{11} \ a_1 \ \dots \ a_n \ (D \ a_1 \ \dots \ a_n)) \ \dots \ (T_{1m_1} \ a_1 \ \dots \ a_n \ (D \ a_1 \ \dots \ a_n)) \\ \dots \\ \text{Cons}_p(T_{p1} \ a_1 \ \dots \ a_n \ (D \ a_1 \ \dots \ a_n)) \ \dots \ (T_{pm_p} \ a_1 \ \dots \ a_n \ (D \ a_1 \ \dots \ a_n)) \end{array}$$

First notice that there are p constructors: each constructor has a type. This is rather hidden in the Haskell syntax above so lets make it explicit:

$$\text{Cons}_i :: (T_{i1} \ a_1 \ \dots \ a_n \ (D \ a_1 \ \dots \ a_n)) \rightarrow \dots \rightarrow (T_{im_i} \ a_1 \ \dots \ a_n \ (D \ a_1 \ \dots \ a_n)) \rightarrow (D \ a_1 \ \dots \ a_n)$$

2. Notice also that we are using types T_{ij} which we are assuming have already been introduced. Furthermore, we shall assume that with each of these already introduced types there is a MapT_{ij} function:

$$\frac{f :: x \rightarrow y}{\text{MapT}_{ij} \ f :: (T_{ij} \ a_1 \ \dots \ a_n \ x) \rightarrow (T_{ij} \ a_1 \ \dots \ a_n \ y)}.$$

3. Now we know we must construct the fold function: in this construction each constructor turns into a function. It is useful to think about the types of these functions and of how the fold will be written. The types of the functions are obtained by replacing the type of the data being defined by an arbitrary type, b :

$$f_i :: (T_{i1} \ a_1 \ \dots \ a_n \ b) \rightarrow \dots \rightarrow (T_{im_i} \ a_1 \ \dots \ a_n \ b) \rightarrow b$$

This makes the type of the fold look:

$$\begin{array}{l} ((T_{11} \ a_1 \ \dots \ a_n \ b) \rightarrow \dots \rightarrow (T_{1m_1} \ a_1 \ \dots \ a_n \ b) \rightarrow b) \\ \rightarrow \dots \\ \text{foldD} :: \rightarrow ((T_{p1} \ a_1 \ \dots \ a_n \ b) \rightarrow \dots \rightarrow (T_{pm_p} \ a_1 \ \dots \ a_n \ b) \rightarrow b) \\ \rightarrow (D \ a_1 \ \dots \ a_n) \\ \rightarrow b \end{array}$$

This just sorts out the typing involved but you must be clear on this to start with!

4. Now the next step is to represent elements of this data type in the λ -calculus and the basic idea is to represent it as an abstracted fold. This means that in the end you will know that the actual fold becomes:

$$\text{foldD } c_1 \ \dots \ c_p \ d = d \ c_1 \ \dots \ c_p.$$

Of course you are done if you can write down the constructors as then you can inductively construct representing elements of the datatype (and so be able to define inductively what it means to be a element of this datatype).

5. So what do the constructors look like?

$$\text{Cons}_i \ y_1 \ \dots \ y_{m_i} := \lambda \ c_1 \ \dots \ c_p. c_i(\text{MapT}_{i1}(\lambda y. y \ c_1 \ \dots \ c_p) y_1) \ \dots \ (\text{MapT}_{im_i}(\lambda y. y \ c_1 \ \dots \ c_p) y_{m_i})$$

The use of the map functions may initially be surprising but, recall, that one needs to transmit the “internal” names of the constructors to the occurrences of the datatype inside each type: this is what the map does.

So with this pattern in mind it is possible to write down representative elements of any datatype. Clearly, as the datatype becomes more complex so the structure of the elements becomes more complex. Of course, this is exactly why you do not want to write down such elements but rather be able to rely on the pattern!

3.7 Rose trees

A small example to illustrate the necessity of the map in describing the constructors. Consider the datatype:

$$\text{data Rose } f \ v = \begin{cases} \text{Var } v \\ \text{Rs } f \ (\text{List } (\text{Rose } a \ v)) \end{cases}$$

This is a Rose tree with “function symbols” at the nodes and “variables” at the leaves. The constructors, using the above process, are:

$$\begin{aligned} \text{Var } x &= \lambda v \ r. vx \\ \text{Rs } a \ l &= \lambda v \ r. r \ a \ (\text{MapList } (\lambda t. t \ v \ r) \ l) \end{aligned}$$

4 General recursion using fixed points

So far we have only shown how to represent all the (higher-order) primitive recursive functions. What we have not done is to mimic the ability a program has to recursively call itself or to iterate (possibly forever). Fundamental in this is the use of fixed point combinators.

4.1 Fixed points

A fixed point combinator Y has the property that for every f we have $Yf = f(Yf)$. this can be used to encode a recursive call to a function. As suppose we define $gx = Mg$ where M is some combinator which reuses the function g recursively then we can define $g = Y(\lambda gx.Mg)$ this means:

$$\begin{aligned} gN &:= (Y(\lambda gx.Mg))N \\ &\rightarrow (\lambda x.M(Y(\lambda gx.Mg)))N \\ &\rightarrow M[N/x](Y(\lambda gx.Mg)) \end{aligned}$$

Note how this presents the body of the recursive call to act on the argument N once the Y combinator has been unwrapped if we do a leftmost outermost reduction.

There are two well known examples of fixed point combinators (although there are infinitely many!):

$$\begin{aligned} Y &= \lambda f.(Af)(Af) \\ Af &= \lambda x.f(xx) \\ Y' &= A'A' \\ A' &= \lambda xy.y(xxy) \end{aligned}$$

The first is due to Curry the second is due to Turing. To show that Y is a fixed point combinator we have:

$$\begin{aligned} (Yf) &:= (Af)(Af) := (\lambda x.f(xx))(Af) \\ &\rightarrow f((Af)(Af)) := f(Yf) \end{aligned}$$

To show that Y' is a fixed point combinator we have:

$$\begin{aligned} Y'f &:= (\lambda xy.y(xxy))(\lambda xy.y(xxy))f \\ &\rightarrow (\lambda y.y(Y'y))f \\ &\rightarrow f(Y'f) \end{aligned}$$

How do we know that these two expressions are not equivalent?

$$\begin{aligned} Y &= \lambda f.(\lambda x.f(xx))(\lambda x.f(xx)) \quad \text{due to Curry} \\ Y' &= (\lambda xy.y(xxy))(\lambda xy.y(xxy)) \quad \text{due to Turing} \end{aligned}$$

To see that they are not the same look at the β -reductions of Y setting $Af = \lambda x.f(xx)$

$$Y \rightarrow \lambda f.f((Af)(Af)) \rightarrow \lambda f.f(f((Af)(Af))) \rightarrow \dots \rightarrow \lambda f.f^n((Af)(Af)) \rightarrow \dots$$

Now consider Turing's fixed point combinator set, $X = \lambda x f.f(xxf)$, then

$$Y' = XX \rightarrow \lambda f.f(XXf) \rightarrow \dots \rightarrow \lambda f.f^n(XXf) \rightarrow \dots$$

so they never have a common β -reduction and so they cannot be equal.

Here is a challenge: can you prove that there are infinitely many distinct fixed point?

4.2 Using fixed points for general recursion

(1) Rather than define `append` using a fold we can implement it recursively as

$$\text{append } x \ y := \text{case } x \ \text{of} \ \left| \begin{array}{ll} \text{nil} & \mapsto y \\ \text{cons } a \ as & \mapsto \text{cons } a \ (\text{append } as \ y) \end{array} \right.$$

This translates to

$$\text{append} := Y(\lambda f \ x \ y. \text{case } x \ \text{of} \ \left| \begin{array}{ll} \text{nil} & \mapsto y \\ \text{cons } a \ as & \mapsto \text{cons } a \ (f \ as \ y) \end{array} \right. \ .)$$

Note that the correct behaviour of this relies on the `case` construct which when the first argument is `nil` will not recall the function.

(2) To implement the recursive (non-terminating) function $\mathbf{f} \ n := n : (\mathbf{f} \ (n+1))$ we set $\mathbf{f} = Y(\lambda f n.n : (f(n+1)))$ reducing gives the desired recursive behavior:

$$\begin{aligned} \mathbf{f} \ n &:= (Y(\lambda f n.n : (f(n+1))))n \\ &\rightarrow (\lambda f n.n : (f(n+1)))(Y(\lambda f n.n : (f(n+1))))n \\ &\rightarrow n : ((Y(\lambda f n.n : (f(n+1))))n + 1) \\ &=: n : (\mathbf{f}(n+1)) \end{aligned}$$

5 Computability and the λ -calculus

When people first started studying computability it seemed very natural for them to focus on the question of whether a function $\mathbb{N}^n \rightarrow \mathbb{N}$ on the natural numbers was “computable” – in the sense that it could be calculated by machine. Thus, the first notions of computability emerged as those functions *on the natural numbers* which were computable. To determine the power of the λ -calculus people therefore considered the question of which functions on the natural numbers were “represented”.

A partial function $F : \mathbb{N}^n \rightarrow \mathbb{N}$ is represented in the λ -calculus if there is a λ -term ψ_F such that whenever $F(m_1, \dots, m_n) = m_0$ then $\psi_F m_1 \dots m_n =_{\beta} m_0$ (where m_i is the Church numeral for m_i) and, furthermore, when $F(m_1, \dots, m_n)$ is undefined, then the corresponding λ -term does not have a normal form.

Kleene developed a simple theorem which we shall use to show that all recursive functions can be represented in the λ -calculus:

Theorem 5.1 *The computable (partial) functions on the natural numbers are exactly those generated by primitive recursion and (pure) “minimization” (Kleene’s μ -operator).*

We shall now develop these ideas briefly.

5.1 Primitive recursion

The primitive recursive functions are (total) functions $f : \mathbb{N}^n \rightarrow \mathbb{N}$ (this include *constants* such as zero : $\mathbb{N}^0 = 1 \rightarrow \mathbb{N}$) which are generated by:

Basic functions: The following functions are primitive recursive:

- (a) zero : $1 \rightarrow \mathbb{N}$
- (b) succ : $\mathbb{N} \rightarrow \mathbb{N}$
- (c) All projections $\pi_i^n : \mathbb{N}^n \rightarrow \mathbb{N}; (x_1, \dots, x_n) \mapsto x_i$ (here $1 \leq i \leq n$).

