The Useful MAM, a Reasonable Implementation of the Strong λ -Calculus

Beniamino Accattoli

► To cite this version:

Beniamino Accattoli. The Useful MAM, a Reasonable Implementation of the Strong λ -Calculus. 23rd International Workshop on Logic, Language, Information, and Computation (WoLLIC 2016), Aug 2016, Puebla, Mexico. pp.1 - 21, 10.1007/978-3-662-52921-8_1 . hal-01425534

HAL Id: hal-01425534

<https://hal.inria.fr/hal-01425534>

Submitted on 3 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Useful MAM, a Reasonable Implementation of the Strong λ -Calculus

Beniamino Accattoli

INRIA & LIX, École Polytechnique
`beniamino.accattoli@inria.fr`

Abstract. It has been a long-standing open problem whether the strong λ -calculus is a reasonable computational model, i.e. whether it can be implemented within a polynomial overhead with respect to the number of β -steps on models like Turing machines or RAM. Recently, Accattoli and Dal Lago solved the problem by means of a new form of sharing, called *useful sharing*, and realised via a calculus with explicit substitutions. This paper presents a new abstract machine for the strong λ -calculus based on useful sharing, the *Useful Milner Abstract Machine*, and proves that it reasonably implements leftmost-outermost evaluation. It provides both an alternative proof that the λ -calculus is reasonable and an improvement on the technology for implementing strong evaluation.

1 Introduction

The higher-order computational model of reference is the λ -calculus, that comes in two variants, *weak* or *strong*. Introduced at the inception of computer science as a mathematical approach to computation, it later found applications in the theoretical modelling of programming languages and, more recently, proof assistants. The weak λ -calculus is the backbone of functional languages such as LISP, Scheme, OCAML, or Haskell. It is *weak* because evaluation does not enter function bodies and, usually, terms are assumed to be closed. By removing these restrictions one obtains the *strong* λ -calculus, that underlies proof assistants like Coq, Isabelle, and Twelf, or higher-order logic programming languages such as λ -prolog or the Edinburgh Logical Framework. Higher-order features nowadays are also part of mainstream programming languages like Java or Python.

The abstract, mathematical character is both the advantage and the drawback of the higher-order approach. The advantage is that it enhances the modularity and the conciseness of the code, allowing to forget about low-level details at the same time. The drawback is that the distance from low-level details makes its complexity harder to analyse, in particular its main computational rule, called *β -reduction*, at first sight is not an atomic operation. In particular, β can be nasty, and make the program grow at an exponential rate. The number of β -steps, then, does not even account for the time to write down the result, suggesting that it is not a reasonable cost model. This is the *size-explosion problem* [6], and affects both the weak and the strong λ -calculus.

The λ -Calculus is Reasonable, Indeed A cornerstone of the theory is that, nonetheless, in the weak λ -calculus the number of β -steps *is* a reasonable cost model for time complexity analyses [9,25,14], where *reasonable* formally means that it is polynomially related to the cost model of RAM or Turing machines.

For the strong λ -calculus, the techniques developed for the weak one do not work, as wilder forms of size-explosion are possible. A natural candidate cost model from the theory of λ -calculus is the number of (Lévy) optimal parallel steps, but it has been shown by Asperti and Mairson that such a cost model is not reasonable [8].

It is only very recently that the strong case has been solved by Accattoli and Dal Lago [6], who showed that the number of leftmost-outermost β -steps to full normal form is a reasonable cost model. The proof of this result relies on two theoretical tools. First, the *Linear Substitution Calculus* (LSC), an expressive and simple decomposition of the λ -calculus via linear logic and rewriting theory, developed by Accattoli and Kesner [3] as a variation over a calculus by Robin Milner [24]. Second, *useful sharing*, a new form of shared evaluation introduced by Accattoli and Dal Lago on top of the LSC. Roughly, the LSC is a calculus where the meta-level operation of substitution used by β -reduction is internalised and decomposed in micro steps, *i.e.* it is what is usually called a calculus with *explicit substitutions*. The further step is to realise that some of these micro substitution steps are useless: they do not lead to the creation of other β -redexes, their only aim is to unshare the result and provide the full normal form. Useful evaluation then performs only those substitution steps that are useful, *i.e.* not useless. By avoiding useless unsharing steps, it computes a shared representation of the normal form of size linear in the number of steps, whose unsharing may cause an exponential blow up in size. This is how the size-explosion problem is circumvented, see [6] for more explanations.

This Paper In this paper we provide an alternative proof that the strong λ -calculus is reasonable (actually only of the hard half, that is the simulation of λ -calculus on RAM, the other half being much easier, see [5]), by replacing the LSC with the *Useful Milner Abstract Machine*. The aim of the paper is threefold:

1. *Getting Closer To Implementations*: the LSC decomposes β -reduction in micro-steps but omits details about the search for the next redex to reduce. Moreover, in [6] useful sharing is used as a sort of black box on top of the LSC. Switching to abstract machines provides a solution closer to implementations and internalises useful sharing.
2. *The First Reasonable Strong Abstract Machine*: the literature on abstract machines for strong evaluation is scarce (see below) and none of the machines in the literature is reasonable. This work thus provides an improvement of the technology for implementing strong evaluation.
3. *Alternative Proof*: the technical development in [6] is sophisticated, because a second aim of that paper is to connect some of the used tools (namely useful sharing and the subterm property) with the seemingly unrelated notion of *standardisation* from rewriting theory. Here we provide a more basic, down-to-earth approach, not relying on advanced rewriting theory.

The Useful MAM The Milner Abstract Machine (MAM) is a variant with just one *global environment* of the Krivine Abstract Machine (KAM), introduced in [1] by Accattoli, Barenbaum, and Mazza. The same authors introduce in [2] the Strong MAM, *i.e.* the extension of the MAM to strong evaluation, that is a version with just one global environment of Cregut’s Strong KAM [13], essentially the only other abstract machine for strong (call-by-name) evaluation in the literature. Both are not reasonable. The problem is that these machines do not distinguish between useful and useless steps.

The Useful MAM introduced in this paper improves the situation, by refining the Strong MAM. The principle is quite basic, let us sketch it. Whenever a β -redex $(\lambda x.t)u$ is encountered, the Strong MAM adds an entry $[x \leftarrow u]$ to the environment E . The Useful MAM, additionally, executes an auxiliary machine on u —the Checking Abstract Machine (Checking AM)—to establish its usefulness. The result of this check is a label l that is attached to the entry $[x \leftarrow u]^l$. Later on, when an occurrence of x is found, the Useful MAM replaces x with u only if the label on $[x \leftarrow u]^l$ says that it is useful. Otherwise the machine backtracks, to search for the next redex to reduce.

The two results of the paper are:

1. *Qualitative* (Theorem 2): the Useful MAM correctly and completely implements leftmost-outermost (LO for short) β -evaluation—formally, the two are weakly bisimilar.
2. *Quantitative* (Theorem 5): the Useful MAM is a reasonable implementation, *i.e.* the work done by both the Useful MAM and the Checking AM is polynomial in the number of LO β -steps and in the size of the initial term.

Related Work Beyond Crégut’s [12,13] and Accattoli, Barenbaum, and Mazza’s [2], we are aware of only two other works on strong abstract machines, García-Pérez, Nogueira and Moreno-Navarro’s [22] (2013), and Smith’s [27] (unpublished, 2014). Two further studies, de Carvalho’s [11] and Ehrhard and Regnier’s [19], introduce strong versions of the KAM but for theoretical purposes; in particular, their design choices are not tuned towards implementations (*e.g.* rely on a naïve parallel exploration of the term). Semi-strong machines for call-by-value (*i.e.* dealing with weak evaluation but on open terms) are studied by Grégoire and Leroy [23] and in a recent work by Accattoli and Sacerdoti Coen [4] (see [4] for a comparison with [23]). More recent work by Dénès [18] and Boutiller [10] appeared in the context of term evaluation in Coq. None of the machines for strong evaluation in the literature is reasonable, in the sense of being polynomial in the number of β -steps. The machines developed by Accattoli and Sacerdoti Coen in [4] are reasonable, but they are developed in a semi-strong setting only. Another difference between [4] and this work is that call-by-value simplifies the treatment of usefulness because it allows to compute the labels for usefulness while evaluating the term, that is not possible in call-by-name.

Global environments are explored by Fernández and Sifakakis in [20], and used in a minority of works, *e.g.* [25,17]. Here we use the terminology for abstract machines coming from the distillation technique in [1], related to the *refocusing*

semantics of Danvy and Nielsen [16] and introduced to revisit the relationship between the KAM and weak linear head reduction pointed out by Danos and Regnier [15]. We do not, however, employ the distillation technique itself.

Proofs All proofs have been omitted. Those of the main lemmas and theorems concerning the Useful MAM can be found in the appendix. The other ones can be found in the longer version on the author's web page.

2 λ -Calculus and Leftmost-Outermost Evaluation

The syntax of the λ -calculus is given by the following grammar for terms:

$$\lambda\text{-TERMS} \quad t, u, w, r ::= x \mid \lambda x. t \mid tu.$$

We use $t\{x \leftarrow u\}$ for the usual (meta-level) notion of substitution. An abstraction $\lambda x. t$ binds x in t , and we silently work modulo α -equivalence of bound variables, *e.g.* $(\lambda y. (xy))\{x \leftarrow y\} = \lambda z. (yz)$. We use $\text{fv}(t)$ for the set of free variables of t .

