Submitted to the International Conference on Rewriting Techniques and Applications
http://rewriting.loria.fr/rta/

INFINITARY REWRITING: FOUNDATIONS REVISITED

STEFAN KAHRS

University of Kent at Canterbury, School of Computing
E-mail address: S.M.Kahrs@kent.ac.uk

ABSTRACT. Infinitary Term Rewriting allows to express infinitary terms and infinitary reductions that converge to them. As their notion of transfinite reduction in general, and as binary relations in particular two concepts have been studied in the past: strongly and weakly convergent reductions, and in the last decade research has mostly focused around the former.

Finitary rewriting has a strong connection to the equational theory of its rule set: if the rewrite system is confluent this (implies consistency of the theory and) gives rise to a semi-decision procedure for the theory, and if the rewrite system is in addition terminating this becomes a decision procedure. This connection is the original reason for the study of these properties in rewriting.

For infinitary rewriting there is barely an established notion of an equational theory. The reason this issue is not trivial is that such a theory would need to include some form of "getting to limits", and there are different options one can pursue. These options are being looked at here, as well as several alternatives for the notion of reduction relation and their relationships to these equational theories.

1. Introduction

Infinitary rewriting deals with infinite terms, which are defined through the metric completion of finite terms through some metric. In the simplest case (metric d_∞) this is equivalent to a co-inductive definition of terms, i.e. the set of *infinitary terms* $Ter^\infty(\Sigma)$ is the *largest set* such that every t in this set has some root symbol F taken from the signature Σ and n direct subterms t_i ($1 \leq i \leq n$) that are all in $Ter^\infty(\Sigma)$, where n is the arity of F as defined by the signature. In other words, infinitary terms are defined co-inductively through the way they unfold, without a guarantee that this unfolding ever comes to an end. Infinite terms are indeed those where it does not.

Metric completion is a general-purpose semantic construction on metric spaces which "adds" to a metric space limits to all its Cauchy-sequences, in the sense that there is a dense isometric embedding of the orginal space into a complete metric space. Using metric completion with other metrics on terms than d_∞ can restrict the infinite terms under consideration (but not add others) [6], as $Ter^\infty(\Sigma)$ can also be seen as a final co-algebra. For the purposes of this paper the choice of metric is (largely) immaterial, i.e. as long it allows for infinite terms at all.

1998 ACM Subject Classification: F.4.2.
Key words and phrases: Infinitary Rewriting, Equivalence Relation, Model.

Reductions that only involve finite terms can still approximate infinite terms "in the limit" by moving arbitrarily close (in terms of the distance function) to some such infinite term.

That is the idea behind infinitary rewriting. The problem is then to decide *how* exactly the limit is reached. For transfinite reductions we may expect something similar to metric completion, some kind of closure operation, allowing us to reach limit terms in the limit of... *what*? Here we encounter our first problem: what is the structure to which we need to add limits and what kind of limits?

Traditionally [3], the answer to this is: (i) the structure to which we add limits are (infinite) reduction sequences. As these are functions with ordinals as their domain, we can ensure that the limits are connected to "what came before" by demanding that these functions are (ii) *continuous*. The result of these choices is the notion of (weakly) convergent reduction $\twoheadrightarrow_w$. There are alternatives to continuity; we can strengthen the condition further, or weaken it by aiming for something like *adherence* rather than convergence. We are also not forced to stick to reduction sequences as our starting point: in particular, we may just as well use the reduction relation as a whole.

The initial interest in the subject of (finitary) term rewriting was triggered by a number of observations that led to the Knuth-Bendix completion procedure [9], deriving (if and when it succeeds) a decision procedure for a given equational theory. These observations establish links between the many-step rewrite relation and the equational theory. What happens to those links in infinitary rewriting?

Before we can even pose this question meaningfully we need a notion of an equational theory for infinitary rewriting. There is none in the literature, although there are ideas approximating it. One is the specific equivalence relation $\sim_{hc}$ [8] (the so-called equivalence "modulo hyper-collapsing terms"). Its construction is not easy to generalise as it relies upon being used in connection with orthogonal iTRS: in the presence of overlapping rules the top of a hyper-collapsing term could form a redex with the context, undermining the idea that such terms are unsolvable in the sense of the λ-calculus [1]. Another idea comes from the concept of equational model [2, 3, 10], since a notion of syntactic equivalence can be derived from that: $t =_R u$ iff the equation $t = u$ holds in all equational models of R. As we shall see, the notion of equational theory derived from such classes of models is quite strong.

The reason the definition of the equational theory is an issue at all is this: on finite terms we can define the equivalence relation $=_R$ as the *smallest equivalence relation* that includes all the rules, is substitutive, and for which $F(t_1, \ldots, t_n) =_R F(u_1, \ldots, u_n)$ holds whenever $\forall i.\, t_i =_R u_i$ and F is an n-ary function symbol from the signature. But for infinitary terms this definition is not suitable, because it does not allow for any form of limit-taking: from $A =_R B$ we could not deduce that the infinite term $t = C(A, t)$ is $=_R$-equivalent to the infinite term $u = C(B, u)$, since the equivalence closure only permits a *finite number* of equation applications. For the same reason, the transfinite reduction relation $\twoheadrightarrow$ (in any of the variations we consider) would not be a subrelation of $=_R$.

Thus some form of limit-taking needs to be incorporated into a suitable notion of equational theory. There are several ways in which one can allow for limit-taking that lead to different notions of equivalence:

- the conservative approach: form the equivalence closure of $\twoheadrightarrow$, for any given notion of transfinite reduction relation $\twoheadrightarrow$;

- defining an equivalence relation: take the *blue* and equivalence closure of the single-step rewrite relation $\rightarrow$, where "*blue*" is a placeholder for properties that ensure that certain limits are included in the relation.
- inducing an equivalence from a class of models

The conservative approach is pragmatic and makes the connection to the reduction relation straightforward, but it denies rules true equational status. For example, it does not permit us to form an infinite sequence by applying rules forwards as well as backwards and conclude that the limit of the sequence is equivalent to its beginning.

2. Preliminaries

We need several notions from Topology and Infinitary Rewriting. This section contains those definitions, to make the paper self-contained.

A property P on subsets of a set A is said to be *closable* if the intersection $\bigcap_{i\in I} A_i$ of any family of subsets $A_i \subseteq A, i \in I$ that each satisfy P itself satisfies the property. As I may be empty, A has to satisfy P too. If P is closable we can form the P-closure of any subset $K \subset A$, the smallest subset of A that contains K and satisfies P. In particular, this concept will be used on relations, viewing them as sets of pairs. Clearly, the conjunction of closable properties is closable.

