Abstract models of transfinite reductions

Bahr, Patrick

Published in:
Proceedings of the 21st International Conference on Rewriting Techniques and Applications

DOI:
10.4230/LIPIcs.RTA.2010.49

Publication date:
2010

Document version
Publisher's PDF, also known as Version of record

Citation for published version (APA):
Bahr, P. (2010). Abstract models of transfinite reductions. In C. Lynch (Ed.), *Proceedings of the 21st International Conference on Rewriting Techniques and Applications* (pp. 49-66). Schloss Dagstuhl - Leibniz-Zentrum für Informatik. Leibniz International Proceedings in Informatics, Vol.. 6
https://doi.org/10.4230/LIPIcs.RTA.2010.49

International Conference on Rewriting Techniques and Applications 2010 (Edinburgh), pp. 49-66
http://rewriting.loria.fr/rta/

ABSTRACT MODELS OF TRANSFINITE REDUCTIONS

PATRICK BAHR

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

ABSTRACT. We investigate transfinite reductions in abstract reduction systems. To this end, we study two abstract models for transfinite reductions: a metric model generalising the usual metric approach to infinitary term rewriting and a novel partial order model. For both models we distinguish between a weak and a strong variant of convergence as known from infinitary term rewriting. Furthermore, we introduce an axiomatic model of reductions that is general enough to cover all of these models of transfinite reductions as well as the ordinary model of finite reductions. It is shown that, in this unifying axiomatic model, many basic relations between termination and confluence properties known from finite reductions still hold. The introduced models are applied to term rewriting but also to term graph rewriting. We can show that for both term rewriting as well as for term graph rewriting the partial order model forms a conservative extension to the metric model.

1. Introduction

The study of infinitary term rewriting, introduced by Dershowitz et al. [Der91], is concerned with reductions of possibly infinite length. To formalise the concept of transfinite reductions, a variety of different models were investigated in the last 20 years. The most thoroughly studied model is the metric model, both its weak [Der91] and its strong [Ken95] variant. Other models, using for example general topological spaces [Rod98] or partial orders [Cor93, Blo04], were mostly considered to pursue specific purposes. Within these models many fundamental properties do not depend on the particular structure of terms, e.g. the property that strongly converging reductions in the metric model have countable length. Moreover, when studying these different approaches to transfinite reductions, one realises that they often share many basic properties, e.g. in how reductions can be composed and decomposed.

The purpose of this paper is to study transfinite reductions on an abstract level using several different models. This includes a metric model (Section 5) as well as a novel partial order model (Section 6), each of which induces a weak and a strong variant of convergence. Moreover, we introduce an *axiomatic model* of *transfinite abstract reduction systems* (Section 4) which captures the fundamental properties of transfinite reductions. This model

1998 ACM Subject Classification: F.4.2, F.1.1.
Key words and phrases: infinitary rewriting, metric, partial order, abstract reduction system, axiomatic, term rewriting, graph rewriting.

Proceedings of the 21st International Conference on Rewriting Techniques and Applications, Edinburgh, July, 2010
Editors: Christopher Lynch
LIPIcs - Leibniz International Proceedings in Informatics. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany
Digital Object Identifier: 10.4230/LIPIcs.RTA.2010.49

subsumes both variants of the metric and the partial order model, respectively, as well as ordinary finite reductions. In fact, we formulate these more concrete models in terms of the axiomatic model, which simplifies their presentation and their analysis substantially. To illustrate this, we reformulate well-known termination and confluence properties in the unifying axiomatic model and show that this then yields the corresponding standard termination and confluence properties for standard finite term rewriting resp. infinitary term rewriting. Additionally, we also prove that basic relations between these properties known from the finite setting also hold in this more general setting.

Lastly, we briefly mention that our models can be applied to term graph rewriting [Bar87] (Section 7) which yields the first formalisation of infinitary term graph rewriting. Moreover, we show that the partial order model is in fact superior to the metric model, at least for interesting cases like terms and term graphs: It can model convergence as in the metric model but additionally allows to distinguish between different levels of divergence.

Related Work. The idea of investigating transfinite reductions on an abstract level was first pursued by Kennaway [Ken92] by studying strongly convergent reductions in an abstract metric framework similar to ours. In this paper we will show that almost all of Kennaway's positive results (except countability of strong convergence) already hold in our more general axiomatic framework, and that countability already holds for strongly continuous reductions.

Kahrs [Kah07] investigated a more concrete model in which he considered weakly convergent reductions in term rewriting systems parametrised by the metric on terms. Although being parametric in the metric space, the results of Kahrs are tied to term rewriting and are for example not applicable to term graph rewriting [Bah09].

The use of partial orders and their notion of *limit inferior* for transfinite reductions is inspired by Blom [Blo04] who studied strongly convergent reductions in lambda calculus using a partial order and compared this to the ordinary metric model of strongly convergent reductions.

2. Preliminaries

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

Transfinite Sequences. We use $\alpha, \beta, \gamma, \lambda, \iota$ to denote ordinal numbers. A *transfinite sequence* (or simply called *sequence*) S of length α in a set A, written $(a_\iota)_{\iota<\alpha}$, is a function from α to A with $\iota \mapsto a_\iota$ for all $\iota \in \alpha$. We use $|S|$ to denote the length α of S. If α is a limit ordinal, then S is called *open*. Otherwise, it is called *closed*. If α is a finite ordinal, then S is called *finite*. Otherwise, it is called *infinite*. For a finite sequence $(a_\iota)_{\iota<n}$, we also write $\langle a_0, a_1, \dots, a_{n-1}\rangle$.

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 *prefix* of a sequence T, denoted $S \leq T$ if there is a sequence S' with $S \cdot S' = T$. The prefix of T of length β is denoted $T|_\beta$. The relation $\leq$ forms a complete semilattice.

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), $(M, \mathbf{d})$ is called an *ultrametric space.* If a sequence $(a_\iota)_{\iota<\alpha}$ in a metric space converges to an element a, we write $\lim_{\iota\to\alpha} a_\iota$ to denote a. A sequence $(a_\iota)_{\iota<\alpha}$ in a metric space 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.

Partial Orders. A *partial order* $\leq$ on a class A is a binary relation on A that is *transitive*, *reflexive*, and *antisymmetric.* A partial order $\leq$ on A is called a *complete semilattice* if it has a *least element*, every *directed subset* D of A has a *least upper bound* (*lub*) $\bigsqcup D$ in A, and every subset of A having an upper bound in A also has a least upper bound in A. Hence, complete semilattices also admit a *greatest lower bound* (*glb*) $\bigsqcap B$ for every *non-empty* subset B of A. In particular, this means that for any non-empty sequence $(a_\iota)_{\iota<\alpha}$ in a complete semilattice, its *limit inferior*, defined by $\liminf_{\iota\to\alpha} a_\iota = \bigsqcup_{\beta<\alpha}\left(\bigsqcap_{\beta\leq\iota<\alpha} a_\iota\right)$, always exists. A partial order is called a *linear order* if $a \leq b$ or $b \leq a$ holds for each pair of elements a, b. A linearly ordered subclass of a partially ordered class is also called a *chain.*

Term Rewriting Systems. Instead of finite terms, we consider the set $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ of *infinitary terms* over some *signature* Σ and a countably infinite set $\mathcal{V}$ of variables. We consider $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ as a superset of the set $\mathcal{T}(\Sigma, \mathcal{V})$ of *finite terms.* For a term $t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ we use the notation $\mathcal{P}(t)$ to denote the *set of positions* in t. For terms $s, t \in \mathcal{T}^\infty(\Sigma, \mathcal{V})$ and a position $\pi \in \mathcal{P}(t)$, we write $t|_\pi$ for the *subterm* of t at π, and $t[s]_\pi$ for the term t with the subterm at π replaced by s.

On $\mathcal{T}^\infty(\Sigma, \mathcal{V})$ a distance function $\mathbf{d}$ can be defined by $\mathbf{d}(s,t) = 0$ if $s = t$ and $\mathbf{d}(s,t) = 2^{-k}$ if $s \neq t$, where k is the minimal depth at witch s and t differ. The pair $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$ is known to form a complete ultrametric space [Arn80]. *Partial terms*, i.e. terms over signature $\Sigma_\bot = \Sigma \uplus \{\bot\}$, can be endowed with a relation $\leq_\bot$ by defining $s \leq_\bot t$ iff s can be obtained from t by replacing some subterm occurrences in t by $\bot$. The pair $(\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V}), \leq_\bot)$ is known to form a complete semilattice [Kah93].

