

An operational semantics for R⁵RS Scheme

Jacob Matthews
University of Chicago
jacobm@cs.uchicago.edu

Robert Bruce Findler
University of Chicago
robby@cs.uchicago.edu

Abstract

This paper presents an operational semantics for the core of Scheme. Our specification improves over the existing R⁵RS denotational specification in four ways. First, it is more complete, since it contains *eval*, **quote**, and *dynamic-wind*. Second, it models multiple values in a way that does not require changes to unrelated parts of the language. Third, it provides a more faithful model of Scheme’s undefined order of evaluation. Finally, it is executable, because it is encoded as a program in PLT Redex, a domain-specific language for writing operational semantics. The executable specification allows others to experiment with our specification and allows us to build a specification test suite, which improves our confidence that our system is a faithful model of Scheme.

In addition to contributing a specification of Scheme, this paper presents several novel modeling techniques for Felleisen Hieb-style rewriting semantics that we discovered while developing our R⁵RS Scheme semantics. All are applicable to a wider range of problems than the specific uses we have for them, and the fact that they combine seamlessly in our full R⁵RS model shows that they scale to real languages.

1. Introduction

The Revised⁵ Report on the Algorithmic Language Scheme [15], R⁵RS, provides an informal, English specification of Scheme and a denotational model of a core Scheme language. The denotational specification is more precise than the informal specification, but is also incomplete with respect to it. For instance, the formal specification does not present the top-level mentioned throughout the informal specification, and is missing key procedures such as *dynamic-wind* and *eval* whose inclusion could have a significant impact on the formalism. While that is not necessarily a problem — the measure of a model is not its completeness but its ability to clearly and accurately explain its subject — Gasbichler et al’s recent explanation of the difficulties involving dynamic contexts and threads [12], for instance, demonstrate that the formal model is insufficient for some important questions.

In this paper we give a new treatment of the R⁵RS formal semantics that models more of the language described in the informal semantics section than the formal semantics section in the R⁵RS Scheme document does. Rather than extending the denotational semantics with extra constructs, we present an alternate specification as a small-step operational semantics. We do this for two major reasons. First, to make the semantics natively executable: operational semantics are much more amenable to direct execution than denotational semantics. Second, to allow for nondeterminism and non-confluence: small-step operational semantics are particularly well-

suited for modeling programming languages with nondeterministic and nonconfluent behavior. We make important use of nondeterminism in our model, as we will explain in section 2.

As a side benefit of using a small-step operational encoding, we can use PLT Redex [17], a domain-specific language for context-sensitive term-rewriting systems, to give a directly executable operational encoding for our model. PLT Redex provides a graphical browser for exploring reduction graphs and allows us to maintain a large test suite of terms and their expected normal forms that we can run whenever we change any reduction rules. This test suite increases our confidence that our model is a faithful representation of Scheme.

While writing our model, we developed new techniques for modeling some of Scheme’s features. In the rest of our paper we first introduce those techniques in isolation to explain our models for particular Scheme features, and then combine them into a single unified model. In section 2 we show how to use nondeterminism to model Scheme’s unspecified application order; in section 3 we show a novel technique for modeling multiple return values; in section 4 we give a model for **quote** and *eval*; and in section 5 we give a model for *call/cc* in the presence of *dynamic-wind*. Finally in section 6 we combine all those models along with several other more straightforward features: **if**, *cons* and cons-cell mutation, variable-arity procedures, *apply*, and an object-identity-sensitive notion of *eqv?* equality.

We will assume the reader has a basic familiarity with context-sensitive reduction semantics. Readers unfamiliar with this system may wish to consult Felleisen and Flatt’s monograph [5] or Wright and Felleisen [24] for a thorough introduction or our previous work with Flatt and Felleisen [17] for a somewhat lighter one. We should also emphasize before we proceed that this semantics still leaves out many important Scheme features — among them the numeric tower, the top-level environment, and macros — but that it models more features than the Report’s formal semantics does and is more suitable for extension.

2. Unspecified application order

In evaluating a procedure call, the R⁵RS document deliberately leaves unspecified the order in which arguments are evaluated, but section 4.1.3 specifies that

the effect of any concurrent evaluation of the operator and operand expressions is constrained to be consistent with some sequential order of evaluation. The order of evaluation may be chosen differently for each procedure call.

In the formal semantics section, the authors explain how they model this ambiguity:

[w]e mimic [the order of evaluation] by applying arbitrary permutations permute and unpermute . . . to the arguments in a call before and after they are evaluated. This is not quite

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

Sixth Workshop on Scheme and Functional Programming. September 25, 2005, Tallinn, Estonia.

Copyright © 2005 Jacob Matthews and Robert Bruce Findler.

$ \begin{aligned} p & ::= (\mathbf{store} ((x\ v) \cdots) e) \\ e & ::= (e\ e \cdots) \mid (\mathbf{set!}\ x\ e) \mid (\mathbf{begin}\ e\ e \cdots) \mid v \\ v & ::= (\mathbf{lambda}\ (x \cdots) e) \mid n \end{aligned} $	$ \begin{aligned} C & ::= (v \cdots C\ e \cdots) \mid (\mathbf{set!}\ x\ C) \mid (\mathbf{begin}\ C\ e\ e \cdots) \mid [] \\ x & ::= \text{identifiers, store locations for mutable bindings} \\ n & ::= \text{numbers} \end{aligned} $
$ (\mathbf{store} ((x_1\ v_1) \cdots) C[(\mathbf{lambda}\ (x_2 \cdots) e)\ v_2 \cdots]) \rightarrow $	$ (\mathbf{store} ((x_1\ v_1) \cdots (x'_2\ v_2) \cdots) C[e[x'_2 \cdots / x_2 \cdots]]) \quad (\text{MAPP}) $ <p style="font-size: small; margin-left: 20px;">(#$x_2 = \#v_2$, each x'_2 fresh)</p>
$ (\mathbf{store} ((x_1\ v_1) \cdots) C[(\mathbf{lambda}\ (x_2 \cdots) e)\ v_2 \cdots]) \rightarrow $	$ \mathbf{error}: \text{wrong number of arguments} \quad (\text{MAPPERR}) $ <p style="font-size: small; margin-left: 20px;">(#$x_2 \neq \#v_2$)</p>
$ (\mathbf{store} ((x_1\ v_1) \cdots (x\ v)\ (x_2\ v_2) \cdots) C[(\mathbf{set!}\ x\ v')]) \rightarrow $	$ (\mathbf{store} ((x_1\ v_1) \cdots (x\ v')\ (x_2\ v_2) \cdots) C[0]) \quad (\text{MSET}) $
$ (\mathbf{store} ((x_1\ v_1) \cdots (x\ v)\ (x_2\ v_2) \cdots) C[x]) \rightarrow $	$ (\mathbf{store} ((x_1\ v_1) \cdots (x\ v)\ (x_2\ v_2) \cdots) C[v]) \quad (\text{MLOOKUP}) $
$ (\mathbf{store} ((x\ v) \cdots) C[(\mathbf{begin}\ v\ e_1\ e_2 \cdots)]) \rightarrow $	$ (\mathbf{store} ((x\ v) \cdots) C[(\mathbf{begin}\ e_1\ e_2 \cdots)]) \quad (\text{MSEQ}) $
$ (\mathbf{store} ((x\ v) \cdots) C[(\mathbf{begin}\ e)]) \rightarrow $	$ (\mathbf{store} ((x\ v) \cdots) C[e]) \quad (\text{MTRIVSEQ}) $
$ (\mathbf{store} ((x\ v) \cdots) C[(-\ \lceil n \rceil)]) \rightarrow $	$ (\mathbf{store} ((x\ v) \cdots) C[\lceil -n \rceil]) \quad (\text{MNEG}) $

Figure 1. Core Scheme with mutation

right since it suggests, incorrectly, that the order of evaluation is constant throughout a program ... [section 7.2]

In this section we present an operational technique that captures the intended semantics more faithfully. We begin by considering a core Scheme with arbitrary arity procedures, **set!**, numbers, and negation, but with a fixed left-to-right order of evaluation for applications, as shown in figure 1. It is a minor variation of Felleisen and Hieb’s Λ_S [6]. A program consists of a store that associates variable names to values and an expression, where expressions are built up of numbers, arbitrary-arity lambda terms and applications, **set!**, and **begin** expressions, and a built-in negation operator. MAPP gives the rule for application of a procedure to fully-evaluated arguments: make one fresh identifier x'_i for each formal parameter x_i , introduce a new binding in the store for each x'_i associating it with the corresponding argument v_i in the application, and then rewrite the application as the procedure’s body with each occurrence of an x_i rewritten into the corresponding x'_i (in this figure as in all figures in this paper, we will use vertically-centered ellipses \cdots to indicate any number of occurrences, including zero, of the preceding element). MAPPERR gives the rule for procedures applied to the wrong number of arguments: rewrite the term in its entirety to an error message, which halts the program immediately because it abandons the application’s original context. MSET rewrites to the constant 0 but also replaces the value associated with the given identifier in the store with the given replacement. (We choose to have **set!** return the constant 0 in this semantics as a “quick and dirty” unique value; in the examples that follow 0 never appears in any program term except as the result of assignment.) MLOOKUP replaces an identifier with its associated value in the store when that value becomes necessary (*i.e.*, when it appears as a redex in an evaluation context). MSEQ drops the first subexpression in a **begin** when there are more expressions to evaluate, and MTRIVSEQ drops the **begin** when there is only one expression to evaluate. The last rule, MNEG, simply negates its argument (the notation $\lceil n \rceil$ indicates the syntactic representation corresponding to the mathematical number n).

The order of evaluation is determined by the grammar for evaluation contexts (C). The first production of the grammar specifies that evaluation of a sub-expression of an application only takes place when all of the sub-expressions to its left are values (or have been reduced to values). If we replace that first production with this one:

$$C ::= (e \cdots C\ v \cdots) \mid \dots$$

the semantics would specify a right-to-left order instead.

Either of these choices results in a system with unique decomposition. That is, each term can only be split into an evaluation context and a reducible sub-expression in one way (unless it is stuck

or an answer). Accordingly, there is at most one way to reduce any expression.

To model a language with unspecified order of operations instead, we can use a reduction system with non-unique decomposition to model the choice. We might be tempted to use this definition of evaluation contexts:

$$C ::= (e \cdots C\ e \cdots) \mid \dots$$

Since this definition allows the hole to appear in any subexpression of an application, this simple program that negates 1, negates 2, and then applies a trivial procedure to the results

$$((\mathbf{lambda}\ (x\ y)\ y)\ (-\ 1)\ (-\ 2))$$

can be split into an evaluation context with either $(- 1)$ or $(- 2)$ as the reducible expression.

