Stack-based Typed Assembly Language

Citation

Morrisett, Greg, Karl Crary, Neal Glew, and David Walker. 2003. Stack-based typed assembly language. *Journal of Functional Programming* 13: 957-959.

Published Version

<http://dx.doi.org/10.1017/S0956796802004446>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3710253>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Stack-Based Typed Assembly Language *

Greg Morrisett Karl Crary Neal Glew David Walker

Cornell University

Abstract

In previous work, we presented a *Typed Assembly Language* (TAL). TAL is sufficiently expressive to serve as a target language for compilers of high-level languages such as ML. Moreover, TAL’s type system is powerful enough to encode source-level abstractions such as closures and abstract data types, yet flexible enough to admit most traditional low-level optimizations.

The compiler presented in our original work was based on a *continuation-passing style* (CPS) transform, which eliminated the need for a control stack by heap allocating activation records. However, modern architectures and most compilers are based on stack allocation of these records. This paper presents STAL, an extension of TAL with stack constructs and stack types to support the stack allocation style. We show that STAL is sufficiently expressive to support languages such as Java, Pascal, and ML; constructs such as exceptions and displays; and optimizations such as tail call elimination and callee-saves registers.

This paper makes two additional contributions. First, it formalizes the typing connection between CPS-based compilation and stack-based compilation. The latter is seen as a continuation-passing style where continuation closures are unboxed by the caller and the continuation’s environment is placed on the stack. Second, it illustrates how STAL can formally model calling conventions by specifying them as formal translations of source function types to STAL types.

1 Introduction and Motivation

Statically typed source languages have efficiency and software engineering advantages over their dynamically typed counterparts. Modern type-directed compilers [19, 25, 7, 32, 20, 28, 12] exploit the properties of typed languages more extensively than their predecessors by preserving type information computed in the front end through a series of typed intermediate languages. These compilers use types to direct sophisticated transformations such as closure conversion [18, 31, 17, 4, 21], region inference [8], subsumption elimination [9, 11], and unboxing [19, 24, 29]. In many cases, without types, these transformations are less effective or simply impossible. Furthermore, the type translation partially specifies the corresponding term translation and often captures the critical concerns in an elegant and succinct fashion. Strong type systems not only describe but also enforce many important invariants. Consequently, developers of type-based compilers may invoke a type-checker after each code transformation, and if the output fails to type-check, the developer

*This material is based on work supported in part by the AFOSR grant F49620-97-1-0013, ARPA/RADC grant F30602-96-1-0317, ARPA/AF grant F30602-95-1-0047, and AASERT grant N00014-95-1-0985. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not reflect the views of these agencies.

knows that the compiler contains an internal error. Although type-checkers for decidable type systems will not catch all compiler errors, they have proven themselves valuable debugging tools in practice [22].

Despite the numerous advantages of compiling with types, until recently, no compiler propagated type information through the final stages of code generation. The TIL/ML compiler, for instance, preserves types through approximately 80% of compilation but leaves the remaining 20% untyped. Many of the complex tasks of code generation including register allocation and instruction scheduling are left unchecked and types cannot be used to specify or explain these low-level code transformations.

These observations motivated our exploration of very low-level type systems and corresponding compiler technology. In Morrisett et al. [23], we presented a *typed assembly language* (TAL) and proved that its type system was sound with respect to an operational semantics. We demonstrated the expressiveness of this type system by sketching a type-preserving compiler from an ML-like language to TAL. The compiler ensured that well-typed source programs were always mapped to well-typed assembly language programs and that they preserved source level abstractions such as user-defined abstract data types and closures. Furthermore, we claimed that the type system of TAL did not interfere with many traditional compiler optimizations including inlining, loop-unrolling, register allocation, instruction selection, and instruction scheduling.

However, the compiler we presented was critically based on a *continuation-passing style* transform, which eliminated the need for a control stack. In particular, activation records were represented by heap-allocated closures as in the SML of New Jersey compiler (SML/NJ) [5, 3]. For example, Figure 1 shows the TAL code our heap-based compiler would produce for the recursive factorial computation. Each function takes an additional argument which represents the control stack as a continuation closure. Instead of “returning” to the caller, a function invokes its continuation closure by jumping directly to the code of the closure, passing the environment of the closure and the result in registers.

Allocating continuation closures on the heap has many advantages over a conventional stack-based implementation. First, it is straightforward to implement control primitives such as exceptions, first-class continuations, or user-level lightweight coroutine threads when continuations are heap allocated [3, 31, 34]. Second, Appel and Shao [2] have shown that heap allocation of closures can have better space properties, primarily because it is easier to share environments. Third, there is a unified memory management mechanism (namely the garbage collector) for allocating and collecting all kinds of objects, including stack frames. Finally, Appel and Shao [2] have argued that, at least for SML/NJ, the locality lost by heap-allocating stack frames is negligible.

Nevertheless, there are also compelling reasons for providing support for stacks. First, Appel and Shao’s work did not consider imperative languages, such as Java, where the ability to share environments is greatly reduced nor did it consider languages that do not require garbage collection. Second, Tarditi and Diwan [14, 13] have shown that with some cache architectures, heap allocation of continuations (as in SML/NJ) can have substantial overhead due to a loss of locality. Third, stack-based activation records can have a smaller memory footprint than heap-based activation records. Finally, many machine architectures have hardware mechanisms that expect programs to behave in a stack-like fashion. For example, the Pentium Pro processor has an internal stack that it uses to predict return addresses for procedures so that instruction pre-fetching will not be stalled [16]. The internal stack is guided by the use of call/return primitives which use the standard

control stack.

Clearly, compiler writers must weigh a complex set of factors before choosing stack allocation, heap allocation, or both. The target language must not constrain these design decisions. In this paper, we explore the addition of a stack to our typed assembly language in order to give compiler writers the flexibility they need. Our stack typing discipline is remarkably simple, but powerful enough to compile languages such as Pascal, Java, or ML without adding high-level primitives to the assembly language. More specifically, the typing discipline supports stack allocation of temporary variables and values that do not escape, stack allocation of procedure activation frames, exception handlers, and displays, as well as optimizations such as callee-saves registers. Unlike the JVM architecture [20], our system does not constrain the stack to have the same size at each control-flow point, nor does it require new high-level primitives for procedure call/return. Instead, our assembly language continues to have low-level RISC-like primitives such as loads, stores, and jumps. However, source-level stack allocation, general source-level stack pointers, general pointers into either the stack or heap, and some advanced optimizations cannot be typed.

A key contribution of the type structure is that it provides a unifying declarative framework for specifying procedure calling conventions regardless of the allocation strategy. In addition, the framework further elucidates the connection between a heap-based continuation-passing style compiler, and a conventional stack-based compiler. In particular, this type structure makes explicit the notion that the only differences between the two styles are that, instead of passing the continuation as a boxed, heap-allocated tuple, a stack-based compiler passes the continuation unboxed in registers and the environments for continuations are allocated on the stack. The general framework makes it easy to transfer transformations developed for one style to the other. For instance, we can easily explain the callee-saves registers of SML/NJ [5, 3, 1] and the callee-saves registers of a stack-based compiler as instances of a more general CPS transformation that is independent of the continuation representation.

2 Overview of TAL and CPS-Based Compilation

In this section, we briefly review our original proposal for typed assembly language (TAL) and sketch how a polymorphic functional language, such as ML, can be compiled to TAL in a continuation-passing style, where continuations are heap-allocated.

Figure 2 gives the syntax for TAL. Programs (P) are triples consisting of a heap, register file, and instruction sequence. Heaps map labels to heap values which are either tuples of word-sized values or code sequences. Register files map registers to word-sized values. Instruction sequences are lists of instructions terminated by either a `jmp` or `halt` instruction. The context Δ binds the free type variables of Γ in $\forall[\Delta].\Gamma$, and of both Γ and I in `code` $[\Delta]\Gamma.I$. The instruction `unpack` $[\alpha, r], v$ binds α in the following instructions. We consider syntactic objects to be equivalent up to alpha-conversion, and consider label assignments, register assignments, heaps, and register files equivalent up to reordering of labels and registers. Register names do not alpha-convert. The notation $\bar{X}$ denotes a sequence of zero or more X s, and $|\cdot|$ denotes the length of a sequence.

The instruction set consists mostly of conventional RISC-style assembly operations, including arithmetic, branches, loads, and stores. One exception, the `unpack` instruction, strips the quantifier from the type of an existentially typed value and introduces a new type variable into scope. On an un-

$(H, \{\}, I)$ where

$H =$

l_fact:

$\text{code}[\langle \rangle, \text{r2:} \text{int}, \text{r3:} \tau_k \rangle.$

$\text{bneq r2, l_nonzero}$

$\text{unpack } [\alpha, \text{r3}], \text{r3}$

ld r4, r3(0)

ld r1, r3(1)

mov r2, 1

jmp r4

% zero branch: call k (in r3) with 1

% project k code

% project k environment

% jump to k with $\{\text{r1} = \text{env}, \text{r2} = 1\}$

l_nonzero:

$\text{code}[\langle \rangle, \text{r2:} \text{int}, \text{r3:} \tau_k \rangle.$

sub r4, r2, 1

$\text{malloc r5}[\text{int}, \tau_k]$

st r5(0), r2

st r5(1), r3

$\text{malloc r3}[\forall[\langle \rangle, \tau_k^1], \text{r2:} \text{int}, \langle \text{int}^1, \tau_k^1 \rangle]$

mov r2, l_cont

st r3(0), r2

st r3(1), r5

mov r2, r4

$\text{mov r3, pack}[\langle \text{int}^1, \tau_k^1 \rangle, \text{r3}] \text{ as } \tau_k$

jmp l_fact

% $n - 1$

% create environment for cont in r5

% store n into environment

% store k into environment

% create cont closure in r3

% store cont code

% store environment $\langle n, k \rangle$

% arg := $n - 1$

% abstract the type of the environment

% recursive call with $\{\text{r1} = \text{env}, \text{r2} = n - 1, \text{r3} = \text{cont}\}$

l_cont:

$\text{code}[\langle \text{int}^1, \tau_k^1 \rangle, \text{r2:} \text{int} \rangle.$

ld r3, r1(0)

ld r4, r1(1)

mul r2, r3, r2

$\text{unpack } [\alpha, \text{r4}], \text{r4}$

ld r3, r4(0)

ld r1, r4(1)

jmp r3

% r2 contains $(n - 1)!$

% retrieve n

% retrieve k

% $n \times (n - 1)!$

% unpack k

% project k code

% project k environment

% jump to k with $\{\text{r1} = \text{env}, \text{r2} = n!\}$

l_halt:

$\text{code}[\langle \rangle, \text{r2:} \text{int} \rangle.$

mov r1, r2

$\text{halt}[\text{int}]$

% halt with result in r1

and $I =$

$\text{malloc r1}[]$

$\text{malloc r2}[]$

$\text{malloc r3}[\forall[\langle \rangle, \text{r2:} \text{int}], \langle \rangle]$

mov r4, l_halt

st r3(0), r4

st r3(1), r2

mov r2, 6

$\text{mov r3, pack}[\langle \rangle, \text{r3}] \text{ as } \tau_k$

jmp l_fact

% create an empty environment $\langle \rangle$

% create another empty environment

% create halt closure in r3

% store cont code

% store environment $\langle \rangle$

% load argument (6)

% abstract the type of the environment

% begin fact with $\{\text{r1} = \langle \rangle, \text{r2} = 6, \text{r3} = \text{haltcont}\}$

and $\tau_k = \exists \alpha. \forall [\langle \rangle, \text{r1:} \alpha, \text{r2:} \text{int} \rangle^1, \alpha^1 \rangle$

Figure 1: Typed Assembly Code for Factorial

<i>types</i>	$\tau ::= \alpha \mid \text{int} \mid \forall[\Delta].\Gamma \mid \langle \tau_1^{\varphi_1}, \dots, \tau_n^{\varphi_n} \rangle \mid \exists\alpha.\tau$
<i>initialization flags</i>	$\varphi ::= 0 \mid 1$
<i>label assignments</i>	$\Psi ::= \{\ell_1:\tau_1, \dots, \ell_n:\tau_n\}$
<i>type assignments</i>	$\Delta ::= \cdot \mid \alpha, \Delta$
<i>register assignments</i>	$\Gamma ::= \{r_1:\tau_1, \dots, r_n:\tau_n\}$
<i>registers</i>	$r ::= \mathbf{r1} \mid \dots \mid \mathbf{rk}$
<i>word values</i>	$w ::= \ell \mid i \mid ?\tau \mid w[\tau] \mid \text{pack } [\tau, w] \text{ as } \tau'$
<i>small values</i>	$v ::= r \mid w \mid v[\tau] \mid \text{pack } [\tau, v] \text{ as } \tau'$
<i>heap values</i>	$h ::= \langle w_1, \dots, w_n \rangle \mid \text{code}[\Delta]\Gamma.I$
<i>heaps</i>	$H ::= \{\ell_1 \mapsto h_1, \dots, \ell_n \mapsto h_n\}$
<i>register files</i>	$R ::= \{r_1 \mapsto w_1, \dots, r_n \mapsto w_n\}$
<i>instructions</i>	$\iota ::= \text{aop } r_d, r_s, v \mid \text{bop } r, v \mid \text{ld } r_d, r_s(i) \mid \text{malloc } r[\bar{\tau}] \mid$ $\text{mov } r_d, v \mid \text{st } r_d(i), r_s \mid \text{unpack } [\alpha, r_d], v \mid$
<i>arithmetic ops</i>	$\text{aop} ::= \text{add} \mid \text{sub} \mid \text{mul}$
<i>branch ops</i>	$\text{bop} ::= \text{beq} \mid \text{bneq} \mid \text{bgt} \mid \text{blt} \mid \text{bgte} \mid \text{blte}$
<i>instruction sequences</i>	$I ::= \iota; I \mid \text{jmp } v \mid \text{halt } [\tau]$
<i>programs</i>	$P ::= (H, R, I)$

Figure 2: Syntax of TAL

typed machine, this is implemented by an ordinary move. The other non-standard instruction is `malloc`, which is explained below. Evaluation is specified as a deterministic rewriting system that takes programs to programs (see Morrisett et al. [23] for details).

The types for TAL consist of type variables, integers, tuple types, existential types, and polymorphic code types. Tuple types contain initialization flags (either 0 or 1) that indicate whether or not components have been initialized. For example, if register r has type $\langle \text{int}^0, \text{int}^1 \rangle$, then it contains a label bound in the heap to a pair that can contain integers, where the first component may not have been initialized, but the second component has. In this context, the type system allows the second component to be loaded, but not the first. If an integer value is stored into $r(0)$ then afterwards r has the type $\langle \text{int}^1, \text{int}^1 \rangle$, reflecting the fact that the first component is now initialized. The instruction `malloc` $r[\tau_1, \dots, \tau_n]$ heap-allocates a new tuple with uninitialized fields and places its label in register r .

Code types $(\forall[\alpha_1, \dots, \alpha_n].\Gamma)$ describe code blocks $(\text{code}[\alpha_1, \dots, \alpha_n]\Gamma.I)$, which are instruction sequences, that expect a register file of type Γ and in which the type variables $\alpha_1, \dots, \alpha_n$ are held abstract. In other words, Γ serves as a register file pre-condition that must hold before control may be transferred to the code block. Code blocks have no post-condition because control is either terminated via a `halt` instruction or transferred to another code block.

The type variables that are abstracted in a code block provide a means to write polymorphic code

sequences. For example, the polymorphic code block

```
code[α]{r1:α, r2:∀[].{r1:⟨α1, α1⟩}}.
  malloc  r3[α, α]
  st      r3(0), r1
  st      r3(1), r1
  mov     r1, r3
  jmp     r2
```

roughly corresponds to a CPS version of the SML function $\text{fn } (x:\alpha) \Rightarrow (x, x)$. The block expects upon entry that register **r1** contains a value of the abstract type α , and **r2** contains a return address (or continuation label) of type $\forall[].\{r1 : \langle \alpha^1, \alpha^1 \rangle\}$. In other words, the return address requires register **r1** to contain an initialized pair of values of type α before control can be returned to this address. The instructions of the code block allocate a tuple, store into the tuple two copies of the value in **r1**, move the pointer to the tuple into **r1** and then jump to the return address in order to “return” the tuple to the caller. If the code block is bound to a label ℓ , then it may be invoked by simultaneously instantiating the type variable and jumping to the label (*e.g.*, `jmp  $\ell[int]$`).

Source languages like ML have nested higher-order functions that might contain free variables and thus require *closures* to represent functions. At the TAL level, we represent closures as a pair consisting of a code block label and a pointer to an environment data structure. The type of the environment must be held abstract in order to avoid typing difficulties [21], and thus we *pack* the type of the environment and the pair to form an existential type.

All functions, including continuation functions introduced during CPS conversion, are thus represented as existentials. For example, once CPS converted, a source function of type $\text{int} \rightarrow \langle \rangle$ has type $(\text{int}, (\langle \rangle \rightarrow \text{void})) \rightarrow \text{void}$.¹ After closures are introduced, the code will have type:

$$\exists \alpha_1. \langle (\alpha_1, \text{int}, \exists \alpha_2. \langle (\alpha_2, \langle \rangle) \rightarrow \text{void}, \alpha_2 \rangle) \rightarrow \text{void}, \alpha_1 \rangle$$

Finally, at the TAL level the function will be represented by a value with the type:

$$\exists \alpha_1. \langle \forall []. \{r1:\alpha_1, r2:\text{int}, r3:\exists \alpha_2. \langle \forall []. \{r1:\alpha_2, r2:\langle \rangle\}^1, \alpha_2^1\}^1, \alpha_1^1 \rangle \rangle$$

Here, α_1 is the abstracted type of the closure’s environment. The code for the closure requires that the environment be passed in register **r1**, the integer argument in **r2**, and the continuation in **r3**. The continuation is itself a closure where α_2 is the abstracted type of its environment. The code for the continuation closure requires that the environment be passed in **r1** and the unit result of the computation in **r2**.

To apply a closure at the TAL level, we first use the **unpack** operation to open the existential package. Then the code and the environment of the closure pair are loaded into appropriate registers, along with the argument to the function. Finally, we use a jump instruction to transfer control to the closure’s code.

Figure 1 gives the CPS-based TAL code for the following ML expression which computes six factorial:

```
let fun fact n = if n = 0 then 1 else n * (fact(n - 1)) in fact 6 end
```

¹The *void* return types are intended to suggest the non-returning aspect of CPS code.

<i>types</i>	$\tau ::= \dots \mid ns$
<i>stack types</i>	$\sigma ::= \rho \mid nil \mid \tau::\sigma$
<i>type assignments</i>	$\Delta ::= \dots \mid \rho, \Delta$
<i>register assignments</i>	$\Gamma ::= \{r_1:\tau_1, \dots, r_n:\tau_n, \mathbf{sp}:\sigma\}$
<i>word values</i>	$w ::= \dots \mid w[\sigma] \mid ns$
<i>small values</i>	$v ::= \dots \mid v[\sigma]$
<i>register files</i>	$R ::= \{r_1 \mapsto w_1, \dots, r_n \mapsto w_n, \mathbf{sp} \mapsto S\}$
<i>stacks</i>	$S ::= nil \mid w::S$
<i>instructions</i>	$\iota ::= \dots \mid \mathbf{salloc} \ n \mid \mathbf{sfree} \ n \mid \mathbf{sld} \ r_d, \mathbf{sp}(i) \mid \mathbf{sst} \ \mathbf{sp}(i), r_s$

Figure 3: Additions to TAL for Simple Stacks

3 Adding Stacks to TAL

In this section, we show how to extend TAL to achieve a Stack-based Typed Assembly Language (STAL). Figure 3 defines the new syntactic constructs for the language. In what follows, we informally discuss the dynamic and static semantics for the modified language, leaving formal treatment to Appendix A.

Operationally, we model stacks (S) as lists of word-sized values. Uninitialized stack slots are filled with nonsense (ns). Register files now include a distinguished register, $\mathbf{sp}$, which represents the current stack. There are four new instructions that manipulate the stack. The $\mathbf{salloc} \ n$ instruction places n words of nonsense on the top of the stack. In a conventional machine, assuming stacks grow towards lower addresses, an $\mathbf{salloc}$ instruction would correspond to subtracting n from the current value of the stack pointer. The $\mathbf{sfree} \ n$ instruction removes the top n words from the stack, and corresponds to adding n to the current stack pointer. The $\mathbf{sld} \ r, \mathbf{sp}(i)$ instruction loads the i^{th} word of the stack into register r , whereas the $\mathbf{sst} \ \mathbf{sp}(i), r$ stores register r into the i^{th} word. Note, the instructions $\mathbf{ld}$ and $\mathbf{st}$ cannot be used with the stack pointer.

