Infinitary Term Graph Rewriting

Patrick Bahr
Department of Computer Science,
University of Copenhagen
Universitetsparken 1, 2100 Copenhagen, Denmark
paba@diku.dk

5th July 2011

Abstract

Term graph rewriting provides a formalism for implementing term rewriting in an efficient manner by avoiding duplication. Infinitary term rewriting has been introduced to study infinite term reduction sequences. Such infinite reductions can be used to reason about lazy evaluation. In this paper, we combine term graph rewriting and infinitary term rewriting thereby addressing both components of lazy evaluation: non-strictness and sharing. Moreover, we show how our theoretical underpinnings, based on a metric space and a complete semilattice, provides a unified framework for both term rewriting and term graph rewriting. This makes it possible to study the correspondences between these two worlds. As an example, we show how the soundness of term graph rewriting w.r.t. term rewriting can be extended to the infinitary setting.

Introduction

Infinitary term rewriting [14] extends the theory of term rewriting by giving a meaning to transfinite reductions instead of dismissing them as undesired and meaningless artifacts. *Term graphs*, on the other hand, allow to explicitly represent and reason about sharing and recursion [2] by dropping the restriction to a tree structure that we have for terms. Apart from that, term graphs also provide a finite representation of certain infinite terms, viz. *rational terms*. As Kennaway et al. [13, 15] have shown, this can be leveraged in order to finitely represent restricted forms of infinitary term rewriting using *term graph rewriting*.

However, in order to benefit from this, we need to know for which class of term rewriting systems the set of rational terms is closed under (normalising) reductions. One such class of systems – a rather restrictive one – is the class of *regular equation systems* [9] which consist of rules having only constants on their left-hand side. Having an understanding of infinite reductions over term graphs could help to investigate closure properties of rational terms in the setting of infinitary term rewriting.

By studying infinitary calculi of term graph rewriting, we can also expect to better understand calculi with explicit sharing and/or recursion. Due to the lack of finitary confluence of these systems, Ariola and Blom [1] resort to a notion of skew confluence in order to be able to define infinite normal forms. An appropriate infinitary calculus could provide a direct approach to define infinite normal forms.

Historically, the theory of infinitary term rewriting is mostly based on the metric space of terms [14]. Its notion of convergence captures "well-behaved" transfinite reductions. A more structured approach, based on the complete semilattice structure of terms, yields a conservative extension of the metric calculus of infinitary term rewriting [5] that allows local divergence.

In previous work [6], we have carefully devised a complete metric space and a complete semilattice of term graphs in order to investigate different modes of convergence for term graphs. The resulting theory allows to treat infinitary term rewriting as well as graph rewriting in the same theoretical framework. While the devised metric and partial order on term graphs manifests the same compatibility that is known for terms [5], it is too restrictive as we will illustrate.

In this paper, we follow a different approach by taking the arguably simplest generalisation of the metric space and the complete semilattice of terms to term graphs. While the notion of convergence in these structures has some oddities which makes them somewhat incompatible, we will show that these incompatibilities vanish once we move from the weak notion of convergence that was considered in [6] to the much more well-behaved strong notion of convergence [16]. More concretely, we will show that, w.r.t. strong convergence, the metric calculus of infinitary term graph rewriting is the *total fragment* of the partial order calculus of infinitary term graph rewriting.

We show that our simple approach to infinitary term graph rewriting yields simple limit constructions that makes them easy to relate to the limit constructions on terms. As a result of that we are able to generalise the soundness result as well as a limited completeness result for term graph rewriting [15] to the infinitary setting.

1 Preliminaries

We assume the reader to be familiar with the basic theory of ordinal numbers, orders and topological spaces [11], as well as term rewriting [19]. In the following, we briefly recall the most important notions.

1.1 Sequences

We use $\alpha, \beta, \gamma, \lambda, \iota$ to denote ordinal numbers. A *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 $(a_i)_{i<n}$ we also use the notation $\langle a_0, a_1, \ldots, a_{n-1}\rangle$. 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 $S \le T$ (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 $\le$ forms a complete semilattice. 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 $f\colon \beta \to \alpha$ such that $b_\iota = a_{f(\iota)}$ for all $\iota < \beta$. The subsequence S is called *finial* if f is cofinal, i.e. if for each $\iota < \beta$ there is some $\gamma < \alpha$ with $f(\gamma) \ge \iota$.

1.2 Metric Spaces

A pair $(M, \mathbf{d})$ is called a *metric space* if $\mathbf{d}\colon M \times M \to \mathbb{R}^+_0$ 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 $\lim_{\iota\to\alpha} a_\iota$, 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 $\lim_{\iota\to\lambda} a_\iota = a_\lambda$ 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.

Note that the limit of a converging sequence is preserved by taking cofinal subsequences:

Proposition 1.1 (invariance of the limit)**.** *Let $(a_i)_{i<\alpha}$ be a sequence in a metric space $(A, \mathbf{d})$. If $\lim_{\iota\to\alpha} a_\iota = a$ then $\lim_{\iota\to\beta} b_\iota = a$ for any cofinal subsequence $(b_i)_{i<\beta}$ of $(a_i)_{i<\alpha}$.*

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*. 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 partial order* (*cpo*) if it has a least element and each directed set D has a *least upper bound* (*lub*) $\bigsqcup D$. A cpo $(A, \leq)$ is called a *complete semilattice* if every *non-empty* set B has *greatest lower bound* (*glb*) $\bigsqcap B$. In particular, this means that for any non-empty sequence $(a_\iota)_{\iota<\alpha}$ in a complete semilattice, its *limit inferior*, defined by $\liminf_{\iota\to\alpha} a_\iota = \bigsqcup_{\beta<\alpha} \left(\bigsqcap_{\beta\leq\iota<\alpha} a_\iota\right)$, always 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 1.2 (invariance of the limit inferior)**.** *Let $(a_\iota)_{\iota<\alpha}$ be a sequence in a partially ordered set and $(b_\iota)_{\iota<\beta}$ a non-empty suffix of $(a_\iota)_{\iota<\alpha}$. Then $\liminf_{\iota\to\alpha} a_\iota = \liminf_{\iota\to\beta} b_\iota$.*

Proof. We have to show that $\bigsqcup_{\gamma<\alpha} \bigsqcap_{\gamma\leq\iota<\alpha} a_\iota = \bigsqcup_{\beta\leq\gamma<\alpha} \bigsqcap_{\gamma\leq\iota<\alpha} a_\iota = \overline{a}'$ holds for each $\beta < \alpha$. Let $b_\gamma = \bigsqcap_{\gamma\leq\iota<\alpha} a_\iota$ for each $\gamma < \alpha$, $A = \{b_\gamma \mid \gamma < \alpha\}$ and $A' = \{b_\gamma \mid \beta \leq \gamma < \alpha\}$. Note that $\overline{a} = \bigsqcup A$ and $\overline{a}' = \bigsqcup A'$. Because $A' \subseteq A$, we have that $\overline{a}' \leq \overline{a}$. On the other hand, since $b_\gamma \leq b_{\gamma'}$ for $\gamma \leq \gamma'$, we find, for each $b_\gamma \in A$, some $b_{\gamma'} \in A'$ with $b_\gamma \leq b_{\gamma'}$. Hence, $\overline{a} \leq \overline{a}'$. Therefore, due to the antisymmetry of $\leq$, we can conclude that $\overline{a} = \overline{a}'$. $\square$

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 $\alpha = 0$, $\left(\prod_{\iota<\alpha'} S_\iota\right) \cdot S_{\alpha'}$ if $\alpha = \alpha' + 1$, and $\bigsqcup_{\gamma<\alpha} \prod_{\iota<\gamma} S_\iota$ if α is a limit ordinal.

1.4 Terms

Since we are interested in the infinitary calculus of term rewriting, we consider the set $\mathcal{T}^\infty(\Sigma)$ of *infinitary terms* (or simply *terms*) over some *signature* Σ. A *signature* Σ is a countable set of symbols. Each symbol f is associated with its arity $\mathsf{ar}(f) \in \mathbb{N}$, and we write $\Sigma^{(n)}$ for the set of symbols in Σ which have arity n. The set $\mathcal{T}^\infty(\Sigma)$ is defined as the *greatest* set T such that $t \in T$ implies $t = f(t_1, \dots, t_k)$, where $f \in \Sigma^{(k)}$, and $t_1, \dots, t_k \in T$. For each constant symbol $c \in \Sigma^{(0)}$, we write c for the term $c()$. We consider $\mathcal{T}^\infty(\Sigma)$ as a superset of the set $\mathcal{T}(\Sigma)$ of *finite terms*. For a term $t \in \mathcal{T}^\infty(\Sigma)$ we use the notation $\mathcal{P}(t)$ to denote the *set of positions* in t. $\mathcal{P}(t)$ is the least subset of $\mathbb{N}^*$ such that $\langle\rangle \in \mathcal{P}(t)$ and $\pi \cdot \langle i \rangle \in \mathcal{P}(t)$ if $t = f(t_1, \dots, t_k)$ with $0 \le i < k$. For terms $s, t \in \mathcal{T}^\infty(\Sigma)$ and a position $\pi \in \mathcal{P}(t)$, we write $t|_\pi$ for the *subterm* of t at π, $t(\pi)$ for the function symbol in t at π, and $t[s]_\pi$ for the term t with the subterm at π replaced by s. A position is also called an *occurrence* if the focus lies on the subterm at that position rather than the position itself.

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

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

That is, $\mathsf{sim}(s,t)$ is the minimal depth at which s and t differ, resp. ∞ if $s = t$. Based on this, a distance function $\mathbf{d}$ can be defined by $\mathbf{d}(s,t) = 2^{-\mathsf{sim}(s,t)}$, where we interpret $2^{-\infty}$ as 0. The pair $(\mathcal{T}^\infty(\Sigma), \mathbf{d})$ is known to form a complete ultrametric space [3]. *Partial terms*, i.e. terms over signature $\Sigma_\bot = \Sigma \uplus \{\bot\}$ with $\bot$ a fresh constant symbol, can be endowed with a binary relation $\le_\bot$ by defining $s \le_\bot t$ iff s can be obtained from t by replacing some subterm occurrences in t by $\bot$. Interpreting the term $\bot$ as denoting “undefined”, $\le_\bot$ can be read as “is less defined than”. The pair $(\mathcal{T}^\infty(\Sigma_\bot), \le_\bot)$ is known to form a complete semilattice [10]. To explicitly distinguish them from partial terms, we call terms in $\mathcal{T}^\infty(\Sigma)$ *total*.

1.5 Term Rewriting Systems

For term rewriting systems, we have to consider terms with variables. To this end, we assume a countably infinite set $\mathcal{V}$ of variables and extend a signature Σ to a signature $\Sigma_\mathcal{V} = \Sigma \uplus \mathcal{V}$ with variables in $\mathcal{V}$ as nullary symbols. Instead of $\mathcal{T}^\infty(\Sigma_\mathcal{V})$ we also write $\mathcal{T}^\infty(\Sigma, \mathcal{V})$. 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 $l \to r$ with $l \in \mathcal{T}(\Sigma,\mathcal{V}) \setminus \mathcal{V}$ and $r \in \mathcal{T}^\infty(\Sigma,\mathcal{V})$ such that all variables in r are contained in l. Note that the left-hand side must be a finite term [14]! We usually use x, y, z and primed resp. indexed variants thereof to denote variables in $\mathcal{V}$.

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

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

Instead of $s \to_\mathcal{R} t$, we sometimes write $s \to_{\pi,\rho} t$ in order to indicate the applied rule ρ and the position π, or simply $s \to t$. 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 $\rho_1\colon l_1 \to r_1$, $\rho_2\colon l_2 \to r_2$ be rules in a TRS $\mathcal{R}$ with variables renamed apart. 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 $t \in \mathcal{T}(\Sigma,\mathcal{V})$ 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 Infinitary Term Rewriting

Before pondering over the right approach to an infinitary calculus of term graph rewriting, we want to provide a brief overview of infinitary term graph rewriting [14, 5, 8]. This should give a insight into the different approaches to deal with infinite reductions.

A *(transfinite) reduction* in a term rewriting system $\mathcal{R}$, is a sequence $S = (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\alpha}$ of rewriting steps in $\mathcal{R}$. The reduction S is called *weakly m-continuous*, written $S\colon\ t_0 \xrightarrow{m} \dots$, if the sequence of terms $(t_\iota)_{\iota<\widehat{\alpha}}$ is continuous, i.e. $\lim_{\iota\to\lambda} t_\iota = t_\lambda$ for each limit ordinal $\lambda < \alpha$. The reduction S is said to *weakly m-converge* to a term t, written $S\colon\ t_0 \xrightarrow{m} t$, if it is weakly m-continuous and $\lim_{\iota\to\alpha} t_\iota = t$.

For strong convergence, also the positions π_ι at which reductions take place are taken into consideration: A reduction S is called strongly *m-continuous*, written $S\colon\ t_0 \xrightarrow{m}\!\!\!\!\!\!\to \dots$, if it is weakly m-continuous and the depths of redexes $(|\pi_\iota|)_{\iota<\lambda}$ tend to infinity for each limit ordinal $\lambda < \alpha$, i.e. $liminf_{\iota\to\lambda} |\pi_\iota| = \omega$. A reduction S is said to strongly m-converge to t, written $S\colon\ t_0 \xrightarrow{m}\!\!\!\!\!\!\to t$, if it weakly m-converges to t and the depths of redexes $(|\pi_\iota|)_{\iota<\lambda}$ tend to infinity for each limit ordinal $\lambda \leq \alpha$.

Example 2.1. Consider the term rewriting system $\mathcal{R}$ containing the rule $\rho_1\colon\ a : x \to b : a : x$, where $:$ is a binary symbol that we write infix and assume to associate to the right. That is, the right-hand side of the rule is parenthesised as $b : (a : x)$. Think of the $:$ symbol as the list constructor *cons*. Using ρ_1. we have the infinite reduction

$$S\colon\ a : c \to b : a : c \to b : b : a : c \to \dots$$

The position at which two consecutive terms differ moves deeper and deeper during the reduction S. Hence, S weakly m-converges to the infinite term s satisfying the equation $s = b : s$, i.e. $s = b : b : b : \dots$. Since also the position at which the reductions take place moves deeper and deeper, S also strongly m-converges to s.

Now consider a TRS with the slightly different rule $\rho_2\colon\ a : x \to a : b : x$. This yields a reduction

$$S'\colon\ a : c \to a : b : c \to a : b : b : c \to \dots$$

The reduction S' weakly m-converges to the term $s' = a : b : b : \dots$. However, since in each step in S' takes place at the root, it is not strongly m-converging.

Strong m-convergence is determined by the depth of the redexes only. The metric space is only used to determine the limit term.

Proposition 2.2 ([4, Prop. 5.5])**.** *Let $S = (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\lambda}$ be a strongly m-continuous open reduction in a TRS. Then S is strongly m-convergent iff the sequence $(|\pi|_\iota)_{\iota<\lambda}$ of redex depths tends to infinity.*

In the partial order model of infinitary rewriting, convergence is modelled by the limit inferior: A reduction $S = (t_\iota \to_{\pi_\iota} t_{\iota+1})_{\iota<\alpha}$ of *partial terms* is called *weakly p-continuous*, written $S\colon\ t_0 \xrightarrow{p} \dots$, if $\liminf_{\iota<\lambda} t_\iota = t_\lambda$ for each limit ordinal $\lambda < \alpha$. The reduction S is said to *weakly p-converge* to a term t, written $S\colon\ t_0 \xrightarrow{p} t$, if it is weakly p-continuous and $\liminf_{\iota<\widehat{\alpha}} t_\iota = t$.

Again, for strong convergence, the positions π_ι at which reductions take place are taken into consideration. In particular, we consider for a reduction step $t_\iota \to_{\pi_\iota} t_{\iota+1}$ the *reduction context* $c_\iota = t_\iota[\perp]_\pi$. To indicate the reduction context of a reduction step, we also write $t_\iota \to_{c_\iota} t_{\iota+1}$. A reduction $S = (t_\iota \to_{c_\iota} t_{\iota+1})_{\iota<\alpha}$ is called *strongly p-continuous*, written $S\colon t_0 \xrightarrow{p}\!\!\!\!\!\rightarrow \dots$, if $\liminf_{\iota<\lambda} c_\iota = t_\lambda$ for each limit ordinal $\lambda < \alpha$. The reduction S is said to *weakly p-converge* to a term t, written $S\colon t_0 \xrightarrow{p}\!\!\!\!\!\rightarrow t$, if it is weakly p-continuous and either T is closed with $t = t_\alpha$, or $\liminf_{\iota<\widehat{\alpha}} c_\iota = t$.

The distinguishing feature of the partial order approach is that, given a complete semilattice, each continuous reduction also converges. This provides a conservative extension to m-convergence that allows rewriting modulo *meaningless terms* [5] by essentially mapping those parts of the reduction to $\perp$ that are divergent according to the metric model.

Intuitively, weak p-convergence on terms describes an approximation process. To this end, the partial order $\leq_\perp$ captures a notion of *information preservation*, i.e. $s \leq_\perp t$ iff t contains at least the same information as s does but potentially more. A monotonic sequence of terms $t_0 \leq_\perp t_1 \leq_\perp \dots$ thus approximates the information contained in $\bigsqcup_{i<\omega} t_i$. Given this reading of $\leq_\perp$, the glb $\bigsqcap T$ of a set of terms T captures the common (non-contradicting) information of the terms in T. Leveraging this, a sequence that is not necessarily monotonic can be turned into a monotonic sequence $t_j = \bigsqcap_{j\leq i<\omega} s_j$ such that each t_j contains exactly the information that remains stable in $(s_i)_{i<\omega}$ from j onwards. Hence, the limit inferior $\liminf_{i\to\omega} s_i = \bigsqcup_{j<\omega} \bigsqcap_{j\leq i<\omega} s_i$ is the term that contains the accumulated information that eventually remains stable in $(s_i)_{i<\omega}$. This is expressed as an approximation of the monotonically increasing information that remains stable from some point on. For the strong variant, instead of the terms s_ι, the reduction contexts c_ι are considered. Each reduction context c_ι is an underapproximation of the shared structure $s_\iota \sqcap s_{\iota+1}$ between two consecutive terms $s_\iota, s_{\iota+1}$.

Example 2.3. Reconsider the system from Example 2.1. The reduction S also weakly and strongly p-converges to s. Its sequence of stable information $\perp : \perp \leq_\perp b : \perp : \perp \leq_\perp b : b : \perp : \perp \leq_\perp \dots$ approximates s. The same also applies to the stricter underapproximation $\perp \leq_\perp b : \perp \leq_\perp b : b : \perp \leq_\perp \dots$ by reduction contexts. Now consider the rule ρ_1 together with the rule $\rho_3\colon\ b : x \to a : b : x$. Starting with the same term, but applying the two rules alternately at the root, we obtain the reduction sequence

$$T\colon\ a : c \to b : a : c \to a : b : a : c \to b : a : b : a : c \to \dots$$

Now the differences between two consecutive terms occur right below the root symbol " : ". Hence, T does not even weakly m-converge. This, however, only affects the left argument of " : ". Following the right argument position, the bare list structure becomes eventually stable. The sequence of stable information $\perp : \perp \leq_\perp \perp : \perp : \perp \leq_\perp \perp : \perp : \perp : \perp \leq_\perp \dots$ approximates the term $t = \perp : \perp : \perp \dots$. Hence, T weakly p-converge to t. Since each reduction takes place at the root, each reduction context is $\perp$. Therefore, T strongly p-converges to the term $\perp$.

Note that in both the metric and the partial order setting continuity is simply the convergence of every proper prefix:

Proposition 2.4 ([4])**.** *Let $S = (t_\iota \to t_{\iota+1})_{\iota<\alpha}$ be a reduction in a TRS. Then S is strongly m-continuous iff every proper prefix $S|_\beta$ strongly m-converges to t_β The same holds for strong p-continuity/-convergence and weak counterparts.*

Moreover, the relation between m- and p-convergence illustrated in the examples above is characteristic: p-convergence is a conservative extension of m-convergence.

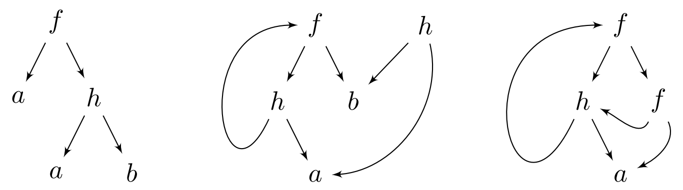

(a) $f(a, h(a, b))$. (b) A graph. (c) A term graph.

Figure 1: Example for a tree representation of a term; generalisation to (term) graphs.

Theorem 2.5 (total p-convergence = m-convergence)**.** *For every reduction S in a TRS the following equivalences hold:*

(i) $S\colon s \hookrightarrow_{p} t$ is total iff $S\colon s \hookrightarrow_{m} t$, and (ii) $S\colon s \twoheadrightarrow_{p} t$ is total iff $S\colon s \twoheadrightarrow_{m} t$.

The same also holds for continuity instead of convergence.

Kennaway [12] and Bahr [4] investigated abstract models of infinitary rewriting based on metric spaces resp. partially ordered sets. We will take these abstract models as a basis to formulate a theory of infinitary term graph reductions. The key question that we have to address is what an appropriate metric space resp. partial order on term graphs looks like.

3 Graphs and Term Graphs

This section provides the basic notions for term graphs and more generally for graphs. Terms over a signature, say Σ, can be thought of as rooted trees whose nodes are labelled with symbols from Σ. Moreover, in these trees a node labelled with a k-ary symbol is restricted to have out-degree k and the outgoing edges are ordered. In this way the i-th successor of a node labelled with a symbol f is interpreted as the root node of the subtree that represents the i-th argument of f. For example, consider the term $f(a, h(a, b))$. The corresponding representation as a tree is shown in Figure 1a.

In term graphs, the restriction to a tree structure is abolished. The notion of term graphs we are using is taken from Barendregt et al. [7].

Definition 3.1 (graph)**.** Let Σ be a signature. A *graph* over Σ is a tuple $g = (N, \mathsf{lab}, \mathsf{suc})$ consisting of a set N (of *nodes*), a *labelling function* $\mathsf{lab}\colon N \to \Sigma$, and a *successor function* $\mathsf{suc}\colon N \to N^*$ such that $|\mathsf{suc}(n)| = \mathsf{ar}(\mathsf{lab}(n))$ for each node $n \in N$, i.e. a node labelled with a k-ary symbol has precisely k successors. If $\mathsf{suc}(n) = \langle n_0, \dots, n_{k-1} \rangle$, then we write $\mathsf{suc}_i(n)$ for n_i. Moreover, we use the abbreviation $\mathsf{ar}_g(n)$ for the arity $\mathsf{ar}(\mathsf{lab}(n))$ of n.

Example 3.2. Let $\Sigma = \{f/2, h/2, a/0, b/0\}$ be a signature. The graph over Σ, depicted in Figure 1b, is given by the triple $(N, \mathsf{lab}, \mathsf{suc})$ with $N = \{n_0, n_1, n_2, n_3, n_4\}$, $\mathsf{lab}(n_0) = f, \mathsf{lab}(n_1) = \mathsf{lab}(n_4) = h, \mathsf{lab}(n_2) = b, \mathsf{lab}(n_3) = a$ and $\mathsf{suc}(n_0) = \langle n_1, n_2 \rangle, \mathsf{suc}(n_1) = \langle n_0, n_3 \rangle, \mathsf{suc}(n_2) = \mathsf{suc}(n_3) = \langle \rangle, \mathsf{suc}(n_4) = \langle n_2, n_3 \rangle$.

Definition 3.3 (path, reachability)**.** Let $g = (N, \mathsf{lab}, \mathsf{suc})$ be a graph and $n, n' \in N$.

(i) A *path* in g from n to n' is a finite sequence $(p_i)_{i<l}$ in $\mathbb{N}$ such that either

- $n = n'$ and $(p_i)_{i<l}$ is empty, i.e. $l = 0$, or

- $0 \leq p_0 < \mathsf{ar}_g(n)$ and the suffix $(p_i)_{1 \leq i < l}$ is a path in g from $\mathsf{suc}_{p_0}(n)$ to n'.

(ii) If there exists a path from n to n' in g, we say that n' is *reachable* from n in g.

Definition 3.4 (term graph)**.** Given a signature Σ, a *term graph* g over Σ is a tuple $(N, \mathsf{lab}, \mathsf{suc}, r)$ consisting of an *underlying* graph $(N, \mathsf{lab}, \mathsf{suc})$ over Σ whose nodes are all reachable from the *root node* $r \in N$. The class of all term graphs over Σ is denoted $\mathcal{G}^\infty(\Sigma)$. We use the notation N^g, lab^g, suc^g and r^g to refer to the respective components N,lab, suc and r of g. Given a graph or a term graph h and a node n in h, we write $h|_n$ to denote the sub-term graph of h rooted in g.

Example 3.5. Let $\Sigma = \{f/2, h/2, c/0\}$ be a signature. The term graph over Σ, depicted in Figure 1c, is given by the quadruple $(N, \mathsf{lab}, \mathsf{suc}, r)$, where $N = \{r, n_1, n_2, n_3\}$, $\mathsf{suc}(r) = \langle n_1, n_2\rangle$, $\mathsf{suc}(n_1) = \langle r, n_3\rangle$, $\mathsf{suc}(n_2) = \langle n_1, n_3\rangle$, $\mathsf{suc}(n_3) = \langle\rangle$ and $\mathsf{lab}(r) = \mathsf{lab}(n_2) = f$, $\mathsf{lab}(n_1) = h$, $\mathsf{lab}(n_3) = c$.

Paths in a graph are not absolute but relative to a starting node. In term graphs, however, we have a distinguished root node from which each node is reachable. Paths relative to the root node are central for dealing with term graphs:

Definition 3.6 (position, depth, cyclicity, tree)**.** Let $g \in \mathcal{G}^\infty(\Sigma)$ and $n \in N$.

(i) A *position* of n is a path in the underlying graph of g from r^g to n. The set of all positions in g is denoted $\mathcal{P}(g)$; the set of all positions of n in g is denoted $\mathcal{P}_g(n)$.[1]

(ii) The *depth* of n in g, denoted $\mathsf{depth}_g(n)$, is the minimum of the lengths of the positions of n in g, i.e. $\mathsf{depth}_g(n) = \min \{|\pi| \mid \pi \in \mathcal{P}_g(n)\}$.

(iii) For a position $\pi \in \mathcal{P}(g)$, we write $\mathsf{node}_g(\pi)$ for the unique node $n \in N^g$ with $\pi \in \mathcal{P}_g(n)$ and $g(\pi)$ for its symbol $\mathsf{lab}^g(n)$.

(iv) A position $\pi \in \mathcal{P}(g)$ is called *cyclic* if there are paths $\pi_1 < \pi_2 \leq \pi$ with $\mathsf{node}_g(\pi_1) = \mathsf{node}_g(\pi_2)$. The non-empty path π' with $\pi_1 \cdot \pi' = \pi_2$ is then called a *cycle* of $\mathsf{node}_g(\pi_1)$. A position that is not cyclic is called *acyclic*.

(v) The term graph g is called a *term tree* if each node in g has exactly one position.

Note that the labelling function of graphs – and thus term graphs – is *total*. In contrast, Barendregt et al. [7] considered *open* (term) graphs with a *partial* labelling function such that unlabelled nodes denote holes or variables. This is reflected in their notion of homomorphisms in which the homomorphism condition is suspended for unlabelled nodes.

3.1 Homomorphisms

Instead of a partial node labelling function, we chose a *syntactic* approach that is closer to the representation in terms: Variables, holes and "bottoms" are represented as distinguished syntactic entities. We achieve this on term graphs by making the notion of homomorphisms dependent on a distinguished set of constant symbols Δ for which the homomorphism condition is suspended:

[1] The notion/notation of positions is borrowed from terms: Every position π of a node n corresponds to the subterm represented by n occurring at position π in the unravelling of the term graph to a term.

Definition 3.7 (Δ-homomorphism)**.** Let Σ be a signature, $\Delta \subseteq \Sigma^{(0)}$, and $g, h \in \mathcal{G}^\infty(\Sigma)$.

