

**Yale University
Department of Computer Science**

A Modular Semantics for Compiler Generation ¹

Sheng Liang

YALEU/DCS/TR-1067
February 1995

¹This work was supported by the Advanced Research Project Agency and the Office of Naval Research under Arpa Order 8888, Contract N00014-92-C-0153.

Abstract

We propose a modular specification of programming language semantics to facilitate semantics-directed compiler generation. *Modular monadic semantics* allows us to define a language with a rich set of features from reusable building blocks, each implementing a single feature. We achieve modularity by using *monads* and *monad transformers* to structure denotational semantics, and still retain the power of equational reasoning. The added benefits of reasoning in monadic style include useful laws for programming language constructs, better structured proofs and more general results. To demonstrate, we present an axiomatization of environments and use it to prove the correctness of a well-known compilation technique. The monadic approach also helps to generate efficient code for various target languages with different sets of built-in features. In addition to the standard semantics, a variety of program analysis tools, such as type checkers and abstract interpreters, can also be incorporated in our framework.

1 Introduction

We propose a modular semantics which allows language designers to:

- automatically generate compilers from semantic descriptions,
- add (or remove) programming language features without causing global changes to the existing specification,
- prove the correctness of compilation strategies,
- generate code for a variety of target languages with different sets of built-in features, and
- express various aspects of programming language semantics (in addition to core semantics, for example, type systems and optimization techniques).

Our goal is similar to that of Action Semantics [32] and related approaches (by, for example, Lee [26]). It has long been recognized that traditional denotational semantics [39] is not suitable for compiler generation for a number of crucial reasons [26], among which is the lack of modularity and extensibility.

We take advantage of a new development in programming language theory, called *monads* [29], to structure denotational semantics so that it achieves a high level of modularity and extensibility. Figure 1 shows how our modular monadic semantics is organized. Language designers specify semantic modules by using a set of basic operations. The expression “ $e_1 := e_2$ ”, for example, is interpreted by “ $\text{Assign}(v_1, v_2)$ ”, where v_1 and v_2 are the results of e_1 and e_2 . Each function in the intermediate layer (such as “ Assign ”) is in turn implemented using “kernel-level” primitive operations (such as “ update ”).

While it is a well-known practice to base programming language semantics on a kernel-language, the novelty of our approach lies in how the kernel-level primitive operations are organized. In our framework, depending on how much support the upper layers need, any set of primitive operations can be put together in a modular way using an abstraction mechanism called *monad transformers* [29] [27]. Monad transformers provide the power needed to represent the abstract notion of programming language features, but still allow us to access low-level semantic details. In fact, since monad transformers are defined as higher-order functions, our monadic semantics is no more than a structured version of denotational semantics, and all conventional reasoning methods (such as β substitution) apply.

The source language we consider in this paper has a variety of features, including both call-by-name and call-by-value versions of functions and variable bindings:

$e ::=$	$n \mid e_1 + e_2$	(arithmetic operations)
	$v \mid \lambda_n v.e \mid \lambda_v v.e \mid e_1 e_2$	(cbn and cbv functions)
	$\text{let}_n v = e_1 \text{ in } e_2 \mid \text{let}_v v = e_1 \text{ in } e_2$	(let bindings)
	callcc	(first-class continuations)
	$e_1 := e_2 \mid \text{ref } e \mid \text{deref } e$	(imperative features)
	$e_1 \rightarrow e_2, e_3$	(conditionals)

Our monadic semantics is developed with automatic compiler generation in mind. We will investigate how an interpreter based on the modular monadic semantics can be turned into a compiler. Major overheads of interpreters include run-time dispatch on the abstract syntax, explicit variable lookups, dynamic type checking, and the inability to perform most optimizations. We address each of these problems in this paper. In Section 2, we will define a compositional high-level semantics for our source language which guarantees that we can unfold all recursive calls to the evaluator, and thus avoid the overhead of dispatching on the abstract syntax tree. In section 3, we show how monad laws and axioms can be used to optimize intermediate code. To demonstrate