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 this notion of a TRS is standard in infinitary rewriting [Ken03], but deviates from standard TRSs as it allows infinitary terms on the right-hand side of rules.

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

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

We write $s \to_{\pi,\rho} t$ in order to indicate the applied rule ρ and the position π.

3. Abstract Reduction Systems

In order to analyse transfinite reductions on an abstract level, we consider *abstract reduction systems (ARS)*. In ARSs, the principal items of interest are the reduction steps of the system. Therefore, the structure of the individual objects on which the reductions are performed is neglected. This abstraction is usually modelled by a pair (A, R) consisting of a set A of objects and a binary relation R on A describing the possible reductions on the objects. The ARS induced by a TRS $\mathcal{R}$ is then simply the pair $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), R)$ with $(s, t) \in R$ iff $s \to_\mathcal{R} t$.

In the setting of infinitary rewriting, however, this model is not appropriate. Instead, we need a model which *reifies* the reduction steps of the system since the semantics of transfinite reductions does not only depend on the objects involved in the reduction but also on *how* each reduction step is performed – at least when we consider *strong convergence*. However, it is not always possible to reconstruct how a reduction was performed given only the starting and end object of it due to so-called *syntactic accidents* [Lév78]: Consider the term rewrite rule $\rho\colon\ f(x) \to x$ and the term $f(f(x))$. The rule ρ can be applied both at root position $\langle\rangle$ and at position $\langle 0\rangle$ of $f(f(x))$. In both cases the resulting term is $f(x)$.

Therefore, we rather choose a model in which reduction steps are "first-class citizens" [Ter03] similarly to morphisms in a category:

Definition 3.1 (abstract reduction system)**.** An *abstract reduction system (ARS)* $\mathcal{A}$ is a quadruple $(A, \Phi, \mathsf{src}, \mathsf{tgt})$ consisting of a set of *objects* A, a set of *reduction steps* Φ, and *source* and *target functions* $\mathsf{src}\colon\ \Phi \to A$ and $\mathsf{tgt}\colon\ \Phi \to A$, respectively. We write $\varphi\colon\ a \to_\mathcal{A} b$ whenever there are $\varphi \in \Phi$, $a, b \in A$ such that $\mathsf{src}(\varphi) = a$ and $\mathsf{tgt}(\varphi) = b$.

In order to define the semantics of a TRS in terms of an ARS we only need to define an appropriate notion of a reduction step:

Definition 3.2 (operational semantics of TRSs)**.** Let $\mathcal{R} = (\Sigma, R)$ be a TRS. The ARS *induced* by $\mathcal{R}$, denoted $\mathcal{A}_\mathcal{R}$, is given by $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt})$, where $\Phi = \{(s, \pi, \rho, t) \mid s \to_{\pi,\rho} t\}$, $\mathsf{src}(\varphi) = s$ and $\mathsf{tgt}(\varphi) = t$, for each $\varphi = (s, \pi, \rho, t) \in \Phi$.

A *reduction* in this setting is simply a sequence of reduction steps in an ARS such that consecutive steps are "compatible":

Definition 3.3 (reduction)**.** A sequence $S = (\varphi_\iota)_{\iota<\alpha}$ of reduction steps in an ARS $\mathcal{A}$ is called a *reduction* if there is a sequence of objects $(a_\iota)_{\iota<\widehat{\alpha}}$ in the underlying set A, where $\widehat{\alpha} = \alpha$ if S is open, and $\widehat{\alpha} = \alpha + 1$ if S is closed, such that $\varphi_\iota\colon\ a_\iota \to a_{\iota+1}$ for all $\iota < \alpha$. For such a sequence, we also write $(\varphi\colon\ a_\iota \to a_{\iota+1})_{\iota<\alpha}$ or simply $(a_\iota \to a_{\iota+1})_{\iota<\alpha}$. The reduction S is said to start in a_0, and if S is closed, it is said to end in a_α. If S is finite, we write $S\colon\ a_0 \to^*_\mathcal{A} a_\alpha$. We use the notation $\mathsf{Red}(\mathcal{A})$ to refer to the class of all *non-empty reductions* in $\mathcal{A}$.

Observe that the empty sequence $\langle\rangle$ is always a reduction, and that $\langle\rangle$ starts and ends in a for every object a of the ARS. Also note that this notion of reductions alone does only make sense for sequences of length at most ω. For longer reductions, the ω-th step is not related to the preceding steps of the reduction:

Example 3.4. In the TRS consisting of the rules $a \to f(a)$ and $b \to g(b)$ the following constitutes a valid reduction of length $\omega \cdot 2$:

$$S\colon\ a \to f(a) \to f(f(a)) \to f(f(f(a))) \to\ \dots\ b \to g(b) \to g(g(b)) \to g(g(g(b))) \to\ \dots$$

The second half of the reduction is completely unrelated to the first half. The reason for this issue is that the ω-th reduction step $b \to g(b)$ has no immediate predecessor.

The above problem can occur for all reduction steps indexed by a limit ordinal. For successor ordinals, this is not a problem as by Definition 3.3 the $(\iota+1)$-st step is required to start in the object that the ι-th step ends in. Meaningful definitions for reductions of length beyond ω have to include an appropriate notion of *continuity* which bridges the gaps caused by limit ordinals. Exploring different variants of such a notion of continuity is the topic of the subsequent sections.

4. Transfinite Abstract Reduction Systems

In the last section we have seen that we need a notion of continuity in order to obtain a meaningful model of transfinite reductions. In this section we introduce an axiomatic framework for convergence in which we can derive a corresponding notion of continuity.

The resulting notion of continuity is quite natural and resembles the definition of continuity of real-valued functions: A reduction is *continuous* if every proper prefix *converges* to the object the subsequent suffix is starting in. In order to use this idea, we need to endow an ARS with a notion of convergence:

Definition 4.1 (transfinite abstract reduction system). A *transfinite abstract reduction system (TARS)* $\mathcal{T}$ is a tuple $(A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathsf{conv})$, such that

(i) $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt})$ is an ARS, called the *underlying ARS* of $\mathcal{T}$, and
(ii) $\mathsf{conv}\colon \mathsf{Red}(\mathcal{A}) \rightharpoonup A$ is a partial function, called *notion of convergence*, which satisfies the following two axioms:

$$\mathsf{conv}(\langle\varphi\rangle) = \mathsf{tgt}(\varphi) \quad \text{for all } \varphi \in \Phi \qquad \text{(STEP)}$$

$$\begin{gathered}\mathsf{conv}(S) = a \text{ and } \mathsf{conv}(T) = b \iff \mathsf{conv}(S \cdot T) = b \\ \text{for all } a, b \in A,\ S, T \in \mathsf{Red}(\mathcal{A}) \text{ with } T \text{ starting in } a.\end{gathered} \qquad \text{(CONCATENATION)}$$

That is, we require convergence to include single reduction steps and to be preserved by both composition and decomposition.

Axiom (CONCATENATION) is, in fact, quite comprehensive. But we can split it up into two axioms whose conjunction is equivalent to it:

$$\mathsf{conv}(S) = a \implies \mathsf{conv}(S \cdot T) = \mathsf{conv}(T) \qquad \text{(COMPOSITION)}$$

$$\mathsf{conv}(S \cdot T) \text{ defined} \implies \mathsf{conv}(S) = a \qquad \text{(CONTINUITY)}$$

where S and T range over reductions in $\mathsf{Red}(\mathcal{A})$ with T starting in $a \in A$.

