

Denotational Semantics Using an Operationally-Based Term Model

Mitchell Wand* Gregory T. Sullivan*

College of Computer Science

Northeastern University

360 Huntington Avenue, 161CN

Boston, MA 02115, USA

voice: (617) 373-2462, fax: (617) 373-5121

{wand,gregs}@ccs.neu.edu

<http://www.ccs.neu.edu/home/{wand,gregs}>

Abstract

We introduce a method for proving the correctness of transformations of programs in languages like Scheme and ML. The method consists of giving the programs a denotational semantics in an operationally-based term model in which interaction is the basic observable, and showing that the transformation is meaning-preserving. This allows us to consider correctness for programs that interact with their environment without terminating, and also for transformations that change the internal store behavior of the program. We illustrate the technique on one of the Meyer-Sieber examples, and we use it to prove the correctness of assignment elimination for Scheme. The latter is an important but subtle step for Scheme compilers; we believe ours is the first proof of its correctness.

1 Introduction

Compilers for higher-order languages typically perform elaborate program transformations in order to improve performance. Such transformations often change the storage behavior of the program. For concreteness, let us consider a simple language that, like Scheme, allocates a new mutable cell for each parameter on each procedure call. Then we should have an equivalence

$$((\text{lambda } (x) x) e) \equiv e \quad (1)$$

even though the left-hand program allocates one more cell than the right-hand program. We should expect a good compiler to perform this optimization, and we would expect the semantics of the language to justify it.

*Work supported by the National Science Foundation under grants number CCR-9304144 and CCR-9404646. Current address: Gregory T. Sullivan, Fulton 414, Computer Science Department, Boston College, Chestnut Hill, MA 02167. E-mail: gregs@cs.bc.edu ; WWW: <http://www.cs.bc.edu/~gregs>.

Unfortunately, in the denotational semantics of Scheme, the equivalence above is unsound, because these programs could be executed in a continuation that tests the store in ways that are not expressible in Scheme itself. It is possible to construct denotational semantics in which some equivalences like (1) are true, but they are quite complex, and none extend far enough to encompass real languages [MS88, Ole85, OT95, Ode94].

We seek to prove the correctness of such source-level transformations and translations in compilers for languages like Scheme and ML. This setting is challenging in a number of respects:

- We seek methods that are applicable to both untyped and typed languages, including Scheme, ML, and Idealized Algol, and also to translations from one language to another.
- We seek methods that will be sound even for reasoning about programs like servers that do not terminate.
- We seek methods that will be sound even for transformations that change the behavior of a heap-allocated store.

Our approach to such equivalences is to give the language a denotational semantics in an *event-based model*. An event-based model is a term model embedded in a process calculus. Observable behavior is taken from the process calculus; equality of terms in the model is taken to be contextual equivalence as measured by this observable behavior.

We illustrate the technique on one of the Meyer-Sieber examples [MS88], and we use it to prove the correctness of assignment elimination, an important optimization used in compilers for Scheme [KKsR⁺86].

We believe that this paper makes two contributions: first, it formulates a framework that handles all the requirements above. Second, it presents the first known proof of correctness of assignment elimination.

Following a review of related work in Section 2, we begin, in Section 3, by defining a metalanguage. This metalanguage is an extended PCF with recursive types, a store, and a CPS I/O system. The semantics of our

source languages will be specified by syntax-directed translations into this metalanguage. The recursive types in the metalanguage allow us to translate both typed and untyped languages easily.

In Sections 3.2 and 3.3, we define the behavior of a metalanguage term, and define two terms to be *contextually equivalent* (denoted $\cong_{\forall C}$) if they have the same I/O behavior when inserted into every possible program context. By identifying input/output as the observable actions of the system, we avoid considering the store as part of the “final answer” of a computation. Indeed, we avoid the notion of a “final answer” at all, so that we can study equivalence of terms that have infinite behavior [CG95, TW96].

Contextual equivalence is, by the form of its definition, a *congruence* – that is, it is preserved by all term constructors of the language. Furthermore, it is easy to show (using the confluence property for PCF), that if $M =_{pcf} N$, then $M \cong_{\forall C} N$. Hence the structure obtained by taking equivalence classes of open terms under $\cong_{\forall C}$ is a model of typed lambda-calculus. This is our operationally-based term model. Because this is a model, ordinary reasoning using β -conversion is valid in the model. This distinguishes it from treatments like [MT91], in which only the weaker β_v rule is valid.

Unfortunately, the quantification over all contexts makes $\cong_{\forall C}$ difficult to establish directly. We must consider not only contexts in which the terms are executed, but also contexts that pass the terms around, put many copies of them in the store, etc. Our major tool for easing this proof burden is a “context lemma” (Theorem 4.2) stating that two terms are contextually equivalent if and only if they are equivalent in a much smaller set of contexts we call *extensional contexts*; for types constructed with \rightarrow , these are the same as head contexts.

We then turn to the specification of source languages. We can specify a semantics in this model by giving a compositional (syntax-directed) translation $\llbracket - \rrbracket_{\mathcal{L}}$ from a source language L to terms of the metalanguage; the meaning of a source language expression is the $\cong_{\forall C}$ -equivalence class of the translated term. To show two source language expressions e_1 and e_2 have the same meaning, we must show that their translations are equal, i.e., that $\llbracket e_1 \rrbracket_{\mathcal{L}} \cong_{\forall C} \llbracket e_2 \rrbracket_{\mathcal{L}}$. To show a translation Φ from L to L' is correct, we equip L' with a similar translation $\llbracket - \rrbracket_{\mathcal{L}'}$ and show that $\llbracket e \rrbracket_{\mathcal{L}} \cong_{\forall C} \llbracket \Phi(e) \rrbracket_{\mathcal{L}'}$.

Our translations will produce terms in continuation-passing style. This means that the translation will not be fully abstract for stack-based languages like Idealized Algol. However, we are primarily concerned with languages like Scheme and ML that already have continuations, and CPS translation is well-suited to these languages.

Section 5 illustrates this technique using an example from [MS88].

In Section 6, we use the technique to prove the correctness of the assignment elimination transformation for Scheme [KKsR⁺86]. We conclude with consideration of open problems and future work.

2 Related Work

Operationally-based term models were introduced in [TW96]. Here we extend [TW96] by adding types and a heap, and by considering applications to program transformations. The formulation of input-output both here and in [TW96] is based on Gordon’s CPS input-output system [Gor94].

Observational or contextual equivalence for programming languages was introduced by [Plo77], adapting the work of Morris [Mor68] for the λ -calculus. Plotkin’s formulation, like almost all other work for sequential languages, dealt with conversion to a constant. Context lemmas for proving contextual equivalence were introduced by Milner [Mil77] in a typed functional language. Gordon [Gor95] extended the result to include recursive types with potentially infinite behavior.

The work most closely related to ours is that of Mason and Talcott [MT91] and Felleisen [Fel87]. They considered contextual equivalence for a number of untyped call-by-value languages using operational techniques, and proved context lemmas for these languages. Our approach differs from theirs in that we propose a single metalanguage into which a variety of source languages can be translated; neither Mason and Talcott nor Felleisen consider issues of translation. Our Context Lemma (Theorem 4.2) also extends their results by considering input-output effects rather than final answers and by validating full β -conversion rather than their β -value.

Crole and Gordon [CG95] give a denotational semantics for a typed call-by-value language with input-output by translation into a metalogic \mathcal{M} similar in purpose, though not in structure, to our metalanguage. Their translation uses a monadic style in place of our continuation-passing style. Their treatment does not include a store, and it is not clear whether their system can be extended to deal with latently recursively typed languages like Scheme.

Most work on store semantics [MS88, Ole85, OT95, Ode94] deals with stores in stack-discipline languages, and are complicated by the need to model that restriction. Most of this work depends crucially on the type structure of the language; our work is applicable to untyped or latently typed languages as well. [PS93] treats heap-allocated names without content or mutation.

Bisimulation was introduced by [Par81]. Abramsky [Abr90] used bisimulation to define a λ -model, but his observables were very different from ours. The concept of a relation preserved under computation is probably much older. It was used for program equivalence in [MT91, WO92, ORW95]; [Sta94] uses a similar notion extended to become a logical relation.

Assignment elimination was introduced in the Orbit compiler [KKsR⁺86] and used in [KH89, CH94]. It was omitted from [GSR95] because the authors were unable to prove its correctness [Ram96].

3 An Operationally-Based Term Model

3.1 Syntax

Values in the model are equivalence classes of terms in an extended λ -calculus that we call the *metalanguage*.

