

MODULARITY OF CONVERGENCE AND STRONG CONVERGENCE IN INFINITARY REWRITING

STEFAN KAHRS

University of Kent, Department of Computer Science, Canterbury CT2 7NF
e-mail address: S.M.Kahrs@kent.ac.uk

ABSTRACT. Properties of Term Rewriting Systems are called *modular* iff they are preserved under (and reflected by) disjoint union, i.e. when combining two Term Rewriting Systems with disjoint signatures. *Convergence* is the property of Infinitary Term Rewriting Systems that all reduction sequences converge to a limit. *Strong Convergence* requires in addition that redex positions in a reduction sequence move arbitrarily deep.

In this paper it is shown that both Convergence and Strong Convergence are modular properties of non-collapsing Infinitary Term Rewriting Systems, provided (for convergence) that the term metrics are *granular*. This generalises known modularity results beyond metric d_∞ .

1. INTRODUCTION

Modular properties of Term Rewriting are properties that are preserved under (and reflected by) the disjoint union of signatures. They are of particular interest to reason about rewrite systems by *divide and conquer*.

Since Toyama showed that confluence *is* [22] and strong normalisation *is not* [23] modular, the modularity of many further properties have been investigated. Examples of other modular properties are: uniqueness of normal forms [14], weak normalisation, simple termination; examples of non-modular properties are: completeness, strong confluence. Modularity proofs for these properties are sometimes complex [22], sometimes simple: the modularity of weak confluence follows directly from the critical pair lemma [2]. Many modularity results come with syntactic restrictions, e.g. completeness is modular in the absence of non-left-linear rules [21] and strong normalisation in the absence of collapsing rules [16].

One aspect of Term Rewriting is that it can be viewed as a computational model for Functional Programming. Lazy Functional Programming exhibits a phenomenon that ordinary Term Rewriting does not really capture: reductions that converge to an infinite result in infinitely many steps. Partly for that reason, the concept of Infinitary Term Rewriting was devised in [5]; technically, it arises from the finite version by equipping

1998 ACM Subject Classification: F.4.2.

Key words and phrases: infinitary rewriting, convergence, modularity.

A shorter version of this paper, only concerned with metric d_∞ , has appeared in the proceedings of RTA 2009 [10].

the signature with a metric d_m and using as its universe of terms $Ter^m(\Sigma)$, the metric completion of the metric space $(Ter(\Sigma), d_m)$, where $Ter(\Sigma)$ is the set of finite terms.

It is common to only consider the metric d_∞ (which goes back to [1]) in infinitary term rewriting, ever since its beginning in [5]; d_∞ sets the distance $d_\infty(t, t')$ to 2^{-k} where k is the length of the shortest position at which the two terms t and t' have different subterm-roots. This also coincides with a straightforward co-inductive definition of infinitary terms [13]. The metric completion of d_∞ plays a special role as it is a final co-algebra and thus contains *all* infinite terms.

However, this is not the only metric of interest: for instance, the functional programming language Haskell permits to equip constructors with strictness annotations [8]. These annotations change the universe of terms by ensuring that certain infinite terms are impossible to construct. For example:

```
data LTree = Null | Bin LTree Int !LTree
```

The exclamation mark before the third argument of the constructor `Bin` labels that argument as strict. For Haskell, this means that to construct a value for `Bin t1 n t2` the expression `t2` needs to be evaluated first. As a consequence it is impossible to construct an `LTree` with an infinite right spine. The same effect can be achieved by modifying the distance function [9]; in this case, we could define

$$d_m(\text{Bin}(a, b, c), \text{Bin}(x, y, z)) = \max(d_m(c, z), d_m(b, y), \frac{1}{2} \cdot d_m(a, x))$$

With this, the distance between two finite `LTree` values would be 1 if their right spines have different lengths. Therefore, all Cauchy-sequences would eventually have right spines of a fixed finite length with the corresponding implication for metric completion.

Metrics of this form are exactly the *granular* metrics mentioned in title and abstract. They distinguish between strict and lazy argument positions, which correspond to multiplicative factors of 1 and $\frac{1}{2}$ (respectively) in the *components* of their associated *ultra-metric maps*. The associated ultra-metric maps tell us how to compute distances between terms with the same root symbol, and the components split this task into different argument positions. The reasons to treat these metrics separately are (i) the practical interest in them and (ii) that their restricted form is particularly well-behaved.

Infinitary Term Rewriting views transfinite reductions as reduction sequences that are also convergent sequences, w.r.t. its metric. In a reduction sequence each term rewrites to its successor. An infinitary Term Rewriting System (iTRS) is called *convergent* iff all its reduction sequences converge.

Before [10] there were few positive results about convergence in the literature, instead research has focussed on *strong convergence* which in addition requires that redex positions in a reduction sequence cannot stay short: if we replace each contracted redex with a random term we would still have a converging sequence.

It is known that strong normalisation is not modular for finitary Term Rewriting Systems. The well-known counter example from [23] has a strongly normalising TRS with rule $F(0, 1, x) \rightarrow F(x, x, x)$, and another one with rules $G(x, y) \rightarrow x$, $G(x, y) \rightarrow y$; their combination fails to be strong normalising as the term $F(0, 1, G(0, 1))$ reduces to itself in three steps. As we shall see, for the discrete metric strong convergence and strong normalisation coincide and therefore modularity of strong convergence can fail for other term metrics than d_∞ . The example also shows that “finite convergence”, the convergence of

all reduction sequences starting from a finite term, fails to be modular in both weak and strong form [19].

Simonsen [19] proved that strong convergence is a modular property for left-linear iTRSs (for metric d_∞); the left-linearity condition is indeed unnecessary [10]. Proving such results for ordinary convergence is more complicated.

Example 1.1. $R_1 = \{F(x, x, y) \rightarrow F(x, y, x)\}$ $R_2 = \{0 \rightarrow S(0)\}$

Both systems R_1 and R_2 are converging. In the combination we have reductions like $F(0, 0, 0) \rightarrow F(0, 0, 1) \rightarrow F(0, 1, 0) \rightarrow F(1, 1, 0) \rightarrow F(1, 0, 1) \rightarrow F(1, 1, 1)$, with $1 = S(0)$. If we project position 2 (or 3) from this reduction sequence, thus extracting a sequence of elements from the other system then we get a sequence where the values go down as well as up, i.e. on the partial order induced by R_2 : $0, 0, 1, 1, 0, 1, 1, \dots$. For an overall convergence proof we need to demonstrate why that projected sequence is supposed to be converging, and that is not trivial.

If a reduction sequence in the disjoint union of two iTRSs is converging but not strongly converging then its top layer “remains active” throughout the sequence and in particular it can rearrange its principal subterms, as happened with 0 and 1 in Example 1.1. Informally, principal subterms are those subterms where first a change of signature occurs.

These rearrangements are strongly constrained though by the fact that the iTRS associated with the top layer is itself converging. In particular, it can be shown that the subterms at a principal position in the limit form a so-called *focussed sequence*; these have peculiar properties: (i) any infinite subsequence of a focussed sequence contains an infinite weak reduction sequence; (ii) any non-converging focussed sequence contains a subsequence which is a non-converging weak reduction sequence. Here, “weak reduction sequences” are reduction sequences with gaps: an element of such a sequence does not necessarily reduce to its neighbour in a single reduction step, this can take any number of steps. For abstract reduction systems over Hausdorff spaces, convergence of all reduction sequences and convergence of all focussed sequences are equivalent properties (Proposition 3.8).

This implies that any reduction sequence in the combined iTRS must converge “layer by layer”, and therefore is converging w.r.t. the metric d_∞ . Thus such a sequence can only fail to converge if the infinite term the sequence is converging to (under metric d_∞) does not exist in the metric completion derived from the combined system’s own metric d_m . For example, if one system has the rule $F(x) \rightarrow G(x)$ and the other a unary function symbol H then the infinite term $t = F(H(t))$ may exist under the combined metric whilst $u = G(H(u))$ may not, in which case the combined system would fail to converge. In granular term metrics such rules can be dismissed as they instantly break convergence; in the example, $a = F(a)$ would and $b = G(b)$ would not exist.

2. PRELIMINARIES

To make the terminology of this paper self-contained, this section contains a reminder of some basic definitions and facts about topological and metric spaces [6, 4], as well as sequences over them.

A *topological space* is a pair (S, \mathcal{O}) where S is a set and \mathcal{O} a subset of $\wp(S)$ such that it is closed under finite intersections and arbitrary unions, and $\emptyset, S \in \mathcal{O}$. The elements of \mathcal{O} are called *open sets*, their complements w.r.t. S are called *closed sets*. A function $f : A \rightarrow B$ between topological spaces is called continuous iff $f^{-1}(X)$ is open whenever X is open.

An *open base* B is a subset of $\wp(S)$; its associated topology \mathcal{O} is the smallest set such that $B \subseteq \mathcal{O}$ and that (S, \mathcal{O}) is a topological space.

A *neighbourhood* of $x \in S$ is a set $N \subseteq S$ for which there is an $A \in \mathcal{O}$ such that $x \in A \subseteq N$. A topological space is called *Hausdorff* iff any two distinct elements of S have disjoint neighbourhoods. It is called *discrete* if $\mathcal{O} = \wp(S)$.

A *sequence* in a topological space (S, \mathcal{O}) is a continuous function $f : \alpha \rightarrow S$, where α is an ordinal. The topology on ordinals α is defined as follows: for every $x < \alpha$ the sets $\{y \mid y < x\}$ and $\{y \mid x < y\}$ are open and form an *open base*, i.e. all other open sets arise through finite intersections and arbitrary unions of these sets. The sequence f is called *open* if α is a limit ordinal, otherwise it is *closed*.

An open sequence $f : \alpha \rightarrow S$ *diverges* iff it cannot be extended to a sequence $f' : \alpha + 1 \rightarrow S$ where $\forall \gamma < \alpha. f'(\gamma) = f(\gamma)$. Otherwise it *converges*. For example, if S is the set $\{0, 1\}$ and $f : \omega \rightarrow S$ is given as $f(n) = n \bmod 2$ then f diverges for the discrete topology on S ($\mathcal{O} = \wp(S)$), but it converges on the other 3 topologies one can define on S .

A *subsequence* of f is a strictly monotonic and continuous function $g : \beta \rightarrow \alpha$. g is called *cofinal* [7] iff $\forall \gamma. \gamma < \alpha \Rightarrow \exists \zeta. \zeta < \beta \wedge g(\zeta) \geq \gamma$. We say that a subsequence g converges iff the sequence $f \circ g$ does. Clearly, an open sequence converges iff all its subsequences do. Note: (i) cofinal subsequences compose: if g is a cofinal subsequence of f and h a cofinal subsequence of g then $g \circ h$ is a cofinal subsequence of f ; (ii) subsequences are *strictly normal* functions on ordinals [15], i.e. their application distributes over suprema.

An open sequence $f : \alpha \rightarrow S$ *converges to* c iff it stays eventually within any neighbourhood of c . Notice that this notion coincides with a neighbourhood-based definition of being “continuous at α ” which implies that if f converges to c then it converges, and dually if it converges it must converge to some value.

A *metric space* is a pair (M, d) where M is a set and $d : M \times M \rightarrow \mathcal{R}$ is a distance function with the properties $d(x, y) = 0 \iff x = y$ and $d(x, y) + d(x, z) \geq d(y, z)$ which implies that d is commutative and returns only positive values [4]. Metric spaces are topological spaces: $A \subseteq M$ is open iff $\forall x \in A. \exists \epsilon > 0. \forall y \in M. d(x, y) < \epsilon \Rightarrow y \in A$. Metric spaces are Hausdorff. A special case of metric spaces are *ultra-metric spaces* which have a stronger triangle inequality: $\max(d(x, y), d(x, z)) \geq d(y, z)$.

The *diameter* of a subset $A \subseteq M$ is the supremum of all distances between elements of A . The *diameter* of an open sequence $f : \alpha \rightarrow M$ is the smallest $\epsilon \geq 0$ such that $\forall \beta < \alpha. \exists \gamma_1, \gamma_2 \geq \beta. d(f(\gamma_1), f(\gamma_2)) \geq \epsilon$.

