

One Logic to Use Them All

Jean-Christophe Filliâtre^{1,2,3}

¹ CNRS

² LRI, Univ. Paris-Sud, Orsay, F-91405

³ INRIA Saclay – Île-de-France, Orsay, F-91893

Abstract. Deductive program verification is making fast progress these days. One of the reasons is a tremendous improvement of theorem provers in the last two decades. This includes various kinds of automated theorem provers, such as ATP systems and SMT solvers, and interactive proof assistants. Yet most tools for program verification are built around a single theorem prover. Instead, we defend the idea that a collaborative use of several provers is a key to easier and faster verification.

This paper introduces a logic that is designed to target a wide set of theorem provers. It is an extension of first-order logic with polymorphism, algebraic data types, recursive definitions, and inductive predicates. It is implemented in the tool Why3, and has been successfully used in the verification of many non-trivial programs.

1 Introduction

The idea behind deductive program verification [20] is to break down the correctness of a program to a set of logical formulas and to prove them valid. A tremendous improvement of theorem provers in the last two decades now allows this idea to scale up. Projects such as CompCert [31] and seL4 [26] show how interactive proof assistants (resp. Coq and Isabelle) can be successfully used to tackle large program verifications. Even more impressive is the progress in automated provers¹, notably the so-called *SMT revolution*. Designed with program verification in mind, SMT solvers have led to powerful verification tools, such as VCC [39], Frama-C [23], Dafny [29], or VeriFast [25], to mention only a few. Yet we note that most of these tools are built on top of a single automated prover (*e.g.* Z3 in the case of VCC, Dafny, or VeriFast) or a single dedicated prover to handle a specific program logic (*e.g.* B method [1] or KIV [37]).

We rather defend the idea that program verification should be done using as many theorem provers as possible, including those that were not necessarily designed with program verification in mind (*e.g.* ATP systems). We are developing the tool Why3 [10,22] to implement this idea. Both the specification logic of Why3 and its programming language, WhyML, are used as intermediate languages in tools such as Frama-C [23], Krakatoa [33], Easycrypt [4], and GNATprove [14]. WhyML is also used directly to implement data structures and algorithms that can later be translated to executable OCaml code. Our gallery of

¹ We use *automated provers* to denote both SMT solvers and ATP systems.

verified programs (<http://toccata.lri.fr/gallery/>) currently contains more than 80 entries. These examples show the benefits of our approach. Indeed, it is often the case that several theorem provers are successfully used in proving all goals, but none can prove all of them by itself. In particular, an interactive theorem prover can be used to discharge a complex lemma (for instance one requiring induction), while the remaining goals can all be discharged automatically.

Designing a logic to target a wide set of theorem provers is not that easy. We have come up with a *logic of compromise*, that is not as rich as the logic of proof assistants such as Coq or PVS, yet richer than the usual logic of automated theorem provers. Our logic is an extension of first-order logic with rank-1 polymorphism, algebraic data types, recursive definitions, and inductive predicates. The purpose of this paper is to define this logic (Section 2), and to explain the processes by which Why3 translates it to the input format of fifteen theorem provers (Section 3). Finally, Section 4 presents experimental results obtained in the context of program verification. We conclude with related work and perspectives.

2 A Polymorphic First-Order Logic

This section describes the logic of Why3 as faithfully as possible. Earlier work [11] only describes the extension of first-order logic with polymorphism. Here we also consider algebraic data types, recursive definitions, and inductive predicates.

2.1 Syntax

A type symbol is simply a name $\mathbf{t}$ and a type arity $n \in \mathbb{N}$. We write it $\mathbf{t}\langle\alpha_1, \dots, \alpha_n\rangle$ where the α_i are type variables. Names $\alpha_1, \dots, \alpha_n$ are not relevant but we adopt an homogeneous syntax for all symbols. We note d_t such a declaration:

$$d_t ::= \mathbf{t}\langle\alpha, \dots, \alpha\rangle \quad \text{type symbol declaration}$$

When $n = 0$, we say that $\mathbf{t}$ is a *monomorphic type symbol* and we simply write $\mathbf{t}$ instead of $\mathbf{t}\langle\rangle$. In the following, let Σ_T be a set of type symbols. We assume that Σ_T contains at least the two monomorphic type symbols `int` and `real`. Types are built from type variables and type symbols:

$$\begin{array}{ll} \tau ::= \alpha & \text{type variable} \\ \quad | \mathbf{t}\langle\tau, \dots, \tau\rangle & \text{type symbol application} \end{array}$$

We note $tv(\tau)$ the set of type variables of type τ . When $tv(\tau) = \emptyset$, we say that τ is a *sort*.

