Extended Static Checking of Call-by-Value Functional Programs

Francois Pottier
INRIA
Francois.Pottier@inria.fr

Yann Régis-Gianas
INRIA
Yann.Regis-Gianas@inria.fr

Abstract

We present a Hoare logic for a call-by-value programming language equipped with recursive, higher-order functions, algebraic data types, and a polymorphic type system in the style of Hindley and Milner. It is the theoretical basis for a tool that extracts proof obligations out of programs annotated with logical assertions. These proof obligations, expressed in a typed, higher-order logic, are discharged using off-the-shelf automated or interactive theorem provers. Although the technical apparatus that we exploit is by now standard, its application to call-by-value functional programming languages appears to be new, and (we claim) deserves attention. As a sample application, we check the partial correctness of a balanced binary search tree implementation.

Categories and Subject Descriptors D.3.3 [*Programming Languages*]: Language Constructs and Features—Data types and structures; Polymorphism; Procedures, functions, and subroutines; Recursion; F.3.1 [*Logics and Meanings of Programs*]: Specifying and Verifying and Reasoning about Programs—Assertions; Mechanical verification

General Terms Theory

Keywords Hoare logic, extended static checking

1. Introduction

Hoare logic [24, 30, 14] is a discipline for annotating programs with logical formulae, known as assertions, and for extracting logical formulae, known as proof obligations, out of such annotated programs. The validity of the proof obligations, which can be verified either manually or mechanically, entails the correctness of the annotated program. That is, it guarantees that the assertions are correct static predictions of the program's dynamic behavior. The process of constructing and checking proof obligations is sometimes known as "extended static checking" [19].

Hoare logic was originally designed for a "while language", that is, a simple imperative programming language, equipped with an iteration construct and a fixed number of global, mutable variables. Recursive, higher-order procedures were the subject of much attention in the late 1970's and early 1980's [12, 5, 15, 26, 27]. More recently, heap-allocated, mutable data structures, as well as object-oriented features, have been deeply investigated. This has led to the development of practical specification languages and tools targeting, for instance, C [21], Java [9, 22, 41], and C# [6].

We would like to put forth the thesis that this traditional focus on imperative programming languages has been, to some extent, detrimental: it has consumed a great amount of energy, while comparatively little effort was being devoted to the key features that will be required in order for the methodology to scale up, such as modularity and abstraction. We would also like to raise a question: since functional programs are significantly easier to check for correctness, why hasn't this activity become routine in the functional programming community, forty years after Floyd and Hoare's seminal papers?

On the cost of imperative programming There are several reasons why functional programming can be considered superior to imperative programming [32]. One of them is that functional programs are easier to reason about. In other words, there is a cost to reasoning about state.

In a typical modern imperative programming language, all heap-allocated data is mutable. As a result, instead of reasoning in terms of high-level entities such as, say, pairs, lists, trees, etc., programmers are forced to reason in terms of a view of the heap as a graph. More concretely, they must write down and prove formulae that involve mappings of memory addresses to memory blocks [42, 21].

The possibility of aliasing means that, whenever some memory block is written, the memory that is accessible through every type-compatible pointer is potentially affected. This makes it difficult to reason about the effects of a single write operation, and creates the problem of representation exposure [18, 36]. In order to address this issue, researchers have developed linear types and regions [25], ownership types [11], and separation logic [51], among other approaches.

Our research agenda We do not claim that the above issues are not worth investigating: on the contrary, they are quite fascinating. However, it is a pity that we do not, today, have mature tools for checking the correctness of functional programs. This explains why, in this paper, we study a Hoare logic for (call-by-value) functional programs without state.

The programs that we are interested in checking rely heavily on (possibly higher-order) functions, algebraic data structures, and type polymorphism. We claim that it is quite easy to extract succinct and natural proof obligations out of such programs, provided, of course, that they are annotated with specifications.

There are two benefits to be reaped by not reasoning about state. As far as the user is concerned, this leads to simpler specifications and proof obligations. As far as the implementor is concerned, this saves a large part of the "implementation budget", which can then be spent on features such as type polymorphism, type abstraction, and modularity. The importance of these features cannot be overstated: in the end, the key to success is the ability to develop and check program components independently.

Today, to the best of our knowledge, no tool exists that can extract proof obligations out of a call-by-value functional program and pass them on to an off-the-shelf, interactive or automated, theorem prover. (See §8 for a more precise discussion of related work.) We would like to fill this gap.

Contribution In this paper, we present the design of a typed, polymorphic, higher-order programming language, where programs can be annotated with assertions expressed in a typed, polymorphic, higher-order logic. We define a procedure for extracting proof obligations out of programs, and show that it is sound. A prototype tool [52] has been developed, which works in conjunction with the interactive theorem prover Coq [56], with the automated first-order theorem prover Ergo [13], or with both at once. This tool has been used to check the partial correctness of several non-trivial data structure implementations, including balanced binary search trees and persistent double-ended queues [33]. We hope to publish detailed accounts of these implementations in the future.

Highlights of our approach Here are some of the key technical features of our approach.

We focus on partial correctness. We do not require programs to terminate, and do not generate proof obligations to ensure termination. It is up to the user to determine which properties of the code are of sufficient interest to deserve proof, and to insert assertions where desired. At one extreme, a program that contains no assertions leads to no proof obligations. There is no cost to be paid up front for using our methodology.

Our preconditions are prescriptive: it is impossible to call a function unless its precondition F_1 holds. A descriptive interpretation of preconditions can be simulated by using the precondition **true** and the postcondition $F_1 \Rightarrow F_2$. This allows unconditional invocation, and states that the function's result must satisfy F_2 if its argument satisfies F_1.

Values, programs, types, and logical formulae are distinct syntactic categories. Proofs do not necessarily appear within programs: proof obligations are delegated to an external theorem prover, which may or may not require or produce explicit proof terms.

We do not embed values, programs, or formulae within types. Thus, our types are first-order terms: they include type variables, parameterized algebraic data types, and function types, just as in ML. As a result, type inference in the style of Milner [44] is possible, and implemented in our tool [52]. Type inference does not generate any proof obligations. We do not have dependent types, such as lists indexed with an integer length [58], but simulate them as follows. Instead of declaring that x has type $list\, n$, we declare that x has type $list$, and assert the logical formula $length(x) = n$, where the function $length$ is inductively defined at the logical level.

Formulae can refer to values, but not to expressions. This is important, because values are pure, whereas expressions are potentially impure. Although our logic cannot explicitly reason about state, it is nevertheless soundly applicable to programs that involve non-termination, non-determinism, input/output, or mutable state. (Reading an input stream, or dereferencing a pointer to mutable storage, can be viewed as non-deterministic operations.) In that case, it allows establishing properties that do not depend on the behavior of any impure operation. This means, for instance, that we can prove the partial correctness of a functional program even if it has been instrumented with possibly impure debugging, profiling, or logging instructions.

In our programming language, functions, which are potentially impure, are values, so they can appear within formulae. But what does it mean for a formula to refer to a computational function f of type, say, $\tau_1 \longrightarrow \tau_2$? Our answer is to view f, at the logical level, as a pair of predicates, which represent f's precondition and postcondition. In other words, when used within a formula, f has type (roughly) $(\tau_1 \rightarrow \textbf{prop}) \times (\tau_1 \rightarrow \tau_2 \rightarrow \textbf{prop})$. The two pair projections, written **pre** and **post**, can be used to refer to the pair components. That is, $\textbf{pre}(f)$ and $\textbf{post}(f)$ offer lightweight notations for referring to f's precondition and postcondition. When f is a known (**let**-bound) function, this mechanism can be viewed merely as offering abbreviations for known formulae. However, when f is unknown (λ- or $\forall$-bound), it becomes key to writing natural specifications for higher-order functions (§7.5).

			Logical Types
θ	$::=$	α	*Variable*
	$\mid$	$d\,\bar\theta$	*Data*
	$\mid$	$\theta \rightarrow \theta$	*Function*
	$\mid$	$\theta \times \theta$	*Product*
	$\mid$	**prop**	*Proposition*
ς	$::=$	$\forall\bar\alpha.\,\theta$	*Scheme*
			Logical Type Environments
Δ	$::=$	$\emptyset$	*Nil*
	$\mid$	$\Delta, (x : \varsigma)$	*Variable*
	$\mid$	$\Delta, \bar\alpha$	*Type Variables*
			Logical Terms
t, F, P	$::=$	$x\,\bar\theta$	*Variable*
	$\mid$	$D\,\bar\theta\,(t,\ldots,t)$	*Data*
	$\mid$	$\lambda(x:\theta).t$	*Abstraction*
	$\mid$	$t(t)$	*Application*
	$\mid$	(t, t)	*Product*
	$\mid$	π_1	*Projection (also written* **pre***)*
	$\mid$	π_2	*Projection (also written* **post***)*
	$\mid$	**true**	*Truth*
	$\mid$	$t = t$	*Equality*
	$\mid$	$t \wedge t$	*Conjunction*
	$\mid$	$\neg\, t$	*Negation*
	$\mid$	$\forall(x:\varsigma).t$	*Universal Quantification*
	$\mid$	$\forall\bar\alpha.t$	*Type Quantification*

Figure 1. The logic (syntax)

In summary, although the technical apparatus that we exploit is by now standard, we believe that it is worth drawing attention to the combination of power and simplicity offered by our technical choices. If extended with a suitable module system, and equipped with a compilation path down to, say, Objective Caml [38], our tool could be used to construct correct purely functional program components, possibly for use within larger, partly imperative programs.

Outline of the paper The paper is laid out as follows. First, we briefly introduce a higher-order logic, in which assertions and proof obligations are expressed (§2). Then, we present the syntax and call-by-value semantics of a core functional programming language whose expressions carry explicit assertions (§3). We describe the type system, as well as the procedure for extracting proof obligations out of programs (§4). We present a few extensions of the language (§5) and discuss how proof obligations are transformed for submission to external theorem provers (§6). Last, we present a few excerpts of our balanced binary search tree implementation (§7) and review related work (§8).