A *Cauchy-sequence* over (M, d) is a function $f : \omega \rightarrow M$ such that $\forall \epsilon > 0. \exists k. \forall m, n \geq k. d(f(m), f(n)) < \epsilon$. Thus Cauchy-sequences are precisely those with a diameter of 0. The metric space is called *complete* iff all its Cauchy-sequences converge; each metric space M has a completion M^\bullet , which is unique up to isomorphism. In infinitary term rewriting, metric completion is used to generate infinite terms from finite terms. A function $f : M \rightarrow N$ between metric spaces (M, d_M) and (N, d_N) is called *uniformly continuous* iff $\forall \epsilon > 0. \exists \delta > 0. \forall x \in M. \forall y \in M. d_M(x, y) < \delta \Rightarrow d_N(f(x), f(y)) < \epsilon$. A uniformly continuous function $f : M \rightarrow N$ has a unique continuous extension to $M^\bullet \rightarrow N^\bullet$. Continuous functions preserve converging sequences, uniformly continuous functions also preserve Cauchy-sequences. A special case are *non-expansive* functions, characterised by the property $d_N(f(x), f(y)) \leq d_M(x, y)$. If even $d_N(f(x), f(y)) = d_M(x, y)$ then f is an *isometric embedding*. Isometric embeddings both preserve and reflect Cauchy-sequences, and — if the spaces are complete — converging sequences.

3. REDUCTION SYSTEMS AND CONVERGENCE

A *topological abstract reduction system* (short: TARS) is a structure $(S, \mathcal{O}, \rightarrow)$ such that (S, \mathcal{O}) is a topological space and (S, \rightarrow) an abstract reduction system, i.e. \rightarrow is a binary relation on S . In the world of infinitary terms the underlying topological spaces are typically ultra-metric spaces, see [9, Proposition 3] and [20, Section 4.3]. The reason for the added generality of this section is that (i) the notion of convergence is a topological concept, (ii) certain convergence results about abstract reduction systems can be proven at this more general level, and (iii) not all forms of infinitary rewriting derive from metric spaces, e.g. for infinitary λ -calculus Böhm-Trees relate to the Scott-topology [3].

A reduction sequence in a TARS $(S, \mathcal{O}, \rightarrow)$ is a sequence $f : \alpha \rightarrow S$ such that $\forall \beta. \beta + 1 < \alpha \Rightarrow f(\beta) \rightarrow f(\beta + 1)$. Recall that being a sequence entails the continuity of f .

For infinitary term rewriting we are specifically interested in TARS where the underlying topological space is a metric space. This gives rise to the notion of a *metric abstract reduction systems* (short: MARS) (M, d, \rightarrow) where (M, d) is a metric space, and \rightarrow a relation on M .

A TARS is called *convergent* iff all its open reduction sequences converge. Otherwise we call it *divergent*.

Proposition 3.1. *If the open sequence $f : \alpha \rightarrow S$ diverges then for every $s \in S$ there is a neighbourhood N_s , an ordinal β_s and a cofinal subsequence $w_s : \beta_s \rightarrow \alpha$ of f such that $\forall \gamma < \beta_s. f(w_s(\gamma)) \notin N_s$.*

Proof. Otherwise, there would be an s (possibly many) for which these premises are not true. In that case f converges to s . □

For topological spaces that are not Hausdorff it is possible that sequences converge to more than one value. For example, consider ω with the cofinite topology: a subset $A \subseteq \omega$ is open iff either A is empty or $\omega \setminus A$ is finite. In that topology the identity function on ω converges to all numbers, because the sequence $1, 2, 3, 4, \dots$ will eventually stay within any cofinite set of natural numbers.

Weak reduction \rightarrow_w of a TARS is defined as follows: $t \rightarrow_w u$ iff there exists a closed reduction sequence $f : \alpha + 1 \rightarrow S$ with $f(0) = t$ and $f(\alpha) = u$. Clearly, the relation \rightarrow_w is transitive. We also write \rightarrow_{ww} for the weak reduction of the TARS $(S, \mathcal{O}, \rightarrow_w)$. Moreover, we call reduction sequences of the TARS $(S, \mathcal{O}, \rightarrow_w)$ *weak reduction sequences* of the original TARS $(S, \mathcal{O}, \rightarrow)$. These are the earlier mentioned “reduction sequences with gaps”.

Proposition 3.2. *Let $(S, \mathcal{O}, \rightarrow)$ be a convergent TARS, where (S, \mathcal{O}) is Hausdorff. If $t \rightarrow_{ww} u$ then $t \rightarrow_w u$.*

Proof. We can prove the result by induction on the indexing ordinal α of the witnessing sequence $f : \alpha + 1 \rightarrow S$ for $t \rightarrow_{ww} u$. The base case $\alpha = 0$ is trivial and if α is a successor ordinal the result follows by the induction hypothesis and transitivity of \rightarrow_w .

Otherwise, α is a limit ordinal. For any $\beta < \alpha$ we have $t \rightarrow_{ww} f(\beta)$ and thus by the induction hypothesis $t \rightarrow_w f(\beta)$. This reduction has an associated indexing function $g_\beta : \gamma_\beta \rightarrow S$. The increasing sequence of ordinals γ_β (with β approaching α) must converge to some ordinal γ [17, page 290]. W.l.o.g. we can assume that this is a limit ordinal as otherwise almost all reductions in the image of f would be empty. Moreover the functions g_β agree on their common domain. In their limit they thus extend to a function $g : \gamma \rightarrow S$ which is an open reduction sequence. As the original TARS is convergent g can be extended to a closed sequence $g' : \gamma + 1 \rightarrow S$. Since both $f(\alpha)$ and $g'(\gamma)$ are limits of the sequence $f(\beta)$ with β approaching α we must have $g'(\gamma) = f(\alpha)$ and the result follows. □

It is easiest to explain the meaning of Proposition 3.2 by showing how the property fails if we drop the condition that the TARS is convergent. Consider the following example, given as an infinitary string rewriting system:

Example 3.3. $\{BE \rightarrow CSE, AC \rightarrow AB, BS \rightarrow SB, SC \rightarrow CS\}$

In this system we have for any n :

$$ABS^n E \rightarrow^* AS^n BE \rightarrow AS^n CSE \rightarrow^* ACS^{n+1} E \rightarrow ABS^{n+1} E$$

Thus $ABS^n E \rightarrow^* ABS^{n+1} E$ and also $ABS^n E \twoheadrightarrow_w ABS^{n+1} E$. Starting from $n = 0$ we get $ABE \twoheadrightarrow_{ww} ABS^\infty$ as the result is the limit of all $ABS^n E$. However, we do not have $ABE \twoheadrightarrow_w ABS^\infty$ as intermediate reduction results in the sequences for $ABS^n E \twoheadrightarrow_w ABS^{n+1} E$ will change the B to an S . The example does not contradict the proposition as the system is not converging: we have $ABS^n E \twoheadrightarrow_w AS^n BE \twoheadrightarrow_w ABS^{n+1} E \twoheadrightarrow_w \dots$

The requirement that the topological space is Hausdorff is not merely used in (the last part of) the proof, the property can indeed fail for non-Hausdorff-spaces.

Example 3.4. We define a TARS E with carrier set ω and the following sets as an open base:

- $S_n = \{m \mid m \geq n\}$, for any $n \in \omega$
- the set of all even numbers

Its relation \rightarrow is given as $n \rightarrow n + 1$, for all n .

For the topology of E every element has a smallest neighbourhood, which is subset of all its neighbourhoods. For an odd number n it is the set S_n , for an even number n it is S_n intersected with the set of all even numbers.

Since $S_1 = \omega \setminus \{0\}$, sequences converge to 1, unless they have a cofinal subsequence that is constantly 0. In that case the sequence might converge to 0, but only if it does not contain a cofinal subsequence in which all numbers are odd. As a whole the TARS is convergent as all its reduction sequences converge to 1.

But the proposition about the weak reduction relation does not hold here: we have $2 \twoheadrightarrow_{ww} 0$ but not $2 \twoheadrightarrow_w 0$. The reason is that the reduction sequence $2 \rightarrow 3 \rightarrow 4 \rightarrow \dots$ converges to all odd numbers, but to no even numbers. By omitting every other element the sequence gives rise to the weak reduction sequence $2 \twoheadrightarrow_w 4 \twoheadrightarrow_w 6 \twoheadrightarrow_w \dots$ which converges to all numbers, including 0; but there is no reduction sequence starting from 2 that would converge to 0.

Aside: it is possible that \twoheadrightarrow_w and \twoheadrightarrow_{ww} coincide without $(S, \mathcal{O}, \rightarrow)$ being convergent — e.g. we can choose $\rightarrow = S \times S$.

Proposition 3.5. *Let $(S, \mathcal{O}, \rightarrow)$ be a convergent TARS where (S, \mathcal{O}) is Hausdorff. Then $(S, \mathcal{O}, \twoheadrightarrow_w)$ is also convergent.*

Proof. Using the same construction as in the proof of Proposition 3.2 we can expand an open \twoheadrightarrow_w -reduction into an open \rightarrow -reduction. The convergence assumption gives us a limit to the latter which must also be a limit to the former. \square

As the proof is based on Proposition 3.2 we still require that the topology is Hausdorff, and the property can indeed fail for non-Hausdorff spaces, as Example 3.4 provides again a counterexample: we can define a weak reduction sequence of length ω as follows: $0 \rightarrow 1 \rightarrow 2 \twoheadrightarrow_w 4 \twoheadrightarrow_w 6 \twoheadrightarrow_w \dots$ which converges to 0. By iterating this reduction we can construct a weak open reduction sequence $f : \omega^2 \rightarrow \omega$ such that for all limit ordinals $\lambda < \omega^2$ we have

$f(\lambda) = 0 \wedge f(\lambda + 1) = 1 \wedge f(\lambda + n + 2) = 2n + 2$. This fails to converge under the topology of E , as both $g : \omega \rightarrow \omega^2, g(n) = \omega \cdot n$ and $h : \omega \rightarrow \omega^2, h(n) = \omega \cdot n + 1$ are cofinal; $f \circ g$ and $f \circ h$ are constantly 0 and 1, respectively, and thus miss the smallest neighbourhoods of all other numbers.

Corollary 3.6. *A TARS is convergent if and only if all its weak reduction sequences converge, provided its topology is Hausdorff.*

Proof. If the TARS is not convergent then the result follows because any reduction sequence is also a weak reduction sequence. Otherwise Proposition 3.5 applies. \square

Given a TARS $(S, \mathcal{O}, \rightarrow)$ and a sequence $f : \alpha \rightarrow S$, the predicate $\text{Foc}_{f,\alpha}$ on α is defined as follows:

$$\text{Foc}_{f,\alpha}(\beta) \iff \forall \gamma. \alpha > \gamma \geq \beta \Rightarrow \exists \zeta. \alpha > \zeta \wedge \forall \kappa. \alpha > \kappa \geq \zeta \Rightarrow f(\gamma) \twoheadrightarrow_w f(\kappa)$$

We have: (i) if $\text{Foc}_{f,\alpha}(\beta)$ then $\text{Foc}_{f,\alpha}(\gamma)$ for all $\gamma \geq \beta$; (ii) if $\text{Foc}_{f,\alpha}(\beta)$ then $\exists \gamma > \beta. f(\beta) \twoheadrightarrow_w f(\gamma)$. The sequence f is called *focussed* iff $\exists \beta. \text{Foc}_{f,\alpha}(\beta)$.

In words: a sequence is focussed iff every of its elements sufficiently close to α weakly reduces to all its elements sufficiently close to α . In particular, all weak reduction sequences are focussed, but not vice versa.

Of interest are specifically *open* focussed sequences. Closed sequences ($\alpha = \alpha' + 1$) are trivially focussed: to prove $\exists \beta. \text{Foc}_{f,\alpha}(\beta)$ we can choose $\beta = \alpha'$ which determines $\gamma = \alpha'$, and ζ can also be chosen as α' which in turn determines $\kappa = \alpha'$, and $f(\alpha') \twoheadrightarrow_w f(\alpha')$ holds by reflexivity of \twoheadrightarrow_w .

Proposition 3.7. *Every cofinal subsequence g of a focussed sequence f has itself a cofinal subsequence h such that $f \circ g \circ h$ is a cofinal weak reduction sequence.*

Proof. Proposition 4 in [10] claimed this for MARS, but the proof works unchanged for TARS, with no constraints on the topology. \square

Proposition 3.8. *A TARS converges iff all its focussed sequences converge, provided the topology is Hausdorff.*

Proof. Since all weak reduction sequences are focussed the \Leftarrow implication follows from Corollary 3.6. Now suppose $f : \alpha \rightarrow S$ is a divergent focussed sequence; from it, we need to construct a divergent reduction sequence, and by Corollary 3.6 it suffices if it is a weak reduction sequence.

