Perpetual Reductions and Strong Normalization in Orthogonal Term Rewriting Systems

Zurab Khasidashvili

CWI

P.O. Box 4079, 1009 AB Amsterdam, The Netherlands

Abstract

We design a strategy that for any given term t in an Orthogonal Term Rewriting System (OTRS) constructs a longest reduction starting from t if t is strongly normalizable, and constructs an infinite reduction otherwise. We define some classes of OTRSs for which the strategy is easily computable. We develop a method for finding the least upper bound of lengths of reductions starting from a strongly normalizable term. We give also some applications of our results.

AMS Subject Classification (1991): 68Q42.
CR Subject Classification (1991): F4.2, F4.1.
Keywords & Phrases: orthogonal term rewriting systems, perpetual reductions, strong normalization, reduction strategies.
Note: Part of this work was completed during an enjoyable visit of the author at CWI in the summer of 1993.

1. Introduction

It is shown in O'Donnell [9] that the innermost strategy is perpetual for orthogonal term rewriting systems (OTRSs). That is, contraction of innermost redexes gives an infinite reduction of a given term whenever such a reduction exists. In fact, a strategy that only chooses redexes that do not erase any other redex is perpetual. Moreover, one can even reduce redexes whose erased arguments are strongly normalizable (Klop [7]). For the lambda-calculus, a more subtle perpetual strategy was invented in Barendregt et al. [1]. However, none of these strategies work for all Orthogonal Combinatory Reduction Systems (OCRSs), i.e., OTRSs with bound variables and a substitution mechanism [6].

We design a perpetual strategy for orthogonal term rewriting systems that works also for all orthogonal combinatory reduction systems. Perpetual reductions are interesting because termination of a perpetual reduction starting from a term t implies strong normalization of t (i.e., termination of all reductions starting from t). Our aim is not only to construct an infinite reduction of any given term t whenever it exist, but also to construct a longest reduction if all reductions starting from t are finite. Thus we will be able to characterize the complexity of computations of terms. The idea is that in order to construct a perpetual reduction one should try to avoid erasure of (infinite) redexes. On the other hand, in order to construct a longest possible reduction, one should delay contraction of a redex until it will no longer be possible to duplicate it by reducing an outer redex. The two conditions agree if in each term s one contracts a *limit* redex, which is defined as follows: choose in s an *unabsorbed* redex u_1, i.e., a redex whose descendants never appear inside arguments of other redexes; choose an erased argument s_1 of u_1 that is not in normal form; choose in s_1 an unabsorbed redex u_2, and so on, as long as possible. The last chosen redex is a limit redex of s.

An unabsorbed redex exists in any term not in normal form, but there is no general algorithm to find one. So we define some classes of OTRSs, such as persistent, inside-creating, non-absorbing, non-left-absorbing, and non-right-absorbing systems, for which the unabsorbed redexes are easy to find. For example, in non-left-absorbing systems, no subterm can be absorbed to the left of the contracted redex, so the leftmost-outermost redexes are unabsorbed. (In particular, the λ-calculus and the combinatory logic are non-left-absorbing.) Unabsorbed redexes can easily be found also in the wide class of strongly sequential OTRSs [3].

We develop a method for proving that the reductions constructed according to our perpetual strategy are indeed the longest, and for finding their lengths. Our method is similar to the Nederpelt's method [8] invented to reduce proofs of strong normalization to proofs of weak normalization (i.e., existence of a normal form). For any OTRS R, we define the corresponding non-erasing OTRS R_μ, called the μ-extension of R. We add fresh function symbols μ^n of arity n ($n = 0, 1, \ldots$) in the alphabet of R. For any R-rule $r : t \to s$, we keep the erased variables of t in the right-hand side of each corresponding R_μ-rule $r_\mu : t' \to s'$ as "μ-erased" arguments of s'. Since this transformation affects the structure of redex-creation in R, we have to introduce infinitely many R_μ-rules for each R-rule. This helps to have a natural translation of R-reductions into R_μ-reductions and vice-versa. Finally, we keep also all μ-symbols of t' as μ-erased symbols in s', since they can be used as "counters" of steps in longest normalizing reductions. We then show that the least upper bound of lengths of R-reductions starting from a term o coincides with the number of μ-occurrences in the R_μ-normal form of o. To find this number, sometimes it is not necessary to do actual transformation of t. We show this for the case of persistent TRSs.

Another consequence is that a term t is strongly normalizable in R iff it is weakly normalizable in R_μ; this result holds also on the level of OTRSs: an OTRS R is strongly normalizing iff its μ-extension R_μ is weakly normalizing [6]. Therefore, for any class of OTRSs that is closed under μ-extension, i.e., contains the μ-extension of each of its elements, one can prove undecidability of weak normalization if undecidability of strong normalization is known, and prove decidability of strong normalization if decidability of weak normalization is known. For example, all the above classes of OTRSs are closed under μ-extension. We describe some applications of our results in section 3. The main results are obtained in section 2.

2. Perpetual strategies in OTRSs

We recall some basic notions of TRS theory; one can find comprehensive introductions to the subject in [2] and [7]. A *TRS* is a pair (Σ, R), where the alphabet Σ consists of variables and function symbols and R is a set of rewrite rules r of the form $t \to s$. The left-hand side t is any term different from a variable, and the term s may only contain variables that occur in t. An *r-redex* u is obtained from t by substituting arbitrary terms for the variables in t, and the corresponding instance of s is the *contractum* of u. *Arguments* of u are subterms of u that correspond to variables of t, and the rest is the *pattern* of u. Subterms of u rooted at the pattern are called the *pattern-subterms* of u. The arguments, pattern, and pattern-subterms are defined analogously in the contractum of u. A TRS is *orthogonal* if it is left-linear and non-ambiguous, i.e., patterns of redexes can never overlap in a term.

