Math. Struct. in Comp. Science (1997), *vol.* 7, *pp.* 507–541. Printed in the United Kingdom
© 1997 Cambridge University Press

A computational formalization for partial evaluation

JOHN HATCLIFF[†] and OLIVIER DANVY[‡]

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

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

Received 4 December 1995; revised August 1996

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. It also 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.

1. Introduction

Partial evaluation is a program-transformation technique for specializing programs. It is based on propagating constant values and folding constant expressions (Consel and Danvy 1993; Jones *et al.* 1993). 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 two-phase organization thus makes it clear that the more parts of a source program are static, the more the source program can be specialized. Over the last five years, there have been investigations into various pre-transformations that 'improve binding times' (making binding-time analysis classify more parts as 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 metalanguage.

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

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

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

computations, which are 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 show 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 improvement 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, which 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. We will 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 is in general 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 (Bondorf and Danvy 1991), 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 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

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

above (Bondorf and Danvy 1991), 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

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

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:

$$\begin{array}{l}\lambda z\,.\,\mathsf{let}\ a \Leftarrow z \times z\\ \quad \mathsf{in}\ (\lambda(x, y)\,.\,\mathsf{let}\ v_1 \Leftarrow x + 1\\ \qquad\qquad \mathsf{in}\ \mathsf{let}\ v_2 \Leftarrow y + y\\ \qquad\qquad\quad \mathsf{in}\ v_1 - v_2)\ @\ (10, a).\end{array}$$

Specializing this term 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

$$\begin{array}{l}\lambda z\,.\,\mathsf{let}\ a \Leftarrow z \times z\\ \quad \mathsf{in}\ \mathsf{let}\ v_1 \Leftarrow 10 + 1\\ \qquad \mathsf{in}\ \mathsf{let}\ v_2 \Leftarrow a + a\\ \qquad\quad \mathsf{in}\ v_1 - v_2\,.\end{array}$$

The static addition is performed, which leads to

$$\begin{array}{l}\lambda z\,.\,\mathsf{let}\ a \Leftarrow z \times z\\ \quad \mathsf{in}\ \mathsf{let}\ v_1 \Leftarrow \mathsf{unit}\ 11\\ \qquad \mathsf{in}\ \mathsf{let}\ v_2 \Leftarrow a + a\\ \qquad\quad \mathsf{in}\ v_1 - v_2\,.\end{array}$$

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 (Jones *et al.* 1993, 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 improvements in the static data flow and the static control flow of

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

a source program. The data-flow aspects have been recently clarified (Danvy *et al.* 1995; Danvy *et al.* 1996), but the control-flow aspects still remain a research topic (Bondorf 1992; Bondorf and Dussart 1994; Consel and Danvy 1991; Holst and Gomard 1991; Lawall and Danvy 1994).

The program calculus for Moggi's computational metalanguage includes *monadic laws*. We will now 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

$$\lambda f\,.\,((\lambda x\,.\,2)\mathbin{@}(f\mathbin{@}1))+3$$

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

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

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

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

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' (Bondorf 1992; Jones *et al.* 1993; Lawall and Danvy 1994). Essentially they mimic one-pass CPS transformations (Danvy and Filinski 1992).

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

$$\begin{array}{l}\lambda f\,.\,\mathsf{let}\, v_1 \Leftarrow \mathsf{let}\, v_2 \Leftarrow f\mathbin{@}1\\ \qquad\qquad\quad \mathsf{in}\,(\lambda x\,.\,\mathsf{unit}\,2)\mathbin{@}v_2\\ \quad \mathsf{in}\, v_1+3\,.\end{array}$$

Static reduction of the β-redex yields

$$\begin{array}{l}\lambda f\,.\,\mathsf{let}\, v_1 \Leftarrow \mathsf{let}\, v_2 \Leftarrow f\mathbin{@}1\\ \qquad\qquad\quad \mathsf{in\ unit}\,2\\ \quad \mathsf{in}\, v_1+3\,.\end{array}$$

Reassociating the let expression yields

$$\begin{array}{l}\lambda f\,.\,\mathsf{let}\, v_2 \Leftarrow f\mathbin{@}1\\ \quad \mathsf{in\ let}\, v_1 \Leftarrow \mathsf{unit}\,2\\ \qquad \mathsf{in}\, v_1+3\,.\end{array}$$

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Unfolding the inner let and statically reducing the addition yields

$$\lambda f\,.\,\mathsf{let}\,v_2 \Leftarrow f\ @\ 1\ \mathsf{in}\ \mathsf{unit}\,5\,.$$

The result essentially coincides with the residual program above.

1.3. *Computational effects*

Since Similix (Bondorf and Danvy 1991), 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, splitting the store into 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 (Dussart and Thiemann 1996).

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, which applies the monadic laws. After a review of related work in Section 7, Section 8 gives conclusions.

2. The computational meta-language

2.1. *Syntax*

Figure 1 presents the language Λ_{ml}, which is based on Moggi's *computational meta-language* (Moggi 1991)†. The typing system of Λ_{ml} captures the distinction between *values* (terms with no remaining computation steps) and *computations* (terms with remaining computation steps). Types of the form nat and $\tau_1 \to \tau_2$ are *value types*. Accordingly, the rules for numerals and abstractions belong to the introduction rules for value types. Types of the form $\widetilde{\tau}$ are *computation types*. All functions have computational co-domains.

† We have added fix, succ, pred, and if0 to Moggi's published description (Moggi 1991).

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

$$
\begin{array}{rcl}
e & \in & \mathit{Terms}[\Lambda_{ml}] \\
e & ::= & x \mid \ulcorner n \urcorner \mid \mathsf{succ}\, e \mid \mathsf{pred}\, e \mid \mathsf{if0}\, e_1\, e_2\, e_3 \mid \lambda x.\, e \mid e_0 \mathbin{@} e_1 \mid \mathsf{fix}\, x.\, e \mid \\
 & & \mathsf{unit}\, e \mid \mathsf{let}\, x \Leftarrow e_1\, \mathsf{in}\, e_2
\end{array}
$$

$$
\begin{array}{rcl}
\tau & \in & \mathit{Types}[\Lambda_{ml}] \\
\tau & ::= & \mathsf{nat} \mid \tau_1 \to \widetilde{\tau_2} \mid \widetilde{\tau}
\end{array}
\qquad\qquad
\begin{array}{rcl}
\Gamma & \in & \mathit{Assumptions}[\Lambda_{ml}] \\
\Gamma & ::= & \cdot \mid \Gamma, x : \tau
\end{array}
$$

$$
\Gamma \vdash_{ml} \ulcorner n \urcorner : \mathsf{nat}
\qquad
\frac{\Gamma \vdash_{ml} e : \mathsf{nat}}{\Gamma \vdash_{ml} \mathsf{succ}\, e : \widetilde{\mathsf{nat}}}
\qquad
\frac{\Gamma \vdash_{ml} e : \mathsf{nat}}{\Gamma \vdash_{ml} \mathsf{pred}\, e : \widetilde{\mathsf{nat}}}
$$

$$
\Gamma \vdash_{ml} x : \Gamma(x)
\qquad
\frac{\Gamma \vdash_{ml} e_1 : \mathsf{nat} \quad \Gamma \vdash_{ml} e_2 : \widetilde{\tau} \quad \Gamma \vdash_{ml} e_3 : \widetilde{\tau}}{\Gamma \vdash_{ml} \mathsf{if0}\, e_1\, e_2\, e_3 : \widetilde{\tau}}
$$

$$
\frac{\Gamma, x : \tau_1 \vdash_{ml} e : \widetilde{\tau_2}}{\Gamma \vdash_{ml} \lambda x.\, e : \tau_1 \to \widetilde{\tau_2}}
\qquad
\frac{\Gamma \vdash_{ml} e_0 : \tau_1 \to \widetilde{\tau_2} \quad \Gamma \vdash_{ml} e_1 : \tau_1}{\Gamma \vdash_{ml} e_0 \mathbin{@} e_1 : \widetilde{\tau_2}}
\qquad
\frac{\Gamma, x : \widetilde{\tau} \vdash_{ml} e : \widetilde{\tau}}{\Gamma \vdash_{ml} \mathsf{fix}\, x.\, e : \widetilde{\tau}}
$$

monadic constructs:

$$
\frac{\Gamma \vdash_{ml} e : \tau}{\Gamma \vdash_{ml} \mathsf{unit}\, e : \widetilde{\tau}}
\qquad
\frac{\Gamma \vdash_{ml} e_1 : \widetilde{\tau_1} \quad \Gamma, x : \tau_1 \vdash_{ml} e_2 : \widetilde{\tau_2}}{\Gamma \vdash_{ml} \mathsf{let}\, x \Leftarrow e_1\, \mathsf{in}\, e_2 : \widetilde{\tau_2}}
$$

Fig. 1. The computational meta-language Λ_{ml}

$$
\begin{array}{rcl}
\mathsf{succ}\, \ulcorner n \urcorner & \longmapsto_i & \mathsf{unit}\, \ulcorner n+1 \urcorner \\
(\lambda x.\, e_0) \mathbin{@} e_1 & \longmapsto_i & e_0[x := e_1] \\
\mathsf{fix}\, x.\, e & \longmapsto_i & e[x := \mathsf{fix}\, x.\, e] \\
\mathsf{let}\, x \Leftarrow \mathsf{unit}\, e_1\, \mathsf{in}\, e_2 & \longmapsto_i & e_2[x := e_1]
\end{array}
\qquad
\begin{array}{rcl}
\mathsf{pred}\, \ulcorner n+1 \urcorner & \longmapsto_i & \mathsf{unit}\, \ulcorner n \urcorner \\
\mathsf{pred}\, \ulcorner 0 \urcorner & \longmapsto_i & \mathsf{unit}\, \ulcorner 0 \urcorner \\
\mathsf{if0}\, \ulcorner 0 \urcorner\, e_2\, e_3 & \longmapsto_i & e_2 \\
\mathsf{if0}\, \ulcorner n+1 \urcorner\, e_2\, e_3 & \longmapsto_i & e_3
\end{array}
$$

$$
\frac{e \longmapsto_i e'}{E_i[e] \longmapsto_i E_i[e']}
\qquad\qquad
E_i \quad ::= \quad \mathsf{let}\, x \Leftarrow [\cdot]\, \mathsf{in}\, e_2
$$

Fig. 2. Interpretation steps for Λ_{ml}

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Thus, all applications have a computation type, which captures 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†. Intuitively, $\mathsf{unit}\, e$ is a trivial computation that simply yields the value of e. $\mathsf{let}\, x \Leftarrow e_1\, \mathsf{in}\, e_2$ 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.* (Barendregt 1984). We write $e_1 \equiv e_2$ when e_1 and e_2 are α-equivalent.

