

Infinitary Combinatory Reduction Systems

(Extended Abstract)

Jeroen Ketema¹ and Jakob Grue Simonsen²

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 (i λ c) [3, 4]. 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 i λ c 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, 5, 6].

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 i λ c:

CRSs	iTRSs/i λ c
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 i λ c 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 Sect. 2 we give some preliminary definitions, and in Sect. 3, we define infinite (meta-)terms and substitutions. Thereafter, in Sect. 4 we define infinitary rewriting and prove compression, and in Sect. 5 we investigate complete developments. Finally, in Sect. 6 we give directions for further research.

2 Preliminaries

Prior knowledge of CRSs [6] 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 $\alpha, \beta, \gamma, \dots$

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, \dots, s_n are finite meta-terms, then $Z(s_1, \dots, s_n)$ is a finite meta-term,
4. if $f \in \Sigma$ has arity n and s_1, \dots, s_n are finite meta-terms, then $f(s_1, \dots, s_n)$ 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 $Pos(s)$, is the set of finite strings over \mathbb{N} , with ϵ the empty string, such that:

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

Given $p, q \in \mathcal{Pos}(s)$, we say that p is a *prefix* of q , denoted $p \leq q$, if there exists an $r \in \mathcal{Pos}(s)$ 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 λc , 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 [2,4,7]. 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 $\text{root}(s|_p) \neq \text{root}(t|_p)$. In λc , 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 $s[x \rightarrow y]$ 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 $p \in \mathcal{Pos}(s) \cap \mathcal{Pos}(t)$ such that:*

1. if $p = \epsilon$, then $\text{root}(s) \neq \text{root}(t)$,
2. if $p = i \cdot q$ for $i \geq 1$, then $\text{root}(s) = \text{root}(t)$ 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 $s'[x_1 \rightarrow y]$ and $t'[x_2 \rightarrow y]$, 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 $p \in \mathcal{Pos}(s)$. Define:*

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

Note that meta-variables are not counted by D . Changing the second clause to $D(Z(t_1, \dots, t_n), i \cdot p') = 1 + D(t_i, p')$ 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 $\lambda\mathcal{c}$.

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

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

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 $\lambda\mathcal{c}$, 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 $\text{root}(s|_q)$ 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:

$$\begin{aligned} & Z(Z(\dots(Z(\dots)))) \\ & Z_1(Z_2(\dots(Z_n(\dots)))) \end{aligned}$$

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:

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

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 $\lambda\mathcal{c}$ 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 $\lambda x_1.\lambda x_2 \dots \lambda x_n \dots$ are disallowed.

The terms can now be defined as in the finite case [3,5,6]. 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, 6], except that coinduction is employed instead of induction. This is identical to what is done in the case of iTRSs and $\text{i}\lambda\text{c}$ with respect to the finite systems they are based on. In the definitions we use \mathbf{x} and \mathbf{t} as a short-hands for respectively the sequences x_1, \dots, x_n and t_1, \dots, t_n with $n \geq 0$. We assume n fixed in the next two definitions.

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

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

The above definition implicitly takes into account the variable convention [8] 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 $\underline{\lambda}x_1, \dots, x_n.s$ or $\underline{\lambda}\mathbf{x}.s$, with s a term, such that:

$$(\underline{\lambda}\mathbf{x}.s)(t_1, \dots, t_n) = s[\mathbf{x} := \mathbf{t}]. \quad (1)$$

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

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

This rule can be seen a *parallel β -rule*. That is, a variant of the β -rule from $\text{i}\lambda\text{c}$ which substitutes for multiple variables simultaneously. The root of $(\underline{\lambda}\mathbf{x}.s)$ 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. $\bar{\sigma}(Z(s_1, \dots, s_m)) = \sigma(Z)(\bar{\sigma}(s_1), \dots, \bar{\sigma}(s_m))$,
4. $\bar{\sigma}(f(s_1, \dots, s_m)) = f(\bar{\sigma}(s_1), \dots, \bar{\sigma}(s_m))$.