Axiom (COMPOSITION) states that the composition of reductions preserves the convergence behaviour whereas (CONTINUITY) ensures that every notion of convergence already includes continuity. To see the latter we need to define convergence and continuity in TARSs:

Definition 4.2 (convergence, continuity). Let $\mathcal{T} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathsf{conv})$ be a TARS and $S \in \mathsf{Red}(\mathcal{T})$ a non-empty reduction starting in $a \in A$. S is said to *converge* to $b \in A$, written $S\colon a \twoheadrightarrow_{\mathcal{T}} b$, if $\mathsf{conv}(S) = b$. S is said to be *continuous*, written $S\colon a \twoheadrightarrow_{\mathcal{T}} \ldots$, if for every two $S_1, S_2 \in \mathsf{Red}(\mathcal{T})$ with $S = S_1 \cdot S_2$, we have that S_1 converges to the object S_2 is starting in. If S is continuous but not converging, then S is called *divergent*. For the empty reduction $\langle\rangle$, we define to have $\langle\rangle\colon a \twoheadrightarrow_{\mathcal{T}} a$ and $\langle\rangle\colon a \twoheadrightarrow_{\mathcal{T}} \ldots$ for all $a \in A$, i.e. $\langle\rangle$ is always convergent and continuous. To indicate the length α of a reduction we use the notation $\twoheadrightarrow^{\alpha}_{\mathcal{T}}$. For some object $a \in A$, we write $\mathsf{Cont}(\mathcal{T}, a)$ and $\mathsf{Conv}(\mathcal{T}, a)$ to denote the class of all continuous resp. convergent reductions in $\mathcal{T}$ starting in a.

Axiom (CONTINUITY) is equivalent to the statement that every converging reduction is also continuous. That is, only meaningful – i.e. continuous – reductions can be convergent. This is a natural model which is in particular also adopted in the theory of infinitary term rewriting [Ken03].

Returning to Example 3.4, we can see that for S to be continuous the prefix $S|_\omega$ has to converge to b. However, as one might expect, all notions of convergence for TRSs we will introduce in this paper agree on that $S|_\omega$ converges to f^ω.

Since for closed reductions not only does convergence imply continuity, but also the converse holds true, we have the following proposition:

Proposition 4.3 (convergence of closed reductions)**.** *Let $\mathcal{T}$ be a TARS and S a closed reduction in $\mathcal{T}$. Then S is continuous iff S is converging.*

Proof. The "if" direction follows from (CONTINUITY). The "only if" direction is trivial if S is empty and follows from (STEP) if S has length one. Otherwise, S is of the form $T \cdot \varphi$. Since φ is converging by (STEP) and T is converging by (CONTINUITY), S is converging due to (COMPOSITION). ∎

It is obvious from the definition that a well-defined notion of convergence has to include at least all finite (non-empty) reductions. In fact, the trivial notion of convergence which consists of precisely the *finite reductions* is the least notion of convergence w.r.t. set inclusion of its domain:

Definition 4.4 (finite convergence)**.** Let $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt})$ be an ARS. Then the *finite convergence* of $\mathcal{A}$ is the TARS $\mathcal{A}^f = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathsf{conv})$, where conv is defined by $\mathsf{conv}(S) = b$ iff $S\colon\ a \to^*_{\mathcal{A}} b$. That is, $\mathsf{conv}(S)$ is undefined iff S is infinite.

The TARS given above can be easily checked to be well-defined, i.e. conv satisfies the axioms given in Definition 4.1. We then obtain for every reduction S that $S\colon\ a \to^*_{\mathcal{A}} b$ iff $S\colon\ a \twoheadrightarrow_{\mathcal{A}^f} b$. This shows that TARSs merely provide a generalisation of what is considered to be a well-formed reduction.

Defining conv for the finite convergence was simple. In general, however, it is quite cumbersome to define, as a notion of convergence has to already comprise the corresponding notion of continuity, i.e. satisfy (CONTINUITY). We can avoid this by defining for each partial function $\mathsf{conv}\colon \mathsf{Red}(\mathcal{A}) \rightharpoonup A$ its *continuous core* $\overline{\mathsf{conv}}\colon \mathsf{Red}(\mathcal{A}) \rightharpoonup A$. For each non-empty reduction $S = (a_\iota \to a_{\iota+1})_{\iota<\alpha}$ in $\mathcal{A}$ we define

$$\overline{\mathsf{conv}}(S) = \begin{cases} \mathsf{conv}(S) & \text{if } \forall 0 < \beta < \alpha \quad \mathsf{conv}(S|_\beta) = a_\beta \\ \text{undefined} & \text{otherwise} \end{cases}$$

We then have the following lemma:

Lemma 4.5 (continuous core)**.** *Let $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt})$ be an ARS and $\mathsf{conv}\colon \mathsf{Red}(\mathcal{A}) \rightharpoonup A$ a partial function satisfying* (STEP) *and* (COMPOSITION)*. Then $\overline{\mathsf{conv}}$ satisfies* (STEP) *and* (CONCATENATION)*, i.e. $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \overline{\mathsf{conv}})$ is a TARS.*

Proof. Straightforward. ∎

Next we have a look at transfinite versions of well-known termination and confluence properties. The basic idea for lifting these properties to the setting of transfinite reductions is to replace finite reductions, i.e. $\to^*$, with transfinite reductions, i.e. $\twoheadrightarrow$.

Applied to the properties *confluence* (CR), *normalisation* (WN), and the *unique normal form property w.r.t. reduction* ($\mathsf{UN}_\to$) we obtain the following transfinite properties:

- CR^∞: If $b \twoheadleftarrow a \twoheadrightarrow c$, then $b \twoheadrightarrow d \twoheadleftarrow c$.
- WN^∞: For each a, there is a normal form b with $a \twoheadrightarrow b$.
- $\mathsf{UN}^\infty_\to$: If $b \twoheadleftarrow a \twoheadrightarrow c$ and b, c are normal forms, then $b = c$.

For properties involving convertibility, i.e. $\leftrightarrow^*$, one has to be more careful. The seemingly straightforward formalisation using transfinite reductions in the *symmetric closure* of the underlying ARS does not work since we do not have a notion of convergence for the symmetric closure. Even if we had one, as in the more concrete models that use a metric space or a partial order, the resulting transfinite convertibility relation would not be symmetric [Bah09].

We therefore follow the approach of Kennaway [Ken92]:

Definition 4.6 (transfinite convertibility)**.** Let $\mathcal{T}$ be a TARS, and a, b objects in $\mathcal{T}$. The objects a and b are called *transfinitely convertible*, written $a \twoheadleftarrow\!\!\twoheadrightarrow_\mathcal{T} b$, whenever there is a finite sequence of objects $a_0, \dots, a_n$, $n \geq 0$, in $\mathcal{T}$ such that $a_0 = a$, $a_n = b$, and, for each $0 \leq i < n$, we have $a_i \twoheadrightarrow_\mathcal{T} a_{i+1}$ or $a_i \twoheadleftarrow_\mathcal{T} a_{i+1}$. The minimal n of such a sequence is called the length of $a \twoheadleftarrow\!\!\twoheadrightarrow_\mathcal{T} b$.

This definition of transfinite convertibility is in some sense not "fully transfinite": For two objects to be transfinitely convertible, there has to be a transfinite "reduction" which may only finitely often changes its direction. However, with this definition, transfinite convertibility is an equivalence relation as desired, and we can establish an alternative characterisation of CR^∞ analogously to the original finite version:

Proposition 4.7 (alternative characterisation of CR^∞)**.** *Let $\mathcal{T}$ be a TARS.*

$$\mathcal{T} \text{ is } \mathsf{CR}^\infty \qquad \Longleftrightarrow \qquad \text{Whenever } a \twoheadleftarrow\!\!\twoheadrightarrow b, \text{ then } a \twoheadrightarrow c \twoheadleftarrow b.$$

Proof. The argument is the same as for finite reductions: The "if" direction is trivial, and the "only if" direction can be proved by an induction on the length of $a \twoheadleftarrow\!\!\twoheadrightarrow b$. ∎

With the definition of transfinite convertibility in place, we can define the transfinite versions of the *normal form property* (NF) and the *unique normal form property* (UN):

- NF^∞: For each object a and normal form b with $a \twoheadleftarrow\!\!\twoheadrightarrow b$, we have $a \twoheadrightarrow b$.
- UN^∞: All normal forms a, b with $a \twoheadleftarrow\!\!\twoheadrightarrow b$ are identical.

The above definition of NF^∞ differs from that of Kennaway et al. [Ken95] who, instead of $a \twoheadleftarrow\!\!\twoheadrightarrow b$, use $a \twoheadleftarrow c \twoheadrightarrow b$ as the precondition. One can, however, easily show that both definitions are equivalent.

