Transfinite Reductions in Orthogonal Term Rewriting Systems

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

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

[2,4]*Department of Software Technology, CWI,*
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

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

Abstract

We define the notion of transfinite term rewriting: rewriting in which terms may be infinitely large and rewrite sequences may be of any ordinal length. For orthogonal rewrite systems, some fundamental properties known in the finite case are extended to the transfinite case. Among these are the Parallel Moves lemma and the Unique Normal Form property. The transfinite Church-Rosser property (CR^∞) fails in general, even for orthogonal systems, including such well-known systems as Combinatory Logic. Syntactic characterisations are given of some classes of orthogonal TRSs which do satisfy CR^∞. We also prove a weakening of CR^∞ for all orthogonal systems, in which the property is only required to hold up to a certain equivalence relation on terms. Finally, we extend the theory of needed reduction from the finite to the transfinite case. The reduction strategy of needed reduction is normalising in the finite case, but not in the transfinite case. To obtain a normalising strategy, it is necessary and sufficient to add a requirement of fairness. Parallel outermost reduction is such a strategy.

AMS Subject Classification (1991): 68Q42
CR Subject Classification (1991): F1.1, F4.1, F4.2
Keywords & Phrases: Orthogonal term rewriting system, infinitary rewriting, strongly converging reduction, infinite Church-Rosser property, normalising reduction strategy.
Note: All authors were partially sponsored by SEMAGRAPH, ESPRIT Basic Research Action 3074, and SEMAGRAPH II, ESPRIT working group 6345. The first author was also partially supported by a SERC Advanced Fellowship, and by SERC Grant no. GR/F 91582.

1. Introduction

This paper extends the established theory of term rewriting to include infinite terms and infinite rewriting sequences. At first sight, such an extension might appear to be only of theoretical interest. However, it arises naturally from the use of term rewriting in functional programming, in several ways. One can write an expression whose normal form is, intuitively speaking, an infinite term — the list of all prime numbers, for example. Such infinite normal forms are approached as the limits of infinitely long reduction sequences. Infinite terms and rewrite sequences also arise when considering the correspondence between graph rewriting and term rewriting. Graph rewriting extends term rewriting with the idea of sharing subterms and is an important implementation technique for functional languages. Some implementations use cyclic graphs in order to make certain optimisations. For example, the Y combinator from combinatory logic can be used to implement functions defined by recursion. An efficient representation of the graph rewrite rule for the Y-combinator involves cyclic graphs. Cyclic graphs correspond to certain infinite terms; rewriting cyclic graphs corresponds to infinite computations on terms. Hence to study the soundness of graph rewriting as an implementation of term rewriting requires extending term rewriting with infinite terms and rewrite sequences.

The infinite computations which are useful in such contexts are the computations whose successive intermediate results reveal increasingly larger parts of the final outcome. The parts of the intermediate results that approximate the final result should be final as well, in the sense that in the remainder of the computation no rewrite step will be made in these approximating fragments. So not only should the outcome be the limit of the intermediate results (cf. example 1.1.2), but it should also be the limit of the final fragments of the intermediate results (cf. example 1.1.3). We will call such infinite reductions *strongly converging*.

EXAMPLE 1.1 Some examples of infinite computations.

1. With the rules $A \rightarrow B$ and $B \rightarrow A$, the sequence $A \rightarrow B \rightarrow \ldots$ is not converging.
2. With the rule $A(x) \rightarrow A(B(x))$, the sequence $A(x) \rightarrow A(B(x)) \rightarrow A(B(B(x))) \rightarrow A(B(B(B(x)))) \rightarrow \cdots$ converges to $A(B^\omega)$, but no part of any intermediate result is final. Each reduction is performed at the root of the (syntax tree of the) term.
3. With the rule $A \rightarrow S(A)$, the sequence $A \rightarrow S(A) \rightarrow S(S(A)) \rightarrow S(S(S(A))) \rightarrow \cdots$ converges to S^ω, which is the limit of the successive final fragments [ ], $S[\ ]$, $S(S([\ ]))$, $S(S(S([\ ])))$, ...

If an infinite computation reaches a limit which is not yet a normal form, we can compute further. So, infinite reductions may be transfinite.

In this paper we will concentrate on strongly converging reductions and reconsider the basic theory of orthogonal term rewriting systems. In a companion paper [KKSdVar] we consider the relationship between cyclic graph rewriting and transfinite term rewriting.

1.1 Overview

In Section 2 we give the basic definitions of finite and infinite terms, and of finite reduction sequences. In Section 3 we introduce Cauchy convergent and strongly convergent transfinite reduction sequences. Cauchy convergence is the well-known topological concept, applied to the topological space of finite and infinite trees. We find a stronger notion of convergence to be more useful, for reasons that will be discussed in the final section. Section 4 considers the notions of residual, projection, and Lévy-equivalence for transfinite reduction sequences, and notes that these concepts require the notion of strong convergence rather than Cauchy convergence for their definitions. Section 5 shows that every transfinite strongly convergent reduction sequence is equivalent to one of length at most ω. Section 6 considers the Church-Rosser property. In orthogonal rewrite systems, this property holds for finitary reduction, but in general does not hold for infinitary reduction. Necessary and sufficient conditions are given for the infinitary Church-Rosser property to hold in an orthogonal system. Section 7 considers properties relating to the uniqueness of normal forms in orthogonal systems. Section 8 considers reduction strategies, generalising Huet and Lévy's results concerning needed reduction in orthogonal rewrite systems. Finally, in Section 9 we compare our concepts and results with Farmer and Watro's studies of combinator reduction with cycles, and the studies of Dershowitz, Kaplan, and Plaisted on Cauchy-convergent transfinite reduction.

Table 1 summarizes the main properties of transfinite reduction, for both strong convergence and Cauchy convergence.

	Cauchy converging reductions	strongly converging reductions
Residuals can be defined	NO (3.3)	YES (4.1)
Projection of reductions	NO (3.3)	YES (4.8)
Compressing Lemma	NO [Far89], [Der91], (3.8)	YES (5.1)
Inf. Church-Rosser Property	NO (6.2)	NO (6.2)
Fair reductions result in	ω-normal forms [Der91]	normal forms (8.20)
Unique normal form properties	Partial yes (Section 9)	YES (7.15)
Unique ω-normal forms	NO (6.2)	NO (6.2)

Table 1. Basic facts for infinitary orthogonal term rewrite systems.

1.2 Acknowledgments

We acknowledge Nachum Dershowitz, Aart Middeldorp, Vincent van Oostrom and an anonymous referee for valuable comments at various stages of the paper, and Paul Taylor for the use of his commutative diagram package and his help in using it.

2. Preliminaries on term rewriting systems

First we recall briefly the basic concepts of finitary term rewriting systems. For ample introductions the reader is referred to [DJ90] and [Klo92].

Then we define infinitary term rewriting systems. As infinite terms we will consider only terms which, when considered as trees, have the property that each node is at finite distance from the root. Our notion of infinitary term rewrite system generalises finitary term rewriting systems:

- The set of terms over a signature is extended by the infinite terms over that signature.
- The right-hand side of a rewrite rule may be an infinite term.

2.1 Finitary Term Rewriting Systems

A *finitary term rewriting system* over a signature Σ is a pair $(Ter(\Sigma), R)$ consisting of the set $Ter(\Sigma)$ of finite terms over the signature Σ and a set of rewrite rules $R \subseteq Ter(\Sigma) \times Ter(\Sigma)$.

The *signature* Σ consists of a countably infinite set Var of variables $(x, y, z, \ldots)$ and a non-empty set of function symbols $(A, B, C, \ldots, F, G, \ldots)$ of various finite arities ≥ 0. Constants are function symbols with arity 0. The set $Ter(\Sigma)$ of *finite terms* $(t, s, \ldots)$ over Σ is the smallest set containing the variables and closed under function application.

The set $O(t)$ of *positions* (or occurrences) of a term $t \in Ter(\Sigma)$ is defined by induction on the structure of t as follows: $O(t) = \{\lambda\}$ if t is a variable, and $O(t) = \{\lambda\} \cup \{i \cdot u | 1 \leq i \leq n \text{ and } u \in O(t_i)\}$, if t is of the form $F(t_1, \ldots, t_n)$. If $u \in O(t)$ then the subterm $t|u$ at position u is defined as follows: $t|\lambda = t$ and $F(t_1, \ldots, t_n)|i \cdot u = t_i|u$. The *depth* of a subterm of t at position u is the length of u. Two positions are *disjoint* if neither is a prefix of the other. Two subterms of the same term are disjoint if their positions are disjoint.

Contexts are terms in $Ter(\Sigma \cup \{\Box\})$, in which the special constant $\Box$, denoting an empty place, occurs exactly once. Contexts are denoted by $C[\,]$ and the result of substituting a term t in place of $\Box$ is $C[t] \in Ter(\Sigma)$. A *proper* context is a context not equal to $\Box$.

Maps $\sigma : Var \to Ter(\Sigma)$ satisfying the equation $\sigma(F(t_1, \ldots, t_n)) = F(\sigma(t_1), \ldots, \sigma(t_n))$ are called *substitutions*.

The set R of *rewrite rules* contains pairs (l, r) of terms in $Ter(\Sigma)$, written as $l \to r$, such that the left-hand side l is not a variable and every variable occurring in r also occurs in l. The result l^σ of the application of the substitution σ to the term l is an *instance* of l. A *redex* (reducible expression) is an instance of a left-hand side of a rewrite rule. A reduction step $t \to s$ is a pair of terms of the form $C[l^\sigma] \to C[r^\sigma]$, where $l \to r$ is a rewrite rule in R. The *position* of the redex in the term $C[l^\sigma]$ is the position of the occurrence of $\Box$ in $C[\,]$. The positions of t which are *pattern-matched* by the reduction step are those of the form $u \cdot v$, where u is the position of the redex and v is the position of an occurrence of a function symbol in l. Two redexes in the same term are *disjoint* if their positions are disjoint. Concatenating reduction steps we get *finite reduction sequences* or *infinite* reduction sequences. A finite reduction sequence has an initial term t and a final term t'; we say that the sequence *reduces* t to t', and that t is *reducible* to t'. t is *convertible with* t' if there is a finite series of forwards and reverse reduction sequences of the form $t \to^* t_1 \leftarrow^* t_2 \ldots \to^* t_n \leftarrow^* t'$. (That is, convertibility is the reflexive transitive symmetric closure of $\to$.)

A *normal form* is a term containing no redexes. A term t *has a normal form* n if t is reducible to n and n is a normal form.

2.2 Infinitary Terms

An *infinitary term rewriting system* over a signature Σ is a pair $(Ter^\infty(\Sigma), R)$ consisting of the set $Ter^\infty(\Sigma)$ of finite and infinite terms over the signature Σ and a set of rewrite rules $R \subseteq Ter(\Sigma) \times Ter^\infty(\Sigma)$. It takes some elaboration to define the set $Ter^\infty(\Sigma)$ of finite and infinite terms.

The set $Ter(\Sigma)$ of finite terms for a signature Σ can be provided with an metric $d : Ter(\Sigma) \times Ter(\Sigma) \to [0, 1]$. The *distance* $d(t, s)$ between two terms t and s is 0, if t and s are equal, and 2^{-k}, otherwise, where $k \in \omega$ is the largest natural number such that all nodes of s and t at depth less than or equal to k have the same label. The set of infinitary terms $Ter^\infty(\Sigma)$ is the metric completion of $Ter(\Sigma)$. (This is all well known, see for instance [AN80]). Note that in infinitary TRSs, arities are still assumed to be finite, as are the left-hand sides of rewrite rules.

The notions of substitution, context, redex, reduction step, and normal form as defined for finitary term rewriting systems generalise trivially to the set of infinitary terms $Ter^\infty(\Sigma)$. Transfinite reduction sequences, that is, those of length greater than ω — the main subject of this paper — will be introduced in the next chapter.

We shall henceforth drop the word "infinitary"; all TRSs we consider in this paper will be infinitary.

2.3 Orthogonal term rewriting systems

The following properties for finitary term rewriting systems extend verbatim to infinitary term rewriting systems:

Definition 2.1 Let R be a finitary or an infinitary TRS.

1. A rewrite rule $l \to r$ of R is *left-linear* if no variable occurs more than once in the left-hand side l.

2. R is *non-overlapping* if for any two left-hand sides s and t, any position u in t, and any substitutions σ and $\tau : Var \to Ter(\Sigma)$ it holds that if $(t|u)^\sigma = s^\tau$ then either $t|u$ is a variable or t and s are left-hand sides of the same rewrite rule and $u = \lambda$ (i.e. non-variable

parts of different rewrite rules do not overlap and non-variable parts of the same rewrite rule overlap only entirely).

3. R is *orthogonal* if its rules are left-linear and non-overlapping.

Unless explicitly stated otherwise our results in this paper will concern orthogonal infinitary term rewrite systems. It is well known (cf. [Ros73, Klo92]) that finitary orthogonal term rewriting systems satisfy the *finitary Church-Rosser property*, i.e., ${}^*\!\!\leftarrow \circ \rightarrow^* \subseteq \rightarrow^* \circ {}^*\!\!\leftarrow$, where $\rightarrow^*$ is the transitive, reflexive closure of the relation $\rightarrow$. It is not difficult to see that infinitary orthogonal term rewriting systems with finite right-hand sides inherit this finitary Church-Rosser property. Without that condition they may not.

3. Strongly converging reductions

In [DKP91] Dershowitz, Kaplan and Plaisted have introduced (transfinite) Cauchy converging reductions. These are transfinite reduction sequences whose elements form a transfinite Cauchy sequence in the complete metric space $Ter^\infty(\Sigma)$.

Definition 3.1 A *transfinite reduction sequence* consists of a function f whose domain is an ordinal α, such that f maps each $\beta < \alpha$ to a reduction step $f_\beta \rightarrow f_{\beta+1}$. The *length* of the sequence is α if α is a limit ordinal, otherwise it is $\alpha - 1$. When α is a limit ordinal, the sequence is called *open*, otherwise it is *closed*.

Note that for limit ordinals $\beta < \alpha$, the above definition does not stipulate any relationship between f_β and the earlier terms in the sequence. A first requirement to ensure this is that the earlier terms converge to f_β (see Figure 1).

Definition 3.2 A reduction sequence as denoted above is *Cauchy continuous* if the sequence of terms $\{f_\beta | \beta < \alpha\}$ is a continuous function from α (with the usual topology on ordinals) to the metric space $Ter^\infty(\Sigma)$. More explicitly, for every limit ordinal $\lambda < \alpha$, it is required that $\forall \epsilon > 0\, \exists \beta < \lambda\, \forall \gamma (\beta < \gamma < \lambda \rightarrow d(t_\gamma, t_\lambda) < \epsilon)$. The sequence is *Cauchy convergent* if it is a closed sequence.