2.1. Topology

A *topological space* is a pair $(S, \mathcal{O})$ where S is a set and $\mathcal{O}$ a subset of $\wp(S)$ such that it is closed under finite intersections and arbitrary unions, and $\emptyset, S \in \mathcal{O}$. The elements of $\mathcal{O}$ are called *open* sets, their complements w.r.t. S are called *closed* sets. The *closure* of a subset $A \subset S$ is the intersection of all closed sets that contain A, and is written $\mathrm{Cl}(A)$. A *neighbourhood* of $x \in S$ is a set $N \subseteq S$ for which there is an $A \in \mathcal{O}$ such that $x \in A \subseteq N$. A point $z \in S$ is called *discrete* iff $\{z\}$ is open.

A topological space S is called T_0 iff $\forall x, y \in S.\, x \in \mathrm{Cl}(\{y\}) \wedge y \in \mathrm{Cl}(\{x\}) \Rightarrow x = y$. It is called T_1 iff all singleton sets are closed. It is T_2 (or Hausdorff) iff any two distinct points in S have disjoint neighbourhoods.

A function $f : A \rightarrow B$ between topological spaces is called continuous iff $f^{-1}(X)$ is open whenever X is open. A function $f : A \rightarrow B$ between topological spaces is an *open map* iff it preserves open sets; it is a *closed map* iff it preserves closed sets. A relation R between topological spaces A and B is called *lower semi-continuous* (or lsc) if R^{-1} preserves open sets, and *upper semi-continuous* (or usc) if R^{-1} preserves closed sets.

A subset F of a topological space is called *compact* iff whenever $\bigcup_{i\in I} A_i \supseteq F$, where each A_i is open, then there is a finite subset $J \subseteq I$ such that $\bigcup_{i\in J} A_i \supseteq F$.

A metric space is a set (M, d) where M is a set and $d : M \times M \rightarrow \mathcal{R}$ a distance function such that for all $x, y, z \in M$: (i) $d(x, y) = 0 \iff x = y$ and (ii) $d(x, z) \leq d(x, y) + d(z, y)$. For an ultra-metric space (ii) is replaced by the stronger condition $d(x, z) \leq \max(d(x, y), d(z, y))$. The topology of a metric space is defined as follows: $A \subseteq M$ is open iff $\forall x \in A.\, \exists \epsilon > 0. \forall y \in M.\, d(x, y) < \epsilon \Rightarrow y \in A$.

The metric completion of (M, d) is the unique (up to isomorphism) metric space $(M^\bullet, d^\bullet)$, with a function $e : M \rightarrow M^\bullet$, such that e preserves distances and $\mathrm{Cl}(e(M)) = M^\bullet$.

A function between metric spaces A and B is *uniformly continuous* iff $\forall.\, \epsilon > 0.\, \exists \delta > 0.\, \forall x, x' \in A.\, d_A(x, x') < \delta \Rightarrow d_B(f(x), f(x')) < \epsilon$. Uniformly continuous functions have unique continuous extensions to the respective metric completions. A special case are *non-expansive* functions where $\delta = \epsilon$.

2.2. Infinitary Term Rewriting

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

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

$$\begin{array}{rcll} [\![x]\!]^{\rho}_A & = & \rho(x), & \text{if } x \in \mathit{Var} \\ [\![F(t_1, \ldots, t_n)]\!]^{\rho}_A & = & F_A([\![t_1]\!]^{\rho}_A, \ldots, [\![t_n]\!]^{\rho}_A) & \end{array}$$

Infinitary terms are defined through a metric completion process. For this paper we focus on the metric d_∞ which is defined as follows inductively on finite terms: $d_\infty(t, u) = 1$ iff t and u have different roots; $d_\infty(t, t) = 0$; otherwise, $d_\infty(F(t_1, \ldots, t_n), F(u_1, \ldots, u_n)) = 1/2 {*} \max_{1 \le i \le n} d_\infty(t_i, u_i)$. $(\mathit{Ter}(\Sigma), d_\infty)$ is an ultra-metric space and we write $(\mathit{Ter}^\infty(\Sigma), d_\infty)$ for its metric completion, which is also a Σ-algebra [6]. This can be generalised for other metrics d_m with their completions $\mathit{Ter}^m(\Sigma)$ [6].

A rewrite rule is pair (l, r), usually written $l \to r$, such that $l \in \mathit{Ter}(\Sigma) \setminus \mathit{Var}$, $r \in \mathit{Ter}^\infty(\Sigma)$ and all variables occurring in r also occur in l.

An iTRS is a pair (Σ, R) where Σ is a signature and R a set of rewrite rules for that signature. The rewrite step relation $\to_R$ relates $C[\sigma(l)] \to_R C[\sigma(r)]$, where $l \to r$ is a rewrite rule in R, σ a substitution, and $C[\,]$ a context. Substitution application and context application are uniquely derived from their respective concepts on finite terms [6], and can also be defined in terms of the $[\![\_]\!]_{\mathit{Ter}^\infty(\Sigma)}$ interpretation.

3. Transfinite Sequences

In the following we are looking at several notions of transfinite reduction relations $X(\to_R)$. These are all functions of the single step reduction relation $\to_R$. We call $X(\to_R)$ *infinitarily transitive* if $X(\to_R) = X(X(\to_R))$, i.e. if it is a fixpoint of X. Usually, we take $\to_R$ to be clear from the context and write $\twoheadrightarrow_x$ for $X(\to_R)$ and $\twoheadrightarrow_{xx}$ for $X(X(\to_R))$. We also write $\twoheadleftarrow_x$ for $\twoheadrightarrow_x^{-1}$, $\downarrow_x$ for $\twoheadrightarrow_x ; \twoheadleftarrow_x$, and $\uparrow_x$ for $\twoheadleftarrow_x ; \twoheadrightarrow_x$.

The reason the property "infinitarily transitive" is desirable is similar to wanting that $\to_R^*$ is the same as $(\to_R^*)^*$; the property also features in the proofs of [7].

Transfinite sequences of terms can be defined as functions from an ordinal (the index domain) to the set of (infinitary) terms, viewing ordinals as von Neumann ordinals, i.e. the ordinal α is the set of all ordinals strictly smaller than α. Reduction sequences of an iTRS are those where neighbouring elements are within the single-step reduction relation, i.e. if $f : \alpha \to \mathit{Ter}^\infty(\Sigma)$ is our reduction sequence then $f(n) \to_R f(n+1)$, provided $n + 1 < \alpha$.

This works fine for *finite* sequences. For infinite sequences this definition fails to put any constraints whatsoever what happens at $f(\lambda)$, for limit ordinals λ.

3.1. Standard solution: weak convergence

The traditional choice to fix this is to demand that the function f is continuous w.r.t. the usual order topology. Effectively, this means that $f(\lambda)$ must be the (unique) limit of $f(\gamma)$, as γ approaches λ from below. To express it in terms of distances:

$$\forall\epsilon > 0.\, \exists\gamma.\, \forall\gamma'.\, \gamma \leq \gamma' < \lambda \Rightarrow d_\infty(f(\gamma'), f(\lambda)) < \epsilon$$

If the indexing set of a transfinite reduction sequence is a successor ordinal then we have a *closed sequence*, because it is a sequence with a last element: if $\alpha + 1$ is the domain of f then the last element of the sequence is $f(\alpha)$.

