Contents

Chapter 1

Infinitary Rewriting

1.1. Introduction

In the previous chapters terms were finite and the focus was on finite computations. Computations are supposed to produce an outcome in a finite amount of time. However, we can imagine infinite computations: think of a computation calculating the decimal expansion of $\sqrt{2}$. Infinite terms also make an intuitive sense in lazy functional programming. Lazy evaluation allows functional programmers to write down expressions which denote infinite data structures, such as a list of all primes. Although only a finite number will ever be evaluated, one can consider an infinite computation as tending to a limit in which the entire infinite list has been built. The following rewrite rules illustrate this.

$$
\begin{array}{rcl}
\mathit{filter}(x:y,0,m) & \to & 0:\mathit{filter}(y,m,m) \\
\mathit{filter}(x:y,s(n),m) & \to & x:\mathit{filter}(y,n,m) \\
\mathit{sieve}(0:y) & \to & \mathit{sieve}(y) \\
\mathit{sieve}(s(n):y) & \to & s(n):\mathit{sieve}(\mathit{filter}(y,n,n)) \\
\mathit{nats}(n) & \to & n:\mathit{nats}(s(n)) \\
\mathit{primes} & \to & \mathit{sieve}(\mathit{nats}(s(s(0))))
\end{array}
$$

This is a version of the sieve of Eratosthenes in which the filter function replaces struck-out members of the list by zeroes, which are subsequently deleted by the sieve function. (It is a natural implementation of the hand-calculation one would perform on a blackboard, avoids the necessity of doing any arithmetic other than counting, and is as efficient as the usual method of performing division tests. Despite its simplicity it appears to be novel.)

Although the lazy order of evaluation will produce each prime after computing only a finite part of the list of all natural numbers, the expression $\mathit{nats}(s(s(0)))$ can be thought of as denoting the infinite list of all natural numbers. Longer and longer evaluations of the term approximate closer and closer to that value.

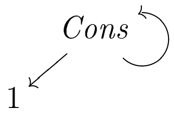

Figure 1.1: A cyclic graph

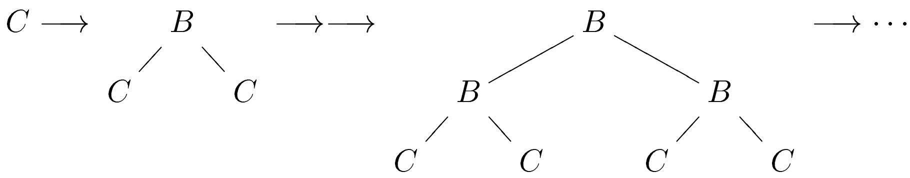

Figure 1.2: An infinite reduction

Another situation in which infinite terms may arise is in graph rewriting. Representing terms by directed acyclic graphs allows repeated subterms to be represented by pointers to the same subgraph, with a saving of both space and time, since reductions performed within that subgraph are equivalent to two or more reductions of the corresponding term. However, once one begin to use graphs to represent terms, the possibility arises of using cyclic graphs to represent infinite terms. The infinite list $Cons(1, Cons(1, Cons(1, \ldots)))$ can be more compactly represented by the cyclic graph in Figure 1.1, which we can write with the notation $x : Cons(1, x)$. The relation between cyclic graphs and the rewriting of infinite terms is explored in Kennaway et al. [1994].

It is therefore natural to inquire whether one can extend the theory of rewriting to allow infinite terms and infinite reduction sequences which converge to limits.

As an example of the kind of infinite reductions we have in mind, consider the rule

$$C \to B(C, C)$$

Parallel outermost reduction of the term C gives an infinite reduction to a binary tree, as shown in Figure 1.2. This limit is reached in infinitely many reduction steps and is clearly a normal form. However leftmost innermost reduction leads in infinitely many steps to the term on the right in Figure 1.3. This limit is not a normal form and can be reduced further, which gives a reduction sequence of length greater than ω.

In this chapter we will give the basic definitions and properties of infinite terms and infinite reduction sequences, for both term rewrite systems and the lambda calculus. We will then study confluence properties in orthogonal systems, which turns out to be significantly more complicated than in the finitary case. In general, these systems are only confluent up to an identifi-

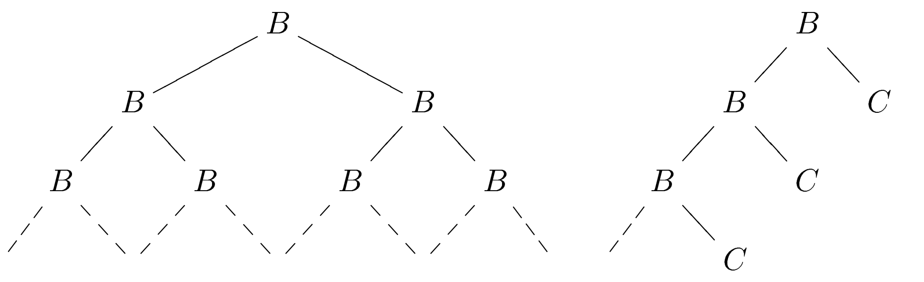

Figure 1.3: Two infinite terms

cation of a certain class of terms. The breakdown of confluence leads us to consider the concept of a meaningless term, which further suggests a link with the lambda-calculus concept of Böhm reduction (Barendregt [1984]), and to denotational semantics for TRSs.

1.2. Infinitary term rewriting systems

As in the finitary setting, an infinitary term rewriting system will be given by a signature and a set of rewrite rules, but the signature is now understood to generate a class of finite and infinite terms. The signature may also contain infinitely many symbols.

In this setting it may seem natural to consider rewrite rules which are pairs of possibly infinite terms. However, based on the computational intuition that performing a reduction should require only a finite amount of work, we will consider rewrite rules with finite left- and right-hand sides only. We also restrict attention to left-linear systems, since the test for syntactic equality required by a non-left-linear rule cannot always be done in finite time when terms are infinite.

But first, we must define infinite terms. A finite term can be described as the set of its positions, together with a function from that set to the function symbol at each position. The set of positions must be finite and prefix-closed, and if position u maps to symbol F, then $u \cdot i$ must be in the set of positions if and only if i is between 1 and the arity of F. The class of finite and infinite terms can be obtained simply by dropping the requirement that there be only finitely many positions. Each position is still finite, but the set of positions can be unbounded in length. This yields infinite terms such as those illustrated in Figure 1.3. We allow the signature to be an infinite set, but still require that the arity of each symbol be finite, although the set of arities is allowed to be unbounded.

A more formal mathematical construction of the set of finite and infinite terms can be given by defining a metric on finite terms and taking the metric completion. The set $Ter(\Sigma)$ of finite terms over Σ can be provided with a metric $d : Ter(\Sigma) \times Ter(\Sigma) \to [0,1]$. Define $d(t,t') = 0$ if the terms are identical, otherwise $d(t,t') = 2^{-k}$, where k is the length of the shortest position at which the two terms differ.

1.2.1. THEOREM. *The metric completion of $Ter(\Sigma)$ is isomorphic to $Ter^\infty(\Sigma)$.*

PROOF. The definition of the metric can be applied verbatim to $Ter^\infty(\Sigma)$ and yields a metric on that set. Since the metric completion is unique up to isomorphism, we need only show that $Ter^\infty(\Sigma)$ with this metric is complete.

Suppose we have a Cauchy sequence of finite or infinite terms $\{ t_i \mid i \in \mathcal{N} \}$. Let p_i be the set of positions of t_i, and f_i map p_i to the function symbols of t_i. If $d(t_i,t_j) < 2^{-k}$, then p_i and p_j must have the same members of lengths up to k, and f_i and f_j must agree on those members. Define $p = \bigcup_i \bigcap_{j \geq i} p_j$, which by the above must be the same as $\bigcap_i \bigcup_{j \geq i} p_j$. For each $u \in p$, there must be an i such that for $j \geq i$, $u \in p_j$ and $f_j(u) = f_i(u)$. Define $f(u) = f_i(u)$. The conditions for p and f to define a (possibly infinite) term t are easily verified, as is the property that if $d(t_i,t_j) < 2^{-k}$, then $d(t_i,t) < 2^{-k}$. t is therefore the limit of the sequence. ⊠

$Ter^\infty(\Sigma)$ includes the finite terms, and infinite terms like those of Figures 1.2 and 1.3. Note that there are no nodes at the end of an infinite path in a term. The position of each node in a term is a finite string, and we do not need to consider infinite positions. There is in this framework no such term as $A(A(A(\ldots$ *infinitely many A's* $\ldots(B)\ldots)))$; such terms have no obvious computational meaning and do not appear in the metric completion of $Ter(\Sigma)$.

1.2.2. DEFINITION. An *infinitary reduction rule* (or rewrite rule) is a pair (t,s) where $t \in Ter(\Sigma)$ and $s \in Ter^\infty(\Sigma)$, such that t is not a variable, and every variable in s occurs in t.

An *infinitary Term Rewriting System*, or *iTRS*, is a pair $\mathcal{R} = (\Sigma, R)$ of a signature Σ and a set of infinitary reduction rules R.

Many concepts generalise immediately to the infinitary setting: context, position, redex, reduction step, normal form left-linear rule, non-duplicating rule, and (non-)overlapping rules, orthogonality. Many results also carry over.

Proofs must sometimes be changed, since structural induction over terms is not well-founded for infinite terms. In many cases, such inductions can be replaced by inductions over positions of a term, since positions are always finite. For example, to prove that two terms are equal, in the finitary case one can prove that they have the same symbol and arity at the root, and assume by induction that their corresponding subterms are equal. In the infinitary case, one could instead prove that for every position common to both terms, the terms have the same symbol and arity at that position. Induction over the length of positions, with the base case being the root position, establishes that all corresponding finite prefixes of the terms are identical, hence the terms themselves are identical. Similarly, if a proof requires the construction of some term, in the finitary case it can be constructed from the leaves upwards or from the root downward. In the infinitary case, only the latter construction is valid.

We shall henceforth drop the word 'infinitary'; all TRSs we consider in this chapter will be infinitary. The TRSs of previous chapters will be referred to as finitary TRSs.

1.2.3. Exercise. Recall that finitary orthogonal TRSs are finitely confluent, that is, $\twoheadrightarrow \cdot \twoheadleftarrow \; \subseteq \; \twoheadleftarrow \cdot \twoheadrightarrow$. Give an example of an orthogonal iTRS which is not finitely confluent. Show that if all its reduction rules are finitary, then an orthogonal iTRS is finitely confluent.

1.3. Strongly converging reductions

As well as infinite terms, we wish to consider infinitely long reduction sequences which may converge to limits.

In our construction of infinite terms, we required each occurrence of a function symbol to be at finite depth, i.e. distance from the root, considering terms containing positions of length ω or greater to be without computational meaning. A similar criticism might be raised against reduction sequences longer than ω. However, the infinite terms as we have defined them do not require us to consider any larger set, since the limit of any sequence of terms in $Ter^{\infty}(\Sigma)$ is another term in $Ter^{\infty}(\Sigma)$. Terms with function symbols at infinite depth do not arise. In contrast, there is no reason to expect the limit of an infinite reduction sequence to be in normal form. Reducing a redex in the limit term at once gives a sequence of length $\omega + 1$. Thus reductions of length greater than ω cannot be ruled out of consideration. In general, a reduction sequence might be of any ordinal length (although we will later prove that according to the definition we are about to give, they will never be of uncountable length). As for their computational meaning, for orthogonal systems the Compression Property (Theorem 1.7.1) shows that they are equivalent to sequences of length no longer than ω.

1.3.1. Definition. A *transfinite reduction sequence* of length α, where α is any ordinal number, is a sequence of reduction steps $(t_\beta \to t_{\beta+1})_{\beta<\alpha}$. In the step $t_\beta \to t_{\beta+1}$, let the redex be reduced at position u_β of t_β. Let d_β be the length of u_β (called the *depth* of the redex).

The sequence is *weakly continuous* if, for every limit ordinal $\lambda < \alpha$, the distance $d(t_\beta, t_\lambda)$ tends to 0 as β approaches λ from below.

It is *strongly continuous* if in addition, d_β tends to infinity as β approaches λ from below.

It is *weakly convergent* or *strongly convergent* if the corresponding continuity condition holds for every limit ordinal $\lambda \leq \alpha$. (Thus convergence is only distinct from continuity when α is itself a limit ordinal.)

We write $t \twoheadrightarrow\!\!\!\rightarrow t'$ for a strongly convergent reduction of any finite or infinite length, and $t \rightarrow^\alpha t'$ for a strongly convergent reduction of ordinal length α; $t \rightarrow^{\leq\alpha} t'$ and $t \rightarrow^{<\alpha} t'$ denote strongly convergent reductions whose lengths satisfy the indicated condition. As in previous chapters, the notation $t \twoheadrightarrow t'$ still refers to reduction sequences that are at most finitely long (of possibly infinite terms).

1.3.2. Remark. Readers familiar with ordinal topology will recognise that a weakly continuous or convergent sequence is one which is continuous or convergent with respect to the usual topology on ordinals and the metric on terms.

This concise definition may become clearer by the light of some examples. Firstly, all finitely long reductions are (trivially) strongly convergent. Now consider a reduction $t_0 \rightarrow t_1 \rightarrow \ldots$ of length ω. For this to be weakly convergent to a limit t_ω requires that for every depth d, there is some n such that all the terms of the sequence after t_n are identical down to depth d. For it to be strongly convergent requires, in addition, that not only are those terms identical down to depth d, but that the position of each reduction from that point onwards has depth greater than d.

Figures 1.4 and 1.5 give an example of the distinction between weak and strong convergence. Both of these sequences are weakly convergent, since larger and larger prefixes of the terms are identical as one proceeds along the sequences. The first is not strongly convergent, since every reduction step takes place at depth 0. The second is strongly convergent, since the depths of the reduction steps tend to infinity.

Now consider a reduction $t \rightarrow^\omega t' \rightarrow^\omega t''$, of length $\omega+\omega$. This is weakly or strongly convergent if and only if both of its ω-long halves has that property. In general any concatenation of finitely many weakly or strongly convergent sequences will be weakly or strongly convergent respectively.

For a reduction $t_0 \rightarrow^\omega t_1 \rightarrow^\omega \ldots$ of length ω^2 to be weakly or strongly convergent, it is necessary for each of the ω-long segments to be so. This is not sufficient: we also require that the continuity condition hold at ω^2, and for strong convergence, the depth condition also. This amounts to the condition that for every depth d, there is an $\alpha < \omega^2$ such that all terms in the sequence $t_\alpha \rightarrow^{\omega^2} t_{\omega^2}$ must be identical down to depth d, and all reduction steps in that sequence must be at depths greater than d.

