

A Systematic Study of Functional Language Implementations

Rémi Douence and Pascal Fradet

INRIA / IRISA

Campus de Beaulieu, 35042 Rennes Cedex, France

douence@irisa.fr fradet@irisa.fr

Abstract: We introduce a unified framework to describe, relate, compare and classify functional language implementations. The compilation process is expressed as a succession of program transformations in the common framework. At each step, different transformations model fundamental choices. A benefit of this approach is to structure and decompose the implementation process. The correctness proofs can be tackled independently for each step and amount to proving program transformations in the functional world. This approach also paves the way to formal comparisons by making it possible to estimate the complexity of individual transformations or compositions of them. Our study aims at covering the whole known design space of sequential functional languages implementations. In particular, we consider call-by-value, call-by-name and call-by-need reduction strategies as well as environment and graph-based implementations. We describe for each compilation step the diverse alternatives as program transformations. In some cases, we illustrate how to compare or relate compilation techniques, express global optimizations or hybrid implementations. We also provide a classification of well-known abstract machines.

Key-words: functional programming, compilers, optimization, program transformation, combinators

1 INTRODUCTION

One of the most studied issues concerning functional languages is their implementation. Since Landin's seminal proposal, 30 years ago [31], a plethora of new abstract machines or compilation techniques have been proposed. The list of existing abstract machines includes the SECD [31], the Cam [10], the CMCM [36], the Tim [20], the Zam [32], the G-machine [27] and the Krivine-machine [11]. Other implementations are not described via an abstract machine but as a collection of transformations or compilation techniques such as compilers based on continuation passing style (CPS) [2][22][30][52]. Furthermore, numerous papers present optimizations often adapted to a specific abstract machine or a specific approach [3][8][28]. Looking at this myriad of distinct works, obvious questions spring to mind: what are the fundamental choices? What are the respective benefits of these alternatives? What are precisely the common points and differences between two compilers? Can a particular optimization, designed for machine *A*, be adapted to machine *B*? One finds comparatively very few papers devoted to these questions. There have been studies of the relationship between two individual machines [37][43] but, to the best of our knowledge, no global approach to study implementations.

The goal of this paper is to fill this gap by introducing a unified framework to describe, relate, compare and classify functional language implementations. Our approach is to express the whole compilation process as a succession of program transformations. The common framework considered here is a hierarchy of intermediate languages all of which are subsets of the lambda-calculus. Our description of an implementation consists of a series of transformations $\Lambda \xrightarrow{T_1} \Lambda_1 \xrightarrow{T_2} \dots \xrightarrow{T_n} \Lambda_n$, each one compiling a particular task by mapping an expression from one intermediate language into another. The last language Λ_n consists of functional expressions that can be seen as assembly code (essentially, combinators with explicit sequencing and calls). For each step, different transformations are designed to represent fundamental choices or optimizations. A benefit of this approach is to structure and decompose the implementation process. Two seemingly disparate implementations can be found to share some compilation steps. This approach also has interesting payoffs as far as correctness proofs and comparisons are concerned. The correctness of each step can be tackled independently and amounts to proving a program transformation in the functional world. Our approach also paves the way to formal comparisons by estimating the complexity of individual transformations or compositions of them.

We concentrate on pure λ -expressions and our source language Λ is $E ::= x \mid \lambda x.E \mid E_1 E_2$. Most fundamental choices can be described using this simple language. The two steps which cause the greatest impact on the compiler are the implementation of the reduction strategy (searching for the next redex) and the environment management (compilation of the β -reduction). Other steps include the implementation of control transfers (calls & returns), the implementation of closure sharing and update (implied by the call-by-need strategy), the representation of components like the data stack or environments and various optimizations.

In Section 2 we describe the framework used to model the compilation process. In Section 3, we present the alternatives to compile the reduction strategy (i.e. call-by-value and call-by-name). The compilation of control used by graph reducers is peculiar. A separate section (3.3) is dedicated to this point. Section 3 ends with a comparison of two compilation techniques of call-by-value and a study of the relationship between the compilation of control in the environment and graph-based models. Section 4 (resp. Section 5) describes the different options to compile the β -reduction (resp. the control transfers). Call-by-need is nothing but call-by-name with redex sharing and update and we present in Section 6 how it can be expressed in our framework. Section 7 embodies our study in a taxonomy of classical functional implementations. In Section 8, we outline some extensions and applications of the framework. Section 9 is devoted to a review of related work and Section 10 concludes by indicating directions for future research.

In order to alleviate the presentation, some more involved material such as proofs, variants of transformations and other technical details have been kept out of the main text. We refer the motivated reader to the (electronically published) appendix. References to the appendix are noted “[Ⓧ]”. A previous conference paper [16] concentrates on call-by-value and can be used as a short introduction to this work. Additional details can also be found in two companion technical reports ([17], [18]) and a PhD thesis [19].

2 GENERAL FRAMEWORK

Each compilation step is represented by a transformation from an intermediate language to another one that is closer to machine code. In this paper, the whole implementation process is described via a transformation sequence $\Lambda \xrightarrow{T_1} \Lambda_s \xrightarrow{T_2} \Lambda_e \xrightarrow{T_3} \Lambda_k \xrightarrow{T_4} \Lambda_h$ starting with Λ and involving four intermediate languages (very close to each other). This framework possesses several benefits:

- It has a *strong formal basis*. Each intermediate language can be seen either as a formal system with its own conversion rules or as a subset of the λ -calculus by defining its constructs as λ -expressions. The intermediate languages share many laws and properties; the most important being that every reduction strategy is normalizing. These features facilitate program transformations, correctness proofs and comparisons.
- It is (relatively) *abstract*. Since we want to model completely and precisely implementations, the intermediate languages must come closer to an assembly language as we progress in the description. The framework nevertheless possesses many abstract features which do not lessen its precision. The combinators of the intermediate languages and their conversion rules allow a more abstract description of notions such as instructions, sequencing, stacks, ... than an encoding as λ -expressions. As a consequence, the compilation of control is expressed more abstractly than using CPS expressions and the implementation of components (e.g. data stack, environment stack, ...) is a separate step.
- It is *modular*. Each transformation implements one compilation step and can be defined independently from the former steps. Transformations implementing different steps are freely composed to specify implementations. Transformations implementing the same step represent different choices and can be compared.
- It is *extendable*. New intermediate languages and transformations can be defined and inserted into the transformation sequence to model new compilation steps (e.g. register allocation).

2.1 Overview

The first step is the compilation of control which is described by transformations from Λ to Λ_s . The intermediate language Λ_s (Figure 1) is defined using the combinators \circ , **push_s** and a new form of λ -abstraction $\lambda_{s,x}.E$. Intuitively, \circ is a sequencing operator and $E_1 \circ E_2$ can be read “evaluate E_1 then evaluate E_2 ”, **push_s** E returns E as a result and $\lambda_{s,x}.E$ binds the previous intermediate result to x before evaluating E . The pair (**push_s**, λ_s) specifies a component (noted s) storing intermediate results (e.g. a data stack). So, **push_s** and λ_s can be seen as “store” and “fetch” in s .

The most notable syntactic feature of Λ_s is that it rules out unrestricted applications. Its main property is that the choice of the next weak redex is not relevant anymore: all weak redexes are needed. This is the key point to view transformations from Λ to Λ_s as compiling the evaluation strategy.

Transformations from Λ_s to Λ_e are used to compile the β -reduction. The language Λ_e excludes unrestricted uses of variables which are now only needed to define macro-combina-

tors. The encoding of environment management is made possible using the new pair $(\mathbf{push}_e, \lambda_e)$. They behave exactly as \mathbf{push}_s and λ_s ; they just act on a (at least conceptually) different component e (e.g. a stack of environments).

Λ_s	$E ::= x \mid E_1 \circ E_2 \mid \mathbf{push}_s E \mid \lambda_s x.E$
Λ_e	$E ::= x \mid E_1 \circ E_2 \mid \mathbf{push}_s E \mid \lambda_s x.E \mid \mathbf{push}_e E \mid \lambda_e x.E$
Λ_k	$E ::= x \mid E_1 \circ E_2 \mid \mathbf{push}_s E \mid \lambda_s x.E \mid \mathbf{push}_e E \mid \lambda_e x.E \mid \mathbf{push}_k E \mid \lambda_k x.E$
Λ_h	$E ::= x \mid E_1 \circ E_2 \mid \mathbf{push}_s E \mid \lambda_s x.E \mid \mathbf{push}_e E \mid \lambda_e x.E \mid \mathbf{push}_k E \mid \lambda_k x.E \mid \mathbf{push}_h E \mid \lambda_h x.E$

Figure 1 The intermediate languages

Transformations from Λ_e to Λ_k describe the compilation of control transfers. The language Λ_k makes calls and returns explicit. It introduces the pair $(\mathbf{push}_k, \lambda_k)$ which specifies a component k storing return addresses.

The last transformations from Λ_k to Λ_h adds a memory component in order to express closure sharing and updating. The language Λ_h introduces the pair $(\mathbf{push}_h, \lambda_h)$ which specifies a global heap h . The expressions of this last language can be read as assembly code.

2.2 Conversion Rules

The substitution and the notion of free or bound variables are the same as in λ -calculus. The basic combinators can be given different definitions (possible definitions are given in 2.5). We do not pick specific ones up at this point; we simply impose the associativity of sequencing and that the combinators satisfy the equivalent of β and η -conversions (Figure 2).

(assoc)	$(E_1 \circ E_2) \circ E_3 = E_1 \circ (E_2 \circ E_3)$
(β_i)	$(\mathbf{push}_i F) \circ (\lambda_i x.E) = E[F/x]$
(η_i)	$\lambda_i x.(\mathbf{push}_i x \circ E) = E$ <i>if x does not occur free in E</i>

Figure 2 Conversion rules in Λ_i (for $i \in \{s, e, k, h\}$)

We consider only reduction rules corresponding to the classical β -reduction:

$$(\mathbf{push}_i F) \circ (\lambda_i x.E) \rightarrow E[F/x]$$

As with all standard implementations, we are only interested in modeling weak reductions. In our framework, a weak redex is a redex that does not occur inside an expression of the form $\mathbf{push}_i E$ or $\lambda_i x.E$. Weak reduction does not reduce under \mathbf{push}_i 's or λ_i 's and, from

here on, we write “redex” (resp. reduction, normal form) for weak redex (resp. weak reduction, weak normal form).

The following example illustrates β_i -reduction (note that $\mathbf{push}_s F \circ \lambda_{s,z}.G$ is not a (weak) redex of the global expression).

$$\begin{aligned} & \mathbf{push}_e E \circ \mathbf{push}_s (\mathbf{push}_s F \circ \lambda_{s,z}.G) \circ \lambda_{s,x}.\lambda_{e,y}.\mathbf{push}_s(\mathbf{push}_e y \circ x) \\ & \quad \blacktriangleright \mathbf{push}_e E \circ \lambda_{e,y}.\mathbf{push}_s(\mathbf{push}_e y \circ \mathbf{push}_s F \circ \lambda_{s,z}.G) \\ & \quad \blacktriangleright \mathbf{push}_s (\mathbf{push}_e E \circ \mathbf{push}_s F \circ \lambda_{s,z}.G) \end{aligned}$$

Any two redexes are clearly disjoint and the β_i -reductions are left-linear so the term rewriting system is orthogonal hence confluent [29]. Alternatively, it is very easy to show that the relation \blacktriangleright is strongly confluent therefore confluent ^③. Furthermore, any redex is needed (a rewrite cannot suppress a redex) thus

Property 1 *All Λ_i reduction strategies are normalizing.*

This property is the key point to view transformations from Λ to Λ_s as compiling the reduction order.

2.3 A Typed Subset

All the expressions of the intermediate languages can be given a meaning as λ -expressions (Section 2.5). Using conversion rules such as (assoc) the same expression can be represented differently. For example, one can write equivalently

$$\mathbf{push}_s E_1 \circ (\mathbf{push}_s E_2 \circ \lambda_{s,x}.\lambda_{s,y}.E_3) \quad \text{or} \quad (\mathbf{push}_s E_1 \circ \mathbf{push}_s E_2) \circ \lambda_{s,x}.\lambda_{s,y}.E_3$$

This flexibility is very useful to transform or reshape the code. However, unrestricted transformations may lose information about the structure of the expression. Many laws and transformations (see e.g. laws (L2) and (L3) in Section 2.4 or transformation $\mathcal{H}c$ in Section 6.1) rely on the fact that a subexpression denotes a result (i.e. can be reduced to an expression of the form $\mathbf{push}_i E$) or a function (i.e. can be reduced to an expression of the form $\lambda_{i,x}.E$). If we allow subexpressions such as $(\mathbf{push}_s E_1 \circ \mathbf{push}_s E_2)$ which neither denote a result nor a function, less laws and transformations can be expressed. It is therefore convenient to restrict Λ_i using a type system (Figure 3).

$$\frac{}{\Gamma \cup \{x:\sigma\} \vdash x:\sigma} \quad \frac{\Gamma \vdash E : \sigma}{\Gamma \vdash \mathbf{push}_i E : R_i\sigma} \quad \frac{\Gamma \cup \{x:\sigma\} \vdash E : \tau}{\Gamma \vdash \lambda_{i,x}.E : \sigma \rightarrow_i \tau} \quad \frac{\Gamma \vdash E_1 : R_i\sigma \quad \Gamma \vdash E_2 : \sigma \rightarrow_i \tau}{\Gamma \vdash E_1 \circ E_2 : \tau}$$

Figure 3 Λ_i typed subset ($\Lambda_i^{\mathcal{C}}$) (for $i \in \{s,e,k,h\}$)

The restrictions enforced by the type system are on how results and functions are combined in Λ_i . For example, the composition $E_1 \circ E_2$ is restricted so that E_1 denotes a result (i.e.

has type $R_i\sigma$, R_i being a type constructor) and E_2 denotes a function. The type system restricts the set of normal forms (which in general includes expressions such as $\mathbf{push}_i E_1 \circ \mathbf{push}_j E_2$) and we have the following natural facts [Ⓟ]

Property 2 - If a closed expression $E:R_i\sigma$ has a normal form then $E \xrightarrow{*} \mathbf{push}_i V$

- If a closed expression $E:\sigma \rightarrow_i \tau$ has a normal form then $E \xrightarrow{*} \lambda_{\tau}x.F$

So, the reduction of any well-typed expression $A \circ F$ either reaches an expression of the form $\mathbf{push}_i A' \circ \lambda_{\tau}x.F'$ or loops.

Our transformations implementing compilation steps will produce well-typed expressions denoting results and, during all the compilation process, the compiled program will be well-typed. Typing is used to maintain some structure in the expression and does not impose any restrictions on source λ -expressions [Ⓟ]. It should be regarded as a syntactic tool not a semantic one. Ill-typed Λ_i -expressions have a meaning in terms of λ -expressions as well (see Section 2.5).

2.4 Laws

This framework possesses a number of algebraic laws that are useful to transform the functional code or to prove the correctness or equivalence of program transformations such as

$$\text{If } x \text{ does not occur free in } F \quad (\lambda_{\tau}x.E) \circ F = \lambda_{\tau}x.(E \circ F) \quad (\text{L1})$$

$$\text{For all } E_1:R_i\sigma, \text{ if } x \text{ does not occur free in } E_2 \quad E_1 \circ (\lambda_{\tau}x.(E_2 \circ E_3)) = E_2 \circ (E_1 \circ (\lambda_{\tau}x.E_3)) \quad (\text{L2})$$

$$\text{For all } E_1:R_i\sigma, E_2:R_j\tau \text{ and } x \neq y \quad E_1 \circ (E_2 \circ (\lambda_{\tau}x.\lambda_{\tau}y.E_3)) = E_2 \circ (E_1 \circ (\lambda_{\tau}y.\lambda_{\tau}x.E_3)) \quad (\text{L3})$$

These rules [Ⓟ] permit code to be moved inside or outside function bodies or to invert the evaluation order of two intermediate results (which is correct because we consider only purely functional expressions). To illustrate the conversion rules at work, let us prove the law (L1). Note that x does not occur free in $(\lambda_{\tau}x.E)$ nor, by hypothesis, in F and

$$\begin{aligned} (\lambda_{\tau}x.E) \circ F &= \lambda_{\tau}x.\mathbf{push}_i x \circ ((\lambda_{\tau}x.E) \circ F) && (\eta_i) \\ &= \lambda_{\tau}x.((\mathbf{push}_i x \circ (\lambda_{\tau}x.E)) \circ F) && (\text{assoc}) \\ &= \lambda_{\tau}x.(E[x/x] \circ F) && (\beta_i) \\ &= \lambda_{\tau}x.(E \circ F) && (\text{subst}) \end{aligned}$$

Even if using some rules or laws (e.g. (assoc) or (L1)) may lead to untyped programs, we still can use them as long as the final program is well-typed. For example, a closed and well-typed expression

$$(\mathbf{push}_s V \circ (\lambda_{s,x}.\mathbf{push}_s E)) \circ (\lambda_{s,y}.F)$$

can be transformed using (assoc) and (L1) into the well-typed expression

$$\mathbf{push}_s V \circ \lambda_s x. (\mathbf{push}_s E \circ (\lambda_s y. F))$$

To simplify the presentation, we often omit parentheses and write for example $\mathbf{push}_i E \circ \lambda_i x. F \circ G$ for $(\mathbf{push}_i E) \circ (\lambda_i x. (F \circ G))$. We also use syntactic sugar such as tuples (x_1, \dots, x_n) and simple pattern-matching $\lambda_i(x_1, \dots, x_n). E$.

2.5 Instantiation

The intermediate languages Λ_i are subsets of the λ -calculus made of combinators. An important point is that we do not have to give a precise definition to combinators. We just assume that they respect properties (β_i) , (η_i) and (assoc). Definitions can be chosen only after the last compilation step. This feature allows us to shift from the β_i -reduction in Λ_i to a state-machine-like expression reduction. Moreover, it permits to specify the implementation of components independently from the other steps. For example, we may eventually choose to implement the data component s and the environment component e either as a single stack or as two separate ones. We present in Section 7 an example of instantiation for the Cam.

In order to provide some intuition, we nevertheless give here some possible definitions in terms of standard λ -expressions. The most natural definition for the sequencing combinator is $\circ = \lambda abc. a (b c)$, that is $E_1 \circ E_2 = \lambda c. E_1 (E_2 c)$. The (fresh) variable c can be seen as a continuation and implements the sequencing.

The pairs of combinators $(\lambda_i, \mathbf{push}_i)$ can be seen as encoding a component of an underlying abstract machine and their definitions as specifying the state transitions. A sequence of code such as $\mathbf{push}_i E_1 \circ \dots \circ \mathbf{push}_i E_n \circ \dots$ suggests that the underlying machine must possess a component i (such as a stack, a list, a tree or a vector) in order to store intermediate results. We can choose to keep the components separate or merge (some of) them.

Keeping all the components separate leads to the following possible definitions (c, s, e, k, h being fresh variables):

$$\begin{array}{ll} \mathbf{push}_s N = \lambda c. \lambda s. \lambda e. \lambda k. \lambda h. c (s, N) e k h & \lambda_s x. X = \lambda c. \lambda (s, x). \lambda e. \lambda k. \lambda h. X c s e k h \\ \mathbf{push}_e N = \lambda c. \lambda s. \lambda e. \lambda k. \lambda h. c s (e, N) k h & \lambda_e x. X = \lambda c. \lambda s. \lambda (e, x). \lambda k. \lambda h. X c s e k h \\ \mathbf{push}_k N = \lambda c. \lambda s. \lambda e. \lambda k. \lambda h. c s e (k, N) h & \lambda_k x. X = \lambda c. \lambda s. \lambda e. \lambda (k, x). \lambda h. X c s e k h \\ \mathbf{push}_h N = \lambda c. \lambda s. \lambda e. \lambda k. \lambda h. c s e k (h, N) & \lambda_h x. X = \lambda c. \lambda s. \lambda e. \lambda k. \lambda (h, x). X c s e k h \end{array}$$

Then, the reduction (using classical β -reduction and normal order) of our expressions can be seen as state transitions of an abstract machine with five components (code, data stack, environment stack, control stack, heap), e.g.:

$$\mathbf{push}_s N C S E K H \rightarrow C (S, N) E K H$$

$$\mathbf{push}_h N C S E K H \rightarrow C S E K (H, N)$$

According to the definition of \circ the rewriting rule for sequencing is

$$(E_1 \circ E_2) C S E K H \rightarrow E_1 (E_2 C) S E K H$$

Note that C plays the role of a continuation. A code can be seen as a state transformer of type

$$(data \rightarrow env \rightarrow control \rightarrow heap \rightarrow Ans) \rightarrow data \rightarrow env \rightarrow control \rightarrow heap \rightarrow Ans$$

To be reduced, a code is applied to an initial continuation (e.g. **id**), initial (empty) data, environment and control components and an initial heap.

Keeping some components separate brings new properties such as

$$\mathbf{push}_i E \circ \mathbf{push}_j F = \mathbf{push}_j F \circ \mathbf{push}_i E \quad \text{if } i \neq j$$

allowing code motion and simplifications.

A second option is to merge all the components. The underlying abstract machine has only two components (the code and a data-environment-control-heap stack). Possible definitions are:

$$\mathbf{push}_s N = \mathbf{push}_e N = \mathbf{push}_k N = \mathbf{push}_h N = \lambda c. \lambda z. c (z, N)$$

$$\lambda_s x. X = \lambda_e x. X = \lambda_k x. X = \lambda_h x. X = \lambda c. \lambda (z, x). X c z$$

and the reduction of expressions is of the form $\mathbf{push}_i N C Z \rightarrow C (Z, N)$ for $i \in \{s, e, k, h\}$

Let us point out that our use of the term “abstract machines” should not suggest a layer of interpretation. The abstraction only consists of the use of components and generic code. At the end of the compilation process, we get realistic assembly code and the “abstract machines” resemble real machines.

3 COMPILATION OF CONTROL

We focus here on the compilation of the call-by-value and the call-by-name reduction strategies. Call-by-need is only a refinement of call-by-name involving redex sharing and update. It is described in Section 6. We first present the two main choices taken by environment-based implementations. Following Peyton Jones’ terminology [42], these two options are named the *eval-apply* model (presented in Section 3.1) and the *push-enter* model (presented in Section 3.2). The graph-based implementations use an interpretative implementation of the reduction strategy. They are presented in Section 3.3. Finally, we compare the eval-apply and the push-enter schemes for call-by-value and we relate environment machines and graph reducers.

3.1 The Eval-ApPLY Model

In the eval-apply model, a λ -abstraction is considered as a result and the application of a function to its argument is an explicit operation. This model is the most natural choice to implement call-by-value where functions can be evaluated as arguments.

3.1.1 Call-by-value

In this scheme, applications $E_1 E_2$ are compiled by evaluating the argument E_2 , the function E_1 and finally applying the result of E_1 to the result of E_2 . Normal forms denote results; so λ -abstractions and variables (which, in strict languages, are always bound to normal forms) are transformed into results (i.e. $\mathbf{push}_s E$). The compilation of right-to-left call-by-value is formalized by the transformation $\mathcal{V}a$ in Figure 4.

This compilation choice is taken by the SECD machine [31] and the Tabac compiler [22]. The rules can be explained intuitively by reading “return the value” for \mathbf{push}_s , “evaluate” for $\mathcal{V}a$, “then” for \circ and “apply” for \mathbf{app} . Even if environment management will be tackled only in Section 4, it is also useful to keep in mind that a Λ_s -expression returning a function (such as $\mathbf{push}_s(\lambda_s x. E)$) will involve building a closure (i.e. a data structure containing the function and an environment recording the values of its free variables).

$$\begin{aligned} \mathcal{V}a &: \Lambda \rightarrow \Lambda_s \\ \mathcal{V}a \llbracket x \rrbracket &= \mathbf{push}_s x \\ \mathcal{V}a \llbracket \lambda x. E \rrbracket &= \mathbf{push}_s (\lambda_s x. \mathcal{V}a \llbracket E \rrbracket) \\ \mathcal{V}a \llbracket E_1 E_2 \rrbracket &= \mathcal{V}a \llbracket E_2 \rrbracket \circ \mathcal{V}a \llbracket E_1 \rrbracket \circ \mathbf{app} \quad \text{with } \mathbf{app} = \lambda_s f.f \end{aligned}$$

Figure 4 Compilation of right-to-left call-by-value in the eval-apply model ($\mathcal{V}a$)

Strictly speaking, $\mathcal{V}a$ does not enforce a right-to-left evaluation ($\mathcal{V}a \llbracket E_1 \rrbracket$ could be reduced before $\mathcal{V}a \llbracket E_2 \rrbracket$). However, after instantiation, the normal order of reductions will enforce the sequencing nature of “ \circ ”. It is easy to check that $\mathcal{V}a$ produces well-typed expressions of result type $R_s \sigma$ [Ⓣ].