This also gives us a way of defining the relation $\twoheadrightarrow_w$: $t \twoheadrightarrow_w u$ iff there is some ordinal α and some closed reduction sequence $f : \alpha+1 \to \mathit{Ter}^\infty(\Sigma)$ such that $f(0) = t$ and $f(\alpha) = u$.

In the infinitary rewriting literature this is often called "weak convergence" [8], where "strong convergence" requires that the sequences converges solely due to the positioning of redexes, i.e. if $f(\gamma) = C_\gamma[\sigma_\gamma(l_\gamma)] \to_R f(\gamma+1) = C_\gamma[\sigma_\gamma(r_\gamma)]$ we can build a new sequence $g(\gamma) = C_\gamma[x]$ (i.e. replacing all contracted redexes with the variable x); f is then strongly converging if g converges too, and to the same limit as f. Expressing strong convergence as a function of $\to_R$ (rather than R) would require to recover a minimal rule set from the relation $\to_R$ — a slightly delicate issue that goes beyond the scope of this paper.

Example 3.1. Consider the iTRS with rules

$$\begin{array}{rcl} F(A,x) & \to & F(B,D(x)) \\ B & \to & A \end{array}$$

We have $F(A,x) \to_R^2 F(A,D(x))$, but we do not have $F(A,x) \twoheadrightarrow_w F(A,D^\infty)$, because A would change to B at every other step. If we add the rule $F(A,x) \to F(A,D(x))$ to the system then $\to_R^*$ does not change, but $\twoheadrightarrow_w$ would change and now include $F(A,x) \twoheadrightarrow_w F(A,D^\infty)$.

Proposition 3.2. *$\twoheadrightarrow_w$ is in general not infinitarily transitive (though it is on converging iTRSs* [7]*).*

Proof. See example 3.1. Since $F(A,x) \to_R^2 F(A,D(x))$ we also have $F(A,x) \twoheadrightarrow_w F(A,D(x))$ and consequently $F(A,x) \twoheadrightarrow_{ww} F(A,D^\infty)$. ∎

Incidentally, this contradicts theorem 1(c) in [2].

3.2. Adherence

Instead of asking for convergence we can ask for adherence: $t \twoheadrightarrow_a u$ is defined like $t \twoheadrightarrow_w u$, except for one thing: instead of requiring that the witnessing indexing function f is continuous at limit ordinals λ we require that it is "adherent": This is in a certain sense a concept dual to convergence, because instead of demanding that a sequence is eventually always within a neighbourhood, the definition asks instead that it always eventually goes there. Formally:

$$\forall\epsilon > 0.\, \forall\gamma < \lambda.\, \exists\gamma'.\, \gamma \leq \gamma' < \lambda \wedge d_\infty(f(\gamma'), f(\lambda)) < \epsilon$$

The difference is that adherence merely requires that (a neighbourhood of) an accumulation point is visited by the sequence for index positions arbitrarily close to λ, without demanding that the sequence stays there. Intuitively, adherence requires that a cofinal subsequence of a reduction sequence converges to the limit, allowing for other terms in the sequence as computational noise.

The result of this is that a sequence can adhere to more than one limit. Certainly, any sequence converging to a limit adheres to it and therefore:

Proposition 3.3. $\twoheadrightarrow_w \subseteq \twoheadrightarrow_a$

Clearly, $\twoheadrightarrow_a$ is (finitary) transitive, $\twoheadrightarrow_a;\twoheadrightarrow_a \subseteq \twoheadrightarrow_a$, because the adherence condition never stops adherent sequences from being concatenated. One peculiarity of adherence over convergence is that the notion is less sensitive to the notion of the single-step relation, in the sense that if $\rightarrow_R^* = \rightarrow_S^*$ then the adherence relations of $\rightarrow_R$ and $\rightarrow_S$ are identical too. Also:

Proposition 3.4. $\twoheadrightarrow_a$ *is infinitarily transitive.*

Proof. Let $f : \alpha + 1 \rightarrow Ter^\infty(\Sigma)$ be the sequence witnessing $t \twoheadrightarrow_{aa} u$. We need to show $t \twoheadrightarrow_a u$. This can be proved by induction on α.

If $\alpha = 0$ then $t = u$ and the result follows by reflexivity.

If $\alpha = \beta + 1$ then there is a t' such that $t \twoheadrightarrow_{aa} t' \twoheadrightarrow_a u$, where the sequence witnessing $t \twoheadrightarrow_{aa} t'$ has length β. By induction hypothesis $t \twoheadrightarrow_a t'$, and the result follows by (finitary) transitivity of $\twoheadrightarrow_a$.

If α is a limit ordinal then the restriction of f to domain α has a cofinal subsequence [7] $g : \beta \rightarrow \alpha$ such that $f \circ g$ converges to $f(\alpha)$. Once we expand every $\twoheadrightarrow_a$ step in f to a new sequence h we have that $f = h \circ g'$ for some cofinal subsequence g' of h and thus $f \circ g = (h \circ g') \circ g = h \circ (g' \circ g)$ where $g' \circ g$ is a cofinal subsequence of h. ∎

By implication proposition 3.4 also shows that the inclusion $\twoheadrightarrow_w \subset \twoheadrightarrow_a$ is (in general) proper, as witnessed by example 3.1.

Adherence can also be characterised as follows: let W be the function mapping a relation $\rightarrow_R$ to its weak convergence relation $\twoheadrightarrow_w$. Then $\twoheadrightarrow_a$ is the least fixpoint of W that contains the single-step relation.

4. Relations

Instead of completing sequences by adding limits or accumulation points, we can define $\twoheadrightarrow$ more directly through closable properties of relations. There are the following notions of interest:

4.1. Pointwise Closure

We can view relations as set-valued functions, and add limits to their range. This leads to the following concept:

Definition 4.1. A relation R between topological spaces is called *pointwise closed* iff the sets $R^x = \{y \mid x\, R\, y\}$ are all closed.

Proposition 4.2. *Being pointwise closed is a closable property of relations.*

Proof. Let $A = \bigcap_i R_i$. Then $A^x = \{y \mid x\, A\, y\} = \{y \mid \forall i.\, x\, R_i\, y\} = \bigcap_i R_i^x$. Hence A^x is an intersection of closed sets and therefore closed. ∎

This allows us to use pointwise closure as a relation-constructing property.

Definition 4.3. The relation $P(\to_R) = \twoheadrightarrow_p$ is defined as the smallest reflexive, transitive and pointwise closed relation containing $\to_R$.

That $\twoheadrightarrow_p$ is infinitarily transitive is trivial by construction. We can explain $t \twoheadrightarrow_p u$ as "t can rewrite to something arbitrarily close to u", but if we want to get any closer we may have to start all over again from t.

