

Basic Research in Computer Science

BRICS RS-96-34

Hatcliff & Danvy: A Computational Formalization for Partial Evaluation

A Computational Formalization for Partial Evaluation

(Extended Version)

John Hatcliff
Olivier Danvy

BRICS Report Series

RS-96-34

ISSN 0909-0878

October 1996

**Copyright © 1996, 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 publications in the BRICS
Report Series. 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 World Wide
Web and anonymous FTP:**

`http://www.brics.dk/
ftp://ftp.brics.dk/`

This document in subdirectory RS/96/34/

A Computational Formalization for Partial Evaluation (extended version) *

John Hatcliff

Olivier Danvy

BRICS[†]

Oklahoma State University [‡]

University of Aarhus [§]

Abstract

We formalize a partial evaluator for Eugenio Moggi’s computational metalanguage. This formalization gives an evaluation-order independent view of binding-time analysis and program specialization, including a proper treatment of call unfolding, and enables us to express the essence of “control-based binding-time improvements” for let expressions. Specifically, we prove that the binding-time improvements given by “continuation-based specialization” can be expressed in the metalanguage via monadic laws.

* A shorter version of this paper (essentially, a version without the proofs of Appendix B) will appear in a special issue of *Mathematical Structures in Computer Science* devoted to the 1995 Workshop on Logic, Domains, and Programming Languages (LDPL’95) at which this work was presented. The essential results of this work obtained during the fall of 1994 and refined throughout 1995 and 1996 while the first author was at the University of Copenhagen.

[†]Basic Research in Computer Science,
Centre of the Danish National Research Foundation.

[‡]Department of Computer Science
219 Mathematical Sciences, Stillwater, OK 74078-1053, USA.
E-mail: hatcliff@a.cs.okstate.edu

[§]Department of Computer Science
Ny Munkegade, Building 540, DK-8000 Aarhus C, Denmark.
E-mail: danvy@brics.dk

1 Introduction

Partial evaluation is a program-transformation technique for specializing programs, based on propagating constant values and folding constant expressions [11, 42]. Over the last ten years, it has been organized as a two-phase process: binding-time analysis and program specialization. Binding-time analysis classifies which parts of the source program can be computed statically, i.e., at partial-evaluation time. Program specialization carries out these static computations and residualizes the other computational steps. This organization in two phases thus makes it clear that the more parts of a source program are static, the more this source program is specialized. Over the last five years, various pre-transformations have been investigated that “improve binding times” (making binding-time analysis classify more parts are static) and thus increase specialization.

Our goal is to formalize both the two-phase process of partial evaluation and binding-time improvements, using Moggi’s computational meta-language.

Moggi’s computational metalanguage distinguishes *values* (terms with no remaining computation steps) and *computations* (terms with remaining computation steps). This makes it possible to express a variety of evaluation strategies (including call-by-name and call-by-value). In addition, programs can be parameterized with various notions of computations, expressed as computational monads. Monads are often claimed to make reasoning about programs easier — but very few applications actually exploit the monadic laws.

In this paper, we illustrate that the computational meta-language can be a useful framework for partial evaluation. It allows a clear distinction between static computation steps (to be performed at specialization time) and dynamic computation steps (to be residualized at specialization time and performed at run time). Moreover, it allows an evaluation-order independent view of binding-time analysis and program specialization.

To this end, we present a PCF-like version of the computational metalanguage, and give it a structural operational semantics. We specify binding-time analysis as a non-standard type inference, and then a specializer using structural operational semantics. We prove the correctness of the binding-time analysis and of the specializer.

The computational metalanguage enables us to formalize existing techniques in partial evaluation — namely the linguistic device of `let` insertion, and the partial-evaluation techniques of control-based binding-time improve-

ment and continuation-based specialization, which we achieve by incorporating the monadic laws into our binding-time analysis and our specializer. We prove the equivalence between a continuation-based partial evaluator and our partial evaluator that applies the monadic laws. This formalizes binding-time improvements independently of any particular evaluation order.

We believe that the computational metalanguage also enables new insights and techniques, e.g., to process computational effects. Let us first review each of these points, before outlining the rest of the paper.

1.1 Call unfolding and computation duplication

A partial evaluator unfolds function calls. Call unfolding is thus necessary but in general it is unsound under call-by value. It is necessary to expose opportunities for constant propagation and folding. It is unsound under call-by-value because it may alter termination properties and duplicate computations. Since Similix [7], virtually all call-by-value partial evaluators insert a `let` expression for every dynamic (i.e., unknown) parameter, at each unfolding point. This insertion ensures sound call unfolding: termination properties are preserved and both code and computation duplications are avoided.

Moggi’s computational metalanguage specifies the order of computations through `let` expressions. By expressing source programs in this metalanguage, `let` expressions appear naturally. By residualizing these `let` expressions, sound call unfolding is achieved naturally.

For example, the call-by-value program

$$\lambda z. (\lambda(x, y). (x + 1) - (y + y)) @ (10, z \times z)$$

can be specialized as follows. The inner β -redex can be reduced (noting application with the infix operator “@”), and the left-most addition can be computed. The residual program reads:

$$\lambda z. 11 - ((z \times z) + (z \times z))$$

It is unsatisfactory because the computation of $z \times z$ has been duplicated. As mentioned above [7], call-by-value partial evaluators insert residual `let` expressions to name dynamic expressions that should not be duplicated. With such a strategy, the residual program reads:

$$\lambda z. \text{let } a = z \times z \text{ in } 11 - (a + a)$$

This solution of inserting let expressions is effective, but ad hoc, in that there were no corresponding let expressions in the source program. In contrast, **let** expressions are an integral part of Moggi’s computational metalanguage. Using this metalanguage, the source program can be rewritten as follows.

$$\begin{aligned} &\lambda z . \text{let } a \Leftarrow z \times z \\ &\quad \text{in } (\lambda(x, y) . \text{let } v_1 \Leftarrow x + 1 \\ &\quad \quad \text{in let } v_2 \Leftarrow y + y \\ &\quad \quad \quad \text{in } v_1 - v_2) @ (10, a) \end{aligned}$$

Specializing this program amounts to β -reducing function applications (which is *always* sound because their argument is a value, not a computation), δ -reducing static operations, and unfolding **let** expressions that bind a unit computation.

Presently, the outer **let** expression is not unfolded because z is dynamic (its value is unknown) and the function application is β -reduced. This leads to:

$$\begin{aligned} &\lambda z . \text{let } a \Leftarrow z \times z \\ &\quad \text{in let } v_1 \Leftarrow 10 + 1 \\ &\quad \quad \text{in let } v_2 \Leftarrow a + a \\ &\quad \quad \quad \text{in } v_1 - v_2 \end{aligned}$$

The static addition is performed, which leads to:

$$\begin{aligned} &\lambda z . \text{let } a \Leftarrow z \times z \\ &\quad \text{in let } v_1 \Leftarrow \text{unit } 11 \\ &\quad \quad \text{in let } v_2 \Leftarrow a + a \\ &\quad \quad \quad \text{in } v_1 - v_2 \end{aligned}$$

The **let** expression declaring v_1 is unfolded, and there is no other opportunity for static reduction. The result essentially coincides with the residual program above.

1.2 Control-based binding-time improvements

The structure of source programs influences the precision of binding-time analysis and thus the effectiveness of program specialization. A panoply of “binding-time improvements” aiming at restructuring source programs has been developed [42, Chapter 12] but they have not been formalized.

Essentially, a partial evaluator propagates static values into static contexts, where this propagation gives rise to a static computation. Binding-time analysis thus refines the usual notions of data flow and control flow into

a static and a dynamic data flow, and a static and a dynamic control flow. Data-flow and control-flow binding-time improvements respectively amount to improve the static data flow and the static control flow of a source program. The data-flow aspects have been recently clarified [17, 18], but the control-flow aspects still remain a research topic [6, 8, 10, 39, 45].

The program calculus for Moggi’s computational metalanguage includes *monadic laws*. Let us show how a partial evaluator applying monadic laws to restructure source programs captures the essence of control-based binding-time improvements.

For example, the program

$$\lambda f . ((\lambda x . 2) @ (f @ 1)) + 3$$

can be specialized as follows. The inner β -redex can be reduced but because its argument is dynamic, a let expression must be inserted. The result reads:

$$\lambda f . (\text{let } x = f @ 1 \text{ in } 2) + 3$$

It would be unsound to unfold the let expression and discard $f @ 1$ from the residual program because f could have a computational effect (e.g., divergence). A partial evaluator needs to perform some contorsions to move the context of the let expression $[\cdot] + 3$ in its body, yielding

$$\lambda f . \text{let } x = f @ 1 \text{ in } 5$$

Such restructurings are referred to as “control-based binding-time improvements” since they alter the control flow of the program so as to improve binding times. They are usually achieved by maintaining an explicit representation of control, using continuations. Specializers incorporating these techniques are known as “continuation-based specializers” [6, 42, 45]. Essentially they mimic one-pass CPS transformations [15].

We formalize this technique with the associativity of let expressions in Moggi’s computational metalanguage. Encoding the term above in the metalanguage yields:

$$\begin{aligned} \lambda f . \text{let } v_1 \Leftarrow \text{let } v_2 \Leftarrow f @ 1 \\ \text{in } (\lambda x . \text{unit } 2) @ v_2 \\ \text{in } v_1 + 3 \end{aligned}$$

Static reduction of the β -redex yields:

$$\begin{aligned} \lambda f . \text{let } v_1 \Leftarrow \text{let } v_2 \Leftarrow f @ 1 \\ \text{in unit } 2 \\ \text{in } v_1 + 3 \end{aligned}$$

Reassociating the `let` expression yields:

$$\begin{aligned} &\lambda f . \text{let } v_2 \Leftarrow f @ 1 \\ &\quad \text{in let } v_1 \Leftarrow \text{unit } 2 \\ &\quad \quad \text{in } v_1 + 3 \end{aligned}$$

Unfolding the inner `let` and statically reducing the addition yields:

$$\begin{aligned} &\lambda f . \text{let } v_2 \Leftarrow f @ 1 \\ &\quad \text{in unit } 5 \end{aligned}$$

The result essentially coincides with the residual program above.

1.3 Computational effects

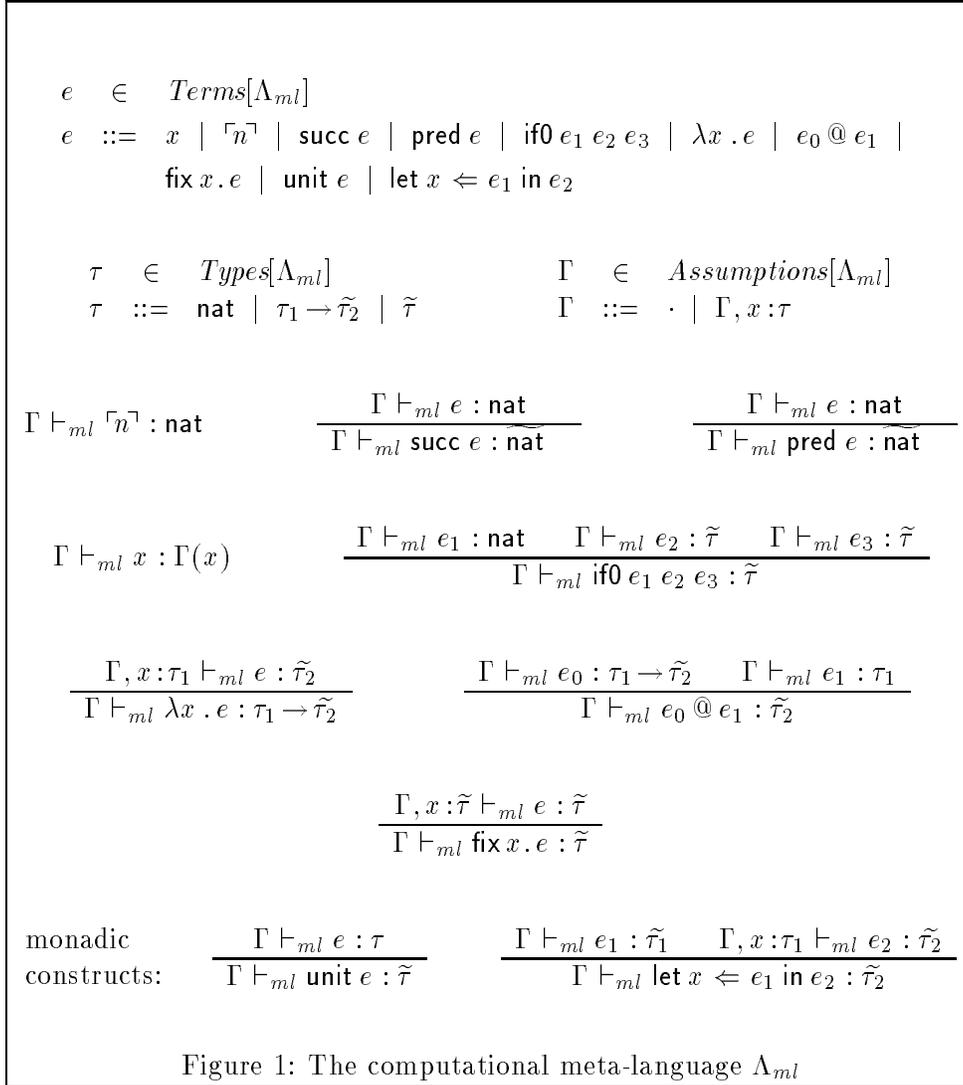
Since Similix [7], partial evaluators classify computational effects (I/O, state, etc.) as dynamic computations. The specializer’s only duty is to maintain their order and to ensure that none disappears or is duplicated.

The computational metalanguage’s *raison d’être* is to provide a modular specification of computational effects. Using this medium, we originally planned to provide a sound treatment of side-effects that would be more effective than systematic residualization. We envisioned for example to split the store in regions: some side-effecting operations could then be classified as static and performed statically (e.g., a symbol table in an interpreter could be processed statically in a compiler) while others would be classified as dynamic and delayed until runtime (e.g., an error message associated with division by 0). This line of work is currently being pursued independently by Dussart and Thiemann [22].

1.4 This paper

The rest of this paper is organized as follows. Section 2 presents a PCF-like version of the computational metalanguage, Λ_{ml} . We give it a structural operational semantics. Section 3 presents an offline partial evaluator for Λ_{ml} . Section 4 addresses call unfolding. In Section 5, we return to the monadic laws and incorporate them into the specializer — adjusting the binding-time analysis accordingly. In Section 6, we prove the equivalence between a continuation-based partial evaluator and our partial evaluator that applies the monadic laws. After a review of related work in Section 7, Section 8 concludes.

Appendix A gives the details for translating call-by-name and call-by-value PCF into Λ_{ml} . Appendix B gives proof details.



2 The computational meta-language

2.1 Syntax

Figure 1 presents the language Λ_{ml} based on Moggi's *computational meta-language* [48].¹ The typing system of Λ_{ml} captures the distinction be-

¹We have added **fix**, **succ**, **pred**, and **if0** to Moggi's published description [48].

tween *values* (terms with no remaining computation steps) and *computations* (terms with remaining computation steps). Types of the form \mathbf{nat} and $\tau_1 \rightarrow \tau_2$ are *value types*. Accordingly, the rules for numerals and abstractions belong to the introduction rules for value types. Types of the form $\tilde{\tau}$ are *computation types*. All functions have computational co-domains. Thus, all applications have a computation type — capturing the fact that evaluating a function application always requires one or more computational steps.

The *monadic constructs* are used to make the computational process explicit.² Intuitively, $\mathbf{unit} \ e$ is a trivial computation that simply yields the value of e . $\mathbf{let} \ x \leftarrow e_1 \ \mathbf{in} \ e_2$ forces the evaluation of e_1 . If that evaluation terminates, the resulting value is substituted for x in e_2 , and evaluation continues with the modified version of e_2 .

We identify terms up to renaming of bound variables (i.e., up to α -equivalence) and use standard notation and conventions for substitution, free variables, contexts, etc. [2]. We write $e_1 \equiv e_2$ when e_1 and e_2 are α -equivalent.

We represent simultaneous substitutions $e[x_1 := e_1, \dots, x_n := e_n]$ using a substitution function $\sigma : \mathit{Terms}[\Lambda_{ml}] \rightarrow \mathit{Terms}[\Lambda_{ml}]$ whose application is denoted $e\sigma$. A substitution σ is *closed* if, for all $x \in \text{dom } \sigma$, $x\sigma$ is a closed term.

We write $\Gamma \vdash_{ml} e_1, e_2 : \tau$ when both $\Gamma \vdash_{ml} e_1 : \tau$ and $\Gamma \vdash_{ml} e_2 : \tau$. A closed substitution σ is *compatible* with an assumption Γ , if for all $x \in \text{dom } \Gamma$, $\cdot \vdash_{ml} x\sigma : \Gamma(x)$. *Programs* are closed terms with type $\widetilde{\mathbf{nat}}$.

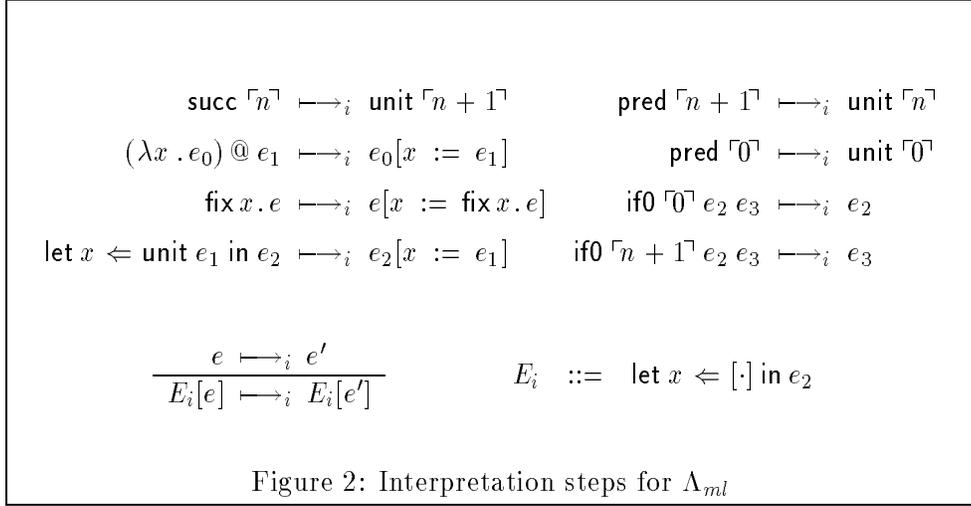
2.2 Operational semantics

Figure 2 presents single-step evaluation rules for Λ_{ml} .³ Axioms such as $\mathbf{succ} \ \lceil n \rceil \mapsto_i \ \mathbf{unit} \ \lceil n + 1 \rceil$ define basic computation steps. The single inference rule describes contexts in which evaluation steps may occur.

The following lemma gives evaluation properties for closed terms at each type τ . The intuition is that each well-typed term (a) is a canonical term of

²The exact connection to the structure of a monad can be found in Moggi’s original work [48, page 61].

³Moggi originally gave a categorical interpretation for his computational meta-language [48]. Crole and Pitts extended Moggi’s work with a fix-point operator and an associated logic [13]. Later, Gordon developed an elegant operational theory for the meta-language with a fixpoint operator and inductive and coinductive types [31]. Λ_{ml} is basically a sub-language of Gordon’s — except that we include directly type \mathbf{nat} and associated operations whereas Gordon constructs them via inductive types. Here, we follow Gordon and our previous work [34, 36] and present a structural operational semantics for Λ_{ml} .



the corresponding type, or (b) can undergo an evaluation step that preserves typing. In particular, there are no “stuck” [56, p. 151] terms.⁴

Lemma 1 (interpretation properties)

1. If $\cdot \vdash_{ml} e : \mathbf{nat}$ then $e \equiv \ulcorner n \urcorner$ for some number n .
2. If $\cdot \vdash_{ml} e : \tau_1 \rightarrow \tilde{\tau}_2$ then $e \equiv \lambda x . e'$ and $x : \tau_1 \vdash_{ml} e' : \tilde{\tau}_2$.
3. If $\cdot \vdash_{ml} e : \tilde{\tau}$ then exactly one of the following statements holds:
 - (a) $e = \text{unit } e'$ and $\cdot \vdash_{ml} e' : \tau$.
 - (b) $e \mapsto_i e'$ and $\cdot \vdash_{ml} e' : \tilde{\tau}$.

Proof: by induction over the height of the derivation of $\cdot \vdash_{ml} e : \tau$ relying on the property that if $\Gamma, x : \tau_1 \vdash_{ml} e_0 : \tau_0$ and $\Gamma \vdash_{ml} e_1 : \tau_1$ then $\Gamma \vdash_{ml} e_0[x := e_1] : \tau_0$. See the proof of Lemma 2 (which generalizes the current lemma) given in Appendix B. ■

It is easy to check that $e \mapsto_i e'$ and $e \mapsto_i e''$ implies $e' \equiv e''$. This justifies the definition of the following (partial) function in terms of the reflexive transitive closure of \mapsto_i .

⁴For example, the untypable term $\ulcorner 3 \urcorner @ \ulcorner 2 \urcorner$ is stuck: it cannot undergo an evaluation step and it is not a proper canonical term for any type.

Definition 1 (interpreter)

For $\cdot \vdash_{ml} e : \widetilde{\text{nat}}$,

$$\text{int } e = \lceil n \rceil \quad \text{iff} \quad e \xrightarrow{i}^* \text{unit } \lceil n \rceil$$

We write $\text{int } e \downarrow$ when $\text{int } e$ is defined and $\text{int } e \uparrow$ when $\text{int } e$ is undefined. As a consequence of Lemma 1, $\text{int } e \uparrow$ implies that e heads an infinite sequence of computation steps.

