

# Metacomputation-based Compiler Architecture

William L. Harrison    Samuel N. Kamin

Department of Computer Science  
University of Illinois, Urbana-Champaign  
Urbana, Illinois 61801-2987  
{harrison,kamin}@cs.uiuc.edu

## Abstract

*This paper presents a modular and extensible style of language specification based on metacomputations. This style uses two monads to factor the static and dynamic parts of the specification, thereby staging the specification and achieving strong binding-time separation. Because metacomputations are defined in terms of monads, they can be constructed modularly and extensibly using monad transformers. A number of language constructs are specified: expressions, control-flow, imperative features, and block structure. Metacomputation-style specification lends itself to semantics-directed compilation, which we demonstrate by creating a modular compiler for a block-structured, imperative while language.*

*Keywords:* Compilers, Partial Evaluation, Semantics-Based Compilation, Programming Language Semantics, Monads, Monad Transformers, Pass Separation.

## 1 Introduction

*Metacomputations*—computations that produce computations—arise naturally in the compilation of programs. Figure 1 illustrates this idea. The source language program  $s$  is taken as input by the compiler, which produces a target language program  $t$ . So, compiling  $s$  produces another computation—namely, the computation of  $t$ . Observe that there are two entirely distinct notions of computation here: the compilation of  $s$  and the execution of  $t$ . The reader will recognize this distinction as the classic separation of static from dynamic. Thus, *staging* is an instance of metacomputation.

**The main contributions of this paper are:** (1) *Compiler architecture based on metacomputations:* Metacomputation-based compiler architecture yields substantially simpler language definitions than in [8], while still retaining its modular “mix and match” approach to compiler construction. Combining the metacomputation-based “reusable compiler building blocks” is also much simpler than combining those in [8] (as is proving their correctness). (2) *A modular and extensible method of staging denotational specifications based on metacomputations:* A style of language specification based on metacomputation is proposed in which the static and dynamic parts of a language specification are factored into distinct monads[7, 13, 16, 24]. (3) *Direct-style*

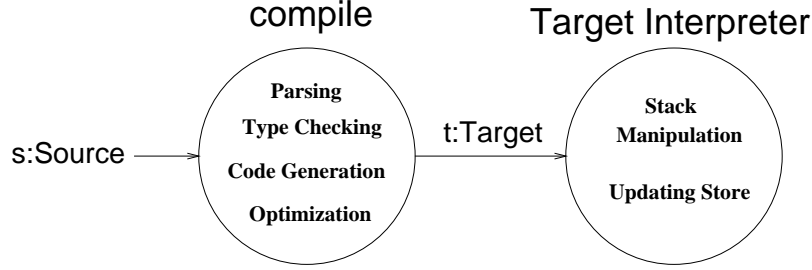


Figure 1: Handwritten compiler as metacomputation

---

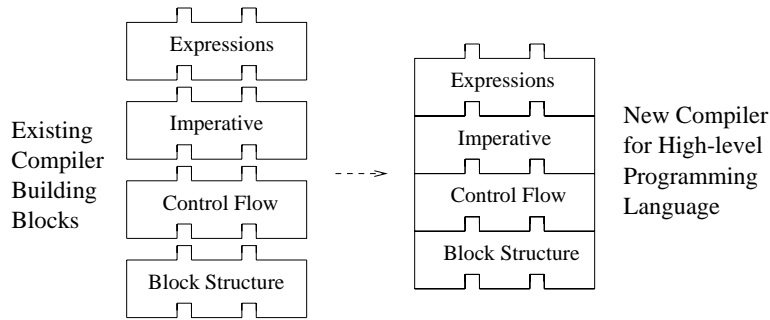


Figure 2: Modular Compilers: Existing compiler building blocks combine to make new compiler

---

*specifications*: instead of writing all specifications in continuation-passing style, here we write in direct style, invoking the CPS monad transformer only when needed. This naturally simplifies many of the equations, and although less essential than (1) and (2), it also helps to make the approach more practical.

We believe this style of language specification may have many uses, but in this paper we concentrate on one: modular compilation. *Modular compilers* are compilers built from building blocks that represent *language features* rather than *compilation phases*, as illustrated in Figure 2. Espinosa [7] and Liang & Hudak [13] showed how to construct modular *interpreters* using the notion of *monads* [7, 13, 16, 24] — or, more precisely, *monad transformers*.

The current authors built on those ideas to produce modular *compilers* in [8]. However, there the notion of *staging*, though conceptually at the heart of the approach, was not *explicit* in the compiler building blocks we constructed. As in traditional monadic semantics, the monadic structure was useful in creating the domains, but those domains, once constructed, were “monolithic;” that is, they gave no indication of which parts were for dynamic aspects of the computation and which for static aspects. The result was awkwardness in communicating between these aspects of the domain, which meant that “gluing together” compiler blocks was sometimes delicate. However, metacomputation-based compiler architecture *completely* alleviates this awkwardness, so that combining compiler blocks is simply a matter of applying the appropriate monad transformers.

Indeed, metacomputation is *purposely* avoided in [7, 13, 8]. A key aspect of that work is that monad transformers are used to create the single monad used to interpret or compile the

language. The problem that inspired it was that monads don't compose nicely. Given monads  $M$  and  $M'$ , the composed monad  $M \circ M'$  — corresponding to an  $M$ -computation that produces an  $M'$  computation — usually does not produce the “right” monolithic domain. However, there may exist monad transformers  $T_M$  and  $T_{M'}$  such that  $T_M \text{ ld} = M$  and  $T_{M'} \text{ ld} = M'$ , where  $(T_M \circ T_{M'}) \text{ ld}$  does give the “right” domain. The difference between composing monads and composing monad transformers is what makes these approaches work — monad transformers are a way to *avoid* metacomputation.

In this paper, we show that, for some purposes, metacomputation may be exactly what one wants: *Defining a compiler block via the metacomputation of two monads gives an effective representation of staging.* We are not advocating abandoning monad transformers: the two monads can be constructed using them, with the attendant advantages of that approach. We are simply saying that having two monads — what might be called the *static* and *dynamic* monads — and composing them seems to give the “right” domain for modular compilation.