Proposition 4.4. $\twoheadrightarrow_a \subseteq \twoheadrightarrow_p$.

Proof. By induction on the indexing ordinals for the sequences witnessing $t \twoheadrightarrow_a u$. The interesting case for $f : \alpha + 1 \to \mathit{Ter}^\infty(\Sigma)$ is when α is a limit ordinal. By induction hypothesis, $f(0) \twoheadrightarrow_p f(\gamma)$, for all $\gamma < \alpha$. The restriction of f to α has to contain a subsequence that converges to $f(\alpha)$. But then $f(\alpha)$ has to be in the closure of the $f(\gamma)$ and as $\twoheadrightarrow_p$ is pointwise closed the result follows. ∎

However, the relations $\twoheadrightarrow_a$ and $\twoheadrightarrow_p$ are not always the same:

Example 4.5.

$$\begin{array}{rcl} A & \to & B(A) \\ A & \to & C \\ B(C) & \to & D(C) \\ B(D(x)) & \to & D(D(x)) \end{array}$$

In this system we have $A \to_R^* D^n(C)$ for any finite n and therefore $A \twoheadrightarrow_p D^\infty$. But there is no single adherent (or convergent) sequence that can build up to that limit; as soon as a D appears in a reduct of A the reduction sequence is guaranteed to terminate.

Usually we can construct $\twoheadrightarrow_p$ more directly, as the pointwise closure of $\to_R^*$ (call it $\twoheadrightarrow_{p0}$).

Theorem 4.6. *If $\to_R$ is lsc then $\twoheadrightarrow_p = \twoheadrightarrow_{p0}$.*

Proof. It suffices to show that $\twoheadrightarrow_{p0}$ is transitive. Since $\to_R$ is lsc so is $\to_R^*$ [6]. Suppose $A \twoheadrightarrow_{p0} B \twoheadrightarrow_{p0} C$.

We need to show $A \twoheadrightarrow_{p0} C$, which means that for any $\epsilon > 0$ there is a C_ϵ such that $A \to_R^* C_\epsilon$ and $d_\infty(C_\epsilon, C) < \epsilon$. Because $\to_R^*$ is lsc and $B \to_R^* C$ for every $\epsilon > 0$ there is a $\delta > 0$ such that for all B' with $d_\infty(B', B) < \delta$ there is a $C_\epsilon(B')$ with $B' \to_R^* C_\epsilon(B')$ and $d_\infty(C_\epsilon(B'), C) < \epsilon$. Since $A \twoheadrightarrow_{p0} B$ we can find B_δ with $d_\infty(B_\delta, B) < \delta$ and $A \to_R^* B_\delta$. Hence $A \to_R^* C_\epsilon(B_\delta)$. ∎

In [6] a number of conditions are given under which the relation $\to_R$ is *uniformly lsc*, for a variety of term metrics. For the much weaker condition that $\to_R$ is lsc it suffices to require that the rules are left-linear, and in that case this is even independent of the term metric.

4.2. Topological Closure

The pointwise closure add limits to a relation at the "result side", and stays in this respect still very much within the intuition behind infinitary rewriting. Going beyond that and allowing the input side to change as well leads to fairly unintuitive relations.

For example, another closable property on relations between topological spaces A and B is that their set of pairs (their graph) is closed in the product space $A \times B$.

Definition 4.7. The relation $C(\rightarrow_R) = \twoheadrightarrow_t$ is the smallest reflexive and transitive relation containing $\rightarrow_R$ such that its graph is closed.

This means: if t_n and u_n are sequences converging to t and u, respectively, and if for all i: $t_i \twoheadrightarrow_t u_i$, then $t \twoheadrightarrow_t u$.

Clearly, closed relations are also pointwise closed and therefore $\twoheadrightarrow_p \subseteq \twoheadrightarrow_t$. Again, the inclusion is proper:

Example 4.8.

$$\begin{array}{rcl} \mathrm{LEQ}(0, x) & \rightarrow & T \\ \mathrm{LEQ}(S(x), 0) & \rightarrow & F \\ \mathrm{LEQ}(S(x), S(y)) & \rightarrow & \mathrm{LEQ}(x, y) \end{array}$$

The infinite term $t = \mathrm{LEQ}(S^\infty, S^\infty)$ only reduces to itself, in a single step, and thus also $\forall u.\ t \twoheadrightarrow_p u \Rightarrow t = u$. But we also have $t \twoheadrightarrow_t T$ and $t \twoheadrightarrow_t F$, because the sequences $a_n = LEQ(S^n(0), S^n(0))$ and $b_n = LEQ(S^{n+1}(0), S^n(0))$ both converge to t, but $a_n \twoheadrightarrow_t T$ and $b_n \twoheadrightarrow_t F$.

In contrast to $\twoheadrightarrow_p$ we usually cannot construct $\twoheadrightarrow_t$ as the topological closure of $\rightarrow_R^*$ (call it $\twoheadrightarrow_{t0}$), because that relation is often not transitive:

Example 4.9. Add the rule $\mathrm{INF} \rightarrow S(\mathrm{INF})$ to example 4.8. Then $\mathrm{LEQ}(\mathrm{INF}, \mathrm{INF}) \twoheadrightarrow_{t0} \mathrm{LEQ}(S^\infty, S^\infty) \twoheadrightarrow_{t0} T$, but we do not have $\mathrm{LEQ}(\mathrm{INF}, \mathrm{INF}) \twoheadrightarrow_{t0} T$.

5. Notions of Equivalence

When using rewrite relations to (semi-)decide an equivalence we want that equivalent terms have common reducts. Hence:

Definition 5.1. A pre-order $\twoheadrightarrow_x$ is called a *semi-decider* for an equivalence $=_E$ iff (i) $\twoheadrightarrow_x \subseteq =_E$ and (ii) $=_E \subseteq \downarrow_x$.

The first condition gives us soundness (if terms have common reducts they are equivalent), the second completeness. If $=_E$ is the equivalence closure of $\twoheadrightarrow_x$ then (i) is trivial and (ii) is equivalent to infinitary confluence, $\uparrow_x \subseteq \downarrow_x$. However, for infinitary rewriting this is a big "if".

In the presence of infinite terms, ordinary congruence relations fail to capture what is needed for equational reasoning in infinitary rewriting as equivalence closure is an inductive concept, not a coinductive one. This problem shows up in two separate ways: (i) for including transfinite reductions in the equivalence, and (ii) for allowing infinitely many subterm changes in a term of infinite size.

However, any equivalence relation $\sim$ on a topological space A induces a canonical topology on the quotient $A/\!\sim$: a set of equivalence classes is open iff their union is open in the topology of A. This condition is the finest topology that makes the projection map $[\_]_\sim : A \rightarrow A/\!\sim$ continuous. This also means that if we have *any* converging sequence $f(n)$ in A then $[f(n)]_\sim$ is converging in $A/\!\sim$.