2. The underlying logic

2.1 Syntax

We rely on a mostly standard higher-order logic [4] whose types and terms appear in Figure 1. Types θ include type variables, pa-

$$
\begin{array}{rcll}
(\Delta,(x:\varsigma))(x) & = & \varsigma & \\
(\Delta,(x_1:\varsigma))(x_2) & = & \Delta(x_2) & \text{if } x_1 \# x_2 \\
(\Delta,\bar\alpha)(x) & = & \Delta(x) & \text{if } \bar\alpha \# \Delta(x)
\end{array}
$$

$$
\frac{\Delta(x) = \forall\bar\alpha.\,\theta}{\Delta \vdash x\,\bar\theta : [\bar\alpha \mapsto \bar\theta]\theta}
\qquad
\frac{D : \forall\bar\alpha.\,\theta_1 \times \ldots \times \theta_n \to d\,\bar\alpha \qquad \forall i \quad \Delta \vdash t_i : [\bar\alpha \mapsto \bar\theta]\theta_i}{\Delta \vdash D\,\bar\theta\,(t_1,\ldots,t_n) : d\,\bar\theta}
$$

$$
\frac{\Delta,(x:\theta_1) \vdash t : \theta_2}{\Delta \vdash \lambda(x:\theta_1).t : \theta_1 \to \theta_2}
\qquad
\frac{\Delta \vdash t_1 : \theta_1 \to \theta_2 \qquad \Delta \vdash t_2 : \theta_1}{\Delta \vdash t_1(t_2) : \theta_2}
$$

$$
\frac{\forall i \quad \Delta \vdash t_i : \theta_i}{\Delta \vdash (t_1,t_2) : \theta_1 \times \theta_2}
\qquad
\frac{\Delta \vdash t : \theta_1 \times \theta_2}{\Delta \vdash \pi_i(t) : \theta_i}
\qquad
\frac{}{\Delta \vdash \mathbf{true} : \mathbf{prop}}
$$

$$
\frac{\forall i \quad \Delta \vdash t_i : \theta}{\Delta \vdash t_1 = t_2 : \mathbf{prop}}
\qquad
\frac{\forall i \quad \Delta \vdash t_i : \mathbf{prop}}{\Delta \vdash t_1 \wedge t_2 : \mathbf{prop}}
$$

$$
\frac{\Delta \vdash t : \mathbf{prop}}{\Delta \vdash \neg\, t : \mathbf{prop}}
\qquad
\frac{\Delta,(x:\varsigma) \vdash t : \mathbf{prop}}{\Delta \vdash \forall(x:\varsigma).t : \mathbf{prop}}
\qquad
\frac{\Delta,\bar\alpha \vdash t : \mathbf{prop}}{\Delta \vdash \forall\bar\alpha.t : \mathbf{prop}}
$$

Figure 2. The logic (type system)

rameterized inductive types, function types, product types, and the type **prop** of logical propositions. In the following, the syntax of terms is extended with standard syntactic sugar for falsity, disjunction, implication, equivalence, existential quantification, etc.

The typing rules appear in Figure 2. In general, we write t for terms of arbitrary type. We write F for *formulae*, that is, terms of type **prop**, and P for *predicates*, that is, terms of type $\theta \to$ **prop**. The binary operator $\#$, used in several definitions, expresses the fact that two objects have no common free names.

Our logic is not simply-typed. Because our computational language (§3) is polymorphic, and because we wish to lift every computational value up to the logical level, we need polymorphism at the logical level as well. For this reason, we have logical type schemes $\varsigma ::= \forall\bar\alpha.\,\theta$, and every use of a variable x is explicitly annotated with a type vector $\bar\theta$, which states how the type scheme associated with x is instantiated. For this reason also, the syntax of formulae offers universal quantification over type variables, so a fact can be asserted to hold at all types.

The logic offers parameterized inductive types. We assume that each inductive type constructor d carries a fixed integer arity, and that every application $d\,\bar\theta$ is arity-consistent. We further assume that d comes with a finite number of data constructors D, each of which is assigned a type scheme of the form:

$$\forall\bar\alpha.\,\theta_1 \times \ldots \times \theta_n \to d\,\bar\alpha$$

We impose a positivity condition [48], which is informally summed up as follows: in the above type scheme, the type constructor d (or any type constructor whose definition is mutually recursive with the definition of d) must not appear under the left-hand side of an arrow within $\theta_1,\ldots,\theta_n$.

Although there is an introduction form for inductive types, namely the application of a data constructor D, no elimination form is provided here. We can get away with this omission because the process of *extracting* proof obligations, which is the focus of the present paper, requires no such forms. Of course, when it comes to *discharging* proof obligations, that is, proving theorems, then inductive definitions and proofs become necessary.

			Computational Types
τ	$::=$	α	*Variable*
	$\mid$	$d\,\bar\tau$	*Data*
	$\mid$	$\tau \stackrel{\cdot}{\longrightarrow} \tau$	*Function*
σ	$::=$	$\forall\bar\alpha.\,\tau$	*Scheme*
			Computational Type Environments
Γ	$::=$	$\emptyset$	*Nil*
	$\mid$	$\Gamma,(x:\sigma)$	*Variable*
	$\mid$	$\Gamma,\bar\alpha$	*Type Variables*
	$\mid$	Γ,F	*Assumption*
			Values
v	$::=$	$x\,\bar\tau$	*Variable*
	$\mid$	$D\,\bar\tau\,(v,\ldots,v)$	*Data*
	$\mid$	**fun** $f(x:\tau/F) : (x:\tau/F) = e$	*Recursive Function*
			Patterns
p	$::=$	$x\,\bar\tau$	*Variable*
	$\mid$	$D\,\bar\tau\,(p,\ldots,p)$	*Data*
			Expressions
e	$::=$	v	*Value*
	$\mid$	$v(v)$	*Function Application*
	$\mid$	**let** $(x\,\bar\alpha:\tau/F) = e$ **in** e	*Local Binding*
	$\mid$	**case** v **of** c	*Pattern Matching*
			Cases
c	$::=$	$\emptyset$	*Nil*
	$\mid$	$(p \mapsto e)\ []\ c$	*Cons*

Figure 3. The computation language (syntax)

2.2 Interpretation

We assume that the logic is equipped with a consistent interpretation in a model [4], so that, in particular, every well-formed, closed formula evaluates, in the model, to a truth value. We assume that the model validates $\beta\eta\iota$-equivalence. These abstract hypotheses about the model suffice to establish the soundness of our system (§4.8).

3. The computational language

3.1 Syntax

The syntax of our programming language appears in Figure 3. It is equipped with an ML-style type system [44], so types τ and type schemes σ are distinguished. Types include type variables, parameterized algebraic data types, and function types. We write $\stackrel{\cdot}{\longrightarrow}$ for the computational function type constructor, so as to distinguish it from the logical function type constructor, written $\to$ (Figure 1).

We impose a syntactic separation between values and expressions, and require both operands of the function application operator, as well as **case** scrutinees, to be values. This imposes a style, reminiscent of A-normal form, where the result of every intermediate computation is named via a **let** construct. Of course, such a style is quite user-unfriendly, so, in practice, we offer an unrestricted surface language, and automatically translate it down to the kernel language described here.

The language supports type inference in the style of Hindley and Milner. However, in this paper, we are not concerned with type inference, so we work with explicitly-typed programs. This is visible (i) in the syntax of values and patterns, where variables and data constructors are annotated with vectors of types that indicate how polymorphic type schemes are instantiated, (ii) at **fun** and **let** constructs, where bound variables are annotated with types, and (iii) at **let** constructs, where a vector of type variables $\bar{\alpha}$ can be explicitly bound.

A function definition takes the general form:

$$\textbf{fun}\, f(x_1 : \tau_1/F_1) : (x_2 : \tau_2/F_2) = e$$

The keyword $/$ should be read "where". Every function is recursive, so that f is bound within e. The formal parameter x_1 is bound within the precondition F_1, within the postcondition F_2, and within e. The variable x_2, which stands for the result of the function, is bound within the postcondition F_2. We require every function to be annotated with an explicit precondition and postcondition (if missing, **true** is assumed).

A local variable definition takes the general form:

$$\textbf{let}\,(x\,\bar{\alpha} : \tau/F) = e_1 \textbf{ in } e_2$$

The local variable x is bound within F and within e_2. The type variables $\bar{\alpha}$ are bound within τ, F, and e_1. The proposition F serves a postcondition for e_1. If it is missing, a default postcondition is assumed, whose definition is deferred to §3.3.

A case analysis takes the general form:

$$\textbf{case}\, v \textbf{ of } c$$

Here, c is a possibly empty sequence of cases (i.e., branches). Each branch is of the form $(p \mapsto e)$, where the variables that appear in the pattern p are bound within e. Patterns must be linear.

3.2 Lifting computational entities to the logical level

In a Hoare logic, formulae refer to values. That is, if x is bound, at the computational level, by a **fun**, **let**, or **case** construct, then it is possible for a formula F, embedded in the code within the scope of x, to refer to x. This raises two questions: first, if x has computational type τ, what is its logical type, to be used when typechecking F? Second, if, for the purposes of evaluation, x is substituted with a computational value v, what is the corresponding logical value, to be used when interpreting F?

We answer these questions by lifting both computational types and computational values up to the logical level (Figure 4). That is, to each computational type τ, we associate a logical type $\lceil\tau\rceil$, and to each computational value v, we associate a logical term $\lceil v\rceil$, with the intended property that if v has computational type τ, then $\lceil v\rceil$ has logical type $\lceil\tau\rceil$.

