



Basic Research in Computer Science

A Syntactic Correspondence between Context-Sensitive Calculi and Abstract Machines

Małgorzata Biernacka
Olivier Danvy

BRICS Report Series

RS-05-38

ISSN 0909-0878

December 2005

**Copyright © 2005, Małgorzata Biernacka & Olivier Danvy.
BRICS, Department of Computer Science
University of Aarhus. All rights reserved.**

**Reproduction of all or part of this work
is permitted for educational or research use
on condition that this copyright notice is
included in any copy.**

**See back inner page for a list of recent BRICS Report Series publications.
Copies may be obtained by contacting:**

**BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK-8000 Aarhus C
Denmark
Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk**

**BRICS publications are in general accessible through the World Wide
Web and anonymous FTP through these URLs:**

`http://www.brics.dk`
`ftp://ftp.brics.dk`
This document in subdirectory RS/05/38/

A Syntactic Correspondence between Context-Sensitive Calculi and Abstract Machines

Małgorzata Biernacka and Olivier Danvy

BRICS*

Department of Computer Science

University of Aarhus†

December 27, 2005

Abstract

We present a systematic construction of environment-based abstract machines from context-sensitive calculi of explicit substitutions, and we illustrate it with ten calculi and machines for applicative order with an abort operation, normal order with generalized reduction and call/cc, the lambda-mu-calculus, delimited continuations, stack inspection, proper tail-recursion, and lazy evaluation. Most of the machines already exist but have been obtained independently and are only indirectly related to the corresponding calculi. All of the calculi are new and they make it possible to directly reason about the execution of the corresponding machines.

In connection with the functional correspondence between evaluation functions and abstract machines initiated by Reynolds, the present syntactic correspondence makes it possible to construct reduction-free normalization functions out of reduction-based ones, which was an open problem in the area of normalization by evaluation.

Revised version of BRICS RS-05-22.

*Basic Research in Computer Science (www.brics.dk),
funded by the Danish National Research Foundation.

†IT-parken, Aabogade 34, DK-8200 Aarhus N, Denmark.
Email: {mbiernac,danvy}@brics.dk

Contents

1	Introduction	1
1.1	Calculi and machines	1
1.2	Calculi of explicit substitutions and environment-based machines	2
1.3	Calculi of explicit substitutions for computational effects and environment-based machines	2
1.4	Overview	2
2	Preliminaries	3
2.1	Reduction semantics	3
2.2	Context-sensitive reduction	3
2.3	Our base calculus of closures: $\lambda\hat{\rho}$	4
3	The $\lambda\hat{\rho}\mathcal{A}$-calculus	5
3.1	A reduction semantics for applicative order and abort	6
3.1.1	Decomposition	7
3.1.2	Context-sensitive contraction	7
3.1.3	Plugging	8
3.1.4	One-step reduction	8
3.1.5	Reduction-based evaluation	8
3.2	A pre-abstract machine	8
3.3	A staged abstract machine	10
3.4	An eval/apply abstract machine	11
3.5	Transition compression	11
3.6	An environment machine	12
3.7	Correctness	12
3.8	Conclusion	12
4	The $\lambda\hat{\rho}\mathcal{K}$-calculus	13
4.1	The language of $\lambda\hat{\rho}\mathcal{K}$	13
4.2	Notion of context-sensitive reduction	13
4.3	Krivine's machine	14
4.4	Formal correspondence	14
5	The $\lambda\hat{\rho}\mu$-calculus	14
5.1	The language of $\lambda\hat{\rho}\mu$	15
5.2	Notion of context-sensitive reduction	15
5.3	An eval/apply abstract machine	15
5.4	Formal correspondence	16
6	Delimited continuations	16
6.1	The $\lambda\hat{\rho}\mathcal{S}$ -calculus	17
6.1.1	The language of $\lambda\hat{\rho}\mathcal{S}$	17
6.1.2	The eval/apply/meta-apply abstract machine	17
6.1.3	Notion of context-sensitive reduction	18
6.1.4	Formal correspondence	18
6.1.5	The CPS hierarchy	18
6.2	The $\lambda\hat{\rho}\mathcal{F}$ -calculus	19
6.2.1	The language of $\lambda\hat{\rho}\mathcal{F}$	19

6.2.2	Notion of context-sensitive reduction	19
6.2.3	The eval/apply abstract machine	20
6.2.4	Formal correspondence	20
6.2.5	A hierarchy of control delimiters	20
6.3	Conclusion	20
7	Stack inspection	22
7.1	The $\lambda_{\widehat{\rho}_{\text{sec}}}$ -calculus	22
7.1.1	The language of $\lambda_{\widehat{\rho}_{\text{sec}}}$	22
7.1.2	Notion of context-sensitive reduction	22
7.1.3	An eval/apply abstract machine	23
7.1.4	Formal correspondence	23
7.2	Properly tail-recursive stack inspection	23
7.2.1	The storeless cm machine	24
7.2.2	The underlying calculus $\lambda_{\widehat{\rho}_{\text{sec}}}^{\text{cm}}$	24
7.3	State-based properly tail-recursive stack inspection	25
7.3.1	The unzipped storeless cm machine	25
7.3.2	The language of $\lambda_{\widehat{\rho}_{\text{sec}}}^{\text{ncm}}$	26
7.3.3	Notion of context-sensitive reduction	26
7.3.4	Formal correspondence	27
7.4	Conclusion	27
8	A calculus for proper tail-recursion	27
8.1	A simplified version of Clinger’s abstract machine	27
8.2	The language of $\lambda_{\widehat{\rho}_{\text{ptr}}}$	28
8.3	Notion of context-sensitive reduction	28
8.4	Formal correspondence	28
9	A lazy calculus of closures	29
9.1	The language of $\lambda_{\widehat{\rho}_1}$	29
9.2	Notion of context-sensitive reduction	29
9.3	An eval/apply abstract machine	29
9.4	Formal correspondence	30
10	Conclusion	30

1 Introduction

How does one construct a new semantic artifact?

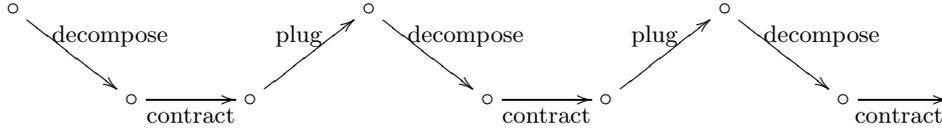
1.1 Calculi and machines

Sixty-five years ago, the λ -calculus was introduced [15]. Forty-five years ago, its expressive power was observed to be relevant for computing [67,81]. Forty years ago, a first abstract machine for the λ -calculus was introduced [61]. Thirty years ago, calculi and abstract machines were formally connected [70]. Twenty years ago, a calculus format—reduction semantics—with an explicit representation of reduction contexts was introduced [40]. Today calculi and abstract machines are standard tools to study programming languages. Given a calculus, it is by now a standard activity to design a corresponding abstract machine and to prove its correctness [42].

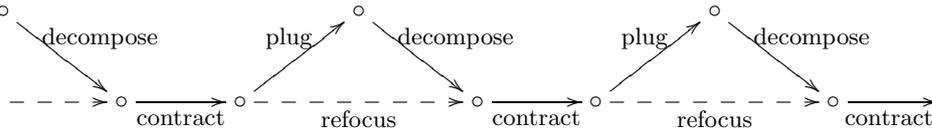
From calculus to machine by refocusing and transition compression: Recently, Danvy and Nielsen have pointed out that the reduction strategy for a calculus actually determines the structure of the corresponding machine [36]. They present a method for constructing an abstract machine out of a reduction semantics satisfying the unique-decomposition property. In such a reduction semantics, a non-value term is reduced by

1. decomposing it (uniquely) into a redex and its context,
2. contracting the redex, and
3. plugging the contractum into the reduction context,

yielding a new term. A reduction-based evaluation function is defined using the reflexive and transitive closure of the one-step reduction function:



Danvy and Nielsen have observed that the intermediate terms, in the composition of plug and decompose, could be avoided by fusing the composition into a ‘refocus’ function:



The resulting ‘refocused’ evaluation function is defined as the reflexive and transitive closure of refocusing and contraction.

The refocus function takes the form of a state-transition function, i.e., an abstract machine. The refocused evaluation function therefore also takes the form of an abstract machine. Compressing its intermediate transitions (i.e., short-circuiting them) yields abstract machines that are often independently known: for example, for the pure λ -calculus with normal order, the resulting abstract machine is a substitution-based version of the Krivine machine (i.e., a push/enter machine); for the pure λ -calculus with applicative order, the resulting abstract machine is Felleisen et al.’s CK machine (i.e., an eval/apply machine). Refocusing has also been applied to the term language of the free monoid, yielding a reduction-free normalization function [26], and to context-based CPS transformations, improving them from quadratic time to operating in one pass [36].

1.2 Calculi of explicit substitutions and environment-based machines

Twenty years ago, Curien observed that while most calculi use actual substitutions, most implementations use closures and environments [22]. He then developed a calculus of closures, $\lambda\rho$ [23], thereby launching the study of explicit substitutions [1, 24, 51, 65, 75].

From calculus to machine by refocusing, transition compression, and closure unfolding: Recently, we have applied the refocusing method to $\lambda\hat{\rho}$, a minimal extension of $\lambda\rho$ where one can express single-step computations; we added an unfolding step to make the machine operate not on a closure, but on a term and its environment [8]. We have shown how $\lambda\hat{\rho}$ with left-to-right applicative order directly corresponds to the CEK machine [43], how $\lambda\hat{\rho}$ with normal order directly corresponds to the Krivine machine [20, 23], how $\lambda\hat{\rho}$ with generalized reduction directly corresponds to the original version of Krivine’s machine [59], and how $\lambda\hat{\rho}$ with right-to-left applicative order directly corresponds to the ZINC abstract machine [64]. All of these machines are environment-based and use closures.

1.3 Calculi of explicit substitutions for computational effects and environment-based machines

Twenty years ago, Felleisen introduced reduction semantics—a version of small-step operational semantics with an explicit representation of reduction contexts—in order to provide calculi for control and state [40, 43]. In these calculi, reduction rules are not oblivious to their reduction context; on the contrary, they are context sensitive in that the context takes part in some reduction steps, e.g., for call/cc. Reduction semantics are in wide use today, e.g., to study the security technique of stack inspection [16, 47, 71].

From calculus to machine by refocusing, transition compression, and closure unfolding: In this article, we apply the refocusing method to context-sensitive extensions of $\lambda\hat{\rho}$ accounting for a variety of computational effects. We present ten calculi of closures and the corresponding environment-based machines. What is significant here is that each machine is mechanically derived from the corresponding calculus (instead of designed and then proved correct) and also that each machine directly corresponds to this calculus (instead of indirectly via an ‘unload’ function at the end of each run [70] or via a compilation / decompilation scheme in the course of execution [51]).

1.4 Overview

We successively consider call by name: Krivine’s machine with call/cc (Section 4) and the $\lambda\mu$ -calculus (Section 5); call by value: static and dynamic delimited continuations (Section 6), stack inspection (Section 7), and proper tail-recursion (Section 8); and call by need (Section 9). Towards this end, we first present the $\lambda\hat{\rho}$ -calculus and the notion of context-sensitive reduction (Section 2) and then a detailed walkthrough for the $\lambda\hat{\rho}$ -calculus with an abort operation (Section 3). The subsequent repetitiveness of Sections 4 to 9 (language, notion of context-sensitive reduction, abstract machine, and formal correspondence) is a deliberate feature, not an inadvertent presentational bug: we apply the same simple method to many situations that have been separately studied so far. Each of these sections can therefore be read independently of the others.

2 Preliminaries

2.1 Reduction semantics

A reduction semantics [40, 42] consists of a grammar of terms from a source language, syntactic notions of value and redex, a collection of contraction rules, and a reduction strategy. This reduction strategy is embodied in a grammar of reduction contexts (terms with a hole as induced by the compatibility rules) and a plug function mapping a term and a context into a new term. One-step reduction of a non-value term consists in

1. decomposing the term into a redex and a reduction context,
2. contracting the redex, and
3. plugging the contractum in the reduction context.

In some reduction semantics, non-value terms are uniquely decomposed into a redex and a context. Decomposition can then be implemented as a function mapping a non-value term to a redex and a reduction context. Danvy and Nielsen have shown that together with the unique decomposition property, the following property of a reduction semantics is sufficient to define a decomposition function by induction on terms and reduction contexts [36, Figure 2, page 8]:

Property 1. *For each syntactic construct building a term out of n subterms, there is a number $0 \leq i \leq n$ and a fixed traversal order of subterms encoded in the grammar of the reduction contexts for this construct such that the holes of these contexts are precisely the positions of the i subterms. Furthermore, a term with all the chosen i subterms reduced to values is either a value or a potential redex (i.e., an actual redex or a “stuck term”) but not both.*

If the redexes do not overlap, the contraction rules can be implemented as a function.

2.2 Context-sensitive reduction

Traditional specifications of one-step reduction as the compatible closure of a notion of reduction provide a *local* characterization of a computation step in the form of a redex.¹ This local characterization is not fit for non-local reductions such as one involving a control operator capturing all its surrounding context in one step, or a global state. For these, one needs a notion of *context-sensitive* reduction, i.e., a binary relation defined both on redexes and on their reduction context instead of only on redexes [40].

This relation is given by context-sensitive contraction rules of the form $\langle r, C \rangle \rightarrow \langle c', C' \rangle$, where $\langle r, C \rangle$ denotes the decomposition of a program into a potential redex r and its context C .

A one-step reduction relation for a given notion of context-sensitive reduction is defined as follows: A program p reduces in one step to p' if decomposing p yields $\langle r, C \rangle$, reducing $\langle r, C \rangle$ yields $\langle c', C' \rangle$, and plugging c' into C' yields p' .

A context-sensitive reduction implicitly assumes a decomposition of the entire program, and therefore it cannot be used locally. One way to recover compatibility in the context-sensitive setting is to add explicit local control delimiters to the language (see Section 6 for an illustration). For a language without explicit control delimiters (as the $\lambda\hat{\rho}$ -calculus with call/cc), there is an implicit global control delimiter around the program [41].

¹For example, a potential redex in the λ -calculus is the application of a value to a term. If the value is a λ -abstraction, the potential redex is an actual one and it can be β -reduced. If no reduction rule applies, the potential redex is not an actual one and the program is stuck [70].

2.3 Our base calculus of closures: $\lambda\hat{\rho}$

Since Landin [61], most abstract machines implementing variants and extensions of the λ -calculus use closures and environments, and the substitution of terms for free variables is thus delayed until a variable is reached in the evaluation process. This implementation technique has motivated the study of calculi of explicit substitutions [1, 23, 75] to mediate between the traditional abstract specifications of the λ -calculus and its traditional concrete implementations [51].

To derive an abstract machine for evaluating λ -terms, a *weak* calculus of explicit substitutions suffices. The first (and simplest) of such calculi was Curien's calculus of closures $\lambda\rho$ [23]. Although this calculus is not expressive enough to model full normalization, it is suitable for evaluating a λ -term, i.e., to produce the corresponding weak head normal form. Its operational semantics is specified using multi-step reductions, but its syntax is too restrictive to allow single-step computations, which is what we need to apply the refocusing method. For this reason, in our earlier work [8], we have proposed a minimal extension of $\lambda\rho$ with one-step reduction rules, the $\lambda\hat{\rho}$ -calculus.

The language of $\lambda\hat{\rho}$ is as follows:

$$\begin{array}{ll} \text{(terms)} & t ::= i \mid \lambda t \mid t t \\ \text{(closures)} & c ::= t[s] \mid c c \\ \text{(substitutions)} & s ::= \bullet \mid c \cdot s \end{array}$$

(For comparison, $\lambda\rho$ does not have the $c c$ production.)

We use de Bruijn indices for variables in a term ($i \geq 1$). A closure is a term equipped with a substitution, i.e., a list of closures to be substituted for free variables in the term. Programs are closures of the form $t[\bullet]$ where t does not contain free variables.

The notion of reduction in the $\lambda\hat{\rho}$ -calculus is given by the following rules:

$$\begin{array}{ll} \text{(Var)} & i[c_1 \cdots c_j] \xrightarrow{\hat{\rho}} c_i \quad \text{if } i \leq j \\ \text{(Beta)} & ((\lambda t)[s]) c \xrightarrow{\hat{\rho}} t[c \cdot s] \\ \text{(Prop)} & (t_0 t_1)[s] \xrightarrow{\hat{\rho}} (t_0[s]) (t_1[s]) \end{array}$$

We write $s(i)$ for the i th element of the substitution s considered as a list. (So $[c_1 \cdots c_j](i) = c_i$ if $1 \leq i \leq j$.)

