
PARTIAL ORDER INFINITARY TERM REWRITING *

PATRICK BAHR

Department of Computer Science, University of Copenhagen
Universitetsparken 5, 2100 Copenhagen, Denmark
URL: <http://www.diku.dk/~paba>
e-mail address: paba@di.ku.dk

ABSTRACT. We study an alternative model of infinitary term rewriting. Instead of a metric on terms, a partial order on partial terms is employed to formalise convergence of reductions. We consider both a weak and a strong notion of convergence and show that the metric model of convergence coincides with the partial order model restricted to total terms. Hence, partial order convergence constitutes a conservative extension of metric convergence, which additionally offers a fine-grained distinction between different levels of divergence.

In the second part, we focus our investigation on strong convergence of orthogonal systems. The main result is that the gap between the metric model and the partial order model can be bridged by extending the term rewriting system by additional rules. These extensions are the well-known Böhm extensions. Based on this result, we are able to establish that – contrary to the metric setting – orthogonal systems are both infinitarily confluent and infinitarily normalising in the partial order setting. The unique infinitary normal forms that the partial order model admits are Böhm trees.

INTRODUCTION

Infinitary term rewriting [14] extends the theory of term rewriting by giving a meaning to transfinite rewriting sequences. Its formalisation [9] is chiefly based on the metric space of terms as studied by Arnold and Nivat [2]. Other models for transfinite reductions, using for example general topological spaces [22] or partial orders [7, 6], were mainly considered to pursue quite specific purposes and have not seen nearly as much attention as the metric model. In this paper we introduce a novel foundation of infinitary term rewriting based on the partially ordered set of partial terms [12]. We show that this model of infinitary term rewriting is superior to the metric model. This assessment is supported by two findings: First, the partial order model of infinitary term rewriting conservatively extends the metric model. That is, anything that can be done in the metric model can be achieved in the partial order model as well by simply restricting it to the set of total terms. Secondly, unlike the metric model, the partial order model provides a fine-grained distinction between different

2012 ACM CCS: [Theory of computation]: Formal languages and automata theory—Formalisms—Rewrite systems.

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

* Parts of this paper have appeared in the proceedings of RTA 2010 [5].

levels of divergence and exhibits nice properties like infinitary confluence and normalisation of orthogonal systems.

The defining core of a theory of infinitary term rewriting is its notion of convergence for transfinite reductions: which transfinite reductions are “admissible” and what is their final outcome. In this paper we study both variants of convergence that are usually considered in the established theory of metric infinitary term rewriting: weak convergence [9] and strong convergence [16]. For both variants we introduce a corresponding notion of convergence based on the partially ordered set of partial terms.

The first part of this paper is concerned with comparing the metric model and the partial order model both in their respective weak and strong variants. In both cases, the partial order approach constitutes a conservative extension of the metric approach: a reduction in the metric model is converging iff it is converging in the partial order model and only contains total terms.

In the second part we focus on strong convergence in orthogonal systems. To this end we reconsider the theory of meaningless terms of Kennaway et al. [17]. In particular, we consider Böhm extensions. The Böhm extension of a term rewriting system adds rewrite rules which admit contracting meaningless terms to \perp . The central result of the second part of this paper is that the additional rules in Böhm extensions close the gap between partial order convergence and metric convergence. More precisely, we show that reachability w.r.t. partial order convergence in a term rewriting system coincides with reachability w.r.t. metric convergence in the corresponding Böhm extension.

From this result we can easily derive a number of properties for strong partial order convergence in orthogonal systems:

- Infinitary confluence,
- infinitary normalisation, and
- compression, i.e. each reduction can be compressed to length at most ω

The first two properties exhibit another improvement over the metric model which does not have either of these. Moreover, it means that each term has a unique infinitary normal form – its Böhm tree.

The most important tool for establishing these results is provided by a notion of complete developments that we have transferred from the metric approach to infinitary rewriting [16]. We show, that the final outcome of a complete development is unique and that, in contrast to the metric model, the partial order model admits complete developments for any set of redex occurrences. To this end, we use a technique similar to paths and finite jumps known from metric infinitary term rewriting [14, 21].

Outline. After providing the basic preliminaries for this paper in Section 1, we will briefly recapitulate the metric model of infinitary term rewriting including meaningless terms and Böhm extensions in Section 2. In Section 3, we introduce our novel approach to infinitary term rewriting based on the partial order on terms. In Section 4, we compare both models and establish that the partial order model provides a conservative extension of the metric model. In the remaining part of this paper, we focus on the strong notion of convergence. In Section 5, we establish a theory of complete developments in the setting of partial order convergence. This is then used in Section 6 to prove the equality of reachability w.r.t. partial order convergence and reachability w.r.t. metric convergence in the Böhm extension. Finally, we evaluate our results and point to interesting open questions in Section 7.

1. PRELIMINARIES

We assume the reader to be familiar with the basic theory of ordinal numbers, orders and topological spaces [13], as well as term rewriting [24]. In the following, we briefly recall the most important notions.

1.1. 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*. For a finite sequence $(a_i)_{i < n}$ or a sequence $(a_i)_{i < \omega}$ of length ω , we also use the notation $\langle a_0, a_1, \dots, a_{n-1} \rangle$ respectively $\langle a_0, a_1, \dots \rangle$. In particular, $\langle \rangle$ denotes an empty sequence.

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 binary relation \leq forms a complete semilattice (see Section 1.3 below). Similarly, a sequence S is a (proper) *suffix* of a sequence T if there is a (non-empty) sequence S' with $S' \cdot S = T$.

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 $S|_{[0, \alpha]} = S|_\alpha$ for each sequence S and ordinal $\alpha \leq |S|$.

1.2. Metric Spaces. A pair (M, \mathbf{d}) is called a *metric space* if $\mathbf{d}: M \times M \rightarrow \mathbb{R}_0^+$ is a function satisfying $\mathbf{d}(x, y) = 0$ iff $x = y$ (identity), $\mathbf{d}(x, y) = \mathbf{d}(y, x)$ (symmetry), and $\mathbf{d}(x, z) \leq \mathbf{d}(x, y) + \mathbf{d}(y, z)$ (triangle inequality), for all $x, y, z \in M$. If \mathbf{d} instead of the triangle inequality, satisfies the stronger property $\mathbf{d}(x, z) \leq \max\{\mathbf{d}(x, y), \mathbf{d}(y, z)\}$ (strong triangle), then (M, \mathbf{d}) is called an *ultrametric space*. Let $(a_\iota)_{\iota < \alpha}$ be a sequence in a metric space (M, \mathbf{d}) . The sequence $(a_\iota)_{\iota < \alpha}$ *converges* to an element $a \in M$, written $\lim_{\iota \rightarrow \alpha} a_\iota$, if, for each $\varepsilon \in \mathbb{R}^+$, there is a $\beta < \alpha$ such that $\mathbf{d}(a, a_\iota) < \varepsilon$ for every $\beta < \iota < \alpha$; $(a_\iota)_{\iota < \alpha}$ is *continuous* if $\lim_{\iota \rightarrow \lambda} a_\iota = a_\lambda$ for each limit ordinal $\lambda < \alpha$. The sequence $(a_\iota)_{\iota < \alpha}$ is called *Cauchy* if, for any $\varepsilon \in \mathbb{R}^+$, there is a $\beta < \alpha$ such that, for all $\beta < \iota < \iota' < \alpha$, we have that $\mathbf{d}(m_\iota, m_{\iota'}) < \varepsilon$. A metric space is called *complete* if each of its non-empty Cauchy sequences converges.

1.3. Partial Orders. A *partial order* \leq on a set A is a binary relation on A that is *transitive*, *reflexive*, and *antisymmetric*. The pair (A, \leq) is then called a *partially ordered set*. We use $<$ to denote the strict part of \leq , i.e. $a < b$ iff $a \leq b$ and $b \not\leq a$. A sequence $(a_\iota)_{\iota < \alpha}$ in (A, \leq) is called a (*strict*) *chain* if $a_\iota \leq a_\gamma$ (resp. $a_\iota < a_\gamma$) for all $\iota < \gamma < \alpha$. A subset D of the underlying set A is called *directed* if it is non-empty and each pair of elements in D has an upper bound in D . A partially ordered set (A, \leq) is called a *complete semilattice* if it has a *least element*, every *directed subset* D of A has a *least upper bound* (*lub*) $\bigsqcup D$, 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) $\prod B$ 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 $\liminf_{\iota \rightarrow \alpha} a_\iota = \prod_{\beta < \alpha} \left(\prod_{\beta \leq \iota < \alpha} a_\iota \right)$, exists.

It is easy to see that the limit inferior of closed sequences is simply the last element of the sequence. This is, however, only a special case of the following more general proposition:

Proposition 1.1 (invariance of the limit inferior). *Let $(a_\iota)_{\iota < \alpha}$ be a sequence in a complete semilattice and $(b_\iota)_{\iota < \beta}$ a non-empty suffix of $(a_\iota)_{\iota < \alpha}$. Then $\liminf_{\iota \rightarrow \alpha} a_\iota = \liminf_{\iota \rightarrow \beta} b_\iota$.*

Proof. Let $a = \liminf_{\iota \rightarrow \alpha} a_\iota$ and $b = \liminf_{\iota \rightarrow \beta} b_\iota$. Since $(b_\iota)_{\iota < \beta}$ is a suffix of $(a_\iota)_{\iota < \alpha}$, there is some $\delta < \alpha$ such that $b_\iota = a_{\delta + \iota}$ for all $\iota < \beta$. Hence, we know that $a = \prod_{\gamma < \alpha} \prod_{\gamma \leq \iota < \alpha} a_\iota$ and $b = \prod_{\delta \leq \gamma < \alpha} \prod_{\gamma \leq \iota < \alpha} a_\iota$. Let $c_\gamma = \prod_{\gamma \leq \iota < \alpha} a_\iota$ for each $\gamma < \alpha$, $A = \{c_\gamma \mid \gamma < \alpha\}$ and $B = \{c_\gamma \mid \delta \leq \gamma < \alpha\}$. Note that $a = \prod A$ and $b = \prod B$. Because $B \subseteq A$, we have that $b \leq a$. On the other hand, since $c_\gamma \leq c_{\gamma'}$ for $\gamma \leq \gamma'$, we find, for each $c_\gamma \in A$, some $c_{\gamma'} \in B$ with $c_\gamma \leq c_{\gamma'}$. Hence, $a \leq b$. Therefore, due to the antisymmetry of \leq , we can conclude that $a = b$. \square

Note that the limit in a metric space has the same behaviour as the one for the limit inferior described by the proposition above. However, one has to keep in mind that – unlike the limit – the limit inferior is not invariant under taking cofinal subsequences!

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 $\langle \rangle$ if $\alpha = 0$, $(\prod_{\iota < \alpha'} S_\iota) \cdot S_{\alpha'}$ if $\alpha = \alpha' + 1$, and $\prod_{\gamma < \alpha} \prod_{\iota < \gamma} S_\iota$ if α is a limit ordinal. For instance, the concatenation $\prod_{i < \omega} \langle i, i + 1 \rangle$ yields the sequence $\langle 0, 1, 1, 2, 2, \dots \rangle$ of length ω , and the concatenation $\prod_{\iota < \alpha} \langle \iota \rangle$, for any ordinal α , yields the sequence $(\iota)_{\iota < \alpha}$.

1.4. Terms. Unlike in the traditional – i.e. finitary – 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. A *signature* Σ is a countable set of symbols. Each symbol f is associated with its arity $\text{ar}(f) \in \mathbb{N}$, and we write $\Sigma^{(n)}$ for the set of symbols in Σ which have arity n . 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 $t = f(t_0, \dots, t_{k-1})$, where $f \in \Sigma^{(k)}$, and $t_0, \dots, t_{k-1} \in T$. A symbol $c \in \Sigma^{(0)}$ of arity 0 is also called a *constant symbol*, and we use the shorthand c to denote a term $c()$. We consider $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ as a superset of the set $\mathcal{T}(\Sigma, \mathcal{V})$ of *finite terms*.

For each term $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$, we define the *set of positions* in t , denoted $\mathcal{P}(t)$, as the smallest set of finite sequences in \mathbb{N} such that $\langle \rangle \in \mathcal{P}(t)$, and $\langle i \rangle \cdot \pi \in \mathcal{P}(t)$ whenever $t = f(t_0, \dots, t_{k-1})$, $i < k$, and $\pi \in \mathcal{P}(t_i)$. Given a position $\pi \in \mathcal{P}(t)$, we define the *subterm* of t at π , denoted $t|_\pi$, by recursion on π as follows: $t|_{\langle \rangle} = t$, and $f(t_0, \dots, t_{k-1})|_{\langle i \rangle \cdot \pi} = t_i|_\pi$. Moreover, we write $t(\pi)$ for the symbol in t at π , i.e. $t(\pi) = f$ if $t|_\pi = f(t_0, \dots, t_{k-1})$ and $t(\pi) = v$ if $t|_\pi = v \in \mathcal{V}$. For terms $s, t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ and a position $\pi \in \mathcal{P}(t)$, we write $t[s]_\pi$ for the term t with the subterm at π replaced by s , i.e.

$$t[s]_{\langle \rangle} = s, \quad \text{and} \quad f(t_0, \dots, t_{k-1})[s]_{\langle i \rangle \cdot \pi} = f(t_0, \dots, t_{i-1}, t_i[s]_\pi, t_{i+1}, \dots, t_{k-1}).$$

Note that while the set of terms $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ is defined coinductively, the set of positions of a term is defined inductively. Consequently, the subterm at a position and substitution at a position are defined by recursion.

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 $\pi_2 \leq \pi_1$.

A *context* is a “term with holes”, which are represented by a distinguished variable \square . We write $C[\dots]$ for a context with at least one occurrence of \square , and $C\langle \dots \rangle$ for a context with zero more occurrences of \square . $C[t_1, \dots, t_n]$ denotes the result of replacing the occurrences of \square in C (from left to right) by t_1, \dots, t_n . $C\langle t_1, \dots, t_n \rangle$ is defined accordingly.

A *substitution* σ is a mapping from \mathcal{V} to $\mathcal{T}^\infty(\Sigma, \mathcal{V})$. Its *domain*, denoted $\text{dom}(\sigma)$, is the set $\{x \in \mathcal{V} \mid \sigma(x) \neq x\}$ of variables not mapped to itself by σ . Substitutions are uniquely extended to functions from $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ to $\mathcal{T}^\infty(\Sigma, \mathcal{V})$: $\sigma(f(t_1, \dots, t_n)) = f(\sigma(t_1), \dots, \sigma(t_n))$ for $f \in \Sigma^{(n)}$ and $t_1, \dots, t_n \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$. Instead of $\sigma(s)$, we shall also write $s\sigma$.

On $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ a similarity measure $\text{sim}(\cdot, \cdot) \in \mathbb{N} \cup \{\infty\}$ can be defined by setting

$$\text{sim}(s, t) = \min \{|\pi| \mid \pi \in \mathcal{P}(s) \cap \mathcal{P}(t), s(\pi) \neq t(\pi)\} \cup \{\infty\} \quad \text{for } s, t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$$

That is, $\text{sim}(s, t)$ is the minimal depth at which s and t differ, or ∞ if $s = t$. Based on this, a distance function \mathbf{d} can be defined by $\mathbf{d}(s, t) = 2^{-\text{sim}(s, t)}$, where we interpret $2^{-\infty}$ as 0. Note that $0 \leq \mathbf{d}(s, t) \leq 1$. In particular, $\mathbf{d}(s, t) = 0$ iff s and t coincide, and $\mathbf{d}(s, t) = 1$ iff s and t differ at the root. The pair $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$ is known to form a complete ultrametric space [2]. *Partial terms*, i.e. terms over signature $\Sigma_\perp = \Sigma \uplus \{\perp\}$ with \perp a fresh constant symbol, can be endowed with a binary 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 . Interpreting the term \perp as denoting “undefined”, \leq_\perp can be read as “is less defined than”. The pair $(\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V}), \leq_\perp)$ is known to form a complete semilattice [12]. 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. With this, \leq_\perp can be characterised alternatively by $s \leq_\perp t$ iff $s(\pi) = t(\pi)$ for all $\pi \in \mathcal{P}_\perp(s)$. To explicitly distinguish them from partial terms, we call terms in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ *total*.

1.5. Term Rewriting Systems. 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 occurring in r also occur in l . Note that this notion of a TRS deviates slightly from the standard notion of TRSs in the literature on infinitary rewriting [14] in that it allows infinite terms on the left-hand side of rewrite rules! This generalisation will be necessary to accommodate Böhm extensions, which are introduced later in Section 2.2. TRSs having only finite left-hand sides are called *left-finite*.

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

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

Instead of $s \rightarrow_{\mathcal{R}} t$, we sometimes write $s \rightarrow_{\pi, \rho} t$ in order to indicate the applied rule ρ and the position π , or simply $s \rightarrow t$. 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 $r\sigma$.

Let $\rho: l \rightarrow r$ be a term rewrite rule. The *pattern* of ρ is the context $l\sigma$, where σ is the substitution $\{x \mapsto \square \mid x \in \mathcal{V}\}$ that maps all variables to \square . If t is a ρ -redex, then the pattern P of ρ is also called the *redex pattern* of t w.r.t. ρ . When referring to the occurrences in a pattern, occurrences of the symbol \square are neglected.

Let $\rho_1: l_1 \rightarrow r_1$, $\rho_2: l_2 \rightarrow r_2$ be rules in a TRS \mathcal{R} . The rules ρ_1, ρ_2 are said to *overlap* if there is a non-variable position π in l_1 such that $l_1|_\pi$ and l_2 are unifiable and π is not the root position $\langle \rangle$ in case ρ_1, ρ_2 are renamed copies of the same rule. A TRS is called *non-overlapping* if none of its rules overlap. A term t is called *linear* if each variable occurs at most once in t . The TRS \mathcal{R} is called *left-linear* if the left-hand side of every rule in \mathcal{R} is linear. It is called *orthogonal* if it is left-linear and non-overlapping.

2. METRIC INFINITARY TERM REWRITING

In this section we briefly recall the metric model of infinitary term rewriting [16] 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 shall develop correspondence theorems (Theorem 4.9 and Theorem 4.12) that relate convergence in the metric model and convergence in the partial order model. Specifically, these correspondence results show that the two notions of convergence coincide if we restrict ourselves to total terms.

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

Definition 2.1 (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 $(\varphi_\iota)_{\iota < \alpha}$ of reduction steps in \mathcal{R} such that there is a sequence $(t_\iota)_{\iota < \hat{\alpha}}$ of terms, with $\hat{\alpha} = \alpha$ if S is open and $\hat{\alpha} = \alpha + 1$ if S is closed, such that $\varphi_\iota: t_\iota \rightarrow t_{\iota+1}$. If S is finite, we write $S: t_0 \rightarrow^* t_\alpha$.

This definition of reductions is a straightforward generalisation of finite reductions. As an example consider the TRS with the single rule $a \rightarrow f(a)$. In this system we get a reduction $S: a \rightarrow^* f(f(f(a)))$ of length 3:

$$S = \langle \varphi_0: a \rightarrow f(a), \varphi_1: f(a) \rightarrow f(f(a)), \varphi_2: f(f(a)) \rightarrow f(f(f(a))) \rangle$$

In a more concise notation we write

$$S: a \rightarrow f(a) \rightarrow f(f(a)) \rightarrow f(f(f(a)))$$

Clearly, we can extend this reduction arbitrarily often which results in the following infinite reduction of length ω :

$$T: a \rightarrow f(a) \rightarrow f^2(a) \rightarrow f^3(a) \rightarrow f^4(a) \rightarrow \dots$$

However, this is as far we can go with this simple definition of reductions. As soon as we go beyond ω , we get reductions which do not make sense. For example, consider the following reduction:

$$T \cdot S: a \rightarrow f(a) \rightarrow f^2(a) \rightarrow f^3(a) \rightarrow f^4(a) \rightarrow \dots a \rightarrow f(a) \rightarrow f(f(a)) \rightarrow f(f(f(a)))$$

The reduction T of length ω can be extended by an arbitrary reduction, e.g. by the reduction S . The notion of reductions according to Definition 2.1 is only meaningful if restricted to reductions of length at most ω . The problem is that the ω -th step in the reduction, viz. the second step of the form $a \rightarrow f(a)$ in the example above, is completely independent of all previous steps since it does not have an immediate predecessor. This issue occurs at each limit ordinal number. 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 the next section we

will present the well-know metric approach to this. Later in Section 3, we will introduce a novel approach using partial orders.

2.1. Metric Convergence. In this section we consider two notions of *convergence* based on the metric on terms as defined in Section 1.4. We consider both the weak [9] and the strong [16] variant known from the literature. Related to this notion of convergence is a corresponding notion of *continuity*. In order to distinguish both from the partial order model that we will introduce in Section 3 we will use the names *weak* resp. *strong m-convergence* and *weak* resp. *strong m-continuity*.

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 convergence resp. continuity, which does not only depend on the sequence of terms that are derived within the reduction but does also depend on the positions where the contractions take place:

Definition 2.2 (*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

- (i) *weakly m-continuous* in \mathcal{R} , written $S: t_0 \xrightarrow{m} \mathcal{R} \dots$, if $\lim_{i \rightarrow \lambda} t_i = t_\lambda$ for each limit ordinal $\lambda < \alpha$.
- (ii) *strongly m-continuous* in \mathcal{R} , written $S: t_0 \xrightarrow{m} \mathcal{R} \dots$, if it is weakly *m-continuous* and for each limit ordinal $\lambda < \alpha$, the sequence $(|\pi_i|)_{i < \lambda}$ of contraction depths tends to infinity.
- (iii) *weakly m-converging* to t in \mathcal{R} , written $S: t_0 \xrightarrow{m} \mathcal{R} t$, if it is weakly *m-continuous* and $t = \lim_{i \rightarrow \hat{\alpha}} t_i$.
- (iv) *strongly m-converging* to t in \mathcal{R} , written $S: t_0 \xrightarrow{m} \mathcal{R} t$, if it is strongly *m-continuous*, weakly *m-converges* to t and, in case that S is open, $(|\pi_i|)_{i < \alpha}$ tends to infinity.

Whenever $S: t_0 \xrightarrow{m} \mathcal{R} t$ or $S: t_0 \xrightarrow{m} \mathcal{R} t$, we say that t is weakly resp. strongly *m-reachable* from t_0 in \mathcal{R} . By abuse of notation we use $\xrightarrow{m} \mathcal{R}$ and $\xrightarrow{m} \mathcal{R}$ as a binary relation to indicate weakly resp. strongly *m-reachability*. In order to indicate the length of S and the TRS \mathcal{R} , we write $S: t_0 \xrightarrow{m} \mathcal{R}^\alpha t$ resp. $S: t_0 \xrightarrow{m} \mathcal{R}^\alpha t$. The empty reduction $\langle \rangle$ is considered weakly/strongly *m-continuous* and *m-convergent* for any identical start and end term, i.e. $\langle \rangle: t \xrightarrow{m} \mathcal{R} t$ for all $t \in \mathcal{T}(\Sigma, \mathcal{V})$.

From the above definition it is clear that strong *m-convergence* implies both weak *m-convergence* and strong *m-continuity* and that both weak *m-convergence* and strong *m-continuity* imply weak *m-continuity*, respectively. This is indicated in Figure 1. It is important to recognise that *m-convergence* implies *m-continuity*. Hence, only meaningful, i.e. *m-continuous*, reductions can be *m-convergent*.

