

INFINITARY COMBINATORY REDUCTION SYSTEMS: CONFLUENCE*

JEROEN KETEMA^a AND JAKOB GRUE SIMONSEN^b

^a Research Institute of Electrical Communication, Tohoku University, 2-1-1 Katahira, Aoba-ku, Sendai 980-8577, Japan
e-mail address: jketema@nue.riec.tohoku.ac.jp

^b Department of Computer Science, University of Copenhagen (DIKU), Universitetsparken 1, 2100 Copenhagen Ø, Denmark
e-mail address: simonsen@diku.dk

ABSTRACT. We study confluence in the setting of higher-order infinitary rewriting, in particular for infinitary Combinatory Reduction Systems (iCRSs). We prove that fully-extended, orthogonal iCRSs are confluent modulo identification of hypercollapsing sub-terms. As a corollary, we obtain that fully-extended, orthogonal iCRSs have the normal form property and the unique normal form property (with respect to reduction). We also show that, unlike the case in first-order infinitary rewriting, almost non-collapsing iCRSs are not necessarily confluent.

CONTENTS

1. Introduction	2
1.1. Overview and roadmap to confluence	3
2. Preliminaries	4
2.1. Terms, meta-terms, and positions	4
2.2. Valuations	6
2.3. Rewrite rules and reductions	7
2.4. Developments	11
2.5. Tiling diagrams	11
3. Projection pairs	12
4. Confluence	14
4.1. Hypercollapsingness	15
4.2. Confluence modulo	18

1998 ACM Subject Classification: D.3.1, F.3.2, F.4.1, F.4.2.

Key words and phrases: term rewriting, higher-order computation, combinatory reduction systems, lambda-calculus, infinite computation, confluence, normal forms.

* Parts of this paper have previously appeared as [11].

^a This author was partially funded by the Netherlands Organisation for Scientific Research (NWO) under FOCUS/BRICKS grant number 642.000.502.

4.3. Almost non-collapsingness	25
5. Normal form properties	26
6. Conclusion and suggestions for future work	27
Acknowledgement	27
References	27
Appendix A. Proof of Lemma 2.26	29

1. INTRODUCTION

This paper is part of a series outlining the fundamental theory of higher-order infinitary rewriting in the guise of *infinitary Combinatory Reduction Systems* (iCRSs). In preliminary papers [10, 11] we outlined basic motivation and definitions, and gave a number of introductory results. Moreover, we lifted a number of results from first-order infinitary rewriting to the setting of iCRSs. In particular, staple results such as compression and existence of complete developments of sets of redexes (subject to certain conditions) were proved.

The purpose of iCRSs is to extend infinitary term rewriting to encompass higher-order rewrite systems. This allows us, for instance, to reason about the behaviour of the well-known `map` functional when it is applied to infinite lists. The `map` functional and the usual constructors and destructors for lists can be represented by the below iCRS:

$$\begin{aligned}
 \text{map}([z]F(z), \text{cons}(X, XS)) &\rightarrow \text{cons}(F(X), \text{map}([z]F(z), XS)) \\
 \text{map}([z]F(z), \text{nil}) &\rightarrow \text{nil} \\
 \text{hd}(\text{cons}(X, XS)) &\rightarrow X \\
 \text{tl}(\text{cons}(X, XS)) &\rightarrow XS
 \end{aligned}$$

Systems such the above may satisfy certain simple criteria: being *orthogonal* (rules do not overlap syntactically) and *fully-extended* (if a variable is bound, then every meta-variable in its scope must be applied to it). We show that systems satisfying these two criteria are confluent modulo identification of a certain class of ‘meaningless’ subterms: Subterms that are *hypercollapsing*. As an example, `map` above, when applied to any infinite list `cons(s0, cons(s1, cons(...)))`, will yield identical results no matter how it is computed, except when applied to lists that will never yield a proper result *irrespective* of the evaluation order.

A succinct description for researchers familiar with infinitary rewriting: In the current paper, we employ the methods developed in previous papers to show that fully-extended, orthogonal iCRSs are confluent modulo identification of hypercollapsing subterms. As a corollary, we obtain that fully-extended, orthogonal iCRSs have the normal form property and the unique normal form property (with respect to reduction). Finally, we show that, unlike the case in first-order infinitary rewriting, almost non-collapsing iCRSs are not necessarily confluent.

Parts of this paper have previously appeared as [11]; the current paper *corrects* the results of that paper and *extends* them: We now allow rules with infinite right-hand sides, not only finite right-hand sides. The present paper requires some of the results proved in the previously published, peer-reviewed papers [10, 11]. A much-updated and extended version of these results is available as [13].

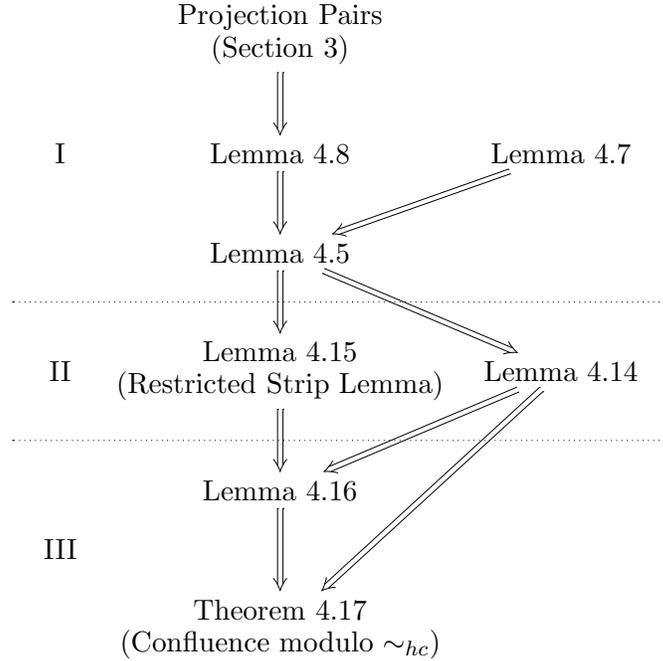


Figure 1: Roadmap to confluence

1.1. Overview and roadmap to confluence. The contents of the paper are as follows: Section 2 introduces preliminary notions. Section 3 on projection pairs recapitulates in an abstract way the fundamental results on *essential rewrite steps*, the primary method used to prove confluence in the higher-order infinitary setting. Section 4 provides proofs of our main results on confluence. Section 5 considers the normal form property, the unique normal form property, and the unique normal form with respect to reduction property. Section 6 concludes.

The main result of the paper is Theorem 4.17: Fully-extended, orthogonal iCRSs are confluent modulo identification of hypercollapsing subterms. To aid the reader we give a roadmap of the most important auxiliary results leading up to that theorem in Figure 1.

The auxiliary results are divided into three parts that all depend on the concept of projection pairs (and also the results of Sections 2.4 and 2.5 concerning developments and tiling diagrams, although not depicted explicitly). Part I, forming Section 4.1, relates hypercollapsing subterms and so-called hypercollapsing reductions (Lemma 4.5). These reductions simplify the reasoning regarding hypercollapsing subterms in the face of the arbitrary reductions that occur in the context of any confluence theorem.

Part II, forming the first half of Section 4.2, considers reductions that do not affect hypercollapsing subterms and establishes a Strip Lemma for such reductions. Although not depicted in Figure 1, Part II also establishes — in Proposition 4.12 — that the relation obtained by replacing the hypercollapsing subterms of a term by other hypercollapsing terms yields an equivalence relation (denoted by \sim_{hc}).

Part III, forming the latter half of Section 4.2, establishes our confluence theorem. The bulk of the work in this part consist in proving that confluence holds in case reductions that do not affect hypercollapsing subterms are considered (Lemma 4.16). Constructing tiling diagrams, the proof heavily depends on the restricted Strip Lemma established in Part II, and thus follows the lines of earlier confluence proofs [20]. However, the proof also contains a completely novel ingredient: The constructed tiling diagrams are in a sense incomplete and must be superimposed to effectively complete each other. The main result is established in Theorem 4.17.

2. PRELIMINARIES

We presuppose a working knowledge of the basics of ordinary finitary term rewriting [20]. The basic theory of infinitary Combinatory Reduction Systems has been laid out in [10, 11], and we give only the briefest of definitions in this section. Full proofs of all results may be found in the above-mentioned papers. Moreover, the reader familiar with [13] may safely skip this section; this section is essentially an abstract of that paper.

Throughout, infinitary Term Rewriting Systems are invariably abbreviated as iTRSs and infinitary λ -calculus is abbreviated as i λ c. Moreover, we denote the first infinite ordinal by ω , and arbitrary ordinals by α , β , γ , and so on. We use \mathbb{N} to denote the set of natural numbers, starting from zero.

2.1. Terms, meta-terms, and positions. We assume a signature Σ , each element of which has finite arity. We also assume a countably infinite set of variables and, for each finite arity, a countably infinite set of meta-variables of that arity. Countably infinite sets suffice, given that we can employ ‘Hilbert hotel’-style renaming.

The (infinite) *meta-terms* are defined informally in a top-down fashion by the following rules, where s and s_1, \dots, s_n are again meta-terms:

- (1) each variable x is a meta-term,
- (2) if x is a variable and s is a meta-term, then $[x]s$ is a meta-term,
- (3) if Z is a meta-variable of arity n , then $Z(s_1, \dots, s_n)$ is a meta-term,
- (4) if $f \in \Sigma$ has arity n , then $f(s_1, \dots, s_n)$ is a meta-term.

We consider meta-terms modulo α -equivalence.

A meta-term of the form $[x]s$ is called an *abstraction*. Each occurrence of the variable x in s is *bound* in $[x]s$, and each subterm of s is said to occur in the *scope* of the abstraction. If s is a meta-term, we denote by $root(s)$ the root symbol of s . Following the definition of meta-terms, we define $root(x) = x$, $root([x]s) = [x]$, $root(Z(s_1, \dots, s_n)) = Z$, and $root(f(s_1, \dots, s_n)) = f$.

The set of *terms* is defined as the set of all meta-terms without meta-variables. Moreover, a *context* is defined as a meta-term over $\Sigma \cup \{\square\}$ where \square is a fresh nullary function symbol and a *one-hole* context is a context in which precisely one \square occurs. If $C[\square]$ is a one-hole context and s is a term, we obtain a term by replacing \square by s ; the new term is denoted by $C[s]$.

Replacing a hole in a context does not avoid the capture of free variables: A free variable x in s is bound by an abstraction over x in $C[\square]$ in case \square occurs in the scope of the abstraction. This behaviour is not obtained automatically when working modulo α -equivalence: It is *always* possible find a representative from the α -equivalence class of

$C[\square]$ that does not capture the free variables in s . Therefore, we will always work with *fixed* representatives from α -equivalence classes of contexts. This convention ensures that variables will be captured properly.

Remark 2.1. Capture avoidance is disallowed for contexts as we do not want to lose variable bindings over rewrite steps in case: (i) an abstraction occurs in a context, and (ii) a variable bound by the abstraction occurs in a subterm being rewritten. Note that this means that the representative employed as the context must already be fixed *before* performing the actual rewrite step.

As motivation, consider λ -calculus: In the term $\lambda x.(\lambda y.x)z$, contracting the redex inside the context $\lambda x.\square$ yields $\lambda x.x$, whence the substitution rules for contexts should be such that

$$(\lambda x.\square)\{(\lambda y.x)z/\square\} \rightarrow_{\beta} \lambda x.x.$$

If we assumed capture avoidance in effect for contexts, we would have an α -conversion in the rewrite step, whence

$$(\lambda x.\square)\{(\lambda y.x)z/\square\} \rightarrow_{\beta} \lambda w.x,$$

which is clearly wrong.

Formally, meta-terms are defined by taking the metric completion of the set of finite meta-terms, the set inductively defined by the above rules. The distance between two terms is either taken as 0, if the terms are α -equivalent, or as 2^{-k} with k the minimal depth at which the terms differ, also taking into account α -equivalence. By definition of metric completion, the set of finite meta-terms is a subset of the set of meta-terms. Moreover, the metric on finite meta-terms extends uniquely to a metric on meta-terms.

Example 2.2. Any finite meta-term, e.g. $[x]Z(x, f(x))$, is a meta-term. We also have that $Z'(Z'(Z'(\dots)))$ is a meta-term, as is $Z_1([x_1]x_1, Z_2([x_2]x_2, \dots))$.

The meta-terms $[x]Z(x, f(x))$ and $[y]Z(y, f(y))$ have distance 0 and the meta-terms $[x]Z(x, f(x))$ and $[y]Z(y, f(z))$ have distance $\frac{1}{8}$.