A one step *reduction* in which a redex u in a term o is contracted is written $o \xrightarrow{u} e$ or $o \to e$. We write $P : o \twoheadrightarrow e$ if P is a *reduction* of o to e comprising 0 or more steps. A term t is called *weakly* (resp. *strongly*) *normalizable* if t has a normal form (resp. if any reduction starting from t is terminating). An OTRS R is *weakly* (resp. *strongly*) *normalizing* if any term in R is weakly (resp. strongly) normalizable. We use t, s, e, o for terms, u, v, w for redexes, and P, Q for reductions. $|P|$ denotes the length of P. We write $s \subseteq t$ if s is a subterm of t.

For a given OTRS R we now define its "μ-extension" R_μ: for each R-rule $t \to s$, we have a set of R_μ-rules of the form $t' \to \mu^l(\mu^0, \ldots, \mu^0, x_{i_1}, \ldots, x_{i_k}, s)$, where μ^n is a fresh n-ary function symbol; t is obtained from t' by removing all except the last arguments of μ-symbols occurring in t' (we write $t = [t']_\mu$); and $x_{i_1}, \ldots, x_{i_k}$ are all variables of t that do not occur in s. If a term e has a normal form in R_μ, then all R_μ-reductions of e are finite, since their lengths can not exceed the number $\|o\|_\mu$ of μ-occurrences in R_μ-normal form o of e. (Indeed, for any R_μ-reduction $P : e \twoheadrightarrow e'$, we have $e' \twoheadrightarrow o$; hence $|P| \leq \|e'\|_\mu \leq \|o\|_\mu$.) For any R-reduction $Q : t_1 \to t_2 \to \ldots \to t_n$ we construct a corresponding R_μ-reduction $Q_\mu : s_1 = t_1 \to s_2 \to \ldots \to s_n$ such that $[s_i]_\mu = t_i$. So in order to prove that a term t in R is strongly normalizable it is enough to prove that t has a normal form in R_μ. This is the idea of Nederpelt's method. Now if s_n is an R_μ-normal form of s_1, then $\|s_n\|_\mu$ is an upper bound of lengths of R-reductions starting from t_1. Thus the length of Q is maximal if s_n is the R_μ-normal form of s_1 whenever t_n is the normal form of t_1, and $|Q_\mu| = \|s_n\|_\mu$, i.e., each step of Q_μ increases the number of μ-occurrences exactly by 1. This is achieved by contracting the limit redexes only. Indeed, in this case the old μ-occurrences do not duplicate, and the only new μ-symbol created in each step is the head symbol of the contractum.

Definition 2.1 The *μ-extension* (Σ_μ, R_μ) of an OTRS (Σ, R) is defined as follows:

1. $\Sigma_\mu = \Sigma \cup \{\mu^n \mid n = 0, 1, \ldots\}$, where μ^n is a fresh n-ary function symbol. For any subterm $s = \mu^{n+1}(t_1, \ldots, t_n, t_0)$ of a term t over Σ_μ, the arguments $t_1, \ldots, t_n$, as well as subterms and symbols in $t_1, \ldots, t_n$ and the head-symbol μ itself, are called *μ-erased* or more precisely *μ'-erased*, where μ' is the occurrence of the head symbol of s in t. The argument t_0 is called *μ'-main*. Symbols and subterms in t that are not μ-erased are called *μ-main*. We denote by $[t]_\mu$ the term obtained from t by removing all μ-erased symbols.

2. R_μ is the set of all rules of the form $r_\mu : t' \to s'$ such that

(a) there is a rule $r : t \to s$ in R such that $[t']_\mu = t$;

(b) the term t' is linear (i.e., no variable appears twice or more in t');

(c) the head symbol of t' is not a μ-symbol, i.e., it coincides with the head symbol of t;

(d) the μ-erased arguments of each occurrence μ' of a μ-symbol in t' are variables, and the μ'-main argument is not a variable (i.e., it contains a function symbol from Σ or a μ-symbol);

(e) if $x_1, \ldots, x_n$ are all μ-main variables of t' (from left to right), $y_1, \ldots, y_m$ are all μ-erased variables of t', and $x_{i_1}, \ldots, x_{i_p}$ are all variables among $x_1, \ldots, x_n$ that do not occur in s, then

$$s' = \mu^l(\overbrace{\mu^0, \ldots, \mu^0}^{k}, y_1, \ldots, y_m, x_{i_1}, \ldots, x_{i_p}, s),$$

where k is the number of occurrences of μ-symbols in t' and $l = k + m + p + 1$. For any r_μ-redex $u = t'\theta$, we call arguments that correspond to $x_{i_1}, \ldots, x_{i_l}$ *quasi-erased* arguments of u, and call the arguments that correspond to other variables from $x_1, \ldots, x_n$ *quasi-main*. R_μ and R are called μ-corresponding OTRSs, and r_μ and r are called corresponding rules in R_μ and R.

Example 2.1 Let $R = \{r : f(a, x) \to b\}$. Then R_μ-rules have the form

$$\begin{aligned} f(\mu^k(x_1, \ldots, x_{k-1}, \mu^l(y_1, \ldots, y_{l-1}, \ldots \mu^m(z_1, \ldots, z_{m-1}, a) \ldots)), x) \to \\ \mu(\mu^0, \ldots, \mu^0, x_1, \ldots, x_{k-1}, y_1 \ldots, y_{l-1}, \ldots, z_1 \ldots, z_{m-1}, x, b) \end{aligned}$$

For example, $r_\mu : f(\mu^2(y, \mu^2(z, a), x) \to \mu^6(\mu^0, \mu^0, y, z, x, b))$ is an R_μ-rule. For any r_μ-redex $t = f(\mu^2(o, \mu^2(s, a)), e)$, $[t]_\mu = f(a, e)$ is an r-redex, $t' = \mu(\mu^0, \mu^0, o, s, e, b)$ is the contractum of t, and $[t']_\mu = b$ is the contractum of $f(a, e)$.