A program becomes *stuck* if it attempts to execute:

- $\mathbf{sfree} \ n$ and the stack does not contain at least n words,
- $\mathbf{sld} \ r, \mathbf{sp}(i)$ and the stack does not contain at least $i + 1$ words or else the i^{th} word of the stack is ns , or
- $\mathbf{sst} \ \mathbf{sp}(i), r$ and the stack does not contain at least $i + 1$ words.

As in the original TAL, the typing rules for the modified language prevent well-formed programs from becoming stuck.

Stacks are described by *stack types* (σ), which include nil and $\tau::\sigma$. The latter represents a stack of the form $w::S$ where w has type τ and S has type σ . Stack slots filled with nonsense have type ns . Stack types also include stack type variables (ρ) which may be used to abstract the tail of a stack type. The ability to abstract stacks is critical for supporting procedure calls and is discussed in detail later.

As before, the register file for the abstract machine is described by a register file type (Γ) mapping registers to types. However, Γ also maps the distinguished register **sp** to a stack type σ . Finally, code blocks and code types support polymorphic abstraction over both types and stack types.

One of the uses of the stack is to save temporary values during a computation. The general problem is to save on the stack n registers, say r_1 through r_n , of types τ_1 through τ_n , perform some computation e , and then restore the temporary values to their respective registers. This would be accomplished by the following instruction sequence where the comments (delimited by %) show the stack's type at each step of the computation.

		% σ
salloc	n	% $ns::ns::\dots::ns::\sigma$
sst	sp (0), r_1	% $\tau_1::ns::\dots::ns::\sigma$
:		
sst	sp ($n-1$), r_n	% $\tau_1::\tau_2::\dots::\tau_n::\sigma$
code for e		% $\tau_1::\tau_2::\dots::\tau_n::\sigma$
sld	r_1 , sp (0)	% $\tau_1::\tau_2::\dots::\tau_n::\sigma$
:		
sld	r_n , sp ($n-1$)	% $\tau_1::\tau_2::\dots::\tau_n::\sigma$
sfree	n	% σ

If, upon entry, r_i has type τ_i and the stack is described by σ , and if the code for e leaves the state of the stack unchanged, then this code sequence is well-typed. Furthermore, the typing discipline does not place constraints on the order in which the stores or loads are performed.

It is straightforward to model higher-level primitives, such as **push** and **pop**. The former can be seen as simply **salloc** 1 followed by a store to **sp**(0), whereas the latter is a load from **sp**(0) followed by **sfree** 1. Also, a “jump-and-link” or “call” instruction which automatically moves the return address into a register or onto the stack can be synthesized from our primitives. To simplify the presentation, we did not include these instructions in STAL; a practical implementation, however, would need a full set of instructions appropriate to the architecture.

The stack is commonly used to save the current return address, and temporary values across procedure calls. Which registers to save and in what order is usually specified by a compiler-specific calling convention. Here we consider a simple calling convention where it is assumed there is one integer argument and one unit result, both of which are passed in register **r1**, and the return address is passed in the register **ra**. When invoked, a procedure may choose to place temporaries on the stack as shown above, but when it jumps to the return address, the stack should be in the same state as it was upon entry. Naively, we might expect the code for a function obeying this calling convention to have the following STAL type:

$$\forall[].\{\mathbf{r1}:int, \mathbf{sp}:\sigma, \mathbf{ra}:\forall[].\{\mathbf{r1}:\langle \rangle, \mathbf{sp}:\sigma\}\}$$

Notice that the type of the return address is constrained so that the stack must have the same shape upon return as it had upon entry. Hence, if the procedure pushes any arguments onto the stack, it must pop them off.

However, this typing is unsatisfactory for two reasons. The first problem is that there is nothing preventing the procedure from popping off values from the stack and then pushing new values (of

the appropriate type) onto the stack. In other words, the caller's stack frame is not protected from the function's code. The second problem is much worse: such a function can only be invoked from states where the stack is exactly described by σ . This effectively prevents invocation of the procedure from two different points in the program. In particular, there is no way for the procedure to push its return address on the stack and jump to itself.

The solution to both problems is to abstract the type of the stack using a stack type variable:

$$\forall[\rho].\{\mathbf{r1}:int, \mathbf{sp}:\rho, \mathbf{ra}:\forall[].\{\mathbf{r1}:int, \mathbf{sp}:\rho\}\}$$

To invoke a function with this type, the caller must instantiate the bound stack type variable ρ with the current type of the stack. As before, the function can only jump to the return address when the stack is in the same state as it was upon entry. However, the first problem above is addressed because the type checker treats ρ as an abstract stack type while checking the body of the code. Hence, the code cannot perform an `sfree`, `sld`, or `sst` on the stack. It must first allocate its own space on the stack, only this space may be accessed by the function, and the space must be freed before returning to the caller.² The second problem is solved because the stack type variable may be instantiated in different ways. Hence multiple call sites with different stack states, including recursive calls, may now invoke the function. In fact, a recursive call will usually instantiate the stack variable with a different type than the original call because unless it is a tail call, it will need to store its return address on the stack.

Figure 4 gives stack-based code for the factorial example of the previous section. The function is invoked by moving its environment (an empty tuple) into `r1`, the argument into `r2`, and the return address label into `ra` and jumping to the label `l_fact`. Notice that the nonzero branch must save the argument and current return address on the stack before jumping to the `fact` label in a recursive call. It is interesting to note that the stack-based code is quite similar to the heap-based code of Figure 1. Indeed, the code remains in a continuation-passing style, but instead of passing the continuation as a heap-allocated tuple, the environment of the continuation is passed in the stack pointer and the code of the continuation is passed in the return address register.

To more fully appreciate the correspondence, consider the type of the TAL version of `l_fact` from Figure 1:

$$\forall[].\{\mathbf{r1}:\langle\rangle, \mathbf{r2}:int, \mathbf{r3}:\exists\alpha.\langle\forall[].\{\mathbf{r1}:\alpha, \mathbf{r2}:int\}^1, \alpha^1\rangle\}$$

We could have used an alternative approach where we pass the components of the continuation closure unboxed in separate registers. To do so, the caller must unpack the continuation and the function must abstract the type of the continuation's environment resulting in a quantifier rotation:

$$\forall[\alpha].\{\mathbf{r1}:\langle\rangle, \mathbf{r2}:int, \mathbf{r3}:\forall[].\{\mathbf{r1}:\alpha, \mathbf{r2}:int\}, \mathbf{r4}:\alpha\}$$

Now, it is clear that the STAL code, which has type

$$\forall[\rho].\{\mathbf{r1}:\langle\rangle, \mathbf{r2}:int, \mathbf{ra}:\forall[].\{\mathbf{sp}:\rho, \mathbf{r1}:int\}, \mathbf{sp}:\rho\}$$

is essentially the same! Indeed, the only difference between a CPS-based compiler, such as SML/NJ, and a conventional stack-based compiler, is that for the latter, continuation environments are allocated on a stack. Our type system describes this well-known connection elegantly.

²Some intuition on this topic may be obtained from Reynolds' theorem on parametric polymorphism [27] but a formal proof is difficult.

$(H, \{\text{sp} \mapsto \text{nil}\}, I)$ where

$H =$

```

l_fact:    code[ρ]{r1:⟨⟩, r2:int, sp:ρ, ra:τρ}.    % main entry point for factorial
           bneq  r2, l_nonzero[ρ]                  % if n = 0 continue
           mov   r1, 1                               % result is 1
           jmp   ra                                   % return
l_nonzero: code[ρ]{r1:⟨⟩, r2:int, sp:ρ, ra:τρ}.    % code for n ≠ 0
           sub   r3, r2, 1                           % n - 1
           salloc 2                                  % save n and return address to stack
           sst   sp(0), r2
           sst   sp(1), ra
           mov   r2, r3                               % recursive call fact(n - 1)
           mov   ra, l_cont[ρ]
           jmp   l_fact[int::τρ::ρ]
l_cont:    code[ρ]{r1:int, sp:int::τρ::ρ}.          % return point of recursive call
           sld   r2, sp(0)                            % restore n and return address
           sld   ra, sp(1)
           sfree 2
           mul   r1, r2, r1                            % result is n × fact(n - 1)
           jmp   ra                                   % return
l_halt:    code[] {r1:int, sp:nil}.
           halt [int]

I =
           malloc r1[]                                % environment
           mov   r2, 6                                % argument
           mov   ra, l_halt                            % return address for initial call
           jmp   l_fact[nil]

```

and $\tau_\rho = \forall[], \{r1 : \text{int}, \text{sp} : \rho\}$

Figure 4: STAL Factorial Example

Our techniques can be applied to other calling conventions and do not appear to inhibit most optimizations. For instance, tail calls can be eliminated in CPS simply by forwarding a continuation closure to the next function. If continuations are allocated on the stack, we have the mechanisms to pop the current activation frame off the stack and to push any arguments before performing the tail call. Furthermore, the type system is expressive enough to type this resetting and adjusting for any kind of tail call, not just a self tail call. As another example, some CISC-style conventions place the environment, the argument(s), and return address on the stack, and return the result on the stack. With this convention, the factorial code would have type:

$$\forall[\rho].\{\mathbf{sp}:\langle\rangle::int::\forall[]\{\mathbf{sp}:int::\rho\}::\rho\}$$

Callee-saves registers (registers whose values must be preserved across function calls) can be handled in the same fashion as the stack pointer. In particular, the function holds abstract the type of the callee-saves register and requires that the register have the same type upon return. For instance, if we wish to preserve register $\mathbf{r3}$ across a call to factorial, we would use the type:

$$\forall[\rho, \alpha].\{\mathbf{r1}:\langle\rangle, \mathbf{r2}:int, \mathbf{r3}:\alpha, \mathbf{ra}:\forall[]\{\mathbf{sp}:\rho, \mathbf{r1}:int, \mathbf{r3}:\alpha\}, \mathbf{sp}:\rho\}$$

Translating this type back in to a boxed, heap allocated closure, we obtain:

$$\forall[\alpha].\{\mathbf{r1}:\langle\rangle, \mathbf{r2}:int, \mathbf{r3}:\alpha, \mathbf{ra}:\exists\beta.\langle\forall[]\{\mathbf{r1}:\beta, \mathbf{r2}:int, \mathbf{r3}:\alpha\}^1, \beta^1\rangle\}$$

This is the type of the callee-saves approach of Appel and Shao [1]. Thus we see how our correspondence enables transformations developed for heap-based compilers to be used in traditional stack-based compilers and vice versa. The generalization to multiple callee-saves registers and other calling conventions should be clear. Indeed, we have found that the type system of STAL provides a concise way to declaratively specify a variety of calling conventions.