Contexts. One-hole contexts C and the *plugging* $C\langle t \rangle$ of a term t into a context C are defined by:

CONTEXTS	PLUGGING
$C ::= \langle \cdot \rangle \mid \lambda x. C \mid Ct \mid tC$	$\langle \cdot \rangle \langle t \rangle := t \quad (Cu) \langle t \rangle := C \langle t \rangle u$ $(\lambda x. C) \langle t \rangle := \lambda x. C \langle t \rangle \quad (uC) \langle t \rangle := uC \langle t \rangle$

As usual, plugging in a context can capture variables, *e.g.* $(\lambda y. (\langle \cdot \rangle y)) \langle y \rangle = \lambda y. (yy)$. The plugging $C\langle C' \rangle$ of a context C' into a context C is defined analogously. A context C is *applicative* if $C = C' \langle \langle \cdot \rangle \bar{u} \rangle$ for some C' and $\bar{u}$.

We define β -reduction $\rightarrow_\beta$ as follows:

RULE AT TOP LEVEL	CONTEXTUAL CLOSURE
$(\lambda x. t)u \mapsto_\beta t\{x \leftarrow u\}$	$C \langle t \rangle \rightarrow_\beta C \langle u \rangle \quad \text{if } t \mapsto_\beta u$

A term t is a *normal form*, or simply *normal*, if there is no u such that $t \rightarrow_\beta u$, and it is *neutral* if it is normal and it is not of the form $\lambda x. u$ (*i.e.* it is not an abstraction). The *position* of a β -redex $C \langle t \rangle \rightarrow_\beta C \langle u \rangle$ is the context C in which it takes place. To ease the language, we will identify a redex with its position. A *derivation* $d : t \rightarrow^k u$ is a finite, possibly empty, sequence of reduction steps. We write $|t|$ for the size of t and $|d|$ for the length of d .

Leftmost-Outermost Derivations The left-to-right outside-in order on redexes is expressed as an order on positions, *i.e.* contexts.

Definition 1 (Left-to-Right Outside-In Order).

1. The outside-in order:
 - (a) Root: $\langle \cdot \rangle \prec_O C$ for every context $C \neq \langle \cdot \rangle$;
 - (b) Contextual closure: If $C \prec_O C'$ then $C'' \langle C \rangle \prec_O C'' \langle C' \rangle$ for any C'' .
2. The left-to-right order: $C \prec_L C'$ is defined by:

- (a) Application: If $C \prec_p t$ and $C' \prec_p u$ then $Cu \prec_L tC'$;
- (b) Contextual closure: If $C \prec_L C'$ then $C''\langle C \rangle \prec_L C''\langle C' \rangle$ for any C'' .
- 3. The left-to-right outside-in order: $C \prec_{LO} C'$ if $C \prec_O C'$ or $C \prec_L C'$:

The following are a few examples. For every context C , it holds that $\langle \cdot \rangle \not\prec_L C$. Moreover $(\lambda x. \langle \cdot \rangle)t \prec_O (\lambda x. (\langle \cdot \rangle u))t$ and $(\langle \cdot \rangle)t u \prec_L (wt)\langle \cdot \rangle$.

Definition 2 (LO β -Reduction). Let t be a λ -term and C a redex of t . C is the leftmost-outermost β -redex (LO β for short) of t if $C \prec_{LO} C'$ for every other β -redex C' of t . We write $t \rightarrow_{LO\beta} u$ if a step reduces the LO β -redex.

The next immediate lemma guarantees that we defined a total order.

Lemma 1 (Totality of $\prec_{LO}$). If $C \prec_p t$ and $C' \prec_p t$ then either $C \prec_{LO} C'$ or $C' \prec_{LO} C$ or $C = C'$. Therefore, $\rightarrow_{LO\beta}$ is deterministic.

LO Contexts For the technical development of the paper we need two characterisations of when a context is the position of the LO β -redex. The first, following one, is used in the proofs of Lemma 5.2 and Lemma 6.4.

Definition 3 (LO Contexts). A context C is LO if

- 1. Right Application: whenever $C = C'\langle tC'' \rangle$ then t is neutral, and
- 2. Left Application: whenever $C = C'\langle C''t \rangle$ then $C'' \neq \lambda x.C'''$.

The second characterisation is inductive, and it used to prove Lemma 10.3.

Definition 4 (iLO Context). Inductive LO β (or iLO) contexts are defined by induction as follows:

$$\begin{array}{c} \frac{}{\langle \cdot \rangle \text{ is iLO}} (ax\text{-iLO}) \quad \frac{C \text{ is iLO} \quad C \neq \lambda x.C'}{Ct \text{ is iLO}} (@l\text{-iLO}) \\ \frac{C \text{ is iLO}}{\lambda x.C \text{ is iLO}} (\lambda\text{-iLO}) \quad \frac{t \text{ is neutral} \quad C \text{ is iLO}}{tC \text{ is iLO}} (@r\text{-iLO}) \end{array}$$

As expected,

Lemma 2 ($\rightarrow_{LO\beta}$ -steps and Contexts). Let t be a λ -term and C a redex in t . C is the LO β redex in t iff C is LO iff C is iLO.

3 Preliminaries on Abstract Machines

We study two abstract machines, the Useful MAM (Fig. 4) and an auxiliary machine called the Checking AM (Fig. 2).

The Useful MAM is meant to implement LO β -reduction strategy via a decoding function $\underline{\cdot}$ mapping machine states to λ -terms. Machine states s are given by a code $\bar{t}$, that is a λ -term t not considered up to α -equivalence (which is why it is over-lined), and some data-structures like stacks, frames, and environments. The data-structures are used to implement the search for the next LO-redex and

a form of micro-steps substitution, and they decode to evaluation contexts for $\rightarrow_{\text{LO}\beta}$. Every state s decodes to a term $\underline{s}$, having the shape $C_s\langle\bar{t}\rangle$, where $\bar{t}$ is the code currently under evaluation and C_s is the evaluation context given by the data-structures.

The Checking AM tests the usefulness of a term (with respect to a given environment) and outputs a label with the result of the test. It uses the same states and data-structures of the Useful MAM.

The Data-Structures First of all, our machines are executed on *well-named* terms, that are those α -representants where all variables (both bound and free) have distinct names. Then, the data-structures used by the machines are defined in Fig. 1, namely:

- *Stack* π : it contains the arguments of the current code;
- *Frame* F : a second stack, that together with π is used to walk through the term and search for the next redex to reduce. The items ϕ of a frame are of two kinds. A variable x is pushed on the frame F whenever the machines starts evaluating under an abstraction λx . A *head argument context* $\bar{t}\Diamond\pi$ is pushed every time evaluation enters in the right subterm $\bar{u}$ of an application $\bar{t}\bar{u}$. The entry saves the left part $\bar{t}$ of the application and the current stack π , to restore them when the evaluation of the right subterm $\bar{u}$ is over.
- *Global Environment* E : it is used to implement micro-step evaluation (*i.e.* the substitution on a variable occurrence at the time), storing the arguments of β -redexes that have been encountered so far. Most of the literature on abstract machines uses *local environments* and *closures*. Having just one global environment E removes the need for closures and simplifies the machine. On the other hand, it forces to use explicit α -renamings (the operation $\bar{t}^\alpha$ in $\rightsquigarrow_{\mathbf{e}_{red}}$ and $\rightsquigarrow_{\mathbf{e}_{abs}}$ in Fig. 4), but this does not affect the overall complexity, as it speeds up other operations, see [1]. The entries of E are of the form $[x\leftarrow\bar{t}]^l$, *i.e.* they carry a label l used to implement usefulness, to be explained later on in this section. We write $E(x) = [x\leftarrow\bar{t}]^l$ when E contains $[x\leftarrow\bar{t}]^l$ and $E(x) = \perp$ when in E there are no entries of the form $[x\leftarrow\bar{t}]^l$.

The Decoding Every state s decodes to a term $\underline{s}$ (see Fig. 3), having the shape $C_s\langle\bar{t}\downarrow_E\rangle$, where

- $\bar{t}\downarrow_E$ is a λ -term, roughly obtained by applying to the code the substitution induced by the global environment E . More precisely, the operation $\bar{t}\downarrow_E$ is called *unfolding* and it is properly defined at the end of this section.
- C_s is a context, that will be shown to be a LO context, obtained by decoding the stack π and the dump F and applying the unfolding. Note that, to improve readability, π is decoded in postfix notation for plugging.

The Transitions According to the distillation approach of [1] we distinguish different kinds of transitions, whose names reflect a proof-theoretical view, as machine transitions can be seen as cut-elimination steps [7,1]:

- *Multiplicatives* $\rightsquigarrow_{\mathbf{m}}$: they fire a β -redex, except that if the argument is not a variable then it is not substituted but added to the environment;
- *Exponentials* $\rightsquigarrow_{\mathbf{e}}$: they perform a clashing-avoiding substitution from the environment on the single variable occurrence represented by the current code. They implement micro-step substitution.
- *Commutatives* $\rightsquigarrow_{\mathbf{c}}$: they locate and expose the next redex according to the LO evaluation strategy, by rearranging the data-structures.