Note that all finitely long reduction sequences are trivially Cauchy convergent.

We denote a Cauchy convergent sequence of length α starting from a term t and ending as a term t' by $t \rightarrow^c_\alpha t'$. A Cauchy continuous open sequence of length α may be denoted by $t \rightarrow^c_\alpha \ldots$. To indicate that the length of some sequence is not more than α, we may write $t \rightarrow^c_{\leq\alpha} t'$ or $t \rightarrow^c_{\leq\alpha} \ldots$. When we do not wish to explicitly indicate the length of the sequence, we write $t \rightarrow^c_\infty t'$ for any finite or transfinite closed sequence, and $t \rightarrow^c_\infty \ldots$ for an open one. We will often consider the terms of a sequence $t \rightarrow^c_\alpha t'$ to be indexed by the ordinals from 0 to α, and write $t_0 \rightarrow^c_\alpha t_\alpha$.

For orthogonal TRSs, the relationship between t_α for a limit ordinal α and the earlier terms in the Cauchy converging sequence is not strong enough to generalise the notion of residual, which plays a fundamental role in the theory of finitary rewriting in orthogonal TRSs. Consider for example the following TRS.

Example 3.3 Let R be the orthogonal TRS given by the rules:

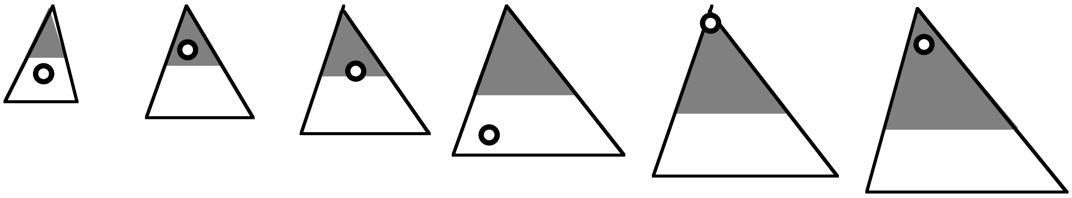

Figure 1. Cauchy converging sequence: reduction activity may occur anywhere.

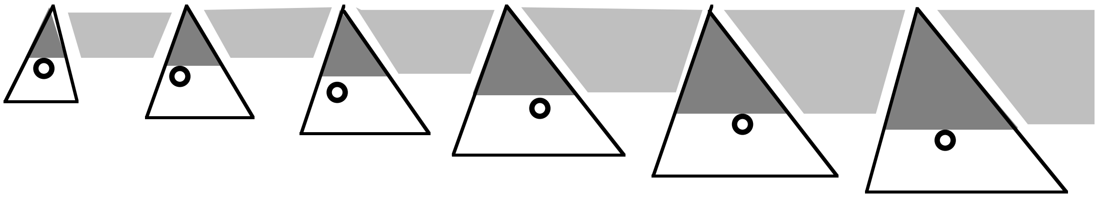

Figure 2. Strongly converging sequence: reductions occur deeper and deeper.

$$\begin{cases} A(x,y) \to A(y,x) \\ C \to D \end{cases}$$

The following infinite reduction in R is Cauchy converging:

$$A(C,C) \to A(C,C) \to A(C,C) \to \ldots \to_\omega A(C,C).$$

In contrast, the following reduction:

$$A(C,D) \to A(D,C) \to A(C,D) \to \ldots$$

has no limit.

The example shows that in the reduction $A(C,C) \to_\omega A(C,C)$ one cannot determine which occurrence of C in the original term $A(C,C)$ the left occurrence of C in the limit $A(C,C)$ is a residual of.

What is needed to extend the notion of residual to the limit of a sequence is that for any residual there will be a moment during the reduction such that no symbols of the redex pattern itself, nor any symbol between the root of the term and the residual will be rewritten. These considerations lead to the notion of transfinite *strongly* converging reductions.

Informally, we need to add the condition that as a sequence approaches a limit ordinal, the depth of the redexes which are reduced at each step tends to infinity (see Figure 2).

Definition 3.4 Given a Cauchy continuous open or closed reduction $t_0 \to^c_\alpha \ldots$ or $t_0 \to^c_\alpha t_\alpha$, let d_β be the depth of the redex reduced in $t_\beta \to t_{\beta+1}$ for $\beta < \alpha$. The sequence is *strongly continuous* (written $t_0 \to_\alpha t_\alpha$) if for every limit ordinal $\lambda < \alpha$, the sequence $\{d_\beta | \beta < \lambda\}$ tends to infinity. If the reduction is closed then it is *strongly convergent*.

Note that all finitely long reduction sequences are trivially strongly convergent.

As for Cauchy convergent reduction, we use the notations $t \to_{\leq\alpha} t'$ and $t \to_\infty t'$ for strongly convergent reductions of length up to α or of arbitrary length, respectively.

From a computational point of view the length of convergent reductions can be rather unrealistic. For example, the TRS with the single rule $A \to A$ allows Cauchy converging reductions $A \to^c_\alpha A$ for any ordinal α, even uncountable ordinals. In contrast, the length of a strongly converging reduction is always countable, which is a corollary of the following lemma.

LEMMA 3.5 *In any TRS, if $t_0 \to_\lambda t_\lambda$ is strongly convergent, then the number of steps in $t_0 \to_\lambda t_\lambda$ reducing a redex at depth $\leq n$ is finite for any n.*

PROOF. Assume $t_0 \to_\lambda t_\lambda$ is strongly convergent. If it contracts any redex at a depth no more than n, then by strong convergence there is a largest $\alpha < \lambda$ such that the step $t_\alpha \to t_{\alpha+1}$ is such a reduction step. Consider the initial segment $t_0 \to_\alpha t_\alpha$, and repeat the argument. Since there are no infinite descending chains of ordinals, this process stops in finitely many steps. □

COROLLARY 3.6 *Every strongly converging reduction has countable length.* □

EXERCISE 3.7 Given any countable ordinal α construct a strongly converging reduction of length α in the Binary Tree TRS given by the single rule $C \to B(C, C)$. When α is a countable limit ordinal, choose this reduction so as to converge to normal form.

In fact, we will later prove that in a left-linear TRS, every strongly converging reduction is equivalent to a strongly converging reduction of length at most ω (the Compressing Lemma, 5.1). It is equivalent not only in the sense that it has the same endpoints, but that it can be seen as doing the same reductions in a different order. This implies that the computational power of transfinite reduction is not increased by considering reductions longer than ω; the full power of the notion is already attained for sequences of length ω. In contrast, the compression property does not hold for Cauchy converging reductions, as shown by the following example of Farmer and Watro [FW91]:

EXAMPLE 3.8 Let R be the orthogonal TRS given by the rules:

$$\begin{cases} G(x, B) \to G(F(x), B) \\ B \to C \end{cases}$$

The term $G(A, B)$ cannot reduce to $G(F^\omega, C)$ in ω many or less steps. The shortest Cauchy-convergent reduction from $G(A, B)$ to C has length $\omega + 1$:

$$G(A, B) \to G(F(A), B) \to G(F(F(A)), B) \to_\omega G(F^\omega, B) \to G(F^\omega, C)$$

In this TRS, the right-hand sides of the rules are not all normal forms. If it is desired to find a counterexample that does have that property, that is easily arranged. Replace the first rule in the above system by $G(x, B, y) \to G(F(x), y, y)$, and consider reducing the term $G(A, B, B)$ to $G(F^\omega, C, B)$.

From now on we will concentrate on the theory of strongly converging reductions. In section 9 we will compare our results with results of Dershowitz, Kaplan and Plaisted in [DKP91] who have based their study on Cauchy converging reductions.

4. Projection of strongly converging reductions

In this section we extend the notions of residuals, developments, projection of reduction sequences, and Lévy-equivalence to the transfinite situation.

4.1 Complete developments

We saw in Example 3.3 that the notion of the residual of a subterm by a Cauchy convergent reduction in an orthogonal TRS is problematic. In this section we demonstrate that subterms do have well-defined residuals by strongly convergent reduction sequences, and that the well-known apparatus of Lévy-equivalence and projection of reduction sequences can almost entirely be duplicated for the infinitary theory based on strong convergence.

Definition 4.1 Let $S = t_0 \to_\alpha t_\alpha$ be a strongly convergent sequence. Let U be a set of positions of t. The *residuals* or *descendants* of U by S are the positions of t_α defined by induction thus.

- If $\alpha = 0$, then S is the empty sequence and U/S 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 $u/R = \{u\}$. If u is one of the positions pattern-matched by R, then $u/R = \emptyset$. 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 $u \in U$.
- If $\alpha = \beta + 1$, then $U/S = (U/(t_0 \to_\beta t_\beta))/(t_\beta \to t_{\beta+1})$.
- If α is a limit ordinal, then $u \in U/S$ iff $u \in U/(t_0 \to_\beta t_\beta)$ for all large enough $\beta < \alpha$.

In an orthogonal system, if $\mathcal{R}$ is a set of redexes of t_0, whose positions are the set U, then by $\mathcal{R}/S$ we denote the set of redexes of t_α at positions in U/S. Orthogonality ensures that these are redexes, and that a residual of a redex r is a redex of the same rule as r. We further note that when the length of S is a limit ordinal α, and u is the position of a member of U/S and is also the position of a redex, 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 the unique residual of the redex at u in t_β.

Sets of disjoint redexes will be technically important. The following proposition is immediate from the above definition.

Proposition 4.2 *If $\mathcal{R}$ is a set of disjoint redexes, then so is $\mathcal{R}/S$ for any sequence S.* □

Definition 4.3 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 $\beta \leq \alpha$ 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 developments is said to be *complete*.

Note that strong convergence is part of the definition of a development. For example, given the rule $I(x) \to x$ and the term $I(I(I(\cdots)))$ (which we abbreviate to I^ω), any reduction which attempts to reduce every redex of $I(I(I(\cdots)))$ is Cauchy convergent but not strongly convergent. The set of all redexes of $I(I(I(\cdots)))$ therefore has no complete development.

Proposition 4.4 *Every set of pairwise disjoint redexes has a complete development. The final term is uniquely determined by the set of redexes.*

Proof. Let the redexes be at positions u_1, u_2, ... in a term t. Then each redex has exactly one residual by reduction of any other redex in the set, and the residual is at exactly the same position as the original redex. The operation of replacing each redex by its reduct can therefore be carried out for each redex independently of the others, and it is clear that the resulting term does not depend on the order in which the reductions are made. Since arities are finite, only finitely many reductions can be performed at any given finite depth, implying that reduction of all the redexes, in whatever order, is strongly convergent. □

Definition 4.5 A *collapsing* rewrite rule is a rule whose right-hand side is a variable. That variable is the *collapsing variable* of the rule, and its unique position in the left-hand side of the rule is the *collapsing position* of the rule A *collapsing* redex is a redex of a collapsing rule. A *collapsing string* of redexes is a finite or infinite set of collapsing redexes in a term at positions u_0, $u_0 \cdot u_1$, $u_0 \cdot u_1 \cdot u_2$, ..., such that when i is positive and less than the number of redexes, the collapsing position of the rule matching occurrence $u_0 \cdots u_{i-1}$ is u_i. If the set is infinite, it is called a *collapsing tower*. An example of a collapsing tower is the term I^ω mentioned above.

Proposition 4.6 *(Complete developments.) In an orthogonal system, let $\mathcal{R}$ be a set of redexes in a term t. If $\mathcal{R}$ contains no collapsing tower, then $\mathcal{R}$ has a complete development. (In particular, every finite set of redexes has a complete development.) Every complete development of $\mathcal{R}$, if any, ends at the same term.*

Proof. Let n be the minimum depth of any member of $\mathcal{R}$. Since all function symbols have finite arity, there can only be finitely many members of $\mathcal{R}$ at depth n. Therefore by the standard theory of finitary rewriting, we can perform a complete development of those redexes (the fact that the whole term may be infinitely large does not affect things). We now consider the residuals of the other members of $\mathcal{R}$ by this complete development, and perform the same operation. Repeating this, either we reach in a finite number of stages a term containing no residuals of $\mathcal{R}$, or we continue for ω stages. In the first case, the reduction sequence is finitely long, therefore strongly convergent.

Consider the second case. In the first stage of the construction, we reduced all members of $\mathcal{R}$ at depth n. At the second stage, there can only be residuals of $\mathcal{R}$ at depth n if some redex reduced at the first stage was a collapsing redex. In general, a residual of $\mathcal{R}$ at the k-th stage can only be at depth n if there was initially a chain of collapsing redexes of length k starting from a redex at depth n. Since $\mathcal{R}$ contains no infinitely long chains of collapsing redexes, and there can only be finitely many finite chains starting at depth n, there is an upper bound, say m, on the length of those chains. After m stages of the construction, all remaining residuals of $\mathcal{R}$ must be at depths strictly greater than n. Repeating the argument shows that the development we have constructed is strongly convergent.

To prove completeness, note that every outermost member of $\mathcal{R}$ will eventually be reduced. If some member r of $\mathcal{R}$ is contained inside n other members of $\mathcal{R}$, then when (a residual of) the outermost of them is reduced, either a residual of r will be erased, or it will be contained in $n-1$ other members of the set of residuals of $\mathcal{R}$. Therefore it must eventually either be erased or become an outermost member, and in the latter case is eventually reduced. Therefore the limit term cannot contain any residual of $\mathcal{R}$, and so the development is complete.

We now prove by induction that every finite prefix of the final term depends only on $\mathcal{R}$ and not on the development. It is sufficient to prove that the function symbol at the root of the final term depends only on $\mathcal{R}$.

$$
\begin{tikzcd}
t_0 \arrow[r, dotted, no head] \arrow[d, "\mathcal{R}"'] & t_\beta \arrow[r, "u_\beta"] \arrow[d, "\mathcal{R}_\beta"', "\leq\omega"] & t_{\beta+1} \arrow[r, dotted, no head] \arrow[d, "\mathcal{R}_{\beta+1}"', "\leq\omega"] & t_\alpha \arrow[d, "\mathcal{R}_\alpha"', "\leq\omega"] \\
s_0 \arrow[r, dotted, no head] & s_\beta \arrow[r, "u_\beta/s_\beta"', "\leq\omega"] & s_{\beta+1} \arrow[r, dotted, no head] & s_\alpha
\end{tikzcd}
$$

Figure 3. The infinitary strip lemma

If no member of $\mathcal{R}$ is at the root of t, then the root of the final term is the same as the root of t.