At first glance, this appears to be a faithful model of R^5RS Scheme. It is not. Consider this application of two **set!** expressions in a store binding x to 1 .

$$\begin{aligned}
& (\mathbf{store}\ ((x\ 1)) \\
& \quad ((\mathbf{set!}\ x\ (-\ x)) \\
& \quad (\mathbf{set!}\ x\ (-\ x))))
\end{aligned}$$

In Scheme, this program should always reduce to the application of zero to zero with x set to 1 in the store (and then get stuck). According to R^5RS , no matter which of the application’s subterms is reduced first, the result should be that x is negated twice. If we just modify evaluation contexts as above, however, we allow other interleavings. The problem is that that definition of evaluation contexts would allow a different argument of the same application to take one step of computation every step of the way, which may produce an outcome that could not be reached by any sequential ordering.

We discovered this problem while experimenting with that reduction system in PLT Redex. We encoded the erroneous reduction system in PLT Redex and automatically generated the reduction sequence for the above term, shown in figure 2. The first term is shown on the left. The top-most and the bottom-most paths correspond to the two sequential orderings and result in the proper store. In the middle section, the two assignments are interleaved, resulting in -1 being left in the store.

With that in mind, we can design a more sophisticated strategy that captures unspecified evaluation order but only allows sequential orderings. Figure 3 shows the necessary revisions to core Scheme to support R^5RS -style procedure applications (each replaces the appropriate rule from figure 1 — the other rules in that figure are unchanged). The basic idea is to use non-deterministic choice to pick a sub-expression to reduce only when we have not already committed to reducing some other subexpression. To achieve

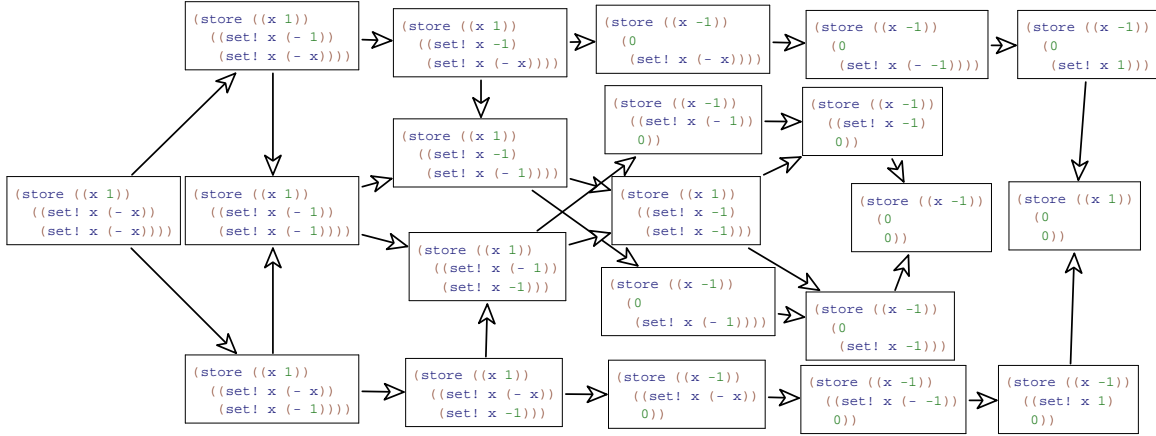


Figure 2. Interleavings possible with an erroneous unspecified-application-order model

$$\begin{aligned} \text{inert} &::= v^\circ \mid e \\ C &::= (\text{inert} \cdots C^\circ \text{inert} \cdots) \mid \dots \end{aligned}$$

$$\begin{aligned} (\text{store } (\dots) C[(\text{inert} \cdots e \text{inert} \cdots)]) &\rightarrow (\text{store } (\dots) C[(\text{inert} \cdots e^\circ \text{inert} \cdots)]) && \text{(UMARK)} \\ (\text{store } (\dots) C[(\text{lambda } (x \dots) e)^\circ v^\circ \dots]) &\rightarrow (\text{store } (\dots (x' v) \dots) C[e[x' \dots / x \dots]]) && \text{(UAPP)} \\ &&& \text{(#}x = \#v, \text{ each } x' \text{ fresh)} \\ (\text{store } (\dots) C[(-^\circ \lceil n^\circ \rceil)]) &\rightarrow (\text{store } (\dots) C[\lceil -n \rceil]) && \text{(UNEG)} \end{aligned}$$

Figure 3. Revisions to core Scheme to support unspecified application order

that effect, we introduce the non-terminal *inert* and the notion of a marked expression, denoted with the \circ superscript. (These marks are not an extension to the general term-rewriting framework — e° and C° are just alternate typesettings of (mark e) and (mark C .) Marks identify chosen expressions: only marked expressions may be reduced, and only one reducible marked expression may appear in any application at one time. The *inert* production stands for terms in which evaluation may not occur, i.e., unmarked expressions (those expressions we have not tried to evaluate yet) and marked values (those expressions we have already finished reducing). We add the UMARK reduction rule that marks an arbitrary unmarked expression in an application on the condition that every other expression is inert, and we modify the MAPP and MNEG rules to apply only to fully-marked applications, becoming the UAPP and UNEG rules.

Figure 4 (also generated by PLT Redex) shows how our new system evaluates the term from figure 2. The initial term appears in the center on the left. That term is an application, so the first reduction either marks the first sub-expression or the second. If the first subexpression is marked, evaluation continues down to the bottom of the figure, over to the right and back up to the middle. If the second is marked, evaluation proceeds up, over, and back to the middle. In both paths there are a few other application expressions to evaluate, leading to smaller separations. Eventually, all of the terms join back together and the final result in the store is I , as shown in the center on the right.

One should not take that example to mean that this language has any kind of confluence property, however. Consider this program:

```
((lambda (choice)
  ((lambda (x y) choice)
   (set! choice 1)
   (set! choice 2)))
 0)
```

It will either will return either I or 2 , depending on the order of evaluation. This is the way we want it; the model's nonconfluence reflects the underspecification of R^5RS Scheme rather than a technical bug in our model. It does, however, always make progress. We formalize this with the following theorem statement:

THEOREM 2.1. *For any closed program p in the language of figure 3, either $p \rightarrow p'$, where p' is also closed, $p \rightarrow e$ where e is some error indicator, or p is of the form $(\text{store } ((x v) \dots) v)$.*

Proof is contained in the first author's master's thesis [16].

This technique has other uses besides giving semantics for unspecified application evaluation orders. In general, it is useful for modeling any kind of delimited nondeterminism, where evaluation may proceed arbitrarily but only at certain points in a program. This is useful for modeling unspecified behaviors and for complex non-deterministic features such as threads.

3. Multiple return values

R^5RS Scheme provides a facility for expressions to evaluate to multiple or no values rather than just a single value. The procedure *values* builds multiple values and *call-with-values* accepts multiple values. Unlike tuples in SML and Haskell, multiple values are not themselves values. For example, this program

```
(define (f x) (values (+ x x) (* x x)))
(define (g x y) y)
(g (f 3))
```

produces an error, since procedure application expects each of its arguments to be a single value (and the result of f is two values). Instead, the programmer must use *call-with-values* to catch multiple values. It expects a thunk as its first argument, applies the thunk, catches any number of values that thunk produces, and applies them

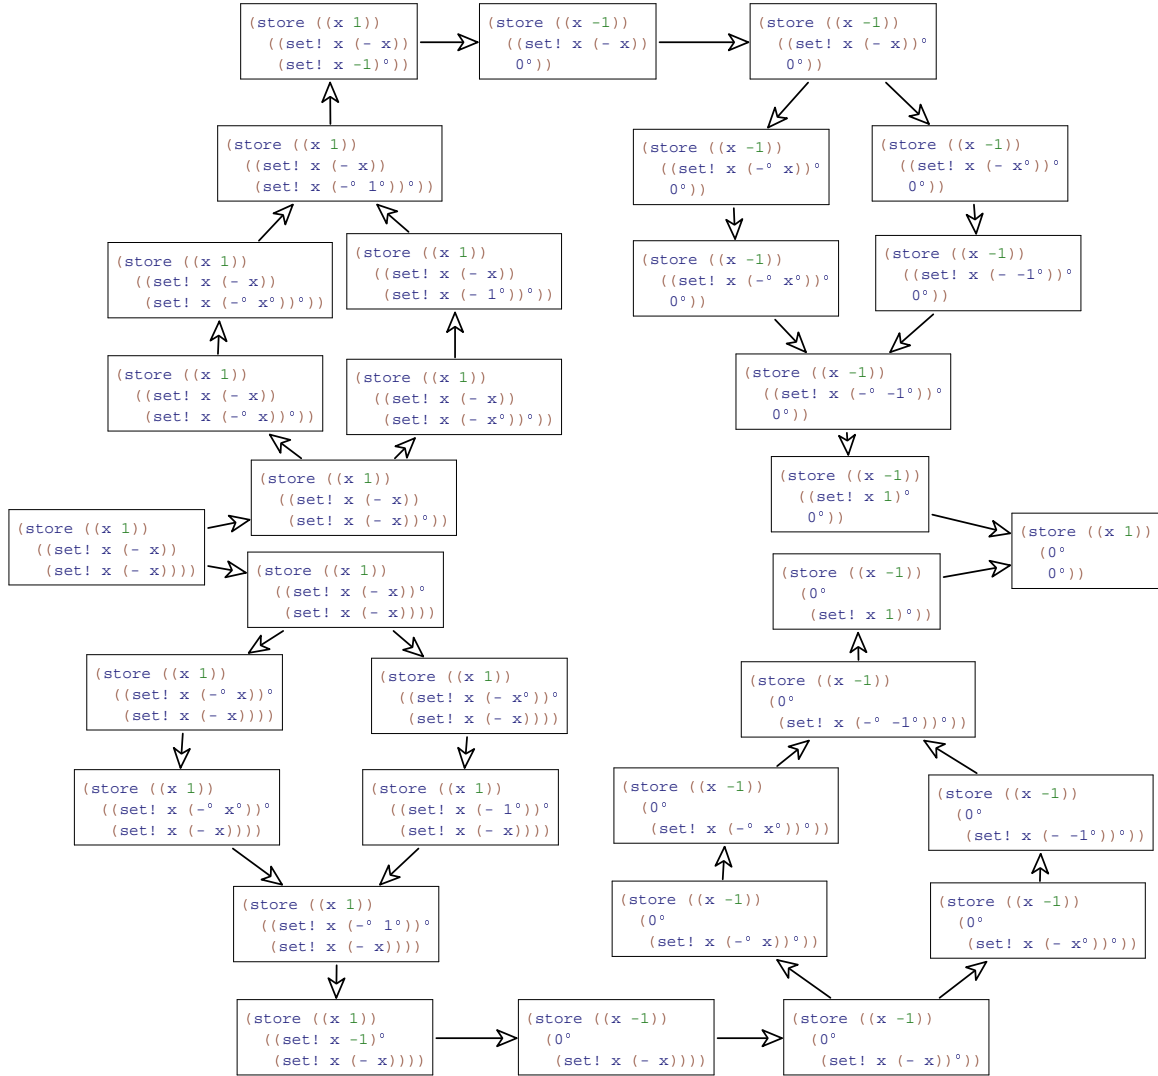


Figure 4. Evaluation in the unspecified-application-order model

to its second argument. So, a programmer could supply f 's results to g like this:

`(call-with-values (lambda () (f 3)) g)`

In addition, there is no difference between *values* applied to a single argument and that argument by itself, so `(g (values 6) (values 9))` is the same as `(g 6 9)`.