We represent simultaneous substitutions $e[x_1 := e_1, \ldots, x_n := e_n]$ using a substitution function $\sigma : \mathit{Terms}[\Lambda_{ml}] \to \mathit{Terms}[\Lambda_{ml}]$ whose application is denoted $e\sigma$. A substitution σ is *closed* if, for all $x \in \mathrm{dom}\ \sigma$, $x\sigma$ 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 $x \in \mathit{dom}\ \Gamma$, $\cdot \vdash_{ml} x\sigma : \Gamma(x)$. *Programs* are closed terms with type $\widetilde{\mathsf{nat}}$.

2.2. *Operational semantics*

Figure 2 presents single-step evaluation rules for Λ_{ml}‡. Axioms such as $\mathsf{succ}\ulcorner n\urcorner \longmapsto_i \mathsf{unit}\ulcorner n+1\urcorner$ 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 the corresponding type, or (b) can undergo an evaluation step that preserves typing. In particular, there are no 'stuck' (Plotkin 1975, p. 151) terms§.

Lemma 1. (interpretation properties)

— If $\cdot \vdash_{ml} e : \mathsf{nat}$, then $e \equiv \ulcorner n\urcorner$ for some number n.
— If $\cdot \vdash_{ml} e : \tau_1 \to \widetilde{\tau_2}$, then $e \equiv \lambda x.\, e'$ and $x : \tau_1 \vdash_{ml} e' : \widetilde{\tau_2}$.
— If $\cdot \vdash_{ml} e : \widetilde{\tau}$, then exactly one of the following statements holds:

1 $e = \mathsf{unit}\, e'$ and $\cdot \vdash_{ml} e' : \tau$.

2 $e \longmapsto_i e'$ and $\cdot \vdash_{ml} e' : \widetilde{\tau}$.

Proof. The proof is 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 $\Gamma \vdash_{ml} e_0[x := e_1] : \tau_0$. □

† The exact connection to the structure of a monad can be found in Moggi's original work (Moggi 1991, page 61).

‡ Moggi originally gave a categorical interpretation for his computational meta-language (Moggi 1991). Crole and Pitts extended Moggi's work with a fix-point operator and an associated logic (Crole and Pitts 1992). Later, Gordon developed an elegant operational theory for the meta-language with a fixpoint operator and inductive and coinductive types (Gordon 1993). Λ_{ml} is basically a sub-language of Gordon's, except that we include directly type nat and associated operations, whereas Gordon constructs them via inductive types. Here, we follow Gordon and our previous work (Hatcliff 1994; Hatcliff and Danvy 1994) and present a structural operational semantics for Λ_{ml}.

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

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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

Definition 1. (Interpreter) For $\cdot \vdash_{ml} e : \widetilde{\mathsf{nat}}$,

$$int\, e = \ulcorner n \urcorner \quad iff \quad e \longmapsto_i^* \mathsf{unit} \ulcorner n \urcorner.$$

We write $int\, e\!\downarrow$ when $int\, e$ is defined and $int\, e\!\uparrow$ when $int\, e$ is undefined. As a consequence of Lemma 1, $int\, e\!\uparrow$ 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[†].

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,

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

Definition 3. (Operational equivalence) For $\Gamma \vdash_{ml} e_1, e_2 : \tau$,

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

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

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, $\twoheadrightarrow$ denotes the reflexive, transitive closure of $\longrightarrow$, and $=_{ml}$ denotes the smallest equivalence relation generated by $\longrightarrow$ (Barendregt 1984).

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 (Gordon 1993, Chapters 3,4,5). □

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 (Hatcliff 1994; Hatcliff and Danvy 1994), we proved the correctness of an encoding $\mathscr{E}_n$ of a call-by-name version of Λ into Λ_{ml}, and an encoding $\mathscr{E}_v$ of a call-by-value version of Λ into Λ_{ml}. These encodings can be found in Appendix A.2[‡].

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

[‡] These encodings are based on similar encodings given by *e.g.*, Moggi (Moggi 1991) and Wadler (Wadler 1992b).

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

$$\begin{array}{rcl@{\qquad}rcl}
\mathsf{succ}\ulcorner n\urcorner & \longrightarrow & \mathsf{unit}\ulcorner n+1\urcorner & \mathsf{pred}\ulcorner n+1\urcorner & \longrightarrow & \mathsf{unit}\ulcorner n\urcorner \\
(\lambda x.e_0)\mathbin{@} e_1 & \longrightarrow & e_0[x := e_1] & \mathsf{pred}\ulcorner 0\urcorner & \longrightarrow & \mathsf{unit}\ulcorner 0\urcorner \\
\mathsf{fix}\,x.e & \longrightarrow & e[x := \mathsf{fix}\,x.e] & \mathsf{if0}\ulcorner 0\urcorner\, e_2\, e_3 & \longrightarrow & e_2 \\
 & & & \mathsf{if0}\ulcorner n+1\urcorner\, e_2\, e_3 & \longrightarrow & e_3
\end{array}$$

Monadic reductions:

$$\begin{array}{lrcll}
(let.\beta) & \mathsf{let}\ x \Leftarrow \mathsf{unit}\ e_1\ \mathsf{in}\ e_2 & \longrightarrow & e_2[x := e_1] & \\
(let.\eta) & \mathsf{let}\ x \Leftarrow e\ \mathsf{in}\ \mathsf{unit}\ x & \longrightarrow & e & \\
(let.assoc) & \mathsf{let}\ x_2 \Leftarrow \mathsf{let}\ x_1 \Leftarrow e_1 & \longrightarrow & \mathsf{let}\ x_1 \Leftarrow e_1 & x_1 \notin FV(e_3) \\
 & \quad \mathsf{in}\ e_2 & & \mathsf{in\ let}\ x_2 \Leftarrow e_2 & \\
 & \mathsf{in}\ e_3 & & \quad \mathsf{in}\ e_3 &
\end{array}$$

Fig. 3. Notions of reduction *ml* for Λ_{ml}

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†. Applications are translated as follows:

$$\begin{array}{lrcl}
\textit{Call-by-name:} & \mathscr{E}_n\langle\!\langle e_0 \mathbin{@} e_1\rangle\!\rangle & = & \mathsf{let}\ x_0 \Leftarrow \mathscr{E}_n\langle\!\langle e_0\rangle\!\rangle\ \mathsf{in}\ x_0 \mathbin{@} \mathscr{E}_n\langle\!\langle e_1\rangle\!\rangle \\
\textit{Call-by-value:} & \mathscr{E}_v\langle\!\langle e_0 \mathbin{@} e_1\rangle\!\rangle & = & \mathsf{let}\ x_0 \Leftarrow \mathscr{E}_v\langle\!\langle e_0\rangle\!\rangle\ \mathsf{in\ let}\ x_1 \Leftarrow \mathscr{E}_v\langle\!\langle e_1\rangle\!\rangle\ \mathsf{in}\ x_0 \mathbin{@} x_1\,.
\end{array}$$

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 next 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 that 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.

† Alternative 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 (Hatcliff and Danvy 1994, Section 4).

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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 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)†.

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^{bt}_{ml} 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^{bt}_{ml} to residualize the result of evaluating an eliminable term. This allows static computation to occur in a residual context‡. A term $w \in \Lambda^{bt}_{ml}$ is *completely residual* if it consists of only underlined constructs and identifiers. *Residual-terms*$[\Lambda^{bt}_{ml}]$ denotes the set of completely residual terms. Intuitively, the specializer will output completely

† Following other formal treatments of partial evaluation (Jones *et al.* 1993; Palsberg 1993; Wand 1993), we simplify the presentation by omitting folding and generalization strategies.

‡ We 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 (Danvy *et al.* 1995; Danvy *et al.* 1996). This definition is also interesting on its own (Danvy 1996).

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

$$
\begin{array}{lcl}
w & \in & \mathit{Terms}[\Lambda_{ml}^{bt}] \\
w & ::= & x \mid \\
 & & \ulcorner n \urcorner \mid \mathsf{succ}\, w \mid \mathsf{pred}\, w \mid \mathsf{if0}\, w_1\, w_2\, w_3 \mid \lambda x.\, w \mid w_0 \mathbin{@} w_1 \mid \mathsf{fix}\, x.\, w \mid \\
 & & \mathsf{unit}\, w \mid \mathsf{let}\, x \Leftarrow w_1\, \mathsf{in}\, w_2 \mid \\
 & & \underline{n} \mid \underline{\mathsf{succ}}\, w \mid \underline{\mathsf{pred}}\, w \mid \underline{\mathsf{if0}}\, w_1\, w_2\, w_3 \mid \underline{\lambda} x.\, w \mid w_0 \mathbin{\underline{@}} w_1 \mid \underline{\mathsf{fix}}\, x.\, w \mid \\
 & & \underline{\mathsf{unit}}\, w \mid \underline{\mathsf{let}}\, x \Leftarrow w_1\, \underline{\mathsf{in}}\, w_2 \mid \mathsf{lift}\, w
\end{array}
$$

Fig. 4. The binding-time annotated meta-language Λ_{ml}^{bt}

$$
\begin{array}{lcl@{\qquad}lcl@{\qquad}lcl}
\varphi_{\mathsf{nat}} & \in & \mathit{STypes}[\mathsf{nat}] & \varphi_{\tau_1 \to \widetilde{\tau}_2} & \in & \mathit{STypes}[\tau_1 \to \widetilde{\tau}_2] & \varphi_{\widetilde{\tau}} & \in & \mathit{STypes}[\widetilde{\tau}] \\
\varphi_{\mathsf{nat}} & ::= & \mathsf{s} \mid \mathsf{d} & \varphi_{\tau_1 \to \widetilde{\tau}_2} & ::= & \varphi_{\tau_1} \to \varphi_{\widetilde{\tau}_2} \mid \mathsf{d} & \varphi_{\widetilde{\tau}} & ::= & \widetilde{\varphi}_{\tau} \mid \mathsf{d}
\end{array}
$$

$$
\begin{array}{lcl@{\qquad}lcl}
\tau & \in & \mathit{Types}[\Lambda_{ml}^{bt}] & \Gamma & \in & \mathit{Assumptions}[\Lambda_{ml}^{bt}] \\
\tau & ::= & \mathsf{nat} \mid \tau_1 \to \widetilde{\tau}_2 \mid \widetilde{\tau} & \Gamma & ::= & \cdot \mid \Gamma, x : \tau[\varphi_\tau]
\end{array}
$$

Fig. 5. Specialization types and binding-time assumptions for Λ_{ml}^{bt}

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 $\varphi = \mathsf{d}$; otherwise φ is *static*.