Weak convergence (also called *Cauchy convergence*) is studied by Dershowitz et al. [1991]. However, without imposing further restrictions it does not have good properties. From a computational point of view it is intuitively

$$\begin{array}{ccccccc}
A & \longrightarrow & A & \longrightarrow & A & \longrightarrow \cdots & A \\
| & & | & & | & & | \\
B & & C & & C & & C \\
 & & | & & | & & | \\
 & & B & & C & & C \\
 & & & & | & & | \\
 & & & & B & & C \\
 & & & & & & | \\
 & & & & & & \vdots
\end{array}$$

Figure 1.4: Weakly but not strongly convergent reduction with the rule $A(x) \to A(C(x))$

$$\begin{array}{ccccccc}
A & \longrightarrow & B & \longrightarrow & B & \longrightarrow \cdots & B \\
 & & | & & | & & | \\
 & & A & & B & & B \\
 & & & & | & & | \\
 & & & & A & & B \\
 & & & & & & | \\
 & & & & & & B \\
 & & & & & & | \\
 & & & & & & \vdots
\end{array}$$

Figure 1.5: Strongly convergent reduction with the rule $A \to B(A)$

lacking, since a reduction sequence may weakly converge to a limit even although every step in the sequence is performed at the root, and the term can be thought of as still changing, even although in the limit it is being reduced to itself. More precisely, weakly convergent reduction in orthogonal systems does not have the Compression Property (Theorem 1.7.1). The construction of tiling diagrams and projection (Definition 1.6.1) also cannot be performed, because when one passes to a weakly convergent limit, information about the subterm structure is lost. This is illustrated by the following two examples.

1.3.3. Example. Consider the rules $A(x, y) \to A(y, x)$ and $C \to D$. The term $A(C, C)$ reduces to itself in one step, and therefore also by a weakly convergent reduction of length ω. However, it is impossible to say which of the two occurrences of C in the limit term descend from which of the two occurrences in the initial term. This can be seen by reducing the left-hand occurrence to D in the initial term, and performing the same sequence of reductions. This gives a divergent sequence $A(D, C) \to A(C, D) \to A(D, C) \to A(C, D) \to \ldots$.

1.3.4. Example. Consider the rule $I(x) \to x$ and the term $I(I(I(\ldots)))$, which we will abbreviate to I^ω. Suppose we try to perform a complete development of each of the infinitely many redexes in this term, outermost first. We obtain the sequence $I^1(I^2(I^3(\ldots))) \to I^2(I^3(I^4(\ldots))) \to I^3(I^4(I^5(\ldots))) \to \ldots$, where we have attached labels to show how each subterm is derived from the subterms of the previous term. The sequence weakly converges to $I(I(I(\ldots)))$. In the process, every redex of the initial term gets reduced, yet in the limit, we appear to still have all of them left. Where did they come from?

For this reason we do not further consider weak continuity or convergence. When we talk of reduction sequences, they will be assumed to be strongly continuous, and if they have limits, they will be assumed to strongly converge to them. The condition on depths of reduction which strong convergence adds was first stated (for sequences of length up to ω^2) by Farmer and Watro [1990].

1.3.5. Exercise.

1. Show that all finite reductions are strongly converging.
2. Give an example of a non-converging reduction.
3. Show that all reductions in the binary tree iTRS given by the rule $C \to B(C, C)$ are strongly convergent and that they can be of any countable length.
4. Show that all reductions in the iTRS with the single rule $A \to A$ are weakly convergent and that for each ordinal there is a weakly converging reduction of that length. Prove that no infinite reductions are strongly convergent.
5. Describe the normal form by strongly convergent transfinite reduction of $F(A, B)$, given the rule $F(x, y) \to G(x, F(y, H(x, y)))$.
6. Consider the iTRS $J(x) \to J(x)$: Let $J^\omega \to^\omega_{SC} J^\omega$ abbreviate the strongly convergent reduction $J^\omega \to J^\omega \to \ldots$ in which the redex contracted at the n-th step is at depth n. Why is the long sequence $J^\omega \to^\omega_{SC} J^\omega \to^\omega_{SC} J^\omega \to^\omega_{SC} J^\omega \ldots$ of length ω^2 not strongly converging?

1.3.6. Exercise. Show that in any iTRS, a reduction sequence is strongly convergent if and only if for every natural number n, the number of steps of the sequence which reduce a redex at depth less than n is finite.

Use the previous exercise to prove:

1.3.7. Proposition. *Every strongly converging reduction has countable length.*

The major difference between finitary and transfinite reduction is that for orthogonal systems, the transfinite Church-Rosser property fails. That is, $s \twoheadleftarrow\!\!\twoheadleftarrow\!\!\twoheadleftarrow \cdot \twoheadrightarrow\!\!\twoheadrightarrow\!\!\twoheadrightarrow t$ does not imply $s \twoheadrightarrow\!\!\twoheadrightarrow\!\!\twoheadrightarrow \cdot \twoheadleftarrow\!\!\twoheadleftarrow\!\!\twoheadleftarrow t$. As this fact has inspired many of our results about transfinite rewriting, we exhibit here the canonical counterexample.

Take the rules $A(x) \to x$ and $B(x) \to x$, and the term $A(B(A(B(\ldots))))$. This can be transfinitely reduced to both $A(A(A(\ldots)))$ and $B(B(B(\ldots)))$, each of which reduces only to itself.

The failure of the Church-Rosser property is not confined to specially constructed pathological cases, but occurs also for such standard systems as combinatory logic, or the rules for describing lists with head and tail operators.

1.3.8. EXERCISE. 1. Find a counterexample to the Church-Rosser property for the rules $Head(Cons(x,y)) \to x$ and $Tail(Cons(x,y)) \to y$.

2. Combinatory logic has the rules $Sxyz \to xz(yz)$ and $Kxy \to x$. Find a finite closed term t and two infinite reductions $t \twoheadrightarrow\!\!\!\!\!\rightarrow t_0$ and $t \twoheadrightarrow\!\!\!\!\!\rightarrow t_1$ such that t_0 and t_1 have no common reduct.

3. With just the rule $Kxy \to x$, find a closed term t and two infinite reductions $t \twoheadrightarrow\!\!\!\!\!\rightarrow t_0$ and $t \twoheadrightarrow\!\!\!\!\!\rightarrow t_1$ such that t_0 and t_1 have no common reduct.

In Theorem 1.8.2 we will see that examples like these are the only way in which the Church-Rosser property fails.

1.4. Infinitary lambda calculus

The theory of infinite rewriting can be also carried out for lambda calculus, and as most of the results and proofs are common to both, we shall now describe infinitary lambda calculus.

Infinite lambda expressions can be defined in the same way as infinite terms are defined. We take the function symbols to be the binary application operator, the unary symbol λx and the nullary symbol x for each variable x. However, the question of α-conversion raises a slight complication. It is customary to define finite lambda expressions up to α-equivalence. In the infinitary case, we must deal with renamings of infinite sets of variables, and substitutions applied to infinitely large terms. The desired result is that we can write expressions using explicit variables, but assume that α-conversion is implicitly performed to avoid variable capture. To achieve this formally requires a certain amount of care.

The most direct way to proceed is to define infinite lambda terms exactly as for ordinary terms, without regard to α-conversion. Then define α-equivalence for finite and infinite terms, and finally define the distance between two α-equivalence classes as the infimum of the distance between pairs of respective members.

Alternatively, we can define the distance metric on α-equivalence classes of finite terms, and then take the metric completion. This gives us a space in which the finite elements are α-equivalence classes of finite terms, and in which the infinite elements can be identified with the α-equivalence classes of infinite terms as previously defined.

To carry out the first method, we must define α-equivalence for infinite terms. We will assume that a textual replacement of a variable x by a variable x' without regard to binding is unproblematic, even for infinite terms. We denote this operation by $[x \to x']$.

1.4.1. DEFINITION. A *conflict* between (finite or infinite) lambda terms s and t is a common occurrence u such that one of the following holds:

1. $u = \langle\,\rangle$ and s and t are not identical variables, not both applications, and not both abstractions.
2. $u = n \cdot v$, $s = s_1 s_2$, $t = t_1 t_2$, and v is a conflict of s_n and t_n ($n = 1$ or 2).
3. $u = 1 \cdot v$, $s = \lambda x.s'$, $t = \lambda x'.t'$, 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'.

s and t are called *α-equivalent* if there is no conflict between them.

This definition also gives us a convenient way of defining the metric.

1.4.2. DEFINITION. If s and t have no conflict, then $d(s,t) =_{df} 0$. Otherwise, $d(s,t) =_{df} 2^{-n}$, where n is the minimum depth of any conflict between them.

1.4.3. EXERCISE.

1. Prove that for finite terms, the above definition of α-equivalence coincides with the classical definition.
2. Prove that d gives a metric on α-equivalence classes of terms, by establishing the following:
   (a) If s and s' are α-equivalent, then $d(s,t) = d(s',t)$ for every term t.
   (b) $d(s,t) = 0$ if and only if s and t are α-equivalent.

The second method of defining the space of infinite terms begins by ignoring issues of α-equivalence and considering the metric on finite terms given by considering them as terms over a signature in which λx is a unary operator for every variable x. Call this metric d'. Then define a metric on α-equivalence classes of finite terms by defining $d'(S,T) = \min\{\, d'(s,t) \mid s \in S \wedge t \in T \,\}$ for α-equivalence classes S and T. Finally, take the metric completion of the space.

1.4.4. EXERCISE. Prove that this gives a space isomorphic to the space of α-equivalence classes of finite and infinite terms with metric d.

In the finitary theory, it is customary to introduce the variable convention at this point, to avoid dealing explicitly with α-conversion. This convention stipulates that the free variables of a terms are distinct from its bound variables, and that when a substitution of a term for a variable in another term is made, bound variables are implicitly renamed to avoid variable capture. In the infinitary case it is slightly more complicated, since if there are only countably many variables, a term might include all of them free, leaving no room to pick new variables. We solve this problem by the 'Hilbert hotel' trick. If all of the variables are in use at some point in an argument, and we still need more, then an implicit renaming of every variable x_i by x_{2i} is performed, which frees an infinite number of variables for reuse.

In section 1.2 we remarked on the way that inductive constructions on infinitary terms must proceed coinductively from the root downwards rather than inductively from the leaves upwards. As an example, here is a formal definition of substitution and a proof of the substitution lemma.

1.4.5. Definition. The result of substituting N for x in M, notated $M[x := N]$, is defined to be the term whose positions and symbols are as follows.

The positions of $M[x := N]$ are the positions of M other than the free occurrences of x, and the positions $u \cdot v$ where $M|_u = x$ and v is a position of N.

In the first case, the root symbol of $M[x := N]|_u$ is that of $M|_u$. In the second case, the root symbol of $M[x := N]|_{u \cdot v}$ is that of $N|_v$.

1.4.6. Lemma (Substitution Lemma, cf. Barendregt [1984], 2.1.16.). *If $x \neq y$ and x is not free in L, then*

$$M[x := N][y := L] \quad \equiv \quad M[y := N][x := N[y := L]]$$

Proof. Let u be any common position of two terms of the above form. We will show by induction on the length of u that u is not a conflict between those terms.

We proceed by cases of the form of M. If M is x then both sides are $N[y := L]$. If M is y, then both sides are L. If M is any other variable, both sides are M. In all these cases, both sides are identical and therefore have no conflict.

If M is $\lambda z.M_1$, we may assume by the variable convention that z is not x or y and is not free in N. Then the two terms are $\lambda z.M_1[x := N][y := L]$ and $\lambda z.M_1[y := N][x := N[y := L]]$. $\langle\,\rangle$ is not a conflict of these terms, and $i \cdot v$ is a conflict if and only if $i = 1$ and v is a conflict of $M_1[x := N][y := L]$ and $M_1[y := N][x := N[y := L]]$. By induction, it cannot be.

If $M = M_1 M_2$, then the two terms are $(M_1[x := N][y := L])(M_2[x := N][y := L])$ and $(M_1[y := N][x := N[y := L]])(M_2[y := N][x := N[y := L]])$. $\langle\,\rangle$ is not a conflict of these terms, and $i \cdot v$ is a conflict if and only if $i = 1$ and

v is a conflict of $M_1[x := N][y := L]$ and $M_1[y := N][x := N[y := L]]$, or $i = 2$ and v is a conflict of $M_2[x := N][y := L]$ and $M_2[y := N][x := N[y := L]]$. By induction, neither alternative is possible. ⊠

β- and η-reduction are defined for infinitary lambda calculus in terms of substitution in the same way as for the finitary calculus. However, η-reduction has a property that is unsatisfactory from a computational point of view. To determine whether a term $\lambda x.(Mx)$ is an η-redex, one may have to inspect the whole of M to determine whether x occurs free. If M is infinite, this is an infinite amount of work that must be completed before the redex can be reduced. In effect, the pattern of the η-rule is infinite. A β-redex, in contrast, can be recognised by seeing just two nodes of the parse tree: an application with an abstract on its left. Although a β-redex may require an infinite amount of work to reduce, one may devise computational schemes (such as explicit substitution Abadi et al. [1991]) which allow any finite part of the reduct to be computed in finite time.

From now on, we will deal with iTRSs and infinitary lambda calculus together. The primary differences between them will be seen to arise from the more complicated way in which in the lambda calculus, residuals of a single redex can become nested inside each other, whereas they always remain disjoint in term rewriting.

By 'orthogonal systems' we mean orthogonal (i)TRSs and the lambda calculus; by 'left-linear systems' we mean left-linear (i)TRSs and the lambda calculus.

1.4.7. Exercise. Based on the exercises in the previous section, or otherwise, find a counterexample to the Church-Rosser property for the infinitary lambda calculus with β-reduction.

1.5. Residuals and developments

The concepts of descendant, residual, and development extend in a straightforward manner to the transfinite context. For a single-step reduction with infinite terms, there is nothing new, but to deal with infinite sequences we must define how to take the limit of a sequence of sets of residuals.

1.5.1. Definition. Let $\sigma : t_0 \to^\alpha t_\alpha$ be a strongly convergent sequence. Let U be a set of positions of t_0. The set of *residuals* or *descendants* of U by σ is a set of positions of t_α denoted by U/σ, and defined by induction as follows.

- If $\alpha = 0$, then σ is the empty sequence and U/σ is just U.
- If $\alpha = 1$, let the redex R reduced in $t_0 \to t_1$ be at position v. Take any $u \in U$. If v is not a proper prefix of u, then $u/R = \{u\}$.

(For term rewriting.) If u is one of the positions pattern-matched by R, that is, a function symbol in the left-hand side of R is matched to the node of the term at u, then $u/R = \varnothing$. Otherwise, if $u = v \cdot w \cdot x$, where w is the position of a variable in the left-hand side of R, then u/R consists of all positions of the form $v \cdot w' \cdot x$, where w' is a position of the same variable in the right-hand side of R. U/R is the union of all u/R for $u \in U$.

(For lambda calculus.) If $u = v \cdot 1$, then $u/R = \varnothing$. If $u = v \cdot 1 \cdot 1 \cdot w$, then $u/R = \{v \cdot w\}$. If $u = v \cdot 2 \cdot w$, then u/R is the set of all $u \cdot w' \cdot w$, where $u \cdot w'$ is the position of any occurrence of the bound variable of the abstraction of the redex.

- If $\alpha = \beta + 1$, then $U/\sigma = (U/(t_0 \rightarrow_\beta t_\beta))/(t_\beta \rightarrow t_{\beta+1})$.
- If α is a limit ordinal, then $u \in U/\sigma$ iff $u \in U/(t_0 \rightarrow_\beta t_\beta)$ for all large enough $\beta < \alpha$.

1.5.2. Exercise. Prove that in the last case of the above definition, the condition 'for all large enough $\beta < \alpha$' can equivalently be replaced by 'for arbitrarily large $\beta < \alpha$'.