To model multiple values, R⁵RS Scheme's formal semantics models continuations as functions from an arbitrary number of values to a final answer. The informal semantics says that "except for continuations created with the *call-with-values* procedure, all continuations take exactly one value" [15, section 6.4]. The formal semantics reflects this by checking the opposite property: in every context that expects a single value, it uses a helper function, *single*, to ensure that only a single value appears. This indirect checking impacts the entire semantics: it requires every continuation to accept any number of arguments initially and requires a call to *single* at every point where a continuation would be restricted.

Our semantic model captures the difference between contexts that accept multiple values and contexts that reject multiple values

directly. Our strategy is distilled in figure 5. That figure contains a pure core Scheme extended with *values*, and **apply-values**, a syntactic form that has as its operands an expression that must evaluate to a procedure and another expression that may evaluate to any number of values, and calls the procedure with those values as arguments. We use **apply-values** in this section rather than *call-with-values* because the resulting model is clearer and both **apply-values** and *call-with-values* can be defined simply in terms of each other in R⁵RS Scheme:

`(define (call-with-values thunk f)
 (apply-values f (thunk)))`

`(define-syntax apply-values
 (syntax-rules ()
 [(- f vs-expr)
 (call-with-values (lambda () vs-expr) f)]))`

Our model uses a modest addition to the standard reduction-semantics formalism. We extend the notation so that holes have names (written as subscripts) but otherwise behave as unnamed

e	$::=$	$(e\ e\ \dots) \mid x \mid v \mid (\mathbf{apply-values}\ e\ e)$	
v	$::=$	$(\mathbf{lambda}\ (x\ \dots)\ e) \mid \mathit{values}$	
C	$::=$	$[\]_{\circ} \mid (v\ \dots\ C_{\circ}\ e\ \dots) \mid (\mathbf{apply-values}\ C_{\circ}\ e) \mid (\mathbf{apply-values}\ v\ C_{*})$	
C_{\circ}	$::=$	$[\]_{\circ} \mid C$	
C_{*}	$::=$	$[\]_{*} \mid C$	
$C_{\circ}[(\mathbf{lambda}\ (x\ \dots)\ e)\ v\ \dots]_{\circ}$	\rightarrow	$C_{\circ}[e[x\ \dots/v\ \dots]]$	(VAPP)
		$(\#v = \#x)$	
$C_{\circ}[(\mathbf{lambda}\ (x\ \dots)\ e)\ v\ \dots]_{\circ}$	\rightarrow	error: wrong number of arguments	(VAPPERR)
		$(\#v \neq \#x)$	
$C_{\circ}[(\mathbf{apply-values}\ v_1\ (\mathit{values}\ v_2\ \dots))]_{\circ}$	\rightarrow	$C_{\circ}[(v_1\ v_2\ \dots)]$	(VAPPVALS)
$C_{\circ}[v]_{*}$	\rightarrow	$C_{\circ}[(\mathit{values}\ v)]$	(VPROMOTE)
$C_{\circ}[(\mathit{values}\ v)]_{\circ}$	\rightarrow	$C_{\circ}[v]$	(VDEMOTE)
$C_{\circ}[(\mathit{values}\ v\ \dots)]_{\circ}$	\rightarrow	error: expected a single value	(VDEMOTERR)
		$(\#v \neq 1)$	

Figure 5. Pure core Scheme with multiple values

holes do. The context-matching syntax is now annotated with names as well, restricting legal decompositions to those where the hole has the same name.

In figure 5 we use this feature to give three distinct names to holes, indicated with subscripts. $[\]_{\circ}$ indicates a hole in which any expression should reduce to an element of v , $[\]_{*}$ indicates a hole in which any expression should reduce to $(\mathit{values}\ v\ \dots)$, and $[\]_{\circ}$ indicates a hole in which either result is acceptable. There are three parallel context nonterminals. The context C_{\circ} produces an element of v , C_{*} produces $(\mathit{values}\ v\ \dots)$, and C might produce either.

Since each subexpression of an application is expected to produce a single value, the evaluation context inside an application is C_{\circ} . For the same reason, the evaluation context for the first subexpression of **apply-values** is C_{\circ} . The evaluation context for the second subexpression, however, is C_{*} because it is expected to produce multiple values.

Since procedure applications (defined by the VAPP and VAPPERR reductions) and **apply-values** uses (defined by the VAPPVALS reduction) may produce a single value or $(\mathit{values}\ v\ \dots)$, they take place in $[\]_{\circ}$ holes. VPROMOTE, promotes a single value v to $(\mathit{values}\ v)$. Because of the subscript $*$ on the hole, it applies only when multiple values are expected. VDEMOTE demotes a single value inside values to just the value, and VDEMOTERR signals an error if values does not return exactly one value. These two rules apply only when a values expression appears where a single value is expected. All reductions take place in C_{\circ} to ensure that the final result of any program is a single value. If we wanted to allow any number of values as the final result of a program we could replace C_{\circ} with C_{*} in all rules.

To get a sense of how evaluation proceeds, consider this reduction sequence:

$$\begin{aligned}
& ((\mathbf{lambda}\ (y)\ y) \\
& \quad (\mathbf{apply-values}\ (\mathbf{lambda}\ (x)\ (\mathit{values}\ x))\ I)) \\
\rightarrow & ((\mathbf{lambda}\ (y)\ y) \\
& \quad (\mathbf{apply-values}\ (\mathbf{lambda}\ (x)\ (\mathit{values}\ x)) \\
& \quad \quad (\mathit{values}\ I))) \quad (\text{VPROMOTE}) \\
\rightarrow & ((\mathbf{lambda}\ (y)\ y) \\
& \quad ((\mathbf{lambda}\ (x)\ (\mathit{values}\ x))\ (\mathit{values}\ I))) \quad (\text{VAPPVALS}) \\
\rightarrow & ((\mathbf{lambda}\ (y)\ y) \\
& \quad ((\mathbf{lambda}\ (x)\ (\mathit{values}\ x))\ I)) \quad (\text{VDEMOTE})
\end{aligned}$$

$$\begin{aligned}
& \rightarrow ((\mathbf{lambda}\ (y)\ y)\ (\mathit{values}\ I)) \quad (\text{VAPP}) \\
& \rightarrow ((\mathbf{lambda}\ (y)\ y)\ I) \quad (\text{VDEMOTE}) \\
& \rightarrow I \quad (\text{VAPP})
\end{aligned}$$

First, the VPROMOTE applies and promotes I into $(\mathit{values}\ I)$ because it appears as the second argument of an **apply-values** expression. Then VAPPVALS applies, followed by VAPP. Then the term $(\mathit{values}\ I)$ is used as an argument to a procedure, so VDEMOTE applies and converts it to the single value I . Finally VAPP applies and the result is I .

The erroneous expression from the beginning of this section signals an error due to the VDEMOTERR rule.

$$\begin{aligned}
& (g\ (f\ 3)) \\
& \rightarrow \dots \\
& \rightarrow (g\ (\mathit{values}\ 3\ 9)) \\
& \rightarrow \mathbf{error:}\ \text{expected a single value}
\end{aligned}$$

The evaluation contexts and the three promotion and demotion rules are all that we need to add multiple values to the language. Furthermore, the extension of adding names to holes does not significantly complicate proof of progress for the system, and so we can prove the following theorem reasonably straightforwardly [16]:

THEOREM 3.1. *For any closed program p in the language of figure 5, either $p \rightarrow p'$, where p' is also closed, $p \rightarrow e$ where e is an error indicator, or p is of the form $(\mathbf{store}\ ((x\ v)\ \dots)\ v)$.*

Proof is contained in the first author's master's thesis [16].

The strategy described in this section can be used whenever the notion of a fully-evaluated subterm is different in different parts of a program. For instance, it can be used to model embedded sublanguages such as regular-expressions, format strings, and SQL commands, which could help develop theoretical underpinnings for work like Herman and Meunier's static analysis of embedded languages [14].

4. Quote and Eval

Scheme inherits the meta-programming facilities *eval* and **quote** from Lisp [22]. The **quote** operator turns a program into data and the *eval* procedure turns that data back into a program. When quoted, a program is represented as a list of lists and symbols, where lists represent parenthesized sequences and symbols represent identifiers. For example, **(quote (lambda (x) x))** is a three el-

