

Notes on Simply Typed Lambda Calculus*

Ralph Loader[†]

February, 1998

The purpose of this course is to provide an introduction to λ -calculi, specifically the simply typed lambda calculus (λ^\rightarrow).

λ -calculi are formalisms that are useful in computer science. They are languages that express both *computational* and *logical* information.

Computational information in that they can be seen as functional programming languages, or more realistically, a solid core on which to build a functional language.

Logical information in two ways. First, typed λ -calculi can be used directly as logics—these are ‘intuitionistic type theories’ such as the calculus of constructions (e.g., used in the ‘Lego’ theorem proving software developed in this department). Second, typed λ -calculi underly the term structure of higher order logics.

There are many different λ -calculi—everyone working in the field has their own pet language(s)—but they generally fall into a few classes of broadly similar languages. This course does not provide a wide ranging survey, but instead aims to explain a good range of theoretical techniques, most of which are applicable to a much wider range of languages than we will actually study.

The first half of the course gives a fairly thorough introduction to fairly traditional and fundamental results about λ -calculi. The second half of the course gives a detailed, quantitative analysis of λ^\rightarrow .

The book [3] covers much of the material in this course, especially sections 1–3. In section 1, we look at the untyped λ -calculus; [1] is a fairly comprehensive work on this. The article [2] is a fairly comprehensive work on pure type systems, which include the simply typed λ -calculus and impredicative intuitionistic type theories.

*Version 1.8 dated 1998/02/16 15:10:19.

[†]Department of Computer Science, University of Edinburgh. loader@dcs.ed.ac.uk

References

- [1] Hendrik P. Barendregt. *The Lambda Calculus*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, 1981 (revised edition 1984).
- [2] Hendrik P. Barendregt. Lambda calculi with types. In S. Abramsky, D. M. Gabbai, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume II. Oxford University Press, 1997.
- [3] Jean-Yves Girard, Yves Lafont, and Paul Taylor. *Proofs and Types*, volume 7 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, 1989.

Contents

1	Untyped Lambda Calculus	4
1.1	Syntax	4
1.2	Occurrences	5
1.3	Alpha Equivalence	6
1.4	Substitution	7
1.5	Beta Conversion	8
1.6	Some Examples—Recursive Functions	9
1.7	Church-Rosser (Confluence) Property	12
2	Simply Typed Lambda Calculus	14
2.1	Syntax	14
2.2	Basic Properties	16
2.3	Models	17
3	Strong Normalisation	17
3.1	Preliminaries	18
3.2	The Main Induction	20
3.3	A Combinatorial Proof	22
4	Quantitative Normalisation	23
4.1	Complete Developments	23
4.2	Redex Depths	24
4.3	Calculating Upper Bounds	25
5	Quantitative Strong Normalisation	26
5.1	Translations	26
5.2	Reduction of Translations	27
5.3	Translations and Strong Normalisation	28
6	Lower Bounds for Normalisation	30
6.1	Higher Order Propositional Classical Logic	31
6.2	Encoding the Boolean Hierarchy	32
6.3	Encoding the Logic	35
A	PCL_ω is not Elementary Recursive	35
A.1	Basic Encoding in PCL_ω	35
A.2	Encoding Turing Machines	37

1 Untyped Lambda Calculus

A typed λ -calculus has a language in two parts: the language of types, and the language of terms. In order to introduce some basic notions, we shall first look at the **untyped λ -calculus** Λ , so that we can consider terms independently of types.

1.1 Syntax

The set \mathcal{T} terms of the untyped λ -calculus is given by the following:

- A countably infinite set \mathcal{V} of variables. Each variable is a term.
- If t and r are terms, then so is tr .
- If s is a term, and \mathbf{x} is a variable, then $\lambda\mathbf{x}.s$ is a term.

Each term has a unique representation in one of the above forms, and the set of terms is the smallest set closed under the three clauses above.

In short, this is given by the following grammar:

$$\mathcal{T} := \mathcal{V} \mid \mathcal{T}\mathcal{T} \mid \lambda\mathcal{V}.\mathcal{T}$$

Induction and unique reading of terms are given as follows.

Induction: Suppose that S is a set (of terms) such that (a) $\mathbf{x} \in S$ for each variable \mathbf{x} , (b) if $t, r \in S$ then $tr \in S$ and (c) if $s \in S$ then $\lambda\mathbf{x}.s \in S$. Then S contains all terms.

Unique reading: Each term is exactly one of the following: a variable, an application, or an abstraction. If $tr = t'r'$ then $t = t'$ and $r = r'$. If $\lambda\mathbf{x}.s = \lambda\mathbf{x}'.s'$ then $\mathbf{x} = \mathbf{x}'$ and $s = s'$. (Also, if two variables are equal as terms, then they are equal as variables—this one is confusing.)

These two principles enable us to define functions on \mathcal{T} by recursion:

Recursion: Let S be a set, and $F_{\text{var}} : \mathcal{V} \rightarrow S$, $F_{\text{app}} : S \times S \rightarrow S$ and $F_{\text{abs}} : \mathcal{V} \times S \rightarrow S$ be functions. Then there is a unique function $F : \mathcal{T} \rightarrow S$ such that

$$F(\mathbf{x}) = F_{\text{var}}(\mathbf{x}), \quad F(tr) = F_{\text{app}}(F(t), F(r)), \quad F(\lambda\mathbf{x}.s) = F_{\text{abs}}(\mathbf{x}, F(s)).$$

As an application of the above, take S to be the set of natural numbers, and let $l_{\text{var}}(\mathbf{x}) = 1$, $l_{\text{app}}(m, n) = n + m + 1$ and $l_{\text{abs}}(\mathbf{x}, m) = m + 1$. Then the function l given by recursion satisfies:

$$l(\mathbf{x}) = 1, \quad l(tr) = l(t) + l(r) + 1, \quad l(\lambda\mathbf{x}.s) = l(s) + 1.$$

The number $l(r)$ is the **length** of the term r .

Another instance of recursion: Let $S = \mathcal{P}_{\text{fin}}(\mathcal{V})$, the set of finite sets of variables. Let $\text{FV}_{\text{var}}(\mathbf{x}) = \{\mathbf{x}\}$, $\text{FV}_{\text{app}}(c, a) = c \cup a$ and $\text{FV}_{\text{abs}}(\mathbf{x}, b) = b - \{\mathbf{x}\}$. Then the function FV given by recursion satisfies:

$$\text{FV}(\mathbf{x}) = \{\mathbf{x}\}, \quad \text{FV}(tr) = \text{FV}(t) \cup \text{FV}(r), \quad \text{FV}(\lambda\mathbf{x}.s) = \text{FV}(s) - \{\mathbf{x}\}.$$

The finite set $\text{FV}(r)$ is called the set of **free variables** occurring in r . A term t is **closed** if $\text{FV}(t) = \emptyset$.

The **bound variables** of a term are given by the following equations:

$$\text{BV}(\mathbf{x}) = \emptyset, \quad \text{BV}(tr) = \text{BV}(t) \cup \text{BV}(r), \quad \text{BV}(\lambda\mathbf{x}.s) = \{\mathbf{x}\} \cup \text{BV}(s).$$

The set of **sub-terms** of a term r is the set defined by:

$$\begin{aligned} \text{sub}(\mathbf{x}) &= \{\mathbf{x}\}, & \text{sub}(\lambda\mathbf{x}.s) &= \{\lambda\mathbf{x}.s\} \cup \text{sub}(s), \\ \text{sub}(tr) &= \text{sub}(t) \cup \text{sub}(r) \cup \{tr\}. \end{aligned}$$

1.2 Occurrences

Consider the term $r = \mathbf{x}\mathbf{x}$. The variable \mathbf{x} occurs twice in r . In order to reason about such situations, we shall give a formal definition of the notion of occurrence of a sub-term or variable, in terms of a path π telling us how to descend into a term. The reader who has an intuitive grasp of the notion of occurrence may wish to skip its formalisation...

The set of **paths** is $\{R, L, *\}^{<\omega}$, the set of finite sequences of R s, L s and $*$ s. For each term t we define $\text{path}(t)$, the set of **paths in t** , by

$$\begin{aligned} \text{path}(\mathbf{x}) &= \{\langle \rangle\} & \text{path}(\lambda\mathbf{x}.s) &= \{\langle *, \bar{\sigma} \mid \langle \bar{\sigma} \rangle \in \text{path}(s)\} \cup \{\langle \rangle\} \\ \text{path}(tr) &= \{\langle L, \bar{\sigma} \mid \langle \bar{\sigma} \rangle \in \text{path}(t)\} \cup \{\langle R, \bar{\sigma} \mid \langle \bar{\sigma} \rangle \in \text{path}(r)\} \cup \{\langle \rangle\} \end{aligned}$$

We define a function occ such that if $t \in \mathcal{T}$ and $\sigma \in \text{path}(t)$, then $\text{occ}(t, \sigma)$ is a sub-term of t .

$$\begin{aligned} \text{occ}(t, \langle \rangle) &= t \\ \text{occ}(tr, \langle L, \bar{\sigma} \rangle) &= \text{occ}(t, \langle \bar{\sigma} \rangle) & \text{occ}(tr, \langle R, \bar{\sigma} \rangle) &= \text{occ}(r, \langle \bar{\sigma} \rangle) \\ \text{occ}(\lambda\mathbf{x}.s, \langle *, \bar{\sigma} \rangle) &= \text{occ}(s, \langle \bar{\sigma} \rangle) \end{aligned}$$

A path σ is an **occurrence of s in r** if $\text{occ}(r, \sigma) = s$.

When it is convenient (i.e., nearly always) we shall abuse our terminology, confusing a sub-term with a particular sub-term occurrence. E.g., referring to an “occurrence s in r ”, rather than a “sub-term s of r and some occurrence σ of s in r ”.

Exercise 1.1 Refine the notions of free and bound variables of a term to free occurrences and bound occurrences of a variable in a term.

Suppose that σ is a path in r , that τ is a path in s , where s is the sub-term at σ in r , and t is the sub-term at τ in s . Show that the concatenation $\sigma\tau$ is a path in r , and that t is the sub-term of r at $\sigma\tau$.

Given a bound occurrence ρ of a variable \mathbf{x} in r , the **binder** of ρ is the longest initial sub-sequence σ of ρ , such that $\text{occ}(r, \sigma)$ is the form $\lambda\mathbf{x}. s$.

Exercise 1.2 Give a definition of the binder of a bound variable occurrence by recursion on the term r .

1.3 Alpha Equivalence

One of the most basic operations we need to consider is the replacement of a variable \mathbf{y} in a term r by another term s . This operation is called **substitution**, and is denoted $r[s/\mathbf{y}]$.

Unfortunately, there is an immediate complication in the definition of this. Consider the terms $r = \lambda\mathbf{x}. \mathbf{x} \mathbf{y}$ and $s = \mathbf{x} \mathbf{x}$. Then simply replacing \mathbf{y} by s in r gives the term $r' = \lambda\mathbf{x}. \mathbf{x} (\mathbf{x} \mathbf{x})$. The free variable \mathbf{x} in s has been ‘captured’ by the λ -abstraction, making it a bound variable in r' . This is undesirable for our purposes.

The solution is to first rename the bound variable \mathbf{x} in r to some other variable, before carrying out the replacement. E.g., renaming \mathbf{x} to \mathbf{z} in r gives the term $\lambda\mathbf{z}. \mathbf{z} \mathbf{y}$. Then replacing \mathbf{y} by $s = \mathbf{x} \mathbf{x}$ gives $\lambda\mathbf{z}. \mathbf{z} (\mathbf{x} \mathbf{x})$. The variable \mathbf{x} is free in this term, as desired.

The formalisation of the ‘renaming of bound variables in a term’ is called **α -equivalence**. Two terms r_1 and r_2 are α -equivalent ($r_1 \equiv_\alpha r_2$) if one can be obtained from the other by (a sensible) renaming of bound variables.

We could define α -equivalence in terms of a step-by-step renaming of bound variables in a term. Instead, we define α -equivalence in terms of occurrences. Two terms r_1 and r_2 are α -equivalent if the following hold:

1. $\text{path}(r_1) = \text{path}(r_2)$.
2. The sets of free variable occurrences in r_1 and r_2 are equal. The sets of bound variable occurrences in r_1 and r_2 are equal.
3. If σ is a free variable occurrence in r_1 (and by 2., in r_2 also), then $\text{occ}(r_1, \sigma) = \text{occ}(r_2, \sigma)$.
4. If σ is a bound variable occurrence in r_1 and r_2 , then the binder of σ in r_1 is equal to the binder of σ in r_2 .

Exercise 1.3 Show that α -equivalence is an equivalence relation.

Let r_1 and r_2 be terms with $\text{path}(r_1) = \text{path}(r_2)$. (1) Take $\sigma \in \text{path}(r_1) = \text{path}(r_2)$ and let s_1 and s_2 be the corresponding sub-terms occurring at σ . Show that s_1 and s_2 are either both variables, or both applications, or both abstractions. Show that the two parts of 2. in the definition of \equiv_α are equivalent for r_1 and r_2 .

Exercise 1.4 Give a formal definition of what it means to replace a bound variable \mathbf{x} by \mathbf{y} in a term r . Give a definition of α -equivalence in terms of this replacement, and prove that it gives the same relation as our definition. You must be careful to avoid variable capture (e.g., avoid renaming $\lambda\mathbf{x}.\lambda\mathbf{y}.\mathbf{x}$ to $\lambda\mathbf{x}.\lambda\mathbf{x}.\mathbf{x}$).

The exercise below shows that every term has plenty of nice α -equivalents.

Exercise 1.5 A term t is **regular** if it satisfies the condition that for every variable \mathbf{x} , there is at most one occurrence in t of a sub-term in the form $\lambda\mathbf{x}.s$, and no variable is both free and bound in t .

Let t_0 be a term, and let S be a finite set of variables (more generally, such that $\mathcal{V} - S$ is infinite). Show that there is a regular $t \equiv_\alpha t_0$ such that no member of S is bound in t .

When it is convenient, we consider the set $\mathcal{T}/\equiv_\alpha$ of α -equivalence classes of terms, rather than the set \mathcal{T} of terms. Often we will not distinguish between a term and its α -equivalence class, and write such things as $\lambda\mathbf{x}.\mathbf{x} = \lambda\mathbf{y}.\mathbf{y}$.