Function and predicate symbols are declared, possibly with polymorphic types, as follows:

$$\begin{array}{ll} d_f ::= f\langle\alpha, \dots, \alpha\rangle(\tau, \dots, \tau) : \tau & \text{function symbol declaration} \\ d_p ::= p\langle\alpha, \dots, \alpha\rangle(\tau, \dots, \tau) & \text{predicate symbol declaration} \end{array}$$

In the following, Σ_F denotes a set of function symbols and Σ_P a set of predicate symbols. We assume that Σ_P contains at least a polymorphic predicate $=\langle\alpha\rangle(\alpha, \alpha)$ that denotes equality. Terms and formulas are then built from function and predicate symbols according to the syntax given in Fig. 1–3. Note that syntax for terms and formulas are mutually recursive, since a conditional term expression **if** f **then** t_1 **else** t_2 involves a formula f .

We note $fv(t)$ (resp. $fv(f)$, $fv(p)$) the set of free variables of a term t (resp. a formula f , a pattern p). Definitions of $fv(t)$ and $fv(f)$ are standard, variables being bound by **let**, quantifiers, and patterns. A pattern p binds all variables in $fv(p)$, which is defined as follows:

$$\begin{aligned} fv(x_\tau) &= \{x_\tau\} \\ fv(f(\tau_1, \dots, \tau_m)(p_1, \dots, p_n)) &= fv(p_1) \cup \dots \cup fv(p_n) \\ fv(_) &= \emptyset \\ fv(p_1 \mid p_2) &= fv(p_1) \cup fv(p_2) \\ fv(p \text{ as } x_\tau) &= fv(p) \cup \{x_\tau\} \end{aligned}$$

The type checking rule for pattern $p_1 \mid p_2$, given in next section, imposes that $fv(p_1)$ and $fv(p_2)$ are equal. Thus the definition above actually simplifies to $fv(p_1 \mid p_2) = fv(p_1) = fv(p_2)$.

A *signature* Σ denotes a triple $(\Sigma_T, \Sigma_F, \Sigma_P)$. A *context* Γ extends a signature Σ with definitions for some of its symbols. A definition d introduces either algebraic data types, recursive definitions, or inductive predicates, as follows:

d	datatype a with ... with a	algebraic data types
	recursive δ with ... with δ	recursive definitions
	inductive i with ... with i	inductive predicates
a	$d_t = d_f \mid \dots \mid d_f$	algebraic data type
δ	function $f\langle\alpha, \dots, \alpha\rangle(x_\tau, \dots, x_\tau) : \tau = t$	function definition
	predicate $p\langle\alpha, \dots, \alpha\rangle(x_\tau, \dots, x_\tau) = f$	predicate definition
i	$d_p = f \mid \dots \mid f$	inductive predicate

Example. Let Σ_T be the following set of type symbols:

$$\Sigma_T = \{\text{int}, \text{real}, \text{nat}, \text{list}\langle\alpha\rangle, \text{pair}\langle\alpha_1, \alpha_2\rangle\}.$$

Let Σ_F and Σ_P be the following sets of function and predicate symbols:

$$\begin{aligned} \Sigma_F &= \{O : \text{nat}, S(\text{nat}) : \text{nat}, \\ &\quad \text{Nil}\langle\alpha\rangle : \text{list}\langle\alpha\rangle, \text{Cons}\langle\alpha\rangle(\alpha, \text{list}\langle\alpha\rangle) : \text{list}\langle\alpha\rangle, \text{length}\langle\alpha\rangle(\text{list}\langle\alpha\rangle) : \text{nat}, \\ &\quad \text{Pair}\langle\alpha_1, \alpha_2\rangle(\alpha_1, \alpha_2) : \text{pair}\langle\alpha_1, \alpha_2\rangle\} \\ \Sigma_P &= \{=\langle\alpha\rangle(\alpha, \alpha), \text{even}(\text{nat})\} \end{aligned}$$