The types in the metalanguage, ranged over by τ , are potentially infinite trees which respect the following structure:

$$\begin{aligned} o &::= \text{unit} \mid \text{bool} \mid \text{int} \mid \text{ref}(\tau) \mid pr \\ \tau &::= o \mid \tau \rightarrow \tau \mid \tau + \tau \mid \tau \times \tau \end{aligned}$$

The base types o include *int* for integers, *bool* for booleans, *ref*(τ) for locations with contents of type τ , and *unit* for the singleton type. Type *pr* is for processes, which may do I/O.

We allow infinite types, so we will need to define operations on types coinductively instead of inductively. We write τ^* for the infinite type

$$\text{unit} + (\tau \times (\text{unit} + (\tau \times \dots)))$$

of lists of elements of type τ .

The syntax of our metalanguage is given by the following grammar:

$$M ::= x \mid n \mid c \mid l^\tau \mid \lambda x:\tau. M \mid MM$$

with x ranging over a countably infinite set of variables and l ranging over a countably infinite set of location constants. Locations l^τ are annotated with types and are assigned type *ref*(τ). The typing rules for the language are standard. The meta-operation *locs*(M) returns the set of locations occurring in M and is easily defined via structural induction.

The metalanguage comes equipped with PCF-like constants and the usual constants for sums and products. More interesting are the constants for stores and processes, whose types are given in Figure 1. Their behavior is considered in Section 3.2.

In the remainder of this section, we turn this language into a model of the lambda-calculus by defining the basic observable to be the input-output behavior of a closed term of type *pr*. We then define two terms to be equivalent (\cong_{VC}) iff they have the same observable behavior in any closing context $C[-]$ of type *pr*. It follows immediately that \cong_{VC} is a congruence that respects β -reduction, so the equivalence classes of metalanguage terms under \cong_{VC} form a model of typed lambda-calculus.

3.2 Operational Semantics and Observable Equivalence

We define the operational semantics of the metalanguage in three steps. The reduction semantics of the PCF fragment of the metalanguage defines a relation \rightarrow_{pcf} between closed well-typed terms. The reduction relation \rightarrow_{sto} is defined on states, includes PCF reduction within term components, and gives transition rules

for states with store operations in head position. Finally, we define a labeled transition relation, $\xrightarrow{a}_{\text{io}}$, between states which reduce under $\rightarrow_{\text{sto}}^*$ to states headed by I/O operations.

The definition of \rightarrow_{pcf} is the usual call-by-name reduction [Plo77] and is omitted. Figure 2 defines the transition relation \rightarrow_{sto} between states. A *state*, denoted $\langle M ; \Sigma \rangle$, is a term-store pair where M is a closed term of type *pr* and store Σ is a finite map from locations to terms. All locations l^τ in the domain of Σ must map to closed terms of type τ .

The first rule of Figure 2 incorporates PCF reduction into \rightarrow_{sto} ; if the term component of a state may take a single PCF transition, the state pair may take the corresponding transition. The new operator allocates a new location in the store, initializes the fresh location to its first argument, and invokes its second argument with the new location. The *deref* operator invokes its second argument with the value in the store at the location given by its first argument. The *update* operator destructively updates the location given by its first argument to be the term given by its second argument and then invokes its third argument. The reflexive, transitive closure of \rightarrow_{sto} is denoted by $\rightarrow_{\text{sto}}^*$.

A labeled transition system $\xrightarrow{a}_{\text{io}}$ for the I/O operators is defined in Figure 3. The *write* operator writes its integer argument and invokes its continuation. The *read* operator reads an integer and invokes its continuation with the number read. We make termination observable by having a *stop* operator that takes a \checkmark transition to a stopped state $\langle \text{stopped} ; \Sigma' \rangle$, from which there are no transitions.

Our notion of observable equivalence on states is \approx_{sim} , the standard notion of (strong) bisimulation for this labelled transition system.¹

It is easy to show that consistent renaming of locations preserves observable equivalence. Two states are equal up to a renaming of location constants, denoted α_{loc} , if there is a relation Θ between locations satisfying the following properties:

1. if $(i, j) \in \Theta$ and $(i', j') \in \Theta$, then $(i = i' \Leftrightarrow j = j')$
2. $M\Theta = N$
3. $\text{locs}(M) \subseteq \text{dom}(\Theta)$
4. if $(l, l') \in \Theta$ then either
 - (a) $l \notin \text{dom}(\Sigma_M)$ and $l \notin \text{dom}(\Sigma_N)$ or
 - (b) $\Sigma_M(l)\Theta = \Sigma_N(l')$ and $\text{locs}(\Sigma_M(l)) \subseteq \text{dom}(\Theta)$

Lemma 3.1

$$\alpha_{\text{loc}} \subseteq \lesssim_{\text{sim}}$$

PROOF SKETCH: Use coinduction on the definition of \approx_{sim} .

¹Alternatively, we could have made \rightarrow_{sto} into a silent transition, and used weak bisimulation. We believe the results would be the same; the proofs would look somewhat different.

$\text{write} : \text{int} \rightarrow \text{pr} \rightarrow \text{pr},$	write an integer to stdout
$\text{read} : (\text{int} \rightarrow \text{pr}) \rightarrow \text{pr},$	read integer from stdin
$\text{stop} : \text{pr}, \text{stopped} : \text{pr},$	halt, halted
$\text{new}^\tau : \tau \rightarrow (\text{ref}(\tau) \rightarrow \text{pr}) \rightarrow \text{pr},$	get fresh loc, initialize to 1st arg
$\text{deref}^\tau : \text{ref}(\tau) \rightarrow (\tau \rightarrow \text{pr}) \rightarrow \text{pr},$	dereference a location
$\text{update}^\tau : \text{ref}(\tau) \rightarrow \tau \rightarrow \text{pr} \rightarrow \text{pr},$	update a location
$\text{eq?}_{\text{loc}}^\tau : \text{ref}(\tau) \rightarrow \text{ref}(\tau) \rightarrow \text{bool}$	compare locations for equality

Figure 1: Metalanguage Constants for Store and I/O Operations

$$\begin{aligned}
M \rightarrow_{\text{pcf}} M' &\Rightarrow \langle\langle M ; \Sigma \rangle\rangle \rightarrow_{\text{sto}} \langle\langle M' ; \Sigma \rangle\rangle \\
\langle\langle \text{new}^\tau MN ; \Sigma \rangle\rangle &\rightarrow_{\text{sto}} \langle\langle N l^\tau ; \Sigma[l^\tau \mapsto M] \rangle\rangle && \text{if } l^\tau \notin \text{dom}(\Sigma) \cup \text{locs}(\text{rng}(\Sigma), M, N) \\
\langle\langle \text{deref}^\tau l^\tau M ; \Sigma \rangle\rangle &\rightarrow_{\text{sto}} \langle\langle MN ; \Sigma \rangle\rangle && \text{if } \Sigma(l^\tau) = N \\
\langle\langle \text{update}^\tau l^\tau MN ; \Sigma \rangle\rangle &\rightarrow_{\text{sto}} \langle\langle N ; \Sigma[l^\tau \mapsto M] \rangle\rangle && \text{if } l^\tau \in \text{dom}(\Sigma)
\end{aligned}$$

Figure 2: Operational Semantics for Store Operators

3.3 Term Equivalence

We say that two terms are *contextually equivalent* if they have the same I/O behavior when inserted into every possible program context:²

Definition 3.2 (contextual equivalence, $\cong_{\forall C}$)
If $\Gamma \vdash M, N : \tau$, then $\Gamma \vdash M \cong_{\forall C} N : \tau$ iff for all closing contexts $C[-]$ of type pr that respect Γ , and for all stores Σ ,

$$\langle\langle C[M] ; \Sigma \rangle\rangle \approx_{\text{sim}} \langle\langle C[N] ; \Sigma \rangle\rangle$$

Contextual equivalence is, by the form of its definition, a *congruence* – that is, it is preserved by all term constructors of the language. Furthermore, it is easy to show (using the confluence property for PCF), that if $M =_{\text{pcf}} N$, then $M \cong_{\forall C} N$. Hence the structure obtained by taking equivalence classes of open terms³ under $\cong_{\forall C}$ is a model of typed lambda-calculus. This is the operationally-based term model we will use. Because this is a model, ordinary reasoning using β -conversion is valid in the model.

Definition 3.2 might seem unduly restrictive, because it requires terms to behave identically not only in every program context, but in every store. For terms that contain no location constants, however, the store adds no discriminative power:

Lemma 3.3 *If $\Gamma \vdash M, N : \tau$, and M and N contain no location constants, then $\Gamma \vdash M \cong_{\forall C} N : \tau$ iff for all*

²respecting Γ means that the context binds the free variables of M with binders of the correct type.

