Reinterpreting Compression in Infinitary Rewriting

Jeroen Ketema

Faculty EEMCS, University of Twente
PO Box 217, 7500 AE Enschede, the Netherlands
`j.ketema@ewi.utwente.nl`

Abstract

Departing from a computational interpretation of compression in infinitary rewriting, we view compression as a degenerate case of standardisation. The change in perspective comes about via two observations: (a) no compression property can be recovered for non-left-linear systems and (b) some standardisation procedures, as a 'side-effect', yield compressed reductions.

1998 ACM Subject Classification F.1.1 Models of Computation, F.3.2 Semantics of Programming Languages, F.4.2 Grammars and Other Rewriting Systems

Keywords and phrases term rewriting, infinitary rewriting, compression, standardisation

Digital Object Identifier 10.4230/LIPIcs.RTA.2012.209

Category Regular Research Paper

1 Introduction

One of the most fundamental properties studied in infinitary rewriting is the so-called *compression property*. Roughly, the property states that for every reduction of transfinite length a 'similar' reduction can be found of length at most ω (the first infinite ordinal). Consider, e.g., the binary function symbol f and the rules $a \to g(a)$ and $b \to g(b)$. We have the following reduction of length $\omega + \omega$:

$$f(a,b) \to f(g(a),b) \to \cdots \to f(g^n(a),b) \to \cdots \\ f(g^\omega,b) \to f(g^\omega,g(b)) \to \cdots \to f(g^\omega,g^n(b)) \to \cdots f(g^\omega,g^\omega)\,.$$

By interleaving the a- and b-steps, we can compress this reduction to obtain a 'similar' reduction of length ω:

$$f(a,b) \to f(g(a),b) \to f(g(a),g(b)) \to \cdots \to f(g^n(a),g^n(b)) \to \cdots f(g^\omega,g^\omega)\,.$$

This second reduction has a very appealing property: We can obtain an arbitrarily good approximation of the final term by rewriting $f(a,b)$ a sufficient, *finite* number of times. As we are now in the realm of finite rewriting, it can be said that compression gives computational meaning to infinitary rewriting (see also [8], although room is left there for other interpretations than a computational one).

There are, however, two problems with the aforementioned computational interpretation. First, the compression property does not apply to all rewrite systems, while it can be argued that every rewrite system computes something. In particular, the property can fail for systems with non-left-linear rules. Consider, e.g., the non-left-linear rule $f(x,x) \to c$. This rule, in combination with the rules $a \to g(a)$ and $b \to g(b)$ from above, yields the standard counterexample to compression for non-left-linear systems [5, 8, 7]; the following reduction is of length $\omega + 1$ and cannot be compressed:

$$f(a,b) \to f(g(a),b) \to f(g(a),g(b)) \to \cdots \to f(g^n(a),g^n(b)) \to \cdots f(g^\omega,g^\omega) \to c\,.$$

23rd International Conference on Rewriting Techniques and Applications (RTA'12).
Editor: A. Tiwari; pp. 209–224
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Essentially, we need ω steps to obtain $f(g^\omega, g^\omega)$ before we can rewrite to c.

The second problem with the computational interpretation has to do with the fact that infinitary rewriting is susceptible to a similar computational treatment [10] as that of the real numbers in computable analysis [15]. We can think of both terms and reductions as Turing Machines. As such, any term along every transfinite reduction can be approximated with arbitrary precision in finite time (by executing the Turing Machines). Of course, such a computational treatment is not perfect: Although terms and reductions can be represented by Turning Machines, we cannot compute whether a reduction step that occurs in any such representation is indeed 'valid', as computing the validity of a match would take infinite time in case a non-left-linear rule is employed. Instead, with each representation of a reduction one should provide a 'certificate' (i.e. a proof) showing the validity of the matches that occur.

In the current paper we address both points of critique regarding the computational interpretation. To address the first, i.e. the lack of compression in non-left-linear systems, it would suffice if we could establish a *generalised* compression property: For each system a countable ordinal α might exist such that each reduction within that system can be compressed to one of length at most α [11, 12]. Unfortunately, as we will show, such a generalised compression property fails even for very simple systems. Hence, and as is also needed to address the second point of critique, a different interpretation of compression — one outside the realm of computability — is warranted for.

We reinterpret compression as a degenerate case of standardisation. Such a reinterpretation is not unexpected: Likewise to connections that exist between the equivalence of reductions and standardisation in finite rewriting [14], connections can be drawn between equivalence and compression in infinitary rewriting [8]. To enable our reinterpretation, we provide the first ever standardisation procedures for infinitary Term Rewriting Systems (iTRSs). We will show that these procedures, as a 'side-effect', yield compressed reductions.

The paper now proceeds as follows: In Section 2, we state a number of preliminaries needed in the remainder of the paper. We discuss the generalised compression property and its failure in Section 3. In Section 4, we formulate two standardisation procedures. Finally, in Section 5, we conclude.

2 Preliminaries

We briefly review some basic facts regarding infinitary Term Rewriting Systems (iTRSs); see [8, 7] for more detailed accounts. Throughout, we denote the first infinite ordinal by ω and arbitrary ordinals by α, β, γ, and so on; $[\alpha, \beta)$ denotes a left-closed, right-open interval of ordinals and $(\alpha, \beta]$ a right-open, left-closed interval. By $\mathbb{N}$ we denote the natural numbers including zero.

Terms. Let Σ be a signature, each element of which has finite arity, and let V be a countably infinite set of variables. The set of (finite and infinite) terms is commonly defined by metric completion [2, 8, 7]. Here, we give the shorter, but equivalent, definition from [3, 9].

► **Definition 2.1.** The set of *terms* $\mathcal{T}er(\Sigma, V)$ is defined coinductively such that x is a term for each $x \in V$ and if $f(t_1, \ldots, t_n)$ is a term, then $f \in \Sigma$ is n-ary and $t_1, \ldots, t_n$ are terms.

Substitutions over terms are defined by interpreting the usual definition coinductively [8, 7]. For the *root symbol*, $\text{root}(t)$, of a term t we have $\text{root}(x) = x$ and $\text{root}(f(t_1, \ldots, t_n)) = f$.

The set of positions $\mathcal{P}\text{os}(t)$ of a term t is a set of *finite* strings over $\mathbb{N}$ representing the 'locations' of subterms in t [8, 7]. Denoting the empty string by ϵ, we have $\mathcal{P}\text{os}(x) = \{\epsilon\}$

and $\mathcal{P}os(f(t_1,\ldots,t_n)) = \{\epsilon\} \cup \bigcup_{1\le i\le n}\{i\cdot p \mid p \in \mathcal{P}os(t_i)\}$. If p is a position of t, then $t|_p$ denotes the *subterm* of t at position p; we have $t|_\epsilon = t$ and $f(t_1,\ldots,t_i,\ldots,t_n)|_{i\cdot p} = t_i|_p$. By $t[s]_p$ we denote the *replacement* of the subterm at position p in t by s; we have $t[s]_\epsilon = s$ and $f(t_1,\ldots,t_i,\ldots,t_n)[s]_{i\cdot p} = f(t_1,\ldots,t_i[s]_p,\ldots,t_n)$. The *length* of p is denoted $|p|$. There exists a well-founded order $<$ on positions: $p < q$ iff p is a proper prefix of q. We write $\le$ for the reflexive closure of $<$. If neither $p \le q$ nor $q \le p$, then p and q are *parallel* and we write $p \parallel q$. The concatenation of positions p and q is denoted by $p \cdot q$.

Rewrite Rules and Reductions. Rewrite rules and iTRSs are defined as in the finite case, except that the finiteness restriction on the right-hand side of rewrite rules is dropped:

▶ **Definition 2.2.** A *rewrite rule* is a pair of terms (l, r), denoted $l \to r$, with l finite and such all variables that occur in r also occur in l. A rewrite rule is *left-linear*, if each variable occurs at most once in l.