4 Exceptions

We now consider how to implement exceptions in STAL. We will find that a calling convention for function calls in the presence of exceptions may be derived from the heap-based CPS calling convention, just as was the case without exceptions. However, implementing this calling convention will require that the type system be made more expressive by adding *compound* stack types. This additional expressiveness will turn out to have uses beyond exceptions, allowing most compiler-introduced uses of pointers into the midst of the stack.

4.1 Exception Calling Conventions

In a heap-based CPS framework, exceptions are implemented by passing two continuations: the usual continuation and an *exception continuation*. Code raises an exception by jumping to the latter. For an integer to unit function, this calling convention is expressed as the following TAL type (ignoring the outer closure and environment):

$$\forall[].\{\mathbf{r1}:int, \mathbf{ra}:\exists\alpha_1.\langle\forall[].\{\mathbf{r1}:\langle\rangle, \mathbf{r2}:\alpha_1\}^1, \alpha_1^1\rangle, \mathbf{re}:\exists\alpha_2.\langle\forall[].\{\mathbf{r1}:exn, \mathbf{r2}:\alpha_2\}^1, \alpha_2^1\rangle\}$$

Again, the caller might unpack the continuations:

$$\forall[\alpha_1, \alpha_2].\{\mathbf{r1}:int, \mathbf{ra}:\forall[].\{\mathbf{r1}:\alpha_1, \mathbf{r2}:\langle\rangle\}, \mathbf{ra}':\alpha_1, \mathbf{re}:\forall[].\{\mathbf{r1}:\alpha_2, \mathbf{r2}:exn\}, \mathbf{re}':\alpha_2\}$$

Then the caller might attempt to place the continuation environments on stacks, as before:

$$\forall[\rho_1, \rho_2].\{\mathbf{r1}:int, \mathbf{ra}:\forall[].\{\mathbf{sp}:\rho_1, \mathbf{r1}:\langle\rangle\}, \mathbf{sp}:\rho_1, \mathbf{re}:\forall[].\{\mathbf{sp}:\rho_2, \mathbf{r1}:exn\}, \mathbf{sp}':\rho_2\} \quad (\text{illegal})$$

Unfortunately, this calling convention uses two stack pointers, and STAL has only one stack.³ Observe, though, that the exception continuation's stack is necessarily a tail of the ordinary continuation's stack. This observation leads to the following calling convention for exceptions with stacks:

$$\forall[\rho_1, \rho_2].\{\mathbf{sp}:\rho_1 \circ \rho_2, \mathbf{r1}:int, \mathbf{ra}:\forall[].\{\mathbf{sp}:\rho_1 \circ \rho_2, \mathbf{r1}:\langle\rangle\}, \mathbf{re}':ptr(\rho_2), \mathbf{re}:\forall[].\{\mathbf{sp}:\rho_2, \mathbf{r1}:exn\}\}$$

This type uses some two new constructs we now add to STAL (see Figure 5). When σ_1 and σ_2 are stack types, the stack type $\sigma_1 \circ \sigma_2$ is the result of appending the two types. Thus, in the above type, the function is presented with a stack with type $\rho_1 \circ \rho_2$, all of which is expected by the regular continuation, but only a tail of which (ρ_2) is expected by the exception continuation. Since ρ_1 and ρ_2 are quantified, the function may still be used for any stack so long as the exception continuation accepts some tail of that stack.

To raise an exception, the exception is placed in **r1** and the control is transferred to the exception continuation. This requires cutting the actual stack down to just that expected by the exception continuation. Since the length of ρ_1 is unknown, this can not be done by **sfree**. Instead, a pointer to the desired position in the stack is supplied in **re'**, and is moved into **sp**. The type $ptr(\sigma)$ is the type of pointers into the stack at a position where the stack has type σ . Such pointers are obtained simply by moving **sp** into a register.

4.2 Compound Stacks

The additional syntax to support exceptions is summarized in Figure 5. The new type constructors were discussed above. The word value $ptr(i)$ is used by the operational semantics to represent pointers into the stack; the element pointed to is i words from the bottom of the stack. (See Figure 7 for details.) Of course, on a real machine, these would be implemented by actual pointers. The instructions **mov r_d, sp** and **mov sp, r_s** save and restore the stack pointer, and the instructions **sld r_d, r_s(i)** and **sst r_d(i), r_s** allow for loading from and storing to pointers.

The introduction of pointers into the stack raises a delicate issue for the typing system. When the stack pointer is copied into a register, changes to the stack are not reflected in the type of the copy, and can invalidate a pointer. Consider the following incorrect code:

```
% begin with sp :  $\tau::\sigma$ , sp  $\mapsto w::S$  ( $\tau \neq ns$ )
mov r1, sp    % r1 :  $ptr(\tau::\sigma)$ 
sfree 1      % sp :  $\sigma$ , sp  $\mapsto S$ 
salloc 1     % sp :  $ns::\sigma$ , sp  $\mapsto ns::S$ 
sld r2, r1(0) % r2 :  $\tau$  but r2  $\mapsto ns$ 
```

³Some language implementations use a separate exception stack; with some minor modifications, this calling convention would be satisfactory for such implementations.

<i>types</i>	τ	$::=$	$\dots \mid ptr(\sigma)$
<i>stack types</i>	σ	$::=$	$\dots \mid \sigma_1 \circ \sigma_2$
<i>word values</i>	w	$::=$	$\dots \mid ptr(i)$
<i>instructions</i>	ι	$::=$	$\dots \mid \text{mov } r_d, \text{sp} \mid \text{mov } \text{sp}, r_s \mid \text{sld } r_d, r_s(i) \mid \text{sst } r_d(i), r_s$

Figure 5: Additions to TAL for Compound Stacks

When execution reaches the final line, **r1** still has type $ptr(\tau::\sigma)$, but this type is no longer consistent with the state of the stack; the pointer in **r1** points to *ns*.

To prohibit erroneous loads of this sort, the type system requires that the pointer r_s be *valid* in the instructions $\text{sld } r_d, r_s(i)$, $\text{sst } r_d(i), r_s$, and $\text{mov } \text{sp}, r_s$. An invariant of our system is that the type of **sp** always describes the current stack, so using a pointer into the stack will be sound if that pointer's type is consistent with **sp**'s type. Suppose **sp** has type σ_1 and r has type $ptr(\sigma_2)$, then r is valid if σ_2 is a tail of σ_1 (formally, if there exists some σ' such that $\sigma_1 = \sigma' \circ \sigma_2$). If a pointer is invalid, it may be neither loaded from nor moved into the stack pointer. In the above example the load will be rejected because **r1**'s type $\tau::\sigma$ is not a tail of **sp**'s type, $ns::\sigma$.

4.3 Using Compound Stacks

Recall the type for a function in the presence of exceptions:

$$\forall[\rho_1, \rho_2].\{\text{sp}:\rho_1 \circ \rho_2, \text{r1}:int, \text{ra}:\forall[.]\{\text{sp}:\rho_1 \circ \rho_2, \text{r1}:\langle\rangle\}, \text{re}':ptr(\rho_2), \text{re}:\forall[.]\{\text{sp}:\rho_2, \text{r1}:exn\}\}$$

An exception may be raised within the body of such a function by restoring the handler's stack from **re'** and jumping to the handler. A new exception handler may be installed by copying the stack pointer to **re'** and making forthcoming function calls with the stack type variables instantiated to *nil* and $\rho_1 \circ \rho_2$. Calls that do not install new exception handlers would attach their frames to ρ_1 and pass on ρ_2 unchanged.

Since exceptions are probably raised infrequently, an implementation could save a register by storing the exception continuation's code pointer on the stack, instead of in its own register. If this convention were used, functions would expect stacks with the type $\rho_1 \circ (\tau_{\text{handler}}::\rho_2)$ and exception pointers with the type $ptr(\tau_{\text{handler}}::\rho_2)$ where $\tau_{\text{handler}} = \forall[.]\{\text{sp}:\rho_2, \text{r1}:exn\}$.

This last convention illustrates a use for compound stacks that goes beyond implementing exceptions. We have a general tool for locating data of type τ amidst the stack by using the calling convention:

$$\forall[\rho_1, \rho_2].\{\text{sp}:\rho_1 \circ (\tau::\rho_2), \text{r1}:ptr(\tau::\rho_2), \dots\}$$

One application of this tool would be for implementing Pascal with displays. The primary limitation of this tool is that if more than one piece of data is stored amidst the stack, although quantification may be used to avoid specifying the precise locations of that data, function calling conventions would have to specify in what *order* data appears on the stack. It appears that this limitation could be removed by introducing a limited form of intersection type, but we have not yet explored the ramifications of this enhancement.

5 Related and Future Work

Our work is partially inspired by Reynolds [26], which uses functor categories to “replace continuations by instruction sequences and store shapes by descriptions of the structure of the run-time stack.” However, Reynolds was primarily concerned with using functors to express an intermediate language of a semantics-based compiler for Algol, whereas we are primarily concerned with type structure for general-purpose target languages.

Stata and Abadi [30] formalize the Java bytecode verifier’s treatment of subroutines by giving a type system for a subset of the Java Virtual Machine language. In particular, their type system ensures that for any program control point, the Java stack is of the same size each time that control point is reached during execution. Consequently, procedure call must be a primitive construct (which it is in JVM). In contrast, our treatment supports polymorphic stack recursion, and hence procedure calls can be encoded with existing assembly-language primitives.