(i) A function $\phi\colon N^g \to N^h$ is called *homomorphic* in $n \in N^g$ if the following holds:

$$\mathsf{lab}^g(n) = \mathsf{lab}^h(\phi(n)) \qquad \text{(labelling)}$$
$$\phi(\mathsf{suc}_i^g(n)) = \mathsf{suc}_i^h(\phi(n)) \quad \text{for all } 0 \le i < \mathsf{ar}_g(n) \qquad \text{(successor)}$$

(ii) A Δ*-homomorphism* ϕ from g to h, denoted $\phi\colon g \to_\Delta h$, is a function $\phi\colon N^g \to N^h$ that is homomorphic in n for all $n \in N^g$ with $\mathsf{lab}^g(n) \notin \Delta$ and satisfies $\phi(r^g) = r^h$.

It should be obvious that we get the usual notion of homomorphisms on term graphs if $\Delta = \emptyset$. The Δ-nodes can be thought of as holes in the term graphs which can be filled with other term graphs. For example, if we have a distinguished set of variable symbols $\mathcal{V} \subseteq \Sigma^{(0)}$, we can use $\mathcal{V}$-homomorphisms to formalise the matching step of term graph rewriting which requires the instantiation of variables.

Proposition 3.8 (Δ-homomorphism preorder)**.** *The Δ-homomorphisms on $\mathcal{G}^\infty(\Sigma)$ form a category which is a preorder. That is, there is at most one Δ-homomorphism from one term graph to another.*

Proof. The identity Δ-homomorphism is obviously the identity mapping on the set of nodes. Moreover, an easy equational reasoning reveals that the composition of two Δ-homomorphisms is again a Δ-homomorphism. Associativity of this composition is obvious as Δ-homomorphisms are functions.

In order to show that the category is a preorder assume that there are two Δ-homomorphisms $\phi_1, \phi_2\colon g \to_\Delta h$. We prove that $\phi_1 = \phi_2$ by showing that $\phi_1(n) = \phi_2(n)$ for all $n \in N^g$ by induction on the depth of n.

Let $\mathsf{depth}_g(n) = 0$, i.e. $n = r^g$. By the root condition, we have that $\phi_1(r^g) = r^h = \phi_2(r^g)$. Let $\mathsf{depth}_g(n) = d > 0$. Then n has a position $\pi \cdot \langle i \rangle$ in g such that $\mathsf{depth}_g(n') < d$ for $n' = \mathsf{node}_g(\pi)$. Hence, we can employ the induction hypothesis for n' to obtain the following:

$$\begin{aligned}
\phi_1(n) &= \mathsf{suc}_i^h(\phi_1(n')) && \text{(successor condition for } \phi_1\text{)}\\
&= \mathsf{suc}_i^h(\phi_2(n')) && \text{(ind. hyp.)}\\
&= \phi_2(n) && \text{(successor condition for } \phi_2\text{)}
\end{aligned}$$

□

As a consequence, each Δ-homomorphism is both monic and epic, and whenever there are two Δ-homomorphisms $\phi\colon g \to_\Delta h$ and $\psi\colon h \to_\Delta g$, they are inverses of each other, i.e. Δ*-isomorphisms*. If two term graphs are Δ*-isomorphic*, we write $g \cong_\Delta h$.

Note that injectivity is in general different from both being monic and the existence of left-inverses. The same holds for surjectivity and being epic resp. having right-inverses. However, each Δ-homomorphism is a Δ-isomorphism iff it is bijective.

For the two special cases $\Delta = \emptyset$ and $\Delta = \{\sigma\}$, we write $\phi\colon g \to h$ resp. $\phi\colon g \to_\sigma h$ instead of $\phi\colon g \to_\Delta h$ and call ϕ a homomorphism resp. σ-homomorphism. The same convention applies to Δ-isomorphisms.

Lemma 3.9 (homomorphisms are surjective)**.** *Every homomorphism $\phi\colon g \to h$, with $g, h \in \mathcal{G}^\infty(\Sigma)$, is surjective.*

Proof. Follows from an easy induction on the depth of the nodes in h. $\square$

Note that a bijective Δ-homomorphism is not necessarily a Δ-isomorphism. To realise this, consider two term graphs g, h, each with one node only. Let the node in g be labelled with a and the node in h with b then the only possible a-homomorphism from g to h is clearly a bijection but not an a-isomorphism. On the other hand, bijective homomorphisms are isomorphisms.

Lemma 3.10 (bijective homomorphisms are isomorphisms)**.** *Let $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\phi\colon g \to h$. Then the following are equivalent*

(a) ϕ is an isomorphism.

(b) ϕ is bijective.

(c) ϕ is injective.

Proof. The implication (a) $\Rightarrow$ (b) is trivial. The equivalence (b) $\Leftrightarrow$ (c) follows from Lemma 3.9. For the implication (b) $\Rightarrow$ (a), consider the inverse ϕ^{-1} of ϕ. We need to show that ϕ^{-1} is a homomorphism from h to g. The root condition follows immediately from the root condition for ϕ. Similarly, an easy equational reasoning reveals that the fact that ϕ is homomorphic in N^g implies that ϕ^{-1} is homomorphic in N^h $\square$

3.2 Canonical Term Graphs

In this section, we introduce a canonical representation of isomorphism classes of term graphs. We use a well-known trick to achieve this [18]. As we shall see at the end of this section, this will also enable us to construct term graphs modulo isomorphism very easily.

Definition 3.11 (canonical term graph)**.** A term graph g is called *canonical* if $n = \mathcal{P}_g(n)$ holds for each $n \in N^g$. That is, each node is the set of its positions in the term graph. The set of all canonical term graphs over Σ is denoted $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$.

This structure allows a convenient characterisation of Δ-homomorphisms:

Lemma 3.12 (characterisation of Δ-homomorphisms)**.** *For $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$, a function $\phi\colon N^g \to N^h$ is a Δ-homomorphism $\phi\colon g \to_\Delta h$ iff the following holds for all $n \in N^g$:*

$$\text{(a) } n \subseteq \phi(n), \quad \text{and} \qquad \text{(b) } \mathsf{lab}^g(n) = \mathsf{lab}^h(\phi(n)) \quad \text{whenever} \quad \mathsf{lab}^g(n) \notin \Delta.$$

Proof. For the "only if" direction, assume that $\phi\colon g \to_\Delta h$. (b) is the labelling condition and is therefore satisfied by ϕ. To establish (a), we show the equivalent statement

$$\forall \pi \in \mathcal{P}(g).\ \forall n \in N^g.\ \pi \in n \implies \pi \in \phi(n)$$

We do so by induction on the length of π: If $\pi = \langle\rangle$, then $\pi \in n$ implies $n = r^g$. By the root condition, we have $\phi(r^g) = r^h$ and, therefore, $\pi = \langle\rangle \in r^h$. If $\pi = \pi' \cdot \langle i\rangle$, then let $n' = \mathsf{node}_g(\pi')$. Consequently, $\pi' \in n'$ and, by induction hypothesis, $\pi' \in \phi(n')$. Since $\pi = \pi' \cdot \langle i\rangle$, we have $\mathsf{suc}^g_i(n') = n$. By the successor condition we can conclude $\phi(n) = \mathsf{suc}^h_i(\phi(n'))$. This and $\pi' \in \phi(n')$ yields that $\pi' \cdot \langle i\rangle \in \phi(n)$.

For the "if" direction, we assume (a) and (b). The labelling condition follows immediately from (b). For the root condition, observe that since $\langle\rangle \in r^g$, we also have $\langle\rangle \in \phi(r^g)$. Hence, $\phi(r^g) = r^h$. In order to show the successor condition, let $n, n' \in N^g$ and $0 \le i < \mathsf{ar}_g(n)$ such that $\mathsf{suc}^g_i(n) = n'$. Then there is a position $\pi \in n$ with $\pi \cdot \langle i\rangle \in n'$. By (a), we can conclude that $\pi \in \phi(n)$ and $\pi \cdot \langle i\rangle \in \phi(n')$ which implies that $\mathsf{suc}^h_i(\phi(n)) = \phi(n')$. $\square$

By Proposition 3.8, there is at most one Δ-homomorphism between two term graphs. The lemma above uniquely defines this Δ-homomorphism: If there is a Δ-homomorphism from g to h, it is defined by $\phi(n) = n'$, where n' is the unique node $n' \in N^h$ with $n \subseteq n'$.

Remark 3.13. Note that the lemma above is also applicable to non-canonical term graphs. It simply has to be rephrased such that instead of just referring to a node n, its set of positions $\mathcal{P}_g(n)$ is referred to whenever the "inner structure" of n is used.

The set of nodes in a canonical term graph forms a partition of the set of positions. Hence, it defines an equivalence relation on the set of positions. For a canonical term graph g, we write $\sim_g$ for this equivalence relation on $\mathcal{P}(g)$. According to Remark 3.13, we can extend this to arbitrary term graphs: $\pi_1 \sim_g \pi_2$ iff $\mathsf{node}_g(\pi_1) = \mathsf{node}_g(\pi_2)$. The characterisation of Δ-homomorphisms can thus be recast to obtain the following lemma that characterises the *existence* of Δ-homomorphisms:

Lemma 3.14 (characterisation of Δ-homomorphisms). *Given $g, h \in \mathcal{G}^\infty(\Sigma)$, there is a Δ-homomorphism $\phi\colon g \to_\Delta h$ iff, for all $\pi, \pi' \in \mathcal{P}(g)$, we have*

$$\text{(a) } \pi \sim_g \pi' \quad \Longrightarrow \quad \pi \sim_h \pi', \text{ and} \qquad \text{(b) } g(\pi) = h(\pi) \quad \text{whenever} \quad g(\pi) \notin \Delta.$$

Proof. W.l.o.g. we assume g and h to be canonical. For the "only if" direction, assume that ϕ is a Δ-homomorphism from g to h. Then we can use the properties (a) and (b) of Lemma 3.12, which we will refer to as (a') and (b') to avoid confusion. In order to show (a), assume $\pi \sim_g \pi'$. Then there is some node $n \in N^g$ with $\pi, \pi' \in n$. (a') yields $\pi, \pi' \in \phi(n)$ and, therefore, $\pi \sim_g \pi'$. To show (b), we assume some $\pi \in \mathcal{P}(g)$ with $g(\pi) \notin \Delta$. Then we can reason as follows:

$$g(\pi) = \mathsf{lab}^g(\mathsf{node}_g(\pi)) \overset{\text{(b')}}{=} \mathsf{lab}^h(\phi(\mathsf{node}_g(\pi))) \overset{\text{(a')}}{=} \mathsf{lab}^h(\mathsf{node}_h(\pi)) = h(\pi)$$

For the converse direction, assume that both (a) and (b) hold. Define the function $\phi\colon N^g \to N^h$ by $\phi(n) = n'$ iff $n \subseteq n'$ for all $n \in N^g$ and $n' \in N^h$. To see that this is well-defined, we show at first that, for each $n \in N^g$, there is at most one $n' \in N^h$ with $n \subseteq n'$. Suppose there is another node $n'' \in N^h$ with $n \subseteq n''$. Since $n \neq \emptyset$, this implies $n' \cap n'' \neq \emptyset$. Hence, $n' = n''$. Secondly, we show that there is at least one such node n'. Choose some $\pi^* \in n$. Since then $\pi^* \sim_g \pi^*$ and, by (a), also $\pi^* \sim_h \pi^*$ holds, there is some $n' \in N^h$ with $\pi^* \in n'$. For each $\pi \in n$, we have $\pi^* \sim_g \pi$ and, therefore, $\pi^* \sim_h \pi$ by (a). Hence, $\pi \in n'$. So we know that ϕ is well-defined. By construction, ϕ satisfies (a'). Moreover, because of (b), it is also easily seen to satisfy (b'). Hence, ϕ is a homomorphism from g to h. $\square$

Intuitively, (a) means that h has at least as much sharing of nodes as g has, whereas (b) means that h has at least the same non-Δ-symbols as g.

Corollary 3.15 (characterisation of Δ-isomorphisms). *Given $g, h \in \mathcal{G}^\infty(\Sigma)$, the following holds:*

(i) $\phi\colon N^g \to N^h$ *is a Δ-isomorphism* *iff* *for all* $n \in N^g$
(a) $\mathcal{P}_h(\phi(n)) = \mathcal{P}_g(n)$, *and* *(b)* $\mathsf{lab}^g(n) = \mathsf{lab}^h(\phi(n))$ *or* $\mathsf{lab}^g(n), \mathsf{lab}^h(\phi(n)) \in \Delta$.

(ii) $g \cong_\Delta h$ *iff* *(a)* $\sim_g \, = \, \sim_h$, *and* *(b)* $g(\pi) = h(\pi)$ *or* $g(\pi), h(\pi) \in \Delta$.

Proof. Immediate consequence of Lemma 3.12 resp. Lemma 3.14 and Proposition 3.8. $\square$

From (ii) we immediately obtain the following equivalence:

Corollary 3.16. *Given $g, h \in \mathcal{G}^\infty(\Sigma)$ and $\sigma \in \Sigma^{(0)}$, we have $g \cong h$ iff $g \cong_\sigma h$.*

Now we can revisit the notion of canonical term graphs using the above characterisation of Δ-isomorphisms. We will define a function $\mathcal{C}(\cdot)\colon \mathcal{G}^\infty(\Sigma) \to \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ that maps a term graph to its canonical representation. To this end, let $g = (N, \mathsf{lab}, \mathsf{suc}, r)$ be a term graph and define $\mathcal{C}(g) = (N', \mathsf{lab}', \mathsf{suc}', r')$ as follows:

$$N' = \{\mathcal{P}_g(n) \mid n \in N\} \qquad r' = \mathcal{P}_g(r)$$
$$\mathsf{lab}'(\mathcal{P}_g(n)) = \mathsf{lab}(n) \qquad \mathsf{suc}'_i(\mathcal{P}_g(n)) = \mathcal{P}_g(\mathsf{suc}_i(n)) \qquad \text{for all } n \in N, 0 \leq i < \mathsf{ar}_g(n)$$

$\mathcal{C}(g)$ is obviously a well-defined canonical term graph. With this definition we indeed capture the idea of a canonical representation of isomorphism classes:

Proposition 3.17 (canonical partial term graphs are a canonical representation)**.** *Given $g \in \mathcal{G}^\infty(\Sigma)$, the term graph $\mathcal{C}(g)$ canonically represents the equivalence class $[g]_\cong$. More precisely, it holds that*

$$\textit{(i) } [g]_\cong = [\mathcal{C}(g)]_\cong, \textit{ and} \qquad \textit{(ii) } [g]_\cong = [h]_\cong \quad \textit{iff} \quad \mathcal{C}(g) = \mathcal{C}(h).$$

In particular, we have, for all canonical term graphs g, h, that $g = h$ iff $g \cong h$.

Proof. Straightforward consequence of Corollary 3.15. □

Remark 3.18. Δ-homomorphisms can be naturally lifted to $\mathcal{G}^\infty(\Sigma)/_\cong$: We say that two Δ-homomorphisms $\phi\colon g \to_\Delta h$, $\phi'\colon g' \to_\Delta h'$, are isomorphic, written $\phi \cong \phi'$ iff there are isomorphisms $\psi_1\colon g \cong g'$ and $\psi_2\colon h' \cong h$ such that $\phi = \psi_2 \circ \phi' \circ \psi_1$. Given a Δ-homomorphism $\phi\colon g \to_\Delta h$ in $\mathcal{G}^\infty(\Sigma)$, $[\phi]_\cong\colon [g]_\cong \to_\Delta [h]_\cong$ is a Δ-homomorphism in $\mathcal{G}^\infty(\Sigma)/_\cong$. These Δ-homomorphisms then form a category which can easily be show to be isomorphic to the category of Δ-homomorphisms on $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ via the mapping $[\cdot]_\cong$.

Corollary 3.15 has shown that term graphs can be characterised up to isomorphism by only giving the equivalence $\sim_g$ and the labelling $g(\cdot)\colon \pi \mapsto g(\pi)$. This observation gives rise to the following definition:

Definition 3.19 (labelled quotient tree)**.** A *labelled quotient tree* over signature Σ is a triple $(P, l, \sim)$ consisting of a non-empty set $P \subseteq \mathbb{N}^*$, a function $l\colon P \to \Sigma$, and an equivalence relation $\sim$ on P that satisfies the following conditions for all $\pi, \pi' \in P$ and $i \in \mathbb{N}$:

$$\pi \cdot \langle i \rangle \in P \quad \Longrightarrow \quad \pi \in P \quad \text{and} \quad i < \mathsf{ar}(l(\pi)) \qquad \text{(reachability)}$$
$$\pi \sim \pi' \quad \Longrightarrow \quad \begin{cases} l(\pi) = l(\pi') & \text{and} \\ \pi \cdot \langle j \rangle \sim \pi' \cdot \langle j \rangle & \text{for all } \ j < \mathsf{ar}(l(\pi)) \end{cases} \qquad \text{(congruence)}$$

The following lemma confirms that labelled quotient trees uniquely characterise any term graph up to isomorphism:

Lemma 3.20. *Each term graph $g \in \mathcal{G}^\infty(\Sigma)$ induces a* canonical labelled quotient tree *$(\mathcal{P}(g), g(\cdot), \sim_g)$ over Σ. Vice versa, for each labelled quotient tree $(P, l, \sim)$ over Σ there is a unique canonical term graph $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ whose canonical labelled quotient tree is $(P, l, \sim)$, i.e. $\mathcal{P}(g) = P$, $g(\pi) = l(\pi)$ for all $\pi \in P$, and $\sim_g = \sim$.*

Proof. The first part is trivial: $(\mathcal{P}(g), g(\cdot), \sim_g)$ satisfies the conditions from Definition 3.19.

Let $(P, l, \sim)$ be a labelled quotient tree. Define the term graph $g = (N, \mathsf{lab}, \mathsf{suc}, r)$ by

$$
\begin{aligned}
N &= P/_{\sim} & \mathsf{lab}(n) &= f \quad \text{iff} \quad \exists \pi \in n.\, l(\pi) = f \\
r = n \quad &\text{iff} \quad \langle\rangle \in n & \mathsf{suc}_i(n) &= n' \quad \text{iff} \quad \exists \pi \in n.\, \pi \cdot \langle i \rangle \in n'
\end{aligned}
$$

The functions lab and suc are well-defined due to the congruence condition satisfied by $(P, l, \sim)$. Since P is non-empty and closed under prefixes, it contains $\langle\rangle$. Hence, r is well-defined. Moreover, by the reachability condition, each node in N is reachable from the root node. An easy induction proof shows that $\mathcal{P}_g(n) = n$ for each node $n \in N$. Thus, g is a well-defined canonical term graph. The canonical labelled quotient tree of g is obviously $(P, l, \sim)$. Whenever there are two canonical term graphs with labelled quotient tree $(P, l, \sim)$, they are isomorphic due to Corollary 3.15 and, therefore, have to be identical by Proposition 3.17. $\square$

Labelled quotient trees provide a valuable tool for constructing canonical term graphs. Nevertheless, the original graph representation remains convenient for practical purposes as it allows a straightforward formalisation of term graph rewriting and provides a finite representation of finite cyclic term graphs which induce an infinite labelled quotient tree.

Before we continue, it is instructive to make the correspondence between terms and term graphs clear. Note, that there is an obvious one-to-one correspondence between canonical term *trees* and terms. For example, the term tree depicted in Figure 1a corresponds to the term $f(a, h(a, b))$. We thus consider the set of terms $\mathcal{T}^\infty(\Sigma)$ to be the subset of canonical term trees of $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$.

With this correspondence in mind, we can define the *unravelling* of a term graph g as the unique term t such that there is a homomorphism $\phi\colon\, t \to g$. The unravelling of cyclic term graphs yields infinite terms, e.g. in Figure 5 on page 30, the term h_ω is the unravelling of the term graph g_2. We use the notation $\mathcal{U}(g)$ for the unravelling of g.

Another convenience for dealing with term graphs is a linear notation that makes it easy to write down (canonical) term graphs instead of using the formal definition or a drawing. The notation that we use is based on the linear notation for graphs by Barendregt et al.[7]:

Definition 3.21. Let Σ be a signature, $\mathcal{N}$ a countably infinite set (of names) disjoint from Σ and $\widehat{\Sigma}$ a signature such $n \in \widehat{\Sigma}^{(0)}$ and $f, f^n \in \widehat{\Sigma}^{(k)}$ for each $n \in \mathcal{N}$, $k \in \mathbb{N}$ and $f \in \Sigma^{(k)}$. A *linear notation for a canonical term graph* in $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ is a term $t \in \mathcal{T}^\infty(\widehat{\Sigma})$ such that for each $n \in \mathcal{N}$ that occurs in t, there is exactly one occurrence of a function symbol of the form $f^{[n]}$ in t.

For each such linear notation t we define the corresponding canonical term graph g as follows: Consider the term tree representation of t with the root node r. Redirect every edge to a node labelled n to the unique node labelled $f^{[n]}$. Then, change all labellings of the form f^n to f. After removing all nodes not reachable from the node r, define g as the canonical term graph of the thus obtained term graph rooted in r.

We use n, m and primed resp. indexed variants thereof to denote names in $\mathcal{N}$.

Intuitively, in a linear notation for a term graph, a subterm n denotes a pointer to a subterm with the corresponding name n, i.e. a subterm of the form $f^n(t_1, \dots, t_k)$.

Example 3.22. Consider the term graph in Figure 1c. This term graph can be described by the linear notation $f^{[n_1]}(h^{[n_2]}(n_1, c^{[n_3]}), f(n_2, n_3))$. On the other hand, $f^{[n_1]}(n_1, n_2)$ and $f(a^{[n]}, b^{[n]})$ are not valid linear notations.

Note that every term $t \in \mathcal{T}^\infty(\Sigma)$ is a linear notation for the corresponding term tree in $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$.

4 Partial Order on Term Graphs

In this section, we want to establish a partial order suitable for formalising convergence of sequences of canonical term graphs similarly to *p*-convergence on terms.

In previous work, we have studied several different partial orders on term graphs and the notion of convergence they induce [6]. All of these partial orders have in common that they are based on $\perp$-homomorphisms. This approach is founded on the observation that if we consider terms as term trees, then $\perp$-homomorphisms characterise the partial order on terms:

$$s \leq_\perp t \iff \text{there is a } \perp\text{-homomorphism } \phi\colon\ s \to_\perp t.$$

Thus $\perp$-homomorphisms constitute the ideal tool to define a partial order on partial term graphs, i.e. term graphs over the signature $\Sigma_\perp = \Sigma \uplus \{\perp\}$.

In this paper, we focus on the simplest among these partial orders on term graphs:

Definition 4.1. The relation $\leq^\mathcal{G}_\perp$ on $\mathcal{G}^\infty(\Sigma_\perp)$ is defined as follows: $g \leq^\mathcal{G}_\perp h$ iff there is a $\perp$-homomorphism $\phi\colon\ g \to_\perp h$.

Proposition 4.2 (partial order $\leq^\mathcal{G}_\perp$)**.** *The relation $\leq^\mathcal{G}_\perp$ is a partial order on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$.*

Proof. Transitivity and reflexivity of $\leq^\mathcal{G}_\perp$ follows immediately from Proposition 3.8. For antisymmetry, consider $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$ with $g \leq^\mathcal{G}_\perp h$ and $h \leq^\mathcal{G}_\perp g$. Then, by Proposition 3.8, $g \cong_\perp h$. This is equivalent to $g \cong h$ by Corollary 3.16 from which we can conclude $g = h$ using Proposition 3.17. $\square$

In our previous attempts to formalise convergence on term graphs [6], this partial order was rejected as the induced notion of convergence manifests some unintuitive behaviour. However, as we will show in Section 7.4, theses quirks will vanish when we move to strong convergence.

Before we study the properties of the partial order $\leq^\mathcal{G}_\perp$, it is helpful to make its characterisation in terms of labelled quotient trees explicit:

Corollary 4.3 (characterisation of $\leq^\mathcal{G}_\perp$)**.** *Let $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$. Then $g \leq^\mathcal{G}_\perp h$ iff the following conditions are met:*

(a) $\pi \sim_g \pi' \quad \implies \quad \pi \sim_h \pi' \quad$ *for all* $\pi, \pi' \in \mathcal{P}(g)$

(b) $g(\pi) = h(\pi) \quad$ *for all* $\pi \in \mathcal{P}(g)$ *with* $g(\pi) \in \Sigma$.

Proof. This follows immediately from Lemma 3.14. $\square$

Note that the partial order $\leq_\perp$ on terms is entirely characterised by (b). That is, the partial order $\leq^\mathcal{G}_\perp$ is simply the partial order $\leq_\perp$ on its underlying tree structure (i.e. its unravelling) plus the preservation of sharing as stipulated by (a).

Next, we will show that the partial order on term graphs has the properties that make it suitable as a basis for *p*-convergence, i.e. that it forms a complete semilattice. At first we show its cpo structure:

Theorem 4.4. *The relation $\leq^{\mathcal{G}}_{\bot}$ is a complete partial order on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$. In particular, it has the least element $\bot$, and the least upper bound of a directed set G is given by the following labelled quotient tree $(P, l, \sim)$:*

$$P = \bigcup_{g \in G} \mathcal{P}(g) \qquad\quad \sim\, = \bigcup_{g \in G} \sim_g \qquad\quad l(\pi) = \begin{cases} f & \text{if } f \in \Sigma \text{ and } \exists g \in G.\ g(\pi) = f \\ \bot & \text{otherwise} \end{cases}$$

Proof. The least element of $\leq^{\mathcal{G}}_{\bot}$ is obviously $\bot$. Hence, it remains to be shown that each each directed subset of $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$ has a least upper bound. To this end, suppose that G is a directed subset of $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$. We define a canonical term graph $\overline{g}$ by giving the labelled quotient tree $(P, l, \sim)$

In order to show that the canonical term graph $\overline{g}$ given by the labelled quotient tree $(P, l, \sim)$ above is indeed the lub of G, we will make extensive use of Corollary 4.3. Therefore, we use (a) and (b) to refer to the conditions mentioned there.

At first we need to show that l is indeed well-defined. For this purpose, let $g_1, g_2 \in G$ and $\pi \in \mathcal{P}(g_1) \cap \mathcal{P}(g_2)$ with $g_1(\pi), g_2(\pi) \in \Sigma$. Since G is directed, there is some $g \in G$ such that $g_1, g_2 \leq^{\mathcal{G}}_{\bot} g$. By (b), we can conclude $g_1(\pi) = g(\pi) = g_2(\pi)$.

Next we show that $(P, l, \sim)$ is indeed a labelled quotient tree. Recall that $\sim$ needs to be an equivalence relation. For the reflexivity, assume that $\pi \in P$. Then there is some $g \in G$ with $\pi \in \mathcal{P}(g)$. Since $\sim_g$ is an equivalence relation, $\pi \sim_g \pi$ must hold and, therefore, $\pi \sim \pi$. For the symmetry, assume that $\pi_1 \sim \pi_2$. Then there is some $g \in G$ such that $\pi_1 \sim_g \pi_2$. Hence, we get $\pi_2 \sim_g \pi_1$ and, consequently, $\pi_2 \sim \pi_1$. In order to show transitivity, assume that $\pi_1 \sim \pi_2, \pi_2 \sim \pi_3$. That is, there are $g_1, g_2 \in G$ with $\pi_1 \sim_{g_1} \pi_2$ and $\pi_2 \sim_{g_2} \pi_3$. Since G is directed, we find some $g \in G$ such that $g_1, g_2 \leq^{\mathcal{G}}_{\bot} g$. By (a), this implies that also $\pi_1 \sim_g \pi_2$ and $\pi_2 \sim_g \pi_3$. Hence, $\pi_1 \sim_g \pi_3$ and, therefore, $\pi_1 \sim \pi_3$.

For the reachability condition, let $\pi \cdot \langle i \rangle \in P$. That is, there is a $g \in G$ with $\pi \cdot \langle i \rangle \in \mathcal{P}(g)$. Hence, $\pi \in \mathcal{P}(g)$, which in turn implies $\pi \in P$. Moreover, $\pi \cdot \langle i \rangle \in \mathcal{P}(g)$ implies that $i < \mathsf{ar}(g(\pi))$. Since $g(\pi)$ cannot be a nullary symbol and in particular not $\bot$, we obtain that $l(\pi) = g(\pi)$. Hence, $i < \mathsf{ar}(l(\pi))$.

