

Portable and high-level access to the stack with
Continuation Marks

A dissertation presented
by

John Clements

to
the College of Computer and Information Science

In partial fulfillment of the requirements for the degree of
Doctor of Philosophy

in the field of

Computer Science

Northeastern University
Boston, Massachusetts

February 8, 2006

©2005
John Brinckerhoff Clements
All Rights Reserved

Abstract

My dissertation presents and defends the thesis that lightweight stack inspection improves the implementation of intensional programming tools and programming language extensions.

Intensional programming tools, such as debuggers, need to observe and change the function call behavior of programs as they run. Intensional programming language extensions, such as aspect-oriented programming, behave similarly. Traditionally, these tools and extensions are granted privileged access to the evaluator's implementation details, particularly the stack. This design can stall further development, particularly the extension of these tools to new platforms and new versions of the evaluator.

To solve some of these problems, we propose an alternate architecture. An extension of the language, "continuation marks," allows these tools to obtain runtime information without special privilege, and to operate as annotators on the source language. We claim that this language feature is a simple addition to many runtime systems. We further claim that this architecture results in tools and extensions that are more portable and well-defined.

In order to support our claims, we have successfully applied it to three different problems. First, we designed and implemented a stepping debugger. Second, we showed how a security feature, stack inspection, may be expressed. Finally, we consider the implementation of aspects, and particularly those aspects that require information about control flow.

A Failed Dedication

No dedication could possibly do justice to the debt I owe to my wife, Anika.

Contents

Contents	vii
List of Figures	ix
1 Introduction	1
2 Continuation Marks	3
2.1 Intuition	3
2.2 The Importance of Tail-Calling	6
2.3 An Implementation in Scheme	9
3 A Stepper built with Continuation Marks	17
3.1 A Stepper for Scheme	17
3.2 The Continuation Mark Architecture	18
3.3 Modeling a Stepper	24
3.4 Correctness	29
3.5 Notes on Proof Methods	50
3.6 The Pragmatics of a Stepper	51
3.7 Related Work	53
4 Stack Inspection implemented with Continuation Marks	55
4.1 Stacks, Security, and Tail Calls	55
4.2 The λ_{sec} Language	57
4.3 Tail-calling	60
4.4 An Abstract Machine for λ_{sec}	61
4.5 An Alternative Implementation	64
4.6 Space Consumption	68
4.7 An implementation using Continuation Marks	72
4.8 The Richer Models of Java and .Net	72
4.9 Related Work	73
5 Conclusion	75
5.1 Results	75
5.2 Future Directions	76

A Stepper Appendices	77
A.1 Annotation and Reconstruction Definitions	77
A.2 Evaluator Fragments	83
B Stack Inspection Appendices	89
B.1 Equivalence of Fournet and Gordon’s evaluator and the one presented herein	89
C Aspect-Oriented Programming using Continuation Marks	93
C.1 Aspect-Oriented Programming	93
C.2 Implementing AOP	95
C.3 CFlow and Tail-Calling	96
Bibliography	97

List of Figures

2.1	A first example using continuation marks	5
2.2	An example using continuation marks with different keys	5
2.3	A hierarchy of graphical classes	7
2.4	Marks in tail position	8
2.5	The relationship between continuation-marks and tail-recursion	9
2.6	The Annotation for a Lambda in Trace	11
2.7	The expansion of parameterize	13
2.8	The macro that expands with-handlers	15
3.1	Syntax	20
3.2	Machine Definitions	21
3.3	Expression Reductions	22
3.4	Value Reductions	23
3.5	The Stepper Architecture	25
3.6	Annotation Functions	26
3.7	Configuration Annotation Functions	27
3.8	Reconstruction Functions	29
3.9	Debug Value Template	30
3.10	The Simulation Lemma	31
3.11	Steps taken in reduction of if to test position	36
3.12	Steps taken in reduction of cons1 continuation	37
3.13	Steps taken in reduction of wcm continuation	38
3.14	Steps taken in reduction of if continuation leading to an error	40
3.15	Steps taken in reduction of cons2 continuation	41
3.16	Steps taken in reduction of car continuation leading to an error	42
3.17	Steps taken in reduction of empty continuation	43
3.18	Steps taken in reduction of empty continuation with no further definitions	44
3.19	Steps taken in reduction of loaded annotated definitions	45
3.20	The correspondence between steps taken and steps emitted	47
3.21	Transition Categories	47
4.1	The FG Machine	62
4.2	Computing permissions as a finite state automaton	65
4.3	The CM Machine	66

C.1 A simple piece of advice	94
C.2 Capturing context information using continuation marks	95

Chapter 1

Introduction

Debugging tools and certain language extensions need to observe and change programs as they run. Steppers, debuggers, and profilers all gather information about running programs. Stack inspection extracts runtime data to verify security properties. Aspect-oriented-programming uses dynamic context to decide which code to evaluate.

Each of these needs to inspect control information. The traditional solution in each case is to “hard-wire” low-level access to the inner machinery of the evaluator. As a result of this choice, these tools are fragile and often difficult to port. This architecture also makes it extremely difficult to specify the meaning of these tools or language features in a high-level way.

Our thesis is that this Faustian compromise is unnecessary. We can give these tools the information they need through a narrow and portable interface without binding them to a particular set of machine details. Our dissertation supports this thesis with the introduction of *continuation marks*, a language feature that allows programs to observe control information without exposing details of the evaluator’s implementation, and a demonstration of their power.

Using continuation marks, tools such as debuggers and profilers may be designed at the level of the language itself, rather than as approximate translations from exposed machine state to configurations of a higher-level source language. Furthermore, language features such as stack inspection and aspect-oriented programming may be specified as transformations from a source language with these features to a simpler language with just continuation-marks. In essence, we show that they may be specified as “macros,” rather than language features.

We have described the meaning of continuation marks by adding them to a reduction semantics and also to a simple register machine. The register machine we use is quite general, and therefore we believe that continuation marks may be added to a wide range of existing languages, including both Java and C#. Moreover, we believe that adding continuation marks to C# may allow us to lift the current restriction that tail-calls cannot be handled properly in the presence of security tracking.

This dissertation consists of five chapters, including the introduction. The second chapter describes continuation marks; their syntax, their semantics, and also the pragmatics of implementing them in a working evaluator for the Scheme programming language. The third chapter builds upon the stated semantics to show how a stepper may be specified and implemented as a source-to-source transformation for programs coupled with a runtime reconstructor. The fourth chapter illustrates how stack inspection may be implemented using continuation marks. A fifth chapter concludes the dissertation.

Chapter 2

Continuation Marks

This chapter introduces continuation marks. It begins by providing an intuition, both for the reasons that motivated their creation and for their actual behavior. The next section explains the importance of tail-calling language evaluators. The third section details the implementation of continuation marks that appears in MzScheme, and shows how they are used to accomplish various tasks.

2.1 Intuition

Continuation marks provide a means to observe control behavior. More specifically, they allow the innermost part of the control context—the evaluation of a method or procedure, for instance—to obtain information about what other contexts are currently active. To put it another way, they allow a program to inspect its own call stack.

However, continuation marks are also designed *not* to reveal any information not available to those contexts themselves. To take one example, continuation marks should not expose the way in which the evaluator lays out the activation record.

More subtly, continuation marks should not make it possible to observe control information that is a consequence of an evaluator implementation or a subsequent code rewriting. For instance, a program should not be able to use continuation marks to count the total number of frames on the stack, as this would allow the program to observe certain optimizations—thereby rendering them unsound.

Continuation marks strike a balance between these two goals by limiting the possible observations to values already available to the program, and by linking the visibility and duration of these values to notions of language definition, rather than to compiler representations.

In any abstract machine, a continuation consists of a sequence of frames. Continuation marks provide a means to label these frames with program values, permitting a run-time observation of the dynamic program state. A language

with continuation marks includes two new operations; **w-c-m**, short for **with-continuation-mark**, and **c-c-m**, short for **current-continuation-marks**. The first places a mark on the most recent frame, and the second retrieves the marks associated with all frames currently on the stack.

Continuation marks are intended as a means of instrumenting programs without changing their behavior. If the program being instrumented—henceforth the *source* program—uses continuation marks itself, some mechanism is needed to prevent the later marks from interfering with the earlier ones. Marks therefore are associated with *keys*, and marks associated with different keys are mutually transparent.

This dissertation describes a number of different models for languages that include continuation marks. Each one is tailored for a different purpose, and each one contains a slightly different definition of continuation marks. The following two definitions, taken from chapter 3, present the intuition behind continuation marks clearly and concisely.

(**w-c-m** *key mark-expr body-expr*) : *key* is chosen from a fixed but infinite set of keys. The *mark-expr* and *body-expr* are arbitrary expressions. The *mark-expr* evaluates to a mark-value, which is then associated with the given key in the continuation’s innermost (that is, most recently added) frame. If the current top frame already has a value associated with this key, the new value replaces it. Finally, *body-expr* is evaluated. Its value becomes the result of the entire **w-c-m** expression.

(**c-c-m** *key ...*) : This expression collects values associated with the given set of keys by traversing the frames in the continuation from innermost to outermost. Any frame whose mark table contains one or more of the named keys results in an association list containing (quoted) keys and values. Frames without any of the named keys are entirely unrepresented in the result of **c-c-m**.

Abstract machines vary in the granularity of their continuation frames. Models such as Landin’s SECD machine have coarse-grained continuation frames, where a frame captures the state of a procedure call’s progress, including a list of remaining instructions and a stack of intermediate results. The continuation (or “dump”, to use Landin’s term) is extended when a procedure call occurs. Other models, such as Felleisen’s CEK machine, use finer-grained continuation frames that capture the information necessary to evaluate a single expression. Every expression whose computation would require an “intermediate value” in the SECD machine generates an additional continuation frame in the CEK machine. This latter model vastly simplifies the specification of continuation marks, and we shall define continuation-mark primitives exclusively with respect to these models.

Figure 2.1 shows a simple example using continuation marks. In this example, a simple **if** expression has been instrumented with two mark-placing

```

(w-c-m k 'around-if
 (if (w-c-m k 'around-test
      (begin (display (c-c-m k))
              false)
      3
      4)))

```

↪

```
(((k 'around-test)) ((k 'around-if)))
```

4

Figure 2.1: A first example using continuation marks

operations—one around the whole **if**, and one around the **if**'s test. In addition, the test expression (in this case, simply **false**) is wrapped in a **begin** that displays all marks with the key k that are associated with the continuation stack. Since the context of the test expression is different from that of the **if** expression, both marks remain.

```

(w-c-m k 'around-if
 (w-c-m k2 'another-around
  (if (w-c-m k 'around-test
      (begin (display (c-c-m k2))
              false)
      3
      4)))

```

↪

```
(((k2 'another-around)))
```

4

Figure 2.2: An example using continuation marks with different keys

Figure 2.2 shows a similar example that uses two different keys. Since the newly added **w-c-m** uses a different key ($k2$) than the other **w-c-ms**, its marks and the existing marks are mutually transparent. In particular, the **c-c-m** call shows only one frame, because only one mark with this key was placed on the continuation.

The example and description given thus far suggest a simple implementation of continuation marks. In particular, it would seem that a program with continuation marks could be rewritten into one without them by wrapping every expression with a “push” to a stack kept on the side, and a corresponding “pop” of the mark stack on every expression’s return. Aside from the appalling expense of such an approach, there are two key reasons that this approach is infeasible: first, the importance of tail-calling in language implementations, and second, the importance of library code.

2.2 The Importance of Tail-Calling

The Past

There are many models for the behavior of computers in evaluating programs. Some of them, like Plotkin’s β_v variant of the lambda-calculus, model procedure application using a substitution of the procedure’s body for the application. Others, like Landin’s SECD machine, model procedure application using a dump that grows on every procedure call.

One difference between these two models arises when, for instance, a procedure’s f ’s body is simply an application of another function g . Using the β_v rule guarantees that the body of f is replaced by the body of g , and the resulting term is no larger than it would have been had the original call been to g . In the SECD machine, on the other hand, the call to f and f ’s call to g each generate a stack frame, and the computation of g now proceeds in an environment where the stack is larger than it would have been in a direct call to g .

This “extra frame” in the SECD machine does not affect the result of the computation—the result of the call to g is returned unchanged to the caller of f —but the extra frame consumes memory, and a simple loop of such direct calls may exhaust memory.

This behavior of the SECD machine, and more broadly of many compilers and evaluators, is principally a consequence of the early evolution of computers and computer languages; procedure calls, and particularly recursive procedure calls, were added to many computer languages long after constructs like branches, jumps, and assignment. They were thought to be expensive, inefficient, and esoteric.

Guy Steele [43] outlined this history and disrupted the canard of expensive procedure calls, showing how inefficient procedure call mechanisms could be replaced with simple “JUMP” instructions, making the space complexity of recursive procedures equivalent to that of any other looping construct. This work applies with little change to the evaluators of the present day. The key is to ensure that calls made when “no work remains to be done” in the current procedure do not add a frame to the stack. We use the name “tail-calling” or “properly tail-recursive” for evaluators that have this property.

The Present

In the past twenty years, inductively defined data structures (lists, trees, and the like) have become increasingly prevalent. Felleisen et al. [14]—among others—observe that inductive data definitions give rise to inductively defined operations on such data (sorting, searching, etc.), and that these correspond directly to recursive functions. This approach narrows the distance between specification and implementation, and greatly simplifies the task of teaching students how to design operations on these data types.

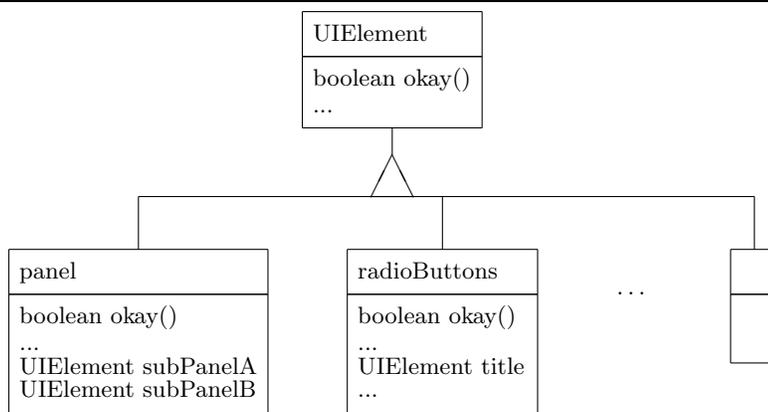


Figure 2.3: A hierarchy of graphical classes

Another push in this direction has come from object-oriented programming. In a pure object-oriented model, the code that operates on data is contained in methods associated with the classes that make up the data. That is, a computation over a piece of data composed of elements of two classes will naturally result in two method bodies, one in each class. If these classes’ instances refer to each other, then the methods defined for the data will naturally refer to the corresponding method in the other class.

For example, consider a hierarchy of GUI elements. Figure 2.3 shows a part of what such a hierarchy might look like. Each element of the hierarchy has an `okay()` method that checks recursively to make sure that this element and all enclosed elements are well-formed. The natural implementation of the `okay()` method in each class would involve recursive calls to the `okay()` methods of the contained objects. Since many of these objects will contain one or two sub-elements, many of the recursive calls will be tail calls—that is, calls with no work remaining to be done. Unfortunately, the lack of tail-calling in Java makes this pattern of recursive calls needlessly wasteful of memory, leaving memory-conscious programmers with no choice but to maintain a stack explicitly in order to translate this traversal into a loop. Note that a “Visitor”-style traversal suffers from exactly the same problem.

Finally, a key early motivation for the creation of object-oriented languages is the elimination of mutation.

Though OOP came from many motivations, two were central.
 ... [T]he small scale one was to find a more flexible version of
 assignment, and then to try to eliminate it altogether.

— Alan Kay, *History of Smalltalk* (1993)

Unfortunately, the absence of tail-calling in Java means that programmers must code traversals of self-referential data (trees, lists, stacks, etc.) using looping

constructs and mutation. That is, a tail-calling language would allow a purer OO-style traversal with the *added* benefit of eliminating mutation.

Dismayingly, the majority of compilers and evaluators—even for object-oriented languages—still lack the simple change that would make the space consumed by recursive procedures equivalent to that consumed by a loop.

Ironically, the very compilers that fail to provide tail-calling to the programs they compile now use intermediate representations that are derived from transformations such as CPS [12], ANF [20], and SSA [29]. These transformations are based on the premise that code sequences are in essence nested tail-calling procedures; unfortunately, this insight has not risen to the surface of the language, allowing programmers to use the lightweight calling convention that their optimizing compilers do.

The Future

Awareness of the need for tail-calling compilers and interpreters is growing. Some languages, such as Scheme [28], have included a tail-calling requirement for many years. That is, all “correct” implementations of Scheme must evaluate certain divergent programs without unbounded space growth.

More recently, Microsoft’s .NET framework [34] has begun to address the need for tail-calling languages by providing a tail-calling primitive. Unfortunately, the use of the tail-calling primitive comes with restrictions; in particular, the use of the native stack inspection mechanism disables the tail-calling behavior. In fact, I and my co-authors demonstrate a possible alternative to this restriction which is described in chapter 4. Finally, this tail-calling primitive turns out to be prohibitively slower than its traditional push-and-pop alternative, due apparently to a runtime check for stack-inspection information.

The Continuation Marks

The importance of tail-calling highlights the need for continuation marks. Since continuation marks observe the behavior of the evaluator rather than anticipating it, the continuation mark architecture can be applied without foreknowledge of the stack’s behavior.

```

(w-c-m if-example 'outer
 (if true
  (w-c-m if-example 'inner
   (c-c-m 'if-example))
  13))
  ↪
(((if-example inner)))

```

Figure 2.4: Marks in tail position

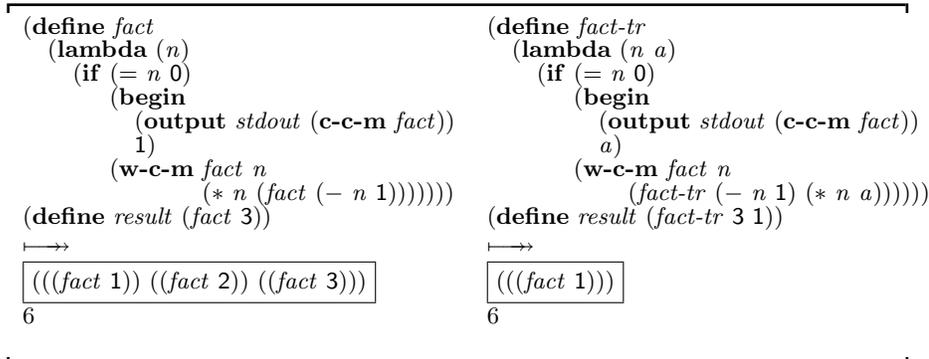


Figure 2.5: The relationship between continuation-marks and tail-recursion

To see this, we consider several simple examples using continuation marks in Scheme. In figure 2.4, a **w-c-m** is wrapped around the **if**, and another one around the enclosed **then** clause. The **then** clause (and also the **else** clause) are in tail position with respect to the **if**, because the **then** clause replaces the whole **if** when the test evaluates to **true**. That is, the context of the **if** will be the same as the context of the **then** clause. As a result, the inner mark and the outer mark are both applied to the same context, and the inner mark overwrites the outer one. This example differs from the earlier example of figure 2.1 in that the earlier example showed a wrapping around the *test* clause, which is not in tail position.

The two programs in fig. 2.5 illustrate how a programmer might instrument a factorial function with these constructs. The boxed text represents the program's output. Both definitions implement a factorial function that marks its continuation at the recursive call site and reports the continuation-mark list before returning. The one in the left column is properly recursive, the one on the right is tail-recursive. The values of V are the outputs that applications of their respective functions produce. For the properly recursive program on the left, the value shows that the continuation contains three mark frames. For the tail-recursive variant, only one continuation mark remains; the others have been overwritten during the evaluation.

2.3 An Implementation in Scheme

Matthew Flatt's MzScheme is a bytecode interpreter written in **C**. It maintains three separate stacks; the standard **C** call stack, a scheme stack for activation variables and application arguments, and a third one for continuation marks.

When a mark is set using **w-c-m**, MzScheme scans the frames of the continuation-mark stack associated with the topmost activation record to find existing marks with this key. If one exists, its binding is mutated. Otherwise, a new frame is added to this stack. MzScheme's implementation for **w-c-m**

takes time linear in the number of keys used for marks on the topmost frame. This has not been a bottleneck in practice, because programs typically use a small number of keys. Alternative implementations could reduce the asymptotic bound, but would likely result in higher average times, due to increased overhead. Since this stack is copied when a continuation is captured, a mutation is not visible to other references to this continuation.

When the primitive procedure **c-c-m** is called, MzScheme scans the full stack for instances of the given key or keys and collects the results in a list. This operation is linear in the size of the continuation-mark stack. MzScheme also includes a *continuation-mark-set-first* primitive that identifies only the topmost mark associated with a given key. By caching values, this operation can in principle be implemented in amortized-constant time. The former is typically useful in operations like breakpoints where time is not an issue, where the latter is useful in more time-sensitive applications, like checking permissions or looking up exception handlers.

MzScheme uses continuation marks to implement a wide variety of tools and language features. These include tools such as the stepper, the errortrace facility, the profiler, a trace engine, and a prototype debugger. Continuation marks also form the basis for language features such as parameters and exceptions.

The Stepper

Chapter 3 explores the stepper in depth. Continuation marks are the technology that enables the implementation of the stepper with respect to a high-level model.

Errortrace

DrScheme provides an errortrace facility which is moderately fast and shows source-position backtraces when errors occur. That is, an uncaught exception comes with a set of continuation marks captured when the exception occurred. The user has the option of displaying these as a list of source positions (along with the actual text at these positions), where each element in the list corresponds to a continuation frame. An alternative display links the source positions of the continuation frames with a chain of arrows, superimposed upon the source text in the editor window.

The annotation associated with errortrace is simple. Each expression is wrapped in a **w-c-m** that associates the source position with an errortrace key. If expression *B* is tail with respect to expression *A*, then the mark placed by the wrapper around *B* will replace the one associated with *A*. This guarantees that only the source positions associated with current continuation frames will be present in the current set of marks.

Profiler

DrScheme provides a profiler which reports the time spent in each procedure. Additionally, the profiler can supply information about call paths; this information is collected through an annotation that uses continuation marks to capture calling information.

In fact, the profiler re-uses the Errortrace annotator. Since the two tools use distinct keys for their continuation marks, these two annotations are mutually transparent.

Trace

```

...
[lambda-clause-abstraction
 (lambda (clause)
  (kernel-syntax-case clause #f
   [(arglist . bodies)
    (let-values
     ([[arglist-proper improper?] (arglist-flatten #'arglist)])
      (if name-guess
         #'(arglist
             (with-continuation-mark
              #,calltrace-key
              'unimportant
              (begin (let ([call-depth (length (continuation-mark-set→list
                                                (current-continuation-marks)
                                                #,calltrace-key))])
                      (#,print-call-trace
                       (quote-syntax #,name-guess)
                       #,(syntax-original? name-guess)
                       (#,stx-protector-stx #,(make-stx-protector stx))
                       (list #,@arglist-proper)
                       #,improper?
                       call-depth))
                    #,@(recur-on-sequence (syntax->list #'bodies))))))
          #'(arglist #,@(recur-on-sequence (syntax->list #'bodies))))))
    [else
     (error 'expr-syntax-object-iterator
            "unexpected (case-)lambda clause: ~a"
            (syntax-object->datum stx))]]])
...

```

Figure 2.6: The Annotation for a Lambda in Trace

MzScheme includes a very simple trace utility that instruments each procedure with code that displays its arguments. Additionally, each procedure's

body is wrapped in a constant with-continuation-mark; this allows the trace code to determine its call depth, and indent the output accordingly.

This tiny utility is nevertheless a clear illustration of the value of continuation marks. Keeping track of this information without continuation marks would require changing the calling convention, exposing the format of the stack, or the use of fluids—which are in fact built using continuation marks.

Figure 2.6 shows code taken directly from the definition of the annotation used for the Trace utility. The function *lambda-clause-abstraction* accepts MzScheme syntax representing a **lambda** clause (without the **lambda** itself) and wraps its body in a **with-continuation-mark** (abbreviated in this chapter as **w-c-m**). The depth of the recursion is computed at runtime by taking the length of the list of continuation marks revealed by **current-continuation-marks** (abbreviated in this chapter as **c-c-m**).

Debugger

DrScheme has a prototype debugger, called MzTake, built upon yet another annotation using continuation marks. MzTake’s annotation is similar to Error-trace, but in addition captures the values bound to variables.

Parameters

Often, programs are parameterized by certain global attributes—an output port, for instance, or a current directory. This is particularly true of languages used for scripting, where interaction with the file system or other program invocations is common. Moreover, these attributes are typically assigned in a dynamic-extent style, where the value of an attribute is valid during the dynamic extent of a call, rather than behaving like a lexically scoped variable.

The traditional implementation of these attributes is simply to have a mutable table of global variables. For instance, the UNIX `execve()` call includes an `envp` parameter that indicates what values for these environment variables should be given to a child process. Each process then has its own copy of the environment variables, which it may manage as it wishes.

Languages like Scheme feature a more sophisticated and fluid interaction between computations. First, a program may consist of many threads, each with its own values for these attributes. Beyond this, Scheme’s continuation mechanisms create additional opportunities for confusion, where a computation may migrate from one thread to another. A model in which programmers have to manually anticipate all possible continuation and thread operations in order to properly update the dynamic values of these attributes would be difficult for small programs and unmanageable at large scales.

The Scheme language defines a **dynamic-wind** primitive that is designed to address this need [28]. The **dynamic-wind** form allows a program to specify actions to be taken whenever a block begins or is entered by invoking a continuation, and likewise to specify actions to be taken whenever a block ends or is exited via continuation. Fluid bindings are built on top of dynamic-wind,

by specifying a swap between the existing and a “saved” binding on both entry and exit from the specified scope. These language forms are frequently sufficient, but suffer from at least three problems.

First, they require additional computation. That is, when a continuation is invoked, the evaluator must check to see what *after* actions are currently pending. Likewise, entering a continuation requires a check to see what *before* actions are required. This is particularly onerous for systems that use continuations as the basis for thread-switching.

Secondly, the body of a **dynamic-wind** is not in tail position with respect to the **dynamic-wind** itself. This means that a loop whose body is a **dynamic-wind**—implementing an exception-handler, for instance—will rapidly exhaust memory.

Third, the existing fluid-let mechanism is incompatible with context switches other than continuation jumps. That is, since the unwinding of a fluid requires the execution of code, a simple halt to a running thread will not typically cause the execution of the unwinding code.

What is needed instead is a passive association between continuations and the bindings of these attributes, and this is precisely what continuation-marks provide. Designed by Matthew Flatt and Michael Sperber, MzScheme’s **parameterize** form associates a new binding with the current continuation, using a continuation-mark. The ‘current binding’ associated with a given parameter is therefore obtained by inspecting the current continuation marks. This definition, more declarative in style, ensures that the maintenance of the fluid invariants—that is, that a given binding is visible only in a certain region of code—is enforced without need of explicit mutation.

```
(parameterize ([a 13])
  (+ 3 4))
```

expands to

```
(with-continuation-mark parameterization-key
  (extend-parameterization
    (continuation-mark-set-first #f parameterization-key)
    a
    13)
  (+ 3 4))
```

Figure 2.7: The expansion of **parameterize**

Figure 2.7 shows the expansion of a simple **parameterize** expression in MzScheme. The result is a **with-continuation-mark** that associates the parameterization key with an extended parameterization. The parameterization-key is known only by the syntax system, to ensure that marks associated with parameters are orthogonal to all prior annotations. The primitive procedure

continuation-mark-set-first is an optimized version of **current-continuation-marks** that extracts only the most recent continuation mark with a given key. The *extend-parameterization* extends this parameterization by mapping the given parameter to a new binding.

When a parameter's value is needed, MzScheme uses *continuation-mark-set-first* to obtain the most recent parameterization, and looks up the parameter's value in this table.

Exceptions

Exception-handlers are much like parameters. Placing an exception handler associates pairs of procedures with the continuation (each one a predicate and a handler). This exception handler should have the dynamic extent of its body, and should not incur additional cost on entry to and exit from the block it controls. It is therefore natural to implement exception-handling using parameters, and therefore using continuation marks.

Figure 2.8 shows the macro definition for *with-handlers* defined by Matthew Flatt for MzScheme. The core of the expansion is the **parameterize** that installs the new procedure as the *current-exception-handler*. Note also that the pair of **with-continuation-mark** expressions using the *break-enabled* key essentially open a “gap” in the continuation where breaks are disabled, so that a procedure invoking the escape continuation *k* is evaluated in an environment where breaks are disabled.

```

(quasisyntax/loc str
  (let ([l (list (cons pred handler) ...)]
        [body (lambda () expr1 expr ...)])
    ;; Capture current break parameterization, so we can use it to
    ;; evaluate the body
    (let ([bpz (continuation-mark-set-first #f break-enabled-key)])
      ;; Disable breaks here, so that when the exception handler jumps
      ;; to run a handler, breaks are disabled for the handler
      (with-continuation-mark
       break-enabled-key
       (make-thread-cell #f)
       ((call/ec
        (lambda (k)
          ;; Restore the captured break parameterization for
          ;; evaluating the 'with-handlers' body. In this
          ;; special case, no check for breaks is needed,
          ;; because bpz is quickly restored past call/ec.
          ;; Thus, 'with-handlers' can evaluate its body in
          ;; tail position.
          (with-continuation-mark
           break-enabled-key
           bpz
           (parameterize
            ([current-exception-handler
             (lambda (e)
              (k
               (lambda ()
                 (let loop ([l l])
                   (cond
                    [(null? l)
                     (raise e)]
                    [((caar l) e)
                     #,(if disable-break?
                          #'(begin0
                            ((cdar l) e)
                            (with-continuation-mark
                             break-enabled-key
                             bpz
                             (check-for-break)))
                          #'(with-continuation-mark
                             break-enabled-key
                             bpz
                             (begin
                               (check-for-break)
                               ((cdar l) e))))))
                    [else
                     (loop (cdr l))])))
                (call-with-values body
                 (lambda args (lambda () (apply values args))))))))))))))

```

Figure 2.8: The macro that expands **with-handlers**

Chapter 3

A Stepper built with Continuation Marks

The earliest motivation for the introduction of continuation marks was the need for a Scheme stepper. This chapter provides a detailed model of a language with continuation marks, the definition of a stepper as an annotation from source programs to instrumented ones, and a proof that the resulting program acts as a correct stepper, showing the same evaluation steps that the original program produces.¹

3.1 A Stepper for Scheme

Our DrScheme programming environment [19] provides an algebraic stepper for Scheme. It explains a program's execution as a sequence of reduction steps based on the ordinary laws of algebra for the functional core [2, 37] and more general algebraic laws for the rest of the language [17]. An algebraic stepper is particularly helpful for teaching: students often have trouble with the key concepts of recursion, and seeing the evaluation of the program in a step-by-step fashion can enlighten them. Selective uses can also provide excellent information for complex debugging situations.

Our stepper implementation must satisfy several requirements. First, it must faithfully match the evaluator. That is, running the program in the stepper must produce the same result as running the compiled program. Second, it must be portable and high-level. That is, it should not be tied to a particular machine architecture, or even to a particular implementation of its language.

One possible strategy involves formulating the language being stepped as a set of reduction rules, and coding the stepper as a translation of those rules into a program. Unfortunately, this is immensely labor-intensive, requiring a second implementation of the language evaluator. Even if the stepper succeeds in duplicating the meaning of the language, subsequent changes to the language

¹This chapter was published in a much earlier form at the European Symposium on Programming [9].

itself will almost certainly find their way into the principal, non-stepping evaluator/compiler before the stepper is updated. The result is typically a slow divergence between the semantics represented by the compiler and the semantics represented by the stepper.