³Actually, terms open with respect to an extended type assignment Γ^* [Gun92, p. 55].

closing contexts $C[-]$ of type pr that respect Γ ,

$$\langle\langle C[M] ; \emptyset \rangle\rangle \approx_{\text{sim}} \langle\langle C[N] ; \emptyset \rangle\rangle$$

PROOF SKETCH: The forward direction is immediate. For the reverse direction, given a context C and store Σ , it is easy to construct a context C' that builds Σ and then executes M in context C .

Our plan is to specify the semantics of our source language in this model by giving a compositional translation $\llbracket - \rrbracket$ from the source language to terms of the metalanguage; the meaning of a source language expression is the $\cong_{\forall C}$ -equivalence class of the translated term. To show two source language expressions e_1 and e_2 have the same meaning, we must show that their translations lie in the same equivalence class, i.e., that $\llbracket e_1 \rrbracket \cong_{\forall C} \llbracket e_2 \rrbracket$.

4 Extensional Equivalence

The quantification over all contexts makes $\cong_{\forall C}$ difficult to establish directly. We must consider not only contexts in which the terms are executed, but also contexts that pass the terms around, put many copies of them in the store, etc.

Our major tool for easing this proof burden is a “context lemma” which states that two terms are contextually equivalent if and only if they are equivalent in a smaller set of contexts we call *extensional contexts*.

Extensional contexts are motivated by the observation that for each type constructor $(\rightarrow, +, \times)$, there is a natural way of determining whether two terms of such a type “behave the same way.” At a function type $\alpha \rightarrow \beta$, two terms are extensionally equal if they take any term

$$\begin{array}{c}
\frac{\langle\langle M ; \Sigma \rangle\rangle \rightarrow_{sto}^* \langle\langle \text{write } n \ M' ; \Sigma' \rangle\rangle}{\langle\langle M ; \Sigma \rangle\rangle \xrightarrow{!n}_{io} \langle\langle M' ; \Sigma' \rangle\rangle} \quad \frac{\langle\langle M ; \Sigma \rangle\rangle \rightarrow_{sto}^* \langle\langle \text{read } M' ; \Sigma' \rangle\rangle}{\langle\langle M ; \Sigma \rangle\rangle \xrightarrow{?n}_{io} \langle\langle M' \ n ; \Sigma' \rangle\rangle} \\
\\
\frac{\langle\langle M ; \Sigma \rangle\rangle \rightarrow_{sto}^* \langle\langle \text{stop} ; \Sigma' \rangle\rangle}{\langle\langle M ; \Sigma \rangle\rangle \xrightarrow{\vee}_{io} \langle\langle \text{stopped} ; \Sigma' \rangle\rangle}
\end{array}$$

Figure 3: Semantics for I/O Operators

of type α to equivalent terms of type β . For product types, both projections should be equivalent, and for sum types, the case operator should produce equivalent terms. We can view these type-based contexts (application, projection, and case) as *type deconstructors*, as they decompose complex types into terms of their component types. If our only complex types were function types, then the extensional contexts would be exactly the head contexts of Milner’s context lemma.

Figure 4 gives the grammar for extensional contexts. Each extensional context $T[-\tau] : o$ is a context with a single hole of type τ that produces base type o .

We will say that two closed terms are extensionally equivalent iff the terms have equivalent base-type behavior in every extensional context. Open terms will be extensionally equivalent if all their substitution instances are. Equivalence for closed terms at base types, denoted \cong_o , is defined simply: Equivalent terms of type pr must be in \approx_{sim} for all stores Σ . For base types other than pr , terms must PCF-reduce to the same constant. For base type $ref(\tau)$ this means the same location constant. We omit the formal definition in this abstract. In the general case, the formal definition becomes:

Definition 4.1 (extensional equivalence, \cong_{ext})
If $\Gamma \vdash M, N : \tau$, then $\Gamma \vdash M \cong_{ext} N : \tau$ iff for all closing substitutions σ which respect Γ , and for all extensional contexts $T[-\tau] : o$,

$$T[M\sigma] \cong_o T[N\sigma]$$

This definition is similar to Gordon’s notion of equivalence [Gor95]. His bisimilarity condition requires terms to be equivalent in exactly our extensional contexts; our base types correspond to his active types.

The main result of this section is that extensional and contextual equivalence coincide:

Theorem 4.2 (context lemma)

$$\Gamma \vdash M \cong_{ext} N : \tau \Leftrightarrow \Gamma \vdash M \cong_{VC} N : \tau$$

PROOF SKETCH: First, we show that when extensionally equivalent closed terms are substituted into a term with a single free variable, the resulting terms are extensionally equivalent. This is the most complex portion of the proof, as it involves the inductive definition of substitution, the coinductive definition of \approx_{sim} , and the inductive definition of \rightarrow_{sto}^* . It is then easy to show

that \cong_{ext} is a congruence, and hence that $\cong_{ext} \subseteq \cong_{VC}$. The reverse inclusion follows by a standard argument [TW96]. \square

Extensional equivalence is not the same as using logical relations; it imposes a lesser proof burden than methods based on logical relations, and also extends to untyped or recursively typed languages. Extensional equivalence requires only that terms behave equivalently in the *same* context. This contrasts with the logical-relations approach, in which terms are required to behave equivalently in all *equivalent* contexts. This in turn leads to non-monotonicity which is remedied only by restriction to finite types. That extensionally-equivalent terms do behave equivalently in equivalent contexts follows from Theorem 4.2.

5 Meyer-Sieber Examples

We now leave the metalanguage and turn to applications. We begin by considering the examples from [MS88]. Our techniques suffice to do all these examples; we give an illustrative one here.

Example 4 from [MS88] says:

The block below always diverges.

```

begin
  new x; new y;
  procedure Twice; begin y := 2 * contents(y) end;
  x := 0; y := 0;
  Q(Twice);           % Q is declared elsewhere
  if contents(x) = 0 then diverge fi
end

```

In a language with I/O, the above block is *not* necessarily equivalent to *diverge*, as the procedure Q might do I/O. So we will undertake to prove the equivalence of the following two blocks:

```

begin
  new x; new y;
  procedure Twice; begin y := 2 * contents(y) end;
  x := 0; y := 0;
  Q(Twice);           % Q is declared elsewhere
  if contents(x) = 0 then skip else diverge fi
end

```

and

$T[-_o]:o \quad ::= [-]$	base types
$T[-_{\tau_1 \rightarrow \tau_2}]:o \quad ::= T[[-] M]:o$	for all closed $M:\tau_1$
$T[-_{\tau_1 \times \tau_2}]:o \quad ::= T[\pi_1[-]]:o \mid T[\pi_2[-]]:o$	
$T[-_{\tau_1 + \tau_2}]:o \quad ::= T[\text{case}_{\tau_1, \tau_2}^\gamma [-] M_1 M_2]:o$	for all $\gamma, M_1:\tau_1 \rightarrow \gamma, M_2:\tau_2 \rightarrow \gamma$

Figure 4: Grammar for Extensional Contexts

```

begin
  new  $x$ ; new  $y$ ;
  procedure  $\text{Twice}$ ; begin  $y := 2 * \text{contents}(y)$  end;
   $x := 0$ ;  $y := 0$ ;
   $Q(\text{Twice})$ ;           %  $Q$  is declared elsewhere
end

```

where *skip* does nothing.

Our general plan for proving the equivalence of two source-program expressions, e_1 and e_2 , is to specify a compositional translation $\llbracket - \rrbracket$ from the source language into our metalanguage, and then to show that $\llbracket e_1 \rrbracket \cong_{\text{ext}} \llbracket e_2 \rrbracket$.

The semantics of this Algol fragment is given by a straightforward CPS translation [App92, Plo75] into our metalanguage, where all blocks get translated to terms of type $pr \rightarrow pr$. After some simplification, the translation yields the following metalanguage term for the left-hand side:

$$\llbracket lhs \rrbracket = (\lambda \kappa:pr. \text{new } 0 (\lambda x:ref(int). \text{new } 0 (\lambda y:ref(int). Q \text{ Twice} (\text{deref } x (\lambda v_x. \text{zero? } v_x \kappa \Omega_{pr}))))))$$

where *Twice* is defined as:

$$\text{Twice} \stackrel{\text{def}}{=} (\lambda \kappa':pr. \text{deref } y (\lambda v_y. \text{update } y (* 2 v_y) \kappa'))$$

and where Ω_τ is the canonical divergent term $Y(\lambda x:\tau.x)$.

The $\llbracket lhs \rrbracket$ term takes a run-time continuation κ , allocates fresh locations for x and y initialized to 0, and then continues with x and y bound to the fresh locations. The body of the block invokes the unknown function Q with the closure *Twice* as its argument and with a continuation that checks if x is zero and continues with κ if so or diverges if not. The entire block has type $pr \rightarrow pr$.