Example 5.2. Consider the iTRS with the single rule $C \to S(C)$. Take as equivalence $\approx$ the congruence closure of the equation $C = S(C)$. The reduction sequence $C \to S(C) \to S(S(C)) \to ...$ converges to S^∞. By continuity, the sequence $[C]_\approx, [S(C)]_\approx, [S(S(C))]_\approx, ...$ converges to $[S^\infty]_\approx$. However, $[C]_\approx = [S(C)]_\sim = [S(S(C))]_\approx, ...$ and $[C]_\approx \neq [S^\infty]_\approx$.

Example 5.2 shows that quotient spaces can have very poor separation properties, e.g. in the example $Ter^\infty(\Sigma)/\approx$ is not T_1. These separation properties closely correspond to properties of equivalence relations, in the sense that they indirectly provide recipes for adding limits to an equivalence. This leads to the following concepts:

Definition 5.3. An equivalence relation $\sim$ on a topological space A is called *weakly separating*, iff:

$$\forall x, y \in A.\, x \in \mathrm{Cl}([y]_\sim) \wedge y \in \mathrm{Cl}([x]_\sim) \Rightarrow x \sim y$$

Proposition 5.4. *$A/\sim$ is a T_0 space iff $\sim$ is weakly separating.*

Proof. Let $\sim$ be weakly separating and $[x]_\sim$ and $[y]_\sim$ be accumulation points of each other. Then $[x]_\sim \in \mathrm{Cl}(\{[y]_\sim\})$ which is equivalent to $[x]_\sim \subseteq \mathrm{Cl}([y]_\sim)$, hence $x \in \mathrm{Cl}([y]_\sim)$; the same argument gives $y \in \mathrm{Cl}([x]_\sim)$. As $\sim$ is weakly separating $x \sim y$ and so $[x]_\sim = [y]_\sim$.

If $\sim$ is not weakly separating then any witnessing counterexample is also a counterexample against $A/\sim$ being T_0. ∎

Given any relation R, we can form the "weakly separating equivalence closure" due to the following property:

Proposition 5.5. *Being weakly separating is a closable property.*

Proof. Clearly, the intersection $\sim = \bigcap_i \sim_i$ gives another equivalence where each equivalence class $[a]_\sim$ is the intersection of the equivalence classes $[a]_{\sim_i}$. Now assume $x \in \mathrm{Cl}([y]_\sim)$ and $y \in \mathrm{Cl}([x]_\sim)$. $x \in \mathrm{Cl}([y]_\sim) = \mathrm{Cl}(\bigcap_i [y]_{\sim_i}) \subseteq \bigcap_i \mathrm{Cl}([y]_{\sim_i})$. Hence, for all i, $x \in \mathrm{Cl}([y]_{\sim_i})$, and by the dual argument $y \in \mathrm{Cl}([x]_{\sim_i})$. Since each $\sim_i$ is weakly separating this implies $x \sim_i y$, for all i, and so $x \sim y$. ∎

T_0 is a very weak form of separation and we do not have that $\twoheadrightarrow_w$ would be included in the weakly separating equivalence closure of its rules. Often, the weakly separating equivalence closure makes no difference, but there are cases where it does:

Example 5.6. Consider the following specification of equality and logical negation:

$$\begin{array}{rcl}
E(x,x) & = & T \\
E(0,S(x)) & = & F \\
E(S(x),0) & = & F \\
E(S(x),S(y)) & = & E(x,y) \\
N(T) & = & F \\
N(F) & = & T
\end{array}$$

Instantiating the first equation we have $E(S^\infty, S^\infty) = T$. We can also derive $E(S^n(0), S^\infty) = F$, for any finite n. Thus $[T]_\sim$ is in the closure of $[F]_\sim$, where $\sim$ is the equivalence closure of $\to_R$. Moreover, for any finite n, $N(E(S^n(0), S^\infty)) = N(F) = T$; hence the closure of $[T]_\sim$ will also contain $N(E(S^\infty, S^\infty)) = N(T) = F$. Therefore, the weakly separating closure of $\sim$ will relate T and F.

Without the first equation the weakly separating closure would not change the relation: $E(S^\infty, S^\infty)$ would be in an equivalence class of its own; both the closures of $[T]_\sim$ and $[F]_\sim$ would contain that class, but not vice versa.

Definition 5.7. An equivalence relation $\sim$ on a topological space A is called *separating* iff all its equivalence classes are closed.

Again, we can use this concept as a closure principle:

Proposition 5.8. *Being separating is a closable property.*

Proof. The equivalence classes of $\sim$ are closed iff $\sim$ is pointwise closed, hence this follows from proposition 4.2. ∎

Example 5.9. If we remove the first equation $E(x, x) = T$ from example 5.6 then T and F would not be related by the weakly separationg closure of $\sim$, but they would be by the separating closure: since $E(S^\infty, S^\infty)$ is in the closure of both $[T]_\sim$ and $[F]_\sim$, T, F and $E(S^\infty, S^\infty)$ would all be equivalent under the separating closure of $\sim$.

Thus "separating" is a much stronger property than "weakly separating". If an iTRS (over metric d_∞) contains a collapsing rule $C[x, \ldots, x] \to x$ then the sequence $x_0 = x$, $x_{n+1} = C[x_n, \ldots, x_n]$ converges to a limit C^∞. The same is true if we start the sequence with $x_0 = y$ instead. Hence, $x \sim y$ if $\sim$ is the separating closure of $\to_R$. For other metrics d_m collapsing rules may not cause that problem, as the sequence x_n could be diverging under d_m.

In finitary term rewriting, the derivability of $x =_R y$ is used as the standard criterion for inconsistency. For infinitary rewriting (over metric d_∞) this becomes trivial for separating equivalences: the separating closure of $\to_R$ contains the pair (x, y) iff R contains a collapsing rule. The reason: if it does contain a collapsing rule then the previous argument applies, if it does not then each set $\{x\}$ remains an equivalence class of its own, since x is a discrete point in $\mathit{Ter}^\infty(\Sigma)$: thus $\{x\}$ is closed and no set not containing x has it in its closure. In particular, even example 5.6 is consistent despite $T \sim F$.

Separating equivalences characterise T_1 spaces:

Proposition 5.10. *$A/\!\sim$ is T_1 iff $\sim$ is a separating equivalence.*

Proof. Folklore[5, p. 207]. ∎

In a T_1 space, a sequence that is eventually constant can only converge to that constant, because all singleton sets in a T_1 space are closed. That also means that if all elements of a converging sequence are equivalent to each other then that also applies to the limit. More generally:

Theorem 5.11. *Let $\sim$ be the separating equivalence closure of $\to_R$. Then $\twoheadrightarrow_p \subseteq \sim$.*

Proof. Immediate, because $\sim$ is pointwise closed, reflexive, transitive, and contains $\to_R$, and $\twoheadrightarrow_p$ is by definition the smallest such relation. ∎

Theorem 5.12. *A pre-order $\twoheadrightarrow_x$ is a semi-decider for the separating closure of $\twoheadrightarrow_x$ iff (i) $\uparrow_x \subseteq \downarrow_x$ and (ii) $\downarrow_x$ is pointwise closed.*

Proof. Let $=_x$ be the separating closure of $\twoheadrightarrow_x$. The "only if" part of the theorem is trivial.