The correctness of $\mathcal{V}a$ is stated by Property 3 which establishes that the reduction ($\xrightarrow{*}$) of transformed programs simulates the call-by-value reduction ($\xrightarrow{\text{cbv}}$) of source λ -expressions [Ⓣ]. As it is standard, we consider that the source program (i.e. the global expression) is a closed Λ -expression.

Property 3 For all closed Λ -expression E , $E \xrightarrow{\text{cbv}} V$ if and only if $\mathcal{V}a \llbracket E \rrbracket \xrightarrow{*} \mathcal{V}a \llbracket V \rrbracket$

It is clearly useless to store a function to apply it immediately after. This optimization is expressed by the following law

$$\mathbf{push}_s E \circ \mathbf{app} = E \quad (\mathbf{push}_s E \circ \lambda_s f.f =_{\beta_s} f[E/f] = E) \quad (\text{L4})$$

Example. Let $E \equiv (\lambda x.x)((\lambda y.y)(\lambda z.z))$; after simplifications, we get:

$$\begin{aligned} \mathcal{V}a \llbracket E \rrbracket &\equiv \mathbf{push}_s(\lambda_s z. \mathbf{push}_s z) \circ (\lambda_s y. \mathbf{push}_s y) \circ (\lambda_s x. \mathbf{push}_s x) \\ &\quad \rightarrow \mathbf{push}_s(\lambda_s z. \mathbf{push}_s z) \circ (\lambda_s x. \mathbf{push}_s x) \end{aligned}$$

$$\blacktriangleright \text{push}_s(\lambda_s z. \text{push}_s z) \equiv \mathcal{V}a \llbracket \lambda z. z \rrbracket$$

The source expression has two redexes $(\lambda x. x)((\lambda y. y)(\lambda z. z))$ and $(\lambda y. y)(\lambda z. z)$ but only the latter can be chosen by a call-by-value strategy. In contrast, $\mathcal{V}a \llbracket E \rrbracket$ has only the compiled version of $(\lambda y. y)(\lambda z. z)$ as redex. The illicit (in call-by-value) reduction $E \rightarrow (\lambda y. y)(\lambda z. z)$ cannot occur within $\mathcal{V}a \llbracket E \rrbracket$. This illustrates the fact that the reduction strategy has been compiled and that the choice of redex in Λ_s is not semantically relevant. \square

The law (L4) is central in the implementation of *uncurrying* (see e.g. [2]). To illustrate a simple case of uncurrying, let us take the case of a function applied to all of its arguments $(\lambda x_1 \dots \lambda x_n. E_0) E_1 \dots E_n$, then

$$\begin{aligned} & \mathcal{V}a \llbracket (\lambda x_1 \dots \lambda x_n. E_0) E_1 \dots E_n \rrbracket \\ &= \mathcal{V}a \llbracket E_n \rrbracket \circ \dots \circ \mathcal{V}a \llbracket E_1 \rrbracket \circ \text{push}_s(\lambda_s x_1 \dots (\text{push}_s(\lambda_s x_n. \mathcal{V}a \llbracket E_0 \rrbracket) \dots)) \circ \text{app} \circ \dots \circ \text{app} \end{aligned}$$

using (L4), (assoc) and (L1) this expression can be simplified into

$$= \mathcal{V}a \llbracket E_n \rrbracket \circ \dots \circ \mathcal{V}a \llbracket E_1 \rrbracket \circ (\lambda_s x_1. \lambda_s x_2 \dots \lambda_s x_n. \mathcal{V}a \llbracket E_0 \rrbracket)$$

All the **app** combinators have been statically removed. In doing so, we have avoided the construction of n intermediary closures corresponding to the n unary functions denoted by $\lambda x_1 \dots \lambda x_n. E_0$. An important point to note is that, in Λ_s , $\lambda_s x_1 \dots \lambda_s x_n. E$ denotes always a n -ary function, that is to say a function that will be applied to at least n arguments (otherwise there would be **push**'s between the λ_s 's).

There exist several variants of $\mathcal{V}a$ such as $\mathcal{V}a_L$ (used by the Cam) which implements a left-to-right call-by-value or $\mathcal{V}a_r$ (used by the SML-NJ compiler) which does not assume a data stack and disallows several pushes in a row^①.

3.1.2 Call-by-name

For call-by-name in the eval-apply model, applications $E_1 E_2$ are compiled by returning E_2 , evaluating E_1 and finally applying the evaluated function to the unevaluated argument. This choice is implemented by the call-by-need version of the Tabac compiler [22] and it is described by the transformation $\mathcal{N}a$ in Figure 5.

$$\begin{aligned} \mathcal{N}a : \Lambda &\rightarrow \Lambda_s \\ \mathcal{N}a \llbracket x \rrbracket &= x \\ \mathcal{N}a \llbracket \lambda x. E \rrbracket &= \text{push}_s(\lambda_s x. \mathcal{N}a \llbracket E \rrbracket) \\ \mathcal{N}a \llbracket E_1 E_2 \rrbracket &= \text{push}_s(\mathcal{N}a \llbracket E_2 \rrbracket) \circ \mathcal{N}a \llbracket E_1 \rrbracket \circ \text{app} \quad \text{with } \text{app} = \lambda_s f.f \end{aligned}$$

Figure 5 Compilation of call-by-name in the eval-apply model ($\mathcal{N}a$)

The correctness of $\mathcal{N}a$ is stated by Property 4 which establishes that the reduction of transformed expressions ($\xrightarrow{*}$) simulates the call-by-name reduction ($\xrightarrow{\text{cbn}}$) of source λ -expressions.

Property 4 For all closed Λ -expression E , $E \xrightarrow{\text{cbn}} V$ if and only if $\mathcal{N}a \llbracket E \rrbracket \xrightarrow{*} \mathcal{N}a \llbracket V \rrbracket$

Example. Let $E \equiv (\lambda x.x)((\lambda y.y)(\lambda z.z))$; after simplifications, we get:

$$\begin{aligned} \mathcal{N}a \llbracket E \rrbracket &\equiv \text{push}_s(\text{push}_s(\text{push}_s(\lambda_s z.z)) \circ \lambda_s y.y) \circ \lambda_s x.x \\ &\Rightarrow \text{push}_s(\text{push}_s(\lambda_s z.z)) \circ \lambda_s y.y \\ &\Rightarrow \text{push}_s(\lambda_s z.z) \equiv \mathcal{V}a \llbracket \lambda z.z \rrbracket \end{aligned}$$

The illicit (in call-by-name) reduction $E \rightarrow (\lambda x.x)(\lambda z.z)$ cannot occur within $\mathcal{N}a \llbracket E \rrbracket$. \square

Like $\mathcal{V}a$, the transformation $\mathcal{N}a$ has a variant which does not assume a data stack (i.e. disallows several pushes in a row) $\textcircled{3}$.

3.2 The Push-Enter Model

In the eval-apply model, the straightforward compilation of a function expecting n arguments produces a code building n closures. In practice, much of this overhead can be removed by uncurrying but this optimization is not always possible for functions passed as arguments. The main motivation of the push-enter model is to avoid useless closure build-ings. In the push-enter model, unevaluated functions are applied right away and application is an implicit operation.

3.2.1 Call-by-value

Instead of evaluating the function and its argument and then applying the results as in the eval-apply model, another solution is to evaluate the argument and to apply the unevaluated function right away. With call-by-value, a function can also be evaluated as an argument. In this case it cannot be immediately applied but must be returned as a result. In order to detect when its evaluation is over, there has to be a way to distinguish if its argument is present or absent: this is the role of *marks*. After a function is evaluated, a test is performed: if there is a mark, the function is returned as a result (and a closure is built), otherwise the argument is present and the function is applied. This technique avoids building some closures but at the price of performing dynamic tests. It is implemented in Zinc [32].

The mark ε is supposed to be a value that can be distinguished from others. Functions are transformed into $\text{grab}_s E$ which satisfies the reduction rules

$$\text{push}_s \varepsilon \circ \text{grab}_s E \Rightarrow \text{push}_s E$$

that is, a mark is present and the function E is returned and

$$\text{push}_s V \circ \text{grab}_s E \Rightarrow \text{push}_s V \circ E \quad (V \neq \varepsilon)$$

that is, no mark is present and the function E is applied to its argument V .

The combinator \mathbf{grab}_s and the mark ε can be defined in $\Lambda_s^{\text{①}}$. In practice, \mathbf{grab}_s is implemented using a conditional testing the presence of a mark. The transformation for right-to-left call-by-value is described in Figure 6.

$$\begin{aligned} \mathcal{V}m &: \Lambda \rightarrow \Lambda_s \\ \mathcal{V}m \llbracket x \rrbracket &= \mathbf{grab}_s x \\ \mathcal{V}m \llbracket \lambda x. E \rrbracket &= \mathbf{grab}_s (\lambda_s x. \mathcal{V}m \llbracket E \rrbracket) \\ \mathcal{V}m \llbracket E_1 E_2 \rrbracket &= \mathbf{push}_s \varepsilon \circ \mathcal{V}m \llbracket E_2 \rrbracket \circ \mathcal{V}m \llbracket E_1 \rrbracket \end{aligned}$$

Figure 6 Compilation of right-to-left call-by-value in the push-enter model ($\mathcal{V}m$)

The correctness of $\mathcal{V}m$ is stated by Property 5.

Property 5 For all closed Λ -expression E , $E \xrightarrow{\text{cbv}} V$ if and only if $\mathcal{V}m \llbracket E \rrbracket \xrightarrow{*} \mathcal{V}m \llbracket V \rrbracket$

Example. Let $E \equiv (\lambda x.x)((\lambda y.y)(\lambda z.z))$ then after simplifications

$$\begin{aligned} \mathcal{V}m \llbracket E \rrbracket &\equiv \mathbf{push}_s \varepsilon \circ \mathbf{push}_s (\lambda_s z. \mathbf{grab}_s z) \circ (\lambda_s y. \mathbf{grab}_s y) \circ (\lambda_s x. \mathbf{grab}_s x) \\ &\quad \Rightarrow \mathbf{push}_s \varepsilon \circ \mathbf{grab}_s (\lambda_s z. \mathbf{grab}_s z) \circ (\lambda_s x. \mathbf{grab}_s x) \\ &\quad \Rightarrow \mathbf{push}_s (\lambda_s z. \mathbf{grab}_s z) \circ (\lambda_s x. \mathbf{grab}_s x) \\ &\quad \Rightarrow \mathbf{grab}_s (\lambda_s z. \mathbf{grab}_s z) \equiv \mathcal{V}m \llbracket \lambda z.z \rrbracket \quad \square \end{aligned}$$

As before, when a function $\lambda x_1 \dots \lambda x_n. E$ is known to be applied to n arguments, the code can be optimized to save n dynamic tests. Actually, it appears that $\mathcal{V}m$ is subject to the same kind of optimizations as $\mathcal{V}a$. Uncurrying and related optimizations can be expressed based on the reduction rules of \mathbf{grab}_s and (L2).

It would not make much sense to consider a left-to-right strategy here. The whole point of this approach is to prevent building some closures by testing if the argument is present. Therefore the argument must be evaluated before the function. However, other closely related transformations using marks exist ^①.

3.2.2 Call-by-name

Contrary to call-by-value, the most natural choice to implement call-by-name is the push-enter model. In call-by-name, functions are evaluated only when applied to an argument. Functions do not have to be considered as results. This option is taken by Tim [20], the Krivine machine [11] and graph-based implementations (see Section 3.3.2). The transformation $\mathcal{N}m$ formalizes this choice; it is described in Figure 7.

$$\begin{aligned}
\mathcal{N}m &: \Lambda \rightarrow \Lambda_s \\
\mathcal{N}m \llbracket x \rrbracket &= x \\
\mathcal{N}m \llbracket \lambda x.E \rrbracket &= \lambda_s x. \mathcal{N}m \llbracket E \rrbracket \\
\mathcal{N}m \llbracket E_1 E_2 \rrbracket &= \mathbf{push}_s(\mathcal{N}m \llbracket E_2 \rrbracket) \circ \mathcal{N}m \llbracket E_1 \rrbracket
\end{aligned}$$

Figure 7 Compilation of call-by-name in the push-enter model ($\mathcal{N}m$)

Variables are bound to arguments which must be evaluated when accessed. Functions are not returned as results but assume that their argument is present. Applications are transformed by returning the unevaluated argument to the function. The correctness of $\mathcal{N}m$ is stated by Property 6.

Property 6 For all closed Λ -expression E , $E \xrightarrow{cm} V$ if and only if $\mathcal{N}m \llbracket E \rrbracket \xrightarrow{*} \mathcal{N}m \llbracket V \rrbracket$

Example. Let $E \equiv (\lambda x.x)((\lambda y.y)(\lambda z.z))$ then

$$\begin{aligned}
\mathcal{N}m \llbracket E \rrbracket &\equiv \mathbf{push}_s(\mathbf{push}_s(\lambda_s z.z) \circ \lambda_s y.y) \circ \lambda_s x.x \\
&\quad \rightarrow \mathbf{push}_s(\lambda_s z.z) \circ \lambda_s y.y \\
&\quad \rightarrow \lambda_s z.z \equiv \mathcal{N}m \llbracket \lambda z.z \rrbracket \quad \square
\end{aligned}$$

Arguably, $\mathcal{N}m$ is the simplest way to compile call-by-name. However, it makes the compilation of call-by-need problematic. After the evaluation of an unevaluated expression bound to a variable (i.e. a closure), a call-by-need implementation updates it by its normal form. Contrary to $\mathcal{N}a$, $\mathcal{N}m$ makes it impossible to distinguish results of closures (which have to be updated) from regular functions (which are applied right away). This problem is solved, as in $\mathcal{V}m$, with the help of marks. We come back to this issue in Section 6.

Transformations from Λ to Λ_s share the goal of compiling control with CPS transformations [21][47]. Actually, with a properly chosen instantiation of the combinators, the transformation $\mathcal{V}a_f$ is nothing but Fischer's CPS transformation [21]^①. As for CPS-expressions, it is also possible to design an inverse transformation [15] mapping Λ_s -expressions back to Λ -expressions^②.

3.3 Graph Reduction

Graph-based implementations manipulate a graph representation of the source λ -expression. The reduction consists of rewriting the graph more or less interpretatively. One of the motivations of this approach is to elegantly represent sharing which is ubiquitous in call-by-need implementations. So, even if call-by-value can be envisaged, well-known graph-based implementations only consider call-by-need. In the following, we focus on the push-enter mod-

el for call-by-name which is largely adopted by existing graph reducers. Its refinement into call-by-need is presented in Section 6.2.2.

3.3.1 Graph building

As before, the compilation of control is expressed by transformations from Λ to Λ_s . However, this step is now divided in two parts: the graph construction, then its reduction via an interpreter. The transformation \mathcal{G} (Figure 8) produces an expression which builds a graph (for now, only a tree) when reduced.

$$\begin{aligned} \mathcal{G} &: \Lambda \rightarrow \Lambda_s \\ \mathcal{G} \llbracket x \rrbracket &= \mathbf{push}_s x \circ \mathbf{mkVar}_s \\ \mathcal{G} \llbracket \lambda x. E \rrbracket &= \mathbf{push}_s(\lambda_s x. \mathcal{G} \llbracket E \rrbracket) \circ \mathbf{mkFun}_s \\ \mathcal{G} \llbracket E_1 E_2 \rrbracket &= \mathcal{G} \llbracket E_2 \rrbracket \circ \mathcal{G} \llbracket E_1 \rrbracket \circ \mathbf{mkApp}_s \end{aligned}$$

Figure 8 Generic graph building code (\mathcal{G})

The three new combinators \mathbf{mkVar}_s , \mathbf{mkFun}_s and \mathbf{mkApp}_s take their arguments from the s component and return graph nodes (respectively variable, function and application nodes) on s . The following condition formalizes the fact that the reduction of $\mathcal{G} \llbracket E \rrbracket$ is just the graph construction which terminates and yields a result in the s component.

$$(\mathit{CondG}) \quad \text{For all } \Lambda\text{-expression } E, \mathcal{G} \llbracket E \rrbracket \xrightarrow{*} \mathbf{push}_s V$$

The graph is scanned and reduced using a small interpreter denoted by the combinator \mathbf{unwind}_s . After the compilation of control, the global expression is of the form $\mathcal{G} \llbracket E \rrbracket \circ \mathbf{unwind}_s$. This transformation is common to all the graph reduction schemes we describe. The push-enter or eval-apply models of the compilation of call-by-value or call-by-name can be specified simply by defining the interactions of \mathbf{unwind}_s with the three graph builders \mathbf{mkVar}_s , \mathbf{mkFun}_s and \mathbf{mkApp}_s .

3.3.2 Call-by-name: the push-enter model

This option is defined by the three following conditions:

$$\begin{aligned} (\mathcal{GN}m1) \quad & (E \circ \mathbf{mkVar}_s) \circ \mathbf{unwind}_s = E \circ \mathbf{unwind}_s \\ (\mathcal{GN}m2) \quad & V \circ (\mathbf{push}_s F \circ \mathbf{mkFun}_s) \circ \mathbf{unwind}_s = (V \circ F) \circ \mathbf{unwind}_s \\ (\mathcal{GN}m3) \quad & (E_2 \circ E_1 \circ \mathbf{mkApp}_s) \circ \mathbf{unwind}_s = E_2 \circ E_1 \circ \mathbf{unwind}_s \end{aligned}$$

These conditions can be explained intuitively as:

- ($\mathcal{GN}1$) The reduction of a variable node amounts to reducing the graph which has been bound to the variable. The combinator \mathbf{mkVar}_s may seem useless since it is bypassed by \mathbf{unwind}_s . However, when call-by-need is considered, \mathbf{mkVar}_s is needed to implement updating without losing sharing properties. As the combinator \mathbf{I} in [53], it represents indirection nodes.
- ($\mathcal{GN}2$) The reduction of a function node amounts to applying the function to its argument and to reducing the resulting graph. This rule makes the push-enter model clear. The reduction of the function node does not return the function F as a result, but immediately applies it.
- ($\mathcal{GN}3$) The reduction of an application node amounts to storing the argument graph and to reducing the function graph.

Figure 9 presents one possible instance of the graph combinators.

$$\begin{aligned} \mathbf{mkVar}_s &= \lambda_s x. \mathbf{push}_s x \\ \mathbf{mkFun}_s &= \lambda_s f. \mathbf{push}_s (\lambda_s a. (\mathbf{push}_s a \circ f) \circ \mathbf{unwind}_s) \\ \mathbf{mkApp}_s &= \lambda_s x_1. \lambda_s x_2. \mathbf{push}_s (\mathbf{push}_s x_2 \circ x_1) \\ \mathbf{unwind}_s &= \mathbf{app} = \lambda_s x. x \end{aligned}$$

Figure 9 Instantiation of graph combinators according to \mathcal{GN} (option node-as-code)

Here, the graph is not encoded by data structures but by code performing the needed actions. For example, \mathbf{mkFun}_s takes a function f and returns a code (i.e. builds a closure) that will evaluate the function f applied to its argument a using \mathbf{unwind}_s whereas \mathbf{mkApp}_s takes two expressions x_1 and x_2 and returns a code that will apply x_1 to x_2 . This encoding simplifies the interpreter which just has to trigger a code; that is, \mathbf{unwind}_s is just an application. It is easy to check that these definitions verify the conditions (CondG), ($\mathcal{GN}1$), ($\mathcal{GN}2$), and ($\mathcal{GN}3$). Moreover, the definition of \mathbf{mkVar}_s (the identity function in Λ_s) makes it clear that indirection chains can be collapsed. That is to say,

$$\forall E \in \Lambda, \mathcal{G} \llbracket E \rrbracket \circ \mathbf{mkVar}_s = \mathcal{G} \llbracket E \rrbracket \quad (\text{L5})$$

With this combinator instantiation, the graph is represented by closures. More classical representations, based on data structures, are mentioned in Section 3.3.3. The correctness of \mathcal{G} with respect to conditions \mathcal{GN} is stated by Property 7^D.

Property 7 Let (CondG), ($\mathcal{GN}1$), ($\mathcal{GN}2$), ($\mathcal{GN}3$) and (L5) hold, then for all closed Λ -expression E , if $E \xrightarrow{\text{cbn}} V$ then $\mathcal{G} \llbracket E \rrbracket \circ \mathbf{unwind}_s = \mathcal{G} \llbracket V \rrbracket \circ \mathbf{unwind}_s$.

Compared to the corresponding properties for the previous transformations ($\mathcal{V}a$, $\mathcal{N}a$, $\mathcal{V}m$, $\mathcal{N}m$), Property 7 is expressed using equality instead of reduction ($\xrightarrow{*}$). This is because the normal form of $\mathcal{G} \llbracket E \rrbracket \circ \mathbf{unwind}_s$ may contain indirections nodes (\mathbf{mkVar}_s) and is not, in

general, syntactically identical to $\mathcal{G} \llbracket V \rrbracket \circ \mathbf{unwind}_s$. Actually, \mathcal{G} verifies a stronger (but less easily formalized) property than Property 7: $\mathcal{G} \llbracket E \rrbracket \circ \mathbf{unwind}_s$ reduces to an expression X which, after removal of indirection chains, is syntactically equal to the graph of $\mathcal{G} \llbracket V \rrbracket$.

Example. Let $E \equiv (\lambda x.x)((\lambda y.y)(\lambda z.z))$ and

$$I_w \equiv (\lambda_s a. (\mathbf{push}_s a \circ (\lambda_s w. \mathbf{push}_s w \circ \mathbf{mkVar}_s)) \circ \mathbf{unwind}_s) \text{ then}$$

$$\mathcal{G} \llbracket E \rrbracket \circ \mathbf{unwind}_s \equiv (\mathcal{G} \llbracket \lambda z.z \rrbracket \circ \mathcal{G} \llbracket \lambda y.y \rrbracket \circ \mathbf{mkApp}_s) \circ \mathcal{G} \llbracket \lambda x.x \rrbracket \circ \mathbf{mkApp}_s \circ \mathbf{unwind}_s$$

- $\xrightarrow{*} \mathbf{push}_s(\mathbf{push}_s(\mathbf{push}_s I_z \circ I_y) \circ I_x) \circ \mathbf{unwind}_s$
- $\xrightarrow{*} \mathbf{push}_s(\mathbf{push}_s I_z \circ I_y) \circ (\lambda_s a. (\mathbf{push}_s a \circ (\lambda_s x. \mathbf{push}_s x \circ \mathbf{mkVar}_s)) \circ \mathbf{unwind}_s)$
- $\xrightarrow{*} \mathbf{push}_s(\mathbf{push}_s I_z \circ I_y) \circ \mathbf{unwind}_s$
- $\xrightarrow{*} \mathbf{push}_s I_z \circ (\lambda_s a. (\mathbf{push}_s a \circ (\lambda_s y. \mathbf{push}_s y \circ \mathbf{mkVar}_s)) \circ \mathbf{unwind}_s)$
- $\xrightarrow{*} (\mathbf{push}_s I_z \circ \mathbf{mkVar}_s) \circ \mathbf{unwind}_s \xrightarrow{*} \mathbf{push}_s I_z \circ \mathbf{unwind}_s$

In this example, there is no indirection chain and the result is syntactically equal to the graph of the source normal form. That is, $\mathbf{push}_s I_z \circ \mathbf{unwind}_s$ is exactly $\mathcal{G} \llbracket \lambda z.z \rrbracket \circ \mathbf{unwind}_s$ after the few reductions corresponding to graph construction.