Suppose a non-collapsing member of $\mathcal{R}$ is at the root of t. Let the rule for that redex be $C_l[x_1, \ldots, x_n] \to C_r[x_{i_1}, \ldots, x_{i_n}]$. Then t has the form $C_l[t_1, \ldots, t_n]$, and the limit of any complete development of $\mathcal{R}$ must have the form $C_r[t'_{i_1}, \ldots, t'_{i_n}]$, where each t'_{i_j} is the limit of a complete development, starting from t_j, of those members of $\mathcal{R}$ in t_j.

If t contains a collapsing string of n redexes starting at its root, then by some point in the sequence, all of those redexes must be reduced. Let t' be the subterm of t at the foot of the collapsing string. Then the limit of any complete development of $\mathcal{R}$ must be identical to the limit of a complete development, starting from t', of all redexes of $\mathcal{R}$ in t'. By hypothesis, none of these redexes is a collapsing redex at the root of t', so one of the previous cases applies to t'. $\square$

COROLLARY 4.7 *Let $\mathcal{R}$ and $\mathcal{R}'$ be two sets of redexes of the same term t. If $\mathcal{R} \cup \mathcal{R}'$ does not contain a collapsing tower, then the sets of redexes $\mathcal{R}/\mathcal{R}'$ and $\mathcal{R}'/\mathcal{R}$ have complete developments, and they end at the same term.*

PROOF. The sequences $\mathcal{R}\cdot(\mathcal{R}'/\mathcal{R})$ and $\mathcal{R}'\cdot(\mathcal{R}/\mathcal{R}')$ are both complete developments of $\mathcal{R}\cup\mathcal{R}'$. $\square$

4.2 The Strip Lemma

In this section we will prove a generalisation to infinitary orthogonal term rewriting of the Strip Lemma for finitary orthogonal term rewriting.

LEMMA 4.8 THE INFINITARY STRIP LEMMA. *Let $S = t_0 \to_\alpha t_\alpha$ be a strongly converging reduction in an infinitary orthogonal TRS. Let t_0 reduce to s_0 by complete development of a set of disjoint redexes $\mathcal{R}$ of t_0. Let t_α reduce to s_α by complete development of $\mathcal{R}/S$. Then there is a strongly convergent reduction of s_0 to s_α.*

PROOF. We prove this by induction on α, constructing the diagram shown in Figure 3.

For $\alpha = 0$ this is trivial.

Let $\alpha = \beta + 1$, and assume the diagram has been constructed up to t_β and s_β. Then we must construct the rightmost square of Figure 4. This is an instance of Corollary 4.7, applied to the single redex reduced in $t_\beta \to t_{\beta+1}$ and the set $\mathcal{R}/(t_0 \to_\beta t_\beta)$.

$$\begin{tikzcd}
t_0 \arrow[r, dash, dotted] \arrow[d, "\mathcal{R}"', "\leq\omega"] & t_\beta \arrow[r, "u_\beta"] \arrow[d, "\mathcal{R}_\beta"', "\leq\omega"] & t_{\beta+1} \arrow[d, "\mathcal{R}_{(\beta+1)}"', "\leq\omega"] \\
s_0 \arrow[r, dash, dotted] & s_\beta \arrow[r, "u_\beta/s_\beta"', "\leq\omega"] & s_{\beta+1}
\end{tikzcd}$$

Figure 4. Successor case

Let α be a limit ordinal. Assume that the lemma holds for all $\beta < \alpha$. There are two possibilities: either there exists a $\beta < \alpha$ such that there are no residuals of $\mathcal{R}$ in t_β, or there does not. In the first case, we define $s_\gamma = t_\gamma$ for all γ with $\beta \leq \gamma < \alpha$. It follows that $s_0 \to \cdots$ strongly converges to s_α.

Otherwise, assume there is no such β.

We will show that for any d, the entire sequence from s_β to $s_{\beta+1}$ takes place at depth at least d, for all sufficiently large β — that is, that $s_0 \to \ldots$ is strongly convergent.

Consider the complete development of $\mathcal{R}_\alpha$, and the effect this has on the depths of nodes of t_α. We will show that for every depth d, there exists a depth d' such that for every node n of t of depth at least d', every residual of n in s_α has depth at least d'. To do this we will construct for each node of t_α a lower bound on the depth of its residuals in s_α. Given any node n of t_α, label n with $d(n) - p(n)$, where $d(n)$ is the depth of n, and $p(n)$ is the number of nodes between n and the root of t, not counting n itself, which are part of the pattern of any member of $\mathcal{R}_\alpha$. Every residual of n in s_α must have depth at least $d(n) - p(n)$, since $\mathcal{R}_\alpha$ will lift a redex by the greatest possible amount if it collapses its pattern to one of the leaf nodes of the pattern.

Now consider the sequence of labels one sees on any path from the root of t_α. The sequence is obviously increasing (not not necessarily strictly increasing). If the path is infinite, the sequence is also unbounded. This follows from the disjointness of $\mathcal{R}$: there can be at most one redex on any path, and its left-hand side must be finite.

Since t_α is finitely branching, it follows from König's lemma that there can be only finitely many nodes of t_α whose labels are less than a given value (since otherwise the subtree containing the infinite set of such nodes would have to contain an infinite path with bounded labels).

Now choose any d, and then choose d' such that for every node of t_α of depth at least d', all its residuals in s_α have depth at least d. Choose β so that every step in $t_\beta \to t_\alpha$ takes place at depth at least d'. Then the prefix of t_β to depth d' is a prefix of all later t_γ. Therefore if a node of t_β is at depth d' or greater, then its residuals by complete development of $\mathcal{R}_\beta$ must also be at depths at least d in s_β.

This demonstrates not only the strong convergence of the bottom side of Figure 3, but also that its limit is s_α. □

$$
\begin{tikzcd}[column sep=small]
 & & & & T & & & & \\
S & t_{0,0} \arrow[r] \arrow[d] & t_{0,1} \arrow[d] & \cdots\cdots & t_{0,\delta} \arrow[r] \arrow[d] & t_{0,\delta+1} \arrow[d] & \cdots\cdots & t_{0,\beta} \arrow[d] & \\
 & t_{1,0} \arrow[r] & t_{1,1} & \cdots\cdots & t_{1,\delta} \arrow[r] & t_{1,\delta+1} & \cdots\cdots & t_{1,\beta} & \\
 & \vdots & \vdots & & \vdots & \vdots & & \vdots & \\
 & t_{\gamma,0} \arrow[r] \arrow[d] & t_{\gamma,1} \arrow[d] & \cdots\cdots & t_{\gamma,\delta} \arrow[r, "T_{\gamma,\delta}"] \arrow[d, "S_{\gamma,\delta}"'] & t_{\gamma,\delta+1} \arrow[d] & \cdots\cdots & t_{\gamma,\beta} \arrow[d] & S/T \\
 & t_{\gamma+1,0} \arrow[r] & t_{\gamma+1,1} & \cdots & t_{\gamma+1,\delta} \arrow[r] & t_{\gamma+1,\delta+1} & \cdots & t_{\gamma+1,\beta} & \\
 & \vdots & \vdots & & \vdots & \vdots & & \vdots & \\
 & t_{\alpha,0} \arrow[r] & t_{\alpha,1} & \cdots\cdots & t_{\alpha,\delta} \arrow[r] & t_{\alpha,\delta+1} & \cdots\cdots & t_{\alpha,\beta} & \\
 & & & & T/S & & & &
\end{tikzcd}
$$

Figure 5. Projection of reductions

4.3 Lévy equivalent transfinite reductions

Lévy defined for the lambda calculus a notion of two reduction sequences with the same endpoints doing "the same work". We extend this notion to the transfinite setting. All systems are assumed to be orthogonal.

DEFINITION 4.9 Let $S = t_0 \to_\alpha t_\alpha$ and $T = t'_0 \to_\beta t'_\beta$ be two strongly convergent reduction sequences with $t_0 = t_0$. The *projection* of S over T, when it exists, is denoted by S/T, and is constructed by induction on the lengths of S and T.

We construct the following objects. See Figure 5.

- A doubly infinite sequence of terms $t_{i,j}$, where $0 \leq i \leq \alpha$ and $0 \leq j \leq \beta$.
- A set of redexes $\mathcal{S}_{i,j}$ of $t_{i,j}$, having a complete development $S_{i,j}$ from $t_{i,j}$ to $t_{i+1,j}$, where $0 \leq i < \alpha$ and $0 \leq j \leq \beta$.
- A set of redexes $\mathcal{T}_{i,j}$ of $t_{i,j}$, having a complete development $T_{i,j}$ from $t_{i,j}$ to $t_{i,j+1}$, where $0 \leq i \leq \alpha$ and $0 \leq j < \beta$.

Let S_j be the concatenation of the sequences $S_{i,j}$ for all i, and let T_i be the concatenation of the sequences $T_{i,j}$ for all j. We also write $S_{[i,i'],j}$ for the concatenation of all $S_{i'',j}$ for $i \leq i'' < i'$, and similarly $T_{i,[j,j']}$. For example, $T_{\gamma,[0,\delta]}$ is the horizontal sequence in Figure 5 from $t_{\gamma,0}$ to the term $t_{\gamma,\delta}$ in the centre of the diagram.

We shall say that the construction has been carried out to (γ, δ) if $\mathcal{S}_{i,j}$ and $\mathcal{T}_{i,j}$ have been defined and $S_{[0,\gamma],j}$ and $T_{i,[0,\delta]}$ are strongly convergent, for all i and j such that $0 \leq i \leq \gamma$ and $0 \leq j \leq \delta$.

Define $S_{i,0}$ to be the ith step of S, the complete development of a single redex $\mathcal{S}_{i,0}$. Similarly, take $T_{0,j}$ to be the jth step of T.

It is immediate that the construction has been carried out to $(0,0)$.

Suppose that the construction has been carried out to (γ,δ). We extend it to $(\gamma+1,\delta)$ by defining for $j \leq \delta$, $\mathcal{S}_{\gamma,j} = \mathcal{S}_{0,j}/T_{\gamma,[0,j]}$. Each of the sets $\mathcal{S}_{\gamma,j}$ is a set of pairwise disjoint redexes, and hence has a strongly convergent complete development $S_{\gamma,j}$. For $j \leq \delta$, define $\mathcal{T}_{\gamma+1,j} = \mathcal{T}_{\gamma,j}/S_{\gamma,j}$. $\mathcal{T}_{\gamma+1,j}$ similarly has a strongly convergent complete development $T_{\gamma+1,j}$. Symmetrically, the construction may be carried out to $(\gamma,\delta+1)$.

This leaves the case of extending the construction to (γ,δ) where both γ and δ are limit ordinals, and the construction has been carried out for all (γ',δ') where $\gamma' \leq \gamma$ and $\delta' \leq \delta$, and at least one inequality is strict. In this case, all that is needed is to define $T_{\gamma,\delta}$. This is taken to be the common limit of the sequences $S_{[0,\gamma],\delta}$ and $T_{\gamma,[0,\delta]}$, provided these sequences are strongly convergent and have the same limit. If this fails to be the case (and we will later see that this can happen), then the construction cannot be carried out to (γ,δ).

If the construction can be carried out to (α,β), then we define $S/T = S_{[0,\alpha],\beta}$. S/T is the *projection* of S over T. We symmetrically define $T/S = T_{\alpha,[0,\beta]}$, the projection of T over S.

Definition 4.10 Two sequences S and T with the same initial term and the same final term are *Lévy-equivalent*, written $S \sim_L T$, if S/T and T/S both exist and are empty.

A square of reductions of the form S, T, S/T, and T/S will be called the *projection square* or *tile* based on S and T. It is *elementary* if S and T are both complete developments (in which case S/T and T/S will also be so).

Note that the order in which the whole of Figure 5 is constructed is not completely determined. However, it is easy to see that the final result is the same. The following equivalent description of the figure without reference to the order of its construction makes this clear. It consists of elementary tiles based on $S_{\gamma,\delta}$ and $T_{\gamma,\delta}$, where $S_{\gamma,\delta} = S_{\gamma,0}/T_{\gamma,[0,\delta]}$ and $T_{\gamma,\delta} = T_{0,\delta}/S_{[0,\gamma],\delta}$, with all the rows and columns being strongly convergent and all pairs of reductions of the form $S_{[0,\gamma],\delta}$ and $T_{\gamma,[0,\delta]}$ having the same endpoint.

The next two propositions list some basic properties of projection and Lévy-equivalence. Their proofs are immediate from the construction.

Proposition 4.11

1. *S/S is empty. Hence $S \sim_L S$.*
2. *If S/T and T/S exist, they have the same final term.*
3. *Let $\mathcal{R}$ be a set of redexes in the common initial term of S and T. If $S \sim_L T$ then $\mathcal{R}/S = \mathcal{R}/T$.*
4. *All complete developments of the same set of redexes are Lévy-equivalent.*
5. *$(S \cdot T)/U = (S/U) \cdot (T/(S/U))$; if either exists then so does the other.*
6. *$U/(S \cdot T) = (U/S)/T$; if either exists then so does the other.*
7. *$S \cdot (T/S) \sim_L T \cdot (S/T)$; if either exists then so does the other.*
8. *Let S, T, and U be sequences with suitable endpoints. If $S \sim_L T$, then $U/S = U/T$ if both exist, and $S/U \sim_L T/U$ if both exist.*

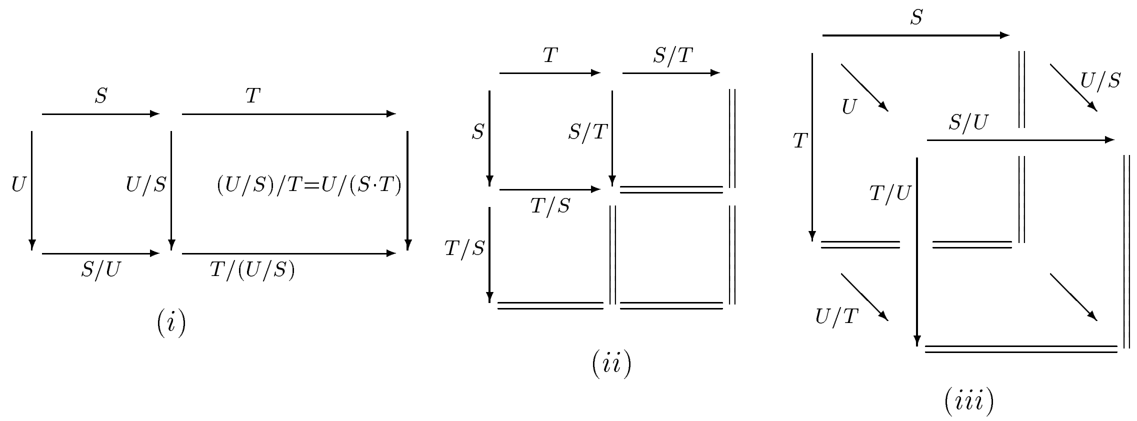

Figure 6. Projection laws.

