Infinitary Combinatory Reduction Systems

Jeroen Ketema[1] and Jakob Grue Simonsen[2]

[1] Department of Computer Science, Vrije Universiteit Amsterdam
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands
jketema@cs.vu.nl

[2] Department of Computer Science, University of Copenhagen (DIKU)
Universitetsparken 1, DK-2100 Copenhagen Ø, Denmark
simonsen@diku.dk

Abstract. We define infinitary combinatory reduction systems (iCRSs). This provides the first extension of infinitary rewriting to higher-order rewriting. We lift two well-known results from infinitary term rewriting systems and infinitary λ-calculus to iCRSs:

1. every reduction sequence in a fully-extended left-linear iCRS is compressible to a reduction sequence of length at most ω, and
2. every complete development of the same set of redexes in an orthogonal iCRS ends in the same term.

1 Introduction

One of the main reasons for the initial research in infinitary rewriting was to have a model of lazy or stream-based programming languages easily accessible to people familiar with term rewriting. Two notions of infinitary rewriting were developed: infinitary (first-order) term rewriting systems (iTRSs) [1–3] and infinitary λ-calculus (iλc) [3–5]. However, the standard notion of rewriting employed to model higher-order programs is *higher-order* rewriting, and thus goes beyond λ-calculus. The absence of a general notion of infinitary higher-order rewriting thus constitutes a gap in the arsenal of the rewriting theorist bent on modelling lazy or stream-based languages.

In the present paper we aim to plug this gap by investigating infinitary higher-order rewriting.

As for iTRSs and iλc some finitary system needs to be chosen as a starting point. We choose the notion of higher-order rewriting most familiar to the authors, namely combinatory reduction systems (CRSs) [3, 6, 7].

The definition of infinitary combinatory reduction systems (iCRSs) consists of a combination of the usual four-stage definition of CRSs and the corresponding four-stage definition of iTRSs and iλc:

CRSs	**iTRSs/iλc**
1a. Meta-terms	
1b. Terms	1. Infinite terms
2. Substitutions	2. Substitutions
3. Rewrite rules	3. Rewrite rules
4. Rewrite relation	4. Rewrite relation

Given the definition of iCRSs, we seek to answer two of the most pertinent questions asked for any notion of infinitary rewriting:

1. Are reduction sequences compressible to reduction sequences of length at most ω?
2. Do complete developments of the same set of redexes end in the same term?

For iTRSs these questions have positive answers under assumption of respectively left-linearity and orthogonality. For iλc the same holds as long as the η-rule is not introduced. Apart from the definition of iCRSs, the main contribution of this paper is that similar positive answers can be given in the case of iCRSs.

The remainder of this paper is organised as follows. In Section 2 we give some preliminary definitions, and in Section 3, we define infinite (meta-)terms and substitutions. Thereafter, in Section 4 we define infinitary rewriting and prove compression, and in Section 5 we investigate complete developments. Finally, in Section 6 we give directions for further research.

2 Preliminaries

Prior knowledge of CRSs [7] and infinitary rewriting [3] is not required, but will greatly improve the reader's understanding of the text.

Throughout the paper 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. Countably infinite sets are sufficient, given that we can employ 'Hilbert hotel'-style renaming. We denote the first infinite ordinal by ω, and arbitrary ordinals by $\alpha, \beta, \gamma, \ldots$.

The set of *finite meta-terms* is defined as follows:

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

A finite meta-term of the form $[x]s$ is called an *abstraction*. Each occurrence of the variable x in s is *bound* in $[x]s$. If s is a finite meta-term, we denote by $root(s)$ the root symbol of s.

The *set of positions* of a finite meta-term s, denoted $\mathcal{P}os(s)$, is the set of finite strings over $\mathbb{N}$, with ϵ the empty string, such that:

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

Given $p, q \in \mathcal{P}os(s)$, we say that p is a *prefix* of q, denoted $p \leq q$, if there exists an $r \in \mathcal{P}os(s)$ such that $p \cdot r = q$. If $r \neq \epsilon$, then we say that the prefix is *strict* and we write $p < q$. Moreover, if neither $p < q$ nor $q < p$, then we say that p and q are *parallel*, which we denote $p \parallel q$. We denote by $s|_p$ the subterm of s at position p.

3 (Meta-)Terms and Substitutions

In iTRSs and iλc, terms are defined by means of introducing a metric on the set of finite terms and subsequently taking the completion of the metric. That is, taking the least set of objects containing the set finite terms such that every Cauchy sequence converges [2,4,8]. Intuitively, in such a metric, two terms s and t are close to each other if the first 'conflict' between them occurs 'deep' according to some depth measure. In iTRSs, a conflict is a position p such that $root(s|_p) \neq root(t|_p)$. In i$\lambda$c, a conflict is defined similarly, but also takes into account α-equivalence. The metric, denoted $d(s, t)$, is defined as 0 when no conflict occurs between s and t and otherwise as 2^{-k}, where k denotes the minimal depth such that a conflict occurs between s and t.

To define terms and meta-terms for iCRSs, we first define the notions of a conflict and α-equivalence for finite meta-terms. In the definition we denote by $s[x \to y]$ the replacement in s of the occurrences of the free variable x by the variable y.

Definition 3.1. *Let s and t be finite meta-terms. A* conflict *of s and t is a position $p \in \mathcal{P}os(s) \cap \mathcal{P}os(t)$ such that:*

1. *if $p = \epsilon$, then $root(s) \neq root(t)$,*
2. *if $p = i \cdot q$ for $i \geq 1$, then $root(s) = root(t)$ and q is a conflict of $s|_i$ and $t|_i$,*
3. *if $p = 0 \cdot q$, then $s = [x_1]s'$ and $t = [x_2]t'$ and q is a conflict of $s'[x_1 \to y]$ and $t'[x_2 \to y]$, where y does not occur in either s' or t'.*

The finite meta-terms s and t are α-equivalent *if no conflict exists [4].*

We next define the depth measure D.

Definition 3.2. *Let s be a meta-term and $p \in \mathcal{P}os(s)$. Define:*

$$\begin{aligned} D(s, \epsilon) &= 0 \\ D(Z(t_1, \ldots, t_n), i \cdot p') &= D(t_i, p') \\ D([x]t, 0 \cdot p') &= 1 + D(t, p') \\ D(f(t_1, \ldots, t_n), i \cdot p') &= 1 + D(t_i, p') \end{aligned}$$

Note that meta-variables are not counted by D. Changing the second clause to $D(Z(t_1, \ldots, t_n), i \cdot p') = 1 + D(t_i, p')$ yields the 'usual' depth measure, which counts the number of symbols in a position.

The measure D is employed in the definition of the metric, which is defined precisely as in the case of iTRSs and iλc.

Definition 3.3. *Let s and t be meta-terms. The metric d is defined as:*

$$d(s,t) = \begin{cases} 0 & \textit{if } s \textit{ and } t \textit{ are } \alpha\textit{-equivalent} \\ 2^{-k} & \textit{otherwise,} \end{cases}$$

where k is the minimal depth with respect to the measure D such that a conflict occurs between s and t.

Following precisely the definition of terms in the case of iTRSs and iλc, we define the meta-terms.

Definition 3.4. *The set of* meta-terms *over a signature Σ is the metric completion of the set of finite meta-terms with respect to the metric d.*

Note that, by definition of metric completion, the set of finite meta-terms is a subset of the set of meta-terms.

The notions of a set of positions and a subterm of a finite meta-term carry over directly to the meta-terms, we use the same notation in both cases.

The metric completion allows precisely those meta-terms such that the depth measure D increases to infinity along all infinite paths in the meta-term. Thus, by the definition of D and d, no meta-term has a subterm s such that there exists an infinite string p over $\mathbb{N}$ with the property that each finite prefix q of p is a position of s with $root(s|_q)$ a meta-variable. Informally, *no meta-term has an infinite chain of meta-variables.*

Examples of candidate 'meta-terms' that are disallowed by the definition of meta-term are:

$$Z(Z(\ldots(Z(\ldots))))$$
$$Z_1(Z_2(\ldots(Z_n(\ldots))))$$

A construction that *is* allowed is an infinite number of *finite* chains of meta-variables 'guarded' by abstractions or function symbols. For example, the following is allowed:

$$[x_1]Z_1([x_2]Z_2(\ldots([x_n]Z_n(\ldots))))$$

If we had wanted to include 'meta-terms' with infinite chains of meta-variables we should have used the usual depth measure on finite meta-terms instead of the measure D.

We explain the reason for the exclusion of meta-terms with infinite chains of meta-variables after the definition of substitutions. The idea of the exclusion of certain meta-terms comes from iλc where it is possible to define subsets of the set of infinite λ-terms by slightly changing the notion of the depth measure on which the metric is based [4]. It is, for example, possible to define a subset in which no λ-terms with infinite chains of λ-abstractions occur, i.e., subterms of the form $\lambda x_1.\lambda x_2 \ldots \lambda x_n \ldots$ are disallowed.

The terms can now be defined as in the finite case [3,6,7]. The only difference is that meta-terms now occur in the definition instead of finite meta-terms.

Definition 3.5. *The set of* terms *is the largest subset of the set of meta-terms, such that no meta-variables occur in the meta-terms.*

Note that the definition of meta-terms, as defined by the measure D, only restricts meta-terms containing meta-variables, not meta-terms *without* meta-variables. Hence, the set of terms is independent of the use of either D in Definition 3.3 or the usual depth measure. As a consequence, both the set of (infinite) first-order terms and the set of (infinite) λ-terms are easily shown to be included in the set of terms.

We next define substitutions. The required definitions are the same as in the case of CRSs [3, 7], except that coinduction is employed instead of induction. This is identical to what is done in the case of iTRSs and iλc with respect to the finite systems they are based on. In the definitions we use $\boldsymbol{x}$ and $\boldsymbol{t}$ as a short-hands for respectively the sequences $x_1, \ldots, x_n$ and $t_1, \ldots, t_n$ with $n \geq 0$. We assume n fixed in the next two definitions.

Definition 3.6. *A* substitution *of the terms* $\boldsymbol{t}$ *for distinct variables* $\boldsymbol{x}$ *in a term* s*, denoted* $s[\boldsymbol{x} := \boldsymbol{t}]$*, is coinductively defined as:*

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

The above definition implicitly takes into account the variable convention [9] in the third clause to avoid the binding of free variables by the abstraction.

