Centrum voor Wiskunde en Informatica
REPORT*RAPPORT*

Infinitary lambda calculus

R. Kennaway, J.W. Klop, R. Sleep and F.-J. de Vries

Computer Science/Department of Software Technology

CS-R9535 1995

Report CS-R9535
ISSN 0169-118X

CWI
P.O. Box 94079
1090 GB Amsterdam
The Netherlands

Infinitary Lambda Calculus

Richard Kennaway[1], Jan Willem Klop[2], Ronan Sleep[3] and Fer-Jan de Vries[4]

[1,3] *School of Information Systems, University of East Anglia, Norwich NR4 7TJ, UK*

[2] *Department of Software Technology, CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands*

[4] *NTT Communication Science Laboratories, Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-02, Japan*

[1]jrk@sys.uea.ac.uk, [2]jwk@cwi.nl, [3]mrs@sys.uea.ac.uk, [4]ferjan@cwi.nl

Abstract

In a previous paper we have established the theory of transfinite reduction for orthogonal term rewriting systems. In this paper we perform the same task for the lambda calculus.

From the viewpoint of infinitary rewriting, the Böhm model of the lambda calculus can be seen as an infinitary term model. In contrast to term rewriting, there are several different possible notions of infinite term, which give rise to different Böhm-like models, which embody different notions of lazy or eager computation.

AMS Subject Classification (1991): 03B40
CR Subject Classification (1991): F4.1
Keywords & Phrases: infinitary lambda calculus, Böhm models
Note: The authors were partially supported by SEMAGRAPH II (ESPRIT working group 6345). Richard Kennaway was also supported by a SERC Advanced Fellowship, and by SERC Grant no. GR/F 91582. From July 1995 Fer-Jan de Vries will be at the Hitachi Advanced Research Laboratory, Hatoyama, Saitama 350-03, Japan.

1. Introduction

Infinitary rewriting is a natural generalisation of finitary rewriting which extends it with the notion of computing towards a possibily infinite limit. Such limits naturally arise in the semantics of lazy functional languages, in which it is possible to write and compute with expressions which intuitively denote infinite data structures, such as a list of all the integers. If the limit of a reduction sequence still contains redexes, then it is natural to consider sequences whose length is longer than ω — in fact, sequences of any ordinal length.

Infinitary rewriting also arises from computations with terms implemented as graphs. Such implementations suggest the possibility of using cyclic graphs, which correspond in a natural way to infinite terms. Finite computations on cyclic graphs correspond to infinite computations on terms.

The infinitary theory also suggests new ways of dealing with some of the concepts that arise in the finitary theory, such as notions of undefinedness of terms. In this connection, Berarducci and Intrigila ([Ber, BI94]) have independently developed an infinitary lambda calculus and applied it to the study of consistency problems in the finitary lambda calculus.

In [KKSdV95] we developed the basic theory of transfinite reduction for orthogonal term rewrite systems. In this paper we perform the same task for the lambda calculus. In contrast to the situation for term rewriting, in lambda calculus there turn out to be several different possible domains of infinite terms. These give rise to different Böhm-like models of the calculus corresponding to different notions of laziness.

This paper is a revised version of the paper "Infinitary lambda calculus and Böhm models" appearing in the Proceedings of the Conference on Rewriting Techniques and Applications, Kaiserslautern, April 1995. It differs from that version primarily by a rewriting of section 6.

2. Basic definitions

2.1 *Finitary lambda calculus*

We assume familiarity with the lambda calculus, or as we shall refer to it here, the finitary lambda calculus. [Bar84] is a standard reference. The syntax is simple: there is a set Var of variables; an expression or term E is either a variable, an abstraction $\lambda x.E$ (where x is called the bound variable and E the body), or an application $E_1 E_2$ (where E_1 is called the rator and E_2 the rand). This is the pure lambda calculus — we do not have any built-in constants nor any type system.

As customary, we identify α-equivalent terms with each other, and consider bound variables to be silently renamed when necessary to avoid name clashes.

2.2 *What is an infinite term?*

Drawing lambda expressions as syntax trees gives an immediate and intuitive notion of infinite terms: they are just infinite trees. Formally, we can define this set as the metric completion of the space of finite trees with a well-known (ultra-)metric. The larger the common prefix of two trees, the more similar they are, and the closer together they may be considered to be. First, some terminology. A *position* or *occurrence* is a finite string of positive integers. Given a term M and a position u, the term $M|u$, when it exists, is a subterm of M defined inductively thus:

$$
\begin{aligned}
M|\langle\rangle &= M \\
(\lambda x.M)|1 \cdot u &= M|u \\
(MN)|1 \cdot u &= M|u \\
(MN)|2 \cdot u &= N|u
\end{aligned}
$$

$M|u$ is called the subterm of M at u, and when this is defined, u is called a position of M. The *syntactic depth* of u is its length.

Two positions u and v are *disjoint* if neither is a prefix of the other. Two redexes are disjoint if their positions are. A set of positions or redexes is disjoint if every two distinct members are.

Given two distinct terms M and N, let l be the length of the shortest position u such that $M|u$ and $N|u$ are both defined, and are either of different syntactic types or are distinct variables. Then the larger l is, the more similar are M and N. The distance between M and N is defined to be 2^{-l}. Denote this measure by $d^s(M,N)$. $d^s(M,M)$ is defined to be 0. This is the *syntactic metric*. It is easily proved that it is a metric on the set of finite terms. In fact, it is an ultrametric, i.e. $d^s(M,N) \leq max(d^s(M,P), d^s(P,N))$, although this will not be important. The completion of this metric space adds the infinite terms. We call this set Λ^s.