$ \begin{aligned} e & ::= (e \ e \ \dots) \mid v \mid x \\ E & ::= [] \mid (v \ \dots \ E \ e \ \dots) \\ v & ::= (\mathbf{lambda} \ (x \ \dots) \ e) \mid (\mathbf{quote} \ sy) \\ & \quad \mid p \mid \mathbf{null} \mid n \mid \mathbf{prim} \mid \#\mathbf{t} \mid \#\mathbf{f} \\ \mathbf{prim} & ::= \mathit{eval} \mid \mathit{cons} \mid \mathit{car} \mid \mathit{cdr} \mid \mathit{eqv?} \\ p & ::= \text{pointers} \\ x & ::= \text{program variables} \\ & \quad (\text{members of } sy \text{ except } \mathbf{lambda}, \mathbf{quote}, \mathbf{ccons}) \end{aligned} $	$ \begin{aligned} s & ::= (s \ \dots) \mid n \mid sy \\ & \quad \mid (s \ \dots \ \mathbf{dot} \ sy) \mid (s \ \dots \ \mathbf{dot} \ n) \\ S & ::= [] \mid (e \ \dots \ S \ s \ \dots) \\ & \quad \mid (\mathbf{lambda} \ (x \ \dots) \ S) \\ & \quad \mid (\mathbf{ccons} \ v \ S) \mid (\mathbf{ccons} \ S \ s) \\ n & ::= \text{numbers} \\ sf & ::= (p \ (\mathit{cons} \ v \ v)) \\ sy & ::= \text{names of symbols} \\ & \quad (\text{identifiers except } \mathbf{dot}) \end{aligned} $
$ \begin{aligned} (\mathbf{store} \ (sf_1 \ \dots) \ E[(\mathit{cons} \ v_1 \ v_2)]) & \quad \rightarrow \quad (\mathbf{store} \ (sf_1 \ \dots \ (p \ (\mathit{cons} \ v_1 \ v_2))) \ E[p]) & \quad (\mathbf{ECONS}) \\ & \quad (p \ \text{fresh}) \\ (\mathbf{store} \ (sf_1 \ \dots \ (p \ (\mathit{cons} \ v_a \ v_d)) \ sf_2 \ \dots) \ E[(\mathit{car} \ p)]) & \quad \rightarrow \quad (\mathbf{store} \ (sf_1 \ \dots \ (p \ (\mathit{cons} \ v_a \ v_d)) \ sf_2 \ \dots) \ E[v_a]) & \quad (\mathbf{ECAR}) \\ (\mathbf{store} \ (sf_1 \ \dots \ (p \ (\mathit{cons} \ v_a \ v_d)) \ sf_2 \ \dots) \ E[(\mathit{cdr} \ p)]) & \quad \rightarrow \quad (\mathbf{store} \ (sf_1 \ \dots \ (p \ (\mathit{cons} \ v_a \ v_d)) \ sf_2 \ \dots) \ E[v_d]) & \quad (\mathbf{ECDR}) \\ (\mathbf{store} \ (sf_1 \ \dots) \ E[(\mathit{eqv?} \ p \ p)]) & \quad \rightarrow \quad (\mathbf{store} \ (sf_1 \ \dots) \ E[\#\mathbf{t}]) & \quad (\mathbf{EEQV1}) \\ (\mathbf{store} \ (sf_1 \ \dots) \ E[(\mathit{eqv?} \ p_1 \ p_2)]) & \quad \rightarrow \quad (\mathbf{store} \ (sf_1 \ \dots) \ E[\#\mathbf{f}]) & \quad (\mathbf{EEQV2}) \\ & \quad (p_1 \neq p_2) \\ (\mathbf{store} \ (sf \ \dots) \ E[(\mathbf{lambda} \ (x \ \dots) \ e) \ v \ \dots]) & \quad \rightarrow \quad (\mathbf{store} \ (sf \ \dots) \ E[e[x \ \dots / v \ \dots]]) & \quad (\mathbf{EAPP}) \\ & \quad (\#x = \#v) \\ (\mathbf{store} \ (sf \ \dots) \ S[(\mathbf{quote} \ sexp_1 \ sexp_2 \ \dots)]) & \quad \rightarrow \quad (\mathbf{store} \ (sf \ \dots) \ S[(\mathbf{ccons} \ sexp_1 \ (\mathbf{quote} \ sexp_2))]) & \quad (\mathbf{EQUOTESEQ}) \\ (\mathbf{store} \ (sf \ \dots) \ S[(\mathbf{quote} \ ())]) & \quad \rightarrow \quad (\mathbf{store} \ (sf \ \dots) \ S[\mathbf{null}]) & \quad (\mathbf{EQUOTENULL}) \\ (\mathbf{store} \ (sf \ \dots) \ S[(\mathbf{quote} \ n)]) & \quad \rightarrow \quad (\mathbf{store} \ (sf \ \dots) \ S[n]) & \quad (\mathbf{EQUOTENUM}) \\ (\mathbf{store} \ (sf \ \dots) \ S[(\mathbf{ccons} \ v_1 \ v_2)]) & \quad \rightarrow \quad (\mathbf{store} \ (sf \ \dots \ (p \ (\mathit{cons} \ v_1 \ v_2))) \ S[p]) & \quad (\mathbf{EQUOTEPAIR}) \\ & \quad (p \ \text{fresh}) \\ (\mathbf{store} \ (sf \ \dots) \ E[(\mathit{eval} \ v)]) & \quad \rightarrow \quad (\mathbf{store} \ (sf \ \dots) \ E[\mathcal{R} \ [[(sf \ \dots), v]]]) & \quad (\mathbf{EEVAL}) \end{aligned} $	
$ \begin{aligned} \mathcal{R} : (p \mapsto (\mathit{cons} \ v \ v)) \times v \mapsto s \\ \mathcal{R} \ [[S, \mathbf{null}]] & = () \\ \mathcal{R} \ [[S, n]] & = n \\ \mathcal{R} \ [[S, \#\mathbf{t}]] & = \#\mathbf{t} \\ \mathcal{R} \ [[S, \#\mathbf{f}]] & = \#\mathbf{f} \\ \mathcal{R} \ [[S, (\mathbf{quote} \ sy)]] & = sy \\ \mathcal{R} \ [[S, p]] & = \mathcal{C} \ [[\mathcal{R} \ [[v_a]], \mathcal{R} \ [[v_d]]]] \\ & \quad \text{where } S \text{ binds } p \text{ to } (\mathit{cons} \ v_a \ v_d) \end{aligned} $	$ \begin{aligned} \mathcal{C} : s \times s \mapsto s \\ \mathcal{C} \ [[sexp_1, (sexp_2 \ \dots)]] & = (sexp_1 \ sexp_2 \ \dots) \\ \mathcal{C} \ [[sexp_1, n]] & = (sexp_1 \ \mathbf{dot} \ n) \\ \mathcal{C} \ [[sexp_1, sy]] & = (sexp_1 \ \mathbf{dot} \ sy) \\ \mathcal{C} \ [[sexp_1, \#\mathbf{t}]] & = (sexp_1 \ \mathbf{dot} \ \#\mathbf{t}) \\ \mathcal{C} \ [[sexp_1, \#\mathbf{f}]] & = (sexp_1 \ \mathbf{dot} \ \#\mathbf{f}) \end{aligned} $

Figure 6. Core Scheme, extended with eval and quote

ement list whose first and third elements are symbols and whose second element is a list of one element:

$$\begin{aligned}
& (\mathit{cons} \ (\mathbf{quote} \ \mathbf{lambda}) \\
& \quad (\mathit{cons} \ (\mathit{cons} \ (\mathbf{quote} \ x) \ \mathbf{null}) \\
& \quad \quad (\mathit{cons} \ (\mathbf{quote} \ x) \ \mathbf{null})))
\end{aligned}$$

R⁵RS suggests (but does not require) that quoted data be allocated only once, before the program runs. In systems with that behavior (including all Scheme implementations we tested), this program returns $\#\mathbf{t}$:

$$\begin{aligned}
& ((\mathbf{lambda} \ (f) \ (\mathit{eqv?} \ (f) \ (f))) \\
& \quad (\mathbf{lambda} \ () \ (\mathbf{quote} \ (x))))
\end{aligned}$$

since the thunk passed as f returns the same result each time it is called.

Our core Scheme calculus for modeling *eval* and **quote** is shown in figure 6. (Note that this model simplifies R⁵RS Scheme’s *eval* procedure in that it does not accept an environment argument.) To ensure that a datum behind a **quote** is inserted into the store only once, the rewriting system is structured in two tiers roughly corresponding to “compile-time” and “run-time.” Initially, programs are just viewed as uncompiled s-expressions (elements of the s non-terminal; note that we write dotted pairs with **dot** rather than a period to avoid meta-circular confusion in our PLT Redex implementation), which in particular include programs with quoted lists. Reduction rules that apply to these uncompiled expressions do not evaluate them, but instead compile them into program expressions that do not contain quoted lists (elements of the e nonterminal).

Evaluation reductions only apply to a program after it has been completely compiled.

Each program consists of a store and an expression. Program expressions (e) can be applications, values, or identifiers. Evaluation contexts (E) dictate that evaluation takes place in a left to right order inside application expressions. The values (v) are procedures, quoted symbols, pointers (to cons cells), null, numbers, primitive operations, and booleans.

The first group of evaluation rules (from ECONS to EAPP) correspond to the language’s runtime semantics, and show how the list primitives and procedure application behave. ECONS models the application of *cons* to arguments by allocating a new pair in the store; and *car* and *cdr* select the first and second values in a pair by rules ECAR and ECDR. EEQV1 and EEQV2 give *eqv?*’s semantics; it compares pointers for literal syntactic equality (and, for this language, operates only on pairs). As in the previous systems we have presented, procedure application is modeled by rule EAPP as substitution. Since each reduction takes place in an *evaluation* (rather than *compilation*) context, they will only apply to programs that are completely compiled.

The second group of rules (from EQUOTESEQ to EQUOTEPAIR) apply at compile-time and show how to compile a program by rewriting quoted constants into locations in the store. If those rules used the E context and quoted s-expressions were legal expressions, **quote** would merely be a short-hand notation for building lists at run-time and the above program would return $\#\mathbf{f}$, which would not capture our intended semantics.

$ \begin{aligned} p &::= (\mathbf{store} ((x\ v) \dots) (\mathbf{dw} (dws \dots) e)) \\ e &::= \dots \mid (\mathbf{push} (x\ e\ e)) \mid (\mathbf{pop}) \\ v &::= \dots \mid \mathit{dynamic-wind} \mid \mathit{call/cc} \\ dws &::= (x\ e) \end{aligned} $	$ \begin{aligned} PC &::= (\mathbf{store} ((x\ v) \dots) DC) \\ DC &::= (\mathbf{dw} ((dws \dots) C)) \\ C &::= (\text{as in figure 1}) \end{aligned} $
$ PC[(\mathit{dynamic-wind} (\mathbf{lambda} ()\ e_1) (\mathbf{lambda} ()\ e_2) (\mathbf{lambda} ()\ e_3))] $	$ \rightarrow PC[(\mathbf{begin}\ e_1 (\mathbf{push}\ (x_1\ e_1\ e_3)) (\mathbf{lambda}\ (x_2) (\mathbf{begin}\ (\mathbf{pop})\ e_3\ x_2)) e_2))] \quad (\text{DWWIND}) $
$ PC[(\mathbf{dw} (dws \dots) C[(\mathbf{push}\ x_2\ e_1\ e_2)])] $	$ \rightarrow PC[(\mathbf{dw} (dws \dots (x_2\ e_1\ e_2)) C[0])] \quad (\text{DWPUSH}) $
$ PC[(\mathbf{dw} (dws_1 \dots dws_n) C[(\mathbf{pop})])] $	$ \rightarrow PC[(\mathbf{dw} (dws_1 \dots) C[0])] \quad (\text{DWPOP}) $
$ PC[(\mathbf{dw} (dws_1 \dots) C[(\mathit{call/cc}\ v_1)])] $	$ \rightarrow PC[(\mathbf{dw} (dws_1 \dots) C[(v_1 (\mathbf{lambda}\ (x) (\mathbf{throw}\ (dws_1 \dots) C[x])])])] \quad (\text{DWCALLCC}) $
$ PC[(\mathbf{dw} (dws_1 \dots) C[(\mathbf{throw}\ (dws_2 \dots) e_1)])] $	$ \rightarrow PC[(\mathbf{dw} (dws_2 \dots) C[(\mathbf{begin}\ \mathcal{F} [(dws_2 \dots), (dws_1 \dots) e_1])])] \quad (\text{DWTHROW}) $
$ \mathcal{F} [(x_1\ e_1\ e_2)\ dws_1 \dots, (x_1\ e_3\ e_4)\ dws_2 \dots] $	$ = \mathcal{F} [(dws_1 \dots), (dws_2 \dots)] $
$ \mathcal{F} [(x_1\ e_1\ e_2) \dots, (x_2\ e_3\ e_4) \dots] $	$ = (\mathbf{begin}\ e_2 \dots_r\ e_3 \dots) \quad (x_1 \neq x_2) $

Figure 7. Additions to figure 1 to support call/cc and dynamic-wind

Instead, the second group of rewriting rules eliminate **quote**, turning s-expressions into Scheme programs. Though we have presented them second, these rules will actually come first in reduction sequences, making reduction sequences follow a two-phase pattern where the EQUOTE rules apply in the first phase and the evaluation rules apply in the second phase. Intuitively, programs in this first phase are arbitrary s-expressions and values are Scheme programs, whereas second-phase programs are Scheme expressions and values are Scheme values. This parallelism can be seen particularly clearly in the definition of the evaluation contexts for application expressions. In S , a rewrite may occur once all of the s-expressions to the left have become Scheme programs. In E , a rewrite may occur once all of the expressions to the left have become values. So, for the program above, the only rewriting rules that apply are those that rewrite the thunk's body. Once it contains only a pointer to a store value, the outer application can proceed.

