

ON CONSTRUCTOR REWRITE SYSTEMS AND THE LAMBDA-CALCULUS*

UGO DAL LAGO AND SIMONE MARTINI

Università di Bologna, and INRIA Sophia Antipolis
e-mail address: {dallago, martini}@cs.unibo.it

ABSTRACT. We prove that orthogonal constructor term rewrite systems and λ -calculus with weak (i.e., no reduction is allowed under the scope of a λ -abstraction) call-by-value reduction can simulate each other with a linear overhead. In particular, weak call-by-value beta-reduction can be simulated by an orthogonal constructor term rewrite system in the same number of reduction steps. Conversely, each reduction in a term rewrite system can be simulated by a constant number of beta-reduction steps. This is relevant to implicit computational complexity, because the number of beta steps to normal form is polynomially related to the actual cost (that is, as performed on a Turing machine) of normalization, under weak call-by-value reduction. Orthogonal constructor term rewrite systems and λ -calculus are thus both polynomially related to Turing machines, taking as notion of cost their natural parameters.

INTRODUCTION

Implicit computational complexity is a young research area, whose main aim is the description of complexity phenomena based on language restrictions, and not on external measure conditions or on explicit machine models. It borrows techniques and results from mathematical logic (model theory, recursion theory, and proof theory) and in doing so it has allowed the incorporation of aspects of computational complexity into areas such as formal methods in software development and programming language design. The most developed area of implicit computational complexity is probably the model theoretic one — finite model theory being a very successful way to describe complexity classes. In the design of programming language tools (e.g., type systems), however, syntactical techniques prove more useful. In the last years we have seen much work restricting recursive schemata and developing general proof theoretical techniques to enforce resource bounds on programs. Important achievements have been the characterizations of several complexity classes by means of limitations of recursive definitions (e.g., [BC92, Lei95]) and, more recently, by using the “*light*” fragments of linear logic [Gir98]. Moreover, rewriting techniques such as recursive path orderings and the interpretation method have been proved useful in the field [MM00]. By borrowing the terminology from software design technology, we may dub

1998 ACM Subject Classification: F.4.1.

Key words and phrases: lambda calculus, term rewriting, implicit computational complexity.

* This paper is an extended version of [DLM09], appeared in the proceedings of ICALP 2009.

this area as implicit computational complexity *in the large*, aiming at a broad, global view on complexity classes. We may have also an implicit computational complexity *in the small* — using logic to study single machine-free models of computation. Indeed, many models of computations do not come with a natural cost model — a definition of cost which is both intrinsically rooted in the model of computation, and, at the same time, it is polynomially related to the cost of implementing that model of computation on a standard Turing machine. The main example is the λ -calculus: the most natural intrinsic parameter of a computation is its number of beta-reductions, but this very parameter bears no relation, in general, with the actual cost of performing that computation, since a beta-reduction may involve the duplication of arbitrarily big subterms¹. What we call implicit computational complexity in the small, therefore, gives complexity significance to notions and results for computation models where such natural cost measures do not exist, or are not obvious. In particular, it looks for cost-explicit simulations between such computational models.

The present paper applies this viewpoint to the relation between λ -calculus and orthogonal constructor term rewrite systems (OCRSSs in the following). We will prove that these two computational models simulate each other with a linear overhead. That each OCRSS could be simulated by λ -terms and beta-reduction is well known, in view of the availability, in λ -calculus, of fixed-point operators, which may be used to solve the mutual recursion expressed by first-order rewrite rules. Here (Section 3) we make explicit the complexity content of this simulation, by showing that any first-order rewriting of n steps can be simulated by kn beta steps, where k depends on the specific rewrite system but *not* on the size of the involved terms. Crucial to this result is the encoding of constructor terms using Scott’s schema for numerals [Wad80]. Indeed, Parigot [Par90] (see also [PR93]) shows that in the pure λ -calculus Church numerals do not admit a predecessor working in a constant number of beta steps. Moreover, Splawski and Urzyczyn [SU99] show that it is unlikely that our encoding could work in the typed context of System F.

Section 2 studies the converse – the simulation of (weak) λ -calculus reduction by means of OCRSSs. We give an encoding of λ -terms into a (first-order) constructor term rewrite system. We write $[\cdot]_{\Phi}$ for the map returning a first-order term, given a λ -term; $[M]_{\Phi}$ is, in a sense, a complete defunctionalization of the λ -term M , where any λ -abstraction is represented by an atomic constructor. This is similar, although not technically the same, to the use of supercombinators (e.g., [Jon87]). We show that λ -reduction is simulated step by step by first-order rewriting (Theorem 2.9).

As a consequence, taking the number of beta steps as a cost model for weak λ -calculus is equivalent (up to a linear function) to taking the number of rewritings in OCRSSs systems. This is relevant to implicit computational complexity “in the small”, because the number of beta steps to normal form is polynomially related to the actual cost (that is, as performed on a Turing machine) of normalization, under weak call-by-value reduction. This has been established by Sands, Gustavsson, and Moran [SGM02], by a fine analysis of a λ -calculus implementation based on a stack machine. OCRSSs and λ -calculus are thus both *reasonable* machines (see the “invariance thesis” in [vEB90]), taking as notion of cost their natural, intrinsic parameters.

¹ In full beta-reduction, the size of the duplicated term is indeed arbitrary and does not depend on the size of the original term the reduction started from. The situation is much different with weak reduction, as we will see.

As a byproduct, in Section 4 we sketch a different proof of the cited result in [SGM02]. Instead of using a stack machine, we show how we could implement constructor term rewriting via term graph rewriting. In term graph rewriting we avoid the explicit duplication and substitution inherent to rewriting (and thus also to beta-reduction) and, moreover, we exploit the possible sharing of subterms. A more in-depth study of the complexity of (constructor) graph rewriting and its relations with (constructor) term rewriting can be found in another paper by the authors [DLM10].

In Section 5, we show how to obtain the same results of the previous sections when call-by-name replaces call-by-value as the underlying strategy in the λ -calculus.

1. PRELIMINARIES

The language we study is the pure untyped λ -calculus endowed with weak (that is, we never reduce under an abstraction) call-by-value reduction.

Definition 1.1. The following definitions are standard:

- *Terms* are defined as follows:

$$M ::= x \mid \lambda x.M \mid MM,$$

where x ranges a denumerable set Υ . Λ denotes the set of all λ -terms. We assume the existence of a fixed, total, order on Υ ; this way $\text{FV}(M)$ will be a sequence (without repetitions) of variables, not a set. A term M is said to be *closed* if $\text{FV}(M) = \varepsilon$, where ε is the empty sequence.

- Values are defined as follows:

$$V ::= x \mid \lambda x.M.$$

- Weak call-by-value reduction is denoted by \rightarrow_v and is obtained by closing call-by-value reduction under any applicative context:

$$\frac{}{(\lambda x.M)V \rightarrow_v M\{V/x\}} \quad \frac{M \rightarrow_v N}{ML \rightarrow_v NL} \quad \frac{M \rightarrow_v N}{LM \rightarrow_v LN}$$

Here M, N, L range over terms, while V ranges over values.

- The length $|M|$ of M is defined as follows, by induction on M : $|x| = 1$, $|\lambda x.M| = |M| + 1$ and $|MN| = |M| + |N| + 1$.

Weak call-by-value reduction enjoys many nice properties. In particular, the one-step diamond property holds and, as a consequence, the number of beta steps to normal form (if any) is invariant on the reduction order [DLM08] (this justifies the way we defined reduction, which is slightly more general than Plotkin's one [Pl75]). It is then meaningful to define $\text{Time}_v(M)$ as *the* number of beta steps to normal form (or ω if such a normal form does not exist). This cost model will be referred to as the *unitary* cost model, since each beta (weak call-by-value) reduction step counts for 1 in the global cost of normalization. Moreover, notice that α -conversion is not needed during reduction of closed terms: if $M \rightarrow_v N$ and M is closed, then the reduced redex will be in the form $(\lambda x.L)V$, where V is a *closed* value. As a consequence, arguments are always closed and open variables cannot be captured. Suppose M has n free variables $x_1 \leq \dots \leq x_n$, and that N_1, \dots, N_n are lambda-terms. The term $M\{N_1/x_1, \dots, N_n/x_n\}$ is sometimes denoted simply with $M(N_1, \dots, N_n)$, taking advantage of the implicit order between the variables.

The following lemma gives us a generalization of the fixed-point (call-by-value) combinator (but observe the explicit limit k on the reduction length, in the spirit of implicit computational complexity in the small):

Lemma 1.2 (call-by-value fixpoint combinator). *For every natural number n , there are terms H_1, \dots, H_n and a natural number m such that for any sequence of values V_1, \dots, V_n and for any $1 \leq i \leq n$:*

$$H_i V_1 \dots V_n \rightarrow_v^k V_i (\lambda x. H_1 V_1 \dots V_n x) \dots (\lambda x. H_n V_1 \dots V_n x),$$

where $k \leq m$.

Proof. The terms we are looking for are simply the following:

$$H_i \equiv M_i M_1 \dots M_n$$

where, for every $1 \leq j \leq n$,

$$M_j \equiv \lambda x_1 \dots \lambda x_n. \lambda y_1 \dots y_n. y_j (\lambda z. x_1 x_1 \dots x_n y_1 \dots y_n z) \dots (\lambda z. x_n x_1 \dots x_n y_1 \dots y_n z).$$

The natural number m is simply $2n$. \square

We only consider orthogonal and constructor rewriting in this paper. A *constructor term rewrite system* is a pair $\Xi = (\Sigma_\Xi, \mathcal{R}_\Xi)$ where:

- Symbols in the signature Σ_Ξ can be either *constructors* or *function symbols*, each with its arity.
 - Terms in $\mathcal{C}(\Xi)$ are those built from constructors and are called *constructor terms*.
 - Terms in $\mathcal{P}(\Xi, \Upsilon)$ are those built from constructors and variables and are called *patterns*.
 - Terms in $\mathcal{T}(\Xi)$ are those built from constructor and function symbols and are called *closed terms*.
 - Terms in $\mathcal{V}(\Xi, \Upsilon)$ are those built from constructors, functions symbols and variables in Υ and are dubbed *terms*.
- Rules in \mathcal{R}_Ξ are in the form $\mathbf{f}(s_1, \dots, s_n) \rightarrow_\Xi t$ where \mathbf{f} is a function symbol, $s_1, \dots, s_n \in \mathcal{P}(\Xi, \Upsilon)$ and $t \in \mathcal{V}(\Xi, \Upsilon)$. We here consider orthogonal rewrite systems only, i.e. we assume that no distinct two rules in \mathcal{R}_Ξ are overlapping and that every variable appears at most once in the lhs of any rule in \mathcal{R}_Ξ . Moreover, we assume that reduction is call-by-value, i.e. the substitution triggering any reduction must assign *constructor* terms to variables. This restriction is anyway natural in constructor rewriting.

For any term t in a OCRS, $|t|$ denotes the number of symbol occurrences, while $|t|_{\mathbf{f}}$ denotes the number of occurrences of the symbol \mathbf{f} in t . Similarly to λ -terms, if t contains instances of n variables $x_1 \leq \dots \leq x_n$, the term $t\{u_1/x_1, \dots, u_n/x_n\}$ is sometimes denoted simply with $t(u_1, \dots, u_n)$.

2. FROM λ -CALCULUS TO CONSTRUCTOR TERM REWRITING

In this section, we will prove that the λ -calculus, in the form introduced in Section 1, can be seen as a OCRS. This result will be spelled out as follows.

- An OCRS Φ on a signature Σ_Φ will be defined, together with two maps $[\cdot]_\Phi : \Lambda \rightarrow \mathcal{V}(\Sigma_\Phi, \Upsilon)$ and $\langle \cdot \rangle_\Lambda : \mathcal{V}(\Sigma_\Phi, \Upsilon) \rightarrow \Lambda$. These two maps are *not* bijections. However, $[\cdot]_\Phi$ is injective, and $\langle \cdot \rangle_\Lambda \circ [\cdot]_\Phi$ is the identity.
- The concept of canonicity for terms in $\mathcal{V}(\Phi, \Upsilon)$ will be defined. Moreover, the set of canonical terms will be shown to include $[\Lambda]_\Phi$ and to be closed by reduction.