Composition If $f : \mathbb{N}^n \rightarrow \mathbb{N}$ and $(g_i : \mathbb{N}^m \rightarrow \mathbb{N})_{i=1, \dots, n}$ are primitive recursive functions then their composite h is primitive recursive:

$$h : \mathbb{N}^m \rightarrow \mathbb{N}; (x_1, \dots, x_m) \mapsto f(g_1(x_1, \dots, x_m), \dots, g_n(x_1, \dots, x_m))$$

Primitive recursion: If $g : \mathbb{N}^n \rightarrow \mathbb{N}$ and $h : \mathbb{N}^{n+2} \rightarrow \mathbb{N}$ are primitive recursive functions then $f : \mathbb{N}^{n+1} \rightarrow \mathbb{N}$ defined by:

$$\begin{aligned} f(\text{zero}, x_1, \dots, x_n) &= g(x_1, \dots, x_n) \\ f(\text{succ}(n), x_1, \dots, x_n) &= h(f(n, x_1, \dots, x_n), n, x_1, \dots, x_n) \end{aligned}$$

is primitive recursive.

All the basic arithmetic functions (addition, multiplication, exponentiation, predecessor, monus, ..) are all primitive recursive. Indeed primitive recursion already contains so much power that any reasonable (terminating) program will be primitive recursive. in particular, all the complexity

classes such as polynomial time, polynomial space, exponential-time are all within primitive recursion.

We can quite easily simulate primitive recursion in the λ -calculus. At this stage all except for the last requirement, that is actually producing a new function using primitive recursion, are clearly satisfied by the λ -calculus. Here is how we can produce the function f defined by primitive recursion (assuming that g and h are represented) using the fold:

$$f(n, x_1, \dots, x_n) = \pi_1^2(\text{foldNat}\langle g(x_1, \dots, x_n), 0 \rangle(\lambda z. \langle h(\pi_1^2 z, \pi_2^2 z, x_1, \dots, x_n), \pi_2^2 z + 1 \rangle))$$

Notice the use of pairs in this definition.

Primitive recursive and total recursive functions

Primitive recursive functions are along way from representing all the computable (total) functions. This can be demonstrated by the using Cantor's diagonal argument.

Consider all the primitive recursive functions $\mathbb{N} \rightarrow \mathbb{N}$: they can certainly be enumerated (here one must resort to some sort of Gödel numbering to establish this: we shall – for the moment – be happy to assume that such exists):

$$f_0, f_1, f_2, f_3, \dots$$

We allow repetitions in this list – in fact we have to allow repetitions – but we can guarantee that every primitive recursive function will appear. Now define the function

$$g : \mathbb{N} \rightarrow \mathbb{N}; i \mapsto f_i(i) + 1$$

then g is not in the collection as for all i we have $g(i) = f_i(i) + 1 \neq f_i(i)$. In this way, so long as the enumeration was computable, we can build a new computable function from our old collection. This sketched argument shows that it is not possible to generate an enumeration of all the *total* computable functions. Somewhat surprisingly – considering the total functions are a subset – it *is* possible to enumerate all the partial computable functions. Notice that Cantor's diagonal argument cannot be applied to an enumeration of partial functions as $f_i(i)$ may not be defined for many values of i .

The diagonal argument can be used in a number of slightly more sophisticated ways: for example if we wish to build a function which eventually grows faster than any function in our enumeration we can use the diagonal argument to do so. By “eventually grows faster” (or “majorizes”) we mean a function g such that for each f in the enumerable collection there is an $m \in \mathbb{N}$ such that for all $m' > m$ $g(m') > f(m')$. Using Cantor's argument we can define g by:

$$g(m) = \max\{f_i(m) \mid i \leq m\} + 1$$

the function g thus defined has $g(m) > f_n(m)$ for all $n < m$. Clearly, to build such fast growing functions we need not enumerate *all* the functions in the set but merely a subset which dominates any given functions in the set: this allows a more economical description and this idea underlay the development of Ackermann's function.

It turns out just using folds one can in the λ -calculus represent far more than just primitive recursive functions: one can represent all the *higher-order* primitive recursive functions. This is a *much* larger class. In particular, one can implement Ackermann's function as we now show.

Here is a first definition of Ackermann's function:

$$A(m, n) = \begin{cases} n + 1 & \text{if } m = 0 \\ A(m - 1, 1) & \text{if } m > 0 \text{ and } n = 0 \\ A(m - 1, A(m, n - 1)) & \text{if } m > 0 \text{ and } n > 0. \end{cases}$$

However, we can re-express this in a higher-order fashion:

$$\begin{aligned} \text{A Zero} &= \text{Succ} \\ \text{A (Succ } m) &= \text{IT (A } m) \\ \text{IT } f \text{ Zero} &= f (\text{Succ Zero}) \\ \text{IT } f (\text{Succ } n) &= f (\text{IT } f n) \end{aligned}$$

which translates into folding on the natural numbers as:

$$A m n := (\text{FoldNat Succ } [\lambda f. \text{FoldNat } (\lambda x. f (\text{Succ Zero})) (\lambda h x. f (h x)) n] m n)$$

This also show how to represent Ackermann's function using higher-order primitive recursion ...

We now know that the λ -calculus can represent all the primitive recursive functions and much more: namely, the higher-order primitive recursive functions. Furthermore, we know that there are computable total functions which are definitely not primitive recursive functions or indeed higher-order primitive recursive!

5.2 Minimization

To capture all computable functions Kleene introduced the following minimization operator over the class of all primitive recursive functions:

$$\frac{g : \mathbb{N}^{n+1} \rightarrow \mathbb{N}}{\text{mini } g : \mathbb{N}^n \rightarrow \mathbb{N}}$$

where $\text{mini } g$ is, in general, a partial function defined by:

$$(\text{mini } g)x_1 \dots x_n := \mu n. g(n, x_1, \dots, x_n) := \min\{n \mid g(n, x_1, \dots, x_n) = 0\}$$

Here we regard μn as binding n in $g(n, x_1, \dots, x_n)$.

To give a concrete feel for the operation of minimization consider a function g with values:

$$\begin{aligned} g(0, 4) &= 22 \\ g(1, 4) &= 1 \\ g(2, 4) &= 14 \\ g(3, 4) &= 0 \\ g(4, 4) &= 119 \\ &\dots \end{aligned}$$

then $\mu n. g(n, 4) = 3$. Importantly, $x \mapsto \mu n. g(n, x)$ can be a partial map, because g may never hit zero. The use of the binder μ is a useful notation as sometimes it is not the first argument over which we wish to minimize but rather some other argument. In this case we can write:

$$\mu n. g(x_1, n, x_2, \dots, x_n) = \min\{n \mid g(x_1, n, x_2, \dots, x_n) = 0\}$$

and, thus, we treat μ as a binding operator which binds variables much as λ -abstraction does.

Let us make some simple observations on the minimization operator used on primitive recursive functions:

Primitive recursive functions using minimization:

First, we observe that we can re-express any primitive recursive function as a minimization. Let $f : \mathbb{N}^k \rightarrow \mathbb{N}$ be any primitive recursive function then set

$$g : \mathbb{N}^{k+1} \rightarrow \mathbb{N}; (n, x_1, \dots, x_k) \mapsto (f(x_1, \dots, x_k) \dot{-} n) + (n \dot{-} f(x_1, \dots, x_k))$$

then $\mu n.g(n, x_1, \dots, x_k) = f(x_1, \dots, x_k)$. This means minimization can express all primitive recursive functions as a minimization of a primitive recursive function.

Composition of minimized functions

Now something harder! Consider a composition of minimized primitive recursive functions

$$\begin{aligned} &(\lambda x_1 \dots x_k. \mu n. g_0(n, x_1, \dots, x_k)) \\ &(\mu m_1. g_1(m_1, y_1, \dots, y_p)) \\ &\dots \\ &(\mu m_k. g_k(m_k, y_1, \dots, y_p)) \end{aligned}$$

we wish to represent it as a single minimization over a single primitive recursive function. First we will simplify notation a bit by representing y_1, \dots, y_p by a single variable y so that we have a slightly more manageable term:

$$(\lambda x_1 \dots x_k. \mu n. g_0(n, x_1, \dots, x_k))(\mu m_1. g_1(m_1, y)) \dots (\mu m_k. g_k(m_k, y))$$

To achieve this composition we will employ an enumeration of \mathbb{N}^{k+1} satisfying a special property.

We shall write the enumeration as $e_p : \mathbb{N} \rightarrow \mathbb{N}^p$, and when $e_p(m) = (r_1, \dots, r_p)$ we will write $e_{p,i}(m) = r_i$ as the function which gives the i^{th} -coordinate. This enumeration to solve our problem must satisfy the following condition: when $e_p(m) = (r_1, \dots, r_p)$ then, whenever $r'_i \leq r_i$, for each $i = 1, \dots, p$, then there must be an $m' \leq m$ with $e_p(m') = (r'_1, \dots, r'_p)$.

This certainly does not determine the enumeration and, in fact, there are a number of such enumerations which are primitive recursive. Let us consider one such enumeration which is due to Cantor. We need two functions one the inverse of the other:

$$e_2^{-1} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}; (i, j) \mapsto \frac{1}{2}(i+j)(i+j+1) + j$$

This looks like

$e_2^{-1}(i, j)$	0	1	2	3	4	...
0	0	2	5	9	14	...
1	1	4	8	13	...	
2	3	7	12	...		
3	6	11	...			
4	10	...				
⋮						

Notice that the condition which must be satisfied requires that the numbers in the rectangle determined by (i, j) , that is every member of $\{e_2^{-1}(i', j') \mid i' \leq i \ \& \ j' \leq j\}$ must be less or equal to $e_2^{-1}(i, j)$. This is clearly the case with this enumeration and it can be easily generalized to arbitrary dimensions.

