Typing First-Class Continuations in ML

Bruce F. Duba
Rice University

Robert Harper
Carnegie Mellon University

David MacQueen
AT&T Bell Laboratories

Abstract

An extension of Standard ML with continuation primitives similar to those found in Scheme is considered. A number of alternative type systems are discussed, and several programming examples are given. The semantics of type assignment for a small, purely functional fragment of the language is presented, for which both a Milner-style soundness theorem and an observational soundness theorem may be established.

1 Introduction

First-class continuations are a simple and natural way to provide access to the flow of evaluation in functional languages. The ability to seize the "current continuation" (control state of the evaluator) provides a simple and natural basis for defining numerous higher-level constructs such as coroutines [15], exceptions [41], and engines [6], for writing interpreters and compilers [30, 34], and for organizing run-time support for multiple threads of control [40, 27]. Tractable logics for reasoning about program equivalence in the presence of first-class continuations in an untyped setting have been developed [8, 9, 38]. Recent studies have focused on questions of typing for first-class continuations [12, 11, 14] and their impact on "full abstraction" results [32].

The subject of this paper is the extension of Standard ML with primitives for first-class continuations similar to those found in Scheme. The two new primitives are `callcc`, for *call with current continuation*, which takes a function as argument, and calls it with the current continuation, and `throw`, which takes a continuation and a value, and passes the value to that continuation.

This paper is organized as follows. In Section 2 we give an informal presentation of the extension of ML with continuation primitives, and illustrate their use in a programming example. We also discuss the role of continuations in the implementation of Standard ML of New Jersey, and some problems that they raised. In Section 3 we present a formal system of type assignment for a small functional fragment of ML. A denotational semantics for this fragment is given in Section 3, and the semantics of type assignment is considered. The main results are a Milner-style soundness theorem ("well-typed programs cannot go wrong") and an observational soundness theorem ("convergent programs of type `int` yield integers"). Finally, in Section 4 we give an operational semantics for the language in the "natural semantics" style of Plotkin and Kahn [25, 3]. The operational presentation illustrates the extent to which the definition of Standard ML [22] would have to be changed to accomodate the proposed extension.

We are grateful to Stephen Brookes, Matthias Felleisen, Andrzej Filinski, Timothy Griffin, and Olin Shivers for their comments and suggestions.

2 Adding Continuations to Standard ML

A first-class continuation is an abstraction that evolved from various nonstandard control structures such as Landin's J-operator [19], Reynold's `escape` [30], label variables in Gedanken [29] and PAL [7], and from the semantic analyses of general control structures, including jumps [35]. Scheme [37] originally introduced a binding construct `(catch x body)` that captured its own expression continuation and bound it to the variable `x`, with the expression `body` as the scope of the binding. The continuation represents the "rest of the computation," and behaves as a function that takes the value of the expression as its argument and yields the final result of the evaluation of the remainder of the program.

In a typical implementation the final result is passed to the interactive top-level, which prints the result, and continues by evaluating the next expression.

In 1982 Will Clinger and Dan Friedman [4] noticed that the act of capturing the current continuation did not require a special variable binding form, but could be performed by a primitive operation, called `call-with-current-continuation`, or `call/cc` for short, whose argument was a function that would be applied to the captured continuation, so that `(catch x body)` becomes `(call/cc (lambda (x) body))` in Scheme. (This is an example of the well-known technique of replacing a special variable binding form with an operation acting on a function, so that variable binding is handled solely by lambda abstraction.)

In an untyped language there is not much to choose between the functional and binding forms of continuation-capturing construct. However, in the context of an ML-like type system, the two differ substantially. To understand the distinction, it is helpful to consider the interaction between typing and the invocation of a captured continuation. There are two main points. First, continuations arise in a program only by capturing the evaluation context of some expression: there are no expression forms denoting continuations (in contrast to the language considered by Filinski [12].) Therefore continuations expect values of the type of the expression whose evaluation context the continuation represents. Second, the invocation of a captured continuation discards the current evaluation context, passing a value to the captured, instead of the current, continuation. Although the passed value must be consistent with the argument type of the continuation, the result type is unconstrained since invocations of continuations do not return to the evaluation context. (For similar reasons the exception-raising construct of Standard ML has arbitrary result type.)

For example, if `k` is bound to a continuation expecting an integer value, we may invoke `k` in several incompatible type contexts, as in the following expression[1]

```
if b then [ (k 3) + 1) ] else 5 :: (k 4)
```

Here `k` is invoked in two contexts, one expecting an integer, the other expecting an integer list. Since continuation invocations never return, it makes sense to regard this as a well-typed expression (of type `int list`).

The incorporation of continuation primitives in ML involves making two, related decisions, namely the continuation-capturing construct and the continuation-invoking construct. Since ML is a typed language, continuations should be values of some type, say τ `cont`, the type of continuations expecting values of type τ. The continuation-capturing constructs may then be given typing rules as follows. The functional form, written `callcc` in keeping with the ML lexical conventions, may be assigned any type of the form

$$(\tau\ \texttt{cont} \to \tau) \to \tau$$

since the body may either invoke the passed continuation, or else return normally. Written polymorphically, the type of `callcc` is then

$$\forall\alpha.(\alpha\ \texttt{cont} \to \alpha) \to \alpha.$$

The variable-binding form, written `let` k `in` e, has the following typing rule:

$$\frac{A, k{:}\tau\ \texttt{cont} \vdash e : \tau}{A \vdash \texttt{let}\ k\ \texttt{in}\ e : \tau}$$

where A is a type assignment giving types to the set of free variables of e.