$t ::=$	c_{int}	literal integer constant
	c_{real}	literal real constant
	x_τ	variable
	$f\langle\tau, \dots, \tau\rangle(t, \dots, t)$	function symbol application
	let $x_\tau = t$ in t	local binding
	if f then t else t	conditional expression
	match t with $p \rightarrow t \mid \dots \mid p \rightarrow t$ end	pattern matching

Fig. 1. Syntax for terms

$f ::=$	$\mathbf{p}\langle\tau, \dots, \tau\rangle(t, \dots, t)$	predicate symbol application
	$\forall x_\tau. f$	universal quantification
	$\exists x_\tau. f$	existential quantification
	$f \wedge f$	conjunction
	$f \vee f$	disjunction
	$f \Rightarrow f$	implication
	$f \Leftrightarrow f$	equivalence
	not f	negation
	true	tautology
	false	absurdity
	let $x_\tau = t$ in f	local binding
	if f then f else f	conditional expression
	match t with $p \rightarrow f \mid \dots \mid p \rightarrow f$ end	pattern matching

Fig. 2. Syntax for formulas

$p ::=$	x_τ	variable
	$f\langle\tau, \dots, \tau\rangle(p, \dots, p)$	constructor application
	-	catch all
	$p \mid p$	or pattern
	p as x_τ	binding

Fig. 3. Syntax for patterns

Let Γ be the extension of the signature above with the following set of definitions:

```

datatype nat = O | S(nat)
datatype list< $\alpha$ > = Nil< $\alpha$ > | Cons< $\alpha$ >( $\alpha$ , list< $\alpha$ >)
recursive function length< $\alpha$ >( $l_{\text{list}\langle\alpha\rangle}$ ) : nat =
  match  $l_{\text{list}\langle\alpha\rangle}$  with Nil< $\alpha$ >  $\rightarrow$  O | Cons< $\alpha$ >( $\_$ ,  $r_{\text{list}\langle\alpha\rangle}$ )  $\rightarrow$  S(length( $r_{\text{list}\langle\alpha\rangle}$ )) end
inductive even(nat) =
  even(O) |  $\forall n_{\text{nat}}. \text{even}(n_{\text{nat}}) \Rightarrow \text{even}(S(S(n_{\text{nat}})))$ 

```

Here and below, we omit type annotations when they are obvious from the context.

2.2 Type Checking

For a signature Σ to be well-formed, any type expression τ occurring in a declaration of Σ_F or Σ_P must be well-typed, that is $\Sigma \vdash \tau$. Additionally, the free type variables of τ must be included in the type variables $\langle\alpha_1, \dots, \alpha_m\rangle$ of the declaration. In the following, we only consider signatures that are well-formed. Let Γ be a context extending a signature Σ with definitions. For Γ to be well-formed, any symbol from a definition must appear in Σ , and no symbol can be defined more than once.

Type checking of types, terms, formulas, and patterns is straightforward. Judgements are $\Sigma \vdash \tau$ (type τ is well-formed in Σ), $\Sigma \vdash t : \tau$ (term t is well-formed and of type τ), $\Sigma \vdash f$ (formula f is well-formed), and $\Sigma \vdash p : \tau$ (pattern p is well-formed and of type τ), respectively. Typing rules are given in Fig. 4–7. Note that, contrary to Hindley-Milner type system, the **let** binder does not generalize the type of the term that is bound. Thus polymorphism is introduced by symbols, but not by local definitions. In addition to these rules, we also check that, for any pattern matching expression **match** t **with** $p_1 \rightarrow . \mid \dots \mid p_n \rightarrow .$ **end** (in terms or formulas), t has type $\mathbf{t}\langle\tau_1, \dots, \tau_m\rangle$ where type $\mathbf{t}$ is defined as an algebraic data type and patterns $p_1, \dots, p_n$ cover any possible value for term t .

