

PARTIAL ORDER INFINITARY TERM REWRITING AND BÖHM TREES

PATRICK BAHR

Department of Computer Science, University of Copenhagen
Universitetsparken 1, 2100 Copenhagen, Denmark
URL: <http://www.diku.dk/~paba>
E-mail address: paba@diku.dk

ABSTRACT. We investigate an alternative model of infinitary term rewriting. Instead of a metric, a partial order on terms is employed to formalise (strong) convergence. We compare this partial order convergence of orthogonal term rewriting systems to the usual metric convergence of the corresponding Böhm extensions. The Böhm extension of a term rewriting system contains additional rules to equate so-called root-active terms. The core result we present is that reachability w.r.t. partial order convergence coincides with reachability w.r.t. metric convergence in the Böhm extension. This result is used to show that, unlike in the metric model, orthogonal systems are infinitarily confluent and infinitarily normalising in the partial order model. Moreover, we obtain, as in the metric model, a compression lemma. A corollary of this lemma is that reachability w.r.t. partial order convergence is a conservative extension of reachability w.r.t. metric convergence.

1. Introduction

The study of infinitary term rewriting as a discipline to investigate infinitely long reductions on terms is mostly based on a metric model [Der91]. Other models, using for example general topological spaces [Rod98] or partial orders [Cor93, Blo04], were mainly considered to pursue quite specific purposes. Since in the metric model, even for orthogonal systems, infinitary rewriting lacks a number of important properties such as *compression* and *infinitary confluence* [Sim04], a stricter variant of convergence, so-called *strong convergence* [Ken95] was considered.

However, even strong convergence does not provide infinitary confluence for all orthogonal term rewriting systems and does not admit *complete developments* for arbitrary sets of redexes. This has been resolved by introducing a notion of meaningless terms [Ken99]. Having this notion, a term rewriting system can be augmented with rules which essentially allow rewriting meaningless terms to a fresh constant \perp . When starting with an orthogonal system, the resulting system, called *Böhm extension*, is both infinitarily normalising and infinitarily confluent w.r.t. strong convergence.

1998 ACM Subject Classification: F.4.2.

Key words and phrases: infinitary term rewriting, Böhm trees, partial order, confluence, normalisation.



In this paper we present a partial order model of strongly convergent reductions. We show that for orthogonal systems, reachability with this notion of convergence is equivalent to reachability according to metric strong convergence of the corresponding Böhm extensions w.r.t. the least set of meaningless terms, the *root-active* terms. As corollaries we thus obtain infinitary confluence and infinitary normalisation of partial order convergence. Moreover, we can show that this model also enjoys the compression property and admits arbitrary complete developments.

Related Work. This study of strong partial order convergence is inspired by Blom [Blo04] who investigated strong partial order convergence in lambda calculus and compared it to strong metric convergence. Similarly to our findings for orthogonal term rewriting systems, Blom has shown for lambda calculus that reachability in the metric model coincides with reachability in the partial order model modulo equating so-called 0-undefined terms.

Also Corradini [Cor93] studied a partial order model. However, he uses it to develop a theory of parallel reductions which allows simultaneous contraction of a set of mutually independent redexes of left-linear rules. To this end, Corradini defines the semantics of redex contraction in a non-standard way by allowing a partial matching of left-hand sides. Our definition of complete developments also provides, at least for orthogonal systems, a notion of parallel reductions but does so using the standard semantics.

2. Preliminaries

We assume the reader to be familiar with the basic theory of ordinal numbers, orders and topological spaces [Kel55], as well as term rewriting [Ter03]. In the following, we briefly recall the most important notions.

Transfinite Sequences. We use $\alpha, \beta, \gamma, \lambda, \iota$ to denote ordinal numbers. A *transfinite sequence* (or simply called *sequence*) S of length α in a set A , written $(a_\iota)_{\iota < \alpha}$, is a function from α to A with $\iota \mapsto a_\iota$ for all $\iota \in \alpha$. We use $|S|$ to denote the length α of S . If α is a limit ordinal, then S is called *open*. Otherwise, it is called *closed*. If α is a finite ordinal, then S is called *finite*. Otherwise, it is called *infinite*.

The *concatenation* $(a_\iota)_{\iota < \alpha} \cdot (b_\iota)_{\iota < \beta}$ of two sequences is the sequence $(c_\iota)_{\iota < \alpha + \beta}$ with $c_\iota = a_\iota$ for $\iota < \alpha$ and $c_{\alpha + \iota} = b_\iota$ for $\iota < \beta$. A sequence S is a (proper) *prefix* of a sequence T , denoted $S \leq T$ (resp. $S < T$), if there is a (non-empty) sequence S' with $S \cdot S' = T$. The prefix of T of length β is denoted $T|_\beta$. The relation \leq forms a complete semilattice.

Let $S = (a_\iota)_{\iota < \alpha}$ be a sequence. A sequence $T = (b_\iota)_{\iota < \beta}$ is called a *subsequence* of S if there is a monotone function $f: \beta \rightarrow \alpha$ such that $b_\iota = a_{f(\iota)}$ for all $\iota < \beta$. To indicate this, we write S/f for the subsequence T . If $f(\iota) = f(0) + \iota$ for all $\iota < \beta$, then S/f is called a *segment* of S . That is, T is a segment of S iff there are two sequences T_1, T_2 such that $S = T_1 \cdot T \cdot T_2$. We write $S|_{[\beta, \gamma)}$ for the segment S/f , where $f: \alpha' \rightarrow \alpha$ is the mapping defined by $f(\iota) = \beta + \iota$ for all $\iota < \alpha'$, with α' the unique ordinal with $\gamma = \beta + \alpha'$. Note that in particular $S|_\alpha = S|_{[0, \alpha)}$ for each sequence S and ordinal $\alpha \leq |S|$.

Partial Orders. A *partial order* \leq on a set A is a binary relation on A that is *transitive*, *reflexive*, and *antisymmetric*. A partial order \leq on A is called a *complete semilattice* if it has a *least element*, every *directed subset* D of A has a *least upper bound* (*lub*) $\sqcup D$, and every subset of A having an upper bound also has a least upper bound. Hence, complete semilattices also admit a *greatest lower bound* (*glb*) $\sqcap B$ for every *non-empty* subset B of A . In particular, this means that for any non-empty sequence $(a_\iota)_{\iota < \alpha}$ in a complete semilattice, its *limit inferior*, defined by $\liminf_{\iota \rightarrow \alpha} a_\iota = \sqcup_{\beta < \alpha} \left(\sqcap_{\beta \leq \iota < \alpha} a_\iota \right)$, always exists.

With the prefix order \leq on sequences we can generalise concatenation to arbitrary sequences of sequences: Let $(S_\iota)_{\iota < \alpha}$ be a sequence of sequences in a common set. The concatenation of $(S_\iota)_{\iota < \alpha}$, written $\prod_{\iota < \alpha} S_\iota$, is recursively defined as the empty sequence ε if $\alpha = 0$, $(\prod_{\iota < \alpha'} S_\iota) \cdot S_{\alpha'}$ if $\alpha = \alpha' + 1$, and $\sqcup_{\gamma < \alpha} \prod_{\iota < \gamma} S_\iota$ if α is a limit ordinal.

Term Rewriting Systems. Unlike in the traditional framework of term rewriting, we consider the set $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ of *infinitary terms* (or simply *terms*) over some *signature* Σ and a countably infinite set \mathcal{V} of variables. The set $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ is defined as the *greatest* set T such that, for each element $t \in T$, we either have $t \in \mathcal{V}$ or $t = f(t_1, \dots, t_k)$, where $k \geq 0$, $f \in \Sigma^{(k)}$, and $t_1, \dots, t_k \in T$. We consider $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ as a superset of the set $\mathcal{T}(\Sigma, \mathcal{V})$ of *finite terms*. For a term $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ we use the notation $\mathcal{P}(t)$ to denote the *set of positions* in t . For terms $s, t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ and a position $\pi \in \mathcal{P}(t)$, we write $t|_\pi$ for the *subterm* of t at π , $t(\pi)$ for the symbol in t at π , and $t[s]_\pi$ for the term t with the subterm at π replaced by s . Two terms s and t are said to *coincide* in a set of positions $P \subseteq \mathcal{P}(s) \cap \mathcal{P}(t)$ if $s(\pi) = t(\pi)$ for all $\pi \in P$. A position is also called an *occurrence* if the focus lies on the subterm at that position rather than the position itself. Two positions π_1, π_2 are called *disjoint* if neither $\pi_1 \leq \pi_2$ nor $\pi_2 \leq \pi_1$.