For a reduction to be weakly *m-continuous*, each open *proper* prefix of the underlying sequence $(t_i)_{i < \hat{\alpha}}$ of terms must converge to the term following next in the sequence – or, equivalently, $(t_i)_{i < \hat{\alpha}}$ must be continuous. For strong *m-continuity*, additionally, the depth at which contractions take place has to tend to infinity for each of the reduction's open proper prefixes. The convergence properties do only differ from the continuity properties in that they require the above conditions to hold for *all* open prefixes, i.e. including the whole reduction itself unless it is closed. For example, considering the rule $a \rightarrow f(a)$, the reduction $g(a) \rightarrow g(f(a)) \rightarrow g(f(f(a))) \rightarrow \dots$ strongly *m-converges* to the infinite term $g(f^\omega)$. The first step takes place at depth 1, the next step at depth 2 and so on. 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




FIGURE 1. Relation between continuity and convergence properties.

trivially strongly m -continuous but is now not strongly m -convergent since every step in this reduction takes place at depth 0, i.e. the sequence of reduction depths does not tend to infinity. However, the reduction still weakly m -converges to $g(f^\omega)$.

In contrast to the strong notions of continuity and convergence, the corresponding weak variants are independent from the rules that are applied during the reduction. What makes strong m -convergence (and -continuity) strong is the fact that it employs a conservative overapproximation of the differences between consecutive terms in the reduction. For weak m -convergence the distance $\mathbf{d}(t_\iota, t_{\iota+1})$ between consecutive terms in a reduction $(t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \lambda}$ has to tend to 0. For strong m -convergence the depth $|\pi_\iota|$ of the reduction steps has to tend to infinity. In other words, $2^{-|\pi_\iota|}$ has to tend to 0. Note that $2^{-|\pi_\iota|}$ is a conservative overapproximation of $\mathbf{d}(t_\iota, t_{\iota+1})$, i.e. $2^{-|\pi_\iota|} \geq \mathbf{d}(t_\iota, t_{\iota+1})$. So strong m -convergence is simply weak m -convergence w.r.t. this overapproximation of \mathbf{d} [4]. If this approximation is actually precise, i.e. coincides with the actual value, both notions of m -convergence coincide.

Remark 2.3. The notion of m -continuity can be defined solely in terms of m -convergence [4]. More precisely, we have for each reduction $S = (t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \alpha}$ that S is weakly m -continuous iff every (open) proper prefix of $S|_\beta$ weakly m -converges to t_β . Analogously, strong m -continuity can be characterised in terms of strong m -convergence. An easy consequence of this is that m -converging reductions are closed under concatenation, i.e. $S: s \rightsquigarrow t, T: t \rightsquigarrow u$ implies $S \cdot T: s \rightsquigarrow u$ and likewise for strong m -convergence.

For the most part our focus in this paper is set on strong m -convergence and its partial order correspondent that we will introduce in Section 3. Weak m -convergence is well-known to be rather unruly [23]. Strong convergence is far more well-behaved [16]. Most prominently, we have the following Compression Lemma [16] which in general does not hold for weak m -convergence:

Theorem 2.4 (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 a strongly m -converging reduction can be approximated arbitrarily accurately by a finite reduction:

Corollary 2.5 (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 $s \mathop{m}\!\!\!\rightarrow t$. By Theorem 2.4, there is a reduction $S: s \mathop{m}\!\!\!\rightarrow^{\leq \omega} t$. If S is of finite length, then we are done. If $S: s \mathop{m}\!\!\!\rightarrow^{\omega} t$, then, by strong 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 . \square

As a special case of the above corollary, we obtain that $s \mathop{m}\!\!\!\rightarrow t$ implies $s \rightarrow^* t$ whenever t is a finite term.

An important difference between m -converging reductions and finite reductions is the confluence of orthogonal systems. In contrast to finite reachability, m -reachability of orthogonal TRSs – even in its strong variant – does not necessarily have the diamond property, i.e. orthogonal systems are confluent but not infinitarily confluent [16]:

Example 2.6 (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 $S: t \mathop{m}\!\!\!\rightarrow g^\omega$ and $T: t \mathop{m}\!\!\!\rightarrow f^\omega$ by successively contracting all ρ_1 -resp. ρ_2 -redexes. However, there is no term s such that $g^\omega \mathop{m}\!\!\!\rightarrow s \mathop{m}\!\!\!\leftarrow f^\omega$ (or $g^\omega \mathop{m}\!\!\!\rightarrow s \mathop{m}\!\!\!\leftarrow f^\omega$) as both g^ω and f^ω can only be rewritten to themselves, respectively.

In the following section we discuss a method for obtaining an appropriate notion of transfinite reachability based on strong m -reachability which actually has the diamond property.

2.2. Meaningless Terms and Böhm Trees. At the end of the previous section we have seen that orthogonal TRSs are in general not infinitarily confluent. However, as Kennaway et al. [16] have shown, orthogonal TRSs are infinitarily confluent modulo so-called *hyper-collapsing* terms – in the sense that two forking strongly m -converging reductions $t \mathop{m}\!\!\!\rightarrow t_1, t \mathop{m}\!\!\!\rightarrow t_2$ can always be extended by two strongly m -converging reductions $t_1 \mathop{m}\!\!\!\rightarrow t_3, t_2 \mathop{m}\!\!\!\rightarrow t'_3$ such that the resulting terms t_3, t'_3 only differ in the hyper-collapsing subterms they contain.

This result was later generalised by Kennaway et al. [17] to develop an axiomatic theory of *meaningless terms*. Intuitively, a set of meaningless terms in this setting consists of terms that are deemed meaningless since, from a term rewriting perspective, they cannot be distinguished from one another and they do not contribute any information to any computation. Kennaway et al. capture this by a set of axioms that characterise a set of meaningless terms. For orthogonal TRSs, one such set of terms, in fact the least such set, is the set of *root-active* terms [17]:

Definition 2.7 (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 2.6.

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 the set \mathcal{RA} instead of the original more general axiomatic treatment by Kennaway et al. [17].

Since, denotationally, root-active terms cannot be distinguished from each other it is appropriate to equate them [17]. This can be achieved 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 [17]:

Definition 2.8 (Böhm extension). Let $\mathcal{R} = (\Sigma, R)$ be a TRS, 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 strongly 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 .

The definition of \mathcal{U}_\perp is quite subtle and deserves further attention before we move on. According to the definition, a term t is in \mathcal{U}_\perp iff the term obtained from t by replacing occurrences of \perp in t by terms from \mathcal{U} is also in \mathcal{U} . More illuminating, however, is the converse view, i.e. how to construct a term in \mathcal{U}_\perp from a term in \mathcal{U} . First of all, any term in \mathcal{U} is also in \mathcal{U}_\perp . Secondly, we may obtain a term in \mathcal{U}_\perp by taking a term $t \in \mathcal{U}$ and replacing any number of subterms of t that are in \mathcal{U} by \perp . For Böhm extensions, this means that we may contract any term $t \in \mathcal{U}$ to \perp , even if we already contracted some proper subterms of t to \perp before.

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.

Remark 2.9 (closure under substitution). Note that, for orthogonal TRSs, $\mathcal{R}\mathcal{A}$ is closed under substitutions and, hence, so is $\mathcal{R}\mathcal{A}_\perp$ [17]. Therefore, whenever $C[t] \rightarrow_{\mathcal{R}\mathcal{A}, \perp} C[\perp]$, we can assume that $t \in \mathcal{R}\mathcal{A}_\perp$.

With the additional rules provided by the Böhm extension, we gain infinitary confluence of orthogonal systems:

Theorem 2.10 (infinitary confluence of Böhm-converging reductions, [17]). *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. $s_1 \leftarrow_{\mathcal{B}}^m t \xrightarrow{\mathcal{B}}^m s_2$ implies $s_1 \xrightarrow{\mathcal{B}}^m t' \leftarrow_{\mathcal{B}}^m s_2$.*

The lack of confluence for strongly 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 2.6, we can see that g^ω and f^ω can be rewritten to \perp as both terms are root-active.

In fact, w.r.t. Böhm-convergence, every term of an orthogonal TRS has a normal form:

Theorem 2.11 (infinitary normalisation of Böhm-converging reductions, [17]). *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. $\mathcal{R}\mathcal{A}$) [17].

The rest of this paper is concerned with establishing an alternative to the metric notion of convergence based on the partial order on terms that is equivalent to the Böhm extension approach.

3. PARTIAL ORDER INFINITARY REWRITING

In this section we introduce an alternative model of infinitary term rewriting which uses the partial order on terms to formalise convergence of transfinite reductions. To this end we will turn to partial terms which, like in the setting of Böhm extensions, have an additional constant 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 3.1. Consider the TRS consisting of the rules $h(x) \rightarrow h(g(x))$, $b \rightarrow g(b)$ and the term $t = f(h(a), b)$. 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 .

The reduction S weakly m -converges to the term $f(h(g^\omega), g^\omega)$. But it does not *strongly* 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. Within the partial order model we will still be able to obtain that S weakly converges to $f(h(g^\omega), g^\omega)$ but we will also obtain that it strongly converges to the term $f(\perp, g^\omega)$. That is, we will be able to identify that the reduction S strongly converges except at position $\langle 0 \rangle$, the first argument of f .

3.1. Partial Order Convergence. In order to formalise continuity and convergence in terms of the complete semilattice $(\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V}), \leq_\perp)$ instead of the complete metric space $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$, we move from the limit of the metric space to the limit inferior of the complete semilattice:

Definition 3.2 (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 $\mathcal{R}_\perp = (\Sigma_\perp, R)$. The reduction S is called

- (i) *weakly p -continuous* in \mathcal{R} , written $S: t_0 \xrightarrow{\mathcal{R}} \dots$, if $\liminf_{\iota \rightarrow \lambda} t_\iota = t_\lambda$ for each limit ordinal $\lambda < \alpha$.
- (ii) *strongly p -continuous* in \mathcal{R} , written $S: t_0 \xrightarrow{\mathcal{R}} \dots$, 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 , which we indicate by writing $\varphi_\iota: t_\iota \rightarrow_{c_\iota} t_{\iota+1}$.
- (iii) *weakly p -converging* to t in \mathcal{R} , written $S: t_0 \xrightarrow{\mathcal{R}} t$, if it is weakly p -continuous and $t = \liminf_{\iota \rightarrow \hat{\alpha}} t_\iota$.
- (iv) *strongly p -converging* to t in \mathcal{R} , written $S: t_0 \xrightarrow{\mathcal{R}} t$, if it is strongly p -continuous and S is closed with $t = t_{\alpha+1}$ or $t = \liminf_{\iota \rightarrow \alpha} c_\iota$.

Whenever $S: t_0 \xrightarrow{\mathcal{R}} t$ or $S: t_0 \xrightarrow{\mathcal{R}} t$, we say that t is weakly resp. strongly p -reachable from t_0 in \mathcal{R} . By abuse of notation we use $\xrightarrow{\mathcal{R}}$ and $\xrightarrow{\mathcal{R}}$ as a binary relation to indicate weak resp. strong p -reachability. In order to indicate the length of S and the TRS \mathcal{R} , we write $S: t_0 \xrightarrow{\mathcal{R}^\alpha} t$ resp. $S: t_0 \xrightarrow{\mathcal{R}^\alpha} t$. The empty reduction $\langle \rangle$ is considered weakly/strongly p -continuous and p -convergent for any start and end term, i.e. $\langle \rangle: t \xrightarrow{\mathcal{R}} t$ for all $t \in \mathcal{T}(\Sigma, \mathcal{V})$.

The definitions of weak p -continuity and weak p -convergence are straightforward “translations” from the metric setting to the partial order setting replacing the limit $\lim_{\iota \rightarrow \alpha}$ by the limit inferior $\liminf_{\iota \rightarrow \alpha}$. On the other hand, the definitions of the strong counterparts

seem a bit different compared to the metric model: Whereas strong m -convergence simply adds a side condition regarding the depth $|\pi_\iota|$ of the reduction steps, strong p -convergence is defined in a different way compared to the weak variant. Instead of the terms t_ι of the reduction, it considers the contexts $c_\iota = t_\iota[\pi_\iota]_\perp$. However, one can surmise some similarity due to the fact that the partial order model of strong convergence indirectly takes into account the position π_ι of each reduction step as well. Moreover, for the sake of understanding the intuition of strong p -convergence it is better to compare the contexts c_ι rather with the glb of two consecutive terms $t_\iota \sqcap t_{\iota+1}$ instead of the term t_ι itself. The following proposition allows precisely that.

Proposition 3.3 (limit inferior of open sequences). *Let $(a_\iota)_{\iota < \lambda}$ be an open sequence in a complete semilattice. Then it holds that $\liminf_{\iota \rightarrow \lambda} a_\iota = \liminf_{\iota \rightarrow \lambda} (a_\iota \sqcap a_{\iota+1})$.*

Proof. Let $\bar{a} = \liminf_{\iota \rightarrow \lambda} a_\iota$ and $\hat{a} = \liminf_{\iota \rightarrow \lambda} (a_\iota \sqcap a_{\iota+1})$. Since $a_\iota \sqcap a_{\iota+1} \leq a_\iota$ for each $\iota < \lambda$, we have $\hat{a} \leq \bar{a}$. On the other hand, consider the sets $\bar{A}_\alpha = \{a_\iota \mid \alpha \leq \iota < \lambda\}$ and $\hat{A}_\alpha = \{a_\iota \sqcap a_{\iota+1} \mid \alpha \leq \iota < \lambda\}$ for each $\alpha < \lambda$. Of course, we then have $\prod \bar{A}_\alpha \leq a_\iota$ for all $\alpha \leq \iota < \lambda$, and thus also $\prod \bar{A}_\alpha \leq a_\iota \sqcap a_{\iota+1}$ for all $\alpha \leq \iota < \lambda$. Hence, $\prod \bar{A}_\alpha$ is a lower bound of \hat{A}_α which implies that $\prod \bar{A}_\alpha \leq \prod \hat{A}_\alpha$. Consequently, $\bar{a} \leq \hat{a}$ and, due to the antisymmetry of \leq , we can conclude that $\bar{a} = \hat{a}$. \square

With this in mind we can replace $\liminf_{\iota \rightarrow \lambda} t_\iota$ in the definition of weak p -convergence resp. p -continuity with $\liminf_{\iota \rightarrow \lambda} t_\iota \sqcap t_{\iota+1}$. From there it is easier to see the intention of moving from $t_\iota \sqcap t_{\iota+1}$ to the context $t_\iota[\perp]_{\pi_\iota}$ in order to model strong convergence:

What makes the notion of strong p -convergence (and p -continuity) *strong*, similar to the notion of strong m -convergence (resp. m -continuity), 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 a reduction step $\varphi_\iota: t_\iota \rightarrow_{\pi_\iota} t_{\iota+1}$. More specifically, we have that $t_\iota[\perp]_{\pi_\iota} \leq_\perp t_\iota \sqcap t_{\iota+1}$. That is, as in the metric model of strong convergence, the difference between two consecutive terms is overapproximated by using the position of the reduction step as an indicator. Likewise, strong p -convergence is simply weak p -convergence w.r.t. this underapproximation of $t_\iota \sqcap t_{\iota+1}$ [4]. If this approximation is actually precise, i.e. coincides with the actual value, both notions of p -convergence coincide.

Remark 3.4. As for the metric model, also in the partial order model, continuity can be defined solely in terms of convergence [4]. More precisely, we have for each reduction $S = (t_\iota \rightarrow t_{\iota+1})_{\iota < \alpha}$ that S is weakly p -continuous iff every (open) proper prefix of $S|_\beta$ weakly p -converges to t_β . Analogously, strong p -continuity can be characterised in terms of strong p -convergence. An easy consequence of this is that p -converging reductions are closed under concatenation, i.e. $S: s \hookrightarrow t, T: t \hookrightarrow u$ implies $S \cdot T: s \hookrightarrow u$ and likewise for strong p -convergence.

In order to understand the difference between weak and strong p -convergence let us look at a simple example:

Example 3.5. Consider the TRS with the single rule $f(x, y) \rightarrow f(y, x)$. This rule induces the following reduction:

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

S simply alternates between the terms $f(a, f(g(a), g(b)))$ and $f(a, f(g(b), g(a)))$ by swapping the arguments of the inner f occurrence. The reduction is depicted in Figure 2. The picture




FIGURE 2. Reduction with stable context.


 (a) Limit w.r.t. weak p -convergence.

 (b) Limit w.r.t. strong p -convergence.

 FIGURE 3. Limits of a p -converging reduction.

illustrates the parts of the terms that remain *unchanged* and those that remain completely *untouched* by the corresponding reduction step by using a lighter resp. a darker shade of grey. The unchanged part corresponds to the glb of the two terms of a reduction step, viz. for the first step

$$f(a, f(g(a), g(b))) \sqcap f(a, f(g(b), g(a))) = f(a, f(g(\perp), g(\perp)))$$

By symmetry, the glb of the terms of the second step is the same one. It is depicted in Figure 3a. Let $(t_i)_{i < \omega}$ be the sequence of terms of the reduction S . By definition, S weakly p -converges to $\liminf_{i \rightarrow \omega} t_i$. According to Proposition 3.3, this is equal to $\liminf_{i \rightarrow \omega} (t_i \sqcap t_{i+1})$. Since $t_i \sqcap t_{i+1}$ is constantly $f(a, f(g(\perp), g(\perp)))$, the reduction sequence weakly p -converges to $f(a, f(g(\perp), g(\perp)))$.

Similarly, the part of the term that remains untouched by the reduction step corresponds to the context. For the first step, this is $f(a, \perp)$. It is depicted in Figure 3b. By definition, S strongly p -converges to $\liminf_{i \rightarrow \omega} c_i$ for $(c_i)_{i < \omega}$ the sequence of contexts of S . As one can see in Figure 2, the context constantly remains $f(a, \perp)$. Hence, S strongly p -converges to $f(a, \perp)$. The example sequence is a particularly simple one as both the glbs $t_i \sqcap t_{i+1}$ and the contexts c_i remain stable. In general, this is not necessary, of course.

One can clearly see from the definition that, as for their metric counterparts, weak resp. strong p -convergence implies weak resp. strong p -continuity. In contrast to the metric model, however, also the converse implication holds! Since the partial order \leq_{\perp} on partial terms forms a complete semilattice, the limit inferior is defined for any non-empty sequence

of partial terms. Hence, any weakly resp. strongly p -continuous reduction is also weakly resp. strongly p -convergent. This is a major difference to m -convergence/ $-$ continuity. Nevertheless, p -convergence constitutes 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 3.1 and Example 3.5. In fact, as we will show in Section 4, whenever there are no \perp ’s involved, i.e. if there is no “local divergence”, m -convergence and p -convergence coincide – both in the weak and the strong variant.

Recall that strong m -continuity resp. m -convergence implies weak m -continuity resp. m -convergence. This is not the case in the partial order setting. The reason for this is that strong p -convergence resp. p -continuity is defined differently compared to its weak variant. It uses the contexts instead of the terms in the reduction, whereas in the metric setting the strong notion of convergence is a mere restriction of the weak counterpart as we have observed earlier.

Example 3.6. Consider the TRS consisting of the rules $\rho_1: h(x) \rightarrow h(g(x)), \rho_2: f(x) \rightarrow g(x)$ and the reductions

$$S: f(h(a)) \rightarrow_{\rho_1} f(h(g(a))) \rightarrow_{\rho_1} f(h(g(g(a)))) \rightarrow_{\rho_1} \dots \quad \text{and} \quad T: f(\perp) \rightarrow_{\rho_2} g(\perp)$$

Then the reduction

$$S \cdot T: f(h(a)) \rightarrow_{\rho_1} f(h(g(a))) \rightarrow_{\rho_1} f(h(g(g(a)))) \rightarrow_{\rho_1} \dots f(\perp) \rightarrow_{\rho_2} g(\perp)$$

is clearly both strongly p -continuous and $-$ convergent. On the other hand it is neither weakly p -continuous nor $-$ convergent for the simple fact that S does not weakly p -converge to $f(\perp)$ but to $f(h(g^\omega))$.

Nevertheless, by observing that $\liminf_{\iota \rightarrow \alpha} c_\iota \leq_\perp \liminf_{\iota \rightarrow \alpha} t_\iota$ since $c_\iota \leq_\perp t_\iota$ for each $\iota < \alpha$, we obtain the following weaker relation between weak and strong p -convergence:

Proposition 3.7. *Let \mathcal{R} be a left-linear TRS with $s \xrightarrow{\mathcal{R}} t$. Then there is a term $t' \geq_\perp t$ with $s \xrightarrow{\mathcal{R}} t'$.*

Proof. Let $S = (\varphi_\iota: t_\iota \rightarrow_{\rho_\iota} t_{\iota+1})_{\iota < \alpha}$ be a reduction strongly p -converging to t_α . By induction we construct for each prefix $S|_\beta$ of S a reduction $S'_\beta = (\varphi'_\iota: t'_\iota \rightarrow_{\rho_\iota} t'_{\iota+1})_{\iota < \beta}$ weakly p -converging to a term t'_β such that $t_\iota \leq_\perp t'_\iota$ for each $\iota \leq \alpha$. The proposition then follows from the case where $\beta = \alpha$.

The case $\beta = 0$ is trivial. If $\beta = \gamma + 1$, then by induction hypothesis we have a reduction $S'_\gamma: t'_\gamma \xrightarrow{\mathcal{R}} t'_\gamma$. Since $t_\gamma \leq_\perp t'_\gamma$ and t_γ is a ρ_γ -redex, also t'_γ is a ρ_γ -redex due to the left-linearity of \mathcal{R} . Hence, there is a reduction step $\varphi'_\gamma: t'_\gamma \rightarrow t'_\beta$. One can easily see that then $t_\beta \leq_\perp t'_\beta$. Hence, $S'_\beta = S'_\gamma \cdot \langle \varphi'_\gamma \rangle$ satisfies desired conditions.

If β is a limit ordinal, we can apply the induction hypothesis to obtain for each $\gamma < \beta$ a reduction $S'_\gamma = (\varphi'_\iota: t'_\iota \rightarrow_{\rho_\iota} t'_{\iota+1})_{\iota < \gamma}$ that weakly p -converges to $t'_\gamma \geq_\perp t_\gamma$. Hence, according to Remark 3.4, $S'_\beta = (\varphi'_\iota: t'_\iota \rightarrow_{\rho_\iota} t'_{\iota+1})_{\iota < \beta}$ is weakly p -continuous. Therefore, we obtain that S'_β weakly p -converges to $t'_\beta = \liminf_{\iota \rightarrow \beta} t'_\iota$. Moreover, since $c_\iota \leq_\perp t_\iota$ and $t_\iota \leq_\perp t'_\iota$ for each $\iota < \beta$, we can conclude that

$$t_\beta = \liminf_{\iota \rightarrow \beta} c_\iota \leq_\perp \liminf_{\iota \rightarrow \beta} t_\iota \leq_\perp \liminf_{\iota \rightarrow \beta} t'_\iota = t'_\beta. \quad \square$$

And indeed, returning to Example 3.6, we can see that there is a reduction

$$f(h(a)) \rightarrow_{\rho_1} f(h(g(a))) \rightarrow_{\rho_1} f(h(g(g(a)))) \rightarrow_{\rho_1} \dots f(h(g^\omega)) \rightarrow_{\rho_2} g(h(g^\omega))$$

that, starting from $f(h(a))$, weakly p -converges to $g(h(g^\omega))$ which is strictly larger than $g(\perp)$.

A simple example shows that left-linearity is crucial for the above proposition:

Example 3.8. Let \mathcal{R} be a TRS consisting of the rules

$$\rho_1: a \rightarrow a, \quad \rho_2: b \rightarrow b, \quad \rho_3: f(x, x) \rightarrow c.$$

We then get the strongly p -converging reduction

$$f(a, b) \rightarrow_{\rho_1} f(a, b) \rightarrow_{\rho_2} f(a, b) \rightarrow_{\rho_1} f(a, b) \rightarrow_{\rho_2} \dots f(\perp, \perp) \rightarrow_{\rho_3} c$$

Yet, there is no reduction in \mathcal{R} that, starting from $f(a, b)$, weakly p -converges to c .

