Infinitary lambda calculus

J.R. Kennaway [a], J.W. Klop [b], M.R. Sleep [a] and F.J. de Vries [c,1]

[a] *School of Information Systems, University of East Anglia, Norwich NR4 7TJ, U.K.*

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

[c] *Hitachi Advanced Research Laboratory, Hatoyama, Saitama 350-03, Japan*

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.

1 Introduction

In this paper we extend to the lambda calculus the theory of transfinite term rewriting developed in [8].

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 terms which intuitively denote infinite

[1] Email addresses: `jrk@sys.uea.ac.uk`, `jwk@cwi.nl`, `mrs@sys.uea.ac.uk`, and `ferjan@cwi.nl`. The authors were partially supported by SEMAGRAPH II (ESPRIT working group 6345). Richard Kennaway was also supported by an EPSRC Advanced Fellowship, and by EPSRC Grant no. GR/F 91582. Fer-Jan de Vries was at NTT Communication Science Laboratories, Kyoto 619-02, Japan when this work was initiated.

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 ([3,4]) have independently developed an infinitary lambda calculus and applied it to the study of consistency problems in the finitary lambda calculus.

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. [2] is a standard reference. There are three syntactic classes of terms: variables, which form a set Var, abstractions, which have the form $\lambda x.E$ where $x \in \mathrm{Var}$ and E is a term, and applications, which have the form $E_1 E_2$.

x is called the bound variable and E the body of the abstraction $\lambda x.E$. E_1 is called the rator and E_2 the rand of the application $E_1 E_2$. This is the pure lambda calculus — we do not have any built-in constants nor any type system.

As customary, we identify alpha-equivalent terms with each other, and consider bound variables to be silently renamed when necessary to avoid name clashes. Grouping is indicated by parentheses, with application taken to be left-associative by default.

We consider terms in the abstract to be trees, and identify subterms by their position in the containing term.

A *position* or *occurrence* is a finite string of 1's and 2's. The empty position is denoted $\langle\rangle$. Given a finite term t and a position u, the term $t|u$, when it exists, is a subterm of t defined inductively thus:

$$t|\langle\rangle = t$$

$$(\lambda x.t) \mid 1 \cdot u = t|u$$
$$(st) \mid 1 \cdot u = s|u$$
$$(st) \mid 2 \cdot u = t|u$$

$t|u$ is called the subterm of t at u, and when this is defined, u is called a position of t. The *syntactic depth* of u is its length. Note that in a term $\lambda x.t$ we regard the binding occurrence of x as labelling the root of the term, rather than being a proper subterm. This is merely a technical convenience.

If a position u is a prefix of v we write $u \leq v$. Two positions u and v are *disjoint* if neither is a prefix of the other. A set of positions or redexes is disjoint if every two distinct members are. The set of all positions of t is denoted $pos(t)$.

Let t be a term and u an occurrence of t. If $t|u = \lambda x.s$, u is said to be a *binding occurrence* of x in t. For an occurrence v of t such that $v > u$, x is *bound by* u at v if there is no occurrence w such that $u < w \leq v$ and w is also a binding occurrence of x. If v is an occurrence of x, it is called a *bound occurrence* of x. An occurrence of x in t which is not bound is *free*.

2.2 What is an infinite term?

Drawing lambda terms as syntax trees gives an immediate and intuitive notion of infinite terms: they are just infinite trees, and can be formally defined as the points added by taking the metric completion of the space of finite terms, subject to a suitable metric, as was done in [8] for infinitary term rewriting. However, we will want to identify with each other terms which differ only in the names of their bound variables. This is a little more complicated in the infinitary setting than in the finite setting. There are several approaches we can take. Firstly, we may define infinite terms by a metric space completion, and then take the quotient of this space by alpha-equivalence. Alternatively, we can start with finite terms modulo alpha-equivalence, and then construct the infinite terms by the metric space completion. Finally, we could avoid the notion of alpha-equivalence altogether by using the de Bruijn representation of lambda terms.

Although de Bruijn notation avoids alpha-equivalence, it significantly complicates the definition of substitution (as demonstrated in [5]). Instead, we shall adopt both of the first two alternatives above.

To begin with, then, we ignore concerns about renaming of variables. We define a metric on the set of finite terms thus: $d(s,t) = 0$ if $s = t$; otherwise, $d(s,t) = 2^{-n}$, where n is the length of the shortest common occurrence u of s and t such that s and t differ at u. The latter concept means this: s and t differ at u if none of the following hold: $s|u$ and $t|u$ are the same variable, $s|u$

and $t|u$ are both applications, of $s|u$ and $t|u$ are both abstractions having the same bound variable. The completion of this metric space adds the infinite terms. These look like infinite trees, in which each node is at a finite distance from the root of the tree. A tree can contain infinitely long branches, but there is no node at the end of an infinite branch.

Now we define alpha-equivalence of finite or infinite terms. Firstly, we require a limited notion of substitution: Given a term t and two variables x and x', $t[x \to x']$ is the term resulting from replacing every free occurrence of x in t by x'.

Now we define s and t to be *alpha-equivalent* if they have no conflict, where a conflict between s and t is a common occurrence u at which they "look different" in the sense that one of the following holds:

(i) $u = \langle\rangle$ and s and t are not identical variables, not both applications, and not both abstractions.
(ii) $u = n \cdot v$, $s = s_1 s_2$, $t = t_1 t_2$, and v is a conflict of s_n and t_n ($n = 1$ or 2).
(iii) $u = 1 \cdot v$, $s = \lambda x.s'$, $t = \lambda x'.t'$, and v is a conflict of $s'[x \to x'']$ and $t'[x' \to x'']$, where x'' is a variable not occurring in s' or t'.

It is routine to show that this is an equivalence relation. To justify taking the quotient of the metric space by the equivalence, we must show (1) that given two distinct equivalence classes, there is a positive lower bound on the distance between a member of one and a member of the other, and (2) that the triangle inequality is satisfied. Both of these are immediate from the following lemma, which also gives a direct definition of the metric on the space of equivalence classes.

Lemma 1 *Let S and T be distinct alpha-equivalence classes. Choose any s and s' in S and t and t' in T. Then $d(s,t) = d(s',t')$.*

Proof. By symmetry, it is enough to prove $d(s,t) = d(s,t')$. Let u be an occurrence of minimal length where s and t conflict. Since t and t' are alpha-equivalent, they have the same set of occurrences, and at each common occurrence, their syntactic class is the same. We proceed by induction on u.

(i) If u is empty, it is immediate from the definition of conflict that s and t' conflict at u.
(ii) Suppose that $u = n \cdot v$, and s, t are applications. Then t' is also an application and $s|n$ and $t|n$ conflict at v. By induction, $s|n$ and $t'|n$ conflict at v.
(iii) Suppose that $u = 1 \cdot v$, $s = \lambda x.s'$, and $t = \lambda y.r$. Then t' has the form $\lambda y'.r'$, with $r'[y' \to z]$ alpha-equivalent to $r[y \to z]$ (for z not occurring in any of these terms). Then $s'[x \to z]$ and $r[y \to z]$ conflict at v. for

a variable z chosen so as not to occur in s, t, or t'. $r'[y' \to z]$ is alpha-equivalent to $r[y \to z]$. By induction, $s'[x \to z]$ and $r'[y' \to z]$ conflict at v.

□

Here are some examples of this metric on the space of equivalence classes (for which we use the same notation d).

$$\begin{aligned}
d(\lambda x.x, \lambda y.y) &= 0 \\
d(\lambda x.x, x) &= 1 \\
d(\lambda x.x, \lambda y.yy) &= d(\lambda x.x, \lambda x.xx) = \frac{1}{2} \\
d(\lambda x.\lambda y.x, \lambda y.\lambda x.x) &= d(\lambda x.\lambda y.x, \lambda x.\lambda y.y) = \frac{1}{4}
\end{aligned}$$

We could instead construct the desired space by starting from the conventional notion of alpha-equivalence on finite terms, and defining a metric on the set of finitary alpha-equivalence classes in terms of the length of the shortest conflict, as above. The metric completion of this space gives the same space of alpha-equivalence classes of finite and infinite terms.

Henceforth, we will not mention alpha-equivalence directly. When we write particular terms, we understand them as representatives of their classes, and take alpha-conversions to be performed implicitly.

We are now ready to define substitution and beta-reduction. In finitary lambda calculus, substitution may be defined by an induction over the structure of terms. Such an induction is not well-founded for infinite terms, but we will see that with a little care, definitions can still be given in the inductive style.