On $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ a distance function \mathbf{d} can be defined by $\mathbf{d}(s, t) = 0$ if $s = t$ and $\mathbf{d}(s, t) = 2^{-k}$ if $s \neq t$, where k is the minimal depth at which s and t differ. The pair $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$ is known to form a complete ultrametric space [Arn80]. *Partial terms*, i.e. terms over signature $\Sigma_\perp = \Sigma \uplus \{\perp\}$, can be endowed with a relation \leq_\perp by defining $s \leq_\perp t$ iff s can be obtained from t by replacing some subterm occurrences in t by \perp . The pair $(\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V}), \leq_\perp)$ is known to form a complete semilattice [Kah93]. For a partial term $t \in \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$ we use the notation $\mathcal{P}_\perp(t)$ and $\mathcal{P}_\Sigma(t)$ for the set $\{\pi \in \mathcal{P}(t) \mid t(\pi) \neq \perp\}$ of non- \perp positions resp. the set $\{\pi \in \mathcal{P}(t) \mid t(\pi) \in \Sigma\}$ of positions of function symbols. To explicitly distinguish them from partial terms, we call terms in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ *total*.

A *term rewriting system* (TRS) \mathcal{R} is a pair (Σ, R) consisting of a signature Σ and a set R of *term rewrite rules* of the form $l \rightarrow r$ with $l \in \mathcal{T}^\infty(\Sigma, \mathcal{V}) \setminus \mathcal{V}$ and $r \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ such that all variables in r are contained in l . Note that this notion of a TRS deviates slightly from the standard notion of TRSs in the literature on infinitary rewriting [Ken03] in that it allows infinite terms on the left-hand side of rewrite rules! This generalisation will be necessary to accommodate Böhm extensions. TRSs having only finite left-hand sides are called *left-finite*.

A term t is called *linear* if each variable occurs at most once in t . A TRS \mathcal{R} is called *left-linear* if the left-hand side of every rule in \mathcal{R} is linear. A TRS \mathcal{R} is called *orthogonal* if it is left-linear and has no *critical pairs*.

As in the finitary case, every TRS \mathcal{R} defines a rewrite relation $\rightarrow_{\mathcal{R}}$:

$$s \rightarrow_{\mathcal{R}} t \iff \exists \pi \in \mathcal{P}(s), l \rightarrow r \in R, \sigma: s|_\pi = l\sigma, t = s[r\sigma]_\pi$$

We write $s \rightarrow_{\pi, \rho} t$ in order to indicate the applied rule ρ and the position π . The subterm $s|_{\pi}$ is called a ρ -*redex* or simply *redex*, $r\sigma$ its *contractum*, and $s|_{\pi}$ is said to be *contracted* to $r\sigma$.

3. Metric Infinitary Term Rewriting

In this section we briefly recall the metric model of infinitary term rewriting [Ken95] and some of its properties. We will use the metric model in two ways: Firstly, it will serve as a yardstick to compare the partial order model to. But most importantly, we will use known results for metric infinitary rewriting and transfer them to the partial order model. In order to accomplish the latter, we will make use of Theorem 5.6 which we shall present at the end of Section 5.

At first we have to make clear what a *reduction* in our setting of infinitary rewriting is:

Definition 3.1 (reduction (step)). Let \mathcal{R} be a TRS. A *reduction step* φ in \mathcal{R} is a tuple (s, π, ρ, t) such that $s \rightarrow_{\pi, \rho} t$; we also write $\varphi: s \rightarrow_{\pi, \rho} t$. A *reduction* S in \mathcal{R} is a sequence $(\varphi_i)_{i < \alpha}$ of reduction steps such that there is a sequence $(t_i)_{i < \hat{\alpha}}$ of terms, with $\hat{\alpha} = \alpha$ if S is open, $\hat{\alpha} = \alpha + 1$ if S is closed, such that $\varphi_i: t_i \rightarrow t_{i+1}$. If S is finite, we write $S: t_0 \rightarrow^* t_{\alpha}$.

Note that this notion of reductions does only make sense for sequences of length at most ω . For longer reductions, the ω -th step is not related to the preceding steps of the reduction. This holds in general for all reductions steps indexed by a limit ordinal. An appropriate definition of a reduction of length beyond ω requires a notion of continuity to bridge the gaps that arise at limit ordinals. In this section we look at the notion of *strong continuity* modelled by the metric on terms. Since we are not interested in weak continuity [Der91] here, we refer to this notion simply as *continuity*, or *m-continuity* to distinguish it from continuity in the partial order model that we will present in Section 5.

It is important to understand that a reduction is a *sequence of reduction steps* rather than just a sequence of terms. This is crucial for a proper definition of strong continuity, which also depends on where contractions take place:

Definition 3.2 (*m-continuity/-convergence*). Let \mathcal{R} be a TRS and $S = (\varphi_i: t_i \rightarrow_{\pi_i} t_{i+1})_{i < \alpha}$ a non-empty reduction in \mathcal{R} . The reduction S is called *m-continuous* if $\lim_{i \rightarrow \lambda} t_i = t_{\lambda}$, and the sequence $(|\pi_i|)_{i < \lambda}$ of contraction depths tends to infinity for each limit ordinal $\lambda < \alpha$. Provided it is *m-continuous*, S is said to *m-converge* to t , written $S: t_0 \xrightarrow{m} t$, if S is closed and $t = t_{\alpha}$ or if $(|\pi_i|)_{i < \alpha}$ tends to infinity and $t = \lim_{i \rightarrow \alpha} t_i$. In this case we also say that t is *m-reachable* from t_0 . In order to indicate the length of S and the TRS \mathcal{R} , we write $S: t_0 \xrightarrow{m, \mathcal{R}}^{\alpha} t$. The empty reduction ε is considered *m-continuous* and *m-convergent* for any start and end term, i.e. $\varepsilon: t \xrightarrow{m} \dots$ and $\varepsilon: t \xrightarrow{m} t$ for all $t \in \mathcal{T}(\Sigma, \mathcal{V})$.

For a reduction to be *m-continuous*, each open *proper* prefix of the underlying sequence of terms must converge to the term following next in the sequence. Additionally, the depth at where contractions take place has to tend to infinity for each of the reduction's open proper prefixes. In contrast, *m-convergence* requires the above conditions to hold for *all* open prefixes, i.e. including the whole reduction itself provided it is open. For example, considering the rule $a \rightarrow f(a)$, the reduction $g(a) \rightarrow g(f(a)) \rightarrow g(f(f(a))) \rightarrow \dots$ *m-converges* to the infinite term $g(f^{\omega})$. Note that *m-convergence* implies *m-continuity*. Hence, only meaningful, i.e. *m-continuous*, reductions can be *m-convergent*. On the other hand not every *m-continuous* reduction is also *m-convergent*. Having the rule $g(x) \rightarrow g(f(x))$

instead, the reduction $g(a) \rightarrow g(f(a)) \rightarrow g(f(f(a))) \rightarrow \dots$ is trivially m -continuous but is now not m -convergent.

If we only want to express that there is some reduction S with, say, $S: s \rightsquigarrow t$, then we simply write $s \rightsquigarrow t$. An example for this notation can be seen in the following phrasing of the Compression Lemma [Ken95]:

Theorem 3.3 (Compression Lemma). *For each left-linear, left-finite TRS, $s \rightsquigarrow t$ implies $s \rightsquigarrow^{\leq \omega} t$.*

As an easy corollary we obtain that the final term of an m -converging reduction can be approximated arbitrarily accurately by a finite reduction:

Corollary 3.4 (finite approximation). *Let \mathcal{R} be a left-linear, left-finite TRS and $s \rightsquigarrow t$. Then, for each depth $d \in \mathbb{N}$, there is a finite reduction $s \rightarrow^* t'$ such that t and t' coincide up to depth d , i.e. $\mathbf{d}(t, t') < 2^{-d}$.*

Proof. Assume $s \rightsquigarrow t$. By Theorem 3.3, there is a reduction $S: s \rightsquigarrow^{\leq \omega} t$. If S is of finite length, then we are done. If $S: s \rightsquigarrow^{\omega} t$, then, by m -convergence, there is some $n < \omega$ such that all reductions steps in S after n take place at a depth greater than d . Consider $S|_n: s \rightarrow^* t'$. It is clear that t and t' coincide up to depth d . ■

An important difference of m -converging reductions and finite reductions is the confluence of orthogonal systems. In contrast to finite reachability, m -reachability of orthogonal TRSs does not necessarily have the diamond property, i.e. orthogonal systems are confluent but not infinitarily confluent [Ken95]:

Example 3.5 (failure of infinitary confluence). Consider the orthogonal TRS consisting of the *collapsing* rules $\rho_1: f(x) \rightarrow x$ and $\rho_2: g(x) \rightarrow x$ and the infinite term $t = g(f(g(f(\dots))))$. We then obtain the reductions $S: t \rightsquigarrow g^\omega$ and $T: t \rightsquigarrow f^\omega$ by successively contracting all ρ_1 - resp. ρ_2 -redexes. However, there is no term s such that $g^\omega \rightsquigarrow s \rightsquigarrow^m f^\omega$ as both g^ω and f^ω can only be rewritten to themselves, respectively.