Observing termination of terms in program contexts gives the following notion of operational approximation, which in turn, induces a notion of operational equivalence.⁵

Definition 2 (operational approximation)

For $\Gamma \vdash_{ml} e_1, e_2 : \tau$, and for all contexts C such that $C[e_1]$ and $C[e_2]$ are programs,

$$e_1 \preceq e_2 \quad \text{iff} \quad \text{int } C[e_1] \downarrow \text{ implies } \text{int } C[e_2] \downarrow$$

Definition 3 (operational equivalence)

For $\Gamma \vdash_{ml} e_1, e_2 : \tau$,

$$e_1 \approx e_2 \quad \text{iff} \quad e_1 \preceq e_2 \text{ and } e_2 \preceq e_1$$

Note that if $\cdot \vdash_{ml} e_1, e_2 : \widetilde{\text{nat}}$ and $e_1 \approx e_2$, then $\text{int } e_1$ and $\text{int } e_2$ are both undefined, or else both are defined and there exists a number n such that $\text{int } e_1 \equiv \lceil n \rceil \equiv \text{int } e_2$.

2.3 Equational reasoning

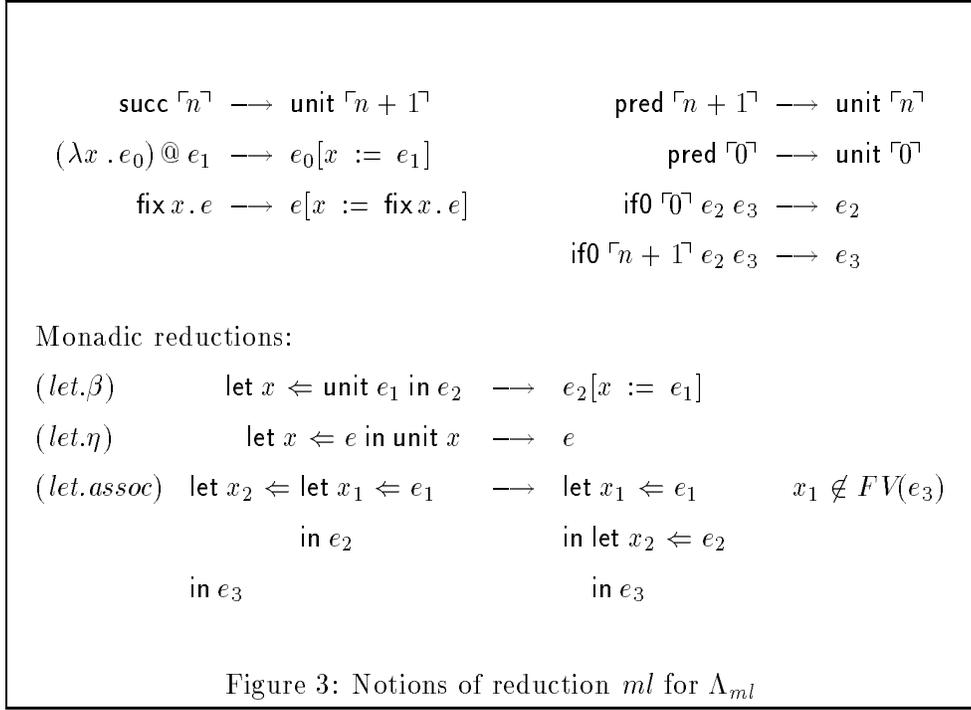
Figure 3 presents notions of reduction for the computational meta-language Λ_{ml} . \longrightarrow also denotes construct-compatible one-step reduction, \longrightarrow^* denotes the reflexive, transitive closure of \longrightarrow , and $=_{ml}$ denotes the smallest equivalence relation generated by \longrightarrow [2].

Theorem 1 (soundness of calculus) For $\Gamma \vdash_{ml} e_1, e_2 : \tau$,

$$e_1 =_{ml} e_2 \Rightarrow e_1 \approx e_2$$

Proof: The proof follows by a straightforward adaptation of Gordon's operational theory [31, Chapters 3,4,5]. ■

⁵It is sufficient to observe termination since one can always distinguish between numerals using conditional expressions.



2.4 Encoding evaluation strategies in Λ_{ml}

A variety of evaluation strategies for a PCF-like language Λ can be encoded in Λ_{ml} due to Λ_{ml} 's explicit distinction between values and computations. In previous work [34, 36], we proved the correctness of an encoding \mathcal{E}_n of a call-by-name version of Λ into Λ_{ml} , and an encoding \mathcal{E}_v of a call-by-value version of Λ into Λ_{ml} . These encodings can be found in Appendix A.2.⁶

In summary, call-by-name Λ functions are encoded as functions from computations to computations; call-by-value Λ functions are encoded as functions from values to computations.⁷ Applications are translated as follows:

$$\text{Call-by-name: } \mathcal{E}_n \langle e_0 @ e_1 \rangle = \text{let } x_0 \Leftarrow \mathcal{E}_n \langle e_0 \rangle \text{ in } x_0 @ \mathcal{E}_n \langle e_1 \rangle$$

$$\text{Call-by-value: } \mathcal{E}_v \langle e_0 @ e_1 \rangle = \text{let } x_0 \Leftarrow \mathcal{E}_v \langle e_0 \rangle \text{ in let } x_1 \Leftarrow \mathcal{E}_v \langle e_1 \rangle \text{ in } x_0 @ x_1$$

⁶These encodings are based on similar encodings given by e.g., Moggi [48] and Wadler [67].

⁷Alternative call-by-value encodings exist that are similar to the encoding of call-by-value procedures in the call-by-name programming language Algol 60 [36, Section 4].

Thus, evaluation of argument expressions is not forced in the call-by-name encoding, but is forced in the call-by-value encoding.

In general (i.e., in all encodings), a **let** is inserted around each computation step — making the computational structure of a program explicit. This property is crucial to the evaluation-order independent treatment of binding-time analysis and program specialization presented in the following section.

3 An offline partial evaluator for Λ_{ml}

A partial evaluator takes a *source program* p and a subset s of p 's input, and produces a *residual program* p_s which is specialized with respect to s . The correctness of the partial evaluator implies that running p_s on p 's remaining input d gives the same result as running p on the complete input s and d . The data s and d are often referred to as *static* and *dynamic* data (respectively) since s is fixed at specialization time whereas one may supply various data d during runs of p_s .

The specialized program p_s is obtained from p by evaluating constructs that depend only on s , while rebuilding constructs that may depend on dynamic data. “Offline” partial evaluation accomplishes this in two phases: (1) a *binding-time analysis* phase, and (2) a *specialization* phase.

1. **Binding-time analysis:** Given assumptions about which program inputs are static and dynamic, binding-time analysis constructs an annotated program where each program construct is annotated with a *specialization directive* and a *specialization type*.
 - **Specialization directives:** A construct is assigned a directive of *eliminable* if it depends only on static data and thus can be completely evaluated during the specialization phase. A construct is assigned a directive of *residual* if it may depend on dynamic data and thus must be reconstructed in the specialization phase.
 - **Specialization types:** The specialization types (*a.k.a. binding times*) are the carriers of information during the analysis phase. They describe the “knownness” or the “unknownness” of expressions. This information is used to determine the specialization directives assigned to constructs. For example, if the argument of a destructor construct has a specialization type indicating that

$$\begin{aligned}
w &\in \text{Terms}[\Lambda_{ml}^{bt}] \\
w &::= x \mid \\
&\quad \ulcorner n \urcorner \mid \text{succ } w \mid \text{pred } w \mid \text{if0 } w_1 w_2 w_3 \mid \lambda x . w \mid w_0 @ w_1 \mid \\
&\quad \text{fix } x . w \mid \text{unit } w \mid \text{let } x \Leftarrow w_1 \text{ in } w_2 \mid \\
&\quad \underline{n} \mid \underline{\text{succ}} w \mid \underline{\text{pred}} w \mid \underline{\text{if0}} w_1 w_2 w_3 \mid \underline{\lambda} x . w \mid w_0 @ \underline{w}_1 \mid \\
&\quad \underline{\text{fix}} x . w \mid \underline{\text{unit}} w \mid \underline{\text{let}} x \Leftarrow w_1 \underline{\text{in}} w_2 \mid \text{lift } w
\end{aligned}$$

Figure 4: The binding-time annotated meta-language Λ_{ml}^{bt}

it is unknown, then that construct cannot be evaluated at specialization time and must be given a *residual* directive.

2. **Specialization:** During the specialization phase, the specializer simply follows the directives assigned during binding-time analysis: eliminable constructs are evaluated (and thus eliminated); residual constructs are reconstructed (and thus appear in the residual program).⁸

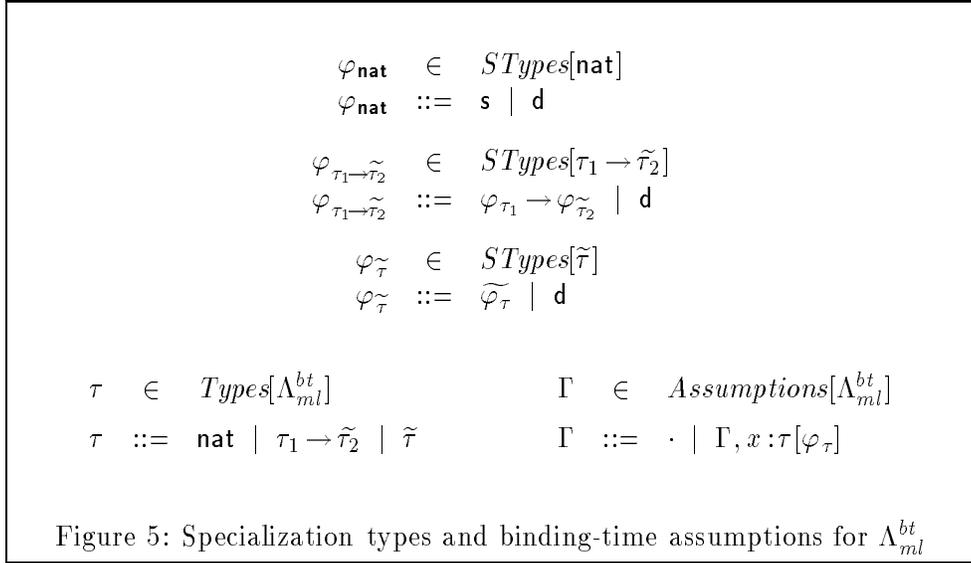
3.1 Binding-time analysis

We first outline how binding-time information is expressed in a program annotated with specialization directives and types. Following this, we give a binding-time logic that determines which annotations are appropriate for a given source language term.

3.1.1 Specialization directives

A binding-time analysis for Λ_{ml} associates each Λ_{ml} term with a term in the annotated language Λ_{ml}^{bt} of Figure 4. Eliminable terms are non-underlined; residual terms are underlined. Identifiers are not annotated since the appropriate information can be determined from the environment. A coercion construct `lift` is added to Λ_{ml}^{bt} to residualize the result of evaluating an elim-

⁸Following other formal treatments of partial evaluation [42, 54, 68], we simplify the presentation by omitting folding and generalization strategies.



inable term. This allows static computation to occur in a residual context.⁹ A term $w \in \Lambda_{ml}^{bt}$ is *completely residual* if it consists of only underlined constructs and identifiers. $Residual\text{-}terms[\Lambda_{ml}^{bt}]$ denotes the set of completely residual terms. Intuitively, the specializer will output completely residual terms — all eliminable constructs will have been evaluated (this will be proved in Section 3.2).

3.1.2 Specialization types

Figure 5 presents a τ -indexed family of specialization types ($STypes$) for Λ_{ml}^{bt} (we omit type indices on specialization types when they can be inferred from the context). A specialization type φ is *dynamic* if $\varphi = \mathbf{d}$; otherwise φ is *static*.

- $\mathbf{s} \in STypes[\mathbf{nat}]$ will tag expressions of type \mathbf{nat} that are guaranteed to evaluate to known data (i.e., numerals).
- $\mathbf{d} \in STypes[\mathbf{nat}]$ will tag expressions of type \mathbf{nat} whose evaluation may depend on unknown data and thus cannot be guaranteed to evaluate

⁹We restrict lifting to base types for simplicity, as we wish to formalize control-flow binding-time improvements. Defining **lift** at higher types would enable data-flow binding-time improvements, as investigated elsewhere [17, 18]. This definition is also interesting on its own [14].

to numerals. However, because `nat` is a value type, one does know that the tagged expression will denote a value when dynamic data is supplied.

- $\varphi_{\tau_1} \rightarrow \varphi_{\tilde{\tau}_2} \in STypes[\tau_1 \rightarrow \tilde{\tau}_2]$ will tag expressions of type $\tau_1 \rightarrow \tilde{\tau}_2$ that are guaranteed to evaluate to known data (i.e., abstractions).
- $d \in STypes[\tau_1 \rightarrow \tilde{\tau}_2]$ will tag expressions of type $\tau_1 \rightarrow \tilde{\tau}_2$ whose evaluation may depend on unknown data and thus cannot be guaranteed to evaluate to an abstraction. However, because $\tau_1 \rightarrow \tilde{\tau}_2$ is a value type, one does know that the tagged expression will denote a value when dynamic data is supplied.
- $\tilde{\varphi} \in STypes[\tilde{\tau}]$ will tag expressions of type $\tilde{\tau}$ that are guaranteed to either diverge or to evaluate to a trivial computation (i.e., `unit e` for some $e \in \Lambda_{ml}$ of type τ).
- $d \in STypes[\tilde{\tau}]$ will tag expressions of type $\tilde{\tau}$ whose evaluation may depend on unknown data and thus cannot be guaranteed to diverge or evaluate to a trivial computation.

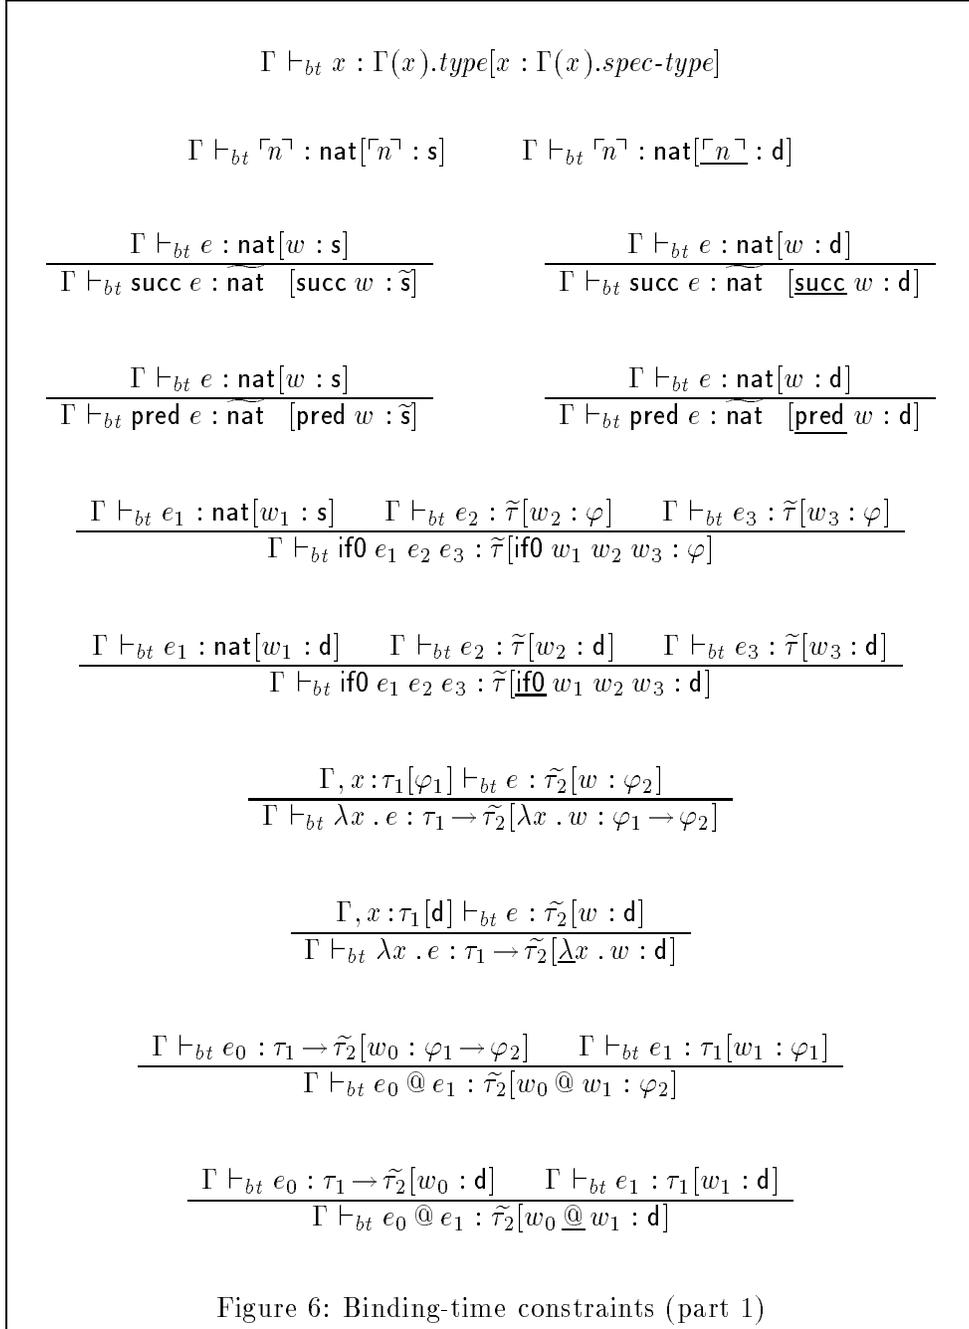
Figure 5 also presents type assumptions for Λ_{ml}^{bt} variables. If $\Gamma, x : \tau[d]$ then x is a *dynamic variable*; if $\Gamma, x : \tau[\varphi]$ where φ is static then x is a *static variable*. $\Gamma(x).type$ and $\Gamma(x).spec\text{-}type$ project types and specialization types from an assumption for $x \in dom \Gamma$.

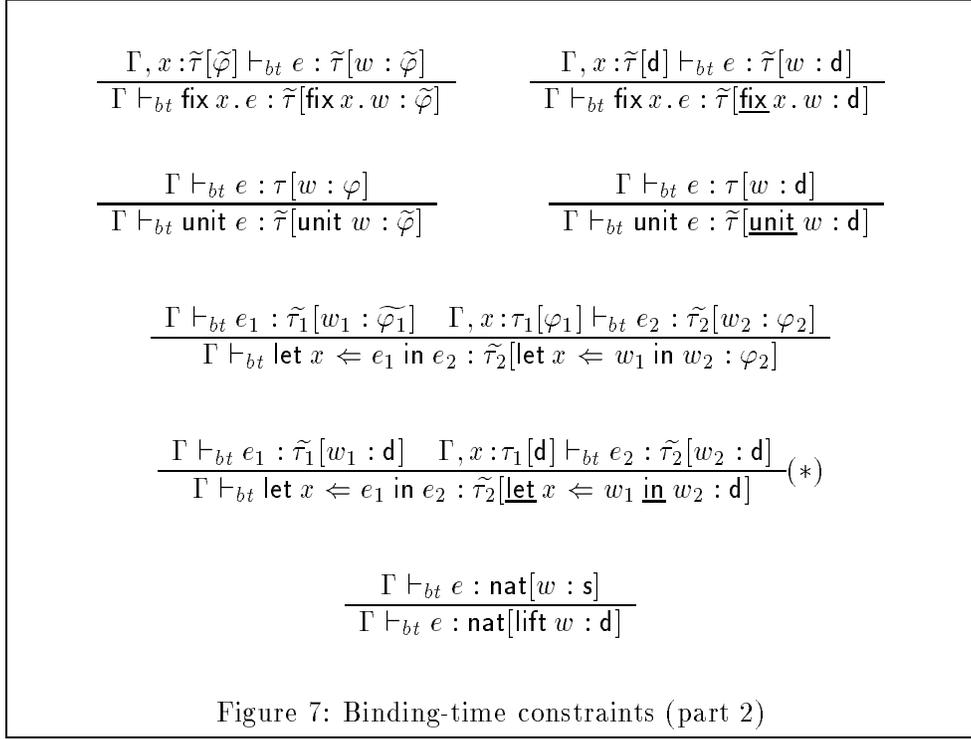
3.1.3 Binding-time analysis specification

Figures 6 and 7 present rules for deriving judgments of the form

$$\Gamma \vdash_{bt} e : \tau[w : \varphi_\tau].$$

Derivable judgements specify constraints that an actual binding-time analysis algorithm must satisfy. Intuitively, if $\Gamma \vdash_{bt} e : \tau[w : \varphi_\tau]$, then given initial binding-time assumptions Γ that indicate which free variables are static or dynamic, a binding-time analysis algorithm maps $e \in Terms[\Lambda_{ml}]$ of type $\tau \in Types[\Lambda_{ml}]$ to a directive annotated term $w \in Terms[\Lambda_{ml}^{bt}]$ of specialization type $\varphi_\tau \in STypes[\tau]$. Specifying the analysis in this way allows one to reason about correctness of the analysis independently of the actual algorithm — proving correctness with respect to the constraints implies that any algorithm satisfying the constraints will be correct.





$[\cdot] : Terms[\Lambda_{ml}^{bt}] \rightarrow Terms[\Lambda_{ml}]$ is an erasing function that removes annotations and `lift` constructs. $[\cdot] : Assumptions[\Lambda_{ml}^{bt}] \rightarrow Assumptions[\Lambda_{ml}]$ simply drops specialization types appearing in assumptions. $[\sigma]$ represents the component-wise application of $[\cdot] : Terms[\Lambda_{ml}^{bt}] \rightarrow Terms[\Lambda_{ml}]$.

