

Verifying Time Bounds for General Function Pointers

Robert Dockins^{1,*} and Aquinas Hobor^{2,**}

¹ Princeton University

² National University of Singapore

Abstract. We develop a logic of explicit time resource bounds for a language with function pointers and semantic assertions. We apply our logic to examples containing nontrivial “higher-order” uses of function pointers and we prove soundness with respect to a standard operational semantics. Our core technique is very compact and may be applicable to other resource bounding problems, and is the first application of step-indexed models in which the outermost quantifier is existential instead of universal. Our results are machine checked in Coq.

1 Introduction

We define a minimal Halting Assert Language with two distinctive features: function pointers and semantic assertions. Semantic assertions are program commands that assert the truth of a **formula in logic** at a program point. Although semantic assertions have runtime behavior equivalent to `skip`, they are useful during static analysis (*e.g.*, [BCD⁺05]) and as a mechanism for ensuring that the intermediate states of programs meet set invariants. Semantic assertions may seem benign, but their inclusion in a language with function pointers leads to an unpleasant contravariant circularity. Other domains containing a similar contravariant circularity in their semantic models (*e.g.*, concurrency with first-class locks, self-modifying code) are often quite complex in ways unrelated to the circularity. Therefore, HAL is a kind of test bed for semantic techniques that we believe may be applicable in richer settings in the future.

We design a program logic of explicit time resource bounds for HAL. Programs verified in our logic are guaranteed to satisfy all invariants given in assert statements at runtime and are verified against an explicit bound on the number of function calls they make before safely halting. We hope this kind of logic will be applicable to real-time systems, where one is interested in concrete bounds rather than simple termination.

We are unaware of any other logic of resource bounds for languages containing either function pointers or the kind of contravariant circularity present in HAL. We can reason about programs that exhibit nontrivial use of function pointers including mutually recursive function groups and higher-order functions. Each recursive group is verified as a whole and combined into proofs of whole-program termination, which makes the logic compositional. Higher-order functions are verified independently of the context in which they will be used, and we are able to apply such functions to themselves without trouble (*e.g.*, `map of map`). Our semantic model demonstrates how step-indexed models can be applied to logics of resource bounds in a compact manner.

* Supported in part by NSF CNS-0910448 and AFOSR FA9550-09-1-0138.

** Supported by a Lee Kuan Yew Postdoctoral Fellowship.

Contributions. We design a language containing function pointers and semantic assertions, develop an associated logic of time resource bounds, and apply the logic to some example programs. We build a semantic model for the Hoare judgment and prove the logic sound with respect to the model. Our results are checked in Coq and available at:

<http://msl.cs.princeton.edu/termination/>

Associated material. Interested readers can find more details, particularly on the semantic models, in unreviewed previous versions of this work [DH10].

2 An Introduction to HAL

$\chi(\tau) \equiv$ <ul style="list-style-type: none"> $x := \ell$ $x_3 := (x_1, x_2)$ $x_2 := x_1.1$ $x_2 := x_1.2$ $\chi_1(\tau); \chi_2(\tau)$ ifnil x then $\chi_1(\tau)$ else $\chi_2(\tau)$ call x return assert P 	<ul style="list-style-type: none"> load constant ℓ into x allocate a fresh pair project first component project second component sequence two commands test if $x = 0$ and branch call function pointer x return from function semantic assertion, wherein $P : \tau$
$\phi(\tau) \equiv \text{label} \rightarrow \chi(\tau)$	<ul style="list-style-type: none"> parametrized program

Fig. 1. Parameterized commands and programs

We present definitions for HAL commands and programs in Figure 1. Our syntax is parameterized over the type of assertions τ (*i.e.*, the metatype of our parameterized syntax is $\text{Type} \rightarrow \text{Type}$ instead of Type). Most of our commands are unexciting: load a constant, create (allocate) a fresh pair, project the first and second components, sequence, and branch-if-0. Note that the subcommands for sequences and branches are parameterized over the same type variable τ . Our **call** instruction is noteworthy because x is a **variable** instead of a constant function label—*i.e.*, x is a function pointer. Our functions do not take explicit arguments; instead, a programmer must establish an ad-hoc calling convention. Our **return** command is standard. The unusual command is the semantic assertion **assert**; here P has the type of the argument τ . A parameterized program $\phi(\tau)$ is a partial function from code labels to parameterized commands.

We give the basic semantic definitions for HAL in Figure 2a. We use natural numbers for program variables (for readability we use r_i instead of i for concrete program variables in our examples). We also use natural numbers for code labels (addresses). We define values as trees having labels as leaves. A store (*a.k.a.* register bank) is a partial function from variables to values. A measure is a partial function from stores to natural numbers; we will require measures to decrease during function calls. A predicate is (essentially) a function from pairs of program and store to truth values \mathcal{T} (Prop in Coq). A command is a specialization of a parameterized command χ with predicate; a stack is a list of commands. A program is a partial function from labels to commands—*i.e.*, $\text{program} = \phi(\text{predicate})$. Notice that the metatypes predicate, command, and program contain a contravariant cycle. The real semantic definition for predicate, which is similar in flavor but with the pleasing addition of being sound, is the subject of §5.1.

<table border="0"> <tr><td>variable</td><td>$x \equiv \mathbb{N}$</td></tr> <tr><td>label</td><td>$\ell \equiv \mathbb{N}$</td></tr> <tr><td>value</td><td>$v \equiv \text{label} + (\text{value} \times \text{value})$</td></tr> <tr><td>store</td><td>$\rho \equiv \text{variable} \rightarrow \text{value}$</td></tr> <tr><td>measure</td><td>$t \equiv \text{store} \rightarrow \mathbb{N}$</td></tr> <tr><td>predicate</td><td>$P \approx (\text{program} \times \text{store}) \rightarrow \mathcal{T}$</td></tr> <tr><td>command</td><td>$c \equiv \chi(\text{predicate})$</td></tr> <tr><td>stack</td><td>$s \equiv \text{list command}$</td></tr> <tr><td>program</td><td>$\Psi \equiv \text{label} \rightarrow \text{command}$</td></tr> </table>	variable	$x \equiv \mathbb{N}$	label	$\ell \equiv \mathbb{N}$	value	$v \equiv \text{label} + (\text{value} \times \text{value})$	store	$\rho \equiv \text{variable} \rightarrow \text{value}$	measure	$t \equiv \text{store} \rightarrow \mathbb{N}$	predicate	$P \approx (\text{program} \times \text{store}) \rightarrow \mathcal{T}$	command	$c \equiv \chi(\text{predicate})$	stack	$s \equiv \text{list command}$	program	$\Psi \equiv \text{label} \rightarrow \text{command}$	<table border="0"> <tr><td>\top, \perp</td><td>truth and falsehood</td></tr> <tr><td>$P \wedge Q, P \vee Q$</td><td>conjunction and disjunction</td></tr> <tr><td>$P \Rightarrow Q, \neg P$</td><td>implication and negation</td></tr> <tr><td>$\forall a : \tau, \exists a : \tau$</td><td>impredicative quantification</td></tr> <tr><td>$\mu X. P$</td><td>equirecursive predicate</td></tr> <tr><td>$x \Downarrow v$</td><td>variable x evaluates to value v</td></tr> <tr><td>$[x \leftarrow v]P$</td><td>P will hold if x is updated to v</td></tr> <tr><td>$\text{closed}(P)$</td><td>P holds on all stores</td></tr> <tr><td>$P \vdash Q$</td><td>entailment</td></tr> <tr><td>$\langle t \rangle$</td><td>measure t on the current store</td></tr> <tr><td>funptr $\ell t [\mathbb{P}] [\mathbb{Q}]$</td><td>terminating function pointer</td></tr> </table>	\top, \perp	truth and falsehood	$P \wedge Q, P \vee Q$	conjunction and disjunction	$P \Rightarrow Q, \neg P$	implication and negation	$\forall a : \tau, \exists a : \tau$	impredicative quantification	$\mu X. P$	equirecursive predicate	$x \Downarrow v$	variable x evaluates to value v	$[x \leftarrow v]P$	P will hold if x is updated to v	$\text{closed}(P)$	P holds on all stores	$P \vdash Q$	entailment	$\langle t \rangle$	measure t on the current store	funptr $\ell t [\mathbb{P}] [\mathbb{Q}]$	terminating function pointer
variable	$x \equiv \mathbb{N}$																																								
label	$\ell \equiv \mathbb{N}$																																								
value	$v \equiv \text{label} + (\text{value} \times \text{value})$																																								
store	$\rho \equiv \text{variable} \rightarrow \text{value}$																																								
measure	$t \equiv \text{store} \rightarrow \mathbb{N}$																																								
predicate	$P \approx (\text{program} \times \text{store}) \rightarrow \mathcal{T}$																																								
command	$c \equiv \chi(\text{predicate})$																																								
stack	$s \equiv \text{list command}$																																								
program	$\Psi \equiv \text{label} \rightarrow \text{command}$																																								
\top, \perp	truth and falsehood																																								
$P \wedge Q, P \vee Q$	conjunction and disjunction																																								
$P \Rightarrow Q, \neg P$	implication and negation																																								
$\forall a : \tau, \exists a : \tau$	impredicative quantification																																								
$\mu X. P$	equirecursive predicate																																								
$x \Downarrow v$	variable x evaluates to value v																																								
$[x \leftarrow v]P$	P will hold if x is updated to v																																								
$\text{closed}(P)$	P holds on all stores																																								
$P \vdash Q$	entailment																																								
$\langle t \rangle$	measure t on the current store																																								
funptr $\ell t [\mathbb{P}] [\mathbb{Q}]$	terminating function pointer																																								