Positions of meta-terms are defined by considering such terms in a top-down fashion. Given a meta-term s , its *set of positions*, denoted $\mathcal{Pos}(s)$, is the set of finite strings over \mathbb{N} , with ϵ the empty string, such that:

- (1) if $s = x$ for some variable x , then $\mathcal{Pos}(s) = \{\epsilon\}$,
- (2) if $s = [x]t$, then $\mathcal{Pos}(s) = \{\epsilon\} \cup \{0 \cdot p \mid p \in \mathcal{Pos}(t)\}$,
- (3) if $s = Z(t_1, \dots, t_n)$, then $\mathcal{Pos}(s) = \{\epsilon\} \cup \{i \cdot p \mid 1 \leq i \leq n, p \in \mathcal{Pos}(t_i)\}$,
- (4) if $s = f(t_1, \dots, t_n)$, then $\mathcal{Pos}(s) = \{\epsilon\} \cup \{i \cdot p \mid 1 \leq i \leq n, p \in \mathcal{Pos}(t_i)\}$.

The *depth* of a position p , denoted $|p|$, is the number of characters in p . Given $p, q \in \mathcal{Pos}(s)$, we write $p \leq q$ and say that p is a *prefix* of q , if there exists an $r \in \mathcal{Pos}(s)$ such that $p \cdot r = q$. If $r \neq \epsilon$, we also write $p < q$ and say that the prefix is *strict*. Moreover, if neither $p \leq q$ nor $q \leq p$, we say that p and q are *parallel*, which we write as $p \parallel q$.

We denote by $s|_p$ the subterm of s that occurs *at* position $p \in \mathcal{Pos}(s)$. Moreover, if $q \in \mathcal{Pos}(s)$ and $p < q$, we say that the subterm at position p occurs *above* q . Finally, if $p > q$, then we say that the subterm occurs *below* q .

Below we introduce a restriction on meta-terms called the *finite chains property*, which enforces the proper behaviour of valuations. Intuitively, a *chain* is a sequence of contexts in a meta-term occurring ‘nested right below each other’.

Definition 2.3. Let s be a meta-term. A *chain* in s is a sequence of (context, position)-pairs $(C_i[\square], p_i)_{i < \alpha}$, with $\alpha \leq \omega$, such that for each $(C_i[\square], p_i)$:

- (1) if $i + 1 < \alpha$, then $C_i[\square]$ has *one* hole and $C_i[t_i] = s|_{p_i}$ for some term t_i , and
 - (2) if $i + 1 = \alpha$, then $C_i[\square]$ has *no* holes and $C_i[\square] = s|_{p_i}$,
- and such that $p_{i+1} = p_i \cdot q_i$ for all $i + 1 < \alpha$ where q_i is the position of the hole in $C_i[\square]$.

If $\alpha < \omega$, respectively $\alpha = \omega$, then the chain is called *finite*, respectively *infinite*.

Observe that at most one \square occurs in any context $C_i[\square]$ in a chain. In fact, \square only occurs in $C_i[\square]$ if $i + 1 < \alpha$; if $i + 1 = \alpha$, we have $C_i[\square] = s|_{p_i}$.

2.2. Valuations. We next define valuations, the iCRS analogue of substitutions as defined for iTRSs and i λ c. As it turns out, the most straightforward and liberal definition of meta-terms has rather poor properties: Applying a valuation need not necessarily yield a well-defined term. Therefore, we also introduce an important restriction on meta-terms: the *finite chains property*. This property will also prove crucial in obtaining positive results later in the paper.

Essentially, the definitions are the same as in the case of CRSs [17, 23], except that the interpretation of the definition is top-down (due to the presence of infinite terms and meta-terms). Below, we use \vec{x} and \vec{t} as short-hands for, respectively, the sequences x_1, \dots, x_n and t_1, \dots, t_n with $n \geq 0$. Moreover, we assume n fixed in the next two definitions.

Definition 2.4. A *substitution* of terms \vec{t} for distinct variables \vec{x} in a term s , denoted $s[\vec{x} := \vec{t}]$, is defined as:

- (1) $x_i[\vec{x} := \vec{t}] = t_i$,
- (2) $y[\vec{x} := \vec{t}] = y$, if y does not occur in \vec{x} ,
- (3) $([y]s')[\vec{x} := \vec{t}] = [y](s'[\vec{x} := \vec{t}])$,
- (4) $f(s_1, \dots, s_m)[\vec{x} := \vec{t}] = f(s_1[\vec{x} := \vec{t}], \dots, s_m[\vec{x} := \vec{t}])$.

The above definition implicitly takes into account the usual variable convention [1] in the third clause to avoid the binding of free variables by the abstraction. We now define substitutes (adopting this name from Kahrs [5]) and valuations.

Definition 2.5. An *n-ary substitute* is a mapping denoted $\underline{\lambda}x_1, \dots, x_n.s$ or $\underline{\lambda}\vec{x}.s$, with s a term, such that:

$$(\underline{\lambda}\vec{x}.s)(t_1, \dots, t_n) = s[\vec{x} := \vec{t}]. \quad (2.1)$$

The intention of a substitute is to ensure that proper ‘housekeeping’ of substitutions is observed when performing a rewrite step. Reading Equation (2.1) from left to right yields a rewrite rule:

$$(\underline{\lambda}\vec{x}.s)(t_1, \dots, t_n) \rightarrow s[\vec{x} := \vec{t}].$$

The rule can be seen as a *parallel β -rule*. That is, a variant of the β -rule from (infinitary) λ -calculus which simultaneously substitutes multiple variables.

Definition 2.6. Let σ be a function that maps meta-variables to substitutes such that, for all $n \in \mathbb{N}$, if Z has arity n , then so does $\sigma(Z)$.

A *valuation* induced by σ is a relation $\bar{\sigma}$ that takes meta-terms to terms such that:

- (1) $\bar{\sigma}(x) = x$,
- (2) $\bar{\sigma}([x]s) = [x](\bar{\sigma}(s))$,
- (3) $\bar{\sigma}(Z(s_1, \dots, s_m)) = \sigma(Z)(\bar{\sigma}(s_1), \dots, \bar{\sigma}(s_m))$,
- (4) $\bar{\sigma}(f(s_1, \dots, s_m)) = f(\bar{\sigma}(s_1), \dots, \bar{\sigma}(s_m))$.

Similar to Definition 2.4, the above definition implicitly takes into account the variable convention, this time in the second clause, to avoid the binding of free variables by the abstraction.

The definition of a valuation yields a straightforward two-step way of applying it to a meta-term: In the first step each subterm of the form $Z(t_1, \dots, t_n)$ is replaced by a subterm of the form $(\underline{\lambda}\vec{x}.s)(t_1, \dots, t_n)$. In the second step Equation (2.1) is applied to each of these subterms.

In the case of (finite) CRSs, valuations are always (everywhere defined) *maps* taking each meta-term to a unique term [15, Remark II.1.10.1]. This is no longer the case when infinite meta-terms are considered. For example, given the meta-term $Z(Z(\dots Z(\dots)))$ and applying any map that satisfies $Z \mapsto \underline{\lambda}x.x$, we obtain $(\underline{\lambda}x.x)((\underline{\lambda}x.x)(\dots(\underline{\lambda}x.x)(\dots)))$. Viewing Equation (2.1) as a rewrite rule, this ‘ $\underline{\lambda}$ -term’ reduces only to itself and never to a *term*, as required by the definition of valuations (for more details, see [10]). To mitigate this problem a subset of the set of meta-terms is introduced in [10].

Definition 2.7. Let s be a meta-term. A *chain of meta-variables* in s is a chain in s , written $(C_i[\square], p_i)_{i < \alpha}$ with $\alpha \leq \omega$, such that for each $i < \alpha$ it is the case that $C_i[\square] = Z(t_1, \dots, t_n)$ with $t_j = \square$ for exactly one $1 \leq j \leq n$.

The meta-term s is said to satisfy the *finite chains property* if no infinite chain of meta-variables occurs in s .

Example 2.8. The meta-term $[x_1]Z_1([x_2]Z_2(\dots[x_n]Z_n(\dots)))$ satisfies the finite chains property. The meta-terms $Z(Z(\dots Z(\dots)))$ and $Z_1(Z_2(\dots Z_n(\dots)))$ do not.

From [10] we now have the following result:

Proposition 2.9. *Let s be a meta-term satisfying the finite chains property and let $\bar{\sigma}$ a valuation. There is a unique term that is the result of applying $\bar{\sigma}$ to s . \square*

2.3. Rewrite rules and reductions. Having defined terms and valuations, we move on to define rewrite rules and reductions.

2.3.1. Rewrite rules. We give a number of definitions that are direct extensions of the corresponding definitions from CRS theory.

Definition 2.10. A finite meta-term is a *pattern* if each of its meta-variables has distinct bound variables as its arguments. Moreover, a meta-term is *closed* if all of its variables occur bound.

We next define rewrite rules and iCRSs. The definitions are identical to the definitions in the finite case, with exception of the restrictions on the right-hand sides of the rewrite rules: The finiteness restriction is lifted and the finite chains property is put in place.

Definition 2.11. A *rewrite rule* is a pair (l, r) , denoted $l \rightarrow r$, where l is a finite meta-term and r is a meta-term, such that:

- (1) l is a pattern with a function symbol at the root,
- (2) all meta-variables that occur in r also occur in l ,
- (3) l and r are closed, and
- (4) r satisfies the finite chains property.

The meta-terms l and r are called, respectively, the *left-hand side* and the *right-hand side* of the rewrite rule.

An *infinitary Combinatory Reduction System (iCRS)* is a pair $\mathcal{C} = (\Sigma, R)$ with Σ a signature and R a set of rewrite rules.

With respect to the left-hand sides of rewrite rules, it is always the case that only finite chains of meta-variables occur, as the left-hand sides are finite.

We now define rewrite steps.

Definition 2.12. A *rewrite step* is a pair of terms (s, t) , denoted $s \rightarrow t$, adorned with a one-hole context $C[\square]$, a rewrite rule $l \rightarrow r$, and a valuation $\bar{\sigma}$ such that $s = C[\bar{\sigma}(l)]$ and $t = C[\bar{\sigma}(r)]$. The term $\bar{\sigma}(l)$ is called an $l \rightarrow r$ -*redex*, or simply a *redex*. The redex *occurs* at position p and depth $|p|$ in s , where p is the position of the hole in $C[\square]$.

A position q of s is said to occur in the *redex pattern* of the redex at position p if $q \geq p$ and if there does not exist a position q' with $q \geq p \cdot q'$ such that q' is the position of a meta-variable in l .

For example, $f([x]Z(x), Z') \rightarrow Z(Z')$ is a rewrite rule, and $f([x]h(x), a)$ rewrites to $h(a)$ by contracting the redex of the rule $f([x]Z(x), Z') \rightarrow Z(Z')$ occurring at position ϵ , i.e. at the root.

We now mention some standard restrictions on rewrite rules that we need later in the paper:

Definition 2.13. A rewrite rule is *left-linear*, if each meta-variable occurs at most once in its left-hand side. Moreover, an iCRS is *left-linear* if all its rewrite rules are.

Definition 2.14. Let s and t be finite meta-terms that have no meta-variables in common. The meta-term s *overlaps* t if there exists a non-meta-variable position $p \in \mathcal{Pos}(s)$ and a valuation $\bar{\sigma}$ such that $\bar{\sigma}(s|_p) = \bar{\sigma}(t)$.

Two rewrite rules *overlap* if their left-hand sides overlap and if the overlap does not occur at the root when two copies of the same rule are considered. An iCRS is *orthogonal* if all its rewrite rules are left-linear and no two (possibly the same) rewrite rules overlap.

In case the rewrite rules $l_1 \rightarrow r_1$ and $l_2 \rightarrow r_2$ overlap at position p , it follows that p cannot be the position of a bound variable in l_1 . If it were, we would obtain for some valuation $\bar{\sigma}$ and variable x that $\bar{\sigma}(l_1|_p) = x = \bar{\sigma}(l_2)$, which would imply that l_2 does not have a function symbol at the root, as required by the definition of rewrite rules.

Moreover, it is easily seen that if two left-linear rules overlap in an infinite term, there is also a finite term in which they overlap. As left-hand sides are *finite* meta-terms, we may appeal to standard ways of deeming CRSs orthogonal by inspection of their rules. We shall do so informally on several occasions in the remainder of the paper.

Definition 2.15. A rewrite rule is *collapsing* if the root of its right-hand side is a meta-variable. Moreover, a redex and a rewrite step are *collapsing* if the employed rewrite rule is. A rewrite step is *root-collapsing* if it is collapsing and occurs at the root of a term.

Definition 2.16. A pattern is *fully-extended* [4, 21], if, for each of its meta-variables Z and each abstraction $[x]s$ having an occurrence of Z in its scope, x is an argument of that occurrence of Z . Moreover, a rewrite rule is *fully-extended* if its left-hand side is and an iCRS is *fully-extended* if all its rewrite rules are.

Example 2.17. The pattern $f(g([x]Z(x)))$ is fully-extended. Hence, so is the rewrite rule $f(g([x]Z(x))) \rightarrow h([x]Z(x))$. The pattern $g([x]f(Z(x), Z'))$, with Z' occurring in the scope of the abstraction $[x]$, is not fully-extended as x does not occur as an argument of Z' .

2.3.2. *Transfinite reductions.* We can now define transfinite reductions. The definition is equivalent to those for iTRSs and λc [8, 6].