We actually need the inverse of this function which we can write as:

$$e_2 : \mathbb{N} \rightarrow \mathbb{N} \times \mathbb{N}; n \rightarrow (i, j) \quad \text{where} \quad \begin{array}{l} c = \lfloor \sqrt{2n} - \frac{1}{2} \rfloor \\ i = n - c(c+1) \\ j = c - i + 2 \end{array}$$

Given such an enumeration now one can define

$$q(r, y) = g(e_1^{k+1}(r), h_1(e_2^{k+1}(r), y), \dots, h_k(e_{k+1}^{k+1}(r), y)) \\ + \sum_{j=1}^k g_j(e_j^{k+1}(r), y)$$

Note that $q(m, y) = 0$ only if each component of the sum is zero – in which case the composite function will be defined. It is now easily checked that the composite required above can be written as

$$\mu n. h(n, y) := \mu n. g(e_{0,k+1}(n), h_1(e_{1,k+1}(n), \dots, h_k(e_{k,k+1}(n), y))).$$

This allows us to rewrite any composite of minimized primitive recursive functions as a single primitive recursive function with exactly one minimization. This is, thus, a “normal form” for recursive functions:

Proposition 5.2 *Every partial recursive function can be expressed in the form $\mu n. p(n, x_1, \dots, x_k)$ where p is a primitive recursive function.*

5.2.1 Minimization in the λ -calculus

Minimization is often characterized as the ability to do “unbounded search”. This is because we can implement the μ -operator by

$$\begin{aligned} \mu n. g(n, x) &= \text{if } g(0, x) == 0 \text{ then } 0 \\ &\quad \text{elseif } g(1, x) == 0 \text{ then } 1 \\ &\quad \text{elseif } g(2, x) == 0 \text{ then } 2 \\ &\quad \dots \end{aligned}$$

where we simply search the numerals in ascending order for a zero of the function. This, of course can be implemented using general recursion as:

$$\begin{aligned} f(x) &= h(0, x) \quad \text{where} \\ h(n, x) &= \text{if } g(n, x) == 0 \text{ then } n \text{ else } h(n+1, x) \end{aligned}$$

Because general recursion can be implemented in the λ -calculus, as discussed above, we can certainly implement this function! This allows us to conclude that all computable functions can be written in the λ -calculus.

Remark: An alternative – and perhaps a simpler – way to show that all computable functions can be programmed in the λ -calculus is to directly program a Turing machine! The fact that we have lists allows us to do this in a quite straightforward manner: thus we have two ways of verifying that all computable functions are present.

The attraction of the above method is that it introduced two families of functions which did not require general recursion in their definition: primitive recursive and higher-order primitive recursive functions.

6 Undecidability of the λ -calculus

The purpose of this section is to prove what is sometimes known as the Scott-Curry undecidability theorem for the λ -calculus: it is an analogue of Rice’s theorem for partial recursive functions. It has the consequence that every non-trivial computable (or recursive) property P on λ -terms which is closed under β -equality is recursively undecidable.

A property on the terms of the λ -calculus is **non-trivial** in case there are terms A and B such that $P(A) = \text{True}$ and $P(B) = \text{False}$. Being **closed to β -equality** means that if $P(N)$ is true and $N =_\beta M$ then $P(M)$ must also hold. Clearly for this to make sense P must be a predicate on terms (that is the syntax) rather than the “meaning” of those terms: being closed to equality means it is also a predicate on the meaning.

Examples of a properties which are close to equality and which we know are non-trivial include:

1. Consider $P(X) \Leftrightarrow X =_\beta N$ for some is fixed term N . This is obviously closed to equality but it is also non-trivial as the property is true of N itself but cannot be true of both True and False as we know they are not β -equal. If there was a *recursive* predicate P then this would mean equality of λ -terms would be decidable: conversely if there were no such recursive predicate – as the Scott-Curry theorem states – this would mean that equality of λ -terms is undecidable.
2. Consider $P(X) \Leftrightarrow X =_\beta N \wedge \text{nf}(N)$: that is P is true of X precisely when X has a normal form. We know Ω does not have a normal form and that True , being in normal form, certainly has a normal form: thus, this property is non-trivial. The confluence of β -reduction, tells us that if two terms are equal and one has a normal form then the other must also have a normal form. Thus this P is closed to β -equality. Having a normal form is equivalent to asking whether a program (evaluated by name) terminates ... so we would not expect this to be decidable.
3. A λ -term N is “solvable” when there are M_1, \dots, M_n (for $n \in \mathbb{N}$ – so n could be 0 such that $N M_1 \dots M_n = I$ (where $I = \lambda x.x$). Equivalently – and this is a theorem – N is solvable if and only if it has a head normal form. This is asking that the reduction of the term will produce *some* information. Consider the predicate $P(X) \Leftrightarrow \text{solvable}(X)$.

This predicate is non-trivial: as $P(I)$ is certainly true (use $n = 0!$) but also $\Omega := \lambda x.x x$ is equally clearly not solvable. To show the predicate is closed with respect to β -equality, note that any term M β -equal to a term N which is solvable must itself be solvable because if $N M_1 \dots M_n =_\beta I$ then as $M =_\beta N$, certainly $M M_1 \dots M_n =_\beta I$. Thus, M is solvable.

Using Scott-Curry’s theorem we can conclude that being solvable is undecidable.

4. The negation of any non-trivial property which is closed to β -equality is automatically non-trivial and closed to equality. Thus, for example we may observe that being insolvable or having no normal form is undecidable as well.

The first example may seem rather counter-intuitive as surely, if one chooses a term M in normal form, say True , one can simply reduce the other term to normal form! However, a little thought about this will make one realize that there is a serious flaw: if the term is not equal to the other term it may not even have a normal form so that it is not clear when one should abandon

the attempt to reduce the term to a normal form. If one abandons the attempt too soon you may actually miss that it does have a normal form which, worse, may actually equal `True`.

This also has the consequence that there can be no computable total function which determines a number corresponding to terms such that, if the term could be reduced to normal form, the process could be completed in less steps than that number. Clearly access to such a number would allow one to decide the equality and indeed whether there was a normal form. Thus, no such number can be computed. However, there is definitely a way to associate such numbers with terms: simply do the reductions and count! Here is the point: this is a partial function.

We shall state the Scott-Curry theorem in its positive form:

Theorem 6.1 (Scott-Curry)

Given any recursive predicate P on the terms of the λ -calculus, for which there are terms A and B with $P(\underline{A}) = \neg P(\underline{B})$, then there are terms A' and B' with $A' =_{\beta} B'$ and $P(\underline{A'}) = \neg P(\underline{B'})$.

For the proof we represent λ -terms internally in the λ -calculus by using the datatype of λ -terms:

```
data Lambda a = Var a
              | App (Lambda a) (Lambda a)
              | Abst a (Lambda a)
```

We may represent the variables by natural numbers. We shall write a term so translated into the λ -calculus as \underline{N} . It is clear that these syntactic terms have an application which we will write as

$$\underline{N} \bullet \underline{M} := \text{App } \underline{N} \ \underline{M}$$

and, furthermore, that so defined $\underline{N} \bullet \underline{M} = \underline{N \ M}$. Recall we are just translating syntax!

Furthermore, we can define a “translation” function T in the λ -calculus such that $T \underline{N} = \underline{\underline{N}}$. This is surprisingly simple as it is a primitive recursive function. By mapping over the λ -term we can replace the variables by their internal forms. To translate a Church numeral into its internal representation we simply fold (for natural numbers) over the Church numeral replacing `Zero` with `Zero` and `Succ` with `$\lambda x.$ Succ \bullet x` . It then remains to translate the λ -term itself: we can achieve this by a fold (for the datatype of λ -terms) which replaces `Var` with `$\lambda x.$ Var \bullet x` , `App` with `$\lambda xy.$ App \bullet x \bullet y` , and `Abst` with `$\lambda xy.$ Abst \bullet x \bullet y` .

This amounts to setting

$$T := \lambda t. \text{Fold}_{\lambda} (\lambda x. \underline{\text{Var}} \bullet x) (\lambda xy. \underline{\text{App}}) (\lambda xy. \underline{\text{Abst}} \bullet x \bullet y) (\text{Map}_{\lambda} (\text{Fold}_{\mathbb{N}} \underline{\text{Zero}} (\lambda x. \underline{\text{Succ}} \bullet x)) t).$$

If the predicate P is recursive we may represent it as a λ -term P which normalizes on *all* the terms \underline{N} of the datatype of λ -terms and produces either `True` or `False` depending on whether P is satisfied or not. As the predicate is non-trivial, we then have, without loss of generality, that $P \underline{A} = \text{True}$ and $P \underline{B} = \text{False}$ for some λ -terms A and B .

Lemma 6.2 (Second recursion theorem)

Given any F then there is a λ -term X such that $X =_{\beta} F \underline{X}$.

PROOF: Define

$$H := \lambda x.F(x \bullet (Tx))$$

where T is the translation function discussed above. This allows:

$$H\underline{H} =_{\beta} F(\underline{H} \bullet (T\underline{H})) =_{\beta} F(\underline{H} \bullet (\underline{H})) =_{\beta} F(\underline{H}\underline{H})$$

and so we may set $X := H\underline{H}$. □

PROOF: (of theorem 6.1) To complete the proof set

$$F := \lambda x.\text{if } Px \text{ then } B \text{ else } A$$

and X to be such that $X =_{\beta} F\underline{X}$ then

If $P\underline{X} = \text{True}$: We get $X =_{\beta} F\underline{X} =_{\beta} B$ so that setting $A' = X$ and $B' = B$ for the theorem provides $P(\underline{A}') = P(\underline{X}) = \text{True}$ and $P(\underline{B}') = P(\underline{B}) = \text{False}$ while $X = \beta B$ proving the theorem.

If $P\underline{X} = \text{False}$: We get $X = F\underline{X} = A$ so that $A' = A$ and $B' = X$ for the theorem $P(\underline{A}') = P(\underline{A}) = \text{True}$ and $P(\underline{B}') = P(\underline{X}) = \text{False}$ while $X = \beta A$ proving the theorem. □

Notice that Theorem 6.1 says something even stronger:

Corollary 6.3 *Given any two disjoint sets of terms which are non-empty and closed to β -equality then there is no recursive predicate which separates the sets.*

7 Evaluation of λ -terms

An evaluation of a λ -term essentially involves repeatedly doing β -reductions until a desired normal form is reached. Of course, doing random reductions of a λ -term may not result in a normal form even if there is one. Recall that $\Omega := (\lambda x.xx) (\lambda x.xx)$ has a never terminating reduction sequence and this means that in the term

$$\text{If True False } \Omega =: (\lambda xyz.x y z)(\lambda xy.x)(\lambda xy,y)((\lambda x.x x) (\lambda x.x x))$$

which clearly “should” evaluate to **False** – a normal form – we must clearly avoid trying to evaluate Ω .

7.1 Basic evaluation strategies

Evaluation strategies may be broadly classified by how one searches for the next **redex**, this is an occurrence of where a β -reduction step could be applied. Here are some basic strategies:

Innermost: Before evaluating a term (with a redex at the root) ensure that all its subterms have been evaluated.

Outermost: Reduce a term with a redex at the root before evaluating any of its subterms.

