

Stefan Kahrs

Infinitary Rewriting: Meta-Theory and Convergence

the date of receipt and acceptance should be inserted later

Abstract When infinitary rewriting was introduced by Kaplan et. al. [9] at the beginning of the 1990s, its term universe was explained as the metric completion of a metric on finite terms. The motivation for this connection to topology was that it allowed to import other well-studied notions from metric spaces, in particular the notion of convergence as a replacement for normalisation.

This paper generalises the approach by parameterising it with a term metric, and applying the process of metric completion not only to terms but also to operations on and relations between terms. The resulting meta-theory is studied, leading to a revised notion of infinitary rewrite system. For these systems a method is devised to prove their convergence.

1 Introduction

Infinitary rewriting is a variation of term rewriting that studies infinite terms and reduction sequences of infinite length. The subject had been introduced by Der-showitz and Kaplan at POPL 1989 [9].

There are different ways to introduce infinite terms. In the view that terms are functions from tree domains to symbols [4], simply dropping the demand that a tree domain is finite would permit infinite terms. Alternatively, infinite term can be introduced through the ideal completion of a partial order on finite terms; we even have some influence on which infinite terms we (do not) want by choosing the order with care.

Another alternative which goes back to at least [1] is now standard: define a metric that measures how different two terms are, typically $d(t, t') = 2^{-k}$ where k is the length of the shortest position at which the two terms t and t' differ; the

S. Kahrs

Department of Computer Science, University of Kent at Canterbury, Canterbury CT2 7NF,
United Kingdom, Tel.: 00441227827146, Fax: 00441227762811
E-mail: S.M.Kahrs@kent.ac.uk

metric completion of the term set produces infinite terms. Beside infinite terms this approach offers a topology and a notion of convergence. Again, the choice of metric allows to prevent the creation of infinite terms we do not want.

We can view convergence as a generalisation of termination: the only converging sequences over metrics with discrete topologies are the ones that remain fixed after finitely many steps. On finite terms, all term metrics have a discrete topology — we shall see later (proposition 4) why. In the presence of infinite terms, a convergent sequence may not reach a normal form but it can approximate an infinite one in transfinitely many steps. Of particular interest to a user of infinitary rewriting is then the question: “is my system *as a whole* convergent?”.

Although the original papers by Kaplan et. al. defined the concept of a system (an abstract reduction system over a metric space) being converging, their proof methods aimed lower for infinitary rewriting, as they excluded reduction sequences starting with infinite terms. This restriction hints at a fundamental problem with their approach: they permit too many terms or too many rules to achieve convergence proofs that apply generally. More recently, Terese [21] chose not even to define the concept of a converging metric ARS; they focused instead on the notion of *strong convergence* but only w.r.t. reduction sequences, not systems.

The approach used in this paper goes back to the original notion of semantic convergence and shows how infinitary rewrite systems can be proved convergent w.r.t. this notion.

Very early on, papers on infinitary rewriting allowed transfinite reduction sequences *beyond* ordinal ω . This is not followed here: the ordinal ω is intrinsically tied to the notion of metric completion and the relations studied here are “continuous” w.r.t. the term topology. Of course, one could work instead with topological notions of convergence, based on filters or nets, see e.g. [20], and use a matching notion of completion. In this paper I stick to the ω -case, at least in part because it allows me to make use of known properties of metric completion.

In summary, this paper provides the following:

- the preliminary section 2 recalls concepts from rewriting and topology; a couple of basic propositions about metric spaces are proved that are used later but which I could not find in standard textbooks;
- section 4 introduces a general concept of what a term metric actually *is*, based on the notion of ultra-metric map (section 3);
- section 5 studies fundamental properties of infinitary terms, i.e. the terms that arise through metric completion;
- section 6 re-defines the notions of infinitary rewrite rule and infinitary TRS;
- section 7 shows how (and which) relations can be lifted through metric completion — these results have wider applications than infinitary rewriting and are stated in a more general way;
- section 8 studies under which conditions infinitary TRSs match the requirements of the previous section;
- section 9 reduces convergence proofs for infinitary TRSs to proofs that certain finitary TRS are Cauchy;
- section 10 shows a method to prove finitary TRSs to be Cauchy;
- section 11 shows the method working on an example;
- sections 12 and 13 conclude, with 12 pointing out a number of possible alternative approaches.

2 Preliminaries

We rely on notations and terminology from both Term Rewriting and Topology which are introduced in this section. For the former to the conventions and notations from [25] are used, experts in the area may therefore want to skip that section.

Rewriting

We write $R : A \leftrightarrow B$ to declare a relation R between sets A and B . Given a binary relation R , we write R^{-1} for the relation $x R^{-1} y \iff y R x$. We write $R ; S$ for the composition of the relations R and S , i.e. $x(R ; S)y \iff \exists z. xRz \wedge zSy$. A *quasi ordering* on A is a reflexive and transitive relation on A .

An *Abstract Reduction System* is a structure $\mathcal{A} = (A, \rightarrow)$ where A is a set and \rightarrow a binary relation on A . Given an ordinal α , a *sequence of length α* in the set S is a function $f : \alpha \rightarrow S$, viewing α as a von Neumann ordinal, i.e. identifying it with the set of all smaller ordinals. A *reduction sequence* of length α in an ARS $(A, >)$ is a sequence f of length α in A such that $\forall n \in \alpha. (n+1) \in \alpha \Rightarrow f(n) > f(n+1)$. An ARS is *terminating*, or *strongly normalising*, if it has no reduction sequences of length ω .

Given an ARS (A, \rightarrow) and an equivalence relation \approx on A the ARS A/\approx has as objects the equivalence classes $[t]_{\approx}$ of A , and the relation \rightarrow/\approx is defined as:

$$[t]_{\approx} \rightarrow/\approx [u]_{\approx} \iff t \approx ; \rightarrow ; \approx u$$

One says that t rewrites to u *modulo \approx* .

A *signature* is a pair $\Sigma = (\mathcal{F}, \#)$ where \mathcal{F} is a set (of functions symbols) and $\# : \mathcal{F} \rightarrow \mathcal{N}$ the function assigning each symbol its arity. As notational convention function symbols are written as upper-case letters. The infinite set of *variables* is called *Var* and particular variables are referred to by lower-case letters. The *set of (finite) terms* over Σ is indicated as $Ter(\Sigma)$ and is defined inductively: (i) $Var \subset Ter(\Sigma)$, (ii) if $F \in \mathcal{F}$ and $\#(F) = n$ and $t_1, \dots, t_n \in Ter(\Sigma)$ then $F(t_1, \dots, t_n) \in Ter(\Sigma)$. Parentheses are dropped when $n = 0$. The set of variables occurring in the term t is called $var(t)$. The *root* of a term t is either t , if $t \in Var$, or F , if $t = F(t_1, \dots, t_n)$.

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 $h(\llbracket t \rrbracket_A^\rho) = \llbracket t \rrbracket_B^{h \circ \rho}$.

Given any Σ -algebra A , $\llbracket A \rrbracket$ is the Σ -algebra with carrier set $(Var \rightarrow A) \rightarrow A$ and $F_{\llbracket A \rrbracket}(a_1, \dots, a_n)(\rho) = F_A(a_1(\rho), \dots, a_n(\rho))$. The function $\rho : Var \rightarrow \llbracket A \rrbracket$ given

as $\rho(x)(f) = f(x)$ makes it possible to interpret terms with variables. We write $[[t]]_{[[A]]}$ as shorthand for $[[t]]_{[[A]]}^\rho$.

$Ter(\Sigma)$ is itself a Σ -algebra with $F_{Ter(\Sigma)} = F$, the so-called term algebra. A substitution is a function $\sigma : Var \rightarrow Ter(\Sigma)$ the domain of which is extended to $Ter(\Sigma)$ by requiring it to be a homomorphism. The set of all substitutions over signature Σ is called $\Theta(\Sigma)$. A binary relation R on $Ter(\Sigma)$ is called *substitutive*, if and only if

$$\forall t, u \in Ter(\Sigma). \forall \theta \in \Theta(\Sigma). t R u \Rightarrow \theta(t) R \theta(u)$$

This concept is typically used to form the *substitutive closure* of a relation.

Subterm positions are finite sequences of natural numbers. The empty sequence is denoted $\langle \rangle$, otherwise $i \cdot p$ is the prefixing a sequence p with the number (or sequence of numbers) i . The set of positions of a term t , $Pos(t)$ is inductively defined as follows: (i) $\langle \rangle \in Pos(t)$, (ii) $\{i \cdot p \mid 1 \leq i \leq \#(F) \wedge p \in Pos(t_i)\} \subset Pos(F(t_1, \dots, t_{\#(F)}))$.

The *subterm* of t at position p , notation: $t|_p$, is defined as follows:

$$\begin{aligned} t|_{\langle \rangle} &= t \\ F(t_1, \dots, t_n)|_{i \cdot p} &= t_i|_p \end{aligned}$$

Updating a term t at position p with term u , notation: $t[u]_p$, is defined as follows:

$$\begin{aligned} t[u]_{\langle \rangle} &= u \\ F(t_1, \dots, t_n)[u]_{i \cdot p} &= F(t_1, \dots, t_{i-1}, t_i[u]_p, t_{i+1}, \dots, t_n) \end{aligned}$$

For positions p and q we write $p \leq q$ iff $\exists r. p \cdot r = q$, i.e. if p is a prefix of q . Two positions p and q are *independent*, $p \parallel q$, if $\neg(p \leq q) \wedge \neg(q \leq p)$.

A Σ -context $C[\]$ is a term $t \in Ter(\Sigma)$ together with a position $p \in Pos(t)$, called the *hole*; we often write $C[u]$ instead of $t[u]_p$, keeping the position implicit. An n -ary context $C[\]$ is a term t together with n mutually independent positions. Here, we define

$$t[t_1, \dots, t_{k+1}]_{p_1, \dots, p_{k+1}} = (t[t_1, \dots, t_k]_{p_1, \dots, p_k})[t_{k+1}]_{p_{k+1}}$$

and typically abbreviate $t[t_1, \dots, t_k]_{p_1, \dots, p_k}$ as $C[t_1, \dots, t_k]$, leaving the positions implicit and understanding them to be lexicographically ordered.

A binary relation R on $Ter(\Sigma)$ is called *compatible* iff

$$\forall s, t, u \in Ter(\Sigma). \forall p \in Pos(s). t R u \Rightarrow s[t]_p R s[u]_p$$

Again, this is used mostly to form the compatible closure of a relation; the compatible closure of R is expressed as R^\odot .

A relation R is called a *rewrite relation* [2] if it is both compatible and substitutive. Note that the compatible closure of a substitutive relation is substitutive, and that the substitutive closure of a compatible relation is compatible.

A *rewrite rule* (over Σ) is a pair of terms $t \in Ter(\Sigma)$ and $u \in Ter(\Sigma)$, written $t \rightarrow u$, such that (i) $t \notin Var$, and (ii) $var(u) \subseteq var(t)$. We call t the left-hand and u the right-hand side of the rule. A term is called *linear* if no variable occurs in it more than once; a rewrite rule is called *left-linear* if its left-hand side is linear.

A *term rewriting system* (short: TRS) consists of a signature Σ and a set of rewrite rules R over that signature. The associated ARS of a term rewriting system (Σ, R) is $(Ter(\Sigma), \rightarrow_R)$ where \rightarrow_R is the compatible and substitutive closure of R .

Topological Spaces

Regarding topological and metric spaces we use notation and terminology mostly taken from [23], and occasionally from [5, 15, 22]. A topological space is a set X together with a function $\text{Cl} : \wp(X) \rightarrow \wp(X)$ satisfying the following properties for all $A, B \subseteq X$:

$$\begin{aligned}\text{Cl}(\emptyset) &= \emptyset \\ A &\subseteq \text{Cl}(A) \\ \text{Cl}(\text{Cl}(A)) &= \text{Cl}(A) \\ \text{Cl}(A \cup B) &= \text{Cl}(A) \cup \text{Cl}(B)\end{aligned}$$

Any set of the form $\text{Cl}(A)$ is called *closed*, and a set $B \subseteq X$ is called *open* iff $X \setminus B$ is closed. A function between topological spaces $f : A \rightarrow B$ is called *continuous* iff its inverse image of any closed set is closed, i.e. $f^{-1}(\text{Cl}(V)) = \text{Cl}(f^{-1}(\text{Cl}(V)))$; beware that the two occurrences of Cl on the right-hand side of that equation refer to (potentially) different topologies.

A topology on a set A is called *discrete* iff every subset of A is open (which is the case iff all singleton sets are open). Note that functions between topological spaces are always continuous if their domain is discrete.

A *neighbourhood* of a point $a \in A$ is a set $B \subseteq A$ such that there is an open set C such that $a \in C \subseteq B$. An *accumulation point* is a point $a \in A$ such that every neighbourhood of a is an infinite set. Thus, a topological space is discrete if and only if it has no accumulation points.

A topological space (X, Cl) is called *compact* if for any family of open sets $S_i, i \in I$ such that $X = \bigcup_{i \in I} S_i$ then $X = \bigcup_{i \in J} S_i$ for some finite subset J of I .

Metric Spaces

A *metric space* is a set M together with a function $d : M \times M \rightarrow \mathcal{R}$ satisfying the following formulae [5] for all x, y , and z :

$$\begin{aligned}d(x, y) &= 0 \iff x = y \\ d(y, z) &\leq d(x, y) + d(x, z)\end{aligned}$$

The *open ε -ball* of an element $x \in M$ is the set of elements at distance smaller than ε : $B_\varepsilon(x) = \{y \in M \mid d(x, y) < \varepsilon\}$. The *closed ε -ball* is $\overline{B}_\varepsilon(x) = \{y \in M \mid d(x, y) \leq \varepsilon\}$. We sometimes view these as binary relations, i.e. $x B_\varepsilon y \iff x \in B_\varepsilon(y) \iff d(x, y) < \varepsilon$. Every metric space induces a topological structure: a set $A \subseteq M$ is open iff $\forall a \in A. \exists \varepsilon > 0. B_\varepsilon(a) \subseteq A$.

The *diameter* of a subset S of a metric space (M, d) is the supremum of the distances in S . A metric space is called *bounded* if it has a finite diameter. The metric spaces of interest to this paper all have a diameter of 1.