At the other end of the spectrum, systems like `gdb` allow a direct observation of the evaluation performed by the compiled program. In this case, the semantics of the stepper is likely to match the semantics of the compiled program closely. Unfortunately, linking a stepper’s definition to knowledge of the evaluator’s internals is expensive, fragile, and non-portable; every new back-end must include a corresponding piece of the stepper, and any change to the evaluator’s internals must be accompanied by a change to the stepper. Furthermore, it make a proof of correctness essentially impossible.

In our stepper, we balance these requirements by using continuation marks. The stepper transforms a source program into one that uses continuation marks to store information about its current source position and bindings. This annotated program also gathers all current continuation marks at each step, and uses this information to display the state of execution. Since the stepper is defined as a source-to-source translation, the correctness of the debugger depends only on the correctness of the evaluator with respect to its stated semantics, and not on low-level invariants of the evaluator’s internal structures. Since the action of the debugger is independent of the evaluator’s representation of runtime data, the debugger’s proof of correctness is similarly independent. This makes a proof of correctness feasible, and indeed straightforward.

We choose a single language as both source and target of the debugging transformation; that is, the program being debugged may itself use continuation marks. This means, among other things, that multiple “layers” of annotation—a profiler and a debugger, for instance—can coexist.

In the next section, we present a model of a language with continuation marks, and define the annotation and reconstruction functions that make up the stepper. Section 3.4 shows that our stepper satisfies a correctness criterion; in particular, that the steps emitted by the annotated program correspond to those of the original program’s evaluation. Section 3.6 discusses some of the obstacles we overcame in the translation of the model to a concrete implementation. Finally, we consider some of the related work.

3.2 The Continuation Mark Architecture

A traditional debugger interrupts the evaluation of a program to provide information about the state of a program’s evaluation, resumes the computation, stops again, and so on. The generated sequence of evaluation states induces a semantics. A debugger is *adequate* if this induced semantics represents steps occurring in the defined semantics of the language. An adequate debugger is called a *stepper* if there is a bijection between the two sequences of steps.

This comparison is only possible for a language with a semantics that models evaluation as a sequence of steps; that is, a small-step semantics. We choose

the CEKD machine because this model closely matches the implementation strategy of current evaluators and allows a translation to virtual machines such as the JVM [33] and the .NET semantics [24]. In the CEKD machine, an evaluation consists of a sequence of Expression and Value configurations. Each Expression configuration is a 4-tuple of values for machine registers: C , E , K , and D . The first specifies the current expression and is akin to the program counter in other accounts of machine-level computation. The second is an environment, which maps the free variables of the current expression to values. The third is a continuation, or control stack, whose records indicate what is to be done with the computed value. The fourth stores the values associated with top-level variables. Value configurations are represented as 3-tuples— V , K , and D —where the current expression and environment are replaced by a value.

To simulate the steps taken by such a machine, a stepper must be able to reproduce the content of these four machine registers. The stepper can capture most of the machine state directly. More precisely, the current expression (C), the content of the current environment (E), and the values defined for the top-level bindings (D) may all be exposed by a systematic program instrumentation (or *annotation*) with suitable *print*-style expressions or calls to a stepper coroutine.

The only significant obstacle is the reconstruction of the control stack, K . To address this, most debuggers are granted privileged, extra-lingual access to the machine state. This privilege generally involves exposing the format of the compiler’s stack representations and impedes later changes to these protocols. DrScheme avoids this trap through the use of continuation marks. Rather than building a debugger that can see every detail of the compiler’s data structures, we end up with a debugger that can see only the values stored earlier by the annotated program itself.

Also, the continuation mark mechanism does not allow a program to inspect the control stack associated with uninstrumented code, and therefore does not compromise security. That is, the stepper sees foreign program components as “black boxes”, and does not expose details of their operation.

The Model

The language of our model, λ_{cm} , is based on the λ_v -calculus, augmented with the language forms needed to express the stepper for a functional language with definitions. Its syntax appears in figure 3.1. It includes the **w-c-m** and **c-c-m** constructs, and the **output** expression models the output of values on a set of ports. List primitives are included to simplify the presentation of the stepper’s model. Since the (**list** ...) form is syntactic shorthand for a nested series of (**cons** ...) expressions, **list** does not appear in the machine definitions that follow. Similarly, a **let** form is used as an abbreviation for a simple application.

A program in λ_{cm} consists of a nonempty sequence of top-level definitions. Variable references are statically partitioned into top-level and lexical refer-

```

P = (define f C) (define f C) ...
C ∈ Exprs = n | 'x | p | (C C ...) | (lambda (x ...) C) | x | f
           | (let (x C1) C2) ;; syntax for ((lambda (x) C2) C1)
           | (cons C C) | null
           | (list C ...) ;; syntax for (cons C (cons ... null))
           | (car C) | (cdr C)
           | true | false | (if C C C)
           | (w-c-m k C C) | (c-c-m k ...)
           | (output j C)
x ∈ Identifiers(lexical)
f ∈ Identifiers(top-level)
j ∈ Ports
k ∈ Keys
n ∈ Numbers
p ∈ Primops

```

Figure 3.1: Syntax

ences. We use the metavariable f to represent a top-level variable, and the metavariable x to refer to a lexical one. In our typeless setting, forward references to top-level variables present no problems, except that the evaluation of a reference to a variable that has not yet been defined causes a run-time error.

The λ_{cm} language includes Lisp-style symbols, written as 'x. The set of symbols is taken to include (at least) the names of the lexical variables, the names of the top-level variables, the names of the primitive operators, the continuation mark keys, and the names of the output ports.

We present the semantics of the λ_{cm} language using the CEKD register machine. The data definitions for this machine appear in figure 3.2. Configurations in this machine take on one of four forms: the Expression and Value configurations represent a running machine, and the Error and Finished configurations represent a halted one. In the Expression configuration, the machine is searching for an expression to reduce. In the Value configuration, the machine has a value that must be supplied to a control context. The Error configuration arises as the result of a failed run-time check, and the Finished configuration occurs upon successful completion of evaluation.

More precisely, the Expression configuration contains an expression, C , an environment E in which to evaluate that expression, a continuation pair containing a continuation K and a mark table M , and a top-level environment D . The Value configuration contains a value V , a continuation pair $\langle K, M \rangle$, and a top-level environment D . The lexical environment E is discarded, as it

$$\begin{aligned}
S \in \text{Configs} &= \langle C, E, \langle K, M \rangle, D \rangle && \text{(expression configurations)} \\
&| \langle V, \langle K, M \rangle, D \rangle && \text{(value configurations)} \\
&| \langle E \rangle && \text{(finished)} \\
&| \langle \text{error} \rangle \\
V \in \text{Values} &= \langle \text{num} : n \rangle | \langle \text{str} : s \rangle | \langle \text{sym} : x \rangle | \langle \text{prim} : p \rangle | \langle \text{clo} : \langle x, \dots \rangle, C, E \rangle \\
&| \langle \text{pair} : V, V \rangle | \langle \text{null} \rangle | \langle \text{true} \rangle | \langle \text{false} \rangle \\
E \in \text{Environments} &= \text{Identifier} \rightarrow \text{Value} \\
K \in \text{Continuations} &= \langle \text{app} : \langle V, \dots \rangle, \langle C, \dots \rangle, E, \langle K, M \rangle \rangle \\
&| \langle \text{cons1} : C, E, \langle K, M \rangle \rangle | \langle \text{cons2} : V, \langle K, M \rangle \rangle \\
&| \langle \text{car} : \langle K, M \rangle \rangle | \langle \text{cdr} : \langle K, M \rangle \rangle \\
&| \langle \text{if} : C, C, E, \langle K, M \rangle \rangle \\
&| \langle \text{wcm} : k, C, E, \langle K, M \rangle \rangle \\
&| \langle \text{out} : j, \langle K, M \rangle \rangle \\
&| \langle \rangle \\
M \in \text{Marks} &= \text{Identifier} \rightarrow \text{Value} \\
D \in \text{Topenvs} &= \langle E, f, \langle \langle f, C \rangle, \dots \rangle \rangle \\
\delta \in \text{Primop} &\times \langle \text{Value}, \dots \rangle \rightarrow \text{Value}
\end{aligned}$$

Figure 3.2: Machine Definitions

cannot affect the evaluation of the program. the Finished configuration contains the top-level environment that is the result of evaluating all the program's definitions, and the Error configuration contains nothing.

Figures 3.3 and 3.4 show the transition rules for expression configurations and value configurations, respectively. The majority of these are conventional. Subexpressions of applications and pair constructors are evaluated left-to-right. Primitive applications are implemented using a partial function δ mapping primitive names and tuples of values to values. Empty environments and mark tables are written as \emptyset .

Transitions producing values—**true**, for instance—discard the mark table associated with the top continuation frame, replacing it with \emptyset . This does not change the behavior of the machine, because this mark table is guaranteed to be removed (without being observed) during the following transition. Preserving this mark table, then, would needlessly complicate the proof by introducing information that is difficult to capture and does not affect the computation. Eliminating this mark table does not affect the asymptotic memory behavior of an implementation.²

²Since all value configurations therefore contain empty mark tables, the left-hand-sides of the value reductions all show an empty mark table.

$$\begin{aligned}
\langle n, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle \langle \text{num} : n \rangle, \langle K, \emptyset \rangle, D \rangle \\
\langle x, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle \langle \text{sym} : x \rangle, \langle K, \emptyset \rangle, D \rangle \\
\langle p, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle \langle \text{prim} : p \rangle, \langle K, \emptyset \rangle, D \rangle \\
\langle \langle C_1 C_2 \dots \rangle, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle C_1, E|_{fv(C_1)}, \\
&\quad \langle \langle \text{app} : \langle \rangle, \langle C_2, \dots \rangle, E|_{fv(C_2, \dots)}, \langle K, M \rangle \rangle, \emptyset \rangle, \\
&\quad D \rangle \\
\langle \langle \text{lambda } (x \dots) C \rangle, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle \langle \text{clo} : \langle x, \dots \rangle, C, E|_{fv(C) \setminus \{x, \dots\}}, \langle K, \emptyset \rangle, D \rangle \\
\langle x, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle V, \langle K, \emptyset \rangle, D \rangle \text{ where } E(x) = V \\
\langle f, E, \langle K, M \rangle, \langle E', f_1, \langle \langle f_2, C_2 \rangle, \dots \rangle \rangle \rangle &\xrightarrow{\tau} \begin{cases} \langle V, \langle K, \emptyset \rangle, \langle E', f_1, \langle \langle f_2, C_2 \rangle, \dots \rangle \rangle \rangle \\ \text{if } E'(f) = V \\ \langle \text{error} \rangle \text{ otherwise} \end{cases} \\
\langle \langle \text{cons } C_1 C_2 \rangle, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle C_1, E|_{fv(C_1)}, \\
&\quad \langle \langle \text{cons1} : C_2, E|_{fv(C_2)}, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \\
\langle \text{null}, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle \langle \text{null} \rangle, \langle K, \emptyset \rangle, D \rangle \\
\langle \langle \text{car } C \rangle, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle C, E|_{fv(C)}, \langle \langle \text{car} : \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \\
\langle \langle \text{cdr } C \rangle, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle C, E|_{fv(C)}, \langle \langle \text{cdr} : \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \\
\langle \text{true}, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle \langle \text{true} \rangle, \langle K, \emptyset \rangle, D \rangle \\
\langle \text{false}, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle \langle \text{false} \rangle, \langle K, \emptyset \rangle, D \rangle \\
\langle \langle \text{if } C_1 C_2 C_3 \rangle, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle C_1, E|_{fv(C_1)}, \\
&\quad \langle \langle \text{if} : C_2, C_3, E|_{fv(C_2, C_3)}, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \\
\langle \langle \text{w-c-m } k C_1 C_2 \rangle, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle C_1, E, \langle \langle \text{wcm} : k, C_2, E|_{fv(C_2)}, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \\
\langle \langle \text{c-c-m } k \dots \rangle, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle V, \langle K, \emptyset \rangle, D \rangle \text{ where } V = \pi_{\vec{k}}(\langle K, M \rangle) \\
\langle \langle \text{output } j C \rangle, E, \langle K, M \rangle, D \rangle &\xrightarrow{\tau} \langle C, E|_{fv(C)}, \langle \langle \text{out} : j, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle
\end{aligned}$$

where

$$\pi_{\vec{k}}(\langle K_1, M_1 \rangle) = [\phi_{\vec{k}}(M) \mid M \in M_1 \dots M_n \text{ and } \phi_{\vec{k}}(M) \neq \langle \text{null} \rangle]$$

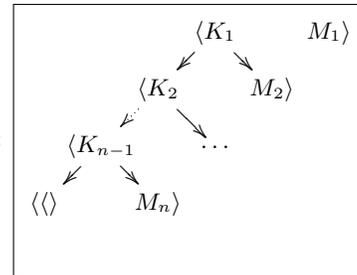
in which $K_1 = \langle \dots, \langle K_2, M_2 \rangle \rangle$

\vdots

and $K_m = \langle \dots, \langle K_{m+1}, M_{m+1} \rangle \rangle$ that is:

\vdots

up to $K_n = \langle \rangle$



and

$$\phi_{\vec{k}}(M) = [(\text{list } 'k V) \mid k \in \vec{k} \text{ and } M(k) = V]$$

Figure 3.3: Expression Reductions

$$\begin{array}{l}
\langle V_n, \langle \langle \text{app} : \langle V_1, \dots, V_{n-1} \rangle, \langle C_1, C_2, \dots \rangle, E, \langle K, M \rangle \rangle, M_0 \rangle, D \rangle \xrightarrow{\tau} \\
\langle C_1, E|_{fv(C_1)}, \langle \langle \text{app} : \langle V_1, \dots, V_n \rangle, \langle C_2, \dots \rangle, E|_{fv(C_2, \dots)}, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \\
\langle V_n, \langle \langle \text{app} : \langle V_1, \dots, V_{n-1} \rangle, \langle \rangle, \\
E, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \xrightarrow{\tau} \begin{cases} \langle C, E''|_{fv(C)}, \langle K, M \rangle, D \rangle \\ \text{if } V_1 = \langle \text{clo} : \langle x_1 \dots x_{n-1} \rangle, C, E' \rangle \\ \text{and } E'' = E'[x_1 \mapsto V_2] \dots [x_{n-1} \mapsto V_n] \\ \langle V, \langle K, \emptyset \rangle, D \rangle \text{ if } V_1 = \langle \text{prim} : p \rangle \\ \text{and } \delta(p, \langle V_2, \dots, V_n \rangle) = V \\ \langle \text{error} \rangle \text{ otherwise} \end{cases} \\
\langle V, \langle \langle \text{cons1} : C, E, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \xrightarrow{\tau} \langle C, E|_{fv(C)}, \langle \langle \text{cons2} : V, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \\
\langle V_2, \langle \langle \text{cons2} : V_1, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \xrightarrow{\tau} \langle \langle \text{pair} : V_1, V_2 \rangle, \langle K, \emptyset \rangle, D \rangle \\
\langle V, \langle \langle \text{car} : \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \xrightarrow{\tau} \begin{cases} \langle V_1, \langle K, \emptyset \rangle, D \rangle \text{ if } V = \langle \text{pair} : V_1, V_2 \rangle \\ \langle \text{error} \rangle \text{ otherwise} \end{cases} \\
\langle V, \langle \langle \text{cdr} : \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \xrightarrow{\tau} \begin{cases} \langle V_2, \langle K, \emptyset \rangle, D \rangle \text{ if } V = \langle \text{pair} : V_1, V_2 \rangle \\ \langle \text{error} \rangle \text{ otherwise} \end{cases} \\
\langle V, \langle \langle \text{if} : C_1, C_2, E, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \xrightarrow{\tau} \begin{cases} \langle C_1, E|_{fv(C_1)}, \langle K, M \rangle, D \rangle \text{ if } V = \langle \text{true} \rangle \\ \langle C_2, E|_{fv(C_2)}, \langle K, M \rangle, D \rangle \text{ if } V = \langle \text{false} \rangle \\ \langle \text{error} \rangle \text{ otherwise} \end{cases} \\
\langle V, \langle \langle \text{wcm} : k, C, E, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \xrightarrow{\tau} \langle C, E|_{fv(C)}, \langle K, M[k \mapsto V] \rangle, D \rangle \\
\langle V, \langle \langle \text{out} : j, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \xrightarrow{\langle j, V \rangle} \langle \langle \text{false} \rangle, \langle K, \emptyset \rangle, D \rangle \\
\langle V, \langle \langle \rangle, \emptyset \rangle, \langle E, f, \langle \langle f_1, C_1 \rangle, \dots \rangle \rangle \rangle \xrightarrow{\tau} \begin{cases} \langle E[f \mapsto V] \rangle \text{ if } \langle \langle f_1, C_1 \rangle, \dots \rangle = \langle \rangle \\ \langle C, \emptyset, \langle \rangle, \langle E[f \mapsto V], f_1, \langle \langle f_2, C_2 \rangle, \dots \rangle \rangle \rangle \\ \text{if } \langle \langle f_1, C_1 \rangle, \dots \rangle = \langle \langle f_1, C_1 \rangle, \langle f_2, C_2 \rangle, \dots \rangle \end{cases}
\end{array}$$

Figure 3.4: Value Reductions

As described before, a **w-c-m** expression contains a statically chosen key k and an expression C_1 whose evaluated result is associated with k in the mark table associated with the continuation. The machine then evaluates the body, C_2 . This second evaluation is in tail position; the machine does not create a new continuation frame in which to evaluate C_2 . For simplicity, the extension of a mark table M is described using the familiar notation of environment extension. In an implementation, overwritten bindings are unreachable and may be removed, preserving the tail-recursive memory behavior of the annotated source.

The **c-c-m** expression provides the means to extract the values associated with one or more keys in the mark tables of a continuation.³ The function $\pi_{\vec{k}}$, specified in figure 3.3, describes the list that is the result of mark extraction. This list is constructed by restricting the mark table associated with each

³The given mechanism accommodates **c-c-m** expressions with zero keys, but the result is always the null list.

continuation frame to the set of keys specified in the **c-c-m** expression, and eliding all empty restrictions. The list comprehensions and explicit uses of **list** in the definitions of $\pi_{\bar{k}}$ and $\phi_{\bar{k}}$ are taken to construct chains of $\langle \text{pair} : V, V \rangle$ values.

We use a labeled transition system [35] with $\langle \text{port}, \text{value} \rangle$ pairs to model output. The port j in this pair designates an output stream for the value. We must use a model with separate ports in order to annotate a source program into a target program that produces output, but does not disrupt the program's existing output. Expressions other than the **output** expression produce no output, indicated in the model with a τ -label. When we write \mapsto with no superscript, it does not mean that there is no output, but rather that the output is not pertinent.

The evaluator is safe-for-space [11]. That is, every reduction maintains the invariant that each environment contains bindings only for variables that occur free in the associated terms.⁴ This invariant is explicitly enforced for all reductions, even those where it is already known to hold. This explicit and uniform enforcement greatly simplifies the proofs later.

Multi-step evaluation \mapsto^* is defined as the transitive, reflexive closure of the relation \mapsto . That is, we say that $S_1 \mapsto^* S_n$ if $n = 1$ or there exist S_1, \dots, S_n such that $S_i \mapsto S_{i+1}$ for $1 \leq i < n$.

Every Value or Expression configuration has exactly one successor configuration. This follows from a simple case analysis: the left-hand sides of the transition definitions are mutually exclusive and collectively exhaustive, following the structure of either the C register (in the case of an Expression configuration) or the K register (in the case of a Value configuration).

The function \mathcal{L} loads a program as an initial configuration:

$$\mathcal{L}[(\mathbf{define} \ f_1 \ C_1) (\mathbf{define} \ f_2 \ C_2) \dots] = \langle C_1, \emptyset, \langle \langle \rangle, \emptyset \rangle, \langle \emptyset, f_1, \langle f_2, C_2 \rangle \dots \rangle \rangle$$

Evaluation of a program is a partial function:

$$\text{Eval}[P] = \begin{cases} \langle E \rangle & \text{if } \mathcal{L}(P) \mapsto \langle E \rangle \\ \langle \text{error} \rangle & \text{if } \mathcal{L}(P) \mapsto \langle \text{error} \rangle \end{cases}$$

Programs may produce output. The Trace function gathers the output at a given port as a sequence of configurations:⁵

$$\begin{aligned} \text{Trace}_j(S) &= \langle \rangle \\ \text{Trace}_j(S_1 \mapsto S_2 \mapsto \dots \mapsto S_n) &= \begin{cases} V :: \text{Trace}_j(S_2 \mapsto \dots \mapsto S_n) & \text{if } S_1 \xrightarrow{\langle j, V \rangle} S_2 \\ \text{Trace}_j(S_2 \mapsto \dots \mapsto S_n) & \text{otherwise} \end{cases} \end{aligned}$$

3.3 Modeling a Stepper

The stepper works by annotating a source program with code fragments that produce debugging outputs. These debugging outputs are mapped back to ma-

⁴The function fv computes the free variables of an expression or expressions.

⁵The infix operator $::$ constructs a list from a list element and a list.

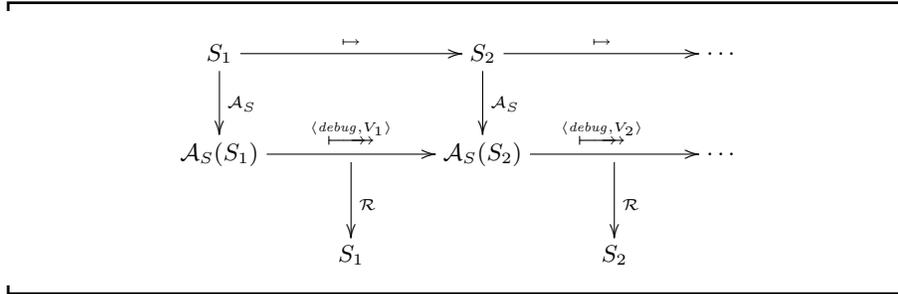


Figure 3.5: The Stepper Architecture

chine configurations. The stepper is correct if the configurations produced by the reconstruction of the debugging output match the steps taken in the evaluation of the original program. No other communication between the stepper and the runtime system is required.

Figure 3.5 illustrates the induction hypothesis that is the basis for our correctness proof: each given configuration maps by \mathcal{A}_S to a corresponding configuration that occurs in the evaluation of the annotated program, and \mathcal{R} maps the debugging outputs to configurations that match those of the original machine.

The next three subsections explain the details of our stepper model. Subsection 3.3 describes the source-to-source transformation that annotates the given program into one that emits state information. Subsection 3.3 outlines the configuration annotation that is the basis for the induction hypothesis. Subsection 3.3 explains the mapping from the values that the annotated program outputs back to the configurations of the original program.

Annotation

The stepper must preserve the behavior defined by the source program. For this reason, the annotation must use a mark key, an output port, and a set of lexical identifiers that do not occur in the source program. Since mark keys, ports, and identifiers are statically evident in the source program, this is not difficult. For simplicity, then, we assume that the source program does not use the mark key *debug*, the output port *debug*, or any of a set of lexical identifiers we shall denote as t_1, t_2, \dots or t_{dc} . We say that the predicate `Clean()` holds for programs, configurations, environments, and mark tables that fulfill these requirements. A conflict may be resolved by α -renaming.

Figure 3.6 summarizes the functions that perform the annotation. We explain each in the following subsections. Their full definitions appear in appendix A.1.

$\text{Annotate} : \text{Program} \rightarrow \text{Program}$	Annotate a program
$\mathcal{A}_{C_{\vec{k}}} : \text{Expr} \rightarrow \text{Expr}$	Annotate an expression
$\mathcal{E}_{\vec{k}} : \text{Expr} \rightarrow \text{Expr}$	Wrap an expression in a mark-destroying w-c-m
$\mathcal{B}_{\vec{k}} : \text{Expr} \times \text{Expr} \rightarrow \text{Expr}$	Wrap a breakpoint around an expression
$\mathcal{Q} : \text{Expr} \rightarrow \text{Expr}$	Quote an expression

Figure 3.6: Annotation Functions

Outermost Annotation: Annotate

The outermost annotation function, `Annotate`, is a simple wrapper for the main annotation function. It adds a top-level breakpoint to allow reconstruction of the final step in a definition’s evaluation.

In order to accurately reconstruct the context, the breakpoint expression’s **w-c-m** must gather all mark values placed by the source program. However, there is no way to precisely compute the possible contexts in which an expression may appear. Our model chooses the simplest conservative approximation to this set by capturing mark values associated with *all* keys used in the source program at each breakpoint. This is the set \vec{k} that parameterizes the functions \mathcal{A}_C , \mathcal{W} , \mathcal{E} , and \mathcal{B} .

Annotation: \mathcal{A}_C and \mathcal{W}

The functions \mathcal{A}_C and \mathcal{W} define annotation on expressions. The function \mathcal{W} takes the form of a structural induction over the list of non-tail subexpressions. For each non-tail subexpression in the argument, \mathcal{W} results in a **w-c-m**, a breakpoint, and a term that evaluates the subexpression.

The function \mathcal{A}_C is a wrapper for \mathcal{W} that maps a language term onto a language-independent representation. That is, it separates an expression into a label, a set of non-tail subexpressions, and a single tail expression. This function encapsulates the knowledge about the source language needed for the annotation, and extending the language with a new form requires adding a clause to the definition of \mathcal{A} .

Expression Breakpoints: \mathcal{E}

The function \mathcal{E} wraps each expression in a place-holding **w-c-m** and a breakpoint around an annotated expression. The trivial **w-c-m** serves to overwrite the mark associated with the prior expression’s evaluation. In the absence of this **w-c-m**, spurious information is retained on tail calls.

$\mathcal{A}_{S_{\vec{k}}} : \text{Configuration} \rightarrow \text{Configuration}$	Annotate a configuration
$\mathcal{A}_{K_{\vec{k}}} : \text{Continuation} \times \text{Topenv} \rightarrow \text{Continuation}$	Annotate a continuation
$\mathcal{A}_{V_{\vec{k}}} : \text{Value} \rightarrow \text{Value}$	Annotate a value
$\mathcal{A}_{E_{\vec{k}}} : \text{Env} \rightarrow \text{Env}$	Annotate an environment
$\mathcal{A}_{D_{\vec{k}}} : \text{Topenv} \rightarrow \text{Topenv}$	Annotate a top-level environment
$\mapsto_s : \text{Expr} \times \text{Env} \times \text{Env} \rightarrow \text{Value}$	Static Evaluation

Figure 3.7: Configuration Annotation Functions

Breakpoint syntax: \mathcal{B}

The breakpoint function \mathcal{B} adds an **output** expression to its argument. The breakpoints of our model emit output rather than halting execution. Nevertheless, we retain the breakpoint name as the most natural one.

The \mathcal{B} function accepts two arguments:

- a ‘to-be-wrapped’ expression, and
- a ‘to-be-output’ expression.

As the names suggest, the resulting expression outputs the result of the second argument before evaluating the first.

The function \mathcal{B} is the only one whose result depends on the set of keys that parameterizes every annotation function. The mark values associated with the named keys are captured at the breakpoint, along with the mark values associated with the *debug* key by the stepper’s own **w-c-ms**.

Quoting: \mathcal{Q}

The \mathcal{Q} function maps source expressions to expressions whose evaluated result encodes those source expressions. The only requirements are that the expressions produced by \mathcal{Q} evaluate to values that can be unambiguously mapped back to the source expressions, and secondarily that the quoted expressions are evaluated in minimal time. In our model, we settle for a function that uses linear time in the size of the expression. This is reduced to a constant time in our actual implementation, as we discuss later.

Configuration Annotation

The statement of the inductive hypothesis is based on a mapping from configurations to configurations, \mathcal{A}_S . This mapping relies upon the mapping from expressions to expressions (\mathcal{A}_C) described earlier, along with annotation functions for continuations (\mathcal{A}_K), values (\mathcal{A}_V) environments (\mathcal{A}_E), and top-level environments (\mathcal{A}_D). Figure 3.7 summarizes these functions. We describe each

of them in the following subsections, and their full definitions appear in appendix A.1.

Configuration Annotation: \mathcal{A}_S

Configuration annotation is defined as the application of the appropriate annotation function to each register of the machine.

Continuation Annotation: \mathcal{A}_K

The function \mathcal{A}_K maps continuations to annotated continuations. It is a homomorphism over chains of continuation frames. Each continuation frame in the original is represented by an *app* frame in the annotated continuation. The structure of this application frame follows from the definition of expression annotation, taking into account the fact that the code fragments introduced by annotation will be in a partially evaluated state.

Like the function \mathcal{A}_C , the function \mathcal{A}_k relies upon a language-independent abstraction, \mathcal{W}_K .

Continuation annotation must also produce mark tables containing values that are the result of evaluating **w-c-m** expressions. Fortunately, the mark expressions produced by annotation are part of a strongly normalizing subset of the language and may therefore be mapped to their evaluated counterparts.⁶

The annotated continuation also depends upon the top-level environment, D , as this environment is encoded in the outermost continuation frame.

Value Annotation: \mathcal{A}_V

Value annotation is defined by the function \mathcal{A}_V . Closures in annotated values contain annotated bodies. Annotated values are otherwise identical to their unannotated counterparts.

Environment Annotation: \mathcal{A}_E

The function \mathcal{A}_E defines the annotation of environments and mark tables. It follows directly from the annotation of values.

Top-level Environment Annotation: \mathcal{A}_D

The function \mathcal{A}_D annotates top-level environments. The annotated definitions depend upon the names of the prior definitions, as well as the current and remaining definitions.

$\text{Reconstruct} : \text{Value} \rightarrow \text{Configuration}$	Reconstruct a configuration
$\mathcal{R}_K : \text{Value} \rightarrow \text{Continuation}$	Reconstruct a continuation
$\mathcal{R}_E : \text{Value} \rightarrow \text{Expr}$	Reconstruct an environment

Figure 3.8: Reconstruction Functions

Reconstruction

The annotated program outputs a sequence of values to the *debug* port. The Reconstruct function maps this sequence of values back to a sequence of configurations. Figure 3.8 summarizes the reconstruction functions. Their full definitions appear in appendix A.1.

Throughout these definitions, we condense the text by using the input forms rather than the output forms. That is, we write **(cons A null)** rather than $\langle \text{pair}: A, \langle \text{null} \rangle \rangle$. This is possible because evaluation defines a context-independent bijection for terms containing only constants and **cons**. In a similar manner, we write \mathcal{Q}^{-1} to map values back to terms. Since \mathcal{Q} 's inverse would naturally map terms to terms, \mathcal{Q}^{-1} in fact denotes the composition of the mapping from values back to terms and the true inverse of \mathcal{Q} . A similar convention is used for the function \mathcal{Q}_D^{-1} . Finally, the function \mathcal{A}_V^{-1} maps annotated values back to their unannotated counterparts. Inspection of \mathcal{A}_V shows that it (and the other annotation functions) are injective, making the inverse mapping a function. Further, note that the set of keys \vec{k} is not needed for the inverse mapping, and therefore that \mathcal{A}_V^{-1} , \mathcal{A}_C^{-1} , \mathcal{B}^{-1} , \mathcal{E}^{-1} , and \mathcal{W}^{-1} are well-defined without it.

At each step, the value output by the annotated program corresponds to a single configuration. This value is structured as a list of lists. Figure 3.9 shows a template for this value.

Rebuilding a configuration from such values is a straightforward recursive process. The `Reconstruct()` function produces either a value or an expression configuration using the information in the first element of the output value. The \mathcal{R}_k function reconstructs the configuration's continuation by recursive descent on the remainder of the output value.

3.4 Correctness

This section contains two results. First, the two non-interference theorems show that the annotated program produces the same results as the source program. Second, the correctness theorem demonstrates that the stepper is correct.

