Böhm Reduction in Infinitary Term Graph Rewriting Systems

Patrick Bahr

IT University of Copenhagen, Denmark
paba@itu.dk

Abstract

The confluence properties of lambda calculus and orthogonal term rewriting do not generalise to the corresponding infinitary calculi. In order to recover the confluence property in a meaningful way, Kennaway et al. [11, 10] introduced Böhm reduction, which extends the ordinary reduction relation so that 'meaningless terms' can be contracted to a fresh constant ⊥. In previous work, we have established that Böhm reduction can be instead characterised by a different mode of convergences of transfinite reductions that is based on a partial order structure instead of a metric space.

In this paper, we develop a corresponding theory of Böhm reduction for term graphs. Our main result is that partial order convergence in a term graph rewriting system can be truthfully and faithfully simulated by metric convergence in the Böhm extension of the system. To prove this result we generalise the notion of residuals and projections to the setting of infinitary term graph rewriting. As ancillary results we prove the infinitary strip lemma and the compression property, both for partial order and metric convergence.

1998 ACM Subject Classification F.4.2, F.1.1

Keywords and phrases infinitary rewriting, term graphs, Böhm trees

Digital Object Identifier 10.4230/LIPIcs.FSCD.2017.8

1 Introduction

A meaningless term (originally called 'undefined element' by Barendregt [6]) is a term that in some sense does not provide any information because it cannot be suitably distinguished from other meaningless terms. In their seminal work on infinitary lambda calculus, Kennaway et al. [11] recognised that infinitary confluence of the infinitary lambda calculus can be established if all meaningless terms are equated. This idea can be elegantly expressed by extending the reduction relation with rules of the form $t \to \bot$ for all meaningless terms of a certain kind. The resulting reduction was coined Böhm reduction. Later, Kennaway et al. [10] applied this idea to first order term rewriting as well.

In previous work [1, 2], we have introduced an alternative approach to deal with meaningless terms that leaves the original reduction relation intact (i.e. no additional rules of the form $t \to \bot$ are needed) but instead changes the underlying model of convergence. Infinitary rewriting, both in the lambda calculus and first-order term rewriting, originally has been based on a metric space to model convergence of transfinite reductions. We showed that if we change the underlying structure from a metric space to a partial order, the resulting infinitary term rewriting system is – under mild assumptions – equivalent to the metric-based system extended to Böhm reduction [1].

In this paper we seek out to develop a corresponding theory of Böhm reduction for term graph rewriting systems in the sense of Barendregt et al. [7] and use it to compare metric- and partial order-based notions of convergence in a similar way. To this end, we use the theory of infinitary term graph rewriting that was shown to be sound and complete w.r.t. infinitary term rewriting [3]. In this paper, we investigate this approach to infinitary term graph rewriting further and study confluence and convergence properties as well as the relation between the partial order-based and the metric-based mode of convergence.

The main result of this paper is that either mode of convergence can be simulated by the other one if we add rules of the form $g \to \bot$, where g is a root-active term graph. This construction is analogous to the *Böhm reduction* outlined above, and our findings mirror a corresponding result in infinitary term rewriting [2].

As our main proof method we develop a theory of residuals and projections. Some of the theory becomes considerably simpler – compared to infinitary term rewriting – simply because redexes cannot be duplicated by term graph rewriting but are 'shared' instead. On the other hand, many proofs become more tedious as the application of rewrite rules is more complicated than in term rewriting. The theory is put to use in proving the infinitary strip lemma and the compression property, which form the cornerstones in the proof of the main result.

The remainder of this paper is structured as follows: In Section 2, we introduce basic notions of term graphs. Then, in Section 3, we present our infinitary term graph rewriting calculi including their fundamental properties. In Section 4, we develop the theory of residuals and projections, which we then apply to prove the infinitary strip lemma as well the compression lemma; the full account of this development is given in Appendix A. Finally, in Section 5, we use these results in order to prove the equivalence of partial order convergence and metric convergence modulo Böhm extensions as described above. Many proofs in this paper are abridged or have been omitted due to lack of space. The remaining proofs can be found in the extended version of this paper [5].

2 Term Graphs and Modes of Convergence

In this section, we briefly present our notion of term graphs (based on Barendregt et al.[7]) together with the metric and the partial order that are used to formalise infinitary term graph rewriting. For a more thorough exposition, the reader is referred to previous work [3, 4].

Sequences. A *sequence* over a set A of length α is a mapping from an ordinal α into A and is written as $(a_\iota)_{\iota<\alpha}$, which indicates the mapping $\iota \mapsto a_\iota$; the notation $|(a_\iota)_{\iota<\alpha}|$ denotes the length α of the sequence. A sequence is called *open* if its length is a limit ordinal; otherwise it is called *closed*. If $(a_\iota)_{\iota<\alpha}$ is finite it is also written as $\langle a_0, \dots, a_{\alpha-1}\rangle$; in particular, $\langle\rangle$ denotes the empty sequence. A^* denotes the set of finite sequences over A. We write $S \cdot T$ for the *concatenation* of two sequences S and T; S is called a (*proper*) *prefix* of T, denoted $S \leq T$ (resp. $S < T$) if there is a (non-empty) sequence S' such that $S \cdot S' = T$. The uniquely determined prefix of a sequence S of length $\beta < |S|$ is denoted by $S|_\beta$.

Graphs and Term Graphs. A signature Σ is a finite set of symbols together with an associated arity function $\mathsf{ar}(\cdot)$. A *graph* over Σ is a triple $g = (N, \mathsf{lab}, \mathsf{suc})$ 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. The graph g is called *finite* whenever the underlying set N of nodes is finite. If $\mathsf{suc}(n) = \langle n_0, \dots, n_{k-1}\rangle$, then we write $\mathsf{suc}_i(n)$ for n_i. The successor function suc defines, for each node n, directed edges from n to $\mathsf{suc}_i(n)$. A path from a node m to a node

n in a graph is a finite sequence $\langle e_0, \ldots, e_l \rangle$ of numbers such that $n = \mathsf{suc}_{e_l}(\ldots \mathsf{suc}_{e_0}(m))$, i.e. n is reached from m by first taking the e_0-th edge, then the e_1-th edge etc.

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 term graph g is called *finite* if the underlying graph is finite. The class of all term graphs over Σ is denoted $\mathcal{G}^\infty(\Sigma)$; the class of all finite term graphs over Σ is denoted $\mathcal{G}(\Sigma)$. 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 n.

Paths, Positions, Term Trees. Let $g \in \mathcal{G}^\infty(\Sigma)$ and $n \in N^g$. 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 $\mathcal{P}_g(n)$. 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. $\mathsf{depth}_g(n) = \min \{ |\pi| \mid \pi \in \mathcal{P}_g(n) \}$. 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 $\mathsf{lab}^g(\mathsf{node}_g(\pi))$, i.e. the labelling of g at π. The term graph g is called a *term tree* if each node in g has exactly one position.

Homomorphisms. The notion of homomorphisms is central for dealing with term graphs. For greater flexibility, we will parametrise this notion by a set of constant symbols Δ for which the homomorphism condition is suspended. This will allow us to deal with variables and partiality appropriately. Given $g, h \in \mathcal{G}^\infty(\Sigma)$, a Δ-*homomorphism* ϕ from g to h, denoted $\phi\colon g \to_\Delta h$, is a function $\phi\colon N^g \to N^h$ with $\phi(r^g) = r^h$ that satisfies the following equations for all for all $n \in N^g$ with $\mathsf{lab}^g(n) \notin \Delta$:

$$\mathsf{lab}^g(n) = \mathsf{lab}^h(\phi(n)) \qquad \text{(labelling)}$$

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

For $\Delta = \emptyset$, we get the usual notion of homomorphisms on term graphs and from that the notion of isomorphisms. Nodes labelled with symbols in Δ can be thought of as holes in the term graphs (e.g. variables or $\bot$). Homomorphisms also give us a way to describe differences in sharing: given two term graphs g and h, we say that g has *less sharing than* h, written $g \le^{\mathsf{S}} h$, if there is a homomorphism $\phi\colon g \to h$.

Canonical Form, Unravelling, Bisimilarity. We do not want to distinguish between isomorphic term graphs. Therefore, we use a well-known trick [13] to obtain canonical representatives of isomorphism classes: a term graph g is called *canonical* if $n = \mathcal{P}_g(n)$ holds for each $n \in N^g$. The set of all (finite) canonical term graphs over Σ is denoted $\mathcal{G}^\infty_\mathcal{C}(\Sigma)$ (resp. $\mathcal{G}_\mathcal{C}(\Sigma)$). For each term graph $h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$, its *canonical representative* $\mathcal{C}(h)$ is obtained from h by replacing each node n in h by $\mathcal{P}_h(n)$.

We consider the set of terms $\mathcal{T}^\infty(\Sigma)$ as 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 $\mathcal{U}(g)$ such that there is a homomorphism $\phi\colon \mathcal{U}(g) \to g$. Two term graphs g, h are called *bisimilar*, denoted $g \simeq h$, if $\mathcal{U}(g) = \mathcal{U}(h)$.

Labelled Quotient Tree. We shall use an alternative representation that describes term graphs uniquely up to isomorphism. To this end, we define the binary relation $\sim_g$ on positions in the term graph g as follows: $\pi_1 \sim_g \pi_2$ iff $\mathsf{node}_g(\pi_1) = \mathsf{node}_g(\pi_2)$. That is, positions are related if they lead to the same node. The triple $(\mathcal{P}(g), g(\cdot), \sim_g)$, called *labelled quotient tree*, consisting of the mapping $g(\cdot)\colon \mathcal{P}(g) \to \Sigma$ and the binary relation $\sim_g$ over $\mathcal{P}(g)$ as defined above characterises the term graph g up to isomorphism. In particular, each canonical term graph is uniquely determined by exactly one labelled quotient tree. The name is derived from the fact that $(\mathcal{P}(g), g(\cdot))$ describes a labelled tree and $\sim_g$ is a congruence on the set of nodes in this tree.

Metric Space. Next we present the partial order and the metric on term graphs that give us the modes of convergence needed for infinitary rewriting. The metric $\mathbf{d}_\dagger$ on term graphs is defined analogously to the metric that is used in infinitary term rewriting [8]. We define $\mathbf{d}_\dagger(g,h) = 0$ if $g = h$ and otherwise $\mathbf{d}_\dagger(g,h) = 2^{-d}$, where d is the minimal depth at which g and h differ. More precisely, d is defined as the largest number e such that g and h become isomorphic if all nodes at depth e are relabelled with a fresh symbol $\bot$ and their outgoing edges are removed (along with all nodes that become unreachable). We can give a concise characterisation of limits in the resulting metric space using labelled quotient trees:

▶ **Theorem 2.1** ([3, 4]). *$(\mathcal{G}^\infty_\mathcal{C}(\Sigma), \mathbf{d}_\dagger)$ is a complete ultrametric space, and the limit of each Cauchy sequence $(g_\iota)_{\iota<\alpha}$ is given by the labelled quotient tree $(P, l, \sim)$ with*

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

$$l(\pi) = g_\beta(\pi) \quad \text{if } \exists\beta < \alpha \forall\beta \leq \iota < \alpha\colon g_\iota(\pi) = g_\beta(\pi) \qquad \text{for all } \pi \in P$$

Intuitively, the limit of a Cauchy sequence $(g_\iota)_{\iota<\alpha}$ is the tree consisting of all nodes that become eventually stable in $(\mathcal{U}(g_\iota))_{\iota<\alpha}$ (i.e. remain unchanged in the unravelling from some point onwards), but quotiented to a graph by sharing all nodes that eventually remain shared in $(\mathcal{U}(g_\iota))_{\iota<\alpha}$. An example is depicted in Figure 1c.

