

How to Add Threads to a Sequential Language Without Getting Tangled Up

Martin Gasbichler Eric Knauer
Institut für Informatik
Universität Tübingen

{gasbichl,knauer}@informatik.uni-tuebingen.de

Michael Sperber*
sperber@deinprogramm.de

Richard A. Kelsey
Ember Corporation
kelsey@s48.org

Abstract

It is possible to integrate Scheme-style first-class continuations and threads in a systematic way. We expose the design choices, discuss their consequences, and present semantical frameworks that specify the behavior of Scheme programs in the presence of threads. While the issues concerning the addition of threads to Scheme-like languages are not new, many questions have remained open. Among the pertinent issues are the exact relationship between continuations and the `call-with-current-continuation` primitive, the interaction between threads, first-class continuations, and `dynamic-wind`, the semantics of dynamic binding in the presence of threads, and the semantics of thread-local store. Clarifying these issues is important because the design decisions related to them have profound effects on the programmer's ability to write modular abstractions.

1 What's in a Continuation?

Scheme [21] was one of the first languages to endorse `call-with-current-continuation` as a primitive. `Call-with-current-continuation` (or `call/cc`, for short) is an essential ingredient in the implementation of a wide range of useful abstractions, among them non-local control flow, exception systems, coroutines, non-deterministic computation, and Web programming session management. So much is often repeated, non-controversial and clear.

Nowadays, even the name `call-with-current-continuation` is confusing. It suggests erroneously that `call/cc` applies its argument to a reified version of the current continuation—the meta-level object the underlying machine uses to remember what should happen with the value of the expression currently being evaluated.

*For parts of this work, Michael Sperber was supported by the Institut für Informatik at Universität Tübingen.

Permission to make digital or hard copies, to republish, to post on servers or to redistribute to lists all or part of this work is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To otherwise copy or redistribute requires prior specific permission.

Fourth Workshop on Scheme and Functional Programming. November 7, 2003, Boston, Massachusetts, USA.

Copyright 2003 Martin Gasbichler, Eric Knauer, Michael Sperber and Richard Kelsey.

The denotational semantics presented in R^5RS [21] supports this impression. Here is a slightly simplified version:

$$cwcc : E^* \rightarrow K \rightarrow C \quad [\text{call-with-current-continuation}]$$
$$cwcc = \text{onearg } (\lambda \varepsilon \kappa . \text{apply } \varepsilon ((\lambda \varepsilon^* \kappa' . \kappa \varepsilon^*) \text{ in } E) \kappa)$$

The reified value passed to the argument of `cwcc` is the function $\lambda \varepsilon^* \kappa' . \kappa \varepsilon^*$ —essentially an eta-expanded version of κ , the current continuation as handled by the semantics. Calling this function merely re-installs or *reflects* κ as the current continuation. With this definition, the distinction between the *escape procedure*—the procedure passed to `call/cc`'s argument and the actual meta-level continuation is largely academic.

Unfortunately, the semantics for `call/cc` given in R^5RS is not correct, as noted in a “Clarifications and corrections” appendix to the published version: an R^5RS -compliant `call/cc` must also execute *thunks* along the branches of the control tree as introduced by the `dynamic-wind` primitive [18] added to Scheme in R^5RS . Even in pre- R^5RS Scheme, the escape procedure would typically re-install previously captured values for the current input and output ports. Thus, the escape procedure created by `call/cc` performs actions *in addition* to installing a captured continuation. Hence, the name `call-with-current-continuation` is misleading.

`Dynamic-wind` allows enhancing and constraining first-class continuations: (`dynamic-wind before thunk after`) calls *thunk* (a procedure of no parameters), ensuring that *before* (also a *thunk*) is always called before the program enters the application of *thunk*, and that *after* is called after the program has left it. Therefore, escape procedures created by `call/cc` must also call the *after* and *before* *thunks* along the paths leading from the current node in the control tree to the target tree. This creates a significant distinction between an escape procedure and its underlying continuation.

This distinction has created considerable confusion: Specifically, continuations are suitable abstractions for building thread systems [37], and this suggests that escape procedures are, too. However, a thread system based on R^5RS `call/cc` will run *before* and *after* *thunks* introduced by `dynamic-wind` upon every context switch, which leads to semantic and pragmatic problems in addition to the common conceptual misunderstandings noted by Shivers [32]. Moreover, other common abstractions, such as dynamic binding and thread-local storage, interact in sometimes surprising ways with threads and first-class continuations, depending on their exact semantics in a given system.

Thus, the integration of first-class continuations with `dynamic-wind`, concurrency and parallelism, along with associated functionality such as dynamic binding and thread-local storage form a puzzle: Most of the pieces have long been on the table, but there is little published documentation on how all of them fit together in a systematic way, which often causes confusion for users and implementors alike. With this paper, we try to make the pieces fit, and close some of the remaining gaps.

Here are the contributions of our work:

- We discuss some of the pertinent semantic properties of `dynamic-wind`, specifically as they relate to the implementation of dynamic binding.
- We discuss design issues for thread systems in Scheme-like languages, and how different design choices affect program modularity.
- We present a systematic treatment of two abstractions for thread-aware programming: `thread-wind` extends the context switch operation, and `thread-local` storage implements extensible processor state.
- We present a denotational semantics of R⁵RS `call/cc` and `dynamic-wind`.
- We clarify the relationship between threads and `call/cc/dynamic-wind` by presenting a transition semantics based on the CEK machine [6] equivalent to the denotational semantics, and extending this semantics by simple models for threads and multiprocessing.

Overview: Section 2 gives an account of `call/cc` as present in (sequential) Scheme, and its interaction with `dynamic-wind`. Section 3 lists some specific design issues pertinent to the addition of threads to Scheme and describes their impact on the ability to write modular programs. More issues arise during implementation; Section 4 discusses these. Section 5 describes facilities for thread-aware programming. Section 6 presents semantic characterizations of Scheme with `dynamic-bind` and threads. Related work is discussed in Section 7; Section 8 concludes.

2 Call/cc As We Know It

In this section, we give an informal overview of the behavior of the R⁵RS Scheme version of `call/cc`. Specifically, we discuss the interaction between `call/cc` and the current dynamic environment implicit in R⁵RS, and the interaction between `call/cc` and `dynamic-wind`. We also explain how these interactions affect possible implementations of an extensible dynamic environment.

2.1 The current dynamic environment

R⁵RS [21] implies the presence of a *current dynamic environment* that contains bindings for the current input and output ports. Scheme’s I/O procedures default to these ports when they are not supplied explicitly as arguments. Also, the program can retrieve the values of the bindings via the `current-input-port` and `current-output-port` procedures. “Dynamic” in this context means that the values for the program behave as if the current dynamic environment were implicitly passed as an argument with each procedure application. In this interpretation, `with-input-from-file` and `with-output-to-file` each call its argument with a newly created dynamic environment containing a new binding, and `current-{input,output}-port` retrieve

```
(define *dynamic-env* (lambda (v) (cdr v)))

(define (make-fluid default) (cons 'fluid default))

(define (fluid-ref fluid) (*dynamic-env* fluid))

(define (shadow env var val)
  (lambda (v)
    (if (eq? v var)
        val
        (env var))))

(define (bind-fluid fluid val thunk)
  (let ((old-env *dynamic-env*)
        (new-env (shadow *dynamic-env* fluid val)))
    (set! *dynamic-env* new-env)
    (let ((val (thunk)))
      (set! *dynamic-env* old-env)
      val)))
```

Figure 1. Dynamic binding via dynamic assignment

the values introduced by the most recent, still active application of these procedures. The interpretation of the current dynamic environment as an implicit argument means that dynamic environments are effectively associated with continuations. Specifically, reflecting a previously reified continuation also means returning to the dynamic environment which was current at the time of the reification.¹

It is often useful to be able to introduce new dynamic bindings [24, 16] in addition to `current-{input,output}-port`, for example to implement exception handling. However, as the dynamic environment is implicit (and not reifiable), a program cannot extend it. Fortunately, it is possible to simulate extending the dynamic environment with first-class procedures by keeping the current dynamic environment in a global variable, and simply save and restore it for new bindings—a technique known as *dynamic assignment* [11].

Figure 1 shows naive code for dynamic assignment. `make-fluid` creates a fluid represented as a pair consisting of the symbol `fluid` as its `car` and the default value as its `cdr`. `*dynamic-env*` holds the current dynamic environment, represented as a procedure mapping a fluid to its value. The initial function in `*dynamic-env*` extracts the default value of a fluid. `shadow` makes a new dynamic environment from an old one, shadowing one binding with a new one. `bind-fluid` remembers the old value of `*dynamic-env*`, sets it to a new one created by `shadow`, calls `thunk`, and restores the old value. (The code ignores the issue of multiple return values for simplicity.) The `fluid-ref` procedure looks up a fluid binding in the dynamic environment, returning its value.

Unfortunately, `bind-fluid` does not implement the implicit-argument semantics in the presence of `call/cc`: it is possible for the `thunk` argument to `bind-fluid` to reflect a previously reified continuation which will then inherit the current dynamic environment, rather than the dynamic environment current at the time of reification. For implementing the implicit-argument semantics, it is necessary to capture the current value of `*dynamic-env*` at the time of reification, and re-set it to that value upon reflection.

¹Note that this behavior is not mandated by R⁵RS. However, existing Scheme code often assumes it [23].




Figure 2. Control tree and dynamic-wind

2.2 Dynamic-wind

While the naive implementation of dynamic assignment does not have the desirable semantics, it is possible to implement a version that does, via the Scheme primitive `dynamic-wind`. (Very roughly, `(dynamic-wind before thunk after)` ensures that *before* is called before every control transfer into the application of `thunk`, and *after* is called after every control transfer out of it. Here is a new version of `bind-fluids` that utilizes `dynamic-wind` to get the correct semantics:

```
(define (bind-fluid fluid val thunk)
  (let ((old-env *dynamic-env*)
        (new-env (shadow *dynamic-env* fluid val)))
    (dynamic-wind
     (lambda () (set! *dynamic-env* new-env))
     thunk
     (lambda () (set! *dynamic-env* old-env)))))
```

The behavior of `dynamic-wind` is based on the intuition that the continuations active in a program which uses `call/cc` form a tree data structure called the *control tree* [18]: each continuation corresponds to a singly-linked list of frames, and the continuations reified by a program may share frames with each other and/or with the current continuation. Reflecting a previously reified continuation means making a different node of the tree the current continuation. A Scheme program handles the current control tree node in much the same way as the dynamic environment. Together, they constitute the *dynamic context*.

Conceptually, `(dynamic-wind before thunk after)` annotates the continuation of the call to *thunk* with *before* and *after*. Calling an escape procedure means travelling from the current node in the control tree to the node associated with the previously reified continuation. This means ascending from the current node to the nearest common ancestor of the two nodes, calling the *after* thunks along the way, and then descending down to the target node, calling the *before* thunks. Figure 2 shows such a path in the control tree.

Using `dynamic-wind` for implementing dynamic binding assures that part of the global state—the value of `*dynamic-env*`, in this case—is set up to allow the continuation to run correctly. This works well for dynamic binding, as changes to `*dynamic-env*` are always easily reversible. However, in some situations a continuation might not be able to execute correctly because global state has changed in an irreversible way. Figure 3 shows a typical code fragment which employs `dynamic-wind` to ensure that the program will close an input port immediately after a set of port operations

```
(let ((port (open-input-file file-name)))
  (dynamic-wind
   (lambda ()
     (if (not port)
         (error "internal error"))))
   (lambda () (read from port))
   (lambda ()
     (close-input-port port)
     (set! port #f))))
```

Figure 3. Restricting the use of escape procedures

has completed (in the *after* thunk) as well as preventing the program from inadvertently entering the code that performs file I/O after the close has happened. Moreover, the *before* thunk prevents the port access code from being re-entered because the port operations are likely to have caused irreversible state changes.² Thus, three main uses for `dynamic-wind` emerge [18]:

1. extending the dynamic context associated with continuations (as in `bind-fluid`)
2. releasing resources used by a region of code after that code has completed (as in Figure 3)
3. preventing the reification of a continuation because its dynamic context cannot be recreated (as in Figure 3)

Item #2 is akin to the default or `finally` clauses of exception handling systems or to the `unwind-protect` facilities in some languages. The unlimited extent of escape procedures created by `call/cc` makes the more general `dynamic-wind` necessary.

The presence of `dynamic-wind` requires a more careful handling of terminology when it comes to continuations: We call the process of turning the meta-level continuation into an object-level value *reification*, and the reverse—re-installing a previously reified continuation—*reflection*. The process of creating an escape procedure (by `call/cc`) is a *capture*; this includes reifying the current continuation. Conversely, *invoking* the escape procedure travels to the target point in control space, installs the dynamic environment, and then reflects the continuation.

3 Design Requirements for Thread Systems

In this section, we consider some of the design issues that arise when adding threads to a higher-order language. We assume that the thread system features a `spawn` operation. `Spawn` starts a new thread and calls *thunk* (a thunk) in that new thread. The thread terminates once *thunk* returns:

```
(spawn thunk) procedure
```

The presence of `spawn` in a language with `call/cc`, `dynamic-wind`, and dynamic binding exposes a number of language design choices, as these features interact in potentially subtle ways. Specifically, the ability to migrate continuations between threads, and the interaction between dynamic binding and threads fundamentally affect the ability to write modular programs.

3.1 Migrating continuations

A Scheme program can invoke an escape procedure in a thread different from the one where it was captured. Notably, this scenario

²Even though it might be possible to redo changes on a file port, this is usually impossible with, say, a network connection.

occurs in multithreaded web servers which use `call/cc` to capture the rest of a partially completed interaction between the server and a client: typically, the server will create a new thread for each new request and therefore must be able to invoke the escape procedure that was captured in the thread which handled the connection belonging to the previous step in the interaction [28].

In MrEd, the Scheme platform on which PLT’s web server is based, continuations are “local to a thread”—only the thread that created an escape procedure can invoke it, forcing the web server to associate a fixed thread with a session [14].³ While this may seem like a technical restriction with a purely technical solution, this scenario exposes serious general modularity issues: Modules may communicate escape procedures, and tying an escape procedure to a thread restricts the implementation choices for a client which needs to invoke an escape procedure created by another module. If the escape procedure is thread-local, the client cannot even tell if invoking it might make the program fail; all it knows is that the invocation will *definitely* fail if performed in a freshly created thread.

Once continuations are allowed to migrate between threads, additional questions arise. In particular, the use of certain abstractions might make the continuation sensitive to migration, which is usually not what the programmer intended.

3.2 Dynamic binding and the thread system

Consider the following program fragment:

```
(define f (make-fluid 'foo))

(bind-fluid f 'bar
 (spawn
  (lambda ()
    (display (fluid-ref f))))))
```

Should the program print `foo` or should it print `bar`? This is a well-known design issue with thread systems [13]. The general question is this: Should a newly spawned thread inherit the dynamic environment from the original thread—or, more precisely, from the continuation of the call to `spawn`—or should it start with an empty dynamic environment, assuming the default values for all fluids?⁴

For at least two dynamic entities, inheritance does not make sense: the current control tree node and, if present, the current exception handler, as they both conceptually reach back into the part of the control tree belonging to the original thread. Thus, it is unclear what should happen if the new thread ever tries to travel back to that part of the tree. (For `dynamic-wind`, we discuss another closely related issue in Section 4.1.) Instead, a newly spawned thread must start with a fresh current exception handler and an empty control tree.