- Reduction of canonical terms will be shown to simulate weak call-by-value reduction on λ -terms, via $\langle \cdot \rangle_\Lambda$. Conversely, the dynamics of λ -terms is proved to simulate rewriting of constructor terms again through $\langle \cdot \rangle_\Lambda$.

Altogether, the three ingredients above implies that Φ is a sound and complete way of implementing call-by-value β -reduction.

Let us start by defining Φ and the two functions allowing to translate terms Φ into λ -terms and, conversely, λ -terms back into terms of Φ . Canonicity can already be defined.

Definition 2.1 (The OCRS Φ , Canonicity). The OCRS Φ is defined as a set of rules \mathcal{R}_Φ over an infinite signature Σ_Φ . In particular:

- The signature Σ_Φ includes the binary function symbol **app** and constructor symbols $\mathbf{c}_{x,M}$ for every $M \in \Lambda$ and every $x \in \Upsilon$. The arity of $\mathbf{c}_{x,M}$ is the length of $\mathbf{FV}(\lambda x.M)$. To every term $M \in \Lambda$ we can associate a term $[M]_\Phi \in \mathcal{V}(\Phi, \Upsilon)$ as follows:

$$\begin{aligned} [x]_\Phi &= x; \\ [\lambda x.M]_\Phi &= \mathbf{c}_{x,M}(x_1, \dots, x_n), \text{ where } \mathbf{FV}(\lambda x.M) = x_1, \dots, x_n; \\ [MN]_\Phi &= \mathbf{app}([M]_\Phi, [N]_\Phi). \end{aligned}$$

Observe that if M is closed, then $[M]_\Phi \in \mathcal{T}(\Phi)$.

- The rewrite rules in \mathcal{R}_Φ are all the rules in the following form:

$$\mathbf{app}(\mathbf{c}_{x,M}(x_1, \dots, x_n), x) \rightarrow [M]_\Phi,$$

where $\mathbf{FV}(\lambda x.M) = x_1, \dots, x_n$.

- To every term $t \in \mathcal{V}(\Phi, \Upsilon)$ we can associate a term $\langle t \rangle_\Lambda \in \Lambda$ as follows:

$$\begin{aligned} \langle x \rangle_\Lambda &= x \\ \langle \mathbf{app}(u, v) \rangle_\Lambda &= \langle u \rangle_\Lambda \langle v \rangle_\Lambda \\ \langle \mathbf{c}_{x,M}(t_1, \dots, t_n) \rangle_\Lambda &= (\lambda x.M)\{\langle t_1 \rangle_\Lambda/x_1, \dots, \langle t_n \rangle_\Lambda/x_n\} \end{aligned}$$

where $\mathbf{FV}(\lambda x.M) = x_1, \dots, x_n$.

- A term $t \in \mathcal{T}(\Phi)$ is *canonical* if either $t \in \mathcal{C}(\Phi)$ or $t = \mathbf{app}(u, v)$ where u and v are themselves canonical.

Notice that the signature Σ_Φ contains an infinite number of constructors.

Example 2.2. Consider the λ -term $M = (\lambda x.xx)(\lambda y.yy)$. $[M]_\Phi$ is $t \equiv \mathbf{app}(\mathbf{c}_{x,xx}, \mathbf{c}_{y,yy})$. Moreover, $t \rightarrow \mathbf{app}(\mathbf{c}_{y,yy}, \mathbf{c}_{y,yy}) \equiv u$, as expected. We have $u \rightarrow u$. Both t and u are canonical. Finally, $\langle u \rangle_\Lambda = (\lambda y.yy)(\lambda y.yy)$.

The map $[\cdot]_\Phi$ is injective, but not surjective. However:

Lemma 2.3. For every λ -term $M \in \Lambda$, $\langle [M]_\Phi \rangle_\Lambda = M$.

Proof. By induction on M :

- If $M = x$, then

$$\langle [M]_\Phi \rangle_\Lambda = \langle [x]_\Phi \rangle_\Lambda = \langle x \rangle_\Lambda = x.$$

- If $M = NL$, then

$$\langle [M]_\Phi \rangle_\Lambda = \langle \mathbf{app}([N]_\Phi, [L]_\Phi) \rangle_\Lambda = \langle [N]_\Phi \rangle_\Lambda \langle [L]_\Phi \rangle_\Lambda = NL.$$

- If $M = \lambda y.N$, then

$$\langle [M]_\Phi \rangle_\Lambda = \langle \mathbf{c}_{y,N}(x_1, \dots, x_n) \rangle_\Lambda = (\lambda y.N)\{x_1/x_1, \dots, x_n/x_n\} = \lambda y.N = M.$$

This concludes the proof. \square

Canonicity holds for terms in Φ obtained as images of (closed) λ -terms via $[\cdot]_\Phi$. Moreover, canonicity is preserved by reduction in Φ :

Lemma 2.4. *For every closed $M \in \Lambda$, $[M]_\Phi$ is canonical. Moreover, if t is canonical and $t \rightarrow u$, then u is canonical.*

Proof. $[M]_\Phi$ is canonical for any $M \in \Lambda$ by induction on the structure of M (which, by hypothesis, is either an abstraction or an application NL where both N and L are closed). We can further prove that $v = [M]_\Phi\{t_1/x_1, \dots, t_n/x_n\}$ is canonical whenever $t_1, \dots, t_n \in \mathcal{C}(\Phi)$ and x_1, \dots, x_n includes all the variables in $\mathbf{FV}(M)$:

- If $M = x_i$, then $v = t_i$, which is clearly canonical.
- If $M = NL$, then

$$\begin{aligned} v &= [NL]_\Phi\{t_1/x_1, \dots, t_n/x_n\} \\ &= \mathbf{app}([N]_\Phi\{t_1/x_1, \dots, t_n/x_n\}, [L]_\Phi\{t_1/x_1, \dots, t_n/x_n\}) \end{aligned}$$

which is canonical, by IH.

- If $M = \lambda y.N$, then

$$\begin{aligned} v &= [\lambda y.N]_\Phi\{t_1/x_1, \dots, t_n/x_n\} \\ &= \mathbf{c}_{y,N}(x_{i_1}, \dots, x_{i_m})\{t_1/x_1, \dots, t_n/x_n\} \\ &= \mathbf{c}_{y,N}(t_{i_1}, \dots, t_{i_m}) \end{aligned}$$

which is canonical, because each t_i (and hence also v) is in $\mathcal{C}(\Phi)$.

This implies the rhs of any instance of a rule in \mathcal{R}_Φ is canonical. As a consequence, u is canonical whenever $t \rightarrow u$ and t is canonical. This concludes the proof. \square

For canonical terms, being a normal form is equivalent to being mapped to a normal form via $\langle \cdot \rangle_\Lambda$. This is not true, in general: take as a counterexample $\mathbf{c}_{x,y}(\mathbf{app}(\mathbf{c}_{z,z}, \mathbf{c}_{z,z}))$, which corresponds to $\lambda x.(\lambda z.z)(\lambda z.z)$ via $\langle \cdot \rangle_\Lambda$.

Lemma 2.5. *A canonical term t is a normal form iff $\langle t \rangle_\Lambda$ is a normal form.*

Proof. If a canonical t is a normal form, then t does not contain the function symbol \mathbf{app} and, as a consequence, $\langle t \rangle_\Lambda$ is an abstraction, which is always a normal form. Conversely, if $\langle t \rangle_\Lambda$ is a normal form, then t is not in the form $\mathbf{app}(u, v)$, because otherwise $\langle t \rangle_\Lambda$ will be a (closed) application, which cannot be a normal form. But since t is canonical, $t \in \mathcal{C}(\Phi)$, which only contains terms in normal form. \square

The following substitution lemma will be useful later.

Lemma 2.6 (Substitution). *For every term $t \in \mathcal{V}(\Phi, \Upsilon)$ and every $t_1, \dots, t_n \in \mathcal{C}(\Phi)$,*

$$\langle t\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda = \langle t \rangle_\Lambda\{\langle t_1 \rangle_\Lambda/x_1, \dots, \langle t_n \rangle_\Lambda/x_n\}$$

whenever x_1, \dots, x_n includes all the variables in t .

Proof. By induction on t :

- If $t = x_i$, then

$$\begin{aligned} \langle t\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda &= \langle x_i\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda \\ &= \langle x_i \rangle_\Lambda \\ &= x_i\{\langle t_1 \rangle_\Lambda/x_1, \dots, \langle t_n \rangle_\Lambda/x_n\} \\ &= t\{\langle t_1 \rangle_\Lambda/x_1, \dots, \langle t_n \rangle_\Lambda/x_n\}. \end{aligned}$$

- If $t = \mathbf{app}(u, v)$, then

$$\begin{aligned}
\langle t\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda &= \langle \mathbf{app}(u, v)\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda \\
&= \langle \mathbf{app}(u\{t_1/x_1, \dots, t_n/x_n\}, v\{t_1/x_1, \dots, t_n/x_n\}) \rangle_\Lambda \\
&= \langle u\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda \langle v\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda \\
&= \langle u \rangle_\Lambda \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \} \langle v \rangle_\Lambda \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \} \\
&= \langle u \rangle_\Lambda \langle v \rangle_\Lambda \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \} \\
&= \langle \mathbf{app}(u, v) \rangle_\Lambda \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \} \\
&= \langle t \rangle_\Lambda \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \}.
\end{aligned}$$

- If $t = \mathbf{c}_{y,N}(u_1, \dots, u_m)$, then

$$\begin{aligned}
\langle t\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda &= \langle \mathbf{c}_{y,N}(u_1, \dots, u_m)\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda \\
&= \langle \mathbf{c}_{y,N}(u_1\{t_1/x_1, \dots, t_n/x_n\}, \dots, u_m\{t_1/x_1, \dots, t_n/x_n\}) \rangle_\Lambda \\
&= (\lambda y.N)\{ \langle u_1\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda / x_{i_1} \\
&\quad \dots, \\
&\quad \langle u_m\{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda / x_{i_m} \} \\
&= (\lambda y.N)\{ \langle u_1 \rangle_\Lambda \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \} / x_{i_1} \\
&\quad \dots, \\
&\quad \langle u_m \rangle_\Lambda \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \} / x_{i_m} \} \\
&= ((\lambda y.N)\{ \langle u_1 \rangle_\Lambda / x_1, \dots, \langle u_m \rangle_\Lambda / x_{i_m} \}) \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \} \\
&= \langle \mathbf{c}_{y,N}(u_1, \dots, u_m) \rangle_\Lambda \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \} \\
&= \langle t \rangle_\Lambda \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \}.
\end{aligned}$$

This concludes the proof. \square

Two of the previous lemmas imply that if $M \in \Lambda$, $t_1, \dots, t_n \in \mathcal{C}(\Phi)$ and x_1, \dots, x_n includes all the variables in $\mathbf{FV}(M)$, then:

$$\langle [M]_\Phi \{t_1/x_1, \dots, t_n/x_n\} \rangle_\Lambda = M \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \}. \quad (2.1)$$

Reduction in Φ can be simulated by reduction in the λ -calculus, provided the starting term is canonical.