Partial Order. To define a partial order on term graphs, we consider signatures of the form $\Sigma_\bot$ that extend a signature Σ with a fresh constant symbol $\bot$. We call term graphs over $\Sigma_\bot$ *partial*, and term graphs over Σ *total*. We then use Δ-homomorphisms with $\Delta = \{\bot\}$ – also called *$\bot$-homomorphisms* – to define the *simple partial order* $\leq^\mathsf{S}_\bot$ on $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$ as follows: $g \leq^\mathsf{S}_\bot h$ iff there is a $\bot$-homomorphism $\phi\colon s \to_\bot t$. Using labelled quotient trees, we get the following alternative characterisation:

▶ **Corollary 2.2** (characterisation of $\leq^\mathsf{S}_\bot$, [3, 4]). *Let $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$. Then $g \leq^\mathsf{S}_\bot h$ iff, for all $\pi, \pi' \in \mathcal{P}(g)$, we have*

$$(a)\ \pi \sim_g \pi' \implies \pi \sim_h \pi' \qquad \text{and} \qquad (b)\ g(\pi) = h(\pi) \quad \text{if } g(\pi) \in \Sigma.$$

The partially ordered set $(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot), \leq^\mathsf{S}_\bot)$ forms a *complete semilattice*, i.e. it has a least element $\bot$, each directed set D in $(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot), \leq^\mathsf{S}_\bot)$ has a *least upper bound* (*lub*) $\bigsqcup D$, and every *non-empty* set B in $(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot), \leq^\mathsf{S}_\bot)$ has *greatest lower bound* (*glb*) $\bigsqcap B$. In particular, this means that for any sequence $(g_\iota)_{\iota<\alpha}$ in $(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot), \leq^\mathsf{S}_\bot)$, its *limit inferior*, defined by $\liminf_{\iota\to\alpha} g_\iota = \bigsqcup_{\beta<\alpha} \left(\bigsqcap_{\beta\leq\iota<\alpha} g_\iota\right)$, exists.

▶ **Theorem 2.3** ([3, 4]). *$(\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot), \leq^\mathsf{S}_\bot)$ is a complete semilattice. In particular, the limit* *inferior of a sequence* $(g_\iota)_{\iota<\alpha}$ *is given by the labelled quotient tree* $(P, \sim, l)$ *with*

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

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

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

The limit inferior generalises the limit of Cauchy sequences to arbitrary sequences. Similarly to the limit, the limit inferior of a sequence $(g_\iota)_{\iota<\alpha}$ is also the tree consisting of all nodes that become eventually stable in $(\mathcal{U}(g_\iota))_{\iota<\alpha}$, quotiented to a graph by sharing nodes that become eventually shared in $(g_\iota)_{\iota<\alpha}$. But since $(g_\iota)_{\iota<\alpha}$ is not necessarily Cauchy, some nodes never become stable and are thus replaced by $\bot$-nodes in the limit inferior construction. For an example, consider the sequence $(g_\iota)_{\iota<\alpha}$ from Figure 1c, but where the label f is replaced by h in g_0, g_2, g_4 etc. The resulting sequence is not Cauchy anymore; its limit inferior can be obtained from g_ω in Figure 1c, by replacing the label f with $\bot$.

Another example of a complete semilattice is the prefix order $\leq$ on sequences, which allows us to generalise concatenation as follows: Let $(S_\iota)_{\iota<\alpha}$ be a sequence of sequences over some set A. The concatenation of $(S_\iota)_{\iota<\alpha}$, written $\prod_{\iota<\alpha} S_\iota$, is recursively defined as the empty sequence $\langle\rangle$ if $\alpha = 0$, $\left(\prod_{\iota<\beta} S_\iota\right) \cdot S_\beta$ if $\alpha = \beta + 1$, and $\bigsqcup_{\gamma<\alpha} \prod_{\iota<\gamma} S_\iota$ if α is a limit ordinal.

3 Term Graph Rewriting

In this paper, we adopt the term graph rewriting framework of Barendregt et al. [7]. To represent placeholders in rewrite rules, we use variables – in a manner that is very similar to term rewrite rules. To this end, we consider a signature $\Sigma_\mathcal{V} = \Sigma \uplus \mathcal{V}$ that extends the signature Σ with a countably infinite set $\mathcal{V}$ of nullary variable symbols.

▶ **Definition 3.1** (term graph rewriting systems). 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 are reachable from l or r. We write ρ_l and ρ_r to denote the left- and right-hand side of ρ, respectively, i.e. the term graph $g|_l$ and $g|_r$, respectively. Additionally, we require that for each variable $v \in \mathcal{V}$ there is at most one node n in g labelled v and that n is different from l but still reachable from l. ρ is called *left-linear* (resp. *left-finite*) if ρ_l is a term tree (resp. is finite). A *term graph rewriting system (GRS)* $\mathcal{R}$ is a pair (Σ, R) consisting of a signature Σ and a set R of term graph rules over Σ. $\mathcal{R}$ is called left-linear (resp. left-finite) if each rule of $\mathcal{R}$ is left-linear (resp. left-finite).

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 rewrite rules that every variable occurring on the right-hand side must also occur on the left-hand side.

▶ **Example 3.2.** Figure 1a shows two term graph rules which both unravel to the term rule $\rho \colon Y\,x \to x\,(Y\,x)$. Note that sharing of nodes is used both to refer from the right-hand side to variables on the left-hand side, and in order to simulate duplication.

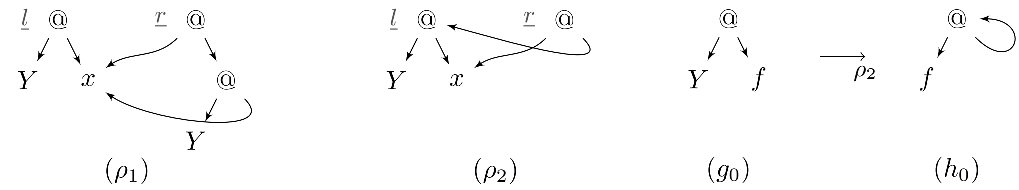

(a) Term graph rules that unravel to $Y\,x \to x\,(Y\,x)$. **(b)** A single ρ_2-step.

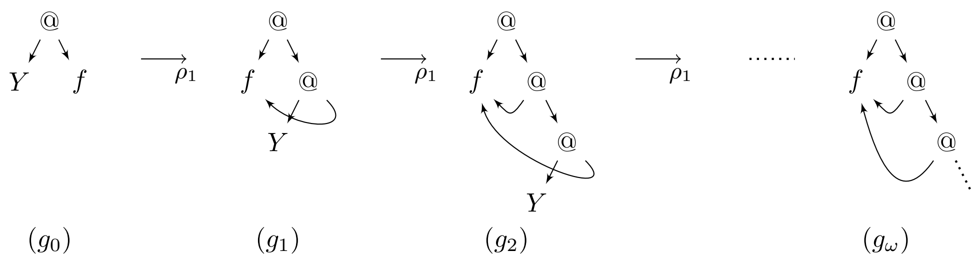

(c) A strongly m-convergent term graph reduction over ρ_1.

Figure 1 Implementation of the fixed point combinator as a term graph rewrite rule.

The notion of unravelling term graphs to terms straightforwardly extends to term graph rules: The *unravelling* of a term graph rule ρ, denoted $\mathcal{U}(\rho)$, is the term rule $\mathcal{U}(\rho_l) \to \mathcal{U}(\rho_r)$. The unravelling of a GRS $\mathcal{R} = (\Sigma, R)$, denoted $\mathcal{U}(\mathcal{R})$, is the TRS $(\Sigma, \{\mathcal{U}(\rho) \mid \rho \in R\})$.

Without going into all details of the construction, we describe the application of a rewrite rule ρ with root nodes l and r to a term graph g in four steps: at first a suitable sub-term graph $g|_n$ of g rooted in some node n of g is *matched* against the left-hand side of ρ. This matching amounts to finding a $\mathcal{V}$-homomorphism ϕ from the left-hand side ρ_l to $g|_n$. The term graph $g|_n$ is called a *redex*, and the pair (n, ρ) is called a *redex occurrence* in g; abusing notation we write (π, ρ) for the redex occurrence $(\mathsf{node}_g(\pi), \rho)$. The $\mathcal{V}$-homomorphism ϕ instantiates variables in the rule with sub-term graphs of the redex. In the second step, nodes and edges in ρ that are not in ρ_l are copied into g, such that each edge pointing to a node m in ρ_l is redirected to $\phi(m)$. In the next step, all edges pointing to the root n of the redex are redirected to the root n' of the *contractum*, which is either r or $\phi(r)$, depending on whether r has been copied into g or not (because it is reachable from l in ρ). Finally, all nodes not reachable from the root of (the now modified version of) g are removed. With h the result of the this construction, we obtain a *pre-reduction step* $\psi\colon\, g \mapsto_{n,\rho,n'} h$ from g to h.

Figure 1b and 1c illustrate how the two rules in Figure 1a are applied to a term graph.

In order to define convergence on infinite reductions, we require that all term graphs are in canonical form. Therefore, we define a reduction step as a pre-reduction step as described above, where both term graphs have been turned into their canonical form:

▶ **Definition 3.3** (reduction steps). Let $\mathcal{R} = (\Sigma, R)$ be GRS, $\rho \in R$ and $g, h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ with $n \in N^g$ and $m \in N^h$. A tuple $\phi = (g, n, \rho, m, h)$ is called a *reduction step*, written $\phi\colon\, g \to_{n,\rho,m} h$, if there is a pre-reduction step $\phi'\colon\, g' \mapsto_{n',\rho,m'} h'$ with $\mathcal{C}(g') = g$, $\mathcal{C}(h') = h$, $n = \mathcal{P}_{g'}(n')$, and $m = \mathcal{P}_{h'}(m')$. We use the shorthand notation $\phi\colon\, g \to_{n,\rho} h$ and $\phi\colon\, g \to_n h$ if $\phi\colon\, g \to_{n,\rho,m} h$ for some m (and ρ). We write $\phi\colon\, g \to_\mathcal{R} h$ to indicate $\mathcal{R}$.

In this paper, we focus on the strong variant of convergence [3]. This variant of convergence takes into account the position of contracted redexes. For metric convergence, only the depth of the contracted redex is needed; for the partial order variant, we need an appropriate notion of reduction contexts, which is provided with the help of local truncations:

▶ **Definition 3.4** (local truncation). Let $g \in \mathcal{G}^\infty(\Sigma_\bot)$ and $M \subseteq N^g$. The *local truncation* of g at M, denoted $g\backslash M$, is obtained from g by labelling all nodes in M with $\bot$ and removing all outgoing edges from nodes in M (also removing all nodes that thus become unreachable from the root). Instead of $g\backslash\{n\}$ and $g\backslash\{\mathsf{node}_g(\pi)\}$, we also write $g\backslash n$ and $g\backslash\pi$, respectively.