For all other dynamic bindings, it is unclear whether a single inheritance strategy will satisfy the needs of all programs. For many entities typically held in fluids, it makes sense for a new thread to inherit dynamic bindings from the thread which spawned it:

- Scsh [31], tries to maintain an analogy between threads and Unix processes, and keeps Unix process resources in fluids [13]. In Scsh, a special `fork-thread` operation acts like

³This restriction will be lifted in a future version of MrEd.

⁴The issue becomes more subtle with SRFI-18-like thread systems [5] with separate `make-thread` and `thread-start!` operations. Whose dynamic environment should the new thread inherit?

```
(define (current-dynamic-context)
  (let ((pair (call-with-current-continuation
              (lambda (c) (cons #f c)))))
    (if (car pair)
        (call-with-values (car pair) (cdr pair))
        (cdr pair))))

(define (with-dynamic-context context thunk)
  (call-with-current-continuation
   (lambda (done)
     (context (cons thunk done)))))

(define (spawn thunk)
  (let ((context (current-dynamic-context)))
    (spawn
     (lambda ()
       (with-dynamic-context context thunk)))))
```

Figure 4. Reifying and reflecting the dynamic context

`spawn`, but has the new thread inherit the values of the process resources from the original thread.

- MzScheme [9] provides abstractions for running a Scheme program in a protected environment, thus providing operating-system-like capabilities [10]. Some of the entities controlling the encapsulation of such programs are held in fluids (called *parameters* in MzScheme), such as the current custodian that controls resource allocation and destruction. Children of an encapsulated thread inherit the custodian of the parent so that shutting down the custodian will kill the encapsulated thread along with all of its children.
- Generally, programmers might expect (`spawn f`) to behave as similarly as possible to (`f`). This is especially likely if the programmer uses threads to exploit parallelism, in a similar way to using futures [15], and thus merely wants to offload parts of the computation to a different processor.

3.3 Dynamic binding and modularity

The issue of fluid inheritance is most pertinent when a program module keeps mutable data in fluids. Specifically, consider the following scenario: Program module A creates and uses fluids holding mutable state. The fluids might be exported directly, or module A might provide `with-f` abstractions roughly like the following:

```
(define f (make-fluid default))

(define (with-f value thunk)
  (bind-fluid f (... value ...) thunk))
```

A client of module A might want to create multiple threads, and use the abstractions of module A from several of them. Generally, the client might need to control the sharing of state held in `f` for each new thread it creates in the following ways:

1. getting A’s default dynamic bindings,
2. creating a new binding for A by using `with-f` in the new thread, or
3. inheriting the current thread’s dynamic environment.

If each thread starts up with a fresh dynamic environment, this degree of control is available:

1. Starting a new thread with a fresh dynamic environment means that it will get default bindings for all fluids.

2. Explicitly creating new bindings is possible via `with-f`.
3. It is still possible to implement a variant of `spawn` that does cause the new thread to inherit the dynamic environment from the thread which created it.

Figure 4 shows how to achieve the last of these: As `call/cc` captures the dynamic context, it is possible to reify and reflect it, along with the dynamic environment, through escape procedures. `Current-dynamic-context` uses `call/cc` to create an escape procedure associated with the current dynamic context, and packages it up as the `cdr` of a pair. The `car` of that pair is used to distinguish between a normal return from `call/cc` and one from `with-dynamic-context` which runs a thunk with the original continuation—and, hence, the original dynamic context—in place and restores its own continuation after the thunk has finished. With the help of these two abstractions, `spoon` (for a “fluid-preserving fork operation for threads,” a term coined by Alan Bawden) starts a new thread which inherits the current dynamic environment.⁵

Note that `spoon` causes the new thread to inherit the *entire* dynamic context, including the current control tree node, and the current exception handler (if the Scheme system supports exception handling in the style of ML.) This can lead to further complications [1]. Also, inheritance is not the only possible solution to the security requirements of MrEd: Thread systems based on nested engines [2] such as that of Scheme 48 allow defining custom schedulers. Here, a scheduler has full control over the initial dynamic environment of all threads spawned with it.

4 Implementing Concurrency

The previous section has already stated some of the design requirements and choices for an implementation of threads in a language with first-class continuations. Additional issues emerge when actually implementing threads in the presence of `call/cc` and `dynamic-wind`. In particular, many presentations of thread systems build threads *on top* of the language, using `call/cc` to implement the context switch operation. However, this choice incurs undesirable complications (especially in the presence of multiprocessing) when compared to the alternative—implementing threads primitively and building the sequential language on top.

4.1 Dynamic-wind vs. the context switch

The presence of `dynamic-wind` makes `call/cc` less suitable for implementing context-switch-like abstractions like coroutines or thread systems: Uses of `dynamic-wind` may impose restrictions on the use of the escape procedures incompatible with context switching.⁶ Consider the code from Figure 3. This code should continue to work correctly if run under a Scheme system with threads—say, in thread X. However, if the context switch operation of thread system is implemented using ordinary `call/cc`, each context switch out of thread X means ascending up the control tree to the scheduler

⁵The same trick is applicable to promises which exhibit the same issues, and which also do not capture the dynamic context: the fluid-preserving versions of `delay` and `force` would be called `freeze` and `thaw`.

⁶Note that this is an inherent issue with the generality of `call/cc`: `Call/cc` allows capturing contexts which simply are not restorable because they require access to non-restorable resources. Providing a version of `call/cc` which does not capture the dynamic context would violate the invariants guaranteed by `dynamic-wind` and break most code which uses it.

(whose continuation frames constitute the shared part of tree)—executing all *after* thunks of all `dynamic-wind` operations active within the current thread. The next context switch back into thread X will then run the *before* thunks, which in this case will make the program fail. Naturally, this is unacceptable.

Moreover, if every context switch would run `dynamic-wind before` and *after* thunks, the program would expose the difference between a virtualized thread system running on a uniprocessor and a multiprocessor where multiple threads can be active without any context switch: If each thread ran on a different processor, no continuations would ever be captured or invoked for a context switch, so a context switch would never cause `dynamic-wind` thunks to run.

Thus, building a thread system on top of R⁵RS `call/cc` leads to complications and invalidates common uses of `dynamic-wind`. (Similar complications occur in the presence of ML-style exception handling [1].) Hence, a more reasonable approach for implementations is to build threads natively into the system, and build `call/cc` and `dynamic-wind` on top of it. In this scenario, each newly spawned thread starts with an empty dynamic context.

4.2 Dynamic binding vs. threads

In the presence of threads, the implementation of dynamic binding that keeps the current dynamic environment in a global variable no longer works: all threads share the global variable, and, consequently, any application of `bind-fluid` is visible in other threads, violating the intended semantics. Therefore, it is necessary to associate each thread with its own dynamic environment. Here are some possible implementation strategies:

1. pass the dynamic environment around on procedure calls as an implicit argument
2. keep looking for dynamic bindings in the `*dynamic-env*` global variable, and change the value of this variable upon every context switch, always setting it to the dynamic environment associated with the current thread
3. like #2, but keep the dynamic environment in the thread data structure, and always access that instead of a global variable

#1 incurs overhead for every single procedure call; considering that access and binding of fluid variables is relatively rare, this is an excessive cost rarely taken by actual implementations. #2 is incompatible with multiprocessing, as multiple threads can access fluid variables without intervening context switches. #3 is viable.

All of these strategies require what is known as “deep binding” in the Lisp community—`fluid-ref` always looks up the current value of a fluid variable in a table, and only reverts to the top-level value stored in the fluid itself when the table does not contain a binding. Many Lisp implementations have traditionally favored “shallow binding” that manages dynamic bindings by mutating the fluid objects themselves. With shallow binding, access to a fluid variable is simply dereferencing the fluid object; no table searching is necessary. However, this technique is also fundamentally incompatible with multiprocessing because it mutates global state.

4.3 Virtual vs. physical processors

The previous two sections have shown that a multiprocessor thread system can potentially expose differences in implementation strategies for dynamic binding, as well as different ways of dealing with

`dynamic-wind`. These differences all concern the notion of “what a thread is”—specifically, if a thread encompasses the dynamic context, or if it is an exterior, global entity.

A useful analogy is viewing a thread as a virtual processor [32] running on a physical processor. In this view, the dynamic context and the dynamic environment are akin to processor registers. In a multiprocessor implementation of threads, each physical processor indeed must keep those values in locations separate from that of the other processors. Each of these processors can then run multiple threads, swapping the values of these registers on each context switch. (This corresponds to Shivers’s notion of “continuation = abstraction of processor state” as the entity being swapped upon a context switch [32].) In this model, a thread accessing these registers cannot distinguish whether it is running in a uniprocessor or a multiprocessor system.

5 Thread-Aware Programming

The previous two sections have focused on protecting sequential programs from the adverse effects resulting from the presence of threads, and on decoupling previously present sequential abstractions such as `dynamic-wind` and dynamic binding from the thread system as far as possible. However, the implementations of low-level abstractions occasionally benefit from access to the guts of the thread system. Two abstractions provide this access in a systematic way: the `thread-wind` operation allows running code local to a thread upon context-switch operations, and *thread-local cells* are an abstraction for managing thread-local storage. However, the use of these facilities requires great care to avoid unexpected pitfalls.

5.1 Extending the context switch operation

Accessing state like the `dynamic-wind` context or the dynamic environment through processor registers is convenient and fast. However, as the scheduler needs to swap the values of these registers on each context switch, they are not easily extensible: each new register requires an addition to the context-switch operation. Also, it is occasionally desirable that a thread is able to specify code to be run whenever control enters or exits that thread, thus making the context switch operation extensible. (Originally, `dynamic-wind` had precisely that purpose, but, as pointed out in Section 4.1, this is not reasonable in light of current usage of `dynamic-wind`.) Therefore, we propose a new primitive:

```
(thread-wind before thunk after)           procedure
```

In a program with only a single thread, `thread-wind` acts exactly like `dynamic-wind`: *before*, *thunk*, and *after* are thunks; they run in sequence, and the `thread-wind` application returns whatever *thunk* returns. Moreover, *before* gets run upon each control transfer into the application *thunk*, and *after* gets run after each transfer out of it. Unlike with `dynamic-wind`, however, during the dynamic extent of the call to *thunk*, every context switch out of the thread runs the *after* thunk, and every context switch back in runs the *before* thunk.

`Thread-wind` is a low-level primitive; its primary intended purpose is to control parts of the processor state not managed by the underlying, primitive thread system. For example, in a uniprocessor setting, it is possible to continue treating the variable `*dynamic-env*` as a sort of register, and implement `bind-fluid` correctly by using `thread-wind` instead of `dynamic-wind`:

```
(define (bind-fluid fluid val thunk)
```

```
(let ((old-env *dynamic-env*)
      (new-env (shadow *dynamic-env* fluid val)))
  (thread-wind
   (lambda () (set! *dynamic-env* new-env))
   thunk
   (lambda () (set! *dynamic-env* old-env))))))
```

The semantics of `thread-wind` extends smoothly to the escape procedure migration scenario: in this case, before the program installs the new continuation, it runs the active `thread-wind` *after* thunks of the current thread, and the active *before* thunks of the continuation being reflected.

Ideally, the *before* and *after* thunks are transparent to the running thread in the sense that running *after* invalidates whatever state changes *before* has performed. Still, it is possible to use `thread-wind` to set up more intrusive code to be run on context switches, such as profiling, debugging, or benchmarking.

5.2 Thread-local storage

The version of `bind-fluid` using `thread-wind` still is not correct in the presence of multiprocessing, as all processors share the value of `*dynamic-env*`. For correctly implementing dynamic binding, another conceptual abstraction is needed: *thread-local storage*. Thread-local storage is available through *thread-local cells* or *thread cells* for short. Here is the interface to thread-local cells:

```
(make-thread-cell default)                 procedure
(thread-cell-ref thread-cell)             procedure
(thread-cell-set! thread-cell value)      procedure
```

`Make-thread-cell` creates a reference to a thread cell with default value *default*, `thread-cell-ref` fetches its current value, and `thread-cell-set!` sets it. Any mutations of a thread cell are only visible in the thread which performs them. A thread cell acts like a table associating each thread with a value which defaults to *default*; `thread-cell-ref` accesses the table entry belonging to the current thread, and `thread-cell-set!` modifies it.

With thread cells, it is possible to implement dynamic binding correctly in the presence of multiprocessing: `*dynamic-env*`, instead of being bound directly to the environment, is now a thread cell:

```
(define *dynamic-env*
  (make-thread-cell (lambda (v) (cdr v))))

(define (make-fluid default) (cons 'fluid default))

(define (fluid-ref fluid)
  ((thread-cell-ref *dynamic-env*) fluid))

(define (bind-fluid fluid val thunk)
  (let ((old-env (thread-cell-ref *dynamic-env*))
        (new-env (shadow
                  (thread-cell-ref *dynamic-env*)
                  fluid val)))
    (dynamic-wind
     (lambda ()
       (thread-cell-set! *dynamic-env* new-env))
     thunk
     (lambda ()
       (thread-cell-set! *dynamic-env* old-env))))))
```

5.3 Modularity issues

While thread-local storage is a useful low-level abstraction, its use in programs imposes restrictions which may have an adverse effect on modularity. Consider the scenario from Section 3.3 with “dynamic binding/environment” replaced by “thread-local storage”: module *A* creates and uses thread-local cells. This makes it much harder and potentially confusing for the client to use threads and control the sharing of among them. Here are the three choices for dynamic binding, revisited for thread-local storage:

1. New threads get a fresh thread-local store with default values for the thread-local variables—in this respect, they behave similarly to dynamic bindings.
2. Since thread-local storage is specifically not about binding, a *with-f*-like abstraction may not be feasible.
3. Inheritance of the thread-local storage is not easily possible for a new thread, as escape procedures do not capture the thread-local store.

Thus, if module *A* uses the thread-local store, the client has essentially no control over how *A* behaves with respect to the threads. This is unfortunate as the client might use threads for any number of reasons that in turn require different sharing semantics.

Especially the migration of escape procedures between threads raises troublesome questions with no obvious answer: As the escape procedure does not install the thread-local store from the thread which reified it, a solution to option #3—unsharing module *A*’s state between the old and the new thread—becomes impossible. On the other hand, if the escape procedure were closed over the thread-local store, it would need to capture a *copy* of the store—otherwise, the name “thread-local storage” would be inappropriate, and the ensuing sharing semantics would carry more potential for confusion and error. Capturing the copy raises the next question: At what time should the program create the copy? At the time of capture, at the time of creating the new thread, or at the time of invocation of the escape procedure? The only feasible solution to the dilemma would be to make the thread-local store itself reifiable. However, it is unclear whether this abstraction would have benefits that outweigh the potential for confusion, and the inflexibility of abstractions which use thread-local storage in restricting ways.

Note that none of these problems manifest themselves in the implementation of dynamic binding presented in the previous section: the *dynamic-wind* thunks ensure that the **dynamic-env** thread-local-cell always holds the dynamic environment associated with the current continuation. Consequently, it seems that thread-local storage is a natural means for building other (still fairly low-level) abstractions such as dynamic binding, but rarely appropriate for use in higher-level abstractions or in applications.

6 Semantics

This section provides semantic specifications for a subset of Scheme with *dynamic-wind* and threads. We start with a version of the R^5RS denotational semantics which describes the behavior of *dynamic-wind*. We then formulate a transition semantics equivalent to the denotational semantics, which in turn forms the basis for a semantics for a concurrent version of the Scheme subset. This concurrent semantics specifies the interaction between *dynamic-wind* and threads. (We have also formulated a semantics which accounts for multiprocessing and for *thread-wind* which we have relegated to Appendix A. The appendix also contains an

augmented version of the entire R^5RS semantics.) Moreover, we present a version of the denotational semantics with an explicit dynamic environment, and show that implementing the dynamic environment indirectly with dynamic assignment and *dynamic-wind* is indeed equivalent to propagating it directly in the semantics, thus demonstrating the utility of the semantics.