The choice of functional or binding form of continuation-capturing construct depends on the definition of the type τ `cont`. We consider two possibilities: regard a continuation as a function that is invoked by application, or regard a continuation as a new form of value that is invoked by a special primitive. In the first case the type τ `cont` is rendered as a functional type, whereas in the second it is introduced as a new primitive type. We consider each in turn.

If continuations are to be regarded as functions, some provision must be made for ensuring that the result type is allowed to vary according to context. This suggests the following polymorphic typing:

$$\texttt{callcc} : \forall\alpha.\forall\beta.((\alpha \to \beta) \to \alpha) \to \alpha$$

But since `k` is lambda-bound in the expression `callcc(fn k => ...)` this does not give us the freedom to instantiate the polymorphic type variable β independently at each applied occurrence of `k` within the body of the abstraction. Instead we are forced to choose a *single* type for β suitable for all applications of k, ruling out examples such as the one considered above.

There are two ways to proceed. One involves formally moving the quantifier over β inward (which could be formally justified by the observation that β occurs in a positive position in the type expression), yielding the typing

$$\texttt{callcc} : \forall\alpha.((\alpha \to \forall\beta.\beta) \to \alpha) \to \alpha,$$

then replacing the type $\forall\beta.\beta$, which is not the type of any defined value, by a new primitive type `void`, resulting in the typing

$$\texttt{callcc} : \forall\alpha.((\alpha \to \texttt{void}) \to \alpha) \to \alpha.$$

[1] We (temporarily) use ordinary function application notation to indicate invocation of a continuation.

The type τ `cont` is then defined to be the type $\tau \to$ `void`. To match the type of a continuation invocation (*i.e.*, `void`) with its context we could either view void as a subtype of all types and use a subsumption rule (which introduces many of the complexities of subtyping into the type system), or we can simply introduce a polymorphic coercion function

$$\texttt{ignore :}\quad \texttt{void} \to \alpha$$

and surround applications of continuations with a call to `ignore`, as in

```
if b then [ ignore(k 3) + 1 ]
     else 5 :: ignore(k 4)
```

(where `k` is an `int cont`).

An alternative to `ignore` is afforded by exploiting the polymorphic type system of ML, using `letcc` instead of `callcc`. The idea is to take type τ `cont` to stand for the polymorphic type $\forall\alpha.\tau \to \alpha$, leading to the following typing rule for `letcc`:

$$\frac{A, k : \forall\alpha.\tau \to \alpha \vdash e : \tau}{A \vdash \texttt{letcc}\ k\ \texttt{in}\ e : \tau}$$

Since k is assigned a polymorphic type, the result type, α, may be chosen freely on a case-by-case basis. This rule is consistent with the ML type system in that `let`-like constructs admit assignment of polymorphic types to the bound identifier. Since types of the form $(\forall\alpha.\tau \to \alpha) \to \tau$ lie outside of the scope of the ML type system, this method may not be adapted to `callcc`. It is here that the two constructs differ in an ML-like setting.

Another way of typing continuations, and the one currently adopted in Standard ML of New Jersey [2], is to abandon the view that continuations are functions in the ordinary sense and to consider τ `cont` as a primitive type with an operation `throw` for invoking a continuation. The type of `throw` is given by

$$\texttt{throw} : \forall\alpha.\forall\beta.(\alpha\ \texttt{cont}) \to (\alpha \to \beta),$$

and hence `throw` is essentially a coercion that turns a continuation into a function, reintroducing a separate instance of the type parameter β at each invocation of the continuation. Our example becomes

```
if b then [ (throw k 3) + 1 ]
     else 5 :: (throw k 4)
```

where the first and second occurrences of throw receive the types `int cont` $\to$ `int` $\to$ `int` and `int cont` $\to$ `int` $\to$ `int list`, respectively. From a programming point of view, there is not much to distinguish this approach from the use of empty types together with `ignore`: the expression `throw M N` could be provided as syntactic sugar for `ignore (M N)`. Conversely, `ignore M` could be taken as standing for `throw` κ_0 `M`, where κ_0 is the initial continuation. Sine M must be of type `void`, it cannot return a value, and hence the abort to top level will never in fact occur.

It would seem, then that there are essentially two alternatives for representing continuations in ML: as polymorphic functions, using `letcc` as the capturing construct, and values of a new primitive type, using `throw` to invoke them. Although the two are equivalent for the purely functional fragment of ML, the approach based on a primitive type of continuations is better-behaved in the context of the full Standard ML language than is the polymorphic approach. The problem is that current schemes for introducing references (assignable cells) in ML preclude the possibility of storing objects of polymorphic type. For instance, if the identity function is stored into a cell, then a single instance of its polymorphic type must be chosen for all subsequent retrievals: it's polymorphic character is lost. (See Tofte's thesis [39] for further discussion of this point.) Thus if continuations were represented as functions of polymorphic result type, then the result type would have to be fixed at the time that the continuation is *stored*, rather than *invoked*, significantly limiting their utility. The approach based on a primitive type of continuations does not suffer from this limitation, and is therefore to be preferred for Standard ML.

We are thus led to the following simple signature for supporting first-class continuations in Standard ML:

```
type α cont
val callcc : (α  cont → α) →  α
val throw : α  cont →  α  →  β
```

(We could just as well have taken `letcc` as primitive, but since there is no advantage in doing so, it is simpler to introduce `callcc` as a new constant of polymorphic type.)

Some examples will suggest how first-class continuations are used in practice. The simplest and earliest use of continuations was to provide an escape function, as in the following function that scans a list of integers for a negative element and returns it via a continuation if one is found.

```
fun catch_neg exit x =
    if x < 0 then
       throw exit (SOME x)
    else ()
fun find_neg(l:int list):int option =
    callcc(
      fn exit =>
         ( app (catch_neg exit) l);
           NONE )
```

The function app takes two arguments, a function and a list, and applies the function to every element of the list. (The type τ option has as values NONE, and SOME v, where v is a value of type τ.)

