

# Proving the Correctness of Storage Representations

Mitchell Wand and Dino P. Oliva\*

College of Computer Science  
Northeastern University  
360 Huntington Avenue, 161CN  
Boston, MA 02115, USA  
wand@corwin.ccs.northeastern.edu  
oliva@corwin.ccs.northeastern.edu

## Abstract

Conventional techniques for semantics-directed compiler derivation yield abstract machines that manipulate trees. However, in order to produce a real compiler, one has to represent these trees in memory. In this paper we show how the technique of *storage-layout relations* (introduced by Hannan [7]) can be applied to verify the correctness of storage representations in a very general way. This technique allows us to separate denotational from operational reasoning, so that each can be used when needed. As an example, we show the correctness of a stack implementation of a language including dynamic catch and throw. The representation uses static and dynamic links to thread the environment and continuation through the stack. We discuss other uses of these techniques.

## 1 Introduction

Typical semantically-derived compiler systems [2, 3, 7, 18, 19] translate the parse trees of the source language into the language of an abstract byte-code machine. This abstract byte code machine typically manipulates trees that represent the environment, continuation, local stack, etc. However, in order to turn this into a real compiler, one has to represent these trees in machine storage. There are a variety of strategies for doing this [8], but this representation step remains a critical and untrustworthy part of compiler-development technology. Indeed, in two recent compiler-correctness projects [9, 10], it turned out to be the most difficult part of the task.

Hannan [7] has shown how, under certain restrictions, one can automatically transform an abstract machine into a storage machine. In this paper we show how similar techniques can be used to verify the correctness of storage representations in a far more general way. As in [18, 19, 20], we do not attempt to derive these representations automatically. Instead, our goal is to build a framework in which com-

piler writers can prove the correctness of arbitrarily clever human-invented representations.

The representations are formalized by the use of representations that we call *storage-layout relations*. Storage-layout relations allow a compiler writer to specify the representation decisions he or she has made. Storage-layout relations were introduced in [7]. This work extends that of [7] by liberating it from the particular architecture considered there, and by considering more complex representation strategies, such as multiple terms threaded through a single stack.

As an application, we consider a simple (not necessarily functional) language with continuation semantics and **catch** and **throw** operators. The language is designed to obey *stack discipline*: that is, the language can be implemented using only a stack. Our goal is to prove the correctness of this assertion. In fact, we will show that an ordinary stack representation, with static and dynamic chain pointers, constitutes a representation of the three trees (display, local stack, and continuation) of the abstract bytecode machine.

We first give a semantics and a simple byte-code compiler using the techniques of [3]. This compiler produces code for an abstract machine that manipulates combinator trees representing the environment, continuation, and a local stack. We then consider two successively more concrete implementations of this abstract machine. In the first version, we represent the display, local stack, and continuation in a single stack kept in memory, while the program remains as a tree. In the second version, we represent programs linearly as well.

This example shows how we can separate denotational from operational reasoning in a compiler correctness proof, using each to best advantage: the translation from source code to combinator trees is proved correct using structural induction and denotational reasoning, the translation from the abstract machine to the concrete machines is proved correct using induction on length of execution and operational reasoning. The link between the two regimes is obtained using general results about the completeness of reduction strategies in the  $\lambda$ -calculus. We believe that this strategy will be of general use.

\*Work supported by the National Science Foundation and DARPA under grants CCR-9002253 and CCR-9014603

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

1992 ACM LISP & F.P.-6/92/CA

© 1992 ACM 0-89791-483-X/92/0006/0151...\$1.50

## 2 Comparison with Previous Work

Early proofs of compiler correctness, such as [11], did the translation from source to target code in a single step. In modern terminology, such proofs used a source semantics in a direct style and a target semantics in a continuation style or an operational style.

The restriction to direct-style semantics made it infeasible to use these techniques to prove the correctness of compilers for languages with non-local jumps. Therefore, in the 1970's, attention turned to correctness proofs in which the source language was given using continuation semantics. Such proofs typically involved complex arguments about inclusive predicates that related the domains in the source and target semantics [12, 17]. These proofs were quite complex because they sought to model reasoning that was essentially operational (induction on reduction sequences) by methods that were essentially denotational (induction on approximations to inverse limit domains).

At the beginning of the 1980's, Wand [18] showed how to avoid the use of inclusive predicates by modelling the target machine in the same semantic domains as the source semantics. In this paradigm, one showed that the meanings of the source and target programs were the same. Because the semantics was compositional, this could be done by a simple structural induction. This approach was refined by Clinger [3] to prove the correctness of a compiler for Scheme.

In this approach, the target language was a sublanguage of the  $\lambda$ -calculus, designed to be closed under reduction. The target machine simply had to simulate the reduction of the target program. One could use any desired mechanism to do this. The simulation proof could use operational reasoning. The reduction strategy was known to be complete, so the machine was guaranteed to reduce the target program to a normal form if one existed.

The Wand-Clinger approach simplified the inclusive-predicate approach in two ways. first, it allowed the denotational reasoning to be done in a single  $\lambda$ -model that gave denotations to both source and target terms, and second, it separated operational from denotational reasoning, so each could be used where appropriate. This paper extends the operational part of the proof by showing how the tree-rewriting systems can be implemented in a conventional storage model, and how that implementation can be proved correct using operational reasoning.

Raoult and Sethi [14] and Schmidt [16] considered the issue of single-threading: when a particular term (or class of terms) could be represented by a single destructively-updated quantity, but they did not consider the representation of such quantities in a linear store. Hannan [7] showed how certain classes of tree-rewriting systems could be transformed into a stored-program form, including the linear representation of programs, much like our stored-program machine in Section 8, but he represented all mutable data either in the heap or in multiple stacks, rather than considering the single-stack representation considered here.

## 3 Source Language

Our source language is a simple language with continuation semantics. Its values are integers, recursive procedures, and continuations. Operations on integers are addition and conditional branch (on zero). Operations on continuations are

catch and throw. Catch captures the current continuation and binds it to an identifier. Throw invokes a bound continuation on a value, discarding the current continuation.

Procedures are recursive, using a local identifier for self-reference. Any kind of value may be passed to a procedure, but procedures may return only integers, so that the language can be implemented using only a stack. The semantics is given relative to an unspecified domain of command continuations, so that non-functional primitives may be added.