Having these transfinite properties, we can establish some relations between them analogously to the setting of finite reductions:

Proposition 4.8 (confluence properties)**.** *For every TARS, the following implications hold:*

(i) $\mathsf{CR}^\infty \implies \mathsf{NF}^\infty \implies \mathsf{UN}^\infty \implies \mathsf{UN}^\infty_\to$

(ii) $\mathsf{WN}^\infty \;\&\; \mathsf{UN}^\infty_\to \implies \mathsf{CR}^\infty$

Proof. The arguments are the same as for their finite variants. ∎

Also when formulating a transfinite version of the termination property, we have to be careful. In fact, several different formalisations of transfinite termination can be found in the literature [Ken92, Rod98, Klo05].

We suggest a notion of transfinite termination which we belief is a direct generalisation of finite termination. Recall that an object a in an ARS is terminating iff there is no infinite reduction starting in a. From this we can see that for finite reductions, we can make use of infinite reductions as a meta-concept for defining finite termination. A corresponding meta-concept for transfinite reductions is provided by the class $\mathsf{Conv}(\mathcal{T}, a)$ of converging reductions starting in a ordered by the prefix order $\leq$. The analogue of an infinite reduction, which witnesses finite non-termination, is an unbounded chain in $\mathsf{Conv}(\mathcal{T}, a)$, which witnesses transfinite non-termination:

Definition 4.9 (transfinite termination)**.** Let $\mathcal{T}$ be a TARS. An object a in $\mathcal{T}$ is said to be *transfinitely terminating* (SN^∞) if each chain in $\mathsf{Conv}(\mathcal{T}, a)$ has an upper bound in $\mathsf{Conv}(\mathcal{T}, a)$. The TARS $\mathcal{T}$ itself is called *transfinitely terminating* (SN^∞) if every object in $\mathcal{T}$ is.

The following alternative characterisation of SN^∞ will be useful for comparing our definition to other formalisations of SN^∞ in the literature:

Proposition 4.10 (transfinite termination)**.** *An object a in a TARS $\mathcal{T}$ is SN^∞ iff*
(a) $\mathsf{Cont}(\mathcal{T}, a) \subseteq \mathsf{Conv}(\mathcal{T}, a)$*, and*
(b) every chain in $\mathsf{Conv}(\mathcal{T}, a)$ *is a set.*

Proof. Note that (b) is equivalent to the statement that, for every chain C in $\mathsf{Conv}(\mathcal{T}, a)$, there is an upper bound on the length of the reductions in C.

We show the "only if" direction by proving its contraposition: If (a) is violated, then there is a divergent reduction $S\colon a \twoheadrightarrow \dots$. Hence, the set of all proper prefixes of S forms a chain in $\mathsf{Conv}(\mathcal{T}, a)$ which has no upper bound. Consequently, a is not SN^∞. If (b) is violated, transfinite non-termination of a follows immediately.

For the "if" direction, consider an arbitrary chain C in $\mathsf{Conv}(\mathcal{T}, a)$. Because of (b), C has a lub S. For each proper prefix $S' < S$, there has to be an extension $S'' \geq S$ in C. Since S'' is converging, so is S'. Consequently, S is continuous and, therefore, also convergent, due to (a). Hence, S is an upper bound for C in $\mathsf{Conv}(\mathcal{T}, a)$. ∎

The above characterisation shows that there are two different reasons for transfinite non-termination: Diverging reductions and reductions that can be extended indefinitely. This characterisation of termination closely resembles that of Rodenburg [Rod98] which, however, additionally to (a) and instead of (b) requires an upper bound on the length of reductions. This is too restrictive, since an object, in which for each ordinal α a reduction of length α to a normal form starts, is not transfinitely terminating according to Rodenburg's definition.[1] An example witnessing this difference to our definition can be devised straightforwardly.

In order to verify that our formalisation of SN^∞ is appropriate, we have to make sure that it implies WN^∞:

Proposition 4.11 (SN^∞ is stronger than WN^∞)**.** *For every TARS $\mathcal{T}$, it holds that SN^∞ implies WN^∞ for every object in $\mathcal{T}$.*

[1]In fact, in an earlier draft of this paper we adopted Rodenburg's definition. We thank the anonymous referee who pointed out the mentioned issue.

Proof. We prove the contraposition of the implication using Proposition 4.10. For this purpose, let $\mathcal{T}$ be an TARS and a some object in $\mathcal{T}$ that is not WN^∞. We show that then (a) or (b) of Proposition 4.10 is violated. For this purpose, we assume (a) and show that then (b) does not hold. To this end we define a function f on the class On of ordinal numbers such that, for each $\alpha \in \mathsf{On}$, (1) $f(\alpha)$ is a converging reduction of length α starting in a and (2) $f(\alpha)$ is a proper extension of $f(\iota)$ for all $\iota < \alpha$, i.e. $f(\alpha) > f(\iota)$. Hence, the class $\{f(\alpha) \mid \alpha \in \mathsf{On}\}$ is a chain in $\mathsf{Conv}(\mathcal{T}, a)$ which is not a set since f is a bijection from the proper class On to $\{f(\alpha) \mid \alpha \in \mathsf{On}\}$. The construction of f is justified by the principle of transfinite recursion, and the properties (1) and (2) are established by transfinite induction.

For $\alpha = 0$, both (1) and (2) are trivial. Let α be a successor ordinal $\beta + 1$. By induction hypothesis, we have $f(\beta)\colon a \twoheadrightarrow^{\beta} b$ for some b. Since a is not WN^∞, b cannot be a normal form. Hence, there is a step $\varphi\colon b \to b'$ in $\mathcal{M}$. Define $f(\alpha) = f(\beta) \cdot \langle \varphi \rangle$. That is, $f(\alpha)\colon a \twoheadrightarrow^{\alpha} b'$ which shows (1). (2) follows from the induction hypothesis since $f(\beta) < f(\alpha)$.

Let α be a limit ordinal. Since, by the induction hypothesis, (2) holds for all $f(\beta)$, we have that $F = \{f(\beta) \mid \beta < \alpha\}$ is a directed set. Hence, $f(\alpha) = \bigsqcup F$ is well-defined. Consequently, all elements in F are proper prefixes of $f(\alpha)$. This shows (2) and, additionally, it shows that $f(\alpha)$ is a reduction of length α starting in a. Since, by the induction hypothesis $f(\beta)$ is converging for each $\beta < \alpha$, we have that $f(\alpha)$ is continuous. Due to (a), $f(\alpha)$ is also convergent, which shows (1). ∎

Note that the transfinite properties we have introduced are equivalent to their finite counterpart if we consider the finite convergence of an ARS. This shows that the transfinite properties that we have given here are in fact generalisations of their original finite versions to the setting of TARS. Moreover, all counterexamples known from the finite setting carry over to the setting of transfinite reductions. This means, for example, that the implications shown in Proposition 4.11 and Proposition 4.8 are in fact strict as they are in the setting of finite reductions.

There are also many interrelations between finite properties which do not hold in the transfinite setting. Notable examples are Newman's Lemma and the implication from sub-commutativity to confluence. Counterexamples for these and other interrelations are given by Kennaway [Ken92].

5. Metric Model of Transfinite Reductions

The most common model of infinitary term rewriting is based on the complete ultra-metric space of $\mathcal{T}^\infty(\Sigma, \mathcal{V})$. One usually distinguishes between two different variants in this context: A weak variant [Der91], which only takes into account the metric space, and a strong variant [Ken95], which stipulates additional restrictions on the applications of rewrite rules in order to obtain a more well-behaved notion of convergence.

At first we introduce the abstract theory of metric reduction systems. Afterwards, we describe how this can be applied to term rewriting.

Definition 5.1 (metric reduction system)**.** A *metric reduction system (MRS)* $\mathcal{M}$ is a tuple $(A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt})$, such that

(i) $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt})$ is an ARS, called the *underlying ARS* of $\mathcal{M}$,
(ii) $\mathbf{d}\colon A \times A \to \mathbb{R}^+_0$ is a function such that $(A, \mathbf{d})$ is a metric space,
(iii) $\mathsf{hgt}\colon \Phi \to \mathbb{R}^+$ is a function, called the *height function*, and

(iv) if $\varphi\colon a \to_{\mathcal{A}} b$, then $\mathbf{d}(a,b) \le \mathsf{hgt}(\varphi)$.