The translation of the right-hand side, $\llbracket rhs \rrbracket$, is exactly the same as the left-hand side, except that the $(\text{deref } x \dots)$ continuation is replaced by simply κ .

To prove that the two translations are operationally equivalent, we appeal to the context lemma and set out to prove that for all terms K of type $pr \rightarrow pr$ and all stores Σ ,

$$\langle\langle \llbracket lhs \rrbracket K ; \Sigma \rangle\rangle \approx_{\text{sim}} \langle\langle \llbracket rhs \rrbracket K ; \Sigma \rangle\rangle$$

The left-hand side will take a number of steps to the following state:

$$\langle\langle Q \text{ Twice} (\text{deref } x (\lambda v_x. \text{zero? } v_x \kappa \Omega_{pr})) \rrbracket [l_x/x, l_y/y] ; \Sigma[l_x \mapsto 0, l_y \mapsto 0] \rangle\rangle \quad (l_x, l_y \text{ fresh})$$

Similarly, the right-hand side goes to:

$$\langle\langle (Q \text{ Twice } K) \rrbracket [l_x/x, l_y/y] ; \Sigma[l_x \mapsto 0, l_y \mapsto 0] \rangle\rangle$$

up to α_{loc} .

We now construct a relation S , which includes the two states above, and we prove that S is a bisimulation.

$$S \stackrel{\text{def}}{=} \{ (\langle\langle M ; \Sigma \rangle\rangle \sigma_1, \langle\langle M ; \Sigma \rangle\rangle \sigma_2) \mid \begin{array}{l} fv(M, rng(\Sigma)) \subseteq \{z\}, \\ \sigma_1 = [(\text{deref } l_x (\lambda v_x. \text{zero? } v_x \kappa \Omega_{pr})) / z], \\ \sigma_2 = [K / z], \\ \Sigma(l_x) = 0, \\ l_x \notin locs(M, K, rng(\Sigma)) \end{array} \}$$

This simulation S characterizes the relevant aspects of the example – namely that x 's location l_x is never updated (or, more exactly, that the only references to l_x are within a term of the form $(\text{deref } l_x \dots)$).

The pair of states in which we are interested is clearly in S , with $M = (Q \text{ Twice } z)$. Because l_x is allocated fresh, we know that Q and K contain no references to l_x . The proof that S is a bisimulation is straightforward.

This example illustrates both the power of this technique to prove equivalences but also its limitations: our translation does not validate the original equivalence from [MS88]; hence the translation is not fully abstract. Since, however, the languages that we are primarily interested in, like Scheme and ML, do have input-output behavior and do not obey stack discipline, the failure of full abstraction for Idealized Algol is not so important for our intended applications.

6 Assignment Elimination

For a more substantial example, we consider the problem of assignment elimination in Scheme. While Scheme has a rich set of expressed values, its variables denote only locations, because a Scheme program can always mutate a variable by performing a *set!* on it. The semantics of Scheme accomplishes this by automatically allocating storage for each parameter at every procedure call, and by automatically dereferencing each variable reference [RC⁺86].

Many Scheme compilers perform a preliminary pass called assignment elimination, in which Scheme is translated into an intermediate language in which the denoted values include expressed values as well, and in

$$\begin{array}{ll}
e ::= x \mid n \mid (\text{lambda } (x \dots) e) \mid (e \ e \dots) & \text{expressions} \\
\mid (\text{set! } x \ e) \mid (\text{letrec } ((x \ e) \dots) e) & \\
p ::= \text{run } e & \text{programs}
\end{array}$$

where x ranges over identifiers and n ranges over integers.

Figure 5: A subset of Scheme

$$\begin{aligned}
\llbracket n \rrbracket_{scm} &= \lambda \kappa: Val \rightarrow pr. in_N \ n \ \kappa \\
\llbracket x \rrbracket_{scm} &= \lambda \kappa: Val \rightarrow pr. \text{deref } x \ \kappa \\
\llbracket (\text{lambda } (x_1 \dots x_n) e) \rrbracket_{scm} &= \lambda \kappa: Val \rightarrow pr. \\
&\quad in_F (\lambda v^*: Val^*. \lambda \kappa': Val \rightarrow pr. \\
&\quad \quad \text{new}^{Val}(\text{listref}_1 \ v^*) (\lambda x_1: ref(Val). \dots \\
&\quad \quad \text{new}^{Val}(\text{listref}_n \ v^*) (\lambda x_n: ref(Val). \llbracket e \rrbracket_{scm} \ \kappa')) \ \kappa \\
\llbracket (e_0 \ e_1 \dots e_n) \rrbracket_{scm} &= \lambda \kappa: Val \rightarrow pr. \llbracket e_0 \rrbracket_{scm} (\lambda f: Val. \\
&\quad \llbracket e_1 \rrbracket_{scm} (\lambda v_1: Val. \dots \\
&\quad \llbracket e_n \rrbracket_{scm} (\lambda v_n: Val. \\
&\quad \quad out_F \ f \ (\lambda f': Val^* \rightarrow (Val \rightarrow pr) \rightarrow pr. \\
&\quad \quad \quad f' \langle v_1, \dots, v_n \rangle \kappa)) \dots) \\
\llbracket (\text{set! } x \ e) \rrbracket_{scm} &= \lambda \kappa: Val \rightarrow pr. \llbracket e \rrbracket_{scm} (\lambda v: Val. \text{update } x \ v \ (\kappa \ \Omega_{Val})) \\
\llbracket (\text{letrec } ((x_1 \ e_1) \dots (x_n \ e_n)) e) \rrbracket_{scm} &= \lambda \kappa: Val \rightarrow pr. \\
&\quad \text{new } \Omega_{Val} (\lambda x_1: ref(Val). \dots \text{new } \Omega_{Val} (\lambda x_n: ref(Val). \\
&\quad \quad \llbracket e_1 \rrbracket_{scm} (\lambda v_1: Val. \text{update } x_1 \ v_1 \dots \\
&\quad \quad \llbracket e_n \rrbracket_{scm} (\lambda v_n: Val. \text{update } x_n \ v_n \ (\llbracket e \rrbracket_{scm} \ \kappa)) \dots) \\
\llbracket \text{run } e \rrbracket_{scm} &= \langle\langle C_0 \llbracket e \rrbracket_{scm} \rrbracket; \emptyset \rangle\rangle
\end{aligned}$$

Figure 6: CPS translation from Scheme subset to metalanguage

which allocation and dereferencing are explicit. Variables that are demonstrably never mutated are bound directly to expressed values, and those that may be mutated are bound to explicitly-allocated cells [KKsR⁺86].

Although this translation is generally considered unproblematic, its correctness is difficult to prove, because it radically changes the store behavior of the program. Here we sketch the first known proof of correctness for this transformation. We first describe the source and target language of the transformation. We formulate the translation, and then we sketch the proof itself.

6.1 Source Language Semantics

We will consider the subset of Scheme given by the grammar in Figure 5. The language consists of expressions e and programs p .

We specify the semantics of this Scheme subset by a straightforward CPS translation (Figure 6). The main difference from an ordinary denotational semantics is that it does not introduce an environment; instead, it leaves free variables in the source term as free variables in the translation.

The type Val of expressed values is given as a recursive sum:

$Val = N + F$	expressed values
$N = int$	integers
$F = Val^* \rightarrow (Val \rightarrow pr) \rightarrow pr$	procedures

The denoted values are of type $ref(Val)$. The translation of any expression is a term of type $(Val \rightarrow pr) \rightarrow pr$, given that all the free variables have type $ref(Val)$.

For each summand S of the Val type, we use CPS injection and extraction combinators, $in_S: S \rightarrow (Val \rightarrow pr) \rightarrow pr$ and $out_S: Val \rightarrow (S \rightarrow pr) \rightarrow pr$. We also use $\langle -, \dots, - \rangle$ and $listref_i$ for the evident list operations. All these can be easily defined in the metalanguage.

The rule for an abstraction $(\text{lambda } (x_1 \dots x_n) e)$ injects into the Val type a closure. When the closure is invoked, storage is allocated for each formal, the cells are initialized, and the body is executed in an environment in which the formals are bound to the new locations.