Another common application is to implement coroutines. Here an interesting typing issue arises. A common technique is to resume a coroutine by passing the continuation of the current coroutine as the argument to the continuation representing the resumed coroutine. If state is the type representing the state of a coroutine, this leads naively to the circular identification

```
state = state cont
```

We cannot solve this identity directly, but we can use a datatype declaration to define the type state recursively. This is illustrated by the following example of a pair of coroutines, one producing and the other consuming a sequence of integers.

```
datatype state = S of state cont
fun resume(S k: state) : state =
    callcc(fn k': state cont =>
              throw k (S k'))
val buf = ref 0
fun produce(n: int, cons: state) =
    (buf := n; produce(n+1, resume(cons)))
fun consume(prod: state) =
    (print(!buf); consume(resume prod))
fun pinit (n: int) : state =
    callcc(fn k : state cont =>
              produce(n,S k))
fun prun () = consume(pinit(0))
```

Closely related to coroutines are lightweight processes or threads, and continuations have been used as the basis for several implementations of process facilities for Standard ML of New Jersey, some of which use preemptive scheduling [27, 5, 26, 36].

Another use of continuations is to provide a clean, typed interface for asynchronous signal handling [28]. In Standard ML of New Jersey the type of a signal handler is (int * unit cont) -> unit cont, where the argument is a pair consisting of a count of pending signals of the kind being handled, and the continuation representing the interrupted process. The continuation returned by the signal handler is typically used to resume the interrupted process after signals have been unmasked, but it can also provide an alternative continuation like aborting the computation and returning to top-level. The signature of the signal handling facilities is:

```
signature SIGNALS =
  sig
    datatype signal =
        SIGHUP | SIGINT | SIGQUIT | ...
    val setHdlr : (signal * handler -> unit
    val inqHdlr : signal -> handler
    val maskSignals : bool -> unit
  end
```

where *handler* stands for the type

```
(int * unit cont) -> unit cont) option.
```

The function setHdlr associates a handler with a signal; inqHdlr returns the handler associated with a signal; maskSignals masks out all catchable signals.

Remark 2.1 *We definitely need an example that stores continuations into a ref cell since the discussion about typing hinges on it. (RH)*

In an interactive system such as Standard ML of New Jersey, the context of a continuation has to be properly defined and controlled to avoid subverting the type system through continuations. This could happen in the following way: in the context of a top-level evaluation of an expression of type int we capture a continuation (of type unit cont say) and store it in a global reference. Then we later invoke this stored continuation in the context of another top level evaluation of type boolean. This latter context receives an integer value from the stored continuation and tries to interpret it as a boolean value. To prevent this, we use a stack of timestamps that allow us to detect the invocation of a continuation outside of its proper context.

Remark 2.2 *The discussion of the interactive front-end needs to be expanded. I do not understand the problem about the front-end. My understanding of it is that in the current (at least, recent) implementation the continuation invoked by a top-level* callcc *is the "wrong" one: instead of branching back to complete the* val*-binding for* it*, it instead does something else. In any case, more examples and explanation are needed here. (RH)*

Another interesting issue is the relation between continuations and exception handling. The evaluation context of an expression in ML actually consists of two continuations, one for the "normal" value return, the other for the exception handling context. The continuation-passing primitives discussed above provide access only to the normal value continuation, and not to the exception handling continuation. It would be possible to add a second primitive, say throwRaise : α cont $\rightarrow$ (exn $\rightarrow$ β) that could be used to throw an exception value to the exception component of the continuation. This primitive is currently not provided because it could be used in conjunction with the asynchronous signal handling facilities to cause asynchronous exceptions, and for various reasons these are to be discouraged.

Remark 2.3 *What are the various reasons? Could you spell out in more detail why the two continuations are needed? Describe exception raising and handling in terms of the exception continuation. Is there a use for* `throwRaise`*? Also, how about* `throwExn`*? It might be worth explaining how handlers are bound, and to explain the interaction between continuations and references in the implementation of exceptions. (RH)*

3 A Denotational Semantics of Typing

In this section we study the soundness of type assignment for a small, purely functional, monomorphic fragment of ML extended with primitives for first-class continuations. Explicit treatment of polymorphism and `let` is omitted since the main issues do not involve polymorphism.

3.1 Type Assignment

Consider the following language:

$$M \quad ::= \quad x \mid \lambda x.M \mid M\,N \mid \texttt{letcc}\,x\,\texttt{in}\,M \mid \texttt{throw}\,M\,N$$

The variables x, y, and z range over a set of variables, and M, N, P, Q, and R range over the set of terms. The expression $\texttt{letcc}\,x\,\texttt{in}\,M$ binds x in M. We use `letcc` in place of `callcc` to facilitate comparison between the various approaches to type assignment.

Let b range over some set B of base types. Type expressions are defined by the following grammar:

$$\tau \quad ::= \quad b \mid \tau_1 \to \tau_2 \mid \tau\ \mathit{cont}$$

A *typing context* is a partial function Γ mapping some finite set of variables $\mathrm{dom}(\Gamma)$ to type expressions. A *typing assertion* is a triple $\Gamma \vdash M : \tau$, where τ is a type expression.

