Modular Denotational Semantics for Compiler Construction *

Sheng Liang   Paul Hudak

Yale University, Department of Computer Science
New Haven, CT 06520-8285
{liang-sheng, hudak}@cs.yale.edu

Abstract. We show the benefits of applying *modular monadic semantics* to compiler construction. Modular monadic semantics allows us to define a language with a rich set of features from reusable building blocks, and use program transformation and equational reasoning to improve code. Compared to denotational semantics, reasoning in monadic style offers the added benefits of highly modularized proofs and more widely applicable 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 facilitates generating code in various target languages with different sets of built-in features.

1 Introduction

We propose a modular semantics which allows language designers to add (or remove) programming language features without causing global changes to the existing specification, derive a compilation scheme from semantic descriptions, prove the correctness of program transformation and compilation strategies, and generate code in various target languages with different built-in features.

Our goals are similar to that of Action Semantics [21] and related efforts by, for example, Wand [27], Lee [16], Appel & Jim [1], and Kelsey & Hudak [14]. None of the existing approaches are completely satisfactory in achieving the above goals. For example, it has long been recognized that traditional denotational semantics [24] is not suitable for compiler generation for a number of crucial reasons [16], among which is the lack of modularity and extensibility.

We take advantage of a new development in programming language theory, a monadic approach [19] to structured denotational semantics, that achieves a high level of modularity and extensibility. 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:

$$
\begin{array}{ll}
e ::= n \mid e_1 + e_2 & \text{(arithmetic operations)} \\
\quad \mid v \mid \lambda v.e \mid (e_1\ e_2)_n \mid (e_1\ e_2)_v & \text{(cbn and cbv functions)} \\
\quad \mid \text{callcc} & \text{(first-class continuations)} \\
\quad \mid e_1 := e_2 \mid \text{ref } e \mid \text{deref } e & \text{(imperative features)}
\end{array}
$$

Assignments
:=

Function Calls
lambda

rdEnv inEnv
callcc
update
err

Continuations
callcc

Arithmetic Ops
+ , - , * , /

pure lambda calculus
error reporting
store
continuations
environment
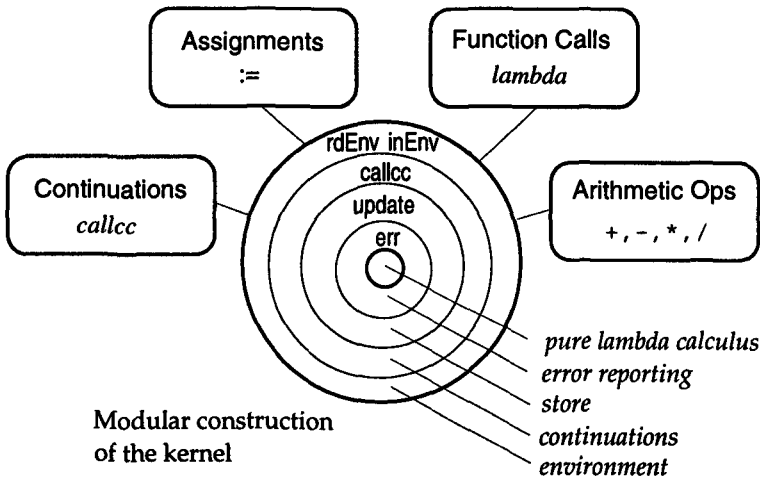
Modular construction
of the kernel

Fig. 1. The organization of modular monadic semantics

Figure 1 shows how our modular monadic semantics is organized. Language designers specify semantic modules by using a set of "kernel-level" operations. The expression "$e_1 := e_2$", for example, is interpreted by the low-level primitive operation "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* [19] [17]. 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.

We will investigate how an interpreter based on the modular monadic semantics can be turned into a compiler. 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 the reasoning powers of monad transformers, in Section 4 we generalize Wand's [28] 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 target different languages.