Observe that the basic syntactical operations of abstraction and application are well-defined on α -classes, and that we may prove that a predicate holds for all α -classes by structural induction. However, the unique-reading property does not hold for α -classes, as the α -class of a λ -abstraction $\lambda\mathbf{x}.s$ does not determine the variable \mathbf{x} . Hence, to define a function on α -classes by structural recursion, we must carry out the recursion on terms, and verify that the function on terms leads to a well-defined function on α -classes.

1.4 Substitution

Let r and s be terms, and \mathbf{y} be a variable, such that no variable is both free in s and bound in r , and \mathbf{y} is not bound in r . The substitution $r[s/\mathbf{y}]$ is defined by recursion on r , as follows:

- $\mathbf{y}[s/\mathbf{y}] = s$, and $\mathbf{x}[s/\mathbf{y}] = \mathbf{x}$ for $\mathbf{x} \neq \mathbf{y}$.
- $(tr)[s/\mathbf{y}] = (t[s/\mathbf{y}])(r[s/\mathbf{y}])$.

- $(\lambda \mathbf{x}. r)[s/\mathbf{y}] = \lambda \mathbf{x}.(r[s/\mathbf{y}])$.

The following exercise shows that the substitution above gives a well defined, and total, operation on α -equivalence classes.

Exercise 1.6 Show that the paths in $r[s/\mathbf{y}]$ are (a) the paths ρ in r , (b) concatenations $\rho\sigma$ where ρ is a free occurrence of \mathbf{y} in r and σ is a path in s . Express the sub-term at a path in $r[s/\mathbf{y}]$ in terms of sub-terms at paths in r and s .

Show that if $r_1 \equiv_\alpha r_2$ and $s_1 \equiv_\alpha s_2$ are such that $r_1[s_1/\mathbf{y}]$ and $r_2[s_2/\mathbf{y}]$ are both defined, then these two terms are α -equivalent.

Show that for any r , s and \mathbf{y} , there is $r' \equiv_\alpha r$ such that $r'[s/\mathbf{y}]$ is defined.

Exercise 1.7 Show that if a variable \mathbf{x}' is neither free nor bound in s' , then $s \equiv_\alpha s'$ iff $\lambda \mathbf{x}. s \equiv_\alpha \lambda \mathbf{x}'. s'[\mathbf{x}'/\mathbf{x}]$.

Exercise 1.8 Formulate a notion of **simultaneous substitution**, carrying out several substitutions at once: $s[r_1/\mathbf{x}_1 \dots r_m/\mathbf{x}_m]$. Note that in general $s[r_1/\mathbf{x}_1, r_2/\mathbf{x}_2] \neq s[r_1/\mathbf{x}_1][r_2/\mathbf{x}_2]$.

1.5 Beta Conversion

The intuition behind the syntactical constructs of abstraction and application, is that $\lambda \mathbf{x}. s$ is the function F defined by $F \mathbf{x} = s$ (i.e., $F(r)$ is given by replacing \mathbf{x} by r in s), and that tr is the result of applying the function t to argument r . This intuition is captured by the notion of β -reduction.

A **redex** is a term in the form $(\lambda \mathbf{x}. s)r$. The **basic β -reduction** relation is given by $(\lambda \mathbf{x}. s)r \rightarrow_b s'[r/\mathbf{x}]$, whenever $s' \equiv_\alpha s$ is such that $s'[r/\mathbf{x}]$ is defined.

The **one step β -reduction** relation \rightarrow_{β_1} is given by allowing a basic β -reduction of redex occurrence. In other words, \rightarrow_{β_1} is the least relation containing \rightarrow_b and closed under

$$\frac{t \rightarrow_{\beta_1} t'}{tr \rightarrow_{\beta_1} t'r}, \quad \frac{r \rightarrow_{\beta_1} r'}{tr \rightarrow_{\beta_1} tr'}, \quad \frac{s \rightarrow_{\beta_1} s'}{\lambda \mathbf{x}. s \rightarrow_{\beta_1} \lambda \mathbf{x}. s'}$$

The **n step β -reduction** relations \rightarrow_{β_n} are given inductively by $r \rightarrow_{\beta_0} r$ and, if $r \rightarrow_{\beta_n} r' \rightarrow_{\beta_1} r''$, then $r \rightarrow_{\beta_{n+1}} r''$. **β -reduction** \rightarrow_β is the least pre-order containing \rightarrow_{β_1} , so that $r \rightarrow_\beta r'$ iff there is n such that $r \rightarrow_{\beta_n} r'$. **β -conversion** (\equiv_β) is the least equivalence relation containing \rightarrow_{β_1} .

Exercise 1.9 For each of the reduction relations $\longrightarrow_{\bullet}$, give an analogous definition of $\longrightarrow_{\bullet}$ as a relation on α -equivalence classes. Show that you have made the correct definition, by showing that for α -classes r and r' , the following are equivalent (a) $r \longrightarrow_{\bullet} r'$, (b) there are $r_0 \in r$ and $r'_0 \in r'$ such that $r_0 \longrightarrow_{\bullet} r'_0$, and (c) for every $r_0 \in r$ there is $r'_0 \in r'$ such $r_0 \longrightarrow_{\bullet} r'_0$.

Show that (a) if $s \longrightarrow_{\beta_1} s'$ then $s[r/x] \longrightarrow_{\beta_1} s'[r/x]$, and (b) if $r \longrightarrow_{\beta} r'$, then $s[r/x] \longrightarrow_{\beta} s[r'/x]$.

A term is called **normal** if no redexes occur in it. A term is **weakly normalising** (WN) if it β -reduces to a normal term. A term r is **strongly normalising** (SN) if there is no infinite sequence $r = r_0 \longrightarrow_{\beta} r_1 \longrightarrow_{\beta} \dots$. Given a term t , we define $\text{sn}(t) = \sup\{n \mid \exists t' \text{ s.t. } t \longrightarrow_{\beta_n} t'\} \in \mathcal{N} \cup \{\infty\}$.

Exercise 1.10 Show that every strongly normalising term is weakly normalising, but not conversely. Show that $\text{sn}(t) < \infty$ iff t is SN.¹

1.6 Some Examples—Recursive Functions

We give some examples of terms, and their behaviour under reduction.

The **Church Numerals** are an encoding of the natural numbers, given by

$$\ulcorner n \urcorner = \lambda f. \lambda x. f^n(x)$$

where $f^n(x)$ indicates the usual iterated application: $f^0(x) = x$, $f^{n+1}(x) = f(f^n(x))$. A term F **encodes** a function f if $F \ulcorner i \urcorner \longrightarrow_{\beta} \ulcorner f(i) \urcorner$ for each n .² More generally, F **encodes** a function f of N arguments if

$$F \ulcorner i_1 \urcorner \dots \ulcorner i_N \urcorner \longrightarrow_{\beta} \ulcorner f(i_1 \dots i_N) \urcorner$$

for all $i_1 \dots i_N$.

The successor function $n \mapsto n+1$ is encoded by

$$S = \lambda n. \lambda f. \lambda x. f(n f x),$$

so that $S \ulcorner n \urcorner \longrightarrow_{\beta} \ulcorner n+1 \urcorner$. Addition is given by

$$\text{add} = \lambda m. \lambda n. \lambda f. \lambda x. m f (n f x),$$

so that $\text{add} \ulcorner n \urcorner \ulcorner m \urcorner \longrightarrow_{\beta} \ulcorner n + m \urcorner$, and multiplication by

$$\text{mult} = \lambda m. \lambda n. \lambda f. m(n f).$$

¹Note that one direction of this is not constructive. Everything else in these notes is constructive.

²By theorem 1.18 (Church-Rosser), a term encodes at most one function.

There is a **zero test**

$$?_0 = \lambda n. \lambda x. \lambda y. n (\lambda z. y) x$$

with the property that $?_0 \ulcorner 0 \urcorner r s \longrightarrow_{\beta} r$ and $?_0 \ulcorner n+1 \urcorner r s \longrightarrow_{\beta} s$.

There is the **Church pairing**, with

$$\text{pair} = \lambda x. \lambda y. \lambda z. z x y \quad \pi_1 = \lambda p. p (\lambda x. \lambda y. x) \quad \pi_2 = \lambda p. p (\lambda x. \lambda y. y)$$

so that

$$\pi_1 (\text{pair } r s) \longrightarrow_{\beta} r \quad \text{and} \quad \pi_2 (\text{pair } r s) \longrightarrow_{\beta} s.$$

This enables us to encode **primitive recursion**. Let

$$\text{step} = \lambda p. p (\lambda x. \lambda y. \text{pair } (S x) (\mathbf{f} x y))$$

so that $\text{step } (\text{pair } \ulcorner n \urcorner a) \longrightarrow_{\beta} \text{pair } \ulcorner n+1 \urcorner (\mathbf{f} a)$. Defining

$$\text{rec}_p = \lambda n. n \text{ step } (\text{pair } \ulcorner 0 \urcorner \mathbf{x})$$

we have

$$\text{rec}_p \ulcorner 0 \urcorner \longrightarrow_{\beta} \text{pair } \ulcorner 0 \urcorner \mathbf{x}$$

and, if $\text{rec}_p \ulcorner n \urcorner \longrightarrow_{\beta} \text{pair } \ulcorner n \urcorner a$, then

$$\text{rec}_p \ulcorner n+1 \urcorner \longrightarrow_{\beta} \text{pair } \ulcorner n+1 \urcorner (\mathbf{f} \ulcorner n \urcorner a).$$

Letting $\text{rec} = \lambda \mathbf{f}. \lambda \mathbf{x}. \lambda n. \pi_1 (\text{rec}_p \mathbf{f} \mathbf{x} n)$, we have that

$$\text{rec } \mathbf{f} a \ulcorner 0 \urcorner \equiv_{\beta} a \quad \text{rec } \mathbf{f} a \ulcorner n+1 \urcorner \equiv_{\beta} \mathbf{f} \ulcorner n \urcorner (\text{rec } \mathbf{f} a \ulcorner n \urcorner).$$

Not every term has a normal form. For example, letting

$$\Omega = (\lambda \mathbf{x}. \mathbf{x} \mathbf{x}) (\lambda \mathbf{x}. \mathbf{x} \mathbf{x}),$$

we have that $\Omega \longrightarrow_{\beta_1} \Omega \longrightarrow_{\beta_1} \dots$. Some terms without normal forms are useful, as the example following shows. Programming languages usually allow recursive function definitions: $F \mathbf{x} = t$, where t is some expression containing both \mathbf{x} and F . In λ -calculus notation, $F = \lambda \mathbf{x}. t$. Of course, we cannot solve this up to equality; the next best thing is to find a term F such that $F \longrightarrow_{\beta} \lambda \mathbf{x}. t$. More generally, letting $G = \lambda F. \lambda \mathbf{x}. t$, it suffices to find Y_G such that $Y_G \longrightarrow_{\beta} G Y_G$. For example, take

$$Y_G = (\lambda \mathbf{x}. G (\mathbf{x} \mathbf{x})) (\lambda \mathbf{x}. G (\mathbf{x} \mathbf{x})).$$

Exercise 1.11 Check all claims made above.

Exercise 1.12 Find a **fixed-point combinator** Y that β -reduces to the term $\lambda f. f(Y f)$. (Given Y , we could have taken the Y_G above to be $Y G$. You can also go the other way round, and find G such that $Y = Y_G$ has the required property.)

Using Y , we can encode minimisation. If f is a function on the natural numbers such that $f(i) = 0$ for some i , then μf is the least such i . Similarly, let $\mu_{\geq n} f$ be the least $i \geq n$ such that $f(i) = 0$, if such exists. This is given recursively by:

$$\mu_{\geq n} f = \begin{cases} n & \text{if } f(n) = 0 \\ \mu_{\geq n+1} f & \text{if } f(n) \neq 0 \end{cases}$$

Hence, we want a term μ_{\geq} , such that $\mu_{\geq} n f \rightarrow_{\beta} ?_0 (f n) n (\mu_{\geq} (S n) f)$. Letting

$$\mu_{\geq} = Y (\lambda M. \lambda n. \lambda f. ?_0 (f n) n (M (S n) f))$$

does the trick, and then let $\mu = \mu_{\geq} \ulcorner 0 \urcorner$. With these definitions, if F encodes a function f , and f has a zero, then

$$\mu F \rightarrow_{\beta} \ulcorner \mu f \urcorner.$$

Exercise 1.13 The **total recursive functions** are the least set of functions (of arbitrary arity) on the natural numbers, that

1. Contains the constant functions, the successor function and the projections $(x_1 \dots x_m) \mapsto x_i$ (whenever $1 \leq i \leq m$).
2. Is closed under composition: $h(\bar{x}) = f(g_1(\bar{x}) \dots g_n(\bar{x}))$.
3. Is closed under primitive recursion: $h(\bar{x}, 0) = f(\bar{x})$, $h(\bar{x}, n + 1) = g(\bar{x}, n, f(\bar{x}, n))$.
4. Is closed under minimisation, whenever this gives a *total* function: if for each \bar{x} there is n such that $f(\bar{x}, n) = 0$, then take $g(\bar{x})$ to be the least such n .

Show that every total recursive function can be encoded by a λ -term.

It was conjectured by Church that the computable functions on the natural numbers are precisely those that can be encoded by a λ -term. This is the famous **Church's Thesis**. The exercise above is one half of showing that the functions encodable in Λ coincide with the recursive functions. Unsurprisingly, the functions on the natural numbers encodable in Λ coincide with the functions computable in lots of other mathematical models of computability.

Theorem 1.14 *The predicate of being weakly normalising is not decidable.*

PROOF SKETCH: Suppose that WN were decidable. Then we could define a computable function F on λ -terms such that $F(t)$ is the normal form of t if there is one, and $F(t) = \lambda \mathbf{x}. \mathbf{x}$ otherwise. Then F would be a total computable function majorising every total recursive function, which is not possible. \square

1.7 Church-Rosser (Confluence) Property

Our reduction relation is not deterministic—given a term r , it may contain several redexes, and hence several different ways of reducing it.

However, there is a certain determinism to the intuition behind our calculus—function application gives a unique value. This is expressed locally in the fact that there is (up to \equiv_α) only one way of reducing a given redex. It can also be expressed globally in the confluence, or Church-Rosser, property: if $r \longrightarrow_\beta r_1$ and $r \longrightarrow_\beta r_2$, then there is r' such that $r_1 \longrightarrow_\beta r'$ and $r_2 \longrightarrow_\beta r'$. In particular, a term has at most one normal form.