Definition 2 *For terms t and t' and a variable x, the term $t[x := t']$ is defined thus:*

$$\begin{aligned}
x[x := t'] &= t' \\
y[x := t'] &= y \\
(t_1 t_2)[x := t'] &= (t_1[x := t'])(t_2[x := t']) \\
(\lambda x.t)[x := t'] &= \lambda x.t \\
(\lambda y.t)[x := t'] &= \lambda y.(t[x := t''[y := z]])
\end{aligned}$$

where in the last case, z is a variable not occurring in $\lambda y.t$ or t'.

This definition can be justified by reading it as an algorithm for constructing the term $t[x := t']$ from the root downwards, rather than an inductive construction of $t[x := t']$ from its proper subterms. Although it may take an infinite amount of work to construct the whole term, each finite part of the result depends on only a finite part of t and t'. Technically, the construction could be expressed as a definition of the set of occurrences of $t[x := t']$, together with a definition of the syntactic class of the subterm of $t[x := t']$ at each such occurrence. This would cast the induction as being over the length of an occurrence. This is well-founded, since even for infinite terms, every occurrence is by definition finite.

The following generalises a standard property of substitution to infinite terms (cf. [2], 2.1.16).

Lemma 3 (Substitution Lemma.) *If x and y are distinct, and x is not free in r, then $t[x := s][y := r] = t[y := r][x := s[y := r]]$.*

Proof. A term is completely determined by the set of its positions, and the symbols which it has at those positions. By definition, $t[x := s]$ is the term having the following positions and symbols (writing $S(t, u)$ for the symbol at position u in t, and assuming alpha-conversion to ensure that x and y do not occur bound in any of the terms under consideration):

(i) u where $u \in pos(t)$ and $t|u \neq x$. $S(t[x := s], u) = S(t, u)$.
(ii) $u.v$ where $u \in pos(t)$, $t|u = x$, $v \in pos(s)$. $S(t[x := s], u.v) = S(s, v)$.

From this we can calculate the positions and symbols of $t[x := s][y := r]$ and $t[y := r][x := s[y := r]]$. An analysis of cases establishes that the two terms are identical. □

Definition 4 *A* beta redex *of a term t is a subterm of the form $(\lambda x.t')t''$, at some position u. The operation of* beta-reduction *of that redex replaces it by $t'[x := t'']$. If the resulting term is s, we write $t \to_u s$. The u may be omitted when it is unimportant. A finite sequence of reductions leading from t to s is written $t \to^* s$.*

The above is the natural equivalent for lambda calculus of the definition of infinite terms which we used in our study of transfinite term rewriting. However for lambda calculus there are further considerations. Consider the term ${}^{\omega}x = (((\ldots x)x)x)x$. 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

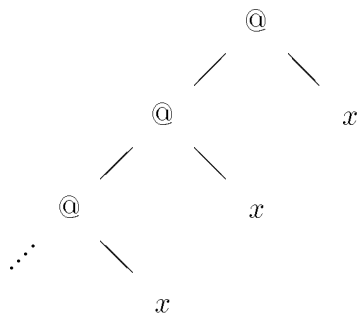

Fig. 1. A paradoxical infinite term

is, there are no terms $t_1, \ldots, t_n$ such that $tt_1 \ldots t_n$ reduces to I). The problem is that application is in some sense 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 5 *Given a term t and a position u of t, the* applicative depth *of the subterm of t at u, if it exists, is defined by:*

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

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 $\lambda x.[\,]$, $[\,]t$, and $t[\,]$ whether the depth of the hole is equal to or one greater than the depth of the whole term. 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 6 *Given a term t a position u of t, and a string of three binary digits abc, there is an associated measure of depth D^{abc}:*

$$\begin{aligned} D^{abc}(t, \langle\rangle) &= 0 \\ D^{abc}(\lambda x.t, 1 \cdot u) &= a + D^{abc}(t, u) \\ D^{abc}(st, 1 \cdot u) &= b + D^{abc}(s, u) \\ D^{abc}(st, 2 \cdot u) &= c + D^{abc}(t, u) \end{aligned}$$

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 [1], in which abstraction is considered lazy — $\lambda x.t$ is meaningful even when t 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 (as we will see when we define infinite reductions) no infinite reduction sequences converging to limits — the usual finitary lambda calculus.

We will only specify the depth measure when necessary. Some of our results will apply uniformly to all eight infinitary lambda calculi, others are restricted. In the final section we will find that three of them give rise to different Böhm-like transfinite term models of the lambda calculus, one of which is the usual Böhm model.

Lemma 7 *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. (Note that its metric and topology are not the subspace metric and topology.)*

Hence by König's Lemma a term of Λ^{abc} contains only finitely many nodes at any given depth, and a fortiori *finitely many redexes.* $\square$

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 [8]. There are three problems. Firstly, a certain property which is important for attaching computational meaning to reduction sequences longer than ω fails.

Definition 8 *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.*

For the above notion of convergent sequence, 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 $C \to (\lambda x.C)(y(B^\omega))$. A_0 cannot be reduced to $(\lambda x.C)(y(B^\omega))$ 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 in 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.

The 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 $A(C) \to A(B(C)) \to A(B(B(C)))\ldots A(B^\omega)$. 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 [8]: 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 9 *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 $a \to b$ 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 $\mu \leq \alpha$, $\lim_{\beta\to\mu} d_\beta = \infty$, where d_β is the depth of the redex reduced 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. We finish this section with explicit notation to indicate the length of reduction sequences.

Definition 10 *(i) $s \to^* t$ denotes a reduction from s to t of finite length.*
(ii) $s \to^\alpha t$ denotes a reduction from s to t of ordinal length α.
(iii) $s \to^{\leq\alpha} t$ denotes a reduction from s to t of ordinal length at most α.
(iv) $s \to^\infty t$ denotes a reduction from s to t of some unspecified finite or infinite length.

3 Descendants and residuals

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

Definition 11 *Let u be a position of t, and let there be a redex $(\lambda x.s)r$ of t at v, reduction of which gives a term t'. The set of* descendants *of u by this reduction, u/v, and its* trace *$u/\!\!/v$ are defined by cases.*

- *If $u \not\geq v$ then $u/v = u/\!\!/v = \{u\}$.*
- *If $u = v$ or $u = v \cdot 1$ then $u/v = \emptyset$ and $u/\!\!/v = \{v\}$.*
- *If $u = v \cdot 2 \cdot w$ then $u/v = u/\!\!/v = \{v \cdot p \cdot w \mid p$ is a free occurrence of x in $s\}$.*
  If $u = v \cdot 1 \cdot 1 \cdot w$ then $u/v = u/\!\!/v = \{v \cdot w\}$.

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

Consider the term $(\lambda x.xx)w((\lambda z.y)w)$. This has redexes at positions 1 and 2. Here are some examples of descendants and traces.

$$\epsilon/1 = \epsilon/\!/1 = \{\epsilon\}$$

$$\begin{gathered} 1/1 = \{\,\} \qquad 1/\!/1 = \{1\} \\ 1{\cdot}1/1 = 1{\cdot}1/\!/1 = \{\,\} \\ 1{\cdot}2/1 = 1{\cdot}1/\!/1 = \{1{\cdot}1, 1{\cdot}2\} \\ 1{\cdot}1{\cdot}1/1 = 1{\cdot}1{\cdot}1/\!/2 = \{1\} \\ 2{\cdot}2/2 = 2{\cdot}2/\!/2 = \{\,\} \end{gathered}$$

The only difference between descendant and trace is that the root of a redex has no descendant with respect to itself, but is its own trace with respect to itself.

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

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

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

Definition 13 *Let U be a set of positions of t, and let S be a reduction sequence from t to t'. If S is empty then $U/S = U/\!/S = U$.*

For a reduction sequence of the form $S \cdot r$ where r is a single step, $U/(S \cdot r) = (U/S)/r$. If the length of S is a limit ordinal α then $U/S = \lim_{\beta<\alpha} U/S_\beta$. (The existence of this limit is proved below.) $U/\!/S$ is defined similarly.

Lemma 14 *Strong convergence of S ensures that the limit in Definition 13 exists.*

Proof. Suppose that U/S' has been defined for every proper initial segment S' of a strongly convergent sequence S of limit ordinal length μ. Write U_α for U/S', where S' is the initial segment of S of length α. Write U_α/d for the subset of U_α of positions of depth at most d. Choose any depth d. By strong convergence of S, there is an ordinal $\alpha < \mu$ such that every reduction of S after step α occurs at a depth greater than d. Such a reduction cannot add or remove positions of depth at most d. Therefore when $\alpha \leq \beta < \mu$,

$U_\beta/d = U_\alpha/d$. Hence $\liminf\{U_\beta/d \mid \beta < \mu\} = U_\alpha/d = \limsup\{U_\beta/d \mid \beta < \mu\}$. Since this holds for all d, $\liminf\{U_\beta \mid \beta < \mu\} = \limsup\{U_\beta \mid \beta < \mu\}$. $\square$

Lemma 15 *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.*

Proof. This is trivial if S is empty.