The contributions of this paper are:

- proposing a monad-based modular approach to semantics directed compiler generation,
- applying monad laws to program transformation,
- presenting a monadic style axiomatization of environments,
- demonstrating that reasoning in monadic style enables us to better structure proofs and obtain more general results than in denotational semantics, and
- taking advantage of monad transformer properties (for example, naturality of liftings) to utilize target language features.

We present our results in the traditional denotational semantics style [24], augmented with a Haskell-like [10] type declaration syntax to express monads as type constructors. We use the denotation semantics notation because it is more succinct than a real programming language such as Haskell. No prior knowledge of monads is assumed.

2   A Modular Monadic Semantics

In this section we use some of the results from our earlier work on modular interpreters [17] 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:

$E : Term \rightarrow Compute\ Value$

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:

then   : $Compute\ a \rightarrow (a \rightarrow Compute\ b) \rightarrow Compute\ b$
return : $a \rightarrow Compute\ a$

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.

In the standard semantics, *Value* is the domain sum of basic values and functions. Functions map computations to computations:[2]

type *Fun*   = $Compute\ Value \rightarrow Compute\ Value$
type *Value* = $Int + Bool + Addr + Fun + \ldots$

---

[2] This generality allows us to model both call-by-name and call-by-value.

The standard semantics for arithmetic expressions is as follows:

$$E[\![n]\!] \quad\quad = \text{return}(n \text{ in } Value)$$
$$E[\![e_1 + e_2]\!] = E[\![e_1]\!] \text{ then } \lambda v_1.$$
$$\quad\quad\quad E[\![e_2]\!] \text{ then } \lambda v_2.$$
$$\quad\quad\quad \text{if checkType}(v_1, v_2) \text{ then err } \textit{"type error"}$$
$$\quad\quad\quad\quad\quad\quad\quad \text{else return}((v_1 \mid Int) + (v_2 \mid Int) \text{ in } Value)$$

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

err : *String* $\rightarrow$ *Compute a*

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

"$E[\![n]\!]$" just returns the number n (injected into the *Value* domain) as the result of a trivial computation. To evaluate "$e_1 + e_2$", we evaluate e_1 and e_2 in turn, and then sum the results. 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.*

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* $\rightarrow$ *Compute Value*
rdEnv   : *Compute Env*
inEnv   : *Env* $\rightarrow$ *Compute Value* $\rightarrow$ *Compute Value*

The standard semantics for functions is as follows:

$$E[\![v]\!] \quad\quad\quad = \text{rdEnv then } \lambda\rho.\rho\ v$$
$$E[\![\lambda v.e]\!] \quad\quad = \text{rdEnv then } \lambda\rho.\text{return}(\lambda c.\text{inEnv } (\rho[c/v])\ E[\![e]\!])$$
$$E[\![(e_1\ e_2)_n]\!] = E[\![e_1]\!] \text{ then } \lambda f.\text{rdEnv then } \lambda\rho.f(\text{inEnv } \rho\ E[\![e_2]\!])$$
$$E[\![(e_1\ e_2)_v]\!] = E[\![e_1]\!] \text{ then } \lambda f.E[\![e_2]\!] \text{ then } \lambda v.f(\text{return } v)$$

The difference between call-by-value and call-by-name is clear: the former reduces the argument before invoking the function, whereas the latter packages the argument with the current environment to form a closure.

To simplify the presentation somewhat, we assume that imperative features can be defined using the primitive monad combinator:

update : (*Store* $\rightarrow$ *Store*) $\rightarrow$ *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 store-manipulating functions can guarantee that the store is never duplicated (see, for example, [26]).