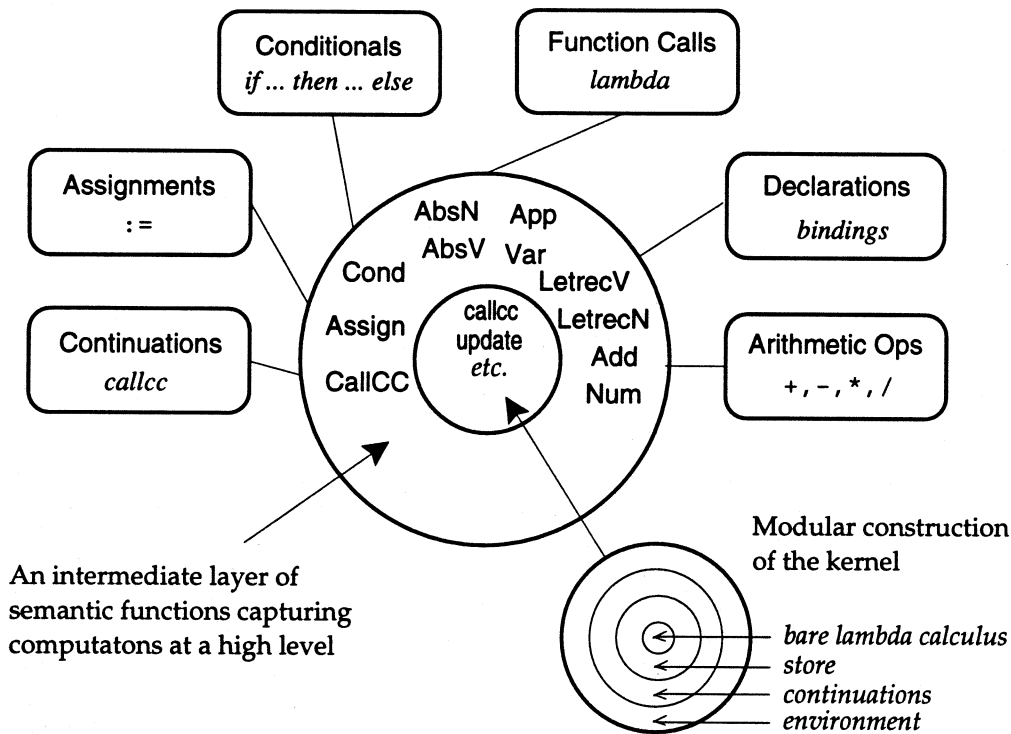


Figure 1: The organization of modular monadic semantics

the reasoning powers of monad transformers, in Section 4 we generalize Wand's [43] proof of the correctness of a well-known technique to overcome the overhead of dynamic variable lookups by transforming variables in the source language into variables in the meta-language. In section 5 we discuss how to utilize the modularity provided by monad transformers to generate efficient code for a variety of target languages. In section 6 we address how to specify non-standard semantics (for type checking and abstract interpretation) for the source language.

Although we have not actually generated a compiler from a monadic semantics, we believe that the issues addressed in this paper are important steps towards that goal.

The major contribution of this paper is to propose a modular and easy to reason about approach to semantics-directed compiler generation. Along the way we demonstrate that reasoning in monadic style enables us to better structure proofs and obtain more general results than in denotational semantics. Such insights are potentially useful in verifying large and evolving language implementations [14].

We present our results in the traditional denotational semantics style [39], augmented with a Haskell-like [19] type declaration syntax to express monads as type constructors. No prior knowledge of monads is assumed.

2 A Modular Monadic Semantics

This section uses some of the results from our earlier work on modular interpreters [27] to define a modular semantics for our source language.

2.1 A High-level Monadic Semantics

Traditional denotational semantics maps, say, a term, an environment and a continuation to an answer. In contrast, monadic semantics maps terms to computations, where the details of the environment, store, etc. are “hidden”. Specifically, our semantic evaluation function E has type:

$$E : \text{Term} \rightarrow \text{Compute Value}$$

where *Value* denotes the result of the computation. The type constructor *Compute* is called a *monad*. It abstracts away the details of a computation, exposing only the result type. We will define monads more formally later, but for now we note that *Compute* comes equipped with two basic operations:

$$\begin{aligned} \text{then} & : \text{Compute } a \rightarrow (a \rightarrow \text{Compute } b) \rightarrow \text{Compute } b \\ \text{return} & : a \rightarrow \text{Compute } a \end{aligned}$$

We usually write “then” in an infix form. Intuitively, “ c_1 then $\lambda v.c_2$ ” is a computation that first computes c_1 , binds the result to v , and then computes c_2 . “Return v ” is a trivial computation that simply returns v as result.

Monadic combinators are semantic functions returning a computation, such as:

$$\begin{aligned} \text{Num} & : \text{Int} \rightarrow \text{Compute Value} \\ \text{Add} & : (\text{Value}, \text{Value}) \rightarrow \text{Compute Value} \end{aligned}$$

where “Num(n)” just returns the number n (injected into the *Value* domain) as the result of a trivial computation, and “Add(v_1, v_2)” is a computation that adds v_1 and v_2 , and returns the result. Using these monadic combinators, we can define a (very tiny) arithmetic semantics as follows:

$$\begin{aligned} E[n] & = \text{Num}(n) \\ E[e_1 + e_2] & = E[e_1] \text{ then } \lambda v_1. E[e_2] \text{ then } \lambda v_2. \text{Add}(v_1, v_2) \end{aligned}$$

To evaluate “ $e_1 + e_2$ ”, we evaluate e_1 and e_2 in turn, and then pass the results to Add. In denotational semantics, the interpretations for arithmetic expressions are slightly different depending on whether we are passing an environment around, or whether we write in direct or continuation-passing styles. In contrast, our monadic semantics for arithmetic expressions stays the same no matter what details of computation (e.g., continuations, environments, states) are captured in the underlying monad.

The following are high-level specifications of a few other building blocks, those for functions, declarations, first-class continuations, assignments, and conditionals:

$$\begin{aligned} \text{Var} & : \text{Name} \rightarrow \text{Compute Value} \\ \text{AbsN} & : (\text{Name}, \text{Compute Value}) \rightarrow \text{Compute Value} \\ \text{AbsV} & : (\text{Name}, \text{Compute Value}) \rightarrow \text{Compute Value} \\ \text{App} & : (\text{Value}, \text{Compute Value}) \rightarrow \text{Compute Value} \end{aligned}$$

$$\begin{aligned} E[v] & = \text{Var}(v) \\ E[\lambda_n v. e] & = \text{AbsN}(v, E[e]) && \text{(call-by-name)} \\ E[\lambda_v v. e] & = \text{AbsV}(v, E[e]) && \text{(call-by-value)} \\ E[e_1 e_2] & = E[e_1] \text{ then } \lambda f. \text{App}(f, E[e_2]) \end{aligned}$$