Leftmost: Reduce redexes to the left of the term before redexes on the right of the term: leftmost evaluation

Rightmost: Reduce redexes to the right of the term before redexes on the left of the term: rightmost evaluation.

A basic **by-value** strategy for evaluation is *rightmost innermost* : this means amongst the innermost redexes (those nearest the leaves of the λ -term) one always chooses the rightmost to reduce. Adopting this strategy means that, when there are still redexes to reduce, there is always exactly one redex which is the rightmost innermost redex. Clearly, however, a by-value reduction strategy will run into trouble when evaluating the term above!

A basic **by-name** evaluation strategy is a *leftmost outermost* strategy: this means that amongst the outermost redexes (i.e. those closest to the root) one should choose the leftmost redex. This, for the λ -calculus is called **normal order evaluation**. This strategy has the important property (which we prove below) that if a term has a normal form this strategy will find that normal form.

Evaluation strategies can also be distinguished by where they are designed to stop. The two strategies by-value and by-name only halt when a normal form is reached. However, often it is useful to stop much earlier and this gives another dimension to the rewriting:

Normal form: A rewriting strategy which only halts on a normal form is sometimes called a “strong” rewriting strategy.

Head normal form: A rewriting strategy designed to terminate on a head normal form is called a “head” rewriting strategy. Such a strategy repeatedly rewrites the head redex (see more below).

Weak head normal form: A rewriting strategy designed to terminate on a weak head' normal form is sometimes called a “weak” rewriting strategy. These strategies are the (basically) the ones usually used by programming languages.

A term is in *normal form* if it contains no redexes (that is places where a β -reduction can be performed). A term is in *head normal form* if it does not contain a **head redex**. A head redex is one of the form

$$\lambda x_1 \dots x_n. (\underline{\lambda} z. M) N N_1 \dots N_m$$

where $n, m \in \mathbb{N}$ (which means they can both be zero). Notice that a head redex is necessarily a leftmost outermost redex and that it is found by burrowing under the outer λ -abstractions and then following the “application chain” to left until the “head term” (which cannot be an application) is reached. A term is, thus in **head** normal form if when searching for the head redex one encounters as the head term a variable. This means a term in head normal form looks like:

$$\lambda x_1 \dots x_n. x_r N_1 \dots N_m$$

Notice that if a term is in head normal form any subsequent rewritings will have to take place inside the N_i and, thus, will not effect the “head structure” of the term:

$$\lambda x_1 \dots x_n. x_r ? \dots ?$$

and this means that one can publish this head structure as an intermediate result. This is important in modelling potentially infinite computations which can give partial answers as they develop.

A term is in *weak head normal form* if it is a λ -abstraction or it is head normal form. This means that a weak head normal form has one of the two forms:

$$\lambda x. N \quad \text{or} \quad x N_1 N_2 \dots N_m$$

As before we may characterize a weak head normal form as a term which does not have a **weak head redex**: this is a redex of the form:

$$(\underline{\lambda} x. M) N N_1 N_2 \dots N_m$$

where again m can be zero. Thus, in searching for a weak head normal form one simply follows the application chain to the left to find a redex (with no burrowing under λ -abstractions). Thus, unlike for a head normal form, one never looks inside λ -abstractions. This is also the basic strategy lazy functional languages like Haskell use for searching for the next rewrite, however, importantly, they also hold the term as an acyclic graph rather than a tree: this means that subterms are not duplicated on performing β -reductions but rather shared. This makes their rewriting strategies “optimal” in the sense that they perform the minimum number of β -reductions.

This may suggest that lazy rewriting is always the best option, however, this is not necessarily the case as there is an overhead associated with managing the acyclic graph and, in particular, in recording when a rewrite has been performed.

7.2 Standardization

Here we are mostly concerned with the theoretical properties of normal order reduction: namely that it always manages to find a normal form if there is one. To establish this we shall prove the standardization theorem in some detail. We shall follow a proof by Ryo Kashima, [6], which uses the notion of the “leftmost index” of a β reduction:

$$M \xrightarrow[\beta(n)]{} N$$

A β reduction has index n if its redex is the n redex when one counts redexes left to right (outer first) across the term. This index more precisely is defined by:

Definition 7.1 A leftmost indexed β -reduction $M \xrightarrow[\beta(n)]{} N$ is defined by:

[LT.1] $(\lambda x.A)B \xrightarrow[\beta(1)]{} A[B/x]$.

[LT.2] If $A \xrightarrow[\beta(n)]{} B$ then if A is not an abstraction $AP \xrightarrow[\beta(n)]{} BP$ while if A is an abstraction then $AP \xrightarrow[\beta(n+1)]{} BP$.

[LT.3] Let $r(M)$ denote the number of redexes in a term. If $A \xrightarrow[\beta(n)]{} B$ then if P is not an abstraction $PA \xrightarrow[\beta(r(P)+n)]{} PB$ while if P is an abstraction then $PA \xrightarrow[\beta(1+r(P)+n)]{} PB$.

[LT.4] If $A \xrightarrow[\beta(n)]{} B$ then $\lambda x.A \xrightarrow[\beta(n)]{} \lambda x.B$.

Clearly every reduction $M \xrightarrow[\beta]{} N$ has a leftmost index, so there is an $n \geq 1$ with $M \xrightarrow[\beta(n)]{} N$.

Definition 7.2

- (i) A reduction sequence $A_0 \xrightarrow[\beta(n_1)]{} A_1 \xrightarrow[\beta(n_2)]{} \dots \xrightarrow[\beta(n_r)]{} A_r$ is **standard** in case $n_1 \leq n_2 \leq \dots \leq n_r$.
- (ii) A reduction sequence as above is a **leftmost outermost reduction** or **normal order reduction** in case $n_1 = n_2 = \dots = n_r = 1$ that is $A_0 \xrightarrow[\beta(1)]{}^* A_r$.

This puts normal order reduction, as introduced earlier, on a more formal footing as the reduction strategy is precisely to rewrite the redexes of leftmost index 1 repeatedly.

Our aim is now to prove:

Theorem 7.3 (Standardization)

- (i) If $M \xrightarrow[\beta]{}^* N$ then there is a standard reduction sequence $M \xrightarrow[\beta(n_1)\dots\beta(n_r)]{} N$ (with $n_1 \leq \dots \leq n_r$).
- (ii) If $M \xrightarrow[\beta]{}^* N$ has N in normal form then there is a leftmost outermost reduction $M \xrightarrow[\beta(1)]{}^* N$.

We observe that the second part can be obtained fairly straightforwardly from the first part. In a standard reduction if one skips over a redex, that is one does a step

$$A_i \xrightarrow{\beta(n)} A_{i+1} \quad \text{for } i > 1,$$

then one is obliged to leave any skipped redex unreduced in the term for the remainder of the reduction. Thus, if there is such a step in a standard reduction the end result must still contain the redexes: so the result cannot be in normal form. Thus, any standard reduction which ends in a normal form must be a normal order reduction! This also means that, if there is a normal form, then the normal order reduction will find it.

To prove this theorem we introduce two relations (the first of which is familiar!):

Definition 7.4

(i) A **weak head reduction** is a reduction of the form

$$(\lambda x.A_0)A_1 \dots A_n \xrightarrow{h} A_0[A_1/x]A_2 \dots A_n$$

(where $n \geq 1$).

(ii) Define the **standard relation** on terms inductively as follows:

[S.1] If $L \xrightarrow{h^*} x$ is a weak head reduction to a variable then $L \xrightarrow{h^*(\cdot)} x$ is in the standard relation.

[S.2] If $L \xrightarrow{h^*} AB$ is a weak head reduction to an application and $A \xrightarrow{s_1} A'$ and $B \xrightarrow{s_2} B'$ are in the standard relation then $L \xrightarrow{h^*(s_1 \ s_2)} A'B'$ is also in the standard relation.

[S.3] If $L \xrightarrow{h^*} \lambda x.A$ is a weak head reduction to an abstraction and $A \xrightarrow{s} A'$ is in the standard relation then $L \xrightarrow{h^*(\lambda.s)} \lambda x.A'$ is in the standard relation.

We shall denote a standard relation by $L \xrightarrow{s} M$ where s is a structured expression built inductively by the rules above.

Weak head reductions are always standard reductions, in fact, they are normal order or leftmost outermost reductions: so $\xrightarrow{h^*} C \xrightarrow{\beta(1)^*} \cdot$. That $t_1 \rightarrow t_2$ is in the standard relation does not immediately mean there is a standard reduction as, for example, the order of doing the reductions of $h^*(s_1 \ s_2)$ is not specified ... that is while s_1 and s_2 may correspond to standard order reductions, we have not specified how they are performed when in parallel. It is not hard to see that one can order the reductions of terms which are in the standard relation to create a standard reduction:

Lemma 7.5 If $L \xrightarrow{s} B$ is in the standard relation then there is a standard reduction $L \xrightarrow{\bar{s}} B$.

PROOF: We do a structural induction on the form of the standard relation:

[ws.1] If $L \xrightarrow{h^*} x$ then this is a standard reduction of the form $A \xrightarrow{\beta(1)^*} x$.

[ws.2] If $L \xrightarrow{h^*(s_1, s_2)} A'B'$ is weakly standard relation then inductively we have standard reductions $A \xrightarrow{\tilde{s}_1} A'$ and $B \xrightarrow{\tilde{s}_2} B'$ but this gives a standard reduction

$$L \xrightarrow{\beta(1)^*} AB \xrightarrow{\tilde{s}_1} A'B \xrightarrow{\tilde{s}_2} A'B'.$$

Note that the indexes in \tilde{s}_2 must be shifted by $r(A')$ when A' is not an abstraction and $r(A') + 1$ when it is an abstraction.

[ws.3] If $L \xrightarrow{h^*(\lambda.s)} \lambda x.A'$ is in the standard relation then inductively there is a corresponding standard reduction $A \xrightarrow{\tilde{s}} A'$ and then $L \xrightarrow{\beta(1)^*} \lambda xA \xrightarrow{\lambda.\tilde{s}} \lambda x.A'$ is standard. □

We now develop some properties of weak head reductions and the standard relation:

Lemma 7.6

- (i) If $M \xrightarrow{h^*} N$ is a weak head reduction then $M P \xrightarrow{h^*P} N P$ is a weak head reduction.
- (ii) If $M \xrightarrow{h^*} N$ is a weak head reduction then $M[P/x] \xrightarrow{h^*} N[P/x]$ is a weak head reduction.
- (iii) $M \rightarrow M$ is in the standard relation for each term M .
- (iv) If $L \xrightarrow{h^*} M$ is a weak head reduction and $M \xrightarrow{s} N$ is in the standard relation then $L \xrightarrow{h^*(s)} N$ is in the standard relation.
- (v) If $M \rightarrow N$ and $P \rightarrow Q$ are in the standard relation then $M[P/x] \rightarrow N[Q/x]$ is in the standard relation.