Finally, the one-step reduction relation (i.e., the compatible closure of the notion of reduction) extends the notion of reduction with the following rules:

$$\begin{array}{ll} \text{(L-Comp)} & \frac{c_0 \xrightarrow{\hat{\rho}} c'_0}{c_0 c_1 \xrightarrow{\hat{\rho}} c'_0 c_1} \\ \text{(R-Comp)} & \frac{c_1 \xrightarrow{\hat{\rho}} c'_1}{c_0 c_1 \xrightarrow{\hat{\rho}} c_0 c'_1} \\ \text{(Sub)} & \frac{c_i \xrightarrow{\hat{\rho}} c'_i}{t[c_1 \cdots c_i \cdots c_j] \xrightarrow{\hat{\rho}} t[c_1 \cdots c'_i \cdots c_j]} \quad \text{for } i \leq j \end{array}$$

Specific, deterministic reduction strategies can be obtained by restricting the compatibility rules. In the following sections, we consider two such strategies:

1. the normal-order strategy obtained by discarding the (Sub) and (R-Comp) rules, where the grammars of values, redexes and reduction contexts are specified as follows:

$$\begin{array}{ll}
\text{(values)} & v ::= (\lambda t)[s] \\
\text{(redexes)} & r ::= v c \mid i[s] \mid (t_0 t_1)[s] \\
\text{(reduction contexts)} & C ::= [] \mid C[[] c],
\end{array}$$

and

2. the left-to-right applicative-order strategy obtained in the usual way by discarding the (Sub) rule and restricting the (Beta) and (R-Comp) rules, where the grammars of values, redexes and reduction contexts are specified as follows:

$$\begin{array}{ll}
\text{(values)} & v ::= (\lambda t)[s] \\
\text{(redexes)} & r ::= v v \mid i[s] \mid (t_0 t_1)[s] \\
\text{(reduction contexts)} & C ::= [] \mid C[[] c] \mid C[v []]
\end{array}$$

All of the calculi presented in this article extend the $\lambda\hat{\rho}$ -calculus. For each of them, we define a suitable notion of reduction, denoted \rightarrow_X , where X is a subscript identifying a particular calculus. For each of them, we then define a one-step reduction relation as the composition of: decomposing a non-value closure into a redex and a reduction context, contracting a (context-sensitive) redex, and then plugging the resulting closure into the resulting context. Finally, we define the evaluation relation (denoted \rightarrow_X^*) using the reflexive, transitive closure of one-step reduction, i.e., we say that c evaluates to c' if $c \rightarrow_X^* c'$ and c' is a value closure. We define the convertibility relation between closures as the smallest equivalence relation containing \rightarrow_X^* . If two closures c and c' are convertible, they behave similarly under evaluation (i.e., either they both evaluate to the same value, or they both diverge).

3 The $\lambda\hat{\rho}\mathcal{A}$ -calculus

As an illustration, we present a detailed and systematic derivation of an abstract machine for call-by-value evaluation in the λ -calculus with an abort operation, starting from the specification of the applicative-order reduction strategy in the $\lambda\hat{\rho}$ -calculus with an abort operation. We follow the steps outlined by Biernacka, Danvy, and Nielsen [8, 36]:

Section 3.1: We specify the applicative-order reduction strategy in the form of a reduction semantics, i.e., with a one-step reduction function specified as decomposing a non-value term into a reduction context and a redex, contracting this redex, and plugging the contractum into the context. As is traditional, we also specify evaluation as the transitive closure of one-step reduction.

Section 3.2: We replace the combination of plugging and decomposition by a refocus function that iteratively goes from redex site to redex site in the reduction sequence. The resulting ‘refocused’ evaluation function is the transitive closure of the refocus function and takes the form of a ‘pre-abstract machine.’

Section 3.3: We merge the definitions of the transitive closure and the refocus function into a ‘staged abstract machine’ that implements the reduction rules and the compatibility rules of the $\lambda\hat{\rho}$ -calculus with two separate transition functions.

Section 3.4: We inline the transition function implementing the reduction rules. The result is an eval/apply abstract machine consisting of an ‘eval’ transition function dispatching on closures and an ‘apply’ transition function dispatching on contexts.

Section 3.5: Observing that in a reduction sequence, an (App) reduction step is always followed by a decomposition step, we coalesce these two steps into one. Observing that in a reduction sequence, an (Abort) reduction step is always followed by a decomposition step, we coalesce these two steps into one. This shortcut makes the resulting abstract machine dispatch on terms rather than on closures, and enables the following step.

Section 3.6: We unfold the data type of closures, making the abstract machine operate over two components—a term and a substitution—instead of over one—a closure. The substitution component is the traditional environment of environment machines, and the resulting machine is an environment machine. This machine coincides with the CEK machine with an abort operator [42].

In Section 3.7, we state the correctness of the resulting CEK machine with respect to evaluation in the $\lambda\hat{\rho}$ -calculus with an abort operation.

3.1 A reduction semantics for applicative order and abort

A reduction semantics for applicative-order reduction in the $\lambda\hat{\rho}\mathcal{A}$ -calculus builds on the applicative-order strategy presented in Section 2.3. The grammar of terms and closures contains additional productions for the abort operation:

$$\begin{array}{ll} \text{(Term)} & t ::= \dots \mid \mathcal{A} t \\ \text{(Closure)} & c ::= \dots \mid \mathcal{A} c \end{array}$$

The reduction semantics is context-sensitive and its contraction rules contain the contractions of the $\lambda\hat{\rho}$ -calculus (here stated in the context-sensitive form):

$$\begin{array}{ll} \text{(Var)} & \langle i[c_1 \dots c_j], C \rangle \rightarrow_{\mathcal{A}} \langle c_i, C \rangle \quad \text{if } i \leq j \\ \text{(Beta)} & \langle ((\lambda t)[s]) c, C \rangle \rightarrow_{\mathcal{A}} \langle t[c \cdot s], C \rangle \\ \text{(Prop)} & \langle (t_0 t_1)[s], C \rangle \rightarrow_{\mathcal{A}} \langle (t_0[s]) (t_1[s]), C \rangle \end{array}$$

as well as two new contractions for the abort operation:

$$\begin{array}{ll} \text{(Prop}_{\mathcal{A}}) & \langle (\mathcal{A} t)[s], C \rangle \rightarrow_{\mathcal{A}} \langle \mathcal{A} (t[s]), C \rangle \\ \text{(Abort)} & \langle \mathcal{A} v, C \rangle \rightarrow_{\mathcal{A}} \langle v, [] \rangle \end{array}$$

Finally, the grammar of reduction contexts reads as follows:

$$\text{(Context)} \quad C ::= [] \mid C[[]] c \mid C[v []] \mid C[\mathcal{A} []]$$

This reduction semantics satisfies the conditions stated in Section 2.1. We can therefore define the following three functions:

$$\begin{array}{l} \text{decompose} : \text{Closure} \rightarrow \text{Value} + (\text{Redex} \times \text{Context}) \\ \text{contract} : \text{Redex} \times \text{Context} \rightarrow \text{Closure} \times \text{Context} \\ \text{plug} : \text{Closure} \times \text{Context} \rightarrow \text{Closure} \end{array}$$

3.1.1 Decomposition

We define `decompose` as a transition function over closures and reduction contexts:

$$\begin{aligned}
& \text{decompose} : \text{Closure} \rightarrow \text{Value} + (\text{Redex} \times \text{Context}) \\
& \text{decompose } c = \text{decompose}'(c, []) \\
\\
& \text{decompose}' : \text{Closure} \times \text{Context} \rightarrow \text{Value} + (\text{Redex} \times \text{Context}) \\
& \text{decompose}'(i[s], C) = (i[s], C) \\
& \text{decompose}'((\lambda t)[s], C) = \text{decompose}'_{\text{aux}}(C, (\lambda t)[s]) \\
& \text{decompose}'((t_0 t_1)[s], C) = ((t_0 t_1)[s], C) \\
& \text{decompose}'(c_0 c_1, C) = \text{decompose}'(c_0, C[[] c_1]) \\
& \text{decompose}'((\mathcal{A} t)[s], C) = ((\mathcal{A} t)[s], C) \\
& \text{decompose}'(\mathcal{A} c, C) = \text{decompose}'(c, C[\mathcal{A} []]) \\
\\
& \text{decompose}'_{\text{aux}} : \text{Context} \times \text{Value} \rightarrow \text{Value} + (\text{Redex} \times \text{Context}) \\
& \text{decompose}'_{\text{aux}}([], v) = v \\
& \text{decompose}'_{\text{aux}}(C[[] c], v) = \text{decompose}'(c, C[v []]) \\
& \text{decompose}'_{\text{aux}}(C[v' []], v) = (v' v, C) \\
& \text{decompose}'_{\text{aux}}(C[\mathcal{A} []], v) = (\mathcal{A} v, C)
\end{aligned}$$

The main decomposition function, `decompose`, uses two auxiliary transition functions that work according to Property 1:

- `decompose'` is passed a closure and a reduction context. It dispatches on the closure and iteratively builds the reduction context:
 - if the current closure is a value, then `decompose'`_{aux} is called to inspect the context;
 - if the current closure is a redex, then a decomposition is found; and
 - otherwise, a subclosure of the current closure is chosen to be visited in a new context.
- `decompose'`_{aux} is passed a reduction context and a value. It dispatches on the reduction context:
 - if the current context is empty, then the value is the result of the function;
 - if the top constructor of the context is that of a function application, the actual parameter is decomposed in a new context; and
 - otherwise, a redex has been found.

The decomposition function is total.

3.1.2 Context-sensitive contraction

We define `contract` by cases, as a straightforward implementation of the contraction rules:

$$\begin{aligned}
& \text{contract} : \text{Redex} \times \text{Context} \rightarrow \text{Closure} \times \text{Context} \\
& \text{contract}(i[v_1 \cdots v_m], C) = (v_i, C) \\
& \text{contract}((t_0 t_1)[s], C) = ((t_0[s]) (t_1[s]), C) \\
& \text{contract}(((\lambda t)[s]) v, C) = (t[v \cdot s], C) \\
& \text{contract}((\mathcal{A} t)[s], C) = (\mathcal{A} (t[s]), C) \\
& \text{contract}(\mathcal{A} v, C) = (v, [])
\end{aligned}$$

In general, the contraction function is partial because of stuck terms.

3.1.3 Plugging

We define `plug` by structural induction over the reduction context. It iteratively peels off the context and thus also takes the form of a transition function:

$$\begin{aligned} \text{plug} &: \text{Closure} \times \text{Context} \rightarrow \text{Closure} \\ \text{plug}(c, []) &= c \\ \text{plug}(c_0, C[[] c_1]) &= \text{plug}(c_0 c_1, C) \\ \text{plug}(c_1, C[c_0 []]) &= \text{plug}(c_0 c_1, C) \\ \text{plug}(c, C[\mathcal{A} []]) &= \text{plug}(\mathcal{A} c, C) \end{aligned}$$

The plugging function is total.

3.1.4 One-step reduction

Given these three functions, we can define the following one-step reduction function that tests whether a closure is a value or can be decomposed into a redex and a context, that contracts this redex together with its context, and that plugs the contractum in the resulting context:

$$\begin{aligned} \text{reduce} &: \text{Closure} \rightarrow \text{Closure} \\ \text{reduce } c &= \text{case } \text{decompose } c \\ &\quad \text{of } \quad v \Rightarrow v \\ &\quad \mid (r, C) \Rightarrow \text{plug}(c', C') \quad \text{where } (c', C') = \text{contract}(r, C) \end{aligned}$$

In general, the one-step reduction function is partial because of the contraction function.

The following proposition is a consequence of the unique-decomposition property.

Proposition 1. *For any non-value closure c and for any closure c' , $c \rightarrow_{\mathcal{A}} c' \Leftrightarrow \text{reduce } c = c'$.*

3.1.5 Reduction-based evaluation

Finally, we can define evaluation as the reflexive, transitive closure of one-step reduction. For simplicity, we use `decompose` to test whether a value has been reached:

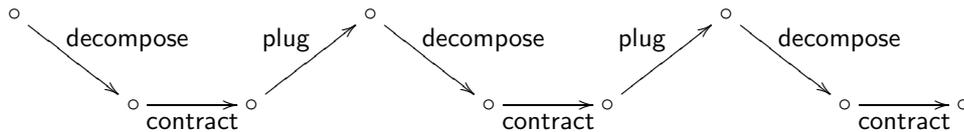
$$\begin{aligned} \text{iterate} &: \text{Value} + (\text{Redex} \times \text{Context}) \rightarrow \text{Value} \\ \text{iterate } v &= v \\ \text{iterate}(r, C) &= \text{iterate}(\text{decompose}(\text{plug}(c', C'))) \\ &\quad \text{where } (c', C') = \text{contract}(r, C) \\ \\ \text{evaluate} &: \text{Term} \rightarrow \text{Value} \\ \text{evaluate } t &= \text{iterate}(\text{decompose}(t[\bullet])) \end{aligned}$$

This evaluation function is partial because of the one-step reduction function and also because a reduction sequence might not terminate.

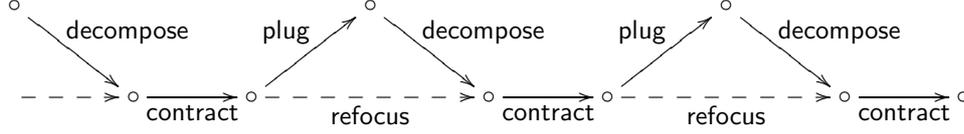
Proposition 2. *For any closed term t and any value v , $t[\bullet] \rightarrow_{\mathcal{A}}^* v \Leftrightarrow \text{evaluate } t = v$.*

3.2 A pre-abstract machine

The reduction sequence implemented by evaluation can be depicted as follows:



At each step, an intermediate term is constructed by the function `plug`; it is then immediately decomposed by the subsequent call to `decompose`. In earlier work [36], Danvy and Nielsen pointed out that the composition of `plug` and `decompose` (which are total) could be replaced by a more efficient function, `refocus` (which is also total), that would directly go from redex site to redex site in the reduction sequence:



The essence of refocusing for a reduction semantics satisfying the unique decomposition property is captured in the following proposition:

Proposition 3 (Danvy and Nielsen [26,36]). *For any closure c and reduction context C ,*

$$\text{decompose}(\text{plug}(c, C)) = \text{decompose}'(c, C)$$

In words: refocusing amounts to continuing the decomposition of the given contractum in the given context.

The definition of the `refocus` function is therefore a clone of that of `decompose'`. In particular, it involves an auxiliary function `refocusaux` and takes the form of two state-transition functions, i.e., of an abstract machine:

$$\text{refocus} : \text{Closure} \times \text{Context} \rightarrow \text{Value} + (\text{Redex} \times \text{Context})$$

$$\text{refocus}(i[s], C) = (i[s], C)$$

$$\text{refocus}((\lambda t)[s], C) = \text{refocus}_{\text{aux}}(C, (\lambda t)[s])$$

$$\text{refocus}((t_0 t_1)[s], C) = ((t_0 t_1)[s], C)$$

$$\text{refocus}(c_0 c_1, C) = \text{refocus}(c_0, C[[] c_1])$$

$$\text{refocus}((\mathcal{A} t)[s], C) = ((\mathcal{A} t)[s], C)$$

$$\text{refocus}(\mathcal{A} c, C) = \text{refocus}(c, C[\mathcal{A} []])$$

$$\text{refocus}_{\text{aux}} : \text{Context} \times \text{Value} \rightarrow \text{Value} + (\text{Redex} \times \text{Context})$$

$$\text{refocus}_{\text{aux}}([], v) = v$$

$$\text{refocus}_{\text{aux}}(C[[] c], v) = \text{refocus}(c, C[v []])$$

$$\text{refocus}_{\text{aux}}(C[v' []], v) = (v' v, C)$$

$$\text{refocus}_{\text{aux}}(C[\mathcal{A} []], v) = (\mathcal{A} v, C)$$

In this abstract machine, the configurations are pairs containing a closure and a context; the final transitions are specified by the first, third, and fifth clauses of `refocus` and by the first, third, and fourth clauses of `refocusaux`; and the initial transition is specified by two clauses of the corresponding 'refocused' evaluation function, which reads as follows:

$$\text{iterate} : \text{Value} + (\text{Redex} \times \text{Context}) \rightarrow \text{Value}$$

$$\text{iterate } v = v$$

$$\text{iterate}(r, C) = \text{iterate}(\text{refocus}(r', C'))$$

$$\text{where } (r', C') = \text{contract}(r, C)$$

$$\text{evaluate} : \text{Term} \rightarrow \text{Value}$$

$$\text{evaluate } t = \text{iterate}(\text{refocus}(t[\bullet], []))$$

(For the initial call to `iterate`, we have exploited the double equality $\text{decompose } (t[\bullet]) = \text{decompose } (\text{plug } (t[\bullet], [])) = \text{refocus } (t[\bullet], [])$.)

This evaluation function computes the transitive closure of `refocus` using the auxiliary function `iterate` as a trampoline [48]. Due to the non-tail call to `refocus` in `iterate`, we refer to this evaluation function as a ‘pre-abstract machine.’