Fig. 2a. Basic semantic definitions

Fig. 2b. A variety of predicates

Fig. 2. Semantic definitions and a selection of useful predicates

We give a variety of predicates in Figure 2b. We have constants (\top, \perp) and the standard logical connectives ($\wedge, \vee, \Rightarrow, \neg$). Of note is that our quantification (\forall, \exists) is **impredicative**—that is, the metavariable τ ranges over all of the types in our metalogic ($\tau : \text{Type}$ in Coq), including of course the type predicate itself. We provide equirecursive μ to describe recursive program invariants. The assertion $x \Downarrow v$ means that the variable x evaluates to value v in the current store. We write $[x \leftarrow v]P$ to mean that the predicate P will be true if the current store is updated so that variable x maps to value v ; $[x \leftarrow v]$ is therefore a kind of modal operator—the modality of store update. We define another modal operator, $\text{closed}(P)$, meaning P holds on all stores.

We write $P \vdash Q$ for predicate entailment. We also introduce a notational convenience for reasoning about measures in the context of a predicate. Since a predicate is more-or-less a function taking (among other things) a store ρ as an argument, and since a measure t is a **partial** function from stores to \mathbb{N} , it is straightforward to evaluate $t(\rho)$ and then compare the result against other naturals with the usual operators $=, <, \text{etc.}$ To indicate this kind of evaluation and comparison, we will write *e.g.*, “ $\langle t \rangle < n$ ”—that is, evaluate t with the current store and require that the result be less than n . When $t(\rho)$ is not defined, terms containing $\langle t \rangle$ are equivalent to \perp .

The assertion of particular interest is the terminating function pointer assertion “funptr $\ell t [\mathbb{P}] [\mathbb{Q}]$ ”, wherein ℓ is a function address, t is a termination measure, \mathbb{P} is a precondition, and \mathbb{Q} is a postcondition. The precondition \mathbb{P} and postcondition \mathbb{Q} are actually functions from some shared type A to predicate, *i.e.*, $\mathbb{P} = \lambda a : A. (\dots)$ and $\mathbb{Q} = \lambda a : A. (\dots)$. The function pointer assertion is actually of form “funptr $\ell A t [\mathbb{P}] [\mathbb{Q}]$ ” but A has been elided to simplify the presentation. When funptr $\ell t [\mathbb{P}] [\mathbb{Q}]$ holds:

1. The program has code c at address ℓ (recall programs are functions from labels to commands); this is why we need predicates to take programs as arguments.
2. When c is called from a context with an initial store ρ , if $t(\rho)$ is defined, then c makes at most $t(\rho)$ function calls before returning to its caller.
3. If $t(\rho)$ is defined, then for all a , if the assertion $\mathbb{P}(a)$ holds prior to executing c , then the assertion $\mathbb{Q}(a)$ will hold when c returns. The parameter a is thus able to relate pre- and postconditions to each other over the function call without auxiliary state.

3 Total Correctness for HAL

Our program logic is divided into two parts. We provide a set of Hoare rules in §3.1 for verifying commands in HAL; a strength of our approach is that these are natural. A second set of rules, in §3.2, show how to use the verification of a function’s body to show that the function satisfies its specification, and can thus be called by other functions.

3.1 Hoare Rules

Our Hoare judgment is a six-place relation written $\Gamma, R \vdash_n \{P\} c \{Q\}$, where P, Q , and R are predicates (assertions), Γ is a closed predicate that only looks at programs, n is a natural number and c is a command. We defer the formal model until §6, but the informal meaning is straightforward. P, c , and Q are the standard precondition, instruction, and postcondition triple common in Hoare logics. The return assertion R is the postcondition of the current function; R must hold before the function can `return`. We collect funptr assertions in Γ . Finally, starting from precondition P , n is an upper bound on the number of function calls c will execute before it terminates. Our logic is powerful enough that all of these parameters (including the time bound n) can take logical variables instead of concretes; indeed, the second example from §4 demonstrates that we can verify higher-order polymorphic functions independently of their call sites.

We present the Hoare rules for total correctness in Figure 3; they are mostly unsurprising. The four rules `Hlabel`, `Hcons`, `Hfetch1`, and `Hfetch2` are the standard weakest precondition forms for local variable updates for constants, fresh pairs, and first/second projections respectively. Since these rules do not make any function calls $n = 0$.

The sequence rule `Hseq` looks standard; the key point is that the upper bounds on the subcommands c_1 and c_2 are summed for the sequence. For the conditional rule `Hif`, we require that both c_1 and c_2 share the same bound n , which is then used for the whole. If the natural bounds differ, one increases the lower bound of the lower via weakening.

The weakening/consequence rule `Hweaken` allows covariance in the preconditions (P, P') and contravariance in the postconditions (Q, Q') and return conditions (R, R'). The function assertions (Γ, Γ') are related covariantly and incorporated in the other entailments in the most general way. We allow the upper bound on the number of function calls (n, n') to increase during weakening since the bound is not strict.

Although semantic assertions caused significant headaches in the semantic model due to the contravariance outlined in §2, the `Hassert` rule is pleasingly direct. We simply ensure that the precondition P (including the function assertions in Γ) entails the assertion Q . Since the `assert` command does not make any function calls, we can use $n = 0$ for the upper bound. The `Hreturn` rule requires that the precondition match the return assertion. After a function `returns` the remainder of the function is not executed, so we provide the postcondition \perp . Since `return` does not make any function calls $n = 0$.