An *infinitary Term Rewriting System (iTRS)* is a pair $\mathcal{R} = (\Sigma, R)$ with Σ a signature and R a set of rewrite rules over Σ. An iTRS is *left-linear* if all its rewrite rules are.

Rewrite steps are now defined as usual:

▶ **Definition 2.3.** A *rewrite step* is a pair of terms $s \to t$ adorned with a position p and a rewrite rule $l \to r$ such that $s = s[\sigma(l)]_p$ and $t = s[\sigma(r)]_p$ for some substitution σ. The term $\sigma(l)$ is called an $l \to r$*-redex*. The redex *occurs* at position p and depth $|p|$ in s.

The previous gives sufficient background to define strongly convergent reductions [8, 7] (the most common notion of reduction in infinitary rewriting; see [6] for further discussion of notions of reduction in infinitary rewriting).

▶ **Definition 2.4.** A *strongly convergent reduction* of ordinal length α, denoted $t_0 \twoheadrightarrow^{\alpha} t_\alpha$, is a pair consisting of sequence of *terms* $(t_\beta)_{\beta<\alpha+1}$ and a sequence of *steps* $(p_\beta, l_\beta \to r_\beta)_{\beta<\alpha}$ such that for all $\beta < \alpha$ it holds that (a) $t_\beta \to t_{\beta+1}$ is a rewrite step adorned with the position p_β and the rewrite rule $l_\beta \to r_\beta$, and (b) if β is a limit ordinal, then t_γ converges to t_β and $|p_\gamma|$ tends to infinity when γ approaches β from below.

Here, a sequence of terms $(t_\gamma)_{\gamma<\beta}$ is said to converge to a term t_β whenever the depth up to which t_γ and t_β are identical increases as γ approaches β.

In case we are only interested in an upper bound, respectively a lower bound, α on the length of a strongly convergent reduction we write $s \twoheadrightarrow^{\le\alpha} t$, respectively $s \twoheadrightarrow^{\ge\alpha} t$. Moreover, in case the length is irrelevant, respectively finite, we write $s \twoheadrightarrow t$, respectively $s \to^* t$. We say that a reduction $s \twoheadrightarrow t$ is *maximal* if there does not exist a term t' such that $t \to t'$. In other words, in this case t is a *normal form*.

▶ **Example 2.5.** Consider the rewrite rule $a \to g(a)$ from the introduction. The following is a strongly convergent, maximal reduction with normal form g^ω:

$$a \to g(a) \to g(g(a)) \to \cdots \to g^n(a) \to \cdots g^\omega\,.$$

Replacing the final term of the reduction by a breaks strong convergence, as the sequence $(g^n(a))_{n<\omega}$ does not converge to a; it only converges to g^ω.

Consider now the rule $c \to c$. We can construct the following reduction of length ω:

$$c \to c \to \cdots \to c \to \cdots c\,.$$

Although a sequence in which all terms are equal to c obviously converges to c, this reduction is not strongly convergent, as the depth of the contracted redexes does not tend to infinity along the reduction.

Compression. Having defined iTRSs and strongly convergent reductions we can now state the classical compression property for left-linear systems [8, 7]:

▶ **Theorem 2.6.** *Let $\mathcal{R}$ be a left-linear iTRS. For each $s \twoheadrightarrow t$ there exists a reduction $s \twoheadrightarrow^{\leq\omega} t$.*

Thus, for every reduction in a left-linear iTRS we can find a reduction with the same initial and final term that is of length at most ω. In the case of *orthogonal* iTRSs the above can be strengthened [8]: a reduction $s \twoheadrightarrow^{\leq\omega} t$ exists which is Lévy equivalent to $s \twoheadrightarrow t$.

3 Generalised Compression

Assuming compression should hold equally for all reductions within a rewrite system, the only obvious generalisation of the compression property is the one from [12]:

▶ **Definition 3.1.** Let α be a countable ordinal, an iTRS $\mathcal{R}$ satisfies the *α-compression property* iff it holds for each $s \twoheadrightarrow t$ that there exists a reduction $s \twoheadrightarrow^{\leq\alpha}$ t.

Since all strongly convergent reductions have countable length [8, 7], only countable ordinals make sense in the definition; the property holds trivially for any uncountable ordinal. By Theorem 2.6, we have that ω-compression holds for every left-linear iTRS.

▶ Remark. The α-compression property is related to the notion of α-closedness from [5]. The difference between the notions is two-fold: First, α-compression allows for both finite and infinite terms as the starting terms of reductions, while α-closedness only allows for finite terms as starting terms. Second, while we consider strongly convergent reductions, α-closedness is concerned with the more general, but less well-behaved, Cauchy convergent or weakly convergent reductions (which dispose of the depth requirement that is part of Definition 2.4 [5, 8]).

Based on the known counterexamples to compression from the literature, one may conjecture that for every iTRS it is possible to find an ordinal α such that α-compression holds: With regard to the standard counterexample from the introduction, [5] states without proof that for *finite* starting terms all reductions are compressible to length at most $\omega + \omega$. A similar property holds for $\lambda\beta\eta$-calculus — the prototypical higher-order system not satisfying the compression property. As can be inferred from [12, Lemma 5], all reductions in $\lambda\beta\eta$-calculus also compress to reductions of length at most $\omega + \omega$.

As we will show below, the above conjecture is false: We uncover two systems that do not satisfy α-compression for any countable α. The first system, discussed in Section 3.1, is — remarkably enough — the one from the standard counterexample. The second system, discussed in Section 3.2, refutes the above conjecture in even stronger ways: We exhibit terms that have unique strongly convergent reductions starting from them which are *incompressible*:

▶ **Definition 3.2.** Let $\mathcal{R}$ be an iTRS, a reduction $s \twoheadrightarrow^{\alpha} t$ is *incompressible* if all reductions $s \twoheadrightarrow t$ are of length at least α.

Considering incompressible reductions mitigates a possible point of critique regarding the proof in Section 3.1: We only show that the considered reductions are at least of a certain length, we do not show that they cannot be compressed precisely up to that length.

3.1 Compression in the Standard Counterexample

Consider the iTRS with as its sole rule $f(x, x) \to c$, the non-left-linear rule from the standard counterexample. We have the following:

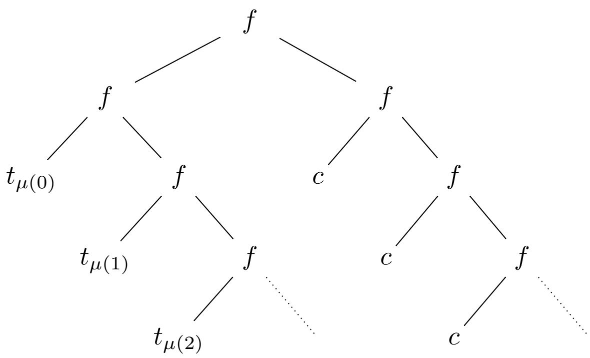

Figure 1 A limit term from the proof of Lemma 3.3.

▶ **Lemma 3.3.** *For every countable ordinal α, there is a term t_α such that $t_\alpha \twoheadrightarrow c$ exists and is of length at least $\alpha + 1$ and such that no $t_\alpha \twoheadrightarrow c$ exists of length less than $\alpha + 1$.*

Proof. We prove the lemma by transfinite induction over the ordinal α. In case $\alpha = 0$, define $t_\alpha = f(c, c)$. As we have $f(c, c) \to c$, we are done.

In case $\alpha = \beta + 1$, we have by the induction hypothesis that a strongly convergent reduction exists from t_β to c. Set $t_\alpha = f(t_\beta, c)$. The reduction $t_\alpha \twoheadrightarrow^{\geq \beta+1} f(c, c) \to c$ is obviously strongly convergent, as $t_\beta \twoheadrightarrow c$ is. Moreover, as the redex at the root is only created after at least $\beta + 1$ steps — for $t_\beta \twoheadrightarrow c$ consists of at least $\beta + 1$ steps — we have that $t_\alpha \twoheadrightarrow c$ is of length at least $\alpha + 1 = \beta + 2$.