PROOF:

- (i) This is immediate!
- (ii) Here $M = (\lambda y.A_0)A_1 \dots A_n \xrightarrow{h} A_0[A_1/y]A_2 \dots A_n = N$ substituting P for x gives:
$$M[P/x] = (\lambda y.A_0[P/x])A_1[P/x] \dots A_n[P/x] \xrightarrow{h} A_0[P/x][A_1[P/x]/y]A_2[P/x] \dots A_n[P/x] = N[P/x]$$
- (iii) We may induce on the structure of the λ -term. If it is a variable then, using [S.1] we certainly have $x \xrightarrow{h^*} x$ whence $x \rightarrow x$ is in the standard relation. If the term is an application AB then by structural induction we know $A \xrightarrow{\varepsilon} A$ and $B \xrightarrow{\varepsilon} B$ are in the standard relation so we have $AB \xrightarrow{h^*} AB \xrightarrow{(\varepsilon, \varepsilon)} AB$ in the standard relation. The argument is similarly when the term is a λ -abstraction.
- (iv) This is immediate as all the standard relations start with an arbitrary weak head reduction, so one can just append the head reductions.

(v) We structurally induce on the structure of $M \xrightarrow[s]{}$ N :

[S.1] If $M \xrightarrow[h^*]{}$ $y = N$ then, when $x \neq y$, then the substitution has no effect and, as the head reduction by itself is in the standard relation (by (iii)), the result follows. The case remaining is when P is substituted into x but in that case we can use (iv) to get the result.

[S.2] If $M \xrightarrow[h^*]{}$ $AB \xrightarrow[(s_1, s_2)]{}$ $A'B' = N$ then we may use the structural assumption that we have $A[P/x] \xrightarrow[s_a]{}$ $A'[Q/x]$ and $B[P/x] \xrightarrow[s_b]{}$ $B'[Q/x]$ this gives

$$M[P/x] \xrightarrow[h^*]{}$$

$$(AB)[P/x] = A[P/x]B[P/x] \xrightarrow[(s_a, s_b)]{}$$

$$A'[Q/x] B'[Q/x] = N[Q/x]$$

as desired.

[S.3] If $M \xrightarrow[h^*]{}$ $\lambda y.A \xrightarrow[\lambda.s]{}$ $\lambda y.A'$ then substituting and using the induction hypothesis on $A \xrightarrow[s]{}$ A' we have:

$$M[P/x] \xrightarrow[h^*]{}$$

$$(\lambda y.A)[P/x] = \lambda y.(A[P/x]) \xrightarrow[\lambda.s_a]{}$$

$$\lambda y.(A'[Q/x]) = (\lambda y.A')[Q/x]$$

This completes the proof of (v). □

Finally to complete the proof of the standardization theorem we have:

Lemma 7.7 *If $L \rightarrow M$ is in the standard relation and $M \xrightarrow[\beta]{}$ N then $L \rightarrow N$ is in the standard relation.*

PROOF: We do a structural induction on the form of the standard relation:

[S.1] If M is a variable there cannot be any subsequent β -reduction and we are done.

[S.2] If $L \xrightarrow[h^*(s_1, s_2)]{}$ $A B$ then there are two cases to consider: if the β reduction is at the first leftmost redex (so that A is an abstraction) or if the index is greater than one. In the latter case the β -reduction falls inside either A or B and we can use our inductive hypothesis on $s_i\beta$ so we are done. In the former case we have:

$$L \xrightarrow[h^*]{}$$

$$(\lambda x.B)C \xrightarrow[h]{\beta(1)} B[C/x]$$

$$\lambda.s_1, s_2 \downarrow \qquad \qquad \qquad \downarrow (s_1[s_2/x])$$

$$(\lambda x.B')C' \xrightarrow[\beta(1)]{}$$

$$B'[C'/x]$$

where the square is Lemma 7.6 (v).

[S.3] $L \xrightarrow[h^*(\lambda s)]{}$ $\lambda x.A$ then the subsequent β -reduction is below the lambda abstraction and we can use our inductive hypothesis. □

This lemma puts – by a straight forward induction on the length of the reduction – all reduction sequences in the standard relation and so, by Lemma 7.5, can be replaced by a standard reduction.

8 Combinatory algebra

One may think that the λ -calculus is a remarkably simple system for expressing computability but it turns out there is an even simpler system!! This is often called “combinatory logic” but it is a very simple algebraic system: it has a binary operation, called application, and two constants \mathbf{k} and \mathbf{s} . These are subject to the following equations which we orient as rewrites:

$$\begin{aligned} (\mathbf{k} \bullet x) \bullet y &\rightarrow x \\ ((\mathbf{s} \bullet x) \bullet y) \bullet z &\rightarrow ((x \bullet z) \bullet (y \bullet z)) \end{aligned}$$

These rewrite rules form an orthogonal left-linear rewriting system and thus are a confluent rewriting system. Furthermore standard reductions (leftmost outermost) are guaranteed to find normal forms. It is also a non-trivial system as $\mathbf{s} \bullet \mathbf{k} \neq \mathbf{k}$. The system, in fact, satisfies some quite remarkable properties which we shall briefly explore below.

This remarkable system was invented by Moses Schönfinkel, a Russian mathematician who was born in 1889 and worked in Hilbert’s group at the University of Göttingen from 1914-1924. He presented these ideas on the 7th of December in 1920. Schönfinkel returned to Moscow in the mid 1920s where he fell in ill and eventually died in poverty in 1942. It is said that, after his death, his neighbours (suffering from harsh war-time conditions) burned his papers to heat their homes!

To access the properties of a combinatory algebra it is first useful to consider an **applicative system**, that is is a set A with a single binary operation

$$_ \bullet _ : A \times A \rightarrow A; (x, y) \mapsto x \bullet y.$$

Clearly a combinatory algebra is an example of a applicative system as one can simply forget about \mathbf{k} and \mathbf{s} . An applicative system is said to be functionally complete if whenever there is a polynomial expression $p(x_1, \dots, x_n)$ in n variable

$$p \rightarrow x_1|x_2|\dots|x_n|a \in A|p \bullet p$$

then there is an element \widehat{p} such that for every substitution of the variables by elements of A

$$(\dots((\widehat{p} \bullet x_1) \bullet \dots) \bullet x_n = p(x_1, \dots, x_n)$$

Proposition 8.1 *Combinatory algebra is functionally complete. Furthermore, a combinatory complete applicative system is a combinatory algebra*

PROOF: The proof involves the introduction of an abstraction mechanism which is a mechanism for building constants with the desired property for functional completeness. Suppose, therefore, we have a polynomial p which involves free variables $X = x_1, \dots, x_n$ and then define $\lambda^*x_i.p$ as the following polynomial in which the variable x_i has been removed:

$$\begin{aligned} \lambda^*x.x &= (\mathbf{s} \bullet \mathbf{k}) \bullet \mathbf{k} \\ \lambda^*x.M &= \mathbf{k} \bullet M \quad x \notin M \\ \lambda^*x.(p_1 \bullet p_2) &= (\mathbf{s} \bullet (\lambda^*x.p_1)) \bullet (\lambda^*x.p_2) \end{aligned}$$

Observe that $(\lambda * x.p) \bullet M = p[M/x]$ which can be seen by a structural induction:

$$\begin{aligned}
(\lambda^* x.x) \bullet M &= ((s \bullet k) \bullet k) \bullet M \\
&= (k) \bullet M \bullet (k) \bullet M = M = x[M/x] \\
(\lambda^* x.z) \bullet M &= (k \bullet z) \bullet M = z = z[M/x] \\
\lambda^* x.(p_1 \bullet p_2) \bullet M &= ((s \bullet (\lambda^* x.p_1)) \bullet (\lambda^* x.p_2)) \bullet M \\
&= ((\lambda^* x.p_1) \bullet M) \bullet ((\lambda^* x.p_2) \bullet M) \\
&= p_1[M/x] \bullet p_2[M/x] = (p_1 \bullet p_2)[M/x]
\end{aligned}$$

Also we observe that $(\lambda^* x.p)[M/y] = (\lambda^* x.p[M/y])$ when $x \notin FV(M)$ again a proof by structural induction is required which we leave to the reader. This now shows that we can take

$$\widehat{p} := (\lambda^* x_1.(\lambda^* x_2 \dots (\lambda^* x_n.p) \dots))$$

to obtain:

$$\begin{aligned}
&(((\lambda^* x_1.(\lambda^* x_2 \dots (\lambda^* x_n.p) \dots)) \bullet M_1) \bullet M_2) \bullet \dots \bullet M_n \\
&= (((\lambda^* x_2.(\dots (\lambda^* x_n.p) \dots)[M_1/x_1]) \bullet M_2) \bullet \dots) \bullet M_n \\
&= ((\lambda^* x_2.(\dots (\lambda^* x_n.p[M_1/x_1]) \dots)) \bullet M_2) \bullet \dots \bullet M_n \\
&= p[M_1/x_1, \dots, M_n/x_n]
\end{aligned}$$

For the converse it suffices to show that a combinatory algebra which is functionally complete has a k and an s . However, the equations for these combinators are functional completeness equations so this must be so! \square

These observations suggest that in combinatory algebra, following the λ -calculus

$$\Omega = (\lambda^* x.x \bullet x) \bullet (\lambda^* x.x \bullet x)$$

should have a non-terminating reduction. This is easily checked to be the case. However, it should not be imagined that the two systems are equivalent. The following equality implication

$$N = M \not\Rightarrow \lambda^* x.M = \lambda^* x.N$$

is true in the λ -calculus but is definitely not true in combinatory algebra. Thus $(k \bullet x) \bullet y = x$ but

$$\begin{aligned}
\lambda^* x.(k \bullet x) \bullet y &= (s \bullet (\lambda^* x.k \bullet x)) \bullet (\lambda^* x.y) \\
&= (s \bullet ((s \bullet (\lambda^* x.k)) \bullet (\lambda^* x.x))) \bullet (\lambda^* x.y) \\
&= (s \bullet ((s \bullet (k \bullet k)) \bullet ((s \bullet k) \bullet k))) \bullet (k \bullet y) \\
&\neq (s \bullet k) \bullet k = \lambda^* x.x
\end{aligned}$$

So that combinatory logic is much weaker. Combinatory logic is important as a system as it one of the simplest systems in which all (partial) computable functions can be simulated. The encoding technique follows the techniques of the λ -calculus.