If the metric of an MRS $\mathcal{M}$ is an ultrametric, then $\mathcal{M}$ is called an *ultrametric reduction system (URS)*. Furthermore, an MRS is referred to as *complete* if the underlying metric space is complete. We use the notation $\varphi\colon a \to_h b$ to indicate that $\mathsf{hgt}(\varphi) = h$.

The definition of metric reduction systems follows the idea of *metric abstract reduction systems* investigated by Kennaway [Ken92]. The essential difference between our approach and that of Kennaway is the use of abstract reduction systems with reified reduction steps instead of a family of binary relations. Moreover, unlike Kennaway, we do not restrict ourselves to complete ultrametric spaces. This will allow us to distinguish in which circumstances completeness or an ultrametric is necessary and in which not.

Before continuing the discussion of the abstract model, let us have a look at how TRSs fit into it:

Definition 5.2 (MRS semantics of TRSs). Let $\mathcal{R} = (\Sigma, R)$ be a TRS. The MRS *induced* by $\mathcal{R}$, denoted $\mathcal{M}_\mathcal{R}$, is given by $(\mathcal{T}^\infty(\Sigma,\mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt})$, where $(\mathcal{T}^\infty(\Sigma,\mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt})$ is the ARS $\mathcal{A}_\mathcal{R}$ induced by $\mathcal{R}$, $\mathbf{d}$ is the metric on $\mathcal{T}^\infty(\Sigma,\mathcal{V})$, and hgt is defined as

$$\mathsf{hgt}(\varphi) = 2^{-|\pi|}, \text{ where } \varphi\colon t \to_{\pi,\rho} t'.$$

One can easily check that $\mathcal{M}_\mathcal{R}$ indeed forms an MRS for each TRS $\mathcal{R}$. In fact, since the metric on $\mathcal{T}^\infty(\Sigma,\mathcal{V})$ is a *complete ultrametric* [Arn80], $\mathcal{M}_\mathcal{R}$ is a *complete URS*.

Next we define for each MRS two notions of convergence:

Definition 5.3 (convergence in MRSs). Let $\mathcal{M} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt})$ be an MRS. The *weak convergence* of $\mathcal{M}$, denoted $\mathcal{M}^w$, is the TARS given by the tuple $(A, \Phi, \mathsf{src}, \mathsf{tgt}, \overline{\mathsf{conv}^w})$, where $\mathsf{conv}^w(S) = \lim_{\iota\to\widehat{\alpha}} a_\iota$ for a reduction $S = (a_\iota \to a_{\iota+1})_{\iota<\alpha}$. The *strong convergence* of $\mathcal{M}$, denoted $\mathcal{M}^s$, is the TARS given by the tuple $(A, \Phi, \mathsf{src}, \mathsf{tgt}, \overline{\mathsf{conv}^s})$, where $\mathsf{conv}^s(S) = \lim_{\iota\to\widehat{\alpha}} a_\iota$ for a reduction $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ if S is closed or $\lim_{\iota\to\alpha} h_\iota = 0$; otherwise it is undefined.

The notions of convergence defined above yield precisely the weakly converging [Der91] resp. the strongly converging [Ken95] reductions typically considered in the literature on infinitary term rewriting [Ken03].

From the definition we can immediately derive that strong convergence implies weak convergence. Hence, also strong continuity implies weak continuity.

Note that the height function hgt provides an overapproximation $\mathsf{hgt}(\varphi)$ of the real distance $\mathbf{d}(a,b)$ between the objects a, b involved in a reduction step $\varphi\colon a \to b$. Intuitively, speaking, the difference between weak and strong convergence is that, in the latter variant, the underlying sequence of objects $(a_\iota)_{\iota<\widehat{\alpha}}$ has to converge for the overapproximation provided by hgt as well. In fact, if it is a precise approximation, then weak and strong convergence coincide:

Fact 5.4 (equivalence of weak and strong convergence). *Let $\mathcal{M} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt})$ be an MRS with $\mathsf{hgt}(\varphi) = \mathbf{d}(a,b)$ for every reduction step $\varphi\colon a \to b \in \Phi$. Then for each reduction S in $\mathcal{M}$ we have*

$$(i)\ S\colon a \twoheadrightarrow_{\mathcal{M}^w} \dots \text{ iff } S\colon a \twoheadrightarrow_{\mathcal{M}^s} \dots, \quad \text{and} \quad (ii)\ S\colon a \twoheadrightarrow_{\mathcal{M}^w} b \text{ iff } S\colon a \twoheadrightarrow_{\mathcal{M}^s} b.$$

Proof. We only need to show that conv^s and conv^w coincide for $\mathcal{M}$. For closed reductions this is trivial. Let $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ be an open reduction. If $\mathsf{conv}^w(S)$ is undefined, then so is $\mathsf{conv}^s(S)$. If $\mathsf{conv}^w(S)$ is defined, then the sequence $(a_\iota)_{\iota<\alpha}$ converges and is therefore Cauchy. Consequently, the sequence $(\mathbf{d}(a_\iota, a_{\iota+1}))_{\iota<\alpha}$ tends to 0 which implies that also $(h_\iota)_{\iota<\alpha}$ tends to 0 as $h_\iota = \mathbf{d}(a_\iota, a_{\iota+1})$ for each $\iota < \alpha$. Thus, $\mathsf{conv}^s(S) = \mathsf{conv}^w(S)$. ∎

It is instructive to see how hgt provides an overapproximation of the distance function for the example of terms: It assumes that the metric distance between redex and contractum is maximal. That is, the height function only provides a precise approximation if every redex has a root symbol different from the one of its contractum as it is the case for the rule $\rho_1\colon\ c \to g(c)$: The reduction $f(c) \to_{\rho_1} f(g(c)) \to_{\rho_1} f(g(g(c))) \to_{\rho_1} \dots$ converges both weakly and strongly to $f(g^\omega)$. For the rule $\rho_2\colon\ f(x) \to f(g(x))$ this is not the case; both redex and contractum have the same root symbol f. The reduction $f(c) \to_{\rho_2} f(g(c)) \to_{\rho_2} f(g(g(c))) \to_{\rho_2} \dots$ now converges weakly to $f(g^\omega)$ but is not strongly converging.

Note that this also shows the need for reifying reduction steps since in a system containing both ρ_1 and ρ_2 a reduction of the shape $f(c) \to f(g(c)) \to f(g(g(c))) \to \dots$ can be strongly convergent or not, depending on which rules are applied. Similarly, with only a single rule $\rho_3\colon\ g(x) \to g(g(x))$ a reduction of the shape $g(c) \to g(g(c)) \to g(g(g(c))) \to \dots$ can be strongly converging or not, depending on *where* ρ_3 is applied.

The reason for considering strong convergence is that it is considerably more well-behaved [Ken95] than weak convergence [Sim04]. However, weak convergence in the systems characterised in Fact 5.4 inherit the nice properties of strong convergence. For TRSs these systems are precisely those for which the root-symbol of each right-hand side is a function symbol different from the root symbol of the corresponding left-hand side.