Definition 3.7. *An* n-ary substitute *is a mapping denoted* $\underline{\lambda} x_1, \ldots, x_n.s$ *or* $\underline{\lambda}\boldsymbol{x}.s$*, with* s *a term, such that:*

$$(\underline{\lambda}\boldsymbol{x}.s)(t_1, \ldots, t_n) = s[\boldsymbol{x} := \boldsymbol{t}]\,. \tag{1}$$

Reading Equation (1) from left to right gives rise to the rewrite rule

$$(\underline{\lambda}\boldsymbol{x}.s)(t_1, \ldots, t_n) \to s[\boldsymbol{x} := \boldsymbol{t}]\,.$$

This rule can be seen a *parallel* β*-rule*. That is, a variant of the β-rule from iλc which substitutes for multiple variables simultaneously. The root of $(\underline{\lambda}\boldsymbol{x}.s)$ is called the $\underline{\lambda}$-abstraction and the root of the left-hand side of the parallel β-rule is called the $\underline{\lambda}$-application.

Definition 3.8. *A* valuation $\bar{\sigma}$ *is an extension of a function* σ *which assigns* n*-ary substitutes to* n*-ary meta-variables. It is coinductively defined as:*

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

Similar to Definition 3.6, the above definition implicitly takes into account the variable convention in the second clause to avoid the binding of free variables by the abstraction.

Thus, applying a substitution means applying a valuation and proceeds in two steps: In the first step each subterm of the form $Z(t_1, \ldots, t_n)$ is replaced by a subterm of the form $(\underline{\lambda}\boldsymbol{x}.s)(t_1, \ldots, t_n)$. In the second step Equation (1) is applied to each subterm of the form $(\underline{\lambda}\boldsymbol{x}.s)(t_1, \ldots, t_n)$ as introduced in the first step.

In the light of the rewrite rule introduced just below Definition 3.7 the second step can be viewed as a complete development of the parallel β-redexes introduced in the first step. This is obviously a complete development in a variant of iλc. The variant has the parallel β-rule and a signature containing the $\underline{\lambda}$-application, the $\underline{\lambda}$-abstraction, the abstractions, the meta-variables, and the elements of Σ.

As in the finite case [6, Remark II.1.10.1], we need to prove that the application of a valuation to a meta-term yields a unique term.

Proposition 3.9. *Let s be a meta-term and $\bar{\sigma}$ a valuation. There exists a unique term that is the result of applying $\bar{\sigma}$ to s.*

Proof (Sketch). That the first step in applying $\bar{\sigma}$ to s has a unique result is an immediate consequence of being defined coinductively. We denote the result of the first step by s_σ. The set of parallel β-redexes in s_σ is denoted $\mathcal{U}$.

To prove that the second step also has a unique result we employ the rewriting terminology as introduced above. Although omitted, the definitions of a development and a complete development can be easily derived from the iλc definitions.

Note that to repeatedly rewrite the root of s_σ by means of the parallel β-redex, the root must look like

$$(\underline{\lambda}\boldsymbol{x}.x_i)(t_1, \ldots, t_n)\,,$$

with $1 \leq i \leq n$ and t_i again such a redex. This is only possible if there exists in s_σ an infinite chain of such redexes which starts at the root. However, this requires an infinite chain of meta-variables to be present in s, which is not allowed by the definition of meta-terms. Thus, the root can only be rewritten finitely often in a development. Applying the same reasoning to the roots of the subterms, gives that a complete development is obtained by reducing the redexes in $\mathcal{U}$ in an outside-in fashion. As all parallel β-redexes occur in $\mathcal{U}$ and as no $\underline{\lambda}$-applications and $\underline{\lambda}$-abstractions occur in s the result of the complete development, which we denote $\bar{\sigma}(s)$, is necessarily a term.

To show that each complete development ends in $\bar{\sigma}(s)$, note that we can view each parallel β-redex $(\underline{\lambda}x_1, \ldots, x_n.s)(t_1, \ldots, t_n)$ as a sequence of β-redexes:

$$(\lambda x_1(\ldots((\lambda x_n.s)t_n)\ldots))t_1\,.$$

This means that each complete development in our variant of iλc corresponds to a complete development in iλc extended with some function symbols. As each complete development in iλc ends in the same term, a result independent of added function symbols, the complete developments of the second step must also end in the same term. Hence, $\bar{\sigma}(s)$ is unique. □

Let us now see why we excluded 'meta-terms' with infinite chains of meta-variables from Definition 3.4. Consider the 'meta-term'

$$Z(Z(\ldots(Z(\ldots)))) \,.$$

Applying the valuation that assigns to Z the substitute $\underline{\lambda}x.x$ yields:

$$(\underline{\lambda}x.x)((\underline{\lambda}x.x)(\ldots((\underline{\lambda}x.x)(\ldots))))$$

which has no complete development, as no matter how many parallel β-redexes are contracted, it reduces only to itself and not to a term. This is inadequate, as rewrite steps in iCRSs need to relate terms to terms.

The previous problem does not depend on only a single meta-variable being present in the 'meta-term'. The same behaviour can occur with different meta-variables of different arities. In that case, we can define a valuation that assigns $\underline{\lambda}\boldsymbol{x}.y$ to each meta-variable Z in the 'meta-term' with y in $\boldsymbol{x}$ such that y corresponds to an argument of Z which is a chain of meta-variables.

The above 'meta-term' still has the nice property that it exhibits confluence with respect to the parallel β-rule. Unfortunately, there are 'meta-terms' that do not have this property. Consider a signature with constants a and b and also consider the 'meta-term'

$$Z(a, Z(b, Z(a, Z(b, Z(\ldots))))) \,.$$

Applying the valuation that assigns to Z the substitute $\underline{\lambda}xy.y$ yields the '$\underline{\lambda}$-term':

$$(\underline{\lambda}xy.y)(a, (\underline{\lambda}xy.y)(b, (\underline{\lambda}xy.y)(a, (\underline{\lambda}xy.y)(b, (\underline{\lambda}xy.y)(\ldots))))) \,,$$

which is also depicted in Figure 1. The term reduces by means of two different developments to the $\underline{\lambda}$-terms:

$$(\underline{\lambda}xy.y)(a, (\underline{\lambda}xy.y)(a, (\underline{\lambda}xy.y)(a, (\underline{\lambda}xy.y)(a, (\underline{\lambda}xy.y)(\ldots))))) \,,$$

as depicted in Figure 2, and:

$$(\underline{\lambda}xy.y)(b, (\underline{\lambda}xy.y)(b, (\underline{\lambda}xy.y)(b, (\underline{\lambda}xy.y)(b, (\underline{\lambda}xy.y)(\ldots))))) \,,$$

as depicted in Figure 3. These last two $\underline{\lambda}$-terms have no common reduct with respect to parallel β-reduction. They reduce only to themselves. Note that this problem also occurs in iλc [4, Section 4].

Concluding, when we allow 'meta-terms' with infinite chains of meta-variables we have two problems. First, substitution in such a 'meta-term' does not always yield a term. Second, substitution may yield distinct results, none of which are terms. We can overcome these problems by not allowing infinite chains of meta-variables to occur in meta-terms, as shown in Proposition 3.9.

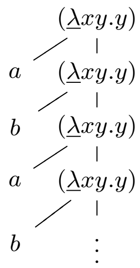

Fig. 1.

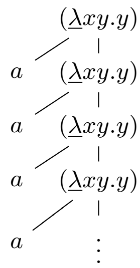

Fig. 2.

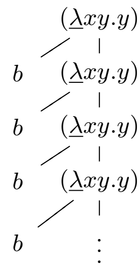

Fig. 3.

4 Infinitary Rewriting

We continue to combine the definitions of iTRSs and iλc and those of CRSs. We start with a definition that comes directly from CRS theory.

Definition 4.1. *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 its variables occur bound.*

We next define rewrite rules and iCRSs. In analogy to the rewrite rules of iTRSs, the definition is identical to the one in the finitary case, but without the finiteness restriction on the right-hand sides of the rewrite rules [1, 2].

Definition 4.2. *A* rewrite rule *is a pair (l, r), denoted $l \to r$, where l is a finite meta-term and r is a meta-term, such that:*

1. *l is a pattern and of the form $f(s_1, \ldots, s_n)$ with $f \in \Sigma$ of arity n,*
2. *all meta-variables that occur in r also occur in l, and*
3. *l and r are closed.*

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

As the rewrite rules of iTRSs and iλc only have finite chains of meta-variables when their rules are considered as rewrite rules in the above sense, it follows easily that iTRSs and iλc are iCRSs.

A context is a term over $\Sigma \cup \{\Box\}$ where $\Box$ is a fresh constant. One-hole contexts are defined in the usual way. We now define redexes and rewrite steps.

Definition 4.3. *Let $l \to r$ be a rewrite rule. Given a valuation $\bar{\sigma}$, the term $\bar{\sigma}(l)$ is called a $l \to r$*-redex. *If $s = C[\bar{\sigma}(l)]$ for some context $C[\Box]$ with $\bar{\sigma}(l)$ a $l \to r$-redex and p the position of the hole in $C[\Box]$, then an $l \to r$-redex, or simply a redex,* occurs *at position p and depth $D(s, p)$ in s. A* rewrite step *is a pair (s, t), denoted $s \to t$, such that a $l \to r$-redex occurs in $s = C[\bar{\sigma}(l)]$ and such that $t = C[\bar{\sigma}(r)]$.*

We can now define what a transfinite reduction sequence is. The definition copies the definition from iTRSs and iλc verbatim [2, 4].

Definition 4.4. *A* transfinite reduction sequence *of ordinal length α is a sequence of terms $(s_\beta)_{\beta<\alpha+1}$ such that $s_\beta \to s_{\beta+1}$ for all $\beta < \alpha$. For each rewrite step $s_\beta \to s_{\beta+1}$, let d_β denote the depth of the contracted redex. The reduction sequence is* weakly convergent *or* Cauchy convergent *if for every ordinal $\gamma \leq \alpha$ the distance between t_β and t_γ tends to 0 as β approaches γ from below. The reduction sequence is* strongly convergent *if it is weakly convergent and if d_β tends to infinity as β approaches γ from below.*

Notation 4.5. By $s \twoheadrightarrow^\alpha t$, respectively $s \twoheadrightarrow^{\leq\alpha} t$, we denote a *strongly convergent* transfinite reduction sequence of ordinal length α, respectively of ordinal length less than or equal to α. By $s \twoheadrightarrow t$ we denote a *strongly convergent* transfinite reduction sequence of arbitrary ordinal length and by $s \to^* t$ we denote a reduction sequence of finite length.

As in [2–4], we prefer to reason about strongly converging reduction sequences. This ensures that we can restrict our attention to reduction sequences of length at most ω by the so-called *compression property*. To prove the property we need the following lemma and definitions.