In the following sections we discuss two different methods for obtaining an appropriate notion of transfinite reachability which actually has the diamond property.

4. Meaningless Terms and Böhm Trees

Meaningless terms, as formalised by Kennaway et al. [Ken99], are terms which can be considered meaningless because, from a term rewriting perspective, they cannot be distinguished from one another and they do not contribute any information to any computation. For orthogonal TRSs, one such set of terms, in fact the least such set, is the set of *root-active* terms [Ken99]:

Definition 4.1 (root-activeness). Let \mathcal{R} be a TRS and $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$. Then t is called *root-active* if for each reduction $t \rightarrow^* t'$, there is a reduction $t' \rightarrow^* s$ to a redex s . The set of all root-active terms of \mathcal{R} is denoted $\mathcal{RA}_{\mathcal{R}}$ or simply \mathcal{RA} if \mathcal{R} is clear from the context.

Intuitively speaking, as the name already suggests, root-active terms are terms that can be contracted at the root arbitrarily often, e.g. the terms f^ω and g^ω from Example 3.5.

In this paper we are only interested in this particular set of meaningless terms. So for the sake of brevity we restrict our discussion in this section to \mathcal{RA} instead of the original more general axiomatic treatment by Kennaway et al. [Ken99].

Since, operationally, root-active terms cannot be distinguished from each other it is appropriate to equate them [Ken99]. This can be done by introducing a new constant symbol \perp and making each root-active term equal to \perp . By adding rules which enable rewriting root-active terms to \perp , this can be encoded into an existing TRS [Ken99]:

Definition 4.2 (Böhm extension). Let \mathcal{R} be a TRS over Σ , and $\mathcal{U} \subseteq \mathcal{T}^\infty(\Sigma, \mathcal{V})$.

- (i) A term $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ is called a \perp, \mathcal{U} -instance of a term $s \in \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$ if t can be obtained from s by replacing each occurrence of \perp in s with some term in \mathcal{U} .
- (ii) \mathcal{U}_\perp is the set of terms in $\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$ that have a \perp, \mathcal{U} -instance in \mathcal{U} .
- (iii) The *Böhm extension* of \mathcal{R} w.r.t. \mathcal{U} is the TRS $\mathcal{B}_{\mathcal{R}, \mathcal{U}} = (\Sigma_\perp, R \cup B)$, where

$$B = \{t \rightarrow \perp \mid t \in \mathcal{U}_\perp \setminus \{\perp\}\}$$

We write $s \rightarrow_{\mathcal{U}, \perp} t$ for a reduction step using a rule in B . If \mathcal{R} and \mathcal{U} are clear from the context, we simply write \mathcal{B} and \rightarrow_\perp instead of $\mathcal{B}_{\mathcal{R}, \mathcal{U}}$ and $\rightarrow_{\mathcal{U}, \perp}$, respectively.

A reduction that is m -converging in the Böhm extension \mathcal{B} is called *Böhm-converging*. A term t is called *Böhm-reachable* from s if there is a Böhm-converging reduction from s to t .

Note that, for orthogonal TRSs, $\mathcal{R}\mathcal{A}$ is closed under substitutions and, hence, so is $\mathcal{R}\mathcal{A}_\perp$ [Ken99]. Therefore, whenever $C[t] \rightarrow_{\mathcal{R}\mathcal{A}, \perp} C[\perp]$, we can assume that $t \in \mathcal{R}\mathcal{A}_\perp$.

It is at this point where we, in fact, need the generality of allowing infinite terms on the left-hand side of rewrite rules: The additional rules of a Böhm extension allow possibly infinite terms $t \in \mathcal{U}_\perp \setminus \{\perp\}$ on the left-hand side.

Theorem 4.3 (infinitary confluence of Böhm-converging reductions, [Ken99]). *Let \mathcal{R} be an orthogonal, left-finite TRS. Then the Böhm extension \mathcal{B} of \mathcal{R} w.r.t. $\mathcal{R}\mathcal{A}$ is infinitarily confluent, i.e. $s_1 \xleftarrow{m}_{\mathcal{B}} t \xrightarrow{m}_{\mathcal{B}} s_2$ implies $s_1 \xrightarrow{m}_{\mathcal{B}} t' \xleftarrow{m}_{\mathcal{B}} s_2$.*

The lack of confluence for m -converging reductions is resolved in Böhm extensions by allowing (sub-)terms, which were previously not joinable, to be contracted to \perp . Returning to Example 3.5, g^ω and f^ω can be rewritten to \perp as both terms are root-active.

Theorem 4.4 (infinitary normalisation of Böhm-converging reductions, [Ken99]). *Let \mathcal{R} be an orthogonal, left-finite TRS. Then the Böhm extension \mathcal{B} of \mathcal{R} w.r.t. $\mathcal{R}\mathcal{A}$ is infinitarily normalising, i.e. for each term t there is a \mathcal{B} -normal form Böhm-reachable from t .*

This means that each term t of an orthogonal, left-finite TRS \mathcal{R} has a unique normal form in $\mathcal{B}_{\mathcal{R}, \mathcal{R}\mathcal{A}}$. This normal form is called the *Böhm tree* of t (w.r.t. $\mathcal{R}\mathcal{A}$) [Ken99].

5. Partial Order Infinitary Rewriting

In this section we define an alternative model of infinitary term rewriting which uses the partial order on terms to formalise (strong) convergence of transfinite reductions. To this end we will turn to partial terms which, like in the setting of Böhm extensions, have an additional special symbol \perp . The result will be a more fine-grained notion of convergence in which, intuitively speaking, a reduction can be diverging in some positions but at the same time converging in other positions. The “diverging parts” are then indicated by a \perp -occurrence in the final term of the reduction:

Example 5.1. Consider the TRS consisting of the rules $h(x) \rightarrow h(g(x)), c \rightarrow g(c)$ and the term $t = f(a, c)$. In this system, we have the reduction

$$S: f(h(a), b) \rightarrow f(h(g(a)), b) \rightarrow f(h(g(a)), g(b)) \rightarrow f(h(g(g(a))), g(b)) \rightarrow \dots$$

which alternately contracts the redex in the left and in the right argument of f .

Reduction S does not m -converge as the depth at which contractions are performed does not tend to infinity. However, this does only happen in the left argument of f , not in the other one. With the notion of p -convergence, we will discover S to be p -converging to the term $f(\perp, g^\omega)$:

Definition 5.2 (p -continuity/ p -convergence). Let $\mathcal{R} = (\Sigma, R)$ be a TRS and $S = (\varphi_\iota: t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \alpha}$ a non-empty reduction in $\mathcal{R}' = (\Sigma_\perp, R)$. The reduction S is called p -continuous if $\liminf_{\iota \rightarrow \lambda} c_\iota = t_\lambda$ for each limit ordinal $\lambda < \alpha$, where $c_\iota = t_\iota[\perp]_{\pi_\iota}$. Each c_ι is called the *context* of the reduction step φ_ι . Provided it is p -continuous, S is said to p -converge to t , written $S: t_0 \xrightarrow{p} t$ if S is closed and $t = t_{\alpha+1}$ or if $t = \liminf_{\iota \rightarrow \alpha} c_\iota$. In this case we also say that t is p -reachable from t_0 . In order to indicate the length of S , we write $S: t_0 \xrightarrow{p}^\alpha t$. The empty reduction ε is considered p -continuous and p -convergent for any start and end term.

