

REWRITE, REWRITE, REWRITE, REWRITE, REWRITE, ...*

Nachum DERSHOWITZ[†]

*Department of Computer Science, University of Illinois at Urbana-Champaign,
1304 West Springfield Ave., Urbana, IL 61801-2987, U.S.A.*

Stéphane KAPLAN[‡]

*Department of Computer Science, Bar-Ilan University, Ramat-Gan 52100, Israel
Laboratoire de Recherche en Informatique, Université de Paris-Sud, Orsay 91405,
France*

David A. PLAISTED[§]

*Department of Computer Science, University of North Carolina at Chapel Hill,
Chapel Hill, NC 27514-3175, U.S.A.*

Communicated by

Received

Revised

Abstract. We study properties of rewrite systems that are not necessarily terminating, but allow instead for transfinite derivations that have a limit. In particular, we give conditions for the existence of a limit and for its uniqueness and relate the operational and algebraic semantics of infinitary theories. We also consider sufficient completeness of hierarchical systems.

Is there no limit?

—Job 16:3

1. Introduction

Rewrite systems are sets of directed equations used to compute by repeatedly replacing equal terms in a given formula, as long as possible. For one approach to their use in computing, see [23]. The theory of rewriting is an outgrowth of the study of the lambda calculus and combinatory logic, and

*Preliminary versions [6, 7] of ideas in this paper were presented at the *Sixteenth ACM Symposium on Principles of Programming Languages*, Austin, TX (January 1989) and at the *Sixteenth EATCS International Colloquium on Automata, Languages and Programming*, Stresa, Italy (July 1989).

[†]Supported in part by the U.S. National Science Foundation under Grant DCR 85-13417.

[‡]Supported in part by the European Communities under ESPRIT project 432 (METEOR).

[§]Supported in part by the U.S. National Science Foundation under Grant CCR 88-02282.

has important applications in abstract data type specifications and functional programming languages. For surveys of the theory of rewriting, see [13, 16], or [8]; our notations conform to the latter.

A key property for rewrite systems is that every term rewrites to a unique normal form. This is usually decomposed into two requirements: “normalization”, which ensures that at least one normal form always exists; and “confluence”, which ensures that there can never be more than one normal form.

In this paper, we consider systems that have *infinite* terms as normal forms. Such systems are not terminating in the classical sense; instead we develop a notion of “ ω -converging”, the property that any (infinite) derivation has a limit (not necessarily a normal form).¹ Under certain conditions, if a system is ω -converging then it is also “ ω -normalizing”, that is, there is a limit that is in normal form. We then investigate “ ω -confluence”, a property that ensures uniqueness of normal forms. Together, these properties imply the existence of a (potentially infinite) unique normal form for any input term, which can be viewed as the “value” of the term initiating the derivation. As a programming language, rewrite systems have the full power of Turing machines; hence, these properties, and others, are easily shown to be undecidable in general. Our results may have implications for stream-based programming languages.

The next section defines the basic concept of normal forms in the context of transfinite chains. Section 3 considers properties of rewrite systems and the length of their derivations. Section 4 characterizes derivations that lead to normal forms. Section 5 presents methods for establishing ω -normalization and Section 6, for ω -confluence. With these operational notions in place, Section 7 gives “algebraic” semantics to infinite rewriting; it is followed by a section on the semantics of hierarchically typed systems. We conclude with a brief discussion.

Algebraic semantics involving infinite terms were developed by [3, 4, 21, 19a] from a different perspective (particularly, the formal meaning of recursive program schemata). The work in [9, 19, 27] is more closely related to our approach. Recent developments, following our preliminary work [6, 7] in this area, include [2, 10, 15].

2. Infinite chains

We are interested in properties of binary relations. Let \rightarrow denote any binary relation, and \rightarrow^* its reflexive-transitive closure. We use \leftarrow and $^*\leftarrow$

¹This property should not be confused with (finite) “convergence” of (finitely) terminating and confluent systems (to unique normal forms)—in the terminology of [8].

for their respective inverses. A relation \rightarrow over a set S is said to be *finitely terminating* if there exist no infinite chains $s_0 \rightarrow s_1 \rightarrow \cdots \rightarrow s_n \rightarrow \cdots$ of elements s_i in S ; it is *finitely confluent* if for any elements s, t, u in S such that $u \rightarrow^* s$ and $u \rightarrow^* t$, it is the case that $s \rightarrow^* v$ and $t \rightarrow^* v$ for some v in S . Confluence may be expressed as the set-theoretic inclusion $\rightarrow^* \circ \rightarrow^* \subseteq \rightarrow^* \circ \rightarrow^*$, where \circ denotes composition of relations.

This paper concerns limits of transfinite sequences of elements. Let $(s_\beta)_{\beta < \alpha}$ be a finite or transfinite sequence of elements of a topological space S , indexed by ordinals β less than some ordinal number α . We say that t is the *limit* of the sequence, written $\lim_{\beta \rightarrow \alpha} s_\beta = t$, if for any neighborhood V of t , there exists an ordinal $\gamma < \alpha$ such that s_β is in V for all β between γ and α .²

Infinite “chains” are defined by transfinite induction:

Definition 1. Given a binary relation \rightarrow on a topological space, its α -*iterate* \rightarrow^α , for given ordinal α , is defined as follows:

- (a) if $\alpha = 0$, then \rightarrow^α is the identity relation;
- (b) if α is a successor ordinal $\beta + 1$, then $\rightarrow^\alpha = \rightarrow^\beta \cup (\rightarrow^\beta \circ \rightarrow)$;
- (c) if α is a limit ordinal, then $s_0 \rightarrow^\alpha t$ if $s_0 \rightarrow^\beta t$ for some $\beta < \alpha$ or if there exist elements $(s_\beta)_{\beta < \alpha}$ such that $s_\gamma \rightarrow^\beta s_\beta$ for all $\gamma < \beta < \alpha$ and $\lim_{\beta \rightarrow \alpha} s_\beta = t$.

Definition 2. An α -*chain* for a binary relation \rightarrow over a topological space and ordinal α is a finite or transfinite sequence $(s_\beta)_{\beta < \alpha}$ such that $s_\gamma \rightarrow^\beta s_\beta$ for all $\gamma < \beta < \alpha$.

In particular, $s \rightarrow^\omega t$, for s and t in S , if $s \rightarrow^* t$ or if there is a chain $s = s_0 \rightarrow s_1 \rightarrow \cdots \rightarrow s_n \rightarrow \cdots$ such that the limit of the s_n , as n goes to infinity ($n \rightarrow \infty$), is t .

Definition 3. A binary relation \rightarrow over a topological space is α -*closed* if $\rightarrow^\beta = \rightarrow^\alpha$ for any ordinal $\beta \geq \alpha$.

Lemma 1. A binary relation \rightarrow over a metric space is ω -*closed* if $\rightarrow^\omega \circ \rightarrow \subseteq \rightarrow^\omega$.

Proof. Transfinite induction establishes that $\rightarrow^\alpha = \rightarrow^\omega$ for successor ordinals and limit ordinals $\alpha \geq \omega$:

The hypothesis means that $\rightarrow^{\omega+1} = \rightarrow^\omega$.

Suppose α is a successor ordinal $\beta + 1 > \omega$. By induction, $\rightarrow^\beta = \rightarrow^\omega$. Hence, $\rightarrow^\alpha = \rightarrow^\omega \circ \rightarrow$, which, by hypothesis, is contained in \rightarrow^ω (which, by definition, is contained in \rightarrow^α).

²We use this notation for limits instead of the more precise $\lim_{\beta \rightarrow \alpha^-}$.

Suppose $\alpha > \omega$ is a limit ordinal and $s_0 \rightarrow^\alpha s_\alpha$ for elements s_0 and s_α of the space. There must be an ordinal $\gamma < \alpha$ indexing an element of the transfinite sequence $s_0 \rightarrow^\gamma s_\gamma \rightarrow^\alpha s_\alpha$ such that the distance (in the metric space) between s_γ and s_α is less than $1/4$. By induction, $s_0 \rightarrow^\omega s_\omega$. Far enough along this sequence, there is an element s_{β_1} of distance less than $1/2$ from s_α . We have $s_0 \rightarrow^* s_{\beta_1} \rightarrow^\alpha s_\alpha$, and this construction can be continued to form an ω -chain $s_0 \rightarrow^* s_{\beta_1} \rightarrow^* s_{\beta_2} \rightarrow \dots$ with limit s_α . \square

Definition 4. A binary relation \rightarrow over a topological space is α -converging, for ordinal α , if for any α -chain $(s_\beta)_{\beta < \alpha}$ of elements s_β of S , the limit $\lim_{\beta \rightarrow \alpha} s_\beta$ exists.

Definition 5. A binary relation \rightarrow over a topological space S is α -normalizing, for ordinal α , if for every s_0 in S there is an α -chain $(s_\beta)_{\beta < \alpha}$ of elements of S (starting with s_0), such that the limit $s_\alpha = \lim_{\beta \rightarrow \alpha} s_\beta$ exists and if $s_\alpha \rightarrow s'$, then $s' = s_\alpha$. It is *uniquely α -normalizing* if it is α -normalizing and for every element s_0 of S , and every such α -chain, the α -normal form s_α is the same.

By $\alpha \leftarrow$, we will denote the inverse of the relation \rightarrow^α .

Definition 6. A binary relation \rightarrow over a topological space is α -confluent, for ordinal α , if $\alpha \leftarrow \circ \rightarrow^\alpha \subseteq \rightarrow^\alpha \circ \alpha \leftarrow$.

Note that α -confluence implies that there can be at most one α -normal form, since were one element s to lead to two limits u and v , then by confluence $u \rightarrow^\alpha \circ \alpha \leftarrow v$, which for α -normal forms u and v could only be if $u = v$. Without α -normalization, it could be that all α -chains have limits that are not α -normal forms. For an α -converging relation all chains must “end”, but none need end in α -normal forms.