For the definition of our subset of Scheme, Mini-Scheme, we employ the same terminology, and, where possible, the same notation as R^5RS . (See Appendix D for details.) As compared to the language covered by the R^5RS semantics, a procedure has a fixed number of parameters and returns a single value, a procedure body consists of a single expression, procedures do not have an identifying location, evaluation is always left-to-right, and *if* forms always specify both branches. Mini-Scheme does, however, feature assignment, *call-with-current-continuation* and *dynamic-wind*. Here is the expression syntax of Mini-Scheme:

$$\text{Exp} \longrightarrow \text{K} \mid \text{I} \mid (\text{E}_0 \text{ E}^*) \mid (\text{lambda } (\text{I}^*) \text{ E}_0) \\ \mid (\text{if } \text{E}_0 \text{ E}_1 \text{ E}_2) \mid (\text{set! } \text{I } \text{E})$$

6.1 Denotational semantics

The semantic domains are analogous to those in R^5RS with changes according to the restrictions of Mini-Scheme—expression continuations always take one argument. The definition of E^* now needs special multi-argument *argument continuations*.

$\phi \in F = (E^* \rightarrow P \rightarrow K \rightarrow C)$	procedure values
$\kappa \in K = E \rightarrow C$	expression continuations
$\kappa' \in K' = E^* \rightarrow C$	argument continuations
$\omega \in P = (F \times F \times P) + \{\text{root}\}$	dynamic points

In addition, *P* is the domain for *dynamic points* which are nodes in the control tree: *root* is the root node, and all other nodes consist of two thunks and a parent node. Figure 5 shows the semantics for Mini-Scheme expressions. It is completely analogous to the R^5RS version of the \bar{E} function; the only addition is the propagation of the current dynamic point. The auxiliary functions are analogous to their R^5RS counterparts, apart from a change in *apply* to take dynamic points into account:

$$\text{apply} : E \rightarrow E^* \rightarrow P \rightarrow K \rightarrow C$$

$$\text{apply} = \lambda \varepsilon \omega \kappa . \varepsilon \in F \rightarrow (\varepsilon \mid F) \varepsilon^* \omega \kappa, \text{wrong "bad procedure"}$$

Here is a version of the *cwcc* primitive implementing *call-with-current-continuation* which respects *dynamic-wind*:

$$\text{cwcc} : E^* \rightarrow P \rightarrow K \rightarrow C \quad [\text{call-with-current-continuation}]$$

$$\text{cwcc} =$$

$$\text{onearg } (\lambda \varepsilon \omega \kappa . \varepsilon \in F \rightarrow$$

$$\quad (\text{apply } \varepsilon$$

$$\quad \quad \langle (\lambda \varepsilon^* \omega' \kappa' .$$

$$\quad \quad \quad \text{travel } \omega' \omega (\kappa (\varepsilon^* \downarrow 1))) \text{ in } E)$$

$$\quad \quad \omega \kappa),$$

$$\quad \text{wrong "bad procedure argument"})$$

The escape procedure captures the dynamic point, and, when called, “travels” from the current dynamic point to it, running the *after* and *before* thunks in the process, before actually installing the continuation. Here is the definition of *travel*:

$$\text{travel} : P \rightarrow P \rightarrow C \rightarrow C$$

$$\text{travel} = \lambda \omega_1 \omega_2 . \text{travelpath } (\text{path } \omega_1 \omega_2)$$

The *travelpath* function performs the actual travelling along a sequence of thunks and dynamic points, running each thunk with the

$$\begin{array}{l}
E : \text{Exp} \rightarrow \mathcal{U} \rightarrow \mathcal{P} \rightarrow \mathcal{K} \rightarrow \mathcal{C} \\
E^* : \text{Exp}^* \rightarrow \mathcal{U} \rightarrow \mathcal{P} \rightarrow \mathcal{K}' \rightarrow \mathcal{C} \\
E[\![\mathcal{K}]\!] = \lambda \rho \omega \kappa . \text{send}(\mathcal{K}[\![\mathcal{K}]\!]) \kappa \\
E[\![\mathcal{I}]\!] = \lambda \rho \omega \kappa . \text{hold}(\text{lookup } \rho \mathcal{I}) \\
\quad (\lambda \varepsilon . \varepsilon = \text{undefined} \rightarrow \\
\quad \quad \text{wrong "undefined variable",} \\
\quad \quad \text{send } \varepsilon \kappa) \\
E[\![\text{if } E_0 \ E_1 \ E_2]\!] = \\
\quad \lambda \rho \omega \kappa . E[\![E_0]\!] \rho \omega (\lambda \varepsilon . \text{truish } \varepsilon \rightarrow E[\![E_1]\!] \rho \omega \kappa, \\
\quad \quad E[\![E_2]\!] \rho \omega \kappa) \\
E[\![\text{set! } \mathcal{I} \ E]\!] = \\
\quad \lambda \rho \omega \kappa . E[\![E]\!] \rho \omega (\lambda \varepsilon . \text{assign}(\text{lookup } \rho \mathcal{I}) \\
\quad \quad \varepsilon \\
\quad \quad (\text{send unspecified } \kappa)) \\
E[\![\mathcal{E}_0 \ E^*]\!] = \\
\quad \lambda \rho \omega \kappa . E^*[\![\mathcal{E}_0]\!] \rho \omega (\lambda \varepsilon \omega' \kappa' . \# \varepsilon^* = \# \mathcal{I}^* \rightarrow \\
\quad \quad \text{tievals}(\lambda \alpha^* . (\lambda \rho' . E[\![E]\!] \rho' \omega' \kappa') \\
\quad \quad \quad (\text{extends } \rho \mathcal{I}^* \alpha^*)) \\
\quad \quad \varepsilon^*, \\
\quad \quad \text{wrong "wrong number of arguments"}) \\
E[\![\mathcal{E}_0 \ E^*]\!] = \lambda \rho \omega \kappa' . \kappa' \langle \rangle \\
E^*[\![E_0 \ E^*]\!] = \lambda \rho \omega \kappa' . E[\![E_0]\!] \rho \omega (\lambda \varepsilon \omega' . E^*[\![E^*]\!] \rho \omega (\lambda \varepsilon^* . \kappa' (\langle \varepsilon_0 \rangle \varepsilon^*)))
\end{array}$$

Figure 5. Semantics of Mini-Scheme expressions

corresponding dynamic point in place:

$$\begin{array}{l}
\text{travelpath} : (\mathcal{P} \times \mathcal{F})^* \rightarrow \mathcal{C} \rightarrow \mathcal{C} \\
\text{travelpath} = \lambda \pi^* \theta . \# \pi^* = 0 \rightarrow \theta, \\
\quad ((\pi^* \downarrow 1) \downarrow 2) \langle \rangle ((\pi^* \downarrow 1) \downarrow 1) \\
\quad (\lambda \varepsilon^* . \text{travelpath}(\pi^* \uparrow 1) \theta)
\end{array}$$

The *path* function accepts two dynamic points and prefixes the journey between the two to its continuation argument:

$$\begin{array}{l}
\text{path} : \mathcal{P} \rightarrow \mathcal{P} \rightarrow (\mathcal{P} \times \mathcal{F})^* \\
\text{path} = \lambda \omega_1 \omega_2 . (\text{pathup } \omega_1 (\text{commonancest } \omega_1 \omega_2)) \S \\
\quad (\text{pathdown } (\text{commonancest } \omega_1 \omega_2) \omega_2)
\end{array}$$

The *commonancest* function finds the lowest common ancestor of two dynamic points in the control tree. Leaving aside its definition for a moment, *pathup* ascends in the control tree, picking up *after* thinks, and *pathdown* descends, picking up *before* thinks:

$$\begin{array}{l}
\text{pathup} : \mathcal{P} \rightarrow \mathcal{P} \rightarrow (\mathcal{P} \times \mathcal{F})^* \\
\text{pathup} = \\
\quad \lambda \omega_1 \omega_2 . \omega_1 = \omega_2 \rightarrow \langle \rangle, \\
\quad \langle (\omega_1, \omega_1 \mid (\mathcal{F} \times \mathcal{F} \times \mathcal{P}) \downarrow 2) \rangle \S \\
\quad (\text{pathup } (\omega_1 \mid (\mathcal{F} \times \mathcal{F} \times \mathcal{P}) \downarrow 3) \omega_2)
\end{array}$$

$$\begin{array}{l}
\text{pathdown} : \mathcal{P} \rightarrow \mathcal{P} \rightarrow (\mathcal{P} \times \mathcal{F})^* \\
\text{pathdown} = \\
\quad \lambda \omega_1 \omega_2 . \omega_1 = \omega_2 \rightarrow \langle \rangle, \\
\quad (\text{pathdown } \omega_1 (\omega_2 \mid (\mathcal{F} \times \mathcal{F} \times \mathcal{P}) \downarrow 3)) \S \\
\quad \langle (\omega_2, \omega_2 \mid (\mathcal{F} \times \mathcal{F} \times \mathcal{P}) \downarrow 1) \rangle
\end{array}$$

The *commonancest* function finds the lowest common ancestor of two dynamic points:

$$\begin{array}{l}
\text{commonancest} : \mathcal{P} \rightarrow \mathcal{P} \rightarrow \mathcal{P} \\
\text{commonancest} = \\
\quad \lambda \omega_1 \omega_2 . \text{the only element of} \\
\quad \quad \{\omega' \mid \omega' \in (\text{ancestors } \omega_1) \cap (\text{ancestors } \omega_2), \\
\quad \quad \text{pointdepth } \omega' \geq \text{pointdepth } \omega'' \\
\quad \quad \forall \omega'' \in (\text{ancestors } \omega_1) \cap (\text{ancestors } \omega_2)\}
\end{array}$$

$$\begin{array}{l}
\text{pointdepth} : \mathcal{P} \rightarrow \mathbb{N} \\
\text{pointdepth} = \\
\quad \lambda \omega . \omega = \text{root} \rightarrow 0, 1 + (\text{pointdepth } (\omega \mid (\mathcal{F} \times \mathcal{F} \times \mathcal{P}) \downarrow 3))
\end{array}$$

The *ancestors* function computes the set of ancestors of a node (including the node itself):

$$\begin{array}{l}
\text{ancestors} : \mathcal{P} \rightarrow \mathcal{P}\mathcal{P} \\
\text{ancestors} =
\end{array}$$

$$\lambda \omega . \omega = \text{root} \rightarrow \{\omega\}, \{\omega\} \cup (\text{ancestors } (\omega \mid (\mathcal{F} \times \mathcal{F} \times \mathcal{P}) \downarrow 3))$$

The *dynamic-wind* primitive calls its first argument, then calls its second argument with a new node attached to the control tree, and then calls its third argument:

$$\begin{array}{l}
\text{dynamicwind} : E^* \rightarrow \mathcal{P} \rightarrow \mathcal{K} \rightarrow \mathcal{C} \\
\text{dynamicwind} = \\
\quad \text{threearg } (\lambda \varepsilon_1 \varepsilon_2 \varepsilon_3 \omega \kappa . (\varepsilon_1 \in \mathcal{F} \wedge \varepsilon_2 \in \mathcal{F} \wedge \varepsilon_3 \in \mathcal{F}) \rightarrow \\
\quad \quad \text{apply } \varepsilon_1 \langle \rangle \omega (\lambda \zeta^* . \\
\quad \quad \quad \text{apply } \varepsilon_2 \langle \rangle ((\varepsilon_1 \mid \mathcal{F}, \varepsilon_3 \mid \mathcal{F}, \omega) \text{ in } \mathcal{P}) \\
\quad \quad \quad (\lambda \varepsilon^* . \text{apply } \varepsilon_3 \langle \rangle \omega (\lambda \zeta^* . \kappa \varepsilon^*))), \\
\quad \quad \text{wrong "bad procedure argument"})
\end{array}$$

6.2 Transition Semantics

The denotational semantics is an awkward basis for incorporating concurrency. We therefore formulate a transition semantics [27] based on the CEK machine [6] based on the denotational semantics which is amenable to the addition of concurrency. Figure 6 shows the semantics. We deliberately use the functional environment and store *mutatis mutandis* and the same letters from the denotational semantics to simplify the presentation.

The \longrightarrow relation describes transitions between states. Two kinds of state exist: either the underlying machine is about to start evaluating an expression, or it must return a value to the current continuation. The former kind is represented by a tuple $\langle \sigma, \langle E, \rho, \omega, \kappa \rangle \rangle$ where σ is the current store, E is the expression to be evaluated, ρ is the current environment, ω is the current dynamic point, and κ is the continuation of E . The latter kind of state is a tuple $\langle \sigma, \langle \kappa, \omega, \varepsilon \rangle \rangle$; σ , κ , and ω are as before, and ε is the value being passed to κ . The notable addition to the CEK machine is the *path* continuation which tracks the *dynamic-wind* *after* and *before* thinks that still need to run before returning to the “real” continuation.

6.3 Adding concurrency to the semantics

Figure 7 extends the sequential transition semantics by concurrency with preemptive scheduling in a way similar to the semantic specification of Concurrent ML [29]. The relation \Longrightarrow operates on tuples, each of which consists of the global store and a *process set* containing the running threads. Each process is represented by a unique identifier l and a state γ which is the state of the sequential semantics, sans the store. The *newid* function allocates an unused process identifier. The first rule adds concurrency. The second rule (added to the sequential semantics) describes the behavior of *spawn*: the program must first evaluate *spawn*’s argument and pass the result