Definition 2.18. A *transfinite reduction* with domain $\alpha > 0$ is a sequence of terms $(s_\beta)_{\beta < \alpha}$ adorned with a rewrite step $s_\beta \rightarrow s_{\beta+1}$ for each $\beta+1 < \alpha$. In case $\alpha = \alpha' + 1$, the reduction is *closed* and of length α' . In case α is a limit ordinal, the reduction is called *open* and of length α . The reduction is *weakly continuous* or *Cauchy continuous* if, for every limit ordinal $\gamma < \alpha$, the distance between s_β and s_γ tends to 0 as β approaches γ from below. The reduction is *weakly convergent* or *Cauchy convergent* if it is weakly continuous and closed.

Intuitively, an open transfinite reduction is lacking a well-defined final term, while a closed reduction does have such a term.

As in [8, 6, 7], we prefer to reason about strongly convergent reductions.

Definition 2.19. Let $(s_\beta)_{\beta < \alpha}$ be a transfinite reduction. For each rewrite step $s_\beta \rightarrow s_{\beta+1}$, let d_β denote the depth of the contracted redex. The reduction is *strongly continuous* if it is weakly continuous and if, for every limit ordinal $\gamma < \alpha$, the depth d_β tends to infinity as β approaches γ from below. The reduction is *strongly convergent* if strongly continuous and closed.

Example 2.20. Consider the rewrite rule $f([x]Z(x)) \rightarrow Z(f([x]Z(x)))$ and observe that $f([x]x) \rightarrow f([x]x)$. Define $s_\beta = f([x]x)$ for all $\beta < \omega \cdot 2$. The reduction $(s_\beta)_{\beta < \omega \cdot 2}$, where in each step we contract the redex at the root, is open and weakly continuous. Adding the term $f([x]x)$ to the end of the reduction yields a weakly convergent reduction. Both reductions are of length $\omega \cdot 2$.

The above reduction is not strongly continuous as all contracted redexes occur at the root, i.e. at depth 0. In addition, it cannot be extended to a strongly convergent reduction. However, the following reduction

$$f([x]g(x)) \rightarrow g(f([x]g(x)) \rightarrow \dots \rightarrow g^n(f([x]g(x))) \rightarrow g^{n+1}(f([x]g(x))) \rightarrow \dots$$

is open and strongly continuous. Extending the reduction with the term g^ω , where g^ω is shorthand for the infinite term $g(g(\dots g(\dots)))$, yields a strongly convergent reduction. Both reductions are of length ω .

Notation 2.21. By $s \twoheadrightarrow^\alpha t$, respectively $s \twoheadrightarrow^{\leq \alpha} t$, we denote a *strongly convergent* reduction of ordinal length α , respectively of ordinal length at most α . By $s \twoheadrightarrow t$ we denote a *strongly convergent* reduction of arbitrary ordinal length and by $s \twoheadrightarrow^* t$ we denote a reduction of finite length. Reductions are usually ranged over by capital letters such as D , S , and T . The concatenation of reductions S and T is denoted by $S;T$.

Note that the concatenation of any finite number of strongly convergent reductions yields a strongly convergent reduction. For strongly convergent reductions, the following is proved in [10].

Lemma 2.22. *If $s \twoheadrightarrow t$, then the number of steps contracting redexes at depths less than $d \in \mathbb{N}$ is finite for any d and $s \twoheadrightarrow t$ has countable length.* \square

The following result [10] shows that, as in other forms of infinitary rewriting, reductions can always be ‘compressed’ to have length at most ω :

Theorem 2.23 (Compression). *For every fully-extended, left-linear iCRS, if $s \twoheadrightarrow^\alpha t$, then $s \twoheadrightarrow^{\leq \omega} t$. \square*

2.3.3. Descendants and residuals. The twin notions of *descendants* and *residuals* formalise, respectively, “what happens” to positions and redexes across reductions. Across a rewrite step, the only positions that can have descendants are those that occur outside the redex pattern of the contracted redex and that are not positions of the variables bound by abstractions in the redex pattern. Across a reduction, the definition of descendants follows from the notion of a descendant across a rewrite step, employing strong convergence in the limit ordinal case. We do not appeal to further details of the definitions in the remainder of this paper and these details are hence omitted. For the full definitions we refer the reader to [10].

Notation 2.24. Let $s \twoheadrightarrow t$. Assume $P \subseteq \mathcal{P}os(s)$ and \mathcal{U} a set of redexes in s . We denote the descendants of P across $s \twoheadrightarrow t$ by $P/(s \twoheadrightarrow t)$ and the residuals of \mathcal{U} across $s \twoheadrightarrow t$ by $\mathcal{U}/(s \twoheadrightarrow t)$. Moreover, if $P = \{p\}$ and $\mathcal{U} = \{u\}$, then we also write $p/(s \twoheadrightarrow t)$ and $u/(s \twoheadrightarrow t)$. Finally, if $s \twoheadrightarrow t$ consists of a single step contracting a redex u , then we sometimes write \mathcal{U}/u .

2.3.4. Reducts. In addition to descendants and residuals we need a notion of a reduct of a subterm.

Definition 2.25. Let $s_0 \twoheadrightarrow^\alpha s_\alpha$. Moreover, let $p_0 \in \mathcal{P}os(s_0)$ and $p_\alpha \in \mathcal{P}os(s_\alpha)$. The subterm $s_\alpha|_{p_\alpha}$ is called a *reduct* of $s_0|_{p_0}$ if for every $\beta \leq \alpha$ there exists a position q_β in s_β with $q_\alpha = p_\alpha$ such that:

- if $\beta = 0$, then $q_0 = p_0$,
- if $\beta = \beta' + 1$, then $q_\beta = q_{\beta'}$ unless $s_{\beta'} \rightarrow s_{\beta'+1}$ contracts a redex strictly above $q_{\beta'}$ in which case $q_\beta \in q_{\beta'}/(s_{\beta'} \rightarrow s_{\beta'+1})$, and
- if β is a limit ordinal, then $q_\beta = q_\gamma$ for all large enough $\gamma < \beta$.

A position $q \geq p_\alpha$ in s_α is said to occur in a *reduct* $s_\alpha|_{p_\alpha}$ of $s_0|_{p_0}$ if, for all positions $p_\alpha < p' \leq q$ in s_α , the subterm $s_\alpha|_{p'}$ is a reduct of a subterm strictly below p_0 in s_0 .

The above notion generalises the usual notion of a reduct. The usual notion is obtained by taking the root position for every q_β . There is a slight difference between reducts and descendants: Contracting a redex at a position p yields a reduct at position p , while p does not have a descendant.

Employing the above definition, we obtain the following property with respect to bound variables; a proof can be found in Appendix A.

Lemma 2.26. *Let $s_0 \twoheadrightarrow^\alpha s_\alpha$ and suppose u_α and v_α in s_α are residuals of redexes in s_0 . Denote for all $\gamma \leq \alpha$ by u_γ and v_γ , respectively, the unique redexes at positions p_γ and q_γ in s_γ of which u_α and v_α are residuals. Assume for all $\gamma < \alpha$ that if the step $s_\gamma \rightarrow s_{\gamma+1}$ contracts a redex at prefix position of q_γ then the redex is a residual of a redex in s_0 . Then, given that a variable bound by an abstraction in the redex pattern of u_α occurs in v_α , it follows that (a) $p_0 < q_0$ and (b) q_α occurs in the reduct $s_\alpha|_{p_\alpha}$ of $s_0|_{p_0}$. \square*

Observe that, as nestings of subterms can only be created by substitution of bound variables, the above lemma precludes nestings from occurring in reducts unless the conditions in the lemma are met.

2.4. Developments. We need some basic facts about developments which we recapitulate now.

Assuming in the remainder of this section that every iCRS is *orthogonal* and that s is a term and \mathcal{U} a set of redexes in s , we first define developments:

Definition 2.27. A *development* of \mathcal{U} is a strongly convergent reduction such that each step contracts a residual of a redex in \mathcal{U} . A development $s \twoheadrightarrow t$ is called *complete* if $\mathcal{U}/(s \twoheadrightarrow t) = \emptyset$. Moreover, a development is called *finite* if $s \twoheadrightarrow t$ is finite.

A complete development of a set of redexes does not necessarily exist in the infinite case. Consider for example the rule $f(Z) \rightarrow Z$ and the term f^ω . The set of all redexes in f^ω does not have a complete development: After any (partial) development a residual of a redex in f^ω always remains at the root of the resulting term. Hence, any complete development will have an infinite number of root-steps and hence is not strongly convergent.

Although complete developments do not always exist, the following results can still be obtained [11], where we write $s \Rightarrow^{\mathcal{U}} t$ for the reduction $s \twoheadrightarrow t$ if it is a complete development of the set of redexes \mathcal{U} in s .

Lemma 2.28. *If \mathcal{U} has a complete development and if $s \twoheadrightarrow t$ is a (not necessarily complete) development of \mathcal{U} , then $\mathcal{U}/(s \twoheadrightarrow t)$ has a complete development.* \square

Lemma 2.29. *Let s be a term and \mathcal{U} a set of redexes in s . If \mathcal{U} is finite, then it has a finite complete development.* \square

Proposition 2.30. *Let \mathcal{U} and \mathcal{V} be sets of redexes in s such that \mathcal{U} has a complete development $s \Rightarrow^{\mathcal{U}} t$ and \mathcal{V} is finite. The following diagram commutes:*

$$\begin{array}{ccc} s & \xrightarrow{\mathcal{V}} & t' \\ \Downarrow \mathcal{U} & & \Downarrow \mathcal{U}/(s \Rightarrow^{\mathcal{V}} t') \\ t & \xrightarrow{\mathcal{V}/(s \Rightarrow^{\mathcal{U}} t)} & s' \end{array}$$

We remark that we do not use the full power of the above proposition: In the current paper \mathcal{V} is always a singleton set.

2.5. Tiling diagrams. Tiling diagrams are defined as follows.

Definition 2.31. A *tiling diagram* of two strongly convergent reductions $S : s_{0,0} \rightarrow^\alpha s_{\alpha,0}$ and $T : s_{0,0} \rightarrow^\beta s_{0,\beta}$ is a rectangular arrangement of strongly convergent reductions as depicted in Figure 2 such that (1) each reduction $S_{\gamma,\delta} : s_{\gamma,\delta} \twoheadrightarrow s_{\gamma+1,\delta}$ is a complete development of a set of redexes of $s_{\gamma,\delta}$, and similarly for $T_{\gamma,\delta} : s_{\gamma,\delta} \twoheadrightarrow s_{\gamma,\delta+1}$, (2) the leftmost vertical reduction is S and the topmost horizontal reduction is T , and (3) for each γ and δ the set of redexes developed in $S_{\gamma,\delta}$ is the set of residuals of the redex contracted in $s_{\gamma,0} \rightarrow s_{\gamma+1,0}$ across the (strongly convergent) reduction $T_{\gamma,[0,\delta]} : s_{\gamma,0} \rightarrow s_{\gamma,1} \rightarrow \cdots s_{\gamma,\delta}$ (symmetrically for $T_{\gamma,\delta}$).

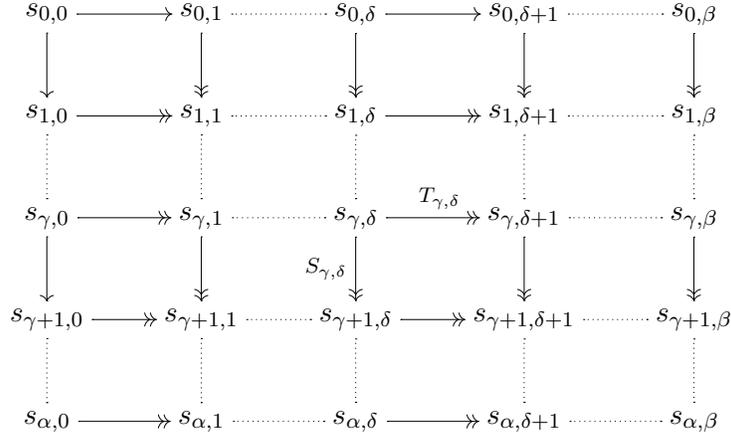


Figure 2: A tiling diagram

For $S_{[0,\alpha],\beta}$ we usually write S/T and we call this reduction the *projection* of S across T (similarly for $T_{\alpha,[0,\beta]}$ and T/S). Moreover, if T consists of a single step contracting a redex u , we also write S/u (symmetrically T/u).

Given two strongly convergent reductions, even in the case these where one of these is *finite*, a tiling need not exist, witness e.g. the failure of the Strip Lemma in [6]. To cope with this issue later in the paper we employ the following theorem from [11] in combination with the results from Section 2.4. The theorem, which is valid for *orthogonal* iCRSs, extends Theorem 12.6.5 from [7]: In [7] it is assumed that S and T are reductions of limit ordinal length; in this paper, S and T may be reductions of arbitrary ordinal length.