The first sequence of reductions corresponds to the graph construction. Then \mathbf{unwind}_s scans the (leftmost) spine (the first \mathbf{push}_s represents an application node). The graph representing the function $(\lambda x.x)$ is applied. The result is the application node $\mathbf{push}_s(\mathbf{push}_s I_z \circ I_y)$ which is scanned by \mathbf{unwind}_s . Then, the reduction proceeds in the same way until it reaches the normal form. \square

Because of the interpretative essence of the graph reduction, a naive implementation of call-by-need is possible without introducing marks (as opposed to $\mathcal{N}m$ in Section 3.2.2). Such a scheme performs many useless updates some of which can be detected by simple syntactic criteria or a sharing analysis. An optimized implementation, performing selective updates, can be defined by introducing marks. These two points are presented in Section 6.2.2.

3.3.3 Other choices

A graph and its associated reducer can be seen as an abstract data type with different implementations [41]. We have already used one encoding that represents nodes by code (i.e. closures). Another natural solution is to represent the graph by a data structure. It amounts to introducing three data constructors **VarNode**, **FunNode** and **AppNode** and to defining the interpreter \mathbf{unwind}_s by a case expression. A refinement, exploited by the G-machine, is to enclose in nodes the code to be executed when it is unwound. Adding code in data structures comes very close to the solution using closures described in Figure 9. The interpreter \mathbf{unwind}_s can just execute the code and does not have to perform a dynamic test. In any case, the new combinator definitions should still verify the $\mathcal{G}\mathcal{N}m$ properties in order to implement a push-enter model of the compilation of call-by-name.

By far, the most common use of graph reduction is the implementation of call-by-need in the push-enter model. However, the eval-apply model or the compilation of call-by-value can be expressed as well. These choices are specified by redefining the interactions of **unwind**_s with the three graph builders (**mkVar**_s, **mkFun**_s, **mkApp**_s). In each case, it amounts to defining new properties like ($\mathcal{GN}1$), ($\mathcal{GN}2$), and ($\mathcal{GN}3$).

More details on these alternate choices can be found in [18].

3.4 Comparisons

We compare the efficiency of codes produced by transformations $\mathcal{V}a$ (eval-apply CBV) and $\mathcal{V}m$ (push-enter CBV). Then, we exhibit the precise relationship between the environment and graph approaches. In particular, it is shown how to derive the transformation $\mathcal{N}m$ from \mathcal{G} and the properties ($\mathcal{GN}i$). We take only these two examples to show the advantages of a unified framework in terms of formal comparisons. It should be clear that such comparisons could be carried on for other transformations and compilation steps.

3.4.1 $\mathcal{V}a$ versus $\mathcal{V}m$

Let us first emphasize that our comparisons focus on finding complexity upper bounds. They do not take the place of benchmarks which are still required to take into account complex implementation aspects (e.g. interactions with memory cache or the garbage collector).

A code produced by $\mathcal{V}m$ builds less closures than the corresponding $\mathcal{V}a$ -code. Since a mark can be represented by one bit (e.g. in a bit stack parallel to the data stack), $\mathcal{V}m$ is likely to be, on average, more efficient with respect to space resources. Concerning time efficiency, the size of compiled expressions provides a first approximation of the cost entailed by the encoding of the reduction strategy (assuming **push**_s, **grab**_s and **app** have a constant time implementation). It is easy to show that code expansion is linear with respect to the size of the source expression. More precisely, for $\mathcal{V}\chi = \mathcal{V}a$ or $\mathcal{V}m$, we have

$$\text{If } \text{Size}(E) = n \text{ then } \text{Size}(\mathcal{V}\chi[E]) < 3n.$$

This upper bound can be reached by taking for example $E \equiv \lambda x.x \dots x$ (n occurrences of x). A more thorough investigation is possible by associating costs with the different combinators encoding the control: *push* for the cost of “pushing” a variable or a mark, *clos* for the cost of building a closure (i.e. **push**_s E), *app* and *grab* for the cost of the corresponding combinators. If we take n_λ for the number of λ -abstractions and n_v for the number of occurrences of variables in the source expression, we have

$$\text{Cost}(\mathcal{V}a[E]) = n_\lambda \text{ clos} + n_v \text{ push} + (n_v - 1) \text{ app}$$

and

$$\text{Cost}(\mathcal{V}m[E]) = (n_\lambda + n_v) \text{ grab} + (n_v - 1) \text{ push}$$

The benefit of $\mathcal{V}m$ over $\mathcal{V}a$ is to sometimes replace a (useless) closure construction by a test. When a closure has to be built, $\mathcal{V}m$ involves a useless test compared to $\mathcal{V}a$. So if *clos* is comparable to the cost of a test (for example, when returning a closure amounts to building a

pair as in Section 4.1.2) $\mathcal{V}m$ will produce more expensive code than $\mathcal{V}a$. If closure building is not a constant time operation (as in Section 4.1.3) $\mathcal{V}m$ can be arbitrarily better than $\mathcal{V}a$. Actually, it can change the program complexity in contrived cases. In practice, however, the situation is not so clear. When no mark is present, **grab**_s is implemented by a test followed by an **app**. If a mark is present, the test is followed by a **push**_s (i.e. a closure building for λ -abstractions). So, we have

$$Cost(\mathcal{V}m \llbracket E \rrbracket) = (n_\lambda + n_v) test + p (n_\lambda + n_v) app + p n_\lambda clos + p n_v push + (n_v - 1) push$$

with p (resp. p) representing the likelihood ($p+p=1$) of the presence (resp. absence) of a mark which depends on the program. The best situation for $\mathcal{V}m$ is when no closure has to be built, that is $p=0$ and $p=1$. If we take some reasonable hypothesis such as $test=app$ and $n_\lambda < n_v < 3n_\lambda$, we find that the cost of closure construction must be 3 to 5 times more costly than app or $test$ to make $\mathcal{V}m$ advantageous. With less favorable odds such as $p=p=1/2$, $clos$ must be worth 7 or 8 app .

We are led to conclude that $\mathcal{V}m$ should be considered only when closure building is potentially costly (such as the $\mathcal{A}c2$ transformation in Section 4.1.3 which builds closures by copying part of the environment). Even so, tests may be too costly in practice compared to the construction of small closures. The best way would probably be to perform an analysis to detect cases when $\mathcal{V}m$ is profitable. Such information could be taken into account to get the best of each approach. We present in [17] how $\mathcal{V}a$ and $\mathcal{V}m$ could be mixed.

3.4.2 Environment machine versus graph reducer

Even if their starting points are utterly different, graph reducers and environment machines can be related. This has been done for specific implementations such as [43] which shows how to transform a G-machine into a Tim. We focus here on the compilation of control and compare the transformation $\mathcal{N}m$ with the $\mathcal{G}\mathcal{N}m$ approach to graph reduction.

The two main departures of graph reduction from the environment approach are

- *The potentially useless graph constructions.* For example, the rule $\mathcal{G} \llbracket E_1 E_2 \rrbracket = \mathcal{G} \llbracket E_2 \rrbracket \circ \mathcal{G} \llbracket E_1 \rrbracket \circ \mathbf{mkApp}_s$ builds a graph for E_2 even if E_2 is never reduced (i.e. if it is not needed). On the other hand, $\mathcal{N}m$ suspends all operations (such as variable instantiation) on E_2 by building a closure ($\mathcal{N}m \llbracket E_1 E_2 \rrbracket = \mathbf{push}_s (\mathcal{N}m \llbracket E_2 \rrbracket) \circ \mathcal{N}m \llbracket E_1 \rrbracket$).
- *The interpretative nature of graph reduction.* Even in the “node-as-code” instantiation, each application node (**mkApp**_s) is “interpreted” by **unwind**_s. In the environment family, no interpreter is needed and this approach can be seen as the specialization of the interpreter **unwind**_s according to the source graph built by $\mathcal{G} \llbracket \rrbracket$.

In order to formalize these two points, we first change the rule for graph building in the case of applications by

$$\mathcal{G} \llbracket E_1 E_2 \rrbracket = \mathbf{push}_s (\mathcal{G} \llbracket E_2 \rrbracket \circ \mathbf{unwind}_s) \circ \mathcal{G} \llbracket E_1 \rrbracket \circ \mathbf{mkApp}_s$$

This corresponds to a lazy graph construction where the graph argument is built only if needed. In particular, variables will be bound to unbuilt graphs. This new kind of graph entails replacing property $(\mathcal{GN}1)$ with

$$(\mathcal{GN}1) \quad (\mathbf{push}_s E \circ \mathbf{mkVar}_s) \circ \mathbf{unwind}_s = E$$

We can now show that $\mathcal{N} \llbracket E \rrbracket$ is merely the specialization of \mathbf{unwind}_s with respect to the graph of E ; that is

$$\mathcal{N} \llbracket E \rrbracket = \mathcal{G} \llbracket E \rrbracket \circ \mathbf{unwind}_s$$

For example, the specialization for the application case is:

$$\begin{aligned} & \mathcal{G} \llbracket E_1 E_2 \rrbracket \circ \mathbf{unwind}_s \\ &= \mathbf{push}_s (\mathcal{G} \llbracket E_2 \rrbracket \circ \mathbf{unwind}_s) \circ \mathcal{G} \llbracket E_1 \rrbracket \circ \mathbf{mkApp}_s \circ \mathbf{unwind}_s \quad (\text{unfolding } \mathcal{G}) \\ &= \mathbf{push}_s (\mathcal{G} \llbracket E_2 \rrbracket \circ \mathbf{unwind}_s) \circ \mathcal{G} \llbracket E_1 \rrbracket \circ \mathbf{unwind}_s \quad (\mathcal{GN}3) \\ &= \mathbf{push}_s (\mathcal{N} \llbracket E_2 \rrbracket) \circ \mathcal{N} \llbracket E_1 \rrbracket \quad (\text{induction hypothesis}) \\ &= \mathcal{N} \llbracket E_1 E_2 \rrbracket \quad (\text{folding } \mathcal{N}) \quad \square \end{aligned}$$

This property shows that, as far as the compilation of control is concerned, environment based transformations are more efficient than their graph counterpart. However, optimized graph reducers avoid as much as possible interpretative scans of the graph or graph building and are similar to environment-based implementations.

4 COMPILATION OF THE β -REDUCTION

This compilation step implements the substitution using transformations from Λ_s to Λ_e . These transformations are akin to abstraction algorithms and consist of replacing variables with combinators. Compared to Λ_s , Λ_e adds the pair $(\mathbf{push}_e, \lambda_e)$ encoding an environment component and it uses variables only to define combinators. Graph reducers use specific (usually environment-less) transformations. We express in our framework the **SKI** abstraction algorithm (Section 4.2).

4.1 Environment Based Abstractions

In the λ -calculus, the β -reduction is defined as a textual substitution. In environment-based implementations, substitutions are compiled by storing the value to be substituted in a data structure (an environment). Values are then accessed in the environment only when needed. This technique can be compared with the activation records used by imperative language compilers. The main choice is using list-like (shared) environments or vector-like (copied) environments. For the latter choice, there are several transformations depending when the environments are copied.

4.1.1 A generic abstraction

The denotational-like transformation $\mathcal{A}g$ (Figure 10) is a generic abstraction which will be specialized to model several choices in the following sections. It introduces an environment where the values of variables are stored and fetched from. The transformation is done with respect to a compile-time environment ρ (initially empty for a closed expression). We note x_i the variable occurring at the i th entry in the environment.

$$\begin{aligned} \mathcal{A}g : \Lambda_s &\rightarrow env \rightarrow \Lambda_e \\ \mathcal{A}g \llbracket E_1 \circ E_2 \rrbracket \rho &= \mathbf{dupl}_e \circ \mathcal{A}g \llbracket E_1 \rrbracket \rho \circ \mathbf{swap}_{se} \circ \mathcal{A}g \llbracket E_2 \rrbracket \rho \\ \mathcal{A}g \llbracket \mathbf{push}_s E \rrbracket \rho &= \mathbf{push}_s (\mathcal{A}g \llbracket E \rrbracket \rho) \circ \mathbf{mkclos} \\ \mathcal{A}g \llbracket \lambda_s x. E \rrbracket \rho &= \mathbf{mkbind} \circ \mathcal{A}g \llbracket E \rrbracket (\rho.x) \\ \mathcal{A}g \llbracket x_i \rrbracket (\dots((\rho.x_i), x_{i-1}) \dots, x_0) &= \mathbf{access}_i \circ \mathbf{appclos} \end{aligned}$$

Figure 10 A generic abstraction ($\mathcal{A}g$)

$\mathcal{A}g$ needs six new combinators to express environment saving and restoring (\mathbf{dupl}_e , \mathbf{swap}_{se}), closure building and calling (\mathbf{mkclos} , $\mathbf{appclos}$), access to values (\mathbf{access}_i) and adding a binding (\mathbf{mkbind}).

The first combinator pair (\mathbf{dupl}_e , \mathbf{swap}_{se}) is defined in Λ_e by

$$\mathbf{dupl}_e = \lambda_e e. \mathbf{push}_e e \circ \mathbf{push}_e e \quad \mathbf{swap}_{se} = \lambda_s x. \lambda_e e. \mathbf{push}_s x \circ \mathbf{push}_e e$$

Note that \mathbf{swap}_{se} is needed only if s and e are implemented by a single component. In our approach, this choice is made in the final implementation step (see Section 2.5). If eventually e and s are implemented by, say, two distinct stacks then new algebraic simplifications become valid; in particular \mathbf{swap}_{se} can be removed (its definition as a λ -expression will be the identity function).

The closure combinators (\mathbf{mkclos} , $\mathbf{appclos}$) can have different definitions in Λ_e as long as they satisfy the condition

$$(\mathbf{push}_e E \circ \mathbf{push}_s X \circ \mathbf{mkclos}) \circ \mathbf{appclos} \xrightarrow{+} \mathbf{push}_e E \circ X$$

That is, evaluating a closure made of the function X and environment E amounts to evaluating X with the environment E . For example, two possible definitions are

$$\begin{aligned} \mathbf{mkclos} &= \lambda_s x. \lambda_e e. \mathbf{push}_s(x, e) & \mathbf{appclos} &= \lambda_s(x, e). \mathbf{push}_e e \circ x \\ \text{or } \mathbf{mkclos} &= \lambda_s x. \lambda_e e. \mathbf{push}_s(\mathbf{push}_e e \circ x) & \mathbf{appclos} &= \mathbf{app} = \lambda_s x. x \end{aligned}$$

The first option uses pairs and is, in a way, more concrete than the other one. The second option abstracts from representation considerations. It simplifies the expression of correctness properties and it will be used in the rest of the paper.

In the same way, the environment combinators (**mkbind**, **access_i**) can have several instantiations in Λ_e . Let us note **combⁱ** the sequence **comb** $\circ \dots \circ$ **comb** (*i* times), then the definitions of **mkbind** and **access_i** must satisfy the condition

$$(\mathbf{push}_s X_0 \circ \dots \circ \mathbf{push}_s X_i \circ \mathbf{push}_e E \circ \mathbf{mkbind}^{i+1}) \circ \mathbf{access}_i^+ \rightarrow \mathbf{push}_s X_i$$

This property simply says that adding *i*+1 bindings X_i, \dots, X_0 in an environment *E* then accessing the *i*th value is equivalent to returning directly X_i . Examples of definitions for **mkbind** and **access_i** appear in Figure 11 and Figure 12.

The transformation $\mathcal{A}g$ can be optimized by adding the rules

$$\mathcal{A}g \llbracket E \circ \mathbf{app} \rrbracket \rho = \mathcal{A}g \llbracket E \rrbracket \rho \circ \mathbf{appclos}$$

$$\mathcal{A}g \llbracket \lambda_s x. E \rrbracket \rho = \mathbf{pop}_{se} \circ \mathcal{A}g \llbracket E \rrbracket \rho \quad \text{if } x \text{ not free in } E \quad \text{with } \mathbf{pop}_{se} = \lambda_e e. \lambda_s x. \mathbf{push}_e e$$

Variables are bound to closures stored in the environment. With the original rules, $\mathcal{A}g \llbracket \mathbf{push}_s x_i \rrbracket$ would build yet another closure. This useless “boxing”, which may lead to long indirection chains, is avoided by the following rule:

$$\mathcal{A}g \llbracket \mathbf{push}_s x_i \rrbracket (\dots((\rho, x_i), x_{i-1}), \dots, x_0) = \mathbf{access}_i$$

Whether this new rule duplicates the closure or only its address depends on the memory management (Section 6). In call-by-need, one has to make sure that **access_i** returns the address of the closure since closure duplication may entail a loss of sharing.

4.1.2 Shared environments

A first choice is to instantiate $\mathcal{A}g$ with linked environments. The structure of the environment is a tree of closures and a closure is added to the environment in constant time. On the other hand, a chain of links has to be followed when accessing a value. The access time complexity is $O(n)$ where *n* is the number of λ_s 's from the occurrence of the variable to its binding λ_s (i.e. its de Bruijn index). This specialization, noted $\mathcal{A}s$, is used by the Cam [10], the SECD [31] and the strict and lazy versions of the Krivine machine [32][11].

Specializing $\mathcal{A}g$ into $\mathcal{A}s$ amounts to defining the environment combinators as follows

mkbind = $\lambda_e e. \lambda_s x. \mathbf{push}_e(e, x)$	access_i = fstⁱ \circ snd
<i>with</i> cⁱ = c $\circ \dots \circ$ c (<i>i</i> times)	fst = $\lambda_e(e, x). \mathbf{push}_e e$ snd = $\lambda_e(e, x). \mathbf{push}_s x$

Figure 11 Combinator instantiation for shared environments ($\mathcal{A}s$)

Example. $\mathcal{A}_s \llbracket \lambda_s x_1. \lambda_s x_0. \text{push}_s E \circ x_1 \rrbracket \rho = \mathbf{mkbind} \circ \mathbf{mkbind} \circ \mathbf{dupl}_e \circ$

$$\mathbf{push}_s (\mathcal{A}_s \llbracket E \rrbracket ((\rho, x_1), x_0)) \circ \mathbf{mkclos} \circ \mathbf{swap}_{se} \circ \mathbf{access}_1 \circ \mathbf{appclos}$$

Two bindings are added ($\mathbf{mkbind} \circ \mathbf{mkbind}$) to the current environment and the x_1 access is coded by $\mathbf{access}_1 = \mathbf{fst} \circ \mathbf{snd}$. \square

The correctness of \mathcal{A}_s is stated by Property 8 [Ⓜ].

Property 8 For all closed well-typed Λ_s -expression E , $\mathbf{push}_e () \circ \mathcal{A}_s \llbracket E \rrbracket () = E$

4.1.3 Copied environments

Another choice is to provide a constant access time. In this case, the structure of the environment must be a vector of closures. A code copying the environment (a $O(\text{length } \rho)$ operation) has to be inserted in \mathcal{A}_g in order to avoid links. This scheme is less prone to space leaks since it permits suppressing useless variables during copies.

The macro-combinator **Copy** ρ produces code performing this copy according to ρ 's structure.

$$\mathbf{Copy} (\dots ((\rho, x_n), \dots, x_0)) = (\mathbf{dupl}_e \circ \mathbf{access}_n \circ \mathbf{swap}_{se}) \circ \dots$$

$$\circ (\mathbf{dupl}_e \circ \mathbf{access}_1 \circ \mathbf{swap}_{se}) \circ \mathbf{access}_0 \circ \mathbf{push}_e () \circ \mathbf{mkbind}^{n+1}$$

The combinators \mathbf{dupl}_e and \mathbf{swap}_{se} are needed to pass the environment to each \mathbf{access}_i which will store each value of the environment in s . With all the values in s , a fresh copy of the environment can be built (using $\mathbf{push}_e () \circ \mathbf{mkbind}^{n+1}$). If we still see the structure of the environment as a tree of closures, the effect of **Copy** ρ is to prevent sharing. Environments can thus be represented by vectors. The combinator \mathbf{mkbind} now adds a binding in a vector and \mathbf{access}_i becomes a constant time operation (Figure 12).

$$\mathbf{mkbind} = \lambda_e e. \lambda_s x. \mathbf{push}_e (e[\text{next}] := x) \qquad \mathbf{access}_i = \lambda_e e. \mathbf{push}_s (e[i])$$

where $e[\text{next}] := x$ adds the value x in the first empty cell of the vector e

Figure 12 Combinators instantiation for abstraction with copied environments (\mathcal{A}_c)

The index *next* designates the first free cell in the vector. It can be statically computed as the rank of the variable (associated with the \mathbf{mkbind} occurrence) in the static environment ρ . For example, in

$$\mathcal{A}_c \llbracket \lambda_s y. E \rrbracket (((\rho, x_2), x_1), x_0) = \mathbf{mkbind} \circ \mathcal{A}_c \llbracket E \rrbracket (((\rho, x_2), x_1), x_0), y)$$

we have $\text{next} = \text{rank } y \text{ } (((\rho, x_2), x_1), x_0), y) = 4$, and y is stored in the fourth cell of the environment. The maximum size of each vector can be statically calculated too. To simplify the presentation, we leave these administrative tasks implicit.

There are several abstractions according to the time of the copies. We present them by indicating only the rules that differ from $\mathcal{A}g$. A first solution (Figure 13) is to copy the environment just before adding a new binding (as in [20][46]). From the first compilation step we know that n-ary functions $(\lambda_s x_1 \dots \lambda_s x_n . E)$ are fully applied and cannot be shared: they need only one copy of the environment. The overhead is placed on function entry and closure building remains a constant time operation. The transformation $\mathcal{A}c1$ produces (possibly oversized) environments which can be shared by several closures but only as a whole. So, there must be an indirection when accessing the environment. The environment ρ represents ρ restricted to variables occurring free in the subexpression E .

$$\mathcal{A}c1 \llbracket \lambda_s x_i \dots \lambda_s x_0 . E \rrbracket \rho = \mathbf{Copy} \rho \circ \mathbf{mkbind}^{i+1} \circ \mathcal{A}c1 \llbracket E \rrbracket (\dots(\rho, x_i) \dots, x_0)$$

Figure 13 Copy at function entry ($\mathcal{A}c1$)

Example. $\mathcal{A}c1 \llbracket \lambda_s x_1 \lambda_s x_0 . \mathbf{push}_s E_1 \circ x_j \rrbracket \rho = \mathbf{Copy} \rho \circ \mathbf{mkbind}^2 \circ \mathbf{dupl}_e \circ$

$$\mathbf{push}_s (\mathcal{A}c1 \llbracket E \rrbracket ((\rho, x_1), x_0))) \circ \mathbf{mkclos} \circ \mathbf{swap}_{se} \circ \mathbf{access}_1 \circ \mathbf{appclos}$$

The code builds a vector environment made of a specialized copy of the previous environment and two new bindings (\mathbf{mkbind}^2); the x_1 access is now coded by a constant time \mathbf{access}_1 . □

A second solution (Figure 14) is to copy the environment when building and opening closures (as in [22]). The copy at opening time is necessary in order to be able to add new bindings in contiguous memory (the environment has to remain a vector). The transformation $\mathcal{A}c2$ produces environments which cannot be shared but may be accessed directly (they can be packaged with a code pointer to form a closure).

$$\mathcal{A}c2 \llbracket \mathbf{push}_s E \rrbracket \rho = \mathbf{Copy} \rho \circ \mathbf{push}_s (\mathbf{Copy} \rho \circ \mathcal{A}c2 \llbracket E \rrbracket \rho) \circ \mathbf{mkclos}$$

Figure 14 Copy at closure building and opening ($\mathcal{A}c2$)

A refinement of this last option, the $\mathcal{A}c3$ abstraction[Ⓣ], is to copy the environment only when building closures. Variations of $\mathcal{A}c3$ are used in the SML-NJ compiler [2] and the spineless tagless G-machine [42]. In order to be able to add new bindings after closure opening, an additional local environment is needed.