Proof. Items 1–3 follow immediately from the construction. For item 4, let S and T be two complete developments of a set R of redexes. If the construction of Figure 5 is applied to S and T, every reduction step in the resulting diagram reduces a residual of a member of R. Since S and T are complete developments, the right and bottom sides of the diagram are therefore empty. Therefore $S \sim_L T$. Items 5 and 6 describe the situation of Figure 6(i), and are immediate from the definition of the construction. Item 7 follows from the first and Figure 6(ii). Item 8 is illustrated by Figure 6(iii). □

Finally, we note that the concept of Lévy-equivalence casts further light on the proof of the Strip Lemma.

Proposition 4.12 *In Figure 3, the sequences $t_0 \to t_\alpha \to s_\alpha$ and $t_0 \to s_0 \to s_\alpha$ are Lévy-equivalent.*

Proof. The use of the Strip Lemma in the construction of Figure 5 immediately implies that Figure 3 is a special case of Figure 5. The two sequences are therefore $S \cdot (\mathcal{R}/S)$ and $\mathcal{R} \cdot (S/\mathcal{R})$, which by Proposition 4.11, are equivalent. □

5. The Compressing Lemma

In this section we prove the Compressing Lemma for left-linear TRSs. This states that for any reduction of length greater than ω, there is an equivalent reduction of length at most ω, equivalent in the sense that it not only has the same endpoints, but is Lévy-equivalent.

The lemma reassures us that transfinite strongly converging reductions are just "slow" variants of the strongly converging reductions of length at most ω, and that the power of transfinite reduction does not depend on being able to execute reductions of length greater than ω. These might otherwise be felt to be computationally meaningless, whereas reductions of length ω are naturally interpreted as approaching arbitrarily close to their limit with enough finite computation.

Lemma 5.1 Compressing Lemma. *For any strongly converging reduction of length α in a left-linear TRS, there is a strongly converging reduction of length at most ω with the same* *endpoints. If the TRS is orthogonal, that reduction can be chosen to be Lévy-equivalent to the given reduction.*

PROOF. For finitely long reductions this is trivial. Let the given reduction be $t_0 \twoheadrightarrow_{\omega+\alpha} t_{\omega+\alpha}$ for some α. We shall prove the theorem by induction on α.

- The case $\alpha = 0$ is trivial.
- Let $\alpha = \beta + 1$. By the induction hypothesis applied to β, we can first compress the initial segment $t_0 \twoheadrightarrow_\beta t_\beta$ down to length at most ω. It then suffices to show that any reduction $t_0 \twoheadrightarrow_{\omega+1} t_{\omega+1}$ can be compressed into $t_0 \twoheadrightarrow_{\leq\omega} t_{\omega+1}$, Lévy-equivalent in the case of an orthogonal TRS.

  Let $t_0 \twoheadrightarrow_{\omega+1} t_{\omega+1}$ be a strongly converging reduction. Let the redex R_ω contracted in $t_\omega \to t_{\omega+1}$ have depth d_ω and position u_ω. Put $d = d_\omega + h$, where h is the height of the left-hand side of the rule associated with the redex R_ω. By strong convergence there exists an N such that for $n \geq N$ the depth of the redex R_n contracted in $t_n \to t_{n+1}$ is larger than d. This implies that in every t_n with $n \geq N$, there is a redex at u_ω, by the same rule as R_ω, and that R_ω is the unique residual of each of these redexes by the sequence from t_n to t_ω. Note that left-linearity is essential here, but not orthogonality, since by construction R_ω and R_k are far enough apart that they cannot conflict.

  We will now construct a strongly converging reduction $t_0 \twoheadrightarrow_{\leq\omega} t_{\omega+1}$. For the first N steps we take $t_0 \twoheadrightarrow_N t_N$. Then when $N \leq k \leq \omega$ we reduce $t_k \to s_k$ by contracting the redex R_ω at u in t_k. By construction, the redex R_k reduced in $t_k \to t_{k+1}$ is either disjoint from R_ω or below it.

  Define reductions $s_k \twoheadrightarrow_{\leq\omega} s_{k+1}$ by contracting the copies of R_k resulting from the contraction of R_ω in t_k. It is easy to see that the resulting reduction sequence $s_N \twoheadrightarrow_{\leq\omega} s_{N+1} \twoheadrightarrow_{\leq\omega} s_{N+2} \twoheadrightarrow_{\leq\omega} \cdots$ strongly converges to $s_\omega \equiv t_{\omega+1}$. If the right-hand side of the rule reduced by R_ω is finite, each of the sequences $s_k \twoheadrightarrow_{\leq\omega} s_{k+1}$ must also be finite, and the reduction of s_N to $t_{\omega+1}$ has length at most ω, proving the lemma. However, it is possible that the right-hand side of the rule is infinite. We deal with this possibility by showing that the reduction of s_N to $t_{\omega+1}$ (which might have length up to ω^2) can be reordered so as to bring its length down to at most ω.

  Every step of the reduction of t_N to t_ω takes place at depth at least d. Therefore the reduction proceeds independently in each of the subterms of t_N whose root is at depth exactly d, and it can be seen as an interleaving of a finite set of reduction sequences, one in each such subterm. Consider what the reduction of R_ω in t_N does to each of these subterms. Those not contained in R_ω are left unchanged, but those contained in R_ω may be replaced by any number of copies of themselves, possibly infinitely many. s_N can therefore be reduced to s_ω by applying to each such copy of a subterm of t_N the reduction sequence which that subterm is subject to in the reduction of t_N to t_ω. We thus obtain possibily countably infinitely many reductions, each of length at most ω, of subterms of s_N, such that the combination of all these reductions reduces s_N to s_ω. To obtain a reduction of s_N to s_ω or length at most ω, we need only interleave these reduction sequences.

  It follows directly from this construction that in an orthogonal system, the projection of the original sequence over the compressed sequence is empty, and vice versa. Therefore they are Lévy-equivalent.

$$
\begin{tikzcd}
t_0 \arrow[r, "0{,}0"'] & s_0 \arrow[r, "\leq\omega{,}0"'] \arrow[d, "*{,}0"] & r_0 \arrow[r, "\lambda{,}0"'] \arrow[d, "\mu_1{,}0"] & t_{\omega+\lambda} \arrow[d, equal] \\
& s_1 \arrow[r, "\leq\omega{,}1"'] \arrow[d, "*{,}1"] & r_1 \arrow[r, "\lambda_1{,}1"'] \arrow[d, "\mu_2{,}1"] & t_{\omega+\lambda} \arrow[d, equal] \\
& s_2 \arrow[r, "\leq\omega{,}2"'] & r_2 \arrow[r, "\lambda_2{,}2"'] & t_{\omega+\lambda} \\
& \vdots & \vdots & \vdots \\
& t_{\omega+\lambda} \arrow[r, equal] & t_{\omega+\lambda} \arrow[r, equal] & t_{\omega+\lambda}
\end{tikzcd}
$$

Figure 7. Compressing Lemma, limit case

- Finally, let $\alpha = \lambda$ for some limit ordinal λ. We make the following claim, to be proved later.

  Claim 5.2 (Cf. [FW91].) Given a strongly converging reduction of the form

  $$s_d \rightarrow_{\leq\omega,d} r_d \rightarrow_{\leq\lambda,d} t$$

  there exists also one of the form

  $$s_d \rightarrow_d^* s_{d+1} \rightarrow_{\leq\omega,d+1} r_{d+1} \rightarrow_{\leq\lambda,d+1} t$$

  Continuing the main proof, write the given reduction as $t_0 \rightarrow_{\leq\omega,0} t_\omega \rightarrow_{\lambda,0} t_{\omega+\lambda}$, and let $s_0 = t_0$, $r_0 = t_\omega$, $t = t_{\omega+\lambda}$. We may then repeatedly apply Claim 5.2 so as to construct Figure 7. Clearly the reduction sequence $s_0 \rightarrow_0^* s_1 \rightarrow_1^* s_2 \rightarrow_2^* \cdots$ strongly converges to the limit $t_{\omega+\lambda}$.

  For orthogonal systems, we will show in proving Claim 5.2 that the two routes from top left to bottom right in each of the rectangles in Figure 7 bounded by s_d, s_{d+1}, and $t_{\omega+\lambda}$ are Lévy-equivalent. It follows that for such systems, the sequences $t_0 \rightarrow s_0 \rightarrow_0^* s_1 \rightarrow^\infty t_{\omega+\lambda}$ and $t_0 \rightarrow s_0 \rightarrow_0^* r_0 \rightarrow_{\lambda,0} t_{\omega+\lambda}$ are Lévy equivalent.

  It remains to prove Claim 5.2.

  First we observe that by the definition of strong convergence, for any limit ordinal λ, any strongly converging reduction $t_0 \rightarrow_{\lambda,d} t_\lambda$ can be split into $t_0 \rightarrow_{\mu,d} t_\mu \rightarrow_{\lambda',d+1}$, where $\mu < \lambda$ and λ' is a limit ordinal such that $\lambda' \leq \lambda$.

  Next we construct Figure 8.

  The areas marked by (1) are constructed using this observation. For the top right triangle marked (1), we also take β to be less than λ'. This is possible because λ' is a limit ordinal and the upper side of the triangle is strongly converging. We then use the induction hypothesis to construct the triangle marked (2).

  For each triangle marked (1), the two routes round the triangle are identical reduction sequences. By induction we may assume that in an orthogonal system, the two routes round triangle (2) are Lévy-equivalent. Therefore the two routes round the perimeter of Figure 8 are Lévy equivalent. These are also the two routes round the rectangle in Figure 7 bounded by s_d, s_{d+1}, and $t_{\omega+\lambda}$, which as therefore Lévy equivalent, as required. □

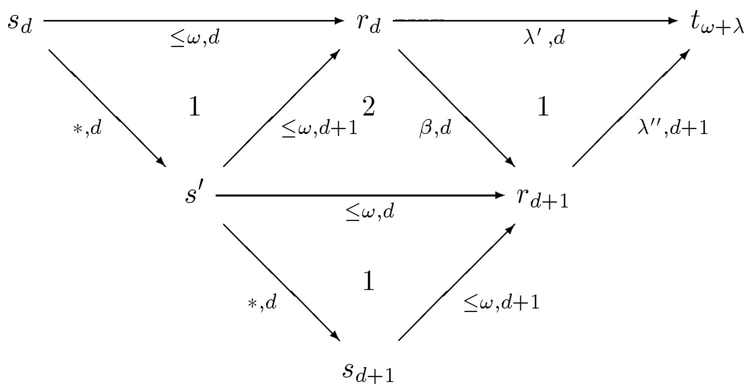

Figure 8.

Note that the condition of left-linearity is essential for this part of compressing lemma, as witnessed by the following counterexample from [DKP91]:

Counterexample 5.3 Let R be the following TRS:

$$\begin{cases} A \to G(A) \\ B \to G(B) \\ F(x,x) \to C \end{cases}$$

The term $F(A,B)$ cannot reduce to C in at most ω many steps: the shortest reduction to C is the strongly converging reduction

$$F(A,B) \to_\omega F(G^\omega, G^\omega) \to C.$$

6. The transfinite Church-Rosser property

The Church-Rosser property with respect to finite reductions holds for infinitary orthogonal TRSs with finite right-hand sides: the usual proofs for the finitary case go through verbatim. We will see that it does not necessarily hold when there are infinite right-hand sides. One can also formulate a version of the Church-Rosser property for transfinite reductions.

Definition 6.1 Let R be an infinitary TRS. R has the *transfinite Church-Rosser* property if For any term t in R and any reductions $t \to_\alpha s$ and $t \to_\beta r$, there exist a term u and reductions $s \to_\gamma u$ and $r \to_\delta u$ (cf. Figure 9.i).

We shall in this section demonstrate to what extent this property holds for orthogonal infinitary TRSs. For such systems, the compression lemma implies that it is equivalent to the following version:

- for any term t and any reductions $t \to_{\leq\omega} s$ and $t \to_{\leq\omega} r$ there exist a term u and reductions $s \to_{\leq\omega} u$ and $r \to_{\leq\omega} u$ (cf. Figure 9.ii).

$$
\begin{tikzcd}
& t \arrow[dl, "\alpha"'] \arrow[dr, "\beta"] & \\
s \arrow[dr, "\gamma"'] & & r \arrow[dl, "\delta"] \\
& u &
\end{tikzcd}
\qquad
\begin{tikzcd}
& t \arrow[dl, "\leq\omega"'] \arrow[dr, "\leq\omega"] & \\
s \arrow[dr, "\leq\omega"'] & & r \arrow[dl, "\leq\omega"] \\
& u &
\end{tikzcd}
\qquad
\begin{tikzcd}
& t \arrow[dl, "\omega"'] \arrow[dr, "\omega"] & \\
s \arrow[dr, "\leq\omega"'] & & r \arrow[dl, "\leq\omega"] \\
& u &
\end{tikzcd}
$$

(i) (ii) (iii)

Figure 9.

The Strip Lemma implies that this already holds when at least one of the two sequences is finitely long. For orthogonal TRSs, the transfinite CR property is therefore equivalent to:

- for any term t and any reductions $t \rightarrow_\omega s$ and $t \rightarrow_\omega r$ there exist a term u and reductions $s \rightarrow_{\leq\omega} u$ and $r \rightarrow_{\leq\omega} u$ (cf. Figure 9.iii).

In contrast to the finitary case, the Church-Rosser property for transfinite reductions in orthogonal TRSs is in general false, for both strongly and Cauchy converging reductions. Even the finite Church-Rosser property can fail. We noted above that it holds when all right-hand sides are finite. We also noted that the Strip Lemma implies that a pair of finite coinitial reduction sequences can always be joined; however, with infinite right-hand sides it is sometimes only possible to join them by infinite sequences. The following counterexamples demonstrate these claims.

Counterexample 6.2

- Rules: $A(x) \rightarrow x$, $B(x) \rightarrow x$, $C \rightarrow A(B(x))$

  Sequences: $C \rightarrow A(B(C)) \rightarrow A(C) \rightarrow A(A(B(C))) \rightarrow A(A(C)) \rightarrow_\omega A^\omega$

  $C \rightarrow A(B(C)) \rightarrow B(C) \rightarrow B(A(B(C))) \rightarrow B(B(C)) \rightarrow_\omega B^\omega$

  Hence $C \rightarrow_{\leq\omega} A^\omega$ as well as $C \rightarrow_{\leq\omega} B^\omega$. But there is no term t such that $A^\omega \rightarrow_{\leq\omega} t \leftarrow_{\leq\omega} B^\omega$ (whether by strongly or Cauchy converging reduction). (This example also demonstrates that in general ω-normal forms are not unique, for both notions of convergence.)