Lemma 2.1 If R is an orthogonal TRS, then so is R_μ.

Proof Any overlap of patterns of two R_μ-redexes in a term t over Σ_μ causes an overlap of patterns of corresponding R-redexes in $[t]_\mu$.

Lemma 2.2 Let t be a term over Σ_μ the head-symbol of which is not a μ-symbol, and let $[t]_\mu = s$. Then t is an r_μ-redex iff s is an r-redex, where r_μ and r are corresponding rules in R_μ and R, respectively. Moreover, if t' is the contractum of t in R_μ and s' is the contractum of s in R, then $[t']_\mu = s'$.

Proof From Definition 2.1 (see also Example 2.1).

Corollary 2.1 Let R be an OTRS and $s_0 \stackrel{u_0}{\to} s_1 \stackrel{u_1}{\to} \ldots$ be a reduction in R. Then, for any term t_0 in R_μ such that $[t_0]_\mu = s_0$, there is a reduction $t_0 \stackrel{v_0}{\to} t_1 \stackrel{v_1}{\to} \ldots$ in R_μ such that $[t_i]_\mu = s_i$ and u_i and v_i are corresponding subterms in s_i and t_i $(i = 0, 1, \ldots)$.

Notation $\|t\|_\mu$ denotes the number of occurrences of μ-symbols in t.

Lemma 2.3 Let t be a term in an OTRS R. If t is weakly normalizable in R_μ, then t is strongly normalizable in R_μ and R.

Proof Let s be an R_μ-normal form of t and $t \to t_1 \to \ldots$ be an R_μ-reduction. By Lemma 2.1 and the Church-Rosser theorem, $t_i \twoheadrightarrow s$ for all $i = 1, 2, \ldots$ It is easy to see that $i \le \|t_i\|_\mu \le \|s\|_\mu$. So t is strongly normalizable in R_μ. Hence, by Corollary 2.1, t is strongly normalizable in R.

Definition 2.2 Let $t \stackrel{u}{\to} s$ and let e be the contractum of u in s. For each argument o of u there are 0 or more arguments of e. We call them (u-)*descendants* of o. Correspondingly, subterms of o have 0 or more *descendants*. An argument of u is called (u)-*erased* if it does not have a descendant, and is called (u)-*main* otherwise. By definition, the *descendant* of each pattern-subterm of u is e. Descendants of all redexes of t except u are also called *residuals*. By definition, u does not have residuals in s. A redex of s is said to be *created* by contracting u or to be an (u)-*new* redex if it is not a residual of a redex of t'. It is clear what is to be meant under *descendants* of subterms that are not in u. The notion of *descendant* and *residual* extend naturally to arbitrary reductions. The *ancestor* relation is the inverse of the descendant relation.

Definition 2.3 We call a redex u *complete* (resp. ∞-*complete*) if erased arguments of u are in normal form (resp. are strongly normalizable). A reduction is *complete* (resp. ∞-*complete*) if it only contracts complete (resp. ∞-*complete*) redexes.

Lemma 2.4 Let $P : t_0 \stackrel{u_0}{\to} t_1 \stackrel{u_1}{\to} \ldots \to t_n$ be a complete reduction in an OTRS (Σ, R), and let $P_\mu : s_0 \stackrel{v_0}{\to} s_1 \stackrel{v_1}{\to} \ldots \to s_n$ be a corresponding R_μ-reduction such that all R_μ-redexes in s_0 are μ-main. Then

(1) for each k $(0 \le k \le n)$, we have $(b)_k$: any R_μ-redex in s_k is μ-main.

(2) if P is normalizing, then so is P_μ.

Proof (1) By induction on k. $(b)_0$ is obvious. Suppose that $(b)_k$ holds and let us show $(b)_{k+1}$. Let $u_k = C[o_1, \ldots, o_m]$ and $v_k = C'[e_1, \ldots, e_m, e'_1, \ldots, e'_p]$, where $e_1, \ldots, e_m$ are μ-main arguments of v_k, which correspond to arguments $o_1, \ldots, o_m$ of u_k, respectively, and $e'_1, \ldots, e'_p$ are μ-erased arguments of v_k. Then the contractum of v_k in R_μ has the form $o' = \mu(\mu^0, \ldots, \mu^0, e'_1, \ldots, e'_p, e_{i_1}, \ldots, e_{i_l}, o)$, where o is the contractum of $C[e_1, \ldots, e_m]$ in (Σ_μ, R) and, for each j, o_{i_j} is u_k-erased. Since v_k and u_k are corresponding subterms of s_k and t_k, and e_i and o_i are corresponding arguments, we have that (α): $[e_i]_\mu = o_i$, $i = 1, \ldots, m$. Since $o_{i_1}, \ldots, o_{i_l}$ are u_k-erased and u_k is complete, we have that (β): $o_{i_1}, \ldots, o_{i_l}$ are R-normal forms. It follows from $(b)_k$ that in $e'_1, \ldots, e'_p$ there are no R_μ-redexes and that R_μ-redexes in s_{k+1}, that are not in o' or are in o, are μ-main. It follows from $(b)_k$, (α), (β), and Lemma 2.2 that $e_{i_1}, \ldots, e_{i_l}$ are R_μ-normal forms. Thus $(b)_{k+1}$ holds and (1) is proved.

(2) By Lemma 2.2 and $(b)_n$.

Recall that a (*sequential*) *strategy* selects a redex to be contracted in any given term. A *complete* (resp. ∞-*complete*) *strategy* contracts a complete (resp. an ∞-complete) redex in each step. A strategy is *perpetual* if it constructs an infinite reduction of any given term whenever such a reduction exists.

Theorem 2.1 A complete strategy is perpetual in orthogonal TRSs.

Proof It is enough to show that if t has a normalizing complete reduction $P : t \twoheadrightarrow t'$, then t is strongly normalizable. Indeed, by Lemma 2.4, the corresponding R_μ-reduction of P is also normalizing. Hence, by Lemma 2.3, t is strongly normalizable in R.