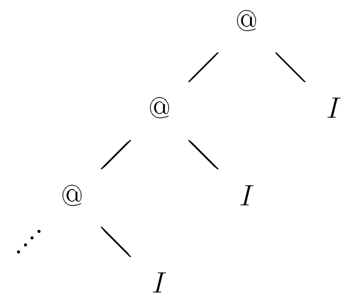

Figure 1.

The above is the definition of infinite terms which we used in our study of transfinite term rewriting, but for lambda calculus the situation is a little more complicated. Consider the term $(((\ldots I)I)I)I$ where $I = \lambda x.x$. See Fig. 1. This term has a combination of properties which is rather strange from the point of view of finitary lambda calculus. By the usual definition of head normal form — being of the form $\lambda x_1 \ldots \lambda x_n.yt_1 \ldots t_m$ — it is not in head normal form. By an alternative formulation, trivially equivalent in the finitary case, it is in head normal form — it has no head redex. It is also a normal form, yet it is unsolvable (that is, there are no terms $N_1, \ldots, N_n$ such that $MN_1 \ldots N_n$ reduces to I). The problem is that application is strict in its first argument, and so an infinitely left-branching chain of applications has no obvious meaning. We can say much the same for an infinite chain of abstractions $\lambda x_1.\lambda x_2.\lambda x_3.\ldots$.

Another reason for reconsidering the definition of infinite terms arises from analogy with term rewriting. In a term such as $F(x, y, z)$, the function symbol F is at syntactic depth 0. If it is curried, that is, represented as $Fxyz$, or explicitly $@(@(@(F, x), y), z)$ (as it would be if we were to translate the term rewrite system into lambda calculus), the symbol F now occurs at syntactic depth 3. We could instead consider it to be at depth zero; more generally, we can define a new measure of depth which deems the left argument of an application to be at the same depth as the application itself, and the body of an abstraction to be at the same depth as the abstraction.

Definition 2.1 Given a term M and a position u of M, the *applicative depth* of the subterm of M at u, if it exists, is defined by:

$$\begin{aligned} D^a(M, \langle\rangle) &= 0 \\ D^a(\lambda x.M, 1 \cdot u) &= D^a(M, u) \\ D^a(MN, 1 \cdot u) &= D^a(M, u) \\ D^a(MN, 2 \cdot u) &= 1 + D^a(N, u) \end{aligned}$$

The associated measure of distance is denoted d^a, and the space of finite and infinite terms Λ^a.

In general, we can specify for each of the three contexts $\lambda x.[]$, $[]M$, and $M[]$ whether the depth of the hole is equal to or one greater than the depth of the whole expression. Syntactic depth sets all three equal to 1. For applicative depth, the three depths are 0, 0, and 1 respectively. This suggests a general definition.

Definition 2.2 Given a term M a position u of M, and a string of three binary digits abc, there is an associated measure of depth D^{abc}:

$$\begin{array}{rcl} D^{abc}(M, \langle\rangle) & = & 0 \\ D^{abc}(\lambda x.M, 1 \cdot u) & = & a + D^{abc}(M, u) \\ D^{abc}(MN, 1 \cdot u) & = & b + D^{abc}(M, u) \\ D^{abc}(MN, 2 \cdot u) & = & c + D^{abc}(N, u) \end{array}$$

The associated measure of distance is denoted d^{abc} and the space of finite and infinite terms Λ^{abc}.

We write Λ^∞, D, or d when we do not need to specify which space of infinite terms, measure of depth, or metric we are referring to. When we refer to certain sets of depth measures, we write e.g. Λ^{**1} to mean all of Λ^{001}, Λ^{011}, Λ^{101}, and Λ^{111}.

We have already seen that $d^s = d^{111}$ and $d^a = d^{001}$. Some of the other measures also have an intuitive significance. d^{101} (*weakly applicative* depth, or d^w) may be associated with the lazy lambda calculus [AO93], in which abstraction is considered lazy — $\lambda x.M$ is meaningful even when M is not. Denote the corresponding set of finite and infinite terms by Λ^w. d^{000} is the discrete metric, the trivial notion in which the depth of every subterm of a term is zero. This gives the discrete metric space of finite terms, no infinite terms, and no reduction sequences converging to infinite terms — the usual finitary lambda calculus.

Many of our results will apply uniformly to all eight infinitary lambda calculi, and we will only specify the depth measure when necessary. In the final section we will find that some of them have unsatisfactory technical properties. The other measures all give rise to different Böhm-like transfinite term models of the lambda calculus.

Lemma 2.3 *Considered as a set, Λ^{abc} is the subset of Λ^{111} consisting exactly of those terms which do not contain an infinite sequence of nodes in which each node is at the same abc-depth as its parent. (Its metric and topology are not the subspace metric and topology.)* □

Both Λ^s and Λ^w contain unsolvable normal forms, such as $\lambda x_1.\lambda x_2.\lambda x_3 \ldots$ In Λ^a every normal form is solvable.

2.3 What is an infinite reduction sequence?

We have spoken informally of convergent reduction sequences but not yet defined them. The obvious definition is that a reduction sequence of length ω converges if the sequence of terms converges with respect to the metric. However, this proves to be an unsatisfactory definition, for the same reasons as in [KKSdV95]. There are two problems. Firstly, a certain property which is important for attaching computational meaning to reduction sequences longer than ω fails.

Definition 2.4 A reduction system admitting transfinite sequences satisfies the *Compression Property* if for every reduction sequence from a term s to a term t, there is a reduction sequence from s to t of length at most ω.