When dealing with complete URSs, strong convergence can be characterised by the height only:

Proposition 5.5 (strong convergence in complete URSs)**.** *Let $\mathcal{M}$ be a complete URS. Every open strongly continuous reduction $(a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ in $\mathcal{M}$ is strongly convergent iff $(h_\iota)_{\iota<\alpha}$ tends to* 0*.*

Proof. The "only if" direction is immediate from the definition of strong convergence. For the "if" direction, assume a strongly continuous reduction $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ with $\lim_{\iota\to\alpha} h_\iota = 0$. Then $\lim_{\iota\to\alpha} \mathbf{d}(a_\iota, a_{\iota+1}) = 0$ which in turn implies that $(a_\iota)_{\iota<\alpha}$ is Cauchy as $\mathbf{d}$ is an ultrametric. Since we have a complete metric space, this means that $(a_\iota)_{\iota<\alpha}$ converges. From this and $\lim_{\iota\to\alpha} h_\iota = 0$ we can conclude that S is strongly converging. ∎

Having a complete URS is crucial for the "if" direction of Proposition 5.5. If $\mathcal{M}$ it is not a URS, the underlying sequence $(a_\iota)_{\iota<\alpha}$ might not be Cauchy:

Example 5.6. Consider the MRS $\mathcal{M}$ in the complete metric (but not ultrametric) space $(\mathbb{R}, \mathbf{d})$ with reduction steps of the form $a \to_b (a+b)$, for each $a \in \mathbb{R}, b \in \mathbb{R}^+$. More formally, $\mathcal{M}$ is defined by $\mathcal{M} = (\mathbb{R}, \mathbb{R} \times \mathbb{R}^+, \mathsf{src}, \mathsf{tgt}, \mathbf{d}, \mathsf{hgt})$ with $\mathsf{src}((a,b)) = a$, $\mathsf{tgt}((a,b)) = a+b$, and $\mathsf{hgt}((a,b)) = b$ for all $(a,b) \in \mathbb{R} \times \mathbb{R}^+$. We then have the following reduction in $\mathcal{M}$:

$$0 \to_1 1 \to_{\frac{1}{2}} \left(1 + \frac{1}{2}\right) \to_{\frac{1}{3}} \left(1 + \frac{1}{2} + \frac{1}{3}\right) \to_{\frac{1}{4}} \dots$$

This reduction is trivially strongly continuous but not strongly convergent even though the sequence $(\frac{1}{1+i})_{i<\omega}$ of heights tends to 0. It is not even weakly converging since the series $\sum_{k=1}^\infty \frac{1}{k}$ is known to be diverging.

On the other hand, if $\mathcal{M}$ is not complete $(a_\iota)_{\iota<\alpha}$ might not converge:

Example 5.7. Consider the TRS $\mathcal{R}$ with the single rule $a \to f(a)$ and the MRS $\mathcal{M}$ which can be obtained from the induced MRS $\mathcal{M}_\mathcal{R}$ by taking $\mathcal{T}(\Sigma, \mathcal{V})$ as the set of objects instead of $\mathcal{T}^\infty(\Sigma, \mathcal{V})$. Then we have the following reduction in $\mathcal{M}$:

$$a \to_1 f(a) \to_{\frac{1}{2}} f(f(a)) \to_{\frac{1}{4}} f(f(f(a))) \to_{\frac{1}{8}} \dots$$

This reduction is trivially strongly continuous but not strongly convergent, even though the sequence $(2^{-i})_{i<\omega}$ of heights tends to 0. The reduction is not even weakly convergent as the sequence $(f^i(a))_{i<\omega}$ does converge to f^ω in the complete ultrametric space $(\mathcal{T}^\infty(\Sigma, \mathcal{V}), \mathbf{d})$ but does not converge in the incomplete ultrametric space $(\mathcal{T}(\Sigma, \mathcal{V}), \mathbf{d})$

From the above characterisation of strong convergence, we can derive the following more general characterisation:

Proposition 5.8 (strong convergence)**.** *Let S be a reduction in an MRS $\mathcal{M}$.*

(i) If S is strongly convergent, then, for any $h \in \mathbb{R}^+$, there are at most finitely many steps in S whose height is greater than h.

(ii) If S is weakly continuous and, for any $h \in \mathbb{R}^+$, there are at most finitely many steps in S whose height is greater than h, then S is strongly continuous. If, additionally, $\mathcal{M}$ is a complete URS, then S is even strongly convergent.

Proof. (i) The proof of Kennaway [Ken92] also works for MRSs.

(ii) Let $S = (a_\iota \to_{h_\iota} a_{\iota+1})_{\iota<\alpha}$ be a reduction in $\mathcal{M}$. Suppose that S is weakly continuous, and that the set $\{\iota \mid h_\iota > h\}$ is finite for each $h \in \mathbb{R}^+$. We have to show that $\lim_{\iota\to\lambda} h_\iota = 0$ for each limit ordinal $\lambda < \alpha$. To this end, let $\varepsilon > 0$. Then choose some h such that $0 < h < \varepsilon$. Since, by hypothesis, the set $\{\iota \mid h_\iota > h\}$ is finite, there is some ordinal $\beta < \lambda$ such that $h_\iota \leq h < \varepsilon$ for all $\beta < \iota < \lambda$. Hence, $\lim_{\iota\to\lambda} h_\iota = 0$.

The second part of (ii) is follows from Proposition 4.3 if S is closed. Otherwise it follows from Proposition 5.5. ∎

The restriction to complete URSs in the second part of (ii) is essential as Example 5.6 and Example 5.7 illustrate.

From this proposition, the following corollary follows as shown by Kennaway [Ken92]:

Corollary 5.9 (countable length of strongly convergent reductions)**.** *In an MRS every strongly convergent reduction has countable length.*

As a result of the above corollary, part (b) of Proposition 4.10 is always satisfied for strong convergence. This makes our definition of SN^∞ equivalent to that of Klop and de Vrijer [Klo05], who considered strong convergence only.

By employing an argument similar to the one used by Klop and de Vrijer [Klo05] for the particular case of infinitary term rewriting, we can generalise Corollary 5.9 to strongly *continuous* reductions, provided we have a complete URS.

Proposition 5.10 (countable length of strongly continuous reductions)**.** *Every strongly continuous reduction in a complete URS has countable length.*

This generalises corresponding results of Kennaway [Ken92] and Klop and de Vrijer [Klo05]. The above proposition is not true for weakly continuous (or convergent) reductions as pointed out by Kennaway [Ken92].

6. Partial Order Model of Transfinite Reductions

The metric model of transfinite reductions has rather restrictive notions of convergence. For example, suppose that we have a TRS consisting of the rules

$$f(x,a) \to f(s(x),b), \quad f(x,b) \to f(s(x),a).$$

Then we can construct the reduction

$$f(0,a) \to f(s(0),b) \to f(s(s(0)),a) \to f(s(s(s(0))),b) \to \dots$$

which is neither strongly nor weakly convergent in terms of its MRS semantics. The culprit is the second argument of the f symbol which constantly changes between a and b. However, excluding this "flickering", the reduction seems to converge somehow. The investigation of *partial reduction systems* is aimed at formalising this relaxation of the notion of convergence. With this tool we will be able to identify $f(s^\omega, \bot)$ as the limit of the reduction above.

To this end, a partially ordered set is employed rather than a metric space, and the limit construction is replaced by the limit inferior.

Definition 6.1 (partial reduction system)**.** A *partial reduction system (PRS)* $\mathcal{P}$ is a tuple $(A, \Phi, \mathsf{src}, \mathsf{tgt}, \leq, \mathsf{cxt})$ such that

(i) $\mathcal{A} = (A, \Phi, \mathsf{src}, \mathsf{tgt})$ is an ARS, called the *underlying ARS* of $\mathcal{P}$,
(ii) $(A, \leq)$ is a partially ordered set,
(iii) $\mathsf{cxt}\colon\ \Phi \to A$ is a function, called the *context function*, and
(iv) if $\varphi\colon\ a \to_{\mathcal{A}} b$, then $\mathsf{cxt}(\varphi) \leq a, b$.

If the partial order $\leq$ is a complete semilattice, then $\mathcal{P}$ is called *complete*. We use the notation $\varphi\colon\ a \to_c b$ to indicate that $\mathsf{cxt}(\varphi) = c$.

Also this model can be applied to TRSs. Note, however, that we have to add a fresh constant symbol $\bot$ to the signature in order to use the partial order $\leq_\bot$:

Definition 6.2 (PRS semantics of TRSs)**.** Let $\mathcal{R} = (\Sigma, R)$ be a TRS. The PRS *induced* by $\mathcal{R}$, denoted $\mathcal{P}_\mathcal{R}$, is given by $(\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt}, \leq_\bot, \mathsf{cxt})$, with $(\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V}), \Phi, \mathsf{src}, \mathsf{tgt})$ the ARS $\mathcal{A}_{\mathcal{R}'}$ induced by the TRS $\mathcal{R}' = (\Sigma_\bot, R)$, $\leq_\bot$ the usual partial order on $\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$, and cxt defined by

$$\mathsf{cxt}(\varphi) = t[\bot]_\pi, \text{ where } \varphi\colon\ t \to_{\pi,\rho} t'.$$

One can easily verify that the context function defined for TRSs satisfies the condition $\mathsf{cxt}(\varphi\colon\ a \to b) \leq a, b$. Since the partial order on terms forms a *complete semilattice*, this means that the PRS $\mathcal{P}_\mathcal{R}$ induced by a TRS $\mathcal{R}$ is always a *complete PRS*.