⁶Since configuration annotation is used only in the proof of correctness, we care only that the function is well-defined, and not that it be easily or efficiently implementable.

```

    ccm-val = (list outer-info frame-info ...)
    outer-info = (list 'expstep Q(C) E) | (list 'valstep V)
    frame-info = (list (list 'debug kont-val) (list V V) ...)
    kont-val = (list V                ; tag
                (list V ...)         ; temp vals
                (list (list Q(C) V) ...) ; binding pairs
                (list Q(C) ...)      ; non-tail exps
                (list Q(C) ...)      ; tail exps
  
```

Figure 3.9: Debug Value Template

Theorem 1 (Result Non-interference) *For any program P containing keys \vec{k} such that $Clean(P)$, $Eval(P) = S$ if and only if $Eval(Annotate(P)) = \mathcal{A}_{S\vec{k}}(S)$*

Theorem 2 (Output Non-interference) *For any program P with keys \vec{k} and for any output port o other than debug, if $\mathcal{L}(P) \mapsto S_n$, then*

$$\mathcal{A}_{V\vec{k}}(Trace_o(\mathcal{L}(P) \mapsto S_n)) = Trace_o(\mathcal{L}(Annotate(P)) \mapsto \mathcal{A}_{S\vec{k}}(S_n))$$

Theorem 3 (Stepper Correctness) *Given a program P containing keys \vec{k} such that $Clean(P)$ (and let $S_1 = \mathcal{L}(P)$), and an evaluation sequence $S_1 \mapsto S_n$ where $n > 1$, then*

$$\begin{aligned} & \text{Reconstruct}(Trace_{debug}(\mathcal{L}(Annotate(S_1)) \mapsto S_n)) \\ &= \begin{cases} \langle S_1, \dots, S_n \rangle & \text{when } S_n \text{ is an Expression configuration, or} \\ \langle S_1, \dots, S_{n-1} \rangle & \text{otherwise.} \end{cases} \end{aligned}$$

All three of these depend on a simulation lemma, a loading lemma, and a cleanliness preservation lemma. The following three subsections prove these lemmas.

Simulation Lemma

In order to prove non-interference and correctness, we must first establish a link between the evaluation of a source program and the evaluation of its annotated counterpart. As the diagram in figure 3.5 shows, the configurations in the evaluation of the source program are related by the annotation function to (a subsequence of) the configurations of the annotated program. The following lemma describes the relationship between the two evaluations.

Lemma 1 (Multi-step Simulation) *Given a configuration S_1 where $Clean(S_1)$ and $S_1 \mapsto S_2$, and any set of keys \vec{k} ,*

$$\mathcal{A}_{S\vec{k}}(S_1) \mapsto \mathcal{A}_{S\vec{k}}(S_2)$$

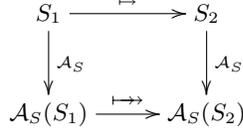


Figure 3.10: The Simulation Lemma

In other words: for every step the original machine takes, the annotated machine takes several steps. No constraint is placed on the set \vec{k} ; the lemma holds for any such set. This argument to the annotation is constrained only when we prove that the reconstructed terms match the steps taken by the original machine.

This lemma is expressed graphically in figure 3.10. This lemma forms the basis for the proof illustrated in figure 3.5.

Several auxiliary lemmas are needed for the proof of lemma 1:

Lemma 2 (Quoted expressions converge) • For any expression C , environment E , continuation pair $\langle K, M \rangle$, and definitions D , there exists V such that $\langle \mathcal{Q}(C), E, \langle K, M \rangle, D \rangle \mapsto \langle V, \langle K, M \rangle, D \rangle$. That is, a quoted expression always converges to a value.

- For any expression C , $\mathcal{Q}(C)$ has no free variables.

Both halves of lemma 2 follow from the definition of \mathcal{Q} ; the resulting expressions contain only **cons**, **null**, and constants.

Lemma 3 (No extra free variables in annotation) For any expression C and key set \vec{k} , $fv(C) = fv(\mathcal{A}_{C\vec{k}}(C))$. That is, annotation does not change the set of free variables.

The proof of lemma 3 follows by structural induction on source terms, and (within this) on left-to-right induction on the non-tail subterms of individual expressions, following the structure of the function \mathcal{A}_C .

In the following sections, we prove the principal lemma, number 1, by considering each class of configuration as the left-hand side (S_1) of the transition.

Finished and Error Configurations

The Finished and Error configurations have no next step. If S_1 is either of these, then it cannot be that $S_1 \mapsto S_2$. The lemma holds trivially for these two cases.

Expression Configurations

The transitions on Expression configurations are determined by the configuration's expression, C . We proceed by case analysis on this expression, with arbitrary E , K , M , and D .

Numbers, n : By the definition of the abstract machine,

$$S_1 = \langle n, E, \langle K, M \rangle, D \rangle \mapsto \langle \langle \text{num} : n \rangle, \langle K, \emptyset \rangle, D \rangle = S_2$$

By the definition of configuration annotation, furthermore,

$$\mathcal{A}_{S_{\bar{k}}}(S_1) = \langle n, \mathcal{A}_{E_{\bar{k}}}(E), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$$

and

$$\mathcal{A}_{S_{\bar{k}}}(S_2) = \langle \langle \text{num} : n \rangle, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$$

So we must show that

$$\begin{aligned} & \langle n, \mathcal{A}_{E_{\bar{k}}}(E), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \\ & \mapsto \langle \langle \text{num} : n \rangle, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \end{aligned}$$

This transition occurs in one step, so the lemma is proved for this case.

Other Constants: Other constants (symbols, primitives, **null**, **true**, and **false**) are precisely analogous. To prove the lemma for these cases, we substitute the corresponding expression and value forms for n and $\langle \text{num} : n \rangle$ in the case above.

Variable references: Variable references are similarly simple. In the case of a successful variable reference (either lexical or top-level), we may substitute the variable reference (x or f) for n and the appropriate environment application for $\langle \text{num} : n \rangle$.

In the case of a top-level reference to an unevaluated binding, we prove the lemma by substituting the variable reference for n , and $\langle \text{Error} \rangle$ for both S_2 and $\mathcal{A}_{S_{\bar{k}}}(S_2)$.

The Lambda Expression: By the definition of the abstract machine,

$$\begin{aligned} S_1 &= \langle \langle \text{lambda } (x \dots) C \rangle, E, \langle K, M \rangle, D \rangle \\ &\mapsto \langle \langle \text{clo} : \langle x, \dots \rangle, C, E|_{fv(C) \setminus \{x, \dots\}} \rangle, \langle K, \emptyset \rangle, D \rangle = S_2 \end{aligned}$$

By the definition of configuration annotation, furthermore,

$$\begin{aligned} \mathcal{A}_{S_{\bar{k}}}(S_1) &= \langle \langle \text{lambda } (x \dots) \mathcal{E}_{\bar{k}}(C) \rangle, \mathcal{A}_{E_{\bar{k}}}(E), \\ & \quad \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \end{aligned}$$

and

$$\mathcal{A}_{S_{\bar{k}}}(S_2) = \langle \langle \text{clo} : \langle x, \dots \rangle, \mathcal{E}_{\bar{k}}(C), \mathcal{A}_{E_{\bar{k}}}(E|_{fv(C) \setminus \{x, \dots\}}) \rangle, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \rangle$$

So we must show that

$$\begin{aligned} & \langle (\mathbf{lambda} (x \dots) \mathcal{E}_{\bar{k}}(C)), \mathcal{A}_{E_{\bar{k}}}(E), \\ & \quad \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \rangle \\ \mapsto & \langle \langle \text{clo} : \langle x, \dots \rangle, \mathcal{E}_{\bar{k}}(C), \mathcal{A}_{E_{\bar{k}}}(E|_{fv(C) \setminus \{x, \dots\}}) \rangle, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \rangle \end{aligned}$$

By the definition of the abstract machine,

$$\begin{aligned} \mathcal{A}_{S_{\bar{k}}}(S_1) \mapsto & \langle \langle \text{clo} : \langle x, \dots \rangle, \mathcal{E}_{\bar{k}}(C), (\mathcal{A}_{E_{\bar{k}}}(E)|_{fv(\mathcal{E}_{\bar{k}}(C)) \setminus \{x, \dots\}}) \rangle, \\ & \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M) \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \rangle \end{aligned}$$

But by the definition of \mathcal{A}_E and lemma 3,

$$\mathcal{A}_{E_{\bar{k}}}(E)|_{fv(\mathcal{E}_{\bar{k}}(C)) \setminus \{x, \dots\}} = \mathcal{A}_{E_{\bar{k}}}(E|_{fv(\mathcal{E}_{\bar{k}}(C)) \setminus \{x, \dots\}}) = \mathcal{A}_{E_{\bar{k}}}(E|_{fv(C) \setminus \{x, \dots\}})$$

So this transition occurs in one step.

The C-c-m Expression: By the definition of the abstract machine,

$$S_1 = \langle (\mathbf{c-c-m} k' \dots), E, \langle K, M \rangle, D \rangle \mapsto \langle \pi_{\bar{k}'}(\langle K, M \rangle), \langle K, \emptyset \rangle, D \rangle = S_2$$

By the definition of configuration annotation, furthermore,

$$\begin{aligned} \mathcal{A}_{S_{\bar{k}}}(S_1) = & \\ & \langle (\mathbf{c-c-m} k' \dots), \mathcal{A}_{E_{\bar{k}}}(E), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \rangle \end{aligned}$$

and

$$\mathcal{A}_{S_{\bar{k}}}(S_2) = \langle \mathcal{A}_{V_{\bar{k}}}(\pi_{\bar{k}'}(\langle K, M \rangle)), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$$

So we must show that

$$\begin{aligned} & \langle (\mathbf{c-c-m} k' \dots), \mathcal{A}_{E_{\bar{k}}}(E), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \rangle \\ \mapsto & \langle \mathcal{A}_{V_{\bar{k}}}(\pi_{\bar{k}'}(\langle K, M \rangle)), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \rangle \end{aligned}$$

By the definition of the abstract machine,

$$\mathcal{A}_{S_{\bar{k}}}(S_1) \mapsto \langle \pi_{\bar{k}'}(\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M) \rangle), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \rangle$$

We can now prove that $\mathcal{A}_{V_{\bar{k}}}(\pi_{\bar{k}'}(\langle K, M \rangle)) = \pi_{\bar{k}'}(\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M) \rangle)$.

By the definition of $\pi_{\bar{k}'}$,

$$\mathcal{A}_{V_{\bar{k}}}(\pi_{\bar{k}'}(\langle K, M \rangle)) = \mathcal{A}_{V_{\bar{k}}}([\phi_{\bar{k}'}(M) | M = M_1 \dots M_n \text{ and } \phi_{\bar{k}'}(M) \neq \langle \text{null} \rangle])$$

for the sequence of marks $M_1 \dots M_n$ contained in the continuation.

Turning to the other half of the equation, we see by the definition of \mathcal{A}_K that the sequence of marks contained in $\langle \mathcal{A}_{K\vec{k}}(K, D), \mathcal{A}_{E\vec{k}}(M) \rangle$ is

$$\begin{aligned} M'_1 &= \mathcal{A}_{E\vec{k}}(M_1) \\ M'_2 &= \mathcal{A}_{E\vec{k}}(M_2)[debug \mapsto V_1] \\ &\dots \\ M'_n &= \mathcal{A}_{E\vec{k}}(M_n)[debug \mapsto V_{n-1}] \\ M'_{n+1} &= \emptyset[debug \mapsto V_n] \end{aligned}$$

for some set of values $V_1 \dots V_n$.

We know that $\text{Clean}(S_1)$ and therefore that $debug \notin k'$, so for all M and V , $\phi_{\vec{k}'}(M[debug \mapsto V]) = \phi_{\vec{k}'}(M)$.

$$\begin{aligned} \text{So } \pi_{\vec{k}'}(\langle \mathcal{A}_{K\vec{k}}(K, D), \mathcal{A}_{E\vec{k}}(M) \rangle) & \\ &= [\phi_{\vec{k}'}(M') | M' = \langle M'_1, \dots, M'_{n+1} \rangle \text{ and } \phi_{\vec{k}'}(M') \neq \langle \text{null} \rangle] \\ &\quad \text{(by the definition of } \pi_{\vec{k}'} \text{)} \\ &= [\phi_{\vec{k}'}(M') | M' = \langle \mathcal{A}_{E\vec{k}}(M_1), \dots, \mathcal{A}_{E\vec{k}}(M_n), \emptyset \rangle \text{ and } \phi_{\vec{k}'}(M') \neq \langle \text{null} \rangle] \\ &\quad \text{(since } \phi_{\vec{k}'}(M[debug \mapsto V]) = \phi_{\vec{k}'}(M) \text{)} \\ &= [\phi_{\vec{k}'}(M') | M' = \langle \mathcal{A}_{E\vec{k}}(M_1), \dots, \mathcal{A}_{E\vec{k}}(M_n) \rangle \text{ and } \phi_{\vec{k}'}(M') \neq \langle \text{null} \rangle] \\ &\quad \text{(because } \phi_{\vec{k}'}(\emptyset) = \langle \text{null} \rangle \text{)} \\ &= [\phi_{\vec{k}'}(\mathcal{A}_{E\vec{k}}(M)) | M = \langle M_1, \dots, M_n \rangle \text{ and } \phi_{\vec{k}'}(M) \neq \langle \text{null} \rangle] \\ &\quad \text{(because } \phi_{\vec{k}'}(M) = \langle \text{null} \rangle \text{ iff } \phi_{\vec{k}'}(\mathcal{A}_{E\vec{k}}(M)) = \langle \text{null} \rangle \\ &\quad \text{for all } k, k', \text{ and } M \text{)} \\ &= [\mathcal{A}_{V\vec{k}}(\phi_{\vec{k}'}(M)) | M = \langle M_1, \dots, M_n \rangle \text{ and } \phi_{\vec{k}'}(M) \neq \langle \text{null} \rangle] \\ &\quad \text{(because } \phi_{\vec{k}'}(\mathcal{A}_{E\vec{k}}(M)) = \mathcal{A}_{V\vec{k}}(\phi_{\vec{k}'}(M)) \text{ for all } M, k, \text{ and } k' \text{)} \\ &= \mathcal{A}_{V\vec{k}}[\phi_{\vec{k}'}(M) | M = \langle M_1, \dots, M_n \rangle \text{ and } \phi_{\vec{k}'}(M) \neq \langle \text{null} \rangle] \\ &\quad \text{(by the definition of } \mathcal{A}_{V\vec{k}} \text{)} \end{aligned}$$

So this transition occurs in one step.

All Other Expressions The remaining expressions all have non-tail expressions. The proof cases for each of these differ only in minor ways. We therefore present only the proof for a representative class of expressions, the **if** expressions.

By the definition of the abstract machine,

$$S_1 = \langle \langle \text{if } C_1 \ C_2 \ C_3 \rangle, E, \langle K, M \rangle, D \rangle \\ \mapsto \langle C_1, E|_{fv(C_1)}, \langle \langle \text{if } : C_2, C_3, E|_{fv(C_2, C_3)}, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle = S_2$$

By the definition of configuration annotation, furthermore,

$$\mathcal{A}_{S_{\vec{k}}}(S_1) = \langle C', \mathcal{A}_{E_{\vec{k}}}(E), \langle \mathcal{A}_{K_{\vec{k}}}(K, D), \mathcal{A}_{E_{\vec{k}}}(M)[debug \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\vec{k}}}(D) \rangle$$

where

$$C' = (\mathbf{w-c-m} \ \text{debug} \\ (\mathbf{list} \ \text{'if} \ \mathbf{null} \ (\mathbf{list} \ (\mathbf{list} \ \text{'x} \ x) \ \dots) \ \mathbf{;;} \ \text{for } x \in fv(C_2, C_3) \\ \mathbf{null} \ (\mathbf{list} \ \mathcal{Q}(C_2) \ \mathcal{Q}(C_3)) \ \mathbf{null}) \\ (\mathbf{let} \ ([t_1 \ \mathcal{E}_{\vec{k}}(C_1)]) \\ (\mathbf{begin} \\ (\mathbf{output} \ \text{debug} \ (\mathbf{list} \ (\mathbf{list} \ \text{'valstep} \ t_1 \\ (\mathbf{c-c-m} \ \text{debug} \ k \ \dots))) \ \mathbf{;;} \ \text{for } k \in \vec{k} \\ (\mathbf{if} \ t_1 \ \mathcal{E}_{\vec{k}}(C_2) \ \mathcal{E}_{\vec{k}}(C_3))))))$$

and

$$\mathcal{A}_{S_{\vec{k}}}(S_2) = \langle \mathcal{A}_{C_{\vec{k}}}(C_1), \mathcal{A}_{E_{\vec{k}}}(E)|_{fv(C_1)}, \\ \langle \langle \mathbf{app} : \langle \langle \mathbf{clo} : \langle t_1 \rangle, C', \mathcal{A}_{E_{\vec{k}}}(E)|_{fv(C_2, C_3)} \rangle \rangle, \langle \rangle, \emptyset, \\ \langle \mathcal{A}_{K_{\vec{k}}}(K, D), \mathcal{A}_{E_{\vec{k}}}(M)[debug \mapsto V'] \rangle \rangle, \emptyset[debug \mapsto \langle \text{false} \rangle] \rangle, \\ \mathcal{A}_{D_{\vec{k}}}(D) \rangle$$

where

$$C' = (\mathbf{begin} \ (\mathbf{output} \ \text{debug} \\ (\mathbf{list} \ (\mathbf{list} \ \text{'valstep} \ t_1) \ (\mathbf{c-c-m} \ \text{debug} \ k \ \dots))) \\ (\mathbf{if} \ t_1 \ \mathcal{E}_{\vec{k}}(C_2) \ \mathcal{E}_{\vec{k}}(C_3)))$$

and V' is the result of evaluating the mark term.

We must show that $\mathcal{A}_{S_{\vec{k}}}(S_1) \mapsto \mathcal{A}_{S_{\vec{k}}}(S_2)$. Figure 3.11 shows the key steps in this evaluation in abbreviated form (where $_C$ refers to an omitted expression, and so forth). The numeric indices refer to the individual steps taken by the machine. The vast majority of the elided steps concern the evaluation of the mark expressions, and the precise number of such steps depends on the choice of C_1 , C_2 , C_3 , and D . The time taken by these steps is reduced to near-constant time in our implementation, where mark values are represented as **lambda** terms that go in one step to closures. The final step is equal to $\mathcal{A}_{S_{\vec{k}}}$, and the lemma holds for this case (and, by extension, for all other expression configurations).

Value Configurations

Value reductions make up the final group of configurations in the proof of lemma 1. Value reductions fall into four sub-categories, based on the topmost

```

0  <<(w-c-m debug (list ...) (let ((t1 -C)) -C)),
    AEk̄(E), <AKk̄(K, D), AEk̄(M)[debug ↦ -V]>, ADk̄(D)>
35 <<(let ((t1 Ek̄(C1))) (begin -C -C)),
    -E, <AKk̄(K, D), AEk̄(M)[debug ↦ -V]>, ADk̄(D)>
38 <Ek̄(C1), -E, <<app : -K>, ∅>, ADk̄(D)>
41 <<(begin (output debug -C) ACk̄(C1)),
    -E, <<app : -K>, ∅[debug ↦ -V]>, ADk̄(D)>
44 <<(output debug (list -C -C)), -E, <<app : -K>, ∅>, ADk̄(D)>
67 <<false>, <<app : -K>, ∅>, ADk̄(D)>
68 <ACk̄(C1), -E|fv(C1), <<app : -K>, ∅[debug ↦ -V]>, ADk̄(D)>

```

Figure 3.11: Steps taken in reduction of if to test position

frame of the current continuation. In the first case, the topmost frame has additional non-tail expressions to evaluate. In the second, the configuration's value is the last non-tail expression but the frame has a tail subexpression. In the third, the frame produces a value directly. In the fourth, the continuation stack is empty.

Non-tail Expressions Remaining The first case includes application with remaining non-tail subexpressions and also the *cons1* continuation frame. We present the proof for the *cons1* continuation frame.

By definition of the abstract machine,

$$\begin{aligned}
S_1 &= \langle V, \langle \langle \text{cons1} : C, E, \langle K, M \rangle \rangle, M_0 \rangle, D \rangle \\
&\mapsto \langle C, E|_{fv(C)}, \langle \langle \text{cons2} : V, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle = S_2
\end{aligned}$$

By the definition of configuration annotation, furthermore,

$$\begin{aligned}
\mathcal{A}_{S_{k̄}}(S_1) &= \langle \mathcal{A}_{V_{k̄}}(V), \\
&\quad \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_1 \rangle, C', \mathcal{A}_{E_{k̄}}(E) \rangle \rangle, \langle \rangle, \emptyset, \\
&\quad \langle \mathcal{A}_{K_{k̄}}(K, D), \mathcal{A}_{E_{k̄}}(M)[\text{debug} \mapsto V'] \rangle \rangle, \emptyset \rangle, \mathcal{A}_{D_{k̄}}(D) \rangle
\end{aligned}$$

where

$$\begin{aligned}
C' &= (\text{begin} (\text{output debug} (\text{list} (\text{list 'valstep } t_1) (\text{c-c-m debug } k \dots))) \\
&\quad (\text{w-c-m debug} \\
&\quad \quad (\text{list 'cons} (\text{list } t_1) \text{ null null null null}) \\
&\quad \quad (\text{let} ((t_2 E_{k̄}(C))) \\
&\quad \quad \quad (\text{begin} (\text{output debug} (\text{list} (\text{list 'valstep } t_2) \\
&\quad \quad \quad \quad (\text{c-c-m debug } k \dots)))) \\
&\quad \quad \quad (\text{cons } t_1 t_2))))))
\end{aligned}$$

and V' is the result of evaluating the mark term.

Also by definition of configuration annotation,

```

0  ⟨ $\mathcal{A}_{V_{\bar{k}}}(V), \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
1  ⟨⟨begin (output debug  $\_C$ ) (w-c-m debug  $\_C$   $\_C$ )⟩,
    $\_E, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
4  ⟨⟨output debug (list  $\_C$   $\_C$ )⟩,  $\_E, \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
23 ⟨⟨false⟩, ⟨⟨app :  $\_K$ ⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
24 ⟨⟨w-c-m debug (list  $\dots$ ) (let ( $t_2$   $\_C$ )  $\_C$ )⟩,
    $\_E, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
55 ⟨⟨let ( $t_2$   $\mathcal{E}_{\bar{k}}(C)$ ) (begin  $\_C$   $\_C$ )⟩,
    $\_E, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
58 ⟨ $\mathcal{E}_{\bar{k}}(C), \_E, \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
61 ⟨⟨begin (output debug  $\_C$ )  $\mathcal{A}_{C_{\bar{k}}}(C)$ ⟩,
    $\_E, \langle \langle \text{app} : \_K \rangle, \emptyset[\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
64 ⟨⟨output debug (list  $\_C$   $\_C$ )⟩,  $\_E, \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
87 ⟨⟨false⟩, ⟨⟨app :  $\_K$ ⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
88 ⟨ $\mathcal{A}_{C_{\bar{k}}}(C), \_E|_{fv(C)}, \langle \langle \text{app} : \_K \rangle, \emptyset[\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 

```

Figure 3.12: Steps taken in reduction of cons1 continuation

$$\mathcal{A}_{S_{\bar{k}}}(S_2) = \langle \mathcal{A}_{C_{\bar{k}}}(C), \mathcal{A}_{E_{\bar{k}}}(E), \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_2 \rangle, C' \rangle, \emptyset[t_1 \mapsto \mathcal{A}_{V_{\bar{k}}}(V)] \rangle \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V'] \rangle, \emptyset[\text{debug} \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$$

where

$$C' = (\text{begin} (\text{output} \textit{debug} (\text{list} (\text{list} \textit{valstep} t_2) (\text{c-c-m} \textit{debug} k \dots))) (\text{cons} t_1 t_2))$$

and

$$V' = \langle \text{pair} : \langle \text{sym} : \text{cons} \rangle, \langle \text{pair} : \langle \text{pair} : \mathcal{A}_{V_{\bar{k}}}(V), \langle \text{null} \rangle \rangle, \langle \text{pair} : \langle \text{null} \rangle, \langle \text{null} \rangle \rangle \rangle \rangle \rangle \rangle$$

Figure 3.12 shows the key steps in the reduction of the *cons1* continuation frame. The final step is equal to $\mathcal{A}_{S_{\bar{k}}}(S_2)$, and the lemma holds for this case.

Tail Expression Remaining The second class of Value configurations are those whose continuation has no more non-tail expressions but results in a new tail subexpression. This case includes the *if* and *wcm* continuations, as well as the application continuation with no more non-tail subexpressions, and a

```

0  ⟨ $\mathcal{A}_{V_{\bar{k}}}(V), \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
1  ⟨⟨begin (output debug  $\_C$ ) (w-c-m  $k \_C \_C$ )⟩⟩,
     $\_E, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
4  ⟨⟨output debug (list  $\_C \_C$ )⟩⟩,  $\_E, \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
23 ⟨⟨false⟩⟩, ⟨⟨app :  $\_K$ ⟩⟩,  $\emptyset, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
24 ⟨⟨w-c-m  $k t_1 \mathcal{E}_{\bar{k}}(C)$ ⟩⟩,  $\_E, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
25 ⟨ $t_1, \_E, \langle \langle \text{wcm} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
26 ⟨ $\mathcal{A}_{V_{\bar{k}}}(V), \langle \langle \text{wcm} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
27 ⟨ $\mathcal{E}_{\bar{k}}(C), \_E, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V][k \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
30 ⟨⟨begin (output debug  $\_C$ )  $\mathcal{A}_{C_{\bar{k}}}(C)$ ⟩⟩,
     $\_E, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[k \mapsto \_V][\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
33 ⟨⟨output debug (list  $\_C \_C$ )⟩⟩,  $\_E, \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
56 ⟨⟨false⟩⟩, ⟨⟨app :  $\_K$ ⟩⟩,  $\emptyset, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
57 ⟨ $\mathcal{A}_{C_{\bar{k}}}(C), \_E|_{fv(C)}, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[k \mapsto \_V][\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 

```

Figure 3.13: Steps taken in reduction of wcm continuation

non-primitive first argument. We take the *wcm* continuation as representative of this group, and then consider the error that may result from an *if*.

By definition of the abstract machine,

$$S_1 = \langle V, \langle \langle \text{wcm} : k, C, E, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \mapsto \langle C, E, \langle K, M[k \mapsto V] \rangle, D \rangle = S_2$$

By the definition of configuration annotation, furthermore,

$$\mathcal{A}_{S_{\bar{k}}}(S_1) = \langle \mathcal{A}_{V_{\bar{k}}}(V), \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_1 \rangle, C', \mathcal{A}_{E_{\bar{k}}}(E) \rangle \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V'] \rangle \rangle, \mathcal{A}_{E_{\bar{k}}}(M_0) \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$$

where

$$C' = (\mathbf{begin} (\mathbf{output debug} (\mathbf{list} (\mathbf{list} \text{'valstep } t_1) (\mathbf{c-c-m debug } k \dots))) (\mathbf{w-c-m } k t_1 \mathcal{E}_{\bar{k}}(C)))$$

and V' is the result of evaluating the mark term.

Also by definition of configuration annotation,

$$\mathcal{A}_{S_{\bar{k}}}(S_2) = \langle \mathcal{A}_{C_{\bar{k}}}(C), \mathcal{A}_{E_{\bar{k}}}(E), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[k \mapsto \mathcal{A}_{V_{\bar{k}}}(V)][\text{debug} \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$$

Figure 3.13 shows the key steps in the reduction of the *wcm* continuation frame. The final step is equal to $\mathcal{A}_{S_{\bar{k}}}(S_2)$, and the lemma holds for this case.

The reduction of an *if* continuation may also result in an error, when the configuration's value is neither true nor false. We take as a representative non-boolean value the number 0.

By definition of the abstract machine,

$$S_1 = \langle \langle \text{num} : 0 \rangle, \langle \langle \text{if} : C_2, C_3, E, \langle K, M \rangle, \emptyset \rangle, D \rangle \mapsto \langle \text{error} \rangle = S_2$$

By the definition of configuration annotation, furthermore,

$$\begin{aligned} \mathcal{A}_{S_{\bar{k}}}(S_1) = \langle \langle \text{num} : 0 \rangle, \\ \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_1 \rangle, C', \mathcal{A}_{E_{\bar{k}}}(E) \rangle \rangle, \langle \rangle, \emptyset, \\ \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V'] \rangle \rangle, \mathcal{A}_{E_{\bar{k}}}(M_0) \rangle, \\ \mathcal{A}_{D_{\bar{k}}}(D) \rangle \end{aligned}$$

where

$$C' = (\mathbf{begin} (\mathbf{output} \text{ debug } (\mathbf{list} (\mathbf{list} \text{ 'valstep } t_1) (\mathbf{c-c-m} \text{ debug } k \dots))) \\ (\mathbf{if} \ t_1 \ \mathcal{E}_{\bar{k}}(C_2) \ \mathcal{E}_{\bar{k}}(C_3)))$$

and V' is the result of evaluating the mark term.

Also by the definition of configuration annotation,

$$\mathcal{A}_{S_{\bar{k}}}(S_2) = \langle \text{error} \rangle$$

Figure 3.14 shows the key steps in the reduction of the *if* continuation frame to an error. The final step is equal to $\mathcal{A}_{S_{\bar{k}}}(S_2)$, and the lemma holds for this case.

New Value Reductions The third class includes all those whose continuation frames have no remaining non-tail subexpressions. This corresponds to the *cons2*, *car*, *cdr*, and *out* continuation frames, and also to the application frame when a primitive is invoked. We take the *cons2* continuation as a representative of this class, and then consider the error that may result from the evaluation of a *car* continuation.

By definition of the abstract machine,

$$S_1 = \langle V_2, \langle \langle \text{cons2} : V_1, \langle K, M \rangle, \emptyset \rangle, D \rangle \mapsto \langle \langle \text{pair} : V_1, V_2 \rangle, \langle K, \emptyset \rangle, D \rangle = S_2$$

By definition of configuration annotation, furthermore,

$$\begin{aligned} \mathcal{A}_{S_{\bar{v}}}(S_1) = \langle \mathcal{A}_{V_{\bar{k}}}(V_2), \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_2 \rangle, C', \emptyset[t_1 \mapsto \mathcal{A}_{V_{\bar{k}}}(V_1)] \rangle \rangle, \langle \rangle, \emptyset, \\ \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V'] \rangle \rangle, \mathcal{A}_{E_{\bar{k}}}(M_0) \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \end{aligned}$$

where

$$C' = (\mathbf{begin} (\mathbf{output} \text{ debug } (\mathbf{list} (\mathbf{list} \text{ 'valstep } t_2) (\mathbf{c-c-m} \text{ debug } k \dots))) \\ (\mathbf{cons} \ t_1 \ t_2))$$

```

0  ⟨⟨num : 0⟩, ⟨⟨app :  $_{-K}$ ⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
1  ⟨⟨begin (output debug  $_{-C}$ ) (if  $_{-C}$   $_{-C}$   $_{-C}$ )⟩,
    $_{-E}$ ,  $\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \mathit{_{-V}}] \rangle, \mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
4  ⟨⟨output debug (list  $_{-C}$   $_{-C}$ )⟩,  $_{-E}$ ,  $\langle \langle \text{app} : \mathit{_{-K}} \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
23 ⟨⟨false⟩,  $\langle \langle \text{app} : \mathit{_{-K}} \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
24 ⟨⟨if  $t_1$   $\mathcal{E}_{\bar{k}}(C_2)$   $\mathcal{E}_{\bar{k}}(C_3)$ ⟩,  $_{-E}$ ,  $\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \mathit{_{-V}}] \rangle, \mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
25 ⟨ $t_1$ ,  $_{-E}$ ,
   ⟨⟨if :  $\mathcal{E}_{\bar{k}}(C_2)$ ,  $\mathcal{E}_{\bar{k}}(C_3)$ ,  $_{-E}$ ,
     ⟨ $\mathcal{A}_{K_{\bar{k}}}(K, D)$ ,  $\mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \mathit{_{-V}}]$ ⟩⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
26 ⟨⟨num : 0⟩,
   ⟨⟨if :  $\mathcal{E}_{\bar{k}}(C_2)$ ,  $\mathcal{E}_{\bar{k}}(C_3)$ ,  $_{-E}$ ,
     ⟨ $\mathcal{A}_{K_{\bar{k}}}(K, D)$ ,  $\mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \mathit{_{-V}}]$ ⟩⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
27 ⟨error⟩

```

Figure 3.14: Steps taken in reduction of if continuation leading to an error

and

$$V' = \langle \text{pair} : \langle \text{sym} : \text{cons} \rangle, \langle \text{pair} : \langle \text{pair} : \mathcal{A}_{V_{\bar{k}}}(V_1), \langle \text{null} \rangle \rangle, \langle \text{pair} : \langle \text{null} \rangle, \langle \text{null} \rangle \rangle \rangle \rangle \rangle \rangle$$

Also by definition of configuration annotation,

$$\mathcal{A}_{S_{\bar{v}}}(S_2) = \langle \langle \text{pair} : \mathcal{A}_{V_{\bar{k}}}(V_1), \mathcal{A}_{V_{\bar{k}}}(V_2) \rangle, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$$

Figure 3.15 shows the key steps in the reduction of the *cons2* continuation frame. The final step is equal to $\mathcal{A}_{S_{\bar{k}}}(S_2)$, and the lemma holds for this case.