$$\begin{array}{l}
\sigma \in S_d = L \rightarrow (E_d \times T) \qquad \omega \in P_d = (F_d \times F_d \times P_d) + \{root\} \\
State_d = S_d \times PState_d \qquad \phi \in F_d = \langle c1 \ \rho, I^*, E \rangle \\
PState_d = (E_d \times U \times P \times K_d) \mid (K_d \times P_d \times E_d) \qquad \varepsilon \in E_d = \dots \mid M \mid F_d \\
\kappa \in K_d = \text{stop} \mid \langle \text{cnd } E_1, E_2, \rho, \omega, \kappa \rangle \mid \langle \text{app } \langle \dots, \varepsilon, \bullet, E, \dots \rangle, \rho, \omega, \kappa \rangle \mid \langle \text{set! } \alpha, \kappa \rangle \mid \langle \text{cwcc } \kappa \rangle \\
\mid \langle \text{dw } \bullet, E_1, E_2, \rho, \omega, \kappa \rangle \mid \langle \text{dw } \varepsilon_0, \bullet, E_2, \rho, \omega, \kappa \rangle \mid \langle \text{dw } \varepsilon_0, \varepsilon_1, \bullet, \rho, \omega, \kappa \rangle \\
\mid \langle \text{dwe } \varepsilon_1, \varepsilon_2, \varepsilon_3, \rho, \omega, \kappa \rangle \mid \langle \text{dwe } \varepsilon, \rho, \omega, \kappa \rangle \mid \langle \text{return } \varepsilon, \kappa \rangle \mid \langle \text{path } (\omega, \phi)^*, \varepsilon, \kappa \rangle \\
\langle \sigma, \langle I, \rho, \omega, \kappa \rangle \rangle \mapsto \langle \sigma, \langle \kappa, \omega, \sigma(\text{lookup } \rho \ I) \downarrow 1 \rangle \rangle \\
\langle \sigma, \langle K, \rho, \omega, \kappa \rangle \rangle \mapsto \langle \sigma, \langle \kappa, \omega, K[[K]] \rangle \rangle \\
\langle \sigma, \langle (\text{lambda } (I^*) \ E_0), \rho, \omega, \kappa \rangle \rangle \mapsto \langle \sigma, \langle \kappa, \omega, \langle c1 \ \rho, I^*, E_0 \rangle \rangle \rangle \\
\langle \sigma, \langle (\text{if } E_0 \ E_1 \ E_2), \rho, \omega, \kappa \rangle \rangle \mapsto \langle \sigma, \langle E_0, \rho, \omega, \langle \text{cnd } E_1, E_2, \rho, \omega, \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle (E_0 \ E^*), \rho, \omega, \kappa \rangle \rangle \mapsto \langle \sigma, \langle E_0, \rho, \omega, \langle \text{app } \langle \bullet, E^* \rangle, \rho, \omega, \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle (\text{set! } I \ E), \rho, \omega, \kappa \rangle \rangle \mapsto \langle \sigma, \langle E, \rho, \omega, \langle \text{set! } (\text{lookup } \rho \ I), \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle (\text{call/cc } E), \omega, \kappa \rangle \rangle \mapsto \langle \sigma, \langle E, \rho, \omega, \langle \text{cwcc } \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle (\text{dynamic-wind } E_0 \ E_1 \ E_2), \rho, \omega, \kappa \rangle \rangle \mapsto \langle \sigma, \langle E_0, \rho, \omega, \langle \text{dw } \bullet, E_1, E_2, \rho, \omega, \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle (\text{cnd } E_1, E_2, \rho, \omega, \kappa), \omega', \text{false} \rangle \rangle \mapsto \langle \sigma, \langle E_2, \rho, \omega, \kappa \rangle \rangle \\
\langle \sigma, \langle (\text{cnd } E_1, E_2, \rho, \omega, \kappa), \omega', \varepsilon \rangle \rangle \mapsto \langle \sigma, \langle E_1, \rho, \omega, \kappa \rangle \rangle \quad \text{if } \varepsilon \neq \text{false} \\
\langle \sigma, \langle (\text{app } \langle \dots, \varepsilon_i, \bullet, E_{i+2}, \dots \rangle, \rho, \omega, \kappa), \omega', \varepsilon_{i+1} \rangle \rangle \mapsto \langle \sigma, \langle E_{i+2}, \rho, \omega, \langle \text{app } \langle \dots, \varepsilon_i, \varepsilon_{i+1}, \bullet, \dots \rangle, \rho, \omega, \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle (\text{app } \langle \varepsilon_0, \dots, \varepsilon_{n-1}, \bullet \rangle, \rho, \omega, \kappa), \omega', \varepsilon_n \rangle \rangle \mapsto \langle \sigma[\varepsilon_1/\alpha_1] \dots [\varepsilon_n/\alpha_n], \langle E_0, \rho_0[\alpha_1/I_1] \dots [\alpha_n/I_n], \omega, \kappa \rangle \rangle \\
\text{if } \varepsilon_0 = \langle c1 \ \rho_0, (I_1, \dots, I_n), E_0 \rangle, \alpha_1 = \text{new } \sigma \mid L, \alpha_2 = \text{new } \sigma[\varepsilon_1/\alpha_1] \mid L, \dots \\
\langle \sigma, \langle (\text{app } \langle \varepsilon_0, \bullet \rangle, \rho, \omega, \kappa), \omega', \varepsilon_1 \rangle \rangle \mapsto \langle \sigma, \langle (\text{path } (\text{path } \omega \omega'), \varepsilon_1, \kappa'), \omega, \text{unspecified} \rangle \rangle \quad \text{if } \varepsilon_0 = \langle \text{cont } \omega', \kappa' \rangle \\
\langle \sigma, \langle (\text{set! } \alpha, \kappa), \omega', \varepsilon \rangle \rangle \mapsto \langle \sigma[\varepsilon, \text{true}]/\alpha], \langle \kappa, \omega, \text{unspecified} \rangle \rangle \\
\langle \sigma, \langle (\text{cwcc } \kappa), \omega, \varepsilon \rangle \rangle \mapsto \langle \sigma[\langle \text{cont } \omega \rangle / \text{new } \sigma], \langle E_0, \rho_0[\text{new } \sigma/I], \omega, \kappa \rangle \rangle \quad \text{if } \varepsilon = \langle c1 \ \rho_0, (I), E_0 \rangle \\
\langle \sigma, \langle (\text{dw } \bullet, E_1, E_2, \rho, \omega, \kappa), \omega', \varepsilon_0 \rangle \rangle \mapsto \langle \sigma, \langle E_1, \rho, \omega, \langle \text{dw } \varepsilon_0, \bullet, E_2, \rho, \omega, \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle (\text{dw } \varepsilon_0, \bullet, E_2, \rho, \omega, \kappa), \omega', \varepsilon_1 \rangle \rangle \mapsto \langle \sigma, \langle E_2, \rho, \omega, \langle \text{dw } \varepsilon_1, \varepsilon_2, \bullet, \rho, \omega, \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle (\text{dw } \varepsilon_0, \varepsilon_1, \bullet, \rho, \omega, \kappa), \omega', \varepsilon_2 \rangle \rangle \mapsto \langle \sigma, \langle E_0, \rho_0, \omega, \langle \text{dwe } \varepsilon_0, \varepsilon_1, \varepsilon_2, \rho, \omega, \kappa \rangle \rangle \rangle \quad \text{if } \varepsilon_0 = \langle c1 \ \rho_0, \langle \rangle, E_0 \rangle \\
\langle \sigma, \langle (\text{dwe } \varepsilon_0, \varepsilon_1, \varepsilon_2, \rho, \omega, \kappa), \omega', \varepsilon'_0 \rangle \rangle \mapsto \langle \sigma, \langle E_1, \rho_1, (\varepsilon_0, \varepsilon_2, \omega), \langle \text{dwe } \varepsilon_2, \rho, \omega, \kappa \rangle \rangle \rangle \quad \text{if } \varepsilon_1 = \langle c1 \ \rho_1, \langle \rangle, E_1 \rangle \\
\langle \sigma, \langle (\text{dwe } \varepsilon_2, \rho, \omega, \kappa), \omega', \varepsilon'_1 \rangle \rangle \mapsto \langle \sigma, \langle E_2, \rho_2, \omega, \langle \text{return } \varepsilon'_1, \kappa \rangle \rangle \rangle \quad \text{if } \varepsilon_2 = \langle c1 \ \rho_2, \langle \rangle, E_2 \rangle \\
\langle \sigma, \langle (\text{return } \varepsilon, \kappa), \omega', \varepsilon \rangle \rangle \mapsto \langle \sigma, \langle \kappa, \omega, \varepsilon \rangle \rangle \\
\langle \sigma, \langle (\text{path } \langle \rangle, \varepsilon, \kappa), \omega', \varepsilon \rangle \rangle \mapsto \langle \sigma, \langle \kappa, \omega, \varepsilon \rangle \rangle \\
\langle \sigma, \langle (\text{path } \langle (\omega_0, \varepsilon_0), (\omega_1, \varepsilon_1), \dots \rangle, \varepsilon, \kappa), \omega', \varepsilon \rangle \rangle \mapsto \langle \sigma, \langle E_0, \rho_0, \omega_0, \langle \text{path } \langle (\omega_1, \varepsilon_1), \dots \rangle, \varepsilon, \kappa \rangle \rangle \rangle \quad \text{if } \varepsilon_0 = \langle c1 \ \rho_0, \langle \rangle, E_0 \rangle
\end{array}$$

Figure 6. Transition semantics for Mini-Scheme

to the `spwn` continuation. Once that happens, the third rule describes the creation of a new thread with an empty control tree and an empty continuation. The last rule removes a thread from the system once it has reached the empty continuation.

6.4 Relating the semantics

To relate the operational and the denotational semantics, we first define an evaluation function for the transition semantics:

$$eval(E, \rho, \omega, \kappa, \sigma) = \varepsilon \quad \text{if } \langle \sigma, \langle E, \rho, \omega, \kappa \rangle \rangle \mapsto^* \langle \sigma', \langle \text{stop}, \omega', \varepsilon \rangle \rangle$$

To actually prove the evaluation functions equivalent, their arguments and the result need to be equivalent in some sense. We conjecture that defining relations between the semantic domains in the spirit of [30, Section 12.6] provides us with the right notion of equivalence. Using R_{cont} to relate continuations, R_{dp} for dynamic points, R_{store} for stores, and R_* for values including errors and \perp , the equation we would like to hold is:

PROPOSITION 1. *For any Mini-Scheme expression E and environment ρ , if $\langle \hat{\kappa}, \kappa \rangle \in R_{cont}$, $\langle \hat{\omega}, \omega \rangle \in R_{dp}$, and $\langle \hat{\sigma}, \sigma \rangle \in R_{store}$, then*

$$\langle eval(E, \rho, \hat{\omega}, \hat{\kappa}, \hat{\sigma}), E[[E]]\rho\omega\kappa\sigma \rangle \in R_*$$

We commit the actual definition of the relations and the proof of the proposition to future work.

6.5 Semantics for dynamic binding

This section extends the denotational semantics for Mini-Scheme with a dynamic environment. We use the denotational semantics for `dynamic-wind` to prove the indirect implementation of dynamic binding from Section 2.2 correct. The new semantics requires a dynamic environment domain and extends the semantic domain for procedures and dynamic points by a dynamic environment:

$$\begin{array}{ll}
\phi \in F = (E^* \rightarrow P \rightarrow D \rightarrow K \rightarrow C) & \text{procedure values} \\
\psi \in D = E \rightarrow E & \text{dynamic environments} \\
\omega \in P = (F \times F \times P \times D) + \{root\} & \text{dynamic points}
\end{array}$$

The initial dynamic environment is $\psi_{init} = \lambda \varepsilon. (\varepsilon \mid E_p \downarrow 2)$. The dynamic environment is threaded through the evaluation exactly like the dynamic point. (Revised evaluation functions are in Appendix B.) All previous definitions can be adapted *mutatis mutandis* except for `dynamicwind` which needs to insert the dynamic environment into this created point and `travelpath` which calls the thinks with the environment from the point:

$$\begin{array}{l}
\text{dynamicwind} : E^* \rightarrow P \rightarrow D \rightarrow K \rightarrow C \\
\text{dynamicwind} = \text{threearg} \\
(\lambda \varepsilon_1 \varepsilon_2 \varepsilon_3 \omega \psi \kappa. (\varepsilon_1 \in F \wedge \varepsilon_2 \in F \wedge \varepsilon_3 \in F) \rightarrow \\
\text{apply } \varepsilon_1 \langle \rangle \omega \psi \\
(\lambda \zeta^*. \text{apply } \varepsilon_2 \langle \rangle ((\varepsilon_1 \mid F, \varepsilon_3 \mid F, \omega, \psi) \text{ in } P) \psi \\
(\lambda \varepsilon^*. \text{apply } \varepsilon_3 \langle \rangle \omega \psi (\lambda \zeta^*. \kappa \varepsilon^*))), \\
\text{wrong "bad procedure argument"})
\end{array}$$

$$\text{travelpath} : (P \times F)^* \rightarrow C \rightarrow C$$

ι	\in	\mathbf{I}	process IDs
$\tau = \langle \iota, \gamma \rangle$	\in	$\mathbf{Z} = \mathbf{I} \times \text{PState}_d$	processes
Ψ	\in	$\mathcal{P}^{\text{fin}} \mathbf{I}$	process sets

$$\frac{\langle \sigma, \gamma \rangle \mapsto \langle \sigma', \gamma' \rangle}{\langle \sigma, \Psi \cup \{ \langle \iota, \gamma \rangle \} \rangle \Longrightarrow \langle \sigma', \Psi \cup \{ \langle \iota, \gamma' \rangle \} \rangle}$$

$$\langle \sigma, \langle \text{spawn } E \rangle, \rho, \omega, \kappa \rangle \mapsto \langle \sigma, \langle E, \rho, \omega, \langle \text{spwn } \kappa \rangle \rangle \rangle$$

$$\langle \sigma, \Psi \rangle \Longrightarrow \langle \sigma, \Psi' \cup \{ \langle \iota, \langle \kappa, \text{unspecified} \rangle \}, \langle \text{newid } \Psi, \langle E, \rho, \text{root}, \text{stop} \rangle \} \rangle \quad \text{if } \Psi = \Psi' \cup \{ \langle \iota, \langle \text{spwn } \kappa, \varepsilon \rangle \} \}, \varepsilon = \langle \text{cl } \rho, \langle \rangle, E \rangle$$

$$\langle \sigma, \Psi \cup \{ \langle \iota, \langle \text{stop}, \varepsilon \rangle \} \rangle \Longrightarrow \langle \sigma, \Psi \rangle$$

Figure 7. Concurrent evaluation

$$\text{travelpath} = \lambda \pi^* \theta. \# \pi^* = 0 \rightarrow \theta,$$

$$((\pi^* \downarrow 1) \downarrow 2) \langle \rangle ((\pi^* \downarrow 1) \downarrow 1) \langle \rangle ((\pi^* \downarrow 1) \downarrow 1) \langle \rangle 4$$

$$(\lambda \varepsilon^*. \text{travelpath} (\pi^* \dagger 1) \theta)$$

The only additions are the definitions for creating, referencing, and binding dynamic variables:

$$E_d[\langle \text{make-fluid } E \rangle] = E_d[\langle \text{cons 'fluid } E \rangle]$$

$$\text{fluidref} : E^* \rightarrow P \rightarrow D \rightarrow K \rightarrow C$$

$$\text{fluidref} = \text{onearg } (\lambda \varepsilon \omega \psi \kappa. \text{send } (\psi \varepsilon) \kappa)$$

$$\text{bindfluid} : E^* \rightarrow P \rightarrow D \rightarrow K \rightarrow C$$

$$\text{bindfluid} =$$

$$\text{threearg } (\lambda \varepsilon_1 \varepsilon_2 \varepsilon_3 \omega \psi \kappa. \varepsilon_3 \in \mathbf{F} \rightarrow (\varepsilon_3 \mid \mathbf{F}) \langle \rangle \sigma \psi [\varepsilon_1 / \varepsilon_2] \kappa,$$

wrong “bad procedure argument”)

Again, we relate the semantics—there is only space for an informal outline of the actual proof; to abbreviate the presentation, we use value identifiers (lower-case or greek) in place of expressions evaluating to the corresponding values.