With the kernel-level function:

callcc : ((*Value* $\rightarrow$ *Compute Value*) $\rightarrow$ *Compute Value*) $\rightarrow$ *Compute Value*

The semantics of "callcc" is a function expecting another function as an argument, to which the current continuation will be passed:

$E[\![\text{callcc}]\!] = \text{return}(\lambda f.f \text{ then } \lambda f'.\text{callcc}(\lambda k.f'(\lambda a.a \text{ then } k)))$

Our high-level monadic semantics somewhat resembles to action semantics, except that it uses only "then" and "return" to thread computations. Together with primitive 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 an effort to give a high-level view of the source language semantics.

We require that a semantics specified in terms of monadic combinators 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 inductive 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 the abstract syntax tree.

2.2 Constructing the *Compute* Monad

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

err   : *String* $\rightarrow$ *Compute a*
rdEnv   : *Compute Env*
inEnv   : *Env* $\rightarrow$ *Compute Value* $\rightarrow$ *Compute Value*
update : (*Store* $\rightarrow$ *Store*) $\rightarrow$ *Compute Store*
callcc : ((*Value* $\rightarrow$ *Compute Value*) $\rightarrow$ *Compute Value*) $\rightarrow$ *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 features 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* $= a$

$\text{return}_{Id}\,x = x$
$c\,\text{then}_{Id}\,f = f\,c$

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* $= s \rightarrow m\,(s, a)$

<div style="border:1px solid">

Continuation:

type $ContT\ r\ m\ a = (a \to m\ r) \to m\ r$

$\text{return}_{(ContT\ r\ m)}\ a = \lambda k.ka$

$c\ \text{then}_{(ContT\ r\ m)}\ f = \lambda k.c(\lambda a.fak)$

$\text{callcc}\ f = \lambda k.f(\lambda a.\lambda k'.ka)k$

Environment:

type $EnvT\ e\ m\ a = e \to m\ a$

$\text{return}_{(EnvT\ e\ m)}\ a = \lambda \rho.\ \text{return}_m\ a$

$c\ \text{then}_{(EnvT\ e\ m)}\ k = \lambda \rho.c\rho\ \text{then}_m\ \lambda a.ka\rho$

$\text{rdEnv} = \lambda \rho.\ \text{return}_m\ \rho$

$\text{inEnv}\ \rho\ c = \lambda \rho'.c\ \rho$

Errors:

type $Error\ a = Ok\ a \mid Error\ String$

type $ErrorT\ m\ a = m\ (Error\ a)$

$\text{return}_{(ErrorT\ m)}\ a = \text{return}_m\ \cdot Ok$

$c\ \text{then}_{(ErrorT\ m)}\ k = c\ \text{then}_m\ \lambda a.\ \textbf{case}\ a\ \textbf{of}$
$$(Ok\ x) \to kx$$
$$(Error\ msg) \to \text{return}_m(Error\ msg)$$

$\text{err} = \text{return}_m\ \cdot Error$

</div>

Fig. 2. Monad transformers

$\text{return}_{(StateT\ s\ m)}\ x = \lambda s.\ \text{return}_m(s, x)$

$c\ \text{then}_{(StateT\ s\ m)}\ k = \lambda s_0.c\ s_0\ \text{then}_m\ \lambda(s_1, a).k\ a\ s_1$

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

type $StateT\ s\ Id\ a = s \to Id\ (s, a)$
$$= s \to (s, a)$$

$\text{return}_{(StateT\ s\ Id)}\ x = \lambda s.\ \text{return}_{Id}(s, x)$
$$= \lambda s.(s, x)$$

$c\ \text{then}_{(StateT\ s\ Id)}\ k = \lambda s_0.c\ s_0\ \text{then}_{Id}\ \lambda(s_1, a).k\ a\ s_1$
$$= \lambda s_0.\textbf{let}\ (s_1, a) = m\ s_0\ \textbf{in}\ k\ a\ s_1$$

Note that "$StateT\ s\ Id$" is the standard state monad found, for example, in Wadler's work [25].