1.5.3. Definition. In an orthogonal system, let $\sigma : t_0 \rightarrow^\alpha t_\alpha$ be a strongly convergent sequence. Let R be a set of redexes of t_0, whose positions are the set U. R/σ denotes the set of redexes of t_α at positions in U/σ.

In the above definition, orthogonality ensures that there is a redex at each position in U/σ, and that each redex is a redex of the same rule as its ancestor.

1.5.4. Lemma. *In a left-linear system, let σ be a reduction whose length is a limit ordinal. Every redex in the final term of σ is a residual of a redex in some earlier term.*

Proof. Let there be a redex at position u in the final term. Let d be the depth of the left-hand side of a rule which matches at u. By strong convergence, every step of σ after some point must be at depth at least $|u| + d$. In every term from that point on, there must be a redex of the same rule at u. The redex in the final term is a residual of all of them. ⊠

When the length of σ is a limit ordinal α, and there is a redex at a position u of the final term, then for all sufficiently large $\beta < \alpha$, u is the position of a redex of the same rule in t_β, and the redex at u in t_α is its unique residual. This fact depends on the finiteness of left-hand sides.

The following proposition is immediate from the definition of residuals.

1.5.5. Proposition. *In an iTRS, if $\mathcal{R}$ is a set of pairwise disjoint redexes, then so is $\mathcal{R}/\sigma$ for any strongly converging reduction sequence σ. For the lambda calculus this is true when σ is a single step and $\mathcal{R}$ is a single redex.*

1.5.6. Definition. A *development* of a set of redexes $\mathcal{R}$ of a term t_0 is a strongly converging reduction $t_0 \to^\alpha t_\alpha$ such that for any $\beta \le \alpha$ the step $t_\beta \to t_{\beta+1}$ reduces a residual of a member of $\mathcal{R}$ by $t_0 \to_\beta t_\beta$. If t_α contains no residual of $\mathcal{R}$, then the development is said to be *complete*.

Our examples of the failure of the Church-Rosser property also demonstrate that the complete developments theorem fails: not every set of redexes has a complete development. For example, given the rule $I(x) \to x$ and the term $I(I(I(\ldots)))$, the set of redexes in this term has no complete development, since any reduction which attempts to reduce all of the redexes will not be strongly convergent. (Cf. Example 1.3.4.) However, we can prove a restricted version of the theorem.

Let S be a set of redexes in a term t or an orthogonal iTRS or lambda calculus. We will begin by exhibiting a construction of which we will later prove that, when it exists, it constructs the final term of every complete development of S, and when it does not, S has no complete development.

We first informally describe the basic idea, in the setting of iTRSs. We use the notion of a *path* in a term, which in the lambda calculus has been used in Asperti et al. [1994] and Asperti and Laneve [1995] for characterising Lévy's redex families. A path starts from the root of t and proceeds from each node to an immediate descendant, until we encounter the root of a redex in S, at some position u. All the nodes traversed before this point will be nodes of the result $\mathcal{F}(t, S)$ of a complete development of S. u will also be a position of $\mathcal{F}(t, S)$, but to determine the symbol there we must jump to the root of the term r, the right-hand side of the rule $l \to r$ that applies at u. We continue by traversing r, with each non-variable node we reach corresponding to a node of $\mathcal{F}(t, S)$. When we reach an occurrence in r of a variable x, we must jump back to t at the position corresponding to the occurrence of x in l.

For lambda expressions, the method is similar. When we reach the root of a β-redex in S, we jump to the root of its body; when we encounter within the body an occurrence of the bound variable, we must jump to the root of the right-hand side of the redex.

This construction will fail if we find ourselves making an infinite sequence of jumps without passing along an edge of t or of a right-hand side. It is precisely in such cases that S will fail to have a complete development.

We now give a formal description and a proof that it yields a term with the claimed properties.

1.5.7. Definition. For a given term t in an iTRS and a set of redexes S of t, define a *path* as a certain kind of finite sequence of nodes and edges. Each node is labelled by a term and a position of that term, and each edge is either unlabelled or labelled by an integer. The term labelling a node is either t or the right-hand side of some rule.

Every path of t begins with a node labelled by $(t, \langle\rangle)$. Suppose the final node of a path is labelled with (s, u). The path may be extended in any of the following ways.

1. Either s is t, and u is not the position of a redex in S, or s is a right-hand side t', and u is any position of t'. If $u \cdot i$ is a position of t or t' respectively, then s can be extended with an edge labelled i and a node labelled $(t, u \cdot i)$ or $(t', u \cdot i)$ respectively.

2. s is t, and u is the position of a redex in S, the right-hand side of whose rule is t'. s can be extended by an unlabelled edge and a node labelled with $(t', \langle\rangle)$.

3. s is a right-hand side t', and $t'|_u$ is a variable x. Let the last occurrence of t in a node label of s be with position v. v will be the position of a redex in S having a rule $t'' \to t'$. s may be extended by an unlabelled edge and a node labelled with $(t, v \cdot w)$, where w is the position of the (unique) occurrence of x in t''.

For a lambda expression t, the construction is similar, the cases being:

1. s is t, and u is not the position of a redex in S. If $u \cdot i$ is a position of t, then s can be extended with an edge labelled i and a node labelled $(t, u \cdot i)$.

2. s is t, and u is the position of a β-redex in S. s can be extended by an unlabelled edge and a node labelled by $(t, u \cdot 1 \cdot 1)$.

3. s is t, and u is the position of a variable of t bound by a redex in S at position v. s can be extended by an unlabelled edge and a node labelled by $(t, v \cdot 2)$.

Let $\mathcal{P}(t, S)$ be the set of such paths. Define $U(t, S)$ to be the set of sequences of edge labels of the paths. For any $u \in U(t, S)$, let s be a path of maximal length for which u is its sequence of edge labels. s is uniquely determined. It may be finite or infinite. If s is always finite, we will say that S has *finite jumps*. When s is finite, it has a final node. Let its label be (t', v). Define $F(u)$ to be the symbol of t' at v. Thus the domain of F is all of $U(t, S)$ if and only if S has finite jumps.

1.5.8. Proposition. *If the domain of F is all of $U(t, S)$, then $U(t, S)$ and F together define a term, which we will denote by $\mathcal{F}(t, S)$.*

Proof. $U(t, S)$ is prefix-closed by definition. $u \cdot i$ is in $U(t, S)$ if and only if i is no more than the arity of $F(u)$. Therefore the set of positions $U(t, S)$ and the labelling F define a term. ⊠

1.5.9. PROPOSITION (Finite Jumps Developments Theorem). *In an orthogonal system, let S be a set of redexes in a term t having finite jumps.*

1. *Every complete development of S has $\mathcal{F}(t, S)$ as its final term.*
2. *For any position u of t, the set of descendants of u by a complete development of S is independent of the choice of development.*
3. *The same is true for the residuals of any redex of t.*
4. *S has a complete development.*

PROOF. 1. Let t and S be as described. Suppose that we are given a complete development of S. We will show by induction that for every term t_α in the sequence, containing residuals S_α of S, $\mathcal{P}(t_\alpha, S_\alpha)$ is obtainable from $\mathcal{P}(t, S)$ by deleting some unlabelled edges (or perhaps none of them), and hence that S_α has the finite jumps property, and $\mathcal{F}(t_\alpha, S_\alpha) = \mathcal{F}(t, S)$.

For $t_0 = t$ this is immediate.

For the successor case, let $t_{\alpha+1}$ result from reducing a redex s in t_α. By tracing through the $\mathcal{P}$ construction, one can (tediously) verify that $\mathcal{P}(t_{\alpha+1}, S_{\alpha+1})$ is obtainable from $\mathcal{P}(t_\alpha, S_\alpha)$ by deleting some unlabelled edges. Therefore $S_{\alpha+1}$ has the finite jumps property, and $\mathcal{F}(t_{\alpha+1}, S_{\alpha+1}) = \mathcal{F}(t_\alpha, S_\alpha) = \mathcal{F}(t, S)$.

For the limit case, assume that the claim holds up to a limit ordinal λ. $\mathcal{P}(t_\lambda, S_\lambda)$ results from $\mathcal{P}(t, S)$ by deleting all the unlabelled edges that were deleted in any of the previous $\mathcal{P}(t_\alpha, S_\alpha)$. Therefore S_λ has the finite jumps property, and $\mathcal{F}(t_\lambda, S_\lambda) = \mathcal{F}(t, S)$.

We conclude that this holds for the final term t'. But then $\mathcal{F}(t, S) = \mathcal{F}(t', \varnothing) = t'$, and t' is thus independent of the development.

2,3. We leave it to the reader to carry through a similar argument to show how the descendant and residual relationships between t and t' can also be determined from $\mathcal{P}(t, S)$, and are therefore also independent of the development.

4. Let $p \in \mathcal{P}(t, S)$. Let $r \in S$ be a redex encountered in p. Let $d(p, r)$ be the number of labelled edges from the beginning of p up to r. This is the depth at which some residual of r can be reduced in an outermost development of S (as is easily proved by induction along the sequence of redexes preceding r in p). Conversely, every residual of r reduced in an outermost development is reduced at a depth obtained in this way. Since $\mathcal{P}(t, S)$ is finitely branching, there are only finitely many occurrences (r, p) of redexes r in $\mathcal{P}(t, S)$ having a given value of $d(r, p)$. Therefore in

an infinite outermost development, the depth of reductions must tend to infinity. That is, the sequence is strongly convergent. Outermost reduction can therefore be continued until a complete development is obtained.

⊠

A concrete example of a class of sets of redexes having the finite jumps property are the *finitely nesting* sets of redexes, i.e. those sets of redexes whose positions contain no infinite chain $u_1 < u_2 < u_3 < \ldots$.

1.5.10. Proposition (Finitely Nesting Developments Theorem). *Let S be a set of finitely nesting redexes in a term t.*

1. *Every development of S is strongly convergent.*
2. *If S and T are sets of finitely nesting redexes in a term t, then T/S is finitely nesting, where T/S is the set of residuals of $T \cup S$ after a complete development of S.*

Proof. 1. A non-strongly convergent reduction must contain a subsequence of reductions all at the same position. In an iTRS, this is impossible for the residuals of a set of finitely nesting redexes, since a reduction in an iTRS can never change the nesting relationship between two redexes. For lambda-calculus, this argument is insufficient, since reduction of one redex can cause residuals of non-nested redexes to be nested. To deal with this, consider the following relation on S: $r_1 < r_2$ if either r_1 contains r_2, or there is a redex $r_3 \in S$ of the form $\lambda x.t_1t_2$ such that r_1 is contained in t_1, r_1 contains x free, and r_2 is contained in t_2. We extend this to its transitive closure. Since $r_1 < r_2$ implies that their positions u_1 and u_2 satisfy $u_1 < u_2$ in the lexicographic partial order, this relation on S is (the irreflexive part of) a partial order. If a residual of r_1 can contain a residual of r_2 in some development of S, then $r_1 < r_2$. Therefore if this relation has no infinite chains, S cannot have a non-strongly convergent development. That it has no infinite chains follows from Lemma 1.5.11.

2. T/S is well-defined because of Proposition 1.5.9. Thus it is sufficient to demonstrate that the set of residuals remaining after any development of a finitely nesting set S is finitely nesting. This follows from the same partial ordering as in the first part: that partial ordering includes all the nesting relationships that can hold among residuals of S.

⊠

1.5.11. Lemma. *Let T be a finitely branching tree. Let N be a subset of the nodes of T. The prefix ordering restricted to N is a partial ordering which*

$$\begin{tikzcd} t_{0,0} \arrow[r, twoheadrightarrow, "H"] \arrow[d, twoheadrightarrow, "V"'] & t_{0,1} \arrow[d, twoheadrightarrow, "V/H"] \\ t_{1,0} \arrow[r, twoheadrightarrow, "H/V"'] & t_{1,1} \end{tikzcd}$$

Figure 1.6: Elementary tile

we will denote $n_0 < n_1$. Define a second relation on N by $n_0 <' n_1$ if n_0 is before n_1 in the lexicographic ordering of positions, and some common prefix of n_0 and n_1 is in N. Then if $<$ has no infinite chains, neither does $< \cup <'$.

PROOF. By König's lemma, the absence of infinite chains in N implies that N under the prefix ordering consists of a finite or infinite set of finite trees. The definition of $<'$ implies that when $n_0 <' n_1$, they belong to the same such tree. The same is trivially true of $<$. Therefore every chain of $< \cup <'$ is contained in one of those trees, and is finite. ⊠

The above theorems may be generalisable to various types of higher-order rewriting. Lemma 1.5.11 was required to establish them for lambda calculus, due to the possibility of nesting of residuals. Higher-order systems can in general create more complicated forms of nesting, and would likely require stronger forms of the lemma. The remainder of the proofs of Propositions 1.5.9 and 1.5.10 are largely independent of the details of rewriting.

The importance of the Finite Jumps and Finitely Nesting Developments Theorems is that they allow us to construct *elementary tiles*: diagrams of the form of Figure 1.6. H and V are complete developments of finitely nesting sets $\mathcal{H}$ and $\mathcal{V}$. H/V is a complete development of the finitely nesting set $\mathcal{H}/V$, and V/H is defined similarly. The theorems tell us that H/V and V/H have the same final term, since $H \cdot (V/H)$ and $V \cdot (H/V)$ are both complete developments of the finitely nesting set $\mathcal{H} \cup \mathcal{V}$.

1.5.12. LEMMA. *Let U be a set of positions of a term t and $S : t \twoheadrightarrow t'$ be a reduction sequence. Suppose that every step of S has depth greater than d. Then U and U/S have exactly the same members of depth at most d.*

PROOF. This is immediate from the fact that a reduction cannot affect any part of the term at lesser depth than itself. ⊠

1.6. Tiling diagrams

Extending to transfinite rewriting the notion of projection of one sequence over another requires some extra work to deal with limit ordinals. Unlike the finitary case, projections need not exist, even in orthogonal systems, for the same reason that the Church-Rosser property does not always hold.

In this section we reconstruct the theory of projections and tiling diagrams for transfinite rewriting. In the next section we will use these results to establish approximate Church-Rosser properties, and exact properties for a restricted class of orthogonal iTRSs.

1.6.1. DEFINITION. A *tiling diagram* for two strongly converging reductions $H : t_{0,0} \to^{\alpha} t_{\alpha,0}$ and $V : t_{0,0} \to^{\beta} t_{0,\beta}$ consists of a rectangular arrangement of strongly convergent reduction sequences as in Figure 1.7, subject to certain conditions.

- Each component reduction $H_{\gamma,\delta} : t_{\gamma,\delta} \twoheadrightarrow t_{\gamma,\delta+1}$ is a complete development of a set of redexes $\mathcal{H}_{\gamma,\delta}$ of $t_{\gamma,\delta}$. Each $V_{\gamma,\delta}$ is similarly related to a set $\mathcal{V}_{\gamma,\delta}$.
- $\mathcal{H}_{0,\delta}$ is the redex reduced in step δ of H, and $\mathcal{V}_{\gamma,0}$ is the redex reduced in step γ of V.
- Let $H_{\gamma,[\delta,\delta']} : t_{\gamma,\delta} \twoheadrightarrow t_{\gamma,\delta'}$ be the (strongly convergent) concatenation of all $H_{\gamma,\delta''}$ for $\delta \leq \delta'' < \delta'$. Define $V_{[\gamma,\gamma'],\delta}$ similarly.