In case α is a limit ordinal, choose any bijection μ from $\mathbb{N}$ to α and define $t_\alpha = f(\phi_\mu(0), \psi)$ (see also Figure 1), where:

$$\phi_\mu(n) = f(t_{\mu(n)}, \phi_\mu(n + 1))$$
$$\psi = f(c, \psi)$$

By the induction hypothesis we have for every $n \in \mathbb{N}$ that a strongly convergent reduction of length at least $\mu(n) + 1$ starts from $t_{\mu(n)}$. Reduce each $t_{\mu(n)}$ to c, possibly interleaving the reductions in the different subterms. Doing so, we obtain $f(\psi, \psi)$, which we can further reduce to c. The reduction to $f(\psi, \psi)$ is strongly convergent, for suppose not, then an infinite number of reduction steps occurs at a certain fixed depth [7, Exercise 12.3.6]. By the pigeonhole principle for infinite sets and since the terms $t_{\mu(n)}$ occur at increasingly greater depths, this means there exists an m such that an infinite number of steps occurs in $t_{\mu(m)}$ at a fixed position. Hence, the reduction from $t_{\mu(m)}$ is not strongly convergent, contradiction. As the reductions from the different $t_{\mu(n)}$ are independent and as a reduction of length at least $\beta + 1$ occurs for every ordinal $\beta < \alpha$, the constructed reduction has length at least $\alpha + 1$. No reduction of length $\gamma < \alpha$ exists, otherwise for infinitely many $t_{\mu(n)}$ with reductions of length at least $\beta + 1 > \gamma$ there also exist reductions of length at most γ, contradicting the induction hypothesis. ◀

Remark that the rules $a \to g(a)$ and $b \to g(b)$ from the standard counterexample do not affect the above result, as a and b do not occur in the defined terms. Hence, the lemma also holds for the iTRS from the standard counterexample.

By the above, we immediately have:

▶ **Theorem 3.4.** *There exists an iTRS such that the α-compression property fails to hold for every countable ordinal α.*

3.2 Unique Incompressible Reductions

Although the result from the previous section establishes that compression, even in a generalised form, cannot be carried over to arbitrary iTRSs, the result is not completely satisfactory: We only uncovered reductions that are of *at least* a certain length; we did not show that these reductions cannot be compression at all (up to the provided lower bounds on their lengths). To rectify this situation, we study a second system with a non-left-linear rule in the current section. For this second system, we will define for every countable ordinal α terms t_{ν_α} with starting from them unique maximal and incompressible reductions of length *at least* α. As a corollary we will obtain that each t_{ν_α} also has a unique incompressible reduction of length *precisely* α starting from it.

Before we give the actual proofs, we first describe our second system and give some examples.

Signature. The signature Σ of our second system consists of five symbols:

$$\Sigma = \{f, f', h, h', k\}\,,$$

where f and f' are ternary and such that all other symbols are unary. Of these symbols, f is the most important and is the root symbol of our unique non-left-linear rule. The symbol h will be inserted above the redexes we are going to contract to ensure that the redexes occur at sufficient depth to guarantee strong convergence. Moreover, the symbols f' and h' will allow us to construct terms that are already of the required 'shape', but in which no redexes occur (due to the presence of the primes). Finally, the function symbol k will help us to convert 'primed' into 'unprimed' terms.

Rewrite Rules. Our system has four rewrite rules: One non-left-linear rule and three rules to convert 'primed' into 'unprimed' terms. The non-left-linear rule is as follows:

$$f(x, x, y) \to h(k(y)) \tag{1}$$

That is, we match the first two arguments of f.

Contracting an f-redex introduces an h to ensure that the term substituted for y will eventually — after contracting of a number of k-redexes — occur at the same depth as before contraction of the f-redex.

As already mentioned, the rules associated with k convert 'primed' into 'unprimed' terms. The rules are as follows:

$$k(f'(x, y, z)) \to f(k(x), y, z) \tag{2}$$

$$k(h'(x)) \to h(k(x)) \tag{3}$$

$$k(h(x)) \to h(x) \tag{4}$$

The first two rules remove a prime and then recurse. The third rule ensures that the conversion terminates once we encounter a non-primed h. Although a similar terminating rule could be introduced for f, i.e. $k(f(x, y, z)) \to f(x, y, z)$, the terms considered below are such that this rule is never needed and, hence, we prefer omit the rule.

▶ Remark. Observe that the first two rules associated with k move the k to a lower position and that the last rule removes the k all together. Hence, if we would instead use $f(x, x, y) \to k(y)$ as our f-rule, it not longer holds that the (converted) term substituted for y eventually occurs at the same depth d as before application of the f-rule; it will occur at depth $d - 1$ instead. The equal depth is, however, needed to ensure that our reductions are strongly convergent. For this reason the h occurs at the root of the right-hand side of our f-rule.

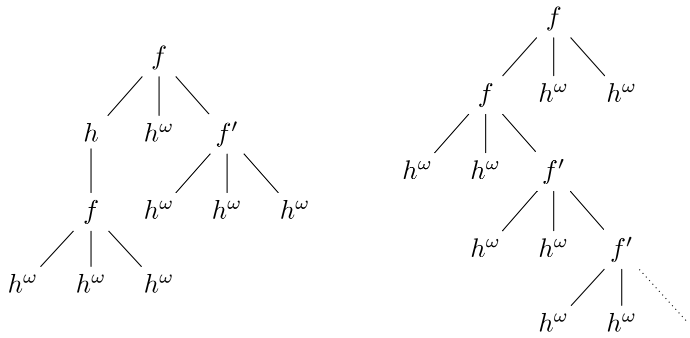

(a) Example 3.5. **(b)** Example 3.6.

Figure 2 The initial terms from Examples 3.5 and 3.6.

Two Examples. Before proving the central result of this section, let us consider two concrete incompressible reductions that make use of the above rules.

▶ **Example 3.5.** Consider the following term (also depicted in Figure 2(a)):

$$s = f(h(f(h^\omega, h^\omega, h^\omega)), h^\omega, f'(h^\omega, h^\omega, h^\omega)) .$$

Underlining contracted redexes, we have precisely one maximal reduction starting from s (consisting of seven steps):

$$\begin{aligned} f(h(\underline{f(h^\omega, h^\omega, h^\omega)}), h^\omega, f'(h^\omega, h^\omega, h^\omega)) &\to f(h(h(\underline{k(h^\omega)})), h^\omega, f'(h^\omega, h^\omega, h^\omega)) \\ &\to \underline{f(h^\omega, h^\omega, f'(h^\omega, h^\omega, h^\omega))} \to h(k(\underline{f'(h^\omega, h^\omega, h^\omega)})) \\ &\to h(f(\underline{k(h^\omega)}, h^\omega, h^\omega)) \to h(\underline{f(h^\omega, h^\omega, h^\omega)}) \to h(h(\underline{k(h^\omega)})) \to h^\omega . \end{aligned}$$

This reduction cannot be compressed to a reduction of less than seven steps: For each step the redex contracted in that step is created in the step immediately preceding it.

With regard to the above example, note that, if we fix the ordinal α to be 3, we obtain a bijection ν_α^{-1} between α and the depths of the α $(= 3)$ f-steps in the reduction: $\nu_\alpha^{-1}(0) = 2$, $\nu_\alpha^{-1}(1) = 0$, and $\nu_\alpha^{-1}(2) = 1$; the first f-step occurs at depth 2, the second at depth 0, and the third at depth 1. Bijections like ν_α^{-1} will be central in the construction of reductions such as the above. As the inverse of the bijection will be more important below, we choose to write ν_α^{-1} here instead of ν_α.