Theorem 2.32. *Let S and T be strongly convergent reductions starting from the same term. Suppose that the tiling diagram for S and T exists except that it is unknown if S/T and T/S are strongly convergent and end in the same term. The following are equivalent:*

- (1) *The tiling diagram of S and T can be completed, i.e. S/T and T/S are strongly convergent and end in the same term.*
- (2) *S/T is strongly convergent.*
- (3) *T/S is strongly convergent.*

□

3. PROJECTION PAIRS

For the confluence result, we shall employ a technique by van Oostrom [22], combining the concept of essentiality from [14, 3] with a termination technique from [19, 18]. We give an abstract formulation of the technique in terms of so-called *projection pairs*; the formulation is taken from [9] and extends the more primitive notions from [12]. Please note that the main definitions given below do not occur in [12], and the reader is thus advised to review them carefully.

We require an auxiliary definition:

Definition 3.1. Let s and t be terms and $P \subseteq \mathcal{Pos}(s)$. The set P is a *prefix set* of s if P is *finite* and if all prefixes of positions in P are also in P . Moreover, t *mirrors* s in P , if for all $p \in P$ it holds that $p \in \mathcal{Pos}(t)$ and $root(t|_p) = root(s|_p)$ (modulo α -equivalence).

Van Oostrom's technique uses a termination argument on a prefix set P and a reduction D that consists of a finite sequence of complete developments starting from a term s . The crux of the termination argument is, as always, some measure μ over a well-founded order that decreases across the sequence of developments.

The technique hinges on projecting D across a single rewrite step starting from s . If the rewrite step occurs in some specific prefix, Q , of s , it is called *essential*; otherwise it is called *inessential*. Projecting D across the step to obtain a new sequence D' , one shows by case analysis that the measure is always non-increasing, but decreases strictly if the step is essential. The specific prefix, Q , is obtained from a prefix set P of the final term of D by a map ε mapping P to Q . The pair (μ, ε) is called a *projection pair*.

Intuition done, we now proceed to give precise definitions:

Definition 3.2. Given a well-founded order \prec on a set O , a *projection pair* is a pair (μ, ε) of maps over finite sequences of complete developments D and prefix sets P of the final term of the chosen D such that:

- $\mu_P(D)$ maps to an element of O , and
- $\varepsilon_P(D)$ maps to a prefix set of the initial term of D ,

and such that if D' is a sequence of complete developments strictly shorter than D with P' a prefix set of the final term of D' , then $\mu_{P'}(D') \prec \mu_P(D)$.

The map μ is the measure and ε is the map for prefix sets. The measure requires a sequence that is strictly shorter than D to map to a smaller element in the well-founded order. Although of a technical nature, this property is easily obtained in case tuples are used to define the well-founded order and the tuples are first compared length-wise and next lexicographically.

We can now define (in)essentiality as follows:

Definition 3.3. Let (μ, ε) be a projection pair. If D is a finite sequence of complete developments and P is a prefix set of the final term of D , then a position p of, respectively a redex u in, the initial term of D is called *essential* for P if p , respectively the position of u , occurs in $\varepsilon_P(D)$. A position, respectively a redex, is called *inessential* otherwise.

The existence of the projection mentioned above can now be formulated as the soundness of a projection pair:

Definition 3.4. Let \prec be a well-founded order on a set O . A projection pair (μ, ε) is *sound* if for every finite sequence of complete development D , prefix set P of the final term of D , and $s \rightarrow t$, with s the initial term of D , it holds that:

- (1) if $s \rightarrow t$ consists of a *single* step contracting a redex u at an essential position, with no residual in u/D occurring at a position in P , then there exists a D' such that $\mu_P(D') \prec \mu_P(D)$, and
- (2) if $s \rightarrow t$ consists of one or more steps and *only* contracts redexes at *inessential* positions, then there exists a D' such that $\mu_P(D') = \mu_P(D)$ and $\varepsilon_P(D') = \varepsilon_P(D)$,

where in both cases D' is a finite sequence of complete developments with initial term t such that the final term of D' mirrors the final one of D in P .

The restriction in the first clause that no residual in u/D occurs in P ensures that the projection preserves P . Together, the clauses formalise the intuition behind ε , i.e. that P only depends on positions in $\varepsilon_P(D)$. The map is constant for reductions contracting only redexes outside $\varepsilon_P(D)$ and, obviously, any term in such a reduction mirrors all the other terms in $\varepsilon_P(D)$.

Remark 3.5. The first clause of Definition 3.4 deals neither with reductions where residuals from u/D occur in P nor with infinite reductions. In the next section, we deal with the first by means of the restriction on strictly shorter sequences of complete developments and with the second by means of strong convergence.

We have the following theorem, proved in [12]:

Theorem 3.6. *For each fully-extended, orthogonal iCRS a sound projection pair exists.*

4. CONFLUENCE

We will now present our confluence result. To start, recall that confluence in general does not hold for iTRSs, even under assumption of orthogonality [8]. As every iTRS can be seen as a fully-extended iCRS, it follows that fully-extended, orthogonal iCRSs are in general not confluent either.

In case of iTRSs two approaches are known for restoring confluence [8], namely (1) identifying all subterms that disrupt confluence, and (2) restricting the rewrite rules that are allowed. Identifying all subterms that disrupt confluence leads to the definition of so-called hypercollapsing subterms and yields the result that orthogonal iTRSs are confluent modulo these subterms. Restricting the rules that are allowed yields results regarding almost non-collapsing iTRSs.

Considering only *fully-extended, orthogonal* iCRSs, we next prove that such iCRSs are also confluent modulo hypercollapsing subterms, where a term s is called hypercollapsing if for every $s \rightarrow t$ we have that $t \rightarrow t'$ where t' has a collapsing redex at the root. This not only generalises the result for iTRSs but also a similar result for λc [6]. Regrettably, the proofs for iTRSs and λc from [7] cannot be lifted to the general higher-order case: For iTRSs the proof hinges on the Strip Lemma and for λc it hinges on the notion of head reduction, both of which fail to properly generalise to iCRSs. To circumvent these problems, we employ the measure defined in the previous section.

As an added benefit, we are able to overcome a small infelicity in the similar proof for λc in [7]. There, Lemma 12.8.14 treats reductions outside hypercollapsing subterms in a way similar to our Lemma 4.16; however, for λc , the induction step in the proof of [7] can apparently only be carried out if a stronger induction hypothesis is assumed than the one given — the two resulting reductions should be outside hypercollapsing subterms. The general result for iCRSs given in the present paper subsumes the one for λc .

Apart from confluence modulo, we show in Section 4.3 that the positive result that an iTRS is confluent iff it is almost non-collapsing cannot be trivially lifted to iCRSs.

Remark 4.1. On a historical note: Courcelle [2] observed similar problems with confluence while trying to define second-order substitutions on infinite trees. He circumvented these problems by requiring rules to be non-collapsing. In a general setting such as ours this would be too harsh a restriction.

4.1. Hypercollapsingness. We now proceed to define a particularly troublesome kind of reduction and term.

Definition 4.2. A *hypercollapsing reduction* is an open strongly continuous reduction with an infinite number of root-collapsing steps.

Thus, a hypercollapsing reduction is a particular example of a transfinite reduction of some limit ordinal length α that cannot be extended to a strongly *convergent* reduction — the term s_α is undefined. Note that, writing $(s_\beta)_{\beta < \alpha}$ for a hypercollapsing reduction sequence, we have that every initial sequence $(s_\beta)_{\beta < \gamma+1}$ with $\gamma < \alpha$ is strongly convergent.

Example 4.3. Hypercollapsing reductions are known even in the first-order case where we have, e.g. (in the syntax of iCRSs) the rewrite rule $f(Z) \rightarrow Z$ and the term f^ω from which there is the hypercollapsing reduction

$$f^\omega \rightarrow f^\omega \rightarrow \dots$$

which is obtained by repeatedly contracting the redex at the root.

For an example in more higher-order spirit, consider the rule $g([x]Z(x)) \rightarrow Z([x]Z(x))$. From the term $g([x]g(x))$ there is the hypercollapsing reduction

$$g([x]g(x)) \rightarrow g([x]g(x)) \rightarrow \dots$$

which is again obtained by repeatedly contracting the redex at the root.

The crucial definition is now the following:

Definition 4.4. A term s is said to be *hypercollapsing* if, for all terms t with $s \twoheadrightarrow t$, there exists a term t' with $t \twoheadrightarrow t'$ such that t' has a collapsing redex at the root.

It is not hard to see that a hypercollapsing term has a hypercollapsing reduction starting from it; the converse, however, is much more difficult, and is contained in the following lemma, to the proof of which we devote the remainder of the section.

Lemma 4.5. *Let s be a term. If there is a hypercollapsing reduction starting from s , then s is hypercollapsing.*

To start, we observe that hypercollapsing reductions satisfy a ‘compression’ property:

Lemma 4.6. *Let s be a term. If there is a hypercollapsing reduction starting from s , then there is a hypercollapsing reduction of length ω starting from it.*

Proof. By definition, we may write a hypercollapsing reduction starting from s as:

$$s = s_0 \twoheadrightarrow s'_0 \rightarrow s_1 \twoheadrightarrow s'_1 \rightarrow s_2 \twoheadrightarrow \dots,$$

where $s'_i \rightarrow s_{i+1}$ is root-collapsing and no root-collapsing steps occur in $s_i \twoheadrightarrow s'_i$ for all $i \in \mathbb{N}$.

We inductively define a hypercollapsing reduction of length ω :

$$s = t_0 \twoheadrightarrow^* t'_0 \rightarrow t_1 \twoheadrightarrow^* t'_1 \rightarrow t_2 \twoheadrightarrow^* \dots,$$

where for all $i \in \mathbb{N}$ it holds that $t'_i \rightarrow t_{i+1}$ is root-collapsing and that $t_i \twoheadrightarrow^* t'_i$ is finite and without root-collapsing steps. First, define $t_0 = s_0 = s$. Next, assume we have defined a term t_i with $t_i \twoheadrightarrow s_i$. Compression of $t_i \twoheadrightarrow s_i \twoheadrightarrow s'_i \rightarrow s_{i+1}$ yields a reduction $t_i \twoheadrightarrow^* t'_i \rightarrow t_{i+1} \twoheadrightarrow^{\leq \omega} s_{i+1}$ with $t'_i \rightarrow t_{i+1}$ root-collapsing and $t_i \twoheadrightarrow^* t'_i$ finite and without root-collapsing steps. Thus, there is a hypercollapsing reduction with the required properties. \square

The following is the iCRS analogue of Lemma 12.8.4 in [7] for iTRSs and strengthening for i λ c:

Lemma 4.7. *Let s and t be terms with $s \rightarrow t$. If there is a hypercollapsing reduction starting in s , then there is a hypercollapsing reduction starting in t .*

Proof. Define $s_0 = s$, $t_0 = t$, and suppose u is the redex contracted in $s \rightarrow t$. By Lemma 4.6, we may write the hypercollapsing reduction starting in s_0 as:

$$s_0 \rightarrow^* s'_0 \rightarrow s_1 \rightarrow^* s'_1 \rightarrow s_2 \rightarrow^* \dots,$$

where for all $i \in \mathbb{N}$, we have that $s'_i \rightarrow s_{i+1}$ is root-collapsing and $s_i \rightarrow^* s'_i$ is finite and without root-collapsing steps. By repeated application of Proposition 2.30, we obtain the following diagram:

$$\begin{array}{ccccccccc} s_0 & \xrightarrow{*} & s'_0 & \longrightarrow & s_1 & \xrightarrow{*} & s'_1 & \longrightarrow & s_2 & \xrightarrow{*} & \dots \\ \downarrow u & & \Downarrow \mathcal{U}'_0 & & \Downarrow \mathcal{U}_1 & & \Downarrow \mathcal{U}'_1 & & \Downarrow \mathcal{U}_2 & & \Downarrow \dots \\ t_0 & \twoheadrightarrow & t'_0 & \twoheadrightarrow & t_1 & \twoheadrightarrow & t'_1 & \twoheadrightarrow & t_2 & \twoheadrightarrow & \dots \end{array}$$

Write S_i for $s_i \rightarrow^* s'_i \rightarrow s_{i+1} \rightarrow^* \dots$ and T_i for $t_i \twoheadrightarrow t'_i \twoheadrightarrow t_{i+1} \twoheadrightarrow \dots$. If we can show for each $i \in \mathbb{N}$ that a root-collapsing step occurs in T_i , then an infinite number of root-collapsing steps occurs in T_0 , implying that the reduction is hypercollapsing.