For the "if" part, $=_x$ can be computed by repeatedly (and alternatingly) applying the equivalence closure and pointwise closure, starting with $\twoheadrightarrow_x$. The result can be shown by induction on the number of closures needed to derive a particular equation $t =_x u$.

Suppose $t =_x u$ is derived from an equivalence closure with $t = t_1 =_x t_2 =_x, \ldots, =_x t_n = u$, where each equation $t_i =_x t_{i+1}$ has a shorter derivation. Then by induction hypothesis $t_i \downarrow_x t_{i+1}$, and by repeatedly applying condition (i) we have $t \downarrow_x u$.

Alternatively $t =_x u$ is derived from the pointwise closure, i.e. u is in the closure of the set of all t_i with $t =_x t_i$ and where these equations have shorter derivations. Then, for all i, $t \downarrow_x t_i$ by induction hypothesis and $t \downarrow_x u$ by condition (ii). ∎

Although trivial in the proof, the "only if" part of the theorem shows that confluence alone (condition (i)) is insufficient to make $\twoheadrightarrow_x$ a semi-decider. Confluence properties of $\twoheadrightarrow_x$ aside (for $\twoheadrightarrow_a$ and $\twoheadrightarrow_p$ this is entirely new territory), when is its joinability relation pointwise closed? A sufficient condition is the following:

Theorem 5.13. *If $\twoheadrightarrow_x$ is pointwise closed and usc then $\downarrow_x$ is pointwise closed.*

Proof. Let $t \twoheadrightarrow_x a_i \twoheadleftarrow_x u_i$ with $i \in I$ for some index set I. Let $u \in \mathrm{Cl}(\{u_i \mid i \in I\})$.

Let $A = \mathrm{Cl}(\{a_i \mid i \in I\})$ then $t \twoheadrightarrow_x a$ for any $a \in A$ since $\twoheadrightarrow_x$ is pointwise closed. The set $\twoheadrightarrow_x^{-1}(A)$ clearly contains all the u_i and as $\twoheadrightarrow_x$ is usc it must be a closed set, so it does contain u too. Hence $u \twoheadrightarrow_x a'$ for some $a' \in A$ and thus $t \downarrow_x u$. ∎

In the previous examples in this section the separating closure of $\rightarrow_R$ was not just pointwise closed but also closed. While this is often the case there are exceptions, which means that we do not have that $\twoheadrightarrow_t$ is always included in the separating closure of $\rightarrow_R$; for example, if the original equivalence just contained $F(A^n(x)) \sim F(B^n(x))$ for all n then the separating closure would not add $F(A^\infty) \sim F(B^\infty)$, so it would not be closed.

Definition 5.14. An equivalence relation $\sim$ on a topological space A is called *strongly separating* iff its graph is closed.

Proposition 5.15. *Being strongly separating is a closable property.*

Proof. Trivial, as any intersection of closed sets is closed. ∎

Equally trivially, the strongly separating closure of $\rightarrow_R$ contains $\twoheadrightarrow_t$, by the analogous argument to theorem 5.11. Strongly separating equivalences (almost) characterise Hausdorff spaces:

Proposition 5.16. *If $A/\!\sim$ is T_2 then $\sim$ is strongly separating. If $\sim$ is strongly separating and $[\_]_\sim$ is an open map then $A/\!\sim$ is T_2. Or: if A is compact and $\sim$ is strongly separating then $A/\!\sim$ is T_2.*

Again, these are well-established results in topology. The first part of Proposition 5.16 can influence our choice for the class of algebraic models: if all our algebraic models are Hausdorff (e.g. metric spaces) then their induced equational theory is automatically strongly separating. Regarding the second part, open maps are functions that map open sets to open sets. For infinitary rewriting, the map $[\_]_\sim$ (where $\sim$ is the strongly separating equivalence closure of $\rightarrow_R$) is typically not open though, e.g.:

Example 5.17. Take the iTRS with rule: $F(x) \rightarrow G(x, x)$. Take as open set an ϵ-ball around $F(S^\infty)$. Let $\sim$ be the strongly separating equivalence closure of $\rightarrow_R$. The union of the $\sim$-equivalence classes of all terms in the ϵ-ball contains $G(S^\infty, S^\infty)$; it is not open, because it does not contain $G(S^\infty, S^n)$ for any finite n.

However, one can argue that strongly separating equivalence characterise T_2 separation for infinitary rewriting in most cases, because of the third part of Proposition 5.16 and the following:

Theorem 5.18. *If the signature Σ is finite the set of ground terms of $Ter^\infty(\Sigma)$ is compact.*

Proof. Let $f : \omega \to Ter^\infty(\Sigma)$ be a sequence of infinitary ground terms. Because Σ is finite at least one function symbol F occurs infinitely often as root symbol in f. Within the infinite subsequence of f that always has F as root, infinitely many times at least one function symbol G will occur infinitely many times in the first argument position. Iterating this argument leads to an accumulation point of f which is a ground term. ∎

This argument fails for any term metric m such that $Ter^m(\Sigma)$ (the metric completion of terms under d_m [6]) is not homeomorphic to $Ter^\infty(\Sigma)$, because sequences of finite ground terms that converge under d_∞ but not under d_m have no accumulation point in $Ter^m(\Sigma)$.

Without the finiteness of Σ the pigeon-hole principle in the proof fails. The restriction to ground terms is necessary as there are infinitely many variables: take the associative and commutative theory of a binary function symbol G, then the equivalence class of $G(x_1, G(x_2, \ldots))$ is not compact.

Theorem 5.19. *A pre-order $\twoheadrightarrow_x$ is a semi-decider for the strongly separating closure of $\twoheadrightarrow_x$ iff (i) $\uparrow_x \subseteq \downarrow_x$ and (ii) $\downarrow_x$ is closed.*

Proof. Analogous to Theorem 5.12. ∎

Again, this raises the issue under which conditions $\downarrow_x$ is closed.

Theorem 5.20. *If $\twoheadrightarrow_x$ is closed and usc then $\downarrow_x$ is closed.*

Proof. Let $a_i \twoheadrightarrow_x c_i \twoheadleftarrow_x b_i$ for all $i \in I$ and let a and b limits of the sequences a_n and b_n, respectively. We need to show $a \downarrow_x b$.