  Then $\mathcal{H}_{\gamma,\delta} = \mathcal{H}_{0,\delta}/V_{[0,\gamma],\delta}$ and $\mathcal{V}_{\gamma,\delta} = \mathcal{V}_{\gamma,0}/H_{\gamma,[0,\delta]}$.

A partial tiling diagram for H and V is defined similarly, except that the diagram is only required to include H and V, and if it includes $t_{\gamma,\delta}$, it also includes $t_{\gamma',\delta'}$, $H_{\gamma,\delta'}$, and $V_{\gamma',\delta}$ for all $\gamma' < \gamma$ and $\delta' < \delta$. The conditions must hold whenever all the objects they refer to exist.

If H and V have a tiling diagram, then the strongly convergent sequences along the bottom and right edges will be denoted by H/V and V/H respectively.

Not all pairs of sequences have a tiling diagram, as illustrated by the counterexample to the Church-Rosser property. The next few theorems establish some sufficient conditions for a tiling diagram to exist.

1.6.2. LEMMA. *Let U be a set of positions of redexes in a term t. For each $u \in U$, define l_u to be the maximum depth of any variable occurrence in the left-hand side of the rule for the redex at u in s_α. Then as $|u| + l_u$ tends to infinity, so does $|u|$.*

PROOF. In an orthogonal iTRS and in infinitary lambda calculus, distinct redexes of the same term must be at different positions. There are only finitely many positions u of a given depth, and so only finitely many values of $|u| + l_u$ for any given $|u|$. ⊠

$$
\begin{tikzcd}
 & & & H & & & \\
V & t_{0,0} \arrow[r] \arrow[d] & t_{0,1} \arrow[r, dash, dashed] \arrow[d, ->>] & t_{0,\delta} \arrow[r] \arrow[d, ->>] & t_{0,\delta+1} \arrow[r, dash, dashed] \arrow[d, ->>] & t_{0,\beta} \arrow[d, ->>] & V/H \\
 & t_{1,0} \arrow[r, ->>] \arrow[d, dash, dashed] & t_{1,1} \arrow[r, dash, dashed] \arrow[d, dash, dashed] & t_{1,\delta} \arrow[r, ->>] \arrow[d, dash, dashed] & t_{1,\delta+1} \arrow[r, dash, dashed] \arrow[d, dash, dashed] & t_{1,\beta} \arrow[d, dash, dashed] & \\
 & t_{\gamma,0} \arrow[r, ->>] \arrow[d] & t_{\gamma,1} \arrow[r, dash, dashed] \arrow[d, ->>] & t_{\gamma,\delta} \arrow[r, ->>, "H_{\gamma,\delta}"] \arrow[d, ->>, "V_{\gamma,\delta}"'] & t_{\gamma,\delta+1} \arrow[r, dash, dashed] \arrow[d, ->>] & t_{\gamma,\beta} \arrow[d, ->>] & \\
 & t_{\gamma+1,0} \arrow[r, ->>] \arrow[d, dash, dashed] & t_{\gamma+1,1} \arrow[r, dash, dashed] \arrow[d, dash, dashed] & t_{\gamma+1,\delta} \arrow[r, ->>] \arrow[d, dash, dashed] & t_{\gamma+1,\delta+1} \arrow[r, dash, dashed] \arrow[d, dash, dashed] & t_{\gamma+1,\beta} \arrow[d, dash, dashed] & \\
 & t_{\alpha,0} \arrow[r, ->>] & t_{\alpha,1} \arrow[r, dash, dashed] & t_{\alpha,\delta} \arrow[r, ->>] & t_{\alpha,\delta+1} \arrow[r, dash, dashed] & t_{\alpha,\beta} & \\
 & & & H/V & & &
\end{tikzcd}
$$

Figure 1.7: Tiling diagram

Note that since arities and left-hand sides are unbounded, the condition that U is a set of positions of a single term is essential, although t plays no other part in the lemma. For lambda calculus, the lemma is trivial, since l_u is always 2.

The next lemma is one of the results which holds only for orthogonal iTRSs and fails for lambda calculus.

1.6.3. THEOREM (The Strip Lemma). *1. In an orthogonal iTRS, let H and V be strongly convergent reductions with the same initial term, with V being one step long. Then H and V have a tiling diagram.*

2. In lambda calculus, (1) will hold given the additional assumption that H/V is strongly convergent.

PROOF. 1. If H is empty this is trivial, and if the theorem holds for H of length β, then by the Finitely Nesting Developments Theorem it also holds for H of length $\beta + 1$.

Let H have limit ordinal length β. Let V be one step, reducing a redex at position u. Choose a depth d. By Lemma 1.6.2, there is a d' such that for $u \in \mathcal{V}_\beta$, $|u| + l_u > d'$, then $|u| > d$. Let l be the maximum depth of any left-hand side for the redexes of t_α at positions $u \in \mathcal{V}_\beta$ having length not more than d'. Since there are finitely many such redexes, l is finite. Choose γ such that every step of $H_{0,[\gamma,\beta]}$ has depth at least $d'+l$. Then for any $\epsilon \geq \gamma$, the positions in $\mathcal{V}_\epsilon$ of depth at most d' are the same as those of $\mathcal{V}_\beta$, and the redexes there are redexes by the same rules. These redexes are residuals of a single redex, therefore they are pairwise disjoint (something which need not be true in lambda calculus). Therefore every member of $\mathcal{H}_\epsilon$ is contained in at most one member of $\mathcal{V}_\epsilon$ and so every member of $\mathcal{H}_\epsilon/\mathcal{V}_\epsilon$ must have depth at least d. It follows that every step of $H_{1,[\gamma,\beta)}$ has depth at least d. Therefore H/V is strongly convergent, and converges to t_β.

2. Suppose that V/H has no strongly convergent complete development. Then for some position u, there is a development of V/H performing infinitely many steps at u. Let u have depth d. Since H and H/V are strongly convergent, there is a α such that every step of these sequences from $t_{0,\alpha}$ or $t_{1,\alpha}$ respectively has depth at least d.

   The sets $\mathcal{V}_\beta$ for $\beta \geq \alpha$ must have complete developments which perform larger and larger numbers of reductions at u as β increases. This implies that for arbitrarily large values of β, H_β/V_β must include a redex at u/V_β. Because H is strongly convergent, u/V_β must be eventually constant as β increases. But this contradicts the strong convergence of H/V.

$\boxtimes$

Here is a counterexample showing that the iTRS part of the Strip Lemma fails for the lambda calculus. Let t be a finite term with the property that $tx \to^* x(tx)$. Such a t is easily constructed by means of a fixed point operator. Consider the term tr, where $r = \lambda x.(\lambda y.x)w$. This reduces to $r(tr)$, and then to $\lambda y.(tr)w$, and so in infinitely many steps to the infinite term s such that $s \equiv (\lambda y.s)w$. The term tr also reduces in one step to $t(\lambda x.x)$. These cannot have any common reduct, because w is free in every reduct of s, but every reduct of $t(\lambda x.x)$ is closed. Notice that the infinite chain of redexes in s are all residuals of the single redex $(\lambda y.x)w$. If we try to construct the Strip Lemma tiling diagram, we find that H/V consists of an infinite repetition of the sequence $t(\lambda x.x) \to^* (\lambda x.x)(t(\lambda x.x)) \to t(\lambda x.x)$, which is not strongly convergent.

1.6.4. THEOREM. *Let H and V be strongly convergent reductions with the same initial term. Using the notation of Definition 1.6.1, suppose that a tiling diagram exists for $H_{0,[0,\delta]}$ and $V_{[0,\gamma],0}$ with $\gamma < \alpha$. In an orthogonal iTRS, a tiling diagram exists for $H_{0,[0,\delta]}$ and $V_{[0,\gamma+1],0}$. In lambda calculus, it will exist provided that $H_{0,[0,\delta]}/V_{[\gamma,\gamma+1],0}$ is strongly convergent.*

PROOF. Apply the Strip Lemma to the reductions $H_{\gamma,[0,\delta]}$ and $V_{[\gamma,\gamma+1],0}$. $\boxtimes$

1.6.5. THEOREM. *In an orthogonal iTRS or lambda calculus, let H and V be strongly convergent reductions with the same initial term, whose lengths are limit ordinals β and α respectively. Suppose that all of the data for a tiling diagram for H and V exists, except that $t_{\alpha,\beta}$ is not known to exist, and the sequences H/V and V/H are not known to be strongly convergent, although every proper initial segment of them is. Then the following conditions are equivalent:*

1. *The tiling diagram can be completed, i.e. H/V and V/H are strongly convergent and have the same limit.*
2. *H/V is strongly convergent.*
3. *V/H is strongly convergent.*
4. *The sequence of rows is uniformly strongly convergent, by which is meant that for every depth d there is a $\delta < \beta$ such that for all $\gamma < \alpha$, every step of $H_{\gamma,[\delta,\beta]}$ has depth at least d.*
5. *The sequence of columns is uniformly strongly convergent.*

PROOF. Suppose that the rows are uniformly strongly convergent. Choose a depth d. There is a $\delta < \beta$ such that for all $\gamma < \alpha$, every redex in $\mathcal{H}_{\gamma,\delta}$ has depth at least d. Since $\mathcal{H}_{\alpha,\delta}$ is the limit of $\mathcal{H}_{\gamma,\delta}$ for $\gamma < \alpha$, it follows that every redex in $\mathcal{H}_{\alpha,\delta}$ has depth at least d, and therefore so does every step of $H_{\alpha,[\delta,\beta]}$. Therefore H/V is strongly convergent. This proves that (4) implies (2), and therefore by symmetry that (5) implies (3).

Suppose that V/H is strongly convergent. Choose any depth d. Then for some γ, every step of $V_{[\gamma,\alpha],\beta}$ has depth at least d. By strong convergence of each row of the diagram, we can find a δ such that every step of $H_{\gamma,[\delta,\beta]}$ has depth at least d. Therefore by Lemma 1.5.12, every horizontal and vertical step in the diagram below and to the right of (γ, δ) has depth at least d. This demonstrates horizontal and vertical uniform strong convergence, strong convergence of the right and bottom edges, and their convergence to the same limit. This proves that (3) implies (4), (5), and (1).

It is clear that (1) implies all the others. Combining these implications shows that all the items are equivalent. ⊠

1.6.6. EXERCISE. Consider the counterexamples given earlier to the Church-Rosser property. Show that they satisfy the hypothesis of Theorem 1.6.5, and that all of the enumerated properties are false.

1.6.7. THEOREM. *Let H and V be reduction sequences having the tiling diagram of Figure 1.7. Let u be a position of $t_{0,0}$. Then the set of descendants of u by any path in the diagram from $(0, 0)$ to (γ, δ) is independent of the path or of the choice of complete developments for the edges of the tiles. The same holds for the residuals of any redex r of $t_{0,0}$.*

PROOF. By Proposition 1.5.10 this is true for each path consisting of a single tile-edge. The theorem follows. ⊠

1.7. The Compression Lemma

To perform an infinite number of computational steps to obtain some final result may reasonably be thought of as an idealisation of the notion of computing indefinitely to obtain more and more of the final result. To perform an infinite number of steps, and then to perform more steps, looks less reasonable. Reduction sequences of length greater than ω might therefore be considered to have no computational significance. The Compression Property gives such sequences computational meaning. This property, which holds for left-linear iTRSs and for lambda calculus, is that for every sequence of length greater than ω, there is another sequence with the same endpoints of length at most ω.

1.7.1. THEOREM (The Compression Lemma). *In every left-linear iTRS, if* $s \to^{\alpha} t$ *then* $s \to^{\leq\omega} t$.

PROOF. This is proved by induction on α. It is trivial for $\alpha < \omega$.

Suppose $\alpha = \omega+1$. We have $s_0 \to^{\omega} s_{\omega} \to t$. The redex reduced in the final step must be the unique residual by $s_n \to^{\omega} s_{\omega}$ of a redex r in s_n at the same position u. Every step of $s_n \to^{\omega} s_{\omega}$ has depth greater than d plus the depth of the pattern of r. Thus there can be no conflicts between that sequence and the residuals of r, so its projection over r can be constructed exactly as in the orthogonal case. It can be considered as an interleaving of a number of copies of subsequences of that sequence, each of length at most ω, and acting on disjoint parts of s'_n. That interleaving might have length ω^2, but there clearly also exists an interleaving of length at most ω.

Suppose $\alpha = \beta+1 > \omega+1$, and the theorem has been proved for β. Then by applying the theorem for the first β steps, we reduce this to the case of $\omega+1$.

Suppose α is a limit ordinal greater than ω, and the theorem has been proved for all smaller ordinals. Then $\alpha = \omega + \lambda$ where λ is a limit ordinal. Let d be the minimum depth of any step of the given sequence. We can construct Figure 1.8, where the first label on each arrow is a lower bound on the depth of its reductions, and the second indicates its length. Subdiagrams (1), (2), and (4) exist by strong convergence. Subdiagram (3) exists by the Compression Property for ordinals less than λ.

This transforms the sequence into a finite initial segment of depth d ending at some term s', followed by a final segment from u to t of length at most $\omega+\lambda$ and of depth at least $d+1$. The latter segment is not necessarily shorter than the sequence we started with, in the sense of number of steps, but it is

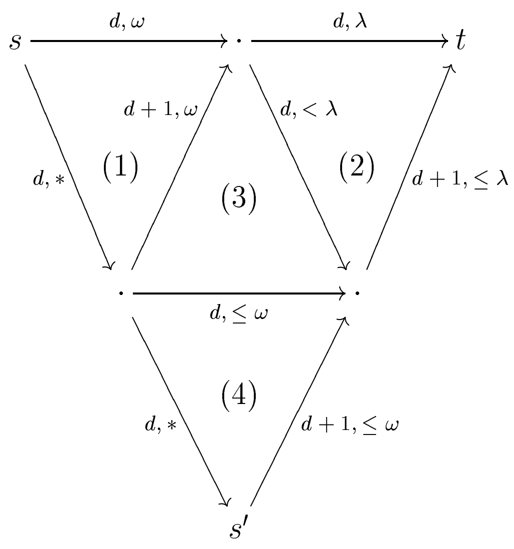

Figure 1.8: $\omega + \lambda$-compression

shorter in terms of the metric, lying entirely within a ball of diameter 2^{-d-1}. We can repeat this construction starting with the reduction of s' to t. By concatenating the finite segments $s \rightarrow^{d,*} s' \rightarrow^{d+1,*} s' \rightarrow^{d+2,*} \ldots$ we obtain a sequence of length at most ω strongly converging to t. ⊠

1.7.2. Theorem. *In lambda calculus, if $s \rightarrow^{\alpha} t$ then $s \rightarrow^{\leq\omega} t$.*

Proof. The only change required to the proof of the preceding theorem is in establishing the case $\alpha = \omega + 1$, since the argument in every other case is independent of the details of rewriting.