Definition 2.4 (1) A TRS is called *non-erasing* if left- and right-hand sides of each rule in it contain occurrences of the same variables.

(2) A TRS is called *weakly innermost normalizing* [9] if each term has a normal form reachable by an innermost reduction.

Corollary 2.2 (Church) Let R be a non-erasing OTRS. Then R is weakly normalizing iff it is strongly normalizing.

Corollary 2.3 (O'Donnell [9]) Let R be an OTRS. Then R is strongly normalizing iff it is weakly innermost normalizing.

Theorem 2.2 (Klop [7]) An ∞-complete strategy is perpetual in orthogonal TRSs.

Proof It is enough to prove that if t_0 has a normalizing ∞-complete reduction $P : t_0 \xrightarrow{u_0} t_1 \rightarrow \ldots \rightarrow t_n$, then t_0 is strongly normalizable in R. Since u_i is ∞-complete, there is a complete normalizing reduction $Q : t_0 \twoheadrightarrow t_1 \twoheadrightarrow \ldots \twoheadrightarrow t_n$. Now by Theorem 2.1, t_0 is strongly normalizable.

The following propositions are obtained in Klop [6]. The proof of Proposition 2.2 is, however, not correct, since it establishes weak normalization of R_μ only for terms of R, not for all terms of R_μ.

Proposition 2.1 (Klop [6]) A term t in an OTRS R is strongly normalizable iff t is weakly normalizable in R_μ.

Proof ($\Rightarrow$) From Lemma 2.4. ($\Leftarrow$) From Lemma 2.3.

Proposition 2.2 (Klop [6]) An OTRS R is strongly normalizing iff R_μ is weakly normalizing.

Proof ($\Rightarrow$) Let t' be a term in R_μ. We prove that t' is weakly normalizable by induction on the length of t'. By the induction assumption, all μ-erased subterms of t' are weakly normalizable in R_μ. Let t^* be obtained from t' by their reduction to normal forms. By Lemma 2.4, t^* is weakly normalizable in R_μ, since $[t^*]_\mu$ is strongly normalizable in R. Hence, t' is weakly normalizable. ($\Leftarrow$) From Lemma 2.3.

Definition 2.5 ([4]) We call a subterm s of a term t *unabsorbed in a reduction* $P : t \twoheadrightarrow e$ if the descendants of s do not appear inside redex-arguments of terms in P, and call s *absorbed in* P otherwise. We call s *unabsorbed in* t if it is unabsorbed in any reduction starting from t, and *absorbed in* t otherwise.

In [3], Huet and Lévy introduced the notion of *external redex* of a term and proved that each term not in normal form possesses an external redex. It is easy to show that a redex $u \subseteq t$ is unabsorbed iff u is external in t. Thus we have the following lemma; a short direct proof of it can be found in [4].

Lemma 2.5 In any term t not in normal form there is an unabsorbed redex.

Definition 2.6 Let u_l be a redex in a term t defined as follows: choose an unabsorbed redex u_1 in t; choose an erased argument s_1 of u_1 that is not in normal form (if any); choose in s_1 an unabsorbed redex u_2, and so on, as long as possible. Let $u_1, s_1, u_2, \ldots, u_l$ be such a sequence. Then we call u_l a *limit redex* and call $u_1, s_1, u_2, \ldots, u_l$ a *limit sequence* of t.

It follows from Lemma 2.5 that in any term not in normal form there is a limit redex. We call a reduction *limit* if each contracted redex in it is limit, and call a strategy *limit* if in any term not in normal form it contracts a limit redex.

Lemma 2.6 Let u be a limit redex in t and $P : t \twoheadrightarrow e$. Then there is no new redex in e that contains a descendant of u in its argument.

Proof Let $u_1, s_1, u_2, \ldots, u_l$ be the limit sequence of t with $u_l = u$. We prove by induction on $|P|$ that (a): descendants of redexes $u_1, \ldots, u_l$ do not appear inside arguments of new redexes. If $|P| = 0$, then (a) is obvious. So let $P : t \twoheadrightarrow e' \xrightarrow{v} e$, let o be a descendant of u in e, and let o' be its ancestor in e'. It follows from the induction assumption that each redex $u_i\,(i = 1, \ldots, l-1)$ has exactly one residual u'_i in e' (because contraction of a residual of any of the redexes $u_1, \ldots, u_{l-1}$ erases the descendant of u), there is no new redex in e' that contains o' in its argument, and o is the only descendant of u. Thus if there is a new redex w in e that contains the residual u''_i of some u_i in its argument, then it must be created by v. If $v \not\subseteq u'_1$, then w contains u''_i in its argument iff it contains the residual of u'_1 in its argument, but this is impossible since u_1 is unabsorbed. Thus $v \subseteq u'_1$. Let k be the maximal number such that v is in u'_k and let s'_k be the descendant of s_k in e'. Then v is in s'_k and contains u'_{k+1}. Let $Q : s_k \twoheadrightarrow s''_k$ consist of steps of P that are made in descendants of s_k. Then the residual of u_{k+1} is in an argument of the new redex $w \subseteq s''_k$. But this is impossible since u_{k+1} is unabsorbed in s_k. Thus (a) is valid and the lemma is proved.

Lemma 2.7 Let (Σ, R) be an OTRS, $P : t_0 \xrightarrow{u_0} t_1 \xrightarrow{u_1} \ldots \to t_n$ be a limit reduction in R and $P_\mu : s_0 = t_0 \xrightarrow{v_0} s_1 \xrightarrow{v_1} \ldots \to s_n$ be its corresponding reduction in R_μ.

(1) For each k $(0 \leq k \leq n)$ the following holds:

$(a)_k$: $\|s_k\|_\mu = k$;

$(b)_k$: each redex $v'_k \subseteq s_k$ is μ-main in s_k;

$(c)_k$: in quasi-main arguments of any redex v''_k in s_k there are no μ-symbols.