9 The typed λ -calculus

The typed λ -calculus, as opposed to the “untyped” λ -calculus – which is what we described above – has some rather striking and important properties. Notably, the pure typed λ -calculus has a word problem which is decidable and a rewriting system which is terminating. Once one adds datatypes one obtains a very expressive system which is still terminating but can be viewed as a basic “strong” (in the sense of terminating) programming language.

9.1 The simply typed λ -calculus

We shall start with a fragment which has both λ -abstraction and pairing which assumes a set of atomic types, \mathcal{A} , from which one can build all the types using the rules:

$$\boxed{\begin{array}{c} \frac{A \in \mathcal{A}}{A \text{ type}} \qquad \frac{}{1 \text{ type}} \\ \frac{X \text{ type } Y \text{ type}}{X \times Y \text{ type}} \qquad \frac{X \text{ type } Y \text{ type}}{X \rightarrow Y \text{ type}} \end{array}}$$

One can then describe the term judgements for the system as follows:

$$\boxed{\begin{array}{c} \frac{A \in \mathcal{A}}{x : A, \Gamma \vdash x : A} \text{ projection} \\ \frac{}{\Gamma \vdash () : 1} \text{ empty tuple} \qquad \frac{\Gamma \vdash t_1 : T_1 \quad \Gamma \vdash t_2 : T_2}{\Gamma \vdash (t_1, t_2) : T_1 \times T_2} \text{ pairing} \\ \frac{\Gamma \vdash t : T_0 \times T_1}{\Gamma \vdash \pi_0 t : T_0} \text{ first} \qquad \frac{\Gamma \vdash t : T_0 \times T_1}{\Gamma \vdash \pi_1 t : T_1} \text{ second} \\ \frac{x : A, \Gamma \vdash t : B}{\Gamma \vdash \lambda x.t : A \rightarrow B} \text{ abstraction} \qquad \frac{\Gamma \vdash t : A \rightarrow B \quad \Gamma \vdash s : A}{\Gamma \vdash t s : B} \text{ application} \end{array}}$$

There are then just three equalities which can all be viewed as rewrites:

$$\begin{aligned} \pi_0(t_0, t_1) &\rightarrow t_0 \\ \pi_1(t_0, t_1) &\rightarrow t_1 \\ (\lambda x.t)s &\rightarrow t[s/x] \end{aligned}$$

Proposition 9.1 *These rewrites form a confluent system.*

We sketch a proof which uses a marking argument which extends the rewriting of the λ -calculus with projections for pairs. The argument in fact shows that adding explicit pairs to the untyped λ -calculus results in a confluent system.

We start by introducing marked terms $\pi_i(t_0, t_1)$ and, as before, note that there is a by-value evaluation of marked terms:

$$\begin{aligned}
V(x) &= x && \text{for } x \text{ a variable} \\
V(\pi_1) &= \pi_i && \text{for unmarked projections} \\
V(\lambda x.N) &= \lambda x.(V(N)) \\
V(NM) &= V(N)V(M) \\
V((N, M)) &= (V(N), V(M)) \\
V(\pi_i(M_0, M_1)) &= V(M_i) \\
V((\lambda x.M)N) &= V(M)[V(N)/x]
\end{aligned}$$

It then suffices, for the marking argument to go through, to show that

$$\begin{array}{ccc}
t_0 & \longrightarrow & t_1 \\
* \downarrow & & \downarrow * \\
V(t_0) & \dashrightarrow & V(t_1)
\end{array}$$

As before we may consider the single step rewriting initially to be at the root and then to be in context. In the first case there are two cases: when the single rewrite is marked and when it is not. If it is marked it is easily seen that $V(t_0) = V(t_1)$ We shall reconsider the case when there is an unmarked projection or β -reduction at the root:

$$\begin{array}{ccc}
\pi_0(t_1, t_2) & \longrightarrow & t_2 \\
* \downarrow & & \downarrow * \\
\pi_0(V(t_1), V(t_2)) & \longrightarrow & V(t_2)
\end{array}
\quad
\begin{array}{ccc}
(\lambda x.M)N & \longrightarrow & M[n/x] \\
* \downarrow & & \downarrow * \\
(\lambda x.V(M))V(N) & \longrightarrow & V(M)[V(N)/x] = V(M[N/x])
\end{array}$$

These diagram show these one-step diagrams work provided we have the substitution equality (which we leave to the reader). For a rewriting in context we must show:

$$\begin{array}{ccc}
C[[N]] & \longrightarrow & C[[M]] \\
* \downarrow & & \downarrow * \\
V(C[[N]]) & \dashrightarrow & V(C[[M]])
\end{array}$$

and again we have to work by structural induction on the context. It is not hard to show, as before, the useful fact that $V(C[[t]]) = V(C[[V(t)])$. Now if there is no marked material in the context the result is immediate as no rewriting of the context happens. This means there are two important inductive cases corresponding to building the context with a marked term. A marked projection redex or a marked β -reduction. For a marked projection redex the context must be in one of the terms of the pair: for example $\pi_i(t_0, C[[N]])$. When the projection is to the coordinate which does not contain the context his reduces to $V(t_0)$ and the rewrite step $N \rightarrow M$ is completely removed. So the remaining case is when the projection is to the coordinate of the context: but in this case

we have:

$$\begin{array}{ccc}
\pi_1(t_0, C[[N]]) & \longrightarrow & \pi_i(t_0, C[[M]]) \\
\downarrow * & \text{ind} & \downarrow * \\
\pi_1(V(t_0), V(C[[V(N)]))) & \xrightarrow{*} & \pi_i(V(t_0), V(C[[V(M)]))) \\
\downarrow & & \downarrow \\
V(C[[V(N)]]) & \xrightarrow{*} & V(C[[V(M)]])
\end{array}$$

The argument for a β -reduction is then similar to the argument in the untyped λ -calculus.

9.2 Normalization

We shall first prove that in the simply typed λ -calculus that each term has a normal form. This is called the normalization theorem, but note that it presumes a prescribed order of reduction in order to reach the normal form. Thus, having normalization does not rule out the possibility that, using some strange rewriting order, one could avoid ever reach a normal form. If it is the case that no matter which rewriting order one uses one is *guaranteed* to reach a normal form then the system would be said to be *strongly* normalizing. In fact, the simply typed λ -calculus *is* strongly normalizing – and we shall return to prove this shortly. Meanwhile we will prove normalization.

We prove normalization by showing that there is a series of rewrites which is guaranteed to terminate and, thus, end in a normal form. The technical part of the argument uses a termination argument based on well-founded relations. We start therefore by reviewing well-founded relations.

9.2.1 Well-founded relations and termination

A rewriting system may be viewed as a relation on a set (usually a set of terms): it is **terminating** if the rewriting relation is *well-founded*. A relation $\rightarrow \subseteq A \times A$ is **well-founded** in case every non-empty set has a minimal element: that is for each subset $S \subseteq A$ there is an $s \in S$ such that there is no $s' \in S$ with $s \rightarrow s'$. This means, equivalently, that in any subset S there are no infinitely descending chains of the form

$$t \rightarrow t_1 \rightarrow t_2 \rightarrow t_3 \rightarrow \dots$$

as such a chain would not have a minimal element and conversely given an S with no minimal element one would be able to choose an infinitely descending chain.

The β -reduction relation on λ -terms is *not* a well-founded relation: a minimal element of this reduction relation is exactly a normal form and we know there are terms that lack a normal form and have an infinite reduction sequence. For example:

$$\Omega \xrightarrow{\beta} \Omega \xrightarrow{\beta} \Omega \xrightarrow{\beta} \dots$$

We start by observing:

Lemma 9.2 *If a relation \rightarrow is well-founded its transitive closure, $\xrightarrow{+}$, is also well-founded.*

PROOF: Suppose S is a non-empty subset we must find a minimal element in the set with respect to $\xrightarrow{+}$. Notice that this minimal element remains minimal when we up-close the set

$$\uparrow S = \{t \mid t \xrightarrow{+} t', t' \in S\}.$$

Now this set has a minimal element, s_0 , with respect to the original relation, as it is, by assumption, well-founded.

The claim is that s_0 must also be minimal with respect to $\xrightarrow{+}$ in the up-closure. Suppose for contradiction that $s_0 \xrightarrow{+} t$ and $t \in \uparrow S$ then either $s_0 \rightarrow t$ – which cannot be as s_0 is minimal – or $s_0 \xrightarrow{+} t' \rightarrow t$ but then $t' \in \uparrow S$ and, thus, s_0 is not minimal with respect to \rightarrow . \square

In a rewriting system, in which the number of possible rewrites leaving any term is always finite (i.e. each term contains a finite number of redexes), when the rewrite relation is well founded then there is for each term t a bound $\beta(t) \in \mathbb{N}$ on the number of rewrites which a chain leaving that term can have. Note that having a bound is a strictly stronger property than being well-founded. However, in many terminating rewrite systems, finding an explicit bound (that is a natural number) on the length of any chain of rewrites leaving a term can be quite straightforward.

Lemma 9.3 *A rewriting system \mathcal{R} is terminating if and only if there is a well-founded set $(W, <)$ and a map $\alpha : T_\Omega(X) \rightarrow W$ such that:*

- $\alpha(r_i) > \alpha(c_i)$ for each $r_i \rightarrow c_i \in \mathcal{R}$;
- If $\alpha(t) > \alpha(t')$ then $\alpha(t[\sigma]) > \alpha(t'[\sigma])$;
- If $\alpha(t_i) > \alpha(t'_i)$ then $\alpha(f(t_1, \dots, t_i, \dots, t_n)) > \alpha(f(t_1, \dots, t'_i, \dots, t_n))$.

PROOF: If the rewriting system is terminating then $t \rightarrow t'$ is well founded so we take α to be the identity map. Conversely given such an α suppose $\emptyset \neq S \subseteq T_\Omega(X)$ then the set $\{\alpha(s) \mid s \in S\} \subseteq W$ is non-empty and therefore has a minimal element $\alpha(t_0)$. However if $t_0 \rightarrow t'_0$ is a reduction step which has $t'_0 \in S$ then $\alpha(t'_0) < \alpha(t_0)$ contradicting the minimality of $\alpha(t_0)$. Thus, t_0 is a minimal element in S with respect to the reduction order showing this order is well-founded. \square

Thus, to demonstrate that a rewrite system terminates it suffices to exhibit a map α satisfying the conditions of this lemma.