We take as given a map *type* assigning a type to each constant. (In a polymorphic system this map would assign a type scheme to each constant.) The type assignment rules for the above language are as follows:

$$\Gamma \vdash x : \Gamma(x) \qquad (\textsc{var})$$

$$\frac{\Gamma[x:\tau_1] \vdash M : \tau_2}{\Gamma \vdash \lambda x.M : \tau_1 \to \tau_2} \qquad (\textsc{abs})$$

$$\frac{\Gamma \vdash M : \tau \to \tau' \quad \Gamma \vdash N : \tau}{\Gamma \vdash M\,N : \tau'} \qquad (\textsc{app})$$

$$\frac{\Gamma[x:\tau\ \mathit{cont}] \vdash M : \tau}{\Gamma \vdash \texttt{letcc}\,x\,\texttt{in}\,M : \tau} \qquad (\textsc{letcc})$$

$$\frac{\Gamma \vdash M : \tau\ \mathit{cont} \quad \Gamma \vdash N : \tau}{\Gamma \vdash \texttt{throw}\,M\,N : \tau'} \qquad (\textsc{throw})$$

3.2 Denotational Semantics

We give here a denotation semantics for the untyped language, in preparation for establishing a Milner-type soundness theorem. The value space is given by the following domain equations:

$$\begin{array}{rcl} \mathit{Val} & = & \mathit{Bas}(\mathit{BaseVal}) + \mathit{Clsr}(\mathit{Val} \to (\mathit{Cont} \to \mathit{Ans})) + \\ & & \mathit{Cnt}(\mathit{Cont}) \\ \mathit{Cont} & = & \mathit{Val} \to \mathit{Ans} \\ \mathit{Ans} & = & \mathit{Value}(\mathit{Val}) + \mathit{Wrong}() \end{array}$$

An answer is either a value or a special token *wrong* representing a "run-time" type error.

Let $\mathit{Env} = \mathit{Var} \to \mathit{Val}$. Let the variables u, v, w, m, and n range over Val, the variable κ range over Cont, and the variable ρ range over Env.

The definition of the meaning function $[\![\cdot]\!] : \mathit{Env} \to \mathit{Cont} \to \mathit{Ans}$ is given by induction on the structure of expressions as follows:

$$\begin{array}{rcl} [\![x]\!]\,\rho\kappa & = & \kappa(\rho x) \\ [\![\lambda x.M]\!]\,\rho\kappa & = & \kappa(\mathit{clsr}(\lambda u'.\lambda\kappa'.\,[\![M]\!]\,\rho[x{:=}u']\kappa')) \\ [\![M\,N]\!]\,\rho\kappa & = & [\![M]\!]\,\rho(\lambda m.\,[\![N]\!]\,\rho(\lambda n. \\ & & \mathit{let}\ \mathit{clsr}(f){=}m\ \mathit{in}\ fn\kappa\ \mathit{else}\ \mathit{wrong})) \\ [\![\texttt{letcc}\,x\,\texttt{in}\,M]\!]\,\rho\kappa & = & [\![M]\!]\,\rho[x{:=}\,\mathit{cnt}(\kappa)]\kappa \\ [\![\texttt{throw}\,M\,N]\!]\,\rho\kappa & = & [\![M]\!]\,\rho(\lambda m.\,[\![N]\!]\,\rho(\lambda n. \\ & & \mathit{let}\ \mathit{cnt}(\kappa'){=}m\ \mathit{in}\ \kappa' n\ \mathit{else}\ \mathit{wrong})) \end{array}$$

3.3 Soundness

To state the semantic soundness theorem, we require two definitions, one for a value to be of a given type, written $v : \tau$, and one for a continuation to accept values of a given type, written $\kappa :: \tau$. These are defined simultaneously by induction on the structure of τ as follows:

1. $v : \tau$ holds iff $v = \perp$ or
   (a) $\tau = b$ and $v = \mathit{bas}(w)$ with w a value of base type b.
   (b) $\tau = \tau_1 \to \tau_2$ and $v = \mathit{clsr}(f)$ and for all values v_1 and continuations κ_2, if $v_1 : \tau_1$ and $\kappa_2 :: \tau_2$, then $f(v_1)(\kappa_2) \neq \mathit{wrong}$.

(c) $\tau = \tau_1\ cont$ and $v = cnt(\kappa)$ and $\kappa :: \tau_1$.

2. $\kappa :: \tau$ holds iff $\kappa(v) \neq wrong$ for all v such that $v : \tau$.

To check that this in fact is a proper definition, simply "expand" the definition of $\kappa :: \tau$ in the definition of $v : \tau$, and check that the membership relation is used only on subsidiary types of the given type. The relation $v : \tau$ is extended to environments pointwise: $\rho : \Gamma$ iff for all $x \in \mathrm{dom}(\Gamma)$, $\rho(x) : \Gamma(x)$.

Theorem 3.1 (Soundness) *If $\Gamma \vdash M : \tau$ and $\rho : \Gamma$ and $\kappa :: \tau$, then $[\![M]\!]\rho\kappa \neq wrong$.*

The proof is by a straightforward induction on the structure of the typing derivation.

As usual in continuation semantics, we need to provide an initial continuation. In order to preserve the soundness of typing the initial continuation must be chosen in such a way that it yields a non-*wrong* result for an argument of the type of the expression. An obvious choice is the continuation κ_0 defined by $\kappa_0(v) = value(v)$, which is essentially the identity function. In an implementation the initial continuation might print the result value of the computation, and hence must be chosen on a case-by-case basis, according to the type of the expression being evaluated.

3.4 Observational Soundness

In contrast to languages with a "direct" semantics, the soundness theorem does not yield positive results about typing. In particular, we may not conclude that the value of an arbitrary expression of a base type (*e.g.*, *int*) yields a result of the expected form (*e.g.*, a numeral). A natural attempt to obtain such results from the soundness theorem proceeds by choosing the continuation argument to be the function which yields *wrong* except on values of the desired type. However, this overlooks the fact that in a language with continuation-passing primitives, the result of evaluating an open expression need not be given in terms of its continuation argument. This suggests that the best we can hope for is a "observational soundness" theorem that yields positive results about closed expressions of base type.