(2) If P is normalizing, then so is P_μ.

Proof (1) $(a)_0 - (c)_0$ are obvious. Suppose that $(a)_k - (c)_k$ hold and let us show $(a)_{k+1} - (c)_{k+1}$. Let $u_k = C[o_1, \ldots, o_q]$ and $v_k = C'[e_1, \ldots, e_q, e'_1, \ldots, e'_m]$, where $e_1, \ldots, e_q$ are μ-main arguments of v_k (which correspond to arguments $o_1, \ldots, o_q$ of u_k respectively) and $e'_1, \ldots, e'_m$ are μ-erased arguments of v_k. Since u_k and v_k are corresponding redexes in t_k and s_k, we have $[v_k]_\mu = u_k$ and hence (α): $[e_i]_\mu = o_i$ for all $i = 1, \ldots, q$. Let $o_{i_1}, \ldots, o_{i_l}$ be u_k-erased arguments and $o_{j_1}, \ldots, o_{j_p}$ be u_k-main arguments. Then contractum of v_k in R_μ has the following form: $o' = \mu(\mu^0, \ldots, \mu^0, e'_1, \ldots, e'_m, e_{i_1}, \ldots, e_{i_l}, o)$, where o is the contractum of $C[e_1, \ldots, e_n]$ in (Σ_μ, R). Since u_k is limit, (β): $o_{i_1}, \ldots, o_{i_l}$ are in R-normal form. By $(c)_k$, (γ): there are no occurrences of μ-symbols in $e_{j_1}, \ldots, e_{j_p}, o$. (Hence o coincides with the contractum of u_k.) It follows from $(\alpha), (\beta)$, $(b)_k$, and Lemma 2.2 that (δ): $e_{i_1}, \ldots, e_{i_l}$ are in R_μ-normal form.

By (γ), $\|o'\|_\mu = \|v_k\|_\mu + 1$. Hence $\|s_{k+1}\|_\mu = \|s_k\|_\mu + 1 = k + 1$, i.e., $(\alpha)_{k+1}$ holds.

If $v'_{k+1} \not\subseteq o'$, then $(b)_k$ implies that v'_{k+1} is μ-main. If $v'_{k+1} \subseteq o'$, then, by $(b)_k$, $v'_{k+1} \not\subseteq e'_1, \ldots, e'_m$ (since ancestors of $e'_1, \ldots, e'_m$ are μ-erased arguments of v_k) and, by (δ), $v'_{k+1} \not\subseteq e_{i_1}, \ldots, e_{l_1}$. Hence $v'_{k+1} \subseteq o$ and v'_{k+1} is μ-main by (γ). Now $(b)_{k+1}$ is proved.

If $o' \cap v''_{k+1} = \emptyset$, then $(c)_{k+1}$ follows immediately from $(c)_k$. If $v''_{k+1} \subseteq o'$, then as we have shown above (for v'_{k+1}), $v''_{k+1} \subseteq o$ and $(c)_{k+1}$ follows from (γ). Suppose now that o' is a proper subterm of v''_{k+1} and v''_{k+1} has an v_k-ancestor v^*_k in s_k for which v''_{k+1} is a residual. Let u^*_k be the corresponding redex of v^*_k in t_k (it exists, because, by $(b)_k$, v^*_k is μ-main). Obviously, u_k is a proper subterm of u^*_k and since u_k is limit, it must be in an erased argument of u^*_k. Hence v_k is in a quasi-erased argument of v^*_k. Therefore o' is in a quasi-erased argument of v''_{k+1} and the quasi-main arguments of v''_{k+1} coincide with the corresponding quasi-main arguments of v^*_k. Thus, by $(c)_k$, in the quasi-main arguments of v''_{k+1} there are no occurrences of μ-symbols. To prove $(c)_{k+1}$, it remains to consider the case when o' is a proper subterm of v''_{k+1} and v''_{k+1} is created by v_k. If in quasi-main arguments of v''_{k+1} there are μ-symbols, then in main arguments of corresponding redex u''_{k+1} in s_{k+1}, which is also a u_k-new redex, there are descendants of redexes contracted in P. (Since v_k is μ-main, o' and hence u_{k+1} are also μ-main.) But each redex contracted in P is a limit redex. Thus, by Lemma 2.6, their descendants can not occur in arguments of new redexes. Hence, also in this case, there are no μ-symbols in quasi-main arguments of v''_{k+1}, and $(c)_{k+1}$ is valid. Now (1) is proved.

(2) By Lemma 2.2 and $(b)_n$.

Theorem 2.3 A limit strategy is perpetual in OTRSs. Moreover, if a term t in an OTRS R is strongly normalizable, then a limit strategy constructs a longest normalizing reduction starting from t, and its length coincides with the number of μ-occurrences in an R_μ-normal form of t.

Proof If a limit R-reduction P starting from t is normalizing, then by Lemma 2.7 its corresponding R_μ-reduction also is normalizing. Hence, by Lemma 2.3, t is strongly normalizable in R. Thus, the limit strategy is perpetual. Now, if t is strongly normalizable, Q is a normalizing R-reduction, and s is an R_μ-normal form of t, then $|Q|$ =(by Corollary 2.1)= $|Q_\mu| \leq$ (by the CR property of R_μ)$\leq \|s\|_\mu$ =(by Lemma 2.7)= $|P|$. Thus, P has the maximal length among all reductions of t to normal form.

We now give a direct proof of the fact that limit reductions are the longest. The advantage of the previous proof is that it additionally gives a characterization of lengths of longest reductions.

Lemma 2.8 Let $P : t_0 \xrightarrow{u_0} t_1 \xrightarrow{u_1} \ldots$ and $t_0 \xrightarrow{u} s_0$, where u is a limit redex. Then there is a reduction Q starting from s_0 such that $|Q| \geq |P| - 1$.