Lemma 4.6. *If $s \twoheadrightarrow t$, then the number of steps contracting redexes at depths less than $d \in \mathbb{N}$ is finite for any d.*

Proof. This is exactly the proof of [2, Lemma 3.5]. □

Definition 4.7. *A rewrite rule $l \to r$ is* left-linear*, if each meta-variable occurs at most once in l. Moreover, an iCRS is* left-linear *if all its rewrite rules are left-linear.*

Definition 4.8. *A pattern is* fully-extended *[10, 11], if, for each of its meta-variables Z, and each abstraction $[x]$ having Z in its scope, x is an argument of Z. Moreover, an iCRS is* fully-extended *if the left-hand sides of all rewrite rules are fully-extended.*

Left-linearity and fully-extendedness ensure no redex is created by either making two subterms equal in an infinite number of steps or by erasing some variable in an infinite number of steps.

Theorem 4.9 (Compression). *For every fully-extended, left-linear iCRS, if $s \twoheadrightarrow^\alpha t$, then $s \twoheadrightarrow^{\leq\omega} t$.*

Proof (Sketch). Let $s \twoheadrightarrow^\alpha t$, and proceed by ordinal induction on α. By [3, Theorem 12.7.1] it suffices to show that the theorem holds for $\alpha = \omega + 1$: The cases where α is 0, a limit ordinal, or a successor ordinal greater than $\omega + 1$ do not depend on the definition of rewriting.

For $\alpha = \omega+1$ it follows by Lemma 4.6 that we can write $s \twoheadrightarrow^\alpha t$ as $s \to^* s' \twoheadrightarrow^\omega s'' \to t$, such that all rewrite steps in $s' \twoheadrightarrow^\omega s''$ occur below the meta-variable positions of the redex contracted in the step of $s'' \to t$. By fully-extendedness and left-linearity it follows that a redex of which the redex contracted in $s'' \to t$ is a residual occurs in s'. Hence, we can contract the redex in s', which yields a term t'.

The result now follows if we can construct a strongly convergent reduction sequence $t' \twoheadrightarrow^{\leq \omega} t$. To construct such a reduction sequence, assume $t_0 = t'$ and construct for each $d > 0$ a reduction sequence $t_{d-1} \to^* t_d$ where all rewrite steps occur at depths greater or equal to $d - 1$, and where $d(t_d, t) \leq 2^{-d}$. That the construction of these reduction sequences is possible follows by a proof that is similar to the proof of compression for iλc [4]. Using the fact that only finite chains of meta-variables occur in meta-terms is essential to the proof. By the requirements on the constructed reduction sequences, it follows that $t_0 \to^* t_1 \to^* \ldots \to^* t_{d-1} \to^* t_d \to^* \ldots t$ is a strongly convergent reduction sequence of length at most ω. As $s \to^* t'$, we then have that $s \twoheadrightarrow^{\leq \omega} t$, as required. □

A complete proof of the above theorem is given in Appendix A.

The assumptions of left-linearity and fully-extendedness cannot be omitted from the previous theorem. In case of left-linearity this follows from the iTRS counterexample in [2]. In case of fully-extendedness this follows from the infinitary $\lambda\beta\eta$-calculus in which reduction sequences occur that are not compressible to reduction sequences of length at most ω [3, 4, 12]. The η-rule is not fully-extended.

5 Developments

In this section we prove that each complete development of the same set of redexes in an orthogonal iCRS ends in the same term. As all the left-hand sides of the rewrite rules in iCRSs are finite, the definition of orthogonality carries over immediately from CRSs.

Definition 5.1. *Let $R = \{l_i \to r_i \mid i \in I\}$ be a set of rewrite rules.*

1. *R is* non-overlapping *if it holds that:*
   - *each $l_i \to r_i$-redex that occurs at a position p in an $l_j \to r_j$-redex with $i \neq j$ occurs such that there exists a position $q \leq p$ with $q \in \mathcal{P}os(l_j)$ and $root(l_j|_p)$ a meta-variable,*
   - *likewise for $p \neq \epsilon$ and $i = j$.*
2. *R is* orthogonal *if it is left-linear and non-overlapping.*
3. *An iCRS is* orthogonal *if its set or rewrite rules is orthogonal.*

In the remainder of this section we assume an orthogonal iCRS, a term s, and a set $\mathcal{U}$ of redexes in s.

5.1 Descendants and Residuals

Before we can consider developments, we need to define descendants and residuals. The definition of descendant across a rewrite step $\bar{\sigma}(l) \to \bar{\sigma}(r)$ follows the definition of substitution, and is thus defined in two steps. The first step defines descendants in $\bar{\sigma}(r)$ where only the valuation is applied and not Equation (1). The second step defines descendants across application of Equation (1).

Given that the second step of the substitution is just a complete development in a variant of iλc, the second step in the definition of descendants is just a variant of descendants in iλc [3, 4]. For this reason, the step is not made explicit here.

We next give a definition of the first step. In the definition we denote by 0 the position of the subterm on the left-hand side of a $\underline{\lambda}$-application and also the position of the body of a $\underline{\lambda}$-abstraction. By $1, \ldots, n$ we denote the positions of the subterms on the right-hand side of the $\underline{\lambda}$-application. This means that $(\underline{\lambda}\boldsymbol{x}.s)(t_1, \ldots, t_n)|_0 = (\underline{\lambda}\boldsymbol{x}.s)$, $\underline{\lambda}\boldsymbol{x}.s|_0 = s$, and $Z(t_1, \ldots, t_n)|_i = (\underline{\lambda}\boldsymbol{x}.s)(t_1, \ldots, t_n)|_i = t_i$ for $1 \leq i \leq n$. We denote by $\bar{\sigma}(l) \to r_\sigma$ the rewrite step $\bar{\sigma}(l) \to \bar{\sigma}(r)$ when only the first step of the substitution applied to r.

Definition 5.2. *Let $l \to r$ be a rewrite rule, $\bar{\sigma}$ a valuation, and $p \in \mathcal{P}os(\bar{\sigma}(l))$. Suppose $u : \bar{\sigma}(l) \to r_\sigma$. The set $p/^1u$ is defined as follows:*

- *if a position $q \in \mathcal{P}os(l)$ exists such that $p = q \cdot q'$ and $root(l|_q) = Z$, then define $p/^1u = \{p' \cdot 0 \cdot 0 \cdot q' \mid p' \in P\}$ with $P = \{p' \mid root(r|_{p'}) = Z\}$,*
- *if no such position exists, then define $p/^1u = \emptyset$.*

Note that $\mathcal{P}os(r) \subseteq \mathcal{P}os(r_\sigma)$ by the notation of positions in subterms of the form $(\underline{\lambda}\boldsymbol{x}.s)(t_1, \ldots, t_n)$. From this it follows that $P \subseteq \mathcal{P}os(r_\sigma)$.

We can now give a complete definition of a descendant across a rewrite step.

Definition 5.3. *Let $u : C[\bar{\sigma}(l)] \to C[\bar{\sigma}(r)]$ be a rewrite step, such that p is the position of the hole in $C[\Box]$, and let $q \in \mathcal{P}os(C[\bar{\sigma}(l)])$. The set of* descendants *of q* across *u, denoted q/u, is defined as $q/u = \{q\}$ in case $p \parallel q$ or $p < q$. In case $q = p \cdot q'$, it is defined as $q/u = \{p \cdot q'' \mid p'' \in Q\}$, where Q is the set of descendants of $q'/^1u'$ with $u' : \bar{\sigma}(l) \to r_\sigma$ across complete development of the parallel β-redexes in r_σ.*

Descendants across a reduction sequence are defined as for iTRSs and iλc.

Definition 5.4. *Let $s_0 \twoheadrightarrow^\alpha s_\alpha$ and let $P \subseteq \mathcal{P}os(s_0)$. The set of* descendants *of P across $s_0 \twoheadrightarrow^\alpha s_\alpha$, denoted $P/(s_0 \twoheadrightarrow^\alpha s_\alpha)$, is defined as follows:*

- *if $\alpha = 0$, then $P/(s_0 \twoheadrightarrow^\alpha s_\alpha) = P$,*
- *if $\alpha = 1$, then $P/(s_0 \to s_1) = \bigcup_{p \in P} p/(s_0 \to s_1)$,*
- *if $\alpha = \beta + 1$, then $P/(s_0 \twoheadrightarrow^{\beta+1} s_{\beta+1}) = (P/(s_0 \twoheadrightarrow^\beta s_\beta))/(s_\beta \to s_{\beta+1})$,*
- *if α is a limit ordinal, then $p \in P/(s_0 \twoheadrightarrow^\alpha s_\alpha)$ iff $p \in P/(s_0 \twoheadrightarrow^\beta s_\beta)$ for all large enough $\beta < \alpha$.*

By orthogonality, if there exists a redex at a position p using a rewrite rule $l \to r$ that is not contracted in rewrite step and if p has descendants across the step, then there exists a redex at each descendant of p also employing the rule $l \to r$. Hence, there exists a well-defined notion of *residual* by strongly convergent reduction sequences. We overload the notation $\cdot/\cdot$ to denote both the descendant and the residual relation.

Let $s \to t$ by contraction of some redex u in s. We sometimes write $\mathcal{U}/u$ instead of $\mathcal{U}/(s \to t)$, where $\mathcal{U}$ is a set of redexes of s.

5.2 Complete Developments

We now define developments. Recall that we assume we are working in an orthogonal iCRS and that $\mathcal{U}$ is a set of redexes in a term s.

Definition 5.5. *A* development *of $\mathcal{U}$ is a strongly convergent reduction sequence such that each step contracts a residual of a redex in $\mathcal{U}$. A development $s \twoheadrightarrow t$ is* complete *if $\mathcal{U}/(s \twoheadrightarrow t) = \emptyset$.*

To prove that each complete development of the same set of redexes ends in the same term, we extend the technique of the Finite Jumps Developments Theorem [3] to orthogonal iCRSs. The theorem employs notions of paths and path projections. In essence, paths and path projections are 'walks' through terms starting at the root and proceeding to greater and greater depths. An important property of paths and path projections is that when a walk encounters a redex to be contracted in a development, a 'jump' is made to the right-hand side of the employed rewrite rule. It continues there until a meta-variable is encountered, at which point a jump back to the original term occurs.

In the following definition, we denote by p_u the position of the redex u in s. Moreover, we say that a variable x is *bound by a redex u* when x is bound by an abstraction $[x]$ which occurs in the left-hand side of the rewrite rule employed in u.