It holds if S is one step long, since if S reduces at u, then u/S is by definition empty, and if S does not reduce at u, where u is the position of a redex, then every subterm of t' at u/S is an application node whose left descendant is a lambda node, i.e. a redex.

Hence if the lemma holds for reductions of length α, it holds for reductions of length $\alpha + 1$.

If the length of S is a limit ordinal, then every member of U/S is a member of U/S' for long enough proper initial segments S' of S. If the lemma is true for every such S', then every member of U/S is the position of a redex in every term of S from some point onwards, hence is the position of a redex in the limit of S. $\square$

Definition 16 *The redexes at U/S in the preceding lemma are called the* residuals *of the redexes at U.*

Definition 17 *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 18 *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 w 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 $(\lambda x.x)t$) are the only such positions of t_β. If $w > v$, and the redex at v is $(\lambda x.s)t$, then there is a unique position in either s or t 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 steps of U contributing to u. $\square$

Lemma 19 *Let $t_0 \to t'_0$ by reduction at a position u, and $t_0 \to^\alpha t_\alpha$ by a sequence of reductions, none of which is at any prefix of u. Then t_α reduces by a reduction at u to a term t'_α, and $t' \to^\infty s'$.*

If no step of $t_0 \to^\alpha t_\alpha \to_u t'_\alpha$ except the last contributes to the position u in t'_α, then neither does any step except the first in $t_0 \to_u t'_0 \to^\infty t'_\alpha$.

Proof. It is immediate that each t_β has a redex at u. Let t'_β be the result of reducing it. To prove the theorem it is sufficient to construct a reduction of each t_β to $t_{\beta+1}$, and to show that their concatenation is strongly convergent.

Let r_β be the redex which is reduced to obtain t_{beta+1} from t_β. This redex has a set of residuals R_β by the redex at u. These are at pairwise disjoint positions, therefore they can be reduced in increasing order of depth, to obtain a strongly convergent reduction whose limit is clearly $t'_{\beta+1}$. The depth of all of these residuals is at least the depth of r minus 2. (Equality will hold, for example, in depth measure 001 if r is inside the body of the rator of the redex at u.) Since the reduction of t_0 to t_α is strongly convergent, this implies that the concatenation of the reductions from each t_β to $t_{\beta+1}$ is also strongly convergent.

Suppose that some step of $t'_0 \to^\infty t'_\alpha$ contributed to u. By Theorem 18, there are only finitely many such steps, therefore there is a last one. Since no step of $t_0 \to^\alpha t_\alpha$ contributes to u in t'_α, This can only be possible if the step of $t_0 \to^\alpha t_\alpha$ corresponding to the last contributing step of $t'_0 \to^\alpha t'_\alpha$ is at position $u \cdot 1 \cdot 1$. But then the reduction from that step onwards has the form $C[(\lambda x.((\lambda y.a)b))c] \to_{u \cdot 1 \cdot 1} C[(\lambda x.b[y := a])c] \to^\infty C'[(\lambda x.b'[y := a'])c'] \to_u b'[y := a'][x := c']$. But by definition, the first step contributes to u in the final term. $\square$

Corollary 20 *Given the hypothesis of Theorem 18, there is a reduction of t_0 to t_α in which all the steps contributing to the position u occur before all the other steps.*

Proof. Almost immediate from Lemma 19. Let the first step of the given sequence which contributes to u be a redex r at position u', and let the preceding segment of the sequence be S. r must be the unique residual by S of a redex in t_0 at u', and no step of S can be at any prefix of u. This establishes the conditions of Lemma 19. Repeating the argument for all of the finitely many steps contributing to u yields the corollary. $\square$

4 Developments

Definition 21 *A* development *of a set of redexes R of a term 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. However, we can show that when complete developments exist, the limit is determined by the set of redexes being developed.

Theorem 22 *(Infinite Developments Theorem.) Complete developments of the same set of redexes end at the same term.*

The proof of this theorem fills the remainder of this section.

In the finitary case it can be proved by showing that (1) it is true for a set of redexes at pairwise disjoint positions, (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 of its developments being strongly convergent.

We proceed instead by picking out a certain well-behaved class of developments, the outside-in developments, and proving that we can restrict attention to this class.

Properties of outside-in developments then allow one to use (1) and (2) to construct a "tiling diagram" for an outside-in development and any complete development, and to show that the right and bottom edges of the diagram exist and are empty. This shows that the two developments converge to the same limit.

Definition 23 *An* outside-in *development of a set of redexes U is a development which at each step reduces a redex which is outermost among the residuals of U and of minimum depth.*

Note that in Λ^{111}, minimum depth implies outermost, making the outermost condition unnecessary, but this is not so for the other depth measures.

Lemma 24 *An outside-in development has length at most ω.*

Proof. Consider a strongly convergent outside-in development of length ω. Then for every depth d, there is some point in the sequence after which no reductions are performed at depth d, and hence no residuals of U exist at depth d. Therefore the limit contains no residuals of U, and the development is complete. $\square$

We shall now obtain a syntactic description of outside-in developments, which will be achieved in Lemma 33.

Definition 25 *Given a term t and a set U of redexes of t, a* level path *of U is a sequence of positions $u_0, u_1, \ldots$ of maximal length, constructed as follows. u_0 is the position of an outermost redex of U of minimum depth. If u_n has been defined, then the path is continued by one of the following cases.*

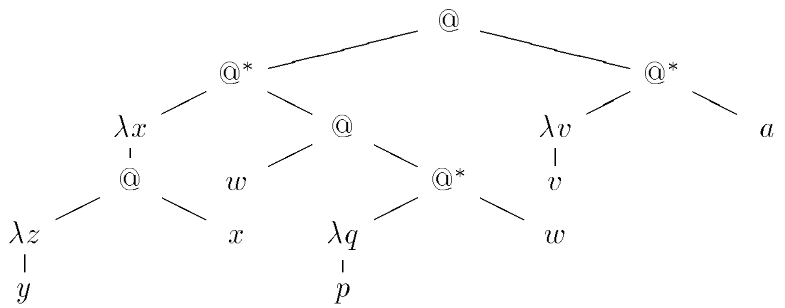

Fig. 2. Example for Definition 25

(i) *If u_n is the position of a redex of U, then $u_{n+1} = u_n \cdot 1$ and $u_{n+2} = u_n \cdot 1 \cdot 1$. That is, u_{n+1} is the position of the abstraction node of the redex and u_{n+2} is the position of the body of the abstraction.*

(ii) *Otherwise, if u_n is the position of a variable x, and x is bound by a redex $(\lambda x.t')t''$ in U at position u_i for some $i < n$, then $u_{n+1} = u_i \cdot 2$. That is, u_{n+1} is the current position of the subterm which will be substituted for x when that redex is reduced.*

(iii) *Otherwise, if u_n has an immediate descendant whose depth is the same as the depth of u_n, then u_{n+1} can be chosen to be any such descendant. (A choice only arises if u_n is an application node and the depth measure is $*00$.)*

(iv) *Otherwise, the path stops at u_n.*

In Figure 2, let U contain the three starred redexes. Assume the depth measure 100. There are four level paths of U:

$$\langle 1,\ 1{\cdot}1,\ 1{\cdot}1{\cdot}1,\ 1{\cdot}1{\cdot}1{\cdot}1 \rangle$$

$$\langle 1,\ 1{\cdot}1,\ 1{\cdot}1{\cdot}1,\ 1{\cdot}1{\cdot}1{\cdot}2,\ 1{\cdot}2,\ 1{\cdot}2{\cdot}1 \rangle$$
$$\langle 1,\ 1{\cdot}1,\ 1{\cdot}1{\cdot}1,\ 1{\cdot}1{\cdot}1{\cdot}2,\ 1{\cdot}2,\ 1{\cdot}2{\cdot}2,\ 1{\cdot}2{\cdot}2{\cdot}1,\ 1{\cdot}2{\cdot}2{\cdot}1{\cdot}1 \rangle$$
$$\langle 2,\ 2{\cdot}1,\ 2{\cdot}1{\cdot}1,\ 2{\cdot}2 \rangle$$

Lemma 26 *(i) No node can occur twice in the same level path.*
(ii) *An infinite level path of U contains (the root nodes of) infinitely many redexes of U.*

Proof. The first follows from the fact that each node is strictly below or to the right of every preceding node in the path.