Let $s \rightarrow^{\omega} s_1 \rightarrow t$ be a sequence of length $\omega + 1$. We can split $s \rightarrow^{\omega} s_1$ into $s \rightarrow^{*} s_2 \rightarrow^{\omega} s_1$ in such a way that the redex reduced in $s_1 \rightarrow t$ is present at the same position in s_2, and every step of $s_2 \rightarrow^{\omega} s_1$ takes place at greater depth. All that is then required is to reduce that redex in s_2 to give a term s_3, and demonstrate that the remainder of the sequence $s_2 \rightarrow^{\omega} s'$ can be performed beginning from s_3 and permuted into an order of length not more than ω. The detailed construction can be found in Kennaway et al. [1997]. ⊠

It can be proved that the compressed sequence is Lévy equivalent to the given sequence, in the sense that a tiling diagram for the two sequences exists, and its right and bottom sides are empty. Thus the shorter sequence is just a reordering of the longer one.

To demonstrate the necessity of left-linearity for the Compression property, consider the system whose rules are $F(x) \rightarrow A(F(x))$ and $G(x,x) \rightarrow H$. The term $G(F(C), F(D))$ reduces in ω steps to the term $G(A(A(A(\ldots))), A(A(A(\ldots))))$, and then in one step to H. But there is no reduction of $G(F(C), F(D))$ to H in fewer than $\omega + 1$ steps.

Compression is technically useful, as it can be used to simplify proofs about transfinite rewrite systems.

Compression fails for weakly convergent reduction. An example is the rules $A(x) \to A(B(x))$ and $B(x) \to C(x)$. These allow a weakly convergent reduction $A(D) \to A(B(D)) \to A(B(B(D))) \to \ldots A(B^\omega) \to A(C(B^\omega))$ of length $\omega + 1$. There is clearly no reduction of $A(D)$ to $A(C(B^\omega))$ in ω or fewer steps.

In lambda calculus, compression fails for $\beta\eta$-reduction Let $M = \lambda x.M'xx$, where $M' = Y\lambda mx.y(mx)$ and $Y = \lambda f.(\lambda x.f(xx))(\lambda x.f(xx))$. Then M contains no eta-redex, but the subexpression $(M'x)$ reduces in ω steps to the term $y^\omega = y(y(y(\ldots)))$, following which we can eta-reduce $\lambda x.y^\omega x$ to y^ω. M does not $\beta\eta$-reduce to this term in ω or fewer steps. The infinite reduction of $(M'x)$ pushes the free variable x down to infinity, creating an η-redex in the limit.

1.8. A partial Church-Rosser property

We shall now show that the failure of the Church-Rosser property is due to a certain class of terms, the hypercollapsing terms.

1.8.1. DEFINITION. A *collapsing* redex is a redex whose reduct is the residual of a proper subterm of the redex. A *collapsing* rule is a rewrite rule, every redex of which is collapsing. A term t is *hypercollapsing* if every term it can reduce to can be reduced to a collapsing redex. $\mathcal{H}$ denotes the set of hypercollapsing terms.

For term rewrite systems, a collapsing rule is one whose right-hand side is a variable. For lambda calculus, the beta and eta rules are collapsing.

1.8.2. THEOREM. *If an iTRS has two or more collapsing rules, or it has a single collapsing rule whose left-hand side contains a variable not occurring on the right, then it is not Church-Rosser.*

PROOF. In the first case, let the left-hand sides of the two rules be $C_1[x]$ and $C_2[y]$, where x and y are the variables on their respective right-hand sides. In the second case, let $C[x, y]$ be the left-hand side, where x is the variable occurring on the right and y is variable occurring on the left but not the right. Then define $C_1[x] = C[x, a]$ and $C_2[x] = C[x, b]$ where a and b are distinct variables. Now consider the term $C_1[C_2[C_1[C_2[\ldots]]]]$. This transfinitely reduces to both $C_1[C_1[C_1[\ldots]]]$ and $C_2[C_2[C_2[\ldots]]]$, which are distinct but reduce only to themselves. ⊠

We now establish a limited version of the Church-Rosser property.

1.8.3. DEFINITION. We write $s \xrightarrow{in\,\mathcal{H}} t$ for a reduction step in a hypercollapsing subterm of s, and $s \xrightarrow{out\,\mathcal{H}} t$ for a reduction step outside all hypercollapsing subterms of s. Every reduction step is of exactly one of these types. The corresponding types of reduction sequence are written $s \overset{in\,\mathcal{H}}{\twoheadrightarrow} t$ and $s \overset{out\,\mathcal{H}}{\twoheadrightarrow} t$.

1.8.4. DEFINITION. A *collapsing tower* of redexes in a term is a set of collapsing redexes $\{r_1, r_2, \ldots\}$ such that the residual of any r_{i+1} by r_i is at the same position as r_i.

1.8.5. LEMMA. *Let R be a set of redexes in a term t, and $r \in R$. If R/r contains a collapsing tower, so does R.*

PROOF. This is trivial for term rewriting, due to the fact that reductions cannot create new containment relations between residuals. For lambda calculus it is almost trivial: a single reduction step has only a limited capacity to create containment relations. If u, v, and w are positions of t, with respective residuals u', v', and w' in t' such that $u' < v' < w'$, then either $u < v$ or $v < w$. An infinite chain of containment relations therefore cannot be created in a single step. ⊠

1.8.6. REMARK. Note that it is possible for a collapsing tower of new residuals to be created in a single step, as demonstrated by the term $(\lambda x.(x(x(x(\ldots)))))(\lambda z.z)$.

1.8.7. LEMMA. *If set of redexes has no complete development, then it contains a collapsing tower.*

PROOF. Let R be a set of redexes with no complete development. Perform a maximally long development of R, reducing a residual of minimum depth at each step. The depths of successive steps are thus non-decreasing. If they increased to infinity, this would give a strongly convergent complete development, contrary to hypothesis. Therefore after finitely many steps, we must reach a term t' such that every subsequent step is at depth d. Furthermore, there must be a position u of t' at depth d at which infinitely many reductions are performed. This implies that there is a collapsing tower at u in t'. Let R' be the set of residuals of R in t'. By Lemma 1.8.5, R has a collapsing tower. ⊠

1.8.8. LEMMA. *In an orthogonal system, if $s \overset{out\,\mathcal{H}}{\twoheadrightarrow} t_0$ and $s \overset{out\,\mathcal{H}}{\twoheadrightarrow} t_1$, then for some q, $t_0 \twoheadrightarrow q$ and $t_1 \twoheadrightarrow q$.*

PROOF. Let the two given reductions be S and T. We will prove that S and T have a tiling diagram. We know this is true when S and T are at most one step each. An inductive argument over ordinals requires us to prove it when S is one step and T has limit ordinal length, and when both S and T have limit ordinal length. We may assume that for S and T shorter than those being considered, the theorem is known to hold.

The first of these cases is covered by the Strip Lemma for iTRSs, which does not require the *out* $\mathcal{H}$ restriction. For lambda calculus, we must show that the set of residuals of (the redex) S by T has a strongly convergent complete development. Suppose it does not. By Lemma 1.8.7, S/T must contain a collapsing tower. Let it be at position u. Consider a non-empty final segment $T' : t \stackrel{out\,\mathcal{H}}{\twoheadrightarrow\!\!\!\rightarrow} t'$ of T which performs no reductions at positions which are prefixes of u (possible by strong convergence and limit length of T). Then $t|_u \stackrel{out\,\mathcal{H}}{\twoheadrightarrow\!\!\!\rightarrow} t'|_u$. Since $t'|_u$ is hypercollapsing, so is $t|_u$. Therefore $t|_u \stackrel{out\,\mathcal{H}}{\twoheadrightarrow\!\!\!\rightarrow} t'|_u$ is the empty sequence and $t|_u = t'|_u$. But then $t|_u$ contains a set of residuals of r having no complete development, contradicting the inductive assumption.

We now consider the case where S and T both have limit ordinal length. Assume the notation of Definition 1.6.1 for the terms and sequences appearing in their tiling diagram. We have to show that the right and bottom sides of the diagram are strongly convergent and have the same limit.

Suppose that this is false. By Theorem 1.6.5, this implies that neither the right nor the bottom edge is strongly convergent, and the diagram is neither vertically nor horizontally uniformly convergent. From this we will demonstrate the presence of a reduction within a hypercollapsing subterm in the top edge, contradicting the hypotheses.

From the non-convergence properties, we can construct a zig-zag reduction path through the diagram, starting from the top left corner, and proceeding through infinitely many horizontal and vertical segments, which is not strongly convergent. See Figure 1.9.

Let d be the smallest depth for which such a path can be found containing infinitely many steps of depth d. Then some lower right segment of the diagram will contain no steps of depth less than d. By choosing a suitable subterm of the top left term of depth d, and considering the reduction sequences restricted to that subterm, we can assume without loss of generality that $d = 0$. Finally, by symmetry we can assume that the path contains infinitely many vertical steps of depth 0.

When the right edge of an elementary tile contains a redex of depth zero, either the left edge also contains such a redex, or if it does not, the top edge must contain a collapsing redex of depth zero, as in Figure 1.10. Since S contains no steps of depth zero, wherever there is a vertical redex of depth 0 in the term $t_{(i,j)}$, there must be a $j' < j$ such that $\mathcal{T}_{i,j'}$ includes a collapsing redex of depth 0. See Figure 1.11.

We now change the path so as to begin with $S_{[0,i],0} \cdot T_{i,[0,j]} \cdot S_{i,j}$. This gives us a path which contains a collapsing redex at depth zero, and then continues to later perform infinitely many steps at depth zero. Consider the subdiagram whose top left corner is at $(i+1, j)$. We have a path from $(i+1, j)$ containing

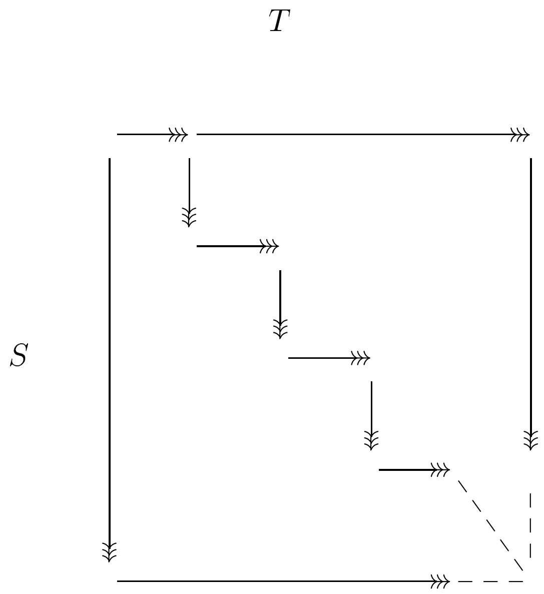

Figure 1.9: Non-strongly-convergent zig-zag path in a tiling diagram

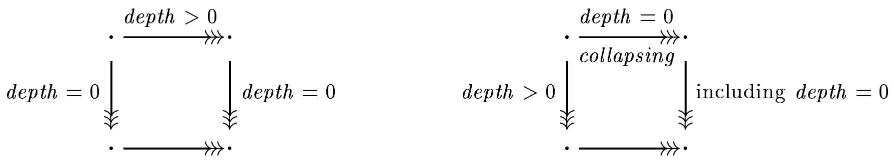

Figure 1.10: The arising of residuals of depth zero

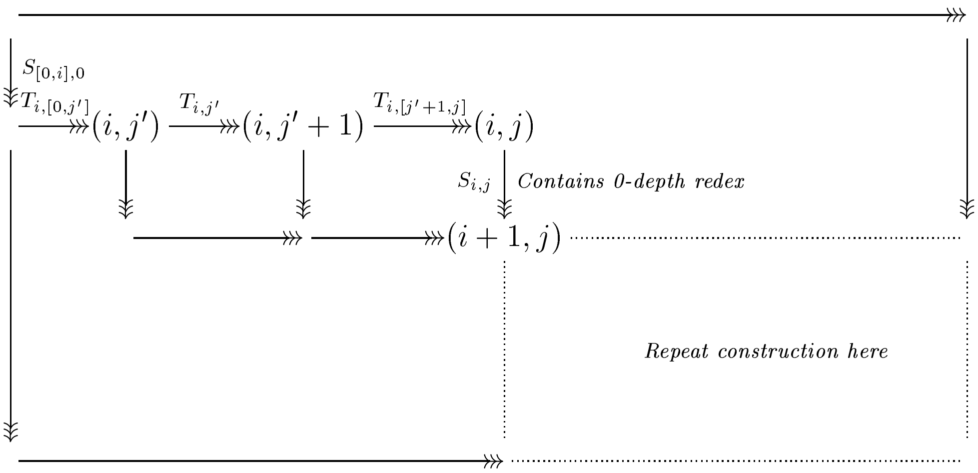

Figure 1.11: Construction for Lemma 1.8.8

infinitely many redexes of depth zero, and therefore a similar path in the diagram from (i', j) for some $i' \geq i$ such that the left edge of that diagram contains no steps of depth zero. Thus we can repeat the argument to show that we can find a path from $(i+1, j)$ onwards that also contains a collapsing step of depth 0.

Therefore the path contains infinitely many collapsing steps of depth 0. This implies that $t_{0,0}$ is hypercollapsing, contradicting the assumption that S and T reduce no redexes in hypercollapsing subterms. ⊠

Note that in both iTRSs and lambda calculus the reductions from t_0 and t_1 to q in the above lemma need not be *out* $\mathcal{H}$-reductions, because there is no particular reason that a residual of an *out* $\mathcal{H}$ redex by an *out* $\mathcal{H}$ redex would be *out* $\mathcal{H}$.

1.8.9. Corollary. *An orthogonal iTRS containing no collapsing rules is transfinitely Church-Rosser.*

Proof. Since there are no hypercollapsing terms, $\xrightarrow{out\ \mathcal{H}}\!\!\!\!\twoheadrightarrow\!\!\!\!\twoheadrightarrow$ coincides with $\twoheadrightarrow\!\!\!\!\twoheadrightarrow$. ⊠

In the next section we will establish more refined Church-Rosser properties in a more general setting.

1.9. Meaningless terms

1.9.1. Concepts of meaninglessness

The hypercollapsing terms do not have any obvious meaning. The function symbols A and B in the canonical counterexample both represent the identity function, and the two limit terms $A(A(A(\ldots)))$ and $B(B(B(\ldots)))$ can both be interpreted as least fixed points of the identity. The terms in the other examples are of a similar nature.

In a denotational semantics the hypercollapsing terms would usually be mapped to $\perp$ (called 'bottom'), the undefined element of a semantic domain. So it is natural to consider them 'meaningless' and to identify them with each other. When we do this, we find that a Church-Rosser property is restored, and that direct Lévy equivalence is once again an equivalence relation. These properties are not unique to the set of hypercollapsing terms. The same is true if we instead take certain other sets of terms and identify them with each other. For example, we may define the *root-active* terms to be those terms, each reduct of which can be reduced to a redex. Let $\mathcal{R}$ be the set of root-active terms. This set is also a candidate for a notion of 'undefinedness' in a term rewrite system.

In Kennaway et al. [1999] we have axiomatised some properties which such classes of terms should have, and proved that for classes satisfying them, certain standard theorems relating to the concept of undefinedness can be proved, theorems which in the past have only been proved for particular notions of undefinedness. These theorems are Genericity, an approximate Church-Rosser property, Relative Consistency, and the existence and uniqueness of Böhm trees. The last of these gives rise to denotational semantics for orthogonal term rewrite systems.