What makes this notion of p -convergence *strong*, similar to the notion of m -convergence we are considering here, is the choice of taking the contexts $t_\iota[\perp]_{\pi_\iota}$ for defining the limit behaviour of reductions instead of the whole terms t_ι . The context $t_\iota[\perp]_{\pi_\iota}$ provides a conservative underapproximation of the shared structure $t_\iota \sqcap t_{\iota+1}$ of two consecutive terms t_ι and $t_{\iota+1}$. In fact, $t_\iota[\perp]_{\pi_\iota} \leq_\perp t_\iota \sqcap t_{\iota+1}$. Returning to Example 5.1, we can observe that with the weaker notion of p -convergence, i.e. using t_ι instead of $t_\iota[\perp]_{\pi_\iota}$ for the limit behaviour, reduction S would p -converge to $f(h(g^\omega), g^\omega)$ instead of $f(\perp, g^\omega)$.

This approach is analogous to the metric notion of strong convergence which requires $|\pi_\iota|$ to tend to infinity, i.e. $2^{-|\pi_\iota|}$ to tend to 0. However, $2^{-|\pi_\iota|}$ is an overapproximation of the actual difference $\mathbf{d}(t_\iota, t_{\iota+1})$ of two consecutive terms t_ι and $t_{\iota+1}$, i.e. $2^{-|\pi_\iota|} \geq \mathbf{d}(t_\iota, t_{\iota+1})$.

Note that we have to consider reductions over the extended signature Σ_\perp , i.e. reductions containing partial terms. Thus, from now on, we assume reductions in a TRS over Σ to be implicitly over Σ_\perp . When we want to make it explicit that a reduction S contains only total terms, we say that S is *total*. When we say that $S: s \xrightarrow{p} t$ is total, we mean that both the reduction S and the final term t are total.¹

Due to the partial order \leq_\perp on partial terms being a complete semilattice, the limit inferior is defined for any sequence of partial terms. Hence, any p -continuous reduction is also p -convergent. This is one of the major differences to m -convergence/ m -continuity. Nevertheless, p -convergence is a meaningful notion of convergence. The final term of a p -convergent reduction contains a \perp subterm at each position at which the reduction is “locally diverging” as we have seen in Example 5.1. We will call these positions *volatile*:

Definition 5.3 (volatility). Let \mathcal{R} be a TRS and $S = (t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \lambda}$ an open p -converging reduction in \mathcal{R} . A position π is said to be *volatile* in S if, for each ordinal $\beta < \lambda$, there is some $\beta \leq \gamma < \lambda$ such that $\pi_\gamma = \pi$. If π is volatile in S and no proper prefix of π is volatile in S , then π is called *outermost-volatile*.

In Example 5.1 the position 0 is outermost-volatile in the reduction S . One can show that \perp subterms are indeed created precisely at outermost-volatile positions [Bah09]:

Lemma 5.4 (\perp subterms in open reductions). *Let \mathcal{R} be a TRS and $S = (t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \lambda}$ an open reduction in \mathcal{R} p -converging to t_λ . Then, for every position π , we have the following:*

- (i) *If π is volatile in S , then $\pi \notin \mathcal{P}_\perp(t_\lambda)$.*

¹Note that if S is open, the final term t is not explicitly contained in S . Hence, the totality of S does not necessarily imply the totality of t .

- (ii) $t_\lambda(\pi) = \perp$ iff
 - (a) π is outermost-volatile in S , or
 - (b) there is some $\beta < \lambda$ such that $t_\beta(\pi) = \perp$ and $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \lambda$.
- (iii) Let t_ι be total for all $\iota < \lambda$. Then $t_\lambda(\pi) = \perp$ iff π is outermost-volatile in S .

From this we can deduce that the absence of volatile positions is equivalent to the totality of a p -converging reduction:

Lemma 5.5 (total reductions). *Let \mathcal{R} be a TRS, s a total term in \mathcal{R} , and $S: s \xrightarrow{\mathcal{R}} t$. $S: s \xrightarrow{p} t$ is total iff no prefix of S has a volatile position.*

Proof. The “only if” direction follows straightforwardly from Lemma 5.4.

We prove the “if” direction by induction on the length of S . If $|S| = 0$, then the totality of S follows from the assumption of s being total. If $|S|$ is a successor ordinal, then the totality of S follows from the induction hypothesis since single reduction steps preserve totality. If $|S|$ is a limit ordinal, then the totality of S follows from the induction hypothesis using Lemma 5.4. ■

The following theorem is the central tool for transferring results for m -convergent reductions to the realm of p -convergence:

Theorem 5.6 (total p -convergence = m -convergence). *For every reduction S in a TRS, $S: s \xrightarrow{p} t$ is total iff $S: s \xrightarrow{m} t$.*

We won’t go into the details of the proof of Theorem 5.6 here but instead refer to [Bah09]. The key for the proof are the following two observations: At first, the limit inferior and the limit of a sequence of total terms coincide whenever the limit exists or the limit inferior is a total term. Secondly, for each open m -converging reduction $S = (\varphi: t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \lambda}$ the limit inferior of the sequence of terms $(t_\iota)_{\iota < \lambda}$ coincides with the limit inferior of the sequence of contexts $(t_\iota[\perp]_{\pi_\iota})_{\iota < \lambda}$ since the \perp ’s in $(t_\iota[\perp]_{\pi_\iota})_{\iota < \lambda}$ are “pushed down” deeper and deeper due to the m -convergence of S .

6. Complete Developments

There are several methods to show (finitary) confluence of orthogonal systems. A quite instructive technique uses notions of *residuals* and *complete developments* [Ter03]. Intuitively speaking, the residuals of a set of redexes are the remains of this set of redexes after a reduction, and a complete development of a set of redexes is a reduction which only contracts residuals of these redexes and ends in a term with no residuals. Kennaway et al. [Ken95] have lifted these notions to (metric) infinitary term rewriting. However, in contrast to the finitary setting, complete developments do not always exist in infinitary orthogonal term rewriting, e.g. for the term f^ω from Example 3.5 and the set of all redex occurrences in it.

In this section we define residuals and complete developments in the setting of partial order infinitary term rewriting and show that complete developments do always exist for orthogonal TRSs and converge to a unique term. Having this, we can show the Infinitary Strip Lemma which is a crucial tool for proving our main result. However, since the proofs of these results are rather technical and tedious we will not provide the full proofs here but rather refer the interested reader to the author’s thesis [Bah09] where detailed proofs for all results in this section can be found.

At first we need to formalise the notion of residuals. It is virtually equivalent to the definition for m -convergence by Kennaway et al. [Ken95]:

Definition 6.1 (descendants, residuals). Let \mathcal{R} be a TRS, $S: t_0 \xrightarrow{\mathcal{R}}^\alpha t_\alpha$, and $U \subseteq \mathcal{P}_{\perp}(t_0)$. The *descendants* of U by S , denoted $U//S$, is the set of positions in t_α inductively defined as follows:

- (a) If $\alpha = 0$, then $U//S = U$.
- (b) If $\alpha = 1$, i.e. $S: t_0 \rightarrow_{\pi, \rho} t_1$ for some $\rho: l \rightarrow r$, take any $u \in U$ and define the set R_u as follows: If $\pi \not\leq u$, then $R_u = \{u\}$. If u is in the pattern of the ρ -redex, i.e. $u = \pi \cdot \pi'$ with $\pi' \in \mathcal{P}_{\Sigma}(l)$, then $R_u = \emptyset$. Otherwise, i.e. if $u = \pi \cdot w \cdot x$, with $l|_w \in \mathcal{V}$, then $R_u = \{\pi \cdot w' \cdot x \mid r|_{w'} = l|_w\}$. Define $U//S = \bigcup_{u \in U} R_u$.
- (c) If $\alpha = \alpha' + 1$, then $U//S = (U//S|_{\alpha'})//\varphi_{\alpha'}$, where $S = (\varphi_\iota)_{\iota < \alpha}$.
- (d) If α is a limit ordinal, then $U//S = \mathcal{P}_{\perp}(t_\alpha) \cap \liminf_{\iota \rightarrow \alpha} U//S|_\iota$
That is, $u \in U//S$ iff $u \in \mathcal{P}_{\perp}(t_\alpha)$ and $\exists \beta < \alpha \forall \beta \leq \iota < \alpha: u \in U//S|_\iota$

If, in particular, U is a set of redex occurrences, then $U//S$ is also called the set of *residuals* of U by S . Moreover, by abuse of notation, we write $u//S$ instead of $\{u\}//S$.