Most of the time we will use the characterisation of local truncations in terms of labelled quotient trees instead of the definition above:

▶ **Lemma 3.5** ([3]). *For each $g \in \mathcal{G}^\infty(\Sigma_\bot)$ and $M \subseteq N^g$, the local truncation $g\backslash M$ has the following labelled quotient tree $(P, l, \sim)$:*

$$P = \{\pi \in \mathcal{P}(g) \mid \forall \pi' < \pi \colon \mathsf{node}_g(\pi') \notin M\} \qquad l(\pi) = \begin{cases} g(\pi) & \text{if } \mathsf{node}_g(\pi) \notin M \\ \bot & \text{if } \mathsf{node}_g(\pi) \in M \end{cases}$$
$$\sim \; = \; \sim_g \cap \, P \times P$$

Now we have everything in place to define our notions of convergence:

▶ **Definition 3.6** ([3]). Let $\mathcal{R} = (\Sigma, R)$ be a GRS.

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

(ii) A *reduction* in $\mathcal{R}$, is a sequence $(\phi_\iota\colon g_\iota \to_\mathcal{R} g_{\iota+1})_{\iota<\alpha}$ of rewrite steps in $\mathcal{R}$.

(iii) Let $S = (\phi_\iota\colon g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$ be a reduction in $\mathcal{R}$. S is *m-continuous* in $\mathcal{R}$ if $\lim_{\iota\to\lambda} g_\iota = g_\lambda$ and $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\lambda}$ tends to infinity for each limit ordinal $\lambda < \alpha$. S *m-converges* to g in $\mathcal{R}$, denoted $S\colon g_0 \overset{m}{\twoheadrightarrow}_\mathcal{R} g$, if it is m-continuous and either S is closed with $g = g_\alpha$ or S is open with $g = \lim_{\iota\to\alpha} g_\iota$ and $(\mathsf{depth}_{g_\iota}(n_\iota))_{\iota<\alpha}$ tending to infinity.

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

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

We have the following correspondence between m- and p-convergence:

▶ **Theorem 3.7** ([3]). *Let $\mathcal{R}$ be a GRS and S a reduction in $\mathcal{R}_\bot$. We then have that*

$$S\colon g \overset{m}{\twoheadrightarrow}_\mathcal{R} h \qquad \textit{iff} \qquad S\colon g \overset{p}{\twoheadrightarrow}_\mathcal{R} h \textit{ in } \mathcal{G}^\infty_\mathcal{C}(\Sigma).$$

Most of our results will be restricted to GRSs with (weakly) non-overlapping rules:

▶ **Definition 3.8** ((weakly) non-overlapping, [7]). Let (n, ρ) and (n', ρ') be redex occurrences in a term graph g, with corresponding matching $\mathcal{V}$-homomorphisms $\phi\colon \rho_l \to g|_n$ and $\phi'\colon \rho'_l \to g|_{n'}$.

(i) (n, ρ) and (n', ρ') are called *disjoint* if $n' \notin \phi(N)$ and $n \notin \phi'(N')$, where N and N' are the non-variable nodes in ρ_l and ρ'_l, respectively.

(ii) (n, ρ) and (n', ρ') are called *weakly disjoint* if they are either (a) disjoint, or (b) $n = n'$ and contracting both redexes results in isomorphic term graphs, i.e. $g \mapsto_{n,\rho} h \cong h' \leftarrow\!\shortmid_{n',\rho'} g$.

A GRS $\mathcal{R}$ is *non-overlapping* resp. *weakly non-overlapping* if for every term graph g in $\mathcal{R}$, every two distinct redex occurrences are disjoint resp. weakly disjoint. A GRS that is non-overlapping and left-linear is called *orthogonal*.

Below we summarise the soundness and completeness property of infinitary term graph rewriting in terms of infinitary term rewriting.

▶ **Theorem 3.9** (soundness & completeness, [3]). *Let $\mathcal{R}$ be a left-finite GRS.*

(i) If $\mathcal{R}$ is left-linear, then $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ implies $\mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$.
(ii) If $\mathcal{R}$ is orthogonal, then $\mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t$ implies $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ and $t \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$.

The complete semilattice structure that underlies the definition of p-convergence ensures that every p-continuous reduction also p-converges – a property that distinguishes it from m-convergence. In other words, any open well-formed reduction can be uniquely completed to a closed well-formed reduction in the partial order model. A consequence of Theorem 3.7 is that a p-convergent reduction that does not m-converges must produce nodes labelled $\bot$. In the following, we analyse this formation of $\bot$-nodes and characterise it in terms of volatile positions, which are positions repeatedly contracted in a reduction:

▶ **Definition 3.10** (volatility). Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\lambda}$ be an open p-converging reduction. A position π is said to be *volatile* in S if, for each $\alpha < \lambda$, there is some $\alpha \leq \beta < \lambda$ such that $\pi \in n_\beta$. If π is volatile and no proper prefix of it is volatile in S, then π is called *outermost-volatile* in S.

Moreover, we need to characterise positions that are affected by rewrite steps:

▶ **Definition 3.11.** Let π be a position and n a node in a term graph g. Then π is said to *pass through* n in g if there is a prefix $\pi' \leq \pi$ with $\pi' \in \mathcal{P}_g(n)$, and π is said to *properly pass through* n in g if there is a proper prefix $\pi' < \pi$ with $\pi' \in \mathcal{P}_g(n)$.

Using Lemma 3.5 and Theorem 2.3 we can give the following characterisation of the formation $\bot$-nodes, where $\mathcal{P}_{\not\bot}(g)$ denotes the positions of nodes in a term graph $g \in \mathcal{G}^\infty(\Sigma_\bot)$ that are not labelled with $\bot$:

▶ **Lemma 3.12** (volatility). *Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\lambda}$ be an open reduction p-converging to g_λ. Then, for every position π, we have the following:*

(i) If π is volatile in S, then $\pi \notin \mathcal{P}_{\not\bot}(g_\lambda)$.
(ii) $g_\lambda(\pi) = \bot$ iff (a) π is outermost-volatile in S, or (b) there is some $\alpha < \lambda$ such that $g_\alpha(\pi) = \bot$ and, for all $\alpha \leq \iota < \lambda$, π does not pass through n_ι in g_ι.

Volatile positions give us the vocabulary to formulate the following variant of Theorem 3.7:

▶ **Corollary 3.13.** *For every GRS $\mathcal{R}$, $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$, and reduction S in $\mathcal{R}_\bot$, we have that $S\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ and no open prefix of S has a volatile position iff $S\colon g \overset{m}{\twoheadrightarrow}_{\mathcal{R}} h$.*

Proof. This follows straightforwardly from Theorem 3.7 using Lemma 3.12 (ii). ◀

4 Residuals and Projections

In this section, we develop the theory of residuals and projections for infinitary term graph rewriting.[1] We then use this machinery to prove the infinitary strip lemma and

[1] This section is heavily abridged; see Appendix A for the full theory and all proofs.

the compression lemma for both p- and m-convergence. We start by recapitulating the basic definitions and properties of residuals and projections for single reduction steps from Barendregt et al. [7].

Given two disjoint redex occurrences (n, ρ) and (n', ρ') in a term graph g, with matching $\mathcal{V}$-homomorphisms ϕ and ϕ', respectively, and a pre-reduction step $g \mapsto_{n,\rho} h$, we know that either n' is not a node in h, or there is a redex occurrence (n', ρ') in h [7]. This finding motivates the definition of residuals and projections:

▶ **Definition 4.1** (reduction step residuals,[7])**.** Let $\psi\colon g \to_{n,\rho} h$ be a reduction step, $\overline{\psi}\colon \overline{g} \mapsto_{\overline{n},\rho} \overline{h}$ the underlying pre-reduction step, and (n', ρ') a redex occurrence in g weakly disjoint from (n, ρ); let $\overline{n'}$ be the node corresponding to n' in $\overline{g}$, i.e. $\overline{n'} = \phi(n')$, where ϕ is the isomorphism from g to $\overline{g}$.

(i) The *residual* of (n', ρ') by ψ, denoted $(n', \rho')//\psi$, is either
  (a) the empty set $\emptyset$ if (n', ρ') and (n, ρ) are not disjoint or $\overline{n'} \notin N^{\overline{h}}$, or
  (b) $\mathcal{P}_{\overline{h}}(\overline{n}')$ if (n', ρ') and (n, ρ) are disjoint and $\overline{n'} \in N^{\overline{h}}$.

(ii) The *projection* of the reduction step $\psi'\colon g \to_{n',\rho'} h'$ by ψ, denoted ψ'/ψ, is either
  (a) the empty reduction if $(n', \rho')//\psi = \emptyset$, or
  (b) the single step reduction contracting the ρ-redex rooted in $(n', \rho')//\psi$ in h otherwise.

Note that the residual $(n', \rho')//\psi$ is either the empty set or a node in h, namely $\mathcal{P}_{\overline{h}}(\overline{n}')$. This property generalises to residuals by reductions of arbitrary length:

▶ **Definition 4.2** (residuals)**.** Let $\mathcal{R}$ be a weakly non-overlapping GRS, $S\colon g_0 \overset{p}{\twoheadrightarrow}_{\mathcal{R}} g_\alpha$, and (n, ρ) a redex occurrence in g_0 with ρ a rule in $\mathcal{R}$. The *residual* of (n, ρ) by S, denoted $(n, \rho)//S$, is inductively defined on the length of S as follows:

- S is empty: $(n, \rho)//S = n$
- $S = T \cdot \langle\psi\rangle$: $(n, \rho)//S = \begin{cases} \emptyset & \text{if } (n, \rho)//T = \emptyset \\ (m, \rho)//\psi & \text{if } (n, \rho)//T = m \neq \emptyset \end{cases}$
- S is open: $(n, \rho)//S = \mathcal{P}_{\not\perp}(g_\alpha) \cap \liminf_{\iota \to \alpha} (n, \rho)//S|_\iota$,
  that is $\quad \pi \in (n, \rho)//S \quad$ iff $\quad \pi \in \mathcal{P}_{\not\perp}(g_\alpha)$ and $\exists \beta < \alpha \forall \beta \leq \iota < \alpha\colon \pi \in (n, \rho)//S|_\iota$.

Note that since m-convergence is just a special case of p-convergence, according to Theorem 3.7, the definition of residuals also applies to m-convergent reductions. For *open* m-convergent reductions, however, we can simplify the characterisation by omitting the explicit requirement that a residual position has to be in $\mathcal{P}_{\not\perp}(g_\alpha)$.