3.3 A staged abstract machine

To transform the pre-abstract machine into an abstract machine, we distribute the calls to `iterate` from the definitions of `evaluate` and of `iterate` to the definitions of `refocus` and `refocusaux`:

$$\begin{aligned}
&\text{evaluate} : \text{Term} \rightarrow \text{Value} \\
&\text{evaluate } t = \text{refocus } (t[\bullet], []) \\
\\
&\text{iterate} : \text{Value} + (\text{Redex} \times \text{Context}) \rightarrow \text{Value} \\
&\text{iterate } v = v \\
&\text{iterate } (r, C) = \text{refocus } (c', C') \\
&\quad \text{where } (c', C') = \text{contract } (r, C) \\
\\
&\text{refocus} : \text{Closure} \times \text{Context} \rightarrow \text{Value} \\
&\text{refocus } (i[s], C) = \text{iterate } (i[s], C) \\
&\text{refocus } ((\lambda t)[s], C) = \text{refocus}_{\text{aux}} (C, (\lambda t)[s]) \\
&\text{refocus } ((t_0 t_1)[s], C) = \text{iterate } ((t_0 t_1)[s], C) \\
&\text{refocus } (c_0 c_1, C) = \text{refocus } (c_0, C[[] c_1]) \\
&\text{refocus } ((\mathcal{A} t)[s], C) = \text{iterate } ((\mathcal{A} t)[s], C) \\
&\text{refocus } (\mathcal{A} c, C) = \text{refocus } (c, C[\mathcal{A} []]) \\
\\
&\text{refocus}_{\text{aux}} : \text{Context} \times \text{Value} \rightarrow \text{Value} \\
&\text{refocus}_{\text{aux}} ([], v) = \text{iterate } v \\
&\text{refocus}_{\text{aux}} (C[[] c], v) = \text{refocus } (c, C[v []]) \\
&\text{refocus}_{\text{aux}} (C[v' []], v) = \text{iterate } (v' v, C) \\
&\text{refocus}_{\text{aux}} (C[\mathcal{A} []], v) = \text{iterate } (\mathcal{A} v, C)
\end{aligned}$$

The resulting definitions of `evaluate`, `iterate`, `refocus`, and `refocusaux` are that of four mutually recursive transition functions that form an abstract machine. In this abstract machine, the configurations are pairs of a closure and a context, the initial transition is specified by `evaluate`, and the final transition in the first clause of `iterate`. The compatibility rules are implemented by `refocus` and `refocusaux`, and the reduction rules by the call to `contract` in the second clause of `iterate`. We can make this last point even more manifest by inlining `contract` in the definition of `iterate`:

$$\begin{aligned}
&\text{iterate} : \text{Value} + (\text{Redex} \times \text{Context}) \rightarrow \text{Value} \\
&\text{iterate } v = v \\
&\text{iterate } (i[v_1 \cdots v_m], C) = \text{refocus } (v_i, C) \\
&\text{iterate } ((t_0 t_1)[s], C) = \text{refocus } ((t_0[s]) (t_1[s]), C) \\
&\text{iterate } (((\lambda t)[s]) v, C) = \text{refocus } (t[v \cdot s], C) \\
&\text{iterate } ((\mathcal{A} t)[s], C) = \text{refocus } (\mathcal{A} (t[s]), C) \\
&\text{iterate } (\mathcal{A} v, C) = \text{refocus } (v, [])
\end{aligned}$$

By construction, the machine therefore separately implements the reduction rules (with `iterate`) and the compatibility rules (with `refocus` and `refocusaux`); for this reason, we refer to it as a ‘staged abstract machine’ [51].

3.4 An eval/apply abstract machine

As already observed by Danvy and Nielsen in their work on refocusing, inlining `iterate` yields an eval/apply abstract machine [66]. Inlining the calls to `iterate` in the staged abstract machine yields the following eval/apply machine, where `refocus` (the ‘eval’ transition function) dispatches on closures and `refocusaux` (the ‘apply’ function) dispatches on contexts:

$$\begin{aligned}
& \text{evaluate} : \text{Term} \rightarrow \text{Value} \\
& \text{evaluate } t = \text{refocus } (t[\bullet], []) \\
\\
& \text{refocus} : \text{Closure} \times \text{Context} \rightarrow \text{Value} \\
& \text{refocus } (i[v_1 \cdots v_m], C) = \text{refocus } (v_i, C) \\
& \text{refocus } ((\lambda t)[s], C) = \text{refocus}_{\text{aux}} (C, (\lambda t)[s]) \\
& \text{refocus } ((t_0 t_1)[s], C) = \text{refocus } ((t_0[s]) (t_1[s]), C) \\
& \text{refocus } (c_0 c_1, C) = \text{refocus } (c_0, C[[] c_1]) \\
& \text{refocus } ((\mathcal{A} t)[s], C) = \text{refocus } (\mathcal{A} (t[s]), C) \\
& \text{refocus } (\mathcal{A} c, C) = \text{refocus } (c, C[\mathcal{A} []]) \\
\\
& \text{refocus}_{\text{aux}} : \text{Context} \times \text{Value} \rightarrow \text{Value} \\
& \text{refocus}_{\text{aux}} ([], v) = v \\
& \text{refocus}_{\text{aux}} (C[[] c], v) = \text{refocus } (c, C[v []]) \\
& \text{refocus}_{\text{aux}} (C[(\lambda t)[s] []], v) = \text{refocus } (t[v \cdot s], C) \\
& \text{refocus}_{\text{aux}} (C[\mathcal{A} []], v) = \text{refocus } (v, [])
\end{aligned}$$

3.5 Transition compression

In three cases, the eval/apply machine of Section 3.4 yields a configuration that uniquely determines the next one. Let us shortcut these “corridor transitions:”

- The machine only produces an application of closures through an (App) reduction step (third clause of `refocus`). We observe that in a reduction sequence, an (App) reduction step is always followed by a decomposition step (fourth clause of `refocus`). As a shortcut, we coalesce the two consecutive transitions into one, replacing the third and fourth clauses of `refocus` with the following one:

$$\text{refocus } ((t_0 t_1)[s], C) = \text{refocus } (t_0[s], C[[] (t_1[s])])$$

- The machine only produces an application of \mathcal{A} through an (Abort) reduction step (second-to-last clause of `refocus`). We observe that in a reduction sequence, an (Abort) reduction step is always followed by a decomposition step (last clause of `refocus`). As a shortcut, we coalesce the two consecutive transitions into one, replacing the last two clauses of `refocus` with the following one:

$$\text{refocus } (\mathcal{A} (t[s]), C) = \text{refocus } (t[s], C[\mathcal{A} []])$$

- The configuration `refocus (v, [])` always yields the configuration `refocusaux ([], v)`, which triggers a final transition. As a shortcut, we coalesce these three consecutive transitions into one final transition, replacing the last clause of `refocusaux` with the following one:

$$\text{refocus}_{\text{aux}} (C[\mathcal{A} []], v) = v$$

3.6 An environment machine

Because of the two first compressions of Section 3.5, `refocus` is now defined by structural induction over terms instead of by structural induction over closures. Its type can therefore be refined to $(\text{Term} \times \text{Substitution}) \times \text{Context} \rightarrow \text{Value} + (\text{Redex} \times \text{Context})$, in effect unfolding closures into a term and a substitution. We thus replace each closure in the compressed definition of `evaluate` and `refocus` by its unfolding. Flattening $(\text{Term} \times \text{Substitution}) \times \text{Context}$ into $\text{Term} \times \text{Substitution} \times \text{Context}$ yields the following abstract machine:

$$\begin{aligned}
& \text{(Value)} && v ::= (\lambda t, s) \\
& \text{(Context)} && C ::= [] \mid C[[](t, s)] \mid C[v []] \mid C[\mathcal{A} []] \\
& \text{evaluate} && : \text{Term} \rightarrow \text{Value} \\
& \text{evaluate } t && = \text{refocus}(t, \bullet, []) \\
& \text{refocus} : \text{Term} \times \text{Substitution} \times \text{Context} && \rightarrow \text{Value} \\
& \text{refocus}(i, v_1 \cdots v_m, C) && = \text{refocus}(v_i, C) \\
& \text{refocus}(\lambda t, s, C) && = \text{refocus}_{\text{aux}}(C, (\lambda t, s)) \\
& \text{refocus}(t_0 t_1, s, C) && = \text{refocus}(t_0, s, C[[](t_1, s)]) \\
& \text{refocus}(\mathcal{A} t, s, C) && = \text{refocus}(t, s, C[\mathcal{A} []]) \\
& \text{refocus}_{\text{aux}} : \text{Context} \times \text{Value} && \rightarrow \text{Value} \\
& \text{refocus}_{\text{aux}}([], v) && = v \\
& \text{refocus}_{\text{aux}}(C[[](t, s)], v) && = \text{refocus}(t, s, C[v []]) \\
& \text{refocus}_{\text{aux}}(C[(\lambda t, s) []], v) && = \text{refocus}(t, v \cdot s, C) \\
& \text{refocus}_{\text{aux}}(C[\mathcal{A} []], v) && = v
\end{aligned}$$

We observe that this machine coincides with the CEK machine extended with an abort operator—an extension that was designed as such [42]. In particular, the substitution component assumes the role of the environment.

3.7 Correctness

We state the correctness of the final result—the CEK machine—with respect to evaluation in the $\lambda\hat{\rho}\mathcal{A}$ -calculus.

Theorem 1. *For any closed term t in $\lambda\hat{\rho}\mathcal{A}$,*

$$t[\bullet] \rightarrow_{\mathcal{A}}^* (\lambda t') [s'] \quad \text{if and only if} \quad \text{evaluate } t = (\lambda t', s').$$

Proof. The proof relies on the correctness of refocusing [36], and the (trivial) meaning preservation of each of the subsequent transformations. \square

The theorem states that the CEK machine is correct in the sense that it computes closed weak head normal forms, and that it realizes the applicative-order strategy in the $\lambda\hat{\rho}\mathcal{A}$ -calculus, which makes it a call-by-value machine [70]. Furthermore, each of the intermediate abstract machines is also correct with respect to call-by-value evaluation in the $\lambda\hat{\rho}\mathcal{A}$ -calculus.

3.8 Conclusion

We have presented a detailed and systematic derivation of an abstract machine for call-by-value evaluation in the λ -calculus with an abort operation. We started from the specification of the applicative-order reduction strategy in the $\lambda\hat{\rho}$ -calculus with an abort operation and we finished with a formal correspondence between the calculus and the abstract machine.

In the next six sections, we apply the same method in a variety of computational situations, each of which has been separately studied and reported in the literature. Our presentation is structured around four elements: language, notion of context-sensitive reduction, abstract machine, and formal correspondence. The rest is mechanical and therefore omitted.

4 The $\lambda\widehat{\rho}\mathcal{K}$ -calculus

The Krivine machine is probably the best-known abstract machine implementing the normal-order reduction strategy in the λ -calculus [30]. In our previous work [8], we have pointed out that Krivine’s original machine [59] does not coincide with the Krivine Machine As We Know It [21, 23] in that it implements generalized instead of ordinary β -reduction: indeed Krivine’s machine reduces the term $(\lambda\lambda t) t_1 t_2$ in one step whereas the Krivine machine reduces it in two steps. Furthermore, an extension of the archival version of Krivine’s machine [60, Section 3] also caters for call/cc (noted \mathcal{K} below).

In our previous work [8], we have presented the calculus corresponding to the original version of Krivine’s machine. This machine uses closures and an environment and correspondingly, the calculus is one of explicit substitutions, $\lambda\widehat{\rho}$.

Here, we present the calculus corresponding to the archival version of Krivine’s machine with \mathcal{K} . This machine also uses closures and an environment. Correspondingly, the calculus is one of explicit substitutions, $\lambda\widehat{\rho}\mathcal{K}$. We build on top of Krivine’s language of terms by specifying syntactic categories of closures and substitutions as shown below. The calculus is tied to a particular reduction strategy. Here, like Krivine, we consider the normal-order reduction strategy and therefore call by name [70].

4.1 The language of $\lambda\widehat{\rho}\mathcal{K}$

The abstract syntax of the language is as follows:

(terms)	$t ::= i \mid \lambda^n t \mid tt \mid \mathcal{K}t$
(closures)	$c ::= t[s] \mid cc \mid \mathcal{K}c \mid \ulcorner C \urcorner$
(values)	$v ::= (\lambda^n t)[s] \mid \ulcorner C \urcorner$
(substitutions)	$s ::= \bullet \mid c \cdot s$
(reduction contexts)	$C ::= [] \mid C[[] c] \mid C[\mathcal{K}[]]$

A nested λ -abstraction of the form $\lambda^n t$ is to be understood as a syntactic abbreviation for $\underbrace{\lambda\lambda \dots \lambda}_n t$, where t is not a λ -abstraction.

In $\lambda\widehat{\rho}\mathcal{K}$, a value is either a closure with a λ -abstraction in the term part, or the representation of a reduction context captured by \mathcal{K} .

4.2 Notion of context-sensitive reduction

The notion of reduction is specified by the set of rules shown below. The rules (Var) and (Prop) are as in the $\lambda\widehat{\rho}$ -calculus, and (Beta⁺) supersedes the (Beta) rule by performing a generalized β -reduction in one step:

(Var)	$\langle i[c_1 \dots c_j], C \rangle \rightarrow_{\mathcal{K}} \langle c_i, C \rangle$ if $i \leq j$
(Beta ⁺)	$\langle (\lambda^n t)[s], C[[\dots [] c_n] \dots] c_1 \rangle \rightarrow_{\mathcal{K}} \langle t[c_1 \dots c_n \cdot s], C \rangle$
(Beta _C)	$\langle \ulcorner C' \urcorner, C[[] c] \rangle \rightarrow_{\mathcal{K}} \langle c, C' \rangle$

$$\begin{aligned}
(\text{Prop}) \quad & \langle (t_0 t_1)[s], C \rangle \rightarrow_{\mathcal{K}} \langle (t_0[s]) (t_1[s]), C \rangle \\
(\text{Prop}_{\mathcal{K}}) \quad & \langle (\mathcal{K}t)[s], C \rangle \rightarrow_{\mathcal{K}} \langle \mathcal{K}(t[s]), C \rangle \\
(\mathcal{K}_{\lambda}) \quad & \langle \mathcal{K}((\lambda t)[s]), C \rangle \rightarrow_{\mathcal{K}} \langle t[\ulcorner C \urcorner \cdot s], C \rangle \\
(\mathcal{K}_C) \quad & \langle \mathcal{K}\ulcorner C \urcorner, C \rangle \rightarrow_{\mathcal{K}} \langle \ulcorner C \urcorner \ulcorner C \urcorner, C \rangle
\end{aligned}$$

The three last rules account for call/cc: the first is an ordinary propagation rule, and the two others describe a continuation capture. In the first case, the current continuation is captured and passed to a function, and in the second, it is captured and passed to an already captured continuation.

4.3 Krivine's machine

Refocusing, compressing the intermediate transitions, and unfolding the data type of closures mechanically yields the following environment-based machine:

$$\begin{aligned}
\langle i, s, C \rangle & \Rightarrow_{\mathcal{K}} \langle t', s', C \rangle \text{ if } s(i) = (t', s') \\
\langle i, s, C \rangle & \Rightarrow_{\mathcal{K}} \langle C, \ulcorner C \urcorner \rangle \text{ if } s(i) = \ulcorner C \urcorner \\
\langle \lambda^n t, s, C \rangle & \Rightarrow_{\mathcal{K}} \langle C, (\lambda^n t, s) \rangle \\
\langle t_0 t_1, s, C \rangle & \Rightarrow_{\mathcal{K}} \langle t_0, s, C[[\] (t_1, s)] \rangle \\
\langle \mathcal{K}t, s, C \rangle & \Rightarrow_{\mathcal{K}} \langle t, s, C[\mathcal{K}[\]] \rangle \\
\langle [\], v \rangle & \Rightarrow_{\mathcal{K}} v \\
\langle C[[\dots [[\] c_n] \dots] c_1], (\lambda^n t, s) \rangle & \Rightarrow_{\mathcal{K}} \langle t, c_1 \dots c_n \cdot s, C \rangle \\
\langle C[[\] (t, s)], \ulcorner C \urcorner \rangle & \Rightarrow_{\mathcal{K}} \langle t, s, C' \rangle \\
\langle C[[\] \ulcorner C \urcorner], \ulcorner C \urcorner \rangle & \Rightarrow_{\mathcal{K}} \langle C', \ulcorner C \urcorner \rangle \\
\langle C[\mathcal{K}[\]], v \rangle & \Rightarrow_{\mathcal{K}} \langle C[[\] \ulcorner C \urcorner], v \rangle
\end{aligned}$$

This machine coincides with the extension of Krivine's machine with \mathcal{K} —an extension which was designed as such and not connected to any calculus [60, Section 3].

4.4 Formal correspondence

Proposition 4. *For any term t in the $\lambda\widehat{\rho}\mathcal{K}$ -calculus,*

$$t[\bullet] \rightarrow_{\mathcal{K}}^* v \text{ if and only if } \langle t, \bullet, [\] \rangle \Rightarrow_{\mathcal{K}}^* v.$$

The $\lambda\widehat{\rho}\mathcal{K}$ -calculus therefore directly corresponds to the archival version of Krivine's machine with call/cc.

5 The $\lambda\widehat{\rho}\mu$ -calculus

In this section we present a calculus of closures that extends Parigot's $\lambda\mu$ -calculus [69] and the corresponding call-by-name abstract machine obtained by refocusing.

We want to compare our derived abstract machine with an existing one designed by de Groote [38] and therefore we adapt his syntax, which differs from Parigot's in that arbitrary terms can be abstracted by μ (not only named ones). In addition, de Groote presents a calculus of explicit substitutions built on top of the $\lambda\mu$ -calculus, and uses it to prove the correctness of his machine. We show that a $\lambda\widehat{\rho}$ -like calculus of closures is enough to model evaluation in the $\lambda\mu$ -calculus and to derive the same abstract machine as de Groote.

The $\lambda\mu$ -calculus is typed, and suitable typing rules can be given to the calculus of closures we present below. The reduction rules we show satisfy the subject reduction property, and in consequence, the machine we derive operates on typed terms. To remain concise, we omit all the typing considerations and concentrate on the syntactic correspondence between the calculus and the machine.

5.1 The language of $\lambda\hat{\rho}\mu$

We use de Bruijn indices for both the λ -bound variables and the μ -bound variables. The two kinds of variables are represented using the same set of indices, which leads one to an abstract machine with one environment [38]. Alternatively, we could use two separate sets of indices, which would then yield two environments in the resulting machine (one for each kind of variable).