A counterexample to the Compression Property is easily found in Λ^s. Let $A_n = (\lambda x.A_{n+1})(B^n(x))$ and $B = (\lambda x.y)z$. Then $A_0 \to^\omega C$ where $C = (\lambda x.C)(B^\omega)$, and $C \to (\lambda x.C)(yB^\omega)$. A_0 cannot be reduced to $(\lambda x.C)(yB^\omega)$ in ω or fewer steps. (We do not know if the Compression Property holds for the above notion of convergence in Λ^a or Λ^w.)

The second difficulty with this notion of convergence is that taking the limit of a sequence loses certain information about the relationship between subterms of different terms in the sequence. Consider the term I^ω of Λ^a, and the infinite reduction sequence starting from this term which at each stage reduces the outermost redex: $I^\omega \to I^\omega \to I^\omega \to \ldots$ All the terms of this sequence are identical, so the limit is I^ω. However, each of the infinitely many redexes contained in the original term is eventually reduced, yet the limit appears to still have all of them. It is not possible to say that any redex in the limit term arises from any of the redexes in the previous terms in the sequence.

A third difficulty arises when we consider translations of term rewriting systems into the lambda calculus. Even when such a translation preserves finitary reduction, it may not preserve Cauchy convergent reduction. Consider the term rewrite rule $A(x) \to A(B(x))$. This gives a Cauchy convergent term rewrite sequence $A(C) \to A(B(C)) \to A(B(B(C)))\ldots$ If one tries to translate this by defining $A_\lambda = Y(\lambda f.\lambda x.f(Bx))$ (for some λ-term B), where Y is Church's fixed point operator $\lambda f.(\lambda x.f(xx))(\lambda x.f(xx))$, then the resulting sequence will have an accumulation point corresponding to the term $A(B^\omega)$, but will not be Cauchy convergent. The reason is that what is a single reduction step in the term rewrite system becomes a sequence of several steps in the lambda calculus, and while the first and last terms of that sequence may be very similar, the intermediate terms are not, destroying convergence.

The remedy for all these problems is the same as in [KKSdV95]: besides requiring that the sequence of terms converges, we also require that the depths of the redexes which the sequence reduces must tend to infinity.

Definition 2.5 A *pre-reduction* sequence of length α is a function ϕ from an ordinal α to reduction steps of Λ^∞, and a function τ from $\alpha+1$ to terms of Λ^∞, such that if $\phi(\beta)$ is $a \to^r b$ then $a = \tau(\beta)$ and $b = \tau(\beta+1)$. Note that in a pre-reduction sequence, there need be no relation between the term $\phi(\beta)$ and any of its predecessors when β is a limit ordinal.

A pre-reduction sequence is a *Cauchy convergent reduction sequence* if τ is continuous with respect to the usual topology on ordinals and the metric on Λ^∞.

It is a *strongly convergent reduction sequence* if it is Cauchy convergent and if, for every limit ordinal $\lambda \leq \alpha$, $\lim_{\beta\to\lambda} d_\beta = \infty$, where d_β is the depth of the redex reduces by the step $\phi(\beta)$. (The measure of depth is the one appropriate to each version of Λ^∞.)

If α is a limit ordinal, then an *open* pre-reduction sequence is defined as above, except that the domain of τ is α. If τ is continuous, the sequence is *Cauchy continuous*, and if the condition of strong convergence is satisfied at each limit ordinal less than α, it is *strongly continuous.*

When we speak of a reduction sequence, we will mean a strongly continuous reduction sequence unless otherwise stated. Different measures of depth give different notions of strong continuity and convergence.

3. Descendants and residuals

3.1 Descendants

When a reduction $M \to N$ is performed, each subterm of M gives rise to certain subterms of N — its descendants — in an intuitively obvious way. Everything works in almost exactly the same way as for finitary lambda calculus.

Definition 3.1 Let u be a position of t, and let there be a redex $(\lambda x.M)N$ of t at v, reduction of which gives a term t'. The set of *descendants* of u by this reduction, u/v, is defined by cases.

- If $u \not\geq v$ then $u/v = \{u\}$.
- If $u = v$ or $u = v \cdot 1$ then $u/v = \emptyset$.
- If $u = v \cdot 2 \cdot w$ then $u/v = \{v \cdot y \cdot w \mid y \text{ is a free occurrence of } x \text{ in } M\}$. If $u = v \cdot 1 \cdot w$ then $u/v = \{v \cdot w\}$.

The *trace* of u by the reduction at v, $u /\!\!/ v$, is defined in the same way, except for the second case: if $u = v$ or $u = v \cdot 1$ then $u /\!\!/ v = \{v\}$.

For a set of positions U, $U/v = \bigcup\{u/v \mid u \in U\}$ and $U /\!\!/ v = \bigcup\{u /\!\!/ v \mid u \in U\}$.

The notions of descendant and trace can be extended to reductions of arbitrary length, but first we must define the notion of the limit of an infinite sequence of sets.

Definition 3.2 Let $S = \{S_\beta \mid \beta < \alpha\}$ be a sequence of sets, where α is a limit ordinal. Define

$$\liminf S = \bigcup_{\beta \to \alpha} \bigcap_{\beta < \gamma < \alpha} S_\gamma \qquad \limsup S = \bigcap_{\beta \to \alpha} \bigcup_{\beta < \gamma < \alpha} S_\gamma$$

When $\liminf S = \limsup S$, write $\lim S$ or $\lim_{\beta \to \alpha} S_\beta$ for both.