As announced (§1), computational functions are reflected, at the logical level, as pairs of a precondition and postcondition. This is made explicit in the lifting of computational function types:

$$\lceil\tau_1 \stackrel{\cdot}{\longrightarrow} \tau_2\rceil = (\lceil\tau_1\rceil \to \textbf{prop}) \times (\lceil\tau_1\rceil \to \lceil\tau_2\rceil \to \textbf{prop})$$

The first component of the pair, which represents the function's precondition, is abstracted over the function's argument, while the second component, which represents the postcondition, is abstracted over both argument and result.

As a result of this definition, if f is bound, at the computational level, to a function of type $\tau_1 \stackrel{\cdot}{\longrightarrow} \tau_2$, then a formula embedded within the code, in the scope of f, views f as a pair of predicates, and can refer to $\textbf{pre}(f)$ and $\textbf{post}(f)$. (Recall that, as per Figure 1, **pre** and **post** are sugar for the projections π_1 and π_2.)

Values of computational function type (that is, λ-abstractions) are lifted up to the logical level in a way that is consistent with this definition. A function's precondition and postcondition alone determine how it is lifted: its code is ignored. (The conformance of a function's body to its declared pre- and postcondition is checked, of course, via a proof obligation: see rule FUN in Figure 7.) This reflects a philosophy in which the only way of reasoning about the behavior of a function is via its specification: it is not possible, in our approach, to reason directly about the code of a function.

In order to lift algebraic data types, we lift every algebraic data type definition into an isomorphic inductive type definition. So, for every computational-level algebraic data type constructor d, there must be a logical-level inductive type constructor, also written d, of identical arity. For every computational-level data constructor

$$D : \forall\bar{\alpha}.\, \tau_1 \times \ldots \times \tau_n \to d\,\bar{\alpha},$$

there must be a logical-level data constructor

$$D : \forall\bar{\alpha}.\, \lceil\tau_1\rceil \times \ldots \times \lceil\tau_n\rceil \to d\,\bar{\alpha}.$$

Due to the manner in which computational function types are lifted, the positivity condition (§2) requires the type constructor d to not appear *under any side* of a computational arrow within $\tau_1, \ldots, \tau_n$. This can be a limitation (§9).

3.3 Inferring strongest postconditions

In order to simplify the definition of the procedure that extracts proof obligations, we have required every **let** construct to carry an explicit postcondition for its left-hand sub-expression (§3.1). In practice, however, annotating every **let** construct would be quite unpleasant, so it is desirable to construct a reasonable postcondition when the user does not provide one.

Ideally, the formula that we should construct in such a situation is the *strongest postcondition* of the left-hand sub-expression. Our logic is, in fact, sufficiently powerful to express strongest postconditions for every construct in our programming language. For instance, the strongest postcondition for a value v is $\lambda x.(x = \lceil v\rceil)$. The strongest postcondition for a function application $v_1(v_2)$ is $\textbf{post}(\lceil v_1\rceil)(\lceil v_2\rceil)$. We could go on and explain how to construct strongest postconditions for **let** and **case** constructs. However, in these two cases, they would be complex formulae, involving existential quantification and disjunction.

Eventually, the postconditions carried by **let** constructs become part of proof obligations, where they appear as hypotheses. For this reason, we do not want them to be too complex: we wish to produce simple, comprehensible proof obligations.

Our answer to this issue is to construct strongest postconditions for values and function applications, as suggested above, but not for **let** and **case** constructs: instead, we rely on the user-provided postcondition, if there is one, or use the trivial postcondition **true**, otherwise.

It is possible, via a conversion to A-normal form, to ensure that the left-hand sub-expression of a **let** construct is never another **let** construct. So, the only case where our simple-minded approach may call for an explicit, user-provided annotation is that of a **let** construct whose left-hand sub-expression is a **case** construct.

3.4 Notions of substitution

Neither types nor formulae influence execution, but they are kept around in order to make a subject reduction statement and prove the soundness of our Hoare logic. So, the operational semantics reduces expressions that contain explicit types and formulae. To ensure that these annotations remain consistent as expressions are transformed, we must define a few slightly non-standard notions of substitution.

A single type variable α can appear within logical types as well as within computational types. Similarly, a single variable x can appear within formulae as well as within expressions. For this reason, we write $[\alpha \mapsto \tau]$ for the substitution that replaces every free occurrence of α at the computational level with τ and every free occurrence of α at the logical level with $\lceil\tau\rceil$. Similarly, we write $[x \mapsto v]$ for the substitution that replaces every free occurrence of x at the computational level with v and every free occurrence of x at the logical level with $\lceil v\rceil$.

$$
\begin{array}{lrcl}
\textit{\textbf{Types}} & \lceil \alpha \rceil & = & \alpha \\
& \lceil d\,\bar{\tau} \rceil & = & d\,\lceil \bar{\tau} \rceil \\
& \lceil \tau_1 \stackrel{\cdot}{\longrightarrow} \tau_2 \rceil & = & (\lceil \tau_1 \rceil \rightarrow \textbf{prop}) \times (\lceil \tau_1 \rceil \rightarrow \lceil \tau_2 \rceil \rightarrow \textbf{prop}) \\
\textit{\textbf{Type schemes}} & \lceil \forall \bar{\alpha}.\, \tau \rceil & = & \forall \bar{\alpha}.\, \lceil \tau \rceil \\
\textit{\textbf{Type environments}} & \lceil \emptyset \rceil & = & \emptyset \\
& \lceil \Gamma, (x : \sigma) \rceil & = & \lceil \Gamma \rceil, (x : \lceil \sigma \rceil) \\
& \lceil \Gamma, \bar{\alpha} \rceil & = & \lceil \Gamma \rceil, \bar{\alpha} \\
& \lceil \Gamma, F \rceil & = & \lceil \Gamma \rceil \\
\textit{\textbf{Values}} & \lceil x\,\bar{\tau} \rceil & = & x\,\lceil \bar{\tau} \rceil \\
& \lceil D\,\bar{\tau}\,(v_1, \ldots, v_n) \rceil & = & D\,\lceil \bar{\tau} \rceil\,(\lceil v_1 \rceil, \ldots, \lceil v_n \rceil) \\
& \lceil \textbf{fun}\, f(x_1 : \tau_1 / F_1) : (x_2 : \tau_2 / F_2) = e \rceil & = & (\lambda(x_1 : \lceil \tau_1 \rceil).F_1, \lambda(x_1 : \lceil \tau_1 \rceil).\lambda(x_2 : \lceil \tau_2 \rceil).F_2)
\end{array}
$$

Figure 4. Lifting computational types and values to the logical level

$$
\begin{array}{rcll}
v_1(v_2) & \rightarrow & [x \mapsto v_2][f \mapsto v_1]e & \text{if } v_1 \text{ is } \textbf{fun}\, f(x : \tau/F) : (\ldots) = e \\
\textbf{let}\,(x\,\bar{\alpha} : \tau/F) = v \textbf{ in } e & \rightarrow & [x \mapsto \Lambda\bar{\alpha}.v]e & \\
\textbf{case}\, v \textbf{ of } (p \mapsto e) \,[]\, c & \rightarrow & [p \mapsto v]e & \text{if } [p \mapsto v] \text{ is defined} \\
& \rightarrow & \textbf{case}\, v \textbf{ of } c & \text{otherwise} \\
\textbf{let}\,(x\,\bar{\alpha} : \tau/F) = e_1 \textbf{ in } e_2 & \rightarrow & \textbf{let}\,(x\,\bar{\alpha} : \tau/F) = e_1' \textbf{ in } e_2 & \text{if } e_1 \rightarrow e_1'
\end{array}
$$

Figure 5. Operational semantics

We have annotated **let** constructs with explicit type abstractions and occurrences of variables with explicit type applications. As a result, contracting a **let**-redex requires contracting type-level β-redexes as well. In order to do so, we write $[x \mapsto \Lambda\bar{\alpha}.v]$ for a substitution that replaces every variable occurrence of the form $x\,\bar{\tau}$ with $[\bar{\alpha} \mapsto \bar{\tau}]v$. Again, this replacement is performed at both computational and logical levels, up to a lifting operation in the latter case.

Last, the notation $[x \mapsto v]$, which denotes a substitution of a value for a variable, is extended to the notation $[p \mapsto v]$, which, when p does not match v, is undefined, and, when p does match v, denotes a simultaneous substitution of values for variables, as follows. The formal definition is:

$$[D\,\bar{\tau}\,(p_1, \ldots, p_n) \mapsto D\,\bar{\tau}\,(v_1, \ldots, v_n)]$$

stands for

$$[p_1 \mapsto v_1] \cup \ldots \cup [p_n \mapsto v_n]$$

Because patterns are linear, this is a union of substitutions whose domains are pairwise disjoint.

3.5 Operational semantics

A standard small-step, call-by-value operational semantics appears in Figure 5. There are three kinds of redexes (β, **let**, and **case**) and one evaluation context (the left-hand side of a **let** construct). An expression is *stuck* if it is irreducible and not a value. It is easy to check that an expression is stuck if and only if it contains, within an evaluation context, a sub-expression of the form $v_1(v_2)$, where v_1 is not a syntactic function, or of the form **case** v **of** $\emptyset$.

4. The type system and proof system

We now equip the computational language with an ML-style type system and with a proof system (a Hoare logic), which can be viewed as an algorithm for extracting proof obligations out of well-typed programs. For the sake of succinctness, both are described using a single set of judgements, which assert at once that a program is well-typed and is annotated with consistent formulae.

$$
\frac{F \text{ valid}}{\emptyset \models F}
\qquad
\frac{\Gamma \models \forall(x : \lceil\sigma\rceil).F}{\Gamma, (x : \sigma) \models F}
\qquad
\frac{\Gamma \models \forall\bar{\alpha}.F}{\Gamma, \bar{\alpha} \models F}
\qquad
\frac{\Gamma \models F_1 \Rightarrow F_2}{\Gamma, F_1 \models F_2}
$$

Figure 6. Emitting proof obligations

In practice, our tool [52] first checks that the program is well-typed, and, at the same time, infers any omitted type annotations. Then, a set of proof obligations, expressed in our typed higher-order logic, is extracted. The fact that the program (including embedded formulae) is well-typed guarantees that the proof obligations are in turn well-typed.