Tofte and others [8, 33] have developed an allocation strategy involving regions. Regions are lexically scoped containers that have a LIFO ordering on their lifetimes, much like the values on a stack. As in our approach, polymorphic recursion on abstracted region variables plays a critical role. However, unlike the objects in our stacks, regions are variable-sized, and objects need not be allocated into the region which was most recently created. Furthermore, there is only one allocation mechanism in Tofte’s system (the stack of regions) and no need for a garbage collector. In contrast, STAL only allows allocation at the top of the stack and assumes a garbage collector for heap-allocated values. However, the type system for STAL is considerably simpler than the type system of Tofte et al., as it requires no effect information in types.

Bailey and Davidson [6] also describe a specification language for modeling procedure calling conventions and checking that implementations respect these conventions. They are able to specify features such as a variable number of arguments that our formalism does not address. However, their model is explicitly tied to a stack-based calling convention and does not address features such as exception handlers. Furthermore, their approach does not integrate the specification of calling conventions with a general-purpose type system.

Although our type system is sufficiently expressive for compilation of a number of source languages, it falls short in several areas. First, it cannot support general pointers into the stack because of the ordering requirements; nor can stack and heap pointers be unified so that a function taking a tuple argument can be passed either a heap-allocated or a stack-allocated tuple. Second, threads and advanced mechanisms for implementing first-class continuations such as the work by Hieb et al. [15] cannot be modeled in this system without adding new primitives.

However, we claim that the framework presented here is a practical approach to compilation. To substantiate this claim, we are constructing a compiler called TALC that maps the KML programming language [10] to a variant of STAL described here, suitably adapted for the Intel IA32 architecture. We have found it straightforward to enrich the target language type system to include support for other type constructors, such as references, higher-order constructors, and recursive types. The compiler uses an unboxed stack allocation style of continuation passing, as discussed in this paper.

Although we have discussed mechanisms for typing stacks at the assembly language level, our techniques generalize to other languages. The same mechanisms, including the use of polymorphic

recursion to abstract the tail of a stack, can be used to introduce explicit stacks in higher level calculi. An intermediate language with explicit stacks would allow control over allocation at a point where more information is available to guide allocation decisions.

6 Summary

We have given a type system for a typed assembly language with both a heap and a stack. Our language is flexible enough to support the following compilation techniques: CPS using both heap allocation and stack allocation, a variety of procedure calling conventions, displays, exceptions, tail call elimination, and callee-saves registers.

A key contribution of the type system is that it makes procedure calling conventions explicit and provides a means of specifying and checking calling conventions that is grounded in language theory. The type system also makes clear the relationship between heap allocation and stack allocation of continuation closures, capturing both allocation strategies in one calculus.

References

- [1] Andrew Appel and Zhong Shao. Callee-saves registers in continuation-passing style. *Lisp and Symbolic Computation*, 5:189–219, 1992.
- [2] Andrew Appel and Zhong Shao. An empirical and analytic study of stack vs. heap cost for languages with closures. *Journal of Functional Programming*, 1(1), January 1993.
- [3] Andrew W. Appel. *Compiling with Continuations*. Cambridge University Press, 1992.
- [4] Andrew W. Appel and Trevor Jim. Continuation-passing, closure-passing style. In *Sixteenth ACM Symposium on Principles of Programming Languages*, pages 293–302, Austin, January 1989.
- [5] Andrew W. Appel and David B. MacQueen. Standard ML of New Jersey. In Martin Wirsing, editor, *Third International Symposium on Programming Language Implementation and Logic Programming*, pages 1–13, New York, August 1991. Springer-Verlag. Volume 528 of *Lecture Notes in Computer Science*.
- [6] Mark Bailey and Jack Davidson. A formal model of procedure calling conventions. In *Twenty-Second ACM Symposium on Principles of Programming Languages*, pages 298–310, San Francisco, January 1995.
- [7] Lars Birkedal, Nick Rothwell, Mads Tofte, and David N. Turner. The ML Kit (version 1). Technical Report 93/14, Department of Computer Science, University of Copenhagen, 1993.
- [8] Lars Birkedal, Mads Tofte, and Magnus Vejlstrup. From region inference to von Neumann machines via region representation inference. In *Twenty-Third ACM Symposium on Principles of Programming Languages*, pages 171–183, St. Petersburg, January 1996.
- [9] Val Breazu-Tannen, Thierry Coquand, Carl A. Gunter, and Andre Scedrov. Inheritance as implicit coercion. *Information and Computation*, 93:172–221, 1991.
- [10] Karl Crary. *KML Reference Manual*. Department of Computer Science, Cornell University, 1996.
- [11] Karl Crary. Foundations for the implementation of higher-order subtyping. In *ACM SIGPLAN International Conference on Functional Programming*, pages 125–135, Amsterdam, June 1997.
- [12] Allyn Dimock, Robert Muller, Franklyn Turbak, and J. B. Wells. Strongly typed flow-directed representation transformations. In *ACM SIGPLAN International Conference on Functional Programming*, pages 85–98, Amsterdam, June 1997.
- [13] Amer Diwan, David Tarditi, and Eliot Moss. Memory subsystem performance of programs using copying garbage collection. In *Twenty-First ACM Symposium on Principles of Programming Languages*, pages 1–14, January 1994.
- [14] Amer Diwan, David Tarditi, and Eliot Moss. Memory system performance of programs with intensive heap allocation. *ACM Transactions on Computer Systems*, 13(3):244–273, August 1995.

- [15] Robert Hieb, R. Kent Dybvig, and Carl Bruggeman. Representing control in the presence of first-class continuations. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 66–77, June 1990. Published as *SIGPLAN Notices*, 25(6).
- [16] Intel Corporation. *Intel Architecture Optimization Manual*. Intel Corporation, P.O. Box 7641, Mt. Prospect, IL, 60056-7641, 1997.
- [17] David Kranz, R. Kelsey, J. Rees, P. R. Hudak, J. Philbin, and N. Adams. ORBIT: An optimizing compiler for Scheme. In *Proceedings of the ACM SIGPLAN '86 Symposium on Compiler Construction*, pages 219–233, June 1986.
- [18] P. J. Landin. The mechanical evaluation of expressions. *Computer J.*, 6(4):308–20, 1964.
- [19] Xavier Leroy. Unboxed objects and polymorphic typing. In *Nineteenth ACM Symposium on Principles of Programming Languages*, pages 177–188, Albuquerque, January 1992.
- [20] Tim Lindholm and Frank Yellin. *The Java Virtual Machine Specification*. Addison-Wesley, 1996.
- [21] Y. Minamide, G. Morrisett, and R. Harper. Typed closure conversion. In *Twenty-Third ACM Symposium on Principles of Programming Languages*, pages 271–283, St. Petersburg, January 1996.
- [22] G. Morrisett, D. Tarditi, P. Cheng, C. Stone, R. Harper, and P. Lee. The TIL/ML compiler: Performance and safety through types. In *Workshop on Compiler Support for Systems Software*, Tucson, February 1996.
- [23] Greg Morrisett, David Walker, Karl Cray, and Neal Glew. From System F to typed assembly language. In *Twenty-Fifth ACM Symposium on Principles of Programming Languages*, San Diego, January 1998. Extended version published as Cornell University technical report TR97-1651, November 1997.
- [24] Gregory Morrisett. *Compiling with Types*. PhD thesis, Carnegie Mellon University, 1995. Published as CMU Technical Report CMU-CS-95-226.
- [25] Simon L. Peyton Jones, Cordelia V. Hall, Kevin Hammond, Will Partain, and Philip Wadler. The Glasgow Haskell compiler: a technical overview. In *Proc. UK Joint Framework for Information Technology (JFIT) Technical Conference*, July 1993.
- [26] John Reynolds. Using functor categories to generate intermediate code. In *Twenty-Second ACM Symposium on Principles of Programming Languages*, pages 25–36, San Francisco, January 1995.
- [27] John C. Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing '83*, pages 513–523. North-Holland, 1983. Proceedings of the IFIP 9th World Computer Congress.
- [28] Z. Shao. An overview of the FLINT/ML compiler. In *Workshop on Types in Compilation*, Amsterdam, June 1997. ACM SIGPLAN. Published as Boston College Computer Science Dept. Technical Report BCCS-97-03.
- [29] Zhong Shao. Flexible representation analysis. In *ACM SIGPLAN International Conference on Functional Programming*, pages 85–98, Amsterdam, June 1997.
- [30] Raymie Stata and Martín Abadi. A type system for java bytecode subroutines. In *Twenty-Fifth ACM Symposium on Principles of Programming Languages*, San Diego, January 1998.
- [31] Guy L. Steele Jr. Rabbit: A compiler for Scheme. Master's thesis, MIT, 1978.
- [32] D. Tarditi, G. Morrisett, P. Cheng, C. Stone, R. Harper, and P. Lee. TIL: A type-directed optimizing compiler for ML. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 181–192, Philadelphia, May 1996.
- [33] Mads Tofte and Jean-Pierre Talpin. Implementation of the typed call-by-value λ -calculus using a stack of regions. In *Twenty-First ACM Symposium on Principles of Programming Languages*, pages 188–201, January 1994.
- [34] Mitchell Wand. Continuation-based multiprocessing. In *Proceedings of the 1980 LISP Conference*, pages 19–28, August 1980.