Likewise, we can also generalise the notion of projections (cf. Figure 2). The basis for this generalisation is that given a reduction $S\colon g_0 \overset{p}{\twoheadrightarrow}_{\mathcal{R}} g_\alpha$ in a weakly non-overlapping GRS $\mathcal{R}$, and (n, ρ) a redex occurrence in g_0, we have that if $(n, \rho)//S = m$ is non-empty, then (m, ρ) is a redex occurrence in g_α:

▶ **Definition 4.3** (projections)**.** Let $\mathcal{R}$ be a weakly non-overlapping GRS, $\phi\colon g \to_{n,\rho} h$ a reduction step in $\mathcal{R}$, and $S = (\psi_\iota\colon g_\iota \to g_{\iota+1})_{\iota<\alpha}$ a p-converging reduction in $\mathcal{R}$. The *projection* of ϕ by S, denoted ϕ/S, is (a) the empty reduction if $(n, \rho)//S = \emptyset$, and (b) the single step reduction contracting the ρ-redex rooted in $(n, \rho)//S$ in h otherwise. The *projection* of S by ϕ, denoted S/ϕ, is defined as the concatenation $\prod_{\iota<\alpha} \psi_\iota / (\phi/S|_\iota)$.

One can show that projections commute for both m- and p-convergent reductions given that one of the reductions is finite:

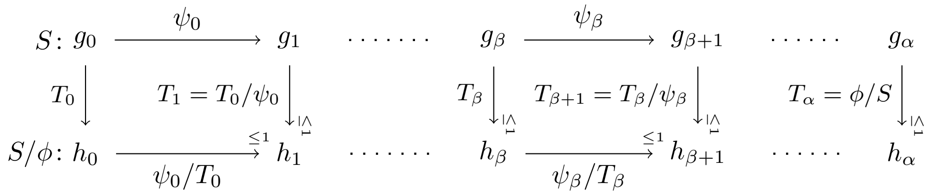

Figure 2 The Infinitary Strip Lemma.

▶ **Theorem 4.4** (infinitary strip lemma: p-convergence)**.** *Let $\mathcal{R}$ be a weakly non-overlapping GRS, $\phi\colon g_0 \to_{n,\rho} h_0$ a reduction step in $\mathcal{R}$, $S\colon g_0 \overset{p}{\twoheadrightarrow}_{\mathcal{R}} g_\alpha$, and $\phi/S\colon g_\alpha \to^{\leq 1}_{\mathcal{R}} h_\alpha$. Then we have that $S/\phi\colon h_0 \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h_\alpha$.*

Note that the strip lemma for term graph rewriting is simpler than for term rewriting as a redex has at most one residual and, thus, we do not have to deal with complete developments. The proof of the strip lemma constructs the commuting diagram shown in Figure 2. For the basic squares we can use the result of Barendregt et al. [7] who showed that projections of single reduction steps commute.

From the above infinitary strip lemma, one can derive the corresponding variant for m-convergence using Corollary 3.13 fairly easily:

▶ **Theorem 4.5** (infinitary strip lemma: m-convergence)**.** *Let $\mathcal{R}$ be a weakly non-overlapping GRS, $\phi\colon g_0 \to_{n,\rho} h_0$ a reduction step in $\mathcal{R}$, $S\colon g_0 \overset{m}{\twoheadrightarrow}_{\mathcal{R}} g_\alpha$, and $\phi/S\colon g_\alpha \to^{\leq 1} h_\alpha$. Then we have that $S/\phi\colon h_0 \overset{m}{\twoheadrightarrow}_{\mathcal{R}} h_\alpha$.*

The definition of projections can be generalised to projections S/T of arbitrary pairs of reductions S, T in the obvious way (by extending Figure 2 vertically). While we conjecture that these general projections of p-convergent reductions commute as well, which means that we have infinitary confluence, the same cannot be said for m-convergence: the counterexample of Kennaway et al. [9] applies here as well.

The infinitary strip lemmas are a powerful tool as we shall see. Below we will apply them to prove that reductions can be compressed to length at most ω – a useful property in its own right. To this end, we need the two lemmas below. The first one states that any redex obtained by an open reduction must already occur in an earlier term graph, which is subsequently unaffected by reduction. The second lemma states that also all positions within the redex itself remain untouched.

▶ **Lemma 4.6.** *Given an open reduction $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\lambda}$ p-converging to g_λ and a redex occurrence (π, ρ) in g_λ with ρ left-finite, there is a position $\pi \in \mathcal{P}(g_\lambda)$ and some $\alpha < \lambda$ such that (π, ρ) is a redex occurrence in g_ι, and π does not pass through n_ι in g_ι for any $\alpha \leq \iota < \lambda$.*

▶ **Lemma 4.7.** *Let $S\colon g \overset{p}{\twoheadrightarrow} h$ be a p-converging reduction in a weakly non-overlapping GRS $\mathcal{R}$ and (n, ρ) a redex occurrence in g. For each $\pi \in \mathcal{P}_g(n)$ such that π does not pass through the root of any redex contracted in S, we have that $\pi \in (n, \rho)//S$.*

Proof. Straightforward induction on the length of S. ◀

The proof of the full compression property for p-convergent reductions is tricky. For now, we only show that infinite, closed reductions can be compressed. This property will turn out to be sufficient for our purposes and later in Section 5, we can use our main result to extend it to full compression much more easily.

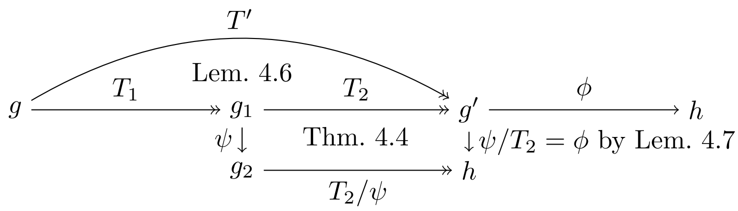

Figure 3 Compression of closed transfinite reductions.

▶ **Proposition 4.8** (compression of closed transfinite reductions). *Let $S\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ in a weakly non-overlapping, left-finite GRS $\mathcal{R}$. Then there is a reduction $T\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ that is finite or open but not longer than S.*

Proof sketch. We proceed by induction on the length of S. The only non-trivial case is where $S = S' \cdot \langle\phi\rangle$ with $S'\colon g \overset{p}{\twoheadrightarrow} g'$ and $\phi\colon g' \to h$. By induction hypothesis there is a finite or open reduction $T'\colon g \overset{p}{\twoheadrightarrow} g'$ of length at most $|S'|$. If T' is finite, then so is $T' \cdot \langle\phi\rangle\colon g \overset{p}{\twoheadrightarrow} h$. Otherwise, let (π, ρ) be the redex occurrence contracted in ϕ and construct the diagram illustrated in Figure 3, where ψ contracts (π, ρ) in g_1. This gives us a reduction $T_3 = T_1 \cdot \langle\psi\rangle \cdot T_2/\psi$ with $T_3\colon g \overset{p}{\twoheadrightarrow} h$ and $|T_3| < |S|$. Thus, we may apply the induction hypothesis to T_3 to obtain a finite or open reduction $T\colon g \overset{p}{\twoheadrightarrow} h$ ◀

The above proof carries over to m-convergent reductions by using Theorem 4.5 instead of Theorem 4.4. Moreover, we can strengthen it to obtain full compression for m-convergent reductions:

▶ **Proposition 4.9.** *Let $S\colon g \overset{m}{\twoheadrightarrow} h$ in a weakly non-overlapping, left-finite GRS. Then there is a reduction $T\colon g \overset{m}{\twoheadrightarrow} h$ of length at most ω.*

Proof. By the proof of Lemma 5.1 in [9] it suffices to show the property for $|S| = \omega + 1$, which can be done analogously to Proposition 4.8 but using Theorem 4.5 instead of Theorem 4.4. ◀

We conclude this section by deriving a compression property for reductions p-converging to $\bot$. To this end, we need the following lemma, which states that any term graph that reduces to $\bot$ must also reduce to a redex:

▶ **Lemma 4.10.** *For each reduction $S\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \bot$ in a weakly non-overlapping, left-finite GRS $\mathcal{R}$ with $g \neq \bot$, we find a finite reduction $g \to^*_{\mathcal{R}} h$ to a redex h.*

Proof sketch. There is at least one step in S contracting a redex at the root, i.e. a proper prefix T of S p-converges to a redex. By induction on the length of T, we show that there is a finite reduction from g to a redex: By Proposition 4.8, we may assume that T is finite or open. If T is finite, we are done. Otherwise, we use Lemma 4.6 to find a proper prefix of T that p-converges to a redex. The induction hypothesis then yields the finite reduction to a redex. ◀

Given this property we can compress any reduction to $\bot$ to a length of at most ω:

▶ **Proposition 4.11.** *For each reduction $S\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \bot$ in a weakly non-overlapping, left-finite GRS $\mathcal{R}$, there is a reduction $T\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \bot$ of length at most ω.*

Proof. Let $g_0 \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \bot$ with $g_0 \neq \bot$. Then we may apply Lemma 4.10 to obtain a reduction $g_0 \to^* h_0 \to g_1$ whose last rewrite step is at the root. By Theorem 4.4, there is also a reduction $g_1 \overset{p}{\twoheadrightarrow} \bot$. Hence, we may repeat this construction to obtain a reduction of the form $g_0 \to^* h_0 \to g_1 \to^* h_1 \to g_2 \to^* \dots$ Either the construction stops at some $i < \omega$ because $g_i = \bot$ in which case we have found a finite reduction $T\colon g_0 \to^* \bot$, or there is no such i with $g_i = \bot$ and we have found a reduction T of length ω with a volatile position $\langle\rangle$. Hence, $T\colon g_0 \overset{p}{\twoheadrightarrow} \bot$ according to Lemma 3.12. ◀

5 Böhm Reduction

Recall Theorem 3.7, which states that p-convergence and m-convergence coincide if we restrict ourselves to total term graphs. In this section, we show that the remaining gap between p- and m-convergence is bridged by adding rewrite rules that contract certain term graphs directly to $\bot$, thereby simulating reductions of the form $g \overset{p}{\twoheadrightarrow} \bot$. We give two characterisations of such term graphs:

▶ **Definition 5.1.** Let $\mathcal{R}$ be a GRS. A partial term graph g in $\mathcal{R}$ is called *fragile* if there is an open reduction $S\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \bot$. A total term graph g in $\mathcal{R}$ is called *root-active* if for each reduction $g \to^*_{\mathcal{R}} h$ there is a reduction $h \to^*_{\mathcal{R}} h'$ to a redex h'. We write $\mathcal{RA}^{\mathcal{R}}$, or simply $\mathcal{RA}$, to denote the set of root-active terms in $\mathcal{R}$.

As it turns out the above two concepts – fragility and root-activeness – coincide on total term graphs. The following observation will help us to establish that:

▶ **Corollary 5.2.** *A total term graph g in a weakly non-overlapping, left-finite GRS $\mathcal{R}$ is fragile in $\mathcal{R}$ iff there is a reduction $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \bot$.*

Proof. The "only if" direction follows by definition, whereas the "if" direction follows from Proposition 4.11 and the fact that total term graphs cannot reduce to $\bot$ in finitely many steps. ◀

▶ **Proposition 5.3.** *Let g be a total term graph in a weakly non-overlapping, left-finite GRS $\mathcal{R}$. Then g is root-active iff g is fragile.*

Proof. If g is root-active, then we can construct a reduction of length ω that infinitely often contracts a redex at the root and thus p-converges to $\bot$. For the converse direction assume some finite reduction $g \to^* h$. If g is fragile, then there is a reduction $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \bot$, according to Corollary 5.2. By iterating Theorem 4.4, we thus find a reduction $h \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \bot$. Moreover, since g is total, so is h. Hence, by Corollary 5.2, h is fragile, too. That means, according to Lemma 4.10 that there is a finite reduction from h to a redex. ◀

To bridge the gap between p- and m-convergence, we adopt the notion of Böhm extensions from term rewriting [10], which is a construction that extends TRSs by rules of the form $t \to \bot$. The definition on GRS is analogous:

▶ **Definition 5.4** (Böhm extension)**.** Let $\mathcal{R} = (\Sigma, R)$ be a GRS, and $\mathcal{U} \subseteq \mathcal{G}^\infty_\mathcal{C}(\Sigma)$.