Proposition 1. *A binary relation over a topological space is uniquely α -normalizing, for ordinal α , if it is α -normalizing and α -confluent.*

3. Infinite derivations

We are particularly interested in relations over *terms*. Let $\mathcal{T}(\mathcal{F}, \mathcal{X})$, or just \mathcal{T} , denote a set of *finite* (first-order) terms containing function symbols and constants from some finite vocabulary (signature) \mathcal{F} and variables from some denumerable set \mathcal{X} . Let $\mathcal{T}^\infty(\mathcal{F}, \mathcal{X})$ (or just \mathcal{T}^∞) denote the set of finite and *infinite* terms over the same vocabulary and variable set. (Infinite terms are like infinite ordered trees, rooted at their outermost symbol, with finite outdegree at each node. Their nodes are all connected by paths of finite length, though some paths may be of length ω .) The set of finite

ground (variable-free) terms is $\mathcal{G}(\mathcal{F}) (= \mathcal{T}(\mathcal{F}, \emptyset))$, or just \mathcal{G} ; the set of finite and infinite ground terms is $\mathcal{G}^\infty(\mathcal{F})$, or just \mathcal{G}^∞ . In the sequel, we use lower-case Greek letters for ordinals, s and t for terms in \mathcal{T}^∞ , x and y for variables in \mathcal{X} , and various lower-case Latin letters for operators and constants.

A *position* in (finite or infinite) terms may be represented as a *finite* sequence of positive integers, giving the path to that position in the ordered-tree representation of terms. The empty sequence Λ denotes the outermost (topmost) position. The subterm of a term t in \mathcal{T}^∞ rooted at position p will be denoted $t|_p$, and—in particular—the i th immediate subterm of t is $t|_i$. The *depth* $|p|$ of a position p is the length of the sequence representing it, which corresponds to the level of the subterm in the tree representation of the term. *Disjoint* positions are such that neither is below the other.

A *distance* d is defined on \mathcal{T}^∞ and \mathcal{G}^∞ as follows: Let $d(s, t) = 1/2^{v(s, t)}$, where the natural number $v(s, t)$ is the smallest depth of a symbol occurrence at which terms s and t differ—with the convention that $d(t, t) = 0$. Finite and infinite (first-order or ground) terms, with this distance, form a complete ultra-metric space [21].

A *rewrite system* R is a *finite* family of pairs (l, r) of (finite) terms of $\mathcal{T}(\mathcal{F}, \mathcal{X})$, each written in the form $l \rightarrow r$; we will assume that all variables appearing on a right-hand side r also appear on the corresponding left-hand side l . A system R defines a *rewrite relation* \rightarrow_R over \mathcal{T}^∞ as follows: For t in \mathcal{T}^∞ , we say that t *rewrites* via R to t' , and write $t \rightarrow_R t'$ (or simply $t \rightarrow t'$), if there exists a rule $l \rightarrow r$ in R , a context $c[\cdot]_p$, where c is a term in \mathcal{T}^∞ and p is a position in c , and a substitution $\sigma : \mathcal{X} \rightarrow \mathcal{T}^\infty$ such that $t = c[l\sigma]_p$ (the subterm $t|_p$ of t is an instance of the left-hand side l) and $t' = c[r\sigma]_p$ (t' is the result of replacing the subterm at p with the corresponding instance $r\sigma$ of the right-hand side). A position p in t at which a rewrite can take place is called a *redex*. We use $|R|$ to denote the maximum depth of a left-hand side of a system R .

Since we are interested here primarily in sequences of rewrites issuing from finite terms t_0 (unlike [4, 10]), we will restrict our attention to that case:³

Definition 7. A *derivation* of length α , for ordinal α , is a finite or transfinite sequence of (finite or infinite) terms t_β in \mathcal{T}^∞ , such that t_0 is a finite term in \mathcal{T} and $(t_\beta)_{\beta < \alpha}$ is an α -chain for \rightarrow_R .

In particular, if (t_β) is of length $\omega + 1$, then, for any depth d , there is a point N such that for all n , $N \leq n < \omega$, the distance between t_n and the limit t_ω is no more than $1/2^d$.

³For many of the results we report on here, this assumption is not critical.

For example, the system

$$a \rightarrow f(a) \tag{1}$$

has a derivation

$$a \rightarrow_R f(a) \rightarrow_R f(f(a)) \rightarrow_R \cdots \rightarrow_R f(f \dots f(a) \dots) \rightarrow_R \cdots f^\omega$$

of length $\omega + 1$, where the limit f^ω of the chain $(f^n(a))_{n < \omega}$ is the infinite term $f(f(f(\dots)))$, composed of infinitely many occurrences of the unary symbol f , and in which the constant a no longer occurs.

We will say that a rewrite system R is α -closed if $s \rightarrow^\beta t$ implies $s \rightarrow^\alpha t$, for all finite terms s in \mathcal{T} , finite or infinite terms t in \mathcal{T}^∞ , and ordinals $\beta \geq \alpha$. In particular, ω -closure means that all (finite or infinite) terms derivable from a finite term are limits of derivations containing only finite terms. Finitely terminating systems are ω -closed.

The systems

$$\begin{aligned} a &\rightarrow b \\ b &\rightarrow a \end{aligned} \tag{2}$$

and

$$\begin{aligned} f(x) &\rightarrow g(f(x)) \\ g(x) &\rightarrow f(g(x)) \end{aligned} \tag{3}$$

are ω -closed. The following system is not ω -closed:

$$\begin{aligned} a &\rightarrow g(a) \\ b &\rightarrow g(b) \\ f(x, x) &\rightarrow c \end{aligned} \tag{4}$$

We have

$$f(a, b) \rightarrow_R^\omega f(g^\omega, g^\omega) \rightarrow_R c$$

but $f(a, b) \not\rightarrow_R^\omega c$. In fact, this system is closed only at $\omega \times 2$, since there are longer derivations like:

$$f(f(f(a, b), c), c) \rightarrow_R^\omega f(f(f(g^\omega, g^\omega), c), c) \rightarrow_R^2 f(c, c) \rightarrow_R c$$

Closure may occur farther up the ordinal hierarchy, too; for example, the system

$$\begin{aligned} a &\rightarrow g(a) \\ b &\rightarrow g(b) \\ f(x, x) &\rightarrow h(f(a, b)) \end{aligned} \tag{5}$$

has the following derivation:

$$f(a, b) \rightarrow_R^{\omega \times i} h^i(f(a, b)) \rightarrow_R^{\omega^2} h^\omega$$

The remainder of this section is devoted to ω -closure of rewrite systems. Subsequent sections deal with properties of normal forms.

Definition 8. A rewrite system R is *top-terminating* if there are no derivations of length ω with infinitely many rewrites at the topmost position Λ .

A rewrite system R is said to be *left-linear* if the left-hand side l of each rule $l \rightarrow r$ in R has at most one occurrence of any variable. As we will see in the proof of the following theorem, with left-linearity, a derivation $t_0 \rightarrow_R^\alpha t_\infty \rightarrow_R t'_\infty$ can be simulated by a derivation $t_0 \rightarrow_R^* t' \rightarrow_R^\alpha t'_\infty$, provided none of the steps in $t_0 \rightarrow_R^\alpha t_\infty$ are above or only slightly below the position of the step $t_\infty \rightarrow_R t'_\infty$.

Theorem 1. *If R is a left-linear top-terminating rewrite system, then R is ω -closed.*

This is essentially the same as the $\omega\omega$ -Lemma in [10], where only top-terminating derivations are considered.⁴

Proof. By Lemma 1, it is sufficient to show that $\rightarrow_R^\omega \circ \rightarrow_R \subseteq \rightarrow_R^\omega$. Let $t_0 \rightarrow_R t_1 \rightarrow_R t_2 \rightarrow_R \cdots \rightarrow_R^\omega t_\infty$, and suppose $t_\infty \rightarrow_R t'_\infty$ at position p via rule $l \rightarrow r$ with substitution σ . Since R is top-terminating, there exists an index N such that, for $n \geq N$, $t_n \rightarrow_R t_{n+1}$ at position p_n deeper than $|p| + |R|$, which is within the “variable-part” of $l \rightarrow r$. Then, on account of left-linearity, t_n is rewritable at $t_n|_{p_n}$ via $l \rightarrow r$ to the term $t'_n = t_n[r\sigma]_p$.

If p_n is at or below the position of a variable of l (that is, if the redex is a subterm of the term matching the variable), then $t'_n \rightarrow_R^* t'_{n+1}$, again due to the left-linearity of R . Otherwise, p and p_n are disjoint positions in t_n , and $t'_n \rightarrow_R t'_{n+1}$. In either case:

$$\begin{array}{ccccccc}
 t_0 & \rightarrow_R & \cdots & \rightarrow_R & t_{N-1} & \rightarrow_R & t_N & \rightarrow_R & t_{N+1} & \rightarrow_R & \cdots & \rightarrow_R^\omega & t_\infty \\
 & & & & & & R \downarrow & & R \downarrow & & & & R \downarrow \\
 & & & & & & t'_N & \rightarrow_R^* & t'_{N+1} & \rightarrow_R^* & \cdots & & t'_\infty
 \end{array}$$

Clearly, $\lim_{n \rightarrow \omega} t'_n = t'_\infty$, which proves that $t_0 \rightarrow_R^\omega t'_\infty$. \square

Neither hypotheses in this theorem suffices by itself. Example (4) is a non-left-linear, top-terminating system, for which the theorem does not hold. Left-linearity turns out to be crucial (as in [4], but cf. [11]), and throughout this paper, we deal exclusively with left-linear systems. Unfortunately, left-linearity is insufficient. The following is an example (from

⁴Our preliminary versions of this claim omitted the top-termination requirement.

[10]) of a left-linear non-top-terminating system that is not ω -closed:

$$\begin{array}{l} a \rightarrow b \\ f(x, a) \rightarrow f(g(x), a) \end{array} \quad (6)$$

We have

$$f(c, a) \rightarrow_R^\omega f(g^\omega, a) \rightarrow_R f(g^\omega, b)$$

but $f(c, a) \not\rightarrow_R^\omega f(g^\omega, b)$.

A term t is said to *overlap* a term t' if t —after renaming all its variables so as not to conflict with those in t' —unifies with a non-variable subterm of t' ; a system is *non-overlapping* if no left-hand side overlaps another (or itself at a proper subterm). A system that is both left-linear and non-overlapping is called *orthogonal* (or “regular”). Orthogonal systems are important since they are always confluent and can have at most one normal form [22], but since they need not be finitely terminating, there may be terms having no normal form.