Lemma 2.7. *If t is canonical and $t \rightarrow u$, then $\langle t \rangle_\Lambda \rightarrow_v \langle u \rangle_\Lambda$.*

Proof. Consider the (instance of the) rewrite rule which turns t into u . Let it be

$$\mathbf{app}(\mathbf{c}_{y,M}(t_1, \dots, t_n), v) \rightarrow [M]_\Phi \{t_1/x_1, \dots, t_n/x_n, v/y\}.$$

Clearly,

$$\langle \mathbf{app}(\mathbf{c}_{y,M}(t_1, \dots, t_n), v) \rangle_\Lambda = ((\lambda y.M)\{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \}) \langle v \rangle_\Lambda$$

while, by (2.1):

$$\langle [M]_\Phi \{t_1/x_1, \dots, t_n/x_n, v/y\} \rangle_\Lambda = M \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n, \langle v \rangle_\Lambda / y \}$$

which implies the thesis. \square

Conversely, call-by-value reduction in the λ -calculus can be simulated in Φ :

Lemma 2.8. *If $M \rightarrow_v N$, t is canonical and $\langle t \rangle_\Lambda = M$, then $t \rightarrow u$, where $\langle u \rangle_\Lambda = N$.*

Proof. Let $(\lambda x.L)V$ be the redex fired in M when rewriting it to N . There must be a corresponding subterm v of t such that $\langle v \rangle_\Lambda = (\lambda x.L)V$. Then

$$v = \mathbf{app}(\mathbf{c}_{x,P}(t_1, \dots, t_n), w),$$

where $\langle \mathbf{c}_{x,P}(t_1, \dots, t_n) \rangle_\Lambda = \lambda x.L$. and $\langle w \rangle_\Lambda = V$. Observe that, by definition,

$$\langle \mathbf{c}_{x,P}(t_1, \dots, t_n) \rangle_\Lambda = (\lambda x.P)\{\langle t_1 \rangle_\Lambda/x_1, \dots, \langle t_n \rangle_\Lambda/x_n\}$$

where $\mathbf{FV}(P) = x_1, \dots, x_n$. Since t is canonical, $t_1, \dots, t_n \in \mathcal{C}(\Phi)$. Moreover, since V is a value, w itself is in $\mathcal{C}(\Phi)$. This implies

$$\mathbf{app}(\mathbf{c}_{x,P}(t_1, \dots, t_n), w) \rightarrow [P]_\Phi\{t_1/x_1, \dots, t_n/x_n, w/x\}.$$

By (2.1):

$$\begin{aligned} \langle [P]_\Phi\{t_1/x_1, \dots, t_n/x_n, w/x\} \rangle_\Lambda &= P\{\langle t_1 \rangle_\Lambda/x_1, \dots, \langle t_n \rangle_\Lambda/x_n, \langle w \rangle_\Lambda/x\} \\ &= (P\{\langle t_1 \rangle_\Lambda/x_1, \dots, \langle t_n \rangle_\Lambda/x_n\})\{\langle w \rangle_\Lambda/x\} \\ &= (\lambda x.L)\{V/x\}. \end{aligned}$$

This concludes the proof. \square

The previous lemmas together imply the following theorem, by which λ -calculus normalization can be mimicked (step-by-step) by reduction in Φ :

Theorem 2.9 (Term Reducibility). *Let $M \in \Lambda$ be a closed term. The following two conditions are equivalent:*

1. $M \rightarrow_v^n N$ where N is in normal form;
2. $[M]_\Phi \rightarrow^n t$ where $\langle t \rangle_\Lambda = N$ and t is in normal form.

Proof. Suppose $M \rightarrow_v^n N$, where N is in normal form. Then, by applying Lemma 2.8, we obtain a term t such that $[M]_\Phi \rightarrow^n t$ and $\langle t \rangle_\Lambda = N$. By Lemma 2.4, t is canonical and, by Lemma 2.5, it is in normal form. Now, suppose $[M]_\Phi \rightarrow^n t$ where $\langle t \rangle_\Lambda = N$ and t is in normal form. By applying n times Lemma 2.7, we obtain $\langle [M]_\Phi \rangle_\Lambda \rightarrow_v^n \langle t \rangle_\Lambda = N$. But $\langle [M]_\Phi \rangle_\Lambda = M$ by Lemma 2.3 and N is a normal form by Lemma 2.5, since $[M]_\Phi$ and t are canonical by Lemma 2.4. \square

There is another nice property of Φ , that will be crucial in proving the main result of this paper:

Proposition 1 (Subterm Property). *For every $M \in \Lambda$, for every t with $[M]_\Phi \rightarrow^* t$ and for every occurrence of a constructor $\mathbf{c}_{x,N}$ in t , N is a subterm of M .*

Proof. Assume $[M]_\Phi \rightarrow^n t$ and proceed by induction on n . \square

Example 2.10. Let us consider the λ -term $M = (\lambda x.(\lambda y.x)x)(\lambda z.z)$. Notice that

$$M \rightarrow_v (\lambda y.(\lambda z.z))(\lambda z.z) \rightarrow_v \lambda z.z.$$

Clearly $[M]_\Phi = \mathbf{app}(\mathbf{c}_{x,(\lambda y.x)x}, \mathbf{c}_{z,z})$. Moreover:

$$\mathbf{app}(\mathbf{c}_{x,(\lambda y.x)x}, \mathbf{c}_{z,z}) \rightarrow \mathbf{app}(\mathbf{c}_{y,x}(\mathbf{c}_{z,z}), \mathbf{c}_{z,z}) \rightarrow \mathbf{c}_{z,z}.$$

For every constructor $\mathbf{c}_{w,N}$ occurring in any term in the previous reduction sequence, N is a subterm of M .

A remark on Φ is now in order. Φ is an infinite OCRS, since Σ_Φ contains an infinite amount of constructor symbols and, moreover, there are infinitely many rules in \mathcal{R}_Φ . As a consequence, what we have presented here is an embedding of the (weak, call-by-value) λ -calculus into an infinite OCRS. Consider, now, the following scenario: suppose the λ -calculus is used to write a *program* M , and suppose that inputs to M form an infinite set of λ -terms Θ which can anyway be represented by a finite set of constructors in Φ . In this scenario, Proposition 1 allows to conclude the existence of finite subsets of Σ_Φ and \mathcal{R}_Φ such that *every* MN (where $N \in \Theta$) can be reduced via Φ by using only constructors and rules in those *finite* subsets. As a consequence, we can see the above schema as one that puts any program M in correspondence to a *finite* OCRS. Finally, observe that assuming *data* to be representable by a finite number of constructors in Φ is reasonable. Scott's scheme [Wad80], for example, allows to represent any term in a given free algebra in a finitary way, e.g. the natural number 0 becomes $[0] \equiv \mathbf{c}_{y,\lambda z.z}$ while $n + 1$ becomes $[n + 1] \equiv \mathbf{c}_{y,\lambda z.yx}([n])$. Church's scheme, on the other hand, does not have this property.

2.1. An Example. Consider the lambda terms $M = \lambda x.\lambda y.xyxy$ and $N = \lambda x.\lambda y.yxy$. It is easy to verify that:

$$\begin{aligned} L &\equiv (MN)M \rightarrow_v^2 (NM)N, \\ P &\equiv (NM)N \rightarrow_v^2 (MN)M. \end{aligned}$$

Therefore, both L and P diverge. Now:

$$\begin{aligned} [L]_\Phi &\equiv \mathbf{app}(\mathbf{app}(\mathbf{c}_{x,\lambda y.xyxy}, \mathbf{c}_{x,\lambda y.yxy}), \mathbf{c}_{x,\lambda y.xyxy}) \\ &\rightarrow t \equiv \mathbf{app}(\mathbf{c}_{y,xyx}(\mathbf{c}_{x,\lambda y.yxy}), \mathbf{c}_{x,\lambda y.xyxy}) \\ &\rightarrow \mathbf{app}(\mathbf{app}(\mathbf{c}_{x,\lambda y.yxy}, \mathbf{c}_{x,\lambda y.xyxy}), \mathbf{c}_{x,\lambda y.yxy}) \\ &\equiv [P]_\Phi. \end{aligned}$$

Similarly, $[P]_\Phi \rightarrow^2 [L]_\Phi$. Observe that along the computation we reach the term t , which is not the image of any λ -term. However, all constructor terms in the reduction are canonical and, moreover, $\langle t \rangle_\Lambda$ is $(\lambda y.NyN)M$, the lambda term found along the reduction from L to P .

3. FROM CONSTRUCTOR TERM REWRITING TO THE λ -CALCULUS

In this section, we will show that one rewriting step of any constructor rewrite system can be simulated by a fixed number of weak call-by-value beta-reductions. As an easy consequence, λ -calculus will be shown to efficiently simulate any OCRS. During this section we will assume fixed an OCRS Ξ over a finite signature Σ_Ξ . Let $\mathbf{c}_1, \dots, \mathbf{c}_g$ be the constructors of Ξ and let

$\mathbf{f}_1, \dots, \mathbf{f}_h$ be the function symbols of Ξ . We will describe several constructions, which work independently of Ξ (they only depends on the arity of the symbols).

- A map $\langle\langle \cdot \rangle\rangle_\Lambda : \mathcal{C}(\Xi) \rightarrow \Lambda$ can be defined by recursion on the structure of the input. The map can be extended to constructors of Ξ (which are *not* terms by themselves), in such a way that for every \mathbf{c}_i , the lambda term $\langle\langle \mathbf{c}_i \rangle\rangle_\Lambda$ “computes” $\langle\langle \mathbf{c}_i(t_1 \dots t_{ar(\mathbf{c}_i)}) \rangle\rangle_\Lambda$ when fed with $\langle\langle t_1 \rangle\rangle_\Lambda \dots \langle\langle t_{ar(\mathbf{c}_i)} \rangle\rangle_\Lambda$, for any $t_1, \dots, t_{ar(\mathbf{c}_i)} \in \mathcal{C}(\Xi)$. (See Definition 3.1.)
- Defining a map analogous to $\langle\langle \cdot \rangle\rangle_\Lambda$, but acting on closed terms (and not only on constructor terms) is more delicate. Indeed, a term $\mathbf{f}_i(t_1 \dots t_{ar(\mathbf{c}_i)})$ does not necessarily rewrite to a constructor term, even if it does *not* diverge — the rewrite rules of Ξ are not necessarily exhaustive and a deadlock can be reached. To handle this case we define a lambda term $\perp \in \Lambda$, which will represent any deadlocked term.
- Now a map $[\cdot]_\Lambda : \mathcal{T}(\Xi) \rightarrow \Lambda$ can be defined, in such a way that $[t]_\Lambda$ reduces to $\langle\langle u \rangle\rangle_\Lambda$ (where $u \in \mathcal{C}(\Xi)$) if t has normal form u , but $[t]_\Lambda$ reduces to \perp if t rewrites to a deadlock. The map $[\cdot]_\Lambda$ is defined compositionally, that is to say:

$$\begin{aligned} [\mathbf{c}(t_1, \dots, t_{ar(\mathbf{c}_i)})]_\Lambda &= [\mathbf{c}_i]_\Lambda [t_1]_\Lambda \dots [t_{ar(\mathbf{c}_i)}]_\Lambda; \\ [\mathbf{f}_i(t_1, \dots, t_{ar(\mathbf{f}_i)})]_\Lambda &= [\mathbf{f}_i]_\Lambda [t_1]_\Lambda \dots [t_{ar(\mathbf{f}_i)}]_\Lambda. \end{aligned}$$

In other words, $[\cdot]_\Lambda$ is completely specified by its behavior on constructors and function symbols.

- While defining $[\mathbf{c}]_\Lambda$ is relatively easy (Definition 3.2 and Lemma 3.3), $[\mathbf{f}]_\Lambda$ requires a form of pattern matching to be implemented in the λ -calculus (Lemma 3.4 and Definition 3.5).
- The complete simulation is stated in Theorem 3.6. The example in 3.1 may be used along the section to clarify the definitions.

We will first concentrate on constructor terms, encoding them as λ -terms using Scott’s schema [Wad80].

Definition 3.1. • Constructor terms can be easily put in correspondence with λ -terms by way of a map $\langle\langle \cdot \rangle\rangle_\Lambda$ defined by induction as follows:

$$\langle\langle \mathbf{c}_i(t_1 \dots, t_n) \rangle\rangle_\Lambda \equiv \lambda x_1 \dots \lambda x_g. \lambda y. x_i \langle\langle t_1 \rangle\rangle_\Lambda \dots \langle\langle t_n \rangle\rangle_\Lambda.$$

- The function $\langle\langle \cdot \rangle\rangle_\Lambda$ can be extended to a map on constructors:

$$\langle\langle \mathbf{c}_i \rangle\rangle_\Lambda \equiv \lambda x_1 \dots \lambda x_{ar(\mathbf{c}_i)}. \lambda y_1 \dots \lambda y_g. \lambda z. y_i x_1 \dots x_{ar(\mathbf{c}_i)}.$$

Trivially, if t_1, \dots, t_n are in $\mathcal{C}(\Phi)$, $\langle\langle \mathbf{c}_i \rangle\rangle_\Lambda \langle\langle t_1 \rangle\rangle_\Lambda \dots \langle\langle t_n \rangle\rangle_\Lambda$ rewrites to $\langle\langle \mathbf{c}_i(t_1 \dots t_n) \rangle\rangle_\Lambda$ in $ar(\mathbf{c}_i)$ steps.