(i) A $\mathcal{U}$-*instance* of a term graph $h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$ is a term graph $g \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ that is obtained from h by replacing each occurrence of $\bot$ in g with some term graph in $\mathcal{U}$, i.e. there is a set $M \subseteq N^g$ with $g|_m \in \mathcal{U}$ for all $m \in M$, and $h \cong g\backslash M$.

(ii) $\mathcal{U}_\bot$ is the set of term graphs in $\mathcal{G}^\infty_\mathcal{C}(\Sigma_\bot)$ that have a $\mathcal{U}$-instance in $\mathcal{U}$. In other words, $\mathcal{U}_\bot$ contains all those term graphs that can be obtained by taking a term graph g from $\mathcal{U}$ and replacing some sub-term graphs of g that are themselves in $\mathcal{U}$ with $\bot$.

(iii) The *Böhm extension* of $\mathcal{R}$ w.r.t. $\mathcal{U}$ is the GRS $\mathcal{B} = (\Sigma_\perp, R \cup B)$, where

$$B = \left\{ (g \uplus \perp, r^g, r^\perp) \,\middle|\, g \in \mathcal{U}_\perp \setminus \{\perp\} \right\}.$$

That is, B consists of rules with left-hand side $g \in \mathcal{U}_\perp \setminus \{\perp\}$ and right-hand side $\perp$. The rules in B are called *$\perp$-rules w.r.t. $\mathcal{U}$* and we write $g \to_\perp h$ for a reduction step using such a rule in B and call it a *$\perp$-step*.

In the remainder of this section we prove that $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ is equivalent to $g \overset{m}{\twoheadrightarrow}_{\mathcal{B}} h$, where $\mathcal{B}$ is the Böhm extension of $\mathcal{R}$ w.r.t. $\mathcal{RA}^{\mathcal{R}}$.

The semantics of term graph rewriting makes the behaviour of Böhm extensions slightly different compared to term rewriting. Not only term graphs in $\mathcal{U}_\perp$ are contracted to $\perp$ but also term graphs that have more sharing than those in $\mathcal{U}_\perp$:

▶ **Lemma 5.5.** *Let $g \to_{n,\rho,m} h$ be a reduction step of a $\perp$-rule ρ w.r.t. a set of term graphs $\mathcal{U} \subseteq \mathcal{G}^\infty_\mathcal{C}(\Sigma)$. Then there is some $g' \in \mathcal{U}_\perp \setminus \{\perp\}$ with $g' \leq^{\mathsf{S}} g|_n$ and $h = g\backslash n$.*

Proof. The equality $h = g\backslash n$ follows from fact that the right-hand side of ρ is by definition $\perp$. Since the rewrite step takes place at node n in g, we find a matching $\mathcal{V}$-homomorphism $\phi\colon\ \rho_l \to_\mathcal{V} g|_n$. By definition of $\perp$-rules, the left-hand side ρ_l of ρ is some term graph $g' \in \mathcal{U}_\perp \setminus \{\perp\}$. Hence, $\phi\colon\ g' \to_\mathcal{V} g|_n$. Since term graphs in $\mathcal{U}$ do not contain variables, g' does not contain variables either. Therefore, ϕ is a homomorphism. Consequently, $g' \leq^{\mathsf{S}} g|_n$. ◀

In general, this is a problem as root-active term graphs are not closed under increase of sharing. Consider the following example:

▶ **Example 5.6.**

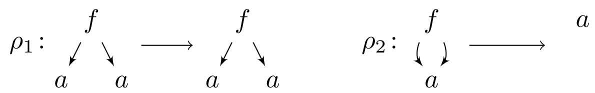

In the GRS consisting of the two rules above, the left-hand side g of ρ_1 is root-active while the left-hand side h of ρ_2 is not, even though $g \leq^{\mathsf{S}} h$. However, if we consider orthogonal systems, this phenomenon cannot occur:

▶ **Lemma 5.7.** *Let $\mathcal{R}$ be an orthogonal, left-finite GRS and g, h two partial term graphs in $\mathcal{R}$ that are bisimilar. Then $g \overset{p}{\twoheadrightarrow} \perp$ iff $h \overset{p}{\twoheadrightarrow} \perp$.*

Proof. As bisimilarity is symmetric we only need to show one direction. Assume that $g \simeq h$ and that $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \perp$. By Theorem 3.9(i), we find a reduction $\mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \perp$, since $\mathcal{U}(\perp) = \perp$. Since $g \simeq h$, we know that $\mathcal{U}(g) = \mathcal{U}(h)$, which means that $\mathcal{U}(h) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \perp$. Since $\perp$ is a normal form in $\mathcal{U}(\mathcal{R})$, we find, according to Theorem 3.9(ii), a reduction $h \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \perp$. ◀

Thus, fragility and, by Proposition 5.3, root-activeness is preserved by bisimilarity. By a similar argument, we have preservation by p-converging reductions as well:

▶ **Lemma 5.8.** *Let $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ and $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \perp$ be a reduction in an orthogonal, left-finite GRS $\mathcal{R}$. Then there is a reduction $h \overset{p}{\twoheadrightarrow} \perp$.*

Proof. By Theorem 3.9 (i), we have $\mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h)$ and $\mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \perp$. Since $\mathcal{R}$ is orthogonal and left-finite, so is $\mathcal{U}(\mathcal{R})$. Because orthogonal, left-finite TRSs are known to be infinitary confluent w.r.t. p-convergence [2], we know that there is a reduction $\mathcal{U}(h) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \perp$. Since $\perp$ is a normal form in $\mathcal{U}(\mathcal{R})$, we may apply Theorem 3.9 (ii) to obtain a term graph reduction $h \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \perp$. ◀

Next, we show that for each $\mathcal{RA}$-instance g of a term graph h, we have $g \overset{p}{\twoheadrightarrow} h$.

▶ **Lemma 5.9.** *If g is a total term graph in a GRS $\mathcal{R}$ that is an $\mathcal{RA}$-instance of a term graph h, then $g \overset{p}{\twoheadrightarrow} h$.*

Proof sketch. We know that $h = g\backslash M$ for some set of nodes M and $g|_m \in \mathcal{RA}$ for all $m \in M$. We then construct a reduction $S\colon g_0 \overset{p}{\twoheadrightarrow}^{\omega} g_1 \overset{p}{\twoheadrightarrow}^{\omega} g_2 \overset{p}{\twoheadrightarrow}^{\omega} \dots g_\omega$ starting in $g_0 = g$ and p-converging to g_ω. Each reduction $g_i \overset{p}{\twoheadrightarrow}^{\omega} g_{i+1}$ in S rewrites a root-active sub-term graph $g|_m$ to $\bot$. If M is finite, $g_\omega = h$ follows easily. Otherwise, one can show that $\liminf_{i \to \omega} g_{i+1} = g_\omega$, which implies $h \leq^{\mathsf{S}}_{\bot} g_\omega$ since $h \leq^{\mathsf{S}}_{\bot} g_i$ for all $i < \omega$. Using Corollary 2.2 we can then show with the help of Theorem 2.3 that $g_\omega \leq^{\mathsf{S}}_{\bot} h$. Hence $S\colon g \overset{p}{\twoheadrightarrow} h$. ◀

According to Proposition 5.3, each term graph $g \in \mathcal{RA}$ is characterised by a reduction $g \overset{p}{\twoheadrightarrow} \bot$. With the above lemma, this property generalises to $\mathcal{RA}_\bot$.

▶ **Proposition 5.10.** *In orthogonal, left-finite GRSs, we have $g \in \mathcal{RA}_\bot$ iff $g \overset{p}{\twoheadrightarrow} \bot$.*

Proof. If $g \in \mathcal{RA}_\bot$, then there is some $h \in \mathcal{RA}$ that is an $\mathcal{RA}$-instance of g. According to Lemma 5.9, we thus find a reduction $h \overset{p}{\twoheadrightarrow} g$. By Proposition 5.3, there is a reduction $h \overset{p}{\twoheadrightarrow} \bot$. Applying Lemma 5.8, we find a reduction $g \overset{p}{\twoheadrightarrow} \bot$.

For the converse direction we show that if $g \overset{p}{\twoheadrightarrow} \bot$ and $h \in \mathcal{G}^\infty_\mathcal{C}(\Sigma)$ is an $\mathcal{RA}$-instance of g, then h is root-active. By Lemma 5.9, we find a reduction $g \overset{p}{\twoheadrightarrow} h$, which means, according to Lemma 5.8, that there is a reduction $h \overset{p}{\twoheadrightarrow} \bot$. By Corollary 5.2, we know that h is fragile, which implies, by Proposition 5.3, that h is root-active. ◀

Finally, we have everything in place to prove our main result:

▶ **Theorem 5.11.** *Let $\mathcal{R}$ be an orthogonal, left-finite GRS and $\mathcal{B}$ its Böhm extension w.r.t. $\mathcal{RA}$. Then we have that $g \overset{p}{\twoheadrightarrow}_\mathcal{R} h$ iff $g \overset{m}{\twoheadrightarrow}_\mathcal{B} h$.*

Proof sketch. $\mathcal{B}$ is a GRS over the signature $\Sigma' = \Sigma \uplus \{\bot\}$, i.e. term graphs containing $\bot$ are considered total in $\mathcal{B}$, which justifies our use of Corollary 3.13 and Theorem 3.7 below.

Given a reduction $S\colon g \overset{m}{\twoheadrightarrow}_\mathcal{B} h$, we know that, by Theorem 3.7, $S\colon g \overset{p}{\twoheadrightarrow}_\mathcal{B} h$, too. We construct a reduction T from S by replacing each $\bot$-step $\widehat{g} \to_{\bot,n} \widehat{h}$ in S by a reduction $S'\colon \widehat{g} \overset{p}{\twoheadrightarrow}_\mathcal{R} \widehat{h}$. For each such $\bot$-step there is, by Lemma 5.5, some $\overline{g} \in \mathcal{RA}_\bot \setminus \{\bot\}$ with $\overline{g} \leq^{\mathsf{S}} \widehat{g}|_n$ and $\widehat{h} = \widehat{g}\backslash n$. Hence, by Proposition 5.10 and Lemma 5.7, we find a reduction $\widehat{g}|_n \overset{p}{\twoheadrightarrow}_\mathcal{R} \bot$. By embedding this reduction in $\widehat{g}$ at node n, we obtain the desired reduction $S'\colon \widehat{g} \overset{p}{\twoheadrightarrow}_\mathcal{R} \widehat{h}$. Using Theorem 2.3, one can show that the thus obtained reduction T p-converges to h.