The proof of Church-Rosser is based on a quite simple idea—if r contains two redexes, then the reduction of one redex does not stop the reduction of the other redex, and it does not matter in what order the reductions are carried out. This makes it quite simple to show that if $r \longrightarrow_{\beta_1} r_1, r_2$, then there is r' such that $r_1, r_2 \longrightarrow_\beta r'$. However, note the β_1 - and β - reductions in this. To extend this result to confluence for β -reduction, we must be careful to avoid an infinite regress. The proof we give (due to Hyland) replaces $\longrightarrow_{\beta_1}$ with a ‘parallel’ reduction, which simultaneously reduces any collection of redexes in a term.

Let R be a relation, and \leq be the least pre-order containing R . R is **confluent**, if whenever $a R a_1$ and $a R a_2$, there is a' such that $a_1 R a'$ and $a_2 R a'$. R is **weakly confluent**, if whenever $a R a_1, a_2$, there is a' such that $a_1, a_2 \leq a'$.

Exercise 1.15 Let R be a relation and \leq the least pre-order containing R . Show that if R is confluent, then so is \leq . Find a weakly confluent R such that \leq is not confluent.

The **parallel reduction** relation, \longrightarrow_p , is the relation on α -classes given inductively by:

$$\frac{}{\mathbf{x} \longrightarrow_p \mathbf{x}} \quad \frac{t \longrightarrow_p t' \quad r \longrightarrow_p r'}{tr \longrightarrow_p t'r'} \quad \frac{s \longrightarrow_p s'}{\lambda \mathbf{x}. s \longrightarrow_p \lambda \mathbf{x}. s'}$$

$$\frac{s \longrightarrow_p s' \quad r \longrightarrow_p r'}{(\lambda \mathbf{x}. s') r \longrightarrow_p s'[r'/\mathbf{x}]}$$

The lemma below follows easily from the definitions of the various reduction relations.

Lemma 1.16 *If $r \longrightarrow_{\beta_1} r'$ then $r \longrightarrow_p r'$. If $r \longrightarrow_p r'$, then $r \longrightarrow_{\beta} r'$. The least pre-order containing \longrightarrow_p is \longrightarrow_{β} . If $s \longrightarrow_p s'$ and $r \longrightarrow_p r'$, then $s[r/\mathbf{x}] \longrightarrow_p s'[r'/\mathbf{x}]$. \square*

Lemma 1.17 \longrightarrow_p *is confluent.*

PROOF: We show that if t, t_1 and t_2 are α -classes, with $t \longrightarrow_p t_1$ and $t \longrightarrow_p t_2$, then, then there is t' such that $t_1 \longrightarrow_p t'$ and $t_2 \longrightarrow_p t'$. We perform the construction by induction on t .

The non-trivial step is for t in the form $(\lambda \mathbf{x}. s) r$. Then (for $i = 1, 2$) there are s_i and r_i such that $s \longrightarrow_p s_i$, $r \longrightarrow_p r_i$, and either $t_i = (\lambda \mathbf{x}. s_i) r_i$ or $t_i = s_i[r_i/\mathbf{x}]$. By the induction hypothesis, there are r' and s' such that (for $i = 1, 2$) $r_i \longrightarrow_p r'$ and $s_i \longrightarrow_p s'$. Let $t' = s'[r'/\mathbf{x}]$. Then we have that both $(\lambda \mathbf{x}. s_i) r_i \longrightarrow_p t'$ and $s_i[r_i/\mathbf{x}] \longrightarrow_p t'$, so that $t_1, t_2 \longrightarrow_p t'$.

The other cases are easy: e.g. if t is in the form $s r$, but is not a redex, then the t_i are in the form $s_i r_i$ where $s \longrightarrow_p s_1, s_2$ and $r \longrightarrow_p r_1, r_2$. Applying the induction hypothesis, there are s' and r' such that $s_1, s_2 \longrightarrow_p s'$ and $r_1, r_2 \longrightarrow_p r'$. Letting $t' = s' r'$, we have $t_1, t_2 \longrightarrow_p t'$, as required. \square

Theorem 1.18 (Church-Rosser) *The relation \longrightarrow_{β} is confluent (on α -classes). \square*

Corollary 1.19 *If a term β -reduces to a normal form, then the normal form is unique (up to \equiv_{α}). \square*

Exercise 1.20 (Standardisation) The **standard** reduction relation \longrightarrow_s is given by:

$$\frac{r_1 \longrightarrow_s r'_1 \quad \dots \quad r_m \longrightarrow_s r'_m}{\mathbf{x} r_1 \dots r_m \longrightarrow_s \mathbf{x} r'_1 \dots r'_m} \quad \frac{s[r_0/\mathbf{x}] r_1 \dots r_m \longrightarrow_s t}{(\lambda \mathbf{x}. s) r_0 r_1 \dots r_m \longrightarrow_s t}$$

$$\frac{s \longrightarrow_s s' \quad r_1 \longrightarrow_s r'_1 \quad \dots \quad r_m \longrightarrow_s r'_m}{(\lambda \mathbf{x}. s) r_1 \dots r_m \longrightarrow_s (\lambda \mathbf{x}. s) r'_1 \dots r'_m}$$

Show that if $r \longrightarrow_s r'$, then $r \longrightarrow_{\beta} r'$, and give a characterisation of the reduction sequences obtained (in terms of what redexes are reduced in what

order). Show that if r' is normal, then $r \rightarrow_{\beta} r'$ by a reduction sequence that always reduces the leftmost redex (appropriately defined).

Show that \rightarrow_s is closed under the following:

$$\frac{s \rightarrow_s s' \quad r \rightarrow_s r'}{s[r/x] \rightarrow_s s'[r'/x]} \quad \frac{t \rightarrow_s (\lambda x. s') r'}{t \rightarrow_s s'[r'/x]} \quad \frac{t \rightarrow_s t' \rightarrow_{\beta 1} t''}{t \rightarrow_s t''}$$

and that \rightarrow_s and \rightarrow_{β} are the same relation.

Conclude that if $r \rightarrow_{\beta} r'$, then r reduces to r' with a reduction sequence in some special form. In particular, if $r \rightarrow_{\beta} r'$, and r' is normal, then leftmost reduction reduces r to r' . These are the **standardisation theorems**.

The notion of a term F encoding a (total) function f , is extended to **F encoding** a partial function f , by requiring that if $f(n)$ is undefined, then $F \uparrow n \uparrow$ does not β -reduce to a λ -abstraction.³

In exercise 1.13, we referred to only *total* recursive functions. This is because we didn't have results useful for showing partiality as encoded above. The standardisation results in the previous exercise are such results.

Exercise 1.21 The **partial recursive functions** are given by removing the condition (in exercise 1.13) that functions given by minimisation are total. Show that every partial recursive function can be encoded by a λ -term. You will need to take care of the fact that mathematical functions are strict, in that $f(x_1 \dots x_m)$ is undefined if any of the x_i are, while the obvious encodings in Λ are not strict.

2 Simply Typed Lambda Calculus

So far, we have dealt with the *untyped* λ -calculus. There is no restriction on how we can use objects—for example, any value can be used as a function. In real life, this is undesirable; what does it mean to apply the number 5 to the function \cos ? Type systems are one way of (statically) ensuring that operations are only applied to appropriate objects. We shall introduce a type system for the λ -calculus, obtaining the **simply typed λ -calculus**, λ^{\rightarrow} for short.

2.1 Syntax

The syntax of λ^{\rightarrow} is presented in stages.

³Perhaps not having a normal form would seem more intuitive. However, the stronger condition given is easier to work with.

1. The types are defined.
2. The pre-terms are defined. Pre-terms are essentially terms of Λ , and include things that we do not wish to consider in λ^\rightarrow , such as $\lambda\mathbf{x}:o. \mathbf{x} \mathbf{x}$.
3. A system of typing judgements is given, that identifies the terms—those pre-terms that we wish to consider in λ^\rightarrow .

We take an arbitrary non-empty set \mathcal{G} of **ground** types. For concreteness and simplicity, we shall take $\mathcal{G} = \{o\}$; there is no reason why some other (larger) set should not be used. Types are then given by the following grammar:

$$\tau := \mathcal{G} \mid \tau \rightarrow \tau.$$

In order to avoid ambiguity in λ^\rightarrow , we decorate λ -abstractions with types. The grammar for **pre-terms** of λ^\rightarrow is:

$$\mathcal{T} := \mathcal{V} \mid \mathcal{T} \mathcal{T} \mid \lambda\mathcal{V}:\tau. \mathcal{T}.$$

The notions such as α -equivalence, substitution and reduction can be carried across more or less verbatim from Λ to λ^\rightarrow . The definition of α -equivalence must be modified to ensure that the type decorations are preserved, e.g., so that $\lambda\mathbf{x}:A. \mathbf{x} \equiv_\alpha \lambda\mathbf{y}:B. \mathbf{y}$ if and only if $A = B$. Our proof of the Church-Rosser theorem still works for the pre-terms of the typed calculus.

A **context** is a function mapping a finite set of variables into the set of types. Contexts will be written in the form $\mathbf{x}_1:A_1 \dots \mathbf{x}_m:A_m$. They should be read as a typing declaration: the context $\mathbf{x}_1:A_1 \dots \mathbf{x}_m:A_m$ indicates that each variable \mathbf{x}_i will be considered to have the corresponding type A_i . We present a ternary relation $\Gamma \vdash s : B$ on contexts Γ , pre-terms s and types B by the following clauses:

$$\frac{}{\mathbf{x}_1:A_1 \dots \mathbf{x}_m:A_m \vdash \mathbf{x}_i : A_i} \quad \text{for } 1 \leq i \leq m$$

$$\frac{\Gamma \vdash f : A \rightarrow B \quad \Gamma \vdash r : A}{\Gamma \vdash fr : B}$$

$$\frac{\Gamma \mathbf{x}:A \vdash s : B}{\Gamma \vdash \lambda\mathbf{x}. s : A \rightarrow B}$$

A pre-term s is said to be a **term** (of type B in context Γ) if $\Gamma \vdash s : B$.

We also define the relation $\Gamma \vdash s : B$ for α -classes s of pre-terms, using identical clauses.

Exercise 2.1 *Show that the relations \vdash for pre-terms, and for α -classes, are related as follows: an α -class s has $\Gamma \vdash s : B$ derivable if and only if there is $s_0 \in s$ such that $\Gamma \vdash s_0 : B$ is derivable.*

2.2 Basic Properties

Lemma 2.2 (Renaming) *If $\Gamma z:C \Delta \vdash s : B$ for some α -class s , and z' is not in the context, then $\Gamma z':C \Delta \vdash s[z'/z] : B$.*

PROOF: We in fact show that simultaneous renaming is possible, in that if $x_1:A_1 \dots x_m:A_m \vdash s : B$, and $x'_1 \dots x'_m$ is a sequence of distinct variables, then $x'_1:A_1 \dots x'_m:A_m \vdash s[x'_1/x_1 \dots x'_m/x_m] : B$.

We show this by induction on derivations. We do the case of a derivation ending

$$\frac{x_1:A_1 \dots x_m:A_m \ x:A \vdash s : B}{x_1:A_1 \dots x_m:A_m \vdash \lambda x:A. s : A \rightarrow B.}$$

Let x' be a variable distinct from the x'_i . By the induction hypothesis,

$$x'_1:A_1 \dots x'_m:A_m \ x':A \vdash s[x'_1/x_1 \dots x'_m/x_m, x'/x] : B$$

and hence

$$x'_1:A_1 \dots x'_m:A_m \vdash \lambda x':A. (s[x'_1/x_1 \dots x'_m/x_m, x'/x]) : B$$

are derivable. The equation

$$\lambda x':A. (s[x'_1/x_1 \dots x'_m/x_m, x'/x]) = (\lambda x:A. s)[x'_1/x_1 \dots x'_m/x_m]$$

holds, and the induction step follows. \square

Lemma 2.3 (Weakening) *Let s be an α -class such that $\Gamma \Delta \vdash s : B$. Let z be a variable not in $\Gamma \Delta$. Then $\Gamma z:C \Delta \vdash s : B$*

PROOF: By the renaming lemma, it suffices to show that there is z such that $\Gamma z:C \Delta \vdash s : B$. This can be shown by an easy induction on derivations. \square

Lemma 2.4 (Uniqueness of Derivations) *If s is a pre-term, with $\Gamma \vdash s : B$, then the derivation of this fact is unique.* \square

Lemma 2.5 (Unicity of Typing) *If s is an α -class, with $\Gamma \vdash s : B$ and $\Gamma \vdash s : B'$, then $B = B'$.* \square

Exercise 2.6 *Prove the two lemmas above.*

Lemma 2.7 (Substitution) *Let r and s be α -classes such that $\Gamma \vdash r : A$ and $\Gamma x:A \Delta \vdash s : B$. Then $\Gamma \Delta \vdash s[r/x] : B$.* \square

Exercise 2.8 *Prove it.* Use induction on a derivation of $\Gamma x:A \Delta \vdash s : B$.

Lemma 2.9 (Subject Reduction) *Let s be an α -class with $\Gamma \vdash s : B$, and $s \longrightarrow_{\beta} s'$. Then $\Gamma \vdash s' : B$ also.*

PROOF: The case when $s \longrightarrow_b s'$ reduces to the substitution lemma above. Closing under the term constructs of application and abstraction shows that the lemma holds for $s \longrightarrow_{\beta_1} s'$. An induction shows that the lemma holds for \longrightarrow_{β} also. \square

2.3 Models

We give a notion of model of λ^{\rightarrow} . What is given here is far from the most general notion of model possible. A **model** is a map \mathcal{M} assigning a set $\mathcal{M}(o)$ to each ground type. We define the **interpretation** $\llbracket \cdot \rrbracket$, first for types, and then for terms. Put $\llbracket o \rrbracket = \mathcal{M}(o)$, and $\llbracket A \rightarrow B \rrbracket = \llbracket B \rrbracket^{\llbracket A \rrbracket}$ (the set of all functions from $\llbracket A \rrbracket$ to $\llbracket B \rrbracket$).

Given a context $\Gamma = \mathbf{x}_1:A_1 \dots \mathbf{x}_m:A_m$, we define a **valuation** σ on Γ to be a function defined on the \mathbf{x}_i such that $\sigma(\mathbf{x}_i) \in \llbracket A_i \rrbracket$ for $1 \leq i \leq m$. Given another variable \mathbf{x} and $a \in \llbracket A \rrbracket$, we define $\sigma[\mathbf{x} := a](\mathbf{x}_i) = \sigma(\mathbf{x}_i)$ and $\sigma[\mathbf{x} := a](\mathbf{x}) = a$, so that $\sigma[\mathbf{x} := a]$ is a valuation on $\Gamma \mathbf{x}:A$.