A function between metric spaces $f : A \rightarrow B$ is called *uniformly continuous* iff there is a function \hat{f} on the strictly positive real numbers such that

$$\forall \varepsilon > 0. B_{\hat{f}(\varepsilon)} ; f \subseteq f ; B_\varepsilon$$

where “;” is relational composition. As we shall see later, this formulation suitably generalises. The function \hat{f} witnesses the uniform continuity of f , and we can w.l.o.g. assume that it is weakly monotonic. Moreover, f is called *non-expansive* iff $\hat{f}(\varepsilon) \leq \varepsilon$, for all ε .

Given a metric space (M, d) , a *Cauchy sequence* in this space is a sequence f of length ω in M such that:

$$\forall \varepsilon > 0. \exists q. \forall m, n. m \geq q \wedge n \geq q \Rightarrow d(f(m), f(n)) < \varepsilon$$

This is equivalent to saying that f is uniformly continuous, with respect to the metric $d(m, n) = |\frac{1}{m} - \frac{1}{n}|$. A sequence f of length ω in M is called *converging* to $a \in M$ if

$$\forall \varepsilon > 0. \exists q. \forall m. m \geq q \Rightarrow d(f(m), a) < \varepsilon$$

and f is called *converging* if an $a \in M$ exists to which f converges. A metric space is *complete* iff every Cauchy sequence converges.

Every metric space has a unique completion, up to isometry, which we will call M^\bullet for this metric space¹. We can characterise the metric completion M^\bullet as follows: (i) M^\bullet is a complete metric space, (ii) there is an isometric embedding $e : M \rightarrow M^\bullet$ and (iii) the closure of $e(M)$ in M^\bullet is M^\bullet (M is *dense*). An *isometric embedding* is a distance-preserving function between metric spaces (note that this implies injectivity); an *isometry* is a bijective isometric embedding.

All points in $M^\bullet \setminus e(M)$ are accumulation points; moreover, these are the only accumulation points in M^\bullet if M is discrete. In the following, M will be regarded as a subset of M^\bullet .

Uniformly continuous functions between metric spaces can be uniquely lifted to their metric completions, i.e. metric completion is a functor on the category of metric spaces (as objects) and uniformly continuous functions (as morphisms). Moreover, a witness function for f^\bullet can be constructed as $\hat{f}^\bullet(\varepsilon) = \hat{f}(\varepsilon)/2$.

Given a set A and a bounded metric space (B, d) the function space $A \rightarrow B$ has the metric $d(f, g) = \sup_{x \in A} d(f(x), g(x))$ (and is itself bounded).

Proposition 1 *Let B be a bounded metric space. Then $(A \rightarrow B)^\bullet \cong A \rightarrow (B^\bullet)$.*

Proof Left to right: if f_n is a Cauchy sequence in $A \rightarrow B$ then $y_n = f_n(x)$ is a Cauchy sequence in B , for any $x \in A$; we can set $f(x) = \lim_{n \rightarrow \infty} f_n(x)$. Right to left: if $f : A \rightarrow (B^\bullet)$ then each $f(x)$ is approximated by a Cauchy sequence y_x in B ; we can construct $f_n : A \rightarrow B$ as $f_n(x) = y_x(m)$ where we pick m such that $\forall k. k \geq m. d(y_x(k), f(x)) < 2^{-n}$. Clearly, $d(f, f_n) \leq 2^{-n}$ and thus the sequence f_n converges to f . \square

This is surely a standard result for metric spaces, but the standard literature [22, 14] only shows the weaker result that the function space is complete if the codomain is — which does not say anything about the completion of a function space whose codomain is not complete.

Proposition 2 *Let A, B and C be metric spaces such that C is bounded.*

If $f : A \rightarrow (B \rightarrow C)$ is uniformly continuous and pointwise uniformly continuous (i.e. each $f(x)$ is uniformly continuous in $B \rightarrow C$) then there is a unique function $f^\bullet : A^\bullet \rightarrow (B^\bullet \rightarrow C^\bullet)$ extending f that is continuous and pointwise continuous.

¹ A traditional notation is M^* but this would lead to notational clashes when we lift relations to metric completions.

Proof Metric completion extends f to a function $f' : A^\bullet \rightarrow (B \rightarrow C)^\bullet$ which we can view (using proposition 1) as a function $f'' : A^\bullet \rightarrow (B \rightarrow C^\bullet)$. The function f'' is still uniformly continuous and it is still pointwise continuous — because Cauchy sequences of continuous functions converge to continuous functions, see [22, page 209]. Given a function $cB : B^\bullet \rightarrow (\mathcal{N} \rightarrow B)$ that maps every b to some Cauchy sequence that converges to b , and a fixed $a \in A^\bullet$ we can construct functions $f_n : B^\bullet \rightarrow C^\bullet$ as follows: $f_n(b) = f''(a)(cB(b)(\lceil \frac{1}{cB(b)(f''(2^{-n}))} \rceil))$. The construction ensures that $d(f_n(b), f_{n+k}(b)) < 2^{-n}$ (independent of b) which makes it a Cauchy sequence in $B^\bullet \rightarrow C^\bullet$, and as this space is complete the sequence converges to a limit: this limit is $f^\bullet(a)$; it is itself continuous, because the sequence f_n is continuously convergent [14, 28.9.5]. The uniqueness of f^\bullet follows from more general topological properties [15, page 54]. \square

Note: the premise in proposition 2 is strictly weaker than to require that f is uniformly continuous as a function of type $A \times B \rightarrow C$, because the category of metric spaces and uniformly continuous functions is not Cartesian-closed. In particular, function application itself is (in general) not uniformly continuous in $[A \rightarrow B] \times A \rightarrow B$, if we understand $[A \rightarrow B]$ to be the set of uniformly continuous functions from A to B .

A metric space (M, d) is called an *ultra-metric* space if it satisfies the stronger inequality $\forall x, y, z. d(y, z) \leq \max(d(x, y), d(x, z))$. In an ultra-metric space (M, d) , each B_ε becomes an equivalence relation, and a sufficient condition for an ω -sequence to be Cauchy is that the distances between adjacent elements converge to 0. If M is an ultra-metric space then so is M^\bullet .

The category of ultra-metric spaces and non-expansive functions is Cartesian-closed [23]. The product construction derived from that gives also categorical products in the categories of metric spaces (objects) with either continuous, uniformly continuous, or non-expansive functions as morphisms: the product of metric spaces (A, d_A) and (B, d_B) is $(A \times B, d_{A \times B})$ where

$$d_{A \times B}((a_1, b_1), (a_2, b_2)) = \max(d_A(a_1, a_2), d_B(b_1, b_2)).$$

For every set there is a metric d_α defined as $d_\alpha(t, u) = 1 \iff t \neq u$, the so-called *discrete metric* or *trivial metric*; the resulting metric space is always complete with the discrete power-set topology. The converse does not hold, i.e. there are many other metrics with discrete topologies (sometimes, the literature calls them discrete as well), and we shall encounter several of them in this paper.

Since every subset of an (ultra-) metric space gives rise to an (ultra-) metric space the notion of compactness generalises to arbitrary subsets of an (ultra-) metric space. In this context, compact sets are always closed.

3 Ultra-Metric Maps

A function f on the non-negative reals is called *metric-preserving* [6] if for any metric space (M, d) , $(M, f \circ d)$ is a metric space as well. All metric-preserving functions are *amenable*: $f(x) = 0 \iff x = 0$ and *subadditive*: $f(x + y) \leq f(x) + f(y)$. If f is subadditive then $f(n \cdot x) \leq n \cdot f(x)$ and $\frac{f(x)}{n} \leq f(\frac{x}{n})$ for any positive

integer n . Notice that if f is a metric-preserving function then the identity function id is a uniformly continuous function from $(M, f \circ d)$ to (M, d) with $\widehat{\text{id}} = f$. Monotonicity is a sufficient (but not necessary) condition for an amenable and subadditive function to be metric-preserving.

Examples are: the ceiling function $\lceil _ \rceil$ (defined as $\lceil r \rceil = \min\{n \mid n \in \mathcal{N}, n \geq r\}$), and for any $0 < \alpha < 1$ both multiplication and exponentiation with α , e.g. halving and square root.

For any function f on $[0, 1]$ we can define $f^\#(y) = \sup\{x \mid f(x) \leq y\}$. By construction, $f^\#$ is weakly monotonic. We also have $f^\#(f(x)) \geq x$, and this becomes an equality if f is strictly monotonic. If f is continuous then $f(f^\#(x)) \leq x$, and this becomes an equality if f is surjective (or, by the intermediate value theorem, if $f(0) \leq x \leq f(1)$). If f is monotonic and continuous then f and $f^\#$ form a Galois connection.

Notice that if f is continuous and amenable then $f^\#$ is amenable. If f is a continuous metric-preserving function then the identity function is also uniformly continuous in the other direction, from (M, d) to $(M, f \circ d)$ (making the two metrics equivalent), with $\widehat{\text{id}} = f^\#$ — provided f is injective; otherwise, we can set the witness $\widehat{\text{id}}(\varepsilon) = f^\#(\varepsilon/2)$. The ceiling function is not continuous, our other examples of metric-preserving functions are.

Of particular interest in this paper are functions that preserve *ultra-metrics*. These are functions that are amenable and monotonic, they need in general not be subadditive, e.g. squaring is not metric-preserving as it is not subadditive, but it preserves ultra-metrics.

An *ultra-metric map* (short: umm) is an n -ary function $m : [0, 1]^n \rightarrow [0, 1]$ such that (i) it is monotonic: $x_1 \leq y_1 \wedge \dots \wedge x_n \leq y_n \Rightarrow m(x_1, \dots, x_n) \leq m(y_1, \dots, y_n)$ and (ii) it is amenable, i.e. $m(x_1, \dots, x_n) = 0 \iff x_1 = 0 \wedge \dots \wedge x_n = 0$. Ultra-metric maps are closed under composition, i.e. if f is an n -ary umm and k_1, \dots, k_n are p -ary umms then $f \circ \langle k_1, \dots, k_n \rangle$ is a p -ary umm.

The concept of being “subadditive” is extended to n -ary functions as follows: f is subadditive iff $\forall a_1, \dots, a_n, b_1, \dots, b_n. f(a_1 + b_1, \dots, a_n + b_n) \leq f(a_1, \dots, a_n) + f(b_1, \dots, b_n)$.

The *components* of an n -ary umm f are the functions $\tilde{f}_i : [0, 1] \rightarrow [0, 1]$ defined as $\tilde{f}_i(x) = f(0^{i-1}, x, 0^{n-i})$. Each component is itself a umm. The *kernel* of a umm is the function $\tilde{f}(x_1, \dots, x_n) = \max_{1 \leq i \leq n} \tilde{f}_i(x_i)$. Again, the kernel of a umm is itself a umm. Each umm is pointwise greater or equal than its kernel; a umm is called *simple* iff it is equal to its kernel. Thus, a simple umm is determined by its components.

For an n -ary ($n > 0$) umm f we set $f^\#(y) = \sup\{(x_1, \dots, x_n) \mid f(x_1, \dots, x_n) \leq y\}$. We still have $f^\#(f(x_1, \dots, x_n)) \geq (x_1, \dots, x_n)$. If f is continuous then $f^\#$ is amenable in the extended sense that $f^\#(x) = (0, \dots, 0) \iff x = 0$. If f is a simple umm then $f^\#(x) = (f_1^\#(x), \dots, f_n^\#(x))$, and in addition $f(f^\#(x)) \leq x$ if f is also continuous, which again gives us a Galois connection.

For n -ary f the unary function Δf is defined as $\Delta f(x) = f(x, \dots, x)$. If f is continuous then so is Δf and thus $\Delta f^\#$ is amenable.

4 Term Metrics

A *term metric* for a signature Σ is a Σ -algebra m where the carrier set is $[0, 1]$ and each F_m is an ultra-metric map. A term metric is called *simple* (continuous, subadditive) iff all its F_m are simple (continuous, subadditive) ultra-metric maps. The reason for the name “term metric” is that this gives rise to a distance function d_m on $Ter(\Sigma)$ as follows:

$$\begin{aligned} d_m(t, t) &= 0 \\ d_m(t, u) &= 1, \text{ if } \text{root}(t) \neq \text{root}(u) \\ d_m(F(t_1, \dots, t_n), F(u_1, \dots, u_n)) &= F_m(d_m(t_1, u_1), \dots, d_m(t_n, u_n)) \end{aligned}$$

The final equation also means that d_m is a Σ -algebra homomorphism, from the product algebra $Ter(\Sigma) \times Ter(\Sigma)$ to m ; this implies:

$$d_m(\theta(t), \sigma(t)) = \llbracket t \rrbracket_m^{x \mapsto d_m(\theta(x), \sigma(x))}.$$

The function d_m is indeed a distance in the metric sense, even more:

Proposition 3 *Each d_m is an ultra-metric on $Ter(\Sigma)$, bounded by 1.*

Proof By induction on the term structure. First, that d_m is bounded by 1 is trivial by construction.

Second, we need to show that $d_m(t, u) = 0 \iff t = u$. The first and second equation clearly comply. For the third, we assume that the property holds on the subterms of t and u . The result follows from F_m being amenable.

Finally, the strong triangular property. Consider three terms a , b and c . We have to show $d_m(a, b) \leq \max(d_m(a, c), d_m(b, c))$. If a and b have different root symbols then the root of c is different to at least one of the two, take w.l.o.g. that to be b . In this case the inequation to be proven becomes $1 \leq \max(d_m(a, c), 1)$ which is trivially true.

Now suppose that a and b have the same root symbol. If the root of c differs from that then the inequation becomes $d_m(a, b) \leq 1$ which is always true. Finally, assume that all 3 terms have the same root symbol F and that $n = \#(F)$

$$\begin{aligned} &\max(d_m(a, c), d_m(b, c)) \\ &= \max(F_m(d_m(a_1, c_1), \dots, d_m(a_n, c_n)), F_m(d_m(b_1, c_1), \dots, d_m(b_n, c_n))) \\ &= F_m(\max(d_m(a_1, c_1), d_m(b_1, c_1)), \dots, \max(d_m(a_n, c_n), d_m(b_n, c_n))) \\ &\geq F_m(d_m(a_1, b_1), \dots, d_m(a_n, b_n)) \\ &= d_m(a, b) \end{aligned}$$

□

Ultra-metric Σ -algebras have some applications in domain theory [24, section 4.3] but these constrain algebra operations to be non-expansive, and this is not always satisfied by the operations and metrics of interest here.