Take any element $c \in S$. Proposition 3.1 gives us a cofinal subsequence $w_c : \beta_c \rightarrow \alpha$ of f witnessing the lack of convergence of f to c . We can apply Proposition 3.7 to f and w_c , which gives us a cofinal subsequence $h_c : \gamma_c \rightarrow \beta_c$ such that $g_c = f \circ w_c \circ h_c$ is a weak reduction sequence.

If g_c is divergent we are done. Otherwise it converges to a limit e . Again, we have a cofinal subsequence w_e witnessing that f does not converge to e . From this we can construct another cofinal weak reduction sequence h that alternates between elements from g_c and $f \circ w_e$. This cannot converge to e because of the cofinal overlap with $f \circ w_e$, but if h converged to any other e' then because of the cofinal overlap with g_c all neighbourhoods of e and e' intersect, contradicting the Hausdorff proviso. \square

The argument of the proof is essentially the same as in [10], but reworded to avoid reference to distances. In [10] it was said that it was “not clear” if the proposition could be generalised to non-Hausdorff spaces. It actually cannot, as our earlier counterexample to a generalisation of Proposition 3.5 would also refute a generalisation of Proposition 3.8.

4. FINITE TERMS AND INFINITARY TERMS

A *signature* is a pair $\Sigma = (\mathcal{F}, \#)$ where \mathcal{F} is a set (of function symbols) and $\# : \mathcal{F} \rightarrow \mathcal{N}$ is the function assigning each symbol its arity. We assume an infinite set Var of *variables*, disjoint from \mathcal{F} . The set of finite terms over Σ is called $Ter(\Sigma)$ and it is defined to be the smallest set such that (i) $Var \subset Ter(\Sigma)$ and (ii) $F(t_1, \dots, t_n) \in Ter(\Sigma)$ whenever $F \in \mathcal{F} \wedge \#(F) = n \wedge \{t_1, \dots, t_n\} \subset Ter(\Sigma)$. The *root symbol* of a term $F(t_1, \dots, t_n)$ is F , the root symbol of a variable x is x .

We write \mathcal{N}^* for the free monoid over the natural numbers (i.e. finite words), with neutral element λ and infix \cdot as monoid multiplication. The set of positions $\text{Pos}(t)$ of a finite term t is the smallest subset of \mathcal{N}^* such that: (i) $\lambda \in \text{Pos}(t)$, (ii) $\{i \cdot q \mid 1 \leq i \leq n, q \in \text{Pos}(t_i)\} \subset \text{Pos}(F(t_1, \dots, t_n))$. Positions are used for selecting or replacing subterms. For selecting the subterm at position p of a term t we use the notation t/p . We extend the domain of this function by setting $t/p = x$ if $p \notin \text{Pos}(t)$ which allows to extend this notation pointwise to sequences of terms: $(f/p)(\alpha) = f(\alpha)/p$. For replacing the subterm at position p of a term t by u we use the notation $t[u]_p$; if $p \notin \text{Pos}(t)$ then we set $t[u]_p = t$. Positions are partially ordered by the *prefix-ordering*: $p \preceq p \cdot q$.

A Σ -*algebra* is a set A together with functions $F_A : A^n \rightarrow A$ for every $F \in \mathcal{F}$ with $\#(F) = n$. A *valuation* into A is a function $\rho : Var \rightarrow A$. Any Σ -algebra A determines an interpretation function $\llbracket _ \rrbracket_A : Ter(\Sigma) \times (Var \rightarrow A) \rightarrow A$ as follows:

$$\begin{aligned} \llbracket x \rrbracket_A^\rho &= \rho(x), & \text{if } x \in Var \\ \llbracket F(t_1, \dots, t_n) \rrbracket_A^\rho &= F_A(\llbracket t_1 \rrbracket_A^\rho, \dots, \llbracket t_n \rrbracket_A^\rho) \end{aligned}$$

Given two Σ -algebras A and B , a Σ -algebra *homomorphism* from A to B is a function $h : A \rightarrow B$ such that $h(F_A(a_1, \dots, a_n)) = F_B(h(a_1), \dots, h(a_n))$. If $h : A \rightarrow B$ is a homomorphism then $h(\llbracket t \rrbracket_A^\rho) = \llbracket t \rrbracket_B^{h \circ \rho}$.

A *context* is a pair of a term t and a position $p \in \text{Pos}(t)$. If $C[\]$ is the context (t, p) and u a term the context application $C[u]$ is defined as $t[u]_p$. A substitution is a function $\theta : Var \rightarrow Ter(\Sigma)$ which is uniquely extended to a function $\bar{\theta} : Ter(\Sigma) \rightarrow Ter(\Sigma)$ by requiring that $\bar{\theta}$ is a Σ -algebra homomorphism.

An *ultra-metric map* (short: umm) is an n -ary function $f : [0, 1]^n \rightarrow [0, 1]$ such that (i) $f(x_1, \dots, x_n) = 0 \iff x_1 = 0 \wedge \dots \wedge x_n = 0$ and (ii) it distributes over max: $f(\max(x_1, y_1), \dots, \max(x_n, y_n)) = \max(f(x_1, \dots, x_n), f(y_1, \dots, y_n))$. Each n -ary umm f can be expressed in the form $f(x_1, \dots, x_n) = \max_{1 \leq i \leq n} \tilde{f}_i(x_i)$, where the functions $\tilde{f}_i : [0, 1] \rightarrow [0, 1]$, called the *components* of f , are also umms. A *term metric* for a signature Σ is a Σ -algebra m with carrier set $[0, 1]$, such that for each $F \in \mathcal{F}$ the function F_m is a umm of the same arity. A term metric is called *continuous* iff all its ultra-metric maps are continuous. In this paper, we shall only consider continuous term metrics.

For every term metric m , the distance function $d_m : Ter(\Sigma) \times Ter(\Sigma) \rightarrow [0, 1]$ is defined as follows: $d_m(t, t) = 0$, $d_m(t, u) = 1$ if t and u have different root symbols, and $d_m(F(t_1, \dots, t_n), F(u_1, \dots, u_n)) = F_m(d_m(t_1, u_1), \dots, d_m(t_n, u_n))$. Some fundamental results about term metrics from [9] are: d_m is an ultra-metric; the topology induced by d_m is

discrete. Moreover, the set of *infinitary terms* over Σ and m is called $Ter^m(\Sigma)$ and defined as the metric completion of the metric space $(Ter(\Sigma), d_m)$. $Ter^m(\Sigma)$ is a Σ -algebra if m is continuous.

Each position p in $\text{Pos}(t)$ has a unique umm $(t, p)_m$ associated with a term metric m :

$$\begin{aligned} (t, \lambda)_m &= id \\ (F(t_1, \dots, t_n), i \cdot q)_m &= \tilde{F}_{m,i} \circ (t_i, q)_m \end{aligned}$$

This associated umm gives us the distance function for position p , in the sense that: $d_m(t[u]_p, t[s]_p) = (t, p)_m(d_m(u, s))$. As the definition only depends on a finite part of t it is even defined for infinite terms t which are not in $Ter^m(\Sigma)$.

The notions of contexts, substitutions, positions and the operations on them lift canonically from $Ter(\Sigma)$ to $Ter^m(\Sigma)$, at least for continuous term metrics, see [9].

Proposition 4.1. *The subterm replacement operation $t[u]_p$ is non-expansive in t , for any term metric.*

Proof. By induction on the length of p . Base case: $d_m(t[u]_\lambda, t'[u]_\lambda) = d_m(u, u) = 0 \leq d_m(t, t')$. Otherwise $p = i \cdot q$. If t and t' have different root symbols then $d_m(t, t') = 1 \geq d_m(t[u]_p, t'[u]_p)$. Finally, let $t = F(t_1, \dots, t_n)$ and $t' = F(t'_1, \dots, t'_n)$. Then

$$d_m(t, t') = \max_{1 \leq j \leq n} \tilde{F}_{m,j}(d_m(t_j, t'_j)) = \max(\delta, \tilde{F}_{m,i}(d_m(t_i, t'_i)))$$

where $\delta = \max_{1 \leq j \leq n, j \neq i} \tilde{F}_{m,j}(d_m(t_j, t'_j))$. Moreover,

$$d_m(t[u]_{i \cdot q}, t'[u]_{i \cdot q}) = \max(\delta, \tilde{F}_{m,i}(d_m(t_i[u]_q, t'_i[u]_q)))$$

Induction hypothesis on q gives us: $d_m(t_i, t'_i) \geq d_m(t_i[u]_q, t'_i[u]_q)$, so the result follows by monotonicity of the functions $\tilde{F}_{m,i}$ (as they are umms) and \max . \square

Aside: subterm replacement $t[u]_p$ is non-expansive in u if and only if the term metric m is non-expansive, i.e. if all functions F_m for $F \in \mathcal{F}$ are non-expansive.

4.1. Characterising Convergence in Term Metric m . The ϵ -positions of a term t w.r.t. to term metric m , the set $\text{Pos}_\epsilon^m(t)$, is defined as $\{p \in \text{Pos}(t) \mid (t, p)_m(1) \geq \epsilon\}$. Clearly, $\text{Pos}_\epsilon^m(t)$ is prefix-closed: if $q \cdot w = p$ then $(t, q)_m(1) \geq (t, q)_m((t/q, w)_m(1)) = (t, q \cdot w)_m(1) \geq \epsilon$.

The properties of $\text{Pos}_\epsilon^m(t)$ can be used to characterise whether an infinite term (from $Ter^\infty(\Sigma)$) exists (in $Ter^m(\Sigma)$).

Proposition 4.2. *If $t \in Ter^m(\Sigma)$ then $\text{Pos}_\epsilon^m(t)$ is always finite.*

Proof. Let $t_\epsilon \in Ter(\Sigma)$ be a finite term such that $d_m(t, t_\epsilon) = \eta_1 < \epsilon$ — such finite terms always exist, by construction of the metric completion.

Let $p \in \text{Pos}_\epsilon^m(t)$ and $y \in \text{Var}$ be fresh. Then $d_m(t, t[y]_p) = (t, p)_m(1) \geq \epsilon$, by definition of $\text{Pos}_\epsilon^m(t)$. We also have $d_m(t[y]_p, t_\epsilon[y]_p) = \eta_2 \leq \eta_1 < \epsilon$, because subterm replacement is non-expansive. Thus, using the ultra-metric properties of d_m , we know $\epsilon \leq (t, p)_m(1) = d_m(t, t[y]_p) \leq \max(\eta_1, \max(d_m(t_\epsilon, t_\epsilon[y]_p), \eta_2))$. As both η_i are smaller than ϵ this is equivalent to $d_m(t, t[y]_p) \leq d_m(t_\epsilon, t_\epsilon[y]_p)$, and the dual argument shows that these distances must be equal. As that value is non-zero, p must be in $\text{Pos}(t_\epsilon)$. This shows $\text{Pos}_\epsilon^m(t) \subseteq \text{Pos}(t_\epsilon)$ and as the latter is finite, so is the former. \square

Thus, for any infinitary term under any term metric there are only finitely many subterm positions which — if changed — could change the overall term by ϵ or more. This observation plays a role when we reason about divergence.

An immediate consequence of Proposition 4.2 is that if $\text{Pos}_\epsilon^m(t)$ is not finite (for any $\epsilon > 0$) then $t \notin \text{Ter}^m(\Sigma)$. This gives us a criterion to rule out certain infinite terms. In fact, the criterion is complete:

Proposition 4.3. *Let $t \in \text{Ter}^\infty(\Sigma)$. If $\text{Pos}_\epsilon^m(t)$ is finite, for all ϵ , then $t \in \text{Ter}^m(\Sigma)$.*

Proof. We need to create a Cauchy-sequence (w.r.t. d_m) of finite terms converging to t . To make the n -th term t_n of the sequence to be within distance of $1/n$ of the limit t we obtain it by replacing all subterms of t in minimal (w.r.t. \preceq) positions not in $\text{Pos}_\epsilon^m(t)$ (for $\epsilon = 1/n$) by the variable x . As $\text{Pos}_\epsilon^m(t)$ is finite its elements have a maximal length k and the longest position in t_n is (at most) of length $k + 1$, and thus by König's Lemma t_n must be finite too. \square

An important concept when studying convergence is the equivalence relation that two terms have “the same function symbols up to position p ”. Formally, this is defined by induction on the length of the positions as follows:

$$t \stackrel{\lambda}{\sim} u \\ F(t_1, \dots, t_n) \stackrel{i\text{-}q}{\sim} F(u_1, \dots, u_n) \iff t_i \stackrel{q}{\sim} u_i$$

Clearly, if $a \stackrel{p}{\sim} b$ then $(a, p)_m = (b, p)_m$; see also [9].