- To represent an error value, we use the λ -term $\perp \equiv \lambda x_1 \dots \lambda x_g. \lambda y. y$. A λ -term which is either \perp or in the form $\langle\langle t \rangle\rangle_\Lambda$ is denoted with metavariables like X or Y .

The map $\langle\langle \cdot \rangle\rangle_\Lambda$ defines encodings of constructor terms. For function symbols our goal is defining another map $[\cdot]_\Lambda$ returning a λ -term given any term t in $\mathcal{T}(\Xi)$, in such a way that $t \rightarrow^* u$ and $u \in \mathcal{C}(\Xi)$ implies $[t]_\Lambda \rightarrow_v^* \langle\langle u \rangle\rangle_\Lambda$. Moreover, $[t]_\Lambda$ should rewrite to \perp whenever the rewriting of t causes an error (i.e. t has a normal form containing a function symbol). First of all, we define the λ -term $[\mathbf{c}_i]_\Lambda$ corresponding to a constructor \mathbf{c}_i .

Definition 3.2. • For every $1 \leq i \leq g$, for every $0 \leq m \leq ar(\mathbf{c}_i)$, and for every sequence of variables x_1, \dots, x_m , define the λ -term CON_{x_1, \dots, x_m}^i by induction on $ar(\mathbf{c}_i) - m$:

$$\begin{aligned} CON_{x_1, \dots, x_{ar(\mathbf{c}_i)}}^i &\equiv \lambda y_1 \dots \lambda y_g. z. y_i x_1 \dots x_{ar(\mathbf{c}_i)}; \\ \forall m : 0 \leq m < ar(\mathbf{c}_i) \quad CON_{x_1, \dots, x_m}^i &\equiv \lambda y. y N_{1,i}^m \dots N_{g,i}^m L_i^m; \end{aligned}$$

where:

$$\begin{aligned} N_{j,i}^m &\equiv \lambda z_1 \dots \lambda z_{ar(\mathbf{c}_j)} \cdot (\lambda x_{m+1} \cdot CON_{x_1, \dots, x_{m+1}}^i) CON_{z_1, \dots, z_{ar(\mathbf{c}_j)}}^j; \\ L_i^m &\equiv \lambda z_{m+2} \dots \lambda z_{ar(\mathbf{c}_i)} \cdot \perp. \end{aligned}$$

- For every $1 \leq i \leq g$, the λ -term $[\mathbf{c}_i]_\Lambda$ is CON_ε^i .

We need to prove that $[\mathbf{c}_i]_\Lambda$ does what it is supposed to do. We show something slightly stronger:

Lemma 3.3. *There is a constant $n \in \mathbb{N}$ such that for any i , for any m , and for any $\langle\langle t_1 \rangle\rangle_\Lambda, \dots, \langle\langle t_{ar(\mathbf{c}_i)} \rangle\rangle_\Lambda$ in $\mathcal{C}(\Xi)$:*

$$CON_{x_1, \dots, x_m}^i \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_m \rangle\rangle_\Lambda / x_m \} \langle\langle t_{m+1} \rangle\rangle_\Lambda \dots \langle\langle t_{ar(\mathbf{c}_i)} \rangle\rangle_\Lambda \rightarrow^k \langle\langle \mathbf{c}_i(t_1 \dots t_{ar(\mathbf{c}_i)}) \rangle\rangle_\Lambda$$

where $k \leq n$, and

$$CON_{x_1, \dots, x_m}^i \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_m \rangle\rangle_\Lambda / x_m \} X_{m+1} \dots X_{ar(\mathbf{c}_i)} \rightarrow^l \perp$$

where $l \leq n$, whenever X_j is either $\langle\langle t_j \rangle\rangle_\Lambda$ or \perp but at least one among $X_{m+1} \dots X_{ar(\mathbf{c}_i)}$ is \perp .

Proof. We proceed by induction on $ar(\mathbf{c}_i) - m$:

- If $m = ar(\mathbf{c}_i)$, then

$$\begin{aligned} &CON_{x_1, \dots, x_{ar(\mathbf{c}_i)}}^i \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_{ar(\mathbf{c}_i)} \rangle\rangle_\Lambda / x_{ar(\mathbf{c}_i)} \} \\ &\equiv (\lambda y_1 \dots \lambda y_g y_i x_1 \dots x_{ar(\mathbf{c}_i)}) \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_{ar(\mathbf{c}_i)} \rangle\rangle_\Lambda / x_{ar(\mathbf{c}_i)} \} \\ &\equiv \lambda y_1 \dots \lambda y_g \cdot y_i \langle\langle t_1 \rangle\rangle_\Lambda \dots \langle\langle t_{ar(\mathbf{c}_i)} \rangle\rangle_\Lambda \\ &\equiv \langle\langle \mathbf{c}_i(t_1, \dots, t_{ar(\mathbf{c}_i)}) \rangle\rangle_\Lambda. \end{aligned}$$

- If $m < ar(\mathbf{c}_i)$, we use the following abbreviations:

$$\begin{aligned} P_{j,i}^m &\equiv N_{j,i}^m \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_m \rangle\rangle_\Lambda / x_m \}; \\ Q_j^m &\equiv L_j^m \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_m \rangle\rangle_\Lambda / x_m \}. \end{aligned}$$

Let's distinguish two cases:

- If $X_{m+1} \equiv \perp$, then:

$$\begin{aligned} &CON_{x_1, \dots, x_m}^i \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_m \rangle\rangle_\Lambda / x_m \} X_{m+1} \dots X_{ar(\mathbf{c}_i)} \\ &\rightarrow_v (\perp P_{1,i}^m \dots P_{g,i}^m Q_i^m) X_{m+2} \dots X_{ar(\mathbf{c}_i)} \\ &\rightarrow_v^* Q_i^m X_{m+2} \dots X_{ar(\mathbf{c}_i)} \\ &\rightarrow_v^* \perp \end{aligned}$$

• Let X_{m+1} be $\langle\langle t_{m+1} \rangle\rangle_\Lambda$, where $t_{m+1} \equiv \mathbf{c}_j(u_1, \dots, u_{ar(\mathbf{c}_j)})$. Then:

$$\begin{aligned}
& CON_{x_1, \dots, x_m}^i \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_m \rangle\rangle_\Lambda / x_m \} X_{m+1} \dots X_{ar(\mathbf{c}_i)} \\
\rightarrow_v & \langle\langle \mathbf{c}_j(u_1, \dots, u_{ar(\mathbf{c}_j)}) \rangle\rangle_\Lambda P_{1,i}^m \dots P_{g,i}^m Q_i^m X_{m+2} \dots X_{ar(\mathbf{c}_i)} \\
\rightarrow_v^* & P_{j,i}^m \langle\langle u_1 \rangle\rangle_\Lambda \dots \langle\langle u_{ar(\mathbf{c}_j)} \rangle\rangle_\Lambda X_{m+2} \dots X_{ar(\mathbf{c}_i)} \\
\rightarrow_v^* & (\lambda x_{m+1}. CON_{x_1, \dots, x_{m+1}}^i \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_m \rangle\rangle_\Lambda / x_m \}) \\
& (CON_{z_1, \dots, z_{ar(\mathbf{c}_j)}}^j \{ \langle\langle u_1 \rangle\rangle_\Lambda / y_1, \dots, \langle\langle t_{ar(\mathbf{c}_j)} \rangle\rangle_\Lambda / y_{ar(\mathbf{c}_j)} \}) X_{m+2} \dots X_{ar(\mathbf{c}_i)} \\
\rightarrow_v^* & (\lambda x_{m+1}. CON_{x_1, \dots, x_{m+1}}^i \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_m \rangle\rangle_\Lambda / x_m \}) \\
& \langle\langle \mathbf{c}_j(u_1, \dots, u_{ar(\mathbf{c}_j)}) \rangle\rangle_\Lambda X_{m+2} \dots X_{ar(\mathbf{c}_i)} \\
\rightarrow_v^* & CON_{x_1, \dots, x_{m+1}}^i \{ \langle\langle t_1 \rangle\rangle_\Lambda / x_1, \dots, \langle\langle t_{m+1} \rangle\rangle_\Lambda / x_{m+1} \} X_{m+2} \dots X_{ar(\mathbf{c}_i)}
\end{aligned}$$

and, by the inductive hypothesis, the last term in the reduction sequence reduces to the correct normal form. The existence of a natural number n with the prescribed properties is clear by observing that none of the reductions above have a length which depends on the parameters $\langle\langle t_1 \rangle\rangle_\Lambda, \dots, \langle\langle t_m \rangle\rangle_\Lambda$ and $X_{m+1} \dots X_{ar(\mathbf{c}_i)}$.

This concludes the proof. \square

Interpreting function symbols is more difficult, since we have to “embed” the reduction rules into the λ -term interpreting the function symbol. To do that, we need a preliminary result to encode pattern matching. More specifically, suppose $\alpha_1, \dots, \alpha_n$ are non-overlapping sequences of patterns of the same length m , i.e. that for every sequence of constructor terms t_1, \dots, t_m there is at most one i with $1 \leq i \leq n$ such that t_1, \dots, t_m unifies with the patterns in α_i . Then, we need to build a λ -term $PAT_{\alpha_1, \dots, \alpha_n}^m$ which, when fed with m (encodings of) constructor terms and n values, perform pattern matching and select the “right” value, or returns \perp if none of $\alpha_1, \dots, \alpha_n$ unifies with the constructor terms in input.

Lemma 3.4 (Pattern matching). *Let $\alpha_1, \dots, \alpha_n$ be non-overlapping sequences of patterns of the same length m . Then there are a term $PAT_{\alpha_1, \dots, \alpha_n}^m$ and an integer l such that for every sequence of values V_1, \dots, V_n , if $\alpha_i = s_1, \dots, s_m$ then*

$$\begin{aligned}
& PAT_{\alpha_1, \dots, \alpha_n}^m \langle\langle s_1(t_1^1, \dots, t_1^{k_1}) \rangle\rangle_\Lambda \dots \langle\langle s_m(t_m^1, \dots, t_m^{k_m}) \rangle\rangle_\Lambda V_1 \dots V_n \\
& \rightarrow_v^k V_i \langle\langle t_1^1 \rangle\rangle_\Lambda \dots \langle\langle t_1^{k_1} \rangle\rangle_\Lambda \dots \langle\langle t_m^1 \rangle\rangle_\Lambda \dots \langle\langle t_m^{k_m} \rangle\rangle_\Lambda,
\end{aligned}$$

where $k \leq l$, whenever the t_i^j are constructor terms. Moreover,

$$PAT_{\alpha_1, \dots, \alpha_n}^m X_1, \dots, X_m V_1 \dots V_n \rightarrow_v^k \perp,$$

where $k \leq l$, whenever X_1, \dots, X_m do not unify with any of the sequences $\alpha_1, \dots, \alpha_n$ or any of the X_1, \dots, X_m is itself \perp .

Proof. We go by induction on $a = \sum_{i=1}^n \|\alpha_i\|$, where $\|\alpha_i\|$ is the number of constructors occurrences in patterns inside α_i :

- If $a = 0$ and $n = 0$, then we should always return \perp :

$$PAT_\varepsilon^m \equiv \lambda x_1. \dots \lambda x_m. \perp.$$

- If $a = 0$ and $n > 0$, then $n = 1$ and α_1 is simply a sequence of variables x_1, \dots, x_m , because the α_i are assumed to be non-overlapping. Then PAT_{x_1, \dots, x_m}^m is a term defined by induction on m which returns \perp only if one of its first m arguments is \perp and otherwise returns its $m + 1$ -th argument applied to its first m arguments.

- If $a \geq 1$, then there must be integers i and j with $1 \leq i \leq m$ and $1 \leq j \leq n$ such that

$$\alpha_j = s_1, \dots, s_{i-1}, \mathbf{c}_k(r_1, \dots, r_{ar(\mathbf{c}_k)}), s_{i+1}, \dots, s_m$$

for a constructor \mathbf{c}_k and for some patterns s_p and some r_q . Now, for every $1 \leq p \leq n$ and for every $1 \leq j \leq g$ we define sequences of patterns β_p^j and values W_p^j as follows:

- If

$$\alpha_p = s_1, \dots, s_{i-1}, \mathbf{c}_j(q_1, \dots, q_{ar(\mathbf{c}_j)}), s_{i+1} \dots s_m$$

then β_p^j is defined to be the sequence

$$s_1, \dots, s_{i-1}, q_1, \dots, q_{ar(\mathbf{c}_j)}, s_{i+1}, \dots, s_m.$$

Moreover, W_p is simply the identity $\lambda x.x$.