The abstract syntax of the language is specified as follows:

(terms)	$t ::= i \mid \lambda t \mid t t \mid \mu t \mid [i]t$
(closures)	$c ::= t[s] \mid c c$
(values)	$v ::= (\lambda t)[s]$
(substitutions)	$s ::= \bullet \mid C \cdot s \mid c \cdot s$
(reduction contexts)	$C ::= [] \mid C[[] c]$

We consider only closed λ -terms, and $i \geq 0$. Bound variables are indexed starting with 1, and a (free) occurrence of a variable 0 indicates a distinguished toplevel continuation (similar to \mathbf{tp} in Ariola et al.'s setting [6]). A substitution is a non-empty sequence of either closures—to be substituted for λ -bound variables, or captured reduction contexts—to be used when accessing μ -bound variables.

Programs are closures of the form $t[[] \cdot \bullet]$, where the empty context is to be substituted for the toplevel continuation variable 0.

5.2 Notion of context-sensitive reduction

The notion of reduction extends that of the $\lambda\hat{\rho}$ with two rules: (Mu), which captures the entire reduction context and stores it in the substitution, and (Rho), which reinstates a captured context when a continuation variable is applied:

(Beta)	$\langle (\lambda t)[s], C[[] c] \rangle \rightarrow_{\mu} \langle t[c \cdot s], C \rangle$
(Var)	$\langle i[s], C \rangle \rightarrow_{\mu} \langle c, C \rangle \quad \text{if } s(i) = c$
(Prop)	$\langle (t_0 t_1)[s], C \rangle \rightarrow_{\mu} \langle (t_0[s]) (t_1[s]), C \rangle$
(Mu)	$\langle (\mu t)[s], C \rangle \rightarrow_{\mu} \langle t[C \cdot s], [] \rangle$
(Rho)	$\langle ([i]t)[s], [] \rangle \rightarrow_{\mu} \langle t[s], C \rangle \quad \text{if } s(i) = C$

5.3 An eval/apply abstract machine

Refocusing, compressing the intermediate transitions, and unfolding the data type of closures mechanically yields the following environment-based machine:

$\langle \lambda t, s, C \rangle \Rightarrow_{\mu} \langle C, (\lambda t, s) \rangle$
$\langle i, s, C \rangle \Rightarrow_{\mu} \langle t', s', C \rangle \quad \text{if } s(i) = (t', s')$
$\langle t_0 t_1, s, C \rangle \Rightarrow_{\mu} \langle t_0, s, C[[] (t_1, s)] \rangle$
$\langle \mu t, s, C \rangle \Rightarrow_{\mu} \langle t, C \cdot s, [] \rangle$
$\langle [i]t, s, [] \rangle \Rightarrow_{\mu} \langle t, s, C \rangle \quad \text{if } s(i) = C$

$$\begin{aligned} \langle [], v \rangle &\Rightarrow_{\mu} v \\ \langle C[[\] c], (\lambda t, s) \rangle &\Rightarrow_{\mu} \langle t, c \cdot s, C \rangle \end{aligned}$$

This machine coincides with de Groote’s final abstract machine [38, p. 24], except that instead of traversing the environment as a list, it directly fetches the right substitutee for a given index i .

5.4 Formal correspondence

Proposition 5. *For any term t in the $\lambda\hat{\rho}\mu$ -calculus,*

$$t[[\] \cdot \bullet] \rightarrow_{\mu}^* v \quad \text{if and only if} \quad \langle t, [] \cdot \bullet, [] \rangle \Rightarrow_{\mu}^* v.$$

The $\lambda\hat{\rho}\mu$ -calculus therefore directly corresponds to de Groote’s abstract machine for the $\lambda\mu$ -calculus.

6 Delimited continuations

Continuations have been discovered multiple times [73], but they acquired their name for describing jumps [82], using what is now known as continuation-passing style (CPS) [80]. A full-fledged control operator, J [62, 84], however, existed before CPS, providing first-class continuations in direct style. Continuations therefore existed before CPS, and so one could say that it was really CPS that was discovered multiple times.

Conversely, delimited continuations, in the form of the traditional success and failure continuations [76], have been regularly used in artificial-intelligence programming [14, 54, 83] for generators and backtracking. They also occur in the study of reflective towers [79], where the notions of meta-continuation [89] and of “jumpy” vs. “pushy” continuations [34] arose. A full-fledged delimited control operator, $\#$ (pronounced “prompt”), however, was introduced independently of CPS and of reflective towers, to support operational equivalence in λ -calculi with first-class control [41, 44]. Only subsequently were control delimiters connected to success and failure continuations [32].

The goal of this section is to provide a uniform account of delimited continuations. Three data points are in presence—a calculus and an abstract machine, both invented by Felleisen [41], and an extension of CPS, as discovered by Danvy and Filinski [32]:

Calculus: As we show below, an explicit-substitutions version of Felleisen’s calculus of dynamic delimited continuations can be refocused into his extension of the CEK machine, which uses closures and an environment.

Abstract machine: As we have shown elsewhere [11], Felleisen’s extension of the CEK machine is not in defunctionalized form (at least for the usual notion of defunctionalization [35, 72]); it needs some adjustment to be so, which leads one to a dynamic form of CPS that threads a state-like trail of delimited contexts.

CPS: Defunctionalizing Danvy and Filinski’s continuation-based evaluator yields an environment-based machine [7], and we present below the corresponding calculus of static delimited continuations.

The syntactic correspondence makes it possible to directly compare (1) the calculi of dynamic and of static delimited continuations, (2) the extended CEK machine and the machine corresponding to the calculus of static delimited continuations and to the continuation-based evaluator, and (3) the evaluator corresponding to the extended CEK

machine and the continuation-based evaluator. In other words, rather than having to relate heterogeneous semantic artifacts such as a calculus with actual substitutions, an environment-based machine, and a continuation-based evaluator, we are now in position to directly compare two calculi, two abstract machines, and two continuation-based evaluators.

We address static delimited continuations in Section 6.1 and dynamic delimited continuations in Section 6.2. In both cases, we consider the left-to-right applicative-order reduction strategy and therefore left-to-right call by value.

6.1 The $\lambda\hat{\rho}\mathcal{S}$ -calculus

The standard λ -calculus is extended with the control operator shift (written \mathcal{S}) that captures the current delimited continuation and with the control delimiter reset (written $\langle \cdot \rangle$) that initializes the current delimited continuation.

6.1.1 The language of $\lambda\hat{\rho}\mathcal{S}$

The abstract syntax of the language is as follows:

(terms)	$t ::= i \mid \lambda t \mid tt \mid \mathcal{S}t \mid \langle t \rangle$
(closures)	$c ::= t[s] \mid cc \mid \mathcal{S}c \mid \langle c \rangle \mid \ulcorner C \urcorner$
(values)	$v ::= (\lambda t)[s] \mid \ulcorner C \urcorner$
(substitutions)	$s ::= \bullet \mid c \cdot s$
(contexts)	$C_1 ::= [] \mid C_1[v []] \mid C_1[[] c] \mid C_1[\mathcal{S}[]]$
(meta-contexts)	$C_2 ::= \bullet \mid C_1 \cdot C_2$

For readability, we write $C_1 \cdot C_2$ rather than $C_2[\langle C_1[] \rangle]$.

The control operator \mathcal{S} captures the current delimited context and replaces it with the empty context. The control delimiter $\langle \cdot \rangle$ initializes the current delimited context, saving the then-current one onto the meta-context. When a captured delimited context is resumed, the current delimited context is saved onto the meta-context. When the current delimited context completes, the previously saved one, if there is any, is resumed; otherwise, the computation terminates. This informal description paraphrases the definitional interpreter for shift and reset, which has two layers of control—a current delimited continuation (akin to a success continuation) and a meta-continuation (akin to a failure continuation), as arises naturally when one CPS-transforms a direct-style evaluator twice [32]. Elsewhere [7], we have defunctionalized this interpreter into an environment-based machine, which we present next.

6.1.2 The eval/apply/meta-apply abstract machine

The environment-based machine is in “eval/apply/meta-apply” form (to build on Peyton Jones’s terminology [66]) because the continuation is defunctionalized into a context and the corresponding apply transition function, and the meta-continuation is defunctionalized into a meta-context (here a list of contexts) and the corresponding meta-apply transition function:

$$\begin{aligned}
\langle i, s, C_1, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle t', s', C_1, C_2 \rangle \quad \text{if } s(i) = (t', s') \\
\langle \lambda t, s, C_1, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle C_1, (\lambda t, s), C_2 \rangle \\
\langle t_0 t_1, s, C_1, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle t_0, s, C_1[[] (t_1, s)], C_2 \rangle \\
\langle \mathcal{S}t, s, C_1, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle t, s, C_1[\mathcal{S}[]], C_2 \rangle \\
\langle \langle t \rangle, s, C_1, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle t, s, [], C_1 \cdot C_2 \rangle
\end{aligned}$$

$$\begin{aligned}
\langle [], v, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle C_2, v \rangle \\
\langle C_1[[]](t, s), v, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle t, s, C_1[v []], C_2 \rangle \\
\langle C_1[(\lambda t, s) []], v, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle t, v \cdot s, C_1, C_2 \rangle \\
\langle C_1[\ulcorner C_1^\urcorner []], v, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle C_1', v, C_1 \cdot C_2 \rangle \\
\langle C_1[\mathcal{S}[]], (\lambda t, s), C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle t, \ulcorner C_1^\urcorner \cdot s, [], C_2 \rangle \\
\langle C_1[\mathcal{S}[]], \ulcorner C_1^\urcorner, C_2 \rangle &\Rightarrow_{\mathcal{S}} \langle C_1', \ulcorner C_1^\urcorner, [] \cdot C_2 \rangle \\
\langle \bullet, v \rangle &\Rightarrow_{\mathcal{S}} v \\
\langle C_1 \cdot C_2, v \rangle &\Rightarrow_{\mathcal{S}} \langle C_1, v, C_2 \rangle
\end{aligned}$$

We have observed that this machine is in the range of refocusing, transition compression, and closure unfolding for the following calculus $\lambda\hat{\rho}\mathcal{S}$.

6.1.3 Notion of context-sensitive reduction

The $\lambda\hat{\rho}\mathcal{S}$ -calculus uses two layers of contexts: C_1 and C_2 . A non-value closure is decomposed into a redex, a context C_1 , and a meta-context C_2 , and the notion of reduction is specified by the following rules:

$$\begin{aligned}
(\text{Var}) \quad &\langle i[c_1 \cdots c_j], C_1, C_2 \rangle \rightarrow_{\mathcal{S}} \langle c_i, C_1, C_2 \rangle \quad \text{if } i \leq j \\
(\text{Beta}) \quad &\langle ((\lambda t)[s]) v, C_1, C_2 \rangle \rightarrow_{\mathcal{S}} \langle t[v \cdot s], C_1, C_2 \rangle \\
(\text{Beta}_C) \quad &\langle \ulcorner C_1^\urcorner v, C_1, C_2 \rangle \rightarrow_{\mathcal{S}} \langle \langle C_1'[v] \rangle, C_1, C_2 \rangle \\
(\text{Prop}) \quad &\langle (t_0 t_1)[s], C_1, C_2 \rangle \rightarrow_{\mathcal{S}} \langle (t_0[s]) (t_1[s]), C_1, C_2 \rangle \\
(\text{Prop}_{\mathcal{S}}) \quad &\langle (\mathcal{S} t)[s], C_1, C_2 \rangle \rightarrow_{\mathcal{S}} \langle \mathcal{S}(t[s]), C_1, C_2 \rangle \\
(\text{Prop}_{\langle \cdot \rangle}) \quad &\langle \langle t \rangle [s], C_1, C_2 \rangle \rightarrow_{\mathcal{S}} \langle \langle t[s] \rangle, C_1, C_2 \rangle \\
(\mathcal{S}_\lambda) \quad &\langle \mathcal{S}((\lambda t)[s]), C_1, C_2 \rangle \rightarrow_{\mathcal{S}} \langle t[\ulcorner C_1^\urcorner \cdot s], [], C_2 \rangle \\
(\mathcal{S}_C) \quad &\langle \mathcal{S}\ulcorner C_1^\urcorner, C_1, C_2 \rangle \rightarrow_{\mathcal{S}} \langle \ulcorner C_1^\urcorner \ulcorner C_1^\urcorner, [], C_2 \rangle \\
(\text{Reset}) \quad &\langle \langle v \rangle, C_1, C_2 \rangle \rightarrow_{\mathcal{S}} \langle v, C_1, C_2 \rangle
\end{aligned}$$

Since none of the contractions depends on the meta-context, it is evident that the notion of reduction $\rightarrow_{\mathcal{S}}$ is compatible with meta-contexts, but it is not compatible with contexts, due to \mathcal{S}_λ and \mathcal{S}_C . The $\langle \cdot \rangle$ construct therefore delimits the parts of non-value closures in which context-sensitive reductions may occur, and partially restores the compatibility of reductions. In particular, $\langle \langle t[s] \rangle, C_1, C_2 \rangle$ is decomposed into $\langle t[s], [], C_1 \cdot C_2 \rangle$ in the course of decomposition towards a context-sensitive redex.

6.1.4 Formal correspondence

Proposition 6. *For any term t in the $\lambda\hat{\rho}\mathcal{S}$ -calculus,*

$$t[\bullet] \rightarrow_{\mathcal{S}}^* v \quad \text{if and only if} \quad \langle t, \bullet, [], \bullet \rangle \Rightarrow_{\mathcal{S}}^* v.$$

The $\lambda\hat{\rho}\mathcal{S}$ -calculus therefore directly corresponds to the abstract machine for shift and reset.

6.1.5 The CPS hierarchy

Iterating the CPS transformation on a direct-style evaluator for the λ -calculus gives rise to a family of CPS evaluators. At each iteration, one can add shift and reset to the new inner

layer. The result forms a CPS hierarchy of static delimited continuations [32, 37] which Filinski has shown to be able to represent layered monads [46]. Recently, Kameyama has proposed an axiomatization of the CPS hierarchy [56]. Elsewhere [7], we have studied its defunctionalized counterpart and the corresponding hierarchy of calculi.

6.2 The $\lambda\widehat{\rho}\mathcal{F}$ -calculus

The standard λ -calculus is extended with the control operator \mathcal{F} that captures a segment of the current context and with the control delimiter prompt (noted $\#$) that initializes a new segment in the current context.

6.2.1 The language of $\lambda\widehat{\rho}\mathcal{F}$

The abstract syntax of the language is as follows:

(terms)	$t ::= i \mid \lambda t \mid tt \mid \mathcal{F}t \mid \#t$
(closures)	$c ::= t[s] \mid cc \mid \mathcal{F}c \mid \#c \mid \ulcorner C \urcorner$
(values)	$v ::= (\lambda t)[s] \mid \ulcorner C \urcorner$
(substitutions)	$s ::= \bullet \mid c \cdot s$
(reduction contexts)	$C ::= [] \mid C[[]c] \mid C[v[]] \mid C[\mathcal{F}[]] \mid C[\#[]]$

6.2.2 Notion of context-sensitive reduction

The control operator \mathcal{F} captures a segment of the current context up to a mark. The control delimiter $\#$ sets a mark on the current context. When a captured segment is resumed, it is composed with the current context. For the rest, the notion of reduction is as usual:²

(Var)	$\langle i[c_1 \cdots c_j], C \rangle \rightarrow_{\mathcal{F}} \langle c_i, C \rangle$ if $i \leq j$
(Beta)	$\langle ((\lambda t)[s])v, C \rangle \rightarrow_{\mathcal{F}} \langle t[v \cdot s], C \rangle$
(Beta _C)	$\langle \ulcorner C \urcorner v, C \rangle \rightarrow_{\mathcal{F}} \langle C'[v], C \rangle$
(Prop)	$\langle (t_0 t_1)[s], C \rangle \rightarrow_{\mathcal{F}} \langle (t_0[s]) (t_1[s]), C \rangle$
(Prop _{\mathcal{F}})	$\langle (\mathcal{F}t)[s], C \rangle \rightarrow_{\mathcal{F}} \langle \mathcal{F}(t[s]), C \rangle$
(Prop _{$\#$})	$\langle (\#t)[s], C \rangle \rightarrow_{\mathcal{F}} \langle \#(t[s]), C \rangle$
($\mathcal{F}\lambda$)	$\langle \mathcal{F}((\lambda t)[s]), C[\#C'] \rangle \rightarrow_{\mathcal{F}} \langle t[\ulcorner C \urcorner \cdot s], C \rangle$ if C' contains no mark
($\mathcal{F}C$)	$\langle \mathcal{F}\ulcorner C'' \urcorner, C[\#C'] \rangle \rightarrow_{\mathcal{F}} \langle \ulcorner C'' \urcorner \ulcorner C \urcorner, C \rangle$ if C' contains no mark
(Prompt)	$\langle \#v, C \rangle \rightarrow_{\mathcal{F}} \langle v, C \rangle$

Alternatively, we could specify the reduction rules using two layers of contexts, similarly to the $\lambda\widehat{\rho}\mathcal{S}$ -calculus [7, 10, 11]. The difference between the two calculi would then be only in the rule (Beta_C):

$$\text{(Beta}_C\text{)} \quad \langle \ulcorner C_1 \urcorner v, C_1, C_2 \rangle \rightarrow_{\mathcal{F}} \langle C_1'[v], C_1, C_2 \rangle$$

where there is no delimiter around $C'[v]$. As in the previous case of the $\lambda\widehat{\rho}\mathcal{S}$ -calculus, such two-layered decomposition makes it evident that the contraction rules are compatible with the meta-context, since it is isolated by the use of a control delimiter.

²The original version of \mathcal{F} does not reduce its argument first, but its followers do. The present version of \mathcal{F} does likewise here, for a more direct comparison with \mathcal{S} .