To model *eval*, we use a technique similar to Muller's *reify* [18]. The \mathcal{R} metafunction accepts a value and turns it back into a program (the \mathcal{C} function is used by \mathcal{R} ; it is just the syntactic analogue of *cons*). Once \mathcal{R} completes, evaluation continues as usual. Of course, reification may produce an s-expression containing **quote**. In that case, the quote rules apply and put quoted data into the store before evaluation continues.¹

¹Most Scheme systems share quoted data even across calls to *eval*. For example, our semantics produces $\#f$ for the following program, but most Schemes produce $\#t$.

$$\begin{aligned}
&((\mathbf{lambda}\ (f) \\
&\quad (eqv?\ (f) \\
&\quad\quad (\mathit{eval}\ (\mathbf{cons}\ 'quote\ (\mathbf{cons}\ (f)\ '()))))) \\
&(\mathbf{lambda}\ ()\ '(x)))
\end{aligned}$$

We can adapt the definition of \mathcal{R} to handle this by special handling of quoted forms during reification:

$$\mathcal{R} [(S, p_1)] = v \quad \text{if } S \text{ maps } p_1 \text{ to } (\mathbf{cons}\ (\mathbf{quote}\ quote)\ p_2) \text{ and maps } p_2 \text{ to } (\mathbf{cons}\ v\ '()).$$

which causes our semantics to produce $\#t$ for the above example, but this technique does not scale to a full Scheme that includes macros.

As mentioned above, the *eval* we present here and in section 6 is not as full-featured as the *eval* of the R⁵RS informal description because it does not accept an environment argument. Modeling an *eval* that took an environment argument would be somewhat more involved but would essentially require only running *eval*ed programs in an alternate store.

The technique used in this section applies generally to languages in which computation of a term proceeds in multiple phases that must be considered together — it is not sufficient in our case to write a preprocessor that moves quoted data in a program into the store because that program could call *eval* at runtime. Scheme's macros are similar in this respect, so the technique shown here could be used as a basis for modeling them. Staged and partial evaluation could also be modeled using this technique.

5. Call/cc and dynamic-wind

Scheme's *dynamic-wind* feature for annotating the dynamic extent of a procedure call with entry and exit code that run whenever the program flows into or out of that extent, either through normal program evaluation or through the invocation of continuation objects made by *call/cc* (the latter situation being the more interesting one, of course). Unfortunately, though *dynamic-wind* has a large impact on the meaning of continuation objects *call/cc* produces, the R⁵RS formal semantics does not include any mention of it and models *call/cc* without respect to it. Here we will show how it works in the context of the core Scheme with mutation presented in section 2. Our strategy for modeling these new features is based heavily on earlier treatments [4, 10, 12].

The language in figure 7 consists of the core Scheme with mutation as shown in figure 1 augmented with *call/cc* and *dynamic-wind*. The basic strategy we take is to maintain a stack of all *dynamic-wind* calls entered but not yet exited, which we call the dynamic-wind stack. When we capture a continuation, we record the current dynamic-wind stack. When we throw to a continuation object, we use the difference between the current dynamic-wind stack and that recorded dynamic-wind stack to determine which *pre* and *post* thunks need to be called.

p	::=	(store ((ptr sv) ...) (dw (dws ...) e))	dws	::=	(x cp cp)
e	::=	(e e ...) (if e e e) (if e e) (set! x e) (begin e e ...)	sv	::=	v (%cons v v) lam mulam
		(throw x dws ... EC[e]) (push (x e e) e) (pop e)	s	::=	(s ...) (s ... dot nss) nss
		lam mulam v x	nss	::=	number #t #f [variable except dot]
lam	::=	(lambda (x ...) e e ...)	SC	::=	[] (e ... SC s ...)
$mulam$::=	(lambda (x ... dot x) e e ...)			(if SC s s) (if e SC s) (if e e SC)
v	::=	fun nonfun			(if SC s) (if e SC)
fun	::=	cp mp %cons %null? %pair?			(set! x SC)
		%car %cdr %set-car! %set-cdr! %list			(begin SC s ...) (begin e e ... SC s ...)
		%+ %- %/ %* %call/cc			(throw x dws ... SC) (push (x SC s) s)
		%dynamic-wind %values %call-with-values			(push (x e SC) s) (push (x e e) SC) (pop SC)
		%eqv? %apply %eval			(lambda (x ...) SC s ...)
$nonfun$::=	pp number %null #t #f			(lambda (x ...) e e ... SC s ...)
		(quote symbol) unspecified			(lambda (x ... dot x) SC s ...)
					(lambda (x ... dot x) e e ... SC s ...)
					(ccons SC s) (ccons v SC)
PC	::=	(store ((ptr sv) ...) DC)	var	::=	[variable except dot and keywords]
DC	::=	(dw (dws ...) EC _o)	x	::=	[variable names]
EC	::=	[] (inert ... EC _o ^o inert ...)	pp	::=	[pair pointers]
		(if EC _o e e) (if EC _o e) (set! x EC _o)	cp	::=	[closure pointers]
		(begin EC e e ...)	mp	::=	[μ closure pointers]
		(%call-with-values ^o (cww-mark EC*) v ^o)	ptr	::=	x pp cp mp
EC_o	::=	[] _o EC			
EC_*	::=	[] _* EC			
$inert$::=	e v ^o			

Figure 8. Grammar for full Scheme semantics

That strategy is formally encoded in three parts. First, we add a dynamic-wind stack to each program context. It contains one dynamic context frame (dws) for each annotated dynamic extent in which the current evaluation is taking place. A dynamic context frame is a triple consisting of a unique identifier and the pre and $post$ thunks of the corresponding *dynamic-wind* call. The unique identifier allows us to disambiguate multiple dynamic evaluations of the same syntactic appearance of a *dynamic-wind* expression. Second, we add the primitive procedure value *dynamic-wind* to the set of values, which expects each of its three arguments to evaluate to a thunk. Then using the DWWIND rule it invokes its pre thunk, pushes a dynamic context frame onto the stack with a fresh identifier and its own pre and $post$ thunks, evaluates its second thunk, pops its dynamic context frame off the stack, evaluates its $post$ thunk, and finally returns the value its second thunk evaluated to. To allow the program to manipulate the stack, we introduce the **push** and **pop** forms and their associated reduction rules DWPUSH and DWPOP. The former pushes a new dynamic context frame onto the end of the stack, and the latter pops the last context frame off the stack (and then evaluates to the trivial value θ , which is never used). These two forms are intended to be used only by *dynamic-wind*, never by the programmer directly.

The third piece is *call/cc*. When *call/cc* is called, the DW-CALLCC rule builds a continuation object that consists of a procedure of one argument, a fresh identifier we will call x . That procedure's body is a **throw** form that consists of the current dynamic stack and the expression formed by plugging x into the hole of the evaluation context where the application of *call/cc* itself was found. A **throw** form is itself evaluated using the DWTHROW rule by discarding the evaluation context in which it was found, replacing the dynamic stack with its own stored dynamic stack, and replacing the entire program body with a specially-constructed **begin** expression built by the \mathcal{T} metafunction (where T stands for "trim," because it trims away the common context frames leaving only the suffixes whose pre- or post-thunks need to be executed). That function compares its first argument, the dynamic-wind stack of the dynamic context being exited, with its second argument, the dynamic-wind stack of the context being entered. The first rule in its definition

simply discards any common prefix the two stacks may have, which correspond to dynamic extents that were never left or entered during the transitions from the time the continuation object was created and the time it was invoked. Then, once the two stacks have been trimmed to the point where they have distinct heads, the metafunction produces a **begin** expression consisting of applications of all the $post$ thunks from \mathcal{T} 's first argument, invoked in order, followed by all the pre thunks from \mathcal{T} 's second argument, invoked in reverse order (which we indicate with the special notation \dots_r , indicating a sequence being expanded out backwards).

6. Operational semantics for R⁵RS Scheme

This section combines the techniques from sections 2 through 5 with other known techniques for modeling programming languages to build a model of R⁵RS Scheme that includes all the features from those sections along with **if** and booleans, mathematical operations (but not the numeric tower), list constructors, selectors, mutators and predicates, μ -lambda procedures², *apply*, and object identity-based equivalence. Although this section appears large and complex at first, it is mostly just a simple combination of the previous four sections.

This specification is executable, and the figures presented in this section were automatically generated from the source code that implements the specification. Since an executable specification was an explicit goal of our work, we have made some modeling choices that may not be obvious at first. For example, there are many expressions whose return values are explicitly unspecified in the R⁵RS Scheme document, such as the result of a **set!** expression. A non-executable specification might model the evaluation of those expressions using the rule schema

$$\forall v. PC[unspecified] \rightarrow PC[v]$$

²Procedures declared with an improper list of formal arguments described in section 4.1.4 of the Report that accept an arbitrary number of arguments beyond a certain minimum. The name dates back at least to Indiana University's Scheme 84 system where MULAMBDA was a keyword used to declare procedures that accepted any number of arguments and collected them in a list [11].

meaning that an unspecified term reduces to any value. Instead, we model unspecified results with a special value *unspecified* that has no associated reduction rules and will cause programs that inspect it to get stuck. We also chose to ignore out-of-memory errors. These would be easy to add at the expense of a additional clutter when visualizing traces: reductions from each allocation site to the out-of-memory error would suffice.