We define $\llbracket s \rrbracket_{\sigma} \in \llbracket B \rrbracket$ whenever $\Gamma \vdash s : B$ and σ is a valuation on Γ , by

$$\llbracket \mathbf{x} \rrbracket_{\sigma} = \sigma(\mathbf{x}) \quad \llbracket tr \rrbracket_{\sigma} = \llbracket t \rrbracket_{\sigma}(\llbracket r \rrbracket_{\sigma})$$

and, for $a \in \llbracket A \rrbracket$,

$$\llbracket \lambda \mathbf{x}:A. s \rrbracket_{\sigma}(a) = \llbracket s \rrbracket_{\sigma[\mathbf{x}:=a]}.$$

The following lemma shows that our semantics gives some meaning to β -reduction, as well as to syntax.

Lemma 2.10 *If $\Gamma \vdash r : A$ and $\Gamma \mathbf{x}:A \vdash s : B$, then*

$$\llbracket s[r/\mathbf{x}] \rrbracket_{\sigma} = \llbracket s \rrbracket_{\sigma[\mathbf{x} := \llbracket r \rrbracket_{\sigma}]}.$$

If $\Gamma \vdash s : B$ and $s \longrightarrow_{\beta} s'$, then $\llbracket s \rrbracket_{\sigma} = \llbracket s' \rrbracket_{\sigma}$. \square

Exercise 2.11 *Prove the lemma.* Be careful about contexts!

3 Strong Normalisation

The set-theoretic semantics of λ^{\rightarrow} of section 2.3 gives a sense in which λ^{\rightarrow} meaningful—it formalises the intuition that $A \rightarrow B$ is a type of functions,

that terms of a function type are functions, and that terms are values of some sort. The semantics gives meaning to β -reduction, by lemma 2.10.

The meaning given to λ^\rightarrow by the semantics is *static* and *external*. Static in the sense that terms are represented by values, rather than by computations; it gives no sense of the term t computing the value $\llbracket t \rrbracket$. External in the sense that terms are given meaning in terms of some structure outside of the syntax.

The normalisation theorem proved in this section gives meaning to λ^\rightarrow in a way that is *dynamic* and *internal*, by showing that every term is strongly normalising. This is dynamic in the sense that the meaning of a term is given by a computational process: the reduction of a term t to a normal form v is in a series of steps $t = t_0 \longrightarrow_{\beta_1} t_1 \longrightarrow_{\beta_1} \cdots \longrightarrow_{\beta_1} t_n = v$. These steps can be seen as the successive states of a computer running the program t to reach a value v . Strong normalisation is *internal*, in that it gives meaning to the calculus without reference to another structure.

A common way to prove the termination of some sequence t_0, t_1, \dots is to find some measure $|t_i| \in \mathcal{N}$ that is strictly decreasing in i . Later, we will carry out such a construction, in order to give a quantitative analysis of strong normalisation for λ^\rightarrow . However, such a proof is non-trivial. A first attempt might be to take $|t|$ to be the length of the term t . This fails as reduction of a redex $(\lambda x. s) r$ to $s[r/x]$ can increase length. In the proof of strong normalisation that follows, we avoid this problem by taking into account the meaning of terms: we show that a term $f : A \rightarrow B$ is not just SN, but that f maps SN terms of type A to SN terms of type B .

Normalisation for λ^\rightarrow was first proved by Turing, although Gentzen had earlier given similar cut-elimination results for sequent calculi. The proof method we use is due to Tait. Girard extended it to prove normalisation for F , the second order λ -calculus, and since then it has been used to prove normalisation for just about everything.

3.1 Preliminaries

We start by noting some properties of SN, the set of strong normalising terms of Λ . We leave it as an exercise to rewrite these for λ^\rightarrow .

Exercise 3.1 *If $t \in \text{SN}$ and $t \longrightarrow_{\beta} t'$, then $t' \in \text{SN}$. If $t' \in \text{SN}$ for all t' such that $t \longrightarrow_{\beta_1} t'$, then $t \in \text{SN}$. If $t \in \text{SN}$ and t' is a sub-term of t , then $t' \in \text{SN}$.*

Lemma 3.2 *The set SN is closed under the following:*

$$\frac{r_1 \quad \dots \quad r_m}{\mathbf{x} r_1 \dots r_m} \quad \frac{s}{\lambda \mathbf{x}. s} \quad \frac{s[r_0/\mathbf{x}] r_1 \dots r_m \quad r_0}{(\lambda \mathbf{x}. s) r_0 r_1 \dots r_m} \quad (1)$$

PROOF: The first two clauses are more or less trivial. We do the last clause. Recall that $\text{sn}(t)$ is the l.u.b. of the length of the reduction sequences starting from t . Suppose that $s[r_0/\mathbf{x}] r_1 \dots r_m \in \text{SN}$ and $r_0 \in \text{SN}$.

We show by induction on the number $\text{sn}(s[r_0/\mathbf{x}] r_1 \dots r_m) + \text{sn}(r_0)$, that $(\lambda \mathbf{x}. s) r_0 r_1 \dots r_m \in \text{SN}$.

Suppose $(\lambda \mathbf{x}. s) r_0 r_1 \dots r_m \longrightarrow_{\beta_1} t$. We show that t is SN by considering the possibilities for t .

If $t = s[r_0/\mathbf{x}] r_1 \dots r_m$, then $t \in \text{SN}$.

If $t = (\lambda \mathbf{x}. s) r'_0 r_1 \dots r_m$ where $r_0 \longrightarrow_{\beta_1} r'_0$, then we get the β -reduction $s[r_0/\mathbf{x}] r_1 \dots r_m \longrightarrow_{\beta} s[r'_0/\mathbf{x}] r_1 \dots r_m$. This gives the inequalities

$$\text{sn}(s[r'_0/\mathbf{x}] r_1 \dots r_m) \leq \text{sn}(s[r_0/\mathbf{x}] r_1 \dots r_m) \quad \text{and} \quad \text{sn}(r'_0) < \text{sn}(r_0).$$

By the induction hypothesis, $t \in \text{SN}$.

Otherwise, the redex reduced is in s or one of $r_1 \dots r_m$, and we have $t = (\lambda \mathbf{x}. s') r_0 r'_1 \dots r'_m$ where $s[r_0/\mathbf{x}] r_1 \dots r_m \longrightarrow_{\beta_1} s'[r_0/\mathbf{x}] r'_1 \dots r'_m$, and applying the induction hypothesis gives $t \in \text{SN}$. \square

Exercise 3.3 *Refine the proof of lemma 3.2 to show that if t can be derived by using the clauses of (1) n times, then $\text{sn}(t) \leq n$.*

Show that SN is the least set closed under the three clauses (1)—i.e., that every SN term can be derived using those clauses.

Exercise 3.4 *Show that if \mathbf{x} is free in s , and $s[r_0/\mathbf{x}] r_1 \dots r_m \in \text{SN}$, then $(\lambda \mathbf{x}. s) r_0 \dots r_m \in \text{SN}$.*

The λI -calculus Λ_I is given in exactly the same way as Λ , except that in all λ -abstractions $\lambda \mathbf{x}. s$, the variable \mathbf{x} is required to be free in s . *Show that every weakly normalising Λ_I term is strongly normalising.* Hint: use the standardisation results of exercise 1.20 and the first part of this exercise.

We wish to make the following definition, recursive in types:

- Red_o is the set of strongly normalising terms of type o , for any ground type o .
- $\text{Red}_{A \rightarrow B}$ is the set of terms f of type $A \rightarrow B$, such that for any $r \in \text{Red}_A$, we have $f r \in \text{Red}_B$.

Unfortunately, this does not quite work for our purposes, as the contexts of f and r might be different. One way of avoiding this issue is to reformulate λ^\rightarrow slightly, so that variables have types pre-assigned, rather than being declared in contexts. Another way is to simply keep track of contexts; we therefore make the following definition of **the sets of reducible terms**:

$$\text{Red}_o = \{(\Gamma, r) \mid r \in \text{SN}, \Gamma \vdash r : o\}$$

for ground types o , and

$$\text{Red}_{A \rightarrow B} = \{(\Gamma, f) \mid \Gamma \vdash f : A \rightarrow B, (\Gamma\Delta, f r) \in \text{Red}_B \forall (\Gamma\Delta, r) \in \text{Red}_A\}.$$

Lemma 3.5 1. If $(\Gamma, s) \in \text{Red}_B$ then $s \in \text{SN}$.

2. If $\Gamma \vdash \mathbf{x} r_1 \dots r_m : B$, and $r_1, \dots, r_m \in \text{SN}$, then

$$(\Gamma, \mathbf{x} r_1 \dots r_m) \in \text{Red}_B.$$

3. If $(\Gamma, s[r_0/\mathbf{x}] r_1 \dots r_m) \in \text{Red}_B$, $\Gamma \vdash r_0 : A$ and $r_0 \in \text{SN}$, then

$$(\Gamma, (\lambda \mathbf{x}:A. s) r_0 r_1 \dots r_m) \in \text{Red}_B.$$

PROOF: We prove the properties simultaneously by induction on the type B . For $B = o$, 1. is trivial, and 2. and 3. are properties of SN. Suppose that $B = C \rightarrow D$.

For 1., let \mathbf{z} be a variable not in Γ , so that by the induction hypothesis of 2. for C , $(\Gamma \mathbf{z}:C, \mathbf{z}) \in \text{Red}_C$. As $(\Gamma, s) \in \text{Red}_{C \rightarrow D}$, we have $(\Gamma \mathbf{z}:C, s \mathbf{z}) \in \text{Red}_D$, so that by the induction hypothesis of 1. for D , $s \mathbf{z} \in \text{SN}$ and hence $s \in \text{SN}$ also.

For 2., take any $(\Gamma\Delta, r_{m+1}) \in \text{Red}_C$. By the IH 1. for C , we have that $r_{m+1} \in \text{SN}$, so that by the IH 2. for D , $(\Gamma\Delta, \mathbf{x} r_1 \dots r_{m+1}) \in \text{Red}_D$. The definition of $\text{Red}_{C \rightarrow D}$ now gives $(\Gamma, \mathbf{x} r_1 \dots r_m) \in \text{Red}_B$.

For 3., take any $(\Gamma\Delta, r_{m+1}) \in \text{Red}_C$. Then, $(\Gamma\Delta, s[r_0/\mathbf{x}] r_1 \dots r_{m+1}) \in \text{Red}_D$, so that by the IH, $(\Gamma\Delta, (\lambda \mathbf{x}:A. s) r_0 r_1 \dots r_{m+1}) \in \text{Red}_D$. This shows that $(\Gamma, (\lambda \mathbf{x}:A. s) r_0 r_1 \dots r_m) \in \text{Red}_B$. \square

Exercise 3.6 Check that if $(\Gamma, s) \in \text{Red}_B$, then $(\Gamma \mathbf{x}:A, s) \in \text{Red}_B$ also.

3.2 The Main Induction

We prove that terms of λ^\rightarrow are SN, by showing that they are members of the reducibility sets defined above. The definition of $\text{Red}_{A \rightarrow B}$ as the set of those terms which give a function from Red_A to Red_B , means that the sets Red_A form something like a model of λ^\rightarrow . Our proof that every term is reducible will be a structural induction bearing a more than passing similarity to the definition of the interpretation of a term in a model.

Lemma 3.7 *Suppose $y_1:B_1 \dots y_n:B_n \vdash r : A$. For any $(\Gamma, s_j) \in \text{Red}_{B_j}$ ($1 \leq j \leq n$), we have $(\Gamma, r[s_1/y_1 \dots s_n/y_n]) \in \text{Red}_A$.*

PROOF: We write \bar{s}/\bar{y} for $s_1/y_1 \dots s_n/y_n$. The proof is by induction on the term r .

For r a variable, the lemma is trivial.

Consider an application tr where $r : A$ and $t : A \rightarrow B$. The induction hypotheses are that $(\Gamma, t[\bar{s}/\bar{y}]) \in \text{Red}_{A \rightarrow B}$ and $(\Gamma, r[\bar{s}/\bar{y}]) \in \text{Red}_A$. The definition of $\text{Red}_{A \rightarrow B}$ gives $(\Gamma, (tr)[\bar{s}/\bar{y}]) \in \text{Red}_B$, as required.

Consider a λ -abstraction, $\lambda x:A. b : A \rightarrow B$. Take any $(\Gamma\Gamma', a) \in \text{Red}_A$. For each j , $(\Gamma\Gamma', s_j) \in \text{Red}_{B_j}$, so that by the induction hypothesis,

$$(\Gamma\Gamma', b[\bar{s}/\bar{y}, a/x]) \in \text{Red}_B.$$

Also, $a \in \text{SN}$ as $a \in \text{Red}_A$, so that by 3. of lemma 3.5, we have

$$(\Gamma\Gamma', (\lambda x:A. b)[\bar{s}/\bar{y}] a) \in \text{Red}_B.$$

By the definition of $\text{Red}_{A \rightarrow B}$, we get $(\Gamma, (\lambda x:A. b)[\bar{s}/\bar{y}]) \in \text{Red}_{A \rightarrow B}$. \square

Theorem 3.8 *Every term of λ^\rightarrow is SN.*

PROOF: Given $y_1:B_1 \dots y_n:B_n \vdash r : A$, take $\Gamma = y_1:B_1 \dots y_n:B_n$ and $s_j = y_j$ in the lemma above. This gives $(\Gamma, r) = (\Gamma, r[\bar{y}/\bar{y}]) \in \text{Red}_A$, and so $r \in \text{SN}$. \square

Exercise 3.9 Gödel's system T is the following extension of λ^\rightarrow : take the ground type to be \mathcal{N} , and introduce constructs 0 , S and R with the following rules:

$$\frac{}{\Gamma \vdash 0 : \mathcal{N}} \quad \frac{\Gamma \vdash n : \mathcal{N}}{\Gamma \vdash S(n) : \mathcal{N}}$$

$$\frac{\Gamma \vdash n : \mathcal{N} \quad \Gamma \vdash a : A \quad \Gamma, \mathbf{x}:A \vdash r : A}{\Gamma \vdash R(n, a, \mathbf{x}.r)}$$

The type \mathcal{N} is intended to represent the natural numbers, with 0 and S being zero and the successor function. The operator R gives definitions by recursion: $R(n, a, \mathbf{x}. f \mathbf{x})$ is intended to be $f^n(a)$. This is expressed by the following reduction rules:

$$R(0, a, \mathbf{x}. r) \longrightarrow a \quad R(S(n), a, \mathbf{x}. r) \longrightarrow r[R(n, a, \mathbf{x}. r)/\mathbf{x}].$$

Adapt the proof of theorem 3.8 to prove SN for Gödel's T .