3.2. Strong p -Convergence. In this paper we are mainly focused on the strong notion of convergence. To this end, the rest of this section will be concerned exclusively with strong p -convergence. We will, however, revisit weak p -convergence in Section 4 when comparing it to weak m -convergence.

Note that in the partial order model 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 a strongly p -convergent reduction $S: s \twoheadrightarrow t$ is total, we mean that both the reduction S and the final term t are total.¹

In order to understand the behaviour strong p -convergence, we need to look at how the lub and glb of a set of terms looks like. The following two lemmas provide some insight.

Lemma 3.9 (lub of terms). *For each $T \subseteq \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$ and $t = \bigsqcup T$, the following holds*

- (i) $\mathcal{P}(t) = \bigcup_{s \in T} \mathcal{P}(s)$
- (ii) $t(\pi) = f$ iff there is some $s \in T$ with $s(\pi) = f$ for each $f \in \Sigma \cup \mathcal{V}$, and position π .

Proof. Clause ((i)) follows straightforwardly from clause ((ii)). The “if” direction of ((ii)) follows from the fact that if $s \in T$, then $s \leq_\perp t$ and, therefore, $s(\pi) = f$ implies $t(\pi) = f$. For the “only if” direction assume that no $s \in T$ satisfies $s(\pi) = f$. Since, $s \leq_\perp t$ for each $s \in T$, we have $\pi \notin \mathcal{P}_\perp(s)$ for each $s \in T$. But then $t' = t[\perp]_\pi$ is an upper bound of T with $t' <_\perp t$. This contradicts the assumption that t is the least upper bound of T . \square

Lemma 3.10 (glb of terms). *Let $T \subseteq \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$ and P a set of positions closed under prefixes such that all terms in T coincide in all positions in P , i.e. $s(\pi) = t(\pi)$ for all $\pi \in P$ and $s, t \in T$. Then the glb $\bigsqcap T$ also coincides with all terms in T in all positions in P .*

Proof. Construct a term s such that it coincides with all terms in T in all positions in P and has \perp at all other positions. That is, given an arbitrary term $t \in T$, we define s as the unique term with $s(\pi) = t(\pi)$ for all $\pi \in P$, and $s(\pi \cdot \langle i \rangle) = \perp$ for all $\pi \in P$ with $\pi \cdot \langle i \rangle \in \mathcal{P}(t) \setminus P$. Then s is a lower bound of T . By construction, s coincides with all terms in T in all positions in P . Since $s \leq_\perp \bigsqcap T$, this property carries over to $\bigsqcap T$. \square

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 .

Following the two lemmas above, we can observe that – intuitively speaking – the limit inferior $\liminf_{\iota \rightarrow \alpha} t_\iota$ of a sequence of terms is the term that contains those parts that become *eventually stable* in the sequence. Remaining holes in the term structure are filled with ' \perp 's. Let us see what this means for strongly p -converging reductions:

Lemma 3.11 (non- \perp symbols in open reductions). *Let $\mathcal{R} = (\Sigma, R)$ be a TRS and $S: s \xrightarrow{\mathcal{R}}^\lambda t$ an open reduction with $S = (t_\iota \rightarrow_{\pi_\iota, c_\iota} t_{\iota+1})_{\iota < \lambda}$. Then the following statements are equivalent for all positions π :*

- (a) $t(\pi) \neq \perp$.
- (b) *there is some $\alpha < \lambda$ such that $c_\alpha(\pi) = t(\pi) \neq \perp$ for all $\alpha \leq \iota < \lambda$.*
- (c) *there is some $\alpha < \lambda$ such that $t_\alpha(\pi) = t(\pi) \neq \perp$ and $\pi_\alpha \not\leq \pi$ for all $\alpha \leq \iota < \lambda$.*
- (d) *there is some $\alpha < \lambda$ such that $t_\alpha(\pi) \neq \perp$ and $\pi_\alpha \not\leq \pi$ for all $\alpha \leq \iota < \lambda$.*

Proof. At first consider the implication from ((a)) to ((b)). To this end, let $t(\pi) \neq \perp$ and $s_\gamma = \prod_{\gamma \leq \iota < \lambda} c_\iota$ for each $\gamma < \lambda$. Note that then $t = \bigsqcup_{\gamma < \lambda} s_\gamma$. Applying Lemma 3.9 yields that there is some $\alpha < \lambda$ such that $s_\alpha(\pi) = t(\pi)$. Moreover, for each $\alpha \leq \iota < \lambda$, we have $s_\alpha \leq_\perp c_\iota$ and, therefore, $s_\alpha(\pi) = c_\iota(\pi)$. Consequently, we obtain $c_\iota(\pi) = t(\pi)$ for all $\alpha \leq \iota < \lambda$.

Next consider the implication from ((b)) to ((c)). Let $\alpha < \lambda$ be such that $c_\alpha(\pi) = t(\pi) \neq \perp$ for all $\alpha \leq \iota < \lambda$. Recall that $c_\iota = t_\iota[\perp]_{\pi_\iota}$ for all $\iota < \lambda$. Hence, the fact that $c_\alpha(\pi) \neq \perp$ for all $\alpha \leq \iota < \lambda$ implies that $t_\alpha(\pi) = c_\alpha(\pi)$ and that $\pi_\alpha \not\leq \pi$ for all $\alpha \leq \iota < \lambda$. Since $c_\alpha(\pi) = t(\pi) \neq \perp$, we also have $t_\alpha(\pi) = t(\pi) \neq \perp$.

The implication from ((c)) to ((d)) is trivial.

Finally, consider the implication from ((d)) to ((a)). For this purpose, let $\alpha < \lambda$ be such that (1) $\pi \in \mathcal{P}_\perp(t_\alpha)$ and (2) $\pi_\alpha \not\leq \pi$ for all $\alpha \leq \iota < \lambda$. Consider the set P consisting of all positions in t_α that are prefixes of π . P is obviously closed under prefixes and, because of (2), all terms in the set $T = \{c_\iota \mid \alpha \leq \iota < \lambda\}$ coincide in all positions in P . According to Lemma 3.10, also $s_\alpha = \prod T$ coincides with all terms in T in all positions in P . Since $\pi \in P$ and $c_\alpha \in T$, we thereby obtain that $c_\alpha(\pi) = s_\alpha(\pi)$. As we also have $t_\alpha(\pi) = c_\alpha(\pi)$, due to (2), and $\pi \in \mathcal{P}_\perp(t_\alpha)$ we can infer that $\pi \in \mathcal{P}_\perp(s_\alpha)$. Since $s_\alpha \leq_\perp t$, we can then conclude $\pi \in \mathcal{P}_\perp(t)$. \square

The above lemma is central for dealing with strongly p -convergent reductions. It also reveals how the final term of a strongly p -convergent reduction is constructed. According to the equality of ((a)) and ((c)), the final term has the non- \perp symbol f at some position π iff some term t_α in the reduction also had this symbol f at this position π and no reduction after that term occurred at π or above. In this way, the final outcome of a strongly p -convergent reduction consists of precisely those parts of the intermediate terms which become *eventually persistent* during the reduction, i.e. are from some point on not subjected to contraction any more.

Now we turn to a characterisation of the parts that are not included in the final outcome of a strongly p -convergent reduction, i.e. those that do not become persistent. These parts are either omitted or filled by the placeholder \perp . We will call these positions *volatile*:

Definition 3.12 (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 3.1 the position $\langle 0 \rangle$ is outermost-volatile in the reduction S .

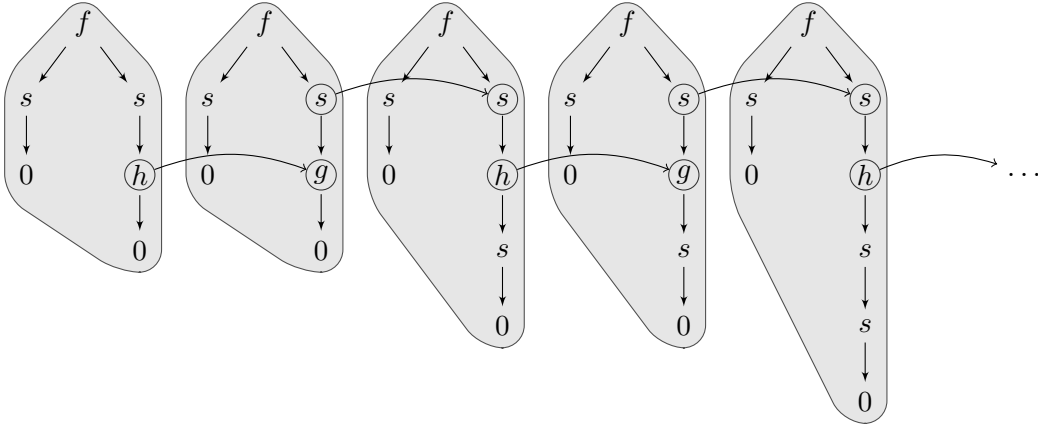


FIGURE 4. Reduction with two nested volatile positions.

Example 3.13 (volatile positions). Consider the TRS \mathcal{R} consisting of the rules

$$\rho_1: h(x) \rightarrow g(x), \quad \rho_2: s(g(x)) \rightarrow s(h(s(x)))$$

\mathcal{R} admits the following reduction S of length ω :

$$\begin{aligned} S: f(s(0), s(h(0))) &\rightarrow_{\rho_1} f(s(0), s(g(0))) \rightarrow_{\rho_2} f(s(0), s(h(s(0)))) \\ &\rightarrow_{\rho_1} f(s(0), s(g(s(0)))) \rightarrow_{\rho_2} f(s(0), s(h(s(s(0)))))) \end{aligned}$$

The reduction S p -converges to $f(s(0), \perp)$, i.e. we have $S: f(s(0), s(h(0))) \xrightarrow{\omega}_{\mathcal{R}} f(s(0), \perp)$. Figure 4 illustrates the reduction indicating the position of each reduction step by two circles and a reduction arrow in between. One can clearly see that both $\pi_1 = \langle 1 \rangle$ and $\pi_2 = \langle 1, 0 \rangle$ are volatile in S . Again and again reductions take place at π_1 and π_2 . Since these are the only volatile positions and π_1 is a prefix of π_2 , we have that π_1 is an outermost-volatile position in S .

As we shall see later in Section 6, volatility is closely related to root-active terms: if a reduction has a volatile position π , then we find a term in the reduction with a root-active subterm at π . Conversely, from each root-active term starts a reduction with volatile position $\langle \rangle$ (cf. Proposition 6.9). This connection between volatility and root-activeness is the cornerstone of the correspondence between p -convergence and Böhm-convergence that we prove in Section 6.

The following lemma shows that \perp symbols are produced precisely at outermost-volatile positions in open reduction.

Lemma 3.14 (\perp subterms in open reductions). *Let $S = (t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \alpha}$ an open reduction p -converging to t_α in some TRS. Then, for every position π , we have the following:*

- (i) *If π is volatile in S , then $\pi \notin \mathcal{P}_\perp(t_\alpha)$.*
- (ii) *$t_\alpha(\pi) = \perp$ iff*
 - (a) *π is outermost-volatile in S , or*
 - (b) *there is some $\beta < \alpha$ such that $t_\beta(\pi) = \perp$ and $\pi_\beta \not\leq \pi$ for all $\beta \leq \iota < \alpha$.*
- (iii) *Let t_ι be total for all $\iota < \alpha$. Then $t_\alpha(\pi) = \perp$ iff π is outermost-volatile in S .*

Proof. ((i)) This follows from Lemma 3.11, in particular the equivalence of ((a)) and ((c)).

((ii)) At first consider the “only if” direction. To this end, suppose that $t_\alpha(\pi) = \perp$. In order to show that then ((a)) or ((b)) holds, we will prove that ((b)) must hold true whenever ((a)) does not hold. For this purpose, we assume that π is not outermost-volatile in S . Note that no proper prefix π' of π can be volatile in S as this would imply, according to clause ((i)), that $\pi' \notin \mathcal{P}_\perp(t_\alpha)$ and, therefore, $\pi \notin \mathcal{P}(t_\alpha)$. Hence, π is also not volatile in S . In sum, no prefix of π is volatile in S . Consequently, there is an upper bound $\beta < \alpha$ on the indices of reduction steps taking place at π or above. But then $t_\beta(\pi) = \perp$ since otherwise Lemma 3.11 would imply that $t_\alpha(\pi) \neq \perp$. This shows that ((b)) holds.

For the converse direction, we will show that both ((a)) and ((b)) independently imply that $t_\alpha(\pi) = \perp$:

((a)) Let π be outermost-volatile in S . By clause ((i)), this implies $\pi \notin \mathcal{P}_\perp(t_\alpha)$. Hence, it remains to be shown that $\pi \in \mathcal{P}(t_\alpha)$. If $\pi = \langle \rangle$, then this is trivial. Otherwise, π is of the form $\pi' \cdot i$. Since all proper prefixes of π are not volatile, there is some $\beta < \alpha$ such that $\pi_\beta = \pi$ and $\pi_\iota \not\leq \pi'$ for all $\beta \leq \iota < \alpha$. This implies that $\pi \in \mathcal{P}(t_\beta)$. Hence, $t_\beta(\pi') = f$ is a symbol having an arity of at least $i + 1$. Consequently, according to Lemma 3.11, also $t_\alpha(\pi') = f$. Since f 's arity is at least $i + 1$, also $\pi = \pi' \cdot i \in \mathcal{P}(t_\alpha)$.

((b)) Let $\beta < \alpha$ such that $t_\beta(\pi) = \perp$ and $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \alpha$. According to Proposition 1.1, we have that $t_\alpha = \bigsqcup_{\beta \leq \gamma < \alpha} \prod_{\gamma \leq \iota < \alpha} c_\iota$. Define $s_\gamma = \prod_{\gamma \leq \iota < \alpha} c_\iota$ for each $\gamma < \alpha$. Since from β onwards no reduction takes place at π or above, it holds that all c_ι , for $\beta \leq \iota < \alpha$, coincide in all prefixes of π . By Lemma 3.10, this also holds for all s_ι and c_ι with $\beta \leq \iota < \alpha$. Since $c_\beta(\pi) = t_\beta(\pi) = \perp$, this means that $s_\iota(\pi) = \perp$ for all $\beta \leq \iota < \alpha$. Recall that $t_\alpha = \bigsqcup_{\beta \leq \gamma < \alpha} s_\gamma$. Hence, according to Corollary 3.9, we can conclude that $t_\alpha(\pi) = \perp$.

((iii)) is a special case of ((ii)): If each t_ι , $\iota < \alpha$, is total, then ((b)) cannot be true. \square

Clause ((ii)) shows that a \perp subterm in the final term can only have its origin either in a preceding term which already contains this \perp which then becomes stable, or in an outermost-volatile position. That is, it is exactly the outermost-volatile positions that generate ' \perp 's.

We can apply this lemma to Example 3.13: As we have seen, the position $\pi_1 = \langle 1 \rangle$ is outermost-volatile in the reduction S mentioned in the example. Hence, S strongly p -converges to a term that has, according to Lemma 3.14, the symbol \perp at position π_1 . That is, S strongly p -converges to $f(s(0), \perp)$.

This characterisation of the final outcome of a p -converging reduction clearly shows that the partial order model captures the intuition of strong convergence in transfinite reductions even though it allows that every continuous reduction is also convergent: The final outcome only represents the parts of the reduction that *are* converging. Locally diverging parts are cut off and replaced by \perp .

In fact, the absence of such local divergence, or volatility, as we call it here, is equivalent to the absence of \perp :

Lemma 3.15 (total reductions). *Let \mathcal{R} be a TRS, s a total term in \mathcal{R} , and $S: s \xrightarrow{p} \mathcal{R} t$. $S: s \xrightarrow{p} \mathcal{R} t$ is total iff no prefix of S has a volatile position.*

Proof. The “only if” direction follows straightforwardly from Lemma 3.14.

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 3.14. \square

Moreover, as we shall show in the next section, if local divergences are excluded, i.e. if total reductions are considered, both the metric model and the partial order model coincide.

4. COMPARING m -CONVERGENCE AND p -CONVERGENCE

In this section we want to compare the metric and the partial order model of convergence. In particular, we shall show that the partial order model is only a conservative extension of the metric model: If we only consider total reductions, i.e. reductions over terms in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$, then m -convergence and p -convergence coincide both in their weak and strong variant.

The first and rather trivial observation to this effect is that already on the level of single reduction steps the partial order model conservatively extends the metric model:

Fact 4.1. Let $\mathcal{R} = (\Sigma, R)$ be a TRS, $\mathcal{R}_\perp = (\Sigma_\perp, R)$, and $s, t \in \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$. Then we have

$$s \rightarrow_{\mathcal{R}, \pi} t \quad \text{iff} \quad s \rightarrow_{\mathcal{R}_\perp, \pi} t \text{ and } s \text{ is total.}$$

The next step is to establish that the underlying structures that are used to formalise convergence exhibit this behaviour as well. That is, the limit inferior in the complete semilattice $(\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V}), \leq_\perp)$ is conservative extension of the limit in the complete metric space $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$. More precisely, we want to have that for a sequence $(t_\iota)_{\iota < \alpha}$ in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$

$$\liminf_{\iota \rightarrow \alpha} t_\iota = \lim_{\iota \rightarrow \alpha} t_\iota \quad \text{whenever} \quad \begin{array}{l} \lim_{\iota \rightarrow \alpha} t_\iota \text{ is defined, or} \\ \liminf_{\iota \rightarrow \alpha} t_\iota \text{ is a total term.} \end{array}$$

Note that, as a corollary, the above property implies that $\lim_{\iota \rightarrow \alpha} t_\iota$ is defined iff $\liminf_{\iota \rightarrow \alpha} t_\iota$ is a total term. In Section 4.1 we shall establish the above property. This result is then used in Section 4.2 in order to show the desired property that p -convergence is a conservative extension of m -convergence in both their respective weak and strong variant.

4.1. Complete Semilattice vs. Complete Metric Space. In order to compare the complete semilattice of partial terms with the complete metric space of term, it is convenient to have an alternative characterisation of the similarity $\mathbf{sim}(s, t)$ of two terms s, t , which in turn provides an alternative characterisation of the metric \mathbf{d} on terms. To this end we use the *truncation* of a term at a certain depth. This notion was originally used by Arnold and Nivat [2] to show that the \mathbf{d} is a complete ultrametric on terms:

Definition 4.2 (truncation). Let $d \in \mathbb{N} \cup \{\infty\}$ and $t \in \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$. The *truncation* $t|d$ of t at depth d is defined inductively on d as follows

$$\begin{aligned} t|0 &= \perp & t|\infty &= t \\ t|d+1 &= \begin{cases} t & \text{if } t \in \mathcal{V} \\ f(t_1|d, \dots, t_k|d) & \text{if } t = f(t_1, \dots, t_k) \end{cases} \end{aligned}$$

More concisely we can say that the truncation of a term t at depth d replaces all subterms at depth d with \perp . From this we can easily establish the following two properties of the truncation:

Proposition 4.3 (truncation). *For each two $s, t \in \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$ we have*

- (i) $t|d \leq_\perp t$ for all $d \in \mathbb{N} \cup \{\infty\}$.
- (ii) $s|d \leq_\perp t$ implies $s|d = t|d$ for all $d \in \mathbb{N} \cup \{\infty\}$ given s is total.

(iii) $s|d = t|d$ for all $d \in \mathbb{N}$ iff $s = t$.

Proof. Straightforward. \square

Recall that the similarity of two terms is the minimal depth at which they differ resp. ∞ if they are equal. However, saying that two terms differ at a certain minimal depth d is the same as saying that the truncation of the two terms at that depth d coincide. This provides an alternative characterisation of similarity:

Proposition 4.4 (characterisation of similarity). *For each pair $s, t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ we have*

$$\text{sim}(s, t) = \max \{d \in \mathbb{N} \cup \{\infty\} \mid s|d = t|d\}$$

Proof. Straightforward. \square

We can use this characterisation to show the first part of the compatibility of the metric and the partial order:

Lemma 4.5 (metric limit equals limit inferior). *Let $(t_\iota)_{\iota < \alpha}$ be a convergent sequence in $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$. Then $\lim_{\iota \rightarrow \alpha} t_\iota = \liminf_{\iota \rightarrow \alpha} t_\iota$.*

Proof. If α is a successor ordinal, this is trivial. Let α be a limit ordinal, $\hat{t} = \lim_{\iota \rightarrow \alpha} t_\iota$, and $\bar{t} = \liminf_{\iota \rightarrow \alpha} t_\iota$. Then for each $\varepsilon \in \mathbb{R}^+$ there is a $\beta < \alpha$ such that $\mathbf{d}(\hat{t}, t_\iota) < \varepsilon$ for all $\beta \leq \iota < \alpha$. Hence, for each $d \in \mathbb{N}$ there is a $\beta < \alpha$ such that $\text{sim}(\hat{t}, t_\iota) > d$ for all $\beta \leq \iota < \alpha$. According to Proposition 4.4, $\text{sim}(\hat{t}, t_\iota) > d$ implies $\hat{t}|d = t_\iota|d$, which, according to Proposition 4.3, implies $\hat{t}|d \leq_\perp t_\iota$. Therefore, $\hat{t}|d$ is a lower bound of $T_\beta = \{t_\iota \mid \beta \leq \iota < \alpha\}$, i.e. $\hat{t}|d \leq_\perp \prod T_\beta$. Since $\bar{t} = \bigsqcup_{\beta < \alpha} \prod T_\beta$, we also have that $\prod T_\beta \leq_\perp \bar{t}$. By transitivity, we obtain $\hat{t}|d \leq_\perp \bar{t}$ for each $d \in \mathbb{N}$. Since \hat{t} is total, we can thus conclude, according to Proposition 4.3, that $\hat{t} = \bar{t}$. \square

Before we continue, we want introduce another characterisation of similarity which bridges the gap to the partial order \leq_\perp . In order to follow this approach, we need to define the \perp -depth of a term $t \in \mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$. It is the minimal depth of an occurrence of the subterm \perp in t :

$$\perp\text{-depth}(t) = \min \{|\pi| \mid t(\pi) = \perp\} \cup \{\infty\}$$

Intuitively, the glb $s \sqcap t$ of two terms s, t represents the common structure that both terms share. The similarity $\text{sim}(s, t)$ is a much more condensed measure. It only provides the depth up to which the terms share a common structure. Using the \perp -depth we can directly condense the glb $s \sqcap t$ to the similarity $\text{sim}(s, t)$:

Proposition 4.6 (characterisation of similarity). *For each pair $s, t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ we have*

$$\text{sim}(s, t) = \perp\text{-depth}(s \sqcap t)$$

Proof. Follows from Lemma 3.10. \square

We can employ this alternative characterisation of similarity to show the second part of the compatibility of the metric and the partial order:

Lemma 4.7 (total limit inferior implies Cauchy). *Let $(t_\iota)_{\iota < \alpha}$ be a sequence in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ such that $\liminf_{\iota \rightarrow \alpha} t_\iota$ is total. Then $(t_\iota)_{\iota < \alpha}$ is Cauchy.*

Proof. For α a successor ordinal this is trivial. For the case that α is a limit ordinal, suppose that $(t_\iota)_{\iota < \alpha}$ is not Cauchy. That is, there is an $\varepsilon \in \mathbb{R}^+$ such that for all $\beta < \alpha$ there is a pair $\beta < \iota, \iota' < \alpha$ with $\mathbf{d}(t_\iota, t_{\iota'}) \geq \varepsilon$. Hence, there is a $d \in \mathbb{N}$ such that for all $\beta < \alpha$ there is a pair $\beta < \iota, \iota' < \alpha$ with $\mathbf{sim}(t_\iota, t_{\iota'}) \leq d$, which, according to Proposition 4.6, is equivalent to $\perp\text{-depth}(t_\iota \sqcap t_{\iota'}) \leq d$. That is,

$$\text{for each } \beta < \alpha \text{ there are } \beta < \iota, \iota' < \alpha \text{ with } \perp\text{-depth}(t_\iota \sqcap t_{\iota'}) \leq d \quad (1)$$