Traditionally, infinitary term rewriting has focussed on one particular distance function d_∞ . The same distance function arises through the definition of a term metric: we define the term metric ∞ in which all n -ary ultra-metric maps are given as $\infty_n(x_1, \dots, x_n) = \frac{1}{2} \cdot \max(x_1, \dots, x_n)$. It plays a fundamental role as $\text{Ter}^\infty(\Sigma)$ includes *all* infinite terms:

Proposition 4.4. *Any Cauchy-sequence in $(\text{Ter}(\Sigma), d_m)$ is also a Cauchy-sequence in $(\text{Ter}(\Sigma), d_\infty)$.*

Proof. Let $f : \omega \rightarrow \text{Ter}(\Sigma)$ be that sequence.

Assume it was not Cauchy under metric d_∞ . Then there is a maximal (under the prefix ordering) position p such that f/p changes its root symbol infinitely often. By maximality, the symbols above p are fixed for all sufficiently large k , i.e. $f(k) \stackrel{p}{\sim} f(k+1)$, and infinitely often have $f(k)/p$ and $f(k+1)/p$ different root symbols.

This implies (see proposition 6 in [9]) that there is a fixed umm C_m such that for all sufficiently large k

$$d_m(f(k), f(k+1)) \geq C_m(d_m(f(k)/p, f(k+1)/p)).$$

Whenever we have a change of function symbol at position p the expression on the right becomes the fixed value $C_m(1)$. As this happens infinitely often we have a contradiction to the pre-condition that f is Cauchy under metric d_m . Therefore the assumption was wrong and f is Cauchy under d_∞ as well. \square

Corollary 4.5. *For any term metric m , there is a unique continuous extension of the identity function from finite terms to $\text{Ter}^m(\Sigma) \rightarrow \text{Ter}^\infty(\Sigma)$.*

Proof. The identity is continuous on finite terms between any two term metrics, because the topologies on finite terms are discrete for any term metric (proposition 4 in [9]). Proposition

4.4 allows to extend the function to infinite terms with domain $Ter^m(\Sigma)$, and such extensions to metric completions are always unique when they exist. \square

Proposition 4.4 shows that $Ter^\infty(\Sigma)$ “contains all infinite terms”, because any infinite term arises as the limit of a Cauchy-sequence under some metric d_m , and the corollary indicates that the construction is continuous. Thus we can view any element of $Ter^m(\Sigma)$ also as an element of $Ter^\infty(\Sigma)$. It also means that any sequence in $Ter^m(\Sigma)$ that fails to converge under d_∞ will necessarily also fail to converge under d_m .

Aside: for finite signatures there is a simpler argument to establish Corollary 4.5: the identity function is uniformly continuous (from d_m to d_∞). However, this is no longer true for infinite signatures, e.g. if a signature contained infinitely many unary function symbols $F\langle n \rangle$ with umms $F\langle n \rangle_m(x) = \frac{x}{n}$.

Proposition 4.6. *Let f be a sequence over $Ter^m(\Sigma)$ such that f is converging under metric d_∞ to $t \in Ter^m(\Sigma)$. Then f is also converging to t under metric d_m .*

Proof. As $t \in Ter^m(\Sigma)$, the set $\text{Pos}_\epsilon^m(t)$ is finite for any ϵ . As f is converging to t under d_∞ , there must be a γ_0 such that $\forall \gamma \geq \gamma_0. \forall p \in \text{Pos}_\epsilon^m(t). t \stackrel{p}{\approx} f(\gamma)$. This implies that t and all $f(\gamma)$ are within ϵ -distance, w.r.t. to metric d_m . Thus f stays eventually within the ϵ -ball of t , and f must converge to t as this holds for any ϵ . \square

Therefore, if a sequence of finite terms f converges under metric d_∞ to a limit $t \in Ter^\infty(\Sigma)$, the question whether f converges under d_m reduces to the problem whether $t \in Ter^m(\Sigma)$ or not. If $t \in Ter^m(\Sigma)$, Proposition 4.6 applies and so f converges; if $t \notin Ter^m(\Sigma)$ then f cannot converge under d_m , because if it did converge to $t' \neq t$ then f would converge under d_∞ to both t and t' — which is impossible in a Hausdorff space.

Let $t \in Ter^\infty(\Sigma)$. An *infinite path* of t is a strictly monotonic function $f : \omega \rightarrow \text{Pos}(t)$, i.e. $n < k \Rightarrow f(n) \prec f(k)$. The path f is called *full* iff the length of $f(n)$ is always n . We say that f converges w.r.t. term metric m iff the sequence of umms $(t, f(n))_m$ converges to the constant 0 function.

Lemma 4.7. *Let f be an infinite path of t . Then f converges w.r.t. m iff the sequence $(t, f(n))_m(1)$ converges to 0.*

Proof. If $(t, f(n))_m$ converges to the constant 0 function then it will converge to 0 when applied to argument 1. If $(t, f(n))_m(1)$ converges to 0 then by monotonicity $(t, f(n))_m(\epsilon)$ must converge to 0 at (at least) the same rate, and therefore $(t, f(n))_m$ converges to the constant 0 function. \square

Proposition 4.8. *Let $t \in Ter^\infty(\Sigma)$. Then $t \in Ter^m(\Sigma)$ iff all infinite paths of t converge w.r.t. term metric m .*

Proof. Suppose path f does not converge. Then by 4.7 the sequence $(t, f(n))_m(1)$ does not converge to 0. Thus, $(t, f(n))_m(1) \geq \epsilon > 0$ for infinitely many n . As each of these is in $\text{Pos}_\epsilon^m(t)$, this set is infinite and $t \notin Ter^m(\Sigma)$ follows from Proposition 4.2.

Now assume $t \notin Ter^m(\Sigma)$. Then by proposition 4.3 $\text{Pos}_\epsilon^m(t)$ must be infinite for some ϵ . We can iteratively construct a full non-converging infinite path through $\text{Pos}_\epsilon^m(t)$ by the pigeon hole principle: set $f(0) = \lambda$; as the symbol at $t/f(n)$ has finite arity, there must be at least one i such that infinitely many positions in $\text{Pos}_\epsilon^m(t)$ have $f(n) \cdot i$ as prefix, and we can set $f(n+1) = f(n) \cdot i$. As all values in the range of f are in $\text{Pos}_\epsilon^m(t)$, the sequence $(t, f(n))_m(1)$ stays above ϵ and is therefore not converging to 0. \square

5. INFINITARY TERM REWRITING

An *infinitary term rewriting system* (iTRS) is a triple (Σ, m, R) such that (i) Σ is a signature, (ii) m is a term metric for Σ , and (iii) $R \subseteq \text{Ter}^m(\Sigma) \times \text{Ter}^m(\Sigma)$ is a set of rewrite rules such that for all $(l, r) \in R$ we have $l \notin \text{Var}$ and all variables in r also occur in l . We write $l \rightarrow r$ if $(l, r) \in R$.

Note: this definition permits infinite terms on both left-hand and right-hand sides of rules, as well as repeated variables in the left-hand side. For better continuity properties of the rewrite relation one would typically require in addition that each l is a finite linear term. Moreover, [9] required all rules to be depth-preserving (see below). However, for the purposes of this paper these restrictions are not necessary.

A relation S on $\text{Ter}^m(\Sigma)$ is called *substitutive* iff $s S t$ implies $\bar{\theta}(s) S \bar{\theta}(t)$, for any substitution θ . It is called *compatible* iff $s S t$ implies $C[s] S C[t]$ for any context $C[\]$.

The single-step rewrite relation \rightarrow_R of an iTRS (Σ, m, R) is defined as the compatible and substitutive closure of R , i.e. $C[\bar{\theta}(l)] \rightarrow_R C[\bar{\theta}(r)]$ whenever $(l, r) \in R$, $C[\]$ is a context, and θ is a substitution.

Aside: there is a slight issue with the definition of “compatible”. For better continuity properties one could require in addition that a compatible relation should be reflexive on infinite terms. If t is an infinite term then a rewrite step $t[l]_p \rightarrow_R t[r]_p$ should vanish in the limit, as p goes arbitrarily deep. This is necessary to ensure that the relation \rightarrow_R is upper semi-continuous, and that it is closed. For the purposes of this paper the distinction is immaterial as the presence of reflexive rewrite steps does not affect the convergence relations.

We associate with an iTRS (Σ, m, R) the MARS $(\text{Ter}^m(\Sigma), d_m, \rightarrow_R)$; in particular, the iTRS is called convergent iff its associated MARS is.

In the following we will be using a generalisation of the notion of *depth-preservation* [12] to arbitrary term metrics. In term metric ∞ , the depth of a variable x in a term t is the length l of the shortest position p in t such that $t/p = x$. Equivalently, that information is also present in the distance $d_\infty(t, t[y/x]) = (t, p)_\infty(1)$ (for fresh y), because that distance is 2^{-l} .

For other term metrics that figure is insufficient to fully capture the notion of depth of a variable in a term, because it only tells us what happens with distances when terms with different roots replace the x . More generally, we are interested in the distances $d_m(t[u/x], t[s/x])$ for varying u and s ; the dependence on u and s can be replaced by a dependence on $d_m(u, s)$. Thus the generalised *variable depth* of x in t w.r.t. term metric m , $vdepth(x, m, t)$, is an ultra-metric map of type $m \rightarrow m$, given as:

$$vdepth(x, m, t)(y) = \llbracket t \rrbracket_m^{\rho(y)} \quad \text{where } \rho(y)(z) = \begin{cases} y & \text{if } z = x \\ 0 & \text{otherwise} \end{cases}$$

A rewrite step $t \rightarrow u$ preserves the depth of x iff $vdepth(x, m, t) \geq vdepth(x, m, u)$, in the pointwise order on $\text{Var} \rightarrow m$. Overall, a relation R on $\text{Ter}^m(\Sigma)$ is called *depth-preserving* iff $\forall t, u. t R u \Rightarrow \llbracket t \rrbracket \geq \llbracket u \rrbracket$ which means that all its rewrite steps preserve the depth of all variables. Here, the order \geq is the pointwise order on the function space, inherited (twice) from m .

An iTRS (for term metric ∞) is *strongly converging* iff for all open reduction sequences of length α no redex position p is reduced arbitrarily close to α . Equivalent to this [24] is the

notion of top-termination, i.e. that no reduction sequence contracts a root redex infinitely often.

It is possible to define the property without referring to redex positions though: the *indirected version* of an iTRSs (Σ, m, R) is defined as follows. Its signature is Σ extended by a fresh unary function symbol I ; for each rule $l \rightarrow r$ in R it has a rule $l \rightarrow I(r)$; in addition, it has the rule $I(x) \rightarrow x$. We extend the term metric m by setting $I_m(x) = x$.

Proposition 5.1. *An iTRS over term metric ∞ is strongly converging iff its indirected version is converging.*

Proof. Any reduction sequence f of the original iTRS can be mapped into a reduction sequence f' of the indirected version by splitting the step into two parts, where the second part removes the just introduced I -symbol. If f is not top-terminating then f' fails to converge, as it would have infinitely many changes of root-symbol.

Dually, a reduction sequence g of the indirected version can be mapped into a (shorter) sequence g^+ of the original system, by removing all I -symbols. However, g can only fail to converge by an infinitely often occurring function symbol change at some minimal position p . That means though that g must apply rules other than $I(x) \rightarrow x$ at position p infinitely often, and thus g^+ would apply those rules at p or higher infinitely often which means that g^+ is not strongly converging. \square

The reason that this proposition is limited to term metric ∞ is simply that the notion of strong convergence is only *defined* for metric d_∞ . The idea behind strong convergence is that redex contractions must contribute less and less to the diameter of a sequence as redex positions go deeper. For other term metrics the movement of the redex positions alone does not suffice, similarly as modifications of the metric for the infinitary λ -calculus [11, page 709f] affect the notion of “root-active”.

For other term metrics than ∞ we *define* strong convergence as the convergence of the indirected iTRS.

Consequence: consider the term metric *id* with $F_{id}(x_1, \dots, x_n) = \max(x_1, \dots, x_n)$ for all function symbols F . This is a discrete metric and there are no infinite terms: $Ter^{id}(\Sigma) = Ter(\Sigma)$. Thus, convergence means that reduction sequences are eventually constant, and strong convergence coincides with strong normalisation.

This also shows that strong convergence is not in general a modular property, simply because strong normalisation is not. However, strong normalisation is modular in the absence of collapsing rules [16], so we will consider the corresponding constraint here.