\approx relates a pair of a dynamic point ω and a dynamic environment ψ in the direct implementation with a dynamic points in the indirect implementation $\hat{\omega}$, where $\alpha_\psi = \rho * \text{dynamic-env}^*$. $\langle \omega, \psi \rangle \approx \hat{\omega}$ iff $\psi = \psi_{\text{init}}$ and $\omega = \hat{\omega}$ or all of:

$$\omega = \langle \varepsilon_{1,1}, \varepsilon_{2,1}, \langle \dots, \langle \varepsilon_{1,i}, \varepsilon_{2,i}, \omega', \psi \rangle, \dots \rangle, \psi \rangle$$

$$\hat{\omega} = \langle \varepsilon_{1,1}, \varepsilon_{2,1}, \langle \dots, \langle \varepsilon_{1,i}, \varepsilon_{2,i}, \langle \varepsilon_1, \varepsilon_2, \hat{\omega}' \rangle \rangle \rangle \rangle$$

$$\langle \omega', (\omega' \downarrow 4) \rangle \approx \hat{\omega}'$$

$$\varepsilon_1 = \lambda \varepsilon^* \omega \kappa. \lambda \sigma. \text{sendunspecified } \kappa \sigma [\langle \psi, \text{true} \rangle / \alpha_\psi]$$

$$\varepsilon_2 = \lambda \varepsilon^* \omega \kappa. \lambda \sigma. \text{sendunspecified } \kappa \sigma [\langle (\omega' \downarrow 4), \text{true} \rangle / \alpha_\psi]$$

PROPOSITION 2. *If either α_ψ holds the value of $(\text{lambda } (v) (\text{cdr } v))$ and $\psi = \psi_{\text{init}}$, or $\psi = \psi_{\text{init}} [f / \varepsilon] \dots$ and α_ψ holds the value of $(\text{shadow } \dots (\text{extend } * \text{dynamic-env}^* f \varepsilon) \dots)$, then $\forall E : E[\langle \text{fluid-ref } E \rangle] \rho \omega = E_d[\langle \text{fluid-ref } E \rangle] \rho \omega' \psi$*

THEOREM 1. *$E[\langle E \rangle] \rho \hat{\omega} \hat{\kappa} \hat{\sigma} = E_d[\langle E \rangle] \rho \omega \kappa \sigma \psi$ holds if*

$$\langle \omega, \psi \rangle \approx \hat{\omega}$$

$$\hat{\sigma} \alpha_\psi = \psi$$

$$\hat{\sigma} \varepsilon = \sigma \varepsilon \text{ for } \varepsilon \neq \alpha_\psi$$

$$\hat{\kappa} = \lambda v \sigma. \kappa v \sigma [\psi / \alpha_\psi]$$

The proof is by structural induction on E . The relevant cases are:

Case $E = (\text{bind-fluid } f \varepsilon \varepsilon_t)$: Let E_0 be the body of ε_t . By Proposition 2, the definitions of `bind-fluid` and `dynamic-wind`,

the denotation of $E[\langle E \rangle] \rho \hat{\kappa} \hat{\omega} \hat{\sigma}$ is $E[\langle E_0 \rangle] \rho \hat{\kappa}' \hat{\omega}' \hat{\sigma}'$ with

$$\hat{\sigma}' \alpha_{\psi'} = \psi [\varepsilon / f]$$

$$\hat{\omega}' = \langle \varepsilon'_1, \varepsilon'_2, \hat{\omega} \rangle$$

$$\varepsilon'_1 = \lambda \varepsilon^* \omega \kappa. \lambda \sigma. \text{sendunspecified } \kappa \sigma [\langle \psi [\varepsilon / f], \text{true} \rangle / \alpha_\psi]$$

$$\varepsilon'_2 = \lambda \varepsilon^* \omega \kappa. \lambda \sigma. \text{sendunspecified } \kappa \sigma [\langle \psi, \text{true} \rangle / \alpha_\psi]$$

$$\hat{\kappa}' = \lambda v \omega. \hat{\kappa} v \sigma [\langle \psi, \text{true} \rangle / \alpha_\psi]$$

In the direct case the denotation of $E_d[\langle E \rangle] \rho \omega \sigma \psi$ is $E_d[\langle E_0 \rangle] \rho \omega \sigma \psi [\varepsilon / f]$. The denotations of E_0 are equal by the induction hypothesis because $\langle \omega, \psi [f / v] \rangle \approx \hat{\omega}'$.

Case $E = (\text{call/cc } E_0)$: For the direct implementation, the escape procedure is $\lambda \varepsilon \omega' \psi' \kappa'. \text{travel } \omega' \omega (\kappa \varepsilon)$; the continuation is closed over the dynamic environment ψ . For the indirect implementation, the denotation is $\lambda \varepsilon \hat{\omega}' \kappa'. \text{travel } \omega' \hat{\omega} (\kappa \varepsilon)$. We show that the denotations of the escape procedures are equal if $\langle \omega', \psi' \rangle \approx \hat{\omega}'$ by case analysis of the application’s dynamic point:

1. $\hat{\omega}' = \hat{\omega}$ or $\hat{\omega}' = \langle \dots, \hat{\omega} \rangle$ and none of the intermediate dynamic points was generated by a `bind-fluid`. This corresponds to an application of the escape procedure within the body of the `call/cc` without an intermediate `bind-fluid`. This means that α_ψ is not modified and remains equal to ψ . In both cases *travelpath* evaluates all thunks with dynamic environment ψ .
2. $\hat{\omega}$ is an ancestor of $\hat{\omega}'$, w.l.o.g. $\hat{\omega}' = \langle \dots, \langle \varepsilon'_1, \varepsilon'_2, \hat{\omega} \rangle \rangle$ where $\langle \varepsilon'_1, \varepsilon'_2, \hat{\omega} \rangle$ was introduced by a `bind-fluid`. Then *commonancest* $\hat{\omega}' \hat{\omega} = \hat{\omega}$. This means *pathdown* (*commonancest* $\hat{\omega}' \hat{\omega}$) $\hat{\omega} = \langle \rangle$ and *travelpath* is applied to *pathup* $\hat{\omega}' \hat{\omega} = \langle \dots, \langle \hat{\omega}', \varepsilon'_2 \rangle \rangle$. ε'_2 sets $\sigma \alpha_{\psi'}$ to ψ because $\langle \omega, \psi \rangle \approx \hat{\omega}$. The definition of \approx ensures that the intermediate thunks are applied in equal dynamic environments.
3. Otherwise, the common ancestor is some other dynamic point ω^a i.e. $\hat{\omega} = \langle \dots, \langle \varepsilon_1, \varepsilon_2, \omega^a \rangle \rangle$. Then, *travel* $\omega' \hat{\omega} = \text{travelpath}((\text{pathup } \omega' \omega^a) \$(\text{pathdown } \omega^a \hat{\omega}))$. The second part of the argument sequence, *pathdown* $\omega^a \hat{\omega}$, is equal to $\langle \dots, \langle \hat{\omega}, \varepsilon_1 \rangle \rangle$. That is, *travelpath* will call ε_1 as last function of the sequence, which sets $\alpha_{\psi'}$ to ψ . Again the intermediate thunks are applied with identical dynamic environments because of \approx .

7 Related Work

R⁵RS [21] contains information on the history of `call/cc` in Scheme, which was part of the language (initially under a different name) from the beginning. `Dynamic-wind` was originally suggested by Richard Stallman, and reported by Friedman and Haynes [18]. Friedman and Haynes make the terminological distinction between the “plain” continuations that are just reified meta-

level continuations, and “cobs” (“continuation objects”), the actual escape procedures, which may perform work in addition to replacing the current meta-level continuation by another.

Dynamic-wind first appeared in the Scheme language definition in R⁵RS. Sitaram, in the context of the *run* and *fcontrol* control operators, associates “prelude” and “postlude” procedures with each continuation delimiter. This mechanism is comparable to *dynamic-wind* [34]. Dybvig et al. also describe a similar but more general mechanism called *process filters* for subcontinuations [19]. Filinski uses *call/cc* to transparently implement layered monads [8]. He shows how to integrate multiple computational effects by defining the relevant operators in terms of *call/cc* and state, and then re-defining *call/cc* to be compatible with the new operators. Filinski notes that that this is similar in spirit to the redefinition of *call/cc* to accommodate *dynamic-wind*, and the goals of Filinski’s work and of *dynamic-wind* are fundamentally similar.

The implementation of thread systems using *call/cc* goes back to Wand [37]. Haynes, Friedman, and others further develop this approach to implementing concurrency [17, 2]. It is also the basis for the implementation of Concurrent ML [29]. Shivers rectifies many of the misunderstandings concerning the relationship between (meta-level) continuations and threads [32]. Many implementors have since noted that *call/cc* is an appropriate explicative aid for understanding threads, but that it is not the right tool for implementing them, especially in the presence of *dynamic-wind*.

Dynamic binding goes back to early versions of Lisp [35]. Even though the replacement of dynamic binding by lexical binding was a prominent contribution of early Scheme, dynamic binding has found its way back into most implementations of Scheme and Lisp.

The inheritance issue for the dynamic environment also appears in the implementation of parallelism via futures, as noted in Feeley’s Ph.D. thesis [3] and Moreau’s work on the semantics of dynamic binding [26]. In the context of parallelism, inheritance is important because the *future* construct [15] is ideally a transparent annotation. This notion causes considerable complications for *call/cc*; Moreau investigates the semantical issues [25]. Inheritance is also a natural choice for concurrency in purely functional languages: in the Glasgow implementation of Concurrent Haskell, a new thread inherits the implicit parameters [24] from its parent. Most implementations of Common Lisp which support threads seem to have threads inherit the values of special (dynamically scoped) variables and share their values with all other threads.

The situation is different in concurrent implementations of Scheme: Scheme is not a purely functional language, and threads are typically not a transparent annotation for achieving parallelism. Therefore, Scheme implementations supporting threads and dynamic binding have made different choices: In MzScheme [9], fluid variables (called *parameters*) are inherited; mutations to parameters are only visible in the thread that performs them. The upcoming version of Gambit-C has inheritance, but parameters refer to shared cells [4]. Fluids in Scheme 48 [22] are not inherited, and do not support mutation. Scsh [33] supports a special kind of *thread fluid* [13] where inheritance can be specified upon creation. Discussion on the inheritance and sharing issues has often been controversial [4].

There is a considerable body of work on the interaction of parallelism and continuations (even though the term concurrency is often used): Parallel Scheme implementations have traditionally offered annotation-style abstractions for running computations on

other processors, such as parallel procedure calls or futures [15]. These annotations are normally transparent in purely functional programs without *call/cc*. Implementors have tried to make them transparent even in the presence of *call/cc* [20], which makes it necessary (and sensible) to have reified continuations span multiple threads. However, none of the implementations behaves intuitively in all cases, and none maintains transparency when the program executes side effects. Hieb et al. [19] alleviate this problem by proposing the use of delimited continuations—so-called *subcontinuations*—to express intuitive behavior. All of this work is largely orthogonal to ours which is largely concerned with concurrency as a programming paradigm. However, in our view, this confirms our conclusion that comingling threads and continuations leads to undesirable complications.

8 Conclusion

Combining first-class continuations, *dynamic-wind*, dynamic binding, and concurrency in a single functional language is akin to walking a minefield. The design space exhibits many peculiarities, and its size is considerable; existing systems occupy different places within it. Some design choices lead to semantic or implementation difficulties, others impact the programmer’s ability to write modular multithreaded programs. In general, the discussion about the correct way to combine these facilities has been plagued by controversy and confusion. In this paper, we have examined the interactions between them in a systematic way. The most important insights are:

- It is better to build first-class continuations and *dynamic-wind* on top of a native thread system rather than building the thread system on top of continuations.
- Decoupling threads from the sequential part of the programming language leads to clean semantic specifications and easier-to-understand program behavior.
- Abstractions for thread-aware programming are useful, but their use can have a negative impact on modularity and thus requires great care.
- The semantic interaction between threads and dynamic binding in Scheme is easiest to explain when newly created threads start with a fresh dynamic context. Even though this design option is not current practice in many systems, it also offers the greatest flexibility when writing modular abstractions which use threads and dynamic binding.

Our work opens a number of avenues for further research. In particular, an equational specification for *dynamic-wind* in the style of Felleisen and Hieb’s framework [7] would be very useful. This could also be the basis for characterizing “benevolent” uses of *dynamic-wind* and *thread-wind* that do not interfere with *call/cc* in undesirable ways.

Acknowledgments: We thank Will Clinger, Olivier Danvy, Marc Feeley, Matthias Felleisen, and Matthew Flatt for patiently answering our questions. The reviewers of ICFP 2003 and the Scheme Workshop also provided helpful feedback.

9 References

- [1] Edoardo Biagioni, Ken Cline, Peter Lee, Chris Okasaki, and Chris Stone. Safe-for-space threads in Standard ML. *Higher-Order and Symbolic Computation*, 11(2):209–225, December 1998.

- [2] R. Kent Dybvig and Robert Hieb. Engines from continuations. *Computer Languages*, 14(2):109–123, 1989.
- [3] Marc Feeley. *An Efficient and General Implementation of Futures on Large Scale Shared-Memory Multiprocessors*. PhD thesis, Brandeis University, 1993.
- [4] Marc Feeley. SRFI 39: Parameter objects. <http://srfi.schemers.org/srfi-39/>, December 2002.
- [5] Marc Feeley. SRFI 18: Multithreading support. <http://srfi.schemers.org/srfi-18/>, March 2001.
- [6] Matthias Felleisen and Daniel Friedman. Control operators, the SECD-machine, and the λ -calculus. In *Formal Description of Programming Concepts III*, pages 193–217. North-Holland, 1986.
- [7] Matthias Felleisen and Robert Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 10(2):235–271, 1992.
- [8] Andrzej Filinski. Representing layered monads. In Alexander Aiken, editor, *Proceedings of the 1999 ACM SIGPLAN Symposium on Principles of Programming Languages*, pages 175–188, San Antonio, Texas, USA, January 1999. ACM Press.
- [9] Matthew Flatt. *PLT MzScheme: Language Manual*. Rice University, University of Utah, January 2003. Version 203.
- [10] Matthew Flatt, Robert Bruce Findler, Shriram Krishnamurthi, and Matthias Felleisen. Programming languages as operating systems. In Peter Lee, editor, *Proc. International Conference on Functional Programming 1999*, pages 138–147, Paris, France, September 1999. ACM Press, New York.
- [11] Daniel P. Friedman, Mitchell Wand, and Christopher T. Haynes. *Essentials of Programming Languages*. MIT Press and McGraw-Hill, 2nd edition, 2001.
- [12] Martin Gasbichler and Michael Sperber. Final shift for call/cc: Direct implementation of shift and reset. In Simon Peyton-Jones, editor, *Proc. International Conference on Functional Programming 2002*, Pittsburgh, PA, USA, October 2002. ACM Press, New York.
- [13] Martin Gasbichler and Michael Sperber. Processes vs. user-level threads in Scsh. In Olin Shivers, editor, *Proceedings of the Third Workshop on Scheme and Functional Programming*, Pittsburgh, October 2002.
- [14] Paul Graunke, Shriram Krishnamurthi, Steve Van Der Hoeven, and Matthias Felleisen. Programming the Web with high-level programming languages. In David Sands, editor, *Proceedings of the 2001 European Symposium on Programming*, Lecture Notes in Computer Science, pages 122–136, Genova, Italy, April 2001. Springer-Verlag.
- [15] Robert H. Halstead, Jr. A language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems*, 7(4):501–538, 1985.
- [16] David R. Hanson and Todd A. Proebsting. Dynamic variables. In *Proceedings of the 2001 Conference on Programming Language Design and Implementation*, pages 264–273, Snowbird, UT, June 2001. ACM Press.
- [17] Christopher T. Haynes and Daniel P. Friedman. Abstracting timed preemption with engines. *Computer Languages*, 12(2):109–121, 1987.
- [18] Christopher T. Haynes and Daniel P. Friedman. Embedding continuations in procedural objects. *ACM Transactions on Programming Languages and Systems*, 9(4):582–598, October 1987.
- [19] Robert Hieb, R. Kent Dybvig, and C. W. Anderson, III. Sub-continuations. *Lisp and Symbolic Computation*, 7(1):x, 1994.
- [20] Morry Katz and Daniel Weise. Continuing into the future: On the interaction of futures and first-class continuations. In *Proceedings of the 1990 ACM Conference on Lisp and Functional Programming*, pages 176–184, Nice, France, 1990. ACM Press.
- [21] Richard Kelsey, William Clinger, and Jonathan Rees. Revised⁵ report on the algorithmic language Scheme. *Higher-Order and Symbolic Computation*, 11(1):7–105, 1998.
- [22] Richard Kelsey and Jonathan Rees. *Scheme 48 Reference Manual*, 2002. Part of the Scheme 48 distribution at <http://www.s48.org/>.
- [23] Richard Kelsey and Michael Sperber. SRFI 34: Exception handling for programs. <http://srfi.schemers.org/srfi-34/>, December 2002.
- [24] Jeffery R. Lewis, John Launchbury, Erik Meijer, and Mark B. Shields. Implicit parameters: Dynamic scoping with static types. In Tom Reps, editor, *Proc. 27th Annual ACM Symposium on Principles of Programming Languages*, pages 108–118, Boston, MA, USA, January 2000. ACM Press.
- [25] Luc Moreau. The semantics of future in the presence of first-class continuations and side-effects. Technical Report M95/3, University of Southampton, November 1995.
- [26] Luc Moreau. A syntactic theory for dynamic binding. *Higher-Order and Symbolic Computation*, 11(3):233–279, 1998.
- [27] Gordon Plotkin. A structural approach to operational semantics. Technical Report DAIMI FN-19, Aarhus University, Denmark, 1981.
- [28] Christian Queinnec. The influence of browsers on evaluators or, continuations to program Web servers. In Philip Wadler, editor, *Proc. International Conference on Functional Programming 2000*, pages 23–33, Montreal, Canada, September 2000. ACM Press, New York.
- [29] John H. Reppy. *Concurrent Programming in ML*. Cambridge University Press, 1999.
- [30] John C. Reynolds. *Theories of Programming Languages*. Cambridge University Press, 1998.
- [31] Olin Shivers. A Scheme Shell. Technical Report TR-635, Massachusetts Institute of Technology, Laboratory for Computer Science, April 1994.
- [32] Olin Shivers. Continuations and threads: Expressing machine concurrency directly in advanced languages. In Olivier Danvy, editor, *Proceedings of the Second ACM SIGPLAN Workshop on Continuations*, number NS-96-13 in BRICS Notes, Paris, France, January 1997. Dept. of Computer Science, Aarhus, Denmark.
- [33] Olin Shivers, Brian D. Carlstrom, Martin Gasbichler, and Mike Sperber. *Scsh Reference Manual*, 2003. Available from <http://www.scsh.net/>.
- [34] Dorai Sitaram. *Models of Control and Their Implications for Programming Language Design*. PhD thesis, Rice University, Houston, Texas, April 1994.
- [35] Guy L. Steele Jr. and Richard P. Gabriel. The evolution of Lisp. In Jean E. Sammet, editor, *History of Programming*

Languages II, pages 231–270. ACM, New York, April 1993. SIGPLAN Notices 3(28).