For the congruence condition, assume that $\pi_1 \sim \pi_2$ and that $l(\pi_1) = f$. If $f \in \Sigma$, then there are $g_1, g_2 \in G$ with $\pi_1 \sim_{g_1} \pi_2$ and $g_2(\pi_1) = f$. Since G is directed, there is some $g \in G$ such that $g_1, g_2 \leq^{\mathcal{G}}_{\bot} g$. Hence, by (a) resp. (b), we have $\pi_1 \sim_g \pi_2$ and $g(\pi_1) = f$. Using Lemma 3.20 we can conclude that $g(\pi_2) = g(\pi_1) = f$ and that $\pi_1 \cdot \langle i \rangle \sim_g \pi_2 \cdot \langle i \rangle$ for all $i < \mathsf{ar}(g(\pi_1))$. Because $g \in G$, it holds that $l(\pi_2) = f$ and that $\pi_1 \cdot \langle i \rangle \sim \pi \cdot \langle i \rangle$ for all $i < \mathsf{ar}(l(\pi_1))$. If $f = \bot$, then also $l(\pi_2) = \bot$, for if $l(\pi_2) = f'$ for some $f' \in \Sigma$, then, by the symmetry of $\sim$ and the above argument (for the case $f \in \Sigma$), we would obtain $f = f'$ and, therefore, a contradiction. Since $\bot$ is a nullary symbol, the remainder of the condition is vacuously satisfied.

This shows that $(P, l, \sim)$ is a labelled quotient tree which, by Lemma 3.20, uniquely defines a canonical term graph. In order to show that the thus obtained term graph $\overline{g}$ is an upper bound for G, we have to show that $g \leq^{\mathcal{G}}_{\bot} \overline{g}$ by establishing (a) and (b). This is an immediate consequence of the construction.

In the final part of this proof, we will show that $\overline{g}$ is the least upper bound of G. For this purpose, let $\hat{g}$ be an upper bound of G, i.e. $g \leq^{\mathcal{G}}_{\bot} \hat{g}$ for all $g \in G$. We will show that $\overline{g} \leq^{\mathcal{G}}_{\bot} \hat{g}$ by establishing (a) and (b). For (a), assume that $\pi_1 \sim \pi_2$. Hence, there is some $g \in G$ with $\pi_1 \sim_g \pi_2$. Since, by assumption, $g \leq^{\mathcal{G}}_{\bot} \widehat{g}$, we can conclude $\pi_1 \sim_{\widehat{g}} \pi_2$ using (a). For (b), assume $\pi \in P$ and $l(\pi) = f \in \Sigma$. Then there is some $g \in G$ with $g(\pi) = f$. Applying (b) then yields $\widehat{g}(\pi) = f$ since $g \leq^{\mathcal{G}}_{\bot} \widehat{g}$. □

The following proposition shows that the partial order $\leq_{\bot}^{\mathcal{G}}$ also admits glbs of arbitrary non-empty sets:

Proposition 4.5. *In the partially ordered set $(\mathcal{G}_{\mathcal{C}}^{\infty}(\Sigma_\bot), \leq_{\bot}^{\mathcal{G}})$ every non-empty set has a glb. In particular, the glb of a non-empty set G is given by the following labelled quotient tree $(P, l, \sim)$:*

$$P = \left\{ \pi \in \bigcap_{g \in G} \mathcal{P}(g) \;\middle|\; \forall \pi' < \pi \exists f \in \Sigma_\bot \forall g \in G : g(\pi') = f \right\}$$

$$l(\pi) = \begin{cases} f & \text{if } \forall g \in G : f = g(\pi) \\ \bot & \text{otherwise} \end{cases} \qquad\qquad \sim = \bigcap_{g \in G} \sim_g \cap P \times P$$

Proof. At first we need to prove that $(P, l, \sim)$ is in fact a well-defined labelled quotient tree. That $\sim$ is an equivalence relation follows straightforwardly from the fact that each $\sim_g$ is an equivalence relation.

Next, we show the reachability and congruence properties from Definition 3.19. In order to show the reachability property, assume some $\pi \cdot \langle i \rangle \in P$. Then, for each $\pi' \leq \pi$ there is some $f_{\pi'} \in \Sigma_\bot$ such that $g(\pi') = f_{\pi'}$ for all $g \in G$. Hence, $\pi \in P$. Moreover, we have in particular that $i < \mathsf{ar}(f_\pi) = \mathsf{ar}(l(\pi))$.

For the congruence condition, assume that $\pi_1 \sim \pi_2$. Hence, $\pi_1 \sim_g \pi_2$ for all $g \in G$. Consequently, we have for each $g \in G$ that $g(\pi_1) = g(\pi_2)$ and that $\pi_1 \cdot \langle i \rangle \sim_g \pi_2 \cdot \langle i \rangle$ for all $i < \mathsf{ar}(g(\pi_1))$. We distinguish two cases: At first assume that there are some $g_1, g_2 \in G$ with $g_1(\pi_1) \neq g_2(\pi_1)$. Hence, $l(\pi_2) = \bot$. Since, we also have that $g_1(\pi_2) = g_1(\pi_1) \neq g_2(\pi_1) = g_2(\pi_2)$ we can conclude that $l(\pi_2) = \bot = l(\pi_1)$. Since $\mathsf{ar}(\bot) = 0$ we are done for this case. Next, consider the alternative case that there is some $f \in \Sigma_\bot$ such that $g(\pi_1) = f$ for all $g \in G$. Consequently, $l(\pi_1) = f$ and since also $g(\pi_2) = g(\pi_1) = f$ for all $g \in G$, we can conclude that $l(\pi_2) = f = l(\pi_1)$. Moreover, we obtain from the initial assumption for this case, that $\pi_1 \cdot \langle i \rangle, \pi_2 \cdot \langle i \rangle \in P$ for all $i < \mathsf{ar}(f)$ which implies that $\pi_1 \cdot \langle i \rangle \sim \pi_2 \cdot \langle i \rangle$ for all $i < \mathsf{ar}(f) = \mathsf{ar}(l(\pi_1))$.

Next, we show that the term graph $\overline{g}$ defined by $(P, l, \sim)$ is a lower bound of G, i.e. that $\overline{g} \leq_{\bot}^{\mathcal{G}} g$ for all $g \in G$. By Lemma 3.14, it suffices to show $\sim \cap P \times P \subseteq \sim_g$ and $l(\pi) = g(\pi)$ for all $\pi \in P$ with $l(\pi) \in \Sigma$. Both conditions follow immediately from the construction of $\overline{g}$.

Finally, we show that $\overline{g}$ is the greatest lower bound of G. To this end, let $\widehat{g} \in \mathcal{G}_{\mathcal{C}}^{\infty}(\Sigma_\bot)$ with $\widehat{g} \leq_{\bot}^{\mathcal{G}} g$ for each $g \in G$. We will show that then $\widehat{g} \leq_{\bot}^{\mathcal{G}} \overline{g}$ using Lemma 3.14. At first we show that $\mathcal{P}(\widehat{g}) \subseteq P$. Let $\pi \in \mathcal{P}(\widehat{g})$. We know that $\widehat{g}(\pi') \in \Sigma$ for all $\pi' < \pi$. According to Lemma 3.14, using the assumption that $\widehat{g} \leq_{\bot}^{\mathcal{G}} g$ for all $g \in G$, we obtain that $g(\pi') = \widehat{g}(\pi')$ for all $\pi' < \pi$. Consequently, $\pi \in P$. Next, we show part (a) of Lemma 3.14. Let $\pi_1, \pi_2 \in \mathcal{P}(\widehat{g}) \subseteq P$ with $\pi_1 \sim_{\widehat{g}} \pi_2$. Hence, using the assumption that $\widehat{g}$ is a lower bound of G, we have $\pi_1 \sim_g \pi_2$ for all $g \in G$ according to Lemma 3.14. Consequently, $\pi_1 \sim \pi_2$. For part (b) of Lemma 3.14 let $\pi \in \mathcal{P}(\widehat{g}) \subseteq P$ with $\widehat{g}(\pi) = f \in \Sigma$. Using Lemma 3.14, we obtain that $g(\pi) = f$ for all $g \in G$. Hence, $l(\pi) = f$. ☐

From this we can immediately derive the complete semilattice structure of $\leq_{\bot}^{\mathcal{G}}$:

Theorem 4.6. *The partially ordered set $(\mathcal{G}_{\mathcal{C}}^{\infty}(\Sigma_\bot), \leq_{\bot}^{\mathcal{G}})$ forms a complete semilattice.*

Proof. Follows from Theorem 4.4 and Proposition 4.5. ☐

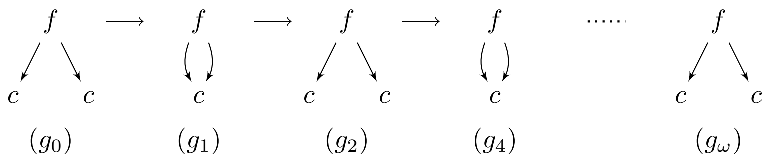

Figure 2: Limit inferior in the presence of acyclic sharing.

In particular, this means that the limit inferior is defined for every sequence of term graphs. Moreover, from the constructions given in Theorem 4.4 and Proposition 4.5, we can derive the following direct construction of the limit inferior:

Corollary 4.7. *The limit inferior of a sequence $(g_\iota)_{\iota<\alpha}$ over $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$ is given by the following labelled quotient tree $(P, \sim, l)$:*

$$P = \bigcup_{\beta<\alpha} \{\pi \in \mathcal{P}(g_\beta) \mid \forall \pi' < \pi \forall \beta \leq \iota < \alpha \colon g_\iota(\pi') = g_\beta(\pi')\}$$

$$\sim \; = \left(\bigcup_{\beta<\alpha} \bigcap_{\beta\leq\iota<\alpha} \sim_{g_\iota} \right) \cap P \times P$$

$$l(\pi) = \begin{cases} g_\beta(\pi) & \text{if } \exists \beta < \alpha \forall \beta \leq \iota < \alpha \colon g_\iota(\pi) = g_\beta(\pi) \\ \perp & \text{otherwise} \end{cases} \qquad \text{for all } \pi \in P$$

In particular, given $\beta < \alpha$ and $\pi \in \mathcal{P}(g_\beta)$, we have that $g(\pi) = g_\beta(\pi)$ if $g_\iota(\pi') = g_\beta(\pi')$ for all $\pi' \leq \pi$ and $\beta \leq \iota < \alpha$.

Example 4.8. Figure 5c on page 30 illustrates a sequence of term graphs $(h_\iota)_{\iota<\omega}$. Except for the edge to the root that closes a cycle each term graph h_ι as a tree structure. Since this edge is pushed down as the sequence progresses, it vanishes in the the limit inferior of $(h_\iota)_{\iota<\omega}$, depicted as h_ω in Figure 5c.

Changing acyclic sharing on the other hand exposes an oddity of the partial order $\leq^\mathcal{G}_\perp$. Let $(g_\iota)_{\iota<\omega}$ be the sequence of term graphs illustrated in Figure 2. The sequence alternates between g_0 and g_1 which differ only in the sharing of the two arguments of the f function symbol. Hence, there is an obvious homomorphism from g_0 to g_1 and we thus have $g_0 \leq^\mathcal{G}_\perp g_1$. Therefore, g_0 is the greatest lower bound of every suffix of $(g_\iota)_{\iota<\omega}$, which means that $\liminf_{\iota\to\omega} g_\iota = g_0$.

5 Metric Spaces

In this section, we shall define a metric space on canonical term graphs. We base our approach to defining a metric distance on the definition of the metric distance $\mathbf{d}$ on terms.

Originally, Arnold and Nivat [3] used a truncation $t \restriction d$ of terms to define the metric on terms. The truncation of a term t at depth d replaces all subterms at depth d by $\perp$:

$$t \restriction 0 = \perp, \quad f(t_1, \ldots, t_k) \restriction d + 1 = f(t_1 \restriction d, \ldots, t_k \restriction d), \quad t \restriction \infty = t$$

The similarity of two terms, on which the metric distance $\mathbf{d}$ is based, can thus be characterised via truncations:

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

We will adopt this approach for term graphs as well. To this end, we will first define abstractly what a truncation on term graphs is and how a metric distance can be derived from it. Then we show a concrete truncation and show that the induced metric space is in fact complete. We will conclude the section by showing that the metric space we considered is robust in the sense that it is invariant under small changes to the definition of truncation.

5.1 Truncation Functions

As we have seen above, the truncation on terms is a function that, depending on a depth value d, transforms a term t to a term $t{\upharpoonright}d$. We shall generalise this to term graphs and stipulate some axioms that ensure that we can derive a metric distance in the style of Arnold and Nivat [3]:

Definition 5.1 (truncation function)**.** A family $\tau = (\tau_d \colon \mathcal{G}^\infty(\Sigma_\perp) \to \mathcal{G}^\infty(\Sigma_\perp))_{d\in\mathbb{N}\cup\{\infty\}}$ of functions on term graphs is called a *truncation function* if it satisfies the following properties for all $g, h \in \mathcal{G}^\infty(\Sigma_\perp)$ and $d \in \mathbb{N} \cup \{\infty\}$:

(a) $\tau_0(g) \cong \perp$, (b) $\tau_\infty(g) \cong g$, and (c) $\tau_d(g) \cong \tau_d(h) \implies \tau_e(g) \cong \tau_e(h)$ for all $e < d$.

Note that from axioms (b) and (c) it follows that truncation functions must be defined modulo isomorphism, i.e. $g \cong h$ implies $\tau_d(g) \cong \tau_d(h)$ for all $d \in \mathbb{N} \cup \{\infty\}$.

Given a truncation function, we can define a distance measure in the style of Arnold and Nivat:

Definition 5.2 (truncation-based similarity/distance)**.** Let τ be a truncation function. The *τ-similarity* is the function $\mathsf{sim}_\tau \colon \mathcal{G}^\infty(\Sigma_\perp) \times \mathcal{G}^\infty(\Sigma_\perp) \to \mathbb{N} \cup \{\infty\}$ defined by

$$\mathsf{sim}_\tau(g,h) = \max\{d \in \mathbb{N} \cup \{\infty\} \mid \tau_d(g) \cong \tau_d(h)\}$$

The *τ-distance* is the function $\mathbf{d}_\tau \colon \mathcal{G}^\infty(\Sigma_\perp)\times\mathcal{G}^\infty(\Sigma_\perp) \to \mathbb{R}^+$ defined by $\mathbf{d}_\tau(g,h) = 2^{-\mathsf{sim}_\tau(g,h)}$, where $2^{-\infty}$ is interpreted as 0.

Observe, that the similarity $\mathsf{sim}_\tau(g,h)$ induced by a truncation function τ is well-defined since the axiom (a) of Definition 5.1 guarantees that the set $\{d \in \mathbb{N} \cup \{\infty\} \mid \tau_d(g) \cong \tau_d(h)\}$ is not empty. The following proposition confirms that the τ-distance restricted to $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ is indeed an ultrametric:

Proposition 5.3 (truncation-based ultrametric)**.** *For each truncation function τ, the τ-distance $\mathbf{d}_\tau$ constitutes an ultrametric on $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$.*

Proof. The identity resp. the symmetry condition follow by

$$\mathbf{d}_\tau(g,h) = 0 \iff \mathsf{sim}_\tau(g,h) = \infty \iff \tau_\infty(g) \cong \tau_\infty(h) \overset{(*)}{\iff} g \cong h \overset{\text{Prop. 3.17}}{\iff} g = h, \quad \text{and}$$
$$\mathbf{d}_\tau(g,h) = 2^{-\mathsf{sim}_\tau(g,h)} = 2^{-\mathsf{sim}_\tau(h,g)} = \mathbf{d}_\tau(h,g).$$

The equivalence $(*)$ is valid by axiom (b) of Definition 5.1. For the strong triangle condition we have to show that

$$\mathsf{sim}_\tau(g_1, g_3) \geq \min\{\mathsf{sim}_\tau(g_1, g_2), \mathsf{sim}_\tau(g_2, g_3)\}.$$

With $d = \min\{\mathsf{sim}_\tau(g_1, g_2), \mathsf{sim}_\tau(g_2, g_3)\}$ we have, by axiom (c) of Definition 5.1, that $\tau_d(g_1) \cong \tau_d(g_2)$ and $\tau_d(g_2) \cong \tau_d(g_3)$. Since we have that $\tau_d(g_1) \cong \tau_d(g_3)$ then, we can conclude that $\mathsf{sim}_\tau(g_1, g_2) \geq d$. □

Given their particular structure, we can reformulate the definition of Cauchy sequences and convergence in metric spaces induced by truncation functions in terms of the truncation function itself:

Lemma 5.4. *For each truncation function τ, each $g \in (\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\tau)$, and each sequence $(g_\iota)_{\iota<\alpha}$ in $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\tau)$ the following holds:*

- *(i) $(g_\iota)_{\iota<\alpha}$ is Cauchy iff for each $d \in \mathbb{N}$ there is some $\beta < \alpha$ such that $\tau_d(g_\gamma) \cong \tau_d(g_\iota)$ for all $\beta \leq \gamma, \iota < \alpha$.*
- *(ii) $(g_\iota)_{\iota<\alpha}$ converges to g iff for each $d \in \mathbb{N}$ there is some $\beta < \alpha$ such that $\tau_d(g) \cong \tau_d(g_\iota)$ for all $\beta \leq \iota < \alpha$.*

Proof. We only show (i) as (ii) is essentially the same. For "only if" direction assume that $(g_\iota)_{\iota<\alpha}$ is Cauchy and that $d \in \mathbb{N}$. We then find some $\beta < \alpha$ such that $\mathbf{d}_\tau(g_\gamma, g_\iota) < 2^{-d}$ for all $\beta \leq \gamma, \iota < \alpha$. Hence, we obtain that $\mathsf{sim}_\tau(g_\gamma, g_\iota) > d$ for all $\beta \leq \gamma, \iota < \alpha$. That is, $\tau_e(g_\gamma) \cong \tau_e(g_\iota)$ for some $e > d$. According to axiom (c) of Definition 5.1, we can then conclude that $\tau_d(g_\gamma) \cong \tau_d(g_\iota)$ for all $\beta \leq \gamma, \iota < \alpha$.

For the "if" direction assume some $\varepsilon \in \mathbb{R}^+$. Then there is some $d \in \mathbb{N}$ with $2^{-d} \leq \varepsilon$. By the initial assumption we find some $\beta < \alpha$ with $\tau_d(g_\gamma) \cong \tau_d(g_\iota)$ for all $\beta \leq \gamma, \iota < \alpha$, i.e. $\mathsf{sim}_\tau(g_\gamma, g_\iota) \geq d$. Hence, we have that $\mathbf{d}_\tau(g_\gamma, g_\iota) = 2^{\mathsf{sim}_\tau(g_\gamma, g_\iota)} < 2^{-d} \leq \varepsilon$ for all $\beta \leq \gamma, \iota < \alpha$. □

5.2 The Strict Truncation and its Metric Space

In this section, we consider a straightforward truncation function that simply cuts off all nodes at the given depth d.

Definition 5.5 (strict truncation)**.** Let $g \in \mathcal{G}^\infty(\Sigma_\perp)$ and $d \in \mathbb{N} \cup \{\infty\}$. The *strict truncation* $g{\restriction}d$ of g at d is a term graph defined by

$$N^{g\restriction d} = \left\{n \in N^g \,\middle|\, \mathsf{depth}_g(n) \leq d\right\} \qquad r^{g\restriction d} = r^g$$

$$\mathsf{lab}^{g\restriction d}(n) = \begin{cases} \mathsf{lab}^g(n) & \text{if } \mathsf{depth}_g(n) < d \\ \perp & \text{if } \mathsf{depth}_g(n) = d \end{cases} \qquad \mathsf{suc}^{g\restriction d}(n) = \begin{cases} \mathsf{suc}^g(n) & \text{if } \mathsf{depth}_g(n) < d \\ \langle\rangle & \text{if } \mathsf{depth}_g(n) = d \end{cases}$$

Figure 3 on page 26 shows a term graph g and its strict truncation at depth 2. Note that a node can get truncated even though its successor is retained.

One can easily see that the truncated term graph $g{\restriction}d$ is obtained from g by relabelling all nodes at depth d to $\perp$, removing all their outgoing edges and then removing all nodes that thus become unreachable from the root. This makes the strict truncation a straightforward generalisation of the truncation on terms.

The strict truncation indeed induces a truncation function:

Proposition 5.6. *Let $\restriction$ be the function with $\restriction_d(g) = g\restriction d$. Then $\restriction$ is a truncation function.*

Proof. (a) and (b) of Definition 5.1 follow immediately from the construction of the truncation. For (c) assume that $g\restriction d \cong h\restriction d$. Let $0 \leq e < d$ and let $\phi\colon g\restriction d \to h\restriction d$ be the witnessing isomorphism. Note that strict truncations preserve the depth of nodes, i.e. $\mathsf{depth}_{g\restriction d}(n) = \mathsf{depth}_g(n)$ for all $n \in N^{g\restriction d}$. This can be shown by a straightforward induction on $\mathsf{depth}_g(n)$. Moreover, by Corollary 3.15 also isomorphisms preserve the depth of nodes. Hence,

$$\mathsf{depth}_h(\phi(n)) = \mathsf{depth}_{h\restriction d}(\phi(n)) = \mathsf{depth}_{g\restriction d}(n) = \mathsf{depth}_g(n) \quad \text{for all } n \in N^{g\restriction d}$$

Restricting ϕ to the nodes in $g\restriction e$ thus yields an isomorphism from $g\restriction e$ to $h\restriction e$. □

Next we show that the metric space $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\restriction)$ that is induced by the truncation function $\restriction$ is in fact complete. To do this, we give a characterisation of the strict truncation in terms of labelled quotient trees.

Lemma 5.7 (labelled quotient tree of a strict truncation)**.** *Let $g \in \mathcal{G}^\infty(\Sigma_\bot)$ and $d \in \mathbb{N}\cup\{\infty\}$. The strict truncation $g\restriction d$ is uniquely determined up to isomorphism by the labelled quotient tree $(P, l, \sim)$ with*

(a) $P = \{\pi \in \mathcal{P}(g) \,|\, \forall \pi_1 < \pi \exists \pi_2 \sim_g \pi_1 \text{ with } |\pi_2| < d\}$,

(b) $l(\pi) = \begin{cases} g(\pi) & \text{if } \exists \pi' \sim_g \pi \text{ with } |\pi'| < d \\ \bot & \text{otherwise} \end{cases}$

(c) $\sim\, = \,\sim_g \cap\, P \times P$

Proof. We just have to show that $(P, l, \sim)$ is the canonical labelled quotient tree induced by $g\restriction d$. Then the lemma follows from Lemma 3.20. The case $d = \infty$ is trivial. In the following we assume that $d \in \mathbb{N}$.

Before continuing the proof, note that

$$\text{for each } \pi \in \mathcal{P}(g\restriction d) \text{ we have that } \pi \in \mathcal{P}(g) \text{ and } \mathsf{node}_{g\restriction d}(\pi) = \mathsf{node}_g(\pi). \qquad (*)$$

This can be shown by an induction on the length of π: The case $\pi = \langle\rangle$ is trivial. If $\pi = \pi'\cdot\langle i\rangle$, let $n = \mathsf{node}_{g\restriction d}(\pi')$ and $m = \mathsf{node}_{g\restriction d}(\pi)$. Hence, $m = \mathsf{suc}^{g\restriction d}_i(n)$ and, by construction of $g\restriction d$, also $m = \mathsf{suc}^g_i(n)$. Since by induction hypothesis $n = \mathsf{node}_g(\pi')$, we can thus conclude that $\pi \in \mathcal{P}(g)$ and that $\mathsf{node}_g(\pi) = m = \mathsf{node}_{g\restriction d}(\pi)$.

(a) $P = \mathcal{P}(g\restriction d)$. For the "$\subseteq$" direction let $\pi \in P$. To show that $\pi \in \mathcal{P}(g\restriction d)$, assume a $\pi_1 < \pi$ and let $n = \mathsf{node}_g(\pi_1)$. Since $\pi \in P$, there is some $\pi_2 \sim_g \pi_1$ with $|\pi_2| < d$. That is, $\mathsf{depth}_g(n) < d$. Therefore, we have that $n \in N^{g\restriction d}$ and $\mathsf{suc}^{g\restriction d}(n) = \mathsf{suc}^g(n)$. Hence, each node on the path π in g is also a node in $g\restriction d$ and has the same successor nodes as in g. That is, $\pi \in \mathcal{P}(g\restriction d)$.

For the "$\supseteq$" direction, assume some $\pi \in \mathcal{P}(g\restriction d)$. By $(*)$, π is also a position in g. To show that $\pi \in P$, let $\pi_1 < \pi$. Since only nodes of depth smaller than d can have a successor node in $g\restriction d$, the node $\mathsf{node}_{g\restriction d}(\pi_1)$ in $g\restriction d$ is at depth smaller than d. Hence, there is some $\pi_2 \sim_{g\restriction d} \pi_1$ with $|\pi_2| < d$. Because $\pi_2 \sim_{g\restriction d} \pi$ implies that $\pi_2 \sim_g \pi$, we can conclude that $\pi \in P$.

(b) $l(\pi) = g(\pi)$ for all $\pi \in P$. Let $\pi \in P$ and $n = \mathsf{node}_g(\pi)$. We distinguish two cases. At first suppose that there is some $\pi' \sim_g \pi$ with $|\pi'| < d$. Then $l(\pi) = g(\pi)$.

Since $n = \mathsf{node}_g(\pi')$, we have that $\mathsf{depth}_g(n) < d$. Consequently, $\mathsf{lab}^{g\restriction d}(n) = \mathsf{lab}^g(n)$ and, therefore, $g\restriction d(n) = g(\pi) = l(\pi)$. In the other case that there is no $\pi' \sim_g \pi$ with $|\pi| < d$, we have $l(\pi) = \bot$. This also means that $\mathsf{depth}_g(n) = d$. Consequently, $g\restriction d(\pi) = \mathsf{lab}^{g\restriction d}(n) = \bot = l(\pi)$.

(c) $\sim\, = \,\sim_{g\restriction d}$. Using the fact that $P = \mathcal{P}(g\restriction d)$, we can conclude for all $\pi_1, \pi_2 \in P$ that

$$\pi_1 \sim_{g\restriction d} \pi_2 \iff \mathsf{node}_{g\restriction d}(\pi_1) = \mathsf{node}_{g\restriction d}(\pi_2) \overset{(*)}{\iff} \mathsf{node}_g(\pi_1) = \mathsf{node}_g(\pi_2) \iff \pi_1 \sim_g \pi_2$$

□

Notice that a position π is retained by a truncation, i.e. $\pi \in P$, iff each node that π passes through is at a depth lower than d (and is thus not truncated or relabelled).

From this characterisation we immediately obtain the following relation between a term graph and its strict truncations:

Corollary 5.8. *Let $g \in \mathcal{G}^\infty(\Sigma_\bot)$ and $d \in \mathbb{N} \cup \{\infty\}$. Then*

(i) $\pi \in \mathcal{P}(g)$ iff $\pi \in \mathcal{P}(g\restriction d)$ for all π with $|\pi| \leq d$, and

(ii) $g\restriction d(\pi) = g(\pi)$ for all $\pi \in \mathcal{P}(g)$ with $|\pi| < d$.

Proof. Using the reflexivity of $\sim_g$, (i) follows immediately from Lemma 5.7 (a), and (ii) follows immediately from Lemma 5.7 (b). □

We can now show that the metric space induced by the strict truncation is complete:

Theorem 5.9. *The metric space $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\restriction)$ is complete. In particular, each Cauchy sequence $(g_\iota)_{\iota<\alpha}$ in $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\restriction)$ converges to the canonical term graph given by the following labelled quotient tree $(P, l, \sim)$:*

$$P = \liminf_{\iota\to\alpha} \mathcal{P}(g_\iota) = \bigcup_{\beta<\alpha} \bigcap_{\beta\leq\iota<\alpha} \mathcal{P}(g_\iota) \qquad \sim\, = \liminf_{\iota\to\alpha} \sim_{g_\iota} = \bigcup_{\beta<\alpha} \bigcap_{\beta\leq\iota<\alpha} \sim_{g_\iota}$$

$$l(\pi) = g_\beta(\pi) \quad \text{for some } \beta < \alpha \text{ with } g_\iota(\pi) = g_\beta(\pi) \text{ for each } \beta \leq \iota < \alpha \qquad \text{for all } \pi \in P$$

Proof. We need to check that $(P, l, \sim)$ is a well-defined labelled quotient tree. At first we show that l is a well-defined function on P. In order to show that l is functional, assume that there are $\beta_1, \beta_2 < \alpha$ such that there is a π with $g_\iota(\pi) = g_{\beta_k}(\pi)$ for all $\beta_k \leq \iota < \alpha$, $= 1, 2$. but then $g_{\beta_1}(\pi) = g_\beta(\pi) = g_{\beta_2}(\pi)$ for $\beta = \max\{\beta_1, \beta_2\}$.