Definition 3.3 Let U be a set of positions of t, and let S be a reduction sequence from t to t'. For a reduction sequence of the form $S \cdot r$ where r is a single step, $U/(S \cdot r) = (U/S) \cdot r$. If the length of S is a limit ordinal α then $U/S = \lim_{\beta \to \alpha} U/S_\beta$.

$U /\!\!/ S$ is defined similarly.

Strong convergence of S ensures that the above limit exists.

Lemma 3.4 *Let U be a set of positions of redexes of t, and let S be a reduction from t to t'. Then there is a redex at every member of U/S.* □

Definition 3.5 The redexes at U/S in the preceding lemma are called the *residuals* of the redexes at U.

Definition 3.6 Let u and v be positions of the initial and final terms respectively of a sequence S. If $v \in u /\!\!/ S$, we also say that u *contributes* to v (via S). If there is a redex at v, then u *contributes* to that redex if u contributes to v or $v \cdot 1$.

We do not define descendants, traces, residuals, and contribution for Cauchy convergent reductions, which is not surprising given the examples of section 2.3.

The next theorem establishes the computational meaning of transfinite sequences, by showing that every finite part of the limit of such a sequence depends on only a finite amount of the work occurring in the sequence.

THEOREM 3.7 *For any strongly convergent sequence $t_0 \to^\alpha t_\alpha$ and any position u of t_α, the set of all positions of all terms in the sequence which contribute to u is finite, and the set of all reduction steps contributing to u is finite.*

PROOF. For each t_β in the sequence, we construct the set U_β of positions of t_β contributing to u, and prove that it is finite. We also show that there are only finitely many different such sets, hence their union is finite.

Suppose $U_{\beta+1}$ is finite, and $t_\beta \to t_{\beta+1}$ reduces a redex at position v. Let $w \in U_{\beta+1}$. If w and v are disjoint, or $w < v$, then w is the only position of t_β contributing to v in $t_{\beta+1}$. If $w = v$, then v, $v \cdot 1$, $v \cdot 1 \cdot 1$, and possibly $v \cdot 2$ (if the redex has the form $(\lambda x.x)N$) are the only such positions. If $w > v$, and the redex at v is $(\lambda x.M)N$, then there is a unique position in either M or N which contributes to w. In each case, the set of positions is finite, hence U_β, which is the union of those sets for all $w \in U_{\beta+1}$, is finite.

Suppose U_β is defined and finite for a limit ordinal β. By strong convergence and the finiteness of U_β, there is a final segment of $t_0 \to^\beta t_\beta$, say from t_γ to t_β, in which every step is at a depth more than 2 greater than the depth of every member of U. It follows that each U_δ for $\gamma \leq \delta < \beta$ is equal to U_β, and is therefore finite.

Finitely many repetitions of the above argument suffice to calculate U_β for all β, demonstrating that there are only finitely many different such sets, and all of them are finite.

Each reduction step contributing to u takes place at a prefix of a position in some U_β. By strong convergence, only finitely many steps can take place at any one position, therefore there are only finitely many such steps. □

3.2 Developments

DEFINITION 3.8 A *development* of a set of redexes R of a term M is a sequence in which every step reduces some residual of some member of R by the previous steps of the sequence. It is *complete* if it is strongly convergent and the final term contains no residual of any member of R.

Not every set of redexes has a complete development. Λ^{**1} contains the term $I^\omega = (\lambda x.x)((\lambda x.x)((\lambda x.x)(\ldots)))$. Every attempt to reduce all the redexes in this term must give a reduction sequence containing infinitely many reduction steps at the root of the term, which is not strongly convergent by any notion of depth. Note that the set consisting of every redex at odd syntactic depth has a complete development, as does the set consisting of every redex at even syntactic depth, but their union does not. In every other version of Λ^∞ except 000 (the finitary calculus) the term $(\lambda x.((\lambda x.((\lambda x.(\ldots))z))z))z$ behaves in a similar manner.

THEOREM 3.9 *Complete developments of the same set of redexes end at the same term.*

Proof. *(Outline.)* In the finitary case one proves this by showing that (1) it is true for a set of pairwise disjoint redexes, (2) it is true for any pair of redexes, and (3) all developments are finite. The result then follows by an application of Newman's Lemma.

In the infinitary case, (1) and (2) are still true, and indeed obvious, but (3) is of course false. The situation is complicated by the fact that a set of redexes can have a strongly convergent complete development without all its developments being strongly convergent.

One proceeds instead by picking out one particular development of the given set of redexes, analogous to the "standard" development defined in finitary rewriting, such that the set has a strongly convergent complete development if and only if its standard development is complete. Properties of the standard development then allow one to use (1) and (2) to construct a "tiling diagram" for the standard development and any other complete development, and to show that the right and bottom edges of the diagram are empty. This shows that the two developments converge to the same limit. □

In the finitary case, the existence of complete developments can be used to prove the Church-Rosser property. In the infinitary case, we have seen that complete developments do not always exist. As a result, the Church-Rosser property does not hold. An example which works for depth measures 1** and *1* is the infinite term which may be described thus: $M = (\lambda x.M')y$, $M' = (\lambda x.M)z$. This can be reduced in infinitely many steps to $M_y = (\lambda x.M_y)y$ or to $M_z = (\lambda x.M_z)z$, which clearly have no common reduct. For depth measures *1* and **1, the term $M = KM'K$, $M' = KMI$, where $K = \lambda x.\lambda y.x$ behaves similarly. We shall see later, however, that the Church-Rosser property does hold up to equality of a certain class of "meaningless" terms.