▶ **Example 3.6.** Consider the following system of recursive equations, with the solution of s depicted in Figure 2(b):

$$\begin{aligned} s &= f(f(h^\omega, h^\omega, t), h^\omega, h^\omega) \\ t &= f'(h^\omega, h^\omega, t) \end{aligned}$$

As in the previous example, we have exactly one maximal reduction starting from s (in this case one of length $\omega + 2$):

$$\begin{aligned} f(\underline{f(h^\omega, h^\omega, t)}, h^\omega, h^\omega) &\to f(h(\underline{k(t)}), h^\omega, h^\omega) \to f(h(f(\underline{k(h^\omega)}, h^\omega, t)), h^\omega, h^\omega) \\ &\to f(h(\underline{f(h^\omega, h^\omega, t)}), h^\omega, h^\omega) \to f(h^2(\underline{k(t)}), h^\omega, h^\omega) \to \cdots \\ &\to f(h^n(\underline{k(t)}), h^\omega, h^\omega) \to \cdots \underline{f(h^\omega, h^\omega, h^\omega)} \to h(\underline{k(h^\omega)}) \to h^\omega . \end{aligned}$$

Like the reduction from the previous example, the above reduction cannot be compressed, as each contracted redex is created in the step immediately preceding it.

As before, we have with regard to the above example that a bijection ν_α^{-1} exists from an ordinal α (in this case $\omega + 1$) to the depths of the α f-steps that occur along the reduction: $\nu_\alpha^{-1}(i) = i + 1$ for all $i < \omega$ and $\nu_\alpha^{-1}(\omega) = 0$.

Initial Term Construction. Following the above examples, our incompressible reductions will be defined by specifying initial terms and showing that each of these initial terms allows for exactly one maximal reduction (which cannot be compressed).

To define the initial terms, we first introduce a function $\rho : \mathcal{T}er(\Sigma, V) \to \mathcal{T}er(\Sigma, V)$ which maps 'unprimed' terms to 'primed' ones:

$$\rho(f(s_1, s_2, s_3)) = f'(\rho(s_1), s_2, s_3)$$
$$\rho(h(s)) = \begin{cases} h'(\rho(s)) & \text{if } s \neq h^\omega \\ h^\omega & \text{if } s = h^\omega \end{cases}$$

for all $s_1, s_2, s_3, s \in \mathcal{T}er(\Sigma, V)$.

The rewrite rules for k 'cancel out' ρ:

► **Lemma 3.7.** *Let $t \in \mathcal{T}er(\Sigma, V)$. It holds that $k(\rho(t)) \twoheadrightarrow^{\leq\omega} t$. Moreover, if $t|_{1^n} = h^\omega$ for some $n \in \mathbb{N}$, then $k(\rho(t)) \twoheadrightarrow^{\leq n} t$.*

Proof. By definition of ρ, we have that exactly one maximal strongly convergent reduction can be defined starting from $\rho(t)$ and employing only k-rules. Also by definition of ρ, the reduction is finite in case $t|_{1^n} = h^\omega$ for some $n \in \mathbb{N}$. That the final term of the reduction is t follows by structural induction over the positions of t, observing that $k(\rho(f(s_1, s_2, s_3))) \to f(k(\rho(s_1)), s_2, s_3)$ and that, moreover, $k(\rho(h(s))) \to h(k(\rho(s)))$ in case $s \neq h^\omega$, and $k(\rho(h(s))) \to h^\omega$ in case $s = h^\omega$. ◄

We can now define the initial terms of our incompressible reductions, where it is strongly suggested that the reader satisfies him- or herself of the fact that the initial terms from Examples 3.5 and 3.6 can be constructed by employing this definition.

► **Definition 3.8.** Let α be a countable ordinal and $\nu_\alpha : D \to \alpha$ a bijection with $D \subseteq \mathbb{N}$. The term t_{ν_α} is defined by $\tau_{\nu_\alpha}(0, 0, \alpha)$ where $\tau_{\nu_\alpha} : \mathbb{N} \times (\alpha + 1)^2 \to \mathcal{T}er(\Sigma, V)$ is such that:

$$\tau_{\nu_\alpha}(d, \delta, \gamma) = \begin{cases} f(\tau_{\nu_\alpha}(d+1, \delta, \nu_\alpha(d)), \\ \qquad h^\omega, \rho \circ \tau_{\nu_\alpha}(d+1, \nu_\alpha(d)+1, \gamma)) & \text{if } d \in D \text{ and } \nu_\alpha(d) \in [\delta, \gamma) \\ h(\tau_{\nu_\alpha}(d+1, \delta, \gamma)) & \text{otherwise} \end{cases}$$

With regard to the f-steps from the reductions we will be constructing, ν_α^{-1} indicates the depth at which each of these steps occurs. The first parameter of τ_{ν_α} specifies the depth of the subterm of t_{ν_α} we are currently constructing; the second and third parameter indicate the range of f-steps that will occur within this subterm.

The first clause of τ_{ν_α} defines a subterm at depth d with an f at the root in case an f-step should occur at depth d and is within the current range (i.e. $\nu_\alpha(d) \in [\delta, \gamma)$). Moreover, the range is split over the two occurrences of τ_{ν_α}: All earlier f-steps at greater depths occur in the first argument of f and all later f-steps at greater depths occur in the third argument. Note that the second argument of f is always h^ω and, hence, f is not the root of a redex unless the first argument is also equal to h^ω. Moreover, note that no redexes occur in the third argument of f due to the occurrence of ρ.

The second clause of τ_{ν_α} defines a subterm at depth d with h at the root in case no f-step from within the current (possibly empty) range occurs at d.

Every subterm $\tau_{\nu_\alpha}(d, \delta, \gamma)$, in particular $t_{\nu_\alpha} = \tau_{\nu_\alpha}(0, 0, \alpha)$, satisfies the following property.

▶ **Lemma 3.9.** *Suppose $\tau_{\nu_\alpha}(d, \delta, \gamma)$ is such that for all $d' \geq d$ it holds that either $d' \notin D$ or $\nu_\alpha(d') \notin [\delta, \gamma)$, then $\tau_{\nu_\alpha}(d, \delta, \gamma) = h^\omega$.*

Proof. By induction on $d' \geq d$, where we have for every $d' \geq d$ that the second clause from the definition of τ_{ν_α} applies. ◀

We also have:

▶ **Lemma 3.10.** *For every $\tau_{\nu_\alpha}(0, \delta, \gamma)$ with $\nu_\alpha(d) \in [\delta, \gamma)$ for some $d \in D$ there exists an $n \in \mathbb{N}$ such that $\tau_{\nu_\alpha}(0, \delta, \gamma)|_{1^n} = f(h^\omega, h^\omega, t)$ for some term t.*

Proof. Suppose not, then either $\tau_{\nu_\alpha}(0, \delta, \gamma) = h^\omega$ or eventually always the first clause from the definition of τ_{ν_α} applies. In the first case, no d exists such that $\nu_\alpha(d) \in [\delta, \gamma)$, contradiction. In the second case, there exists an infinite chain of depths $d_1 < d_2 < \ldots < d_n < \ldots$ with $[\delta, \nu_\alpha(d_1)) \supsetneq [\delta, \nu_\alpha(d_2)) \supsetneq \ldots \supsetneq [\delta, \nu_\alpha(d_n)) \supsetneq \ldots$ Hence, $\nu_\alpha(d_1) > \nu_\alpha(d_2) > \ldots > \nu_\alpha(d_n) > \ldots$ However, this is an infinite descending chain of ordinals, which cannot exist, again a contradiction. ◀

Incompressible Reductions. Having introduced all necessary ingredients, we can now prove the central theorem of this section.

▶ **Theorem 3.11.** *For every ν_α, the term t_{ν_α} has an unique maximal strongly convergent reduction starting from it which is incompressible and of length at least length α.*

Proof. We prove by transfinite induction that $t_{\nu_\alpha} = \tau_{\nu_\alpha}(0, 0, \alpha)$ reduces to $\tau_{\nu_\alpha}(0, \kappa, \alpha)$ with $\kappa \leq \alpha$ in at least κ rewrite steps. In case $\kappa = 0$, the result is immediate.