It seems clear, on the basis of operational intuitions, that evaluation of a complete program either goes wrong, or else passes a value to the initial continuation. Unfortunately these operational intuitions do not seem to transfer readily to the setting of the untyped denotational semantics given above. We give here a brief sketch of the argument. The main idea is to prove that complete programs are "non-escaping" in the sense that their denotation is either *wrong*, or is determined as a function of the initial continuation. Since the semantics of certain closed expression involves the semantics of open sub-expressions, we must in fact prove a stronger result that takes account of environments and intermediate continuations. This entails extending the "non-escaping" property to arbitrary values, and to do so appears to require an inclusive predicate argument similar to that considered by Reynolds [31]. Given that such a predicate exists, we may choose the initial continuation as discussed above, and conclude, by the non-escaping property, that well-typed closed terms of base type evaluate to values of that type. Although it seems plausible that the required predicate exists, we have not proved this, and turn instead to a more straightforward argument.

The need for an inclusive predicate argument can be traced to the rich structure of the semantic domain needed to interpret untyped programs. By exploiting the type structure of the language, the complexity of reflexive domains can be avoided (while still admitting extensions such as *fix*). The fundamental idea is to adapt the methods of Meyer and Wand [20], and make use of standard results of the typed λ-calculus to obtain the desired result. The main idea is to define the meaning of a term M in our illustrative language as the meaning of its *cps transform*, $\overline{M}$, defined as follows:

$$\begin{array}{rcl}
\overline{x} & = & \lambda\kappa.\kappa\, x \\
\overline{\lambda x.M} & = & \lambda\kappa.\kappa\,(\lambda x.\overline{M}) \\
\overline{M N} & = & \lambda\kappa.\overline{M}\,(\lambda m.\overline{N}\,(\lambda n.mn\kappa)) \\
\overline{\texttt{letcc}\, x\, \texttt{in}\, M} & = & \lambda\kappa.[\kappa/x]\overline{M}\,\kappa \\
\overline{\texttt{throw}\, M\, N} & = & \lambda\kappa.\overline{M}\,(\lambda m.\overline{N}\,(\lambda n.mn))
\end{array}$$

(In fact, the Standard ML of New Jersey compiler is based on a similar transformation [1].)

To relate the type of a term M to the type of its cps transform, we associate to each type τ of the illustrative language, a simple type τ^* given by

$$\begin{array}{rcl}
b^* & = & b \\
(\sigma\to\tau)^* & = & \sigma^*\to((\tau^*\to\alpha)\to\alpha) \\
(\tau\ cont)^* & = & \tau^*\to\alpha
\end{array}$$

where α is a fixed, but arbitrary, base type. The function $(\ )^*$ is extended to typing contexts pointwise: if Γ is a typing context, then $\Gamma^*(x)$ is the typing context that assigns to each variable $x \in \mathrm{dom}(\Gamma)$ the simple type $\Gamma(x)^*$.

The following lemma relates the type of a term in the illustrative language to the type of its cps transform:

Lemma 3.2 *If $\Gamma \vdash M : \tau$, then $\Gamma^* \vdash^{\lambda^{\to}} M : (\tau^*\to\alpha)\to\alpha$.*

(Here $\vdash^{\lambda^{\to}}$ denotes derivability in the simply-typed λ-calculus.) The proof is a simple induction on the structure of M, taking account of the definition of $\overline{M}$.

Let $\mathcal{A}$ be any model of the simply-typed λ-calculus with carrier set A^τ for each type τ and application operation $\cdot : A^{\sigma\to\tau} \to A^\sigma \to A^\tau$. We say that an environment ρ *matches* a typing context Γ iff $\rho(x) \in A^{\Gamma(x)}$ for each $x \in \mathrm{dom}(\Gamma)$. It is a standard result that if $\Gamma \vdash^{\lambda^\to} M : \tau$ and ρ matches Γ, then $\mathcal{A}[\![M]\!]\rho \in A^\tau$ [23]. Hence if we define $[\![M]\!]$ to be $\mathcal{A}[\![\overline{M}]\!]$, then $[\![M]\!]\rho$ is in the set $A^{(\tau^*\to\alpha)\to\alpha}$ whenever ρ matches Γ^*. Now if M is a closed term of base type b, we may take ρ to be the empty environment, and $\alpha = b$, to obtain $[\![M]\!]\rho \in A^{(b\to b)\to b}$. Now since the identity function on type b is denotable by a λ term, and since $\mathcal{A}$ is a model of the typed λ calculus, we may apply this to id_b, the initial continuation, to obtain $[\![M]\!]\rho \cdot id_b \in A^b$. In other words, the value of M, when applied to the identity as initial continuation, is a value of base type b. In the case of cpo-based models (which are needed to interpret arbitrary recursion) with base types interpreted as flat cpo's, the result is either $\perp$, or a "true" value of the base type. Note that this argument does not extend to higher types τ since the definition of τ^* for higher types involves α, and hence we may not simply choose α to coincide with τ.

3.5 Continuations as Functions

Two alternative type systems given in the introduction rely on the representation of continuations as functions. We consider here the semantics of typing for the system based on empty types, and for the system based on regarding continuations as functions of polymorphic result type.

To express the idea that a continuation is a function that never returns, we introduce a type, `void`, with no values, and regard a continuation accepting values of type τ as a function of type $\tau \to$ `void`. This leads to the following typing rule for `letcc`:

$$\frac{\Gamma[x : \tau \to \texttt{void}] \vdash M : \tau}{\Gamma \vdash \texttt{letcc}\, x \,\texttt{in}\, M : \tau} \qquad (\textsc{callcc-empty})$$

Since the result of application of a continuation is now `void`, we must introduce, in compensation, a map `ignore` witnessing the inclusion of `void` into every type τ:

$$\frac{\Gamma \vdash M : \texttt{void}}{\Gamma \vdash \texttt{ignore}\ M : \tau} \qquad (\textsc{give-up})$$

(These two rules have the form of Pierce's Law and false elimination, respectively; see [14] for further discussion.) The expression `throw` $M\,N$ is now defined as `ignore` (MN). For the polymorphic variant, the typing rule for `letcc` appears in the introduction.