The language is given by

$$E ::= Id \mid (E + E) \mid (\text{if } E \ E \ E) \mid (\text{catch } Id \ E) \\ \mid (\text{throw } Id \ E) \mid (E \ E) \mid (\text{rec } Id(Id) \ E)$$

The semantics of the language is given by a valuation  $\mathcal{E}[-]$ . We typically think of  $\mathcal{E}[-]$  as a syntactic transformation, whose output is a  $\lambda$ -term of the appropriate type. That is, if we write  $\underline{T}$  for the set of  $\lambda$ -terms of type  $T$ , then we think of  $\mathcal{E}$  as having type  $E \rightarrow \underline{U} \rightarrow \underline{K} \rightarrow \underline{C}$ , where  $U$ ,  $K$ , and  $C$  are types given by

$$C = \text{Command Continuations} \\ V = \text{Int} + [V \rightarrow K \rightarrow C] + [V \rightarrow C] \\ K = V \rightarrow C \\ U = Id \rightarrow V$$

The lattice-theoretical details are not relevant, since we are dealing with  $\lambda$ -terms rather than denotations; we assume only the existence of a non-trivial model with this type theory. We write the summands of  $V$  as  $\langle \text{int}, n \rangle$ ,  $\langle \text{proc}, p \rangle$  or  $\langle \text{cont}, \kappa \rangle$ . We assume the existence of combinators that can distinguish these pairwise; this is easy to do by coding

The semantics is given by

$$\mathcal{E}[I] = \lambda\rho\kappa.\kappa(\rho[I]) \\ \mathcal{E}[(e_1 + e_2)] = \lambda\rho\kappa.\mathcal{E}[e_1]\rho(\lambda v_1.\mathcal{E}[e_2]\rho(\lambda v_2 \ \text{sum } v_1 v_2 \kappa)) \\ \mathcal{E}[(\text{if } e_1 \ e_2 \ e_3)] = \lambda\rho\kappa.\mathcal{E}[e_1]\rho(\lambda v_1.v_1 = \langle \text{int}, 0 \rangle \rightarrow \\ \mathcal{E}[e_3]\rho\kappa, \\ \mathcal{E}[e_2]\rho\kappa) \\ \mathcal{E}[(\text{rec } F(I) \ e)] = \lambda\rho\kappa.\kappa(\text{proc}, \\ (\text{fix } (\lambda p \lambda v \kappa.\mathcal{E}[e](\rho[(\text{proc}, p])/F.v/I]) \\ (\lambda v.\text{return } \kappa v)))) \\ \mathcal{E}[(e_1 \ e_2)] = \lambda\rho\kappa.\mathcal{E}[e_1]\rho(\lambda v_1.\mathcal{E}[e_2]\rho \\ (\lambda v_2.\text{apply}_p \ v_1 v_2 \kappa)) \\ \mathcal{E}[(\text{catch } I \ e)] = \lambda\rho\kappa.\mathcal{E}[e]\rho[\langle \text{cont}, \kappa \rangle/I]\kappa \\ \mathcal{E}[(\text{throw } I \ e)] = \lambda\rho\kappa.\mathcal{E}[e]\rho(\lambda v_1.\mathcal{E}[I]\rho(\lambda v_2 \ \text{apply}_k \ v_1 v_2))$$

The operators *sum*, *apply<sub>p</sub>*, *apply<sub>k</sub>*, and *return* all check their arguments for membership in the appropriate summand. In particular, *return* only allows procedures to return integers; this prevents upward funargs and allows a stack-discipline implementation.

$$\text{sum} = \lambda v_1 \ v_2 \ \kappa.(v_1 = \langle \text{int}, n_1 \rangle \rightarrow \\ (v_2 = \langle \text{int}, n_2 \rangle \rightarrow \\ \kappa(\langle \text{int}, n_1 + n_2 \rangle), \\ (\text{error "Non-integer second arg"}), \\ (\text{error "Non-integer first arg"}))$$

$$\text{apply}_k = \lambda v_1 \ v_2.(v_2 = \langle \text{cont}, \kappa \rangle \rightarrow \\ (v_1 = \langle \text{int}, n \rangle \rightarrow \kappa(v_1), \\ (\text{error "Non-integer second arg"}), \\ (\text{error "Non-continuation first arg"}))$$



we can give the bytecode machine an operational semantics by thinking of it as a 4-register machine and regarding each of the equations for the instruction combinators as rewriting rules. The resulting abstract machine is shown in Figure 1. We can state the soundness of this machine as follows:

**Theorem 2** *If  $\langle \pi, u, \zeta, \kappa \rangle \Longrightarrow \langle \pi', u', \zeta', \kappa' \rangle$ , then  $\pi u \zeta \kappa \rightarrow \pi' u' \zeta' \kappa'$ .*

It can also be shown that the machine is complete with respect to either call-by-value or call-by-name reduction [13], since these reduction strategies coincide on the terms manipulated by the machine, and the machine emulates these reduction strategies.

So far this is just an example of the method set out in [3, 18] (see also [5] for some more extended examples). We next turn to our main topic: the representation of these trees in a linear store.

## 5 Stack Representation

This abstract machine simulates the reduction of the terms produced by the compiler. However, it is still a tree manipulation system. In order to implement this abstract machine on a real computer, one must represent these trees in memory. It is the correctness of this representation that is our main concern in this paper.

In particular, our language was carefully designed to obey *stack discipline*, that is, the language can be implemented using only a stack. Our goal is to prove the correctness of this assertion. In fact, we will show that an ordinary stack representation, with static and dynamic chain pointers, constitutes a representation of the three trees (display, local stack, and continuation) of the abstract bytecode machine.

We will do this in two stages. In the first stage, we will represent the display, local stack, and continuation in a single stack kept in memory, whereas the program  $\pi$  will remain as a tree. In the second stage, we will represent programs linearly as well.


Our concrete machine will be a 5-tuple

$$\langle \pi, up, sp, kp, \sigma \rangle$$

consisting of a program  $\pi$ , a display pointer  $up$ , a stack pointer  $sp$ , a continuation pointer  $kp$ , and a store  $\sigma$ . The pointers are addresses in the store, satisfying  $sp \geq kp > up$ . We assume that each cell of the store can hold a tag (**proc** or **int**), a pointer, and occasionally a program. (The program will be condensed to a pointer in the second step).


We now sketch the representation of the abstract machine's quantities in the concrete machine.

1. Values (integers or closures) are represented as single words in memory, including their tags. Continuations bound by **catch** are represented as single words containing a tag and a pointer to the continuation structure.
2. A local stack ( $v_n \dots v_{n-1} :: \dots :: v_1 :: \langle \rangle$ ) is specified by giving upper and lower bounds as indices, and is represented as a series of words in the stack




This notation is intended to mean that  $ub$  and  $lb$  are pointers into the stack, that  $ub$  points higher in the stack than  $lb$ , and that the values from position  $ub$ , moving downward, are  $v_n, v_{n-1}, \dots, v_1$ .