For any term metric m there is another Σ -algebra \underline{m} , also with carrier set $[0, 1]$, where its operations are defined as follows:

$$F_{\underline{m}}(a_1, \dots, a_n) = \min_{1 \leq i \leq n} \tilde{F}_{m,i}(a_i)$$

Here, $\tilde{F}_{m,i}$ is the i -th component of F_m , and the minimum is set to be 1 if $n = 0$. Although n -ary functions in \underline{m} are not amenable, it is the case that $\llbracket t \rrbracket_{\underline{m}}^{\rho} > 0$, provided $\rho(x) > 0$ for all x . Algebra \underline{m} is useful for propagating distances:

Lemma 1 $F_m(a_1, \dots, a_n) < F_{\underline{m}}(b_1, \dots, b_n) \Rightarrow \forall 1 \leq i \leq n. a_i < b_i$

Proof

$$\begin{aligned} F_m(a_1, \dots, a_n) < F_{\underline{m}}(b_1, \dots, b_n) &\implies \tilde{F}_m(a_1, \dots, a_n) < F_{\underline{m}}(b_1, \dots, b_n) \\ &\iff \max_{1 \leq i \leq n} \tilde{F}_{m,i}(a_i) < F_{\underline{m}}(b_1, \dots, b_n) \\ &\iff \max_{1 \leq i \leq n} \tilde{F}_{m,i}(a_i) < \min_{1 \leq i \leq n} \tilde{F}_{m,i}(b_i) \\ &\implies \forall 1 \leq i \leq n. \tilde{F}_{m,i}(a_i) < \tilde{F}_{m,i}(b_i) \\ &\implies \forall 1 \leq i \leq n. a_i < b_i \end{aligned}$$

□

In particular, one can use \underline{m} to show that all terms are discrete points, for any term metric m .

Proposition 4 *Every term metric m gives rise to a discrete topology.*

Proof The statement means that for every term t there is a constant $c > 0$ such that $d_m(t, u) < c \Rightarrow t = u$. We can set $c = \llbracket t \rrbracket_{\underline{m}}^{k1}$, where $k1$ is the constant-1 function.

If t is a variable or a function symbol with of arity 0 then $c = 1$. Anything closer than that distance has the same root symbol as t and hence is equal to t .

If $t = F(t_1, \dots, t_n)$ then $\llbracket t \rrbracket_{\underline{m}}^{k1} = F_{\underline{m}}(\llbracket t_1 \rrbracket_{\underline{m}}^{k1}, \dots, \llbracket t_n \rrbracket_{\underline{m}}^{k1})$. Now, if $d_m(t, u) < c \leq 1$ then u has the form $F(u_1, \dots, u_n)$, and we get overall, using lemma 1 and the induction hypotheses on the subterms:

$$\begin{aligned} d(t, u) < \llbracket t \rrbracket_{\underline{m}}^{k1} &\iff d_m(F(t_1, \dots, t_n), F(u_1, \dots, u_n)) < \llbracket t \rrbracket_{\underline{m}}^{k1} \\ &\iff F_m(d_m(t_1, u_1), \dots, d_m(t_n, u_n)) < F_{\underline{m}}(\llbracket t_1 \rrbracket_{\underline{m}}^{k1}, \dots, \llbracket t_n \rrbracket_{\underline{m}}^{k1}) \\ &\implies \forall i. 1 \leq i \leq n \Rightarrow d_m(t_i, u_i) < \llbracket t_i \rrbracket_{\underline{m}}^{k1} \\ &\iff \forall i. 1 \leq i \leq n \Rightarrow t_i = u_i \\ &\iff F(t_1, \dots, t_n) = F(u_1, \dots, u_n) \end{aligned}$$

□

Each metric d_m gives us a metric completion $Ter^m(\Sigma)$ of the metric space $(Ter(\Sigma), d_m)$, adding the limits of Cauchy sequences. It depends on the term metric which, if any, “infinite terms” are added by this process.

Here are some examples of term metrics that people have used before, albeit not expressed in the framework presented here.

- The term metric ∞ sets $F_{\infty}(a_1, \dots, a_n) = \frac{1}{2} \cdot \max_{1 \leq i \leq n} a_i$; the set $Ter^{\infty}(\Sigma)$ contains “all” infinite terms: it is terminal in the category of $Ter^m(\Sigma)$ objects and (partial) Σ -homomorphisms as morphisms.
- The trivial term metric id sets $F_{id}(a_1, \dots, a_n) = \max_{1 \leq i \leq n} a_i$; distance between any two distinct terms is 1, and hence $(Ter(\Sigma), d_{id})$ is already complete; thus this gives no infinite terms at all.

- A “lazy signature” (Σ, Λ) [19,18] has a predicate $\Lambda \subset \mathcal{F} \times \mathcal{N}$ such that $\Lambda(F, i)$ signifies that F is “lazy” in its i -th argument position. Given this, the simple term metric g is defined through its components: $\tilde{F}_{g,i}(a) = a/2$ if $\Lambda(F, i)$, and $\tilde{F}_{g,i}(a) = a$, otherwise. The term universe $Ter^g(\Sigma)$ corresponds very closely to the set of constructor values we get in Haskell [16] when any constructor F is declared to be strict in all argument positions for which $\Lambda(F, i)$ does not hold.

These are not the only sensible possibilities, in particular one could restrict the set of infinite terms further than d_g manages. We define the following term metrics as variants of d_g : they are all simple term metrics, defined through their components; we assume a laziness predicate Λ and keep $\tilde{F}_{m,i}(a) = (a/2)$ for lazy argument positions. Otherwise, if $\neg\Lambda(F, i)$ we set

- for term metric c : $\tilde{F}_{c,i}(a) = \lceil a \rceil$;
- for the term metric r : $\tilde{F}_{r,i}(a) = \sqrt{a}$;
- for the term metric d : $\tilde{F}_{d,i}(a) = \min(2 \cdot a, 1)$.

Under the term metric c the subterms at a strict position must be finite terms. Under term metric r we can only iterate a context $C[\]$ to create an infinite term $C^\infty = C[C^\infty]$ if all argument positions leading to the hole of C are lazy. Under term metric d any infinite path through an infinite term must (eventually) cross arbitrarily more lazy than strict argument positions.

All these examples are simple term metrics, i.e. their umms are equal to their kernels. What simple term metrics have in common is that terms can “grow” independently in independent positions; in particular, if for a binary context $C[\]$ the terms $C[t, x]$ and $C[y, u]$ exist in $Ter^m(\Sigma)$ then so does $C[t, u]$. It is possible to define a non-simple term metric for which this is not true. The examples are also all subadditive — a property that is essential for certain aspects of infinitary rewriting.

Proposition 5 *For term metric m , function symbol F is uniformly continuous on the metric space $(Ter(\Sigma), d_m)$ if F_m is continuous.*

Proof The domain for an n -ary F is $(Ter(\Sigma), d_m)^n$ which is the metric space $(Ter(\Sigma)^n, d_m^n)$, where $d_m^n(\langle a_1, \dots, a_n \rangle, \langle b_1, \dots, b_n \rangle) = \max_{1 \leq i \leq n} d_m(a_i, b_i)$. We can simply construct the witness function \hat{F} as follows: $\hat{F}(\varepsilon) = \Delta F_m^\#(\varepsilon/2)$. To show that this is a uniformity witness:

$$\begin{aligned}
& d_m^n(\langle a_1, \dots, a_n \rangle, \langle b_1, \dots, b_n \rangle) < \Delta F_m^\#(\varepsilon/2) \\
& \iff \forall i. 1 \leq i \leq n \Rightarrow d_m(a_i, b_i) < \Delta F_m^\#(\varepsilon/2) \\
& \implies F_m(d_m(a_1, b_1), \dots, d_m(a_n, b_n)) \leq \Delta F_m(\Delta F_m^\#(\varepsilon/2)) \\
& \iff d_m(F(a_1, \dots, a_n), F(b_1, \dots, b_n)) \leq \Delta F_m(\Delta F_m^\#(\varepsilon/2)) \\
& \implies d_m(F(a_1, \dots, a_n), F(b_1, \dots, b_n)) \leq \varepsilon/2 < \varepsilon
\end{aligned}$$

□

Continuity of F_m is only used in the proof for the step $\Delta F_m(\Delta F_m^\#(\varepsilon/2)) \leq \varepsilon/2$ — thus continuity of ΔF_m suffices as condition; in fact, it even suffices if ΔF_m is merely continuous at 0, but for this claim the proof would need a different witness function.

Corollary 1 *If m is a continuous term metric for signature Σ then $Ter^m(\Sigma)$ is a Σ -algebra.*

Proof Since all term-building functions are uniformly continuous they uniquely lift to the completed metric spaces. \square

Notice that proposition 5 views n -ary functions as functions from the n -ary product space. Proposition 2 would suggest to use curried functions instead, for lifting n -ary functions to the metric completion; the difference is only slight though as this corresponds to replacing each algebra-map of the term metric with its kernel.

Non-continuous term metrics are not necessarily regarded as “evil”, and may well be worth serious study — but they are certainly more awkward to work with. If (unary) F_m is not continuous at 0 then the function symbol F can indeed not be applied to infinite terms, in the following sense: if $Ter^m(\Sigma)$ contains an accumulation point u and $g : \mathcal{N} \rightarrow Ter(\Sigma)$ is any Cauchy sequence converging to u then $F \circ g$ is never a Cauchy sequence. An example for a non-continuous term metric is d_c .

5 Operations on Infinitary Terms

Given a signature Σ and a term metric m , the *infinitary terms* are the elements in $Ter^m(\Sigma)$; the *infinite terms* are the accumulation points in $Ter^m(\Sigma)$. Operations on infinitary terms are mostly defined here as uniformly continuous functions operating on finite terms, which thus have a unique lifting.

In particular, this applies to the function $[[_]]_{[m]} : Ter(\Sigma) \rightarrow (Var \rightarrow m) \rightarrow m$: this function is non-expansive for simple m and then it is justified to use this notation for infinitary terms as well.

The set of ε -positions of a term $t \in Ter^m(\Sigma)$, $Pos_\varepsilon^m(t)$ is defined as $\bigcap \{Pos(u) \mid u \in Ter(\Sigma), d_m(t, u) < \varepsilon\}$. All positions of a term, $Pos^m(t)$ are the union of these: $\bigcup \{Pos_\varepsilon^m(t) \mid \varepsilon > 0\}$.

The set of infinitary substitutions is defined as $\Theta(\Sigma)^m = Var \rightarrow Ter^m(\Sigma)$. The definition of d_m is extended to substitutions using the function space metric, i.e. $d_m(\sigma, \theta) = \sup_{x \in Var} d(f(x), g(x))$. Notice that it can make a difference here whether we regard the domain of substitutions as terms or as variables; in the former case some metrics (such as d_r and d_c) would make the substitution space discrete.

Substitution application on *finite terms* (as an operation in $\Theta(\Sigma) \rightarrow Ter(\Sigma) \rightarrow Ter(\Sigma)$) is pointwise non-expansive and, provided the metric is continuous, also uniformly continuous. Proposition 2 then allows us to lift substitution application uniquely to infinitary terms and infinitary substitutions. Notice that this is a case in which it would *not* suffice to consider substitution application as a function in $\Theta(\Sigma) \times Ter(\Sigma) \rightarrow Ter(\Sigma)$, because there are continuous term metrics (an example is term metric r) for which substitution application is not uniformly continuous in this domain.

W.r.t. to non-continuous metrics, substitution application can be undefined, e.g. under metric d_c when $t = F(x)$, F is strict in its argument and θ maps x to an infinite term. Nevertheless it is still possible to view it as a partial function

— where $\theta(t)$ is defined iff a sequence $\theta_n(t)$ converges, where θ_n is a Cauchy sequence of finitary substitutions converging to θ .

A relation R on infinitary terms is called *substitutive* if $t R u$ implies $\theta(t) R \theta(u)$ for all $\theta \in \Theta(\Sigma)^m$ for which $\theta(t)$ and $\theta(u)$ are defined.

For (finite) unary contexts the metric d_m is extended as follows:

$$d_m(C[\]_p, D[\]_q) = \max(d_m(C[x], D[x]), d(p, q))$$

where the metric on positions is discrete, i.e. different positions are at distance 1. Context application, seen as a function in $(Ter(\Sigma) \times \mathcal{N}^*) \rightarrow Ter(\Sigma) \rightarrow Ter(\Sigma)$, is non-expansive and, provided the metric is continuous, pointwise uniformly continuous. Again this allows to apply proposition 2 and generalise context application to infinitary terms; for non-continuous metrics context application is a partial function. An infinitary context is an element in $(Ter(\Sigma) \times \mathcal{N}^*)^\bullet$ (which is isomorphic to $Ter^m(\Sigma) \times \mathcal{N}^*$).

For every n -ary finitary context $C[\]$ there is an n -ary ultra-metric map C_m defined as follows: $C_m(a_1, \dots, a_n) = \llbracket C[x_1, \dots, x_n] \rrbracket_m^{\rho[x_i \mapsto a_i]}$, where the variables x_i do not occur in $C[\]$, $\rho(x) = 0$ for all x , and the notation $\rho[x_i \mapsto a_i]$ updates ρ at these variables. Clearly:

$$d_m(C[t_1, \dots, t_n], C[u_1, \dots, u_n]) = C_m(d_m(t_1, u_1), \dots, d_m(t_n, u_n)).$$

If m is a continuous term metric then C_m is uniformly continuous. Notice that for unary contexts, $C_m(d_m(t, u)) = d_m(C[t], C[u]) < C_m(\varepsilon)$ implies $d_m(t, u) < \varepsilon$.

We can express the property that two terms do not differ up to a certain position formally as follows. There is a family of equivalence relations $\overset{p}{\sim}$ (indexed by positions p), defined as follows:

$$\begin{aligned} t &\overset{\langle \rangle}{\sim} u \\ F(t_1, \dots, t_n) &\overset{i \cdot q}{\sim} F(u_1, \dots, u_n) \iff t_i \overset{q}{\sim} u_i \end{aligned}$$

We also write $C \sim D$ for contexts C and D , to express (i) their holes are at the same position p , and (ii) $C[x] \overset{p}{\sim} D[x]$.

Proposition 6 *Let C, D be contexts such that $C \sim D$.*

- (i) $C_m = D_m$;
- (ii) $d_m(C[t], C[u]) \leq d_m(C[t], D[u])$;
- (iii) $d_m(C[x], D[x]) \leq d_m(C[t], D[u])$.