4. The truncation theorem

Some results about the finitary lambda calculus can be transferred to the infinitary setting by using finite approximations to infinite terms.

Definition 4.1 A $\Lambda_\perp$ term is a term of the version of lambda calculus obtained by adding $\perp$ as a new symbol. $\Lambda_\perp^\infty$ is defined from $\Lambda_\perp$ as Λ^∞ is from Λ.

The terms of $\Lambda_\perp^\infty$ have a natural partial ordering, defined by stipulating that $\perp \leq t$ for all t, and that application and abstraction are monotonic.

A *truncation* of a term t is any term t' such that $t' \leq t$. We may also say that t' is weaker than t, or t is stronger than t'.

Theorem 4.2 *Let $t_0 \rightarrow^\alpha t_\alpha$ be a reduction sequence. Let s_α be a prefix of t_α, and for $\beta < \alpha$, let s_β be the prefix of t_β contributing to s_α. Then for any term r_0 such that $s_0 \leq r_0$ there is a reduction sequence $r_0 \rightarrow^{\leq\alpha} r_\alpha$ such that:*

1. *For all β, s_β is a prefix of r_β.*
2. *If $t_\beta \rightarrow t_{\beta+1}$ is performed at position u and contributes to s_α, then $r_\beta \rightarrow r_{\beta+1}$ by reduction at u.*
3. *If $t_\beta \rightarrow t_{\beta+1}$ is performed at position u and does not contribute to s_α, then $r_\beta = r_{\beta+1}$.* □

As an example of the use of this theorem, we demonstrate that Λ^∞ is conservative over the finitary calculus, for terms having finite normal forms.

COROLLARY 4.3 *If $t \to^\infty s$ and s' is a finite prefix of s, then t is reducible in finitely many steps to a term having s' as a prefix. In particular, if t is reducible to a finite term, it is reducible to that term in finitely many steps.*

PROOF. From Theorems 4.2 and 3.7. □

COROLLARY 4.4 *If a finite term is reducible to a finite normal form, it is reducible to that normal form in the finitary lambda calculus.* □

5. The Compressing Lemma

One of our justifications for the interest of infinite terms and sequences is to see them as limits of finite terms and sequences. From this point of view, the computational meaning may be obscure of a sequence of length longer than ω — which performs an infinite amount of work and then doing some more work. We therefore wish to be assured that every reduction sequence of length greater than ω is equivalent to one of length no more than ω, in the sense of having the same initial and final term. This allows us to freely use sequences longer than ω without losing computational relevance.

THEOREM 5.1 (Compressing Lemma.) *In Λ^∞, for every strongly convergent sequence there is a strongly convergent sequence with the same endpoints whose length is at most ω.*

PROOF. The corresponding theorem of [KKSdV95] shows that the case of a sequence of length $\omega + 1$ implies the whole theorem, and the proof is not dependent on the details of rewriting — it is valid for any abstract transfinite reduction system (as defined in [Ken92]).

Suppose we have a reduction of the form $S_{\omega+1} = s_0 \to^\omega s_\omega \to_d s_{\omega+1}$, where the final step rewrites a redex at depth d. By strong convergence of the first ω steps, the sequence must have the form $s_0 \to^* C[(\lambda x.M)N, M_1, \ldots, M_n] \to^\omega_{d+1} C[(\lambda x.M')N', M'_1, \ldots, M'_n] \to_d C[M'[x := N']]$. where the context $C[\ldots]$ is a prefix of every term of the sequence from some point onwards, and all its holes are at depth d. The reduction of $C[(\lambda x.M)N, M_1, \ldots, M_n]$ to $C[(\lambda x.M')N', M'_1, \ldots, M'_n]$ consists of an interleaving of reductions of M to M', N to N', and each M_i to M'_i of length at most ω. Conversely, any reductions of lengths at most ω starting from M, N, and each M_i can be interleaved to give a reduction of length at most ω starting from $C[(\lambda x.M)N, M_1, \ldots, M_n]$. The theorem will therefore be established if, given reductions of M to M' and N to N' of length at most ω, we can construct a reduction from $(\lambda x.M)N$ to $M'[x := N']$ of length at most ω. This can be done by first reducing $(\lambda x.M)N$ to $M[x := N]$, and then interleaving a reduction of M to M' and reductions of all the copies of N to N' in a strongly convergent way. The details are simple to work out. □

REMARK 5.2 The Compressing Lemma is false for $\beta\eta$-reduction. For a counterexample, let $M = Y(\lambda f.\lambda x.I(fx))$. Then $\lambda x.Mxx \to^\omega \lambda x.I(I(I(...)))x \to_\eta I(I(I(...)))$. However, $\lambda x.Mxx$ is not reducible in ω steps or fewer to $I(I(I(...)))$.

This is not surprising. The η-rule requires testing for the absence of the bound variable in the body of the abstraction; if the abstraction is infinite, this is an infinite task, and such discontinuities are to be expected.

6. Head normal forms and Böhm trees

We earlier gave counterexamples to the Church-Rosser property for all the infinitary lambda calculi, and remarked the property does hold up to equality of a certain set of "meaningless" terms. Here we define and study that class.

In the finitary calculus, one has the concept of the Böhm tree of a term, which from the infinitary perspective can be regarded as its normal form with respect to infinitary reduction together with a rule allowing subterms having no head normal form to be rewritten to the symbol $\bot$. A head normal form is simply a term of the form $\lambda x_1 \ldots \lambda x_n . y t_1 \ldots t_m$.