The denotational semantics must be changed to reflect the representation of continuations as functions. The domain equation for values is simplified to

$$Val = Bas(BaseVal) + Clsr(Val \to Cont \to Ans).$$

and the semantic equations become:

$$\begin{array}{rcl}
[\![x]\!]\rho\kappa & = & \kappa(\rho x) \\
[\![\lambda x.M]\!]\rho\kappa & = & \kappa(clsr(\lambda u'.\lambda\kappa'.[\![M]\!]\rho[x{:=}u']\kappa')) \\
[\![MN]\!]\rho\kappa & = & [\![M]\!]\rho(\lambda m.[\![N]\!]\rho(\lambda n. \\
 & & \quad let\ clsr(f){=}m\ in\ f\ n\ \kappa\ ow\ wrong)) \\
[\![\texttt{letcc}\, x \,\texttt{in}\, M]\!]\rho\kappa & = & [\![M]\!]\rho[x{:=}\,clsr(\lambda v.\lambda\kappa'.\kappa(v))]\kappa \\
[\![\texttt{ignore}\ M]\!]\rho\kappa & = & [\![M]\!]\rho\kappa
\end{array}$$

A continuation is represented by a closure that ignores its continuation argument [30]. The definition of the relation $v : \tau$ remains essentially the same, ignoring the clause for continuation types. Since `void` is a base type, a value v is of type `void` iff $v = \perp$: there are no terminating values of type `void`.

The proof of the soundness theorem for the system with an empty type relies on two facts. First, we must show that a continuation κ regarded as a closure has type $\tau \to$ `void` whenever $\kappa :: \tau$. Suppose that $v : \tau$ and that $\kappa' ::$ `void`. Then $(\lambda v.\lambda\kappa'.\kappa(v))\, v\, \kappa' = \kappa(v)$. Now since $\kappa :: \tau$ and $v : \tau$, it follows that $\kappa(v) \neq wrong$, as required. Second, we must show that `ignore` M may be assigned an arbitrary type whenever M is of type `void`. But if $\kappa :: \tau$, then $\kappa ::$ `void` since $\perp : \tau$ for any τ, and hence the result follows by induction. For the case of polymorphic continuations, we need only remark that $\forall t.\tau(t)$ is defined by intersection over all monotypes, and note that the above argument shows that if $\kappa :: \tau$, then $\lambda v.\lambda\kappa'.\kappa(v) : \tau{\to}\tau'$ for any type τ' (not just `void`).

4 Extending the Definition of Standard ML

An operational semantics for our language presented in the relational semantics style of [22] may be obtained from the denotational semantics by a process of "defunctionalization" [30] whereby closures and continuations are represented by finitary objects that are then "interpreted" on an argument, rather than simply applied to it.

The "finitary objects" of the operational semantics are defined by the following grammar:

$$\begin{array}{rcl}
V & ::= & \textsc{bas}(B) \mid \textsc{clsr}(x, M, E) \mid \textsc{cnt}(K) \\
K & ::= & \textsc{top} \mid \textsc{farg}(N, E, K) \mid \textsc{app}(V, K) \mid \\
 & & \textsc{targ}(N, E) \mid \textsc{thr}(V) \\
A & ::= & \textsc{valu}(V) \mid \textsc{wrong}
\end{array}$$

where B is a basic value (of unspecified structure), and E is a finite function mapping variables to finitary values.

The dynamic operational semantics is defined in terms of two judgement forms, $E; K \vdash M \Rightarrow A$ and $V \vdash K \Rightarrow A$. The derivation rules for these judgements are as follows:

$$\frac{E(x) \vdash K \Rightarrow A}{E; K \vdash x \Rightarrow A} \qquad (\textsc{o-var})$$

$$\frac{\textsc{clsr}(x, M, E) \vdash K \Rightarrow A}{E; K \vdash \lambda x.M \Rightarrow A} \qquad (\textsc{o-abs})$$

$$\frac{E; \textsc{farg}(N, E, K) \vdash M \Rightarrow A}{E; K \vdash M N \Rightarrow A} \qquad (\textsc{o-apply})$$

$$\frac{E[x{=}\,\textsc{cnt}(K)]; K \vdash M \Rightarrow A}{E; K \vdash \texttt{letcc}\, x\, \texttt{in}\, M \Rightarrow A} \qquad (\textsc{o-letcc})$$

$$\frac{E; \textsc{targ}(N, E) \vdash M \Rightarrow A}{E; K \vdash \texttt{throw}\, M\, N \Rightarrow A} \qquad (\textsc{o-throw})$$

$$V \vdash \textsc{top} \Rightarrow \textsc{valu}(V) \qquad (\textsc{o-top})$$

$$\frac{E; \textsc{app}(V, K) \vdash N \Rightarrow A}{V \vdash \textsc{farg}(N, E, K) \Rightarrow A} \qquad (\textsc{o-farg})$$

$$\frac{E[x{=}V]; K \vdash M \Rightarrow A}{V \vdash \textsc{app}(\textsc{clsr}(x, M, E), K) \Rightarrow A} \qquad (\textsc{o-app})$$

$$\frac{E; \textsc{thr}(V) \vdash N \Rightarrow A}{V \vdash \textsc{targ}(N, E) \Rightarrow A} \qquad (\textsc{o-targ})$$

$$\frac{V \vdash K \Rightarrow A}{V \vdash \textsc{thr}(\textsc{cnt}(K)) \Rightarrow A} \qquad (\textsc{o-thr})$$

Here the initial continuation, designated TOP, is axiomatized as the trivial insertion of values into answers. Other choices are possible.