— $\mathsf{s} \in \mathit{STypes}[\mathsf{nat}]$ will tag expressions of type nat that are guaranteed to evaluate to known data (*i.e.*, numerals).

— $\mathsf{d} \in \mathit{STypes}[\mathsf{nat}]$ will tag expressions of type nat whose evaluation may depend on unknown data and thus cannot be guaranteed to evaluate 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.

— $\varphi_{\tau_1} \to \varphi_{\widetilde{\tau}_2} \in \mathit{STypes}[\tau_1 \to \widetilde{\tau}_2]$ will tag expressions of type $\tau_1 \to \widetilde{\tau}_2$ that are guaranteed to evaluate to known data (*i.e.*, abstractions).

— $\mathsf{d} \in \mathit{STypes}[\tau_1 \to \widetilde{\tau}_2]$ will tag expressions of type $\tau_1 \to \widetilde{\tau}_2$ whose evaluation may depend on unknown data and thus cannot be guaranteed to evaluate to an abstraction. However, because $\tau_1 \to \widetilde{\tau}_2$ is a value type, one does know that the tagged expression will denote a value when dynamic data is supplied.

— $\widetilde{\varphi} \in \mathit{STypes}[\widetilde{\tau}]$ will tag expressions of type $\widetilde{\tau}$ that are guaranteed to either diverge or to evaluate to a trivial computation (*i.e.*, $\mathsf{unit}\, e$ for some $e \in \Lambda_{ml}$ of type τ).

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

— $\mathsf{d} \in \mathit{STypes}[\widetilde{\tau}]$ will tag expressions of type $\widetilde{\tau}$ 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^{bt}_{ml} variables. If $\Gamma, x:\tau[\mathsf{d}]$, then x is a *dynamic variable*; if $\Gamma, x:\tau[\varphi]$ where φ is static, then x is a *static variable*. $\Gamma(x).\mathit{type}$ and $\Gamma(x).\mathit{spec\text{-}type}$ project types and specialization types from an assumption for $x \in \mathit{dom}\ \Gamma$.

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 $e \in \mathit{Terms}[\Lambda_{ml}]$ of type $\tau \in \mathit{Types}[\Lambda_{ml}]$ to a directive annotated term $w \in \mathit{Terms}[\Lambda^{bt}_{ml}]$ of specialization type $\varphi_\tau \in \mathit{STypes}[\tau]$. 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.

$\lfloor\cdot\rfloor : \mathit{Terms}[\Lambda^{bt}_{ml}] \to \mathit{Terms}[\Lambda_{ml}]$ is an erasing function that removes annotations and lift constructs. $\lfloor\cdot\rfloor : \mathit{Assumptions}[\Lambda^{bt}_{ml}] \to \mathit{Assumptions}[\Lambda_{ml}]$ simply drops specialization types appearing in assumptions. $\lfloor\sigma\rfloor$ represents the component-wise application of $\lfloor\cdot\rfloor : \mathit{Terms}[\Lambda^{bt}_{ml}] \to \mathit{Terms}[\Lambda_{ml}]$. A closed Λ^{bt}_{ml} substitution σ is *compatible* with an assumption $\Gamma \in \mathit{Assumptions}[\Lambda^{bt}_{ml}]$, if for all $x \in \mathit{dom}\ \Gamma$, $\cdot \vdash_{ml} \lfloor x\sigma \rfloor : \Gamma(x).\mathit{type}[x\sigma : \Gamma(x).\mathit{spec\text{-}type}]$. For expressions $e \in \mathit{Terms}[\Lambda_{ml}]$, Γ_{bt} is a *binding-time assumption for* $\Gamma \vdash_{ml} e : \tau$, if $\lfloor\Gamma_{bt}\rfloor = \Gamma$. $\Gamma_{bt} = \Gamma_s \cup \Gamma_d$ 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

$$\mathit{bta} : \mathit{Assumptions}[\Lambda^{bt}_{ml}] \to \mathit{Terms}[\Lambda_{ml}] \to \mathit{Terms}[\Lambda^{bt}_{ml}]$$

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

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

The binding-time analysis rules specify a one-to-many relation between Λ_{ml} terms and Λ^{bt}_{ml} terms. Thus, there are many valid binding-time analysis functions, each varying in pratical 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 (Jones *et al.* 1993). Our specification abstracts away from these orthogonal implementation issues.

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

3.2. *Specialization*

To specialize a binding-time analyzed program term w such that $\Gamma_{bt} = \Gamma_s \cup \Gamma_d \vdash_{bt} e : \widetilde{\mathsf{nat}}[w : \mathsf{d}]$, one supplies known data into the static variables via a substitution σ_s

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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

$$\Gamma \vdash_{bt} \ulcorner n \urcorner : \mathsf{nat}[\ulcorner n \urcorner : \mathsf{s}] \qquad \Gamma \vdash_{bt} \ulcorner n \urcorner : \mathsf{nat}[\underline{\ulcorner n \urcorner} : \mathsf{d}]$$

$$\frac{\Gamma \vdash_{bt} e : \mathsf{nat}[w : \mathsf{s}]}{\Gamma \vdash_{bt} \mathsf{succ}\, e : \widetilde{\mathsf{nat}} \quad [\mathsf{succ}\, w : \widetilde{\mathsf{s}}]} \qquad \frac{\Gamma \vdash_{bt} e : \mathsf{nat}[w : \mathsf{d}]}{\Gamma \vdash_{bt} \mathsf{succ}\, e : \widetilde{\mathsf{nat}} \quad [\underline{\mathsf{succ}}\, w : \mathsf{d}]}$$

$$\frac{\Gamma \vdash_{bt} e : \mathsf{nat}[w : \mathsf{s}]}{\Gamma \vdash_{bt} \mathsf{pred}\, e : \widetilde{\mathsf{nat}} \quad [\mathsf{pred}\, w : \widetilde{\mathsf{s}}]} \qquad \frac{\Gamma \vdash_{bt} e : \mathsf{nat}[w : \mathsf{d}]}{\Gamma \vdash_{bt} \mathsf{pred}\, e : \widetilde{\mathsf{nat}} \quad [\underline{\mathsf{pred}}\, w : \mathsf{d}]}$$

$$\frac{\Gamma \vdash_{bt} e_1 : \mathsf{nat}[w_1 : \mathsf{s}] \qquad \Gamma \vdash_{bt} e_2 : \widetilde{\tau}[w_2 : \varphi] \qquad \Gamma \vdash_{bt} e_3 : \widetilde{\tau}[w_3 : \varphi]}{\Gamma \vdash_{bt} \mathsf{if0}\, e_1\, e_2\, e_3 : \widetilde{\tau}[\mathsf{if0}\, w_1\, w_2\, w_3 : \varphi]}$$

$$\frac{\Gamma \vdash_{bt} e_1 : \mathsf{nat}[w_1 : \mathsf{d}] \qquad \Gamma \vdash_{bt} e_2 : \widetilde{\tau}[w_2 : \mathsf{d}] \qquad \Gamma \vdash_{bt} e_3 : \widetilde{\tau}[w_3 : \mathsf{d}]}{\Gamma \vdash_{bt} \mathsf{if0}\, e_1\, e_2\, e_3 : \widetilde{\tau}[\underline{\mathsf{if0}}\, w_1\, w_2\, w_3 : \mathsf{d}]}$$

$$\frac{\Gamma, x : \tau_1[\varphi_1] \vdash_{bt} e : \widetilde{\tau_2}[w : \varphi_2]}{\Gamma \vdash_{bt} \lambda x\,.\, e : \tau_1 \to \widetilde{\tau_2}[\lambda x\,.\, w : \varphi_1 \to \varphi_2]}$$

$$\frac{\Gamma, x : \tau_1[\mathsf{d}] \vdash_{bt} e : \widetilde{\tau_2}[w : \mathsf{d}]}{\Gamma \vdash_{bt} \lambda x\,.\, e : \tau_1 \to \widetilde{\tau_2}[\underline{\lambda} x\,.\, w : \mathsf{d}]}$$

$$\frac{\Gamma \vdash_{bt} e_0 : \tau_1 \to \widetilde{\tau_2}[w_0 : \varphi_1 \to \varphi_2] \qquad \Gamma \vdash_{bt} e_1 : \tau_1[w_1 : \varphi_1]}{\Gamma \vdash_{bt} e_0 \mathbin{@} e_1 : \widetilde{\tau_2}[w_0 \mathbin{@} w_1 : \varphi_2]}$$

$$\frac{\Gamma \vdash_{bt} e_0 : \tau_1 \to \widetilde{\tau_2}[w_0 : \mathsf{d}] \qquad \Gamma \vdash_{bt} e_1 : \tau_1[w_1 : \mathsf{d}]}{\Gamma \vdash_{bt} e_0 \mathbin{@} e_1 : \widetilde{\tau_2}[w_0 \mathbin{\underline{@}} w_1 : \mathsf{d}]}$$

$$\frac{\Gamma, x : \widetilde{\tau}[\widetilde{\varphi}] \vdash_{bt} e : \widetilde{\tau}[w : \widetilde{\varphi}]}{\Gamma \vdash_{bt} \mathsf{fix}\, x.\, e : \widetilde{\tau}[\mathsf{fix}\, x.\, w : \widetilde{\varphi}]} \qquad \frac{\Gamma, x : \widetilde{\tau}[\mathsf{d}] \vdash_{bt} e : \widetilde{\tau}[w : \mathsf{d}]}{\Gamma \vdash_{bt} \mathsf{fix}\, x.\, e : \widetilde{\tau}[\underline{\mathsf{fix}}\, x.\, w : \mathsf{d}]}$$

Fig. 6. Binding-time constraints (part 1)