When one considers the various forms of infinitary calculus, one sees that in Λ^{001}, the head normal forms are precisely the terms which do not have a redex at depth 0. An equivalent characterisation is that they are the terms which cannot be reduced to a term having a redex at depth 0. The equivalence does not hold for some of the other measures of depth. We take the latter as more important, and call it *0-stability.*

Definition 6.1 A *0-redex* of a term is a beta redex or an occurrence of $\bot$ at depth 0. A term of $\Lambda^\infty_\bot$ is *0-stable* if it cannot be beta reduced to a term containing a 0-redex. It is *0-active* it cannot be beta reduced to a 0-stable term.

For $\Lambda^{000}_\bot$, 0-stability is the same as being in normal form and not containing $\bot$. For $\Lambda^{001}_\bot$, 0-stability is the same as being in head normal form and not containing $\bot$ in the place of the head variable.

We now generalise the traditional concept of Böhm reduction.

Definition 6.2 *Böhm reduction* is reduction in $\Lambda^\infty_\bot$ by the β rule and the $\bot$ rule, viz. $M \to \bot$ if M is 0-active and not $\bot$. We write $\to_{\mathcal{B}}$ for Böhm reduction and $\to_\bot$ for reduction by the $\bot$-rule alone.

A *Böhm tree* is a normal form of $\Lambda^\infty_\bot$ with respect to Böhm reduction.

We will show that for some depth measures, every term has a unique Böhm normal form. However, for this it is essential that the 0-active terms are closed under substitution. This is not so for the measures $**0$, as shown by the term $(x\Omega)$, where $\Omega = (\lambda x.xx)(\lambda x.xx)$. This is 0-active, but its instance $(KI\Omega)$ reduces to the 0-stable term I.

Lemma 6.3 *For depth measures $**1$, the set of 0-active terms is closed under substitution.*

Proof. Suppose that t is 0-active. Consider any instance $\theta(t)$ of t and any reduction $\theta(t) \to^\infty_\beta s'$. We shall prove that s' is not 0-stable, which implies that $\theta(t)$ is 0-active.

Begin by imitating the reduction of $\theta(t)$ to s' on t. Let r' be a term in the former sequence and r the corresponding term of the constructed sequence. There will be a set of disjoint positions U_r of r such that r and r' differ only in the subterms at U_r. Initially, this set will be the set of positions of free variables of t which are substituted for in t'.

If the step starting from r' is within a subterm in U_r, then we omit that step from the constructed sequence. If the redex of r' is at a position u such that no prefix of $u \cdot 1$ is in U_r, then the redex is present in r also, and may be reduced. Finally, the redex may be at a position u such that $u \cdot 1$ is in U_r. This means that the redex node is outside the subterms at U_r, but its rator node is in one of those subterms. In this case, both u and $u \cdot 1$ are positions

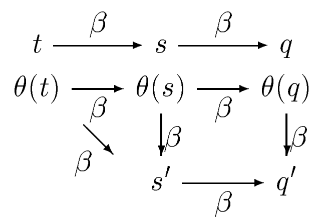

Figure 2.

of r, the former is an application, but the latter may not be an abstraction, and therefore there may be no beta redex at u in r. We omit this reduction step from the constructed sequence, add u to U_r, and omit from U_r every position of which u is a proper prefix, to obtain the set of positions relating the next pair of corresponding terms of the sequences.

The result is to reduce t to a term s which differs from s' only in subterms at positions in a set U_s, such that for each $u \in U_s$, $s|u$ has the form $xt_1 \dots t_n (n \geq 0)$ where x is free in s, and s' is a reduct of a substitution instance of s.

Furthermore, $\theta(t)$ is reducible to $\theta(s)$ (by performing exactly the same reductions that reduce t to s), and $\theta(s) \rightarrow_\beta^\infty s'$ by reductions taking place entirely within the subterms at U_r for the terms r in the sequence from t to s.

By hypothesis, s is not 0-stable. Therefore it is beta reducible to a term q containing a 0-redex. By continuing the construction above, we can obtain the remaining reductions of Fig. 2, where q and q' differ in the same manner that s and s' differed. (To reduce clutter, all arrows in this and similar figures represent reductions of arbitrary length.)

Because the depth measure is $**1$, the subterms of q at U_q, being all of the form $xt_1 \dots t_n$, cannot contain any 0-redexes of q, nor the abstraction node of a 0-redex. Therefore $\theta(q)$ must contain a 0-redex at the same position as q does. The reduction of $\theta(q)$ to q' is performed entirely within subterms in U_q, therefore q' also contains a 0-redex at the same position. Thus s' is not 0-stable. □

Definition 6.4 Two terms t and s are *equivalent* if they differ from each other only at a set of positions U such that for all $u \in U$, $t|u$ and $s|u$ are 0-active.