The canonical well-founded set is the natural numbers with the usual relation $n < m$.

If P is a set with a well-founded relation then $P^* = \text{List}(P)$ with the lexicographic relation is also well-founded. The lexicographic relation is given by setting $[]$ to be less than everything and $x : xs < y : ys$ if either $x < y$ or if $x = y$ and $xs < ys$ inductively.

Lemma 9.4 *The lexicographic relation on P^* derived from a well-founded relation on P is itself well-founded.*

PROOF: The lexicographic relation is well-founded as given any $X \subset P^*$ if the empty list is in there is automatically a minimal element otherwise choose an element $x_0 \in P$ which is minimal in the first entry and consider all the elements with the same first entry. We argue that this set has a minimal element. Consider the tails $\text{Tail}_{x_0}(X) = \{xs \mid x_0 : xs \in X\}$, if there is a minimal element

in this set then there is a minimal element in X . In fact, if we continue this process of removing minimal first elements, one of two things can happen: either the set reduces to a one element set – and then we have found a minimal element – or we reach a set which has the empty list and again we have found a minimal element. \square

Another important source of well-founded relations uses an induced relation on “bags” which was introduced by Dershowitz and Manna.

A bag is an unordered list: this means repetitions are allowed: $\{\{x, y, x\}\} = \{\{y, x, x\}\} \neq \{\{x, y\}\}$ but the order in which the elements occur does not matter. The set $\mathbf{Bag}(P)$ of bags of elements of a set P , with a well-founded relation, can itself be endowed with a well-founded relation. One defines one bag to be less than another in case:

- $\{\{x\}\} \sqcup b_1 < \{\{x\}\} \sqcup b_2$ if $b_1 < b_2$;
- $\{\{x_1, \dots, x_n\}\} \sqcup b_1 < \{\{y\}\} \sqcup b_2$ for $n \in \mathbb{N}$ whenever each $x_i < y$ and $b_1 = b_2$ or $b_1 < b_2$.

Notice that the empty bag $\{\{\}\}$ is minimal because given any bag we can repeatedly remove all elements using the second rule with $n = 0$

Thus we may determine whether one bag is (strictly) less than the other by first removing the elements in common and then for each element in the bigger bag removing all elements in the smaller bag which are strictly less than it. The smaller bag will be emptied by this process.

Intuitively to construct an element strictly smaller than a given b we are allowed to pick an element from the bag and either remove it or replace it with a bag of elements each of which is individually strictly smaller than the element removed. It is not so obvious that this process must always end as you may replace the element with a very large bag of smaller elements!

Proposition 9.5 *If P is a well founded set then the induced relation on $\mathbf{Bag}(P)$ given above is well founded.*

PROOF: To prove this is well-founded is not so easy (see Dershowitz and Manna). Here is a sketch: suppose there is an infinite descending chain of bags $b_0 < b_1 < b_2 < \dots$ then we may build a forest with roots $x \in b_0$ the children of x are the bag $\{\{x_1, \dots, x_n\}\}$ of elements which eventually replace x that is $b_i < b_{i+1}$ (for some i) introduces this strict replacement of x (if this never happens then x is a leaf). Every step in the sequence must do one (or more) replacements of this nature. However, this tree is finitely branching and has all its paths of finite length so itself is finite (König’s lemma). This means the sequence itself must be finite. \square

9.2.2 The argument for weak normalization

The ideas in this section are a slight modification of the ideas in [5].

To show weak normalization we shall use the following measures:

Type height: The height of a type $\mathbf{hgt}(T)$ is defined by:

$$\begin{aligned} \mathbf{hgt}(1) = \mathbf{hgt}(A) &= 1 \quad \text{for } A \text{ atomic (that is } A \in \mathcal{A}) \\ \mathbf{hgt}(T_1 \times T_2) = \mathbf{hgt}(T_1 \rightarrow T_2) &= \max(\mathbf{hgt}(T_1), \mathbf{hgt}(T_2)) + 1. \end{aligned}$$

Redex weight: The redex weight of a term is zero unless it starts with a redex:

$$\begin{aligned} \text{red}(\pi_i(t_1, t_2)) &= \text{hgt}(T_1 \times T_2) \text{ where } (t_1, t_2) : T_1 \times T_2 \\ \text{red}((\lambda x.t)s) &= \text{hgt}(T \rightarrow S) \text{ where } (\lambda x.t) : T \rightarrow S \\ \text{red}(t) &= 0 \text{ otherwise} \end{aligned}$$

Finally the bag of redex weights of a term is:

$$\text{rbag}(t) = \{\text{red}(s) \mid s \ll t \ \& \ \text{red}(s) \neq 0\}$$

where $s \ll t$ means s is a subterm of t . We observe:

Lemma 9.6

- (i) Suppose $(\lambda x.t)s \xrightarrow{\beta} t[s/x]$ then provided the weight of this (top) redex is greater than any redex contained in s then $\text{rbag}((\lambda x.t)s) > \text{rbag}(t[s/x])$
- (ii) Suppose $\pi_i(t_0, t_1) \rightarrow t_i$ then $\text{rbag}(\pi_i(t_0, t_1)) > \text{rbag}(t_i)$.
- (iii) Suppose $C[[r]] \rightarrow C[[c]]$ where $\text{rbag}(r) > \text{rbag}(c)$ then $\text{rbag}(C[[r]]) > \text{rbag}(C[[c]])$.

PROOF:

- (i) Suppose $(\lambda x.t)s \xrightarrow{\beta} t[s/x]$ then the redexes in $t[s/x]$ are:
 - redexes of t with s substituted for x , which have the same weight as those in t so are dominated;
 - redexes of s : these may be proliferated when s is substituted for x in t , however, their weight is less than the weight of the top redex by assumption so they are dominated by that weight.
 - new redexes created by the substitution of the form st' . The weight of these redexes is also dominated by the weight of the top redex as if $s : U$ then $\lambda x.t$ must have type $U \rightarrow V$ for some type V and $\text{hgt}(U) < \text{hgt}(U \rightarrow V)$.
- (ii) Suppose $\pi_i(t_0, t_1) \rightarrow t_i$ then $\text{rbag}(\pi_i(t_0, t_1)) > \text{rbag}(t_i)$ as the weights of the redexes of the term eliminated together with that of the top redex are removed.
- (iii) This is also immediate as weights of the redexes in the context are unchanged.

□

Whenever we have a redex such that $\text{red}(t) > \text{red}(s)$ for all $s \neq t$ and $s \ll t$ then the lemma tells us that by reducing that redex we can descend in the bag ordering. As we can always find such a redex – any redex with no redexes below it will do – we can prove that the term must eventually reduce a normal form if one repeatedly reduces such a redex. In particular, note that the by-value reduction strategy *always* has this form. This means every term has a normal form and this means:

Corollary 9.7 *In the typed λ -calculus (with products) both by-value and normal-order reduction always terminate.*

Normal order reduction terminates because, if there is a normal form it will be found by a normal order reduction! (Although we have not shown this for products!)

While knowing that terms have a normal form is important even more important is to know that whatever order the reduction is done in that it will always terminate. This, however, requires a more sophisticated argument.

9.2.3 The argument for strong normalization

To prove strong normalization we shall employ a much stronger argument introduced by Tait and modified by Girard. The idea of the proof is to introduce for each type a family of terms called the *reducible* terms. These terms are then proven to satisfy three conditions – called by Girard the “candidats de reductibilité”. Finally it is shown by induction that all terms are reducible.

We define $\text{RED}(T)$ for a type T to be the set of terms defined by:

- [Red.1] For an atomic type $A \in \mathcal{A}$, $t \in \text{RED}(A)$ if and only if it is strongly normalizing;
- [Red.2] For a product type, $T \times S$, $t \in \text{RED}(T \times S)$ if and only if $\pi_0(t) \in \text{RED}(T)$ and $\pi_1(t) \in \text{RED}(S)$;
- [Red.3] For an arrow type $T \rightarrow S$, $u \in \text{RED}(T \rightarrow S)$ if and only if for every $t \in \text{RED}(T)$, $ut \in \text{RED}(S)$.

Say that a term is *neutral* if it is not a pair, $\langle u, v \rangle$, or an abstraction, $\lambda x.t$. Thus, a neutral term is either a variable, x , a projection, $\pi_0(t)$ or $\pi_1(t)$, or an application, $t s$. We shall now show that the reducible terms satisfy the following three conditions:

- [CR.1] If $t \in \text{RED}(T)$ then t is strongly normalizable;
- [CR.2] If $t \in \text{RED}(T)$ and $t \xrightarrow{*} t'$ then $t' \in \text{RED}(T)$;
- [CR.3] If t is neutral and, for every reduction step $t \rightarrow t'$, $t' \in \text{RED}(T)$ then $t \in \text{RED}(T)$.

In particular, importantly, [CR.3] implies that a neutral term, t , which is in normal form is always reducible.

Proposition 9.8 *Reducible terms satisfy [CR.1]-[CR.2].*

PROOF: We work by structural induction:

Atomic types: [CR.1] and [CR.2] are obviously true. For [CR.3] notice that if every one step reduction of t takes one to a strongly normalizing term t' then the term t is certainly strongly normalizing and so in $\text{RED}(A)$.

Product types: A term of product type is reducible if its projections are.

- [**CR.1**] If $t \in \text{RED}(U \times V)$ then $\pi_0(t) \in \text{RED}(U)$ and $\pi_1(t) \in \text{RED}(V)$ by assumption [**CR.1**] hold for U and V so both $\pi_0(t)$ and $\pi_1(t)$ are strongly normalizing. Thus, t is strongly normalizing.
- [**CR.2**] If $t \in \text{RED}(U \times V)$ and $t \xrightarrow{*} t'$ then $\pi_0(t) \xrightarrow{*} \pi_0(t')$ and $\pi_1(t) \xrightarrow{*} \pi_1(t')$ but by assumption [**CR.2**] holds for U and V so that $\pi_0(t') \in \text{RED}(U)$ and $\pi_1(t') \in \text{RED}(V)$ and so $t' \in \text{RED}(U \times V)$.
- [**CR.3**] Now let t be neutral and such that for every one step reduction $t \rightarrow t'$ we have $t' \in \text{Red}(U \times V)$ then $\pi_0(t') \in \text{RED}(U)$ and $\pi_1(t') \in \text{RED}(V)$. We must show $\pi_0(t)$ and $\pi_1(t)$ are reducible. However, consider a one step reduction of these terms, as t is neutral such a step must be to $\pi_0(t')$ or $\pi_1(t')$ and these are reducible. Now using the inductive assumption this means $\pi_0(t)$ and $\pi_1(t)$ are reducible and thus that, as required, t is reducible.