6.2.3 The eval/apply abstract machine

Refocusing, compressing the intermediate transitions, and unfolding the data type of closures mechanically yields the following environment-based machine:

$$\begin{aligned}
\langle i, s, C \rangle &\Rightarrow_{\mathcal{F}} \langle t', s', C \rangle \quad \text{if } s(i) = (t', s') \\
\langle \lambda t, s, C \rangle &\Rightarrow_{\mathcal{F}} \langle C, (\lambda t, s) \rangle \\
\langle t_0 t_1, s, C \rangle &\Rightarrow_{\mathcal{F}} \langle t_0, s, C[\](t_1, s) \rangle \\
\langle \mathcal{F} t, s, C \rangle &\Rightarrow_{\mathcal{F}} \langle t, s, C[\mathcal{F}[\]] \rangle \\
\langle \#t, s, C \rangle &\Rightarrow_{\mathcal{F}} \langle t, s, C[\#[\]] \rangle \\
\\
\langle C[\](t, s), v \rangle &\Rightarrow_{\mathcal{F}} \langle t, s, C[v[\]] \rangle \\
\langle C[(\lambda t, s) \]], v \rangle &\Rightarrow_{\mathcal{F}} \langle t, v \cdot s, C \rangle \\
\langle C[\ulcorner C' \urcorner \]], v \rangle &\Rightarrow_{\mathcal{F}} \langle C' \circ C, v \rangle \\
\langle C[\#[C'[\mathcal{F}[\]]], (\lambda t, s) \rangle &\Rightarrow_{\mathcal{F}} \langle t, \ulcorner C' \urcorner \cdot s, C \rangle \\
&\quad \text{where } C' \text{ contains no mark} \\
\langle C[\#[C'[\mathcal{F}[\]]], \ulcorner C'' \urcorner \rangle &\Rightarrow_{\mathcal{F}} \langle C'' \circ C, \ulcorner C'' \urcorner \rangle \\
&\quad \text{where } C' \text{ contains no mark} \\
\langle C[\#[\]], v \rangle &\Rightarrow_{\mathcal{F}} \langle C, v \rangle
\end{aligned}$$

This machine coincides with Felleisen's extension of the CEK machine—an extension which was designed as such [41, Section 3].

6.2.4 Formal correspondence

Proposition 7. *For any term t in the $\lambda\hat{\rho}\mathcal{F}$ -calculus,*

$$t[\bullet] \rightarrow_{\mathcal{F}}^* v \quad \text{if and only if} \quad \langle t, \bullet, [\] \rangle \Rightarrow_{\mathcal{F}}^* v.$$

This proposition parallels Felleisen's second correspondence theorem [41, p. 186]. The $\lambda\hat{\rho}\mathcal{F}$ -calculus therefore directly corresponds to Felleisen's extension of the CEK machine.

6.2.5 A hierarchy of control delimiters

As described by Sitaram and Felleisen [78], one could have not one but several marks in the context and have control operators capture segments of the current context up to a particular mark. For these marks not to interfere in programming practice, they need to be organized hierarchically, forming a hierarchy of control delimiters [78, Section 5]. Alternatively, one could iterate Biernacki et al.'s dynamic CPS transformation [11] to give rise to a hierarchy of dynamic delimited continuations with a functional (CPS) counterpart. Except for the work of Gunter et al. [50] and more recently of Dybvig et al. [39], this area is little explored.

6.3 Conclusion

The syntactic correspondence has made it possible to exhibit the calculus corresponding to static delimited continuations as embodied in the functional idiom of success and failure continuations and more generally in the CPS hierarchy, and to show that (the explicit-substitutions version of) Felleisen's calculus of dynamic delimited continuations corresponds to his extension of the CEK machine [41]. Elsewhere, we present the abstract machine [7] and the evaluator [32] corresponding to static delimited continuations and an

evaluator [11] corresponding to dynamic delimited continuations. We are now in position to compare them pointwise:

- From a calculus point of view, it seems to us that one is better off with layered contexts because it is immediately obvious whether a notion of reduction is compatible with them (see Section 6.1.3); a context containing marks is less easy to treat. Otherwise, the difference between static and dynamic delimited continuations is tiny (see Section 6.2.2), and located in the rule (Beta_C) .
- From a machine point of view, separating between the current delimited context and the other ones is also simpler, as it avoids linear searches, copies, and concatenations (in this respect, efficient implementations, e.g., with an Algol-style display, in effect separate between the current delimited context and the other ones).
- From the point of view of CPS, the abstract machine for dynamic delimited continuations is not in defunctionalized form whereas the abstract machine for static delimited continuations is (and corresponds to a evaluator in CPS). Conversely, defunctionalizing a CPS evaluator provides design guidelines, whereas without CPS, one is on one's own, and locally plausible choices may have unforeseen global consequences which are then taken as the norm. Two cases in point:
 1. in Lisp, it was locally plausible to push both formal and actual parameters at function-call time, and to pop them at return time, but this led to dynamic scope since variable lookup then follows the dynamic link; and
 2. here, it was locally plausible to concatenate a control-stack segment to the current control stack ("From this, we learn that an empty context adds no information." [45, p. 58]), but this led to dynamic delimited continuations since capturing a segment of a concatenated context then gives access to beyond the concatenation point.

Granted, a degree of dynamism makes it possible to write compact programs (e.g., a compositional breadth-first traversal without a data-queue accumulator *and* in direct style [12]), but it is very difficult to reason about them and they are not necessarily more efficient.

- From the point of view of expressiveness, for example, in Lisp, one can simulate the static scope of Scheme by making each lambda-abstraction a "funarg" and in Scheme, one can simulate the dynamic scope of Lisp by threading an environment of fluid variables in a state-monad fashion. Similarly, static delimited continuations can be simulated using dynamic ones by delimiting the extent of each captured continuation [10], and dynamic delimited continuations can be simulated using static ones by threading a trail of contexts in a state-monad fashion [11, 39, 58, 77]. As to which should be used by default, the question then reduces to which behavior is the norm and which should be simulated if it is needed.

In summary, the calculi, the abstract machines, and the evaluators all differ. In one approach, continuations are dynamically composed by concatenating their representations [45] and in the other, continuations are statically composed through a meta-continuation. These differences result from distinct designs: Felleisen and his colleagues started from a calculus and wanted to go "beyond continuations" [44], and therefore beyond CPS, whereas Danvy and Filinski were aiming at a CPS account of delimited control, one that has turned out to be not without practical, semantical, and logical content.

7 Stack inspection

This section addresses Fournet and Gordon’s λ_{sec} -calculus, which formalizes security enforcement by stack inspection [47]. We first present a calculus of closures built on top of the λ_{sec} -calculus, and we construct the corresponding environment-based machine. This machine is a storeless version of the fg machine presented by Clements and Felleisen [16, Figure 1]. (We consider the issue of store-based machines in Section 8.) This machine is not properly tail-recursive, and so Clements and Felleisen presented another machine—the cm machine—which does implement stack inspection in a properly tail-recursive manner [16, Figure 2]. The cm machine builds on Clinger’s formalization of proper tail-recursion (see Section 8) and it is therefore store-based; we considered its storeless version here, and we present the corresponding calculus of closures. We show how the tail-optimization of the cm machine is reflected in the calculus. Finally, we turn to the unzipped version of the cm machine [4] and we present the corresponding state-based calculus of closures.

7.1 The $\lambda_{\widehat{\rho}_{\text{sec}}}$ -calculus

7.1.1 The language of $\lambda_{\widehat{\rho}_{\text{sec}}}$

(terms)	$t ::= i \mid \lambda t \mid tt \mid \mathbf{grant} R \mathbf{in} t \mid \mathbf{test} R \mathbf{then} t \mathbf{else} t \mid R[t] \mid \mathbf{fail}$
(closures)	$c ::= t[s] \mid cc \mid \mathbf{grant} R \mathbf{in} c \mid \mathbf{test} R \mathbf{then} c \mathbf{else} c \mid R[c]$
(values)	$v ::= (\lambda t)[s] \mid \mathbf{fail}$
(substitutions)	$s ::= \bullet \mid c \cdot s$
(reduction contexts)	$C ::= [] \mid C[[] c] \mid C[v []] \mid C[\mathbf{grant} R \mathbf{in} []] \mid C[R[[]]]$
(permissions)	$R \subseteq \mathcal{P}$

The set of terms consists of λ -terms and four constructs for handling different levels of security specified in a set \mathcal{P} : $\mathbf{grant} R \mathbf{in} t$ grants the permissions R to t ; $\mathbf{test} R \mathbf{then} t_0 \mathbf{else} t_1$ proceeds to evaluate t_0 if permissions R are available, and otherwise t_1 ; a frame $R[t]$ restricts the permissions of t to R ; and finally, \mathbf{fail} aborts the computation.

7.1.2 Notion of context-sensitive reduction

Given the predicate $\mathcal{OK}_{\text{sec}}(R, C)$ checking whether the permissions R are available within the context C ,

$$\frac{}{\mathcal{OK}_{\text{sec}}(\emptyset, C)} \quad \frac{}{\mathcal{OK}_{\text{sec}}(R, [])} \quad \frac{\mathcal{OK}_{\text{sec}}(R, C)}{\mathcal{OK}_{\text{sec}}(R, C[[] c])} \quad \frac{\mathcal{OK}_{\text{sec}}(R, C)}{\mathcal{OK}_{\text{sec}}(R, C[v []])} \quad \frac{R \subset R' \quad \mathcal{OK}_{\text{sec}}(R, C)}{\mathcal{OK}_{\text{sec}}(R, C[R' [[]]])} \quad \frac{\mathcal{OK}_{\text{sec}}(R \setminus R', C)}{\mathcal{OK}_{\text{sec}}(R, C[\mathbf{grant} R' \mathbf{in} []])}$$

the notion of reduction is given by the following set of rules:

$$\begin{aligned} (\text{Var}) \quad & \langle i[v_1 \cdots v_j], C \rangle \rightarrow_{\text{sec}} \langle v_i, C \rangle \quad \text{if } i \leq j \\ (\text{Beta}) \quad & \langle ((\lambda t)[s]) v, C \rangle \rightarrow_{\text{sec}} \langle t[v \cdot s], C \rangle \\ (\text{Prop}) \quad & \langle (t_0 t_1)[s], C \rangle \rightarrow_{\text{sec}} \langle (t_0[s]) (t_1[s]), C \rangle \\ (\text{Prop}_G) \quad & \langle (\mathbf{grant} R \mathbf{in} t)[s], C \rangle \rightarrow_{\text{sec}} \langle \mathbf{grant} R \mathbf{in} t[s], C \rangle \\ (\text{Prop}_F) \quad & \langle (R[t])[s], C \rangle \rightarrow_{\text{sec}} \langle R[t[s]], C \rangle \end{aligned}$$

$$\begin{aligned}
(\text{Prop}_T) \quad & \langle (\text{test } R \text{ then } t_0 \text{ else } t_1)[s], C \rangle \rightarrow_{\text{sec}} \langle \text{test } R \text{ then } t_0[s] \text{ else } t_1[s], C \rangle \\
(\text{Frame}) \quad & \langle R[v], C \rangle \rightarrow_{\text{sec}} \langle v, C \rangle \\
(\text{Grant}) \quad & \langle \text{grant } R \text{ in } v, C \rangle \rightarrow_{\text{sec}} \langle v, C \rangle \\
(\text{Test}_1) \quad & \langle \text{test } R \text{ then } c_1 \text{ else } c_2, C \rangle \rightarrow_{\text{sec}} \langle c_1, C \rangle \quad \text{if } \mathcal{OK}_{\text{sec}}(R, C) \\
(\text{Test}_2) \quad & \langle \text{test } R \text{ then } c_1 \text{ else } c_2, C \rangle \rightarrow_{\text{sec}} \langle c_2, C \rangle \quad \text{otherwise} \\
(\text{Fail}) \quad & \langle \text{fail}[s], C \rangle \rightarrow_{\text{sec}} \langle \text{fail}, [] \rangle
\end{aligned}$$

The only context-sensitive rules are (Test_1) and (Test_2) , which perform a reduction step after inspecting the entire context C , and (Fail) which aborts the computation.

7.1.3 An eval/apply abstract machine

Refocusing, compressing the intermediate transitions, and unfolding the data type of closures mechanically yields the following environment-based machine:

$$\begin{aligned}
& \langle i, s, C \rangle \Rightarrow_{\text{sec}} \langle C, s(i) \rangle \\
& \langle \lambda t, s, C \rangle \Rightarrow_{\text{sec}} \langle C, (\lambda t, s) \rangle \\
& \langle t_0 t_1, s, C \rangle \Rightarrow_{\text{sec}} \langle t_0, s, C[[]](t_1, s) \rangle \\
& \langle \text{grant } R \text{ in } t, s, C \rangle \Rightarrow_{\text{sec}} \langle t, s, C[\text{grant } R \text{ in } []] \rangle \\
& \langle \text{test } R \text{ then } t_0 \text{ else } t_1, s, C \rangle \Rightarrow_{\text{sec}} \langle t_0, s, C \rangle \quad \text{if } \mathcal{OK}_{\text{sec}}(R, C) \\
& \langle \text{test } R \text{ then } t_0 \text{ else } t_1, s, C \rangle \Rightarrow_{\text{sec}} \langle t_1, s, C \rangle \quad \text{otherwise} \\
& \langle \text{fail}, s, C \rangle \Rightarrow_{\text{sec}} \text{fail} \\
& \langle R[t], s, C \rangle \Rightarrow_{\text{sec}} \langle t, s, C[R[[]]] \rangle \\
& \langle [], v \rangle \Rightarrow_{\text{sec}} v \\
& \langle C[[]](t, s), v \rangle \Rightarrow_{\text{sec}} \langle t, s, C[v []] \rangle \\
& \langle C[(\lambda t, s) []], v \rangle \Rightarrow_{\text{sec}} \langle t, v \cdot s, C \rangle \\
& \langle C[\text{grant } R \text{ in } []], v \rangle \Rightarrow_{\text{sec}} \langle C, v \rangle \\
& \langle C[R[[]]], v \rangle \Rightarrow_{\text{sec}} \langle C, v \rangle
\end{aligned}$$

This machine is a storeless version of Clements and Felleisen's fg machine [16, Figure 1].

7.1.4 Formal correspondence

Proposition 8. *For any term t in the $\lambda\hat{\rho}_{\text{sec}}$ -calculus,*

$$t[\bullet] \rightarrow_{\text{sec}}^* v \quad \text{if and only if} \quad \langle t, \bullet, [] \rangle \Rightarrow_{\text{sec}}^* v.$$

The $\lambda\hat{\rho}_{\text{sec}}$ -calculus therefore directly corresponds to the storeless version of the fg machine.

7.2 Properly tail-recursive stack inspection

On the ground that the fg machine is not properly tail-recursive, Clements and Felleisen presented a new, properly tail-recursive, machine—the cm machine [16, Figure 2]—thereby debunking the folklore that stack inspection is incompatible with proper tail recursion. Below, we consider the storeless version of the cm machine and we present the underlying calculus of closures.

7.2.1 The storeless cm machine

The cm machine operates on a λ_{sec} -term, an environment, and an evaluation context enriched with updatable permission tables (written m below):

$$\text{(stack frames)} \quad C ::= m[] \mid C[[](c, m)] \mid C[(v, m)[[]]]$$

A permission table is a partial function with a finite domain from a set of permissions \mathcal{P} to the set $\{\perp = \text{not granted}, \top = \text{granted}\}$. A permission table with the empty domain is written ε .

Given the predicate $\mathcal{OK}_{\text{sec}}^{\text{cm}}(R, C)$,

$$\frac{}{\mathcal{OK}_{\text{sec}}^{\text{cm}}(\emptyset, C)} \quad \frac{R \cap m^{-1}(\perp) = \emptyset}{\mathcal{OK}_{\text{sec}}^{\text{cm}}(R, m[])} \\ \frac{R \cap m^{-1}(\perp) = \emptyset \quad \mathcal{OK}_{\text{sec}}^{\text{cm}}(R \setminus m^{-1}(\top), C)}{\mathcal{OK}_{\text{sec}}^{\text{cm}}(R, C[[](c, m)])} \quad \frac{R \cap m^{-1}(\perp) = \emptyset \quad \mathcal{OK}_{\text{sec}}^{\text{cm}}(R \setminus m^{-1}(\top), C)}{\mathcal{OK}_{\text{sec}}^{\text{cm}}(R, C[(v, m)[[]])}$$

the transitions of the storeless cm machine read as follows:

$$\begin{aligned} \langle i, s, C \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \langle C, s(i) \rangle \\ \langle \lambda t, s, C \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \langle C, (\lambda t, s) \rangle \\ \langle t_0 t_1, s, C \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \langle t_0, s, C[[](t_1, s), \varepsilon] \rangle \\ \langle \text{grant } R \text{ in } t, s, C \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \langle t, s, C[R \mapsto \top] \rangle \\ \langle \text{test } R \text{ then } t_0 \text{ else } t_1, s, C \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \langle t_0, s, C \rangle \text{ if } \mathcal{OK}_{\text{sec}}^{\text{cm}}(R, C) \\ \langle \text{test } R \text{ then } t_0 \text{ else } t_1, s, C \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \langle t_1, s, C \rangle \text{ otherwise} \\ \langle \text{fail}, s, C \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \text{fail} \\ \langle R[t], s, C \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \langle t, s, C[\overline{R} \mapsto \perp] \rangle \\ \langle m[], v \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} v \\ \langle C[[](t, s), m], v \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \langle t, s, C[(v, \varepsilon)[[]] \rangle \\ \langle C[(\lambda t, s), m][[]], v \rangle &\Rightarrow_{\text{sec}}^{\text{cm}} \langle t, v \cdot s, C \rangle \end{aligned}$$

where $\overline{R} = \mathcal{P} \setminus R$ and $C[R \mapsto v]$ is a modification of the permission table in the context C obtained by granting or restricting the permissions R , depending on v .