Both exponential and commutative transitions are invisible on the λ -calculus. Garbage collection is here simply ignored, or, more precisely, it is encapsulated at the meta-level, in the decoding function.

Labels for Useful Sharing A label l for a code in the environment can be of three kinds. Roughly, they are:

- *Neutral*, or $l = \text{neu}$: it marks a neutral term, that is always useless as it is β -normal and its substitution cannot create a redex, because it is not an abstraction;
- *Abstraction*, or $l = \text{abs}$: it marks an abstraction, that is a term that is at times useful to substitute. If the variable that it is meant to replace is applied, indeed, the substitution of the abstraction creates a β -redex. But if it is not applied, it is useless.
- *Redex*, or $l = \text{red}$: it marks a term that contains a β -redex. It is always useful to substitute these terms.

Actually, the explanation we just gave is oversimplified, but it provides a first intuition about labels. In fact in an environment $[x \leftarrow \bar{t}]^l : E$ it is not really $\bar{t}$ that has the property mentioned by its label, rather the term $\bar{t} \downarrow_E$ obtained by unfolding the rest of the environment on $\bar{t}$. The idea is that $[x \leftarrow \bar{t}]^{\text{red}}$ states that it is useful to substitute $\bar{t}$ to *later on* obtain a redex inside it (by potential further substitutions on its variables coming from E). The precise meaning of the labels will be given by Definition 6, and the properties they encode will be made explicit by Lemma 11.

A further subtlety is that the label *red* for redexes is refined as a pair (red, n) , where n is the number of substitutions in E that are needed to obtain the LO redex in $\bar{t} \downarrow_E$. Our machines never inspect these numbers, they are only used for the complexity analysis of Sect. 5.2.

Grafting and Unfoldings The unfolding of the environment E on a code $\bar{t}$ is defined as the recursive *capture-allowing* substitution (called *grafting*) of the entries of E on $\bar{t}$.

Definition 5 (Grafting and Environment Unfolding). *The operation of grafting $\bar{t}\{\{x \leftarrow \bar{u}\}\}$ is defined by*

$$\begin{aligned} (\overline{w\bar{r}})\{\{x \leftarrow \bar{u}\}\} &:= \overline{w}\{\{x \leftarrow \bar{u}\}\}\bar{r}\{\{x \leftarrow \bar{u}\}\} & (\lambda y.\overline{w})\{\{x \leftarrow \bar{u}\}\} &:= \lambda y.\overline{w}\{\{x \leftarrow \bar{u}\}\} \\ x\{\{x \leftarrow \bar{u}\}\} &:= \bar{u} & y\{\{x \leftarrow \bar{u}\}\} &:= y \end{aligned}$$

Frames	$F ::= \epsilon \mid F : \phi$	Stacks	$\pi ::= \epsilon \mid \bar{t} : \pi$
Frame Items	$\phi ::= \bar{t} \Diamond \pi \mid x$	Phases	$\varphi ::= \blacktriangledown \mid \blacktriangle$
Labels	$l ::= \text{abs} \mid (\text{red}, n \in \mathbb{N}) \mid \text{neu}$	Environments	$E ::= \epsilon \mid [x \leftarrow \bar{t}]^l : E$

Fig. 1: Grammars.

Given an environment E we define the unfolding of E on a code $\bar{t}$ as follows:

$$\bar{t} \downarrow_{\epsilon} := \bar{t} \quad \bar{t} \downarrow_{[x \leftarrow \bar{u}]^l : E} := \bar{t} \{ \{x \leftarrow \bar{u}\} \} \downarrow_E$$

or equivalently as:

$$\begin{aligned} (\bar{u}\bar{w}) \downarrow_E &:= \bar{u} \downarrow_E \bar{w} \downarrow_E & x \downarrow_{[x \leftarrow \bar{u}]^l : E'} &:= \bar{u} \downarrow_{E'} \\ (\lambda x. \bar{u}) \downarrow_E &:= \lambda x. \bar{u} \downarrow_E & x \downarrow_{[y \leftarrow \bar{u}]^l : E'} &:= x \downarrow_{E'} & x \downarrow_{\epsilon} &:= x \end{aligned}$$

For instance, $(\lambda x. y) \downarrow_{[y \leftarrow x x]^{neu}} = \lambda x. (xx)$. The unfolding is extended to contexts as expected (*i.e.* recursively propagating the unfolding and setting $\langle \cdot \rangle \downarrow_E = E$).

Let us explain the need for grafting. In [2], the Strong MAM is decoded to the LSC, that is a calculus with explicit substitutions, *i.e.* a calculus able to represent the environment of the Strong MAM. Matching the representation of the environment on the Strong MAM and on the LSC does not need grafting but it is, however, a quite technical affair. Useful sharing adds many further complications in establishing such a matching, because useful evaluation computes a shared representation of the normal form and forces some of the explicit substitutions to stay under abstractions. The difficulty is such, in fact, that we found much easier to decode directly to the λ -calculus rather than to the LSC. Such an alternative solution, however, has to push the substitution induced by the environment through abstractions, which is why we use grafting.

Lemma 3 (Properties of Grafting and Unfolding).

1. If the bound names of t do not appear free in u then $t\{x \leftarrow u\} = t\{\{x \leftarrow u\}\}$.
2. If moreover they do not appear free in E then $t \downarrow_E \{x \leftarrow u \downarrow_E\} = t\{x \leftarrow u\} \downarrow_E$.

4 The Checking Abstract Machine

The Checking Abstract Machine (Checking AM) is defined in Fig. 2. It starts executions on states of the form $(\epsilon, \bar{t}, \epsilon, E, \blacktriangledown)$, with the aim of checking the usefulness of $\bar{t}$ with respect to the environment E , *i.e.* it walks through $\bar{t}$ and whenever it encounters a variable x it looks up its usefulness in E .

The Checking AM has six commutative transitions, noted $\rightarrow_{c_i}$ with $i = 1, \dots, 6$, used to walk through the term, and five output transitions, noted $\rightarrow_{o_j}$ with $j = 1, \dots, 5$, that produce the value of the test for usefulness, to be later used by the Useful MAM. The exploration is done in two alternating phases, evaluation $\blacktriangledown$ and backtracking $\blacktriangle$. Evaluation explores the current code towards

Frame	Code	Stack	Env	Ph		Frame	Code	Stack	Env	Ph
F	$\bar{t}\bar{u}$	π	E	$\blacktriangledown$	$\rightarrow_{c_1}$	F	$\bar{t}$	$\bar{u} : \pi$	E	$\blacktriangledown$
F	$\lambda x.\bar{t}$	$\bar{u} : \pi$	E	$\blacktriangledown$	$\rightarrow_{o_1}$	output ($red, 1$)				
F	$\lambda x.\bar{t}$	ϵ	E	$\blacktriangledown$	$\rightarrow_{c_2}$	$F : x$	$\bar{t}$	ϵ	E	$\blacktriangledown$
F	x	π	E	$\blacktriangledown$	$\rightarrow_{o_2}$	output ($red, n+1$)				
						if $E(x) = [x \leftarrow \bar{t}]^{(red, n)}$				
F	x	$\bar{u} : \pi$	E	$\blacktriangledown$	$\rightarrow_{o_3}$	output ($red, 2$)				
						if $E(x) = [x \leftarrow \bar{t}]^{abs}$				
F	x	π	E	$\blacktriangledown$	$\rightarrow_{c_3}$	F	x	π	E	$\blacktriangle$
						if $E(x) = \perp$ or $E(x) = [x \leftarrow \bar{t}]^{neu}$ or $(E(x) = [x \leftarrow \bar{t}]^{abs} \text{ and } \pi = \epsilon)$				
$F : x$	$\bar{t}$	ϵ	E	$\blacktriangle$	$\rightarrow_{c_4}$	F	$\lambda x.\bar{t}$	ϵ	E	$\blacktriangle$
$F : \bar{t} \diamond \pi$	$\bar{u}$	ϵ	E	$\blacktriangle$	$\rightarrow_{c_5}$	F	$\bar{t}\bar{u}$	π	E	$\blacktriangle$
F	$\bar{t}$	$\bar{u} : \pi$	E	$\blacktriangle$	$\rightarrow_{c_6}$	$F : \bar{t} \diamond \pi$	$\bar{u}$	ϵ	E	$\blacktriangledown$
ϵ	$\bar{t}\bar{u}$	ϵ	E	$\blacktriangle$	$\rightarrow_{o_4}$	output neu				
ϵ	$\lambda x.\bar{t}$	ϵ	E	$\blacktriangle$	$\rightarrow_{o_5}$	output abs				

Fig. 2: The Checking Abstract Machine (Checking AM).

$$\begin{array}{l}
\begin{array}{l}
\epsilon := \langle \cdot \rangle \\
\bar{u} : \pi := \langle \langle \cdot \rangle \bar{u} \rangle \pi \\
F : \bar{t} \diamond \pi := F \langle \langle \bar{t} \rangle \rangle \pi \\
F : x := F \langle \lambda x. \langle \cdot \rangle \rangle
\end{array}
\quad
\begin{array}{l}
C_s := F \langle \pi \rangle \downarrow_E \\
s := F \langle \langle \bar{t} \rangle \pi \rangle \downarrow_E = C_s \langle \bar{t} \rangle \downarrow_E \\
\text{where } s = (F, \bar{t}, \pi, E)
\end{array}
\end{array}$$

Fig. 3: Decoding.

the head, storing in the stack and in the frame the parts of the code that it leaves behind. Backtracking comes back to an argument that was stored in the frame, when the current head has already been checked. Note that the Checking AM never modifies the environment, it only looks it up.