6.1 Grammar

The grammar for R⁵RS Scheme programs is given in figure 8. In that figure, a program (given by the *p* nonterminal) consists of a store, a dynamic-wind stack, and an expression. The *e* nonterminal gives expressions, which in addition to standard Scheme core forms can be **throw**, **push** and **pop**, as in section 5. Values (*v*) are either procedures or non-procedure values, but notice that syntactic **lambda** terms are not values themselves. Instead, procedure values (*fun*) can be references to procedures in the store (*cp* and *mp*) or the built-in procedures, while the **lambda** form, as we will see, places new procedure values into the store when evaluated. Non-procedure values (*nonfun*) include pair pointers, numbers, *null*, booleans, symbols, and the unspecified value.

As in section 4, we write dotted pairs (as in the parameter list of a μ -lambda) with **dot** rather than a period to avoid meta-circular confusion in our PLT Redex implementation.

Section 6 of the R⁵RS Scheme specification indicates that primitive procedures are bound to names in the initial environment, but that those names can be mutated during the course of a program. To model that, we use special names with *##* prefixes to indicate the actual built-in procedures, and we bind those values to their *##*-less names in the initial store:

```
(store ((list ##list) (cons ##cons) (car ##car) (cdr ##cdr)
      (pair? ##pair?) (null ##null) (null? ##null?)
      (set-car! ##set-car!) (set-cdr! ##set-cdr!)
      (+ ##+) (- ##-) (/ ##/) (* ##*)
      (call/cc ##call/cc) (dynamic-wind ##dynamic-wind)
      (values ##values) (call-with-values ##call-with-values)
      (eqv? ##eqv?) (apply ##apply) (eval ##eval)) ...)
```

There are three different contexts we will make use of: program evaluation contexts, dynamic-wind contexts, and expression contexts. Each program evaluation context (PC) contains a store, and a dynamic-wind context. Each dynamic-wind context (DC) contains a dynamic-wind stack and an expression context. Expression contexts (EC) are the contexts in which program evaluation takes place; they allow evaluation in marked sub-expressions of an application (as in section 2), the test positions of **if** expressions, in **set!** expressions and in the first position in a **begin** (as long as there are at least two expressions in the **begin**). The evaluation context for *##call-with-values* is explained in section 6.7. The EC_o and EC_{*} evaluation contexts and *inert* work like C_o and C_{*} and *inert* from section 3.

The *dws* non-terminal corresponds to one frame of dynamic-wind context information and its use is explained in section 5. The *sv* non-terminal generates values that appear in the store.

S-expressions (*s* and *nss*) and s-expression contexts (SC) correspond to s-expressions and s-expression contexts from section 4. There are more possible s-expression contexts in the full language because there are more possible syntactic forms.

Finally, the *x* nonterminal represents both program variables and binding locations, and the *pp*, *cp*, and *mp* nonterminals represent pointers to pairs, fixed-arity procedures, and variable-arity procedures, respectively. The *ptr* non-terminal is a short-hand for terms that index into the store. One subtle point here is that the *v* production produces *pp*, *cp*, and *mp* but not *x*. Those variables are not included because free variables are not values and bound variables

have to be dereferenced before use, so neither qualifies as an irreducible value.

6.2 Relations

In the remaining figures, we will make heavy use of various reduction relation symbols. The basic reduction relation we will use is \rightarrow , which indicates that the program term on the left reduces in one step to the term on the right. We also use two other relations to aid in the system's readability, defined in terms of the \rightarrow relation:

- $(e_1 \dots e_n) \mapsto^\circ e'$ iff $\text{PC}[(e_1^\circ \dots e_n^\circ)] \rightarrow \text{PC}[e']$
The application on the left reduces to the term on the right in a program context, assuming that all of the expressions in the application are marked.
- $e \rightarrow^e \text{error}: s$ iff $\text{PC}[e] \rightarrow \text{error}: s$
The term on the left signals an error, halting the program immediately.

6.3 Basic syntactic forms

Figure 9 shows rules for the basic syntactic forms. For the **if** form, if the test position evaluates to anything other than *##f*, the term rewrites to its “then” subexpression. If the test position evaluates to *##f*, it rewrites to its “else” subexpression, if present, *unspecified* otherwise. For the **begin** form, the evaluation contexts defined in figure 8 ensure that the first term of a **begin** expression containing at least two expressions is evaluated fully; then these rules cause **begin** expression that consists of a fully-evaluated value followed by one or more expressions to rewrite to a new **begin** expression with the initial value dropped. These rules also specify that a **begin** form with only a single expression reduces immediately to that expression, even if that expression is not yet a value.

Because our model does not take into account R⁵RS Scheme's numeric tower, we model its numeric operations in terms of true mathematical functions. We assume that we can identify the true number represented by each numeric term and model each numeric procedure by performing the appropriate mathematical operation on those true numbers: + is modeled by summation on the represented numbers, * is modeled by multiplication, and so on.

6.4 Cons and cons-cell mutation

The rules for constructing new *cons* cells are given in figure 10. Since all cons cells are mutable and therefore can be distinguished even when they hold identical values, we cannot allow *##cons v v* to be a value itself. Instead, the *##cons* rule introduces a new pair into the store and reduces to a pointer to that new pair. The *##list v₁ ...* rule rewrites to *((lambda x x) v₁ ...)*, taking advantage of the μ -lambda application rules described in the next subsection.

Figure 11 gives rules for *car* and *cdr*. Application of either procedure to a pair pointer rewrites to the contents of the appropriate field in the pair being pointed to. If either selector is applied to a non-pair value, the term rewrites to an error message.

The predicates in figure 12 are similarly simple. The *##pair?* procedure reduces to *##t* if its argument is identifiable as a pair pointer and *##f* otherwise. The *##null?* procedure reduces to *##t* if and only if it is supplied with the built-in null value.

Figure 13 gives rules for *set-car!* and *set-cdr!*, for cons-cell mutation. The *##set-car!* and *##set-cdr!* rules are the same as the *car* and *cdr* rules, respectively, except that instead of reducing to the current value of appropriate component of the pair being pointed to, they replace that component with the given replacement then rewrite to an unspecified value.

6.5 Procedures and assignable variables

The rules in figure 14 handle variable lookup and variable assignment: a binding pointer is replaced with its value in the store when

$\text{PC}[\text{if } v_1 e_1 e_2]$	\rightarrow	$\text{PC}[e_1]$ ($v_1 \neq \#f$)	$\text{PC}[(\text{begin } v e_1 e_2 \dots)]$	\rightarrow	$\text{PC}[(\text{begin } e_1 e_2 \dots)]$
$\text{PC}[\text{if } \#f e_1 e_2]$	\rightarrow	$\text{PC}[e_2]$	$\text{PC}[(\text{begin } e_1)]$	\rightarrow	$\text{PC}[e_1]$
$\text{PC}[\text{if } v_1 e_1]$	\rightarrow	$\text{PC}[e_1]$ ($v_1 \neq \#f$)	$(+ \lceil n \rceil \dots)$	\mapsto°	$\lceil \Sigma n \dots \rceil$
$\text{PC}[\text{if } \#f e_1]$	\rightarrow	$\text{PC}[\text{unspecified}]$	$(- \lceil n_1 \rceil \lceil n_2 \rceil \dots)$	\mapsto°	$\lceil n_1 - (\Sigma n_2 \dots) \rceil$
			$(- \lceil n \rceil)$	\mapsto°	$\lceil -n \rceil$
			$(* \lceil n \rceil \dots)$	\mapsto°	$\lceil \Pi n \dots \rceil$
			$(/ \lceil n_1 \rceil \lceil n_2 \rceil \dots)$	\mapsto°	$\lceil n_1 / (\Pi n_2 \dots) \rceil$
			$(/ \lceil n_1 \rceil)$	\mapsto°	$\lceil I / n_1 \rceil$

Figure 9. Basic syntactic forms

$(\text{store } ((ptr_1 sv_1) \dots)$ $\text{DC}[(\# \%cons^\circ v_{car}^\circ v_{cdr}^\circ)])$	\rightarrow	$(\text{store } ((ptr_1 sv_1) \dots (p_i (\# \%cons v_{car} v_{cdr})))$ $\text{DC}[p_i]$ (p_i fresh)
$\text{PC}[(\# \%list^\circ v_1^\circ \dots)]$	\rightarrow	$\text{PC}[(\text{lambda } (\text{dot } l) l)^\circ v_1^\circ \dots]$

Figure 10. List constructors

$(\text{store } ((ptr_1 sv_1) \dots$ $(pp_i (\# \%cons v_{car} v_{cdr}))$ $(ptr_{i+1} sv_{i+1}) \dots)$ $\text{DC}[(\# \%car^\circ pp_i^\circ)])$	\rightarrow	$(\text{store } ((ptr_1 sv_1) \dots$ $(pp_i (\# \%cons v_{car} v_{cdr}))$ $(ptr_{i+1} sv_{i+1}) \dots)$ $\text{DC}[v_{car}])$	$(\# \%null? \# \%null)$	\mapsto°	$\#t$
$(\# \%car^\circ v_i^\circ)$	\rightarrow^e	error: can't take car of non-pair ($v_i \notin pp$)	$(\# \%null? v_i)$	\mapsto°	$\#f$ ($v_i \neq \# \%null$)
$(\text{store } ((ptr_1 sv_1) \dots$ $(pp_i (\# \%cons v_{car} v_{cdr}))$ $(ptr_{i+1} sv_{i+1}) \dots)$ $\text{DC}[(\# \%cdr^\circ pp_i^\circ)])$	\rightarrow	$(\text{store } ((ptr_1 sv_1) \dots$ $(pp_i (\# \%cons v_{car} v_{cdr}))$ $(ptr_{i+1} sv_{i+1}) \dots)$ $\text{DC}[v_{cdr}])$	$(\# \%pair? pp)$	\mapsto°	$\#t$
$(\# \%cdr^\circ v_i^\circ)$	\rightarrow^e	error: can't take cdr of non-pair ($v_i \notin pp$)	$(\# \%pair? v_i)$	\mapsto°	$\#f$ ($v_i \notin pp$)

Figure 11. List accessors

Figure 12. List predicates

dereferenced, and mutation of a binding pointer is represented by replacing the value pointed to by the update. **lambda** is the only binding form in this semantics, so the rules for procedure calls are the only ones that introduce new bindings. Procedure calls are modeled by two features: closure introduction and procedure application.

The rules in figure 15 govern the introduction of closure values into the store. Like cons cells, procedures are not values, but pointers to them are; procedures are modeled this way so that we can model *eqv?* more accurately. The allocation rule for fixed-arity procedures is straightforward. The allocation for μ -lambda procedures always puts two procedures into the store: a stub μ -lambda procedure whose body contains a call to an ordinary procedure, and an ordinary procedure that contains the original μ -lambda's body expressions.