The most important rule is `Hcall`, for verifying a function pointer call. The precondition P has five conjuncts. First, the variable x must point to a code label ℓ . Second, ℓ must be a function pointer to some code with termination measure t , function precondition \mathbb{P}_ℓ , and function postcondition \mathbb{Q}_ℓ . Third, the termination measure t must be defined on the current store, and evaluate to some n . That is, starting from the current store, the function ℓ will make no more than n function calls before `returning`. Fourth, the function precondition \mathbb{P}_ℓ must hold when applied to some a . Finally, we require that the function postcondition \mathbb{Q}_ℓ , when applied to the same a , implies the postcondition

$$\begin{array}{c}
\frac{}{\Gamma, R \vdash_0 \{[x \leftarrow \ell]Q\} x := \ell \{Q\}} \text{Hlabel} \\
\frac{}{\Gamma, R \vdash_0 \{x_1 \Downarrow v_0 \wedge x_2 \Downarrow v_1 \wedge [x_3 \leftarrow (v_0, v_1)]Q\} x_3 := (x_1, x_2) \{Q\}} \text{Hcons} \\
\frac{}{\Gamma, R \vdash_0 \{x_1 \Downarrow (v_1, v_2) \wedge [x_2 \leftarrow v_1]Q\} x_2 := x_{1.1} \{Q\}} \text{Hfetch1} \\
\frac{}{\Gamma, R \vdash_0 \{x_1 \Downarrow (v_1, v_2) \wedge [x_2 \leftarrow v_2]Q\} x_2 := x_{1.2} \{Q\}} \text{Hfetch2} \\
\frac{\Gamma, R \vdash_n \{P\} c_1 \{Q\} \quad \Gamma, R \vdash_{n'} \{Q\} c_2 \{S\}}{\Gamma, R \vdash_{n+n'} \{P\} c_1 ; c_2 \{S\}} \text{Hseq} \\
\frac{\Gamma, R \vdash_n \{P_1\} c_1 \{Q\} \quad \Gamma, R \vdash_n \{P_2\} c_2 \{Q\}}{\Gamma, R \vdash_n \{(x \Downarrow 0 \wedge P_1) \vee ((\neg(x \Downarrow 0)) \wedge P_2)\} \text{ifnil } x \text{ then } c_1 \text{ else } c_2 \{Q\}} \text{Hif} \\
\frac{n \leq n' \quad \Gamma' \wedge R \vdash R' \quad \Gamma' \wedge P' \vdash P \quad \Gamma' \vdash \Gamma \quad \Gamma' \wedge Q \vdash Q' \quad \Gamma, R \vdash_n \{P\} c \{Q\}}{\Gamma', R' \vdash_{n'} \{P'\} c \{Q'\}} \text{Hweaken} \\
\frac{\Gamma \wedge P \vdash Q}{\Gamma, R \vdash_0 \{P\} \text{assert } Q \{P\}} \text{Hassert} \\
\frac{}{\Gamma, R \vdash_0 \{R\} \text{return } \{\perp\}} \text{Hreturn} \\
\frac{P \equiv x \Downarrow \ell \wedge \text{funptr } \ell t [\mathbb{P}_\ell] [\mathbb{Q}_\ell] \wedge \langle t \rangle = n \wedge \mathbb{P}_\ell(a) \wedge \text{closed}(\mathbb{Q}_\ell(a) \Rightarrow Q)}{\Gamma, R \vdash_{n+1} \{P\} \text{call } x \{Q\}} \text{Hcall}
\end{array}$$

Fig. 3. Hoare rules

$$\begin{array}{c}
\frac{}{\Psi : \top} \text{Vstart} \\
\frac{\Psi : \Gamma \quad \forall a, n. \left((\Gamma \wedge \text{funptr } \ell t [\lambda a'. \mathbb{P}(a') \wedge \langle t \rangle < n] [\mathbb{Q}]), \right. \\
\quad \left. \mathbb{Q}(a) \vdash_n \{\mathbb{P}(a) \wedge \langle t \rangle = n\} \Psi(\ell) \{\perp\} \right)}{\Psi : (\Gamma \wedge \text{funptr } \ell t [\mathbb{P}] [\mathbb{Q}])} \text{Vsimple} \\
\frac{\Psi : \Gamma \quad \begin{array}{l} I'(b, n) \equiv \forall(\ell, t, \mathbb{P}, \mathbb{Q}) \in \Phi(b). \text{funptr } \ell t [\lambda a. \mathbb{P}(a) \wedge \langle t \rangle < n] [\mathbb{Q}] \\ \forall b. \forall(\ell, t, \mathbb{P}, \mathbb{Q}) \in \Phi(b). \left(\forall a, n. \right. \\ \left. (\Gamma \wedge I'(b, n)), \mathbb{Q}(a) \vdash_n \{\mathbb{P}(a) \wedge \langle t \rangle = n\} \Psi(\ell) \{\perp\} \right) \end{array}}{\Psi : (\Gamma \wedge \forall b. \forall(\ell, t, \mathbb{P}, \mathbb{Q}) \in \Phi(b). \text{funptr } \ell t [\mathbb{P}] [\mathbb{Q}])} \text{Vfull}
\end{array}$$

Fig. 4. Single and mutually recursive function verification

Q in all stores. It is insufficient to assert that $Q_\ell(a) \Rightarrow Q$ in the current store (*i.e.*, the store before the function call); we must know that the implication will still hold after the function call is completed. The metavariable a is chosen to relate the function pre- and postconditions to each other over the call. Consider the pair

$$\begin{aligned} \mathbb{P}_\ell &\equiv \lambda(x, v). (r_0 \Downarrow 4) \wedge ((x \neq r_0) \Rightarrow x \Downarrow v) \\ \mathbb{Q}_\ell &\equiv \lambda(x, v). (r_0 \Downarrow 8) \wedge ((x \neq r_0) \Rightarrow x \Downarrow v) \end{aligned}$$

If the caller needs to be sure that the invariant $r_{15} \Downarrow (16, (23, 42))$ is preserved over the function then he sets $a = (r_{15}, (16, (23, 42)))$. The key point of the HCall rule is that if we satisfy P then we can verify a function pointer call with a bound of $n + 1$ calls.

Precondition generator. Our update rules are in weakest-precondition style and our predicates include general quantification. Our Coq development defines a precondition generator that computes P from R , n , c , and Q . We use such a generator to cut down on the tedium of mechanically verifying the example programs presented in §4.

3.2 Function Verification

The whole-function verification rules given in Figure 4 form the heart of our program logic. Although the symbol count is daunting, the core idea is natural.

Functions are normally verified one at a time, although mutually recursive function groups are verified as a set. One begins with rule Vstart, which says that program Ψ has specification \top (*i.e.*, that none of the functions in Ψ have been verified to terminate). The Vsimple and Vfull rules verify the addition of terminating function specifications into the context Γ . Vsimple is sufficient to handle simple recursive functions that take non-polymorphic function pointers as arguments. Vfull handles mutually recursive function groups and polymorphic function pointers; Vsimple is just a special case of Vfull. After verifying the first function/group, one continues with the next function/group with another application of Vsimple/Vfull until all of Ψ has been verified.

The Vsimple rule assumes that the program Ψ already has specification Γ ; we wish to add the specification for the function at ℓ using termination measure t , precondition \mathbb{P} , and postcondition \mathbb{Q} . The main requirement is the second premise. We must verify, using the H-rules, that for any n and a , the function body $\Psi(\ell)$ meets the specification

$$\dots, \mathbb{Q}(a) \vdash_n \{ \mathbb{P}(a) \wedge \langle t \rangle = n \} \Psi(\ell) \{ \perp \}$$

That is, starting from a state that satisfies the precondition $\mathbb{P}(a)$ and in which the termination measure t evaluates to n , the function will `return` in a state satisfying $\mathbb{Q}(a)$ after having made no more than n function calls. We use \perp as the postcondition since the function is not allowed to “fall off the bottom”. The key to doing recursive functions is how we set up the function specifications: we verify $\Psi(\ell)$ using the previously-verified function specifications in Γ as well as a modified specification for ℓ itself:

$$\text{funptr } \ell t \ [\lambda a'. \mathbb{P}(a') \wedge \langle t \rangle < n] \ [\mathbb{Q}]$$

That is, the function body $\Psi(\ell)$ is can call to other functions specified in Γ as well as recursive calls to itself *as long as the termination measure decreases*.