6. EXAMPLES, COUNTEREXAMPLES AND NEAR COUNTEREXAMPLES

In this section we will be looking at examples that show why the modularity results for convergence come with side conditions. The examples here are using metric d_∞ ; to consider other metrics at all we would have to first explain what the disjoint union of two iTRSs with arbitrary metrics is, see the next section.

6.1. Collapsing Rules. Particularly problematic w.r.t. convergence are collapsing rules, i.e. rules that have a variable as their right-hand side. An iTRS immediately fails to be strongly convergent if a collapsing rule is present. The reason is this: any such rule takes

the form $C[x, \dots, x] \rightarrow x$, the equation $t = C[t, \dots, t]$ has a (unique) solution in $Ter^\infty(\Sigma)$, and $t \xrightarrow{\lambda} t$ gives an infinite reduction sequence with contractions at the root.

The weaker notion of convergence can coexist with collapsing rules, but barely: if $C[x] \rightarrow^* x$ holds for any other contexts $C[\]$ than those of the form $C[y] = F^n(y)$ (for varying n but fixed F) then such an iTRS is not convergent either. In particular, the presence of two rules $F(x) \rightarrow x$ and $G(x) \rightarrow x$ means that the terms $t = F(G(t))$, $u = G(F(u))$ reduce to each other and thus give rise to a non-converging reduction sequence. The example also shows that convergence (on its own) is not a modular property.

Convergence can also be broken if only one of the two systems has a collapsing rule.

Example 6.1. System R has rule $G(H(x)) \rightarrow G(x)$, system S has the rule $F(x) \rightarrow x$. In the combined system there is the term $t = F(H(t))$ and the non-convergent reduction $G(t) \rightarrow G(H(t)) \rightarrow G(t)$.

6.2. Building Infinite Results. Typically we can observe convergence (or its absence) at ω -indexed reduction sequences. In the presence of non-left-linear rules this can be rather different though.

Example 6.2. (from [24]) $\{E \rightarrow S(E), F \rightarrow S(F), G(x, x) \rightarrow G(E, F)\}$

This fails to converge, but this failure first occurs at ω^2 -indexed reduction sequences. We have $G(E, F) \rightarrow^* G(S^n(E), S^m(F))$ for any m, n and in the limit $G(E, F) \rightarrow_w G(S^\infty, S^\infty) \rightarrow G(E, F)$. This does not provide any convergence problems at ordinals of the form $\omega \cdot k$ though; however, when we attempt to extend a reduction sequence to ω^2 we find both $G(E, F)$ and $G(S^\infty, S^\infty)$ in the image of the reduction sequence at any of its neighbourhoods.

This is a common pattern for establishing non-convergence: any sequence of length α that begins and ends in the same value, but that visits a different value in between, gives rise to a non-converging sequence of length $\omega \cdot \alpha$ by simply repeating that sequence ω times. That construction works in any Hausdorff space.

Example 6.3. In Example 6.2 we had $E \rightarrow^* S^n(E)$ for any n , and $E \rightarrow_w S^\infty$ in the limit. However, one can have an iTRS with the former but not the latter property, e.g. if we replace the first rule by these four:

$$\{E \rightarrow Z, E \rightarrow H(E), H(Z) \rightarrow S(Z), H(S(x)) \rightarrow S(S(x))\}$$

Patterns like this complicate the reasoning about (weak) convergence, because in a combined system an outer layer may place in a position p different reducts of E , and thus the term in position p could converge at a limit to S^∞ , showing that a principal subterm at a limit may not be a reduct of an earlier principal subterm. This happens when we combine this system with the following:

$$\{J(K(x, y)) \rightarrow J(y)\}$$

Let $t = K(E, t)$ and $u = K(S^\infty, u)$: we have $t = K(E, K(E, \dots)) \rightarrow_w K(Z, K(S(Z), \dots))$, and thus $J(t) \rightarrow_w J(K(Z, K(S(Z), \dots)))$. Continuing the reduction from there by repeatedly applying the rule $J(K(x, y)) \rightarrow J(y)$ we have overall $J(t) \rightarrow_w J(u)$.

What is remarkable about this is that E is the only principal subterm of $J(t)$ and it does not reduce to S^∞ , the only principal subterm of $J(u)$. However, this second iTRS already

fails to converge: given the term $s = K(J(K(x, y)), s)$ we can construct the non-converging sequence $J(s) \rightarrow J(K(J(y)), s) \rightarrow J(s)$.

This construction is by no means arbitrary, as we shall see later: if a reduction sequence f in the combined system “substantially rearranges” principal subterms then one can construct a non-converging reduction sequence in the iTRS to which the root symbol belongs, where the role of principal subterms is played by some fixed terms $l \neq r$ with $l \rightarrow r$.

6.3. Beyond Ordinary iTRSs. What happens to convergence if we remove any constraints on rewrite rules, i.e. if we allow left-hand sides to be variables (expansive rules), and right-hand sides to contain extra variables the left-hand sides does not?

One of the main reasons these kind of rules are excluded in finite term rewriting is that either kind of rule immediately breaks the strong normalisation property. Another reason is more pragmatic: the presence of such rules requires a very different implementation.

Expansive rules immediately break strong convergence, but they may not break weak convergence. The rule $x \rightarrow F(x)$ on its own is weakly convergent, and if disjointly added to a strongly convergent iTRS the result would be weakly convergent. On the other hand, the presence of a second expansive rule $y \rightarrow G(y)$ breaks weak convergence: $x \rightarrow F(x) \rightarrow G(F(x)) \rightarrow F(G(F(x))) \rightarrow \dots$. Similar to collapsing rules, the presence of two different function symbols in an expansive rule breaks convergence: in that case the right-hand side must have the form $C[F(t_1, \dots, t_n)]_p$ where the root symbol of $C[\]$ is different from F . The reduction sequence that applies the rule alternatively in positions λ and p is non-converging, as the function symbol in position p changes every time.

Rules with extra variables may not only be convergent, they may even be strongly convergent. An example for that is the rule $F(x) \rightarrow G(y)$. This is not only convergent on its own, it preserves convergence if disjointly added to any convergent iTRS. A rule like $l \rightarrow y$, where the right-hand side *is* an extra variable, can only be convergent if l is a ground term, as otherwise two different substitution instances of l could rewrite to each other. If l is ground then $l \rightarrow y$ can only be convergent for (metrics homeomorphic to) d_∞ , as the rule can be used to gradually build terms of any infinite shape.

But that does not mean that extra variables are unproblematic: rules with extra variables always fail to be depth-preserving, for any term metric. They also break many proof techniques, as they enable the rank of terms to go up by rewriting.

7. THE DISJOINT UNION OF iTRSs

A *signature morphism* between signatures $(\mathcal{F}_1, \#_1)$ and $(\mathcal{F}_2, \#_2)$ is a function $\sigma : \mathcal{F}_1 \rightarrow \mathcal{F}_2$ that preserves arities. This gives rise to a translation between finite terms $\bar{\sigma}$ defined as $\bar{\sigma}(F(t_1, \dots, t_n)) = \sigma(F)(\bar{\sigma}(t_1), \dots, \bar{\sigma}(t_n))$, and $\bar{\sigma}(x) = x$ for $x \in \text{Var}$.

A *metric signature morphism* between (Σ_1, m_1) and (Σ_2, m_2) is a signature morphism $\sigma : \Sigma_1 \rightarrow \Sigma_2$ such that for all $F \in \mathcal{F}_1$ we have: $F_{m_1} = \sigma(F)_{m_2}$. For a metric signature morphism σ the function $\bar{\sigma}$ is a homeomorphic embedding as it preserves distances: $d_{m_2}(\bar{\sigma}(t), \bar{\sigma}(u)) = d_{m_1}(t, u)$, and therefore it extends to infinitary terms.

Notice that signature morphisms preserve the *term metric* m , not just the derived notion of the metric d_m . Different term metrics can give rise to identical topologies:

Example 7.1. If Σ comprises a single unary function symbol F then both term metrics $F_{id}(x) = x$ and $F_r(x) = \sqrt{x}$ give rise to identical metrics and therefore produce the same set

$Ter^{id}(\Sigma) = Ter^r(\Sigma) = Ter(\Sigma)$, with no infinite terms at all. But, combined with another unary symbol G with term metric $G_\infty(x) = x/2$ the two versions of F behave differently: the term $t = G(F(t))$ exists with F_{id} and does not with F_r .

An *iTRS morphism* from (Σ_1, m_1, R_1) to (Σ_2, m_2, R_2) is a metric signature morphism $\sigma : (\Sigma_1, m_1) \rightarrow (\Sigma_2, m_2)$ such that $\bar{\sigma}(R_1) \subseteq R_2$, where $\bar{\sigma}$ is canonically extended to rules and sets of rules.

Infinitary term rewriting systems and their morphisms form a cocomplete category. The *disjoint union* $R + S$ of two iTRSs R and S is defined to be the coproduct in that category. This contains all symbols from both systems (possibly renamed if there was a conflict) and all rules from both systems whilst preserving the umms for all function symbols. As distances are preserved by iTRS morphisms we shall use the metric d_m of the coproduct also for distances of the constituent iTRSs.

The preservation of the metric by the injection into the coproduct also settles the “reflected by” part of modularity properties:

Theorem 7.2. *The properties of convergence and strong convergence are reflected by the disjoint union of iTRSs.*

Proof. Convergence: assume $R + S$ is converging, let ι_1 be the coproduct injection from R to $R + S$. Any open reduction sequence f of R is mapped by ι_1 into an open reduction sequence of $R + S$ which by assumption is converging. As ι_1 is an isometric embedding (and as its domain is complete) it reflects the convergence back to f .

Strong convergence: any counterexample to strong convergence is clearly preserved by ι_1 , as the map does not affect redex positions. \square

A *signature extension* of an iTRS R is the disjoint union of R with another iTRS with empty set of rules.

One might expect that signature extensions preserve most properties of interest, but that is questionable. We can even have that a signature extension turns a strongly converging iTRS into one that is not converging at all.

Example 7.3. Consider the iTRS with rule $F(F(x)) \rightarrow G(x)$ under term metric ∞ — this is strongly converging. If we extend the signature with a function symbol H with associated umm $H_m(x) = \min(1, 2 \cdot x)$ then convergence is broken: the term $t = F(F(H(t)))$ exists under the combined metric, but repeatedly applying the rule to this term could only give us $u = G(H(u))$ which does not exist in that metric.

There are two “culprits” contributing to this: (i) the signature extension used an expansive umm, (ii) the iTRS was not depth-preserving. In this case, the counterexample can be reconstructed with non-expansive metrics:

Example 7.4. Modify Example 7.3 by setting $F_m(x) = x^2$, $G_m(x) = x$, $H_m(x) = \min(x, \frac{1}{2})$. In this case, neither of the two iTRSs permits infinite terms, but their combination does; the counterexample is preserved, as t does and u does not exist under this metric.

8. AUXILIARY DEFINITIONS AND OBSERVATIONS

When studying the disjoint union of term rewriting systems R and S it is useful to split a term into *layers* of the participating systems, and distinguish its *principal subterms*.

The principal subterm positions of a term t , $\text{PPos}(t) \subset \text{Pos}(t)$ have the following properties: p is principal if (i) the root of t/p is a function symbol belonging to a different signature than the root of t ; (ii) no proper prefix of p is principal. The top positions of a term t , $\text{TPos}(t) \subset \text{Pos}(t)$ are defined as follows: $\text{TPos}(t) = \{p \in \text{Pos}(t) \mid \forall q. q \prec p \Rightarrow q \notin \text{PPos}(t)\}$. Thus top positions are either parallel to principal subterms or they are their prefixes.

A *principal path* of t is an infinite path f where $f(0) = \lambda$, $f(n+1) = f(n) \cdot p_n$ such that $p_n \in \text{PPos}(t/f(n))$. The modularity of convergence is tightly connected to the question whether principal paths that arise through rewriting are convergent.