Proof Straightforward induction on the depth of C . □

Lemma 2 *If $d_m(C[t]_p, u) < C_m(\varepsilon)$ then there is a context $D \sim C$ and a term u' such that $u = D[u']$ and $d_m(t, u') < \varepsilon$.*

Proof The proof goes by induction on the length of p . The context $D[\]$ is $u[\]_p$. The base case $p = \langle \rangle$ is trivial. Otherwise $p = i \cdot q$, and let $C'[\]_q = C[\]_i$. Since $d_m(C[t]_p, u) < C_m(\varepsilon) \leq 1$, the term u must have the same root symbol as $C[\]$, call it F . Hence $d_m(C[t]_p, u) = F_m(a_1, \dots, a_n)$ where $a_j = d_m(C[t]_j, u|_j)$. In particular, $C[t]_i = C'[t]$. This implies: $(\tilde{F}_m)_i(a_i) \leq d_m(C[t]_p, u) < C_m(\varepsilon) = (\tilde{F}_m)_i(C'_m(\varepsilon))$. By monotonicity: $a_i = d_m(C'[t], u|_i) < C'_m(\varepsilon)$. By induction hypothesis $u|_i[\]_q \sim C'$ and there is a term u' such that $u|_i = D'[u']$ and $d_m(t, u') < \varepsilon$. This also implies $u[\]_p \sim C$ by definition of this relation. □

The purpose of the rather awkward looking lemma 2 is to reason about distances in the situation where we put a redex inside a context and then move away from the result by a specific distance.

Corollary 2 *Let $C[\]_p, D[\]_q$ be finite. If $d_m(C[\]_p, D[\]_q) < C_m(1)$ then $C_m = D_m$.*

Consequently, infinitary contexts also have a unique umm: let $D[\]_p$ be a (unary) infinitary context, then $D_m = C_m$ for any finite context $C[\]_p$ close enough to $D[\]_p$, i.e. $d_m(C[x], D[x]) < d_m(D[x], D[y])$.

A relation R on infinitary terms is called *compatible* if $a R b$ implies $C[a] R C[b]$ for all contexts $C[\]$ for which both $C[a]$ and $C[b]$ are defined. The notation R^\odot denotes the compatible closure of R .

6 Infinitary Rules

A term $t \in \text{Ter}^m(\Sigma)$ is called a *pattern* if there is a constant $t^m > 0$ such that:

$$\forall u \in \text{Ter}^m(\Sigma). \forall \sigma \in \Theta(\Sigma)^m. d_m(u, \sigma(t)) \leq t^m \Rightarrow \exists \theta \in \Theta(\Sigma)^m. \theta(t) = u$$

Notice that patterns are necessarily finite terms, because infinite terms are arbitrarily close to some finite terms. The constant t^m (if it exists) is the same as the one constructed in the proof of proposition 4. Moreover, for continuous term metrics, patterns must be linear terms: if a pattern were of the form $C[x, x]$ then $C[t, t]$ can be made arbitrarily close to $C[t, t']$ with $t \neq t'$, and $C[t, t']$ is not a substitution instance of $C[x, x]$. In non-continuous term metrics non-linear patterns are possible, provided (and for simple metrics this is a sufficient condition) that each repeated variable occurs somewhere in a non-continuous position.

An *infinitary rewrite rule* (over Σ , w.r.t. to term metric m) is a pair of terms $t \in \text{Ter}^m(\Sigma)$ and $u \in \text{Ter}^m(\Sigma)$, written $t \rightarrow u$, such that (i) $t \notin \text{Var}$ is a pattern, (ii) $\llbracket u \rrbracket_{\llbracket m \rrbracket} \leq \llbracket t \rrbracket_{\llbracket m \rrbracket}$, where the partial order \leq on functions is the pointwise order, inherited from $[0, 1]$, and (iii) if m is not continuous or not simple then u is finite.

Explanation: for the discrete term metric *id* the second condition is equivalent to the familiar constraint for finite rules that all variables of the right-hand side occur on the left-hand side as well. Moreover, this condition is indeed implied by (ii) for *any* term metric m : suppose some variable x occurred in u but not in t . Consider the function $f : \text{Var} \rightarrow [0, 1]$ with $f(x) = 1, x \neq y \Rightarrow f(y) = 0$: then $\llbracket t \rrbracket_{\llbracket m \rrbracket}(f) = 0$ but $\llbracket u \rrbracket_{\llbracket m \rrbracket}(f) > 0$.

Such a semantic re-interpretation of the condition that the variables of the right occur on the left is not new, see [17] for the situation in higher-order rewriting.

Condition (ii) can be difficult to check, especially for non-simple term metrics. Simple term metrics allow to check this condition variable by variable: in that case, each variable is associated with a unary umm, and these have to be “larger” on the left. In term metric ∞ the condition can be expressed as follows: if $u|_p \in \text{Var}$ then there exists $q \in \text{Pos}(t)$ such that $t|_q = u|_p$ and the length of p is not shorter than the length of q . As a consequence, “collapsing rules” (where u is itself a variable) are not allowed under term metric ∞ .

Condition (iii) has a double purpose: for non-simple term metrics it ensures that condition (ii) is well-defined; for non-continuous term metrics it ensures that any context that can be applied to (instances of) the left-hand side of a rule can

also be applied to the corresponding instances of its right-hand side. A relation R on $Ter^m(\Sigma)$ is called *context-safe* (*substitution-safe*) if $t R u$ implies that if $C[t]$ ($\theta(t)$) is defined then $C[u]$ ($\theta(u)$) is defined as well.

Lemma 3 *Let $t \rightarrow u$ be an infinitary rule w.r.t. term metric m . The finite relation $\{(t, u)\}$ is substitution-safe, and its substitutive closure is context-safe.*

Proof The lemma is trivial for continuous term metrics (because substitution and context application are then total operations), so assume that m is not continuous and that u is therefore finite (by condition (iii)). Assume that $\theta(t)$ is defined.

Consider a Cauchy-sequence of finitary substitutions θ_n that converges to θ . The sequence $\theta_n(t)$ converges to $\theta(t)$.

For any $i, j \in \mathcal{N}$ and any term s : $d_m(\theta_i(s), \theta_j(s)) = \llbracket s \rrbracket_{\llbracket m \rrbracket}(f_{ij})$, where $f_{ij}(x) = d_m(\theta_i(x), \theta_j(x))$. In particular:

$$d_m(\theta_i(t), \theta_j(t)) = \llbracket t \rrbracket_{\llbracket m \rrbracket}(f_{ij}) \geq \llbracket u \rrbracket_{\llbracket m \rrbracket}(f_{ij}) = d_m(\theta_i(u), \theta_j(u)),$$

where the \geq step follows from condition (ii) of being a rule. This means that $\theta_n(u)$ is itself a Cauchy sequence, and hence $\theta(u)$ is defined.

Now consider a Cauchy sequence $C_n[\]$ of finite contexts approximating C . Then $C_i[\theta_i(t)]$ is a Cauchy sequence approximating $C[\theta(t)]$, where for all but finitely many i $d_m(C_i[\theta_i(s)], C_j[\theta_j(s)]) = C_m(d_m(\theta_i(s), \theta_j(s)))$; because C_m is a umm it is monotonic and hence the distances between elements i and j in sequence $C_i[\theta_i(t)]$ are pointwise greater or equal than the corresponding distances in $C_i[\theta_i(u)]$ which is therefore a Cauchy sequence as well. \square

A relation on $Ter^m(\Sigma)$ is called an *infinitary rewrite relation* if it is lsc, pointwise compact, substitutive and compatible.

An *infinitary term rewrite system* consists of a signature Σ , a term metric m for Σ , and a finite set of infinitary rewrite rules, relative to Σ and m . Its associated ARS is $(Ter^m(\Sigma), \rightarrow_R)$ where \rightarrow_R is the compatible, substitutive and reflexive closure of the relation given by the rules.

The motivation for these definitions has to be delayed for a little while, as some of this rests on a number of technical results, on relations and their interaction with metric completion.

7 Continuous Relations

We would like to lift relations between metric spaces V and W to relations between their metric completions V^\bullet and W^\bullet . To be able to do this in a systematic and unambiguous way, we need some structural properties for such relations which the lifting needs to preserve, in analogy to (uniform) continuity of functions.

There are different notions of continuity for relations around. In a nutshell, the problem is the following: a *function* f between topological spaces is continuous iff f^{-1} maps open sets to open sets, and that is the case iff f^{-1} maps closed sets to closed sets. Relations also have an associated inverse image function, but for them these two conditions are not the same.

In particular, for any relation $R : V \leftrightarrow W$ there is a function $R^+ : 2^W \rightarrow 2^V$ defined as $R^+(X) = V - R^{-1}(W - X) = \{v \in V \mid R(v) \subseteq X\}$. Note that this function

coincides with R^{-1} iff the relation R happens to be a total function: we always have $R^+(R(A)) \supseteq A$ and $R^+(W) = V$. Clearly, R^+ preserves openness iff R^{-1} preserves closedness and vice versa.

To lift relations systematically, we view relations as set-valued functions, moving from a relation $R : V \leftrightarrow W$ to a function $\lceil R \rceil : V \rightarrow 2^W$. Dually, if $f : V \rightarrow 2^W$ we write $\lfloor f \rfloor : V \leftrightarrow W$ for the corresponding relation. Since we are operating with metric spaces, this requires a metric on 2^W . Given a bounded metric space $\mathcal{M} = (M, d_M)$ (we assume w.l.o.g. that the bound is 1), the metric space $2^{\mathcal{M}}$ has as elements the closed subsets of \mathcal{M} , and their distance d_H is defined as follows:

$$\begin{aligned} d_H(A, B) &= \max(d_L(A, B), d_L(B, A)) \\ d_L(\emptyset, \emptyset) &= 0 \\ d_L(\emptyset, B) &= 1, \text{ if } B \neq \emptyset \\ d_L(A, B) &= \sup_{a \in A} \inf_{b \in B} d_M(a, b), \text{ if } A \neq \emptyset \end{aligned}$$

Aside: this is sometimes defined *without* the empty set as member of $2^{\mathcal{M}}$, because that simplifies the rest of the definition. However, that modification would only permit to model relations that are entire.

This is the Hausdorff metric, see [22, page 214], giving us for each bounded metric space \mathcal{M} another metric space $2^{\mathcal{M}}$; this construction extends to a functor: given any function $f : \mathcal{M} \rightarrow \mathcal{P}$ the function $2^f : 2^{\mathcal{M}} \rightarrow 2^{\mathcal{P}}$ is defined by $2^f(X) = \text{Cl}(\{f(x) \mid x \in X\})$. If f is continuous (uniformly continuous, non-expansive, an isometric embedding, an isometry) then so is 2^f .

A fundamental property of the Hausdorff construction is its relation to completeness:

Proposition 7 *Let \mathcal{M} be any bounded metric space. $2^{(\mathcal{M}^\bullet)}$ and $(2^{\mathcal{M}})^\bullet$ are isometric.*

Proof The literature, e.g. [22, page 407] and [14, page 124ff], focuses on showing that the power-set construction preserves metric completeness. However, this means that $2^{(\mathcal{M}^\bullet)}$ and $(2^{(\mathcal{M}^\bullet)})^\bullet$ are isometric — which is only half the proof. For the other half we need to show that $(2^{\mathcal{M}})^\bullet$ and $(2^{(\mathcal{M}^\bullet)})^\bullet$ are isometric as well. First notice that both power-set construction and metric completion (as functors) preserve isometric embeddings; this means (together with the previous argument) that there is an isometric embedding from $(2^{\mathcal{M}})^\bullet$ to $2^{(\mathcal{M}^\bullet)}$, but we need to show that this is onto. For each set $X \in 2^{(\mathcal{M}^\bullet)}$ and each $x \in X$ there is a Cauchy sequence $x' : \mathcal{N} \rightarrow \mathcal{M}$ converging to x , and a function $\hat{x} : \mathcal{R} \rightarrow \mathcal{N}$ with $\forall m, n \in \mathcal{N}. m, n \geq \hat{x}(\varepsilon) \Rightarrow d_M(x'(m), x'(n)) < \varepsilon$. With this we can synchronise all Cauchy sequences and form $X_n = \text{Cl}(\{x'(\hat{x}(2^{-n})) \mid x \in X\})$ which gives us a Cauchy sequence in $2^{\mathcal{M}}$ the limit of which is X . \square

Remark: the reason for using $\hat{x}(2^{-n})$ rather than simply n in the construction of X_n is that the latter could fail to turn X_n into a Cauchy sequence if X is an infinite set. The given construction guarantees that $d_H(X, X_i) \leq 2^{-i}$.

Together with the metric completion functor this gives us a method to lift set-valued continuous functions to metric completions — where we can view them as relations again.

Theorem 1 *Let $f : V \rightarrow 2^W$ be uniformly continuous. There is a unique uniformly continuous function $f' : V^\bullet \rightarrow 2^{(W^\bullet)}$ such that for all $x \in V$ we have $f'(x) \cap W = f(x)$.*

Proof We can always lift f to the metric completions: $f^\bullet : V^\bullet \rightarrow (2^W)^\bullet$. We then post-compose f^\bullet with the isometry from proposition 7 and get the desired map f' . The uniqueness of the metric lifting gives us the uniqueness of f' . \square

The function f and f' can easily fail to coincide on values in V , because sets that are closed in W may no longer be closed in W^\bullet ; for $x \in V$ the value of $f'(x)$ is $\text{Cl}(f(x))$ — the closure of the set $f(x)$ in W^\bullet . This is the reason for the intersection with W in the theorem.

In order to be able to model a relation R as a set-valued function in this topology we need that it is “pointwise closed”, and that its associated function $\lceil R \rceil$ is continuous. These conditions are vacuously satisfied when we consider relations between metric spaces with discrete topologies. A condition that is much better behaved than “pointwise closed” is “pointwise compact” — this uses the same metric on sets. The biggest advantage of this notion is that the relational composition of two relations that are pointwise compact is itself pointwise compact. On discrete topologies a set is compact iff it is finite, and thus a pointwise compact relation (with discrete codomain) is finitely branching.

We would like to express continuity and especially uniform continuity more directly in terms of the relation rather than indirectly through its associated set-valued function. There are a couple of relevant properties of relations. A relation R is called *lower semi-continuous* (short: lsc) iff $R^{-1}(A)$ is open for any open set A . It is called *upper semi-continuous* (short: usc) iff $R^{-1}(A)$ is closed for any closed A . In [13], the lsc relations were called continuous, while [8] reserve the term for relations R that are not only lsc, but also usc, and in addition finitely branching. These terminology decisions are tied to various topologies (or metrics) on power-set domains, for example the exponential topology, which has the same carrier set as the exponential metric, but its topology can differ.