Given an assumption $\Gamma \in Assumptions[\Lambda_{ml}^{bt}]$, a closed Λ_{ml}^{bt} substitution σ is *compatible* with Γ if

$$\text{for all } x \in \text{dom } \Gamma, \cdot \vdash_{ml} [x\sigma] : \Gamma(x).type[x\sigma : \Gamma(x).spec\text{-}type].$$

For expressions $e \in Terms[\Lambda_{ml}]$, Γ_{bt} is a *binding-time assumption* for $\Gamma \vdash_{ml} e : \tau$, if $[\Gamma_{bt}] = \Gamma$. $\Gamma_{bt} = \Gamma_s \cup \Gamma_d$ represents the splitting of a binding-time assumption Γ_{bt} into its (necessarily disjoint) static and dynamic variable assumptions.

Definition 4 (binding-time analysis)

A binding-time analysis is a function

$$bta : Assumptions[\Lambda_{ml}^{bt}] \rightarrow Terms[\Lambda_{ml}] \rightarrow Terms[\Lambda_{ml}^{bt}]$$

such that when Γ_{bt} is a binding-time assumption for $\Gamma \vdash_{ml} e : \widetilde{\mathbf{nat}}$,

$$\Gamma_{bt} \vdash_{bt} e : \widetilde{\mathbf{nat}}[(bta \Gamma_{bt} e) : \mathbf{d}].$$

The binding-time analysis rules specify a one-to-many relation between Λ_{ml} terms and Λ_{ml}^{bt} terms. Thus, there are many valid binding-time analysis functions — each varying in practical effectiveness. One usually desires an analysis that gives as many eliminable terms as possible. However, it may be useful to deliberately annotate some eliminable terms as residual to ensure that specialization terminates [42]. Our specification abstracts away from these orthogonal implementation issues.

Conversely, an induction over the derivation of $\Gamma \vdash_{ml} e : \tau[w : \varphi]$ shows that $e \equiv [w]$. Thus, every well-annotated Λ_{ml}^{bt} term has exactly one Λ_{ml} term related to it — the Λ_{ml} term is the erasure of the Λ_{ml}^{bt} term.

3.2 Specialization

To specialize a binding-time analyzed program term w such that

$$\Gamma_{bt} = \Gamma_s \cup \Gamma_d \vdash_{bt} e : \widetilde{\mathbf{nat}}[w : \mathbf{d}]$$

one supplies known data into the static variables via a substitution σ_s compatible with Γ_s . The resulting term $w\sigma_s$ is then specialized using the rules of Figure 8 (which define the single-step specialization function \longrightarrow_s).

- The interpretation rules allow interpreter steps to reduce eliminable (i.e., non-underlined) terms. This corresponds to the idea that partial evaluators are often described as having an interpreter component.
- The axiom `lift` coerces a static numeral to a dynamic numeral. This allows static values of type `nat` to occur in dynamic contexts.
- The compilation rules direct the activities of the specializer in dynamic contexts E_s . This corresponds to the fact that the non-interpretive component of the specializer simply constructs terms that appear in the residual program. In contexts E_s , r ranges over completely residual terms (i.e., terms containing only underlined constructs and free dynamic variables).

The following lemma gives specialization properties for the Λ_{ml}^{bt} terms satisfying judgements of the form $\Gamma_d \vdash_{bt} e : \tau[w : \varphi_\tau]$. The intuition is that

Interpretation rules:

$$\begin{array}{ll}
\text{succ } \ulcorner n \urcorner \mapsto_i \text{unit } \ulcorner n + 1 \urcorner & \text{pred } \ulcorner n + 1 \urcorner \mapsto_i \text{unit } \ulcorner n \urcorner \\
(\lambda x . w_0) @ w_1 \mapsto_i w_0[x := w_1] & \text{pred } \ulcorner 0 \urcorner \mapsto_i \text{unit } \ulcorner 0 \urcorner \\
\text{fix } x . w \mapsto_i w[x := \text{fix } x . w] & \text{if0 } \ulcorner 0 \urcorner w_2 w_3 \mapsto_i w_2 \\
\text{let } x \Leftarrow \text{unit } w_1 \text{ in } w_2 \mapsto_i w_2[x := w_1] & \text{if0 } \ulcorner n + 1 \urcorner w_2 w_3 \mapsto_i w_3
\end{array}$$

$$\frac{w \mapsto_i w'}{E_i[w] \mapsto_i E_i[w']} \quad E_i ::= \text{let } x \Leftarrow [\cdot] \text{ in } w_2$$

$$\frac{w \mapsto_i w'}{w \mapsto_s w'}$$

Compilation rules:

$$\text{lift } \ulcorner n \urcorner \mapsto_s \underline{\ulcorner n \urcorner} \quad \frac{w \mapsto_s w'}{E_s[w] \mapsto_s E_s[w']}$$

$$\begin{aligned}
E_s ::= & \text{pred } [\cdot] \mid \text{succ } [\cdot] \mid \text{if0 } [\cdot] w_2 w_3 \mid \text{if0 } r_1 [\cdot] w_3 \mid \text{if0 } r_1 r_2 [\cdot] \mid \\
& \lambda x . [\cdot] \mid [\cdot] @ w_1 \mid r_0 @ [\cdot] \mid \text{fix } x . [\cdot] \mid \\
& \text{unit } [\cdot] \mid \text{let } x \Leftarrow [\cdot] \text{ in } w_2 \mid \\
& \text{let } x \Leftarrow r_1 \text{ in } [\cdot] \text{ ...where } r_i \in \text{Residual-terms}[\Lambda_{ml}^{bt}]
\end{aligned}$$

Figure 8: Specialization rules for Λ_{ml}^{bt}

each well-annotated term w (a) is a canonical annotated term of the appropriate type, or (b) is completely residual, or (c) can undergo a specialization step which preserves annotations. Moreover, each specialization step $w \mapsto_s w'$ maintains Λ_{ml} convertibility on the corresponding erasures (i.e., $[w] =_{ml} [w']$).

Lemma 2 (specialization properties)

1. If $\Gamma_d \vdash_{bt} e : \text{nat}[w : s]$ then $w \equiv \ulcorner n \urcorner \equiv e$ for some number n .
2. If $\Gamma_d \vdash_{bt} e : \tau_1 \rightarrow \tau_2[w : \varphi_1 \rightarrow \varphi_2]$ then $w \equiv \lambda x . w'$, $e \equiv \lambda x . [w']$ and

- $\Gamma_d, x : \tau_1[\varphi_1] \vdash_{bt} [w'] : \tau_2[w' : \varphi_2]$ for some $w' \in \text{Terms}[\Lambda_{ml}^{bt}]$.
3. If $\Gamma_d \vdash_{bt} e : \tilde{\tau}[w : \tilde{\varphi}]$ then exactly one of the following statements holds:
 - (a) $w \equiv \text{unit } w'$, $e \equiv \text{unit } [w']$ and $\Gamma_d \vdash_{bt} [w'] : \tau[w' : \varphi]$ for some $w' \in \text{Terms}[\Lambda_{ml}^{bt}]$, or
 - (b) $w \mapsto_i w'$ and $\Gamma_d \vdash_{bt} [w'] : \tilde{\tau}[w' : \tilde{\varphi}]$ and $e \equiv [w] \longrightarrow_{ml} [w']$.
 4. If $\Gamma_d \vdash_{bt} e : \text{nat}[w : \mathbf{d}]$ then exactly one of the following statements holds:
 - (a) $w \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$, or
 - (b) $w \mapsto_s w'$, $\Gamma_d \vdash_{bt} [w'] : \tau[w' : \mathbf{d}]$, and $e \equiv [w] \equiv [w'] \equiv \ulcorner n \urcorner$ for some number n .
 5. If $\Gamma_d \vdash_{bt} e : \tau_1 \rightarrow \tau_2[w : \mathbf{d}]$ then exactly one of the following statements holds:
 - (a) $w \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$, or
 - (b) $w \mapsto_s w'$ and $\Gamma_d \vdash_{bt} [w'] : \tau[w' : \mathbf{d}]$ and $e \equiv [w] \longrightarrow_{ml} [w']$.
 6. If $\Gamma_d \vdash_{bt} e : \tilde{\tau}[w : \mathbf{d}]$ then exactly one of the following statements holds:
 - (a) $w \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$, or
 - (b) $w \mapsto_s w'$ and $\Gamma_d \vdash_{bt} [w'] : \tilde{\tau}[w' : \mathbf{d}]$ and $e \equiv [w] \longrightarrow_{ml} [w']$.

Proof: by induction over the height of the derivation of $\Gamma_d \vdash_{bt} e : \tau[w : \varphi]$ relying on the property that if we have both $\Gamma, x : \tau_1[\varphi_1] \vdash_{bt} e_0 : \tau_0[w_0 : \varphi_0]$ and $\Gamma \vdash_{bt} e_1 : \tau_1[w_1 : \varphi_1]$, then $\Gamma \vdash_{bt} e_0[x := e_1] : \tau_0[w_0[x := w_1] : \varphi_0]$. See Appendix B for details. \blacksquare

It is easy to check that if $w \mapsto_s w'$ and $w \mapsto_s w''$ then $w' \equiv w''$. Furthermore, by Lemma 2, well-annotated terms $r \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$ cannot undergo further specialization steps. This justifies the definition of the following (partial) function in terms of the reflexive, transitive closure of \mapsto_s .

Definition 5 (specializer)

For $\Gamma_d \vdash_{bt} e : \text{nat}[w : \mathbf{d}]$,

$$\text{spec } w = r \quad \text{iff} \quad w \mapsto_s^* r$$

where $r \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$.

3.3 Correctness of binding-time analysis and specialization

A binding-time analysis is correct if it always produces sound directives. Directives are unsound if they direct the specialized to attempt the reduction of a non-redex. The last component of Lemma 2 implies that any binding-time analysis satisfying our constraints is correct: when an annotated term (with only dynamic free variables) is specialized, the specialized always finds a redex to contract, or terminates because only residual components are left. In other words, the specialized never “sticks” on a non-redex.

Theorem 2 (correctness of binding-time analysis)

For all $\Gamma_d \vdash_{bt} e : \widetilde{\text{nat}}[w : d]$, exactly one of the following statements holds:

1. $\text{spec } w = r$ where $r \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$, or
2. $\text{spec } w \downarrow$ and w heads an infinite sequence of specialization steps.

Proof: follows from the definition of *spec* (Definition 5) and Lemma 2. See Appendix B for details. ■

This statement of binding-time correctness is analogous to statements of binding-time correctness for the λ -calculus [42, 54, 68].

A specialized is sound if its steps reflect a meaning-preserving transformation of the source program.

Theorem 3 (soundness of specialization)

For all $\Gamma_d \vdash_{bt} e : \widetilde{\text{nat}}[w : d]$,

$$\text{spec } w \downarrow \quad \text{implies} \quad e \approx [\text{spec } w]$$

Proof: Given the definition of *spec* (Definition 5) and Lemma 2, one has $e =_{ml} [\text{spec } w]$ by induction over the number of specialization steps. The soundness of the Λ_{ml} calculus (Theorem 1) then gives the desired result. ■

3.4 Partial evaluation

An offline partial evaluator for Λ_{ml} is obtained by composing binding-time analysis with specialization. When supplied with a source expression e , a binding-time assumption Γ_{bt} , and a collection of “known” data σ_s , the partial evaluator specializes the program to the known data as directed by the binding-time analysis.

Definition 6 (partial evaluator)

Let bta be a correct binding-time analysis, $\Gamma_{bt} = \Gamma_s \cup \Gamma_d$ be a binding-time assumption for $\Gamma \vdash_{ml} e : \widetilde{\mathbf{nat}}$, and σ_s be a closed substitution compatible with Γ_s , then

$$pe\ e\ \Gamma_{bt}\ \sigma_s \stackrel{\text{def}}{=} [spec(bta\ \Gamma_{bt}\ e)\ \sigma_s]$$

The correctness theorem for the partial evaluator is analogous to Kleene's S_n^m -theorem: partial evaluation computes a program specialized to known input data (σ_s) such that running the specialized program on the remaining input data (σ_d) yields a result observationally equivalent to the result of running the source program on the complete input data.

Theorem 4 (correctness of partial evaluation)

Let $\Gamma_{bt} = \Gamma_s \cup \Gamma_d$ be a binding-time assumption for $\Gamma \vdash_{ml} e : \widetilde{\mathbf{nat}}$, let σ_s be a closed substitution compatible with Γ_s , and let σ_d be a closed substitution compatible with $[\Gamma_d]$. If $(pe\ e\ \Gamma_{bt}\ \sigma_s)\downarrow$, then

$$e[\sigma_s]\sigma_d \approx (pe\ e\ \Gamma_{bt}\ \sigma_s)\sigma_d$$

Proof: follows from Theorem 3 and the definition of specialization (Definition 5) and partial evaluation (Definition 6). See Appendix B for details. ■

4 Call unfolding

At the core of a partial evaluator lies call unfolding. This basic transformation enables static information to be propagated across procedure boundaries. Call unfolding is implemented using the copy rule. Thus it is only sound in a call-by-name setting such as Launchbury's [44]. In a call-by-value setting, call unfolding is unsound (i.e., it does not preserve observational equivalence).

For example, consider the following Λ -term (under call-by-value), occurring in a static context where e_1 is dynamic.

$$\mathbf{pred}\ ((\lambda x.\ \ulcorner 43 \urcorner)\ @\ e_1)$$

Unfolding the inner call is unsound in general because e_1 may diverge. Yet a partial evaluator such as Lambda-Mix would unfold it [29, 30].

Through a systematic insertion of let-expressions, partial evaluators such as Similix or Schism ensure sound call unfolding [7, 9]. They would unfold the term above into the following term

$$\mathbf{pred} (\text{let } x = e_1 \text{ in } 43)$$

and would residualize the let expression to preserve the termination properties of the source term. Now a partial evaluator only needs to move the context $\mathbf{pred} [\cdot]$ inside the let expression to enable a static reduction that yields the expected answer. This may require a binding-time improvement, as investigated in Section 5. The \mapsto_x steps below represent how the above expression would be treated e.g., in Similix.

$$\begin{array}{lll} & \mathbf{pred} (\text{let } y = e_1 \text{ in } (\lambda x . 43) @ y) & \\ \mapsto_x & \mathbf{pred} (\text{let } y = e_1 \text{ in } 43) & \dots \text{unfold call} \\ \mapsto_x & \text{let } y = e_1 \text{ in } \mathbf{pred} 43 & \dots \text{binding-time improvement} \\ \mapsto_x & \text{let } y = e_1 \text{ in } 42 & \dots \text{static computation} \end{array}$$

In Λ_{ml} , evaluation steps only occur in **let** contexts and thus their result is named. Encodings of Λ into Λ_{ml} (such as those in Appendix A.2) insert **let** constructs at all evaluation contexts (rather than in an ad hoc manner as above). Because of the more natural placement of **let** expressions, the binding-time improvement step in the evaluation above is avoided — specialization steps are simply reductions in the Λ_{ml} calculus.

For example, the call-by-value encoding of Appendix A.2 yields the following term (after performing some initial *let.β* reductions).

$$\begin{aligned} e \equiv & \text{let } y_1 \Leftarrow \mathcal{E}_v \langle e_1 \rangle \\ & \text{in let } y_2 \Leftarrow (\lambda x . \text{unit } \ulcorner 43 \urcorner) @ y_1 \\ & \text{in pred } y_2 \end{aligned}$$

Assuming that e occurs in a static context and that $\mathcal{E}_v \langle e_1 \rangle$ is dynamic, binding-time analysis would associate e with the following annotated term (where w_1 is the annotation of $\mathcal{E}_v \langle e_1 \rangle$).

$$\begin{aligned} w \equiv & \underline{\text{let}} y_1 \Leftarrow w_1 \\ & \underline{\text{in}} \text{let } y_2 \Leftarrow (\lambda x . \text{unit } \ulcorner 43 \urcorner) @ y_1 \\ & \text{in pred } y_2 \end{aligned}$$

Specialization reduces the inner **let** expression and yield the expected answer.

$$\begin{aligned}
& w \\
\longrightarrow_s & \text{let } y_1 \Leftarrow \mathcal{E}_v\langle e \rangle \text{ in } \text{let } y_2 \Leftarrow \text{unit } \ulcorner 43 \urcorner \text{ in } \text{pred } y_2 \\
\longrightarrow_s & \text{let } y_1 \Leftarrow \mathcal{E}_v\langle e \rangle \text{ in } \text{pred } 43 \\
\longrightarrow_s & \text{let } y_1 \Leftarrow \mathcal{E}_v\langle e \rangle \text{ in } \text{unit } 42
\end{aligned}$$

Finally, another advantage of phrasing partial evaluation in terms of Λ_{ml} is that concepts of proper **let**-insertion and binding-time improvements (as discussed in the following section) can be presented independently of evaluation order. For example, our formalization of these concepts still remains valid if one considers e.g., a mixed evaluation strategy [16].

5 Control-based binding-time improvements

The partial evaluator defined in Section 3 is sensitive to the *structure* of source programs. Consider the two following Λ_{ml} -convertible expressions.

$$\begin{aligned}
e_1 & \equiv \text{let } v_1 \Leftarrow (\text{let } v_2 \Leftarrow x_2 \text{ in } \text{succ } x_1) \text{ in } e && \dots \text{where } v_2 \notin FV(e) \\
e_2 & \equiv \text{let } v_2 \Leftarrow x_2 \text{ in } \text{let } v_1 \Leftarrow \text{succ } x_1 \text{ in } e
\end{aligned}$$

Assume $\Gamma, x_1 : \mathbf{nat}, x_2 : \widetilde{\mathbf{nat}} \vdash_{ml} e_1, e_2 : \tilde{\tau}$, and let $\Gamma_{bt} = \Gamma'_{bt}, x_1 : \mathbf{nat}[s], x_2 : \widetilde{\mathbf{nat}}[d]$ be binding-time assumptions for e_1 and e_2 (i.e., x_1 is a static variable, x_2 is a dynamic variable).

Given the binding-time rules of Figures 6 and 7, e_1 must map to the following expression w_1 for some w (i.e., $\Gamma \vdash_{bt} e_1 : \tilde{\tau}[w_1 : d]$ is derivable).

$$w_1 \equiv \text{let } v_1 \Leftarrow (\text{let } v_2 \Leftarrow x_2 \text{ in } \text{succ } \text{lift } x_1) \text{ in } w$$

On the other hand, e_2 may map to the following expression w_2 for some w' (i.e., $\Gamma \vdash_{bt} e_2 : \tilde{\tau}[w_2 : d]$ is derivable).

$$w_2 \equiv \text{let } v_2 \Leftarrow x_2 \text{ in } \text{let } v_1 \Leftarrow \text{succ } x_1 \text{ in } w'$$

In w_1 , the **let** associated with v_1 is forced to be residual. In w_2 , it can be eliminable — which may enable further static computations, particularly within a loop.

This example captures in a nutshell the need for *binding-time improvements*: meaning-preserving transformations over source programs that enable more source expressions to be classified as static (and thus to be computed away statically, which yields more efficient specialized programs). In this example, e_1 can be “binding-time improved” by transforming it to e_2 using the *let.assoc* reduction of Λ_{ml} calculus (see Figure 3).

5.1 Flattening before binding-time analysis

The example above illustrates that the *let.assoc* reduction gives a useful binding-time improvement. This (along with the fact that *let.assoc* is confluent and strongly normalizing [34]) suggests a general strategy for improving binding times: before analyzing the binding times of a program, map this program to its *let.assoc* normal form.

The following example, however, illustrates that this strategy does not discover all binding-time improvements associated with the *let.assoc* reduction. *During specialization*, unfolding a call may expose flattening opportunities that are not apparent in the source program.

For example, assume that the following variant of e_1 above is in *let.assoc* normal form.

$$e_3 \equiv \text{let } v_1 \Leftarrow ((\lambda z. \text{let } v_2 \Leftarrow x_2 \text{ in succ } x_1) @ \ulcorner 0 \urcorner) \text{ in } e \\ \dots \text{where } v_2 \notin FV(e)$$

Using the same binding-time assumptions Γ_{bt} as above, e_3 maps to w_3 below for some w'' (i.e., the *let* associated with v_2 is forced to be residual).

$$w_3 \equiv \underline{\text{let}} v_1 \Leftarrow ((\lambda z. \underline{\text{let}} v_2 \Leftarrow x_2 \text{ in succ lift } x_1) @ \ulcorner 0 \urcorner) \underline{\text{in}} w''$$

Specialization (after supply static data $\ulcorner 42 \urcorner$ for x_1 includes the following step.

$$\begin{aligned} & \underline{\text{let}} v_1 \Leftarrow ((\lambda z. \underline{\text{let}} v_2 \Leftarrow x_2 \text{ in succ lift } \ulcorner 42 \urcorner) @ \ulcorner 0 \urcorner) \underline{\text{in}} w'' \\ \mapsto_s & \underline{\text{let}} v_1 \Leftarrow (\underline{\text{let}} v_2 \Leftarrow x_2 \text{ in succ lift } \ulcorner 42 \urcorner) \underline{\text{in}} w'' \end{aligned}$$

Note that this last term has the same problematic form as w_1 above.