The reason for arranging the system this way is so that when a μ -lambda procedure is applied, we can rewrite it into a corresponding call to the fixed-arity code pointer and thereby use the same reduction for both kinds of applications. The rules in figure 16 show this and the rest of the rules for application in detail. The first rule shows how marks are placed in applications, which is just as in section 2. Application of a procedure pointer to arguments is mod-

eled by creating one new binding pointer in the store per formal argument where the value being pointed to by each pointer is the argument supplied in the appropriate position, and rewriting to the procedure's body with these new bound-variable pointers substituted for occurrences of the formal arguments.

Application of a μ -lambda allocates a list for its extra arguments, applies the initial portion of the arguments as usual, and applies the extra arguments into the last argument of the procedure that actually contains the body expressions. The function \mathcal{L} used here is a metafunction that builds the syntax of a *cons*-list from its arguments:

$$\begin{aligned} \mathcal{L} \llbracket x y \dots \rrbracket &= (\# \%cons x \mathcal{L} \llbracket y \dots \rrbracket) \\ \mathcal{L} \llbracket \rrbracket &= \# \%null \end{aligned}$$

The last rules specify the behavior of Scheme's *apply* procedure which accepts a procedure and an arbitrary number of arguments, the last of which must be a list. It calls the procedure with the arguments and the contents of the list as subsequent arguments. To model it, the first two *#%apply* rules flatten out the argument list and, when the list is exhausted, reduce to a normal application.

$\begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (pp_i (\#\%cons\ v_{car}\ v_{cdr})) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[(\#\%set-car!^\circ\ pp_i^\circ\ v_{new}^\circ)] \\ &(\#\%set-car!^\circ\ v_1^\circ\ v^\circ) \end{aligned} \quad \rightarrow \quad \begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (pp_i (\#\%cons\ v_{new}\ v_{cdr})) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[\text{unspecified}] \\ &\text{error: can't set-car! on a non-pair} \\ &(v_1 \notin pp) \end{aligned}$	$\begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (pp_i (\#\%cons\ v_{car}\ v_{cdr})) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[(\#\%set-cdr!^\circ\ pp_i^\circ\ v_{new}^\circ)] \\ &(\#\%set-cdr!^\circ\ v_1^\circ\ v^\circ) \end{aligned} \quad \rightarrow \quad \begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (pp_i (\#\%cons\ v_{car}\ v_{new})) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[\text{unspecified}] \\ &\text{error: can't set-cdr! on a non-pair} \\ &(v_1 \notin pp) \end{aligned}$
--	--

Figure 13. Cons cell mutation

$\begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (x_i sv_i) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[x_i] \\ &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (x_i sv_i) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[(\text{set! } x_i\ v_{new})] \end{aligned} \quad \rightarrow \quad \begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (x_i sv_i) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[sv_i] \\ &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (x_i v_{new}) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[\text{unspecified}] \end{aligned}$	$\begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (x_i sv_i) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[(\text{set! } x_i\ v_{new})] \end{aligned} \quad \rightarrow \quad \begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (x_i sv_i) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[sv_i] \\ &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (x_i v_{new}) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[\text{unspecified}] \end{aligned}$
--	--

Figure 14. Variable mutation and lookup

$\begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad DC[lam_i]) \\ &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad DC[(\text{lambda } (x_1 \dots \text{dot } x_r)\ e_1\ e_2 \dots)]) \end{aligned} \quad \rightarrow \quad \begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots (cp_i\ lam_i)) \\ &\quad DC[cp_i]) \\ &(cp_i\ \text{fresh}) \\ &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (mp_i (\text{lambda } (x_1 \dots \text{dot } x_r)\ (cp_i\ x_1 \dots x_r))) \\ &\quad (cp_i (\text{lambda } (x_1 \dots \text{dot } x_r)\ e_1\ e_2 \dots))) \\ &(mp_i,\ cp_i\ \text{fresh}) \end{aligned}$	
--	--

Figure 15. Procedure introduction

$\begin{aligned} &PC[(\text{inert}_1 \dots e_i\ \text{inert}_{i+1} \dots)] \\ &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (cp_i (\text{lambda } (x_1 \dots) e_{body1}\ e_{body2} \dots)) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[(cp_i^\circ\ v_{arg1}^\circ \dots)] \end{aligned} \quad \rightarrow \quad \begin{aligned} &PC[(\text{inert}_1 \dots e_i^\circ\ \text{inert}_{i+1} \dots)] \\ &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (cp_i (\text{lambda } (x_1 \dots) e_{body1}\ e_{body2} \dots)) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots \\ &\quad (x_{arg2}\ v_{arg1}) \dots)) \\ &DC[(\text{begin } e_{body1}\ e_{body2} \dots)[x_1 \dots /x_{arg2} \dots]]) \\ &(\#x_{arg} = \#v_{arg}, x_{arg2} \dots \text{fresh}) \end{aligned}$	
$\begin{aligned} &(\text{store } ((ptr sv) \dots \\ &\quad (cp_i (\text{lambda } (x_1 \dots) e\ e \dots)) \\ &\quad (ptr sv) \dots)) \\ &DC[(cp_i^\circ\ v_{arg1}^\circ \dots)] \end{aligned} \quad \rightarrow \quad \text{error: arity mismatch} \\ &(\#x_{arg} \neq \#v_{arg})$	
$\begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (mp_i (\text{lambda } (x_1 \dots \text{dot } y)\ (cp_t\ x_1 \dots y))) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[(mp_i\ v_{n1}^\circ \dots v_R^\circ \dots)] \end{aligned} \quad \rightarrow \quad \begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (mp_i (\text{lambda } (x_1 \dots \text{dot } y)\ (cp_t\ x_1 \dots y))) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[(cp_t\ v_{n1}^\circ \dots \mathcal{L} [\ v_R^\circ \dots])]) \\ &(\#x = \#v_n) \end{aligned}$	
$\begin{aligned} &(\text{store } ((ptr sv) \dots \\ &\quad (mp_i (\text{lambda } (x_1 \dots \text{dot } x)\ (cp\ x \dots))) \\ &\quad (ptr sv) \dots)) \\ &DC[(mp_i^\circ\ v_{arg1}^\circ \dots)] \end{aligned} \quad \rightarrow \quad \text{error: too few arguments} \\ &(\#x_{arg} < \#v_{arg})$	
$\text{nonfun}^\circ\ v^\circ \dots \quad \rightarrow^e \quad \text{error: can't apply non-function}$	
$\begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (pp_i (\#\%cons\ v_{car}\ v_{cdr})) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[(\#\%apply^\circ\ v_f^\circ\ v_{arg1}^\circ \dots pp_i^\circ)] \end{aligned} \quad \rightarrow \quad \begin{aligned} &(\text{store } ((ptr_1 sv_1) \dots \\ &\quad (pp_i (\#\%cons\ v_{car}\ v_{cdr})) \\ &\quad (ptr_{i+1}\ sv_{i+1}) \dots)) \\ &DC[(\#\%apply^\circ\ v_f^\circ\ v_{arg1}^\circ \dots v_{car}^\circ\ v_{cdr}^\circ)] \end{aligned}$	
$\text{(\#\%apply } v_f\ v_{arg1} \dots \#\%null) \quad \mapsto^\circ \quad (v_f\ v_{arg1} \dots)$	
$\text{(\#\%apply}^\circ\ v_f^\circ\ v_{arg1}^\circ \dots v_{last}^\circ) \quad \rightarrow^e \quad \text{error: apply must take a list as its last argument} \\ (v_{last} \notin pp \cup \{\#\%null\})$	

Figure 16. Procedure application

```

(store ((ptrs svs) ...)
  (dw (dws1 ...))
  EC1[(#%call/cco v1o)])
→ (store ((ptrs svs) ...)
  (dw (dws1 ...))
  EC1[(v1 (lambda (dot args)
    (throw xk dws1 ...
    EC1[(begin xk (#%apply #%values args))]))])])
    (x, xk fresh)

(store ((ptrs svs) ...)
  (dw (dws1 ...))
  EC1[(#%dynamic-windo cp1o cp2o cp3o)])
→ (store ((ptrs svs) ...)
  (dw (dws1 ...))
  EC1[(begin (cp1)
    (push (x1 cp1 cp3)
    ((lambda (x2) (pop (begin (cp3) x2))
    (cp2))))])])
    (x1, x2 fresh)

```

Figure 17. Call/cc and dynamic-wind

```

(store ((ptrs svs) ...)
  (dw (dws1 ...))
  EC1[(push dws2 enext)])
→ (store ((ptrs svs) ...)
  (dw (dws2 dws1 ...))
  EC1[enext])

(store ((ptrs svs) ...)
  (dw (dws1 dws2 ...))
  EC1[(pop enext)])
→ (store ((ptrs svs) ...)
  (dw (dws2 ...))
  EC1[enext])

(store ((ptrs svs) ...)
  (dw (dws1 ...))
  EC1[(throw xk dws2 ... EC2[e2])])
→ (store ((ptrs svs) ...)
  (dw (dws2 ...))
  (begin T [(dws2 ...), (dws1 ...)]
  EC2[e2]))

```

Figure 18. Call/cc and dynamic-wind support

PC[v ₁]*	→	PC[(#%values ^o v ₁ ^o)]
PC[(#%values ^o v ₁ ^o) _o]	→	PC[v ₁]
PC[(#%values ^o v ₁ ^o ...) _o]	→	error: wrong number of values (#v ₁ ≠ 1)
(#%call-with-values v _{vals} v _{fun})	↦ ^o	(#%call-with-values ^o (cww-mark (v _{vals}) v _{fun} ^o)
PC[(#%call-with-values ^o (cww-mark (#%values ^o v _{arg} ^o ...)) v _{fun} ^o)]	→	PC[(v _{fun} ^o v _{arg} ^o ...)]
(#%call-with-values ^o v _i ^o ...)	→ ^e	error: arity mismatch (#v _i ≠ 2)

Figure 19. Multiple values and call-with-values

(#%eqv? pp _i pp _i)	↦ ^o	#t
(#%eqv? cp _i cp _i)	↦ ^o	#t
(#%eqv? number ₁ number ₁)	↦ ^o	#t
(#%eqv? v ₁ v ₁)	↦ ^o	#t
PC[(#%eqv? ^o v ₁ ^o v ₂ ^o)]	→	PC[#f] (v ₁ ≠ v ₂)
(#%eqv? ^o v ₁ ^o ...)	→ ^e	error: arity mismatch (#v ₁ ≠ 2)

Figure 20. Eqv and equivalence

6.6 Call/cc

Our technique for modeling *call/cc* and *dynamic-wind*, shown in figures 17 and 18, is essentially the technique from section 5. Apart from the change of using procedure pointers rather than the literal source text of procedures as required to model equality (see section 6.8), the only substantial change is that the continuation procedures in this model accept any number of arguments. The trimming metafunction \mathcal{T} is the same function defined in section 5.