Let us explain the transitions. First the commutative ones:

- $\rightarrow_{c_1}$: the code is an application $\bar{t}\bar{u}$ and the machine starts exploring the left subterm $\bar{t}$, storing $\bar{u}$ on top of the stack π .
- $\rightarrow_{c_2}$: the code is an abstraction $\lambda x.\bar{t}$ and the machine goes under the abstraction, storing x on top of the frame F .
- $\rightarrow_{c_3}$: the machine finds a variable x that either has no associated entry in the environment (if $E(x) = \perp$) or its associated entry $[x \leftarrow \bar{t}]^l$ in the environment is useless. This can happen if either $l = neu$, *i.e.* substituting $\bar{t}$ would only lead to a neutral term, or $l = abs$, *i.e.* substituting $\bar{t}$ would provide an abstraction, but the stack is empty, and so it is useless to substitute the abstraction because no β -redexes will be obtained. Thus the machine switches to the backtracking phase ($\blacktriangle$), whose aim is to undo the frame to obtain a new subterm to explore.
- $\rightarrow_{c_4}$: it is the inverse of $\rightarrow_{c_2}$, it puts back on the code an abstraction that was previously stored in the frame.
- $\rightarrow_{c_5}$: backtracking from the evaluation of an argument $\bar{u}$, it restores the application $\bar{t}\bar{u}$ and the stack π that were previously stored in the frame.

- $\rightarrow_{\blacktriangle c_6}$: backtracking from the evaluation of the left subterm $\bar{t}$ of an application $\bar{t}\bar{u}$, the machine starts evaluating the right subterm (by switching to the evaluation phase $\blacktriangledown$) with an empty stack ϵ , storing on the frame the pair $\bar{t}\blacktriangledown\pi$ of the left subterm and the previous stack π .

Then the output transitions:

- $\rightarrow_{o_1}$: the machine finds a β -redex, namely $(\lambda x.\bar{t})\bar{u}$ and thus outputs a label saying that it requires only one substitution step (namely substituting the term the machine was executed on) to eventually find a β -redex.
- $\rightarrow_{o_2}$: the machine finds a variable x whose associated entry $[x \leftarrow \bar{t}]^{(red,n)}$ in the environment is labeled with (red, n) , and so outputs a label saying that it takes $n + 1$ substitution steps to eventually find a β -redex (n plus 1 for the term the machine was executed on).
- $\rightarrow_{o_3}$: the machine finds a variable x whose associated entry $[x \leftarrow \bar{t}]^{abs}$ in the environment is labeled with abs , so $\bar{t}$ is an abstraction, and the stack is non-empty. Since substituting the abstraction will create a β -redex, the machine outputs a label saying that it takes two substitution steps to obtain a β -redex, one for the term the machine was executed on and one for the abstraction $\bar{t}$.
- $\rightarrow_{o_4}$: the machine went through the whole term, that is an application, and found no redex, nor any redex that can be obtained by substituting from the environment. Thus that term is neutral and so the machine outputs the corresponding label.
- $\rightarrow_{o_5}$: as for the previous transition, except that the term is an abstraction, and so the output is the abs label.

The fact that commutative transitions only walk through the code, without changing anything, is formalised by the following lemma, that is crucial for the proof of correctness of the Checking AM (forthcoming Theorem 1).

Lemma 4 (Commutative Transparency).

Let $s = (F, \bar{u}, \pi, E, \varphi) \rightsquigarrow_{c_{1,2,3,4,5,6}} (F', \bar{u}', \pi', E, \varphi') = s'$. Then

1. Decoding Without Unfolding: $\underline{F}\langle\langle\bar{u}\rangle\pi\rangle = \underline{F'}\langle\langle\bar{u}'\rangle\pi'\rangle$, and
2. Decoding With Unfolding: $\underline{s} = \underline{s'}$.

For the analysis of the properties of the Checking AM we need a notion of well-labeled environment, *i.e.* of environment where the labels are consistent with their intended meaning. It is a technical notion also providing enough information to perform the complexity analysis, later on. Moreover, it includes two structural properties of environments: 1) in $[x \leftarrow \bar{t}]^l$ the code $\bar{t}$ cannot be a variable, and 2) there cannot be two entries associated to the same variables.

Definition 6 (Well-Labeled Environments). Well-labeled global environments E are defined by

1. Empty: ϵ is well-labeled;
2. Inductive: $[x \leftarrow \bar{t}]^l : E'$ is well-labeled if E' is well-labeled, x is fresh with respect to $\bar{t}$ and E' , and

- (a) Abstractions: if $l = \text{abs}$ then $\bar{t}$ and $\bar{t}|_{E'}$ are normal abstractions;
- (b) Neutral Terms: if $l = \text{neu}$ then $\bar{t}$ is an application and $\bar{t}|_{E'}$ is neutral.
- (c) Redexes: if $l = (\text{red}, n)$ then $\bar{t}$ is not a variable, $\bar{t}|_{E'}$ contains a β -redex. Moreover, $\bar{t} = C\langle\bar{u}\rangle$ with C a LO context and
 - if $n = 1$ then $\bar{u}$ is a β -redex,
 - if $n > 1$ then $\bar{u} = x$ and $E' = E'' : [y \leftarrow \bar{u}]^l : E'''$ with
 - if $n > 2$ then $l = (\text{red}, n - 1)$
 - if $n = 2$ then $l = (\text{red}, 1)$ or $(l = \text{abs}$ and C is applicative).

Remark 1. Note that by the definition it immediately follows that if $E = E' : [x \leftarrow \bar{t}]^{(\text{red}, n)} : E''$ is well-labeled then the length of E'' , and thus of E , is at least n . This fact is used in the proof of Theorem 3.1.

The study of the Checking AM requires some terminology and two invariants. A state s is *initial* if it is of the form $(\epsilon, \bar{t}, \epsilon, E, \varphi)$ with E well-labeled and it is *reachable* if there are an initial state s' and a Checking AM execution $\rho : s' \multimap^* s$. Both invariants are used to prove the correctness of the Checking AM: the *normal form invariant* to guarantee that codes labeled with *neu* and *abs* are indeed normal or neutral, while the *decoding invariant* is used for the redex labels.

Lemma 5 (Checking AM Invariants). *Let $s = F \mid \bar{u} \mid \pi \mid E \mid \varphi$ be a Checking AM reachable state and E be a well-labeled environment.*

1. Normal Form:
 - (a) Backtracking Code: if $\varphi = \blacktriangle$, then $\bar{u}|_{E'}$ is normal, and if π is non-empty, then $\bar{u}|_{E'}$ is neutral;
 - (b) Frame: if $F = F' : \bar{w} \diamond \pi' : F''$, then $\bar{w}|_{E'}$ is neutral.
2. Decoding: C_s is a LO context.

Finally, we can prove the main properties of the Checking AM, *i.e.* that when executed on $\bar{t}$ and E it provides a label l to extend E with a consistent entry for $\bar{t}$ (*i.e.* such that $[x \leftarrow \bar{t}]^l : E$ is well-labeled), and that such an execution takes time linear in the size of $\bar{t}$.

Theorem 1 (Checking AM Properties). *Let $\bar{t}$ be a code and E a global environment.*

1. Determinism and Progress: the Checking AM is deterministic and there always is a transition that applies;
2. Termination and Complexity: the execution of the Checking AM on $\bar{t}$ and E always terminates, taking $O(|\bar{t}|)$ steps, moreover
3. Correctness: if E is well-labeled, x is fresh with respect to E and $\bar{t}$, and l is the output then $[x \leftarrow \bar{t}]^l : E$ is well-labeled.