- Rules: $D(x, y) \rightarrow y$, $C \rightarrow D(A, D(B, C))$

  Sequences:

  $C \rightarrow D(A, D(B, C)) \rightarrow D(A, C) \rightarrow_2 D(A, D(A, C)) \rightarrow_2 D(A, D(A, D(A, C))) \rightarrow \ldots$

  $C \rightarrow D(A, D(B, C)) \rightarrow D(B, C) \rightarrow_2 D(B, D(B, C)) \rightarrow_2 D(B, D(B, D(B, C))) \rightarrow \ldots$

  It is not possible to join the limits of these two sequences.

- Rules: $A \rightarrow B$, $F(x) \rightarrow G(x, G(x, G(x, \ldots)))$.

  Sequences:

  $F(A) \rightarrow G(A, G(A, G(A, \ldots)))$

  $F(A) \rightarrow F(B)$

The terms $G(A, G(A, G(A, \ldots)))$ and $F(B)$ can both be reduced to $G(B, G(B, G(B, \ldots)))$, but not by a pair of finite reduction sequences.

In the first two counterexamples, we only include the last rewrite rule in order to have examples in which the reduction sequences start from finite terms. As with Example 3.8, we can in addition arrange that none of the right-hand sides of the rules contain redexes. In the first example, replace the third rule by $C(D, x) \to A(B(C(x, x)))$, and consider sequences starting from $C(D, D)$. More interestingly, we can exhibit the phenomenon within combinatory logic (CL), which is the TRS having nullary symbols S, K, and binary application (which we will indicate, using the usual convention, by left-associative juxtaposition), and rules $Sxyz \to xz(yz)$ and $Kxy \to x$. A counterexample to the transfinite Church-Rosser property for this system will be constructed following Theorem 6.10.

We shall, however, prove several slightly weaker forms of the Church-Rosser property for orthogonal systems. One approach will be to restrict the form of rules. The above counterexamples suggest that collapsing rules are the source of the problem, and this is indeed the case. We shall establish restrictions on the form of collapsing rules which will restore the Church-Rosser property. Another approach is based on the notion that the terms in the above counterexamples are in an intuitive sense meaningless. They not only have no normal forms, they have no head normal forms, in the sense that every term they can be reduced to can be further reduced to a redex. In an orthogonal system, such a term cannot "make a difference" to any surrounding context (a notion which will be formalised later). If we identify together all such "meaningless" terms, we find that the Church-Rosser property is restored for all orthogonal systems.

6.1 Depth-preserving orthogonal Term Rewriting Systems

In this section and the next we consider two classes of orthogonal TRS in which the infinitary Church-Rosser property holds for strongly converging sequences. In this section we will consider a very restrictive condition on the form of rewrite rules and prove the Church-Rosser property for such systems. This will serve as a stepping-stone to proving a more general form of the property which requires less restrictive conditions.

Definition 6.3 A *depth-preserving* TRS is a left-linear TRS such that for all rules the depth of any variable in a right-hand side is greater than or equal to the depth of the same variable in the corresponding left-hand side.

Theorem 6.4 *A depth-preserving orthogonal TRS has the infinitary Church-Rosser property for strongly converging sequences.*

Proof. Let $t_{0,0} \to t_{0,1} \to \ldots \to_\omega t_{0,\omega}$ and $t_{0,0} \to t_{1,0} \to \ldots \to_\omega t_{\omega,0}$ be strongly convergent.

Construct the projection diagram for the two sequences as in Figure 10. The Strip Lemma suffices to construct all of this diagram save for the bottom-right-hand corner. We must prove that the right-hand side and the bottom side of the diagram are strongly convergent and have the same limit. Each of the other rows and columns of the figure is strongly convergent. By the depth-preserving property it holds for all $m < \omega$ and $n \le \omega$ that the depth of the redexes reduced in $t_{n,m} \to_\omega t_{n,m+1}$, being the residuals of the redex $R_{0,m}$ in $t_{0,m} \to t_{0,m+1}$, is at least the depth of $R_{0,m}$ itself. Because $t_{0,0} \to t_{0,1} \to \ldots \to_\omega t_{0,\omega}$ is strongly convergent, it follows that $t_{\omega,0} \to_{\le\omega} t_{\omega,1} \to_{\le\omega} t_{\omega,2} \ldots$ is strongly converging. Let us call its limit $t_{\omega,\omega}$.

$$\begin{tikzcd}
t_{0,0} \arrow[r] \arrow[d] & t_{0,1} \arrow[r] \arrow[d, "\leq\omega"] & \cdots & t_{0,\omega} \arrow[d, "\leq\omega"] \\
t_{1,0} \arrow[r, "\leq\omega"] \arrow[d] & t_{1,1} \arrow[r, "\leq\omega"] \arrow[d, "\leq\omega"] & \cdots & t_{1,\omega} \arrow[d, "\leq\omega"] \\
\vdots & \vdots & & \vdots \\
t_{\omega,0} \arrow[r, "\leq\omega"'] & t_{\omega,1} \arrow[r, "\leq\omega"'] & \cdots &
\end{tikzcd}$$

Figure 10.

In the same way the terms $t_{n,\omega}$ are part of a strongly converging sequence. The limit of this sequence is also equal to $t_{\omega,\omega}$, as can be seen with the following argument.

Let $\epsilon > 0$. Because $(t_{\omega,n})_{n<\omega}$ is a Cauchy sequence, there is an N_1 such that if $N_1 \leq m < \omega$ then $d(t_{\omega,m}, t_{\omega,\omega}) < \frac{1}{3}\epsilon$.

Because $t_{0,0} \twoheadrightarrow_{\leq\omega} t_{\omega,0}$ is strongly converging, there is an N_2 such that for $n \geq N_2$ we have that $2^{-d_n} < \frac{1}{3}\epsilon$ where d_n is the depth of the redex R_n reduced at step $t_{n,0} \rightarrow t_{n+1,0}$. Since the residuals of this redex R_n occur at least at the same depth, and since the TRS R is depth-preserving, we get $d(t_{n,m}, t_{\omega,m}) < \frac{1}{3}\epsilon$ for all $m < \omega$ and $n \geq N_2$.

For similar reasons there is an N_3 such that if $n < \omega$ and $N_3 \leq m < \omega$ then $d(t_{n,\omega}, t_{n,m}) < \frac{1}{3}\epsilon$.

Therefore taking N to be the maximum of N_1, N_2 and N_3, for $n \geq N$ we find by the triangle inequality that

$$\begin{aligned}
d(t_{n,\omega}, t_{\omega,\omega}) &\leq d(t_{n,\omega}, t_{n,n}) + d(t_{n,n}, t_{\omega,n}) + d(t_{\omega,n}, t_{\omega,\omega}) \\
&\leq \frac{1}{3}\epsilon + \frac{1}{3}\epsilon + \frac{1}{3}\epsilon \\
&\leq \epsilon.
\end{aligned}$$

□

6.2 Almost non-collapsing orthogonal Term Rewriting Systems

DEFINITION 6.5

- A rewrite rule is a collapsing rule if its right hand side is a variable. The *arity* of a collapsing rule is the number of different variables that occur in its left hand side.
- A TRS R is *non-collapsing* if all its rewrite rules are non-collapsing.
- A TRS is *almost non-collapsing* if all its rewrite rules are non-collapsing, but for at most a single collapsing rule, with the left-hand side of that rule containing no variables other than the one which is its right-hand side.

We will show that an orthogonal TRS satisfies the Church-Rosser property for strongly converging reductions if and only if it is almost non-collapsing. The proofs will use a variant of Park's notion of hiaton (cf. [Par83]) in order to transform every almost non-collapsing system into a depth-preserving one, and deriving the Church-Rosser property for the given system from that for the transformed system.

The basic idea is to replace a depth-decreasing rule like $A(x, B(y)) \to C(y)$ by a depth-preserving variant $A(x, B(y)) \to C(\epsilon(y))$. In order to keep the rewrite rules applicable to terms involving hiatons, we also have to add more variants such as $A(x, \epsilon^m(B(y))) \to C(\epsilon^{m+1}(y))$ for $m > 0$. By adding to an arbitrary TRS enough depth-preserving variants of its rewrite rules, we transform it into a depth-preserving TRS.

Definition 6.6 Let R be a TRS based on the alphabet Σ. Let Σ_ϵ be the extension of Σ with a fresh unary symbol ϵ.

1. Let the ϵ-hiding function $\rho : Ter^\infty(\Sigma_\epsilon) \to Ter^\infty(\Sigma)$ be partially defined by induction as follows:

   (a) $\rho(x) = x$,
   (b) $\rho(f(t_1, \ldots, t_n)) = f(\rho(t_1), \ldots, \rho(t_n))$ for f in Σ and $t_i \in Ter^\infty(\Sigma_\epsilon)$ for $0 \leq i \leq n$,
   (c) $\rho(\epsilon(t)) = \rho(t)$ for $t \in Ter^\infty(\Sigma_\epsilon)$.

   ρ is well-defined on terms in $Ter^\infty(\Sigma_\epsilon)$ containing no infinite string of ϵs.

2. A term $t \in Ter^\infty(\Sigma_\epsilon)$ is an *ϵ-variant* of a term $s \in Ter^\infty(\Sigma)$ if $\rho(t) = s$, that is, if hiding the ϵs in t results in s.

3. An ϵ-variant of a rule $l \to r$ is a pair of terms (l_ϵ, r_ϵ) such that

   (a) $\rho(l_\epsilon) = l$.
   (b) $\rho(r_\epsilon) = r$.
   (c) the root symbol of l_ϵ is not ϵ.
   (d) l_ϵ does not contain a subterm of the form $\epsilon(x)$ for any variable x.
   (e) the root symbol of r_ϵ is not ϵ unless r is a variable,

4. An ϵ-completion R^ϵ of R has alphabet Σ_ϵ. Its rules are depth-preserving ϵ-variants of rules of R, chosen thus: for each rule $l \to r$ of R, and each ϵ-variant l_ϵ of l satisfying the conditions of item 3, R^ϵ includes exactly one depth-preserving ϵ-variant of the rule having l_ϵ as its left-hand side. We denote reduction in R^ϵ by $\to^\epsilon$. Note that R^ϵ is not uniquely defined; the precise choice of R^ϵ will not be significant.

The following lemma is immediate.

Lemma 6.7 *The ϵ-completion of an orthogonal TRS is depth-preserving and orthogonal.* □

Lemma 6.8 *Let R be a non-collapsing orthogonal TRS.*

1. *Let t_ϵ be an ϵ-variant of a term t of R. If t_ϵ strongly ϵ-converges in ω steps to some term s in R^ϵ, then s does not contain a branch ending in an infinite string of ϵs.*

2. *Let t_0 be an ϵ-variant of some term s_0. If $t_0 \to^\epsilon_\omega t_\omega$ is a strongly converging reduction in R^ϵ, then so is $s_0 \to_\omega s_\omega$ in R, where $s_i = \rho(t_i)$ for $0 \leq i \leq \omega$.*

3. *Let $t_0 \to_\omega t_\omega$ be a strongly converging reduction in R. Let s_0 be an ϵ-variant of t_0. Then there exists a strongly converging reduction $s_0 \to^\epsilon_\omega s_\omega$ in R^ϵ such that each s_i is an ϵ-variant of the corresponding t_i, and similarly for the reduction rules used.*

PROOF.

1. Since there are no collapsing rules, a string of ϵs can only be made longer by a reduction occurring at the parent node of its topmost ϵ. Strong convergence implies that only finitely many such reductions can be made, and therefore that an infinite string of ϵs cannot be created.

2. Since t_0 is an ϵ-variant it does not contain an infinite string of ϵs. Neither do any of the t_i for $i \in \omega$, nor t_ω itself by item 1. Hence, $\rho(t_n)$ is well-defined for all $0 \leq n \leq \omega$.

   Because there are no infinite strings of ϵs in t_ω, every infinite path from the root of t_ω must contain infinitely many occurrences of members of Σ. Note also that t_ω is necessarily an infinite term.

   Since by 1, t_ω contains no infinite string of ϵs, it must contain occurrences of members of Σ at arbitrarily great depth.

   Given any finite number k, consider those occurrences v of t_ω, such that the path from the root to v contains at least k occurrences of symbols in Σ. By the preceding remarks, there must be at least one such occurrence. Let N_k be the minimum length of all such v. Because there are no infinite strings of ϵs, N_k must tend to infinity with k. Since $t_0 \rightarrow_\omega t_\omega$ is strongly converging there exists for any $k > 0$ an N such that for $n > N$, the depth of the redex reduced in $t_{n-1} \rightarrow t_n$ is at least N_k. This implies that the corresponding redex in $s_{n-1} \rightarrow s_n$ is at depth at least k, and hence $s_0 \rightarrow_\omega s_\omega$ is strongly convergent.

3. The ϵ-variant s_0 of t_0 contains the corresponding ϵ-variant of the redex reduced in t_0. Apply an ϵ-variant of the corresponding rule. The resulting reduction satisfies the required properties. □

THEOREM 6.9 *Every non-collapsing orthogonal TRS satisfies the infinitary Church-Rosser property for strongly converging reductions.*

PROOF. Let R be an orthogonal TRS. Construct its ϵ-completion R^ϵ. By Theorem 6.4 the depth-preserving orthogonal TRS R^ϵ satisfies the infinitary Church-Rosser property. So if we start with two strongly converging reductions $t \rightarrow_{\leq\omega} s_1$ and $t \rightarrow_{\leq\omega} s_2$, then by Lemma 6.8(3) these reductions lift to two strongly converging reductions in R^ϵ, let us say $t \rightarrow^\epsilon_{\leq\omega} r_1$ and $t \rightarrow^\epsilon_{\leq\omega} r_2$. By Theorem 6.4 there exists a join u for the two lifted reductions such that $r_1 \rightarrow^\epsilon_{\leq\omega} u$ as well as $r_2 \rightarrow^\epsilon_{\leq\omega} u$. Erasing all ϵs using Lemma 6.8(2) we see that the term $\rho(u)$ is the join in R of $t \rightarrow_{\leq\omega} s_1$ and $t \rightarrow_{\leq\omega} s_2$.

THEOREM 6.10 *An orthogonal TRS satisfies the infinitary Church-Rosser property for strongly converging reductions if and only if it is almost non-collapsing.*

PROOF.