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

$$\frac{\Gamma \vdash_{bt} e : \tau[w : \varphi]}{\Gamma \vdash_{bt} \mathsf{unit}\, e : \widetilde{\tau}[\mathsf{unit}\, w : \widetilde{\varphi}]} \qquad \frac{\Gamma \vdash_{bt} e : \tau[w : \mathsf{d}]}{\Gamma \vdash_{bt} \mathsf{unit}\, e : \widetilde{\tau}[\underline{\mathsf{unit}}\, w : \mathsf{d}]}$$

$$\frac{\Gamma \vdash_{bt} e_1 : \widetilde{\tau_1}[w_1 : \widetilde{\varphi_1}] \quad \Gamma, x : \tau_1[\varphi_1] \vdash_{bt} e_2 : \widetilde{\tau_2}[w_2 : \varphi_2]}{\Gamma \vdash_{bt} \mathsf{let}\, x \Leftarrow e_1 \,\mathsf{in}\, e_2 : \widetilde{\tau_2}[\mathsf{let}\, x \Leftarrow w_1 \,\mathsf{in}\, w_2 : \varphi_2]}$$

$$\frac{\Gamma \vdash_{bt} e_1 : \widetilde{\tau_1}[w_1 : \mathsf{d}] \quad \Gamma, x : \tau_1[\mathsf{d}] \vdash_{bt} e_2 : \widetilde{\tau_2}[w_2 : \mathsf{d}]}{\Gamma \vdash_{bt} \mathsf{let}\, x \Leftarrow e_1 \,\mathsf{in}\, e_2 : \widetilde{\tau_2}[\underline{\mathsf{let}}\, x \Leftarrow w_1 \,\underline{\mathsf{in}}\, w_2 : \mathsf{d}]}\,(*)$$

$$\frac{\Gamma \vdash_{bt} e : \mathsf{nat}[w : \mathsf{s}]}{\Gamma \vdash_{bt} e : \mathsf{nat}[\mathsf{lift}\, w : \mathsf{d}]}$$

Fig. 7. Binding-time constraints (part 2)

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 $\longmapsto_s$).

— 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^{bt}_{ml} terms satisfying judgements of the form $\Gamma_d \vdash_{bt} e : \tau[w : \varphi_\tau]$. The intuition is that 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 that preserves annotations. Moreover, each specialization step $w \longmapsto_s w'$ maintains Λ_{ml} convertibility on the corresponding erasures (*i.e.*, $\lfloor w \rfloor =_{ml} \lfloor w' \rfloor$).

Lemma 2. (specialization properties)

— If $\Gamma_d \vdash_{bt} e : \mathsf{nat}[w : \mathsf{s}]$, then $w \equiv \ulcorner n \urcorner \equiv e$ for some number n.
— If $\Gamma_d \vdash_{bt} e : \tau_1 \to \tau_2[w : \varphi_1 \to \varphi_2]$, then $w \equiv \lambda x.\, w'$, $e \equiv \lambda x.\, \lfloor w' \rfloor$ and $\Gamma_d, x : \tau_1[\varphi_1] \vdash_{bt} \lfloor w' \rfloor : \tau_2[w' : \varphi_2]$ for some $w' \in \mathit{Terms}[\Lambda^{bt}_{ml}]$.
— If $\Gamma_d \vdash_{bt} e : \widetilde{\tau}[w : \widetilde{\varphi}]$, then exactly one of the following statements holds:

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Interpretation rules:

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

$$\frac{w \longmapsto_i w'}{E_i[w] \longmapsto_i E_i[w']} \qquad E_i \quad ::= \quad \mathsf{let}\,x \Leftarrow [\cdot]\,\mathsf{in}\,w_2$$

$$\frac{w \longmapsto_i w'}{w \longmapsto_s w'}$$

Compilation rules:

$$\mathsf{lift}\,\ulcorner n\urcorner \longmapsto_s \underline{\ulcorner n\urcorner} \qquad \frac{w \longmapsto_s w'}{E_s[w] \longmapsto_s E_s[w']}$$

$$\begin{array}{rcl}
E_s & ::= & \underline{\mathsf{pred}}\,[\cdot] \mid \underline{\mathsf{succ}}\,[\cdot] \mid \underline{\mathsf{if0}}\,[\cdot]\,w_2\,w_3 \mid \underline{\mathsf{if0}}\,r_1\,[\cdot]\,w_3 \mid \underline{\mathsf{if0}}\,r_1\,r_2\,[\cdot] \mid \\
 & & \underline{\lambda}x\,.\,[\cdot] \mid [\cdot]\,\underline{@}\,w_1 \mid r_0\,\underline{@}\,[\cdot] \mid \underline{\mathsf{fix}}\,x.\,[\cdot] \mid \\
 & & \underline{\mathsf{unit}}\,[\cdot] \mid \underline{\mathsf{let}}\,x \Leftarrow [\cdot]\,\underline{\mathsf{in}}\,w_2 \mid \underline{\mathsf{let}}\,x \Leftarrow r_1\,\underline{\mathsf{in}}\,[\cdot] \quad \textit{...where } r_i \in \textit{Residual-terms}[\Lambda^{bt}_{ml}]
\end{array}$$

Fig. 8. Specialization rules for Λ^{bt}_{ml}

1 $w \equiv \mathsf{unit}\,w'$, $e \equiv \mathsf{unit}\,\lfloor w'\rfloor$ and $\Gamma_d \vdash_{bt} \lfloor w'\rfloor : \tau[w' : \varphi]$ for some $w' \in \textit{Terms}[\Lambda^{bt}_{ml}]$, or

2 $w \longmapsto_i w'$ and $\Gamma_d \vdash_{bt} \lfloor w'\rfloor : \widetilde{\tau}[w' : \widetilde{\varphi}]$ and $e \equiv \lfloor w\rfloor \longrightarrow_{ml} \lfloor w'\rfloor$.

— If $\Gamma_d \vdash_{bt} e : \mathsf{nat}[w : \mathsf{d}]$, then exactly one of the following statements holds:

1 $w \in \textit{Residual-terms}[\Lambda^{bt}_{ml}]$, or

2 $w \longmapsto_s w'$, $\Gamma_d \vdash_{bt} \lfloor w'\rfloor : \tau[w' : \mathsf{d}]$, and $e \equiv \lfloor w'\rfloor \equiv \lfloor w\rfloor \equiv \ulcorner n\urcorner$ for some number n.

— If $\Gamma_d \vdash_{bt} e : \tau_1 \to \tau_2[w : \mathsf{d}]$, then exactly one of the following statements holds:

1 $w \in \textit{Residual-terms}[\Lambda^{bt}_{ml}]$, or

2 $w \longmapsto_s w'$ and $\Gamma_d \vdash_{bt} \lfloor w'\rfloor : \tau[w' : \mathsf{d}]$ and $e \equiv \lfloor w\rfloor \longrightarrow_{ml} \lfloor w'\rfloor$.

— If $\Gamma_d \vdash_{bt} e : \widetilde{\tau}[w : \mathsf{d}]$, then exactly one of the following statements holds:

1 $w \in \textit{Residual-terms}[\Lambda^{bt}_{ml}]$, or

2 $w \longmapsto_s w'$ and $\Gamma_d \vdash_{bt} \lfloor w'\rfloor : \widetilde{\tau}[w' : \mathsf{d}]$ and $e \equiv \lfloor w\rfloor \longrightarrow_{ml} \lfloor w'\rfloor$.

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Proof. by induction over the height of the derivation of $\Gamma_d \vdash_{bt} e : \tau[w : \varphi]$ and relying on the property that 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 $\Gamma \vdash_{bt} e_0[x := e_1] : \tau_0[w_0[x := w_1] : \varphi_0]$. □

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

Definition 5. (Specializer) For $\Gamma_d \vdash_{bt} e : \widetilde{\mathsf{nat}}[w : \mathsf{d}]$,

$$\mathit{spec}\, w = r \quad \mathit{iff} \quad w \longmapsto_s^* r$$

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

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 specializer 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 specializer always finds a redex to contract, or terminates because only residual components are left. In other words, the specializer never 'sticks' on a non-redex.

Theorem 2. (Correctness of binding-time analysis) For all $\Gamma_d \vdash_{bt} e : \widetilde{\mathsf{nat}}[w : \mathsf{d}]$, exactly one of the following statements holds:

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

Proof. The proof follows from the definition of *spec* (Definition 5) and Lemma 2. □

This statement of binding-time correctness is analogous to statements of binding-time correctness for the λ-calculus (Jones *et al.* 1993; Palsberg 1993; Wand 1993).

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

Theorem 3. (Soundness of specialization) For all $\Gamma_d \vdash_{bt} e : \widetilde{\mathsf{nat}}[w : \mathsf{d}]$,

$$\mathit{spec}\, w\downarrow \quad \mathit{implies} \quad e \approx \lfloor \mathit{spec}\, w \rfloor.$$

Proof. Given the definition of *spec* (Definition 5) and Lemma 2, one has $e =_{ml} \lfloor \mathit{spec}\, w \rfloor$ 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.

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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 $\Gamma \vdash_{ml} e : \widetilde{\mathsf{nat}}$, and σ_s be a closed substitution compatible with Γ_s, then

$$pe\ e\ \Gamma_{bt}\ \sigma_s \stackrel{\text{def}}{=} \lfloor spec\ (bta\ \Gamma_{bt}\ e)\sigma_s \rfloor.$$

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 $\Gamma \vdash_{ml} e : \widetilde{\mathsf{nat}}$, let σ_s be a closed substitution compatible with Γ_s, and let σ_d be a closed substitution compatible with $\lfloor \Gamma_d \rfloor$. If $(pe\ e\ \Gamma_{bt}\ \sigma_s)\downarrow$, then

$$e\lfloor \sigma_s \rfloor \sigma_d \approx (pe\ e\ \Gamma_{bt}\ \sigma_s)\sigma_d.$$

Proof. The proof follows from Theorem 3 and the definition of specialization (Definition 5) and partial evaluation (Definition 6). □

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 (Launchbury 1991). 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:

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

Unfolding the inner call is unsound in general because e_1 may diverge. Yet a partial evaluator such as Lambda-Mix would unfold it (Gomard 1992; Gomard and Jones 1991).