Frame	Code	Stack	Env	Ph		Frame	Code	Stack	Env	Ph
F	$\bar{t}\bar{u}$	π	E	$\blacktriangledown$	$\rightsquigarrow_{c_1}$	F	$\bar{t}$	$\bar{u} : \pi$	E	$\blacktriangledown$
F	$\lambda x.\bar{t}$	$y : \pi$	E	$\blacktriangledown$	$\rightsquigarrow_{m_1}$	F	$\bar{t}\{x \leftarrow y\}$	π	E	$\blacktriangledown$
F	$\lambda x.\bar{t}$	$\bar{u} : \pi$	E	$\blacktriangledown$	$\rightsquigarrow_{m_2}$	F	$\bar{t}$	π	$[x \leftarrow \bar{u}]^l : E$	$\blacktriangledown$
if $\bar{u}$ is not a variable and l is the output of the Checking AM on $\bar{u}$ and E										
F	$\lambda x.\bar{t}$	ϵ	E	$\blacktriangledown$	$\rightsquigarrow_{c_2}$	$F : x$	$\bar{t}$	ϵ	E	$\blacktriangledown$
F	x	π	E	$\blacktriangledown$	$\rightsquigarrow_{red}$	F	$\bar{t}^\alpha$	π	E	$\blacktriangledown$
if $E(x) = [x \leftarrow \bar{t}]^{(red, n)}$										
F	x	$\bar{u} : \pi$	E	$\blacktriangledown$	$\rightsquigarrow_{abs}$	F	$\bar{t}^\alpha$	$\bar{u} : \pi$	E	$\blacktriangledown$
if $E(x) = [x \leftarrow \bar{t}]^{abs}$										
F	x	π	E	$\blacktriangledown$	$\rightsquigarrow_{c_3}$	F	x	π	E	$\blacktriangle$
if $E(x) = \perp$ or $E(x) = [x \leftarrow \bar{t}]^{neu}$ or $(E(x) = [x \leftarrow \bar{t}]^{abs} \text{ and } \pi = \epsilon)$										
$F : x$	$\bar{t}$	ϵ	E	$\blacktriangle$	$\rightsquigarrow_{c_4}$	F	$\lambda x.\bar{t}$	ϵ	E	$\blacktriangle$
$F : \bar{t} \Diamond \pi$	$\bar{u}$	ϵ	E	$\blacktriangle$	$\rightsquigarrow_{c_5}$	F	$\bar{t}\bar{u}$	π	E	$\blacktriangle$
F	$\bar{t}$	$\bar{u} : \pi$	E	$\blacktriangle$	$\rightsquigarrow_{c_6}$	$F : \bar{t} \Diamond \pi$	$\bar{u}$	ϵ	E	$\blacktriangledown$

$\bar{t}^\alpha$ is any code α -equivalent to $\bar{t}$ such that it is well-named and its bound names are fresh with respect to those in the other machine components.

Fig. 4: The Useful Milner Abstract Machine (Useful MAM).

5 The Useful Milner Abstract Machine

The Useful MAM is defined in Fig. 4. It is very similar to the Checking AM, in particular it has exactly the same commutative transitions, and the same organisation in evaluating and backtracking phases. The difference with respect to the Useful MAM is that the output transitions are replaced by micro-step computational rules that reduce β -redexes and implement useful substitutions. Let us explain them:

- *Multiplicative Transition* $\rightsquigarrow_{m_1}$: when the argument of the β -redex $(\lambda x.\bar{t})y$ is a variable y then it is immediately substituted in $\bar{t}$. This happens because 1) such substitution are not costly and 2) because in this way the environment stays compact, see also Remark 2 at the end of the paper.
- *Multiplicative Transition* $\rightsquigarrow_{m_2}$: if the argument $\bar{u}$ is not a variable then the entry $[x \leftarrow \bar{u}]^l$ is added to the environment. The label l is obtained by running the Checking AM on $\bar{u}$ and E .
- *Exponential Transition* $\rightsquigarrow_{red}$: the environment entry associated to x is labeled with (red, n) thus it is useful to substitute $\bar{t}$. The idea is that in at most n additional substitution steps (shuffled with commutative steps) a β -redex will be obtained. To avoid variable clashes the substitution α -renames $\bar{t}$.
- *Exponential Transition* $\rightsquigarrow_{abs}$: the environment associates an abstraction to x and the stack is non empty, so it is useful to substitute the abstraction (again, α -renaming to avoid variable clashes). Note that if the stack is empty the machine rather backtracks using $\rightsquigarrow_{c_3}$.

The Useful MAM starts executions on *initial states* of the form $(\epsilon, \bar{t}, \epsilon, \epsilon)$, where $\bar{t}$ is such that any two variables (bound or free) have distinct names, and any other component is empty. A state s is *reachable* if there are an initial state s' and a Useful MAM execution $\rho : s' \rightsquigarrow^* s$, and it is *final* if no transitions apply.

5.1 Qualitative Analysis

The results of this subsection are the correctness and completeness of the Useful MAM. Four invariants are required. The *normal form* and *decoding invariants* are exactly those of the Checking AM (and the proof for the commutative transitions is the same). The *environment labels invariant* follows from the correctness of the Checking AM (Theorem 1.2). The *name invariant* is used in the proof of Lemma 7.

Lemma 6 (Useful MAM Qualitative Invariants). *Let $s = F \mid \bar{u} \mid \pi \mid E \mid \varphi$ be a state reachable from an initial term $\bar{t}_0$. Then:*

1. Environment Labels: E is well-labeled.
2. Normal Form:
 - (a) Backtracking Code: if $\varphi = \blacktriangle$, then $\bar{u} \downarrow_E$ is normal, and if π is non-empty, then $\bar{u} \downarrow_E$ is neutral;
 - (b) Frame: if $F = F' : \bar{w} \Diamond \pi' : F''$, then $\bar{w} \downarrow_E$ is neutral.
3. Name:
 - (a) Substitutions: if $E = E' : [x \leftarrow \bar{t}] : E''$ then x is fresh wrt $\bar{t}$ and E'' ;
 - (b) Abstractions and Evaluation: if $\varphi = \blacktriangledown$ and $\lambda x. \bar{t}$ is a subterm of $\bar{u}$, π , or π' (if $F = F' : \bar{w} \Diamond \pi' : F''$) then x may occur only in $\bar{t}$;
 - (c) Abstractions and Backtracking: if $\varphi = \blacktriangle$ and $\lambda x. \bar{t}$ is a subterm of π or π' (if $F = F' : \bar{w} \Diamond \pi' : F''$) then x may occur only in $\bar{t}$.
4. Decoding: C_s is a LO context.

We can now show how every single transition projects on the λ -calculus, and in particular that multiplicative transitions project to LO β -steps.

Lemma 7 (One-Step Weak Simulation, Proof at Page 17). *Let s be a reachable state.*

1. Commutative: if $s \rightsquigarrow_{c_{1,2,3,4,5,6}} s'$ then $\underline{s} = \underline{s}'$;
2. Exponential: if $s \rightsquigarrow_{e_{red}, e_{abs}} s'$ then $\underline{s} = \underline{s}'$;
3. Multiplicative: if $s \rightsquigarrow_{m_1, m_2} s'$ then $\underline{s} \rightarrow_{LO\beta} \underline{s}'$.

We also need to show that the Useful MAM computes β -normal forms.

Lemma 8 (Progress, Proof at Page 18). *Let s be a reachable final state. Then $\underline{s}$ is β -normal.*

The theorem of correctness and completeness of the machine with respect to $\rightarrow_{LO\beta}$ follows. The bisimulation is *weak* because transitions other than $\rightsquigarrow_m$ are invisible on the λ -calculus. For a machine execution ρ we denote with $|\rho|$ (resp. $|\rho|_x$) the number of transitions (resp. x -transitions for $x \in \{m, e, c, \dots\}$) in ρ .

Theorem 2 (Weak Bisimulation, Proof at Page 18). *Let s be an initial Useful MAM state of code $\bar{t}$.*

1. Simulation: for every execution $\rho : s \rightsquigarrow^* s'$ there exists a derivation $d : \underline{s} \rightarrow_{LO\beta}^* \underline{s}'$ such that $|d| = |\rho|_m$;
2. Reverse Simulation: for every derivation $d : \bar{t} \rightarrow_{LO\beta}^* u$ there is an execution $\rho : s \rightsquigarrow^* s'$ such that $\underline{s}' = u$ and $|d| = |\rho|_m$.

5.2 Quantitative Analysis

The complexity analyses of this section rely on two additional invariants of the Useful MAM, the subterm and the environment size invariants.

The subterm invariant bounds the size of the duplicated subterms and it is crucial. For us, $\bar{u}$ is a subterm of $\bar{t}$ if it does so up to variable names, both free and bound. More precisely: define t^- as t in which all variables (including those appearing in binders) are replaced by a fixed symbol $*$. Then, we will consider u to be a subterm of t whenever u^- is a subterm of t^- in the usual sense. The key property ensured by this definition is that the size $|\bar{u}|$ of $\bar{u}$ is bounded by $|\bar{t}|$.

Lemma 9 (Useful MAM Quantitative Invariants). *Let $s = F \mid \bar{u} \mid \pi \mid E \mid \varphi$ be a state reachable by the execution ρ from the initial code $\bar{t}_0$.*

1. Subterm:
 - (a) Evaluating Code: if $\varphi = \blacktriangledown$, then $\bar{u}$ is a subterm of $\bar{t}_0$;
 - (b) Stack: any code in the stack π is a subterm of $\bar{t}_0$;
 - (c) Frame: if $F = F' : \bar{w} \Diamond \pi' : F''$, then any code in π' is a subterm of $\bar{t}_0$;
 - (d) Global Environment: if $E = E' : [x \leftarrow \bar{w}]^l : E''$, then $\bar{w}$ is a subterm of $\bar{t}_0$;
2. Environment Size: the length of the global environment E is bound by $|\rho|_m$.

The proof of the polynomial bound of the overhead is in three steps. First, we bound the number $|\rho|_e$ of exponential transitions of an execution ρ using the number $|\rho|_m$ of multiplicative transitions of ρ , that by Theorem 2 corresponds to the number of LO β -steps on the λ -calculus. Second, we bound the number $|\rho|_c$ of commutative transitions of ρ by using the number of exponential transitions and the size of the initial term. Third, we put everything together.

Multiplicative vs Exponential Analysis This step requires two auxiliary lemmas. The first one essentially states that commutative transitions *eat* normal and neutral terms, as well as LO contexts.