Proposition 8 *Let $R : A \leftrightarrow B$ be lsc and pointwise closed. Let $a : \mathcal{N} \rightarrow A$ and $b : \mathcal{N} \rightarrow B$ be Cauchy sequences converging to $a' \in A$ and $b' \in B$, respectively, such that $\forall n. a(n) R b(n)$. Then $a' R b'$.*

Proof Since R is lsc it is in particular lsc at a' which means that for any $\varepsilon > 0$ there is a $\delta > 0$ such that $d_A(a', x) < \delta$ and $x R y$ implies that there is a b'' such that $a' R b''$ and $d_B(b'', y) < \varepsilon$. Since $d_A(a', a(n))$ converges to 0 there must exist b''_n with $a' R b''_n$ and $d_B(b''_n, b(n)) < 2^{-k}$ for any k . Thus, because R is pointwise closed, we must have $a' R b''$ as well. \square

For the purposes of lifting relations to their metric completion, it will not suffice to merely use *semi-continuous* relations, because even in the special case of continuous functions is the lifting not always possible (or not unique). In other words, a notion of *uniform* (semi-) continuity for relations is needed — which should coincide with uniform continuity of the associated set-valued functions. This is achieved by adapting the earlier notion: R is called *uniformly lsc* iff there is a function \hat{R} on the strictly positive real numbers such that

$$\forall \varepsilon > 0. B_{\hat{R}(\varepsilon)} ; R \subseteq R ; B_\varepsilon$$

It is easy to see that “uniformly lsc” implies “lsc”. The property *uniformly usc* can be expressed as:

$$\forall \varepsilon > 0. R^{-1} ; B_{\hat{R}(\varepsilon)} \subseteq B_\varepsilon / R$$

where $x(S/R)y \iff \forall z. y R z \Rightarrow x S z$, see [3, page 99]. This property is not used in the following.

There is a strong correspondence between uniformly continuous set-valued functions and uniformly lsc relations.

Lemma 4 *If a function $f : V \rightarrow 2^W$ is uniformly continuous then $\lfloor f \rfloor$ is uniformly lower semi-continuous.*

Proof We show that $R = \lfloor f \rfloor$ is uniformly lsc with witness $\hat{R}(\varepsilon) = \hat{f}(\varepsilon)$.

$$\begin{aligned} x (B_{\hat{R}(\varepsilon)} ; R) y &\iff \exists x'. d(x, x') < \hat{R}(\varepsilon) \wedge x' R y \\ &\iff \exists x'. d(x, x') < \hat{f}(\varepsilon) \wedge y \in f(x') \\ &\implies \exists x'. d(f(x), f(x')) < \varepsilon \wedge y \in f(x') \\ &\implies \exists x'. \sup_{v \in f(x')} (\inf_{w \in f(x')} d(v, w)) < \varepsilon \wedge y \in f(x') \\ &\implies \exists x'. (\forall v \in f(x'). \inf_{w \in f(x')} d(v, w) < \varepsilon) \wedge y \in f(x') \\ &\implies \inf_{w \in f(x)} d(y, w) < \varepsilon \\ &\iff \exists w. w \in f(x) \wedge d(y, w) < \varepsilon \\ &\iff \exists w. x R w \wedge w B_\varepsilon y \\ &\iff x (R ; B_\varepsilon) y \end{aligned}$$

The step from line 2 to 3 uses the premise that f is uniformly continuous, 3 to 4 one half of the definition of d_H ; 4 to 5 is an equivalence if $f(x')$ is finite (if the relation is finitely branching) but it is always an implication. \square

An implication in this direction is what we might have expected. Slightly surprisingly, the implication also holds in the other direction:

Lemma 5 *If a relation $R : V \leftrightarrow W$ is uniformly lower semi-continuous and pointwise closed then $\lceil R \rceil$ is uniformly continuous.*

Proof Because R is pointwise closed, $\lceil R \rceil$ indeed inhabits our semantic domain, mapping each element to a closed set.

We first show that $d(a, b) < \hat{R}(\varepsilon)$ implies that $d_L(\lceil R \rceil(b), \lceil R \rceil(a)) \leq \varepsilon$:

$$\begin{aligned} d(a, b) < \hat{R}(\varepsilon) &\implies \forall x. b R x \Rightarrow \exists y. a R y \wedge d(x, y) < \varepsilon \\ &\iff \forall x \in \lceil R \rceil(b). \exists y \in \lceil R \rceil(a). d(x, y) < \varepsilon \\ &\iff \forall x \in \lceil R \rceil(b). \inf_{y \in \lceil R \rceil(a)} d(x, y) < \varepsilon \\ &\implies \sup_{x \in \lceil R \rceil(b)} (\inf_{y \in \lceil R \rceil(a)} d(x, y)) \leq \varepsilon \\ &\iff d_L(\lceil R \rceil(b), \lceil R \rceil(a)) \leq \varepsilon \end{aligned}$$

The first step unravels the relation-algebraic statement of uniformly lsc. The introduction of the supremum can lose precision if R is not finitely branching — this is the reason for the \leq instead of $<$. From this we can now prove the lemma:

$$\begin{aligned} d(a, b) < \hat{R}(\varepsilon) &\iff d(a, b) < \hat{R}(\varepsilon) \wedge d(b, a) < \hat{R}(\varepsilon) \\ &\implies d_L(\lceil R \rceil(b), \lceil R \rceil(a)) \leq \varepsilon \wedge d_L(\lceil R \rceil(a), \lceil R \rceil(b)) \leq \varepsilon \\ &\iff d_H(\lceil R \rceil(b), \lceil R \rceil(a)) \leq \varepsilon \\ &\implies d_H(\lceil R \rceil(b), \lceil R \rceil(a)) < \varepsilon \cdot 2 \end{aligned}$$

Hence, for $f = \lceil R \rceil$, we can set $\hat{f}(\varepsilon) = \hat{R}(\frac{\varepsilon}{2})$, giving us a witness function for the uniform continuity of f . \square

Thus both lemmas together give us the following nice characterisation:

Theorem 2 *A pointwise closed relation $R : V \leftrightarrow W$ (between metric spaces V and W , where W is bounded) is uniformly lsc if and only if its associated set-valued function $\lceil R \rceil : V \rightarrow 2^W$ is uniformly continuous.*

Note: it appears unlikely that something as fundamental as that is a new result, but I could not find it anywhere. Kuratowski's results about the exponential *topology* are ever so slightly different, e.g. in that setting continuous functions are both lsc and usc [22, page 173].

In the following, the notation R^\bullet is also used to describe the lifting of a pointwise closed and uniformly lsc relation R from $V \leftrightarrow W$ to $V^\bullet \rightarrow W^\bullet$.

Proposition 9 *Some useful observations about uniformly lower semi-continuous relations. Uniformly lsc relations are closed under:*

1. *binary union*
2. *composition*
3. *product, i.e. if $p : Z \leftrightarrow A$ and $q : Z \leftrightarrow B$ then $\langle p, q \rangle : Z \leftrightarrow A \times B$.*

Proof Note that in all three cases we need to construct a new witness function as well.

1. Let R and S be uniformly lsc and $T = R \cup S$. We set $\hat{T}(\varepsilon) = \min(\hat{R}(\varepsilon), \hat{S}(\varepsilon))$.

$$\begin{aligned} B_{\hat{T}(\varepsilon)}; T &= (B_{\hat{T}(\varepsilon)}; R) \cup (B_{\hat{T}(\varepsilon)}; S) \subseteq (B_{\hat{R}(\varepsilon)}; R) \cup (B_{\hat{S}(\varepsilon)}; S) \\ &\subseteq (R; B_\varepsilon) \cup (S; B_\varepsilon) = T; B_\varepsilon \end{aligned}$$

2. This time let $T = R; S$, and $\hat{T}(\varepsilon) = \hat{R}(\hat{S}(\varepsilon))$.

$$\begin{aligned} B_{\hat{T}(\varepsilon)}; T &= B_{\hat{R}(\hat{S}(\varepsilon))}; (R; S) = (B_{\hat{R}(\hat{S}(\varepsilon))}; R); S \\ &\subseteq (R; B_{\hat{S}(\varepsilon)}); S = R; (B_{\hat{S}(\varepsilon)}; S) \\ &\subseteq R; (S; B_\varepsilon) = T; B_\varepsilon \end{aligned}$$

3. We set $\widehat{\langle p, q \rangle}(\varepsilon) = \min(\hat{p}(\varepsilon), \hat{q}(\varepsilon))$ and get:

$$\begin{aligned} B_{\widehat{\langle p, q \rangle}(\varepsilon)}; \langle p, q \rangle &= \langle B_{\widehat{\langle p, q \rangle}(\varepsilon)}; p, B_{\widehat{\langle p, q \rangle}(\varepsilon)}; q \rangle \subseteq \langle B_{\hat{p}(\varepsilon)}; p, B_{\hat{q}(\varepsilon)}; q \rangle \\ &\subseteq \langle p; B_\varepsilon, q; B_\varepsilon \rangle = \langle p, q \rangle; B_\varepsilon \end{aligned}$$

□

Although uniform lsc is preserved by binary union, it is not (in general) preserved by arbitrary union. The arbitrary union of lsc relations is always lsc [22, page 179], but uniformity can be lost through that process. We would have to set the witness function $\widehat{\bigcup R_i}(\varepsilon)$ to be $\inf(\widehat{R_i}(\varepsilon))$ and this infimum could be 0. This can already happen when we form the transitive closure of a uniformly lsc relation, because $R^* = \bigcup R^m$ and each R^i is uniformly continuous (a consequence of proposition 9). An example is the function $f : \mathcal{R} \rightarrow \mathcal{R}$ with $f(x) = 2 \cdot x$; this is a uniformly continuous function, relative to the usual distance metric, and hence it clearly is uniformly lsc when viewed as a relation, but the transitive closure of f is not uniformly lsc: for a given ε we would need to find a δ such that $2^m \cdot \delta < \varepsilon$, which is possible for a finite number of m 's but not if we need a δ that works for all m .

An important special case of uniformly semi-continuous relations is the following: a relation R is called *eventually non-expansive below α* (short: *ene*), where $\alpha > 0$, iff $\forall \varepsilon < \alpha. B_\varepsilon ; R \subseteq R ; B_\varepsilon$. It is *strictly ene* if this holds with respect to closed ε -balls: $\forall \varepsilon \leq \alpha. \overline{B_\varepsilon} ; R \subseteq R ; \overline{B_\varepsilon}$. Clearly, if R is ene below α then it is uniformly lsc, with witness function $\hat{R}(\varepsilon) = \min(\varepsilon, \alpha)$. For bounded metrics one can assume w.l.o.g. that $\alpha = \hat{R}(\delta)$, where δ is the diameter of the metric space; in this case we can leave “below α ” implicit. Moreover: if R is ene below α and pointwise closed then R^\bullet is also ene below α .

Relations that are (eventually) non-expansive are also closed under union, composition and finite products, but there are further operations under which they are closed. In particular, non-expansive relations are closed under *arbitrary* union (and infinite products); this is not true for ene relations in general, but the transitive closure of an ene relation is always ene.

8 Continuity of Rewriting

Linking the definitions of the rewrite relations of finite and infinitary term rewriting systems we would expect that two constructions should be strongly related: given a *finite* TRS, the rules of which also match the constraints for *infinitary* TRSs (w.r.t. some term metric m) we can either:

- view it as a finite TRS and lift its (finitary) rewrite relation using theorem 1
- view it as an infinite TRS and construct its (infinitary) rewrite relation directly

This is only meaningful if the relation on finite terms is uniformly lsc, because otherwise there is no canonical lifting; it also has to be pointwise compact, but this condition is implied by the constraint to finitely many rewrite rules.

Proposition 10 *Let $R : Ter(\Sigma) \leftrightarrow Ter(\Sigma)$ be any rewrite relation which is uniformly lsc and pointwise compact. Then R^\bullet is an infinitary rewrite relation.*

Proof The properties of lifting ensure that R^\bullet is uniformly lsc and pointwise compact. It is clearly closed under finite contexts and finite substitutions. The application of infinitary substitutions and contexts arises as the limits of finitary substitution and context application. Then apply proposition 8. □

The substitutive closure is particularly well-behaved:

Lemma 6 *The substitutive closure of a rule $t \rightarrow u$ of an infinitary TRS is strictly ene.*

Proof Let R be the substitutive closure of this rule. It is strictly ene below $\llbracket t \rrbracket_m^{k1}$:

Suppose $a R b$ with $a = \theta(t)$, $b = \theta(u)$. Let $s \in \text{Ter}^m(\Sigma)$ with $d_m(s, a) \leq \varepsilon < \llbracket t \rrbracket_m^{k1}$. Hence $d_m(s, \theta(t)) < \llbracket t \rrbracket_m^{k1}$ and by the pattern property of t this means that $s = \sigma(t)$ for some σ ; so: $s R \sigma(u)$. But $d_m(\sigma(u), \theta(u)) = \llbracket u \rrbracket_{\llbracket m \rrbracket}(f) \leq \llbracket t \rrbracket_{\llbracket m \rrbracket}(f) = d_m(\sigma(t), \theta(t)) \leq \varepsilon$, where f is the function $f(x) = d_m(\sigma(x), \theta(x))$. \square

The proof works unchanged to show that similarly the substitutive closure of a finitary rule is strictly ene as a relation on finite terms, w.r.t. any term metric m for which the rule classifies as an infinitary rule.

Aiming for rewrite relations that are lsc but not uniformly lsc is not very satisfactory, because lsc is a weak property for relations on infinitary terms. One can observe this as follows:

Proposition 11 *Let R be any relation on $\text{Ter}^m(\Sigma)$ that is compatible and reflexive. Then it is lsc as well.*

Proof Let $t \in \text{Ter}^m(\Sigma)$ and $\varepsilon > 0$. We have to find $\delta > 0$ such that $d_m(t, t') < \delta \wedge u R v$ implies that there is a u' with $d_m(u', v) < \varepsilon \wedge t R u' \vee t = u'$. We can set $\delta = \min\{C_m(\varepsilon) \mid p \in \text{Pos}_\varepsilon^m(t), C[\] = t[\]_p\}$, and for this δ it suffices to pick $u' = t$. \square

Explanation: any infinitary term t has only finitely many positions at which changing the subterm at that position deviates from t with ε or more. We can “protect” these positions by translating them into a (safe) distance for t ; applying R at other positions will stay within ε -distance of t .

To get something stronger the metric has to have certain properties.