Proof Using the Church-Rosser theorem, we can construct the following diagram

$$\begin{tikzcd}
t_0 \arrow[r, "u_0"] \arrow[d, "u = P_0"'] & t_1 \arrow[r, "u_1"] \arrow[d, "P_1"'] & t_2 \arrow[r, dashed, no head] \arrow[d, "P_2"'] & {} \\
s_0 \arrow[r, "Q_0"'] & s_1 \arrow[r, "Q_1"'] & s_2 \arrow[r, dashed, no head] & {}
\end{tikzcd}$$

where in $P_i : t_i \twoheadrightarrow s_i$ all residuals of u in t_i are contracted, and in $Q_i : s_i \twoheadrightarrow s_{i+1}$ all P_i-residuals of u_i are contracted. It follows from Lemma 2.6 that u has at most one residual in each term t_i, which is complete. Thus if there is an i such that u_i is a residual of u, then Q_i is empty, $s_j = t_j$ for all $j > i$, and Q_j contracts exactly one redex. Otherwise, each reduction Q_k contracts at least one redex. Hence, we can take the concatenation of $Q_0, Q_1, \ldots$ for Q.

Corollary 2.4 The limit strategy is perpetual in Orthogonal TRSs, and the limit reductions are longest among reductions starting from a strongly normalizable term.

3. Applications

3.1 Recursive perpetual strategies for some classes of OTRSs

We now define some classes of OTRSs for which the limit strategy is efficient.

Definition 3.1 (1) Let $t \xrightarrow{u} s$ in an OTRS R, and let $v \subseteq s$ be a new redex. We call v *generated* if its pattern is in the pattern of the contractum of u. We call an OTRS R *persistent* (PTRS) if, for each reduction step in R, any created redex is generated.

(2) We call an OTRS *outside-creating* if, for any reduction step $t \xrightarrow{u} s$, any new redex in s contains at least one symbol above the contractum of u, and call it *inside-creating* if any new redex in s is inside the contractum of u.

(3) We call an OTRS *non-absorbing* if, for any reduction step $t \xrightarrow{u} s$, the arguments of any new redex in s are in the contractum of u.

(4) We call an OTRS *non-left-absorbing* (resp. *non-right-absorbing*) if, for any reduction step $t \xrightarrow{u} s$, any argument of a created redex in s is inside the contractum of u or to the right (resp. to the left) of it.

Proposition 3.1 (1) Let t be a term in a non-absorbing OTRS. Then any outermost redex in t is unabsorbed.

(2) Let t be a term in a non-left absorbing (resp. non-right-absorbing) OTRS. Then the leftmost-outermost (resp. the rightmost-outermost) redex in t is unabsorbed.

Proof From Definitions 2.5 and 3.1.

Thus a limit redex can be found efficiently in non-absorbing, non-left-absorbing, and non-right-absorbing systems. Note that left-normal OTRSs [9] (where in left-hand sides of rules function symbols precede variables), and Combinatory Logic in particular, are non-left-absorbing. Persistent and inside-creating systems are non-absorbing.

Proposition 3.2 Outside-creating OTRSs are strongly normalizing.
Proof Let t be a term in an outside-creating OTRS R. It is easy to see that R_μ is outside-creating as well. Thus, by Lemma 2.3, it is enough to prove that t has a normal form in R. Let $P : t \to t_1 \to \dots$ be an innermost reduction in which any created redex is contracted immediately after creation (each step creates at most one redex). Since a created redex is strictly above the contractum of a contracted redex, the number of redexes of terms in P is decreasing. Thus P is finite.

Remark 3.1 The above proposition is equivalent to Corollary 4.10 of van Raamsdonk [10] stating that all "superdevelopments" in OTRSs are finite.

3.2 On decidability of weak and strong normalization
Definition 3.2 We call a class $\Re$ of OTRSs *closed under μ-extension* if $\Re$ contains the μ-extension of each of its elements.

Proposition 3.3 Let $\Re$ be a class of OTRSs closed under μ-extension. Then decidability of weak normalization for OTRSs in $\Re$ implies decidability of strong normalization for OTRSs in $\Re$.
Proof For any $R \in \Re$, weak normalization is decidable for $R_\mu \in \Re$. Hence, by Proposition 2.2, strong normalization is decidable for R.

It is easy to see that all classes of OTRSs, defined in Definition 3.1, and the class of strongly sequential OTRSs [3] are closed under μ-extension. Note that although, for any OTRS R, R_μ contains infinitely many rules, it is decidable whether a term in R_μ is an R_μ-redex. Thus the decidability question makes sense for R_μ. We show in [5] that weak normalization is decidable in Persistent TRSs. Hence, by Proposition 3.3, strong normalization is also decidable for Persistent TRSs. For inside-creating TRSs, decidability of weak and strong normalization is open. For OTRSs in general, undecidability of strong normalization follows from undecidability of the (uniform) halting problem. Thus, weak normalization is also undecidable.

3.3 The least upper bound of lengths of reductions in persistent TRSs
We now present an algorithm for finding the least upper bound of lengths of coinitial reductions in persistent TRSs, which does not need to make an actual transformation of an input term. We first recall some results from [5].
Notation We write $t = (t_1//e_1, ..., t_k//e_k)e$ if $e_1, ..., e_k$ are non-overlapping proper subterms in e, and t is obtained from e by their replacement with $t_1, ..., t_k$, respectively.

Definition 3.3 We call a subterm s in t *free* if s is not a proper pattern-subterm of a redex in t.

Lemma 3.1 Let $t = (e//o)s$ in a persistent TRS, let e be free in t, and let s and e be strongly normalizable. Then t is strongly normalizable.
Proof sketch By persistency, all descendants of e remain free; therefore, infinitely many steps of an infinite reduction starting from t must be performed inside (descendants of) e or outside e, meaning that symbols of e do not belong to patterns of the redexes contracted outside e. But e and s are strongly normalizable. Therefore, t does not possess an infinite reduction.