Lemma 10. *Let $s = F \mid \bar{t} \mid \pi \mid E \mid \blacktriangledown$ be a state and E be well-labeled. Then*

1. If $\bar{t} \downarrow_E$ is a normal term and $\pi = \epsilon$ then $s \rightsquigarrow_c^* F \mid \bar{t} \mid \pi \mid E \mid \blacktriangle$.
2. If $\bar{t} \downarrow_E$ is a neutral term then $s \rightsquigarrow_c^* F \mid \bar{t} \mid \pi \mid E \mid \blacktriangle$.
3. If $\bar{t} = C\langle \bar{u} \rangle$ with $C \downarrow_E$ a LO context then there exist F' and π' such that $s \rightsquigarrow_c^* F' \mid \bar{u} \mid \pi' \mid E \mid \blacktriangledown$;

The second lemma uses Lemma 10 and the environment labels invariant (Lemma 6.1) to show that the exponential transitions of the Useful MAM are indeed useful, as they head towards a multiplicative transition, that is towards β -redexes.

Lemma 11 (Useful Exponentials Lead to Multiplicatives). *Let s be a reachable state such that $s \rightsquigarrow_{e_{(red,n)}} s'$.*

1. If $n = 1$ then $s' \rightsquigarrow_c^* \rightsquigarrow_m s''$;

2. If $n = 2$ then $s' \rightsquigarrow_c^* \rightsquigarrow_{e_{abs}} \rightsquigarrow_m s''$ or $s' \rightsquigarrow_c^* \rightsquigarrow_{e_{(red,1)}} s''$;
3. If $n > 1$ then $s' \rightsquigarrow_c^* \rightsquigarrow_{e_{(red,n-1)}} s''$.

Finally, using the environment size invariant (Lemma 9.2) we obtain the local boundedness property, that is used to infer a quadratic bound via a standard reasoning (already employed in [6]).

Theorem 3 (Exponentials vs Multiplicatives, Proof at Page 19). *Let s be an initial Useful MAM state and $\rho : s \rightsquigarrow^* s'$.*

1. Local Boundedness: if $\sigma : s' \rightsquigarrow^* s''$ and $|\sigma|_m = 0$ then $|\sigma|_e \leq |\rho|_m$;
2. Exponentials are Quadratic in the Multiplicatives: $|\rho|_e \in O(|\rho|_m^2)$.

Commutative vs Exponential Analysis The second step is to bound the number of commutative transitions. Since the commutative part of the Useful MAM is essentially the same as the commutative part of the Strong MAM of [2], the proof of such bound is essentially the same as in [2]. It relies on the subterm invariant (Lemma 9.1).

Theorem 4 (Commutatives vs Exponentials, Proof at Page 20). *Let $\rho : s \rightsquigarrow^* s'$ be a Useful MAM execution from an initial state of code t . Then:*

1. Commutative Evaluation Steps are Bilinear: $|\rho|_{\blacktriangledown c} \leq (1 + |\rho|_e) \cdot |t|$.
2. Commutative Evaluation Bounds Backtracking: $|\rho|_{\blacktriangle c} \leq 2 \cdot |\rho|_{\blacktriangledown c}$.
3. Commutative Transitions are Bilinear: $|\rho|_c \leq 3 \cdot (1 + |\rho|_e) \cdot |t|$.

The Main Theorem Putting together the matching between LO β -steps and multiplicative transitions (Theorem 2), the quadratic bound on the exponentials via the multiplicatives (Theorem 3.2) and the bilinear bound on the commutatives (Theorem 4.3) we obtain that the number of the Useful MAM transitions to implement a LO β -derivation d is at most quadratic in the length of d and linear in the size of t . Moreover, the subterm invariant (Lemma 9.1) and the analysis of the Checking AM (Theorem 1.2) allow to bound the cost of implementing the execution on RAM.

Theorem 5 (Useful MAM Overhead Bound, Proof at Page 20). *Let $d : t \rightarrow_{LO\beta}^* u$ be a leftmost-outermost derivation and ρ be the Useful MAM execution simulating d given by Theorem 2.2. Then:*

1. Length: $|\rho| = O((1 + |d|^2) \cdot |t|)$.
2. Cost: ρ is implementable on RAM in $O((1 + |d|^2) \cdot |t|)$ steps.

Remark 2. Our bound is quadratic in the number of the LO β -steps but we believe that it is not tight. In fact, our transition $\rightsquigarrow_{m_1}$ is a standard optimisation, used for instance in Wand's [28] (section 2), Friedman et al.'s [21] (section 4), and Sestoft's [26] (section 4), and motivated as an optimization about *space*. In Sands, Gustavsson, and Moran's [25], however, it is shown that it lowers the overhead for *time* from quadratic to linear (with respect to the number of β -steps) for call-by-name evaluation in a weak setting. Unfortunately, the simple proof used in [25] does not scale up to our setting, nor we have an alternative proof that the overhead is linear. We conjecture, however, that it does.

References

1. Accattoli, B., Barenbaum, P., Mazza, D.: Distilling Abstract Machines. In: ICFP 2014. pp. 363–376 (2014)
2. Accattoli, B., Barenbaum, P., Mazza, D.: A Strong Distillery. In: APLAS 2015. pp. 231–250 (2015)
3. Accattoli, B., Bonelli, E., Kesner, D., Lombardi, C.: A Nonstandard Standardization Theorem. In: POPL. pp. 659–670 (2014)
4. Accattoli, B., Coen, C.S.: On the Relative Usefulness of Fireballs. In: LICS 2015. pp. 141–155 (2015)
5. Accattoli, B., Dal Lago, U.: On the invariance of the unitary cost model for head reduction. In: RTA. pp. 22–37 (2012)
6. Accattoli, B., Lago, U.D.: (Leftmost-Outermost) Beta Reduction is Invariant, Indeed. *Logical Methods in Computer Science* 12(1) (2016)
7. Ariola, Z.M., Bohannon, A., Sabry, A.: Sequent calculi and abstract machines. *ACM Trans. Program. Lang. Syst.* 31(4) (2009)
8. Asperti, A., Mairson, H.G.: Parallel beta reduction is not elementary recursive. In: POPL. pp. 303–315 (1998)
9. Blelloch, G.E., Greiner, J.: Parallelism in sequential functional languages. In: FPCA. pp. 226–237 (1995)
10. Boutiller, P.: De nouveaux outils pour manipuler les inductif en Coq. Ph.D. thesis, Université Paris Diderot - Paris 7 (2014)
11. de Carvalho, D.: Execution time of lambda-terms via denotational semantics and intersection types. *CoRR* abs/0905.4251 (2009)
12. Crégut, P.: An abstract machine for lambda-terms normalization. In: LISP and Functional Programming. pp. 333–340 (1990)
13. Crégut, P.: Strongly reducing variants of the Krivine abstract machine. *Higher-Order and Symbolic Computation* 20(3), 209–230 (2007)
14. Dal Lago, U., Martini, S.: The weak lambda calculus as a reasonable machine. *Theor. Comput. Sci.* 398(1-3), 32–50 (2008)
15. Danos, V., Regnier, L.: Head linear reduction. Tech. rep. (2004)
16. Danvy, O., Nielsen, L.R.: Refocusing in reduction semantics. Tech. Rep. RS-04-26, BRICS (2004)
17. Danvy, O., Zerny, I.: A synthetic operational account of call-by-need evaluation. In: PPDP. pp. 97–108 (2013)
18. Dénès, M.: Étude formelle d’algorithmes efficaces en algèbre linéaire. Ph.D. thesis, Université de Nice - Sophia Antipolis (2013)
19. Ehrhard, T., Regnier, L.: Böhm trees, Krivine’s machine and the Taylor expansion of lambda-terms. In: CiE. pp. 186–197 (2006)
20. Fernández, M., Sifakas, N.: New developments in environment machines. *Electr. Notes Theor. Comput. Sci.* 237, 57–73 (2009)
21. Friedman, D.P., Ghuloum, A., Siek, J.G., Winebarger, O.L.: Improving the lazy krivine machine. *Higher-Order and Symbolic Computation* 20(3), 271–293 (2007)
22. García-Pérez, Á., Nogueira, P., Moreno-Navarro, J.J.: Deriving the full-reducing krivine machine from the small-step operational semantics of normal order. In: PPDP. pp. 85–96 (2013)
23. Grégoire, B., Leroy, X.: A compiled implementation of strong reduction. In: (ICFP ’02). pp. 235–246 (2002)
24. Milner, R.: Local bigraphs and confluence: Two conjectures. *Electr. Notes Theor. Comput. Sci.* 175(3), 65–73 (2007)

25. Sands, D., Gustavsson, J., Moran, A.: Lambda calculi and linear speedups. In: The Essence of Computation, Complexity, Analysis, Transformation. Essays Dedicated to Neil D. Jones. pp. 60–84 (2002)
26. Sestoft, P.: Deriving a lazy abstract machine. J. Funct. Program. 7(3), 231–264 (1997)
27. Smith, C.: Abstract machines for higher-order term sharing, Presented at IFL 2014
28. Wand, M.: On the correctness of the krivine machine. Higher-Order and Symbolic Computation 20(3), 231–235 (2007)

Proofs of the Main Lemmas and Theorems

Proof of One-Step Weak Bisimulation Lemma (Lemma 7, p. 13)

1. *Commutative*: the proof is exactly as the one for the Checking AM (Lemma 4.2), that can be found in the longer version of this paper on the author’s webpage.
2. *Exponential*:
 - **Case** $s = (F, x, \pi, E, \blacktriangledown) \rightsquigarrow_{\mathbf{e}_{red}} (F, \bar{t}^\alpha, \pi, E, \blacktriangledown) = s'$ **with** $E(x) = [x \leftarrow \bar{t}]^{(red, n)}$. Then $E = E' : [x \leftarrow \bar{t}]^{(red, n)} : E''$ for some environments E' , and E'' . Remember that terms are considered up to α -equivalence.

$$\underline{s} = C_{s'} \langle x \downarrow_E \rangle = C_{s'} \langle \bar{t} \downarrow_{E''} \rangle = C_{s'} \langle \bar{t} \downarrow_E \rangle = \underline{s}'$$

In the chain of equalities we can replace $\bar{t} \downarrow_{E''}$ with $\bar{t} \downarrow_E$ because by well-labeledness the variables bound by E' are fresh with respect to $\bar{t}$.