3.3 A Combinatorial Proof

We give another proof of SN for λ^\rightarrow that is less abstract in the sense that it doesn't require the predicates Red_A . The combinatorial methods we use here are less general—for instance, cannot be applied to Gödel's T —but can be formalised in a logic of first order arithmetic, such as Peano Arithmetic. The technique is in essence that used by Gentzen to prove cut-elimination for LK, the classical sequent calculus. Later, we shall give an even less abstract proof, one that is formalisable in primitive recursive arithmetic.

We shall call a term **inductively normalising** (IN) if it is derivable by the three clauses (1). By lemma 3.2, the inductively normalising terms are strongly normalising.

We shall call a type A **substitutive** if IN is closed under the following clause:

$$\frac{\Gamma \vdash r : A \quad \Gamma \mathbf{x}:A \Delta \vdash s : B \quad r \in \text{IN} \quad s \in \text{IN}}{s[r/\mathbf{x}] \in \text{IN}} \quad \forall \Gamma, r, \mathbf{x}, \Delta, s, B. \quad (2)$$

The type A is called **applicative** if IN is closed under:

$$\frac{\Gamma \vdash t : A \rightarrow B \quad \Gamma \vdash r : A \quad t \in \text{IN} \quad r \in \text{IN}}{tr \in \text{IN}} \quad \forall \Gamma, t, B, r. \quad (3)$$

Lemma 3.10 *Every substitutive type is applicative.*

PROOF: Let A be a substitutive type. We show that IN is closed under (3) by induction on the fact that $t \in \text{IN}$.

If $t = \mathbf{x} r_1 \dots r_m$ with $r_1, \dots, r_m \in \text{IN}$, then as $r \in \text{IN}$ also, we have $tr = \mathbf{x} r_1 \dots r_m r \in \text{IN}$.

If $t = \lambda \mathbf{x}:A. s$ with $s \in \text{IN}$, then as A is substitutive, we have $s[r/\mathbf{x}] \in \text{IN}$ and hence $tr = (\lambda \mathbf{x}:A. s) r \in \text{IN}$.

If $t = (\lambda \mathbf{z}:C. s) r_0 \dots r_m$ where $s[r_0/\mathbf{z}] r_1 \dots r_m \in \text{IN}$ and $r_0 \in \text{IN}$, then by the induction hypothesis, $s[r_0/\mathbf{z}] r_1 \dots r_m r \in \text{IN}$. As also $r_0 \in \text{IN}$, we get $tr = (\lambda \mathbf{z}:C. s) r_0 \dots r_m r \in \text{IN}$. \square

Lemma 3.11 *Suppose that $A = A_1 \rightarrow \dots \rightarrow A_M \rightarrow o$, where each A_i is applicative. Then A is substitutive.*

PROOF: We show that IN is closed under (2) by induction on the fact that $s \in \text{IN}$.

If $s = \mathbf{x} s_1 \dots s_m$ where $s_1, \dots, s_m \in \text{IN}$, then by the induction hypothesis, $s_i[r/\mathbf{x}] \in \text{IN}$ ($1 \leq i \leq m$). The term $s_i[r/\mathbf{x}]$ has type A_i , which is applicative, so that $s[r/\mathbf{x}] = r s_1[r/\mathbf{x}] \dots s_m[r/\mathbf{x}] \in \text{IN}$.

If $s = \mathbf{y} s_1 \dots s_m$ where $s_1, \dots, s_m \in \text{IN}$ and $\mathbf{y} \neq \mathbf{x}$, then by the induction hypothesis, $s_i[r/\mathbf{x}] \in \text{IN}$ ($1 \leq i \leq m$). Then $s[r/\mathbf{x}] = \mathbf{y} s_1[r/\mathbf{x}] \dots s_m[r/\mathbf{x}] \in \text{IN}$.

If $s = \lambda \mathbf{y}:B. s'$ where $s' \in \text{IN}$, then by the induction hypothesis, $s'[r/\mathbf{x}] \in \text{IN}$ so that $s[r/\mathbf{x}] = \lambda \mathbf{y}:B. s'[r/\mathbf{x}] \in \text{IN}$.

If $s = (\lambda \mathbf{y}:B. s') r_0 \dots r_m$ where $s'[r_0/\mathbf{y}] r_1 \dots r_m \in \text{IN}$ and $r_0 \in \text{IN}$, then by the induction hypothesis,

$$s'[\frac{r}{\mathbf{x}}][r_0[\frac{r}{\mathbf{x}}]/\mathbf{y}] (r_1[\frac{r}{\mathbf{x}}]) \dots (r_m[\frac{r}{\mathbf{x}}]) = (s'[r_0/\mathbf{y}] r_1 \dots r_m)[r/\mathbf{x}] \in \text{IN},$$

and $r_0[r/\mathbf{x}] \in \text{IN}$, so that $s[r/\mathbf{x}] = (\lambda \mathbf{y}:B. s'[\frac{r}{\mathbf{x}}])(r_0[\frac{r}{\mathbf{x}}]) \dots (r_m[\frac{r}{\mathbf{x}}]) \in \text{IN}$. \square

Theorem 3.12 *Every term of λ^\rightarrow is SN.*

PROOF: An induction on types using lemmas 3.10 and 3.11 shows that every type is applicative. Together with the definition of IN, this shows that IN is closed under all the typing rules of λ^\rightarrow , so that induction on the derivation of $\Gamma \vdash s : B$ gives $s \in \text{IN}$. \square

4 Quantitative Normalisation

In this section we give a *quantitative* analysis of normalisation in λ^\rightarrow . We shall give an upper bound for these questions: given a term t , how many reduction steps need to be carried out to reach normal form? how big is the normal form? The results we give were originally proven by Statman.

Later, we will also give lower bounds, showing that any algorithm that gives any information about the normal form of a term t , has a run-time comparable to that needed to reduce terms to normal form.

4.1 Complete Developments

Our analysis starts with the following notion of the **complete development** $\mathcal{D}(t)$ of a term t . Intuitively, $\mathcal{D}(t)$ is given by reducing every redex in t , but not any of the new redexes created by reducing the original redexes. Formally,

$$\mathcal{D}(\mathbf{x}) = \mathbf{x} \quad \mathcal{D}(\lambda \mathbf{x}:A. s) = \lambda \mathbf{x}:A. \mathcal{D}(s) \quad \mathcal{D}((\lambda \mathbf{x}:A. s) r) = \mathcal{D}(s)[\mathcal{D}(r)/\mathbf{x}]$$

and

$$\mathcal{D}(\mathbf{x} r) = \mathbf{x} \mathcal{D}(r) \quad \mathcal{D}(t s r) = \mathcal{D}(t s) \mathcal{D}(r).$$

The last two clauses are equivalent to $\mathcal{D}(t r) = \mathcal{D}(t) \mathcal{D}(r)$ for t not a λ -abstraction.

Exercise 4.1 Show that if $r \rightarrow_{\beta_1} s$, then $s \rightarrow_{\beta} \mathcal{D}(r) \rightarrow_{\beta} \mathcal{D}(s)$. Show that if $r \rightarrow_{\beta_n} s$, then $s \rightarrow_{\beta} \mathcal{D}^n(s)$, and (a) give another proof of the Church-Rosser property, and (b) show that if r is weakly normalising, then $\mathcal{D}^N(r)$ is normal for some r .

4.2 Redex Depths

Given $\Gamma \vdash s : B$, and a sub-term occurrence t in the term s , there is a corresponding sequent in the form $\Gamma \Delta \vdash t : C$ in the (unique) derivation of $\Gamma \vdash s : B$. We call C the **type** of t (or $t : C$).

Exercise 4.2 Formalise the notion of the sequent corresponding to an occurrence above. Show that for an α -class s , and an occurrence in $s_0 \in s$, that the type of the occurrence does not depend on s_0 .

Given a redex $(\lambda x:A. s) r$ in t , its **redex type** is the type of $\lambda x:A. s$.⁴

If $t \rightarrow_{\beta_1} t'$ by reducing a redex $(\lambda x:A. s) r$ with redex type $A \rightarrow B$, then the redexes in t' can be classified as follows:

1. Those redexes in t' which correspond to redexes in t .
2. If s is in the form $\lambda x':A'. s'$, and the redex occurs in a position in the form $(\lambda x:A. s) r r'$ in t , then $s[r/x] r' = (\lambda x':A'. s'[r/x]) r'$ is a redex in t' , with redex type B .
3. If r is in the form $\lambda z:C. r'$, and there is an occurrence $x s'$ in s , then there is the redex $r s'[r/x] = (\lambda z:C. r')(s'[r/x])$ in t' , with redex type A .

The redex types in t' are thus among: 1. the redex types of redexes in t , 2. the type A , and 3. the type B . Considering instead $\mathcal{D}(t)$, we don't get any redexes of class 1., and so:

Lemma 4.3 Let $\Gamma \vdash t : C$. The redex type of a redex in $\mathcal{D}(t)$ is either A or B , where $A \rightarrow B$ is a redex type in t .

PROOF: A redex in $s[r/x]$ has redex type which is either a redex type in s or r , or type A , where $r : A$.

We use induction over the clauses defining \mathcal{D} . Only two of the clauses can create new redexes. For, $\mathcal{D}((\lambda x:A. s) r) = \mathcal{D}(s)[\mathcal{D}(r)/x]$, the redexes of $\mathcal{D}((\lambda x:A. s) r)$ are either from $\mathcal{D}(s)$, from $\mathcal{D}(r)$, or are created by the substitution and have redex type A . For $\mathcal{D}(t s r) = \mathcal{D}(t s) \mathcal{D}(r)$, each redex

⁴This is confusing, because the redex is a sub-term, and the type of the sub-term is not the redex type.

of $\mathcal{D}(t s r)$ is either from $\mathcal{D}(t s)$, $\mathcal{D}(r)$, or is $\mathcal{D}(t s) \mathcal{D}(r)$. In the latter case, $\mathcal{D}(t s)$ must be λ -abstraction of some type B , so that $t s$ must be a redex with redex type $A \rightarrow B$. \square

The **depth** $d(A)$ of a type A is defined by $d(o) = 0$ and $d(A \rightarrow B) = 1 + \max(d(A), d(B))$. The **redex depth** of a term t is the maximum of the depths of the redex types in t . The lemma above gives us another proof that every term is weakly normalising:

Corollary 4.4 *Let $\Gamma \vdash s : B$, and s have redex depth N . Then $\mathcal{D}^N(s)$ is normal.*

4.3 Calculating Upper Bounds

Recall that the **length** $l(s)$ of a term is defined by $l(\mathbf{x}) = 1$, $l(\lambda \mathbf{x}. A. s) = 1 + l(s)$ and $l(t r) = 1 + l(t) + l(r)$. We also define the **depth** $d(s)$ of a term s by $d(\mathbf{x}) = 1$, $d(\lambda \mathbf{x}. A. s) = 1 + d(s)$ and $d(t r) = 1 + \max(d(t), d(r))$.

Lemma 4.5 1. $l(s) < 2^{d(s)}$. $d(s[r/\mathbf{x}]) < d(s) + d(r)$.

2. $d(\mathcal{D}(s)) \leq l(s)$. $l(\mathcal{D}(s)) \leq 2^{l(s)}$.

3. If s has m redexes, then $s \rightarrow_{\beta} \mathcal{D}(s)$ by a sequence of m one-step reductions.

4. s has at most $\frac{l(s)}{2}$ redexes. \square

Exercise 4.6 *Prove the above.*

We define the **repeated exponential** 2_N^m by $2_0^m = m$ and $2_{N+1}^m = 2^{2_N^m}$.

Theorem 4.7 *For $\Gamma \vdash s : B$, let N be the redex depth of s . Then the normal form of s has length $\leq 2_N^{l(s)}$, and s (if not normal) has a reduction sequence to normal form with length $\leq 2_{N-1}^{l(s)}$.*

PROOF: By corollary 4.4, $\mathcal{D}^N(s)$ is the normal form of s . That $l(\mathcal{D}^N(s)) \leq 2_N^{l(s)}$ follows by induction using 2. of the previous lemma.

We show that $s \rightarrow_{\beta} \mathcal{D}^N(s)$ in $\leq 2_{N-1}^{l(s)}$ steps by induction on N . For $N = 1$, we have that $s \rightarrow_{\beta} \mathcal{D}(s)$ in $\leq \frac{1}{2}l(s) \leq 2_0^{l(s)}$ steps. By 3. and 4. of the lemma above, we have that $\mathcal{D}^N(s) \rightarrow_{\beta} \mathcal{D}^{N+1}(s)$ in $\leq \frac{1}{2}l(\mathcal{D}^N(s)) \leq \frac{1}{2}2_N^{l(s)}$ steps. Then $s \rightarrow_{\beta} \mathcal{D}^N(s) \rightarrow_{\beta} \mathcal{D}^{N+1}(s)$ in at most

$$2_{N-1}^{l(s)} + \frac{1}{2}2_N^{l(s)} = \frac{1}{2}(2 \cdot 2_{N-1}^{l(s)} + 2_N^{l(s)}) \leq \frac{1}{2}(2^{2_{N-1}^{l(s)}} + 2_N^{l(s)}) = 2_N^{l(s)}$$

steps. □

Exercise 4.8 The **rank** of a type is defined by $\text{rank}(o) = 0$, $\text{rank}(A \rightarrow B) = \max(\text{rank}(A) + 1, \text{rank}(B))$. The **redex rank** of a term is the maximum of the ranks of the redex types in the term. Modify the definition of \mathcal{D} so that

$$\mathcal{D}((\lambda \mathbf{x}_1 \dots \lambda \mathbf{x}_m. s) r_1 \dots r_m) = \mathcal{D}(s) [\mathcal{D}(r_1) / \mathbf{x}_1 \dots \mathcal{D}(r_m) / \mathbf{x}_m].$$

Refine theorem 4.7 to have N the redex rank instead of the redex depth.

5 Quantitative Strong Normalisation

In the previous section we gave an upper bound for the lengths of reduction sequences for a particular reduction strategy. In this section we will give an upper bound for $\text{sn}(t)$, the maximum lengths of any reduction sequences starting from t .