Starting from different properties a collection of abstractions can be systematically derived from $\mathcal{A}g$. Some of these abstractions are new, some have already been used in well-known implementations. For example, starting from the equation $\mathcal{A}g_s \llbracket E \rrbracket \rho = \mathbf{swap}_n \circ \mathcal{A}g \llbracket E \rrbracket \rho$ one can derive the swap-less transformation $\mathcal{A}g_s$. With this variation, the references to environments stay at a fixed distance from the bottom of the stack until they are popped (the references are no more **swapped**). These variations introduce different environment ma-

nipulation schemes avoiding stacks elements reordering (swap-less), environment duplication (dupl-less), environment building (mkbind-less) or closure building (mkclos-less)[Ⓞ].

4.1.4 Comparison

Assuming each basic combinator can be implemented in constant time, the size of the abstracted expressions gives an approximation of the overhead entailed by the encoding of the β -reduction. It is easy to show that $\mathcal{A}s$ entails a code expansion which is quadratic with respect to the size of the source expression. More precisely

$$\text{if } \text{Size}(E) = n \text{ then } \text{Size}(\mathcal{A}s(\mathcal{V}a \llbracket E \rrbracket)) \leq n_\lambda n_v + 6n + 6$$

with n_λ the number of λ -abstractions and n_v the number of variable occurrences ($n = n_\lambda + n_v$) of the source expression. This expression reaches a maximum with $n_v = (n-1)/2$. This upper bound can be approached with, for example, $\lambda x_1 \dots \lambda x_{n_\lambda} . x_1 \dots x_{n_\lambda}$. The product $n_\lambda n_v$ indicates that the efficiency of $\mathcal{A}s$ depends equally on the number of accesses (n_v) and their length (n_λ). For $\mathcal{A}c1$ we have

$$\text{if } \text{Size}(E) = n \text{ then } \text{Size}(\mathcal{A}c1(\mathcal{V}a \llbracket E \rrbracket)) \leq 6n_\lambda^2 - 6n_\lambda + 7n + 6$$

which makes clear that the efficiency of $\mathcal{A}c1$ is not dependent of accesses. The two transformations have the same complexity order, nevertheless one may be more adapted than the other to individual source expressions. These complexities highlight the main difference between shared environments that favors building, and copied environments that favors access. Let us point out that these bounds are related to the quadratic growth implied by Turner's abstraction algorithm [53]. Balancing expressions reduces this upper bound to $O(n \log n)$ [28]. It is very likely that this technique could also be applied to λ -expressions to get a $O(n \log n)$ complexity for environment management.

The abstractions can be compared according to their memory usage too. $\mathcal{A}c2$ copies the environment for every closure, where $\mathcal{A}c1$ may share a bigger copy. So, the code generated by $\mathcal{A}c2$ consumes more memory and implies frequent garbage collections whereas the code generated by $\mathcal{A}c1$ may create space leaks and needs special tricks to plug them (see [43] section 4.2.6).

4.2 A SKI Abstraction Algorithm

Some abstraction algorithms do not use the environment notion, but encode separately every substitution. A simple algorithm [13] uses only three combinators $\{\mathbf{S}, \mathbf{K}, \mathbf{I}\}$ but is inefficient with respect to code expansion. Different refinements, which use extended combinators families (e.g. $\{\mathbf{S}, \mathbf{K}, \mathbf{I}, \mathbf{B}, \mathbf{C}, \mathbf{S}', \mathbf{B}', \mathbf{C}'\}$), have been proposed [28][53][54]. They usually lower the complexity of code expansion from exponential with $\{\mathbf{S}, \mathbf{K}, \mathbf{I}\}$ to quadratic or even $O(n \log n)$. We describe only the SKI abstraction algorithm in this paper. It should be clear that the optimized versions could be expressed as easily in our framework.

It is possible to define a transformation $\mathcal{SKI} \llbracket E \rrbracket x$ that can be applied to all Λ_s -expressions ([18]). In particular, it can be composed with the transformations for the compilation

of the graph reduction control (Section 3.3). The resulting code, although correct, does not accurately model the classical compilation scheme of the SKI-machine. The easiest way to model it precisely is to define a transformation specialized to graph code (Figure 15).

$$\begin{aligned}
SKI: \Lambda_s &\rightarrow var \rightarrow \Lambda_e \\
SKI \llbracket E \rrbracket x &= E \circ (\mathbf{push}_s \mathbf{Ks} \circ \mathbf{mkFun}_s) \circ \mathbf{mkApp}_s && x \text{ not free in } E \\
SKI \llbracket E_1 \circ E_2 \circ \mathbf{mkApp}_s \rrbracket x & \\
&= SKI \llbracket E_1 \rrbracket x \circ (SKI \llbracket E_2 \rrbracket x \circ (\mathbf{push}_s \mathbf{Ss} \circ \mathbf{mkFun}) \circ \mathbf{mkApp}_s) \circ \mathbf{mkApp}_s \\
SKI \llbracket \mathbf{push}_s (\lambda_s y. E) \circ \mathbf{mkFun}_s \rrbracket x &= SKI \llbracket SKI \llbracket E \rrbracket y \rrbracket x \\
SKI \llbracket \mathbf{push}_s x \circ \mathbf{mkVar}_s \rrbracket x &= \mathbf{push}_s \mathbf{Is} \circ \mathbf{mkFun}_s
\end{aligned}$$

Figure 15 Abstraction SKI (SKI)

The **Ss**, **Ks** and **Is** combinators build or select a graph. They can be defined as

$$\begin{aligned}
\mathbf{Ss} &= \lambda_s e_2. \lambda_s e_1. \lambda_s x. (\mathbf{push}_s x \circ \mathbf{push}_s e_1 \circ \mathbf{mkApp}_s) \circ (\mathbf{push}_s x \circ \mathbf{push}_s e_2 \circ \mathbf{mkApp}_s) \circ \mathbf{mkApp}_s \\
\mathbf{Ks} &= \lambda_s e. \lambda_s x. \mathbf{push}_s e && \mathbf{Is} = \lambda_s x. \mathbf{push}_s x
\end{aligned}$$

In the same way, the transformation \mathcal{A}_{dsb} (a **dupl**-less, **swap**-less and **mkbind**-less abstraction algorithm) can be specialized for graph code ([18]). It would then precisely model the classical abstraction of the G-machine ([27]).

5 COMPILATION OF CONTROL TRANSFERS

A conventional machine executes linear sequences of basic instructions. In our framework, reducing expressions of the form **appclos** $\circ E$ involves evaluating a closure and then returning to E . We have to make calls and returns explicit. We present here two solutions.

$$\begin{aligned}
S: \Lambda_e &\rightarrow \Lambda_k && \text{with } i \equiv s, e \\
S \llbracket E_1 \circ E_2 \rrbracket &= \mathbf{push}_k (S \llbracket E_2 \rrbracket) \circ S \llbracket E_1 \rrbracket \\
S \llbracket \mathbf{push}_i E \rrbracket &= \mathbf{push}_i (S \llbracket E \rrbracket) \circ \mathbf{rts}_i && \text{with } \mathbf{rts}_i = \lambda_p x. \lambda_k k. \mathbf{push}_i x \circ k \\
S \llbracket \lambda_p x. E \rrbracket &= \lambda_p x. S \llbracket E \rrbracket \\
S \llbracket x \rrbracket &= x
\end{aligned}$$

Figure 16 General compilation of control transfers (S)

The first solution, adopted by most implementations, is to save the return address on a call stack k . The transformation \mathcal{S} (Figure 16) saves the code following the function call using \mathbf{push}_k , and returns to it with \mathbf{rts}_i ($= \lambda_r.x.\lambda_k.f.\mathbf{push}_i x \circ f$ and $i \equiv s$ or e) when the function ends. Intuitively these combinators can be seen as implementing a control stack. Compared to Λ_e , Λ_k -expressions do not have $\mathbf{appclos} \circ E$ code sequences. The correctness \mathcal{S} of is stated by Property 9.

Property 9 *For all closed well-typed Λ_e -expression E and N a normal form,*

$$\text{if } E \xrightarrow{*} N \text{ then } \mathcal{S}[\![E]\!] \xrightarrow{*} \mathcal{S}[\![N]\!]$$

An optimized version of \mathcal{S} for the different previous transformations could easily be derived. For example, we get

$$\mathcal{S}[\![\mathbf{dupl}_e \circ E_1 \circ \mathbf{swap}_{se} \circ E_2]\!] = \mathbf{dupl}_e \circ \mathbf{push}_k (\mathbf{swap}_{se} \circ \mathcal{S}[\![E_2]\!]) \circ \mathbf{swap}_{ke} \circ \mathcal{S}[\![E_1]\!]$$

The second solution is to use a transformation \mathcal{S} between the control and the abstraction phases ($\mathcal{S}: \Lambda_s \rightarrow \Lambda_s$). It transforms the expression into CPS. The continuation k encodes return addresses and will be abstracted as an ordinary variable. Let us present only two transformation rules

$$\mathcal{S}[\![\mathbf{push}_s E]\!] = \lambda_s.k.\mathbf{push}_s (\mathcal{S}[\![E]\!]) \circ k$$

$$\mathcal{S}[\![E_1 \circ E_2]\!] = \lambda_s.k.\mathbf{push}_s (\mathbf{push}_s k \circ \mathcal{S}[\![E_2]\!]) \circ \mathcal{S}[\![E_1]\!]$$

The first one replaces returns by continuation calls, and the second rule encodes the return stack of \mathcal{S} by a continuation composition. This solution is used in the SML-NJ compiler [2].

6 SHARING AND UPDATES

The call-by-need strategy is an optimization of the call-by-name strategy which shares and updates closures. In order to express sharing, we introduce a memory component to store closures. The evaluation of an unevaluated argument amounts to accessing a closure in the memory, to reducing it and to updating the memory with its normal form. This way, every argument is reduced at most once. The new intermediate language Λ_h adds to Λ_k the combinator pair $(\mathbf{push}_h, \lambda_h)$ which specifies a memory component h . This component is represented and accessed via a heap pointer. A first transformation \mathcal{Hc} from Λ_k to Λ_h threads the component h in which closures are allocated and accessed. Then we express updating and present several options specific to graph reduction.

6.1 Introduction of a Heap

The transformation \mathcal{Hc} (Figure 17) introduces a new component h , which encodes a heap threaded through the expression. Throughout the reduction of such an expression, there is only one reference to the heap (i.e. h is single-threaded [48]).

The transformed expression $\mathcal{H}c \llbracket E \rrbracket$ takes the heap as an argument and returns the heap as result. The last two rules of $\mathcal{H}c$ are responsible for making closure allocation and access explicit. In our framework, constructions of updatable closures are of the form $\mathbf{push}_s E$ with $E : \mathbf{R}_s \sigma$, and accesses of updatable closures are of the form $x : \mathbf{R}_s \tau$ where x is bound by a λ_s . These rules use two contexts. The context $\mathbf{Store}[E]$ can be read as: allocate a new cell in the heap, write the code E in this cell, return its address a and the heap. The context $\mathbf{Call}[E]$ can be read as: access the expression stored in the heap in the cell of address E , then reduce it (with the heap as an argument). Henceforth, the argument of a function is a closure address rather than the closure itself. A closure address is represented by an integer and the heap is represented by a pair made of a list of written cells and the address of the next free cell $((tail, \{add, val\}), free)$. The initial empty heap is noted \mathbf{emptyH} and is defined as $(((), 0))$. The three combinators \mathbf{alloc} , \mathbf{write} and \mathbf{read} perform basic heap manipulations. Since h is single-threaded, these combinators can be implemented efficiently as constant time operators on a mutable data structure.

$$\begin{array}{ll}
\mathcal{H}c : \Lambda_k \rightarrow \Lambda_h & \text{with } i \equiv s, e \text{ or } k \quad \text{and } h \text{ a fresh variable} \\
\mathcal{H}c \llbracket E_1 \circ E_2 \rrbracket = \mathcal{H}c \llbracket E_1 \rrbracket \circ \mathcal{H}c \llbracket E_2 \rrbracket & \\
\mathcal{H}c \llbracket \lambda_{\tau} x. E \rrbracket = \lambda_h h. \lambda_{\tau} x. \mathbf{push}_h h \circ \mathcal{H}c \llbracket E \rrbracket & \text{with } i \equiv s, e \text{ or } k \\
\mathcal{H}c \llbracket \mathbf{push}_i E \rrbracket = \mathbf{Store}[\mathcal{H}c \llbracket E \rrbracket] & \text{if } i \equiv s \text{ and } E : \mathbf{R}_s \sigma \\
= \lambda_h h. \mathbf{push}_i (\mathcal{H}c \llbracket E \rrbracket) \circ \mathbf{push}_h h & \text{otherwise } (i \equiv s, e \text{ or } k) \\
\mathcal{H}c \llbracket x \rrbracket = \mathbf{Call}[x] & \text{if } x : \mathbf{R}_s \tau \text{ bound by } \lambda_s x. \\
= x & \text{otherwise}
\end{array}$$

with $\mathbf{Store}[E] \equiv \lambda_h h. \mathbf{push}_h h \circ \mathbf{alloc} \circ \lambda_h h. \lambda_s a.$

$$\begin{array}{l}
\mathbf{push}_s E \circ \mathbf{push}_s a \circ \mathbf{push}_h h \circ \mathbf{write} \circ \lambda_h h. \mathbf{push}_s a \circ \mathbf{push}_h h \\
\mathbf{Call}[E] \equiv \lambda_h h. \mathbf{push}_s E \circ \mathbf{push}_h h \circ \mathbf{read} \circ \lambda_s y. \mathbf{push}_h h \circ y \\
\mathbf{alloc} = \lambda_h (heap, free). \mathbf{push}_s free \circ \mathbf{push}_h (heap, free+1) \\
\mathbf{write} = \lambda_h (heap, free). \lambda_s add. \lambda_s val. \mathbf{push}_h ((heap, \{add, val\}), free) \\
\mathbf{read} = \lambda_h ((heap, \{add_1, val\}), free). \lambda_s add_2. \\
\text{if } add_1 = add_2 \text{ then } \mathbf{push}_s val \text{ else } \mathbf{push}_h (heap, free) \circ \mathbf{push}_s add_2 \circ \mathbf{read}
\end{array}$$

Figure 17 Introducing a heap where closures are allocated and accessed ($\mathcal{H}c$)

We can apply the transformation $\mathcal{H}c$ to get new versions of the combinators introduced by the previous compilation steps. When a combinator neither create nor call a closure, the

transformation $\mathcal{H}c$ threads the heap without interaction. For example, for the combinator **dupl_e** introduced by the abstraction $\mathcal{A}g$, we get

$$\mathbf{dupl}_{eh} = \mathcal{H}c \llbracket \mathbf{dupl}_e \rrbracket = \lambda_h h. \lambda_e e. \mathbf{push}_e e \circ \mathbf{push}_e e \circ \mathbf{push}_h h$$

On the other hand, combinators such as **appclos** and **mkclos** create or call closures. So, their transformed definitions use **Call** and **Store**:

$$\mathbf{appclos}_h = \mathcal{H}c \llbracket \mathbf{appclos} \rrbracket = \mathcal{H}c \llbracket \lambda_s x.x \rrbracket = \lambda_h h. \lambda_s x. \mathbf{push}_h h \circ \mathbf{Call}[x]$$

$$\mathbf{mkclos}_h = \mathcal{H}c \llbracket \mathbf{mkclos} \rrbracket = \lambda_h h. \lambda_s x. \lambda_e e. \mathbf{push}_h h \circ \mathbf{Store}[\lambda_h h. \mathbf{push}_e e \circ \mathbf{push}_h h \circ x]$$

6.2 Updating

The transformation $\mathcal{H}c$ only makes memory management explicit. A heap stored closure is still reduced every time it is accessed. The call-by-need strategy updates the heap allocated closures with their normal forms.

The main choice is whether the update is performed by the caller (i.e. by the code from which the closure is accessed) or by the callee (i.e. by the code of the closure itself). The caller update scheme updates a closure every time it is accessed, when the callee-update scheme updates closures only the first time they are accessed: once in normal form, other accesses will not entail further (useless) updates. This last scheme is more efficient and is implemented by all the realistic, environment-based implementations. We model here only callee updates.

6.2.1 Callee update

In order to have self updating closures, the transformation $\mathcal{U}callee$ (Figure 18) changes the rule of $\mathcal{H}c$ for **push_s E**. It introduces a combinator **updt** which takes as its arguments the heap h , the address b of the result, and the address a of the closure to be updated. It returns the address b and a new heap where the cell a contains an indirection to b . The combinator **swap_{sh}** reorders the address x and the heap.

$$\mathcal{U}callee : \Lambda_k \rightarrow \Lambda_h \qquad \text{with } E : \mathbf{R}_s \sigma$$

$$\mathcal{U}callee \llbracket \mathbf{push}_s E \rrbracket = \mathbf{Store}[\mathbf{push}_s a \circ \mathbf{swap}_{sh} \circ \mathcal{U}callee \llbracket E \rrbracket \circ \mathbf{updt}]$$

with $\mathbf{swap}_{sh} = \lambda_s a. \lambda_h h. \mathbf{push}_s a \circ \mathbf{push}_h h$

and $\mathbf{updt} = \lambda_h h. \lambda_s b. \lambda_s a. \mathbf{push}_s (\lambda_h h. \mathbf{push}_s b \circ \mathbf{push}_h h) \circ \mathbf{push}_s a \circ \mathbf{push}_h h \circ \mathbf{write}$

$$\circ \lambda_h h. \mathbf{push}_s b \circ \mathbf{push}_h h$$

Figure 18 Callee closure update ($\mathcal{U}callee$)

A closure is allocated in the heap when it is created as in $\mathcal{H}c$, but its code is modified. The closure now stores its own address ($\mathbf{push}_s a$), and its evaluation is followed by \mathbf{updt} . Note that a is a variable bound in the context $\mathbf{Store}[]$ (see the definition of \mathbf{Store}) and denotes the address of a fresh allocated cell. Of course, when E is already (syntactically) in normal form the simple rule $\mathcal{U}callee \llbracket \mathbf{push}_s E \rrbracket = \mathbf{Store}[\mathcal{U}callee \llbracket E \rrbracket]$ suffices. Thus, a closure is updated at most once (i.e. after the first access) because the compiled code of its normal form ($\mathcal{H}c \llbracket \mathbf{push}_s N \rrbracket$) contains no \mathbf{updt} .

The callee update scheme can be used with $\mathcal{N}m$. However, as noted in Section 3.2.2, marks have to be inserted in expressions to suspend the reduction before performing an update. The rule for λ -abstractions becomes

$$\mathcal{N}m \llbracket \lambda x. E \rrbracket = \mathbf{grab}_s(\lambda_s x. \mathcal{N}m \llbracket E \rrbracket)$$

and $\mathcal{U}callee$ is specialized for the push-enter model as follows:

$$\begin{aligned} \mathcal{U}callee \llbracket \mathbf{push}_s E \rrbracket = & \\ & \mathbf{Store}[\mathbf{push}_s a \circ \mathbf{swap}_{sh} \circ \mathbf{push}_s \varepsilon \circ \mathbf{swap}_{sh} \circ \mathcal{U}callee \llbracket E \rrbracket \circ \mathbf{updt} \circ \mathbf{resume}_h] \\ & \text{with } \mathbf{resume}_h = \lambda_h h. \lambda_s x. \mathbf{push}_s h \circ \mathbf{grab}_h x \end{aligned}$$

An evaluation context is isolated by inserting a mark ε after the update address ($\mathbf{push}_s a$); and \mathbf{resume}_h resumes the reduction once the update has been performed. The combinator \mathbf{grab}_h is defined by $\mathcal{H}c \llbracket \mathbf{grab}_s \rrbracket$. Marks are used by Tim[20], Clean [46], the Krivine Machine [11] and the spineless-tagless G-machine [42]. The codes produced by $\mathcal{N}a$ and $\mathcal{N}m$ have the same update opportunities. As in call-by-name, the call-by-need version of $\mathcal{N}m$ may prevent from building unnecessary intermediate closures.

6.2.2 Updating and graph reduction

The previous transformations can be employed to transform the call-by-name graph reduction schemes into call-by-need. Here, we present two updating techniques (spineless and spine variations) that have been introduced for the G-machine.

The spineless G-machine [8] updates only selected application nodes. Unwinding application nodes entails stacking either their address (updatable) or only the argument address (non updatable). So, in general, the complete leftmost spine of the graph does not appear in the stack. The code must annotate updatable nodes and marks are necessary to dynamically detect when an update must be performed. Updatable nodes are distinguished using the combinator \mathbf{mkAppS}_s which has the same definition as \mathbf{mkApp}_s , and \mathbf{mkFun}_s must be redefined to detect marks:

$$\begin{aligned} \mathbf{mkAppS}_s &= \lambda_s x_1. \lambda_s x_2. \mathbf{push}_s (\mathbf{push}_s x_2 \circ x_1) \\ \mathbf{mkFun}_s &= \lambda_s f. \mathbf{push}_s (\mathbf{grab}_s(\lambda_s a. (\mathbf{push}_s a \circ f) \circ \mathbf{unwind}_s)) \end{aligned}$$

The transformation $\mathcal{U}callee$ for the push-enter model can be applied to the graph constructors. For \mathbf{mkAppS}_s we get

$$\mathcal{U}call_{ee} \llbracket \mathbf{mkAppS}_s \rrbracket = \lambda_h h. \lambda_s x_1. \lambda_s x_2. \mathbf{Store}[\mathbf{push}_s a \circ \mathbf{swap}_{sh} \circ \mathbf{push}_s \varepsilon \circ \mathbf{swap}_{sh} \circ \mathcal{U}call_{ee} \llbracket \mathbf{push}_s x_2 \circ x_1 \rrbracket \circ \mathbf{updt} \circ \mathbf{resume}_h]$$

As suggested in Section 3.3.2, the use of marks is not mandatory to express updating in the G-machine [27] where graph building and graph reduction are separate steps. Application nodes must stack their address as they are unwound, then updates can be systematically inserted between each graph building and reduction step. However, this naive scheme (that we call the spine variation) cannot be expressed using the previous transformations. Indeed, the canonical definition of \mathbf{mkApp}_s for $\mathcal{G}\mathcal{M}$ is

$$\mathbf{mkApp}_s = \lambda_s x_1. \lambda_s x_2. \mathbf{push}_s (\mathbf{push}_s x_2 \circ x_1) \quad \text{where } \mathbf{push}_s x_2 \circ x_1: \sigma_1 \rightarrow_s \sigma_2$$

Since $\mathcal{H}c$ shares only expressions of the form $\mathbf{push}_s E$ with $E: \mathbf{R}, \sigma$, application nodes will not be considered for updating with this definition of \mathbf{mkApp}_s . In order to model the G-machine scheme, a new transformation should be defined (see $\mathcal{U}spine$ in [18]).

The introduction of the threaded memory component in our functional intermediate code makes formal manipulations more complicated. For example, a property ensuring that the reduction of $\mathcal{H}c \llbracket E \rrbracket$ simulates the reduction of E , should use a decompilation transformation in order to replace the addresses in reduced expressions by their actual values which lie in the heap. This prevented us from finding a simple and convincing formulation of correctness properties for the transformations presented in this section.

7 CLASSICAL FUNCTIONAL IMPLEMENTATIONS

The description of the compilation process is now complete. A compiler can be described by a simple composition of transformations. Figure 19 states the main design choices structuring several classical implementations. There are cosmetic differences between our descriptions and the real implementations. Some descriptions of the literature leave the compilation of control transfers implicit (e.g. the Cam and Tim). Also, some extensions and optimizations are not described here.