Even orthogonal systems need not be ω -closed, when they are not top-terminating, as can be seen from the following (rather disconcerting) example:

$$\begin{array}{l} a \rightarrow b \\ f(x, y) \rightarrow f(x, g(x, y)) \end{array} \quad (7)$$

It allows an “outermost” derivation

$$f(a, c) \rightarrow_R^\omega f(a, g(a, g(a, \dots))) \rightarrow_R f(b, g(a, g(a, \dots)))$$

but $f(a, c) \not\rightarrow_R^\omega f(b, g(a, g(a, \dots)))$.

On the other hand, neither left-linearity, nor top-termination is necessary for ω -closure; just consider the simple rule

$$f(x, x) \rightarrow f(g(x), g(x)) \quad (8)$$

Even with left-linearity, ω -closure does not imply top-termination, as can be seen from the following trivial system:

$$x \rightarrow f(x) \quad (9)$$

4. Infinite normal forms

Definition 9. An α -normal form of a term s in \mathcal{T}^∞ , for a rewrite system R and ordinal α , is a term t in \mathcal{T}^∞ , such that $s \rightarrow_R^\alpha t$ and $t \rightarrow_R t'$ only if $t' = t$ (for $t' \in \mathcal{T}^\infty$).

Note that this does not imply that a normal form cannot be rewritten at all, but rather that it may rewrite only to itself. This allows a system like

$$f(x) \rightarrow f(f(x)) \quad (10)$$

to compute an ω -normal form f^ω of any term $f(t)$, just as Example (1) computes the ω -normal form f^ω of a .

Accordingly, we say that a system is α -converging if every derivation of length α issuing from a finite term in \mathcal{T} has a limit in \mathcal{T}^∞ , that it is α -normalizing if every finite term t in \mathcal{T} admits an α -normal form t_∞ in \mathcal{T}^∞ , and that it is *uniquely α -normalizing* if every finite term has exactly one α -normal form.⁵ Similarly, we say that a system is α -confluent if, for all finite terms u in \mathcal{T} and (finite or infinite) terms s and t in \mathcal{T}^∞ , $s \xrightarrow{\alpha} u \rightarrow^\alpha t$ implies the existence of a (possibly infinite) term v in \mathcal{T}^∞ such that $s \rightarrow^\alpha v \xrightarrow{\alpha} t$.

System (9) is ω -closed, non-finitely-terminating, and has no finite normal forms; yet it is both ω -converging and uniquely ω -normalizing. Systems (2) and (3), though ω -closed, are not ω -converging; System (4) is $(\omega \times 2)$ -closed, $(\omega \times 2)$ -converging, and ω -confluent, but has no ω -normal forms.

If a standard rewriting system is (finitely) normalizing and finitely confluent, then any finite term t has exactly one finite normal form, which can be taken as its “value”. Similarly, infinite normal forms can be considered the “value” of a term, when they are unique and lend themselves to approximation. Analogous to the finite case, combining existence of ω -normal forms with ω -confluence, gives uniqueness of ω -normal forms. By definition, any finitely-terminating system is ω -converging.

For systems, like (6), that are ω -normalizing, certain derivations always lead to normal forms. A “fair” computation is a derivation for which no redex persists forever. More precisely:

Definition 10. A derivation $t_0 \rightarrow_R t_1 \rightarrow_R \cdots \rightarrow_R t_n \rightarrow_R \cdots$ is *fair* if whenever there is a rule $l \rightarrow r$ in R and position p such that, for all n past some N , the subterm $t_n|_p$ is an instance $l\sigma_n$ of l , then (at least) one of the rule applications $t_n \rightarrow_R t_{n+1}$ ($n \geq N$) is an application of $l \rightarrow r$ at p .

Theorem 2. *Let R be a left-linear rewrite system. If a term t_0 in \mathcal{T} admits an ω -normal form t_∞ in \mathcal{T}^∞ , then there exists a fair derivation $t_0 \rightarrow_R t_1 \rightarrow_R \cdots \xrightarrow{\omega}_R t_\infty$ with limit t_∞ .*

⁵This property should not be confused with the (finite) “unique normal form” (i.e. at most one normal form) of property (finitely) confluent systems—in the terminology of [16].

Of course, unfair derivations can also need lead to normal forms. For example, either rule in

$$\begin{array}{l} f(x) \rightarrow f(f(x)) \\ f(x) \rightarrow f(f(f(x))) \end{array} \quad (11)$$

can be forever ignored.

Proof. Suppose that $t_0 \rightarrow_R t_1 \rightarrow_R \cdots \xrightarrow{\omega}_R t_\infty$, and t_∞ is a normal form. If the derivation is not fair, then for some point N' , position p , and rewrite rule r , the rule must be continually applicable at p in the subsequence $(t_n)_{n \geq N'}$, though not actually applied. Let $N \geq N'$ be such that for all $n \geq N$, we have $d(t_n, t_\infty) \leq 1/2^{|p|+|R|+|R|}$. Let t'_n denote the result of applying r to t_n at p . On account of the low positions of the rewrites, any changes incurred by the steps past N take place in the variable part of r . With left-linearity, this implies that the given derivation from t_N can be mimicked by a derivation issuing from t'_N :

$$\begin{array}{ccccccc} t_0 & \rightarrow_R & \cdots & \rightarrow_R & t_{N-1} & \rightarrow_R & t_N & \rightarrow_R & t_{N+1} & \rightarrow_R & \cdots & \xrightarrow{\omega}_R & t_\infty \\ & & & & & & r \downarrow & & r \downarrow & & & & r \downarrow \\ & & & & & & t'_N & \xrightarrow{*}_R & t'_{N+1} & \xrightarrow{*}_R & \cdots & & t_\infty \end{array}$$

Though the same rule also applies to t_∞ , since t_∞ is a normal form, it must be that the result of rewriting t_∞ is t_∞ itself. Because essentially the same rewrites are being applied to the t'_n , $d(t'_n, t_\infty) \leq 1/2^{|p|}$ for all $n \geq N$ and, moreover, $\lim_{n \rightarrow \omega} t'_n = t_\infty$. This process may be repeated, beginning at some $t'_{n'}$ (for $n' > N$) such that $d(t'_{n'}, t_\infty) \leq 1/2^{|p|+1}$, to obtain a fair derivation with t_∞ as the limit. \square

Fair derivations compute normal forms:

Theorem 3. *Let R be a left-linear rewrite system. For any fair derivation $t_0 \rightarrow_R t_1 \rightarrow_R \cdots \xrightarrow{\omega}_R t_\infty$, the limit t_∞ , if it exists, is an ω -normal form of t_0 .*

Proof. Suppose that t_∞ is not a normal form and that $t_\infty \rightarrow_R t'_\infty$ via some rule r at some position p . For all n greater or equal to some N , $d(t_n, t_\infty) \leq 1/2^{|p|+|R|}$. But with left-linearity, r may be applied at position p of each t_n ($n \geq N$), contradicting the fairness of the derivation. \square

With a weaker notion of fairness, in which applying any rule at or above the position of the persisting redex is fair enough, the limit is not necessarily a normal form. The second rule of System (7), for example, applied repeatedly to $f(a, c)$, leads to $f(a, g(a, g(a, \dots)))$, to which the first rule can still be applied.

Without left-linearity, even with top-termination, the limit of a fair derivation need not be a normal form. For example, the only normal form of terms in $\mathcal{G}(\{f, g, c\})$ with the ω -closed system

$$\begin{array}{lcl} c & \rightarrow & g(c) \\ f(x, x) & \rightarrow & c \end{array} \quad (12)$$

is g^ω , though there is a fair derivation

$$f(g(c), c) \rightarrow_R f(g(g(c)), c) \rightarrow_R f(g(g(c)), g(c)) \rightarrow_R \cdots \rightarrow_R^\omega f(g^\omega, g^\omega)$$

And, without left-linearity, even with ω -closure, there are normal forms that cannot be obtained fairly. For example, there is a derivation

$$h(f(g(c), g(c))) \rightarrow_R h(f(g(g(c)), g(g(c)))) \rightarrow_R \cdots \rightarrow_R^\omega h(f(g^\omega, g^\omega))$$

for

$$\begin{array}{lcl} f(x, x) & \rightarrow & f(g(x), g(x)) \\ h(f(x, g(y))) & \rightarrow & h(f(g(x), y)) \end{array} \quad (13)$$

that is not even weakly fair (in the above sense), but once the second rule is applied the normal form $h(f(g^\omega, g^\omega))$ is unreachable.

Since fair derivations must end in normal forms for left-linear ω -converging systems, it follows that at least one normal form exists:

Corollary. *If R is a left-linear ω -converging rewrite system, then R is ω -normalizing.*

Examples (2) and (4) demonstrate the need for both requirements.

5. Existence of normal forms

If a system R is finitely terminating, then any finite term has at least one finite normal form. For a survey of methods for establishing finite termination of rewrite systems, see [5]. In this section, we weaken this demand, and consider, instead, ω -converging systems, for which every derivation of length ω has a limit. We concentrate on special cases that are of practical importance.

A top-terminating system need not be finitely terminating. For instance, System (1) is top-terminating, while (10) is not. Neither is finitely terminating. The system

$$\begin{aligned} f(x) &\rightarrow x \\ g(x) &\rightarrow x \end{aligned} \tag{14}$$

is finitely terminating, and, hence, top-terminating for *finite* terms, but note that some derivations issuing from *infinite* terms like $f(g(f(g(\dots))))$ have no limit.

Proposition 2. *If R is a top-terminating rewrite system, then it is ω -converging.*

The converse does not hold. For example, System (10) is not top-terminating.