The top layer of a term is comprised of all the positions reachable from the root without a change of signature, and is thus “cut off” at its principal subterm positions. This leaves the layer with gaps at the cut-off points; therefore, the top layer is technically defined as the function that recreates a term if suitable fill material for the gaps is provided: The *top-layer* of a term t is a function $[t]$ with type $(\text{PPos}(t) \rightarrow \text{Ter}^m(\Sigma)) \rightarrow \text{Ter}^m(\Sigma)$ where we write $[t]_\xi$ for the application of $[t]$ to a function $\xi : \text{PPos}(t) \rightarrow \text{Ter}^m(\Sigma)$. The top layer has the following properties: above principal subterm positions it agrees with t :

$$\forall \xi. \forall p \in \text{PPos}(t). t \stackrel{p}{\sim} [t]_\xi$$

In positions parallel to principal subterms, t and $[t]$ coincide:

$$\forall \xi. \forall q \in \text{Pos}(t). (\forall p \in \text{PPos}(t). q \parallel p) \Rightarrow t/q = [t]_\xi/q$$

And in principal subterm positions themselves the gaps are filled with the function ξ :

$$\forall p \in \text{PPos}(t). [t]_\xi/p = \xi(p).$$

The standard notation [11] for layers in finite terms is $C[t_1, \dots, t_n]$. Elsewhere [19] this has been adapted for infinitary terms as $C[t_i]_{i \in I}$, to use some (possibly infinite) set I as indexing set for those positions. The main reason for departing from that is that the $[t]$ notation avoids the need of a separate indexing set by using $\text{PPos}(t)$ itself, as $[t]_\xi = t[\xi(p)]_{p \in \text{PPos}(t)}$.

We can define a distance between top-layers: $d([t], [u]) = d_m([t]_{kx}, [u]_{kx})$ where kx is the function that constantly returns x , where x is a fresh variable.

Lemma 8.1. *Let $t \in \text{Ter}^m(\Sigma)$. Let $f : \alpha \rightarrow \text{PPos}(t) \rightarrow \text{Ter}^m(\Sigma)$ be a function such that for each $p \in \text{PPos}(t)$ the functions $g_p : \alpha \rightarrow \text{Ter}^m(\Sigma)$ defined as $g_p(\beta) = f(\beta)(p)$ are converging sequences. Then $h : \alpha \rightarrow \text{Ter}^m(\Sigma)$ defined as $h(\beta) = [t]_{f(\beta)}$ is a converging sequence.*

Proof. Suppose t is a finite term. Then $\text{PPos}(t)$ is finite, so f is essentially an n -tuple of converging sequences which is also a converging sequence of n -tuples. Moreover, $[t]$ is uniformly continuous, as all term-forming umms are (Proposition 5 in [9]), and as uniformly continuous functions preserve Cauchy-sequences the result follows.

Otherwise (i.e. t is infinite), for any $\epsilon > 0$ there is a finite term t_ϵ such that $d_m(t, t_\epsilon) < \epsilon$. W.l.o.g. we can assume $\text{PPos}(t_\epsilon) \subseteq \text{PPos}(t)$. Thus the previous argument applies to t_ϵ and f , i.e. the function h_ϵ defined as $h_\epsilon(\beta) = [t_\epsilon]_{f(\beta)}$ is a converging sequence. Moreover, as subterm replacement is non-expansive (Proposition 4.1) the functions h and h_ϵ are pointwise within ϵ -distance. Thus h stays eventually within an ϵ -ball, and — as this holds for any ϵ — h must converge. \square

From this we get a compositionality property for converging sequences:

Corollary 8.2. *Let $f : \alpha \rightarrow \text{Ter}^m(\Sigma)$ be a sequence such that $[f]$ converges to $[u]$ and f/p converges for any $p \in \text{PPos}(u)$. Then f converges.*

Proof. We can apply Lemma 8.1 to u and f/p , giving a converging sequence $h : \alpha \rightarrow Ter^m(\Sigma)$ with $h(\beta) = \lceil u \rceil_{p \mapsto f(\beta)/p}$. Moreover, the distances between h and f are pointwise the same as the distances between their top-layers, which means that f must converge to the limit of h . \square

The (possibly infinite) rank of an infinitary term t is defined as follows:

$$rank(t) = \sup\{1 + rank(t/p) \mid p \in PPos(t)\}$$

From now the reasoning will be about the disjoint union of two non-collapsing iTRSs R and S with rewrite relation \rightarrow_{RS} and combined signature Σ , and for the top layer of any reduction (step or sequence) it is assumed that it is situated in system R , with signature Σ_R .

In the following we will be using a construction in which all principal subterms of a term t at finite rank n are replaced by a fixed term u , $t[n \searrow u]$. Formally:

$$\begin{aligned} t[0 \searrow u] &= u \\ t[n+1 \searrow u] &= \lceil t \rceil_\xi \quad \text{where } \xi(p) = (t/p)[n \searrow u] \end{aligned}$$

This notation is also extended to sequences of terms: if $f : \alpha \rightarrow Ter^m(\Sigma)$ is continuous then $f[n \searrow u] : \alpha \rightarrow Ter^m(\Sigma)$ is defined pointwise as $f[n \searrow u](\beta) = f(\beta)[n \searrow u]$. Notice that the function $t \mapsto t[n \searrow u]$ is non-expansive and therefore preserves the continuity of the sequence.

Lemma 8.3. *Let $f : \alpha \rightarrow Ter^m(\Sigma)$ be a reduction sequence. Let $u \in Ter^m(\Sigma)$ and $n \in \mathcal{N}$ be arbitrary. Then the sequence $f[n \searrow u]$ is a reduction sequence of the reflexive closure of \rightarrow_{RS} .*

Proof. The absence of collapsing rules means that principal subterm replacement at level n commutes with reduction steps at level $m < n$, even for contraction of non-left-linear redexes (even if rules with extra variables were permitted). Limits are also preserved because the operation is non-expansive. Reduction steps at lower level become reflexive steps. \square

Any reduction sequence of the MARS $(M, d, \rightarrow^=)$ can be turned into a (possibly shorter) reduction sequence of (M, d, \rightarrow) . Thus, the presence of reflexive steps in $f[n \searrow u]$ is merely a technicality. The full generality of Lemma 8.3 will be used later; but of special interest is the case $n = 1$:

Corollary 8.4. *Let $f : \alpha \rightarrow Ter^m(\Sigma)$ be a reduction sequence. Let $u \in Ter^m(\Sigma_R)$. Then $f[1 \searrow u]$ is a reduction sequence of the reflexive closure of \rightarrow_R .*

Proof. As all terms of the sequence are by construction Σ_R -terms each \rightarrow_{RS} -step is a \rightarrow_R -step. \square

We define the principal positions of a reduction sequence $f : \alpha \rightarrow Ter^m(\Sigma)$ as follows:

$$PPos(f) = \bigcup_{\beta < \alpha} \bigcap_{\beta < \gamma < \alpha} PPos(f(\gamma))$$

Lemma 8.5. *Let $f : \alpha \rightarrow Ter^m(\Sigma)$ be an open reduction sequence. If R is convergent then the top layers of f converge to some top layer $\lceil u \rceil$ where $PPos(u) = PPos(f)$.*

Proof. Since the iTRS of the top layer is convergent then so is $f[1 \searrow x]$, for some fresh variable x , and let its limit be t . Let c be any non-variable term in the other iTRS then by substitutivity of weak convergence we also have $f[1 \searrow x][c/x] = f[1 \searrow c] \rightarrow_w t[x/c]$, and we can set $u = t[x/c]$ since $\lceil u \rceil$ does not depend on the choice of c . Clearly, any principal position in $t[x/c]$ must be principal for $f(\beta)[1 \searrow c]$ for all γ in the open interval (β, α) for some β . \square

9. MODULARITY OF CONVERGENCE

In Example 6.3 we have seen how a top-layer can through rearrangement of its principal subterms interfere with the convergence at subterm positions. However, the example that rearranged principal subterms “in a substantial way” was divergent.

We will show that this is not an accident. The idea behind this is the following: if f is an open reduction sequence then so is $g = f[1 \searrow l]$ and if l is a term belonging to the (converging) top-layer then g must converge. Moreover, if $l \rightarrow r$ (in the interesting case $l \neq r$) then g must still converge if we interleave it with the occasional reduction of $l \rightarrow r$, provided we preserve non-left-linear redexes along the way.

It is impossible to have both l and r occur infinitely often in a fixed position p , because that would create a divergence with diameter of at least $d_\infty(l, r) \cdot 2^{-|p|}$ in $Ter^\infty(\Sigma)$, which by Proposition 4.4 gives divergence for $Ter^m(\Sigma)$ as well. That means that two situations are ruled out: (i) that the top-layer reduction drags eventually infinitely many different subterms into position p — as in $f[1 \searrow l]$ each l could be reduced once it appears in position p ; (ii) at the same time infinitely many terms appearing in position p are “descendants” of a principal subterm t and infinitely many others are not. In that case we would create a non-converging sequence by replacing all weak reducts of t in f by r , and all other principal subterms by l . With both these scenarios ruled out one can show that the sequence of subterms in position p must be a focussed sequence.

These ideas will be formalised in a slightly different way that avoids the tracing of redexes and the rather cumbersome notion of “descendant” — which is impossible anyway for non-strong reductions [18].

9.1. Predicate Sequences. A *predicate sequence* is an ordinal-indexed sequence of predicates s on terms such that $\beta \leq \gamma \wedge t \rightarrow_w u \wedge s(\gamma)(u) \Rightarrow s(\beta)(t)$. A predicate sequence is meant to be used alongside a reduction sequence of the same length, w.r.t. the same signature and term metric. The idea behind predicate sequences is that they are used as a tool to decide when to reduce an l to r in $f[1 \searrow l]$; or rather: some principal subterms will be replaced by l , some by r , and the construction makes sure that reduction steps are preserved. So the motivation behind the definition of a predicate sequence is this: once we have decided to replace a principal subterm t by r (rather than l) then (i) we cannot change our mind *later*, i.e. for a larger indexing ordinal, and (ii) all “descendants” of t need to be replaced by r too, which we can safely over-approximate as all *weak reducts* of t .

Let $f : \alpha \rightarrow Ter^m(\Sigma)$ be a reduction sequence and p be an ω -sequence of positive integers. The function $f_p : \alpha \rightarrow Ter^m(\Sigma) \rightarrow \mathcal{B}$ maps ordinals to predicates on terms, and is defined as follows:

$$f_p(\beta)(t) \iff \exists \gamma. \beta \leq \gamma < \alpha \wedge p \in \text{PPos}(f(\gamma)) \wedge t \rightarrow_w f(\gamma)/p$$

Thus the predicate f_p is true for terms that weakly reduce to terms which appear in principal position p later in the sequence.

Lemma 9.1. *f_p is a predicate sequence.*

Proof. Assume $\beta \leq \gamma$ and $t \rightarrow_w u$ and $f_p(\gamma)(u)$. $f_p(\gamma)(u)$ means that there is a $\zeta \geq \gamma$ such that p is a principal position in $f(\zeta)$ and $u \rightarrow_w f(\zeta)/p$. In order to show that $f_p(\beta)(t)$ we can choose the same ζ as the witness; $\beta \leq \zeta$ holds because $\beta \leq \gamma \leq \zeta$; $t \rightarrow_w f(\zeta)/p$ holds because $t \rightarrow_w u \rightarrow_w f(\zeta)/p$, and \rightarrow_w is transitive. \square

Explanation: the function f_p watches the subterms in position p , where the ordinal parameter can be seen as a notion of time; f_p is true for a term if a weak reduct of the term will “in the future” appear in position p . The only way f_p can (in time) switch from true to false is when at some point in the reduction sequence the term in position p is replaced by another that is not a reduct of the previous. This is only possible if a redex in a position above p is contracted, i.e. a redex of the top layer; in the simulated sequence this redex can be employed to replace an r in position p by an l taken from elsewhere.

Another predicate sequence we shall be using is not dependent on the ordinal, and defined as: $k_t(\beta)(u) \iff \neg(t \rightarrow_w u)$. Thus k_t is constantly true for all terms that are not weak reducts of t .

Lemma 9.2. *k_t is a predicate sequence.*

Proof. Let $r \rightarrow_w s$ and $k_t(\gamma)(s)$. Thus $\neg(t \rightarrow_w s)$. To show $k_t(\beta)(r)$ we need to show $\neg(t \rightarrow_w r)$. But $t \rightarrow_w r$ gives $t \rightarrow_w r \rightarrow_w s$ and so by transitivity of \rightarrow_w the contradiction $t \rightarrow_w s$. \square

Given a rule $l \rightarrow r$ in R and a predicate sequence s we define a dependent function $\xi_s : \Pi t : Ter^m(\Sigma). \alpha \rightarrow \text{PPos}(t) \rightarrow Ter^m(\Sigma)$ as follows:

$$\xi_s(t)(\beta)(p) = \begin{cases} l & \text{if } s(\beta)(t/p) \\ r & \text{otherwise} \end{cases}$$

Explanation: the reason this is a dependent function is that the type of the third parameter $\text{PPos}(t)$ depends on the first parameter, the term t . Thus $\xi_s(t)(\gamma)$ is a function of type $\text{PPos}(t) \rightarrow Ter^m(\Sigma)$, which is exactly the type of function needed to fill the holes in $[t]$.