Let us describe precisely our modeling of the categorical abstract machine and state the differences with the description in [10]. The Cam implements the left-to-right call-by-value strategy using the eval-apply model and has linked environments. In our framework, this is expressed as $\mathcal{CAM} = \mathcal{A}s \bullet \mathcal{V}a_L$. By simplifying this composition of transformations, we get:

$$\begin{aligned} \mathcal{CAM} \llbracket x_i \rrbracket \rho &= \mathbf{fst}^i \circ \mathbf{snd} \\ \mathcal{CAM} \llbracket \lambda x. E \rrbracket \rho &= \mathbf{push}_s (\mathbf{mkbnd} \circ (\mathcal{CAM} \llbracket E \rrbracket (\rho, x))) \circ \mathbf{mkclos} \\ \mathcal{CAM} \llbracket E_1 E_2 \rrbracket \rho &= \mathbf{dupl}_e \circ (\mathcal{CAM} \llbracket E_1 \rrbracket \rho) \circ \mathbf{swap}_{se} \circ (\mathcal{CAM} \llbracket E_2 \rrbracket \rho) \circ \mathbf{appclos}_L \\ \text{with } \mathbf{appclos}_L &= \lambda_s x. \lambda_s f. \mathbf{push}_s x \circ f \end{aligned}$$

To illustrate its output, let us consider the expression $E \equiv (\lambda x. x)((\lambda y. y)(\lambda z. z))$, then

$$\mathcal{CAM} \llbracket E \rrbracket = \mathbf{dupl}_e \circ \mathbf{push}_s C_1 \circ \mathbf{mkclos} \circ \mathbf{swap}_{se} \circ \mathbf{dupl}_e \circ \mathbf{push}_s C_1 \circ \mathbf{mkclos} \circ \mathbf{swap}_{se}$$

$$\circ \text{push}_s C_1 \circ \text{mkclos} \circ \text{appclos}_L \circ \text{appclos}_L$$

$$\text{with } C_1 \equiv \text{mkbind} \circ \text{snd}$$

The code is made of two linear code sequences, each of them composed of combinators which can be implemented by a few standard assembly instructions. The minor step consisting of naming code fragments has been left implicit. By instantiating the combinators, we get the rules of the machine. In the Cam, both components s and e are merged; the instantiation is therefore:

$$\circ = \lambda abc. a (b c) \quad \text{push}_s N = \text{push}_e N = \lambda c. \lambda z. c (z, N) \quad \lambda_s x. X = \lambda_e x. X = \lambda c. \lambda (z, x). X c z$$

The definitions of the (macro) combinators follow. For example:

$$\text{dupl}_e = \lambda_e e. \text{push}_e e \circ \text{push}_e e = \lambda c. \lambda (z, e). c ((z, e), e)$$

$$\text{mkbind} = \lambda_e e. \lambda_s x. \text{push}_e (e, x) = \lambda c. \lambda ((z, e), x). c (z, (e, x))$$

$$\text{snd} = \lambda_e (e, x). \text{push}_s x = \lambda c. \lambda (z, (e, x)). c (z, x)$$

If these combinators are considered as the basic instructions of an abstract machine, their definitions imply the following state transitions:

$$\text{dupl}_e \quad C (Z, E) \quad \rightarrow \quad C ((Z, E), E)$$

$$\text{mkbind} \quad C ((Z, E), X) \quad \rightarrow \quad C (Z, (E, X))$$

$$\text{snd} \quad C (Z, (E, X)) \quad \rightarrow \quad C (Z, X)$$

The **fst**, **snd**, **dupl_e** and **swap_{se}** combinators correspond to Cam's **Fst**, **Snd**, **Push** and **Swap**. The sequence **push_s** (E) \circ **mkclos** is equivalent to Cam's **Cur**(E). The only difference comes from the place of **mkbind** (at the beginning of each closure in our case). Shifting this combinator to the place where the closures are evaluated and merging it with **appclos_L**, we get $\lambda_s(x, e). \text{push}_e e \circ \text{mkbind} \circ x$, which is exactly Cam's sequence **Cons;App**.

Figure 19 gathers our modelings of 13 implementations of strict or lazy functional languages. It refers to a few transformations not described in this paper but which can be found in [17] and [18].

Let us quickly review the differences between Figure 19 and real implementations. The Clean implementation is based on graph rewriting, however the final code is similar to environment machines (for example, a closure is encoded by an n -ary node). Our replica is an environment machine that we believe is close. However, the numerous optimizations and especially the lack of clear description ([46] details only examples of final code) makes it difficult to precisely determine the compilation choices.

<i>Compiler</i>	<i>Transformations</i>					<i>Components</i>
	$\Lambda \rightarrow$	$\Lambda_s \rightarrow$	$\Lambda_e \rightarrow$	$\Lambda_k \rightarrow$	Λ_h	
<i>Cam</i>	$\mathcal{V}a_L$	$\mathcal{A}s$	Id	Id		$s \equiv e$
<i>Clean</i>	$\mathcal{N}ml$	$\mathcal{A}c1$	Id	$\mathcal{U}Callee$		$s \ e \ k \ h$
<i>G-machine</i>	$\mathcal{G}\mathcal{N}m$	$\mathcal{A}c1_{dsb}$	Id	$\mathcal{U}Spine$		$s \ e \ k \ h$
<i>Spineless G-machine</i>	$\mathcal{G}\mathcal{N}ml$	$\mathcal{A}c1_{dsb}$	Id	$\mathcal{U}Callee$		$s \ e \ k \ h$
<i>Spineless tagless G-machine</i>	$\mathcal{N}ml$	$\mathcal{A}c3$	Id	$\mathcal{U}Callee$		$(s \equiv k) \ e \ h$
<i>Mak</i> (cbn)	$\mathcal{N}ml$	$\mathcal{A}s$	Id	$\mathcal{U}Callee$		$s \equiv e \equiv k \ h$
<i>Mak_s</i> (cbv)	$\mathcal{V}m$	$\mathcal{A}s$	S	Id		$s \equiv e \equiv k$
<i>SECD</i>	$\mathcal{V}a$	$\mathcal{A}s$	S	Id		$s \ (e \equiv k)$
<i>SKI-machine</i>	$\mathcal{G}\mathcal{N}m$	SKI	Id	$\mathcal{U}Spine$		$s \ h$
<i>SML-NJ</i>	$\mathcal{V}a_f$	$\mathcal{A}c3$	S	Id		$s \ e \ (\text{registers})$
<i>Tabac</i> (cbv)	$\mathcal{V}a$	$\mathcal{A}c2_{dsb}$	S	Id		$(s \equiv e) \ k$
<i>Tabac</i> (cbn)	$\mathcal{N}a$	$\mathcal{A}c2_{dsb}$	S	$\mathcal{U}Hybrid$		$(s \equiv e) \ k \ h$
<i>Tim</i>	$\mathcal{N}ml$	$\mathcal{A}c1_m$	Id	$\mathcal{U}Callee$		$s \ e \ k \ h$

Figure 19 Several classical compilation schemes

The G-machine [27] and the spineless G-machine [8] perform only one test for all the arguments of the function (by comparing the arity of the function with the activation record size) whereas our **grab_s** combinator performs a test for every argument. So, an n-ary combinator **grabs_n** should be introduced.

The spineless tagless G-machine [42] uses also an n-ary version of **grab_s** and a local and a global environment. The abstraction with two environments ($\mathcal{A}c_3$ in our framework) is not directly compatible with **grab_s** and extra environment copies must be inserted. The simplest way to model faithfully the real machine would be to introduce an specialized abstraction algorithm.

The Grab instruction of the Krivine abstract machine (Mak) [11][32] is a combination of our **grab_s** (in fact, a recursive version ^①) and **mkbind** combinators.

The SECD machine [31] saves environments a bit later than in our scheme. Furthermore, the control stack and the environment stack are grouped into a component called

“dump”. The data stack is also (uselessly) saved in the dump. Actually, our replica is closer to the idealized version derived in [24].

The SKI-machine [53] reduces a graph made of combinators **S**, **K**, **I** and application nodes. The graph representing the source expression is totally built at compile time. The machine is made of a recursive interpreter and a data stack to store the unwound spine. Our modeling is close to the somewhat informal description of the SKI-machine in [53].

The SML-NJ compiler [2] uses only the heap which is represented in our framework by a unique environment e . It also includes registers and numerous optimizations not described here.

The Tabac compiler is a by-product of our work in [22] and has greatly inspired this study. It implements strict or non-strict languages by program transformations. Tabac integrated many optimizations that we have not described here.

Our call-by-name Tim description is accurate according to [20]. The environment copying included in the transformation \mathcal{Acl} have the same effect as the preliminary lambda-lifting phase of Tim. A n-ary **grab** _{s} should be added to our call-by-need version.

8 EXTENSIONS AND APPLICATIONS

Our framework is powerful enough to handle realistic languages and to model optimizing compilers or hybrid implementations. We illustrate each point in turn. We first present the integration of constants, primitive operators and data structures, then we take an example of how to express a classical global optimization and finally we describe a hybrid transformation.

8.1 Constants, Primitive Operators and Data Structures

We have only considered pure λ -expressions because most fundamental choices can be described through this simple language. Realistic implementations also deal with constants, primitive operators and data structures which are easily taken into account in our framework.

Concerning basic constants, one question is whether results of basic type are returned in s or another component (**push** _{b} , λ_b) is introduced. The latter has the advantage of marking a difference between pointers and values which can be exploited by the garbage collector. But in this case, precise type information must also be available at compile time to transform variables and λ -abstractions correctly. In a polymorphic setting, this information is not available in general (a variable x of polymorphic type α can be bound to anything) so constants, functions and data structures must be stored in s . The fix-point operator, the conditional and primitive operators acting on basic values are introduced in our language in a straightforward way. The compilation of control using the eval-apply model for these constructs is described in Figure 20.

A naive compilation of β -reduction for letrec expressions yields a code building a closure at each recursive call. Two optimizations exist. The first one consists in building a circu-

lar environment or graph. A second optimization for environment based machines is to implement recursive calls to statically known functions by a jump to their address[Ⓓ].

$$\begin{aligned}
\mathcal{V}[\text{letrec } f = E] &= \mathbf{push}_s (\lambda_s f. \mathcal{V}[\![E]\!]) \circ \mathbf{Y}_s \quad \text{with } \mathbf{push}_s F \circ \mathbf{Y}_s \Rightarrow \mathbf{push}_s(\mathbf{push}_s F \circ \mathbf{Y}_s) \circ F \\
\mathcal{V}[n] &= \mathbf{push}_s n \\
\mathcal{V}[\text{if } E_1 \text{ then } E_2 \text{ else } E_3] &= \mathcal{V}[\![E_1]\!] \circ \mathbf{cond}_s (\mathcal{V}[\![E_2]\!], \mathcal{V}[\![E_3]\!]) \\
\text{with } \quad \mathbf{push}_s \text{True} \circ \mathbf{cond}_s (E, F) &\Rightarrow E \quad \text{and} \quad \mathbf{push}_s \text{False} \circ \mathbf{cond}_s (E, F) \Rightarrow F \\
\mathcal{V}[\![E_1 + E_2]\!] &= \mathcal{V}[\![E_2]\!] \circ \mathcal{V}[\![E_1]\!] \circ \mathbf{plus}_s \quad \text{with } \mathbf{push}_s n_2 \circ \mathbf{push}_s n_1 \circ \mathbf{plus}_s \Rightarrow \mathbf{push}_s n_1 + n_2 \\
\mathcal{V}[\![\text{head}]\!] &= \mathbf{head}_s \quad \text{with } \mathbf{head}_s = \lambda_s (\text{tag}, h, t). \mathbf{push}_s h \\
\mathcal{V}[\![\text{cons } E_1 E_2]\!] &= \mathcal{V}[\![E_2]\!] \circ \mathcal{V}[\![E_1]\!] \circ \mathbf{cons}_s \quad \text{with } \mathbf{cons}_s = \lambda_s h. \lambda_s t. \mathbf{push}_s (\text{tag}, h, t)
\end{aligned}$$

Figure 20 An extension with constants, primitive operators and lists

As far as data structures are concerned, we can choose to represent them using tags or higher-order functions [20]. Figure 20 describes a possible extension using the data stack to store constants and tagged cells of lists. It just indicates one simple way to accommodate data structures in our framework. The efficient implementation of data structures brings a whole new collection of choices (see e.g. [42]) and optimizations (see e.g. [23] [51]). A thorough description of this subject is beyond the scope of this paper.

Until now, we considered only pure λ -expressions and the typing of the source language was not an issue. When constants and data structures are taken into account two cases arise depending on the typing policy of the source language. If the source language is statically typed then the code produced by our transformation does not need to be modified (however, supporting polymorphism efficiently involves new and specific optimizations such as unboxing of floats and tuples [33]). For dynamically typed languages, functions, constants and data structures must carry a type information which will be checked by combinators or primitive operators at run time.

8.2 Optimizations

Let us take the example of the optimization brought by strictness analysis in call-by-need implementations. It changes the evaluation order and, more interestingly, avoids some thunks using unboxing [9]. If we assume that a strictness analysis has annotated the code $\underline{E}_1 E_2$ if E_1 denotes a strict function and \underline{x} if the variable is defined by a strict λ -abstraction then $\mathcal{N}a$ can be optimized as follows

$$\mathcal{N}a \llbracket \underline{x} \rrbracket = \mathbf{push}_s x \qquad \mathcal{N}a \llbracket \underline{E}_1 E_2 \rrbracket = \mathcal{N}a \llbracket \underline{E}_2 \rrbracket \circ \mathcal{N}a \llbracket \underline{E}_1 \rrbracket \circ \mathbf{app}$$

Underlined variables are known to be already evaluated; they are represented as unboxed values. For example, without any strictness information, the expression

$$(\lambda x. x+1) 2$$

is compiled into $\mathbf{push}_s (\mathbf{push}_s 2) \circ (\lambda_s x. x \circ \mathbf{push}_s 1 \circ \mathbf{plus}_s)$.

The code $\mathbf{push}_s 2$ will be represented as a closure and evaluated by the call x ; it is the boxed representation of 2. With strictness annotations we have

$$\mathbf{push}_s 2 \circ (\lambda_s x. \mathbf{push}_s x \circ \mathbf{push}_s 1 \circ \mathbf{plus}_s)$$

and the evaluation is the same as with call-by-value (no closure is built). Actually, more general forms of unboxing (as in [33] or [44]) and optimizations (e.g. let-floating [45]) could be expressed as well.

8.3 Hybrid Implementations

The study of the different options showed that there is no universal best choice. It is natural to strive to get the best of each world. Our framework makes intricate hybridizations and related correctness proofs possible. It is for example possible to mix the eval-apply and push-enter models and to design a $\mathcal{V}a\text{-}\mathcal{V}m$ hybrid transformation ([17]). Here, we describe how to mix shared and copied environments. We suppose that a static analysis has produced an annotated code indicating the chosen mode for each subexpression.

One solution could be to use coercion functions to fit the environment into the chosen structure (list or vector). Instead, we describe a more sophisticated solution (Figure 21) which allows lists and vectors to coexist within environments (as in [50]). Motivations for this feature may be to optimize run time using vectors (resp. links) when access (resp. closure building) cost is predominant or to optimize space usage by using a copy scheme (e.g. vectors) when it eliminates a space leak which would be introduced by linking environments.

$$\begin{aligned} \mathcal{M}ix\mathcal{A} \llbracket \lambda_s x. E^{\theta, \oplus} \rrbracket \rho &= \mathbf{Mix} \rho \theta \circ \mathbf{mkbind}^{\oplus} \circ \mathcal{M}ix\mathcal{A} \llbracket E \rrbracket (\theta \oplus x) \\ \mathcal{M}ix\mathcal{A} \llbracket x_i \rrbracket (\dots(\rho, \rho_i), \dots, \rho_0) &= \mathbf{access}_i^l \circ \mathcal{M}ix\mathcal{A} \llbracket x_i \rrbracket \rho_i && \text{with } x_i \text{ in } \rho_i \\ \mathcal{M}ix\mathcal{A} \llbracket x_i \rrbracket [\rho; \rho_i; \dots; \rho_0] &= \mathbf{access}_i^v \circ \mathcal{M}ix\mathcal{A} \llbracket x_i \rrbracket \rho_i && \text{with } x_i \text{ in } \rho_i \\ \mathcal{M}ix\mathcal{A} \llbracket x_i \rrbracket (\dots(\rho, x_i), \dots, x_0) &= \mathbf{access}_i^l \circ \mathbf{appclos} \\ \mathcal{M}ix\mathcal{A} \llbracket x_i \rrbracket [\rho; x_i; \dots; x_0] &= \mathbf{access}_i^v \circ \mathbf{appclos} \\ \text{with } \mathbf{access}_i^l \text{ (resp. } \mathbf{access}_i^v \text{) is the } \mathbf{access}_i \text{ version which accesses a list (resp. a vector)} \end{aligned}$$

Figure 21 Hybrid Abstraction (extract)

Each λ -abstraction is annotated by a new mixed environment structure θ and \oplus ($\in \{v, l\}$) which indicates how to bind the current value (as a vector “ v ” or as a link “ l ”). Mixed structures are built by \mathbf{mkbind}^v , \mathbf{mkbind}^l and the macro-combinator \mathbf{Mix} which copies and restructures the environment ρ according to the annotation θ (Figure 21). Paths to values are now expressed by sequences of \mathbf{access}_i^l and \mathbf{access}_i^v . The abstraction algorithm distinguishes vectors from lists in the compile time environment using constructors “ $:$ ” and “ $,$ ”.

9 RELATED WORK

We review in this section the different formalisms used in the description of functional implementations: the λ -calculus, λ -calculi with explicit substitutions, combinators, monads. We also present papers comparing specific implementations and the related area of semantic-directed compiler derivation.

Our approach and this paper stem from our previous work on compilation of functional languages by program transformation [22]. Our goal then was to show that the whole implementation process could be described in the functional framework. The two main steps were the compilation of control using a CPS conversion and the compilation of the β -reduction using indexed combinators that could be seen as basic instructions on a stack. We remained throughout within the λ -calculus and did not have to introduce an ad-hoc abstract machine. We described only one particular implementation; our main motivation was to make correctness proofs of realistic implementations simpler not to describe and compare various implementation techniques. The SML-NJ compiler has also been described using program transformations including CPS and closure conversions [2]. Other compilers use the CPS transformation to encode the reduction strategy within the λ -calculus [30][52]. Encoding implementation issues within the λ -calculus leads to complex expressions (e.g. sequencing is coded as a composition of continuations). The constructors \mathbf{push}_p , \circ and λ_i make our framework more abstract and simplify the expressions. The instantiation of these constructors as λ -expressions provides an interesting new implementation step (Section 2.5): the choice of the number and the representation of the components of the underlying abstract machine are kept apart. Within the λ -calculus, one has to choose before describing an implementation whether, for example, data and environments are stored in two separate components or in a single one.

The de Bruijn λ -calculus [14], which uses indices instead of variables, has been used as an intermediate language by several abstract machines. As we saw in Section 4.1.2, a de Bruijn index can be seen as the address of a value in the run-time environment. A collection of formalisms, the λ -calculi with explicit substitutions, emphasize also the environment management and can be seen as calculi of closures [1]. These calculi help formal reasoning on substitution and make some implementation details explicit. However, important implementation choices such as the representation of the environments (lists or vectors) are, in general, not tackled in these formalisms. Hardin & *al.* [25] introduce $\lambda\sigma_w$, a weak λ -calculus with explicit substitutions, which can serve as the output language of functional compilers. They describe several abstract machines in this framework. However, their goal is to exhibit the common points of implementations not to model precisely existing implementations. Another variant, $\lambda\sigma_w^a$ [7], can describe sharing and eases the proofs concerning memory

management. The $\lambda\sigma_w^a$ -expressions stay at a higher level than real machine code since, for example, sharing is modeled by formal labels and parallel reductions.

A closely related framework used as intermediate language is combinatory logic [13]. Combinators have been used to encode the compilation of the β -reduction. Some compilation issues, such as the representation of environments, are usually not dealt with. Different set of combinators, such as $\{\mathbf{S}, \mathbf{K}, \mathbf{I}, \mathbf{B}, \mathbf{C}\}$ [53], have been used to define abstraction algorithms for graph reducers [28][36]. The categorical combinators [12] have been used in environment machines such as the Cam [10] and the Krivine machine [4].

Arising from different roots, our first intermediate language Λ_s is surprisingly close to Moggi's computational metalanguage [40]. In particular, we may interpret the monadic construct $[E]$ as $\mathbf{push}_s E$ and $(\mathbf{let } x \leftarrow E_1 \mathbf{ in } E_2)$ as $E_1 \circ \lambda_x.x.E_2$ and get back the monadic laws ($\mathbf{let}.\beta$), ($\mathbf{let}.\eta$) and (\mathbf{ass}). The monadic framework is more abstract. For example, one can write monadic expressions such as

$$\mathbf{let } _ \leftarrow \mathbf{writeStack}(X) \mathbf{ in } (\mathbf{let } e \leftarrow \mathbf{readEnv}() \mathbf{ in } E)$$

whereas, in our formalism, we need to reorder data and environment with a **swap** combinator:

$$\mathbf{push}_s X \circ \mathbf{swap}_{se} \circ \lambda_e.e.E$$

These administrative combinators allow us to merge several components in the instantiation step. The abstract features of monads can be an hindrance to express low level implementation details and to get closer to a machine code. For example, the monadic call-by-value CPS expression $(\mathbf{let } a \leftarrow A \mathbf{ in } (\mathbf{let } f \leftarrow F \mathbf{ in } [f a]))$ evaluates the argument A , the function F and returns the application $(f a)$, but does not state if the application is reduced before it is returned. In Λ_s , we disallow unrestricted applications and make the previous reduction explicit with an **app** combinator. A key feature of our approach is to describe and structure the compilation process as a composition of individualized transformations. The monadic framework does not appear to be well suited to this purpose since monads are notoriously difficult to compose. Liang & *al.* [35] needs complex parametrized monads to describe and compose different compilation steps. The difficulties to compose monads and to represent low level details are serious drawbacks with respect to our goals. Overall, the monadic framework is a general tool to structure functional programs [55] whereas our small framework has been tailor-made to describe implementations.

Besides benchmarks, few functional language implementations have been compared. Some particular compilation steps have been studied. For example, [28] compare different abstraction algorithms and [26] expresses CPS transformations in the monadic framework. A few works explore the relationship between two abstract machines such as CMC and Tim [37] and Tim and the G-machine [43]. Their goal is to show the equivalence between seemingly very different implementations. CMC and Tim are compared by defining transformations between the states of the two machines. The comparison of Tim and the G-machine is more informal but highlights the relationship between an environment machine and a graph reducer. Also, let us mention Asperti [4] who provides a categorical understanding of the Krivine machine and an extended Cam and Crégut [11] who has studied the relationship be-

tween the Tim and the Krivine machine. All these implementation comparisons focus on particular compilation steps or machines but do not define a global approach to compare implementations.

Related work also includes the derivation of abstract machines from denotational or operational semantics. Starting from a denotational semantics with continuations, Wand [56] compiles the β -reduction using combinators and linearizes expressions in sequences of abstract code. The semantics of the program is translated into a sequence representing the code and a machine to execute it. In our approach, semantics or machines do not appear explicitly. Hannan [24] and Sestoft [49] start from a “big step” (natural) operational semantics, incrementally suppress ambiguities (e.g. impose a left-to-right reduction order) and refine complex operation (e.g. β -reduction), until they get a “small step” (structural) operational semantics. Some of the refinement steps have to deal with operations specific to their framework (e.g. suppressing unification). Meijer [38] uses program algebra to calculate some simple compilers from a denotational semantics via a series of refinements. All these derivation techniques aim at providing a methodology to formally develop implementations from semantics. Their focus is on the refinement process and correction issues and, usually, they describe the derivation of a single implementation. Not surprisingly, the derived compilers do not model precisely existing implementations. They are best described as idealized than sophisticated or optimized implementations. Comparisons of implementation choices seem harder with a description based on semantics refinement than with a description by program transformations. Also, some choices seem difficult to naturally obtain by derivation (e.g. the push-enter model for call-by-value). On the other hand, these semantics based methodologies can potentially be applied to any language that can be described in their semantics framework.