Similar to Definition 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 $Z(t_1, \dots, t_n)$ is replaced by a subterm of the form $(\underline{\lambda}\mathbf{x}.s)(t_1, \dots, t_n)$. In the second step Eq. (1) is applied to each subterm of the form $(\underline{\lambda}\mathbf{x}.s)(t_1, \dots, t_n)$ 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 λc . 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 [5, 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 λc definitions.

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

$$(\underline{\lambda}\mathbf{x}.x_i)(t_1, \dots, t_n),$$

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 $(\underline{\lambda}x_1, \dots, x_n.s)(t_1, \dots, t_n)$ as a sequence of β -redexes:

$$(\lambda x_1(\dots((\lambda x_n.s)t_n)\dots))t_1.$$

This means that each complete development in our variant of λc corresponds to a complete development in λc extended with some function symbols. As each complete development in λc 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. \square

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

$$Z(Z(\dots(Z(\dots))))).$$

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

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

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 $\underline{\lambda}x.y$ to each meta-variable Z in the ‘meta-term’ with y in 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’

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

Applying the valuation that assigns to Z the substitute $\underline{\lambda}xy.y$ yields the ‘ $\underline{\lambda}$ -term’ of Fig. 1. It reduces by means of two different developments to the $\underline{\lambda}$ -terms of Fig. 2 and Fig. 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 $i\lambda c$ [4, Sect. 4].




Fig. 1.




Fig. 2.




Fig. 3.

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.

4 Infinitary Rewriting

We continue to combine the definitions of iTRSs and $i\lambda c$ 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 $l \rightarrow r$, where l is a finite meta-term and r is a meta-term, such that:*

1. l is a pattern and of the form $f(s_1, \dots, s_n)$ 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 $i\lambda c$ 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 $i\lambda c$ are iCRSs.

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

Definition 4.3. *Let $l \rightarrow r$ be a rewrite rule. Given a valuation $\bar{\sigma}$, the term $\bar{\sigma}(l)$ is called a $l \rightarrow r$ -redex. If $s = C[\bar{\sigma}(l)]$ for some context $C[\square]$ with $\bar{\sigma}(l)$ a $l \rightarrow r$ -redex and p the position of the hole in $C[\square]$, then an $l \rightarrow r$ -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 $s \rightarrow t$, such that a $l \rightarrow r$ -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 $i\lambda c$ 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 $s_\beta \rightarrow s_{\beta+1}$ for all $\beta < \alpha$. For each rewrite step $s_\beta \rightarrow s_{\beta+1}$, 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 $s \twoheadrightarrow^* t$ 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 $l \rightarrow r$ 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 [9, 10], 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 $s \twoheadrightarrow^* s' \twoheadrightarrow^\omega s'' \rightarrow t$, such that all rewrite steps in $s' \twoheadrightarrow^\omega s''$ occur below the meta-variable positions of the redex contracted in the step of $s'' \rightarrow t$. By fully-extendedness and left-linearity it follows that a redex of which the redex contracted in $s'' \rightarrow t$ 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 $t_{d-1} \twoheadrightarrow^* t_d$ 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 λc [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 $t_0 \twoheadrightarrow^* t_1 \twoheadrightarrow^* \dots \twoheadrightarrow^* t_{d-1} \twoheadrightarrow^* t_d \twoheadrightarrow^* \dots t$ is a strongly convergent reduction sequence of length at most ω . As $s \twoheadrightarrow^* t'$, we then have that $s \twoheadrightarrow^{\leq \omega} t$, as required. □

The previous theorem does not hold in general for iCRSs that are not left-linear or fully-extended. For left-linearity, this follows from the iTRS counterexample in [2]. For 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]. 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 $R = \{l_i \rightarrow r_i \mid i \in I\}$ be a set of rewrite rules.*

1. *R is non-overlapping if it holds that:*
 - *each $l_i \rightarrow r_i$ -redex that occurs at a position p in an $l_j \rightarrow r_j$ -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 of 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 $\bar{\sigma}(l) \rightarrow \bar{\sigma}(r)$ 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 Eq. (1). The second step defines descendants across application of Eq. (1).