Clauses (a), (b) and (c) are as in the finitary setting. Clause (d) lifts the definition to the infinitary setting. However, the only difference to the definition of Kennaway et al. is, that we consider partial terms here. Yet, for technical reasons, the notion of descendants has to be restricted to non- \perp occurrences. Since \perp cannot be a redex, this is not a restriction for residuals, though.

As for finitary rewriting and metric infinitary rewriting, we have that residuals are always redexes and are pairwise disjoint if the original redexes are:

Proposition 6.2 ((disjoint) residuals). *Let \mathcal{R} be an orthogonal TRS, $S: s \xrightarrow{\mathcal{R}} t$ and U a set of redex occurrences in s . Then the following holds:*

- (i) $U//S$ is a set of redex occurrences in t .
- (ii) If the occurrences in U are pairwise disjoint, then so are the occurrences in $U//S$.

The property of residuals being redexes is, in fact, crucial for the concept of complete developments as it requires all residuals to be eventually contracted:

Definition 6.3 ((complete) development). Let \mathcal{R} be an orthogonal TRS, s a partial term in \mathcal{R} , and U a set of redex occurrences in s .

- (i) A *development* of U in s is a p -converging reduction $S: s \xrightarrow{\mathcal{R}}^\alpha t$ in which each reduction step $\varphi_\iota: t_\iota \rightarrow_{\pi_\iota} t_{\iota+1}$ contracts a redex at $\pi_\iota \in U//S|_\iota$ for $\iota < \alpha$.
- (ii) A development $S: s \xrightarrow{\mathcal{R}} t$ of U in s is called *complete*, denoted $S: s \xrightarrow{\mathcal{R}}_U t$, if $U//S = \emptyset$.

This is a straightforward generalisation of complete developments known from the finitary setting and coincides with the corresponding formalisation for metric infinitary rewriting [Ken95] if restricted to total terms. However, unlike in the metric setting, partial order infinitary rewriting admits complete developments for any orthogonal system:

Proposition 6.4 (complete developments). *Let \mathcal{R} be an orthogonal TRS, t a partial term in \mathcal{R} , and U a set of redex occurrences in t . Then U has a complete development in t .*

This result follows from the fact that every p -continuous reduction is also p -converging. Proving the above proposition simply amounts to devising a reduction strategy which eventually contracts all redexes. A *parallel-outermost* strategy achieves this.

Next we need to show that the final term of a complete development is uniquely defined by the initial set of redex occurrences U . Using a technique of *paths and jumps* similar to the one described by Kennaway and de Vries [Ken03], we can define for each partial term t ,




Figure 1: The Infinitary Strip Lemma.

set of redex occurrences U in t , and orthogonal TRS \mathcal{R} , a term $\mathcal{F}(t, U, \mathcal{R})$ that is the final term of a complete development of U in t :

Proposition 6.5 (unique p -convergence of complete developments). *Let \mathcal{R} be an orthogonal TRS, t a partial term in \mathcal{R} , and U a set of redex occurrences in t . Then each complete development of U in t p -converges to $\mathcal{F}(t, U, \mathcal{R})$.*

We can use the above result in order to show that descendants by complete developments are uniquely defined. To achieve this, one can use the well-known labelling technique that keeps track of descendants by means of syntactic methods (e.g. see [Ter03]):

Proposition 6.6 (unique descendants of complete developments). *Let \mathcal{R} be an orthogonal TRS, t a partial term in \mathcal{R} , and U a set of redex occurrences in t . Then, for each set $V \subseteq \mathcal{P}_{\setminus 1}(t)$ and two complete developments S and T of U in t , respectively, it holds that $V // S = V // T$.*

As a corollary we obtain that complete developments enjoy the diamond property:

Corollary 6.7 (diamond property of complete developments). *Let \mathcal{R} be an orthogonal TRS and $t \xrightarrow{U} t_1$ and $t \xrightarrow{V} t_2$ be two complete developments of U respectively V in t . Then t_1 and t_2 are joinable by complete developments $t_1 \xrightarrow{V // U} t'$ and $t_2 \xrightarrow{U // V} t'$.*

The result of this effort of analysing complete developments is the Infinitary Strip Lemma for p -convergence:

Proposition 6.8 (Infinitary Strip Lemma). *Let \mathcal{R} be an orthogonal TRS, $S: t_0 \xrightarrow{p} t_\alpha$, and $T: t_0 \xrightarrow{U} s_0$ a complete development of a set U of disjoint redex occurrences in t_0 . Then t_α and s_0 are joinable by $S/T: s_0 \xrightarrow{p} s_\alpha$ and a complete development $T/S: t_\alpha \xrightarrow{U // S} s_\alpha$.*

The idea of the construction of S/T and T/S is illustrated in Figure 1. Each U_i is the set of residuals of U by the reduction $S|_i$. Each arrow in the diagram represents a complete development of the indicated set of redex occurrences. In particular, each v_i indicates the redex occurrence contracted in the i -th step of S . The construction uses an induction on the length α of the horizontal reduction S . The case $\alpha = 0$ is trivial. For α a successor ordinal, the statement follows from the induction hypothesis using Corollary 6.7. For α a limit ordinal we can make use of the fact that, by Proposition 6.2 each U_i is a set of pairwise disjoint redex occurrences. The constructed reduction S/T is called the *projection* of S by T . Likewise, T/S is called the *projection* of T by S .

7. p -convergence and Böhm-convergence

In this section we shall show the core result of this paper: For orthogonal, left-finite TRSs, p -reachability and Böhm-reachability w.r.t. \mathcal{RA} coincide. As corollaries of that, leveraging the properties of Böhm-convergence, we obtain both infinitary normalisation and infinitary confluence of orthogonal systems in the partial order model. Moreover, we will show that p -convergence also satisfies the compression property.

The central step of the proof of the equivalence of both models of infinitary rewriting is an alternative characterisation of root-active terms which is captured by the following definition:

Definition 7.1 (destructiveness, fragility). Let \mathcal{R} be a TRS.

- (i) A reduction $S: t \twoheadrightarrow s$ is called *destructive* if ε is a volatile position in S .
- (ii) A partial term t in \mathcal{R} is called *fragile* if a destructive reduction starts in t .

Looking at the definition, fragility seems to be a more general concept than root-activeness: A term is fragile iff it admits a reduction in which infinitely often a redex at the root is contracted. For orthogonal TRSs, root-active terms are characterised in almost the same way. The difference is that only total terms are considered and that the stipulated reduction contracting infinitely many root redexes has to be of length ω . However, we shall show the set of total fragile terms to be equal to the set of root-active terms by establishing a compression lemma for destructive reductions.

Using Lemma 5.4 we can immediately derive the following alternative characterisations:

Fact 7.2 (destructiveness, fragility). Let \mathcal{R} be a TRS.

- (i) A reduction $S: s \twoheadrightarrow t$ is destructive iff S is open and $t = \perp$.
- (ii) A partial term t in \mathcal{R} is fragile iff there is an open p -convergent reduction $t \twoheadrightarrow \perp$.

Using this, we can establish that any p -convergent reduction can be simulated by a Böhm-convergent reduction w.r.t. total, fragile terms:

Proposition 7.3 (p -reachability implies Böhm-reachability). Let \mathcal{R} be a TRS, \mathcal{U} the set of fragile terms in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, and \mathcal{B} the Böhm extension of \mathcal{R} w.r.t. \mathcal{U} . Then, for each p -convergent reduction $s \twoheadrightarrow_{\mathcal{R}} t$, there is a Böhm-convergent reduction $s \twoheadrightarrow_{\mathcal{B}} t$.

Proof. Assume that there is a reduction $S = (t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \alpha}$ in \mathcal{R} that p -converges to t_α . We will construct an m -convergent reduction $T: t_0 \twoheadrightarrow_{\mathcal{B}} t_\alpha$ in \mathcal{B} by removing reduction steps in S that take place at or below outermost-volatile positions of some prefix of S and replace them by \rightarrow_{\perp} -steps.