For the second, since no term can contain infinitely many nodes at the same depth, in an infinite level path there must be infinitely many places where the depth increases. This only happens in case (ii) of the definition. But each time

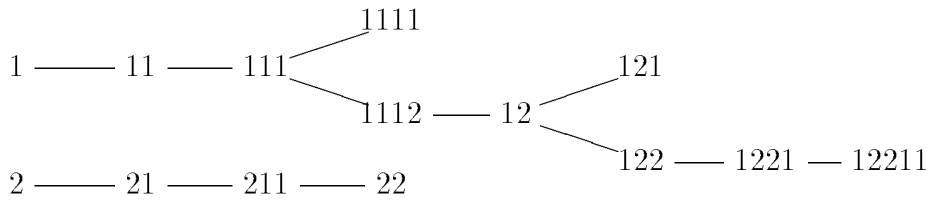

Fig. 3. A level forest

that case is applied, a different variable, and hence a different redex on the path, must be involved. □

Lemma 27 *Let U be a set of redexes in a term t, and let U' be a subset of U consisting of all the redexes of U lying on a level path of U. Then in an outside-in development of U', every reduction is performed at the same depth.*

Proof. An outside-in development of U' must begin by reducing an outermost redex r of the level path. It is immediate from the construction of a level path that if it is not the only member of U', then U'/r is a set of redexes also lying on a level path, and that its outermost member is at the same depth as r. □

The next lemma is immediate, and is the motivation for the notion of level path.

Lemma 28 *Let r and r' be consecutive redexes in U on a level path of U. Then r' has a residual by r whose depth is the same as the depth of r.*

Let a redex r contain a redex r'. Then r' has a residual by r whose depth is equal to the depth of r if and only if r and r' are consecutive redexes on some level path of $\{r, r'\}$. □

Definition 29 *Let U be a set of redexes in a term t. The set of level paths of U can be arranged into a forest, the* level forest *of U, denoted $Lf(U)$. This forest has one node for every different non-empty initial segment of every level path of U, arranged according to the prefix ordering.*

The depth *of the level forest is the minimum depth of any position occurring in it (which is the same as the minimum depth of any redex of U).*

The level forest of the earlier example is drawn horizontally in Figure 3.

Lemma 30 *The level forest of U is finitely branching, that is, it has finitely many component trees, each of which is finitely branching.*

Proof. U can have only finitely many redexes of minimum depth, therefore the level forest has finitely many component trees. Each node of the forest can have at most two descendants (and then only in depth measures $*00$, as noted in the definition of level path). $\square$

Lemma 31 *Let r be a redex of U whose position does not occur in $Lf(U)$. Then $Lf(U/r) = Lf(U)$.*

Proof. It is immediate from the definition of $Lf(U)$ that r cannot contain any position occurring in $Lf(U)$. From this the lemma follows. $\square$

Lemma 32 *Let $Lf(U)$ be finite. Let r be a redex of U whose position occurs in $Lf(U)$.*

If at least one redex of U in $Lf(U)$ has a residual by r, then $Lf(U/r)$ and $Lf(U)$ have the same depth. Furthermore, if $Lf(U)$ is finite then $Lf(U/r)$ is smaller and if $Lf(U)$ is infinite then $Lf(U/r)$ is infinite.

If no redex of U in $Lf(U)$ has a residual by r, and U/r is nonempty, then the depth of $Lf(U/r)$ is greater than the depth of $Lf(U)$.

Proof. The depth claims follow from Lemma 28.

We prove the rest by examining the effect of reducing r on each of the level paths of U.

Consider any level path of U not containing r. Clearly, it will be unaffected by the reduction of r, and must be a level path of U/r.

Now consider any level path of U which contains $r = (\lambda x.t')t''$. Let the position of r be u. If r is the only redex on that level path, then reducing r eliminates the path from the level forest, and hence deletes at least one leaf node. This cannot happen if the path is infinite, since an infinite path must contain infinitely many redexes in U.

If r is not the only redex of U on the path, then reducing r removes at least two nodes from the path (the root and the abstraction node of r). If the path contains a free occurrence of x, and therefore continues from there with the root of t'', then the path is shortened by a third node, since after reducing r the parent of that occurrence of x will be immediately followed by (a descendant of) the root of t''. No other nodes are removed from the path. Reducing r cannot add any nodes to the end of the path.

Therefore reducing r deletes or shortens by a finite amount every path it lies on, does not delete any infinite path, and does not affect any other path. The lemma follows. $\Box$

From the above lemmas, we obtain the following picture of outside-in developments.

Lemma 33 *An outside-in development of a set of redexes U consists of a sequence of developments $S_0 \cdot S_1 \cdot S_2 \cdots$ in which*

(i) S_i is an outside-in development of the redexes in the level forest of the initial term of S_i, complete if finite.
(ii) Each S_i consists of reductions all at the same depth d_i.
(iii) The sequence of d_i is strictly increasing.
(iv) Either the sequence consists of infinitely many finite segments and is strongly convergent, or it consists of finitely many finite segments, or it consists of finitely many finite segments followed by a non-strongly convergent infinite segment.

Proof. From Lemma 32, an outside-in development of U begins with a development S_0 of the redexes U_0 contained in the level forest of U. By the same lemma, these reductions all happen at the same depth, the depth d_0 of the level forest. That development is infinite if and only if the level forest is infinite; in that case all the claims hold.

If S_0 is finite, i.e. is a complete development of U_0, let the last step of S_0 be a reduction of a redex r in a term t. Each member of U/S_0 must be a residual by r of a redex r' of t which is not in the level forest at that point. From the definition of the level forest it follows that every residual of r' by r has a depth greater than r. Hence the depth of the level forest of U/S_0 is $d_1 > d_0$. By iterating the argument the lemma follows. $\Box$

Lemma 34 *If $Lf(U)$ is infinite then U has no complete development.*

Proof. Let the minimum depth of any redex of U be d. By Lemmas 31 and 32, $Lf(U/r)$ is infinite for any $r \in U$ and its outermost redexes are at depth d. Therefore after any number of reductions, there is still a residual of U at depth d, and therefore for the development to be complete, there must be a later reduction at depth d. Hence the development cannot be strongly convergent. $\Box$

Lemma 35 *Let S be a complete development of U, and let r be an outermost member of U of minimal depth. Then U/r has a complete development with the same final term as S.*

Proof. Let r be at position u, of depth d. Since it is of minimal depth among members of U, and S is a development of U, the set of residuals of u by any initial segment of S must be either $\{u\}$ or empty, and will be empty for a long enough segment. Furthermore, we can construct a complete development of U/r by performing for each step r' of S, a complete development of $r'/(r/T)$, where T is the initial segment of S before r'. Since r/T is at most one step, $r'/(r/T)$ is a set of redexes at pairwise disjoint positions, hence trivially has a complete development. Furthermore, since r/T is empty for long enough T, this complete development of U must contain a final segment of S, hence has the same limit. □

Proof of Theorem 22. Let U be a set of redexes of a term t having a complete development S. Choose an outside-in development S' of U of maximal length. Decompose S' into $S_0 \cdot S_1 \cdots$ according to Lemma 33. Apply Lemma 35 to S and each step of S', to obtain for each i for which S_i is finite a complete development of $U/(S_0 \cdots S_i)$ ending with the same term as S. Every reduction in any development of that set is at a depth of at least d_{i+1}. Therefore the distance between the final terms of S and S' is not more than 2^{-i-1}.

If S' is finite, then for some i, S_i is empty, and so S and S' have the same limit.

If S' is infinite and strongly convergent, then the sequence of d_i is infinite and tends to infinity, and so the distance between the final terms of S and S' is zero.

Finally, if S' is infinite and not strongly convergent, then for some i, S_i is infinite and $U/(S_0 \cdots S_{i-1})$ has an infinite level forest. By Lemma 34 it has no complete development. But $S/(S_0 \cdots S_{i-1})$ is a complete development of that set, contradiction.

Therefore every complete development has the same limit as every maximal outside-in development, and hence all complete developments have the same limit. □

This proof yields as a corollary a new proof of the finite developments theorem of finitary lambda calculus. Consider the characterisation of outside-in developments in Lemma 33. The segment S_0 must be finite, since the level forest of U is finite. Let $r = (\lambda x.t)t'$ be the redex reduced in the last step of S_0.