4.1 Environments

The syntax of type environments Γ appears in Figure 3. As is standard, type environments bind variables and type variables. Environments also contain assumptions, that is, formulae that become hypotheses when proof obligations are emitted. An environment of the form Γ, F is well-formed when F has type **prop** under $\lceil\Gamma\rceil$.

4.2 Proof obligations

A proof obligation is a judgement of the form $\Gamma \models F$, where F has type **prop** under $\lceil\Gamma\rceil$. This judgement is defined by the rules of Figure 6. When read from bottom to top, the rules can be understood as translating a proof obligation to a closed formula of type **prop**, which must be valid for the judgement to hold. The validity of a formula is decided via an external theorem prover.

4.3 Judgements

The proof system is defined via three judgements, which state properties about values, patterns, and expressions, respectively:

Values	$\Gamma \vdash v : \tau$	(Figure 7)
Patterns	$\Gamma \vdash p : \tau$	(Figure 8)
Expressions	$\Gamma \vdash e : \tau\ \{P\}$	(Figure 9)

$$\begin{array}{rcll}(\Gamma,(x:\sigma))(x) & = & \sigma & \\ (\Gamma,(x_1:\sigma))(x_2) & = & \Gamma(x_2) & \text{if } x_1 \# x_2 \\ (\Gamma,\bar\alpha)(x) & = & \Gamma(x) & \text{if } \bar\alpha \# \Gamma(x) \\ (\Gamma,F)(x) & = & \Gamma(x) & \end{array}$$

VAR
$$\frac{\Gamma(x) = \forall\bar\alpha.\tau}{\Gamma \vdash x\,\bar\tau : [\bar\alpha \mapsto \bar\tau]\tau}$$

DATA
$$\frac{D : \forall\bar\alpha.\tau_1 \times \ldots \times \tau_n \to d\,\bar\alpha \qquad \forall i \quad \Gamma \vdash v_i : [\bar\alpha \mapsto \bar\tau]\tau_i}{\Gamma \vdash D\,\bar\tau\,(v_1,\ldots,v_n) : d\,\bar\tau}$$

FUN
$$\frac{\begin{array}{c} f \# F_1, F_2 \\ \lceil \Gamma, (x_1:\tau_1)\rceil \vdash F_1 : \textbf{prop} \\ \lceil \Gamma, (x_1:\tau_1),(x_2:\tau_2)\rceil \vdash F_2 : \textbf{prop} \\ \Gamma, (f : \tau_1 \stackrel{\cdot}{\longrightarrow} \tau_2), f = \lceil \textbf{fun}\, f \ldots\rceil, (x_1:\tau_1), F_1 \\ \vdash e : \tau_2 \ \{\lambda(x_2 : \lceil\tau_2\rceil).F_2\}\end{array}}{\Gamma \vdash \textbf{fun}\, f(x_1:\tau_1/F_1) : (x_2:\tau_2/F_2) = \ e : \tau_1 \stackrel{\cdot}{\longrightarrow} \tau_2}$$

Figure 7. The computation language (proof system: values)

PAT-VAR
$$(x:\tau) \vdash x : \tau$$

PAT-DATA
$$\frac{D : \forall\bar\alpha.\tau_1 \times \ldots \times \tau_n \to d\,\bar\alpha \qquad \forall i \quad \Gamma_i \vdash p_i : [\bar\alpha \mapsto \bar\tau]\tau_i}{\Gamma_1,\ldots,\Gamma_n \vdash D\,\bar\tau\,(p_1,\ldots,p_n) : d\,\bar\tau}$$

Figure 8. The computation language (proof system: patterns)

VALUE
$$\frac{\Gamma \vdash v : \tau \qquad \Gamma \models P(\lceil v\rceil)}{\Gamma \vdash v : \tau\ \{P\}}$$

APP
$$\frac{\Gamma \vdash v_1 : \tau_1 \stackrel{\cdot}{\longrightarrow} \tau_2 \qquad \Gamma \vdash v_2 : \tau_1 \qquad \Gamma \models \textbf{pre}(\lceil v_1\rceil)(\lceil v_2\rceil) \qquad \Gamma \models \textbf{post}(\lceil v_1\rceil)(\lceil v_2\rceil) \Rightarrow P}{\Gamma \vdash v_1(v_2) : \tau_2\ \{P\}}$$

LET
$$\frac{\begin{array}{c} x \# P \\ \lceil \Gamma, \bar\alpha, (x:\tau_1)\rceil \vdash F : \textbf{prop} \\ \Gamma, \bar\alpha \vdash e_1 : \tau_1\ \{\lambda(x:\lceil\tau_1\rceil).F\} \\ \Gamma, (x : \forall\bar\alpha.\tau_1), \forall\bar\alpha.[x \mapsto x\,\bar\alpha]F \vdash e_2 : \tau_2\ \{P\}\end{array}}{\Gamma \vdash \textbf{let}\,(x\,\bar\alpha : \tau_1/F) = e_1\ \textbf{in}\ e_2 : \tau_2\ \{P\}}$$

CASE-NIL
$$\frac{\Gamma \vdash v : \tau \qquad \Gamma \models \textbf{false}}{\Gamma \vdash \textbf{case}\, v\, \textbf{of}\ \emptyset : \tau'\ \{P\}}$$

CASE-CONS
$$\frac{\begin{array}{c} \Gamma \vdash v : \tau \qquad \Gamma' \vdash p : \tau \qquad p \# v, P \\ \Gamma, \Gamma', \lceil v\rceil = \lceil p\rceil \vdash e : \tau'\ \{P\} \\ \Gamma, (\forall\Gamma'.\lceil v\rceil \neq \lceil p\rceil) \vdash \textbf{case}\, v\, \textbf{of}\ c : \tau'\ \{P\}\end{array}}{\Gamma \vdash \textbf{case}\, v\, \textbf{of}\ (p \mapsto e)\ []\ c : \tau'\ \{P\}}$$

Figure 9. The computation language (proof system: expressions)

4.4 Values

The judgement $\Gamma \vdash v : \tau$ (Figure 7) states that, under the type environment Γ, the value v has type τ. No precondition or postcondition appear in the judgement. Indeed, because values require no computation, they never have a precondition. Furthermore, because all values can be lifted up to the logical level, they don't need an explicit postcondition: the strongest possible postcondition of a value v is simply equality with $\lceil v\rceil$.

Rules VAR and DATA are straightforward. Rule FUN is more complex. Two premises require the precondition F_1 and postcondition F_2 to be well-formed formulae, under appropriate environments. The last premise, which spans two lines, checks that the function's body conforms to the function's specification. In order to do so, the type environment is extended with bindings for f and x_1. It is also extended with the hypothesis

$$f = \lceil \textbf{fun}\, f \ldots\rceil,$$

which by definition of lifting (Figure 4) is synonymous for

$$f = (\lambda(x_1 : \lceil\tau_1\rceil).F_1, \lambda(x_1 : \lceil\tau_1\rceil).\lambda(x_2 : \lceil\tau_2\rceil).F_2).$$

This hypothesis gives meaning to occurrences of **pre**(f) and **post**(f) within the body of the function, allowing recursive calls to f to be checked. Last, the environment is also extended with the precondition F_1, which means that, within the body of the function, the precondition is assumed to hold. Under this extended environment, the body of the function is required to produce a value that meets the postcondition $\lambda(x_2 : \lceil\tau_2\rceil).F_2$.

It is not difficult to see that $\Gamma \vdash v : \tau$ implies $\lceil\Gamma\rceil \vdash \lceil v\rceil : \lceil\tau\rceil$. This property is required for the typing rules to construct only well-formed formulae.

4.5 Patterns

The judgement $\Gamma \vdash p : \tau$ (Figure 8) states that a value of type τ can safely be matched against the pattern p, giving rise to (exactly) the bindings described by Γ. These bindings are monomorphic (see PAT-VAR). Because patterns are linear, the type environments $\Gamma_1,\ldots,\Gamma_n$ in PAT-DATA have disjoint domains.

4.6 Expressions

The judgement $\Gamma \vdash e : \tau\ \{P\}$ (Figure 9) states that, under the type environment Γ, the expression e has type τ and (if it terminates) produces a value whose logical reflection satisfies the predicate P. In such a judgement, P has type $\lceil\tau\rceil \to$ **prop** under $\lceil\Gamma\rceil$.

Rule VALUE directly reflects this intended meaning: the judgement $\Gamma \vdash v : \tau\ \{P\}$ holds if and only if v has type τ under Γ and its logical reflection $\lceil v\rceil$ provably satisfies P under the hypotheses found in Γ. The premise $\Gamma \models P(\lceil v\rceil)$ is a proof obligation.

Rule APP requires the function v_1 and its actual argument v_2 to have matching computational types. Furthermore, it emits two proof obligations, stating that (i) the actual argument must satisfy the function's precondition, and (ii) the function's postcondition must imply the desired postcondition P. In the last premise, we write $P' \Rightarrow P$, where P' and P have type $\lceil\tau_2\rceil \to$ **prop**, for $\forall(x : \lceil\tau_2\rceil).(P'(x) \Rightarrow P(x))$, where x is fresh for P' and P.

Rule LET checks that e_1 has type τ_1 and that e_1 complies with the postcondition F. Then, the rule performs type generalization, in the style of Milner [44], so that e_2 is checked under the assignment $(x : \forall\bar\alpha.\tau_1)$. The hypothesis F is changed into $\forall\bar\alpha.[x \mapsto x\,\bar\alpha]F$, so as to reflect the fact that x now has polymorphic type.