Given $S\colon g \overset{p}{\twoheadrightarrow}_\mathcal{R} h$, we construct a reduction $T\colon g \overset{p}{\twoheadrightarrow}_\mathcal{B} h$, without any volatile positions. For each open prefix $S|_\lambda$ with an outermost-volatile position π, we find some $\beta < \lambda$ such that no step between β and λ takes place strictly above π. We then remove all reduction steps between β and λ at π or below and replace them with a single $\bot$-step $g_\beta \to_\bot g'_\beta$, which is justified by Proposition 5.10 and Lemma 5.5. Using Lemma 3.12 (ii), one can show that the resulting reduction T p-converges to the same term graph h. By construction, no prefix of T contains a volatile position. Thus, we may apply Corollary 3.13 to conclude $T\colon g \overset{m}{\twoheadrightarrow}_\mathcal{B} h$. ◀

Using the above correspondence, we can leverage the compression property for m-converging reductions to obtain full compression for p-converging reductions:

▶ **Proposition 5.12.** *For every reduction $S\colon g \overset{p}{\twoheadrightarrow}_\mathcal{R} h$ in an orthogonal, left-finite GRS $\mathcal{R}$, there is a reduction $T\colon g \overset{p}{\twoheadrightarrow}_\mathcal{R} h$ of length at most ω.*

Proof sketch. According to Theorem 5.11, $g \overset{p}{\twoheadrightarrow}_\mathcal{R} h$ implies $g \overset{m}{\twoheadrightarrow}_\mathcal{B} h$. One can show that the latter reduction can be reordered to the form $g \overset{m}{\twoheadrightarrow}_\mathcal{R} g' \overset{m}{\twoheadrightarrow}_\bot h$ that performs the $\bot$-steps at the very end (cf. Lemma 27 from Kennaway et al. [10]). By Proposition 4.9 there is a reduction $S\colon g \overset{m}{\twoheadrightarrow}_{\mathcal{R}} g'$ of length at most ω. Moreover, we can show that there is a reduction $T\colon g' \overset{m}{\twoheadrightarrow}_{\bot} h$ of length at most ω (cf. Lemma 7.2.4 from Ketema [12]). As in the proof of Theorem 5.11, we can replace each application of a $\bot$-rule $r \to \bot$ in T with a reduction derived from a corresponding reduction $r \overset{p}{\twoheadrightarrow}_{\mathcal{R}} \bot$, which according to Proposition 4.11 has length at most ω. The thus obtained reduction $T'\colon g' \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ has length at most $\omega \cdot \omega$. If S is finite, then we interleave the reduction steps in T' to obtain a reduction $T''\colon g' \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ of length at most ω, and thus we get a reduction $S \cdot T''\colon g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ of length at most ω. Otherwise, if S is of length ω, then we can interleave the steps in T' into S as shown in the successor case of the proof of the Compression Lemma in [9] to obtain a reduction $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h$ of length ω. ◀

Using the above compression result, we can strengthen the correspondence result of Theorem 3.7 for orthogonal GRSs as follows:

▶ **Corollary 5.13.** *Given an orthogonal, left-finite GRS $\mathcal{R}$ and two total term graphs g, h in $\mathcal{R}$, we have $g \overset{m}{\twoheadrightarrow} h$ iff $g \overset{p}{\twoheadrightarrow} h$.*

Proof. The "only if" direction follows from Theorem 3.7. By Proposition 5.12, we may assume that $g \overset{p}{\twoheadrightarrow} h$ is not longer than ω. Since g is total, and totality is preserved by reduction steps, we may apply Theorem 3.7 to conclude that $g \overset{m}{\twoheadrightarrow} h$. ◀

That is, reachability between total term graphs is independent from the choice between m- and p-convergence.

6 Concluding Remarks

Böhm extensions already entail some technical complications in term rewriting, which require some care, e.g. the additional rewrite rules may have infinite left-hand sides (which breaks the precondition of the compression lemma for example). In term graph rewriting we get additional complications: a redex may have sharing that is different from the rule's left-hand side that it instantiates. This phenomenon motivated the restriction to left-linear systems as we illustrated in Example 5.6. However, we conjecture that the issue illustrated in Example 5.6 does not occur in weakly non-overlapping systems – making the left-linearity restriction superfluous.

For the proof of our main result in Section 5, we also moved from weakly non-overlapping to non-overlapping systems, which made it possible to leverage the soundness and completeness properties from Theorem 3.9 in the proofs of Lemma 5.7 and Lemma 5.8. We conjecture that this additional restriction is not essential and merely simplified the proof at these two points.

A question that remains unanswered is whether orthogonal GRSs are confluent w.r.t. p-convergence. We conjecture that this is the case, but the technical difficulties that we already encountered in the proof of the infinitary strip lemmas appear to multiply when analysing the general case of constructing a tiling diagram.

Note that confluence of p-converging term graph reductions *modulo bisimilarity* can be easily obtained using the soundness and completeness properties from Theorem 3.9. Given two reductions $g \overset{p}{\twoheadrightarrow}_{\mathcal{R}} h_i$, $i \in \{1, 2\}$ in a left-finite orthogonal GRS $\mathcal{R}$, we have $\mathcal{U}(g) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} \mathcal{U}(h_i)$. Since $\mathcal{U}(\mathcal{R})$ is normalising and confluent w.r.t. p-convergence [2], we thus find reductions $\mathcal{U}(h_i) \overset{p}{\twoheadrightarrow}_{\mathcal{U}(\mathcal{R})} t$ to a normal form t. By completeness, we then have reductions $h_i \overset{p}{\twoheadrightarrow}_{\mathcal{R}} g_i$ with $\mathcal{U}(g_i) = t$, i.e. $g_1 \simeq g_2$. Due to the correspondence result of Theorem 5.11, this confluence property also carries over to m-convergence in the corresponding Böhm extension.

References

1. P. Bahr. Abstract models of transfinite reductions. In C. Lynch, editor, *RTA 10*, volume 6 of *LIPIcs*, pages 49–66, 2010.
2. P. Bahr. Partial order infinitary term rewriting and Böhm trees. In C. Lynch, editor, *RTA*, volume 6 of *LIPIcs*, pages 67–84, 2010.
3. P. Bahr. Infinitary term graph rewriting is simple, sound and complete. In A. Tiwari, editor, *RTA*, volume 15 of *LIPIcs*, pages 69–84, 2012.
4. P. Bahr. Convergence in infinitary term graph rewriting systems is simple. Unpublished manuscript, available from the author's website, 2013.
5. P. Bahr. Böhm reduction in infinitary term graph rewriting systems. Companion report, available from the author's website, 2017.
6. Henk P Barendregt. Representing 'undefined' in the lambda calculus. *Journal of Functional Programming*, 2:367–374, 1992.
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 *PARLE*, pages 141–158, 1987.
8. R. Kennaway and F.-J. de Vries. Infinitary rewriting. In Terese, editor, *Term Rewriting Systems*, chapter 12, pages 668–711. Cambridge University Press, 2003.
9. R. Kennaway, J.W. Klop, M.R. Sleep, and F.-J. de Vries. Transfinite reductions in orthogonal term rewriting systems. *Inf. Comput.*, 119(1):18–38, 1995.
10. R. Kennaway, V. van Oostrom, and F.-J. de Vries. Meaningless terms in rewriting. *J. Funct. Logic Programming*, 1999(1):1–35, February 1999.
11. Richard Kennaway, Jan Willem Klop, M R Sleep, and Fer-Jan de Vries. Infinitary lambda calculus. *Theoretical Computer Science*, 175(1):93–125, 1997.
12. J. Ketema. *Böhm-Like Trees for Rewriting*. PhD thesis, Vrije Universiteit Amsterdam, 2006.
13. D. Plump. Term graph rewriting. In H. Ehrig, G. Engels, H.-J. Kreowski, and G. Rozenberg, editors, *Handbook of Graph Grammars and Computing by Graph Transformation*, volume 2, pages 3–61. World Scientific Publishing Co., Inc., 1999.

A Residuals and Projections

In this appendix we give the full proofs for the theory of residuals and projections.

A.1 Preliminaries

In this section we list some properties from previous work [4] that are used for developing the theory of residuals and projections.

▶ **Lemma A.1** ([4]). *Given $g, h \in \mathcal{G}^\infty(\Sigma)$, there is a $\phi\colon g \to_\Delta h$ iff for all $\pi, \pi' \in \mathcal{P}(g)$,*

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

▶ **Definition A.2** ([4]). A position $\pi \in \mathcal{P}(g)$ in a term graph $g \in \mathcal{G}^\infty(\Sigma)$ is called *redundant* if there are $\pi_1, \pi_2 \in \mathcal{P}(g)$ with $\pi_1 < \pi_2 < \pi$ such that $\pi_1 \sim_g \pi_2$. A position that is not redundant is called *essential*. The set of all essential positions of g are denoted $\mathcal{P}^e(g)$.

Intuitively, the set of essential positions of a term graph is a minimal set of positions that still describes its structure (up to isomorphism) completely. In particular, any repetition due to cycles is omitted. The following proposition confirms that essential positions are indeed sufficient to describe the full structure of a term graph (up to isomorphism):

▶ **Proposition A.3** ([4]). *Given two term graphs $g, h \in \mathcal{G}^\infty(\Sigma)$, there is a Δ-homomorphism $\phi\colon g \to_\Delta h$ iff, for all $\pi, \pi' \in \mathcal{P}^e(g)$, we have*

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

▶ **Lemma A.4** ([4]). *Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\lambda}$ be an open reduction p-converging to g_λ.*

(i) If there is some $\alpha < \lambda$ such that $\pi \in \mathcal{P}(g_\alpha)$ and, for all $\alpha \leq \iota < \lambda$, π does not pass through n_ι in g_ι, then $g_\iota(\pi) = g_\alpha(\pi)$ for all $\alpha \leq \iota \leq \lambda$.
(ii) If $\pi \in \mathcal{P}_{\not\perp}(g_\lambda)$, then there is some $\alpha < \lambda$ such that, for all $\alpha \leq \iota < \lambda$, $g_\iota(\pi) = g_\lambda(\pi)$ and π does not pass through n_ι in g_ι.

▶ **Proposition A.5** ([4]). *A term graph $g \in \mathcal{G}^\infty(\Sigma)$ is finite iff $\mathcal{P}^e(g)$ is finite.*

▶ **Lemma A.6** ([4]). *Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\lambda}$ be an open reduction p-converging to g_λ and $\pi_1 \sim_{g_\lambda} \pi_2$. Then there is some $\alpha < \lambda$ such that $\pi_1 \sim_{g_\iota} \pi_2$ for all $\alpha \leq \iota < \lambda$.*

A.2 Residuals

We first mention the key properties from Barendregt et al. [7] that motivate the definition of residuals and projections:

▶ **Proposition A.7** (pre-reduction step residuals, [7]). *Let (n, ρ) and (n', ρ') be disjoint redex occurrences in a term graph g, with matching $\mathcal{V}$-homomorphisms ϕ and ϕ', respectively, and let $g \mapsto_{n,\rho} h$. Then n' is not a node in h, or there is a redex occurrence (n', ρ') in h.*

▶ **Proposition A.8** (reduction step projections, [7]). *Given two reduction steps $\psi\colon g \to h$ and $\psi'\colon g \to h'$ contracting two weakly disjoint redex occurrences, there are two reductions $\psi'/\psi\colon h \to^{\leq 1} g'$ and $\psi/\psi'\colon h' \to^{\leq 1} g'$.*

The following proposition confirms the claim that, for m-convergent reductions, the definition can be simplified by omitting the requirement that residual positions have to be in the set of non-$\perp$ positions of the final term graph.