- If

$$\alpha_p = s_1, \dots, s_{i-1}, \mathbf{c}_s(q_1, \dots, q_{ar(\mathbf{c}_s)}), s_{i+1} \dots s_m$$

where $s \neq j$ then β_p^j and W_p^j are both undefined.

- Finally, if

$$\alpha_p = s_1, \dots, s_{i-1}, x, s_{i+1} \dots s_m$$

then β_p^j is defined to be the sequence

$$s_1, \dots, s_{i-1}, x_1, \dots, x_{ar(\mathbf{c}_j)}, s_{i+1}, \dots, s_m.$$

and W_p^j is the following λ -term

$$\lambda x. \lambda y_1. \dots \lambda y_t. x_1. \dots \lambda x_{ar(\mathbf{c}_k)}. \lambda z_1. \dots \lambda z_u. x y_1 \dots y_t (\langle \langle \mathbf{c}_j \rangle \rangle_{\Lambda} x_1 \dots x_{ar(\mathbf{c}_j)}) z_1 \dots z_u$$

where t is the number of variables in s_1, \dots, s_{i-1} and u is the number of variables in s_{i+1}, \dots, s_m .

As a consequence, for every $1 \leq j \leq g$, we can find a natural number t_j and a sequence of pairwise distinct natural numbers i_1, \dots, i_{t_j} such that $\beta_{i_1}^j, \dots, \beta_{i_{t_j}}^j$ are exactly the sequences which can be defined by the above construction. We are now able to formally define $PAT_{\alpha_1, \dots, \alpha_n}^m$; it is the term

$$\lambda x_1. \dots \lambda x_m. \lambda y_1. \dots \lambda y_n. ((x_i Z_1 \dots Z_g Z_{\perp}) x_1 \dots x_{i-1} x_{i+1} \dots x_m) y_1 \dots y_n$$

where

$$\forall 1 \leq j \leq g. Z_j \equiv \lambda z_1. \dots \lambda z_{ar(\mathbf{c}_j)}. \lambda x_1. \dots \lambda x_{i-1}. \lambda x_{i+1}. \dots \lambda x_m. \lambda y_1. \dots \lambda y_n.$$

$$PAT_{\beta_{i_1}^j, \dots, \beta_{i_{t_j}}^j}^{m-1+ar(\mathbf{c}_j)} x_1 \dots x_{i-1} z_1 \dots z_{ar(\mathbf{c}_j)} x_{i+1} \dots x_m (W_{i_1}^j y_{i_1}) \dots (W_{i_{t_j}}^j y_{i_{t_j}})$$

$$Z_{\perp} \equiv \lambda x_1. \dots \lambda x_{i-1}. \lambda x_{i+1}. \dots \lambda x_m. \lambda y_1. \dots \lambda y_n. \perp$$

Notice that, for every j , $a > \sum_{v=1}^{t_j} \|\beta_v^j\|$. Moreover, for every j any β_v^j has the same length $m - 1 + ar(\mathbf{c}_j)$. This justifies the application of the induction hypothesis above. Informally, $PAT_{\alpha_1, \dots, \alpha_n}^m$ first do some case analysis based on the shape of its i -th argument. Based on the topmost constructor in it, one between $Z_1, \dots, Z_h, Z_{\perp}$ is selected which itself do the rest of the pattern matching by way of $PAT_{\beta_{i_1}^j, \dots, \beta_{i_{t_j}}^j}^{m-1+ar(\mathbf{c}_j)}$.

This concludes the proof. \square

Once a general form of pattern matching is available in the λ calculus, we may define the λ -term $[\mathbf{f}_i]_\Lambda$ interpreting the function symbol \mathbf{f}_i .

Definition 3.5. For every function symbol \mathbf{f}_i , let

$$\mathbf{f}_i(\alpha_i^1) \rightarrow t_i^1, \dots, \mathbf{f}_i(\alpha_i^{n_i}) \rightarrow t_i^{n_i}$$

be the rules for \mathbf{f}_i . Moreover, suppose that the variables appearing in the patterns in α_i^j are $z_i^{j,1}, \dots, z_i^{j,m_i,j}$. Observe that the sequences $\alpha_i^1, \dots, \alpha_i^{n_i}$ all have the same length m . Recall that we have a signature with function symbols $\mathbf{f}_1, \dots, \mathbf{f}_h$. For any $1 \leq i \leq h$ the λ -term $[\mathbf{f}_i]_\Lambda$ interpreting \mathbf{f}_i is defined to be:

$$H_i V_1 \dots V_h$$

where

$$\begin{aligned} V_i &\equiv \lambda x_1 \dots \lambda x_h. \lambda y_1 \dots \lambda y_{ar(\mathbf{f}_i)}. PAT_{\alpha_i^1, \dots, \alpha_i^{n_i}}^m y_1 \dots y_{ar(\mathbf{f}_i)} W_i^1 \dots W_i^{n_i}; \\ W_i^j &\equiv \lambda z_1 \dots \lambda z_{m_i,j}. \langle t_i^j \rangle_\Lambda; \end{aligned}$$

whenever $1 \leq i \leq h$ and $1 \leq j \leq n_i$ and $\langle \cdot \rangle_\Lambda$ is defined by induction as follows:

$$\begin{aligned} \langle x \rangle_\Lambda &= x; \\ \langle \mathbf{c}_i(t_1, \dots, t_{ar(\mathbf{c}_i)}) \rangle_\Lambda &= [\mathbf{c}_i]_\Lambda \langle t_1 \rangle_\Lambda \dots \langle t_{ar(\mathbf{c}_i)} \rangle_\Lambda; \\ \langle \mathbf{f}_i(t_1, \dots, t_{ar(\mathbf{f}_i)}) \rangle_\Lambda &= x_i \langle t_1 \rangle_\Lambda \dots \langle t_{ar(\mathbf{f}_i)} \rangle_\Lambda. \end{aligned}$$

We have now implicitly defined how the map $[\cdot]_\Lambda$ behaves on any term in $\mathcal{V}(\Xi, \Upsilon)$:

$$\begin{aligned} [x]_\Lambda &= x; \\ [\mathbf{c}_i(t_1, \dots, t_{ar(\mathbf{c}_i)})]_\Lambda &= [\mathbf{c}_i]_\Lambda [t_1]_\Lambda \dots [t_{ar(\mathbf{c}_i)}]_\Lambda; \\ [\mathbf{f}_i(t_1, \dots, t_{ar(\mathbf{f}_i)})]_\Lambda &= [\mathbf{f}_i]_\Lambda [t_1]_\Lambda \dots [t_{ar(\mathbf{f}_i)}]_\Lambda. \end{aligned}$$

Theorem 3.6. *There is a natural number k such that for every function symbol \mathbf{f} and for every $t_1, \dots, t_{ar(\mathbf{f})} \in \mathcal{C}(\Xi)$, the following three implications hold, where u stands for $\mathbf{f}(t_1, \dots, t_{ar(\mathbf{f})})$ and M stands for $[\mathbf{f}]_\Lambda \langle t_1 \rangle_\Lambda \dots \langle t_{ar(\mathbf{f})} \rangle_\Lambda$:*

- *If u rewrites to $v \in \mathcal{C}(\Xi)$ in n steps, then M rewrites to $\langle v \rangle_\Lambda$ in at most kn steps.*
- *If u rewrites to a normal form $v \notin \mathcal{C}(\Xi)$, then M rewrites to \perp .*
- *If u diverges, then M diverges.*

Proof. By an easy combinatorial argument following from the definition of $[\cdot]_\Lambda$. Actually, a slightly stronger statement should be proved to make the proof formal: there is a natural number k such that for every $u \in \mathcal{V}(\Xi, \Upsilon)$, for every $t_1, \dots, t_m \in \mathcal{C}(\Xi)$ (where m is the number of distinct variables in u), the following three implications hold, where M stands for $[u]_\Lambda$.

- *If $u(t_1, \dots, t_m)$ rewrites to $v \in \mathcal{C}(\Xi)$ in n steps, then $M(\langle t_1 \rangle_\Lambda, \dots, \langle t_m \rangle_\Lambda)$ rewrites to $\langle v \rangle_\Lambda$ in at most $kn|u|$ steps.*
- *If $u(t_1, \dots, t_m)$ rewrites to a normal form $v \notin \mathcal{C}(\Xi)$, then $M(\langle t_1 \rangle_\Lambda, \dots, \langle t_m \rangle_\Lambda)$ rewrites to \perp .*
- *If $u(t_1, \dots, t_m)$ diverges, then $M(\langle t_1 \rangle_\Lambda, \dots, \langle t_m \rangle_\Lambda)$ diverges.*

The first statement can be proved by induction on n . The second and third one are quite easy. \square

Clearly, the constant k in Theorem 3.6 depends on Ξ , but is independent on the particular term u .

3.1. An Example. In this section, we will describe the encoding of a concrete OCRS called ADD as a set of λ -terms. The signature Σ_{ADD} contains two constructor symbols $\mathbf{0}$ and \mathbf{s} , with arity 0 and 1 (respectively), and a single function symbol \mathbf{add} of arity 2. The only two rules in \mathcal{R}_{ADD} are the following:

$$\begin{aligned} \mathbf{add}(\mathbf{0}, x) &\rightarrow x; \\ \mathbf{add}(\mathbf{s}(x), y) &\rightarrow \mathbf{s}(\mathbf{add}(x), y). \end{aligned}$$

Let us construct first some λ -terms in the image of $\langle\langle\rangle\rangle_\Lambda$:

$$\begin{aligned} \perp &= \lambda x. \lambda y. \lambda z. z; \\ \langle\langle \mathbf{0} \rangle\rangle_\Lambda &= \lambda x. \lambda y. \lambda z. x; \\ \langle\langle \mathbf{s}(\mathbf{0}) \rangle\rangle_\Lambda &= \lambda x. \lambda y. \lambda z. y \langle\langle \mathbf{0} \rangle\rangle_\Lambda; \\ \langle\langle \mathbf{s}(\mathbf{s}(\mathbf{0})) \rangle\rangle_\Lambda &= \lambda x. \lambda y. \lambda z. y \langle\langle \mathbf{s}(\mathbf{0}) \rangle\rangle_\Lambda; \\ &\vdots \\ \langle\langle \mathbf{s} \rangle\rangle_\Lambda &= \lambda w. \lambda x. \lambda y. \lambda z. yw. \end{aligned}$$

We now take a look at $[\mathbf{s}]_\Lambda$. By definition:

$$\begin{aligned} [\mathbf{s}]_\Lambda &\equiv CON_\varepsilon^2 \equiv \lambda y. y N_{1,2}^0 N_{2,2}^0 L_2^0 \\ &\equiv \lambda y. y ((\lambda x_1. CON_{x_1}^2) CON_\varepsilon^1) (\lambda z_1. (\lambda x_1. CON_{x_1}^2) (CON_{z_1}^2)) \perp. \end{aligned}$$

This λ -term indeed “simulates” the successor constructor, when fed with an input. Suppose $u \in \mathcal{C}(ADD)$, then:

$$\begin{aligned} [\mathbf{s}]_\Lambda \langle\langle \mathbf{0} \rangle\rangle_\Lambda &\rightarrow^4 (\lambda x_1. CON_{x_1}^2) CON_\varepsilon^1 \equiv (\lambda x_1. CON_{x_1}^2) \langle\langle \mathbf{0} \rangle\rangle_\Lambda \\ &\rightarrow \langle\langle \mathbf{s}(\mathbf{0}) \rangle\rangle_\Lambda; \\ [\mathbf{s}]_\Lambda \langle\langle \mathbf{s}(u) \rangle\rangle_\Lambda &\rightarrow^4 (\lambda z_1. (\lambda x_1. CON_{x_1}^2) (CON_{z_1}^2)) \langle\langle u \rangle\rangle_\Lambda \rightarrow (\lambda z_1. (\lambda x_1. CON_{x_1}^2) \langle\langle \mathbf{s}(u) \rangle\rangle_\Lambda \\ &\rightarrow \langle\langle \mathbf{s}(\mathbf{s}(u)) \rangle\rangle_\Lambda; \\ [\mathbf{s}]_\Lambda \perp &\rightarrow^4 \perp. \end{aligned}$$

Finally, consider \mathbf{add} , the only function symbol of Σ_{ADD} . By definition:

$$[\mathbf{add}]_\Lambda \equiv H_1 V_1 \equiv H_1 (\lambda x_1. \lambda y_1. \lambda y_2. PAT_{(\mathbf{0}, x), (\mathbf{s}(x), y)} y_1 y_2 W_1^1 W_i^2)$$

It is easy to verify that, by Lemma 3.4,

$$\begin{aligned} [\mathbf{add}]_\Lambda \perp \langle\langle t \rangle\rangle_\Lambda &\rightarrow^* \perp; \\ [\mathbf{add}]_\Lambda \langle\langle \mathbf{s}(\mathbf{0}) \rangle\rangle_\Lambda \langle\langle \mathbf{s}(\mathbf{s}(\mathbf{0})) \rangle\rangle_\Lambda &\rightarrow^* \langle\langle \mathbf{s}(\mathbf{s}(\mathbf{s}(\mathbf{0}))) \rangle\rangle_\Lambda. \end{aligned}$$

4. GRAPH REPRESENTATION

The previous two sections proved the main simulation result of the paper. To complete the picture, we show in this section that the unitary cost model for the (weak call-by-value) λ -calculus (and hence the number of rewriting in a OCRSs) is polynomially related to the actual cost of implementing those reductions². We do so by introducing term graph rewriting, following [BEG⁺86] but adapting the framework to call-by-value constructor rewriting. Contrarily to what we did in Section 2, we will stay abstract here: our attention will not be restricted to the particular graph rewrite system that is needed to implement reduction in the λ -calculus.