Rule CASE-NIL emits the proof obligation $\Gamma \models$ **false**, which requires the conjunction of hypotheses found within Γ to be inconsistent. This ensures that a **case** construct with zero branches is never executed.

Rule CASE-CONS requires the value v and the pattern p to have a common type τ. The environment Γ' collects the variables bound by p, together with their types. Under the hypothesis that a certain instance of p matches v, which is expressed by extending Γ with Γ' and with the hypothesis $\lceil v \rceil = \lceil p \rceil$, the branch e must have the desired type τ' and meet the desired postcondition P. Furthermore, under the hypothesis that no instance of p matches v, which is written $\forall\Gamma'.\lceil v \rceil \neq \lceil p \rceil$, the remaining branches must have type τ' and meet the postcondition P. (Our use of $\lceil p \rceil$ plays on the fact that patterns form a subset of values, a welcome but unessential property.)

When checking a **case** construct with n branches, the $(k+1)$-th branch is checked under the assumption that none of the patterns $p_1, \ldots, p_k$ match the value v. In particular, for $k = n$, the conjunction of all hypotheses of the form $(\forall\Gamma'_i.\lceil v \rceil \neq \lceil p_i \rceil)$ is required to be inconsistent. This ensures that control cannot fall off the end of a **case** construct, or, in other words, that the case analyses are exhaustive. Today's ML and Haskell compilers implement a sound approximation to this check, using a purely syntactic criterion. We also implement this syntactic criterion: when it succeeds, emitting a proof obligation is unnecessary.

4.7 Algorithmic reading

The judgement $\Gamma \vdash e : \tau\ \{P\}$ defines an algorithm for generating proof obligations. All four parameters of the judgement, namely Γ, e, τ, and P, are inputs of the algorithm, which attempts to build a derivation of the judgement by starting at the root of the expression e and working its way down into the sub-expressions of e. As the algorithm descends, entering **fun**, **let**, and **case** constructs, the environment Γ grows, accumulating new bindings and assumptions. At the same time, the postcondition P is propagated down, in a very straightforward process. At **let** constructs, this propagation process relies on the (default or user-provided, see §3.3) annotation in order to determine which postcondition must be propagated into the left-hand sub-expression. The output of the algorithm consists of the proof obligations, of the form $\Gamma \models F$, carried by the leaves of the derivation (see VALUE, APP, and CASE-NIL).

4.8 Soundness

The soundness of our type system and proof system is established in a standard, syntactic manner. It states that the types and logical assertions carried by a program are a sound approximation of its dynamic semantics.

Lemma 4.1 (Environment Weakening) $\Gamma_1, F, \Gamma_2 \vdash e : \tau\ \{P\}$ *and* $\Gamma_1 \models F$ *imply* $\Gamma_1, \Gamma_2 \vdash e : \tau\ \{P\}$. ◇

Lemma 4.2 (Postcondition Weakening) $\Gamma \vdash e : \tau\ \{P_1\}$ *and* $\Gamma \models P_1 \Rightarrow P_2$ *imply* $\Gamma \vdash e : \tau\ \{P_2\}$. ◇

Lemma 4.3 (Type Substitution) *Let* ϕ *stand for* $[\bar\alpha \mapsto \bar\tau]$. *Then,* $\Gamma_1, \bar\alpha, \Gamma_2 \vdash e : \tau\ \{P\}$ *and* $\bar\alpha \# dom(\Gamma_2)$ *imply*

$$\Gamma_1, \phi(\Gamma_2) \vdash \phi(e) : \phi(\tau_2)\ \{\phi(P)\}$$ ◇

Lemma 4.4 (Value Substitution) *Let* ρ *stand for* $[x \mapsto \Lambda\bar\alpha.v]$. *Then,* $\Gamma_1, (x : \forall\bar\alpha.\tau_1), \Gamma_2 \vdash e : \tau_2\ \{P\}$ *and* $\Gamma_1, \bar\alpha \vdash v : \tau_1$ *and* $x \notin dom(\Gamma_2)$ *imply*

$$\Gamma_1, \rho(\Gamma_2) \vdash \rho(e) : \tau_2\ \{\rho(P)\}$$ ◇

Lemma 4.5 (Pattern Matching) *Let* $\emptyset \vdash v : \tau$ *and* $\Gamma' \vdash p : \tau$ *and* $p \# v$. *Then,* $[p \mapsto v]$ *is defined if and only if the formula* $\exists\Gamma'.\lceil v \rceil = \lceil p \rceil$ *is valid.* ◇

Theorem 4.6 (Subject Reduction) $\Gamma \vdash e : \tau\ \{P\}$ *and* $e \to e'$ *imply* $\Gamma \vdash e' : \tau\ \{P\}$. ◇

Theorem 4.7 (Progress) $\emptyset \vdash e : \tau\ \{P\}$ *implies that* e *is either reducible or a value* v *such that* $P(\lceil v \rceil)$ *is valid.* ◇

5. A few extensions

Extra assertions The following construct allows inserting an assertion at an arbitrary point in the code:

$$\textbf{assert}\ F\ \textbf{in}\ e$$

This construct requires F to hold: a proof obligation is emitted. It has no computational content: dynamically, it behaves like e. Here, it can be viewed as syntactic sugar for a **let** construct. It is particularly useful when our tool is used in conjunction with an automated theorem prover: if the theorem prover fails to discharge a proof obligation, the user can use **assert** to cut the proof into smaller, easier steps (if the proof obligation is in fact valid) or to find out what is wrong with the specification (if the proof obligation is in fact invalid).

The construct **absurd**, which requires **false** to hold, marks a piece of code as inaccessible. Here, it can be viewed as syntactic sugar for a **case** construct with zero branches.

Ghost variables and ghost parameters It is sometimes desirable to explicitly introduce a *ghost variable*, that is, a name for a witness to an existentially quantified hypothesis. For this purpose, we suggest writing

$$\textbf{let logic}\ x : \theta / F\ \textbf{in}\ e$$

This construct binds x within F and e. It requires the assertion $\exists(x : \theta).F$ to hold, and introduces F as a new hypothesis into the context. Assertions embedded within e can refer to x, and their proofs can exploit the hypothesis F. However, computational occurrences of x within e are forbidden, since "**let logic**" has no computational content.

Similarly, it is sometimes desirable to abstract a function with respect to a ghost parameter x, like this:

$$\textbf{fun}\ f(\textbf{logic}\ x : \theta)(x_1 : \tau_1/F_1) : (x_2 : \tau_2/F_2) = e$$

This construct binds x within F_1, F_2, and e. Note that θ can be an arbitrary logical type, so this extension allows explicitly abstracting a function with respect to a proposition or predicate, if desired (see §7.5). Ghost variables and ghost parameters can in principle be viewed as syntactic sugar and translated away [46]. In a realistic implementation, however, they should be primitive notions.

6. Interfacing with external theorem provers

6.1 Coq

Our typed, higher-order logic is easily embedded within the Calculus of Inductive Constructions, which underlies Coq [56]. As a result, exporting proof obligations towards Coq is a simple matter of pretty-printing.

Coq is an interactive theorem prover. In order to discharge a proof obligation, the user writes a proof script. An open problem is how to maintain these scripts as the source code of the program evolves. The location in the code where a proof obligation arises might change. The statement of a proof obligation might change as well. Perhaps a solution would be to allow only explicitly-stated, explicitly-named, lemmas to be proved interactively, and to rely solely on an automated theorem prover for discharging anonymous proof obligations, possibly by appeal to an explicit lemma.

6.2 Ergo

Ergo [13] is a fully automated theorem prover for a typed, polymorphic, first-order logic. Its design is partly inspired by Simplify [17]. However, Ergo's logic is typed and polymorphic, whereas Simplify's is untyped. This makes Ergo superior, from our point of view, to Simplify. Indeed, provided our proof obligations lie in the first-order fragment of our logic, they can be directly exported towards Ergo. If, on the other hand, we wished to use Simplify, we would have to encode our typed, polymorphic logic into Simplify's untyped logic. Such encodings have been studied [39], but are complex and costly. Of course, the trivial encoding that erases all types is unsound.

In addition to first-order logic, Ergo has native support for linear arithmetic and for the theory of constructors (that is, function symbols f such that $f(x) = f(y)$ implies $x = y$). The latter is useful for reasoning efficiently about algebraic data structures.

In the general case, our proof obligations are most naturally expressed in a higher-order logic, as shown in this paper. However, higher-order logic can be encoded into first-order logic. A standard encoding introduces "apply" predicates that help simulate β-conversion [34].

Perhaps surprisingly, in our case, this encoding can be made to look fairly natural. The symbols **pre** and **post**, which so far have stood for the pair projections, can be turned into predicates and simulate not only projection, but also application. Furthermore, we can make **pre** a binary predicate and **post** a ternary predicate, avoiding curried function applications. That is, instead of the higher-order formula:

$$f = (\lambda(x_1 : \lceil\tau_1\rceil).F_1, \lambda(x_1 : \lceil\tau_1\rceil).\lambda(x_2 : \lceil\tau_2\rceil).F_2),$$

we can write:

$$\begin{aligned} &\forall(x_1 : \lceil\tau_1\rceil).(\mathbf{pre}(f, x_1) \Leftrightarrow F_1) \\ \wedge\ &\forall(x_1 : \lceil\tau_1\rceil).\forall(x_2 : \lceil\tau_2\rceil).(\mathbf{post}(f, x_1, x_2) \Leftrightarrow F_2) \end{aligned}$$

The pair and the three λ-abstractions have been η-expanded, and the projection and application symbols have been fused into applications of **pre** and **post**. Provided F_1 and F_2 are first-order formulae, this is a first-order formula.