Given that the second step of the substitution is just a complete development in a variant of $i\lambda c$, the second step in the definition of descendants is just a variant of descendants in $i\lambda c$ [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 $1, \dots, n$ we denote the positions of the subterms on the right-hand side of the $\underline{\lambda}$ -application. This means that $(\underline{\lambda}\mathbf{x}.s)(t_1, \dots, t_n)|_0 = (\underline{\lambda}\mathbf{x}.s)$, $\underline{\lambda}\mathbf{x}.s|_0 = s$, and $Z(t_1, \dots, t_n)|_i = (\underline{\lambda}\mathbf{x}.s)(t_1, \dots, t_n)|_i = t_i$ for $1 \leq i \leq n$. We denote by $\bar{\sigma}(l) \rightarrow r_\sigma$ the rewrite step $\bar{\sigma}(l) \rightarrow \bar{\sigma}(r)$ when only the first step of the substitution applied to r .

Definition 5.2. *Let $l \rightarrow r$ be a rewrite rule, $\bar{\sigma}$ a valuation, and $p \in \mathcal{P}os(\bar{\sigma}(l))$. Suppose $u : \bar{\sigma}(l) \rightarrow r_\sigma$. 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 $\mathcal{Pos}(r) \subseteq \mathcal{Pos}(r_\sigma)$ by the notation of positions in subterms of the form $(\lambda \mathbf{x}.s)(t_1, \dots, t_n)$. From this it follows that $P \subseteq \mathcal{Pos}(r_\sigma)$.

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

Definition 5.3. Let $u : C[\bar{\sigma}(l)] \rightarrow C[\bar{\sigma}(r)]$ be a rewrite step, such that p is the position of the hole in $C[\square]$, and let $q \in \mathcal{Pos}(C[\bar{\sigma}(l)])$. 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 $q'/^1 u'$ with $u' : \bar{\sigma}(l) \rightarrow r_\sigma$ across complete development of the parallel β -redexes in r_σ .

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

Definition 5.4. Let $s_0 \twoheadrightarrow^\alpha s_\alpha$ and let $P \subseteq \mathcal{Pos}(s_0)$. 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 $P/(s_0 \twoheadrightarrow s_1) = \bigcup_{p \in P} p/(s_0 \rightarrow s_1)$,
- if $\alpha = \beta + 1$, then $P/(s_0 \twoheadrightarrow^{\beta+1} s_{\beta+1}) = (P/(s_0 \twoheadrightarrow^\beta s_\beta))/(s_\beta \rightarrow s_{\beta+1})$,
- 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 $l \rightarrow r$ 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 $l \rightarrow r$. 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.

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 .

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 $p \in \mathcal{Pos}(s)$ or (r, p, q) with r a right-hand side of a rewrite rule, $p \in \mathcal{Pos}(r)$, and $q = p_u$ with $u \in \mathcal{U}$. 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 $i \in \mathcal{Pos}(s|_p) \cap \mathbb{N}$ 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 $l \rightarrow r$ 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 an abstraction $[x]$ occurring in the left-hand side of the rule $l \rightarrow r$ of a redex $u \in \mathcal{U}$, 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 , $\text{root}(r|_{p'}) = Z$, 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 $\text{root}(r|_p)$ is not a meta-variable, then for some $i \in \mathcal{Pos}(r|_p) \cap \mathbb{N}$ the node n' is labelled $(r, p \cdot i, p_u)$ and the edge from n to n' is labelled i ,
2. if $\text{root}(r|_p)$ is a meta-variable Z , then the node n' is labelled $(s, q \cdot q')$ and the edge from n to n' is unlabelled, such that $l \rightarrow r$ is the rewrite rule employed in u , q is the position of u in s , and 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 P as a (possibly infinite) sequence of alternating nodes and edges $P = n_1 e_1 n_2 \dots$

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

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

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

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, $s \rightarrow t$ is a complete development.