3. A run-time environment (display) extension of the form  $ext_R \langle v, \langle \text{proc}, \text{closure } u' \pi' \rangle \rangle u'$  is specified by a pointer  $p$  into the stack:



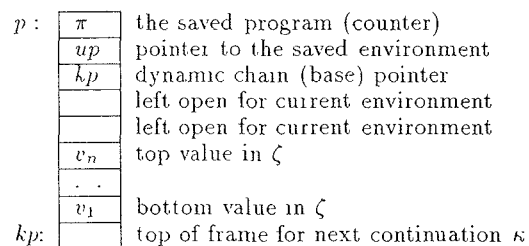
where  $p'$  is a pointer to the display  $u'$  (ie, a static chain pointer). Here the second slot is used both for the closure that is bound to the local name for the procedure and for the static chain that is in the closure.

4. A run-time environment (display) extension of the form  $ext_R \langle \langle \text{cont}, (\text{cont } \pi' u \zeta \kappa) \rangle, \langle \rangle \rangle u$  is specified by a pointer  $p$  into the stack.



where  $\langle \text{cont}, p' \rangle$  represents  $(\text{cont } \pi' u \zeta \kappa)$  (according to item 1 above) and  $p''$  is a pointer to the display  $u$  (ie, a static chain pointer).

5. The continuation  $(\text{cont } \pi u \zeta \kappa)$  is represented by a pointer  $p$  into the stack as follows:



## 6 Storage Layout Relations

This pictorial specification is adequate for compiler writers, but it is not quite formal enough to allow us to do proofs. We formalize these pictures as a family of relations, called *storage layout relations*, that tell precisely when a pointer into a stack corresponds to a given abstract structure. In general, a storage layout relation will be a ternary relation  $\sigma \models p \simeq m$ , which means that pointer  $p$  corresponds to abstract structure  $m$  in store  $\sigma$ . We will consistently put the concrete object on the left of the equivalence, and read  $\simeq$  as "corresponds to."

We will have four storage layout relations, one each for values, stacks, environments, and continuations. These will be defined by simultaneous induction on the terms for values, stacks, environments, and continuations. Since these are relations on terms (not their denotations), the existence of the relations is assured.

The definition of the relations is now a straightforward transcription of the data in the diagrams above.

Abstract machine:

$$\begin{aligned}
& \langle (\text{fetch } l \pi), u, \zeta, \kappa \rangle \Longrightarrow \langle \pi, u, (u(l) :: \zeta), \kappa \rangle \\
& \langle (\text{add } \pi), u, (\langle \text{int}, n_2 \rangle :: \langle \text{int}, n_1 \rangle :: \zeta), \kappa \rangle \Longrightarrow \langle \pi, u, (\langle \text{int}, n_1 + n_2 \rangle :: \zeta), \kappa \rangle \\
& \langle (\text{brz } \pi' \pi), u, v :: \zeta, \kappa \rangle \Longrightarrow \langle (\text{choose } v \pi' \pi), u, \zeta, \kappa \rangle \\
& \langle (\text{save-cont } \pi' \pi), u, \zeta, \kappa \rangle \Longrightarrow \langle \pi, (\text{ext}_R \langle \langle \text{cont}, (\text{cont } \pi' u \zeta \kappa) \rangle, \langle \rangle) u, \\
& \quad \langle \rangle, (\text{cont } \pi' u \zeta \kappa) \rangle \\
& \langle (\text{throw}), u, \langle \text{cont}, (\text{cont } \pi' u' \zeta' \kappa') \rangle :: \langle \text{int}, n \rangle :: \zeta, \kappa \rangle \Longrightarrow \langle \pi', u', (\langle \text{int}, n \rangle :: \zeta'), \kappa' \rangle \\
& \quad \langle (\text{close } \pi' \pi), u, \zeta, \kappa \rangle \Longrightarrow \langle \pi, u, (\langle \text{proc}, \text{closure } u \pi' \rangle :: \zeta), \kappa \rangle \\
& \quad \langle (\text{rts}), u, (\langle \text{int}, n \rangle :: \zeta), (\text{cont } \pi' u' \zeta' \kappa') \rangle \Longrightarrow \langle \pi', u', (\langle \text{int}, n \rangle :: \zeta'), \kappa' \rangle \\
& \quad \langle (\text{jsr } \pi), u, (v \cdot \langle \text{proc}, \text{closure } u' \pi' \rangle :: \zeta), \kappa \rangle \Longrightarrow \langle \pi', \text{ext}_R \langle v, \langle \text{proc}, \text{closure } u' \pi' \rangle \rangle u', \\
& \quad \langle \rangle, (\text{cont } \pi u \zeta \kappa) \rangle
\end{aligned}$$

Auxiliaries:

$$(\text{choose } v \pi' \pi) = (v = \langle \text{int}, 0 \rangle) - \pi, \pi'$$

Figure 1: Actions of the abstract machine.

• Relating Pointers and Tagged Values

$$\begin{aligned}
\sigma \models_V p \simeq v & \iff \\
& \text{either } v = \langle \text{int}, n \rangle = \sigma(p) \\
& \text{or } v = \langle \text{proc}, (\text{closure } u \pi) \rangle \\
& \quad \text{and } \sigma(p) = \langle \text{proc}, (\text{closure } u \pi) \rangle \\
& \quad \text{and } u p < p \text{ and } \sigma \models_V u p \simeq u \\
& \text{or } v = \langle \text{cont}, (\text{cont } \pi u \zeta \kappa) \rangle \\
& \quad \text{and } \sigma(p) = \langle \text{cont}, p' \rangle \\
& \quad \text{and } p' < p \\
& \quad \text{and } \sigma \models_K p' \simeq (\text{cont } \pi u \zeta \kappa)
\end{aligned}$$

• Relating Pointers and Stacks

$$\begin{aligned}
\sigma \models_Z \langle p, p' \rangle \simeq \zeta & \iff \\
& \text{either } \zeta = \langle \rangle \text{ and } p = p' \\
& \text{or } \zeta = (z :: \zeta') \text{ and } p > p' \text{ and } \sigma \models_V p \simeq z \\
& \quad \text{and } \sigma \models_Z \langle p-1, p' \rangle \simeq \zeta'
\end{aligned}$$