Proof. If R is top-terminating, then after a finite number of rewrites, no more rewrites are applied at the top, and the outermost symbol in the remainder of the derivation is fixed. The same argument can then be applied to the subterms to show that the rewrites must occur deeper and deeper, and, hence, that the sequence has a limit. \square

By the corollary to Theorem 3, top-termination, with left-linearity, guarantees existence of ω -normal forms. Left-linearity and top-termination do not, however, suffice for all limits to be normal forms:

$$\begin{aligned} a &\rightarrow b \\ c &\rightarrow f(a, c) \end{aligned} \tag{15}$$

Theorem 4. *Let \succ be a well-founded (partial) ordering on the set \mathcal{T} of finite terms, and let \succeq be a compatible quasi-ordering. Let R and S be two rewrite systems such that $\rightarrow_R \subseteq \succ$, $\rightarrow_S \subseteq \succeq$, and S is α -converging. Then the combined system $R \cup S$ is α -converging.*

A *quasi-ordering* is a reflexive and transitive binary relation. As usual, $s \succ t$ if $s \succeq t \not\succeq s$; also, $s \sim t$ means $s \succeq t \succeq s$. By “compatible”, we mean here that $\succ \circ \succ \subseteq \succ \circ \succeq$.

Proof. Let $t_0 \rightarrow_{R \cup S} t_1 \rightarrow_{R \cup S} \dots$ be a derivation in the combined system. Then, $t_0 \succeq t_1 \succeq t_2 \succeq \dots$. Were there an infinite number of R steps in the derivation, then by compatibility, there would be an infinite sequence of terms $t_0 \succ t'_1 \succ t'_2 \succ \dots$, contradicting well-foundedness. Hence, there must be a point t_N in the derivation after which only S steps appear, and by the fact that S is α -converging, the derivation has a limit. \square

As an example, if $R \cup S$ is

$$\begin{array}{l} h(g(x), f(y)) \rightarrow h(x, y) \\ a \rightarrow g(a) \end{array} \quad (16)$$

we can compare finite terms by comparing the total number of occurrences of the symbol f in them. Applying the finitely terminating rule $h(g(x), f(y)) \rightarrow h(x, y)$ reduces this number; applying the ω -converging rule $a \rightarrow g(a)$ effects no change. (This example is an adaptation of the finite termination method of [18].)

In some cases, one can use the transitive closure of \rightarrow_R for \succ and the reflexive-transitive closure of \rightarrow_S for \succeq . A rewrite system is *right-linear* if the right-hand sides have at most one occurrence of each variable.

Corollary. *Let R and S be two rewrite systems. If R is left-linear and finitely terminating, S is right-linear and α -converging, and the right-hand sides of S and left-hand sides of R do not overlap, then $R \cup S$ is α -converging.*

Under the stated circumstances, the “commutation” property $\rightarrow_S \circ \rightarrow_R \subseteq \rightarrow_R \circ \rightarrow_S^*$ holds; see [24].

This corollary does not have wide applicability, since the left-hand sides of finitely-terminating R may not refer to symbols used by S to construct an infinite structure. It does not, for example, apply to Example (16), since $g(a)$ unifies with the subterm $g(x)$ of the left-hand side of R .

To provide semantic methods of proving existence of limits, we define top-termination orderings for proofs of top-termination that are analogous to the well-founded quasi-orderings used to show finite termination (see [5]).

Definition 11. A quasi-ordering \succeq over a set of terms \mathcal{T} is a *top-termination ordering* if it has the replacement property: $s|_i \succeq t|_i$, for $i = 1, \dots, n$ implies $s \succeq t$ for all $s = f(s|_1, \dots, s|_n)$ and $t = f(s|_1, \dots, s|_n)$ in \mathcal{T} , and if its strict part \succ is well-founded.

Such orderings are not hard to devise. What is significant is what happens near the top of the term. For example, one can define a top-termination ordering on terms that is induced by a given quasi-ordering \succeq on operators in \mathcal{F} : if $f \succ g$ in the operator ordering, then any term $s = f(s|_1, \dots, s|_m)$ is greater than $t = g(t|_1, \dots, t|_n)$ in the term ordering; while if $f \sim g$, then s and t are equivalent. We always have $f(s|_1, \dots, s|_n) \sim f(t|_1, \dots, t|_n)$ in this ordering on terms.

Theorem 5. *A rewrite system R over $\mathcal{T}(\mathcal{F}, \mathcal{X})$ is top-terminating if there exists a top-termination ordering \succeq such that $l\sigma \succ r\sigma$ for all rules $l \rightarrow r$ in R and finite substitutions $\sigma : \mathcal{X} \rightarrow \mathcal{T}$.*

Proof. Suppose $l\sigma \succ r\sigma$ for a top-termination ordering \succeq . For an infinite rewrite sequence $t_0 \rightarrow_R t_1 \rightarrow_R t_2 \rightarrow_R \dots$, if $t_i \rightarrow_R t_{i+1}$ at top position Λ , then $t_i \succ t_{i+1}$, and if $t_i \rightarrow_R t_{i+1}$ at an inner position, then $t_i \succeq t_{i+1}$. Thus, an infinite number of *top* rewrites $t_i \rightarrow_R t_{i+1}$ would contradict the no infinite descending sequence property of top-termination orderings. \square

In the remainder of this section, we deal with *constructor* rewrite systems, by which we mean a set of rules R over a vocabulary $\mathcal{F} = \mathcal{F}_0 \cup \mathcal{F}_1$, such that no term in $\mathcal{G}(\mathcal{F}_0)$ is rewritable. In such systems, the left-hand sides must always contain a *non-constructor* symbol from \mathcal{F}_1 .

Theorem 6. *Let R be a left-linear constructor rewrite system. Suppose there is a quasi-ordering \succeq , the strict part of which is well-founded, with the following properties:*

(a) *If f is a constructor in \mathcal{F}_0 , then $s = f(\dots, s|_i, \dots) \succeq s|_i$ in the quasi-ordering (for any immediate subterm $s|_i$ of s); if f is a non-constructor in \mathcal{F}_1 , then $f(\dots, s|_i, \dots) \succ s|_i$.*

(b) *If $s \rightarrow_R t$ then $s \succeq t$ in the quasi-ordering; if $s \rightarrow_R t$ at its topmost redex but t has a non-constructor at its topmost position, then $s \succ t$.*

(c) *All finite ground non-constructor terms in $\mathcal{G} \setminus \mathcal{G}(\mathcal{F}_0)$ are reducible.*

Then R is top-terminating for ground terms. Moreover, all ω -normal forms of ground terms are constructor terms in $\mathcal{G}^\infty(\mathcal{F}_0)$.

Proof. Were the system not top-terminating, then there would be an infinite derivation $t_0 \rightarrow_R t_1 \rightarrow_R \dots$ with an infinite number of rewrites at the top, and, hence, with non-constructors at the top of each of its terms. Then, by (b), there would be an infinite descending sequence in \succ , contradicting its well-foundedness.

Next, we show that, for any ground term s in \mathcal{G} and for any integer $k \geq 0$, there is a term t in \mathcal{G} such that $s \rightarrow_R^* t$ and t has only constructors at depths less than k . This is by induction on s , using the well-founded ordering \succ , and for equivalent terms, by induction on k .

Let s be of the form $f(s|_1, \dots, s|_n)$. If f is a constructor in \mathcal{F}_0 , then $s \succeq s|_i$, and by the inductive hypotheses for each $s|_i$ there is a t_i such that $s|_i \rightarrow_R^* t_i$ and t_i has only constructors above level $k - 1$.

Suppose, then, that f is not a constructor. Then, by (a), the $s \succ s|_i$, so we can apply the induction hypothesis to each $s|_i$, and get arbitrarily many constructors at the top levels of those subterms. Let m be larger than the maximum depth of a left-hand side in R . Thus, $s|_i \rightarrow_R^* t_i$ for t_i with only constructors at the top m levels. By assumption (c) and left-linearity, any sufficiently deep term, headed by a non-constructor, must be rewritable at the topmost position Λ (since applicability of a rule cannot depend on anything below level m). So, $s \rightarrow_R^* f(t_1, \dots, t_n) \rightarrow_R t$, where the last rewrite is at the top.

Now, consider the consequences of (b). If t has a non-constructor at the top, then $s \succ t$, and we apply the induction hypothesis to t . If t has a constructor at the top, then $s \succeq t$, but t is strictly greater than its subterms, and we apply the induction hypothesis to them, as above. \square

For instance, let $\mathcal{F}_0 = \{a\}$ and $\mathcal{F}_1 = \{f, g\}$. The system:

$$\begin{aligned} f(a) &\rightarrow a \\ f(x) &\rightarrow g(f(g(x))) \end{aligned} \quad (17)$$

has normal forms a and g^ω . For the quasi-ordering, we use $s \succeq t$ if s has at least as many occurrences of the non-constructor f as does t .

We conclude this section with another method of showing that the limit of a derivation is a constructor term.

The *nesting level* $\ell(t)$ of non-constructors in a ground term t in \mathcal{G} is defined as follows:

- If t is a constructor constant, then $\ell(t) = 0$.
- If t is a non-constructor constant, then $\ell(t) = 1$.
- If t is $f(t_1, \dots, t_n)$ and f is a constructor, then $\ell(t) = \max\{\ell(t_1), \dots, \ell(t_n)\}$.
- If t is $f(t_1, \dots, t_n)$ and f is not a constructor, then $\ell(t) = 1 + \max\{\ell(t_1), \dots, \ell(t_n)\}$.

We say that a constructor rewrite system R *does not increase* the nesting level of non-constructors if, for all rules $l \rightarrow r$ in R and substitutions σ , $\ell(l\sigma) \geq \ell(r\sigma)$. This condition can be checked syntactically by noting the nesting of function symbols above each variable.

Theorem 7. *Suppose that R is a left-linear top-terminating constructor rewrite system that does not increase the nesting level of non-constructors. Suppose further that all finite ground non-constructor terms in $\mathcal{G} \setminus \mathcal{G}(\mathcal{F}_0)$ are reducible. Then all ω -normal forms of ground terms in \mathcal{G} are constructor terms in $\mathcal{G}^\infty(\mathcal{F}_0)$.*

Proof. Let $t_0 \rightarrow_R t_1 \rightarrow_R t_2 \rightarrow_R \dots$ be an infinite derivation. Since R does not increase the depth of nesting of non-constructors, $\ell(t_0) \geq \ell(t_1) \geq \ell(t_2) \geq \dots$. It can be shown that the non-constructors must get farther and farther apart in the terms t_0, t_1, t_2, \dots , and so the limit is a constructor term. Note that if all finite ground non-constructor terms in $\mathcal{G} \setminus \mathcal{G}(\mathcal{F}_0)$ are reducible by a left-linear system, then infinite ground non-constructor terms in $\mathcal{G}^\infty \setminus \mathcal{G}^\infty(\mathcal{F}_0)$ are too. The limit is irreducible, since R is top-terminating. Therefore, the limit cannot have a non-constructor symbol, since all such terms are reducible. \square

System (17) falls under this theorem. On the other hand, if we have a depth increasing rule, instead, as in:

$$\begin{aligned} f(a) &\rightarrow a \\ f(c(x)) &\rightarrow f(f(c(x))) \end{aligned} \quad (18)$$

then there are non-constructor ω -normal forms, such as f^ω .