10 CONCLUSION

Let us review the implementation choices encountered in our study. The most significant choice for the compilation of control is using the eval-apply model ($\mathcal{V}a, \mathcal{N}a$) or the push-enter model ($\mathcal{V}m, \mathcal{N}m$). There are other minor options such as stackless variations ($\mathcal{V}a_p, \mathcal{N}a_p$) or right-to-left vs. left-to-right call-by-value. We have shown that the transformations employed by graph reducers can be seen as interpretative versions of the environment-based transformations. For the compilation of β -reduction, the main choice is using environment-less (e.g. SKI) abstraction algorithms, list-like (shared) environments ($\mathcal{A}s$) or vector-like (copied) environments ($\mathcal{A}c$). For the latter choice, there are several transformations depending on the way environments are copied ($\mathcal{A}c1, \mathcal{A}c2, \mathcal{A}c3$). Actually, a complete family of generic transformations modeling different managements of the environment stack can be derived. For control transfers, one can introduce a return address stack or use CPS conversion. Self updatable closures (i.e. callee update) is the standard way to implement updating but graph reduction brings other options.

Our approach focuses on (but is not restricted to) the description and comparison of fundamental options. The transformations are designed to model a precise compilation step; they are generic with respect to the other steps. It is then not surprising that, often, simple compositions of transformations do not model accurately real implementations whose de-

sign is more ad-hoc than generic. In most cases, the differences are nevertheless superficial and it is sufficient to specialize the transformations to obtain existing implementations.

The use of program transformations appears to be well suited to precisely and completely model the compilation process. Many standard optimizations (uncurrying, unboxing, hoisting, peephole optimizations) can be expressed as program transformations as well. This unified framework simplifies correctness proofs. For example, we do not introduce explicitly an abstract machine and therefore we do not have to prove that its operational definition is coherent with the semantics of the language (as in [47] and [34]). Program transformations makes it possible to reason about the efficiency of the produced code as well as about the complexity of transformations themselves. Actually, these advantages appear clearly before the last compilation step. The introduction of a threaded state seriously complicates program manipulations and correctness proofs. This is not surprising because our final code is similar to a real assembly code.

Our main goal was to structure and clarify the design space of functional language implementations. The exploration is still far from complete. There are still many avenues for further research:

- It would be interesting to give a concrete form to our framework by implementing all the transformations presented. This compiler construction workbench would make it possible to implement a wide variety of implementations just by composing transformations. This would be useful to try completely new associations of compilation choices and to assess the implementations and optimizations in practice.
- A last step towards high quality machine code would be the modeling of register allocation. This could be done via the introduction of another component: a vector of registers.
- A systematic description of standard optimizations and program transformations should be undertaken. A benefit would be to clarify the impact of a program transformation depending on the implementation choices. Let us consider, for example, λ -lifting, a controversial transformation [27][39]. Intuitively, λ -lifting can be beneficial for implementations using linked environments. Indeed, in this case, its effect is to shorten accesses to variables by performing copies. Whether the gain is worth the cost depends on how many times a variable is accessed. We believe that this question could be studied and settled in our framework. Also, proving the correctness of optimizations based on static analyses is a difficult (and largely neglected) problem [9]. Expressing these optimizations as program transformations in our unified framework should make this task easier.
- Another research direction is the design of hybrid transformations (mixing several compilation schemes). We hinted at a solution to mix copied and linked environments in Section 8.3 and a solution to mix the eval-apply and the push-enter model in [18]. Others hybrid transformations as well as the analyses needed to make these transformations worthwhile have yet to be devised. Without the help of a formal framework, such transformations would be quite difficult to design and prove correct. The description of previously unknown compositions of transformations, the mechanical derivation of new abstraction algorithms and hybrid transformations all indicate that our approach can also suggest new implementation techniques.

- Many interesting formal comparisons of transformations remain to be done. At the moment, we have just compared a few couples of transformations ($\mathcal{V}a$ and $\mathcal{V}m$, $\mathcal{N}a$ and $\mathcal{N}m$ [18], $\mathcal{A}s$ and $\mathcal{A}cI$). It might be the case that a specific choice for a compilation step designates a best candidate for the compilation of another step. This could be established by comparing compositions of transformations (e.g. $\mathcal{A}s \bullet \mathcal{V}a$ and $\mathcal{A}cI \bullet \mathcal{V}a$).

We believe that the accomplished work already shows that our framework is expressive and powerful enough to tackle these problems.

ACKNOWLEDGMENTS

We are grateful to Luke Hornof and Daniel Le Métayer for their comments on several versions of this paper. Thanks are also due to the referees who provided helpful suggestions.

REFERENCES

- [1] M. Abadi, L. Cardelli, P. L. Curien and J. J. Levy. Explicit Substitutions. In Proc of *POPL'1990*, pp. 31-46, 1990.
- [2] A. W. Appel. *Compiling with Continuations*. Cambridge University Press. 1992.
- [3] G. Argo. Improving the three instruction machine. In *Proc. of FPCA'89*, pp. 100-115, 1989.
- [4] A. Asperti. A categorical understanding of environment machines. *Journal of Functional Programming*, 2(1), pp. 23-59, 1992.
- [5] H.P. Barendregt. *The lambda calculus. Its syntax and semantics*. North-Holland, 1981.
- [6] H.P. Barendregt and K. Hemerik. Types in lambda calculi and programming languages. In *3rd European Symp. on Programming*, LNCS Vol. 432, pp. 1-35, 1990.
- [7] Z. Benaissa, P. Lescanne and K. H. Rose: Modeling Sharing and Recursion for Weak Reduction Strategies Using Explicit Substitution. In Proc of PLILP'1996, pp. 393-407, 1996.
- [8] G. Burn, S.L. Peyton Jones and J.D. Robson. The spineless G-machine. In *Proc. of LFP'88*, pp. 244-258, 1988.
- [9] G. Burn and D. Le Métayer. Proving the correctness of compiler optimisations based on a global analysis. *Journal of Functional Programming*, 6(1), pp. 75-110, 1996.
- [10] G. Cousineau, P.-L. Curien and M. Mauny, The categorical abstract machine. *Science of Computer Programming*, 8(2), pp. 173-202, 1987.
- [11] P. Crégut. *Machines à environnement pour la réduction symbolique et l'évaluation partielle*. Thèse de l'université de Paris VII, 1991.
- [12] P. L. Curien. *Categorical Combinators, Sequential Algorithms and Functional Programming*. Birkhauser, 1993.
- [13] H.B. Curry and R. Feys. *Combinatory Logic*, Vol. 1, North-Holland. 1958.
- [14] N. G. De Bruijn. λ -calculus notation with nameless dummies: a tool for automatic formula manipulation, with application to Church Rosser theorem. In *Indagationes mathematicae*, 34, pp. 381-392, 1972.
- [15] O. Danvy. Back to direct style. In *Proc. of ESOP'92*, LNCS Vol. 582, pp. 130-150, 1992.
- [16] R. Douence and P. Fradet. Towards a taxonomy of functional language implementations. In *Proc of PLILP'95*, LNCS 982, pp. 27-44, 1995.

- [17] R. Douence and P. Fradet. A taxonomy of functional language implementations. Part I: call-by-value. *INRIA research report 2783*, Jan. 1996.
- [18] R. Douence and P. Fradet. A taxonomy of functional language implementations. Part II: call-by-name, call-by-need, and graph reduction. *INRIA research report 3050*, Nov. 1996.
- [19] R. Douence. *Décrire et comparer les mises en œuvre de langages fonctionnels*. PhD Thesis, University of Rennes I, 1996.
- [20] J. Fairbairn and S. Wray. Tim: a simple, lazy abstract machine to execute supercombinators. In *Proc of FPCA'87*, LNCS 274, pp. 34-45, 1987.
- [21] M. J. Fischer. Lambda-calculus schemata. In *Proc. of the ACM Conf. on Proving Properties about Programs*, Sigplan Notices, Vol. 7(1), pp. 104-109, 1972.
- [22] P. Fradet and D. Le Métayer. Compilation of functional languages by program transformation. *ACM Trans. on Prog. Lang. and Sys.*, 13(1), pp. 21-51, 1991.
- [23] C. Hall. Using Hindley-Milner type inference to optimise list representation. In *Proc. of LFP'94*, pp. 162-172, 1994.
- [24] J. Hannan. From operational semantics to abstract machines. *Math. Struct. in Comp. Sci.*, 2(4), pp. 415-459, 1992.
- [25] T. Hardin, L. Maranget and B. Pagano. Functional back-ends within the lambda-sigma calculus. In *Proc. of ICFP'1996*, pp. 25-33, 1996.
- [26] J. Hatcliff and O. Danvy. A generic account of continuation-passing styles. In *Proc. of POPL'94*, pp. 458-471, 1994.
- [27] T. Johnsson. *Compiling Lazy Functional Languages*. PhD Thesis, Chalmers University, 1987.
- [28] M. S. Joy, V. J. Rayward-Smith and F. W. Burton. Efficient combinator code. *Computer Languages*, 10(3), 1985.
- [29] J.W. Klop. Term rewriting systems. In *Handbook of Logic in Computer Science*. Vol. 2, pp. 2-108, Oxford University Press, 1992.
- [30] D. Kranz, R. Kesley, J. Rees, P. Hudak, J. Philbin, and N. Adams. ORBIT: An optimizing compiler for Scheme. *SIGPLAN Notices*, 21(7), pp. 219-233, 1986.
- [31] P. J. Landin. The mechanical evaluation of expressions. *The Computer Journal*, 6(4), pp. 308-320, 1964.
- [32] X. Leroy. The Zinc experiment: an economical implementation of the ML language. *INRIA Technical Report 117*, 1990.
- [33] X. Leroy. Unboxed objects and polymorphic typing. In *ACM Symp. on Princ. of Prog. Lang.*, pp. 177-188, 1992.
- [34] D. Lester, *Combinator Graph Reduction: a congruence and its application*. PhD thesis, Oxford University, 1989.
- [35] S. Liang, P. Hudak, M. P. Jones. Monad transformers and modular interpreters. In *Proc. of POPL'95*, pp. 333-343, 1995.
- [36] R. D. Lins. Categorical multi-combinators. In *Proc. of FPCA'87*, LNCS 274, pp. 60-79, 1987.
- [37] R. Lins, S. Thompson and S.L. Peyton Jones. On the equivalence between CMC and TIM. *Journal of Functional Programming*, 4(1), pp. 47-63, 1992.
- [38] E. Meijer, *Calculating Compilers*. PhD thesis, Nijmegen University, 1992.
- [39] E. Meijer and R. Paterson. Down with lambda lifting. *Unpublished paper*. copies available at: erik@cs.kun.nl, 1991.
- [40] E. Moggi. Notions of computation and monads. *Inf. and Comp.*, 93:55-92, 1991.

- [41] S.L. Peyton-Jones. *The implementation of functional programming languages*. Prentice Hall, 1987.
- [42] S.L. Peyton Jones. Implementing lazy functional languages on stock hardware: the spineless tagless G-machine. *Journal of Func. Prog.*, 2(2):127-202, 1992.
- [43] S. L. Peyton Jones and D. Lester. *Implementing functional languages, a tutorial*. Prentice Hall, 1992.
- [44] S. L. Peyton Jones and J. Launchbury. Unboxed values as first class citizens in a non-strict functional language. In *Proc. of FPCA'91*, LNCS 523, pp. 636-666, 1991.
- [45] S. L. Peyton Jones, W. Partain and A. Santos. Let-floating: moving bindings to give faster programs. In *Proc. of ICFP'96*, pp. 1-12, 1996.
- [46] R. Plasmeijer and M. van Eekelen. *Functional Programming and Parallel Graph Rewriting*. Addison-Wesley, 1993.
- [47] G. D. Plotkin: Call-by-Name, Call-by-Value and the lambda-Calculus. In *TCS 1*(2), pp. 125-159, 1975.
- [48] D. Schmidt. *Denotational Semantics, a Methodology for Language Development*. W.C. Brown Publishers, 1986.
- [49] P. Sestoft. Deriving a lazy abstract machine. *Technical Report 1994-146*, Technical University of Denmark, 1994.
- [50] Z. Shao and A. Appel. Space-efficient closure representations. In *Proc. of LFP'94*, pp. 150-161, 1994.
- [51] Z. Shao, J. Reppy and A. Appel. Unrolling lists. In *Proc. of LFP'94*, pp. 185-195, 1994.
- [52] G.L. Steele. Rabbit: a compiler for scheme. AI-TR-474, MIT, Cambridge, Mass., 1978.
- [53] D.A. Turner. A new implementation technique for applicative languages. *Soft. Pract. and Exper.*, 9, pp. 31-49, 1979.
- [54] D.A. Turner. Another algorithm for bracket abstraction. *J. of Symbolic Logic*, 44, pp. 267-270, 1979.
- [55] P. Wadler. The essence of functional programming. In *Proc. of POPL'92*, pp. 1-14, 1992.
- [56] M. Wand. Deriving target code as a representation of continuation semantics. *ACM Trans. on Prog. Lang. and Sys.*, 4(3), pp. 496-517, 1982.

APPENDIX

a. The strong confluence of the β_j -reduction is evident. The important point to note is that different redexes are always disjoint. Therefore, an expression E with two redexes R_1, R_2 can always be written as $C[R_1][R_2]$ ($C[\]$ being a context) and two different reductions

$$E \twoheadrightarrow F_1 \text{ and } E \twoheadrightarrow F_2$$

can be seen as $C[R_1][R_2] \twoheadrightarrow C[N_1][R_2]$ and $C[R_1][R_2] \twoheadrightarrow C[R_1][N_2]$

with $E \equiv C[R_1][R_2]$, N_1 and N_2 the reduced redexes (i.e. $F_1 \equiv C[N_1][R_2]$ and $F_2 \equiv C[R_1][N_2]$). Then clearly, the expression $G \equiv C[N_1][N_2]$ is such that $F_1 \twoheadrightarrow G$ and $F_2 \twoheadrightarrow G$

b. Proof of Property 2 and other typing issues.

For simplicity reasons, we implicitly assume that the source language can be typed using a standard type system. Let us note however that we could allow reflexive types (e.g. using a type system similar to $\lambda\mu$ -Curry [6]) to type any source expression and its compiled version. For example, the expression $\lambda x.x x$ would have type $\mu\alpha.\alpha \rightarrow \beta$ whereas its compiled form using, for example, $\mathcal{V}a$ (Section 3.1) is $\mathbf{push}_s(\lambda_s x.\mathbf{push}_s x \circ x)$ and would have type $R_s(\mu\alpha.\alpha \rightarrow_s R_s\beta)$. Typing in Λ_i does not impose any restrictions on source λ -expressions. The restrictions enforced by the type system are on how results and functions are combined in Λ_i .

In order to prove Property 2, we must first show a subject reduction property

Property 10 *If $E \xrightarrow{*} F$ then $\Gamma \vdash E : \tau \Rightarrow \Gamma \vdash F : \tau$*

Proof. It is clearly sufficient to show the property for one reduction step. The proof for the inductive rules such as $E \twoheadrightarrow N \Rightarrow E \circ F \twoheadrightarrow N \circ F$ is obvious. The interesting rule is the β_j -reduction and the proof boils down to the proof of $\Gamma \vdash F : \sigma$ and $\Gamma \cup \{x:\sigma\} \vdash E : \tau \Rightarrow \Gamma \vdash E[F/x] : \tau$. This is shown by structural induction.

- $E \equiv x$ then $\sigma \equiv \tau$ and $x[F/x] \equiv F$ so $\Gamma \vdash F : \sigma \Rightarrow \Gamma \vdash E[F/x](\equiv F) : \tau (\equiv \sigma)$
- $x \notin E$ (i.e. $E \equiv y \neq x$ or $E \equiv \lambda_z x.E'$) then $\Gamma \cup \{x:\sigma\} \vdash E : \tau \Rightarrow \Gamma \vdash E[F/x] (\equiv E) : \tau$
- $E \equiv \lambda_z z.E'$ ($z \neq x$) then

$$\Gamma \cup \{x:\sigma\} \vdash \lambda_z z.E' : \tau (\equiv \tau_1 \rightarrow_i \tau_2) \Leftrightarrow \Gamma \cup \{x:\sigma\} \cup \{z:\tau_1\} \vdash E' : \tau_2$$

since $z \neq x$, $\Gamma \cup \{z:\tau_1\} \cup \{x:\sigma\} \vdash E' : \tau_2$ and since the definition of substitution enforces z not to occur free in F (by variable renaming or convention) $\Gamma \vdash F : \sigma \Rightarrow \Gamma \cup \{z:\tau_1\} \vdash F : \sigma$. So, by induction hypothesis, $\Gamma \cup \{z:\tau_1\} \vdash E'[F/x] : \tau_2$ which implies $\Gamma \vdash \lambda_z z.E'[F/x] : \tau_1 \rightarrow_i \tau_2$.

- $E \equiv E_1 \circ E_2$ then

$\Gamma \cup \{x:\sigma\} \vdash E_1 \circ E_2 : \tau \Rightarrow \Gamma \cup \{x:\sigma\} \vdash E_1 : R_i \tau_1$ and $\Gamma \cup \{x:\sigma\} \vdash E_2 : \tau_1 \rightarrow_i \tau$. Using $\Gamma \vdash F : \sigma$ and the induction hypothesis we get $\Gamma \cup \{x:\sigma\} \vdash E_1[F/x] : R_i \tau_1$ and $\Gamma \cup \{x:\sigma\} \vdash E_2[F/x] : \tau_1 \rightarrow_i \tau$ so $\Gamma \cup \{x:\sigma\} \vdash (E_1 \circ E_2)[F/x] : \tau$

- $E \equiv \mathbf{push}_i E'$ then

$\Gamma \cup \{x:\sigma\} \vdash \mathbf{push}_i E' : \tau (\equiv R_i \tau_1) \Rightarrow \Gamma \cup \{x:\sigma\} \vdash E' : \tau_1 \Rightarrow \Gamma \cup \{x:\sigma\} \vdash E'[F/x] : \tau_1$ (by induction hypothesis) $\Rightarrow \Gamma \cup \{x:\sigma\} \vdash \mathbf{push}_i E'[F/x] : R_i \tau_1 (\equiv \tau)$ \square

We also have the following property:

Property 11 *A closed expression $E:\tau$ is either canonical (i.e. $E \equiv \mathbf{push}_i V$ or $\lambda_x.F$) or reducible.*

Proof. Structural induction. We have to show that an expression $E_1^{R_i \sigma} \circ E_2^{\sigma \rightarrow_i \tau}$ is reducible. If $E_1 \equiv \mathbf{push}_i E$ then either $E_2 \equiv \lambda_x.F$ (and $E_1 \circ E_2$ is a redex) or $E_2 \equiv E'_2 \circ E''_2$ and by hypothesis E_2 has a redex (thus $E_1 \circ E_2$ is reducible). Otherwise $E_1 \equiv E'_1 \circ E''_1$ and by hypothesis E_1 has a redex (thus $E_1 \circ E_2$ is reducible). \square

Property 2 is a direct consequence of the two previous properties. If $E:R_i \tau$ has a normal form N then $E \xrightarrow{*} N$. By Property 10, $N:R_i \tau$ and by Property 11 (N is not reducible) $N \equiv \mathbf{push}_i V$, so $E \xrightarrow{*} \mathbf{push}_i V$. Same thing with $E:\sigma \rightarrow_i \tau$ \square

Another consequence of the type system, is that the reduction of typed closed expressions can be specified by the following natural semantics:

$$\frac{E_1 \triangleright \mathbf{push}_i V \quad E_2 \triangleright \lambda_x.F \quad F[V/x] \triangleright N}{E_1 \circ E_2 \triangleright N} \quad (\text{with } N \text{ a normal form})$$

and we have

Property 12 *For all typed closed expression E $E \xrightarrow{*} N \Leftrightarrow E \triangleright N$ (with N a normal form)*

Proof. Induction on the reduction tree. Evident if E is canonical (by the implicit rule $N \triangleright N$). If $E \equiv E_1 \circ E_2$, since all reduction strategies are normalizing :

$$\begin{aligned} E \xrightarrow{*} N &\Leftrightarrow E_1 \xrightarrow{*} \mathbf{push}_i V \text{ and } E_2 \xrightarrow{*} \lambda_x.F \text{ and } F[V/x] \xrightarrow{*} N && (\text{Property 2}) \\ &\Leftrightarrow E_1 \triangleright \mathbf{push}_i V \text{ and } E_2 \triangleright \lambda_x.F \text{ and } F[V/x] \triangleright N && (\text{by induction hypothesis}) \\ &\Leftrightarrow E \triangleright N && \square \end{aligned}$$

C. Laws (L2) and (L3)

As stated, laws (L2) and (L3) are valid only within the corresponding of a classical consistent extension of the λ -calculus. Our framework comprises the two additional rules:

- (Ω) If the closed expressions M and N do not have a (weak) normal form then $M = N$
- (ω) Let $\Gamma \cup \{z:\sigma\} \vdash M, N:\tau$ if for all closed expression $\vdash Z:\sigma, M[Z/z] = N[Z/z]$ then $M = N$

Intuitively, the motivation behind this extension is that our only concern is that two equal terms behave the same during the reduction. That is, we accept to replace an expression by another as long as they are equal after their free variables are instantiated or to replace a looping expression by another looping expression.

One may wonder whether the rules (assoc), (β_i), (η_i), (Ω), (ω) define a consistent theory. Recall that the meaning of Λ_i expressions are defined in terms of λ -expressions (Section 2.5). It is sufficient to verify after the instantiation that these rules are valid in a consistent theory of λ -calculus. With all the instantiations we have considered, it is easy to check that these rules are valid in the lambda theory $\mathcal{H}\omega$ (according to Barendregt terminology [5]). If we write $\llbracket E \rrbracket$ the λ -expression obtained after instantiation of a Λ_i -expression E , then it amounts to showing that if $E = F$ in Λ_i then $\llbracket E \rrbracket = \llbracket F \rrbracket$ in $\mathcal{H}\omega$. The theory $\mathcal{H}\omega$ is defined by the classic laws of the λ -calculus but also identifies unsolvable terms (a more general case than terms without weak normal form) (see [5] chapters 16 and 17).