Finally, we perform the following checks on definitions.

Algebraic Data Types. Within a block of mutually-defined algebraic data types **datatype** a_1 **with** ... **with** a_k , each algebraic data type definition a_i must have the form

$$\begin{aligned}
 \mathbf{t}_i\langle\alpha_1, \dots, \alpha_m\rangle &= \mathbf{f}_1\langle\alpha_1, \dots, \alpha_m\rangle(\dots) : \mathbf{t}_i\langle\alpha_1, \dots, \alpha_m\rangle \\
 &\mid \dots \\
 &\mid \mathbf{f}_l\langle\alpha_1, \dots, \alpha_m\rangle(\dots) : \mathbf{t}_i\langle\alpha_1, \dots, \alpha_m\rangle
 \end{aligned}$$

that is, all constructors $\mathbf{f}_j$ share the same type parameters as type $\mathbf{t}_i$ and all have return type $\mathbf{t}_i\langle\alpha_1, \dots, \alpha_m\rangle$. Additionally, we require that all types $\mathbf{t}_i$ are *inhabited* according to the following definition: Type $\mathbf{t}_i$ is inhabited if it has at least one constructor $\mathbf{f}_j\langle\alpha_1, \dots, \alpha_m\rangle(\tau_1, \dots, \tau_{n_j})$ such that all types τ_l are inhabited. Since types τ_l may recursively involve types from data type definitions, such a definition is to be understood as a least fixed point. Types from Σ that are

$$\frac{\mathbf{t}\langle\alpha_1, \dots, \alpha_m\rangle \in \Sigma \quad \Sigma \vdash \tau_i}{\Sigma \vdash \mathbf{t}\langle\tau_1, \dots, \tau_m\rangle}$$

Fig. 4. Typing rule for types

$$\frac{\frac{\frac{\Sigma \vdash c_{\text{int}} : \text{int}}{\Sigma \vdash c_{\text{real}} : \text{real}} \quad \frac{\Sigma \vdash \tau}{\Sigma \vdash x_\tau : \tau} \quad \begin{array}{l} f\langle\alpha_1, \dots, \alpha_m\rangle(\tau'_1, \dots, \tau'_n) : \tau \in \Sigma \quad \Sigma \vdash \tau_i \\ \sigma = \{\alpha_1 \mapsto \tau_1, \dots, \alpha_m \mapsto \tau_m\} \quad \Sigma \vdash t_i : \sigma(\tau'_i) \end{array}}{\Sigma \vdash f\langle\tau_1, \dots, \tau_m\rangle(t_1, \dots, t_n) : \sigma(\tau)} \\ \frac{\Sigma \vdash t_1 : \tau \quad \Sigma \vdash t_2 : \tau_2}{\Sigma \vdash \text{let } x_\tau = t_1 \text{ in } t_2 : \tau_2} \quad \frac{\Sigma \vdash f \quad \Sigma \vdash t_1 : \tau \quad \Sigma \vdash t_2 : \tau}{\Sigma \vdash \text{if } f \text{ then } t_1 \text{ else } t_2 : \tau} \\ \frac{\Sigma \vdash t : \mathbf{t}\langle\tau_1, \dots, \tau_m\rangle \quad \Sigma \vdash p_i : \mathbf{t}\langle\tau_1, \dots, \tau_m\rangle \quad \Sigma \vdash t_i : \tau}{\Sigma \vdash \text{match } t \text{ with } p_1 \rightarrow t_1 \mid \dots \mid p_n \rightarrow t_n \text{ end} : \tau}$$