The next section explains the advantages for modular compilation of metacomputation-based language specification over the monolithic style. Section 3 reviews the most relevant related work. In Section 4, we review the theory of monads and monad transformers and their use in language specification. Section 5 presents a case study in metacomputation-style language specification; its subsections present metacomputation-style specifications for expressions, control flow, block structure, and imperative features, respectively. Section 6 shows how to combine these compiler building blocks into a compiler for the combined language, and presents a compiler and an example compilation. Section 7 discusses the impact of metacomputation-based specification on compiler correctness. Finally, Section 8 summarizes this work and outlines future research.

## 2 Why Metacomputations?

In this section, we will describe at a high level why two monads are better than one for modular compilation. Using metacomputations instead of a single monolithic monad simplifies the use of the “code store” (defined below) in the specification of reusable compiler building blocks.

In [8], we borrowed a technique from denotational semantics[23] for modeling jumps, namely storing command continuations in a “code store” and denoting “`jump L`” as “execute the continuation at label  $L$  in the code store.” Viewing command continuations as machine code is a common technique in semantics-directed compilation[25, 22]. Because our language specifications were in monadic style, it was a simple matter to add label generator and code store states to the underlying monad. Indeed, the primary use for monads in functional programming seems to be that of adding state-like features to purely functional languages and programs[24, 20], and the fact that we structured our monads in [8] with monad transformers made adding the new states simple.

The use of a code store is integral to the modular compilation technique described in [8]. We use it to compile control-flow and procedures, and the presence of the code store in our language specifications allowed us to make substantial improvements over Reynolds[22] (e.g., avoiding infinite programs through jumps and labels). Yet the mixing of static with dynamic data into one “monolithic” monad causes a number of problems with using the code store. Consider the program “`if b then (if b' then c)`”. Compiling the outer “`if`” with initial continuation `halt` and label 0 will result in the continuation “`[[if b' then c]; halt`” being stored at label 0 and the label counter being incremented. The problem here is that trying to compile this continuation

via partial evaluation will fail. Why? Because having been *stored* rather than *executed*, it will not have access to the next label 1. Instead, the partial evaluator will try to increment a (dynamic) variable rather than an actual (static) integer, and this will cause an error (e.g., a partial evaluator can evaluate “1+1” but not “x+1”). In [8], the monolithic style specifications forced all static data to be explicitly passed to stored command continuations, although this was at the expense of modularity. In fact to compile **if-then-else**, the snapback operator had to be used. These complications also make reasoning about compilers constructed in [8] difficult. We shall demonstrate in Section 5 that using metacomputations results in vastly simpler compiler specifications than in [8] and that this naturally makes them easier to reason about.

### 3 Related work

Espinosa [7] and Hudak, Liang, and Jones [13] use monad transformers to create modular, extensible interpreters. Liang [12, 14] addresses the question of whether compilers can be developed similarly, but since he does not compile to machine language, many of the issues we confront—especially staging—do not arise.

A syntactic form of metacomputation can be found in the two-level  $\lambda$ -calculus of Nielson[19]. Two-level  $\lambda$ -calculus contains two distinct  $\lambda$ -calculi—representing the static and dynamic *levels*. Expressions of mixed level, then, have strongly separated binding times by definition. Nielson[18] applies two-level  $\lambda$ -calculus to code generation for a typed  $\lambda$ -calculus, and Nielson[19] presents an algorithm for static analysis of a typed  $\lambda$ -calculus which converts one-level specifications into two-level specifications. Mogensen[15] generalizes this algorithm to handle variables of mixed binding times. The present work offers a semantic alternative to the two-level  $\lambda$ -calculus. We formalize distinct levels (in the sense of Nielson[19]) as distinct monads, and the resulting specifications have all of the traditional advantages of monadic specifications (reusability, extensibility, and modularity). While our binding time analysis is not automatic as in [19, 15], we consider a far wider range of programming language features than they do.

Danvy and Vestergaard [5] show how to produce code that “looks like” machine language, by expressing the source language semantics in terms of machine language-like combinators (e.g., “popblock”, “push”). When the interpreter is closed over these combinators, partial evaluation of this closed term with respect to a program produces a completely *dynamic* term, composed of a sequence of combinators, looking very much like machine language. This approach is key to making the monadic structure useful for compilation.

Reynolds’ [22] demonstration of how to produce efficient code in a compiler derived from the functor category semantics of an Algol-like language was an original inspiration for this study. Our approach to compilation improves on Reynolds’s in two ways: it is monad-structured—that is, built from interchangeable parts—and it includes jumps and labels where Reynolds simply allowed code duplication and infinite programs.

### 4 Monads and Monad Transformers

In this section, we review the theory of monads [16, 24] and monad transformers [7, 13]. Readers familiar with these topics may skip the section.

A *monad* is a type constructor  $\mathbf{M}$  together with a pair of functions (obeying certain algebraic

---

**Identity Monad  $Id$ :**
$$\begin{aligned} \text{id } \tau &= \tau \\ \mathbf{unit}_{\text{id}} x &= x \\ x \star_{\text{id}} f &= f x \end{aligned}$$
**Environment Monad Transformer  $\mathcal{T}_{\text{Env}}$ :**
$$\begin{aligned} M'\tau &= \mathcal{T}_{\text{Env}} \text{ Env } M \tau = \text{Env} \rightarrow M\tau \\ \mathbf{unit}_{M'} x &= \lambda \rho : \text{Env}. \mathbf{unit}_M x \\ x \star_{M'} f &= \lambda \rho : \text{Env}. (x \rho) \star_M (\lambda a. f a \rho) \\ \text{lift}_{M\tau \rightarrow M'\tau} x &= \lambda \rho : \text{Env}. x \\ \mathbf{rdEnv} : M' \text{Env} &= \lambda \rho : \text{Env}. \mathbf{unit}_{M'} \rho \\ \mathbf{inEnv}(\rho : \text{Env}, x : M'\tau) &= \lambda \_ . (x \rho) : M'\tau \end{aligned}$$
**CPS Monad Transformer  $\mathcal{T}_{\text{CPS}}$ :**
$$\begin{aligned} M'\tau &= \mathcal{T}_{\text{CPS}} \text{ ans } M \tau = \\ & (\tau \rightarrow M \text{ ans}) \rightarrow M \text{ ans} \\ \mathbf{unit}_{M'} x &= \lambda \kappa. \kappa x \\ x \star_{M'} f &= \lambda \kappa. x(\lambda a. f a \kappa) \\ \text{lift}_{M\tau \rightarrow M'\tau} x &= \star_M \\ \mathbf{callcc} : ((a \rightarrow Mb) \rightarrow Ma) \rightarrow Ma \\ \mathbf{callcc} f &= \lambda \kappa. f(\lambda a. \lambda \_ . \kappa a) \kappa \end{aligned}$$
**State Monad Transformer  $\mathcal{T}_{\text{St}}$ :**
$$\begin{aligned} M'\tau &= \mathcal{T}_{\text{St}} \text{ store } M \tau = \text{store} \rightarrow M(\tau \times \text{store}) \\ \mathbf{unit}_{M'} x &= \lambda \sigma : \text{store}. \mathbf{unit}_M(x, \sigma) \\ x \star_{M'} f &= \lambda \sigma_0 : \text{store}. (x \sigma_0) \star_M (\lambda(a, \sigma_1). f a \sigma_1) \\ \text{lift}_{M\tau \rightarrow M'\tau} x &= \lambda \sigma. x \star_M \lambda y. \mathbf{unit}_M(y, \sigma) \\ \mathbf{update}(\Delta : \text{store} \rightarrow \text{store}) &= \lambda \sigma. \mathbf{unit}_M(\bullet, \Delta \sigma) \\ \mathbf{getStore} &= \lambda \sigma. \mathbf{unit}_M(\sigma, \sigma) \end{aligned}$$