We will do this by a method of translations. Give a term t , we will find a new term $|t| : o$, such that $\text{sn}(t)$ is in some sense computed by $|t|$. Combined with the results of the previous section, this will give an upper bound for $\text{sn}(t)$. In order to make the translation work, given a term f of function type $A \rightarrow B$, we need not only $|f|$, but also a term mapping $|r|$ to $|f r|$. We do this by first defining functionals $\llbracket r \rrbracket : A$ whenever $r : A$, and operators $\aleph_A : A \rightarrow o$ mapping $\llbracket r \rrbracket$ to $|r|$. The technique is due to Loader, who used it to prove normalisation results. Gandy earlier gave a somewhat similar translation, in order to derive strong normalisation from weak.

5.1 Translations

We introduce quite a lot of notation. Take two arbitrary variables Σ and C . Let Ξ be the context $\Sigma : o \rightarrow o \rightarrow o \ C : o$. We define operators $\Sigma_A(\cdot, \cdot)$ and $\aleph_A(\cdot)$, and terms C_A by recursion on types A :

$$\Sigma_o(i, a) = \Sigma i a, \quad \aleph_o(i) = i, \quad C_o = C$$

and, taking \mathbf{x} to avoid variable capture,

$$\Sigma_{A \rightarrow B}(i, f) = \lambda \mathbf{x} : A. \Sigma_B(i, f \mathbf{x}),$$

$$\aleph_{A \rightarrow B}(f) = \aleph_B(f C_A).$$

$$C_{A \rightarrow B} = \lambda \mathbf{x} : A. \Sigma_B(\aleph_{A \mathbf{x}}, C_B),$$

We have the following derived rules:

$$\frac{\Gamma \vdash i : o \quad \Gamma \vdash r : A}{\Xi \Gamma \vdash \Sigma_A(i, r) : A} \quad \frac{\Gamma \vdash r : A}{\Xi \Gamma \vdash \aleph_A(r) : A} \quad \overline{\Xi \Gamma \vdash C_A : A}.$$

Given a term r , we define $\langle r \rangle$ by recursion on r :

$$\langle \mathbf{x} \rangle = \mathbf{x} \quad \langle t r \rangle = \langle t \rangle \langle r \rangle \quad \langle \lambda \mathbf{x} : A. s \rangle = \lambda \mathbf{x} : A. \Sigma_A(\aleph_A \mathbf{x}, \langle s \rangle).$$

This gives us the following derived rule:

$$\frac{\Gamma \vdash s : B}{\Xi \Gamma \vdash \langle s \rangle : B}.$$

We also have that $\langle \cdot \rangle$ commutes with substitution:

$$\langle s[r/\mathbf{x}] \rangle = \langle s \rangle [\langle r \rangle / \mathbf{x}].$$

Given $\mathbf{x}_1 : A_1 \dots \mathbf{x}_m : A_m \vdash s : B$, we define

$$|s| = \aleph_B \langle s \rangle [C_{A_1} / \mathbf{x}_1 \dots C_{A_m} / \mathbf{x}_m],$$

so that $\Xi \vdash |s| : o$.

5.2 Reduction of Translations

The relationship between the term $|s|$ and the length of reduction sequences starting from s is derived by examining the behaviour of $|s|$ under reduction.

Lemma 5.1

$$\aleph_A(\Sigma_A(i, r)) \longrightarrow_{\beta} \Sigma i (\aleph_A r) \tag{4}$$

$$|\lambda \mathbf{x} : A. s| \longrightarrow_{\beta} \Sigma (\aleph_A C_A) |s| \tag{5}$$

If $\Gamma \vdash \mathbf{x} : A_1 \rightarrow \dots \rightarrow A_m \rightarrow B$ and $\Gamma \vdash r_i : A_i$ ($1 \leq i \leq m$), then

$$|\mathbf{x} r_1 \dots r_m| \longrightarrow_{\beta} \Sigma |r_1| (\dots (\Sigma |r_m| (\aleph_B C_B)) \dots). \tag{6}$$

$$|(\lambda \mathbf{x} : A. s) r_0 \dots r_m| \longrightarrow_{\beta} \Sigma |r_0| |s[r_0/\mathbf{x}] r_1 \dots r_m| \tag{7}$$

PROOF: For (4), we use induction on types. The base case holds by definition, and the induction step is

$$\begin{aligned} \aleph_{A \rightarrow B}(\Sigma_{A \rightarrow B}(i, f)) &= \aleph_B(\Sigma_{A \rightarrow B}(i, f) C_A) \\ &\longrightarrow_{\beta} \aleph_B(\Sigma_B(i, f C_A)) \\ &\longrightarrow_{\beta} \Sigma i (\aleph_B(f C_A)) = \Sigma i (\aleph_{A \rightarrow B} f) \end{aligned}$$

We denote the substitution in the definition of $|\cdot|$ by σ , so that $|s| = \aleph_B(|s|)[\sigma]$. For (5), we have

$$\begin{aligned}
|\lambda \mathbf{x}:A. s| &= \aleph_{A \rightarrow B}(\lambda \mathbf{x}:A. \Sigma_B(\aleph_{A\mathbf{x}}, (|s|))[\sigma]) \\
&\longrightarrow_{\beta} \aleph_B(\Sigma_B(\aleph_{AC_A}, (|s|)[\sigma, C_A/\mathbf{x}])) \\
&\longrightarrow_{\beta} \Sigma(\aleph_{AC_A})(\aleph_B(|s|)[\sigma, C_A/\mathbf{x}]) \\
&= \Sigma(\aleph_{AC_A})|s|.
\end{aligned}$$

For (6),

$$\begin{aligned}
|\mathbf{x} r_1 \dots r_m| &= \aleph_B(C_A(|r_1|)[\sigma] \dots (|r_m|)[\sigma]) \\
&\longrightarrow_{\beta} \aleph_B(\Sigma_B(\aleph_{A_1}(|r_1|)[\sigma], \dots, \Sigma_B(\aleph_{A_m}(|r_m|)[\sigma], C_B) \dots)) \\
&= \aleph_B(\Sigma_B(|r_1|, \dots, \Sigma_B(|r_m|, C_B) \dots)) \\
&\longrightarrow_{\beta} \Sigma|r_1|(\dots(\Sigma|r_m|(\aleph_B C_B) \dots)),
\end{aligned}$$

where the last step uses (4) m times. For (7),

$$\begin{aligned}
(|(\lambda \mathbf{x}:A. s) r_0 \dots r_m|) &= (\lambda \mathbf{x}:A. \Sigma_C(\aleph_{A\mathbf{x}}, (|s|)))(|r_0|) \dots (|r_m|) \\
&\longrightarrow_{\beta} \Sigma_C(\aleph_A(|r_0|), (|s|)[(|r_0|/\mathbf{x})])(|r_1|) \dots (|r_m|) \\
&\longrightarrow_{\beta} \Sigma_B(\aleph_A(|r_0|), (|s|)[(|r_0|/\mathbf{x})])(|r_1|) \dots (|r_m|) \\
&= \Sigma_B(\aleph_A(|r_0|), (|s[r_0/\mathbf{x}] r_1 \dots r_m|)).
\end{aligned}$$

where $C = A_1 \rightarrow \dots \rightarrow A_m \rightarrow B$. Applying the definition $|t| = \aleph_B(|t|)[\sigma]$ to the above gives:

$$\begin{aligned}
(|(\lambda \mathbf{x}:A. s) r_0 \dots r_m|) &\longrightarrow_{\beta} \aleph_B(\Sigma_B(\aleph_A(|r_0|), (|s[r_0/\mathbf{x}] r_1 \dots r_m|)))[\sigma] \\
&\longrightarrow_{\beta} \Sigma(\aleph_A(|r_0|)[\sigma])(\aleph_B(|s[r_0/\mathbf{x}] r_1 \dots r_m|)[\sigma]) \\
&= \Sigma|r_0||s[r_0/\mathbf{x}] r_1 \dots r_m|
\end{aligned}$$

as required. \square

For us, the crucial facts in the lemma above are that $|\mathbf{x} r_1 \dots r_m|$ is expressed in terms of the $|r_i|$, $|\lambda \mathbf{x}:A. s|$ in terms of $|s|$, and $|\lambda \mathbf{x}:A. s) r_0 \dots r_m|$ in terms of $|r_0|$ and $|s[r_0/\mathbf{x}] r_1 \dots r_m|$. This will enable us to relate the size of the normal form of $|t|$ to the number of instances of the clauses of (1) needed to derive t , and hence, by exercise 3.3, to $\text{sn}(t)$.

5.3 Translations and Strong Normalisation

Lemma 5.2 *Let μ be a natural number valued function defined on those terms n with $\Xi \vdash n : o$, such that (a) if $n \longrightarrow_{\beta} n'$, then $\mu(n) = \mu(n')$, and (b) $\mu(\Sigma n m) \geq \mu(n) + \mu(m) + 1$. Then $\text{sn}(t) \leq \mu|t|$ for any term t .*

PROOF: By exercise 3.3, it suffices to show that

$$\mu|\mathbf{x} r_1 \dots r_m| \geq \mu|r_1| + \dots + \mu|r_m| + 1, \quad \mu|\lambda\mathbf{x}:A. s| \geq \mu|s| + 1$$

and

$$\mu|(\lambda\mathbf{x}:A. s) r_0 \dots r_m| \geq \mu|s[r_0/\mathbf{x}] r_1 \dots r_m| + |r_0| + 1.$$

These three inequalities follow immediately from the two properties of μ and 2., 3. and 4. of lemma 5.1. \square

One such μ is given by letting $\mu(s)$ be the length of the normal form of s . Combined with theorem 4.7, this gives:

Theorem 5.3 *There is a constant K such that, for any term s of λ^\rightarrow , we have $\text{sn}(s) \leq 2_N^{K \cdot M \cdot l(s)}$, where M (resp. N) is the maximum of the lengths (resp. depths) of the types of the sub-terms of s .*

PROOF: By inspection, there is K such that $l|s| \leq K \cdot M \cdot l(s)$. The redexes in $|s|$ are sub-types of the types of sub-terms of s . Letting μ be the length of normal form function, and applying theorem 4.7 and lemma 5.2 gives the result. \square

Exercise 5.4 Let \mathcal{M} be the model of λ^\rightarrow with $\mathcal{M}(o)$ the set of natural numbers. By choosing a suitable valuation on Ξ , reprove strong normalisation using lemma 5.2. This is a quite general technique for proving strong normalisation.

Exercise 5.5 Refine theorem 5.3 to give $\text{sn}(s) \leq 2_N^{K \cdot l(s)}$, where N is redex rank of s . Hints:

The translation we used is more sophisticated than is required if we just wish to bound $\text{sn}(t)$. Show that taking the C_A to be variables, and $\mathfrak{N}_A(r) = C_{A \rightarrow o} r$, in fact gives $\text{sn}(t)$ bounded by the length the normal form of $\langle t \rangle$.

Additionally, modify $\Sigma_{A \rightarrow B}$ so that $\Sigma_{A \rightarrow B}(i, \lambda\mathbf{x}:A. s) = \lambda\mathbf{x}:A. \Sigma_B(i, s)$ (the original definition β -reduces to this). Also, check that, if the t for which we wish to bound $\text{sn}(t)$ has no sub-term of type A , then we can take $\Sigma_{A \rightarrow B}(i, f) = f$.

These modifications ensure that the types of redexes in $\langle t \rangle$ are exactly the types of redexes in t , so that the N of our bound can be taken to be the redex rank of t (using exercise 4.8).

To show that the exponent $K \cdot M \cdot l(t)$ can be reduced to $K \cdot l(t)$, carefully β -expand (i.e., replace $s[r/\mathbf{x}]$ with $(\lambda\mathbf{x}:A. s)r$) the term $\langle t \rangle$. Doing this correctly should give a term with length bounded by $K \cdot l(t)$, and hence give the bound $\text{sn}(t) \leq 2_N^{K \cdot l(t)}$.

6 Lower Bounds for Normalisation

The upper bounds given in the previous two sections are not too bad, as the following example shows.

For any number n and type A , let n_A be the term $\lambda f:A \rightarrow A. \lambda x:A. f^n x$ of type $(A \rightarrow A) \rightarrow A \rightarrow A$. It is easy to show that

$$m_{A \rightarrow A} n_A = (n^m)_A.$$

Let the term $\text{big}_{A,n}$ be given by

$$\text{big}_{A,0} = 0_A \quad \text{big}_{A,n+1} = \text{big}_{A \rightarrow A,n} 2_A.$$

By induction on n , we see that

$$\text{big}_{A,n} \longrightarrow_{\beta} (2_n^0)_A.$$

The term $\text{big}_{o,n}$ has redex rank $n + 2$, and length $O(n)$, so that the upper bound on the size of the normal form of $\text{big}_{o,n}$ given in section 4 is $2_{n+2}^{O(n)}$.

A program P that takes a natural number as input is called **elementary recursive** (or Kalmar elementary) if there is K such that $P(n)$ has run time $\leq 2_K^n$. A function f is called **elementary recursive** if it is computed by an elementary recursive program.⁵

The example above shows that the function computing the normal form of a typed λ -term is not elementary recursive, while the bounds of the previous sections show that that function is only just outside of elementary recursive.

In the following, we shall prove stronger versions of this, showing that any function giving non-trivial information about normal forms of terms is not elementary recursive.

For example, consider the closed terms of type $o \rightarrow o \rightarrow o$. There are just two closed normal forms of this type: $\lambda x:o. \lambda y:o. x$ and $\lambda x:o. \lambda y:o. y$. Computing the normal form of a closed term of type $o \rightarrow o \rightarrow o$ by β -reduction is not elementary recursive. However, this leaves the possibility that there is an elementary recursive program computing the same function by a different method. We shall show that this is not the case:

Theorem 6.1 *There is no elementary recursive program computing the normal form of closed terms of type $o \rightarrow o \rightarrow o$. \square*

This extends easily to give the no-non-trivial-information result:

⁵Note that due to the size of the functions 2_K^n , it is pretty much irrelevant what formalism is used to express programs and their run-times—even an exponential difference in efficiency between two formalisms can be accommodated by changing K .

Corollary 6.2 *Suppose that $\Gamma \vdash s_1 : A$, $\Gamma \vdash s_2 : A$. Let P be a program such that if $\Gamma \vdash t : A$ and $t \longrightarrow_{\beta} s_i$ then $P(t)$ returns i ($i = 1, 2$). Then P is not elementary recursive.*

PROOF (of corollary from theorem): Let P be a such a program. We find a program Q computing the normal of closed terms of type $o \rightarrow o \rightarrow o$ as follows.