The reduction of an *car* continuation may also result in an error, when the configuration's value is not a pair. We take as a representative illegal value the number 0.

By definition of the abstract machine,

$$S_1 = \langle \langle \text{num} : 0 \rangle, \langle \langle \text{car} : \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \mapsto \langle \text{error} \rangle = S_2$$

By the definition of configuration annotation, furthermore,

$$\mathcal{A}_{S_{\bar{k}}}(S_1) = \langle \langle \text{num} : 0 \rangle, \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_1 \rangle, C' \rangle, \emptyset \rangle, \langle \rangle, \emptyset \rangle, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto V'] \rangle, \mathcal{A}_{E_{\bar{k}}}(M_0) \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$$

```

0  ⟨ $\mathcal{A}_{V_{\bar{k}}}(V_2), \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
1  ⟨⟨begin (output debug  $\_C$ ) (cons  $\_C$   $\_C$ )⟩⟩,
     $\_E, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
4  ⟨⟨output debug (list  $\_C$   $\_C$ )⟩⟩,  $\emptyset[t_2 \mapsto \_V], \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
23 ⟨⟨false⟩⟩,  $\langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
24 ⟨⟨cons  $t_1$   $t_2$ ⟩⟩,  $\_E, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
25 ⟨ $t_1, \emptyset[t_1 \mapsto \_V],$ 
    ⟨⟨cons1 :  $t_2, \emptyset[t_2 \mapsto \_V],$ 
    ⟨ $\mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
26 ⟨ $\mathcal{A}_{V_{\bar{k}}}(V_1),$ 
    ⟨⟨cons1 :  $t_2, \emptyset[t_2 \mapsto \_V],$ 
    ⟨ $\mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
27 ⟨ $t_2, \emptyset[t_2 \mapsto \_V],$ 
    ⟨⟨cons2 :  $\mathcal{A}_{V_{\bar{k}}}(V_1),$ 
    ⟨ $\mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
28 ⟨ $\mathcal{A}_{V_{\bar{k}}}(V_2), \langle \langle \text{cons2} : \mathcal{A}_{V_{\bar{k}}}(V_1), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto \_V] \rangle \rangle, \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 
29 ⟨⟨pair :  $\mathcal{A}_{V_{\bar{k}}}(V_1), \mathcal{A}_{V_{\bar{k}}}(V_2) \rangle, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle$ 

```

Figure 3.15: Steps taken in reduction of cons2 continuation

where

$$C' = (\text{begin} (\text{output } \text{debug} (\text{list} (\text{list } \text{'valstep } t_1) (\text{c-c-m } \text{debug } k \dots))) (\text{car } t_1))$$

and V' is the result of evaluating the mark term.

Also by the definition of configuration annotation,

$$\mathcal{A}_{S_{\bar{k}}}(S_2) = \langle \text{error} \rangle$$

Figure 3.16 shows the key steps in the reduction of the *car* continuation frame to an error. The final step is equal to $\mathcal{A}_{S_{\bar{k}}}(S_2)$, and the lemma holds for this case.

Empty Stack The fourth and final case occurs when the continuation stack is empty. There are two sub-cases; one that arises when some definitions remain to be evaluated, and one that occurs when no definitions remain.

Suppose on the one hand that D is of the form $\langle E, f, \langle \langle f_1, C_1 \rangle, \dots \rangle \rangle$. By the definition of the abstract machine,

$$\begin{aligned} S_1 &= \langle V, \langle \langle \rangle, \emptyset \rangle, \langle E, f, \langle \langle f_1, C_1 \rangle, \langle f_2, C_2 \rangle, \dots \rangle \rangle \rangle \\ &\mapsto \langle C_1, \emptyset, \langle \langle \rangle, \emptyset \rangle, \langle E[f \mapsto V], f_1, \langle \langle f_2, C_2 \rangle, \dots \rangle \rangle \rangle = S_2 \end{aligned}$$

By definition of configuration annotation, furthermore,

```

0  ⟨⟨num : 0⟩, ⟨⟨app :  $_{-K}$ ⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
1  ⟨⟨begin (output debug  $_{-C}$ ) (car  $_{-C}$ )⟩,
   ⟨ $\emptyset[t_1 \mapsto \_V]$ , ⟨ $\mathcal{A}_{K_{\bar{k}}}(K, D)$ ,  $\mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \_V]$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
4  ⟨⟨output debug (list  $_{-C}$   $_{-C}$ )⟩,  $\emptyset[t_1 \mapsto \_V]$ , ⟨⟨app :  $_{-K}$ ⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
23 ⟨⟨false⟩, ⟨⟨app :  $_{-K}$ ⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
24 ⟨⟨car  $t_1$ ⟩,  $\emptyset[t_1 \mapsto \_V]$ , ⟨ $\mathcal{A}_{K_{\bar{k}}}(K, D)$ ,  $\mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \_V]$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
25 ⟨ $t_1$ ,  $\emptyset[t_1 \mapsto \_V]$ , ⟨⟨car : ⟨ $\mathcal{A}_{K_{\bar{k}}}(K, D)$ ,  $\mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \_V]$ ⟩⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
26 ⟨⟨num : 0⟩, ⟨⟨car : ⟨ $\mathcal{A}_{K_{\bar{k}}}(K, D)$ ,  $\mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \_V]$ ⟩⟩,  $\emptyset$ ⟩,  $\mathcal{A}_{D_{\bar{k}}}(D)$ ⟩
27 ⟨error⟩

```

Figure 3.16: Steps taken in reduction of car continuation leading to an error

$$\mathcal{A}_{S_{\bar{k}}}(S_1) = \langle \mathcal{A}_{V_{\bar{k}}}(V), \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_1 \rangle, C' \rangle, \emptyset \rangle \rangle, \langle \rangle, \emptyset, \langle \langle \rangle, \emptyset[debug \mapsto V'] \rangle \rangle, \mathcal{A}_{E_{\bar{k}}}(M_0) \rangle, \mathcal{A}_{D_{\bar{k}}}(\langle E, f, \langle \langle f_1, C_1 \rangle, \langle f_2, C_2 \rangle, \dots \rangle \rangle) \rangle$$

where

$$C' = (\text{begin} (\text{output debug} (\text{list} (\text{list 'valstep } t_1) (\text{c-c-m debug } k \dots))) t_1)$$

and V' is the result of evaluating the debug mark.

Also by the definition of configuration annotation,

$$\mathcal{A}_{S_{\bar{k}}}(S_2) = \langle \mathcal{A}_{C_{\bar{k}}}(C_1), \emptyset, \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_1 \rangle, C' \rangle, \emptyset \rangle \rangle, \langle \rangle, \emptyset, \langle \langle \rangle, \emptyset[debug \mapsto V'] \rangle \rangle, \emptyset[debug \mapsto \langle \text{false} \rangle] \rangle, \mathcal{A}_{D_{\bar{k}}}(\langle E[f \mapsto V], f_1, \langle \langle f_2, C_2 \rangle, \dots \rangle \rangle) \rangle$$

where

$$C' = (\text{begin} (\text{output debug} (\text{list} (\text{list 'valstep } t_1) (\text{c-c-m debug } k \dots))) t_1)$$

and V' is the result of evaluating the debug mark.

Figure 3.17 shows the key steps in the reduction of the empty continuation frame. The final step is equal to $\mathcal{A}_{S_{\bar{k}}}(S_2)$, and the lemma holds for this case.

Suppose on the other hand that D is of the form $\langle E, f, \langle \rangle \rangle$. By the definition of the abstract machine,

$$S_1 = \langle V, \langle \langle \rangle, \emptyset \rangle, \langle E, f, \langle \rangle \rangle \rangle \mapsto \langle E[f \mapsto V] \rangle = S_2$$

By definition of configuration annotation, furthermore,

```

0   ⟨AVk̄(V), ⟨⟨app : -K⟩, ∅⟩, ADk̄(⟨-E, f, ⟨⟨f1, C1⟩, ⟨f2, C2⟩, ...⟩⟩⟩)
1   ⟨⟨begin (output debug -C) t1⟩,
    ∅[t1 ↦ -V], ⟨⟨⟩, ∅[debug ↦ -V]]⟩, ADk̄(⟨-E, f, ⟨⟨f1, C1⟩, ⟨f2, C2⟩, ...⟩⟩)
4   ⟨⟨output debug (list -C -C)⟩,
    ∅[t1 ↦ -V], ⟨⟨app : -K⟩, ∅⟩, ADk̄(⟨-E, f, ⟨⟨f1, C1⟩, ⟨f2, C2⟩, ...⟩⟩)
23  ⟨⟨false⟩, ⟨⟨app : -K⟩, ∅⟩, ADk̄(⟨-E, f, ⟨⟨f1, C1⟩, ⟨f2, C2⟩, ...⟩⟩)
24  ⟨t1, ∅[t1 ↦ -V], ⟨⟨⟩, ∅[debug ↦ -V]]⟩, ADk̄(⟨-E, f, ⟨⟨f1, C1⟩, ⟨f2, C2⟩, ...⟩⟩)
25  ⟨AVk̄(V), ⟨⟨⟩, ∅⟩, ADk̄(⟨-E, f, ⟨⟨f1, C1⟩, ⟨f2, C2⟩, ...⟩⟩)
26  ⟨⟨w-c-m debug (list ...) (let ((t1 -C)) -C)⟩,
    ∅, ⟨⟨⟩, ∅⟩, ADk̄(⟨-E[f ↦ -V], f1, ⟨⟨f2, C2⟩, ...⟩⟩)
73  ⟨⟨let ((t1 (w-c-m debug -C -C))) (begin -C -C)⟩,
    ∅, ⟨⟨⟩, ∅[debug ↦ -V]]⟩, ADk̄(⟨-E[f ↦ -V], f1, ⟨⟨f2, C2⟩, ...⟩⟩)
76  ⟨⟨w-c-m debug false (begin -C -C)⟩,
    ∅, ⟨⟨app : -K⟩, ∅⟩, ADk̄(⟨-E[f ↦ -V], f1, ⟨⟨f2, C2⟩, ...⟩⟩)
79  ⟨⟨begin (output debug -C) ACk̄(C1)⟩,
    ∅, ⟨⟨app : -K⟩, ∅[debug ↦ -V]]⟩, ADk̄(⟨-E[f ↦ -V], f1, ⟨⟨f2, C2⟩, ...⟩⟩)
82  ⟨⟨output debug (list -C -C)⟩,
    ∅, ⟨⟨app : -K⟩, ∅⟩, ADk̄(⟨-E[f ↦ -V], f1, ⟨⟨f2, C2⟩, ...⟩⟩)
105 ⟨⟨false⟩, ⟨⟨app : -K⟩, ∅⟩, ADk̄(⟨-E[f ↦ -V], f1, ⟨⟨f2, C2⟩, ...⟩⟩)
106 ⟨ACk̄(C1)⟩,
    ∅, ⟨⟨app : -K⟩, ∅[debug ↦ -V]]⟩, ADk̄(⟨-E[f ↦ -V], f1, ⟨⟨f2, C2⟩, ...⟩⟩)

```

Figure 3.17: Steps taken in reduction of empty continuation

$$\mathcal{A}_{S_{\bar{k}}}(S_1) = \langle \mathcal{A}_{V_{\bar{k}}}(V), \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_1 \rangle, C' \rangle, \emptyset \rangle \rangle, \langle \rangle, \emptyset, \langle \langle \rangle, \emptyset[\text{debug} \mapsto V'] \rangle \rangle, \mathcal{A}_{E_{\bar{k}}}(M_0) \rangle, \langle \mathcal{A}_{E_{\bar{k}}}(E), f, \langle \rangle \rangle \rangle$$

where

$$C' = (\text{begin} (\text{output debug} (\text{list} (\text{list 'valstep } t_1) (\text{c-c-m debug } k \dots))) t_1)$$

and V' is the result of evaluating the mark term.

Also by definition of configuration annotation,

$$\mathcal{A}_{S_{\bar{k}}}(S_2) = \langle \mathcal{A}_{E_{\bar{k}}}(E)[f \mapsto \mathcal{A}_{V_{\bar{k}}}(V)] \rangle$$

Figure 3.18 shows the key steps in the reduction of the $\mathcal{A}_{S_{\bar{k}}}(S_1)$ configuration. The final step is equal to $\mathcal{A}_{S_{\bar{k}}}(S_2)$, and the lemma holds for this case.

This concludes our sketch of the proof of lemma 1.

```

0  <mathcal{A}_{V\bar{k}}(V), \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \langle \mathcal{A}_{E\bar{k}}(E), f, \langle \rangle \rangle \rangle
1  <mathbf{(begin (output debug \_C) t_1)},
   <math\emptyset[t_1 \mapsto \_V], \langle \langle \rangle, \emptyset[\text{debug} \mapsto \_V] \rangle, \langle \mathcal{A}_{E\bar{k}}(E), f, \langle \rangle \rangle \rangle
4  <mathbf{(output debug (list \_C \_C))},
   <math\emptyset[t_1 \mapsto \_V], \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \langle \mathcal{A}_{E\bar{k}}(E), f, \langle \rangle \rangle \rangle
23 <math\langle \text{false} \rangle, \langle \langle \text{app} : \_K \rangle, \emptyset \rangle, \langle \mathcal{A}_{E\bar{k}}(E), f, \langle \rangle \rangle \rangle
24 <math\langle t_1, \emptyset[t_1 \mapsto \_V], \langle \langle \rangle, \emptyset[\text{debug} \mapsto \_V] \rangle, \langle \mathcal{A}_{E\bar{k}}(E), f, \langle \rangle \rangle \rangle
25 <math\langle \mathcal{A}_{V\bar{k}}(V), \langle \langle \rangle, \emptyset \rangle, \langle \mathcal{A}_{E\bar{k}}(E), f, \langle \rangle \rangle \rangle
26 <math\langle \mathcal{A}_{E\bar{k}}(E)[f \mapsto \_V] \rangle

```

Figure 3.18: Steps taken in reduction of empty continuation with no further definitions

Loading Lemma

An additional lemma addresses the loader, \mathcal{L} .

Lemma 4 (Loading Lemma) *For any program P containing keys \bar{k} ,*

$$\mathcal{L}(\text{Annotate}(P)) \mapsto \mathcal{A}_{S\bar{k}}(\mathcal{L}(P))$$

proof: The program P must be of the form $(\text{define } f_1 C_1) (\text{define } f_2 C_2) \dots$.
By the definition of \mathcal{L} and Annotate,

$$\mathcal{L}(\text{Annotate}(P)) = \langle C', \emptyset, \langle \langle \rangle, \emptyset \rangle, \mathcal{A}_{D\bar{k}}(\langle \emptyset, f_1, \langle \langle f_2, C_2 \rangle, \dots \rangle \rangle) \rangle$$

where

$$\begin{aligned}
C' = & (\text{w-c-m debug (list 'outermost} \\
& \quad \text{null} \\
& \quad \text{null} \\
& \quad \text{null} \\
& \quad (\text{list 'f}_1 (\text{list 'cons 'f}_2 \mathcal{Q}(C_2))) \\
& \quad \text{null}) \\
& (\text{let } ((t_1 (\text{w-c-m debug false} \\
& \quad (\text{begin (output debug (list (list 'expstep } \mathcal{Q}(C_1) \text{ null)} \\
& \quad \quad (\text{c-c-m debug k } \dots))) \\
& \quad \quad \mathcal{A}_{C\bar{k}}(C_1)))))) \\
& (\text{begin (output debug (list (list 'valstep } t_1) (\text{c-c-m debug k } \dots))) \\
& \quad t_1)))
\end{aligned}$$

By the definition of \mathcal{A}_S and \mathcal{L} ,

that there is some sequence of configurations S_1, \dots, S_n where $S_1 = \mathcal{L}(P)$ and $S_i \mapsto S_{i+1}$ for $1 \leq i < n$ and $S_n = S$.

Lemma 4 states that $\mathcal{L}(\text{Annotate}(P)) \mapsto \mathcal{A}_{S_{\bar{k}}}(\mathcal{L}(P)) = \mathcal{A}_{S_{\bar{k}}}(S_1)$. Lemma 5 shows that $\text{Clean}(S_1)$.

Lemmas 1 and 5 now show through induction that for any i where $1 \leq i < n$, $\mathcal{A}_{S_{\bar{k}}}(S_i) \mapsto \mathcal{A}_{S_{\bar{k}}}(S_{i+1})$.

Inspection of the definition of \mathcal{A}_S shows that if $\mathcal{A}_S(S)$ is a final state, then so is S .

Therefore, $\text{Eval}(\text{Annotate}(P)) = \mathcal{A}_{S_{\bar{k}}}(S_n)$.

To see the right-to-left implication, suppose that $\text{Eval}(\text{Annotate}(P)) = \mathcal{A}_{S_{\bar{k}}}(S)$. Now, either $\text{Eval}(P)$ is defined or it is not. If it is not defined, it must be the case that evaluation of P diverges, and in this case lemma 1 tells us that the evaluation of $\text{Annotate}(P)$ must also diverge, and since the \mapsto relation is a function, this contradicts our claim that $\text{Eval}(\text{Annotate}(P)) = \mathcal{A}_{S_{\bar{k}}}(S)$. On the other hand, suppose that $\text{Eval}(P)$ exists and is some S' . In this case, the left-to-right implication tells us that $\text{Eval}(\text{Annotate}(P)) = \mathcal{A}_{S_{\bar{k}}}(S')$. Since \mapsto is a function and so is Annotate , we know that $S = S'$, and we're finished.

Proof Sketch of Theorem 2: Inspection shows that output on port o occurs exactly on the reduction of configurations of the form $\langle V, \langle \langle \text{out} : o, \langle K, M_1 \rangle \rangle, M_0 \rangle, D \rangle$. Case analysis of the abstract machine and of configuration annotation shows that the annotations of such configurations also produce output on port o , and that the output emitted is $\mathcal{A}_{V_{\bar{k}}}(V)$.

Conversely, if $S_i \mapsto S_{i+1}$ and S_i is not of the form shown above, then case analysis of the abstract machine and of configuration annotation (along with Lemma 1) show that the configuration $\mathcal{A}_{S_{\bar{k}}}(S_i)$ reduces in one or more steps to $\mathcal{A}_{S_{\bar{k}}}(S_{i+1})$ without producing output on port o .

Finally, observe that the reduction of $\mathcal{L}(\text{Annotate}(P))$ reduces in one or more steps to $\mathcal{A}_{S_{\bar{k}}}(S_1)$ without producing output on any port other than *debug*.

Induction on the length of the length of the evaluation sequence n therefore demonstrates that $\mathcal{A}_{V_{\bar{k}}}(\text{Trace}_o(S_1 \mapsto S_n)) = \text{Trace}_o(\mathcal{L}(\text{Annotate}(P)) \mapsto \mathcal{A}_{S_{\bar{k}}}(S_n))$

Correctness

Furthermore, the system functions as a stepper. That is, applying the reconstruction function to the sequence of debug values emitted by the annotated program produces the same sequence of steps as the evaluation of the original program. Theorem 3 states this precisely.

To prove the theorem, we use lemma 1 to divide the evaluation of the annotated program into segments corresponding to the steps taken by the source program.

Figure 3.20 illustrates the timing and correspondence of debug outputs with the steps taken by the original program, and shows why it is that the final configuration in the source series appears only when it is an Expression configuration. In particular, the debug output corresponding to an Expression configuration occurs in the segment preceding that annotated Expression

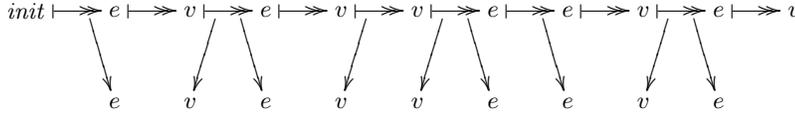


Figure 3.20: The correspondence between steps taken and steps emitted

Kind	Outputs	Contents
$E \rightarrow V$	0	constants, lambda,c-c-m , variables
$E \rightarrow E$	1	all expressions with subexpressions
$V \rightarrow E$	1	<i>cons2</i> group, empty continuation with exps remaining
$V \rightarrow E$	2	<i>cons1</i> group, <i>wcm</i> group
$V \rightarrow \text{Finished}$	1	empty continuation with no more expressions
$V \rightarrow \text{Error}$	1	arity error, <i>car/cdr/if</i> errors
$E \rightarrow \text{Error}$	0	bad reference to top-level variable

Figure 3.21: Transition Categories

configuration, whereas the debug output corresponding to a Value configuration occurs in the segment *following* that annotated Value configuration. This means, for example, that the segment leading from an annotated Expression configuration to an annotated Value configuration emits no debug outputs, and that the segment leading from a Value configuration to an Expression configuration emits two debug outputs.

The initial configuration in the source sequence must be an Expression configuration, by the definition of \mathcal{L} . This initial Expression configuration is emitted by the annotated program by the sequence of configurations leading from $\mathcal{L}(\text{Annotate}(P))$ to $\mathcal{A}_{S\bar{k}}(S_1)$.

The result of reconstruction is a simple sequence of configurations, and therefore does not contain the outputs produced by the original execution. However, the earlier observation of the one-to-one correspondence between outputs and configurations of the form $\langle V, \langle \langle \text{out} : o, \langle K, M_1 \rangle \rangle, M_0 \rangle, D \rangle$ shows that these outputs may be recovered directly from the reconstructed sequence of configurations.

We prove the theorem by considering each possible transition in the source sequence, and showing that the segment that is its annotated counterpart produces the expected output. As in the proof of lemma 1, we can group transitions into classes. The grouping used for this proof is a coarsening of the classes used in that proof. The table in figure 3.21 summarizes these larger groups, and indicates which of the earlier classes fall into each group.

The proof proceeds by case analysis. All of the cases depend upon two lemmas; one that states that reconstruction applied to the parent of a continuation's annotation recovers that continuation, and one that states that the definitions may be recovered from the annotated continuation.

Lemma 6 (Continuation Reconstruction) *For any K with keys \vec{k} and any D where $\text{Clean}(K)$, and $\text{Clean}(D)$,*

$$\mathcal{R}_K(\pi_{\vec{k}}(\langle K', M' \rangle)) = K$$

where $\mathcal{A}_{K\vec{k}}(K, D) = \langle \dots, \langle K', M' \rangle \rangle$ and $k' = \langle \text{debug}, k, \dots \rangle$.

The proof of this lemma proceeds by induction on the length of the continuation. In each case, the result of the annotation is a $\langle K, M \rangle$ where information about each continuation frame is stored in a *debug* mark attached to the parent continuation. Discarding the outermost continuation, extracting these marks with the π function and applying \mathcal{R}_K to the result yields the original continuation pair.

Lemma 7 (Definition Reconstruction) *For any K with keys \vec{k} and any D where $\text{Clean}(K)$ and $\text{Clean}(D)$, and any M at all,*

$$\mathcal{R}_D(\pi_{\vec{k}}(\langle \mathcal{A}_{K\vec{k}}(K, D), M \rangle)) = D$$

where $k' = \langle \text{debug}, k, \dots \rangle$.

The definitions passed to \mathcal{A}_K are stored in the outermost continuation. The function R_D retrieves the definitions stored in a continuation by recurring on the continuation's mark list.

Given these lemmas, the cases in the proof of theorem 3 are quite similar to each other. We take as a representative the value configuration that has a topmost *wcm* continuation frame. This case illustrates the output of both Value and Expression configurations. Figure 3.13, from the proof of the earlier simulation lemma, serves to illustrate the key steps in this reduction segment.

As before, the definition of the abstract machine shows that

$$S_1 = \langle V, \langle \langle \text{wcm} : k, C, E, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \mapsto \langle C, E, \langle K, M[k \mapsto V] \rangle, D \rangle = S_2$$

By the definition of annotation and the abstract machine, the corresponding segment $\mathcal{A}_{S\vec{k}}(S_1) \mapsto \mathcal{A}_{S\vec{k}}(S_2)$ contains exactly two configurations that produce output. We must show that the first one produces an output that goes by reconstruction to S_1 , and that the second one produces an output that goes by reconstruction to S_2 .

By the definition of annotation and the abstract machine, the first such configuration is

$$S'_1 = \langle V_1, \langle \langle \text{out} : \text{debug}, \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_{dc} \rangle, (\mathbf{w-c-m} \ k \ t_1 \ \mathcal{E}_{\vec{k}}(C)) \rangle, _E \rangle \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K\vec{k}}(K, D), \mathcal{A}_{E\vec{k}}(M)[\text{debug} \mapsto V_2] \rangle, \emptyset \rangle, \emptyset \rangle, \mathcal{A}_{D\vec{k}}(D) \rangle \rangle$$

where V_2 is the current debug mark, and

$$V_1 = \langle \text{pair} : \langle \text{pair} : \langle \text{sym} : \text{valstep} \rangle, \langle \text{pair} : \mathcal{A}_{V_{\bar{k}}}(V), \langle \text{null} \rangle \rangle \rangle, \\ \langle \text{pair} : \pi_{\bar{k}'}(\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V_3] \rangle), \\ \langle \text{null} \rangle \rangle \rangle$$

where

$$V_3 = (\text{list } 'wcm \\ \text{null} \\ (\text{list } 'x x) \dots ;; x \in fv(C) \\ \text{null} \\ (\text{list } 'k \mathcal{Q}(C)) \\ \text{null})$$

As before, we have used the input syntax, including the **list** abbreviation, to make this value more readable. By the earlier observation, the value V_1 appears on the *debug* output port. By the definition of reconstruction,

$$\begin{aligned} \mathcal{R}_S(V_1) &= \langle \mathcal{A}_V^{-1}(\mathcal{A}_{V_{\bar{k}}}(V)), \langle \mathcal{R}_K(\pi_{\bar{k}'}(\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V_3] \rangle)), \emptyset \rangle, \\ &\quad \mathcal{R}_D(\pi_{\bar{k}'}(\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V_3] \rangle)) \\ &= \langle V, \langle \mathcal{R}_K(\pi_{\bar{k}'}(\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V_3] \rangle)), \emptyset \rangle, \\ &\quad \mathcal{R}_D(\pi_{\bar{k}'}(\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V_3] \rangle)) \\ &= \langle V, \langle \mathcal{R}_K(\pi_{\bar{k}'}(\langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[\text{debug} \mapsto V_3] \rangle)), \emptyset \rangle, D \rangle \\ &= \langle V, \langle \mathcal{R}_K(\mathbf{cons}(\mathbf{cons}(\text{list } 'debug V_3) \phi_{\bar{k}'}(\mathcal{A}_{E_{\bar{k}}}(M))) \\ &\quad \pi_{\bar{k}'}(\langle K_1, M_1 \rangle)), \emptyset \rangle, D \rangle \\ &\quad \text{where } \mathcal{A}_{K_{\bar{k}}}(K, D) = \langle \dots, \langle K_1, M_1 \rangle \rangle \\ &= \langle V, \langle \langle wcm : k, C, E, \langle \mathcal{R}_K(\pi_{\bar{k}'}(\langle K_1, M_1 \rangle)), \mathcal{R}_E(\phi_{\bar{k}'}(\mathcal{A}_{E_{\bar{k}}}(M))) \rangle \rangle, \emptyset \rangle, D \rangle \\ &= \langle V, \langle \langle wcm : k, C, E, \langle \mathcal{R}_K(\pi_{\bar{k}'}(\langle K_1, M_1 \rangle)), M \rangle \rangle, \emptyset \rangle, D \rangle \\ &= \langle V, \langle \langle wcm : k, C, E, \langle K, M \rangle \rangle, \emptyset \rangle, D \rangle \end{aligned}$$

So the result of reconstruction is the desired configuration S_1 .

The second output configuration is

$$S_2' = \langle V_1, \\ \langle \text{out} : \text{debug}, \\ \langle \langle \text{app} : \langle \langle \text{clo} : \langle t_{dc} \rangle, \mathcal{A}_{C_{\bar{k}}}(C), -E \rangle \rangle, \langle \rangle, \emptyset, \\ \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[k \mapsto \mathcal{A}_{V_{\bar{k}}}(V)][\text{debug} \mapsto \langle \text{false} \rangle] \rangle, \emptyset \rangle, \emptyset \rangle, \\ \mathcal{A}_{D_{\bar{k}}}(D) \rangle \rangle$$

where

$$\begin{aligned}
V_1 &= (\mathbf{list} \ (\mathbf{list} \ 'expstep \\
&\quad \mathcal{Q}(C) \\
&\quad (\mathbf{list} \ 'x \ x) \ \dots) \ ;\; \text{for } x \in fv(C) \\
&\quad \pi_{\vec{k}'}(\langle \mathcal{A}_{K_{\vec{k}'}}(K, D), \mathcal{A}_{E_{\vec{k}'}}(M)[k \mapsto \mathcal{A}_{V_{\vec{k}'}}(V)][debug \mapsto \langle \text{false} \rangle] \rangle)) \\
&= (\mathbf{list} \ (\mathbf{list} \ 'expstep \\
&\quad \mathcal{Q}(C) \\
&\quad (\mathbf{list} \ 'x \ x) \ \dots) \ ;\; \text{for } x \in fv(C) \\
&\quad (\mathbf{cons} \ (\mathbf{cons} \ (\mathbf{list} \ 'debug \ \langle \text{false} \rangle) \\
&\quad \quad (\mathbf{cons} \ (\mathbf{list} \ 'k \ \mathcal{A}_{V_{\vec{k}'}}(V)) \\
&\quad \quad \quad \phi_{\vec{k}'}(\mathcal{A}_{E_{\vec{k}'}}(M)))) \\
&\quad \pi_{\vec{k}'}(\langle K_1, M_1 \rangle)))
\end{aligned}$$

where

$$\mathcal{A}_{K_{\vec{k}'}}(K, D) = \langle \dots \langle K_1, M_1 \rangle \rangle$$

By the definition of reconstruction,

$$\begin{aligned}
\mathcal{R}_S(V_1) &= \langle C, E, \langle \mathcal{R}_K(\langle K_1, M_1 \rangle), \mathcal{R}_E(\mathbf{cons} \ (\mathbf{list} \ 'k \ \mathcal{A}_{V_{\vec{k}'}}(V)) \ \phi_{\vec{k}'}(\mathcal{A}_{E_{\vec{k}'}}(M))) \rangle, \\
&\quad \mathcal{R}_D \pi_{\vec{k}'}(\langle K_1, M_1 \rangle) \rangle \\
&= \langle C, E, \langle \mathcal{R}_K(\langle K_1, M_1 \rangle), M[k \mapsto V] \rangle, \mathcal{R}_D \pi_{\vec{k}'}(\langle K_1, M_1 \rangle) \rangle \\
&= \langle C, E, \langle K, M[k \mapsto V] \rangle, \mathcal{R}_D \pi_{\vec{k}'}(\langle K_1, M_1 \rangle) \rangle \\
&= \langle C, E, \langle K, M[k \mapsto V] \rangle, D \rangle
\end{aligned}$$

So the result of reconstruction is the desired configuration S_1 .