In case $\kappa = \lambda + 1$, we have by the induction hypothesis that t_{ν_α} reduces to $\tau_{\nu_\alpha}(0, \lambda, \alpha)$ in at least λ steps. Hence, it suffices to show that $\tau_{\nu_\alpha}(0, \lambda, \alpha)$ reduces to $\tau_{\nu_\alpha}(0, \kappa, \alpha)$ in at least one step. As $\lambda < \alpha$, we have $\nu_\alpha(d) \in [\lambda, \alpha)$ for some $d \in D$ and, hence, by Lemma 3.10, a redex occurs at a position 1^n for some $n \in \mathbb{N}$. In fact, by definition of τ_{ν_α}, $n = \nu_\alpha^{-1}(\lambda)$. There are now two cases to consider depending on the value of $\nu_\alpha^{-1}(\kappa)$: We either have $\nu_\alpha^{-1}(\kappa) < n$ or $\nu_\alpha^{-1}(\kappa) > n$ (equality is impossible, as ν_α is a bijection). We consider each of these cases in turn.

- In case $\nu_\alpha^{-1}(\kappa) < n$, the redex at position 1^n is defined by $\tau_{\nu_\alpha}(n, \lambda, \kappa)$. Hence, since $\kappa = \lambda + 1$, it follows by definition of τ_{ν_α} and Lemma 3.9 that $\tau_{\nu_\alpha}(n, \lambda, \kappa) = f(h^\omega, h^\omega, h^\omega)$. We have:

  $$f(h^\omega, h^\omega, h^\omega) \to h(k(h^\omega)) \to h^\omega \,.$$

  Thus, $\tau_{\nu_\alpha}(n, \lambda, \kappa)$ reduces to $\tau_{\nu_\alpha}(n, \kappa, \kappa) = h^\omega$ in two steps and the result follows by observing that $\tau_{\nu_\alpha}(0, \lambda, \alpha)$ and $\tau_{\nu_\alpha}(0, \kappa, \alpha)$ are identical except for the subterm at 1^n.
- In case $\nu_\alpha^{-1}(\kappa) > n$, the redex at position 1^n is defined by $\tau_{\nu_\alpha}(n, \lambda, \iota)$ for some $\iota > \kappa$. Hence, the subterm at position 1^n is of the form $f(h^\omega, h^\omega, \rho \circ \tau_{\nu_\alpha}(n+1, \kappa, \iota))$. We have

  $$f(h^\omega, h^\omega, \rho \circ \tau_{\nu_\alpha}(n+1, \kappa, \iota) \to h(k(\rho \circ \tau_{\nu_\alpha}(n+1, \kappa, \iota)))$$

  and by Lemma 3.7, this further reduces to $h(\tau_{\nu_\alpha}(n+1, \kappa, \iota)) = \tau_{\nu_\alpha}(n, \kappa, \iota)$. By these facts and the observation that $\tau_{\nu_\alpha}(0, \lambda, \alpha)$ and $\tau_{\nu_\alpha}(0, \kappa, \alpha)$ are identical except for the subterm at position 1^n, the result follows.

In case κ is a limit ordinal, observe that whenever λ approaches κ from below, we have that the redexes occur at increasingly greater depths. In case of the f-redexes, this follows as the position at which the λth f-redex occurs is equal to $1^{\nu_\alpha^{-1}(\lambda)}$ and as ν_α is a bijection. In case of the k-redexes, we have that only finitely many of these redexes are contracted after each f-step. Moreover, all these redexes occur at depths greater than $\nu_\alpha^{-1}(\lambda)$. Hence, as the redexes occur at increasingly greater depths, it follows by the induction hypothesis that we have a strongly convergent reduction which passes through each $\tau_{\nu_\alpha}(0, \lambda, \alpha)$ with $\lambda < \kappa$. As $\tau_{\nu_\alpha}(0, \lambda, \alpha)$ and $\tau_{\nu_\alpha}(0, \kappa, \alpha)$ are identical except for some subterm at a position 1^n, where n eventually increases as λ increases (as ν_α is a bijection), $\tau_{\nu_\alpha}(0, \kappa, \alpha)$ is the final term of the strongly convergent reduction, concluding this case.

Our theorem follows once we observe that each redex contracted in the constructed reduction (of length at least α) is created in the step immediately preceding it. ◀

We now immediately have:

▶ **Corollary 3.12.** *For every countable ordinal α, there exists a strongly convergent reduction of length* at least *α that is incompressible.*

In fact, as each prefix of each of the above reductions is also incompressible, we obtain the following by the same reasoning regarding redex creation:

▶ **Corollary 3.13.** *For every countable ordinal α, there exists a strongly convergent reduction of length* precisely *α that is incompressible.*

4 Standardisation

Having established that compression does not generalise to systems with non-left-linear rules, we will next reinterpret compression in *left-linear* systems as a degenerate case of standardisation. To be able to do this, we first need to define standard reductions.

Starting from finite rewriting, we could define a standard reduction as one that contracts redexes in leftmost-outermost order [14, Section 8.5], where a position p_1 is said to occur to the left of a position p_2 if $p_1 = q \cdot n_1 \cdot p_1'$ and $p_2 = q \cdot n_2 \cdot p_2'$ with $n_1 < n_2$ for some q (note that $p_1 \parallel p_2$). However, this notion of standardisation is ill-suited to our purposes, as infinite terms exist that do not have a leftmost redex.

▶ **Example 4.1.** Consider the term which is the unique solution of the equation $s = f(s, a)$ and suppose we have at our disposal the rewrite rule $a \to b$. The term s does not have a leftmost-outermost redex: For each $n \in \mathbb{N}$, an outermost redex occurs at position $1^n \cdot 2$. However, no redex is leftmost, as also for each $n \in \mathbb{N}$ the redex at position $1^{n+1} \cdot 2$ occurs to the left of the redex at position $1^n \cdot 2$.

Dropping the requirement that the order in standard reductions should be leftmost, we can alternatively consider *parallel standard reductions* [14, Definition 8.5.6], which allow for more freedom with regard to redexes that occur at parallel positions. The definition is as follows, where we make explicit the case distinction that implicitly occurs in [14, Definition 8.5.6]:

▶ **Definition 4.2.** Let $t_0 \twoheadrightarrow t_\alpha$ with $(p_\beta, l_\beta \to r_\beta)_{\beta<\alpha}$ the sequence of rewrite steps of $t_0 \twoheadrightarrow t_\alpha$. The reduction $t_0 \twoheadrightarrow t_\alpha$ is *parallel standard* iff for every $\beta < \alpha$ either:

- $p_\beta \parallel p_\kappa$ or $p_\beta \leq p_\kappa$ for all $\beta < \kappa < \alpha$, or
- $p_\beta = p_\kappa \cdot p_\beta'$ with $p_\beta' \in \{q \in \mathcal{P}\mathrm{os}(l_\kappa) \mid \mathrm{root}(l_\kappa|_q) \in \Sigma\}$ and $\kappa = \min\{\gamma \in (\beta, \alpha) \mid p_\beta > p_\gamma\}$.

Hence, either (a) every step after the βth step should occur parallel to or below p_β, or (b) the position p_β should occur either in the redex pattern of the first redex that occurs at a position p_κ above p_β (i.e. $p_\beta = p_\kappa \cdot q$ with q a non-variable position of l_κ). Thus, in the second case, contracting the redex in the βth step helps to create the redex pattern of the redex contracted in the κth step.

As in the finite case, parallel standard reductions are not necessarily unique: Given the rule $a \to h(a)$, we have that both $f(\underline{a}, a) \to f(h(a), \underline{a}) \to f(h(a), h(a))$ and $f(a, \underline{a}) \to f(\underline{a}, h(a)) \to f(h(a), h(a))$ are parallel standard reductions from $f(a, a)$ to $f(h(a), h(a))$. We will further discuss the issue of uniqueness in Section 4.2.

▶ Remark. Compressed reductions do not need to be parallel standard: Consider the rules $a \to h(a)$ and $f(x) \to g(x)$. The reduction $f(a) \to f(h(a)) \to g(h(a))$ is compressed, as it has length $\leq \omega$. However, it is not parallel standard, as the a-redex that is being contracted occurs below the contracted f-redex, while it is not part of the redex pattern of the f-redex.