• Relating Pointers and Environments

$$\begin{aligned}
\sigma \models_U p \simeq u & \iff \\
& \text{either } u = \text{emptydisplay} \text{ and } p = -1 \\
& \text{or } u = \text{ext}_R \langle v, \langle \text{proc}, \text{closure } u' \pi' \rangle \rangle u' \\
& \quad \text{and } \sigma \models_V p \simeq v \\
& \quad \text{and } \sigma \models_V p-1 \simeq \langle \text{proc}, \text{closure } u' \pi' \rangle \\
& \quad \text{and } \sigma(p-1).u < p \\
& \quad \text{and } \sigma \models_V \sigma(p-1).u \simeq u' \\
& \text{or } u = (\text{ext}_R \langle \langle \text{cont}, (\text{cont } \pi' u' \zeta \kappa) \rangle, \langle \rangle \rangle u') \\
& \quad \text{and } \sigma(p) = \langle \text{cont}, p' \rangle \\
& \quad \text{and } p' = p + 3 \\
& \quad \text{and } \sigma \models_K p' \simeq (\text{cont } \pi u \zeta \kappa) \\
& \quad \text{and } \sigma(p-1).u < p \\
& \quad \text{and } \sigma \models_V \sigma(p-1).u \simeq u'
\end{aligned}$$

Here we use the notation  $\sigma(p).u$  to denote the contents of environment-pointer field of location  $p$  of  $\sigma$ . We

will use similar notation for field selectors throughout. Note that this relation specifies that the static chain is kept in the same field of the second word of the environment representation, whichever variation (**proc** or **cont**) is used.

The  $p+3$  case arises when a *save-cont* instruction is executed. this instruction pushes a new continuation on the stack, using the format in item 5 above, and binds the continuation in the vacant environment frame.

• Relating Pointers and Continuations

$$\begin{aligned}
\sigma \models_K p \simeq \kappa & \iff \\
& \text{either } \kappa = \text{initcont} \text{ and } p = 0 \\
& \text{or } \kappa = (\text{cont } \pi u \zeta \kappa') \\
& \quad \text{and } \sigma(p) = \pi \\
& \quad \text{and } \sigma(p-1) < p \text{ and } \sigma \models_U \sigma(p-1) \simeq u \\
& \quad \text{and } \sigma(p-2) < p \text{ and } \sigma \models_K \sigma(p-2) \simeq \kappa' \\
& \quad \text{and } \sigma \models_Z (p-5, \sigma(p-2)) \simeq \zeta
\end{aligned}$$

## 7 The Proof

We begin by stating some useful lemmas about these relations. The most important of these lemmas state that the changing locations in the stack above a pointer does not change the quantity that the pointer represents. Let  $\sigma =_p \sigma'$  denote the proposition  $\forall x, 0 \leq x \leq p \implies \sigma(x) = \sigma'(x)$ . Then we have the following:

- Lemma 1**
1. If  $\sigma =_p \sigma'$  &  $\sigma \models_V p \simeq v$  then  $\sigma' \models_V p \simeq v$ .
  2. If  $\sigma =_p \sigma'$  &  $\sigma \models_Z \langle p, p' \rangle \simeq \zeta$  then  $\sigma' \models_Z \langle p, p' \rangle \simeq \zeta$ .
  3. If  $\sigma =_{p+3} \sigma'$  &  $\sigma \models_U u p \simeq u$  then  $\sigma' \models_U u p \simeq u$ .
  4. If  $\sigma =_{kp} \sigma'$  &  $\sigma \models_K k p \simeq \kappa$  then  $\sigma' \models_K k p \simeq \kappa$ .

*Proof:* By induction on the definition of  $\simeq$ . For stacks, the condition  $\sigma \models_Z \langle p, p' \rangle$  implies  $p' \leq p$ . For environments, we need to write  $up+3$  because the first slot might contain a continuation, we potentially can point 3 slots up in the stack.  $\square$

We can now define the correspondence between abstract and concrete machine states

**Definition 1** We say a concrete machine state  $\langle \pi, up, sp, kp, \sigma \rangle$  corresponds to an abstract machine state  $\langle \pi', u, \zeta, \kappa \rangle$  (written  $\langle \pi, up, sp, kp, \sigma \rangle \simeq \langle \pi', u, \zeta, \kappa \rangle$ ) if and only if the following conditions are satisfied:

1.  $\pi = \pi'$ ,
2.  $\sigma \models_U up \simeq u$ ,
3.  $\sigma \models_Z \langle sp, kp \rangle \simeq \zeta$ , and
4.  $\sigma \models_K kp \simeq \kappa$ .

This says (among other things) that the local stack  $\zeta$  is kept at the top of the stack, and is immediately followed by the representation of  $\kappa$ ; the display  $u$  is threaded through the stack but need not be near the top.

With this definition of correspondence between machine states, we can work out the action of our instructions on the concrete machine. Figure 2 shows the most complex example, the *jsr* instruction. Here  $sp'$ ,  $up'$ , and  $kp'$  denote the new values of the concrete machine registers. Since  $sp' = kp'$ , the new value of  $\zeta$  is the empty stack  $\langle \rangle$

A diagram such as that in Figure 2 can be translated into a formal representation more suitable for formal reasoning, similar to the definition of the abstract machine in Section 4. This is done in Figure 3. Here we use  $v.kp$  to extract the continuation field from a stored continuation, and  $v.\pi$  to extract the program field from a stored closure.

We can show that the following invariant is preserved by the machine:

**Lemma 2** In all machine states,  $sp \geq kp \geq up + 3$ .

The correctness of the auxiliary *find* is also shown:

**Lemma 3** If  $\sigma \models_U up \simeq u$  then

$$\sigma \models_V (\text{find } l \text{ up } \sigma) \simeq u(l)$$

Now we can show the main theorem, which asserts that the concrete machine correctly simulates the abstract machine:

**Theorem 3** Let  $A_1$  and  $A_2$  be states of the abstract machine and  $C_1$  and  $C_2$  be states of the concrete machine. If  $C_1 \simeq A_1$ ,  $A_1 \Longrightarrow A_2$ , and  $C_1 \Longrightarrow C_2$ , then  $C_2 \simeq A_2$

*Proof* By analysis of each instruction. Here we show two cases, *fetch* and *jsr*.