The following proposition states the equivalence of the fg machine and the cm machine with respect to the values they compute:

Proposition 9. *For any term t in the $\lambda_{\widehat{\rho}_{\text{sec}}}$ -calculus,*

$$\langle t, \bullet, [] \rangle \Rightarrow_{\text{sec}}^* v \quad \text{if and only if} \quad \langle t, \bullet, [] \rangle (\Rightarrow_{\text{sec}}^{\text{cm}})^* v.$$

Moreover, it can be shown that each step of the fg machine is simulated by at most one step of the cm machine [16]. At the level of the calculus, this is reflected in the fact that the reduction semantics implemented by the cm machine has fewer reductions than $\lambda_{\widehat{\rho}_{\text{sec}}}$.

7.2.2 The underlying calculus $\lambda_{\widehat{\rho}_{\text{sec}}}^{\text{cm}}$

The calculus corresponding to the storeless cm machine is very close to the $\lambda_{\widehat{\rho}_{\text{sec}}}$ -calculus. The grammar of terms, closures and substitutions is the same, but the reduction contexts (which correspond to the stack frames in the machine) contain permission tables.

Consequently, the functions `plug` and `decompose` are defined in a non-standard way:

$$\begin{aligned}\text{plug } (c, m[]) &= \text{build } (m, c) \\ \text{plug } (c_0, C[[]](c_1, m)) &= \text{plug } (\text{build } (m, c_0 \ c_1), C) \\ \text{plug } (c, C[(v, m)[]]) &= \text{plug } (\text{build } (m, v \ c), C)\end{aligned}$$

where the auxiliary function `build` conservatively constructs a closure based on the permission table of the reduction context:

$$\begin{aligned}\text{build}_G (m, c) &= \begin{cases} c & \text{if } m^{-1}(\top) = \emptyset \\ \text{grant } m^{-1}(\top) \text{ in } c & \text{otherwise} \end{cases} \\ \text{build}_F (m, c) &= \begin{cases} c & \text{if } m^{-1}(\perp) = \emptyset \\ \overline{m^{-1}(\perp)}[c] & \text{otherwise} \end{cases} \\ \text{build } (m, c) &= \text{build}_F (m, \text{build}_G (m, c))\end{aligned}$$

Any closure that is not already a value or a potential redex, can be further decomposed as follows:

$$\begin{aligned}\text{decompose } (c_0 \ c_1, C) &= \text{decompose } (c_0, C[[]](c_1, \varepsilon)) \\ \text{decompose } (\text{grant } R \text{ in } c, C) &= \text{decompose } (c, C[\overline{R} \mapsto \top]) \\ \text{decompose } (R[c], C) &= \text{decompose } (c, C[\overline{R} \mapsto \perp]) \\ \text{decompose } (v, C[[]](c, m)) &= \text{decompose } (c, C[(v, \varepsilon)[]])\end{aligned}$$

The notion of reduction includes most rules of the $\lambda\hat{\rho}_{\text{sec}}$ -calculus, except for (Frame) and (Grant).

From a calculus standpoint, Clements and Felleisen therefore obtained proper tail recursion by changing the computational model (witness the change from $\mathcal{OK}_{\text{sec}}$ to $\mathcal{OK}_{\text{sec}}^{\text{cm}}$) and by simplifying the reduction rules and modifying the compatibility rules.

7.3 State-based properly tail-recursive stack inspection

On the observation that the stack of the cm machine can be unzipped into the usual control stack of the CEK machine and a state-like list of permission tables, Ager et al. have presented an unzipped version of the cm machine (characterizing properly tail-recursive stack inspection as a monad, in passing) [4]. We first present this machine, and then the corresponding calculus of closures.

7.3.1 The unzipped storeless cm machine

The unzipped cm machine operates on a λ_{sec} -term, an environment, and an ordinary evaluation context. In addition, the machine has a read-write security register m holding the current permission table and a read-only security register ms holding a list of outer permission tables. Given the predicate $\mathcal{OK}_{\text{sec}}^{\text{ucm}}(R, m, ms)$,

$$\frac{}{\mathcal{OK}_{\text{sec}}^{\text{ucm}}(\emptyset, m, ms)} \quad \frac{R \cap m^{-1}(\perp) = \emptyset}{\mathcal{OK}_{\text{sec}}^{\text{ucm}}(R, m, \bullet)} \quad \frac{R \cap m^{-1}(\perp) = \emptyset \quad \mathcal{OK}_{\text{sec}}^{\text{ucm}}(R \setminus m^{-1}(\top), m', ms)}{\mathcal{OK}_{\text{sec}}^{\text{ucm}}(R, m, m' \cdot ms)}$$

the transitions of the unzipped storeless cm machine read as follows:

$$\begin{aligned}\langle i, s, m, ms, C \rangle &\Rightarrow_{\text{sec}}^{\text{ucm}} \langle C, s(i), ms \rangle \\ \langle \lambda t, s, m, ms, C \rangle &\Rightarrow_{\text{sec}}^{\text{ucm}} \langle C, (\lambda t, s), ms \rangle \\ \langle t_0 \ t_1, s, m, ms, C \rangle &\Rightarrow_{\text{sec}}^{\text{ucm}} \langle t_0, s, \varepsilon, m \cdot ms, C[[]](t_1, s) \rangle\end{aligned}$$

$$\begin{aligned}
& \langle \text{grant } R \text{ in } t, s, m, ms, C \rangle \Rightarrow_{\text{sec}}^{\text{ucm}} \langle t, s, m[R \mapsto \top], ms, C \rangle \\
& \langle \text{test } R \text{ then } t_0 \text{ else } t_1, s, m, ms, C \rangle \Rightarrow_{\text{sec}}^{\text{ucm}} \langle t_0, s, m, ms, C \rangle \text{ if } \mathcal{OK}_{\text{sec}}^{\text{ucm}}(R, m, ms) \\
& \langle \text{test } R \text{ then } t_0 \text{ else } t_1, s, m, ms, C \rangle \Rightarrow_{\text{sec}}^{\text{ucm}} \langle t_1, s, m, ms, C \rangle \text{ otherwise} \\
& \langle R[t], s, m, ms, C \rangle \Rightarrow_{\text{sec}}^{\text{ucm}} \langle t, s, m[\overline{R} \mapsto \perp], ms, C \rangle \\
& \langle \text{fail}, s, m, ms, C \rangle \Rightarrow_{\text{sec}}^{\text{ucm}} \text{fail} \\
& \langle [], v, \bullet \rangle \Rightarrow_{\text{sec}}^{\text{ucm}} v \\
& \langle C[[]](t, s), v, ms \rangle \Rightarrow_{\text{sec}}^{\text{ucm}} \langle t, s, \varepsilon, ms, C[v[[]]] \rangle \\
& \langle C[(\lambda t, s)[[]], v, m \cdot ms] \rangle \Rightarrow_{\text{sec}}^{\text{ucm}} \langle t, v \cdot s, m, ms, C \rangle
\end{aligned}$$

The following proposition states the equivalence of the cm machine and the unzipped cm machine:

Proposition 10. *For any term t in the $\lambda\widehat{\rho}_{\text{sec}}^{\text{ucm}}$ -calculus,*

$$\langle t, \bullet, [] \rangle (\Rightarrow_{\text{sec}}^{\text{cm}})^* v \text{ if and only if } \langle t, \bullet, \varepsilon, \bullet, [] \rangle (\Rightarrow_{\text{sec}}^{\text{ucm}})^* v.$$

Moreover, it can be shown that each step of the cm machine is simulated by one step of the unzipped cm machine.

7.3.2 The language of $\lambda\widehat{\rho}_{\text{sec}}^{\text{ucm}}$

(terms)	$t ::= i \mid \lambda t \mid tt \mid \text{grant } R \text{ in } t \mid \text{test } R \text{ then } t \text{ else } t \mid R[t] \mid \text{fail}$
(closures)	$c ::= t[s]$
(values)	$v ::= (\lambda t)[s] \mid \text{fail}$
(substitutions)	$s ::= \bullet \mid c \cdot s$
(reduction contexts)	$C ::= [] \mid C[[]]c \mid C[v[[]]]$
(annotated closures)	$\tilde{c} ::= c[m, ms] \mid c\tilde{c} \mid \tilde{c}c \mid \text{fail}$
(annotated values)	$\tilde{v} ::= v[m, ms] \mid \text{fail}$

7.3.3 Notion of context-sensitive reduction

The notion of reduction is specified by the rules below. Compared to the rules of Section 7.1.2, the current permission table and the list of outer permission tables are propagated locally to each closure being evaluated. When a value is consumed, the current permission table is discarded.

The (Prop) reduction rule illustrates the propagation to a subclosure to be evaluated:

$$\langle (t_0 t_1)[s][m, ms], C \rangle \rightarrow_{\text{sec}}^{\text{ucm}} \langle (t_0[s][\varepsilon, m \cdot ms]) (t_1[s]), C \rangle$$

So do half of the other reduction rules:

(Var)	$\langle i[c_1 \cdots c_j][m, ms], C \rangle \rightarrow_{\text{sec}}^{\text{ucm}} \langle c_i[m, ms], C \rangle \text{ if } i \leq j$
(Test ₁)	$\langle (\text{test } R \text{ then } t_0 \text{ else } t_1)[s][m, ms], C \rangle \rightarrow_{\text{sec}}^{\text{ucm}} \langle t_0[s][m, ms], C \rangle \text{ if } \mathcal{OK}_{\text{sec}}^{\text{ucm}}(R, m, ms)$
(Test ₂)	$\langle (\text{test } R \text{ then } t_0 \text{ else } t_1)[s][m, ms], C \rangle \rightarrow_{\text{sec}}^{\text{ucm}} \langle t_1[s][m, ms], C \rangle \text{ otherwise}$
(Fail)	$\langle \text{fail}[s][m, ms], C \rangle \rightarrow_{\text{sec}}^{\text{ucm}} \langle \text{fail}, [] \rangle$

A new reduction rule, however, is now necessary to go from one evaluated subclosure to a subclosure to evaluate:

$$\text{(Switch)} \quad \langle (v[m, ms]) c, C \rangle \rightarrow_{\text{sec}}^{\text{ucm}} \langle v(c[\varepsilon, ms]), C \rangle$$

The (Beta) rule doubles up with discarding the permission table of the actual parameter:

$$\text{(Beta)} \langle ((\lambda t)[s]) (v[m, m' \cdot ms]), C \rangle \rightarrow_{\text{sec}}^{\text{ucm}} \langle t[v \cdot s][m', ms], C \rangle$$

Finally, the (Frame) and (Grant) rules embody the state counterpart of Clements and Felleisen's design to enable proper tail recursion:

$$\text{(Frame)} \quad \langle R[t][s][m, ms], C \rangle \rightarrow_{\text{sec}}^{\text{ucm}} \langle t[s][m[\overline{R} \mapsto \perp], ms], C \rangle$$

$$\text{(Grant)} \quad \langle (\text{grant } R \text{ in } t)[s][m, ms], C \rangle \rightarrow_{\text{sec}}^{\text{ucm}} \langle t[s][m[R \mapsto \top], ms], C \rangle$$

7.3.4 Formal correspondence

Proposition 11. *For any term t in the λ_{sec} -calculus,*

$$t[\bullet] (\rightarrow_{\text{sec}}^{\text{ucm}})^* v \quad \text{if and only if} \quad \langle t, \bullet, \varepsilon, \bullet, [] \rangle (\Rightarrow_{\text{sec}}^{\text{ucm}})^* v.$$

7.4 Conclusion

We have presented three corresponding calculi of closures and machines for stack inspection, showing first how the storeless fg machine reflects the $\lambda_{\widehat{\rho}_{\text{sec}}}$ -calculus, second, how the $\lambda_{\widehat{\rho}_{\text{sec}}}^{\text{cm}}$ -calculus reflects the storeless cm machine, and third, how the $\lambda_{\widehat{\rho}_{\text{sec}}}^{\text{ucm}}$ -calculus reflects the unzipped storeless cm machine. In doing so, we have provided a calculus account of machine design and optimization for stack inspection.

8 A calculus for proper tail-recursion

At PLDI'98 [19], Clinger presented a properly tail-recursive semantics for Scheme in the form of a store-based abstract machine. This machine models the memory-allocation behavior of function calls in Scheme, and Clinger used it to specify in which sense an implementation should not run out of memory when processing a tail-recursive program (such as a program in CPS).

We first present a similar machine for the λ -calculus with left-to-right call-by-value evaluation and assignments. This machine is in the range of refocusing, transition compression, and closure unfolding, and so we next present the corresponding store-based calculus, $\lambda_{\widehat{\rho}_{\text{ptr}}}$.

8.1 A simplified version of Clinger's abstract machine

Our simplified version of Clinger's abstract machine is an eval/apply machine with an environment and a store:

$$\begin{aligned} \langle x, s, C, \sigma \rangle &\Rightarrow_{\text{ptr}} \langle C, \sigma(s(x)), \sigma \rangle \\ \langle \lambda x.t, s, C, \sigma \rangle &\Rightarrow_{\text{ptr}} \langle C, (\lambda x.t, s), \sigma \rangle \\ \langle t_0 t_1, s, C, \sigma \rangle &\Rightarrow_{\text{ptr}} \langle t_0, s, C[[]](t_1, s), \sigma \rangle \\ \langle x := t, s, C, \sigma \rangle &\Rightarrow_{\text{ptr}} \langle t, s, C[\text{upd}(s(x), [])], \sigma \rangle \\ \langle [], v, \sigma \rangle &\Rightarrow_{\text{ptr}} (v, \sigma) \\ \langle C[[]](t, s), v, \sigma \rangle &\Rightarrow_{\text{ptr}} \langle t, s, C[v []], \sigma \rangle \\ \langle C[(\lambda x.t, s) []], v, \sigma \rangle &\Rightarrow_{\text{ptr}} \langle t, (x, \ell) \cdot s, C, \sigma[\ell \mapsto v] \rangle \\ &\quad \text{if } \ell \text{ does not occur within } s, C, v, \sigma \\ \langle C[\text{upd}(\ell, [])], v, \sigma \rangle &\Rightarrow_{\text{ptr}} \langle C, \sigma(\ell), \sigma[\ell \mapsto v] \rangle \end{aligned}$$

Locations ℓ range over an unspecified set of locations. A store σ is a finite mapping from locations to value closures. Denotable values are locations.

Clinger's machine also has a garbage-collection rule [19, Figure 5 and Section 3], but for simplicity we ignore it here.

8.2 The language of $\lambda\hat{\rho}_{\text{ptr}}$

The abstract syntax of the language is as follows:

(terms)	$t ::= x \mid \lambda x.t \mid tt \mid x := t$
(closures)	$c ::= t[s] \mid cc$
(values)	$v ::= (\lambda x.t)[s]$
(substitutions)	$s ::= \bullet \mid (x, \ell) \cdot s$
(red. contexts)	$C ::= [] \mid C[[] c] \mid C[v []] \mid C[\text{upd}(\ell, [])]$
(store)	$\sigma ::= \bullet \mid \sigma[\ell \mapsto v]$
(store closures)	$\tilde{c} ::= c[\sigma]$
(store values)	$\tilde{v} ::= v[\sigma]$
(store contexts)	$\tilde{C} ::= C[\sigma]$

8.3 Notion of context-sensitive reduction

In the rules below, (Var) dereferences the store; (Beta) allocates a fresh location, and extends both the substitution and the store with it; (Prop) is context-insensitive and therefore essentially as in the $\lambda\hat{\rho}$ -calculus; and (Upd) updates the store.

(Var)	$\langle x_i[(x_1, \ell_1) \cdots (x_j, \ell_j)], C[\sigma] \rangle \rightarrow_{\text{ptr}} \langle \sigma(\ell_i), C[\sigma] \rangle \quad \text{if } i \leq j$
(Beta)	$\langle ((\lambda x.t)[s]) v, C[\sigma] \rangle \rightarrow_{\text{ptr}} \langle t[(x, \ell) \cdot s], C[\sigma[\ell \mapsto v]] \rangle$ if ℓ does not occur within s, C, v, σ
(Prop)	$\langle (t_0 t_1)[s], C[\sigma] \rangle \rightarrow_{\text{ptr}} \langle (t_0[s]) (t_1[s]), C[\sigma] \rangle$
(Upd)	$\langle \text{upd}(\ell, v), C[\sigma] \rangle \rightarrow_{\text{ptr}} \langle \sigma(\ell), C[\sigma[\ell \mapsto v]] \rangle$

Refocusing, compressing the intermediate transitions, and unfolding the data type of closures mechanically yields the abstract machine of Section 8.1.

8.4 Formal correspondence

Proposition 12. *For any term t in the $\lambda\hat{\rho}_{\text{ptr}}$ -calculus,*

$$t[\bullet][\bullet] \rightarrow_{\text{ptr}}^* v[\sigma] \quad \text{if and only if} \quad \langle t, \bullet, [], \bullet \rangle \Rightarrow_{\text{ptr}}^* (v, \sigma).$$

The $\lambda\hat{\rho}_{\text{ptr}}$ -calculus therefore directly corresponds to the simplified version of Clinger's properly tail-recursive machine.