This concludes the sketch of the proof of theorem 3.

3.5 Notes on Proof Methods

Each of the cases in the principal lemmas above consists of many steps in an evaluator. In order to prevent the errors that are the hallmark of similar hand-generated proofs, we implemented a simple evaluator in OCAML [32] for the language of the model, extended with certain generalizations, that can perform computation on incomplete program terms. In essence, this program is an ad-hoc theorem prover, designed for a particular language.

Appendix A shows some of the key pieces of this evaluator, including the datatype definitions and the language's generalized evaluator.

The datatype definitions for the language extend the definitions given in section 3.2 with *ANYEXP* and *ANYVAL* variants that represent arbitrary

expressions and values, respectively. We can use these variants to express partially specified terms such as **(if A B C)**.

Furthermore, each expression may have a label that represents knowledge about that term necessary to allow the proof step to proceed. For instance, after proving lemma 2, we may label terms that are known to be the result of a “quote”. Then, we extend the evaluator with a rule stating that such expressions always converge, and may safely be replaced with a fresh *ANYVAL*.

By interactively extending the evaluator with labels that correspond to lemmas that we have proven separately, we may generate proof steps in a systematic way, ruling out a large class of errors.

3.6 The Pragmatics of a Stepper

The model of section 3.2 represents the kernel of our stepper. DrScheme’s stepper is implemented with an annotation based on the one described, and is an entirely unprivileged program.

Adapting the model for the desired implementation, however, requires several adjustments. Most significantly, our stepper displays a series of steps in an algebraic reduction semantics, rather than a register machine. Additionally, the stepper must deal with errors, opaque values, unannotated library code and Scheme macros. Finally, our implementation makes use of several simple but significant optimizations to reduce the memory consumed by the stepper. In the following subsections, we explain each of these differences between the model and our implementation.

Algebraic Reduction Semantics

We believe that the best evaluation models to use with beginning programmers are those that highlight the equivalence of program terms and the algebraic nature of program reduction. For this reason, our stepper displays a series of steps in an algebraic reduction semantics, rather than a sequence of machine configurations.

Fortunately, Flatt and Felleisen have demonstrated the correspondence between the steps taken in a β_v reduction semantics and the steps taken by a register machine of the kind this chapter describes [15]. The stepper uses this correspondence to map the steps shown by reconstruction back to the terms in a reduction sequence. This means that every step the user sees is a complete program.

The sequence of steps taken by a register machine is longer than the chain of steps validated by the reduction semantics. This is because many of the transformations in the register machine—the reduction of the syntactic term **false** to the value $\langle \text{false} \rangle$, for instance—correspond to the identity transformation in the reduction semantics. The stepper filters out such steps.

Exceptions

In our model, a program that causes an error is elaborated into a program that causes the same error. In our implementation, we do not want the program's defect to preclude further debugging. The annotated program therefore intercepts uncaught exceptions, and displays them to the user as a final step in the reduction sequence. Since the stepper stores all steps as they are generated, the user may step backward to determine the exception's cause.

Opaque Values

Since the stepper is an ordinary program and has no privileged access to compiler or run-time data, closures are effectively opaque; there is no run-time operation that accepts a closure and returns its source code or the bindings it is closed over. Rather than compromise the division between stepper and compiler or disallow access entirely, the stepper elaborates **lambda** expressions so that these items' source texts and bindings appear in a weak hash table whose keys are the values themselves.⁷ At reconstruction time, the stored records may then be retrieved for display.

Unannotated Code

Source programs often interact with other program components. These components may consist of large bodies of code, may be available only in compiled form, or may be written in another language. In this situation, the debugger cannot hope to annotate all code involved in a computation.

In fact, the debugger might stop working altogether; a debugger based on an annotation that changes the calling conventions, for example, would simply fail if the annotated program were naïvely linked to unannotated code.

Fortunately, an annotation based on continuation marks fails gracefully. On the one hand, most library procedures behave as atomic primitives; in the stepper's reconstruction, the step preceding the call is followed by the result of the call, and the effect is that these library procedures implicitly extend the source language. On the other hand, some higher-order primitives may call annotated code. In this case, reconstruction must proceed with incomplete information. In these cases, reconstruction simply shows gaps in the program context.

Scheme Macros and Little Languages

DrScheme contains a hygienic source-correlating macro system. This means that users can easily extend the source language with new forms, or change the meaning of existing ones. It also means that before evaluation, every program

⁷Values reachable only through references in weak hash tables may be collected; this guarantees that the hash table itself does not affect the asymptotic memory behavior of the program.

is expanded until it matches a core syntax. So, for instance, Scheme’s **cond** form is expanded into a series of **ifs**.

Rather than try to accommodate every new macro added to the system, our stepper operates on elements of the core language, after the source program is fully expanded. This is a robust solution that accommodates the addition of arbitrary macros to the system.

This means that the stepper can in principle be applied to any language built on top of MzScheme. We explore this further in our work on Little Languages [10, 8].

The penalty for this decision is that reconstruction produces terms in the core language, rather than in the extended language. In a program that makes extensive use of macros, or a program written by a user who is not aware of the language’s core syntax, the output of reconstruction may be confusing.

Fortunately, the source-correlation of the syntax system means that it is possible to translate these terms back into source syntax, in the cases where such translations are known. In DrScheme’s teaching languages, for instance, a defined set of macros provide additional error-checking for beginners, and users may not extend the language themselves.

Little Languages

The Content of Marks

Two simple annotations vastly reduce the memory used by annotated programs in our implementation.

Firstly, the compiler and the execution of the compiled code share the same heap. This means that the annotator can store references to the source code in the source code itself, rather than encoding source code using the quoting function \mathcal{Q} as we do in the model. Furthermore, these references do not require evaluation, as they are already values. So the cost of encoding and evaluating these quoted source terms is eliminated.

Secondly, the implementation also reduces the size of marks by storing only those bindings that are not known to occur in an enclosing mark. So, given a source program of the form **(cons (cons a b) c)**, the outer **cons** will have a mark that captures the value of b . This means that the mark on the inner **cons** does not need to capture the value of b .

3.7 Related Work

The idea of elaborating a program in order to observe its behavior is a familiar one. Early systems included BUGTRAN [18] and EXDAMS [1] for FORTRAN. More recent applications of this technique to higher-order languages include Tolmach’s *sml*d [44], Kellomaki’s PSD [27], and several projects in the lazy FP community [25, 36, 39, 42]. None of these, however, addressed the correctness of the tool—not only that the transformation preserves the meaning of the

program, but also that the information divulged by the elaborated program matches the intended purpose.

Indeed, work on modeling the action of programming environment tools is sparse. Bernstein and Stark [4] put forward the idea of specifying the semantics of a debugger. That is, they specify the actions of the debugger with respect to a low-level machine. We extend this work to show that the tool preserves the semantics and also performs the expected computation.

Kishon, Hudak, and Consel [31] study a more general idea than Bernstein and Stark. They describe a theoretical framework for extending the semantics of a language to include execution monitors. Their work guarantees the preservation of the source language's semantics. Our work extends this (albeit with a loss of generality) with a proof that the information output by the tool is sufficient to reconstruct a source expression.

Bertot [5] describes a semantic framework for relating an intermediate state in a reduction sequence to the original program. Put differently, he describes the semantic foundation for source tracking. In contrast, we exploit a practical implementation of source tracking by Matthew Flatt for our implementation of the stepper. Bertot's work does not verify a stepper but simply assumes that the language evaluator *is* a stepper.

Chapter 4

Stack Inspection implemented with Continuation Marks

This chapter looks at the application of continuation marks to stack inspection for security reasons. Specifically, it shows how the stack inspection model of Fournet and Gordon [22] may be implemented as a tail-calling evaluator, using a register machine with a variant of continuation marks.¹

4.1 Stacks, Security, and Tail Calls

Over the last ten years, programming language implementors have spent significant effort on security issues. This effort takes many forms; one is the implementation of a strategy known as stack inspection [46]. It starts from the premise that trusted components may authorize potentially insecure actions for the dynamic extent of some expression, provided that all intermediate calls are made by and to trusted code. When an insecure action is requested, these trusted components must inspect the current stack frames to ensure that it was properly authorized.

In its conventional implementation, stack inspection is incompatible with a traditional language semantics, because it clashes with the well-established idea of modeling function calls with a β or β_v reduction [37]. A β reduction replaces a function's application with the body of that function, where the application's arguments replace the function's parameters. In a language with stack inspection, a β or β_v reduction thus disposes of information that is necessary to evaluate the security primitives.

For this reason, Fournet and Gordon [22] model function calls with a non-standard β reduction. To be more precise, β does not hold as an equation for source terms. Instead, abstraction bodies are wrapped with context-building primitives. Unfortunately, this formalization inhibits a transformation of this semantics into a tail-calling implementation. Fournet and Gordon recognize

¹This paper appeared in an earlier form in Transactions on Programming Languages and Systems [7], and before that at the European Symposium on Programming [6].

this fact and state that “[S]tack inspection profoundly affects the semantics of all programs. In particular, it invalidates [...] tail call optimizations [22].”

This understanding of the stack inspection protocol also pervades the implementation of existing run-time systems. The Java design team, for example, chose not to provide a tail-calling implementation in part because of the perceived incompatibility between tail-calling and stack inspection.² The .NET effort at Microsoft provides a runtime system that is properly tail-calling—except in the presence of security primitives, which disable it. Microsoft’s documentation [34] states that “[t]he current frame cannot be discarded when control is transferred from untrusted code to trusted code, since this would jeopardize code identity security.”

Wallach et al. [47] suggest an alternate implementation of stack inspection that might accommodate tail-calling. They add an argument to each function call that represents the security context as a statement in their belief logic. Statements in this belief logic can be unraveled to determine whether an operation is permitted. However, the details of their memory behavior are opaque. In particular, they do not attempt to present a model in which memory usage can be analyzed.

Our work fills the gap between Fournet and Gordon’s formal model and Wallach’s alternative implementation of stack inspection. Specifically, our security model exploits a novel mechanism for lightweight stack inspection [21], derived from continuation marks. We demonstrate the equivalence between our model and Fournet and Gordon’s, and prove our claims of proper tail-calling. More precisely, our abstract implementation can transform *all* tail calls in the source program into instructions that do not consume any stack (or store) space. Moreover, our abstract implementation represents a relatively minor change to the models used by current implementations, suggesting that these implementations might accommodate tail-calling with minimal effort.

We proceed as follows. First, we derive a CESK machine from Fournet and Gordon’s semantics. Second, we develop a different, but extensionally equivalent CESK machine that uses a variant of continuation marks. Third, we show that our machine uses strictly less space than the machine derived from Fournet and Gordon’s semantics and that our machine uses as much space as Clinger’s canonical tail-calling CESK machine [11].

The chapter consists of nine sections. The second section introduces the λ_{sec} language: its syntax, semantics, and security mechanisms. The third section shows how a pair of tail calls between system and applet code can allocate an unbounded amount of space. In the fourth section, we derive an extensionally equivalent CESK machine from Fournet and Gordon’s semantics; in the fifth section, we modify this machine so that it implements all tail calls in a properly optimized fashion. The sixth section provides a precise analysis of the space consumption of these machines and shows that our new machine is indeed tail-calling. In the seventh section, we discuss the extension of our models for λ_{sec}

²Private communication between Guy Steele and Matthias Felleisen at POPL 1996

to the richer environments of existing languages. The last two sections place our work into context.

4.2 The λ_{sec} Language

Fournet and Gordon work from the λ_{sec} -calculus [38, 41]. This calculus is a simple model of a programming language with security annotations. They present two languages: a source language, in which program components are written, and a target language, which includes an additional form for security annotations. A trusted annotator performs the translation from the source to the target, annotating each component with the appropriate permissions.

In this security model, all code is statically annotated with a given set of permissions, chosen from a fixed set \mathcal{P} . A program component that has permissions R may choose to enable some or all of these permissions. The set of enabled permissions at any point during execution is determined by taking the intersection of the permissions enabled for the caller and the set of permissions contained in the callee’s annotation. That is, a permission is considered enabled only if two conditions are met: first, it must have been legally and explicitly enabled by some calling procedure, and second, all intervening callers must have been annotated with this permission.

A program component consists of a set of permissions and a λ -expression from the source language, (M_s) . This language adds three expressions to the basic call-by-value λ -calculus. The **test** expression checks to see whether a given set of permissions is currently enabled, and branches based on that decision. The **grant** expression enables a privilege, provided that the context endows it with those permissions. Finally, the **fail** expression causes the program to halt immediately, signaling a security failure. Our particular source language also changes the traditional presentation of the λ -calculus by adding an explicit name to each abstraction so that we get concise definitions of recursive procedures.

Syntax

$$\begin{aligned}
 C \in \text{Components} &= \langle R, \lambda_f x. M_s \rangle \\
 M, N &= x \mid M \ N \mid \lambda_f x. M \mid \mathbf{grant} \ R \ \mathbf{in} \ M \\
 &\quad \mid \mathbf{test} \ R \ \mathbf{then} \ M \ \mathbf{else} \ N \mid \mathbf{fail} \mid \underline{R[M]} \\
 x &\in \text{Identifiers} \\
 R &\subseteq \mathcal{P} \\
 V \in \text{Values} &= x \mid \lambda_f x. M
 \end{aligned}$$

The target language (M) adds a framing expression to this source language (underlined in the grammar). A frame specifies the permissions of a component in the source text. To ensure that these framing expressions are present as the program is evaluated, we translate source components into target components by annotating the component’s source term with its permissions. The annotator below performs this annotation, and simultaneously ensures that a **grant**

expression refers only to those permissions to which it is entitled by its source location.

$$\begin{aligned}
 \text{Annotator } \mathcal{A} : 2^{\mathcal{P}} \times M_s &\rightarrow M \\
 \mathcal{A}\langle R, \llbracket x \rrbracket \rangle &= x \\
 \mathcal{A}\langle R, \llbracket \lambda_f x. M \rrbracket \rangle &= \lambda_f x. R[\mathcal{A}\langle R, \llbracket M \rrbracket \rangle] \\
 \mathcal{A}\langle R, \llbracket M N \rrbracket \rangle &= \mathcal{A}\langle R, \llbracket M \rrbracket \rangle \mathcal{A}\langle R, \llbracket N \rrbracket \rangle \\
 \mathcal{A}\langle R, \llbracket \text{grant } S \text{ in } M \rrbracket \rangle &= \text{grant } S \cap R \text{ in } \mathcal{A}\langle R, \llbracket M \rrbracket \rangle \\
 \mathcal{A}\langle R, \llbracket \text{test } S \text{ then } M \text{ else } N \rrbracket \rangle &= \text{test } S \text{ then } \mathcal{A}\langle R, \llbracket M \rrbracket \rangle \text{ else } \mathcal{A}\langle R, \llbracket N \rrbracket \rangle \\
 \mathcal{A}\langle R, \llbracket \text{fail} \rrbracket \rangle &= \text{fail}
 \end{aligned}$$

The annotator \mathcal{A} consumes two arguments: the set of permissions appropriate for the source, and the source code; it produces a target expression. It commutes with all expression constructors except for λ and **grant**. For a λ expression, it adds a frame expression wrapping the body. For a **grant** expression, it replaces the permissions S that the expression specifies with the intersection $S \cap R$. So, if a component containing the expression **grant** $\{a, b\}$ in E were annotated with the permissions $\{b, c\}$, the resulting expression would read **grant** $\{b\}$ in E' , where E' represents the recursive annotation of E .

We adapt Fournet and Gordon’s semantics to our variant of λ_{sec} . Evaluation of programs is specified using a reduction semantics based on evaluation contexts [16]. In such a semantics, every expression is divided into an evaluation context containing a single hole (denoted by \bullet), and a redex. An evaluation context is composed with a redex by replacing the context’s hole with the redex. The choice of evaluation contexts determines where evaluation can occur, and typically the evaluation contexts are chosen to enforce deterministic evaluation; that is, each expression has a unique decomposition into context and redex. Reduction rules in such a semantics take the form “ $E[f] \mapsto E[g]$,” where f is a redex, g is its contractum, and E is the context (which may be observable, as for instance in the **test** rule).

Contexts

$$E = \bullet \mid E M \mid V E \mid \text{grant } R \text{ in } E \mid R[E]$$

Reduction Rules

$$\begin{aligned}
 E[\lambda_f x. M V] &\mapsto E[[\lambda_f x. M / f][V / x]M] \\
 E[R[V]] &\mapsto E[V] \\
 E[\text{grant } R \text{ in } V] &\mapsto E[V] \\
 E[\text{test } R \text{ then } M \text{ else } N] &\mapsto \begin{cases} E[M] & \text{if } \mathcal{OK}\langle R, \llbracket E \rrbracket \rangle \\ E[N] & \text{otherwise} \end{cases} \\
 E[\text{fail}] &\mapsto \text{fail}
 \end{aligned}$$

where

$$\begin{array}{ll}
 \mathcal{OK}\langle\emptyset, [E]\rangle & \\
 \mathcal{OK}\langle R, [\bullet]\rangle & \\
 \mathcal{OK}\langle R, [E[\bullet M]]\rangle & \text{iff } \mathcal{OK}\langle R, [E]\rangle \\
 \mathcal{OK}\langle R, [E[V \bullet]]\rangle & \text{iff } \mathcal{OK}\langle R, [E]\rangle \\
 \mathcal{OK}\langle R, [E[S[\bullet]]]\rangle & \text{iff } R \subseteq S \text{ and } \mathcal{OK}\langle R, [E]\rangle \\
 \mathcal{OK}\langle R, [E[\text{grant } S \text{ in } \bullet]]\rangle & \text{iff } \mathcal{OK}\langle (R - S), [E]\rangle
 \end{array}$$

This semantics is an extension of a standard call-by-value reduction semantics. The hole and the two application contexts are standard and enforce left-to-right evaluation of arguments. The reduction rule for applications is also standard. The added contexts and reduction rules for frame and **grant** expressions are largely transparent; evaluation may proceed inside of either form, and each one disappears when its expression is a value. These expressions affect the evaluation only when a **test** expression occurs as a redex. In this case, the result of the reduction depends on the \mathcal{OK} predicate, which is applied to the current context and the desired permissions.

The \mathcal{OK} predicate recurs over the evaluation context from the inside out, succeeding either when the permissions remaining to check are empty or when the context is exhausted.³ The \mathcal{OK} predicate commutes with both kinds of application context. In the case of a frame annotation, the desired permissions must occur in the frame, and the predicate must succeed recursively. Finally, a **grant** expression removes all permissions it grants from the set of those that need to be checked.

Finally, the Eval function determines the meaning of a source program. A program consists of a list of components. Evaluation is performed by annotating each λ -expression with the permissions of its component, and combining all such expressions into a single application. This application uses the traditional abbreviation of a curried application as a single one.

Definition 1 (Eval)

$$\text{Eval}(C, \dots) = V \text{ if } (\mathcal{A}(C) \ \dots) \mapsto^* V$$

Since the first component is applied to the rest, it is presumed to represent the runtime system, or at least a linker. Eval is undefined for programs that diverge or enter a stuck state.

Minor Differences The semantics we present differs from that of Fournet and Gordon in three ways. First, it reduces programs containing **fail** to a final **fail** state in one step, rather than propagating the **fail** upward one expression at a time. We consider this difference trivial and ignore it. Second, our language includes a named lambda, to simplify the presentation of recursive examples.

³In fact, success on an empty permission set may be derived from the other rules in the definition; the direct statement of this is nevertheless included to simplify understanding.

Since we present an untyped language, a recursive function always has an equivalent form as an application of the Y combinator. Third, our semantics replaces a run-time check in Fournet and Gordon’s semantics with a static check. Appendix B.1 presents a proof of the equivalence of our evaluator and theirs.

4.3 Tail-calling

Chapter 2 details the argument for the importance of tail-calling evaluators. Unfortunately, languages such as Java have no tail-calling implementations, and Microsoft’s CLR supports tail-calling only in the absence of stack inspection.

At first glance, tail call optimization seems inherently incompatible with stack inspection. To see this, consider a mutually recursive loop between applet and library code.

Abbreviations

$$\begin{aligned} \text{UserFn} &\triangleq \lambda_{user\ sys}.sys\ user \\ \text{SystemFn} &\triangleq \lambda_{sys\ user}.user\ sys \\ \mathcal{A}\langle R_A, \llbracket \text{UserFn} \rrbracket \rangle &= \lambda_{user\ sys}.R_A[sys\ user] \\ \mathcal{A}\langle R_S, \llbracket \text{SystemFn} \rrbracket \rangle &= \lambda_{sys\ user}.R_S[user\ sys] \end{aligned}$$

Reduction (w/ Annotations)

$$\begin{aligned} &\mathcal{A}\langle R_A, \llbracket \text{UserFn} \rrbracket \rangle\ \mathcal{A}\langle R_S, \llbracket \text{SystemFn} \rrbracket \rangle \\ \mapsto &R_A[\mathcal{A}\langle R_S, \llbracket \text{SystemFn} \rrbracket \rangle\ \mathcal{A}\langle R_A, \llbracket \text{UserFn} \rrbracket \rangle] \\ \mapsto &R_A[R_S[\mathcal{A}\langle R_A, \llbracket \text{UserFn} \rrbracket \rangle\ \mathcal{A}\langle R_S, \llbracket \text{SystemFn} \rrbracket \rangle]] \\ \mapsto &R_A[R_S[R_A[\mathcal{A}\langle R_S, \llbracket \text{SystemFn} \rrbracket \rangle\ \mathcal{A}\langle R_A, \llbracket \text{UserFn} \rrbracket \rangle]]] \\ \mapsto &R_A[R_S[R_A[R_S[\mathcal{A}\langle R_A, \llbracket \text{UserFn} \rrbracket \rangle\ \mathcal{A}\langle R_S, \llbracket \text{SystemFn} \rrbracket \rangle]]]] \\ &\dots \end{aligned}$$

Reduction (w/o Annotations)

$$\begin{aligned} &\text{UserFn SystemFn} \\ \mapsto &\text{SystemFn UserFn} \\ \mapsto &\text{UserFn SystemFn} \\ \mapsto &\text{SystemFn UserFn} \\ \mapsto &\text{UserFn SystemFn} \\ &\dots \end{aligned}$$

This program consists of two copies of a mutually recursive loop function, one a ‘user’ component and one a ‘system’ component. Each takes the other as an argument, and then calls it, passing itself as the sole argument. To simplify the presentation of the looping functions, we introduce abbreviations for the user and system procedures.

This program is a toy example, but it represents the core of many interactions between user and system code. For instance, any co-routine-style inter-

action between producer and consumer exhibits this behavior—unfortunately, programmers are forced to avoid this powerful and natural style in Java precisely because of the lack of tail-calling. Perhaps the most common examples of this kind of interaction occur in OO-style traversals of data structures, such as the above-mentioned patterns.

The first reduction sequence in the figure shows how λ_{SEC} evaluates the given program, where the two procedures are annotated with their permissions. The context quickly grows without bound in this example. A functional programmer would expect to see a sequence more like the second one. This series is also a reduction sequence in λ_{SEC} , but one that is obtained by evaluating the program’s pure source, without the security annotations.

As Fournet and Gordon point out in their paper, all is not lost. They introduce an additional reduction into their abstract machine that explicitly removes a frame before performing a call. Unfortunately, as they point out, indiscriminate application of this rule changes the semantics of the language. They address this problem with a partial list of circumstances in which the reduction is legal. By casting tail-calling as a specific reduction rather than a property of an abstract machine, Fournet and Gordon fail to realize that a fully tail-recursive implementation of the language is possible. This outlook is similar to the one taken by Microsoft’s CLR, in which a non-allocating tail call requires a dynamic check to ensure that no security information is lost.

4.4 An Abstract Machine for λ_{SEC}

Following Clinger’s work on defining tail-optimized languages via space complexity classes [11], we first reformulate the λ_{SEC} semantics as a CESK machine [16, 15]. We can then measure the space consumed by machine configurations, programs, and machines. Furthermore, we can determine whether the space consumption function of an implementation is in the same complexity class as Clinger’s machine.

The fg machine

We begin with a direct translation of λ_{SEC} ’s semantics into a CESK machine, which we call “frame-generating” or **fg** (see figure 4.1).

A CESK abstract machine takes its names from its four registers: the control string, the environment, the store, and the continuation. It is quite similar to the CEKD machine of chapter 3. It is missing the D register—a λ_{SEC} program does not contain top-level definitions—and it has gained a store, S .⁴

The derivation of a CESK machine from a reduction semantics is straightforward [15]. In particular, the proof of equivalence of the two models is a refinement of Felleisen and Flatt’s proof, which proceeds by a series of transformations from a simple reduction semantics to a register machine. At each

⁴The store in our model is necessitated by Clinger’s model of tail call optimization; a machine with no store can grow without bound due to copying.

The FG Machine

Configurations	= $\langle M, \rho, \sigma, \kappa \rangle \mid \langle V, \rho, \sigma, \kappa \rangle \mid \langle V, \sigma \rangle \mid \text{fail}$
Final Configurations	= $\langle V, \sigma \rangle \mid \text{fail}$
$\kappa \in$ Continuations	= $\langle \rangle \mid \langle \text{push} : M, \rho, \kappa \rangle \mid \langle \text{call} : V, \kappa \rangle$ $\mid \langle \text{frame} : R, \kappa \rangle \mid \langle \text{grant} : R, \kappa \rangle$
$V \in$ Values	= $\langle \text{closure} : M, \rho \rangle$
$\rho \in$ Environments	= Identifiers \rightarrow Locations
α, β	\in Locations
$\sigma \in$ Stores	= Locations \rightarrow Values
empty_{fg}	= $\langle \rangle$
$\langle \lambda_f x.M, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{fg}} \langle \langle \text{closure} : \lambda_f x.M, \rho \rangle, \rho, \sigma, \kappa \rangle$
$\langle x, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{fg}} \langle \sigma(\rho(x)), \rho, \sigma, \kappa \rangle$
$\langle M N, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{fg}} \langle M, \rho, \sigma, \langle \text{push} : N, \rho, \kappa \rangle \rangle$
$\langle R[M], \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{fg}} \langle M, \rho, \sigma, \langle \text{frame} : R, \kappa \rangle \rangle$
$\langle \text{grant } R \text{ in } M, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{fg}} \langle M, \rho, \sigma, \langle \text{grant} : R, \kappa \rangle \rangle$
$\langle \text{test } R \text{ then } M \text{ else } N, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{fg}} \begin{cases} \langle M, \rho, \sigma, \kappa \rangle & \text{if } \mathcal{OK}_{\text{fg}} \langle R, [\kappa] \rangle \\ \langle N, \rho, \sigma, \kappa \rangle & \text{otherwise} \end{cases}$
$\langle \text{fail}, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{fg}} \text{fail}$
$\langle V, \rho, \sigma, \langle \rangle \rangle$	$\mapsto_{\text{fg}} \langle V, \sigma \rangle$
$\langle V, \rho, \sigma, \langle \text{push} : M, \rho', \kappa \rangle \rangle$	$\mapsto_{\text{fg}} \langle M, \rho', \sigma, \langle \text{call} : V, \kappa \rangle \rangle$
$\langle V, \rho, \sigma, \langle \text{call} : V', \kappa \rangle \rangle$	$\mapsto_{\text{fg}} \langle M, \rho' [f \mapsto \beta] [x \mapsto \alpha], \sigma [\alpha \mapsto V] [\beta \mapsto V'], \kappa \rangle$ if $V' = \langle \text{closure} : \lambda_f x.M, \rho' \rangle$ and $\alpha, \beta \notin \text{dom}(\sigma)$
$\langle V, \rho, \sigma, \langle \text{frame} : R, \kappa \rangle \rangle$	$\mapsto_{\text{fg}} \langle V, \rho, \sigma, \kappa \rangle$
$\langle V, \rho, \sigma, \langle \text{grant} : R, \kappa \rangle \rangle$	$\mapsto_{\text{fg}} \langle V, \rho, \sigma, \kappa \rangle$
$\langle V, \rho, \sigma [\beta, \dots \mapsto V', \dots], \kappa \rangle$	$\mapsto_{\text{fg}} \langle V, \rho, \sigma, \kappa \rangle$ if $\{\beta, \dots\}$ is nonempty and β, \dots do not occur in V, ρ, σ , or κ

where

$\mathcal{OK}_{\text{fg}} \langle \emptyset, [\kappa] \rangle$	
$\mathcal{OK}_{\text{fg}} \langle R, [\langle \rangle] \rangle$	
$\mathcal{OK}_{\text{fg}} \langle R, [\langle \text{push} : M, \rho, \kappa \rangle] \rangle$	iff $\mathcal{OK}_{\text{fg}} \langle R, [\kappa] \rangle$
$\mathcal{OK}_{\text{fg}} \langle R, [\langle \text{call} : V, \kappa \rangle] \rangle$	iff $\mathcal{OK}_{\text{fg}} \langle R, [\kappa] \rangle$
$\mathcal{OK}_{\text{fg}} \langle R, [\langle \text{frame} : R', \kappa \rangle] \rangle$	iff $R \subseteq R'$ and $\mathcal{OK}_{\text{fg}} \langle R, [\kappa] \rangle$
$\mathcal{OK}_{\text{fg}} \langle R, [\langle \text{grant} : R', \kappa \rangle] \rangle$	iff $\mathcal{OK}_{\text{fg}} \langle R - R', [\kappa] \rangle$

Figure 4.1: The FG Machine

step, we must strengthen the induction hypothesis by adding a claim about the value of the \mathcal{OK} predicate when applied to the current context.

As a result, most of the steps that can be taken in such a machine correspond either to the reductions of the source semantics or to the mechanical identification of the next expression to be reduced. The first group of reductions in figure 4.1 contains those that refocus the evaluation on subexpressions or translate expressions into their corresponding values. The second, complementary group contains those that fire when a value shows up as the control string, and these correspond both to changes of focus in the control string and to actual reductions. Finally, a machine with a store must also model garbage collection, if its configurations are to be used in space computations. The final reduction therefore provides garbage collection.

To enable comparisons between different machines, the Eval_x function is abstracted over a transition relation and an empty context. Applying this generalized evaluator to \mapsto_{fg} and empty_{fg} yields the evaluation function Eval_{fg} .

In order to ensure that Eval and Eval_{fg} are indeed the same function, the Eval_x function must employ “load” and “unload” functions. The “load” function, \mathcal{L} , coerces the target program to a valid machine configuration. The “unload” function, \mathcal{U} , recursively substitutes values bound in the environment for the variables that represent them.

Definition 2 (Eval_x)

$$\text{Eval}_x(C, \dots) = \mathcal{U}(V, \sigma) \text{ if } \mathcal{L}_x(C, \dots) \mapsto_x^* \langle V, \sigma \rangle$$

where

$$\mathcal{L}_x(\langle \lambda_f x. M_{u0}, R_0 \rangle, \dots) = \langle \langle \mathcal{A}(R_0, \llbracket \lambda_f x. M_{u0} \rrbracket) \dots \rangle, \emptyset, \emptyset, \text{empty}_x \rangle$$

and

$$\mathcal{U}(\langle \text{closure} : M, \{ \langle x_1, \alpha_1 \rangle, \dots, \langle x_n, \alpha_n \rangle \} \rangle, \sigma) = [\mathcal{U}(\sigma(\alpha_1))/x_1] \dots [\mathcal{U}(\sigma(\alpha_n))/x_n] M$$

Theorem 4 (Machine Fidelity) For all (C, \dots) ,

$$\text{Eval}_{\text{fg}}(C, \dots) = V \text{ iff } \text{Eval}(C, \dots) = V$$

The proof proceeds by induction on the length of a reduction sequence.