$$\begin{aligned} \text{LetN} & : (\text{Name}, \text{Compute Value}, \text{Compute Value}) \rightarrow \text{Compute Value} \\ \text{LetV} & : (\text{Name}, \text{Compute Value}, \text{Compute Value}) \rightarrow \text{Compute Value} \end{aligned}$$

$$\begin{aligned} E[\text{let}_n x = e_1 \text{ in } e_2] & = \text{LetN}(x, E[e_1], E[e_2]) && \text{(name substitution)} \\ E[\text{let}_v x = e_1 \text{ in } e_2] & = \text{LetV}(x, E[e_1], E[e_2]) && \text{(value substitution)} \end{aligned}$$

$$\text{CallCC} : \text{Compute Value}$$

$E[\text{callcc}] = \text{CallCC}$

Assign : $(\text{Value}, \text{Value}) \rightarrow \text{Compute Value}$

Ref : $\text{Value} \rightarrow \text{Compute Value}$

Deref : $\text{Value} \rightarrow \text{Compute Value}$

$E[e_1 := e_2] = E[e_1] \text{ then } \lambda l. E[e_2] \text{ then } \lambda r. \text{Assign}(l, r)$

$E[\text{ref } e] = E[e] \text{ then } \lambda v. \text{Ref } v$

$E[\text{deref } e] = E[e] \text{ then } \lambda a. \text{Deref } a$

Cond : $(\text{Value}, \text{Compute Value}, \text{Compute Value}) \rightarrow \text{Compute Value}$

$E[e_1 \rightarrow e_2, e_3] = E[e_1] \text{ then } \lambda b. \text{Cond}(b, E[e_2], E[e_3])$

Our high-level monadic semantics bears a strong resemblance to action semantics, except that it uses only “then” and “return” to thread computations. Together with monadic combinators, these two operations are powerful enough to model various kinds of control flows (e.g., error handling, function calls and callcc) in sequential languages. Like in action semantics, we make a deliberate effort to give a very high-level view of the source language semantics.

We require that a semantics specified in terms of monadic combinators to be *compositional*: the arguments in recursive calls to E are substructures of the argument received on the left-hand side. From a theoretical point of view, it makes induction proofs on program structures possible. In practice, this guarantees that given any abstract syntax tree, we can recursively unfold all calls to the interpreter, effectively removing the run-time dispatch on abstract syntax tree.

Since most of the above semantics clauses are trivial, the intermediate layer of monadic combinators is not strictly necessary. We put them in mainly for the purpose of presentation — from now on we will focus on computations represented as functions, and stay away from syntax. In addition, realistic languages may have many syntactic constructs, but require relatively few monadic combinators. A few other benefits of using an intermediate layer of monad combinators will be mentioned during the course of the paper.

From now on we assume that the source program has been expanded into a core language composed out of “then”, “return” and monadic combinators such as “Add” and “App”. Semantics of the source language is specified by defining the monadic combinators accordingly.

2.2 The Standard Semantics

In the standard semantics, *Value* is the domain sum of basic values and functions. Functions map computations to computations:

type Fun = $\text{Compute Value} \rightarrow \text{Compute Value}$

type Value = $\text{Int} + \text{Bool} + \text{Addr} + \text{Fun} + \dots$

The standard semantics for arithmetic expressions is as follows:

$\text{Num}_s(n) = \text{return}(n \text{ in Value})$

$\text{Add}_s(v_1, v_2) = \text{if checkType}(v_1, v_2) \text{ then err "type error"}$
 $\text{else return}((v_1 | \text{Int}) + (v_2 | \text{Int}) \text{ in Value})$

We use a *primitive monadic combinator* (a semantic function directly supported by the underlying monad *Compute*):

err : $\text{String} \rightarrow \text{Compute } a$

to report type errors. For clarity, from now on we will omit domain injection/projection and type checking.

Function abstractions and applications need access to an environment *Env* which maps variable names to computations, and two more primitive monadic combinators which retrieve the current environment and perform a computation in a given environment, respectively:

```

type Env = Name → Compute Value
rdEnv    : Compute Env
inEnv    : Env → Compute Value → Compute Value

```

The standard semantics for functions is as follows:

```

Vars      = rdEnv then λρ.ρ s
AbsNs(s, c) = rdEnv then λρ. return(λc1.inEnv (ρ[c1/s]) c)
AbsVs(s, c) = rdEnv then λρ. return(λc1.c1 then λv.inEnv (ρ[return v/s]) c)
Apps(f, c)  = rdEnv then λρ.f(inEnv ρ c)

```

The difference between call-by-value and call-by-name is clear: the former reduces the argument before evaluating the function body, whereas the latter does not. In a function application, the argument is packaged up with the current environment to form a closure.

In this paper, we only consider non-recursive let bindings, which can be translated into applications of lambda abstractions:

```

LetNs(s, c1, c2) = Apps(AbsNs(s, c2), c1)
LetVs(s, c1, c2) = Apps(AbsVs(s, c2), c1)

```

Here we see another benefit of using an intermediate layer of monadic combinators such as *App*: Monadic combinators can be defined in terms of each other, while the high-level semantics remains compositional.