To make the newly introduced state accessible, "$StateT\ s$" introduces update on s which applies f to the state, and returns the old state:

$\text{update} : (s \to s) \to StateT\ s\ m\ a$

$\text{update}\ f = \lambda s.\ \text{return}(f\ s, s)$

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.

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:

$\text{err}_{(t\ m)} \quad = \text{lift}_{(t\ m)} \cdot \text{err}_m$
$\text{update}_{(t\ m)} = \text{lift}_{(t\ m)} \cdot \text{update}_m$
$\text{rdEnv}_{(t\ m)} = \text{lift}_{(t\ m)} \text{rdEnv}_m$

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

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

Fig. 3. Liftings

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 [17] for a detailed description.) For example, in the *Compute* monad above, "update f" is: "$\lambda s.Ok\ (f\ s,s)$" 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.

3 Using Monad Laws to Transform Programs

Following the monadic semantics presented in the previous section, by unfolding all calls to the 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 $\lambda\rho$.
return(λc.inEnv $(\rho[c/\text{"}x\text{"}])$
 (rdEnv then $\lambda\rho$.
 $\rho\text{"}x\text{"}$ then λv_1.
 return 1 then λv_2.
 return($v_1 + v_2$))) then λf.
return 2 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 1. A monad M is a type constructor, together with two operations:

then : $M\ a \to (a \to M\ b) \to M\ b$
return : $a \to M\ a$

satisfying the following laws [25]:

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

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.

3.2 Applying Monad Laws to Program Transformation

Monad laws are useful for transforming (and optimizing) monadic style intermediate programs. For example, our compiler translates the expression "$2 + 3$" to:

return 2 then λv_1. return 3 then λv_2. $\text{return}(v_1 + v_2)$

We can apply the left unit law twice, and reduce the above to: "$\text{return}(2 + 3)$", which can of course be further optimized to "return 5".

Each 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 (such as the final reduction of "$2 + 3$" to "5" in the above example), 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 $\lambda \rho$.
($\lambda c.\text{inEnv}\,(\rho[c/"x"])$
   (rdEnv then $\lambda \rho$.
   $\rho"x"$ then λv.
   $\text{return}(v + 1)))$
($\text{return } 2))$

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

3.3 Environment Axioms

We axiomatize the environment-manipulating functions as follows:

Definition 2. Monad M is an **environment monad** if it has two operations: rdEnv and inEnv, which satisfy the following axioms:

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

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).

Proposition 3. *The monads supporting* rdEnv *and* inEnv *constructed using monad transformers* ErrorT, EnvT, StateT *and* ContT *are environment monads.*

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

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

rdEnv then $\lambda\rho.(\lambda c.c$ then $\lambda v.$ 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.

Once again note that the true computation in the original expression "$((\lambda x.x + 1)\ 2)_v$" is left unreduced. In the traditional denotational semantics framework, it is harder to distinguish the redexes introduced by the compilation process from computations in the source program. In the above example, we could safely further reduce the intermediate code:

$(\lambda c.c$ then $\lambda v.$ return$(v+1))($return $2)$
$\Rightarrow$ return 2 then $\lambda v.$ return$(v+1)$ $\qquad (\beta)$
$\Rightarrow$ return 3 $\qquad\qquad\qquad$ (left unit)

However, in general, unrestricted β reduction for arbitrary source programs could result in unwanted compile-time exceptions, such as in "$((\lambda x.10/x)\ 0)_v$."

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 [28], 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 v, we assume there is a corresponding variable name v_m in the meta-language.

Definition 4. A **natural semantics** uses the environment of the meta-language for variables in the source language, and is given as follows:

$$N[\![v]\!] \quad = v_m$$
$$N[\![\lambda v.e]\!] \quad = \mathsf{return}(\lambda v_m.N[\![e]\!])$$
$$N[\![(e_1 \; e_2)_n]\!] = N[\![e_1]\!] \text{ then } \lambda f.f(N[\![e_2]\!])$$
$$N[\![(e_1 \; e_2)_v]\!] = N[\![e_1]\!] \text{ then } \lambda f.N[\![e_2]\!] \text{ then } \lambda v.f(\mathsf{return} \; v)$$

Other source-level constructs, such as $+$, $:=$, and callcc, 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 [28], guarantees that it is safe to implement function calls in the source language using the meta-language environment.