A Formal STAL Semantics

This appendix contains a complete technical description of our calculus, STAL. The STAL abstract machine is very similar to the TAL abstract machine (described in detail in Morrisett et al. [23]). The syntax appears in Figure 6. The operational semantics is given as a deterministic rewriting system in Figure 7. The notation $a[b/c]$ denotes capture avoiding substitution of b for c in a . The notation $a\{b \mapsto c\}$ represents map update:

$$\begin{aligned} \{b \mapsto c', b_1 \mapsto c_1, \dots, b_n \mapsto c_n\}\{b \mapsto c\} &= \{b \mapsto c, b_1 \mapsto c_1, \dots, b_n \mapsto c_n\} \\ \{b_1 \mapsto c_1, \dots, b_n \mapsto c_n\}\{b \mapsto c\} &= \{b \mapsto c, b_1 \mapsto c_1, \dots, b_n \mapsto c_n\} \quad b \notin \{b_1, \dots, b_n\} \end{aligned}$$

To make the presentation simpler for the branching rules, some extra notation is used for expressing sequences of type and stack type instantiations. We introduce a new syntactic class (ψ) for type sequences:

$$\psi ::= \cdot \mid \tau, \psi \mid \sigma, \psi$$

The notation $w[\psi]$ stands for the obvious iteration of instantiations; the substitution notation $I[\psi/\Delta]$ is defined by:

$$\begin{aligned} I[\cdot/\cdot] &= I \\ I[\tau, \psi/\alpha, \Delta] &= I[\tau/\alpha][\psi/\Delta] \\ I[\sigma, \psi/\rho, \Delta] &= I[\sigma/\rho][\psi/\Delta] \end{aligned}$$

The static semantics is similar to TAL's but requires extra judgments for definitional equality of types, stack types, and register file types and uses a more compositional style for instructions. Definitional equality is needed because two stack types (such as $(int::nil) \circ (int::nil)$ and $int::int::nil$) may be syntactically different but represent the same type. The judgments are summarized in Figure 8, the rules for type judgments appear in Figure 9, and the rules for term judgments appear in Figures 10 and 11. The notation Δ', Δ denotes appending Δ' to the front of Δ , that is:

$$\begin{aligned} \cdot, \Delta &= \Delta \\ (\alpha, \Delta'), \Delta &= \alpha, (\Delta', \Delta) \\ (\rho, \Delta'), \Delta &= \rho, (\Delta', \Delta) \end{aligned}$$

As with TAL, STAL is type sound:

Theorem A.1 (Type Soundness) *If $\vdash P$ and $P \mapsto^* P'$ then P' is not stuck.*

This theorem is proved using the following two lemmas.

Lemma A.2 (Subject Reduction) *If $\vdash P$ and $P \mapsto P'$ then $\vdash P'$.*

A well-formed terminal state has the form $(H, R\{\mathbf{r1} \mapsto w\}, \mathbf{halt} \ [\tau])$ where there exists a Ψ such that $\vdash H : \Psi$ and $\Psi; \cdot \vdash w : \tau$ wval.

Lemma A.3 (Progress) *If $\vdash P$ then either P is a well-formed terminal state or there exists P' such that $P \mapsto P'$.*

<i>types</i>	$\tau ::= \alpha \mid \text{int} \mid \text{ns} \mid \forall[\Delta].\Gamma \mid \langle \tau_1^{\varphi_1}, \dots, \tau_n^{\varphi_n} \rangle \mid \exists \alpha. \tau \mid \text{ptr}(\sigma)$
<i>stack types</i>	$\sigma ::= \rho \mid \text{nil} \mid \tau :: \sigma \mid \sigma_1 \circ \sigma_2$
<i>initialization flags</i>	$\varphi ::= 0 \mid 1$
<i>label assignments</i>	$\Psi ::= \{\ell_1:\tau_1, \dots, \ell_n:\tau_n\}$
<i>type assignments</i>	$\Delta ::= \cdot \mid \alpha, \Delta \mid \rho, \Delta$
<i>register assignments</i>	$\Gamma ::= \{r_1:\tau_1, \dots, r_n:\tau_n, \text{sp}:\sigma\}$
<i>registers</i>	$r ::= \text{r1} \mid \dots \mid \text{rk}$
<i>word values</i>	$w ::= \ell \mid i \mid \text{ns} \mid ?\tau \mid w[\tau] \mid w[\sigma] \mid \text{pack } [\tau, w] \text{ as } \tau' \mid \text{ptr}(i)$
<i>small values</i>	$v ::= r \mid w \mid v[\tau] \mid v[\sigma] \mid \text{pack } [\tau, v] \text{ as } \tau'$
<i>heap values</i>	$h ::= \langle w_1, \dots, w_n \rangle \mid \text{code}[\Delta]\Gamma.I$
<i>heaps</i>	$H ::= \{\ell_1 \mapsto h_1, \dots, \ell_n \mapsto h_n\}$
<i>register files</i>	$R ::= \{r_1 \mapsto w_1, \dots, r_n \mapsto w_n, \text{sp} \mapsto S\}$
<i>stacks</i>	$S ::= \text{nil} \mid w :: S$
<i>instructions</i>	$\iota ::= \text{aop } r_d, r_s, v \mid \text{bop } r, v \mid \text{ld } r_d, r_s(i) \mid \text{malloc } r[\vec{r}] \mid$ $\text{mov } r_d, v \mid \text{mov } \text{sp}, r_s \mid \text{mov } r_d, \text{sp} \mid \text{salloc } n \mid$ $\text{sfree } n \mid \text{sld } r_d, \text{sp}(i) \mid \text{sld } r_d, r_s(i) \mid$ $\text{sst } \text{sp}(i), r_s \mid \text{sst } r_d(i), r_s \mid \text{st } r_d(i), r_s \mid$ $\text{unpack } [\alpha, r_d], v$
<i>arithmetic ops</i>	$\text{aop} ::= \text{add} \mid \text{sub} \mid \text{mul}$
<i>branch ops</i>	$\text{bop} ::= \text{beq} \mid \text{bneq} \mid \text{bgt} \mid \text{blt} \mid \text{bgte} \mid \text{blte}$
<i>instruction sequences</i>	$I ::= \iota; I \mid \text{jmp } v \mid \text{halt } [\tau]$
<i>programs</i>	$P ::= (H, R, I)$

Figure 6: Syntax of STAL

$(H, R, I) \mapsto P$ where	
if $I =$	then $P =$
add $r_d, r_s, v; I'$	$(H, R\{r_d \mapsto R(r_s) + \hat{R}(v)\}, I')$ and similarly for mul and sub
beq $r, v; I'$ when $R(r) \neq 0$	(H, R, I') and similarly for bneq , blt , <i>etc.</i>
beq $r, v; I'$ when $R(r) = 0$	$(H, R, I''[\psi/\Delta])$ where $\hat{R}(v) = \ell[\psi]$ and $H(\ell) = \text{code}[\Delta]\Gamma.I''$ and similarly for bneq , blt , <i>etc.</i>
jmp v	$(H, R, I''[\psi/\Delta])$ where $\hat{R}(v) = \ell[\psi]$ and $H(\ell) = \text{code}[\Delta]\Gamma.I'$
ld $r_d, r_s(i); I'$	$(H, R\{r_d \mapsto w_i\}, I')$ where $R(r_s) = \ell$ and $H(\ell) = \langle w_0, \dots, w_{n-1} \rangle$ and $0 \leq i < n$
malloc $r_d[\tau_1, \dots, \tau_n]; I'$	$(H\{\ell \mapsto \langle ?\tau_1, \dots, ?\tau_n \rangle\}, R\{r_d \mapsto \ell\}, I')$ where $\ell \notin H$
mov $r_d, v; I'$	$(H, R\{r_d \mapsto \hat{R}(v)\}, I')$
mov $r_d, \text{sp}; I'$	$(H, R\{r_d \mapsto \text{ptr}(S)\}, I')$
mov $\text{sp}, r_s; I'$	$(H, R\{\text{sp} \mapsto w_j :: \dots :: w_1 :: \text{nil}\}, I')$ where $R(\text{sp}) = w_n :: \dots :: w_1 :: \text{nil}$ and $R(r_s) = \text{ptr}(j)$ with $0 \leq j \leq n$
salloc $n; I'$	$(H, R\{\text{sp} \mapsto \underbrace{ns :: \dots :: ns}_n :: R(\text{sp})\}, I')$
sfree $n; I'$	$(H, R\{\text{sp} \mapsto S\}, I')$ where $R(\text{sp}) = w_1 :: \dots :: w_n :: S$
sld $r_d, \text{sp}(i); I'$	$(H, R\{r_d \mapsto w_i\}, I')$ where $R(\text{sp}) = w_0 :: \dots :: w_{n-1} :: \text{nil}$ and $0 \leq i < n$
sld $r_d, r_s(i); I'$	$(H, R\{r_d \mapsto w_{j-i}\}, I')$ where $R(r_s) = \text{ptr}(j)$ and $R(\text{sp}) = w_n :: \dots :: w_1 :: \text{nil}$ and $0 \leq i < j \leq n$
sst $\text{sp}(i), r_s; I'$	$(H, R\{\text{sp} \mapsto w_0 :: \dots :: w_{i-1} :: R(r_s) :: S\}, I')$ where $R(\text{sp}) = w_0 :: \dots :: w_i :: S$ and $0 \leq i$
sst $r_d(i), r_s; I'$	$(H, R\{\text{sp} \mapsto w_n :: \dots :: w_{j-i+1} :: R(r_s) :: w_{j-i-1} :: \dots :: w_1 :: \text{nil}\}, I')$ where $R(r_d) = \text{ptr}(j)$ and $R(\text{sp}) = w_n :: \dots :: w_1 :: \text{nil}$ and $0 \leq i < j \leq n$
st $r_d(i), r_s; I'$	$(H\{\ell \mapsto \langle w_0, \dots, w_{i-1}, R(r_s), w_{i+1}, \dots, w_{n-1} \rangle\}, R, I')$ where $R(r_d) = \ell$ and $H(\ell) = \langle w_0, \dots, w_{n-1} \rangle$ and $0 \leq i < n$
unpack $[\alpha, r_d], v; I'$	$(H, R\{r_d \mapsto w\}, I'[\tau/\alpha])$ where $\hat{R}(v) = \text{pack}[\tau, w]$ as τ'

$$\text{Where } \hat{R}(v) = \begin{cases} R(r) & \text{when } v = r \\ w & \text{when } v = w \\ \hat{R}(v')[\tau] & \text{when } v = v'[\tau] \\ \text{pack}[\tau, \hat{R}(v')] \text{ as } \tau' & \text{when } v = \text{pack}[\tau, v'] \text{ as } \tau' \end{cases}$$