There is an intriguing parallel between this operational semantics and a call-by-value variant of graph reduction. The idea is that the argument K in $E; K \vdash M \Rightarrow A$ may be thought of as a "marked" spine stack, with the marks indicating whether or not the argument position has been evaluated. Rule O-APPLY "pushes" a node onto the spine stack, marking it as having an unevaluated argument. Rules O-VAR, and O-ABS are "turning points" at which traversal of the expression reverses direction by sending explicit values back up toward the root. This traversal is carried out by the rules for interpreting a continuation. Rules O-FARG and O-TARG cover the case where the argument position is as yet unevaluated, and proceed to evaluate it, marking the node appropriately. Upon return to such a node, the actual application or throw is carried out, either by evaluating the body of the closure in the appropriate environment, or by switching contexts entirely.

The soundness theorem for typing may be proved for the operational semantics by proceeding along much the same lines as for the denotational case. First, we must augment the evaluation relation to include error checking rules that make explicit the notion of "going wrong." These are:

$$\frac{V' \neq \textsc{clsr}(x, M, E)}{V \vdash \textsc{app}(V', K) \Rightarrow \textsc{wrong}} \qquad (\textsc{o-app-x})$$

$$\frac{V' \neq \textsc{cnt}(K')}{V \vdash \textsc{thr}(V', K) \Rightarrow \textsc{wrong}} \qquad (\textsc{o-thr-x})$$

We then define the relations $V : \tau$ and $K :: \tau$ more or less as before, except that instead of relying on the existence of suitable functions in the value space, we appeal directly to the operational semantics.

1. $V : \tau$ holds iff
   - (a) $\tau = b$ and $V = \textsc{valu}(W)$ with W a value of type b;
   - (b) $\tau = \tau_1 \to \tau_2$ and $V = \textsc{clsr}(x, M, E)$ and for every V_1 and K_2, if $V_1 : \tau_1$, $K_2 :: \tau_2$, and $E[x = V_1]; K_2 \vdash M \Rightarrow A$, then $A \neq \textsc{wrong}$;
   - (c) $\tau = \tau_1$ *cont* and $V = \textsc{cnt}(K)$ and $K :: \tau_1$.
2. $K :: \tau$ holds iff for every V such that $V : \tau$, if $V \vdash K \Rightarrow A$, then $A \neq \textsc{wrong}$.

The soundness theorem for the operational semantics is as follows:

Theorem 4.1 *If* $\Gamma \vdash M : \tau$ *and* $E : \Gamma$ *and* $K :: \tau$ *and* $E; K \vdash M \Rightarrow A$, *then* $A \neq \textsc{wrong}$.

In contrast to the domain-theoretic semantics, the proof of the soundness theorem is significantly complicated by the introduction of fixed-point operators for defining recursive functions; see [21] for a careful discussion of a closely-related problem in the setting of natural semantics.

Return-Path: ¡duba@rice.edu¿ Received: from rice.edu by PROOF.ERGO.CS.CMU.EDU id aa03080; 24 Oct 90 17:53:01 EDT Received: from aten.rice.edu by rice.edu (AA07613); Wed, 24 Oct 90 16:52:32 CDT Received: by aten.rice.edu (AA06934); Wed, 24 Oct 90 16:52:54 CDT Date: Wed, 24 Oct 90 16:52:54 CDT From: Bruce Duba ¡duba@rice.edu¿ Message-Id: ¡9010242152.AA06934@aten.rice.edu¿ To: rwh@PROOF.ERGO.CS.CMU.EDU In-Reply-To: Robert.Harper@proof.ergo.cs.cmu.edu's message of Wed, 24 Oct 90 17:38:30 EDT ¡2979.656804310@PROOF.ERGO.CS.CMU.EDU¿ Subject: paper

References

[1] Andrew W. Appel and Trevor Jim. Continuation-passing, closure-passing style. In *Sixteenth ACM Symposium on Principles of Programming Languages*, pages 293–302, Austin, TX, January 1989.

[2] Andrew W. Appel and David MacQueen. Standard ML of New Jersey Reference Manual. (in preparation), 1990.

[3] Dominique Clément, Joëlle Despeyroux, Thierry Despeyroux, Laurent Hascoet, and Gilles Kahn. Natural semantics on the computer. Technical Report RR 416, INRIA, Sophia–Antipolis, France, June 1985.

[4] William Clinger, Daniel P. Friedman, and Mitchell Wand. A scheme for higher-level semantic algebra. In Maurice Nivat and John C. Reynolds, editors, *Algebraic Methods in Semantics*, pages 237–250. Cambridge University Press, Cambridge, 1985.

[5] Eric C. Cooper. An overview of the FOX project. Talk at Bell Laboratories, July 1990.

[6] R. Kent Dybvig and Bob Hieb. Engines from continuations. *Journal of Computer Languages*, 14(2):109–124, 1989.

[7] Arthur Evans. PAL – a language designed for teaching programming linguistics. In *Proc. ACM 23rd National Conference*, pages 395–403, Princeton, 1968. ACM, Brandin Systems Press.

[8] Matthias Felleisen, Daniel Friedman, Eugene Kohlbecker, and Bruce Duba. Reasoning with continuations. In *First Symposium on Logic in Computer Science*. IEEE, June 1986.

[9] Matthias Felleisen, Daniel Friedman, Eugene Kohlbecker, and Bruce Duba. A syntactic theory of sequential control. *Theoretical Computer Science*, 52(3):205–237, 1987.