If: If the TRS has no collapsing rule then it is Church-Rosser by theorem 6.9. Otherwise, let its collapsing rule be $C[x] \rightarrow x$ for some context $C[\,]$ containing no free variables. First, note that the proof of the previous theorem cannot be directly applied in the presence of the rule $C[x] \rightarrow x$. Consider the rules $A(x) \rightarrow I(x)$, $B(x) \rightarrow I(x)$, $I(x) \rightarrow x$. There are obvious reductions of the term $A(B(A(B(\cdots))))$ to both A^ω and B^ω. These lift to reductions ending with $A(\epsilon(A(\epsilon(\cdots))))$ and $\epsilon(B(\epsilon(B(\cdots))))$ respectively. If we now apply the Church-Rosser property of the depth-balanced system, we obtain reductions of these terms to $\epsilon(\epsilon(\epsilon(\epsilon(\cdots))))$, which cannot be lifted to strongly convergent reductions in the original system.

A simple modification of the previous proof establishes the present theorem. We modify the depth-preserving transformation by introducing two versions of ϵ: ϵ itself, and ϵ'. The rule $C[x] \to x$ is replaced by the depth-preserving version $C[x] \to \epsilon'^n(x)$, where n is the depth of the hole in $C[\,]$. The other rules are transformed as before, except that wherever ϵ would appear on the left-hand side in the original transformation, either ϵ or ϵ' is used, in all possible combinations. On the right-hand sides, only ϵ is used. It is easy to see that the resulting system is depth-preserving and orthogonal, and hence that the infinite Church-Rosser property holds.

The distinction between ϵ and ϵ' can be thought of as labelling those occurrences of ϵ which arise from reductions by $C[x] \to \epsilon'^n(x)$.

Now consider two strongly converging reductions $t \to_{\leq\omega} s_1$ and $t \to_{\leq\omega} s_2$. As in the proof of the previous theorem, we obtain in R^ϵ a term u and two strongly converging reductions $r_1 \to^\epsilon_{\leq\omega} u$ and $r_2 \to^\epsilon_{\leq\omega} u$, where r_1 and r_2 are ϵ-variants of s_1 and s_2.

We cannot in general erase every ϵ and ϵ' from these sequences to obtain a join for s_1 and s_2, since u may contain infinite branches consisting only of ϵ and ϵ'. (which we shall call ϵ-branches for short). But we will show that we can transform these sequences in such a way as to eliminate such branches, after which the erasing process can be performed safely.

In every ϵ-branch in u, there must be infinitely many occurrences of ϵ'. This follows for the same reason that in the non-collapsing case, no infinite branch of ϵs can arise.

Now consider an occurrence of ϵ' in an ϵ-branch of u. This must arise from a reduction by the rule $C[x] \to \epsilon'^n(x)$ at some point in each of the sequences $r_1 \to^\epsilon_{\leq\omega} u$ and $r_2 \to^\epsilon_{\leq\omega} u$. This reduction is performed on a subterm of the form $C[T]$, where T reduces to an ϵ-branch. By orthogonality, it is impossible for the reduction of the $C[T]$ redex to be necessary for any later step of the sequence to be possible. If we omit it, the only effect is that certain subterms of the form of $\epsilon'^n(\ldots)$ later in the sequence will be replaced by $C[\ldots]$.

We therefore omit from both $r_1 \to^\epsilon_{\leq\omega} u$ and $r_2 \to^\epsilon_{\leq\omega} u$ every reduction by $C[x] \to \epsilon'^n(x)$ which gives rise to an occurrence of ϵ' in any ϵ-branch of u. This gives a term u' containing no such occurrences of ϵ', and reduction sequences $r_1 \to^\epsilon_{\leq\omega} u'$ and $r_2 \to^\epsilon_{\leq\omega} u'$. These sequences have the property that they contain no ϵ-branch anywhere. They may therefore be lifted to strongly convergent reductions in the original system, providing a strongly convergent joining of the original reduction sequences.

Only if: If the given TRS is not almost non-collapsing, then it contains either two unary collapsing rules, or an n-ary collapsing rule with $n \geq 2$. In either case one can construct a counterexample similar to those of Counterexample 6.2. □

From the argument for the "only if" case of the theorem we can easily construct a counterexample to the Church-Rosser property for CL. Consider the two collapsing redexes KxK and KxS. It is enough to find a term A of CL which reduces to $K(KAK)S$. Then A reduces both to $K(K(K \ldots K)K)K$ and to $K(K(K \ldots S)S)S$, which have no common reduct. Such an A can be constructed by means of the following definitions:

$I = SKK$
$B = S(KS)K$
$C = S(BBS)(KK)$
$Z = B(SI)(SII)$
$T = ZZ$
$A = T(C(BK(CKK))S)$

We then find that $A \to^* K(KAK)S$. Note that T is Turing's fixed point operator, with the property that $Tf \to^* f(Tf)$, and that $C(BK(CKK))S$ has the property that $C(BK(CKK))Sx \to^* K(KxK)S$.

7. Reductions to normal form

In the previous section we showed that the infinitary Church-Rosser property holds for almost non-collapsing orthogonal TRSs. In this section we will show that various infinitary normal form properties hold for arbitrary orthogonal TRSs.

7.1 Unique normal form properties

Definition 7.1

- A TRS has the *normal form property* (NF) if for any term t with a strongly converging reduction to normal form n and any strongly converging reduction from t to s, there exists a strongly converging reduction from s to n:

$$\begin{tikzcd} t \arrow[rr, "\alpha"] \arrow[dr, "\beta"'] & & n \\ & s \arrow[ur, dotted, "\gamma"'] & \end{tikzcd}$$

- A TRS has *unique normal forms* (UN) if convertible normal forms are identical, where t is *convertible with* t' if there is a finite series of forwards and reverse reduction sequences of the form $t \to_\infty t_1 \leftarrow_\infty t_2 \ldots \to_\infty t_n \leftarrow_\infty t'$. (That is, convertibility, for infinitary TRSs, is the reflexive transitive symmetric closure of $\to_\infty$.)

- A TRS has *unique normal forms with respect to reduction* ($\mathrm{UN}^\to$) if all normal forms of a term are identical.

$$\begin{tikzcd} t \arrow[d, shift right, "\alpha"'] \arrow[d, shift left, "\beta"] \\ n_1 \equiv n_2 \end{tikzcd}$$

Note that in the finitary case, the definition of NF could be equivalently stated by taking β to be 1, instead of an arbitrary finite number. Call the property so defined NF^1. In the infinitary setting, NF^1 is weaker than NF, as the following counterexample shows.

Consider the (non-orthogonal) rule system $A(B(x)) \to B(x)$ and $A(B(x)) \to A(A(B(x)))$. In the definition of NF, take t to be $A(B(C))$ and n to be $B(C)$. Then for every finite reduction from t to a term s, there is a finite reduction of s to n, and so the system satisfies NF^1. However, t can be reduced to the infinite term A^ω, which cannot be reduced to $B(C)$. Thus the system does not satisfy NF. Similarly, we may consider the property UN^1, where one of t_i and t_{i+1} must be reducible to the other in a single step. The same example satisfies UN^1 but not UN, since A^ω is a normal form.

The NF property is sometimes stated in a form that allows t to be convertible with s rather than reducible to it; this is equivalent to the definition given here, as is easily seen by an obvious induction on the (finite) number of forward and backward reduction sequences in the conversion of t to s.

As in the finitary case, certain relations among the above properties and CR are immediate.

LEMMA 7.2 *CR* $\Rightarrow$ *NF* $\Rightarrow$ *UN* $\Rightarrow$ $UN^{\rightarrow}$ □

7.2 Reduction modulo equivalence of hypercollapsing terms

In this section, we prove a weakened version of the Church-Rosser theorem. Unlike the earlier Theorem 6.10, this one applies to all orthogonal TRSs, but only requires that the two given sequences be joinable “up to” a certain equivalence relation on terms.

DEFINITION 7.3

- A *hyper-collapsing* reduction sequence is a strongly continuous reduction sequence containing infinitely many collapsing reductions performed at the root.
- A *hyper-collapsing* term is a term from which there is a hyper-collapsing reduction sequence.
- $t \sim_{hc} t'$ (in words: t and t' are *hc-equivalent*) if and only if there is a term t'' and substitutions σ and σ' such that $t = \sigma(t'')$, $t' = \sigma'(t'')$, and σ and σ' are defined on the same set of variables and map them to hyper-collapsing terms.
- A TRS is CR_{hc} if whenever $t_0 \rightarrow_\alpha t_1$ and $t_0 \rightarrow_\beta t_2$, there exist sequences $t_1 \rightarrow_\gamma t_3$ and $t_2 \rightarrow_\delta t_3'$ such that $t_3 \sim_{hc} t_3'$.

THEOREM 7.4 *Every orthogonal TRS is* CR_{hc}.

The proof of this theorem occupies the remainder of the section. We proceed by transforming the given TRS into an almost non-collapsing one, applying the Church-Rosser property for the transformed system, and transforming the resulting reduction sequences back into the reduction sequences required by the theorem.

DEFINITION 7.5 Let R be a TRS. R^I is the almost non-collapsing TRS obtained from R by adding a new unary function symbol I and a rule $I(x) \rightarrow x$, and replacing every collapsing rule $l \rightarrow x$ (where l is a term and x is a variable) by $l \rightarrow I(x)$.

When considering reductions and the equivalence $\sim_{hc}$, we will when necessary indicate by superscripts whether we intend them in R or in R^I. Note that *prima facie*, $\sim_{hc}^R$ and $\sim_{hc}^{R^I}$ are not obviously the same, although we shall later prove them identical. R (and therefore also R^I) will be assumed to be orthogonal.

DEFINITION 7.6 Let t be a term of R^I. If I^ω is not a subterm of t, then t^I is defined to be the result of a complete development of all I-redexes of t.

LEMMA 7.7 *If* $t \rightarrow_\infty^R t'$, *then* $t \rightarrow_\infty^{R^I} t'$.

PROOF. In $t \rightarrow^R_\infty t'$, replace each collapsing step $s \rightarrow s'$ which uses a rule $l \rightarrow x$ by a two-step reduction $s \rightarrow s'' \rightarrow s'$ using the R^I-rules $l \rightarrow I(x)$ and $I(x) \rightarrow x$. Each of these steps takes place at the same depth as $s \rightarrow s'$ did, and so the resulting sequence is still strongly converging. □

LEMMA 7.8 *If t is hyper-collapsing, then there is a hyper-collapsing reduction sequence starting from t of length ω.*

PROOF. Let S be a hyper-collapsing strongly continuous reduction sequence starting from t. S must have the form $S_1 r S_2$, where r is the first root-collapsing step. Apply the Compressing Lemma to S_1, to obtain a sequence S_3 of length at most ω with the same endpoints. By strong convergence, S_3 must have the form $S_4 S_5$, where S_4 is finite and r is the unique residual by S_5 of a root redex r'. Apply the Strip Lemma to r' and S_5 to obtain a sequence $S_6 = S_5/r'$. This transforms the original sequence S into the form $S_4 r' S_6 S_2$, in which S_4 is finite, r' is a root-collapsing reduction, and $S_6 S_2$ is hyper-collapsing (because S_2 is a final segment of S). Repeating the construction on $S_6 S_2$ generates a hyper-collapsing sequence starting from t of length ω. □

LEMMA 7.9 *If t is hyper-collapsing and $t \rightarrow_\infty s$, then s is hyper-collapsing.*

PROOF. By the Compressing Lemma, we may assume that $t \rightarrow_\infty s$ has length at most ω. Choose any strongly continuous hyper-collapsing reduction S starting from t. By Lemma 7.8, S can be chosen to be of length ω.

First consider the case where $t \rightarrow s$ is a single step. Apply the Strip Lemma to $t \rightarrow s$ and each step of S, to generate a reduction sequence $S' = S/(t \rightarrow s)$ from s. If $t \rightarrow s$ is a root reduction, then it will cancel out the first root reduction of S, and hence the final segment of S after that point will also be a final segment of S'. This implies that S' is hyper-collapsing. If, instead, $t \rightarrow s$ is not a root reduction, then its residuals by any initial segment of S must be either a single root redex, or a set of non-root redexes. If the former case ever happens, then by the same argument, S' must contain a final segment of S, and hence be hyper-collapsing. If only the latter case happens, then each root reduction of S must be projected into a root reduction of S' by the same rule. In particular, the root-collapsing reductions of S project to root-collapsing reductions of S'. Hence S' is hyper-collapsing in this case also.

By iterating the above argument, the lemma holds for any finite reduction $t \rightarrow^* s$.

Now consider the case $t \rightarrow_\omega s$. By strong convergence, this sequence contains only finitely many root reductions. We can apply the case already established to obtain a hyper-collapsing sequence starting from a point in $t \rightarrow_\omega s$ after the last root reduction. Otherwise put, we may assume without loss of generality that $t \rightarrow_\omega s$ contains no root reductions.

Apply the Strip Lemma to $t \rightarrow_\omega s$ and each step of S, obtaining a sequence S' starting from s. Because $t \rightarrow_\omega s$ contains no root reductions, an argument similar to that used in the first part of the proof shows that every root reduction of S up to and including the first root-collapsing reduction must be projected to a root reduction of S', and the root-collapsing reduction must project to a root-collapsing reduction. See Figure 11, in which the annotations 'r', 'rc', and '$\neg r$' indicate root reductions, root-collapsing reductions, and non-root reductions respectively. The projection of $t \rightarrow_\omega s$ over that initial segment may contain root reductions. However, we can now repeat the argument for that sequence and the remainder of S. In this way, we generate a hyper-collapsing sequence from s. □

$$\begin{tikzcd}
t \arrow[r, "*"] \arrow[d, "\omega"', "\neg r"] & {} \arrow[r, "r"] \arrow[d, "\leq\omega"', "\neg r"] & {} \arrow[r, "*"] \arrow[d, "\leq\omega"', "\neg r"] & {} \arrow[r, "r"] \arrow[d, "\leq\omega"', "\neg r"] & \cdots \arrow[r, "*"] & {} \arrow[r, "rc"] \arrow[d, "\leq\omega"', "\neg r"] & {} \arrow[r] \arrow[d, "\leq\omega"'] & \cdots \\
s \arrow[r, "\leq\omega"] & {} \arrow[r, "r"] & {} \arrow[r, "\leq\omega"] & {} \arrow[r, "r"] & \cdots \arrow[r, "\leq\omega"] & {} \arrow[r, "rc"] & {} \arrow[r] & \cdots
\end{tikzcd}$$

Figure 11.

LEMMA 7.10 *In any orthogonal TRS, if $t \to_\infty t'$, then there is a reduction $t \to_\infty t''$ containing no reduction steps in hyper-collapsing subterms, such that $t'' \sim_{hc} t'$.*

PROOF. Without loss of generality we may assume that $t \to_\infty t'$ has length at most ω. Write the sequence as $t_0 \to t_1 \to \ldots t_{\leq\omega}$. We will define a sequence $s_0 \to_{\leq 1} s_1 \to_{\leq 1} \ldots s_{\leq\omega}$ which performs no reductions in hyper-collapsing subterms. Take $s_0 = t_0$. Assume s_n has been defined and that $t_n \sim_{hc} s_n$.