Let $C = \mathrm{Cl}(\{c_i \mid i \in I\})$ and $A = \{a' \mid a \twoheadrightarrow_x a'\}$. A is closed because $\twoheadrightarrow_x$ is a closed relation, and therefore $C' = C \cap A$ is closed too. Because $\twoheadrightarrow_x$ is usc the set $\twoheadrightarrow_x^{-1}(C)$ must contain a and thus C' is non-empty. If $\twoheadrightarrow_x^{-1}(C')$ contains infinitely many of the a_i then it also contains the corresponding b_i and (by the usc property) their limit b, hence $a \downarrow_x b$. Otherwise, all elements $c \in C'$ arise as limits of subsequences of the c_i and thus $b \twoheadrightarrow_x c$ for any such c by closedness of $\twoheadrightarrow_x$, and thus again $a \downarrow_x b$. ∎

6. Models

The concept of a model allows to reason semantically about term rewriting. For infinitary rewriting, one can look for two different kinds of models, algebraic and co-algebraic models.

The literature has focussed on algebraic ones which are more difficult to get right — several notions of model in the literature are indeed flawed, see below. The reason they are trickier is that in order for semantic reasoning through an algebraic model A to be sound there needs to be a unique and continuous interpretation of $Ter^\infty(\Sigma)$ into A, and for that it does not suffice for A to be a Σ-algebra, since this does not account for infinite terms.

In [2] the issue of interpreting infinite terms was side-stepped: only interpretations of finite terms were provided a priori (thus rules were not allowed infinite terms in their right-hand sides). In fact, in the special case of equational models (rather than partially ordered ones), the construction in [2] specialises exactly to ordinary Σ-algebras satisfying a set of equations.

A class of algebraic models induces an equivalence relation on $Ter^\infty(\Sigma)$, i.e. t and u are equivalent iff they are the same in any model of that class. When looking at the equivalence derived from all models satisfying a set of equations we get the congruence closure of this set — when looking at models of $Ter(\Sigma)$. This is different for models of $Ter^{\infty}(\Sigma)$, because the extra structure required in this class to guarantee for (unique) interpretations of infinite terms makes more equations true.

In [10] Zantema interpreted iTRSs in weakly monotone Σ-algebras with some extra structure. A special case are ordinary Σ-algebras since the ordering can be chosen to be equality. The extra structure required on a model A comprised there of: (i) a metric d_A; (ii) continuity of f_A for every function symbol f, w.r.t. the topology induced by the metric; (iii) the interpretation of the sequences $\mathrm{trunc}(t,n)$ in A converges for every infinite ground term t. Here, $\mathrm{trunc}(t,n)$ replaces all subterms of t at depth n with the fixed constant c.

As a notion of model this is flawed (see [4]) because the interpretation function from $Ter^{\infty}(\Sigma)$ to A derived from that is in general not continuous. Indeed, even the fix provided in [4] is insufficient because it is only expressed for ground terms, and any infinite term has non-ground terms in any neighbourhood; consequently, if $a \in A$ is arbitrary then the sequence $t_0 = a$, $t_{n+1} = f_A(t_n)$ could converge to a value different from the interpretation of f^{∞}.

This can be fixed by relating the notions of distance in A and $Ter^{\infty}(\Sigma)$: we can require the following condition for the distance function d_A on A:

$$d_A(f_A(a_1,\ldots,a_n), f_A(b_1,\ldots,b_n)) \leq \frac{1}{2} \cdot \max_{1\leq i\leq n} d_A(a_i,b_i)$$

This seemingly arbitrary condition derives as a special case from a more general condition we can set for metric models of other continuous term metrics [6].

Definition 6.1. Given a signature $\Sigma = (\mathcal{F}, \#)$ and continuous term metric m, a *metric model* is a Σ-algebra A, equipped with a metric $d_A : A \times A \to [0,1]$ such that (A, d_A) is a complete metric space and

$$d_A(f_A(a_1,\ldots,a_n), f_A(b_1,\ldots,b_n)) \leq f_m(d_A(a_1,b_1),\ldots,d_A(a_n,b_n))$$

for each $f \in \mathcal{F}, n = \#(f), a_i, b_i \in A$.

Here, f_m is the ultra-metric map associated with function symbol f in term metric m [6]. The condition on the order ensures that each f_A is continuous, but more importantly it leads to the following result:

Lemma 6.2. *Let A be a metric model (for Σ and m). Then the (unique) Σ-algebra homomorphism $[\![\_]\!]_A : Ter(\Sigma) \to A$ is non-expansive.*

Proof. We need to show $d_A([\![t]\!]_A, [\![u]\!]_A) \leq d_m(t,u)$ for all finite terms t and u (note that $Ter(\Sigma)$ is the set of *finite* terms). The argument goes by induction on the size of t. If t is a constant then either $d_m(t,u) = 0$ which implies $t = u$ and thus $[\![t]\!]_A = [\![u]\!]_A$ and $d_A([\![t]\!]_A, [\![u]\!]_A) = 0$; or $d_m(t,u) = 1$ which is an upper bound for d_A.

Otherwise, $t = F(t_1,\ldots,t_n)$. Again, if $d_m(t,u) = 1$ the condition holds because 1 is an upper bound for d_A. If $d_m(t,u) \neq 1$ then u must be of the form $u = F(u_1,\ldots,u_n)$ and $d_m(t,u) = F_m(d_m(t_1,u_1),\ldots,d_m(t_n,u_n))$. Using the induction hypothesis on the t_i, monotonicity of F_m (all ultra-metric maps are monotonic) and the metric model order property we get:

$$\begin{aligned} d_m(t,u) =& d_m(F(t_1,\ldots,t_n),F(u_1,\ldots,u_n))\\ =& F_m(d_m(t_1,u_1),\ldots,d_m(t_n,u_n))\\ \geq& F_m(d_A(\llbracket t_1\rrbracket_A,\llbracket u_1\rrbracket_A),\ldots,d_A(\llbracket t_n\rrbracket_A,\llbracket u_n\rrbracket_A))\\ \geq& d_A(F_A(\llbracket t_1\rrbracket_A,\ldots,\llbracket t_n\rrbracket_A),F_A(\llbracket u_1\rrbracket_A,\ldots,\llbracket u_n\rrbracket_A))\\ =& d_A(\llbracket F(t_1,\ldots,t_n)\rrbracket_A,\llbracket F(u_1,\ldots,u_n)\rrbracket_A)\\ =& d_A(\llbracket t\rrbracket_A,\llbracket u\rrbracket_A) \end{aligned}$$

■

Lemma 6.2 implies that $\llbracket \_ \rrbracket_A$ can be uniquely lifted to the respective metric completions (maintaining continuity), because non-expansive maps are uniformly continuous. Since A is already complete this lifts $\llbracket \_ \rrbracket_A$ to type $Ter^m(\Sigma) \to A$. Therefore we get an interpretation of infinite terms "for free".