Figure 3: The Identity Monad, and Environment, CPS, and State Monad Transformers

---

laws that we omit here):

$$\begin{aligned} \star_M : M\tau &\rightarrow (\tau \rightarrow M\tau') \rightarrow M\tau' \\ \mathbf{unit}_M : \tau &\rightarrow M\tau \end{aligned}$$

A value of type  $M\tau$  is called a  $\tau$ -*computation*, the idea being that it yields a value of type  $\tau$  while also performing some other computation. The  $\star_M$  operation generalizes function application in that it determines how the computations associated with monadic values are combined.  $\mathbf{unit}_M$  defines how a  $\tau$  value can be regarded as a  $\tau$ -computation; it is usually a trivial computation.

To see how monads are used, suppose we wish to define a language of integer expressions containing constants and addition. The standard definition might be:

$$\llbracket e_1 + e_2 \rrbracket = \llbracket e_1 \rrbracket + \llbracket e_2 \rrbracket$$

where  $\llbracket - \rrbracket : \text{Expression} \rightarrow \text{int}$ . However, this definition is inflexible; if expressions needed to look at a store, or could generate errors, or had some other feature not planned on, the equation would need to be changed.

Monads can provide this needed flexibility. To start, we rephrase the definition of  $\llbracket - \rrbracket$  in monadic form (using infix bind  $\star$ , as is traditional) so that  $\llbracket - \rrbracket$  has type  $\text{Expression} \rightarrow M \text{int}$ :

$$\llbracket e_1 + e_2 \rrbracket = \llbracket e_1 \rrbracket \star (\lambda x. \llbracket e_2 \rrbracket \star (\lambda y. \mathbf{add}(x, y)))$$

We must define an operation  $\mathbf{add}$  of type  $\text{int} \times \text{int} \rightarrow M \text{int}$ .

The beauty of the monadic form is that the meaning of  $\llbracket - \rrbracket$  can be reinterpreted in a variety of monads. Monadic semantics separate the *description* of a language from its *denotation*. In this sense, it is similar to *action semantics*[17] and *high-level semantics*[11].

The simplest monad is the identity monad, shown in Figure 3. Given the identity monad, we can define `add` as ordinary addition.  $\llbracket - \rrbracket$  would have type  $Expression \rightarrow int$ .

Perhaps the best known monad is the state monad, which represents the notion of a computation as something that modifies a store:

$$\begin{aligned} M_{St}\tau &= Sto \rightarrow \tau \times Sto \\ x \star f &= \lambda\sigma. \mathbf{let} (x', \sigma') = x\sigma \mathbf{in} f x' \sigma' \\ \mathbf{unit} v &= \lambda\sigma.(v, \sigma) \\ \mathbf{add}(x, y) &= \lambda\sigma.(x + y, \sigma) \end{aligned}$$

The  $\star$  operation handles the bookkeeping of “threading” the store through the computation. Now,  $\llbracket - \rrbracket$  has type  $Expression \rightarrow Sto \rightarrow int \times Sto$ . This might be an appropriate meaning for addition in an imperative language. To define operations that actually have side effects, we can define a function:

$$\begin{aligned} \mathbf{updateSto} &: (Sto \rightarrow Sto) \rightarrow M_{St}\mathbf{void} \\ &: f \mapsto \lambda\sigma.(\bullet, f\sigma) \\ \mathbf{getSto} &: M_{St}Sto \\ &: \lambda\sigma.(\sigma, \sigma) \end{aligned}$$

`updateSto` applies a function to the store and returns a useless value (we assume a degenerate type `void` having a single element, which we denote  $\bullet$ ). `getSto` returns the store.

Now, suppose a computation can cause side effects on two separate stores. One could define a new “double-state” monad  $M_{2St}$ :

$$M_{2St}\tau = Sto \times Sto \rightarrow \tau \times Sto \times Sto$$

that would thread the two states through the computation, with separate `updateSto` and `getSto` operations for each copy of  $Sto$ . One might expect to get  $M_{2St}\tau$  by applying the ordinary state monad twice. Unfortunately,  $M_{St}(M_{St}\tau)$  and  $M_{2St}\tau$  are very different types. This points to a difficulty with monads: they do not compose in this simple manner.

The key contribution of the work [7, 13] on *monad transformers* is to solve this composition problem. When applied to a monad  $M$ , a monad transformer  $\mathcal{T}$  creates a new monad  $M'$ . For example, the state monad transformer,  $\mathcal{T}_{St} store$ , is shown in Figure 3. (Here, the *store* is a type argument, which can be replaced by any value which is to be “threaded” through the computation.) Note that  $\mathcal{T}_{St} Sto \mathbf{ld}$  is identical to the state monad, but here we get a useful notion of composition:  $\mathcal{T}_{St} Sto (\mathcal{T}_{St} Sto \mathbf{ld})$  is equivalent to the two-state monad  $M_{2St}\tau$ . The state monad transformer also provides `updateSto` and `getSto` operations appropriate to the newly-created monad. When composing  $\mathcal{T}_{St} Sto$  with itself, as above, the operations on the “inner” state need to be *lifted* through the outer state monad; this is the main technical issue in [7, 13].