Definition 5.6. *A* path *of s with respect to $\mathcal{U}$ is a sequence of nodes and edges. Each node is labelled either (s, p) with $p \in \mathcal{P}os(s)$ or (r, p, p_u) with r a right-hand side of a rewrite rule, $p \in \mathcal{P}os(r)$, and $u \in \mathcal{U}$. Each directed edge is either unlabelled or labelled with an element of $\mathbb{N}$.*

Every path starts with a node labelled (s, ϵ). If a node n of a path is labelled (s, p) and if it has an outgoing edge to a node n', then:

1. *if the subterm at p is not a redex in $\mathcal{U}$, then for some $i \in \mathcal{P}os(s|_p) \cap \mathbb{N}$ the node n' is labelled $(s, p \cdot i)$ and the edge from n to n' is labelled i,*
2. *if the subterm at p is a redex $u \in \mathcal{U}$ with $l \to r$ the employed rewrite rule, then the node n' is labelled (r, ϵ, p_u) and the edge from n to n' is unlabelled,*
3. *if $s|_p$ is a variable x bound by a redex $u \in \mathcal{U}$ with $l \to r$ the employed rewrite rule, then the node n' is labelled $(r, p' \cdot i, p_u)$ and the edge from n to n' is unlabelled, such that (r, p', p_u) was the last node before n with p_u, $root(r|_{p'}) = Z$, the unique position of Z in l is q, and $l|_{q \cdot i} = x$.*

If a node n of a path is labelled (r, p, p_u) and if it has an outgoing edge to a node n', then:

1. *if $root(r|_p)$ is not a meta-variable, then for some $i \in \mathcal{P}os(r|_p) \cap \mathbb{N}$ the node n' is labelled $(r, p \cdot i, p_u)$ and the edge from n to n' is labelled i,*
2. *if $root(r|_p)$ is a meta-variable Z, then the node n' is labelled $(s, p_u \cdot q')$ and the edge from n to n' is unlabelled, such that $l \to r$ is the rewrite rule employed in u and such that q' is the unique position of Z in l.*

We say that a path is *maximal* if it is not a proper prefix of another path. We write a path Π as a (possibly infinite) sequence of alternating nodes and edges $\Pi = n_1 e_1 n_2 \ldots$.

Definition 5.7. *Let $\Pi = n_1 e_1 n_2 \ldots$ be a path of s with respect to $\mathcal{U}$. The* path projection *of Π is a sequence of alternating nodes and edges $\phi(\Pi) = \phi(n_1)\phi(e_1)\phi(n_2)\ldots$ such that for each node n in Π:*

1. *if n is labelled (s,p), then $\phi(n)$ is unlabelled if $root(s|_p)$ is a redex in $\mathcal{U}$ or a variable bound by some redex in $\mathcal{U}$ and it is labelled $root(s|_p)$ otherwise,*
2. *if n is labelled (r,p,q), then $\phi(n)$ is unlabelled if $root(r|_p)$ is a meta-variable and it is labelled $root(r|_p)$ otherwise.*

For each edge e, if e is labelled i, then $\phi(e)$ has the same label, and if e is unlabelled, then $\phi(e)$ is labelled ϵ.

Example 5.8. Consider the iCRS with the following rewrite rule $l \to r$:

$$f([x]Z(x), Z') \to Z(g(Z(Z'))) \,.$$

Also, consider the terms $s = f([x]g(x), a)$ and $t = g(g(g(a)))$, the meta-term $r = Z(g(Z(Z')))$, and the set $\mathcal{U}$ containing the only redex in s. Obviously, $s \to t$ is a complete development.

The term s has one maximal path with respect to $\mathcal{U}$:

$$\begin{aligned}(s,\epsilon) \to (r,\epsilon,\epsilon) \to (s,10) \xrightarrow{1} (s,101) \to (r,1,\epsilon) \xrightarrow{1} (r,11,\epsilon) \\ \to (s,10) \xrightarrow{1} (s,101) \to (r,111,\epsilon) \to (s,2)\end{aligned}$$

The term t has one maximal path with respect to $\mathcal{U}/\mathcal{U} = \emptyset$:

$$(t,\epsilon) \xrightarrow{1} (t,1) \xrightarrow{1} (t,11) \xrightarrow{1} (t,111) \,.$$

The path projections of the maximal paths are respectively

$$\cdot \xrightarrow{\epsilon} \cdot \xrightarrow{\epsilon} g \xrightarrow{1} \cdot \xrightarrow{\epsilon} g \xrightarrow{1} \cdot \xrightarrow{\epsilon} g \xrightarrow{1} \cdot \xrightarrow{\epsilon} \cdot \xrightarrow{\epsilon} a$$

and

$$g \xrightarrow{1} g \xrightarrow{1} g \xrightarrow{1} a \,.$$

Let $\mathcal{P}(s,\mathcal{U})$ denote the set of path projections of *maximal paths* of s with respect to $\mathcal{U}$. The following result can be witnessed in the above example.

Lemma 5.9. *Let $u \in \mathcal{U}$ and let $s \to t$ be the rewrite step contracting u. There is a bijection between $\mathcal{P}(s,\mathcal{U})$ and $\mathcal{P}(t,\mathcal{U}/u)$. Given a path projection $\phi(\Pi) \in \mathcal{P}(s,\mathcal{U})$, its image under the bijection is acquired from $\phi(\Pi)$ by deleting finite sequences of unlabelled nodes and ϵ-labelled edges from $\phi(\Pi)$.*

Proof (Sketch). By straightforwardly, but very tediously, tracing through the construction of paths, it is evident that the set of maximal paths of t with respect to $\mathcal{U}/u$ can be obtained from the set of maximal paths of s with respect to $\mathcal{U}$ by replacing or deleting nodes of the form (r, p, p_u). If a maximal path of t is obtained from a maximal path of s in this way, then they have identical path projections, except that sequences of ϵ-labelled edges and unlabelled nodes may have been deleted (due to the contraction of u). This establishes the desired bijection. It is easy to see that the sequences of deleted edges can only be infinite if there is an infinite chain of meta-variables in the right-hand side of the rule of u, which is impossible by the definition of meta-terms. □

A complete proof of the above lemma is given in Appendix B.

We next define a property for sets $\mathcal{P}(s,\mathcal{U})$: the finite jumps property. We also define some terminology to relate a term to a set $\mathcal{P}(s,\mathcal{U})$.

Definition 5.10. *The set $\mathcal{U}$ has the* finite jumps property *if no path projection occurring in $\mathcal{P}(s,\mathcal{U})$ contains an infinite sequence of unlabelled nodes and ϵ-labelled edges. Moreover, a term t* matches *$\mathcal{P}(s,\mathcal{U})$ if, for all $\phi(\Pi) \in \mathcal{P}(s,\mathcal{U})$ and all prefixes of $\phi(\Pi)$ ending in a node n labelled f, it holds that $root(t|_p) = f$, where p is the concatenation of the edge labels in the prefix (starting at the first node of $\phi(\Pi)$ and ending at $\phi(n)$).*

We have the following.

Proposition 5.11. *If $\mathcal{U}$ has the finite jumps property, then there exists a unique term, denoted $\mathcal{T}(s,\mathcal{U})$, that matches $\mathcal{P}(s,\mathcal{U})$.*

Proof. Let $\mathcal{P}_p(s,\mathcal{U})$ denote the set of all prefixes of path projections in $\mathcal{P}(s,\mathcal{U})$ such that the concatenation of the edge labels for each prefix is p and such that each prefix ends in a labelled node.

Consider $\mathcal{P}_\epsilon(s,\mathcal{U})$. By the finite jumps property $\mathcal{P}_\epsilon(s,\mathcal{U})$ is non empty and by the definition of paths, $\mathcal{P}_\epsilon(s,\mathcal{U})$ has at most one element. Hence, $\mathcal{P}_\epsilon(s,\mathcal{U})$ is a singleton set. Suppose that the unique prefix in $\mathcal{P}_\epsilon(s,\mathcal{U})$ has f as the label of its only labelled node. It follows that t only matches $\mathcal{P}(s,\mathcal{U})$ if $root(t|_\epsilon) = f$.

Now suppose $\mathcal{P}_p(s,\mathcal{U})$ is a singleton set such that the final labelled node of the unique prefix in the set has label f, where f is either a variable, a function symbol of arity n or an abstraction. In the last two cases, consider $\mathcal{P}_{p\cdot i}(s,\mathcal{U})$ for either $1 \leq i \leq n$ or $i = 0$. By the finite jumps property, the definition of paths, and the fact that $\mathcal{P}_p(s,\mathcal{U})$ is a singleton set, we have that $\mathcal{P}_{p\cdot i}(s,\mathcal{U})$ is a singleton set. Suppose that the unique prefix in $\mathcal{P}_{p\cdot}(s,\mathcal{U})$ has g as the label of its final labelled node. It follows that t only matches $\mathcal{P}(s,\mathcal{U})$ if $root(t|_{p\cdot i}) = g$.

By the fact that all sets $\mathcal{P}_p(s,\mathcal{U})$ are singleton sets there exist terms that match $\mathcal{P}(s,\mathcal{U})$. Moreover, if t is such a term, then we have for all $p \in \mathcal{P}os(t)$ that $\mathcal{P}_p(s,\mathcal{U})$ exists and is a singleton set (look at all the prefixes of p), and if the final labelled node of the unique prefix in such a set has label f, then $root(t|_p) = f$. Hence, the term t is unique. □

The above proof is (almost) identical to the proof of Proposition 12.5.8 in [3].

We can now finally prove the Finite Jumps Developments Theorem:

Theorem 5.12 (Finite Jumps Developments Theorem). *If $\mathcal{U}$ has the finite jumps property, then:*

1. *every complete development of $\mathcal{U}$ ends in $\mathcal{T}(s,\mathcal{U})$,*
2. *for any $p \in \mathcal{P}os(s)$, the set of descendants of p by a complete development of $\mathcal{U}$ is independent of the complete development,*
3. *for any redex u of s, the set of residuals of u by a complete development of $\mathcal{U}$ is independent of the complete development, and*
4. *$\mathcal{U}$ has a complete development.*

Proof. (1) Suppose there is a complete development. We show, by means of ordinal induction, that for every s_α in the complete development with residuals $\mathcal{U}_\alpha = \mathcal{U}/(s \twoheadrightarrow s_\alpha)$ of $\mathcal{U}$, we have that $\mathcal{P}(s_\alpha,\mathcal{U}_\alpha)$ can be obtained from $\mathcal{P}(s,\mathcal{U})$ by deleting finite sequences of unlabelled nodes and ϵ-labelled edges from the elements of $\mathcal{P}(s,\mathcal{U})$.

For $s_0 = s$ this is immediate.