To show that a root-collapsing step occurs in each T_i we distinguish two cases: (1) a root-collapsing step occurs in S_i not contracting a residual of u , and (2) all root-collapsing steps in S_i contract a residual of u . We treat each of these cases in turn:

- (1) Suppose a root-step occurs in S_i that does not contract a residual of u . Thus, there exists a root-collapsing step $s'_j \rightarrow s_{j+1}$ with $j \geq i$ such that the contracted redex, say v , is not a residual of u . Since we have by construction that \mathcal{U}'_j contracts only residuals of u , orthogonality implies that a residual of v occurs at the root of t'_j and that no other residuals of v occur in t'_j . Also by construction, $t'_j \twoheadrightarrow t_{j+1}$ contracts precisely all residuals of v . Hence, $t'_j \twoheadrightarrow t_{j+1}$ is a root-collapsing step.
- (2) Suppose all root-collapsing steps in S_i contract a residual of u (which implies u is a collapsing redex). Moreover, for any term in S_i call a set \mathcal{V} of residuals of u a *root-nesting* if \mathcal{V} is the largest set such that for each redex v in \mathcal{V} there exists a (partial) development of \mathcal{V} that ends in a term with a residual of v at the root (this residual is also a residual of u).

For every term along S_i the root-nesting is finite and non-empty. Finiteness follows as only a finitely many steps occur before each term in S_i and as right-hand sides of rewrite rules satisfy the finite chains condition. Non-emptiness follows as otherwise a root-step occurs that (a) does not contract a residual of u and (b) brings a residual of u to the root. Such a step is by definition root-collapsing, contradicting the assumption that all root-collapsing steps in S_i contract residuals of u .

We make the following claim:

Claim 1. *The number of redexes in a root-nesting eventually increases due to contraction of a step outside the root-nesting.*

To prove the claim, observe that, by definition, any redex inside a root-nesting occurring at a non-root position occurs as an argument of another redex inside the root-nesting. As no redex outside the root-nesting occurs above the root-nesting, the cardinality of a root-nesting can, hence, only decrease by contracting a redex inside the root-nesting.

Now suppose the cardinality of the root-nesting increases only by contracting redexes inside the root-nesting itself. By definition of root-nestings, an increase in cardinality is due in this case to nestings that are created among the redexes already present in the root-nesting. By Lemma 2.26 and the fact that only a finite number of redexes occur above each other redex, only a finite number of nestings occur that increase the cardinality. Hence, as an infinite number of root-collapsing steps occurs in S_i , all of which are in the root-nesting, eventually only decreases can occur, whence, by finiteness of root-nestings, all redexes in the root-nesting must be contracted, contradicting the non-emptiness of root-nestings. This concludes the proof of Claim 1.

By Claim 1, a step outside a root-nesting of S_i occurs that increases the cardinality. The redex contracted in the step, say v , is collapsing and does not contract a residual of u , by definition of root-nestings. Moreover, as the cardinality increases, a (partial) development of residuals of u exists which brings a residual v to the root. As v is not a residual of u , it follows by Lemma 2.28 and the fact that complete developments of residuals of u in terms along S_i exist, that a root-collapsing redex occurs in T_i . Since the redex is a residual of a collapsing redex in S_i which is eventually contracted, a root-collapsing step occurs in T_i .

As required, we have that a root-step occurs in each T_i . Hence, T_0 is a hypercollapsing reduction starting from $t_0 = t$. \square

The next lemma shows that the property of being reducible to a term with a collapsing redex at the root cannot be destroyed by reductions unless they contain a collapsing step at the root themselves. In the proof of the lemma we assume that we have at our disposal a sound projection pair, as is possible by Theorem 3.6.

Lemma 4.8. *If $s \rightarrow t$ has no root-collapsing steps and s reduces to a collapsing redex, then so does t .*

Proof. We show by ordinal induction that every term s_α in $s \rightarrow t$ reduces to a collapsing redex by a finite sequence of complete developments D_α . Denote by P_α the set of positions of the redex pattern at the root of the final term of D_α and remark that this set is a prefix set. To facilitate the induction we also show for each $\beta \leq \alpha$ either that $\mu_{P_\alpha}(D_\alpha) \prec \mu_{P_\beta}(D_\beta)$ or that $\mu_{P_\alpha}(D_\alpha) = \mu_{P_\beta}(D_\beta)$, $\varepsilon_{P_\alpha}(D_\alpha) = \varepsilon_{P_\beta}(D_\beta)$, and $s_\beta \rightarrow s_\alpha$ consists solely of inessential steps.

For $s_0 = s$, it follows by assumption that s_0 reduces to a collapsing redex. In fact, by strong convergence and compression, s_0 reduces to a collapsing redex by a finite reduction D_0 . As any finite reduction can be seen as a finite sequence of complete developments the result follows.

For $s_{\alpha+1}$, there are two cases to consider given the redex u contracted in $s_\alpha \rightarrow s_{\alpha+1}$ depending on the occurrence of a residual of u at the root of the final term of D_α :

- In case no residual of u occurs at the root of the final term of D_α , we discriminate between u being either essential or inessential for P_α . In case u is essential, the result follows by

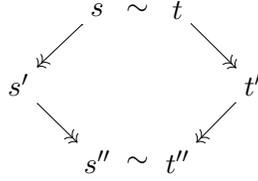


Figure 3: Definition 4.9

the induction hypothesis and Definition 3.4(1). Otherwise, the induction hypothesis and Definition 3.4(2) can be applied, where the assumed reduction consists of a single step.

- In case a residual of u occurs at the root of the final term of D_α , a root-collapsing step not contracting a residual of u occurs somewhere along D_α . Otherwise, a residual of u cannot occur at the root of the final term of D_α , because $s_\alpha \rightarrow s_{\alpha+1}$ is not root-collapsing. Hence, there exists a finite sequence D'_α of complete developments that is shorter than D_α and that has a collapsing redex, other than a residual of u , at the root of its final term. By Definition 3.2, it follows that $\mu_{P'_\alpha}(D'_\alpha) \prec \mu_{P_\alpha}(D_\alpha)$, where P'_α is the set of positions of the redex pattern at the root of the final term of D'_α . The case in which no residual of u occurs at the root of the final term of the complete development now applies and the result follows.

For s_α with α a limit ordinal, it follows by the well-foundedness of \prec and the induction hypothesis that there exist a $\beta < \alpha$ such that for every $\beta < \gamma < \alpha$ we have $\mu_{P_\gamma}(D_\gamma) = \mu_{P_\beta}(D_\beta)$. Hence, since we also have by the induction hypothesis that $\varepsilon_{P_\gamma}(D_\gamma) = \varepsilon_{P_\beta}(D_\beta)$ for all $\beta < \gamma < \alpha$ and that all redexes contracted in $s_\beta \rightarrow s_\gamma$ are inessential, the result follows by strong convergence and Definition 3.4(2), where the assumed reduction is $s_\beta \rightarrow s_\alpha$. \square

We can now prove Lemma 4.5:

Proof of Lemma 4.5. Let $s \rightarrow t$ be arbitrary. By compression and strong convergence, we may write $s \rightarrow^* t' \rightarrow^{\leq \omega} t$ such that all root-reductions occur in $s \rightarrow^* t'$. By repeated application of Lemma 4.7, there exists a hypercollapsing reduction starting from t' . In particular, t' reduces to a collapsing redex. Since $t' \rightarrow t$ contains no steps at the root, Lemma 4.8 yields that t reduces to a collapsing redex, proving that s is hypercollapsing. \square

4.2. Confluence modulo. We now prove confluence modulo identification of hypercollapsing subterms. Confluence modulo is defined as follows:

Definition 4.9. An iCRS is *confluent modulo* an equivalence relation \sim if for all $s \rightarrow s'$ and $t \rightarrow t'$ with $s \sim t$ there exist terms s'' and t'' such that $s' \rightarrow s''$ and $t' \rightarrow t''$ with $s'' \sim t''$ (see Figure 3).

We first show that identification of hypercollapsing subterms yields an equivalence relation. To this end we introduce some notation and show that hypercollapsingness is preserved under replacement of hypercollapsing subterms.

Notation 4.10. We write $s \sim_{hc} t$ if t can be obtained from s by replacing a number of hypercollapsing subterms of s by other hypercollapsing terms.

Proposition 4.11. *Let s and t be terms. If s is hypercollapsing and $s \sim_{hc} t$, then t is hypercollapsing.*

Proof. Let P be the set of positions of hypercollapsing subterms in s that are replaced to obtain t . By definition of s there exists a hypercollapsing reduction S starting from it. The redex patterns employed in the steps of S either occur completely outside or completely inside the reducts of subterms in s at positions in P . This follows by orthogonality and the fact that the subterms at positions in P are hypercollapsing, i.e. each reduct reduces to a term with a collapsing redex at the root. By orthogonality and by the fact that free variables cannot become bound when substituted into the reducts, it does not matter whether any substitutes occur in the reducts.

Omit from S all steps that occur inside the reducts of subterms in s at positions in P to obtain a reduction S' of length α . By definition of S' , together with orthogonality and fully-extendedness, there exists a reduction T of length α starting in t such that for all $\beta \leq \alpha$ we have that the redex pattern and position of the redex contracted in the β th step of both S' and T are identical. Hence, if S' is hypercollapsing, then so is T and the result follows by Lemma 4.5. If S' is not hypercollapsing, then s reduces to a reduct of subterm at a position $p \in P$ and the same holds for T . As the subterm at position p in t is hypercollapsing, there exist a hypercollapsing reduction starting from it. Again, by the fact that free variables cannot get bound when terms are substituted into other terms and by orthogonality, it is irrelevant that any substitutes occur. Hence, T can be prolonged to obtain a hypercollapsing reduction and the result follows again by Lemma 4.5. \square

We can now prove that \sim_{hc} has the required properties:

Proposition 4.12. *The relation \sim_{hc} is an equivalence relation, which is closed under substitution of terms for free variables.*

Proof. We have to prove that the relation is reflexive, symmetric, and transitive. Reflexivity and symmetry are immediate by definition. Transitivity follows by Proposition 4.11.

To see that relation is closed under substitution, consider a hypercollapsing term s and a term t that is a substitution instance of s . By definition of s there exists a hypercollapsing reduction S of length α starting from it. By orthogonality and the fact that no free variables are bound in the terms substituted into s , there exists a reduction T of length α starting from t such that for all $\beta \leq \alpha$ we have that the redex pattern and position of the redex contracted in β th step of both S and T are identical. Hence, since S is hypercollapsing, so is T and the result follows by Lemma 4.5. \square

Introducing some further notation, we next show that we can accurately ‘simulate’ reductions in terms that are \sim_{hc} -related.

Notation 4.13. By $s \rightarrow^{\text{out}} t$ we denote a rewrite step that does not occur inside any hypercollapsing subterm of s .

Lemma 4.14. *Let $s \rightarrow t$ have α steps that occur outside hypercollapsing subterms. If $s \sim_{hc} s'$, then there exists a reduction $s' \rightarrow^{\text{out}} t'$ of length α such that $t \sim_{hc} t'$. Moreover, for all $\beta \leq \alpha$ the redex pattern and position of the redex contracted in the β th step of $s' \rightarrow^{\text{out}} t'$ are identical to those of the β th step of $s \rightarrow t$ that occurs outside a hypercollapsing subterm.*

Proof. Let $s \rightarrow^\gamma t$ and $s \sim_{hc} s'$. We prove the result by ordinal induction on γ .

If $\gamma = 0$, the result is immediate, as an empty reduction is by definition one that only contracts redexes outside hypercollapsing subterms.

If $\gamma = \delta + 1$, assume $s \twoheadrightarrow^\gamma t = s \twoheadrightarrow^\delta s_\delta \rightarrow t$. By the induction hypothesis there exists a term s'_δ such that $s' \twoheadrightarrow^{\text{out}} s'_\delta$ and $s_\delta \sim_{hc} s'_\delta$. There are two possibilities for $s_\delta \rightarrow t$, depending on the contracted redex occurring either outside all hypercollapsing subterms or inside one of them:

- If the redex occurs outside all hypercollapsing subterms, then $s_\delta \sim_{hc} s'_\delta$ together with orthogonality and fully-extendedness implies that a redex employing the same rewrite rule as the redex contracted in $s_\delta \rightarrow t$ occurs at the same position in s'_δ . Moreover, the redex occurs outside all hypercollapsing subterms by Proposition 4.11. Hence, contracting the redex in s'_δ yields a step $s'_\delta \twoheadrightarrow^{\text{out}} t'$. That $t \sim_{hc} t'$ follows by $s_\delta \sim_{hc} s'_\delta$ and the fact that the same rewrite rule is employed in both $s_\delta \rightarrow t$ and $s'_\delta \twoheadrightarrow^{\text{out}} t'$: Clearly, t and t' are identical at all positions p that descend from positions not in hypercollapsing subterms of s_δ or s'_δ . If q is the position of a maximal hypercollapsing subterm of s_δ , then it is also the position of a maximal hypercollapsing subterm of s'_δ and vice versa, by Proposition 4.11. The descendants of q occur at identical positions in t and t' and are hypercollapsing subterms, since $s_\delta \sim_{hc} s'_\delta$ and since \sim_{hc} is closed under substitution. Note, however, that the hypercollapsing subterms are not necessarily maximal.
- If the redex occurs inside a hypercollapsing subterm, then $t \sim_{hc} s_\delta$. Hence, by transitivity of \sim_{hc} we have $t \sim_{hc} s'_\delta$ and we can define $t' = s'_\delta$.