Define the type substitution $B[A/o]$ by $o[A/o] = A$, $(C \rightarrow D)[A/o] = C[A/o] \rightarrow D[A/o]$. Extend to terms by performing the type substitution in each λ -abstraction. Define $Q(t)$ to return $\lambda x:o. \lambda y:o. x$ if $P(t[A/o] s_1 s_2)$ returns 1, and to return $\lambda x:o. \lambda y:o. y$ if $P(t[A/o] s_1 s_2)$ returns 2.

Clearly, if $\vdash t : o \rightarrow o \rightarrow o$, then $Q(t)$ computes the normal form of t , and by theorem 6.1, cannot be elementary recursive. The run-time of Q is only a little more than that of P : if P were elementary recursive then so would be Q . Hence P is not elementary recursive. \square

6.1 Higher Order Propositional Classical Logic

Our proof of theorem 6.1 will proceed via encoding **higher order propositional classical logic** (PCL_{ω}). The language of PCL_{ω} is as follows. There are countably many variables $\mathcal{V}^N = \{x_0^N, x_1^N, \dots\}$ of each rank N . Atomic formulas are x_i^0 and $x_i^N \in x_j^{N+1}$, connectives are \vee (binary) and \neg (unary), and we have quantifiers \exists^N of each rank. As a grammar:

$$\mathcal{F} := \mathcal{V}^0 \mid \mathcal{V}^N \in \mathcal{V}^{N+1} \mid \mathcal{F} \vee \mathcal{F} \mid \neg \mathcal{F} \mid \exists^N \mathcal{V}^N. \mathcal{F}.$$

The logic is classical, so we define other connectives in the usual way, e.g.,

$$\begin{aligned} \phi \wedge \psi &:= \neg(\neg\phi \vee \neg\psi) & \phi \Rightarrow \psi &:= \neg\phi \vee \psi \\ \phi \Leftrightarrow \psi &:= (\phi \Rightarrow \psi) \wedge (\psi \Rightarrow \phi) & \phi \oplus \psi &:= (\phi \wedge \neg\psi) \vee (\neg\phi \wedge \psi) \\ x^N \notin y^{N+1} &:= \neg(x^N \in y^{N+1}) & \forall^N x^N. \phi &:= \neg\exists^N x^N. \neg\phi. \end{aligned}$$

The semantics of PCL_{ω} is given as follows. Let $\mathcal{B} = \{\mathbf{tt}, \mathbf{ff}\}$, and let $\mathcal{B}_0 = \mathcal{B}$, $\mathcal{B}_{N+1} = \mathcal{B}^{\mathcal{B}_N}$. Then the rank N variables x^N are interpreted as ranging over \mathcal{B}_N . The atomic formula $x^N \in y^{N+1}$ is interpreted as $y^{N+1}(x^N)$. The connectives are interpreted in the usual way.

Exercise 6.3 *Formalise the semantics.* Define a **valuation** to be a function σ mapping each variable x^N into \mathcal{B}_N , and then define $\llbracket \phi \rrbracket_{\sigma} \in \mathcal{B}$ for formulae ϕ and valuations σ .

Show that $\llbracket \phi \rrbracket_{\sigma}$ depends only on ϕ and σx for variables x free in ϕ . In particular, $\llbracket \phi \rrbracket$ is well-defined for closed ϕ .

The set \mathcal{B}_N has 2_N^2 elements, so that the obvious algorithm computing $\llbracket \phi \rrbracket$ is non-elementary. In the appendix, we prove that:

Theorem 6.4 *The function $\llbracket \phi \rrbracket$ for closed ϕ is not elementary.* \square

We shall encode \mathcal{B}_N in λ^\rightarrow , in such a way that we can give a reasonably short encoding of iteration over the members of the \mathcal{B}_N , and then of the semantics $\llbracket \phi \rrbracket$. The function giving the encoding will be elementary recursive, which together with theorem 6.4 will give a proof of theorem 6.1.

6.2 Encoding the Boolean Hierarchy

We encode the elements of the sets \mathcal{B}_N as closed terms of λ^\rightarrow . The encoding is 1-many rather than 1-1. We define the types $\mathcal{B} = o \rightarrow o \rightarrow o$, and $\mathcal{B}_0 = \mathcal{B}$, $\mathcal{B}_{N+1} = \mathcal{B}_N \rightarrow \mathcal{B}$. The intention is that elements of the set \mathcal{B}_N will be represented as closed type \mathcal{B}_N terms. We define the relations \sim_N as follows.

For $a \in \mathcal{B}$ and $\Gamma \vdash r : \mathcal{B}$, put $a \sim r$ iff either $a = \mathbf{tt}$ and $r \rightarrow_\beta \lambda x:o. \lambda y:o. x$ or $a = \mathbf{ff}$ and $r \rightarrow_\beta \lambda x:o. \lambda y:o. y$.

Let \sim_0 be \sim , and put $b \sim_{N+1} s$ iff $b(a) \sim sr$ whenever $a \sim_N r$. We let $\mathbf{tt}, \mathbf{ff} : \mathcal{B}$ be given by $\mathbf{tt} = \lambda x:o. \lambda y:o. x$ and $\mathbf{ff} = \lambda x:o. \lambda y:o. y$, so that $\mathbf{tt} \sim \mathbf{tt}$ and $\mathbf{ff} \sim \mathbf{ff}$.

Exercise 6.5 *Show that for each $a \in \mathcal{B}_N$ there is r such that $a \sim_N r$.*

If $\vdash R : (\mathcal{B}_N \rightarrow A \rightarrow A) \rightarrow A \rightarrow A$, then R is called an **N - A -iterator** if it β -reduces to

$$\lambda f:\mathcal{B}_N \rightarrow A \rightarrow A. \lambda x:A. f r_0 (\dots (f r_{2_N^2-1} x) \dots) \quad (8)$$

where the r_i **enumerate** \mathcal{B}_N , in the sense that for each $a \in \mathcal{B}_N$, there is exactly one i such that $a \sim_N r_i$.

We define $r \vee r' = \lambda x:o. \lambda y:o. r x (r' y)$, $\neg r = \lambda x:o. \lambda y:o. r y x$ for terms $r, r' : o$, so that $r \vee r', \neg r : o$ also. Other Boolean connectives are defined in terms of these.

Given a N - \mathcal{B} -iterator R , we define $\exists^N, \forall^N : \mathcal{B}_{N+1} \rightarrow \mathcal{B}$ by

$$\exists^N = \lambda f:\mathcal{B}_{N+1}. R (\lambda x:\mathcal{B}_N. \lambda y:\mathcal{B}. y \vee f x) \mathbf{ff}$$

and

$$\forall^N = \lambda f:\mathcal{B}_{N+1}. R (\lambda x:\mathcal{B}_N. \lambda y:\mathcal{B}. y \wedge f x) \mathbf{tt}.$$

We will write \exists_R^N and \forall_R^N when it is desired to emphasise the iterator.

Exercise 6.6 Let R be a N - \mathcal{B} -iterator. Show that if $a \sim_{N+1} r$ then $\mathbf{tt} \sim \exists_R^N r$ iff $a(b) = \mathbf{tt}$ for some $b \in \mathcal{B}_N$, and $\mathbf{ff} \sim \exists_R^N r$ iff $a(b) = \mathbf{ff}$ for all $b \in \mathcal{B}_N$. Show the corresponding property of \forall_R^N .

An N -**equality test** is a term $\vdash \text{eq} : \mathcal{B}_N \rightarrow \mathcal{B}_N \rightarrow \mathcal{B}$ such that if $a \sim r$, $a' \sim r'$, then $\mathbf{tt} \sim \text{eq} r r'$ if $a = a'$, and $\mathbf{ff} \sim \text{eq} r r'$ if $a \neq a'$. Given an N - \mathcal{B} -iterator R , we define an $N+1$ -equality test

$$\text{eq}_R = \lambda x, y : \mathcal{B}_{N+1}. \forall_R^N (\lambda z : \mathcal{B}_N. x z \Leftrightarrow y z).$$

Exercise 6.7 Prove that eq_R really is an equality test.

Lemma 6.8 Given a type A , let $I_{N,A}$ be the type $(\mathcal{B}_N \rightarrow A \rightarrow A) \rightarrow A \rightarrow A$ of N - A -iterators. Suppose that R is an N - $I_{N+1,A}$ -iterator, and that eq is an N -equality test. Then there is an $N+1$ - A -iterator.

In fact, the existence of N - A -iterators is trivial, since for each $a \in \mathcal{B}_n$ there is r such that $a \sim r$. However, the obvious construction gives iterators of size $O(2_N^2)$ or worse. The construction we give in the proof below will lead to an N - A -iterator that is much smaller.

Observe that a N - I_A -iterator iterates over a $|\mathcal{B}_N|$ element set, while a $N+1$ - A -iterator iterates over a $|\mathcal{B}_{N+1}| = 2^{|\mathcal{B}_N|}$ element set. The idea of the proof below is to define an operator Φ such that if G iterates over a m element set, then ΦG iterates over a $2m$ element set. Then, applying R to Φ , we obtain a term that iterates over a $2^{|\mathcal{B}_N|} = |\mathcal{B}_{N+1}|$ element set.

PROOF: Given a set $S \subset \mathcal{B}_{N+1}$, an S - A -iterator is a term of type $I_{N+1,A}$ which β -reduces to

$$\lambda f : \mathcal{B}_N \rightarrow A \rightarrow A. \lambda x : A. f s_0 (\dots (f s_{m-1} x) \dots)$$

where $s_0 \dots s_{m-1}$ enumerates S .

Given a set $T \subset \mathcal{B}_N$, a T -based A -iterator is a S - A -iterator where

$$S = \{b \in \mathcal{B}_{N+1} \mid b(x) \neq \mathbf{ff} \text{ implies } x \in T\}.$$

A \mathcal{B}_N -based A -iterator is just an $N+1$ - A -iterator. We will find a term

$$\Phi : \mathcal{B}_N \rightarrow I_{N+1,A} \rightarrow I_{N+1,A}$$

such that if $a \sim_N r$, and F is a T -based A -iterator, with $a \notin T$, then $\Phi r F$ is a $T \cup \{a\}$ -based A -iterator.

Given $b \in \mathcal{B}_{N+1}$ and $a \in \mathcal{B}_N$, define $b \cup \{a\} \in \mathcal{B}_{N+1}$ by $(b \cup \{a\})(a) = \mathbf{tt}$, and $(b \cup \{a\})(x) = b(x)$ for $x \neq a$. Given terms $s : \mathcal{B}_{N+1}$ and $r : \mathcal{B}_N$,

define $s \cup \{r\} = \lambda \mathbf{x}:\mathcal{B}_N.(s \mathbf{x}) \vee (\text{eq } r \mathbf{x})$. If $b \sim_{N+1} s$ and $a \sim_N r$, then $b \cup \{a\} \sim_{N+1} s \cup \{r\}$.

If G is an S - A -iterator, for some $S \subset \mathcal{B}_{N+1}$, β -reducing to

$$\lambda \mathbf{f}:\mathcal{B}_{N+1} \rightarrow A \rightarrow A. \lambda \mathbf{x}:A. \mathbf{f} s_0 (\dots (\mathbf{f} s_{m-1} \mathbf{x}) \dots), \quad (9)$$

then

$$\lambda \mathbf{f}:\mathcal{B}_{N+1} \rightarrow A \rightarrow A. \lambda \mathbf{x}:A. G (\lambda \mathbf{y}:\mathcal{B}_{N+1}. \mathbf{f} (\mathbf{y} \cup \{r\})) \mathbf{x} \quad (10)$$

reduces to

$$\lambda \mathbf{f}:\mathcal{B}_{N+1} \rightarrow A \rightarrow A. \lambda \mathbf{x}:A. \mathbf{f} s'_0 (\dots (\mathbf{f} s'_{m-1} \mathbf{x}) \dots), \quad (11)$$

where $s'_i = s_i \cup \{r\}$. If $a \sim r$, then the s'_i enumerate $S' = \{b \cup \{a\} \mid b \in S\}$, so that the term (10) is a S' - A -iterator.

Let Φ be

$$\begin{aligned} & \lambda \mathbf{z}:\mathcal{B}_N. \lambda \mathbf{G}:I_{N+1,A}. \\ & \lambda \mathbf{f}:\mathcal{B}_{N+1} \rightarrow A \rightarrow A. \lambda \mathbf{x}:A. \mathbf{G} (\lambda \mathbf{y}:\mathcal{B}_{N+1}. \mathbf{f} (\mathbf{y} \cup \{\mathbf{z}\})) (\mathbf{G} \mathbf{f} \mathbf{x}). \end{aligned}$$

Given $a \sim_N r$ and an S - A -iterator G with normal form (9), combining (9) and (11) gives that the term $\Phi r G$ reduces to

$$\lambda \mathbf{f}:\mathcal{B}_{N+1} \rightarrow A \rightarrow A. \lambda \mathbf{x}:A. \mathbf{f} s'_0 (\dots (\mathbf{f} s'_{m-1} (\mathbf{f} s_0 (\dots (\mathbf{f} s_{m-1} \mathbf{x}) \dots))) \dots).$$

This shows that $\Phi r G$ is an S'' - A -iterator where $S'' = S \cup S'$. In particular, if G is a T -based A -iterator, then $\Phi r G$ is a $T \cup \{a\}$ - A -iterator, which is the property we desire from Φ .