Definition 6.3 (convergence of PRSs)**.** Let $\mathcal{P} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \leq, \mathsf{cxt})$ be a PRS. The *weak convergence* of $\mathcal{P}$, denoted $\mathcal{P}^w$, is the TARS given by the tuple $(A, \Phi, \mathsf{src}, \mathsf{tgt}, \overline{\mathsf{conv}}^w)$, where $\mathsf{conv}^w(S) = \liminf_{\iota \to \widehat{\alpha}} a_\iota$ for a reduction $S = (a_\iota \to a_{\iota+1})_{\iota<\alpha}$. The *strong convergence* of $\mathcal{P}$, denoted $\mathcal{P}^s$, is the TARS given by the tuple $(A, \Phi, \mathsf{src}, \mathsf{tgt}, \overline{\mathsf{conv}}^s)$, where, for a reduction $S = (a_\iota \to_{c_\iota} a_{\iota+1})_{\iota<\alpha}$, $\mathsf{conv}^s(S) = a_\alpha$ if α is a successor ordinal, and $\mathsf{conv}^s(S) = \liminf_{\iota \to \alpha} c_\iota$ if α is a limit ordinal.

Since the limit inferior is always defined for complete semilattices, we immediately obtain that for complete PRSs, continuity and convergence coincide. That is, a reduction is weakly (resp. strongly) continuous iff it is weakly (resp. strongly) convergent. This fact is the main motivation for considering the partial order model as an alternative to the metric model. As a consequence, part (a) of Proposition 4.10 is always satisfied for complete PRSs.

Returning to the initial example of this section we can now observe that the given reduction sequence weakly converges to $f(s^\omega, \bot)$ and strongly converges to $\bot$.

This example also illustrates a major difference compared to the metric model: In MRSs strong convergence is defined by restricting weak convergence. Hence, if a reduction is both weakly and strongly converging, the final result is the same and strong convergence implies weak convergence. For PRSs, however, strong convergence and weak convergence are defined differently. As a result, unlike for MRSs, strong convergence does not imply weak convergence. In order to obtain this behaviour we have to consider *total reductions*:

Definition 6.4 (total reduction). Let $\mathcal{P}$ be a PRS and $S = (a_\iota \to a_{\iota+1})_{\iota<\alpha}$ a reduction in $\mathcal{P}$. We say that S is *total* if each element a_ι is maximal w.r.t. the partial order of $\mathcal{P}$. If we write S as $S\colon a_0 \twoheadrightarrow_{\mathcal{P}^w} a_\alpha$ or $S\colon a_0 \twoheadrightarrow_{\mathcal{P}^s} a_\alpha$, i.e. the convergence of the reduction is explicitly stated, we additionally require a_α to be maximal for S to be total.

Proposition 6.5 (strong convergence implies weak convergence). *For every total reduction S in a PRS $\mathcal{P}$, it holds that*

(i) $S\colon a \twoheadrightarrow_{\mathcal{P}^s} \dots$ implies $S\colon a \twoheadrightarrow_{\mathcal{P}^w} \dots$, and that
(ii) $S\colon a \twoheadrightarrow_{\mathcal{P}^s} b$ implies $S\colon a \twoheadrightarrow_{\mathcal{P}^w} b$.

Proof. Let $S = (a_\iota \to_{c_\iota} a_{\iota+1})_{\iota<\alpha}$. We only need to show that $\mathsf{conv}^s(S) = \mathsf{conv}^w(S)$ whenever $\mathsf{conv}^s(S)$ is a maximal object in $\mathcal{P}$. If S is closed, this is trivial. If S is open we have $\mathsf{conv}^s(S) = \liminf_{\iota\to\alpha} c_\iota \le \liminf_{\iota\to\alpha} a_\iota = \mathsf{conv}^w(S)$ since, by definition, $c_\iota \le a_\iota$ for each $\iota < \alpha$. Because $\mathsf{conv}^s(S)$ is maximal, we can conclude that $\mathsf{conv}^s(S) = \mathsf{conv}^w(S)$. ∎

Despite this difference to MRSs, the intuition of the distinction between weak and strong convergence remains the same: Like the height in an MRS, the context $\mathsf{cxt}(\varphi)$ in a PRS overapproximates the difference between the objects a, b involved in a reduction step $\varphi\colon a \to b$. More precisely, it underapproximates the shared structure $a \sqcap b$ of a and b, where $a \sqcap b$ denotes the glb of $\{a, b\}$ w.r.t. the partial order of the PRS. This follows from the condition $\mathsf{cxt}(\varphi) \le a, b$ which implies $\mathsf{cxt}(\varphi) \le a \sqcap b$. Likewise, weak and strong convergence coincide if the approximation provided by cxt is precise:

Fact 6.6 (equivalence of weak and strong convergence). *Let $\mathcal{P} = (A, \Phi, \mathsf{src}, \mathsf{tgt}, \le, \mathsf{cxt})$ be a complete PRS with $\mathsf{cxt}(\varphi) = a \sqcap b$ for every reduction step $\varphi\colon a \to b \in \Phi$. Then for each reduction S in $\mathcal{P}$ we have*

$$(i)\ S\colon a \twoheadrightarrow_{\mathcal{P}^w} \dots \text{ iff } S\colon a \twoheadrightarrow_{\mathcal{P}^s} \dots, \quad \text{and} \quad (ii)\ S\colon a \twoheadrightarrow_{\mathcal{P}^w} b \text{ iff } S\colon a \twoheadrightarrow_{\mathcal{P}^s} b.$$

Proof. Analogously to the proof of Fact 5.4 using the observation that $\liminf_{\iota\to\lambda} a_\iota = \liminf_{\iota\to\lambda}(a_\iota \sqcap a_{\iota+1})$ for all open sequences $(a_\iota)_{\iota<\lambda}$ in a complete semilattice. ∎

Again this fact allows us to transfer results for strong convergence [Bah09] to the setting of weak convergence. And as for Fact 5.4 we can derive from Fact 6.6 that weak and strong convergence coincide for TRSs for which the root symbol of each right-hand side is a function symbol different from the root symbol of the corresponding left-hand side.

7. Metric vs. Partial Order Model

The main motivation for the partial order model is to have a more fine-grained notion of convergence. That is, instead of only being able to distinguish converging and diverging reductions, we have intermediate levels between full convergence and full divergence. Since, in complete PRSs, continuous reductions are always convergent, the final object of a reduction S indicates the "level of convergence" according to the partial order on objects. If it is $\bot$, the least element of the partial order, then S can be considered fully diverging. If it is a maximal element, e.g. in $\mathcal{T}^\infty(\Sigma_\bot, \mathcal{V})$ a term not containing $\bot$, then S is fully converging.

Using this intuition, the partial order model also gives rise to a notion of *meaninglessness*: We can consider an object a of a complete PRS meaningless if there is an open reduction from a converging to $\bot$. In fact, for strong convergence in orthogonal TRSs, this concept of meaninglessness coincides with so-called *root-active terms* [Bah10].