To simplify the presentation somewhat, we assume that *Assign_s*, *Deref_s*, and *Ref_s* can be defined using the primitive monad combinator:

```

update : (Store → Store) → Compute Store

```

for some suitably chosen *Store*. We can read the store by passing update the identity function, and change the store by passing it a state transformer. Although update returns the entire state, properly defined *Assign_s*, *Deref_s*, and *Ref_s* can guarantee that the store is never duplicated (see, for example, [36]).

CallCC simply returns a function expecting another function as an argument, to which the current continuation will be passed.

```

callcc : ((Value → Compute Value) → Compute Value) → Compute Value

```

```

CallCCs = return(λf.f then λf'.callcc(λk.f'(λa.a then k)))

```

Conditionals are supported by the meta-language:

```

Conds(v, c1, c2) = if v then c1 else c2

```

2.3 Constructing the Compute Monad

It is clear that monad *Compute* needs to support the following primitive monad combinators:

```

err      : String → Compute a
rdEnv    : Compute Env
inEnv    : Env → Compute Value → Compute Value
update   : (Store → Store) → Compute Store
callcc   : ((Value → Compute Value) → Compute Value) → Compute Value

```

If we follow the traditional denotational semantics approach, now is the time to set up domains and implement the above functions. The major drawback of such monolithic approach is that if we add some source language feature later on, all the functions may have to be redefined.

For the sake of modularity, we start from a simple monad and add more and more features. The simplest monad of all is the identity monad. All it captures is function application:

```

type Id a = a

returnId x = x
c thenId f = f c

```

A *monad transformer* takes a monad, and returns a new monad with added features. For example, “*StateT s*” adds a state *s* to any monad *m*:

```

type StateT s m a = s → m (s, a)

return(StateT s m) x = λs. returnm (s, x)
c then(StateT s m) k = λs0. c s0 thenm λ(s1, a). k a s1

```

To see how monad transformers work, let us apply *StateT* to the identity monad *Id*:

```

type StateT s Id a = s → Id (s, a)
                    = s → (s, a)

return(StateT s Id) x = λs. returnId (s, x)
                    = λs. (s, x)
c then(StateT s Id) k = λs0. c s0 thenId λ(s1, a). k a s1
                    = λs0. let (s1, a) = m s0 in k a s1

```

Note that “*StateT s Id*” is the standard state monad found, for example, in Wadler’s work [40].

To make the newly introduced state accessible, “*StateT s*” introduces *update* on *s* which applies *f* to the state, and returns the old state:

```

update : (s → s) → StateT s m a
update f = λs. return(f s, s)

```

Figure 2 gives the definitions of several other monad transformers, including those for errors (*ErrorT*), continuations (*ContT*) and environments (*EnvT*). Now we can construct *Compute* by applying a series of monad transformers to the base monad *Id*:

```

type Compute a = EnvT Env (ContT Answer (StateT Store (ErrorT Id))) a

```

Env, *Store* and *Answer* are the type of the environment, store and answer, respectively.

One issue remains to be addressed. The *update* function introduced by *StateT* does not work on *Compute*, which contains features added later by other monad transformers. In general, this is the problem of *lifting* operations through monad transformers. Figure 3 gives a brief summary of useful liftings (See [27] for a detailed description.) For example, in the *Compute* monad above, “*update f*” is:

$$\lambda s. \text{Ok } (f \ s, \ s)$$

when first introduced by *StateT*. After *Compute* is finally constructed, “*update f*” becomes:

$$\lambda \rho. \lambda k. \lambda s. k \ s \ (f \ s)$$

In summary, monad transformers allow us to easily construct monads with a certain set of primitive monadic combinators, defined as higher-order functions.

<p>Continuation:</p> <pre> type ContT ans m a = (a → m ans) → m ans return_(ContT ans m) a = λk.k a c then_(ContT ans m) f = λk.c(λa.f a k) callcc f = λk.f(λa.λk'.k a)k </pre> <p>Errors:</p> <pre> type Error a = Ok a Error String type ErrorT m a = m (Error a) return_(ErrorT m) a = return_m .Ok c then_(ErrorT m) k = c then_m λa. case a of (Ok x) → k x (Error msg) → return_m (Error msg) </pre> <p>err = return_m .Error</p>	<p>Environment:</p> <pre> type EnvT e m a = e → m a return_(EnvT e m) a = λρ.return_m a c then_(EnvT e m) k = λρ.c ρ then_m λa.k a ρ rdEnv = λρ.return_m ρ inEnv ρ c = λρ'.c ρ </pre>
---	--

Figure 2: Monad transformers

Every monad transformer t has a function:

$$\text{lift}_{(t\ m)} : m\ a \rightarrow t\ m\ a$$

which embeds a computation in monad m into “ $t\ m$ ”. Functions `err`, `update` and `rdEnv` are easily lifted using `lift`:

$$\begin{aligned} \text{err}_{(t\ m)} &= \text{lift}_{(t\ m)} \cdot \text{err}_m \\ \text{update}_{(t\ m)} &= \text{lift}_{(t\ m)} \cdot \text{update}_m \\ \text{rdEnv}_{(t\ m)} &= \text{lift}_{(t\ m)} \text{rdEnv}_m \end{aligned}$$

Some liftings of `callcc`, `inEnv` and the definitions of `lift` for each monad transformer are listed in the following table:

$t\ m$	$\text{callcc}_{(t\ m)}\ f$	$\text{inEnv}_{(t\ m)}\ \rho\ c$	$\text{lift}_{(t\ m)}\ c$
<i>EnvT e m</i>	$\lambda\rho.\text{callcc}_m(\lambda k.f(\lambda a.\lambda\rho'.k a)\rho)$	$\lambda\rho'.\text{inEnv}_m\ \rho\ (c\rho')$	$\lambda\rho.c$
<i>ContT ans m</i>			$\lambda k.c\ \text{then}_m\ k$
<i>StateT s m</i>	$\lambda s_0.\text{callcc}_m(\lambda k.f(\lambda a.\lambda s_1.k(s_1, a))s_0)$	$\lambda s.\text{inEnv}_m\ \rho\ (cs)$	$\lambda s.c\ \text{then}_m\ \lambda x.\text{return}_m(s, x)$
<i>ErrorT m</i>	$\text{callcc}_m(\lambda k.f(\lambda a.k(\text{Ok } a)))$	$\text{inEnv}_m\ \rho\ c$	$\text{map}_m\ \text{Ok}$

Figure 3: Liftings

3 Using Monad Laws to Transform Programs

Following the monadic semantics presented in the previous section, by unfolding all calls to semantic function E , we can transform source-level programs into monadic-style code. For example, $((\lambda x. x + 1) 2)_v$ is transformed to:

```
rdEnv then λρ.
return(λc.inEnv (ρ[c/"x"])
      (rdEnv then λρ.
        ρ"x" then λv1.
        return 1 then λv2.
        return(v1 + v2))) then λf.
return2 then λv.
f(return v)
```

In this section we formally introduce monads and their laws, and show how to use the laws to simplify the above program.

3.1 Monads and Monad Laws

Definition 3.1 A monad M is a type constructor, together with two operations:

```
then  : M a → (a → M b) → M b
return : a → M a
```

satisfying the following laws [40]:

$$\begin{aligned} (\text{return } a) \text{ then } k &= k a && \text{(left unit)} \\ c \text{ then return} &= c && \text{(right unit)} \\ c_1 \text{ then } \lambda v_1. (c_2 \text{ then } \lambda v_2. c_3) &= (c_1 \text{ then } \lambda v_1. c_2) \text{ then } \lambda v_2. c_3 && \text{(associativity)} \end{aligned}$$

Intuitively, the (left and right) unit laws say that trivial computations can be skipped in certain contexts; and the associativity law captures the very basic property of sequencing, one that we usually take for granted in programming languages.

We can verify, by equational reasoning, for example, that return_{id} and then_{id} satisfy the above laws, and EnvT , ContT etc. indeed transform monads to monads.

Monad laws are important for reasoning about monadic style programs. For example, our compiler translates the expression " $2 + 3$ " to:

```
return2 then λv1. return3 then λv2. return(v1 + v2)
```

We can apply the left unit law twice, and reduce the above to:

```
return(2 + 3)
```

which can of course be further optimized to "return 5".

Every application of a monad law usually corresponds to a number of β reductions. Monad laws allow us to perform β reductions at the right places, and avoid those corresponding to actual computations in the source program, which in turn may lead to non-termination.

Without knowledge about the environment-handling operations inEnv and rdEnv , however, monad laws alone can only simplify the example in the beginning of the section to:

```

rdEnv then λρ.
(λc.inEnv (ρ[c/"x"])
  (rdEnv then λρ.
    ρ"x" then λv.
      return(v + 1)))
(return 2))

```

To further simplify the above program, we need to look at the laws environment-related operations should satisfy.

3.2 Environment Axioms

We axiomatize the environment-manipulating functions as follows:

Definition 3.2 Monad M is an **environment monad** if it has two operations: `rdEnv` and `inEnv`, which satisfy the following axioms:

$$\begin{array}{lll}
(\text{inEnv } \rho) \cdot \text{return} & = & \text{return} & \text{(unit)} \\
\text{inEnv } \rho (c_1 \text{ then } \lambda v.c_2) & = & \text{inEnv } \rho c_1 \text{ then } \lambda v.\text{inEnv } \rho c_2 & \text{(distribution)} \\
\text{inEnv } \rho \text{ rdEnv} & = & \text{return } \rho & \text{(cancellation)} \\
\text{inEnv } \rho (\text{inEnv } \rho' e) & = & \text{inEnv } \rho' e & \text{(overriding)}
\end{array}$$

Intuitively, a trivial computation cannot depend on the environment (the unit law); the environment stays the same across a sequence of computations (the distribution law); the environment does not change between a set and a read if there are no intervening computations (the cancellation law); and an inner environment supercedes an outer one (the overriding law).

As with the monad laws, the environment axioms can be verified by equational reasoning.

Proposition 3.1 *The monads supporting `rdEnv` and `inEnv` constructed using monad transformers `ErrorT`, `EnvT`, `StateT` and `ContT` are environment monads.*

Equipped with the environment axioms, we can further transform the example monadic code to:

```

rdEnv then λρ.
(λc.c then λv. return(v + 1))(return 2)

```

Note that explicit environment accesses have disappeared. Instead, the meta-language environment is directly used to support function calls. This is exactly what many partial evaluators achieve when they transform interpreters to compilers.

4 Using Monad Laws to Reason about Computations

We successfully transformed away the explicit environment in the above example, but can we do the same for arbitrary source programs? It turns out that we can indeed prove such a general result by using monad laws and environment axioms. We follow Wand [43], define a "natural semantics" which translates source language variables to lexical variables in the meta-language, and prove that it is equivalent to the standard semantics.