Let π be an outermost-volatile position of some prefix $S|_\lambda$. Then there is some ordinal $\beta < \lambda$ such that no reduction step between β and λ in S takes place strictly above π , i.e. $\pi_\iota \not\prec \pi$ for all $\beta \leq \iota < \lambda$. Such an ordinal β must exist since otherwise π would not be an outermost-volatile position in $S|_\lambda$. Hence, we can construct a destructive reduction $S': t_\beta|_\pi \twoheadrightarrow \perp$ by taking the subsequence of the segment $S|_{[\beta, \lambda]}$ that contains the reduction steps at π or below. Note that $t_\beta|_\pi$ might still contain the symbol \perp . Since \perp is not relevant for the applicability of rules in \mathcal{R} , each of the \perp symbols in $t_\beta|_\pi$ can be safely replaced by arbitrary total terms, in particular by terms in \mathcal{U} . Let r be a term that is obtained in this way. Then there is a destructive reduction $S'': r \twoheadrightarrow \perp$ that applies the same rules at the same positions as in S' . Hence, $r \in \mathcal{U}$. By construction, r is a \perp, \mathcal{U} -instance of $t_\beta|_\pi$ which means that $t_\beta|_\pi \in \mathcal{U}_\perp$. Additionally, $t_\beta|_\pi \neq \perp$ since there is a non-empty reduction

$S': t_\beta|_\pi \twoheadrightarrow \perp$ starting in $t_\beta|_\pi$. Consequently, there is a rule $t_\beta|_\pi \rightarrow \perp$ in \mathcal{B} . Let T' be the reduction that is obtained from $S|_\lambda$ by replacing the β -th step, which we can assume w.l.o.g. to take place at π , by a step with the rule $t_\beta|_\pi \rightarrow \perp$ at the same position π and removing all reduction steps φ_ι taking place at π or below for all $\beta < \iota < \lambda$. Let t' be the term that the reduction T' p -converges to. t_λ and t' can only differ at position π or below. However, by construction, we have $t'(\pi) = \perp$ and, by Lemma 5.4, $t_\lambda(\pi) = \perp$. Consequently, $t' = t_\lambda$.

This construction can be performed for all prefixes of S and their respective outermost-volatile positions. Thereby, we obtain a p -converging reduction $T: t_0 \twoheadrightarrow_{\mathcal{B}} t_\alpha$ for which no prefix has a volatile position. By Lemma 5.5, T is a total reduction. Note that \mathcal{B} is a TRS over the extended signature $\Sigma' = \Sigma \uplus \{\perp\}$, i.e. terms containing \perp are considered total. Hence, by Theorem 5.6, $T: t_0 \twoheadrightarrow_{\mathcal{B}} t_\alpha$. \blacksquare

This already provides one direction of the equivalence we want to establish. Before we make the next step, we need the following lemma shown by Kennaway et al. [Ken99]:

Lemma 7.4 (postponement of \rightarrow_{\perp} -steps). *Let \mathcal{R} be a left-linear, left-finite TRS and \mathcal{B} some Böhm extension of \mathcal{R} . Then $s \twoheadrightarrow_{\mathcal{B}} t$ implies $s \twoheadrightarrow_{\mathcal{R}} s' \twoheadrightarrow_{\perp} t$ for some term s' .²*

In the next proposition we show that, excluding \perp subterms, the final term of a p -converging reduction can be approximated arbitrarily well by a finite reduction. This corresponds to Corollary 3.4 which establishes finite approximations for m -convergent reductions.

Proposition 7.5 (finite approximation). *Let \mathcal{R} be a left-linear, left-finite TRS and $s \twoheadrightarrow t$. Then, for each finite set $P \subseteq \mathcal{P}_{\perp}(t)$, there is a reduction $s \rightarrow^* t'$ such that t and t' coincide in P .*

Proof. Assume that $s \twoheadrightarrow_{\mathcal{R}} t$. Then, by Proposition 7.3, there is a reduction $s \twoheadrightarrow_{\mathcal{B}} t$, where \mathcal{B} is the Böhm extension of \mathcal{R} w.r.t. the set of total, fragile terms of \mathcal{R} . By Lemma 7.4, there is a reduction $s \twoheadrightarrow_{\mathcal{R}} s' \twoheadrightarrow_{\perp} t$. Clearly, s' and t coincide in $\mathcal{P}_{\perp}(t)$. Let $d = \max\{|\pi| \mid \pi \in P\}$. Since P is finite, d is well-defined. By Corollary 3.4, there is a reduction $s \rightarrow^*_{\mathcal{R}} t'$ such that t' and s' coincide up to depth d and, thus, in particular they coincide in P . Consequently, since s' and t coincide in $\mathcal{P}_{\perp}(t) \supseteq P$, t and t' coincide in P , too. \blacksquare

In order to establish a compression lemma for destructive reductions we need that fragile terms are preserved by finite reductions. We can obtain this from the following more general lemma showing that destructive reductions are preserved by forming projections as constructed in the Infinitary Strip Lemma:

Lemma 7.6 (preservation of destructive reductions by projections). *Let \mathcal{R} be an orthogonal TRS, $S: t_0 \twoheadrightarrow t_\alpha$ a destructive reduction, and $T: t_0 \twoheadrightarrow_U s_0$ a complete development of a set U of disjoint redex occurrences. Then the projection $S/T: s_0 \twoheadrightarrow s_\alpha$ is also destructive.*

Proof. We consider the situation depicted in Figure 1. Since $S: t_0 \twoheadrightarrow t_\alpha$ is destructive, we have, for each $\beta < \alpha$, some $\beta \leq \gamma < \alpha$ such that $v_\gamma = \varepsilon$. If $v_\gamma = \varepsilon$, then also $\varepsilon \in v_\gamma // U_\gamma$ unless $\varepsilon \in U_\gamma$. As by Proposition 6.2, U_γ is a set of pairwise disjoint positions, $\varepsilon \in U_\gamma$ implies $U_\gamma = \{\varepsilon\}$. This means that if $v_\gamma = \varepsilon$ and $\varepsilon \in U_\gamma$, then $U_\iota = \emptyset$ for all $\gamma < \iota < \alpha$. Thus, there is only at most one $\gamma < \alpha$ with $\varepsilon \in U_\gamma$. Therefore, we have, for each $\beta < \alpha$, some $\beta \leq \gamma < \alpha$ such that $\varepsilon \in v_\gamma // U_\gamma$. Hence, T is destructive. \blacksquare

²Strictly speaking, if s is not a total term, i.e. it contains \perp , then we have to consider the system that is obtained from \mathcal{R} by extending its signature to Σ_{\perp} .

As a consequence of this preservation of destructiveness by forming projections, we obtain that the set of fragile terms is closed under finite reductions:

Lemma 7.7 (closure of fragile terms under finite reductions). *In each orthogonal TRS, the set of fragile terms is closed under finite reductions.*

Proof. Let t be a fragile term and $T: t \rightarrow^* t'$ a finite reduction. Hence, there is a destructive reduction starting in t . A straightforward induction proof on the length of T , using Lemma 7.6, shows that there is a destructive reduction starting in t' . Thus, t' is fragile. ■

Now we can show that destructiveness does not need more than ω steps in orthogonal, left-finite TRSs. This property will be useful for proving the equivalence of root-activeness and fragility of total terms as well the Compression Lemma for p -convergent reductions.

Proposition 7.8 (Compression Lemma for destructive reductions). *Let \mathcal{R} be an orthogonal, left-finite TRS and t a partial term in \mathcal{R} . If there is a destructive reduction starting in t , then there is a destructive reduction of length ω starting in t .*

Proof. Let $S: t_0 \xrightarrow{\lambda} \perp$ be a destructive reduction starting in t_0 . Hence, there is some $\alpha < \lambda$ such that $S|_\alpha: t_0 \xrightarrow{\lambda} s_1$, where s_1 is a ρ -redex for some $\rho: l \rightarrow r \in R$. Let P be the set of pattern positions of the ρ -redex s_1 , i.e. $P = \mathcal{P}_\Sigma(l)$. Due to the left-finiteness of \mathcal{R} , P is finite. Hence, by Proposition 7.5, there is a finite reduction $t_0 \rightarrow^* s'_1$ such that s_1 and s'_1 coincide in P . Hence, because \mathcal{R} is left-linear, also s'_1 is a ρ -redex. Now consider the reduction $T_0: t_0 \rightarrow^* s'_1 \rightarrow_{\rho, \varepsilon} t_1$ ending with a contraction at the root. T_0 is of finite length and, according to Lemma 7.7, t_1 is fragile.