The Vfull rule is more general, and correspondingly, more complex. It generalizes the Vsimple rule in two orthogonal ways. First, Vfull is able to verify a mutually recursive set of functions. Second, Vfull is able to verify function specifications where the specifications take *parameters*. The universally-quantified variable b in the Vfull rule represents the specification parameters; b can range over an arbitrary type chosen by the verifier. The variable Φ appearing in the Vfull rule represents a finite set of function specifications, i.e., a set of tuples with a label, a termination measure and a pre- and postcondition. The specifications in Φ represent the set of mutually recursive functions we are going to verify. The quantification over $\Phi(b)$ in the premise of the rule means that we will have to construct a Hoare derivation for each function body represented in Φ . Correspondingly, the quantification in the conclusion means that subsequent verifications may rely on each of the specifications in Φ . In other words, the Vfull rule establishes the specifications of a set of mutually recursive functions simultaneously.

Note that Φ takes an argument; thus the function specifications can depend on the parameter b .³ In the premise of the Vfull rule, the value b is bound *once* and the same b is used to construct the recursive assumptions as is used to construct the verification obligations. In other words, the value of the parameter, b , is a constant throughout the recursion. Contrast this with the value a which connects pre- and postconditions, which is allowed to vary at each recursive call. An interesting case occurs when b is allowed to range over function specifications. In this case, the specifications in Φ take on a *higher-order* flavor. We shall see several examples using this power in the following section.

4 Examples of Verified Programming in HAL

Our logic has three distinctive features: time bounds on recursive function pointers; time bounds on polymorphic function pointers; and semantic assertions. Here we cover two examples, the first demonstrating recursive function pointers, and the second demonstrating polymorphic function pointers. The second example demonstrates the use of a semantic assertion whose truth cannot be practically checked at run-time. In our Coq development we have examples that combine both simultaneously; also see [DH10]. Although HAL is Turing complete, it lacks many features enabling abstraction and ease-of-use. The simplicity of our examples does not indicate a theoretical limitation, but is rather a practical consequence of programming in such an impoverished language.

Example 1: unary addition. Here we examine a simple recursive function which “adds” two lists representing natural numbers in unary notation (lists terminated by the 0 label). The idea is that starting from two unary-encoded naturals in registers r_1 and r_2 , we strip cons cells from r_1 and add them to r_2 until there are no cells left in r_1 , and then return.

To state the specification for the addition function, we need a predicate listnat that relates natural numbers to their unary encoding. Essentially, the number of nested cons cells captures the intended natural number; see figure 5a. We give the specification itself in figure 5c. Note that the pre- and postconditions of the addition function are parameterized by the pair of numbers to be added. We also define the termination measure. Recall that we allow termination measures to be partial functions; we use that power here because addt is only defined when the value in r_1 encodes some natural number.

³ This may be quite a strong dependency; even the type of a which connects the pre- and postconditions can depend on the value of b .

Fig. 5a. Encoding naturals		Fig. 5b. Code c_{add} , loaded at label 1	
$\text{listnat}(0, 0)$ $\text{listnat}(n, v) \rightarrow$ $\quad \text{listnat}(n + 1, (0, v))$		1 assert $(\exists n m. \text{addP}(n, m))$; 2 ifnil r_1 then return ; 3 else 4 $r_3 := r_1.1$; 5 $r_1 := r_1.2$; 6 $r_2 := (r_3, r_2)$; 7 $r_0 := 1$; // address of c_{add} 8 call r_0 ; 9 return ;	
Number	Encoding		
0	0		
1	(0, 0)		
2	(0, (0, 0))		
3	(0, (0, (0, 0)))		

Fig. 5c. Pre- and postcondition; termination measure

$\text{addP}(n, m)$	\equiv	$\exists v_1 v_2. r_1 \Downarrow v_1 \wedge r_2 \Downarrow v_2 \wedge \text{listnat}(n, v_1) \wedge \text{listnat}(m, v_2)$
$\text{addQ}(n, m)$	\equiv	$\exists v_2. r_2 \Downarrow v_2 \wedge \text{listnat}(n + m, v_2)$
$\text{addt}(\rho)$	\equiv	the unique n s.t. $\exists v_1. \rho(r_1) = v_1 \wedge \text{listnat}(n, v_1)$

Fig. 5d. Verification obligation for unary addition (using Vsimple)

$\forall n_1, m_1, n.$	$(\text{funptr } 1 \text{ addt } [\lambda n_2 m_2. \text{addP}(n_2, m_2) \wedge \langle \text{addt} \rangle < n] [\text{addQ}],$
$\text{addQ}(n_1, m_1)$	$\vdash_n \{ \text{addP}(n_1, m_1) \wedge \langle \text{addt} \rangle = n \} \langle \text{code from fig. 5b} \rangle \{ \perp \})$

Fig. 5. Example 1: unary addition.

The code is given in figure 5b. Line 1 simply asserts the precondition of the function. Line 2 tests if the value in register r_1 is nil; if so, we return. Otherwise, we perform one unit of work, which involves shifting one cons cell from r_1 to r_2 . Note lines 7 and 8, where we load the constant label 1 into r_0 and jump to it; this sequence is typical of “static” function calls. Since the code itself is loaded at label 1, this is a recursive call.

The addition function is a simple self-recursive function, so we can to verify it according to its specification using the Vsimple rule. The proof obligation that is generated by Vsimple (after some minor simplifications) is shown in Figure 5, and it is straightforward to use the H-rules of our logic to fulfill this verification obligation proof. In the base case, when r_1 contains the nil value, $n_1 = 0$ and the postcondition follows immediately. In the recursive case, lines 4–6 do the interesting work of shifting a single cons cell, and moves from a state satisfying $\text{addP}(n_1 + 1, m_1)$ to a state satisfying $\text{addP}(n_1, m_1 + 1)$. The recursive call (lines 7–8) thus has its precondition satisfied, and the termination measure (which tracks n_1) has decreased. The final proof obligation is then to show that $\text{addQ}(n_1, m_1 + 1)$ implies $\text{addQ}(n_1 + 1, m_1)$.

Example 2: apply. While the code for the “apply” function is dead simple, the specification is rather subtle. The “apply” function makes essential use of function pointers and thus has a higher-order specification. The basic idea is that one packages together a function label with some additional arguments using a cons cell in register r_0 . Apply unpacks the cons cell and calls the contained function using with the enclosed arguments. We toss in an interesting higher-order **assert** just before the call for fun.