Theorem 5. *Let e be a program in the source language, $E[\![e]\!]$ be its standard semantics in an environment monad, $N[\![e]\!]$ be its natural semantics in the same monad,[3] and ρ be the mapping from the source language variable names v to v_m, we have:*

$$\mathsf{inEnv} \; \rho \; E[\![e]\!] = N[\![e]\!]$$

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.

---

[3] This 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 [28]. Fortunately, the *naturality of liftings* (see our earlier work [17] 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.)

Proof for the Functional Sub-language We can establish the theorem for the functional sub-language by induction on the structure of programs composed out of 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 instead supported dynamically scoped functions:

$$E[\![\lambda v.e]\!] = \text{return}(\lambda c.\text{rdEnv then } \lambda \rho.\text{inEnv } (\rho[c/v]) \ E[\![e]\!])$$

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

Extension to the Complete Language Consider another source language construct "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, 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 ":=".

4.3 Benefits of Reasoning in Monadic Style

Modular Proofs 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. For example, Wand [28] pointed out that he could change the semantics into continuation-based, and prove the theorem, but only by modifying the proofs accordingly.

Modular monadic semantics, on the other hand, offers highly modularized proofs and more general results. This is particularly applicable to real programming languages, which usually carry a large set of features and undergo constantly evolving designs.

Axiomatization of Programming Language Features Denotational semantics captures a computation 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* [20] in action semantics, monads provide an *axiomatic view* of denotational semantics.

The environment axioms provide an answer to the question: "what constitutes an environment?" We are investigating useful axioms for other programming language features, such as exceptions and continuations.

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
return x	x
c_1 then $\lambda x.c_2$	let val $x = c_1$ in c_2 end
update*	ref, !, :=
callcc	callcc
err	raise Err

* ML reference cells support single-threaded states.

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):

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

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'* $a = EnvT\ Env\ (ContT\ Answer\ (StateT\ Store\ (ErrorT\ M_{ML})))\ a$

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 [19]

[17]): 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 any more, it suffices to run the target program on *Compute''*:

type *Compute''* $a = M_{ML}\ a$

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 [15] [7] [17] 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 monad transformers, that the resulting change of ordering does not have unwanted effects on semantics.

6 Related work

Early efforts (e.g., [27]) in semantics-directed compiler generation were based on traditional denotational semantics.

Mosses's *Action Semantics* [21] allows modular specification of programming language semantics, from which efficient compilers can be generated. Action semantics (e.g., [3]) and a related approach by Lee [16] have been successfully used to generate efficient compilers. While action semantics is easy to construct, extend, understand and implement, we note the following comments made by Mosses ([21], 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 approach.

Moggi [19] first used monads and monad transformers to structure denotational semantics.[4] Wadler [26] [25] popularized Moggi's ideas in the functional programming community by using monads in functional programs, in particular, interpreters. This paper is built upon our work on monad-based modular interpreters [17], which in turn follows a series of earlier attempts by Steele [23], Jones and Duponcheel [11], and Espinosa [7].

Moggi [19] raised the issue of reasoning in the monadic framework. Wadler [26] listed the state monad laws. Hudak [9] suggested a more general framework — mutable abstract data types (MADTs) — to reason about states.

Meijer [18] combined the standard initial algebra semantics approach with aspects of Action Semantics to derive compilers from denotational semantics. An interesting area of future research is to combine the nice algebraic properties in Meijer's framework with the modularity offered in ours.

One application of partial evaluation [12] is to automatically generate compilers from interpreters. A partial evaluator has been successfully applied to an action interpreter [2], and similar results can be achieved with monadic interpreters [6].