If γ is a limit ordinal, the result is immediate by strong convergence and the induction hypothesis. \square

Before proving the main theorem of this section, we show that reductions outside hypercollapsing subterms are confluent *modulo* \sim_{hc} . To this end we first prove a restricted variant of the Strip Lemma. It is well-known that the usual Strip Lemma for iTRSs fails for $\text{i}\lambda\text{c}$ [6], and, hence, we see that it must also fail for iCRSs .

Lemma 4.15 (Restricted Strip Lemma). *If $S : s \twoheadrightarrow^{\text{out}} t$ and $T : s \twoheadrightarrow^{\text{out}} t'$, then S/T and T/S exist and end in the same term.*

Proof. Denote the length of S by α . We prove the lemma by ordinal induction on α . Note that, since T contracts a single redex u , we have that T/S is actually a complete development of the residuals of u in t . Obviously, if $\alpha = 0$, then the result follows trivially.

If α is a successor ordinal, then the result is immediate by Proposition 2.30 and the induction hypothesis.

If α is a limit ordinal, then Theorem 2.32 and the induction hypothesis ensure that we only need to show that T/S is strongly convergent. In other words, since T contracts a single redex u , we need to prove that u/S has a strongly convergent complete development. Assume the contrary and observe this implies the rewrite rule employed in T is collapsing, otherwise any development of u/S is strongly convergent.

By assumption, there exists a term t^* in T/S such that from t^* onwards an infinite number of steps occur at a certain depth d and no steps occur above d . Moreover, as function symbols have finite arity, there is a position p at depth d at which an infinite number of steps occur. As T/S contracts only residuals of redexes in t , it follows by Lemma 2.26 that redexes contracted along T/S can only be nested by contracting a residual of a redex, say v , in t such that v occurs above all redexes in t whose residuals are being nested. Hence, since only a finite number of residuals occurs in t above the redex whose residual occurs at position p in t^* , we have by the finite chains condition that the reducts of subterms of t in the subterm at position p in t^* occur in finite chains. Hence, since again by Lemma 2.26

no further nestings can be created among different reducts of the same subterm of t or among reducts of parallel subterms of t , eventually all contracted redexes at p are reducts of a single subterm in t . As there are an infinite number of steps at depth d , this means a hypercollapsing reduction exists starting in a subterm of t , say at position q .

By strong convergence and limit ordinal length of S , we can write $S = S_0; S_1$, where S_0 has successor ordinal length and $S_1 : s^* \rightarrow^{\text{out}} t$ is a non-empty final segment of S contracting no redexes at prefix positions of q . Hence, S_0 has length strictly less than α and $s^*|_q \rightarrow^{\text{out}} t|_q$. As there is a hypercollapsing reduction starting from $t|_q$, it follows by Definition 4.2 that there is also a hypercollapsing reduction starting from $s^*|_q$. But then, by Lemma 4.5, we have that $s^*|_q$ is hypercollapsing, which implies that $s^*|_q \rightarrow^{\text{out}} t|_q$ is empty and that $s^*|_q = t|_q$. Thus, $s^*|_q$ contains a set of descendants of u having no complete development (giving rise to the hypercollapsing reduction starting from $s^*|_q = t|_q$), whence u/S_0 has no complete development. Since S_0 has length strictly less than α , this contradicts the induction hypothesis. Hence, T/S is strongly convergent. \square

Lemma 4.16. *If $s \rightarrow^{\text{out}} t_0$ and $s \rightarrow^{\text{out}} t_1$, then there exist terms t_0^* and t_1^* such that $t_0 \rightarrow^{\text{out}} t_0^*$ and $t_1 \rightarrow^{\text{out}} t_1^*$ with $t_0^* \sim_{hc} t_1^*$.*

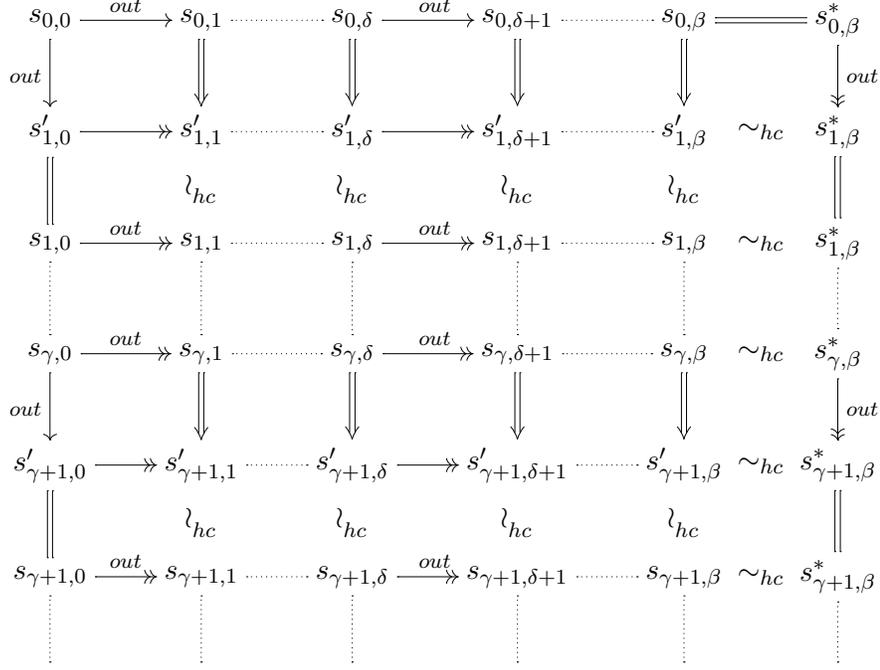
Proof. Let $S : s \rightarrow^{\text{out}} t_0$ and $T : s \rightarrow^{\text{out}} t_1$. By compression and Lemma 4.14 we may assume that both S and T have length at most ω . Suppose S has length $\alpha \leq \omega$ and T has length $\beta \leq \omega$. The proof proceeds in four steps: In the first step two ‘tiling diagrams’ are constructed, yielding (i) a reduction starting in t_0 , and (ii) a reduction starting in t_1 . In the second step a relation is established between the ‘tiles’ of the two diagrams. Employing the relation, it is shown in the third step that the two reductions obtained in the first step are strongly convergent. Finally, in the fourth step it is shown that the final terms of the two strongly convergent reductions are equivalent modulo \sim_{hc} .

Tiling diagrams. Write $S : s_{0,0} \rightarrow^{\text{out}} s_{1,0} \rightarrow^{\text{out}} \dots s_{\gamma,0} \rightarrow^{\text{out}} s_{\gamma+1,0} \rightarrow^{\text{out}} \dots s_{\alpha,0}$ and $T : s_{0,0} \rightarrow^{\text{out}} s_{0,1} \rightarrow^{\text{out}} \dots s_{0,\delta} \rightarrow^{\text{out}} s_{0,\delta+1} \rightarrow^{\text{out}} \dots s_{0,\beta}$ and define $s'_{\gamma,0} = s_{\gamma,0}$ for all $\gamma \leq \alpha$. We inductively construct the ‘tiling diagram’ in Figure 4(a):

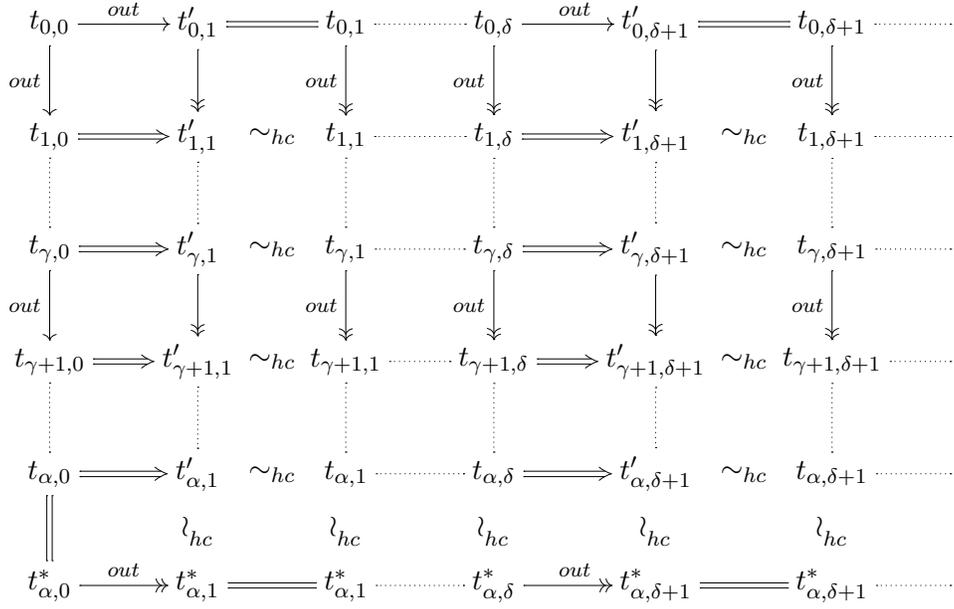
- the tiling of $s_{\gamma,0} \rightarrow^{\text{out}} s'_{\gamma+1,0}$ and $s_{\gamma,0} \rightarrow^{\text{out}} s_{\gamma,\beta}$ exists by Lemma 4.15;
- the reduction $s_{\gamma+1,0} \rightarrow^{\text{out}} s_{\gamma+1,\beta}$ and the equivalences $s_{\gamma+1,\delta} \sim_{hc} s'_{\gamma+1,\delta}$ for all $0 \leq \delta \leq \beta$ exist by Lemma 4.14 and the existence of $s'_{\gamma+1,0} \rightarrow s'_{\gamma+1,\beta}$;
- the reduction $s_{\gamma,\beta} \rightarrow^{\text{out}} s^*_{\gamma+1,\beta}$ and the equivalence $s'_{\gamma+1,\beta} \sim_{hc} s^*_{\gamma+1,\beta}$ exist by Lemma 4.14 and the existence of $s_{\gamma,\beta} \rightarrow s'_{\gamma+1,\beta}$;
- the equivalence $s^*_{\gamma+1,\beta} \sim_{hc} s_{\gamma+1,\beta}$ exists by transitivity of \sim_{hc} and since $s_{\gamma+1,\beta} \sim_{hc} s'_{\gamma+1,\beta} \sim_{hc} s^*_{\gamma+1,\beta}$.

As can be seen in Figure 4(a), the construction yields a reduction S^* starting in $t_1 = s^*_{0,\beta}$ such that all steps in the reduction occur outside hypercollapsing subterms. Note that the constructed diagram is not a tiling diagram in the strict sense of the word: No reduction occurs at the bottom and the diagram consists not only of reductions but also of equivalences modulo hypercollapsing subterms.

To obtain the second ‘tiling diagram’, depicted in Figure 4(b), we write $S : t_{0,0} \rightarrow^{\text{out}} t_{1,0} \rightarrow^{\text{out}} \dots t_{\gamma,0} \rightarrow^{\text{out}} t_{\gamma+1,0} \rightarrow^{\text{out}} \dots t_{\alpha,0}$ and $T : t_{0,0} \rightarrow^{\text{out}} t_{0,1} \rightarrow^{\text{out}} \dots t_{0,\delta} \rightarrow^{\text{out}} t_{0,\delta+1} \rightarrow^{\text{out}} \dots t_{0,\beta}$ and define $t'_{0,\delta} = t_{0,\delta}$ for all $\delta \leq \beta$. The diagram is constructed by vertically repeating the horizontal construction of Figure 4(a). The construction yields a reduction $T^* : t_0 = t^*_{\alpha,0} \rightarrow^{\text{out}} t^*_{\alpha,1} \rightarrow^{\text{out}} \dots t^*_{\alpha,\delta} \rightarrow^{\text{out}} \dots$.



(a)



(b)

Figure 4: The ‘tiling diagrams’ from the proof of Lemma 4.16

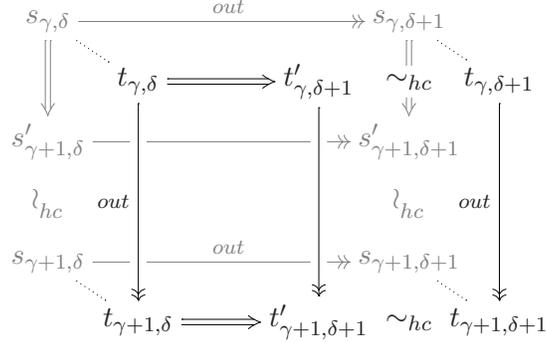


Figure 5: Superimposing the ‘tiles’ of the ‘tiling diagrams’ in Figure 4

Relation. Superimpose the tiles of the constructed ‘tiling diagrams’ as depicted in Figure 5, i.e. $s_{\gamma, \delta}$ and $t_{\gamma', \delta'}$ are superimposed if $\gamma = \gamma'$ and $\delta = \delta'$. Define $s_{0, \delta} = s'_{0, \delta}$, and $t_{\gamma, 0} = t'_{\gamma, 0}$ for all $\gamma \leq \alpha$ and $\delta \leq \beta$. By construction of the ‘tiling diagrams’, no term is superimposed on $s_{\gamma, \beta}$ with $\gamma \leq \alpha$ in case $\beta = \omega$ and similarly for $t_{\alpha, \delta}$ with $\delta < \beta$ in case $\alpha = \omega$.