Fig. 5. Typing rules for terms

$$\frac{\frac{\mathbf{p}\langle\alpha_1, \dots, \alpha_m\rangle(\tau'_1, \dots, \tau'_n) \in \Sigma \quad \Sigma \vdash \tau_i}{\sigma = \{\alpha_1 \mapsto \tau_1, \dots, \alpha_m \mapsto \tau_m\} \quad \Sigma \vdash t_i : \sigma(\tau'_i)}}{\Sigma \vdash \mathbf{p}\langle\tau_1, \dots, \tau_m\rangle(t_1, \dots, t_n)} \\ \frac{\frac{\Sigma \vdash \text{true} \quad \Sigma \vdash \text{false}}{\Sigma \vdash f_1 \circ f_2} \quad \frac{\Sigma \vdash f_1 \quad \Sigma \vdash f_2 \quad \circ \in \{\wedge, \vee, \Rightarrow, \Leftrightarrow\}}{\Sigma \vdash f_1 \circ f_2}}{\Sigma \vdash f_1 \circ f_2} \\ \frac{\Sigma \vdash \tau \quad \Sigma \vdash f}{\Sigma \vdash \forall x_\tau. f} \quad \frac{\Sigma \vdash \tau \quad \Sigma \vdash f}{\Sigma \vdash \exists x_\tau. f} \quad \frac{\Sigma \vdash f}{\Sigma \vdash \text{not } f} \\ \frac{\Sigma \vdash t : \tau \quad \Sigma \vdash f}{\Sigma \vdash \text{let } x_\tau = t \text{ in } f} \quad \frac{\Sigma \vdash f_1 \quad \Sigma \vdash f_2 \quad \Sigma \vdash f_3}{\Sigma \vdash \text{if } f_1 \text{ then } f_2 \text{ else } f_3} \\ \frac{\Sigma \vdash t : \mathbf{t}\langle\tau_1, \dots, \tau_m\rangle \quad \Sigma \vdash p_i : \mathbf{t}\langle\tau_1, \dots, \tau_m\rangle \quad \Sigma \vdash f_i}{\Sigma \vdash \text{match } t \text{ with } p_1 \rightarrow f_1 \mid \dots \mid p_n \rightarrow f_n \text{ end}}$$

Fig. 6. Typing rules for formulas

$$\frac{\frac{\Sigma \vdash \tau}{\Sigma \vdash x_\tau : \tau} \quad \frac{\Sigma \vdash \tau}{\Sigma \vdash _ : \tau} \quad \begin{array}{l} f\langle\alpha_1, \dots, \alpha_m\rangle(\tau'_1, \dots, \tau'_n) : \tau \in \Sigma \quad \Sigma \vdash \tau_i \quad \sigma = \{\alpha_1 \mapsto \tau_1, \dots, \alpha_m \mapsto \tau_m\} \\ i \neq j \Rightarrow fv(p_i) \cap fv(p_j) = \emptyset \quad \Sigma \vdash p_i : \sigma(\tau'_i) \end{array}}{\Sigma \vdash f\langle\tau_1, \dots, \tau_m\rangle(p_1, \dots, p_n) : \sigma(\tau)} \\ \frac{fv(p_1) = fv(p_2) \quad \Sigma \vdash p_i : \tau}{\Sigma \vdash p_1 \mid p_2 : \tau} \quad \frac{x_\tau \notin fv(p) \quad \Sigma \vdash p : \tau}{\Sigma \vdash p \text{ as } x_\tau : \tau}$$

Fig. 7. Typing rules for patterns

not part of data type definitions are assumed to be inhabited. As a consequence of this definition, an algebraic data type must have at least one constructor.

Recursive Definitions. A block of mutually recursive definitions **recursive** δ_1 **with** ... **with** δ_k is well-formed if each definition δ_i is well-typed. A function definition

$$\mathbf{function} \ f\langle\alpha, \dots, \alpha\rangle(x_\tau, \dots, x_\tau) : \tau = t$$

is well-typed if $\Sigma \vdash t : \tau$, and a predicate definition

$$\mathbf{predicate} \ p\langle\alpha, \dots, \alpha\rangle(x_\tau, \dots, x_\tau) = f$$

is well-typed if $\Sigma \vdash f$. Additionally, all free variables (including type variables) in t and f must belong to the symbol declaration. Finally, we require such definitions to *terminate*, according to the following criterion: there must exist a lexicographic order of arguments that guarantees a structural descent (over algebraic data types).