Figure 7: Operational Semantics of STAL

Judgement	Meaning
$\Delta \vdash \tau$	τ is a valid type
$\Delta \vdash \sigma$	σ is a valid stack type
$\vdash \Psi$	Ψ is a valid heap type (no context is used because heap types must be closed)
$\Delta \vdash \Gamma$	Γ is a valid register file type
$\Delta \vdash \tau_1 = \tau_2$	τ_1 and τ_2 are equal types
$\Delta \vdash \sigma_1 = \sigma_2$	σ_1 and σ_2 are equal stack types
$\Delta \vdash \Gamma_1 = \Gamma_2$	Γ_1 and Γ_2 are equal register file types
$\Delta \vdash \tau_1 \leq \tau_2$	τ_1 is a subtype of τ_2
$\Delta \vdash \Gamma_1 \leq \Gamma_2$	Γ_1 is a register file subtype of Γ_2
$\vdash H : \Psi$	the heap H has type Ψ
$\Psi \vdash S : \sigma$	the stack S has type σ
$\Psi \vdash R : \Gamma$	the register file R has type Γ
$\Psi \vdash h : \tau \text{ hval}$	the heap value h has type τ
$\Psi; \Delta \vdash w : \tau \text{ wval}$	the word value w has type τ
$\Psi; \Delta \vdash w : \tau^\varphi \text{ fwval}$	the word value w has flagged type τ^φ (<i>i.e.</i> , w has type τ or w is $?\tau$ and φ is 0)
$\Psi; \Delta; \Gamma \vdash v : \tau$	the small value v has type τ
$\Psi; \Delta; \Gamma \vdash \iota \Rightarrow \Delta'; \Gamma'$	instruction ι requires a context of type $\Psi; \Delta; \Gamma$ and produces a context of type $\Psi; \Delta'; \Gamma'$
$\Psi; \Delta; \Gamma \vdash I$	I is a valid sequence of instructions
$\vdash P$	P is a valid program

Figure 8: Static Semantics of STAL (judgments)

$$\boxed{\Delta \vdash \tau \quad \Delta \vdash \sigma \quad \vdash \Psi \quad \Delta \vdash \Gamma}$$

$$\frac{\Delta \vdash \tau = \tau}{\Delta \vdash \tau} \quad \frac{\Delta \vdash \sigma = \sigma}{\Delta \vdash \sigma} \quad \frac{\cdot \vdash \tau_i}{\vdash \{\ell_1 \mapsto \tau_1, \dots, \ell_n \mapsto \tau_n\}} \quad \frac{\Delta \vdash \Gamma = \Gamma}{\Delta \vdash \Gamma}$$

$$\boxed{\Delta \vdash \tau_1 = \tau_2 \quad \Delta \vdash \sigma_1 = \sigma_2 \quad \Delta \vdash \Gamma_1 = \Gamma_2}$$

$$\frac{\Delta \vdash \tau_2 = \tau_1}{\Delta \vdash \tau_1 = \tau_2} \quad \frac{\Delta \vdash \tau_1 = \tau_2 \quad \Delta \vdash \tau_2 = \tau_3}{\Delta \vdash \tau_1 = \tau_3} \quad \frac{\Delta \vdash \sigma_2 = \sigma_1}{\Delta \vdash \sigma_1 = \sigma_2} \quad \frac{\Delta \vdash \sigma_1 = \sigma_2 \quad \Delta \vdash \sigma_2 = \sigma_3}{\Delta \vdash \sigma_1 = \sigma_3}$$

$$\begin{array}{c} \frac{}{\Delta \vdash \alpha = \alpha} (\alpha \in \Delta) \quad \frac{}{\Delta \vdash \text{int} = \text{int}} \\ \frac{\Delta', \Delta \vdash \Gamma_1 = \Gamma_2}{\Delta \vdash \forall[\Delta']. \Gamma_1 = \forall[\Delta']. \Gamma_2} \quad \frac{\Delta \vdash \tau_i = \tau'_i}{\Delta \vdash \langle \tau_1^{\varphi_1}, \dots, \tau_n^{\varphi_n} \rangle = \langle \tau_1'^{\varphi_1}, \dots, \tau_n'^{\varphi_n} \rangle} \\ \frac{\alpha, \Delta \vdash \tau_1 = \tau_2}{\Delta \vdash \exists \alpha. \tau_1 = \exists \alpha. \tau_2} \quad \frac{}{\Delta \vdash ns = ns} \quad \frac{\Delta \vdash \sigma_1 = \sigma_2}{\Delta \vdash ptr(\sigma_1) = ptr(\sigma_2)} \\ \frac{}{\Delta \vdash \rho = \rho} (\rho \in \Delta) \quad \frac{}{\Delta \vdash nil = nil} \\ \frac{\Delta \vdash \tau_1 = \tau_2 \quad \Delta \vdash \sigma_1 = \sigma_2}{\Delta \vdash \tau_1 :: \sigma_1 = \tau_2 :: \sigma_2} \quad \frac{\Delta \vdash \sigma_1 = \sigma'_1 \quad \Delta \vdash \sigma_2 = \sigma'_2}{\Delta \vdash \sigma_1 \circ \sigma_2 = \sigma'_1 \circ \sigma'_2} \\ \frac{\Delta \vdash \sigma}{\Delta \vdash nil \circ \sigma = \sigma} \quad \frac{\Delta \vdash \sigma}{\Delta \vdash \sigma \circ nil = \sigma} \\ \frac{\Delta \vdash \tau \quad \Delta \vdash \sigma_1 \quad \Delta \vdash \sigma_2}{\Delta \vdash (\tau :: \sigma_1) \circ \sigma_2 = \tau :: (\sigma_1 \circ \sigma_2)} \\ \frac{\Delta \vdash \sigma_1 \quad \Delta \vdash \sigma_2 \quad \Delta \vdash \sigma_3}{\Delta \vdash (\sigma_1 \circ \sigma_2) \circ \sigma_3 = \sigma_1 \circ (\sigma_2 \circ \sigma_3)} \\ \frac{\Delta \vdash \sigma = \sigma' \quad \Delta \vdash \tau_i = \tau'_i}{\Delta \vdash \{\mathbf{sp}:\sigma, r_1 \mapsto \tau_1, \dots, r_n \mapsto \tau_n\} = \{\mathbf{sp}:\sigma', r_1:\tau'_1, \dots, r_n:\tau'_n\}} \end{array}$$

$$\boxed{\Delta \vdash \tau_1 \leq \tau_2 \quad \Delta \vdash \Gamma_1 \leq \Gamma_2}$$

$$\begin{array}{c} \frac{\Delta \vdash \tau_1 = \tau_2}{\Delta \vdash \tau_1 \leq \tau_2} \quad \frac{\Delta \vdash \tau_1 \leq \tau_2 \quad \Delta \vdash \tau_2 \leq \tau_3}{\Delta \vdash \tau_1 \leq \tau_3} \\ \frac{\Delta \vdash \tau_i}{\Delta \vdash \langle \tau_1^{\varphi_1}, \dots, \tau_{i-1}^{\varphi_{i-1}}, \tau_i^1, \tau_{i+1}^{\varphi_{i+1}}, \dots, \tau_n^{\varphi_n} \rangle \leq \langle \tau_1^{\varphi_1}, \dots, \tau_{i-1}^{\varphi_{i-1}}, \tau_i^0, \tau_{i+1}^{\varphi_{i+1}}, \dots, \tau_n^{\varphi_n} \rangle} \\ \frac{\Delta \vdash \sigma = \sigma' \quad \Delta \vdash \tau_i = \tau'_i \quad (\text{for } 1 \leq i \leq n) \quad \Delta \vdash \tau_i \quad (\text{for } n < i \leq m)}{\Delta \vdash \{\mathbf{sp}:\sigma, r_1:\tau_1, \dots, r_m:\tau_m\} \leq \{\mathbf{sp}:\sigma', r_1:\tau'_1, \dots, r_n:\tau'_n\}} (m \geq n) \end{array}$$

Figure 9: Static Semantics of STAL, Judgments for Types

$\vdash P \quad \vdash H : \Psi \quad \Psi \vdash S : \sigma \quad \Psi \vdash R : \Gamma$
--

$$\begin{array}{c}
\frac{\vdash H : \Psi \quad \Psi \vdash R : \Gamma \quad \Psi; \cdot; \Gamma \vdash I}{\vdash (H, R, I)} \quad \frac{\vdash \Psi \quad \Psi \vdash h_i : \tau_i \text{ hval}}{\vdash \{\ell_1 \mapsto h_1, \dots, \ell_n \mapsto h_n\} : \Psi} \quad (\Psi = \{\ell_1 : \tau_1, \dots, \ell_n : \tau_n\}) \\
\\
\frac{\Psi \vdash nil : nil}{\Psi; \cdot \vdash w : \tau \text{ wval}} \quad \frac{\Psi; \cdot \vdash w : \tau \text{ wval} \quad \Psi \vdash S : \sigma}{\Psi \vdash w :: S : \tau :: \sigma} \\
\\
\frac{\Psi \vdash S : \sigma \quad \Psi; \cdot \vdash w_i : \tau_i \text{ wval} \quad (\text{for } 1 \leq i \leq n)}{\Psi \vdash \{\mathbf{sp} \mapsto S, r_1 \mapsto w_1, \dots, r_m \mapsto w_m\} : \{\mathbf{sp} : \sigma, r_1 : \tau_1, \dots, r_n : \tau_n\}} \quad (m \geq n)
\end{array}$$

$\Psi \vdash h : \tau \text{ hval} \quad \Psi; \Delta \vdash w : \tau \text{ wval} \quad \Psi; \Delta \vdash w : \tau^\varphi \text{ fwval} \quad \Psi; \Delta; \Gamma \vdash v : \tau$