In Section 7, we showed storeless variants of two machines for stack inspection (the *fg* and the *cm* machines). The original versions of these machines use a store in the Clinger fashion [16], and we can exhibit their underlying calculi with an explicit representation of the store, as straightforward extensions of the storeless calculi. For conciseness, we do not include them here.

9 A lazy calculus of closures

The store-based account of proper tail-recursion from Section 8 suggests the following lazy calculus of closures, $\lambda\widehat{\rho}_1$.

9.1 The language of $\lambda\widehat{\rho}_1$

The abstract syntax of the language is as follows:

(terms)	$t ::= i \mid \lambda t \mid t t$
(closures)	$c ::= t[s] \mid c \ell \mid \mathbf{upd}(\ell, c)$
(values)	$v ::= (\lambda t)[s]$
(substitutions)	$s ::= \bullet \mid \ell \cdot s$
(reduction contexts)	$C ::= [] \mid C[[\] \ell] \mid C[\mathbf{upd}(\ell, [\])]$
(store)	$\sigma ::= \bullet \mid \sigma[\ell \mapsto v]$
(store closures)	$\tilde{c} ::= c[\sigma]$
(store values)	$\tilde{v} ::= v[\sigma]$
(store contexts)	$\tilde{C} ::= C[\sigma]$

9.2 Notion of context-sensitive reduction

The notion of reduction is specified by the five rules shown below.

(Var ₁)	$\langle i[\ell_1 \cdots \ell_j], C[\sigma] \rangle \rightarrow_1 \langle v, C[\sigma] \rangle$	if $\sigma(\ell_i) = v$
(Var ₂)	$\langle i[\ell_1 \cdots \ell_j], C[\sigma] \rangle \rightarrow_1 \langle \mathbf{upd}(\ell_i, c), C[\sigma] \rangle$	if $\sigma(\ell_i) = c$
(Beta)	$\langle ((\lambda t)[s]) \ell, C[\sigma] \rangle \rightarrow_1 \langle t[\ell \cdot s], C[\sigma] \rangle$	
(App)	$\langle (t_0 t_1)[s], C[\sigma] \rangle \rightarrow_1 \langle (t_0[s]) \ell, C[\sigma[\ell \mapsto t_1[s]]] \rangle$	where ℓ does not occur in s, C, σ
(Upd)	$\langle \mathbf{upd}(\ell, v), C[\sigma] \rangle \rightarrow_1 \langle v, C[\sigma[\ell \mapsto v]] \rangle$	

Variables denote locations, and have two reduction rules, depending on whether the store holds a value or not at that location. In the former case—handled by (Var₁)—the result is this value, the current context, and the current store. In the latter case—handled by (Var₂)—a special closure $\mathbf{upd}(\ell, c)$ is created, indicating that c is a shared computation. When this computation completes and yields a value, the store at location ℓ should be updated with this value, which is achieved by (Upd). Since every argument to an application can potentially be shared, (App) conservatively allocates a new location in the store for such shared closures. (Beta) extends the substitution with this location.

9.3 An eval/apply abstract machine

Refocusing, compressing the intermediate transitions, and unfolding the data type of closures mechanically yields the following store-based machine:³

$\langle i, s, C, \sigma \rangle \Rightarrow_1 \langle C, (\lambda t', s'), \sigma \rangle$	where $s(i) = \ell$ and $\sigma(\ell) = (\lambda t')[s']$
$\langle i, s, C, \sigma \rangle \Rightarrow_1 \langle t', s', C[\mathbf{upd}(\ell, [\])], \sigma \rangle$	where $s(i) = \ell$ and $\sigma(\ell) = t'[s']$
$\langle \lambda t, s, C, \sigma \rangle \Rightarrow_1 \langle C, (\lambda t, s), \sigma \rangle$	
$\langle t_0 t_1, s, C, \sigma \rangle \Rightarrow_1 \langle t_0, s, C[[\] \ell], \sigma[\ell \mapsto (t_1, s)] \rangle$	where ℓ does not occur in s, C, σ

³When a shared closure is to be evaluated, the current context is extended with what is known as an ‘update marker’ in the Three Instruction Machine (denoted $C[\mathbf{upd}(\ell, [\])]$ here).

$$\begin{aligned}
\langle [], v, \sigma \rangle &\Rightarrow_1 \langle v, \sigma \rangle \\
\langle C[[\ell], (\lambda t, s), \sigma] \rangle &\Rightarrow_1 \langle t, \ell \cdot s, C, \sigma \rangle \\
\langle C[\mathbf{upd}(\ell, [])], v, \sigma \rangle &\Rightarrow_1 \langle C, v, \sigma[\ell \mapsto v] \rangle
\end{aligned}$$

This lazy abstract machine coincides with the one derived by Ager et al. out of a call-by-need interpreter for the λ -calculus [3], thereby connecting the present syntactic correspondence between calculi and abstract machines with the functional correspondence between abstract machines and evaluators [2, 4, 7, 28].

9.4 Formal correspondence

Proposition 13. *For any term t in the $\lambda\widehat{\rho}_1$ -calculus,*

$$t[\bullet][\bullet] \rightarrow_1^* v[\sigma] \quad \text{if and only if} \quad \langle t, \bullet, [], \bullet \rangle \Rightarrow_1^* \langle v, \sigma \rangle.$$

The $\lambda\widehat{\rho}_1$ -calculus therefore directly corresponds to call-by-need evaluation [85].

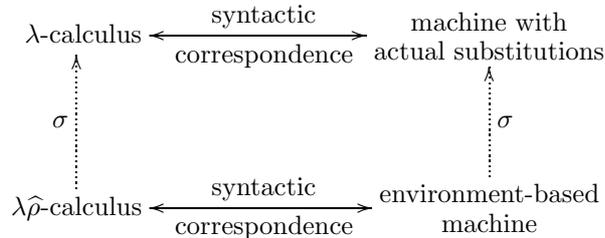
In $\lambda\widehat{\rho}_1$, sharing is made possible through a global heap where actual parameters are stored. On the other hand, a number of other calculi modeling call by need extend the set of terms with a *local* let-like construct, either by statically translating the source language into an intermediate language with explicit indications of sharing (as in Launchbury's approach [63]), or by providing dynamic reduction rules to the same effect (as in Ariola et al.'s calculus [5]). A sequence of let constructs binding variables to shared computations is a local version of a global heap where shared computations are bound to locations; extra reductions are then needed to propagate all the let constructs to the top level.

Another specificity of $\lambda\widehat{\rho}_1$ is that allocation occurs early, i.e., a new cell is allocated in the store every time an application is evaluated. Allocation, however, occurs late in Ariola et al.'s semantics, i.e., a new binding is created only when the operator of the application is known to be a λ -abstraction. Delaying allocation is useful in the presence of strict functions, which we do not consider here.

We can construct a local version of our calculus with either of the store propagated inside closures or of late allocation. From there, one can mechanically derive the corresponding abstract machines.

10 Conclusion

We have presented a series of calculi and abstract machines accounting for a variety of computational effects, making it possible to directly reason about a computation in the calculus and in the corresponding abstract machine (horizontally in the diagram below) and to directly account for actual and explicit substitutions both in the world of calculi and in the world of abstract machines (vertically in the diagram below, where σ maps a closure into the corresponding λ -term and an environment-machine configuration into a configuration in the corresponding machine with actual substitutions):



The correspondence between each calculus and each abstract machine is simple and each can be mechanically built from the other. All of the calculi are new. Many of the abstract machines are known and have been independently designed and proved correct.

The work reported here leads us to drawing the following conclusions.

Curien’s calculus of closures: Once extended to account for one-step reduction, Curien’s calculus of closures directly corresponds to the notions of evaluation (i.e., weak-head normalization) accounted for by environment-based machines, even in the presence of computational effects (state and control).

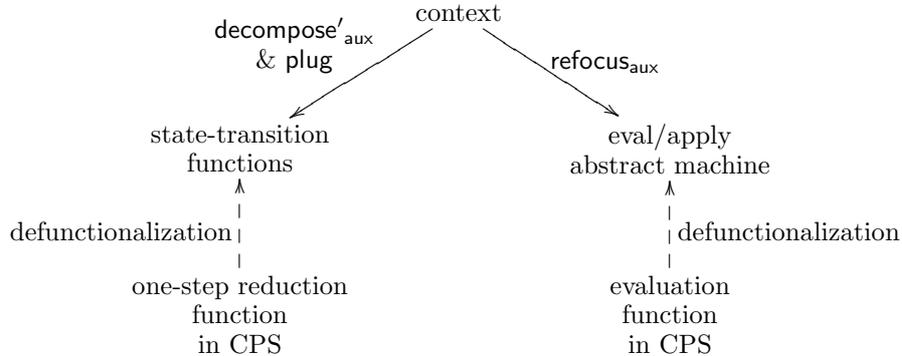
Refocusing: Despite its pragmatic origin—fusing a plug function and a decomposition function in a reduction-based evaluation function to improve its efficiency [36], and in combination with compressing intermediate transitions and unfolding closures, refocusing proves consistently useful to construct reduction-free evaluation functions in the form of abstract machines, even in the presence of computational effects.

Defunctionalization: Despite its practical origin—representing a higher-order function as a data type and a dispatch function [72], defunctionalization proves consistently useful, witness the next item and also the fact that except for the abstract machines for $\lambda\hat{\rho}\mathcal{F}$ and the cm machine, *all* the abstract machines in this article are in defunctionalized form.

Reduction contexts and evaluation contexts: There are three objective reasons—one extensional and two intensional—why contexts are useful as well as, in some sense, unavoidable:

- reduction contexts are in one-to-one correspondence with the compatibility rules of a calculus;
- reduction contexts are the data type of the defunctionalized continuation of a one-step reduction function (as used in a reduction-based (weak-head) normalization function); and
- evaluation contexts are the data type of the defunctionalized continuation of an evaluation function (as used in a reduction-free (weak-head) normalization function).

If nothing else, each of these three reasons has practical value as a guideline for writing the grammar of reduction / evaluation contexts (which can be tricky in practice). But more significantly [27], reduction contexts and evaluation contexts *coincide*, which means that as a data type, they mediate between one-step reduction and evaluation, given an appropriate dispatch function:



Indeed, as initiated by Reynolds [29,72], defunctionalizing a continuation-passing evaluator yields an abstract machine [2–4,7], and as already pointed out above, a vast number of abstract machines are in defunctionalized form [9,11,28].

Together, the syntactic correspondence between calculi and abstract machines and the functional correspondence between abstract machines and evaluators therefore connect apparently distinct approaches to the same computational situations. We already illustrated this connection in Section 6 with delimited continuations; let us briefly illustrate it further with the simpler example of call/cc:

Call/cc was introduced in Scheme [17] as a Church encoding of Reynolds’s escape operator [72]. A typed version of it is available in Standard ML of New Jersey [52] and Griffin has identified its logical content [49]. It is endowed with a variety of specifications: a CPS transformation [33], a CPS interpreter [53,72], a denotational semantics [57], a computational monad [86], a big-step operational semantics [52], the CEK machine [43], calculi in the form of reduction semantics [42], and a number of implementation techniques [18,25,55]—not to mention its call-by-name version that allegedly accounts for the axiom of choice in the archival version of Krivine’s machine [60].

Question: How do we know that all the artifacts in this semantic jungle define the same call/cc?

The elements of answer we contribute here are that the syntactic correspondence links calculi and abstract machines, and the functional correspondence links abstract machines and evaluators. So by construction, all these specifications are inter-derivable and therefore they are consistent.

Normalization by evaluation: Finally, refocusing provides a guideline for constructing reduction-free normalization functions out of reduction-based ones [26]. The reduction-free normalization functions take the form of eval/apply abstract machines, which usually are in defunctionalized form, which paves the way to writing normalization functions as usually encountered in the area of normalization by evaluation [31]. We have illustrated the method here with weak reduction and weak-head normalization (i.e., evaluation), but it also works for strong reduction and normalization, thus linking one-step reduction, abstract machines for strong reduction, and normalization functions.

Acknowledgments

We are grateful to Julia Lawall and Kristian Støvring for commenting a preliminary version of this article.

This work is partially supported by the ESPRIT Working Group APPSEM II (<http://www.appsem.org>) and by the Danish Natural Science Research Council, Grant no. 21-03-0545.

References

- [1] Martín Abadi, Luca Cardelli, Pierre-Louis Curien, and Jean-Jacques Lévy. Explicit substitutions. *Journal of Functional Programming*, 1(4):375–416, 1991.
- [2] Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. A functional correspondence between evaluators and abstract machines. In Dale Miller, editor, *Proceedings of the Fifth ACM-SIGPLAN International Conference on Principles and Practice of Declarative Programming (PPDP'03)*, pages 8–19. ACM Press, August 2003.
- [3] Mads Sig Ager, Olivier Danvy, and Jan Midtgaard. A functional correspondence between call-by-need evaluators and lazy abstract machines. *Information Processing Letters*, 90(5):223–232, 2004. Extended version available as the technical report BRICS RS-04-3.
- [4] Mads Sig Ager, Olivier Danvy, and Jan Midtgaard. A functional correspondence between monadic evaluators and abstract machines for languages with computational effects. *Theoretical Computer Science*, 342(1):149–172, 2005. Extended version available as the technical report BRICS RS-04-28.
- [5] Zena M. Ariola, Matthias Felleisen, John Maraist, Martin Odersky, and Philip Wadler. The call-by-need lambda calculus. In Peter Lee, editor, *Proceedings of the Twenty-Second Annual ACM Symposium on Principles of Programming Languages*, pages 233–246, San Francisco, California, January 1995. ACM Press.
- [6] Zena M. Ariola and Hugo Herbelin. Minimal classical logic and control operators. In Jos C. M. Baeten, Jan Karel Lenstra, Joachim Parrow, and Gerhard J. Woeginger, editors, *Automata, Languages and Programming, 30th International Colloquium (ICALP 2003)*, number 2719 in Lecture Notes in Computer Science, pages 871–885, Eindhoven, The Netherlands, July 2003. Springer.
- [7] Małgorzata Biernacka, Dariusz Biernacki, and Olivier Danvy. An operational foundation for delimited continuations in the CPS hierarchy. *Logical Methods in Computer Science*, 1(2:5):1–39, November 2005. A preliminary version was presented at the Fourth ACM SIGPLAN Workshop on Continuations (CW'04).
- [8] Małgorzata Biernacka and Olivier Danvy. A concrete framework for environment machines. Research Report BRICS RS-05-15, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, May 2005.
- [9] Dariusz Biernacki and Olivier Danvy. From interpreter to logic engine by defunctionalization. In Maurice Bruynooghe, editor, *Logic Based Program Synthesis and Transformation, 13th International Symposium, LOPSTR 2003*, number 3018 in Lecture Notes in Computer Science, pages 143–159, Uppsala, Sweden, August 2003. Springer-Verlag.
- [10] Dariusz Biernacki and Olivier Danvy. A simple proof of a folklore theorem about delimited control. Research Report BRICS RS-05-10, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, March 2005. Accepted for publication in the *Journal of Functional Programming* as a theoretical pearl.

- [11] Dariusz Biernacki, Olivier Danvy, and Kevin Millikin. A dynamic continuation-passing style for dynamic delimited continuations. Research Report BRICS RS-05-16, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, May 2005.
- [12] Dariusz Biernacki, Olivier Danvy, and Chung-chieh Shan. On the dynamic extent of delimited continuations. *Information Processing Letters*, 96(1):7–17, 2005. Extended version available as the technical report BRICS RS-05-13.
- [13] Robert (Corky) Cartwright, editor. *Proceedings of the 1988 ACM Conference on Lisp and Functional Programming*, Snowbird, Utah, July 1988. ACM Press.
- [14] Eugene Charniak, Christopher Riesbeck, and Drew McDermott. *Artificial Intelligence Programming*. Lawrence Earlbaum Associates, 1980.
- [15] Alonzo Church. *The Calculi of Lambda-Conversion*. Princeton University Press, 1941.
- [16] John Clements and Matthias Felleisen. A tail-recursive semantics for stack inspection. *ACM Transactions on Programming Languages and Systems*, 26(6):1029–1052, 2004.
- [17] William Clinger, Daniel P. Friedman, and Mitchell Wand. A scheme for a higher-level semantic algebra. In John Reynolds and Maurice Nivat, editors, *Algebraic Methods in Semantics*, pages 237–250. Cambridge University Press, 1985.
- [18] William Clinger, Anne H. Hartheimer, and Eric M. Ost. Implementation strategies for first-class continuations. *Higher-Order and Symbolic Computation*, 12(1):7–45, 1999.
- [19] William D. Clinger. Proper tail recursion and space efficiency. In Keith D. Cooper, editor, *Proceedings of the ACM SIGPLAN’98 Conference on Programming Languages Design and Implementation*, pages 174–185, Montréal, Canada, June 1998. ACM Press.
- [20] Pierre Crégut. An abstract machine for lambda-terms normalization. In Wand [88], pages 333–340.
- [21] Pierre Crégut. Strongly reducing variants of the Krivine abstract machine. In Danvy [30]. To appear. Journal version of [20].
- [22] Pierre-Louis Curien. *Categorical Combinators, Sequential Algorithms and Functional Programming*, volume 1 of *Research Notes in Theoretical Computer Science*. Pitman, 1986.
- [23] Pierre-Louis Curien. An abstract framework for environment machines. *Theoretical Computer Science*, 82:389–402, 1991.
- [24] Pierre-Louis Curien, Thérèse Hardin, and Jean-Jacques Lévy. Confluence properties of weak and strong calculi of explicit substitutions. *Journal of the ACM*, 43(2):362–397, 1996.
- [25] Olivier Danvy. Formalizing implementation strategies for first-class continuations. In Gert Smolka, editor, *Proceedings of the Ninth European Symposium on Programming*, number 1782 in *Lecture Notes in Computer Science*, pages 88–103, Berlin, Germany, March 2000. Springer-Verlag.