Suppose r' is another residual of U in the term just before reduction of r. If r' were outside r, or in t, then r would not be the last step of S_0 – a residual of r' or of some other member of U in the level forest of U would still be present after reduction of r. The same holds if r' were in t', and x occurred free in t. Therefore such a residual r' can exist only in t', and then only if x is not free in t. This implies that r' has no residuals by r. Therefore after r is reduced, no residuals of U remain, i.e. S_0 is a complete outside-in development of U. Now suppose there was an infinite development S of U. We have proved that there is a complete outside-in development S' of U. The construction of Lemma 35, applied to the case where S is infinite, shows that for every finite initial segment S'' of S' there is an infinite development of U/S''. But if $S'' = S'$, then U/S'' is empty, contradiction. Therefore U has no infinite development.

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 for depth measures $1{*}{*}$ and ${*}1{*}$ is the infinite term which may be described thus: $t = (\lambda x.t')K$, $t' = (\lambda x.t)I$. This can be reduced in infinitely many steps to $t_K = (\lambda x.t_K)K$ or to $t_I = (\lambda x.t_I)I$, which clearly have no common reduct. For depth measures ${*}1{*}$ and ${*}{*}1$, the term $t = Kt'K$, $t' = KtI$, where $K = \lambda x.\lambda y.x$ behaves similarly.

Even the special case known as the Strip Lemma fails. The Strip Lemma says that if $t \to s$ and $t \to^\infty s'$, then s and s' have a common reduct. This is the Church-Rosser property restricted to the case where one of the two reduction sequences consists of a single step. Lemma 19 is a special case of this, but there are counterexamples to the general case. For any infinitary calculus except 001, consider a term $t(\lambda x.(\lambda y.x)w)$, where t is a (finite) term with the property that $tx \to^* x(tx)$ (easily constructed by means of a fixed point operator). This reduces in one step to $t(\lambda x.x)$, or in infinitely many steps to the term s such that $s = (\lambda y.s)w$ (where the '=' is identity, not convertibility, as it is throughout this paper). These cannot have any common reduct, since w is free in every reduct of s, but every reduct of $t(\lambda x.x)$ is closed. There is a similar counterexample for depth 001. Take t as before and consider the term $t(\lambda y.((\lambda x.x)(\lambda w.y)))$. This reduces in one step to $t(\lambda y.\lambda w.y)$, and in infinitely many steps to a term $s = (\lambda x.x)(\lambda y.s)$. Every reduct of s contains the subterm $\lambda x.x$, but no reduct of $t(\lambda y.\lambda w.y)$ does.

Despite these counterexamples, we shall see in Section 7.2 that the Church-Rosser property does hold up to equality of a certain class of "meaningless" terms.

5 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 36 *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 37 *In any Λ^{abc}, let $t_0 \to^\alpha t_\alpha$ 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 $r_0 \to^{\leq\alpha} r_\alpha$ such that:*

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

Proof. Claims (ii) and (iii) immediately suggest an inductive construction of the required reduction sequence. The initial term r_0 is given by hypothesis. Suppose that $r_0 \to^{\leq\beta} r_\beta$ has been constructed. Define $r_{\beta+1}$ by whichever of (ii) or (iii) applies. It is immediate from the definition of contribution that in either case, $s_{\beta+1}$ will be a prefix of $r_{\beta+1}$. Note that the two nodes of the redex-pattern — that is, the application node and the abstraction node — must be either both in s_β or neither in s_β. This guarantees that in case (ii), r_β also has a redex at u.

Suppose that β is a limit ordinal and that each r_γ has been constructed for $\gamma < \beta$. The redex positions in the constructed sequence are a subsequence of those of the original sequence. Since the original sequence is strongly convergent, the constructed sequence converges to a limit r_β. From the definition of contribution, it follows that s_β is the limit of s_γ for $\gamma < \beta$. Since $s_\gamma \leq r_\gamma$, $s_\beta \leq r_\beta$. □

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 38 *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 37 and 18. □

Corollary 39 *If a finite term is reducible to a finite normal form, it is reducible to that normal form in the finitary lambda calculus.* □

6 The compression property

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 of a reduction sequence may be obscure if its length is greater than ω. Such a sequence completes an infinite amount of work and then does 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 40 (Compression 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 [8] 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 [6]).

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 $s_0 \to^* C[(\lambda x.t)r, t_1, \ldots, t_n] \to^\omega_{d+1} C[(\lambda x.t')r', t'_1, \ldots, t'_n] \to_d C[t'[x := r']]$. 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 $C[(\lambda x.t)r, t_1, \ldots, t_n]$ to $C[(\lambda x.t')r', t'_1, \ldots, t'_n]$ consists of an interleaving of reductions of t to t', r to r', and each t_i to t'_i of length at most ω. Conversely, any reductions of lengths at most ω starting from t, r, and each t_i can be interleaved to give a reduction of length at most ω starting from $C[(\lambda x.t)r, t_1, \ldots, t_n]$. The theorem will therefore be established if, given reductions of t to t' and r to r' of length at most ω, we can construct a reduction from $(\lambda x.t)r$ to $t'[x := r']$ of length at most ω. This can be done by first reducing $(\lambda x.t)r$ to $t[x := r]$, and then interleaving a reduction of t to t' and reductions of all the copies of r to r' in a strongly convergent way.

A suitable reduction sequence can be devised as follows. Begin with the term $t[x := r]$. For each i from 0 upwards, we construct a finite number of steps of the desired reduction sequence starting from $t[x := r]$. At each stage, we will have a term t_i which is obtained from some term s in the reduction from t to t', by replacing every free occurrence of x in s by some term in the reduction of r to r' (possibly a different such term for different occurrences of x). Clearly, $t_0 = t[x := r]$ is such a term. Suppose that t_i has been defined, and is constructed from s and various reducts of r as just described. Take a finite initial segment of the reduction of s to t' which include every reduction at depth at most i. (By strong convergence, there must be such a segment.) Suppose that this reduction reduces s to a term s'. Then the same reduction can be performed on t_i to obtain a term t'_i. t'_i is obtained from s' by replacing free occurrences of x in s' by terms occurring in the reduction of r to r'. For each free occurrence of x in s' at any depth $d \leq i$ (and there can be only finitely many of these), let r'' be the reduct of r at that occurrence of t'_i. Perform some finite initial segment of the reduction of r'' to r' on that copy of r'', long enough to ensure that all subsequent reductions in that sequence would take place at a depth greater than i. Again, strong convergence ensures that this is possible.

The resulting term is t_{i+1}. By carrying this out for all i, we obtain a reduction sequence in which after each term t_i, every reduction takes place at a depth of at least i, and is therefore strongly convergent. The limit is clearly $t'[x := r']$. $\square$

The Compressing Lemma is false for beta-eta-reduction. Eta reduction is the rule that reduces $\lambda x.(tx)$ to t if x is not free in t. For a counterexample, let $t = Y(\lambda f.\lambda x.I(fx))$. Then $\lambda x.txx \to^{\omega} \lambda x.I(I(I(...)))x \to_{\eta} I(I(I(...)))$. However, $\lambda x.txx$ 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.

7 The Church-Rosser property and Böhm reduction

7.1 Stable and active terms

At the end of section 4 we gave counterexamples to the Church-Rosser property for all the infinitary lambda calculi, and remarked that the property does hold up to equality of a certain set of "meaningless" terms. Here we define and study that class and prove the claim.

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 $\perp$ which we introduced in Section 5. 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 characterisation as more important, and call it *0-stability.*

Definition 41 *A* 0-redex *of a term is a beta redex or an occurrence of* $\perp$ *at depth 0. A term of* $\Lambda^{\infty}_{\perp}$ *is* 0-stable *if it cannot be beta reduced to a term containing a 0-redex. It is* 0-active *if it cannot be beta reduced to a 0-stable term.*

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

We now generalise the traditional concept of Böhm reduction.

Definition 42 Böhm reduction *is reduction in* $\Lambda^{\infty}_{\perp}$ *by the* β *rule and the* $\perp$ *rule, viz.* $t \to \perp$ *if* t *is 0-active and not* $\perp$*. We write* $\to_{\mathcal{B}}$ *for Böhm reduction and* $\to_{\perp}$ *for reduction by the* $\perp$*-rule alone.*

A Böhm tree *is a normal form of* $\Lambda^{\infty}_{\perp}$ *with respect to Böhm reduction.*

For clarity, we may also write $\to_{\beta}$ *instead of* $\to$ *for beta 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 43 *For depth measures **1, the set of 0-active terms is closed under substitution.*

Proof. See Fig. 4. To reduce clutter, all arrows in this and similar figures represent reductions of arbitrary length.