One of the axioms requires the following definition:

1.9.1. DEFINITION. Let t be an instance $\sigma(l)$ of the left-hand side l of some rewrite rule. t is thus a redex. This redex is said to *overlap* its subterm at position u if u is a non-empty position of l and $l \mid u$ is not a variable.

This definition can be stated for the lambda calculus in more concrete terms: a redex $(\lambda x.s)t$ overlaps just the subterm $\lambda x.s$.

The following are the properties of a set of terms $\mathcal{U}$ which make $\mathcal{U}$ a suitable notion of undefinedness.

Descendants: Every descendant of a subterm in $\mathcal{U}$ is in $\mathcal{U}$.

Overlap: If a redex t overlaps a subterm, and this subterm is in $\mathcal{U}$, then $t \in \mathcal{U}$. For lambda calculus, this is equivalent to: if $\lambda x.s \in \mathcal{U}$ then $(\lambda x.s)t \in \mathcal{U}$.

Hypercollapse: $\mathcal{H} \subseteq \mathcal{U}$.

Root-activeness: $\mathcal{R} \subseteq \mathcal{U}$.

Indiscernability: Define $s \xleftrightarrow{\mathcal{U}} t$ if s can be transformed into t by replacing a set of pairwise disjoint of subterms of s in $\mathcal{U}$ by terms in $\mathcal{U}$. If $s \xleftrightarrow{\mathcal{U}} t$. then $s \in \mathcal{U}$ if and only if $t \in \mathcal{U}$.

These properties are essentially those studied in Kennaway et al. [1999], with some simplifications. Note that since $\mathcal{H} \subseteq \mathcal{R}$, the Root-activeness axiom implies the hypercollapse axiom. For lambda-calculus, both classes are identical, since the β-rule is a collapsing rule. For more details about classes of terms which have these properties, and many examples, we refer the reader to Kennaway et al. [1999] and Kennaway et al. [1997]. The latter paper demonstrates that many notions of meaningless term that have appeared in the lambda calculus literature satisfy the axioms. We will explain here the computational motivation for the axioms.

If computation is performed by reduction, then a meaningless term, or a meaningless subterm of a term, should not allow information to be obtained from it by reduction. Thus, its descendants in any computation should also be meaningless.

To extract information from a term t, one may reduce t; more generally, one may place t in some context $C[\,]$ and reduce that. For example, $\mathit{Print}(t)$ might be intended to reduce to some sort of printable representation of t. In order for this to happen, t must be reducible to a term which can be pattern-matched 'from outside', that is, a term which can be overlapped by a redex. If t is meaningless, it should not be possible to extract information from t, therefore whenever such a t is overlapped by a redex, that redex should also be meaningless. This gives the Overlap property.

The Overlap property can also be related to Knuth-Bendix completion. If we have a rule which rewrites any member of $\mathcal{U}$ to the undefined symbol $\perp$ (as we will do later), then a redex which overlaps a subterm in $\mathcal{U}$ is an example of a conflict between that rule and the rule for the redex. The conflict is resolved if the redex itself is also in $\mathcal{U}$.

We have described above the motivation for considering all the hypercollapsing terms to be meaningless. The root-active terms represent computations which continue indefinitely without ever reaching even a partial final result, since further computation at the root of the term can always take place. These terms may reasonably be considered meaningless. Technically, the Root-activeness axiom ensures the existence of Böhm normal forms for all terms. This class includes all the hypercollapsing terms. For lambda calculus the two classes coincide, since the β-rule is in effect a collapsing rule.

The Indiscernability axiom expresses that the meaningfulness of a term does not depend on the identity of its meaningless subterms.

The following lemma precisely relates the Descendants property to the axioms of Kennaway et al. [1999].

1.9.2. Lemma. *Consider the following two properties of a class of terms $\mathcal{U}$:*

Closure: *$\mathcal{U}$ is closed under reduction.*

Substitution: *$\mathcal{U}$ is closed under substitution.*

For an iTRS, the Descendants property is equivalent to Closure, and for the lambda-calculus, it is equivalent to Closure and Substitution.

Proof. In an iTRS, the possible descendants of a subterm are exactly all its possible reducts. In the lambda-calculus, its possible descendants are exactly the possible reducts of all its substitution instances. ⊠

We now state the properties which follow from these axioms.

1.9.2. Consequences of meaninglessness

Genericity

1.9.3. Definition. A *totally meaningful* term is a term of which no subterm is in $\mathcal{U}$.

Suppose that for every term s in $\mathcal{U}$ and every context $C[\,]$, if $C[s]$ reduces to a totally meaningful term t, then $C[r]$ reduces to t for every r. Then $\mathcal{U}$ is called a *generic* set.

1.9.4. THEOREM. *In a left-linear system, if $\mathcal{U}$ has the Descendants and Overlap properties, it is generic.*

Church-Rosser modulo $\mathcal{U}$

1.9.5. DEFINITION. Given an equivalence relation $\approx$ on the terms of a rewrite system, the reduction relation is said to be *confluent up to* $\approx$ if $s \twoheadleftarrow\!\!\!\twoheadleftarrow\!\!\!\!\leftarrow \cdot \twoheadrightarrow\!\!\!\!\!\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow t$ implies $s \twoheadrightarrow\!\!\!\!\!\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow \cdot \approx \cdot \twoheadleftarrow\!\!\!\twoheadleftarrow\!\!\!\!\leftarrow t$. It is said to be *confluent modulo* $\approx$ if $s \twoheadleftarrow\!\!\!\twoheadleftarrow\!\!\!\!\leftarrow \cdot \approx \cdot \twoheadrightarrow\!\!\!\!\!\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow t$ implies $s \twoheadrightarrow\!\!\!\!\!\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow \cdot \approx \cdot \twoheadleftarrow\!\!\!\twoheadleftarrow\!\!\!\!\leftarrow t$. The reduction relation is said to be *confluent up to (resp. modulo)* $\mathcal{U}$ if it is confluent up to (resp. modulo) $\overset{\mathcal{U}}{=}$.

1.9.6. THEOREM. *In an orthogonal system, if $\mathcal{U}$ has the Descendants, Overlap, Hypercollapse, and Indiscernability properties, the system is confluent modulo $\mathcal{U}$.*

For iTRSs, we can prove a weaker confluence property from just one of the axioms.

1.9.7. THEOREM. *In an orthogonal iTRS, if $\mathcal{U}$ has the Hypercollapse property, the iTRS is confluent up to $\mathcal{U}$.*

Relative Consistency

1.9.8. DEFINITION. A rewriting system is called *consistent* if it contains two normal forms which are not related by $(\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow \cup \twoheadleftarrow\!\!\!\twoheadleftarrow\!\!\!\!\leftarrow)^*$.

For a set of terms $\mathcal{U}$, it is called *$\mathcal{U}$-consistent* if there exist two totally meaningful terms (with respect to $\mathcal{U}$) which are not related by $(\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow \cup \twoheadleftarrow\!\!\!\twoheadleftarrow\!\!\!\!\leftarrow \cup \overset{\mathcal{U}}{=})^*$.

The system is *relatively consistent* with respect to $\mathcal{U}$ if for all totally meaningful terms s and t, $s(\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow \cup \overset{\mathcal{U}}{=} \cup \twoheadleftarrow\!\!\!\twoheadleftarrow\!\!\!\!\leftarrow)^*t$ implies $s(\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow\!\!\!\!\!\twoheadrightarrow \cup \twoheadleftarrow\!\!\!\twoheadleftarrow\!\!\!\!\leftarrow)^*t$.

1.9.9. THEOREM. *In an orthogonal system, if $\mathcal{U}$ has the Descendants, Overlap, Hypercollapse, and Indiscernability properties, the system is relatively consistent with respect to $\mathcal{U}$.*

Böhm Trees

In the lambda calculus, Böhm trees are used to represent 'normal forms' of terms which may be infinite and have undefined subterms (Barendregt [1984]).

Roughly speaking, they are finite or infinite lambda expressions containing no redexes, and which may contain an extra nullary function symbol $\perp$ to represent undefined subterms.

We generalise this notion to iTRSs and infinitary lambda calculus provided with a class of undefined terms. In this setting, Böhm trees will be exactly normal forms with respect to transfinite reduction by the rules of the system, together with an extra rule rewriting undefined terms to $\perp$.

1.9.10. Definition. Given a set of undefined terms $\mathcal{U}$, add a new nullary symbol $\perp$ to an infinitary system, and define $\mathcal{U}_\perp$ to be the set of terms t (in the extended system) for which some replacement of occurrences of $\perp$ by terms in $\mathcal{U}$ can yield a term in $\mathcal{U}$.

Böhm reduction (over a set $\mathcal{U}$) is reduction by the rules of the original system together with the $\perp$-rule: $t \to \perp$ if $t \in \mathcal{U}_\perp$ and $t \neq \perp$. We write $\to_{\mathcal{U},\perp}$ for reduction by the $\perp$-rule, and $\to_{\mathcal{B}_\mathcal{U}}$ for Böhm reduction. When $\mathcal{U}$ is clear from context we may write $\to_\perp$ and $\to_\mathcal{B}$.

A *Böhm tree* over a set $\mathcal{U}$ is a normal form of the extended system with respect to Böhm reduction.

A *Böhm tree* of a term t over $\mathcal{U}$ is a normal form of t by Böhm reduction over $\mathcal{U}$.

1.9.11. Theorem. *In an orthogonal system, if $\mathcal{U}$ has the Root-activeness property, then every term has a Böhm tree. If $\mathcal{U}$ also has the Descendants, Overlap, and Indiscernability properties, then every term has a unique Böhm tree.*

1.9.12. Example. When $\mathcal{U}$ is the set of terms having no head normal form, the Böhm tree of a term as defined above is the classical Böhm tree defined in Barendregt [1984]. When it consists of the terms having no weak head normal form, it gives the Lévy-Longo trees (Lévy [1976], Longo [1983]), and when it consists of the root-active terms, it gives the Berarducci trees Berarducci [1994].

1.9.13. Corollary. *Every orthogonal system has the UN property.*

Undefined terms cast a further light on the relation between weak convergence and strong convergence. Write $t \twoheadrightarrow\!\!\!\!\rightarrow^w s$ for a weakly convergent reduction, and define *weak Böhm reduction* to be the union of weak reduction and $\twoheadrightarrow\!\!\!\!\rightarrow_\perp$.

1.9.14. Theorem. 1. *Assume that $\mathcal{U}$ has the Root-activeness property. If $t \twoheadrightarrow\!\!\!\!\rightarrow^w s$, then the sequence of depths of out $\mathcal{U}$ steps tends to infinity.*

2. *Assume that $\mathcal{U}$ has the Root-activeness, Descendants, Overlap, and Indiscernability properties. If $t \twoheadrightarrow\!\!\!\!\rightarrow^w s$, then there is an s' such that $t \twoheadrightarrow\!\!\!\!\rightarrow s' \stackrel{\mathcal{U}}{=} s$.*

1.9.15. COROLLARY. *Theorem 1.9.11 also holds for weak Böhm reduction.*

PROOF. 1. Given a reduction $t_0 \twoheadrightarrow\!\!\!\!\rightarrow^w t_\alpha$, let u be a position of minimum depth such that every final segment of the sequence contains a reduction at u. Then some final segment, starting from some t_β, will contain no reductions at any proper prefix of u. Therefore $t_\beta|u$ reduces to $t_\alpha|u$ by a weakly convergent reduction with infinitely many steps at the root. By Root-activeness, every term of that sequence is in $\mathcal{U}$, and hence the corresponding steps of the original sequence are not *out* $\mathcal{U}$ steps.

2. Construct a reduction starting from t which omits all the steps of the given sequence which are not *out* $\mathcal{U}$, by methods similar to those of Kennaway et al. [1999]. Completing the argument is left as an exercise.

⊠

1.9.3. Proofs

Our proofs of these properties are in the spirit of Barendregt [1984] (ch. 15) for the finitary lambda calculus. The proofs depend on a collection of commutation properties of certain relations constructed from $\mathcal{U}$.

Similarly to Definition 1.8.3, we write $\xrightarrow{in\,\mathcal{U}}$ and $\xrightarrow{out\,\mathcal{U}}$ to denote reduction steps inside a subterm of $\mathcal{U}$ or outside all such subterms.

1.9.16. DEFINITION. 1. Let A be a set of pairwise disjoint positions of subterms of s in $\mathcal{U}$. $s \xrightarrow{\mathcal{U}}_A t$ if s can be transformed into t by replacing those subterms by arbitrary terms.

2. With A and s as before, $s \xleftrightarrow{\mathcal{U}}_A t$ if s can be transformed into t by replacing the subterms of s at A by terms in $\mathcal{U}$. We omit A where it is not important to identify it. (This coincides with the definition of $\xleftrightarrow{\mathcal{U}}$ introduced in the statement of the Indiscernability property.)

3. $\stackrel{\mathcal{U}}{=}$ is the transitive closure of the union of $\xleftrightarrow{\mathcal{U}}_A$ over all A.

1.9.17. LEMMA. *In any rewrite system, if $\mathcal{U}$ has the Indiscernability property then $\xleftrightarrow{\mathcal{U}}$ is transitive.*

PROOF. Suppose that $r \xleftrightarrow{\mathcal{U}}_A s \xleftrightarrow{\mathcal{U}}_{A'} t$. Let A'' be the set of minimal members of $A \cup A'$. Then the Indiscernability property implies that every subterm of r, s, and t at positions in A'' is in $\mathcal{U}$. From this it follows that $r \xleftrightarrow{\mathcal{U}}_A t$. ⊠

It is not clear to us whether the converse is true. If $C[s]$, s, and t are all in $\mathcal{U}$, and $\xleftrightarrow{\mathcal{U}}$ is transitive, then for any $r \in \mathcal{U}$ we have $r \xleftrightarrow{\mathcal{U}} C[s] \xleftrightarrow{\mathcal{U}} C[t]$

and therefore $r \overset{\mathcal{U}}{\longleftrightarrow} C[t]$. If there are two distinct function symbols such that $\mathcal{U}$ contains members having each of those symbols at their roots, then we can choose r to have a root symbol different from that of $C[t]$, implying that $r \overset{\mathcal{U}}{\longleftrightarrow}_{\{\langle\rangle\}} C[t]$ and therefore $C[t] \in \mathcal{U}$. Alternatively, if $\mathcal{U}$ contains any finite terms, then we can choose r to be a member of $\mathcal{U}$ of which no proper subterm is in $\mathcal{U}$, leading to the same conclusion. These examples suggest that any counterexample to the converse is likely to be highly artificial. However, we currently lack either a proof of the converse or a counterexample.