Inductive Predicates. An inductive definition

$$p\langle\alpha_1, \dots, \alpha_m\rangle(\tau_1, \dots, \tau_n) = f_1 \mid \dots \mid f_k$$

is well-formed if each clause f_i is a closed formula, well-typed *i.e.* $\Sigma \vdash f_i$, that belongs to the following grammar:

$$\begin{array}{lcl} f_0 & ::= & p\langle\alpha_1, \dots, \alpha_m\rangle(t_1, \dots, t_n) \\ & | & f \Rightarrow f_0 \\ & | & \forall x_\tau. f_0 \\ & | & \mathbf{let} \ x_\tau = t \ \mathbf{in} \ f_0 \end{array}$$

Additionally, we require all inductive predicates from that block of inductive definitions to appear in strictly positive positions in the formulas on the left side of $\Rightarrow$, so as to ensure the existence of a least fixpoint.

In the following, we only consider contexts that are well-formed according to the typing rules we just introduced.

2.3 Semantics

We recall that a sort is a monomorphic type, that is, a type τ such that $tv(\tau) = \emptyset$. In the following, we use notation α (resp. s , x , t) for a vector of type variables (resp. sorts, variables, terms). Given a signature, a *pre-interpretation* $\llbracket \cdot \rrbracket$ is defined as follows:

- Each sort s is interpreted as a non-empty domain $\llbracket s \rrbracket$. Sort **int** is interpreted as $\mathbb{Z}$ and sort **real** as $\mathbb{R}$.

- Given a function symbol $f\langle\alpha\rangle(\tau_1, \dots, \tau_n) : \tau$ and sorts $\mathbf{s}$, we interpret the instance $f\langle\mathbf{s}\rangle$ as a function $\llbracket f\langle\mathbf{s}\rangle \rrbracket$ of type

$$\llbracket \sigma(\tau_1) \rrbracket \times \dots \times \llbracket \sigma(\tau_n) \rrbracket \rightarrow \llbracket \sigma(\tau) \rrbracket$$

where σ maps the type variables α to the sorts $\mathbf{s}$.

- Given a predicate symbol $p\langle\alpha\rangle(\tau_1, \dots, \tau_n)$ and sorts $\mathbf{s}$, we interpret the instance $p\langle\mathbf{s}\rangle$ as a function $\llbracket p\langle\mathbf{s}\rangle \rrbracket$ of type

$$\llbracket \sigma(\tau_1) \rrbracket \times \dots \times \llbracket \sigma(\tau_n) \rrbracket \rightarrow \{\perp, \top\}$$

where σ maps the type variables α to the sorts $\mathbf{s}$. For each sort $\mathbf{s}$, the predicate $=\langle\mathbf{s}\rangle$ is interpreted as the equality over $\llbracket \mathbf{s} \rrbracket$.

- For any algebraic data type $t\langle\alpha\rangle$, with constructors $f_1\langle\alpha\rangle, \dots, f_l\langle\alpha\rangle$, we require $\llbracket t\langle\mathbf{s}\rangle \rrbracket$ to be the free algebra generated by $\llbracket f_1\langle\mathbf{s}\rangle \rrbracket, \dots, \llbracket f_l\langle\mathbf{s}\rangle \rrbracket$, that is

$$\text{for } i \neq j, \llbracket f_i\langle\mathbf{s}\rangle \rrbracket(\mathbf{t}) \neq \llbracket f_j\langle\mathbf{s}\rangle \rrbracket(\mathbf{u}) \quad (1)$$

$$\llbracket f_i\langle\mathbf{s}\rangle \rrbracket(\mathbf{t}) = \llbracket f_i\langle\mathbf{s}\rangle \rrbracket(\mathbf{u}) \Rightarrow \mathbf{t} = \mathbf{u} \quad (2)$$

$$\forall x \in \llbracket t\langle\mathbf{s}\rangle \rrbracket, \exists i, \exists \mathbf{t}, x = \llbracket f_i\langle\mathbf{s}\rangle \rrbracket(\mathbf{t}) \quad (3)$$

In the following, isf_i denotes a predicate that identifies values in $\llbracket t\langle\mathbf{s}\rangle \rrbracket$ that are applications of f_i , and $\text{proj}_{f_i,j}$ returns the j -th argument of such an application.