We next prove for all superimposed terms $s_{\gamma, \delta}$ and $t_{\gamma, \delta}$ that $s_{\gamma, \delta} \sim_{hc} s'_{\gamma, \delta} \sim_{hc} t_{\gamma, \delta} \sim_{hc} t'_{\gamma, \delta}$. The proof is by induction on γ and δ . Induction is allowed because $s_{\gamma, \delta}$ and $t_{\gamma, \delta}$ exist for all $\gamma < \alpha$ and $\delta < \beta$:

- In case either $\gamma = 0$ or $\delta = 0$, we have $s_{\gamma, \delta} = s'_{\gamma, \delta} = t_{\gamma, \delta} = t'_{\gamma, \delta}$ by definition. Hence, since \sim_{hc} is an equivalence relation, $s_{\gamma, \delta} \sim_{hc} s'_{\gamma, \delta} \sim_{hc} t_{\gamma, \delta} \sim_{hc} t'_{\gamma, \delta}$.
- In case of $\gamma = \gamma' + 1$ and $\delta = \delta' + 1$, we have by definition of the ‘tiling diagrams’ that $s_{\gamma, \delta} \sim_{hc} s'_{\gamma, \delta}$ and $t_{\gamma, \delta} \sim_{hc} t'_{\gamma, \delta}$. Hence, by transitivity of \sim_{hc} , we obtain the desired result if we can establish $s_{\gamma, \delta} \sim_{hc} t'_{\gamma, \delta}$.

By Lemmas 4.15 and 4.14, as employed in the construction of the ‘tiling diagrams’, $s_{\gamma, \delta'} \xrightarrow{\text{out}} s_{\gamma, \delta}$ is essentially a development of residuals of the redex u contracted in $s_{0, \delta'} \xrightarrow{\text{out}} s_{0, \delta}$ such that no residuals of u in $s_{\gamma, \delta}$ remain *outside* hypercollapsing subterms. Since we have by the induction hypothesis that $s_{\gamma, \delta'} \sim_{hc} t_{\gamma, \delta'}$ and since every step in $s_{\gamma, \delta'} \xrightarrow{\text{out}} s_{\gamma, \delta}$ occurs outside hypercollapsing subterms, it follows by orthogonality and fully-extendedness that there exists a reduction $t_{\gamma, \delta'} \rightarrow t''_{\gamma, \delta}$ such that $s_{\gamma, \delta} \sim_{hc} t''_{\gamma, \delta}$. Since $s_{\gamma, \delta'} \xrightarrow{\text{out}} s_{\gamma, \delta}$ is essentially a development of residuals of u , it follows that $t_{\gamma, \delta'} \rightarrow t''_{\gamma, \delta}$ can be chosen to be a development of residuals of u , i.e. of the redex contracted in $t_{0, \delta'} \rightarrow t'_{0, \delta}$. Moreover, it follows that all residuals of u left in $t''_{\gamma, \delta}$ occur inside hypercollapsing subterms. Hence, since we have by Lemma 2.28 that $t''_{\gamma, \delta} \rightarrow t'_{\gamma, \delta}$, we also have that $t'_{\gamma, \delta} \sim_{hc} t''_{\gamma, \delta}$. But then, by transitivity of \sim_{hc} it follows that $s_{\gamma, \delta} \sim_{hc} t'_{\gamma, \delta}$, as required.

Strong convergence. Employing that $s_{\gamma, \delta} \sim_{hc} s'_{\gamma, \delta} \sim_{hc} t_{\gamma, \delta} \sim_{hc} t'_{\gamma, \delta}$ holds for all superimposed $s_{\gamma, \delta}$ and $t_{\gamma, \delta}$, we next prove that the reduction $S^* : s_{0, \beta}^* \xrightarrow{\text{out}} s_{1, \beta}^* \xrightarrow{\text{out}} \dots s_{\gamma, \beta}^* \xrightarrow{\text{out}} \dots$ in Figure 4(a) is strongly convergent. The proof is by contradiction. Thus, suppose S^* is not strongly convergent. There now exists a position p of minimal depth d such that an infinite number of steps occur at p . As each step in S^* occurs outside hypercollapsing subterms, it follows by minimality of d that from some γ onwards

no redexes are contracted above p and that all redexes contracted at p are of non-collapsing rules. Moreover, by strong convergence of $s_{\gamma,0} \twoheadrightarrow s_{\gamma,\beta}$, there is a δ such that all steps in $s_{\gamma,\delta} \twoheadrightarrow s_{\gamma,\beta}$ also occur below d .

Suppose for some minimal $\kappa \geq \gamma$ that a redex is contracted at some position $q < p$ in either $s_{\kappa,\delta} \Rightarrow s'_{\kappa+1,\delta}$ or $s_{\kappa,\delta} \twoheadrightarrow s_{\kappa,\beta}$. By dependence of the depth of the steps in $s_{\kappa,\delta} \twoheadrightarrow s_{\kappa,\beta}$ on the depth of the steps in $s_{\lambda,\delta} \twoheadrightarrow s_{\lambda,\beta}$ for all $\gamma \leq \lambda < \kappa$, it follows by minimality of κ that the reduction must be $s_{\kappa,\delta} \Rightarrow s'_{\kappa+1,\delta}$. This implies that a redex is also contracted at position q in $s_{\kappa,\beta} \Rightarrow s'_{\kappa+1,\beta}$. Since the redex is by definition not contracted in $s_{\kappa,\beta}^* \twoheadrightarrow^{\text{out}} s_{\kappa+1,\beta}^*$, it follows that the subterm at position q in $s_{\kappa,\beta}^*$ is hypercollapsing. However, as $q < p$, this implies that the infinite number of redexes contracted at position p cannot occur, as redexes in S^* are contracted outside hypercollapsing subterms. Hence, for all $\kappa \geq \gamma$ we have that no reduction $s_{\kappa,\delta} \Rightarrow s'_{\kappa+1,\delta}$ or $s_{\kappa,\delta} \twoheadrightarrow s_{\kappa,\beta}$ contracts a redex at strict prefix position of p .

Since all steps in $s_{\gamma,\delta} \twoheadrightarrow s_{\gamma,\beta}$ occur below d , the above implies that if a redex is contracted at position p in some $s_{\kappa,\beta}^* \twoheadrightarrow^{\text{out}} s_{\kappa+1,\beta}^*$ for minimal $\kappa \geq \gamma$, a redex is also contracted at position p in $s_{\kappa,\delta} \Rightarrow s'_{\kappa+1,\delta}$. Since the contracted redex is of a non-collapsing rule, it follows that the function symbol that occurs at position p in both $s_{\kappa+1,\beta}^*$ and $s'_{\kappa+1,\delta}$ is the root symbol of the next redex contracted at position p . Hence, $s_{\gamma,\delta} \Rightarrow s'_{\gamma+1,\delta} \sim_{hc} s_{\gamma+1,\delta} \Rightarrow s'_{\gamma+2,\delta} \sim_{hc} s_{\gamma+2,\delta} \Rightarrow \dots$ contains an infinite number of steps at position p without any interleaving of collapsing steps at that position. However, as redexes contracted at position p cannot occur inside hypercollapsing subterms by definition of S^* , we have that $t_{0,\delta} \twoheadrightarrow t_{\alpha,\delta}$ also contracts an infinite number of redexes at position p , which is impossible by strong convergence of this reduction, contradiction. Hence, S^* is strongly convergent.

By a similar argument as above it follows that the reduction $T^* : t_{\alpha,0}^* \twoheadrightarrow^{\text{out}} t_{\alpha,1}^* \twoheadrightarrow^{\text{out}} \dots t_{\alpha,\delta}^* \twoheadrightarrow^{\text{out}} \dots$ is strongly convergent.

Equivalence modulo. Since $s_{\gamma,\delta} \sim_{hc} s'_{\gamma,\delta} \sim_{hc} t_{\gamma,\delta} \sim_{hc} t'_{\gamma,\delta}$ for all γ and δ in both ‘tiling diagrams’, the desired result follows by strong convergence. \square

We can now — finally — prove the main result of the paper: confluence modulo \sim_{hc} .

Theorem 4.17. *Fully-extended, orthogonal iCRSs are confluent modulo \sim_{hc} .*

Proof. Let $s \sim_{hc} t$ and assume that $s \twoheadrightarrow s'$ and $t \twoheadrightarrow t'$. Consider the following diagram:

$$\begin{array}{c}
 s \sim_{hc} t \equiv t \\
 \swarrow \quad \downarrow \quad \searrow \quad \downarrow \\
 \begin{array}{ccc}
 s' & \sim_{hc} & t'_0 \\
 \downarrow \text{out} & & \downarrow \text{out} \\
 s'' & \sim_{hc} & t_0^*
 \end{array}
 \quad
 \begin{array}{ccc}
 t_1^* & \sim_{hc} & t' \\
 \downarrow \text{out} & & \downarrow \text{out} \\
 t_1^* & \sim_{hc} & t''
 \end{array}
 \end{array}$$

(1) $s \twoheadrightarrow s'$, (2) $t \twoheadrightarrow t'$, (3) $t \twoheadrightarrow t_1^*$, (4) $s' \twoheadrightarrow s''$, (5) $t' \twoheadrightarrow t''$.

In the diagram, (1) and (2) exist by Lemma 4.14 and (3) exists by Lemma 4.16. Moreover, (4) and (5) also exist by Lemma 4.14. The result now follows by the diagram and transitivity of \sim_{hc} . \square

Example 4.18. The example iCRS from the introduction is confluent modulo \sim_{hc} as it is orthogonal and fully-extended:

$$\begin{aligned}
 \text{map}([z]F(z), \text{cons}(X, XS)) &\rightarrow \text{cons}(F(X), \text{map}([z]F(z), XS)) \\
 \text{map}([z]F(z), \text{nil}) &\rightarrow \text{nil} \\
 \text{hd}(\text{cons}(X, XS)) &\rightarrow X \\
 \text{tl}(\text{cons}(X, XS)) &\rightarrow XS
 \end{aligned}$$

The iCRS consisting of the first two rules above is confluent, because it is confluent modulo \sim_{hc} and contains no collapsing rules.

The iCRS consisting of the infinite set of rules on the form

$$f^n([x]Z(x), Z') \rightarrow g(e^\omega, f^{n+1}([x]Z(x), Z(Z')))) \quad n \geq 1$$

is confluent modulo \sim_{hc} (and confluent, as it does not contain collapsing rules).

4.3. Almost non-collapsingness. We would like to have a characterisation of confluence that appeals only to the syntax of iCRSs without any need to consider equality modulo some relation. The first correct, fundamental confluence result for iTRSs [8] stated that an orthogonal iTRSs is confluent iff it has the property of being ‘almost non-collapsing’: There is at most one rule that is collapsing and the variable at the root of the right-hand side of that rule is the only variable occurring in the left-hand side of that rule.

Unfortunately, this concept does not carry over trivially to iCRSs, when replacing the variables from iTRSs by meta-variables:

Example 4.19. Consider the following rewrite rule, which is almost non-collapsing in the above sense:

$$f([x]Z(x)) \rightarrow Z(f([x]Z(x))).$$

The term $f([x]f([y]x))$ gives rise to the finite reduction

$$f([x]f([y]x)) \rightarrow f([x]x),$$

the final term of which reduces only to itself. However, the following reduction of length ω also exists:

$$f([x]f([y]x)) \rightarrow f([y]f([x]f([y]x))) \rightarrow f([y]f([y]f([x]f([y]x)))) \rightarrow \dots s,$$

where s is solution of the recursive equation $s = f([y]s)$, which is again a term which only reduces to itself. Hence, $f([x]f([y]x))$ reduces to two different terms that only reduce to themselves. In other words, the considered ‘almost non-collapsing’ rewrite rule defines a non-confluent iCRS.

We currently do not know how to give a precise characterisation of the class of confluent iCRSs. From the above example, it is clear that almost non-collapsingness alone does not suffice. It is plausible that the criterion for confluence will be undecidable, even for the class of iCRSs containing only a finite number of rules, all of which have finite right-hand-sides. The above example bears witness of this: It crucially depends on the term $f([x]x)$ being reachable from itself and reachability is of course in general undecidable.

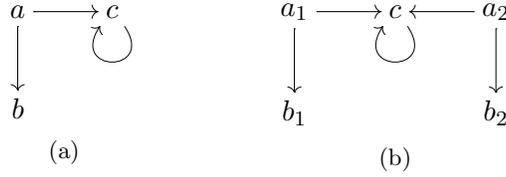


Figure 6: Counterexamples to the reverse of Proposition 5.3

5. NORMAL FORM PROPERTIES

In this section we consider normal forms of iCRSs:

Definition 5.1. A term in an iCRS is a *normal form* if no redexes occur in the term.