Through a systematic insertion of let-expressions, partial evaluators such as Similix or Schism ensure sound call unfolding (Bondorf and Danvy 1991; Consel 1993). They would unfold the term above into the following term

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

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 $\mathsf{pred}\,[\cdot]$ 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 $\longmapsto_x$ steps below represent

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

how the above expression would be treated *e.g.*, in Similix.

$$\begin{array}{lll} & \mathsf{pred}\ (\mathsf{let}\ y = e_1\ \mathsf{in}\ (\lambda x\,.\,43) \mathbin{@} y) & \\ \longmapsto_x & \mathsf{pred}\ (\mathsf{let}\ y = e_1\ \mathsf{in}\ 43) & \textit{...unfold call} \\ \longmapsto_x & \mathsf{let}\ y = e_1\ \mathsf{in}\ \mathsf{pred}\ 43 & \textit{...binding-time improvement} \\ \longmapsto_x & \mathsf{let}\ y = e_1\ \mathsf{in}\ 42 & \textit{...static computation.} \end{array}$$

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):

$$\begin{array}{l} e \equiv \mathsf{let}\ y_1 \Leftarrow \mathscr{E}_v\langle\!\langle e_1 \rangle\!\rangle \\ \quad \mathsf{in}\ \mathsf{let}\ y_2 \Leftarrow (\lambda x\,.\,\mathsf{unit}\ \ulcorner 43 \urcorner) \mathbin{@} y_1 \\ \quad\quad \mathsf{in}\ \mathsf{pred}\ y_2\,. \end{array}$$

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

$$\begin{array}{l} w \equiv \underline{\mathsf{let}}\ y_1 \Leftarrow w_1 \\ \quad \underline{\mathsf{in}}\ \mathsf{let}\ y_2 \Leftarrow (\lambda x\,.\,\mathsf{unit}\ \ulcorner 43 \urcorner) \mathbin{@} y_1 \\ \quad\quad \mathsf{in}\ \mathsf{pred}\ y_2\,. \end{array}$$

Specialization reduces the inner let expression and yields the expected answer:

$$\begin{array}{lll} & w & \\ \longmapsto_s & \underline{\mathsf{let}}\ y_1 \Leftarrow \mathscr{E}_v\langle\!\langle e \rangle\!\rangle\ \underline{\mathsf{in}}\ \mathsf{let}\ y_2 \Leftarrow \mathsf{unit}\ \ulcorner 43 \urcorner\ \mathsf{in}\ \mathsf{pred}\ y_2 & \\ \longmapsto_s & \underline{\mathsf{let}}\ y_1 \Leftarrow \mathscr{E}_v\langle\!\langle e \rangle\!\rangle\ \underline{\mathsf{in}}\ \mathsf{pred}\ 43 & \\ \longmapsto_s & \mathsf{let}\ y_1 \Leftarrow \mathscr{E}_v\langle\!\langle e \rangle\!\rangle\ \mathsf{in}\ \mathsf{unit}\ 42\,. & \end{array}$$

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, for example, a mixed evaluation strategy (Danvy and Hatcliff 1993).

5. Control-based binding-time improvements

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

$$\begin{array}{lll} e_1 & \equiv & \mathsf{let}\ v_1 \Leftarrow (\mathsf{let}\ v_2 \Leftarrow x_2\ \mathsf{in}\ \mathsf{succ}\ x_1)\ \mathsf{in}\ e \qquad \textit{...where } v_2 \notin FV(e) \\ e_2 & \equiv & \mathsf{let}\ v_2 \Leftarrow x_2\ \mathsf{in}\ \mathsf{let}\ v_1 \Leftarrow \mathsf{succ}\ x_1\ \mathsf{in}\ e\,. \end{array}$$

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Assume $\Gamma, x_1 : \mathsf{nat}, x_2 : \widetilde{\mathsf{nat}} \vdash_{ml} e_1, e_2 : \widetilde{\tau}$, and let $\Gamma_{bt} = \Gamma'_{bt}, x_1 : \mathsf{nat[s]}, x_2 : \widetilde{\mathsf{nat}}\mathsf{[d]}$ be binding-time assumptions for e_1 and e_2 (that is, x_1 is a static variable and 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 (that is, $\Gamma \vdash_{bt} e_1 : \widetilde{\tau}[w_1 : \mathsf{d}]$ is derivable):

$$w_1 \equiv \underline{\mathsf{let}}\, v_1 \Leftarrow (\underline{\mathsf{let}}\, v_2 \Leftarrow x_2 \,\underline{\mathsf{in}}\, \underline{\mathsf{succ}}\, \mathsf{lift}\, x_1) \,\underline{\mathsf{in}}\, w.$$

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

$$w_2 \equiv \underline{\mathsf{let}}\, v_2 \Leftarrow x_2 \,\underline{\mathsf{in}}\, \mathsf{let}\, v_1 \Leftarrow \mathsf{succ}\, x_1 \,\mathsf{in}\, w'.$$

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 to yield 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 (Hatcliff 1994)) 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:

$$e_3 \equiv \mathsf{let}\, v_1 \Leftarrow ((\lambda z\,.\,\mathsf{let}\, v_2 \Leftarrow x_2 \,\mathsf{in}\, \mathsf{succ}\, x_1) \mathbin{@} \ulcorner 0 \urcorner) \,\mathsf{in}\, e \qquad \textit{...where } v_2 \notin FV(e).$$

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

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

Specialization (after supplying static data $\ulcorner 42 \urcorner$ for x_1) includes the following step:

$$\begin{array}{ll} & \underline{\mathsf{let}}\, v_1 \Leftarrow ((\lambda z\,.\,\underline{\mathsf{let}}\, v_2 \Leftarrow x_2 \,\underline{\mathsf{in}}\, \underline{\mathsf{succ}}\, \mathsf{lift}\, \ulcorner 42 \urcorner) \mathbin{@} \ulcorner 0 \urcorner) \,\underline{\mathsf{in}}\, w'' \\ \longmapsto_s & \underline{\mathsf{let}}\, v_1 \Leftarrow (\underline{\mathsf{let}}\, v_2 \Leftarrow x_2 \,\underline{\mathsf{in}}\, \underline{\mathsf{succ}}\, \mathsf{lift}\, \ulcorner 42 \urcorner) \,\underline{\mathsf{in}}\, w''. \end{array}$$

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

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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):

$$\begin{array}{ll} & \underline{\mathsf{let}}\, v_1 \Leftarrow ((\lambda z\,.\,\underline{\mathsf{let}}\, v_2 \Leftarrow x_2\ \underline{\mathsf{in}}\ \underline{\mathsf{succ}}\ \mathsf{lift}\ \ulcorner 42 \urcorner)\ @\ \ulcorner 0 \urcorner)\ \underline{\mathsf{in}}\ w'' \\ \longmapsto_s & \underline{\mathsf{let}}\, v_1 \Leftarrow (\underline{\mathsf{let}}\, v_2 \Leftarrow x_2\ \underline{\mathsf{in}}\ \underline{\mathsf{succ}}\ \mathsf{lift}\ \ulcorner 42 \urcorner)\ \underline{\mathsf{in}}\ w'' \\ \longmapsto_f & \underline{\mathsf{let}}\, v_2 \Leftarrow x_2\ \underline{\mathsf{in}}\ \underline{\mathsf{let}}\, v_1 \Leftarrow \underline{\mathsf{succ}}\ \mathsf{lift}\ \ulcorner 42 \urcorner\ \underline{\mathsf{in}}\ w''. \end{array}$$

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.

Since it is the task of binding-time analysis to direct the specializer, the binding-time analysis must also 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:

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

A specializer that incorporates the flattening rule now gives the desired behaviour:

$$\begin{array}{ll} & \mathsf{let}\, v_1 \Leftarrow ((\lambda z\,.\,\underline{\mathsf{let}}\, v_2 \Leftarrow x_2\ \underline{\mathsf{in}}\ \mathsf{succ}\ x_1)\ @\ \ulcorner 0 \urcorner)\ \mathsf{in}\ w''' \\ \longmapsto_s & \mathsf{let}\, v_1 \Leftarrow (\underline{\mathsf{let}}\, v_2 \Leftarrow x_2\ \underline{\mathsf{in}}\ \mathsf{succ}\ \ulcorner 42 \urcorner)\ \mathsf{in}\ w''' \\ \longmapsto_f & \underline{\mathsf{let}}\, v_2 \Leftarrow x_2\ \underline{\mathsf{in}}\ \mathsf{let}\, v_1 \Leftarrow \mathsf{succ}\ \ulcorner 42 \urcorner\ \mathsf{in}\ w''' \\ \longmapsto_s & \underline{\mathsf{let}}\, v_2 \Leftarrow x_2\ \underline{\mathsf{in}}\ \mathsf{let}\, v_1 \Leftarrow \mathsf{unit}\ \ulcorner 43 \urcorner\ \mathsf{in}\ w''' \\ \longmapsto_s & \underline{\mathsf{let}}\, v_2 \Leftarrow x_2\ \underline{\mathsf{in}}\ w'''[v_1\ :=\ \ulcorner 43 \urcorner]\,. \end{array}$$

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 that incorporate the flattening rules. The first flattening rule gives the previously discussed binding-time improvements. The second rule is added for stylistic reasons, since 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

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

$$\frac{\Gamma \vdash^{+}_{bt} e_1 : \widetilde{\tau_1}[w_1 : \mathsf{d}] \qquad \Gamma, x : \tau_1[\mathsf{d}] \vdash^{+}_{bt} e_2 : \widetilde{\tau_2}[w_2 : \varphi]}{\Gamma \vdash^{+}_{bt} \mathsf{let}\ x \Leftarrow e_1\ \mathsf{in}\ e_2 : \widetilde{\tau_2}[\underline{\mathsf{let}}\ x \Leftarrow w_1\ \underline{\mathsf{in}}\ w_2 : \varphi]}(**)$$

Fig. 9. Modified binding-time constraints for Λ^{bt}_{ml}

Flattening rules:

$$\mathsf{let}\ x_2 \Leftarrow (\underline{\mathsf{let}}\ x_1 \Leftarrow w_1\ \underline{\mathsf{in}}\ w_2)\ \mathsf{in}\ w_3 \longmapsto_f \underline{\mathsf{let}}\ x_1 \Leftarrow w_1\ \underline{\mathsf{in}}\ \mathsf{let}\ x_2 \Leftarrow w_2\ \mathsf{in}\ w_3$$

$$\mathsf{let}\ x_2 \Leftarrow (\mathsf{let}\ x_1 \Leftarrow w_1\ \mathsf{in}\ w_2)\ \mathsf{in}\ w_3 \longmapsto_f \mathsf{let}\ x_1 \Leftarrow w_1\ \mathsf{in}\ \mathsf{let}\ x_2 \Leftarrow w_2\ \mathsf{in}\ w_3$$

$$\frac{w \longmapsto_f w'}{w \longmapsto_{s^+} w'}$$

Interpretation rules:

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

$$\frac{w \longmapsto_i w'}{E_i[w] \longmapsto_i E_i[w']} \qquad E_i \quad ::= \quad \mathsf{let}\ x \Leftarrow [\cdot]\ \mathsf{in}\ w_2$$

$$\frac{w \longmapsto_i w'}{w \longmapsto_{s^+} w'} \quad \textit{where } w \not\longmapsto_f w''$$

Compilation rules:

$$\mathsf{lift}\ \ulcorner n \urcorner \longmapsto_{s^+} \underline{\ulcorner n \urcorner} \qquad \frac{w \longmapsto_{s^+} w'}{E_s[w] \longmapsto_{s^+} E_s[w']}$$

$$\begin{array}{rcl}
E_s & ::= & \underline{\mathsf{pred}}\ [\cdot] \mid \underline{\mathsf{succ}}\ [\cdot] \mid \underline{\mathsf{if0}}\ [\cdot]\ w_2\ w_3 \mid \underline{\mathsf{if0}}\ r_1\ [\cdot]\ w_3 \mid \underline{\mathsf{if0}}\ r_1\ r_2\ [\cdot] \mid \\
 & & \underline{\lambda} x\,.\,[\cdot] \mid [\cdot]\ \underline{@}\ w_1 \mid r_0\ \underline{@}\ [\cdot] \mid \underline{\mathsf{fix}}\,x.\,[\cdot] \mid \\
 & & \underline{\mathsf{unit}}\ [\cdot] \mid \underline{\mathsf{let}}\ x \Leftarrow [\cdot]\ \underline{\mathsf{in}}\ w_2 \mid \underline{\mathsf{let}}\ x \Leftarrow r_1\ \underline{\mathsf{in}}\ [\cdot] \qquad \textit{...where } r_i \in \textit{Residual-terms}[\Lambda^{bt}_{ml}]
\end{array}$$

Fig. 10. Modified specialization rules for Λ^{bt}_{ml}

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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^{bt}_{ml} 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 $\Gamma_d \vdash_{bt} e : \widetilde{\tau}[w : \widetilde{\varphi}]$) now reflects the improved treatment of let constructs.

Lemma 3. (specialization properties)

— If $\Gamma_d \vdash^+_{bt} e : \mathsf{nat}[w : \mathsf{s}]$, then $w \equiv \ulcorner n \urcorner \equiv e$ for some number n.

— If $\Gamma_d \vdash^+_{bt} e : \tau_1 \to \tau_2[w : \varphi_1 \to \varphi_2]$, then $w \equiv \lambda x.\, w'$, $e \equiv \lambda x.\, \lfloor w' \rfloor$ and $\Gamma_d, x : \tau_1[\varphi_1] \vdash^+_{bt} \lfloor w' \rfloor : \tau_2[w' : \varphi_2]$ for some $w' \in \mathit{Terms}[\Lambda^{bt}_{ml}]$.

— If $\Gamma_d \vdash^+_{bt} e : \widetilde{\tau}[w : \widetilde{\varphi}]$, then exactly one of the following statements holds:

1 $w \equiv \mathsf{unit}\, w'$, $e \equiv \mathsf{unit}\, \lfloor w' \rfloor$ and $\Gamma \vdash^+_{bt} \lfloor w' \rfloor : \tau[w' : \varphi]$ for some $w' \in \mathit{Terms}[\Lambda^{bt}_{ml}]$, or

2 $w \equiv \underline{\mathsf{let}}\, x \Leftarrow w_1\, \underline{\mathsf{in}}\, w_2$, $e \equiv \mathsf{let}\, x \Leftarrow \lfloor w_1 \rfloor\, \mathsf{in}\, \lfloor w_2 \rfloor$, and $\Gamma_d \vdash^+_{bt} \lfloor w_1 \rfloor : \widetilde{\tau}_1[w_1 : \mathsf{d}]$, $\Gamma_d, x : \tau_1[\mathsf{d}] \vdash^+_{bt} \lfloor w_2 \rfloor : \widetilde{\tau}_2[w_2 : \widetilde{\varphi}]$, for some $w_1, w_2 \in \mathit{Terms}[\Lambda^{bt}_{ml}]$, or

3 $w \longmapsto_i w'$ and $\Gamma_d \vdash^+_{bt} \lfloor w' \rfloor : \widetilde{\tau}[w' : \widetilde{\varphi}]$ and $e \equiv \lfloor w \rfloor \longrightarrow_{ml} \lfloor w' \rfloor$, or

4 $w \longmapsto_f w'$ and $\Gamma_d \vdash^+_{bt} \lfloor w' \rfloor : \widetilde{\tau}[w' : \widetilde{\varphi}]$ and $e \equiv \lfloor w \rfloor \longrightarrow_{ml} \lfloor w' \rfloor$.

— If $\Gamma_d \vdash^+_{bt} e : \mathsf{nat}[w : \mathsf{d}]$, then exactly one of the following statements holds:

1 $w \in \mathit{Residual\text{-}terms}[\Lambda^{bt}_{ml}]$, or

2 $w \longmapsto_{s^+} w'$, $\Gamma_d \vdash^+_{bt} \lfloor w' \rfloor : \tau[w' : \mathsf{d}]$, and $e \equiv \lfloor w' \rfloor \equiv \lfloor w \rfloor \equiv \ulcorner n \urcorner$ for some number n.

— If $\Gamma_d \vdash^+_{bt} e : \tau_1 \to \tau_2[w : \mathsf{d}]$, then exactly one of the following statements holds:

1 $w \in \mathit{Residual\text{-}terms}[\Lambda^{bt}_{ml}]$, or

2 $w \longmapsto_{s^+} w'$ and $\Gamma_d \vdash^+_{bt} \lfloor w' \rfloor : \tau[w' : \mathsf{d}]$ and $e \equiv \lfloor w \rfloor \longrightarrow_{ml} \lfloor w' \rfloor$.

— If $\Gamma_d \vdash^+_{bt} e : \widetilde{\tau}[w : \mathsf{d}]$, then exactly one of the following statements holds:

1 $w \in \mathit{Residual\text{-}terms}[\Lambda^{bt}_{ml}]$, or

2 $w \longmapsto_{s^+} w'$ and $\Gamma_d \vdash^+_{bt} \lfloor w' \rfloor : \widetilde{\tau}[w' : \mathsf{d}]$ and $e \equiv \lfloor w \rfloor \longrightarrow_{ml} \lfloor w' \rfloor$.

Proof. The proof is by induction over the height of the derivation of $\Gamma_d \vdash^+_{bt} e : \tau[w : \varphi]$, relying on the property that 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 $\Gamma \vdash^+_{bt} e_0[x := e_1] : \tau_0[w_0[x := w_1] : \varphi_0]$. □

As before, if $w \longmapsto_{s^+} w'$ and $w \longmapsto_{s^+} w''$, 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, for example, the partial evaluator Similix (Bondorf and Danvy 1991). To substantiate this claim, we recast Bondorf's continuation-based specializer

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Interpretation rules:

$$\langle \mathsf{unit}\, w\,,\, \mathit{nil} \rangle \Downarrow \mathsf{unit}\, w$$

$$\frac{\langle \mathsf{unit}\, \ulcorner n+1 \urcorner,\, \vec{E}_i \rangle \Downarrow a}{\langle \mathsf{succ}\, \ulcorner n \urcorner,\, \vec{E}_i \rangle \Downarrow a} \qquad \frac{\langle \mathsf{unit}\, \ulcorner n \urcorner,\, \vec{E}_i \rangle \Downarrow a}{\langle \mathsf{pred}\, \ulcorner n+1 \urcorner,\, \vec{E}_i \rangle \Downarrow a} \qquad \frac{\langle \mathsf{unit}\, \ulcorner 0 \urcorner,\, \vec{E}_i \rangle \Downarrow a}{\langle \mathsf{pred}\, \ulcorner 0 \urcorner,\, \vec{E}_i \rangle \Downarrow a}$$

$$\frac{\langle w_2\,,\, \vec{E}_i \rangle \Downarrow a}{\langle \mathsf{if0}\, \ulcorner 0 \urcorner\, w_2\, w_3\,,\, \vec{E}_i \rangle \Downarrow a} \qquad \frac{\langle w_3\,,\, \vec{E}_i \rangle \Downarrow a}{\langle \mathsf{if0}\, \ulcorner n+1 \urcorner\, w_2\, w_3\,,\, \vec{E}_i \rangle \Downarrow a}$$

$$\frac{\langle w_0[x := w_1]\,,\, \vec{E}_i \rangle \Downarrow a}{\langle (\lambda x\,.\, w_0) @ w_1\,,\, \vec{E}_i \rangle \Downarrow a} \qquad \frac{\langle w[x := \mathsf{fix}\, x.\, w]\,,\, \vec{E}_i \rangle \Downarrow a}{\langle \mathsf{fix}\, x.\, w\,,\, \vec{E}_i \rangle \Downarrow a}$$

$$\frac{\langle w_1\,,\, \mathsf{let}\, x \Leftarrow \cdot\, \mathsf{in}\, w_2 \;::\; \vec{E}_i \rangle \Downarrow a}{\langle \mathsf{let}\, x \Leftarrow w_1\, \mathsf{in}\, w_2\,,\, \vec{E}_i \rangle \Downarrow a}\ \textit{(Push)}$$

$$\frac{\langle w_2[x := w_1]\,,\, \vec{E}_i \rangle \Downarrow a}{\langle \mathsf{unit}\, w_1\,,\, \mathsf{let}\, x \Leftarrow \cdot\, \mathsf{in}\, w_2 \;::\; \vec{E}_i \rangle \Downarrow a}\ \textit{(Pop)}$$

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

(Bondorf 1992) in terms of Λ_{ml}, and prove that this specializer is equivalent to the $spec^{+}$ of Section 5.2.2.

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 (Nielson and Nielson 1992b)†. The specializer manipulates configurations of the form $\langle w\,,\, \vec{E}_i \rangle$ where $w \in \mathit{Terms}[\Lambda_{ml}^{bt}]$ 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 accumu-

† For 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 (for example, if0) are specialized with respect to the identity continuation (that is, *nil*).