To show that l is total on P, let $\pi \in P$ and $d = |\pi|$. By Lemma 5.4, there is some $\beta < \alpha$ such that $g_\gamma\restriction d+1 \cong g_\iota\restriction d+1$ for all $\beta \leq \gamma, \iota < \alpha$. According to Corollary 5.8, this means that all g_ι for $\beta \leq \iota < \alpha$ agree on positions of length smaller than $d+1$, in particular π. Hence, $g_\iota(\pi) = g_\beta(\pi)$ for all $\beta \leq \iota < \alpha$, and we have $l(\pi) = g_\beta(\pi)$.

One can easily see that $\sim$ is a binary relation on P: If $\pi_1 \sim \pi_2$, then there is some $\beta < \alpha$ with $\pi_1 \sim_{g_\iota} \pi_2$ for all $\beta \leq \iota < \alpha$. Hence, $\pi_1, \pi_2 \in \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$ and thus $\pi_1, \pi_2 \in P$.

Similarly follows that $\sim$ is an equivalence relation on P: To show reflexivity, assume $\pi \in P$. Then there is some $\beta < \alpha$ such that $\pi \in \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$. Hence, $\pi \sim_{g_\iota} \pi$ for all $\beta \leq \iota < \alpha$ and, therefore, $\pi \sim \pi$. In the same way symmetry and transitivity follow from the symmetry and transitivity of $\sim_{g_\iota}$.

Finally, we have to show the reachability and the congruence property from Definition 3.19. To show reachability assume some $\pi \cdot \langle i \rangle \in P$. Then there is some $\beta < \alpha$ such that

$\pi \cdot \langle i \rangle \in \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$. Hence, since then also $\pi \in \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$, we have $\pi \in P$. According to the construction of l, there is also some $\beta \leq \gamma < \alpha$ with $g_\gamma(\pi) = l(\pi)$. Since $\pi \cdot \langle i \rangle \in \mathcal{P}(g_\gamma)$ we can conclude that $i < \mathsf{ar}(l(\pi))$.

To establish congruence assume that $\pi_1 \sim \pi_2$. Consequently, there is some $\beta < \gamma$ such that $\pi_1 \sim_{g_\iota} \pi_2$ for all $\beta \leq \iota < \alpha$. Therefore, we also have for each $\beta \leq \iota < \alpha$ that $\pi_1 \cdot \langle i \rangle \sim_{g_\iota} \pi_2 \cdot \langle i \rangle$ and that $g_\iota(\pi_1) = g_\iota(\pi_2)$. From the former we can immediately derive that $\pi_1 \cdot \langle i \rangle \sim \pi_2 \cdot \langle i \rangle$. Moreover, according to the construction of l, there some $\beta \leq \gamma < \alpha$ such that $l(\pi_1) = g_\gamma(\pi_1) = g_\gamma(\pi_2) = l(\pi_2)$.

This concludes the proof that $(P, l, \sim)$ is indeed a labelled quotient tree. Next, we show that the sequence $(g_\iota)_{\iota<\alpha}$ converges to the thus defined canonical term graph g. By Lemma 5.4, this amounts to giving for each $d \in \mathbb{N}$ some $\beta < \alpha$ such that $g{\restriction}d \cong g_\iota{\restriction}d$ for each $\beta \leq \iota < \alpha$.

To this end, let $d \in \mathbb{N}$. Since $(g_\iota)_{\iota<\alpha}$ is Cauchy, there is, according to Lemma 5.4, some $\beta < \alpha$ such that

$$g_\iota{\restriction}d \cong g_{\iota'}{\restriction}d \quad \text{for all } \beta \leq \iota, \iota' < \alpha. \tag{1}$$

In order to show that this implies that $g{\restriction}d \cong g_\iota{\restriction}d$ for each $\beta \leq \iota < \alpha$, we show that the respective labelled quotient trees of $g{\restriction}d$ and $g_\iota{\restriction}d$ as characterised by Lemma 5.7 coincide. The labelled quotient tree $(P_1, l_1, \sim_1)$ for $g{\restriction}d$ is given by

$$\begin{aligned} P_1 &= \{\pi \in P \mid \forall \pi_1 < \pi \exists \pi_2 \sim \pi_1 : |\pi_2| < d\} \\ \sim_1 &= \sim \cap\, P_1 \times P_1 \end{aligned} \qquad l_1(\pi) = \begin{cases} l(\pi) & \text{if } \exists \pi' \sim \pi : |\pi'| < d \\ \bot & \text{otherwise} \end{cases}$$

The labelled quotient tree $(P_2^\iota, l_2^\iota, \sim_2^\iota)$ for each $g_\iota{\restriction}d$ is given by

$$\begin{aligned} P_2^\iota &= \{\pi \in \mathcal{P}(g_\iota) \mid \forall \pi_1 < \pi \exists \pi_2 \sim_{g_\iota} \pi_1 : |\pi_2| < d\} \\ \sim_2^\iota &= \sim \cap\, P_2^\iota \times P_2^\iota \end{aligned} \qquad l_2^\iota(\pi) = \begin{cases} g_\iota(\pi) & \text{if } \exists \pi' \sim_{g_\iota} \pi : |\pi'| < d \\ \bot & \text{otherwise} \end{cases}$$

Due to (1), all $(P_2^\iota, l_2^\iota, \sim_2^\iota)$ with $\beta \leq \iota < \alpha$ are pairwise equal. Therefore, we write $(P_2, l_2, \sim_2)$ for this common labelled quotient tree. That is, it remains to be shown that $(P_1, l_1, \sim_1)$ and $(P_2, l_2, \sim_2)$ are equal.

(a) $P_1 = P_2$. For the "$\subseteq$" direction let $\pi \in P_1$. If $\pi = \langle\rangle$, we immediately have that $\pi \in P_2$. Hence, we can assume that π is non-empty. Since $\pi \in P_1$ implies $\pi \in P$, there is some $\beta \leq \beta' < \alpha$ with $\pi \in \mathcal{P}(g_\iota)$ for all $\beta' \leq \iota < \alpha$. Moreover this means that for each $\pi_1 < \pi$ there is some $\pi_2 \sim \pi_1$ with $|\pi_2| < d$. That is, there is some $\beta' \leq \gamma_{\pi_1} < \alpha$ such that $\pi_2 \sim_{g_\iota} \pi_1$ for all $\gamma_{\pi_1} \leq \iota < \alpha$. Since there are only finitely many proper prefixes $\pi_1 < \pi$ but at least one, we can define $\gamma = \max\{\gamma_{\pi_1} \mid \pi_1 < \pi\}$ such that we have for each $\pi_1 < \pi$ some $\pi_2 \sim_{g_\gamma} \pi_1$ with $|\pi_2| < d$. Hence, $\pi \in P_2^\gamma = P_2$.

To show the converse direction, assume that $\pi \in P_2$. Then $\pi \in P_2^\iota \subseteq \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$. Hence, $\pi \in P$. To show that $\pi \in P_1$, assume some $\pi_1 < \pi$. Since $\pi \in P_2^\beta$, there is some $\pi_2 \sim_{g_\beta} \pi_1$ with $|\pi_2| < d$. Then $\pi_1 \in P_2$ because P_2 is closed under prefixes and $\pi_2 \in P_2$ because $|\pi_2| < d$. Thus, $\pi_2 \sim_2 \pi_1$ which implies $\pi_2 \sim_{g_\iota} \pi_1$ for all $\beta \leq \iota < \alpha$. Consequently, $\pi_2 \sim \pi_1$, which means that $\pi \in P_1$.

(c) $\sim_1 = \sim_2$. For the "$\subseteq$" direction assume $\pi_1 \sim_1 \pi_2$. Hence, $\pi_1 \sim \pi_2$ and $\pi_1, \pi_2 \in P_1 = P_2$. This means that there is some $\beta \leq \gamma < \alpha$ with $\pi_1 \sim_{g_\gamma} \pi_2$. Consequently, $\pi_1 \sim_2 \pi_2$. For the converse direction assume that $\pi_1 \sim_2 \pi_2$. Then $\pi_1, \pi_2 \in P_2 = P_1$ and $\pi_1 \sim_{g_\iota} \pi_2$ for all $\beta \leq \iota < \alpha$. Hence, $\pi_1 \sim \pi_2$ and we can conclude that $\pi_1 \sim_1 \pi_2$.

(b) $l_1 = l_2$. We show this by proving that, for all $\beta \leq \iota < \alpha$, the condition $\exists \pi' \sim \pi : |\pi'| < d$ from the definition of l_1 is equivalent to the condition $\exists \pi' \sim_{g_\iota} \pi : |\pi'| < d$ from the definition of l_2 and that $l(\pi) = g_\iota(\pi)$ if either condition is satisfied. The latter is simple: Whenever there is some $\pi' \sim \pi$ with $|\pi'| < d$, then $g_\iota(\pi) = l_2^\iota(\pi) = l_2^\beta(\pi) = g_\beta(\pi)$ for all $\beta \leq \iota < \alpha$. Hence, $l(\pi) = g_\beta(\pi) = g_\iota(\pi)$ for all $\beta \leq \iota < \alpha$. For the former, we first consider the "only if" direction of the equivalence. Let $\pi \in P_1$ and $\pi' \sim \pi$ with $|\pi'| < d$. Then also $\pi' \in P_1$ which means that $\pi' \sim_1 \pi$. Since then $\pi' \sim_2 \pi$, we can conclude that $\pi' \sim_{g_\iota} \pi$ for all $\beta \leq \iota < \alpha$. For the converse direction assume that $\pi \in P_2$, $\pi' \sim_{g_\iota} \pi$ and $|\pi'| < d$. Then also $\pi' \in P_2$ which means that $\pi' \sim_2 \pi$ and, therefore, $\pi' \sim \pi$. □

Example 5.10. Reconsider the sequence of term graphs $(h_\iota)_{\iota<\omega}$ Figure 5c on page 30. As we have noticed in Example 4.8, the edge that loops back to the root node is pushed down as the sequence progresses. Thus, we have for each $n \in \mathbb{N}$, that the strict truncations of the term graphs h_ι with $n \leq \iota < \omega$ at depth $n+1$ coincide. Therefore, by Lemma 5.4, $(h_\iota)_{\iota<\omega}$ is Cauchy. In particular, we have that $(h_\iota)_{\iota<\omega}$ converges to h_ω.

The limit inferior induced by $\leq_\bot^{\mathcal{G}}$ showed some curios behaviour as soon as acyclic sharing changes as we have seen in Example 4.8 with the convergence illustrated in Figure 2. This is not the case for the metric $\mathbf{d}_\upharpoonright$. In fact, there is no topological space in which $(g_\iota)_{\iota<\omega}$ from Figure 2 converges to a unique limit.

5.3 Other Truncation Functions and Their Metric Spaces

Generalising concepts from terms to term graphs is not a straightforward matter as we have to decide how to deal with additional sharing that term graphs offer. The definition of strict truncation seems to be an obvious choice for a generalisation of tree truncation. In this section, we shall formally argue that it is in fact the case. More specifically, we show that no matter how we define the sharing of the $\bot$-nodes that fill the holes caused by the truncation, we obtain the same topology. We will then contrast this to the metric that we have used in previous work [6] by showing that small changes to its definition also change the induced topology.

The following lemma is a handy tool for comparing metric spaces induced by truncation functions:

Lemma 5.11. *Let τ, υ be two truncation functions on $\mathcal{G}^\infty(\Sigma_\bot)$ and $f\colon \mathcal{G}^\infty_\mathcal{C}(\Sigma) \to \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ a function on $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$. Then the following are equivalent*

(i) f is a continuous mapping $f\colon (\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\tau) \to (\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\upsilon)$

(ii) For each $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ and $d \in \mathbb{N}$ there is some $e \in \mathbb{N}$ such that

$$\mathsf{sim}_\tau(g,h) \geq e \quad\Longrightarrow\quad \mathsf{sim}_\upsilon(f(g), f(h)) \geq d \quad \textit{for all } h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$$

(iii) For each $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ and $d \in \mathbb{N}$ there is some $e \in \mathbb{N}$ such that

$$\tau_e(g) \cong \tau_e(h) \quad\Longrightarrow\quad \upsilon_d(f(g)) \cong \upsilon_d(f(h)) \quad \textit{for all } h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$$

Proof. Analogous to Lemma 5.4. □

An easy consequence of the above lemma is that if two truncation functions only differ by a constant depth, they induce the same topology:

Proposition 5.12. *Let τ, υ be two truncation functions on $\mathcal{G}^\infty(\Sigma_\bot)$ such that there is a $\delta \in \mathbb{N}$ with $|\mathsf{sim}_\tau(g,h) - \mathsf{sim}_\upsilon(g,h)| \le \delta$ for all $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$. Then $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\tau)$ and $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\upsilon)$ are topologically equivalent, i.e. induce the same topology.*

Proof. We show that the identity function $\mathsf{id}\colon \mathcal{G}^\infty_\mathcal{C}(\Sigma) \to \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ is a homeomorphism from $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\tau)$ to $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\upsilon)$, i.e. both id and id^{-1} are continuous. Due to the symmetry of the setting it suffices to show that id is continuous. To this end, let $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ and $d \in \mathbb{N}$. Define $e = d + \delta$ and assume some $h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ such that $\mathsf{sim}_\tau(g,h) \ge e$. By Lemma 5.11, it remains to be shown that then $\mathsf{sim}_\upsilon(g,h) \ge d$. Indeed, we have $\mathsf{sim}_\upsilon(g,h) \ge \mathsf{sim}_\tau(g,h) - \delta \ge e - \delta = d$. □

This shows that metric spaces induced by truncation functions are essentially invariant under changes in the truncation function bounded by a constant margin.

Remark 5.13. We should point out that the original definition of the metric on terms by Arnold and Nivat [3] was slightly different from the one we showed here. Recall that we defined similarity as the maximum depth of truncation that ensures equality:

$$\mathsf{sim}_\tau(g,h) = \max \{ d \in \mathbb{N} \cup \{\infty\} \mid \tau_d(g) \cong \tau_d(h) \}$$

Arnold and Nivat, on the other hand, defined it as the minimum truncation depth that still shows inequality:

$$\mathsf{sim}'_\tau(g,h) = \min \{ d \in \mathbb{N} \cup \{\infty\} \mid \tau_d(g) \not\cong \tau_d(h) \}$$

However, it is easy to see that either both $\mathsf{sim}_\tau(g,h)$ and $\mathsf{sim}'_\tau(g,h)$ are ∞ or $\mathsf{sim}'_\tau(g,h) = \mathsf{sim}_\tau(g,h) + 1$. Hence, by Proposition 5.12, both definitions yield the same topology.

Proposition 5.12 also shows that two truncation functions induce the same topology if they only differ in way they treat "fringe nodes", i.e. nodes that are introduced in place of the nodes that have been cut off. Since the definition of truncation functions requires that $\tau_0(g) \cong \bot$ and $\tau_\infty(g) \cong g$, we usually do not give the explicit construction of the truncation for the depths 0 and ∞.

Example 5.14. Consider the following variant τ of the strict truncation function $\restriction$. Given a term graph $g \in \mathcal{G}^\infty(\Sigma_\bot)$ and depth $d \in \mathbb{N}^+$ we define the truncation $\tau_d(g)$ as follows:

$$\begin{aligned}
N^g_{<d} &= \left\{ n \in N^g \,\middle|\, \mathsf{depth}_g(n) < d \right\} \\
N^g_{=d} &= \left\{ n^i \,\middle|\, n \in N^g_{<d}, 0 \le i < \mathsf{ar}_g(n), \mathsf{suc}^g_i(n) \notin N^g_{<d} \right\} \\
N^{\tau_d(g)} &= N^g_{<d} \uplus N^g_{=d}
\end{aligned}$$

$$\mathsf{lab}^{\tau_d(g)}(n) = \begin{cases} \mathsf{lab}^g(n) & \text{if } n \in N^g_{<d} \\ \bot & \text{if } n \in N^g_{=d} \end{cases} \qquad \mathsf{suc}^{\tau_d(g)}_i(n) = \begin{cases} \mathsf{suc}^g_i(n) & \text{if } n^i \notin N^g_{=d} \\ n^i & \text{if } n^i \in N^g_{=d} \end{cases}$$

One can easily show that τ is in fact a truncation function. The difference between $\restriction$ and τ is that in the latter we create a fresh node n^i whenever a node n has a successor $\mathsf{suc}^g_i(n)$ that lies at the fringe, i.e. at depth d. Since this only affects the nodes at the fringe and, therefore, only nodes at the same depth d we get the following:

$$\begin{aligned}
g{\restriction}d \cong h{\restriction}d &\quad\Longrightarrow\quad \tau_d(g) \cong \tau_d(h), \text{ and} \\
\tau_d(g) \cong \tau_d(h) &\quad\Longrightarrow\quad g{\restriction}d-1 \cong h{\restriction}d-1.
\end{aligned}$$

Hence, the respectively induced similarities only differ by a constant margin of 1, i.e. we have that $|\mathsf{sim}_{\restriction}(g,h) - \mathsf{sim}_{\tau}(g,h)| = 1$. According to Proposition 5.6, this means that $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_{\restriction})$ and $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\tau)$ are topologically equivalent.

Consider another variant υ of the strict truncation function $\restriction$. Given a term graph $g \in \mathcal{G}^\infty(\Sigma_\perp)$ and depth $d \in \mathbb{N}^+$, we define the truncation $\upsilon_d(g)$ as follows:

$$N^g_{<d} = \left\{ n \in N^g \,\middle|\, \mathsf{depth}_g(n) < d \right\}$$

$$N^g_{=d} = \left\{ n^i \,\middle|\, \begin{array}{rl} n \in N^g, \mathsf{depth}_g(n) = d-1, 0 \le i < \mathsf{ar}_g(n) \text{ with} & \mathsf{suc}^g_i(n) \notin N^g_{<d} \\ \text{or} & n \notin \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n)) \end{array} \right\}$$

$$N^{\upsilon_d(g)} = N^g_{<d} \uplus N^g_{=d}$$

$$\mathsf{lab}^{\upsilon_d(g)}(n) = \begin{cases} \mathsf{lab}^g(n) & \text{if } n \in N^g_{<d} \\ \perp & \text{if } n \in N^g_{=d} \end{cases} \qquad \mathsf{suc}^{\upsilon_d(g)}(n) = \begin{cases} \mathsf{suc}^g_i(n) & \text{if } n^i \notin N^g_{=d} \\ n^i & \text{if } n^i \in N^g_{=d} \end{cases}$$

Also υ forms a truncation function as one can easily show. In addition to creating fresh nodes n^i for each successor that is not in the retained nodes $N^g_{<d}$, the truncation function υ creates such new nodes n^i for each cycle that created by a node just above the fringe. Again, as for the truncation function τ, only the nodes at the fringe, i.e. at depth d are affected by this change. Hence, the respectively induced similarities of $\restriction$ and υ only differ by a constant margin of 1, which makes the metric spaces $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_{\restriction})$ and $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\upsilon)$ topologically equivalent as well.

The robustness of the metric space $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_{\restriction})$ under the changes illustrated above is due to the uniformity of the core definition of the strict truncation which only takes into account the depth. By simply increasing the depth by a constant number, we can compensate for changes in the way fringe nodes are dealt with.

This is much different for the truncation function $g \wr d$ that induces the metric space considered in [6]:

Definition 5.15 (truncation of term graphs)**.** Let $g \in \mathcal{G}^\infty(\Sigma_\perp)$ and $d \in \mathbb{N}$.

(i) Given $n, m \in N^g$, m is an *acyclic predecessor* of n in g if there is an acyclic occurrence $\pi \cdot \langle i \rangle \in \mathcal{P}^a_g(n)$ with $\pi \in \mathcal{P}_g(m)$. The set of acyclic predecessors of n in g is denoted $\mathsf{Pre}^a_g(n)$.

(ii) The set of *retained nodes* of g at d, denoted $N^g_{<d}$, is the least subset M of N^g satisfying the following conditions for all $n \in N^g$:

$$\text{(T1) } \mathsf{depth}_g(n) < d \implies n \in M \qquad \text{(T2) } n \in M \implies \mathsf{Pre}^a_g(n) \subseteq M$$

(iii) For each $n \in N^g$ and $i \in \mathbb{N}$, we use n^i to denote a fresh node, i.e. $\left\{ n^i \,\middle|\, n \in N^g, i \in \mathbb{N} \right\}$ is a set of pairwise distinct nodes not occurring in N^g. The set of *fringe nodes* of g at d, denoted $N^g_{=d}$, is defined as the singleton set $\{r^g\}$ if $d = 0$, and otherwise as the set

$$\left\{ n^i \,\middle|\, \begin{array}{rl} n \in N^g_{<d}, 0 \le i < \mathsf{ar}_g(n) \text{ with} & \mathsf{suc}^g_i(n) \notin N^g_{<d} \\ \text{or} & \mathsf{depth}_g(n) \ge d-1, n \notin \mathsf{Pre}^a_g(\mathsf{suc}^g_i(n)) \end{array} \right\}$$

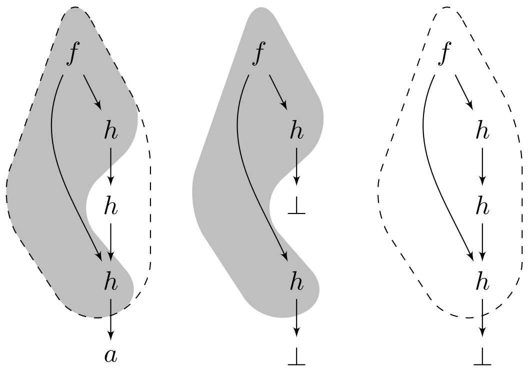

(g) $(g{\restriction}2)$ $(g{\wr}2)$

Figure 3: Comparison to strict truncation.

(iv) The *truncation* of g at d, denoted $g{\wr}d$, is the term graph defined by

$$N^{g\wr d} = N^g_{<d} \uplus N^g_{=d} \qquad r^{g\wr d} = r^g$$

$$\mathsf{lab}^{g\wr d}(n) = \begin{cases} \mathsf{lab}^g(n) & \text{if } n \in N^g_{<d} \\ \bot & \text{if } n \in N^g_{=d} \end{cases} \qquad \mathsf{suc}^{g\wr d}_i(n) = \begin{cases} \mathsf{suc}^g_i(n) & \text{if } n^i \not\in N^g_{=d} \\ n^i & \text{if } n^i \in N^g_{=d} \end{cases}$$

Additionally, we define $g{\wr}\infty$ to be the term graph g itself.

The idea of this definition of truncation is that not only each node at depth $< d$ is kept (T1) but also every acyclic predecessor of such a node (T2). In sum, every node on an acyclic path from the root to a node at depth smaller than d is kept. The difference between the two truncation functions $\restriction$ and $\wr$ are illustrated in Figure 3.

In contrast to $\restriction$, the truncation function $\wr$ is quite vulnerable to small changes:

Example 5.16. Consider the following variant τ of the truncation function $\wr$. Given a term graph $g \in \mathcal{G}^\infty(\Sigma_\bot)$ and depth $d \in \mathbb{N}^+$, we define the truncation $\tau_d(g)$ as follows: The set of retained nodes $N^g_{<d}$ is defined as for the truncation $g{\wr}d$. For the rest we define

$$\begin{aligned} N^g_{=d} &= \left\{ \mathsf{suc}^g_i(n) \,\middle|\, n \in N^g_{<d}, 0 \le i < \mathsf{ar}_g(n), \mathsf{suc}^g_i(n) \not\in N^g_{<d} \right\} \\ N^{\tau_d(g)} &= N^g_{<d} \uplus N^g_{=d} \end{aligned}$$

$$\mathsf{lab}^{\tau_d(g)}(n) = \begin{cases} \mathsf{lab}^g(n) & \text{if } n \in N^g_{<d} \\ \bot & \text{if } n \in N^g_{=d} \end{cases} \qquad \mathsf{suc}^{\tau_d(g)}(n) = \begin{cases} \mathsf{suc}^g(n) & \text{if } n \in N^g_{<d} \\ \langle\rangle & \text{if } n \in N^g_{=d} \end{cases}$$

In this variant of truncation, some sharing of the retained nodes is preserved. Instead of creating fresh nodes for each successor node that is not in the set of retained nodes, we simply keep the successor node. Additionally loops back into the retained nodes are not cut off. This variant of the truncation deals with its retained nodes in essentially the same way as the strict truncation. However, opposed the strict truncation and their variants, this truncation function yields a topology different from the metric space $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\wr)$! To see this, consider the two families of term graphs g_n and h_n indicated in Figure 4. For both

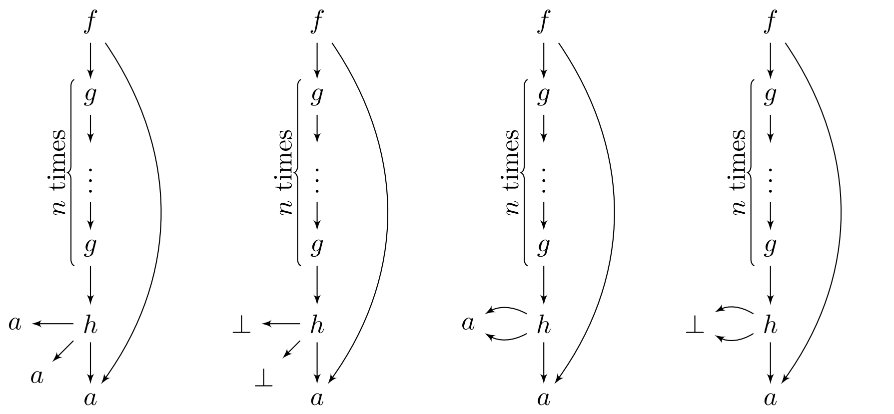

(g_n) $(\tau_2(g_n) = \tau_{n+2}(g_n))$ (h_n) $(\tau_2(h_n) = \tau_{n+2}(h_n))$

Figure 4: Variations in fringe nodes.

families we have that the τ-truncations at depth 2 to $n+2$ are the same, i.e. $\tau_d(g_n) = \tau_2(g_n)$ and $\tau_d(h_n) = \tau_2(h_n)$ for all $2 \leq d \leq n+2$. The same holds for the truncation function $\wr$. Moreover, since the two leftmost successors of the h-node are not shared in g_n, both truncation functions coincide on g_n, i.e. $g_n \wr d = \tau_d(g_n)$. This is not the case for h_n. In fact, they only coincide up to depth 1. However, we have that $h_n \wr d = \tau_d(g_n)$. In total, we can observe that $\mathsf{sim}(g_n, h_n) = n+2$ but $\mathsf{sim}_\tau(g_n, h_n) = 1$. This means, however, that the sequence $\langle g_0, h_0, g_1, h_1, \ldots \rangle$ converges in $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\wr)$ but not in $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\tau)$!

A similar example can be constructed that uses the difference in the way the two truncation functions deal with fringe nodes created by cycles back into the set of retained nodes.

6 Partial Order vs. Metric Space

Recall that p-convergence in term rewriting is a conservative extension of m-convergence (cf. Theorem 2.5). The key property that makes this possible is that for each sequence $(t_\iota)_{\iota<\alpha}$ in $\mathcal{T}^\infty(\Sigma)$, we have that $\lim_{\iota\to\alpha} t_\iota = \liminf_{\iota\to\alpha} t_\iota$ whenever $(t_\iota)_{\iota<\alpha}$ converges, or $\liminf_{\iota\to\alpha} t_\iota \in \mathcal{T}^\infty(\Sigma)$.

Unfortunately, this is not the case for the metric space and the partial order that we consider on term graphs. As we have shown in Example 5.10, the sequence of term graphs depicted in Figure 2 has a total term graph as its limit inferior although it does not converge in the metric space. This example shows that we cannot hope to generalise the compatibility property that we have for terms: Even if a sequence of total term graphs has a total term graph as its limit inferior, it might not converge. However, the other direction of the compatibility does hold true:

Theorem 6.1. *If $(g_\iota)_{\iota<\alpha}$ converges, then* $\lim_{\iota\to\alpha} g_\iota = \liminf_{\iota\to\alpha} g_\iota$.

Proof. In order to prove this property, we will use the construction of the limit resp. the limit inferior of a sequence of term graphs which we have shown in Theorem 5.9 resp. Corollary 4.7.