Fig. 6a. “standard” functions; precondition, postcondition, termination measure for apply

$$\begin{aligned}
\text{csregs}(v_1, v_2, v_3, v_4) &\equiv (r_1 \Downarrow v_1) \wedge (r_2 \Downarrow v_2) \wedge (r_3 \Downarrow v_3) \wedge (r_4 \Downarrow v_4) \\
\text{stdfun}(\ell, t_\ell, P_\ell, Q_\ell) &\equiv \text{funptr } \ell (\lambda \rho. t_\ell(\rho(r_0))) \\
&\quad \left[\lambda(v_1, v_2, v_3, v_4, a). (r_0 \Downarrow v_0) \wedge P_\ell(a)(v_0) \wedge \text{csregs}(v_1, v_2, v_3, v_4) \right] \\
&\quad \left[\lambda(v_1, v_2, v_3, v_4, a). (r_0 \Downarrow v_0) \wedge Q_\ell(a)(v_0) \wedge \text{csregs}(v_1, v_2, v_3, v_4) \right] \\
\text{applyP}(t, P, Q)(a)(v) &\equiv \exists \ell v_2. v = (\ell, v_2) \wedge \text{stdfun}(\ell, t, P, Q) \wedge P(a)(v_2) \\
\text{applyQ}(Q)(a)(v) &\equiv Q(a)(v) \\
\text{applyt}(t)(v) &\equiv t(v) + 1
\end{aligned}$$

Fig. 6b. code c_{app} for apply with abbreviated Hoare triples, loaded at label 1

$$\begin{array}{l}
R \equiv \exists v'_0. r_0 \Downarrow v'_0 \wedge Q(a)(v'_0) \wedge \text{csregs}(v_1, v_2, v_3, v_4) \\
\begin{array}{l}
\vdash_0 \{r_0 \Downarrow (\ell, v) \wedge P(a)(v) \wedge t(v) + 1 = n \wedge \text{stdfun}(\ell, t, P, Q) \wedge \text{csregs}(v_1, v_2, v_3, v_4)\} \\
1 \quad r_5 := r_0.0; \\
\vdash_0 \{r_0 \Downarrow (\ell, v) \wedge r_5 \Downarrow \ell \wedge P(a)(v) \wedge t(v) + 1 = n \wedge \text{stdfun}(\ell, t, P, Q)\} \\
2 \quad r_0 := r_0.1; \\
\vdash_0 \{(r_0 \Downarrow v) \wedge (r_5 \Downarrow \ell) \wedge P(a)(v) \wedge (t(v) + 1 = n) \wedge \text{stdfun}(\ell, t, P, Q)\} \\
3 \quad \text{assert } (\exists \ell', t', P', Q'. (r_5 \Downarrow \ell') \wedge \text{stdfun}(\ell', t', P', Q')) ; \\
\vdash_0 \{(r_0 \Downarrow v) \wedge (r_5 \Downarrow \ell) \wedge P(a)(v) \wedge (t(v) + 1 = n) \wedge \text{stdfun}(\ell, t, P, Q)\} \\
4 \quad \text{call } r_5 ; \\
\vdash_n \{r_0 \Downarrow v' \wedge Q(a)(v')\} \\
5 \quad \text{return} ; \\
\vdash_n \{\perp\}
\end{array}
\end{array}$$

Fig. 6c. Verification obligation for apply (using Vfull, after simplification)

$$\begin{aligned}
&\forall (t, P, Q) (v_1, v_2, v_3, v_4, a) n. \left(\top, R \vdash_n \right. \\
&\quad \left. \{ \exists \ell, v_0. r_0 \Downarrow (\ell, v_0) \wedge \text{stdfun}(\ell, t, P, Q) \wedge P(a)(v_0) \wedge \text{csregs}(v_1, v_2, v_3, v_4) \wedge t(v_0) + 1 = n \} \right. \\
&\quad \left. (\text{code from fig. 6b}) \{ \perp \} \right)
\end{aligned}$$

Fig. 6. Example 2: apply

In order to give a reasonable specification for this function and other higher-order operations, we identify a convenient calling convention. We call functions that adhere to our calling convention “standard” functions. Register r_0 is used for passing function arguments and results. Registers r_1 – r_4 are callee-saves registers (whose values must be preserved over the call) and all other registers are caller-saves. In addition, we require the precondition, postcondition, and termination measure, for standard functions, to be defined only on the argument/return value (the value in r_0). We say a function satisfies $\text{stdfun}(\ell, t, P, Q)$ (where t, P and Q are defined over a single value rather than an entire store) if ℓ is a standard function in the sense just defined.

In the specification for apply (Figure 6) t, P and Q are the *parameters of the specification*; they describe the function that will be called. When it comes time to verify the apply function, we will use the Vfull rule and b will range over tuples (t, P, Q) . This way we can specify and prove correct the apply function in complete isolation, without requiring any static assumptions about the functions it will be passed. **The time bound n in the verification comes from the bound on the input function.** In some later verification, the specification of apply can be instantiated with any function specification

that the verifier knows about. In particular, `apply` can be applied to itself! This would not be a recursive call, in the traditional sense, but is rather a *dynamic* higher-order call.

Termination remains assured because the way the specifications get “stacked” on top of each other. This stacking of function specifications, in general, creates a tree-like structure where the leaves of the tree must be first-order functions (whose specifications do not depend on the specifications of other functions). The whole thing hangs together because there is no way to create a cycle in the tree of function specifications, and thus no way to introduce new, potentially nonterminating, recursion patterns. See the formal development for an example of such “stacked” function applications.

The development also contains an implementation of and verification for the recursive higher-order function `map`, whose termination argument is the sum of the input function’s termination argument applied to each element of the list plus the length of the list (for the recursive calls of `map` itself); for more details see [DH10].

5 Resolving the Circularity in predicates

We have presented a program logic for HAL and demonstrated its use on two examples; we are now ready to move on to the semantics. We divide the task into two parts. In this section we resolve the circularity in predicate—that is, for the assertions of our Hoare tuples—from Figure 2. In §6 we build a model for the program logic itself and prove that programs verified in our logic terminate within the correct time bound.

5.1 Using Indirection Theory to Stratify Through Syntax

The pseudomodel of predicates in Figure 2 fits into the pattern $K \approx F((K \times O) \rightarrow \mathcal{T})$. In this pattern, F is a covariant functor, O is some kind of “flat data”, and K is an object one wishes to model. A cardinality argument shows that there are no solutions in set theory to this pseudoequation, so we instead build an approximate model using indirection theory [HDA10b]. In our case, F is the parameterized program ϕ from Figure 1 and O is just store. Indirection theory “ties the knot” and defines K such that:

$$\begin{array}{llll} K & & \equiv & [\text{ADH10, knot_hered.v}] \\ \text{sq_program} & \check{\Psi} & \equiv & K \\ \text{state} & \sigma & \equiv & \text{sq_program} \times \text{store} \\ \text{predicate} & P & \equiv & \{P : \text{state} \rightarrow \mathcal{T} \mid \text{hereditary}(P)\} \end{array}$$

The construction of the *knot* K is similar to the one given in [HDA10b, §8] but we have enhanced it so that all predicates inside the knot are *hereditary*, a technical property detailed in §5.2. A *squashed program* `sq_program` is simply a knot; a state is a pair of a `sq_program` and a store. A predicate is a hereditary function from states to truth values \mathcal{T} . We write $\sigma \models P$ instead of $P(\sigma)$ when we wish to emphasize that we are thinking of P as an assertion as opposed to a function. What is missing is the relationship between squashed and unsquashed programs. We cannot have an isomorphism or we run into cardinality problems; instead, we get $\text{sq_program} \preceq \mathbb{N} \times \text{program}$, where $X \preceq Y$ means that the “small” type X and the “big” type Y are related by a section/retraction pair witnessed by two functions: $\text{squash} : (\mathbb{N} \times \text{program}) \rightarrow \text{sq_program}$ and $\text{unsquash} : \text{sq_program} \rightarrow (\mathbb{N} \times \text{program})$. The power of indirection theory is that there is a very simple (indeed, categorical) pair of axioms relating `squash` and `unsquash`:

$$\begin{array}{ll} \text{squash}(\text{unsquash}(\check{\Psi})) & = \check{\Psi} \\ \text{unsquash}(\text{squash}(n, \Psi)) & = (n, \text{prog_approx}_n(\Psi)) \end{array} \quad (1)$$