5.2 Flattening during specialization

The solution is to incorporate the flattening rule *let.assoc* of Figure 3 as a specialization step. One might be tempted to simply redefine specialization so that the specialization of w_3 above proceeds as follows (i.e., a flattening step is taken whenever possible).

$$\begin{aligned} & \underline{\text{let}} v_1 \Leftarrow ((\lambda z. \underline{\text{let}} v_2 \Leftarrow x_2 \text{ in succ lift } \ulcorner 42 \urcorner) @ \ulcorner 0 \urcorner) \underline{\text{in}} w'' \\ \mapsto_s & \underline{\text{let}} v_1 \Leftarrow (\underline{\text{let}} v_2 \Leftarrow x_2 \text{ in succ lift } \ulcorner 42 \urcorner) \underline{\text{in}} w'' \\ \mapsto_f & \underline{\text{let}} v_2 \Leftarrow x_2 \text{ in } \underline{\text{let}} v_1 \Leftarrow \text{succ lift } \ulcorner 42 \urcorner \underline{\text{in}} w'' \end{aligned}$$

At this point, however, we are no better off than we were before since the specialization directives indicate that the *let* associated with v_1 must be residualized.

$$\frac{\Gamma \vdash_{bt}^+ e_1 : \tilde{\tau}_1[w_1 : \mathbf{d}] \quad \Gamma, x : \tau_1[\mathbf{d}] \vdash_{bt}^+ e_2 : \tilde{\tau}_2[w_2 : \varphi]}{\Gamma \vdash_{bt}^+ \mathbf{let} x \Leftarrow e_1 \mathbf{in} e_2 : \tilde{\tau}_2[\mathbf{let} x \Leftarrow w_1 \mathbf{in} w_2 : \varphi]} (**)$$

Figure 9: Modified binding-time constraints for Λ_{ml}^{bt}

Since it is the task of binding-time analysis to direct the specializer, the binding-time analysis also must reflect the possible use of the *let.assoc* rule. In the (*) rule of Figure 7, the body of the **let** is forced to have a specialization type (i.e., binding-time) of **d** since it cannot be consumed during specialization. However, after adding the flatten rule as a specialization step, a static expression in the body of a **let** can be consumed. Allowing the body of the **let** to have a static binding-time leads to the following annotation of e_3 above.

$$w_4 \equiv \mathbf{let} v_1 \Leftarrow ((\lambda z . \mathbf{let} v_2 \Leftarrow x_2 \mathbf{in} \mathbf{succ} x_1) @ \ulcorner 0 \urcorner) \mathbf{in} w'''$$

A specializer which incorporates the flattening rule now gives the desired behaviour.

$$\begin{aligned} & \mathbf{let} v_1 \Leftarrow ((\lambda z . \mathbf{let} v_2 \Leftarrow x_2 \mathbf{in} \mathbf{succ} x_1) @ \ulcorner 0 \urcorner) \mathbf{in} w''' \\ \longrightarrow_s & \mathbf{let} v_1 \Leftarrow (\mathbf{let} v_2 \Leftarrow x_2 \mathbf{in} \mathbf{succ} \ulcorner 42 \urcorner) \mathbf{in} w''' \\ \longrightarrow_f & \mathbf{let} v_2 \Leftarrow x_2 \mathbf{in} \mathbf{let} v_1 \Leftarrow \mathbf{succ} \ulcorner 42 \urcorner \mathbf{in} w''' \\ \longrightarrow_s & \mathbf{let} v_2 \Leftarrow x_2 \mathbf{in} \mathbf{let} v_1 \Leftarrow \mathbf{unit} \ulcorner 43 \urcorner \mathbf{in} w''' \\ \longrightarrow_s & \mathbf{let} v_2 \Leftarrow x_2 \mathbf{in} w'''[v_1 := \ulcorner 43 \urcorner] \end{aligned}$$

We formalize the modified binding-time analysis and specializer in the following sections.

5.2.1 Binding-time analysis

Replacing the (*) rule of Figure 7 with the (**) rule in Figure 9 gives improved constraints $\Gamma \vdash_{bt}^+ e : \tau[w : \varphi]$. The following property verifies that the improved constraints are at least as “good” as the previous ones. The example above shows that in many cases the improved constraints are better, i.e., they give more opportunities for static computation.

Property 1 $\Gamma \vdash_{bt} e : \tau[w : \varphi]$ implies $\Gamma \vdash_{bt}^+ e : \tau[w : \varphi]$

Proof: by induction over the derivation of $\Gamma \vdash_{bt} e : \tau[w : \varphi]$. ■

Binding-time analysis is defined as before (Definition 4). bta^+ denotes the analysis functions based on the improved constraints.

5.2.2 Specialization

Figure 10 gives specialization steps which incorporate the flattening rules. The first flattening rule gives the previously discussed binding-time improvements. The second rule is added for stylistic reasons. It simplifies the presentation so that flattening does not occur in an interpreter context E_i but only in specialization contexts E_s . Essentially, the specializer will attempt to apply flattening rules first. Otherwise, specialization proceeds as before. Note that in applying the *let.assoc* reduction, a renaming may be necessary to avoid variable capture (in Figure 3, the condition is that $x_1 \notin FV(e_3)$). We assume that necessary renamings take place when using the flattening rules of Figure 10 and we omit the corresponding formalization.

The following lemma gives specialization properties for the Λ_{ml}^{bt} terms satisfying judgements of the form $\Gamma_d \vdash_{bt}^+ e : \tau[w : \varphi_\tau]$. It only differs from the previous specialization properties (Lemma 2) in that the third component (dealing with judgements of the form $\Gamma_d \vdash_{bt} e : \tilde{\tau}[w : \tilde{\varphi}]$) now reflects the improved treatment of **let** constructs.¹⁰

Lemma 3 (specialization properties)

1. If $\Gamma_d \vdash_{bt}^+ e : \mathbf{nat}[w : s]$ then $w \equiv \ulcorner n \urcorner \equiv e$ for some number n .
2. If $\Gamma_d \vdash_{bt}^+ e : \tau_1 \rightarrow \tau_2[w : \varphi_1 \rightarrow \varphi_2]$ then $w \equiv \lambda x . w'$, $e \equiv \lambda x . [w']$ and $\Gamma_d, x : \tau_1[\varphi_1] \vdash_{bt}^+ [w'] : \tau_2[w' : \varphi_2]$ for some $w' \in \mathit{Terms}[\Lambda_{ml}^{bt}]$.
3. If $\Gamma_d \vdash_{bt}^+ e : \tilde{\tau}[w : \tilde{\varphi}]$ then exactly one of the following statements holds:
 - (a) $w \equiv \mathbf{unit} w'$, $e \equiv \mathbf{unit} [w']$ and $\Gamma \vdash_{bt}^+ [w'] : \tau[w' : \varphi]$ for some $w' \in \mathit{Terms}[\Lambda_{ml}^{bt}]$, or
 - (b) $w \mapsto_i w'$ and $w \not\mapsto_f w''$ and $\Gamma_d \vdash_{bt}^+ [w'] : \tilde{\tau}[w' : \tilde{\varphi}]$ and $e \equiv [w] \twoheadrightarrow_{ml} [w']$, or

¹⁰The current statement of Lemma 3 makes a few format changes and one minor technical correction to the Lemma in the shorter MSCS version of this paper: subcases 3(b) and 3(c) have been switched, and the condition $w \not\mapsto_f w''$ has been added to subcase 3(b) to ensure that only one of the subcases of component 3 hold.

Flattening rules (where $x_1 \notin FV(w_3)$):

$\text{let } x_2 \Leftarrow (\text{let } x_1 \Leftarrow w_1 \text{ in } w_2) \text{ in } w_3 \mapsto_f \text{let } x_1 \Leftarrow w_1 \text{ in let } x_2 \Leftarrow w_2 \text{ in } w_3$

$\text{let } x_2 \Leftarrow (\text{let } x_1 \Leftarrow w_1 \text{ in } w_2) \text{ in } w_3 \mapsto_f \text{let } x_1 \Leftarrow w_1 \text{ in let } x_2 \Leftarrow w_2 \text{ in } w_3$

$$\frac{w \mapsto_f w'}{w \mapsto_{s+} w'}$$

Interpretation rules:

$\text{succ } \ulcorner n \urcorner \mapsto_i \text{unit } \ulcorner n + 1 \urcorner$ $\text{pred } \ulcorner n + 1 \urcorner \mapsto_i \text{unit } \ulcorner n \urcorner$

$(\lambda x . w_0) @ w_1 \mapsto_i w_0[x := w_1]$ $\text{pred } \ulcorner 0 \urcorner \mapsto_i \text{unit } \ulcorner 0 \urcorner$

$\text{fix } x . w \mapsto_i w[x := \text{fix } x . w]$ $\text{if0 } \ulcorner 0 \urcorner w_2 w_3 \mapsto_i w_2$

$\text{let } x \Leftarrow \text{unit } w_1 \text{ in } w_2 \mapsto_i w_2[x := w_1]$ $\text{if0 } \ulcorner n + 1 \urcorner w_2 w_3 \mapsto_i w_3$

$$\frac{w \mapsto_i w'}{E_i[w] \mapsto_i E_i[w']} \quad E_i ::= \text{let } x \Leftarrow [\cdot] \text{ in } w_2$$

$$\frac{w \mapsto_i w'}{w \mapsto_{s+} w'} \quad \text{where } w \not\mapsto_f w''$$

Compilation rules:

$$\text{lift } \ulcorner n \urcorner \mapsto_{s+} \ulcorner n \urcorner \quad \frac{w \mapsto_{s+} w'}{E_s[w] \mapsto_{s+} E_s[w']}$$

$E_s ::= \underline{\text{pred}} [\cdot] \mid \underline{\text{succ}} [\cdot] \mid \underline{\text{if0}} [\cdot] w_2 w_3 \mid \underline{\text{if0}} r_1 [\cdot] w_3 \mid \underline{\text{if0}} r_1 r_2 [\cdot] \mid$
 $\underline{\lambda} x . [\cdot] \mid [\cdot] @ w_1 \mid r_0 @ [\cdot] \mid \underline{\text{fix}} x . [\cdot] \mid$
 $\underline{\text{unit}} [\cdot] \mid \underline{\text{let}} x \Leftarrow [\cdot] \text{ in } w_2 \mid$
 $\underline{\text{let}} x \Leftarrow r_1 \text{ in } [\cdot] \quad \dots \text{where } r_i \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$

Figure 10: Modified specialization rules for Λ_{ml}^{bt}

- (c) $w \equiv \text{let } x \leftarrow w_1 \text{ in } w_2$, $e \equiv \text{let } x \leftarrow [w_1] \text{ in } [w_2]$, and
 $\Gamma_d \vdash_{bt}^+ [w_1] : \tilde{\tau}_1[w_1 : \mathbf{d}]$, $\Gamma_d, x : \tau_1[\mathbf{d}] \vdash_{bt}^+ [w_2] : \tilde{\tau}_2[w_2 : \tilde{\varphi}]$,
for some $w_1, w_2 \in \text{Terms}[\Lambda_{ml}^{bt}]$, or
- (d) $w \mapsto_f w'$ and $\Gamma_d \vdash_{bt}^+ [w'] : \tilde{\tau}[w' : \tilde{\varphi}]$ and $e \equiv [w] \longrightarrow_{ml} [w']$.
4. If $\Gamma_d \vdash_{bt}^+ e : \text{nat}[w : \mathbf{d}]$ then exactly one of the following statements holds:
- (a) $w \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$, or
- (b) $w \mapsto_{s+} w'$, $\Gamma_d \vdash_{bt}^+ [w'] : \tau[w' : \mathbf{d}]$, and $e \equiv [w] \equiv [w'] \equiv \ulcorner n \urcorner$ for some number n .
5. If $\Gamma_d \vdash_{bt}^+ e : \tau_1 \rightarrow \tau_2[w : \mathbf{d}]$ then exactly one of the following statements holds:
- (a) $w \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$, or
- (b) $w \mapsto_{s+} w'$ and $\Gamma_d \vdash_{bt}^+ [w'] : \tau[w' : \mathbf{d}]$ and $e \equiv [w] \longrightarrow_{ml} [w']$.
6. If $\Gamma_d \vdash_{bt}^+ e : \tilde{\tau}[w : \mathbf{d}]$ then exactly one of the following statements holds:
- (a) $w \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$, or
- (b) $w \mapsto_{s+} w'$ and $\Gamma_d \vdash_{bt}^+ [w'] : \tilde{\tau}[w' : \mathbf{d}]$ and $e \equiv [w] \longrightarrow_{ml} [w']$.

Proof: by induction over the height of the derivation of $\Gamma_d \vdash_{bt}^+ e : \tau[w : \varphi]$ relying on the property that if both $\Gamma, x : \tau_1[\varphi_1] \vdash_{bt}^+ e_0 : \tau_0[w_0 : \varphi_0]$ and $\Gamma \vdash_{bt}^+ e_1 : \tau_1[w_1 : \varphi_1]$ then $\Gamma \vdash_{bt}^+ e_0[x := e_1] : \tau_0[w_0[x := w_1] : \varphi_0]$. See Appendix B for details. \blacksquare

As before, if $w \mapsto_{s+} w'$ and $w \mapsto_{s+} w''$ then $w' \equiv w''$. The definitions of specialization and partial evaluation as well as the proofs of correctness of binding-time analysis, specialization, and partial evaluation proceed exactly as before.

6 Equivalence with continuation-based specialization

Not only does the incorporation of the *let.assoc* lead to improved binding times, but the resulting improved specializer captures the essence of control-based binding-time improvements as found in, e.g., the partial evaluator Similix [7]. To substantiate this claim, we recast Bondorf's continuation-based specializer [6] in terms of Λ_{ml} , and prove that this specializer is equivalent to the *spec*⁺ of Section 5.2.2.

Interpretation rules:

$$\begin{array}{c}
\langle \mathbf{unit} \ w, \ nil \rangle \Downarrow \mathbf{unit} \ w \qquad \frac{\langle \mathbf{unit} \ \ulcorner n + 1 \urcorner, \vec{E}_i \rangle \Downarrow a}{\langle \mathbf{succ} \ \ulcorner n \urcorner, \vec{E}_i \rangle \Downarrow a} \\
\\
\frac{\langle \mathbf{unit} \ \ulcorner n \urcorner, \vec{E}_i \rangle \Downarrow a}{\langle \mathbf{pred} \ \ulcorner n + 1 \urcorner, \vec{E}_i \rangle \Downarrow a} \qquad \frac{\langle \mathbf{unit} \ \ulcorner 0 \urcorner, \vec{E}_i \rangle \Downarrow a}{\langle \mathbf{pred} \ \ulcorner 0 \urcorner, \vec{E}_i \rangle \Downarrow a} \\
\\
\frac{\langle w_2, \vec{E}_i \rangle \Downarrow a}{\langle \mathbf{if0} \ \ulcorner 0 \urcorner \ w_2 \ w_3, \vec{E}_i \rangle \Downarrow a} \qquad \frac{\langle w_3, \vec{E}_i \rangle \Downarrow a}{\langle \mathbf{if0} \ \ulcorner n + 1 \urcorner \ w_2 \ w_3, \vec{E}_i \rangle \Downarrow a} \\
\\
\frac{\langle w_0[x := w_1], \vec{E}_i \rangle \Downarrow a}{\langle (\lambda x . w_0) @ w_1, \vec{E}_i \rangle \Downarrow a} \qquad \frac{\langle w[x := \mathbf{fix} \ x . w], \vec{E}_i \rangle \Downarrow a}{\langle \mathbf{fix} \ x . w, \vec{E}_i \rangle \Downarrow a} \\
\\
\frac{\langle w_1, \mathbf{let} \ x \Leftarrow \cdot \ :: \vec{E}_i \rangle \Downarrow a}{\mathbf{in} \ w_2} \qquad \frac{\langle \mathbf{let} \ x \Leftarrow w_1, \vec{E}_i \rangle \Downarrow a}{\mathbf{in} \ w_2} \quad (\text{Push}) \qquad x \notin FV(\vec{E}_i) \\
\\
\frac{\langle w_2[x := w_1], \vec{E}_i \rangle \Downarrow a}{\langle \mathbf{unit} \ w_1, \mathbf{let} \ x \Leftarrow \cdot \ :: \vec{E}_i \rangle \Downarrow a} \quad (\text{Pop}) \\
\mathbf{in} \ w_2
\end{array}$$

Figure 11: Continuation-based specializer for Λ_{ml}^{bt} (part 1)

Compilation rules:

$$\langle \ulcorner n \urcorner, nil \rangle \Downarrow \ulcorner n \urcorner \quad \langle x, nil \rangle \Downarrow x \quad \langle \text{lift } \ulcorner n \urcorner, nil \rangle \Downarrow \ulcorner n \urcorner$$

$$\frac{\langle w, nil \rangle \Downarrow r}{\langle \text{unit } w, nil \rangle \Downarrow \text{unit } r}$$

$$\frac{\langle w, nil \rangle \Downarrow r}{\langle \text{succ } w, nil \rangle \Downarrow \text{succ } r}$$

$$\frac{\langle w, nil \rangle \Downarrow r}{\langle \text{pred } w, nil \rangle \Downarrow \text{pred } r}$$

$$\frac{\langle w, nil \rangle \Downarrow r}{\langle \lambda x . w, nil \rangle \Downarrow \lambda x . r}$$

$$\frac{\langle w, nil \rangle \Downarrow r}{\langle \text{fix } x . w, nil \rangle \Downarrow \text{fix } x . r}$$

$$\frac{\langle w_1, nil \rangle \Downarrow r_1 \quad \langle w_2, nil \rangle \Downarrow r_2}{\langle w_1 @ w_2, nil \rangle \Downarrow r_1 @ r_2}$$

$$\frac{\langle w_1, nil \rangle \Downarrow r_1 \quad \langle w_2, nil \rangle \Downarrow r_2 \quad \langle w_3, nil \rangle \Downarrow r_3}{\langle \text{if0 } w_1 w_2 w_3, nil \rangle \Downarrow \text{if0 } r_1 r_2 r_3}$$

$$\frac{\langle w_1, nil \rangle \Downarrow r_1 \quad \langle w_2, \vec{E}_i \rangle \Downarrow r_2}{\langle \text{let } x \leftarrow w_1 \text{ in } w_2, \vec{E}_i \rangle \Downarrow \text{let } x \leftarrow r_1 \text{ in } r_2} \text{ (Assoc)} \quad x \notin FV(\vec{E}_i)$$

Figure 12: Continuation-based specializer for Λ_{ml}^{bt} (part 2)

Figures 11 and 12 present the continuation-based specializer as a big-step operational semantics. Bondorf’s specializer is expressed denotationally, but there are standard techniques for going from a denotational specification to an operational one [53].¹¹ The specializer manipulates configurations of the form $\langle w, \vec{E}_i \rangle$ where $w \in \text{Terms}[\Lambda_{nil}^{bt}]$ and \vec{E}_i is a stack of interpretation contexts (i.e., \vec{E}_i is a continuation). The intuition is that \vec{E}_i accumulates static contexts (via the *(Push)* rule) and if possible moves them across dynamic **let** expressions to consume static values in the **let** body (via the *(Assoc)* rule). The *(Pop)* rule consumes a static value by reducing the top-most **let** in \vec{E}_i . In dynamic contexts (e.g., in the compilation rules), there is no need to accumulate contexts (i.e., the stack is delimited to be *nil*— which corresponds to the identity continuation).

The free variable restrictions in the *(Push)* and *(Assoc)* rules are analogous to the free variable restrictions on the flattening rules of Figure 10. The following is the fairly obvious extension of the notion of free variables to stacks of interpretation contexts.¹²

$$\begin{aligned} FV(\text{nil}) &= \emptyset \\ FV(\text{let } x \leftarrow \cdot \text{ :: } \vec{E}_i) &= (FV(w_2) - \{x\}) \cup FV(\vec{E}_i) \\ &\quad \text{in } w_2 \end{aligned}$$

Derivations in the continuation-based specializer give rise to context stacks that are *well-formed* in the following sense.

Definition 7 (well-formed context stack)

- *nil* is well-formed; and
- $\text{let } x \leftarrow \cdot \text{ :: } \vec{E}_i$ is well-formed if \vec{E}_i is well-formed and $x \notin FV(\vec{E}_i)$.
in w_2

¹¹For stylistic reasons we have deviated slightly from Bondorf’s presentation, where all the rules for dynamic constructs are expressed in continuation-passing style, and dynamic contexts are thus accumulated as well. This accumulation, however, is not necessary since continuations are only used to move static contexts across dynamic **let** expressions. Based on this observation (made during our initial work, in the Fall of 1994), we have written our rules so that components of dynamic constructs (e.g., **if0**) are specialized with respect to the identity continuation (i.e., *nil*).

¹²The material from the current paragraph down to the statement of Theorem 5 were omitted from the shorter version of this paper due to space constraints. Also, the free variable constraints on the *(Push)* (Figure 11) and *(Assoc)* (Figure 12) rules were not made explicit in the shorter version.

Intuitively, context stacks

$$\vec{E}_i \equiv \text{let } x_1 \Leftarrow \cdot \ :: \ \dots \ :: \ \text{let } x_n \Leftarrow \cdot \ :: \ \text{nil} \\ \text{in } w_1 \qquad \qquad \qquad \text{in } w_n$$

in the continuation-based specializer arise from contexts of the form

$$\text{let } x_n \Leftarrow (\dots \text{let } x_1 \Leftarrow \cdot \text{ in } w_1 \dots) \text{ in } w_n.$$

in the term being specialized. As discussed in Section 5, for all i and for all $j > i$, α -conversion can give $x_i \notin FV(w_j)$.

The following property states that continuation-based specialization does indeed give rise to well-formed contexts. We write $\mathcal{D} \triangleright \langle w, \vec{E}_i \rangle \Downarrow a$ when \mathcal{D} is a derivation of $\langle w, \vec{E}_i \rangle \Downarrow a$.