- [36] Joseph Stoy. *Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory*. MIT Press, 1977.
- [37] Mitchell Wand. Continuation-based multiprocessing. In J. Allen, editor, *Conference Record of the 1980 LISP Conference*, pages 19–28, Palo Alto, 1980. The Lisp Company.

A Multiprocessing and Thread-wind

Figure 8 shows how to extend the sequential transition semantics from Section 6.3 to account for multiprocessing and `thread-wind`: The \longrightarrow relation operates on machine states $\langle \sigma, \Psi, \Pi, \beta \rangle$. As before, σ is the global store, Ψ is still a process set, but contains only the threads currently running on a processor. Π is a *processor map* mapping a processor to a processor state, which is either *idle* or *running* ι for a processor running the thread with ID ι . β is the set of of idle threads waiting to be scheduled on a processor. Each member of this set consists of the thread ID, the dynamic point to return to, and the state of that thread.

The rules for running threads and spawning new ones are much as before, only extended to account for the new machine state components. (The *newid* function now takes both the active and idle process sets as arguments.) The last three rules control the swapping in and swapping out of threads: The first of these prepares a thread for swap-out, prefixing the current continuation with a winding path and a suspend marker. (For simplicity, we allow swapping out only when returning a value to a continuation.) The winding path is obtained by travelling up the control tree, only collecting *after* thunks introduced by `thread-wind`. (The new P domain distinguishes between nodes introduced by `dynamic-wind` and those introduced by `thread-wind` by a new boolean flag.) The subsequent rule actually performs the swapping out once the thread has reached the suspend marker. The last rule swaps a thread back in, prefixing the path back down to the target control node.

The `tpath` continuation works exactly the same as the `path` continuation, with the only exception that a processor running a thread in the midst of `tpath` continuation cannot swap that thread out.

B Semantics for Mini-Scheme with dynamic binding

Figure 9 describes evaluation functions E and E^* for Mini-Scheme with dynamic binding as described in Section 6.5.

C Defining `dynamic-wind` using the continuation monad

The published version of R⁵RS says:

The definition of `call-with-current-continuation` in Section 8.2 is incorrect because it is incompatible with `dynamic-wind`. As shown in Section 4 of [1], however, this incorrect semantics is adequate to define the `shift` and `reset` operators, which can then be used to define the correct semantics of both `dynamic-wind` and `call-with-current-continuation`.

The origin of this comment is unclear, and there is no published (or, to our knowledge, any) implementation of

`call-with-current-continuation` and `dynamic-wind` to support this claim. We work out the details here. Our implementation represents a dynamic point as a pair of a pair of a *before* and an *after* thunk, and the parent point. The root point is represented as the empty list.

```
(define root-point '())

(define root-point? null?)

(define (make-point before after parent)
  (cons (cons before after) parent))

(define point-parent cdr)

(define (point-depth p)
  (if (root-point? p)
      0
      (+ 1 (point-depth (point-parent p)))))

(define point-before caar)

(define point-after cdar)
```

Filinski’s framework for representing monads provides two functions `reify` and `reflect` which mediate between computations and values (the macro `reify*` simply wraps its argument into a thunk to shorten the rest of the examples):

```
(define (reflect meaning)
  (shift k (extend k meaning)))

(define (reify thunk)
  (reset (eta (thunk))))

(define-syntax reify*
  (syntax-rules ()
    ((reify* body ...)
     (reify (lambda () body ...)))))
```

See [12] for a Scheme version of Filinski’s definition of `shift` and `reset` in terms of `call/cc`. The procedures `eta` and `extend` correspond to the usual monadic unit and extension functions. In Haskell, `eta` is known as `return` and `extend` as `bind` or the infix operator `>>=`.

Defining the continuation monad requires defining `eta` and `extend`. The datatype of the plain continuation monad contains a procedure which accepts a continuation as its argument and delivers its result by applying a continuation. The unit operation delivers a value by applying the continuation. The extension operation puts the function into the continuation:

```
(define (eta a)
  (lambda (c) (c a)))

(define (extend k m)
  (lambda (c)
    (m (lambda (v) ((k v) c)))))
```

For actually running programs, an evaluation function which supplies the identity function to its argument comes in handy:

```
(define (eval m)
  ((reify m) (lambda (v) v)))
```

ω	\in	$P = (F \times F \times T \times P) + \{root\}$	dynamic points and thread points
φ	\in	W	processor
ζ	\in	$J = \text{idle} \mid \text{running}$	processor state
Π	\in	$W \rightarrow J$	processor map
π^*	\in	$G = (P \times F)^*$	thread-wind paths
β	\in	$B = P^{\text{fin}}(I \times P \times Y)$	idle threads

$$\begin{array}{c}
\frac{\langle \sigma, \gamma \rangle \mapsto \langle \sigma', \gamma' \rangle}{\langle \sigma, \Psi \cup \{ \langle \iota, \gamma \rangle \}, \Pi, \beta \rangle \mapsto \langle \sigma', \Psi \cup \{ \langle \iota, \gamma' \rangle \}, \Pi, \beta \rangle} \\
\langle \sigma, \Psi \cup \{ \langle \iota, \langle \text{stop}, \omega, \varepsilon \rangle \} \}, \Pi, \beta \rangle \mapsto \langle \sigma, \Psi, \Pi[\text{idle}/\varphi], \beta \rangle \quad \text{if } \Pi\varphi = \text{running } \iota \\
\langle \sigma, \langle \langle \text{spawn } E \rangle, \rho, \omega, \kappa \rangle, \Pi, \beta \rangle \mapsto \langle \sigma, \langle E, \rho, \omega, \langle \text{spwn } \kappa \rangle \rangle, \Pi, \beta \rangle \\
\langle \sigma, \Psi, \Pi, \beta \rangle \mapsto \langle \sigma, \Psi' \cup \{ \langle \iota, \langle \kappa, \text{unspecified} \rangle \} \}, \beta \cup \{ \langle \text{newid } \Psi \beta, \text{root}, \langle E, \rho, \text{root}, \text{stop} \rangle \} \} \\
\quad \text{if } \Psi = \Psi' \cup \{ \langle \iota, \langle \text{spwn } \kappa, \varepsilon \rangle \} \}, \varepsilon = \langle \text{cl } \rho, \langle \rangle, E \rangle \\
\langle \sigma, \Psi \cup \{ \langle \iota, \langle \kappa, \omega, \varepsilon \rangle \} \}, \Pi, \beta \rangle \mapsto \langle \sigma, \Psi \cup \{ \langle \iota, \langle \langle \text{tpath } (\text{pathup}_m \omega), \varepsilon, \langle \text{suspend } \kappa \rangle \rangle, \omega, \text{unspecified} \rangle \} \}, \Pi, \beta \rangle \\
\quad \text{if } \kappa \text{ does not contain tpath} \\
\langle \sigma, \Psi \cup \{ \langle \iota, \langle \langle \text{suspend } \kappa \rangle, \omega, \varepsilon \rangle \} \}, \Pi, \beta \rangle \mapsto \langle \sigma, \Psi, \Pi[\text{idle}/\varphi], \beta \cup \{ \langle \iota, \omega, \langle \kappa, \omega, \varepsilon \rangle \} \} \quad \text{if } \Pi\varphi = \text{running } \iota \\
\langle \sigma, \Psi, \beta \cup \{ \langle \iota, \omega, \langle \kappa, \omega, \varepsilon \rangle \} \} \rangle \mapsto \langle \sigma, \Psi \cup \{ \langle \iota, \langle \langle \text{tpath } (\text{pathdown}_m \omega), \varepsilon, \kappa \rangle \rangle, \omega, \text{unspecified} \rangle \} \}, \Pi[\text{running } \iota/\varphi], \beta \rangle \quad \text{if } \Pi\varphi = \text{idle} \\
\langle \sigma, \langle \langle \text{tpath } \langle \rangle, \varepsilon, \kappa \rangle, \omega', \varepsilon' \rangle \rangle \mapsto \langle \sigma, \langle \kappa, \omega, \varepsilon \rangle \rangle \\
\langle \sigma, \langle \langle \text{tpath } \langle (\omega_0, \varepsilon_0), (\omega_1, \varepsilon_1), \dots \rangle, \varepsilon, \kappa \rangle, \omega', \varepsilon' \rangle \rangle \mapsto \langle \sigma, \langle E_0, \rho_0, \omega_0, \langle \text{tpath } \langle (\omega_1, \varepsilon_1), \dots \rangle, \varepsilon, \kappa \rangle \rangle \rangle \quad \text{if } \varepsilon_0 = \langle \text{cl } \rho_0, \langle \rangle, E_0 \rangle \\
\langle \sigma, \langle \langle \text{thread-wind } E_0 E_1 E_2 \rangle, \rho, \omega, \kappa \rangle \rangle \mapsto \langle \sigma, \langle E_0, \rho, \omega, \langle \text{tw } \bullet, E_1, E_2, \rho, \omega, \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle \langle \text{tw } \bullet, E_1, E_2, \rho, \omega, \kappa \rangle, \omega', \varepsilon_0 \rangle \rangle \mapsto \langle \sigma, \langle E_1, \rho_1, \omega, \langle \text{tw } \varepsilon_0, \bullet, E_2, \rho, \omega, \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle \langle \text{tw } \varepsilon_0, \bullet, E_2, \rho, \omega, \kappa \rangle, \omega', \varepsilon_1 \rangle \rangle \mapsto \langle \sigma, \langle E_2, \rho_2, \omega, \langle \text{tw } \varepsilon_1, \varepsilon_2, \bullet, \rho, \omega, \kappa \rangle \rangle \rangle \\
\langle \sigma, \langle \langle \text{tw } \varepsilon_0, \varepsilon_1, \bullet, \rho, \omega, \kappa \rangle, \omega', \varepsilon_2 \rangle \rangle \mapsto \langle \sigma, \langle E_0, \rho_0, \omega, \langle \text{twe } \varepsilon_0, \varepsilon_1, \varepsilon_2, \rho, \omega, \kappa \rangle \rangle \rangle \quad \text{if } \varepsilon_0 = \langle \text{cl } \rho_0, \langle \rangle, E_0 \rangle \\
\langle \sigma, \langle \langle \text{dwe } \varepsilon_0, \varepsilon_1, \varepsilon_2, \rho, \omega, \kappa \rangle, \omega', \varepsilon'_0 \rangle \rangle \mapsto \langle \sigma, \langle E_1, \rho_0, (\varepsilon_0, \varepsilon_2, \text{false}, \omega), \langle \text{dwe } \varepsilon_2, \rho, \omega, \kappa \rangle \rangle \rangle \quad \text{if } \varepsilon_1 = \langle \text{cl } \rho_1, \langle \rangle, E_1 \rangle \\
\langle \sigma, \langle \langle \text{twe } \varepsilon_0, \varepsilon_1, \varepsilon_2, \rho, \omega, \kappa \rangle, \omega', \varepsilon'_0 \rangle \rangle \mapsto \langle \sigma, \langle E_1, \rho_0, (\varepsilon_0, \varepsilon_2, \text{true}, \omega), \langle \text{dwe } \varepsilon_2, \rho, \omega, \kappa \rangle \rangle \rangle \quad \text{if } \varepsilon_1 = \langle \text{cl } \rho_1, \langle \rangle, E_1 \rangle
\end{array}$$

$\text{pathup}_m : P \rightarrow G$
 $\text{pathup}_m =$
 $\lambda\omega. \omega = \text{root} \rightarrow \langle \rangle,$
 $(\omega \mid (F \times F \times T \times P) \downarrow 3) = \text{true} \rightarrow \langle (\omega, \omega \mid (F \times F \times T \times P) \downarrow 2) \rangle \S (\text{pathup}_m (\omega \mid (F \times F \times T \times P) \downarrow 4)),$
 $(\text{pathup}_m (\omega \mid (F \times F \times T \times P) \downarrow 4))$

$\text{pathdown}_m : P \rightarrow G$
 $\text{pathdown}_m =$
 $\lambda\omega. \omega = \text{root} \rightarrow \langle \rangle,$
 $(\omega \mid (F \times F \times T \times P) \downarrow 3) = \text{true} \rightarrow (\text{pathdown}_m (\omega \mid (F \times F \times T \times P) \downarrow 4)) \langle (\omega, \omega \mid (F \times F \times T \times P) \downarrow 2) \rangle,$
 $(\text{pathdown}_m (\omega \mid (F \times F \times T \times P) \downarrow 4))$

Figure 8. Multiprocessor evaluation

The definition of call/cc is straightforward:

```

(define (call/cc h)
  (reflect
   (lambda (c)
     (let ((k (lambda (v)
                 (reflect (lambda (c-prime) (c v))))))
       ((reify* (h k)) c))))))

```

To incorporate dynamic-wind we pair the continuation function with a dynamic point. Eta still applies the continuation to its argument, while extend supplies the same dynamic point to both of its arguments.

```

(define (eta a)
  (lambda (cdp) ((car cdp) a)))

```

```

(define (extend k m)
  (lambda (cdp)
    (m (cons (lambda (v) ((k v) cdp)) (cdr cdp)))))

```

The evaluation procedure takes a thunk representing the computation as argument, reifies it and applies it to the identity continuation and the root point:

```

(define (eval m)
  ((reify m) (cons (lambda (v) v) root-point)))

```

Dynamic-wind evaluates first evaluates the before thunk. It then evaluates the body thunk with a new dynamic point, before it evaluates the after thunk with a continuation which applies the continuation of the dynamic-wind to the result of the body.