The one-rule system

$$f(x) \rightarrow c(x, f(s(x))) \quad (19)$$

(for constructing a “stream” of “integers” $s^i(0)$), with one non-constructor f , is obviously top-terminating and has constructor normal forms by the above theorem. For example, the ω -normal form of $f(0)$ is $c(0, c(s(0), c(s(s(0)), \dots)))$.

The system

$$\begin{aligned} d(x) &\rightarrow c(a(x, x), d(s(x))) \\ a(0, y) &\rightarrow y \\ a(s(x), y) &\rightarrow s(a(x, y)) \end{aligned} \quad (20)$$

(for computing a stream of even “integers”), with $\mathcal{F}_0 = \{0, s, c\}$, also meets the requirements of the above theorem. On the other hand, the system

$$\begin{aligned} q(x) &\rightarrow c(m(x, x), q(s(x))) \\ a(0, y) &\rightarrow y \\ a(s(x), y) &\rightarrow s(a(x, y)) \\ m(0, y) &\rightarrow 0 \\ m(s(x), y) &\rightarrow a(y, m(x, y)) \end{aligned} \quad (21)$$

(for a stream of “squares”) does not fit into the scheme, since the last rule increases the nesting of non-constructors m and a .

The same idea as in this theorem applies whenever one can show that the nesting level is bounded, even if it temporarily increases.

It can be shown that, in general, top-termination is undecidable, even for left-linear ω -converging systems. Similarly, for top-terminating ω -confluent systems, it is undecidable whether the distance between a term and its normal form is less than a given $\epsilon > 0$. A programming language with infinite constructor normal forms is described in [20].

6. Uniqueness of normal forms

Recall that a relation \rightarrow is ω -confluent if $\omega\leftarrow \circ \rightarrow^\omega \subseteq \rightarrow^\omega \circ \omega\leftarrow$. Confluence is decidable for finitely terminating systems [17, 12], but not for non-finitely-terminating ones. For ω -converging systems, we can use the following variation:

Definition 12. A binary relation \rightarrow over a topological space is *semi- ω -confluent* if $\omega \leftarrow \circ \rightarrow^* \subseteq \rightarrow^\omega \circ \omega \leftarrow$.

A straightforward induction shows that, for ω -closed systems, semi- ω -confluence is equivalent to the more “local” condition: $\omega \leftarrow \circ \rightarrow \subseteq \rightarrow^\omega \circ \omega \leftarrow$.

The notions of ordinary confluence and semi- ω -confluence are independent: the system

$$\begin{array}{l} a \rightarrow f(a) \\ a \rightarrow c \\ f(c) \rightarrow c \end{array} \quad (22)$$

is confluent, but not semi- ω -confluent; the system

$$\begin{array}{l} a \rightarrow b \\ a \rightarrow c \\ b \rightarrow g(b) \\ c \rightarrow g(c) \end{array} \quad (23)$$

is semi- ω -confluent, but not confluent.

Obviously, ω -confluence implies semi- ω -confluence. The converse is not true in general; witness the non- ω -converging, non- ω -confluent, but semi- ω -confluent, rewrite system:

$$\begin{array}{l} a \rightarrow f(a) \\ a \rightarrow g(a) \\ f(x) \rightarrow x \\ g(x) \rightarrow x \end{array} \quad (24)$$

However:

Theorem 8. *An ω -converging semi- ω -confluent binary relation \rightarrow over a metric space has at most one ω -normal form.*

Proof. Let u_∞ and v_∞ be ω -normal-forms of t_0 . Consider an element t_1 at least “halfway” between t_0 and u_∞ . That is, $u_\infty \omega \leftarrow t_1 \overset{*}{\leftarrow} t_0 \rightarrow^\omega v_\infty$ and $d(s, u_\infty) < 1/2$ for every s in the chain from t_1 to u_∞ . By semi-confluence, we have $t_1 \rightarrow^\omega v_\infty$. By the same token, $u_\infty \omega \leftarrow t_1 \rightarrow^* t_2 \rightarrow^\omega v_\infty$, for some t_2 such that $d(s, v_\infty) < 1/4$ for s between t_2 and v_∞ , and $t_2 \rightarrow^\omega u_\infty$. Now, $u_\infty \omega \leftarrow t_3 \overset{*}{\leftarrow} t_2 \rightarrow^\omega v_\infty$, where t_3 is such that $d(s, u_\infty) < 1/8$ for s between t_3 and u_∞ , and so on. Consider the chain $t_0 \rightarrow^* t_1 \rightarrow^* t_2 \rightarrow^* \dots$. Since \rightarrow is ω -converging, this chain has a limit t_∞ , and, clearly, $d(u_\infty, t_\infty) = d(v_\infty, t_\infty) = 0$. Thus, $u_\infty = v_\infty$. \square


A similar argument shows that:

Theorem 9. *An ω -closed, ω -converging binary relation \rightarrow over a metric space is semi- ω -confluent if, and only if, it is ω -confluent.*

The hypothesis of ω -closure may be unnecessary, but is used in the following:

Proof. For the inobvious direction, suppose $*\leftarrow \circ \rightarrow^\omega \subseteq \omega\leftarrow \circ \rightarrow^\omega$. Let t_1 be such that $u_0 \xrightarrow{\omega} t_0 \xrightarrow{*} t_1 \xrightarrow{\omega} v_0$ and $d(s, v_0) < 1/2$ for every s between t_1 to v_0 . By semi-confluence, $u_0 \xrightarrow{\omega} u_1 \xrightarrow{\omega} t_1$. Similarly, $t_1 \xrightarrow{*} t_2 \xrightarrow{\omega} v_1 \xrightarrow{\omega} v_0$, for t_2 such that $d(s, u_1) < 1/4$ for s in $t_2 \xrightarrow{\omega} u_1$, and so on. Since \rightarrow is ω -converging, the ω -chain $t_0 \xrightarrow{*} t_1 \xrightarrow{*} t_2 \xrightarrow{*} \dots$ has a limit t_∞ . The ω -chains $u_0 \xrightarrow{\omega} u_1 \xrightarrow{\omega} u_2 \dots$ and $v_0 \xrightarrow{\omega} v_1 \xrightarrow{\omega} \dots$ are such that the distance between t_{2i} and u_i is no more than $1/4^i$ and between t_{2i+1} and v_i is $1/2^{2i+1}$. Thus, these chains have limits u_∞ and v_∞ , respectively. But $d(u_\infty, t_\infty) = d(v_\infty, t_\infty) = 0$; hence, $u_\infty = v_\infty$. Since \rightarrow is ω -closed, $u_0 \xrightarrow{\omega} u_\infty = v_\infty \xrightarrow{\omega} v_0$, and the relation is ω -confluent.

Pictorially:



□

Recall that orthogonal systems are left-linear and non-overlapping. Such systems are always confluent; they are not, however, ω -confluent.⁶

Theorem 10. *If R is an ω -converging orthogonal rewrite system, then R is semi- ω -confluent.*

⁶Despite claims to the contrary in our earlier work.

Non-left-linear System (4) is not semi- ω -confluent: the term $f(f(a,b), f(a,b))$ rewrites to normal form c in one step and to $f(f(a, g^\omega), f(g^\omega, b))$ in ω steps, but it takes $\omega + 1$ steps to join them. The requirement of ω -convergence may be superfluous.

The notion of *parallel rewriting* is crucial in reasoning about orthogonal systems. We write $\rightarrow_R^{\parallel}$ to denote one multiple rewrite step, applying the same rule and substitution at any number (from 0 to ω inclusive) of disjoint positions, and use $\rightarrow_{r\sigma}^{\parallel}$ to specify the rule (r) and substitution (σ) used.

Proof. Suppose $t_0 \rightarrow_R t_1 \rightarrow_R \cdots \rightarrow_R^\omega t_\infty$, and $t_0 \rightarrow_{r\sigma} t'_0$ by some rule r in R . It is well-known (see [16, 22]) that there exist terms t'_i such that, for all $i < \omega$,

$$\begin{array}{ccc} t_i & \rightarrow_R & t_{i+1} \\ r\sigma \downarrow^{\parallel} & & r\sigma \downarrow^{\parallel} \\ t'_i & \rightarrow_R^{\parallel} & t'_{i+1} \end{array}$$

If one goes far enough along the ω -converging derivations $t_0 \rightarrow_R t_1 \rightarrow_R \cdots \rightarrow_R^\omega t_\infty$ and $t'_0 \rightarrow_R^{\parallel} t'_1 \rightarrow_R^{\parallel} \cdots \rightarrow_R^\omega t'_\infty$, the top parts of the t_i and t'_i will have stabilized, and just like $t_i \rightarrow_R^{\parallel} t'_i$, we have $t_\infty \rightarrow_R^{\parallel} t'_\infty$.

It follows that $t' \rightarrow_R^\omega t'_\infty \leftarrow_R^{\parallel} \cdots \leftarrow_R^{\parallel} t_\infty$, whenever $t' \leftarrow_R^n t \rightarrow_R^\omega t_\infty$, for natural number n . Since R is orthogonal, it is always possible to interleave the (up to $\omega \times n$) steps between t_∞ and t'_∞ so that in fact $t_\infty \rightarrow_R^\omega t'_\infty$. \square

Recall that orthogonal systems need not be ω -converging (Example (2)), nor ω -closed (Example (7)). Indeed, the following non- ω -converging orthogonal system (discovered by [15]) does not have unique ω -normal forms:

$$\begin{array}{lcl} a & \rightarrow & f(g(a)) \\ f(x) & \rightarrow & x \\ g(x) & \rightarrow & x \end{array} \quad (25)$$

The term a has normal forms, f^ω and g^ω .