- **Case** $s = (F, x, \bar{u} : \pi, E, \blacktriangledown) \rightsquigarrow_{\mathbf{e}_{abs}} (F, \bar{t}^\alpha, \bar{u} : \pi, E, \blacktriangledown) = s'$ **with** $E(x) = [x \leftarrow \bar{t}]^{abs}$. The proof that $\underline{s} = \underline{s}'$ is exactly as in the previous case.

3. *Multiplicative*:

- **Case** $s = (F, \lambda x. \bar{t}, y : \pi, E, \blacktriangledown) \rightsquigarrow_{\mathbf{m}_1} (F, \bar{t}\{x \leftarrow y\}, \pi, E, \blacktriangledown) = s'$. Note that $\underline{C_s} = \underline{F} \langle \pi \rangle \downarrow_E$ is LO by the decoding invariant (Lemma 6.4). Note also that by the name invariant (Lemma 6.3b) x can only occur in $\bar{t}$. Then:

$$\begin{aligned} \underline{(F, \lambda x. \bar{t}, y : \pi, E, \blacktriangledown)} &= \underline{F} \langle \langle \lambda x. \bar{t} \rangle y : \pi \rangle \downarrow_E \\ &= \underline{F} \langle \langle (\lambda x. \bar{t}) y \rangle \pi \rangle \downarrow_E \\ &= C_{s'} \langle (\lambda x. \bar{t} \downarrow_E) y \downarrow_E \rangle \\ &\rightarrow_{\text{LO}\beta} C_{s'} \langle \bar{t} \downarrow_E \{x \leftarrow y \downarrow_E\} \rangle \\ &=_{L.6.3b \& L.3.2} C_{s'} \langle \bar{t} \{x \leftarrow y\} \downarrow_E \rangle \\ &= \underline{(F, \bar{t}\{x \leftarrow y\}, \pi, E, \blacktriangledown)} \end{aligned}$$

- **Case** $s = (F, \lambda x. \bar{t}, \bar{u} : \pi, E, \blacktriangledown) \rightsquigarrow_{\mathbf{m}_2} (F, \bar{t}, \pi, [x \leftarrow \bar{u}]^l : E, \blacktriangledown) = s'$ **with** $\bar{u}$ **not a variable**. Note that $\underline{C_{s'}} = \underline{F} \langle \langle \cdot \rangle \pi \rangle \downarrow_E = \underline{F} \downarrow_E \langle \langle \cdot \rangle \pi \downarrow_E \rangle$ is LO by the decoding invariant (Lemma 6.4). Note also that by the name invariant

(Lemma 6.3b) x can only occur in $\bar{t}$. Then:

$$\begin{aligned}
\underline{(F, \lambda x.\bar{t}.\bar{u} : \pi, E, \blacktriangledown)} &= \underline{F \langle \langle \lambda x.\bar{t} \rangle \bar{u} : \pi \rangle \downarrow_E} \\
&= \underline{F \langle \langle (\lambda x.\bar{t})\bar{u} \rangle \pi \rangle \downarrow_E} \\
&= \underline{F \downarrow_E \langle \langle (\lambda x.\bar{t} \downarrow_E) \bar{u} \downarrow_E \rangle \pi \downarrow_E \rangle} \\
&\rightarrow_{L0\beta} \underline{F \downarrow_E \langle \langle \bar{t} \downarrow_E \{x \leftarrow \bar{u} \downarrow_E\} \rangle \pi \downarrow_E \rangle} \\
&=_{L.6.3b \& L.3.2} \underline{F \downarrow_E \langle \langle \bar{t} \{x \leftarrow \bar{u}\} \downarrow_E \rangle \pi \downarrow_E \rangle} \\
&= \underline{F \langle \langle \bar{t} \{x \leftarrow \bar{u}\} \rangle \pi \rangle \downarrow_E} \\
&=_{L.6.3b \& L.3.1} \underline{F \langle \langle \bar{t} \{x \leftarrow \bar{u}\} \rangle \pi \rangle \downarrow_E} \\
&=_{L.6.3b} \underline{F \langle \langle \bar{t} \rangle \pi \rangle \{x \leftarrow \bar{u}\} \downarrow_E} \\
&= \underline{F \langle \langle \bar{t} \rangle \pi \rangle \downarrow_{[x \leftarrow \bar{u}]^l : E}} \\
&= \underline{(F, \bar{t}, \pi, [x \leftarrow \bar{u}]^l : E, \blacktriangledown)} \square
\end{aligned}$$

Proof of the Progress Lemma (Lemma 8, p. 13)

A simple inspection of the machine transitions shows that final states have the form $(\epsilon, \bar{t}, \epsilon, E, \blacktriangle)$. Then by the normal form invariant (Lemma 6.2a) $\underline{s} = \bar{t} \downarrow_E$ is β -normal. $\square$

Proof of the Weak Bisimulation Theorem (Thm 2, p. 13)

- By induction on the length $|\rho|$ of ρ , using the one-step weak simulation lemma (Lemma 7). If ρ is empty then the empty derivation satisfies the statement. If ρ is given by $\sigma : s \rightsquigarrow^* s''$ followed by $s'' \rightsquigarrow s'$ then by *i.h.* there exists $e : \underline{s} \rightarrow_{L0\beta}^* \underline{s''}$ s.t. $|e| = |\sigma|_{\mathbf{m}}$. Cases of $s'' \rightsquigarrow s'$:
 - Commutative or Exponential.* Then $\underline{s''} = \underline{s'}$ by Lemma 7.1 and Lemma 7.2, and the statement holds taking $d := e$ because $|d| = |e| =_{i.h.} |\sigma|_{\mathbf{m}} = |\rho|_{\mathbf{m}}$.
 - Multiplicative.* Then $\underline{s''} \rightarrow_{L0\beta} \underline{s'}$ by Lemma 7.3 and defining d as e followed by such a step we obtain $|d| = |e| + 1 =_{i.h.} |\sigma|_{\mathbf{m}} + 1 = |\rho|_{\mathbf{m}}$.
- We use $\mathbf{nf}_{\text{ec}}(s)$ to denote the normal form of s with respect to exponential and commutative transitions, that exists and is unique because $\rightsquigarrow_c \cup \rightsquigarrow_e$ terminates (termination is given by forthcoming Theorem 3 and Theorem 4, that are postponed because they actually give precise complexity bounds, not just termination) and the machine is deterministic (as it can be seen by an easy inspection of the transitions). The proof is by induction on the length of d . If d is empty then the empty execution satisfies the statement. If d is given by $e : \bar{t} \rightarrow_{L0\beta}^* w$ followed by $w \rightarrow_{L0\beta} u$ then by *i.h.* there is an execution $\sigma : s \rightsquigarrow^* s''$ s.t. $w = \underline{s''}$ and $|\sigma|_{\mathbf{m}} = |e|$. Note that since exponential and commutative transitions are mapped on equalities, σ can be extended as $\sigma' : s \rightsquigarrow^* s'' \rightsquigarrow_{\mathbf{e}_{red}, \mathbf{e}_{abs}, \mathbf{c}_{1,2,3,4,5,6}}^* \mathbf{nf}_{\text{ec}}(s'')$ with $\mathbf{nf}_{\text{ec}}(s'') = w$ and $|\sigma'|_{\mathbf{m}} = |e|$. By the progress property (Lemma 8) $\mathbf{nf}_{\text{ec}}(s'')$ cannot be a final state, otherwise $w = \mathbf{nf}_{\text{ec}}(s'')$ could not reduce. Then $\mathbf{nf}_{\text{ec}}(s'') \rightsquigarrow_{\mathbf{m}} s'$ (the transition is necessarily multiplicative because $\mathbf{nf}_{\text{ec}}(s'')$ is normal with respect to the other transitions). By the one-step weak simulation lemma (Lemma 7.3) $\mathbf{nf}_{\text{ec}}(s'') = w \rightarrow_{L0\beta} \underline{s'}$ and by determinism of $\rightarrow_{L0\beta}$ (Lemma 1)