$$\begin{array}{l}
E_d : \text{Exp} \rightarrow \text{U} \rightarrow \text{P} \rightarrow \text{D} \rightarrow \text{K} \rightarrow \text{C} \\
E_d^* : \text{Exp}^* \rightarrow \text{U} \rightarrow \text{P} \rightarrow \text{D} \rightarrow \text{K} \rightarrow \text{C} \\
E_d[[K]] = \lambda\rho\omega\psi\kappa . \text{send}(K [[K]])\kappa \\
E_d[[I]] = \lambda\rho\omega\psi\kappa . \text{hold}(\text{lookup}\rho I) \\
\quad (\lambda\varepsilon . \varepsilon = \text{undefined} \rightarrow \\
\quad \quad \text{wrong "undefined variable",} \\
\quad \quad \text{send}\varepsilon\kappa) \\
E_d[[\text{if } E_0 E_1 E_2]] = \\
\quad \lambda\rho\omega\psi\kappa . E_d[[E_0]]\rho\omega\psi (\lambda\varepsilon . \text{truish}\varepsilon \rightarrow E_d[[E_1]]\rho\omega\psi\kappa, \\
\quad \quad E_d[[E_2]]\rho\omega\psi\kappa) \\
E_d[[\text{set! } I E]] = \\
\quad \lambda\rho\omega\psi\kappa . E_d[[E]]\rho\omega\psi (\lambda\varepsilon . \text{assign}(\text{lookup}\rho I) \\
\quad \quad \varepsilon \\
\quad \quad (\text{send}\text{unspecified}\kappa)) \\
E_d[[\text{lambda } (I^*) E]] = \\
\quad \lambda\rho\omega\psi\kappa . E_d^*[(\text{E}_0 \text{ E}^*)] = \\
\quad \quad \lambda\rho\omega\psi\kappa . E_d^*(\langle \text{E}_0 \rangle \S \text{E}^*) \\
\quad \quad \rho\omega\psi (\lambda\varepsilon^* . \text{apply}(\varepsilon^* \downarrow 1) (\varepsilon^* \uparrow 1) \omega\psi\kappa) \\
E_d^*[(\text{lambda } (I^*) E)] = \\
\quad \lambda\rho\omega\psi\kappa . \\
\quad \quad \text{send}((\lambda\varepsilon^*\omega'\psi'\kappa' . \#\varepsilon^* = \#I^* \rightarrow \\
\quad \quad \quad \text{tievals}(\lambda\alpha^* . (\lambda\rho' . E_d[[E]]\rho'\omega'\psi'\kappa') \\
\quad \quad \quad (\text{extends}\rho I^*\alpha^*))) \\
\quad \quad \quad \varepsilon^*, \\
\quad \quad \quad \text{wrong "wrong number of arguments"}) \\
\quad \quad \text{in E}) \\
\quad \quad \kappa \\
E_d^*[[\langle \rangle]] = \lambda\rho\omega\psi\kappa . \kappa\langle \rangle \\
E_d^*[[E_0 E^*]] = \\
\quad \lambda\rho\omega\psi\kappa . E_d[[E_0]]\rho\omega\psi (\text{single}(\lambda\varepsilon_0 . E_d^*[[E^*]]\rho\omega\psi (\lambda\varepsilon^* . \kappa ((\varepsilon_0) \S \varepsilon^*))))
\end{array}$$

Figure 9. Semantics of Mini-Scheme with dynamic environment

```

(define (dynamic-wind before thunk after)
  (reflect
   (lambda (cdp)
     ((reify* (before))
      (cons (lambda (v1)
              ((reify* (thunk))
               (cons (lambda (v2)
                       ((reify* (after))
                        (cons (lambda (v3)
                                ((car cdp) v2))
                               (cdr cdp))))
                        (make-point before after
                                   (cdr cdp))))))
              (cdr cdp)))))))

```

Call/cc is responsible for generating an escape procedure which calls the appropriate set of before and after thunks. The following code defers this to the procedure `travel-to-point!`:

```

(define (call/cc h)
  (reflect
   (lambda (cdp)
     (let ((k (lambda (v)
                 (reflect
                  (lambda (cdp-prime)
                    ((reify* (travel-to-point!
                               (cdr cdp-prime)
                               (cdr cdp)))
                     (cons (lambda (ignore)
                             ((car cdp) v))
                            (cdr cdp))))))))
       ((reify* (h k)) cdp))))))

```

`Travel-to-point!` implements an ingenious algorithm invented by Pavel Curtis for Scheme Xerox and used in Scheme 48:

```

(define (travel-to-point! here target)
  (cond ((eq? here target) 'done)
        ((or (root-point? here)
              (and (not (root-point? target))
                   (< (point-depth here)
                       (point-depth target))))
         (travel-to-point! here
                              (point-parent target)))
        (with-point target
         (lambda () ((point-before target))))))

```

```

(else
 (with-point here
  (lambda () ((point-after here))))
 (travel-to-point! (point-parent here)
                   target))))

```

The algorithm seeks the common ancestor by first walking up from lower of the two points until it is at the same level as the other. Then it alternately walks up one step at each of the points until it arrives at the same point, which is the common ancestor. The algorithm runs the *after* thunks walking up the source branch and winds up running the *before* thunks walking up the target branch. The helper procedure `with-point` takes a dynamic point and a thunk as its arguments and evaluates the thunk with the current continuation and the supplied point:

```

(define (with-point point thunk)
  (reflect
   (lambda (cdp)
     ((reify* (thunk))
      (cons (lambda (v) ((car cdp) v)) point))))))

```

D Denotational Semantics

[This is a version of the denotational semantics in *R⁵RS with dynamic-wind*. We have copied the text verbatim, only making the necessary changes to account for the management of dynamic points.]

This section provides a formal denotational semantics for the primitive expressions of Scheme and selected built-in procedures. The concepts and notation used here are described in [36]; the notation is summarized below:

$\langle \dots \rangle$	sequence formation
$s \downarrow k$	k th member of the sequence s (1-based)
$\#s$	length of sequence s
$s \S t$	concatenation of sequences s and t
$s \uparrow k$	drop the first k members of sequence s
$t \rightarrow a, b$	McCarthy conditional “if t then a else b ”
$\rho[x/i]$	substitution “ ρ with x for i ”
$x \text{ in } D$	injection of x into domain D
$x D$	projection of x to domain D

The reason that expression continuations take sequences of values instead of single values is to simplify the formal treatment of pro-

cedure calls and multiple return values.

The boolean flag associated with pairs, vectors, and strings will be true for mutable objects and false for immutable objects.

The order of evaluation within a call is unspecified. We mimic that here by applying arbitrary permutations *permute* and *unpermute*, which must be inverses, to the arguments in a call before and after they are evaluated. This is not quite right since it suggests, incorrectly, that the order of evaluation is constant throughout a program (for any given number of arguments), but it is a closer approximation to the intended semantics than a left-to-right evaluation would be.

The storage allocator *new* is implementation-dependent, but it must obey the following axiom: if $\text{new } \sigma \in L$, then $\sigma(\text{new } \sigma | L) \downarrow 2 = \text{false}$.

The definition of K is omitted because an accurate definition of K would complicate the semantics without being very interesting.

If P is a program in which all variables are defined before being referenced or assigned, then the meaning of P is

$$E[\langle (\text{lambda } (I^*) P') \langle \text{undefined} \rangle \dots \rangle]$$

where I^* is the sequence of variables defined in P , P' is the sequence of expressions obtained by replacing every definition in P by an assignment, $\langle \text{undefined} \rangle$ is an expression that evaluates to *undefined*, and E is the semantic function that assigns meaning to expressions.

D.1 Abstract syntax

$K \in \text{Con}$	constants, including quotations
$I \in \text{Ide}$	identifiers (variables)
$E \in \text{Exp}$	expressions
$\Gamma \in \text{Com} = \text{Exp}$	commands

$\text{Exp} \longrightarrow$	$K \mid I \mid (E_0 E^*)$
	$\mid (\text{lambda } (I^*) \Gamma^* E_0)$
	$\mid (\text{lambda } (I^* . I) \Gamma^* E_0)$
	$\mid (\text{lambda } I \Gamma^* E_0)$
	$\mid (\text{if } E_0 E_1 E_2) \mid (\text{if } E_0 E_1)$
	$\mid (\text{set! } I E)$

D.2 Domain equations

$\alpha \in L$	locations
$v \in \mathbb{N}$	natural numbers
$T = \{\text{false}, \text{true}\}$	booleans
Q	symbols
H	characters
R	numbers
$E_p = L \times L \times T$	pairs
$E_v = L^* \times T$	vectors
$E_s = L^* \times T$	strings
$M = \{\text{false}, \text{true}, \text{null}, \text{undefined}, \text{unspecified}\}$	miscellaneous
$\phi \in F = L \times (E^* \rightarrow P \rightarrow K \rightarrow C)$	procedure values
$\varepsilon \in E = Q + H + R + E_p + E_v + E_s + M + F$	expressed values
$\sigma \in S = L \rightarrow (E \times T)$	stores

$\rho \in U = \text{Ide} \rightarrow L$	environments
$\theta \in C = S \rightarrow A$	command continuations
$\kappa \in K = E^* \rightarrow C$	expression continuations
A	answers
X	errors
$\omega \in P = (F \times F \times P) + \{\text{root}\}$	dynamic points

D.3 Semantic functions

$K : \text{Con} \rightarrow E$
$\bar{E} : \text{Exp} \rightarrow U \rightarrow P \rightarrow K \rightarrow C$
$E^* : \text{Exp}^* \rightarrow U \rightarrow P \rightarrow K \rightarrow C$
$C : \text{Com}^* \rightarrow U \rightarrow P \rightarrow C \rightarrow C$

Definition of K deliberately omitted.

$$E[\llbracket K \rrbracket] = \lambda \rho \omega \kappa . \text{send } (K \llbracket K \rrbracket) \kappa$$

$$E[\llbracket I \rrbracket] = \lambda \rho \omega \kappa . \text{hold } (\text{lookup } \rho I) \\ (\text{single } (\lambda \varepsilon . \varepsilon = \text{undefined} \rightarrow \\ \text{wrong } \text{“undefined variable”}, \\ \text{send } \varepsilon \kappa))$$

$$E[\llbracket (E_0 E^*) \rrbracket] = \\ \lambda \rho \omega \kappa . E^*(\text{permute}(\langle E_0 \rangle \S E^*)) \\ \rho \\ \omega \\ (\lambda \varepsilon^* . ((\lambda \varepsilon^* . \text{apply}(\varepsilon^* \downarrow 1) (\varepsilon^* \uparrow 1) \omega \kappa) \\ (\text{unpermute } \varepsilon^*)))$$

$$E[\llbracket (\text{lambda } (I^*) \Gamma^* E_0) \rrbracket] = \\ \lambda \rho \omega \kappa . \lambda \sigma . \\ \text{new } \sigma \in L \rightarrow \\ \text{send } (\langle \text{new } \sigma | L, \\ \lambda \varepsilon^* \omega' \kappa' . \# \varepsilon^* = \# I^* \rightarrow \\ \text{tievals}(\lambda \alpha^* . (\lambda \rho' . C[\llbracket \Gamma^* \rrbracket] \rho' \omega' (E[\llbracket E_0 \rrbracket] \rho' \omega' \kappa')) \\ (\text{extends } \rho I^* \alpha^*)) \\ \varepsilon^*, \\ \text{wrong } \text{“wrong number of arguments”}) \\ \text{in } E) \\ \kappa \\ (\text{update } (\text{new } \sigma | L) \text{unspecified } \sigma), \\ \text{wrong } \text{“out of memory” } \sigma$$

$$E[\llbracket (\text{lambda } (I^* . I) \Gamma^* E_0) \rrbracket] = \\ \lambda \rho \omega \kappa . \lambda \sigma . \\ \text{new } \sigma \in L \rightarrow \\ \text{send } (\langle \text{new } \sigma | L, \\ \lambda \varepsilon^* \omega' \kappa' . \# \varepsilon^* \geq \# I^* \rightarrow \\ \text{tievalsrest} \\ (\lambda \alpha^* . (\lambda \rho' . C[\llbracket \Gamma^* \rrbracket] \rho' \omega' (E[\llbracket E_0 \rrbracket] \rho' \omega' \kappa')) \\ (\text{extends } \rho (I^* \S (I) \alpha^*)) \\ \varepsilon^* \\ (\# I^*), \\ \text{wrong } \text{“too few arguments”}) \text{ in } E) \\ \kappa \\ (\text{update } (\text{new } \sigma | L) \text{unspecified } \sigma), \\ \text{wrong } \text{“out of memory” } \sigma$$

$$E[\llbracket (\text{lambda } I \Gamma^* E_0) \rrbracket] = E[\llbracket (\text{lambda } (. I) \Gamma^* E_0) \rrbracket]$$

$$E[\llbracket (\text{if } E_0 E_1 E_2) \rrbracket] = \\ \lambda \rho \omega \kappa . E[\llbracket E_0 \rrbracket] \rho \omega (\text{single } (\lambda \varepsilon . \text{truish } \varepsilon \rightarrow E[\llbracket E_1 \rrbracket] \rho \omega \kappa, \\ E[\llbracket E_2 \rrbracket] \rho \omega \kappa))$$

$$E[\langle \text{if } E_0 \ E_1 \rangle] = \lambda \rho \omega \kappa . E[\langle E_0 \rangle] \rho \omega (\text{single}(\lambda \varepsilon . \text{truish } \varepsilon \rightarrow E[\langle E_1 \rangle] \rho \omega \kappa, \text{send unspecified } \kappa))$$

Here and elsewhere, any expressed value other than *undefined* may be used in place of *unspecified*.

$$E[\langle \text{set! } I \ E \rangle] = \lambda \rho \omega \kappa . E[\langle E \rangle] \rho \omega (\text{single}(\lambda \varepsilon . \text{assign}(\text{lookup } \rho \ I, \varepsilon, \text{send unspecified } \kappa)))$$

$$E^*[\langle \rangle] = \lambda \rho \omega \kappa . \kappa \langle \rangle$$

$$E^*[\langle E_0 \ E^* \rangle] = \lambda \rho \omega \kappa . E[\langle E_0 \rangle] \rho \omega (\text{single}(\lambda \varepsilon_0 . E^*[\langle E^* \rangle] \rho \omega (\lambda \varepsilon^* . \kappa (\langle \varepsilon_0 \rangle \ \$ \ \varepsilon^*))))$$

$$C[\langle \rangle] = \lambda \rho \omega \theta . \theta$$

$$C[\langle \Gamma_0 \ \Gamma^* \rangle] = \lambda \rho \omega \theta . E[\langle \Gamma_0 \rangle] \rho \omega (\lambda \varepsilon^* . C[\langle \Gamma^* \rangle] \rho \omega \theta)$$