In our work in [8], we found it convenient to factor the state monad into two parts: the state proper and the address allocator. This was really a “staging transformation,” with the state monad representing dynamic computation and the address allocator static computation, but, as

Standard:	$\text{Dynam} = \text{Id}$	$\begin{aligned} \llbracket -t \rrbracket : \text{Dynam}(int) &= \\ \llbracket t \rrbracket \star_D \lambda i. & \\ \mathbf{unit}_D(-i) & \end{aligned}$
Implementation-oriented/Monolithic:	$\begin{aligned} \text{Dynam} &= \mathcal{T}_{\text{Env}} \text{Addr} (\mathcal{T}_{\text{St}} \text{Sto Id}) \\ \text{Addr} &= int, \text{Sto} = \text{Addr} \rightarrow int \\ \text{Thread}(i : int, a : \text{Addr}) &= \\ \text{updateSto}[a \mapsto i] \star_D \lambda \_ . \text{rdloc}(a) & \\ \text{rdloc}(a) = \text{getSto} \star_D \lambda \sigma . \mathbf{unit}_D(\sigma a) & \end{aligned}$	$\begin{aligned} \mathcal{Mono} \llbracket -t \rrbracket : \text{Dynam}(int) &= \\ \mathcal{Mono} \llbracket t \rrbracket \star_D \lambda i. & \\ \text{rdAddr} \star_D \lambda a. & \\ \text{inAddr}(a + 1) & \\ (\text{Thread}(i, a) \star_D \lambda v . \mathbf{unit}_D(-v)) & \end{aligned}$
Metacomputation:	$\begin{aligned} \text{Dynam} &= \mathcal{T}_{\text{St}} \text{Sto Id} \\ \text{Static} &= \mathcal{T}_{\text{Env}} \text{Addr Id} \end{aligned}$	$\begin{aligned} \mathcal{C} \llbracket -t \rrbracket : \text{Static}(\text{Dynam}(int)) &= \\ \text{rdAddr} \star_S \lambda a. & \\ \text{inAddr}(a + 1) & \\ (\mathcal{C} \llbracket t \rrbracket \star_S \lambda \phi_t : \text{Dynam}(int). & \\ \mathbf{unit}_S \left( \begin{array}{c} \phi_t \star_D \lambda i. \\ \text{Thread}(i, a) \star_D \lambda v. \\ \mathbf{unit}_D(-v) \end{array} \right) & \end{aligned}$

Figure 4: Negation, 3 ways

mentioned earlier, it led to significant complications. In the current paper, we are separating these parts more completely, by viewing compilation as metacomputation.

#### 4.1 A Semantics for Metacomputation

We can formalize this notion of metacomputation using monads[7, 13, 16, 24] and use the resulting framework as a basis for staging computations. Given a monad  $\mathbf{M}$ , the *computations* of type  $a$  is the type  $\mathbf{M}a$ . So given two monads  $\mathbf{M}$  and  $\mathbf{N}$ , the *metacomputations* of type  $a$  is the type  $\mathbf{M}(\mathbf{N}a)$ , because the  $\mathbf{M}$ -computation produces as a value a  $\mathbf{N}$ -computation. This definition is not superfluous; as we have noted,  $\mathbf{M} \circ \mathbf{N}$  is not generally a monad, so metacomputations are generally a different notion altogether from computations.

## 5 A Case Study in Metacomputation-based Compiler Architecture: Modular Compilation for the While Language

In this section, we present several compiler building blocks. In section 6, they will be combined to create a compiler. For the first two of these blocks, we also give monolithic versions, drawn from [8], to illustrate why metacomputation is helpful. Of particular importance to the present work, Section 5.2 presents the reusable compiler building block for control flow, which demonstrates how metacomputation-based compiler architecture solves the difficulties with the monolithic approach we outlined in Section 2.

### 5.1 Integer Expressions Compiler Building Block

Consider the standard monadic-style specification of negation[7, 13, 24] displayed in Figure 4. To use this as a compiler specification for negation, we need to make a more implementation-

---


$$\begin{aligned}
\mathcal{C}[[e_1 + e_2]] : \text{Static}(\text{Dynam } \textit{int}) = & \\
& \text{rdAddr } \star_S \lambda a. \\
& \mathcal{C}[[e_1]] \star_S \lambda \phi_{e_1}. \\
& \text{inAddr } (a + 2) \\
& \mathcal{C}[[e_2]] \star_S \lambda \phi_{e_2}. \\
& \mathbf{unit}_S \left( \begin{array}{l} \phi_{e_1} \star_D \lambda i : \textit{int}. \\ \phi_{e_2} \star_D \lambda j : \textit{int}. \\ \text{Thread}(i, a) \star_D \lambda v_1. \\ \text{Thread}(j, (a + 1)) \star_D \lambda v_2. \\ \mathbf{unit}_D(v_1 + v_2) \end{array} \right)
\end{aligned}$$

Figure 5: Specification for Addition

---

oriented version, which might be defined informally as:

$$[-t] = [[t]] \star_D \lambda i. \text{“Store } i \text{ at } a \text{ and return contents of } a\text{”} \star_D \lambda v. \mathbf{unit}_D(-v)$$

Let us assume that this is written in terms of a monad **Dynam** with bind and unit operations  $\star_D$  and  $\mathbf{unit}_D$ . Observe that this implementation-oriented definition calculates the same value as the standard definition, but it stores the intermediate value  $i$  as well. But where do addresses and storage come from? In [8], we added them to the **Dynam** monad using monad transformers[7, 13] as in the “Implementation-oriented” specification in Figure 4. In that definition, **rdAddr** reads the current top of stack address  $a$ , **inAddr** increments the top of stack, and **Thread** stores  $i$  at  $a$ . The monad (**Dynam**) is used to construct the domain containing both static and dynamic data.

In the “metacomputation”-style specification, we use two monads, **Static**, to encapsulate the static data, and **Dynam** to encapsulate the dynamic data. The meaning of the phrase is a metacomputation—the **Static** monad produces a computation of the **Dynam** monad. Clear separation of binding times is thus achieved. (In our examples, we have set the dynamic parts of the computation in a box for emphasis.)