According to Theorem 5.9, we have that the canonical term graph $\lim_{\iota \to \alpha} g_\iota$ is given by the following labelled quotient tree $(P, \sim, l)$:

$$P = \bigcup_{\beta < \alpha} \bigcap_{\beta \le \iota < \alpha} \mathcal{P}(g_\iota) \qquad \sim = \bigcup_{\beta < \alpha} \bigcap_{\beta \le \iota < \alpha} \sim_{g_\iota}$$
$$l(\pi) = f \quad \text{iff} \quad \exists \beta < \alpha \forall \beta \le \iota < \alpha : g_\iota(\pi) = f$$

We will show that $g = \liminf_{\iota \to \alpha} g_\iota$ induces the same labelled quotient tree.

From Corollary 4.7, we immediately obtain that $\mathcal{P}(g) \subseteq P$. To show the converse direction $\mathcal{P}(g) \supseteq P$, we assume some $\pi \in P$. According to Corollary 4.7, in order to show that $\pi \in \mathcal{P}(g)$, we have to find a $\beta < \alpha$ such that $\pi \in \mathcal{P}(g_\beta)$ and for each $\pi' < \pi$ there is some $f \in \Sigma_\perp$ such that $g_\iota(\pi') = f$ for all $\beta \le \iota < \alpha$.

Because $\pi \in P$, there is some $\beta_1 < \alpha$ such that $\pi \in \mathcal{P}(g_\iota)$ for all $\beta_1 \le \iota < \alpha$. Since $(g_\iota)_{\iota < \alpha}$ converges, it is also Cauchy. Hence, by Lemma 5.4, for each $d \in \mathbb{N}$, there is some $\beta_2 < \alpha$ such that $g_\gamma \upharpoonright d \cong g_\iota \upharpoonright d$ for all $\beta_2 \le \gamma, \iota < \alpha$. Specialising this to $d = |\pi|$, we obtain some $\beta_2 < \alpha$ with $g_\gamma \upharpoonright |\pi| \cong g_\iota \upharpoonright |\pi|$ for all $\beta_2 \le \gamma, \iota < \alpha$. Let $\beta = \max\{\beta_1, \beta_2\}$. Then we have $\pi \in \mathcal{P}(g_\iota)$ and $g_\beta \upharpoonright |\pi| \cong g_\iota \upharpoonright |\pi|$ for each $\beta \le \iota < \alpha$. Hence, for each $\pi' < \pi$, the symbol $f = g_\beta(\pi')$ is well-defined, and, according to Corollary 5.8, we have that $g_\iota(\pi') = f$ for each $\beta \le \iota < \alpha$.

The equalities $\sim = \sim_g$ and $l = g(\cdot)$ follow from Corollary 4.7 as $P = \mathcal{P}(g)$. □

7 Infinitary Term Graph Rewriting

In the previous sections, we have constructed and investigated the necessary metric and partial order structures upon which the infinitary calculus of term graph rewriting that we shall introduce in this section is based. After describing the framework of term graph rewriting that we consider, we will explore different modes convergence on term graphs. In the same way that infinitary term rewriting is based on the abstract notions of m- and p-convergence [4], infinitary term graph rewriting is an instantiation of these abstract modes of convergence to term graphs.

7.1 Term Graph Rewriting Systems

The framework of term graph rewriting that we consider is that of Barendregt et al. [7]. Similarly to term rewriting systems, we have to deal with variables. That is, we consider a signature $\Sigma_\mathcal{V}$ extended with a set of variable symbols $\mathcal{V}$.

Definition 7.1 (term graph rewriting system)**.**

(i) Given a signature Σ, a *term graph rule* ρ over Σ is a triple (g, l, r) where g is a graph over $\Sigma_\mathcal{V}$ and $l, r \in N^g$, such that all nodes in g reachable from l or r. We write ρ_l resp. ρ_r to denote the left- resp. right-hand side of ρ, i.e. the term graph $g|_l$ resp. $g|_r$. Additionally, we require that ρ_l is finite and that for each variable $v \in \mathcal{V}$ there is at most one node n in g labelled v and n is different but still reachable from l.

(ii) A *term graph rewriting system (GRS)* $\mathcal{R}$ is a pair (Σ, R) with Σ a signature and R a set of term graph rules.

The requirement that the root l of the left-hand side is not labelled with a variable symbol is analogous to the requirement that the left-hand side of a term rule is not a variable. Similarly, the restriction that nodes labelled with variable symbols must be reachable from the root of the left-hand side corresponds to the restriction on term rules that every variable occurring on the right-hand side must also occur on the left-hand side.

Term graphs can be used to compactly represent term. This representation of terms is defined by the unravelling. This notion can be extended to term graph rules. Figure 5a illustrates two term graph rules that both represent the term rule $a : x \to b : a : x$ from Example 2.1 to which they unravel.

Definition 7.2 (unravelling of term graph rules)**.** Let ρ be a term graph rule with ρ_l and ρ_r left- resp. right-hand side term graph. The *unravelling* of ρ, denoted $\mathcal{U}(\rho)$ is the term rule $\mathcal{U}(\rho_l) \to \mathcal{U}(\rho_r)$. Let $\mathcal{R} = (\Sigma, R)$ be a GRS. The unravelling of $\mathcal{R}$, denoted $\mathcal{U}(\mathcal{R})$ is the TRS $(\Sigma, \mathcal{U}(R))$ with $\mathcal{U}(R) = \{\mathcal{U}(\rho) \mid \rho \in G\}$.

We will investigate the aspect of how term graph rules simulate their unravellings in Section 8.

The application of a rewrite rule ρ (with root nodes l and r) to a term graph g is performed in four steps: At first a suitable sub-term graph of g rooted in some node n of g is *matched* against the left-hand side of ρ. This amounts to finding a $\mathcal{V}$-homomorphism ϕ from the term graph rooted in l to the sub-term graph rooted in n, the *redex*. The $\mathcal{V}$-homomorphism ϕ allows to instantiate variables in the rule with sub-term graphs of the redex. In the second step, nodes and edges in ρ that are not reachable from l are copied into g, such that edges pointing to nodes in the term graph rooted in l are redirected to the image under ϕ. In the last two steps, all edges pointing to n are redirected to (the copy of) r and all nodes not reachable from the root of (the now modified version of) g are removed.

Definition 7.3 (application of a term graph rewrite rule, [7])**.** Let $\rho = (N^\rho, \mathsf{lab}^\rho, \mathsf{suc}^\rho, l^\rho, r^\rho)$ be a term graph rewrite rule over signature Σ, $g \in \mathcal{G}^\infty(\Sigma)$ and $n \in N^g$. ρ is called *applicable* to g at n if there is a $\mathcal{V}$-homomorphism $\phi\colon\ \rho_l \to_{\mathcal{V}} g|_n$. ϕ is called the *matching $\mathcal{V}$-homomorphism* of the rule application, and $g|_n$ is called a *ρ-redex*. Next, we define the *result* of the application of the rule ρ to g at n using the $\mathcal{V}$-homomorphism ϕ. This is done by constructing the intermediate graphs g_1 and g_2, and the final result g_3.

(i) The graph g_1 is obtained from g by adding the part of ρ not contained in the left-hand side:

$$\begin{aligned} N^{g_1} &= N^g \uplus (N^\rho \setminus N^{\rho_l}) \\ \mathsf{lab}^{g_1}(m) &= \begin{cases} \mathsf{lab}^g(m) & \text{if } m \in N^g \\ \mathsf{lab}^\rho(m) & \text{if } m \in N^\rho \setminus N^{\rho_l} \end{cases} \\ \mathsf{suc}^{g_1}_i(m) &= \begin{cases} \mathsf{suc}^g_i(m) & \text{if } m \in N^g \\ \mathsf{suc}^\rho_i(m) & \text{if } m, \mathsf{suc}^\rho_i(m) \in N^\rho \setminus N^{\rho_l} \\ \phi(\mathsf{suc}^\rho_i(m)) & \text{if } m \in N^\rho \setminus N^{\rho_l}, \mathsf{suc}^\rho_i(m) \in N^{\rho_l} \end{cases} \end{aligned}$$

(ii) Let $n' = \phi(r^\rho)$ if $r^\rho \in N^{\rho_l}$ and $n' = r^\rho$ otherwise. The graph g_2 is obtained from g_1 by redirecting edges ending in n to n':

$$N^{g_2} = N^{g_1} \qquad \mathsf{lab}^{g_2} = \mathsf{lab}^{g_1} \qquad \mathsf{suc}^{g_2}_i(m) = \begin{cases} \mathsf{suc}^{g_1}_i(m) & \text{if } \mathsf{suc}^{g_1}_i(m) \neq n \\ n' & \text{if } \mathsf{suc}^{g_1}_i(m) = n \end{cases}$$

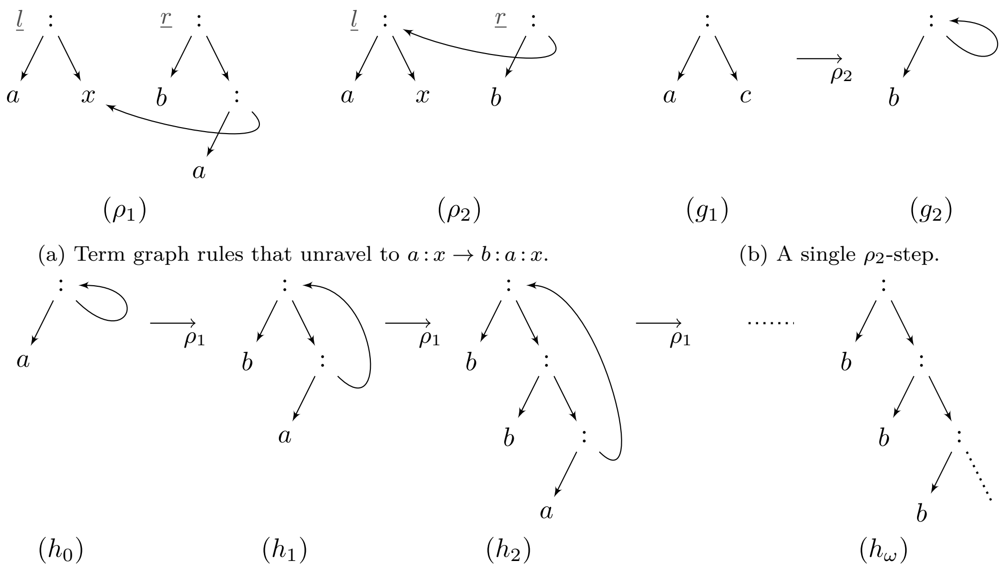

(a) Term graph rules that unravel to $a : x \to b : a : x$.

(b) A single ρ_2-step.

(c) A strongly m-convergent term graph reduction over ρ_1.

Figure 5: Term graph rules and their reductions.

(iii) The term graph g_3 is obtained by setting the root node r', which is r if $l = r^g$, and otherwise r^g. That is, $g_3 = g_2|_{r'}$. This also means that all nodes not reachable from r' are removed.

This induces a reduction step $\psi\colon\ g \to g_3$. In order to indicate the applied rule ρ and the root nodes n, n' of the redex resp. the reduct, we write $\psi\colon\ g \to_{n,\rho,n'} g_3$.

Examples for term graph rewriting steps are shown in Figure 5. We revisit them in more detail in Example 7.7 in the next section.

Note that term graph rules do not provide a duplication mechanism. Each variable is allowed to occur at most once. Duplication must always be simulated by sharing. This means for example that variables that should "occur" on the right-hand side must share the occurrence of that variable on the left-hand side of the rule as seen in the term graph rules in Figure 5a. This sharing can be direct as in ρ_1 or indirect as in ρ_2. For variables that are supposed to be duplicated on the right-hand side, for example in the term rewrite rule $Y\,x \to x\,(Y\,x)$ that defines the fixed point combinator Y in an applicative language, we have to use sharing in order to represent multiple occurrence of the same variable as seen in the corresponding term graph rules in Figure 6a.

As for term graphs, we also give a linear notation for term graph rules:

Definition 7.4 (linear notation of term graph rules)**.** Let Σ be a signature and $\widehat{\Sigma}$ its extension as in Definition 3.21. A linear notation for a term graph rule over Σ is a term rule $\rho\colon s \to t$ over $\widehat{\Sigma}$ such that for each $n \in \mathcal{N}$ that occurs in ρ there is exactly one occurrence of a function symbol of the form $f^{[n]}$ in ρ.

The corresponding term graph rule ρ' is defined as follows: Consider the term tree representations of s and t. Let l and r be the root nodes of s resp. t, and let g be the disjoint union of s and t. In g, redirect every edge to a node labelled n to the unique node labelled $f^{[n]}$ for some $f \in \Sigma$. Then change all labellings of the form $f^{[n]}$ to f. In the resulting graph g' do the following: For each $x \in \mathcal{V}$ occurring in g', redirect all edges to nodes labelled x to a single fresh node labelled x provided $x \neq t$. If $x = t$, then redirect all edges to nodes labelled x to the node r. Let g'' be the thus obtained graph after removing all nodes not reachable from l or r. Then ρ' is the term graph rule (g'', l, r).

As an example, the term graph rules in Figure 5a can be written as $\rho_1\colon\ a:x \to b:a:x$ resp. $\rho_2\colon\ a :^{[n]} x \to b:n$. Also note that each term rule $\rho\colon\ l \to r$ can be interpreted as a linear notation for a term graph rule $\rho'\colon\ l \to r$. This term graph rule ρ' is, in fact, the term graph rule with minimal sharing that unravels to ρ.

7.2 Weak Convergence

We start by first considering weak notions of convergences based on the metric $\mathbf{d}_\upharpoonright$ and the partial order $\leq^{\mathcal{G}}_\bot$:

Definition 7.5. Let $\mathcal{R} = (\Sigma, R)$ be a GRS. Term graphs in $\mathcal{R}$ range over $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$. Thus, we consider the applications of term graph to yield canonical term graphs. That is, when we write $g \to_{n,\rho,n'} h$, we mean that there is a reduction step $g \to_{n,\rho,\widehat{n}} h'$ according to Definition 7.3, such that $h = \mathcal{C}(h')$ and $n' = \phi(\widehat{n})$ for the isomorphism $\phi\colon\ h' \to h$.

(i) A *(transfinite) reduction* in $\mathcal{R}$ is a sequence $(g_\iota \to_{n_\iota,\rho_\iota} g_{\iota+1})_{i<\alpha}$ of rewriting steps in $\mathcal{R}$. If S is finite, we write $S\colon\ g_0 \to^* g_\alpha$.

(ii) Let $S = (g_\iota \to_\mathcal{R} g_{\iota+1})_{\iota<\alpha}$ be a reduction in $\mathcal{R}$. S is *weakly m-continuous*, written $S\colon\ g_0 \xrightarrow{m}_\mathcal{R} \dots$, if the underlying sequence of term graphs $(g_\iota)_{\iota<\widehat{\alpha}}$ is continuous, i.e. $\lim_{\iota\to\lambda} g_\iota = g_\lambda$ for each limit ordinal $\lambda < \alpha$. S *weakly m-converges* to $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ in $\mathcal{R}$, written $S\colon\ g_0 \xrightarrow{m}_\mathcal{R} g$, if it is weakly m-continuous and $\lim_{\iota\to\widehat{\alpha}} g_\iota = g$.

(iii) Let $\mathcal{R}_\bot$ be the GRS $(\Sigma_\bot, R)$ over the extended signature $\Sigma_\bot$ and $S = (g_\iota \to_{\mathcal{R}_\bot} g_{\iota+1})_{\iota<\alpha}$ a reduction in $\mathcal{R}_\bot$. S is weakly p-continuous, written $S\colon\ g_0 \xrightarrow{p}_\mathcal{R} g$, if $\liminf_{\iota<\lambda} g_i = g_\lambda$ for each limit ordinal $\lambda < \alpha$. S *weakly p-converges* to $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$ in $\mathcal{R}$, written $S\colon\ g_0 \xrightarrow{p}_\mathcal{R} g$, if it is weakly p-continuous and $\liminf_{\iota<\widehat{\alpha}} g_i = g$.

Note that we have to extend the signature of $\mathcal{R}$ to $\Sigma_\bot$ for the definition of weak p-convergence. Moreover, since the partial order $\leq^{\mathcal{G}}_\bot$ forms a complete semilattice on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$, weak p-continuity coincides with weak p-convergence:

Proposition 7.6. *In a GRS, every weakly p-continuous reduction is weakly p-convergent.*

Proof. Follows immediately from Corollary 4.7. ☐

Example 7.7. Consider the term graph rule ρ_1 in Figure 5a that unravels to the term rule $a : x \to b : a : x$ from Example 2.1. Starting with the term tree $a : c$, depicted as g_1 in Figure 5b, we obtain the same transfinite reduction as in Example 2.1:

$$S\colon\ a:c \to_{\rho_1} b:a:c \to_{\rho_1} b:b:a:c \to_{\rho_1} \dots\ h_\omega$$

Also in this setting, S both weakly m- and p-converges to the term tree h_ω shown in Figure 5c. Similarly, we can reproduce the weakly p-converging but not weakly m-converging reduction T from Example 2.3. Notice that h_ω is a rational term tree as it can be obtained by unravelling the finite term graph g_2 depicted in Figure 5b. In fact, if we use the rule ρ_2, we can immediately rewrite g_1 to g_2, which unravels to h_ω. In ρ_2, not only the variable x is shared but the whole left-hand side of the rule. This causes each redex of ρ_2 to be *captured* by the right-hand side.

Figure 5c indicates a transfinite reduction starting with a cyclic term graph h_0 that unravels to the rational term $t = a : a : a : \ldots$. This reduction both weakly m- and p-converges to the rational term tree h_ω as well. Again, by using ρ_2 instead of ρ_1, we can rewrite h_0 to the cyclic term graph g_2 in one step.

As for TRSs, we have that weak m-convergence implies weak p-convergence.

Theorem 7.8. *Let S be a reduction in a GRS $\mathcal{R}$.*

$$\textit{If} \quad S\colon g \xrightarrow{m}_{\mathcal{R}} h \qquad \textit{then} \qquad S\colon g \xrightarrow{p}_{\mathcal{R}} h.$$

Proof. Follows straightforwardly from Theorem 6.1. ☐

However, as we have indicated, weak m-convergence is not the total fragment of weak p-convergence as it is the case for TRS. The GRS with the two rules $f(c,c) \to f(c^{[n]}, n)$ and $f(c,c) \to f(c,c)$ yields the reduction sequence shown in Figure 2. This reduction weakly p-converges to $f(c,c)$ but is not weakly m-convergent. However, this peculiar behaviour can be ruled out by considering strong convergence, which is the subject of the following sections.

7.3 Reduction Contexts

The idea of strong convergence is to conservatively approximate the convergence behaviour somewhat independent from the actual rules that are applied. Strong m-convergence in TRSs requires that the depth of the redexes tends to infinity thereby assuming everything at the depth of the redex or below can potentially be affected by a reduction step. Strong p-convergence, on the other hand, uses a better approximation that only assumes that the whole redex can be changed by a reduction not however its siblings. To this end strong p-convergence uses a notion of reduction contexts – essentially the term minus the redex – for the formation of limits. In order to define a suitable notion of strong p-convergence on term graphs, we have to devise a corresponding notion of reduction contexts. In this section we shall devise such a notion and argue for its adequacy.

The following definition provides the basic construction that we use to remove nodes from a term graph:

Definition 7.9 (local truncation)**.** Let $g \in \mathcal{G}^\infty(\Sigma_\perp)$ and $N \subseteq N^g$. The *local truncation* of g by N, denoted $g{\restriction}N$, is given as follows:

$$N^{g\restriction N} \text{ is the least set } M \text{ satisfying } \begin{array}{l} (a)\ r^g \in M, \textit{and} \\ (b)\ n \in M \setminus N \implies \mathsf{suc}^g(n) \subseteq M. \end{array}$$

$$r^{g\restriction N} = r^g \qquad \mathsf{lab}^{g\restriction N} = \begin{cases} \mathsf{lab}^g(n) & \text{if } n \notin N \\ \perp & \text{if } n \notin N \end{cases} \qquad \mathsf{suc}^{g\restriction N}(n) = \begin{cases} \mathsf{suc}^g(n) & \text{if } n \notin N \\ \langle\rangle & \text{if } n \in N \end{cases}$$

By abuse of notation, we write $g{\restriction}n$ instead of $g{\restriction}\{n\}$.

The goal for the rest of this section is to establish that $g \upharpoonright n$ is an adequate notion of reduction context for a reduction step $g \to_n h$ applied at node n in g. According to he abstract notion of strong p-convergence [4], this requires that $g \upharpoonright n \leq_\bot^{\mathcal{G}} g, h$.

The following lemma shows that local truncations only remove positions from a term graph but do not alter them:

Lemma 7.10. *Let $g \in \mathcal{G}^\infty(\Sigma_\bot)$, $N \subseteq N^g$ and $\pi \in \mathcal{P}(g \upharpoonright N)$. Then $\mathsf{node}_g(\pi) = \mathsf{node}_{g \upharpoonright N}(\pi)$.*

Proof. We proceed by induction on the length of π. The case $\pi = \langle\rangle$ follows from the definition $r^{g \upharpoonright N} = r^g$. If $\pi = \pi' \cdot \langle i \rangle$, we can use the induction hypothesis to obtain that $\mathsf{node}_g(\pi') = \mathsf{node}_{g \upharpoonright N}(\pi')$. As $\pi' \cdot \langle i \rangle \in \mathcal{P}(g \upharpoonright N)$, we know that $\mathsf{node}_{g \upharpoonright N}(\pi') \notin N$. We can thus reason as follows:

$$\mathsf{node}_g(\pi) = \mathsf{suc}_i^g(\mathsf{node}_g(\pi')) = \mathsf{suc}_i^g(\mathsf{node}_{g \upharpoonright N}(\pi')) = \mathsf{suc}_i^{g \upharpoonright N}(\mathsf{node}_{g \upharpoonright N}(\pi')) = \mathsf{node}_{g \upharpoonright N}(\pi)$$

$\square$

This leads immediately to the observation that local truncations preserve sharing:

Lemma 7.11 (local truncations preserve sharing)**.** *Let $g \in \mathcal{G}^\infty(\Sigma_\bot)$, $N \subseteq N^g$ and $\pi_1, \pi_2 \in \mathcal{P}(g \upharpoonright N)$. Then $\pi_1 \sim_g \pi_2$ iff $\pi_1 \sim_{g \upharpoonright N} \pi_2$.*

Proof.

$$\begin{aligned} \pi_1 \sim_g \pi_2 &\iff \mathsf{node}_g(\pi_1) = \mathsf{node}_g(\pi_2) \\ &\iff \mathsf{node}_{g \upharpoonright N}(\pi_1) = \mathsf{node}_{g \upharpoonright N}(\pi_2) && \text{(Lemma 7.10)} \\ &\iff \pi_1 \sim_{g \upharpoonright N} \pi_2 \end{aligned}$$

$\square$

Most importantly, we obtain the intuitively expected property that local truncations yield smaller term graphs w.r.t. $\leq_\bot^{\mathcal{G}}$:

Lemma 7.12. *For each $g \in \mathcal{G}^\infty(\Sigma_\bot)$ and $N \subseteq N^g$, we have $g \upharpoonright N \leq_\bot^{\mathcal{G}} g$.*

Proof. We use Corollary 4.3 to show this. (a) follows immediately from Lemma 7.11. For (b), let $\pi \in \mathcal{P}(g \upharpoonright N)$ with $g \upharpoonright N(\pi) \in \Sigma$. Hence, $\mathsf{node}_{g \upharpoonright N}(\pi) \notin N$ and we can reason as follows:

$$g \upharpoonright N(\pi) = \mathsf{lab}^{g \upharpoonright N}(\mathsf{node}_{g \upharpoonright N}(\pi)) = \mathsf{lab}^g(\mathsf{node}_{g \upharpoonright N}(\pi)) \overset{\text{Lem. 7.10}}{=} \mathsf{lab}^g(\mathsf{node}_g(\pi)) = g(\pi).$$

$\square$

The following property summarises the core property that we require for an adequate notion of reduction context: The reduction context of a reduction step is the maximal substructure that is guaranteed to be preserved by the reduction.

Lemma 7.13. *Given a graph reduction step $g \to_{n,\rho,n'} h$ with $g, h \in \mathcal{G}^\infty(\Sigma_\bot)$, we have $g \upharpoonright n \cong h \upharpoonright n'$. The corresponding isomorphism is given by*

$$\phi(m) = \begin{cases} m & \text{if } m \neq n \\ n' & \text{if } m = n \end{cases} \qquad \text{for all } m \in N^{g \upharpoonright n}$$

Proof. At first, observe that n', the root of the reduct, is either a fresh node from ρ, or a node reachable from n in g. Hence, we know that

$$m \in N^{g\restriction n} \setminus \{n\} \quad \text{implies that} \quad m \neq n'. \tag{$*$}$$

In order to prove that $\phi\colon N^{g\restriction n} \to N^{h\restriction n'}$ is well-defined, we have to show that $N^{g\restriction n} \setminus \{n\} \subseteq N^{h\restriction n'}$: Let $m \in N^{g\restriction n} \setminus \{n\}$. We will show by induction on $\mathsf{depth}_{g\restriction n}(m)$, that $m \in N^{h\restriction n'}$.

If $\mathsf{depth}_{g\restriction n}(m) = 0$, then $m = r^{g\restriction n} = r^g = r^h = r^{h\restriction n} \in N^{h\restriction n}$, where $r^g = r^h$ holds because $n \neq r^g$. If $\mathsf{depth}_{g\restriction n}(m) > 0$, then there is some $m' \in N^{g\restriction n}$ with $\mathsf{depth}_{g\restriction n}(m') < \mathsf{depth}_{g\restriction n}(m)$ and $\mathsf{suc}_i^{g\restriction n}(m') = m$ for some $i \in \mathbb{N}$. Hence, $m' \neq n$, which means that also $\mathsf{suc}_i^g(m') = m$ and that, by induction hypothesis, $m' \in N^{h\restriction n'}$. Since, in a graph reduction step, only edges to the redex node n are redirected, we have that $\mathsf{suc}_i^h(m') \neq \mathsf{suc}_i^g(m')$ iff $\mathsf{suc}_i^g(m') = n$. Thus, as $\mathsf{suc}_i^g(m') = m \neq n$, we have $\mathsf{suc}_i^h(m') = \mathsf{suc}_i^g(m') = m$. Moreover, by $(*)$, we know that $m' \neq n'$. Thus, $m = \mathsf{suc}_i^h(m') \in N^{h\restriction n'}$.

Next, we show that ϕ is a homomorphism from $g\restriction n$ to $h\restriction n'$. The root condition is satisfied as follows:

$$\phi(r^{g\restriction n}) = \phi(r^g) = \left\{\begin{matrix} r^g & \text{if } r^g \neq n \\ n' & \text{if } r^g = n \end{matrix}\right\} = r^h = r^{h\restriction n'}.$$

For the labelling and successor condition, assume some $m \in N^{g\restriction n}$. If $m = n$, then $\phi(m) = n'$ and the labelling and successor condition follow immediately from the construction of $g\restriction n$ and $h\restriction n'$. If $m \neq n$, then $\phi(m) = m$ and, by $(*)$, $m \neq n'$. Since the labelling of nodes is not changed by a reduction step, we have

$$\mathsf{lab}^{h\restriction n'}(\phi(m)) = \mathsf{lab}^{h\restriction n'}(m) = \mathsf{lab}^h(m) = \mathsf{lab}^g(m) = \mathsf{lab}^{g\restriction n}(m) = \mathsf{lab}^{g\restriction n}(\phi(m)).$$

For the successor condition, first assume that $\mathsf{suc}_i^g(m) = n$. Then the edge to n is redirected to n' by the reduction step, i.e. $\mathsf{suc}_i^h(m) = n'$, and we have

$$\mathsf{suc}_i^{h\restriction n'}(\phi(m)) = \mathsf{suc}_i^{h\restriction n'}(m) = \mathsf{suc}_i^h(m) = n' = \phi(n) = \phi(\mathsf{suc}_i^g(m)) = \phi(\mathsf{suc}_i^{g\restriction n}(m)).$$

If, on the other hand, $\mathsf{suc}_i^g(m) \neq n$, the edge is retained, i.e. $\mathsf{suc}_i^h(m) = \mathsf{suc}_i^g(m)$, and we have

$$\mathsf{suc}_i^{h\restriction n'}(\phi(m)) = \mathsf{suc}_i^{h\restriction n'}(m) = \mathsf{suc}_i^h(m) = \mathsf{suc}_i^g(m) = \phi(\mathsf{suc}_i^g(m)) = \phi(\mathsf{suc}_i^{g\restriction n}(m)).$$

The injectivity of ϕ follows from the fact that $\phi(m) = n'$ if $m = n$ and that, by $(*)$, $\phi(m) = m \neq n'$ if $m \neq n$. Hence, according Lemma 3.10, ϕ is an isomorphism, i.e. $g\restriction n \cong h\restriction n'$. ☐

As an easy consequence of this, we obtain that $g\restriction n$ is indeed an adequate notion of reduction context.