Although Definition 4.2 alleviates the problem with leftmost redexes as exhibited in Example 4.1, it does allow for standard reductions of length greater than ω.

▶ **Example 4.3.** Consider the term $f(a, a)$ and the rewrite rule $a \to h(a)$ from above. We have the following reduction of length $\omega + \omega$:

$$\begin{aligned} f(a,a) \to f(h(a),a) \to \cdots \to f(h^n(a),a) \to f(h^{n+1}(a),a) \to \cdots f(h^\omega,a) \\ \to f(h^\omega,h(a)) \to \cdots \to f(h^\omega,h^n(a)) \to f(h^\omega,h^{n+1}(a)) \to \cdots f(h^\omega,h^\omega)\,. \end{aligned}$$

The reduction is parallel standard, as the first argument of f is parallel to its second argument.

The situation in the above example can be mitigated by strengthening Definition 4.2. One possibility with regard to the first clause of the definition is to not only consider $p_\beta \parallel p_\kappa$ and $p_\beta \leq p_\kappa$, but additionally $|p_\beta| \leq |p_\kappa|$, which when combined reduces to simply $|p_\beta| \leq |p_\kappa|$. With this restriction the reduction from the above example is no longer valid: contraction of a-redexes in the second argument is postponed even though we are contracting redexes at greater depths in the first argument.

We prefer not to introduce the above strengthening as part of the definition of parallel standard reductions as we wish to stay as close as possible to the definition from finite rewriting. Our standardisation procedure does, however, implicitly use the strengthening.

4.1 Parallel Standardisation

To show that reductions in left-linear systems can be transformed into parallel standard form, we first introduce a variant of parallel standard reduction which is limited to a certain depth.

▶ **Definition 4.4.** Let $t_0 \twoheadrightarrow t_\alpha$ with $(p_\beta, l_\beta \to r_\beta)_{\beta<\alpha}$ the sequence of rewrite steps of $t_0 \twoheadrightarrow t_\alpha$. The reduction $t_0 \twoheadrightarrow t_\alpha$ is *parallel standard up to depth* $d \in \mathbb{N}$ iff for every $\beta < \lambda$ with $\lambda = \max\{\gamma \mid |p_\gamma| < d\}$ either:

- $|p_\beta| \leq |p_\kappa|$ for all $\beta < \kappa \leq \lambda$, or
- $p_\beta = p_\kappa \cdot p'_\beta$ with $p'_\beta \in \{q \in \mathcal{P}\mathrm{os}(l_\kappa) \mid \mathrm{root}(l_\kappa|_q) \in \Sigma\}$ and $\kappa = \min\{\gamma \in (\beta, \lambda] \mid p_\beta > p_\gamma\}$.

Thus, a reduction is parallel standard up to depth d if (a) the reduction is parallel standard up to and including the last step in $t_0 \twoheadrightarrow t_\alpha$ contracting a redex at depth less than d and (b) it incorporates the depth requirement on parallel redexes as mentioned immediately below Example 4.3.

Our standardisation procedure works on a depth-by-depth basis and per depth follows the inversion procedure from [14, Section 8.5.3]. Effectively, this means that per depth, standardisation will be achieved by permuting redexes that are not parallel standard.

Permutation. Write $s \to_{\|} t$ for a reduction contracting (an infinite number of) parallel redexes; such a reduction can always be transformed into a strongly convergent reduction, e.g. by contracting the parallel redexes in order of least depth. We have the following lemma, assuming left-linearity.

▶ **Lemma 4.5** (Step Permutation). *Let $d \in \mathbb{N}$. If $s \to_{\|} t \to s'$ is such that all steps in $s \to_{\|} t$ occur at depth $> d$ and such that $t \to s'$ occurs at a position p with $|p| \geq d$, then there exists a reduction $s \to^* t' \to_{\|} s'$ which is parallel standard up to depth $d+1$ such that all steps occur at depth $\geq d$ and such that the final step of $s \to^* t'$ occurs at position p.*

Proof. As left-hand sides of rewrite rules are finite, only finitely many steps from $s \to_{\|} t$ are required to create the redex at position p contracted in $t \to s'$ and to guarantee parallel standardness; the *required* steps are those at positions q such that either $q \leq p$ or $q = p \cdot q'$ with $q'(\neq \epsilon)$ occurring in the redex pattern of the redex contracted in $t \to s'$. Define $s \to^* t''$ by first contracting all the required redexes in order of least depth and next contracting the redex at position p, which exists by contraction of the required redexes and left-linearity.

Project the remaining redexes from $s \to_{\|} t$ over $s \to^* t''$. By left-linearity and since the original set of redexes is parallel, this projection yields a set of redexes all of which are parallel and occur at depth $> d$, possibly, with exception of a unique redex occurring at position p with $|p| = d$ (in which case $t \to s'$ is collapsing). In this latter case define t' to be the result of contracting the parallel redex at position p in t'', otherwise define t' to be t''. Contracting the remaining parallel redexes yields s', as a projection argument similar to the Strip Lemma for orthogonal iTRSs [8, 7] shows.

It is easily seen that all steps in $s \to^* t' \to_{\|} s'$ occur at depth $\geq d$ and that the final step of $s \to^* t'$ occurs at position p (and depth d in case $|p| = d$). Moreover, by construction it also immediately follows that the reduction is parallel standard up to depth $d+1$. ◀

Write $s \to_{\|}^* t$ for a finite sequence of parallel steps, we can now use the previous result to permute steps in finite reductions, again assuming left-linearity.

▶ **Lemma 4.6** (Reduction Permutation). *Let $n \geq 1$ and $d \in \mathbb{N}$. If $s_0 \to^* s_n$ is such that all steps in $s_0 \to^* s_{n-1}$ occur at depth $> d$ and such that $s_{n-1} \to s_n$ contracts a redex at a position p with $|p| = d$, then there exists a reduction $s_0 \to^* t \to_{\|}^* s_n$ which is parallel standard up to depth $d+1$ such that all steps occur at depth $\geq d$ and such that the last step of $s_0 \to^* t$ is the last step that occurs at position p and depth d.*

Proof. The proof is by induction on the number of steps in $s_0 \to^* s_n$, where in addition we show that the second clause of Definition 4.4 applies to every step at depth $> d$ in $s_0 \to^* t$. In case $n = 1$, the reduction consists of a single step and the result is immediate.

In case $n > 1$, write $s_0 \to s_1 \to^* s_n$. By the induction hypothesis there exists a reduction $s_1 \to^* t' \to_{\|}^* s_n$ which is parallel standard up to $d+1$ and which satisfies the required properties. We construct $s_0 \to^* t$ by permuting a set of parallel redexes such that at any point during the permutation we have a reduction of the form $s_0 \to^* t_1 \to_{\|} t_2 \to^* t'$. Initially, define $s_0 \to^* t_1$ to be empty, $t_1 \to_{\|} t_2$ equal to $s_0 \to s_1$, and $t_2 \to^* t'$ equal to $s_1 \to^* t'$. For each permutation step, write $t_2 \to^* t'$ as $t_2 \to t_2' \to^* t'$ and apply Lemma 4.5 to $t_1 \to_{\|} t_2$ and $t_2 \to t_2'$ to obtain $t_1 \to^* t_1' \to_{\|} t_2'$. In the following permutation step assume that $s_0 \to^* t_1$ is equal to $s_0 \to^* t_1'$, that $t_1 \to_{\|} t_2$ is equal to $t_1' \to_{\|} t_2'$, and that $t_2 \to^* t'$ is equal to $t_2' \to^* t'$. Continue until $t_2 \to^* t'$ is empty and then define t to be t_1 and $t \to_{\|}^* s_n$ to be $t_1 \to_{\|} t' \to_{\|}^* s_n$.