That is, $\text{squash} \circ \text{unsquash}$ is the identity function, and $\text{unsquash} \circ \text{squash}$ is a kind of approximation function. The $\text{prog_approx}_n(\Psi)$ function transforms Ψ by locating all of the $\text{assert}(P)$ statements and replacing them with $\text{assert}(\text{approx}_n(P))$. The core of the approximation is handled by the $\text{approx}_n(P)$ function:

$$\begin{aligned} |\check{\Psi}| &\equiv (\text{unsquash}(\check{\Psi})).1 \\ \text{approx}_n(P) &\equiv \lambda(\check{\Psi}, \rho). \begin{cases} P(\check{\Psi}, \rho) & |\check{\Psi}| < n \\ \perp & |\check{\Psi}| \geq n \end{cases} \end{aligned} \quad (2)$$

First we define the *level* of a squashed program $\check{\Psi}$, written $|\check{\Psi}|$, as the first projection of $\check{\Psi}$'s unsquashing. When a predicate is approximated to level n , its behavior on programs of level **strictly** less than n is unchanged; on all other input it now returns the constant \perp . The approx function is exactly where step-indexed models get both their power (a sound construction) and weakness (a loss of information during approximation).

5.2 Consequences of Approximation

What is the cost of throwing away information during approximation? Ten years after step-indexed models were introduced, the answer is still unclear. Considerable experience has led to an ad-hoc understanding among practitioners of what might be termed *microcosts*—that is, small modifications to the “intuitive” definitions to accommodate the approximation. A large body of previous work has focused on managing and minimizing these microcosts, *e.g.*, via a Gödel-Löb logic of approximation [Ric10].

The fundamental microcost comes from the fact that the approx_n function throws away all behavior on squashed programs of greater than **or equal to** level n . Let P be a predicate contained in (the unsquashing of) a squashed program $\check{\Psi}$ of level n . A consequence of (1) is that P has been approximated to level n —*i.e.*, $P = \text{approx}_n(P)$. What happens if we apply P to a state containing $\check{\Psi}$ itself? A review of (2) proves that the result must be \perp . **A predicate cannot say anything meaningful about the squashed program whence it came.** Instead, we will do the next best thing: make $\check{\Psi}$ a little simpler. We say that $\check{\Psi}$ (or σ) is *approximated* to $\check{\Psi}'$ (or σ'), written $\check{\Psi} \rightsquigarrow \check{\Psi}'$, when

$$\begin{aligned} \check{\Psi} \rightsquigarrow \check{\Psi}' &\equiv \text{let } (n, \Psi) = \text{unsquash}(\check{\Psi}) \text{ in } ((n > 1) \wedge (\check{\Psi}' = \text{squash}(n-1, \Psi))) \\ (\check{\Psi}, \rho) \rightsquigarrow (\check{\Psi}', \rho') &\equiv (\rho = \rho') \wedge (\check{\Psi} \rightsquigarrow \check{\Psi}') \end{aligned} \quad (3)$$

That is, we unsquash $\check{\Psi}$ and then re-squash it to one level lower. Of course, we can only do this when we are not at level 0 to begin with! Since $|\check{\Psi}'| = n-1 < n$, P will be able to judge states containing $\check{\Psi}'$. Every time we pull a predicate out of a squashed program $\check{\Psi}$, we will approximate $\check{\Psi}$ to $\check{\Psi}'$ before we use P .

Repeated approximation leads to a second microcost. Suppose $\sigma \models P$ and $\sigma \rightsquigarrow \sigma'$. We say P is *hereditary*—stable (or monotonic) as σ is approximated—so that $\sigma' \models P$; that is, $\text{hereditary}(P) \equiv \forall \sigma. (\sigma \models P) \rightarrow (\sigma \rightsquigarrow^* \sigma') \rightarrow (\sigma' \models P)$, where we write \rightsquigarrow^* and \rightsquigarrow^+ to mean the reflexive and irreflexive transitive closures, respectively, of \rightsquigarrow . Unfortunately, not all functions from state to \mathcal{T} are hereditary, such as $P_{\text{bad}}(\check{\Psi}, \rho) \equiv |\check{\Psi}| > 5$. The P_{bad} function will be true only while the level of the program is greater than 5; since approximating the state decreases that level, eventually this function will produce only the constant \perp . In our setting, we only consider

predicates that are hereditary. Whenever we give the semantics of a predicate (except for P_{bad} !), we have proved (in Coq) that the definition is hereditary. For most definitions this is not a big deal; however on occasion there can be a significant amount of work that needs to be done. The benefit of only allowing hereditary predicates is that once the definitions are done they are easier to use in the core soundness proofs.

A central question is how these kinds of microcosts become macrocosts—that is, what are the fundamental limitations of step indexing techniques? For some time, it was thought that step-indexed models could not produce the kinds of existential witnesses needed for termination proofs; however, the present work proves otherwise. The practical limitations of step-indexed models remain unknown.

5.3 Models for predicates

We define the standard logical connectives for predicates ($\top, \perp, \wedge, \vee, \Rightarrow, \neg, \forall, \exists, \mu$), as well as entailment ($P \vdash Q$) by a standard intuitionistic lift over the \rightsquigarrow^* relation from the metalogic, *e.g.*, $(\sigma \models P \wedge Q) \equiv (\sigma \models P) \wedge (\sigma \models Q)$ [HDA10a]. As in [HDA10a], we define the modality of approximation ($\triangleright P$), used in the metatheory but not by the end-user, as the boxy operator over \rightsquigarrow^+ : $\sigma \models \triangleright P \equiv \forall \sigma'. \sigma \rightsquigarrow^+ \sigma' \rightarrow \sigma' \models P$. The model for the terminating function pointer assertion `funptr` is complex and is developed in §6; the other domain-specific predicates listed in Figure 2b are simply as follows:

$$\begin{array}{llll}
 (\check{\Psi}, \rho) & \models & x \Downarrow v & \equiv & \rho(x) = v \\
 (\check{\Psi}, \rho) & \models & [x \leftarrow v]P & \equiv & (\check{\Psi}, [x \leftarrow v]\rho) \models P \\
 (\check{\Psi}, \rho) & \models & \text{closed } P & \equiv & \forall \rho'. (\check{\Psi}, \rho') \models P \\
 (\check{\Psi}, \rho) & \models & \langle t \rangle < n & \equiv & t(\rho) < n \quad (\text{etc. e.g., for } \langle t \rangle = n)
 \end{array}$$

For further discussion on the models for predicates see [DH10].

6 A Step-indexed Model for Total Correctness

Soundness of a logic of total correctness (w.r.t. its operational semantics) means that whenever a function in a verified program is run in an initial state satisfying its precondition, it will halt in a state satisfying its postcondition. Our soundness proof follows the program of Appel and Blazy [AB07] which involves: building a semantic model for assertions; defining the meaning of judgments; proving the inference rules of the logic as lemmas; and showing that the judgment semantics implies the desired theorem.

In §5.1 we defined the model for predicates. Because predicates appear in syntax, we needed to do this before defining the operational semantics of the language.

Operational semantics. Most of the operational semantics of our language, involving simple data or control-flow, is straightforward. Here, we only highlight the more interesting portions of the operational semantics; for further details see [DH10].