Under this encoding, the definition of the lifting operation on computational types is modified so that the computational function type constructor is no longer interpreted:

$$\lceil\tau_1 \stackrel{\cdot}{\longrightarrow} \tau_2\rceil = \lceil\tau_1\rceil \stackrel{\cdot}{\longrightarrow} \lceil\tau_2\rceil$$

That is, we make $\stackrel{\cdot}{\longrightarrow}$ an uninterpreted binary type constructor at the logical level, so that the lifting of types becomes the identity. Thus, in the above formula, f has logical type $\tau_1 \stackrel{\cdot}{\longrightarrow} \tau_2$. The type schemes assigned to **pre** and **post** are as follows:

$$\begin{aligned} \mathbf{pre} &: \forall\alpha_1\alpha_2.\,(\alpha_1 \stackrel{\cdot}{\longrightarrow} \alpha_2) \times \alpha_1 \to \mathbf{prop} \\ \mathbf{post} &: \forall\alpha_1\alpha_2.\,(\alpha_1 \stackrel{\cdot}{\longrightarrow} \alpha_2) \times \alpha_1 \times \alpha_2 \to \mathbf{prop} \end{aligned}$$

These declarations are admissible by Ergo. We believe that it should be possible to go a long way with first-order logic alone, even when the program exploits higher-order functions. However, at present, more practical experience is needed in order to support this conjecture.

7. Application: finite sets as binary search trees

As an initial benchmark for our tool [52], we have transcribed Objective Caml's library implementation of finite sets, represented as balanced binary search trees, into our programming language.

7.1 Parameters

In the following, we fix a type "elt" of elements. We assume that an algebraic data type "bool", whose data constructors are "true" and "false", is available. We assume that an equality check over elements, written "=", is given. It is a function of computational type elt $\times$ elt $\stackrel{\cdot}{\longrightarrow}$ bool, whose specification could be written as follows:

$$\mathbf{post}(=, x_1, x_2, b) \Leftrightarrow (b = \text{true} \Leftrightarrow x_1 = x_2)$$

Similarly, we assume that an ordering relation, written "$<$", of logical type elt $\to$ elt $\to$ **prop**, is given, together with an ordering check, also written "$<$", of computational type elt $\times$ elt $\stackrel{\cdot}{\longrightarrow}$ bool, such that the latter decides the former.

In a full-scale programming language, our balanced binary search tree implementation would be a functor, parameterized over the type "elt", the function "=", the relation "$<$", the ordering axioms for "$<$", and the function "$<$".

We assume that an abstract type of sets of elements, written "set elt", is available at the logical level, together with the standard operations (empty set, singleton set, union, membership, etc.) and a number of axioms or theorems that describe the properties of these operations. In a full-scale programming language, this would be provided by a (logical-level) standard library.

type tree =
| Empty : tree
| Node : (elt $\times$ tree $\times$ tree $\times$ int) $\to$ tree

logic fun elements : tree $\to$ set elt =
| Empty $\to \emptyset$
| Node (x, l, r, _) $\to$ {x} $\cup$ elements (l) $\cup$ elements (r)

inductive predicate bst : tree $\to$ **prop** =
| bst (Empty)
| $\forall$ x, l, r. h.
  ($\forall$y. y $\in$ elements (l) $\Rightarrow$ y $<$ x) $\wedge$ bst (l) $\wedge$
  ($\forall$y. y $\in$ elements (r) $\Rightarrow$ x $<$ y) $\wedge$ bst (r) $\wedge$
  $\Rightarrow$ bst (Node (x, l, r, h))

Figure 10. Definitions for binary search trees

7.2 Definitions

Figure 10 contains the definition of the algebraic data type "tree", of the logical-level inductive function "elements", and of the inductive predicate "bst". (The concrete syntax is provisional.) A binary tree is either empty or a binary node, carrying a root element, left and right sub-trees, and a cached measure of the tree's height. Our binary search trees are intended to implement a finite set abstraction. The logical function "elements" maps a binary tree to the finite set that it represents. It is defined by induction over the algebraic data type "tree". The property of being a binary search tree is defined by the inductive predicate "bst".

In the definition of "bst", the types of the universally quantified variables "x", "l", "r", "h", "y" are inferred. The types of the function "elements" and of the predicate "bst" could also be inferred, if desired. In practice, type annotations can always be omitted, except where polymorphic recursion is required.

The definition of "bst" constrains neither the shape of the tree nor the cached height information. This is done by another inductive predicate, named "avl" (not shown). In contrast with the "dependent types" [58, 3, 55] and "generalized algebraic data types" [59] schools, we favor a programming style in which invariants are not necessarily hardwired into data structures at definition time.

7.3 Membership in a binary search tree

Figure 11 shows a function, "member", that checks whether an element "x" is a member of a tree "t". The precondition "bst(t)" requires "t" to be a binary search tree, but does not require it to be balanced, since this is not necessary for correctness. If one wished to (informally) ensure a logarithmic complexity bound, one could strengthen the precondition by adding the requirement "avl(t)". This illustrates how a single data structure can be equipped with multiple invariants, not all of which are necessarily enforced at all times. The postcondition states that the Boolean result tells whether

```
let rec member (x, t)
where bst (t)
returns u where b = true ⇔ x ∈ elements (u)
= match t with
  | Empty → false
  | Node (y, l, r, _) →
    if x = y then
      true
    else if x < y then
      member (x, l)
    else
      member (x, r)
  end
```

Figure 11. Membership in a binary search tree

```
type iterator = list tree

logic fun remaining : iterator → set elt = ...

inductive predicate ok : iterator → prop =
| ok ([])
| ∀ t, ts. elements (t) ∩ remaining (ts) = ∅ ∧
      bst (t) ∧ ok (ts) ⇒ ok (t :: ts)

let iterator (t) returns i
where ok (i) ∧ elements (t) ≡ remaining (i)
= [t]

let next (i)
where ok (i)
returns oix : option (iterator × elt)
where oix = None ⇒ remaining (i) = ∅
      ∀i', x. oix = Some (i', x) ⇒
        remaining (i) ≡ {x} ∪ remaining (i')
        x ∉ remaining (i') ∧ ok (i')
= match i with
  | [] → None
  | Empty :: ts → next (ts)
  | Node (x, l, r, _) :: ts → Some (l :: r :: ts, x)
  end
```

Figure 12. Iterators over binary search trees

"x" is a member of the set implemented by the tree "t". No type annotations are needed in this definition. All types are inferred.

7.4 First-order iteration

We now define and specify first-order, persistent iterators over binary search trees. Their expressive power surpasses that of "fold" (§7.5), yet their specification is simpler.

The implementation appears in Figure 12. An iterator is represented as a list of trees, which can be thought of as a stack in a depth-first traversal of some larger tree.

To an iterator "i" corresponds a set of elements, which we write "remaining(i)". Its inductive definition (omitted) is simply the union of the sets of elements of the trees in the list.

An iterator is well-formed only if the trees that it contains have disjoint sets of elements. This is expressed by the inductive predicate "ok".

The function "iterator" creates an iterator "i" out of a tree "t", and satisfies the postcondition "$\mathrm{ok}(\mathrm{i}) \wedge \mathrm{elements}(\mathrm{t}) \equiv \mathrm{remaining}(\mathrm{i})$", where $\equiv$ stands for extensional equality of sets (which may, or may not, coincide with definitional equality). This initial iterator is simply the singleton list [t].

The function "next", when applied to an iterator "i", returns either nothing or a pair of a new iterator "i'" and an element "x". The postcondition describes how these values are related.

```
let cardinal (t) returns n where n = || elements (t) || =
  let rec count (i, n)
  where ok(i) ∧ n + || remaining (i) || = || elements (t) ||
  returns n' where n' = || elements (t) ||
    = match next (i) with
    | None → n
    | Some (i', _) → count (i', n + 1)
    end
  in count (iterator (t), 0)
```

Figure 13. A sample client of the iterator abstraction

Figure 13 shows how iterators are used. Here, the client is a function that counts the number of elements in a tree. It does not depend on the internals of the tree data structure: it only depends on the specification of iterators, which is expressed in terms of abstract (logical-level) sets. So, this client code could be placed in another module, without access to the definition of trees.

The "cardinal" function performs a loop, expressed as an internal recursive function, with an integer accumulator n. The precondition of this internal function represents the loop invariant: the number of elements counted so far, plus the number of elements remaining to be seen, equals the total number of elements of the set. The postcondition is simply the precondition, specialized to the case where no elements remain.

The precondition of "count" must also state that "i" is an "ok" iterator, even though it does not have to know about the definition of "ok". This is somewhat undesirable. In the future, we will want to allow defining a dependent sum type of the form "i : iterator **where** ok(i)", and exporting it as an abstract type.

The definition of "cardinal" is syntactically somewhat heavy, as it is expressed in our core language. In a full-scale programming language, a more palatable syntax for loops could be introduced, and desugared into recursive functions and iterators. A single formula, the loop invariant, would have to be written down, instead of two formulae in this low-level version of the code.

7.5 Higher-order iteration

We now present a specification of the classic "fold" higher-order function over sets implemented as binary search trees. The specification is rather more complex than that of first-order iterators, for at least two reasons. First, the specification must mention the client's state (the accumulator) and invariant. Second, because the code is not tail-recursive, some information is implicitly encoded within the stack, and a ghost parameter is used to make it explicit in the specification.

The definition "is_invariant" summarizes the required invariant. Its parameters are "inv", the client's invariant, which itself is parameterized over and accumulator and a set of remaining elements; "s_0", a set of remaining elements; "s_1", the initial set of all elements; the client function "f"; and an accumulator "a" ccu. The definition states, in short, that: the set of remaining elements is a subset of the initial set of all elements; the current accumulator and the current set of remaining elements satisfy the client's invariant; and, at any time, if an element "x" is picked among the remaining elements, the invariant guarantees that it is legal to apply "f" to the current accumulator and to "x", and guarantees that the new accumulator thus obtained will still satisfy the invariant.