If $t_n \to t_{n+1}$ is a reduction in a hyper-collapsing subterm of t_N, then take $s_{n+1} = s_n$. Clearly, $t_{n+1} \sim_{hc} s_{n+1}$.

Otherwise, since $t_n \sim_{hc} s_n$, there is a term r and substitutions σ and σ' such that $t_n = \sigma(r)$ and $s_n = \sigma'(r)$, where σ and σ' map the same set of variables to hyper-collapsing terms. The reduction of t_n to t_{n+1} must be at a position of t_n which is also the position of a redex in r. Suppose reducing this redex in r yields r'. Then $t_{n+1} = \sigma(r')$ and $s_n \to \sigma'(r')$. Take $s_{n+1} = \sigma'(r')$; then $t_{n+1} \sim_{hc} s_{n+1}$.

If $t_0 \to_{\leq\omega} t_{\leq\omega}$ is finitely long, the above proves the lemma. If it has length ω, apply the above argument to each step of $t_0 \to_\omega t_\omega$. Every nonempty step $s_n \to s_{n+1}$ is at the same depth as $t_n \to t_{n+1}$, therefore the sequence from s_0 is strongly convergent. We must show that $t_\omega \sim_{hc} s_\omega$.

Consider an outermost hyper-collapsing subterm of t_ω at position u. By strong convergence, from some t_n onwards, no reduction takes place at any position which is at a prefix of u. Therefore $t_n|u$ is also hyper-collapsing, since it reduces to $t_\omega|u$. Since $t_n \sim_{hc} s_n$, $s_n|u$ is hyper-collapsing for all $m > n$. Since $s_n|u$ reduces to $s_\omega|u$, by Lemma 7.9 $s_\omega|u$ is hyper-collapsing. Symmetrically, t_ω has an outermost hyper-collapsing subterm everywhere that s_ω has. Since t_ω and s_ω cannot differ anywhere outside their maximal hyper-collapsing subterms, $t_\omega \sim_{hc} s_\omega$. □

LEMMA 7.11 *Consider a reduction sequence $t \to_\infty^{R^I} t'$ in which no term contains I^ω as a subterm. Then there is a reduction sequence $t^I \to_\infty^R t'^I$.*

PROOF. Consider first a single step $t \to^{R^I} t'$ where neither t nor t' contains I^ω as a subterm. If this step is an I-reduction, then $t^I = t'^I$, and the theorem is satisfied by the empty reduction sequence. If it is a reduction by a rule of the form $l \to I(x)$, then t^I reduces to t'^I by the rule $l \to x$. Otherwise, t^I reduces to t'^I by the same rule that reduced t to t'.

If we apply this to every step of a sequence $t \to_\infty^{R^I} t'$ of arbitrary length, we obtain a reduction sequence of R from t^I to t'^I. It remains to show that this sequence is strongly convergent.

Let λ be a limit ordinal not exceeding the length of the given sequence. t_λ does not contain I^ω as a subterm. By König's lemma it follows that for every d there is an e such that for every position u of t_λ of depth at least e, the number of occurrences of function symbols other than I at positions which are prefixes of u is at least d. This implies that $t_\lambda^I|d = (t_\lambda|e)^I|d$. By strong convergence, there is an n such that when all $n \leq m < \lambda$, $t_m|max(d,e) = t_\lambda|max(d,e)$ and the reduction $t_m \to t_{m+1}$ is at a position deeper than $max(d,e)$. Therefore for such m, $t_m^I|d = t_\lambda^I|d$, and the reduction of t^I_m to t^I_{m+1} is at a depth at least d, demonstrating strong convergence. □

LEMMA 7.12 *A term of R is hyper-collapsing in R if and only if it is hyper-collapsing in R^I.*

PROOF. The forwards implication is immediate from Lemma 7.7. For the converse, let t be hyper-collapsing in R^I. Applying Lemma 7.11 to a hyper-collapsing sequence in R^I from t gives a hyper-collapsing sequence in R from t. □

LEMMA 7.13 *If t and t' are terms of R, then $t \sim^R_{hc} t'$ if and only if $t \sim^{R^I}_{hc} t'$.*

PROOF. Immediate from Lemma 7.12. □

LEMMA 7.14 *Let t and t' be terms of R^I which do not contain I^ω as a subterm. If $t \sim^{R^I}_{hc} t'$, then $t^I \sim^R_{hc} t'^I$.*

PROOF. Let t and t' reduce to s and s' respectively by complete development of all I-redexes which are not contained in hyper-collapsing subterms. (Note that these redexes must be at the same positions in t and t'.) Let u be the position of an outermost hyper-collapsing subterm of s. Then $s|u$ is the residual by the I-reduction of an outermost hyper-collapsing subterm of t at some position v. Since $t \sim_{hc} t'$, t' has an outermost hyper-collapsing subterm at v, and s' must have an outermost hyper-collapsing subterm at u. Symmetrically, everywhere that s' has an outermost hyper-collapsing subterm, so must s. Therefore $s \sim_{hc} s'$. Complete development of the remaining I-redexes in s and s' yields the terms t^I and t'^I respectively. These redexes are contained in hyper-collapsing subterms. Therefore by Lemma 7.9, the outermost hyper-collapsing subterms of s and s' are at the same positions as those of t^I and t'^I. Thus $t^I \sim^{R^I}_{hc} s \sim^{R^I}_{hc} s' \sim^{R^I}_{hc} t'^I$. By Lemma 7.13, $t^I \sim^R_{hc} t'^I$. □

PROOF OF THEOREM 7.4. By Lemma 7.7, $t_0 \to^{R^I}_\infty t_1$ and $t_0 \to^{R^I}_\infty t_2$. By Theorem 6.10, R^I is Church-Rosser, and so there is a t_3 in R^I such that $t_1 \to^{R^I}_\infty t_3$, and $t_2 \to^{R^I}_\infty t_3$.

By Lemma 7.10, there are R^I-reductions $t_1 \to^{R^I}_\infty s'$ and $t_2 \to^{R^I}_\infty s''$, such that $s' \sim_{hc} t_3$ and $s'' \sim_{hc} t_3$. Hence $s' \sim_{hc} s''$.

By Lemma 7.11, there exist R-terms t'_3 and t''_3 and reductions $t_1 \to^R_\infty t'_3$ and $t_2 \to^R_\infty t''_3$ such that s' and s'' reduce respectively to t'_3 and t''_3 by complete development of all I-redexes.

Since $s' \sim^{R^I}_{hc} s''$, Lemma 7.14 implies that $t'_3 \sim^R_{hc} t''_3$, completing the proof of the theorem. □

7.3 Proofs of unique normal form properties

THEOREM 7.15 *Every orthogonal TRS has the NF, UN, and UN$^\to$ properties.*

PROOF. By Theorem 7.4, such a TRS is CR$_{hc}$. Suppose t reduces to normal form n and to a term s. By CR$_{hc}$, s and n reduce to terms s' and n' such that $s' \sim_{hc} n'$. But n is a normal form, therefore $n = n'$. From the definition of $\sim_{hc}$ it is immediate that the only term hc-equivalent to a normal form n is n itself. Therefore $s' = n$. This demonstrates NF, which by Lemma 7.2 also implies UN and UN$^\to$. □

8. Reduction strategies

In this section we generalise to the infinitary setting some existing theory for reduction strategies in finitary term rewriting. We first introduce some terminology regarding infinitary strategies.

Definition 8.1
- A *reduction strategy* for a TRS is a function that maps every term t of the TRS to a (possibly empty) set of finite reductions starting from t.
- For any strategy S, the strategy $quasi(S)$ maps each term t_0 to the set of reductions of the form $t_0 \to_n t_n \to_m t_{n+m}$, where $t_n \to_m t_{n+m} \in S(t_n)$.
- A reduction sequence $t_0 \to^* t_1 \to^* \ldots t_\beta \to^* t_{\beta+1} \ldots t_\alpha$ is generated by a reduction strategy S from the term t_0 if for all $\beta < \alpha$, $t_\beta \to^* t_{\beta+1}$ is a member of $S(t_\beta)$.

Definition 8.2
- A reduction sequence is *transfinitely normalising* if it ends with a normal form.
- A reduction sequence is *normalising* if it ends with a normal form and had length at most ω.
- A reduction strategy is *transfinitely normalising* if it generates a transfinitely normalising sequence from every term which has a normal form.
- A reduction strategy is *normalising* if it generates a normalising sequence from every term which has a normal form.
- A reduction strategy S is (*transfinitely*) *hypernormalising* if $quasi(S)$ is (transfinitely) normalising.

For finitary TRSs, Huet and Lévy have shown that needed reduction is normalising for orthogonal systems [HL79, HL91], where a needed redex of a term is one such that every reduction of the term to normal form reduces at least one residual of the redex. This does not immediately generalise to the infinitary setting. A simple example is provided by the TRS consisting of the single rule: $A \to B(A, A)$. The term A has a normal form which is the infinite binary tree with B at each node. At every finite stage in a reduction starting from A, every redex is needed. However, it is easy to exhibit infinite reductions from A which do not compute the infinite normal form. For example, if we take the leftmost redex at each step, we generate the reduction: $A \to B(A, A) \to B(B(A, A), A) \to B(B(B(A, A), A), A) \to \ldots$ A condition of fairness must be added to obtain normalising strategies in the infinitary setting.

Definition 8.3 Given a reduction sequence S starting from t, a position u of t is *preserved* by S if no reduction step of S is performed at any position $v < u$.

Let R map each term to a set of its redexes. A strongly converging reduction $t \to_{\leq\omega} t'$ is *R-fair* if for every term t'' in the reduction, and every redex r of $R(t'')$, there exists some finite part of the remaining reduction starting at t'' that either reduces some residual of r, or does not preserve r. A reduction is *fair* if it is R-fair, where R is the set of all redexes in the initial term.

We will show that, informally speaking, needed-fair reduction is normalising for infinitary TRSs.

8.1 Neededness and infinitary reduction

For infinitary TRSs the notion of needed redex can be verbatim the same as for finitary TRSs. The definition is implicitly changed by the possibility of infinite normal forms.

DEFINITION 8.4 A redex s of a term t is *needed* if in every strongly converging reduction of t to normal form some residual of s is rewritten.

THEOREM 8.5 *For orthogonal TRSs, in every term having a normal form but not in normal form, there is at least one needed redex.*

PROOF. Huet and Lévy prove this for finite terms in [HL79, HL91]. A study of their proof reveals that it applies equally to infinite terms and strongly convergent reductions to normal form. We only note the few points where the infinitary aspects need some care. Two lemmas need new proofs.

Lemma 3.15 of [HL91] (Lemma 3.11 of [HL79]), stating that every nonempty reduction S has, in Huet and Lévy's terminology, an external redex, is proved by induction on the length of S, and needs a quite different proof to cope with transfinite sequences and infinite terms. First, we recall the definitions of external position and external redex.

Given a reduction sequence S starting from t, a position u of t is *preserved* by S if no reduction step of S is performed at any position $v < u$. A position u of t is *external* for S (written $u \in \mathcal{X}(S)$), if one of the following holds:

- u is preserved by S.
- S has the form $S_1 S_2 S_3$, where S_1 preserves u, S_2 is a single step at address $v < u$, such that u is one of the positions pattern-matched by the redex, and v is external for S_3.

A redex r of t is *reduced by* S if some step of S reduces a descendant of r. We write $\mathcal{R}(S)$ for the set of positions of such redexes of t, and $\mathcal{E}(S) = \mathcal{X}(S) \cap \mathcal{R}(S)$ for the set of positions of *external redexes* of t for S.

For finite reduction sequences the definition of external position coincides with that of [HL91]. Strong convergence implies that for transfinite sequences, the inductive definition is still well-founded, since in a proof that u is an external position, each application of the second case of the definition must choose a different reduction step of the sequence at a position less deep than u. But since u is finitely long, this can happen only finitely many times.

Lemma 3.15 of [HL91] (Lemma 3.11 of [HL79]) asserts that if S is nonempty, so is $\mathcal{E}(S)$. We shall prove this. Assume S is nonempty.

$\mathcal{X}(S)$ is clearly prefix-closed. If it contains all positions of t, then every redex of t (and there must be one, since S is nonempty) is in $\mathcal{E}(S)$. Otherwise, we must prove that $\mathcal{X}(S)$ contains some member of $\mathcal{R}(S)$. So suppose that it does not.

If, throughout S, no member of $\mathcal{X}(S)$ ever became the position of a redex, then S would preserve not only $\mathcal{X}(S)$, but every successor of every member of $\mathcal{X}(S)$, and these successors would also be in $\mathcal{X}(S)$. $\mathcal{X}(S)$ would therefore contain every position of t, contradiction.

Consider the first point S where some member of $\mathcal{X}(S)$ is the position of a redex. There must be a step $t_\alpha \to t_{\alpha+1}$ of S which creates this redex in $t_{\alpha+1}$, since by strong convergence a redex cannot suddenly appear at a limit term. Let the created redex be at position $u \in \mathcal{X}(S)$, and the redex reduced by $t_\alpha \to t_{\alpha+1}$ be at position $v > u$. By hypothesis, v is not in $\mathcal{X}(S)$. Let w be the position and i the integer such that $u \le w < w \cdot i \le v$, $w \in \mathcal{X}(S)$, and $w \cdot i \notin \mathcal{X}(S)$. Since no prefix of w is the position of any redex in $t \to_\alpha t_\alpha$, $w \cdot i$ is preserved by $t \to_\alpha t_\alpha$, and is pattern-matched by the redex at u.

If the redex at u is never reduced, then by orthogonality, $w \cdot i$ is preserved by the tail of S from $t_{\alpha+1}$ onwards. Since it is also preserved by $t \rightarrow_{\alpha+1} t_{\alpha+1}$, it is preserved by S, and is therefore in $\mathcal{X}(S)$, contradiction.

Therefore the redex at u is eventually reduced. Write $S = S_1 S_2 S_3$, where S_2 is the first reduction step performed at u. (S_1 will include the whole of $t \rightarrow_{\alpha+1} t_{\alpha+1}$.) Since $u \in \mathcal{X}(S)$ and u is preserved by $S_1 S_2$, u is in $\mathcal{X}(S_3)$. By orthogonality, $w \cdot i$ is preserved not only by $t \rightarrow_{\alpha+1} t_{\alpha+1}$, but by the whole of S_1. It is pattern-matched by the reduction step S_2, and therefore $w \cdot i \in \mathcal{X}(S)$, contradicting the definition of w and i.