Given a pre-interpretation, a *valuation* v maps each type variable α to a sort $v(\alpha)$, and each variable x_τ to some element of $\llbracket v(\tau) \rrbracket$. Given a pre-interpretation and a valuation v , a term t of type τ is interpreted as an element $\llbracket t \rrbracket_v \in \llbracket v(\tau) \rrbracket$ and a formula f is interpreted as a Boolean value $\llbracket f \rrbracket_v$, according to the definitions given in Fig. 8–9. Note that pattern matching is compiled away as show in Fig. 10. Operator M turns a `match` construct into elementary tests. More precisely, $M(t, p, b, h)$ filters value t against pattern p and returns b in case of success and h in case of failure. Since type checking ensures that any pattern matching is exhaustive, the *error* term cannot actually appear in the result of M . It is only used to simplify the definition and use of function M .

An *interpretation* is a pre-interpretation that is consistent with recursive and inductive definitions, that is:

- For any recursive definition **function** $f\langle\alpha\rangle(\mathbf{x}) : \tau = t$ and any $\mathbf{s}$, we require $\llbracket f\langle\mathbf{s}\rangle \rrbracket$ to be such that, for all $\mathbf{t}$, $\llbracket f\langle\mathbf{s}\rangle \rrbracket(\mathbf{t}) = \llbracket t \rrbracket_v$ where v maps the α to the $\mathbf{s}$ and the $\mathbf{x}$ to the $\mathbf{t}$ (and similarly for a predicate definition).
- For any inductive definition $p\langle\alpha\rangle(\tau) = f_1 \mid \dots \mid f_l$ and any $\mathbf{s}$, we require $\llbracket p\langle\mathbf{s}\rangle \rrbracket$ to be the least predicate such that $\llbracket f_1 \rrbracket_v, \dots, \llbracket f_l \rrbracket_v$ hold where v maps the α to the $\mathbf{s}$.

$$\begin{aligned}
\llbracket n_{\text{int}} \rrbracket_v &= n \\
\llbracket r_{\text{real}} \rrbracket_v &= r \\
\llbracket x_\tau \rrbracket_v &= v(x_\tau) \\
\llbracket f\langle \tau_1, \dots, \tau_m \rangle(t_1, \dots, t_n) \rrbracket_v &= \llbracket f\langle v(\tau_1), \dots, v(\tau_m) \rangle \rrbracket(\llbracket t_1 \rrbracket_v, \dots, \llbracket t_n \rrbracket_v) \\
\llbracket \text{let } x_\tau = t_1 \text{ in } t_2 \rrbracket_v &= \llbracket t_2 \rrbracket_{v[x_\tau \mapsto \llbracket t_1 \rrbracket_v]} \\
\llbracket \text{if } f \text{ then } t_1 \text{ else } t_2 \rrbracket_v &= \llbracket t_1 \rrbracket_v \text{ if } \llbracket f \rrbracket_v = \top, \\
&= \llbracket t_2 \rrbracket_v \text{ otherwise} \\
\llbracket \text{match } t \text{ with } p_1 \rightarrow t_1 \mid \dots \mid p_n \rightarrow t_n \text{ end} \rrbracket_v &= \llbracket M(t, p_1, t_1, M(t, p_2, t_2, \\
&\quad \dots, M(t, p_n, t_n, \text{error}))) \rrbracket_v
\end{aligned}$$