$$\begin{array}{c}
\frac{\Psi; \cdot \vdash w_i : \tau_i^{\varphi_i} \text{ fwval}}{\Psi \vdash \langle w_1, \dots, w_n \rangle : \langle \tau_1^{\varphi_1}, \dots, \tau_n^{\varphi_n} \rangle \text{ hval}} \quad \frac{\Delta \vdash \Gamma \quad \Psi; \Delta; \Gamma \vdash I}{\Psi \vdash \mathbf{code}[\Delta] \Gamma. I : \forall[\Delta]. \Gamma \text{ hval}} \\
\\
\frac{\Delta \vdash \tau_1 \leq \tau_2}{\Psi; \Delta \vdash \ell : \tau_2 \text{ wval}} \quad (\Psi(\ell) = \tau_1) \quad \frac{}{\Psi; \Delta \vdash i : \text{int} \text{ wval}} \\
\\
\frac{\Delta \vdash \tau \quad \Psi; \Delta \vdash w : \forall[\alpha, \Delta']. \Gamma \text{ wval}}{\Psi; \Delta \vdash w[\tau] : \forall[\Delta']. \Gamma[\tau/\alpha] \text{ wval}} \quad \frac{\Delta \vdash \sigma \quad \Psi; \Delta \vdash w : \forall[\rho, \Delta']. \Gamma \text{ wval}}{\Psi; \Delta \vdash w[\sigma] : \forall[\Delta']. \Gamma[\sigma/\rho] \text{ wval}} \\
\\
\frac{\Delta \vdash \tau \quad \Psi; \Delta \vdash w : \tau'[\tau/\alpha] \text{ wval}}{\Psi; \Delta \vdash \mathbf{pack} [\tau, w] \text{ as } \exists \alpha. \tau' : \exists \alpha. \tau' \text{ wval}} \quad \frac{}{\Psi; \Delta \vdash \mathbf{ns} : \mathbf{ns} \text{ wval}} \\
\\
\frac{\Delta \vdash \sigma}{\Psi; \Delta \vdash \mathbf{ptr}(i) : \mathbf{ptr}(\sigma) \text{ wval}} \quad (|\sigma| = i) \\
\\
\frac{\Delta \vdash \tau}{\Psi; \Delta \vdash ?\tau : \tau^0 \text{ fwval}} \quad \frac{\Psi; \Delta \vdash w : \tau \text{ wval}}{\Psi; \Delta \vdash w : \tau^\varphi \text{ fwval}} \\
\\
\frac{}{\Psi; \Delta; \Gamma \vdash r : \tau} \quad (\Gamma(r) = \tau) \quad \frac{\Psi; \Delta \vdash w : \tau \text{ wval}}{\Psi; \Delta; \Gamma \vdash w : \tau} \\
\\
\frac{\Delta \vdash \tau \quad \Psi; \Delta; \Gamma \vdash v : \forall[\alpha, \Delta']. \Gamma'}{\Psi; \Delta; \Gamma \vdash v[\tau] : \forall[\Delta']. \Gamma'[\tau/\alpha]} \quad \frac{\Delta \vdash \sigma \quad \Psi; \Delta; \Gamma \vdash v : \forall[\rho, \Delta']. \Gamma'}{\Psi; \Delta; \Gamma \vdash v[\sigma] : \forall[\Delta']. \Gamma'[\sigma/\rho]} \\
\\
\frac{\Delta \vdash \tau \quad \Psi; \Delta; \Gamma \vdash v : \tau'[\tau/\alpha]}{\Psi; \Delta; \Gamma \vdash \mathbf{pack} [\tau, v] \text{ as } \exists \alpha. \tau' : \exists \alpha. \tau'} \\
\\
\frac{\cdot \vdash \tau_1 = \tau_2 \quad \Psi \vdash h : \tau_2 \text{ hval}}{\Psi \vdash h : \tau_1 \text{ hval}} \quad \frac{\Delta \vdash \tau_1 = \tau_2 \quad \Psi; \Delta \vdash w : \tau_2 \text{ wval}}{\Psi; \Delta \vdash w : \tau_1 \text{ wval}} \quad \frac{\Delta \vdash \tau_1 = \tau_2 \quad \Psi; \Delta; \Gamma \vdash v : \tau_2}{\Psi; \Delta; \Gamma \vdash v : \tau_1}
\end{array}$$

Figure 10: STAL Static Semantics, Term Constructs except Instructions

$$\boxed{\Psi; \Delta; \Gamma \vdash \iota \Rightarrow \Delta'; \Gamma' \quad \Psi; \Delta; \Gamma \vdash I}$$

$$\begin{array}{c}
\frac{\Psi; \Delta; \Gamma \vdash r_s : \text{int} \quad \Psi; \Delta; \Gamma \vdash v : \text{int}}{\Psi; \Delta; \Gamma \vdash \text{aop } r_d, r_s, v \Rightarrow \Delta; \Gamma\{r_d : \text{int}\}} \\
\\
\frac{\Psi; \Delta; \Gamma_1 \vdash r : \text{int} \quad \Psi; \Delta; \Gamma_1 \vdash v : \forall[]. \Gamma_2 \quad \Delta \vdash \Gamma_1 \leq \Gamma_2}{\Psi; \Delta; \Gamma_1 \vdash \text{bop } r, v \Rightarrow \Delta; \Gamma_1} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash r_s : \langle \tau_0^{\varphi_0}, \dots, \tau_{n-1}^{\varphi_{n-1}} \rangle}{\Psi; \Delta; \Gamma \vdash \text{ld } r_d, r_s(i) \Rightarrow \Delta; \Gamma\{r_d : \tau_i\}} \quad (\varphi_i = 1 \wedge 0 \leq i < n) \\
\\
\frac{\Delta \vdash \tau_i}{\Psi; \Delta; \Gamma \vdash \text{malloc } r[\tau_1, \dots, \tau_n] \Rightarrow \Delta; \Gamma\{r : \langle \tau_1^0, \dots, \tau_n^0 \rangle\}} \\
\\
\frac{\Psi; \Delta; \Gamma \vdash v : \tau}{\Psi; \Delta; \Gamma \vdash \text{mov } r_d, v \Rightarrow \Delta; \Gamma\{r_d : \tau\}} \\
\\
\frac{}{\Psi; \Delta; \Gamma \vdash \text{mov } r_d, \text{sp} \Rightarrow \Delta; \Gamma\{r_d : \text{ptr}(\sigma)\}} \quad (\Gamma(\text{sp}) = \sigma) \\
\\
\frac{\Psi; \Delta; \Gamma \vdash r_s : \text{ptr}(\sigma_2) \quad \Delta \vdash \sigma_1 = \sigma_3 \circ \sigma_2}{\Psi; \Delta; \Gamma \vdash \text{mov } \text{sp}, r_s \Rightarrow \Delta; \Gamma\{\text{sp} : \sigma_2\}} \quad (\Gamma(\text{sp}) = \sigma_1) \\
\\
\frac{}{\Psi; \Delta; \Gamma \vdash \text{salloc } n \Rightarrow \Delta; \Gamma\{\text{sp} : \underbrace{ns :: \dots :: ns}_{n} :: \sigma\}} \quad (\Gamma(\text{sp}) = \sigma) \\
\\
\frac{\Delta \vdash \sigma_1 = \tau_0 :: \dots :: \tau_{n-1} :: \sigma_2}{\Psi; \Delta; \Gamma \vdash \text{sfree } n \Rightarrow \Delta; \Gamma\{\text{sp} : \sigma_2\}} \quad (\Gamma(\text{sp}) = \sigma_1) \\
\\
\frac{\Delta \vdash \sigma_1 = \tau_0 :: \dots :: \tau_i :: \sigma_2}{\Psi; \Delta; \Gamma \vdash \text{sld } r_d, \text{sp}(i) \Rightarrow \Delta; \Gamma\{r_d : \tau_i\}} \quad (\Gamma(\text{sp}) = \sigma_1 \wedge 0 \leq i) \\
\\
\frac{\Psi; \Delta; \Gamma \vdash r_s : \text{ptr}(\sigma_3) \quad \Delta \vdash \sigma_1 = \sigma_2 \circ \sigma_3 \quad \Delta \vdash \sigma_3 = \tau_0 :: \dots :: \tau_i :: \sigma_4}{\Psi; \Delta; \Gamma \vdash \text{sld } r_d, r_s(i) \Rightarrow \Delta; \Gamma\{r_d : \tau_i\}} \quad (\Gamma(\text{sp}) = \sigma_1 \wedge 0 \leq i) \\
\\
\frac{\Delta \vdash \sigma_1 = \tau_0 :: \dots :: \tau_i :: \sigma_2 \quad \Psi; \Delta; \Gamma \vdash r_s : \tau}{\Psi; \Delta; \Gamma \vdash \text{sst } \text{sp}(i), r_s \Rightarrow \Delta; \Gamma\{\text{sp} : \tau_0 :: \dots :: \tau_{i-1} :: \tau :: \sigma_2\}} \quad (\Gamma(\text{sp}) = \sigma_1 \wedge 0 \leq i) \\
\\
\frac{\Psi; \Delta; \Gamma \vdash r_d : \text{ptr}(\sigma_3) \quad \Psi; \Delta; \Gamma \vdash r_s : \tau \quad \Delta \vdash \sigma_1 = \sigma_2 \circ \sigma_3 \quad \Delta \vdash \sigma_3 = \tau_0 :: \dots :: \tau_i :: \sigma_4 \quad \Delta \vdash \sigma_5 = \tau_0 :: \dots :: \tau_{i-1} :: \tau :: \sigma_4}{\Psi; \Delta; \Gamma \vdash \text{sst } r_d(i), r_s \Rightarrow \Delta; \Gamma\{\text{sp} : \sigma_2 \circ \sigma_5, r_d : \text{ptr}(\sigma_5)\}} \quad (\Gamma(\text{sp}) = \sigma_1 \wedge 0 \leq i) \\
\\
\frac{\Psi; \Delta; \Gamma \vdash r_d : \langle \tau_0^{\varphi_0}, \dots, \tau_{n-1}^{\varphi_{n-1}} \rangle \quad \Psi; \Delta; \Gamma \vdash r_s : \tau_i}{\Psi; \Delta; \Gamma \vdash \text{st } r_d(i), r_s \Rightarrow \Delta; \Gamma\{r_d : \langle \tau_0^{\varphi_0}, \dots, \tau_{i-1}^{\varphi_{i-1}}, \tau_i^1, \tau_{i+1}^{\varphi_{i+1}}, \dots, \tau_{n-1}^{\varphi_{n-1}} \rangle\}} \quad (0 \leq i < n) \\
\\
\frac{\Psi; \Delta; \Gamma \vdash v : \exists \alpha. \tau}{\Psi; \Delta; \Gamma \vdash \text{unpack } [\alpha, r_d], v \Rightarrow \alpha, \Delta; \Gamma\{r_d : \tau\}} \quad (\alpha \notin \Delta) \\
\\
\frac{\Psi; \Delta; \Gamma \vdash \iota \Rightarrow \Delta'; \Gamma' \quad \Psi; \Delta'; \Gamma' \vdash I}{\Psi; \Delta; \Gamma \vdash \iota; I} \quad \frac{\Delta \vdash \Gamma_1 \leq \Gamma_2 \quad \Psi; \Delta; \Gamma_1 \vdash v : \forall[]. \Gamma_2}{\Psi; \Delta; \Gamma_1 \vdash \text{jmp } v} \quad \frac{\Delta \vdash \tau \quad \Psi; \Delta; \Gamma \vdash \text{r1} : \tau}{\Psi; \Delta; \Gamma \vdash \text{halt } [\tau]}
\end{array}$$

Figure 11: STAL Static Semantics, Instructions