The fg machine is not tail-calling

To see that this implementation of the λ_{SEC} language is not tail-calling, we show the reduction sequence in the fg machine for the program from section 4.3, and verify that the space taken by the configuration is growing without bound.

$$\begin{aligned} \text{UserClo} &\triangleq \langle \text{closure} : \lambda_{\text{user}} \text{sys}. \mathcal{A}(R_A, \llbracket \text{UserFn} \rrbracket), \emptyset \rangle \\ \text{SystemClo} &\triangleq \langle \text{closure} : \lambda_{\text{sys}} \text{user}. \mathcal{A}(R_S, \llbracket \text{SystemFn} \rrbracket), \emptyset \rangle \\ \rho_0 &\triangleq [\text{sys} \mapsto \alpha, \text{user} \mapsto \beta] \\ \sigma_0 &\triangleq [\alpha \mapsto \text{SystemClo}, \beta \mapsto \text{UserClo}] \end{aligned}$$

$$\begin{aligned}
& \langle \mathcal{A}(R_A, \llbracket \text{UserFn} \rrbracket) \mathcal{A}(R_S, \llbracket \text{SystemFn} \rrbracket), \emptyset, \emptyset, \langle \rangle \rangle && (0 \text{ frames}) \\
& \mapsto_{\text{fg}} \langle \mathcal{A}(R_A, \llbracket \text{UserFn} \rrbracket), \emptyset, \emptyset, \langle \text{push} : \mathcal{A}(R_S, \llbracket \text{SystemFn} \rrbracket), \emptyset, \langle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{UserClo}, \emptyset, \emptyset, \langle \text{push} : \mathcal{A}(R_S, \llbracket \text{SystemFn} \rrbracket), \emptyset, \langle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \mathcal{A}(R_S, \llbracket \text{SystemFn} \rrbracket), \emptyset, \emptyset, \langle \text{call} : \text{UserClo}, \langle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{SystemClo}, \emptyset, \emptyset, \langle \text{call} : \text{UserClo}, \langle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle R_A[\text{sys user}], \rho_0, \sigma_0, \langle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{sys user}, \rho_0, \sigma_0, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle && (1 \text{ frame}) \\
& \mapsto_{\text{fg}} \langle \text{sys}, \rho_0, \sigma_0, \langle \text{push} : \text{user}, \rho_0, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{SystemClo}, \rho_0, \sigma_0, \langle \text{push} : \text{user}, \rho_0, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{user}, \rho_0, \sigma_0, \langle \text{call} : \text{SystemClo}, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{UserClo}, \rho_0, \sigma_0, \langle \text{call} : \text{SystemClo}, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \\
& \xrightarrow{2}_{\text{fg}} \langle R_S[\text{user sys}], \rho_0, \sigma_0, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{user sys}, \rho_0, \sigma_0, \langle \text{frame} : R_S, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle && (2 \text{ frames}) \\
& \mapsto_{\text{fg}} \langle \text{user}, \rho_0, \sigma_0, \langle \text{push} : \text{sys}, \rho_0, \langle \text{frame} : R_S, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{UserClo}, \rho_0, \sigma_0, \langle \text{push} : \text{sys}, \rho_0, \langle \text{frame} : R_S, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{sys}, \rho_0, \sigma_0, \langle \text{call} : \text{UserClo}, \langle \text{frame} : R_S, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \rangle \\
& \mapsto_{\text{fg}} \langle \text{SystemClo}, \rho_0, \sigma_0, \langle \text{call} : \text{UserClo}, \langle \text{frame} : R_S, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \rangle \\
& \xrightarrow{7}_{\text{fg}} \langle \text{UserClo}, \rho_0, \sigma_0, && (3 \text{ frames}) \\
& \quad \langle \text{call} : \text{SystemClo}, \langle \text{frame} : R_A, \langle \text{frame} : R_S, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \rangle \rangle \rangle \\
& \xrightarrow{7}_{\text{fg}} \langle \text{SystemClo}, \rho_0, \sigma_0, && (4 \text{ frames}) \\
& \quad \langle \text{call} : \text{UserClo}, \langle \text{frame} : R_S, \langle \text{frame} : R_A, \langle \text{frame} : R_S, \langle \text{frame} : R_A, \langle \rangle \rangle \rangle \rangle \rangle \rangle \rangle \rangle \\
& \dots
\end{aligned}$$

4.5 An Alternative Implementation

How security inspections really work

A close look at λ_{sec} shows that frame ($\ell[\bullet]$) and **grant** contexts affect the computation only when they are observed by a test expression. That is, a program with no test expressions may be simplified by removing all frame and **grant** expressions without changing its meaning. Furthermore, the observations possible with the test expression are limited by the \mathcal{OK} function.

In particular, any sequence of frame and **grant** expressions may be collapsed into a canonical table that provides a partial map from the set of permissions to one of two conditions: ‘no’, indicating that the permission is not granted by the sequence, and ‘grant’, indicating that the permission is granted (and legally so) by some **grant** frame in the sequence.

To derive update rules for this table, we consider evaluation of the \mathcal{OK} function as the recognition of a context-free grammar over the alphabet of frame and **grant** expressions. We start by simplifying the model to one with a single permission. Then each frame is either empty or contains the desired permission. Likewise, there is only one possible grant. All other continuation frames are irrelevant. So a full evaluation context can be seen as an arbitrary string in the alphabet $\Sigma = \{y, n, g\}$, where y and n represent frames that contain or are missing the given permission, and g represents a **grant**. Assume

the ordering of the letters in the word places the outermost frames at the left end of the string.

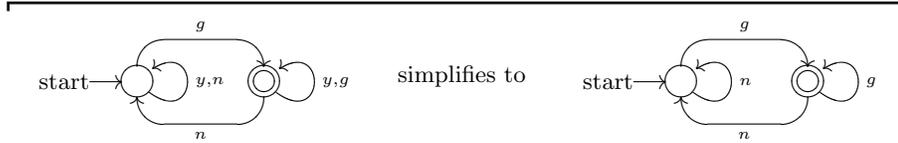


Figure 4.2: Computing permissions as a finite state automaton

With the grammar in place, the \mathcal{OK}_{fg} predicate can easily be interpreted as a finite-state machine that recognizes the regular expression Σ^*gy^* ; that is, a string ending with a grant followed by any number of y 's. The resulting FSA has just two states, one accepting and one non-accepting. A g always transitions to the accepting state, and a n always transitions to the non-accepting state. A y causes a (trivial) transition to the current state.

This last observation leads us to a further simplification of the grammar. Since the presence of the character y does not affect the decision of the FSA, we may ignore the continuation frames that generate them, and consider only the **grant** frames and those security frames that do not include the desired permission. The regular expression indicating the success of \mathcal{OK}_{fg} becomes simply Σ^*g . Figure 4.2 shows these two finite-state automata graphically.

This simplification leads to an insight about the security model of λ_{sec} . In the automaton that λ_{sec} induces, the y may be ignored. In the security model, then, callers with a given permission do not affect the result of a check for that permission. Rather, it is the callers *without* that permission that might change its status, and **grants** of that permission. This suggests that what the security model really tracks is the *absence* of certain permissions. At runtime, then, it is the complement of the permissions attributed to a caller that matters.

Applying the simplified grammar to our reduction semantics allows us to collapse uninterrupted sequences of frame and **grant** expressions that occur in the evaluation context. A substring ending in a g results in an accepting state, a substring ending in an n results in a non-accepting state, and the empty substring does not alter the decision. To extend this to the whole language, we must expand our single-permission state to a full table of permissions.

The cm machine

In the **cm** (continuation-marks) machine, each continuation frame contains a table of permissions, called a *markset*. The evaluation steps for frame and **grant** expressions update the table in the enclosing continuation, rather than increasing the length of the continuation itself. The \mathcal{OK}_{cm} predicate now inspects these marksets, rather than the frame and grant elements of the continuation. Otherwise, the **cm** machine is the same as the **fg** machine.

The CM Machine

$m \in \text{marksets}$	$= \mathcal{P} \rightarrow \{\text{grant}, \text{no}\}$
Configurations	$= \langle M, \rho, \sigma, \kappa \rangle \mid \langle V, \rho, \sigma, \kappa \rangle \mid \langle V, \sigma \rangle \mid \text{fail}$
Final Configurations	$= \langle V, \sigma \rangle \mid \text{fail}$
$\kappa \in \text{continuations}$	$= \langle \text{empty} : m \rangle \mid \langle \text{push} : M, \rho, \kappa, m \rangle \mid \langle \text{call} : V, \kappa, m \rangle$
$V \in \text{Values}$	$= \langle \text{closure} : M, \rho \rangle$
$\rho \in \text{Environments}$	$= \text{Identifiers} \rightarrow \text{Locations}$
$\alpha, \beta \in \text{Locations}$	$\in \text{Locations}$
$\sigma \in \text{Stores}$	$= \text{Locations} \rightarrow \text{Values}$
empty_{cm}	$= \langle \text{empty} : \emptyset \rangle$
$\langle \lambda f x.M, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{cm}} \langle \langle \text{closure} : \lambda f x.M, \rho \rangle, \rho, \sigma, \kappa \rangle$
$\langle x, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{cm}} \langle \sigma(\rho(x)), \rho, \sigma, \kappa \rangle$
$\langle M N, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{cm}} \langle M, \rho, \sigma, \langle \text{push} : N, \rho, \kappa, \emptyset \rangle \rangle$
$\langle R[M], \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{cm}} \langle M, \rho, \sigma, \kappa[\overline{R} \mapsto \text{no}] \rangle$
$\langle \text{grant } R \text{ in } M, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{cm}} \langle M, \rho, \sigma, \kappa[R \mapsto \text{grant}] \rangle$
$\langle \text{test } R \text{ then } M \text{ else } N, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{cm}} \begin{cases} \langle M, \rho, \sigma, \kappa \rangle & \text{if } \mathcal{OK}_{\text{cm}}\langle R, \llbracket \kappa \rrbracket \rangle \\ \langle N, \rho, \sigma, \kappa \rangle & \text{otherwise} \end{cases}$
$\langle \text{fail}, \rho, \sigma, \kappa \rangle$	$\mapsto_{\text{cm}} \text{fail}$
$\langle V, \rho, \sigma, \langle \text{empty} : m \rangle \rangle$	$\mapsto_{\text{cm}} \langle V, \sigma \rangle$
$\langle V, \rho, \sigma, \langle \text{push} : M, \rho', \kappa, m \rangle \rangle$	$\mapsto_{\text{cm}} \langle M, \rho', \sigma, \langle \text{call} : V, \kappa, \emptyset \rangle \rangle$
$\langle V, \rho, \sigma, \langle \text{call} : V', \kappa, m \rangle \rangle$	$\mapsto_{\text{cm}} \langle M, \rho'[f \mapsto \beta][x \mapsto \alpha], \sigma[\alpha \mapsto V][\beta \mapsto V'], \kappa \rangle$ if $V' = \langle \text{closure} : \lambda f x.M, \rho' \rangle$ and $\alpha, \beta \notin \text{dom}(\sigma)$
$\langle V, \rho, \sigma[\beta, \dots \mapsto V, \dots], \kappa \rangle$	$\mapsto_{\text{cm}} \langle V, \rho, \sigma, \kappa \rangle$ if $\{\beta, \dots\}$ is nonempty and β, \dots do not occur in V, ρ, σ , or κ

where

$$\langle \dots, m \rangle[R \mapsto c] = \langle \dots, m[R \mapsto c] \rangle \text{ (pointwise extension)}$$

and

$$\left. \begin{array}{l} \mathcal{OK}_{\text{cm}}\langle \emptyset, \llbracket \kappa \rrbracket \rangle \\ \mathcal{OK}_{\text{cm}}\langle R, \llbracket \langle \text{empty} : m \rangle \rrbracket \rangle \\ \mathcal{OK}_{\text{cm}}\langle R, \llbracket \langle \text{push} : M, \rho, \kappa, m \rangle \rrbracket \rangle \\ \mathcal{OK}_{\text{cm}}\langle R, \llbracket \langle \text{call} : V, \kappa, m \rangle \rrbracket \rangle \end{array} \right\} \begin{array}{l} \text{iff } (R \cap m^{-1}(\text{no}) = \emptyset) \\ \text{iff } (R \cap m^{-1}(\text{no}) = \emptyset) \\ \text{and } \mathcal{OK}_{\text{cm}}\langle R - m^{-1}(\text{grant}), \llbracket \kappa \rrbracket \rangle \end{array}$$

Figure 4.3: The CM Machine

Figure 4.3 shows the definition of the cm machine. Note that the framing operation takes the complement of the set R , in accordance with the insight of the prior section. We assume that the set of permissions is finite. Also, the markset mappings are extended pointwise across sets of permissions; that is, $m[R \rightarrow c](p) = c$ if $p \in R$, and $m(p)$ otherwise.

The Eval_{cm} function is another instance of Eval_x . That is, Eval_{cm} is the same as Eval_{fg} , except that it uses \mapsto_{cm} as its transition function and empty_{cm} as its empty continuation.

The two machines produce the same results.

Theorem 5 (Machine Equivalence) *For all (C, \dots) ,*

$$\text{Eval}_{\text{fg}}(C, \dots) = V \text{ iff } \text{Eval}_{\text{cm}}(C, \dots) = V$$

To prove this theorem, we must show that if the fg machine terminates, the cm machine terminates with the same value, and that if the fg machine does not terminate in a final state, then the cm machine also fails to terminate.

For the purposes of the proof, we assume that no garbage collection steps are taken, because garbage collection cannot affect the result of the evaluation.

Lemma 8 (No Garbage Collection) *For every evaluation sequence in either the fg or cm machine, removing every garbage-collection step produces another legal sequence, and no divergent computation is made finite by such a removal.*

$$\mathcal{T} : C_{\text{fg}} \rightarrow C_{\text{cm}}$$

$$\begin{aligned} \mathcal{T}\langle M, \rho, \sigma, \kappa \rangle &= \langle M, \rho, \sigma, \mathcal{T}(\kappa) \rangle \\ \mathcal{T}\langle V, \rho, \sigma, \kappa \rangle &= \langle V, \rho, \sigma, \mathcal{T}(\kappa) \rangle \\ \mathcal{T}\langle V, \sigma \rangle &= \langle V, \sigma \rangle \\ \mathcal{T}\langle \text{fail} \rangle &= \text{fail} \\ \mathcal{T}\langle \rangle &= \langle \text{empty} : \emptyset \rangle \\ \mathcal{T}\langle \text{push} : M, \rho, \kappa \rangle &= \langle \text{push} : M, \rho, \mathcal{T}(\kappa), \emptyset \rangle \\ \mathcal{T}\langle \text{call} : V, \kappa \rangle &= \langle \text{call} : V, \mathcal{T}(\kappa), \emptyset \rangle \\ \mathcal{T}\langle \text{frame} : R, \kappa \rangle &= \mathcal{T}(\kappa)[\bar{R} \mapsto \text{no}] \\ \mathcal{T}\langle \text{grant} : R, \kappa \rangle &= \mathcal{T}(\kappa)[R \mapsto \text{grant}] \end{aligned}$$

To compare the machines, we introduce the function \mathcal{T} . The function \mathcal{T} maps configurations of the fg machine to configurations of the cm machine. A step in the fg machine corresponds to either no steps or one step in the cm machine.

Lemma 9 (Simulation) *Given a configuration C_{cm} , with $C_{\text{cm}} = \mathcal{T}(C_{\text{fg}})$, one of the following holds:*

1. C_{fg} is either *fail* or $\langle V, \sigma \rangle$
2. C_{fg} and C_{cm} are both *stuck*.

3. $C_{fg} \mapsto_{fg} C'_{fg}$ and $\mathcal{T}(C'_{fg}) = C_{cm}$
4. $C_{fg} \mapsto_{fg} C'_{fg}$ and $C_{cm} \mapsto_{cm} \mathcal{T}(C'_{fg})$

The proof is a case analysis on the four cases and the configurations of the machine. The *fg* machine takes extra steps only when popping frame and *grant* continuations after reducing their arguments to values.

The *cm* machine can always represent a sequence of frame and *grant* expressions with a single markset. The sequence of steps below illustrates this for the divergent mutually-recursive computation shown in section 4.3.

$$\begin{aligned}
 R_S &\triangleq \{b, c\} \\
 R_A &\triangleq \{a, b\}
 \end{aligned}$$

$$\begin{aligned}
 &\langle \mathcal{A}(R_A, \llbracket \text{UserFn} \rrbracket) \mathcal{A}(R_S, \llbracket \text{SystemFn} \rrbracket), \emptyset, \emptyset, \langle \text{empty} : \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle \mathcal{A}(R_A, \llbracket \text{UserFn} \rrbracket), \emptyset, \emptyset, \langle \text{push} : \mathcal{A}(R_S, \llbracket \text{SystemFn} \rrbracket), \emptyset, \langle \text{empty} : \emptyset \rangle, \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle \text{UserClo}, \emptyset, \emptyset, \langle \text{push} : \mathcal{A}(R_S, \llbracket \text{SystemFn} \rrbracket), \emptyset, \langle \text{empty} : \emptyset \rangle, \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle \mathcal{A}(R_S, \llbracket \text{SystemFn} \rrbracket), \emptyset, \emptyset, \langle \text{call} : \text{UserClo}, \langle \text{empty} : \emptyset \rangle, \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle \text{SystemClo}, \emptyset, \emptyset, \langle \text{call} : \text{UserClo}, \langle \text{empty} : \emptyset \rangle, \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle R_A[\text{sys user}], \rho_0, \sigma_0, \langle \text{empty} : \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle \text{sys user}, \rho_0, \sigma_0, \langle \text{empty} : [\{c\} \mapsto \text{no}] \rangle \rangle \\
 \mapsto_{cm} &\langle \text{sys}, \rho_0, \sigma_0, \langle \text{push} : \text{user}, \rho_0, \langle \text{empty} : [\{c\} \mapsto \text{no}] \rangle \rangle \rangle \\
 \mapsto_{cm} &\langle \text{SystemClo}, \rho_0, \sigma_0, \langle \text{push} : \text{user}, \rho_0, \langle \text{empty} : [\{c\} \mapsto \text{no}] \rangle, \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle \text{user}, \rho_0, \sigma_0, \langle \text{call} : \text{SystemClo}, \langle \text{empty} : [\{c\} \mapsto \text{no}] \rangle, \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle \text{UserClo}, \rho_0, \sigma_0, \langle \text{call} : \text{SystemClo}, \langle \text{empty} : [\{c\} \mapsto \text{no}] \rangle, \emptyset \rangle \rangle \\
 \xrightarrow{2}_{cm} &\langle R_S[\text{user sys}], \rho_0, \sigma_0, \langle \text{empty} : [\{c\} \mapsto \text{no}] \rangle \rangle \\
 \mapsto_{cm} &\langle \text{user sys}, \rho_0, \sigma_0, \langle \text{empty} : [\{a, c\} \mapsto \text{no}] \rangle \rangle \\
 \mapsto_{cm} &\langle \text{user}, \rho_0, \sigma_0, \langle \text{push} : \text{sys}, \rho_0, \langle \text{empty} : [\{a, c\} \mapsto \text{no}] \rangle \rangle \rangle \\
 \mapsto_{cm} &\langle \text{UserClo}, \rho_0, \sigma_0, \langle \text{push} : \text{sys}, \rho_0, \langle \text{empty} : [\{a, c\} \mapsto \text{no}] \rangle, \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle \text{sys}, \rho_0, \sigma_0, \langle \text{call} : \text{UserClo}, \langle \text{empty} : [\{a, c\} \mapsto \text{no}] \rangle, \emptyset \rangle \rangle \\
 \mapsto_{cm} &\langle \text{SystemClo}, \rho_0, \sigma_0, \langle \text{call} : \text{UserClo}, \langle \text{empty} : [\{a, c\} \mapsto \text{no}] \rangle, \emptyset \rangle \rangle \\
 \xrightarrow{7}_{cm} &\langle \text{UserClo}, \rho_0, \sigma_0, \langle \text{call} : \text{SystemClo}, \langle \text{empty} : [\{a, c\} \mapsto \text{no}] \rangle, \emptyset \rangle \rangle \\
 \xrightarrow{7}_{cm} &\langle \text{SystemClo}, \rho_0, \sigma_0, \langle \text{call} : \text{UserClo}, \langle \text{empty} : [\{a, c\} \mapsto \text{no}] \rangle, \emptyset \rangle \rangle \\
 &\dots
 \end{aligned}$$

4.6 Space Consumption

In “Proper Tail Recursion and Space Efficiency,” Clinger [11] describes a framework that characterizes the memory behavior of a language implementation as a mapping from programs to the maximum memory that the implementation consumes while evaluating that program. He demonstrates the difference between various named classes of implementation (“tail-recursive,” “safe-for-space,” etc.), and defines asymptotic space complexity classes for each, based on abstract machine definitions.

In order to apply Clinger’s analytic framework of tail recursion to the *fg* and *cm* machines, we must define a memory measure that maps a machine

configuration to a real number. The measure for the **fg** machine is straightforward.

$$\begin{aligned}
\text{space}(\text{fail}) &= 1 \\
\text{space}(\langle V, \sigma \rangle) &= \text{space}(V) + \text{space}(\sigma) \\
\text{space}(\langle M, \rho, \sigma, \kappa \rangle) &= |\text{dom}(\rho)| + \text{space}(\kappa) + \text{space}(\sigma) \\
\text{space}(\langle V, \rho, \sigma, \kappa \rangle) &= \text{space}(V) + |\text{dom}(\rho)| + \text{space}(\kappa) + \\
&\quad \text{space}(\sigma) \\
\text{space}(\langle \text{closure} : \lambda_f x. M, \rho \rangle) &= 1 + |\text{dom}(\rho)| \\
\text{space}(\sigma) &= \sum_{\alpha \in \text{dom}(\sigma)} 1 + \text{space}(\sigma(\alpha)) \\
\text{space}(\langle \rangle) &= 1 \\
\text{space}(\langle \text{push} : M, \rho, \kappa \rangle) &= 1 + |\text{dom}(\rho)| + \text{space}(\kappa) \\
\text{space}(\langle \text{call} : V, \kappa \rangle) &= 1 + \text{space}(V) + \text{space}(\kappa) \\
\text{space}(\langle \text{frame} : R, \kappa \rangle) &= 1 + |R| + \text{space}(\kappa) \\
\text{space}(\langle \text{grant} : R, \kappa \rangle) &= 1 + |R| + \text{space}(\kappa)
\end{aligned}$$

To accommodate the **cm** machine, we extend this function with rules for the size of a markset, and for the size of continuations that contain a markset.

$$\begin{aligned}
\text{space}(\langle \text{empty} : m \rangle) &= 1 + \text{space}(m) \\
\text{space}(\langle \text{push} : M, \rho, \kappa, m \rangle) &= 1 + \text{space}(\rho) + \text{space}(\kappa) + |\text{dom}(m)| \\
\text{space}(\langle \text{call} : V, \kappa, m \rangle) &= 1 + \text{space}(V) + \text{space}(\kappa) + |\text{dom}(m)|
\end{aligned}$$

The space functions \mathcal{S}_{fg} and \mathcal{S}_{cm} are defined as the maximum amount of memory consumed during the evaluation of a program. In order to ensure that unreachable store values do not affect the space function, Clinger defines a “space-efficient” computation as a sequence of steps where the garbage-collection rule is applied as often as possible.

Definition 3 (Space-Efficient Computations) *A space-efficient computation in an implementation x is a finite or countably infinite sequence of configurations $\{C_i\}$ such that*

- *If C_i and C_{i+1} are in the sequence, then $C_i \mapsto_x C_{i+1}$.*
- *If the sequence is finite, then it ends with a final configuration.*
- *If the garbage collection rule is applicable to C_i , then $C_i \mapsto_x C_{i+1}$ by the garbage collection rule.*

Definition 4 (Supremum) *For $R \subseteq \mathfrak{R}$, the supremum $\text{Sup}(R)$ is the least upper bound of R , or ∞ if no such bound exists.*

Definition 5 (Space Consumption S_x) *The space consumption function of an implementation x is $S_x : \text{Program} \rightarrow \mathfrak{R} \cup \{\infty\}$ defined by*

$$S_x(P) = |P| + \sup\{\sup\{\text{space}(\{C_i\})\} \mid \{C_i\} \text{ is a space-efficient computation in } x, \text{ with } C_0 = \mathcal{L}_x(P)\}$$

where $|P|$ is the number of nodes in the abstract syntax tree of P .

Note that the outer ‘supremum’ accommodates the possibility of an implementation that is observationally nondeterministic. This definition is therefore sufficient but overly general.

Following Clinger, we extend the notion of a space function to one of asymptotic space complexity.

Definition 6 (Asymptotic Complexity, $O(f)$) *If A is any set, and $f : A \rightarrow \mathfrak{R} \cup \{\infty\}$, then the asymptotic (upper bound) complexity class of f is $O(f)$, which is defined as the set of all functions $g : A \rightarrow \mathfrak{R} \cup \{\infty\}$ for which there exist real constants c_1 and c_0 such that $c_1 > 0$ and*

$$\forall a \in A. g(a) \leq c_1 f(a) + c_0$$

We can now prove that the asymptotic space consumption of the **fg** machine is strictly greater than that of the **cm** machine. Put differently, the class of implementations for λ_{sec} in $O(S_{\text{cm}})$ is strictly smaller than those in $O(S_{\text{fg}})$.

Theorem 6 (Space Comparison)

$$O(S_{\text{cm}}) \subsetneq O(S_{\text{fg}})$$

Proof Sketch: The proof has two parts. First, we must show that every function in $O(S_{\text{cm}})$ is also in $O(S_{\text{fg}})$. Second, we must show there is a function in $O(S_{\text{fg}})$ that is not in $O(S_{\text{cm}})$.

Since the set $O(S_{\text{cm}})$ takes S_{cm} as its asymptotic upper bound, it suffices for the first half of the proof to show that S_{cm} is in the set $O(S_{\text{fg}})$. That is, for all programs, the maximum space taken while evaluating the program in the **cm** machine is less than or equal to some constant times the space that the **fg** machine takes. For simplicity, we shall choose a loose upper bound, taking as our constant the size of the set of permissions \mathcal{P} .

The translation \mathcal{T} removes frame and grant continuation frames and introduces marksets on each remaining frame, leaving all else untouched. We can therefore show that our bound is satisfied by considering the worst case, in which the **fg** machine contains no frame or grant continuations, and each markset in the **cm** machine contains an entry for every possible permission. In

this case, each frame is increased in size by the size of \mathcal{P} . Since each frame in the **fg** machine is at least of size 1, we may take $|\mathcal{P}|$ as our linear factor to satisfy the bound.

For the second half of the proof, it suffices to show a program whose machine configuration grows without bound in the **fg** machine and does not in the **cm** machine. The example given earlier satisfies this requirement. \square

To prove that our implementation is tail-recursive by Clinger's definition, we must extend his language to include the new forms that appear in λ_{sec} . In order to produce the most stringent possible requirement on space consumption, we propose an implementation that includes a security oracle. That is, we compare ourselves to an implementation that consumes no space at all for the maintenance of stack-trace information. Because of the similarity between the **fg** machine and Clinger's, the easiest way to model this is to consider the space measure obtained by eliminating the size of the markset from the calculation. We therefore define the oracular space measure "space_o," which differs from the existing function only in its rules for continuations.

$$\begin{array}{l} \text{space}_o(\langle \text{empty} : m \rangle) = 1 \\ \text{space}_o(\langle \text{push} : M, \rho, \kappa, m \rangle) = 1 + \text{space}(\rho) + \text{space}_o(\kappa) \\ \text{space}_o(\langle \text{call} : V, \kappa, m \rangle) = 1 + \text{space}(V) + \text{space}_o(\kappa) \end{array}$$

Combining this space measure with the existing **cm** machine, we define the space class S_o of implementations that consume no memory at all for security information. We can now prove that the **cm** machine is tail-recursive.

Theorem 7 (Tail Recursion)

$$O(S_{\text{cm}}) = O(S_o)$$

Proof Sketch To show set equality of asymptotic measures, it suffices to show that $S_{\text{cm}} \in O(S_o)$ and that $S_o \in O(S_{\text{cm}})$.

Since these two space classes use the same machine semantics, we may use the identity function to translate between configurations, and the proof of intensional language equality is trivial.

For the first half of the proof, we must show that the non-oracular measure of a configuration's space ($\text{space}_{\text{cm}}(C)$) is at most k times as large as the space taken by the oracular measurement ($\text{space}_o(C)$). As in the previous proof, we choose as our constant the size of the permission set ($|\mathcal{P}|$). The non-oracular space measure increases the size of a continuation frame by at most $|\mathcal{P}|$, and every continuation frame is of size at least one, so the space taken by a configuration grows by no more than a factor of $|\mathcal{P}|$.

The other direction is much easier, since simple inspection of the rules for the oracular space function reveals that this function is uniformly smaller for a given configuration than the non-oracular space measure. \square

4.7 An implementation using Continuation Marks

The marksets defined by the `cm` machine may be implemented more standard models of continuation marks. Figure 4.2 shows that regarding permissions independently reduces the question of whether a privilege is granted to the simpler one of whether a `grant` or a `no` was more recently seen. Continuation marks can model this behavior simply by using one key for each permission. Following the insight given by that diagram, these marks record only `grant` and `no` states.

To be more precise, we may reformulate the translation from the λ_{sec} language to a language with continuation marks by replacing each frame expression by a nested series of **w-c-m**'s, mapping each of the privileges not granted by the frame expression to `no`. A `grant` expression is replaced by a series of **w-c-m**'s mapping each of the privileges legally granted by the expression to 'grant'. A `test` expression for a given permission translates into a check whether the first element of **(c-c-m k)** for the key k corresponding to the permission is `grant` or `no` (or whether there is any mark corresponding to the permission at all).

4.8 The Richer Models of Java and .Net

The model of Fournet and Gordon is an abstraction of Java's and .Net's security model. Both Java and .Net associate program text with permissions and perform security checks by "walking the stack." Both Java and .Net, however, feature a richer security model than λ_{sec} does.

One difference is that the set of permissions is not fixed at compile or even at load time. In both Java and .Net, every object that subclasses a `Permission` class becomes a permission in the security system.

Also, these systems may permit the mappings from code to permissions to be mutated at runtime. In Java 1.2, however, the default implementation does not allow this. That is, the mapping from code to permissions is performed at class-loading time, and not subsequently altered [23]. We are not aware of such a guarantee in .Net.

These differences, however, do not invalidate the principal invariant that underlies our proposed marriage of security and tail-calling. Specifically, the security information retained by the system can only be observed by a security-checking system call. This means that the representation of this information (and the implementation of the corresponding security-checking primitive) may be changed, as long as these changes preserve the observable behavior of the system.

In addition, all of the systems we have examined share the trait that the frame expressions—or the corresponding constructs in Java or .Net—may be reordered, up to the boundaries established by a `grant` (or .Net's "assert").

The first step in applying the lessons of our work to a given system is to formulate an appropriate model. For languages like Java, Java's JVM, or

.Net’s intermediate language, this model should probably extend λ_{sec} with assignment, objects with subclassing, exceptions, and dynamic loading.

The next step is to formulate the notion of tail-calling in such a system. This is fairly natural for any model that treats function application as the transfer of control to a given source location with an unchanged continuation.

The final step is to extend the system with mechanisms for maintaining security state, and to demonstrate that these additions do not affect the tail-recursive behavior of the model.

Fundamentally, our approach separates the implementation of function calls from the implementation of security primitives. In other words, it decouples the security mechanism from the memory behavior of stack frames. This differs from the existing implementations, which capitalize on unanticipated observations at the machine level and therefore restrict the set of possible language implementations.