The rule for assignment, $(\text{set! } x \ e)$, updates the location bound to x with the value of $\llbracket e \rrbracket_{scm}$. The Scheme semantics calls for the continuation to be invoked with an undefined value, for which we use Ω_{Val} , the canonical nonterminating term of type Val .

$$\begin{aligned}
e &::= x \mid n \mid (\text{lambda } (x \dots) e) \mid (e e \dots) \\
&\quad \mid (\text{letrec } ((x e) \dots) e) \\
&\quad \mid (\text{let } ((x rhs) \dots) e) \\
&\quad \mid (\text{cell-ref } x) \mid (\text{cell-set! } x e) \\
rhs &::= x \mid (\text{make-cell } e)
\end{aligned}$$

Figure 7: Syntax of Scheme_{cell}

$$\begin{aligned}
\llbracket x \rrbracket_{cell} &= \lambda \kappa: Val \rightarrow pr. \kappa x \\
\llbracket (\text{lambda } (x_1 \dots x_n) e) \rrbracket_{cell} &= \lambda \kappa: Val \rightarrow pr. \\
&\quad in_F(\lambda v^*: Val^*. \lambda \kappa': Val \rightarrow pr. \\
&\quad \quad \llbracket e \rrbracket_{cell}[(listref_1 v^*)/x_1, \dots, (listref_n v^*)/x_n] \kappa') \kappa \\
\llbracket (\text{make-cell } e) \rrbracket_{cell} &= \lambda \kappa: ref(Val) \rightarrow pr. \llbracket e \rrbracket_{cell}(\lambda v: Val. \text{new } v \kappa) \\
\llbracket (\text{cell-ref } x) \rrbracket_{cell} &= \lambda \kappa: Val \rightarrow pr. \text{deref } x \kappa \\
\llbracket (\text{cell-set! } x e) \rrbracket_{cell} &= \lambda \kappa: Val \rightarrow pr. \llbracket e \rrbracket_{cell}(\lambda v: Val. \text{update } x v (\kappa \Omega))
\end{aligned}$$

Figure 8: CPS translation from Scheme_{cell} to metalanguage

The rule for recursive definitions follows the Scheme semantics [RC⁺86]. It first allocates all required locations, initializing them to a dummy value, and then updates all of the locations with their contents within a scope that includes bindings from all mutually recursive identifiers x_1, \dots, x_n to their newly allocated locations. Because Scheme allows arbitrary side-effects in the e_i , the usual definition with the Y operator is impossible.

Finally, evaluation of a program runs the expression in a suitable initial context and an empty store. The context $C_0[-]$ contains a preamble that loads the values of the primitives of Scheme into locations in the store, binds identifiers to those locations, and then runs the program in an initial continuation (probably something like $(\lambda v: Val. \text{print}_{Val} v \text{ stop})$ for some suitable metalanguage term print_{Val}). Our definition is flexible enough to accommodate the possibility of primitives that perform input-output or do non-local jumps, like `call/cc`.

This semantics easily verifies the equivalence (1) with which we began this paper.

6.2 Scheme_{cell}

Figure 7 gives the syntax of the intermediate language Scheme_{cell}, which includes new cell operations and no longer contains `set!`. Figure 8 gives the translation from Scheme_{cell} into our metalanguage. Expressed values are still of type *Val*; identifiers denote either values of type *Val* or cells of type *ref(Val)*. It is easy to deduce statically which variables are which; we omit the easy type-checking system.

The rule for variables no longer automatically does a dereference, and the rule for abstraction no longer allocates storage automatically. Instead, cell allocation and dereferencing are moved into the cell operations. The rule for `letrec` is identical to that for $\llbracket - \rrbracket_{scm}$. The rule for `let` does the evident thing and is omitted.

6.3 The Translation

Our goal is to translate our Scheme subset to Scheme_{cell}, using cells only for variables that are potentially mutated. Figure 9 gives the translation. Bound variables that do not appear on the left-hand side of a `set!` expression are translated as is, so no storage is allocated for them; all other variables are translated by explicitly allocating a location at binding time, and by explicitly dereferencing when the variable is used. Variables that are bound by a `letrec` are initialized by mutation, and are therefore treated as cells.

6.4 Correctness of Assignment Elimination

Our goal is to show that assignment elimination is meaning-preserving, that is:

Theorem 6.2 (Correctness of Assignment Elimination) *For any Scheme expression e , and any Γ mapping the free variables of $\llbracket e \rrbracket_{scm}$ and $\llbracket AE(e) \rrbracket_{cell}$ to *ref(Val)*,*

$$\Gamma \vdash \llbracket e \rrbracket_{scm} \cong_{\forall C} \llbracket AE(e) \rrbracket_{cell}: (Val \rightarrow pr) \rightarrow pr$$

PROOF:

Definition 6.1 (Assignment Elimination (AE)) Let $sv(e)$ be the set of free variables of e that appear in e as the left-hand side of a **set!** expression. Define $AE(e)$ to be $AE(\emptyset, e)$, where $AE(V, e)$ is defined as follows. Here V is the set of free variables that are to be bound to values rather than locations.

$$\begin{aligned}
AE(V, x) &= \begin{cases} x & \text{if } x \in V \\ (\text{cell-ref } x) & \text{if } x \notin V \end{cases} \\
AE(V, (\text{set! } x e)) &= (\text{cell-set! } x AE(V, e)) \\
AE(V, (e_0 e_1 \dots e_n)) &= (AE(V, e_0) AE(V, e_1) \dots AE(V, e_n)) \\
AE(V, (\text{lambda } (x_1 \dots x_n) e)) &= (\text{lambda } (x_1 \dots x_n) (\text{let } ((x_1 r_1) \dots (x_n r_n)) AE(V', e))) \\
&\text{where } \begin{cases} r_i = \begin{cases} (\text{make-cell } x_i) & \text{if } x_i \in sv(e) \\ x_i & \text{if } x_i \notin sv(e) \end{cases} \\ V' = V \cup \{x_i \mid x_i \notin sv(e)\} - \{x_i \mid x_i \in sv(e)\} \end{cases} \\
AE(V, (\text{letrec } ((x_1 e_1) \dots (x_n e_n)) e)) &= (\text{letrec } ((x_1 AE(V', e_1)) \dots (x_n AE(V', e_n))) \\
&\quad AE(V', e)) \quad \text{where } V' = V - \{x_1, \dots, x_n\}
\end{aligned}$$

Figure 9: Assignment Elimination

1. We first need to check that

$$\Gamma \vdash \llbracket AE(e) \rrbracket_{cell} : (Val \rightarrow pr) \rightarrow pr$$

This is straightforward.

2. Hence, by the Context Lemma (Theorem 4.2) and the definition of \cong_{ext} , it suffices to show that for all closing substitutions σ mapping free variables to closed terms of type $ref(Val)$, all closed terms K of type $Val \rightarrow pr$, and all stores Σ ,

$$\langle\langle \llbracket e \rrbracket_{scm} \sigma \rangle K ; \Sigma \rangle \approx_{sim} \langle\langle \llbracket AE(e) \rrbracket_{cell} \sigma \rangle K ; \Sigma \rangle$$

3. To show these states are bisimilar, we construct a relation S_{AE} between metalanguage states that characterizes the syntactic differences between (non-transformed) Scheme programs and those programs after assignment elimination, as they reduce under $\xrightarrow{\alpha}_{io}$. This construction involves several stages.
4. As the original program and its AE-translated version run, there are several different ways in which corresponding subterms can differ:

- (a) The translations of Scheme lambda expressions will differ in the bodies of their closures. The left-hand side closure body will contain new operators for all local parameters, whereas the right-hand side will only allocate storage for mutable local parameters. On the left-hand side, references to local variables will always involve store lookup with the **deref** operator, whereas on the right-hand side, references to immutable variables will not do store lookup, as the variable will get its value via substitution. The macros \mathcal{L}_1 and \mathcal{R}_1 in Figure 10

capture the different translations of a Scheme lambda expression. Note that the two equations at the top of the Figure 10 are not definitions; they are instead specifications for the term macros \mathcal{L}_1 and \mathcal{R}_1 . The figure then gives definitions for \mathcal{L}_1 and \mathcal{R}_1 that satisfy these specifications. The same pattern is followed in the remaining four blocks of Figure 10.