Suppose that t is 0-active. Consider any instance $\theta(t)$ of t and any reduction $\theta(t) \to_\beta^\infty s'$. 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 of r such that r and r' differ only in the subterms at U. 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, 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, 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. This means that the redex node is outside the subterms at U, but its rator node is in one of those subterms. In this case, both u and $u \cdot 1$ are positions 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, and omit from U 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', such that for each $u \in U'$, $s|u$ has the form $xt_1 \ldots t_n (n \geq 0)$ 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 $\theta(s) \to_\beta^\infty s'$ by reductions taking place entirely within the subterms at U 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. 4, where q and q' differ in the same manner that s and s' differed.

Because the depth measure is $**1$, the subterms of q at U_q, being all of the

$$\begin{tikzcd}
t \arrow[r, "\beta"] & s \arrow[r, "\beta"] & q \\
\theta(t) \arrow[r, "\beta"'] \arrow[dr, "\beta"'] & \theta(s) \arrow[r, "\beta"'] \arrow[d, "\beta"] & \theta(q) \arrow[d, "\beta"] \\
& s' \arrow[r, "\beta"'] & q'
\end{tikzcd}$$

Fig. 4. Lemma 43

$$\begin{tikzcd}
t \arrow[rr, "\beta"] \arrow[d, "\perp"'] & & t' \arrow[d, dotted, "\perp"'] \\
s \arrow[r, dotted, "\beta"'] & s' \arrow[r, dotted, "\perp"'] & \bullet
\end{tikzcd}$$

Fig. 5. Lemma 45

form $xt_1 \ldots t_n$, 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. $\square$

Definition 44 *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. We write $t \sim s$, or $t \sim_U s$ if we wish to specify U explicitly.*

Lemma 45 *For depth measures 001, 101, and 111, if t and s are equivalent, and $t \to_\beta^\infty t'$, then for some s' equivalent to t', $s \to_\beta^\infty s'$. The latter reduction can be chosen so as to reduce no redexes inside 0-active subterms.*

If t and s are not merely equivalent, but a fortiori *$t \to_\perp^\infty s$, then Fig. 5 can be formed.*

Proof. Assuming the hypotheses, we imitate the reduction of t to t' on s. Suppose we have a step $t_0 \to_\beta t_1$, and a term s_0 equivalent to t_0. If the beta redex is inside one of the 0-active subterms 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 43, 0-active terms are closed under substitution) must be equivalent to t_1. Finally, suppose the redex has the form $(\lambda x.p)q$, where $\lambda x.p$ is one of the 0-active subterms at which t_0 differs from s_0. Let the subterms of s_0 corresponding to $\lambda x.p$ and q be p' and q'. For $\lambda x.p$ to be 0-active, the depth measure must be 001. p' is 0-active. For depth 001, this implies that $p'q'$ is also 0-active (since 0-active terms for depth 001 are just the terms without head normal form). $(\lambda x.p)q$ 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.

When $t \to_{\perp}^{\infty} s$, the above construction yields Fig. 5. $\square$

A counterexample for the depth measures $**0$ follows immediately from the corresponding counterexample for Lemma 43. Take $t = (\lambda x.x\Omega)(KI)$ and $s = (\lambda x.\Omega)(KI)$. Since $x\Omega$ and Ω are 0-active for depths $**0$, t and s are equivalent. However, t reduces to I, but s reduces only to Ω. Neither s nor Ω is equivalent to I.

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.

Lemma 46 *For depth measures 001, 101, and 111:*

(i) If t and s are equivalent, then t is 0-stable if and only if s is 0-stable.
(ii) If t and s are equivalent, then t is 0-active if and only if s is 0-active.
(iii) Lemma 45 also holds when the given reduction of t to t' is a Böhm reduction.

Proof.

(i) 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 45, 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.

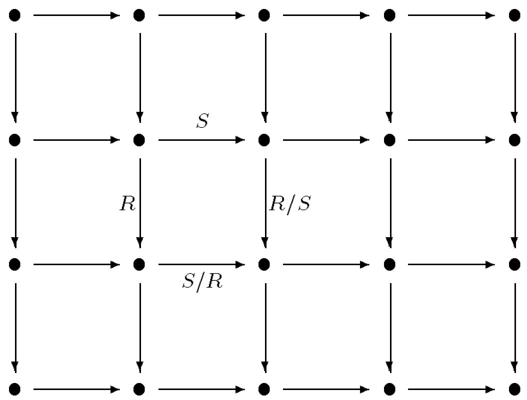

Fig. 6. A tiling diagram

(ii) Suppose s is not 0-active. Then s reduces to a 0-stable term r. By Lemma 45, t reduces to a term equivalent to r, which by part (i) must be 0-stable. Therefore t is not 0-active.

(iii) The proof of Lemma 45 can be extended to handle Böhm reductions, using part (ii) to justify omitting all $\perp$-reductions when constructing the sequence from s.

□

7.2 The Church-Rosser property, up to equality of 0-active subterms

In this section we will prove that the Church-Rosser property holds for depth measures 001, 101, and 111, up to equality of 0-active subterms.

In the finitary case a classic strategy for proving confluence (see [2], Chapter 11) is via the Finite Developments theorem, which says that all complete developments are finite and end in the same term. From this, finitary confluence follows immediately by means of a "tiling diagram" such as Figure 6, in which each square consists of complete developments as indicated in one of the tiles of that figure. The top and left sides of the figure are any two finite coinitial sequences, and the diagram constructs the right and bottom sides, which end with the same term.

In the infinitary case, part of the Finite Developments theorem still holds, Theorem 22, the Infinite Developments Theorem. However, as we saw earlier, infinite terms may contain sets of redexes that have no complete development. We overcome this difficulty by defining a modified version of lambda calculus in which all reductions are strongly convergent, and hence all sets of redexes have a complete development. This allows us to establish an Infinite Developments Theorem for the modified calculus which is strong enough to obtain confluence. Consideration of the relationship between the modified and unmodified calculi then allows an approximate confluence property to be derived, for depth measures 001, 101, and 111. For the other depth measures, counterexamples even to the approximate form of confluence can be given.

Definition 47 *The ϵ^{abc}-calculus is obtained from Λ^{abc} by adding a new unary function symbol ϵ, stipulating that the depth of t in $\epsilon(t)$ is 1, and replacing the beta rule by the family of rules $(\epsilon^n(\lambda x.s))t \to \epsilon^{n+2}(s[x := t])$ for all $n \geq 0$.*

We drop the depth indication when we do not wish to specify it. In the ϵ-calculus, 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 ϵ-calculus is strongly convergent. In particular, every set of redexes has a complete development, which makes for a simple proof that the Church-Rosser property holds exactly, for all measures of depth.

Theorem 48 *(Infinite Developments Theorem for ϵ-calculus.) Complete developments of the same set of ϵ-redexes end at the same term.*

Proof. Similar to the proof of Theorem 22, modifying the notion of level path to disallow a path being extended through an ϵ node. In fact, the proof is somewhat simpler, since concerns about non-strongly convergent maximal developments do not arise. $\square$

Theorem 49 *The ϵ-calculus satisfies the Church-Rosser property, for all depth measures.*

Proof. Suppose we are given two ϵ-reduction sequences starting from the same term. We can consider them as concatenations of complete developments (since any single step is a complete development). Write $\to_R$ for a complete ϵ-development of a set of ϵ-redexes R. Let the sequences be concatenations of developments $t_{\gamma,0} \to_{S_{\gamma,0}} t_{\gamma+1,0}$ and $t_{0,\delta} \to_{R_{0,\delta}} t_{0,\delta+1}$, where $0 \leq \gamma < \alpha$ and $0 \leq \delta < \beta$.

Then define developments $t_{\gamma,\delta} \to_{R_{\gamma,\delta}} t_{\gamma,\delta+1}$ and $t_{\gamma,\delta} \to_{S_{\gamma,\delta}} t_{\gamma+1,\delta}$ inductively as follows.

$R_{\gamma+1,\delta} = R_{\gamma,\delta}/S_{\gamma,\delta}$

$S_{\gamma,\delta+1} = S_{\gamma,\delta}/R_{\gamma,\delta}$

For a limit ordinal $\mu \leq \alpha$, $R_{\mu,\delta} = \lim_{\gamma<\mu} R_{\gamma,\delta}$

For a limit ordinal $\mu \leq \beta$, $S_{\gamma,\mu} = \lim_{\delta<\mu} S_{\gamma,\delta}$

Strong convergence of all reduction sequences ensures that all of these sets of redexes have complete developments, and that the limits exist.

By Theorem 48, $R_{\gamma+1,\delta}$ and $S_{\gamma,\delta+1}$ end at the same term $t_{\gamma+1,\delta+1}$.