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Compilation rules:

$$\langle \underline{\ulcorner n \urcorner}, nil \rangle \Downarrow \underline{\ulcorner n \urcorner} \qquad \langle x, nil \rangle \Downarrow x \qquad \langle \mathsf{lift}\, \ulcorner n \urcorner, nil \rangle \Downarrow \underline{\ulcorner n \urcorner} \qquad \frac{\langle w, nil \rangle \Downarrow r}{\langle \underline{\mathsf{unit}}\, w, nil \rangle \Downarrow \underline{\mathsf{unit}}\, r}$$

$$\frac{\langle w, nil \rangle \Downarrow r}{\langle \underline{\mathsf{succ}}\, w, nil \rangle \Downarrow \underline{\mathsf{succ}}\, r} \qquad \frac{\langle w, nil \rangle \Downarrow r}{\langle \underline{\mathsf{pred}}\, w, nil \rangle \Downarrow \underline{\mathsf{pred}}\, r}$$

$$\frac{\langle w, nil \rangle \Downarrow r}{\langle \underline{\lambda} x.\, w, nil \rangle \Downarrow \underline{\lambda} x.\, r} \qquad \frac{\langle w, nil \rangle \Downarrow r}{\langle \underline{\mathsf{fix}}\, x.\, w, nil \rangle \Downarrow \underline{\mathsf{fix}}\, x.\, r}$$

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

$$\frac{\langle w_1, nil \rangle \Downarrow r_1 \qquad \langle w_2, nil \rangle \Downarrow r_2 \qquad \langle w_3, nil \rangle \Downarrow r_3}{\langle \underline{\mathsf{if0}}\, w_1\, w_2\, w_3, nil \rangle \Downarrow \underline{\mathsf{if0}}\, r_1\, r_2\, r_3}$$

$$\frac{\langle w_1, nil \rangle \Downarrow r_1 \qquad \langle w_2, \vec{E}_i \rangle \Downarrow r_2}{\langle \underline{\mathsf{let}}\, x \Leftarrow w_1\, \underline{\mathsf{in}}\, w_2, \vec{E}_i \rangle \Downarrow \underline{\mathsf{let}}\, x \Leftarrow r_1\, \underline{\mathsf{in}}\, r_2}\ (Assoc)$$

Fig. 12. Continuation-based specializer for Λ^{bt}_{ml} (part 2)

late contexts (*i.e.*, the stack is delimited to be *nil* – which corresponds to the identity continuation).

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

Theorem 5. For $\Gamma_d \vdash^+_{bt} e : \widetilde{\mathsf{nat}}[w : \mathsf{d}]$,

$$\langle w, nil \rangle \Downarrow r \quad \text{iff} \quad w \longmapsto^*_{s^+} r.$$

Proof. The proof is straightforward and follows the standard pattern for relating a big-step and small-step operational semantics (Gunter 1992, p. 111). However, one must strengthen the statement to be proved so as to handle arbitrary stacks of interpretation contexts in the inductive hypothesis. We refer the reader to the extended version of this paper for the full proof (Hatcliff and Danvy 1996). □

Bondorf's presentation (Bondorf 1992) did not include a specification of binding-time analysis. However, our connection of continuation-based specialization and $spec^+$ shows

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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 for specifying denotational semantics modularly (Moggi 1991). This metalanguage has been used for a variety of purposes including the treatment of computational effects in functional languages (Gordon 1993; Filinski 1994; Liang *et al.* 1995; Steele 1994; Wadler 1992a), staging denotational-semantics definitions (Crole and Gordon 1993) and compiler translations (Benton 1995; Hatcliff and Danvy 1994; Sabry and Wadler 1996), and explaining relationships between various constructive logics (Benton *et al.* 1995). In earlier work (Hatcliff 1994; Hatcliff and Danvy 1994), 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 (Nielsen 1997). 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 (Sørensen *et al.* 1994). 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 (Nielsen 1997) and Dussart and Thiemann (Dussart and Thiemann 1996) 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 (Talcott 1992). Sabry and Felleisen demonstrate how a let-based intermediate language (called *A-normal forms*) can give benefits similar to CPS in dataflow analysis (Sabry and Felleisen 1993; Sabry and Felleisen 1994).

Moggi also gave a *computational λ-calculus* λ_c, which 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.* (Flanagan *et al.* 1994) 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 (Jones *et al.* 1989). Since the

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

work of Jones and Gomard (Gomard 1992; Gomard and Jones 1991), binding-time analysis is more often specified using type systems (we have followed this approach in the present work). Palsberg (Palsberg 1993) and Wand (Wand 1993) further clarify the role of such specifications in their work on the correctness of binding-time analysis.

Specializers have mostly been specified as symbolic interpreters in functional style (Jones *et al.* 1993). We note, however, a recent trend (including the present work) to use operational semantics for specifying specializers (Andersen 1994; Danvy *et al.* 1996; Hannan and Miller 1989; Sørensen *et al.* 1994; Sands 1995). 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 (Hatcliff 1995).

Davies and Pfenning have developed a language for expressing staged computation based on the intuistionistic modal logic S4 (Davies and Pfenning 1996; Davies 1996). The type system of this language is strikingly similar to that of Moggi's computational metalanguage. A modality analogous to the $\widetilde{\cdot}$ 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 (Benton *et al.* 1995) 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, but at the price of reducing specialization. This must be expressed in the binding-time analysis by forbidding binding-times such as $\mathsf{d} \to \mathsf{s}$ (*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 (Nielson and Nielson 1992a): 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 $\mathsf{d} \to \mathsf{s}$). As a

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

corollary, the addition is classified as static. The residual program reads

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

In both the above techniques, 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 $\mathsf{d} \rightarrow \mathsf{s}$ 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 (Consel and Danvy 1991). 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 (Holst and Gomard 1991). 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 (Bondorf 1992; Lawall and Danvy 1994), 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 (Nielson 1982; Sabry and Felleisen 1994): 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 the use of 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 (Lawall and Danvy 1994), and Glück and Jørgensen to devise a *multi-level cogen* (Glück and Jørgensen 1995). Thiemann has united both lines of work (Thiemann 1996). 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* (Glück and Klimov 1993) where let-bindings (in the style of Λ_{ml}) are a central feature (Hatcliff and Glück 1996).

We will now briefly attempt a taxonomy of continuation-based partial evaluators. The earliest one is reported in the literature by Bondorf (Bondorf 1992). Its goal is precisely

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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 (Lawall and Danvy 1994). A parallel development is taking place in the 'cogen' approach to partial evaluation (Bondorf and Dussart 1994; Lawall and Danvy 1995; Thiemann 1996). All these continuation-based specializers are specified as symbolic interpreters in functional style. An increasing number of specializers, however, are specified operationally. Here we have presented 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 (Dussart 1997).

As for data-flow binding-time improvements, they arise from insufficient binding-time coercions (Danvy *et al.* 1995; Danvy *et al.* 1996). In particular, the binding-time improvement arising from the presence of booleans and disjoint sums is known as 'The Trick' (Danvy *et al.* 1996; Jones *et al.* 1993). 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 (Hatcliff and Danvy 1994). 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), our let-rearranging rules coincide with the case-rearranging rules (also known as *commuting conversions*) that can be found in partial evaluation (Danvy *et al.* 1996), program extraction (Paulin-Mohring and Werner 1993), and natural-deduction proof theory (Girard *et al.* 1989; Prawitz 1965).

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 (Danvy and Hatcliff 1993). 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 (Nielsen and Sørensen 1995). 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 (Hatcliff 1994).

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

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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

Fig. 13. The language Λ

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 (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.

Appendix 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 given below.

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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

Fig. 14. Call-by-name encoding $\mathscr{E}_n$ into Λ_{ml}

Definition 7. (Values)

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

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†.

† One 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 $\mathscr{E}_v$), evaluation is forced immediately after an argument is received by a function (Hatcliff and Danvy 1994).

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

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

Fig. 15. Call-by-value encoding $\mathscr{E}_v$ into Λ_{ml}

References

Andersen, L. O. (1994) *Program Analysis and Specialization for the C Programming Language*. Ph.D. thesis, DIKU, Computer Science Department, University of Copenhagen (DIKU Rapport 94/19).

Barendregt, H. (1984) *The Lambda Calculus – Its Syntax and Semantics*, North-Holland.

Benton, P. N. (1995) *Strictness Analysis of Lazy Functional Programs*, Ph.D. thesis, Computer Laboratory, University of Cambridge.

Benton, P. N., Bierman, G. M. and de Paiva, V. C. V. (1995) Computational types from a logical perspective I. Technical report 365, Computer Laboratory, University of Cambridge.

Bondorf, A. (1992) Improving binding times without explicit cps-conversion. In: Clinger, W. (ed.) Proceedings of the 1992 ACM Conference on Lisp and Functional Programming, San Francisco, California. *LISP Pointers* **V** (1) 1–10.

Bondorf, A. and Danvy, O. (1991) Automatic autoprojection of recursive equations with global variables and abstract data types. *Science of Computer Programming* **16** 151–195.

Bondorf, A. and Dussart, D. (1994) Improving CPS-based partial evaluation: Writing cogen by hand. In: Sestoft, P. and Søndergaard, H. (eds.) Proceedings of the ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation, Orlando, Florida. Technical Report 94/9, University of Melbourne 1–10.

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Consel, C. (1993) A tour of Schism: A partial evaluation system for higher-order applicative languages. In: Schmidt, D. A. (ed.) *Proceedings of the Second ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, Copenhagen, Denmark, ACM Press 145–154.

Consel, C. and Danvy, O. (1991) For a better support of static data flow. In: Hughes, J. (ed.) Proceedings of the Fifth ACM Conference on Functional Programming and Computer Architecture, Cambridge, Massachusetts. *Springer-Verlag Lecture Notes in Computer Science* **523** 496–519.

Consel, C. and Danvy, O. (1993) Tutorial notes on partial evaluation. In: Graham, S. L. (ed.) *Proceedings of the Twentieth Annual ACM Symposium on Principles of Programming Languages*, Charleston, South Carolina, ACM Press 493–501.

Crole, R. L. and Gordon, A. D. (1993) Factoring an adequacy proof (preliminary report). In: O'Donnell, J. T. and Hammond, K. (eds.) *Functional Programming, Glasgow 1993 Workshops in Computing*, Springer-Verlag 9–25.

Crole, R. L. and Pitts, A. M. (1992) New foundations for fixpoint computations: FIX-hyperdoctrines and the FIX-logic. *Information and Computation* **98** 171–210.

Danvy, O. (1996) Type-directed partial evaluation. In: Steele Jr., G. L. (ed.) *Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Programming Languages*, St. Petersburg Beach, Florida, ACM Press 242–257.

Danvy, O. and Filinski, A. (1992) Representing control, a study of the CPS transformation. *Mathematical Structures in Computer Science* **2** (4) 361–391.

Danvy, O. and Hatcliff, J. (1993) CPS transformation after strictness analysis. *ACM Letters on Programming Languages and Systems* **1** (3) 195–212.