Let $s_\beta = \prod_{\beta \leq \iota < \alpha} t_\iota$. Then $s_\beta \leq_\perp t_\iota \sqcap t_{\iota'}$ for all $\beta \leq \iota, \iota' < \alpha$, which implies $\perp\text{-depth}(s_\beta) \leq \perp\text{-depth}(t_\iota \sqcap t_{\iota'})$. By combining this with (1), we obtain $\perp\text{-depth}(s_\beta) \leq d$. More precisely, we have that

$$\text{for each } \beta < \alpha \text{ there is a } \pi \in \mathcal{P}(s_\beta) \text{ with } |\pi| \leq d \text{ and } s_\beta(\pi) = \perp. \quad (2)$$

Let $\bar{t} = \liminf_{\iota \rightarrow \alpha} t_\iota$. Note that $\bar{t} = \bigsqcup_{\beta < \alpha} s_\beta$. Since, according to Lemma 3.9, $\mathcal{P}(\bar{t}) = \bigcup_{\beta < \alpha} \mathcal{P}(s_\beta)$ we can reformulate (2) as follows:

$$\text{for each } \beta < \alpha \text{ there is a } \pi \in \mathcal{P}(\bar{t}) \text{ with } |\pi| \leq d \text{ and } s_\beta(\pi) = \perp. \quad (2')$$

Since there are only finitely many positions in \bar{t} of length at most d , there is some $\pi^* \in \mathcal{P}(\bar{t})$ such that

$$\text{for each } \beta < \alpha \text{ there is a } \beta \leq \gamma < \alpha \text{ with } s_\gamma(\pi^*) = \perp. \quad (3)$$

Since $s_\beta \leq_\perp s_\gamma$, whenever $\beta \leq \gamma$, we can rewrite (3) as follows:

$$s_\beta(\pi^*) = \perp \text{ for all } \beta < \alpha \text{ with } \pi^* \in \mathcal{P}(s_\beta). \quad (3')$$

Since $\pi^* \in \mathcal{P}(\bar{t})$, we can employ Lemma 3.9 to obtain from (3') that $\bar{t}(\pi^*) = \perp$. This contradicts the assumption that $\bar{t} = \liminf_{\iota \rightarrow \alpha} t_\iota$ is total. \square

The following proposition combines Lemma 4.5 and Lemma 4.7 in order to obtain the desired property that the metric and the partial order are compatible:

Proposition 4.8 (partial order conservatively extends metric). *For every sequence $(t_\iota)_{\iota < \alpha}$ in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ the following holds:*

$$\liminf_{\iota \rightarrow \alpha} t_\iota = \lim_{\iota \rightarrow \alpha} t_\iota \quad \text{whenever} \quad \begin{array}{l} \lim_{\iota \rightarrow \alpha} t_\iota \text{ is defined, or} \\ \liminf_{\iota \rightarrow \alpha} t_\iota \text{ is a total term.} \end{array}$$

Proof. If $\lim_{\iota \rightarrow \alpha}$ is defined, the equality follows from Lemma 4.5. If $\liminf_{\iota \rightarrow \alpha} t_\iota$ is total, the sequence $(t_\iota)_{\iota < \alpha}$ is Cauchy by Lemma 4.7. Then, as the metric space $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$ is complete, $(t_\iota)_{\iota < \alpha}$ converges and we can apply Lemma 4.5 to conclude the equality. \square

4.2. p -Convergence vs. m -Convergence. In the previous section we have established that the metric and the partial order on (partial) terms are compatible in the sense that the corresponding notions of limit and limit inferior coincide whenever the limit is defined or the limit inferior is a total term. As weak m -convergence and weak p -convergence are solely based on the limit in the metric space resp. the limit inferior in the partially ordered set, we can directly apply this result to show that both notions of convergence coincide on total reductions:

Theorem 4.9 (total weak p -convergence = weak m -convergence). *For every reduction S in a TRS the following equivalences hold:*

- (i) $S: s \mathcal{P} \rightarrow \dots$ is total iff $S: s \mathcal{M} \rightarrow \dots$, and
- (ii) $S: s \mathcal{P} \rightarrow t$ is total iff $S: s \mathcal{M} \rightarrow t$.

Proof. Both equivalences follow directly from Proposition 4.8 and Fact 4.1, both of which are applicable as we presuppose that each term in the reduction is total. \square

In order to replicate Theorem 4.9 for the strong notions of convergence, we first need the following two lemmas that link the property of increasing contraction depth to volatile positions and the limit inferior, respectively:

Lemma 4.10 (strong m -convergence). *Let $S = (t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \lambda}$ be an open reduction. Then $(|\pi_\iota|)_{\iota < \lambda}$ tends to infinity iff, for each position π , there is an ordinal $\alpha < \lambda$ such that $\pi_\iota \neq \pi$ for all $\alpha \leq \iota < \lambda$.*

Proof. The “only if” direction is trivial. For the converse direction, suppose that $|\pi_\iota|$ does not tend to infinity as ι approaches λ . That is, there is some depth $d \in \mathbb{N}$ such that there is no upper bound on the indices of reduction steps taking place at depth d . Let d^* be the minimal such depth. That is, there is some $\alpha < \lambda$ such that all reduction steps in $S|_{[\alpha, \lambda)}$ are at depth at least d^* , i.e. $|\pi_\iota| \geq d^*$ holds for all $\alpha \leq \iota < \lambda$. Of course, also in $S|_{[\alpha, \lambda)}$ the indices of steps at depth d^* are not bounded from above. As all reduction steps in $S|_{[\alpha, \lambda)}$ take place at depth d^* or below, $t_\iota|d^* = t_{\iota'}|d^*$ holds for all $\alpha \leq \iota, \iota' < \lambda$. That is, all terms in $S|_{[\alpha, \lambda)}$ have the same set of positions of length d^* . Let $P^* = \{\pi \in \mathcal{P}(t_n) \mid |\pi| = d^*\}$ be this set. Since there is no upper bound on the indices of steps in $S|_{[\alpha, \lambda)}$ taking place at a position in P^* , yet, P^* is finite, there has to be some position $\pi^* \in P^*$ for which there is also no such upper bound. This contradicts the assumption that there is always such an upper bound. \square

Lemma 4.11 (limit inferior of truncations). *Let $(t_\iota)_{\iota < \lambda}$ be a sequence in $\mathcal{T}^\infty(\Sigma_\perp, \mathcal{V})$ and $(d_\iota)_{\iota < \lambda}$ a sequence in \mathbb{N} such that λ is a limit ordinal and $(d_\iota)_{\iota < \lambda}$ tends to infinity. Then $\liminf_{\iota \rightarrow \lambda} t_\iota = \liminf_{\iota \rightarrow \lambda} t_\iota|d_\iota$.*

Proof. Let $\bar{t} = \liminf_{\iota \rightarrow \lambda} t_\iota|d_\iota$ and $\hat{t} = \liminf_{\iota \rightarrow \lambda} t_\iota$. Since, according to Proposition 4.3, $t_\iota|d_\iota \leq_\perp t_\iota$ for each $\iota < \lambda$, we have that $\bar{t} \leq_\perp \hat{t}$. Thus, it remains to be shown that also $\hat{t} \leq_\perp \bar{t}$ holds. That is, we have to show that $\hat{t}(\pi) = \bar{t}(\pi)$ holds for all $\pi \in \mathcal{P}_\perp(\hat{t})$.

Let $\pi \in \mathcal{P}_\perp(\hat{t})$. That is, $\hat{t}(\pi) = f \neq \perp$. Hence, by Lemma 3.9, there is some $\alpha < \lambda$ with $(\prod_{\alpha \leq \iota < \lambda} t_\iota)(\pi) = f$. Let $P = \{\pi' \mid \pi' \leq \pi\}$ be the set of all prefixes of π . Note that $\prod_{\alpha \leq \iota < \lambda} t_\iota \leq_\perp t_\gamma$ for all $\alpha \leq \gamma < \lambda$. Hence, $\prod_{\alpha \leq \iota < \lambda} t_\iota$ and t_γ coincide in all occurrences in P for all $\alpha \leq \gamma < \lambda$. Because $(d_\iota)_{\iota < \lambda}$ tends to infinity, there is some $\alpha \leq \beta < \lambda$ such that $d_\gamma > |\pi|$ for all $\beta \leq \gamma < \lambda$. Consequently, since $t_\gamma|d_\gamma$ and t_γ coincide in all occurrences of length smaller than d_γ for all $\gamma < \lambda$, we have that $t_\gamma|d_\gamma$ and t_γ coincide in all occurrences in P for all $\beta \leq \gamma < \lambda$. Hence, $t_\gamma|d_\gamma$ and $\prod_{\alpha \leq \iota < \lambda} t_\iota$ coincide in all occurrences in P for all $\beta \leq \gamma < \lambda$. Hence, according to Lemma 3.10, $\prod_{\alpha \leq \iota < \lambda} t_\iota$ and $\prod_{\beta \leq \iota < \lambda} t_\iota|d_\iota$ coincide in all occurrences in P . Particularly, it holds that $(\prod_{\beta \leq \iota < \lambda} t_\iota|d_\iota)(\pi) = f$ which in turn implies by Lemma 3.9 that $\bar{t}(\pi) = f$. \square

We now can prove the counterpart of Theorem 4.9 for strong convergences:

Theorem 4.12 (total strong p -convergence = strong m -convergence). *For every reduction S in a TRS the following equivalences hold:*

- (i) $S: s \mathcal{P} \rightarrow \dots$ is total iff $S: s \mathcal{M} \rightarrow \dots$, and

(ii) $S: s \twoheadrightarrow t$ is total iff $S: s \multimap t$.

Proof. It suffices to only prove ((ii)) since ((i)) follows from ((ii)) according to Remark 3.4 resp. Remark 2.3.

Let $S = (\varphi_\iota: t_\iota \rightarrow_{\pi_\iota, c_\iota} t_{\iota+1})_{\iota < \alpha}$ be a reduction in a TRS \mathcal{R}_\perp . We continue the proof by induction on α . The case $\alpha = 0$ is trivial. If α is a successor ordinal $\beta + 1$, we can reason as follows

$$\begin{aligned} S: t_0 \twoheadrightarrow t_\alpha \text{ total} &\text{ iff } S|_\beta: t_0 \twoheadrightarrow t_\beta \text{ and } t_\beta \rightarrow_{\mathcal{R}} t_\alpha && \text{(Remark 3.4, Fact 4.1)} \\ &\text{ iff } S|_\beta: t_0 \multimap t_\beta \text{ and } t_\beta \rightarrow_{\mathcal{R}} t_\alpha && \text{(ind. hyp.)} \\ &\text{ iff } S: t_0 \multimap t_\alpha && \text{(Remark 2.3)} \end{aligned}$$

Let α be a limit ordinal. At first consider the “only if” direction. That is, we assume that $S: t_0 \twoheadrightarrow t_\alpha$ is total. According to Remark 3.4, we have that $S|_\beta: t_0 \twoheadrightarrow t_\beta$ for each $\beta < \alpha$. Applying the induction hypothesis yields $S|_\beta: t_0 \multimap t_\beta$ for each $\beta < \alpha$. That is, following Remark 2.3, we have $S: t_0 \multimap \dots$. Since $c_\iota \leq_\perp t_\iota$ for all $\iota < \alpha$, we have that $t_\alpha = \liminf_{\iota \rightarrow \alpha} c_\iota \leq_\perp \liminf_{\iota \rightarrow \alpha} t_\iota$. Because t_α is total and, therefore, maximal w.r.t. \leq_\perp , we can conclude that $t_\alpha = \liminf_{\iota \rightarrow \alpha} t_\iota$. According to Proposition 4.8, this also means that $t_\alpha = \lim_{\iota \rightarrow \alpha} t_\iota$. For strong m -convergence it remains to be shown that $(|\pi_\iota|)_{\iota < \alpha}$ tends to infinity. So let us assume that this is not the case. By Lemma 4.10, this means that there is a position π such that, for each $\beta < \alpha$, there is some $\beta \leq \gamma < \alpha$ such that the step φ_γ takes place at position π . By Lemma 3.14, this contradicts the fact that t_α is a total term.

Now consider the converse direction and assume that $S: t_0 \multimap t_\alpha$. Following Remark 2.3 we obtain $S|_\beta: t_0 \multimap t_\beta$ for all $\beta < \alpha$, to which we can apply the induction hypothesis in order to get $S|_\beta: t_0 \twoheadrightarrow t_\beta$ for all $\beta < \alpha$ so that we have $S: t_0 \twoheadrightarrow \dots$, according to Remark 3.4. It remains to be shown that $t_\alpha = \liminf_{\iota \rightarrow \alpha} c_\iota$. Since S strongly m -converges to t_α , we have that

- (a) $t_\alpha = \lim_{\iota \rightarrow \alpha} t_\iota$, and that
- (b) the sequence of depths $(d_\iota = |\pi_\iota|)_{\iota < \alpha}$ tends to infinity.

Using Proposition 4.8 we can deduce from ((a)) that $t_\alpha = \liminf_{\iota \rightarrow \alpha} t_\iota$. Due to ((b)), we can apply Lemma 4.11 to obtain

$$\liminf_{\iota \rightarrow \alpha} t_\iota = \liminf_{\iota \rightarrow \alpha} t_\iota | d_\iota \quad \text{and} \quad \liminf_{\iota \rightarrow \alpha} c_\iota = \liminf_{\iota \rightarrow \alpha} c_\iota | d_\iota.$$

Since $t_\iota | d_\iota = c_\iota | d_\iota$ for all $\iota < \alpha$, we can conclude that

$$t_\alpha = \liminf_{\iota \rightarrow \alpha} t_\iota = \liminf_{\iota \rightarrow \alpha} t_\iota | d_\iota = \liminf_{\iota \rightarrow \alpha} c_\iota | d_\iota = \liminf_{\iota \rightarrow \alpha} c_\iota. \quad \square$$

The main result of this section is that we do not lose anything when switching from the metric model to the partial order model of infinitary term rewriting. Restricted to the domain of the metric model, i.e. total terms, both models coincide in the strongest possible sense as Theorem 4.9 and Theorem 4.12 confirm.

At the same time, however, the partial order model provides more structure. Whenever the metric model can only conclude divergence, the partial order model can qualify the degree of divergence. If a reduction p -converges to \perp , it can be considered completely divergent. If it p -converges to a term that only contains \perp as proper subterms, it can be recognised as being only partially divergent with the diverging parts of the reduction indicated by ‘ \perp ’s, whereas complete absence of ‘ \perp ’s then indicates complete convergence.

In the rest of this paper we will put our focus on strong convergence. Theorem 4.12 will be one of the central tools in Section 6 where we shall discover that Böhm-reachability

coincides with strong p -reachability in orthogonal systems. The other crucial tool that we will leverage is the existence and uniqueness of complete developments. This is the subject of the subsequent section.

5. STRONGLY p -CONVERGING COMPLETE DEVELOPMENTS

The purpose of this section is to establish a theory of residuals and complete developments in the setting of strongly p -convergent reductions. 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.

Complete developments are a well-known tool for proving (finitary) confluence of orthogonal systems [24]. It has also been lifted to the setting of strongly m -convergent reductions in order to establish (restricted forms of) infinitary confluence of orthogonal systems [16]. As we have seen in Example 2.6, m -convergence in general does not have this property.

After introducing residuals and complete developments in Section 5.1, we will show in Section 5.2 resp. Section 5.3 that complete developments do always exist and that their final terms are uniquely determined. We then use this in Section 5.4 to show the Infinitary Strip Lemma for strongly p -converging reductions which is a crucial tool for proving our main result in Section 6.

5.1. Residuals. At first we need to formalise the notion of residuals. It is virtually equivalent to the definition for strongly m -convergent reduction by Kennaway et al. [16]:

Definition 5.1 (descendants, residuals). Let \mathcal{R} be a TRS, $S: t_0 \xrightarrow{\mathcal{R}} t_\alpha$, 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. $S: t_0 \xrightarrow{\pi, \rho} t_1$ 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 $\alpha = \beta + 1$, then $U//S = (U//S|_\beta)//S|_{[\beta, \alpha]}$
- (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.

Remark 5.2. One can easily see that the descendants of a set of non- \perp -occurrences is again a set of non- \perp -occurrences. The restriction to non- \perp -occurrences has to be made explicit for the case of open reductions. In fact, without this explicit restriction the definition would yield descendants which might not even be occurrences in the final term t_α of the reduction.

For example, consider the system with the single rule $f(x) \rightarrow x$ and the strongly p -convergent reduction

$$S: f^\omega \rightarrow f^\omega \rightarrow \dots \perp$$

in which each reduction step contracts the redex at the root of f^ω . Consider the set $U = \{\langle \rangle, \langle 0 \rangle, \langle 0, 0 \rangle, \langle 0, 0, 0 \rangle, \dots\}$ of all positions in t^ω . Without the abovementioned restriction, the descendants of U by S would be U itself as the descendants of U by each proper prefix of S is also U . However, none of the positions $\langle 0 \rangle, \langle 0, 0 \rangle, \langle 0, 0, 0 \rangle, \dots \in U$ is even a position in the final term \perp . The position $\langle \rangle \in U$ occurs in \perp , but only as a \perp -occurrence. With the restriction to non- \perp -occurrences we indeed get the expected result $U//S = \emptyset$.

The definition of descendants of open reductions is quite subtle which makes it fairly cumbersome to use in proofs. The lemma below establishes an alternative characterisation which will turn out to be useful in later proofs:

Lemma 5.3 (descendants of open reductions). *Let \mathcal{R} be a TRS, $S: s \xrightarrow{\mathcal{R}}^\lambda t$ and $U \subseteq \mathcal{P}_\perp(s)$, with λ a limit ordinal and $S = (t_\iota \rightarrow_{\pi_\iota, c_\iota} t_{\iota+1})_{\iota < \lambda}$. Then it holds that for each position π*

$$\pi \in U//S \quad \text{iff} \quad \text{there is some } \beta < \lambda \text{ with } \pi \in U//S|_\beta \text{ and } \forall \beta \leq \iota < \lambda \quad \pi_\iota \not\leq \pi.$$

Proof. We first prove the “only if” direction. To this end, assume that $\pi \in U//S$. Hence, it holds that

$$\pi \in \mathcal{P}_\perp(t) \text{ and there is some } \gamma_1 < \lambda \text{ such that } \pi \in U//S|_\iota \text{ for all } \gamma_1 \leq \iota < \lambda \quad (1)$$

Particularly, we have that $t(\pi) \neq \perp$. Applying Lemma 3.11 then yields that

$$\text{there is some } \gamma_2 < \lambda \text{ such that } \pi_\iota \not\leq \pi \text{ for all } \gamma_2 \leq \iota < \lambda \quad (2)$$

Now take $\beta = \max\{\gamma_1, \gamma_2\}$. Then it holds that $\pi \in U//S|_\beta$ and that $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \lambda$ due to (1) and (2), respectively.

Next, consider the converse direction of the statement: Let $\beta < \lambda$ be such that $\pi \in U//S|_\beta$ and $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \lambda$. We will show that $\pi \in U//S$ by proving the stronger statement that $\pi \in U//S|_\gamma$ for all $\beta \leq \gamma \leq \lambda$. We do this by induction on γ .

For $\gamma = \beta$, this is trivial. Let $\gamma = \gamma' + 1 > \beta$. Note that, by definition, $U//S|_\gamma = (U//S|_{\gamma'})//S|_{[\gamma', \gamma)}$. Hence, since for the γ' -th step we have, by assumption, $\pi_{\gamma'} \not\leq \pi$ and for the preceding reduction we have, by induction hypothesis, that $\pi \in U//S|_{\gamma'}$, we can conclude that $\pi \in U//S|_\gamma$.

Let $\gamma > \beta$ be a limit ordinal. By induction hypothesis, we have that $\pi \in U//S|_\iota$ for each $\beta \leq \iota < \gamma$. Particularly, this implies that $\pi \in \mathcal{P}_\perp(t_\beta)$. Together with the assumption that $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \gamma$, this yields that $\pi \in \mathcal{P}_\perp(t_\gamma)$ according to Lemma 3.11. Hence, $\pi \in U//S|_\gamma$. \square

The following lemma confirms the expected monotonicity of descendants:

Lemma 5.4 (monotonicity of descendants). *Let \mathcal{R} be a TRS, $S: s \xrightarrow{\mathcal{R}} t$ and $U, V \subseteq \mathcal{P}_\perp(s)$. If $U \subseteq V$, then $U//S \subseteq V//S$.*

Proof. Straightforward induction on the length of S . \square

This lemma can be generalised such that we can see that descendants are defined “point-wise”:

Proposition 5.5 (pointwise definition of descendants). *Let \mathcal{R} be a TRS, $S: s \xrightarrow{\mathcal{R}} t$ and $U \subseteq \mathcal{P}_{\chi}(s)$. Then it holds that $U//S = \bigcup_{u \in U} u//S$.*

Proof. Let $S = (t_{\iota} \rightarrow_{\pi_{\iota}, c_{\iota}} t_{\iota+1})_{\iota < \alpha}$. For $\alpha = 0$ and $\alpha = 1$, the statement is trivially true. If $\alpha = \alpha' + 1 > 1$, then abbreviate $S|_{\alpha'}$ and $S|_{[\alpha', \alpha)}$ by S_1 and S_2 , respectively, and reason as follows:

$$\begin{aligned} U//S &= (U//S_1)//S_2 \stackrel{IH}{=} \underbrace{\left(\bigcup_{u \in U} \overbrace{u//S_1}^{V_u} \right)}_V //S_2 \stackrel{IH}{=} \bigcup_{u \in V} u//S_2 \\ &= \bigcup_{u \in U} \bigcup_{v \in V_u} v//S_2 \stackrel{IH}{=} \bigcup_{u \in U} V_u//S_2 = \bigcup_{u \in U} (u//S_1)//S_2 = \bigcup_{u \in U} u//S \end{aligned}$$

Let α be a limit ordinal. The “ \supseteq ” direction of the equation follows from Lemma 5.4. For the converse direction, assume that $\pi \in U//S$. By Lemma 5.3, there is some $\beta < \alpha$ such that $\pi_{\iota} \not\leq \pi$ for all $\beta \leq \iota < \alpha$ and $\pi \in U//S|_{\beta}$. Applying the induction hypothesis yields that $\pi \in \bigcup_{u \in U} u//S|_{\beta}$, i.e. there is some $u^* \in U$ such that $\pi \in u^*//S|_{\beta}$. By employing Lemma 5.3 again, we can conclude that $\pi \in u^*//S$ and, therefore, that $\pi \in \bigcup_{u \in U} u//S$. \square

Note that the above proposition fails if we would include \perp -occurrences in our definition of descendants: Reconsider the example in Remark 5.2 and assume we would drop the restriction to non- \perp -occurrences. Then the residuals $u//S$ of each occurrence $u \in U$ would be empty, whereas the residuals $U//S$ of all occurrences would be the root occurrence $\langle \rangle$.