We define a small-step relation $(\check{\Psi}, \rho, s) \mapsto (\check{\Psi}', \rho', s')$, in which ρ and s stand for local variables and stacks and $\check{\Psi}$ is the program text; it is actually modified as the program runs in a very controlled way. Here is the simple rule for loading a label:

$$\frac{}{(\check{\Psi}, \rho, (x := \ell; c) :: s) \mapsto (\check{\Psi}, [x \leftarrow \ell]\rho, c :: s)} \text{Slabel}$$

Note that this rule passes the program $\check{\Psi}$ through without any change. Most of the instructions of the language have this property. Figure 7 lists the two rules of particular

$$\frac{(\tilde{\Psi}, \rho) \models P}{(\tilde{\Psi}, \rho, (\text{assert } P; c) :: s) \mapsto (\tilde{\Psi}, \rho, c :: s)} \text{Sassert}$$

$$\frac{\rho(x) = \ell \quad \text{unsquash}(\tilde{\Psi}) = (n, \Psi) \quad \Psi(\ell) = c' \quad \tilde{\Psi} \rightsquigarrow \tilde{\Psi}'}{(\tilde{\Psi}, \rho, (\text{call } x; c) :: s) \mapsto (\tilde{\Psi}', \rho, (c'; \text{assert } \perp) :: c :: s)} \text{Scall}$$

Fig. 7. Operational Semantics

$$\begin{aligned} \text{configHalts}_n(\tilde{\Psi}, \rho, s) &\equiv \exists \tilde{\Psi}', \rho'. (|\tilde{\Psi}| - |\tilde{\Psi}'| \leq n) \wedge (\tilde{\Psi}, \rho, s) \mapsto^* (\tilde{\Psi}', \rho', \text{nil}) \\ (\tilde{\Psi}, \rho) \models \text{halts}_n s &\equiv \boxed{|\tilde{\Psi}| \geq n \rightarrow \text{configHalts}_n(\tilde{\Psi}, \rho, s)} \\ \text{guards}_n P s &\equiv P \Rightarrow \text{halts}_n s \\ \tilde{\Psi} \models \text{funptr } \ell t [\mathbb{P}] [\mathbb{Q}] &\equiv \exists c. \text{let } (n_\Psi, \Psi) = \text{unsquash}(\tilde{\Psi}) \text{ in} \\ &\quad \Psi(\ell) = c \wedge \forall s, \tilde{\Psi}', n', a. \tilde{\Psi} \rightsquigarrow^+ \tilde{\Psi}' \rightarrow (\forall \rho. (\tilde{\Psi}', \rho) \models \text{guards}_{n'} \mathbb{Q}(a) s) \rightarrow \\ &\quad (\forall \rho n. t(\rho) = n \rightarrow (\tilde{\Psi}', \rho) \models \text{guards}_{n+n'} \mathbb{P}(a) ((c; \text{assert } \perp) :: s)) \\ \Gamma, R \vdash_n \{P\} c \{Q\} &\equiv \forall \tilde{\Psi}, n', k, s. \tilde{\Psi} \models \Gamma \rightarrow (\forall \rho. (\tilde{\Psi}, \rho) \models \text{guards}_{n'} R s) \rightarrow \\ &\quad (\forall \rho. (\tilde{\Psi}, \rho) \models \text{guards}_{n'} Q (k :: s)) \rightarrow (\forall \rho. (\tilde{\Psi}, \rho) \models \text{guards}_{n+n'} P ((c; k) :: s)) \\ \tilde{\Psi}' \models \text{approxodof}(\tilde{\Psi}) &\equiv \tilde{\Psi} \rightsquigarrow^* \tilde{\Psi}' \\ \Psi : \Gamma &\equiv \forall n. ((\text{approxodof}(\text{squash}(n, \Psi)) \wedge \triangleright \Gamma) \vdash \Gamma) \end{aligned}$$

Fig. 8. Terminating function pointers; Hoare tuples; whole-program verification

interest. The first is the Sassert rule, which shows how semantic assertions are checked as the program runs. The second is Scall, which shows what happens at function calls. This is the only rule where the program is modified; what happens is that each assertion in the program text is approximated one level down. We must do this approximation so that assertions in the text of the function body will be able to judge the program. If we did not approximate the program at this point, any assertions in the function body would fail, foiling our desired soundness result. This means that the level of the program $\tilde{\Psi}$ is an upper bound on the number of calls the program can make before getting stuck.

Judgment Definitions. Appel and Blazy build their semantic Hoare triple using the more basic notion of guarding. They say that a predicate “guards” a program stack if, whenever a memory state satisfies the predicate, that stack is safe to run (i.e., will not go wrong). We follow a similar pattern, but use a guards predicate which enforces termination rather than safety. We say that a predicate P guards a stack s at level n , whenever the memory state satisfies P and provided that the program level is at least n , running the stack will eventually terminate. See Figure 8. Notice that there is a clever trick being played here with the definition of halts_n . Halting is not normally a predicate which can be hereditary. As one ages a program, it is able to run for fewer steps and thus might not terminate before it exhausts its level. We work around this issue by saying that a program must only terminate if it has at least level n . As one ages a program, it will eventually cause halts_n to be true vacuously (when its level falls below n).

We use guards in the terminating function pointer assertion, also in Figure 8. Here, ℓ is a program label, t is a *termination measure* (a partial function from stores to \mathbb{N}), and \mathbb{P} and \mathbb{Q} are functions from some type A to assertions. This definition is given in a continuation-oriented style. Whenever we have a stack s which terminates in n

steps when $\mathbb{Q}(x)$ is satisfied, then we know that running the function body of ℓ will terminate in $n + n'$ steps whenever $\mathbb{P}(x)$ is satisfied, and where n' is determined by the measure. Thus, funptr captures the specification of a terminating function. Note the premise $\check{\Psi} \rightsquigarrow^+ \check{\Psi}'$; this is one of the microcosts discussed in §5.2. $\check{\Psi}'$ must be strictly more approximate than $\check{\Psi}$ because stepping over a `call` instruction ages the program.

Both the funptr predicate and the Hoare judgment have a similar flavor: assume the postcondition(s) guard the program continuation point(s) and demonstrate that the precondition guards the extended continuation. This is no accident, as these two definitions interact in interesting ways. First, a funptr is a premise of the Hcall rule and directly provides the required guards fact to satisfy the Hoare judgment. Secondly, The Vfull rule shows how we can use the Hoare judgment to establish funptr facts.

The final definition we need is that of program verification. What this definition means is that we can prove Γ provided that we assume the program under consideration is some squashed version of $\check{\Psi}$ and $\triangleright \Gamma$ (i.e., approximately Γ). The assumption $\triangleright \Gamma$ plays the role of an induction hypothesis and is what allows us to verify recursive functions. The approxedof($\check{\Psi}$) predicate means that the current program is approximated from $\check{\Psi}$.

H- and V-rules. Now that we have finished our semantic definitions, we are prepared to prove the rules of the Hoare logic as lemmas. The proofs are quite straightforward for all of the rules aside from Hcall, which itself is not arduous; for details see [DH10].

The real “magic,” such as it is, happens in the proof of the function verification rule, Vfull. This rule converts Hoare derivations for function bodies into the corresponding funptr assertions on programs containing those function bodies. One supplies Φ , a list containing the precondition, postcondition and termination measure for a group of (potentially) mutually recursive functions. Then, for each function in Φ , one must prove a Hoare derivation of a specific form. The assumptions one is allowed to make are taken from Γ , which contains functions already verified, and from Φ , which allows recursive calls. However, the preconditions in Φ are altered to add a conjunct which strengthens the preconditions by requiring the termination measure to decrease. The return postcondition is, naturally, the postcondition of the function. The precondition is the ordinary function precondition together with the assumption that the termination measure for the initial state is n ; this is what connects the strengthened preconditions of the recursive assumptions with the initial state. The linear postcondition is \perp , which requires the function body to return rather than “falling of the end.” Finally, the Hoare derivation is required to bound the number of function calls by n . This connects the termination measures of the function specifications to their intended semantic meanings.