Lemma 6.5 *For depth measures 001, 101, and 111, if t and s are equivalent, and $t \rightarrow_\beta^\infty t'$, then for some s' equivalent to t', $s \rightarrow_\beta^\infty s'$. The latter reduction can be chosen so as to reduce no redexes inside 0-active subterms.*

Proof. Assuming the hypotheses, we imitate the reduction of t to t' on s. Suppose we have a step $t_0 \rightarrow_\beta t_1$, and a term s_0 equivalent to t_0. If the beta redex is inside one of the 0-active subterm of t_0 at which t_0 differs from s_0, then since 0-active terms are by definition closed under beta reduction, taking $s_1 = s_0$ gives a term equivalent to t_1. If neither the beta redex nor its rator are contained in any of those subterms, the beta redex is present in s_0. Reducing it gives s_1, which (since by Lemma 6.3, 0-active terms are closed under substitution) must be equivalent to t_1. Finally, suppose the redex has the form $(\lambda x.M)N$, where $\lambda x.M$ is one

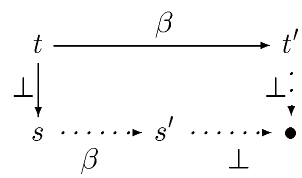

Figure 3.

of the 0-active subterms at which t_0 differs from s_0. Let the subterms of s_0 corresponding to $\lambda x.M$ and N be M' and N'. For $\lambda x.M$ to be 0-active, the depth measure must be 001. M' is 0-active. For depth 001, this implies that $M'N'$ is also 0-active (since 0-active terms for depth 001 are just the terms without head normal form). $(\lambda x.M)N$ is also 0-active. Thus the redex of t_0 is in fact in a 0-active subterm of t_0 corresponding to a 0-active subterm of s_0, reducing this case to one previously considered.

The positions at which reductions are performed in the sequence starting from s are a subsequence of the positions of reductions of the given sequence. Therefore the construction can be continued past limit points of the sequences. □

A counterexample for the depth measure 011 is given by taking $t = (\lambda x.\Omega)y$ and $s = \Omega y$. These are equivalent, since for depth 011, $\lambda x.\Omega$ and Ω are both 0-active. However, $t \to_\beta \Omega$, but s is not beta-reducible to anything equivalent to Ω. The same terms provide counterexamples to all the later theorems which exclude 011.

Corollary 6.6 *For depth measures 001, 101, and 111, given the hypotheses of Lemma 6.5, if $t \to_\perp^\infty s$ then Fig. 3 can be formed.* □

Lemma 6.7 *For depth measures 001, 101, and 111:*

1. *If t and s are equivalent, then t is 0-stable if and only if s is 0-stable.*
2. *If t and s are equivalent, then t is 0-active if and only if s is 0-active.*
3. *Lemma 6.5 also holds when the given reduction of t to t' is a Böhm reduction.*

Proof.

1. Suppose t and s are equivalent, t is 0-stable, and s is not 0-stable. Then s beta reduces to a term r having a 0-redex. By Lemma 6.5, t beta reduces to a term q equivalent to r. If q has a 0-redex, then t is not 0-stable. If q is 0-active, then t is not 0-stable. Thus for t to be 0-stable, q and r must have the same prefix to depth 1, r must have a beta redex $(\lambda x.M)N$ at depth 0, and the corresponding subterm of q must have the form $M'N'$, where both $\lambda x.M$ and M' are 0-active. If the depth measure is $1{*}{*}$, then $\lambda x.M$ cannot be 0-active. If the depth measure is ${*}0{*}$, then M' is at depth 0 in q, so q is not 0-stable, and therefore neither is t. The other depth measures are excluded by hypothesis. Thus in every case, t is not 0-stable.

2. Suppose s is not 0-active. Then s reduces to a 0-stable term r. By Lemma 6.5, t reduces to a term equivalent to r, which by part 1 must be 0-stable. Therefore t is not 0-active.

3. The proof of Lemma 6.5 can be extended to handle Böhm reductions, using part 2 to justify omitting all $\perp$-reductions when constructing the sequence from s.

□

THEOREM 6.8 (The Church-Rosser property, up to equality of 0-active subterms.) *For depth measures 001, 101, and 111: Let $t \rightarrow_\beta^\infty s$ and $t \rightarrow_\beta^\infty s'$. Then there exist equivalent terms r and r', and beta reductions of s to r and s' to r'.*

PROOF.*(Outline.)* The strategy for proving this is the same as that followed in [KKSdV95] in proving the transfinite Church-Rosser property for orthogonal term rewrite systems modulo a class of terms there called "hyper-collapsing".

First, we introduce a new unary function symbol ϵ, and replace the beta rule by a set of rules $(\epsilon^n(\lambda x.M))N \rightarrow \epsilon^{n+2}(M[x := N])$. The depth measure is extended by stipulating that the depth of M in $\epsilon(M)$ is 1. The purpose of this modification to ensure that every residual of a redex is at a depth at least as great as the depth of the redex. From this it follows that every reduction sequence in the new system is strongly convergent. The usual proof of the Church-Rosser property for the finitary calculus, by means of complete developments and tiling diagrams (cf. [Bar84], 11.1), can then be applied to the transfinite case, demonstrating that the modified calculus is transfinitely Church-Rosser.

Finally, we transfer this property to the original calculus. Given two reductions $t \rightarrow_\beta^\infty s$ and $t \rightarrow_\beta^\infty s'$, these correspond in an obvious way to reductions starting from t in the modified calculus, and ending with terms which are versions of s and s' with added occurrences of ϵ. By the Church-Rosser property these can be extended to a common final term. The resulting sequences can then be mapped back to strongly convergent sequences starting from s and s' in the original calculus, provided we omit every step performed in a 0-active subterm. By a version of Lemma 6.5 for the modified calculus, this results in the required reductions of s to r and s' to r'. □

LEMMA 6.9 *For depth measures 001, 101, and 111:*

1. *The set of 0-active terms is closed under Böhm reduction.*

2. *The complement of the set of 0-active terms is closed under beta reduction and $\perp$-reduction.*

PROOF.