When γ and δ are both limit ordinals, we must also show that $t_{\gamma,\delta}$ is well-defined, i.e. that the reduction sequences through the terms $t_{\gamma',\delta}$ ($\gamma' < \gamma$) and $t_{\gamma,\delta'}$ ($\delta' < \delta$) have the same limit. Call the two limits $t^0_{\gamma,\delta}$ and $t^1_{\gamma,\delta}$. By strong convergence of all the sequences, for any $\zeta > 0$, for large enough γ' and δ', all of the following distances are less than ζ: $d(t_{\gamma',0}, t_{\gamma,0})$, $d(t_{0,\delta'}, t_{0,\delta})$, $d(t_{\gamma',\delta}, t^0_{\gamma,\delta})$, and $d(t_{\gamma,\delta'}, t^1_{\gamma,\delta})$. Because residuals of a redex have depth at least as great as the redex, it follows that $d(t_{\gamma',\delta'}, t_{\gamma,\delta'}) < \zeta$ and $d(t_{\gamma',\delta'}, t_{\gamma',\delta}) < \zeta$. Hence $d(t^0_{\gamma,\delta}, t^1_{\gamma,\delta}) < 4\zeta$. Since this is true for all $\zeta > 0$, $t^0_{\gamma,\delta} = t^1_{\gamma,\delta}$.

Thus the inductive construction can be carried out for all ordinals up to α and β, yielding strongly convergent reduction sequences from $t_{\alpha,0}$ and $t_{0,\beta}$ to $t_{\alpha,\beta}$. $\square$

The next step is to transfer the Church-Rosser property of ϵ-calculus to an approximate Church-Rosser property for Λ^∞.

Definition 50 *For any position u of an ϵ-term t, $\epsilon(u,t)$ is the position obtained by omitting every occurrence of 1 in u which corresponds to an occurrence of ϵ in t.*

For an ϵ-term t, if t does not contain the subterm $\epsilon^\omega = \epsilon(\epsilon(\epsilon(\cdots)))$, then $\epsilon(t)$ is the Λ^∞ term resulting from omitting all occurrences of ϵ.

Clearly, when $\epsilon(t)$ is defined and u is a position of t, $\epsilon(u,t)$ is a position of $\epsilon(t)$ having the same symbol as t does at u. Dropping the ϵs from an ϵ-reduction yields a beta reduction. Similarly, beta reductions can be mapped to ϵ-reductions.

Lemma 51 *For any ϵ-reduction $t \to^\infty s$ in which ϵ^ω does not occur in any term, there is a beta reduction $\epsilon(t) \to^\infty \epsilon(s)$ of the same length.*

Where the original sequence contains an ϵ-reduction at position u in a term r, the corresponding beta reduction is at $\epsilon(u,r)$.

For any beta reduction $t \to^\infty s$, and any term t' such that $t = \epsilon(t')$, there is an ϵ-reduction of the same length from t', which is mapped back to the given beta reduction by the preceding mapping. $\square$

Definition 52 *Two ϵ-terms t and t' are equivalent, $t \sim t'$, if, when every occurrence of ϵ^ω is replaced by Ω and every other occurrence of ϵ is omitted,*

$$\begin{tikzcd} t \arrow[r, "*"] \arrow[d, "h"', "*"] & s \arrow[d, dotted, "h"', "*"] \\ t' \arrow[r, dotted, "*"] & s' \end{tikzcd}$$

Fig. 7. Lemma 53

the resulting Λ^∞ terms are equivalent by Definition 44.

For Λ^∞ terms, this definition coincides with Definition 44, and so we may use the same notation.

To proceed further, we must establish some facts about head reduction. The head redex of a term t (if it exists) can be defined as the unique redex of t at depth 0 with respect to the measure 001. (This definition applies whether or not t is a term of Λ^{001}). We write $\to_h$ for the reduction of the head redex.

Lemma 53 *See Fig. 7.*

(i) If $t \to_h^* t'$ and $t \to^\infty s$, then there are reductions $t' \to^\infty s'$ and $s \to_h^* s'$.
(ii) The steps of $s \to_h^* s'$ are in $1-1$ correspondence with a subset of steps of $t \to_h^* t'$, and each step of the former is at the same position as its corresponding step in the latter.
(iii) If the reduction of t to s is finite and has n steps, then the length of $s \to_h^* s'$ is at least the length of $t \to_h^* t'$ minus n.

Proof. For part (i) it is sufficient to prove the case where $t \to_h^* t'$ is exactly one step. By the Compression Property, $t \to^\infty s$ may be assumed to have length at most ω.

A head redex r has at most one residual by reduction of any redex r', If present, that residual is at the same position as r. r' has a set of pairwise disjoint residuals by r. Therefore complete developments of both r/r' and r'/r exist, the former being at most one step. Repeating this construction along every step of $t \to^\infty s$ yields a reduction $t' \to^\infty s'$, such that either $s = s'$ or $s \to_h s'$ by reduction at the same position as r. Finally, the first statement of the lemma follows by repeating this construction for each step of the head reduction from t.

Part (ii) is a corollary of the construction.

For part (iii), suppose that t reduces to s by reduction of a single redex r. Consider the set of residuals R of r by an initial segment of the head reduction of t, and the next head redex r' in the head reduction of t. After each such head

$$
\begin{tikzcd}
t \arrow[r, "*"] \arrow[d, "h"', "\leq\omega"] & t'' \arrow[r, "\leq\omega"] \arrow[d, "h"', "\leq\omega"] & t' \arrow[d, "h"', "\leq\omega"] \\
\vdots & \vdots & \vdots
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
t \arrow[r, "*", "h"'] & t'' \arrow[r, "*"] \arrow[d, "h"', "\omega"] & s \arrow[d, "h"', "\omega"] \\
& \vdots & \vdots
\end{tikzcd}
$$

(a) *(b)*

Fig. 8. Lemma 54

reduction, either no member of R is a head redex, or $R = \{r'\}$, or R is empty. In the first and third cases, there is a single step of $s \rightarrow_h^* s'$ corresponding to r', and in the second case, which can occur at most once (since thereafter r will have no residuals), there is an empty step. Therefore the length of $s \rightarrow_h^* s'$ is at least the length of $t \rightarrow_h^* t'$ minus 1. Repeating the argument for each step of $t \rightarrow^* s$ yields the conclusion. $\square$

Lemma 54 *For depth measures 001, 101, and 111, if t is reducible to a 0-stable term, it is so reducible by a finite head reduction sequence.*

Proof. For depth measures 001 and 101, 0-stability of a term is determined by its prefix to depth 0. By the Compression Property, the reduction of t to t' may be assumed to be of length at most ω, and therefore that t reduces to a 0-stable term t" in finitely many steps.

Now consider the reduction starting from t which at each step reduces the head redex, if that redex is at depth 0, and terminates if there is no such redex. If this sequence is infinite, then by Lemma 53, every term to which t reduces in finitely many steps also has an infinite reduction of the same form. Therefore t cannot reduce in finitely many steps to a 0-stable term, contradicting the 0-stability of t".

For depth measure 111, the 0-stable terms are those which cannot be reduced to a redex. Consider the maximal head reduction sequence of t. Either this performs infinitely many reductions at the root, or it does not.

Consider the first case. See Fig. 8(a). The reduction of t to t' is strongly convergent, and by the Compressing Property can be assumed of length at most ω, so after finitely many steps, it performs no more root reductions. Apply Lemma 53 to the head reduction of t and that finite initial segment. The result is a head reduction of a term t'' performing infinitely many root reductions and a reduction of t'' to t' performing no root reductions. Applying Lemma 53 again yields a head reduction of t' including infinitely many root reductions, contradicting its 0-stability.

Now suppose the head reduction of t performs only finitely many root reductions. If it terminates in finitely many steps, then the resulting term is 0-stable. Suppose it is infinite (and not necessarily strongly convergent). See Fig. 8(b). There is a finite initial segment $t \to_h^* t''$ containing all the root reductions of that sequence. Suppose that t'' were reducible to a redex s. Then it is so reducible in finitely many steps. Apply Lemma 53 to the head reduction of t'' and such a reduction. We obtain an infinite head reduction of s in which every step is performed below the root. But s is a redex, therefore its first head reduction must be at the root, contradiction. $\square$

Lemma 55 *For depth measures 001, 101, and 111, the complement of the set of 0-active terms is closed under reduction.*

Proof. Suppose $t \to^\infty t'$, $t \to^\infty t''$, t' is 0-active, and t'' is 0-stable.