6.7 Multiple values and call-with-values

Multiple values in the full language are nearly identical to multiple values in section 3, and in particular the context arrangement and promotion and demotion rules are the same. Furthermore, even though the present system is much larger than the system presented in section 3, the rules for multiple values are still completely orthogonal to the rules that implement the other features.

There is one twist, though, since rather than the **apply-values** primitive given in section 3, R⁵RS Scheme provides *call-with-values*, so we model it directly. To do so, we have to use the mechanisms described in section 3, along with a new context containing *cww-mark*. A term of the form $(\# \% \text{call-with-values } \text{thunk } f)$ reduces to $(\# \% \text{call-with-values } (\text{cww-mark } (\text{thunk})) f)$; that is, it places a special mark around the application of the thunk to no arguments. At that point the evaluation contexts defined in figure 8 will apply and reduce the applied thunk in a multi-value context. When that reduction sequence yields a result (which will be a multiple-values expression), the entire *call-with-values* expression reduces to the application of the second procedure to those produced values.

6.8 Eqv? and equivalence

Figure 20 shows the rules for *eqv?*. Since all mutable values (and procedures) are allocated in the store, *eqv?* is a simple matter

$(\text{store } ((ptr_1 \ sv_1) \ \dots) \ DC[\#eval^{\circ} \ v_1^{\circ}]))$	→	$(\text{store } ((ptr_1 \ sv_1) \ \dots) \ DC[\mathcal{R} \ [((ptr_1 \ sv_1) \ \dots), \ v_1]]])$
$(\text{store } ((ptr_1 \ sv_1) \ \dots) \ (\text{dw } (dws_1 \ \dots) \ EC_1[SC_1[(\text{quote } (s_1 \ s_2 \ \dots))]]]))$	→	$(\text{store } ((ptr_1 \ sv_1) \ \dots) \ (\text{dw } (dws_1 \ \dots) \ EC_1[SC_1[(\text{ccons } (\text{quote } s_1) \ (\text{quote } (s_2 \ \dots))]]]))$
$(\text{store } ((ptr_1 \ sv_1) \ \dots) \ (\text{dw } (dws_1 \ \dots) \ EC_1[SC_1[(\text{quote } ())]]]))$	→	$(\text{store } ((ptr_1 \ sv_1) \ \dots) \ (\text{dw } (dws_1 \ \dots) \ EC_1[SC_1[\#null]]]))$
$(\text{store } ((ptr_1 \ sv_1) \ \dots) \ (\text{dw } (dws_1 \ \dots) \ EC_1[SC_1[(\text{quote } number_1)]]]))$	→	$(\text{store } ((ptr_1 \ sv_1) \ \dots) \ (\text{dw } (dws_1 \ \dots) \ EC_1[SC_1[number_1]]]))$
$(\text{store } ((ptr_1 \ sv_1) \ \dots) \ (\text{dw } (dws_1 \ \dots) \ EC_1[SC_1[(\text{ccons } v_1 \ v_2)]]]))$	→	$(\text{store } ((ptr_1 \ sv_1) \ \dots \ (pp_1 \ (\#cons \ v_1 \ v_2))) \ (\text{dw } (dws_1 \ \dots) \ EC_1[SC_1[pp_1]]))$ <p style="margin-left: 20px;"><i>(pp_1 fresh)</i></p>

Figure 21. Quote and eval

of checking that the two values supplied have identical syntactic structure (which we indicate here, as PLT Redex does, by repeating the same subscript for both arguments to the *eqv?* procedure to indicate that the two subterms must be identical).

6.9 Quote and eval

The rules for *#eval* and **quote** in figure 21 are essentially the same as the rules for *eval* and **quote** in section 4. The main difference is that the rewriting rules for replacing quote are nested an SC context inside an EC context. This only matters when using *#eval*. In particular, if the call to *#eval* is in some marked context, SC will not match properly, due to the marks. In the smaller calculus, we could get away with just using SC, since it also encompassed evaluation contexts, but here we must be explicit. The reify function (\mathcal{R}) used here is as defined in section 4.

7. Related Work

Reduction semantics has been used to model large programming languages many times and in many different ways. Felleisen’s dissertation [3], which introduced context-sensitive reduction semantics, gives a formulation of a substantially smaller language than the one we present here that he calls “idealized Scheme,” and Felleisen extends that model into the λ -*v-CS* calculus in later work [4]. Since then, reduction semantics have been used to model the cores of many languages including Emacs Lisp [19], MultiLisp [7], Java [9], ML [13, 24] and Concurrent ML [21] among many others. Harper and Stone present a formal semantics for Standard ML that includes a dynamic semantics encoded using a variation on Wright and Felleisen’s notation; it is the largest example of a programming language semantics given in a variant of reduction semantics we have found in the literature (with the possible exception of our own semantics for R⁵RS Scheme).

There has also been extensive work on the semantics of Scheme. Clinger presented an operational semantics for a core Scheme in the development of the notion of space efficiency [2]. Gasbichler, Knauel, Sperber, and Kelsey have presented operational and denotational semantics for *dynamic-wind* [12]. Ramsdell presented a structural operational semantics for Scheme aimed at fixing the unspecified order of argument evaluation problem we discuss in subsection 2 [20]. His model is less complete than ours (for instance, it does not include multiple return values) and is tied much more closely to the R⁵RS Scheme formal semantics. Van Straaten has written an interpreter based on the R⁵RS Scheme denotational se-

mantics [23], but we know of no formal correspondence between his program and the denotational semantics itself.

8. Conclusion

We have presented a semantics for R⁵RS Scheme using context-sensitive reduction semantics developed using PLT Redex. To the best of our knowledge, it formalizes more of the language than any other semantics for the language. In addition it shows how to model R⁵RS Scheme-style multiple return values in an small-step operational semantics setting for the first time, and gives a new model for unspecified sequential evaluation orders that uses nondeterministic choice. In the process, we have introduced several new techniques for modeling programming language features with term rewriting.

PLT Redex and the source code for all the models presented in this paper, including our executable model of R⁵RS Scheme, are available for download at

<http://www.cs.uchicago.edu/~jacobm/r5rs/>

Acknowledgments

Thanks to Kent Dybvig and Matthew Flatt for helpful discussions of the technical details presented here and the inner workings of Chez Scheme [1] and MzScheme [8]. Thanks also to John Reppy and Dave MacQueen and the anonymous reviewers for their helpful suggestions.

References

- [1] Cadence Research Systems. *Chez Scheme Reference Manual*, 1994.
- [2] William D Clinger. Proper tail recursion and space efficiency. In *Proceedings of ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 174–185, June 1998.
- [3] Matthias Felleisen. *The Calculi of Lambda-v-CS Conversion: A Syntactic Theory of Control and State In Imperative Higher-Order Programming Languages*. PhD thesis, Indiana University, 1987.
- [4] Matthias Felleisen. Lambda-v-CS: and extended lambda-calculus for Scheme. In *Proceedings of the Conference on LISP and Functional Programming*, 1988.
- [5] Matthias Felleisen and Matthew Flatt. Programming languages and lambda calculi. Available online: <http://www.cs.utah.edu/plt/publications/plc.pdf>, 2003.
- [6] Matthias Felleisen and Robert Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical*

- Computer Science*, 102:235–271, 1992. Original version in: Technical Report 89-100, Rice University, June 1989.
- [7] Cormac Flanagan and Matthias Felleisen. The semantics of future and an application. *Journal of Functional Programming*, 9:1–31, 1999.
- [8] Matthew Flatt. PLT MzScheme: Language manual. Technical Report TR97-280, Rice University, 1997. <http://www.plt-scheme.org/software/mzscheme/>.
- [9] Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. A programmer’s reduction semantics for classes and mixins. *Formal Syntax and Semantics of Java*, 1523:241–269, 1999. Preliminary version appeared in proceedings of *Principles of Programming Languages*, 1998. Revised version is Rice University technical report TR 97-293, June 1999.
- [10] Daniel P. Friedman and Christopher T. Haynes. Constraining control. In *Proceedings of the ACM Conference Principles of Programming Languages*, 1985.
- [11] Daniel P. Friedman, Christopher T. Haynes, Eugene Kohlbecker, and Mitchell Wand. Scheme 84 interim reference manual. Technical Report 153, Indiana University Computer Science, 1985.
- [12] Martin Gasbichler, Eric Knauer, Michael Sperber, and Richard A. Kelsey. How to add threads to a sequential language without getting tangled up. In *Proceedings of the 2003 Scheme Workshop*, 2003.
- [13] Robert Harper and Mark Lillibridge. Explicit polymorphism and CPS conversion. In *Proceedings of the ACM Conference Principles of Programming Languages*, 1993.
- [14] David Herman and Philippe Meunier. Improving the static analysis of embedded languages via partial evaluation. In *Proceedings of ACM SIGPLAN International Conference on Functional Programming*, pages 16–27, New York, NY, USA, 2004. ACM Press.
- [15] Rickard Kelsey, William Clinger, and Jonathan Rees (Editors). Revised⁵ report of the algorithmic language Scheme. *ACM SIGPLAN Notices*, 33(9):26–76, 1998.
- [16] Jacob Matthews. Operational semantics for Scheme via term rewriting. Technical Report TR-2005-02, University of Chicago, 2005.
- [17] Jacob Matthews, Robert Bruce Finder, Matthew Flatt, and Matthias Felleisen. A visual environment for developing context-sensitive term rewriting systems. In *Proceedings of the International Conference on Rewriting Techniques and Applications (RTA)*, 2004.
- [18] Robert Muller. M-LISP: A representation-independent dialect of LISP with reduction semantics. *ACM Transactions on Programming Languages and Systems*, 14(4), 1992.
- [19] Matthias Neubauer and Michael Sperber. Down with Emacs Lisp: Dynamic scope analysis. In *Proceedings of ACM SIGPLAN International Conference on Functional Programming*, 2001.
- [20] John D. Ramsdell. An operational semantics for Scheme. *Lisp Pointers*, volume 2, April–June 1992.
- [21] John Reppy. *Concurrent Programming in ML*. Cambridge University Press, 1999.
- [22] Gerald Jay Sussman and Jr Guy Lewis Steele. Scheme: An interpreter for extended lambda calculus. Technical Report AI Lab Memo AIM-349, MIT AI Lab, 1975.
- [23] Anton van Straaten. An executable denotational semantics for Scheme. <http://www.appsoptions.com/SchemeDS/>.
- [24] Andrew Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computation*, pages 38–94, 1994. First appeared as Technical Report TR160, Rice University, 1991.