1.9.18. LEMMA. *In a left-linear system, if $\mathcal{U}$ has the Descendants and Indiscernability properties then $s \overset{in\,\mathcal{U}}{\twoheadrightarrow\!\!\!\rightarrow} t$ implies $s \overset{\mathcal{U}}{\longleftrightarrow} t$.*

PROOF. Suppose $s \overset{in\,\mathcal{U}}{\twoheadrightarrow\!\!\!\rightarrow} t$. Let A be the set of positions of maximal subterms of s in $\mathcal{U}$. Suppose that some step of the sequence is performed at a position of which no member of A is a prefix. There must be a first such step $s' \overset{in\,\mathcal{U}}{\longrightarrow} t'$. Let its position be u. Since this is a $\overset{in\,\mathcal{U}}{\longrightarrow}$-reduction, $s'|_u \in \mathcal{U}$. Since all previous steps are within subterms at positions in A, The Descendants property implies that $s'|_u \overset{\mathcal{U}}{\longleftrightarrow} s|_u$. By Indiscernability, $s|_u \in \mathcal{U}$. But this implies that some member of A is a prefix of u, contrary to the choice of u. Therefore the entire reduction of s to t happens within the subterms at positions in A. By Descendants, those subterms of t are also in $\mathcal{U}$, and $s \overset{\mathcal{U}}{\longleftrightarrow} t$. ⊠

1.9.19. LEMMA. *In a left-linear system, assume $\mathcal{U}$ has the Overlap property.*

1. *If $s \overset{\mathcal{U}}{\longleftarrow} \cdot \overset{out\,\mathcal{U}}{\twoheadrightarrow\!\!\!\rightarrow} t$, then $s \twoheadrightarrow\!\!\!\rightarrow \cdot \overset{\mathcal{U}}{\longleftarrow} t$.*
2. *If $s \overset{\mathcal{U}}{\longleftrightarrow} \cdot \overset{out\,\mathcal{U}}{\twoheadrightarrow\!\!\!\rightarrow} t$, then $s \twoheadrightarrow\!\!\!\rightarrow \cdot \overset{\mathcal{U}}{\longleftrightarrow} t$.*
3. *If $s \overset{\mathcal{U}}{=} \cdot \overset{out\,\mathcal{U}}{\twoheadrightarrow\!\!\!\rightarrow} t$, then $s \twoheadrightarrow\!\!\!\rightarrow \cdot \overset{\mathcal{U}}{=} t$.*

PROOF.
1. Suppose $s \overset{\mathcal{U}}{\longleftarrow}_A r \overset{out\,\mathcal{U}}{\twoheadrightarrow\!\!\!\rightarrow} t$. By the Overlap property, each reduction in the sequence $r \overset{out\,\mathcal{U}}{\twoheadrightarrow\!\!\!\rightarrow} t$ must reduce a redex whose pattern lies entirely outside the residuals of the positions in A. This implies that the sequence of positions of reductions in that sequence defines a reduction sequence starting from any s for which $r \overset{\mathcal{U}}{\longrightarrow} s$. (The hypothesis of left-linearity is required to ensure that those positions are still the positions of redexes after the substitution $r \overset{\mathcal{U}}{\longrightarrow} s$.) Corresponding members r' and s' of the two sequences are related by $r' \overset{\mathcal{U}}{\longrightarrow}_{A'} s'$, where A' is the set of residuals of A in r.
2. This is established by a proof similar to the previous item.
3. Immediate from the previous item.

⊠

1.9.20. LEMMA. *In a left-linear system, assume $\mathcal{U}$ has the Descendants and Overlap properties.*

1. *If $s \xrightarrow{in\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow \cdot \xrightarrow{out\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow t$ then $s \xrightarrow{out\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow \cdot \xrightarrow{in\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow t$.*

2. *If $s \twoheadrightarrow\!\!\!\!\rightarrow t$ then $s \xrightarrow{out\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow \cdot \xrightarrow{in\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow t$.*

3. *If $s \twoheadrightarrow\!\!\!\!\rightarrow t$ and t is totally meaningful then $s \xrightarrow{out\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow t$.*

PROOF. 1. Let A be the set of minimal positions of undefined subterms of s. Then the reduction $s \xrightarrow{in\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow r$ consists of an interleaving of reductions of each of those subterms. The proof then follows the same pattern as for that of Lemma 1.9.19(1). The required reductions of the form $s' \xrightarrow{in\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow r'$ are constructed by applying the same reductions to the subterms of s' at A' as were applied to their ancestors in s.

2. If $s \twoheadrightarrow\!\!\!\!\rightarrow t$ consists of an alternation of finitely many $\xrightarrow{in\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow$ and $\xrightarrow{out\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow$ segments, then this is established by a finite number of applications of item (1). If there are infinitely many segments, a more complex argument is required. First, we use the Compression property to assume without loss of generality that $s \twoheadrightarrow\!\!\!\!\rightarrow t$ has length ω, and hence that it can be expressed as an alternation of only ω many $\xrightarrow{in\,\mathcal{U}}\!\!\!\!\rightarrow$ and $\xrightarrow{out\,\mathcal{U}}\!\!\!\!\rightarrow$ segments, each finitely long. We can then use the result of the previous item to construct all of Figure 1.12 except its right-hand edge. The given sequence forms the zig-zag edge of the diagram, and each square is given by item (1). To construct the right edge, the reduction of r to t, we must show that each of the horizontal sequences is strongly convergent, and that their limits can be joined by suitable vertical segments, whose concatenation will strongly converge to t.

   From the construction of item (1), each step of each horizontal sequence of the diagram is at the same depth as the corresponding horizontal step in the zig-zag. Since by hypothesis the latter is strongly convergent, so is the former.

   Each segment of the right edge exists by the same argument as used in item (1).

   Finally, we prove strong convergence of the whole right edge. Choose any depth d. By strong convergence of the given sequence, there is an n such that every step of the zig-zag after the n-th $\xrightarrow{out\,\mathcal{U}}\!\!\!\twoheadrightarrow\!\!\!\!\rightarrow$ segment has depth greater than d. Therefore every step of the top row after the n-th segment has depth greater than d, as do all the horizontal segments below those. Therefore every segment of the right edge after the n-th is the projection of one sequence of depth greater than d over another, and

$$\begin{tikzcd}
s \arrow[r, ->>, "o"'] \arrow[d, "i"', "*"] & \cdot \arrow[r, ->>, "o"'] \arrow[d, ->>, "i"'] & \cdot \arrow[r, ->>, "o"'] \arrow[d, ->>, "i"'] & \cdot \arrow[r, dashed, -] \arrow[d, ->>, "i"'] & r \arrow[d, ->>, "i"'] \\
\cdot \arrow[r, ->>, "*", "o"'] & \cdot \arrow[r, ->>, "o"'] \arrow[d, "i"', "*"] & \cdot \arrow[r, ->>, "o"'] \arrow[d, ->>, "i"'] & \cdot \arrow[r, dashed, -] \arrow[d, ->>, "i"'] & \cdot \arrow[d, ->>, "i"'] \\
& \cdot \arrow[r, ->>, "*", "o"'] & \cdot \arrow[r, ->>, "o"'] \arrow[d, "i"', "*"] & \cdot \arrow[r, dashed, -] \arrow[d, ->>, "i"'] & \cdot \arrow[d, ->>, "i"'] \\
& & \cdot \arrow[r, ->>, "*", "o"'] & \cdot \arrow[r, dashed, -] \arrow[d, dashed, -] & \cdot \arrow[d, dashed, -] \\
& & & & t
\end{tikzcd}$$

Figure 1.12: Diagram for the proof of Lemma 1.9.20(2). $i = in\,\mathcal{U}$, $o = out\,\mathcal{U}$.

therefore has depth greater than d. Therefore the right edge is strongly convergent.

Furthermore, after n segments, the terms of the right edge are within a distance of 2^{-d} of the corresponding terms of the zig-zag. Therefore the right edge has the same limit as the given sequence, t.

3. From the previous item we conclude that $s \xrightarrow{out\,\mathcal{U}}\!\!\!\twoheadrightarrow \cdot \xrightarrow{in\,\mathcal{U}}\!\!\!\twoheadrightarrow t$. But since t is totally meaningful, the Descendants property implies that the $\xrightarrow{in\,\mathcal{U}}\!\!\!\twoheadrightarrow$-reduction must be empty.

$\boxtimes$

Proof of Theorem 1.9.4. Suppose $s \in U$, t is totally meaningful, $C[\,]$ is a context, and $C[s] \twoheadrightarrow t$.

$$\begin{array}{llll}
C[r] \xleftarrow{\mathcal{U}} C[s] \twoheadrightarrow t & \Rightarrow & C[r] \xleftarrow{\mathcal{U}} C[s] \xrightarrow{out\,\mathcal{U}}\!\!\!\twoheadrightarrow t & \text{Lemma 1.9.20(3)} \\
& \Rightarrow & C[r] \twoheadrightarrow \cdot \xleftarrow{\mathcal{U}} t & \text{Lemma 1.9.19(1)}
\end{array}$$

Since t is totally meaningful, $C[r] \twoheadrightarrow t$. $\boxtimes$

1.9.21. Lemma. *In a left-linear system, assume $\mathcal{U}$ has the Descendants, Overlap, and Indiscernability properties.*

1. *If $s \overset{\mathcal{U}}{=} \cdot \twoheadrightarrow t$ then $s \twoheadrightarrow \cdot \overset{\mathcal{U}}{=} t$.*
2. *If $s \overset{\mathcal{U}}{=} \cdot \twoheadrightarrow t$ and t is totally meaningful then $s \xrightarrow{out\,\mathcal{U}}\!\!\!\twoheadrightarrow t$.*

PROOF. 1.

$$\begin{array}{rlll} s \stackrel{\mathcal{U}}{=\!=} \cdot \twoheadrightarrow\!\!\!\rightarrow t & \Rightarrow & s \stackrel{\mathcal{U}}{=\!=} \cdot \xrightarrow{out\,\mathcal{U}}\!\!\!\!\rightarrow\!\!\!\rightarrow \cdot \xrightarrow{in\,\mathcal{U}}\!\!\!\!\rightarrow\!\!\!\rightarrow t & \text{by Lemma 1.9.20(2)} \\ & \Rightarrow & s \twoheadrightarrow\!\!\!\rightarrow \cdot \stackrel{\mathcal{U}}{=\!=} \cdot \xrightarrow{in\,\mathcal{U}}\!\!\!\!\rightarrow\!\!\!\rightarrow t & \text{by Lemma 1.9.19(3)} \\ & \Rightarrow & s \twoheadrightarrow\!\!\!\rightarrow \cdot \stackrel{\mathcal{U}}{=\!=} t & \text{by Lemma 1.9.18} \end{array}$$

2. From the previous item, since a totally meaningful term is related by $\stackrel{\mathcal{U}}{=\!=}$ only to itself.

⊠

1.9.22. LEMMA. *(Postponement.) In a left-linear system, if $s \twoheadrightarrow\!\!\!\rightarrow_{\mathcal{B}} t$ then $s \twoheadrightarrow\!\!\!\rightarrow \cdot \twoheadrightarrow\!\!\!\rightarrow_{\perp} t$. (No properties of $\mathcal{U}$ need be assumed.)*

PROOF. Let the steps of $s \twoheadrightarrow\!\!\!\rightarrow_{\mathcal{B}} t$ be $s_\beta \rightarrow_{\mathcal{B}} s_{\beta+1}$, where $s = s_0$ and $t = s_\alpha$.

Define a new sequence by transfinite induction thus.

Base case: $s'_0 = s_0$.

Successor case: Suppose s'_β has been defined. If $s_\beta \rightarrow_{\mathcal{B}} s_{\beta+1}$ is a $\perp$-reduction, define $s'_{\beta+1} = s'_\beta$. Otherwise let it be a reduction by a rewrite rule applied at position u. Define $s'_\beta \rightarrow s'_{\beta+1}$ by reduction at u.

Limit case: If s'_β has been defined for all β less than a limit ordinal λ, define s'_λ to be the limit of the sequence.

To prove that this defines a reduction sequence, we must show that in the successor case, s'_β has a redex at u, and in the limit case, the limit s'_λ exists. We will also need to know that for all β, $s'_\beta \twoheadrightarrow\!\!\!\rightarrow_{\perp} s_\beta$. These can be proved simultaneously by induction. Clearly, if $s'_\beta \twoheadrightarrow\!\!\!\rightarrow_{\perp} s_\beta$ and $s_\beta \rightarrow s_{\beta+1}$, then s'_β has a redex everywhere that s_β does, and $s'_{\beta+1} \twoheadrightarrow\!\!\!\rightarrow_{\perp} s_{\beta+1}$. Since $s'_\beta \rightarrow s'_{\beta+1}$ takes place at the same place as $s_\beta \rightarrow s_{\beta+1}$, the constructed sequence is strongly convergent, and therefore s'_λ exists.

This also shows that $s'_\alpha \twoheadrightarrow\!\!\!\rightarrow_{\perp} s_\alpha = t$. ⊠

PROOF OF THEOREM 1.9.6. Suppose $s \leftarrow\!\!\!\twoheadleftarrow \stackrel{\mathcal{U}}{=\!=} \twoheadrightarrow\!\!\!\rightarrow t$. We construct Figure 1.13. Subdiagram (1) is given by Descendants, Overlap, Indiscernability, and Lemma 1.9.21(1). Subdiagrams (2) and (3) are given by Descendants, Overlap, and Lemma 1.9.20(2). Subdiagram (4) is given by Hypercollapse and Lemma 1.8.8. Descendants and Indiscernability imply that $\xleftarrow{in\,\mathcal{U}}\!\!\!\!\twoheadleftarrow \subseteq \stackrel{\mathcal{U}}{=\!=}$, so Subdiagrams (5), (6), and (7) follow from the same hypotheses as (1). ⊠

PROOF OF THEOREM 1.9.7. This requires quite a complicated argument. We shall give an outline here and refer the reader to Kennaway et al. [1995] and Kennaway et al. [1997] for the details.

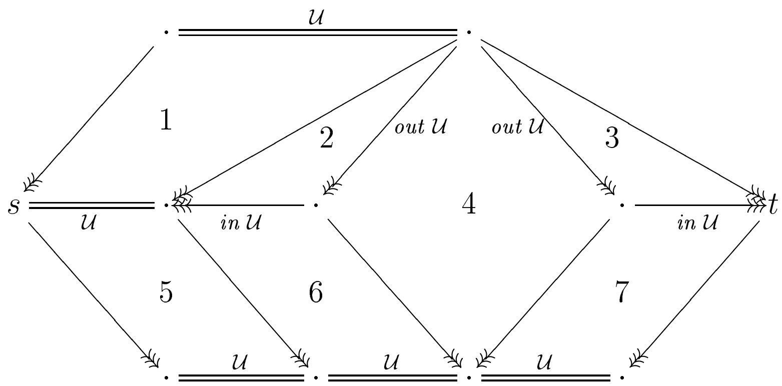

Figure 1.13: Proof of confluence up to $\mathcal{U}$ (Theorem 1.9.6)

First, we transform the system into a system which is not merely non-collapsing, but also depth-preserving, that is, one in which the depth of any occurrence of a variable in the right-hand side of a rule is at least as great as its depth in the left-hand side. This is done by inserting a new nullary function symbol ϵ into the right-hand sides as necessary to make this so, and adding further rules with ϵs on the left-hand side as necessary to preserve the reduction relation. All reductions in such systems are strongly convergent, and orthogonal systems are confluent. Not all terms of the resulting system are translatable back to the original system: subterms of the form $\epsilon(\epsilon(\epsilon(\ldots)))$ cannot be translated. These subterms arise only from reductions in hypercollapsing subterms.