For $s_{\alpha+1}$, we have by the induction hypothesis that $\mathcal{P}(s_\alpha,\mathcal{U}_\alpha)$ can be obtained from $\mathcal{P}(s,\mathcal{U})$ by deleting finite sequences of unlabelled nodes and the ϵ-labelled edges from the elements of $\mathcal{P}(s,\mathcal{U})$. By Lemma 5.9, we have that $\mathcal{P}(s_{\alpha+1},\mathcal{U}_{\alpha+1})$ can be obtained from $\mathcal{P}(s_\alpha,\mathcal{U}_\alpha)$ by deleting finite sequences of unlabelled nodes and ϵ-labelled edges. Hence, $\mathcal{P}(s_{\alpha+1},\mathcal{U}_{\alpha+1})$ can also be obtained from $\mathcal{P}(s,\mathcal{U})$ by deleting finite sequences of unlabelled nodes and ϵ-labelled edges.

For s_α with α a limit ordinal, we have by strong convergence that $\mathcal{P}(s_\alpha,\mathcal{U}_\alpha)$ can be obtained from $\mathcal{P}(s,\mathcal{U})$ by deleting all unlabelled nodes and ϵ-labelled edges deleted in the previous steps. As $\mathcal{P}(s,\mathcal{U})$ has the finite jumps property, $\mathcal{P}(s_\alpha,\mathcal{U}_\alpha)$ can only be obtained by deleting finite sequences of unlabelled nodes and ϵ-labelled edges.

We now have that each $\mathcal{P}(s_\alpha,\mathcal{U}_\alpha)$ has the finite jumps property, as $\mathcal{P}(s,\mathcal{U})$ has the finite jumps property and as each $\mathcal{P}(s_\alpha,\mathcal{U}_\alpha)$ can be obtained by deleting only finite sequences of unlabelled nodes and ϵ-labelled edges.

By Theorem 5.11 it follows for each $\mathcal{P}(s_\alpha,\mathcal{U}_\alpha)$ that there is a unique term $\mathcal{T}(s_\alpha,\mathcal{U}_\alpha)$ that matches it. By inspection of the proof of Proposition 5.11 it is easy to see that the unlabelled nodes and ϵ-labelled edges are irrelevant for the construction of $\mathcal{T}(s_\alpha,\mathcal{U}_\alpha)$. Hence, $\mathcal{T}(s_\alpha,\mathcal{U}_\alpha) = \mathcal{T}(s,\mathcal{U})$ for all α. Moreover, as the above complete development was arbitrary we have that the final term of each complete development is $\mathcal{T}(s,\mathcal{U})$.

(2) Define a labelled version of the assumed iCRS analogous to [6, Section II.2]. The labelled version is orthogonal by [6, Proposition II.2.6]. Each reduction sequence in the labelled version corresponds to a reduction sequence in the original iCRS by removal of all labels. Moreover, given a reduction sequence in the original iCRS and a labelling for the initial term of the sequence, there exists a unique reduction sequence in the labelled version, such that removal of the labels gives the reduction sequence we started out with.

Given a term with some subterms labelled k it is easy to show that the descendants of these subterms across some reduction sequence are precisely the subterms labelled k in the final term. Moreover, these descendants are exactly the descendants obtained in the corresponding unlabelled reduction sequence. The result now follows by the first clause of the current proof when applied to the labelled version of the assumed iCRS.

(3) By the second clause of the current proof and orthogonality of the assumed iCRS.

(4) Identical to the proof of Proposition 12.5.9(iv) in [3], employing Lemma 5.9. ◻

Roughly, the above proof is identical to the proof of Proposition 12.5.9 in [3], except that Lemma 5.9 is employed instead of tracing.

With the Finite Jumps Developments Theorem in hand, we can now precisely characterise the sets of redexes having complete developments. This characterisation seems to be new.

Lemma 5.13. *The set $\mathcal{U}$ has a complete development if and only if $\mathcal{U}$ has the finite jumps property.*

Proof. To prove that the finite jumps property follows if $\mathcal{U}$ has a complete development, suppose $\mathcal{U}$ does not have the finite jumps property. In this case there is a path projection which ends in an infinite sequence of unlabelled nodes and ϵ-labelled edges.

By Lemma 5.9 we have for each step $s \to t$ contracting a redex in $\mathcal{U}$ that there is a surjection from $\mathcal{P}(s,\mathcal{U})$ to $\mathcal{P}(t,\mathcal{U}/u)$ which deletes only finite sequences of unlabelled nodes and ϵ-labelled edges. Hence, for all path projections we have that the nodes and edges left after the contraction of a redex in $\mathcal{U}$ either stay at the same distance from the first node of the path projection in which they occur or move closer to the first node. But then it follows immediately by ordinal induction that a path projection with an infinite sequence of unlabelled nodes and ϵ-labelled edges is present after each development. In particular, such an infinite sequence is present after the complete development. However, by definition of paths and path projections this means that a descendant of a redex in $\mathcal{U}$ is present in the final term of the complete development. But this contradicts the fact that no descendants of redexes in $\mathcal{U}$ exist in the final term of a complete development. Hence, $\mathcal{U}$ has the finite jumps property.

That $\mathcal{U}$ has a complete development if it has the finite jumps property is an immediate consequence of Theorem 5.12(4). ◻

The result we were aiming at now follows easily.

Theorem 5.14. *If $\mathcal{U}$ has a complete development then all complete developments of $\mathcal{U}$ end in the same term.*

Proof. By Lemma 5.13, if $\mathcal{U}$ has a complete development then it has the finite jumps property. But then each complete development of $\mathcal{U}$ ends in the same final term by Theorem 5.12(1). ◻

6 Further Directions

We have defined and proved the first results for iCRSs, but a number of questions that have been answered for iTRSs and iλc remain open: Does there exist a notion of meaningless terms [13] that allows for the construction of Böhm-like trees? Can we prove a partial confluence property [2, 3, 13] showing infinitary confluence up to equivalence of meaningless terms?

Furthermore, can the treatment of iCRS in this paper be extended to the other formats of higher-order rewriting? The fact that CRSs have a clean separation of abstractions (in terms and rewrite rules) and substitutions which is not present in some of the other forms of higher-order rewriting [3] may constitute a stumbling block in this respect.

Finally, it is as yet unclear how to relax the requirement that no infinite chains of meta-variables are allowed in meta-terms while still retaining a meaningful notion of substitution.

References

1. Dershowitz, N., Kaplan, S., Plaisted, D.A.: Rewrite, rewrite, rewrite, rewrite, rewrite, ... Theoretical Computer Science **83** (1991) 71–96
2. Kennaway, R., Klop, J.W., Sleep, R., de Vries, F.J.: Transfinite reductions in orthogonal term rewriting systems. Information and Computation **119** (1995) 18–38
3. Terese: Term Rewriting Systems. Cambridge University Press (2003)
4. Kennaway, J.R., Klop, J.W., Sleep, M., de Vries, F.J.: Infinitary lambda calculus. Theoretical Computer Science **175** (1997) 93–125
5. Berarducci, A.: Infinite λ-calculus and non-sensible models. In Ursini, A., Aglianò, P., eds.: Logic and Algebra. Volume 180 of Lecture Notes in Pure and Applied Mathematics., Marcel Dekker (1996) 339–378
6. Klop, J.W.: Combinatory Reduction Systems. PhD thesis, Rijksuniversiteit Utrecht (1980)
7. Klop, J.W., van Oostrom, V., van Raamsdonk, F.: Combinatory reduction systems: introduction and survey. Theoretical Computer Science **121** (1993) 279–308
8. Arnold, A., Nivat, M.: The metric space of infinite trees. Algebraic and topological properties. Fundamenta Informaticae **3** (1980) 445–476
9. Barendregt, H.P.: The Lambda Calculus: Its Syntax and Semantics. Second edn. Volume 103 of Studies in Logic and the Foundations of Mathematics. Elsevier Science, Amsterdam (1985)
10. Hanus, M., Prehofer, C.: Higher-order narrowing with definitional trees. In Ganzinger, H., ed.: Proceedinings of the 7th International Conference on Rewriting Techniques and Applications (RTA'96), New Brunswick, NJ, USA, July 27-30, 1996. Volume 1103 of Lecture Notes in Computer Science., Springer-Verlag (1996) 138–152
11. van Oostrom, V.: Higher-order families. In Ganzinger, H., ed.: Proceedings of the 7th International Conference on Rewriting Techniques and Applications (RTA '96), New Brunswick, NJ, USA, July 27-30, 1996. Volume 1103 of Lecture Notes in Computer Science., Springer-Verlag (1996) 392–407

12. Severi, P., de Vries, F.J.: An extensional Böhm model. In Tison, S., ed.: Proceedings of the 13th International Conference on Rewriting Techniques and Applications (RTA 2002), Copenhagen, Denmark, July 22-24, 2002. Volume 2378 of Lecture Notes in Computer Science., Springer-Verlag (2002) 159–173
13. Kennaway, R., van Oostrom, V., de Vries, F.J.: Meaningless terms in rewriting. The Journal of Functional and Logic Programming **1** (1999)

A Complete Proof of Theorem 4.9

Proof. Let $s \twoheadrightarrow^{\alpha} t$, and proceed by ordinal induction on α. By [3, Theorem 12.7.1] it suffices to show that the theorem holds for $\alpha = \omega + 1$: The cases where α is 0, a limit ordinal, or a successor ordinal greater than $\omega + 1$ do not depend on the definition of rewriting.

Suppose $\alpha = \omega + 1$ and write

$$s = s_0 \to s_1 \to \dots s_\omega \to s_{\omega+1} = t\,.$$

The redex contracted in $s_\omega \to s_{\omega+1}$, call it u, occurs at a position p and a depth d_u in s_ω. By definition of a rewrite rule, we have that the rewrite rule employed in u, say $l \to r$, has a finite left-hand side. Hence, there exists a depth $d_l \geq d_u$ in s_ω such that the pattern l matched in u, with exception of the meta-variables and the arguments of meta-variables, occurs above depth d_l. By Lemma 4.6 we can now write $s \twoheadrightarrow t$ as

$$s_0 \to^* s_n \twoheadrightarrow s_\omega \to s_{\omega+1}$$

with all rewrite steps in $s_n \twoheadrightarrow s_\omega$ occurring below depth d_l.

By left-linearity and fully-extendedness it follows that a redex v occurs at position p in s_n such that u is the unique residual of v. Contracting v in s_n yields a term t'.

The result now follows if we can construct a strongly convergent reduction sequence $t' \twoheadrightarrow^{\leq\omega} t$. To construct such a reduction sequence, assume $t_0 = t'$ and construct for each $d > 0$ a reduction sequence $t_{d-1} \to^* t_d$ in which all rewrite steps occur at depths greater or equal to $d - 1$ and such that $d(t_d, t) \leq 2^{-d}$. By construction reduction it follows that

$$t_0 \to^* t_1 \to^* \dots \to^* t_{d-1} \to^* t_d \to^* \dots t$$

is a strongly convergent reduction sequence of length at most ω, as required.