Proof of law (L2). Let $z_1:\sigma_1, \dots, z_n:\sigma_n$ the free variables of $E_1 \circ (\lambda_{\dot{x}}x.E_2 \circ E_3)$ then

$\forall Z_1:\sigma_1, \dots, Z_n:\sigma_n$ closed

$$\begin{aligned} & (E_1 \circ (\lambda_{\dot{x}}x.E_2 \circ E_3))[Z_1, \dots, Z_n/z_1, \dots, z_n] \\ &= E_1[Z_1, \dots, Z_n/z_1, \dots, z_n] \circ (\lambda_{\dot{x}}x.E_2 [Z_1, \dots, Z_n/z_1, \dots, z_n] \circ E_3[Z_1, \dots, Z_n/z_1, \dots, z_n]) \end{aligned}$$

If $E_1[Z_1, \dots, Z_n/z_1, \dots, z_n]$ does not have a normal form then both expressions $(E_1 \circ (\lambda_{\dot{x}}x.E_2 \circ E_3))[Z_1, \dots, Z_n/z_1, \dots, z_n]$ and $(E_2 \circ E_1 \circ (\lambda_{\dot{x}}x.E_3))[Z_1, \dots, Z_n/z_1, \dots, z_n]$ will not have normal forms. By (Ω) they are therefore equal and by (ω) we have

$$E_1 \circ (\lambda_{\dot{x}}x.E_2 \circ E_3) = E_2 \circ E_1 \circ (\lambda_{\dot{x}}x.E_3)$$

Otherwise since $E_1[Z_1, \dots, Z_n/z_1, \dots, z_n]$ is closed, we know (Property 2) that there exists N such that $E_1[Z_1, \dots, Z_n/z_1, \dots, z_n] = \mathbf{push}_i N$ so

$$\begin{aligned} & (E_1 \circ (\lambda_{\dot{x}}x.E_2 \circ E_3))[Z_1, \dots, Z_n/z_1, \dots, z_n] \\ &= \mathbf{push}_i N \circ (\lambda_{\dot{x}}x.E_2 [Z_1, \dots, Z_n/z_1, \dots, z_n] \circ E_3[Z_1, \dots, Z_n/z_1, \dots, z_n]) \\ &= E_2[Z_1, \dots, Z_n/z_1, \dots, z_n] \circ E_3 [Z_1, \dots, Z_n/z_1, \dots, z_n][N/x] \quad (\beta_i) \text{ and } x \text{ is not free in } E_2 \\ &= E_2[Z_1, \dots, Z_n/z_1, \dots, z_n] \circ \mathbf{push}_i N \circ (\lambda_{\dot{x}}x.E_3[Z_1, \dots, Z_n/z_1, \dots, z_n]) \quad (\beta_i) \\ &= E_2[Z_1, \dots, Z_n/z_1, \dots, z_n] \circ E_1[Z_1, \dots, Z_n/z_1, \dots, z_n] \circ (\lambda_{\dot{x}}x.E_3[Z_1, \dots, Z_n/z_1, \dots, z_n]) \\ &= (E_2 \circ E_1 \circ (\lambda_{\dot{x}}x.E_3)) [Z_1, \dots, Z_n/z_1, \dots, z_n] \end{aligned}$$

So, for all Z_1, \dots, Z_n closed

$$(E_1 \circ (\lambda_i x. E_2 \circ E_3))(Z_1, \dots, Z_n/z_1, \dots, z_n) = (E_2 \circ E_1 \circ (\lambda_i x. E_3))(Z_1, \dots, Z_n/z_1, \dots, z_n)$$

and by (ω) we have $(E_1 \circ (\lambda_i x. E_2 \circ E_3)) = (E_2 \circ E_1 \circ (\lambda_i x. E_3))$

The proof for law (L3) is similar. \square

d. We show here that $\mathcal{V}a$ yields well-typed expressions.

Property 13 $\forall E \in \Lambda, E \text{ closed} \vdash E: \sigma \Rightarrow \vdash \mathcal{V}a \llbracket E \rrbracket : R_s \bar{\sigma}$ with $\overline{\sigma \rightarrow \tau} = \bar{\sigma} \rightarrow_s R_s \bar{\tau}$ and $\bar{\alpha} = \alpha$ (α type variable)

Proof. We prove the stronger property let E an expression with free variables $\{x_1 \dots x_n\}$ such that $\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash E:\sigma$ then $\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash \mathcal{V}a \llbracket E \rrbracket : R_s \bar{\sigma}$.

Proof. By structural induction.

- $E \equiv x_i \quad \{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash E:\alpha_i$ then $\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash \mathbf{push}_s x_i (\equiv \mathcal{V}a \llbracket x_i \rrbracket) : R_s \alpha_i$
- $E \equiv \lambda z. E' \quad \{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash E:\sigma \rightarrow \tau$ that is $\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \cup \{z:\sigma\} \vdash E':\tau$.

By induction hypothesis, $\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \cup \{z:\sigma\} \vdash \mathcal{V}a \llbracket E' \rrbracket : R_s \tau$

and $\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash \lambda_{s,z}. \mathcal{V}a \llbracket E' \rrbracket : \sigma \rightarrow_s R_s \tau (\equiv \overline{\sigma \rightarrow \tau})$

hence $\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash \mathbf{push}_s (\lambda_{s,z}. \mathcal{V}a \llbracket E' \rrbracket) (\equiv \mathcal{V}a \llbracket \lambda z. E' \rrbracket) : R_s (\overline{\sigma \rightarrow \tau})$

- $E \equiv E_1 E_2 \quad \{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash E_1:\sigma \rightarrow \tau$ and $\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash E_2:\sigma$

By induction hypothesis,

$$\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash \mathcal{V}a \llbracket E_1 \rrbracket : R_s (\overline{\sigma \rightarrow \tau}) \text{ and } \{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash \mathcal{V}a \llbracket E_2 \rrbracket : R_s \bar{\sigma}$$

and $\vdash \mathbf{app} : (\overline{\sigma \rightarrow \tau}) \rightarrow_s (\overline{\sigma \rightarrow \tau})$ thus $\{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash \mathcal{V}a \llbracket E_1 \rrbracket \circ \mathbf{app} : \overline{\sigma \rightarrow \tau}$

$$\text{and } \{x_1:\alpha_1, \dots, x_n:\alpha_n\} \vdash \mathcal{V}a \llbracket E_2 \rrbracket \circ \mathcal{V}a \llbracket E_1 \rrbracket \circ \mathbf{app} : R_s \bar{\tau} \quad \square$$

e. Proof of Property 3.

The proof of Property 3 needs two preliminary lemmas.

A context $\mathbf{X}[\]$ is said to be closed if for all expressions E, F and variable x , $\mathbf{X}[E] [F/x] \equiv \mathbf{X}[E [F/x]]$ (i.e. a closed context does not introduce free variables nor does it bind free variables).

Lemma 14 Let $\mathbf{X}[\], \mathbf{Y}[\], \mathbf{Z}[\][\]$ be closed contexts and \mathcal{T} a transformation such that

$$\mathcal{T} \llbracket x \rrbracket = \mathbf{X} \llbracket x \rrbracket \quad \mathcal{T} \llbracket \lambda x. E \rrbracket = \mathbf{Y} \llbracket \lambda x. \mathcal{T} \llbracket E \rrbracket \rrbracket \quad \mathcal{T} \llbracket E_1 E_2 \rrbracket = \mathbf{Z} \llbracket \mathcal{T} \llbracket E_1 \rrbracket \rrbracket \llbracket \mathcal{T} \llbracket E_2 \rrbracket \rrbracket$$

$$\text{then for all } E \text{ and } F \text{ such that } \mathcal{T} \llbracket F \rrbracket \equiv \mathbf{X} \llbracket F' \rrbracket \quad \mathcal{T} \llbracket E[F/x] \rrbracket \equiv \mathcal{T} \llbracket E \rrbracket \llbracket F'/x \rrbracket$$

Proof. By structural induction.

- $E \equiv x \quad \mathcal{T} \llbracket x[F/x] \rrbracket \equiv \mathcal{T} \llbracket F \rrbracket \equiv \mathbf{X} \llbracket F' \rrbracket \equiv \mathbf{X} \llbracket x[F'/x] \rrbracket \equiv (\mathbf{X} \llbracket x \rrbracket) \llbracket F'/x \rrbracket \equiv \mathcal{T} \llbracket x \rrbracket \llbracket F'/x \rrbracket$
since \mathbf{X} closed
- $x \notin E \quad \mathcal{T} \llbracket E[F/x] \rrbracket \equiv \mathcal{T} \llbracket E \rrbracket \equiv \mathcal{T} \llbracket E \rrbracket \llbracket F'/x \rrbracket$ since \mathcal{T} does not introduce free variables
- $E \equiv \lambda z. E' \quad (z \neq x) \quad \mathcal{T} \llbracket (\lambda z. E')[F/x] \rrbracket \equiv \mathcal{T} \llbracket \lambda z. (E'[F/x]) \rrbracket \equiv \mathbf{Y} \llbracket \lambda z. \mathcal{T} \llbracket E'[F/x] \rrbracket \rrbracket$
 $\equiv \mathbf{Y} \llbracket \lambda z. \mathcal{T} \llbracket E' \rrbracket \llbracket F'/x \rrbracket \rrbracket$ by induction hypothesis
 $\equiv \mathbf{Y} \llbracket \lambda z. \mathcal{T} \llbracket E' \rrbracket \rrbracket \llbracket F'/x \rrbracket$ since \mathbf{Y} closed
 $\equiv \mathcal{T} \llbracket \lambda z. E' \rrbracket \llbracket F'/x \rrbracket$
- $E \equiv E_1 E_2 \quad \mathcal{T} \llbracket (E_1 E_2)[F/x] \rrbracket \equiv \mathcal{T} \llbracket (E_1 \llbracket F'/x \rrbracket) (E_2 \llbracket F'/x \rrbracket) \rrbracket$
 $\equiv \mathbf{Z} \llbracket \mathcal{T} \llbracket E_1 \llbracket F'/x \rrbracket \rrbracket \rrbracket \llbracket \mathcal{T} \llbracket E_2 \llbracket F'/x \rrbracket \rrbracket \rrbracket$
 $\equiv \mathbf{Z} \llbracket \mathcal{T} \llbracket E_1 \rrbracket \llbracket F'/x \rrbracket \rrbracket \llbracket \mathcal{T} \llbracket E_2 \rrbracket \llbracket F'/x \rrbracket \rrbracket$ by induction hypothesis
 $\equiv \mathbf{Z} \llbracket \mathcal{T} \llbracket E_1 \rrbracket \rrbracket \llbracket \mathcal{T} \llbracket E_2 \rrbracket \rrbracket \llbracket F'/x \rrbracket$ since \mathbf{Z} closed
 $\equiv \mathcal{T} \llbracket E_1 E_2 \rrbracket \llbracket F'/x \rrbracket$ □

In particular, the transformation $\mathcal{V}a$ (but also $\mathcal{V}m, \mathcal{N}m, \mathcal{N}a$) verifies the conditions of the lemma. So, we have

$$\mathcal{V}a \llbracket E[F/x] \rrbracket \equiv \mathcal{V}a \llbracket E \rrbracket \llbracket F'/x \rrbracket \quad \text{if } \mathcal{V}a \llbracket F \rrbracket \equiv \mathbf{push}_s F'$$

We will prove Property 3 for the notion of reduction \succ which is equivalent to $\overset{*}{\succ}$ (Property 12). We need the following lemma

Lemma 15 $\forall E \text{ closed} \in \Lambda \quad \mathcal{V}a \llbracket E \rrbracket \succ X \Rightarrow \exists N \in \Lambda \text{ such that } \mathcal{V}a \llbracket N \rrbracket \equiv X$

Proof. If $E \equiv \lambda x. F$ then $N \equiv E$. If $E \equiv E_1 E_2$ then $\mathcal{V}a \llbracket E \rrbracket \equiv \mathcal{V}a \llbracket E_2 \rrbracket \circ \mathcal{V}a \llbracket E_1 \rrbracket \circ \mathbf{app}$. By Property 13 and Property 2 $\mathcal{V}a \llbracket E \rrbracket \succ \mathbf{push}_s X$ so there must be a derivation $\mathcal{V}a \llbracket E_2 \rrbracket \succ \mathbf{push}_s V'$, $\mathcal{V}a \llbracket E_1 \rrbracket \succ \mathbf{push}_s (\lambda_{s,x}. F')$ and $F'[V'/x] \succ \mathbf{push}_s X$. By induction hypothesis, there are V such that $\mathcal{V}a \llbracket V \rrbracket \equiv \mathbf{push}_s V'$ and Z such that $\mathcal{V}a \llbracket Z \rrbracket \equiv \mathbf{push}_s (\lambda_{s,x}. F')$ (i.e. $Z \equiv \lambda x. F$ with $\mathcal{V}a \llbracket F \rrbracket \equiv F'$). So $F'[V'/x] \equiv \mathcal{V}a \llbracket F \rrbracket \llbracket V'/x \rrbracket \equiv \mathcal{V}a \llbracket F[V/x] \rrbracket$ (Lemma 14) and from $\mathcal{V}a \llbracket F[V/x] \rrbracket \succ \mathbf{push}_s X$ we deduce by induction hypothesis that there is N such that $\mathcal{V}a \llbracket N \rrbracket \equiv \mathbf{push}_s X$. □

Call-by-value reduction is described by the following natural operational semantics (with V and N normal forms):

$$\frac{E_1 \xrightarrow{\text{cbv}} \lambda x. F \quad E_2 \xrightarrow{\text{cbv}} V \quad F[V/x] \xrightarrow{\text{cbv}} N}{E_1 E_2 \xrightarrow{\text{cbv}} N}$$

The proof of Property 3 is on the shape of the reduction trees.

Axioms.

(\Rightarrow) If E is not reducible it is of the form $\lambda x.F$ (E is closed) and $\mathcal{V}a \llbracket \lambda x.F \rrbracket \equiv \mathbf{push}_s (\lambda_{s,x}. \mathcal{V}a \llbracket F \rrbracket)$ which is not reducible.

(\Leftarrow) If $\mathcal{V}a \llbracket E \rrbracket$ is not reducible then E is of the form $\lambda x.F$. Indeed, since E is closed, the only alternative would be $E \equiv (\lambda x.F) E_1 \dots E_n$ but then $\mathcal{V}a \llbracket E \rrbracket$ would be reducible (there would be the redex $\mathbf{push}_s (\lambda_{s,x}. \mathcal{V}a \llbracket F \rrbracket) \circ \mathbf{app}$). So E is not reducible.

Induction.

(\Rightarrow) E is reducible, that is, $E \equiv E_1 E_2$, $E_1 \xrightarrow{cbv} \lambda x.F$, $E_2 \xrightarrow{cbv} V$ and $F[V/x] \xrightarrow{cbv} N$. By induction hypothesis, we have $\mathcal{V}a \llbracket E_1 \rrbracket \succ \mathcal{V}a \llbracket \lambda x.F \rrbracket$, $\mathcal{V}a \llbracket E_2 \rrbracket \succ \mathcal{V}a \llbracket V \rrbracket$ and $\mathcal{V}a \llbracket F[V/x] \rrbracket \succ \mathcal{V}a \llbracket N \rrbracket$. Since V is closed $\mathcal{V}a \llbracket V \rrbracket \equiv \mathbf{push}_s V'$ and, by Lemma 14, $\mathcal{V}a \llbracket F \rrbracket [V'/x] \equiv \mathcal{V}a \llbracket F[V/x] \rrbracket$, we have $\mathcal{V}a \llbracket E_2 \rrbracket \succ \mathbf{push}_s V'$, $\mathcal{V}a \llbracket E_1 \rrbracket \circ \mathbf{app} \succ \lambda_{s,x}. \mathcal{V}a \llbracket F \rrbracket$ and $\mathcal{V}a \llbracket F \rrbracket [V'/x] \succ \mathcal{V}a \llbracket N \rrbracket$ therefore, $\mathcal{V}a \llbracket E_1 E_2 \rrbracket \equiv \mathcal{V}a \llbracket E_2 \rrbracket \circ \mathcal{V}a \llbracket E_1 \rrbracket \circ \mathbf{app} \succ \mathcal{V}a \llbracket N \rrbracket$.

(\Leftarrow) $\mathcal{V}a \llbracket E \rrbracket$ is reducible, that is, $E \equiv E_1 E_2$ and $\mathcal{V}a \llbracket E \rrbracket \succ N'$. Since $\mathcal{V}a \llbracket E \rrbracket$ is well-typed (Property 13), the reduction tree must be of the form $\mathcal{V}a \llbracket E_2 \rrbracket \succ \mathbf{push}_s V'$, $\mathcal{V}a \llbracket E_1 \rrbracket \succ \mathbf{push}_s (\lambda_{s,x} F')$ and $F'[V'/x] \succ N'$. By Lemma 15 we know that there is V such that $\mathcal{V}a \llbracket V \rrbracket \equiv \mathbf{push}_s V'$, Z such that $\mathcal{V}a \llbracket Z \rrbracket \equiv \mathbf{push}_s (\lambda_{s,x}. F)$, (i.e. $Z \equiv \lambda x.F$ with $\mathcal{V}a \llbracket F \rrbracket \equiv F'$) and N such that $\mathcal{V}a \llbracket N \rrbracket \equiv N'$. So, by induction hypothesis, $E_1 \xrightarrow{cbv} \lambda x.F$, $E_2 \xrightarrow{cbv} V$. By Lemma 14, $\mathcal{V}a \llbracket F \rrbracket [V'/x] \equiv \mathcal{V}a \llbracket F[V/x] \rrbracket$ and, by induction hypothesis, $F[V/x] \xrightarrow{cbv} N$, thus $E \xrightarrow{cbv} N$. \square

The proofs for the others \mathcal{V} and \mathcal{N} transformations are similar.

f. Variants of $\mathcal{V}a$

The transformation $\mathcal{V}a_L$ which implements a describing left-to-right call-by-value is expressed as $\mathcal{V}a$ except the rule for application which becomes

$$\mathcal{V}a_L \llbracket E_1 E_2 \rrbracket = \mathcal{V}a_L \llbracket E_1 \rrbracket \circ \mathcal{V}a_L \llbracket E_2 \rrbracket \circ \mathbf{app}_L \quad \text{with} \quad \mathbf{app}_L = \lambda_{s,x}. \lambda_{s,f}. \mathbf{push}_s x \circ f$$

This compilation choice is taken by the Cam [10].

Transformations $\mathcal{V}a$ and $\mathcal{V}a_L$ may produce expressions such as $\mathbf{push}_s E_1 \circ \mathbf{push}_s E_2 \circ \dots \circ \mathbf{push}_s E_n \circ \dots$. The reduction of such expressions requires a structure (such as a stack) capable of storing an arbitrary number of intermediate results. Some implementations (such as the SML-NJ compiler [2]) make the choice of not using a data stack and, therefore, disallow several pushes in a row. In this case, the rule for applications of $\mathcal{V}a_L$ should be changed into

$$\mathcal{V}a_f \llbracket E_1 E_2 \rrbracket = \mathcal{V}a_f \llbracket E_1 \rrbracket \circ (\lambda_{s,m}. \mathcal{V}a_f \llbracket E_2 \rrbracket \circ m)$$

Reading the transformation rules as grammar rules, it is clear that $\mathcal{V}a_f$ never produces expressions where two \mathbf{push}_s occur in a row (such as $\mathbf{push}_s A \circ \mathbf{push}_s B$). For these expressions, the component on which \mathbf{push}_s and λ_s act may be a single register. Another possible motivation for this style of transformation (called stackless) is that the produced expressions

now possess a unique redex throughout the reduction. The reduction sequence must be sequential and is unique.

The two variations $\mathcal{V}a_L$ and $\mathcal{V}a_f$ are easily derived from $\mathcal{V}a$ using conversions rules and algebraic laws.

g. Variant of $\mathcal{N}a$

Like $\mathcal{V}a$, transformation $\mathcal{N}a$ may produce expressions such as $\mathbf{push}_s E_1 \circ \dots \circ \mathbf{push}_s E_n$ which require a stack to store intermediate results. To get a stackless variant of $\mathcal{N}a$, the rule for compositions should be changed into:

$$\mathcal{N}a_f \llbracket E_1 E_2 \rrbracket = \mathbf{push}_s (\mathcal{N}a_f \llbracket E_2 \rrbracket) \circ (\lambda_s a. \mathcal{N}a_f \llbracket E_1 \rrbracket) \circ (\lambda_s f. \mathbf{push}_s a \circ f)$$

With this variant, the component on which \mathbf{push}_s and λ_s act may be a single register.

h. The combinator \mathbf{grab}_s and the mark ε can be defined in Λ_s much in the same way that conditional expressions can be defined in pure λ -calculus. A possibility is:

$$\mathbf{grab}_s E \equiv \mathbf{push}_s E \circ \lambda_s x. \lambda_s (m, v). \mathbf{push}_s (\mathbf{push}_s (\mu, x)) \circ \mathbf{push}_s (\mathbf{push}_s v \circ x) \circ m$$

Each argument is associated with a mark in a pair. The mark $\mu \equiv \lambda_s x. \lambda_s y. x$ selects the first alternative (apply the function E) whereas $\varepsilon \equiv (\lambda_s x. \lambda_s y. y, id)$ is a mark (associated with a dummy function id) selecting the second alternative (yield E as result). It is obviously much more efficient to implement \mathbf{grab}_s using the predefined conditional operator provided by the target machine.

i. Variants of $\mathcal{V}m$

For call-by-value, a generic transformation using marks can be described as follows :

$$\mathcal{V}m_g \llbracket x \rrbracket = Xx$$

$$\mathcal{V}m_g \llbracket \lambda x. E \rrbracket = \mathcal{Y}(\lambda_s x. \mathcal{V}m_g \llbracket E \rrbracket)$$

$$\mathcal{V}m_g \llbracket E_1 E_2 \rrbracket = \mathbf{push}_s \varepsilon \circ \mathcal{V}m_g \llbracket E_2 \rrbracket \circ \mathcal{V}m_g \llbracket E_1 \rrbracket$$

$$X, \mathcal{Y} \text{ and } Z \text{ being combinators such that } \mathcal{Y} = X \circ Z$$

$$\mathbf{push}_s \varepsilon \circ \mathcal{Y}(E) \Rightarrow \mathbf{push}_s Z(E)$$

$$\text{and } \mathbf{push}_s V \circ \mathcal{Y}(E) \Rightarrow \mathbf{push}_s V \circ E$$

Figure 22 Generic Compilation of Right-to-Left Call-by-Value with Marks ($\mathcal{V}m_g$)

We get back $\mathcal{V}m$ by taking $\mathcal{Y}=\mathcal{X}=\mathbf{grab}_s$ and $\mathcal{Z}=\mathbf{id}$. The second “canonical” transformation (see [32] page 27) is $\mathcal{V}m'$ with $\mathcal{Y}=\mathcal{Z}=\mathbf{grab}_{sL}$ and $\mathcal{X}=\mathbf{id}$ (i.e. the reduction rule of \mathbf{grab}_{sL} is recursive). By making all the \mathbf{grab}_s explicit in the code, $\mathcal{V}m$ permits more simplifications than the alternative. For example,

$$\mathcal{V}m \llbracket (\lambda x.x x) (\lambda y.E) \rrbracket = \mathbf{push}_s (\lambda_s y. \mathcal{V}m \llbracket E \rrbracket) \circ (\lambda_s x. \mathbf{push}_s x \circ x)$$

(one mark& \mathbf{grab}_s has been simplified), whereas the other transformation $\mathcal{V}m'$ yields $\mathbf{push}_s (\mathbf{grab}_{sL} (\lambda_s y. \mathcal{V}m' \llbracket E \rrbracket)) \circ (\lambda_s x. \mathbf{push}_s \varepsilon \circ x \circ x)$ and \mathbf{grab}_{sL} would be executed twice.

j. Relationship with CPS Conversion

Since CPS expressions have only one redex throughout the reduction, the closest transformations are the stackless ones (i.e. $\mathcal{V}a_f^{\textcircled{1}}$ and $\mathcal{N}a_f^{\textcircled{2}}$). Indeed, if we take the definitions

$$\text{(DEF1)} \quad \lambda_s x.X = \lambda c. \lambda x.X c \quad \mathbf{push}_s N = \lambda c.c N \quad \circ = \lambda a. \lambda b. \lambda c.a (b c)$$

(which satisfy (assoc), (β_s) , and (η_s)) we can rewrite $\mathcal{V}a_f$ as follows:

$$\mathcal{V}a_f \llbracket x \rrbracket = \mathbf{push}_s x = \lambda c.c x \quad \text{(DEF1)}$$

$$\mathcal{V}a_f \llbracket \lambda x.E \rrbracket = \mathbf{push}_s (\lambda_s x. \mathcal{V}a_f \llbracket E \rrbracket) = \lambda c.c (\lambda c. \lambda x. \mathcal{V}a_f \llbracket E \rrbracket c) \quad \text{(DEF1)}$$

$$\mathcal{V}a_f \llbracket E_1 E_2 \rrbracket = \lambda c. \mathcal{V}a_f \llbracket E_1 \rrbracket (\lambda m_1. \mathcal{V}a_f \llbracket E_2 \rrbracket (\lambda m_2. m_1 c m_2)) \quad \text{(DEF1), } (\eta)$$

which is exactly Fischer’s CPS transformation [21].