Staging transformations, first proposed by Jørring and Scherlis [13], are a class of general program transformation techniques for separating a given computation into stages. Monad 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 [14], 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 [22] and Hatcliff and Danvy [8]). We will continue to explore along this direction in order to generate machine-level code from a monadic intermediate form.

7 Conclusions

We have shown that the monadic framework provides good support for high-level extensible specifications, program transformations, reasoning about computations, and code generation in various target languages. 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. Overall, we believe that modular monadic semantics is particularly suitable for compiler construction.

---

[4] More recently, Cenciarelli and Moggi [4] proposed a syntactic approach to modularity in denotational semantics.

anonymous referees for helpful comments on an earlier version of the paper, and Ross Paterson, Tim Sheard and John Launchbury for ideas and discussions.

References

1. Andrew W. Appel and Trevor Jim. Continuation-passing, closure-passing style. In *ACM Symposium on Principles of Programming Languages*, pages 193–302, January 1989.
2. Anders Bondorf and Jens Palsberg. Compiling actions by partial evaluation. In *FPCA '93: Conference on Functional Programming Languages and Computer Architecture, Copenhagen, Denmark*, pages 308–317, New York, June 1993. ACM Press.
3. Deryck F. Brown, Hermano Moura, and David A. Watt. ACTRESS: An action semantics directed compiler generator. In *Proceedings of the 4th International Conference on Compiler Construction, Edinburgh, U.K.*, pages 95–109. Springer-Verlag, 1992. Lecture Notes in Computer Science 641.
4. Pietro Cenciarelli and Eugenio Moggi. A syntactic approach to modularity in denotational semantics. In *Category Theory and Computer Science '93*, 1993.
5. D. Clément, Joëlle Despeyroux, Thierry Despeyroux, and Gilles Kahn. A simple applicative language: Mini-ML. In *Proceedings of the 1986 ACM Symposium on Lisp and Functional Programming*, pages 13–27, 1986.
6. Olivier Danvy, Jürgen Koslowski, and Karoline Malmkjær. Compiling monads. Technical Report CIS-92-3, Kansas State University, December 1991.
7. David Espinosa. Modular denotational semantics. Unpublished manuscript, December 1993.
8. John Hatcliff and Olivier Danvy. A generic account of continuation-passing styles. In *21st ACM Symposium on Principles of Programming Languages (POPL '94), Portland, Oregon*, pages 458–471, New York, January 1994. ACM Press.
9. Paul Hudak. Mutable abstract datatypes. Research Report YALEU/DCS/RR-914, Yale University Department of Computer Science, December 1992.
10. Paul Hudak, Simon Peyton Jones, and Philip Wadler. Report on the programming language Haskell: a non-strict, purely functional language, version 1.2. Technical Report YALEU/DCS/RR-777, Yale University Department of Computer Science, March 1992. Also in ACM SIGPLAN Notices, Vol. 27(5), May 1992.
11. Mark P. Jones and Luc Duponcheel. Composing monads. Research Report YALEU/DCS/RR-1004, Yale University Department of Computer Science, New Haven, Connecticut, December 1993.
12. Neil D. Jones, Peter Sestoft, and Harald Søndergaard. Mix: a self-applicable partial evaluator for experiments in compiler generation. *Lisp and Symbolic Computation*, 2:9–50, 1989.
13. Ulrik Jørring and William Scherlis. Compilers and staging transformations. In *Proceedings Thirteenth ACM Symposium on Principles of Programming Languages, St. Petersburg, Florida*, pages 86–96, 1986.
14. Richard Kelsey and Paul Hudak. Realistic compilation by program transformation. In *ACM Symposium on Principles of Programming Languages*, pages 181–192, January 1989.
15. David J. King and Philip Wadler. Combining monads. In John Launchbury and Patrick Sansom, editors, *Functional Programming, Glasgow 1992*, pages 134–143, New York, 1993. Springer-Verlag.
16. Peter Lee. *Realistic Compiler Generation*. Foundations of Computing. MIT Press, 1989.