To see that we can construct for each $d > 0$ a reduction sequence $t_{d-1} \to^* t_d$ that satisfies the requirements, first note that there exists for some $m \in \mathbb{N}$ a context $C[\Box, \dots, \Box]$ with $m + 1$ holes which all occur a depth greater or equal to d_u, such that we can write $s_n \twoheadrightarrow s_\omega \to s_{\omega+1}$ as

$$C[\sigma(l), s'_1, \dots, s'_m] \twoheadrightarrow C[\sigma'(l), s''_1, \dots, s''_m] \to C[\sigma'(r), s''_1, \dots, s''_m]\,.$$

The context exists as each rewrite step in $s_n \twoheadrightarrow s_\omega$ occurs below depth $d_l > d_u$ and as we have for each initial segment $s_n \to^* s_{n'}$ of $s_n \twoheadrightarrow s_\omega$ that the unique residual of v in $s_{n'}$ occurs at depth d_u, which is also a consequence of the rewrite steps occurring below $d_l > d_u$.

By definition of $C[\Box, \ldots, \Box]$ we have that $t' = C[\sigma(r), s'_1, \ldots s'_m]$. Moreover, as the context stays unchanged in $C[\sigma(l), s'_1, \ldots, s'_m] \twoheadrightarrow C[\sigma'(l), s''_1, \ldots, s''_m]$, we have that this reduction sequence consists of a number of interleaved strongly converging reduction sequences of length at most ω:

- a reduction sequence for each meta-variable Z in l that reduces $\sigma(Z)(\boldsymbol{x})$ to $\sigma'(Z)(\boldsymbol{x})$ for some $\boldsymbol{x}$, and
- a reduction sequence for each $1 \leq i \leq m$ that reduces s'_i to s''_i.

By Lemma 4.6, we have for each of these sequences and each $d \geq 0$ that there exist finite initial segments which only omit rewrite steps that occur at depths greater or equal to d. We denote these sequences $\sigma(Z)(\boldsymbol{x}) \to^* \sigma^d(Z)(\boldsymbol{x})$ and $s'_i \to^* s^d_i$.

Now suppose $d > 0$. We construct in turn for the cases $d < d_u + 1$, $d = d_u + 1$, and $d > d_u + 1$ a reduction sequence $t_{d-1} \to^* t_d$ in which all rewrite steps occur at depths greater or equal to $d - 1$ and such that $d(t_d, t) \leq 2^{-d}$.

In case $d < d_u + 1$, i.e., $d \leq d_u$, define $t_d = t'$. For all $d \leq d_u$ we have that $t_{d-1} = t_d$. Hence, $t_{d-1} \to^* t_d$ and all rewrite steps are at depths greater or equal to $d - 1$. We have that $d(t_d, t) \leq 2^{-d}$ by the fact that the context $C[\Box, \ldots, \Box]$ occurs both in t' and t and the fact that the holes in the context occur below depth $d_u \geq d$.

In case $d = d_u + 1$, we employ $\sigma(Z)(\boldsymbol{x}) \to^* \sigma^1(Z)(\boldsymbol{x})$ and $s'_i \to^* s^1_i$ to construct $t_{d-1} \to^* t_d$, i.e., $t_{d_u} \to^* t_{d_u+1}$. To start, note that for each $\sigma(Z)(\boldsymbol{x})$ in $\sigma(l)$ we have that a number of subterms of the form $\sigma(Z)(\boldsymbol{t})$, for different $\boldsymbol{t}$, occur in $\sigma(r)$. By left-linearity, $\sigma(Z)(\boldsymbol{t}) \to^* \sigma^1(Z)(\boldsymbol{t})$. As only finite chains of meta-variables occur in r, there exists some depth d_r in $\sigma(r)$ such that for each position q of $\sigma(r)$ which is at depth greater than d_r we have that the number of positions $p \leq q$ at which no subterm of the form $\sigma(Z)(\boldsymbol{t})$ occurs is greater or equal to 1. The number of subterms $\sigma(Z)(\boldsymbol{t})$ above depth d_r in $\sigma(r)$ is finite, as each function symbol has finite arity. Reducing the finite number of subterms above depth d_r according to $\sigma(Z)(\boldsymbol{t}) \to^* \sigma^1(Z)(\boldsymbol{t})$ creates in a finite number of steps a term $\sigma_d(r)$. For $\sigma(r) \to^* \sigma_d(r)$ we have that all rewrite steps occur at depths greater or equal to 0 and that $d(\sigma_d(r), \sigma'(r)) \leq 2^{-1}$. This second fact follows by the assumption that all rewrite steps in $\sigma^1(Z)(\boldsymbol{t}) \to^* \sigma'(Z)(\boldsymbol{t})$ occur at depths greater or equal to 1, the assumptions regarding d_r, which are needed in case some subterm $\sigma(Z)(\boldsymbol{t})$ reduces to a term in $\boldsymbol{t}$, and the usage of only initial segments of $\sigma(Z)(\boldsymbol{x}) \to^* \sigma'(Z)(\boldsymbol{x})$, which also occur in $s_n \to^* s_\omega$. Interleaving the reduction sequence $\sigma(r) \to^* \sigma_d(r)$ with the reduction sequences $s'_i \to^* s^1_i$ gives a reduction sequence

$$t_{d_u} = C[\sigma(r), s'_1, \ldots s'_m] \to^* C[\sigma_d(r), s^1_1, \ldots s^1_m] \,.$$

Define $t_d = C[\sigma_d(r), s^1_1, \ldots s^1_m]$. Obviously, $t_{d_u} \to^* t_d$. The requirement $d(t_d, t) \leq 2^{-d}$ follows from $d(\sigma_d(r), \sigma'(r)) \leq 2^{-1}$, the fact that all rewrite steps in $s^1_i \to^* s''_i$ occur at depths greater or equal to 1, and the usage of only initial segments of $s'_i \to^* s''_i$, which also occur in $s_n \to^* s_\omega$.

In case $d > d_u + 1$ assume t_0 reduces to t_{d-1} by applying to some subterms of t_0 the reduction sequences $\sigma(Z)(\boldsymbol{x}) \to^* \sigma^{d-1-d_u}(Z)(\boldsymbol{x})$ and $s'_i \to^* s_i^{d-1-d_u}$. We use $\sigma(Z)(\boldsymbol{x}) \to^* \sigma^{d-d_u}(Z)(\boldsymbol{x})$ and $s'_i \to^* s_i^{d-d_u}$ to define $t_{d-1} \to^* t_d$. As in the previous case we have for each $\sigma(Z)(\boldsymbol{x})$ in $\sigma(l)$ that a number of subterms of the form $\sigma(Z)(\boldsymbol{t})$ occur in $\sigma(r)$ and that by left-linearity $\sigma(Z)(\boldsymbol{t}) \to^* \sigma^{d-d_u}(Z)(\boldsymbol{t})$. Moreover, as only finite chains of meta-variables are allowed in r, there exists some depth d_r in $\sigma(r)$ such that for each position q of $\sigma(r)$ which is at depth greater than d_r we have that the number of positions $p \leq q$ at which no subterm of the form $\sigma(Z)(\boldsymbol{t})$ occurs is greater than $d - d_u$. The number of subterms $\sigma(Z)(\boldsymbol{t})$ above depth d_r in $\sigma(r)$ is finite, as each function symbol has finite arity. Reducing the finite number of residuals above depth d_r according to $\sigma(Z)(\boldsymbol{x}) \to^* \sigma^{d-d_u}(Z)(\boldsymbol{x})$ creates in a finite number of steps a term $\sigma_d(r)$. For this term we have $d(\sigma_d(r), \sigma'(r)) < 2^{d-d_u}$ by the assumption that all rewrite steps in $\sigma^{d-d_u}(Z)(\boldsymbol{t}) \to^* \sigma'(Z)(\boldsymbol{t})$ occur at depths greater or equal to $d - d_u$, the assumptions regarding d_r, which are needed in case some subterm $\sigma(Z)(\boldsymbol{t})$ reduces to a term in $\boldsymbol{t}$, and the usage of only initial segments of $\sigma(Z)(\boldsymbol{x}) \to^* \sigma'(Z)(\boldsymbol{x})$, which also occur in $s_n \to^* s_\omega$. Interleaving the reduction sequence $\sigma(r) \to^* \sigma_d(r)$ with the reduction sequences $s'_i \to^* s_i^{d-d_u}$ which gives a reduction sequence

$$t_0 = C[\sigma(r), s'_1, \ldots s'_m] \to^* C[\sigma_d(r), s_1^{d-d_u}, \ldots s_m^{d-d_u}] \,.$$

Define $t_d = C[\sigma_d(r), s_1^{d-d_u}, \ldots s_m^{d-d_u}]$. The requirement $d(t_d, t) < 2^{-d}$ follows from $d(\sigma_d(r), \sigma'(r)) < 2^{d-d_u}$, the fact that all rewrite steps in $s_i^{d-d_u} \to^* s''_i$ occur at depths greater or equal to $d - d_u$, and the usage of only initial segments of $s'_i \to^* s''_i$, which also occur in $s_n \to^* s_\omega$.

The rewrite sequences used to construct t_{d-1} are initial segments of the reduction sequences used to construct t_d. Hence, we can re-order the reduction sequence $t_0 \to^* t_d$ in such a way that we get a reduction sequence $t_0 \to^* t_{d-1} \to^* t_d$. Now, all rewrite steps in $t_{d-1} \to^* t_d$ occur at depths greater or equal to $d - 1$ by definition of t_{d-1} and t_d and of $\sigma(Z)(\boldsymbol{x}) \to^* \sigma^{d-d_u}(Z)(\boldsymbol{x})$ and $s'_i \to^* s_i^{d-d_u}$.

Hence, give the reduction sequence constructed above we can also construct a strongly converging reduction sequence $t' \twoheadrightarrow^{\leq \omega} t$. Because $s \to^* t'$, we have that $s \twoheadrightarrow^{\leq \omega} t$, as required. □

B Proof of Lemma 5.9

To prove Lemma 5.9 we need a property of path projections. We prove this property first. As in Section 5, we assume in this appendix we are working in an orthogonal iCRS and that $\mathcal{U}$ is a set of redexes in a term s.

Proposition B.1. *Let Ψ denote the set of maximal paths of s with respect to $\mathcal{U}$. The map ϕ defined as:*

$$\phi(\Psi) = \{\phi(\Pi) \mid \Pi \in \Psi\}$$

is a bijection between Ψ and $\mathcal{P}(s, \mathcal{U})$.