- Assume we have  $\langle (\text{fetch } l \pi), up, sp, kp, \sigma \rangle \simeq \langle (\text{fetch } l \pi), u, \zeta, \kappa \rangle$ . We must show that

$$\langle \pi, up, sp+1, kp, \sigma' \rangle \simeq \langle \pi, u, (u(l) :: \zeta), \kappa \rangle$$

where  $\sigma' = \sigma[(\text{lookup } l \text{ up } \sigma)/(sp+1)]$ . We consider each of the conditions for  $\simeq$  in turn:

1.  $\pi = \pi$
2. We have  $sp+1 > up+3$  and  $\sigma \models_U up \simeq u$ . Therefore  $\sigma' \models_U up \simeq u$  by Lemma 1.
3. We have  $\sigma \models_V (\text{find } l \text{ up } \sigma) \simeq u(l)$  and  $\sigma \models_Z \langle sp, kp \rangle \simeq \zeta$ . Hence  $\sigma' \models_Z \langle sp+1, kp \rangle \simeq (u(l) :: \zeta)$
4.  $sp+1 > kp$  and  $\sigma \models_K kp \simeq \kappa$ . Therefore  $\sigma' \models_K kp \simeq \kappa$

- Assume we have  $\langle (\text{jsr } \pi), up, sp, kp, \sigma \rangle \simeq \langle (\text{jsr } \pi), u, (v :: \langle \text{proc}, \text{closure } u' \pi' \rangle :: \zeta), \kappa \rangle$ . We need

$$\langle \sigma(sp-1).\pi, sp, sp+3, sp+3, (\text{push-cont up } sp \text{ kp } \sigma) \rangle \simeq \langle \pi', (\text{ext}_R(v, \langle \text{proc}, \text{closure } u' \pi' \rangle)u'), \langle \rangle, (\text{cont } \pi u \zeta \kappa) \rangle$$

where  $\zeta' = (v :: \langle \text{proc}, \text{closure } u' \pi' \rangle :: \zeta)$  and  $\sigma' = (\text{push-cont up } sp \text{ kp } \sigma)$ . By the definition of *push-cont*,  $\sigma$  and  $\sigma'$  agree on all locations less than or equal to  $sp$ .

1. We have  $\sigma \models_Z \langle sp, kp \rangle \simeq \zeta'$  and  $|\zeta'| \geq 2$ . Hence  $\sigma(sp-1).\pi = \pi'$ .
2.  $\sigma \models_Z \langle sp, kp \rangle \simeq \zeta'$  and  $|\zeta'| \geq 2$  Hence  $\sigma' \models_U sp \simeq (\text{ext}_R(v, \langle \text{proc}, \text{closure } u' \pi' \rangle)u')$ .
3.  $\sigma' \models_Z \langle sp+3, sp+3 \rangle \simeq \langle \rangle$
4.  $\sigma'(sp+3) = \pi$   
 $\sigma'(sp+2) = up+3 < sp+1$  and  $\sigma \models_U up \simeq u$   
 $\implies \sigma' \models_U \sigma'(sp+2) \simeq u$   
 $\sigma'(sp+1) = kp < sp+1$  and  $\sigma \models_K kp \simeq \kappa$   
 $\implies \sigma' \models_K \sigma'(sp+1) \simeq \kappa$   
 $\sigma \models_Z \langle sp, kp \rangle \simeq \zeta'$  and  $|\zeta'| \geq 2 \implies \sigma' \models_Z \langle sp-2, \sigma'(sp+1) \rangle \simeq \zeta$   
 $\implies \sigma' \models_K sp+3 \simeq (\text{cont } \pi u \zeta \kappa)$


$\square$

This theorem shows that the concrete machine and the abstract machine compute in the same way. started in corresponding states, they compute by passing through a sequence of corresponding states. Thus the concrete machine will give the same answer as the abstract machine

This theorem is not surprising, of course. What is new here is a way of formalizing the representation of environments, continuations, etc., so that we can see how even very low-level structures directly represent the quantities in the language definition. Furthermore, this organization allows us to separate denotational from operational reasoning, so that each can be used to best advantage.

## 8 Code Representation

We can similarly represent the program trees  $\pi$  in a linear program store  $M$ . The details here are similar, but easier. Instead of storing a program  $\pi$  in a single word of memory, as we did in the previous machine, we represent  $\pi$  explicitly in a linear program store  $M$ . For example, the program  $(\text{fetch } l \pi')$  would be represented as a location  $p$  in  $M$ .



where  $p+2$  represented  $\pi'$ . We assume that  $M$  has a domain that is disjoint from the data store  $\sigma$  and that its range is

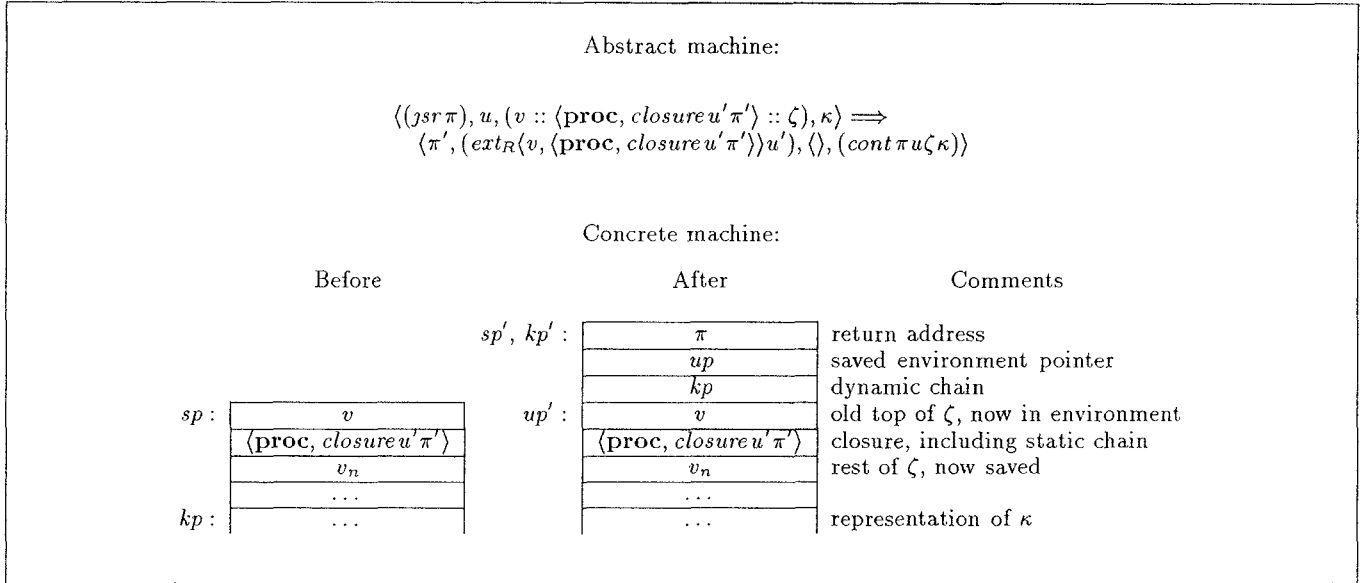


Figure 2: Executing a *jsr* instruction on the abstract and concrete machines.