Figure 5 displays the specification for addition, which is similar to negation. Multiplication and subtraction are defined analogously.

## 5.2 Control-flow Compiler Building Block

We now present an example where separating binding times in specifications with metacomputations has a very significant advantage over the monolithic approach. Consider the three definitions of the conditional **if-then** statement in Figure 6. The first is a dual continuation “control-flow” semantics, found commonly in compilers[2]. If  $B$  is true, then the first continuation,  $[[c]] \star_D \kappa$ , is executed, otherwise  $c$  is skipped and just  $\kappa$  is executed. A more implementation-



---

**Control-Flow:**
$$\text{Dynam} = \mathcal{T}_{\text{CPS}} \text{ void } \text{ld}$$
$$\text{Bool} = \forall \alpha. \alpha \times \alpha \rightarrow \alpha$$
$$\llbracket \text{if } b \text{ then } c \rrbracket : \text{Dynam}(\text{void}) =$$
$$\llbracket b \rrbracket \star_D \lambda B : \text{Bool}.$$
$$\text{callcc } (\lambda \kappa.$$
$$B \langle \llbracket c \rrbracket \star_D \kappa, \kappa \rangle)$$
**Implementation-oriented/Monolithic:**
$$\text{Dynam} = \mathcal{T}_{\text{CPS}} \text{ void } (\mathcal{T}_{\text{St}} \text{ Label } (\mathcal{T}_{\text{St}} \text{ Code } \text{ld}))$$
$$\text{Label} = \text{int}, \text{Code} = \text{void} \rightarrow \text{Dynam } \text{void}$$
$$\text{jump } L = \text{getCode} \star_D (\lambda \Pi : \text{Code}. \Pi L)$$
$$\text{newlabel} : \text{Dynam}(\text{Label}) =$$
$$\text{getLabel} \star_D \lambda l : \text{Label}.$$
$$\text{updateLabel}[L \mapsto L + 1] \star_D \lambda \_.$$
$$\text{unit}_D(l)$$
$$\text{Mono} \llbracket \text{if } b \text{ then } c \rrbracket : \text{Dynam}(\text{void}) =$$
$$\text{Mono} \llbracket b \rrbracket \star_D \lambda B : \text{Bool}.$$
$$\text{newlabel} \star_D \lambda L_\kappa.$$
$$\text{newlabel} \star_D \lambda L_c.$$
$$\text{callcc } (\lambda \kappa.$$
$$\text{newSegment}(L_\kappa, \kappa) \star_D \lambda \_.$$
$$\text{newSegment}(L_c, \text{Mono} \llbracket c \rrbracket \star_D (\text{jump } L_\kappa)) \star_D$$
$$B \langle \text{jump } L_c, \text{jump } L_\kappa \rangle)$$
**Metacomputation:**
$$\text{Dynam} = \mathcal{T}_{\text{CPS}} \text{ void } (\mathcal{T}_{\text{St}} \text{ Code } \text{ld}), \text{Static} = \mathcal{T}_{\text{St}} \text{ Label } \text{ld}$$
$$\text{IfThen} : \text{Dynam}(\text{Bool}) \times \text{Dynam}(\text{void}) \times \text{Label} \times \text{Label} \rightarrow \text{Dynam}(\text{void})$$
$$\text{IfThen}(\phi_B, \phi_c, L_c, L_\kappa) =$$
$$\phi_B \star_D \lambda B : \text{Bool}.$$
$$\text{callcc } (\lambda \kappa.$$
$$\text{updateCode}[L_\kappa \mapsto \kappa] \star_D \lambda \_.$$
$$\text{updateCode}[L_c \mapsto \phi_c \star_D (\text{jump } L_\kappa)] \star_D \lambda \_.$$
$$B \langle \text{jump } L_c, \text{jump } L_\kappa \rangle)$$
$$\mathcal{C} \llbracket \text{if } b \text{ then } c \rrbracket : \text{Static}(\text{Dynam } \text{void}) =$$
$$\mathcal{C} \llbracket b \rrbracket \star_S \lambda \phi_B.$$
$$\mathcal{C} \llbracket c \rrbracket \star_S \lambda \phi_c.$$
$$\text{newlabel} \star_S \lambda L_c.$$
$$\text{newlabel} \star_S \lambda L_\kappa.$$
$$\text{unit}_S (\text{IfThen}(\phi_B, \phi_c, L_c, L_\kappa))$$
Figure 6: **if-then**: 3 ways

---

oriented (informal) specification might be:

$$\llbracket \text{if } b \text{ then } c \rrbracket =$$
$$\llbracket b \rrbracket \star_D \lambda B.$$
$$\text{“get two new labels } L_c, L_\kappa \text{”} \star_D \lambda \langle L_c, L_\kappa \rangle.$$
$$\text{callcc } (\lambda \kappa.$$
$$\text{“store } \kappa \text{ at } L_\kappa, \text{ then } (\llbracket c \rrbracket \star_D (\text{“jump to } L_\kappa \text{”})) \text{ at } L_c \text{”} \star_D \lambda \_.$$
$$B \langle \text{“jump to } L_c \text{”, “jump to } L_\kappa \text{”} \rangle)$$

To formalize this specification, we use a technique from denotational semantics for modeling jumps. We introduce a continuation store, *Code*, and a label state *Label*. A jump to label *L* simply invokes the continuation stored at *L*. The second definition in Figure 6 presents an implementation-oriented specification of **if-then** in monolithic style (that is, where *Code* and *Label* are both added to *Dynam*). Again, this represents our approach in [8].