Proof. By definition $\phi(\Psi) = \mathcal{P}(s,\mathcal{U})$. Hence, ϕ is *surjective*. To prove that ϕ is *injective*, suppose there exists $\Pi, \Pi' \in \Psi$ such that $\phi(\Pi) = \phi(\Pi')$. Let Π^* be the longest shared between Π and Π'. The prefix Π^* is non-empty, as any path of s begins with (s, ϵ). There are now two cases to consider depending on Π^* ending in an edge or a node.

In case Π^* ends in an edge, the next node is uniquely determined by the definition of paths. Hence, as Π and Π' are maximal we can extend Π^* with that unique node, contradiction.

In case Π^* ends in a node, at least one of Π and Π' extends Π^*, otherwise $\Pi = \Pi'$. In case the extension is with an unlabeled edge, the other path must also extend Π^* with an unlabelled edge. This follows by the definition of paths and by Π and Π' being maximal. Otherwise, in case the extension is with an edge labelled i, the other path must also extend Π^* with an edge labelled i. This follows by definition of paths and by $\phi(\Pi) = \phi(\Pi')$. Hence, in case Π^* ends in a node a contradiction also follows and we can conclude that ϕ is injective. □

To prove Lemma 5.9 we define a map θ_u that takes maximal paths Π of s with respect to $\mathcal{U}$ to maximal paths of t with respect to $\mathcal{U}/u$, where $u \in \mathcal{U}$ and $s \to t$ by contraction of u. The definition of $\theta_u(\Pi)$ involves a partial map ψ that has three arguments: a node of Π, a finite string over $\mathbb{N}$, and a partial map from $\mathcal{U} - \{u\}$ to finite strings over $\mathbb{N}$.

We first define ψ. In the definition, given a partial map ρ, we denote by $\rho[x \mapsto y]$ the partial map ρ' defined as:

$$\rho'(z) = \begin{cases} y & \text{if } z = x \\ \rho(z) & \text{otherwise.} \end{cases}$$

Definition B.2. *Let Π be a maximal path of s with respect to $\mathcal{U}$, let $u \in \mathcal{U}$, and let $s \to t$ by contraction of u. Define $\psi(n, q_t, \rho)$ as:*

1. *If n is labelled (s, p) with the subterm at p not a redex in $\mathcal{U}$, then*
   (a) if n has no outgoing edge define $\psi(n, q_t, \rho) = (t, q_t)$,
   (b) if n has an edge labelled i to n' define $\psi(n, q_t, \rho) = (t, q_t) \xrightarrow{i} \psi(n', q_t \cdot i, \rho)$.
2. *If n is labelled (s, p_v) with $v \in \mathcal{U} - \{u\}$ and if n has an unlabelled edge to n', then define $\psi(n, q_t, \rho) = (t, q_t) \to \psi(n', q_t, \rho[v \mapsto q_t])$,*
3. *If n is labelled (s, p_u) and if n has an unlabelled edge to n', then define $\psi(n, q_t, \rho) = \psi(n', q_t, \rho)$.*
4. *If n is labelled (s, p) with $s|_p$ a variable bound by $v \in \mathcal{U} - \{u\}$ and if n has an unlabelled edge to n', then define $\psi(n, q_t, \rho) = (t, q_t) \to \psi(n', \rho(v), \rho)$.*
5. *If n is labelled (s, p) with $s|_p$ a variable bound by u and if n has an unlabelled edge to n', then define $\psi(n, q_t, \rho) = \psi(n', q_t, \rho)$.*
6. *If n is labelled (r, p, p_v) with $r|_p$ not a meta-variable and $v \in \mathcal{U} - \{u\}$, then*
   (a) if n has no outgoing edge define $\psi(n, q_t, \rho) = (r, p, q_t)$,
   (b) if n has an edge labelled i to n' define $\psi(n, q_t, \rho) = (r, p, q_t) \xrightarrow{i} \psi(n', q_t, \rho)$.
7. *If n is labelled (r, p, p_u) with $r|_p$ not a meta-variable, then*
   (a) if n has no outgoing edge define $\psi(n, q_t, \rho) = (t, q_t)$,
   (b) if n has an edge labelled i to n' define $\psi(n, q_t, \rho) = (t, q_t) \xrightarrow{i} \psi(n', q_t \cdot i, \rho)$.

8. *If n is labelled (r, p, p_v) with $r|_p$ a meta-variable Z and $v \in \mathcal{U} - \{u\}$ and if n has an unlabelled edge to n', which is labelled $(s, p_v \cdot q)$, then define $\psi(n, q_t, \rho) = (r, p, q_t) \to \psi(n', q_t \cdot q, \rho)$.*
9. *If n is labelled (r, p, p_u) with $r|_p$ a meta-variable and if n has an unlabelled edge to n', which is labelled $(s, p_u \cdot q)$, then define $\psi(n, q_t, \rho) = \psi(n', q_t, \rho)$.*

Let $\perp$ be completely undefined map. We define the following.

Definition B.3. *Let $u \in \mathcal{U}$ and let Π be a maximal path of s with respect to $\mathcal{U}$. The map θ_u is defined as:*

$$\theta_u(\Pi) = \psi((s, \epsilon), \epsilon, \perp) .$$

Note that $\theta_u(\Pi)$ is calculated by iteration of ψ. After a finite number of iterations, a finite prefix of $\theta_u(\Pi)$ is obtained.

We next show that that θ_u is well-defined: $\theta_u(\Pi)$ is a maximal path of t with respect to $\mathcal{U}/u$.

Proposition B.4. *Let Π be a maximal path of s with respect to $\mathcal{U}$ and let $u \in \mathcal{U}$. For each finite number of iterations of ψ in the calculation of $\theta_u(\Pi)$ the following holds:*

- *Either no nodes and edges have been generated, or the generated nodes and edges, with exception of the edge generated last, form a path, and the edge generated last is a valid one when extending the path to a longer one.*
- *For all defined $\rho(v)$ in the third argument of ψ it holds that $\rho(v) \in \mathcal{P}os(t)$ and that a residual of v occurs at $\rho(v)$ in t.*

Distinguishing on the particular clause of Definition B.2 employed in the last iteration, the following also holds:

(1) *(a) nothing; (b) $q_t \cdot i$ is descendant of $p \cdot i$ and the next node generated is $(t, q_t \cdot i)$, which together with the previously generated nodes and edges forms a path;*
(2) *a residual of v occurs at q_t in t and the next node generated is (r, ϵ, q_t), which together with the previously generated nodes and edges forms a path;*
(3) *$q_t = p_u$ and the next node generated is (t, q_t), which together with the previously generated nodes and edges forms a path;*
(4) *if n' is labelled $(r, p' \cdot i, p_v)$, then the next node generated is $(r, p' \cdot i, \rho(v))$, which together with the previously generated nodes and edges forms a path;*
(5) *there are two subcases:*
    - *if the previous iterations employ clause (3) followed by a number of iterations employing in turn clauses (5) and (9), then $q_t = p_u$;*
    - *if the previous iteration employs clause (1) or if the previous iterations employ clause (7) followed by a number of iterations employing in turn clauses (5) and (9), then $q_t = q'_t \cdot i$ where q'_t is the q_t from clause either clause (1) or (7);*

    in both cases the next node generated is (t, q_t), which together with the previously generated nodes and edges forms a path.

(6) (a) nothing; (b) a residual of v occurs at q_t and the next node generated is $(r, p \cdot i, q_t)$, which together with the previously generated nodes and edges forms a path;
(7) (a) nothing; (b) $q_t \cdot i = p_u \cdot q \cdot i$ where q is a descendant of p across complete development of the parallel β-redexes in r_σ, with r_σ just as before Definition 5.2, and the next node generated is $(t, q_t \cdot i)$, which together with the previously generated nodes and edges forms a path;
(8) $q_t \cdot q$ is a descendant of $p_v \cdot q$ and the next node generated is $(t, q_t \cdot q)$, which together with the previously generated nodes and edges forms a path;
(9) there are two subcases:
- *if the next iteration does* not *employ clause (5), then q_t is a descendant of $p_u \cdot q$;*
- *otherwise, q_t is as in clause (5);*

in both cases the next node generated is (t, q_t), which together with the previously generated nodes and edges forms a path.

In addition, if Π is a maximal path of s with respect to $\mathcal{U}$, then $\theta_u(t)$ is a maximal path of t with respect to $\mathcal{U}/u$.

Proof. Let Π be a maximal path of s with respect to $\mathcal{U}$, let $u \in \mathcal{U}$, and let $s \to t$ by contraction of u. We prove the lemma by induction on the number of iterations of ψ.

Below, when we say that something is a path of t, we implicitly assume that it is a path of t with respect to $\mathcal{U}/u$.

Base case. By definition of ψ we have that only clauses (1), (2), and (3) can be employed in the first iteration. The others clauses either require a bound variable at the root of s, which is impossible, or they require the label of the node in the first argument to be a triple, which is not the case. We deal with each of the possible clauses in turn:

(1) In this case a node labelled (t, ϵ) is generated and possibly an edge labelled i. Obviously, (t, ϵ) is a path. The partial map $\perp$ is unaffected by this clause, thus satisfying the necessary requirements.
As the current iteration implies that no redex from $\mathcal{U}$ occurs at the root of s, we have that $root(t|_\epsilon) = root(s|_\epsilon)$. Hence, in case of the clause (a), the path is maximal just like Π. In case clause (b), the edge labelled i is allowed. Moreover, $i \in \mathcal{P}os(t)$ is a descendant of $i \in \mathcal{P}os(s)$, and the next node generated node is (t, i). This node forms a path together with (t, ϵ) and the edge labelled i. By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(2) In this case a node labelled (t, ϵ) and an unlabelled edge are generated. Obviously, (t, ϵ) is a path. As an orthogonal iCRS is assumed, a redex v' occurs at $\epsilon \in \mathcal{P}os(t)$ which is a residual of v. Hence, $\perp[v \mapsto \epsilon]$ satisfies the necessary requirements.
As $v \in \mathcal{U} - \{u\}$, it holds that $v' \in \mathcal{U}/u$. Hence, the unlabelled edge is allowed, and the next generated node must be (r, ϵ, ϵ), where r is the right-hand side of the rewrite rule employed in v'. This node forms a path together with

(t, ϵ) and the unlabelled edge. By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(3) Obviously, $\perp$ is unaffected by this clause, thus satisfying the necessary requirements. Moreover, $\epsilon \in \mathcal{P}os(t)$ is equal to $\epsilon \in \mathcal{P}os(s)$, and the next generated node must be (t, ϵ), which is a path. By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

Induction step. Assume we have proved the lemma up to some arbitrary number of iterations. We next prove that it also holds in case of one more iteration. We deal with each of the possible clauses in turn:

(1) In this case the only possible clauses employed in the previous iteration are (1), (8), and (9). All other clauses force the label of the node in the first argument to be a triple, which is not the case.
A node labelled (t, q_t) is generated and possibly an edge labelled i. By the clauses possible in the previous iteration, (t, q_t) forms a path together with the previously generated nodes and edges. The partial map ρ is unaffected by this clause, thus satisfying the necessary requirements.
Also by the clauses possible in the previous iteration, q_t is a descendant of p and $root(t|_{q_t}) = root(s|_p)$. Hence, in case of clause (a), the path is maximal just like Π. In case of clause (b), the edge labelled i is allowed. Moreover, $q_t \cdot i \in \mathcal{P}os(t)$ is a descendant of $p \cdot i \in \mathcal{P}os(s)$, and the next node generated is $(t, q_t \cdot i)$. This node forms a path together with previously generated nodes and edges. By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(2) As before, the only possible clauses employed in the previous iteration are (1), (8), and (9). All other clauses force the label of the node in the first argument to be a triple, which is not the case.
A node labelled (t, q_t) and an unlabelled edge are generated. By the clauses possible in the previous iteration, (t, q_t) forms a path together with the previously generated nodes and edges. Moreover, as an orthogonal iCRS is assumed, a redex v' occurs at q_t which is a residual of v. Hence, as ρ satisfies the necessary requirements, $\rho[v \mapsto q_t]$ does so too.
As $v \in \mathcal{U} - \{u\}$, it holds that $v' \in \mathcal{U}/u$. Hence, the unlabelled edge is allowed, and the next node generated is (r, ϵ, q_t), where r is the right-hand side of the rewrite rule employed in v'. This node forms a path together with previously generated nodes and edges. By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(3) In this case the only possible clauses employed in the previous iteration are (1) and (8). All other clauses, except (9), force the label of the node in the first argument to be a triple, which is not the case. Clause (9) is impossible as it requires the redex v to occur above itself in s.
Obviously, ρ is unaffected by this clause, thus satisfying the necessary requirements. Moreover, $q_t = p$, by definition of descendants. By the clauses possible in the previous iteration, the next node generated is (t, q_t) and the node forms a path together with previously generated nodes and edges. By

the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(4) As before, the only possible clauses employed in the previous iteration are (1), (8), and (9). All other clauses force the label of the node in the first argument to be a triple, which is not the case.

A node labelled (t, q_t) and an unlabelled edge are generated. By the clauses possible in the previous iteration, (t, q_t) forms a path together with the previously generated nodes and edges. The partial map ρ is unaffected by this clause, thus satisfying the necessary requirements.

Also by the clauses possible in the previous iteration, q_t is a descendant of p and $root(t|_{q_t})$ is a variable bound by a residual v' of v in t, where by construction $\rho(v)$ is the position of v'. Hence, the unlabelled edge is allowed. That a node labelled $(r, p', \rho(v))$, where r is the right-hand side of the rewrite rule employed in v', has been generated as the last node labelled with $\rho(v)$ follows by definition of ψ. Hence, the next generated node is $(r, p \cdot i, \rho(v))$. This node forms a path together with previously generated nodes and edges. By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(5) As before, the only possible clauses employed in the previous iteration are (1), (8), and (9). All other clauses force the label of the node in the first argument to be a triple, which is not the case.

Obviously, ρ is unaffected by this clause, thus satisfying the necessary requirements. By the clauses in the previous iteration, the requirements of the two subcases are satisfied and the next node generated is (t, q_t). By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(6) In this case the only possible clauses employed in the previous iteration are (2), (4), and (6). Clauses (1), (8), and (9) force the label of the node in the first argument to be a tuple, which is not the case. Clauses (3) and (5) force the v to be equal to u, which is not allowed.

A node labelled (r, p, q_t) is generated and possibly an edge labelled i. By the clauses possible in the previous iteration, (r, p, q_t) forms a path together with the previously generated nodes and edges. The partial map ρ is unaffected by this clause, thus satisfying the necessary requirements.

As r is left unchanged, it holds in case of clause (a), that the path is maximal just like Π. In case of clause (b), the edge labelled i is allowed and that next node generated is $(r, p{\cdot}i, q_t)$. This node forms a path together with previously generated nodes and edges. By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(7) In this case the only possible clauses employed in the previous iteration are (3), (5), and (7). Clauses (1), (8), and (9) force the label of the node in the first argument to be a tuple, which is not the case. Clauses (2) and (4) force a v not equal to u.

A node labelled (t, q_t) is generated and possibly an edge labelled i. By the clauses possible in the previous iteration (t, q_t) forms a path together with the previously generated nodes and edges. The partial map ρ is unaffected by this clause, thus satisfying the necessary requirements.

Also by the clauses possible in the previous iteration, $q_t = p_u \cdot q$ and $root(t|_{q_t}) = root(r_\sigma|_p)$, where q is a descendant of p across complete development of the parallel β-redexes in r_σ, with r_σ just as before Definition 5.2. Hence, in case of clause (a), the path is maximal just like Π. In case of clause (b), the edge labelled i is allowed. Moreover, $q_t \cdot i = p_u \cdot q \cdot i$, and the next node generated is $(t, q_t \cdot i)$. This node forms a path together with previously generated nodes and edges. By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(8) As before, the only possible clauses employed in the previous iteration are (2), (4), and (6). Clauses (1), (8), and (9) force the label of the node in the first argument to be a tuple, which is not the case. Clauses (3) and (5) force the v to be equal to u, which is not allowed.

In this case a node labelled (r, p, q_t) and an unlabelled edge are generated. By the clauses possible in the previous iteration, (r, p, q_t) forms a ptah together with the previously generated nodes and edges. The partial map ρ is unaffected by this clause, thus satisfying the necessary requirements.

As r is left unchanged, the unlabelled edge is allowed. Moreover, as a residual of v occurs at q_t, the next node generated is $(t, q_t \cdot q)$. This node forms a path together with previously generated nodes and edges. By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

(9) As before, the only possible clauses employed in the previous iteration are (3), (5), and (7). Clauses (1), (8), and (9) force the label of the node in the first argument to be a tuple, which is not the case. Clauses (2) and (4) force a v not equal to u.

Obviously, ρ is unaffected by this clause, thus satisfying the necessary requirements. By the clauses in the previous iteration, the requirement that q_t is a descendant of $p_u \cdot q$ is satisfied in the first subcase, and the requirements of clause (5) are satisfied in the second subcase. Moreover, the next node generated is (t, q_t). By the construction of paths the finite prefix of Π considered so far is not maximal, and there is nothing more to show.

There can be no infinite cycle of iterations employing clauses (5) and (9) in the construction of $\theta_u(\Pi)$, since this implies existence of an infinite chain of meta-variables in r. Hence, $\theta_u(\Pi)$ is a well-defined path of t with respect to $\mathcal{U}/u$. In case Π is finite, the induction shows that $\theta_u(\Pi)$ is maximal. In case Π is infinite, then so is $\theta_u(\Pi)$, and we conclude that $\theta_u(\Pi)$ is a maximal path. □

The next lemma relates the maximal paths of s with respect to $\mathcal{U}$ to the maximal paths of t with respect to $\mathcal{U}/u$. In the proof of the lemma we leave out the labels of the explicitly denoted edges.

Proposition B.5. *The map θ_u is a bijection.*

Proof. Let $u \in \mathcal{U}$ and $s \to t$ by contraction of u. By Proposition B.4, we have that θ_u maps maximal paths of s with respect to $\mathcal{U}$ to maximal paths of t with respect to $\mathcal{U}/u$.

To prove that θ_u is *surjective*, let Π_t be a maximal path of t with respect to $\mathcal{U}/u$. We are done if $\Pi_t = \theta_u(\Pi_s)$ for some maximal path Π_s of s with respect to $\mathcal{U}$. Otherwise, Π_t has a finite non-empty prefix Π'_t in common with $\theta_u(\Pi_s)$ for some maximal Π_s of s with respect to $\mathcal{U}$. The prefix is non-empty since any path of t begins with (t, ϵ). Let Π'_t be the longest finite prefix of Π_t such that $\theta_u(\Pi_s) = \Pi'_t \to \ldots$ for some maximal path Π_s. We have $\Pi_s = \Pi'_s \to \ldots$ for some finite path Π'_s. By Proposition B.4 we can extend the prefix Π'_t with a new node precisely when we can extend Π'_s. Hence, Π'_t cannot be the *longest* finite prefix with $\theta_u(\Pi_s) = \Pi'_t \to \ldots$ for some maximal path Π_s in s, since we can extend Π'_s to form a new maximal path with one more node, contradiction. Hence, $\Pi_t = \theta_u(\Pi_s)$ for some maximal path Π_s.

To prove that θ_u is *injective*, suppose there exists two maximal paths Π and Π' of s with respect to $\mathcal{U}$ such that $\theta_u(\Pi) = \theta_u(\Pi')$. Let Π^* be the longest prefix shared between Π and Π'. The prefix Π^* is non-empty, as any path of s begins with (s, ϵ). There are now two cases to consider depending on Π^* ending in an edge or a node.

In case Π^* ends in an edge, the next node is uniquely determined by the definition of paths. Hence, as Π and Π' are maximal we can extend Π^* with that unique node, contradiction.

In case Π^* ends in a node, at least one of Π and Π' extends Π^*, otherwise $\Pi = \Pi'$. In case the extension is with an unlabeled edge, the other path must also extend Π^* with an unlabelled edge. This follows by the definition of paths and by Π and Π' being maximal. Otherwise, in case the extension is with an edge labelled i, the other path must also extend Π^* with an edge labelled i. This follows by definition of paths and by $\theta_u(\Pi) = \theta_u(\Pi')$. Hence, in case Π^* ends in a node a contradiction also follows and we can conclude that θ_u is injective. □

We finally prove Lemma 5.9.

Proof (Lemma 5.9). By Proposition B.5, the map θ_u is a bijection between the maximal paths of s with respect to $\mathcal{U}$ and the maximal paths of t with respect to $\mathcal{U}/u$. By Proposition B.1, a bijection exists between the set of paths and the set of path projections mapping unlabelled edges to ϵ-labelled edges and labelled edges to edges with the same label. Hence, θ_u induces a bijection θ'_u between $\mathcal{P}(s,\mathcal{U})$ and $\mathcal{P}(t,\mathcal{U}/u)$. By examining the construction of θ_u, we see that it only deletes unlabelled edges and nodes corresponding to meta-variables of u and variables bound by u. Moreover, it is evident that if an infinite sequence of nodes and unlabelled edges were deleted, the right-hand side of the rule of u would contain an infinite chain of meta-variables, contradicting the definition of meta-terms. Hence, $\phi(\theta'_u(\Pi))$ can be obtained by $\phi(\Pi)$ by deleting only finite sequences of unlabelled nodes and ϵ-labelled edges, as required. □