D.4 Auxiliary functions

$$\text{lookup} : U \rightarrow \text{Ide} \rightarrow L$$

$$\text{lookup} = \lambda \rho I . \rho I$$

$$\text{extends} : U \rightarrow \text{Ide}^* \rightarrow L^* \rightarrow U$$

$$\text{extends} = \lambda \rho I^* \alpha^* . \#I^* = 0 \rightarrow \rho, \text{extends}(\rho[(\alpha^* \downarrow 1)/(I^* \downarrow 1)])(I^* \uparrow 1)(\alpha^* \uparrow 1)$$

$$\text{wrong} : X \rightarrow C \quad [\text{implementation-dependent}]$$

$$\text{send} : E \rightarrow K \rightarrow C$$

$$\text{send} = \lambda \varepsilon \kappa . \kappa \langle \varepsilon \rangle$$

$$\text{single} : (E \rightarrow C) \rightarrow K$$

$$\text{single} = \lambda \psi \varepsilon^* . \# \varepsilon^* = 1 \rightarrow \psi(\varepsilon^* \downarrow 1), \text{wrong} \text{ "wrong number of return values"}$$

$$\text{new} : S \rightarrow (L + \{\text{error}\}) \quad [\text{implementation-dependent}]$$

$$\text{hold} : L \rightarrow K \rightarrow C$$

$$\text{hold} = \lambda \alpha \kappa \sigma . \text{send}(\sigma \alpha \downarrow 1) \kappa \sigma$$

$$\text{assign} : L \rightarrow E \rightarrow C \rightarrow C$$

$$\text{assign} = \lambda \alpha \varepsilon \theta \sigma . \theta(\text{update } \alpha \varepsilon \sigma)$$

$$\text{update} : L \rightarrow E \rightarrow S \rightarrow S$$

$$\text{update} = \lambda \alpha \varepsilon \sigma . \sigma[(\varepsilon, \text{true})/\alpha]$$

$$\text{tievals} : (L^* \rightarrow C) \rightarrow E^* \rightarrow C$$

$$\text{tievals} = \lambda \psi \varepsilon^* \sigma . \# \varepsilon^* = 0 \rightarrow \psi \langle \rangle \sigma, \text{new } \sigma \in L \rightarrow \text{tievals}(\lambda \alpha^* . \psi(\langle \text{new } \sigma | L \rangle \ \$ \ \alpha^*))(\varepsilon^* \uparrow 1)(\text{update}(\text{new } \sigma | L)(\varepsilon^* \downarrow 1) \sigma), \text{wrong} \text{ "out of memory" } \sigma$$

$$\text{tievalsrest} : (L^* \rightarrow C) \rightarrow E^* \rightarrow N \rightarrow C$$

$$\text{tievalsrest} = \lambda \psi \varepsilon^* \nu . \text{list}(\text{dropfirst } \varepsilon^* \nu)(\text{single}(\lambda \varepsilon . \text{tievals } \psi(\langle \text{takefirst } \varepsilon^* \nu \rangle \ \$ \ \langle \varepsilon \rangle)))$$

$$\text{dropfirst} = \lambda l n . n = 0 \rightarrow l, \text{dropfirst}(l \uparrow 1)(n - 1)$$

$$\text{takefirst} = \lambda l n . n = 0 \rightarrow \langle \rangle, \langle l \downarrow 1 \rangle \ \$ \ (\text{takefirst}(l \uparrow 1)(n - 1))$$

$$\text{truish} : E \rightarrow T$$

$$\text{truish} = \lambda \varepsilon . \varepsilon = \text{false} \rightarrow \text{false}, \text{true}$$

$$\text{permute} : \text{Exp}^* \rightarrow \text{Exp}^* \quad [\text{implementation-dependent}]$$

$$\text{unpermute} : E^* \rightarrow E^* \quad [\text{inverse of permute}]$$

$$\text{applicate} : E \rightarrow E^* \rightarrow P \rightarrow K \rightarrow C$$

$$\text{applicate} = \lambda \varepsilon \varepsilon^* \omega \kappa . \varepsilon \in F \rightarrow (\varepsilon | F \downarrow 2) \varepsilon^* \omega \kappa, \text{wrong} \text{ "bad procedure"}$$

$$\text{onearg} : (E \rightarrow P \rightarrow K \rightarrow C) \rightarrow (E^* \rightarrow P \rightarrow K \rightarrow C)$$

$$\text{onearg} = \lambda \zeta \varepsilon^* \omega \kappa . \# \varepsilon^* = 1 \rightarrow \zeta(\varepsilon^* \downarrow 1) \omega \kappa, \text{wrong} \text{ "wrong number of arguments"}$$

$$\text{twoarg} : (E \rightarrow E \rightarrow P \rightarrow K \rightarrow C) \rightarrow (E^* \rightarrow P \rightarrow K \rightarrow C)$$

$$\text{twoarg} = \lambda \zeta \varepsilon^* \omega \kappa . \# \varepsilon^* = 2 \rightarrow \zeta(\varepsilon^* \downarrow 1)(\varepsilon^* \downarrow 2) \omega \kappa, \text{wrong} \text{ "wrong number of arguments"}$$

$$\text{threearg} : (E \rightarrow E \rightarrow E \rightarrow P \rightarrow K \rightarrow C) \rightarrow (E^* \rightarrow P \rightarrow K \rightarrow C)$$

$$\text{threearg} = \lambda \zeta \varepsilon^* \omega \kappa . \# \varepsilon^* = 3 \rightarrow \zeta(\varepsilon^* \downarrow 1)(\varepsilon^* \downarrow 2)(\varepsilon^* \downarrow 3) \omega \kappa, \text{wrong} \text{ "wrong number of arguments"}$$

$$\text{list} : E^* \rightarrow P \rightarrow K \rightarrow C$$

$$\text{list} = \lambda \varepsilon^* \omega \kappa . \# \varepsilon^* = 0 \rightarrow \text{send null } \kappa, \text{list}(\varepsilon^* \uparrow 1)(\text{single}(\lambda \varepsilon . \text{cons}(\varepsilon^* \downarrow 1, \varepsilon) \kappa))$$

$$\text{cons} : E^* \rightarrow P \rightarrow K \rightarrow C$$

$$\text{cons} = \text{twoarg}(\lambda \varepsilon_1 \varepsilon_2 \omega \sigma . \text{new } \sigma \in L \rightarrow (\lambda \sigma' . \text{new } \sigma' \in L \rightarrow \text{send}(\langle \text{new } \sigma | L, \text{new } \sigma' | L, \text{true} \rangle \text{in } E) \kappa (\text{update}(\text{new } \sigma' | L) \varepsilon_2 \sigma'), \text{wrong} \text{ "out of memory" } \sigma')(\text{update}(\text{new } \sigma | L) \varepsilon_1 \sigma), \text{wrong} \text{ "out of memory" } \sigma)$$

$$\text{less} : E^* \rightarrow P \rightarrow K \rightarrow C$$

$$\text{less} = \text{twoarg}(\lambda \varepsilon_1 \varepsilon_2 \omega \kappa . (\varepsilon_1 \in R \wedge \varepsilon_2 \in R) \rightarrow \text{send}(\varepsilon_1 | R < \varepsilon_2 | R \rightarrow \text{true}, \text{false}) \kappa, \text{wrong} \text{ "non-numeric argument to <"})$$

$$\text{add} : E^* \rightarrow P \rightarrow K \rightarrow C$$

$$\text{add} = \text{twoarg}(\lambda \varepsilon_1 \varepsilon_2 \omega \kappa . (\varepsilon_1 \in R \wedge \varepsilon_2 \in R) \rightarrow \text{send}(\langle \varepsilon_1 | R + \varepsilon_2 | R \rangle \text{in } E) \kappa, \text{wrong} \text{ "non-numeric argument to +"})$$

$$\text{car} : E^* \rightarrow P \rightarrow K \rightarrow C$$

$$\text{car} = \text{onearg}(\lambda \varepsilon \omega \kappa . \varepsilon \in E_p \rightarrow \text{car-internal } \varepsilon \kappa, \text{wrong} \text{ "non-pair argument to car"})$$

$$\text{car-internal} : E \rightarrow K \rightarrow C$$

$$\text{car-internal} = \lambda \varepsilon \omega \kappa . \text{hold}(\varepsilon | E_p \downarrow 1) \kappa$$

$cdr : E^* \rightarrow P \rightarrow K \rightarrow C$ [similar to *car*]

$cdr\text{-internal} : E \rightarrow K \rightarrow C$ [similar to *car-internal*]

$setcar : E^* \rightarrow P \rightarrow K \rightarrow C$

$setcar =$
 $twoarg (\lambda \varepsilon_1 \varepsilon_2 \omega \kappa . \varepsilon_1 \in E_p \rightarrow$
 $(\varepsilon_1 | E_p \downarrow 3) \rightarrow assign (\varepsilon_1 | E_p \downarrow 1)$
 ε_2
 $(send\text{unspecified} \kappa),$
 $wrong \text{ "immutable argument to set-car!"},$
 $wrong \text{ "non-pair argument to set-car!"})$

$eqv : E^* \rightarrow P \rightarrow K \rightarrow C$

$eqv =$
 $twoarg (\lambda \varepsilon_1 \varepsilon_2 \omega \kappa . (\varepsilon_1 \in M \wedge \varepsilon_2 \in M) \rightarrow$
 $send (\varepsilon_1 | M = \varepsilon_2 | M \rightarrow true, false) \kappa,$
 $(\varepsilon_1 \in Q \wedge \varepsilon_2 \in Q) \rightarrow$
 $send (\varepsilon_1 | Q = \varepsilon_2 | Q \rightarrow true, false) \kappa,$
 $(\varepsilon_1 \in H \wedge \varepsilon_2 \in H) \rightarrow$
 $send (\varepsilon_1 | H = \varepsilon_2 | H \rightarrow true, false) \kappa,$
 $(\varepsilon_1 \in R \wedge \varepsilon_2 \in R) \rightarrow$
 $send (\varepsilon_1 | R = \varepsilon_2 | R \rightarrow true, false) \kappa,$
 $(\varepsilon_1 \in E_p \wedge \varepsilon_2 \in E_p) \rightarrow$
 $send ((\lambda p_1 p_2 . ((p_1 \downarrow 1) = (p_2 \downarrow 1) \wedge$
 $(p_1 \downarrow 2) = (p_2 \downarrow 2))) \rightarrow true,$
 $false)$
 $(\varepsilon_1 | E_p)$
 $(\varepsilon_2 | E_p))$
 $\kappa,$
 $(\varepsilon_1 \in E_v \wedge \varepsilon_2 \in E_v) \rightarrow \dots,$
 $(\varepsilon_1 \in E_s \wedge \varepsilon_2 \in E_s) \rightarrow \dots,$
 $(\varepsilon_1 \in F \wedge \varepsilon_2 \in F) \rightarrow$
 $send ((\varepsilon_1 | F \downarrow 1) = (\varepsilon_2 | F \downarrow 1) \rightarrow true, false)$
 $\kappa,$
 $send\ false \kappa)$

$apply : E^* \rightarrow P \rightarrow K \rightarrow C$

$apply =$
 $twoarg (\lambda \varepsilon_1 \varepsilon_2 \omega \kappa . \varepsilon_1 \in F \rightarrow valueslist \varepsilon_2 (\lambda \varepsilon^* . apply \varepsilon^* \omega \kappa),$
 $wrong \text{ "bad procedure argument to apply"})$

$valueslist : E \rightarrow K \rightarrow C$

$valueslist =$
 $\lambda \varepsilon \kappa . \varepsilon \in E_p \rightarrow$
 $cdr\text{-internal} \varepsilon$
 $(\lambda \varepsilon^* . valueslist$
 ε^*
 $(\lambda \varepsilon^* . car\text{-internal}$
 ε
 $(single (\lambda \varepsilon . \kappa (\langle \varepsilon \rangle \S \varepsilon^*))))),$
 $\varepsilon = null \rightarrow \kappa \langle \rangle,$
 $wrong \text{ "non-list argument to values-list"}$

$cwcc : E^* \rightarrow P \rightarrow K \rightarrow C$ [call-with-current-continuation]

$cwcc =$
 $onearg (\lambda \varepsilon \omega \kappa . \varepsilon \in F \rightarrow$
 $(\lambda \sigma . new \sigma \in L \rightarrow$
 $apply \varepsilon$
 $\langle \langle new \sigma | L,$
 $\lambda \varepsilon^* \omega' \kappa' . travel \omega' \omega (\kappa \varepsilon^*) \rangle$
 $in E \rangle$
 ω
 κ
 $(update (new \sigma | L)$

$unspecified$
 $\sigma),$
 $wrong \text{ "out of memory" } \sigma),$
 $wrong \text{ "bad procedure argument"}$

$travel : P \rightarrow P \rightarrow C \rightarrow C$

$travel =$
 $\lambda \omega_1 \omega_2 . travelpath ((pathup \omega_1 (commonancest \omega_1 \omega_2)) \S$
 $(pathdown (commonancest \omega_1 \omega_2) \omega_2))$

$pointdepth : P \rightarrow N$

$pointdepth =$
 $\lambda \omega . \omega = root \rightarrow 0, 1 + (pointdepth (\omega | (F \times F \times P) \downarrow 3))$

$ancestors : P \rightarrow PP$

$ancestors =$
 $\lambda \omega . \omega = root \rightarrow \{\omega\}, \{\omega\} \cup (ancestors (\omega | (F \times F \times P) \downarrow 3))$

$commonancest : P \rightarrow P \rightarrow P$

$commonancest =$
 $\lambda \omega_1 \omega_2 . \text{the only element of}$
 $\{\omega' \mid \omega' \in (ancestors \omega_1) \cap (ancestors \omega_2),$
 $pointdepth \omega' \geq pointdepth \omega''$
 $\forall \omega'' \in (ancestors \omega_1) \cap (ancestors \omega_2)\}$

$pathup : P \rightarrow P \rightarrow (P \times F)^*$

$pathup =$
 $\lambda \omega_1 \omega_2 . \omega_1 = \omega_2 \rightarrow \langle \rangle,$
 $\langle (\omega_1, \omega_1 | (F \times F \times P) \downarrow 2) \rangle \S$
 $(pathup (\omega_1 | (F \times F \times P) \downarrow 3) \omega_2)$

$pathdown : P \rightarrow P \rightarrow (P \times F)^*$

$pathdown =$
 $\lambda \omega_1 \omega_2 . \omega_1 = \omega_2 \rightarrow \langle \rangle,$
 $(pathdown \omega_1 (\omega_2 | (F \times F \times P) \downarrow 3)) \S$
 $\langle (\omega_2, \omega_2 | (F \times F \times P) \downarrow 1) \rangle$

$travelpath : (P \times F)^* \rightarrow C \rightarrow C$

$travelpath =$
 $\lambda \pi^* \theta . \# \pi^* = 0 \rightarrow \theta,$
 $((\pi^* \downarrow 1) \downarrow 2) \langle (\pi^* \downarrow 1) \downarrow 1 \rangle$
 $(\lambda \varepsilon^* . travelpath (\pi^* \dagger 1) \theta)$

$dynamicwind : E^* \rightarrow P \rightarrow K \rightarrow C$

$dynamicwind =$
 $threearg (\lambda \varepsilon_1 \varepsilon_2 \varepsilon_3 \omega \kappa . (\varepsilon_1 \in F \wedge \varepsilon_2 \in F \wedge \varepsilon_3 \in F) \rightarrow$
 $apply \varepsilon_1 \langle \rangle \omega (\lambda \zeta^* .$
 $apply \varepsilon_2 \langle \rangle ((\varepsilon_1 | F, \varepsilon_3 | F, \omega) in P)$
 $(\lambda \varepsilon^* . apply \varepsilon_3 \langle \rangle \omega (\lambda \zeta^* . \kappa \varepsilon^*))),$
 $wrong \text{ "bad procedure argument"}$

$values : E^* \rightarrow P \rightarrow K \rightarrow C$

$values = \lambda \varepsilon^* \omega \kappa . \kappa \varepsilon^*$

$cwnv : E^* \rightarrow P \rightarrow K \rightarrow C$ [call-with-values]

$cwnv =$
 $twoarg (\lambda \varepsilon_1 \varepsilon_2 \omega \kappa . apply \varepsilon_1 \langle \rangle \omega (\lambda \varepsilon^* . apply \varepsilon_2 \varepsilon^* \omega))$