Thus, if we are given two reductions $s \twoheadrightarrow\!\!\!\!\!\rightarrow t_0$ and $s \twoheadrightarrow\!\!\!\!\!\rightarrow t_1$ in the original system, we can transform these into reductions $s \twoheadrightarrow\!\!\!\!\!\rightarrow t'_0$ and $s \twoheadrightarrow\!\!\!\!\!\rightarrow t'_1$ in the transformed system, extend them by confluence to a common endpoint r, and translate the resulting sequences back to sequences $s \twoheadrightarrow\!\!\!\!\!\rightarrow t_0 \twoheadrightarrow\!\!\!\!\!\rightarrow r_0$ and $s \twoheadrightarrow\!\!\!\!\!\rightarrow t_1 \twoheadrightarrow\!\!\!\!\!\rightarrow r_1$ by omitting from $t_i \twoheadrightarrow\!\!\!\!\!\rightarrow r$ all reductions within hypercollapsing subterms. The resulting r_0 and r_1 are related by $\stackrel{\mathcal{H}}{=}$.

Proof of Theorem 1.9.9. The hypotheses, together with Theorem 1.9.6, show that if s and t are in the relation $(\twoheadrightarrow\!\!\!\!\!\rightarrow \cup \stackrel{\mathcal{U}}{=} \cup \leftarrow\!\!\!\!\!\twoheadleftarrow)^*$ then for some s' and t' $s \twoheadrightarrow\!\!\!\!\!\rightarrow s' \stackrel{\mathcal{U}}{=} t' \leftarrow\!\!\!\!\!\twoheadleftarrow t$. But since s and t are totally meaningful, so are s' and t', which must therefore be identical. ⊠

The proof of Theorem 1.9.11 requires extending the set $\mathcal{U}$ to a set $\mathcal{U}_\perp$ of undefined terms over the expanded signature.

1.9.23. Definition. $\mathcal{U}_\perp$ is the set of Böhm terms having a $\perp$-instance in $\mathcal{U}$.

1.9.24. Lemma. *Let $\mathcal{U}$ have the Indiscernability property. If some $\perp$-instance of t is in $\mathcal{U}$, then every $\perp$-instance is.*

For each of the properties of Descendants, Overlap, Hypercollapse, Root-activeness, and Indiscernability, if $\mathcal{U}$ has that property, then so does $\mathcal{U}_\perp$.

Proof. Let t' and t'' be $\perp$-instances of t. Then they differ only by substitution of subterms in $\mathcal{U}$. By Indiscernability, $t' \in \mathcal{U}$ if and only if $t'' \in \mathcal{U}$.

For the second part, we prove the Overlap property as an example. Proofs for the others are equally simple. Let $t \in \mathcal{U}_\perp$, and let $C[t]$ be a redex whose pattern includes the root of t. Let $t' \in \mathcal{U}$ result from a substitution of members of $\mathcal{U}$ for $\perp$ in t. Let $C'[\ ]$ result from $C[\ ]$ by making some substitution of members of $\mathcal{U}$ for occurrences of $\perp$. Then $C'[t']$ is a redex whose pattern overlaps the root of t'. By the Overlap property for $\mathcal{U}$, $C'[t'] \in \mathcal{U}$. Therefore $C[t] \in \mathcal{U}_\perp$. ⊠

Proof of Theorem 1.9.11. From Root-activeness it follows that every $\overset{\mathit{out}\,\mathcal{U}}{\twoheadrightarrow}$-reduction is strongly convergent, and that therefore every term s is reducible by ordinary reduction to a term t containing no redexes outside subterms in $\mathcal{U}$. Reducing all maximal such subterms by the $\perp$-rule gives a Böhm tree. Therefore every term has a Böhm tree.

For uniqueness, see Figure 1.14. Suppose that $s \to_{\mathcal{B}} t_0$ and $s \to_{\mathcal{B}} t_1$. By existence, these reductions can be extended to Böhm normal forms u_0 and u_1. Lemma 1.9.22 gives us subdiagrams (1) and (3) of the figure, where v_0 and v_1 have no redexes (by the ordinary rules) outside subterms in $\mathcal{U}$. Subdiagram (2) is given by Theorem 1.9.6. The *in* $\mathcal{U}$ property of the reductions to w_0 and w_1 follows from the just-mentioned property of v_0 and v_1. By Indiscernability, $\overset{\perp}{\twoheadrightarrow}$ is contained in $\overset{\mathcal{U}_\perp}{=\!=}$. By Descendants, Indiscernability, and Lemma 1.9.18, $\overset{\mathit{in}\,\mathcal{U}}{\twoheadrightarrow}$ is contained in $\overset{\mathcal{U}_\perp}{=\!=}$. By Indiscernability and Lemma 1.9.17, $\overset{\mathcal{U}_\perp}{\longleftrightarrow}$ is transitive. Therefore $u_0 \overset{\mathcal{U}_\perp}{=\!=} u_1$. Since these are Böhm normal forms, $u_0 = u_1$. ⊠

Proof of Corollary 1.9.13. The UN property follows from Theorem 1.9.11, since normal forms are Böhm trees. ⊠

1.9.4. Examples of classes of meaningless terms

For orthogonal systems, Kennaway et al. [1999] presents four candidates for notions of undefinedness in term rewrite systems: the opaque terms, the ω-undefined terms, $\mathcal{H}$, and $\mathcal{R}$.

Call a term *opaque* if no reduct of the term can be overlapped by a redex. That is, no reduct of the term is an instance of a proper non-variable subterm of a left-hand side. This class is motivated by the idea that a meaningful term should be capable of being pattern-matched 'from outside', as described in our justification for the Overlap property. The Overlap property requires

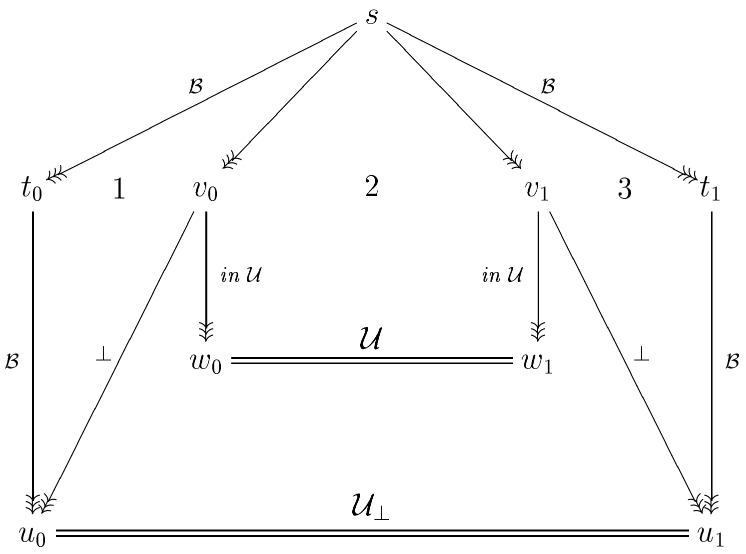

Figure 1.14: Proof of UN for Böhm reduction (Theorem 1.9.11)

that if a meaningless term is pattern-matched, the outer redex must also be meaningless. The class of opaque terms is based on the stronger idea, that the meaningless terms are precisely the terms which can never be pattern-matched.

A finite term is *ω-undefined* if all terms reachable from it can be decomposed into 'redex compatible' parts. We add a nullary symbol ω, and define a partial ordering on terms containing ω by stipulating that $\omega \leq t$ for all t, and that all function symbols are monotonic. Say that a term t (which may contain ω) is a *partial redex* if $t \leq t'$ for some redex t'. Define the ω-rule: $t \to \omega$ if t is a partial redex other than ω. It is easy to show that every finite term t has a unique normal form $\omega(t)$ by this rule. For an infinite term t, define $\omega(t)$ to be the least upper bound of $\omega(t')$ for all finite terms $t' < t$. A term t is ω-undefined if for every reduct t' (by the ordinary rewrite rules of the system) of every instance of t, $\omega(t') = \omega$. This notion is based on ideas from Huet and Lévy [1992]. The ω rule embodies the idea of ignoring the right-hand sides of the given rewrite rules.

In an orthogonal system, each of these classes satisfies all the axioms, except that $\mathcal{H}$ does not always have the Root-activeness property. The properties are all either immediate from the definitions, or follow by a simple argument from orthogonality. The theorems therefore all apply, except that Böhm trees are not necessarily unique for the class $\mathcal{H}$. See Kennaway et al. [1999] for proofs, and counterexamples for left-linear, non-orthogonal systems.

For lambda calculus, the same paper demonstrates that many concepts of undefinedness that have appeared in the literature satisfy all the axioms.

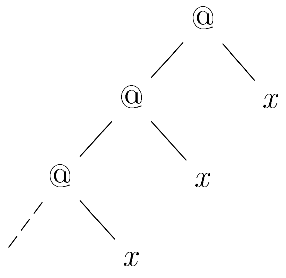

Figure 1.15: A paradoxical infinite term

1.10. A refinement of infinitary lambda calculus

Some of the infinite lambda terms as defined earlier have paradoxical properties. Consider the term $(((\ldots x)x)x)x$. See Fig. 1.15. 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 $\lambda x_1 \ldots \lambda x_n . y t_1 \ldots t_m$ — 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 $t_1, \ldots, t_n$ such that $t\, t_1 \ldots t_n$ 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 $\lambda x_1 . \lambda x_2 . \lambda x_3 . \ldots$

This suggests that we consider a different notion of depth, where all the nodes on the left spine of a term are considered to have depth zero. More generally, we can consider the three different contexts which the immediate subterms of a lambda term can have: the body of an abstraction, and the left and right components of an application. For each of these contexts, we can define the depth of the position of the subterm as either 0 or 1. This gives 8 different notions of depth. We can label each with a string of three numbers *abc*. We then have a depth measure parameterised by a, b, and c measuring the length of any position of a term:

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

Using D^{abc} as the notion of depth in Definition 1.4.2 gives 8 different metrics d^{abc} on the space of finite terms, each of which has a different metric completion. The metric d^{111} is the metric we have been using up to now, and completion gives all possible infinite terms. Completion using each of the other metrics excludes all those infinite terms which contain an infinite path along which the depth does not increase to infinity. At the opposite extreme to D^{111}, D^{000} allows no infinite terms at all and gives the discrete metric d^{000} on finite terms. There are no strongly convergent infinite reductions. This is the ordinary finitary lambda calculus. Of the other six, only two have good properties. Measure 001 considers all nodes on the left spine of a lambda term to have depth zero, and therefore excludes terms having an infinite left spine. The examples we gave earlier of paradoxical terms are all excluded. Measure 101 has connections with the lazy lambda calculus of Abramsky and Ong [1993]. It excludes infinite left-branching chains of applications, but allows infinitely nested lambdas.

Each depth measure gives a corresponding notion of Böhm tree, equivalent to those in Example 1.9.12.

1.10.1. Exercise. In the definition of D^{abc}, we could take a, b, and c to be any non-negative numbers. Prove that the topology of the resulting complete metric space depends only on which of these numbers is non-zero.

The appropriate generalisation of the class of root-active terms is the class of *0-active* terms: those terms whose every reduct reduces to a term having a redex at depth zero. For the depth measures 111, 001, and 101, the class of 0-active terms satisfies all the axioms for a class of undefined terms. This immediately gives us the Genericity property, confluence up to equality of 0-active subterms, Relative Consistency, and the existence and uniqueness of Böhm trees. For every other depth measure, the class violates at least one axiom.

1.10.2. Exercise. For each depth measure and each axiom for a class of meaningless terms, either prove that the class of 0-active terms for that measure satisfies that axiom, or exhibit a counterexample. With these results, verify the above claim.

The counterexample we gave to the Strip Lemma for the lambda calculus with measure 111 also works for the measure 101. It fails for the depth measure 001, since the infinite term $s \equiv (\lambda y.s)w$ does not exist for that measure, but we can find a similar counterexample. Let t be (as with the earlier example) a finite term with the property that $tx \twoheadrightarrow x(tx)$ (by a finitely long β-reduction). Consider the term tr, where $r \equiv \lambda y.((\lambda x.x)(\lambda w.y))$. The term tr reduces to $r(tr)$, and then to $(\lambda x.x)(\lambda w.tr)$, and so in infinitely many steps to an infinite term $s \equiv (\lambda x.x)(\lambda w.s)$. The same term also reduces in one step to $t(\lambda yw.y)$. Every reduct of s contains subterms of the form $\lambda x.x$, because it is not possible to reduce all of these in a strongly convergent reduction. However, no reduct of $t(\lambda yw.y)$ contains such subterms. They therefore have no common reduct, and the Strip Lemma diagram cannot be completed.

Historical remarks

The basic reference on transfinite rewriting is Kennaway et al. [1995] for term rewriting, and Kennaway et al. [1997] for the lambda calculus. The concepts have yet to be extended to higher-order rewrite systems. The link between transfinite term rewriting and term graph rewriting was explored in Kennaway et al. [1994]. The general axioms for notions of undefinedness in rewrite systems were established in Kennaway et al. [1999].

Bibliography

Abadi, M., L. Cardelli, P.-L. Curien and J.-J. Lévy [1991]. Explicit substitutions, *Journal of Functional Programming* **1**(4), pp. 375–416.

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

Asperti, A. and C. Laneve [1995]. Paths, computations and labels in the λ-calculus, *Th. Comp. Sci.* **142**(2), pp. 277–297.

Asperti, A., V. Danos, C. Laneve and L. Regnier [1994]. Paths in the λ-calculus, *Proceedings* 9^{th} *Annual Symposium on Logic in Computer Science,* Paris, France, IEEE Computer Society Press, pp. 426–436.

Barendregt, H.P. [1984]. *The Lambda Calculus. Its Syntax and Semantics*, second, revised edition, North-Holland, Amsterdam.

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

Dershowitz, N., S. Kaplan and D. A. Plaisted [1991]. Rewrite, rewrite, rewrite, rewrite, rewrite, *Theoretical Computer Science* **83**, pp. 71–96.

Farmer, W.M. and R.J. Watro [1990]. Redex capturing in term graph rewriting, *International Journal of Foundations of Computer Science* **1**(4), pp. 369–386.

Huet, G. and J.-J. Lévy [1992]. Computations in orthogonal rewriting systems, Part I + II, *in:* J.L. Lassez and G.D. Plotkin (eds.), *Computational Logic – Essays in Honor of Alan Robinson*, MIT Press, Cambridge MA, pp. 395–443. Update of call-by-need computations in non-ambiguous linear term rewriting systems, 1979.

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

Kennaway, J. R., J. W. Klop, M.R. Sleep and F. J. de Vries [1997]. Infinitary lambda calculus, *Theoretical Computer Science* **175**(1), pp. 93–125.

Kennaway, J.R., J.W. Klop, M.R. Sleep and F.J. de Vries [1994]. On the adequacy of graph rewriting for simulating term rewriting, *ACM Transactions on Programming Languages and Systems* **16**(3), pp. 493–523.

Kennaway, J.R., V. van Oostrom and F.J. de Vries [1999]. Meaningless terms in rewriting, *Journal of Functional and Logic Programming.*

Lévy, J.-J. [1976]. An algebraic interpretation of the $\lambda\beta K$-calculus, and an application of a labelled λ-calculus, *Theoretical Computer Science* **2**(1), pp. 97–114.

Longo, G. [1983]. Set-theoretical models of λ-calculus: theories, expansions, isomorphisms, *Annals of Pure and Applied Logic* **24**(2), pp. 153–188.

Index of Authors

Index of Notation

Index of Subjects