By Lemma 54, the reduction of t to t'' can be assumed to be a finite head reduction. By Lemma 53, t'' and t' are reducible to a common term. But this cannot happen if t' is 0-active and t'' is 0-stable. $\square$

Lemma 56 *For depth measures 001, 101, and 111, for any ϵ-reduction $t \to^\infty s$, there is a beta reduction $t' \to^\infty s'$, such that $t' \sim t$ by replacing every occurrence of ϵ^∞ by a 0-active term and omitting every other occurrence of ϵ, and $s' \sim s$.*

Proof. We construct the reduction inductively. Let the given reduction be of length ω, its nth step being $t_n \to_\epsilon t_{n+1}$ by reduction at u_n.

Construct t'_0 by replacing every occurrence of ϵ^ω in t by Ω and omitting every other occurrence of ϵ.

Suppose t'_n has been defined and is equivalent to t_n. If some prefix of $\epsilon(u_n, t_n)$ is the position of a 0-active subterm of t'_n, then take $t'_{n+1} = t'_n$. Otherwise, $\epsilon(u_n, t_n)$ is the position of a beta redex of t'_n, not contained in a 0-active subterm. Reduce it to obtain t'_{n+1}.

Because 0-active terms are closed under substitution and reduction, $t'_{n+1} \sim t_{n+1}$.

This constructs a reduction of t'_0 in which no step is performed inside any 0-active subterm. Such a sequence must be strongly convergent. Let its limit be t'_ω. Let t''_ω be obtained from t_ω by replacing every occurrence of ϵ^ω by Ω and omitting every other occurrence of ϵ. If t'_ω were not equivalent to t''_ω, then there would be a position u of both terms such that one of t'_ω and t''_ω was 0-active and the other was not, neither subterm being properly contained in a

$$
\begin{tikzcd}
& t \arrow[dl, "\beta"'] \arrow[d, "\epsilon"] \arrow[rr, dash, "\sim"] & & t' \arrow[d, "\epsilon"'] \arrow[dr, "\beta"] & \\
s \arrow[dr, "\beta"'] & q \arrow[l, mapsto, "\epsilon"'] \arrow[dr, "\epsilon"] & & q' \arrow[r, mapsto, "\epsilon"] \arrow[dl, "\epsilon"'] & s' \arrow[dl, "\beta"] \\
& p \arrow[r, dash, "\sim"'] & r \arrow[r, dash, "\sim"'] & p' &
\end{tikzcd}
$$

Fig. 9. Theorem 57: the Church-Rosser property, up to equivalence

0-active subterm. But that would imply that for some finite n, t_n and t'_n were not equivalent, since by Lemma 55 the 0-active terms and their complement are closed under reduction. This is contrary to the construction. □

Theorem 57 *For depth measures 001, 101, and 111: Let $t \sim t'$, $t \to_\beta^\infty s$, and $t' \to_\beta^\infty s'$. Then there exist equivalent terms p and p', and beta reductions of s to p and s' to p'.*

Proof. See Figure 9. Suppose we are given two beta reductions $t \to_\beta^\infty s$ and $t' \to_\beta^\infty s'$, with $t \sim t'$. Lemma 51 constructs ϵ-reductions $t \to_\epsilon^\infty q$ and $t' \to_\epsilon^\infty q'$ such that $\epsilon(q) = s$ and $\epsilon(q') = s'$. By Theorem 49, there are ϵ-reductions from q and q' to a common ϵ-term r.

It remains to map these sequences back to strongly convergent sequences starting from s and s' in the original calculus. The obstacle to applying Lemma 51 is that r may contain ϵ^ω as a subterm. But by Lemma 56, there are beta reductions from s and s' to Λ^∞ terms p and p' respectively, such that each is equivalent to r. Therefore they are equivalent to each other. □

For all other depth measures, the counterexamples provided for Lemma 45 are also counterexamples to this theorem. For 000, failure of this theorem amounts to the well-known fact of the inconsistency of finitary lambda calculus under the additional axiom that terms without normal form are equal ([2], Proposition 2.2.4).

7.3 The Church-Rosser property and Böhm reduction

Lemma 58 *For depth measures 001, 101, and 111:*

(i) The set of 0-active terms is closed under Böhm reduction.

(ii) The complement of the set of 0-active terms is closed under beta reduction and $\perp$-reduction.

Proof.

(i) Immediate from Lemma 46(iii).

(ii) Closure under $\perp$-reduction follows from Lemma 46(ii).

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 57, 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 59 *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 $\perp$. 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 $(\lambda y.y\Omega)(KI) \rightarrow_\beta^* I$ and $(\lambda y.y\Omega)(KI) \rightarrow_\perp \perp(KI)$, 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 $(\lambda x.\Omega)y \rightarrow_\beta \Omega \rightarrow_\perp \perp$ and $(\lambda x.\Omega)y \rightarrow_\perp \Omega y \rightarrow_\perp \perp y$. Both $\perp$ and $\perp y$ are Böhm normal forms. This also refutes the Church-Rosser property of Böhm reduction for these depth measures.

Lemma 60 *For depth measures 001, 101, and 111, $\perp$-reduction is transfinitely Church-Rosser.*

Proof. It is immediate from Lemma 58(i) that if t is $\perp$-reducible to s, it is so reducible by the reduction of a set of $\perp$-redexes at pairwise disjoint positions. Given two $\perp$-reductions $t \rightarrow_\perp^\infty s$ and $t \rightarrow_\perp^\infty s'$, take the set of outermost members of the union of the two associated sets. Reduction of all of these $\perp$-redexes gives a term r which both s and s' are $\perp$-reducible to. □

Theorem 61 *For depth measures 001, 101, and 111, Böhm reduction is transfinitely Church-Rosser.*

$$
\begin{tikzcd}
t \arrow[rrr, "\beta"] \arrow[ddd, "\beta"'] & & & \bullet \arrow[r, "\perp"] \arrow[dd, "\beta"'] & s_0 \arrow[d, "\beta"] \\
& & & & \bullet \arrow[d, "\perp"] \\
& & & \bullet \arrow[r, "\perp"] \arrow[d, "\perp"'] & \bullet \arrow[d, "\perp"] \\
\bullet \arrow[rr, "\beta"] \arrow[d, "\perp"'] & & \bullet \arrow[r, "\perp"] \arrow[d, "\perp"'] & \bullet \arrow[r, "\perp"] \arrow[d, "\perp"'] & \bullet \arrow[d, "\perp"] \\
s_1 \arrow[r, "\beta"'] & \bullet \arrow[r, "\perp"'] & \bullet \arrow[r, "\perp"'] & \bullet \arrow[r, "\perp"'] & \bullet
\end{tikzcd}
$$

Fig. 10. Theorem 61: the Church-Rosser property for Böhm reduction

Proof. Suppose we have two Böhm reductions starting from a term t. By Lemma 46(iii) they can be put into the form $t \to_\beta^\infty \to_\perp^\infty s_0$ and $t \to_\beta^\infty \to_\perp^\infty s_1$.

We then construct Fig. 10. The top left square exists by Theorem 57. The top right and bottom left are given by Lemma 45. The remaining squares follow from Lemma 60. □

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. 36. The usual Böhm model is the model associated with applicative depth, 001. The larger model described by Berarducci ([3]) 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 [1].

8 Acknowledgements

An early stage of this work has been discussed at the NSL workshop in Kanazawa in the autumn of 1994. This paper is a revised and extended version of the paper [7] presented at the conference RTA'95.

We acknowledge insightful discussions with Zurab Khasidashvili, Vincent van

Oostrom, and Femke van Raamsdonk.

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

References

[1] S. Abramsky and L. Ong. Full abstraction in the lazy lambda calculus. *Information and Computation*, 105:159–267, 1993.

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

[3] A. Berarducci. Infinite λ-calculus and non-sensible models. In *Proceedings of the conference in honour of Roberto Magari, Siena*, 1994. To appear in the series on Logic and Algebra, Marcel & Decker, New York.

[4] A. Berarducci and B. Intrigila. Church-Rosser λ-theories, infinite λ-terms and consistency problems. In W. Hodges, M. Hyland, C. Steinhorn, and J. Truss, editors, *Logic Colloquium '93.* Oxford University Press, 1996.

[5] P.L. Curien. *Categorical combinators, sequential algorithms, and functional programming.* Birkhäuser Boston, 2nd edition, 1993.

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

[7] J.R. Kennaway, J.W. Klop, M.R. Sleep, and F.J. de Vries. Infinitary lambda calculus and Böhm models. In *Proc. Conference on Rewriting Techniques and Applications*, pages 257–270, 1995. Lecture Notes in Computer Science, vol. 914.

[8] J.R. Kennaway, J.W. Klop, M.R. Sleep, and F.J. de Vries. Transfinite reductions in orthogonal term rewriting systems. *Information and Computation*, 119(1):18–38, 1995.