Arrow types: Recall a term of arrow type is reducible if applied to any reducible term it is reducible.

- [**CR.1**] Suppose $t \in \text{RED}(U \rightarrow V)$ this means, letting $u \in \text{RED}(U)$ be a reducible term (there is always such as variables are terms and are always reducible) that $tu \in \text{RED}(V)$ but this means tu is strongly normalizing by assumption. This in turn means t must be strongly normalizing.
- [**CR.2**] Suppose $t \xrightarrow{*} t'$ and $t \in \text{RED}(U \rightarrow V)$ then for any $u \in \text{RED}(U)$ we have that $tu \in \text{RED}(V)$ but $tu \xrightarrow{*} t'u$ and by assumption on V this means $t'u$ is reducible. Thus, $t' \in \text{RED}(U \rightarrow V)$.
- [**CR.3**] Not let t be a neutral term and suppose for any $t \rightarrow t'$ that $t' \in \text{RED}(U \rightarrow V)$. Taking any $u \in \text{RED}(U)$ consider $t u$: a one step reduction, as t is neutral, is either of the form $t u \rightarrow t' u$ or $t u \rightarrow t u'$. The former is by assumption reducible the latter case needs some more work. By [**CR.1**] we know that u is strongly normalizing so we can argue for the reducibility of $t u$ by induction on the maximum reduction length for u . If it is 0, that is u is in normal form, we are done as either the reduction is in t or $t u$ has no reductions: but then as this of type V which satisfies [**CR.3**] it must be reducible. Suppose now we have the result for terms of maximum reduction length N and u has maximal reduction length $N + 1$ then applying the one step reduction $t u \rightarrow t u'$ produces the now by hypothesis reducible term $t u'$.

□

Now we proceed to prove that all terms are reducible by proving a slightly stronger statement which allows for the smooth handling of abstractions:

Proposition 9.9 *Any term t with its variables substituted by reducible terms is reducible.*

PROOF: The proof is by structural induction on terms:

Variables: A variable is a reducible term: substituting by a reducible term also produces a reducible term.

Projections: Consider $\pi_0(t)$ (the argument for $\pi_1(t)$ is similar) by inductive assumption $t[u_1/x_1, \dots, u_n/x_n]$ is reducible for any reducible terms u_1, \dots, u_n . However, for $t[u_1/x_1, \dots, u_n/x_n]$ to be reducible it is necessary for its projections to be reducible. Thus, in particular $\pi_0(t[u_1/x_1, \dots, u_n/x_n])$ is reducible.

Pairs: By assumption $\langle t_0, t_1 \rangle$ has t_0 and t_1 satisfying the hypothesis so $t_0[\tilde{u}/\tilde{x}]$ and $t_1[\tilde{u}/\tilde{x}]$ are reducible. To show $\langle t_0, t_1 \rangle[\tilde{u}/\tilde{x}] = \langle t_0[\tilde{u}/\tilde{x}], t_1[\tilde{u}/\tilde{x}] \rangle$ is reducible we must show that $\pi_0(\langle t_0, t_1 \rangle[\tilde{u}/\tilde{x}])$ and $\pi_1(\langle t_0, t_1 \rangle[\tilde{u}/\tilde{x}])$ are reducible. We shall argue for the former. It suffices to consider one step reductions: clearly the step $\pi_0(\langle t_0, t_1 \rangle[\tilde{u}/\tilde{x}]) \rightarrow t_0[\tilde{u}/\tilde{x}]$ gives a reducible term. We must consider reductions within the pair and here we argue on the (maximal) reduction length.

The base case is when the pair is in normal form: here the projection applies and the result follows. The inductive case allows a reduction of the pair to a case in which all one step reductions are reducible and so, by **[CR3]**, it is also reducible!

Applications: Consider $t s$: by assumption $t[\tilde{u}/\tilde{x}]$ and $s[\tilde{u}/\tilde{x}]$ are reducible terms. But $t[\tilde{u}/\tilde{x}]$ is reducible only if applying it to all reducible terms yields a reducible term. Thus, $t[\tilde{u}/\tilde{x}] s[\tilde{u}/\tilde{x}]$ is certainly reducible.

Abstractions: Consider $\lambda x.t$ where, by assumption $t[v/x, \tilde{u}/\tilde{x}]$ is reducible for any substitution by reducible terms. But then the one step reduction $(\lambda x.t)[\tilde{u}/\tilde{x}]v \rightarrow t[v/x, \tilde{u}/\tilde{x}]$ produces a reducible term. The other one step reductions must be internal to $t[\tilde{u}/\tilde{x}]$ or to v . Both are reducible terms so one can argue by induction on the sum of their reduction length.

In the base case (both are in normal form) and one must perform the β -reduction. For the inductive case the term reduces to a term which by **[CR.3]** is reducible. Thus, it is itself reducible.

□

Finally, by choosing the reducible terms to be variables, we then conclude that all terms are reducible. As reducible terms are strongly normalizing, by **[CR.1]**, we obtain the result we sought:

Theorem 9.10 *All terms in the typed λ -calculus with products are strongly normalizing.*

9.3 Higher-order primitive recursive functions

The basic system of the typed λ -calculus with products that we have considered so far is not very expressive. To correct this we shall start adding (inductive) data types. The strategy for strong termination that we shall outline works for *all* inductive data types. However, we shall illustrate the techniques on adding the Booleans and the natural numbers: this gets us to essentially Gödel's system \mathbb{T} which is extremely expressive as it allows all higher-order primitive recursive functions. These are also the functions which are provably total in Peano arithmetic.

We introduce new terms and types into the system:

- Bool and Nat are new atomic types.
- There are three new terms involving Bool:

$$\overline{\Gamma \vdash \text{True} : \text{Bool}} \quad \overline{\Gamma \vdash \text{False} : \text{Bool}}$$

$$\frac{\Gamma \vdash x : T \quad \Gamma \vdash y : T \quad \Gamma \vdash b : \text{Bool}}{\Gamma \vdash \text{FoldBool } x \ y \ b : T}$$

The fold can also be seen as the conditional (or case) for this non-recursive data type: $\text{FoldBool } x \ y \ b := \text{If } b \ x \ y$.

- There are three new terms involving Nat:

$$\frac{\Gamma \vdash \text{Zero} : \text{Nat} \quad \Gamma \vdash \text{Succ} : \text{Nat} \rightarrow \text{Nat} \quad \Gamma \vdash v : T \quad \Gamma \vdash f : T \rightarrow T \quad \Gamma \vdash n : \text{Nat}}{\Gamma \vdash \text{FoldNat } v \ f \ n : T}$$

The fold can be viewed to be an “iterator”. Gödel’s original system had a “recursor” which implemented primitive recursion. The reason for this is that (as discussed above) implementing the “case” construct with just an iterator is rather inefficient (and, in fact, this inefficiency cannot be avoided in a non-higher-order system). However, the two systems are equally expressive and we shall be concerned (here) with only with expressivity (i.e. what functions $f : \text{Nat} \rightarrow \text{Nat}$ can be expressed) and in this regard the systems are equivalent.

These new terms come with the following rewrites:

$$\begin{aligned} \text{FoldBool } x \ y \ \text{True} &\rightarrow x \\ \text{FoldBool } x \ y \ \text{False} &\rightarrow y \\ \text{FoldNat } v \ f \ \text{Zero} &\rightarrow v \\ \text{FoldNat } v \ f \ (\text{Succ } n) &\rightarrow f (\text{FoldNat } v \ f \ n) \end{aligned}$$

It is not hard to see that adding these rewrites to the basic rewrites of the typed λ -calculus with products gives a confluent system.

Theorem 9.11 *The typed λ -calculus with products, Booleans, and natural numbers with rewrites above gives a strongly normalizing system.*

PROOF: (Sketch)

Amazingly the same proof works!!

We extend the notion of neutrality to be any term which is not of the form $\langle u, v \rangle$, $\lambda x.t$, True , False , Zero , or $\text{Succ } n$. Then we have the notion of a reducible term of a type as before and that [CR.1]-[CR.3] still hold.

We then can prove by structural induction that all terms are reducible. Here is the case for FoldNat :

Suppose v , f , and n are reducible then all are strongly normalizing and so one can reason by induction on $|v| + |f| + |n| + \ell(n)$ where $|t|$ is the maximal reduction length of t and $\ell(n)$ is the size of the normal form of n . In a single reduction step $\text{FoldNat } v \ f \ n$ reduces to:

- $\text{FoldNat } v' \ f' \ n'$ which is reducible by induction;
- v if $n = \text{Zero}$, which is reducible;
- $f (\text{FoldNat } v \ f \ m)$ when $n = (\text{Succ } m)$. The induction hypothesis tells us that $\text{FoldNat } v \ f \ m$ is reducible as $\ell(n) = \ell(\text{Succ } m) > \ell(m)$. Thus as f is reducible, $f (\text{FoldNat } v \ f \ m)$ is reducible.

□

The proof that this system is strongly normalizing does not mean that we can determine equality of functions on the natural numbers! The point is that these rewrites allow the evaluation of functions on values but they do not allow the determination of the equality of functions. For example one cannot even prove that $x + y = y + z$, where $x + y := \text{FoldNat } x \text{ Succ } y$ as this latter is already in normal form. Thus, the system is extremely weak in terms of giving a notion of equality but it is, nonetheless, computationally very expressive!

References

- [1] *The λ -calculus: its syntax and semantics*. H. Barendregt, North Holland (1984)
- [2] *Introduction to Combinators and λ -calculus*. J.R. Hindley and J.P. Seldin, London Mathematical Society (1988)
- [3] *Lambda Calculus and Combinators: an introduction*. J.R. Hindley and J.P. Seldin, Cambridge University Press (2008)
- [4] *Term rewriting systems*. J.W. Klop in “Handbook of logic for Computer Science” (Vol. 2) Clarendon Press (1992)
- [5] *Proofs and types*. J-Y Girard, Cambridge University Press (1989) (Translated by Y. Lafont and P.Taylor)
- [6] *A proof of the standardization theorem in λ -calculus* Ryo Kashima (2000)