4.1 A Natural Semantics

We adopt Wand's definition of a *natural semantics* (different from Kahn's [5] notion) to our functional sub-language. For any source language variable name s , we assume there is a corresponding variable name v_s in the meta-language.

Definition 4.1 A *natural semantics* uses the environment of the meta-language for variables in the source language, and is given as follows:

$$\begin{aligned} \text{Var}_n(s) &= v_s \\ \text{AbsN}_n(s, c) &= \text{return}(\lambda v_s. c) \\ \text{AbsV}_n(s, c) &= \text{return}(\lambda v'_s. v'_s \text{ then } \lambda v''_s. \text{return}(\text{return } v''_s)) \text{ then } \lambda v_s. c \\ \text{App}_n(f, c) &= f c \end{aligned}$$

Other monadic combinators, such as CallCC, Assign etc., do not explicitly deal with the environment, and have the same natural semantics as standard semantics.

4.2 Correspondence between Natural and Standard Semantics

The next theorem, a variation of Wand's [43], guarantees that it is safe to implement function calls in the source language using the meta-language environment.

Theorem 4.1 *Let c be a program composed out of then, return, and monad combinators. Let c_s be its standard semantics in an environment monad, c_n be its natural semantics in the same monad,¹ and ρ be the mapping from the source language variable names s to v_s , we have:*

$$\text{inEnv } \rho \ c_s = c_n$$

To emphasize the modularity provided in our framework, we first prove the theorem for the functional sub-language, and then extend the result to the complete language.

4.2.1 Proof for the Functional Sub-language

We can establish the theorem for the functional sub-language by induction on the structure of *computations* composed out of return, then and the monad combinators for variables, lambda abstractions, and function applications. The full proof is given in the Appendix. The basic technique is the same as Wand's, except that in addition to the basic rules of lambda calculus (e.g., β reduction), we also use monad laws and environment axioms.

The proof is possible because both the source language and meta language are lexically scoped. If the source language contained dynamically scoped functions:

$$\text{AbsD}_s(s, c) = \text{return}(\lambda c_1. \text{rdEnv then } \lambda \rho. \text{inEnv } (\rho[c_1/s]) \ c)$$

where the caller-site environment is used within the function body, the theorem would fail to hold.

¹This means that in natural semantics, we are still implicitly passing around an environment, even though it is never used. Thus the theorem as stated does not strictly correspond to Wand's result [43]. Fortunately, the *naturality of liftings* (see our earlier work [27] for details) guarantees that adding and removing a feature does not affect computations which do not use that particular feature. Therefore the theorem still holds if we remove the explicit environment support from the underlying monad in natural semantics. (The next section addresses this in more detail.)

4.2.2 Extension to the Complete Language

Consider another monad combinator `CallCC`. Since in proving the theorem we only used the axioms of environment monads, none of the cases already analyzed need to be proved again. We only have to verify that:

- the monad supporting continuations is still an environment monad, and
- the induction hypothesis holds for `CallCC`.

The former is stated in Proposition 3.1, and can be proved once and for all as we come up with monad transformers. The latter can be easily proved: `CallCC` does not explicitly deal with the environment, and has exactly the same natural semantics as the standard semantics. In addition, it is a trivial computation (see the definition in the last section). Thus the induction hypothesis holds following the unit axiom of environment monads.

Similarly we can extend the theorem to cover other features such as `Assign`.

In denotational semantics, adding a feature may change the structure of the entire semantics, forcing us to redo the induction for every case of abstract syntax. Indeed, Wand [43] pointed out that he could change the semantics into continuation-based, and prove the theorem, but only by modifying the proofs accordingly.

In denotational semantics, a computation is captured as a piece of syntax tree coupled with an environment, a store etc. On the other hand, we view computations as abstract entities with a set of equations. Therefore, like *Semantic Algebras* [31] in action semantics, monads provide an *axiomatic view* to denotational semantics.

5 Targeting Monadic Code

In general, it is more efficient to use target language built-in features instead of monadic combinators defined as higher-order functions. We have seen how the explicit environment can be “absorbed” into the meta-language. This section addresses the question of whether we can do the same for other features, such as the store and continuation.

We can view a target language as having a built-in monad supporting a set of primitive monadic combinators. For example, the following table lists the correspondence between certain ML constructs and primitive monadic combinators:

primitive monadic operators	ML construct
<code>return x</code>	<code>x</code>
<code>c₁ then λx.c₂</code>	<code>let val x = c₁ in c₂ end</code>
<code>update*</code>	<code>ref, !, :=</code>
<code>callcc</code>	<code>callcc</code>
<code>err</code>	<code>raise Err</code>
...	

* ML reference cells support single-threaded states only.