▶ **Proposition A.9.** *Let $\mathcal{R}$ be a weakly non-overlapping GRS $\mathcal{R}$, $S\colon g_0 \overset{m}{\twoheadrightarrow}{}^{\alpha}_{\mathcal{R}} g_\alpha$ open, and (n, ρ) a redex occurrence in g_0 with ρ a rule in $\mathcal{R}$. Then $\liminf_{\iota\to\alpha}(n, \rho)//S|_\iota \subseteq \mathcal{P}_{\not\perp}(g_\alpha)$.*

Proof. Let $n_\iota = (n, \rho)//S|_\iota$ for each $\iota < \alpha$. To prove that $\liminf_{\iota\to\alpha} n_\iota \subseteq \mathcal{P}_{\not\perp}(g_\alpha)$, we assume some $\pi \in \liminf_{\iota\to\alpha} n_\iota$ and show that $\pi \in \mathcal{P}(g_\alpha)$. Then $\pi \in \mathcal{P}_{\not\perp}(g_\alpha)$ follows as g_α is total. Since $\pi \in \liminf_{\iota\to\alpha} n_\iota$, there is some $\beta < \alpha$ such that $\pi \in n_\iota$ for all $\beta \leq \iota < \alpha$. According to Proposition A.11, each n_ι is a node in g_ι, and, therefore, we have that $\pi \in \mathcal{P}(g_\iota)$ for all $\beta \leq \iota < \alpha$. According to Theorem 2.1, this means that $\pi \in \mathcal{P}(g_\alpha)$. ◀

▶ **Lemma A.10.** *Let $\mathcal{R}$ be a weakly non-overlapping GRS $\mathcal{R}$, $S\colon g_0 \overset{m}{\twoheadrightarrow}_{\mathcal{R}} g_\alpha$, and (n, ρ) a redex occurrence in g_0 with ρ a rule in $\mathcal{R}$. If $(n, \rho)//T = \emptyset$ for some prefix T of S, then $(n, \rho)//S = \emptyset$.*

Proof. Let $\alpha = |S|$ and $\beta = |T|$, i.e. $T = S|_\beta$. We show by induction on $\gamma \leq \alpha$ that $(n, \rho)//S|_\gamma = \emptyset$ if $\beta \leq \gamma$.

The case $\gamma \leq \beta$ is trivial. Let $\gamma = \gamma' + 1 > \beta$. Since $\gamma' \geq \beta$, we obtain by induction hypothesis that $(n, \rho)//S|_{\gamma'} = \emptyset$. Hence, $(n, \rho)//S|_\gamma = \emptyset$, too. Let $\gamma > \beta$ be a limit ordinal. According to the induction hypothesis, we know that $(n, \rho)//S|_\iota = \emptyset$ for all $\beta \leq \iota < \gamma$. Hence, $(n, \rho)//S|_\gamma = \emptyset$, too. ◀

The following proposition confirms the that our generalisation of projections to reductions of arbitrary length is well-defined:

▶ **Proposition A.11.** *Let $\mathcal{R}$ be a weakly non-overlapping GRS $\mathcal{R}$, $S\colon g_0 \overset{p}{\twoheadrightarrow}_{\mathcal{R}} g_\alpha$, and (n,ρ) a redex occurrence in g_0 with ρ a rule in $\mathcal{R}$. If $(n,\rho)//S = m$ is non-empty, then (m,ρ) is a redex occurrence in g_α.*

Proof. Let $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\alpha}$, and let $c_\iota = g_\iota \backslash n_\iota$ be the reduction context for each step at $\iota < \alpha$. We proceed by an induction on α.

The case $\alpha = 0$ is trivial. If $\alpha = \beta + 1$, then the statement follows from the induction hypothesis according to Proposition A.7.

Let α be a limit ordinal, and let $m_\iota = (n,\rho)//S|_\iota$ for all $\iota < \alpha$.

Since $m \neq \emptyset$, we know, by Lemma A.10, that $m_\iota \neq \emptyset$ for all $\iota < \alpha$. Hence, we may invoke the induction hypothesis to obtain that (m_ι, ρ) is a redex occurrence in g_ι for each $\iota < \alpha$, which means that we have matching $\mathcal{V}$-homomorphisms $\phi_\iota\colon \rho_l \to_{\mathcal{V}} g_\iota|_{m_\iota}$ for all $\iota < \alpha$.

By definition, m is a set of positions in g_α, but we have to show that m is also a node in g_α. If $\pi_1, \pi_2 \in m$, then there is some $\beta < \alpha$ such that $\pi_1, \pi_2 \in m_\iota$ for all $\beta \leq \iota < \alpha$. Since m_ι is a node in g_ι, we thus have $\pi_1 \sim_{g_\iota} \pi_2$ for all $\beta \leq \iota < \alpha$. Moreover, $\pi_1, \pi_2 \in m$ implies that $\pi_1, \pi_2 \in \mathcal{P}(g_\alpha)$, which means, according to Theorem 2.3, that we can choose β large enough such that $\pi_1, \pi_2 \in \mathcal{P}(c_\iota)$ for all $\beta \leq \iota < \alpha$. By Lemma 3.5, this means that, $\pi_1 \sim_{g_\iota} \pi_2$ implies $\pi_1 \sim_{c_\iota} \pi_2$ for all $\beta \leq \iota < \alpha$. Consequently, by Theorem 2.3, we have that $\pi_1 \sim_{g_\alpha} \pi_2$. Hence, all positions in m are in the same $\sim_{g_\alpha}$-equivalence class, which means that there is some node m' in g_α with $m \subseteq m'$.

Before we show the converse inclusion, we choose some $\pi^* \in m$. By the inclusion $m \subseteq m'$ proved above, we know that then $\pi^* \in m'$ as well. Moreover, there is some $\beta < \alpha$ such that $\pi^* \in m_\iota$ for all $\beta \leq \iota < \alpha$. We assume some $\pi \in m'$ and show that then $\pi \in m$. Since $\pi^*, \pi \in m'$, we know that $\pi^* \sim_{g_\alpha} \pi$, which means, by Theorem 2.3, that we can choose β large enough such that $\pi^* \sim_{c_\iota} \pi$ for all $\beta \leq \iota < \alpha$. According to Lemma 3.5, we thus have that $\pi^* \sim_{g_\iota} \pi$ for all $\beta \leq \iota < \alpha$. Since we know that $\pi^* \in m_\iota$, we can conclude that also $\pi \in m_\iota$ for all $\beta \leq \iota < \alpha$. We then have $\pi \in m$ because the requirement that $\pi \in \mathcal{P}_{\not\perp}(g_\alpha)$ follows from $\pi^* \sim_{g_\alpha} \pi$ and $\pi^* \in \mathcal{P}_{\not\perp}(g_\alpha)$.

By combining both inclusions, we obtain that $m = m'$, i.e. m is a node in g_α.

Before we continue, we shall prove an auxiliary claim. To this end, we pick some $\pi^* \in m$. According to the definition of residuals, we then have that $\pi^* \in \mathcal{P}_{\not\perp}(g_\alpha)$ and that there is some $\beta < \alpha$ with $\pi^* \in m_\iota$ for all $\beta \leq \iota < \alpha$. By Theorem 2.3, the former implies that we can chose β large enough such that $\pi^* \in \mathcal{P}_{\not\perp}(c_\iota)$ for all $\beta \leq \iota < \alpha$. By Lemma 3.5, this means that $c_\iota(\pi^*) = g_\iota(\pi^*)$ and that $\pi^* \notin n_\iota$ for all $\beta \leq \iota < \alpha$. Note that the latter means that $n_\iota \neq m_\iota$ for all $\beta \leq \iota < \alpha$. Since $\mathcal{R}$ is weakly non-overlapping, this implies that the redex occurrences at n_ι and m_ι must be disjoint for all $\beta \leq \iota < \alpha$.

We now proceed to prove the following claim for all $\pi \in \mathcal{P}(\rho_l)$:

$$\pi^* \cdot \pi \in \mathcal{P}(g_\alpha) \text{ and if } \rho_l(\pi) \notin \mathcal{V}, \text{ then } g_\alpha(\pi^* \cdot \pi) = \rho_l(\pi). \tag{1}$$

We prove this claim by induction on the length of π.

If $\pi = \langle\rangle$, then we know, according to the definition of term graph rules, that $\rho_l(\pi) \notin \mathcal{V}$. Hence, using Lemma A.1, we can deduce from the matching $\mathcal{V}$-homomorphisms $\phi_\iota\colon \rho_l \to_{\mathcal{V}} g_\iota|_{m_\iota}$ that $g_\iota|_{m_\iota}(\pi) = \rho_l(\pi)$ for all $\iota < \alpha$. This means that $g_\iota(\pi^* \cdot \pi) = \rho_l(\pi)$ for all $\beta \leq \iota < \alpha$. Moreover, since $\pi^* \cdot \pi = \pi^*$ and $c_\iota(\pi^*) = g_\iota(\pi^*)$, we know that $\pi^* \cdot \pi \in \mathcal{P}(g_\alpha)$ and that $c_\iota(\pi^* \cdot \pi) = \rho_l(\pi)$ for all $\beta \leq \iota < \alpha$. Consequently, by Theorem 2.3, we have that $g_\alpha(\pi^* \cdot \pi) = \rho_l(\pi)$.

If $\pi = \pi' \cdot \langle i \rangle$, then we know that $\rho_l(\pi')$ is not a nullary symbol and, thus, not in $\mathcal{V}$. By applying the induction hypothesis, we then obtain that $\pi^* \cdot \pi' \in \mathcal{P}(g_\alpha)$ and that $g_\alpha(\pi^* \cdot \pi') = \rho_l(\pi')$. Taken together these two facts imply that $\pi^* \cdot \pi \in \mathcal{P}(g_\alpha)$. If $\rho_l(\pi) \notin \mathcal{V}$, then we may apply Lemma A.1, to obtain from the matching $\mathcal{V}$-homomorphisms $\phi_\iota\colon\ \rho_l \to_\mathcal{V} g_\iota|_{m_\iota}$ that $g_\iota|_{m_\iota}(\pi) = \rho_l(\pi)$ for all $\iota < \alpha$. Since $\pi^* \in m_\iota$ for all $\beta \le \iota < \alpha$, we thus have that $g_\iota(\pi^* \cdot \pi) = \rho_l(\pi)$ for all $\beta \le \iota < \alpha$. As we have derived above, the redex occurrences at m_ι and n_ι are disjoint for all $\beta \le \iota < \alpha$. Consequently, $\pi^* \cdot \pi$ does not pass through n_ι in g_ι, which according to Lemma 3.5 implies that $c_\iota(\pi^* \cdot \pi) = g_\iota(\pi^* \cdot \pi)$ for all $\beta \le \iota < \alpha$. The resulting equality $c_\iota(\pi^* \cdot \pi) = \rho_l(\pi)$ for all $\beta \le \iota < \alpha$ together with the fact that $\pi^* \cdot \pi \in \mathcal{P}(g_\alpha)$ yields, by Theorem 2.3, that $g_\alpha(\pi^* \cdot \pi) = \rho_l(\pi)$. That concludes the proof of (1).