We conjecture that the security-policy implementation that results from such an analysis is likely to have much in common with our CESK machines. That is, the security information will be attached to the stack frame of the parent, rather than the stack frame of the child. Garbage-collection-like strategies for the management of such information could be employed to maintain asymptotic memory behavior while delaying run-time costs until the application of a security check.

4.9 Related Work

This chapter is directly inspired by the POPL presentation of a semantics for stack inspection by Fournet and Gordon [22], and by our existing work on continuation marks. The key insight required to apply our earlier work to this area is that continuation marks for security permissions contain negative rather than positive information. Once we understood this, we could derive the rest of the ideas here in a straightforward manner.

Security-Passing Style

Another implementation strategy for stack inspection is due to Wallach et al. [46, 47]. In security-passing style, each procedure accepts an additional argument that represents the security context accumulated thus far.

Essentially, this implementation derives from the observation that the security information computed on some context does not change while that context remains active. Therefore, an implementation can compute the security information (or some representation thereof) once, at each call site, and pass it along during the computation.

As a tool of semantic definition, we believe that security-passing style succeeds. That is, the given transformation, when combined with a semantics for the underlying target language, does specify a meaning for each of the security primitives. However, we find the semantics of Fournet and Gordon to be simpler, as they directly interpret the language of security primitives.

As a tool of implementation, we believe security-passing style leaves something to be desired. In particular, the simple overhead of adding an argument to each call is prohibitive. It is certainly the case that a variety of optimizations may be applied to lower the runtime cost, but in this case we contend that our implementation strategy is a more direct route to a similar optimized machine.

Other Work

Several others [3, 40] have considered the problem of adding tail calls to the JVM, which does support stack inspection. However, none of these specifically addressed stack inspection or security, and all of them made the simplifying assumption that tail-calling was only possible between procedures of the same component; that is, none of them considered calls between code from distinct security domains.

Karjoth [26] presents a semantics for access control in Java 2; his model presents rules for the maintenance of access control information, but leaves the rules for the evaluation of the language itself unspecified. Because he includes rules for matching ‘call’ and ‘return’ expressions, his system cannot be the foundation for a tail-calling implementation.

Erlingsson and Schneider [13] show how to implement the stack inspection primitives with no support from the runtime system. That is, they use an annotation-based approach to transform a program with security primitives into one without them. However, this work winds up simulating the stack on the side, with two unfortunate consequences: exceptions become much more difficult, and tail-calling behavior is destroyed.

Chapter 5

Conclusion

One goal of programming languages research is to explain complex language features in terms of simpler ones. We have done exactly this by introducing continuation marks and by showing how they help to explain tools such as debuggers and features such as stack inspection.

5.1 Results

Continuation marks are useful in a variety of contexts. In chapter 2, we showed how continuation marks form the underpinning of several language features in MzScheme, including exception handlers and parameters. Furthermore, an annotation using continuation marks forms the core of a wide variety of DrScheme tools, including a profiler, an error-tracing exception reporter, a simple debugger, a trace engine, and of course DrScheme's stepper.

In chapter 3, we examined steppers in more detail, formulating claims of correctness for steppers in general and demonstrating that our stepper's model is correct. This chapter also showed how the continuation marks play a key role in decoupling the stepper's definition from the language's semantics. That is, continuation marks allow a stepper to *observe* a program's behavior, rather than defining it.

Chapter 4 examined the possible role of continuation marks in stack inspection. In particular, we showed how continuation marks may be used to endow Fournet and Gordon's stack inspection model with tail-calling behavior, and we suggested a related plan to make tail-calling possible for richer models of stack-inspection behavior.

Finally, appendix C highlights the similarities between aspect-oriented programming and continuation marks. Both of them are designed to allow programs to alter their behavior based on new information about dynamic context, and we suggest how to implement AOP using continuation marks.

Broadly, continuation marks are useful because they allow us to specify models for a broad class of tools and features at a language level, where before these tools and features were specified in terms of an implementation, if

they were specified at all. Continuation marks are therefore a step forward in our understanding of program semantics. Conversely, the models built atop continuation marks suggest new implementation architectures.

5.2 Future Directions

Our work on continuation marks thus far suggests several avenues of future research, ranging from the purely practical to the largely theoretical.

First, our work on stack inspection and tail-calling suggest that continuation marks could be adapted to fit into a framework such as Sun's JVM or Microsoft's Common Language Runtime. Each of these frameworks currently contains a variety of features that might be more cleanly implemented using continuation marks or a related mechanism, and both would benefit from an improved understanding of tail-calling.

Second, our existing Scheme debuggers are research implementations, and not highly optimized. If we wish to convince language implementors that continuation marks are a solid foundation for debuggers and steppers in their languages, we will need to demonstrate that continuation marks can implement a highly optimized debugger for Scheme.

Finally, the declarative nature of continuation marks suggests that they should be applicable to a wide variety of language types, including lazy and concurrent languages for which debuggers are rare. It should also be possible to make use of DrScheme's ProfessorJ embedding of Java to build a stepper for Java without modifying the JVM.

Appendix A

Stepper Appendices

A.1 Annotation and Reconstruction Definitions

Annotation

Annotate : Program \rightarrow Program

$$\text{Annotate}[\langle \text{define } f_1 C_1 \rangle \dots \langle \text{define } f_n C_n \rangle] =$$
$$\langle \text{define } f_i \mathcal{W}_{\vec{k}}(\text{'outermost},$$
$$1,$$
$$\langle C_i \rangle,$$
$$\langle \mathcal{Q}[f_i], \mathcal{Q}_D[\langle \text{define } f_{i+1} C_{i+1} \rangle], \dots, \mathcal{Q}_D[\langle \text{define } f_n C_n \rangle] \rangle,$$
$$\langle f_1, \dots, f_{i-1} \rangle,$$
$$t_1) \dots \quad ;; \text{ for } 1 \leq i \leq n$$

where

$$\vec{k} = \{k | (\mathbf{w-c-m } k \dots) \text{ occurs in } \langle C, \dots \rangle\}$$

$$\mathcal{E}_{\vec{k}} : \text{Expr} \rightarrow \text{Expr}$$

$$\mathcal{E}_{\vec{k}}[C] = (\mathbf{w-c-m} \text{ debug} \quad \mathbf{false} \quad \mathcal{B}_{\vec{k}}(\mathcal{A}_{C_{\vec{k}}}[C], [(list \text{ 'expstep } \mathcal{Q}[C] (list (list \mathcal{Q}[x] x) \dots))])) \quad ;; \text{ for } x \in fv(C)$$

$$\mathcal{B}_{\vec{k}} : \text{Expr} \times \text{Expr} \rightarrow \text{Expr}$$

$$\mathcal{B}_{\vec{k}}[C_1][C_2] = (\mathbf{let} (t_{dc} (\mathbf{output} \text{ debug} \quad (list C_2 (\mathbf{c-c-m} \text{ debug } k \dots)))) \quad ;; \text{ for } k \in \vec{k} \quad C_1)$$

$$\mathcal{Q} : \text{Expr} \rightarrow \text{Expr}$$

$$\begin{aligned} \mathcal{Q}[n] &= n \\ \mathcal{Q}['x] &= 'x \\ \mathcal{Q}[p] &= p \\ \mathcal{Q}[(C_1 \dots)] &= (list \text{ 'app } \mathcal{Q}[C_1] \dots) \\ \mathcal{Q}[(\mathbf{lambda} (x \dots) C)] &= (list \text{ 'lambda } (list \text{ 'x } \dots) \mathcal{Q}[C]) \\ \mathcal{Q}[x] &= (list \text{ 'var } 'x) \\ \mathcal{Q}[f] &= (list \text{ 'topvar } 'f) \\ \mathcal{Q}[(\mathbf{cons} C_1 C_2)] &= (list \text{ 'cons } \mathcal{Q}[C_1] \mathcal{Q}[C_2]) \\ \mathcal{Q}[\mathbf{null}] &= \mathbf{null} \\ \mathcal{Q}[(\mathbf{car} C)] &= (list \text{ 'car } \mathcal{Q}[C]) \\ \mathcal{Q}[(\mathbf{cdr} C)] &= (list \text{ 'cdr } \mathcal{Q}[C]) \\ \mathcal{Q}[\mathbf{true}] &= \mathbf{true} \\ \mathcal{Q}[\mathbf{false}] &= \mathbf{false} \\ \mathcal{Q}[(\mathbf{if} C_1 C_2 C_3)] &= (list \text{ 'if } \mathcal{Q}[C_1] \mathcal{Q}[C_2] \mathcal{Q}[C_3]) \\ \mathcal{Q}[(\mathbf{w-c-m} k C_1 C_2)] &= (list \text{ 'w-c-m } 'k \mathcal{Q}[C_1] \mathcal{Q}[C_2]) \\ \mathcal{Q}[(\mathbf{c-c-m} k \dots)] &= (list \text{ 'ccm } 'k \dots) \\ \mathcal{Q}[(\mathbf{output} j C)] &= (list \text{ 'output } 'j \mathcal{Q}[C]) \end{aligned}$$

$$\mathcal{Q}_D : \text{Definition} \rightarrow \text{Expr}$$

$$\mathcal{Q}_D[(\mathbf{define} f C)] = (list \text{ 'cons } 'f \mathcal{Q}[C])$$

Configuration Annotation

$\mathcal{A}_{S_{\bar{k}}} : \text{Configuration} \rightarrow \text{Configuration}$

$$\begin{aligned} \mathcal{A}_{S_{\bar{k}}}(\langle C, E, \langle K, M \rangle, D \rangle) &= \langle \mathcal{A}_{C_{\bar{k}}}(C), \mathcal{A}_{E_{\bar{k}}}(E), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto \langle false \rangle] \rangle, \\ &\quad \mathcal{A}_{D_{\bar{k}}}(D) \rangle \\ \mathcal{A}_{S_{\bar{k}}}(\langle V, \langle K, \emptyset \rangle, D \rangle) &= \langle \mathcal{A}_{V_{\bar{k}}}(V), \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \emptyset \rangle, \mathcal{A}_{D_{\bar{k}}}(D) \rangle \\ \mathcal{A}_{S_{\bar{k}}}(\langle E \rangle) &= \langle \mathcal{A}_{E_{\bar{k}}}(E) \rangle \\ \mathcal{A}_{S_{\bar{k}}}(\langle error \rangle) &= \langle error \rangle \end{aligned}$$

$\mathcal{A}_{K_{\bar{k}}} : \text{Continuation} \times \text{Topenv} \rightarrow \text{Continuation}$

$$\begin{aligned} \mathcal{A}_{K_{\bar{k}}}(\langle \text{app} : \langle V_1, \dots, V_i \rangle, \langle C_1, \dots, C_j \rangle, E, \langle K, M \rangle \rangle, D) &= \\ \langle \text{app} : \langle V_b \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto V_m] \rangle \rangle & \\ \text{where } (V_b, V_m) = \mathcal{W}_{K_{\bar{k}}}(\text{'app}, \langle V_1, \dots, V_i \rangle, \langle C_1, \dots, C_j \rangle, \langle \rangle, (t_1 \dots t_{i+j}), \mathcal{A}_{E_{\bar{k}}}(E), \emptyset) & \\ \mathcal{A}_{K_{\bar{k}}}(\langle \text{cons1} : C, E, \langle K, M \rangle \rangle, D) &= \\ \langle \text{app} : \langle V_b \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto V_m] \rangle \rangle & \\ \text{where } (V_b, V_m) = \mathcal{W}_{K_{\bar{k}}}(\text{'cons}, \langle \rangle, \langle C \rangle, \langle \rangle, (\mathbf{cons} \ t_1 \ t_2), \mathcal{A}_{E_{\bar{k}}}(E), \emptyset) & \\ \mathcal{A}_{K_{\bar{k}}}(\langle \text{cons2} : V, \langle K, M \rangle \rangle, D) &= \\ \langle \text{app} : \langle V_b \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto V_m] \rangle \rangle & \\ \text{where } (V_b, V_m) = \mathcal{W}_{K_{\bar{k}}}(\text{'cons}, \langle V \rangle, \langle \rangle, \langle \rangle, (\mathbf{cons} \ t_1 \ t_2), \emptyset, \emptyset) & \\ \mathcal{A}_{K_{\bar{k}}}(\langle \text{car} : \langle K, M \rangle \rangle, D) &= \\ \langle \text{app} : \langle V_b \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto V_m] \rangle \rangle & \\ \text{where } (V_b, V_m) = \mathcal{W}_{K_{\bar{k}}}(\text{'car}, \langle \rangle, \langle \rangle, \langle \rangle, (\mathbf{car} \ t_1), \emptyset, \emptyset) & \\ \mathcal{A}_{K_{\bar{k}}}(\langle \text{cdr} : \langle K, M \rangle \rangle, D) &= \\ \langle \text{app} : \langle V_b \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto V_m] \rangle \rangle & \\ \text{where } (V_b, V_m) = \mathcal{W}_{K_{\bar{k}}}(\text{'cdr}, \langle \rangle, \langle \rangle, \langle \rangle, (\mathbf{cdr} \ t_1), \emptyset, \emptyset) & \\ \mathcal{A}_{K_{\bar{k}}}(\langle \text{if} : C_1, C_2, E, \langle K, M \rangle \rangle, D) &= \\ \langle \text{app} : \langle V_b \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto V_m] \rangle \rangle & \\ \text{where } (V_b, V_m) = \mathcal{W}_{K_{\bar{k}}}(\text{'if}, \langle \rangle, \langle \rangle, \langle C_1, C_2 \rangle, (\mathbf{if} \ t_1 \ C_1 \ C_2), \mathcal{A}_{E_{\bar{k}}}(E), \emptyset) & \\ \mathcal{A}_{K_{\bar{k}}}(\langle \text{wcm} : k, C, E, \langle K, M \rangle \rangle, D) &= \\ \langle \text{app} : \langle V_b \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto V_m] \rangle \rangle & \\ \text{where } (V_b, V_m) = \mathcal{W}_{K_{\bar{k}}}(\text{'wcm}, \langle \rangle, \langle \rangle, \langle 'k, C \rangle, (\mathbf{w-c-m} \ k \ t_1 \ C), \mathcal{A}_{E_{\bar{k}}}(E), \emptyset) & \\ \mathcal{A}_{K_{\bar{k}}}(\langle \text{out} : j, \langle K, M \rangle \rangle, D) &= \\ \langle \text{app} : \langle V_b \rangle, \langle \rangle, \emptyset, \langle \mathcal{A}_{K_{\bar{k}}}(K, D), \mathcal{A}_{E_{\bar{k}}}(M)[debug \mapsto V_m] \rangle \rangle & \\ \text{where } (V_b, V_m) = \mathcal{W}_{K_{\bar{k}}}(\text{'out}, \langle \rangle, \langle \rangle, \langle 'j \rangle, (\mathbf{output} \ j \ t_1), \emptyset, \emptyset) & \\ \mathcal{A}_{K_{\bar{k}}}(\langle \rangle, D) &= \\ \langle \text{app} : \langle V_b \rangle, \langle \rangle, \emptyset, \langle \rangle, \emptyset[debug \mapsto V_m] \rangle \rangle & \\ \text{where } (V_b, V_m) & \\ = \mathcal{W}_{K_{\bar{k}}}(\text{'outermost}, \langle \rangle, \langle \rangle, \langle \mathcal{Q}[f], \mathcal{Q}_D[(\mathbf{define} \ f_1 \ C_1)], \dots \rangle, t_1, \emptyset, \mathcal{A}_{E_{\bar{k}}}(E_p)) & \\ \text{and } D = \langle E_p, f, \langle \langle f_1, C_1 \rangle, \dots \rangle \rangle & \end{aligned}$$

$$\begin{aligned} \mathcal{W}_{K\bar{k}} &: (\text{Expr}, \langle \text{Value}, \dots \rangle, \langle \text{Expr}, \dots \rangle, \langle \text{Expr}, \dots \rangle, \text{Expr}, \text{Env}, \text{Env}) \\ &\rightarrow (\text{Value}, \text{Value}) \end{aligned}$$

$$\mathcal{W}_{K\bar{k}}(\ell, \langle V_1, \dots, V_n \rangle, \langle C_1, \dots \rangle, \langle C'_1, \dots \rangle, C'', E, E_t) = (\langle \text{clo} : \langle x \rangle, C_b, E' |_{fv(C_b) \setminus \langle x \rangle} \rangle, V')$$

where

$$\begin{aligned} E' &= E[t_1 \mapsto V_1] \dots \\ (C_m, E, E_t) &\mapsto_s V' \\ \mathcal{W}_{\bar{k}}(\ell, n+1, \langle C_1, \dots \rangle, \langle C'_1, \dots \rangle, \langle f_1, \dots \rangle, C'') \\ &= (\mathbf{w-c-m} \text{ debug } C_m ((\mathbf{lambda} (x) C_b) _)) \\ E_t &= \emptyset[f_1 \mapsto _] \dots \end{aligned}$$

$$\mathcal{A}_{V\bar{k}} : \text{Value} \rightarrow \text{Value}$$

$$\begin{aligned} \mathcal{A}_{V\bar{k}}(\langle \text{num} : n \rangle) &= \langle \text{num} : n \rangle \\ \mathcal{A}_{V\bar{k}}(\langle \text{sym} : x \rangle) &= \langle \text{sym} : x \rangle \\ \mathcal{A}_{V\bar{k}}(\langle \text{prim} : p \rangle) &= \langle \text{prim} : p \rangle \\ \mathcal{A}_{V\bar{k}}(\langle \text{clo} : \langle x, \dots \rangle, C, E \rangle) &= \langle \text{clo} : \langle x, \dots \rangle, \mathcal{E}_{\bar{k}}(C), \mathcal{A}_{E\bar{k}}(E) \rangle \\ \mathcal{A}_{V\bar{k}}(\langle \text{pair} : V_1, V_2 \rangle) &= \langle \text{pair} : \mathcal{A}_{E\bar{k}}(V_1), \mathcal{A}_{E\bar{k}}(V_2) \rangle \\ \mathcal{A}_{V\bar{k}}(\langle \text{null} \rangle) &= \langle \text{null} \rangle \\ \mathcal{A}_{V\bar{k}}(\langle \text{true} \rangle) &= \langle \text{true} \rangle \\ \mathcal{A}_{V\bar{k}}(\langle \text{false} \rangle) &= \langle \text{false} \rangle \end{aligned}$$

$$\mathcal{A}_{E\bar{k}} : \text{Env} \rightarrow \text{Env}$$

$$\mathcal{A}_{E\bar{k}}(\emptyset[x_1 \mapsto V_1] \dots) = \emptyset[x_1 \mapsto \mathcal{A}_{V\bar{k}}(V_1)] \dots$$

$$\mathcal{A}_{D\bar{k}} : \text{Topenv} \rightarrow \text{Topenv}$$

$$\mathcal{A}_{D\bar{k}}(\langle E_t, f, \langle \langle f_1, C_1 \rangle, \dots, \langle f_n, C_n \rangle \rangle \rangle) = \langle \mathcal{A}_{E\bar{k}}(E_t), f, \langle \langle f_1, C'_1 \rangle, \dots, \langle f_n, C'_n \rangle \rangle \rangle$$

where

$$E_t = \emptyset[f'_1 \mapsto _] \dots [f'_m \mapsto _]$$

and

$$\begin{aligned} C'_i &= \mathcal{W}_{K\bar{k}}(\langle \text{outermost}, 1, \langle \mathcal{Q}[[f_i]] \rangle, \mathcal{Q}[(\mathbf{define} \ f_{i+1} \ C_{i+1})] \rangle, \dots, \mathcal{Q}[(\mathbf{define} \ f_n \ C_n)] \rangle), \\ &\quad \langle f'_1, \dots, f'_m, f_1, \dots, f_{i-1} \rangle, t_1 \end{aligned}$$

$\mapsto_s: \text{Expr} \times \text{Env} \times \text{Env} \rightarrow \text{Value}$

$$\begin{aligned}
(n, E, E_t) &\mapsto_s \langle \text{num} : n \rangle \\
('x, E, E_t) &\mapsto_s \langle \text{sym} : x \rangle \\
(p, E, E_t) &\mapsto_s \langle \text{prim} : p \rangle \\
(x, E, E_t) &\mapsto_s E(x) \\
(f, E, E_t) &\mapsto_s E_t(f) \\
((\mathbf{cons} C_1 C_2), E, E_t) &\mapsto_s \langle \text{pair} : V_1, V_2 \rangle \text{ if } (C_1, E, E_t) \mapsto V_1 \text{ and } (C_2, E, E_t) \mapsto V_2 \\
(\mathbf{null}, E, E_t) &\mapsto_s \langle \text{null} \rangle \\
(\mathbf{true}, E, E_t) &\mapsto_s \langle \text{true} \rangle \\
(\mathbf{false}, E, E_t) &\mapsto_s \langle \text{false} \rangle
\end{aligned}$$

Reconstruction

Reconstruct : Value \rightarrow Config

$$\begin{aligned}
\text{Reconstruct} \llbracket (\mathbf{list} (\mathbf{list} \text{'valstep } V_1) V_2) \rrbracket &= \langle \mathcal{A}_V^{-1}(V_1), \langle \mathcal{R}_K(V_2), \emptyset \rangle, \mathcal{R}_D(V_2) \rangle \\
\text{Reconstruct} \llbracket (\mathbf{list} (\mathbf{list} \text{'expstep } V_3 V_4) (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug false}) V_1) V_2)) \rrbracket &= \langle \mathcal{Q}^{-1}(V_3), \mathcal{R}_E(V_4), \langle \mathcal{R}_K(V_2), \mathcal{R}_E(V_1) \rangle, \mathcal{R}_D(V_2) \rangle
\end{aligned}$$

$\mathcal{R}_K: \text{Value} \rightarrow \text{Continuation}$

$$\begin{aligned}
\mathcal{R}_K \llbracket (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug} (\mathbf{list} \text{'app} (\mathbf{list} V_1' \dots) V_3) (\mathbf{list} V_4 \dots) \mathbf{null})) V_1) V_2) \rrbracket &= \langle \text{app} : \langle \mathcal{A}_V^{-1}(V_1'), \dots \rangle, \langle \mathcal{Q}^{-1}(V_4), \dots \rangle, \mathcal{R}_E(V_3), \langle \mathcal{R}_K(V_2), \mathcal{R}_E(V_1) \rangle \rangle \\
\mathcal{R}_K \llbracket (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug} (\mathbf{list} \text{'cons null } V_3) (\mathbf{list} V_4) \mathbf{null})) V_1) V_2) \rrbracket &= \langle \text{cons1} : \mathcal{Q}^{-1}(V_4), \mathcal{R}_E(V_3), \langle \mathcal{R}_K(V_2), \mathcal{R}_E(V_1) \rangle \rangle \\
\mathcal{R}_K \llbracket (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug} (\mathbf{list} \text{'cons} (\mathbf{list} V_4) \text{'_V} \mathbf{null} \mathbf{null})) V_1) V_2) \rrbracket &= \langle \text{cons2} : \mathcal{A}_V^{-1}(V_4), \langle \mathcal{R}_K(V_2), \mathcal{R}_E(V_1) \rangle \rangle \\
\mathcal{R}_K \llbracket (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug} (\mathbf{list} \text{'car null } \text{'_V} \mathbf{null} \mathbf{null})) V_1) V_2) \rrbracket &= \langle \text{car} : \langle \mathcal{R}_K(V_2), \mathcal{R}_E(V_1) \rangle \rangle \\
\mathcal{R}_K \llbracket (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug} (\mathbf{list} \text{'cdr null } \text{'_V} \mathbf{null} \mathbf{null})) V_1) V_2) \rrbracket &= \langle \text{cdr} : \langle \mathcal{R}_K(V_2), \mathcal{R}_E(V_1) \rangle \rangle \\
\mathcal{R}_K \llbracket (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug} (\mathbf{list} \text{'if null } V_3) \mathbf{null} (\mathbf{list} V_4 V_5))) V_1) V_2) \rrbracket &= \langle \text{if} : \mathcal{Q}^{-1}(V_4), \mathcal{Q}^{-1}(V_5), \mathcal{R}_E(V_3), \langle \mathcal{R}_K(V_2), \mathcal{R}_E(V_1) \rangle \rangle \\
\mathcal{R}_K \llbracket (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug} (\mathbf{list} \text{'wcm null } V_3) \mathbf{null} (\mathbf{list} \text{'k } V_4))) V_1) V_2) \rrbracket &= \langle \text{wcm} : k, \mathcal{Q}^{-1}(V_4), \mathcal{R}_E(V_3), \langle \mathcal{R}_K(V_2), \mathcal{R}_E(V_1) \rangle \rangle \\
\mathcal{R}_K \llbracket (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug} (\mathbf{list} \text{'output null } \text{'_V} \mathbf{null} (\mathbf{list} \text{'j}))) V_1) V_2) \rrbracket &= \langle \text{out} : j, \langle \mathcal{R}_K(V_2), \mathcal{R}_E(V_1) \rangle \rangle \\
\mathcal{R}_K \llbracket (\mathbf{cons} (\mathbf{cons} (\mathbf{list} \text{'debug} (\mathbf{list} \text{'outermost null } \text{'_V} \text{'_V} \text{'_V})) \mathbf{null} \mathbf{null})) \rrbracket &= \langle \rangle
\end{aligned}$$

$\mathcal{R}_D : \text{Value} \rightarrow \text{Topenv}$

$$\begin{aligned} \mathcal{R}_D[\langle \text{cons } V_1 (\text{cons } V_2 V_3) \rangle] &= \mathcal{R}_D[\langle \text{cons } V_2 V_3 \rangle] \\ \mathcal{R}_D[\langle \text{cons } (\text{list } (\text{list } \text{'debug } (\text{list } \text{'outermost null } V_1 \text{ null } (\text{list } V_2 V_3 \dots))) \\ \text{null}) \rangle] &= \langle \mathcal{R}_E(V_1), \mathcal{Q}^{-1}(V_2), \langle \langle f_1, C_1 \rangle, \dots \rangle \rangle \end{aligned}$$

where

$$\mathcal{Q}_D^{-1}(V_3) \dots = (\text{define } f_1 C_1) \dots$$

$\mathcal{R}_E : \text{Value} \rightarrow \text{Env}$

$$\mathcal{R}_E[\langle \text{list } (\text{list } V_1 V_2) \dots \rangle] = \emptyset[x \mapsto \mathcal{A}_V^{-1}(V_2)] \dots$$

where

$$\mathcal{Q}^{-1}(V_1) \dots = x \dots$$

A.2 Evaluator Fragments

```
(* ID *)
type id = string
(* SOURCE AST *)
type exp = NUM of int | SYM of string
          | PRIM of string
          | APP of exp list | LAM of (id list) * exp
          | VAR of id | TOPVAR of id
          | CONS of exp * exp | NULL | CAR of exp | CDR of exp
          | TRUE | FALSE | IF of exp * exp * exp
          | WCM of id * exp * exp
          | CCM of (id list)
          | OUTPUT of id * exp
          | ANYEXP of lazylabel
(* SOURCE PROGRAMS *)
and prog = (id * exp) list
(* VALUES *)
and value = NUMVAL of int | SYMVAL of id
           | PRIMVAL of string
           | CLO of (id list * exp * val_env) | PAIR of (value * value)
           | NULLVAL | TRUEVAL | FALSEVAL
           | ANYVAL of lazylabel
and val_env = ENV of (id * value) list
            | ANYENV of (lazylabel * idlist)
            | RESTRICT of ((idlist list) * val_env)
            | EXTEND of ((id * value) * val_env)
and marks = val_env
```

```

and core_kont = APPK of (value list) * (exp list) * val_env * kont
    | CONS1K of exp * val_env * kont
    | CONS2K of value * kont
    | CARK of kont | CDRK of kont
    | IFK of exp * exp * val_env * kont
    | WCMK of id * exp * val_env * kont
    | OUTK of id * kont
    | EMPTYK
    | ANYKONT of lazylabel

and kont = core_kont * marks
and defs = DEFS of val_env * id * ((id * exp) list)
and configuration = EXP_RUNNING of (exp * val_env * kont * defs)
    | VAL_RUNNING of (value * kont * defs)
    | FINISHED of val_env
    | ERROR

and out = OUT of id * value | TAU
and idlist = IDS of id list
    | ANYIDS of lazylabel
    | UNION of (idlist list)
    | DIFF of (idlist * idlist)
    (* invariant: FREE_VARS may never
     * describe sets including "t" variables *)
    | FREE_VARS of lazylabel
and lazylabel = ID of id
    | QUOTED of lazylabel
    | EVALED_QUOTED of lazylabel
    | ANNOTATED of lazylabel
    | CCM_EXP of idlist
    | EVALED_CCM_EXP of (idlist * kont)
    | EVALED_CCM_FRAME of (idlist * val_env)
    | VAR_PAIRS of idlist
    | TOP_VAR_PAIRS of idlist
    | EVALED_VAR_PAIRS of (lazylabel * idlist)
    | EVALED_TOP_VAR_PAIRS of (lazylabel * idlist)
    | MARKS_OF of lazylabel
    | DOMAIN_OF of val_env
    | CLOSED of lazylabel

```

```

and compute c =
  match c with
    | ERROR -> raise Stopped
    | FINISHED _ -> raise Stopped
    | EXP_RUNNING (e,rho,k,d) ->
      exp_compute (e,rho,k,d)
    | VAL_RUNNING (v,k,d) ->
      value_compute(v,k,d)

```

```

and exp_compute (e,rho,k,((DEFS (finished,-,-) as d)) =
  (TAU,
   (let use_val v = VAL_RUNNING(v,(zero_marks k),d)
      (* destroy mark table; not observable *)
    and use_exp e new_core_k = EXP_RUNNING(e,(env_restrict [e] rho),
      (new_core_k,(ENV [])),d)
    in
    (match e with
      NUM(n) -> use_val (NUMVAL n)
    | SYM(id) -> use_val (SYMVAL id)
    | PRIM(prim) -> use_val (PRIMVAL (prim))
    | APP es ->
      (match es with
        [] -> raise (Syntax_error "empty application")
      | e :: esrest -> use_exp e (APPK([],esrest,(env_restrict
          esrest
          rho),k)))
    | LAM(ids,body) ->
      (use_val
       (CLO(ids,body,(env_restrict [e] rho))))
    | VAR(id) ->
      (use_val (env_lookup rho id))
    | TOPVAR(id) ->
      (try
       (use_val (env_lookup finished id))
      with
       Not_found -> ERROR)
    | CONS(e1,e2) -> use_exp e1 (CONSIK(e2,(env_restrict [e2] rho),k))
    | NULL -> use_val NULLVAL
    | CAR(e) -> use_exp e (CARK(k))
    | CDR(e) -> use_exp e (CDRK(k))
    | TRUE -> use_val TRUEVAL
    | FALSE -> use_val FALSEVAL
    | IF(test,e1,e2) -> use_exp test (IFK(e1,e2,(env_restrict [e1;e2] rho),k))
    | WCM (key,mark,body) ->
      (use_exp mark (WCMK (key,body,(env_restrict [body] rho),k)))
    | CCM ids ->
      (use_val (ccm_helper k ids))
    | OUTPUT(id,e) -> use_exp e (OUTK(id,k))
  )

```

```

| ANYEXP (lazylabel) ->
  (match lazylabel with
  (* using quoted-exp lemma: *)
  | (QUOTED l2) -> (use_val (ANYVAL
                           (EVALED_QUOTED l2)))
  (* using closed-configuration lemma: *)
  | (VAR_PAIRS idlist1) ->
    (if (environment_must_contain rho idlist1) then
      (use_val
       (ANYVAL
        (EVALED_VAR_PAIRS ((make_fresh_label ()),
                           idlist1))))
      else
      (raise (Abstract_expression_referenced
              (lazylabel,"eval"))))
  | (TOP_VAR_PAIRS idlist1) ->
    (if (environment_must_contain finished idlist1) then
      (use_val
       (ANYVAL
        (EVALED_TOP_VAR_PAIRS ((make_fresh_label ()),
                               idlist1))))
      else
      (raise (Abstract_expression_referenced
              (lazylabel,"eval"))))
  | (CCM_EXP idlist) ->
    (use_val
     (ANYVAL
      (EVALED_CCM_EXP (idlist,
                       k))))
  | - -> (raise (Abstract_expression_referenced
                 (lazylabel,"eval"))))

```

```

and value_compute(v,(outer_core_k,-),d) =
  (match outer_core_k with
    | APPK(vs,es,rho,k2) ->
      (match es with
        | [] ->
          (let new_vs = (List.append vs [v])
            in
              (match new_vs with
                | [] -> raise (Broken_invariant
                  "can't happen 200404081146")
                | fnv :: args ->
                  (TAU,
                    (match fnv with
                      | CLO(ids,body,rho2) ->
                        (if ((List.length ids) =
                            (List.length args)) then
                          EXP_RUNNING
                            (body,
                              (env_restrict
                                [body]
                                (List.fold_left2
                                  env_extend
                                  rho2
                                  ids
                                  args)),
                                k2,
                                d)
                            else
                              ERROR)
                      | PRIMVAL(prim_name) ->
                        (try (VAL_RUNNING
                            (((prim_table prim_name) args),
                              (zero_marks k2),d))
                          with
                            Prim_error -> ERROR)
                        | _ -> ERROR))))
          | e :: esrest -> (TAU, (EXP_RUNNING(e,
            (env_restrict [e] rho),
            (APPK(List.append vs [v],
              esrest,
              (env_restrict esrest rho),
              k2),
              (ENV [])),
            d))))
        | CONS1K(e,rho,k2) -> (TAU, (EXP_RUNNING(e,(env_restrict [e] rho),
          (CONS2K(v,k2),(ENV [])),d)))
        | CONS2K(vcar,k2) -> (TAU, (VAL_RUNNING
          (PAIR(vcar,v),(zero_marks k2),d)))
        | CARK(k2) ->
          (TAU, (match v with
            | PAIR(car,cdr) -> VAL_RUNNING(car,(zero_marks k2),d)
            | _ -> ERROR))

```

```

| CDRK(k2) ->
  (TAU, (match v with
    | PAIR(car, cdr) ->
      VAL_RUNNING(cdr, (zero_marks k2), d)
    | - -> ERROR))
| IFK(e1, e2, rho, k2) ->
  (TAU, (match v with
    | TRUEVAL ->
      EXP_RUNNING(e1, (env_restrict [e1] rho), k2, d)
    | FALSEVAL ->
      EXP_RUNNING(e2, (env_restrict [e2] rho), k2, d)
    | - -> ERROR))
| WCMK(key, body, rho, (core_k2, m)) ->
  (TAU, (EXP_RUNNING (body,
    (env_restrict [body] rho),
    (core_k2, (add_mark m (key, v))),
    d)))
| OUTK(id, k2) -> (output (id, v)
  (VAL_RUNNING
    (FALSEVAL, (zero_marks k2), d)))
| EMPTYK ->
  (TAU, (match d with
    | DEFS (finished, current_id, (next_id, next_exp)::rest) ->
      EXP_RUNNING(next_exp,
        (ENV []),
        (EMPTYK, (ENV [])),
        (DEFS ((env_extend
          finished
          current_id
          v),
          next_id,
          rest)))
    | DEFS (finished, current_id, []) ->
      FINISHED (env_extend finished current_id v))
| ANYKONT(id) -> (raise (Abstract_continuation_referenced
  (id, "value_compute"))))

```

Appendix B

Stack Inspection Appendices

B.1 Equivalence of Fournet and Gordon’s evaluator and the one presented herein

The model we use for reductions on λ_{sec} differs from that given by Fournet and Gordon in three minor ways, as mentioned in section 4.2. Of these, the only difference that requires explanation is our substitution of a static rewriting step for their dynamic check on **grant** expressions. Briefly, this static check is possible because all program code is a part of some component, and is therefore initially annotated with permissions. Inspection of the semantics shows that a frame expression always accompanies each **grant** expression, and therefore that part of its behavior can be predicted. In particular, the dynamic restriction of a **grant**’s permissions to that of the nearest enclosing frame expression can be performed at the time of annotation.