Proposition 5.6 (uniqueness of ancestors). *Let \mathcal{R} be TRS, $S: s \xrightarrow{\mathcal{R}} t$ and $U, V \subseteq \mathcal{P}_{\chi}(s)$. If $U \cap V = \emptyset$, then $U//S \cap V//S = \emptyset$.*

Proof. We will prove the contraposition of the statement. To this end, suppose that there is some occurrence $w \in U//S \cap V//S$. By Proposition 5.5, there are occurrences $u \in U$ and $v \in V$ such that $w \in u//S \cap v//S$. We will show by induction on the length of S that then $u = v$ and, therefore, $U \cap V \neq \emptyset$. If S is empty, then this is trivial. If S is of successor ordinal length or open, then $u = v$ follows from the induction hypothesis. \square

Remark 5.7. The two propositions above imply that each descendant $u' \in U//S$ of a set U of occurrences is the descendant of a uniquely determined occurrence $u \in U$, i.e. $u' \in u//S$ for exactly one $u \in U$. This occurrence u is also called the *ancestor* of u' by S .

The following proposition confirms a property of descendants that one expects intuitively: The descendants of descendants are again descendants. That is, the concept of descendants is composable.

Proposition 5.8 (descendants of sequential reductions). *Let \mathcal{R} be a TRS, $S: t_0 \xrightarrow{\mathcal{R}} t_1$, $T: t_1 \xrightarrow{\mathcal{R}} t_2$, and $U \subseteq \mathcal{P}_{\chi}(t_0)$. Then $U//S \cdot T = (U//S)//T$.*

Proof. Straightforward proof by induction on the length of T . \square

The following proposition confirms that the disjointness of occurrences is propagated through their descendants:

Proposition 5.9 (disjoint descendants). *The descendants of a set of pairwise disjoint occurrences are pairwise disjoint as well.*

Proof. Let $S: s \xrightarrow{\alpha} t$ and let U be a set of pairwise disjoint occurrences in s . We show that $U//S$ is also a set of pairwise disjoint occurrences by induction on α .

For α being 0, the statement is trivial, and, for α being a successor ordinal, the statement follows straightforwardly from the induction hypothesis. Let α be limit ordinal and suppose that there are two occurrences $u, v \in U//S$ which are not disjoint. By definition, there are ordinals $\beta_1, \beta_2 < \alpha$ such that $u \in U//S|_{\beta_1}$ for all $\beta_1 \leq \iota < \alpha$, and $v \in U//S|_{\beta_2}$ for all $\beta_2 \leq \iota < \alpha$. Let $\beta = \max\{\beta_1, \beta_2\}$. Then we have that $u, v \in U//S|_{\beta}$. This, however, contradicts the induction hypothesis which, in particular, states that $U//S|_{\beta}$ is a set of pairwise disjoint occurrences. \square

For the definition of complete developments it is important that the descendants of redex occurrences are again redex occurrences:

Proposition 5.10 (residuals). *Let \mathcal{R} be an orthogonal TRS, $S: s \xrightarrow{\mathcal{R}} t$ and U a set of redex occurrences in s . Then $U//S$ is a set of redex occurrences in t .*

Proof. Let $S = (t_i \rightarrow_{\pi_i, c_i} t_{i+1})_{i < \alpha}$. We proceed by induction on α . For α being 0, the statement is trivial, and, for α a successor ordinal, the statement follows straightforwardly from the induction hypothesis.

So assume that α is a limit ordinal and that $\pi \in U//S$. We will show that $t|_{\pi}$ is a redex. From Lemma 5.3 we obtain that

$$\text{there is some } \beta < \alpha \text{ with } \pi \in U//S|_{\beta} \text{ and } \pi_{\iota} \not\leq \pi \text{ for all } \beta \leq \iota < \alpha. \quad (1)$$

By applying the induction hypothesis, we get that π is a redex occurrence in t_{β} . Hence, there is some rule $l \rightarrow r \in R$ such that $t_{\beta}|_{\pi}$ is an instance of l .

We continue this proof by showing the following stronger claim:

$$\text{for all } \beta \leq \gamma \leq \alpha \quad t_{\gamma}|_{\pi} \text{ is an instance of } l, \text{ and} \quad (2)$$

$$c_{\iota}|_{\pi} \text{ is an instance of } l \text{ for all } \beta \leq \iota < \gamma \quad (3)$$

For the special case $\gamma = \alpha$ the above claim (2) implies that $t|_{\pi}$ is a redex.

We proceed by an induction on γ . For $\gamma = \beta$, part (2) of the claim has already been shown and (3) is vacuously true. Let $\gamma = \gamma' + 1 > \beta$. According to the induction hypothesis, (2) and (3) hold for γ' . Hence, it remains to be shown that both $t_{\gamma}|_{\pi}$ and $c_{\gamma'}|_{\pi}$ are instances of l . At first consider $c_{\gamma'}|_{\pi}$. Recall that $c_{\gamma'} = t_{\gamma'}[\perp]_{\pi_{\gamma'}}$. At first consider the case where π and $\pi_{\gamma'}$ are disjoint. Then $c_{\gamma'}|_{\pi} = t_{\gamma'}|_{\pi}$. Since, by induction hypothesis, $t_{\gamma'}|_{\pi}$ is an instance of l , so is $c_{\gamma'}|_{\pi}$. Next, consider the case where π and $\pi_{\gamma'}$ are not disjoint. Because of (1), we then have that $\pi < \pi_{\gamma'}$, i.e. there is some non-empty π' with $\pi_{\gamma'} = \pi \cdot \pi'$. Since \mathcal{R} is non-overlapping, π' cannot be a position in the pattern of the redex $t_{\gamma'}|_{\pi}$ w.r.t. l . Therefore, also $c_{\gamma'}|_{\pi}$ is an instance of l . So in either case $c_{\gamma'}|_{\pi}$ is an instance of l . Since $c_{\gamma'} \leq_{\perp} t_{\gamma}$, also $t_{\gamma}|_{\pi}$ is an instance of l .

Let $\gamma > \beta$ be a limit ordinal. Part (3) of the claim follows immediately from the induction hypothesis. Hence, $c_{\iota}|_{\pi}$ is an instance of l for all $\beta \leq \iota < \gamma$. This and (1) implies that all terms in the set $T = \{c_{\iota} | \beta \leq \iota < \gamma\}$ coincide in all occurrences in the set

$$P = \{\pi' | \pi' \leq \pi\} \cup \{\pi \cdot \pi' | \pi' \in \mathcal{P}_{\Sigma}(l)\}$$

P is obviously closed under prefixes. Therefore, we can apply Lemma 3.10 in order to obtain that $\prod T$ coincides with all terms in T in all occurrences in P . Since $\prod T \leq_{\perp} t_{\gamma}$, this property carries over to t_{γ} . Consequently, also $t_{\gamma}|_{\pi}$ is an instance of l . \square

Next we want to establish an alternative characterisation of descendants based on labellings. This is a well-known technique [24] that keeps track of descendants by labelling the symbols at the relevant positions in the initial term. In order to formalise this idea, we need to extend a given TRS such that it can also deal with terms that contain labelled symbols:

Definition 5.11 (labelled TRSs/terms). Let $\mathcal{R} = (\Sigma, R)$ be a TRS.

- (i) The *labelled signature* Σ^{\perp} is defined as $\Sigma \cup \{f^{\perp} \mid f \in \Sigma\}$. The arity of the function symbol f^{\perp} is the same as that of f . The symbols f^{\perp} are called *labelled*; the symbols $f \in \Sigma$ are called *unlabelled*. Terms over Σ^{\perp} are called *labelled terms*. Note that the symbol $\perp \in \Sigma_{\perp}$ has no corresponding labelled symbol \perp^{\perp} in the labelled signature Σ^{\perp} . Likewise, there are no labelled variables.
- (ii) Labelled terms can be projected back to the original unlabelled ones by removing the labels via the projection function $\|\cdot\| : \mathcal{T}^{\infty}(\Sigma^{\perp}, \mathcal{V}) \rightarrow \mathcal{T}^{\infty}(\Sigma, \mathcal{V})$:

$$\|\perp\| = \perp \qquad \|x\| = x \qquad \text{for all } x \in \mathcal{V}, \text{ and}$$

$$\|f^{\perp}(t_1, \dots, t_k)\| = \|f(t_1, \dots, t_k)\| = f(\|t_1\|, \dots, \|t_k\|) \qquad \text{for all } f \in \Sigma^{(k)}$$

- (iii) The *labelled TRS* \mathcal{R}^{\perp} is defined as $(\Sigma^{\perp}, R^{\perp})$, where $R^{\perp} = \{l \rightarrow r \mid \|l\| \rightarrow r \in R\}$.
- (iv) For each rule $l \rightarrow r \in R^{\perp}$, we define its unlabelled original $\|l\| \rightarrow r \in R$.
- (v) Let $t \in \mathcal{T}^{\infty}(\Sigma, \mathcal{V})$ and $U \subseteq \mathcal{P}_{\Sigma}(t)$. The term $t^{(U)} \in \mathcal{T}^{\infty}(\Sigma^{\perp}, \mathcal{V})$ is defined by

$$t^{(U)}(\pi) = \begin{cases} t(\pi) & \text{if } \pi \notin U \\ t(\pi)^{\perp} & \text{if } \pi \in U \end{cases}$$

That is, $\|t^{(U)}\| = t$ and the labelled symbols in $t^{(U)}$ are exactly those at positions in U .

The key property which is needed in order to make the labelling approach work is that any reduction in a left-linear TRS that starts in some term t can be lifted for any labelling t' of t to a unique equivalent reduction in the corresponding labelled TRS that starts in t' :

Proposition 5.12 (lifting reductions to labelled TRSs). *Let $\mathcal{R} = (\Sigma, R)$ be a left-linear TRS, $S = (s_{\iota} \rightarrow_{\rho_{\iota}, \pi_{\iota}} s_{\iota+1})_{\iota < \alpha}$ a reduction strongly p -converging to s_{α} in \mathcal{R} , and $t_0 \in \mathcal{T}^{\infty}(\Sigma^{\perp}, \mathcal{V})$ a labelled term with $\|t_0\| = s_0$. Then there is a unique reduction $T = (t_{\iota} \rightarrow_{\rho'_{\iota}, \pi_{\iota}} t_{\iota+1})_{\iota < \alpha}$ strongly p -converging to t_{α} in \mathcal{R}^{\perp} such that*

- (a) $\|t_{\iota}\| = s_{\iota}$, $\|\rho'_{\iota}\| = \rho_{\iota}$, for all $\iota < \alpha$, and
- (b) $\|t_{\alpha}\| = s_{\alpha}$.

Proof. We prove this by an induction on α . For the case of α being zero, the statement is trivially true. For the case of α being a successor ordinal, the statement follows straightforwardly from the induction hypothesis (the argument is the same as for finite reductions; e.g. consult [24]).

Let α be a limit ordinal. By induction hypothesis, for each proper prefix $S|_{\gamma}$ of S there is a uniquely defined strongly p -convergent reduction T_{γ} in \mathcal{R}^{\perp} satisfying ((a)) and ((b)). Since the sequence $(S|_{\iota})_{\iota < \alpha}$ forms a chain w.r.t. the prefix order \leq , so does the corresponding sequence $(T_{\iota})_{\iota < \alpha}$. Hence the sequence $T = \bigsqcup_{\iota < \alpha} T_{\iota}$ is well-defined. By construction, $T_{\gamma} \leq T$ holds for each $\gamma < \alpha$, and we can use the induction hypothesis to obtain part ((a)) of the

proposition. In order to show $s_\alpha = \|t_\alpha\|$, we prove the two inequalities $s_\alpha \leq_\perp \|t_\alpha\|$ and $s_\alpha \geq_\perp \|t_\alpha\|$:

To show $\|t_\alpha\| \leq_\perp s_\alpha$, we take some $\pi \in \mathcal{P}_\perp(\|t_\alpha\|)$ and show that $\|t_\alpha\|(\pi) = s_\alpha(\pi)$. Let $f = \|t_\alpha\|(\pi)$. That is, either $t_\alpha(\pi) = f$ or $t_\alpha(\pi) = f^\perp$. In either case, we can employ Lemma 3.11 to obtain some $\beta < \alpha$ such that $t_\beta(\pi) = f$ resp. $t_\beta(\pi) = f^\perp$ and $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \alpha$. Since, by ((a)), $s_\beta = \|t_\beta\|$, we have in both cases that $s_\beta(\pi) = f$. By applying Lemma 3.11 again, we get that $s_\alpha(\pi) = f$, too.

Lastly, we show the converse inequality $s_\alpha \leq_\perp \|t_\alpha\|$. For this purpose, let $\pi \in \mathcal{P}_\perp(s_\alpha)$ and $f = s_\alpha(\pi)$. By Lemma 3.11, there is some $\beta < \alpha$ such that $s_\beta(\pi) = f$ and $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \alpha$. Since, by ((a)), $s_\beta = \|t_\beta\|$, we have that $t_\beta(\pi) \in \{f, f^\perp\}$. Applying Lemma 3.11 again then yields that $t_\alpha(\pi) \in \{f, f^\perp\}$ and, therefore, $\|t_\alpha\|(\pi) = f$. \square

Having this, we can establish an alternative characterisation of descendants using labellings:

Proposition 5.13 (alternative characterisation of descendants). *Let \mathcal{R} be a left-linear TRS, $S: s_0 \xrightarrow{\mathcal{R}} s_\alpha$, and $U \subseteq \mathcal{P}_\perp(s_0)$. Following Proposition 5.12, let $T: t_0 \xrightarrow{\mathcal{R}} t_\alpha$ be the unique lifting of S to \mathcal{R}^\perp starting with the term $t_0 = s_0^{(U)}$. Then it holds that $t_\alpha = s_\alpha^{(U//S)}$. That is, for all $\pi \in \mathcal{P}_\perp(s_\alpha)$, it holds that $t_\alpha(\pi)$ is labelled iff $\pi \in U//S$.*

Proof. Let $S = (s_\iota \rightarrow_{\pi_\iota} s_{\iota+1})_{\iota < \alpha}$ and $T = (t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \alpha}$. We prove the statement by an induction on the length α of S . If $\alpha = 0$, then the statement is trivially true. If α is a successor ordinal, then a straightforward argument shows that the statement follows from the induction hypothesis. Here the restriction to left-linear systems is vital.

Let α be a limit ordinal and let $\pi \in \mathcal{P}_\perp(s_\alpha)$. We can then reason as follows:

$$\begin{aligned} t_\alpha(\pi) \text{ is labelled} & \text{ iff } \exists \beta < \alpha: t_\beta(\pi) \text{ is labelled and } \forall \beta \leq \iota < \alpha: \pi_\iota \not\leq \pi & (\text{Lem. 3.11}) \\ & \text{ iff } \pi \in U//S|_\beta \text{ and } \forall \beta \leq \iota < \alpha: \pi_\iota \not\leq \pi & (\text{ind. hyp.}) \\ & \text{ iff } \pi \in U//S & (\text{Lem. 5.3}) \end{aligned}$$

\square

5.2. Constructing Complete Developments. Complete developments are usually defined for (almost) orthogonal systems. This ensures that the residuals of redexes are again redexes. Since we are going to use complete developments for potentially overlapping systems as well, we need to make restrictions on the set of redex occurrences instead:

Definition 5.14 (conflicting redex occurrences). Two distinct redex occurrences u, v in a term t are called *conflicting* if there is a position π such that $v = u \cdot \pi$ and π is a pattern position of the redex at u , or, vice versa, $u = v \cdot \pi$ and π is a pattern position of the redex at v . If this is not the case, then u and v are called *non-conflicting*.

One can easily see that in an orthogonal TRS any pair of redex occurrences is non-conflicting.

Definition 5.15 ((complete) development). Let \mathcal{R} be a left-linear TRS, s a partial term in \mathcal{R} , and U a set of pairwise non-conflicting redex occurrences in s .

- (i) A *development* of U in s is a strongly p -converging reduction $S: s \xrightarrow{\mathcal{R}} t$ in which each reduction step $\varphi_\iota: t_\iota \rightarrow_{\pi_\iota} t_{\iota+1}$ contracts a redex at $\pi_\iota \in U//S|_\iota$.
- (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 [16] if restricted to total terms.

The restriction to non-conflicting redex occurrences is essential in order guarantee that the redex occurrences are independent from each other:

Proposition 5.16 (non-conflicting residuals). *Let \mathcal{R} be a left-linear TRS, s a partial term in \mathcal{R} , U a set of pairwise non-conflicting redex occurrences in s , and $S: s \rightarrow_U t$ a development of U in s . Then also $U//S$ is a set of pairwise non-conflicting redex occurrences.*

Proof. This can be proved by induction on the length of S . The part showing that the descendants are again redex occurrences can be copied almost verbatim from Proposition 5.10. Instead of referring to the non-overlappingness of the system one can refer to the non-conflictingness of the preceding residuals which can be assumed by the induction hypothesis. The part of the induction proof that shows non-conflictingness is analogous to Proposition 5.9. \square

It is relatively easy to show that complete developments of sets of non-conflicting redex occurrences do always exist in the partial order setting. The reason for this is that strongly p -continuous reductions do always strongly p -converge as well. This means that as long as there are (residuals of) redex occurrences left after an incomplete development, one can extend this development arbitrarily by contracting some of the remaining redex occurrences. The only thing that remains to be shown is that one can devise a reduction strategy which eventually contracts (all residuals of) all redexes. The proposition below shows that a parallel-outermost reduction strategy will always yield a complete development in a left-linear system.

Proposition 5.17 (complete developments). *Let \mathcal{R} be a left-linear TRS, t a partial term in \mathcal{R} , and U a set of pairwise non-conflicting redex occurrences in t . Then U has a complete development in t .*


Proof. Let $t_0 = t$, $U_0 = U$ and V_0 the set of outermost occurrences in U_0 . Furthermore, let $S_0: t_0 \xrightarrow{V_0} t_1$ be some complete development of V_0 in t_0 . S_0 can be constructed by contracting the redex occurrences in V_0 in a left-to-right order. This step can be continued for each $i < \omega$ by taking $U_{i+1} = U_i//S_i$, where $S_i: t_i \xrightarrow{V_i} t_{i+1}$ is some complete development of V_i in t_i with V_i the set of outermost redex occurrences in U_i .

Note that then, by iterating Proposition 5.8, it holds that


$$U//S_0 \cdot \dots \cdot S_{n-1} = U_n \quad \text{for all } n < \omega \quad (1)$$

If there is some $n < \omega$ for which $U_n = \emptyset$, then $S_0 \cdot \dots \cdot S_{n-1}$ is a complete development of U according to (1).

If this is not the case, consider the reduction $S = \prod_{i < \omega} S_i$, i.e. the concatenation of all ' S_i 's. We claim that S is a complete development of U . Suppose that this is not the case, i.e. $U//S \neq \emptyset$. Hence, there is some $u \in U//S$. Since all ' U_i 's are non-empty, so are the ' V_i 's. Consequently, all ' S_i 's are non-empty reductions which implies that S is a reduction of limit ordinal length, say λ . Therefore, we can apply Lemma 5.3 to infer from $u \in U//S$ that there is some $\alpha < \lambda$ such that $u \in U//S|_\alpha$ and all reduction steps beyond α do not take place at u or above. This is not possible due to the parallel-outermost reduction strategy that S adheres. \square



(a) Constructing a path in a term.



(b) The constructed path.

FIGURE 5. A path.

This shows that complete developments of any set of redex occurrences do always exist in any (almost) orthogonal system. This is already an improvement over strongly m -converging reductions, which only allow this if no collapsing rules are present or the considered set of redex occurrences does not contain an infinite set of nested collapsing redexes – also known as an *infinite collapsing tower*.

We shall discuss the issue of collapsing rules as well as infinite collapsing towers in more detail in the subsequent section, where we will show that complete developments are also unique in the sense that the final outcome is uniquely determined by the initial set of redexes occurrences.

5.3. Uniqueness of Complete Developments. The goal of this section is to show that the final term of a complete development is uniquely determined by the initial set of redex occurrences U . There are several techniques to show that in the metric model. One of these approaches, introduced by Kennaway and de Vries [14] and detailed by Ketema and Simonsen [20, 19] for infinitary combinatory reduction systems, uses so-called *paths*. Paths are constructed such that they, starting from the root, run through the initial term t of the complete development, and whenever a redex occurrence of the development is encountered, the path jumps to the root of the right-hand side of the corresponding rule and jumps back to the term t when it reaches a variable in the right-hand side.

Figure 5a illustrates this idea. It shows a path in a term t that encounters two redex occurrences of the complete development. As soon as such a redex occurrence is encountered, the path jumps to the right-hand side of the corresponding rule as indicated by the dashed arrows. Then the path runs through the right-hand side. When a variable is encountered, the path jumps back to the position of the term t that matches the variable. This jump is again indicated by a dashed arrow. The path that is obtained by this construction is shown in Figure 5b. With the collection of the thus obtained paths one can then construct the final term of the complete development. This technique – slightly modified – can also be applied in the present setting.

A path consists of nodes, which are connected by edges. We have two kinds of nodes: a node (\top, π) represents a location in the term t and a node (r, π, u) represents a location in the right-hand side r of a rule. These nodes of the form (\top, π) and (r, π, u) encode that the path is currently at position π in the term t resp. r . The additional component u provides the information that the path jumped to the right-hand side r from the redex $t|_u$. Both nodes and the edges between them are labelled. Each node is labelled with the symbol at the current location of the path, unless it is a redex occurrence in t or a variable occurrence in a right-hand side. The labellings of the edges provide information on how the path moves through the terms: a labelling i represents a move along the i -th edge in the term tree from the current location whereas an empty labelling indicates a jump from or to a right-hand side of a rule.

Definition 5.18 (path). Let \mathcal{R} be a left-linear TRS, t a partial term in \mathcal{R} , and U a set of pairwise non-conflicting redex occurrence in t . A U, \mathcal{R} -path (or simply *path*) in t is a sequence of length at most ω containing so-called *nodes* and *edges* in an alternating manner like this:

$$\langle n_0, e_0, n_1, e_1, n_2, e_2, \dots \rangle$$

where the ' n_i 's are nodes and the ' e_i 's are edges. A node is either a pair of the form (\top, π) with $\pi \in \mathcal{P}(t)$ or a triple of the form (r, π, u) with r the right-hand side of a rule in \mathcal{R} , $\pi \in \mathcal{P}(r)$, and $u \in U$. Edges are denoted by arrows \rightarrow . Both edges and nodes might be labelled by elements in $\Sigma_{\perp} \cup \mathcal{V}$ and \mathbb{N} , respectively. We write paths as the one sketched above as

$$n_0 \rightarrow n_1 \rightarrow n_2 \rightarrow \dots$$

or, when explicitly indicating labels, as

$$n_0 \xrightarrow{l_0} n_1 \xrightarrow{l_1} n_2 \xrightarrow{l_2} n_3 \xrightarrow{l_3} n_4 \xrightarrow{l_4} n_5 \xrightarrow{l_5} \dots$$

where empty labels are explicitly given by the symbol \emptyset . If a path has a segment of the form $n \rightarrow n'$, then we say there is an edge from n to n' or that n has an outgoing edge to n' .

Every path starts with the node $(\top, \langle \rangle)$ and is either infinitely long or ends with a node. For each node n having an outgoing edge to a node n' , the following must hold:

- (1) If n is of the form (\top, π) , then
 - (a) $n' = (\top, \pi \cdot i)$ and the edge is labelled by i , with $\pi \cdot i \in \mathcal{P}(t)$ and $\pi \notin U$, or
 - (b) $n' = (r, \langle \rangle, u)$ and the edge is unlabelled, with $t|_u$ a ρ -redex for $\rho: l \rightarrow r \in R$ and $u \in U$.
- (2) If n is of the form (r, π, u) , then
 - (a) $n' = (r, \pi \cdot i, u)$ and the edge is labelled by i , with $\pi \cdot i \in \mathcal{P}(r)$, or
 - (b) $n' = (\top, u \cdot \pi')$ and the edge is unlabelled, with $t|_u$ a ρ -redex for $\rho: l \rightarrow r \in R$, $r|_{\pi}$ a variable, and π' the unique occurrence of $r|_{\pi}$ in l .