Proposition 12 *Let R be any strictly ene and context-safe relation on $\text{Ter}^m(\Sigma)$. If m is subadditive then the compatible and reflexive closure of R is uniformly lsc. In particular, if R is strictly ene below α then for all $\varepsilon < \alpha$:*

$$\overline{B_{\frac{\varepsilon}{\lceil 1/\alpha \rceil}}}; R^\odot \subseteq (R^\odot; \overline{B_{\frac{\varepsilon}{\lceil 1/\alpha \rceil}}}) \cup B_\varepsilon$$

Proof Let R be strictly ene below α . Let $S = id \cup R^\odot$. The function witnessing its uniformity is set as

$$\hat{S}(\varepsilon) = \frac{\varepsilon}{\lceil \frac{1}{\alpha} \rceil}$$

To check that this is indeed a uniformity witness: notice first that $\hat{S}(\varepsilon) \leq \varepsilon$. Now suppose $t R u$ and $d_m(C[t], a) < \hat{S}(\varepsilon)$. Since $C[t] S C[u]$, a b needs to be found such that $a S b$ and $d_m(b, C[u]) < \varepsilon$. There are two cases: (i) $C_m(1) < \varepsilon$, (ii) $C_m(1) \geq \varepsilon$.

In case (i) $d_m(C[t], C[u]) = C_m(d_m(t, u)) \leq C_m(1) < \varepsilon$. Because $d_m(C[t], a) < \hat{S}(\varepsilon) \leq \varepsilon$ the ultra-metric property gives $d_m(a, C[u]) < \varepsilon$. Hence we can pick $b = a$ as $a S a$ by reflexivity.

In case (ii) $\varepsilon \leq C_m(1)$; we abbreviate $k = \lceil 1/\alpha \rceil$ and get:

$$C_m(\alpha) \geq C_m\left(\frac{1}{k}\right) \geq \frac{C_m(1)}{k} \geq \frac{\varepsilon}{k} = \hat{S}(\varepsilon)$$

The first inequation holds by monotonicity of C_m (and $\lceil x \rceil \geq x$, for all x); the second follows from subadditivity of C_m , the third is dividing inequation (ii) by k , the fourth is the definition of $\hat{S}(\varepsilon)$. Hence $d_m(C[t], a) < C_m(\alpha)$ and by lemma 2: $a = D[a']$, $D \sim C$ and $d_m(t, a') < \alpha$. Since $\delta = d_m(t, a') < \alpha$ and R is strictly ene below α it follows that $\overline{B_\delta}; R \subseteq R; \overline{B_\delta}$. Thus there is a b' with $a' R b'$ and $d_m(b', u) \leq \delta$. Because R is context-safe, $D[b']$ is defined and we can set $b = D[b']$; clearly $a S b$. Using proposition 6 one can show that $d_m(D[b'], C[u]) < \varepsilon$ as well:

$$\begin{aligned} d_m(D[b'], C[u]) &\leq \max(d_m(C[u], D[u]), d_m(D[b'], D[u])) \\ &= \max(d_m(C[u], D[u]), C_m(d_m(b', u))) \\ &\leq \max(d_m(C[t], D[a']), C_m(\delta)) \\ &= \max(d_m(C[t], D[a']), d_m(C[t], C[a'])) \\ &= d_m(C[t], D[a']) < \varepsilon. \end{aligned}$$

□

Ultra-metric maps that are not subadditive (such as x^2) can prevent the compatible closure (of the substitutive closure) of a single rewrite rule $t \rightarrow u$ to be uniformly lsc. The reason is: (i) some terms s are closer than distance 1 from t without being substitution instances ($1 > d_m(s, t) > \llbracket t \rrbracket_m^{k1}$); (ii) non sub-additive contexts $C[\]$ can make $C_m(d_m(s, t))$ arbitrarily small but keep $C_m(d_m(t, u)) = 1$.

As before, proposition 12 can be adapted for finite terms and relations, but w.r.t. the same metric.

Theorem 3 *Let (Σ, m, R) be an infinitary rewrite system such that m is subadditive. Then the relation \rightarrow_R of its associated ARS is an infinitary rewrite relation.*

Proof By construction \rightarrow_R is substitutive and compatible. It remains to be shown that it is lsc and pointwise compact. By lemma 6 the substitutive closure of a rule is strictly ene and by lemma 3 it is context-safe, which implies by proposition 12 that its compatible and reflexive closure is uniformly lsc. Since there are only finitely many rules the union of their rewrite relations is still uniformly lsc (proposition 9).

To show that \rightarrow_R is pointwise compact it suffices to show that if $A \subseteq \{u \mid t \rightarrow_R u\}$ is infinite then A contains a Cauchy sequence, and that t is \rightarrow_R -related to the limit of that sequence. The elements of A are all of the form $t[a_{p,i}]_p$, for various $p \in \text{Pos}(t)$ where $t|_p$ is related to $a_{p,i}$ by the substitutive closure of R . Since that relation is finitely branching and A is infinite, A must contain $t[a_{p,i}]_p$ for infinitely many different p . Picking one for each p and arranging them by the length of p gives indeed a Cauchy sequence — with limit t , and $t \rightarrow_R t$ by reflexivity. □

The reason why the rewrite relation of an infinitary TRS is required to be reflexive should be clear from the proof of theorem 3: it is useful for showing that the compatible closure is lsc and also that the relation is pointwise compact.

9 Convergence

A *metric abstract reduction system* (short: MARS) is a structure (M, d, \rightarrow) such that (M, d) is a metric space and (M, \rightarrow) is an abstract reduction system. It is

called *converging* (*Cauchy*) iff any reduction sequence of length ω is converging (Cauchy). Observation: if (M, d) has a discrete topology then (M, d, R) is converging iff the relation $R \setminus id_M$ is terminating.

Proposition 13 *Let (M, d, \rightarrow) be a MARS such that (M, d) is an ultra-metric. Then it is Cauchy iff the irreflexive interior of the relations $\rightarrow/\overline{B_\varepsilon}$ (\rightarrow modulo $\overline{B_\varepsilon}$) is strongly normalising for all $\varepsilon > 0$.*

Proof A reduction sequence f of length ω is Cauchy iff for all ε there is an n such that the set $\{f(k) \mid k \geq n\}$ has a diameter of at most ε . In an ultra-metric this is the case iff, for all $k \geq n$, $d(f(k), f(k+1)) \leq \varepsilon$. If a reduction sequence of \rightarrow fails to be Cauchy, it fails for one particular ε , and in an ultra-metric this means that there is a reduction sequence with infinitely many steps of at least ε -distance. \square

Proposition 13 is a useful tool for convergence proofs, because for an infinitary TRS (Σ, m, R) the MARS $(Ter^m(\Sigma), \rightarrow_R)/\overline{B_\varepsilon}$ can be represented as a relation on finite terms — the termination of which can be checked by traditional means. This is based on another observation of the equivalence relations $\overline{B_\varepsilon}$.

Given an infinitary $t \in Ter^m(\Sigma)$, an *open representative of t at ε* is a finite term $u \in Ter(\Sigma)$, such that (i) $u \in \overline{B_\varepsilon}(t)$, and (ii) $\forall v \in \overline{B_\varepsilon}(t). \exists \theta \in \Theta(\Sigma)^m. \theta(u) = v$; notation: $t \searrow_\varepsilon u$ if u is an open representative of t at ε .

Proposition 14 *Let m be a term metric. For any $\varepsilon > 0$ and any $t \in Ter^m(\Sigma)$ there is a u such that $t \searrow_\varepsilon u$.*

Proof Any $\overline{B_\varepsilon}(t)$ contains a finite term u' . Suppose some $t' \in \overline{B_\varepsilon}(t)$ is not a substitution instance of u' then $u' = C[u_1, \dots, u_n]$, $t' = C[t_1, \dots, t_n]$ for some context $C[\]$ where the roots of t_i and u_i are distinct. Consider the term $u'' = C[x_1, \dots, x_n]$ where the variables x_i are fresh. Clearly, $d_m(t', u'') = d_m(u'', u') = d_m(u', t') \leq \varepsilon$, and both t' and u' are substitution instances of u'' . This cannot be repeated infinitely, because u'' is of smaller size than u' (counting function symbols and repeated variable occurrences). \square

The representatives can be used to express reductions on $\overline{B_\varepsilon}$ equivalence classes, and even modified rewrite rules. Some fundamental properties of representatives:

Proposition 15 *Let $t \searrow_\varepsilon u$. Then for all $p \in Pos(u)$: (i) $\forall s \in \overline{B_\varepsilon}(t). s \stackrel{p}{\sim} u$ and (ii) if m is continuous then $t|_p \searrow_\delta u|_p$ where $\delta = C_m^\#(\varepsilon)$ and $C[\] = u[\]_p$.*

Proof (i) is obvious: since all terms in $\overline{B_\varepsilon}(t)$ are substitution instances of u they must have the same function symbols as u up to its variable positions. For (ii) first note that because of (i) and proposition 6 that C_m is not only the context function of $u[\]_p$ but of any $s[\]_p$ with $s \in \overline{B_\varepsilon}(t)$. Second, consider the distance between $t|_p$ and $u|_p$: $\varepsilon \geq d_m(t, u) \geq d_m(t, t|_p) = C_m(d_m(t|_p, u|_p))$. Applying $C_m^\#$ on both sides gives $C_m^\#(C_m(d_m(t|_p, u|_p))) \leq C_m^\#(\varepsilon)$ which implies $d_m(t|_p, u|_p) \leq C_m^\#(\varepsilon)$. Third, consider any term a with $d_m(a, t|_p) \leq C_m^\#(\varepsilon)$; it needs to be shown that a is a substitution instance of $u|_p$. Since $C_m/C_m^\#$ form a Galois connection for continuous m , the premise $d_m(a, t|_p) \leq C_m^\#(\varepsilon)$ implies $C_m(d_m(a, t|_p)) \leq \varepsilon$, hence $d_m(t|_p, a) \leq \varepsilon$. Thus $\theta(u) = t|_p$ for some substitution θ (as u represents t 's ε -ball) and therefore $\theta(u|_p) = a$. \square

In the following, it is assumed (without loss of generality, merely for simplicity of presentation). that signature Σ contains a function symbol \perp of arity 0. The variables of an equivalence class are those that occur in every term, i.e. $\text{var}(\overline{B_\varepsilon}(t)) = \bigcap \{\text{var}(u) \mid d_m(t, u) \leq \varepsilon\}$. A finite term u is a *closed representative* of r , notation $r \downarrow_\varepsilon u$, if and only if:

$$r \downarrow_\varepsilon u \iff \exists s \in \text{Ter}(\Sigma). r \searrow_\varepsilon s \wedge \theta(s) = u \text{ where}$$

$$\theta(x) = \begin{cases} x & \text{if } x \in \text{var}(\overline{B_\varepsilon}(r)) \\ \perp & \text{if } x \in \text{Var} \setminus \text{var}(\overline{B_\varepsilon}(r)) \end{cases}$$

Explanation: by construction, $r \downarrow_\varepsilon u$ ensures that $u \in \overline{B_\varepsilon}(r)$; moreover, the only variables left in u occur in all terms of the class. In fact, the relation \downarrow_ε is a function, u is unique. Given an infinitary rule $t \rightarrow u$ we write $[t \rightarrow u]_\varepsilon$ for a pair of finite terms (t', u') such that $t \searrow_\varepsilon t' \wedge u \downarrow_\varepsilon u'$ and then either $d_m(t', u') > \varepsilon \wedge u' = u''$ or $d_m(t', u') \leq \varepsilon \wedge u' = t'$.

Note: if $t \rightarrow u$ is an infinitary rule w.r.t. term metric m then $[t \rightarrow u]_\varepsilon$ is a finitary rule (for $\varepsilon < 1$), i.e. t' is not a variable and all variables in u' occur in t' as well; the first follows from $\varepsilon < 1$ (for $\varepsilon = 1$ we always have $[t \rightarrow u]_1 = (x \rightarrow x)$), the second follows from condition (ii) of being an infinitary rule.

The rule $t' \rightarrow u'$ simulates the behaviour of applying rule $t \rightarrow u$ at the root of a term: if $a = \theta(t)$ and $b = \theta(u)$ then there are finite terms a' and b' and a finitary substitution θ' such that: $d_m(a, a') \leq \varepsilon$, $d_m(b, b') \leq \varepsilon$, $\theta'(t') = a'$, $\theta'(u') = b'$.

To simulate the behaviour of the compatible closure one can construct the derived rule $[C[t] \rightarrow C[u]]_\varepsilon$ for any context $C[\]$ and any rule $t \rightarrow u$.

Given an iTRS $A = (\Sigma, m, R)$ the notation $[R]_\varepsilon$ stands for the substitutive closure of the following relation on $\text{Ter}(\Sigma)$: $t' [R]_\varepsilon u' \Leftarrow (t' \rightarrow u') = [C[t] \rightarrow C[u]]_\varepsilon$ for some rule $t \rightarrow u \in R$ and some infinitary context $C[\]$. Notice that $[R]_\varepsilon$ is strictly ene below ε : this holds because any term within ε distance of t' is an instance of t' , and so lemma 6 can be applied.

Proposition 16 *Let $A = (\Sigma, m, R)$ be an iTRS, $\varepsilon > 0$ and f be an ω -sequence in $\text{Ter}^m(\Sigma)$ such that $\forall n \in \omega. f(n) \rightarrow_R t_n \overline{B_\varepsilon} f(n+1)$. Then there is an ω -reduction sequence g of the MARS $(\text{Ter}(\Sigma), d_m, [R]_\varepsilon)$ such that $\forall n. d_m(f(n), g(n)) \leq \varepsilon$.*

Proof For $g(0)$ we can pick any finite term a with $d_m(f(0), a) \leq \varepsilon$, e.g. we can set $f(0) \downarrow_\varepsilon a = g(0)$. Thus $d_m(f(0), g(0)) \leq \varepsilon$.

Since $f(n) \rightarrow_R t_n \overline{B_\varepsilon} f(n+1)$ we have that $f(n) = \sigma(C[t])$ and $t_n = \sigma(C[u])$ for some rule $t \rightarrow u$, some context $C[\]$, and some substitution σ .