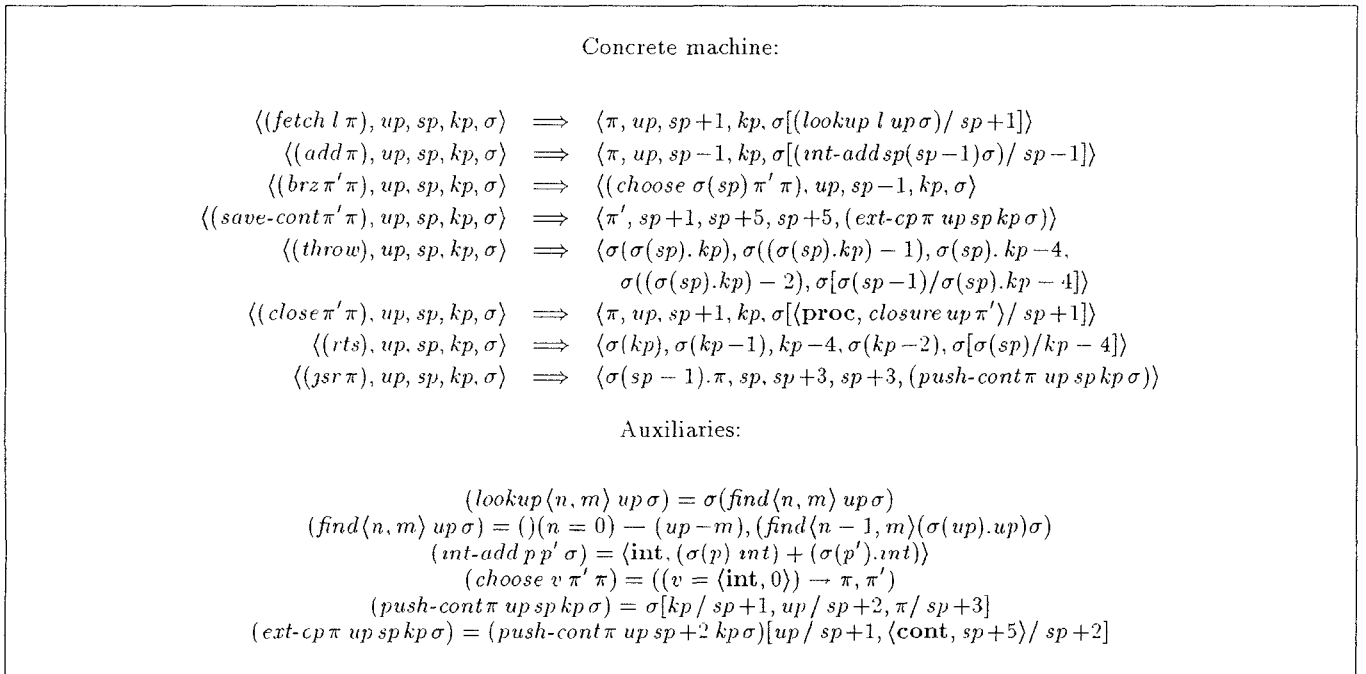


Figure 3: Actions of the concrete machine.

large enough to hold a single opcode or a pointer (either to data or program).

Again, we define a storage layout relation  $M \models_P p \simeq \pi$  as the least fixpoint of a monotonic operator. (Note the presence of *goto*, which prevents the use of structural induction. The *goto* instruction is not necessary, but illustrates the power of the technique).

- Relating pointers and programs

$$\begin{aligned}
M \models_P p \simeq \pi &\iff \\
\text{either } \pi &= (\text{fetch } l \pi') \\
&\text{and } M(p) = \text{fetch} \text{ and } M(p+1) = l \\
&\text{and } M \models_P p+2 \simeq \pi' \\
\text{or } \pi &= (\text{add } \pi') \\
&\text{and } M(p) = \text{add} \text{ and } M \models_P p+1 \simeq \pi' \\
\text{or } \pi &= (\text{brz } \pi' \pi'') \\
&\text{and } M(p) = \text{brz} \\
&\text{and } M \models_P M(p+1) \simeq \pi' \\
&\text{and } M \models_P p+2 \simeq \pi'' \\
\text{or } \pi &= (\text{save-cont } \pi' \pi'') \\
&\text{and } M(p) = \text{save-cont} \\
&\text{and } M \models_P M(p+1) \simeq \pi' \\
&\text{and } M \models_P p+2 \simeq \pi'' \\
\text{or } \pi &= \text{rts} \text{ and } M(p) = \text{rts} \\
\text{or } \pi &= \text{throw} \text{ and } M(p) = \text{throw} \\
\text{or } \pi &= (\text{close } \pi' \pi'') \\
&\text{and } M(p) = \text{close} \\
&\text{and } M \models_P M(p+1) \simeq \pi' \\
&\text{and } M \models_P p+2 \simeq \pi'' \\
\text{or } \pi &= (\text{jsr } \pi') \\
&\text{and } M(p) = \text{jsr} \text{ and } M \models_P p+1 \simeq \pi' \\
\text{or } M(p) &= \text{goto} \text{ and } M \models_P M(p+1) \simeq \pi
\end{aligned}$$

- Relating stored values

Stored values correspond if they are alike except for the representation of programs:

$$\begin{aligned}
M \models_{SV} sv \simeq sv' &\iff \\
\text{either } sv &= \langle \text{proc}, (\text{closure up } p) \rangle \\
&\text{and } sv' = \langle \text{proc}, (\text{closure up } \pi) \rangle \\
&\text{and } M \models_P p \simeq \pi \\
\text{or } sv &= p \text{ and } sv' = \pi \text{ and } M \models_P p \simeq \pi \\
\text{or } sv &= sv'
\end{aligned}$$

- Relating data stores

Data stores correspond up to pointer  $p$  if they correspond location by location for all smaller locations:

$$\begin{aligned}
M \models_p \sigma \simeq \sigma' &\iff \\
0 \leq p' \leq p, M \models_{SV} \sigma(p') &\simeq \sigma'(p')
\end{aligned}$$

We call the machine using this linear program store the *stored-program machine*. In order to describe the machine using simple pattern-matching rules, we specify the stored-program machine using a fetch-execute cycle. The controller of the machine has two states, *fetch* and *execute*. In the *fetch* state, the machine uses 4 registers: an instruction pointer  $ip$ , which is a pointer into the program store, pointers  $up$ ,  $sp$ , and  $kp$ , as in the preceding machine. In the *execute* state, the machine uses in addition an instruction register  $ir$ .

In the *fetch* state, the machine retrieves the opcode to which the instruction pointer points, places it in the instruction register, and goes to the *execute* state:

$$F(ip, up, sp, kp, \sigma) \implies E(ip, M(ip), up, sp, kp, \sigma)$$

The *execute* state then performs a rewrite, dispatching on the contents of the instruction register, and returns to the *fetch* state. The rewriting system for the stored-program machine is shown in Figure 4. This is similar to Figure 3 except for the treatment of pointers to programs.