To argue this equivalence more formally, we model the original semantics with a modified reduction function \mapsto_o . The definition of \mapsto_o differs from that of \mapsto only in that the reference to \mathcal{OK} is replaced by a reference to \mathcal{OK}_o . This modified predicate dynamically restricts the permissions enabled by a **grant** to those appearing in its nearest enclosing frame expression, using an auxiliary `Static` function. The function \mapsto_o therefore models the original semantics of Fournet and Gordon.

Original Permissions Check $\mathcal{OK}_o \subseteq 2^{\mathcal{P}} \times E$

$$\begin{aligned} & \mathcal{OK}_o \langle R, [\bullet] \rangle \\ \mathcal{OK}_o \langle R, [E[\bullet M]] \rangle & \text{ iff } \mathcal{OK}_o \langle R, [E] \rangle \\ \mathcal{OK}_o \langle R, [E[V \bullet]] \rangle & \text{ iff } \mathcal{OK}_o \langle R, [E] \rangle \\ \mathcal{OK}_o \langle R, [E[S[\bullet]]] \rangle & \text{ iff } R \subseteq S \text{ and } \mathcal{OK}_o \langle R, [E] \rangle \\ \mathcal{OK}_o \langle R, [E[\text{grant } S \text{ in } \bullet]] \rangle & \text{ iff } \mathcal{OK}_o \langle R - \text{Static} \langle S, [E] \rangle, [E] \rangle \end{aligned}$$

where

$$\begin{aligned}
\text{Static}\langle R, \llbracket \bullet \rrbracket \rangle &= R \\
\text{Static}\langle R, \llbracket E[\bullet M] \rrbracket \rangle &= \text{Static}\langle R, \llbracket E \rrbracket \rangle \\
\text{Static}\langle R, \llbracket E[V \bullet] \rrbracket \rangle &= \text{Static}\langle R, \llbracket E \rrbracket \rangle \\
\text{Static}\langle R, \llbracket E[S[\bullet]] \rrbracket \rangle &= R \cap S \\
\text{Static}\langle R, \llbracket E[\text{grant } S \text{ in } \bullet] \rrbracket \rangle &= \text{Static}\langle R, \llbracket E \rrbracket \rangle
\end{aligned}$$

To extend the reduction function \mapsto_o to an evaluation function, we must prefix evaluation with an annotator as before. This annotator, \mathcal{A}_o , differs from \mathcal{A} in that it does not restrict the permissions in **grant** expressions to those attached to the entire component.

Definition 7 ((Eval_o))

$$\text{Eval}_o(C, \dots) = V \text{ if } (\mathcal{A}_o(C) \dots) \mapsto_o^* V$$

where

$$\begin{aligned}
\mathcal{A}_o\langle R, \llbracket x \rrbracket \rangle &= x \\
\mathcal{A}_o\langle R, \llbracket \lambda_f x. M \rrbracket \rangle &= \lambda_f x. R[\mathcal{A}_o\langle R, \llbracket M \rrbracket \rangle] \\
\mathcal{A}_o\langle R, \llbracket M N \rrbracket \rangle &= \mathcal{A}_o\langle R, \llbracket M \rrbracket \rangle \mathcal{A}_o\langle R, \llbracket N \rrbracket \rangle \\
\mathcal{A}_o\langle R, \llbracket \text{grant } S \text{ in } M \rrbracket \rangle &= \text{grant } S \text{ in } \mathcal{A}_o\langle R, \llbracket M \rrbracket \rangle \\
\mathcal{A}_o\langle R, \llbracket \text{test } S \text{ then } M \text{ else } N \rrbracket \rangle &= \text{test } S \text{ then } \mathcal{A}_o\langle R, \llbracket M \rrbracket \rangle \text{ else } \mathcal{A}_o\langle R, \llbracket N \rrbracket \rangle \\
\mathcal{A}_o\langle R, \llbracket \text{fail} \rrbracket \rangle &= \text{fail}
\end{aligned}$$

With these definitions in place, we can state the claim that the Eval functions are equal, modulo Fournet and Gordon's notion of contextual equivalence (\equiv_o).

Proposition 1 ((Static Permissions Check)) For any components (A_0, \dots) ,

$$\text{Eval}(A_0, \dots) = U \text{ iff } \text{Eval}_o(A_0, \dots) = V \text{ where } U \equiv_o V$$

Proof Sketch The proof proceeds in two steps. First, we establish the equivalence (modulo contextual equivalence) of Eval_o (the composition of \mathcal{A}_o and \mapsto_o^*) and the composition of \mathcal{A} and \mapsto_o^* . Second, we establish the equivalence of this composition and Eval (the composition of \mathcal{A} and \mapsto^*) when applied to terms that satisfy a new predicate \mathcal{B} , which describes the results of \mathcal{A} .

The two steps require three lemmas:

1. For all programs (A, \dots) , $(\mathcal{A}(A), \dots) \equiv_o (\mathcal{A}_o(A), \dots)$.
2. For all programs (A, \dots) , $\mathcal{B}(\emptyset, \llbracket (\mathcal{A}(A) \dots) \rrbracket)$.
3. For all terms M , $\mathcal{B}(\emptyset, \llbracket M \rrbracket)$ implies that $M \mapsto_o V$ iff $M \mapsto V$.

Lemma 1 implements the first step, lemma 3 corresponds to the third step, and lemma 2 provides the glue.

Lemma 1 Proof Sketch: The translations \mathcal{A}_o and \mathcal{A} differ in their treatment of `grant`; \mathcal{A} restricts the permissions contained in the `grant` to those that appear in the component's permissions, and \mathcal{A} doesn't. This alteration may be derived from the following equations in Fournet and Gordon's contextual equivalence theory:

Selected Equations

$$\begin{aligned}
 (\text{Frame Frame Appl}) : & R_1[R_2[e_1 \ e_2]] \equiv_o R_1[R_2[R_1[R_2[e_1]] \ R_1[R_2[e_2]]]] \\
 (\text{Frame Frame}) : & R_1 \supseteq R_2 \Rightarrow R_1[R_2[e]] \equiv_o R_2[e] \\
 (\text{Frame Grant}) : & R_1[\text{grant } R_2 \text{ in } e] \equiv_o R_1[\text{grant } R_1 \cap R_2 \text{ in } e] \\
 (\text{Frame Grant Frame}) : & R_1 \supseteq R_2 \Rightarrow R_1[\text{grant } R_2 \text{ in } R_3[e]] \equiv_o R_1[R_3[\text{grant } R_2 \text{ in } e]] \\
 (\text{Frame Test Then}) : & R_1 \supseteq R_2 \Rightarrow R_1[\text{test } R_2 \text{ then } e_1 \text{ else } e_2] \\
 & \equiv_o \text{test } R_2 \text{ then } R_1[e_1] \text{ else } R_1[e_2] \\
 (\text{Frame Test Else}) : & \neg(R_1 \supseteq R_2) \Rightarrow R_1[\text{test } R_2 \text{ then } e_1 \text{ else } e_2] \equiv_o R_1[e_2]
 \end{aligned}$$

For a given component $\langle R, M \rangle$, its annotation in Fournet and Gordon's system is $\mathcal{A}_o \langle R, M \rangle$. Since each top-level component expression M must be a λ -expression, at least one frame expression lies outside any `grant` that the component contains. Using the contextual equivalence theory, we can propagate frame expressions inward past any syntactic constructions other than abstraction. The only interesting case is `grant`, where pushing the frame inward requires the application of the Frame-Grant, Frame-Frame, Frame-Grant-Frame, and Frame-Frame equations. Since all abstraction bodies are wrapped in frame expressions with the component's permissions, this calculation leaves each `grant` expression wrapped with a frame expression. Then, the Frame-Grant rule justifies the intersection of the two sets of permissions. Finally, a reversal of the calculation applied thus far may be used to remove the inserted frame expressions. This leaves us with $\mathcal{A}(R, \llbracket M \rrbracket)$. By this reasoning, substituting \mathcal{A} for \mathcal{A}_o in the definition of Eval_o yields contextually equivalent results.

To make the second major step of the proof, we introduce the predicate \mathcal{B} .

Legal Grants Predicate $\mathcal{B} \subseteq 2^{\mathcal{P}} \times M$

$$\begin{aligned}
 & \mathcal{B} \langle R, \llbracket x \rrbracket \rangle \\
 \mathcal{B} \langle R, \llbracket \lambda_f x. M \rrbracket \rangle & \text{ iff } \mathcal{B} \langle \emptyset, \llbracket M \rrbracket \rangle \\
 \mathcal{B} \langle R, \llbracket M \ N \rrbracket \rangle & \text{ iff } \mathcal{B} \langle R, \llbracket M \rrbracket \rangle \ \mathcal{B} \langle R, \llbracket N \rrbracket \rangle \\
 \mathcal{B} \langle R, \llbracket S[M] \rrbracket \rangle & \text{ iff } \mathcal{B} \langle S, \llbracket M \rrbracket \rangle \\
 \mathcal{B} \langle R, \llbracket \text{grant } S \text{ in } M \rrbracket \rangle & \text{ iff } S \subseteq R \text{ and } \mathcal{B} \langle R, \llbracket M \rrbracket \rangle \\
 \mathcal{B} \langle R, \llbracket \text{test } S \text{ then } M \text{ else } N \rrbracket \rangle & \text{ iff } \mathcal{B} \langle R, \llbracket M \rrbracket \rangle \text{ and } \mathcal{B} \langle R, \llbracket N \rrbracket \rangle \\
 & \mathcal{B} \langle R, \llbracket \text{fail} \rrbracket \rangle
 \end{aligned}$$

The predicate \mathcal{B} checks two arguments: a set of permissions and an expression. It is satisfied¹ when all `grant` expressions refer to permissions that appear

¹The statement “ M satisfies \mathcal{B} ” is taken to mean that $\mathcal{B}(\emptyset, \llbracket M \rrbracket)$, or (equivalently) $(\emptyset, \llbracket M \rrbracket) \in \mathcal{B}$.

in the nearest enclosing frame expression, looking no further than the nearest λ boundary.

Lemma 2 Proof Sketch: The predicate \mathcal{B} formulates what the annotator \mathcal{A} enforces; namely, that **grant** expressions refer only to permissions accorded to their components. The proof proceeds by induction on the size of the program. The natural induction hypothesis states that for any R and M , $\mathcal{B}(R, \mathcal{A}(R, \llbracket M \rrbracket))$. That is, the annotation of M with R satisfies \mathcal{B} with permission set R . However, we must strengthen the induction hypothesis for λ -expressions to state that $\mathcal{B}(\emptyset, \mathcal{A}(R, \llbracket \lambda_f x.M \rrbracket))$. In other words, the annotation of a lambda term satisfies \mathcal{B} with the empty permissions set.

Lemma 3 Proof Sketch: We must now prove that the relations \mapsto and \mapsto_o act identically on terms that satisfy \mathcal{B} . First, we show that \mathcal{OK} and \mathcal{OK}_o are equivalent for evaluation contexts formed from expressions that satisfy \mathcal{B} . Second, a subject reduction proof shows that satisfaction of \mathcal{B} is preserved by both \mapsto and \mapsto_o .

Suppose M satisfies \mathcal{B} , and $M = E[N]$. Then for any R , $\mathcal{OK}_o(R, \llbracket E \rrbracket)$ iff $\mathcal{OK}(R, \llbracket E \rrbracket)$. The satisfaction of \mathcal{B} guarantees that **Static** acts as the identity; the permissions attached to a **grant** are already restricted to those occurring in the nearest enclosing frame.

The subject reduction proof is largely mechanical. The only interesting cases are those in which a frame is removed and those involving a β_v reduction. In each case, the key observation is that λ expressions are self-contained. By this we mean that the value of $\mathcal{B}(R, \llbracket \lambda_f x.M \rrbracket)$ does not depend on R at all. Therefore, substituting a value (that is, a lambda expression) that satisfies \mathcal{B} for any other expression does not change an expression that satisfies \mathcal{B} into one that doesn't. Furthermore, this argument applies to the bodies of abstractions as well. That is, if a term before a substitution satisfied \mathcal{B} and contained the term $\lambda_f x.M$, we may conclude that $\mathcal{B}(\emptyset, \llbracket M \rrbracket)$, which implies that for any choice of R , $\mathcal{B}(R, \llbracket M \rrbracket)$ also holds, and therefore that the substitution of the term M for another term preserves satisfaction. Both of the reductions of interest consist entirely of one or more such substitutions, and must therefore preserve satisfaction of \mathcal{B} . This argument applies without modification to both \mapsto and \mapsto_o .

With these two pieces in hand, a simple case analysis shows that \mapsto and \mapsto_o behave identically on terms that satisfy \mathcal{B} .

Taken together, the three lemmas allow us to conclude that our proposition holds. \square

Appendix C

Aspect-Oriented Programming using Continuation Marks

Nearly every programming language has a notion of modularity, and a corresponding unit of organization. Broadly speaking, *Aspect-Oriented Programming* is the idea that for any such division, there will be conceptual elements of functionality—or *concerns*—whose code must be spread across many different units. Aspect-oriented programming, or AOP, proposes that the code for a given concern be gathered in one place—called *advice*—with an accompanying specification that describes how the aspect is to be applied to an existing body of code.

The hope is that this alternative specification of a program will be better able to divide code along the lines dictated by concerns, and therefore will be easier to maintain and extend.

A second goal of aspect-oriented programming is that aspects may be natural units of extension. That is, modifying a program by adding an aspect should be more robust than modifying the program itself.

This chapter contains three sections. In the first, I introduce some of the current aspect-oriented programming tools. In the second, I describe how to implement aspect-oriented language forms using continuation marks and macros. In the third, I discuss the relationship between tail-calling and aspect-oriented programming.

This chapter is partially based on work by Shriram Krishnamurthi and David Tucker [45], who are responsible for the insight connecting AOP and continuation marks.

C.1 Aspect-Oriented Programming

The past few years have seen an explosion of aspect-oriented frameworks and toolsets, but the core ideas of AspectJ [30] remain at the center of most current definitions of AOP.

AOP Concepts

In the AspectJ model, a program’s execution consists of a sequence of execution states. Rather than attempting to give a detailed semantics for the language, AspectJ designates some identifiable execution subsequences as “join points.” Method calls are considered join points, as are constructor calls, field references, field mutations, and a host of other evaluation subranges in the inferred semantics.

In AspectJ, a program’s execution is modified through the application of advice at certain join points. To indicate which advice applies to which join points, AspectJ defines a set of join point specifications, called “pointcut designators” or simply “pointcuts.” These specifications are defined using a specialized sublanguage that allows matching on join points. So, for instance, a pointcut might specify all calls to a particular method, or all calls to constructors matching a certain pattern, and so forth. The language also includes a family of boolean pointcut combinators.

When a program’s execution is seen to match a given pointcut, the corresponding advice is used to modify the program’s execution. This advice may be evaluated before, after, or instead of the existing evaluation sequence. Advice therefore consists of two things: a form—*before*, *after*, *around*, or other less common ones—and a piece of code, much like a method. This code is called with arguments particular to the corresponding pointcut that inform the advice about the execution state. For instance, advice acting on a method call will be called with the method call’s arguments, among other things. The advice’s form determines whether it is called before, after, or instead of (*around*) the specified pointcut. The *around* specification is the most general, and comes with an additional `proceed` binding that the advice may use to trigger the evaluation of the original join point—albeit in a slightly different context.

AOP Examples

```

before(Point p, int x): target(p)
    @@ args(x)
    @@ call(void setX(int)) {
    if (!p.assertX(x)) {
        System.out.println("Illegal value for x"); return;
    }
}

```

Figure C.1: A simple piece of advice

The simplest non-trivial examples of aspect-oriented programming are those that add an orthogonal side-effect to a computation. Figure C.1 shows a simple

example of an invariant check attached to a mutation method. This example comes directly from the AspectJ documentation.

As their name suggests, pointcuts can also match elements of the execution state other than the innermost control point. In particular, AspectJ provides a unary pointcut combinator called *cflow* that applies the join point matching mechanism to higher points on the dynamic call graph. That is, if a pointcut *A* designates a call to a particular method, then *cflow(A)* would designate all join points within the dynamic extent of that call. Using *cflow* and boolean combination, it is possible to specify that a piece of advice, say, should be applied only to calls to a method *m()* that occur within the dynamic extent of a call to *p()*, except for those join points where the call to *p()* was within the dynamic extent of a call another method *q()*.

C.2 Implementing AOP

As the rich language of pointcuts makes clear, the basic intent of AOP is to allow arbitrary interventions to the execution path of a running program. These interventions can halt execution, alter execution, display portions of program state, etc. In its basic mechanism, then, aspect-oriented programming of the AspectJ kind bears a striking resemblance to debugging.

Since continuation marks were devised in order to implement debugging, and more particularly to allow a program to observe nonlocal elements of its dynamic execution state, they are also well-suited as a tool for implementing aspect-oriented programming.

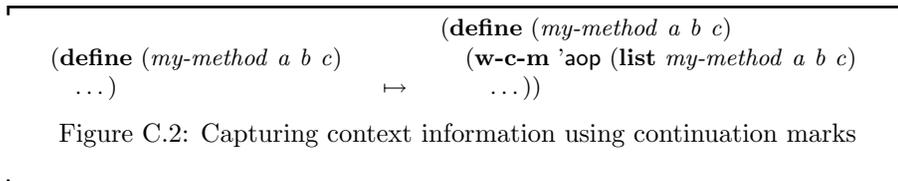


Figure C.2: Capturing context information using continuation marks

The fundamental action of aspect-oriented programming is the matching of the current execution state against a set of pointcut patterns. In order to capture the information needed to perform this matching using continuation marks, it suffices to place a continuation mark on entry to each procedure, storing the procedure itself and the arguments the procedure was called with. Figure C.2 illustrates this transformation. A call to (**c-c-m 'aop**) then retrieves information about all existing context. This represents the dynamic execution state in a form that can be used to match a pointcut.

In the simplest possible implementation of aspect-oriented programming using continuation marks, the resulting program looks much like a stepper. That is, a program is annotated with checks at every step of the program to see whether any of the defined pointcuts match the current execution state. If so, the corresponding piece of advice is evaluated, rather than the current code.

A more sophisticated implementation would remove checks at all points that can be proven never to match one of the given pointcuts. Finally, it would often be the case that many of the mark-placing **w-c-m**'s could be omitted, in cases where the compiler determines that their omission cannot affect the behavior of the program.

C.3 CFlow and Tail-Calling

AOP's *cflow* and MzScheme's continuation marks both make dynamic context visible to a program's inner context. That is, a program using a *cflow* pointcut can cause the meaning of a method $m()$ to be different when it is called in the dynamic extent of a method $p()$. Likewise, continuation marks allow calls to $p()$ to place a mark that is visible to calls to $m()$ that occur within that mark's dynamic extent.

However, a simple mark-placement protocol of the type described in the prior section could fail to behave correctly for tail calls. Suppose, for instance, that the method $p()$ calls a method $q()$ in tail position. If $q()$ then has the same context that the call to $p()$ did, and places a continuation mark on that context, then by the evaluation rules of continuation marks the earlier mark would be replaced, and information about the dynamic extent of $p()$ would be lost.

It is not hard to work around this problem; just as in the case of exception handlers or other parameters, it is straightforward to add to the context using a stack-like extension, as *extend-parameterization* does for parameters.

Nevertheless, I believe that this raises questions that should not be so quickly dismissed: is 'dynamic extent' always the right semantic expression for a *cflow*-like construct? When defining exception handlers, dynamic extent makes perfect sense: "I want to change the behavior of the program for a well-defined period of time." In the case of *cflow*, however, the question has more to do with causality: "Was this call to $m()$ caused by a call to $p()$?" In the context of program execution, this question is not well phrased. After all, it may well be that the call to $m()$ was "caused" by a call to $p()$ that returned before $m()$ was invoked. In AOP, as in debugging, it may be time to take a closer look at execution history as a better metaphor for history than the glimpse of the future afforded by the stack.

Bibliography

- [1] Balzer, R. M. EXDAMS—EXTendable debugging and monitoring system. In *AFIPS 1969 Spring Joint Computer Conference*, volume 34, pages 567–580. AFIPS Press, May 1969.
- [2] Barendregt, H. P. *The Lambda Calculus: Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, revised edition, 1984.
- [3] Benton, N., A. Kennedy and G. Russell. Compiling standard ML to Java bytecodes. In *ACM SIGPLAN International Conference on Functional Programming*, pages 129–140, 1998.
- [4] Bernstein, K. L. and E. W. Stark. Operational semantics of a focusing debugger. In *Eleventh Conference on the Mathematical Foundations of Programming Semantics, Volume 1 of Electronic Notes in Computer Science*. Elsevier, March 1995.
- [5] Bertot, Y. Occurrences in debugger specifications. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, 1991.
- [6] Clements, J. and M. Felleisen. A tail-recursive semantics for stack inspections. In Degano, P., editor, *Proceedings of the 12th European Symposium on Programming*, volume 2618 of *Lecture Notes in Computer Science*, pages 22–37. Springer, 2003.
- [7] Clements, J. and M. Felleisen. A tail-recursive machine with stack inspection. *ACM Transactions on Programming Languages and Systems*, 26(6):1–24, November 2004.
- [8] Clements, J., M. Felleisen, R. Findler, M. Flatt and S. Krishnamurthi. Fostering little languages. *Dr. Dobb's Journal*, March 2004. (Invited Paper).
- [9] Clements, J., M. Flatt and M. Felleisen. Modeling an algebraic stepper. In Sands, D., editor, *Proceedings of the 10th European Symposium on Programming*, volume 2028 of *Lecture Notes in Computer Science*, pages 320–334. Springer, 2001.

- [10] Clements, J., P. Graunke, S. Krishnamurthi and M. Felleisen. Little languages and their programming environments. In *Proceedings of the Monterey Workshop on Engineering Automation for Software Intensive System Integration*, pages 1–18, 2001.
- [11] Clinger, W. D. Proper tail recursion and space efficiency. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 174–185, 1998.
- [12] Danvy, O. and A. Filinski. Representing control: A study of the cps transformation. *Mathematical Structures in Computer Science*, 4:360–391, 1992.
- [13] Erlingsson, U. and F. B. Schneider. IRM enforcement of java stack inspection. In *IEEE Symposium on Security and Privacy*, pages 246–255, 2000.
- [14] Felleisen, M., R. B. Findler, M. Flatt and S. Krishnamurthi. *How To Design Programs*. MIT Press, 2001.
- [15] Felleisen, M. and M. Flatt. Programming languages and their calculi. Unpublished Manuscript. Online at <http://www.ccs.neu.edu/home/matthias/3810-w02/mono.ps.gz>, 1989–2002.
- [16] Felleisen, M. and D. P. Friedman. Control operators, the secd-machine, and the λ -calculus. In Wirsing, M., editor, *Formal Description of Programming Concepts III*, pages 193–217. Elsevier Science Publishers B.V., 1986.
- [17] Felleisen, M. and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 1992.
- [18] Ferguson, H. E. and E. Berner. Debugging systems at the source language level. *Communications of the ACM*, 6(8):430–432, August 1963.
- [19] Findler, R. B., J. Clements, C. Flanagan, M. Flatt, S. Krishnamurthi, P. Steckler and M. Felleisen. Drscheme: A programming environment for Scheme. *Journal of Functional Programming*, 12(2):159–182, 2002.
- [20] Flanagan, C., A. Sabry, B. F. Duba and M. Felleisen. The essence of compiling with continuations. In *PLDI '93: Proceedings of the ACM SIGPLAN 1993 conference on Programming language design and implementation*, pages 237–247, New York, NY, USA, 1993. ACM Press.
- [21] Flatt, M. PLT MzScheme: Language manual. Online at <http://www.plt-scheme.org>, 1995–2002.
- [22] Fournet, C. and A. D. Gordon. Stack inspection: theory and variants. In *ACM SIGPLAN Conference on Principles of Programming Languages*, pages 307–318, 2002.

- [23] Gong, L. *Inside Java 2 Platform Security*. Sun Microsystems, 1999.
- [24] Gordon, A. D. and D. Syme. Typing a multi-language intermediate code. In *ACM SIGPLAN Conference on Principles of Programming Languages*, pages 248–260, 2001.
- [25] Hall, C. and J. O’Donnell. Debugging in a side effect free programming environment. In *ACM SIGPLAN symposium on Language issues in programming environments*, 1985.
- [26] Karjoth, G. An operational semantics of Java 2 access control. In *The Computer Security Foundations Workshop*, pages 224–232, 2000.
- [27] Kellomaki, P. Psd—a portable scheme debugger, February 1995.
- [28] Kelsey, R., W. D. Clinger and J. Rees. Revised⁵ report on the algorithmic language scheme. *SIGPLAN Notices*, 33(9):26–76, 1998.
- [29] Kelsey, R. A. A correspondence between continuation passing style and static single assignment form. *ACM SIGPLAN Notices*, 30(3):13–22, March 1995.
- [30] Kiczales, G., E. Hilsdale, J. Hugunin, M. Kersten, J. Palm and W. Griswold. An overview of aspectj. In *Proceedings of the European Conference on Object-Oriented Programming*, 2001.
- [31] Kishon, A., P. Hudak and C. Consel. Monitoring semantics: A formal framework for specifying, implementing, and reasoning about execution monitors. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 338–352, 1991.
- [32] Leroy, X. *The Objective Caml system release 3.08*, 2004. On the web at <http://caml.inria.fr/pub/docs/manual-ocaml/index.html>.
- [33] Lindholm, T., F. Yellin, B. Joy and K. Walrath. *The Java Virtual Machine Specification*. Addison-Wesley, 1997.
- [34] Microsoft. Common language runtime SDK documentation. Online at <http://www.microsoft.com>. Part of .NET SDK documentation, 2002.
- [35] Milner, R. *Communication and Concurrency*. Prentice Hall, 1989.
- [36] Naish, L. and T. Barbour. Towards a portable lazy functional declarative debugger. In *19th Australasian Computer Science Conference*, 1996.
- [37] Plotkin, G. D. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science*, pages 125–159, 1975.
- [38] Pottier, F., C. Skalka and S. Smith. A systematic approach to static access control. In *European Symposium on Programming*, pages 30–45, 2001.

- [39] Sansom, P. and S. Peyton-Jones. Formally-based profiling for higher-order functional languages. *ACM Transactions on Programming Languages and Systems*, 19(1), January 1997.
- [40] Schinz, M. and M. Odersky. Tail call elimination on the Java virtual machine. In *SIGPLAN BABEL Workshop on Multi-Language Infrastructure and Interoperability*, pages 155–168, 2001.
- [41] Skalka, C. and S. Smith. Static enforcement of security with types. *ACM SIGPLAN Notices*, 35(9):34–45, 2000.
- [42] Sparud, J. and C. Runciman. Tracing lazy functional computations using redex trails. In *Symposium on Programming Language Implementation and Logic Programming*, 1997.
- [43] Steele Jr., G. L. Debunking the “expensive procedure call” myth. In *ACM Conference*, pages 153–162, 1977.
- [44] Tolmach, A. *Debugging Standard ML*. PhD thesis, Department of Computer Science, Princeton University, October 1992.
- [45] Tucker, D. and S. Krishnamurthi. Pointcuts and advice in higher-order languages. In *Proceedings of the 2nd International Conference on Aspect-Oriented Programming*, 2003.
- [46] Wallach, D., D. Balfanz, D. Dean and E. Felten. Extensible security architectures for Java. In *The 16th Symposium on Operating Systems Principles*, pages 116–128, october 1997.
- [47] Wallach, D., E. Felten and A. Appel. The security architecture formerly known as stack inspection: A security mechanism for language-based systems. *ACM Transactions on Software Engineering and Methodology*, 9(4):341–378, October 2000.