Proposition 9.3. *Distributive properties of ξ_s w.r.t. subterm selection and substitution application in top positions:*

$$\begin{aligned} \forall p \in \text{TPos}(t). [t]_{\xi_s(t)(\beta)}/p &= [t/p]_{\xi_s(t/p)(\beta)} \\ \forall u \in Ter^m(\Sigma_R). [\bar{\sigma}(u)]_{\xi_s(\bar{\sigma}(u))(\beta)} &= \bar{\sigma}'(u) \\ \text{where } \forall x. \sigma'(x) &= [\sigma(x)]_{\xi_s(\sigma(x))(\beta)} \end{aligned}$$

Proof. The first property is trivially proven by induction on the length of p . The second is a corollary, because if $u \in Ter^m(\Sigma_R)$ (i.e. only using function symbols from the signature of the top layer) then all variable positions in u become top positions in $\sigma(u)$. \square

The function ξ_s is used to create reducts of $t[1 \searrow l]$ by reducing some of the l to r , guided by the predicate sequence. In particular, it is used to replace principal subterms in a reduction sequence.

Proposition 9.4. *The fundamental properties of the function ξ_s are:*

$$\begin{aligned} t \rightarrow_{RS} u &\Rightarrow [t]_{\xi_s(t)(\beta)} \rightarrow_{\bar{R}} [u]_{\xi_s(u)(\beta)} \\ \beta \leq \gamma &\Rightarrow [t]_{\xi_s(t)(\beta)} \rightarrow_w [t]_{\xi_s(t)(\gamma)} \end{aligned}$$

Proof. First property: let q be the redex position. At positions p parallel to q , $t \stackrel{p}{\sim} u$ thus also $[t]_{\xi_s(t)(\beta)} \stackrel{p}{\sim} [u]_{\xi_s(u)(\beta)}$.

If $p \preceq q$ for some $p \in \text{PPos}(t)$ then certainly $p \in \text{PPos}(u)$ and $[t]_{\xi_s(t)(\beta)}$ and $[u]_{\xi_s(u)(\beta)}$ also coincide in all positions that are prefixes of p . At position p itself these two terms have either l or r . If $\xi_s(u)(\beta)(p) = r$ then there is nothing to show, as both $l \rightarrow_{\bar{R}} r$ and $r \rightarrow_{\bar{R}} r$. If $\xi_s(u)(\beta)(p) = l$ then we know by the definition of ξ_s that $s(\beta)(u/p)$ holds. Since we know $t/p \rightarrow_{RS} u/p$ we can use the predicate sequence property of s to establish $s(\beta)(t/p)$. This means by the definition of ξ_s that $\xi_s(t)(\beta)(p) = l$, and therefore $[t]_{\xi_s(t)(\beta)} = [u]_{\xi_s(u)(\beta)}$ as the terms also agree in position p .

Otherwise the redex q is a top position, $t/q = \sigma(v)$, $u/q = \sigma(w)$ for some rule $v \rightarrow w$ in R . Using Proposition 9.3 we get: $[t]_{\xi_s(t)(\beta)} = C'[\overline{\sigma'}(v)]$ and $[u]_{\xi_s(u)(\beta)} = C'[\overline{\sigma'}(w)]$ for some context C' and some substitution σ' .

So in particular: $[t]_{\xi_s(t)(\beta)} = C'[\overline{\sigma'}(v)] \rightarrow_R C'[\overline{\sigma'}(w)] = [u]_{\xi_s(u)(\beta)}$.

The reason for the second property is that $s(\beta) \Leftarrow s(\gamma)$; consequently, fewer (but not more) principal subterms might be set to l ; again, all these can be reduced to r — this might require ω steps, if the top layer of t is infinite. \square

Given an open reduction sequence $f : \alpha \rightarrow \text{Ter}^m(\Sigma)$ and a predicate sequence s we write $[f]_s : \alpha \rightarrow \text{Ter}^m(\Sigma)$ for the function defined as follows (here n is a finite ordinal and λ a limit ordinal or 0):

$$\begin{aligned} [f]_s(\lambda + 2 \cdot n) &= [f(\lambda + n)]_{\xi_s(f(\lambda+n))(\lambda+n)} \\ [f]_s(\lambda + 2 \cdot n + 1) &= [f(\lambda + n)]_{\xi_s(f(\lambda+n))(\lambda+n+1)} \end{aligned}$$

This definition may require some explanation

- (1) Because f is open, α is a limit ordinal and thus $\lambda + n < \alpha$ implies $\lambda + n + n + 1 < \alpha$, for finite n . Hence α remains as the domain for $[f]_s$.
- (2) That the case distinction into even and odd ordinals is well-defined and covers all cases follows from the normal-form theorem for ordinals [17, page 323].
- (3) The purpose of this definition is the following. A single reduction step $f(\beta) \rightarrow f(\beta + 1)$ of the original sequence is split into two stages:

$$[f(\beta)]_{\xi_s(f(\beta))(\beta)} \rightarrow_w [f(\beta)]_{\xi_s(f(\beta))(\beta+1)} \rightarrow_{\bar{R}} [f(\beta + 1)]_{\xi_s(f(\beta+1))(\beta+1)}$$

This follows from Proposition 9.4. Thus the function $[f]_s$ is a weak reduction sequence.

- (4) As all terms in $[f]_s$ belong to system R the sequence must be convergent — the same argument also explains why $[f]_s$ must be continuous at limit ordinals $\lambda < \alpha$, i.e. why it is a proper sequence.

We can illustrate the construction $[f]_s$ more concretely on an example.

Example 9.5. Consider the two iTRS from Example 6.3; they allowed to eventually rewrite $J(t)$ where $t = K(E, t)$ to $J(u)$ where $u = K(S^\infty, u)$. One way of doing this in $\omega + \omega$ steps is by first rewriting t using rules of the first system: $t \rightarrow_w K(Z, t) \rightarrow_w K(Z, K(S(z), t)) \rightarrow_w K(Z, K(S(Z), K(S(S(Z)), t))) \rightarrow_w \dots$. After reaching a limit at ω we can apply the single rule of the other system repeatedly at the root: if $u_n = K(S^n(Z), u_{n+1})$ then $J(u_n) \rightarrow J(u_{n+1})$, reaching $J(u)$ in the limit.

If this is our sequence f we can form $[f]_{f_p}$, e.g. for $p = 1 \cdot 1$. The initial term of that sequence would be $t_0 = J(a_0)$ where $a_0 = K(l, a_0)$ — for l we can choose $J(K(x, y))$ as it is a left-hand side of a non-trivial rewrite rule. Rewriting E to Z in position p does not change that, because all principal subterms can still rewrite to that term, so $t_1 = t_0$. Once we rewrite the other occurrences of E with the rule $E \rightarrow H(E)$ they no longer weakly reduce to Z , but they still weakly reduce to some term that appears in position p in the second half of sequence f . Therefore, $[f]_{f_p}$ is constantly t_0 for the first ω steps. The step $f(\omega) = J(u_0) \rightarrow J(u_1) = f(\omega + 1)$ is split in $[f]_{f_p}$ into two steps: $[f]_{f_p}(\omega) = J(a_0) \rightarrow J(K(r, a_0)) \rightarrow J(a_0) = [f]_{f_p}(\omega + 2)$. The first step rewrites the topmost l to r , because none of the terms appearing in position p in sequence f henceforth is a weak reduct of Z . All the other principal subterms will still appear in p and thus they will remain l in $[f]_{f_p}$. For all subsequent steps exactly the same happens, as $S(Z)$, $S(S(Z))$, etc. will eventually appear a last time in p . Therefore, in the second half of $[f]_{f_p}$ the values always alternate between $J(a_0)$ and $J(K(r, a_0))$, which makes $[f]_{f_p}$ non-convergent.

The reason a convergent reduction was mapped by this construction to a non-convergent one is that f is convergent “in the wrong way”, i.e. its principal subterms in p do not form a focussed sequence, which is only possible if the system of the top layer is itself non-convergent.

9.2. General Observations.

Lemma 9.6. *Let $f : \alpha \rightarrow \text{Ter}^m(\Sigma)$ be an open reduction sequence and let R be convergent. Let $p \in \text{PPos}(f)$. Then the sequence f/p is focussed.*

Proof. First notice that $[f]_{f_p}$ converges. By construction, that sequence has an l in position p for all even ordinals (including limits), but it reduces that l to r whenever the corresponding subterm in f fails to have any more reducts in position p . This cannot happen arbitrarily close to α , so there is a cut-off point β after which all $f(\gamma)/p$ have reducts $f(\kappa)/p$, for some $\kappa > \gamma$.

Now assume, for some $\zeta > \beta$ there was no cut-off point ϕ such that for all $\psi \geq \phi$, $f(\zeta)/p \rightarrow_w f(\psi)/p$. Then we could construct the weak reduction sequence $[f]_{k_{f(\zeta)/p}}$ in R . This replaces all weak reducts of $f(\zeta)/p$ in principal positions by r and all other principal subterms by l . By the assumption and the previous observation we have both l and r occurring at p arbitrarily close to α which contradicts convergence. Hence the assumption was wrong and f must be focussed. \square

This result means together with the previous propositions about focussed sequences that if reductions in principal subterm positions would converge then so would the reduction as a whole. But why would they converge? We certainly have the following:

Theorem 9.7. *If two non-collapsing iTRSs R and S are convergent and f is a reduction sequence of their disjoint union then $f[n \searrow s]$ is convergent, for any n and s .*

Proof. Convergence is proved by induction on n .

For each $p \in \text{PPos}(f)$ the sequence f/p is focussed by Lemma 9.6, and so is $f[k+1 \searrow s]/p = f/p[k \searrow s]$. Thus by the induction hypothesis and Proposition 3.8 these sequences must converge. This means we can apply Corollary 8.2, establishing the result. \square

Example 9.8. Consider the iTRS under term metric ∞ with rule $F(F(x)) \rightarrow G(x)$, combined with a rule-less iTRS with function symbol H and term metric $H_d(x) = \min(1, 2 \cdot x)$. We can construct the infinite term $t = H(F(F(t)))$ in the combined metric. We can also construct an open reduction sequence g of length ω by starting with t , always reducing its topmost redex. By Theorem 9.7 each $g[n \searrow s]$ is converging, but g itself is not, the distances between $g(k)$ and $g(k+1)$ always remain 1.

In Example 9.8 the sequence was diverging, because the infinite term $u = H(G(u))$ (which exists in $\text{Ter}^\infty(\Sigma)$) towards which the reduction was seemingly heading did not exist in $\text{Ter}^m(\Sigma)$, as its finite approximants do not form Cauchy-sequences under metric d_m . This is always the case:

Proposition 9.9. *If two non-collapsing iTRSs R and S are convergent and f is a reduction sequence of their disjoint union then f converges under the metric d_∞ .*

Proof. From Theorem 9.7 we know that $f[n \searrow s]$ is converging (under metric d_m), and by Proposition 4.4 we also know that it is Cauchy under metric d_∞ . Under that metric the sequences f and $f[n \searrow s]$ are pointwise within distance of 2^{-n} , and therefore f must eventually stay within a ball with diameter 2^{-n} . As this holds for any n , f must be Cauchy (and by completeness: converging) under metric d_∞ . \square

Corollary 9.10. *Convergence is a modular property for non-collapsing iTRSs over metric d_∞ .* \square

This slightly extends the result from [10] as the proof here also covers generalised iTRSs with extra variables. Another consequence of Propositions 9.9 and 4.6 is that the question whether such a sequence f is converging under metric d_m reduces to the question whether its limit t under d_∞ is a term in $\text{Ter}^m(\Sigma)$, and by Proposition 4.8 this is equivalent to whether all its infinite paths converge. However, infinite paths that traverse through only finitely many signature changes would be contained in the positions of some $t[n \searrow s]$ and therefore must converge. Thus it suffices to only consider principal paths.

9.3. Conditions on the Metric. In general, the interplay between the umms of the constituent term metrics can be quite subtle. As Example 7.4 shows, it is even possible that only this interplay gives rise to infinite terms. In this section we are looking at metrics which are much curtailed in that respect.

A term metric g is called *granular* iff each component $F_{i,g}^{\sim}$ of the umm of each function symbol F is either the identity function or the halving function. Clearly, when we form the disjoint union of two iTRSs with granular term metrics the resulting term metric is granular.

As a special case ∞ is granular. Moreover, as indicated in the Introduction, the term universe associated with algebraic types in Haskell with strictness annotations is expressible through granular term metrics.