The relation $[R]_\varepsilon$ contains the rule $[C[t] \rightarrow C[u]]_\varepsilon$, with $C[t] \searrow_\varepsilon t'$ and $C[u] \downarrow_\varepsilon u'$. By definition of \searrow_ε there is a substitution θ with $\theta(t') = C[t]$. Thus also $\sigma(\theta(t')) = f(n)$. Since $C[u] \downarrow_\varepsilon u'$ it follows $d_m(C[u], u') \leq \varepsilon$ and because substitution application is non-expansive also $d_m(t_n, \sigma(u')) = d_m(\sigma(C[u]), \sigma(u')) \leq \varepsilon$; see diagram (top-left triangle). Because $[R]_\varepsilon$ is strictly ene it is uniformly lsc (trivially, it is pointwise closed) and thus can be lifted to infinitary terms — where it remains strictly ene below ε . Since $f(n) [R]_\varepsilon \sigma(u')$ and $d_m(f(n), g(n)) \leq \varepsilon$ there must exist an a_n such that $g(n) [R]_\varepsilon a_n$ and $d_m(a_n, \sigma(u')) \leq \varepsilon$ (bottom-left triangle). We can assume $a_n \in \text{Ter}(\Sigma)$ as all rules in $[R]_\varepsilon$ relate finite terms to finite terms. Overall

we can choose $g(n+1) = a_n$ which gives the following picture:

$$\begin{array}{ccccc}
 f(n) & \xrightarrow{R} & t_n & \xrightarrow{\varepsilon} & f(n+1) \\
 \varepsilon \downarrow & & \varepsilon \downarrow & \nearrow \varepsilon & \varepsilon \downarrow \\
 & & \sigma(u') & & \\
 \varepsilon \downarrow & & \varepsilon \downarrow & \nearrow \varepsilon & \varepsilon \downarrow \\
 g(n) & \xrightarrow{[R]_\varepsilon} & & & g(n+1)
 \end{array}$$

Finally, since d_m is an ultra-metric, \overline{B}_ε is an equivalence relation which allows to conclude $d_m(f(n+1), g(n+1)) \leq \varepsilon$. \square

Corollary 3 *If the rewrite relation of an iTRS (Σ, m, R) is not converging then for some $\varepsilon > 0$ the MARS $(Ter(\Sigma), d_m, [R]_\varepsilon)$ is not Cauchy.*

Proof Any non-converging reduction sequence h has an ε such that (for infinitely many n) $d_m(h(n), h(n+1)) > \varepsilon$; this sequence can be reshaped to match the premise of proposition 16 by combining all consecutive reduction steps within ε -distance as \overline{B}_ε steps. The resulting $[R]_\varepsilon$ steps must preserve the distance. \square

Generally, the number of rules in $[R]_\varepsilon$ is infinite, because there are infinitely many contexts. This does not make it a good candidate for direct proof techniques, especially as ε needs to be chosen as well. However, the relation $[R]_\varepsilon$ can itself be simulated by a *finite* TRS if the term metric is continuous. Given an iTRS $A = (\Sigma, m, R)$ with continuous term metric m the notation $[A]$ stands for the TRS $(\Sigma, \bigcup_{0 < \varepsilon < 1} \{t \rightarrow u\}_\varepsilon \mid (t \rightarrow u) \in R)$.

Forming the union with *all* ε is an over-approximation of what the context functions C_m (and their inverses $C_m^\#$) can do to a specific ε . A desirable side-effect of this construction is that the definition of $[A]$ no longer refers to ε ; thus any ω reduction modulo \overline{B}_ε for any ε can be simulated by the TRS $[A]$.

It is worth illustrating the construction of this TRS at an example. Let R be the following rewrite rule:

$$H(F(x, G(y, z))) \rightarrow K(C, F(D(D(y)), x))$$

considered w.r.t. term metric ∞ . The corresponding TRS $[R]$ consists of the following rules:

$$\begin{aligned}
 H(x_1, x_2) &\rightarrow K(\perp, \perp) \\
 H(F(x_1, x_2)) &\rightarrow K(C, F(\perp, \perp)) \\
 H(F(x, G(y_1, y_2))) &\rightarrow K(C, F(D(\perp), x)) \\
 H(F(x, G(y, z))) &\rightarrow K(C, F(D(D(\perp), x))) \\
 H(F(x, G(y, z))) &\rightarrow K(C, F(D(D(y), x)))
 \end{aligned}$$

These five rules are $[R]_{2^{-k}}$ with k ranging from 1 to 5, in that order. Focusing on ε that are negative powers of two suffices for term metrics ∞ , g and d .

Proposition 17 *Let $A = (\Sigma, m, R)$ be an iTRS with continuous term metric m . If (the MARS of) A is not converging then (the MARS of) $[A]$ is not Cauchy.*

Proof Corollary 3 tells us that $[R]_\varepsilon$ is not Cauchy for some ε ; we are going to show $[R]_\varepsilon \subseteq \rightarrow_{[A]}$. Suppose $a[R]_\varepsilon b$ with $d_m(a, b) > \varepsilon$. Then there is a rule $t \rightarrow u \in R$ and a context $C[\]_p$ and a finitary substitution θ such that $C[t] \searrow_\varepsilon t' \wedge C[u] \downarrow_\varepsilon u'$ and $\theta(t') = a$ and $\theta(u') = b$. In $[A]$ there is the rule $t'' \rightarrow u'' = [t \rightarrow u]_{C_m^\#(\varepsilon)}$. Because m is continuous, proposition 15 implies that $t'|_p$ is a representative of $C[t]_p = t$ at distance $C_m(\varepsilon)$. But so is t'' , and thus there is a variable renaming substitution ρ with $\rho(t'') = t'|_p$. Moreover, $u'' = u'|_p$, because the closed representative is a substitution instance of the open representative, so proposition 15 applies again, and $\rho(u'') = u'$, because $\rho(x) = x$ for all variables $x \in \text{var}(\overline{B_\varepsilon})$. Overall, $a = a[\theta(\rho(t''))]_p \rightarrow_{[A]} a[\theta(\rho(u''))]_p = a[\theta(u'')]_p = b$. \square

10 Proving a TRS to be Cauchy

The contra-positive of proposition 17 gives us a handle to prove an iTRS A to be converging: simply prove that the finite TRS $[A]$ is Cauchy. There is a relatively straightforward technique for managing these proofs which is based on original ideas from [9].

The central idea is the following: If, for a finitary TRS, the substitutive closure of the rules is terminating and no reduction sequence contains infinitely many redex-contractions at position $\langle \rangle$ then the rewrite process moves deeper and deeper inside the terms, and is therefore converging — that is: w.r.t. metric d_∞ . For other metrics the argument does not quite suffice, but it can be adapted: typically termination of a relation is proved by showing that it is included in some other terminating relation $>$. Contexts $C[\]$ that cannot be repeated infinitely many times need to preserve that strict relation $>$.

A umm f is called *shrinking* iff the sequence $a_0 = 1, a_{n+1} = \Delta f(a_n)$ converges to 0. A context $C[\]$ is shrinking (w.r.t. term metric m) iff the umm C_m is shrinking. A term metric m is called *uniform* if the pointwise supremum of all shrinking metric morphisms of the form C_m is itself shrinking.

Proposition 18 *The term metrics $\infty, g, r, d, id,$ and c are all uniform.*

Proof Notice all metric morphisms of concern arise as compositions of the form $f_1 \circ \dots \circ f_k$ where each f_i is either f_m or halving. id has no shrinking umms at all, so the pointwise supremum is the constant 0 function (which is shrinking). For the other metrics, *all* shrinking metric morphisms are multiplications with 2^{-k} for some $k > 0$ (obvious for all but r , see below). Thus, their pointwise supremum is “halving” — which is shrinking.

For r (with $f_r(x) = \sqrt{x}$), we need to show that no shrinking metric morphism involves square root. It suffices to show this for a single occurrence of f_r (one eager position), because these metric morphisms are pointwise lower bounds for the others. Let f be such a metric morphism, i.e. it is of the form $f(x) = 2^{-n} \cdot \sqrt{2^{-k} \cdot x} = 2^{-n-k/2} \cdot \sqrt{x}$, for some fixed k and n . This function fails to shrink for $x \leq 2^{-2n-k}$. \square

The following definition is normally used in the context of termination proofs (see [25, page 253]): a quasi-ordering \succeq on $Ter(\Sigma)$ is called a *reduction quasi-ordering* if

- (i) \succ is strong normalising
- (ii) every substitution is both \succ -monotonic and \succeq -monotonic
- (iii) every function symbol F is \succeq -monotonic w.r.t. the product quasi ordering

Here, \succ is the relation $x \succ y \iff x \succeq y \wedge \neg(y \succeq x)$.

A reduction quasi-ordering \succeq is called *shrink-stable* w.r.t. term metric m if for every shrinking context $C[\]$ we have $t \succ u \Rightarrow C[t] \succ C[u]$.

Theorem 4 *Let (Σ, R) be a TRS, m a uniform term metric and \succeq a shrink-stable reduction quasi-ordering on $Ter(\Sigma)$.*

If R is a subrelation of \succ then $(Ter(\Sigma), d_m, \rightarrow_R)$ is Cauchy.

Proof Since m is uniform there is a shrinking umm h that is the supremum of all shrinking context functions. We show that any ω reduction sequence of \rightarrow is eventually within diameter $h^r(1)$, by induction for all r . The base $r = 0, h^0(1) = 1$ is trivial.

Consider any ω -reduction sequence f of \rightarrow_R . Clearly, it is also a reduction sequence for \succeq , because contexts and substitutions preserve this order. Because \succ is strongly normalising, there is a $k \geq 0$ such that $\forall n \geq k. \neg(f(n) \succ f(n+1))$. Moreover, for each n there are a context $D_n[\]_{p_n}$, a substitution θ_n and a rule $t_n \rightarrow u_n \in R$ such that $f(n) = D_n[\theta_n(t_n)]$ and $f(n+1) = D_n[\theta_n(u_n)]$. Because substitutions preserve the strict order and \succeq is shrink-preserved this means that for all $n \geq k$ context D_n must be non-shrinking.

A position $p \in Pos(f(k))$ is called *stable* if none of the positions $p_j, j \geq k$ is a proper prefix of p . In particular, $f(k) \stackrel{p}{\sim} f(k+x)$. It is maximally stable if it is stable and $p = p_i$ for some i . For any maximally stable p the function $g_p(i) = f(k+i)|_p$ defines an *omega*-reduction sequence (on the reflexive closure of \rightarrow_R). By induction hypothesis, it is eventually within diameter $h^r(1)$, say from k_p . Because $f(k)$ has only finitely many positions it also has only finitely many maximally stable positions q and beyond the maximum of all of their k_q all g_q will be within diameter $h^r(1)$. We can recover the distances within f from the sub-term projections: $d_m(f(k+n), f(k+n+1)) = \max_q (C_q(d_m(g_q(n), g_q(n+1))))$, where q ranges over the maximally stable positions in $f(k)$ and C_q is the context function of the context D_{k+x} for which $p_{k+x} = q$. Notice that for each n , the values $C_q(d_m(g_q(n), g_q(n+1)))$ are non-zero for at most one q , the one for which q is a prefix of p_{k+n} . Because each C_q is pointwise bounded by h it follows that $d_m(f(y), f(y+1)) \leq h(h^r(1)) = h^{r+1}(1)$ for all $y \geq k + \max_q(k_q)$. \square

The argument in the proof of theorem 4 is not fundamentally new (see proposition 5 in [11]) except that the presence of non-shrinking contexts under a term metric complicates matters slightly.

How does one find a quasi-reduction ordering that is shrink-stable? This is typically similar to the task of showing a TRS to be simply terminating, except that at several stages one can use the weak order \succeq where a termination proof would require the strict order \succ . This is best demonstrated at an example.

11 Application Example

The chapter on infinite rewriting in [21] motivates the subject with the following example, an iTRS modelling the sieve of Eratosthenes:

$$\begin{aligned}
\text{Filter}(\text{Cons}(x,y), \text{Zero}, m) &\rightarrow \text{Cons}(\text{Zero}, \text{Filter}(y, m, m)) \\
\text{Filter}(\text{Cons}(x,y), S(n), m) &\rightarrow \text{Cons}(x, \text{Filter}(y, n, m)) \\
\text{Sieve}(\text{Cons}(\text{Zero}, y)) &\rightarrow \text{Sieve}(y) \\
\text{Sieve}(\text{Cons}(S(n), y)) &\rightarrow \text{Cons}(S(n), \text{Sieve}(\text{Filter}(y, n, n))) \\
\text{Nats}(n) &\rightarrow \text{Cons}(n, \text{Nats}(S(n))) \\
\text{Primes} &\rightarrow \text{Sieve}(\text{Nats}(S(S(\text{Zero}))))
\end{aligned}$$

Given that this is such a fundamentally motivating example, one would expect that it is converging, w.r.t. to some term metric. However, it is not, at least not as an iTRS with the definition as in this paper.

Because of the penultimate rule, for it to be converging it is necessary that the component $\text{Cons}_{m,2}$ is shrinking, to allow “infinite lists”. This causes a problem with the third rule, because it lifts variable y out of such a shrinking context; for this to be a proper rule the function Sieve_m would have to be non-continuous at 0, which in turn would prevent us from applying Sieve to any (eventually) infinite lists, but the rewrite system does, with its last two rules.

The problem seems a technicality, caused by our condition (ii) for rewrite rules when the third rule of the system is considered. However, there is indeed a slight problem with this rule. If the metric allows arbitrary infinite lists then rule 3 would rewrite the term $\text{Sieve}(\text{Cons}(\text{Zero}, (\text{Cons}(\text{Zero}, \dots))))$ to itself. Although this does not contradict convergence (reflexive steps never do), it does contradict *strong convergence* [21], which requires that any reduction sequence moves redex positions arbitrarily deep into the terms. This redex would stay happily at position $\langle \rangle$, and the convergence is slightly accidental.

The iTRS is repairable though — it is generally possible to make rules comply with condition (ii) by padding them with “delay” functions. For term metric ∞ the modified rules look like this:

$$\begin{aligned}
\text{Filter}(\text{Cons}(x,y), \text{Zero}, m) &\rightarrow \text{Cons}(\text{Zero}, \text{Filter}(y, m, m)) \\
\text{Filter}(\text{Cons}(x,y), S(n), m) &\rightarrow \text{Cons}(D(x), \text{Filter}(y, n, m)) \\
\text{Sieve}(\text{Cons}(\text{Zero}, y)) &\rightarrow D(\text{Sieve}(y)) \\
\text{Sieve}(\text{Cons}(S(n), y)) &\rightarrow \text{Cons}(S(D(n)), \text{Sieve}(\text{Filter}(y, n, n))) \\
\text{Nats}(n) &\rightarrow \text{Cons}(n, \text{Nats}(S(n))) \\
\text{Primes} &\rightarrow \text{Sieve}(\text{Nats}(S(S(\text{Zero})))) \\
D(\text{Cons}(x,y)) &\rightarrow \text{Cons}(D(x), D(y)) \\
D(S(x)) &\rightarrow S(D(x)) \\
D(\text{Zero}) &\rightarrow \text{Zero}
\end{aligned}$$

Essentially, the function D is the identity function (on streams or numbers), but its appearance on right-hand sides pushes variables further down the term. If Sieve was now applied to an infinite stream of zeros the system would (strongly) converge to the term $D(D(\dots))$.