Whenever one can provide such a Hoare derivation for each function in Φ , then one can conclude that each function referenced in Φ actually respects its contract, and the corresponding funptr facts can be conjoined with Γ in the conclusion of the rule.

The soundness proof for the rule is a little delicate, but the main idea is straightforward; the proof goes by (complete) induction on the value of the termination measure. The Hoare derivation is used to complete the argument, using the induction hypothesis to discharge the premises in the definition of the Hoare judgment.

Total correctness. The final part of the soundness proof is a result that connects the results we obtain using our definitions to a more traditional notion of total correctness.

Theorem 6.1 (Total Correctness). *Suppose $\Psi : \Gamma$, and $\Gamma \vdash \text{funptr } \ell \ t \ [\mathbb{P}] \ [\mathbb{Q}]$. Then for all stores ρ such that $t(\rho) = n$, and $(\text{squash}(n, \Psi), \rho)$ satisfies $\mathbb{P}(a)$ (for some a), executing the function body $\Psi(\ell)$ will terminate in a state satisfying $\mathbb{Q}(a)$.*

Our core semantic ideas (i.e., §6) are compact, requiring only 1,315 lines of Coq; the development may be found at <http://msl.cs.princeton.edu/termination/>.

7 Limitations and Extensions

To highlight the interaction of function pointers and semantic assertions, we have pared down HAL to the bare minimum. We do not consider many other interesting language features such as looping constructs, heap manipulations, or procedure parameters.

Our program logic is somehow simultaneously a bit too weak and a bit too strong. It is too weak in that the upper bound need not be tight, and we make no claims on the lower bound. Our logic is too strong in that the burden of constructing an explicit termination measure may be onerous for someone only concerned with termination. It would be better if one could provide a well-founded relation for each function, hiding the explicit bounds and termination measures under existential binders.

We can relate the precondition to the upper bound so that we can verify, e.g., that a program runs in polynomial bound (e.g., examine the termination measure addt from §4). However, the mechanism to do so is cumbersome; it might be better for this purpose to allow the user to state termination arguments in the ordinals instead of the naturals.

Finally, we have not investigated the completeness of our logic, but have only argued that the logic is *complete enough* to verify interesting programs. Future work may investigate completeness and determine what modifications may be required.

8 Related work

Applications of step-indexing and its alternatives. Step-indexing has been used to prove results in type safety [Ahm04], soundness of program logics [Hob08], and program equivalence [ADR09]. Indirection theory [HDA10b] provides clean axioms for step-indexed models. Domain theory is the classic tool for building semantic models [GHK⁺03]. The price one pays for working in domain theory is the substantial mathematical theory which underlies the discipline. Birkedal *et al.* show how to construct indirection theory in certain categories of ultrametric spaces [BRS⁺].

Predicates in syntax. Semantic assertions are often used in program analysis settings such as BoogiePL [BCD⁺05]. Semantic assertions are one example of a larger class of bookkeeping instructions that embed formulas into program syntax, such as the `makeLock` instruction used in concurrent C minor [Hob08].

Program logics with function pointers. Schwinghammer *et al.*'s recent work on “nested” Hoare triples [SBRY09] combines features of separation logic [Rey02] with the ability to reason about “stored code,” which behaves quite similar to function pointers. It is a logic of partial correctness. The work of Honda *et al.* seems nearest to our own in terms of logical power [HY04]. They provide a logic of total correctness for call-by-value PCF (a λ -calculus with references). The soundness proof goes by a reduction to the π -calculus equipped with a process logic in the rely/guarantee style [Hon04]. Honda *et al.* do not consider embedded semantic assertions or explicit time bounds, but do

consider the issue of completeness [HBY06]. Aspinall et al. have developed a sound and complete program logic for Grail, a Java subset, which reasons about both correctness and resources [ABH⁺07]. Although their system includes a form of virtual method invocation, it is not immediately clear if their formalism allows higher-order behaviors.

9 Conclusion

We have presented a simple language with embedded semantic assertions and function pointers, together with a logic of the total correctness of time resource bounds. Our logic is able to reason about terminating function pointers in a very general way, including polymorphic mutually-recursive function groups. We have proved our logic sound with respect to an operational semantics using step-indexing, thereby demonstrating that step-indexed models are useful for modeling resource logics.

References

- [AB07] Andrew W. Appel and Sandrine Blazy. Separation logic for small-step C minor. In *TPHOLS*, pages 5–21, 2007.
- [ABH⁺07] David Aspinall, Lennart Beringer, Martin Hofmann, Hans-Wolfgang Loidl, and Alberto Momigliano. A program logic for resources. *TCS*, 389(3):411–445, 2007.
- [ADH10] Andrew Appel, Robert Dockins, and Aquinas Hobor. Mechanized Semantic Library. Available at <http://msl.cs.princeton.edu>, 2009–2010.
- [ADR09] Amal Ahmed, Derek Dreyer, and Andreas Rossberg. State-dependent representation independence. In *POPL '09*, pages 340–353, 2009.
- [Ahm04] Amal J. Ahmed. *Semantics of Types for Mutable State*. PhD thesis, Princeton University, Princeton, NJ, November 2004. Tech Report TR-713-04.
- [BCD⁺05] M. Barnett, B. E. Chang, R. DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *FMCO05*, 2005.
- [BRS⁺] L. Birkedal, B. Reus, J. Schwinghammer, K. Støvring, J. Thamsborg, and H. Yang. Step-indexed Kripke models over recursive worlds. Submitted for publication.
- [DH10] Robert Dockins and Aquinas Hobor. A theory of termination via indirection. In *Modelling, Controlling and Reasoning About State*, Dagstuhl Seminar 10351, 2010.
- [GHK⁺03] G. Gierz, K. H. Hofmann, K. Kiemel, J. D. Lawson, M. Mislove, and D. S. Scott. *Continuous Lattices and Domains*, volume 93 of *Encyclopedia of Mathematics and its Applications*. Cambridge U. Press, Cambridge, UK, 2003.
- [HBY06] K. Honda, M. Berger, and N. Yoshida. Descriptive and relative completeness of logics for higher-order functions. In *ALP*, LNCS, pages 360–371, 2006.
- [HDA10a] Aquinas Hobor, Robert Dockins, and Andrew W. Appel. A logical mix of approximation and separation. In *APLAS10*, pages 439–454. Springer ENTCS, 2010.
- [HDA10b] Aquinas Hobor, Robert Dockins, and Andrew W. Appel. A theory of indirection via approximation. In *POPL'10*, pages 171–185, January 2010.
- [Hob08] Aquinas Hobor. *Oracle Semantics*. PhD thesis, Princeton University, 2008.
- [Hon04] Kohei Honda. From process logic to program logic. In *ICFP '04: Proc. of the 9th intl. conference on Functional programming*, pages 163–174, 2004.
- [HY04] Kohei Honda and Nobuko Yoshida. A compositional logic for polymorphic higher-order functions. In *PPDP '04*, pages 191–202, 2004.
- [Rey02] John Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS 2002: IEEE Symposium on Logic in Computer Science*, pages 55–74, July 2002.
- [Ric10] Christopher D. Richards. *The Approximation Modality in Models of Higher-Order Types*. PhD thesis, Princeton University, Princeton, NJ, June 2010.
- [SBRY09] J. Schwinghammer, L. Birkedal, B. Reus, and H. Yang. Nested Hoare triples and frame rules for higher-order store. In *CSL'09/EACSL'09*, pages 440–454, 2009.