Definition 3.4 We call a redex in a PTRS R *trivial* if it is a left-hand side of a rewrite rule in R. We call redexes u and v *similar* if they are instances the left-hand side of the same rule. We call a redex u *finite* if its similar trivial redex is strongly normalizable.

Lemma 3.2 ([5]) A term t in a PTRS is strongly normalizable iff any redex in t is finite.
Proof ($\Rightarrow$) Obviously, any redex u in t is strongly normalizable. So, by Lemma 3.1, the trivial redex similar to u is strongly normalizable as well. ($\Leftarrow$) By induction on the number n of redexes in t. The case $n = 0$ is trivial. So let $n > 0$, u be an innermost redex in t, and $s = (x//u)t$. By the induction assumption, s is strongly normalizable. Since arguments of u are in normal form and u is a finite redex, Lemma 3.1 implies that u is strongly normalizable. Since $t = (u//x)s$, we have again by Lemma 3.1 that t is strongly normalizable.

Definition 3.5 let R be a PTRS and $r \in R$. We call an *r-tree* the maximal tree with rules as nodes and r as the root, such that a redex corresponding to a node has an occurrence in the right-hand side of its ancestor node.

Lemma 3.3 ([5]) An r-redex u in a PTRS R is finite iff the r-tree is finite.
Proof ($\Rightarrow$) If the r-tree has an infinite branch $r_0, r_1, \ldots$, then one can construct infinite reduction of the trivial r-redex $v \xrightarrow{v} t_1 \xrightarrow{v_1} t_2 \xrightarrow{v_2} \ldots$, where v_1 is an r_1-redex created by v, v_2 is an r_2-redex created by v_1, and so on. ($\Leftarrow$) By induction on the height n of the r-tree. If $n = 0$, then contractum s of the trivial r-redex v is normal form. If $n > 1$, then by the induction assumption, each redex in s is finite. Hence, by Lemma 3.2, s is strongly normalizable. Thus, v is also strongly normalizable, i.e., u is finite.

Definition 3.6 Let R be a PTRS.

(1) Let t be a term in R_μ, let $s \subseteq t$, and let $P : t \twoheadrightarrow e$ be the rightmost innermost normalizing R_μ-reduction. Then, by definition, $Mult_\mu(s,t)$ is the number of P-descendants of s in e

(2) Let $u = C[e_1, \ldots, e_n]$ be an r-redex in R_μ, let $s' \subseteq e_i$, let $v = C[o_1, \ldots, o_n]$ be an r-redex (similar to u) with arguments $o_1, \ldots, o_n$ in R_μ-normal form, and let $Q : v \twoheadrightarrow o$ be the rightmost innermost normalizing R_μ-reduction. Then, by definition, $mult_\mu(u,i) = mult_\mu(u,s') = mult_\mu(r,i) = Mult_\mu(o_i, v)$, and $mult_\mu(u) = mult_\mu(r)$ is the number of μ-subterms in o that appear during Q, i.e., that are not descendants of subterms with head-symbol μ from (arguments of) u. Numbers $mult_\mu(u,i)$ and $mult_\mu(r,i)$ are *proper μ-indices* of u and r, and numbers $mult_\mu(u)$ and $mult_\mu(r)$ are *μ-indices* of u and r.

The following lemma implies that the definition is correct.

Lemma 3.4 Let $u = C[e_1, \ldots, e_n]$ and $v = C[o_1, \ldots, o_n]$ be similar redexes with arguments in normal form in a PTRS R_μ, and let $P : u = t_0 \xrightarrow{u_0} t_1 \xrightarrow{u_1} \ldots$ and $Q : v = s_0 \xrightarrow{v_0} s_1 \xrightarrow{v_1} \ldots$ be the rightmost innermost normalizing R_μ-reductions. Then, for each $i = 0, 1, \ldots$, there is an one-to-one correspondence between

1. redexes in t_i and s_i such that corresponding redexes are similar, and u_i and v_i are corresponding.
2. "appeared μ-subterms", which are descendants of contracta of R_μ-redexes in t_i and s_i.

Proof Easy induction on i, using persistency of R_μ.

Lemma 3.5 Let t be a normalizable term in a PTRS R_μ, let $e \subseteq s \subseteq t$, and let s be in R_μ-normal form. Then $Mult_\mu(s,t) = Mult_\mu(e,t)$.
Proof Let $t \twoheadrightarrow o$ be the rightmost innermost normalizing R_μ-reduction. Then the descendants of s in o are disjoint occurrences of s, and each of them contains exactly one descendant of e.

Notation $L(t)$ denotes the least upper bound of lengths of reductions starting from t.

Lemma 3.6 Let t be a strongly normalizable term in a PTRS R and let $u_1, \ldots, u_n$ be all redexes in t. Then

$$L(t) = \sum_{i=1}^{n} Mult_\mu(u_i, t) mult_\mu(u_i)$$

Proof Let $P : t \twoheadrightarrow o$ be the rightmost innermost normalizing R_μ-reduction and let $u_1, \ldots, u_n$ be the enumeration of redexes in t from right to left. In the fragment of P where u_i is reduced to R_μ-normal form, $mult_\mu(u_i)$ new μ-symbols appear (in the beginning of the fragment, all arguments of u_i are in R_μ-normal form). By Lemma 3.5, during the rest of P each of these $mult_\mu(u_i)$ μ-occurrences is duplicated $Mult_\mu(u_i, t)$-times. Hence

$$\|o\|_\mu = \sum_{i=1}^{n} Mult_\mu(u_i, t) mult_\mu(u_i)$$

and the lemma follows from Theorem 2.3.