1. Immediate from Lemma 6.7(3).

2. Closure under $\perp$-reduction follows from Lemma 6.7(2).

   For closure under beta reduction, suppose $t \rightarrow_\beta^\infty s$ and t is not 0-active. Then $t \rightarrow_\beta^\infty r$ for some 0-stable r. By Theorem 6.8, there are beta reductions $s \rightarrow_\beta^\infty q$ and $r \rightarrow_\beta^\infty q'$ such that q and q' are equivalent. Since r is 0-stable, so are q and q', therefore s is not 0-active.

□

Theorem 6.10 *For every depth measure, every term has a Böhm normal form.*

Proof. A term t is either 0-active or not. If it is, it has the Böhm normal form $\bot$. If it is not, then it can be reduced to a 0-stable term s. Repeating the construction recursively on the subterms of s at depth 1 constructs a reduction of t to a term which is stable to every depth, i.e. a Böhm normal form. □

The above proof does not show uniqueness of Böhm normal forms. For three of the possible depth measures, uniqueness does not hold. For 000, a counterexample is the term $(\lambda y.y\Omega)(KI)$, which has the Böhm reductions $(\lambda y.y\Omega)(KI) \to_{\beta}^{*} I$ and $(\lambda y.y\Omega)(KI) \to_{\bot} \bot (KI)$, which have no common reduct. For the measures $01*$, a counterexample is $(\lambda x.\Omega)y$, where $\Omega = (\lambda x.xx)(\lambda x.xx)$. This term has reductions $(\lambda x.\Omega)y \to_{\beta} \Omega \to_{\bot} \bot$ and $(\lambda x.\Omega)y \to_{\bot} \Omega y \to_{\bot} \bot y$. Both $\bot$ and $\bot y$ are Böhm normal forms. This also refutes the Church-Rosser property of Böhm reduction for these depth measures.

Lemma 6.11 *For depth measures 001, 101, and 111, $\bot$-reduction is transfinitely Church-Rosser.*

Proof. It is immediate from Lemma 6.9(1) that if t is $\bot$-reducible to s, it is so reducible by the reduction of a set of $\bot$-redexes at pairwise disjoint positions. Given two $\bot$-reductions $t \to_{\bot}^{\infty} s$ and $t \to_{\bot}^{\infty} s'$, take the set of outermost members of the union of the two associated sets. Reduction of all of these $\bot$-redexes gives a term r which both s and s' are $\bot$-reducible to. □

Theorem 6.12 *For depth measures 001, 101, and 111, Böhm reduction is transfinitely Church-Rosser.*

Proof. Suppose we have two Böhm reductions starting from a term t. By Lemma 6.7(3) they can be put into the form $t \to_{\beta}^{\infty}\to_{\bot}^{\infty} s_0$ and $t \to_{\beta}^{\infty}\to_{\bot}^{\infty} s_1$.

We then construct Fig. 4. The top left square exists by Theorem 6.8. The top right and bottom left are given by Corollary 6.6. The remaining squares follow from Lemma 6.11. □

So for depth measures 001, 101, and 111, every term has a unique Böhm tree. This gives a transfinite term model of lambda calculus, where the objects are the Böhm normal forms, ordered according to Def. 4.1. The usual Böhm model is the model associated with applicative depth, 001. The larger model described by Berarducci ([Ber]) is the one associated with syntactic depth, 111. In this model the 0-stable terms are the root-stable terms, and the 0-active terms are the terms which Berarducci calls mute. The Böhm model for weakly applicative depth, 101, is related to Ong and Abramsky's models for lazy lambda calculus [AO93].

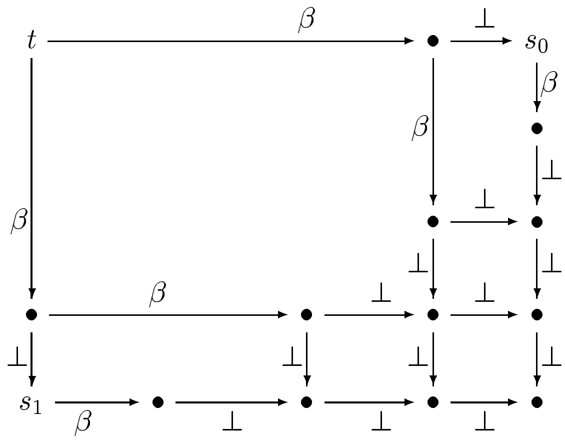

Figure 4. The Church-Rosser property for Böhm reduction

7. Acknowledgements

We acknowledge very useful discussions with Vincent van Oostrom, who drew our attention to the fact that fewer of the depth measures than we realised had good properties.

The figures were composed with Paul Taylor's diagram macros.

References

[AO93] S. Abramsky and L. Ong. Full abstraction in the lazy lambda calculus. *Inf. and Comp.*, 105:159–267, 1993.

[Bar84] H.P. Barendregt. *The Lambda Calculus, its Syntax and Semantics.* North-Holland, 2nd edition, 1984.

[Ber] A. Berarducci. Infinite λ-calculus and non-sensible models. Presented to the conference in honour of Roberto Magari, Siena 1994.

[BI94] A. Berarducci and B. Intrigila. Church-Rosser λ-theories, infinite λ-terms and consistency problems. Dipartimento di Matematica, Università di Pisa, October 1994.

[Ken92] J.R. Kennaway. On transfinite abstract reduction systems. Technical Report CS-9205, CWI, Amsterdam, 1992.

[KKSdV95] J.R. Kennaway, J.W. Klop, M.R. Sleep, and F.J. de Vries. Transfinite reductions in orthogonal term rewriting systems. *Information and Computation*, 1995. To appear. Available by ftp from ftp::/ftp.sys.uea.ac.uk/pub/kennaway/transfinite.{dvi,ps}.Z.