We refer the reader to our [DLM10] for more details on efficient simulations between term graph rewriting and constructor term rewriting, both under innermost (i.e., call-by-value) and outermost (i.e., call-by-name) reduction strategies.

Definition 4.1 (Labelled Graph). Given a signature Σ , a *labelled graph over Σ* consists of a directed acyclic graph together with an ordering on the outgoing edges of each node and a (partial) labelling of nodes with symbols from Σ such that the out-degree of each node matches the arity of the corresponding symbols (and is 0 if the labelling is undefined). Formally, a labelled graph is a triple $G = (V, \alpha, \delta)$ where:

- V is a set of *vertices*.
- $\alpha : V \rightarrow V^*$ is a (total) *ordering function*.
- $\delta : V \rightarrow V$ is a (partial) *labelling function* such that the length of $\alpha(v)$ is the arity of $\delta(v)$ if $\delta(v)$ is defined and is 0 otherwise.

A labelled graph (V, α, δ) is *closed* iff δ is a total function.

Consider the signature $\Sigma = \{\mathbf{f}, \mathbf{g}, \mathbf{h}, \mathbf{p}\}$, where arities of $\mathbf{f}, \mathbf{g}, \mathbf{h}, \mathbf{p}$ are 2, 1, 0, 2 respectively, and $\mathbf{g}, \mathbf{h}, \mathbf{p}$ are constructors. Examples of labelled graphs over the signature Σ are the following ones:



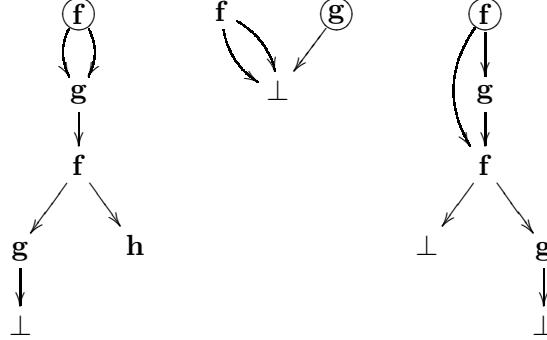
The symbol \perp denotes vertices where the underlying labelling function is undefined (and, as a consequence, no edge departs from such vertices). Their role is similar to the one of variables in terms.

If one of the vertices of a labelled graph is selected as the *root*, we obtain a term graph:

Definition 4.2 (Term Graph). A *term graph*, is a quadruple $G = (V, \alpha, \delta, r)$, where (V, α, δ) is a labelled graph and $r \in V$ is the *root* of the term graph.

²As mentioned in the introduction, see [SGM02] for another proof of this with other means.

The following are graphic representations of some term graphs.



The root is the only vertex drawn inside a circle.

There are some classes of paths which are particularly relevant for our purposes.

Definition 4.3 (Path). A path v_1, \dots, v_n in a labelled graph $G = (V, \alpha, \delta)$ is said to be:

- A *constructor path* iff for every $1 \leq i \leq n$, the symbol $\delta(v_i)$ is a constructor;
- A *pattern path* iff for every $1 \leq i \leq n$, $\delta(v_i)$ is either a constructor symbol or is undefined;
- A *left path* iff $n \geq 1$, the symbol $\delta(v_1)$ is a function symbol and v_2, \dots, v_n is a pattern path.

Definition 4.4 (Homomorphism). A *homomorphism* between two labelled graphs $G = (V_G, \alpha_G, \delta_G)$ and $H = (V_H, \alpha_H, \delta_H)$ over the same signature Σ is a function φ from V_G to V_H preserving the term graph structure. In particular

$$\begin{aligned} \delta_H(\varphi(v)) &= \delta_G(v) \\ \alpha_H(\varphi(v)) &= \varphi^*(\alpha_G(v)) \end{aligned}$$


for any $v \in \text{dom}(\delta)$, where φ^* is the obvious generalization of φ to sequences of vertices. An *homomorphism* between two term graphs $G = (V_G, \alpha_G, \delta_G, r_G)$ and $H = (V_H, \alpha_H, \delta_H, r_H)$ is a homomorphism between $(V_G, \alpha_G, \delta_G)$ and $(V_H, \alpha_H, \delta_H)$ such that $\varphi(r_G) = r_H$. Two labelled graphs G and H are isomorphic iff there is a bijective homomorphism from G to H ; in this case, we write $G \cong H$. Similarly for term graphs.

In the following, we will consider term graphs modulo isomorphism, i.e., $G = H$ iff $G \cong H$. Observe that two isomorphic term graphs have the same graphical representation.

Definition 4.5 (Graph Rewrite Rule). A *graph rewrite rule* over a signature Σ is a triple $\rho = (G, r, s)$ such that:

- G is a labelled graph;
- r, s are vertices of G , called the *left root* and the *right root* of ρ , respectively.
- Any path starting in r is a left path.

The following are examples of graph rewrite rules, assuming a to be a function symbol and b, c, d to be constructors:



Definition 4.6 (Subgraph). Given a labelled graph $G = (V_G, \alpha_G, \delta_G)$ and any vertex $v \in V_G$, the *subgraph of G rooted at v* , denoted $G \downarrow v$, is the term graph $(V_{G \downarrow v}, \alpha_{G \downarrow v}, \delta_{G \downarrow v}, r_{G \downarrow v})$ where

- $V_{G \downarrow v}$ is the subset of V_G whose elements are vertices which are reachable from v in G .
- $\alpha_{G \downarrow v}$ and $\delta_{G \downarrow v}$ are the appropriate restrictions of α_G and δ_G to $V_{G \downarrow v}$.
- $r_{G \downarrow v}$ is v .

Definition 4.7 (Redex). Given a labelled graph G , a *redex* for G is a pair (ρ, φ) , where ρ is a rewrite rule (H, r, s) and φ is a homomorphism between $H \downarrow r$ and G such that for any vertex $v \in V_{H \downarrow r}$ with $v \notin \text{dom}(\delta_{H \downarrow r})$, any path starting in $\varphi(v)$ is a constructor path.


The last condition in the definition of a redex is needed to capture the call-by-value nature of the rewriting process.

Given a term graph G and a redex $((H, r, s), \varphi)$, the result of firing the redex is another term graph obtained by successively applying the following three steps to G :


1. The *build phase*: create an isomorphic copy of the portion of $H \downarrow s$ not contained in $H \downarrow r$, and add it to G , obtaining J . The underlying ordering and labelling functions are defined in the natural way.
2. The *redirection phase*: all edges in J pointing to $\varphi(r)$ are replaced by edges pointing to the copy of s . If $\varphi(r)$ is the root of G , then the root of the newly created graph will be the newly created copy of s . The graph K is obtained.
3. The *garbage collection phase*: all vertices which are not accessible from the root of K are removed. The graph I is obtained.

We will write $G \xrightarrow{(H, r, s)} I$ (or simply $G \rightarrow I$, if this does not cause ambiguity) in this case.


As an example, consider the term graph G and the rewrite rule $\rho = (H, r, s)$:



There is a homomorphism φ from $H \downarrow r$ to G . In particular, φ maps r to the rightmost vertex in G . Applying the build phase and the redirection phase we get J and K as follows:



Finally, applying the garbage collection phase, we get the result of firing the redex (ρ, φ) :



Definition 4.8. A constructor graph rewrite system (CGRS) over a signature Σ consists of a set \mathcal{G} of graph rewrite rules on Σ .

4.1. From Term Rewriting to Graph Rewriting. Any term t over a signature Σ can be turned into a graph $[t]$ in the obvious way: take as $[t]$ the abstract syntax tree of t , where vertices are in one-to-one correspondence with symbol occurrences in t . Conversely, any term graph G over Σ can be turned into a term $\langle G \rangle$ over Σ by simply unfolding the graph, that is applying (the label of) any vertex to (the terms obtained as unfolding of) its sons (remember: we only consider acyclic graphs here). We omit the boring formal definitions of both $[\cdot]$ and $\langle \cdot \rangle$; it is clear that for any term t , $\langle [t] \rangle = t$, while in general $[[\langle G \rangle]]$ is not equal to G , since the sharing present in G is lost during the unfolding.


Definition 4.9. Given a constructor rewriting system \mathcal{R} over Σ , the corresponding constructor graph rewriting system $[\mathcal{R}]$ is defined by translating the terms with $[\cdot]$ and by translating any term rewrite rule $t \rightarrow u$ over Σ into a graph rewrite rule (G, r, s) as follows:

- Take the graphs $[t]$ and $[u]$ (which are trees, in fact).
- From the union of these two trees, share those nodes representing the same variable in t and u . This is G .
- Take r to be the root of t in G and s to be the root of u in G .

As an example, consider the rewrite rule

$$\mathbf{f}(\mathbf{g}(x), y) \rightarrow \mathbf{g}(\mathbf{f}(y, \mathbf{f}(y, x))).$$

Its translation as a graph rewrite rule is the following:



Given a constructor rewriting system \mathcal{R} , it is easy to realize that the following invariant is preserved while performing rewriting in $[\mathcal{R}]$: whenever any vertex v can be reached by two distinct paths starting at the root (i.e., v is *shared*), any path starting at v is a constructor path. A term graph satisfying this invariant is said to be *constructor-shared*.