Finally, we show that $g_\alpha|_m$ is a ρ-redex. To this end we show the existence of a $\mathcal{V}$-homomorphisms $\phi\colon\ \rho_l \to_\mathcal{V} g_\alpha|_m$ using Lemma A.1.

(a) Let $\pi_1 \sim_{\rho_l} \pi_2$. For each $\iota < \alpha$, the matching $\mathcal{V}$-homomorphism $\phi\colon\ \rho_l \to_\mathcal{V} g_\iota|_{m_\iota}$ yields, according to Lemma A.1, that $\pi_1 \sim_{g_\iota|_{m_\iota}} \pi_2$. Consequently, $\pi^* \cdot \pi_1 \sim_{g_\iota} \pi^* \cdot \pi_2$ for all $\beta \le \iota < \alpha$. Since $\pi^* \cdot \pi_1, \pi^* \cdot \pi_2 \in \mathcal{P}(g_\alpha)$ by (1), there is, according to Theorem 2.3, some $\beta \le \beta' < \alpha$ such that $\pi^* \cdot \pi_1, \pi^* \cdot \pi_2 \in \mathcal{P}(c_\iota)$ for all $\beta' \le \iota < \alpha$. Hence, by Lemma 3.5, $\pi^* \cdot \pi_1 \sim_{g_\iota} \pi^* \cdot \pi_2$ implies $\pi^* \cdot \pi_1 \sim_{c_\iota} \pi^* \cdot \pi_2$ for all $\beta' \le \iota < \alpha$. Again using the fact that $\pi^* \cdot \pi_1, \pi^* \cdot \pi_2 \in \mathcal{P}(g_\alpha)$, we can apply Theorem 2.3 to obtain that $\pi^* \cdot \pi_1 \sim_{g_\alpha} \pi^* \cdot \pi_2$. Therefore, $\pi_1 \sim_{g_\alpha|_m} \pi_2$ as $\pi^* \in m$.

(b) Let $\rho_l(\pi) \notin \mathcal{V}$. According to (1), we then have that $g_\alpha(\pi^* \cdot \pi) = \rho_l(\pi)$. Since $\pi^* \in m$, we thus have that $g_\alpha|_m(\pi) = \rho_l(\pi)$.

◀

A.3 Compression Property

In this section, we give the missing proofs for the auxiliary lemmas used to prove the compression property.

Lemma 4.6. *Given an open reduction $S = (g_\iota \to_{n_\iota} g_{\iota+1})_{\iota<\lambda}$ p-converging to g_λ and a redex occurrence (π, ρ) in g_λ with ρ left-finite, there is a position $\pi \in \mathcal{P}(g_\lambda)$ and some $\alpha < \lambda$ such that (π, ρ) is a redex occurrence in g_ι, and π does not pass through n_ι in g_ι for any $\alpha \le \iota < \lambda$.*

Proof of Lemma 4.6. Since (π, ρ) is a redex occurrence in g_λ, there is a matching $\mathcal{V}$-homomorphism $\phi\colon\ \rho_l \to_\mathcal{V} g_\lambda|_\pi$. By Proposition A.3, this means that, for all $\pi_1, \pi_2 \in \mathcal{P}^e(\rho_l)$, we have

$$\pi_1 \sim_{\rho_l} \pi_2 \quad \Longrightarrow \quad \begin{array}{l} \pi \cdot \pi_1 \sim_{g_\lambda} \pi \cdot \pi_2, \text{ and} \\ \rho_l(\pi_1) = g_\lambda(\pi \cdot \pi_1) \quad \text{whenever} \quad \rho_l(\pi_1) \notin \mathcal{V}. \end{array} \tag{1}$$

By definition of term graph rules, we know that $\rho_l(\langle\rangle) \notin \mathcal{V}$. Hence, $g_\lambda(\pi) = \rho_l(\langle\rangle)$ and therefore $g_\lambda(\pi) \neq \bot$. Hence, we may apply Lemma A.4(ii) to obtain some $\alpha < \lambda$ such that $\pi \in \mathcal{P}(g_\iota)$ and π does not pass through n_ι in g_ι for all $\alpha \le \iota < \lambda$. It remains to be shown that there is some $\alpha \le \gamma < \lambda$ such that (π, ρ) is a redex occurrence in g_ι for all $\gamma \le \iota < \lambda$.

Since ρ is left-finite, ρ_l is finite, which means, by Proposition A.5, that $\mathcal{P}^e(\rho_l)$ is finite. Consequently, the set $P = \{\pi \cdot \pi' \mid \pi' \in \mathcal{P}^e(\rho_l)\}$ is finite as well. Hence, we may repeatedly apply Lemma A.6 (once for each pair $\pi \cdot \pi_1, \pi \cdot \pi_2 \in P$) to obtain some $\alpha \le \beta < \lambda$ such that

$$\pi \cdot \pi_1 \sim_{g_\lambda} \pi \cdot \pi_2 \text{ implies } \pi \cdot \pi_1 \sim_{g_\iota} \pi \cdot \pi_2 \text{ for all } \pi_1, \pi_2 \in \mathcal{P}^e(\rho_l) \text{ and } \beta \le \iota < \lambda \tag{2}$$

Likewise, we may repeatedly apply Lemma A.4(ii) to obtain some $\beta \leq \gamma < \lambda$ such that

$$g_\lambda(\pi \cdot \pi_1) = g_\iota(\pi \cdot \pi_1) \text{ for all } \pi_1 \in \mathcal{P}^e(\rho_l) \text{ with } \rho_l(\pi_1) \notin \mathcal{V} \text{ and } \gamma \leq \iota < \lambda \tag{3}$$

Note that we may use Lemma A.4(ii) since rules do not contain $\bot$ and by (1) above we know that $g_\lambda(\pi \cdot \pi_1) = \rho_l(\pi_1)$ for all $\pi_1 \in \mathcal{P}^e(\rho_l)$ with $\rho_l(\pi_1) \notin \mathcal{V}$.

Using both (2) and (3), we can derive from (1), that for all $\pi_1, \pi_2 \in \mathcal{P}^e(\rho_l)$ and $\gamma \leq \iota < \lambda$

$$\pi_1 \sim_{\rho_l} \pi_2 \quad \Longrightarrow \quad \begin{array}{l} \pi \cdot \pi_1 \sim_{g_\iota} \pi \cdot \pi_2, \text{ and} \\ \rho_l(\pi_1) = g_\iota(\pi \cdot \pi_1) \quad \text{whenever} \quad \rho_l(\pi_1) \notin \mathcal{V}. \end{array}$$

By Proposition A.3, the above finding implies the existence of a $\mathcal{V}$-homomorphism $\phi_\iota\colon\ \rho_l \to_\mathcal{V} g_\iota|_\pi$ for all $\gamma \leq \iota < \lambda$, i.e. (π, ρ) is a redex occurrence in g_ι. ◀

Proposition 4.8. *Let $S\colon\ g \overset{p}{\twoheadrightarrow}_\mathcal{R} h$ in a weakly non-overlapping, left-finite GRS $\mathcal{R}$. Then there is a reduction $T\colon\ g \overset{p}{\twoheadrightarrow}_\mathcal{R} h$ that is finite or open but not longer than S.*

Proof of Proposition 4.8. We proceed by induction on the length of S. The only non-trivial case is where $|S|$ is a successor ordinal greater than ω. That is, $S = S' \cdot \langle\phi\rangle$ with $S'\colon\ g \overset{p}{\twoheadrightarrow} g'$ and $\phi\colon\ g' \to h$. By induction hypothesis there is a $T'\colon\ g \overset{p}{\twoheadrightarrow} g'$ of length at most $|S'|$. If T' is finite, then so is $T' \cdot \langle\phi\rangle\colon\ g \overset{p}{\twoheadrightarrow} h$. Otherwise, T' is an open reduction. Let (π, ρ) be the redex occurrence contracted in ϕ. We will construct the diagram illustrated in Figure 3.

According to Lemma 4.6, T' can be factorised into $T_1\colon\ g \overset{p}{\twoheadrightarrow} g_1$ and $T_2\colon\ g_1 \overset{p}{\twoheadrightarrow} g'$ such that (π, ρ) is a redex occurrence in g_1 and π does not pass through the root of any redex contracted in T_2. Consequently, according to Lemma 4.7, $\pi \in (\pi, \rho)//T_2$, which means that the corresponding projection $\psi//T_2$, where $\psi\colon\ g_1 \to g_2$ contracts the redex occurrence (π, ρ) in g_1, coincides with the single step reduction ϕ. According to Theorem 4.4, the projection T_2/ψ is of type $g_2 \overset{p}{\twoheadrightarrow} h$. In sum, we have a reduction $\widehat{T} = T_1 \cdot \langle\psi\rangle \cdot T_2/\psi$ with $\widehat{T}\colon\ g \overset{p}{\twoheadrightarrow} h$. Since by construction T_2/ψ is not longer than T_2 and since T_2 is of limit ordinal length, we know that $|\langle\psi\rangle \cdot T_2/\psi| \leq |T_2|$. Consequently, $\left|\widehat{T}\right| < |S|$. Thus, we may apply the induction hypothesis to $\widehat{T}$ to obtain a reduction $T\colon\ g \overset{p}{\twoheadrightarrow} h$ of finite or limit ordinal length. ◀

Lemma 4.10. *For each reduction $S\colon\ g \overset{p}{\twoheadrightarrow}_\mathcal{R} \bot$ in a weakly non-overlapping, left-finite GRS $\mathcal{R}$ with $g \neq \bot$, we find a finite reduction $g \to^*_\mathcal{R} h$ to a redex h.*

Proof of Lemma 4.10. By Proposition 4.8, we may assume that S is finite or open. If S is finite, then S is non-empty since $g \neq \bot$. Consequently, we have that $S = T \cdot \langle\phi\rangle$ with ϕ a reduction step contracting a redex at the root. That is, T is a finite reduction from g to a redex. If S is open, we can apply Lemma 3.12 (ii), to obtain that either there is a proper prefix T of S that p-converges to $\bot$, or $\langle\rangle$ is volatile in S. The first case is impossible since $\bot$ is a normal form. In the second case, there is a proper prefix T of S that p-converges to a redex. We show by induction on the length of T, that if T p-converges to a redex, then there is a finite reduction $g \to^*_\mathcal{R} h$ to a redex. By Proposition 4.8, we may assume that T is finite or open. In the first case, we are done. In the second case, we may apply Lemma 4.6 to obtain a proper prefix T' of T that p-converges to a redex. We can then apply the induction hypothesis to T' to obtain a finite reduction from g to a redex. ◀