A \emptyset -based A -iterator is given by

$$G_0 = \lambda \mathbf{f}:\mathcal{B}_{N+1} \rightarrow A \rightarrow A. \lambda \mathbf{x}:A. \mathbf{f} (\lambda \mathbf{z}:\mathcal{B}_N. \mathbf{ff}) \mathbf{x}.$$

Let R' be the term $R \Phi G_0$. As R is an N - I_A -iterator, R' reduces to

$$\Phi r_0 (\dots (\Phi r_{m-1} G_0) \dots)$$

where (for some a_i) $a_i \sim r_i$ and the a_i enumerate \mathcal{B}_N . As $\mathcal{B}_N = \emptyset \cup \{a_{m-1}\} \cup \dots \cup \{a_0\}$, we have that R' is a \mathcal{B}_N -based A -iterator, or in other words, a $N+1$ - A -iterator. \square

Lemma 6.9 *There is an elementary recursive function that given N and A , gives a N - A -iterator.*

PROOF: A 0- A -iterator is given explicitly by

$$\lambda \mathbf{f}:\mathcal{B} \rightarrow A \rightarrow A. \lambda \mathbf{x}:A. \mathbf{f} \mathbf{tt} (\mathbf{f} \mathbf{ff} \mathbf{x})$$

and $\lambda \mathbf{x}:\mathcal{B}. \lambda \mathbf{y}:\mathcal{B}. \mathbf{x} \Leftrightarrow \mathbf{y}$ is a 0-equality test. Induction with lemma 6.8 gives N - A -iterators and N -equality tests for all N . We leave it to the reader to check that these constructions are actually elementary recursive. \square

6.3 Encoding the Logic

We encode formulae of PCL_ω as λ^\top terms, in such a way that the terms compute the semantics. Specifically, if ϕ is a closed formula, then $\llbracket \phi \rrbracket \sim \ulcorner \phi \urcorner$. The encoding is elementary recursive, so this will derive theorem 6.1 from theorem 6.4. For each N , let R_N be a N - \mathcal{B} -iterator as given by 6.9. The iterators R_N give us $\exists^N : \mathcal{B}_{N+1} \rightarrow \mathcal{B}$.

Given a formula ϕ with free variables among $\mathbf{x}_1^{N_1} \dots \mathbf{x}_m^{N_m}$, we define $\ulcorner \phi \urcorner$ such that

$$\mathbf{x}_1 : \mathcal{B}_{N_1} \dots \mathbf{x}_m : \mathcal{B}_{N_m} \vdash \ulcorner \phi \urcorner : \mathcal{B}$$

by induction on ϕ . We let $\ulcorner \mathbf{x} \urcorner = \mathbf{x}$ for any variable \mathbf{x} . Let

$$\begin{aligned} \ulcorner \mathbf{x}_i^0 \urcorner &= \mathbf{x}_i & \ulcorner \mathbf{x}_i^{N_i} \in \mathbf{x}_j^{N_i+1} \urcorner &= \mathbf{x}_j(\mathbf{x}_i) \\ \ulcorner \phi \vee \psi \urcorner &= \ulcorner \phi \urcorner \vee \ulcorner \psi \urcorner & \ulcorner \neg \phi \urcorner &= \neg \ulcorner \phi \urcorner \end{aligned}$$

and

$$\ulcorner \exists^N \mathbf{x}^N . \phi \urcorner = \exists^N (\lambda \mathbf{x} : \mathcal{B}_N . \ulcorner \phi \urcorner).$$

Theorem 6.1 follows by checking that $\ulcorner \phi \urcorner$ is well-behaved:

Exercise 6.10 Suppose that σ is a valuation, such that $\sigma(\mathbf{x}_i^{N_i}) \sim_{N_i} r_i$ ($1 \leq i \leq m$). Show that $\llbracket \phi \rrbracket_\sigma \sim \ulcorner \phi \urcorner[r_1/\mathbf{x}_1 \dots r_m/\mathbf{x}_m]$.

Check that $\ulcorner \phi \urcorner$ is an elementary recursive function of ϕ .

A PCL_ω is not Elementary Recursive

We outline a proof of theorem 6.4 that the valuation $\llbracket \cdot \rrbracket$ for the logic PCL_ω is not elementary recursive. We do this by encoding Turing machines with a bound 2_N^2 on their behaviour.

A.1 Basic Encoding in PCL_ω

We first examine the structure of the sets \mathcal{B}_N . \mathcal{B}_N has 2_N^2 elements. We define a bijection $\text{seq}_N : \{0 \dots 2_N^2 - 1\} \rightarrow \mathcal{B}_N$ in a natural way. Let $\text{seq}(0) = \mathbf{ff}$ and $\text{seq}(1) = \mathbf{tt}$ and $\text{seq}_0 = \text{seq}$. Given seq_N , and a number $\sum_{i < 2_N^2} d_i \cdot 2^i$ in binary notation, put

$$\text{seq}_{N+1} \left(\sum_{i < 2_N^2} d_i \cdot 2^i \right) (\text{seq}_N(m)) = \text{seq}(d_m).$$

This considers elements of \mathcal{B}_{N+1} as numbers in binary notation, the order of the digits given by the numbering of \mathcal{B}_N .

We represent a series of predicates as PCL_ω formulae.

Take two binary numbers $d = \sum_{i < n} d_i \cdot 2^i$ and $e = \sum_{i < n} e_i \cdot 2^i$. We have that $d < e$ iff and only if there is a ‘pivot’ $i_0 < n$ such that $d_{i_0} < e_{i_0}$ and $d_i = e_i$ for $i_0 < i < n$. This inspires the following definition of the formulae $a <_N a'$: let $a <_0 a'$ be $\neg a \wedge a'$, and let $b <_{N+1} b'$ be

$$\exists^N \mathbf{x}. (\neg \mathbf{x} \in b \wedge \mathbf{x} \in b' \wedge \forall^N \mathbf{x}'. (\mathbf{x} <_N \mathbf{x}' \Rightarrow (\mathbf{x}' \in b \Leftrightarrow \mathbf{x}' \in b'))).$$

Defining $S_N(a, a') := \forall^N \mathbf{x}. ((a <_N \mathbf{x}) \oplus (\mathbf{x} <_N a'))$ encodes the **successor** relation $a' = a + 1$.

We turn to encoding combinatorial notions. **Equality** $a =_0 b$, at rank 0, is given by $a \Leftrightarrow b$. **Equality** $a =_{N+1} b$ at rank $N+1$, is given by $\forall^N \mathbf{x}. \mathbf{x} \in a \Leftrightarrow \mathbf{x} \in b$.

We encode **finite sets** as follows. For $a_1 \dots a_m$ rank N and b rank $N+1$, the relation $b = \{a_1 \dots a_m\}$ is given by

$$\forall^N \mathbf{x}. (\mathbf{x} \in b \Leftrightarrow (\mathbf{x} =_N a_1 \vee \dots \vee \mathbf{x} =_N a_m)).$$

The **Kuratowski pairing**, $c = (a, b)$ (where $(a, b) = \{\{a\}, \{a, b\}\}$) is given by

$$\exists^{N+1} \mathbf{x}. \exists^{N+1} \mathbf{y}. (c = \{\mathbf{x}, \mathbf{y}\} \wedge \mathbf{x} = \{a\} \wedge \mathbf{y} = \{a, b\}).$$

where c has rank $N+2$ and a, b have rank N .

We can give an efficient encoding of $a =_N n$, where n is some given number, and a has rank N . $a =_0 0$ is $\neg a$, $a =_0 1$ is a , and $a =_0 n$ is $a \wedge \neg a$ for $n \geq 2$. Given $n = 2^{n_1} + \dots + 2^{n_m}$ with the n_i distinct, $b =_{N+1} n$ is

$$\exists^N \mathbf{x}_1 \dots \exists^N \mathbf{x}_m. (b = \{\mathbf{x}_1 \dots \mathbf{x}_m\} \wedge \mathbf{x}_1 = n_1 \wedge \dots \wedge \mathbf{x}_m = n_m).$$

Finally, $b = \langle a_0 \dots a_{m-1} \rangle$, encoding **small tuples**, is

$$\begin{aligned} & \exists^{N+2} \mathbf{p}_0 \dots \exists^{N+2} \mathbf{p}_{m-1}. \exists^N \mathbf{j}_0 \dots \exists^N \mathbf{j}_{m-1}. \\ & b = \{\mathbf{p}_0 \dots \mathbf{p}_{m-1}\} \wedge \bigwedge_{i < m} \mathbf{p}_i = (a_i, \mathbf{j}_i) \wedge \bigwedge_{i < m} \mathbf{j}_i = i \end{aligned}$$

for b rank $N+3$ and the a_i rank N .

Exercise A.1 Check that all the encodings above do what they should.

Exercise A.2 For a, b, c rank $N+1$, encode the predicates $0 \in a$, $2_{N+1}^2 - 1 \in a$, and $c = a + b$.

A.2 Encoding Turing Machines

We apply the constructs above to encode Turing machines. Since we are encoding in the decidable logic PCL_ω , we cannot encode arbitrary runs of Turing machines; instead we must consider a tape of some finite length, fixed for that run.

We will use a slight variant on the usual notion of Turing machine. Instead of the head being at a particular cell on the tape at each step, the head is between two adjacent cells. The action taken depends on the state, and the two cells adjacent to the head; the actions taken are to move either left or right, updating the cell traversed.

Tapes are encoded as two **tape halves**—one to the left of the head and one to the right of the head. In both directions, positions are numbered by some \mathcal{B}_N , starting with 0 adjacent to the head, so that the two tape halves l and r are encoded by elements of \mathcal{B}_{N+1} . The state s is encoded by another element of \mathcal{B}_{N+1} , and the overall configuration of the machine by $\langle l, s, r \rangle$.

We encode various operations. We give four **adjacency** predicates of a pair of tape halves:

$$\begin{aligned} \text{adj}_{\text{tt},\text{tt}}(l, r) &:= 0 \in l \wedge 0 \in r & \text{adj}_{\text{tt},\text{ff}}(l, r) &:= 0 \in l \wedge 0 \notin r \\ \text{adj}_{\text{ff},\text{tt}}(l, r) &:= 0 \notin l \wedge 0 \in r & \text{adj}_{\text{ff},\text{ff}}(l, r) &:= 0 \notin l \wedge 0 \notin r \end{aligned}$$

These check for the four possibilities of the cells next to the tape head. The **shift** predicate $\text{sh}(t, t')$ is given by

$$(\forall^N i. \forall^N j. S(i, j) \Rightarrow (i \in t \Leftrightarrow j \in t')) \wedge 2_N^2 - 1 \notin t.$$

This shifts a tape half along by one, and checks that at the far end, a **ff** comes into, or goes out of, view.

There are four possible actions on tapes: moving either left or right, while updating the cell traversed with either **tt** or **ff**. These are given by the **action** predicates below:

$$\text{act}_{L,\text{tt}}(l, r, l', r') := \text{sh}(l', l) \wedge 0 \in r' \wedge \text{sh}(r, r')$$

$$\text{act}_{R,\text{tt}}(l, r, l', r') := \text{sh}(l, l') \wedge 0 \in l' \wedge \text{sh}(r', r)$$

$$\text{act}_{L,\text{ff}}(l, r, l', r') := \text{sh}(l', l) \wedge 0 \notin r' \wedge \text{sh}(r, r')$$

$$\text{act}_{R,\text{ff}}(l, r, l', r') := \text{sh}(l, l') \wedge 0 \notin l' \wedge \text{sh}(r', r)$$

A rule of the Turing machine is given by a sextuple $(n, \alpha, \beta, d, \gamma, n')$, which is read as follows: if we are in a configuration with state n and with α and

β to the left and right, respectively, of the head, then move in direction $d \in \{L, R\}$, update the cell traversed with γ , and transfer to state n' . The **update predicate** $\text{tr}_{n,\alpha,\beta,d,\gamma,n'}(l, r, s, l', r', s')$ saying that this rule updates the configuration (l, s, r) to give (l', s', r') is

$$s = n \wedge \text{adj}_{\alpha,\beta}(l, r) \wedge \text{act}_{d,\gamma}(l, r, l', r') \wedge s' = n'.$$

Encoding tuples, we also let $\text{tr}_{n,\alpha,\beta,d,\gamma,n'}(c, c')$ be

$$\exists^{N+1} \mathbf{l} \mathbf{s} \mathbf{r} \mathbf{l}' \mathbf{s}' \mathbf{r}'. c = \langle \mathbf{l}, \mathbf{s}, \mathbf{r} \rangle \wedge c' = \langle \mathbf{l}', \mathbf{r}', \mathbf{s}' \rangle \wedge \text{tr}_{n,\alpha,\beta,d,\gamma,n'}(\mathbf{l}, \mathbf{s}, \mathbf{r}, \mathbf{l}', \mathbf{s}', \mathbf{r}').$$

Taking the disjunction over a set of rules $R = \{\rho_1 \dots \rho_M\}$ gives us the **transition relation** $\text{trans}_R(c, c')$:

$$\text{tr}_{\rho_1}(c, c') \vee \dots \vee \text{tr}_{\rho_M}(c, c').$$

The **initial configuration** of the machine is defined to have **ff** everywhere on the tape, and state 0. Hence $\text{init}(c)$ is

$$\exists^{N+1} \mathbf{x}. \mathbf{x} = 0 \wedge c = \langle \mathbf{x}, \mathbf{x}, \mathbf{x} \rangle.$$

A set S of states is **run closed** if it contains the initial configuration, and whenever $c \in S$, there is $c' \in S$ updated from c by some rule. This is given by $\text{closed}_R(S)$:

$$\exists^{N+4} \mathbf{c}. (\text{init}(\mathbf{c}) \wedge \mathbf{c} \in S) \wedge \forall^{N+4} \mathbf{c}. (\mathbf{c} \in S \Rightarrow \exists^{N+4} \mathbf{c}'. (\mathbf{c}' \in S \wedge \text{trans}(\mathbf{c}, \mathbf{c}'))).$$

There can fail to be a run closed set for a Turing machine under the following two conditions: either the machine reaches a state to which no rule applies, or the machine reaches a state with a **tt** on the tape $2_N^2 + 1$ cells away from the head. In these cases we say the machine **2_N^2 -terminates**, and otherwise the machine **loops**.

Lemma A.3 *There is an algorithm that given a Turing machine and N , decides whether or not the machine 2_N^2 -terminates.*

PROOF: There are only finitely many configurations with no **tt** more than 2_N^2 cells from the head. These can be enumerated, and the machine run until either it terminates, or the machine loops. \square

Exercise A.4 *Check that $\llbracket \exists^{N+5} \mathbf{S}. \text{closed}_R(\mathbf{S}) \rrbracket = \mathbf{ff}$ if and only if R 2_N^2 -terminates.*

This gives a proof of theorem 6.4: if there were an elementary recursive decision procedure for PCL_ω , then the encoding above would contradict the following:

Lemma A.5 *There is no elementary recursive algorithm that given a Turing machine and N , decides whether or not the machine 2_N^2 -terminates.*

PROOF SKETCH: Suppose there was such an algorithm.

Let $P(\rho, c, N)$ to be the predicate of a Turing machine ρ terminating from initial configuration c in time 2_N^2 . Then $P(\rho, c, N)$ is elementary recursive. By the usual sort of diagonalisation argument, this is impossible. \square