Definition 6.3. The metric theory of a set of $Ter^m(\Sigma)$-equations E is the set of all pairs $(t,u) \in Ter^m(\Sigma) \times Ter^m(\Sigma)$ such that the equation $t = u$ holds in all metric models (w.r.t. signature Σ and term metric m) that satisfy the equations in E.

Theorem 6.4. *The metric theory of E w.r.t. Σ and m is strongly separating.*

Proof. We need to show that the theory is closed under limits of converging sequences. A sequence of pairs is converging in $Ter^m(\Sigma) \times Ter^m(\Sigma)$ iff the sequences of its first and second projections to $Ter^m(\Sigma)$ are. Suppose p is a sequence of pairs in the metric theory and that it is converging. Let the sequences l_n and r_n be their first and second projections, with limits l and r, respectively. Consider any metric model A of E. Using lemma 6.2 it is clear that the interpretation preserves convergence, hence $\llbracket l_n\rrbracket_A$ and $\llbracket r_n\rrbracket_A$ converge to $\llbracket l\rrbracket_A$ and $\llbracket r\rrbracket_A$. Since each pair $p_i = (l_i, r_i)$ is in the metric theory we have $\llbracket l_i\rrbracket_A = \llbracket r_i\rrbracket_A$ for all i. Therefore these sequences are identical in A and thus $\llbracket l\rrbracket_A = \llbracket r\rrbracket_A$. As this holds for any metric model A the pair (l, r) must be in the metric theory too. ■

In general, we cannot always construct an initial model for E (from its theory), but a sufficient condition is that $Ter^m(\Sigma)$ (restricted to ground terms) is compact.

Theorem 6.5. *Any set of $Ter^\infty(\Sigma)$-equations E has an initial model, if Σ is finite.*

Proof. The proof goes by constructing the model I; most of the argument works any metric m and signature Σ: we can quotient the ground terms of $Ter^m(\Sigma)$ by the metric theory of E and set a distance function d_I as the pointwise supremum of all distance functions of models satisfying E.

Since E is included in its metric theory I clearly satisfies E.

Checking the triangle inequality: $\forall A. d_A(t,u) \leq d_A(s,t)+d_A(s,u)$ implies $\forall A. d_A(t,u) \leq d_I(s,t) + d_I(s,u)$ and thus also $d_I(t,u) = \sup_A(d_A(t,u)) \leq d_I(s,t) + d_I(s,u)$.

Checking the zero-axiom: $d_I(t,u) = 0$ if $d_A(t,u) = 0$ for all A, i.e. if $t = u$ in all models. Hence $t = u$ in I too. If $t = u$ in I then $t = u$ in all A, hence $d_A(t,u) = 0$ in all A and so $d_I(t,u) = 0$.

Checking the metric model inequation:

$$\begin{aligned}
&d_I(f_I([\![a_1]\!]_I,\ldots,[\![a_n]\!]_I), f_I([\![b_1]\!]_I,\ldots,[\![b_n]\!]_I))\\
&\quad= d_I([\![f(a_1,\ldots,a_n)]\!]_I, [\![f(b_1,\ldots,b_n)]\!]_I)\\
&\quad= \sup_A(d_A([\![f(a_1,\ldots,a_n)]\!]_A, [\![f(b_1,\ldots,b_n)]\!]_A))\\
&\quad= \sup_A(d_A(f_A([\![a_1]\!]_A,\ldots,[\![a_n]\!]_A), f_A([\![b_1]\!]_A,\ldots,[\![b_n]\!]_A)))\\
&\quad\le \sup_A f_m(d_A([\![a_1]\!]_A,[\![b_1]\!]_A),\ldots,d_A([\![a_n]\!]_A,[\![b_n]\!]_A))\\
&\quad\le f_m(\sup_A(d_A([\![a_1]\!]_A,[\![b_1]\!]_A),\ldots,\sup_A(d_A([\![a_n]\!]_A,[\![b_n]\!]_A))))\\
&\quad= f_m(d_I([\![a_1]\!]_I,[\![b_1]\!]_I),\ldots,d_I([\![a_n]\!]_I,[\![b_n]\!]_I))
\end{aligned}$$

We still need to show that the metric space (I, d_I) is complete, and for this we need theorem 5.18, and thus specialise the metric. The semantic interpretation $[\![\_]\!]_I$ is continuous, and as its domain is compact and codomain T_2, it is also a closed map. Hence the limit of any Cauchy-sequence in (I, d_I) is in the image of that interpretation. ■

7. Conclusion and Remark

We had a look at large number of alternative notions for transfinite reduction relations and transfinite equivalences. Of particular interest are the relations $\twoheadrightarrow_f$ and $\twoheadrightarrow_t$, as they are the ones most tightly coupled with transfinite equivalences, i.e. the separating and strongly separating equivalence closures of rewrite steps.

Notice that several of the examples show "undesirable" consequences in equational reasoning. This is partly due to the fact that the formation of $Ter^\infty(\Sigma)$ implicitly requires that every function F has a unique fixpoint F^∞. These problems can be addressed by a different choice of term metric that disallows certain infinite terms, because F can have multiple (or no) fixpoints in the absence of F^∞.

References

[1] Hendrik P. Barendregt. *The Lambda-Calculus, its Syntax and Semantics*. North-Holland, 1984.

[2] Nachum Dershowitz and Stéphane Kaplan. Rewrite, Rewrite, Rewrite, Rewrite, Rewrite, ... In *Principles of Programming Languages*, pages 250–259. ACM, 1989.

[3] Nachum Dershowitz, Stéphane Kaplan, and David Plaisted. Rewrite, Rewrite, Rewrite, ... *Theoretical Computer Science*, 83(1):71–96, 1991.

[4] Jörg Endrullis, Clemens Grabmayer, Dimitri Hendriks, Jan Willem Klop, and Roel Vrijer. Proving infinitary normalization. pages 64–82, 2009.

[5] Theordore W. Gamelin and Robert Everist Greene. *Introduction to Topology*. Dover, 1999.

[6] Stefan Kahrs. Infinitary rewriting: Meta-theory and convergence. *Acta Informatica*, 44(2):91–121, May 2007.

[7] Stefan Kahrs. Modularity of convergence in infinitary rewriting. In Ralf Treinen, editor, *Rewriting Techniques and Applications*, volume 5595 of *LNCS*, pages 179–193. Springer, 2009.

[8] Richard Kennaway, Jan Willem Klop, Ronan Sleep, and Fer-Jan de Vries. Transfinite reductions in orthogonal term rewriting systems. *Information and Computation*, 119(1):18–38, 1995.

[9] Donald E. Knuth and Peter B. Bendix. Simple word problems in universal algebra. In J. Leech, editor, *Computational Problems in Abstract Algebra*, pages 263–297. Pergamon, 1970.

[10] Hans Zantema. Normalisation of infinite terms. In *RTA 2008*, volume 5117 of *LNCS*, pages 441–455, 2008.