Danvy, O., Malmkjær, K. and Palsberg, J. (1995) The essence of eta-expansion in partial evaluation. *LISP and Symbolic Computation* **8** (3) 209–227. (An earlier version appeared in the proceedings of the 1994 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation.)

Danvy, O., Malmkjær, K. and Palsberg, J. (1996) Eta-expansion does The Trick. *ACM Transactions on Programming Languages and Systems* (to appear).

Davies, R. (1996) A temporal-logic approach to binding-time analysis. In: *Proceedings of the Eleventh Annual IEEE Symposium on Logic in Computer Science*, New Brunswick, New Jersey, IEEE Computer Society Press 184–195.

Davies, R. and Pfenning, F. (1996) A modal analysis of staged computation. In: Steele Jr., G. L. (ed.) *Proceedings of the Twenty-Third Annual ACM Symposium on Principles of Programming Languages*, St. Petersburg Beach, Florida, ACM Press 258–283.

Dussart, D. (1997) Ph.D. thesis, Katholieke Universiteit Leuven, Leuven, Belgium (forthcoming).

Dussart, D. and Thiemann, P. (1996) Imperative functional specialization. Berichte des Wilhelm-Schickard-Instituts WSI-96-28, Universität Tübingen.

Filinski, A. (1994) Representing monads. In: Boehm, H.-J. (ed.) *Proceedings of the Twenty-First Annual ACM Symposium on Principles of Programming Languages*, Portland, Oregon, ACM Press 446–457.

Flanagan, C., Sabry, A., Duba, B. F. and Felleisen, M. (1994) The essence of compiling with continuations. In: Wall, D. W. (ed.) Proceedings of the ACM SIGPLAN'93 Conference on Programming Languages Design and Implementation, Albuquerque, New Mexico. *SIGPLAN Notices* **28** (6) 237–247.

Girard, J.-Y., Lafont, Y. and Taylor, P. (1989) *Proofs and Types*, Cambridge Tracts in Theoretical Computer Science **7**, Cambridge University Press.

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Glück, R. and Jørgensen. J. (1995) Efficient multi-level generating extensions for program specialization. In: Hermenegildo, M. and Swierstra, S. D. (eds.) Seventh International Symposium on Programming Language Implementation and Logic Programming, Utrecht, The Netherlands. *Springer-Verlag Lecture Notes in Computer Science* **982**.

Glück, R. and Jørgensen. J. (1993) Occam's razor in metacomputation: the notion of a perfect process tree. In: Cousot, P., Falaschi, M., Filè, G. and Rauzy, A. (eds.) Proceedings of the Third International Workshop on Static Analysis WSA'93, Padova, Italy. *Springer-Verlag Lecture Notes in Computer Science* **724** 112–123.

Gomard, C. K. (1992) A self-applicable partial evaluator for the lambda-calculus: Correctness and pragmatics. *ACM Transactions on Programming Languages and Systems* **14** (2) 147–172.

Gomard, C. K. and Jones, N. D. (1991) A partial evaluator for the untyped lambda-calculus. *Journal of Functional Programming* **1** (1) 21–69.

Gordon, A. D. (1993) *Functional Programming and Input/Output*, Ph.D. thesis, Computer Laboratory, University of Cambridge (Technical Report No. 285).

Gunter, C. A. (1992) *Semantics of Programming Languages: Structures and Techniques*, MIT Press.

Hannan, J. and Miller, D. (1989) Deriving mixed evaluation from standard evaluation for a simple functional language. In: van de Snepscheut, J. L. A. (ed.) Proceedings of a Conference on Mathematics of Program Construction, Groningen, The Netherlands. *Springer-Verlag Lecture Notes in Computer Science* **375** 239–255.

Hatcliff, J. (1994) *The Structure of Continuation-Passing Styles*, Ph.D. thesis, Department of Computing and Information Sciences, Kansas State University.

Hatcliff, J. (1995) Mechanically verifying the correctness of an offline partial evaluator. In: Hermenegildo, M. and Swierstra, S. D. (eds.) Seventh International Symposium on Programming Language Implementation and Logic Programming, Utrecht, The Netherlands. *Springer-Verlag Lecture Notes in Computer Science* **982** 279–298.

Hatcliff, J. and Danvy, O. (1994) A generic account of continuation-passing styles. In: Boehm, H.-J. (ed.) *Proceedings of the Twenty-First Annual ACM Symposium on Principles of Programming Languages*, Portland, Oregon, ACM Press 458–471.

Hatcliff, J. and Danvy, O. (1996) A computational formalization for partial evaluation (extended version). Technical Report BRICS RS-96-34, Computer Science Department, Aarhus University.

Hatcliff, J. and Glück, R. (1996) An operational theory of self-applicable on-line program specialization. In: Danvy, O., Glück, R. and Thiemann, P. (eds.) Partial Evaluation, Dagstuhl, Germany. *Springer-Verlag Lecture Notes in Computer Science* **1110** 161–182.

Holst, C. K. and Gomard, C. K. (1991) Partial evaluation is fuller laziness. In: Hudak, P. and Jones, N. D. (eds.) Proceedings of the ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation, New Haven, Connecticut. *SIGPLAN Notices* **26** (9) 223–233.

Jones, N. D., Gomard, C. K. and Sestoft, P. (1993) *Partial Evaluation and Automatic Program Generation*, Prentice Hall International Series in Computer Science.

Jones, N. D., Sestoft, P. and Søndergaard, H. (1989) MIX: A self-applicable partial evaluator for experiments in compiler generation. *LISP and Symbolic Computation* **2** (1) 9–50.

Launchbury, J. (1991) A strongly-typed self-applicable partial evaluator. In: Hughes, J. (ed.) Proceedings of the Fifth ACM Conference on Functional Programming and Computer Architecture, Cambridge, Massachusetts. *Springer-Verlag Lecture Notes in Computer Science* **523** 145–164.

Lawall, J. L. and Danvy, O. (1994) Continuation-based partial evaluation. In: Talcott, C. L. (ed.) Proceedings of the 1994 ACM Conference on Lisp and Functional Programming, Orlando, Florida. *LISP Pointers* **VII** (3).

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Lawall, J. L. and Danvy, O. (1995) 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.)

Liang, S., Hudak, P. and Jones, M. (1995) Monad transformers and modular interpreters. In: Lee, P. (ed.) *Proceedings of the Twenty-Second Annual ACM Symposium on Principles of Programming Languages*, San Francisco, California 333–343.

Moggi, E. (1991) Notions of computation and monads. *Information and Computation* **93** 55–92.

Nielsen, K. (1997) Master's thesis, DIKU, Computer Science Department, University of Copenhagen (forthcoming).

Nielsen, K. and Sørensen, M. H. (1995) Call-by-name CPS-translation as a binding-time improvement. In: Mycroft, A. (ed.) Static Analysis. *Springer-Verlag Lecture Notes in Computer Science* **983** 296–313.

Nielson, F. (1982) A denotational framework for data flow analysis. *Acta Informatica* **18** 265–287.

Nielson, F. and Nielson, H. R. (1992a) *Two-Level Functional Languages*, Cambridge Tracts in Theoretical Computer Science **34**, Cambridge University Press.

Nielson, H. R. and Nielson, F. (1992b) *Semantics with Applications, a formal introduction*, Wiley Professional Computing, John Wiley and Sons.

Palsberg, J. (1993) Correctness of binding-time analysis. In: Jones, N. D. (ed.) *Journal of Functional Programming* **3** (3) (Special issue on Partial Evaluation) 347–363.

Paulin-Mohring, C. and Werner, B. (1993) Synthesis of ML programs in the system Coq. *Journal of Symbolic Computation* **15** 607–640.

Plotkin, G. D. (1975) Call-by-name, call-by-value and the λ-calculus. *Theoretical Computer Science* **1** 125–159.

Prawitz, D. (1965) *Natural Deduction*, Almquist and Wiksell, Uppsala.

Sabry, A. and Felleisen, M. (1993) Reasoning about programs in continuation-passing style. *LISP and Symbolic Computation* **6** (3/4) 289–360.

Sabry, A. and Felleisen, M. (1994) Is continuation-passing useful for data flow analysis? In: Sarkar, V. (ed.) Proceedings of the ACM SIGPLAN'94 Conference on Programming Languages Design and Implementation, Orlando, Florida. *SIGPLAN Notices* **29** (6) 1–12.

Sabry, A. and Wadler, P. (1996) Compiling with reflections. In: Dybvig, R. K. (ed.) *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming*, Philadelphia, Pennsylvania, ACM Press 13–24.

Sands, D. (1995) Proving the correctness of recursion-based automatic program transformations. In: Mosses, P., Nielsen, M., and Schwartzbach, M. (eds.) Proceedings of TAPSOFT '95, Aarhus, Denmark. *Springer-Verlag Lecture Notes in Computer Science* **915** 681–695.

Sørensen, M. H., Glück, R. and Jones, N. (1994) Towards unifying partial evaluation, deforestation, supercompilation, and GPC. In: Sannella, D. (ed.) Proceedings of the Fifth European Symposium on Programming, Edinburgh, Scotland. *Springer-Verlag Lecture Notes in Computer Science* **788** 485–500.

Steele Jr., G, L. (1994) Building interpreters by composing monads. In: Boehm, H.-J. (ed.) *Proceedings of the Twenty-First Annual ACM Symposium on Principles of Programming Languages*, Portland, Oregon, ACM Press 472–492.

Talcott, C. (1992) A theory for program and data type specification. *Theoretical Computer Science* **104** (1) 129–159.

Thiemann, P. (1996) Cogen in six lines. In: Dybvig, R. K. (ed.) *Proceedings of the 1996 ACM SIGPLAN International Conference on Functional Programming*, Philadelphia, Pennsylvania, ACM Press 180–189.

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press

Wadler, P. (1992a) The essence of functional programming (tutorial). In: Appel, A. W. (ed.) *Proceedings of the Nineteenth Annual ACM Symposium on Principles of Programming Languages*, Albuquerque, New Mexico, ACM Press 1–14.

Wadler, P. (1992b) Comprehending monads. *Mathematical Structures in Computer Science* **2** (4) 461–493.

Wand, M. (1993) Specifying the correctness of binding-time analysis. In: Jones, N. D. (ed.) *Journal of Functional Programming* **3** (3) (Special issue on Partial Evaluation) 365–387.

https://doi.org/10.1017/S0960129597002405 Published online by Cambridge University Press