Property 2 (well-formed context stacks are preserved)

If $\mathcal{D} \triangleright \langle w, \vec{E}_i \rangle \Downarrow a$ and \vec{E}_i is a well-formed context stack, then for all subderivations $\mathcal{D}' \triangleright \langle w', \vec{E}'_i \rangle \Downarrow a'$ of \mathcal{D} , \vec{E}'_i is a well-formed context stack.

Proof: by complete induction over the height of \mathcal{D} . The only interesting case is the (*Push*) rule, and the side condition ensures that the property holds for the immediate subderivation. ■

The following theorem establishes the correspondence between the continuation-based specializer and $spec^+$ of Section 5.2.2.

Theorem 5 *For $\Gamma_d \vdash_{bt}^+ e : \widetilde{\text{nat}}[w : d]$,*

$$\langle w, \text{nil} \rangle \Downarrow r \text{ iff } w \mapsto_{s^+}^* r$$

Proof: The proof is straightforward and follows the standard pattern for relating a big-step and small-step operational semantics [32, p. 111]. However, one must strengthen the statement to prove so as to handle arbitrary stacks of interpretation contexts in the inductive hypothesis. The full proof is in Appendix B.3. ■

Bondorf's presentation [6] did not include a specification of binding-time analysis. However, our connection of continuation-based specialization and $spec^+$ shows that the constraints given by $\Gamma \vdash_{bt}^+ e : \tau[w : \varphi]$ judgements are the proper ones (appealing to the correctness of binding-time analysis for $spec^+$).

7 Related work

7.1 Computational meta-languages

Moggi introduced his computational metalanguage as a convenient format to specify denotational semantics modularly [48]. This metalanguage has been used for a variety of purposes including treating computational effects in functional languages [24, 31, 47, 63, 66], staging denotational-semantics definitions [12] and compiler translations [3, 36, 60], and explaining relationships between various constructive logics [4]. In an earlier work [34, 36], we have used this metalanguage to formalize the structure of continuation-passing styles.

Inspired by the present work, Nielsen has recently developed an evaluation-order independent presentation of partial evaluation and deforestation using Moggi’s metalanguage [49]. The distinction between values and computations given by the metalanguage type system (upon which we have relied heavily) lies at the foundation of his work as well. This continues a trend of unifying various program-specialization techniques [62]. Clearly Λ_{ml} stands as a promising testbed for this unification.

As noted above, the metalanguage and monads are commonly used to structure functional programs with computational effects. Both Nielsen [49] and Dussart and Thiemann [22] use monads to structure I/O and state in the partial evaluator itself.

Particular aspects of Moggi’s metalanguage highlighted in the present work also appear in other metalanguages. For example, Talcott uses let-expressions to specify computational steps [64]. Sabry and Felleisen demonstrate how a let-based intermediate language (called *A-normal forms*) can give benefits similar to CPS in dataflow analysis [58, 59].

Moggi also gave a *computational λ -calculus* λ_c . λ_c is essentially a call-by-value λ -calculus extended with equations that capture program equivalences holding under arbitrary monadic effects. It is straightforward to adapt our evaluation-order independent presentation based on Λ_{ml} to a call-by-value presentation based on λ_c . This alternative presentation expresses control-based binding-time improvements via the monadic laws as formulated in λ_c . The reassociation of **let** expressions is again a prominent feature. In fact, the λ_c presentation has strong connections with the work of Flanagan et al. [25] on the essence of compiling with continuations.

7.2 Partial evaluation

7.2.1 Styles of specification

The earliest work on partial evaluation as a two-phase process specified binding-time analysis as an abstract interpretation [43]. Since the work of Jones and Gomard [29, 30], binding-time analysis is more often specified using type systems (we have followed this approach in the present work). Palsberg [54] and Wand [68] further clarify the rôle of such specifications in their work on the correctness of binding-time analysis.

Specializers have mostly been specified as symbolic interpreters in functional style [42]. We note, however, a recent trend (including the present work) to use operational semantics for specifying specializers [1, 18, 33, 61, 62]. In fact, the first author has shown that by emphasizing the logical character of type-based and operational semantics specifications, the correctness of a partial evaluator can be mechanically verified [35].

Davies and Pfenning have developed a language for expressing staged computation based on the intuitionistic modal logic S4 [20, 19]. The type system of this language is strikingly similar to that of Moggi’s computational metalanguage. A modality analogous to the \sim construction of Moggi is used to type *code* objects (in our terminology, objects whose specialization type is *dynamic*). Their language also includes a *let* construct which can act as an “eval function”. However, even though types and terms of both languages are quite similar, the similarity is superficial in the context of our application of the metalanguage. In our setting, staging information is not represented using the modality (as with Davies and Pfenning). Instead it is represented using *specialization types* which are external to the types of the meta-language itself. In our case, the modality is used to distinguish *values* from *computations* — a distinction not captured in Davies and Pfenning’s work. Benton, Bierman, and de Paiva [4] flesh out connections between Moggi’s computational metalanguage and the modal logic S4 in a more general context.

7.2.2 Call unfolding and let insertion

Most specializers ensuring sound call-unfolding under call-by-value adopt the technique of *let* insertion, as discussed in Section 4. One may, however, also enforce soundness by simply not unfolding calls where the argument expression is dynamic, at the price of reducing specialization. This must be expressed in the binding-time analysis by forbidding binding-times such

as $\mathbf{d} \rightarrow \mathbf{s}$ (i.e., a static function mapping a dynamic argument into a static result). This restriction in fact also occurs in Nielson and Nielson’s two-level λ -calculus [52]: the co-domain of a static function should be at least as dynamic as the domain of this function.

In contrast, consider a partial evaluator (e.g., Similix) that (1) ensures sound call-unfolding by **let** insertion, and (2) performs binding-time improvements by relocating static evaluation contexts inside dynamic **let** expressions. This partial evaluator does not constrain the domain and the codomain of static (call-by-value) functions. For example, in the term

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

where d denotes a dynamic integer, the λ -abstraction is classified as a static function mapping a dynamic integer into a static integer (i.e., its binding time is $\mathbf{d} \rightarrow \mathbf{s}$). As a corollary, the addition is classified as static. The residual program reads

$$\text{let } x \Leftarrow d \text{ in } 3.$$

In both techniques above, the possible binding times (i.e., specialization types) are tied to the strategy used to enforce sound call unfolding. A pleasant feature of phrasing partial evaluation in terms of Λ_{ml} is that the characterization of sound call unfolding and possible binding times are orthogonal. The distinction between values and computations in the type system means that static functions with specialization types such as $\mathbf{d} \rightarrow \mathbf{s}$ can *always* be dealt with in a sound manner. It is the encoding into Λ_{ml} where one adopts a technique for a particular evaluation order.

7.2.3 Binding-time improvements

Consel and Danvy observed that a source transformation into continuation-passing style prevented a class of loss of static information across procedures, and they provided a syntactic characterization of this class [10]. Holst and Gomard observed that part of the same effect (intra-procedural and insensitive to call unfolding) could be obtained by “flattening” each source procedure [39]. Bondorf, Danvy, and Lawall observed that a further part of the same effect (accounting for call unfolding, but not crossing specialization points) could be obtained by specific control operations in the specializer itself [6, 45], rather than by CPS-transforming the program before partial evaluation. Our framework achieves this effect by extending the operational semantics of specialization with the *let.assoc* rule. This development

matches contemporary work on determining the effect of the CPS transformation on flow analysis [51, 59]: enriching a direct-style calculus can yield analyses with an added precision that matches the extra precision obtained by the relocation of contexts performed by the CPS transformation.

Section 6 shows how the theory of monads captures the essence of control-based binding-time improvements. This has practical benefits as well — it allows one to avoid using a functional representation of continuations in the specializer (one need only use the *let.assoc* rule). The disadvantage of representing continuations as functions shows up in a self-applicable partial evaluator, as self-application generates programs with many higher-order functions. Such programs are more difficult to reason about, e.g., when searching for binding-time improvements. This difficulty motivated Lawall and Danvy to stick to direct style [45], and Glück and Jørgensen to devise a *multi-level cogen* [27]. Thiemann has united both lines of work [65]. In the present case, and since continuation-passing style can be obtained from monadic style by simply selecting (a term representation of) the continuation monad, our method provides sound guidelines for treating continuations in an offline partial evaluator.

We also expect our technique to be particularly useful in online self-applicable program specializers (e.g., supercompilers). In fact, preliminary work by the first author and Glück on online self-application uses a language called *Sgraph* [28] where **let**-bindings (in the style of Λ_{ml}) are a central feature [37].

Let us briefly attempt a taxonomy of continuation-based partial evaluators. The earliest one is reported in the literature by Bondorf [6]. Its goal is precisely to relocate static contexts across dynamic **let** expressions, as in the CPS transformation. The particular brand of continuation-passing style used for this relocation is expressible in direct style, using control operators (again as in the CPS transformation). In fact, this direct-style specializer turns out to be more efficient in practice [45]. A parallel development is taking place in the “cogen” approach to partial evaluation [8, 46, 65]. All these continuation-based specializers are specified as symbolic interpreters in functional style. An increasing number of specializers, however, are specified operationally. We have presented here such a specializer, which is continuation-based and delimits control to static contexts. This novel feature can be observed to be rippling back into the world of functional specializers, both with explicit continuations and with implicit continuations and the associated control operators [21].

As for data-flow binding-time improvements, they arise from insufficient

binding-time coercions [17, 18]. In particular, the binding-time improvement arising from the presence of booleans and disjoint sums is known as “The Trick” [18, 42]. Because of the nature of disjoint sums, this coercion takes the form of a control-based binding-time improvement: The Trick amounts to duplicating static contexts across dynamic case expressions. Again, continuations can be used to move static contexts across dynamic conditional expressions, duplicating them in the conditional branches. This transformation, however, can also be naturally accomplished using the computational meta-language [36]. It gives rise to an analogue of Figure 9 for `if0`, and to a context duplication in Figure 12. It is interesting to note (this observation is due to Malmkjær) that if we consider a `let` expression as a “unary” case expression (i.e., a case expression with one conditional branch), then our `let`-rearranging rules coincide with the case-rearranging rules (also known as *commuting conversions*) that can be found both in partial evaluation [18], program extraction [55], and natural-deduction proof theory [26, 57].

7.2.4 Evaluation-order independence

We have formulated control-based binding-time improvements via monads using the computational meta-language because it allows an evaluation-order independent view of binding-time analysis and specialization. This appears particularly useful in settings where adopting mixed evaluation strategies (e.g., call-by-name and call-by-value) can be employed to enhance efficiency [16]. In addition, Nielsen and Sørensen have identified situations where transforming sections of call-by-value programs using the call-by-name CPS transformation can increase specialization [50]. The evaluation-order independent partial-evaluation strategy that we have given here seems well-suited for this endeavour since Λ_{ml} allows one to encode mixed evaluation strategies while remaining in direct style [34].

8 Conclusion

Partial evaluation offers a practical way of staging the execution of programs in order to adapt them to the context of their execution. We have identified Moggi’s computational metalanguage as a useful intermediate language for formalizing it. As a first step, we have formalized binding-time analysis and program specialization. Then we have shown how the intermediate language captures the essence of control-based binding-time improvements. Other work gives evidence that the metalanguage can provide an avenue for

$$\begin{aligned}
e &\in \text{Terms}[\Lambda] \\
e &::= x \mid f \mid \ulcorner n \urcorner \mid \text{succ } e \mid \text{pred } e \mid \\
&\quad \text{if0 } e_1 e_2 e_3 \mid \lambda x:\tau . e \mid e_0 @ e_1 \mid \text{fix } f:\tau . e \\
x &\in \text{Identifiers}[\Lambda] \\
f &\in \text{RecIdentifiers}[\Lambda] \\
\tau &\in \text{Types}[\Lambda] \\
\tau &::= \text{nat} \mid \tau_1 \rightarrow \tau_2 \\
\Gamma &\in \text{Assumptions}[\Lambda] \\
\Gamma &::= \cdot \mid \Gamma, x:\tau \mid \Gamma, f:\tau
\end{aligned}$$

Figure 13: The language Λ

(a) incorporating other computational effects into partial evaluation such as I/O and state, and (b) unifying various program-specialization techniques.

Acknowledgements

Thanks are due to Karoline Malmkjær for comments on an earlier version of this paper, and to the anonymous referees for their perceptive reviews.

Most of this work was carried out while the first author was visiting the University of Copenhagen (DIKU), and we thank the DIKU TOPPS group for their interaction, encouragement, and financial support, e.g., to attend LDPL'95 in Darmstadt, Germany.

A The Language Λ

A.1 Syntax

Figure 13 presents the syntax of a PCF-like language Λ . We omit the usual typing rules as well as the call-by-name and call-by-value operational semantics for Λ .

The syntax of Λ includes two categories of identifiers: *Identifiers* are used in λ -bindings; *RecIdentifiers* are used in **fix**-bindings. This distinction is necessary under call-by-value evaluation of Λ , where *Identifiers* will only bind to canonical terms (i.e., values) whereas *RecIdentifiers* may bind to non-canonical terms (i.e., computations). This is reflected in the formal definition of values below.

Definition 8 (Values)

$$\begin{array}{ll} v \in \text{Values}_n[\Lambda] & v \in \text{Values}_v[\Lambda] \\ v ::= \ulcorner n \urcorner \mid \lambda x . e & v ::= x \mid \ulcorner n \urcorner \mid \lambda x . e \end{array}$$

A.2 Encoding Λ evaluation in Λ_{ml}

Figures 14 and 15 give the call-by-name and call-by-value encodings of Λ in Λ_{ml} . In the call-by-name encoding, function arguments are passed unevaluated. This is reflected in the transformation on types, i.e., functions map computations to computations. In the call-by-value encoding, evaluation of function arguments is forced using the **let** construct. This is reflected in the transformation on types, i.e., functions map values to computations.¹³

B Proofs

B.1 Properties of partial evaluation

The section gives the proofs for various partial evaluation properties and theorems occurring in Sections 3 and 4. We begin by considering various properties needed for the proof of Lemma 2 of Section 3.

Property 3 (substitution and binding-time constraints)

If $\Gamma, x : \tau_1[\varphi_1] \vdash_{bt} e_0 : \tau_0[w_0 : \varphi_0]$ and $\Gamma \vdash_{bt} e_1 : \tau_1[w_1 : \varphi_1]$ then

$$\Gamma \vdash_{bt} e_0[x := e_1] : \tau_0[w_0[x := w_1] : \varphi_0].$$

Proof: The proof is a simple induction over the derivation of $\Gamma, x : \tau_1[\varphi_1] \vdash_{bt} e_0 : \tau_0[w_0 : \varphi_0]$ and is omitted. ■

¹³One can also give a call-by-value encoding where functions map computations to computations. Instead of forcing evaluation of function arguments in the application structure (as in \mathcal{E}_v), evaluation is forced immediately after an argument is received by a function [36].

$$\begin{aligned}
\mathcal{E}_n\langle\cdot\rangle &: \Lambda \rightarrow \Lambda_{ml} \\
\mathcal{E}_n\langle v \rangle &= \mathbf{unit} \ \mathcal{E}_n\langle v \rangle \\
\mathcal{E}_n\langle x \rangle &= x \\
\mathcal{E}_n\langle f \rangle &= f \\
\mathcal{E}_n\langle \mathbf{succ} \ e \rangle &= \mathbf{let} \ y \Leftarrow \mathcal{E}_n\langle e \rangle \ \mathbf{in} \ \mathbf{succ} \ y \\
\mathcal{E}_n\langle \mathbf{pred} \ e \rangle &= \mathbf{let} \ y \Leftarrow \mathcal{E}_n\langle e \rangle \ \mathbf{in} \ \mathbf{pred} \ y \\
\mathcal{E}_n\langle \mathbf{if0} \ e_0 \ e_1 \ e_2 \rangle &= \mathbf{let} \ y_0 \Leftarrow \mathcal{E}_n\langle e_0 \rangle \ \mathbf{in} \ \mathbf{if0} \ y_0 \ \mathcal{E}_n\langle e_1 \rangle \ \mathcal{E}_n\langle e_2 \rangle \\
\mathcal{E}_n\langle e_0 \ @ \ e_1 \rangle &= \mathbf{let} \ y_0 \Leftarrow \mathcal{E}_n\langle e_0 \rangle \ \mathbf{in} \ y_0 \ @ \ \mathcal{E}_n\langle e_1 \rangle \\
\mathcal{E}_n\langle \mathbf{fix} \ f . e \rangle &= \mathbf{fix} \ f . \mathcal{E}_n\langle e \rangle \\
\\
\mathcal{E}_n\langle\cdot\rangle &: \mathit{Values}_n[\Lambda] \rightarrow \Lambda_{ml} \\
\mathcal{E}_n\langle \ulcorner n \urcorner \rangle &= \ulcorner n \urcorner \\
\mathcal{E}_n\langle \lambda x . e \rangle &= \lambda x . \mathcal{E}_n\langle e \rangle \\
\\
\mathcal{E}_n\langle \mathbf{nat} \rangle &= \mathbf{nat} \\
\mathcal{E}_n\langle \tau_1 \rightarrow \tau_2 \rangle &= \mathcal{E}_n\langle \tau_1 \rangle \rightarrow \mathcal{E}_n\langle \tau_2 \rangle \\
\mathcal{E}_n\langle \tau \rangle &= \widetilde{\mathcal{E}_n\langle \tau \rangle} \\
\\
\mathcal{E}_n\langle \Gamma, x : \tau \rangle &= \mathcal{E}_n\langle \Gamma \rangle, x : \mathcal{E}_n\langle \tau \rangle \\
\mathcal{E}_n\langle \Gamma, f : \tau \rangle &= \mathcal{E}_n\langle \Gamma \rangle, f : \mathcal{E}_n\langle \tau \rangle
\end{aligned}$$

Figure 14: Call-by-name encoding \mathcal{E}_n into Λ_{ml}

$$\begin{aligned}
\mathcal{E}_v\langle\cdot\rangle & : \Lambda \rightarrow \Lambda_{ml} \\
\mathcal{E}_v\langle v \rangle & = \mathbf{unit} \ \mathcal{E}_v\langle v \rangle \\
\mathcal{E}_v\langle f \rangle & = f \\
\mathcal{E}_v\langle \mathbf{succ} \ e \rangle & = \mathbf{let} \ y \Leftarrow \mathcal{E}_v\langle e \rangle \ \mathbf{in} \ \mathbf{succ} \ y \\
\mathcal{E}_v\langle \mathbf{pred} \ e \rangle & = \mathbf{let} \ y \Leftarrow \mathcal{E}_v\langle e \rangle \ \mathbf{in} \ \mathbf{pred} \ y \\
\mathcal{E}_v\langle \mathbf{if0} \ e_0 \ e_1 \ e_2 \rangle & = \mathbf{let} \ y_0 \Leftarrow \mathcal{E}_v\langle e_0 \rangle \ \mathbf{in} \ \mathbf{if0} \ y_0 \ \mathcal{E}_v\langle e_1 \rangle \ \mathcal{E}_v\langle e_2 \rangle \\
\mathcal{E}_v\langle e_0 \ @ \ e_1 \rangle & = \mathbf{let} \ y_0 \Leftarrow \mathcal{E}_v\langle e_0 \rangle \ \mathbf{in} \ \mathbf{let} \ y_1 \Leftarrow \mathcal{E}_v\langle e_1 \rangle \ \mathbf{in} \ y_0 \ @ \ y_1 \\
\mathcal{E}_v\langle \mathbf{fix} \ f . e \rangle & = \mathbf{fix} \ f . \mathcal{E}_v\langle e \rangle \\
\\
\mathcal{E}_v\langle \cdot \rangle & : \mathit{Values}_v[\Lambda] \rightarrow \Lambda_{ml} \\
\mathcal{E}_v\langle \ulcorner n \urcorner \rangle & = \ulcorner n \urcorner \\
\mathcal{E}_v\langle x \rangle & = x \\
\mathcal{E}_v\langle \lambda x . e \rangle & = \lambda x . \mathcal{E}_v\langle e \rangle \\
\\
\mathcal{E}_v\langle \mathbf{nat} \rangle & = \mathbf{nat} \\
\mathcal{E}_v\langle \tau_1 \rightarrow \tau_2 \rangle & = \mathcal{E}_v\langle \tau_1 \rangle \rightarrow \mathcal{E}_v\langle \tau_2 \rangle \\
\mathcal{E}_v\langle \tau \rangle & = \widetilde{\mathcal{E}_v\langle \tau \rangle} \\
\\
\mathcal{E}_v\langle \Gamma, x : \tau \rangle & = \mathcal{E}_v\langle \Gamma \rangle, x : \mathcal{E}_v\langle \tau \rangle \\
\mathcal{E}_v\langle \Gamma, f : \tau \rangle & = \mathcal{E}_v\langle \Gamma \rangle, f : \mathcal{E}_v\langle \tau \rangle
\end{aligned}$$

Figure 15: Call-by-value encoding \mathcal{E}_v into Λ_{ml}

Property 4 (substitution and erasing)

For all $w_0, w_1 \in \text{Terms}[\Lambda_{ml}^{bt}]$,

$$[w_0[x := w_1]] \equiv [w_0][x := [w_1]]$$

Proof: The proof is a simple induction over the structure of w_0 and is omitted. ■

Proof of Lemma 2: The proof is by induction over the derivation of $\Gamma_d \vdash_{bt} e : \tau[w : \varphi]$. We give illustrative cases. Note that for residual constructs, the desired properties follow directly from the IH. The case for succ w below illustrates this.