Under certain quite natural conditions [Bah09], metric convergence can be considered as the fragment of partial order convergence that only considers full convergence. Vice versa, partial order convergence is a conservative extension to metric convergence which also allows partial convergence. This is, in fact, the case for TRSs:

Theorem 7.1 (PRS semantics of TRSs extends MRS semantics)**.** *For each TRS $\mathcal{R}$, the following holds for each $c \in \{w, s\}$:*

(i) $S\colon a \twoheadrightarrow_{\mathcal{P}^c_\mathcal{R}} \dots$ *is total iff* $S\colon a \twoheadrightarrow_{\mathcal{M}^c_\mathcal{R}} \dots$

(ii) $S\colon a \twoheadrightarrow_{\mathcal{P}^c_\mathcal{R}} b$ *is total iff* $S\colon a \twoheadrightarrow_{\mathcal{M}^c_\mathcal{R}} b$*.*

It has been shown [Bah09] that also on so-called *term graphs*, a generalisation of terms, an appropriate complete ultrametric and complete semilattice can be defined. These concepts generalise the metric and the partial order on terms and allow to define infinitary term graph rewriting in our models of transfinite reductions. Following the framework of term graph rewriting systems (TGRSs) of Barendregt et al. [Bar87] one can show that, at least for weak convergence, the same relation between the partial order and the metric model can be observed:

Theorem 7.2 (PRS semantics of TGRSs extends MRS semantics)**.** *For each TGRS $\mathcal{R}$, the following holds:*

(i) $S\colon a \twoheadrightarrow_{\mathcal{P}^w_\mathcal{R}} \dots$ *is total iff* $S\colon a \twoheadrightarrow_{\mathcal{M}^w_\mathcal{R}} \dots$

(ii) $S\colon a \twoheadrightarrow_{\mathcal{P}^w_\mathcal{R}} b$ *is total iff* $S\colon a \twoheadrightarrow_{\mathcal{M}^w_\mathcal{R}} b$*.*

8. Conclusions

The axiomatic model of transfinite reductions provides a simple framework to formulate and analyse the more concrete models presented here and is yet powerful enough to establish many of their fundamental properties. Moreover, the equivalence of transfinite properties for finite convergence and their respective finite counterparts provides additional evidence for the appropriateness of the definition of these transfinite properties.

Fact 5.4 and Fact 6.6 suggest that the metric and the partial order model have a considerable similarity in their discrimination between weak and strong convergence. This raises the question whether there is an appropriate abstraction of these two models that, in contrast to the axiomatic model, is also able to distinguish between weak and strong convergence.

Theorems 7.1 and 7.2 indicate that the partial order model is superior to the metric model as it is able to express convergence as the metric model but additionally allows to explore different levels of divergence in the metric model. Moreover, these results allow to make use of well-known properties of metric infinitary term rewriting in order to study partial order infinitary term rewriting. This was used in [Bah10] to establish several properties of partial order infinitary orthogonal term rewriting such as compression and convergence.

The models that we presented here can be, of course, easily applied to higher-order rewriting systems [Ket05]. However, in the metric approach to infinitary lambda-calculus [Ken97] one usually considers various different metrics and it is not clear what the corresponding partial orders are which then admit a higher-order version of Theorem 7.1.

Acknowledgements

I would like to thank Jakob Grue Simonsen and the alert anonymous referees for carefully reading earlier drafts of this paper and providing valuable feedback. Especially, I want to thank Bernhard Gramlich for his support and his challenging questions during my work on my master's thesis which made this work possible.

References

[Arn80] A. Arnold and M. Nivat. The metric space of infinite trees. Algebraic and topological properties. *Fundam. Inf.*, 3(4):445–476, 1980.

[Bah09] P. Bahr. *Infinitary Rewriting - Theory and Applications*. Master's thesis, Vienna University of Technology, 2009.
URL `http://www.pa-ba.info/?q=pub/master`

[Bah10] P. Bahr. Partial order infinitary term rewriting and böhm trees. In *RTA 2010*. 2010. cf. this proceedings volume.

[Bar87] H.P. Barendregt, M.C.J.D. van Eekelen, J.R.W. Glauert, Ri. Kennaway, M.J. Plasmeijer, and M.R. Sleep. Term graph rewriting. In Philip C. Treleaven Jaco de Bakker, A. J. Nijman (ed.), *PARLE 1987*, *LNCS*, vol. 259, pp. 141–158. Springer Berlin / Heidelberg, 1987. doi:10.1007/3-540-17945-3_8.
URL `http://www.springerlink.com/content/pw65n058434q1k65/`

[Blo04] S. Blom. An approximation based approach to infinitary lambda calculi. In Vincent van Oostrom (ed.), *RTA 2004*, *LNCS*, vol. 3091, pp. 221–232. Springer Berlin / Heidelberg, 2004. doi:10.1007/b98160.
URL `http://www.springerlink.com/content/4n3gqw43d1bpnldy/`

[Cor93] A. Corradini. Term rewriting in CT_Σ. In Marie-Claude Gaudel and Jean-Pierre Jouannaud (eds.), *TAPSOFT 1993*, *LNCS*, vol. 668, pp. 468–484. Springer Berlin / Heidelberg, 1993. doi:10.1007/3-540-56610-4_83.
URL `http://www.springerlink.com/content/f73r5p2v370220m4/`

[Der91] N. Dershowitz, S. Kaplan, and D.A. Plaisted. Rewrite, rewrite, rewrite, rewrite, rewrite, ... *Theor. Comput. Sci.*, 83(1):71–96, 1991. doi:DOI:10.1016/0304-3975(91)90040-9.
URL `http://www.sciencedirect.com/science/article/B6V1G-45DHJRB-H/2/767b35171dafdfa511dd0463ea25dbdd`

[Kah93] G. Kahn and G.D. Plotkin. Concrete domains. *Theor. Comput. Sci.*, 121(1-2):187–277, 1993. doi:DOI:10.1016/0304-3975(93)90090-G.
URL `http://www.sciencedirect.com/science/article/B6V1G-45FC431-2K/2/6c30777ef97aea14c529418b4d5c5d4a`

[Kah07] S. Kahrs. Infinitary rewriting: meta-theory and convergence. *Acta Inform.*, 44(2):91–121, 2007. doi:10.1007/s00236-007-0043-2.
URL `http://www.springerlink.com/content/gk52386u20857666/`

[Kel55] J.L. Kelley. *General Topology*, *Graduate Texts in Mathematics*, vol. 27. Springer-Verlag, 1955.

[Ken92] R. Kennaway. On transfinite abstract reduction systems. Tech. rep., CWI (Centre for Mathematics and Computer Science), Amsterdam, 1992.

[Ken95] 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. doi:DOI:10.1006/inco.1995.1075.
URL `http://www.sciencedirect.com/science/article/B6WGK-45NJJYB-4W/2/7d48d04a2fe97d6e9e1fc5179f31a488`

[Ken97] R. Kennaway, J.W. Klop, M.R. Sleep, and F.-J. de Vries. Infinitary lambda calculus. *Theor. Comput. Sci.*, 175(1):93–125, 1997. doi:DOI:10.1016/S0304-3975(96)00171-5.
URL `http://www.sciencedirect.com/science/article/B6V1G-3SNTKND-25/2/febf61177fc25aa984b58ee4bb574143`

[Ken03] R. Kennaway and F.-J. de Vries. Infinitary rewriting. In Terese [Ter03], chap. 12, pp. 668–711.
URL `http://amazon.com/o/ASIN/0521391156/`

[Ket05] J. Ketema and J. G. Simonsen. Infinitary combinatory reduction systems. In Jürgen Giesl (ed.), *RTA 2005*, *Lecture Notes in Computer Science*, vol. 3467, pp. 438–452. Springer Berlin / Heidelberg, 2005. doi:10.1007/b135673.
URL `http://www.springerlink.com/content/h96d6qlmuhbad903/`

[Klo05] J.W. Klop and R.C. de Vrijer. Infinitary normalization. In Sergei N. Artëmov, Howard Barringer, Artur S. d'Avila Garcez, Luís C. Lamb, and John Woods (eds.), *We Will Show Them! Essays in Honour of Dov Gabbay*, vol. 2, pp. 169–192. College Publications, 2005.

[Lév78] J.-J. Lévy. *Reductions Correctes et Optimales dans le Lambda-Calcul.* Ph.D. thesis, Université Paris, 1978.

[Rod98] P.H. Rodenburg. Termination and confluence in infinitary term rewriting. *J. Symbolic Logic*, 63(4):1286–1296, 1998.
URL `http://www.jstor.org/stable/2586651`

[Sim04] J.G. Simonsen. On confluence and residuals in Cauchy convergent transfinite rewriting. *Inf. Process. Lett.*, 91(3):141–146, 2004. doi:DOI:10.1016/j.ipl.2004.03.018.
URL `http://www.sciencedirect.com/science/article/B6V0F-4CBVNG4-1/2/d5d0f374f89fd62e07d023512a5b3dfe`

[Ter03] Terese. *Term Rewriting Systems.* Cambridge University Press, 1st edn., 2003.
URL `http://amazon.com/o/ASIN/0521391156/`