Proposition 7.14. *Given a graph reduction step $g \to_{n,\rho,n'} h$, we have $g\restriction n \leq_\bot^{\mathcal{G}} g, h$.*

Proof. Lemma 7.12 yields $g\restriction n \leq_\bot^{\mathcal{G}} g$. By Lemma 7.13 and Lemma 7.12, we get $g\restriction n \cong h\restriction n' \leq_\bot^{\mathcal{G}} h$. ☐

The following lemma provides a convenient characterisation of local truncations in terms of labelled quotient trees:

Lemma 7.15. *For each $g \in \mathcal{G}^\infty(\Sigma_\perp)$ and $n \in N^g$, the local truncation $g{\restriction}n$ has the following canonical labelled quotient tree $(P, l, \sim)$:*

$$P = \{\pi \in \mathcal{P}(g) \,|\, \forall \pi' < \pi\colon \pi' \notin \mathcal{P}_g(n)\} \qquad l(\pi) = \begin{cases} g(\pi) & \text{if } \pi \notin \mathcal{P}_g(n) \\ \perp & \text{if } \pi \in \mathcal{P}_g(n) \end{cases} \quad \text{for all } \pi \in P$$
$$\sim \,=\, \sim_g \cap\, P \times P$$

In particular, given $\pi \in \mathcal{P}(g)$, we have that $g(\pi) = g{\restriction}n(\pi)$ if $\pi' \notin \mathcal{P}_g(n)$ for each $\forall \pi' \leq \pi$.

Proof. The last statement above follows immediately from the preceding characterisation of $(P, l, \sim)$. We will show in the following that $(P, l, \sim)$ is equal to $(\mathcal{P}(g{\restriction}n), g{\restriction}n(\cdot), \sim_{g{\restriction}n})$.

By Lemma 7.10 $\mathcal{P}(g{\restriction}n) \subseteq \mathcal{P}(g)$. Therefore, in order to prove that $\mathcal{P}(g{\restriction}n) \subseteq P$, we assume some $\pi \in \mathcal{P}(g{\restriction}n)$ and show by induction on the length of π that no proper prefix of π is a position of n in g. The case $\pi = \langle\rangle$ is trivial as $\langle\rangle$ has no proper prefixes. If $\pi = \pi' \cdot \langle i \rangle$, we can assume by induction that $\pi' \in P$ since $\pi' \in \mathcal{P}(g{\restriction}n)$. Consequently, no proper prefix of π' is in $\mathcal{P}_g(n)$. It thus remains to be shown that π' itself is not in $\mathcal{P}_g(n)$. Since $\pi' \cdot \langle i \rangle \in \mathcal{P}(g{\restriction}n)$, we know that $\mathsf{suc}_i^{g{\restriction}n}(\mathsf{node}_{g{\restriction}n}(\pi'))$ is defined. Therefore, $\mathsf{node}_{g{\restriction}n}(\pi')$ cannot be n, and since, by Lemma 7.10, $\mathsf{node}_{g{\restriction}n}(\pi') = \mathsf{node}_g(\pi')$, neither can $\mathsf{node}_g(\pi')$. In other words, $\pi' \notin \mathcal{P}_g(n)$.

For the converse direction $P \subseteq \mathcal{P}(g{\restriction}n)$, assume some $\pi \in P$. We will show by induction on the length of π, that then $\pi \in \mathcal{P}(g{\restriction}n)$. The case $\pi = \langle\rangle$ is trivial. If $\pi = \pi' \cdot \langle i \rangle$, then also $\pi' \in P$ which, by induction, implies that $\pi' \in \mathcal{P}(g{\restriction}n)$. Since $\pi \in \mathcal{P}(g)$, according to Lemma 7.10, we have that $\mathsf{ar}_g(\pi') > i$. Let $m = \mathsf{node}_g(\pi')$. According to Lemma 7.10, $m = \mathsf{node}_{g{\restriction}n}(\pi')$. Since $\pi \in P$, we have that $\pi' \notin \mathcal{P}_g(n)$ and thus $m \neq n$. Hence, according to the definition of $g{\restriction}n$, $\mathsf{suc}^{g{\restriction}n}(m) = \mathsf{suc}^g(m)$ which implies that $\mathsf{ar}_{g{\restriction}n}(\pi') > i$. Consequently, $\pi \in \mathcal{P}(g{\restriction}n)$.

The equality $\sim \,=\, \sim_{g{\restriction}n}$ follows directly from Lemma 7.11 and the equality $P = \mathcal{P}(g{\restriction}n)$.

For the equality $l = g{\restriction}n(\cdot)$, consider some $\pi \in \mathcal{P}(g{\restriction}n)$. Since $\mathsf{node}_g(\pi) = n$ iff $\pi \in \mathcal{P}_g(n)$, we can reason as follows:

$$g{\restriction}n(\pi) = \mathsf{lab}^{g{\restriction}n}(\mathsf{node}_{g{\restriction}n}(\pi)) \overset{\text{Lem. 7.10}}{=} \mathsf{lab}^{g{\restriction}n}(\mathsf{node}_g(\pi)) = \begin{cases} g(\pi) & \text{if } \pi \notin \mathcal{P}_g(n) \\ \perp & \text{if } \pi \in \mathcal{P}_g(n) \end{cases}$$

□

7.4 Strong Convergence

Now that we have an adequate notion of reduction context, we can define strong p-convergence on term graphs analogously to strong p-convergence on terms. For strong m-convergence, we simply take the same notion of depth that we already used for the definition of strict truncation and thus the metric space.

Definition 7.16. Let $\mathcal{R}$ be a GRS.

(i) The *reduction context* c of a graph reduction step $\phi\colon g \to_n h$ is the term graph $\mathcal{C}(g{\restriction}n)$. We write $\phi\colon g \to_c h$ to indicate the reduction context of a graph reduction step.

(ii) Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$ be a reduction in $\mathcal{R}$. S is strongly m-continuous in $\mathcal{R}$, denoted $S\colon g_0 \overset{m}{\twoheadrightarrow}_{\mathcal{R}} \dots$, if $\lim_{\iota\to\lambda} g_\iota = g_\lambda$ and $\lim_{\iota\to\lambda} \mathsf{depth}_{g_\iota}(n_\iota) = \infty$ for each limit ordinal $\lambda < \alpha$. S strongly m-converges to g in $\mathcal{R}$, denoted $S\colon g_0 \overset{m}{\twoheadrightarrow}_{\mathcal{R}} g$, if it is strongly m-continuous, $\lim_{\iota\to\widehat{\alpha}} g_\iota = g$, and $\lim_{\iota\to\alpha} \mathsf{depth}_{g_\iota}(n_\iota) = \infty$ in case α is a limit ordinal.

(iii) Let $S = (g_\iota \to_{c_\iota} g_{\iota+1})_{\iota<\alpha}$ be a reduction in $\mathcal{R}$. S is strongly p-continuous in $\mathcal{R}$, denoted $S\colon g_0 \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \dots$, if $\liminf_{\iota\to\lambda} c_\iota = g_\lambda$ for each limit ordinal $\lambda < \alpha$. S strongly p-converges to g in $\mathcal{R}$, denoted $S\colon g_0 \overset{p}{\twoheadrightarrow}_{\mathcal{R}} g$, if it is strongly p-continuous and either $g = \liminf_{\iota\to\alpha} c_\iota$ or $g = g_\alpha$ in case S is closed.

Note that we have to extend the signature of $\mathcal{R}$ to $\Sigma_\perp$ for the definition of strong p-convergence. However, we can obtain the total fragment of strong p-convergence if we restrict ourselves to total term graphs in $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$: A reduction $(g_\iota \to_{\mathcal{R}_\perp} g_{\iota+1})_{\iota<\alpha}$ p-converging to g is called *total* if g as well as each g_ι is total, i.e. an element of $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$.

Since the partial order $\leq^\mathcal{G}_\perp$ forms a complete semilattice on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\perp)$, strong p-continuity coincides with strong p-convergence:

Proposition 7.17. *Every strongly p-continuous reduction is strongly p-convergent.*

Proof. Follows immediately from Corollary 4.7. □

The following technical lemma confirms the intuition that changes during a continuous reduction must be caused by a reduction step that was applied at the position where the difference is observed or above.

Lemma 7.18. *Let $(g_\iota \to_{n_\iota,\rho,m_\iota} g_{\iota+1})_{\iota<\alpha}$ be a strongly p-continuous reduction in a GRS with its reduction contexts $c_\iota = \mathcal{C}(g_\iota \upharpoonright n_\iota)$ such that there are $\beta \leq \gamma < \alpha$ and $\pi \in \mathcal{P}(c_\beta) \cap \mathcal{P}(c_\gamma)$ with $c_\beta(\pi) \neq c_\gamma(\pi)$. Then there is a position $\pi' \leq \pi$ and an index $\beta \leq \iota \leq \gamma$ such that $\pi' \in \mathcal{P}_{g_\iota}(n_\iota)$.*

Proof. Throughout the poof, we can assume that

$$g_\iota(\pi) = c_\iota(\pi) \qquad \text{if } \beta \leq \iota \leq \gamma \text{ and } \pi \in \mathcal{P}(c_\iota). \tag{$*$}$$

If this would not be the case, then, by Lemma 7.15, there is a $\pi' \leq \pi$ such that $\pi' \in \mathcal{P}_{g_\iota}(n_\iota)$, i.e. the statement that we want to prove holds.

We proceed with an induction on γ. The case $\gamma = \beta$ is trivial.

Let $\gamma = \iota + 1 > \beta$. We then consider two cases: If $\pi \notin \mathcal{P}(c_\iota)$, we are done as this together with the assumption $\pi \in \mathcal{P}(c_\gamma)$ implies, by the definition of reduction steps, that $\pi' \in \mathcal{P}_{g_\iota}(n_\iota)$ for some $\pi' < \pi$. If, on the other hand, $\pi \in \mathcal{P}(c_\iota)$, then we can assume that $c_\beta(\pi) = c_\iota(\pi)$ since otherwise the proof goal follows immediately from the induction hypothesis. Consequently, we have that

$$g_\gamma \upharpoonright m_\iota(\pi) \overset{\text{Lem. 7.13}}{=} g_\iota \upharpoonright n_\iota(\pi) = c_\iota(\pi) = c_\beta(\pi) \neq c_\gamma(\pi) \overset{(*)}{=} g_\gamma(\pi)$$

The thus obtained inequality $g_\gamma \upharpoonright m_\iota(\pi) \neq g_\gamma(\pi)$ implies, by Lemma 7.15, that there is a $\pi' \leq \pi$ such that $\pi' \in \mathcal{P}_{g_\gamma}(m_\iota)$. According to Lemma 7.13 there is an isomorphism $\phi\colon g_\iota \upharpoonright n_\iota \to g_\gamma \upharpoonright m_\iota$ with $\phi(n_\iota) = m_\iota$. This means, by Corollary 3.15, that $\mathcal{P}_{g_\iota}(n_\iota) = \mathcal{P}_{g_\gamma}(m_\iota)$. Hence, $\pi' \in \mathcal{P}_{g_\iota}(n_\iota)$.

Let γ be a limit ordinal. By $(*)$, we know that $g_\gamma(\pi) = c_\gamma(\pi) \neq c_\beta(\pi)$. According to Corollary 4.7, the inequality $g_\gamma(\pi) \neq c_\beta(\pi)$ is only possible if there is a $\pi' \leq \pi$ and a $\beta \leq \iota < \gamma$ such that $c_\iota(\pi') \neq c_\beta(\pi')$. Hence, we can invoke the induction hypothesis (for the position π' instead of π) which immediately yields the proof goal. □

By combining the characterisation of the limit inferior from Corollary 4.7 and the characterisation of local truncations from Lemma 7.15, we obtain the following characterisation of the limit of a strongly p-convergent reduction:

Lemma 7.19. *Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$ be an open reduction in a GRS strongly p-converging to g. Then g has the following canonical labelled quotient tree $(P, l, \sim)$:*

$$\begin{aligned}
P &= \bigcup_{\beta<\alpha} \{\pi \in \mathcal{P}(g_\beta) \mid \forall \pi' < \pi \forall \beta \leq \iota < \alpha \colon \pi' \notin \mathcal{P}_{g_\iota}(n_\iota)\} \\
\sim &= \left(\bigcup_{\beta<\alpha} \bigcap_{\beta\leq\iota<\alpha} \sim_{g_\iota} \right) \cap P \times P \\
l(\pi) &= \begin{cases} g_\beta(\pi) & \text{if } \exists \beta < \alpha \forall \beta \leq \iota < \alpha \colon \pi \notin \mathcal{P}_{g_\iota}(n_\iota) \\ \bot & \text{otherwise} \end{cases} \qquad \text{for all } \pi \in P
\end{aligned}$$

In particular, given $\beta < \alpha$ and $\pi \in \mathcal{P}(g_\beta)$, we have that $g(\pi) = g_\beta(\pi)$ if $\pi' \notin \mathcal{P}_{g_\iota}(n_\iota)$ for all $\pi' \leq \pi$ and $\beta \leq \iota < \alpha$.

Proof. The last statement above follows immediately from the preceding characterisation of $(P, l, \sim)$. We will show in the following that $(P, l, \sim)$ is equal to $(\mathcal{P}(g), g(\cdot), \sim_g)$.

Let $c_\iota = \mathcal{C}(g_\iota \upharpoonright n_\iota)$ for each $\iota < \alpha$. At first we show that $\mathcal{P}(g) \subseteq P$. To this end let $\pi \in \mathcal{P}(g)$. Since $g = \liminf_{\iota\to\alpha} c_\iota$, this means, by Corollary 4.7, that

$$\text{there is some } \beta < \alpha \text{ with } \pi \in \mathcal{P}(c_\beta) \text{ and } c_\iota(\pi') = c_\beta(\pi') \text{ for all } \pi' < \pi \text{ and } \beta \leq \iota < \alpha. \quad (1)$$

Since, according to Lemma 7.10, $\mathcal{P}(c_\beta) \subseteq \mathcal{P}(g_\beta)$, we also have $\pi \in \mathcal{P}(g_\beta)$. In order to prove that $\pi \in P$, we assume some $\pi' < \pi$ and $\beta \leq \iota < \alpha$ and show that $\pi' \notin \mathcal{P}_{g_\iota}(n_\iota)$. Since π' is a proper prefix of a position in c_β, we have that $c_\beta(\pi') \in \Sigma$. By (1), also $c_\iota(\pi') \in \Sigma$. Hence, according to Lemma 7.15, $\pi' \notin \mathcal{P}_{g_\iota}(n_\iota)$.

For the converse direction $P \subseteq \mathcal{P}(g)$, we assume some $\pi \in P$ and show that then $\pi \in \mathcal{P}(g)$. Since $\pi \in P$, we have that

$$\text{there is some } \beta < \alpha \text{ with } \pi \in \mathcal{P}(g_\beta) \text{ and } \pi' \notin \mathcal{P}_{g_\iota}(n_\iota) \text{ for all } \pi' < \pi \text{ and } \beta \leq \iota < \alpha. \quad (2)$$

In particular, we have that $\pi' \notin \mathcal{P}_{g_\beta}(n_\beta)$ for all $\pi' < \pi$. Hence, by Lemma 7.15, $\pi \in \mathcal{P}(c_\beta)$. According to Corollary 4.7, it remains to be shown that $c_\gamma(\pi') = c_\beta(\pi')$ for all $\pi' < \pi$ and $\beta \leq \gamma < \alpha$. We will do that by an induction on γ:

The case $\gamma = \beta$ is trivial. For $\gamma = \iota + 1 > \beta$, let $g_\iota \to_{n_\iota,\rho_\iota,n'_\iota} g_\gamma$ be the ι-th reduction step and $\pi' < \pi$. By Lemma 7.13, we then have $c_\iota \cong g_\iota \upharpoonright n_\iota \cong g_\gamma \upharpoonright n'_\iota$. We can thus reason as follows:

$$c_\beta(\pi') \overset{\text{ind. hyp.}}{=} c_\iota(\pi') \overset{\text{Lem. 7.13}}{=} g_\gamma \upharpoonright n'_\iota(\pi') \overset{\text{Lem. 7.15}}{=} g_\gamma(\pi') \overset{\text{Lem. 7.15}}{=} g_\gamma \upharpoonright n_\gamma(\pi') = c_\gamma(\pi')$$

The first application of Lemma 7.15 above is justified by the fact that $\pi' < \pi \in \mathcal{P}(c_\beta)$ and thus $c_\beta(\pi') \neq \bot$. The second application of Lemma 7.15 is justified by (2).

If $\gamma > \beta$ is a limit ordinal, then $g_\gamma = \liminf_{\iota\to\gamma} c_\iota$ and we can apply Corollary 4.7. Since $\pi' \in \mathcal{P}(c_\beta)$ and, by induction hypothesis, $c_\iota(\pi'') = c_\beta(\pi'')$ for all $\pi'' \leq \pi'$, $\beta \leq \iota < \gamma$, we thus obtain that $g_\gamma(\pi') = c_\beta(\pi')$. Since, according to (2), $\pi'' \notin \mathcal{P}_{g_\gamma}(n_\gamma)$ for each $\pi'' \leq \pi'$, we have by Lemma 7.15 that $g_\gamma(\pi') = c_\gamma(\pi')$. Hence, $c_\gamma(\pi') = c_\beta(\pi')$.

The inclusion $\sim_g \subseteq \sim$ follows immediately from Corollary 4.7 and the equality $P = \mathcal{P}(g)$ since $\sim_{c_\iota} \subseteq \sim_{g_\iota}$ for all $\iota < \alpha$ according to Lemma 7.11.

For the reverse inclusion $\sim \subseteq \sim_g$, assume that $\pi_1 \sim \pi_2$. That is, $\pi_1, \pi_2 \in P$ and there is some $\beta_0 < \alpha$ such that $\pi_1 \sim_{g_\iota} \pi_2$ for all $\beta_0 \leq \iota < \alpha$. Since $\pi_1, \pi_2 \in P = \mathcal{P}(g)$, we know, by Corollary 4.7, that there are $\beta_1, \beta_2 < \alpha$ such that $\pi_k \in \mathcal{P}(c_\iota)$ for all $\beta_k \leq \iota < \alpha$. Let $\beta = \max\{\beta_0, \beta_1, \beta_2\}$. For each $\beta \leq \iota < \alpha$, we then obtain that $\pi_1 \sim_{g_\iota} \pi_2$ and $\pi_1, \pi_2 \in \mathcal{P}(c_\iota)$. By Lemma 7.11, this is equivalent to $\pi_1 \sim_{c_\iota} \pi_2$. Applying Corollary 4.7 then yields $\pi_1 \sim_g \pi_2$.

Finally, we show that $l = g(\cdot)$. To this end, let $\pi \in P$. We distinguish two mutually exclusive cases. For the first case, we assume that

$$\text{there is some } \beta < \alpha \text{ such that } c_\iota(\pi) = c_\beta(\pi) \text{ for all } \beta \leq \iota < \alpha. \tag{3}$$

By Corollary 4.7, we know that then $g(\pi) = c_\beta(\pi)$. Next, assume that there is some $\beta' < \alpha$ with $\pi \notin \mathcal{P}_{g_\iota}(n_\iota)$ for all $\beta' \leq \iota < \alpha$. W.l.o.g. we can assume that $\beta = \beta'$. Hence, $l(\pi) = g_\beta(\pi)$. Moreover, since $\pi \notin \mathcal{P}_{g_\beta}(n_\beta)$, we have that $g_\beta(\pi) = c_\beta(\pi)$ according to Lemma 7.15. We thus conclude that $l(\pi) = g_\beta(\pi) = c_\beta(\pi) = g(\pi)$. Now assume there is no such β', i.e. for each $\beta' < \alpha$ there is some $\beta' \leq \iota < \alpha$ with $\pi \in \mathcal{P}_{g_\iota}(n_\iota)$. Consequently, $l(\pi) = \bot$ and, by Lemma 7.15, we have for each $\beta' < \alpha$ some $\beta' \leq \iota < \alpha$ such that $c_\iota(\pi) = \bot$. According to (3), the latter implies that $c_\iota(\pi) = \bot$ for all $\beta \leq \iota < \alpha$. By Corollary 4.7, we thus obtain that $g(\pi) = \bot = l(\pi)$.

Next, we consider the negation of (3), i.e. that

$$\text{for all } \beta < \alpha \text{ there is a } \beta \leq \iota < \alpha \text{ such that } \pi \in \mathcal{P}(c_\iota) \cap \mathcal{P}(c_\beta) \text{ implies } c_\iota(\pi) \neq c_\beta(\pi). \tag{4}$$

By Corollary 4.7, we have that $g(\pi) = \bot$. Since $\pi \in P = \mathcal{P}(g)$, we can apply Corollary 4.7 again to obtain a $\gamma < \alpha$ with $\pi \in \mathcal{P}(c_\iota)$ and $c_\iota(\pi') = c_\gamma(\pi')$ for all $\pi' < \pi$ and $\gamma \leq \iota < \alpha$. Combining this with (4) yields that for each $\gamma \leq \beta < \alpha$ there is a $\beta \leq \iota < \alpha$ with $c_\iota(\pi) \neq c_\beta(\pi)$. According to Lemma 7.18, this can only happen if there is a $\beta \leq \gamma' \leq \iota$ and a $\pi' \leq \pi$ such that $\pi' \in \mathcal{P}_{g_{\gamma'}}(n_{\gamma'})$. Since π has only finitely many prefixes, we can apply the infinite pigeon hole principle to obtain a single prefix $\pi' \leq \pi$ such that for each $\beta < \alpha$ there is some $\beta \leq \iota < \alpha$ with $\pi' \in \mathcal{P}_{g_\iota}(n_\iota)$. However, π' cannot be a proper prefix of π since this would imply that $\pi \notin P$. Thus we can conclude that for each $\beta < \alpha$ there is some $\beta \leq \iota < \alpha$ such that $\pi \in \mathcal{P}_{g_\iota}(n_\iota)$. Hence, $l(\pi) = \bot = g(\pi)$. ☐

The benefit of strong p-convergence over strong m-convergence is that the former has a more fine-grained characterisation of divergence. Strong p-convergence allows for local divergence, i.e. parts of a term graph that do not become persistent along a transfinite reduction. We will call such parts volatile:

Definition 7.20 (volatility)**.** Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$ be an open graph reduction. A position $\pi \in \mathbb{N}^*$ is said to be *volatile* in S if, for each $\beta < \alpha$, there is some $\beta \leq \gamma < \alpha$ such that $\pi \in \mathcal{P}_{g_\gamma}(n_\gamma)$. If π is volatile in S and no proper prefix of π is volatile in S, then π is called *outermost-volatile* in S.

As for infinitary term rewriting [5], local divergence in a strongly p-converging reduction can be characterised by volatile positions:

Lemma 7.21. *Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$ be an open reduction in a GRS strongly p-converging to g. Then, for every $\pi \in \mathbb{N}^*$, we have the following:*

(i) If π is volatile in S, then $\pi \in \mathcal{P}(g)$ implies $g(\pi) = \bot$.

(ii) $g(\pi) = \bot$ *iff*

(a) π *is outermost-volatile in* S*, or*

(b) there is some $\beta < \alpha$ *such that* $g_\beta(\pi) = \bot$ *and* $\pi' \notin \mathcal{P}_{g_\iota}(n_\iota)$ *for all* $\pi' \leq \pi$ *and* $\beta \leq \iota < \alpha$.

(iii) Let g_ι *be total for all* $\iota < \alpha$*. Then* $g(\pi) = \bot$ *iff* π *is outermost-volatile in* S.

Proof. (i) follows immediately from Lemma 7.19.

(ii) At first consider the "only if" direction: Suppose that $g(\pi) = \bot$. We will show that (iib) holds whenever (iia) does not hold. To this end, suppose that π is not outermost-volatile in S. Since $g(\pi) = \bot$, we know that $g(\pi') \in \Sigma$ for all $\pi' < \pi$. By Clause (i), this implies that no prefix of π is volatile. Consequently, π itself is not volatile in S either as it would be outermost-volatile otherwise. Hence, no prefix of π is volatile in S, i.e. there is some $\beta < \alpha$ such that $\pi' \notin \mathcal{P}_{g_\iota}(n_\iota)$ for all $\pi' \leq \pi, \beta \leq \iota < \alpha$. Additionally, by Lemma 7.19, we obtain that $g_\beta(\pi) = g(\pi) = \bot$. That is, (iib) holds.

For the "if" direction, we show that both (iia) and (iib) independently imply that $g(\pi) = \bot$: The implication from (iib) follows immediately from Lemma 7.19. For the implication from (iia), let π be outermost-volatile in S. Since no proper prefix of π is volatile in S, we find some $\beta < \alpha$ such that $\pi' \notin \mathcal{P}_{g_\iota}(n_\iota)$ for all $\pi' < \pi, \beta \leq \iota < \alpha$. Since π itself is volatile, there is some $\beta \leq \gamma < \alpha$ such that $\pi \in \mathcal{P}(g_\gamma)$. As we have, in particular, that $\pi' \notin \mathcal{P}_{g_\iota}(n_\iota)$ for all $\pi' < \pi, \gamma \leq \iota < \alpha$, we have, by Lemma 7.19, that $\pi \in \mathcal{P}(g)$. Consequently, according to Clause (i), we have that $g(\pi) = \bot$.

(iii) is a special case of (ii): If each g_ι is total, then (iib) cannot be true. $\square$

With this in mind, we can characterise total reductions as exactly those that lack volatile positions:

Lemma 7.22 (total reductions)**.** *Let* $\mathcal{R}$ *be a GRS,* g *a total term graph in* $\mathcal{R}$*, and* $S\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$. $S\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ *is total iff no prefix of* S *has a volatile position.*

Proof. The "only if" direction follows straightforwardly from Lemma 7.21.

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 g 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 7.21. $\square$

Next we want to compare strong m- and p-convergence with the ultimate goal of establishing the same relation between them as for term rewriting (cf. Theorem 2.5).

Definition 7.23 (minimal positions)**.** Let $g \in \mathcal{G}^\infty(\Sigma)$ and $n \in N^g$. A position $\pi \in \mathcal{P}_g(n)$ is called minimal if no proper prefix $\pi' < \pi$ is in $\mathcal{P}_g(n)$. The set of all minimal positions of n in g is denoted $\mathcal{P}^m_g(n)$.

Minimal positions have the nice property that they are not affected by term graph reductions:

Lemma 7.24. *Given a term graph reduction step* $g \to_{n,\rho,n'} h$*, we have* $\mathcal{P}^m_g(n) = \mathcal{P}^m_h(n')$.

Proof. We will show that $\mathcal{P}^m_g(n) \subseteq \mathcal{P}^m_h(n')$. The converse inclusion is symmetric. Let $\pi \in \mathcal{P}^m_g(n)$. Hence, $\pi' \notin \mathcal{P}_g(n)$ for all $\pi' < \pi$ which, by Lemma 7.15, implies that $\pi \in \mathcal{P}(g{\restriction}n)$. According to Lemma 7.10, $\pi \in \mathcal{P}_{g{\restriction}n}(n)$. Since, by Lemma 7.13, there is an isomorphism $\phi\colon g{\restriction}n \to h{\restriction}n'$ with $\phi(n) = n'$, we obtain, by Corollary 3.15, that $\pi \in \mathcal{P}_{h{\restriction}n'}(n')$. By Lemma 7.10, this implies $\pi \in \mathcal{P}_h(n')$. Since $\pi \in \mathcal{P}(h{\restriction}n')$, we know, by Lemma 7.15, that $\pi' \notin \mathcal{P}_h(n')$ for all $\pi' < \pi$. Combined, this means that $\pi \in \mathcal{P}^m_h(n')$. □

In order to compare strong m- and p-convergence, we consider positions bounded by a certain depth.

Definition 7.25 (bounded positions)**.** Let $g \in \mathcal{G}^\infty(\Sigma)$ and $d \in \mathbb{N}$. We write $\mathcal{P}_{\leq d}(g)$ for the set $\{\pi \in \mathcal{P}(g) \mid |\pi| \leq d\}$ of positions of length at most π.