It is easy to verify that the monad laws are satisfied in the above context. For example, the ML `let` construct is associative (assuming no unwanted name capturings occur):

$$\boxed{\begin{array}{l} \text{let val } v_2 = \text{let val } v_1 = c_1 \\ \quad \quad \quad \text{in } c_2 \\ \quad \quad \quad \text{end} \\ \text{in } c_3 \\ \text{end} \end{array}} = \boxed{\begin{array}{l} \text{let val } v_1 = c_1 \\ \text{in let val } v_2 = c_2 \\ \quad \quad \quad \text{in } c_3 \\ \quad \quad \quad \text{end} \\ \text{end} \end{array}}$$

Recall (in Section 2) that the *Compute* monad is constructed as:

```
type Compute a = EnvT Env (ContT Answer (StateT Store (ErrorT Id))) a
```

Now we substitute the base monad *Id* with the built-in ML monad (call it M_{ML}):

```
type Compute' a = EnvT Env (ContT Answer (StateT Store (ErrorT MML))) a
```

Note that *Compute'* supports two sets of continuation, state and error handling functions. The monadic code can choose to use the ML built-in ones instead of those implemented as higher-order functions. In addition, all liftings we construct satisfy an important property (called the Naturality of Liftings [29] [27]): adding or deleting a monad transformer does not change the result of programs which do not use its operations. Since none of the monad transformers in *Compute'* is used anymore, it suffices to run the target program on *Compute''*:

```
type Compute'' a = MML a
```

which directly utilizes the more efficient ML built-in features.

The above transformation is possible because ML has a strictly richer set of features than our source language. If the source language requires a non-updatable version of state (for example, for the purpose of debugging), the corresponding state monad transformer will remain, and ensure the state is threaded correctly through all computations. If we instead target our source language to C, both the environment and continuation transformers have to be kept.

Therefore by using a monad with a set of primitive monadic combinators, we can expose the features embedded in the target language. It then becomes clear what is directly supported in the target language, and what needs to be compiled explicitly.

The above process seems trivial, but would have been impossible had we been working with traditional denotational semantics. Various features clutter up and make it hard to determine whether it is safe to remove certain interpretation overhead, and how to achieve that.

Earlier work [25] [11] [27] has shown that the order of monad transformers (in particular, some cases involving *ContT*) has an impact on the resulting semantics. In practice, we need to make sure when we discard the monad transformers, that the resulting change of ordering does not have unwanted effects on semantics.

6 Non-standard Monadic Semantics

In this section we show that by using a different set of monads, and redefining the monadic combinators, we can derive non-standard semantics corresponding to a wide range of language processing tools, in particular, a Hindley-Milner type checker [8], and a higher-order operation count analyzer based on abstract interpretation [7].

Specifying language processing tools as non-standard semantics of monadic combinators has two advantages:

1. We can use equational reasoning to verify the correctness of such tools, and investigate their relationships with the standard semantics.
2. We can share the high-level evaluator between the standard and non-standard semantics.

6.1 A Type Checker

The algorithm is based on a monadic type checker written by Hammond [16]. The difference is that our monad includes the type environment, and is constructed using monad transformers.

The type-checking monad and the value type are:

```

type TyVar      = Int
type Compute a  = EnvT TEnv (StateT Subst (StateT TyVar (ErrorT Id))) a
type Value      = TyVar + (Value → Value)

```

where *TEnv*, *Subst*, and *TyVar* denote the type environment, substitutions, and the counter for generating fresh type variables, respectively. Values here denote type signatures, which, for simplicity, are either a type variable (represented using an integer) or a function type.

We also need the following helper functions:

```

type Maybe a    = Just a + Nothing
lookup          : String → TEnv → Maybe (Value, N + G)
newVar         : Compute Value
subst          : Value → Compute Value
instType       : Value → Compute Value
unify          : Value → Value → Compute Value

```

The environment maps variables to a pair consisting of a type and a flag indicating whether the variable is let-bound or lambda-bound. A type-checking semantics for monadic combinators can be given as (only the interesting cases are listed):

```

Vart(s)        = rdEnv then λρ.
                  case lookup s ρ of
                    Just (t, G) → instType t then λt'.subst t'
                    Just (t, N) → subst t
                    Nothing    → err("undefined var: " ++ s)
AbsNt(s, c)     = rdEnv then λρ.
                  newVar then λt.
                  inEnv (ρ[(t, N)/s] c) then λtc.
                  subst t then λt',
                  return(t' → tc)
Appt(t1, c)    = c then λt2.
                  newVar then λt.
                  unify t1 (t2 → t) then λx.
                  subst t
LetNt(s, c1, c2) = c1 then λt1.
                  subst t1 then λt'1.
                  inEnv (ρ[(t'1, G)/s])c2

```

Note that lambda-bound type variables are marked with "N", whereas let-bound ones are marked with "G", meaning that they can be instantiated later. Unify calls *err* if it fails.

6.2 An Abstract Interpreter

Young [46] has built a parameterized abstract interpreter which does a variety of analyses by redefining primitive functions. It is much harder, however, to reuse part of the *standard* interpreter for abstract interpretation, due to the dramatic differences in evaluation strategies. Hall [15] has parameterized a standard evaluator for abstract interpretation, but she dealt with a very simple language, which, for example, did not need an environment.

In our framework, various abstract interpretation algorithms can be specified by giving a non-standard semantics to monadic combinators. As an example, we derive a simple higher-order operation-count analysis — which is useful for a variety of compiler optimizations — by specifying the monad and value types as follows:

type *Compute a* = *EnvT OpEnv Id a*
type *Value* = *(Int, Value → Value)*

Following Hudak and Young [20], we use a pair to capture the first-order and higher-order properties of a given expression. For the simple analysis we describe below, which does not deal with recursion, it is enough to use a simple environment monad mapping variables to their operation counts. If recursions exist, however, it will be necessary to introduce another environment to cache pending function calls [45], in order to guarantee termination.