$s' = u$. Then the execution ρ defined as σ' followed by $\mathbf{nf}_{ec}(s'') \rightsquigarrow_m s'$ satisfy the statement, as $|\rho|_m = |\sigma'|_m + 1 = |\sigma|_m + 1 = |e| + 1 = |d|$. $\square$

Proof of the Exponentials vs Multiplicatives Theorem (Thm 3, p. 15)

1. We prove that $|\sigma|_e \leq |E|$. The statement follows from the environment size invariant (Lemma 9.2), for which $|E| \leq |\rho|_m$.
If $|\sigma|_e = 0$ it is immediate. Then assume $|\sigma|_e > 0$, so that there is a first exponential transition in σ , *i.e.* σ has a prefix $s' \rightsquigarrow_c^* \rightsquigarrow_e s''$ followed by an execution $\tau : s'' \rightsquigarrow^* s'''$ such that $|\tau|_m = 0$. Cases of the first exponential transition $\rightsquigarrow_e$:
 - Case $\rightsquigarrow_{e_{abs}}$: the next transition is necessarily multiplicative, and so τ is empty. Then $|\sigma|_e = 1$. Since the environment is non-empty (otherwise $\rightsquigarrow_{e_{abs}}$ could not apply), $|\sigma|_e \leq |E|$ holds.
 - Case $\rightsquigarrow_{e_{(red,n)}}$: we prove by induction on n that $|\sigma|_e \leq n$, that gives what we want because $n \leq |E|$ by Remark 1. Cases:
 - $n = 1$) Then τ has the form $s'' \rightsquigarrow_c^* s'''$ by Lemma 11.1, and so $|\sigma|_e = 1$.
 - $n = 2$) Then τ is a prefix of $\rightsquigarrow_c^* \rightsquigarrow_{e_{abs}}$ or $\rightsquigarrow_c^* \rightsquigarrow_{e_{(red,1)}}$ by Lemma 11.2. In both cases $|\sigma|_e \leq 2$.
 - $n > 2$) Then by Lemma 11.3 τ is either shorter or equal to $\rightsquigarrow_c^* \rightsquigarrow_{e_{(red,n-1)}}$, and so $|\sigma|_e \leq 2$, or it is longer than $\rightsquigarrow_c^* \rightsquigarrow_{e_{(red,n-1)}}$, *i.e.* it writes as $\rightsquigarrow_c^*$ followed by an execution τ' starting with $\rightsquigarrow_{e_{(red,n-1)}}$. By *i.h.* $|\tau'| \leq n - 1$ and so $|\sigma|_e \leq n$.
2. This is a standard reasoning: since by local boundedness (the previous point) m -free sequences have a number of e -transitions that are bound by the number of preceding m -transitions, the sum of all e -transitions is bound by the square of m -transitions. It is analogous to the proof of Theorem 7.2.3 in [6]. $\square$

Proof of Commutatives vs Exponentials Theorem (Thm 4, p. 15)

1. We prove a slightly stronger statement, namely $|\rho|_{\blacktriangledown c} + |\rho|_m \leq (1 + |\rho|_e) \cdot |t|$, by means of the following notion of size for stacks/frames/states:

$$\begin{aligned}
 |\epsilon| &:= 0 & |x : F| &:= |F| \\
 |\bar{t} : \pi| &:= |\bar{t}| + |\pi| & |\bar{t} \diamond \pi : F| &:= |\pi| + |F| \\
 |(F, \bar{t}, \pi, E, \blacktriangledown)| &:= |F| + |\pi| + |\bar{t}| & |(F, \bar{t}, \pi, E, \blacktriangle)| &:= |F| + |\pi|
 \end{aligned}$$

By direct inspection of the rules of the machine it can be checked that:

- *Exponentials Increase the Size*: if $s \rightsquigarrow_e s'$ is an exponential transition, then $|s'| \leq |s| + |t|$ where $|t|$ is the size of the initial term; this is a consequence of the fact that exponential steps retrieve a piece of code from the environment, which is a subterm of the initial term by Lemma 9.1;
- *Non-Exponential Evaluation Transitions Decrease the Size*: if $s \rightsquigarrow_a s'$ with $a \in \{m_1, m_2, \blacktriangledown c_1, \blacktriangledown c_2, \blacktriangledown c_3\}$ then $|s'| < |s|$ (for $\blacktriangledown c_3$ because the transition switches to backtracking, and thus the size of the code is no longer taken into account);

- *Backtracking Transitions do not Change the Size*: if $s \rightsquigarrow_a s'$ with $a \in \{\blacktriangle c_4, \blacktriangle c_5, \blacktriangle c_6\}$ then $|s'| = |s|$.

Then a straightforward induction on $|\rho|$ shows that

$$|s'| \leq |s| + |\rho|_e \cdot |t| - |\rho|_{\blacktriangledown c} - |\rho|_m$$

i.e. that $|\rho|_{\blacktriangledown c} + |\rho|_m \leq |s| + |\rho|_e \cdot |t| - |s'|$.

Now note that $|\cdot|$ is always non-negative and that since s is initial we have $|s| = |t|$. We can then conclude with

$$\begin{aligned} |\rho|_{\blacktriangledown c} + |\rho|_m &\leq |s| + |\rho|_e \cdot |t| - |s'| \\ &\leq |s| + |\rho|_e \cdot |t| = |t| + |\rho|_e \cdot |t| = (1 + |\rho|_e) \cdot |t| \end{aligned}$$

- We have to estimate $|\rho|_{\blacktriangle c} = |\rho|_{\blacktriangle c_4} + |\rho|_{\blacktriangle c_5} + |\rho|_{\blacktriangle c_6}$. Note that
 - $|\rho|_{\blacktriangle c_4} \leq |\rho|_{\blacktriangledown c_2}$, as $\rightsquigarrow_{\blacktriangle c_4}$ pops variables from F , pushed only by $\rightsquigarrow_{\blacktriangledown c_2}$;
 - $|\rho|_{\blacktriangle c_5} \leq |\rho|_{\blacktriangle c_6}$, as $\rightsquigarrow_{\blacktriangle c_5}$ pops pairs $\bar{t} \diamond \pi$ from F , pushed only by $\rightsquigarrow_{\blacktriangle c_6}$;
 - $|\rho|_{\blacktriangle c_6} \leq |\rho|_{\blacktriangledown c_3}$, as $\rightsquigarrow_{\blacktriangle c_6}$ ends backtracking phases, started only by $\rightsquigarrow_{\blacktriangledown c_3}$.
 Then $|\rho|_{\blacktriangle c} \leq |\rho|_{\blacktriangledown c_2} + 2|\rho|_{\blacktriangledown c_3} \leq 2|\rho|_{\blacktriangledown c}$.
- We have $|\rho|_c = |\rho|_{\blacktriangledown c} + |\rho|_{\blacktriangle c} \leq_{P.2} |\rho|_{\blacktriangledown c} + 2|\rho|_{\blacktriangledown c} \leq_{P.1} 3 \cdot (1 + |\rho|_e) \cdot |t|$. $\square$

Proof of the Useful MAM Overhead Bound Theorem (Thm 5, p. 15)

- By definition, the length of the execution ρ simulating d is given by $|\rho| = |\rho|_m + |\rho|_e + |\rho|_c$. Now, by Theorem 3.2 we have $|\rho|_e = O(|\rho|_m^2)$ and by Theorem 4.3 we have $|\rho|_c = O((1 + |\rho|_e) \cdot |t|) = O((1 + |\rho|_m^2) \cdot |t|)$. Therefore, $|\rho| = O((1 + |\rho|_e) \cdot |t|) = O((1 + |\rho|_m^2) \cdot |t|)$. By Theorem 2.2 $|\rho|_m = |d|$, and so $|\rho| = O((1 + |d|^2) \cdot |t|)$.
- The cost of implementing ρ is the sum of the costs of implementing the multiplicative, exponential, and commutative transitions. Remember that the idea is that variables are implemented as references, so that environment can be accessed in constant time (*i.e.* they do not need to be accessed sequentially):
 - Commutative*: every commutative transition evidently takes constant time. At the previous point we bounded their number with $O((1 + |d|^2) \cdot |t|)$, which is then also the cost of all the commutative transitions together.
 - Multiplicative*: a $\rightsquigarrow_{m_1}$ transition costs $O(|t|)$ because it requires to rename the current code, whose size is bound by the size of the initial term by the subterm invariant (Lemma 9.1a). A $\rightsquigarrow_{m_2}$ transition also costs $O(|t|)$ because executing the Checking AM on $\bar{u}$ takes $O(|\bar{u}|)$ commutative steps (Theorem 1.2), commutative steps take constant time, and the size of $\bar{u}$ is bound by $|t|$ by the subterm invariant (Lemma 9.1b). Therefore, all together the multiplicative transitions cost $O(|d| \cdot |t|)$.
 - Exponential*: At the previous point we bounded their number with $|\rho|_e = O(|d|^2)$. Each exponential step copies a term from the environment, that by the subterm invariant (Lemma 9.1d) costs at most $O(|t|)$, and so their full cost is $O((1 + |d|) \cdot |t|^2)$ (note that this is exactly the cost of the commutative transitions, but it is obtained in a different way).
 Then implementing ρ on RAM takes $O((1 + |d|) \cdot |t|^2)$ steps. $\square$