This concludes the proof of nonemptiness of $\mathcal{E}(S)$.

Lemma 3.25 of [HL91] (Lemma 3.16 of [HL79]) also needs a different proof, but a simple modification works here. The lemma states that every term having a normal form but not in normal form has an external redex, that is, a redex at a position which is external to every reduction starting from that term. The proof proceeds by induction on the size of the term, applying the inductive hypothesis to the immediate subterms of the given term. For infinite terms, such an induction would not be well-founded. However, it is clear that the induction can be recast as an induction on the stable depth of the term, i.e. the depth d such that no reduction in any sequence starting from that term can be at depth d or less. The only terms that such an induction would miss are the infinite terms in normal form, for which the lemma is trivial. □

Definition 8.6

- A *needed* reduction to normal form is a strongly convergent reduction in which only needed redexes are reduced.
- A strongly convergent reduction to normal form is *quasi-needed* if in between the needed reduction steps at most finitely many other reductions take place.

For the case of needed reduction, the stipulation of strong convergence is redundant. To prove this, we must first introduce the notion of head normal forms and establish some of their properties.

Definition 8.7 A *head normal form* is a term which cannot be reduced to a redex. A term has a head normal form if it can be reduced to a head normal form.

The following equivalent characterisation of terms having head normal forms is useful.

Definition 8.8 A *perpetual* reduction is a strongly continuous reduction containing infinitely many reduction steps performed at the root. A *perpetual* term is a term from which there is a perpetual reduction.

Proposition 8.9 *A term t has a head normal form if and only if it is not perpetual.*

Proof. By a proof analogous to that of Lemma 7.9, we can see that the property of being perpetual is preserved by reduction. Thus if there were a perpetual reduction starting from t, then there would be one starting from every term which t was reducible to, and hence every such term would be reducible to a redex. So t would have no head normal form.

Conversely, if t has no head normal form, then every term to which t is reducible is reducible to a redex. Thus one can construct a perpetual reduction by reducing t to a redex, reducing that redex, reducing the resulting term to a redex, and so on. □

THEOREM 8.10 *A reduction, in which no step is performed inside a subterm having no head normal form, is strongly convergent.*

PROOF. If a reduction is not strongly convergent, then there will be some position which is infinitely often the position of the redex reduced by a reduction step. Choosing an outermost such position, after some finite number of reductions at that position, there would be infinitely many reductions performed at that position while no reduction was performed at any prefix of that position. The subterm at that position would thus have a perpetual reduction, and hence have no head normal form. □

COROLLARY 8.11 *A redex of a term, contained in a subterm having no head normal form, is not needed.*

PROOF. By the previous theorem, every reduction sequence reducing only redexes outside all such subterms is strongly convergent. Therefore by transfinite induction we may construct such a sequence starting from the given term, and extending to as high an ordinal as one chooses, subject to not reaching a normal form. However, all strongly convergent sequences have countable length, and therefore a normal form must be reached. The resulting sequence does not reduce any residual of any redex in any subterm having no normal form, demonstrating that they are not needed. □

COROLLARY 8.12 *Every needed reduction starting from a term having a normal form is strongly converging.*

PROOF. Immediate from Corollary 8.11 and Theorem 8.5. □

Needed reduction is the strategy mapping each term t to the set of one-step reductions starting from t which reduce a needed redex. Corollary 8.12 can be stated in another form.

COROLLARY 8.13 *Needed reduction is transfinitely normalising.* □

DEFINITION 8.14 A converging reduction $t \to_{\leq\omega} t'$ is *needed-fair* if for every term t'' in the reduction, and every needed redex r of t'', there exists some finite part of the remaining reduction starting at t'' that does not preserve r.

To prove that needed-fair reduction is normalising, we must first develop some of the properties of external redexes.

LEMMA 8.15 *If t can be reduced to a redex, then it can be reduced to a redex of the same rule in finitely many steps.*

PROOF. By the Compressing Lemma, if t has a strongly convergent reduction to a redex, then it has such a reduction of length at most ω. If the reduction is finite, the statement is proved. Otherwise, let $t \to_\omega s$ be a reduction of t to a redex. By strong convergence, and the finiteness of left-hand sides, some finite initial segment of this sequence must already reduce t to such a redex. □

$$\begin{tikzcd}
t_0 \arrow[r, "*", "\neg r"'] \arrow[d, "\neg r"'] & t'_0 \arrow[r, "ext"'] \arrow[d, "\neg r"'] & t_1 \arrow[r, "*", "\neg r"'] \arrow[d, "\neg r"'] & t'_1 \arrow[r, "ext"'] \arrow[d, "\neg r"'] & \cdots \\
s_0 \arrow[r, "*", "\neg r"'] & s'_0 \arrow[r, "*", "ext"'] & s_1 \arrow[r, "*", "\neg r"'] & s'_1 \arrow[r, "*", "ext"'] & \cdots
\end{tikzcd}$$

Figure 12.

Lemma 8.16 *If a term is reducible to a redex, then external-fair reduction will reduce it to a redex of the same rule in finitely many steps.*

Proof. Let t_0 have a head normal form, and let $t_0 \to \ldots$ be an external-fair reduction. By Lemma 8.15, there is a finite reduction $t_0 \to^* s_0$ of t_0 to a redex. Choose such a reduction of minimal length. We will construct the tiling diagram of Figure 12. Since $t_0 \to \ldots$ is external-fair, and every term not in normal form has an external redex, the sequence must consist of finite strings of non-external reductions alternating with external steps. Since a root redex is always external, the finite strings of non-external reductions do not perform any root reductions. Neither does the reduction of t_0 to s_0, since that reduction is a minimal length reduction of t_0 to a redex. It follows that $t'_0 \to_\infty s'_0$ and $s_0 \to_\infty s'_0$ must consist of non-root reductions. $s'_0 \to_\infty s_1$, being the projection of an external redex, must be either empty or a single step. If it is a single step, then it is external, and since s'_0 is a redex, the reduction must be of the root redex of s'_0. Therefore the step $t'_0 \to t_1$ must also be a root reduction. If $s'_0 \to_\infty s_1$ is empty, then this can only be because in the tiling diagram, some step of $t'_0 \to^* s'_0$ cancelled out the step $t'_0 \to t_1$. Both steps must be the unique residuals of an external redex of t_0. Therefore $t'_0 \to t_1$ also cancels out $t'_0 \to^* s'_0$, and $t_1 \to^* s_1$ contains no residuals of that redex. Since $t_0 \to^* s_0$ can contain only finitely many external steps, repeating the above argument shows that eventually, some external step $t'_n \to t_{n+1}$ must have a nonempty projection $s'_n \to s_{n+1}$, and hence must be a root reduction. □

Lemma 8.17 *If a term has a head normal form, then external-fair reduction will reduce it to head normal form in finitely many steps.*

Proof. Let t have a head normal form. Either t is in head normal form, or it is reducible to a redex. In the first case the lemma is trivial. In the second, by Lemma 8.16, external-fair reduction will reduce t to a redex in finitely many steps. Since a root-redex is external, external-fair reduction must eventually reduce that redex. If the resulting term is reducible to a redex, then the process is repeated. It can only be repeated finitely often, since by hypothesis t has a head normal form. Therefore external-fair reduction must in finitely many steps reduce t to a term not reducible to a redex, i.e. a head normal form. □

Corollary 8.18 *External-fair reduction is normalising.*

Proof. Let t have a normal form. Then t has a head normal form, and by Lemma 8.17, external-fair reduction starting from t will yield a head normal form t' in finitely many steps.

The external redexes of t' are just the external redexes of the immediate subterms of t', and therefore external-fair reduction of t' will reduce each of those subterms to head normal form in finitely many steps. Continuing the process clearly yields a strongly convergent reduction to normal form of length at most ω. □

Corollary 8.19 *Needed-fair reduction is normalising.*

Proof. Every needed-fair reduction is external-fair. □

Corollary 8.20 *Fair reduction is normalising.*

Proof. Every fair reduction is needed-fair. □

8.2 Other normalising reduction strategies

In the finitary setting the parallel-outermost strategy is hypernormalising. Our example at the beginning of this section shows that in the infinitary setting the parallel-outermost strategy does not guarantee reductions of at most length ω:

Lemma 8.21 *Parallel-outermost reduction is transfinitely hypernormalising.*

Proof. Consider a reduction $\mathcal{R}$ starting from a term which is strongly reducible to normal form. If $\mathcal{R}$ always eventually performs a parallel-outermost reduction, then $\mathcal{R}$ is needed-fair, and Cauchy-converging to normal from, hence strongly converging. Hence parallel-outermost is a transfinitely hypernormalising strategy. (It might not be hypernormalising, since a single parallel-outermost part of the sequence may itself be infinitely long.) □

We would like to have a normalising reduction strategy, that is, one which obtains normal forms in at most ω steps. The depth-increasing strategy has this property.

Definition 8.22 *Depth-increasing reduction* is the following strategy DI. Given a term t_0, for each $n \leq 0$ let t_{n+1} be derived from t_n by complete development of all redexes at positions of depth no more than n. Then $DI(t_0)$ is the set whose only member is the sequence $t_0 \to^* t_1 \to^* t_2 \to^* \ldots$

Theorem 8.23 *Depth-increasing reduction is hypernormalising.*

Proof. A residual of an external redex is always at the same position as the original redex. Therefore for each external redex of t at depth n, $quasi(DI(t))$ will reduce that redex at the nth stage of applying DI. Therefore $quasi(DI)$ is external-fair, so by Corollary 8.18 is normalising. Therefore DI is hypernormalising. □

9. Related work

Farmer and Watro [FW91] are the first to have studied transfinite reduction based on strong convergence. They are primarily concerned with proving the relationship between term graph reduction and term reduction of orthogonal systems such as combinatory logic. The possibility of cyclic graphs, which 'unwind' to infinite terms, requires a consideration of transfinite reduction.

Their results on transfinite reduction are limited to those which they need for that aim, and for example they prove a Compressing Lemma only for reductions of length up to ω^2.

Dershowitz, Kaplan and Plaisted [DKP91] have also studied transfinite reduction, basing their notion on Cauchy convergence rather than strong convergence. We have preferred strong convergence, for two main reasons.

Firstly, the notion of residual does not easily carry over from finitary reduction to Cauchy convergent transfinite reduction. As a result, while the Strip Lemma may hold for Cauchy convergent reduction (we conjecture that it does but have not proved it) the construction of the particular sequences s/r and r/s depends on strong convergence, and a quite different proof would be necessary for the situation of Cauchy convergent reduction, if it holds at all.

Secondly, several of the results of Dershowitz *et al.* depend on conditions stronger than Cauchy convergence. To obtain the Compression Lemma for orthogonal systems (in their terminology, ω-closure), they require the hypothesis of *top-termination*, which requires that, in our terminology, no finite term should have a perpetual reduction. Top-termination in fact implies that every reduction starting from a finite term is strongly converging, and this is the underlying reason that the Compression Lemma holds. The confluence properties of [DKP91] depend on the condition of ω-convergence of the system, which means that every reduction starting from a finite term converges. That paper also considers, instead of normal forms, ω-normal forms, which are terms which may not be in normal form, but which reduce only to themselves. These seem less meaningful as final results of a computation than normal forms, and their results concerning uniqueness of ω-normal forms again depend on the condition of ω-convergence. In contrast, when normal forms and strongly convergent reduction are considered, unique normal form results require no hypotheses beyond orthogonality.

For these reasons, the concept of strong convergence seems to us fundamental to transfinite rewriting.

We also note that for most orthogonal TRSs, the more liberal Cauchy reduction does not yield any new normal forms. We state the following theorem without proof.

THEOREM 9.1 *In an orthogonal TRS for which there is an upper bound on the set of depths of its left-hand sides, if a term has a Cauchy convergent reduction to normal form, then it has a strongly convergent reduction to the same normal form.*

The method of proof is to show that it is possible to omit from the given reduction sequence all steps which are performed within subterms having no head normal form. The resulting sequence is strongly convergent and has the same limit.

The theorem in general fails without the boundedness restriction. A counterexample is given by the rules $A(B^n(C)) \to A(B^{n+1}(C))$ (for $n \geq 0$) $A(C)$ reduces to the normal form $A(B^\omega)$ by Cauchy convergent reduction, but has no normal form by strongly convergent reduction.

REFERENCES

[AN80] A. Arnold and M. Nivat. The metric space of infinite trees: Algebraic and topological properties. *Fundamenta Informatica*, 4:445–476, 1980.

[Boo91] R. V. Book, editor. *Proc. 4th Conference on Rewriting Techniques and Applications*, Springer-Verlag, Lecture Notes in Computer Science 488, Como, Italy, 1991.

[DJ90] N. Dershowitz and J.-P. Jouannaud. Rewrite systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B: Formal Models and Semantics, chapter 15. North-Holland, Amsterdam, 1990.

[DKP91] N. Dershowitz, S. Kaplan, and D. A. Plaisted. Rewrite, rewrite, rewrite, rewrite, rewrite. *Theoretical Computer Science*, 83:71–96, 1991.

[FW91] W. M. Farmer and R. J. Watro. Redex capturing in term graph rewriting. In Book [Boo91], pages 13–24.

[HL79] G. Huet and J.-J. Lévy. Call-by-need computations in non-ambiguous linear term rewriting systems. Technical report, INRIA, 1979.

[HL91] G. Huet and J.-J. Lévy. Computations in orthogonal rewrite systems I and II. In J.-L. Lassez and G. D. Plotkin, editors, *Computational Logic: Essays in Honor of Alan Robinson*, pages 394–443. MIT Press, 1991. (Originally appeared as [HL79].).

[KKSdVar] J. R. Kennaway, J. W. Klop, M. R. Sleep, and F. J. de Vries. The adequacy of term graph rewriting for simulating term rewriting. *ACM TOPLAS*, to appear. An earlier version appeared as chapter 12 of [SPvE93].

[Klo92] J. W. Klop. Term rewriting systems. In S. Abramsky, D. Gabbay, and T. Maibaum, editors, *Handbook of Logic in Computer Science*, volume II, pages 1–116. Oxford University Press, 1992.

[Par83] D. Park. The "fairness problem" and nondeterministic computing networks. In J.W. de Bakker and J. van Leeuwen, editors, *Foundations of Computer Science IV, Part 2*, volume 159 of *Mathematical Centre Tracts*, pages 133–161. CWI, Amsterdam, 1983.

[Ros73] B. K. Rosen. Tree-manipulating systems and Church-Rosser systems. *Journal of the ACM*, 20:160–187, 1973.

[SPvE93] M. R. Sleep, M. J. Plasmeijer, and M. C. J. D. van Eekelen, editors. *Term Graph Rewriting: Theory and Practice*. John Wiley, 1993.