- [26] Olivier Danvy. From reduction-based to reduction-free normalization. In Sergio Antoy and Yoshihito Toyama, editors, *Proceedings of the Fourth International Workshop on Reduction Strategies in Rewriting and Programming (WRS'04)*, number 124 in Electronic Notes in Theoretical Computer Science, pages 79–100, Aachen, Germany, May 2004. Elsevier Science. Invited talk.
- [27] Olivier Danvy. On evaluation contexts, continuations, and the rest of the computation. In Hayo Thielecke, editor, *Proceedings of the Fourth ACM SIGPLAN Workshop on Continuations (CW'04)*, Technical report CSR-04-1, Department of Computer Science, Queen Mary's College, pages 13–23, Venice, Italy, January 2004. Invited talk.
- [28] Olivier Danvy. A rational deconstruction of Landin's SECD machine. In Clemens Grelck, Frank Huch, Greg J. Michaelson, and Phil Trinder, editors, *Implementation and Application of Functional Languages, 16th International Workshop, IFL'04*, number 3474 in Lecture Notes in Computer Science, pages 52–71, Lübeck, Germany, September 2004. Springer-Verlag. Recipient of the 2004 Peter Landin prize. Extended version available as the technical report BRICS RS-03-33.
- [29] Olivier Danvy. Defunctionalized interpreters for higher-order programming languages. In *Preliminary proceedings of the 21st Conference on Mathematical Foundations of Programming Semantics*, Birmingham, UK, May 2005. John Reynolds session.
- [30] Olivier Danvy, editor. *Special Issue on the Krivine Abstract Machine*, Higher-Order and Symbolic Computation. Springer, 2006. In preparation.
- [31] Olivier Danvy and Peter Dybjer, editors. *Proceedings of the 1998 APPSEM Workshop on Normalization by Evaluation (NBE 1998)*, BRICS Note Series NS-98-8, Gothenburg, Sweden, May 1998. BRICS, Department of Computer Science, University of Aarhus.
- [32] Olivier Danvy and Andrzej Filinski. Abstracting control. In Wand [88], pages 151–160.
- [33] Olivier Danvy and Andrzej Filinski. Representing control, a study of the CPS transformation. *Mathematical Structures in Computer Science*, 2(4):361–391, 1992.
- [34] Olivier Danvy and Karoline Malmkjær. Intensions and extensions in a reflective tower. In Cartwright [13], pages 327–341.
- [35] Olivier Danvy and Lasse R. Nielsen. Defunctionalization at work. In Harald Søndergaard, editor, *Proceedings of the Third International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP'01)*, pages 162–174, Firenze, Italy, September 2001. ACM Press. Extended version available as the technical report BRICS RS-01-23.
- [36] Olivier Danvy and Lasse R. Nielsen. Refocusing in reduction semantics. Research Report BRICS RS-04-26, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, November 2004. A preliminary version appears in the informal proceedings of the Second International Workshop on Rule-Based Programming (RULE 2001), Electronic Notes in Theoretical Computer Science, Vol. 59.4.
- [37] Olivier Danvy and Zhe Yang. An operational investigation of the CPS hierarchy. In S. Doaitse Swierstra, editor, *Proceedings of the Eighth European Symposium on*

- Programming*, number 1576 in Lecture Notes in Computer Science, pages 224–242, Amsterdam, The Netherlands, March 1999. Springer-Verlag.
- [38] Philippe de Groote. An environment machine for the lambda-mu-calculus. *Mathematical Structures in Computer Science*, 8:637–669, 1998.
 - [39] R. Kent Dybvig, Simon Peyton-Jones, and Amr Sabry. A monadic framework for subcontinuations. Technical Report 615, Computer Science Department, Indiana University, Bloomington, Indiana, June 2005.
 - [40] Matthias Felleisen. *The Calculi of λ -v-CS Conversion: A Syntactic Theory of Control and State in Imperative Higher-Order Programming Languages*. PhD thesis, Computer Science Department, Indiana University, Bloomington, Indiana, August 1987.
 - [41] Matthias Felleisen. The theory and practice of first-class prompts. In Jeanne Ferrante and Peter Mager, editors, *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Programming Languages*, pages 180–190, San Diego, California, January 1988. ACM Press.
 - [42] Matthias Felleisen and Matthew Flatt. Programming languages and lambda calculi. Unpublished lecture notes. <http://www.ccs.neu.edu/home/matthias/3810-w02/readings.html>, 1989-2003.
 - [43] Matthias Felleisen and Daniel P. Friedman. Control operators, the SECD machine, and the λ -calculus. In Martin Wirsing, editor, *Formal Description of Programming Concepts III*, pages 193–217. Elsevier Science Publishers B.V. (North-Holland), Amsterdam, 1986.
 - [44] Matthias Felleisen, Daniel P. Friedman, Bruce Duba, and John Merrill. Beyond continuations. Technical Report 216, Computer Science Department, Indiana University, Bloomington, Indiana, February 1987.
 - [45] Matthias Felleisen, Mitchell Wand, Daniel P. Friedman, and Bruce F. Duba. Abstract continuations: A mathematical semantics for handling full functional jumps. In Cartwright [13], pages 52–62.
 - [46] Andrzej Filinski. Representing layered monads. In Alex Aiken, editor, *Proceedings of the Twenty-Sixth Annual ACM Symposium on Principles of Programming Languages*, pages 175–188, San Antonio, Texas, January 1999. ACM Press.
 - [47] Cédric Fournet and Andrew D. Gordon. Stack inspection: Theory and variants. *ACM Transactions on Programming Languages and Systems*, 25(3):360–399, May 2003.
 - [48] Steven E. Ganz, Daniel P. Friedman, and Mitchell Wand. Trampolined style. In Peter Lee, editor, *Proceedings of the 1999 ACM SIGPLAN International Conference on Functional Programming*, pages 18–27, Paris, France, September 1999. ACM Press.
 - [49] Timothy G. Griffin. A formulae-as-types notion of control. In Paul Hudak, editor, *Proceedings of the Seventeenth Annual ACM Symposium on Principles of Programming Languages*, pages 47–58, San Francisco, California, January 1990. ACM Press.
 - [50] Carl Gunter, Didier Rémy, and Jon G. Riecke. A generalization of exceptions and control in ML-like languages. In Simon Peyton Jones, editor, *Proceedings of the Seventh ACM Conference on Functional Programming and Computer Architecture*, pages 12–23, La Jolla, California, June 1995. ACM Press.

- [51] Thérèse Hardin, Luc Maranget, and Bruno Pagano. Functional runtime systems within the lambda-sigma calculus. *Journal of Functional Programming*, 8(2):131–172, 1998.
- [52] Robert Harper, Bruce F. Duba, and David MacQueen. Typing first-class continuations in ML. *Journal of Functional Programming*, 3(4):465–484, October 1993.
- [53] Christopher T. Haynes, Daniel P. Friedman, and Mitchell Wand. Continuations and coroutines. In Guy L. Steele Jr., editor, *Conference Record of the 1984 ACM Symposium on Lisp and Functional Programming*, pages 293–298, Austin, Texas, August 1984. ACM Press.
- [54] Carl Hewitt. Control structure as patterns of passing messages. In Patrick Henry Winston and Richard Henry Brown, editors, *Artificial Intelligence: An MIT Perspective*, volume 2, pages 434–465. The MIT Press, 1979.
- [55] Robert Hieb, R. Kent Dybvig, and Carl Bruggeman. Representing control in the presence of first-class continuations. In Bernard Lang, editor, *Proceedings of the ACM SIGPLAN'90 Conference on Programming Languages Design and Implementation*, SIGPLAN Notices, Vol. 25, No 6, pages 66–77, White Plains, New York, June 1990. ACM Press.
- [56] Yuki Yoshi Kameyama. Axioms for delimited continuations in the CPS hierarchy. In Jerzy Marcinkowski and Andrzej Tarlecki, editors, *Computer Science Logic, 18th International Workshop, CSL 2004, 13th Annual Conference of the EACSL, Proceedings*, volume 3210 of *Lecture Notes in Computer Science*, pages 442–457, Karpacz, Poland, September 2004. Springer.
- [57] Richard Kelsey, William Clinger, and Jonathan Rees, editors. Revised⁵ report on the algorithmic language Scheme. *Higher-Order and Symbolic Computation*, 11(1):7–105, 1998.
- [58] Oleg Kiselyov. How to remove a dynamic prompt: Static and dynamic delimited continuation operators are equally expressible. Technical Report 611, Computer Science Department, Indiana University, Bloomington, Indiana, March 2005.
- [59] Jean-Louis Krivine. Un interprète du λ -calcul. Brouillon. Available online at <http://www.pps.jussieu.fr/~krivine/>, 1985.
- [60] Jean-Louis Krivine. A call-by-name lambda-calculus machine. In Danvy [30]. To appear. Available online at <http://www.pps.jussieu.fr/~krivine/>.
- [61] Peter J. Landin. The mechanical evaluation of expressions. *The Computer Journal*, 6(4):308–320, 1964.
- [62] Peter J. Landin. A generalization of jumps and labels. Technical report, UNIVAC Systems Programming Research, 1965. Reprinted in *Higher-Order and Symbolic Computation* 11(2):125–143, 1998, with a foreword [84].
- [63] John Launchbury. A natural semantics for lazy evaluation. In Susan L. Graham, editor, *Proceedings of the Twentieth Annual ACM Symposium on Principles of Programming Languages*, pages 144–154, Charleston, South Carolina, January 1993. ACM Press.

- [64] Xavier Leroy. The Zinc experiment: an economical implementation of the ML language. Rapport Technique 117, INRIA Rocquencourt, Le Chesnay, France, February 1990.
- [65] Pierre Lescanne. From $\lambda\sigma$ to $\lambda\nu$ a journey through calculi of explicit substitutions. In Hans-J. Boehm, editor, *Proceedings of the Twenty-First Annual ACM Symposium on Principles of Programming Languages*, pages 60–69, Portland, Oregon, January 1994. ACM Press.
- [66] Simon Marlow and Simon L. Peyton Jones. Making a fast curry: push/enter vs. eval/apply for higher-order languages. In Kathleen Fisher, editor, *Proceedings of the 2004 ACM SIGPLAN International Conference on Functional Programming*, SIGPLAN Notices, Vol. 39, No. 9, pages 4–15, Snowbird, Utah, September 2004. ACM Press.
- [67] John McCarthy. Recursive functions of symbolic expressions and their computation by machine, part I. *Communications of the ACM*, 3(4):184–195, 1960.
- [68] Peter D. Mosses. A foreword to ‘Fundamental concepts in programming languages’. *Higher-Order and Symbolic Computation*, 13(1/2):7–9, 2000.
- [69] Michel Parigot. $\lambda\mu$ -calculus: an algorithmic interpretation of classical natural deduction. In Andrei Voronkov, editor, *Proceedings of the International Conference on Logic Programming and Automated Reasoning*, number 624 in Lecture Notes in Artificial Intelligence, pages 190–201, St. Petersburg, Russia, July 1992. Springer-Verlag.
- [70] Gordon D. Plotkin. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science*, 1:125–159, 1975.
- [71] François Pottier, Christian Skalka, and Scott Smith. A systematic approach to static access control. *ACM Transactions on Programming Languages and Systems*, 27(2), 2005.
- [72] John C. Reynolds. Definitional interpreters for higher-order programming languages. In *Proceedings of 25th ACM National Conference*, pages 717–740, Boston, Massachusetts, 1972. Reprinted in *Higher-Order and Symbolic Computation* 11(4):363–397, 1998, with a foreword [74].
- [73] John C. Reynolds. The discoveries of continuations. *Lisp and Symbolic Computation*, 6(3/4):233–247, 1993.
- [74] John C. Reynolds. Definitional interpreters revisited. *Higher-Order and Symbolic Computation*, 11(4):355–361, 1998.
- [75] Kristoffer H. Rose. Explicit substitution – tutorial & survey. BRICS Lecture Series LS-96-3, DAIMI, Department of Computer Science, University of Aarhus, Aarhus, Denmark, September 1996.
- [76] Erik Sandewall. An early use of continuations and partial evaluation for compiling rules written in FOPC. *Higher-Order and Symbolic Computation*, 12(1):105–113, 1999.

- [77] Chung-chieh Shan. Shift to control. In Olin Shivers and Oscar Waddell, editors, *Proceedings of the 2004 ACM SIGPLAN Workshop on Scheme and Functional Programming*, Technical report TR600, Computer Science Department, Indiana University, Snowbird, Utah, September 2004.
- [78] Dorai Sitaram and Matthias Felleisen. Control delimiters and their hierarchies. *Lisp and Symbolic Computation*, 3(1):67–99, January 1990.
- [79] Brian C. Smith. Reflection and semantics in Lisp. In Ken Kennedy, editor, *Proceedings of the Eleventh Annual ACM Symposium on Principles of Programming Languages*, pages 23–35, Salt Lake City, Utah, January 1984. ACM Press.
- [80] Guy L. Steele Jr. Rabbit: A compiler for Scheme. Master’s thesis, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, May 1978. Technical report AI-TR-474.
- [81] Christopher Strachey. Fundamental concepts in programming languages. International Summer School in Computer Programming, Copenhagen, Denmark, August 1967. Reprinted in *Higher-Order and Symbolic Computation* 13(1/2):11–49, 2000, with a foreword [68].
- [82] Christopher Strachey and Christopher P. Wadsworth. Continuations: A mathematical semantics for handling full jumps. Technical Monograph PRG-11, Oxford University Computing Laboratory, Programming Research Group, Oxford, England, 1974. Reprinted in *Higher-Order and Symbolic Computation* 13(1/2):135–152, 2000, with a foreword [87].
- [83] Gerald J. Sussman and Guy L. Steele Jr. SCHEME: An interpreter for extended lambda calculus. AI Memo 349, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, December 1975. Reprinted in *Higher-Order and Symbolic Computation* 11(4):405–439, 1998, with a foreword.
- [84] Hayo Thielecke. An introduction to Landin’s “A generalization of jumps and labels”. *Higher-Order and Symbolic Computation*, 11(2):117–124, 1998.
- [85] Jean Vuillemin. Correct and optimal implementations of recursion in a simple programming language. *Journal of Computer and System Sciences*, 9(3):332–354, 1974.
- [86] Philip Wadler. The essence of functional programming (invited talk). In Andrew W. Appel, editor, *Proceedings of the Nineteenth Annual ACM Symposium on Principles of Programming Languages*, pages 1–14, Albuquerque, New Mexico, January 1992. ACM Press.
- [87] Christopher P. Wadsworth. Continuations revisited. *Higher-Order and Symbolic Computation*, 13(1/2):131–133, 2000.
- [88] Mitchell Wand, editor. *Proceedings of the 1990 ACM Conference on Lisp and Functional Programming*, Nice, France, June 1990. ACM Press.
- [89] Mitchell Wand and Daniel P. Friedman. The mystery of the tower revealed: A non-reflective description of the reflective tower. *Lisp and Symbolic Computation*, 1(1):11–38, May 1988.

Recent BRICS Report Series Publications

- RS-05-38** Małgorzata Biernacka and Olivier Danvy. *A Syntactic Correspondence between Context-Sensitive Calculi and Abstract Machines*. December 2005. iii+39 pp. Revised version of BRICS RS-05-22.
- RS-05-37** Gerth Stølting Brodal, Kanela Kaligosi, Irit Katriel, and Martin Kutz. *Faster Algorithms for Computing Longest Common Increasing Subsequences*. December 2005. 16 pp.
- RS-05-36** Dariusz Biernacki, Olivier Danvy, and Chung-chieh Shan. *On the Static and Dynamic Extents of Delimited Continuations*. December 2005. ii+33 pp. To appear in the journal *Science of Computer Programming*. Supersedes BRICS RS-05-13.
- RS-05-35** Kristian Støvring. *Extending the Extensional Lambda Calculus with Surjective Pairing is Conservative*. November 2005. 19 pp. To appear in *Logical Methods in Computer Science*.
- RS-05-34** Henning Korsholm Rohde. *Formal Aspects of Polyvariant Specialization*. November 2005. 27 pp.
- RS-05-33** Luca Aceto, Willem Jan Fokkink, Anna Ingólfssdóttir, and Sumit Nain. *Bisimilarity is not Finitely Based over BPA with Interrupt*. October 2005. 33 pp. This paper supersedes BRICS Report RS-04-24. An extended abstract of this paper appeared in *Algebra and Coalgebra in Computer Science, 1st Conference, CALCO 2005*, Swansea, Wales, 3–6 September 2005, Lecture Notes in Computer Science 3629, pp. 54–68, Springer-Verlag, 2005.
- RS-05-32** Anders Møller, Mads Østerby Olesen, and Michael I. Schwartzbach. *Static Validation of XSL Transformations*. October 2005. 50 pp.
- RS-05-31** Christian Kirkegaard and Anders Møller. *Type Checking with XML Schema in XACT*. September 2005. 20 pp.
- RS-05-30** Karl Krukow. *An Operational Semantics for Trust Policies*. September 2005. 38 pp.
- RS-05-29** Olivier Danvy and Henning Korsholm Rohde. *On Obtaining the Boyer-Moore String-Matching Algorithm by Partial Evaluation*. September 2005. ii+9 pp. To appear in *Information Processing Letters*. This version supersedes BRICS RS-05-14.