One very subtle problem remains: what is “**newSegment**”? One’s first impulse is to define it as a simple update to the *Code* store (i.e.,  $\text{updateCode}[L_\kappa \mapsto \kappa]$ ), but here is where the monolithic

---

$\mathcal{C}[\![e_1 \leq e_2]\!] : \text{Static}(\text{Dynam } Bool) =$   
 $\text{rdAddr } \star_S \lambda a.$   
 $\mathcal{C}[\![e_1]\!] \star_S \lambda \phi_{e_1}.$   
 $\text{inAddr } (a + 2)$   
 $\mathcal{C}[\![e_2]\!] \star_S \lambda \phi_{e_2}.$

$\text{unit}_S \left( \begin{array}{l} \phi_{e_1} \star_D \lambda i : int. \\ \phi_{e_2} \star_D \lambda j : int. \\ \text{Thread}(i, a) \star_D \lambda v_1. \\ \text{Thread}(j, (a + 1)) \star_D \lambda v_2. \\ \text{unit}_D \left( \begin{array}{l} \lambda \langle \kappa_T, \kappa_F \rangle. \\ ((v_1 \leq v_2) \rightarrow \kappa_T, \kappa_F) \end{array} \right) \end{array} \right)$

$\mathcal{C}[\![\text{while } b \text{ do } c]\!] : \text{Static}(\text{Dynam void}) =$   
 $\mathcal{C}[\![b]\!] \star_S \lambda \phi_B.$   
 $\mathcal{C}[\![c]\!] \star_S \lambda \phi_c.$   
 $\text{newlabel } \star_S \lambda L_{test}.$   
 $\text{newlabel } \star_S \lambda L_c.$   
 $\text{newlabel } \star_S \lambda L_\kappa.$

$\text{unit}_S \left( \begin{array}{l} \text{callcc } \lambda \kappa. \\ \phi_B \star_D \lambda B : Bool. \\ \text{updateCode}[L_\kappa \mapsto \kappa] \star_D \lambda \_ . \\ \text{updateCode}[L_c \mapsto \phi_c \star_D (\text{jump } L_{test})] \star_D \\ \text{updateCode}[L_{test} \mapsto \phi_B \star_D \lambda B. ((B \langle \text{jump } L_c, \text{jump } L_\kappa \rangle \bullet))] \star_D \\ \text{jump } L_{test} \end{array} \right)$

Figure 7: Specification for  $\leq$  and **while**

---

approach greatly complicates matters. Because the monolithic specification mixes static and dynamic computation, the continuation  $\kappa$  may contain both kinds of computation. But because it is *stored* and not *executed*,  $\kappa$  will not have access to the current label count and any other static data necessary for proper staging. Therefore, **newSegment** must explicitly pass the current label count and any other static intermediate data structure to the continuation it stores<sup>1</sup>.

The last specification in Figure 6 defines **if-then** as a metacomputation and is much simpler than the monolithic-style specification. Observe that **Dynam** does not include the *Label* store, and so the continuation  $\kappa$  now includes only dynamic computations. Therefore, there is no need to pass in the label count to  $\kappa$ , and so,  $\kappa$  may simply be stored in *Code*. **This is a central advantage of the metacomputation-based specification:** because of the separation of static and dynamic data into two monads, the complications outlined in Section 2 associated with storing command continuations in [8] (e.g., explicitly passing static data and use of a

<sup>1</sup>A full description of **newSegment** is found in [8].

---

$\begin{aligned} \text{Dynam} &= \text{ld} \\ \text{Static} &= \mathcal{T}_{\text{Env}} \text{Env} (\mathcal{T}_{\text{Env}} \text{Addr} \text{ld}) \\ \text{set } a &= \lambda v. \text{updateSto}(a \mapsto v) \\ \text{get } a &= \text{getSto} \star_D \lambda \sigma. \mathbf{unit}_D(\sigma a) \end{aligned}$	$\begin{aligned} \mathcal{C}[\mathbf{new } x \text{ in } c] : \text{Static}(\text{Dynam void}) &= \\ &\text{rdAddr} \star_S \lambda a. \\ &\mathbf{inAddr} (a + 1) \\ &\text{rdEnv} \star_S \lambda \rho. \\ &\mathbf{inEnv} (\rho[x \mapsto \mathbf{unit}_S \langle \text{set } a, \text{get } a \rangle]) \mathcal{C}[c] \end{aligned}$
---	---

Figure 8: Compiler Building Block for Block Structure

---

$\begin{aligned} \text{Dynam} &= \mathcal{T}_{\text{St}} \text{Sto} \text{ld}, \text{Static} = \mathcal{T}_{\text{Env}} \text{Env} \text{ld} \\ \mathcal{C}[c_1; c_2] : \text{Static}(\text{Dynam void}) &= \\ &\mathcal{C}[c_1] \star_S \lambda \phi_{c_1}. \\ &\mathcal{C}[c_2] \star_S \lambda \phi_{c_2}. \\ &\mathbf{unit}_S \boxed{(\phi_{c_1} \star_D \lambda \_ . \phi_{c_2})} \end{aligned}$	$\begin{aligned} \mathcal{C}[x := t] : \text{Static}(\text{Dynam void}) &= \\ &\text{rdEnv} \star_S \lambda \rho. \\ &(\rho x) \star_S \lambda \langle acc, \_ \rangle. \\ &\mathcal{C}[t] \star_S \lambda \phi_t. \\ &\mathbf{unit}_S \boxed{(\phi_t \star_D \lambda i : \text{int}. (acc i))} \end{aligned}$
---	--

Figure 9: Compiler Building Block for Imperative Features

---

*snapback* operator) are *completely* unnecessary.

Figure 7 contains the specifications for  $\leq$  and **while**, which are very similar to the specifications of addition and **if-then**, respectively, that we have seen already.

### 5.3 Block Structure Compiler Building Block

The block structure language includes **new**  $x$  **in**  $c$ , which declares a new program variable  $x$  in  $c$ . The compiler building block for this language appears in Figure 9. The static part of this specification allocates a free stack location  $a$ , and the program variable  $x$  is bound to an accepter-expresser pair[21] in the current environment  $\rho$ . In an *accepter-expresser* pair  $\langle acc, exp \rangle$ ,  $acc$  accepts an integer value and sets the value of its variable to the value, and the expresser  $exp$  simply returns the current value of the variable. **set** and **get** set and return the contents of location  $a$ , respectively.  $c$  is then compiled in the updated environment and larger stack ( $a + 1$ ).

### 5.4 Imperative Features Compiler Building Block

The simple imperative language includes assignment ( $:=$ ) and sequencing ( $;$ ). The compiler building block for this language appears in Figure 9. For sequencing, the static part of the specification compiles  $c_1$  and  $c_2$  in succession, while the dynamic (boxed) part runs them in succession. For assignment, the static part of the specification retrieves the accepter[21]  $acc$  for program variable  $x$  from the current environment  $\rho$  and compiles  $t$ , while the dynamic part calculates the value of  $t$  and passes it to  $acc$ .