Nonetheless:

Corollary. *If R is an ω -converging orthogonal rewrite system, then R is uniquely- ω -normalizing.*

Proof. By Theorem 8, R has at most one ω -normal form. By the corollary to Theorem 3, R has at least one. \square

The semantic methods of Section 5 may be used to establish *ground* ω -confluence of a system R , that is, ω -confluence on the ground terms $\mathcal{G}(\mathcal{F})$, as well as *sufficient completeness*, that is, any ground term in $\mathcal{G}(\mathcal{F})$ has an ω -normal form that is a ground constructor term in $\mathcal{G}(\mathcal{F}_0)$. Using Theorem 6 or 7, one can determine that ω -normal forms of ground terms are always constructor terms in \mathcal{G}^∞ . Let $u[f(s|_1, \dots, s|_m)]_p$ and $u[g(t|_1, \dots, t|_n)]_p$ be two terms (s and t) with distinct constructors f and g in \mathcal{F}_0 embedded in the same constructor context $u[\cdot]_p$ in $\mathcal{T}(\mathcal{F}_0, \mathcal{X})$, where the $s|_i$ and $t|_i$ are arbitrary terms in \mathcal{T} . If R is such that no two such terms are provably equal (by replacement using the symmetric closure of \rightarrow_R), then confluence follows from the normal forms being constructor terms. Without this condition on provability, a simple system like

$$\begin{array}{l} a \rightarrow b \\ a \rightarrow c \end{array} \quad (26)$$

with constructors b and c , would mistakenly be deemed confluent.

Even for left-linear, top-terminating systems, ω -confluence is undecidable. Also, the joinability of critical pairs (as in [17]) is not a sufficient (nor necessary) condition for ω -confluence, as can be seen from the following ω -converging system:

$$\begin{array}{l} a \rightarrow b \\ g(x, a) \rightarrow f(g(x, a)) \\ g(x, b) \rightarrow c \\ f(c) \rightarrow c \end{array} \quad (27)$$

The term $g(x, a)$ rewrites in one step to either $f(g(x, a))$ or $g(x, b)$, both of which rewrite within a few steps to c ; yet, $g(x, a)$ has two distinct ω -normal forms, c and f^ω .

In [4], a condition on normal forms, called “ R -propriety”, is used to establish existence of unique solutions (as functions over \mathcal{T}^∞) to non-constructor operators, defined by a set of recursive rules.

7. Algebraic semantics

In this section, we consider algebraic aspects of infinitary theories—that is, their models—and their connection to operational aspects (namely, ω -rewriting). Since we are interested in infinite computations, it is natural to work with *continuous* models. (We refer the reader to [25, 26] for general references on the topic.) It is also natural to use a topological *completion* process. Alternative notions of completion have been studied in the algebraic framework, leading to different initial models, each with its own abstract properties (see, for instance, [1, 19, 27]).

Since our approach is unusual, we first illustrate the difficulty in assigning an appropriate algebraic semantics in the continuous case. In particular, the “natural” semantics identifies all terms, whenever there is identity function or axiom of idempotence. Consider the equations $q(0) = 0$ and $q(1) = 1$, thinking of q as “squaring”. We have

$$0 = q(0) = q(q(0)) = \dots$$

and

$$1 = q(1) = q(q(1)) = \dots$$

Were equations to carry over to the limits, we would find both 0 and 1 equal to q^ω , which is “inconsistent”. This is why we will work instead with oriented equations, interpreted in a model as *inequations* (but see [2]).

Definition 13. Given a vocabulary \mathcal{F} , a *continuous \mathcal{F} -algebra* consists of:

- A universe M , with a quasi-ordering \preceq , such that each nonempty, strictly increasing ω -chain admits a *least upper bound* (lub) in M .
- A continuous interpretation $f^M : M^n \rightarrow M$, for each f in \mathcal{F} (with arity n). Continuity means $\text{lub}_{i < \omega} f^M(t_i^1, \dots, t_i^n) = f^M(\text{lub}_{i < \omega} t_i^1, \dots, \text{lub}_{i < \omega} t_i^n)$ for any chains $(t_i^1)_{i < \omega}, \dots, (t_i^n)_{i < \omega}$.

Given a continuous \mathcal{F} -algebra M , any assignment $\sigma : \mathcal{X} \rightarrow M$ extends to a homomorphism $\sigma : \mathcal{T}(\mathcal{F}, \mathcal{X}) \rightarrow M$, as usual: if $t = f(t|_1, \dots, t|_n)$ for some f in \mathcal{F} , then $t\sigma = f^M(t|_1\sigma, \dots, t|_n\sigma)$; if $t = x$ for some x in \mathcal{X} , then $t\sigma = x\sigma$.

Definition 14. Given a rewrite system R over $\mathcal{T}(\mathcal{F}, \mathcal{X})$, an *R -model* is a continuous \mathcal{F} -algebra M that satisfies:

- (1) for any rule $l \rightarrow r$ in R , assignment $\sigma : \mathcal{X} \rightarrow M$, and context $c[\cdot]_p$ for \mathcal{T} , the inequality $c[l\sigma]_p \preceq c[r\sigma]_p$ (in M) holds;
- (2) for any assignment $\sigma : \mathcal{X} \rightarrow M$ and derivations $(u_i)_{i < \omega}$ and $(v_i)_{i < \omega}$ such that $\lim_{i \rightarrow \omega} u_i = \lim_{i \rightarrow \omega} v_i$, it is the case that $\text{lub}_{i < \omega} u_i\sigma = \text{lub}_{i < \omega} v_i\sigma$.

The class of all R -models is denoted \mathbf{Ord}_R .

Note that a model need not obey *equality* of left- and right-hand sides, as in the classical case, but, rather, *inequality*. The existence of least upper bounds in (2) comes from the fact that the two sequences (u_i) and (v_i) are increasing, by (1). The class \mathbf{Ord}_R is a non-empty category (cf. Theorem 11).

Definition 15. Let R be an ω -normalizing rewrite system. The *ordered model* $\widehat{\mathcal{G}}_R$ has a universe consisting of all finite ground terms \mathcal{G} and all their (possibly infinite) ω -normal forms, partially ordered by \rightarrow_R^ω . The model $\widehat{\mathcal{T}}_R$ consists of the terms \mathcal{T} and their ω -normal forms, ordered by \rightarrow_R^ω .

Clearly, $\widehat{\mathcal{G}}_R$ and $\widehat{\mathcal{T}}_R$ satisfy both conditions for R -models.

Theorem 11. *The ordered model $\widehat{\mathcal{G}}_R$ is initial in the class \mathbf{Ord}_R of R -models.*

Proof. Let $M \in \mathbf{Ord}_R$; we need a homomorphism $\phi : \widehat{\mathcal{G}}_R \rightarrow M$. For t finite, we must take $\phi[t] = t^M$. For an infinite ω -normal form t_∞ , we have $t_i \rightarrow^\omega t_\infty$, for some derivation $(t_i)_{i < \omega}$. We define: $\phi[t_\infty] = \phi[\text{lub}_{i < \omega} t_i] \stackrel{\text{def}}{=} \text{lub}_{i < \omega} (t_i^M)$. This ϕ is an \mathcal{F} -morphism, since $\phi[f(t_\infty|_1, \dots, t_\infty|_n)] = \text{lub} f^M((t_i|_1)^M, \dots, (t_i|_n)^M) = f^M(\text{lub}(t_i|_1)^M, \dots, \text{lub}(t_i|_n)^M) = f^M(\phi[t_\infty|_1], \dots, \phi[t_\infty|_n])$, where n is the arity of f . Lastly, ϕ is continuous, by construction. \square

We may now extend the definition of $t\sigma$ to infinite terms: For any model M in \mathbf{Ord}_R and infinite term t_∞ in $\widehat{\mathcal{G}}_R$ that is the limit of a derivation $(t_i)_{i < \omega}$, let $t_\infty\sigma \stackrel{\text{def}}{=} \text{lub}_{i < \omega} t_i\sigma$. Thanks to condition (2) of Definition 14, the lub does not depend on the derivation (t_i) leading to t_∞ .

Definition 16. Given terms t and t' in $\widehat{\mathcal{T}}_R$, and an R -model M in \mathbf{Ord}_R , we say that M *obeys* the inequality $t \preceq t'$, and write $M \models t \preceq t'$, if $t\sigma \preceq t'\sigma$ for every assignment $\sigma : \mathcal{X} \rightarrow M$. For a class of models \mathbf{M} , we write $\mathbf{M} \models t \preceq t'$ if $M \models t \preceq t'$, for every M in \mathbf{M} .

Theorem 12. *Let R be an ω -normalizing rewrite system and t, t' be terms in $\widehat{\mathcal{T}}_R$. Then*

$$\mathbf{Ord}_R \models t \preceq t' \text{ if, and only if, } t \rightarrow_R^\omega t'.$$

Proof. Suppose that $\mathbf{Ord}_R \models t \preceq t'$. In particular, $\widehat{\mathcal{T}}_R \models t \preceq t'$, which means that $t \rightarrow_R^\omega t'$. Conversely, suppose that $t \rightarrow_R^\omega t'$. Let $M \in \mathbf{Ord}_R$ and let $\sigma : \mathcal{X} \rightarrow M$. Without loss of generality, we may assume t to be finite, since otherwise it would have to be a normal form and $t = t'$. Hence, there exists a derivation $t = t_0 \rightarrow_R \dots \rightarrow_R t_i \rightarrow_R \dots \rightarrow_R^\omega t'$. By the nature of rewriting, we have $t\sigma \rightarrow_R^* t_i\sigma$ for all i , and, since M is an R -model, the $t_i\sigma$ form a chain in \preceq . Since M is a continuous algebra, $t\sigma \preceq \text{lub}_{i < \omega} t_i\sigma = t'\sigma$. In other words, $M \models t \preceq t'$. \square

Definition 17. The class \mathbf{Eq}_R of *equational* R -models is the subclass of the R -models \mathbf{Ord}_R for which $(c[l\sigma]_p)^M = (c[r\sigma]_p)^M$, for any rule $l \rightarrow r$ in R , substitution $\sigma : \mathcal{X} \rightarrow \mathcal{G}$, and context $c[\cdot]_p$ for \mathcal{G} .