By the permutation procedure and Lemma 4.5, we have that all steps in $s_0 \to^* t \to_{\|}^* s_n$ occur at depth $\geq d$ and that the last step of $s_0 \to^* t$ is the last step that occurs at position p and depth d. Hence, this leaves to show that $s_0 \to^* t$ is parallel standard up to depth $d+1$

such that the second clause of Definition 4.4 applies to any step at depth $> d$: Since all steps occur at depth $\geq d$, the first clause of Definition 4.4 immediately applies to any step at depth d. For any other step it follows by the construction in the proof of Lemma 4.5 and the permutation procedure that any such step either (a) occurs in the redex pattern of a step following it or (b) occurs above it. In the first case, the second clause holds immediately; in the second case, observe that there must be some later step in whose redex pattern the step under consideration occurs (thus implying the second clause), otherwise, tracing usage of redex patterns in contracted redexes, it follows that the second clause does not apply to some of the reduction steps at depth $> d$ in $s_1 \to^* t'$, contradicting the induction hypothesis. ◀

Needed Rewrite Steps. To work on a depth by depth basis in our standardisation theorem, we require a way to establish which redexes are needed for the creation of a redex at a certain depth. To this end we extend from finite rewriting [4] the notion of origins (roughly the inverse of descendants) and we define the related notion of needed steps.

▶ **Definition 4.7.** Let $s \to t$ be adorned with $(p, l \to r)$. If $q \in \mathcal{P}\mathrm{os}(t)$, then the set of *origins of q across $s \to t$*, denoted $(s \to t)\nwarrow q$, is the set of positions of s defined as follows:

- $\{q\}$ if $q \parallel p$ or $q < p$,
- $\{p \cdot q' \mid \mathrm{root}(l|_{q'}) \in \Sigma\}$ if $q = p \cdot p'$ with $\mathrm{root}(r|_{p'}) \in \Sigma$, and
- $\{p \cdot q' \cdot p'' \mid l|_{q'} = r|_{p'}\} \cup \{p \cdot q' \mid l|_{q'} \in \Sigma \text{ and } r|_\epsilon \in V\}$ if $q = p \cdot p' \cdot p''$ with $r|_{p'} \in V$.

If Q is a finite set of positions, then $(s \to t)\nwarrow Q = \bigcup_{q \in Q} (s \to t)\nwarrow q$; if $t_0 \to^* t_n$ is a finite reduction, then $(t_0 \to^* t_n)\nwarrow Q = (t_0 \to^* t_{n-1})\nwarrow((t_{n-1} \to t_n)\nwarrow Q)$. Moreover, if $t_0 \twoheadrightarrow t_\alpha$, then $(t_0 \twoheadrightarrow t_\alpha)\nwarrow Q = (t_0 \twoheadrightarrow t_\gamma)\nwarrow((t_\gamma \to^* t_\beta)\nwarrow Q)$ with γ the largest limit ordinal smaller than or equal to β (and $\gamma = 0$ if no such ordinal exists) and β such that each step in $t_\beta \twoheadrightarrow t_\alpha$ occurs at a position p with $|p| > |q|$ for all $q \in Q$.

Remark that the set of origins of a finite set Q is finite, as left-hand sides of rewrite rules are finite. Moreover, if Q is prefix-closed (i.e. $p < q \in Q$ implies $p \in Q$), then the set of origins is also prefix-closed.

▶ **Example 4.8.** Consider the reduction rules $f(x) \to h(x, x)$ and $g(x) \to x$. We have

$$(f(g(a)) \to h(g(a), g(a)) \to h(a, g(a)))\nwarrow\{\epsilon, 1, 2, 21\}$$
$$= (f(g(a)) \to h(g(a), g(a)))\nwarrow\{\epsilon, 1, 11, 2, 21\} = \{\epsilon, 1, 11\}\,.$$

We now define needed steps.

▶ **Definition 4.9.** Let $t_0 \twoheadrightarrow t_\alpha$ and let Q be a finite, prefix-closed subset of positions of t_α. A step $(p_\beta, l_\beta \to r_\beta)$ is *needed for Q* iff $p_\beta \in (t_\beta \twoheadrightarrow t_\alpha)\nwarrow Q$.

Observe by strong convergence and the definition of origins that only finitely many redexes are needed for finite, prefix-closed subsets of positions of t_α. In case of Example 4.8 both the first and second step are needed for $\{\epsilon, 1\}$; only the first step is needed for $\{\epsilon\}$ and $\{\epsilon, 2, 21\}$.

Standardisation. Employing neededness, we can now prove our standardisation theorem for left-linear iTRSs; the argument derives from the compression argument in [10]. It is, however, necessarily more complicated due to added complexity of standardisation over compression.

▶ **Theorem 4.10.** *Let $s \twoheadrightarrow t$ in a left-linear iTRS. There exists a parallel standard reduction of length at most ω from s to t.*

Proof. We prove by induction over the depth $d \in \mathbb{N}$ that there exists a reduction $s \to^* s_d \twoheadrightarrow t$ which is parallel standard up to depth d and such that all steps in $s_d \twoheadrightarrow t$ occur at depth $\geq d$. In case $d = 0$, define $s_0 = s$. The result is immediate in this case.

In case $d > 0$, it follows by the induction hypothesis that there exists a reduction $s \to^* s_{d-1} \twoheadrightarrow t$ which is parallel standard up to depth $d - 1$ such that all steps in $s_{d-1} \twoheadrightarrow t$ occur at depth $\geq d - 1$. There are now two cases to consider, either no redex in the reduction $s_{d-1} \twoheadrightarrow t$ occurs at depth d or there are such redexes. In the first case, $s \to^* s_{d-1}$ is also parallel standard up to depth d and it suffices to define $s_d = s_{d-1}$. In the remainder of the proof we consider the second case.

Consider the first redex in $s_{d-1} \twoheadrightarrow t$ which occurs at depth d and suppose that the step in which the redex is contracted is $s' \to t'$. Let Q be the smallest prefix-closed set of positions in s' which includes the positions in the redex pattern of the redex contracted in $s' \to t'$. By definition of neededness, finitely many steps from $s_{d-1} \twoheadrightarrow s'$ are needed for Q, all of which occur at depth $> d$ by definition of $s_{d-1} \twoheadrightarrow t$. Moreover, by definition of origins, these needed steps (in their original order) together with the redex contracted in $s' \to t'$ form a finite reduction from s_{d-1} to some term t''. By a projection argument, similar to the Strip Lemma for orthogonal iTRSs [8, 7], it follows that $s_{d-1} \to^* t'' \twoheadrightarrow t$, where the steps at depth d are identical to those of $s_{d-1} \twoheadrightarrow t$. Observe now that Lemma 4.6 can be applied to $s_{d-1} \to^* t''$, as all steps needed for Q occur at depth $> d$ and as $s' \to t'$ occurs at depth d. Let $s_{d-1} \to^* s'_{d-1} \twoheadrightarrow t''$ be the result of applying Lemma 4.6. By definition of $s_{d-1} \to^* s'_{d-1} \twoheadrightarrow t''$ and the observation that the second clause of Definition 4.4 applies to all steps in $s_{d-1} \to^* s'_{d-1}$ at depth $> d$ (see the proof of Lemma 4.6), it follows that $s \to^* s_{d-1} \to^* s'_{d-1}$ is parallel standard up to depth $d + 1$ and that the number of redexes at depth d in $s'_{d-1} \twoheadrightarrow t'' \twoheadrightarrow t$ is equal to the number of redexes at depth d in $t' \twoheadrightarrow t$ (i.e. one less than the number of redexes at depth d in $s' \twoheadrightarrow t$). Hence, we can now repeat the argument with $s'_{d-1} \twoheadrightarrow t$ until no redexes are left at depth d, at which point we have obtained a reduction $s \to^* s_d \twoheadrightarrow t$ which is parallel standard up to depth $d + 1$.

As the redexes considered above occur at increasingly greater depths, we have that the constructed reduction is strongly convergent. Moreover, the reduction is of length at most ω, as $s \to^* s_d$ is finite for each $d \in \mathbb{N}$. Finally, the reduction is parallel standard, as it is parallel standard up to every depth $d \in \mathbb{N}$. ◀

As the above proof applies to arbitrary reductions, the following is now immediate.