Additionally, the nodes of a path are supposed to be labelled in the following way:

- (3) A node of the form (\top, π) is unlabelled if $\pi \in U$ and is labelled by $t(\pi)$ otherwise.
- (4) A node of the form (r, π, u) is unlabelled if $r|_{\pi}$ is a variable and labelled by $r(\pi)$ otherwise.

Remark 5.19. The above definition is actually a coinductive one. This is necessary to also define paths of infinite length. Also in [14] paths are considered to be possibly infinite, although they are defined inductively and are, therefore, finite.

Remark 5.20. Our definition of paths deviates slightly from the usual definition found in the literature [16, 20, 21]: In our setting, term nodes are of the form (\top, π) . The symbol \top is used to indicate that we are in the host term t . In the definitions found in the literature, the term t itself is used for that, i.e. term nodes are of the form (t, π) . Our definition of paths makes them less dependant on the term t they are constructed in. This makes it easier to construct a path in a host term from other paths in different host terms. This will become necessary in the proof of Lemma 5.33. However, we have to keep in mind that the node labels in a path are dependent on the host term under consideration. Thus, the labelling of a path might be different depending on which host term it is considered to be in.

Returning to the schematic example illustrated in Figure 5, we can observe how the construction of a path is carried out: The path starts with a segment in the term t . This segment is entirely regulated by the rule ((a)); all its edges and nodes are labelled according to ((a)) and (3). The jump to the right-hand side r_1 following that initial segment is justified by rule ((b)). This jump consists of a node (\top, u_1) , unlabelled according to (3), corresponding to the redex occurrence u_1 , and an unlabelled edge to the node $(r_1, \langle \rangle, u_1)$, corresponding to the root of the right-hand side r_1 . The segment of the path that runs through the right-hand side r_1 is subject to rule ((a)); again all its nodes and edges are labelled, now according to ((a)) and (4). As soon as a variable is reached in the right-hand side term (in the schematic example it is the variable x) a jump to the main term t is performed as required by rule ((b)). This jump consists of a node (r_1, π, u_1) , unlabelled according to (4), where π is the current position in r_1 , i.e. the variable occurrence, and an unlabelled edge to the node $(\top, u_1 \cdot \pi')$. The position π' is the occurrence of the variable x in the left-hand side. As we only consider left-linear systems, this occurrence is unique. Afterwards, the same behaviour is repeated: A segment in t is followed by a jump to a segment in the right-hand side r_2 which is in turn followed by a jump back to a final segment in t .

Note that paths do not need to be maximal. As indicated in the schematic example, the path ends somewhere within the main term, i.e. not necessarily at a constant symbol or a variable. What the example does not show, but which is obvious from the definition, is that paths can also terminate within a right-hand side. A jump back to the main term is only required if a variable is encountered.

The purpose of the concept of paths is to simulate the contraction of all redexes of the complete development in a locally restricted manner, i.e. only along some branch of the term tree. This locality will keep the proofs more concise and makes them easier to understand once we have grasped the idea behind paths. The strategy to prove our conjecture of uniquely determined final terms is to show that paths can be used to define a term and that a contraction of a redex of the complete development preserves a property of the collection of all paths which ensures that the induced term remains invariant. Then we only have to observe that the induced term of paths in a term with no redexes (in U) is the term itself.

The following fact is obvious from the definition of a path.

Fact 5.21. Let \mathcal{R} be a left-linear TRS, t a partial term in \mathcal{R} , and U a set of redex occurrences in t .

- (i) An edge in a U, \mathcal{R} -path in t is unlabelled iff the preceding node is unlabelled.
- (ii) Any prefix of a U, \mathcal{R} -path in t that ends in a node is also a U, \mathcal{R} -path in t .

As we have already mentioned, collapsing rules and in particular so-called infinite collapsing towers play a significant role in m -convergent reductions as they obstruct complete developments. Also in our setting of p -convergent reductions they are important as they are responsible for volatile positions:

Definition 5.22 (collapsing rules). Let \mathcal{R} be a TRS.

- (i) A rule $l \rightarrow r$ in \mathcal{R} is called *collapsing* if r is a variable. The unique position of the variable r in l is called the *collapsing position* of the rule.
- (ii) A ρ -redex is called *collapsing* if ρ is a collapsing rule.
- (iii) A *collapsing tower* is a non-empty sequence $(u_i)_{i < \alpha}$ of collapsing redex occurrences in a term t such that $u_{i+1} = u_i \cdot \pi_i$ for each $i < \alpha$, where π_i is a collapsing position of the redex at u_i . It is called *maximal* if it is not a proper prefix of another collapsing tower.

One can easily see that, in orthogonal TRSs, maximal collapsing towers in the same term are uniquely determined by their topmost redex occurrence. That is, two maximal collapsing towers $(u_i)_{i < \alpha}, (v_i)_{i < \alpha}$ in the same term are equal iff $u_0 = v_0$.

As mentioned, we shall use the U, \mathcal{R} -paths in a term t in order to define the final term of a complete development of U in t . However, in order to do that, we only need the information that is available from the labellings. The inner structure of nodes is only used for the bookkeeping that is necessary for defining paths. The following notion of traces defines projections to the labels of paths:

Definition 5.23 (trace). Let \mathcal{R} be a left-linear TRS, t a partial term in \mathcal{R} , and U a set of pairwise non-conflicting redex occurrences in t .

- (i) Let Π be a U, \mathcal{R} -path in t . The *trace* of Π , denoted $\text{tr}_t(\Pi)$, is the projection of Π to the labelling of its nodes and edges ignoring empty labels and the node label \perp .
- (ii) $\mathcal{P}(t, U, \mathcal{R})$ is used to denote the set of all U, \mathcal{R} -paths in t that end in a labelled node, or are infinite but have a finite trace. The set of traces of paths in $\mathcal{P}(t, U, \mathcal{R})$ is denoted by $\mathcal{Tr}(t, U, \mathcal{R})$.

By Fact 5.21, the trace of a path is a sequence alternating between elements in $\Sigma \cup \mathcal{V}$ and \mathbb{N} , which, if non-empty, starts with an element in $\Sigma \cup \mathcal{V}$. Moreover, by definition, $\mathcal{Tr}(t, U, \mathcal{R})$ is a set of finite traces of U, \mathcal{R} -paths in t .

As we have mentioned in Remark 5.20, the labelling of a path depends on the host term under consideration. Hence, also the trace of a path is depended on the host term. That is why we need to index the trace mapping $\text{tr}_t(\cdot)$ with the corresponding host term t .

Example 5.24. Consider the term $t = g(f(g(h(\perp))))$ and the TRS \mathcal{R} consisting of the two rules

$$f(x) \rightarrow h(x), \quad h(x) \rightarrow x.$$

Furthermore, let U be the set of all redex occurrences in t , viz. $U = \{\langle 0 \rangle, \langle 0 \rangle^3\}$. The following path Π is a U, \mathcal{R} -path in t :

$$\begin{aligned} (\top, \langle \rangle)^g &\xrightarrow{0} (\top, \langle 0 \rangle)^\emptyset \xrightarrow{\emptyset} (r_1, \langle \rangle, \langle 0 \rangle)^h \xrightarrow{0} (r_1, \langle 0 \rangle, \langle 0 \rangle)^\emptyset \xrightarrow{\emptyset} (\top, \langle 0 \rangle^2)^g \\ &\xrightarrow{0} (\top, \langle 0 \rangle^3)^\emptyset \xrightarrow{\emptyset} (r_2, \langle \rangle, \langle 0 \rangle^3)^\emptyset \xrightarrow{\emptyset} (\top, \langle 0 \rangle^4)^\perp \end{aligned}$$

As a matter of fact, Π is the greatest path of t . Hence, according to Fact 5.21, the set of all prefixes of Π ending in a node is the set of all U, \mathcal{R} -paths in t . Note that since Π itself ends in a labelled node, it is contained in $\mathcal{P}(t, U, \mathcal{R})$. The trace $\text{tr}_t(\Pi)$ of Π is the sequence

$$\langle g, 0, h, 0, g, 0 \rangle$$

Now consider the term $t' = g(f(g(h^\omega)))$ and the set U' of all its redexes, viz. $U' = \{\langle 0 \rangle\} \cup \{\langle 0 \rangle^3, \langle 0 \rangle^4, \dots\}$. Then the following path Π' is a U, \mathcal{R} -path in t' :

$$\begin{aligned} (\top, \langle \rangle)^g &\xrightarrow{0} (\top, \langle 0 \rangle)^\emptyset \xrightarrow{\emptyset} (r_1, \langle \rangle, \langle 0 \rangle)^h \xrightarrow{0} (r_1, \langle 0 \rangle, \langle 0 \rangle)^\emptyset \xrightarrow{\emptyset} (\top, \langle 0 \rangle^2)^g \xrightarrow{0} (\top, \langle 0 \rangle^3)^\emptyset \\ &\xrightarrow{\emptyset} (r_2, \langle \rangle, \langle 0 \rangle^3)^\emptyset \xrightarrow{\emptyset} (\top, \langle 0 \rangle^4)^\emptyset \xrightarrow{\emptyset} (r_2, \langle \rangle, \langle 0 \rangle^4)^\emptyset \xrightarrow{\emptyset} (\top, \langle 0 \rangle^5)^\emptyset \xrightarrow{\emptyset} \dots \end{aligned}$$

Π' is the greatest path of t' . The trace $\text{tr}_{t'}(\Pi')$ of Π' is the sequence

$$\langle g, 0, h, 0, g, 0 \rangle$$

Since Π' is infinitely long but has a finite trace, it is contained in $\mathcal{P}(t', U, \mathcal{R})$.

The lemma below shows that there is a one-to-one correspondence between paths in $\mathcal{P}(t, U, \mathcal{R})$ and their traces in $\mathcal{T}r(t, U, \mathcal{R})$.

Lemma 5.25 ($\text{tr}_t(\cdot)$ is a bijection). *Let \mathcal{R} be an orthogonal TRS, t a partial term in \mathcal{R} , and U a set of redex occurrences in t . $\text{tr}_t(\cdot)$ is a bijection from $\mathcal{P}(t, U, \mathcal{R})$ to $\mathcal{T}r(t, U, \mathcal{R})$.*

Proof. By definition, $\text{tr}_t(\cdot)$ is surjective. Let Π_1, Π_2 be two paths having the same trace. We will show that then $\Pi_1 = \Pi_2$ by an induction on the length of the common trace.

Let $\text{tr}_t(\Pi_1) = \langle \rangle$. Following Fact 5.21, there are two different cases: The first case is that $\Pi_1 = \Pi \cdot (\top, \pi)^\perp$, where the prefix Π corresponds to a finite maximal collapsing tower $(u_i)_{i \leq \alpha}$ starting at the root of t or Π is empty if such a collapsing tower does not exist. If the collapsing tower exists, then

$$\Pi = (\top, u_0)^\emptyset \xrightarrow{\emptyset} (r_0, \langle \rangle, u_0)^\emptyset \xrightarrow{\emptyset} (\top, u_1)^\emptyset \xrightarrow{\emptyset} (r_1, \langle \rangle, u_1)^\emptyset \xrightarrow{\emptyset} \dots \xrightarrow{\emptyset} (\top, u_\alpha)^\emptyset \xrightarrow{\emptyset}$$

But then also Π_2 starts with the prefix $\Pi \cdot (\top, \pi)$ due to the uniqueness of the collapsing tower and the involved rules. In both cases, $\Pi_1 = \Pi_2$ follows immediately.

The second case is that Π_1 is infinite. Then there is an infinite collapsing tower $(u_i)_{i < \omega}$ starting at the root of t . Hence,

$$\Pi_1 = (\top, u_0)^\emptyset \xrightarrow{\emptyset} (r_0, \langle \rangle, u_0)^\emptyset \xrightarrow{\emptyset} (\top, u_1)^\emptyset \xrightarrow{\emptyset} (r_1, \langle \rangle, u_1)^\emptyset \xrightarrow{\emptyset} \dots$$

$\Pi_1 = \Pi_2$ follows from the uniqueness of the infinite collapsing tower.

At first glance one might additionally find a third case where $\Pi_1 = \Pi \cdot (\top, \pi)^\emptyset \xrightarrow{\emptyset} (r, \langle \rangle, \pi)^\perp$ with Π a prefix corresponding to a collapsing tower as in the first case. However, this is not possible as it would require the occurrence of \perp in a rule.

Let $\text{tr}_t(\Pi_1) = f$. Then there are two cases: Either $\Pi_1 = \Pi \cdot (\top, \pi)^f$ or $\Pi_1 = \Pi \cdot (\top, \pi)^\emptyset \xrightarrow{\emptyset} (r, \langle \rangle, \pi)^f$, where the prefix Π corresponds to a finite maximal collapsing tower $(u_i)_{i \leq \alpha}$ starting at the root of t or Π is empty if such a collapsing tower does not exist. The

argument is analogous to the argument employed for the first case of the induction base above.

Finally, we consider the induction step. Hence, there are the two cases: Either $\text{tr}_t(\Pi_1) = T \cdot \langle i \rangle$ or $\text{tr}_t(\Pi_1) = T \cdot \langle i, f \rangle$. For both cases, the induction hypothesis can be invoked by taking two prefixes Π'_1 and Π'_2 of Π_1 and Π_2 , respectively, which both have the trace T and, therefore, are equal according to the induction hypothesis. The argument that the remaining suffixes of Π_1 and Π_2 are equal is then analogous to the argument for two base cases. \square

As mentioned above, the traces of paths contain all information necessary to define a term which we will later identify to be the final term of the corresponding complete development. The following definition explains how such a term, called a *matching term*, is determined:

Definition 5.26 (matching term). Let \mathcal{R} be a left-linear TRS, t a partial term in \mathcal{R} , and U a set of pairwise non-conflicting redex occurrences in t .

- (i) The *position* of a trace $T \in \mathcal{Tr}(t, U, \mathcal{R})$, denoted $\text{pos}(T)$, is the subsequence of T containing only the edge labels. The set of all positions of traces in $\mathcal{Tr}(t, U, \mathcal{R})$ is denoted $\mathcal{PTr}(t, U, \mathcal{R})$.
- (ii) The *symbol* of a trace $T \in \mathcal{Tr}(t, U, \mathcal{R})$, denoted $\text{sym}_t(T)$, is f if T ends in a node label f , and is \perp otherwise, i.e. whenever T is empty or ends in an edge label.
- (iii) A term t' is said to *match* $\mathcal{Tr}(t, U, \mathcal{R})$ if $\mathcal{P}(t') = \mathcal{PTr}(t, U, \mathcal{R})$ and $t'(\text{pos}(T)) = \text{sym}_t(T)$ for all $T \in \mathcal{Tr}(t, U, \mathcal{R})$.

Returning to the definition of paths, one can see that the label of a node is the symbol of the “current” position in a term. Similarly, the label of an edge says which edge in the term tree was taken at that point in the construction of the path. Hence, by projecting to the edge labels, we obtain the “history” of the path, i.e. the position. In the same way we obtain the symbol of that node by taking the label of the last node of the path, provided the corresponding path ends in a non- \perp -labelled node. In the other case that the trace does not end in a node label, the corresponding path either ends in a node labelled \perp or is infinite. As we will see, infinite paths with finite traces correspond to infinite collapsing towers, which in turn yield volatile positions within the complete development. Eventually, these volatile positions will also give rise to \perp subterms.

The following lemma shows that there is also a one-to-one correspondence between the traces in $\mathcal{Tr}(t, U, \mathcal{R})$ and their positions in $\mathcal{PTr}(t, U, \mathcal{R})$:

Lemma 5.27 ($\text{pos}(\cdot)$ is a bijection). Let \mathcal{R} be an orthogonal TRS, t a partial term in \mathcal{R} and U a set of redex occurrences in t . $\text{pos}(\cdot)$ is a bijection from $\mathcal{Tr}(t, U, \mathcal{R})$ to $\mathcal{PTr}(t, U, \mathcal{R})$.

Proof. An argument similar to the one for Lemma 5.25 can be given in order to show that the composition $\text{pos}(\cdot) \circ \text{tr}_t(\cdot)$ is a bijection. Together with the bijectivity of $\text{tr}_s(\cdot)$, according to Lemma 5.25, this yields the bijectivity of $\text{pos}(\cdot)$. \square

Having this lemma, the following proposition is an easy consequence of the definition of matching terms. It shows that matching terms do always exist and are uniquely determined:

Proposition 5.28 (unique matching term). Let \mathcal{R} be an orthogonal TRS, t a partial term in \mathcal{R} , and U a set of redex occurrences in t . Then there is a unique term, denoted $\mathcal{F}(t, U, \mathcal{R})$, that matches $\mathcal{Tr}(t, U, \mathcal{R})$.

Proof. Define the mapping $\varphi: \mathcal{PTr}(t, U, \mathcal{R}) \rightarrow \Sigma_{\perp} \cup \mathcal{V}$ by setting $\varphi(\text{pos}(T)) = \text{sym}_t(T)$ for each trace $T \in \mathcal{Tr}(t, U, \mathcal{R})$. By Lemma 5.27, φ is well-defined. Moreover, it is easy to see from the definition of paths, that $\mathcal{PTr}(t, U, \mathcal{R})$ is closed under prefixes and that φ respects the arity of the symbols, i.e. $\pi \cdot i \in \mathcal{PTr}(t, U, \mathcal{R})$ iff $0 \leq i < \text{ar}(\varphi(\pi))$. Hence, φ uniquely determines a term s with $s(\pi) = \varphi(\pi)$ for all $\pi \in \mathcal{PTr}(t, U, \mathcal{R})$. By construction, s matches $\mathcal{Tr}(t, U, \mathcal{R})$. Moreover, any other term s' matching $\mathcal{Tr}(t, U, \mathcal{R})$ must satisfy $s'(\pi) = \varphi(\pi)$ for all $\pi \in \mathcal{PTr}(t, U, \mathcal{R})$ and is therefore equal to s . \square

It is also obvious that the matching term of a term t w.r.t. an empty set of redex occurrences is the term t itself.

Lemma 5.29 (matching term w.r.t. empty redex set). *For any TRS \mathcal{R} and any partial term t in \mathcal{R} , it holds that $\mathcal{F}(t, \emptyset, \mathcal{R}) = t$.*

Proof. Straightforward. \square

Remark 5.30. Now it only remains to be shown that the matching term stays invariant during a development, i.e. that, for each development $S: t \xrightarrow{\mathcal{P}} t'$ of U , the matching terms $\mathcal{F}(t, U, \mathcal{R})$ and $\mathcal{F}(t', U//S, \mathcal{R})$ coincide. Since the matching term $\mathcal{F}(t, U, \mathcal{R})$ only depends on the set $\mathcal{Tr}(t, U, \mathcal{R})$ of traces, it is sufficient to show that $\mathcal{Tr}(t, U, \mathcal{R})$ and $\mathcal{Tr}(t', U//S, \mathcal{R})$ coincide. The key observation is that in each step $s \rightarrow s'$ in a development the paths in s' differ from the paths in s only in that they might omit some jumps. This can be seen in Figure 5a: In a step $s \rightarrow s'$ of a development, (some residual of) some redex occurrence in U is contracted. In the picture this corresponds to removing the pattern, say l_1 , of the redex and replacing it by the corresponding right-hand side r_1 of the rule. One can see that, except for the jump to and from the right-hand side r_1 the path remains the same.

In order to establish the above observation formally, we need a means to simulate reduction steps in a development directly as an operation on paths. The following definition provides a tool for this.

Definition 5.31 (position and prefix of a path). Let \mathcal{R} be a left-linear TRS, t a partial term in \mathcal{R} , U a set of pairwise non-conflicting redex occurrences in t , and $\Pi \in \mathcal{P}(t, U, \mathcal{R})$.

- (i) Π is said to *contain* a position $\pi \in \mathcal{P}(t)$ if it contains the node (\top, π) .
- (ii) For each $u \in U$, the *prefix* of Π by u , denoted $\Pi^{(u)}$, is defined as Π whenever Π does not contain u and otherwise as the unique prefix of Π that ends in (\top, u) .

Remark 5.32. It is obvious from the definition that each prefix $\Pi^{(u)}$ of a path $\Pi \in \mathcal{P}(t, U, \mathcal{R})$ by an occurrence u is the maximal prefix of Π , that does not contain positions that are proper extensions of u . Hence, if Π contains u , then $\Pi^{(u)}$ is the maximal prefix of Π that only contains prefixes of u (including u itself).

The following lemma is the key step towards proving the invariance of matching terms in developments. It formalises the observation described in Remark 5.30.

Lemma 5.33 (preservation of traces). *Let \mathcal{R} be an orthogonal TRS, t a partial term in \mathcal{R} , U a set of redex occurrences in t , and $S: t \xrightarrow{\mathcal{P}} t'$ a development of U in t . There is a surjective mapping $\vartheta_S: \mathcal{P}(t, U, \mathcal{R}) \rightarrow \mathcal{P}(t', U//S, \mathcal{R})$ such that $\text{tr}_t(\Pi) = \text{tr}_{t'}(\vartheta_S(\Pi))$ for all $\Pi \in \mathcal{P}(t, U, \mathcal{R})$.*

Proof. Let $S = (t_{\iota} \rightarrow_{\pi_{\iota}, c_{\iota}} t_{\iota+1})_{\iota < \alpha}$. We prove the statement by an induction on α .

If $\alpha = 0$, then the statement is trivially true.

Suppose that α is a successor ordinal $\beta + 1$. Let $T: t_0 \xrightarrow{\beta} t_\beta$ be the prefix of S of length β and $\varphi_\beta: t_\beta \rightarrow_{\pi_\beta} t_\alpha$ the last step of S , i.e. $S = T \cdot \langle \varphi_\beta \rangle$. By the induction hypothesis, there is a surjective mapping $\vartheta_T: \mathcal{P}(t, U, \mathcal{R}) \rightarrow \mathcal{P}(t_\beta, U', \mathcal{R})$, with $U' = U//T$ and $\text{tr}_t(\Pi) = \text{tr}_{t_\beta}(\vartheta_T(\Pi))$ for all $\Pi \in \mathcal{P}(t, U, \mathcal{R})$. By a careful case analysis (as done in [21]), one can show that there is a surjective mapping $\vartheta: \mathcal{P}(t_\beta, U', \mathcal{R}) \rightarrow \mathcal{P}(t_\alpha, U'', \mathcal{R})$, with $U'' = U'//\langle \varphi_\beta \rangle = U//S$ and $\text{tr}_{t_\beta}(\Pi) = \text{tr}_{t_\alpha}(\vartheta(\Pi))$ for all $\Pi \in \mathcal{P}(t_\beta, U', \mathcal{R})$. Hence, the composition $\vartheta_S = \vartheta \circ \vartheta_T$ is a surjective mapping from $\mathcal{P}(t, U, \mathcal{R})$ to $\mathcal{P}(t', U//S, \mathcal{R})$ and satisfies $\text{tr}_t(\Pi) = \text{tr}_{t'}(\vartheta_S(\Pi))$ for all $\Pi \in \mathcal{P}(t, U, \mathcal{R})$.