As far as types are concerned we saw that if $E : \sigma$ then $\mathcal{V}a_f \llbracket E \rrbracket : R_s \bar{\sigma}$ with $\bar{\sigma} \rightarrow \tau = \bar{\sigma} \rightarrow_s R_s \bar{\tau}$ and $\alpha = \bar{\alpha}$. We recognize CPS types by giving to R_s and \rightarrow_s the meanings:

$$R_s \sigma = (\sigma \rightarrow Ans) \rightarrow Ans \quad \text{and} \quad \sigma \rightarrow_s R_s \tau = (\tau \rightarrow Ans) \rightarrow \sigma \rightarrow Ans$$

Ans being the distinguished type of answers. Note that if n-ary functions are allowed we should add the rule $\sigma \rightarrow_s (\tau \rightarrow Ans) \rightarrow \upsilon = (\tau \rightarrow Ans) \rightarrow \sigma \rightarrow \upsilon$

k. An inversion transformation for Λ_s -expressions

As for CPS-expressions [15], it is also possible to design an inverse transformation. The transformation $\llbracket \cdot \rrbracket^{-1}$ (Figure 23) can be seen as a generic decompilation transformation and it is easy to show by structural induction that

Property 16 For all Λ -expression E , $\llbracket C \llbracket E \rrbracket \rrbracket^{-1} = E$ (for $C = \mathcal{V}a, \mathcal{V}a_L, \mathcal{V}a_p, \mathcal{N}a, \mathcal{N}a_p, \mathcal{N}m$)

Note that the transformation $\llbracket \cdot \rrbracket^{-1}$ is just a left inverse. In order to get a true inverse transformation, the domain of $\llbracket \cdot \rrbracket^{-1}$ should be restricted to the expressions encoding an evaluation strategy.

$$\begin{array}{ll}
\llbracket \cdot \rrbracket^{-1} : \Lambda_s \rightarrow \Lambda & \\
\llbracket x \rrbracket^{-1} = x & \llbracket \mathbf{push}_s E \rrbracket^{-1} = \llbracket E \rrbracket^{-1} \\
\llbracket \lambda_s x. E \rrbracket^{-1} = \lambda x. \llbracket E \rrbracket^{-1} & \llbracket E_1 \circ E_2 \rrbracket^{-1} = \llbracket E_2 \rrbracket^{-1} \llbracket E_1 \rrbracket^{-1}
\end{array}$$

Figure 23 Back to λ -expressions

L. Proof of Property 7

Call-by-name reduction is described by the following natural operational semantics:

$$\frac{E_1 \xrightarrow{cbn} \lambda x. F \quad F[E_2/x] \xrightarrow{cbn} N}{E_1 E_2 \xrightarrow{cbn} N} \quad N \text{ normal form}$$

The proof of Property 7 is on the shape of the reduction trees. We need two lemmas.

Lemma 17 $\mathcal{G}[\llbracket E \rrbracket \circ \lambda_s x. \mathcal{G}[\llbracket F \rrbracket]] = \mathcal{G}[\llbracket F \rrbracket][\mathcal{G}[\llbracket E \rrbracket] / \mathbf{push}_s x]$

The condition $(Cond\mathcal{G})$ insures that $\mathcal{G}[\llbracket E \rrbracket] = \mathbf{push}_s V$. So, $\mathcal{G}[\llbracket E \rrbracket \circ \lambda_s x. \mathcal{G}[\llbracket F \rrbracket]] = \mathcal{G}[\llbracket F \rrbracket][V/x]$. Using the definition of \mathcal{G} (Figure 8), it is easy to check that a free variable x of an expression $\mathcal{G}[\llbracket E \rrbracket]$ occurs only as $\mathbf{push}_s x$. So, $\mathcal{G}[\llbracket F \rrbracket][V/x] = \mathcal{G}[\llbracket F \rrbracket][\mathbf{push}_s V / \mathbf{push}_s x] = \mathcal{G}[\llbracket F \rrbracket][\mathcal{G}[\llbracket E \rrbracket] / \mathbf{push}_s x]$.

Moreover, using (L5), it is easy to prove by structural induction that

Lemma 18 $\mathcal{G}[\llbracket E_1 \rrbracket][\mathcal{G}[\llbracket E_2 \rrbracket] / \mathbf{push}_s x] = \mathcal{G}[\llbracket E_1 \rrbracket[E_2/x]]$

Axioms.

If E is not reducible, it is of the form $\lambda x. F$ (E is closed). We have then $E \equiv V$ and the property is trivially verified.

Induction.

If E is reducible, that is, $E \equiv E_1 E_2$, $E_1 \xrightarrow{cbn} \lambda x. F$ and $F[E_2/x] \xrightarrow{cbn} N$. By induction hypothesis, we have $\mathcal{G}[\llbracket E_1 \rrbracket] \circ \mathbf{unwind}_s = \mathcal{G}[\llbracket \lambda x. F \rrbracket] \circ \mathbf{unwind}_s$ and $\mathcal{G}[\llbracket F[E_2/x] \rrbracket] \circ \mathbf{unwind}_s = \mathcal{G}[\llbracket N \rrbracket] \circ \mathbf{unwind}_s$. So $\mathcal{G}[\llbracket E_1 E_2 \rrbracket] \circ \mathbf{unwind}_s \equiv \mathcal{G}[\llbracket E_2 \rrbracket] \circ \mathcal{G}[\llbracket E_1 \rrbracket] \circ \mathbf{mkApp}_s \circ \mathbf{unwind}_s$

$$\begin{aligned}
&= \mathcal{G}[\llbracket E_2 \rrbracket] \circ \mathcal{G}[\llbracket E_1 \rrbracket] \circ \mathbf{unwind}_s && (\mathcal{G}\mathcal{N}m3) \\
&= \mathcal{G}[\llbracket E_2 \rrbracket] \circ \mathcal{G}[\llbracket \lambda x. F \rrbracket] \circ \mathbf{unwind}_s && \text{induction hypothesis} \\
&\equiv \mathcal{G}[\llbracket E_2 \rrbracket] \circ \mathbf{push}_s(\lambda_s x. \mathcal{G}[\llbracket F \rrbracket]) \circ \mathbf{mkFun}_s \circ \mathbf{unwind}_s && (\text{def. } \mathcal{G}) \\
&= (\mathcal{G}[\llbracket E_2 \rrbracket] \circ \lambda_s x. \mathcal{G}[\llbracket F \rrbracket]) \circ \mathbf{unwind}_s && (\mathcal{G}\mathcal{N}m2)
\end{aligned}$$

$$\begin{aligned}
&= (\mathcal{G} \llbracket F \rrbracket [\mathcal{G} \llbracket E_2 \rrbracket / \mathbf{push}_s x]) \circ \mathbf{unwind}_s && \text{(Lemma 17)} \\
&= (\mathcal{G} \llbracket F \llbracket E_2/x \rrbracket \rrbracket) \circ \mathbf{unwind}_s && \text{(Lemma 18)} \\
&= \mathcal{G} \llbracket N \rrbracket \circ \mathbf{unwind}_s && \text{induction hypothesis } \square
\end{aligned}$$

m. Proof of Property 8

In order to prove $\mathbf{push}_e() \circ \mathcal{A}_s \llbracket E \rrbracket () = E$, we prove by induction the more general property:

$$\mathbf{push}_e \rho \circ \mathcal{A}_s \llbracket E \rrbracket \rho = E \quad \text{with } \rho = (\dots(), x_n) \dots, x_0 \text{ and } \text{FV}(E) = \{x_0, \dots, x_n\}$$

where $\text{FV}(E)$ is the set of free variables of E .

We will make use of the fact that, if $\text{FV}(E) \subseteq \rho$ then $\mathcal{A}_s \llbracket E \rrbracket \rho$ is closed (easy to check). Note also that it is important that the expression $E_1 \circ E_2$ is well-typed since we use law (L3) which relies on types.

- $E \equiv E_1 \circ E_2$

$$\begin{aligned}
\mathbf{push}_e \rho \circ \mathcal{A}_s \llbracket E_1 \circ E_2 \rrbracket \rho &= \mathbf{push}_e \rho \circ \mathbf{dupl}_e \circ \mathcal{A}_s \llbracket E_1 \rrbracket \rho \circ \mathbf{swap}_{se} \circ \mathcal{A}_s \llbracket E_2 \rrbracket \rho \\
&= \mathbf{push}_e \rho \circ (\mathbf{push}_e \rho \circ \mathcal{A}_s \llbracket E_1 \rrbracket \rho) \circ \lambda_s x. \lambda_e e. \mathbf{push}_s x \circ \mathbf{push}_e e \circ \mathcal{A}_s \llbracket E_2 \rrbracket \rho (\beta_s), (\beta_e) \\
&= \mathbf{push}_e \rho \circ E_1 \circ \lambda_s x. \lambda_e e. \mathbf{push}_s x \circ \mathbf{push}_e e \circ \mathcal{A}_s \llbracket E_2 \rrbracket \rho && \text{by induction hypothesis} \\
&= E_1 \circ \mathbf{push}_e \rho \circ \mathcal{A}_s \llbracket E_2 \rrbracket \rho && \text{(L3), } (\beta_e), (\eta_s) \\
&= E_1 \circ E_2 && \text{by induction hypothesis}
\end{aligned}$$

- $E \equiv \mathbf{push}_s V$

$$\begin{aligned}
\mathbf{push}_e \rho \circ \mathcal{A}_s \llbracket \mathbf{push}_s V \rrbracket \rho &= \mathbf{push}_e \rho \circ \mathbf{push}_s (\mathcal{A}_s \llbracket V \rrbracket \rho) \circ \mathbf{mkclos} \\
&= \mathbf{push}_s (\mathbf{push}_e \rho \circ \mathcal{A}_s \llbracket V \rrbracket \rho) && \mathbf{mkclos} \text{ def. } (\beta_s), (\beta_e) \\
&= \mathbf{push}_s V && \text{by induction hypothesis}
\end{aligned}$$

- $E \equiv \lambda_s x. F$

$$\begin{aligned}
\mathbf{push}_e \rho \circ \mathcal{A}_s \llbracket \lambda_s x. F \rrbracket \rho &= \mathbf{push}_e \rho \circ \mathbf{bind} \circ \mathcal{A}_s \llbracket F \rrbracket (\rho, x) \\
&= \mathbf{push}_e \rho \circ \lambda_e e. \lambda_s y. \mathbf{push}_e (e, y) \circ \mathcal{A}_s \llbracket F \rrbracket (\rho, x) && \mathbf{bind} \text{ def.} \\
&= \mathbf{push}_e \rho \circ \lambda_e e. \lambda_s x. \mathbf{push}_e (e, x) \circ \mathcal{A}_s \llbracket F \rrbracket (\rho, x) && \mathcal{A}_s \llbracket F \rrbracket (\rho, x) \text{ closed and } (\alpha_s) \\
&= \lambda_s x. \mathbf{push}_e (\rho, x) \circ \mathcal{A}_s \llbracket F \rrbracket (\rho, x) && (\beta_e) \\
&= \lambda_s x. F && \text{by induction hypothesis}
\end{aligned}$$

- $E \equiv x_i$

$$\begin{aligned}
\text{push}_s \rho \circ \mathcal{A}3 \llbracket x_i \rrbracket \rho &= \text{push}_e \rho \circ \text{access}_i \circ \text{appclos} && \text{with } \rho = (\dots(((),x_n)\dots,x_0) \\
&= \text{push}_s x_i \circ \text{appclos} && \text{access}_i \text{ def.}, (\beta_s), (\beta_e) \\
&= x_i && \text{appclos def.}, (\beta_s) \quad \square
\end{aligned}$$

n. $\mathcal{A}3$ abstraction algorithm

This refinement consists in copying the environment only when building closures. In order to be able to add new bindings after closure opening, a local environment ρ_L is needed. When a closure is built, the concatenation of the two specialized environments ($\rho_{G++\rho_L}$) is copied. The code for variables has now to specify which environment is accessed. Although the transformation scheme remains similar, every rule must be redefined to take into account the two environments.

$$\begin{aligned}
\mathcal{A}3 \llbracket E_1 \circ E_2 \rrbracket \rho_L \rho_G &= \text{dupl2}_e \circ \mathcal{A}3 \llbracket E_1 \rrbracket \rho_L \rho_G \circ \text{swap2}_{se} \circ \mathcal{A}3 \llbracket E_2 \rrbracket \rho_L \rho_G \\
\mathcal{A}3 \llbracket \text{push}_s E \rrbracket \rho_L \rho_G &= \text{Copy2} (\rho_{G++\rho_L}) \circ \text{push}_s (\text{push}_e () \circ \mathcal{A}3 \llbracket E \rrbracket () \rho_{L++\rho_G}) \circ \text{mkclos} \\
\mathcal{A}3 \llbracket \lambda_s x. E \rrbracket \rho_L \rho_G &= \text{mkbind2} \circ \mathcal{A}3 \llbracket E \rrbracket \rho_L (\rho_G, x) \\
\mathcal{A}3 \llbracket x_i \rrbracket (\dots((\rho_L, x_i), x_{i-1}) \dots, x_0) \rho_G &= \text{getlocal} \circ \text{access}_i \circ \text{appclos} \\
\mathcal{A}3 \llbracket x_i \rrbracket \rho_L (\dots((\rho_G, x_i), x_{i-1}) \dots, x_0) &= \text{getglobal} \circ \text{access}_i \circ \text{appclos} \\
\text{with } \text{dupl2}_e &= \lambda_e e_l. \lambda_e e_g. \text{push}_e e_g \circ \text{push}_e e_l \circ \text{push}_e e_g \circ \text{push}_e e_l \\
\text{swap2}_{se} &= \lambda_s x. \lambda_e e_l. \lambda_e e_g. \text{push}_s x \circ \text{push}_e e_g \circ \text{push}_e e_l \\
\text{mkbind2} &= \lambda_e e_l. \lambda_e e_g. \lambda_s x. \text{push}_e e_g \circ \text{push}_s x \circ \text{push}_e e_l \circ \text{mkbind} \\
\text{getlocal} &= \lambda_e e_l. \lambda_e e_g. \text{push}_e e_l && \text{getglobal} = \lambda_e e_l. \lambda_e e_g. \text{push}_e e_g
\end{aligned}$$

Figure 24 Abstraction with Local Environments ($\mathcal{A}3$ Abstraction)

Local environments are not compatible with $\mathcal{V}m$: $\mathcal{A}3 \llbracket \text{grab}_s E \rrbracket$ would generate two different versions of $\mathcal{A}3 \llbracket E \rrbracket$ since E may appear in a closure or may be applied. This code duplication is obviously not realistic.

0. A family of abstractions algorithms

Starting from different properties, a large family of abstractions can be derived from $\mathcal{A}\mathcal{G}$. These transformations introduce indexed combinators (which are generalizations of previous combinators) and use the arity notion.

Definition 19 An expression E of type $\sigma_1 \rightarrow \dots \rightarrow \sigma_n \rightarrow R_i \sigma$ is said to have arity n .

We present here only the **dupl**-less transformation $\mathcal{A}\mathcal{G}_d$ which suppresses the occurrences of **dupl** _{e} in $\mathcal{A}\mathcal{G} \llbracket E_1 \circ E_2 \rrbracket$. Duplications are postponed until really needed (in closure building or opening). $\mathcal{A}\mathcal{G}_d$ is derived from the equation

$$\mathcal{A}\mathcal{G}_d \llbracket E \rrbracket \rho = \mathbf{copy}_n \circ \mathcal{A}\mathcal{G} \llbracket E \rrbracket \rho \quad (n \text{ arity of } E)$$

Note that $\mathbf{copy}_n = \lambda_e e. \lambda_s x_1 \dots \lambda_s x_n. \mathbf{push}_e e \circ \mathbf{push}_s x_n \circ \dots \circ \mathbf{push}_s x_1 \circ \mathbf{push}_e e$ is a generalized form of **dupl** _{e} ($\mathbf{copy}_0 = \mathbf{dupl}_e$). This abstraction algorithm exploits the sequencing encoded in compositions. Instead of saving and restoring the environment (as in $\mathcal{A}\mathcal{G} \llbracket E_1 \circ E_2 \rrbracket$), it is passed to E_1 which may add new bindings but has to remove them before passing the environment to E_2 .

$$\begin{aligned} \mathcal{A}\mathcal{G}_d : \Lambda_s &\rightarrow env \rightarrow \Lambda_e \\ \mathcal{A}\mathcal{G}_d \llbracket E_1 \circ E_2 \rrbracket \rho &= \mathcal{A}\mathcal{G}_d \llbracket E_1 \rrbracket \rho \circ \mathbf{swap}_{se} \circ \mathcal{A}\mathcal{G}_d \llbracket E_2 \rrbracket \rho \\ \mathcal{A}\mathcal{G}_d \llbracket \mathbf{push}_s E \rrbracket \rho &= \mathbf{push}_s (\mathcal{A}\mathcal{G}_d \llbracket E \rrbracket \rho \circ \mathbf{pop}) \circ \mathbf{mkclos}_d \\ \mathcal{A}\mathcal{G}_d \llbracket \lambda_s x. E \rrbracket \rho &= \mathbf{mkbinding} \circ \mathcal{A}\mathcal{G}_d \llbracket E \rrbracket (\rho, x) \circ \mathbf{brkbinding} \\ \mathcal{A}\mathcal{G}_d \llbracket x_i \rrbracket (\dots ((\rho, x_i), x_{i-1}) \dots, x_0) &= \mathbf{copy}_n \circ \mathbf{access}_i \circ \mathbf{appclos} \quad (n \text{ arity of } x_i) \end{aligned}$$

Figure 25 “Dupl-less” abstraction algorithm ($\mathcal{A}\mathcal{G}_d$)

In the first rule, following the evaluation of $\mathcal{A}\mathcal{G}_d \llbracket E_1 \rrbracket$, the unique current environment is threaded to $\mathcal{A}\mathcal{G}_d \llbracket E_2 \rrbracket$ with the help of \mathbf{swap}_{se} . The second rule builds a closure (using \mathbf{mkclos}_d), duplicating the current environment. The abstraction rule adds (using $\mathbf{mkbinding}$) an argument to the environment then removes it (using $\mathbf{brkbinding}$). Finally, the last rule saves the environment (using \mathbf{copy}_n), before calling the closure. We do not give here the definitions of the new combinators \mathbf{pop} , \mathbf{mkclos}_d and $\mathbf{brkbinding}$; they emerge naturally during the derivation process. This transformation can be used with shared or copied environments. It can change the depth of the environment stack needed to reduce an expression by an order of magnitude. For example, if $E \equiv (\dots (x_n \circ x_{n-1}) \dots \circ x_2) \circ x_1$, the depth of the environment stack will be n for $\mathcal{A}\mathcal{G} \llbracket E \rrbracket \rho$ and 1 for $\mathcal{A}\mathcal{G}_d \llbracket E \rrbracket \rho$.

As with the other derived abstractions, the $\mathcal{A}\mathcal{G}_d$ abstraction is correct by construction. To illustrate how $\mathcal{A}\mathcal{G}_d$ is derived from $\mathcal{A}\mathcal{G}$, let us take the rule for compositions

$$\begin{aligned}
\mathcal{A}g_d \llbracket E_1 \circ E_2 \rrbracket \rho &= \mathbf{copy}_n \circ \mathcal{A}g \llbracket E_1 \circ E_2 \rrbracket \rho && (\mathcal{A}g_d \text{ property}) \\
&= \mathbf{copy}_n \circ (\mathbf{dupl}_e \circ \mathcal{A}g \llbracket E_1 \rrbracket \rho \circ \mathbf{swap}_{se} \circ \mathcal{A}g \llbracket E_2 \rrbracket \rho) && (\text{unfolding}) \\
&= \mathbf{copy}_0 \circ \mathcal{A}g \llbracket E_1 \rrbracket \rho \circ \mathbf{swap}_{se} \circ \mathbf{copy}_{n+1} \circ \mathcal{A}g \llbracket E_2 \rrbracket \rho && (\mathbf{copy}_n, \mathbf{dupl}_e, \mathbf{swap}_{se} \text{ definitions}) \\
&= \mathcal{A}g_d \llbracket E_1 \rrbracket \rho \circ \mathbf{swap}_{se} \circ \mathcal{A}g_d \llbracket E_2 \rrbracket \rho && (\text{folding, } E_1 \text{ is 0-ary and } E_2 \text{ is } n+1\text{-ary})
\end{aligned}$$

This technique allows us to derive realistic abstraction algorithms, where indexed combinators can mimic real stack-machine instructions. In order to compare these different options it would be imperative to determine the cost of each indexed combinator. According to their definition and the components instantiation, some indexed combinators have a constant cost. For example \mathbf{copy}_n boils down to \mathbf{dupl} when s and e are distinct components and a combinator $\mathbf{flush}_n = \lambda_s x_1 \dots \lambda_s x_n. \mathbf{push}_s x_n$ would be implemented as a single instruction on a stack machine.

It is as easy to define swap-less ($\mathcal{A}g_s$), mkbind-less ($\mathcal{A}g_b$), mkclos-less ($\mathcal{A}g_m$) variations or any combination of these [17]. Some of these algorithms can be specialized for shared and copied environments; some are suited to a specific choice. Let us mention Tim [20] which uses a mkclos-less variation of $\mathcal{A}c1$ and Tabac [22] which integrates a dupl-less, swap-less, mkbind-less variation of $\mathcal{A}c2$.

p. Recursion

The rewriting rule for \mathbf{Y}_s is

$$\mathbf{push}_s F \circ \mathbf{Y}_s \Rightarrow \mathbf{push}_s(\mathbf{push}_s F \circ \mathbf{Y}_s) \circ F$$

A naive way to compile the β -reduction for the fixpoint operator is to build a closure at each recursive call (a recursive function can have free variables and a closure must be built). This option can be described by the combinator \mathbf{Y}_e with the rewriting rule

$$\begin{aligned}
&\mathbf{push}_e e \circ \mathbf{push}_s F \circ \mathbf{Y}_e \\
&\Rightarrow \mathbf{push}_e e \circ (\mathbf{push}_e e \circ \mathbf{push}_s (\mathbf{push}_s F \circ \mathbf{Y}_e) \circ \mathbf{mkclos}) \circ F
\end{aligned}$$

This solution builds at each call a closure of the function $(\mathbf{push}_e e \circ \mathbf{push}_s (\mathbf{push}_s F \circ \mathbf{Y}_e) \circ \mathbf{mkclos})$ which is added to the current environment. Recursive calls access these closures and execute them using a sequence of code such as $\mathbf{fst}^i \circ \mathbf{snd} \circ \mathbf{appclos}$.

As the same closure (i.e. same code and environment) is built at each recursive call, a first refinement is to build a circular environment. \mathbf{Y}_e must manipulate directly the store to create a cycle. Recall that the source fixpoint operator is of the form $\mathit{letrec} f = E$, the corresponding Λ_s -expression is of the form $\mathbf{push}_s (\lambda_s f.E) \circ \mathbf{Y}_s$ and therefore the Λ_e -expression is of the form $\mathbf{push}_s (\mathbf{mkbind} \circ E) \circ \mathbf{Y}_e$. The rewriting rule of \mathbf{Y}_e becomes

$$\mathbf{push}_e e \circ \mathbf{push}_s (\mathbf{mkbind} \circ E) \circ \mathbf{Y}_e \Rightarrow \mathit{env} \circ E$$

with $env = \mathbf{push}_e e \circ (\mathbf{push}_e env \circ \mathbf{push}_s E \circ \mathbf{mkclos}) \circ \mathbf{mkbind}$

The closure $(\mathbf{push}_e env \circ \mathbf{push}_s E \circ \mathbf{mkclos})$ is built only once for each series of recursive calls (note that the initial \mathbf{mkbind} has been suppressed). The circular environment env of this closure is made of the environment of the recursive function and the closure itself. When accessing the closure, circularity makes the code reinstall the environment env for free.

A second refinement used in environment based machines is to implement recursive calls to statically known functions by a jump to their address. It is sufficient to replace a recursive call $\mathbf{fst}^i \circ \mathbf{snd} \circ \mathbf{appclos}$ by $\mathbf{fst}^i \circ E$. A recursive call is not anymore the evaluation of a closure, but consists in installing the environment (i.e. the free variables) of the function (\mathbf{fst}^i) and calling its code (E). Of course, in order to get a real code machine, this call should be implemented by a jump to a label. With this solution, recursive functions appear in closures only when they are passed as argument.