- (b) When subterms corresponding to lambda expressions are applied to a run-time continuation, they are injected into the F summand of the Val type and sent to the continuation. The macros \mathcal{L}_2 and \mathcal{R}_2 in Figure 10 express the form of the two sides after they are applied to run-time continuation terms.
- (c) When the closures are used at run-time (i.e., applied to arguments), they are extracted from the Val type and applied to their application-time argument list and continuation. At this point, the differences between the closure bodies are “exposed”, and we see different storage behavior. This stage of execution is expressed by the macros \mathcal{L}_3 and \mathcal{R}_3 in Figure 10.
- (d) As the exposed bodies of corresponding closures reduce, the left-hand side allocates storage for all local variables, whereas the right-hand side only allocates storage for mutable local variables. For immutable local variables, on the left-hand side we have a substitution $[l_i/x_i]$ for each x_i immutable variable (and mutable too, for that matter). On the right-hand side, for the immutable variables x_i , we see value substitutions, $[(listref_i v^*)/x_i]$, where the variable v^* contains the list of incoming arguments. We express this relation by first

Initial Translation:

$$\begin{array}{l|l}
\llbracket (\text{lambda } (x_1 \dots x_m) e) \rrbracket_{scm} = & \llbracket AE(V, (\text{lambda } (x_1 \dots x_m) e)) \rrbracket_{cell} = \\
\mathcal{L}_1 \langle x_1 \dots x_m \rangle \llbracket e \rrbracket_{scm} & \mathcal{R}_1(V', C) \llbracket \llbracket AE(V', e) \rrbracket_{cell} \rrbracket, \text{ with} \\
& V' = V \cup \{x_i \mid x_i \notin sv(e)\} - \{x_i \mid x_i \in sv(e)\} \\
& C = \{x_i \mid x_i \in sv(e)\} \\
\mathcal{L}_1 \langle x_1 \dots x_m \rangle \llbracket P \rrbracket \stackrel{\text{def}}{=} & \mathcal{R}_1(V, C) \llbracket P \rrbracket \stackrel{\text{def}}{=} \\
(\lambda \kappa: Val \rightarrow pr. \text{in}_F(\lambda v^*: Val^*. \lambda \kappa': Val \rightarrow pr. & (\lambda \kappa: Val \rightarrow pr. \text{in}_F(\lambda v^*: Val^*. \lambda \kappa': Val \rightarrow pr. \\
\text{new}^{Val}(\text{listref}_1 v^*) (\lambda x_1: ref(Val). \dots & \text{new}^{Val}(\text{listref}_i v^*) \left. \vphantom{\text{new}^{Val}(\text{listref}_1 v^*)} \right\} \text{ for each } x_i \in C \\
\text{new}^{Val}(\text{listref}_m v^*) (\lambda x_m: ref(Val). & (\lambda x_i: ref(Val). \\
(P \kappa') \dots))) \kappa & ((P[(\text{listref}_i v^*)/x_i]_{x_i \in V} \kappa') \dots)) \kappa)
\end{array}$$

Applied to a continuation K :

$$\begin{array}{l|l}
\langle \langle \mathcal{L}_1 \langle x_1 \dots x_m \rangle \llbracket P \rrbracket K \rangle ; \Sigma \rangle \rightarrow_{pcf}^* & \langle \langle \mathcal{R}_1(V, C) \llbracket P \rrbracket K \rangle ; \Sigma \rangle \rightarrow_{pcf}^* \\
\langle \langle K \mathcal{L}_2 \langle x_1 \dots x_m \rangle \llbracket P \rrbracket \rangle ; \Sigma \rangle & \langle \langle K \mathcal{R}_2(V, C) \llbracket P \rrbracket \rangle ; \Sigma \rangle \\
\mathcal{L}_2 \langle x_1 \dots x_m \rangle \llbracket P \rrbracket \stackrel{\text{def}}{=} & \mathcal{R}_2(V, C) \llbracket P \rrbracket \stackrel{\text{def}}{=} \\
\text{inr}(\lambda v^*: Val^*. \lambda \kappa': Val \rightarrow pr. & \text{inr}(\lambda v^*: Val^*. \lambda \kappa': Val \rightarrow pr. \\
\text{new}^{Val}(\text{listref}_1 v^*) (\lambda x_1: ref(Val). \dots & \text{new}^{Val}(\text{listref}_i v^*) \left. \vphantom{\text{new}^{Val}(\text{listref}_1 v^*)} \right\} \text{ for each } x_i \in C \\
\text{new}^{Val}(\text{listref}_m v^*) (\lambda x_m: ref(Val). & (\lambda x_i: ref(Val). \\
(P \kappa') \dots))) & ((P[(\text{listref}_i v^*)/x_i]_{x_i \in V} \kappa') \dots))
\end{array}$$

Closure Applied to an Argument List Q :

$$\begin{array}{l|l}
\langle \langle \text{out}_F(\mathcal{L}_2 \langle x_1 \dots x_m \rangle \llbracket P \rrbracket) (\lambda f: Val^* \rightarrow (Val \rightarrow pr) \rightarrow pr. & \langle \langle \text{out}_F(\mathcal{R}_2(V, C) \llbracket P \rrbracket) (\lambda f: Val^* \rightarrow (Val \rightarrow pr) \rightarrow pr. \\
f Q K) \rangle ; \Sigma \rangle & f Q K) \rangle ; \Sigma \rangle \\
\text{where } Q: Val^* \text{ and } K: Val \rightarrow pr & \text{where } Q: Val^* \text{ and } K: Val \rightarrow pr \\
\rightarrow_{pcf}^* \langle \langle \mathcal{L}_3 \langle x_1 \dots x_m \rangle \llbracket P Q K \rrbracket \rangle ; \Sigma \rangle & \rightarrow_{pcf}^* \langle \langle \mathcal{R}_3(V, C) \llbracket P Q K \rrbracket \rangle ; \Sigma \rangle \\
\mathcal{L}_3 \langle x_1 \dots x_m \rangle \llbracket P_1 P_2 P_3 \rrbracket \stackrel{\text{def}}{=} & \mathcal{R}_3(V, C) \llbracket P_1 P_2 P_3 \rrbracket \stackrel{\text{def}}{=} \\
\text{new}^{Val}(\text{listref}_1 P_2) (\lambda x_1: ref(Val). \dots & \text{new}^{Val}(\text{listref}_i P_2) \left. \vphantom{\text{new}^{Val}(\text{listref}_1 P_2)} \right\} \text{ for each } x_i \in C \\
\text{new}^{Val}(\text{listref}_m P_2) (\lambda x_m: ref(Val). & (\lambda x_i: ref(Val). \\
(P_1 P_3) \dots))) & ((P_1[(\text{listref}_i P_2)/x_i]_{x_i \in V} P_3) \dots)
\end{array}$$

Figure 10: Definitions of $\mathcal{L}_i, \mathcal{R}_i$ macros

constructing a relation $\overset{AE}{\rightsquigarrow}_V$ between terms that captures the differences between translations of immutable variable references. This is done in Figure 11. The subscript V is a list of variables which are treated as values on the right-hand side. We then construct a relation $\overset{AE}{\rightsquigarrow}_{V|Z}$ between metalanguage terms, defined in Figure 12, that characterizes the differences between terms which were originally related by $\overset{AE}{\rightsquigarrow}_V$ but which have been subjected to run-time substitutions – of locations for variables on the left-hand side and of values for variables on the right-hand side. We then extend $\overset{AE}{\rightsquigarrow}_{V|Z}$ to act on stores pointwise.

5. It is easy to show that if $V = \{x_1, \dots, x_n\}$ and $\Gamma[x_i: ref(Val)]_{i=1}^n \vdash M: \tau$ and

$$\Gamma[x_i: ref(Val)]_{i=1}^n \vdash M \overset{AE}{\rightsquigarrow}_V N: \tau$$

then $\Gamma[x_i: Val]_{i=1}^n \vdash N: \tau$, and that if $\Gamma[x: \tau'] \vdash M \overset{AE}{\rightsquigarrow}_V N: \tau$ and $\Gamma \vdash P \overset{AE}{\rightsquigarrow}_V P': \tau'$ and $x \notin V$, then $\Gamma \vdash M[P/x] \overset{AE}{\rightsquigarrow}_V N[P'/x]: \tau$.