## 6 Combining Compiler Building Blocks

Figure 10 illustrates the process of combining the compiler building blocks for the block structure and control-flow languages. It is important to emphasize that this is much simpler than in [8], in that there is no explicit passing of static data needed. The process is nothing more than

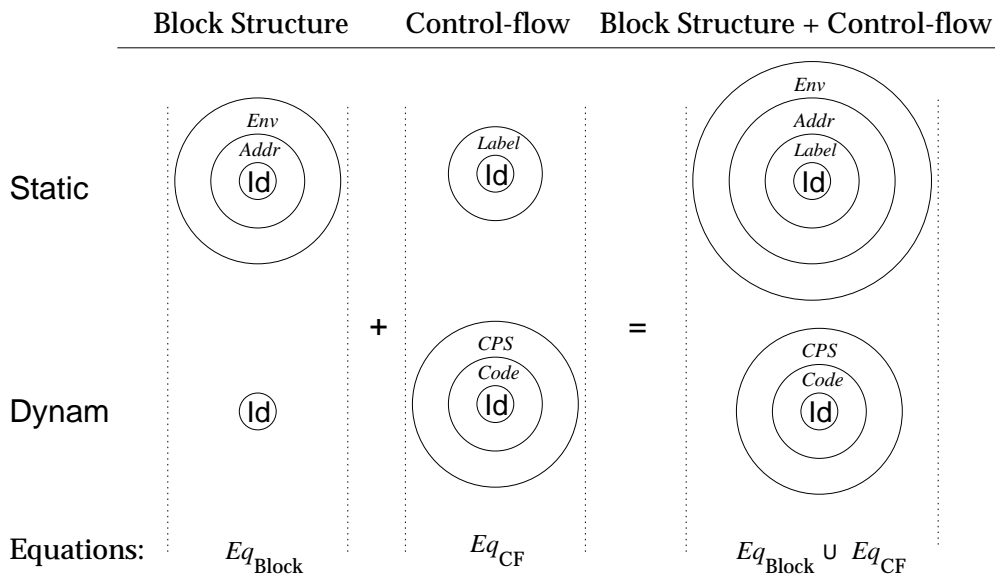


Figure 10: Combining Compiler Building Blocks

applying the appropriate monad transformers to create the **Static** and **Dynam** monads for the combined language. Recall that for the block structure language:

$$\text{Static} = \mathcal{T}_{\text{Env}} \text{Env} (\mathcal{T}_{\text{Env}} \text{Addr} \text{Id}), \text{ and } \text{Dynam} = \text{Id}$$

For the control flow language:

$$\text{Static} = \mathcal{T}_{\text{St}} \text{Label} \text{Id}, \text{ and } \text{Dynam} = \mathcal{T}_{\text{CPS}} \text{void} (\mathcal{T}_{\text{St}} \text{Code} (\mathcal{T}_{\text{St}} \text{Sto} \text{Id}))$$

To combine the compiler building blocks for these languages, one simply combines the respective monad transformers:

$$\text{Static} = \mathcal{T}_{\text{Env}} \text{Env} (\mathcal{T}_{\text{Env}} \text{Addr} (\mathcal{T}_{\text{St}} \text{Label} \text{Id})), \text{ and } \text{Dynam} = \mathcal{T}_{\text{CPS}} \text{void} (\mathcal{T}_{\text{St}} \text{Code} (\mathcal{T}_{\text{St}} \text{Sto} \text{Id}))$$

Now, the specifications for both of the smaller languages,  $Eq_{\text{Block}}$  and  $Eq_{\text{CF}}$ , apply for the “larger” **Static** and **Dynam** monads, and so we have the compiler for the combined language is specified by  $Eq_{\text{Block}} \cup Eq_{\text{CF}}$ .

Code is generated via type-directed partial evaluation[4] using the method of Danvy and Vestergaard[5]. Figure 11 contains the compiler for the while language, and an example program and its pretty-printed compiled version. All that was necessary was to combine the compiler building blocks developed in this section combined as discussed in Section 6.

## 7 Correctness

In this section, we outline an example correctness specification for a reusable compiler building block written in metacomputation style. In particular, we illustrate the advantages w.r.t. compiler correctness of metacomputation-based compiler specifications over the monolithic style specifications of [8] and also of the general usefulness of monads and monad transformers w.r.t.

---

Compiler:

Dynam =  $\mathcal{T}_{\text{CPS}}$  void ( $\mathcal{T}_{\text{St}}$  Code ( $\mathcal{T}_{\text{St}}$  Stold)), Static =  $\mathcal{T}_{\text{Env}}$  Env ( $\mathcal{T}_{\text{Env}}$  Addr ( $\mathcal{T}_{\text{St}}$  Label Id))  
Language = Expressions  $\cup$  Imperative  $\cup$  Control-flow  $\cup$  Block structure  $\cup$  Booleans  
Equations =  $Eq_{\text{Expressions}}$   $\cup$   $Eq_{\text{Imperative}}$   $\cup$   $Eq_{\text{Control-flow}}$   $\cup$   $Eq_{\text{Block structure}}$   $\cup$   $Eq_{\text{Booleans}}$

Source Code:

```
new x in new y in
  x := 5; y := 1;
  while (1 ≤ x) do
    y := y*x; x := x-1;
```

Target Code:

```
0 := 5;                2:  2 := [1];                3:  halt;
1 := 1;                3 := [0];
jump 1;                1 := [2] * [3];
                       2 := [0];
1:  2 := 1;            3 := 1;
3 := [0];              0 := [2] - [3];
BRLEQ [2] [3] 2 3;    jump 1;
```

Figure 11: Compiler for While language and example compilation

---

compiler correctness. Although lack of space makes a full exposition of metacomputation-based compiler correctness impossible here, we hope to convey the basic issues<sup>2</sup>.