Lemma 3.7 Let t be a strongly normalizable term in a PTRS R_μ, let $s \subseteq t$, and let $u_1, \ldots, u_n$ be all redexes in t that contain s in their arguments. Suppose that s is in m_i-th argument of u_i $(i = 1, \ldots, n)$. Then

$$Mult_\mu(s, t) = \prod_{i=1}^{n} mult_\mu(u_i, s) = \prod_{i=1}^{n} mult_\mu(u_i, m_i)$$

Proof Let $P : t \twoheadrightarrow o$ be the rightmost innermost normalizing R_μ-reduction. It follows from Lemma 3.5 that, in the fragment of P in which u_i is reduced to R_μ-normal form, each descendant of s is duplicated $mult_\mu(u_i, s) = mult_\mu(u_i, m_i)$-times. Thus, the lemma is obvious.

Lemma 3.8 Let $u = C[e_1, \ldots, e_k]$ be an r-redex with arguments $e_1, \ldots, e_k$ in normal form in a PTRS R_μ. Then for all $j = 1, \ldots, k$:

$$mult_\mu(u, j) = mult_\mu(r, j) = \sum_{i=1}^{m_j} Mult_\mu(e_{j_i}, o),$$

$$mult_\mu(u) = mult_\mu(r) = \sum_{i=1}^{m} Mult_\mu(u_i, o) mult_\mu(u_i) + 1,$$

where o is the contraction of u in R_μ, $e_{j_1}, \ldots, e_{j_{m_j}}$ are all descendants of e_j in o, and $u_1, \ldots, u_m$ are all redexes in o.

Proof From Definition 3.6 (since all descendants of subterms e_{j_i} are pairwise disjoint).

Theorem 3.1 Let t be a term in a PTRS R. Then the least upper bound $L(t)$ of lengths of reductions starting from t can be found by the following

Algorithm 3.1 Let $r_1, \ldots, r_n$ be all rules in R such that an r_i-redex has an occurrence in t $(i = 1, \ldots, n)$. If the r_i-tree is not finite for at least one i, then $L(t) = \infty$. Otherwise, using Lemmas 3.8 and 3.7, find μ-indices and proper μ-indices of all rules r_i. Finally, using Lemmas 3.6 and 3.7, find $L(t)$.

Proof From Theorem 2.3 and Lemmas 3.2, 3.3, and 3.6-3.8.

3.4 The least upper bound of lengths of developments
Let $R = \{r_i : t_i \to s_i\}$ be an OTRS and let $\underline{R} = \{\underline{r}_i : \underline{t}_i \to s_i\}$, where $\underline{t}_i$ is obtained from t_i by underlining its head-symbol; *terms* in $\underline{R}$ are constructed in the usual way with the restriction that underlined symbols may only occur as head-symbols of redexes. Then, for each development $P : e_0 \to e_1 \to \ldots \to e_n$ of e_0 in R (in which only residuals of redexes from e_0 are contracted), there is a reduction $P' : e'_0 \to e'_1 \to \ldots \to e'_n$ in $\underline{R}$ such that e'_i is obtained from e_i by underlining head-symbols of residuals of redexes from e_0. Obviously, $\underline{R}$ is persistent, since no creation of redexes is possible in it. Thus, to find least upper bounds of developments in R, one can use Algorithm 3.1, which becomes simpler in this case. For any rule $\underline{r} \in \underline{R}$, $mult_\mu(\underline{r}) = 1$, and proper μ-indices of $\underline{r}$ can be found immediately from the right-hand side of $\underline{r}$. Indeed, if $\underline{r} : C[x_1, \ldots, x_n] \to s$, then $mult_\mu(\underline{r}, i) = 1$ if x_i does not occur in s, and $mult(\underline{r}, i)$ coincides with the number of occurrences of x_i in s otherwise.

Acknowledgments I enjoyed discussions with H. Barendregt, J. W. Klop, J.-J. Lévy, L. Maranget, G. Mints, Sh. Pkhakadze, Kh. Rukhaia, and V. Sazonov. I also would like to thank G. Gonthier, G. Tagviashvili, K. Urbaitis, and F. J. de Vries for their help in preparation of this paper.

REFERENCES

1. Barendregt H. P., Bergstra J., Klop J. W., Volken H. Some notes on lambda-reduction, in: Degrees, reductions and representability in the lambda calculus. Preprint no.22, University of Utrect, Department of mathematics, p. 13-53, 1976.
2. Dershowitz N., Jouannaud J.-P. Rewrite Systems. In: J.van Leeuwen ed. Handbook of Theoretical Computer Science, Chapter 6, vol.B, 1990, p. 243-320.
3. Huet G., Lévy J. -J. Computations in Orthogonal Rewriting Systems. In Computational Logic, Essays in Honor of Alan Robinson, ed. by J. -L. Lassez and G. Plotkin, MIT Press, 1991.
4. Khasidashvili Z. Optimal normalization in orthogonal term rewriting systems. In: Proc. of the fifth International Conference on Rewriting Techniques and Applications, Springer LNCS, vol. 690, C. Kirchner, ed. Montreal, 1993, p. 243-258.
5. Khasidashvili Z. On the equivalence of persistent term rewriting systems and recursive program schemes. In: Proc. of the second Israel Symposium on Theory and Computing Systems, Natanya, 1993, p. 240-249.
6. Klop J. W. Combinatory Reduction Systems. Mathematical Centre Tracts n. 127, CWI, Amsterdam, 1980.
7. Klop J. W. Term Rewriting Systems. In: S. Abramsky, D. Gabbay, and T. Maibaum eds. Handbook of Logic in Computer Science, vol. II, Oxford University Press, 1992, p. 1-116.
8. Nederpelt R. P. Strong normalization for a typed lambda-calculus with lambda structured types. Ph. D. Thesis, Eindhoven, 1973.
9. O'Donnell M. J. Computing in systems described by equations. Springer LNCS 58, 1977.
10. Van Raamsdonk F. A simple proof of confluence for weakly orthogonal combinatory reduction Systems. Report CS-R9234, CWI Amsterdam, 1992.