Note that M is in \mathbf{Eq}_R only if M obeys $l \sim r$ (that is, $M \models l \preceq r$ and $M \models r \preceq l$) for all rules $l \rightarrow r$ in R .

Definition 18. Suppose R is a uniquely ω -normalizing rewrite system. The *normal-form* model NF_R consists of the ω -normal forms of the finite terms, ordered in a discrete fashion (that is the quasi-ordering \preceq for NF_R is equality of terms).

Theorem 13. *The normal-form model NF_R is initial in the class \mathbf{Eq}_R of equational R -models.*

Proof. It is clear that NF_R satisfies the conditions on R -models. Let $M \in \mathbf{Eq}_R$; we need a homomorphism $\psi : NF_R \rightarrow M$. Denote by ϕ_{NF} and ϕ_M the homomorphisms (as per Theorem 11) from $\widehat{\mathcal{G}}_R$ to NF_R and from $\widehat{\mathcal{G}}_R$ to M , respectively. We want $\phi_{NF} \circ \psi = \phi_M$. Thus, for any ω -normal form t_∞ that is the limit of a derivation $(t_i)_{i < \omega}$, we must define ψ as follows: $\psi[t_\infty] \stackrel{\text{def}}{=} \text{lub}_{i < \omega} t_i^M$. One can check that such a ψ is well-defined. Moreover, as previously, ψ is a continuous morphism. \square

Corollary. *Let R be a uniquely ω -normalizing rewrite system, and t, t' be terms in $\widehat{\mathcal{T}}_R$. Then*

$$\mathbf{Eq}_R \models t = t' \quad \text{if, and only if,} \quad t \xrightarrow{\omega}_R \circ \xleftarrow{\omega}_R t'.$$

The relationship between the models in \mathbf{Ord}_R , \mathbf{Eq}_R , and \mathbf{Alg}_R (the class of the finite, usual models of the equations represented by R), ordered by inclusion, forms a lattice, as illustrated in Figure 1. The quotient model $\mathcal{G}/=_R$ is initial in the class \mathbf{Alg}_R of algebraic models; the ordered model $\widehat{\mathcal{G}}_R$ is initial in the R -models \mathbf{Ord}_R ; the normal-form model NF_R is initial in the intersection \mathbf{Eq}_R of the two other classes. The trivial model $Triv$ lies at the top of the lattice.

For example, let the vocabulary \mathcal{F} consist of two constants, 0 and 1, and unary operator q . Let R be

$$\begin{array}{l} 0 \rightarrow q(0) \\ 1 \rightarrow q(1) \end{array} \quad (28)$$

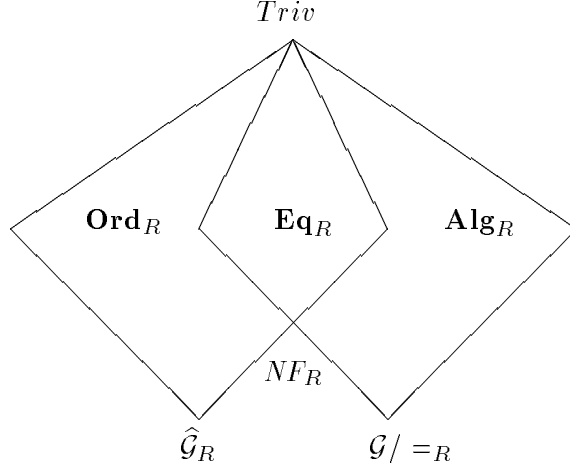
The ordered model $\widehat{\mathcal{G}}_R$ has as universe

$$\{0, q(0), \dots, q^n(0), \dots, 1, q(1), \dots, q^n(1), \dots, q^\omega\}$$

ordered by:

$$\begin{array}{l} 0 \preceq q(0) \preceq \dots \preceq q^n(0) \preceq \dots \preceq q^\omega \\ 1 \preceq q(1) \preceq \dots \preceq q^n(1) \preceq \dots \preceq q^\omega \end{array}$$

The universe NF_R collapses to the trivial model $\{q^\omega\}$. Notice that \mathbf{Ord}_R does not obey $0 \preceq 1$ or $1 \preceq 0$, and that \mathbf{Alg}_R does not obey $0 = 1$, whereas \mathbf{Eq}_R obeys all these.

Figure 1: Lattice of R -models.

Consider the following system R :

$$0 + x \rightarrow x \quad (29)$$

\mathbf{Eq}_R , that is, the set of models that obey $0 + x = x$, contains in particular NF_R , which is exactly the *finite* expressions that do not have subexpressions of the form $0 + x$. Consider now S :

$$x \rightarrow 0 + x \quad (30)$$

\mathbf{Eq}_S has only one trivial model: $NF_S = \{0 + (0 + \dots)\}$. It does *not* contain NF_R . We have the apparent paradox of two classes of models \mathbf{Eq}_R and \mathbf{Eq}_S , “defined” by the same equation $x = 0 + x$, which do not coincide. This stems from the fact that an R -model models an equation *oriented* in a particular way; the orientation influences the class of models \mathbf{Ord}_R , even when the orientation is “forgotten” in \mathbf{Eq}_R .

Definition 19. The class of the ω -reachable models is the subclass of the models M of \mathbf{Ord}_R such that the canonical morphism $\phi_M : \hat{\mathcal{G}}_R \rightarrow M$ is surjective.

The ω -reachable models form a non-empty, complete sublattice of \mathbf{Ord}_R (containing at least $\hat{\mathcal{G}}_R$).

Theorem 14. For any ω -reachable model M , there exists a continuous congruence \equiv_M on $\hat{\mathcal{G}}_R$ such that M is isomorphic to $\hat{\mathcal{G}}_R / \equiv_M$.

By “continuous”, we mean that $t_\infty \equiv_M t$ for any derivation $(t_i)_{i < \omega}$ with limit t_∞ such that $t_i \equiv_M t$ for all i . The proof is as in the finite case, with $t \equiv_M t'$ if, and only if, $\phi_M[t] = \phi_M[t']$.

8. Hierarchical systems

In this section, we consider *typed* systems (cf. [1, 13]), which are more general than the constructor systems of Section 5. A vocabulary (signature) is now a pair $\langle \mathcal{S}, \mathcal{F} \rangle$, where \mathcal{S} stands for a finite family of *sort* names and \mathcal{F} is a finite family of operators on \mathcal{S} . All the definitions given so far extend to the sorted case.

Definition 20. A *hierarchical specification* is a triple $\langle \mathcal{S}, \mathcal{F}, R \rangle$, where the sorts \mathcal{S} are split into a disjoint union $\mathcal{S}_0 \cup \mathcal{S}_1$, the operators \mathcal{F} into $\mathcal{F}_0 \cup \mathcal{F}_1$, and the rewrite system R into $R_0 \cup R_1$. The systems R_0 and R are ω -converging on $\mathcal{T}(\langle \mathcal{S}_0, \mathcal{F}_0 \rangle, \mathcal{X})$ and $\mathcal{T}(\langle \mathcal{S}, \mathcal{F} \rangle, \mathcal{X})$, respectively. Let \mathcal{G}_0 denote the set $\mathcal{G}(\langle \mathcal{S}_0, \mathcal{F}_0 \rangle)$ of ground constructor terms. The class \mathbf{HOrd}_R of *hierarchical models* is the class of models M of \mathbf{Ord}_R such that the restriction of M to the vocabulary $\langle \mathcal{S}_0, \mathcal{F}_0 \rangle$ is isomorphic to $\widehat{\mathcal{G}}_0$.

In the sequel, we suppose that the left-hand sides of the rules of R_1 always contain at least one symbol of \mathcal{F}_1 (otherwise, hierarchical consistency—defined below—could only be satisfied in trivial cases). In Section 5, R_0 was void.

Definition 21. A hierarchical specification is *sufficiently complete* if for every t in $\mathcal{G}(\langle \mathcal{S}_0, \mathcal{F}_0 \rangle)$, there exists a t' in $\widehat{\mathcal{G}}_{0R_0}$ such that $\mathbf{Ord}_R \models t \preceq t'$.

Definition 22. A hierarchical specification is *hierarchically consistent* if for any t in \mathcal{G}_0 and for any t' in $\widehat{\mathcal{G}}_{0R_0}$, $\mathbf{Ord}_R \models t \preceq t'$ if, and only if, $\mathbf{Ord}_{R_0} \models t \preceq t'$.

Note that an infinite term t' in $\widehat{\mathcal{G}}_{0R_0}$ is by definition a normal form for R_0 . It is also a normal form for R , due to the hypothesis about left-hand sides of inequalities of R_1 . Thus, sufficient completeness is equivalent to the existence of a t' such that $t \rightarrow_R^\omega t'$, and hierarchical consistency, to $t \rightarrow_R^\omega t'$ if, and only if, $t \rightarrow_{R_0}^\omega t'$.

These definitions are consonant with those of [27] for their notion of continuous specifications. Sufficient completeness means that any finite term t of an old sort, built with old and (possibly) new operators, is smaller than a (possibly infinite) term t' built with old operators only. Hierarchical completeness means that for two terms t and t' built with old operators only, $t \preceq t'$ holds in the new specification if, and only if, it holds in the old one.

Note also that the above definitions extend, as for finitary specifications, to the case where no new sort is introduced ($\mathcal{S}_1 = \emptyset$); operators of \mathcal{F}_0 are then called *constructors*, and operators of \mathcal{F}_1 are called *derived operators* (or simply “non-constructors”). As before, a *constructor term* is a term containing only constructors; a *non-constructor term* is a term containing at least one non-constructor. For instance, the specification:

constructors:	$a: \rightarrow \text{elem}$
	$c: \text{elem} \times \text{elem} \rightarrow \text{elem}$
derived operator:	$b: \rightarrow \text{elem}$
law:	$b \preceq c(a, b)$

is sufficiently complete in our sense. Note that in the classical, finitary framework [28], it would simply be rejected as being (finitely) incomplete.