Since t_1 is again fragile, the above argument can be iterated arbitrarily often which yields for each $i < \omega$ a finite reduction $T_i: t_i \rightarrow^* t_{i+1}$ whose last step is a contraction at the root. Then the concatenation $T = \prod_{i < \omega} T_i$ of these reductions is a destructive reduction of length ω starting in t_0 . ■

The above proposition bridges the gap between fragility and root-activeness. Whereas the former concept is defined in terms of transfinite reductions, the latter is defined in terms of finite reductions. By Proposition 7.8, however, a fragile term is always finitely reducible to a redex. This is the key to the observation that fragility is not only quite similar to root-activeness but is, in fact, essentially the same concept.

Proposition 7.9 (root-activeness = fragility). *Let \mathcal{R} be an orthogonal, left-finite TRS and t a total term in \mathcal{R} . Then t is root-active iff t is fragile.*

Proof. The “only if” direction is easy: If t is root-active, then there is a reduction S of length ω starting in t with infinitely many steps taking place at the root. Hence, $S: t \xrightarrow{\omega} \perp$ is a destructive reduction, which makes t a fragile term.

For the converse direction we assume that t is fragile and show that, for each reduction $t \rightarrow^* s$, there is a reduction $s \rightarrow^* t'$ to a redex t' . By Lemma 7.7, also s is fragile. Hence, there is a destructive reduction $S: s \xrightarrow{\omega} \perp$ starting in s . According to Proposition 7.8, we can assume that S has length ω . Therefore, there is some $n < \omega$ such that $S|_n: s \rightarrow^* t'$ for a redex t' . ■

Before we prove the missing direction of the equality of p -reachability and Böhm-reachability we need the property that m -convergent reductions consisting only of \rightarrow_\perp -steps can be compressed to length at most ω as well:

Lemma 7.10 (compression of \rightarrow_{\perp} -steps). *Consider the Böhm extension of an orthogonal TRS w.r.t. its root-active terms and $S: s \xrightarrow{m}_{\perp} t$ with $s \in \mathcal{T}^{\infty}(\Sigma, \mathcal{V})$, $t \in \mathcal{T}^{\infty}(\Sigma_{\perp}, \mathcal{V})$. Then there is an m -converging reduction $T: s \xrightarrow{m}_{\perp} t$ of length at most ω that is a complete development of a set of disjoint occurrences of root-active terms in s .*

Proof. The proof is essentially the same as that of Lemma 7.2.4 from Ketema [Ket06]. ■

The important part of the above lemma is the statement that only terms in \mathcal{RA} are contracted instead of the general case where a \rightarrow_{\perp} -step contracts a term in $\mathcal{RA}_{\perp} \supset \mathcal{RA}$.

Finally, we have gathered all tools necessary in order to prove the converse direction of the equivalence of p -reachability and Böhm-reachability w.r.t. root-active terms.

Theorem 7.11 (p -reachability = Böhm-reachability w.r.t. \mathcal{RA}). *Let \mathcal{R} be an orthogonal, left-finite TRS and \mathcal{B} the Böhm extension of \mathcal{R} w.r.t. its root-active terms. Then $s \xrightarrow{p}_{\mathcal{R}} t$ iff $s \xrightarrow{m}_{\mathcal{B}} t$.*

Proof. The “only if” direction follows immediately from Proposition 7.9 and Proposition 7.3.

Now consider the converse direction: Let $s \xrightarrow{m}_{\mathcal{B}} t$ be an m -convergent reduction in \mathcal{B} . W.l.o.g. we assume s to be total. Due to Lemma 7.4, there is a term $s' \in \mathcal{T}^{\infty}(\Sigma, \mathcal{V})$ such that there are m -convergent reductions $S: s \xrightarrow{m}_{\mathcal{R}} s'$ and $T: s' \xrightarrow{m}_{\perp} t$. By Lemma 7.10, we can assume that in $s' \xrightarrow{m}_{\perp} t$ only pairwise disjoint occurrences of root-active terms are contracted. By Proposition 7.9, each root-active term $r \in \mathcal{RA}$ is fragile, i.e. we have a destructive reduction $r \xrightarrow{p}_{\mathcal{R}} \perp$ starting in r . Thus, we can construct a p -converging reduction $T': s' \xrightarrow{p}_{\mathcal{R}} t$ by replacing each step $C[r] \rightarrow_{\perp} C[\perp]$ in T with the corresponding reduction $C[r] \xrightarrow{p}_{\mathcal{R}} C[\perp]$. By combining T' with the m -converging reduction S , which, according to Theorem 5.6, is also p -converging, we obtain the p -converging reduction $S \cdot T': s \xrightarrow{p}_{\mathcal{R}} t$. ■

With this equivalence, p -convergent reductions inherit a number of important properties that are enjoyed by Böhm-convergent reductions:

Theorem 7.12 (infinitary confluence). *Every orthogonal, left-finite TRS is infinitarily confluent. That is, for each orthogonal, left-finite TRS, $s_1 \xleftarrow{p} t \xrightarrow{p} s_2$ implies $s_1 \xrightarrow{p} t' \xleftarrow{p} s_2$.*

Proof. Leveraging Theorem 7.11, this theorem follows from Theorem 4.3. ■

Returning to Example 3.5 again, we can see that the terms g^{ω} and f^{ω} can now be joined by repeatedly contracting the redex at the root which yields destructive reductions $g^{\omega} \xrightarrow{p}_{\perp}$ and $f^{\omega} \xrightarrow{p}_{\perp}$, respectively.

Theorem 7.13 (infinitary normalisation). *Every orthogonal, left-finite TRS is infinitarily normalising. That is, for each orthogonal, left-finite TRS \mathcal{R} and a partial term t in \mathcal{R} , there is an \mathcal{R} -normal form p -reachable from t .*

Proof. This follows immediately from Theorem 7.11 and Theorem 4.4. ■

Combining Theorem 7.12 and Theorem 7.13, we obtain that each term in an orthogonal TRS has a unique normal form w.r.t. p -convergence. Due to Theorem 7.11, this unique normal form is the Böhm tree w.r.t. root-active terms.

Since p -converging reductions in orthogonal TRS can always be transformed such that they consist of a prefix which is an m -convergent reduction and a suffix consisting of nested destructive reductions, we can employ the Compression Lemma for m -convergent reductions (Theorem 3.3) and the Compression Lemma for destructive reductions (Proposition 7.8) to obtain the Compression Lemma for p -convergent reductions:

Theorem 7.14 (Compression Lemma for p -convergent reductions). *For each orthogonal, left-finite TRS, $s \mathcal{P}_{\mathcal{R}} t$ implies $s \mathcal{P}_{\leq \omega} t$.*

Proof. Let $s \mathcal{P}_{\mathcal{R}} t$. According to Theorem 7.11, we have $s \mathcal{M}_{\mathcal{B}} t$ for the Böhm extension \mathcal{B} of \mathcal{R} w.r.t. \mathcal{RA} and, therefore, by Lemma 7.4, we have reductions $S: s \mathcal{M}_{\mathcal{R}} s'$ and $T: s' \mathcal{M}_{\perp} t$. Due to Theorem 3.3, we can assume S to be of length at most ω and, due to Theorem 5.6, to be p -convergent, i.e. $S: s \mathcal{P}_{\mathcal{R}}^{\leq \omega} s'$. If T is the empty reduction, then we are done. If not, then T is a complete development of pairwise disjoint occurrences of root-active terms according to Lemma 7.10. Hence, each step is of the form $C[r] \rightarrow_{\perp} C[\perp]$ for some root-active term r . By Proposition 7.9, for each such term r , there is a destructive reduction $r \mathcal{P}_{\mathcal{R}} \perp$ which we can assume, in accordance with Proposition 7.8, to be of length ω . Hence, each step $C[r] \rightarrow_{\perp} C[\perp]$ can be replaced by the reduction $C[r] \mathcal{P}_{\mathcal{R}}^{\omega} C[\perp]$. Concatenating these reductions results in a reduction $T': s' \mathcal{P}_{\mathcal{R}} t$ of length at most $\omega \cdot \omega$. If $S: s \mathcal{P}_{\mathcal{R}}^{\leq \omega} s'$ is of finite length, we can interleave the reduction steps in T' such that we obtain a reduction $T'': s' \mathcal{P}_{\mathcal{R}}^{\omega} t$ of length ω . Then we have $S \cdot T'': s \mathcal{P}_{\mathcal{R}}^{\omega} t$. If $S: s \mathcal{P}_{\mathcal{R}}^{\leq \omega} s'$ has length ω , we construct a reduction $s \mathcal{P}_{\mathcal{R}} t$ as follows: As illustrated above, T' consists of destructive reductions taking place at some pairwise disjoint positions. These steps can be interleaved into the reduction S resulting into a reduction $s \mathcal{P}_{\mathcal{R}} t$ of length ω . The argument for that is similar to that employed in the successor case of the induction proof of the Compression Lemma of Kennaway et al. [Ken95]. ■