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

$$(s, \epsilon) \rightarrow (r, \epsilon, \epsilon) \rightarrow (s, 10) \rightarrow_1 (s, 101) \rightarrow (r, 1, \epsilon) \rightarrow_1 (r, 11, \epsilon) \\ \rightarrow (s, 10) \rightarrow_1 (s, 101) \rightarrow (r, 111, \epsilon) \rightarrow (s, 2)$$

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

$$(t, \epsilon) \rightarrow_1 (t, 1) \rightarrow_1 (t, 11) \rightarrow_1 (t, 111).$$

The path projections of the maximal paths are respectively

$$\cdot \rightarrow_\epsilon \cdot \rightarrow_\epsilon g \rightarrow_1 \cdot \rightarrow_\epsilon g \rightarrow_1 \cdot \rightarrow_\epsilon g \rightarrow_1 \cdot \rightarrow_\epsilon \cdot \rightarrow_\epsilon a$$

and

$$g \rightarrow_1 g \rightarrow_1 g \rightarrow_1 a.$$

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

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 surjection. 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 definition of meta-terms. \square

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. *If no path projection occurring in $\mathcal{P}(s, \mathcal{U})$ contains an infinite sequences of unlabelled nodes and ϵ -labelled edges, then we say that \mathcal{U} has the finite jumps property. Moreover, we say that a term t matches $\mathcal{P}(s, \mathcal{U})$, if, for all $\phi(P) \in \mathcal{P}(s, \mathcal{U})$, and for all prefixes of $\phi(P)$ ending in a node n labelled f , we have that $\text{root}(t|_p) = f$, where p is the concatenation of the edge labels in the prefix (starting at the first node of $\phi(P)$ 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. The proof is identical to the proof of Proposition 12.5.8 in [3]. \square

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 $p \in \text{Pos}(s)$, 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 (Sketch). The proof is identical to the proof of Proposition 12.5.9 in [3], except that Lemma 5.9 is employed instead of tracing. \square

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 $s \rightarrow t$ 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). \square

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). \square

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 $i\lambda c$ remain open: Does there exist a notion of meaningless terms [11] that allows for the construction of Böhm-like trees? Can we prove a partial confluence property [2, 3, 11] 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, TCS **83** (1991) 71–96
2. Kennaway, R., Klop, J.W., Sleep, R., de Vries, F.J.: Transfinite reductions in orthogonal term rewriting systems. I&C **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. TCS **175** (1997) 93–125
5. Klop, J.W.: Combinatory Reduction Systems. PhD thesis, Rijksuniversiteit Utrecht (1980)
6. Klop, J.W., van Oostrom, V., van Raamsdonk, F.: Combinatory reduction systems: introduction and survey. TCS **121** (1993) 279–308
7. Arnold, A., Nivat, M.: The metric space of infinite trees. Algebraic and topological properties. Fundamenta Informaticae **3** (1980) 445–476
8. Barendregt, H.P.: The Lambda Calculus: Its Syntax and Semantics. Second edn. Elsevier Science (1985)
9. Hanus, M., Prehofer, C.: Higher-order narrowing with definitional trees. In Ganzinger, H., ed.: Proc. of the 7th Int. Conf. on Rewriting Techniques and Applications (RTA'96). Volume 1103 of LNCS., Springer-Verlag (1996) 138–152
10. van Oostrom, V.: Higher-order families. In Ganzinger, H., ed.: Proc. of the 7th Int. Conf. on Rewriting Techniques and Applications (RTA '96). Volume 1103 of LNCS., Springer-Verlag (1996) 392–407
11. Kennaway, R., van Oostrom, V., de Vries, F.J.: Meaningless terms in rewriting. The Journal of Functional and Logic Programming **1** (1999)