$\text{Add}_c(v_1, v_2)$ = $\text{return}(1, \perp)$
 $\text{Var}_c(s)$ = rdEnv then $\lambda\rho. \rho s$
 $\text{AbsN}_c(s, c)$ = rdEnv then $\lambda\rho. \text{return}(1, \lambda c_1. \text{inEnv}(\rho[c_1/s]) c)$
 $\text{AbsV}_c(s, c)$ = rdEnv then $\lambda\rho.$
 $\text{return}(1, \lambda c_1. c_1$ then $\lambda v.$
 $\text{inEnv}(\rho[\text{return}(0, \text{snd } v)/s]) c$ then $\lambda a.$
 $\text{return}(\text{fst } v + \text{fst } a, \text{snd } a))$
 $\text{App}_c(f, c)$ = rdEnv then $\lambda\rho. (\text{snd } f)(\text{inEnv } \rho c)$ then $\lambda a.$
 $\text{return}(\text{fst } f + \text{fst } a + 1, \text{snd } a)$
 $\text{Cond}_c(v, c_1, c_2)$ = c_1 then $\lambda i_1. c_2$ then $\lambda i_2.$
 $\text{return}(\text{fst } v + \text{imax}(\text{fst } i_1, \text{fst } i_2) + 1, \text{fmax}(\text{snd } i_1, \text{snd } i_2))$

Operation counts for all basic operations (such as addition and the creation of a closure) are arbitrarily set to 1. Note that the operations in an argument are counted every time it is evaluated in the body of a call-by-name function, but only once upon entering a call-by-value function. *Imax* simply returns the larger of two integers, whereas *fmax* compares two integer-returning functions, and has to wait for more arguments before it decides.

7 Related work

Due to the problems with traditional denotational semantics, earlier efforts [30] [35] [42] in semantics-directed compiler generation did not produce efficient compilers.

Mosses's *Action Semantics* [32] allows modular specification of programming language semantics, from which efficient compilers can be generated. In action semantics, source language constructs are translated into *actions*, built from a set of primitives and combinators. While action semantics is easy to construct, extend, understand and implement, we note the following comments made by Mosses ([32], page 5):

Although the foundations of action semantics are firm enough, the *theory* for reasoning about actions (and hence about programs) is still rather weak, and needs further development. This situation is in marked contrast to that of denotational semantics, where the theory is strong, but severe pragmatic difficulties hinder its application to realistic programming languages.

Action semantics provided much of the inspiration for our work, which essentially attempts to formulate actions in a denotational semantics framework. Monad transformers roughly correspond to *facets* in action semantics, although issues such as concurrency are beyond the power of our denotational semantics-based approach.

Action semantics [4] [34] [33] and a related approach by Lee [26] have been successfully used to generate efficient compilers. Lee has given a denotational semantics to actions, and Even and Schmidt have presented a categorical model [12] and a type inferencer [13] for action semantics.

Doh and Schmidt [10] have presented a way to assign each action in action semantics specific "analysis functions," such as a type checker and a binding-time analyzer. Such analysis tools can

thus be generated automatically together with the compiler. The differences between our approach and Doh and Schmidt's are:

1. They redefine the operators for threading actions. We do not have to, since "return" and "then" come with any monad.
2. They showed how the type-checking rules for the source language can be automatically derived from the type-checking rules given on actions. Instead, we directly specify the type-checking algorithm. Their approach is clever, and handles subtyping, but does not deal with let-bound polymorphism.

Moggi [29] first used monads and monad transformers to structure denotational semantics. Wadler [41] [40] popularized Moggi's ideas in the functional programming community by using monads to structure functional programs, in particular, interpreters. Our work on monad-based modular interpreters [27] follows a series of earlier attempts by Steele [38], Jones and Duponcheel [21], and Espinosa [11].

Moggi [29] also raised (although did not address in detail) the issue of reasoning in a monadic framework. Wadler [41] listed the laws a state monad should satisfy. Hudak [18] suggested a more general framework — mutable abstract data types (MADTs) — to reason about states.

One application of partial evaluation [22] is to automatically generate compilers from interpreters [44] [2] [6] [28]. A partial evaluator has been successfully applied to an action interpreter [3], and similar results can be achieved with monadic interpreters [9].

Staging transformations, first proposed by Jørring and Scherlis [23], are a class of general program transformation techniques for separating a given computation into stages. Monadic transformers make computational stages somewhat more explicit by separating compile-time features, such as the environment, from run-time features.

Several researchers, including Kelsey and Hudak [24], Appel and Jim [1], and others, have built efficient compilers for higher-order languages by transforming the source language into continuation-passing style (CPS). The suitability of a monadic form as an intermediate form has been observed by many researchers (including, for example, Sabry and Felleisen [37] and Hatcliff and Danvy [17]). We will continue to explore along this direction in order to generate machine-level code from a monadic intermediate form.

8 Conclusions

We have shown that the monadic framework provides good support for high-level extensible specifications, program transformations, and reasoning about computations. Monadic-style proofs are better structured and easier to extend. The modular monadic semantics allows us to have an axiomatized formulation of well-known programming languages features such as environments. Our approach also facilitates generating code for various target languages and constructing program analysis tools.

Overall, we believe that modular monadic semantics is particularly suitable for compiler generation.

Acknowledgement

I would like to thank Paul Hudak for support and advice, Rajiv Mirani, Zhong Shao, Satish Pai, Dan Rabin and the PEPM'95 anonymous referees for helpful comments, and Ross Paterson, Tim Sheard and John Launchbury for ideas and discussions.