The function "fold" is parameterized over two ghost variables, namely the client invariant "inv" and a set "s_0" of remaining el-

```
predicate is_invariant (inv, s0, s1, f, accu) =
  s0 ⊆ s1 ∧ inv (accu, s0) ∧
  ∀ s x s' accu accu'.
    s ⊆ s0 ∧ s = s' ∪ {x} ∧ x ∉ s' ∧ inv (accu, s) ⇒
      pre (f) (accu, x) ∧
      post (f) (accu, x) (accu') ⇒ inv (accu', s')

let rec fold (logic s0, logic inv, accu, t, f)
where is_invariant (inv, s0, elements (t), f, accu) ∧ bst (t)
returns accu' where inv (accu', s0\ elements (t)) =
  match t with
  | Empty → accu
  | Node (x, l, r, _) →
    let accu = fold (s0, inv, accu, l, f) in
    let accu = f (accu, x) in
      fold (s0\ (elements (l) ∪ {x}), inv, accu, r, f)
  end

predicate cardinal_inv (t) =
  λ (accu, s). accu + || s || = || elements (t) ||

let cardinal (t) returns x where x = || elements (t) || =
  fold (elements (t), cardinal_inv (t), 0, t, λ (accu, x). accu + 1)
```

Figure 14. Folding over binary search trees

ements. In the case of first-order iterators, the former was unnecessary because the client retains control over the desired invariant, and the latter was unnecessary because the set of remaining elements was directly expressed as "remaining(i)". Here, the set of remaining elements is implicit in the stack, so a ghost variable must be used in order to refer to it.

This definition is certainly somewhat overwhelming. It shows, at the same time, that it is possible to specify and exploit higher-order functions in our framework, and that there is a cost in complexity to be paid for doing so. More experience is needed before we can tell how easily higher-order functions can be defined and used in practice.

8. Related work

The roots of our work lie in Hoare logic [24, 30]. Extensions of Hoare logic with support for recursive, higher-order procedures were heavily studied in the late 1970's and early 1980's [12, 5, 15, 26, 27]. In particular, the issue of completeness received a lot of attention after Clarke [12] proved that there can be no sound and complete Hoare logic for a programming language equipped with recursive, higher-order procedures and global variables. Clarke's result, however, is based upon the assumption that formulae and proof obligations are expressed in a first-order logic. Damm and Josko [15] point out that, by moving to higher-order logic, it is possible to work around Clarke's negative result. In this paper, we follow Damm and Josko and allow specifications to be expressed in higher-order logic. The intuitive justification for this approach is that, if functions can abstract over functions, then specifications must abstract over specifications.

Our work has been strongly inspired by several existing, practical tools for checking imperative programs [19, 22, 20, 21, 41, 6]. This paper is an attempt to exploit the strengths of these works while steering away from imperative programming and placing renewed emphasis on polymorphism and modularity.

Our method for generating proof obligations is particularly straightforward: it appears in its entirety in Figures 6–9. In comparison with the method used in ESC/Java [23], we avoid a translation to "passive form" because we have no assignments to begin with. We avoid the exponential explosion that could follow from the interplay between sequences and alternatives by requiring sequences (that is, **let** constructs) to carry user-provided postconditions (§3.3).

Our system is not sound with respect to a call-by-name dynamic semantics. There are at least two reasons for this fact. First, some divergent expressions admit **false** as a valid postcondition. If such an expression e_1 is made the first component of a sequence, as in "**let** x/**false** $= e_1$ **in** e_2", then second component e_2 is checked under the assumption **false**. As a result, all of the the proof obligations found within e_2 are vacuously satisfied. This can be sound only if e_2 is never executed, which is the case under call-by-value evaluation, but not under call-by-name evaluation. The second reason is that, in a call-by-name semantics, every type is inhabited by a bottom value, and some types are inhabited by infinite values. This is not reflected in the way we lift computational values and types up to the logical level.

Scott's logic of computable functions [53] interprets λ-terms in a denotational model, where equality implies, or coincides with, observational equivalence. It comes with a set of sound deduction rules, and allows explicit reasoning about divergence and equality of computations. It admits call-by-value and call-by-name variants. It was implemented as early as 1972 by Milner [43]. More recent implementations [2, 8, 45] embed Scott's LCF within some form of higher-order logic. In a somewhat similar vein, Longley and Pollack [40] embed the functional core of Standard ML, via a fully abstract denotational semantics, into higher-order logic.

Our approach is less elaborate: by focusing on partial correctness, by adopting a call-by-value semantics, and by lifting only values, as opposed to expressions, up to the logical level, we are able to ignore non-termination issues entirely, and to work with value spaces that do not have bottom elements or definedness orderings. By contrast, tools or approaches that focus on lazy functional programs, such as Programatica [35, 29], the Cover translator [1], or Honda and Yoshida's logic of higher-order functions [31], require reasoning about non-termination, resulting in proof obligations that can become cluttered with definedness side conditions. The simplicity of our approach comes at a cost: our system can neither establish termination of an expression nor reason about observational equality of expressions.

ESC/Haskell [60] allows annotating Haskell programs with preconditions and postconditions that are also expressed in Haskell. A special-purpose theorem prover, based on symbolic evaluation of Haskell terms, is developed.

The theorem prover Coq [56] can be used as a programming language, in which programs are both developed and proved correct. The Compcert certified compiler [37] offers an example of a large program developed in this style. However, there is some agreement that Coq is not (yet) a convenient programming language: for instance, it only allows writing pure, terminating functions.

The programming language Russell [55] extends Coq with facilities for defining programs annotated with assertions, in the style of Hoare logic. There are many similarities between Russell and our work. One important technical difference is that we separate the typechecking process, which is performed first and remains traditional, and the process of extracting proof obligations, which runs as a second phase, whereas, in Russell, as in Coq, typechecking and proving are one and the same activity. In particular, Russell encourages the use of indexed types, like $list\ n$, so that typechecking can give rise to proof obligations: for instance, supplying an actual argument of type $list\ m$ to a function that expects a formal parameter of type $list\ n$ generates the proof obligation $m = n$. Another difference is that Russell terms are elaborated into Coq terms, whereas we adopt a less foundational approach and are happy to trust an external theorem prover.

Hoare Type Theory [47, 46] is somewhat similar to our system, insofar as it offers decidable basic typechecking and decidable generation of proof obligations. It also shares our use of higher-order logic and our emphasis on polymorphism and abstraction. It is much more ambitious than our proposal, in that it attempts to deal not only with algebraic data types and higher-order functions, but also with heap-allocated, mutable state. As a result, its design and metatheory are considerably more involved.

Some authors [7, 28, 60, 46] allow code to appear in specifications. This is motivated partly by a desire to make formulae executable, so as to allow assertions to be checked at runtime, and partly by fear that, otherwise, a single functionality might have to be implemented twice: once at the computational level, once at the logical level. Our technical and philosophical choice is different: we consider all code as potentially impure, and do not allow code to appear within specifications. We do not check assertions at runtime: if the programmer wishes to insert a runtime check, she must do so explicitly. Furthermore, we believe that, in practice, opportunities for code sharing between computational and logical levels are rare: the oft-cited case of lists is one of only a few situations where implementation and specification coincide.

Indexed types [61, 58] and refinement types [16] rely on so-called indices. Indices are elements of some mathematical domain, such as an arbitrary finite set, or the set of all natural numbers. Types are enriched with constraints over indices, allowing invariants, preconditions, and postconditions to be expressed. The syntax of constraints is carefully restricted so as to ensure that constraint entailment is decidable. This allows proof obligations to be automatically checked. Generalized algebraic data types [59] are also an instance of this idea, where indices are types, that is, first-order terms. The appeal of this approach resides in the high degree of automation that it allows. On the other hand, this comes at the price of a restriction to a decidable logic. In fact, our decision of using a highly expressive, hence undecidable, logic was motivated by our earlier study of generalized algebraic data types [50, 49].

Going beyond indexed types, several programming languages offer full dependent types [3, 10, 54, 57]. By exploiting the Curry-Howard isomorphism, they allow code and proofs to be expressed and combined within a single language. This allows programs to appear more self-contained, but means that a fragment of the programming language must be a consistent logic, and requires mechanisms to assist the user in building proofs. Our design, which relies on an off-the-shelf theorem prover, is more modular.

9. Conclusion

We have presented a simple methodology for extracting proof obligations out of call-by-value functional programs. Our proposed future work includes:

- publishing a useable prototype implementation, equipped with a compilation path down to Objective Caml;
- relaxing our positivity condition (§3.2), which restricts the use of functions within data structures, preventing, for instance, the standard definition of infinite streams;
- internalizing type equality, that is, introducing equations between types into the syntax of formulae, together with suitable conversion rules for exploiting such equations; indeed, we, and other authors [46], have noticed that such an extension would subsume generalized algebraic data types [59];
- studying the issues raised by modularity and mutable state.

Acknowledgement The authors wish to thank the anonymous reviewers of a previous version of this paper for contradicting a false claim and offering useful comments and suggestions.

References

[1] Andreas Abel, Marcin Benke, Ana Bove, John Hughes, and Ulf Norell. Verifying Haskell programs using constructive type theory. In *Haskell workshop*, pages 62–73, September 2005.

[2] Sten Agerholm. A HOL basis for reasoning about functional programs. Technical Report RS-94-44, BRICS, December 1994.

[3] Thorsten Altenkirch, Conor McBride, and James McKinna. Why dependent types matter. Unpublished, April 2005.

[4] Peter B. Andrews. *An introduction to mathematical logic and type theory: to truth through proof*. Academic Press, 1986.

[5] Krzysztof R. Apt. Ten years of Hoare's logic: A survey—part I. *ACM Transactions on Programming Languages and Systems*, 3(4):431–483, 1981.