[10] Matthias Felleisen. The theory and practice of first-class prompts. In *Fifteenth ACM Symposium on Principles of Programming languages*, San Diego, CA, January 1988, pages 180–190.

[11] Andrzej Filinski. Declarative continuations: An investigation of duality in programming language semantics. In *Summer Conference on Category Theory and Computer Science*, volume 389 of *Lecture Notes in Computer Science*, Manchester, UK, 1989. Springer-Verlag.

[12] Andrzej Filinski. Declarative continuations and categorical duality. Master's thesis, University of Copenhagen, Copenhagen, Denmark, August 1989. (DIKU Report 89/11).

[13] Daniel P. Friedman, Christopher T. Haynes, and Eugene Kohlbecker. Programming with continuations. *Program Transformations and Programming Environments*, edited by P. Pepper. Springer-Verlag, Heidelberg, 1985, 263–274.

[14] Timothy Griffin. A formulas-as-types notion of control. In *Seventeenth ACM Symposium on Principles of Programming Languages*, San Francisco, CA, January 1990, pages 47–58.

[15] Christopher T. Haynes, Daniel P. Friedman, and Mitchell Wand. Obtaining coroutines from continuations. *Journal of Computer Languages*, 11:143–153, 1986.

[16] Christopher T. Haynes and Daniel P. Friedman. Embedding continuations in procedural objects. *ACM Transactions on Programming Languages and Systems*, 9:582–598, 1987.

[17] Christopher T. Haynes and Daniel P. Friedman. Abstracting timed preemption with engines. *Journal of Computer Languages*, 12:109–121, 1987.

[18] J. Roger Hindley. The completeness theorem for typing λ terms. *Theoretical Computer Science*, 22:127–134, 1983.

[19] Peter J. Landin. A correspondence between ALGOL–60 and Church's lambda notation. *Communications of the ACM*, 8:89–101, 1965.

[20] Albert R. Meyer and Mitchell Wand. Continuation semantics in typed lambda calculi (summary). In Rohit Parikh, editor, *Logics of Programs*, volume 193 of *Lecture Notes in Computer Science*, pages 219–224. Springer-Verlag, 1985.

[21] Robin Milner and Mads Tofte. Co-induction in relational semantics. Technical Report ECS-LFCS-88-65, Laboratory for the Foundations of Computer Science, Edinburgh University, Edinburgh, October 1988.

[22] Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML.* MIT Press, 1990.

[23] John C. Mitchell. Type systems for programming languages. Technical Report STAN-CS-89-1277, Stanford University, 1989.

[24] Gordon Plotkin. Lambda-definability in the full type hierarchy. In J. P. Seldin and J. R. Hindley, editors, *To H. B. Curry: Essays in Combinatory Logic, Lambda Calculus, and Formalism*, pages 363–373. Academic Press, 1980.

[25] Gordon Plotkin. A structural approach to operational semantics. Technical Report DAIMI–FN–19, Computer Science Department, Aarhus University, 1981.

[26] Norman Ramsey. Concurrent programming in ML. Technical Report CS-TR-262-90, Computer Science Department, Princeton University, 1990.

[27] John Reppy. First-class synchronous operations in Standard ML. Technical Report TR 89-1068, Computer Science Department, Cornell University, Ithaca, NY, December 1989.

[28] John Reppy. Asynchronous signals in Standard ML. Technical Report TR 90-1144, Computer Science Department, Cornell University, Ithaca, NY, August 1990.

[29] John C. Reynolds. GEDANKEN — a simple typeless language based on the principle of completeness and the reference concept. *Communications of the ACM*, 13(5):308–319, May 1970.

[30] John C. Reynolds. Definitional interpreters for higher-order programming languages. In *Conference Record of the 25th National ACM Conference*, pages 717–740, Boston, August 1972. ACM.

[31] John C. Reynolds. On the relation between direct and continuation semantics. In *Proc. ICALP '74*, volume 14 of *Lecture Notes in Computer Science*, pages 141–156. Springer-Verlag, 1974.

[32] Dorai Sitaram and Matthias Felleisen. Reasoning with continuations II: Full abstraction for models of control. In *Proc. 1990 Conference on LISP and Functional Programming*, pages 161–175, June 1990.

[33] Richard Statman. Logical relations and the typed λ-calculus. *Information and Control*, 65:85–97, 1985.

[34] Guy L. Steele Jr. RABBIT: A compiler for SCHEME. Technical Report Memo 474, MIT AI Laboratory, 1978.

[35] Christopher Strachey and Christopher Wadsworth. A mathematical semantics for handling full jumps. Technical Report Technical Monograph PRG–11, Oxford University Computing Laboratory, 1974.

[36] Bernard Sufrin. CSP-style processes in ML. (Private communication), 1989.

[37] Gerald Jay Sussman and Guy Lewis Steele Jr. Scheme: An interpreter for extended lambda calculus. Technical Report Memo No. 349, MIT AI Laboratory, December 1975.

[38] Carolyn Talcott. Rum: An intensional theory of function and control abstractions. In *Proc. 1987 Workshop of Foundations of Logic and Functional Programming*, volume 306 of *Lecture Notes in Computer Science*. Springer-Verlag, 1988.

[39] Mads Tofte. *Operational Semantics and Polymorphic Type Inference.* PhD thesis, Edinburgh University, 1988. Available as Edinburgh University Laboratory for Foundations of Computer Science Technical Report ECS–LFCS–88–54.

[40] Mitchell Wand. Continuation-based multiprocessing. In *Conference Record of the 1980 LISP Conference*, pages 19–28, 1980.

[41] Andrew Wright and Matthias Felleisen. The nature of exceptions in polymorphic languages. Unpublished manuscript, Rice University, 1990.