The following properties relate normal forms and reductions. The properties extend their usual finitary counterparts to infinitary rewriting. Ample motivation for the formulation of the properties can be found in [8]. In the definition, $(\leftarrow \cdot \rightarrow)^*$ denotes the symmetric, transitive, reflexive closure of \rightarrow .

Definition 5.2. Define the following:

- An iCRS has the *normal form property (NF)* if $s (\leftarrow \cdot \rightarrow)^* t$ with t a normal form implies $s \rightarrow t$.
- An iCRS has the *unique normal form property (UN)* if $s (\leftarrow \cdot \rightarrow)^* t$ with s and t normal forms implies $s = t$.
- An iCRS has the *unique normal form property with respect to reduction (UN^\rightarrow)* if $t \leftarrow s \rightarrow t'$ with t and t' normal forms implies $t = t'$.

By the definitions we immediately have:

Proposition 5.3. *NF implies UN, and UN implies UN^\rightarrow .* □

The converse implications of those above do not hold. This can be witnessed by the rewrite systems depicted in Figure 6.

In Figure 6(a) we give a counterexample refuting that UN implies NF: As b is the only normal form, next to all variables, UN is immediate. However, NF does not hold, as there is no reduction $c \rightarrow b$. The rewrite system in Figure 6(b) refutes that UN^\rightarrow implies UN: Since b_1 is the only normal form of a_1 with respect to reduction and since b_2 the only normal form of a_2 , UN^\rightarrow is immediate. However, UN does not hold, as we have $b_1 (\leftarrow \cdot \rightarrow)^* b_2$, while $b_1 \neq b_2$.

The following lemma relates confluence modulo hypercollapsing subterms with the three properties introduced above.

Lemma 5.4. *If an iCRS is confluent modulo hypercollapsing subterms, then NF, UN, and UN^\rightarrow hold.*

Proof. Let $s (\leftarrow \cdot \rightarrow)^* t$ with t a normal form. By induction on the number of changes in the direction of the reductions in $s (\leftarrow \cdot \rightarrow)^* t$ and confluence modulo hypercollapsing subterms it follows that s reduces to a term t' such that $t \sim_{hc} t'$. As no hypercollapsing subterms occur in normal forms, we have $t = t'$. Hence, NF holds and UN and UN^\rightarrow follow by Proposition 5.3. □

In the above proof, confluence is easily substituted for confluence modulo hypercollapsing subterms, yielding the traditional result from finitary rewriting stating that confluence implies NF, UN, and UN^\rightarrow . Moreover, as fully-extended, orthogonal iCRSs are confluent modulo hypercollapsing subterms by Theorem 4.17, the above lemma also gives an affirmative answer to the conjecture posed in [16] stating that fully-extended, orthogonal iCRSs satisfy UN^\rightarrow .

It is not the case that NF implies confluence modulo hypercollapsing subterms. To see this, consider the following four rewrite rules:

$$\begin{array}{ll} a \rightarrow f(b) & b \rightarrow b \\ a \rightarrow g(c) & c \rightarrow c \end{array}$$

No term in which a redex occurs has a normal form. Hence, NF is immediate. However, confluence modulo hypercollapsing subterms does not hold, as a reduces to $f(b)$ and $g(c)$ — both of which only reduce to themselves — and as $f(b) \not\sim_{hc} g(c)$.

6. CONCLUSION AND SUGGESTIONS FOR FUTURE WORK

We have extended confluence modulo identification of hypercollapsing subterms to higher-order infinitary rewriting by employing the proof techniques of earlier papers in the series on iCRSs as well as extending the known proof methods from [7]. Our results properly generalise similar results for iTRSs and $i\lambda c$, and the paper develops and extends the proof methods employed in earlier papers on these subjects.

Two major open questions related to confluence of iCRSs and higher-order infinitary rewriting in general remain. We invite the reader to consider these:

- Can a characterisation be given of the subclass of confluent iCRSs, generalising the first-order result that almost-non-collapsing systems are confluent? As we reason in Section 4.3, a generalisation will likely not be easy to come by.
- The current proof of confluence modulo hypercollapsing subterms requires orthogonality. Is it possible to replace orthogonality by weak orthogonality?

In the greater context of infinitary rewriting, this paper is part of an account of the general theory of iCRSs. We believe that the results and proof methods laid out will contribute to the further development of infinitary rewriting and equational reasoning involving infinite terms.

ACKNOWLEDGEMENT

The authors extend their thanks to the anonymous referees for their diligent work and many comments that have led to substantial improvements in the readability of the paper.

REFERENCES

- [1] H. P. Barendregt. *The Lambda Calculus: Its Syntax and Semantics*. Elsevier Science, revised edition, 1985.
- [2] B. Courcelle. Fundamental properties of infinite trees. *Theoretical Computer Science*, 25(2):95–169, 1983.
- [3] J. Glauert and Z. Khasidashvili. Relative normalization in orthogonal expression reduction systems. In *Proceedings of the International Workshop on Conditional (and Typed) Term Rewriting Systems (CTRS '94)*, volume 968 of *Lecture Notes in Computer Science*, pages 144–165. Springer-Verlag, 1994.

- [4] M. Hanus and C. Prehofer. Higher-order narrowing with definitional trees. In *Proceedings of the 7th International Conference on Rewriting Techniques and Applications (RTA '96)*, volume 1103 of *Lecture Notes in Computer Science*, pages 138–152. Springer-Verlag, 1996.
- [5] S. Kahrs. Compilation of combinatory reduction systems. In *Proceedings of the 1st International Workshop on Higher-Order Algebra, Logic, and Term Rewriting (HOA '93)*, volume 816 of *Lecture Notes in Computer Science*, pages 169–188. Springer-Verlag, 1993.
- [6] J. R. Kennaway, J. W. Klop, M. R. Sleep, and F.-J. de Vries. Infinitary lambda calculus. *Theoretical Computer Science*, 175(1):93–125, 1997.
- [7] R. Kennaway and F.-J. de Vries. Infinitary rewriting. In Terese [20], Chapter 12.
- [8] R. Kennaway, J. W. Klop, R. Sleep, and F.-J. de Vries. Transfinite reductions in orthogonal term rewriting systems. *Information and Computation*, 119(1):18–38, 1995.
- [9] J. Ketema. On normalisation of infinitary combinatory reduction systems. In *Proceedings of the 19th International Conference on Rewriting Techniques and Applications (RTA 2008)*, volume 5117 of *Lecture Notes in Computer Science*, pages 172–186. Springer-Verlag, 2008.
- [10] J. Ketema and J. G. Simonsen. Infinitary combinatory reduction systems. In *Proceedings of the 16th International Conference on Rewriting Techniques and Applications (RTA 2005)*, volume 3467 of *Lecture Notes in Computer Science*, pages 438–452. Springer-Verlag, 2005.
- [11] J. Ketema and J. G. Simonsen. On confluence of infinitary combinatory reduction systems. In *Proceedings of the 12th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2005)*, volume 3835 of *Lecture Notes in Artificial Intelligence*, pages 199–214. Springer-Verlag, 2005.
- [12] J. Ketema and J. G. Simonsen. Infinitary combinatory reduction systems: Normalising reduction strategies, 2008. Draft. Submitted for journal publication.
- [13] J. Ketema and J. G. Simonsen. Infinitary combinatory reduction systems, 2009. Draft. Submitted for journal publication.
- [14] Z. Khasidashvili. Beta-reductions and beta-developments of lambda-terms with the least number of steps. In *Proceedings of the the International Conference in Computer Logic (COLOG '88)*, volume 417 of *Lecture Notes in Computer Science*, pages 105–111. Springer-Verlag, 1988.
- [15] J. W. Klop. *Combinatory Reduction Systems*. PhD thesis, Rijksuniversiteit Utrecht, 1980.
- [16] J. W. Klop and R. de Vrijer. Infinitary normalization. In S. N. Artëmov, H. Barringer, A. S. d'Avila Garcez, L. C. Lamb, and J. Woods, editors, *We Will Show Them: Essays in Honour of Dov Gabbay*, volume 2, pages 169–192. College Publications, 2005.
- [17] J. W. Klop, V. van Oostrom, and F. van Raamsdonk. Combinatory reduction systems: introduction and survey. *Theoretical Computer Science*, 121(1 & 2):279–308, 1993.
- [18] A. Middeldorp. Call by need computations to root-stable form. In *Proceedings of the 24th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '97)*, pages 94–105, 1997.
- [19] R. C. Sekar and I. V. Ramakrishnan. Programming in equational logic: beyond strong sequentiality. *Information and Computation*, 104(1):78–109, 1993.
- [20] Terese, editor. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2003.
- [21] V. van Oostrom. Higher-order families. In *Proceedings of the 7th International Conference on Rewriting Techniques and Applications (RTA '96)*, volume 1103 of *Lecture Notes in Computer Science*, pages 392–407. Springer-Verlag, 1996.
- [22] V. van Oostrom. Normalisation in weakly orthogonal rewriting. In *Proceedings of the 10th International Conference on Rewriting Techniques and Applications (RTA '99)*, volume 1631 of *Lecture Notes in Computer Science*, pages 60–74. Springer-Verlag, 1999.
- [23] F. van Raamsdonk. Higher-order rewriting. In Terese [20], Chapter 11.

APPENDIX A. PROOF OF LEMMA 2.26

We prove Lemma 2.26.

Proof. Suppose that a variable bound by an abstraction in the redex pattern of u_α occurs in v_α . We reason by transfinite induction on α , the length of the reduction $s_0 \rightarrow^\alpha s_\alpha$. In case $\alpha = 0$, the result is immediate since bound variables can only occur below the abstraction by which they are bound and since subterms are substituted for the variables that occur in the subterm at the position of u_0 in s_0 .

In case α is a successor ordinal, suppose either that (a) $p_0 \not\leq q_0$, or that (b) $q_\alpha \geq p_\alpha$, but q_α does not occur in the reduct $s_\alpha|_{p_\alpha}$ of $s_0|_{p_0}$. We have the following:

- In case $p_0 \not\leq q_0$, we have that v_0 does not occur below u_0 . Hence, for some $\beta < \alpha$ a nesting is created by contracting a redex at a prefix position of p_β and q_β and, by definition of valuations, if a variable is bound by the redex pattern of $u_{\beta+1}$, then it cannot occur in $v_{\beta+1}$.
- In case $q_\alpha \geq p_\alpha$, but q_α does not occur in the reduct of $s_0|_{p_0}$ at p_α in s_α , there is a position $p_\alpha < p' \leq q_\alpha$ in s_α such that $s|_{p'}$ is a reduct of a subterm not strictly below $s_0|_{p_0}$. Hence, for some $\beta < \alpha$ a nesting is created by contracting a redex at a prefix position of p_β and q_β . In the term $s_{\beta+1}$ a residual of u_β occurs below $u_{\beta+1}$ and above $v_{\beta+1}$ and, by definition of valuations, if a variable is bound by the redex pattern of $u_{\beta+1}$, then it cannot occur in the residual and, hence, in $v_{\beta+1}$.

Thus, in both cases it follows for some $\beta < \gamma < \alpha$ that $s_\gamma \rightarrow s_{\gamma+1}$ nests a variable bound by u_γ in v_γ . By the definition of valuations, we have for the redex contracted in $s_\gamma \rightarrow s_{\gamma+1}$, say u'_γ at position p'_γ , that

- (1) a variable bound by an abstraction in the redex pattern of u_γ occurs in u'_γ , and that
- (2) a variable bound by an abstraction in the redex pattern of u'_γ occurs in v_γ .

Since it follows by assumption that u'_γ is the residual of a redex u'_δ in s_δ for all $\delta \leq \gamma$, we have by the induction hypothesis that $p'_\delta < q_\delta$ for all $\delta \leq \gamma$. Hence, for every $s_\delta \rightarrow s_{\delta+1}$ with $\delta \leq \gamma$ the contracted redex is a residual of a redex in s_0 in case it occurs at a prefix position of p'_γ . By the induction hypothesis we now have:

- In case $p_0 \not\leq q_0$, it follows that $p_0 < p'_0$ and $p'_0 < q_0$. Hence, $p_0 < q_0$, a contradiction.
- In case $q_\alpha \geq p_\alpha$, but q_α does not occur in the reduct of $s_0|_{p_0}$ at p_α in s_α , it follows that $u'_{\beta+1}$ occurs in the reduct of $s_0|_{p_0}$ at $p_{\beta+1}$ in $s_{\beta+1}$ and that $v_{\beta+1}$ occurs in the reduct of $s_0|_{p'_0}$ at $p'_{\beta+1}$ in $s_{\beta+1}$. Hence, $v_{\beta+1}$ occurs in the reduct of $s_0|_{p_0}$ at $p_{\beta+1}$ in $s_{\beta+1}$, a contradiction.

Hence, the result follows if α is a successor ordinal.

In case α is a limit ordinal, the result is immediate by strong convergence and the induction hypothesis, since residuals occur at finite depth. \square