17. Sheng Liang, Paul Hudak, and Mark Jones. Monad transformers and modular interpreters. In *22nd ACM Symposium on Principles of Programming Languages (POPL '95), San Francisco, California*, New York, January 1995. ACM Press.
18. Erik Meijer. More advice on proving a compiler correct: Improving a correct compiler. Submitted to Journal of Functional Programming.
19. Eugenio Moggi. An abstract view of programming languages. Technical Report ECS-LFCS-90-113, Laboratory for Foundations of Computer Science, University of Edinburgh, Edinburgh, Scotland, 1990.
20. Peter D. Mosses. A basic abstract semantic algebra. In Gilles Kahn, David B. Mac-Queen, and Gordon D. Plotkin, editors, *Semantics of Data Types: International Symposium*, pages 87–107. Springer-Verlag, June 1984. LNCS 173.
21. Peter D. Mosses. *Action Semantics*, volume 26 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1992.
22. Amr Sabry and Matthias Felleisen. Reasoning about programs in continuation-passing style. In *Proceedings of the 1992 ACM Conference on LISP and Functional Programming*, pages 288–298. ACM Press, June 1992.
23. Guy L. Steele Jr. Building interpreters by composing monads. In *Conference Record of POPL '94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, Oregon*, pages 472–492, New York, January 1994. ACM Press.
24. Joseph Stoy. *Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory*. MIT Press, 1977.
25. Philip Wadler. The essence of functional programming. In *Conference Record of the Nineteenth Annual ACM Symposium on Principles of Programming Languages, Albuquerque, New Mexico*, pages 1–14, January 1992.
26. Philip L. Wadler. Comprehending monads. In *Proceedings of the 1990 ACM Conference on Lisp and Functional Programming*, 1990.
27. Mitchell Wand. A semantic prototyping system. *SIGPLAN Notices, ACM Symposium on Compiler Construction*, 19(6):213–221, 1984.
28. Mitchell Wand. A short proof of the lexical addressing algorithm. *Information Processing Letters*, 35:1–5, June 1990.

The proof of theorem 5

The main theorem is then proved by induction on the structure of programs:
Case $\lambda v.e$:

$\mathsf{inEnv}\ \rho\ (E[\![\lambda v.e]\!])$

$= \mathsf{inEnv}\ \rho\ (\mathsf{rdEnv}\ \mathsf{then}\ \lambda\rho.\ \mathsf{return}(\lambda c.\mathsf{inEnv}\ (\rho[c/v])\ E[\![e]\!]))$

$= \mathsf{inEnv}\ \rho\ \mathsf{rdEnv}\ \mathsf{then}\ \lambda\rho'.$

$\quad \mathsf{inEnv}\ \rho\ (\mathsf{return}(\lambda c.\mathsf{inEnv}\ (\rho'[c/v])\ E[\![e]\!]))$ (dist.)

$= \mathsf{return}\ \rho\ \mathsf{then}\ \lambda\rho'.\ \mathsf{return}(\lambda c.\mathsf{inEnv}\ (\rho'[c/v])\ E[\![e]\!])$ (cancel., unit)

$= \mathsf{return}(\lambda c.\mathsf{inEnv}\ (\rho[c/v])\ E[\![e]\!])$ (left unit)

$= \mathsf{return}(\lambda v_m.\mathsf{inEnv}\ (\rho[v_m/v])\ E[\![e]\!])$ (α renaming)

$= \mathsf{return}(\lambda v_m.N[\![e]\!])$ (ind. hypo.)

$= N[\![\lambda v.e]\!]$

Other cases $(v, (e_1\ e_2)_n, (e_1\ e_2)_v)$ can be similarly proved.