Recall that given $\Gamma \vdash_{bt} e : \tau[w : \varphi]$, $e \equiv [w]$ (see Section 3). We use this property often without giving explicit reference in the proof below.

case $\Gamma_d \vdash_{bt} x : \Gamma(x).type[x : \Gamma(x).spec-type]$:

We must have $\Gamma(x).spec-type \equiv \mathbf{d}$ since Γ_d contains only dynamic variables. It is possible that cases 4, 5, and 6 of Lemma 2 apply. But in each case, the conditions are satisfied since $x \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$ (because x is dynamic) by definition of $\text{Residual-terms}[\Lambda_{ml}^{bt}]$, and x cannot undergo a specialization step.

case $\Gamma_d \vdash_{bt} \ulcorner n \urcorner : \text{nat}[\ulcorner n \urcorner : \mathbf{s}]$:

Immediate.

case $\Gamma_d \vdash_{bt} \ulcorner n \urcorner : \text{nat}[\underline{\ulcorner n \urcorner} : \mathbf{d}]$:

Case 4 applies and since $\ulcorner n \urcorner \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$ and $\underline{\ulcorner n \urcorner}$ cannot undergo a specialization step, the conditions are satisfied.

case $\Gamma_d \vdash_{bt} \text{succ } e : \widetilde{\text{nat}} [\text{succ } w : \mathfrak{s}]$
because $\Gamma_d \vdash_{bt} e : \text{nat}[w : \mathbf{s}]$:

Case 3 applies. Since subcase 3(a) cannot hold, we show that subcase 3(b) holds. By IH, $w \equiv \ulcorner n \urcorner$ and thus $\text{succ } \ulcorner w \urcorner \mapsto_i \text{unit } \ulcorner n + 1 \urcorner$. Since we have $\Gamma_d \vdash_{bt} [\text{unit } \ulcorner n + 1 \urcorner] : \widetilde{\text{nat}} [\text{unit } \ulcorner n + 1 \urcorner : \mathfrak{s}]$ and we also have $\text{succ } \ulcorner n \urcorner \equiv [\text{succ } \ulcorner n \urcorner] \xrightarrow{ml} [\text{unit } \ulcorner n + 1 \urcorner]$, the conditions of subcase 3(b) are satisfied.

case $\Gamma_d \vdash_{bt} \text{succ } e : \widetilde{\text{nat}} [\underline{\text{succ}} w : \mathbf{d}]$
because $\Gamma_d \vdash_{bt} e : \text{nat}[w : \mathbf{d}]$:

Case 6 applies, and by IH either subcase 4(a) or subcase 4(b) holds for w .

case 4(a):

Then $w \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$ and w cannot undergo a specialization step. The conditions of 6(a) are satisfied since $\text{succ } w \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$ and note $\text{succ } w$ cannot undergo a specialization step.

case 4(b):

Then $w \mapsto_s w', \Gamma_d \vdash_{bt} [w'] : \tau[w' : \mathbf{d}]$, and $e \equiv [w'] \equiv [w] \equiv \ulcorner n \urcorner$ for some number n . Since we have $\text{succ } w \mapsto_s \text{succ } w', \Gamma_d \vdash_{bt} [\text{succ } w'] : \widetilde{\text{nat}} [\text{succ } w' : \mathbf{d}]$, and $\text{succ } e \equiv [\text{succ } w] \xrightarrow{ml} [\text{succ } w']$, the conditions of 6(b) are satisfied.

case $\Gamma_d \vdash_{bt} \lambda x . e' : \tau_1 \rightarrow \tilde{\tau}_2[\lambda x . w' : \varphi_1 \rightarrow \varphi_2]$
because $\Gamma_d, x : \tau_1[\varphi_1] \vdash_{bt} e' : \tilde{\tau}_2[w' : \varphi_2]$:

Case 2 applies and the conditions are satisfied since $w \equiv \lambda x . w', e \equiv [\lambda x . w'] \equiv \lambda x . [w'] \equiv \lambda x . e'$ (by properties of $[\cdot]$ — see Section 3.1.3), and $\Gamma_d, x : \tau_1[\varphi_1] \vdash_{bt} [w'] : \tau_2[w' : \varphi_2]$ — where w and e are as defined in Lemma 2 case 2.

case $\Gamma_d \vdash_{bt} e_0 @ e_1 : \tilde{\tau}_2[w_0 @ w_1 : \varphi_2]$
because $\Gamma_d \vdash_{bt} e_0 : \tau_1 \rightarrow \tilde{\tau}_2[w_0 : \varphi_1 \rightarrow \varphi_2]$ and $\Gamma \vdash_{bt} e_1 : \tau_1[w_1 : \varphi_1]$:

It follows from the relationship between types and specialization types (see Figure 5) that either $\varphi_2 \equiv \varphi_2'$ (i.e., case 3 applies), or $\varphi_2 \equiv \mathbf{d}$ (i.e., case 6 applies).

Case 3:

Note subcase 3(a) cannot hold, so we show subcase 3(b) holds. By IH, $w_0 \equiv \lambda x . w'_0, e_0 \equiv [\lambda x . w'_0] \equiv \lambda x . [w'_0]$, and $\Gamma_d, x : \tau_1[\varphi_1] \vdash_{bt} [w'_0] : \tau_2[w'_0 : \varphi_2]$. Now $(\lambda x . w'_0) @ w_1 \mapsto_i w'_0[x := w_1]$ and $\Gamma_d \vdash_{bt} [w'_0][x := e_1] : \tau_2[w'_0[x := w_1] : \varphi_2]$ by Property 3.

Finally,

$$\begin{aligned} e_0 @ e_1 &\equiv [(\lambda x . w'_0) @ w_1] \\ &\equiv (\lambda x . [w'_0]) @ [w_1] \\ &\xrightarrow{ml} \\ [w'_0][x := [w_1]] &\equiv [w'_0[x := w_1]] \end{aligned}$$

where the last identity follows by Property 4. Thus, the conditions of subcase 3(b) are satisfied.

Case 6:

Note subcase 6(a) cannot hold since $w_0 @ w_1$ is not residual. To show subcase 6(b) holds, one follows exactly the steps used to show subcase 3(b) above.

case $\Gamma_d \vdash_{bt} \mathbf{fix} x. e : \tilde{\tau}[\mathbf{fix} x. w : \tilde{\varphi}]$
because $\Gamma_d, x : \tilde{\tau}[\tilde{\varphi}] \vdash_{bt} e : \tilde{\tau}[w : \tilde{\varphi}]$:

Case 3 applies, and since subcase 3(a) cannot hold, we only need to show that subcase 3(b) holds. We have $\mathbf{fix} x. w \longrightarrow_i w[x := \mathbf{fix} x. w]$ and $\Gamma_d \vdash_{bt} e[x := \mathbf{fix} x. e] : \tilde{\tau}[w[x := \mathbf{fix} x. w] : \tilde{\varphi}]$ by Property 3.

Finally, $\mathbf{fix} x. e \equiv [\mathbf{fix} x. w] \equiv \mathbf{fix} x. [w]$ and $\mathbf{fix} x. [w] \longrightarrow_{ml} [w][x := \mathbf{fix} x. [w]] \equiv [w[x := \mathbf{fix} x. w]]$ where the last identity follows by Property 4. Thus, the conditions of subcase 3(b) are satisfied.

case $\Gamma_d \vdash_{bt} \mathbf{let} x \leftarrow e_1 \mathbf{in} e_2 : \tilde{\tau}_2[\mathbf{let} x \leftarrow w_1 \mathbf{in} w_2 : \varphi_2]$
because $\Gamma_d \vdash_{bt} e_1 : \tilde{\tau}_1[w_1 : \tilde{\varphi}_1]$ and $\Gamma_d, x : \tau_1[\varphi_1] \vdash_{bt} e_2 : \tilde{\tau}_2[w_2 : \varphi_2]$:

It follows from the relationship between types and specialization types (see Figure 5) that either $\varphi_2 \equiv \tilde{\varphi}'_2$ (i.e., case 3 applies), or $\varphi_2 \equiv \mathbf{d}$ (i.e., case 6 applies).

Case 3:

Note subcase 3(a) cannot hold, so we show subcase 3(b) holds. By IH, case 3 applies for w_1 .

Case 3(a) for w_1 :

Then $w_1 \equiv \mathbf{unit} w'_1$, $e_1 \equiv \mathbf{unit} [w'_1]$ and $\Gamma_d \vdash_{bt} [w'_1] : \tau_1[w'_1 : \varphi_1]$.

Now $\mathbf{let} x \leftarrow \mathbf{unit} w'_1 \mathbf{in} w_2 \longmapsto_i w_2[x := w'_1]$ and

$\Gamma_d \vdash_{bt} [w_2][x := [w'_1]] : \tilde{\tau}_2[w'_1[x := w_1] : \varphi_2]$ by Property 3.

Finally,

$$\begin{aligned} \mathbf{let} x \leftarrow e_1 \mathbf{in} e_2 &\equiv [\mathbf{let} x \leftarrow \mathbf{unit} w'_1 \mathbf{in} w_2] \\ &\equiv \mathbf{let} x \leftarrow \mathbf{unit} [w'_1] \mathbf{in} [w_2] \\ &\xrightarrow{ml} \\ &[w_2][x := [w'_1]] \equiv [w_2[x := w'_1]] \end{aligned}$$

where the last identity follows by Property 4. Thus, the conditions of subcase 3(b) are satisfied.

Case 3(b) for w_1 :

Then $w_1 \mapsto_i w'_1$ and $\Gamma_d \vdash_{bt} [w'_1] : \tilde{\tau}_1[w'_1 : \tilde{\varphi}_1]$ and
 $e_1 \equiv [w_1] \xrightarrow{ml} [w'_1]$.
Therefore, $\text{let } x \leftarrow w_1 \text{ in } w_2 \mapsto_i \text{let } x \leftarrow w'_1 \text{ in } w_2$ and
 $\Gamma_d \vdash_{bt} \text{let } x \leftarrow [w'_1] \text{ in } e_2 : \tilde{\tau}_2[\text{let } x \leftarrow w'_1 \text{ in } w_2 : \varphi_2]$ and

$$\begin{aligned} \text{let } x \leftarrow e_1 \text{ in } e_2 &\equiv [\text{let } x \leftarrow w_1 \text{ in } w_2] \\ &\equiv \text{let } x \leftarrow [w_1] \text{ in } [w_2] \\ &\xrightarrow{ml} \\ \text{let } x \leftarrow [w'_1] \text{ in } [w_2] &\equiv [\text{let } x \leftarrow w'_1 \text{ in } w_2]. \end{aligned}$$

Thus, the conditions of subcase 3(b) are satisfied.

Case 6:

Note subcase 6(a) cannot hold since $\text{let } x \leftarrow w_1 \text{ in } w_2$ is not residual.
To show subcase 6(b) holds, one follows exactly the steps used to show subcase 3(b) above.

case $\Gamma_d \vdash_{bt} e : \mathbf{nat}[\text{lift } w : \mathbf{d}]$
because $\Gamma_d \vdash_{bt} e : \mathbf{nat}[w : \mathbf{s}]$:

Case 4 applies, and note that subcase 4(a) cannot hold since $\text{lift } w$ is not residual. So we show that subcase 4(b) holds. By IH, case 1 applies to w , and therefore $w \equiv \ulcorner n \urcorner \equiv e$ for some number n . Now $\text{lift } \ulcorner n \urcorner \xrightarrow{s} \llbracket n \rrbracket$ and $\Gamma_d \vdash_{bt} \ulcorner n \urcorner : \mathbf{nat}[\llbracket n \rrbracket : \mathbf{d}]$ and $e \equiv \llbracket \text{lift } \ulcorner n \urcorner \rrbracket \equiv \llbracket n \rrbracket \equiv \ulcorner n \urcorner$.

■

We now give the proof of Theorem 2 which establishes the correctness of the binding-time analysis constraints.

Proof of Theorem 2: Clearly, either $\text{spec } w \downarrow$ or $\text{spec } w \uparrow$.

Case $\text{spec } w \downarrow$:

By the definition of spec , there exists an $r \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$ such that $\text{spec } w = r$.

Case $\text{spec } w \uparrow$: In this case we have two possibilities.

1. For all w' such that $w \xrightarrow{s}^* w'$, there exists a w'' such that $w' \xrightarrow{s} w''$ (i.e., w heads an infinite series of specialization steps).

2. There exists an $a \notin \text{Residual-terms}[\Lambda_{ml}^{bt}]$ such that $w \mapsto_s^* a$ and there does not exist an a' such that $a \mapsto_s a'$. However, this cannot be the case since by induction over the number of steps in $w \mapsto_s^* a$, Lemma 2 tells up that either $a \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$ or $a \mapsto_s a'$ for some a' .

■

Given Lemma 2, the correctness of binding-time analysis (Theorem 2), and the soundness of specialization (Theorem 3), we can now establish the correctness of partial evaluation (Theorem 4).

Proof of Theorem 4:

Assume $(pe\ e\ \Gamma_{bt}\ \sigma_s)\downarrow$. The definition of partial evaluation (Definition 6) implies that $spec(bta\ \Gamma_{bt}\ e)\sigma_s\downarrow$.

By definition of binding-time analysis (Definition 4), we have

$$\Gamma_{bt} = \Gamma_s \cup \Gamma_d \vdash_{bt} e : \widetilde{\text{nat}}[(bta\ \Gamma_{bt}\ e) : \mathbf{d}]$$

and by repeated application of Property 3 (substitution and binding-time constraints),

$$\Gamma_d \vdash_{bt} e[\sigma_s] : \widetilde{\text{nat}}[(bta\ \Gamma_{bt}\ e)\sigma_s : \mathbf{d}].$$

Now since $spec(bta\ \Gamma_{bt}\ e)\sigma_s\downarrow$, by induction on the number of specialization steps (applying Lemma 2), we have

$$e[\sigma_s] =_{ml} [spec(bta\ \Gamma_{bt}\ e)\sigma_s].$$

We can now reason as follows.

$$\begin{aligned} & e[\sigma_s] =_{ml} [spec(bta\ \Gamma_{bt}\ e)\sigma_s] \\ \Rightarrow e[\sigma_s]\sigma_d &=_{ml} [spec(bta\ \Gamma_{bt}\ e)\sigma_s]\sigma_d \\ & \dots \text{by substitutivity of ml} \\ \Rightarrow e[\sigma_s]\sigma_d &=_{ml} pe\ e\ \Gamma_{bt}\ \sigma_s\sigma_d \\ & \dots \text{by definition of pe (Definition 6)} \\ \Rightarrow e[\sigma_s]\sigma_d &\approx pe\ e\ \Gamma_{bt}\ \sigma_s\sigma_d \\ & \dots \text{by soundness of ml (Theorem 1)} \end{aligned}$$

■

B.2 Control-based binding-time improvements

The section gives the proofs for various partial evaluation properties and theorems occurring in Sections 5. As before, we need a simple property showing how substitution interacts with the extended binding-time rules.

Property 5 (substitution and binding-time constraints)

If $\Gamma, x : \tau_1[\varphi_1] \vdash_{bt}^+ e_0 : \tau_0[w_0 : \varphi_0]$ and $\Gamma \vdash_{bt}^+ e_1 : \tau_1[w_1 : \varphi_1]$ then

$$\Gamma \vdash_{bt}^+ e_0[x := e_1] : \tau_0[w_0[x := w_1] : \varphi_0].$$

Proof: The proof is a simple induction over the derivation of $\Gamma, x : \tau_1[\varphi_1] \vdash_{bt}^+ e_0 : \tau_0[w_0 : \varphi_0]$ and is omitted. \blacksquare

Proof of Lemma 3: The proof is by induction over the derivation of $\Gamma_d \vdash_{bt}^+ e : \tau[w : \varphi]$. Note that only component 3 of Lemma 3 differs from the previously proven Lemma 2. Therefore we need only rework the cases to which component 3 applies and the cases that rely on component 3 in the inductive hypothesis.

Component 3 applies to cases where the resulting annotated terms include **succ** w , **pred** w , $w_0 @ w_1$ (subcase 3(b) applies), and **unit** w (subcase 3(a) applies). In these cases, the proof is exactly the same as for Lemma 2. Component 3 also applies to **let** $x \Leftarrow w_1$ **in** w_2 . This case must be reworked since it may now satisfy subcase 3(b) or subcase 3(d). In addition, the case **let** $x \Leftarrow w_1$ **in** w_2 can now fall under component 3 whereas in Lemma 2 it could not.

case $\Gamma_d \vdash_{bt}^+ \text{let } x \Leftarrow e_1 \text{ in } e_2 : \tilde{\tau}_2[\text{let } x \Leftarrow w_1 \text{ in } w_2 : \varphi_2]$
 because $\Gamma_d \vdash_{bt}^+ e_1 : \tilde{\tau}_1[w_1 : \tilde{\varphi}_1]$ and $\Gamma_d, x : \tau_1[\varphi_1] \vdash_{bt}^+ e_2 : \tilde{\tau}_2[w_2 : \varphi_2]$:

It follows from the relationship between types and specialization types (see Figure 5) that either $\varphi_2 \equiv \tilde{\varphi}'_2$ (i.e., case 3 applies), or $\varphi_2 \equiv \mathbf{d}$ (i.e., case 6 applies).

Case 3:

Note neither subcase 3(a) nor 3(c) can hold, so we show subcase 3(b) holds or 3(d) holds. By IH, case 3 applies for w_1 .

Case 3(a) for w_1 : (proceeding identically to the proof of Lemma 2).

Then $w_1 \equiv \mathbf{unit} \ w'_1$, $e_1 \equiv \mathbf{unit} \ [w'_1]$ and $\Gamma_d \vdash_{bt}^+ [w'_1] : \tau_1[w'_1 : \varphi_1]$.
 Now $\mathbf{let} \ x \leftarrow \mathbf{unit} \ w'_1 \ \mathbf{in} \ w_2 \mapsto_i w_2[x := w'_1]$ and
 $\Gamma_d \vdash_{bt}^+ [w_2][x := [w'_1]] : \tilde{\tau}_2[w'_1[x := w_1] : \varphi_2]$ by Property 5.
 Finally,

$$\begin{aligned} \mathbf{let} \ x \leftarrow e_1 \ \mathbf{in} \ e_2 &\equiv [\mathbf{let} \ x \leftarrow \mathbf{unit} \ w'_1 \ \mathbf{in} \ w_2] \\ &\equiv \mathbf{let} \ x \leftarrow \mathbf{unit} \ [w'_1] \ \mathbf{in} \ [w_2] \\ &\xrightarrow{ml} \\ [w_2][x := [w'_1]] &\equiv [w_2[x := w'_1]] \end{aligned}$$

where the last identity follows by Property 4. Thus, the conditions of subcase 3(b) are satisfied.

Case 3(b) for w_1 :

If $w_1 \not\equiv \mathbf{let} \ y \leftarrow w_3 \ \mathbf{in} \ w_4$ then the proof proceeds as in Lemma 2.
 That is, $w_1 \mapsto_i w'_1$, $\Gamma_d \vdash_{bt}^+ [w'_1] : \tilde{\tau}_1[w'_1 : \tilde{\varphi}_1]$ and
 $e_1 \equiv [w_1] \xrightarrow{ml} [w'_1]$.
 Therefore, $\mathbf{let} \ x \leftarrow w_1 \ \mathbf{in} \ w_2 \mapsto_i \mathbf{let} \ x \leftarrow w'_1 \ \mathbf{in} \ w_2$ and
 $\Gamma_d \vdash_{bt}^+ \mathbf{let} \ x \leftarrow [w'_1] \ \mathbf{in} \ e_2 : \tilde{\tau}_2[\mathbf{let} \ x \leftarrow w'_1 \ \mathbf{in} \ w_2 : \varphi_2]$ and

$$\begin{aligned} \mathbf{let} \ x \leftarrow e_1 \ \mathbf{in} \ e_2 &\equiv [\mathbf{let} \ x \leftarrow w_1 \ \mathbf{in} \ w_2] \\ &\equiv \mathbf{let} \ x \leftarrow [w_1] \ \mathbf{in} \ [w_2] \\ &\xrightarrow{ml} \\ \mathbf{let} \ x \leftarrow [w'_1] \ \mathbf{in} \ [w_2] &\equiv [\mathbf{let} \ x \leftarrow w'_1 \ \mathbf{in} \ w_2]. \end{aligned}$$

Thus, the conditions of subcase 3(b) are satisfied.

If $w_1 \equiv \mathbf{let} \ y \leftarrow w_3 \ \mathbf{in} \ w_4$, then we show that the conditions of subcase 3(d) are satisfied. First of all, note that only subcase 3(d) can be satisfied since the resulting term must undergo a flattening step. We have

$$\begin{array}{ccc} \mathbf{let} \ x \leftarrow (\mathbf{let} \ y \leftarrow w_3 \ \mathbf{in} \ w_4) & \mapsto_f & \mathbf{let} \ y \leftarrow w_3 \\ \mathbf{in} \ w_2 & & \mathbf{in} \ \mathbf{let} \ x \leftarrow w_4 \\ & & \mathbf{in} \ w_2 \end{array}$$

Now $\Gamma_d, x : \tau_1[\varphi_1] \vdash_{bt}^+ e_2 : \tilde{\tau}_2[w_2 : \varphi_2]$ and it follows that
 $\Gamma_d, y : \tau_3[\varphi_3], x : \tau_1[\varphi_1] \vdash_{bt}^+ e_2 : \tilde{\tau}_2[w_2 : \varphi_2]$ since $y \notin FV(w_2)$.
 Now we also have $\Gamma_d, y : \tau_3[\varphi_3] \vdash_{bt}^+ [w_4] : \tilde{\tau}_1[w_4 : \tilde{\varphi}_1]$.
 From this we have

$$\Gamma_d, y : \tau_3[\varphi_3] \vdash_{bt}^+ \mathbf{let} \ x \leftarrow [w_4] \ \mathbf{in} \ [w_2] : \tilde{\tau}_2[\mathbf{let} \ x \leftarrow w_4 \ \mathbf{in} \ w_2 : \varphi_2].$$

Furthermore,

$$\Gamma_d \vdash_{bt}^+ \text{let } y \Leftarrow [w_3] \text{ in let } x \Leftarrow [w_4] \text{ in } [w_2] : \tilde{\tau}_2[\text{let } y \Leftarrow w_3 \text{ in let } x \Leftarrow w_4 \text{ in } w_2 : \varphi_2].$$

Now

$$\begin{aligned} & \text{let } x \Leftarrow (\text{let } y \Leftarrow e_3 \text{ in } e_4) \text{ in } e_2 \\ \equiv & \text{let } x \Leftarrow (\text{let } y \Leftarrow w_3 \text{ in } w_4) \text{ in } w_2 \\ \equiv & \text{let } x \Leftarrow (\text{let } y \Leftarrow [w_3] \text{ in } [w_4]) \text{ in } [w_2] \\ \xrightarrow{ml} & \text{let } y \Leftarrow [w_3] \text{ in let } x \Leftarrow [w_4] \text{ in } [w_2] \\ \equiv & \text{let } y \Leftarrow w_3 \text{ in let } x \Leftarrow w_4 \text{ in } w_2 \end{aligned}$$

Thus, the conditions of subcase 3(d) are satisfied.