[6] Mike Barnett, K. Rustan M. Leino, and Wolfram Schulte. The Spec# programming system: An overview. In *International Workshop on Construction and Analysis of Safe, Secure and Interoperable Smart devices (CASSIS)*, volume 3362 of *Lecture Notes in Computer Science*. Springer Verlag, 2004.

[7] Mike Barnett, David A. Naumann, Wolfram Schulte, and Qi Sun. 99.44% pure: Useful abstractions in specifications. In *Formal Techniques for Java-like Programs*, 2004.

[8] Frank Bartels, Friedrich von Henke, Holger Pfeifer, and Harald Rueß. Mechanizing domain theory. Ulmer Informatik-Berichte 96-10, Universität Ulm, Fakultät für Informatik, 1996.

[9] Lilian Burdy, Yoonsik Cheon, David Cok, Michael Ernst, Joe Kiniry, Gary T. Leavens, K. Rustan M. Leino, and Erik Poll. An overview of JML tools and applications. *International Journal on Software Tools for Technology Transfer*, 7(3):212–232, June 2005.

[10] Chiyan Chen and Hongwei Xi. Combining programming with theorem proving. In *ACM International Conference on Functional Programming (ICFP)*, September 2005.

[11] David G. Clarke, John M. Potter, and James Noble. Ownership types for flexible alias protection. In *ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 48–64, 1998.

[12] Edmund Clarke. Programming language constructs for which it is impossible to obtain good Hoare axiom systems. *Journal of the ACM*, 26(1):129–147, January 1979.

[13] Sylvain Conchon and Evelyne Contejean. The Ergo automatic theorem prover, 2006. `http://ergo.lri.fr/`.

[14] Patrick Cousot. Methods and logics for proving programs. In *Formal Models and Semantics*, volume B of *Handbook of Theoretical Computer Science*, chapter 15, pages 841–993. Elsevier Science, 1990.

[15] Werner Damm and Bernhard Josko. A sound and relatively* complete axiomatization of Clarke's language L4. In *Logic of Programs*, volume 164 of *Lecture Notes in Computer Science*, pages 161–175. Springer Verlag, 1983.

[16] Rowan Davies. Practical refinement-type checking. Technical Report CMU-CS-05-110, School of Computer Science, Carnegie Mellon University, May 2005.

[17] David Detlefs, Greg Nelson, and James B. Saxe. Simplify: a theorem prover for program checking. *Journal of the ACM*, 52(3):365–473, 2005.

[18] David L. Detlefs, K. Rustan M. Leino, and Greg Nelson. Wrestling with rep exposure. Research Report 156, SRC, July 1998.

[19] David L. Detlefs, K. Rustan M. Leino, Greg Nelson, and James B. Saxe. Extended static checking. Research Report 159, Compaq SRC, December 1998.

[20] Jean-Christophe Filliâtre. Why: a multi-language multi-prover verification tool. Research Report 1366, LRI, Université Paris Sud, March 2003.

[21] Jean-Christophe Filliâtre and Claude Marché. Multi-prover verification of C programs. In *International Conference on Formal Engineer-*

ing Methods (ICFEM), volume 3308 of *Lecture Notes in Computer Science*, pages 15–29. Springer Verlag, November 2004.

[22] Cormac Flanagan, K. Rustan M. Leino, Mark Lillibridge, Greg Nelson, James B. Saxe, and Raymie Stata. Extended static checking for Java. In *ACM Conference on Programming Language Design and Implementation (PLDI)*, pages 234–245, 2002.

[23] Cormac Flanagan and James B. Saxe. Avoiding exponential explosion: generating compact verification conditions. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 193–205, 2001.

[24] R. W. Floyd. Assigning meanings to programs. In *Mathematical Aspects of Computer Science*, volume 19 of *Proceedings of Symposia in Applied Mathematics*, pages 19–32. American Mathematical Society, 1967.

[25] Manuel Fähndrich and Robert DeLine. Adoption and focus: practical linear types for imperative programming. In *ACM Conference on Programming Language Design and Implementation (PLDI)*, pages 13–24, June 2002.

[26] Steven German, Edmund Clarke, and Joseph Halpern. Reasoning about procedures as parameters. In *Logic of Programs*, volume 164 of *Lecture Notes in Computer Science*, pages 206–220. Springer Verlag, 1983.

[27] Andreas Goerdt. A Hoare calculus for functions defined by recursion on higher types. In *Logic of Programs*, volume 193 of *Lecture Notes in Computer Science*, pages 106–117. Springer Verlag, 1985.

[28] Jessica Gronski, Kenneth Knowles, Aaron Tomb, Stephen N. Freund, and Cormac Flanagan. Sage: Hybrid checking for flexible specifications. In *Scheme and Functional Programming*, September 2006.

[29] Thomas Hallgren, James Hook, Mark P. Jones, and Richard Kieburtz. An overview of the Programatica toolset. High Confidence Software and Systems Conference (HCSS), 2004.

[30] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–580, 1969.

[31] Kohei Honda and Nobuko Yoshida. A compositional logic for polymorphic higher-order functions. In *International ACM Conference on Principles and Practice of Declarative Programming (PPDP)*, pages 191–202, August 2004.

[32] John Hughes. Why functional programming matters. *Computer Journal*, 32(2):98–107, 1989.

[33] Haim Kaplan and Robert E. Tarjan. Purely functional, real-time deques with catenation. *Journal of the ACM*, 46(5):577–603, 1999.

[34] Manfred Kerber. How to prove higher order theorems in first order logic. In *International Joint Conferences on Artificial Intelligence*, pages 137–142, 1991.

[35] Richard B. Kieburtz. P-logic: Property verification for Haskell programs. Draft, August 2002.

[36] K. Rustan M. Leino and Greg Nelson. Data abstraction and information hiding. *ACM Transactions on Programming Languages and Systems*, 24(5):491–553, 2002.

[37] Xavier Leroy. Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 42–54, January 2006.

[38] Xavier Leroy, Damien Doligez, Jacques Garrigue, Didier Rémy, and Jérôme Vouillon. *The Objective Caml system*, October 2005.

[39] Stéphane Lescuyer. Codage de la logique du premier ordre polymorphe multi-sortée dans la logique sans sortes. Master's thesis, Master Parisien de Recherche en Informatique, 2006.

[40] John Longley and Randy Pollack. Reasoning about CBV functional programs in Isabelle/HOL. In *International Conference on Theorem Proving in Higher Order Logics (TPHOLs)*, volume 3223 of *Lecture Notes in Computer Science*, pages 201–216. Springer Verlag, September 2004.

[41] Claude Marché, Christine Paulin-Mohring, and Xavier Urbain. The Krakatoa tool for certification of Java/JavaCard programs annotated in JML. *Journal of Logic and Algebraic Programming*, 58(1–2):89–106, 2004.

[42] Farhad Mehta and Tobias Nipkow. Proving pointer programs in higher-order logic. *Information and Computation*, 199(1–2):200–227, 2005.

[43] Robin Milner. Implementation and applications of Scott's logic for computable functions. In *Proceedings of the ACM conference on proving assertions about programs*, pages 1–6, January 1972.

[44] Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17(3):348–375, December 1978.

[45] Olaf Müller, Tobias Nipkow, David von Oheimb, and Oskar Slotosch. HOLCF = HOL + LCF. *Journal of Functional Programming*, 9:191–223, 1999.

[46] Aleksandar Nanevski, Amal Ahmed, Greg Morrisett, and Lars Birkedal. Abstract predicates and mutable ADTs in Hoare type theory. In *European Symposium on Programming (ESOP)*, Lecture Notes in Computer Science. Springer Verlag, March 2007.

[47] Aleksandar Nanevski, Greg Morrisett, and Lars Birkedal. Polymorphism and separation in Hoare type theory. In *ACM International Conference on Functional Programming (ICFP)*, pages 62–73, September 2006.

[48] Christine Paulin-Mohring. Inductive definitions in the system Coq: rules and properties. Research Report RR1992-49, ENS Lyon, 1992.

[49] François Pottier and Yann Régis-Gianas. Stratified type inference for generalized algebraic data types. In *ACM Symposium on Principles of Programming Languages (POPL)*, January 2006.

[50] François Pottier and Yann Régis-Gianas. Towards efficient, typed LR parsers. In *ACM Workshop on ML*, volume 148 of *Electronic Notes in Theoretical Computer Science*, pages 155–180, March 2006.

[51] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *IEEE Symposium on Logic in Computer Science (LICS)*, pages 55–74, 2002.

[52] Yann Régis-Gianas. Extended static checking of call-by-value functional programs: Prototype tool. `http://gallium.inria.fr/~regisgia`, July 2007.

[53] Dana S. Scott. A type-theoretical alternative to ISWIM, CUCH, OWHY. *Theoretical Computer Science*, 121(1–2):411–440, 1993.

[54] Tim Sheard. Putting Curry-Howard to work. In *Haskell workshop*, 2005.

[55] Matthieu Sozeau. Subset coercions in Coq. In *TYPES*, 2006.

[56] The Coq development team. *The Coq Proof Assistant*, 2006.

[57] Edwin Westbrook, Aaron Stump, and Ian Wehrman. A language-based approach to functionally correct imperative programming. In *ACM International Conference on Functional Programming (ICFP)*, pages 268–279, 2005.

[58] Hongwei Xi. *Dependent Types in Practical Programming*. PhD thesis, Carnegie Mellon University, December 1998.

[59] Hongwei Xi, Chiyan Chen, and Gang Chen. Guarded recursive datatype constructors. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 224–235, January 2003.

[60] Dana N. Xu. Extended static checking for Haskell. In *Haskell workshop*, pages 48–59. ACM Press, 2006.

[61] Christoph Zenger. Indexed types. *Theoretical Computer Science*, 187:147–165, 1997.