Now we can define the correspondence between the concrete and stored-program machines. We define the correspondence at the start of the fetch-execute cycle, that is, at the *fetch* state of the stored-program machine. Since the pointers into the data store are exactly the same, all that is required is that the programs correspond, and that the stores correspond up to  $sp$ .

**Definition 2** We say a stored-program machine state  $F(ip, up, sp, kp, \sigma)$  corresponds to a concrete machine state  $\langle \pi', up', sp', kp', \sigma' \rangle$  (written  $M \models F(ip, up, sp, kp, \sigma) \simeq \langle \pi', up, sp, kp, \sigma' \rangle$ ) if and only if:

1.  $M \models_P ip \simeq \pi'$
2.  $up = up'$
3.  $sp = sp'$
4.  $kp = kp'$
5.  $M \models_{sp} \sigma \simeq \sigma'$

Now we can state the main theorem:

**Theorem 4** Let  $C_1$  and  $C_2$  be states of the concrete machine and  $L_1$  and  $L_2$  be *fetch* states of the stored-program machine. If  $M \models L_1 \simeq C_1$ ,  $C_1 \implies C_2$ , and  $L_1 \implies L_2$ , then  $M \models L_2 \simeq C_2$ .

*Proof* By analysis of each instruction. This is tedious but straightforward; since  $M$  is immutable, we do not need free-location lemmas like Lemma 1.  $\square$

If we call the algorithm that translates from program trees  $\pi$  to stored programs the “assembler,” then the correctness of the assembler factors into two parts:

1. Showing that the assembler is correct: that it emits a stored program corresponding to the program tree  $\pi$ .
2. Showing that if a stored program corresponds to a program tree  $\pi$ , then the machines will deliver corresponding outputs.

These two steps are somewhat analogous to the decomposition of the compiler: showing that the compiler is correct by showing it emits code that meets its specification (Theorem 1), and then showing that the bytecode machine simulates reduction of the bytecode term (Theorem 2)

As in the compiler case, the second part of the proof is independent of the assembler algorithm, and that is the part we have shown here; what remains is the first, assembler-dependent part. We hope to report on that elsewhere



Stored-Program machine:

$$\begin{aligned}
F(ip, up, sp, kp, \sigma) &\Longrightarrow E(ip, M(ip), up, sp, kp, \sigma) \\
E(ip, fetch, up, sp, kp, \sigma) &\Longrightarrow F(ip+2, up, sp+1, kp, \sigma[(lookup\ M(ip+1)\ up\ \sigma)/\ sp+1]) \\
E(ip, add, up, sp, kp, \sigma) &\Longrightarrow F(ip+1, up, sp-1, kp, \sigma[(int-add\ sp(sp-1)\ \sigma)/\ sp-1]) \\
E(ip, brz, up, sp, kp, \sigma) &\Longrightarrow F((choose(\sigma(sp))(ip+2)(M(ip+1))), up, sp-1, kp, \sigma) \\
E(ip, save-cont, up, sp, kp, \sigma) &\Longrightarrow F(M(ip+1), sp+1, sp+5, sp+5, (ext-cp(ip+2)\ up\ sp\ kp\ \sigma)) \\
E(ip, throw, up, sp, kp, \sigma) &\Longrightarrow F((\sigma(sp).kp), \sigma(\sigma(sp).kp-1), \sigma(sp).kp-4, \\
&\quad \sigma(\sigma(sp).kp-2), \sigma[\sigma(sp-1)/\sigma(sp).kp-4]) \\
E(ip, close, up, sp, kp, \sigma) &\Longrightarrow F(ip+2, up, sp+1, kp, \sigma[\langle proc, closure\ up\ M(ip+1) \rangle / sp+1]) \\
E(ip, rts, up, sp, kp, \sigma) &\Longrightarrow F(\sigma(kp), \sigma(kp-1), kp-4, \sigma(kp-2), \sigma[\sigma(sp)/kp-4]) \\
E(ip, jsr, up, sp, kp, \sigma) &\Longrightarrow F(\sigma(sp-1).ip, sp, sp+3, sp+3, (push-cont(ip+1)\ up\ sp\ kp\ \sigma))
\end{aligned}$$

Auxiliaries:

$$\begin{aligned}
&(lookup\ \langle n, m \rangle\ up\ \sigma) = \sigma(find\ \langle n, m \rangle\ up\ \sigma) \\
(find\ \langle n, m \rangle\ up\ \sigma) &= (n = 0) \rightarrow (up - m), (find\ \langle n - 1, m \rangle\ (\sigma(up).up\ \sigma)) \\
(int-add\ p\ p'\ \sigma) &= \langle int, (\sigma(p).int) + (\sigma(p').int) \rangle \\
(choose\ v\ ip'\ ip) &= (v = \langle int, 0 \rangle) \rightarrow ip, ip' \\
(push-cont\ ip\ up\ sp\ kp\ \sigma) &= \sigma[kp / sp+1, up / sp+2, ip / sp+3] \\
(ext-cp\ ip\ up\ sp\ kp\ \sigma) &= (push-cont\ ip\ up\ sp+2\ kp\ \sigma)[up / sp+1, \langle cont, sp+5 \rangle / sp+2]
\end{aligned}$$

Figure 4: The fetch/execute cycle of the stored-program machine.

## 9 Partial Correctness and Untagged Representations

The stack machine derivation that we have presented is only one application of these proof strategies. In this section we will sketch another such application.

Our implementation still differs from an “ordinary” implementation of a stack-based language such as C or Pascal in that data values in C or Pascal do not include tags like our `int` or `proc`. Tags like these would impose a considerable run-time penalty in both time and space.<sup>1</sup>

In languages like C or Pascal, the compiler typically uses some form of type-checking to ascertain that primitive operations like procedure call or arithmetic are invoked only on legal quantities. For languages that are not type-safe, such as C, these compile-time checks are of course less than foolproof, and may lead to unpredictable consequences (e.g. core dumps).

Here we have a semantic mismatch: the semantics is given using disjoint sums, with specified error behavior (continuations of the form  $(error \dots)$ ), but the implementation uses non-disjoint sums and has unspecified error behavior. In what sense can we say that an implementation of such a language is correct?