Now, the main result is that, as in the finitary case, a hierarchically consistent and sufficiently complete specification satisfies its hierarchical constraints—in the following sense:

Theorem 15. *If $\langle \mathcal{S}, \mathcal{F}, R \rangle$ is sufficiently complete and hierarchically consistent, then \mathbf{HOrd}_R is a non-empty, complete sub-lattice of \mathbf{Ord}_R . Its initial model is $\widehat{\mathcal{G}}_R$.*

Proof. The proof is essentially as in the finite case. The main difference is in showing that $\widehat{\mathcal{G}}_R$ is actually in \mathbf{HOrd}_R , which we establish as follows:

The restricted model $\widehat{\mathcal{G}}_{0R}$ may be canonically embedded into $\widehat{\mathcal{G}}_{R_0}$. If $t_\infty \in \widehat{\mathcal{G}}_{0R}$ is finite, then t_∞ is in fact in \mathcal{G}_0 and therefore also in $\widehat{\mathcal{G}}_{R_0}$. If t_∞ is infinite, we may write $t_0 \rightarrow_R t_1 \rightarrow_R \dots \rightarrow_R^\omega t_\infty$, where the t_i are terms of $\mathcal{G}^\infty(\langle \mathcal{S}_0, \mathcal{F} \rangle)$. Using sufficient completeness, there exists $t'_\infty \in \widehat{\mathcal{G}}_{0R_0}$ such that $t_\infty \rightarrow_R^\omega t'_\infty$. Since t_∞ is a normal form for R , this means that $t_\infty = t'_\infty$, i.e. that t_∞ belongs to $\widehat{\mathcal{G}}_{0R_0}$.

Thus, $\widehat{\mathcal{G}}_{0R}$ may be seen as a subset of $\widehat{\mathcal{G}}_{0R_0}$. Now, hierarchical consistency shows that it is actually equal to the whole set, and that the orderings induced by R_0 and by R are identical. This finally establishes that $\widehat{\mathcal{G}}_R \in \mathbf{HOrd}_R$. \square

Define a quasi-ordering \preceq^{obs} on $\widehat{\mathcal{G}}_R$ as follows: when restricted to the sorts of \mathcal{S}_0 it is $\rightarrow_{R_0}^\omega$ (or equivalently \rightarrow_R^ω , because of hierarchical consistency), and $t \preceq^{obs} t'$ if, and only if, $c[t]_p \rightarrow_{R_0}^\omega c[t']_p$, for any terms t, t' of sort s in \mathcal{S}_1 in $\widehat{\mathcal{G}}_R$ and context $c[\cdot]_p$ with result in a sort of \mathcal{S}_0 .

Theorem 16. *The quotient $\widehat{\mathcal{G}}_R / \sim^{obs}$ ordered by \preceq^{obs} , is terminal among the ω -generated models of \mathbf{HOrd}_R .*

The proof is classical.

9. Discussion

There are circumstances (such as, stream-based programming) when the intended meaning of a function application is an infinite structure. As nondeterministic rewrite rules are convenient for specifying algebraic properties, we studied the semantics of rewritings that result in infinite terms.

We investigated the following properties of infinite derivations, among a few others:

- *ω -closure*: Any term that can be computed from a finite term in a transfinite number of rewrite steps is the limit of ω steps.
- *top-termination*: Any sequence of ω rewrite steps has at most a finite number of rewrites at the outermost position.
- *ω -normalization*: Every finite term has at least one normal form reachable in ω steps.
- *unique ω -normalization*: Every finite term has exactly one ω -normal form.

A number of counter-examples to conditions that seemed *prima facie* to assure ω -closure or ω -normalization were observed. On the other hand, some sufficient conditions for these properties were found, though we believe they leave room for improvement. In particular, for achieving such properties, left-linearity is important, top-termination is beneficial, while orthogonality is insufficient. With unique ω -normalization, one can assign semantics to rewriting under which function applications that result in the same infinite structure are equal.

There are many alternatives to the definitions we have given that are worth exploring: normal forms can be defined as (possibly infinite) terms to which no rewrite applies (see [15]); the initial terms of derivations are permitted to be infinite (finitely-representable infinite terms are dealt with in [2, 10]); strong convergence properties, like top-termination, can be required (see [10, 11, 15]); transfinite terms (with paths of length greater than ω) can be considered, in which subterms do not necessarily “disappear” after being pushed down ω times (cf. [2]); rewrite rules containing infinite terms (from \mathcal{T}^∞) themselves, and not just applying to infinite terms, can be allowed.

Acknowledgement

We thank John Givler, Jan-Willem Klop, Andrzej Tarlecki, Fer-Jan de Vries, and Ron Watro for their help debugging and improving this paper. All remaining errors are the authors’.

References

- [1] J. Goguen, J. Thatcher and E. Wagner, Initial algebra semantics and continuous algebras, *J. Assoc. Comput. Mach.* **24**(1) (1977), 69–95.
- [2] Y. Chen and M. J. O’Donnell, Infinite terms and infinite rewritings, in: *Extended Abstracts 2d Internat. Workshop on Conditional and Typed Rewriting Systems*, Montreal, Canada (Concordia Univ.), pp. 45–52 (June 1990).
- [3] B. Courcelle, Arbres infinis et systèmes d’équations, *R.A.I.R.O. Informatique Théorique* **13**(1) (1979), 31–48.
- [4] B. Courcelle, Infinite trees in normal form and recursive equations having a unique solution, *Math. Systems Theory* **13** (1979), 131–180.
- [5] N. Dershowitz, Termination of rewriting, *J. Symbolic Computation*, **3**(1&2) (1987), 69–115; **4**(3) (1987), 409–410.
- [6] N. Dershowitz and S. Kaplan, Rewrite, Rewrite, Rewrite, Rewrite, Rewrite, ... in: *Proc. 16th ACM Symp. on Principles of Programming Languages*, Austin, TX, pp. 250–259 (January 1989).
- [7] N. Dershowitz, S. Kaplan and D. A. Plaisted, Infinite normal forms (Preliminary version), in: *Proc. 16th EATCS Internat. Colloquium on Automata, Languages and Programming*, Stresa, Italy, Lecture Notes in Computer Science **372**, pp. 249–262 (Springer, Berlin, 1989); corrigendum distributed at colloquium.
- [8] N. Dershowitz and J.-P. Jouannaud, Rewrite systems, Chap. 6 in *Handbook of Theoretical Computer Science*, vol. B: Formal Models and Semantics, J. van Leeuwen, ed., in press (North-Holland, Amsterdam, 1990).
- [9] H. Ehrig, F. Parisi-Presicce, P. Boehm, C. Rieckhoff, C. Dimitrovici and M. Große-Rhode, Algebraic data type and process specifications based on projection spaces, in *Recent Trends in Data Type Specifications: Selected Papers 5th Workshop on Specification of Abstract Data Types*, Gullane, Scotland, Lecture Notes in Computer Science **332**, pp. 23–43 (Springer, Berlin, 1988).
- [10] W. M. Farmer and R. J. Watro, Redex capturing in term graph rewriting, Technical report M89–36, The Mitre Corporation, Bedford, MA, 1989.
- [11] John S. Givler, Continuous rewriting systems, unpublished note, IBM T. J. Watson Research Center, 1989.
- [12] G. Huet, Confluent reductions: abstract properties and applications to term rewriting systems, *J. Assoc. Comput. Mach.* **27**(4) (1980), 797–821.
- [13] G. Huet and D. C. Oppen, Equations and rewrite rules: A survey, in *Formal Language Theory: Perspectives and Open Problems*, R. Book, ed., pp. 349–405 (Academic Press, New York, 1980).
- [14] S. Kaplan, Conditional rewrite rules, *Theoretical Computer Science* **33** (1984), 175–193.
- [15] J. R. Kennaway, J. W. Klop, M. R. Sleep and F. J. de Vries, Transfinite reductions in orthogonal term rewriting systems, Technical report, Centrum voor Wiskunde en Informatica, Amsterdam, to appear, 1990.
- [16] J. W. Klop, Term rewriting systems: A tutorial, *Bulletin European Assoc. for Theoret. Comput. Sci.* **32**, pp. 143–183 (June 1987).
- [17] D. E. Knuth and P. B. Bendix, Simple word problems in universal algebras, in *Computational Problems in Abstract Algebra*, J. Leech, ed., pp. 263–297 (Pergamon Press, Oxford, 1970).

- [18] Z. Manna and S. Ness, On the termination of Markov algorithms, in: *Proc. 3d Hawaii Internat. Conf. on System Science*, Honolulu, HI, pp. 789–792 (1970).
- [19] B. Möller, On the algebraic specification of infinite objects—Ordered and continuous models of algebraic types, *Acta Informatica* **22**(5) (1984), 537–578.
- [19a] T. Naoi and Y. Inagaki, Algebraic semantics and complexity of term rewriting systems, in: *Proc. 3d Internat. Conf. on Rewriting Techniques and Applications*, Lecture Notes in Computer Science **355**, pp. 311–325 (Springer, Berlin, 1989).
- [20] S. Narain, LOG(F): An optimal combination of logic programming, rewriting and lazy evaluation, Ph.D. thesis, Dept. of Computer Science, Univ. of California, Los Angeles, CA, 1988.
- [21] M. Nivat, On the interpretation of recursive polyadic program schemes, *Symp. Math.* **15**, pp. 255–281 (Rome, 1975).
- [22] M. J. O’Donnell, *Computing in Systems Described by Equations*, Lecture Notes in Computer Science **58** (Springer, Berlin, 1977).
- [23] M. J. O’Donnell, *Equational Logic as a Programming Language*, MIT Press (Cambridge, MA, 1985).
- [24] J.-C. Raoult and J. Vuillemin, Operational and semantic equivalence between recursive programs, *J. Assoc. Comput. Mach.* **27**(4) (1980), 772–796.
- [25] D. Scott, Data types as lattices, *SIAM J. Computing* **5** (1976), 522–587.
- [26] J. Stoy, *Denotational semantics: The Scott-Strachey approach to programming languages*, MIT Press (Cambridge, MA, 1977).
- [27] A. Tarlecki and M. Wirsing, Continuous abstract data types, *Fundamenta Informaticae* **9** (1986), 95–126.
- [28] M. Wirsing, Algebraic specification, in *Handbook of Theoretical Computer Science*, vol. B: Formal Models and Semantics, J. van Leeuwen, ed., in press (North-Holland, Amsterdam, 1990).