Let α be a limit ordinal. By induction hypothesis, there is a surjective mapping $\vartheta_{S|_\iota}$ for each proper prefix $S|_\iota$ of S satisfying $\text{tr}_{t_0}(\Pi) = \text{tr}_{t_\iota}(\vartheta_{S|_\iota}(\Pi))$ for all $\Pi \in \mathcal{P}(t, U, \mathcal{R})$. Let $\Pi \in \mathcal{P}(t, U, \mathcal{R})$ and $\Pi_\iota = \vartheta_{S|_\iota}(\Pi)$ for each $\iota < \alpha$. We define $\vartheta_S(\Pi)$ as follows:

$$\vartheta_S(\Pi) = \liminf_{\iota \rightarrow \alpha} \Pi_\iota^{(\pi_\iota)}$$

At first we have to show that ϑ_S is well-defined, i.e. that $\liminf_{\iota \rightarrow \alpha} \Pi_\iota^{(\pi_\iota)}$ is indeed a path in $\mathcal{P}(t', U//S, \mathcal{R})$, and that it preserves traces. There are two cases to be considered: If there is an outermost-volatile position π in S that is contained in Π_ι whenever $\pi_\iota = \pi$, then there is some $\beta < \alpha$ with $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \alpha$. Hence, $\vartheta_S(\Pi) = \Pi_\beta^{(\pi)}$. By Lemma 3.11 and Lemma 3.14, we have that $\Pi_\beta^{(\pi)} \in \mathcal{P}(t', U//S, \mathcal{R})$, in particular because $t'(\pi) = \perp$. Since the suffix Π' with $\Pi_\beta = \Pi_\beta^{(\pi)} \cdot \Pi'$ follows an infinite collapsing tower and is therefore entirely unlabelled, it cannot contribute to the trace of Π_β . Consequently,

$$\text{tr}_t(\Pi) \stackrel{IH}{=} \text{tr}_{t_\beta}(\Pi_\beta) = \text{tr}_{t'}(\Pi_\beta^{(\pi)}) = \text{tr}_{t'}(\vartheta_S(\Pi)).$$

If, on the other hand, there is no such outermost-volatile position, then either the sequence $(\Pi_\iota^{(\pi_\iota)})_{\iota < \alpha}$ becomes stable at some point or the sequence $(\prod_{\iota < \gamma} \Pi_\iota^{(\pi_\iota)})_{\gamma < \alpha}$ grows monotonically towards the infinite path $\vartheta_S(\Pi)$. In both cases well-definedness and preservation of traces follows easily from the induction hypothesis.

Lastly, we show the surjectivity of ϑ_S . To this end, assume some $\Pi \in \mathcal{P}(t', U//S, \mathcal{R})$. We show the existence of a path $\bar{\Pi} \in \mathcal{P}(t, U, \mathcal{R})$ with $\vartheta_S(\bar{\Pi}) = \Pi$ by distinguishing three cases:

- (a) Π ends in a redex node (r, π, u) . Hence, $u \in U//S$. According to Lemma 5.3, this means that there is some $\beta < \alpha$ such that

$$\pi_\iota \not\leq u \text{ for all } \beta \leq \iota < \alpha. \quad (1)$$

Consequently, all terms in $\{t_\iota \mid \beta \leq \iota < \alpha\}$ coincide in all prefixes of u , and each $v \in U//S$ with $v \leq u$ is in $U//S|_\iota$ for all $\beta \leq \iota < \alpha$. Hence, for all $\beta \leq \gamma < \alpha$ we have $\Pi \in \mathcal{P}(t_\gamma, U//S|_\gamma, \mathcal{R})$ with $\text{tr}_{t'}(\Pi) = \text{tr}_{t_\gamma}(\Pi)$. By induction hypothesis there is for each $\beta \leq \gamma < \alpha$ some $\Pi_\gamma \in \mathcal{P}(t, U, \mathcal{R})$ that is mapped to $\Pi \in \mathcal{P}(t_\gamma, U//S|_\gamma, \mathcal{R})$ by $\vartheta_{S|_\gamma}$ with $\text{tr}_t(\Pi_\gamma) = \text{tr}_{t_\gamma}(\Pi)$. Hence, $\text{tr}_t(\Pi_\gamma) = \text{tr}_{t'}(\Pi)$ which means that all paths Π_γ , with $\beta \leq \gamma < \alpha$, have the same trace in t and are therefore equal according to Lemma 5.25. Let us call this path $\bar{\Pi}$. That is, $\vartheta_{S|_\gamma}(\bar{\Pi}) = \Pi$ for all $\beta \leq \gamma < \alpha$. Since $\pi_\gamma \not\leq u$, we also have $(\vartheta_{S|_\gamma}(\bar{\Pi}))^{(\pi_\gamma)} = \Pi$. Consequently, $\vartheta_S(\bar{\Pi}) = \Pi$.

- (b) Π ends in a term node (\top, π) . Let $f = t'(\pi)$. If $f \neq \perp$, then we can apply Lemma 3.11 to obtain some $\beta < \alpha$ such that $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \alpha$. Then we can reason as in case ((a)) starting from (1). If $f = \perp$, then we have to distinguish two cases according to Lemma 3.14: If there is some $\beta < \alpha$ with $t_\beta(\pi) = \perp$ and $\pi_\iota \not\leq \pi$ for all $\beta \leq \iota < \alpha$, then

we can again employ the same argument as for case ((a)) starting from (1). Otherwise, i.e. if π is an outermost-volatile position in S , then we have some $\beta < \alpha$ such that $\pi_\iota \not\prec \pi$ for all $\beta \leq \iota < \alpha$ and such that

$$\text{for each } \beta \leq \gamma < \alpha \text{ there is some } \gamma \leq \gamma' < \alpha \text{ with } \pi_{\gamma'} = \pi. \quad (2)$$

Hence, we have for each $\beta \leq \gamma < \alpha$ some $\Pi_\gamma \in \mathcal{P}(t_\gamma, U//S|_\gamma, \mathcal{R})$ and an infinite collapsing tower $(u_i)_{i < \omega}$ in $U//S|_\gamma$ with $u_0 = \pi$ such that Π_γ is of the form

$$\Pi \cdot \xrightarrow{\emptyset} (r_0, \langle \rangle, u_0)^\emptyset \xrightarrow{\emptyset} (\top, u_1)^\emptyset \xrightarrow{\emptyset} (r_1, \langle \rangle, u_1)^\emptyset \xrightarrow{\emptyset} \dots$$

Therefore, $\text{tr}_{t_\gamma}(\Pi_\gamma) = \text{tr}_{t'}(\Pi)$. By induction hypothesis there is some $\overline{\Pi}_\gamma \in \mathcal{P}(t, T, \mathcal{R})$ with $\vartheta_{S|_\gamma}(\overline{\Pi}_\gamma) = \Pi_\gamma$ and $\text{tr}_t(\overline{\Pi}_\gamma) = \text{tr}_{t_\gamma}(\Pi_\gamma)$. Hence, $\text{tr}_t(\overline{\Pi}_\gamma) = \text{tr}_{t'}(\Pi)$, i.e. all $\overline{\Pi}_\gamma$ have the same trace in t and are therefore equal according to Lemma 5.25. Let us call this path $\overline{\Pi}$. Since $(\vartheta_{S|_\gamma}(\overline{\Pi}))^{(\pi)} = \Pi_\gamma^{(\pi)} = \Pi$ we can use (2) to obtain that $\vartheta_S(\overline{\Pi}) = \Pi$.

(c) Π is infinite. Hence, Π is of the form

$$\Pi' \cdot (\top, u_0)^\emptyset \xrightarrow{\emptyset} (r_0, \langle \rangle, u_0)^\emptyset \xrightarrow{\emptyset} (\top, u_1)^\emptyset \xrightarrow{\emptyset} (r_1, \langle \rangle, u_1)^\emptyset \xrightarrow{\emptyset} \dots$$

with $(u_i)_{i < \omega}$ an infinite collapsing tower in $U//S$. Consequently, by Lemma 5.3, for each $u_i \in U//S$ there is some $\beta_i < \alpha$ such that

$$u_i \in U//S|_{\beta_i} \text{ and } \pi_{\beta_i} \not\prec u_i \text{ for all } \beta_i \leq \gamma < \alpha. \quad (3)$$

Since $(u_i)_{i < \omega}$ is a chain (w.r.t. the prefix order), we can assume w.l.o.g. that $(\beta_i)_{i < \omega}$ is a chain as well. Following Remark 5.7, we obtain for each $u_i \in U//S$ its ancestor $v_i \in U$ with $v_i//S = u_i$. Let $\overline{\Pi}$ be the unique path in $\mathcal{P}(t, U, \mathcal{R})$ that contains each v_i and for each $j < \omega$ let Π_j be the unique path in $\mathcal{P}(t_{\beta_j}, U//S|_{\beta_j}, \mathcal{R})$ containing each $v_i//S|_{\beta_j}$. Clearly, $\vartheta_{S|_{\beta_j}}(\overline{\Pi}) = \Pi_j$. Note that we have for each $j < \omega$ that all paths $\vartheta_{S|_{\beta_j}}(\overline{\Pi})$ with $\beta_j \leq \iota < \alpha$ coincide in their prefix by u_j , which is a prefix of Π . Since additionally $(u_i)_{i < \omega}$ is a strict chain and because of (3), we can conclude that $\vartheta_S(\overline{\Pi}) = \Pi$. \square

The above lemma effectively establishes the invariance of matching terms during a development. Together with Lemma 5.29 this implies the uniqueness of final terms of complete developments of the same redex occurrences. As a corollary from this, we obtain that descendants are also unique among all complete developments:

Proposition 5.34 (final term and 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 the following holds:*

- (i) *Each complete development of U in t strongly p -converges to $\mathcal{F}(t, U, \mathcal{R})$.*
- (ii) *For each set $V \subseteq \mathcal{P}_\neq(t)$ and two complete developments S and T of U in t , respectively, it holds that $V//S = V//T$.*

Proof. ((i)) Let $S: t \xrightarrow{p} t'$ be a complete development of U in t strongly p -converging to t' . By Lemma 5.33, there is a surjective mapping $\vartheta: \mathcal{P}(t, U, \mathcal{R}) \rightarrow \mathcal{P}(t', U', \mathcal{R})$ with $\text{tr}_t(\Pi) = \text{tr}_{t'}(\vartheta(\Pi))$ for all $\Pi \in \mathcal{P}(t, U, \mathcal{R})$, where $U' = U//S$. Hence, it holds that $\text{Tr}(t, U, \mathcal{R}) = \text{Tr}(t', U', \mathcal{R})$ and, consequently, $\mathcal{F}(t, U, \mathcal{R}) = \mathcal{F}(t', U', \mathcal{R})$. Since S is a complete development of U in t , we have that $U' = \emptyset$ which implies, according to Lemma 5.29, that $\mathcal{F}(t', U', \mathcal{R}) = t'$. Therefore, $\mathcal{F}(t, U, \mathcal{R}) = t'$.

((ii)) Let $t' = t^{(V)}$. By Proposition 5.13, both reductions S and T can be uniquely lifted to reductions S' and T' in \mathcal{R}^\perp , respectively, such that $V//S$ and $V//T$ are determined by the




FIGURE 6. The Infinitary Strip Lemma.

final term of S' and T' , respectively. It is easy to see that also \mathcal{R}^\perp is an orthogonal TRS and that S' and T' are complete developments of U in t' . Hence, we can invoke clause ((i)) of this proposition to conclude that the final terms of S' and T' coincide and that, therefore, also $V//S$ and $V//T$ coincide. \square

By the above proposition, the descendants of a complete development of a particular set of redex occurrences are unique. Therefore, we adopt the notation $U//V$ for the descendants $U//S$ of U by some complete development S of V . According to Proposition 5.17 and Proposition 5.34, $U//V$ is well-defined for any orthogonal TRS.

Furthermore, Proposition 5.34 yields the following corollary establishing the diamond property of complete developments:

Corollary 5.35 (diamond property of complete developments). *Let \mathcal{R} be an orthogonal TRS and $t \xrightarrow{\mathcal{R}_U} t_1$ and $t \xrightarrow{\mathcal{R}_V} t_2$ be two complete developments of U respectively V in t . Then t_1 and t_2 are joinable by complete developments $t_1 \xrightarrow{\mathcal{R}_{V//U}} t'$ and $t_2 \xrightarrow{\mathcal{R}_{U//V}} t'$.*

Proof. By Proposition 5.5, it holds that

$$(U \cup V)//U = U//U \cup V//U = V//U.$$

Let $S: t \xrightarrow{\mathcal{R}_U} t_1$, $T: t \xrightarrow{\mathcal{R}_V} t_2$, $S': t_1 \xrightarrow{\mathcal{R}_{V//U}} t'$ and $T': t_2 \xrightarrow{\mathcal{R}_{U//V}} t''$. By the equation above and Proposition 5.8, we have that $S \cdot S': t \xrightarrow{\mathcal{R}_U} t_1 \xrightarrow{\mathcal{R}_{V//U}} t'$ is a complete development of $U \cup V$. Analogously, we obtain that $T \cdot T': t \xrightarrow{\mathcal{R}_V} t_2 \xrightarrow{\mathcal{R}_{U//V}} t''$ is a complete development of $U \cup V$, too. According to Proposition 5.34, this implies that both $S \cdot S'$ and $T \cdot T'$ strongly p -converge in the same term, i.e. $t' = t''$. \square

In the next section we shall make use of complete developments in order to obtain the Infinitary Strip Lemma for p -converging reductions and a limited form of infinitary confluence for orthogonal systems.

5.4. The Infinitary Strip Lemma. In this section we use the results we have obtained for complete developments in the previous two sections in order to establish that a complete development of a set of pairwise disjoint redex occurrences commutes with any strongly p -converging reduction:

Proposition 5.36 (Infinitary Strip Lemma). *Let \mathcal{R} be an orthogonal TRS, $S: t_0 \xrightarrow{\mathcal{R}}^\alpha t_\alpha$ a strongly p -convergent reduction, and $t_0 \xrightarrow{\mathcal{R}_U} s_0$ a complete development of a set U of pairwise disjoint redex occurrences in t_0 . Then t_α and s_0 are joinable by a reduction $S/T: s_0 \xrightarrow{\mathcal{R}} s_\alpha$ and a complete development $T/S: t_\alpha \xrightarrow{\mathcal{R}_{U//S}} s_\alpha$.*

Proof. We prove this statement by constructing the diagram shown in Figure 6. The U_ι 's in the diagram are sets of redex occurrences: $U_\iota = U//S|_\iota$ for all $0 \leq \iota \leq \alpha$. In particular, $U_0 = U$. All arrows in the diagram represent complete developments of the indicated sets of redex occurrences. Particularly, in each ι -th step of S the redex at v_ι is contracted. We will construct the diagram by an induction on α .

If $\alpha = 0$, then the diagram is trivial. If α is a successor ordinal $\beta + 1$, then we can take the diagram for the prefix $S|_\beta$, which exists by induction hypothesis, and extend it to a diagram for S . The existence of the additional square that completes the diagram for S is affirmed by Corollary 5.35 since $U_{\beta+1} = U_\beta//v_\beta$.

Let α be a limit ordinal. Moreover, let s'_α be the uniquely determined final term of a complete development of U_α in t_α . By induction hypothesis, the diagram exists for each proper prefix of S . Let $T_\iota: s_0 \xrightarrow{U} s_\iota$ denote the reduction at the bottom of the diagram for the reduction $S|_\iota$ for each $\iota < \alpha$. The set of all T_ι is directed. Hence, $T = \bigsqcup_{\iota < \alpha} T_\iota$ exists. Since $T_\iota < T$ for each $\iota < \alpha$, the diagram for S with $T: s_0 \xrightarrow{U} s_\alpha$ at the bottom satisfies almost all required properties. Only the equality of s_α and s'_α remains to be shown.

Note that, by Proposition 5.9, the redex occurrences in U_α are pairwise disjoint. Let $\pi \in U_\alpha$. By Lemma 5.3 and the definition of descendants, there is some $\beta < \alpha$ such that $\pi \in U_\beta$ and $v_\beta \not\leq \pi$ for all $\beta \leq \iota < \alpha$. Hence, for all $\pi' \in v_\iota//U_\iota$ with $\beta \leq \iota < \alpha$, we also have $\pi' \not\leq \pi$. That is, in the remaining reductions $t_\beta \xrightarrow{U} t_\alpha$ and $t_\beta \xrightarrow{U_\beta} s_\beta \xrightarrow{U} s_\alpha$, no reduction takes place at a proper prefix of π . Hence, by Lemma 3.11, t_β coincides with t_α and s_α in all proper prefixes of π . Since in the reduction $t_\alpha \xrightarrow{U_\alpha} s'_\alpha$ also no reduction takes place at a proper prefix of π , we obtain that t_α and s'_α and, thus, also s_α and s'_α coincide in all proper prefixes of π .

Let $\rho: l \rightarrow r$ be the rule for the redex $t_\beta|_\pi$ and $C\langle \dots \rangle, D\langle \dots \rangle$ ground contexts such that $l = C\langle x_1, \dots, x_k \rangle$ and $r = D\langle x_{p(1)}, \dots, x_{p(m)} \rangle$ for some pairwise distinct variables x_1, \dots, x_k and an appropriate mapping $p: \{1, \dots, m\} \rightarrow \{1, \dots, k\}$. Moreover, let t'_1, \dots, t'_k be terms such that $t_\iota = t_\iota[C\langle t'_1, \dots, t'_k \rangle]_\pi$ and $s_\iota = s_\iota[D\langle t'_{p(1)}, \dots, t'_{p(m)} \rangle]_\pi$ for all $\beta \leq \iota \leq \alpha$. The argument in the previous paragraph justifies the assumption of these elements. From β onward, all horizontal reduction steps in the diagram take place within the contexts $t_\iota[\cdot]_\pi$ and $s_\iota[\cdot]_\pi$, respectively, or inside the terms t'_i , and all vertical reductions take place within the contexts $t_\iota[C\langle \dots \rangle]_\pi$ and $s_\iota[D\langle \dots \rangle]_\pi$, respectively. In particular, we have $t_\alpha = t_\alpha[C\langle t'_1, \dots, t'_k \rangle]_\pi$ and $s_\alpha = s_\alpha[D\langle t'_{p(1)}, \dots, t'_{p(m)} \rangle]_\pi$. Let $t_\alpha \rightarrow_\pi t'_\alpha$. This reduction contracts the redex $C\langle t'_1, \dots, t'_k \rangle$ to the term $D\langle t'_{p(1)}, \dots, t'_{p(m)} \rangle$ using rule ρ . Note that a complete development $t_\alpha \xrightarrow{U_\alpha} s'_\alpha$ contracts, besides π , only redex occurrences disjoint with π . Hence, t'_α and s'_α coincide in all extensions of π . Since $t'_\alpha = t_\alpha[D\langle t'_{p(1)}, \dots, t'_{p(k)} \rangle]_\pi$ (and $s_\alpha = s_\alpha[D\langle t'_{p(1)}, \dots, t'_{p(m)} \rangle]_\pi$), we can conclude that s_α and s'_α coincide in all extensions of π .

Since the residual $\pi \in U_\alpha$ was chosen arbitrarily, the above holds for all elements in U_α . That is, s_α and s'_α coincide in all prefixes and all extensions of elements in U_α . It remains to be shown, that they also coincide in positions that are disjoint to all positions in U_α . To this end, we only need to show that t_α and s_α coincide in these positions since the complete development $t_\alpha \xrightarrow{U_\alpha} s'_\alpha$ keeps positions disjoint with all positions in U_α unchanged. Let π be such a position.

Suppose $t_\alpha(\pi) = f \neq \perp$. By Lemma 3.11, there is some $\beta < \alpha$ such that $t_\beta(\pi) = f$ and $v_\beta \not\leq \pi$ for all $\beta \leq \iota < \alpha$. Note that no prefix π' of π is in U_β since otherwise $\pi' \in U_\alpha$, by Lemma 5.3, which contradicts the assumption that π is disjoint to all positions in U_α . Hence,

$s_\beta(\pi) = f$ and $\pi' \not\leq \pi$ for all $\pi' \in v_\iota // U_\iota$ and $\beta \leq \iota < \alpha$, which means that no reduction step in $s_\beta \xrightarrow{p} s_\alpha$ takes place at some prefix of π . Thus, we can conclude, according to Lemma 3.11, that $s_\alpha(\pi) = f$. Similarly, one can show that $s_\alpha(\pi) = f \neq \perp$ implies $t_\alpha(\pi) = f$.

Suppose $t_\alpha(\pi) = \perp$. Hence, according to Lemma 3.14, π is outermost-volatile in S or there is some $\beta < \alpha$ such that $t_\beta(\pi) = \perp$ and $v_\iota \not\leq \pi$ for all $\beta \leq \iota < \alpha$. For the latter case, we can argue as in the case for $t_\alpha(\pi) \neq \perp$ above. In the former case, π is outermost-volatile in T as well. Thus, by applying Lemma 3.14, we obtain that $s_\alpha(\pi) = \perp$. A similar argument can be employed for the reverse direction. \square

The reduction S/T constructed in the proof above is called the *projection* of S by T . Likewise, the reduction T/S is called the *projection* of T by S . As a corollary we obtain the following semi-infinitary confluence result:

Corollary 5.37 (semi-infinitary confluence). *In every orthogonal TRS, two reductions $t \xrightarrow{p} t_2$ and $t \rightarrow^* t_1$ can be joined by two reductions $t_2 \xrightarrow{p} t_3$ and $t_1 \xrightarrow{p} t_3$.*

Proof. This can be shown by an induction on the length of the reduction $t \rightarrow^* t_1$. If it is empty, the statement trivially holds. The induction step follows from Proposition 5.36. \square

In the next section we shall, based on the Infinitary Strip Lemma, show that strong p -reachability coincides with Böhm-reachability, which then yields, amongst other things, full infinitary confluence of orthogonal systems.

6. COMPARING STRONG p -CONVERGENCE AND BÖHM-CONVERGENCE

In this section we shall show the core result of this paper: For orthogonal, left-finite TRSs, strong p -reachability and Böhm-reachability w.r.t. the set \mathcal{RA} of root-active terms 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 strong 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 6.1 (destructiveness, fragility). Let \mathcal{R} be a TRS.

- (i) A reduction $S: t \xrightarrow{p} s$ is called *destructive* if $\langle \rangle$ 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 3.14 we can immediately derive the following alternative characterisations:

Fact 6.2 (destructiveness, fragility). Let \mathcal{R} be a TRS.

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

One has to keep in mind, however, that a closed reduction to \perp is not destructive. Such a notion of destructiveness would include the empty reduction from \perp to \perp , and reductions that end with the contraction of a collapsing redex as, for example, in the single step reduction $f(\perp) \rightarrow \perp$ induced by the rule $f(x) \rightarrow x$. Such reductions do not “produce” the term \perp . They are merely capable of “moving” an already existent subterm \perp by a collapsing rule. In this sense, fragile terms are, according to Lemma 3.15, the only terms which can produce the term \perp . This is the key observation for studying the relation between strong p -convergence and Böhm-convergence.

In order to show that strong p -reachability and Böhm-reachability w.r.t. \mathcal{RA} coincide we will proceed as follows: At first we will show that strong p -reachability implies Böhm-reachability w.r.t. the set of total fragile terms, i.e. the fragile terms in $\mathcal{T}^\infty(\Sigma, \mathcal{V})$. From this we will derive a compression lemma for destructive reductions. We will then use this to show that the set \mathcal{RA} of root-active terms coincides with the set of total fragile terms. From this we conclude that strong p -reachability implies Böhm-reachability w.r.t. \mathcal{RA} . Finally, we then show the other direction of the equality.

6.1. From Strong p -Convergence to Böhm-Convergence. For the first step we have to transform a strongly p -converging reduction in to a Böhm-converging reduction w.r.t. the set of total fragile terms, i.e. a strongly m -converging reduction w.r.t. the corresponding Böhm extension \mathcal{B} . Recall that, by Theorem 4.12, the only difference between strongly p -converging reductions and strongly m -converging reductions is the ability of the former to produce \perp subterms. This happens, according to Lemma 3.14, precisely at volatile positions.

We can, therefore, proceed as follows: Given a strongly p -converging reduction we construct a Böhm-converging reduction by removing reduction steps which cause the volatility of a position in some open prefix of the reduction and then replacing them by a *single* \rightarrow_{\perp} -step.