Case 3(c) for w_1 (i.e., $w_1 \equiv \text{let } y \Leftarrow w_3 \text{ in } w_4$)

and case 3(d) for w_1 (i.e., $w_1 \equiv \text{let } y \Leftarrow w_3 \text{ in } w_4$):

In either of these cases, $\text{let } x \Leftarrow w_1 \text{ in } w_2$ must undergo a flattening step and the proof proceeds as in the case above (i.e., the conditions of subcase 3(d) are satisfied).

Case 6:

Note subcase 6(a) cannot hold since $\text{let } x \Leftarrow w_1 \text{ in } w_2$ is not residual.

To show subcase 6(b) holds, one follows exactly the steps used to show subcase 3(b) or 3(d) for $\text{let } x \Leftarrow w_1 \text{ in } w_2$ above.

case $\Gamma_d \vdash_{bt}^+ \text{let } x \Leftarrow e_1 \text{ in } e_2 : \tilde{\tau}_2[\text{let } x \Leftarrow w_1 \text{ in } w_2 : \varphi_2]$
 because $\Gamma_d \vdash_{bt}^+ e_1 : \tilde{\tau}_1[w_1 : \mathbf{d}]$ and $\Gamma_d, x : \tau_1[\mathbf{d}] \vdash_{bt}^+ e_2 : \tilde{\tau}_2[w_2 : \varphi_2]$:

It follows from the relationship between types and specialization types (see Figure 5) that either $\varphi_2 \equiv \varphi_2'$ (i.e., case 3 applies), or $\varphi_2 \equiv \mathbf{d}$ (i.e., case 6 applies). If case 6 applies, then we are essentially treating the “unextended” form of the binding time rule for residual let’s — the proof proceeds as in Lemma 2. If case 3 applies, then it is immediately obvious that the conditions for subcase 3(c) (and only this subcase) are satisfied.

■

Proof: Part (b) follows immediately from the preconditions and the definition of well-formed context stack. For part (a), we show only the “only if” direction which goes by induction over the structure of the derivation. The “if” direction is similar. We only consider the details of the cases for *(Push)* and *(Pop)* since these are the only rules that modify the context stack. The other cases are trivial or following immediately from the inductive hypothesis.

case *(Push)*:

$$\langle \text{let } x'_1 \Leftarrow w'_1, \vec{E}_{i_1} \text{ ++ } (\text{let } x_1 \Leftarrow \cdot \quad \quad \quad \text{:: } \vec{E}_{i_2}) \rangle \Downarrow a$$

$$\text{in } w'_2 \qquad \qquad \qquad \text{in let } x_2 \Leftarrow w_2 \qquad \qquad \text{in } w_3$$

because

$$\langle w'_1, \text{let } x'_1 \Leftarrow \cdot \quad \quad \quad \text{:: } \vec{E}_{i_1} \text{ ++ } (\text{let } x_1 \Leftarrow \cdot \quad \quad \quad \text{:: } \vec{E}_{i_2}) \rangle \Downarrow a:$$

$$\text{in } w'_2 \qquad \qquad \qquad \text{in let } x_2 \Leftarrow w_2 \qquad \qquad \text{in } w_3$$

By the restriction on the *(Push)* rule,

$$x'_1 \notin FV(\vec{E}_{i_1} \text{ ++ } (\text{let } x_1 \Leftarrow \cdot \quad \quad \quad \text{:: } \vec{E}_{i_2})),$$

$$\text{in let } x_2 \Leftarrow w_2$$

$$\text{in } w_3$$

so this context stack is well-formed.

By IH,

$$\langle w'_1, \text{let } x'_1 \Leftarrow \cdot \quad \quad \quad \text{:: } \vec{E}_{i_1} \text{ ++ } (\text{let } x_1 \Leftarrow \cdot \quad \quad \quad \text{:: } \text{let } x_2 \Leftarrow \cdot \quad \quad \quad \text{:: } \vec{E}_{i_2}) \rangle \Downarrow a$$

$$\text{in } w'_2 \qquad \qquad \qquad \text{in } w_2 \qquad \qquad \qquad \text{in } w_3$$

and thus

$$\langle \text{let } x'_1 \Leftarrow w'_1, \vec{E}_{i_1} \text{ ++ } (\text{let } x_1 \Leftarrow \cdot \quad \quad \quad \text{:: } \text{let } x_2 \Leftarrow \cdot \quad \quad \quad \text{:: } \vec{E}_{i_2}) \rangle \Downarrow a$$

$$\text{in } w'_2 \qquad \qquad \qquad \text{in } w_2 \qquad \qquad \qquad \text{in } w_3$$

case *(Pop)*: if $\vec{E}_{i_1} \neq \text{nil}$ then the result follows from the IH.

If $\vec{E}_{i_1} = \text{nil}$ then we have

$$\langle \text{unit } w_1, \text{let } x_1 \Leftarrow \cdot \quad \quad \quad \text{:: } \vec{E}_{i_2} \rangle \Downarrow a$$

$$\text{in let } x_2 \Leftarrow w_2$$

$$\text{in } w_3$$

because $\langle \text{let } x_2 \Leftarrow w_2[x_1 := w_1] \text{ in } w_3[x_1 := w_1], \vec{E}_{i_2} \rangle \Downarrow a$
...by (Pop)

because $\langle w_2[x_1 := w_1], \text{let } x_2 \Leftarrow \cdot \quad \quad \quad \text{:: } \vec{E}_{i_2} \rangle \Downarrow a$
...by (Push)

Now since $x_1 \notin FV(w_3)$, $\text{let } x_2 \Leftarrow \cdot \text{ in } w_3[x_1 := w_1] \equiv \text{let } x_2 \Leftarrow \cdot \text{ in } w_3$ and we have

$$\begin{aligned} & \langle w_2[x_1 := w_1], \text{let } x_2 \Leftarrow \cdot \text{ in } w_3 \text{ :: } \vec{E}_{i_2} \rangle \Downarrow a \\ & \quad \text{in } w_3 \\ & \dots \text{from above} \\ \text{implies } & \langle \text{unit } w_1, \text{let } x_1 \Leftarrow \cdot \text{ in } w_2 \text{ :: let } x_2 \Leftarrow \cdot \text{ in } w_3 \text{ :: } \vec{E}_{i_2} \rangle \Downarrow a \\ & \dots \text{by (Pop)} \end{aligned}$$

■

The following property states that residual terms are among the canonical terms of specialization.

Property 7 (residual terms are canonical)

For all $r \in \text{Residual-terms}[\Lambda_{ml}^{bt}]$,

$$\langle r, \text{nil} \rangle \Downarrow r.$$

Proof:

It is easy to see that the set $\text{Residual-terms}[\Lambda_{ml}^{bt}] \subset \text{Terms}[\Lambda_{ml}^{bt}]$ is exactly generated by the following grammar

$$\begin{aligned} r ::= & x \mid \ulcorner n \urcorner \mid \text{succ } r \mid \text{pred } r \mid \text{if0 } r_1 r_2 r_3 \mid \lambda x. r \mid r_0 @ r_1 \mid \\ & \text{fix } x. r \mid \text{unit } r \mid \text{let } x \Leftarrow r_1 \text{ in } r_2 \end{aligned}$$

The proof follows by a simple induction over the structure of terms in this grammar. ■

The following property shows that each step in the flattening specializer is appropriately reflected in the continuation-based specializer.

Property 8 If $w \mapsto_{s+} w'$ and $\langle w', \text{nil} \rangle \Downarrow a$, then $\langle w, \text{nil} \rangle \Downarrow a$.

Proof: by induction over the structure of w , proceeding by cases according to the last rule used in $w \mapsto_{s+} w'$. We show only illustrative cases.

case $\text{succ } \ulcorner n \urcorner \mapsto_{s+} \text{unit } \ulcorner n + 1 \urcorner$:

by assumption we have $\langle \text{unit } \ulcorner n + 1 \urcorner, \text{nil} \rangle \Downarrow a$ and thus we can conclude $\langle \text{succ } \ulcorner n \urcorner, \text{nil} \rangle \Downarrow a$.

case $\text{let } x_1 \Leftarrow \text{succ } \ulcorner n \urcorner \text{ in } w_2 \mapsto_{s+} \text{let } x_1 \Leftarrow \text{unit } \ulcorner n + 1 \urcorner \text{ in } w_2$:

by assumption

$$\begin{aligned} & \langle \text{let } x_1 \Leftarrow \text{unit } \ulcorner n + 1 \urcorner \text{ in } w_2, \text{nil} \rangle \Downarrow a \\ \text{thus } & \langle \text{unit } \ulcorner n + 1 \urcorner, \text{let } x_1 \Leftarrow \cdot \text{ :: nil} \rangle \Downarrow a \quad \dots \text{by (Push)} \\ & \text{in } w_2 \end{aligned}$$

Now we need to show

$$\begin{aligned} & \langle \text{let } x_1 \Leftarrow \text{succ } \ulcorner n \urcorner \text{ in } w_2, \text{nil} \rangle \Downarrow a \\ \text{which requires } & \langle \text{succ } \ulcorner n \urcorner, \text{let } x_1 \Leftarrow \cdot \text{ :: nil} \rangle \Downarrow a \\ & \text{in } w_2 \\ & \dots \text{by (Push)} \\ \text{which requires } & \langle \text{unit } \ulcorner n + 1 \urcorner, \text{let } x_1 \Leftarrow \cdot \text{ :: nil} \rangle \Downarrow a \\ & \text{in } w_2 \end{aligned}$$

and we have this from above.

case $\text{let } x_2 \Leftarrow (\text{let } x_1 \Leftarrow w_1 \text{ in } w_2) \text{ in } w_3 \mapsto_{s+} \text{let } x_1 \Leftarrow w_1 \text{ in let } x_2 \Leftarrow w_2 \text{ in } w_3$

where $x_1 \notin FV(w_3)$:

by assumption

$$\begin{aligned} & \langle \text{let } x_1 \Leftarrow w_1 \text{ in let } x_2 \Leftarrow w_2 \text{ in } w_3, \text{nil} \rangle \Downarrow a \\ \text{thus } & \langle w_1, \text{let } x_1 \Leftarrow \cdot \text{ :: nil} \rangle \Downarrow a \\ & \text{in let } x_2 \Leftarrow w_2 \\ & \text{in } w_3 \\ & \dots \text{by (Push)} \\ \text{and } & \langle w_1, \text{let } x_1 \Leftarrow \cdot \text{ :: let } x_2 \Leftarrow \cdot \text{ :: nil} \rangle \Downarrow a \\ & \text{in } w_2 \quad \text{in } w_3 \\ & \dots \text{by Property 6} \end{aligned}$$

Now we need to show

$$\begin{aligned} & \langle \text{let } x_2 \Leftarrow (\text{let } x_1 \Leftarrow w_1 \text{ in } w_2) \text{ in } w_3, \text{nil} \rangle \Downarrow a \\ \text{which requires } & \langle \text{let } x_1 \Leftarrow w_1 \text{ in } w_2, \text{let } x_2 \Leftarrow \cdot \text{ :: nil} \rangle \Downarrow a \\ & \text{in } w_3 \\ & \dots \text{by (Push)} \\ \text{which requires } & \langle w_1, \text{let } x_1 \Leftarrow \cdot \text{ :: let } x_2 \Leftarrow \cdot \text{ :: nil} \rangle \Downarrow a \\ & \text{in } w_2 \quad \text{in } w_3 \\ & \dots \text{by (Push)} \end{aligned}$$

and we have this from above.

case $\text{let } x_2 \Leftarrow (\text{let } x_1 \Leftarrow w_1 \text{ in } w_2) \text{ in } w_3 \mapsto_{s+} \text{let } x_1 \Leftarrow w_1$
 $\text{in let } x_2 \Leftarrow w_2$
 $\text{in } w_3$

where $x_1 \notin FV(w_3)$:

by assumption

$\langle \text{let } x_1 \Leftarrow w_1 \text{ in let } x_2 \Leftarrow w_2 \text{ in } w_3, nil \rangle \Downarrow a$
 thus $\langle w_1, nil \rangle \Downarrow r_1$
 and $\langle \text{let } x_2 \Leftarrow w_2 \text{ in } w_3, nil \rangle \Downarrow r_2$ *by (Assoc)*
 and thus $\langle w_2, \text{let } x_2 \Leftarrow \cdot :: nil \rangle \Downarrow r_2$ *by (Push)*
 $\text{in } w_3$
 and note $a \equiv \text{let } x_1 \Leftarrow r_1 \text{ in } r_2$

Now we need to show

$\langle \text{let } x_2 \Leftarrow (\text{let } x_1 \Leftarrow w_1 \text{ in } w_2) \text{ in } w_3, nil \rangle \Downarrow a$
 which requires $\langle \text{let } x_1 \Leftarrow w_1, \text{let } x_2 \Leftarrow \cdot :: nil \rangle \Downarrow a$
 $\text{in } w_2$ $\text{in } w_3$
...by (Push)
 which requires $\langle w_1, nil \rangle \Downarrow r_1$
 and $\langle w_2, \text{let } x_2 \Leftarrow \cdot :: nil \rangle \Downarrow r_2$
 $\text{in } w_3$
...by (Assoc)

and we have both of these from above.

case $\text{if0 } w_1 w_2 w_3 \mapsto_{s+} \text{if0 } w'_1 w_2 w_3$ because $w_1 \mapsto_{s+} w'_1$:

by assumption

$\langle \text{if0 } w'_1 w_2 w_3, nil \rangle \Downarrow a$
 thus $\langle w'_1, nil \rangle \Downarrow r_1$
 and $\langle w_2, nil \rangle \Downarrow r_2$
 and $\langle w_3, nil \rangle \Downarrow r_3$
 and note $a \equiv \text{if0 } r_1 r_2 r_3$

Now we need to show

$\langle \text{if0 } w_1 w_2 w_3, nil \rangle \Downarrow a$
 which requires $\langle w_1, nil \rangle \Downarrow r_1$ (1)
 and $\langle w_2, nil \rangle \Downarrow r_2$ (2)
 and $\langle w_3, nil \rangle \Downarrow r_3$ (3).

Now (1) follows from IH, and (2) and (3) follow from above.

case $\text{lift } \ulcorner n \urcorner \mapsto_{s+} \ulcorner n \urcorner$:

by assumption we have $\langle \ulcorner n \urcorner, nil \rangle \Downarrow a$ where $a \equiv \ulcorner n \urcorner$ and we have $\langle \text{lift } \ulcorner n \urcorner, nil \rangle \Downarrow \ulcorner n \urcorner$ directly. ■

We can now prove part (1) (as enumerated above) of Theorem 5.

Lemma 4 $w \mapsto_{s^+}^* r$ implies $\langle w, nil \rangle \Downarrow r$.

Proof: by induction over the number of steps n in $w \mapsto_{s^+}^* r$.

case $n = 0$:

we need to show $\langle r, nil \rangle \Downarrow r$ and this follows since residual terms are canonical (Property 7).

case $n = i + 1$: that is, $w \mapsto_{s^+} w' \mapsto_{s^+}^* r$.

by IH, $\langle w', nil \rangle \Downarrow r$ and so by Property 8, $\langle w, nil \rangle \Downarrow r$. ■

Now we turn to part (2) in the proof of Theorem 5. We need to show that, given $\langle w, nil \rangle \Downarrow r$, we have a corresponding sequence of specialization steps in the flattening specializer. As expected, the proof proceeds by induction over the derivation of $\langle w, nil \rangle \Downarrow r$. However, the inductive hypothesis must be strengthened to deal with configurations of the form $\langle w, \vec{E}_i \rangle \Downarrow a$. In this stronger case, we need to show that a term corresponding to w in the interpretation context associated with \vec{E}_i specializes to a .

The following definition shows how context stacks can be unstacked to yield a flattened interpretation context.

Definition 9 (unstacking)

$$\begin{aligned} \text{unstack}(nil) &= [\cdot] \\ \text{unstack}(\text{let } x_1 \leftarrow \cdot \text{ :: } \vec{E}_i) &= \text{let } x_1 \leftarrow \cdot \\ &\text{in } w_2 \qquad \qquad \qquad \text{in } \text{unstack}(\vec{E}_i)[w_2] \end{aligned}$$

For example, unstacking the context stack

$$\begin{aligned} \text{let } x_1 \leftarrow \cdot \text{ :: } \text{let } x_2 \leftarrow \cdot \text{ :: } nil \\ \text{in } w_2 \qquad \qquad \qquad \text{in } w_3 \end{aligned}$$

yields the interpretation context

$$\begin{array}{l} \text{let } x_1 \Leftarrow \cdot \\ \text{in let } x_2 \Leftarrow w_2 \\ \text{in } w_3. \end{array}$$

Note that nested **let**'s appearing in w_2 are not necessarily flattened.

Observation 1 (contexts yielded by unstacking)

For all \vec{E}_i , exactly one of the following holds:

1. $\text{unstack}(\vec{E}_i) = [\cdot]$, or
2. $\text{unstack}(\vec{E}_i) = E'_i$ for some interpretation context E'_i .

The following lemma shows the flattening specialization steps induced by continuation-based specializer derivations.

Lemma 5 *Let \vec{E}_i be a well-formed context stack.*

$$\langle w, \vec{E}_i \rangle \Downarrow a \quad \text{implies} \quad \text{unstack}(\vec{E}_i)[w] \mapsto_{s^+}^* a$$

Proof: by induction over the derivation of $\langle w, \vec{E}_i \rangle \Downarrow a$. Some representative cases are given below. We will implicitly rely on Property 2 which states that all context stacks appearing as subderivations of $\langle w, \vec{E}_i \rangle \Downarrow a$ are well-formed.

case $\langle \text{unit } w, \text{nil} \rangle \Downarrow \text{unit } w$:

$$\text{unstack}(\text{nil})[\text{unit } w] = \text{unit } w \quad \text{and} \quad \text{unit } w \mapsto_{s^+}^* \text{unit } w \quad \text{by reflexivity of } \mapsto_{s^+}^*.$$

case $\langle \text{succ } \ulcorner n \urcorner, \vec{E}_i \rangle \Downarrow a$ because $\langle \text{unit } \ulcorner n + 1 \urcorner, \vec{E}_i \rangle \Downarrow a$:

$$\begin{array}{ll} \text{unstack}(\vec{E}_i)[\text{succ } \ulcorner n \urcorner] & \\ \mapsto_{s^+} \text{unstack}(\vec{E}_i)[\text{unit } \ulcorner n + 1 \urcorner] & \dots \text{relying on Observation 1} \\ \mapsto_{s^+}^* a & \dots \text{by IH} \end{array}$$

case $\langle \text{let } x_1 \Leftarrow w_1 \text{ in } w_2, \vec{E}_i \rangle \Downarrow a$ because $\langle w_1, \text{let } x_1 \Leftarrow \cdot \text{ :: } \vec{E}_i \rangle \Downarrow a$:
 $\text{in } w_2$

Note $x_1 \notin FV(\vec{E}_i)$ by the side-condition and proceed by cases of \vec{E}_i .
 case $\vec{E}_i \equiv \text{nil}$:

$$\begin{aligned}
& \text{unstack}(\vec{E}_i)[\text{let } x_1 \Leftarrow w_1 \text{ in } w_2] \\
&= \text{let } x_1 \Leftarrow w_1 \text{ in } w_2 \quad \dots \text{by def. of unstack} \\
&= \text{unstack}(\text{let } x_1 \Leftarrow \cdot \text{ :: } \vec{E}_i)[w_1] \quad \dots \text{by def. of unstack} \\
&\quad \text{in } w_2 \\
&\xrightarrow{s+^*} a \quad \dots \text{by IH}
\end{aligned}$$

case $\vec{E}_i \equiv \text{let } x_2 \Leftarrow \cdot \text{ :: } \vec{E}'_i$
in w_3

Intuitively, we unfold the definition of *unstack*, apply the flattening rule, and then fold *unstack*.