Local truncations do not change positions bounded by the same depth or above:

Lemma 7.26. *Let $g \in \mathcal{G}^\infty(\Sigma_\perp)$, $n \in N^G$ and $d \leq \mathsf{depth}_g(n)$. Then $\mathcal{P}_{\leq d}(g{\restriction}n) = \mathcal{P}_{\leq d}(g)$.*

Proof. $\mathcal{P}_{\leq d}(g{\restriction}n) \subseteq \mathcal{P}_{\leq d}(g)$ follows from Lemma 7.15. For the converse inclusion, assume some $\pi \in \mathcal{P}_{\leq d}(g)$. Since $|\pi| \leq d \leq \mathsf{depth}_g(n)$, we know for each $\pi' < \pi$ that $|\pi'| < \mathsf{depth}_g(n)$ and thus $\pi' \notin \mathcal{P}_g(n)$. By Lemma 7.15, this implies that π is in $\mathcal{P}(g{\restriction}n)$ and thus also in $\mathcal{P}_{\leq d}(g{\restriction}n)$ □

Additionally, reductions that only contract redexes at a depth d or below do not affect the positions bounded by d.

Lemma 7.27. *Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$ be a strongly p-convergent reduction in a GRS and $d \in \mathbb{N}$ such that $\mathsf{depth}_{g_\iota}(n_\iota) \geq d$ for all $\iota < \alpha$. Then $\mathcal{P}_{\leq d}(g_0) = \mathcal{P}_{\leq d}(g_\iota)$ for all $\iota \leq \alpha$.*

Proof. We prove the statement by an induction on α. The case $\alpha = 0$ is trivial.

For $\alpha = \beta + 1$, let $g_\beta \to_{n_\beta, n'_\beta} g_\alpha$ be the β-th step of S. Due to the induction hypothesis, it suffices to show that $\mathcal{P}_{\leq d}(g_0) = \mathcal{P}_{\leq d}(g_\alpha)$. By Lemma 7.13, $g_\beta{\restriction}n_\beta \cong g_\alpha{\restriction}n'_\beta$, and by Lemma 7.24, $\mathsf{depth}_{g_\alpha}(n'_\beta) = \mathsf{depth}_{g_\beta}(n_\beta) \geq d$. Hence, according to Lemma 7.26, we have both $\mathcal{P}_{\leq d}(g_\beta{\restriction}n_\beta) = \mathcal{P}_{\leq d}(g_\beta)$ and $\mathcal{P}_{\leq d}(g_\alpha{\restriction}n'_\beta) = \mathcal{P}_{\leq d}(g_\alpha)$. We thus obtain the desired equation:

$$\mathcal{P}_{\leq d}(g_0) \overset{\text{ind. hyp.}}{=} \mathcal{P}_{\leq d}(g_\beta) \overset{\text{Lem. 7.26}}{=} \mathcal{P}_{\leq d}(g_\beta{\restriction}n_\beta) \overset{\text{Lem. 7.13}}{=} \mathcal{P}_{\leq d}(g_\alpha{\restriction}n'_\beta) \overset{\text{Lem. 7.26}}{=} \mathcal{P}_{\leq d}(g_\alpha)$$

Lastly, let α be a limit ordinal. By the induction hypothesis, we only need to show $\mathcal{P}_{\leq d}(g_0) = \mathcal{P}_{\leq d}(g_\alpha)$. At first assume $\pi \in \mathcal{P}_{\leq d}(g_\alpha)$. Hence, by Lemma 7.19, there is some $\beta < \alpha$ such that $\pi \in \mathcal{P}(g_\beta)$. Therefore, π is in $\mathcal{P}_{\leq d}(g_\beta)$ and, by induction hypothesis, also in $\mathcal{P}_{\leq d}(g_0)$. Conversely, assume that $\pi \in \mathcal{P}_{\leq d}(g_0)$. Because $\mathsf{depth}_{g_\iota}(n_\iota) \geq d$ for all $\iota < \alpha$, we have that $\pi' \notin \mathcal{P}_{g_\iota}(n_\iota)$ for all $\pi' < \pi$ and $\iota < \alpha$. According to Lemma 7.19, this implies that π is in $\mathcal{P}(g_\alpha)$ and thus also in $\mathcal{P}_{\leq d}(g_\alpha)$. □

The following two lemmas form the central properties that link strong m- and p-convergence:

Lemma 7.28. *Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$ be a strongly p-convergent open reduction in a GRS. If S has no volatile positions then $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\alpha}$ tends to infinity.*

Proof. We will prove the contraposition. To this end, assume that $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\alpha}$ does not tend to infinity. That is, there is some $d \in \mathbb{N}$ such that for each $\gamma < \alpha$ there is a $\gamma \leq \iota < \alpha$ with $\mathsf{depth}_{g_\iota}(n_\iota) \leq d$. Let d^* be the smallest such d. Hence, there is a $\beta < \alpha$ such that $\mathsf{depth}_{g_\iota}(n_\iota) \geq d^*$ for all $\beta \leq \iota < \alpha$. Thus we can apply Lemma 7.27 to the suffix of S starting from β to obtain that $\mathcal{P}_{\leq d^*}(g_\beta) = \mathcal{P}_{\leq d^*}(g_\iota)$ for all $\beta \leq \iota < \alpha$. Since we find for each $\gamma < \alpha$ some $\gamma \leq \iota < \alpha$ with $\mathsf{depth}_{g_\iota}(n_\iota) \leq d^*$, we know that for each $\gamma < \alpha$ there is a $\gamma \leq \iota < \alpha$ and a $\pi \in \mathcal{P}_{\leq d^*}(g_\beta)$ with $\pi \in \mathcal{P}_{g_\iota}(n_\iota)$. Because $\mathcal{P}_{\leq d^*}(g_\beta)$ is finite, the infinite pigeon hole principle yields single $\pi^* \in \mathcal{P}_{\leq d^*}(g_\beta)$ such that for each $\gamma < \alpha$ there is a $\gamma \leq \iota < \alpha$ with $\pi^* \in \mathcal{P}_{g_\iota}(n_\iota)$. That is, π^* is volatile in S. ☐

Lemma 7.29. *Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$ be an open reduction in a GRS strongly p-converging to g. If $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\alpha}$ tends to infinity, then $g \cong \lim_{\iota\to\alpha} g_\iota$.*

Proof. Let $h = \lim_{\iota\to\alpha} g_\iota$ and let $(c_\iota)_{\iota<\alpha}$ be the reduction contexts of S. We will prove that $g \cong h$ by showing that their respective labelled quotient trees coincide.

For the inclusion $\mathcal{P}(g) \subseteq \mathcal{P}(h)$, assume some $\pi \in \mathcal{P}(g)$. According to Corollary 4.7, there is some $\beta < \alpha$ such that $\pi \in \mathcal{P}(c_\beta)$ and $c_\iota(\pi) = c_\beta(\pi)$ for all $\pi' < \pi$ and $\beta \leq \iota < \alpha$. Thus, $\pi \in \mathcal{P}(c_\iota)$ for all $\beta \leq \iota < \alpha$. Since $c_\iota \cong g_\iota \upharpoonright n_\iota$ and, therefore, by Lemma 7.15, $\mathcal{P}(c_\iota) = \mathcal{P}(g_\iota)$, we have that $\pi \in \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$. This implies, by Theorem 5.9, that $\pi \in \mathcal{P}(h)$.

For the converse inclusion $\mathcal{P}(h) \subseteq \mathcal{P}(g)$, assume some $\pi \in \mathcal{P}(h)$. According to Theorem 5.9, there is some $\beta < \alpha$ such that $\pi \in \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$. Since $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\alpha}$ tends to infinity, we find some $\beta \leq \gamma < \alpha$ such that $\mathsf{depth}_{g_\iota}(n_\iota) \geq |\pi|$ for all $\gamma \leq \iota < \alpha$, i.e. $\pi' \notin \mathcal{P}_{g_\iota}(n_\iota)$ for all $\pi' < \pi$. This means, by Lemma 7.19, that $\pi \in \mathcal{P}(g)$.

By Lemma 7.19 and Theorem 5.9, $\sim_g = \sim_h$ follows from the equality $\mathcal{P}(g) = \mathcal{P}(h)$.

In order to show the equality $g(\cdot) = g(\cdot)$, assume some $\pi \in \mathcal{P}(h)$. According to Theorem 5.9, there is some $\beta < \alpha$ such that $h(\pi) = g_\iota(\pi)$ for all $\beta \leq \iota < \alpha$. Additionally, since $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\alpha}$ tends to infinity, there is some $\beta \leq \gamma < \alpha$ such that $\mathsf{depth}_{g_\iota}(n_\iota) > |\pi|$ for all $\gamma \leq \iota < \alpha$, i.e. $\pi \notin \mathcal{P}_{g_\iota}(n_\iota)$. Thus, by Lemma 7.19, $g(\pi) = g_\gamma(\pi)$. Since $h(\pi) = g_\gamma(\pi)$, we can conclude that $g(\pi) = h(\pi)$. ☐

The following property, which relates strong m-convergence and -continuity, follows from the fact that our notion of strong m-convergence on term graphs instantiates the abstract model of strong m-convergence from our previous work [4]:

Lemma 7.30. *Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$ be an open strongly m-continuous reduction in a GRS. If $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\alpha}$ tends to infinity, then S is strongly m-convergent.*

Proof. This is a special case of Proposition 5.5 from [4]. ☐

Now, we have everything in place to prove that strong p-convergence conservatively extends strong m-convergence.

Theorem 7.31. *Let $\mathcal{R}$ be a GRS and S a reduction in $\mathcal{R}$. We then have that*

$$S\colon g \overset{m}{\twoheadrightarrow}_{\mathcal{R}} h \quad \text{iff} \quad S\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h \text{ is total.}$$

Proof. Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$. At first, we prove the "only if" direction by induction on α:

The case $\alpha = 0$ is trivial. If α is a successor ordinal, the statement follows immediately from the induction hypothesis.

Let α be a limit ordinal. Since $S\colon g \overset{m}{\twoheadrightarrow} g_\alpha$, we know that $S|_\gamma\colon g \overset{m}{\twoheadrightarrow} g_\gamma$ for all $\gamma < \alpha$. Hence, we can apply the induction hypothesis in order to obtain that $S|_\gamma\colon g \overset{p}{\twoheadrightarrow} g_\gamma$ for each $\gamma < \alpha$. Consequently, S is strongly p-continuous, which means, by Proposition 7.17, that S strongly p-converges to some term graph h'. However, since S strongly m-converges, we know that $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\alpha}$ tends to infinity. Consequently, we can apply Lemma 7.29 to obtain that $h' = \lim_{\iota\to\alpha} = h$, i.e. $S\colon g \overset{p}{\twoheadrightarrow} h$. Since $S\colon g \overset{m}{\twoheadrightarrow} h$, of course, $S\colon g \overset{p}{\twoheadrightarrow} h$ must be total.

We will also prove the "if" direction by induction on α: Again, the case $\alpha = 0$ is trivial and the case that α is a successor ordinal follows immediately from the induction hypothesis.

Let α be a limit ordinal. Since S is strongly p-convergent, we know that $S|_\gamma\colon g \overset{p}{\twoheadrightarrow} g_\gamma$ is total for each $\gamma < \alpha$. Therefore, we can apply the induction hypothesis to obtain that $S|_\gamma\colon g \overset{m}{\twoheadrightarrow} g_\gamma$ for each $\gamma < \alpha$. Hence, S is strongly m-continuous. Since S is total, we know from Lemma 7.21, that S has no volatile positions. Hence, by Lemma 7.28, $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\alpha}$ tends to infinity. Together with the strong m-continuity of S, this yields, according to Lemma 7.30, that S strongly m-converges to some h'. With Lemma 7.29, we can then conclude that $h' = h$, i.e. $S\colon g \overset{m}{\twoheadrightarrow} h$. $\square$

8 Terms vs. Term Graphs

Term graph rewriting is an efficient implementation technique for term rewriting that uses pointers in order to avoid duplication. This is also used as the basis for the implementation of functional programming languages. A prominent example is the implementation of the fixed point combinator Y defined by the term rule $\rho_0\colon Y\,x \to x\,(Y\,x)$, where we write function application as juxtaposition. Written as a term graph rule $\rho_1\colon Y\,x \to x\,(Y\,x)$ depicted in Figure 6a, we can see that the two occurrences of the variables x on the right-hand side are shared. In fact, since term graph rewriting does not provide a mechanism for duplication, this is the only way to represent non-right linear rules. With the rule ρ_1 applied repeatedly as shown in Figure 6c, more and more pointers to the same occurrence of the function symbol f are created. This reduction strongly m-converges to the infinite term graph $g_\omega = f^{[n]}\,(n\,(n\,(\dots)))$, which has infinitely many edges to the f-node. Note, however, that the term graph rule ρ_1 is not maximally shared. If we apply the maximally shared rule $\rho_2\colon (Y\,x)^{[n]} \to x\,n$, we obtain in one step the cyclic term graph $h_0 = (f\,n)^{[n]}$. This is, in fact, how the fixed point combinator is typically implemented in functional programming languages [17, 20]. Although, the resulting term graphs g_ω and h_0 are different, they both unravel to the same term $f\,(f\,(f\,(\dots)))$.

In this section, we will study the relationship between GRSs and the corresponding TRSs they simulate. In particular, we will show the soundness of GRSs w.r.t. strong convergence and a restricted form of completeness. To this end we make use of the isomorphism between terms and canonical term trees as outlined at the end of Section 3.2.

Note that term trees have an obvious characterisation in terms of their equivalence on positions:

Fact 8.1. *A term graph $g \in \mathcal{G}^\infty(\Sigma)$ is a term tree iff $\sim_g$ is the identity relation, i.e. $\pi_1 \sim_g \pi_2$ iff $\pi_1 = \pi_2$.*

When giving the labelled quotient tree for a term tree t we can thus omit the equivalence $\sim_t$. We refer to the remaining pair $(\mathcal{P}(t), t(\cdot))$ as *labelled tree.*

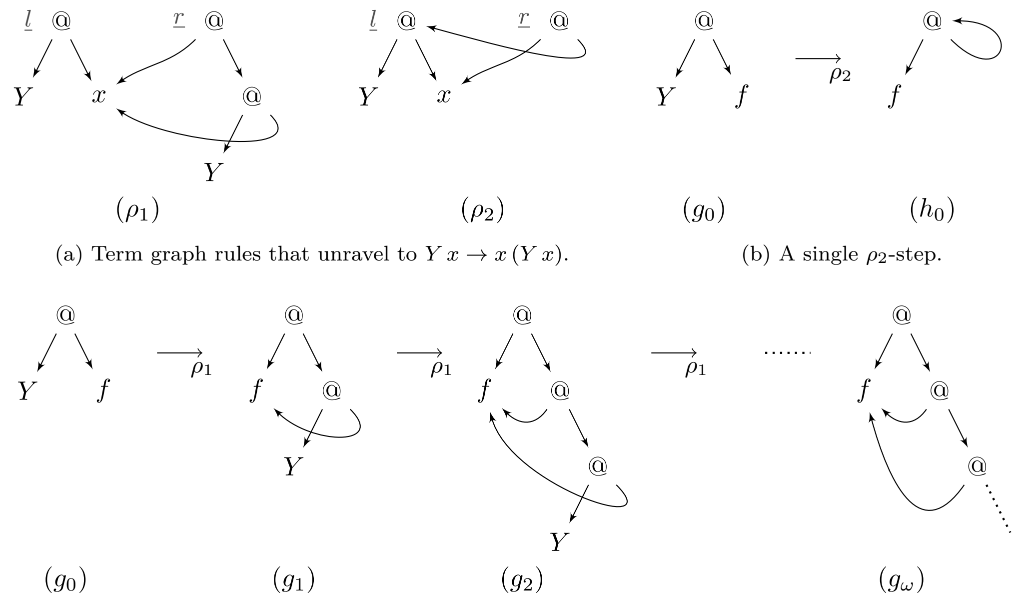

(a) Term graph rules that unravel to $Y\,x \to x\,(Y\,x)$.

(b) A single ρ_2-step.

(c) A strongly m-convergent term graph reduction over ρ_1.

Figure 6: Implementation of the fixed point combinator as a term graph rewrite rule.

Recall that the unravelling $\mathcal{U}(g)$ of a term graph g is the uniquely determined term t such that there is a homomorphism from t to g. Labelled trees give a concrete characterisation of unravellings:

Proposition 8.2. *The unravelling $\mathcal{U}(g)$ of a term graph $g \in \mathcal{G}^\infty(\Sigma)$ is given by the labelled tree (P, l) with $P = \mathcal{P}(g)$ and $l(\pi) = g(\pi)$ for all $\pi \in P$.*

Proof. Since the implicit equivalence $\sim_{\mathcal{U}(g)}$ is reflexive and a subrelation of $\sim_g$, the triple $(P, l, \sim_{\mathcal{U}(g)})$ is a labelled quotient tree. Let t be the term represented by (P, l). By Lemma 3.14, there is a homomorphism from t to g. Thus, $\mathcal{U}(g) = t$. $\square$

Before start investigating the correspondences between term rewriting and term graph rewriting, we need to transfer the notions of left-linearity and orthogonality to GRSs:

Definition 8.3 (left-linearity, orthogonality [7]). Let $\mathcal{R} = (\Sigma, R)$ be a GRS.

(i) A rule $\rho \in R$ is called *left-linear* if its left-hand side ρ_l is a term tree. The GRS $\mathcal{R}$ is called *left-linear* if all its rules are left-linear.

(ii) A ρ-redex g and a ρ'-redex g' in a common term graph, with matching $\mathcal{V}$-homomorphisms ϕ resp. ϕ' are *disjoint*, if $r^g \neq \phi'(n)$ for all non-$\mathcal{V}$ nodes n in ρ'_l and, symmetrically, $r^{g'} \neq \phi(n)$ for all non-$\mathcal{V}$ nodes n in ρ_l. In other words, the root of either redexes must not be matched by the respective other rule.

(iii) The GRS $\mathcal{R}$ is called *non-overlapping* if all its redexes are pairwise disjoint.

(iv) The GRS $\mathcal{R}$ is called *orthogonal* if it is left-linear and non-overlapping.

It is obvious that the unravelling $\mathcal{U}(\mathcal{R})$ of a GRS is left-linear if $\mathcal{R}$ is left-linear, that and $\mathcal{U}(\mathcal{R})$ is orthogonal if $\mathcal{R}$ is orthogonal.

We have to single out a particular kind of term graph redexes that manifest a peculiar behaviour.

Definition 8.4 (circular redex)**.** Let $\rho = (g, l, r)$ be a term graph rule. A ρ-redex is called *circular* if l and r are distinct but the matching $\mathcal{V}$-homomorphism ϕ maps them to the same node, i.e. $l \neq r$ but $\phi(l) = \phi(r)$.

Kennaway et al. [15] show that circular redexes only reduce to themselves:

Proposition 8.5. *For every circular ρ-redex $g|_n$, we have $g \to_{\rho,n} h$.*

However, contracting the unravellings of a circular redex also yields the same term:

Lemma 8.6. *Let g be a term graph with a circular ρ-redex rooted in n. Then $\mathcal{U}(g) \to_{\mathcal{U}(\rho),\pi} \mathcal{U}(g)$ for all $\pi \in \mathcal{P}_g(n)$.*

Proof. Since there is a circular ρ-redex, we know that the right-hand side root r^ρ is reachable from the left-hand side root l^ρ of ρ. Let π^* be a path from l^ρ to r^ρ. Because $g|_n$ is a circular redex, the corresponding matching $\mathcal{V}$-homomorphism maps both l^ρ and r^ρ to n. Since Δ-homomorphisms preserve paths, we thus know that π^* is also a path from n to itself in g. In other words $\pi \in \mathcal{P}_g(n)$ implies $\pi \cdot \pi^* \in \mathcal{P}_g(n)$. Consequently, for each $\pi \in \mathcal{P}_g(n)$ we have that $\mathcal{U}(g)|_\pi = \mathcal{U}(g)|_{\pi\cdot\pi^*}$.

Since there is a path π^* from l^ρ to r^ρ, the unravelling $\mathcal{U}(\rho)$ of ρ is of the form $l \to l|_{\pi^*}$. Hence, we know that each application of $\mathcal{U}(\rho)$ at a position π in some term t replaces the subterm at π with the subterm at $\pi \cdot \pi^*$ in t, i.e. $t \to_{\mathcal{U}(\rho),\pi} t[t|_{\pi\cdot\pi^*}]_\pi$.

Combining the two findings above, we obtain that

$$\mathcal{U}(g) \to_{\mathcal{U}(\rho),\pi} \mathcal{U}(g)\left[\mathcal{U}(g)|_{\pi\cdot\pi^*}\right]_\pi = \mathcal{U}(g)\left[\mathcal{U}(g)|_\pi\right]_\pi = \mathcal{U}(g) \quad \text{for all } \pi \in \mathcal{P}_g(n)$$

$\square$

The following two properties due to Kennaway et al. [15] show how single term graph rewrite steps relate to term reductions in the corresponding unravelling.

Proposition 8.7. *Given a left-linear GRS $\mathcal{R}$ and a term graph g in $\mathcal{R}$, it holds that g is a normal form in $\mathcal{R}$ iff $\mathcal{U}(g)$ is a normal form in $\mathcal{U}(\mathcal{R})$.*

Theorem 8.8. *Let $\mathcal{R}$ be a left-linear GRS with a reduction step $g \to_{n,\rho} h$. Then $S\colon \mathcal{U}(g) \xrightarrow{m}\!\!\!\!\to_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$ such that the depth of every redex reduced in S is greater or equal to $\mathsf{depth}_g(n)$. In particular, if the ρ-redex $g|_n$ is not circular, then S is a complete development of the set of redex occurrences $\mathcal{P}_g(n)$ in $\mathcal{U}(g)$.*

The goal of the following two sections is to generalise the above soundness theorem to strong m- and p-convergence.

8.1 Strong m-Convergence

At first we shall study correspondences w.r.t. strong m-convergence. To this end, we first show that the metric $\mathbf{d}_\restriction$ on term graphs generalises the metric $\mathbf{d}$ on terms.

Lemma 8.9. *Let $t \in \mathcal{T}^\infty(\Sigma_\perp)$ and $d \in \mathbb{N} \cup \{\infty\}$. The strict truncation $t \restriction d$ is given by the labelled tree (P, l) with*

$$\text{(a) } P = \{\pi \in \mathcal{P}(t) \mid |\pi| \leq d\}, \qquad \text{(b) } l(\pi) = \begin{cases} t(\pi) & \text{if } |\pi| < d \\ \perp & \text{if } |\pi| \geq d \end{cases}$$

Proof. Immediate from Lemma 5.7 and Fact 8.1. □

This shows that the metric $\mathbf{d}_\restriction$ restricted to terms coincides with the metric $\mathbf{d}$ on terms. Moreover, we can use this in order to relate the metric distance between term graphs and the metric distance between their unravellings.

Lemma 8.10. *For all $g, h \in \mathcal{G}^\infty(\Sigma)$, we have that $\mathbf{d}_\restriction(g, h) \geq \mathbf{d}_\restriction(\mathcal{U}(g), \mathcal{U}(h))$.*

Proof. Let $d = \mathsf{sim}_\restriction(g, h)$. Hence, $g \restriction d \cong h \restriction d$ and we can assume that the corresponding labelled quotient trees as characterised by Lemma 5.7 coincide. We only need to show that $\mathcal{U}(g) \restriction d \cong \mathcal{U}(h) \restriction d$ since then $\mathsf{sim}_\restriction(\mathcal{U}(g), \mathcal{U}(h)) \geq d$ and thus $\mathbf{d}_\restriction(\mathcal{U}(g), \mathcal{U}(h)) \leq 2^{-d} = \mathbf{d}_\restriction(g, h)$. In order to show this, we show that the labelled trees of $\mathcal{U}(g) \restriction d$ and $\mathcal{U}(h) \restriction d$ as characterised by Lemma 8.9 coincide. For the set of positions we have the following:

$$\begin{aligned}
&\pi \in \mathcal{P}(\mathcal{U}(g) \restriction d) \\
\iff &\pi \in \mathcal{P}(\mathcal{U}(g)), \quad |\pi| \leq d && \text{(Lemma 8.9)} \\
\iff &\pi \in \mathcal{P}(g), \quad |\pi| \leq d && \text{(Proposition 8.2)} \\
\iff &\pi \in \mathcal{P}(g \restriction d), \quad |\pi| \leq d && \text{(Lemma 5.7)} \\
\iff &\pi \in \mathcal{P}(h \restriction d), \quad |\pi| \leq d && (g \restriction d \cong h \restriction d) \\
\iff &\pi \in \mathcal{P}(h), \quad |\pi| \leq d && \text{(Lemma 5.7)} \\
\iff &\pi \in \mathcal{P}(\mathcal{U}(h)), \quad |\pi| \leq d && \text{(Proposition 8.2)} \\
\iff &\pi \in \mathcal{P}(\mathcal{U}(h) \restriction d) && \text{(Lemma 8.9)}
\end{aligned}$$

In order to show that the labellings are equal, consider some $\pi \in \mathcal{P}(\mathcal{U}(g) \restriction d)$ and assume at first that $|\pi| \geq d$. By Lemma 8.9, we then have $(\mathcal{U}(g) \restriction d)(\pi) = \perp = (\mathcal{U}(h) \restriction d)(\pi)$. Otherwise, if $|\pi| < d$, we obtain that

$$\begin{aligned}
(\mathcal{U}(g) \restriction d)(\pi) &\overset{\text{Lem. 8.9}}{=} \mathcal{U}(g)(\pi) \overset{\text{Prop. 8.2}}{=} g(\pi) \overset{\text{Lem. 5.7}}{=} g \restriction d(\pi) \\
&\overset{g \restriction d \cong h \restriction d}{=} h \restriction d(\pi) \overset{\text{Lem. 5.7}}{=} h(\pi) \overset{\text{Prop. 8.2}}{=} \mathcal{U}(h)(\pi) \overset{\text{Lem. 8.9}}{=} (\mathcal{U}(h) \restriction d)(\pi)
\end{aligned}$$

□

This immediately yields that Cauchy sequences are preserved by unravelling:

Lemma 8.11. *If $(g_\iota)_{\iota<\alpha}$ is a Cauchy sequence in $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\restriction)$, then so is $(\mathcal{U}(g_\iota))_{\iota<\alpha}$.*

Proof. This follows immediately from Lemma 8.10. □

Additionally, also limits are preserved by unravellings.

Proposition 8.12. *For every sequence $(g_\iota)_{\iota<\alpha}$ in $(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\restriction)$, we have that $\lim_{\iota \to \alpha} g_\iota = g$ implies $\lim_{\iota \to \alpha} \mathcal{U}(g_\iota) = \mathcal{U}(g)$*

Proof. According to Theorem 5.9, we have that $\mathcal{P}(g) = \liminf_{\iota \to \alpha} \mathcal{P}(g_\iota)$, and that $g(\pi) = g_\beta(\pi)$ for some $\beta < \alpha$ with $g_\iota(\pi) = g_\beta(\pi)$ for all $\beta \leq \iota < \alpha$. By Proposition 8.2, we then obtain $\mathcal{P}(\mathcal{U}(g)) = \liminf_{\iota \to \alpha} \mathcal{P}(\mathcal{U}(g_\iota))$, and that $\mathcal{U}(g)(\pi) = \mathcal{U}(g_\beta)(\pi)$ for some $\beta < \alpha$ with $\mathcal{U}(g_\iota)(\pi) = \mathcal{U}(g_\beta)(\pi)$ for all $\beta \leq \iota < \alpha$. Since by Lemma 8.11, $(\mathcal{U}(g_\iota))_{\iota<\alpha}$ is Cauchy, we can apply Theorem 5.9 to obtain that $\lim_{\iota \to \alpha} \mathcal{U}(g_\iota) = \mathcal{U}(g)$. $\square$

We can now show that term graph reductions are sound w.r.t. reductions in the unravelled system.

Theorem 8.13. *Let $\mathcal{R}$ be a left-linear GRS. If $g \overset{m}{\twoheadrightarrow}_{\mathcal{R}} h$, then $\mathcal{U}(g) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$.*

Proof. Let $S = (g_\iota \to_{d_\iota} g_{\iota+1})_{\iota<\alpha}$ be a reduction strongly m-converging to g_α in $\mathcal{R}$, i.e. $S\colon g_0 \overset{m}{\twoheadrightarrow}_{\mathcal{R}} g_\alpha$. According to Theorem 8.8, there is for each $\iota < \alpha$ a reduction $T_\iota\colon \mathcal{U}(g_\iota) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_{\iota+1})$ such that

$$\text{all steps in } T_\iota \text{ contract a redex at depth} \geq d_\iota. \tag{$*$}$$

Define for each $\delta \leq \alpha$ the concatenation $U_\delta = \prod_{\iota<\delta} T_\iota$. We will show that $U_\delta\colon \mathcal{U}(g_0) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_\delta)$ for each $\delta \leq \alpha$ by induction on δ. The theorem is then obtained by instantiating $\delta = \alpha$.