6. Figure 13 defines the candidate relation, S_{AE} , by abstracting over corresponding pairs of terms which differ as outlined earlier. Each of the distinct ways in which corresponding metalanguage subterms might differ is handled by a substitution pair $(\sigma_{\mathcal{L}}^i, \sigma_{\mathcal{R}}^i)$. Each substitution has a domain $u_1^i, \dots, u_{n_i}^i$, which mark the locations of the corresponding subterms that differ as described in step 4(a)-(c) above.

$$\begin{array}{c}
\Gamma \vdash n \overset{AE}{\rightsquigarrow}_V n:int \quad (\text{CONG-INT}) \\
\\
\Gamma \vdash c \overset{AE}{\rightsquigarrow}_V c:\tau \quad (\tau \text{ from Figure 1}) \quad (\text{CONG-CONST}) \\
\\
\Gamma \vdash l^\tau \overset{AE}{\rightsquigarrow}_V l^\tau:ref(\tau) \quad (\text{CONG-LOC}) \\
\\
\frac{\Gamma[x:\tau'] \vdash M \overset{AE}{\rightsquigarrow}_V N:\tau \quad x \notin V}{\Gamma \vdash (\lambda x:\tau'.M) \overset{AE}{\rightsquigarrow}_V (\lambda x:\tau'.N):\tau' \rightarrow \tau} \quad (\text{CONG-ABS}) \\
\\
\frac{\Gamma \vdash M \overset{AE}{\rightsquigarrow}_V M':\tau' \rightarrow \tau \quad \Gamma \vdash N \overset{AE}{\rightsquigarrow}_V N':\tau'}{\Gamma \vdash (M N) \overset{AE}{\rightsquigarrow}_V (M' N'):\tau} \quad (\text{CONG-APP}) \\
\\
\frac{x \notin V \quad \Gamma(x) = \tau}{\Gamma \vdash x \overset{AE}{\rightsquigarrow}_V x:\tau} \quad (\text{MUTABLE-VAR}) \\
\\
\frac{\Gamma \vdash M \overset{AE}{\rightsquigarrow}_V N:Val \rightarrow pr \quad x \in V \quad \Gamma(x) = ref(Val)}{\Gamma \vdash (deref x M) \overset{AE}{\rightsquigarrow}_V (N x):pr} \quad (\text{IMMUTABLE-VAR})
\end{array}$$

Figure 11: Definition of $\overset{AE}{\rightsquigarrow}_V$

The last section of the definition of S_{AE} requires all of the closure bodies to be $\overset{AE}{\rightsquigarrow}_{V|Z}$ -related. In order to simplify the notation, we assume that all of the bound variables in the states we consider are distinct. This allows us to have a single $\overset{AE}{\rightsquigarrow}_{V|Z}$ relation for the whole simulation relation, rather than different V 's and Z 's for different closures.

7. For any Scheme expression e , any metalanguage context $C[-(Val \rightarrow pr) \rightarrow pr]$, and any store Σ ,

$$\langle\langle C[\llbracket e \rrbracket_{scm}] ; \Sigma \rangle\rangle S_{AE} \langle\langle C[\llbracket AE(e) \rrbracket_{cell}] ; \Sigma \rangle\rangle$$

To establish this, let the Z in the $\overset{AE}{\rightsquigarrow}_{V|Z}$ relation be $\langle \rangle$ (the empty sequence), let the σ^2 and σ^3 substitutions both be empty, let the σ^1 substitutions have elements for the translation of each lambda expression in e , and let the V in $\overset{AE}{\rightsquigarrow}_V$ contain the immutable local variables of the lambda-expressions in e . The proof uses the results of step 5 to deal with binding and substitution.

By choosing C appropriately, this implies that all the states in step 2 above will be in S_{AE} , as desired.

8. We must verify that S_{AE} is a bisimulation relation, using the definition of a bisimulation. This in turn requires an induction on the number of steps to an input-output operation, which is done using a case analysis on the form of M_0 and N_0 ; the interesting cases are when these terms have one of the u_j^i as a head variable. The most delicate step is to confirm that subterms of reference type reduce to the

same location constant when they are in $\overset{AE}{\rightsquigarrow}_{V|Z}$: If $M \overset{AE}{\rightsquigarrow}_{V|Z} N$, $M, N:ref(\tau)$, and $(M, N) \notin Z$, then

$$M \rightarrow_{pcf}^* l \Leftrightarrow N \rightarrow_{pcf}^* l$$

The proof is by induction on the length of the reduction.

9. Thus S_{AE} is a bisimulation relation that contains each of the required pairs of terms, and we are done.

7 Open Problems and Future Work

We believe that these techniques are powerful enough to treat other important optimizations, such as lambda-lifting in Scheme. We would like to consider more such transformations. In particular, we would like to consider the interaction between these techniques and those relying on flow analysis to collect global data about the program.

It is clear that the semantics of Idealized Algol in Section 5 is not fully abstract, because our contexts include procedures with input-output and non-local jumps. On the other hand, full abstraction for our Scheme subset is not a well-formulated question, since the initial context C_0 is unspecified. Our Scheme semantics is as abstract as possible, in that for any Scheme expressions e_1 and e_2 , if $\llbracket e_1 \rrbracket_{scm} \not\equiv_{\forall C} \llbracket e_2 \rrbracket_{scm}$, there is some choice of C_0 that distinguishes them. This remains a matter for further investigation.

Definition 6.3 ($\overset{AE}{\rightsquigarrow}_{V|Z}$ for metalanguage terms) For $V = \langle x_1, \dots, x_n \rangle$, $Z = \langle (l_1, P_1), \dots, (l_k, P_k) \rangle$,

$$\begin{aligned} \Gamma \vdash M \overset{AE}{\rightsquigarrow}_{V|Z} N : \tau \Leftrightarrow & \\ & \wedge k \leq n \\ & \wedge \Gamma(x_i) = \text{ref}(\text{Val}), \text{ for } 1 \leq i \leq n \\ & \wedge P_i : \text{Val}, \text{ for } 1 \leq i \leq k \\ & \wedge \exists M', N' \text{ s.t.} \\ & \quad \wedge M = M'[l_1/x_1, \dots, l_k/x_k] \\ & \quad \wedge N = N'[P_1/x_1, \dots, P_k/x_k] \\ & \wedge \Gamma \vdash M' \overset{AE}{\rightsquigarrow}_V N' : \tau \end{aligned}$$

Definition 6.4 ($\overset{AE}{\rightsquigarrow}_{V|Z}$ for stores)

$$\begin{aligned} \Gamma \vdash \Sigma'_M \overset{AE}{\rightsquigarrow}_{V|Z} \Sigma_N \Leftrightarrow & \\ & \wedge \text{dom}(\Sigma'_M) = \text{dom}(\Sigma_N) \\ & \wedge \forall \tau, l^\tau \in \text{dom}(\Sigma'_M) . \Gamma \vdash \Sigma'_M(l) \overset{AE}{\rightsquigarrow}_{V|Z} \Sigma_N(l) : \tau \end{aligned}$$

Figure 12: The Relations $\overset{AE}{\rightsquigarrow}_{V|Z}$ between Terms and between Stores

$$\begin{aligned} S_{AE} \stackrel{\text{def}}{=} & \{ (\langle M_0 ; \Sigma_M \rangle \sigma_{\mathcal{L}}, \langle N_0 ; \Sigma_N \rangle \sigma_{\mathcal{R}}) \mid \\ & \sigma_{\mathcal{L}} = \sigma_{\mathcal{L}}^1 \cup \sigma_{\mathcal{L}}^2 \cup \sigma_{\mathcal{L}}^3 \\ & \sigma_{\mathcal{R}} = \sigma_{\mathcal{R}}^1 \cup \sigma_{\mathcal{R}}^2 \cup \sigma_{\mathcal{R}}^3 \\ & \text{fv}(M, \text{rng}(\Sigma)) \subseteq \{u_1^1, \dots, u_{n_1}^1, u_1^2, \dots, u_{n_2}^2, u_1^3, \dots, u_{n_3}^3\} \\ (\mathcal{L}_1, \mathcal{R}_1) \text{ pairs} & \begin{cases} \sigma_{\mathcal{L}}^1 = [u_1^1 \mapsto \mathcal{L}_1(X_1^1)[M_1^1], \dots, u_{n_1}^1 \mapsto \mathcal{L}_1(X_{n_1}^1)[M_{n_1}^1]] \\ \sigma_{\mathcal{R}}^1 = [u_1^1 \mapsto \mathcal{R}_1(V_1^1, C_1^1)[N_1^1], \dots, u_{n_1}^1 \mapsto \mathcal{R}_1(V_{n_1}^1, C_{n_1}^1)[N_{n_1}^1]] \\ n_1 = \#\sigma_{\mathcal{L}}^1 = \#\sigma_{\mathcal{R}}^1 \end{cases} \\ (\mathcal{L}_2, \mathcal{R}_2) \text{ pairs} & \begin{cases} \sigma_{\mathcal{L}}^2 = [u_1^2 \mapsto \mathcal{L}_2(X_1^2)[M_1^2], \dots, u_{n_2}^2 \mapsto \mathcal{L}_2(X_{n_2}^2)[M_{n_2}^2]] \\ \sigma_{\mathcal{R}}^2 = [u_1^2 \mapsto \mathcal{R}_2(V_1^2, C_1^2)[N_1^2], \dots, u_{n_2}^2 \mapsto \mathcal{R}_2(V_{n_2}^2, C_{n_2}^2)[N_{n_2}^2]] \\ n_2 = \#\sigma_{\mathcal{L}}^2 = \#\sigma_{\mathcal{R}}^2 \end{cases} \\ (\mathcal{L}_3, \mathcal{R}_3) \text{ pairs} & \begin{cases} \sigma_{\mathcal{L}}^3 = [u_1^3 \mapsto \mathcal{L}_3(X_1^3)[M_{1,1}^3 \ M_{1,2}^3 \ M_{1,3}^3], \dots, u_{n_3}^3 \mapsto \mathcal{L}_3(X_{n_3}^3)[M_{n_3,1}^3 \ M_{n_3,2}^3 \ M_{n_3,3}^3]] \\ \sigma_{\mathcal{R}}^3 = [u_1^3 \mapsto \mathcal{R}_3(V_1^3, C_1^3)[N_{1,1}^3 \ N_{1,2}^3 \ N_{1,3}^3], \dots, u_{n_3}^3 \mapsto \mathcal{R}_3(V_{n_3}^3, C_{n_3}^3)[N_{n_3,1}^3 \ N_{n_3,2}^3 \ N_{n_3,3}^3]] \\ n_3 = \#\sigma_{\mathcal{L}}^3 = \#\sigma_{\mathcal{R}}^3 \end{cases} \\ \exists \Gamma, \Sigma'_M \wedge & \\ \exists V = \langle x_1, \dots, x_n \rangle \wedge & \\ \exists Z = \langle (l_1, P_1), \dots, (l_k, P_k) \rangle, k \leq n \text{ s.t.} & \\ \text{for each } (M, N) \in \{(M_j^i, N_j^i)\} \cup \{(M_0, N_0)\}, & \\ \wedge \Gamma \vdash M \overset{AE}{\rightsquigarrow}_{V|Z} N : \tau & \\ \wedge \Sigma_M = \Sigma'_M[l_1 \mapsto P'_1, \dots, l_k \mapsto P'_k] & \\ \wedge \Gamma \vdash \Sigma'_M \overset{AE}{\rightsquigarrow}_{V|Z} \Sigma_N & \\ \wedge \Gamma \vdash P'_i \overset{AE}{\rightsquigarrow}_{V|Z} P_i, 1 \leq i \leq k & \\ \} & \end{aligned}$$