For term metric ∞ all rules are now “legal”: the lowest nesting depth of a variable on the right-hand side is never smaller than the corresponding value on the left-hand side. To prove that this iTRS S is converging, we can build the finite TRS $[S]$ (since ∞ is continuous) and show that $[S]$ is Cauchy (on finite terms). System $[S]$ contains all the rules of S plus the following:

$$\begin{aligned}
& Filter(x_1, x_2, x_3) \rightarrow Cons(\perp, \perp) \\
& Filter(Cons(x_1, y_1), Zero, m) \rightarrow Cons(Zero, Filter(\perp, \perp, \perp)) \\
& Filter(Cons(x_1, y_1), S(n_1), m) \rightarrow Cons(D(\perp), Filter(\perp, \perp, \perp)) \\
& Sieve(x_1) \rightarrow D(\perp) \\
& Sieve(Cons(x_1, y_1)) \rightarrow D(Sieve(\perp)) \\
& Sieve(x_1) \rightarrow Cons(\perp, \perp) \\
& Sieve(Cons(x_1, y_1)) \rightarrow Cons(S(\perp), Sieve(\perp)) \\
& Sieve(Cons(S(n_1), y)) \rightarrow Cons(S(D(\perp)), Sieve(Filter(\perp, \perp, \perp))) \\
& Nats(n_1) \rightarrow Cons(\perp, \perp) \\
& Nats(n) \rightarrow Cons(n, Nats(\perp)) \\
& Nats(n) \rightarrow Cons(n, Nats(S(\perp))) \\
& Primes \rightarrow Sieve(\perp) \\
& Primes \rightarrow Sieve(Nats(\perp)) \\
& Primes \rightarrow Sieve(Nats(S(\perp))) \\
& Primes \rightarrow Sieve(Nats(S(S(\perp)))) \\
& D(x_1) \rightarrow Cons(\perp, \perp) \\
& D(Cons(x_1, y_1)) \rightarrow Cons(D(\perp), D(\perp)) \\
& D(x_1) \rightarrow S(\perp) \\
& D(S(x_1)) \rightarrow S(D(\perp)) \\
& D(x_1) \rightarrow Zero
\end{aligned}$$

It is very easy to find a reduction ordering for this TRS (being shrink-stable comes for free under term metric ∞) to show that it is Cauchy: the order only compares the root symbols and ignores the rest of the terms. One can view this as an interpretation of terms in the ordinal 4:

$$\begin{array}{lll}
Primes_4 & = 3 & Nats_4(x) = 1 & S_4(x) = 0 \\
Sieve_4(x) & = 2 & D_4(x) = 1 & Zero_4 = 0 \\
Filter_4(x, y, z) & = 1 & Cons_4(x, y) = 0 & \perp_4 = 0
\end{array}$$

This interpretation interprets all left-hand sides as bigger numbers than their right-hand sides. This order is clearly well-founded and preserved by all non-shrinking contexts (as the trivial context is the only one), so $t \succeq u \iff [[t]]_4 \geq [[u]]_4$ is a reduction quasi-ordering that shows (using theorem 4) that $[S]$ is Cauchy and thus by proposition 17 that the iTRS S is converging.

12 Potential Variations

The definitions chosen here do not cover every possible variation one might want to throw at infinitary rewriting, but they go very far and this section discusses some alternatives.

This paper only looked at infinitary rewriting within ordinal ω : all infinite terms arise as metric completions of finite terms, all operations on these terms arise through metric completion of uniformly continuous functions operating on finite terms, and similarly the relations between infinite terms arise as completions of uniformly lsc relations between finite terms. Forcing everything to jump through this completion hoop means that nothing of interest would happen “beyond ω ”, i.e. such iTRSs are “ ω -closed”. This is not the only approach one can take, in fact much of the infinitary rewriting literature [21, 11] operates *directly* on infinitary terms, and extends its studies to relations that are not ω -closed. I would argue that forcing the study of infinitary rewriting to follow a completion process is important, as it is a protection against random concepts and random definitions; however, *metric completion* is not the only mechanism at our disposal, and other completion processes could give sensible notions of transfinite rewriting at larger ordinals.

In all cases, the rewrite relations on infinitary terms were defined to be reflexive. It is possible to deviate from that (and maintain that such relations are uniformly lower semi-continuous and pointwise closed), but not very far: it would suffice to require that such relations are merely reflexive in the neighbourhood of accumulation points, but anything weaker would be problematic.

All iTRSs were required to have only *finitely many* rules. The reason for this constraint is to ensure that the lifting of the rewrite relation from finite to infinitary terms is canonical and unique. This does not mean that an infinite set of rules never has such a canonical and unique lifting, but it would no longer suffice to look at individual rules to establish that.

Rewrite relations were required to be *pointwise compact*. For a uniformly lsc relation R to be liftable it would suffice to impose the weaker condition that it is pointwise closed. However, that stronger condition ensures that lifting is functorial w.r.t. relational composition: $R^\bullet ; S^\bullet = (R ; S)^\bullet$. Moreover, because of the constraint to finitely many rules that was imposed for other reasons anyway, pointwise compactness is guaranteed. Pointwise compactness is not the only invariant one can use to ensure that lifting distributes over relational composition: an alternative would be to require that the inverse relation R^{-1} is uniformly usc.

One condition for pairs of terms (t, u) to qualify as infinitary rules turned out to be very strong, condition (ii): $\llbracket t \rrbracket_{\llbracket m \rrbracket} \geq \llbracket u \rrbracket_{\llbracket m \rrbracket}$. It is possible to relax this requirement in various ways, e.g. an alternative condition would be:

$$\exists k \in \mathcal{N}. \forall f \in \text{Var} \rightarrow m. \quad k \cdot \llbracket t \rrbracket_{\llbracket m \rrbracket}(f) \geq \llbracket u \rrbracket_{\llbracket m \rrbracket}(f).$$

This would still imply that all variables on the right-hand side occur on the left, and for subadditive m the rewrite relation would still be uniformly lsc. In other words, this would give rise to a sensible notion of infinitary rewriting. Moreover, the application example from the previous section would (in its original version) now be legal under metric ∞ . However, the relaxation with a factor k badly affects

convergence proofs, in particular the proof of proposition 16 would irreparably fail.

In [21] (as well as earlier papers on which their chapter is based) the authors largely abandoned the semantic notion of convergence for a stronger variety, called *strong convergence*, because this shows better behaviour w.r.t. confluence problems. A reduction sequence is strongly convergent if redex positions (this is w.r.t. term metric ∞) move eventually arbitrarily deep. First, note that this is generalisable to other term metrics: the “depth” of a rewrite step $C[\theta(t)] \rightarrow C[\theta(u)]$ in term metric m can be seen as the value $C_m(1)$ — which is also the distance between the two terms *if* they have different root symbols. Thus an ω -reduction sequence is strongly convergent if its depths converge to 0. An entire iTRS could be regarded as strongly convergent if all its reduction sequences are strongly convergent.

This has still a very syntactic flavour, because the depths are associated with contexts, and ARSs have such numbers not occur in any other way. However, they might: instead of using ARSs with ordinary relations one could use fuzzy ARSs with *fuzzy relations* (see e.g. [12]) — in a fuzzy set/relation characteristic functions that are $\{0, 1\}$ -valued are replaced with ones that take values in $[0, 1]$. With this we can give a rewrite step its depth as its truth value. In this sense, a strongly convergent reduction sequence would in the limit have reduction steps with truth value 0, i.e. no reduction step at all, and this very much captures the idea of strongly convergent reductions.

Unfortunately, such an approach would mean to redo the entire section on continuous relations from scratch, defining concepts such as “lower semi-continuous fuzzy” relations, etc. It is certainly possible to generalise the Hausdorff metric from sets to fuzzy sets (provided the characteristic functions are continuous), but it opens up further choices: for example, an alternative metric on fuzzy sets is to compare the graphs of their characteristic functions as sets in the Hausdorff metric; in that metric, two fuzzy sets are exactly the same distance apart as their respective complements.

13 Conclusions

We have studied the meta-theory of infinitary rewriting by largely divorcing concrete rewriting from infinite terms and explaining such operations/relations instead through metric completion. Thus, not only infinite terms arise through metric completions, so do rewrite relations on infinite terms. This latter view is novel and required a thorough study of the lifting of relations from metric spaces to their completions. In essence: uniformly lower semi-continuous relations that are pointwise compact can be lifted. What is also novel is the view of regarding a term metric as a Σ -algebra with the carrier set $[0, 1]$.

The investigation has unveiled a variety of areas in rewriting for which the required uniformity is not always forthcoming:

- infinite set of rewrite rules
- non-left-linear rules (except for non-continuous or complete term metrics)
- infinitary right-hand sides cause problems with non-continuous or non-simple term metrics

- the completion w.r.t. a term metric that is not continuous (at 0) does not give a term algebra
- term metrics that are not subadditive may not give rise to uniformly lsc rewrite relations, and thus may have unliftable rewrite relations

Moreover, rewrite rules were restricted beyond non-left-linearity, and an important condition emerged that prevents a certain kind of non-convergent behaviour. The condition $\llbracket t \rrbracket_{\llbracket m \rrbracket} \geq \llbracket u \rrbracket_{\llbracket m \rrbracket}$ is a healthiness condition for infinitary rules; it implies that all variables of the right-hand side occur on the left, and, more importantly, that the same is true for all approximations of the rule.

For the original term metric ∞ this condition forbids (amongst other things) collapsing rules. It has been known (see figure 14 in [11]) that the presence of two different collapsing rules under this metric makes a system necessarily non-convergent on infinitary terms. Collapsing rules are not the sole culprits here, e.g. the rules $F(G(x)) \rightarrow G(x)$, $G(F(x)) \rightarrow F(x)$ would show a similar pattern of non-convergence for $F(G(F(G(\dots))))$. The condition (ii) prevents this particular form of non-convergence.

Moreover, a framework for convergence proofs has been set up that reduces convergence proofs of infinitary systems to Cauchy-ness proofs of certain finite term rewriting systems — provided the term metric is continuous. The Cauchy-ness proofs for finite systems require certain reduction quasi-orderings, which can be set up in similar ways as simplification orderings, although the exact details depend on the term metric involved. Particularly simple is the case of term metric ∞ for which this method was carried through on an example. The technique used is fundamentally the same as in [11], but the mentioned extra condition on rewrite rules ensures that the method is sound to show convergence for *all* reduction sequences, not just those that commence on finite terms.

Acknowledgements

I would like to thank several people who helped me with a number of questions I had on topological issues: an anonymous referee, John Derrick, Alex Simpson, Martin Escardó, and the “Topology Q+A board”.

References

1. Arnold, A., Nivat, M.: Metric interpretations of infinite trees and semantics of nondeterministic recursive programs. *Theoretical Computer Science* **11**(2), 181–205 (1980)
2. Baader, F., Nipkow, T.: *Term Rewriting and All That*. Cambridge University Press (1998)
3. Bird, R., de Moor, O.: *Algebra of Programming*. Prentice Hall (1997)
4. Brainerd, W.S.: Tree generating regular systems. *Information and Control* **14**, 217–231 (1969)
5. Copson, E.: *Metric Spaces*. Cambridge University Press (1968)
6. Corazza, P.: Introduction to metric-preserving functions. *American Mathematical Monthly* **106**(4), 309–323 (1999)
7. Davey, B.A., Priestley, H.A.: *Introduction to Lattices and Order*. Cambridge University Press (1990)
8. Day, J.M., Franklin, S.P.: Spaces of continuous relations. *Mathematische Annalen* **169**, 289–293 (1967)

-
9. Dershowitz, N., Kaplan, S.: Rewrite, Rewrite, Rewrite, Rewrite, Rewrite, ... In: Principles of Programming Languages, pp. 250–259. ACM (1989)
 10. Dershowitz, N., Kaplan, S., Plaisted, D.: Infinite normal forms. In: ICALP, pp. 249–262 (1989)
 11. Dershowitz, N., Kaplan, S., Plaisted, D.: Rewrite, Rewrite, Rewrite, ... Theoretical Computer Science **83**(1), 71–96 (1991)
 12. Goguen, J.: Fuzzy sets. In: E. Mamdani, B. Gaines (eds.) Fuzzy Reasoning and Its Applications, pp. 67–115. Academic Press (1981)
 13. Grimeisen, G.: Continuous relations. *Mathematische Zeitschrift* **127**, 35–44 (1972)
 14. Hahn, H.: Reelle Funktionen. Chelsea Publishing Company, New York (1948)
 15. Jänich, K.: Topology. Springer (1984)
 16. Jones, S.P. (ed.): Haskell 98, Languages and Libraries, The Revised Report. Cambridge University Press (2003)
 17. Kahrs, S.: The variable containment problem. In: Higher-Order Algebra, Logic and Term Rewriting, Second International Workshop, Selected Papers, pp. 109–123 (1995). LNCS 1074
 18. Kampermann, J.: Compilation of term rewriting systems. Ph.D. thesis, University of Amsterdam (1996)
 19. Kampermann, J., Walters, H.: Lazy rewriting and eager machinery. In: Rewriting Techniques and Applications, LNCS, vol. 914, pp. 147–162 (1995)
 20. Kelley, J.: Convergence in topology. *Duke Mathematical Journal* **17**(3), 277–283 (1950)
 21. Kennaway, R., de Vries, F.J.: Term Rewriting Systems, chap. Infinitary Rewriting, pp. 668–711. Cambridge University Press (2003)
 22. Kuratowski, K.: Topology. Academic Press (1966)
 23. Smyth, M.: Topology. In: S. Abramsky, D. Gabbay, T. Maibaum (eds.) Handbook of Logic in Computer Science, Volume 1, pp. 641–762. Oxford University Press (1992)
 24. Stoltenberg-Hansen, V., Tucker, J.: Effective algebras. In: S. Abramsky, D.M. Gabbay, T. Maibaum (eds.) Handbook of Logic in Computer Science, Volume 4, pp. 357–526. Oxford University Press (1995)
 25. Terese (ed.): Term Rewriting Systems. Cambridge University Press (2003)