We can use the Compression Lemma for p -convergent reductions to obtain a stronger variant of Theorem 5.6 for orthogonal TRSs:

Corollary 7.15 (m -reachability = p -reachability of total terms). *Let \mathcal{R} be an orthogonal, left-finite TRS and $s, t \in \mathcal{T}^{\infty}(\Sigma, \mathcal{V})$. Then $s \mathcal{M} t$ iff $s \mathcal{P} t$.*

Proof. The “only if” direction follows immediately from Theorem 5.6. For the “if” direction assume a reduction $S: s \mathcal{P} t$. According to Theorem 7.14, there is a reduction $T: s \mathcal{P}_{\leq \omega} t$. Hence, since s is total and totality is preserved by single reduction steps, $T: s \mathcal{P}_{\leq \omega} t$ is total. Applying Theorem 5.6, yields that $T: s \mathcal{M}_{\leq \omega} t$. ■

8. Conclusions

Infinitary term rewriting in the partial order model provides a more fine-grained notion of convergence. Formally, every meaningful, i.e. p -continuous, reduction is also p -converging. Practically, p -converging reductions can end in a term containing \perp 's indicating positions of “local divergence”. Theorem 5.6 and Corollary 7.15 indicate that the partial model coincides with the metric model but additionally allows a more detailed inspection of non- m -converging reductions. Instead of the coarse discrimination between convergence and divergence provided by the metric model, the partial order model allows different levels between full convergence (a total term as result) and full divergence (\perp as result). Moreover, due to the equivalence to Böhm-reachability, we additionally obtain infinitary normalisation and infinitary confluence for orthogonal systems, which we do not have in the metric model, while still maintaining the compression property. While achieving the same goals as Böhm-extensions, the partial order approach provides an intuitive and more elegant model.

We have only studied strong convergence in this paper. It would be interesting to find out whether the shift to the partial order model has similar benefits for weak convergence, which is known to be rather unruly [Sim04].

Another interesting direction to follow is the ability to finitely simulate transfinite reductions by term graph rewriting. For m -convergence this is possible, at least to some extent [Ken94]. However, we think that a different approach to term graph rewriting, viz. the *double-pushout approach* [Ehr73] or the *equational approach* [Ari96], is more appropriate for p -convergence [Cor97, Bah09].

Acknowledgements

I want to thank Bernhard Gramlich for his constant support during the work on my master's thesis which made this work possible. I am also grateful for the valuable comments of the anonymous referees.

References

- [Ari96] Zena M. Ariola and Jan Willem Klop. Equational term graph rewriting. *Fundam. Inf.*, 26(3-4):207–240, 1996.
- [Arn80] André Arnold and Maurice Nivat. The metric space of infinite trees. Algebraic and topological properties. *Fundam. Inf.*, 3(4):445–476, 1980.
- [Bah09] Patrick Bahr. *Infinitary Rewriting - Theory and Applications*. Master's thesis, Vienna University of Technology, Vienna, 2009.
URL <http://www.pa-ba.info/?q=pub/master>
- [Blo04] Stefan Blom. An approximation based approach to infinitary lambda calculi. In Vincent van Oostrom (ed.), *RTA '04, Lecture Notes in Computer Science*, vol. 3091, pp. 221–232. Springer Berlin / Heidelberg, 2004. doi:10.1007/b98160.
URL <http://www.springerlink.com/content/4n3gqw43d1bpnldy/>
- [Cor93] Andrea Corradini. Term rewriting in CT_{Σ} . In Marie-Claude Gaudel and Jean-Pierre Jouannaud (eds.), *TAPSOFT '93, Lecture Notes in Computer Science*, vol. 668, pp. 468–484. Springer Berlin / Heidelberg, 1993. doi:10.1007/3-540-56610-4_83.
URL <http://www.springerlink.com/content/f73r5p2v370220m4/>
- [Cor97] Andrea Corradini and Frank Drewes. (Cyclic) term graph rewriting is adequate for rational parallel term rewriting. Tech. Rep. TR-14-97, Università di Pisa, Dipartimento di Informatica, 1997.
- [Der91] Nachum Dershowitz, Stéphane Kaplan, and David A. Plaisted. Rewrite, rewrite, rewrite, rewrite, ... *Theor. Comput. Sci.*, 83(1):71–96, 1991. doi:DOI:10.1016/0304-3975(91)90040-9.
URL <http://www.sciencedirect.com/science/article/B6V1G-45DHJRB-H/2/767b35171dafdfa511dd0463ea25dbdd>
- [Ehr73] Hartmut Ehrig, Michael Pfender, and Hans Jürgen Schneider. Graph-grammars: An algebraic approach. In *SWAT '73*, pp. 167–180. IEEE Computer Society, Washington, DC, USA, 1973. doi: <http://dx.doi.org/10.1109/SWAT.1973.11>.
- [Kah93] Gilles Kahn and Gordon D. Plotkin. Concrete domains. *Theor. Comput. Sci.*, 121(1-2):187–277, 1993. doi:DOI:10.1016/0304-3975(93)90090-G.
URL <http://www.sciencedirect.com/science/article/B6V1G-45FC431-2K/2/6c30777ef97aea14c529418b4d5c5d4a>
- [Kel55] John L. Kelley. *General Topology, Graduate Texts in Mathematics*, vol. 27. Springer-Verlag, 1955.
- [Ken94] Richard Kennaway, Jan Willem Klop, M. Ronan Sleep, and Fer-Jan de Vries. On the adequacy of graph rewriting for simulating term rewriting. *ACM Trans. Program. Lang. Syst.*, 16(3):493–523, 1994. doi:<http://doi.acm.org/10.1145/177492.177577>.
- [Ken95] Richard Kennaway, Jan Willem Klop, M. Ronan Sleep, and Fer-Jan de Vries. Transfinite reductions in orthogonal term rewriting systems. *Inform. and Comput.*, 119(1):18–38, 1995. doi:DOI:10.1006/inco.1995.1075.
URL <http://www.sciencedirect.com/science/article/B6WGK-45NJYB-4W/2/7d48d04a2fe97d6e9e1fc5179f31a488>
- [Ken99] Richard Kennaway, Vincent van Oostrom, and Fer-Jan de Vries. Meaningless terms in rewriting. *J. Funct. Logic Programming*, 1999(1):1–35, 1999.

- [Ken03] Richard Kennaway and Fer-Jan de Vries. Infinitary rewriting. In Terese [Ter03], chap. 12, pp. 668–711.
URL <http://amazon.com/o/ASIN/0521391156/>
- [Ket06] Jeroen Ketema. *Böhm-Like Trees for Rewriting*. Ph.D. thesis, Vrije Universiteit Amsterdam, 2006.
URL <http://dare.uvu.vu.nl/handle/1871/9203>
- [Rod98] Pieter Hendrik Rodenburg. Termination and confluence in infinitary term rewriting. *J. Symbolic Logic*, 63(4):1286–1296, 1998.
URL <http://www.jstor.org/stable/2586651>
- [Sim04] Jakob Grue Simonsen. On confluence and residuals in Cauchy convergent transfinite rewriting. *Inf. Process. Lett.*, 91(3):141–146, 2004. doi:DOI:10.1016/j.ipl.2004.03.018.
URL <http://www.sciencedirect.com/science/article/B6V0F-4CBVNG4-1/2/d5d0f374f89fd62e07d023512a5b3dfe>
- [Ter03] Terese. *Term Rewriting Systems*. Cambridge University Press, 1st edn., 2003.
URL <http://amazon.com/o/ASIN/0521391156/>