▶ **Corollary 4.11.** *There exists a standardisation procedure which transforms every reduction of every left-linear iTRS into a parallel standard reduction of length at most ω.*

Hence, standardisation can be used to obtain a compressed reduction. Turning this around and observing that the proof of the compression theorem is also based on permutation of redexes [8, 7], we can also see the compression theorem as a degenerate instance of a standardisation procedure. It is degenerate, as compression is less strict with respect to the order in which redexes are contracted. As such, we have arrived at our reinterpretation of compression.

4.2 Approximating Leftmost-Outermost Standardisation

Although it is not always possible to obtain a reduction which is standard in the sense of the leftmost-outermost order (see Example 4.1), we would still like to approximate this order to ensure that there is a unique standard reduction as in the finite case (see Theorem 8.5.51 and Lemma 8.5.52 in [14]). The natural idea in the context of infinitary rewriting, as also mentioned in [1], is to take into account depth next to the left-to-right order.

▶ **Definition 4.12.** Let $t_0 \twoheadrightarrow t_\alpha$ with $(p_\beta, l_\beta \to r_\beta)_{\beta<\alpha}$ the sequence of rewrite steps of $t_0 \twoheadrightarrow t_\alpha$. The reduction $t_0 \twoheadrightarrow t_\alpha$ is *depth leftmost standard* iff for all $\beta < \lambda < \alpha$:

- if $p_\iota < p_\beta$ for some $\beta < \iota \leq \lambda$, then $p_\beta = p_\kappa \cdot p'_\beta$ with $p'_\beta \in \{q \in \mathcal{P}\mathrm{os}(l_\kappa) \mid \mathrm{root}(l_\kappa|_q) \in \Sigma\}$ and $\kappa = \min\{\gamma \in (\beta, \iota] \mid p_\beta > p_\gamma\}$, and otherwise
- either $|p_\beta| < |p_\lambda|$ or $|p_\beta| = |p_\lambda|$ and p_β occurs to the left of or is equal to p_λ.

Roughly, the above states that the position p_β should either (a) occur in the redex pattern of the first redex that occurs at a position above it or (b) occur above or to the left of any redex that is contracted later.

▶ **Example 4.13.** Consider the term $f(a,a)$ and the rewrite rule $a \to h(a)$ from Example 4.3. The following reduction is depth standard:

$$f(a,a) \to f(h(a),a) \to f(h(a),h(a)) \to \cdots \to f(h^n(a),h^n(a)) \\ \to f(h^{n+1}(a),h^n(a)) \to f(h^{n+1}(a),h^{n+1}(a)) \to \cdots f(h^\omega,h^\omega)\,.$$

Note that the second clause of Definition 4.12 induces a total order on positions: We first order by depth and next we order from left to right. Hence, as the definition considers every initial $t_0 \twoheadrightarrow t_\lambda$, it follows for each $s \twoheadrightarrow t$ that there exists *at most* one depth leftmost standard reduction from s to t. Moreover, we have the following.

▶ **Lemma 4.14.** *Every depth leftmost standard reduction is parallel standard.*

Proof. Immediate by transfinite induction over the steps from the depth leftmost standard reduction, where the first clause of Definition 4.12 implies the second clause of Definition 4.2 and where the second clause of Definition 4.12 implies the first clause of Definition 4.2. ◀

Finally, we have:

▶ **Theorem 4.15.** *Let $s \twoheadrightarrow t$ in a left-linear iTRS. There exists a depth leftmost standard reduction of length at most ω from s to t.*

Proof (Sketch). By Theorem 4.10, there exists a parallel standard reduction from s to t. This parallel standard reduction is not necessarily depth leftmost standard, as certain redexes may violate the second clause of Definition 4.12. However, all these redexes will be parallel to each other and, hence, can be permuted without reservation (as long as we ensure that the finitely many redexes needed to create each redex are also permuted). The resulting reduction will be strongly convergent, as the second clause of Definition 4.12 only requires us to permute redexes that occur at lesser or equal depth than the redexes preceding them. ◀

5 Conclusion and Proposed Research

As shown, the obvious generalisation of the compression property — replacing ω by an arbitrary countable ordinal — does in principle not apply to iTRSs with non-left-linear rules. This result might be considered rather unexpected, as the standard counterexample to compression does satisfy such a property, albeit only for finite starting terms [5], and as does $\lambda\beta\eta$-calculus, the standard counterexample in the higher-order case [12].

Given the above observation and noting that infinitary rewriting is susceptible to a similar computational treatment as the real numbers in computable analysis [10], it seems we can no longer maintain a computational view of the compression property. Hence, we provide a reinterpretation, viz. that compression is a degenerate case of standardisation.

Given that compression and standardisation are closely related and that both have strong ties with notions related to the equivalence of reductions [8, 14], a further study of reduction equivalence in infinitary rewriting seems warranted for. Moreover, a pertinent question is whether our standardisation results can be used to simplify existing proofs in infinitary rewriting, e.g., the proof of confluence modulo hypercollapsing subterms [8, 7]. Finally, it would be interesting to investigate whether both the counterexamples to α-compression and the standardisation theorem extend to infinitary higher-order systems [11], where any counterexample to α-compression then employs a non-fully-extended rule instead of a non-left-linear rule.

Acknowledgements. The author wishes thank Jan Willem Klop, Yoshihito Toyama, and Jaco van de Pol for their continuing support. Moreover, acknowledgements go to the reviewers for their constructive feedback and to Jakob Grue Simonsen for encouraging the author to complete the paper. Finally, the author is indebted to Alastair Donaldson and the Department of Computing at Imperial College London for their financial support.

References

1 Søren Bjerg Andersen and Jakob Grue Simonsen. Term rewriting systems as topological dynamical system. In *Proc. 23rd Int'l Conf. on Rewriting Techniques and Applications (RTA 2012)*, volume 15 of *Leibniz International Proc. in Informatics*, pages 53–68, 2012.

2 André Arnold and Maurice Nivat. The metric space of infinite trees. Algebraic and topological properties. *Fundamenta Informaticae*, 3(4):445–476, 1980.

3 Michael Barr. Terminal coalgebras in well-founded set theory. *Theoretical Computer Science*, 114(2):299–315, 1993.

4 Inge Bethke, Jan Willem Klop, and Roel de Vrijer. Descendants and origins in term rewriting. *Information and Computation*, 159(1-2):59–124, 2000.

5 Nachum Dershowitz, Stéphane Kaplan, and David A. Plaisted. Rewrite, rewrite, rewrite, rewrite, rewrite, ... *Theoretical Computer Science*, 83(1):71–96, 1991.

6 Stefan Kahrs. Infinitary rewriting: Foundations revisited. In *Proc. 21st Int'l Conf. on Rewriting Techniques and Applications (RTA 2010)*, volume 6 of *Leibniz International Proceedings in Informatics*, pages 161–176, 2010.

7 Richard Kennaway and Fer-Jan de Vries. Infinitary rewriting. In Terese [13], Chapter 12, pages 668–711.

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

9 Jeroen Ketema. *Böhm-Like Trees for Rewriting*. PhD thesis, Vrije Universiteit, Amsterdam, 2006.

10 Jeroen Ketema and Jakob Grue Simonsen. Computing with infinite terms and infinite reductions. Unpublished manuscript.

11 Jeroen Ketema and Jakob Grue Simonsen. Infinitary Combinatory Reduction Systems. *Information and Computation*, 209(6):893–926, 2011.

12 Paula Severi and Fer-Jan de Vries. An extensional Böhm model. In *Proc. 13th Int'l Conf. on Rewriting Techniques and Applications (RTA 2002)*, volume 2378 of *Lecture Notes in Computer Science*, pages 159–173, 2002.

13 Terese, editor. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2003.

14 Vincent van Oostrom and Roel de Vrijer. Equivalence of reductions. In Terese [13], Chapter 8, pages 301–474.

15 Klaus Weihrauch. *Computable Analysis: An Introduction*. Springer-Verlag, 2000.