Granular term metrics have comparatively simple convergence properties: a full infinite path f of $t \in \text{Ter}^\infty(\Sigma)$ is converging w.r.t. term metric g iff it contains infinitely many non-strict positions. If f is a path of t then we define the step function as follows:

$$\begin{aligned} \text{step}_g(t, f, 0) &= 0 \\ \text{step}_g(t, f, n+1) &= \begin{cases} \text{step}_g(t, f, n), & \text{if } (t, f(n))_g = (t, f(n+1))_g \\ \text{step}_g(t, f, n) + 1, & \text{otherwise} \end{cases} \end{aligned}$$

Although that definition would technically carry over for other term metrics, for granular term metrics the step-function tells us about the rate of convergence.

Collapsing rules cause problems in term metric ∞ for two separate reasons:

- they allow to make layers merge, so that computations originating in a lower layer can impact on a higher layer;
- the top layer can disappear altogether, so that in terms of overall convergence the computation has made no progress up to the point where that layer collapsed.

With granular term metrics there is a different kind of rule that recreates the second problematic aspect of collapsing rules. This happens with a rule $l \rightarrow r$ when for some variable x , $vdepth(x, r, g)(1) = 1$ and $vdepth(x, l, g)(1) < 1$; such rules are called *pseudo-collapsing*. Note that if all argument positions are lazy ($g = \infty$) this is only possible if $l \rightarrow r$ is a collapsing rule.

Although a pseudo-collapsing rule cannot make a layer itself disappear, it can make it disappear “semantically”, i.e. regarding distances. This has an immediate consequence on convergence:

Proposition 9.11. *If a non-collapsing iTRS contains a pseudo-collapsing rule then it cannot be convergent.*

Proof. Let $l \rightarrow r$ be that rule and $vdepth(x, r, g)(1) = 1$ and $vdepth(x, l, g)(1) < 1$. We can form the infinite term $t = l[t/x]$, because $vdepth(x, l, g)(1) < 1$. For the rewrite step $r[t/x] \rightarrow r[r[t/x]/x]$ we notice that $d_g(r[t/x], r[r[t/x]/x]) = d_g(t, r[t/x])$ because $vdepth(x, r, g)(1) = 1$. This means the rewrite sequence commencing with t has identical distances in each rewrite step. This can only be convergent if they are all 0, which means that the substitution t/x must be a unifier for l and r . This in turn is only possible if $r = x$ which is ruled out by the non-collapsing condition. \square

As pseudo-collapsing rules break convergence we know that such rules must be absent when we combine two non-collapsing convergent iTRSs (over term metric g). Their absence also gives an impact on an invariant of the step-function:

Lemma 9.12. *Let $t \rightarrow_w u$ in the disjoint union of 2 non-collapsing iTRSs with joint term metric g . For every principal path q of u there is a principal path p of t such that $\forall j. \text{step}_g(u, q, j) \geq \text{step}_g(t, p, j)$.*

Proof. Single rewrite step: If q is a principal path not crossing the redex position then q is also in t and the result is immediate. Otherwise there is an i such that $q(i)$ is a position above the contractum, and $q(i+1) = q(i) \cdot v$ the position v is below a variable position of some variable y in the contractum. On the left-hand side of the step there is a path p sharing the first i positions with q , where $p(i+1) = p(i) \cdot w$ such that w is a (depth-minimal) position of y in the redex. As the applied rule is not pseudo-collapsing, if $(t, p(i))_g \neq (t, p(i+1))_g$ then we also must have $(u, q(i))_g \neq (u, q(i+1))_g$.

The transitive closure clearly preserves the property. Transfinite reductions do not change it either as (in the absence of rules with extra variables) limits cannot introduce new principal paths. \square

What Lemma 9.12 says is that the rate of convergence cannot go down when we rewrite, provided we measure it in numbers of layers (rather than actual distances).

Theorem 9.13. *The disjoint union of non-collapsing iTRSs with granular term metrics preserves convergence.*

Proof. Let f be a reduction sequence in that disjoint union. We show that for any $\epsilon > 0$ there is an n such that pointwise $d_m(f(\gamma), f(\gamma)[n \searrow x]) < \epsilon$; by convergence of $f[n \searrow x]$ this implies that f stays eventually within an ϵ -ball, and as this holds for any ϵ , f converges.

We can find the n as follows: first notice that for some k , $\epsilon \geq 2^{-k}$. We need to find an n such that $\text{step}(f(0), p, n) \geq k$ for every principal path p of $f(0)$. As $f(0) \in \text{Ter}^m(\Sigma)$, every principal path p in it is converging which implies $\text{step}(f(0), p, n) \geq k$ for some n , and as $f(0)$ is a finitely branching tree, some n will do for all.

By Lemma 9.12 this property is preserved by rewriting: if $t \rightarrow_w u$ and $\text{step}(t, p, n) \geq k$ for all principal paths p of t then also $\text{step}(u, q, n) \geq k$ for all principal paths q of u . This also implies that u and $u[n \searrow x]$ are within 2^{-k} distance of each other. \square

10. MODULARITY OF STRONG CONVERGENCE

Simonsen showed [19] that *top-termination* (and thus strong convergence) is a modular property of left-linear iTRSs under term metric ∞ . Left-linearity is indeed an unnecessary constraint:

Theorem 10.1. *Over term metric ∞ , the disjoint union of two strongly converging non-collapsing iTRSs is a strongly converging iTRS.*

Proof. Strong convergence implies that the contributing iTRSs are non-collapsing, hence the previous results about convergence apply. According to [24] an iTRS is strongly converging if and only if it contains no reduction sequence that reduces a root redex infinitely often. Suppose f was such a sequence. But $f[1 \searrow x]$ resides in the iTRS R of the top layer; as the construction preserves all top layer reductions it also preserves all root reductions which contradicts the premise that R is strongly converging. \square

This statement does not even need a side condition banning non-collapsing rules, since these break the strong convergence property on the individual iTRSs and are therefore known to be absent. For other metrics this is no longer true: term metric id (all components of all umms are the identity function) rules out infinite terms and makes strong convergence coincide with strong normalisation.

Theorem 10.2. *Let R and S be non-collapsing strongly convergent iTRSs. If $R + S$ is convergent then it is strongly convergent.*

Proof. The same argument as in the proof of Theorem 10.1 applies, i.e. if f is not top-terminating then $f[1 \searrow x]$ is not either. \square

11. CONJECTURES AND OPEN PROBLEMS

I would conjecture that the disjoint union of a convergent iTRS with a strongly convergent one is convergent (under term metric ∞). If the (merely) convergent iTRS contained no collapsing rule then this follows already from the results in [10], but it may hold without that proviso. What makes a proof for this conjecture technically awkward is that collapsing rules allow layers to merge. However, this capability is highly restricted in converging systems.

Convergence requires that reduction sequences of *any length* converge. A related property is α -convergence, which means that reduction sequences up to length α converge, while longer ones may not. Example 6.2 is not convergent in general, but it is $\omega \times n$ -convergent, for finite n . It depends on the ordinal α whether the constructions from this paper would still work unchanged.

The most interesting special case is ω -convergence, and there is a problem with constructing the simulated reduction sequence relative to a predicate sequence: when “time is moved on” the simulation potentially requires ω -many steps to rewrite potentially infinitely many occurrences of l to r . This cannot be done if the overall length of the simulated sequence is limited to ω . Finitely many steps would suffice to move on time if all rules were *finitely linear*, which means that variables only occur a finite number of times on both left- and right-hand sides of rules. While necessary to repair the proof it is unclear whether that constraint is necessary to repair the theorem: are there ω -convergent iTRSs that are not finitely linear and whose disjoint union fails to be ω -convergent?

12. CONCLUSION

For the standard metric d_∞ it has been shown that both weak and strong convergence are modular properties of non-collapsing iTRSs. Otherwise, the iTRSs involved can have various usually undesirable properties: non-left-linear rules, infinite left-hand sides, and infinitely many rules are all permitted.

The proof for strong convergence is perfectly straightforward, the proof for weak convergence is not. At the heart of the latter is a consideration that the reductions in the top layer are limited in the way they can rearrange the principal subterms, in particular that the principal subterms in a fixed position must form a *focussed* sequence. Moreover, abstract reduction systems are converging iff their focussed sequences do, and both observations together give that result.

For other metrics, these arguments still apply to reductions of finite rank (Theorem 9.7), while for reductions of infinite rank there are counterexamples. As a general proof strategy, various convergence arguments have been constructed along the central idea of limiting the pointwise distance between the reduction sequence under investigation and a finite rank approximation of that sequence. These allow to generalise the modularity of convergence and strong convergence of non-collapsing iTRS beyond metric d_∞ , to systems with either granular term metrics or (not elaborated here) arbitrary metrics but depth-preserving rules.

REFERENCES

- [1] A. Arnold and M. Nivat. Metric interpretations of infinite trees and semantics of nondeterministic recursive programs. *Theoretical Computer Science*, 11(2):181–205, June 1980.

- [2] Franz Baader and Tobias Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.
- [3] Hendrik P. Barendregt. *The Lambda-Calculus, its Syntax and Semantics*. North-Holland, 1984.
- [4] E.T. Copson. *Metric Spaces*. Cambridge University Press, 1968.
- [5] Nachum Dershowitz and Stéphane Kaplan. Rewrite, Rewrite, Rewrite, Rewrite, Rewrite, ... In *Principles of Programming Languages*, pages 250–259. ACM, 1989.
- [6] Theodore W. Gamelin and Robert Everist Greene. *Introduction to Topology*. Dover, 1999.
- [7] András Hajnal and Peter Hamburger. *Set Theory*. London Mathematical Society, 1999.
- [8] Simon Peyton Jones, editor. *Haskell 98, Languages and Libraries, The Revised Report*. Cambridge University Press, 2003.
- [9] Stefan Kahrs. Infinitary rewriting: Meta-theory and convergence. *Acta Informatica*, 44(2):91–121, May 2007.
- [10] Stefan Kahrs. Modularity of convergence in infinitary rewriting. In Ralf Treinen, editor, *Rewriting Techniques and Applications*, volume 5595 of *LNCS*, pages 179–193. Springer, 2009.
- [11] Richard Kennaway and Fer-Jan de Vries. *Term Rewriting Systems*, chapter Infinitary Rewriting, pages 668–711. Cambridge University Press, 2003.
- [12] Richard Kennaway, Jan Willem Klop, Ronan Sleep, and Fer-Jan de Vries. Transfinite reductions in orthogonal term rewriting systems. *Information and Computation*, 119(1):18–38, 1995.
- [13] Jeroen Ketema. *Böhm-Like Trees for Rewriting*. PhD thesis, Vrije Universiteit Amsterdam, 2006.
- [14] Aart Middeldorp. Modular aspects of properties of term rewriting systems related to normal forms. In *Rewriting Techniques and Applications*, pages 263–277, 1989. *LNCS* 355.
- [15] Michael D. Potter. *Sets: An Introduction*. Oxford University Press, 1990.
- [16] Michael Rusinowitch. On termination of the direct sum of term-rewriting systems. *Information Processing Letters*, 26:65–70, 1987.
- [17] Waclaw Sierpiński. *Cardinal and Ordinal Numbers*. Polish Scientific Publishers, 1965.
- [18] Jakob Grue Simonsen. On confluence and residuals in cauchy convergent transfinite rewriting. *Information Processing Letters*, 91(3):141–146, 2004.
- [19] Jakob Grue Simonsen. On modularity in infinitary term rewriting. *Information and Computation*, 204(6):957–988, 2006.
- [20] V. Stoltenberg-Hansen and J.V. Tucker. Effective algebras. In S. Abramsky, Dov M. Gabbay, and T.S.E. Maibaum, editors, *Handbook of Logic in Computer Science, Volume 4*, pages 357–526. Oxford University Press, 1995.
- [21] Y. Toyama, J.W. Klop, and H. Barendregt. Termination for the direct sum of left-linear term rewriting systems. In *Rewriting Techniques and Applications*, pages 477–491, 1989. *LNCS* 355. (Revised version to appear in *JACM*).
- [22] Yoshihito Toyama. On the Church-Rosser property for the direct sum of term rewriting systems. *Journal of the ACM*, 34(1):128–143, 1987.
- [23] Yoshito Toyama. Counterexamples to termination for the direct sum of term rewriting systems. *Information Processing Letters*, 25:141–143, 1987.
- [24] Hans Zantema. Normalisation of infinite terms. In *RTA 2008*, volume 5117 of *LNCS*, pages 441–455, 2008.