The intuition of this construction is illustrated in Figure 7. It shows a strongly p -converging reduction of length $\omega \cdot 4$ from s to t . In order to maintain readability, we restrict the attention to a particular branch of the term (tree) as indicated in Figure 7a. The picture shows five positions which are volatile in some open prefix of the reduction. We assume that they are the only volatile positions at least in the considered branch. Note that the positions do not need to occur in all of the terms in the reduction. They might disappear and reappear repeatedly. Each of them, however, appears in infinitely many terms in the reduction, as, by definition of volatility, infinitely many steps take place at each of these positions. In Figure 7b, the prefixes of the reduction that contain a volatile position are indicated by a wavy rewrite arrow pointing to a \perp . The level of an arrow indicates the position which is volatile. A prefix might have multiple volatile positions. For example, both π_2 and π_4 are volatile in the prefix of length ω . But a position might also be volatile for several prefixes. For instance, π_3 is volatile in the prefix of length $\omega \cdot 2$ and the prefix of length $\omega \cdot 4$.

By Lemma 3.14, outermost-volatile positions are responsible for the generation of \perp subterms. By their nature, at some point there are no reductions taking place above outermost-volatile positions. The suffix where this is the case is a *nested* destructive reduction. The subterm where this suffix starts is, therefore, a fragile term and we can replace this suffix with a *single* \rightarrow_{\perp} -step. The segments which are replaced in this way are highlighted by dashed boxes in Figure 7b. As indicated by the dotted lines, this then also includes reduction steps which occur below the outermost-volatile positions. Therefore, also volatile positions which are not outermost are removed as well. Eventually, we obtain a reduction




FIGURE 7. Turning a p -converging reduction into a Böhm-converging reduction.

without volatile positions, which is, by Lemma 3.15, a strongly m -converging reduction in the Böhm extension, i.e. a Böhm-converging reduction in the original system:

Proposition 6.3 (strong 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 strongly p -convergent reduction $s \xrightarrow{\mathcal{R}} t$, there is a Böhm-convergent reduction $s \xrightarrow{\mathcal{B}} t$.*

Proof. Assume that there is a reduction $S = (t_\iota \rightarrow_{\pi_\iota} t_{\iota+1})_{\iota < \alpha}$ in \mathcal{R} that strongly p -converges to t_α . We will construct a strongly m -convergent reduction $T: t_0 \xrightarrow{\mathcal{B}} t_\alpha$ 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 $S': t_\beta|_\pi \xrightarrow{\mathcal{B}} \perp$ 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 $S'': r \xrightarrow{\mathcal{B}} \perp$ 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 $S': t_\beta|_\pi \xrightarrow{\mathcal{B}} \perp$ 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' strongly 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 3.14, $t_\lambda(\pi) = \perp$. Consequently, $t' = t_\lambda$.

This construction can be performed for all prefixes of S and their respective outermost-volatile positions. Thereby, we obtain a strongly p -converging reduction $T: t_0 \xrightarrow{\mathcal{B}} t_\alpha$ for which no prefix has a volatile position. By Lemma 3.15, 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 4.12, $T: t_0 \xrightarrow{\mathcal{B}} t_\alpha$. \square

6.2. From Böhm-convergence to Strong p -Convergence. Next, we establish a compression lemma for destructive reductions, i.e. that each destructive reduction can be compressed to length ω . Before we continue with this, we need to mention the following lemma from Kennaway et al. [17]:

Lemma 6.4 (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 $s \xrightarrow{\mathcal{B}} t$ implies $s \xrightarrow{\mathcal{R}} s' \xrightarrow{\perp} t$ for some term s' .²*

In the next proposition we show that, excluding \perp subterms, the final term of a strongly p -converging reduction can be approximated arbitrarily well by a finite reduction. This corresponds to Corollary 2.5 which establishes finite approximations for strongly m -convergent reductions.

Proposition 6.5 (finite approximation). *Let \mathcal{R} be a left-linear, left-finite TRS and $s \xrightarrow{\mathcal{R}} t$. Then, for each finite set $P \subseteq \mathcal{P}_\perp(t)$, there is a reduction $s \rightarrow_{\mathcal{R}}^* t'$ such that t and t' coincide in P .*

Proof. Assume that $s \xrightarrow{\mathcal{R}} t$. Then, by Proposition 6.3, there is a reduction $s \xrightarrow{\mathcal{B}} t$, where \mathcal{B} is the Böhm extension of \mathcal{R} w.r.t. the set of total, fragile terms of \mathcal{R} . By Lemma 6.4, there is a reduction $s \xrightarrow{\mathcal{R}} s' \xrightarrow{\perp} t$. 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 2.5, 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. \square

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

Proof. We consider the situation depicted in Figure 6. Since $S: t_0 \xrightarrow{\mathcal{R}} t_\alpha$ is destructive, we have, for each $\beta < \alpha$, some $\beta \leq \gamma < \alpha$ such that $v_\gamma = \langle \rangle$. If $v_\gamma = \langle \rangle$, then also $\langle \rangle \in v_\gamma // U_\gamma$ unless $\langle \rangle \in U_\gamma$. As by Proposition 5.9, U_γ is a set of pairwise disjoint positions, $\langle \rangle \in U_\gamma$ implies $U_\gamma = \{\langle \rangle\}$. This means that if $v_\gamma = \langle \rangle$ and $\langle \rangle \in U_\gamma$, then $U_\iota = \emptyset$ for all $\gamma < \iota < \alpha$. Thus, there is only at most one $\gamma < \alpha$ with $\langle \rangle \in U_\gamma$. Therefore, we have, for each $\beta < \alpha$, some $\beta \leq \gamma < \alpha$ such that $\langle \rangle \in v_\gamma // U_\gamma$. Hence, T is destructive. \square

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 6.7 (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 6.6, shows that there is a destructive reduction starting in t' . Thus, t' is fragile. \square

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 as the Compression Lemma for strongly p -convergent reductions.

Proposition 6.8 (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 6.5, 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 $T_0: t_0 \rightarrow^* s'_1 \xrightarrow{\rho, \langle \rangle} t_1$ ending with a contraction at the root. T_0 is of finite length and, according to Lemma 6.7, 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 non-empty 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 . \square

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 6.8, 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 6.9 (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 6.7, also s is fragile. Hence, there is a destructive reduction $S: s \xrightarrow{\omega} \perp$ starting in s . According to Proposition 6.8, 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' . \square

To prove the other direction of the equality of strong p -reachability and Böhm-reachability we need the property that strongly m -convergent reductions consisting only of \rightarrow_{\perp} -steps, i.e. contractions of \mathcal{RA}_{\perp} -terms to \perp , can be compressed to length at most ω as well. In order to show this, we will make use of the following lemma from Kennaway et al. [17]:

Lemma 6.10 (\perp, \mathcal{RA} -instances). *Let \mathcal{RA} be the root-active terms of an orthogonal, left-finite TRS and $t \in \mathcal{T}^{\infty}(\Sigma_{\perp}, \mathcal{V})$. If some \perp, \mathcal{RA} -instance of t is in \mathcal{RA} , then every \perp, \mathcal{RA} -instance of t is.*

Lemma 6.11 (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 a strongly 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 [18].

Let $S = (t_i \rightarrow_{\pi_i} t_{i+1})_{i < \alpha}$ be the mentioned reduction strongly m -converging to t_{α} , and let π be a position at which some reduction step in S takes place. That is, there is some β such that $\pi_{\beta} = \pi$. We will prove by induction on β that $t_0|_{\pi} \in \mathcal{RA}$.

Consider the term $t_{\beta}|_{\pi}$. Since a \rightarrow_{\perp} -rule is applied here, we have, according to Remark 2.9, that $t_{\beta}|_{\pi} \in \mathcal{RA}_{\perp}$. Let $V = \mathcal{P}_{\perp}(t_{\beta}|_{\pi})$. Hence, for each $v \in V$, there is some $\gamma < \beta$ such that $\pi_{\gamma} = \pi \cdot v$. Therefore, we can apply the induction hypothesis and get that $t_0|_{\pi \cdot v} \in \mathcal{RA}$ for all $v \in V$. It is clear that we can obtain $t_0|_{\pi}$ from $t_{\beta}|_{\pi}$ by replacing each \perp -occurrence at $v \in V$ with the corresponding term $t_0|_{\pi \cdot v}$. That is, $t_0|_{\pi}$ is a \perp, \mathcal{RA} -instance of $t_{\beta}|_{\pi}$. Because $t_{\beta}|_{\pi} \in \mathcal{RA}_{\perp}$, there is some \perp, \mathcal{RA} -instance of $t_{\beta}|_{\pi}$ in \mathcal{RA} . Thus, by Lemma 6.10, also $t_0|_{\pi}$ is in \mathcal{RA} . This closes the proof of the claim.

Now let $V = \mathcal{P}_{\perp}(t_{\alpha})$. Clearly, all positions in V are pairwise disjoint. Moreover, for each $v \in V$, there is a step in S that takes place at v . Hence, by the claim shown above, V is a set of occurrences in t_0 of terms in \mathcal{RA} . A complete development of V in t_0 leads to t_{α} and can be performed in at most ω steps by an outermost reduction strategy. \square

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 strong p -reachability and Böhm-reachability w.r.t. root-active terms.

Theorem 6.12 (strong 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 6.9 and Proposition 6.3.

Now consider the converse direction: Let $s \xrightarrow{m}_{\mathcal{B}} t$ be a strongly m -converging reduction in \mathcal{B} . W.l.o.g. we assume s to be total. Due to Lemma 6.4, there is a term $s' \in \mathcal{T}^{\infty}(\Sigma, \mathcal{V})$ such that there are strongly m -converging reductions $S: s \xrightarrow{m}_{\mathcal{R}} s'$ and $T: s' \xrightarrow{m}_{\perp} t$. By Lemma 6.11, we can assume that in $s' \xrightarrow{m}_{\perp} t$ only pairwise disjoint occurrences of root-active terms are contracted. By Proposition 6.9, 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, following Remark 2.9, we can construct a strongly 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 strongly m -converging reduction S , which, according to Theorem 4.12, is also strongly p -converging, we obtain the strongly p -converging reduction $S \cdot T': s \xrightarrow{p}_{\mathcal{R}} t$. \square

6.3. Corollaries. With the equivalence of strong p -reachability and Böhm-reachability established in the previous section, strongly p -convergent reductions inherit a number of important properties that are enjoyed by Böhm-convergent reductions:

Theorem 6.13 (infinitary confluence). *Every orthogonal, left-finite TRS is infinitarily confluent. That is, for each orthogonal, left-finite TRS, $s_1 \ll^p t \Downarrow s_2$ implies $s_1 \Downarrow t' \ll^p s_2$.*

Proof. Leveraging Theorem 6.12, this theorem follows from Theorem 2.10. \square

Returning to Example 2.6 again, we can see that, in the setting of strongly p -converging reduction, the terms g^ω and f^ω can now be joined by repeatedly contracting the redex at the root which yields two destructive reductions $g^\omega \Downarrow \perp$ and $f^\omega \Downarrow \perp$, respectively.

Theorem 6.14 (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 strongly p -reachable from t .*

Proof. This follows immediately from Theorem 6.12 and Theorem 2.11. \square

Combining Theorem 6.13 and Theorem 6.14, we obtain that each term in an orthogonal TRS has a unique normal form w.r.t. strong p -convergence. Due to Theorem 6.12, this unique normal form is the Böhm tree w.r.t. root-active terms.

Since strongly p -converging reductions in orthogonal TRS can always be transformed such that they consist of a prefix which is a strongly m -convergent reduction and a suffix consisting of nested destructive reductions, we can employ the Compression Lemma for strongly m -convergent reductions (Theorem 2.4) and the Compression Lemma for destructive reductions (Proposition 6.8) to obtain the Compression Lemma for strongly p -convergent reductions:

Theorem 6.15 (Compression Lemma for strongly p -convergent reductions). *For each orthogonal, left-finite TRS, $s \Downarrow t$ implies $s \Downarrow^{\leq \omega} t$.*

Proof. Let $s \Downarrow_{\mathcal{R}} t$. According to Theorem 6.12, we have $s \Downarrow_{\mathcal{B}} t$ for the Böhm extension \mathcal{B} of \mathcal{R} w.r.t. \mathcal{RA} and, therefore, by Lemma 6.4, we have reductions $S: s \Downarrow_{\mathcal{R}} s'$ and $T: s' \Downarrow_{\perp} t$. Due to Theorem 2.4, we can assume S to be of length at most ω and, due to Theorem 4.12, to be strongly p -convergent, i.e $S: s \Downarrow_{\mathcal{R}}^{\leq \omega} s'$. 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 6.11. Hence, each step is of the form $C[r] \rightarrow_{\perp} C[\perp]$ for some root-active term r . By Proposition 6.9, for each such term r , there is a destructive reduction $r \Downarrow_{\mathcal{R}} \perp$ which we can assume, in accordance with Proposition 6.8, to be of length ω . Hence, each step $C[r] \rightarrow_{\perp} C[\perp]$ can be replaced by the reduction $C[r] \Downarrow_{\mathcal{R}}^{\omega} C[\perp]$. Concatenating these reductions results in a reduction $T': s' \Downarrow_{\mathcal{R}} t$ of length at most $\omega \cdot \omega$. If $S: s \Downarrow_{\mathcal{R}}^{\leq \omega} s'$ is of finite length, we can interleave the reduction steps in T' such that we obtain a reduction $T'': s' \Downarrow_{\mathcal{R}}^{\omega} t$ of length ω . Then we have $S \cdot T'': s \Downarrow_{\mathcal{R}}^{\omega} t$. If $S: s \Downarrow_{\mathcal{R}}^{\leq \omega} s'$ has length ω , we construct a reduction $s \Downarrow_{\mathcal{R}} t$ 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 $s \Downarrow_{\mathcal{R}} t$ 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. [16]. \square

We do not know whether full orthogonality is essential for the Compression Lemma. However, as for strongly m -convergent reductions, the left-linearity part of it is:

Example 6.16 ([16]). Consider the TRS consisting of the rules $f(x, x) \rightarrow c, a \rightarrow g(a), b \rightarrow g(b)$. Then there is a strongly p -converging reduction

$$f(a, b) \rightarrow f(g(a), b) \rightarrow f(g(a), g(b)) \rightarrow f(g(g(a)), g(b)) \rightarrow \dots f(g^\omega, g^\omega) \rightarrow c$$

of length $\omega + 1$. However, there is no strongly p -converging reduction $f(a, b) \xrightarrow{\leq \omega} c$ (since there is no such strongly m -converging reduction).

We can use the Compression Lemma for strongly p -convergent reductions to obtain a stronger variant of Theorem 4.12 for orthogonal TRSs:

Corollary 6.17 (strong m -reachability = strong 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 $s \xrightarrow{m} t$ iff $s \xrightarrow{p} t$.*

Proof. The “only if” direction follows immediately from Theorem 4.12. For the “if” direction assume a reduction $S: s \xrightarrow{p} t$. According to Theorem 6.15, there is a reduction $T: s \xrightarrow{\leq \omega} t$. Hence, since s is total and totality is preserved by single reduction steps, $T: s \xrightarrow{\leq \omega} t$ is total. Applying Theorem 4.12, yields that $T: s \xrightarrow{m \leq \omega} t$. \square

Notice the similarity of the above corollary with the Compression Lemma. The Compression Lemma states that the reachability relation \xrightarrow{m} (as well as \xrightarrow{p}) is the same whether we consider ordinals beyond ω or not. Analogously, Corollary 6.17 states that the reachability relation \xrightarrow{p} on total terms is the same whether we allow partial convergence in between or not. More apt, however, is the comparison to the following corollary of the Compression Lemma (cf. Corollary 2.5): $s \xrightarrow{m} t$ implies $s \rightarrow^* t$ whenever t is a finite term. In other words, if the final term is finite, we only need finite reductions. Analogously, Corollary 6.17 states that if the initial term and the final term are total, we only need metric convergence.

7. 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. However, p -converging reductions can end in a term containing ‘ \perp ’s indicating positions of local divergence. Theorem 4.9, Theorem 4.12 and Corollary 6.17 show that the partial order 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).

The equivalence of strong p -reachability and Böhm-reachability shows that the differences between the metric and the partial order model can be compensated by simply adding rules that allow to replicate destructive reductions by \rightarrow_{\perp} -steps. By this equivalence, we additionally obtain infinitary normalisation and infinitary confluence for orthogonal systems – a considerable improvement over strong m -convergence. Both strong p -convergence and Böhm-convergence are defined quite differently and have independently justified intentions, yet they still induce the same notion of transfinite reachability. This suggests that this notion of transfinite reachability can be considered a “natural” choice – also because of its properties that admit unique normal forms. Nevertheless, while achieving the same goals as Böhm-extensions, the partial order approach provides a more intuitive and more elegant

model for transfinite reductions as it does not need the clumsily defined “shortcuts” provided by \rightarrow_{\perp} -steps, which depend on allowing infinite left-hand sides in rewrite rules. Vice versa destructive reductions in the partial order model provide a justification for admitting these shortcuts.

Related Work. This study of partial order convergence is inspired by Blom [6] 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 [7] 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 of redex contraction.

Future Work. While we have studied both weak and strong p -convergence and have compared it to the respective metric counterparts, we have put the focus on strong p -convergence. 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 in the metric model [23]. A starting point in this direction would be to find correspondences between weak and strong p -convergence. For example, in the metric setting we have that $s \xrightarrow{\mathcal{R}} t$ implies that there is some t' with $s \xrightarrow{\mathcal{B}} t'$ and $t \xrightarrow{\mathcal{B}} t'$ [14, Theorem 12.9.14]. If we had the analogous correspondence for p -convergence, we would immediately obtain infinitary normalisation and confluence for weak p -convergence.

Moreover, we have focused on orthogonal systems in this paper. It should be easy to generalise our results to almost orthogonal systems. The only difficulty is to deal with the ambiguity of paths when rules are allowed to overlay. This could be resolved by considering equivalence classes of paths instead. The move to weakly orthogonal systems is much more complicated: For strong m -convergence Endrullis et al. [11] have shown that weakly orthogonal systems do not even satisfy the infinitary unique normal form property (UN^{∞}), a property that orthogonal systems do enjoy [16]. Due to Theorem 4.12, this means that also in the setting of strong p -convergence, weakly orthogonal systems do not satisfy UN^{∞} and are therefore not infinitarily confluent either! Endrullis et al. [11] have shown that this can be resolved in the metric setting by prohibiting collapsing rules. However, it is not clear whether this result can be transferred to the partial order setting.

Another interesting direction to follow is the ability to finitely simulate transfinite reductions by term graph rewriting. For strong m -convergence this is possible, at least to some extent [15]. We think that a different approach to term graph rewriting, viz. the *double-pushout approach* [10] or the *equational approach* [1], is more appropriate for the present setting of p -convergence [8, 3].

ACKNOWLEDGEMENTS

I am indebted to Bernhard Gramlich for his constant support during the work on my master's thesis which made this work possible.

REFERENCES

- [1] Zena M Ariola and Jan Willem Klop. Equational term graph rewriting. *Fundam. Inf.*, 26(3-4):207–240, 1996.
- [2] André Arnold and Maurice Nivat. The metric space of infinite trees. algebraic and topological properties. *Fundam. Inf.*, 3(4):445–476, 1980.
- [3] Patrick Bahr. Infinitary rewriting - theory and applications. Master's thesis, Vienna University of Technology, Vienna, September 2009.
- [4] Patrick Bahr. Abstract models of transfinite reductions. In Christopher Lynch, editor, *Proceedings of the 21st International Conference on Rewriting Techniques and Applications*, volume 6 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 49–66, Dagstuhl, Germany, 2010. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [5] Patrick Bahr. Partial order infinitary term rewriting and böhm trees. In Christopher Lynch, editor, *Proceedings of the 21st International Conference on Rewriting Techniques and Applications*, volume 6 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 67–84, Dagstuhl, Germany, 2010. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [6] Stefan Blom. An approximation based approach to infinitary lambda calculi. In Vincent van Oostrom, editor, *Rewriting Techniques and Applications*, volume 3091 of *Lecture Notes in Computer Science*, pages 221–232. Springer Berlin / Heidelberg, 2004.
- [7] Andrea Corradini. Term rewriting in CT_{Σ} . In Marie-Claude Gaudel and Jean-Pierre Jouannaud, editors, *TAPSOFT'93: Theory and Practice of Software Development*, volume 668 of *Lecture Notes in Computer Science*, pages 468–484. Springer Berlin / Heidelberg, 1993.
- [8] Andrea Corradini and Frank Drewes. (Cyclic) term graph rewriting is adequate for rational parallel term rewriting. Technical Report TR-14-97, Università di Pisa, Dipartimento di Informatica, 1997.
- [9] Nachum Dershowitz, Stéphane Kaplan, and David A Plaisted. Rewrite, rewrite, rewrite, rewrite, rewrite, ... *Theor. Comput. Sci.*, 83(1):71–96, 1991.
- [10] Hartmut Ehrig, Michael Pfender, and Hans Jürgen Schneider. Graph-grammars: An algebraic approach. In *14th Annual Symposium on Switching and Automata Theory*, pages 167–180, Washington, DC, USA, 1973. IEEE Computer Society.
- [11] Jörg Endrullis, Clemens Grabmayer, Dimitri Hendriks, Jan Willem Klop, and Vincent van Oostrom. Unique normal forms in infinitary weakly orthogonal rewriting. In Christopher Lynch, editor, *Proceedings of the 21st International Conference on Rewriting Techniques and Applications*, volume 6 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 85–102, Dagstuhl, Germany, 2010. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [12] Joseph A Goguen, James W Thatcher, Eric G Wagner, and Jesse B Wright. Initial algebra semantics and continuous algebras. *J. ACM*, 24(1):68–95, 1977.
- [13] J L Kelley. *General Topology*, volume 27 of *Graduate Texts in Mathematics*. Springer-Verlag, 1955.
- [14] Richard Kennaway and Fer-Jan de Vries. Infinitary rewriting. In Terese [24], chapter 12, pages 668–711.
- [15] 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.
- [16] Richard Kennaway, Jan Willem Klop, M Ronan Sleep, and Fer-Jan de Vries. Transfinite reductions in orthogonal term rewriting systems. *Inf. Comput.*, 119(1):18–38, 1995.
- [17] Richard Kennaway, Vincent van Oostrom, and Fer-Jan de Vries. Meaningless terms in rewriting. *J. Funct. Logic Programming*, 1999(1):1–35, February 1999.
- [18] Jeroen Ketema. *Böhm-Like Trees for Rewriting*. PhD thesis, Vrije Universiteit Amsterdam, 2006.
- [19] Jeroen Ketema and Jakob Grue Simonsen. On confluence of infinitary combinatory reduction systems. In Geoff Sutcliffe and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning*, volume 3835 of *Lecture Notes in Computer Science*, pages 199–214. Springer Berlin / Heidelberg, 2005.

- [20] Jeroen Ketema and Jakob Grue Simonsen. Infinitary combinatory reduction systems: Normalising reduction strategies. *Log. Methods Comput. Sci.*, 6(1):1–35, 2010.
- [21] Jeroen Ketema and Jakob Grue Simonsen. Infinitary combinatory reduction systems. *Inf. Comput.*, 209(6):893–926, 2011.
- [22] Pieter Hendrik Rodenburg. Termination and confluence in infinitary term rewriting. *J. Symbolic Logic*, 63(4):1286–1296, 1998.
- [23] Jakob Grue Simonsen. On confluence and residuals in Cauchy convergent transfinite rewriting. *Inf. Process. Lett.*, 91(3):141–146, 2004.
- [24] Terese. *Term Rewriting Systems*. Cambridge University Press, 1st edition, 2003.