Constructor-sharedness holds for term graphs coming from terms and is preserved by graph rewriting:

Lemma 4.10. *For every closed term t , $[t]$ is constructor-shared. Moreover, if G is closed and constructor-shared and $G \rightarrow I$ in $[\mathcal{R}]$, then I is constructor-shared.*

Proof. The fact that $[t]$ is constructor-shared for every t follows from the way the $[\cdot]$ map is defined: it does not introduce any sharing. Now, suppose G is constructor-shared and

$$G \xrightarrow{(H,r,s)} I$$

where (H, r, s) corresponds to a term rewrite rule $t \rightarrow u$. The term graph J obtained from G by the build phase is itself constructor-shared: it is obtained from G by adding some new nodes, namely an isomorphic copy of the portion of $H \downarrow s$ not contained in $H \downarrow r$. Notice that J is constructor-shared in a stronger sense: any vertex which can be reached from the newly created copy of s by two distinct paths must be a constructor path. This is a consequence of (H, r, s) being a graph rewrite rule corresponding to a term rewrite rule $t \rightarrow u$, where the only shared vertices are those where the labelling function is undefined. The redirection phase preserves itself constructor-sharedness, because only one pointer is redirected (the vertex is labelled by a function symbol) and the destination of this redirection is a vertex (the newly created copy of s) which had no edge incident to it. Clearly, the garbage collection phase preserves constructor-sharedness. \square

Lemma 4.11. *A closed term graph G in $[\mathcal{R}]$ is a normal form iff $\langle G \rangle$ is a normal form.*


Proof. Clearly, if a closed term graph G is in normal form, then $\langle G \rangle$ is a term in normal form, because each redex in G translates to a redex in $\langle G \rangle$. On the other hand, if $\langle G \rangle$ is in normal form, then G is in normal form: each redex in $\langle G \rangle$ translates back to a redex in G . \square

Reduction at the level of graphs correctly simulates reduction at the level of terms, but only if the underlying graphs are constructor shared:

Lemma 4.12. *If G is closed and constructor-shared, and $G \rightarrow I$ in $[\mathcal{R}]$, then $\langle G \rangle \rightarrow \langle I \rangle$ in \mathcal{R} .*

Proof. The fact that each reduction step starting in G can be mimicked in $\langle G \rangle$ is known from the literature. If G is constructor-shared, then the simulation is done in exactly one reduction step, because any redex in a constructor-shared term graph cannot be shared. \square

When G is not constructor-shared, a counterexample can be easily built. Consider the term rewrite rule $\mathbf{f}(\mathbf{h}, \mathbf{h}) \rightarrow \mathbf{h}$ and the following term graph:



It corresponds to $\mathbf{f}(\mathbf{f}(\mathbf{h}, \mathbf{h}), \mathbf{f}(\mathbf{h}, \mathbf{h}))$, and it is not constructor-shared, since the shared vertex a is not a constructor. It rewrites in *one* step to



while the term $\mathbf{f}(\mathbf{f}(\mathbf{h}, \mathbf{h}), \mathbf{f}(\mathbf{h}, \mathbf{h}))$ rewrites to $\mathbf{f}(\mathbf{h}, \mathbf{h})$ in *two* steps.

As can be expected, graph reduction is also complete with respect to term reduction, with the only *proviso* that term graphs must be constructor-shared:

Lemma 4.13. *If $t \rightarrow u$ in \mathcal{R} , G is constructor-shared and $\langle G \rangle = t$, then $G \rightarrow I$ in $[\mathcal{R}]$, where $\langle I \rangle = u$. \square*

Theorem 4.14 (Graph Reducibility). *For every constructor rewrite system \mathcal{R} over Σ and for every term t over Σ , the following two conditions are equivalent:*

1. $t \rightarrow^n u$ in \mathcal{R} , where u is in normal form;
2. $[t] \rightarrow^n G$ in $[\mathcal{R}]$, where G is in normal form and $\langle G \rangle = u$.

Proof. Suppose $t \rightarrow^n u$, where u is in normal form. Then, by applying Lemma 4.13, we obtain a term graph G such that $[t] \rightarrow^n G$ and $\langle G \rangle = u$. By Lemma 4.10, G is constructor-shared and, by Lemma 4.11, it is in normal form. Now, suppose $[t] \rightarrow^n G$ where $\langle G \rangle = u$ and G is in normal form. By applying n times Lemma 4.12, we obtain that $\langle [t] \rangle \rightarrow^n \langle G \rangle = u$. But $\langle [t] \rangle = t$ and u is a normal form by Lemma 4.11, since $[t]$ and G are constructor shared due to Lemma 4.10. \square

There are *term* rewrite systems which are not graph reducible, i.e. for which the two conditions of Theorem 4.14 are not equivalent (see [BEG⁺86]). However, any *orthogonal constructor* rewrite system is graph reducible, due to the strict constraints on the shape of rewrite rules [Plu90]. This result can be considered as a by-product of our analysis, for which graph rewriting is only instrumental.

4.2. Lambda-Terms Can Be Efficiently Reduced by Graph Rewriting. As a corollary of Theorems 4.14 and 2.9, we may reduce λ -terms using term graphs. To this purpose, we apply the construction of the previous section to the OCRS Φ that we defined in Section 2. Let then $\Theta = [\Phi]$:

Corollary 4.15. *Let $M \in \Lambda$ be a closed λ -term. The following two conditions are equivalent:*

1. $M \rightarrow_v^n N$ where N is in normal form;
2. $[[M]_\Phi] \rightarrow^n G$ where $\langle\langle G \rangle\rangle_\Lambda = N$ and G is in normal form. \square

Let us now analyze more closely the combinatorics of graph rewriting in Θ , so that we can obtain information on the efficiency of this simulation.

- Consider a closed λ -term M and a term graph G such that $[[M]_{\Phi}] \rightarrow^* G$. By Proposition 1 and Lemma 4.12, for every constructor $\mathbf{c}_{x,N}$ appearing as a label of a vertex in G , N is a subterm of M .
- As a consequence, if $[[M]_{\Phi}] \rightarrow^* G \rightarrow H$, then the difference $|H| - |G|$ cannot be too big: at most $|M|$. Therefore, if $[[M]_{\Phi}] \rightarrow^n G$ then $|G| \leq (n+1)|M|$. Here, we exploit in an essential way the possibility of sharing constructors.
- Whenever $[[M]_{\Phi}] \rightarrow^n G$, computing a graph H such that $G \rightarrow H$ takes polynomial time in $|G|$, which is itself polynomially bounded by n and $|M|$.

Hence (recall that $Time_v(M)$ is the number of weak call-by-value beta steps to normal form):

Theorem 4.16. *There is a polynomial $p : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that for every λ -term M , the normal form of $[[M]_{\Phi}]$ can be computed in time at most $p(|M|, Time_v(M))$. \square*

This cannot be achieved when using explicit representations of λ -terms. Moreover, reading back a λ -term from a term graph can take exponential time.

We can complement Theorem 4.16 with a completeness statement — any universal computational model with an invariant cost model can be embedded in the λ -calculus with a polynomial overhead. We can exploit for this the analogous result we proved in [DLM08] (Section 4, Theorem 1) — the unitary cost model is easily proved to be more parsimonious than the difference cost model considered in [DLM08].

Theorem 4.17. *Let $f : \Sigma^* \rightarrow \Sigma^*$ be computed by a Turing machine \mathcal{M} in time g . Then, there are a λ -term $N_{\mathcal{M}}$ and a suitable encoding $\ulcorner \cdot \urcorner : \Sigma^* \rightarrow \Lambda$ such that $N_{\mathcal{M}} \ulcorner v \urcorner$ normalizes to $\ulcorner f(v) \urcorner$ in $O(g(|v|))$ beta steps. \square*

The encoding $\ulcorner \cdot \urcorner$ mentioned in the theorem depends only on (the cardinality of) Σ (but not on the Turing machine). Interestingly enough it exploits once again the scheme that we used in Definition 3.1: encode the empty string ε as a zero-ary constructor, and any symbol in Σ as a unary constructor (see [DLM08] for details).

5. VARIATIONS: CALL-BY-NAME REDUCTION

In the previous sections, λ -calculus was endowed with weak call-by-value reduction. The same technique, however, can be applied to weak call-by-name reduction, as we will sketch in this section. Λ is now endowed with a relation \rightarrow_h defined as follows:

$$\frac{}{(\lambda x.M)N \rightarrow_h M\{N/x\}} \qquad \frac{M \rightarrow_h N}{ML \rightarrow_h NL}$$

Similarly to the call-by-value case, $Time_h(M)$ stands for the number of reduction steps to the normal form of M (if any). Since the relation \rightarrow_h is deterministic (i.e., functional), $Time_h(M)$ is well-defined.

We need another OCRS, called Ψ , which is similar to Φ but designed to simulate weak call-by-name reduction:

- The signature Σ_{Ψ} includes the binary function symbol **app** and constructor symbols $\mathbf{c}_{x,M}$ for every $M \in \Lambda$ and every $x \in \Upsilon$, exactly as Σ_{Φ} . Moreover, there is another binary

constructor symbol **capp**. To every term $M \in \Lambda$ we can associate terms $\{M\}_\Psi, [M]_\Psi \in \mathcal{V}(\Psi, \Upsilon)$ as follows:

$$\begin{aligned} \{x\}_\Psi &= x \\ \{\lambda x.M\}_\Psi &= \mathbf{c}_{x,M}(x_1, \dots, x_n), \text{ where } \mathbf{FV}(\lambda x.M) = x_1, \dots, x_n \\ \{MN\}_\Psi &= \mathbf{capp}(\{M\}_\Psi, \{N\}_\Psi) \\ [x]_\Psi &= x \\ [\lambda x.M]_\Psi &= \mathbf{c}_{x,M}(x_1, \dots, x_n), \text{ where } \mathbf{FV}(\lambda x.M) = x_1, \dots, x_n \\ [MN]_\Psi &= \mathbf{app}([M]_\Psi, [N]_\Psi) \end{aligned}$$

Notice that $\{\cdot\}_\Psi$ maps λ -terms to *constructor* terms, while terms obtained via $[\cdot]_\Psi$ can contain function symbols. The “official” translation of a term M is thus $[M]_\Psi$, where only the applications “on the spine” of M are encoded with **app**. All other applications are frozen by the constructor **capp**.

- The rewrite rules in \mathcal{R}_Ψ are all the rules in the following form:

$$\begin{aligned} \mathbf{app}(\mathbf{c}_{z,z}, \mathbf{capp}(w, f)) &\rightarrow \mathbf{app}(w, f) \\ \mathbf{app}(\mathbf{c}_{z,z}, \mathbf{c}_{x,M}(x_1, \dots, x_n)) &\rightarrow \mathbf{c}_{x,M}(x_1, \dots, x_n) \\ \mathbf{app}(\mathbf{c}_{z,w}(\mathbf{capp}(f, g)), h) &\rightarrow \mathbf{app}(f, g) \\ \mathbf{app}(\mathbf{c}_{z,w}(\mathbf{c}_{x,M}(x_1, \dots, x_n)), h) &\rightarrow \mathbf{c}_{x,M}(x_1, \dots, x_n) \\ \mathbf{app}(\mathbf{c}_{y,N}(y_1, \dots, y_m), y) &\rightarrow [N]_\Psi \end{aligned}$$

where M ranges over λ -terms, N ranges over abstractions and applications, $\mathbf{FV}(\lambda x.M) = x_1, \dots, x_n$ and $\mathbf{FV}(\lambda y.N) = y_1, \dots, y_m$. These rewrite rules are said to be *ordinary rules*. We also need the following *administrative* rule:

$$\mathbf{app}(\mathbf{capp}(x, y), z) \rightarrow \mathbf{app}(\mathbf{app}(x, y), z).$$

The CTRS Ψ is more complicated than Φ , because we need to force reduction to happen only in head position. The applications **app** (on the spine) may be fired immediately. Observe, however, that the main rewriting rule (the last of the ordinary ones) is restricted to those $\mathbf{c}_{y,N}$ where N is *not* a variable. When N is a single variable, the corresponding beta redex would be either $(\lambda x.x)L$ or $(\lambda z.w)L$, with w free. In the former case, an application at the top level of L (encoded as a **capp** at this point) would become the top level application of the spine of the reduct: the first ordinary reduction rule handles this case, unfreezing **capp** into **app**. When, on the other hand, the encoded redex is $(\lambda z.w)L$, we do not need to worry for L , which will be discarded, but in the term-reduction we must take care of the eventual substitution that may occur for w : the term substituted for w may have a top level frozen application **capp** that must be converted into an **app** — this is the role of the third ordinary reduction rule. The second and fourth reduction rules just handle the remaining cases (they would be instances of the last ordinary rule if this was not restricted to the non-variable cases). A last remark on the administrative rule. There are never administrative redexes in the translation $[M]_\Psi$ of a term. During reduction, however, by the effect of the other rules a frozen application (a **capp**) may appear on the spine. The administrative rule recognizes this situation and unfreezes the application.

As usual, to every term $t \in \mathcal{V}(\Psi, \Upsilon)$ we can associate a term $\langle t \rangle_\Lambda$:

$$\begin{aligned} \langle x \rangle_\Lambda &= x \\ \langle \mathbf{app}(u, v) \rangle_\Lambda &= \langle \mathbf{capp}(u, v) \rangle_\Lambda = \langle u \rangle_\Lambda \langle v \rangle_\Lambda \\ \langle \mathbf{c}_{x, M}(t_1, \dots, t_n) \rangle_\Lambda &= (\lambda x. M) \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \} \end{aligned}$$

where $\text{FV}(\lambda x. M) = x_1, \dots, x_n$. A term $t \in \mathcal{T}(\Psi)$ is *canonical* if either $t = \mathbf{c}_{x, M}(t_1, \dots, t_n) \in \mathcal{C}(\Psi)$ or $t = \mathbf{app}(u, v)$ where u is canonical and $v \in \mathcal{C}(\Psi)$.