The case $\delta = 0$ is trivial. If $\delta = \delta' + 1$, then we have by induction hypothesis that $U_{\delta'}\colon \mathcal{U}(g_0) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_{\delta'})$. Since $T_{\delta'}\colon \mathcal{U}(g_{\delta'}) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_\delta)$, and $U_\delta = U_{\delta'} \cdot T_{\delta'}$, we have $U_\delta\colon \mathcal{U}(g_0) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_\delta)$.

For the case that δ is a limit ordinal, let $U_\delta = (t_\iota \to_{e_\iota} t_{\iota+1})_{\iota<\beta}$. For each $\gamma < \beta$ we find some $\gamma' < \delta$ with $U_\delta|_\gamma < U_{\gamma'}$. By induction hypothesis, we can assume that $U_{\gamma'}$ is strongly m-continuous. According to Proposition 2.4, this means that the proper prefix $U_\delta|_\gamma$ strongly m-converges to t_γ. This shows that each proper prefix $U_\delta|_\gamma$ of U_δ strongly m-converges to t_γ. Hence, by Proposition 2.4, U_δ is strongly m-continuous.

Since S is strongly m-convergent, $(d_\iota)_{\iota<\delta}$ tends to infinity. By ($*$), also $(e_\iota)_{\iota<\alpha}$ tends to infinity. Hence, U_δ is strongly m-convergent according to Proposition 2.2. Let t be the term U_δ is strongly m-converging to, i.e. $\lim_{\iota \to \beta} t_\iota = t$. Since $(\mathcal{U}(g_\iota))_{\iota<\delta}$ is a cofinal subsequence of $(t_\iota)_{\iota<\beta}$, we have by Proposition 1.1 that $\lim_{\iota \to \delta} \mathcal{U}(g_\iota) = t$. Since S is strongly m-convergent, we also have that $\lim_{\iota \to \delta} g_\iota = g_\delta$. By Proposition 8.12, this yields that $\lim_{\iota \to \delta} \mathcal{U}(g_\iota) = \mathcal{U}(g_\delta)$. Consequently, we have that $t = \mathcal{U}(g_\delta)$, i.e. $U_\delta\colon \mathcal{U}(g_0) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_\delta)$. $\square$

Unfortunately, we will not be able to obtain a full completeness result. Even the weak completeness that was considered by Kennaway et al. [15] for finitary term graph reductions does not hold for infinitary term graph reductions. This weaker completeness property is satisfied by infinitary term graph rewriting iff

$$\mathcal{U}(g) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t \implies \text{there is a term graph } h \text{ with } g \overset{m}{\twoheadrightarrow}_{\mathcal{R}} h, \quad t \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$$

Kennaway et al. [15] consider an informal notion of infinitary term graph rewriting and give a counterexample for the above weak completeness property. This counterexample also applies to strongly m-convergent term graph reductions:

Example 8.14. We consider an infinite alphabet Σ with $\underline{n} \in \Sigma^{(2)}$ for each $n \in \mathbb{N}$. Let g be the term graph depicted in Figure 7a. The root node is labelled $\underline{0}$ and each node labelled $\underline{n}$ has as its left successor itself and as its right successor a node labelled $\underline{n+1}$. Let $\mathcal{R}$ be the GRS that for each natural number n has a rule $\rho_n\colon \underline{n}(x, y) \to \underline{n+1}(x, y)$. A single reduction

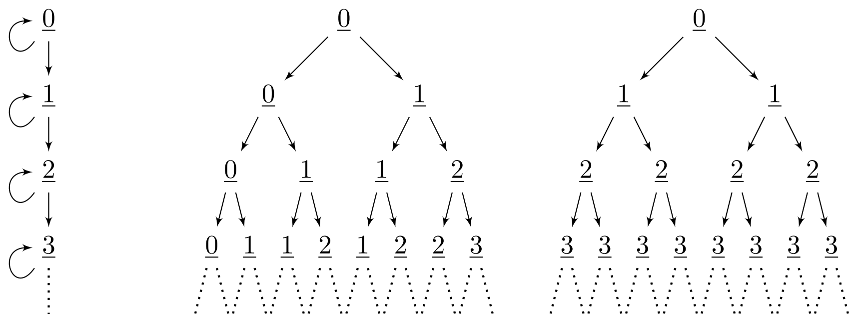

(a) A term graph g. (b) The unravelling of g. (c) Term reduct t of $\mathcal{U}(g)$.

Figure 7: Counterexample for weak completeness.

step in $\mathcal{R}$ increments the label of exactly one node. Figure 7b shows the unravelling of g. In each row of $\mathcal{U}(g)$, the rightmost node has the largest label. However, each node to its left can be incremented by performing finitely many reduction steps in the TRS $\mathcal{U}(\mathcal{R})$ so that it has the same label as the rightmost node. Doing this for each row yields a reduction strongly m-converging to the term t depicted in Figure 7c. Note that for each $n \in \mathbb{N}$, there are only finitely many occurrences of $\underline{n}$ in t. Therefore, also the number of occurrences of labels $\underline{m}$ with $m < n$ is finite for each $n \in \mathbb{N}$. Since it is only possible to obtain a node labelled $\underline{n}$ by repeated reduction on a node labelled $\underline{m}$ with $m < n$, this means that every term t' with $t \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t'$ also has only finitely many occurrences of $\underline{n}$ for any $n \in \mathbb{N}$. On the other hand, there is no strongly m-converging reduction from g to a term graph that unravels to a term with finitely many occurrences of $\underline{n}$ for each $n \in \mathbb{N}$. This is because the structure of g cannot be changed by reductions in $\mathcal{R}$. In particular, the loops in g are maintained.

The above counterexample requires an infinite set of function symbols and rules. Kennaway et al. [15] sketch a variant of this example that gets along with only two function symbols and one rule. However, after closer inspection one can see that this system is not a counterexample! We do not know whether a restriction to finitely many rules may in fact yield weak completeness of infinitary term graph rewriting.

We think, however, that a completeness property w.r.t. normalising reductions can be obtained. To this end, consider the following property of strong m-convergence in TRSs:

Theorem 8.15 ([16])**.** *Every orthogonal TRS has the normal form property w.r.t. strong m-convergence. That is, for each term t with $t \overset{m}{\twoheadrightarrow} t_1$ and $t \overset{m}{\twoheadrightarrow} t_2$, we have $t_1 = t_2$, whenever t_1, t_2 are normal forms.*

We then obtain as a corollary that a term graph has the same normal forms as its unravelling, provided it has one:

Corollary 8.16. *For every orthogonal GRS $\mathcal{R}$, we have that two normalising reductions $g \overset{m}{\twoheadrightarrow}_{\mathcal{R}} h$ and $\mathcal{U}(g) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t$ imply that $t = \mathcal{U}(h)$.*

Proof. According to Theorem 8.13, $g \overset{m}{\twoheadrightarrow}_{\mathcal{R}} h$ implies $\mathcal{U}(g) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$. Note that, by Proposition 8.7, $\mathcal{U}(h)$ is a normal form in $\mathcal{U}(\mathcal{R})$. Hence, the reduction $\mathcal{U}(g) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$ together with $\mathcal{U}(g) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t$ implies $t = \mathcal{U}(h)$ according to Theorem 8.15. $\square$

We conjecture that this can be generalised such that infinitary term graph rewriting is complete w.r.t. normalising reductions. That is, if $\mathcal{U}(g) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t$ with t a normal form of $\mathcal{U}(\mathcal{R})$, then there is a term graph h with $\mathcal{U}(h) = t$ and $g \overset{m}{\twoheadrightarrow}_{\mathcal{R}} h$.

8.2 Strong p-Convergence

In this section, we replicate the results that we have obtained in the preceding section for the case of strong p-convergence. Since p-convergence is a conservative extension of m-convergence, cf. Theorems 2.5 and 7.31, this will in fact generalise the soundness and completeness results for infinitary term graph rewriting.

At first we derive a characterisation of the partial order $\leq_\bot^{\mathcal{G}}$ on terms:

Lemma 8.17. *Given two terms $s, t \in \mathcal{T}^\infty(\Sigma_\bot)$, we have $s \leq_\bot^{\mathcal{G}} t$ iff $s(\pi) = t(\pi)$ for all $\pi \in \mathcal{P}(s)$ with $g(\pi) \in \Sigma$.*

Proof. Immediate from Corollary 4.3 and Fact 8.1. ☐

This shows that the partial order $\leq_\bot^{\mathcal{G}}$ on term graphs generalises the partial order on terms.

From this we easily obtain that the partial order $\leq_\bot^{\mathcal{G}}$ as well as its induced limits are preserved by unravelling:

Theorem 8.18. *In the partially ordered set $(\mathcal{G}_{\mathcal{C}}^\infty(\Sigma_\bot), \leq_\bot^{\mathcal{G}})$ the following holds:*

(i) Given two term graphs g, h, we have that $g \leq_\bot^{\mathcal{G}} h$ implies $\mathcal{U}(g) \leq_\bot^{\mathcal{G}} \mathcal{U}(h)$.

(ii) For each directed set G, we have that $\mathcal{U}\left(\bigsqcup_{g\in G} g\right) = \bigsqcup_{g \in G} \mathcal{U}(g)$.

(iii) For each non-empty set G, we have that $\mathcal{U}\left(\bigsqcap_{g\in G} g\right) = \bigsqcap_{g \in G} \mathcal{U}(g)$.

(iv) For each sequence $(g_\iota)_{\iota<\alpha}$, we have that $\mathcal{U}(\liminf_{\iota\to\alpha} g_\iota) = \liminf_{\iota\to\alpha} \mathcal{U}(g_\iota)$.

Proof. (i) By Corollary 4.3, $g \leq_\bot^{\mathcal{G}} h$ implies that $g(\pi) = h(\pi)$ for all $\pi \in \mathcal{P}(g)$ with $g(\pi) \in \Sigma$. By Proposition 8.2, we then have $\mathcal{U}(g)(\pi) = \mathcal{U}(h)(\pi)$ for all $\pi \in \mathcal{P}(\mathcal{U}(g))$ with $\mathcal{U}(g)(\pi) \in \Sigma$ which, by Lemma 8.17, implies $\mathcal{U}(g) \leq_\bot^{\mathcal{G}} \mathcal{U}(h)$.

By a similar argument (ii) and (iii) follow from the characterisation of least upper bounds and greatest lower bounds in Theorem 4.4 resp. Proposition 4.5 by using Proposition 8.2.

(iv) Follows from (ii) and (iii). ☐

In order to proof the soundness w.r.t. strong p-convergence we need a stronger variant of Theorem 8.8 that does not only make a statement about the depth of the redexes contracted in the term reduction, but also the corresponding reduction contexts.

Theorem 8.19. *Let $\mathcal{R}$ be a left-linear GRS with a reduction step $g \to_c h$. Then there is a non-empty reduction $S = (t_\iota \to_{c_\iota} t_{\iota+1})_{\iota<\alpha}$ with $S\colon \mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$ such that $\mathcal{U}(c) = \bigsqcap_{\iota<\alpha} c_\iota$.*

Proof. By Theorem 8.8, there is a reduction $S\colon \mathcal{U}(g) \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$. At first we assume that the redex $g|_n$ contracted in $g \to_n h$ is not a circular redex. Hence S, is complete development of the set of redex occurrences $\mathcal{P}_g(n)$ in $\mathcal{U}(g)$. By Theorem 2.5, we then obtain $S\colon \mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$. From Lemma 7.15 and Proposition 8.2 it follows that $\mathcal{U}(g{\upharpoonright}n)$

is obtained from $\mathcal{U}(g)$ by replacing each subterm of $\mathcal{U}(g)$ at a position in $\mathcal{P}^m_g(n)$, i.e. a minimal position of n, by $\bot$. Since each step $t_\iota \to_{\pi_\iota} t_{\iota+1}$ in S contracts a redex at a position π_ι that has a prefix in $\mathcal{P}^m_g(n)$, we have $\mathcal{U}(g{\restriction}n) \leq^{\mathcal{G}}_\bot t_\iota[\bot]_{\pi_\iota} = c_\iota$. Moreover, for each $\pi \in \mathcal{P}^m_g(n)$ there is a step at $\iota(\pi) < \alpha$ in S that takes place at π. From Proposition 4.5, it is thus clear that $\mathcal{U}(g{\restriction}n) = \bigsqcap_{\pi\in P} c_{\iota(\pi)}$. Together with $\mathcal{U}(g{\restriction}n) \leq^{\mathcal{G}}_\bot c_\iota$ for all $\iota < \alpha$, this yields $\mathcal{U}(g{\restriction}n) = \bigsqcap_{\iota<\alpha} c_\iota$. Then $\mathcal{U}(c) = \bigsqcap_{\iota<\alpha} c_\iota$ follows from the fact that $c \cong g{\restriction}n$.

If the ρ-redex $g|_n$ contracted in $g \to_{\rho,n} h$ is a circular redex, then $g = h$ according to Proposition 8.5. However, by Lemma 8.6, each $\mathcal{U}(\rho)$-redex at positions in $\mathcal{P}_g(n)$ in $\mathcal{U}(g)$ reduces to itself as well. Hence, we get a reduction $\mathcal{U}(g) \xrightarrow{p}\!\!\!\!\!\rightarrow_{\mathcal{U}(\rho)} \mathcal{U}(h)$ via a complete development of the redexes at the minimal positions $\mathcal{P}^m_g(n)$. The equality $\mathcal{U}(c) = \bigsqcap_{\iota<\alpha} c_\iota$ then follows as for the first case above. ☐

Before we prove the soundness of strongly p-converging term graph reductions, we show the following technical lemma:

Lemma 8.20. *Let $(a_\iota)_{\iota<\alpha}$ be a sequence in a complete semilattice $(A, \leq)$ and $(\gamma_\iota)_{\iota<\delta}$ a strictly monotone sequence in the ordinal α such that $\bigsqcup_{\iota<\delta} \gamma_\iota = \alpha$. Then*

$$\liminf_{\iota\to\alpha} a_\iota = \liminf_{\beta\to\delta} \left(\bigsqcap_{\gamma_\beta\leq\iota<\gamma_{\beta+1}} a_\iota \right).$$

Proof. At first we show that

$$\bigsqcap_{\beta\leq\beta'<\delta} \; \bigsqcap_{\gamma_{\beta'}\leq\iota<\gamma_{\beta'+1}} a_\iota = \bigsqcap_{\gamma_\beta\leq\iota<\alpha} a_\iota \qquad \text{for all } \beta < \delta \tag{$*$}$$

by using the antisymmetry of the partial order $\leq$ on A.

Since for all $\beta \leq \beta' < \delta$, we have that $\bigsqcap_{\gamma_{\beta'}\leq\iota<\gamma_{\beta'+1}} a_\iota \geq \bigsqcap_{\gamma_\beta\leq\iota<\alpha} a_\iota$, we obtain that $\bigsqcap_{\beta\leq\beta'<\delta} \bigsqcap_{\gamma_{\beta'}\leq\iota<\gamma_{\beta'+1}} a_\iota \geq \bigsqcap_{\gamma_\beta\leq\iota<\alpha} a_\iota$

On the other hand, since $(\gamma_\iota)_{\iota<\delta}$ is strictly monotone and $\bigsqcup_{\iota<\delta} \gamma_\iota = \alpha$, we find for each $\gamma_\beta \leq \gamma < \alpha$ some $\beta \leq \beta' < \delta$ such that $\gamma_{\beta'} \leq \gamma < \gamma_{\beta'+1}$ and, thus, $\bigsqcap_{\gamma_{\beta'}\leq\iota<\gamma_{\beta'+1}} a_\iota \leq a_\gamma$. Consequently, we obtain that $\bigsqcap_{\beta\leq\beta'<\delta} \bigsqcap_{\gamma_{\beta'}\leq\iota<\gamma_{\beta'+1}} a_\iota \leq \bigsqcap_{\gamma_\beta\leq\iota<\alpha} a_\iota$.

With the thus obtained equation $(*)$, it remains to be shown that $\bigsqcup_{\beta<\alpha} \bigsqcap_{\beta\leq\iota<\alpha} a_\iota = \bigsqcup_{\beta'<\delta} \bigsqcap_{\gamma'_\beta\leq\iota<\alpha} a_\iota$. Again, we use the antisymmetry of $\leq$.

Since $\bigsqcap_{\iota<\delta} \gamma_\iota = \alpha$, we find for each $\beta < \alpha$ some $\beta' < \delta$ with $\gamma_{\beta'} \geq \beta$. Consequently, we have for each $\beta < \alpha$ some $\beta' < \delta$ with $\bigsqcap_{\beta\leq\iota<\alpha} a_\iota \leq \bigsqcap_{\gamma'_\beta\leq\iota<\alpha} a_\iota$. Hence $\bigsqcup_{\beta<\alpha} \bigsqcap_{\beta\leq\iota<\alpha} a_\iota \leq \bigsqcup_{\beta'<\delta} \bigsqcap_{\gamma'_\beta\leq\iota<\alpha} a_\iota$.

On the other hand, since for each $\beta' < \delta$ there is a $\beta < \alpha$ (namely $\beta = \gamma_\beta$) with $\bigsqcap_{\beta\leq\iota<\alpha} a_\iota = \bigsqcap_{\gamma'_\beta\leq\iota<\alpha} a_\iota$, we also have $\bigsqcup_{\beta<\alpha} \bigsqcap_{\beta\leq\iota<\alpha} a_\iota \geq \bigsqcup_{\beta'<\delta} \bigsqcap_{\gamma'_\beta\leq\iota<\alpha} a_\iota$. ☐

Theorem 8.21. *Let $\mathcal{R}$ be a left-linear GRS. If $g \xrightarrow{p}\!\!\!\!\!\rightarrow_{\mathcal{R}} h$, then $\mathcal{U}(g) \xrightarrow{p}\!\!\!\!\!\rightarrow_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$.*

Proof. Let $S = (g_\iota \to_{c_\iota} g_{\iota+1})_{\iota<\alpha}$ be a reduction strongly p-converging to g_α in $\mathcal{R}$, i.e. $S\colon g_0 \xrightarrow{p}\!\!\!\!\!\rightarrow_{\mathcal{R}} g_\alpha$. According to Theorem 8.19, there is for each $\gamma < \alpha$ a strongly p-converging reduction $T_\gamma\colon \mathcal{U}(g_\gamma) \xrightarrow{p}\!\!\!\!\!\rightarrow_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_{\gamma+1})$ such that

$$\bigsqcap_{\iota<|T_\gamma|} c'_\iota = \mathcal{U}(c_\gamma) \text{ for } (c'_\iota)_{\iota<|T_\gamma|} \text{ the reduction contexts in } T_\gamma. \tag{$*$}$$

Define for each $\delta \leq \alpha$ the concatenation $U_\delta = \prod_{\iota<\delta} T_\iota$. We will show that $U_\delta \colon \mathcal{U}(g_0) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_\delta)$ for each $\delta \leq \alpha$ by induction on δ. The theorem is then obtained for the case $\delta = \alpha$.

The case $\delta = 0$ is trivial. If $\delta = \delta' + 1$, then we have by induction hypothesis that $U_{\delta'} \colon \mathcal{U}(g_0) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_{\delta'})$. Since $T_{\delta'} \colon \mathcal{U}(g_{\delta'}) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_\delta)$, and $U_\delta = U_{\delta'} \cdot T_{\delta'}$, we have $U_\delta \colon \mathcal{U}(g_0) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_\delta)$.

For the case that δ is a limit ordinal, let $U_\delta = (t_\iota \to_{c'_\iota} t_{\iota+1})_{\iota<\beta}$. For each $\gamma < \beta$ we find some $\delta' < \delta$ with $U_\delta|_\gamma < U_{\delta'}$. By induction hypothesis, we can assume that $U_{\delta'}$ is strongly p-continuous. According to Proposition 2.4, this means that the proper prefix $U_\delta|_\gamma$ strongly p-converges to t_γ. This shows that each proper prefix $U_\delta|_\gamma$ of U_δ strongly p-converges to t_γ. Hence, by Proposition 2.4, U_δ is strongly p-continuous.

In order to show that $U_\delta \colon \mathcal{U}(g_0) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g_\delta)$, it remains to be shown that $\liminf_{\iota\to\beta} c'_\iota = \mathcal{U}(g_\delta)$. Since S is strongly p-converging, we know that $\liminf_{\iota\to\delta} c_\iota = g_\delta$. By Theorem 8.18, we thus have $\liminf_{\iota\to\delta} \mathcal{U}(c_\iota) = \mathcal{U}(g_\delta)$. By $(*)$ and the construction of U_δ, the sequence of reduction contexts $(c'_\iota)_{\iota<\beta}$ consists of segments whose glb is the unravelling of a corresponding reduction context c_γ. More precisely, there is a strictly monotone sequence $(\gamma_\iota)_{\iota<\delta}$ with $\gamma_0 = 0$ and $\bigsqcup_{\iota<\delta} \gamma_\iota = \beta$ such that $\mathcal{U}(c_\iota) = \bigsqcap_{\gamma_\iota \leq \gamma < \gamma_{\iota+1}} c'_\gamma$ for all $\iota < \delta$. Thus, we can complete the proof as follows:

$$\mathcal{U}(g_\delta) = \liminf_{\iota\to\delta} \mathcal{U}(c_\iota) = \liminf_{\iota\to\delta} \bigsqcap_{\gamma_\iota \leq \gamma < \gamma_{\iota+1}} c'_\gamma \overset{\text{Lem. 8.20}}{=} \liminf_{\iota\to\beta} c'_\iota$$

□

Note that the counterexample from Example 8.14 is not applicable to strong p-convergence. Since in the considered system every term graph resp. every term is a redex, we can reduce every term graph resp. every term to $\bot$ by an infinite strongly p-converging reduction. We therefore conjecture that the weak completeness property does hold for strongly p-convergent term graph reductions. That is, for every $\mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t$ there is some h with $\mathcal{U}(h) = t$ such that $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$.

However, we can use the confluence of strongly-p-converging term reductions in order to obtain a weak form of completeness for normalising reductions.

Theorem 8.22 ([5])**.** *Every orthogonal term rewriting system is confluent w.r.t. strong p-convergence. That is, $t \overset{p}{\twoheadrightarrow} t_1$ and $t \overset{p}{\twoheadrightarrow} t_2$ implies $t_1 \overset{p}{\twoheadrightarrow} t'$ and $t_2 \overset{p}{\twoheadrightarrow} t'$.*

Corollary 8.23. *For every orthogonal GRS $\mathcal{R}$, we have that two normalising reductions $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} g'$ and $t \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t'$ imply that $t' = \mathcal{U}(g')$.*

Proof. According to Theorem 8.21, $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} g'$ implies $\mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g')$. Note that, by Proposition 8.7, $\mathcal{U}(g')$ is a normal form in $\mathcal{U}(\mathcal{R})$. Hence, the reduction $\mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(g')$ together with $t \overset{m}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t'$ implies $t' = \mathcal{U}(g')$ according to Theorem 8.22. □

9 Discussion

The main contribution of this paper is the establishment of an appropriate calculus of infinitary term graph rewriting. We have shown that strong m-convergence as well as its conservative extension in the form of strong p-convergence provide an adequate theoretical underpinning of such a calculus. The simplicity of the underlying metric resp. partial order structure of term graphs contrasts the intricate structures that we have proposed in our previous work [6]. There, we have been focused exclusively on weak convergence and the peculiar properties of weak convergence made it necessary to carefully define the underlying structures to be quite rigid. As a consequence, a number of intuitively converging term graph reductions do not converge in that calculus. An example is the reduction illustrated in Figure 6c, which in the rigid calculus does not m-converge at all and p-converges only to the partial term graph $\bot^{[n]}\,(n\,(n\,(\dots)))$. In the calculus that we have presented in this paper, this term graph reduction strongly m- and thus p-converges to the term graph $f^{[n]}\,(n\,(n\,(\dots)))$ depicted in Figure 6c.

The new approach that we have presented in this paper – built upon simple generalisations of the metric resp. the partial order on terms to term graphs – is less rigid and captures an intuitive notion of convergence in the form of strong convergence. We have argued for its appropriateness by independently developing two modes of convergence – m- and p-convergence – and showing that both yield the same limits when restricted to total term graphs. Moreover, we have shown the adequacy of our infinitary calculus by establishing its soundness w.r.t. infinitary term rewriting.

We have also made the first steps towards a completeness result by showing that normalising reductions of term graph rewriting systems and their corresponding term rewriting systems are equivalent modulo unravelling. We conjecture that this can be extended to a full completeness property of normalising reductions.

References

[1] Z.M. Ariola and S. Blom. Skew confluence and the lambda calculus with letrec. *Annals of Pure and Applied Logic*, 117(1-3):95–168, 2002.

[2] Z.M. Ariola and J.W. Klop. Lambda calculus with explicit recursion,. *Information and Computation*, 139(2):154 – 233, 1997.

[3] A. Arnold and M. Nivat. The metric space of infinite trees. Algebraic and topological properties. *Fundamenta Informaticae*, 3(4):445–476, 1980.

[4] P. Bahr. Abstract models of transfinite reductions. In C. Lynch, editor, *RTA 2010*, volume 6 of *LIPIcs*, pages 49–66. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2010.

[5] P. Bahr. Partial order infinitary term rewriting and böhm trees. In C. Lynch, editor, *RTA 2010*, volume 6 of *LIPIcs*, pages 67–84. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2010.

[6] Patrick Bahr. Modes of Convergence for Term Graph Rewriting. In Manfred Schmidt-Schauß, editor, *22nd International Conference on Rewriting Techniques and Applications (RTA'11)*, volume 10 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 139–154, Dagstuhl, Germany, 2011. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[7] H.P. Barendregt, M.C.J.D. van Eekelen, J.R.W. Glauert, R. Kennaway, M.J. Plasmeijer, and M.R. Sleep. Term graph rewriting. In Philip C. Treleaven Jaco de Bakker, A. J. Nijman, editor, *PARLE 1987*, volume 259 of *LNCS*, pages 141–158. Springer, 1987.

[8] S. Blom. An approximation based approach to infinitary lambda calculi. In Vincent van Oostrom, editor, *RTA 2004*, volume 3091 of *LNCS*, pages 221–232. Springer, 2004.

[9] B. Courcelle. Fundamental properties of infinite trees. *Theoretical Computer Science*, 25(2):95–169, 1983.

[10] J.A. Goguen, J.W. Thatcher, E.G. Wagner, and J.B. Wright. Initial algebra semantics and continuous algebras. *Journal of the ACM*, 24(1):68–95, 1977.

[11] J.L. Kelley. *General Topology*, volume 27 of *Graduate Texts in Mathematics*. Springer-Verlag, 1955.

[12] R. Kennaway. On transfinite abstract reduction systems. Technical report, CWI (Centre for Mathematics and Computer Science), Amsterdam, 1992.

[13] R. Kennaway. Infinitary rewriting and cyclic graphs. *Electronic Notes in Theoretical Computer Science*, 2:153–166, 1995. SEGRAGRA '95.

[14] R. Kennaway and F.-J. de Vries. Infinitary rewriting. In Terese [19], chapter 12, pages 668–711.

[15] R. Kennaway, J.W. Klop, M.R. Sleep, and F.-J. de Vries. On the adequacy of graph rewriting for simulating term rewriting. *ACM Transactions on Programming Languages and Systems*, 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. *Information and Computation*, 119(1):18–38, 1995.

[17] Simon Peyton-Jones. *The Implementation of Functional Programming Languages.* Prentice Hall, 1987.

[18] D. Plump. Term graph rewriting. In Hartmut Ehrig, Gregor Engels, Hans-Jörg Kreowski, and Grzegorz Rozenberg, editors, *Handbook of Graph Grammars and Computing by Graph Transformation*, volume 2, pages 3–61. World Scientific Publishing Co., Inc., 1999.

[19] Terese. *Term Rewriting Systems.* Cambridge University Press, 1st edition, 2003.

[20] D.A. Turner. A new implementation technique for applicative languages. *Software: Practice and Experience*, 9(1):31–49, 1979.