$$\begin{aligned}
& \text{unstack}(\vec{E}_i)[\text{let } x_1 \Leftarrow w_1 \text{ in } w_2] \\
&= \text{let } x_2 \Leftarrow (\text{let } x_1 \Leftarrow w_1 \text{ in } w_2) \text{ in } \text{unstack}(\vec{E}'_i)[w_3] \\
&\quad \dots \text{by definition of unstack} \\
&\xrightarrow{s+} \text{let } x_1 \Leftarrow w_1 \text{ in let } x_2 \Leftarrow w_2 \text{ in } \text{unstack}(\vec{E}'_i)[w_3] \\
&\quad \dots \text{note } x_1 \notin FV(\vec{E}'_i) \cup FV(w_3) \\
&= \text{let } x_1 \Leftarrow w_1 \text{ in } \text{unstack}(\text{let } x_2 \Leftarrow \cdot \text{ :: } \vec{E}'_i)[w_2] \\
&\quad \text{in } w_3 \\
&\quad \dots \text{by definition of unstack} \\
&= \text{unstack}(\text{let } x_1 \Leftarrow \cdot \text{ :: } \text{let } x_2 \Leftarrow \cdot \text{ :: } \vec{E}'_i)[w_1] \\
&\quad \text{in } w_2 \quad \text{in } w_3 \\
&\quad \dots \text{by definition of unstack} \\
&\equiv \text{unstack}(\text{let } x_1 \Leftarrow \cdot \text{ :: } \vec{E}_i)[w_1] \\
&\quad \text{in } w_2 \\
&\xrightarrow{s+^*} a \quad \dots \text{by IH}
\end{aligned}$$

case $\langle \text{unit } w_1, \text{let } x \Leftarrow \cdot \text{ :: } \vec{E}_i \rangle \Downarrow a$
in w_2

because $\langle w_2[x := w_1], \vec{E}_i \rangle \Downarrow a$:

$$\begin{aligned}
& \text{unstack}(\text{let } x \leftarrow \cdot \text{ :: } \vec{E}_i)[\text{unit } w_1] \\
& \quad \text{in } w_2 \\
& = \quad \text{let } x_1 \leftarrow \text{unit } w_1 \text{ in } \text{unstack}(\vec{E}_i)[w_2] \\
& \quad \quad \dots \text{definition of unstack} \\
& \xrightarrow{s+} \text{unstack}(\vec{E}_i)[w_2][x_1 := w_1] \\
& \equiv \quad \text{unstack}(\vec{E}_i)[w_2[x_1 := w_1]] \\
& \quad \quad \dots \text{since } \text{let } x \leftarrow \cdot \text{ :: } \vec{E}_i \\
& \quad \quad \quad \text{in } w_2 \\
& \quad \quad \dots \text{is a well-formed stack and thus } x_1 \notin FV(\vec{E}_i) \\
& \xrightarrow{s+} a \\
& \quad \quad \dots \text{by IH}
\end{aligned}$$

case $\langle w_1 @ w_2, nil \rangle \Downarrow r_1 @ r_2$
because $\langle w_1, nil \rangle \Downarrow r_1$ and $\langle w_2, nil \rangle \Downarrow r_2$:

$$\begin{aligned}
& \text{unstack}(nil)[w_1 @ w_2] \\
& = \quad w_1 @ w_2 \quad \dots \text{definition of unstack} \\
& \xrightarrow{s+}^* r_1 @ w_2 \quad \dots \text{by IH} \\
& \xrightarrow{s+}^* r_1 @ r_2 \quad \dots \text{by IH}
\end{aligned}$$

case $\langle \text{let } x \leftarrow w_1 \text{ in } w_2, \vec{E}_i \rangle \Downarrow \text{let } x \leftarrow r_1 \text{ in } r_2$ because $\langle w_1, nil \rangle \Downarrow r_1$ and $\langle w_2, \vec{E}_i \rangle \Downarrow r_2$

Note $x_1 \notin FV(\vec{E}_i)$ and proceed by cases of \vec{E}_i .

case $\vec{E}_i \equiv nil$:

$$\begin{aligned}
& \text{unstack}(\vec{E}_i)[\text{let } x_1 \leftarrow w_1 \text{ in } w_2] \\
& = \quad \text{let } x_1 \leftarrow w_1 \text{ in } w_2 \quad \dots \text{by definition of unstack} \\
& \xrightarrow{s+}^* \text{let } x_1 \leftarrow r_1 \text{ in } w_2 \quad \dots \text{by IH} \\
& \xrightarrow{s+}^* \text{let } x_1 \leftarrow r_1 \text{ in } r_2 \quad \dots \text{by IH}
\end{aligned}$$

case $\vec{E}_i \equiv \text{let } x_2 \leftarrow \cdot \text{ :: } \vec{E}_i'$
in w_3

Intuitively, we unfold the definition of *unstack*, apply the flattening rule, and then fold *unstack*.

$$\begin{aligned}
& \text{unstack}(\vec{E}_i)[\underline{\text{let}}\ x_1 \Leftarrow w_1 \ \underline{\text{in}}\ w_2] \\
&= \quad \underline{\text{let}}\ x_2 \Leftarrow (\underline{\text{let}}\ x_1 \Leftarrow w_1 \ \underline{\text{in}}\ w_2) \ \underline{\text{in}}\ \text{unstack}(\vec{E}_i')[w_3] \\
&\quad \dots \text{by definition of unstack} \\
&\xrightarrow{s+} \underline{\text{let}}\ x_1 \Leftarrow w_1 \ \underline{\text{in}}\ \underline{\text{let}}\ x_2 \Leftarrow w_2 \ \underline{\text{in}}\ \text{unstack}(\vec{E}_i')[w_3] \\
&\quad \dots \text{note } x_1 \notin FV(\vec{E}_i') \cup FV(w_3) \\
&\xrightarrow{s+}^* \underline{\text{let}}\ x_1 \Leftarrow r_1 \ \underline{\text{in}}\ \underline{\text{let}}\ x_2 \Leftarrow w_2 \ \underline{\text{in}}\ \text{unstack}(\vec{E}_i')[w_3] \\
&\quad \dots \text{by IH} \\
&= \quad \underline{\text{let}}\ x_1 \Leftarrow r_1 \ \underline{\text{in}}\ \text{unstack}(\underline{\text{let}}\ x_2 \Leftarrow \cdot \ \ :: \ \vec{E}_i')[w_2] \\
&\quad \quad \quad \underline{\text{in}}\ w_3 \\
&\quad \dots \text{by definition of unstack} \\
&\xrightarrow{s+}^* \underline{\text{let}}\ x_1 \Leftarrow r_1 \ \underline{\text{in}}\ r_2 \\
&\quad \dots \text{by IH}
\end{aligned}$$

■

As a corollary, we obtain the proof of part (2) of Theorem 5.

Corollary 1 $\langle w, \text{nil} \rangle \Downarrow r$ implies $w \xrightarrow{s+}^* r$.

Proof: Since nil is a well-formed context stack, by Lemma 5 we have $\text{unstack}(\text{nil})[w] \xrightarrow{s+}^* r$. Now since $\text{unstack}(\text{nil}) = [\cdot]$, we have $w \xrightarrow{s+}^* r$. ■

References

- [1] Lars Ole Andersen. *Program Analysis and Specialization for the C Programming Language*. PhD thesis, DIKU, Computer Science Department, University of Copenhagen, Copenhagen, Denmark, May 1994. DIKU Rapport 94/19.
- [2] Henk Barendregt. *The Lambda Calculus — Its Syntax and Semantics*. North-Holland, 1984.
- [3] P. N. Benton. *Strictness Analysis of Lazy Functional Programs*. PhD thesis, Computer Laboratory, University of Cambridge, Cambridge, England, May 1995.
- [4] P. N. Benton, G. M. Bierman, and V. C. V de Paiva. Computational types from a logical perspective I. Technical report 365, Computer Laboratory, University of Cambridge, Cambridge, England, May 1995.

- [5] Hans-J. Boehm, editor. *Proceedings of the Twenty-First Annual ACM Symposium on Principles of Programming Languages*, Portland, Oregon, January 1994. ACM Press.
- [6] Anders Bondorf. Improving binding times without explicit cps-conversion. In William Clinger, editor, *Proceedings of the 1992 ACM Conference on Lisp and Functional Programming*, LISP Pointers, Vol. V, No. 1, pages 1–10, San Francisco, California, June 1992. ACM Press.
- [7] Anders Bondorf and Olivier Danvy. Automatic autoprojection of recursive equations with global variables and abstract data types. *Science of Computer Programming*, 16:151–195, 1991.
- [8] Anders Bondorf and Dirk Dussart. Improving CPS-based partial evaluation: Writing cogen by hand. In Peter Sestoft and Harald Søndergaard, editors, *Proceedings of the ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation*, Technical Report 94/9, University of Melbourne, Australia, pages 1–10, Orlando, Florida, June 1994.
- [9] Charles Consel. A tour of Schism: A partial evaluation system for higher-order applicative languages. In David A. Schmidt, editor, *Proceedings of the Second ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, pages 145–154, Copenhagen, Denmark, June 1993. ACM Press.
- [10] Charles Consel and Olivier Danvy. For a better support of static data flow. In Hughes [40], pages 496–519.
- [11] Charles Consel and Olivier Danvy. Tutorial notes on partial evaluation. In Susan L. Graham, editor, *Proceedings of the Twentieth Annual ACM Symposium on Principles of Programming Languages*, pages 493–501, Charleston, South Carolina, January 1993. ACM Press.
- [12] Roy L. Crole and Andrew D. Gordon. Factoring an adequacy proof (preliminary report). In John T. O’Donnell and Kevin Hammond, editors, *Functional Programming, Glasgow 1993*, Workshops in Computing, pages 9–25, Ayr, Scotland, 1993. Springer-Verlag.

- [13] Roy L. Crole and Andrew M. Pitts. New foundations for fixpoint computations: FIX-hyperdoctrines and the FIX-logic. *Information and Computation*, 98:171–210, 1992.
- [14] Olivier Danvy. Type-directed partial evaluation. In Guy L. Steele Jr., editor, *Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Programming Languages*, pages 242–257, St. Petersburg Beach, Florida, January 1996. ACM Press.
- [15] Olivier Danvy and Andrzej Filinski. Representing control, a study of the CPS transformation. *Mathematical Structures in Computer Science*, 2(4):361–391, December 1992.
- [16] Olivier Danvy and John Hatcliff. CPS transformation after strictness analysis. *ACM Letters on Programming Languages and Systems*, 1(3):195–212, 1993.
- [17] Olivier Danvy, Karoline Malmkjær, and Jens Palsberg. The essence of eta-expansion in partial evaluation. *LISP and Symbolic Computation*, 8(3):209–227, 1995. An earlier version appeared in the proceedings of the 1994 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation.
- [18] Olivier Danvy, Karoline Malmkjær, and Jens Palsberg. Eta-expansion does The Trick. *ACM Transactions on Programming Languages and Systems*, 1996. To appear.
- [19] Rowan Davies. A temporal-logic approach to binding-time analysis. In *Proceedings of the Eleventh Annual IEEE Symposium on Logic in Computer Science*, pages 184–195, New Brunswick, New Jersey, July 1996. IEEE Computer Society Press.
- [20] Rowan Davies and Frank Pfenning. A modal analysis of staged computation. In Guy L. Steele Jr., editor, *Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Programming Languages*, pages 258–283, St. Petersburg Beach, Florida, January 1996. ACM Press.
- [21] Dirk Dussart. PhD thesis, Katholieke Universiteit Leuven, Leuven, Belgium, 1997. Forthcoming.

- [22] Dirk Dussart and Peter Thiemann. Imperative functional specialization. Berichte des Wilhelm-Schickard-Instituts WSI-96-28, Universität Tübingen, July 1996.
- [23] R. Kent Dybvig, editor. *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming*, Philadelphia, Pennsylvania, May 1996. ACM Press.
- [24] Andrzej Filinski. Representing monads. In Boehm [5], pages 446–457.
- [25] Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. The essence of compiling with continuations. In David W. Wall, editor, *Proceedings of the ACM SIGPLAN'93 Conference on Programming Languages Design and Implementation*, SIGPLAN Notices, Vol. 28, No 6, pages 237–247, Albuquerque, New Mexico, June 1993. ACM Press.
- [26] Jean-Yves Girard, Yves Lafont, and Paul Taylor. *Proofs and Types*, volume 7 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1989.
- [27] Robert Glück and Jesper Jørgensen. Efficient multi-level generating extensions for program specialization. In Hermenegildo and Swierstra [38], pages 259–278.
- [28] Robert Glück and Andrei Klimov. Occam's razor in metacomputation: the notion of a perfect process tree. In Patrick Cousot, Moreno Falaschi, Gilberto Filè, and Antoine Rauzy, editors, *Proceedings of the Third International Workshop on Static Analysis WSA '93*, number 724 in Lecture Notes in Computer Science, pages 112–123, Padova, Italy, September 1993.
- [29] Carsten K. Gomard. A self-applicable partial evaluator for the lambda-calculus: Correctness and pragmatics. *ACM Transactions on Programming Languages and Systems*, 14(2):147–172, 1992.
- [30] Carsten K. Gomard and Neil D. Jones. A partial evaluator for the untyped lambda-calculus. *Journal of Functional Programming*, 1(1):21–69, 1991.
- [31] Andrew D. Gordon. *Functional Programming and Input/Output*. PhD thesis, Computer Laboratory, University of Cambridge, Cambridge, England, February 1993. Technical Report No. 285.

- [32] Carl A. Gunter. *Semantics of Programming Languages: Structures and Techniques*. MIT Press, 1992.
- [33] John Hannan and Dale Miller. Deriving mixed evaluation from standard evaluation for a simple functional language. In J. L. A. van de Snepscheut, editor, *Proceedings of a Conference on Mathematics of Program Construction*, number 375 in Lecture Notes in Computer Science, pages 239–255, Groningen, The Netherlands, 1989.
- [34] John Hatcliff. *The Structure of Continuation-Passing Styles*. PhD thesis, Department of Computing and Information Sciences, Kansas State University, Manhattan, Kansas, June 1994.
- [35] John Hatcliff. Mechanically verifying the correctness of an offline partial evaluator. In Hermenegildo and Swierstra [38], pages 279–298.
- [36] John Hatcliff and Olivier Danvy. A generic account of continuation-passing styles. In Boehm [5], pages 458–471.
- [37] John Hatcliff and Robert Glück. An operational theory of self-applicable on-line program specialization. In Olivier Danvy, Robert Glück, and Peter Thiemann, editors, *Partial Evaluation*, number 1110 in Lecture Notes in Computer Science, pages 161–182, Dagstuhl, Germany, February 1996.
- [38] Manuel Hermenegildo and S. Doaitse Swierstra, editors. *Seventh International Symposium on Programming Language Implementation and Logic Programming*, number 982 in Lecture Notes in Computer Science, Utrecht, The Netherlands, September 1995.
- [39] Carsten K. Holst and Carsten K. Gomard. Partial evaluation is fuller laziness. In Paul Hudak and Neil D. Jones, editors, *Proceedings of the ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, SIGPLAN Notices, Vol. 26, No 9, pages 223–233, New Haven, Connecticut, June 1991. ACM Press.
- [40] John Hughes, editor. *Proceedings of the Fifth ACM Conference on Functional Programming and Computer Architecture*, number 523 in Lecture Notes in Computer Science, Cambridge, Massachusetts, August 1991. Springer-Verlag.

- [41] Neil D. Jones, editor. *Special issue on Partial Evaluation*, Journal of Functional Programming, Vol. 3, Part 3. Cambridge University Press, July 1993.
- [42] Neil D. Jones, Carsten K. Gomard, and Peter Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice Hall International Series in Computer Science. Prentice-Hall, 1993.
- [43] Neil D. Jones, Peter Sestoft, and Harald Søndergaard. MIX: A self-applicable partial evaluator for experiments in compiler generation. *LISP and Symbolic Computation*, 2(1):9–50, 1989.
- [44] John Launchbury. A strongly-typed self-applicable partial evaluator. In Hughes [40], pages 145–164.
- [45] Julia L. Lawall and Olivier Danvy. Continuation-based partial evaluation. In Carolyn L. Talcott, editor, *Proceedings of the 1994 ACM Conference on Lisp and Functional Programming*, LISP Pointers, Vol. VII, No. 3, Orlando, Florida, June 1994. ACM Press.
- [46] Julia L. Lawall and Olivier Danvy. Continuation-based partial evaluation. Technical Report CS-95-178, Computer Science Department, Brandeis University, Waltham, Massachusetts, January 1995. An earlier version appeared in the proceedings of the 1994 ACM Conference on Lisp and Functional Programming.
- [47] Sheng Liang, Paul Hudak, and Mark Jones. Monad transformers and modular interpreters. In Peter Lee, editor, *Proceedings of the Twenty-Second Annual ACM Symposium on Principles of Programming Languages*, pages 333–343, San Francisco, California, January 1995. ACM Press.
- [48] Eugenio Moggi. Notions of computation and monads. *Information and Computation*, 93:55–92, 1991.
- [49] Kristian Nielsen. Master’s thesis, DIKU, Computer Science Department, University of Copenhagen, Copenhagen, Denmark, 1997. Forthcoming.
- [50] Kristian Nielsen and Morten Heine Sørensen. Call-by-name CPS-translation as a binding-time improvement. In Alan Mycroft, editor, *Static Analysis*, number 983 in Lecture Notes in Computer Science, pages 296–313, Glasgow, Scotland, September 1995. Springer-Verlag.

- [51] Flemming Nielson. A denotational framework for data flow analysis. *Acta Informatica*, 18:265–287, 1982.
- [52] Flemming Nielson and Hanne Riis Nielson. *Two-Level Functional Languages*, volume 34 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1992.
- [53] Hanne Riis Nielson and Flemming Nielson. *Semantics with Applications, a formal introduction*. Wiley Professional Computing. John Wiley and Sons, 1992.
- [54] Jens Palsberg. Correctness of binding-time analysis. In Jones [41], pages 347–363.
- [55] Christine Paulin-Mohring and Benjamin Werner. Synthesis of ML programs in the system Coq. *Journal of Symbolic Computation*, 15:607–640, 1993.
- [56] Gordon D. Plotkin. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science*, 1:125–159, 1975.
- [57] Dag Prawitz. *Natural Deduction*. Almqvist and Wiksell, Uppsala, 1965.
- [58] Amr Sabry and Matthias Felleisen. Reasoning about programs in continuation-passing style. *LISP and Symbolic Computation*, 6(3/4): 289–360, December 1993.
- [59] Amr Sabry and Matthias Felleisen. Is continuation-passing useful for data flow analysis? In Vivek Sarkar, editor, *Proceedings of the ACM SIGPLAN'94 Conference on Programming Languages Design and Implementation*, SIGPLAN Notices, Vol. 29, No 6, pages 1–12, Orlando, Florida, June 1994. ACM Press.
- [60] Amr Sabry and Philip Wadler. Compiling with reflections. In Dybvig [23], pages 13–24.
- [61] David Sands. Proving the correctness of recursion-based automatic program transformations. In Peter Mosses, Mogens Nielsen, and Michael Schwartzbach, editors, *Proceedings of TAPSOFT '95*, number 915 in *Lecture Notes in Computer Science*, pages 681–695, Aarhus, Denmark, May 1995.

- [62] Morten Heine Sørensen, Robert Glück, and Neil Jones. Towards unifying partial evaluation, deforestation, supercompilation, and GPC. In Donald Sannella, editor, *Proceedings of the Fifth European Symposium on Programming*, number 788 in Lecture Notes in Computer Science, pages 485–500, Edinburgh, Scotland, April 1994.
- [63] Guy L. Steele Jr. Building interpreters by composing monads. In Boehm [5], pages 472–492.
- [64] Carolyn Talcott. A theory for program and data type specification. *Theoretical Computer Science*, 104(1):129–159, 1992.
- [65] Peter Thiemann. Cogen in six lines. In Dybvig [23], pages 180–189.
- [66] Philip Wadler. The essence of functional programming (tutorial). 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.
- [67] Philip Wadler. Comprehending monads. *Mathematical Structures in Computer Science*, 2(4):461–493, December 1992.
- [68] Mitchell Wand. Specifying the correctness of binding-time analysis. In Jones [41], pages 365–387.

Recent Publications in the BRICS Report Series

- RS-96-34 John Hatcliff and Olivier Danvy. *A Computational Formalization for Partial Evaluation (Extended Version)*. October 1996. 67 pp. To appear in *Mathematical Structures in Computer Science*.
- RS-96-33 Jonathan F. Buss, Gudmund Skovbjerg Frandsen, and Jeffrey Outlaw Shallit. *The Computational Complexity of Some Problems of Linear Algebra*. September 1996. 39 pp. Revised version to appear in *STACS '97: 14th Annual Symposium on Theoretical Aspects of Computer Science Proceedings*, LNCS, 1997.
- RS-96-32 P. S. Thiagarajan. *Regular Trace Event Structures*. September 1996. 34 pp.
- RS-96-31 Ian Stark. *Names, Equations, Relations: Practical Ways to Reason about 'new'*. September 1996. ii+22 pp. To appear in *Typed Lambda Calculi and Applications: 3rd International Conference, TLCA '97 Proceedings*, LNCS, 1997.
- RS-96-30 Arne Andersson, Peter Bro Miltersen, and Mikkel Thorup. *Fusion Trees can be Implemented with AC^0 Instructions only*. September 1996. 8 pp.
- RS-96-29 Lars Arge. *The I/O-Complexity of Ordered Binary-Decision Diagram Manipulation*. August 1996. 35 pp. An extended abstract version appears in Staples, Eades, Kato, and Moffat, editors, *Algorithms and Computation: 6th International Symposium, ISAAC '95 Proceedings*, LNCS 1004, 1995, pages 82–91.
- RS-96-28 Lars Arge. *The Buffer Tree: A New Technique for Optimal I/O Algorithms*. August 1996. 34 pp. This report is a revised and extended version of the BRICS Report RS-94-16. An extended abstract appears in Akl, Dehne, Sack, and Santoro, editors, *Algorithms and Data Structures: 4th Workshop, WADS '95 Proceedings*, LNCS 955, 1995, pages 334–345.
- RS-96-27 Devdatt P. Dubhashi, Volker Priebe, and Desh Ranjan. *Negative Dependence Through the FKG Inequality*. July 1996. 15 pp.