The correctness of a reusable compiler building block for a source language feature is specified by comparing the *compilation semantics*  $\mathcal{C}[-]$  with the *standard semantics*  $[-]$  for that feature. Let us take as an example the conditional **if-then**. Its standard and compilation semantics are presented in Figure 6. A (slightly informal) specification of **if-then** is: If  $L_c \neq L_\kappa$  and  $L_c, L_\kappa$  are unbound in the code store, then

$$\text{IfThen}(\llbracket b \rrbracket, \llbracket c \rrbracket, L_c, L_\kappa) \star_D \lambda\_.\text{initCode} = \llbracket \text{if } b \text{ then } c \rrbracket \star_D \lambda\_.\text{initCode}$$

where  $\text{initCode} = \text{updateCode}(\lambda\_.\Pi)$  for arbitrary constant  $\Pi : \text{Code}$ . Because  $\text{IfThen}(\llbracket b \rrbracket, \llbracket c \rrbracket, L_c, L_\kappa)$  will affect the code store and  $\llbracket \text{if } b \text{ then } c \rrbracket$  will not,  $\text{IfThen}(\llbracket b \rrbracket, \llbracket c \rrbracket, L_c, L_\kappa) \neq \llbracket \text{if } b \text{ then } c \rrbracket$ . But by “masking out” the code store state on both sides with  $\text{initCode}$ —which sets the code store to constant  $\Pi$ —we require that both sides of the above equation have the same action on the value store  $Sto$ .

The above specification is easier to prove than the analogous one in monolithic style because the metacomputation-based definition in Figure 6 just stores the continuation  $\kappa$  while the monolithic-style definition manipulates  $\kappa$  as was outlined in Sections 2 and 5.2. Furthermore, here is an example of how monad transformers help with compiler correctness proofs. Although the above equation holds in  $\text{Dynam} = \mathcal{T}_{\text{CPS}} \text{Void} (\mathcal{T}_{\text{St}} \text{Label} (\mathcal{T}_{\text{St}} \text{Stold}))$ , other monad transformers could be applied to **Dynam** for the purposes of adding new source language features and the specification

---

<sup>2</sup>The interested reader may consult the first author’s forthcoming doctoral dissertation.

would still hold<sup>3</sup>. So, the use of monad transformers in this work yields a kind of proof reuse for metacomputation-based compiler correctness[14].

## 8 Conclusions and Future Work

This paper presents a modular and extensible style of language specification based on metacomputation. This style uses two monads to factor the static and dynamic parts of the specification, thereby staging the specification and achieving strong binding-time separation. Because metacomputations are defined in terms of monads, they can be constructed modularly and extensibly using monad transformers. We exploit this fact to create modular *compilers*.

Future work focuses on two areas: specifying other language constructs like objects, classes, and exceptions; and exploring the use of metacomputations in the semantics of two-level languages.

## Acknowledgements

The authors would like to thank Uday Reddy and the rest of the functional programming research group at UIUC for offering many helpful suggestions that led to significant improvements in the presentation.

## References

- [1] A. V. Aho, R. Sethi, and J. D. Ullman *Compilers: Principles, Techniques, and Tools*, Addison-Wesley, 1986.
- [2] A. Appel, *Modern Compiler Implementation in ML*, Cambridge University Press, New York, 1998.
- [3] A. Appel, *Compiling with Continuations*, Cambridge University Press, New York, 1992.
- [4] O. Danvy, “Type-Directed Partial Evaluation,” *Proceedings of the ACM Conference on the Principles of Programming Languages*, 1996.
- [5] O. Danvy and R. Vestergaard, “Semantics-Based Compiling: A Case Study in Type-Directed Partial Evaluation,” *Eighth International Symposium on Programming Language Implementation and Logic Programming*, 1996, pages 182-497.
- [6] R. Davies and F. Pfenning, “A Modal Analysis of Staged Computation,” *Proceedings of the ACM Conference on the Principles of Programming Languages*, 1996.
- [7] D. Espinosa, “Semantic Lego,” Doctoral Dissertation, Columbia University, 1995.
- [8] W. Harrison and S. Kamin, “Modular Compilers Based on Monad Transformers,” *Proceedings of the IEEE International Conference on Programming Languages*, 1998, pages 122-131.
- [9] N. D. Jones, C. K. Gomard, and P. Sestoft, *Partial Evaluation and Automatic Program Generation*, Prentice-Hall 1993.
- [10] U. Jorring and W. Scherlis, “Compilers and Staging Transformations,” *Proceedings of the ACM Conference on the Principles of Programming Languages*, 1986.
- [11] P. Lee, *Realistic Compiler Generation*, MIT Press, 1989.

---

<sup>3</sup>Given certain fairly weak conditions on the order of monad transformer application. See [12, 13, 14] for details.

- [12] S. Liang, “A Modular Semantics for Compiler Generation,” *Yale University Department of Computer Science Technical Report TR-1067*, February 1995.
- [13] S. Liang, P. Hudak, and M. Jones, Monad Transformers and Modular Interpreters. *Proceedings of the ACM Conference on the Principles of Programming Languages*, 1995.
- [14] S. Liang, “Modular Monadic Semantics and Compilation,” Doctoral Thesis, Yale University, 1997.
- [15] T. Mogensen. “Separating Binding Times in Language Specifications,” *Proceedings of the ACM Conference on Functional Programming and Computer Architecture*, pp 12-25, 1989.
- [16] E. Moggi, “Notions of Computation and Monads,” *Information and Computation* 93(1), pp. 55-92, 1991.
- [17] P. Mosses, *Action Semantics*, Cambridge University Press, 1992.
- [18] H. Nielson and F. Nielson, “Code Generation from two-level denotational metalanguages,” in *Programs as Data Objects*, Lecture Notes in Computer Science **217** (Springer, Berlin, 1986).
- [19] H. Nielson and F. Nielson, “Automatic Binding Time Analysis for a Typed  $\lambda$ -calculus,” *Science of Computer Programming* 10, 2 (April 1988), pp 139-176.
- [20] S. L. Peyton-Jones and Philip Wadler. “Imperative Functional Programming,” *Twentieth ACM Symposium on Principles of Programming Languages*, 1993.
- [21] J. Reynolds. “The Essence of Algol,” *Algorithmic Languages, Proceedings of the International Symposium on Algorithmic Languages*, pp. 345-372, 1981.
- [22] J. Reynolds, “Using Functor Categories to Generate Intermediate Code,” *Proceedings of the ACM Conference on the Principles of Programming Languages*, pages 25–36, 1995.
- [23] J. E. Stoy, *Denotational Semantics: the Scott-Strachey Approach to Programming Language Theory*, MIT Press, 1977.
- [24] P. Wadler, “The essence of functional programming,” *Proceedings of the ACM Conference on the Principles of Programming Languages*, pages 1–14, 1992.
- [25] M. Wand, “Deriving Target Code as a Representation of Continuation Semantics,” *ACM Transactions on Programming Languages and Systems*, Vol. 4, No. 3, pp. 496-517, 1982.