We can model this situation by using untagged storage layout relations and weakening the requirements on the implementation. To get a storage layout relation for an untagged implementation, we replace pattern-matching in the definition of the storage relation by field selection. Thus we

<sup>1</sup>The benefits of tag elimination are less clear for languages that require garbage collection, though even in this case it may be possible to remove all tags [1, 6]

might replace

$$\begin{aligned}
\sigma \models_V p \simeq v &\iff \\
&\text{either } v = \langle int, n \rangle = \sigma(p) \\
&\text{or } v = \langle proc, (closure\ u\ \pi_1) \rangle \\
&\quad \text{and } \sigma(p) = \langle proc, (closure\ up\ \pi_1) \rangle \\
&\quad \text{and } up < p \text{ and } \sigma \models_U up \simeq u
\end{aligned}$$

by

$$\begin{aligned}
\sigma \models_V p \simeq v &\iff \\
&\text{either } v = \langle int, u \rangle \text{ and } \sigma(p) = n \\
&\text{or } v = \langle proc, (closure\ u\ \pi_1) \rangle \\
&\quad \text{and } M \models_F \sigma(p).\pi \simeq \pi_1 \\
&\quad \text{and } \sigma(p).u < p \text{ and } \sigma \models_U \sigma(p).u \simeq u
\end{aligned}$$

We can weaken the requirement on the implementation by saying that if the abstract machine proceeds without invoking an error, then the concrete machine must do so also. If the abstract machine invokes the error continuation, then the behavior of the concrete machine is unspecified. This contract corresponds to the Scheme notion [4] of “is an error.” With this contract, the correctness theorem becomes:

**Theorem 5** *Let  $A_1$  and  $A_2$  be states of the abstract machine and  $C_1$  and  $C_2$  be states of the concrete machine. If  $C_1 \simeq A_1$ ,  $A_1 \Longrightarrow A_2$ ,  $A_2$  is not an error state, and  $C_1 \Longrightarrow C_2$ , then  $C_2 \simeq A_2$ .*

The only change from Theorem 3 is the addition of the phrase “ $A_2$  is not an error state.” The proof goes as before.

## 10 Conclusions and Further Work

We have shown the correctness of a stack implementation of a programming language. The stack was shown to be a

data structure that represented three different terms in the bytecode machine. The techniques generalize those introduced by Hannan [7]. We believe that these techniques will generalize to other crucial representation strategies, such as caching portions of the local stack  $\zeta$  in registers, representing mutable entities with dynamic extent on the stack, and modelling program loops by circular storage structures. These techniques, especially the variant in Section 9, should also be adaptable to model the finite word size of real machines.

## Acknowledgements

Paul Hudak originally suggested the correctness of a stack-discipline language as a useful exercise. Joshua Guttman, John Ramsdell, Vipin Swarup, Larry Monk, and Bill Farmer of the MITRE Corporation provided useful feedback on our Pure PreScheme compiler, which led us to the questions considered in this paper.

## References

- [1] Appel, A.W., "Runtime Tags Aren't Necessary," *Lisp and Symbolic Computation* 2 (1989), 153–162.
- [2] Clément, D., Despeyroux, J., Despeyroux, T., and Kahn, G. "A Simple Applicative Language: Mini-ML" *Proc. 1986 ACM Symp. on Lisp and Functional Programming*, 13–27.
- [3] Clinger, W. "The Scheme 311 Compiler An Exercise in Denotational Semantics," *Conf. Rec. 1984 ACM Symposium on Lisp and Functional Programming* (August, 1984), 356–364
- [4] Clinger, W., and Rees, J., eds. "Revised<sup>4</sup> Report on the Algorithmic Language Scheme", Indiana University Computer Science Department Technical Report No. 341, November, 1991, to appear in *LISP Pointers*, 1992 See also "Revised<sup>3</sup> Report on the Algorithmic Language Scheme", *SIGPLAN Notices* 21,12 (December, 1986), 37–79.
- [5] Friedman, D.P., Wand, M., and Haynes, C.T., *Essentials of Programming Languages*, MIT Press, Cambridge, MA, and McGraw-Hill, Chicago, 1992.
- [6] Goldberg, B. "Tag-Free Garbage Collection for Strongly Typed Programming Languages," *Proc. ACM SIGPLAN '91 Symp. on Programming Language Design and Implementation*, June, 1991, 165–176.
- [7] Hannan, J. "Making Abstract Machines Less Abstract," *Functional Programming Languages and Computer Architecture, 5th ACM Conference* (J. Hughes, ed.) Springer Lecture Notes in Computer Science, Vol. 523 (1991), 618–635.
- [8] Lee, P. *Topics in Advanced Language Implementation*, MIT Press, Cambridge, MA, 1991
- [9] Farmer, W.M., Guttman, J.D., Monk, L.G., Ramsdell, J.D., and Swarup, V. "VLISP. A Verified Language Implementation FY 91 Year-End Report," The MITRE Corporation Technical Report MTR11232, October, 1991.
- [10] Oliva, D.P., and Wand, M., "A Verified Compiler for Pure PreScheme," Final Report for MITRE Contract Number F19628-89-C-001. Included in [9].
- [11] McCarthy, J. and Painter, J. "Correctness of a Compiler for Arithmetic Expressions," in *Proc. Symp. in Appl. Math., Vol. 19, Mathematical Aspects of Computer Science* (J. T. Schwartz, ed.) Amer. Math. Soc., Providence, RI, 1967, 33–41.
- [12] Milne, R. and Strachey C. *A Theory of Programming Language Semantics*, Chapman & Hall, London, and Wiley, New York, 1976
- [13] Plotkin, G.D. "Call-by-Name, Call-by-Value and the  $\lambda$ -Calculus," *Theoret. Comp. Sci.* 1 (1975) 125–159.
- [14] Raoult, J.-C. and Sethi, R. "The Global Storage Needs of a Subcomputation," *Conf. Rec. 11th ACM Symp. on Principles of Programming Languages* (1984), 148–157.
- [15] Reynolds, J.C. *The Craft of Programming*, Prentice-Hall International, 1981.
- [16] Schmidt, D.A. *Denotational Semantics: A Methodology for Language Development* Boston: Allyn and Bacon, 1986.
- [17] Stoy, J.E. *Denotational Semantics. The Scott-Strachey Approach to Programming Language Theory*, MIT Press, Cambridge, MA., 1977.
- [18] Wand, M. "Semantics-Directed Machine Architecture" *Conf. Rec. 9th ACM Symp. on Principles of Prog. Lang.* (1982), 234–241
- [19] Wand, M. "Loops in Combinator-Based Compilers," *Info. and Control* 57, 2–3 (May/June, 1983), 148–164.
- [20] Wand, M. "On the Correctness of Procedure Representations in Higher-Order Assembly Language" *Proc. MFPS '91* (S. Brookes, ed.) Springer Lecture Notes in Computer Science, to appear