Figure 13: Definition of S_{AE}

Our formalism should make it relatively easy to extend our input-output primitives to model communicating programs; how this will change the theoretical results remains to be seen.

8 Conclusions

We have given a denotational semantics for programs with input-output and a heap-allocated store. Because the framework is denotational, it encompasses a variety of source languages, and we need not reprove complex technical results like the Context Lemma each time we change the source language. Because equivalence in the model is based on input-output behavior, we can classify the behavior even of programs that do not terminate. (Space precludes an example here, see [TW96]). Similarly, formulating equivalence in terms of observable input-output behavior sidesteps most of the problems associated with store semantics.

Because equivalence in the meta-language is based on input-output behavior, proofs of equivalence ultimately rely on bisimulation arguments. In our experience, these arguments are intuitively satisfying, since they formalize the notion that the two programs being considered “stay in sync” as they compute. This approach imposes a lesser proof burden than do methods based on logical relations, and also extends to untyped or recursively typed languages.

We have used the technique to obtain a new result: the correctness of assignment elimination in Scheme. We believe that our methods will extend to a variety of transformations used in Scheme and ML compilers.

References

- [Abr90] Samson Abramsky. The lazy lambda calculus. In David A. Turner, editor, *Research Topics in Functional Programming*, pages 65–116. Addison-Wesley, 1990.
- [App92] Andrew W. Appel. *Compiling with Continuations*. Cambridge University Press, Cambridge, 1992.
- [CG95] R. L. Crole and A. D. Gordon. A sound metalogical semantics for input/output effects. In Leszek Pacholski and Jerzy Tiuryn, editors, *Computer science logic: 8th workshop, CSL '94*, volume 933 of *Lecture Notes in Computer Science*, pages 339–353, Berlin, Heidelberg, and New York, 1995. Springer-Verlag.
- [CH94] William D. Clinger and Lars Thomas Hansen. Lambda, the ultimate label, or a simple optimizing compiler for scheme. In *Proc. 1994 ACM Symposium on Lisp and Functional Programming*, pages 128–139, 1994.
- [Fel87] 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.
- [Gor94] Andrew D. Gordon. *Functional Programming and Input/Output*. Cambridge University Press, Cambridge, 1994.
- [Gor95] Andrew D. Gordon. Bisimilarity as a theory of functional programming. In *Proceedings of 11th Conference on Mathematical Foundations of Programming Semantics*, 1995.
- [GSR95] J. D. Guttman, V. Swarup, and J. Ramsdell. The VLISP verified scheme system. *Lisp and Symbolic Computation*, 8(1/2):33–110, 1995.
- [Gun92] Carl A. Gunter. *Semantics of Programming Languages: Structures and Techniques*. MIT Press, Cambridge, MA, 1992.
- [KH89] Richard Kelsey and Paul Hudak. Realistic compilation by program transformation. In *Conf. Rec. 16th ACM Symposium on Principles of Programming Languages*, pages 281–292, 1989.
- [KKsR+86] David A. Kranz, Richard Kelsey, Jonathan A. Rees, Paul Hudak, James Philbin, and Norman I. Adams. Orbit: An optimizing compiler for scheme. In *Proceedings SIGPLAN '86 Symposium on Compiler Construction*, 1986. *SIGPLAN Notices* 21(7), July, 1986, 219–223.
- [Mil77] Robin Milner. Fully abstract models of typed lambda-calculi. *Theoretical Computer Science*, 4:1–22, 1977.
- [Mor68] James H. Morris, Jr. *Lambda Calculus Models of Programming Languages*. PhD thesis, MIT, Cambridge, MA, 1968.
- [MS88] Albert R. Meyer and Kurt Sieber. Towards fully abstract semantics for local variables: Preliminary report. In *Conf. Rec. 15th ACM Symposium on Principles of Programming Languages*, pages 191–203, 1988.
- [MT91] Ian A. Mason and Carolyn L. Talcott. Equivalence in functional languages with effects. *Journal of Functional Programming*, 1:287–327, 1991.
- [Ode94] Martin Odersky. A functional theory of local names. In *Conference Record of POPL '94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 48–59, 1994.
- [Ole85] F. J. Oles. Type categories, functor categories and block structure. In M. Nivat and J. C. Reynolds, editor, *Algebraic Semantics*, pages 543–574. Cambridge University Press, 1985.

- [ORW95] Dino P. Oliva, John D. Ramsdell, and Mitchell Wand. The VLISP verified prescheme compiler. *Lisp and Symbolic Computation*, 8(1/2):111–182, 1995.
- [OT95] P. W. O’Hearn and R. D. Tennent. Parametricity and local variables. *Journal of the ACM*, 42(3):658–709, May 1995.
- [Par81] David Park. Concurrency and automata on infinite sequences. In Peter Deussen, editor, *Theoretical Computer Science: 5th GI-Conference, Karlsruhe*, volume 104 of *Lecture Notes in Computer Science*, pages 167–183, Berlin, Heidelberg, and New York, March 1981. Springer-Verlag.
- [Plo75] Gordon D. Plotkin. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science*, 1:125–159, 1975.
- [Plo77] Gordon D. Plotkin. LCF considered as a programming language. *Theoretical Computer Science*, 5:223–255, 1977.
- [PS93] Andrew M. Pitts and Ian D. B. Stark. Observable properties of higher order functions that create local names, or: What’s new? In *Mathematical Foundations of Computer Science, Proc. 18th Intl. Symp, Gdansk, 1993*, volume 711 of *Lecture Notes in Computer Science*, pages 122–141, Berlin, Heidelberg, and New York, 1993. Springer-Verlag.
- [Ram96] John D. Ramsdell. personal communication, July 1996.
- [RC⁺86] Jonathan A. Rees, William C. Clinger, et al. Revised³ report on the algorithmic language scheme. *SIGPLAN Notices*, 21(12):37–79, December 1986.
- [Sta94] Ian Stark. *Names and Higher-Order Functions*. PhD thesis, University of Cambridge, December 1994. Also published as Technical Report 363, University of Cambridge Computer Laboratory.
- [TW96] Jerzy Tiuryn and Mitchell Wand. Untyped lambda-calculus with input-output. In H. Kirchner, editor, *Trees in Algebra and Programming: CAAP’96, Proc. 21st International Colloquium*, volume 1059 of *Lecture Notes in Computer Science*, pages 317–329, Berlin, Heidelberg, and New York, April 1996. Springer-Verlag.
- [WO92] Mitchell Wand and Dino P. Oliva. Proving the correctness of storage representations. In *Proc. 1992 ACM Symposium on Lisp and Functional Programming*, pages 151–160, 1992.