Lemma 5.1. *For every closed $M \in \Lambda$, $[M]_\Psi$ is canonical.*

Proof. By a straightforward induction on M . □

The obvious variation on Equation (2.1) holds here:

$$\langle [M]_\Psi \{ t_1 / x_1, \dots, t_n / x_n \} \rangle_\Lambda = M \{ \langle t_1 \rangle_\Lambda / x_1, \dots, \langle t_n \rangle_\Lambda / x_n \}. \quad (5.1)$$

Ψ mimics call-by-name reduction in much the same way Φ mimics call-by-value reduction. However, one reduction step in the λ -calculus corresponds to $n \geq 1$ steps in Ψ , although n is kept under control:

Lemma 5.2. *Suppose that $t \in \mathcal{T}(\Psi)$ is canonical and that $t \rightarrow u$. Then there is a natural number n such that:*

1. $\langle t \rangle_\Lambda \rightarrow_h \langle u \rangle_\Lambda$;
2. *There is a canonical term $v \in \mathcal{T}(\Psi)$ such that $u \rightarrow^n v$;*
3. $|w|_{\mathbf{app}} = |u|_{\mathbf{app}} + m$ whenever $u \rightarrow^m w$ and $m \leq n$;
4. $\langle w \rangle_\Lambda = \langle u \rangle_\Lambda$ whenever $u \rightarrow^m w$ and $m \leq n$.

Proof. A term t is said to be *semi-canonical* iff $t = \mathbf{app}(u, v)$, where $v \in \mathcal{C}(\Psi)$ and u is either semi-canonical or is itself an element of $\mathcal{C}(\Psi)$. We now prove that if t is semi-canonical, there there are a natural number n and a canonical term u such that:

- $t \rightarrow^n u$;
- $|v|_{\mathbf{app}} = |t|_{\mathbf{app}} + m$ whenever $t \rightarrow^m v$ and $m \leq n$;
- $\langle v \rangle_\Lambda = \langle t \rangle_\Lambda$ whenever $t \rightarrow^m v$ and $m \leq n$.

We proceed by induction on $|t|$. By definition t is in the form $\mathbf{app}(w, d)$; we have three cases:

- w is semi-canonical. Then, we get what we want by induction hypothesis.
- w is in $\mathcal{C}(\Psi)$ and has the form $\mathbf{c}_{x, M}(t_1, \dots, t_m)$. Then, $n = 0$ and t is itself canonical.
- w is in $\mathcal{C}(\Psi)$ and has the form $\mathbf{capp}(e, f)$. Then

$$t = \mathbf{app}(\mathbf{capp}(e, f), d) \rightarrow \mathbf{app}(\mathbf{app}(e, f), d).$$

Apply now the induction hypothesis to $\mathbf{app}(e, f)$ (since its length is strictly smaller than $|t|$).

We can now proceed as in Lemma 2.7, since whenever t rewrites to u by one of the ordinary rules, u is semi-canonical. □

Lemma 5.3. *A canonical term $t \in \mathcal{T}(\Psi)$ is in normal form iff $\langle t \rangle_\Lambda$ is in normal form.*

Proof. We first prove that any canonical normal form t can be written as $\mathbf{c}_{x, M}(t_1, \dots, t_n)$, where $t_1, \dots, t_n \in \mathcal{C}(\Psi)$. We proceed by induction on t :

- If $t = \mathbf{c}_{x, M}(t_1, \dots, t_n)$, then the thesis holds.

- If $t = \mathbf{app}(u, v)$, then u is canonical and in normal form, hence in the form $\mathbf{c}_{x,M}(t_1, \dots, t_n)$ by induction hypothesis. As a consequence, t is not a normal form, which is a contradiction.

We can now prove the statement of the lemma, by distinguishing two cases:

- If $t = \mathbf{c}_{x,M}(t_1, \dots, t_n)$, where $t_1, \dots, t_n \in \mathcal{C}(\Psi)$, then t is in normal form and $\langle t \rangle_\Lambda$ is an abstraction, hence a normal form.
- If $t = \mathbf{app}(u, v)$, then t cannot be a normal form, since u is canonical and in normal form and, as a consequence, it can be written as $\mathbf{c}_{x,M}(t_1, \dots, t_n)$.

This concludes the proof. \square

Observe that this property holds only if t is canonical: a non-canonical term can reduce to another one (canonical or not) even if the underlying λ -term is a normal form.

Lemma 5.4. *If $M \rightarrow_h N$, t is canonical and $\langle t \rangle_\Lambda = M$, then $t \rightarrow u$, where $\langle u \rangle_\Lambda = N$ and $|u|_{\mathbf{app}} + 1 \geq |t|_{\mathbf{app}}$.*

Proof. Analogous to the one of Lemma 2.8 for the first part of the statement. For the bound on the number of \mathbf{app} , argument similarly to the proof of Lemma 5.2. \square

The slight mismatch between call-by-name reduction in Λ and reduction in Ψ is anyway harmless globally. As we now show, the total number of reduction steps in Ψ is at most two times as large as the total number of call-by-name reduction steps in Λ .

Theorem 5.5 (Term Reducibility). *Let $M \in \Lambda$ be a closed term. The following two conditions are equivalent:*

1. $M \rightarrow_h^n N$ where N is in normal form;
2. $[M]_\Psi \rightarrow^m t$ where $\langle t \rangle_\Lambda = N$ and t is in normal form.

Moreover $n \leq m \leq 2n$.

Proof. Suppose $M \rightarrow_h^n N$, where N is in normal form. M is closed and, by Lemma 5.1, $[M]_\Psi$ is canonical. By iterating over Lemma 5.2 and Lemma 5.4, we obtain the existence of a term t such that $\langle t \rangle_\Lambda = N$, t is in normal form and $[M]_\Psi \rightarrow^m t$, where $m \geq n$ and

$$|t|_{\mathbf{app}} - |[M]_\Psi|_{\mathbf{app}} \geq (m - n) - n.$$

Since $|t|_{\mathbf{app}} = 0$ (t is in normal form), $m \leq 2n$. If $[M]_\Psi \rightarrow^m t$ where $\langle t \rangle_\Lambda = N$ and t is in normal form, then by iterating over Lemma 5.2 we obtain that $M \rightarrow_h^n N$ where $n \leq m \leq 2n$ and N is in normal form. \square

Ξ is the graph rewrite system corresponding to Ψ , in the sense of Section 4. Exactly as for the call-by-value case, computing the normal form of (the graph representation of) any term takes time polynomial in the number of reduction steps to normal form:

Theorem 5.6. *There is a polynomial $p : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that for every λ -term M , the normal form of $[[M]_\Psi]_\Xi$ can be computed in time at most $p(|M|, \text{Time}_h(M))$.*

On the other hand, we cannot hope to *directly* reuse the results in Section 3 when proving the existence of an embedding of OCRSs into weak call-by-name λ -calculus: the same λ -term can have distinct normal forms in the two cases. It is widely known, however, that a continuation-passing translation can be used to simulate call-by-value reduction by call-by-name reduction [Pl075]. The only missing tale is about the relative performances: do terms obtained via the CPS translation reduce (in call-by-name) to their normal forms in a number of steps which is *comparable* to the number of (call-by-value) steps to normal

form for the original terms? We conjecture the answer is “yes”, but we leave the task of proving that to a future work.

6. CONCLUSIONS

We have shown that the most naïve cost models for weak call-by-value and call-by-name λ -calculus (each beta-reduction step has unitary cost) and orthogonal constructor term rewriting (each rule application has unitary cost) are linearly related. Since, in turn, this cost model for λ -calculus is polynomially related to the actual cost of reducing a λ -term on a Turing machine, the two machine models we considered are both *reasonable* machines, when endowed with their natural, intrinsic cost models (see also Gurevich’s opus on Abstract State Machine simulation “at the same level of abstraction”, e.g. [Gur01]). This strong (the embeddings we consider are compositional), complexity-preserving equivalence between a first-order and a higher-order model is the most important technical result of the paper.

Ongoing and future work includes the investigation of how much of this simulation could be recovered either in a typed setting (see [SU99] for some of the difficulties), or in the case of λ -calculus with strong reduction, where we reduce under an abstraction. Novel techniques have to be developed, since the analysis of the present paper cannot be easily extended to these cases.

REFERENCES

- [BC92] Stephen Bellantoni and Stephen Cook. A new recursion-theoretic characterization of the polytime functions. *Computational Complexity*, 2:97–110, 1992.
- [BEG⁺86] H. Barendregt, M. Eekelen, J. Glauert, J. Kennaway, M. Plasmeijer, and M. Sleep. Term graph rewriting. In J. de Bakker, A. Nijman, and P. Treleaven, editors, *Volume II: Parallel Languages on PARLE: Parallel Architectures and Languages Europe*, pages 141–158. Springer-Verlag, 1986.
- [DLM08] Ugo Dal Lago and Simone Martini. The weak lambda-calculus as a reasonable machine. *Theoretical Computer Science*, 398:32–50, 2008.
- [DLM09] Ugo Dal Lago and Simone Martini. On constructor rewrite systems and the lambda-calculus. In *Automata, Languages and Programming, 36th International Colloquium, Proceedings*, volume 5556 of *LNCS*, pages 163–174. Springer, 2009.
- [DLM10] Ugo Dal Lago and Simone Martini. Derivational complexity is an invariant cost model. In *Foundational and Practical Aspects of Resource Analysis, First International Workshop, Proceedings*, volume 6324 of *LNCS*, pages 88–101. Springer, 2010.
- [Gir98] Jean-Yves Girard. Light linear logic. *Information and Computation*, 143(2):175–204, 1998.
- [Gur01] Yuri Gurevich. The sequential ASM thesis. In *Current trends in theoretical computer science*, pages 363–392. World Scientific, 2001.
- [Jon87] Simon Peyton Jones. *The Implementation of Functional Programming Languages*. Prentice Hall, 1987.
- [Lei95] Daniel Leivant. Ramified recurrence and computational complexity I: word recurrence and polytime. In *Feasible Mathematics II*, pages 320–343. Birkhäuser, 1995.
- [MM00] Jean-Yves Marion and Jean-Yves Moya. Efficient first order functional program interpreter with time bound certifications. In *Logic for Programming and Automated Reasoning, 7th International Conference, Proceedings*, volume 1955 of *LNCS*, pages 25–42. Springer, 2000.
- [Par90] Michel Parigot. On the representation of data in lambda-calculus. In *Computer Science Logic, 3rd International Workshop, Proceedings*, volume 440 of *LNCS*, pages 309–321. Springer, 1990.
- [Plo75] Gordon D. Plotkin. Call-by-name, call-by-value and the lambda-calculus. *Theoretical Computer Science*, 1(2):125–159, 1975.
- [Plu90] Detlef Plump. Graph-reducible term rewriting systems. In *Graph-Grammars and Their Application to Computer Science*, volume 532 of *LNCS*, pages 622–636. Springer, 1990.

- [PR93] Michel Parigot and Paul Rozière. Constant time reductions in lambda-calculus. In *Mathematical Foundations of Computer Science 1993, 18th International Symposium, Proceedings*, volume 711 of LNCS, pages 608–617. Springer, 1993.
- [SGM02] D. Sands, J. Gustavsson, and A. Moran. Lambda calculi and linear speedups. In *The Essence of Computation: Complexity, Analysis, Transformation. Essays Dedicated to Neil D. Jones*, number 2566 in LNCS, pages 60–82. Springer, 2002.
- [SU99] Zdzislaw Splawski and Pawel Urzyczyn. Type fixpoints: Iteration vs. recursion. In *Functional Programming, 4th International Conference, Proceedings*, pages 102–113. ACM, 1999.
- [vEB90] Peter van Emde Boas. Machine models and simulation. In *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity (A)*, pages 1–66. MIT Press, 1990.
- [Wad80] Christopher Wadsworth. Some unusual λ -calculus numeral systems. In J.P. Seldin and J.R. Hindley, editors, *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*. Academic Press, 1980.