Fig. 8. Interpretation of terms

$$\begin{aligned}
\llbracket p\langle \tau_1, \dots, \tau_m \rangle(t_1, \dots, t_n) \rrbracket_v &= \llbracket p\langle v(\tau_1), \dots, v(\tau_m) \rangle \rrbracket(\llbracket t_1 \rrbracket_v, \dots, \llbracket t_n \rrbracket_v) \\
\llbracket \forall x_\tau. f \rrbracket_v &= \\
\llbracket \exists x_\tau. f \rrbracket_v &= \\
\llbracket f_1 \circ f_2 \rrbracket_v &= \llbracket f_1 \rrbracket_v \circ \llbracket f_2 \rrbracket_v \text{ where } \circ \in \{\wedge, \vee, \Rightarrow, \Leftrightarrow\} \\
\llbracket \text{not } f \rrbracket_v &= \neg \llbracket f \rrbracket_v \\
\llbracket \text{true} \rrbracket_v &= \top \\
\llbracket \text{false} \rrbracket_v &= \perp \\
\llbracket \text{let } x_\tau = t_1 \text{ in } f_2 \rrbracket_v &= \llbracket f_2 \rrbracket_{v[x_\tau \mapsto \llbracket t_1 \rrbracket_v]} \\
\llbracket \text{if } f_1 \text{ then } f_2 \text{ else } f_3 \rrbracket_v &= \llbracket f_2 \rrbracket_v \text{ if } \llbracket f_1 \rrbracket_v = \top, \\
&= \llbracket f_3 \rrbracket_v \text{ otherwise} \\
\llbracket \text{match } t \text{ with } p_1 \rightarrow f_1 \mid \dots \mid p_n \rightarrow f_n \text{ end} \rrbracket_v &= \llbracket M(t, p_1, f_1, M(t, p_2, f_2, \\
&\quad \dots, M(t, p_n, f_n, \text{error}))) \rrbracket_v
\end{aligned}$$

Fig. 9. Interpretation of formulas

$$\begin{aligned}
M(t, x_\tau, b, h) &= \text{let } x_\tau = t \text{ in } b \\
M(t, f_i(), b, h) &= \text{if } \text{isf}_i(t) \text{ then } b \text{ else } h \\
M(t, f_i(p_1, \dots, p_n), b, h) &= \text{if } \text{isf}_i(t) \text{ then } M(\text{projf}_{i,1}(t), p_1, M(\text{projf}_{i,2}(t), p_2, \dots, h), h) \text{ else } h \\
M(t, -, b, h) &= b \\
M(t, p_1 \mid p_2, b, h) &= M(t, p_1, b, M(t, p_2, b, h)) \\
M(t, p \text{ as } x_\tau, b, h) &= \text{let } x_\tau = t \text{ in } M(t, p, b, h)
\end{aligned}$$

Fig. 10. Interpretation of patterns

A set of closed formulas Δ is *satisfied* by $\llbracket \cdot \rrbracket$ if and only if $\llbracket f \rrbracket_v = \top$ for every $f \in \Delta$ and every valuation v . (Note that only the mapping of type variables in v is relevant here, since formulas are closed.) A set of closed formulas Δ is *satisfiable* if and only if it is satisfied for some interpretation. Given a set of polymorphic closed axioms Δ and a closed formula f to be proved, we may assume that f is monomorphic (otherwise we simply replace type variables in f by fresh type constants). Then we say that formula f is a *logical consequence* of Δ if and only if the set $\Delta, \text{not } f$ is unsatisfiable.

Discussion. It is worth pointing out that our notion of interpretation does not prevent two distinct instances of a polymorphic symbol, say two types $t\langle s_1 \rangle$ and $t\langle s_2 \rangle$ or two functions $f\langle s_1 \rangle$ and $f\langle s_2 \rangle$, to be interpreted in two completely different ways. Even for an algebraic data type, say $\text{list}\langle \alpha \rangle$, we do not require $\text{Nil}\langle s_1 \rangle$ and $\text{Nil}\langle s_2 \rangle$ to be identical. We simply require $\text{list}\langle s_1 \rangle$ to be the free algebra generated by $\text{Nil}\langle s_1 \rangle$ and $\text{Cons}\langle s_1 \rangle$, and similarly $\text{list}\langle s_2 \rangle$ to be the free algebra generated by $\text{Nil}\langle s_2 \rangle$ and $\text{Cons}\langle s_2 \rangle$.

As a consequence, there is nothing wrong with an axiom defining a property of $f\langle \text{int} \rangle$, for some polymorphic function $f\langle \alpha \rangle$, while all other instances are left uninterpreted. Even further, we could have two completely unrelated axioms for $f\langle \text{int} \rangle$ and $f\langle \text{real} \rangle$, *e.g.* that $f\langle \text{int} \rangle$ is the identity over type `int` while $f\langle \text{real} \rangle$ is the square root function.

Simply speaking, everything works as if we were using many-sorted logic with a possibly infinite set of simple sorts (`int`, $\text{list}\langle \text{int} \rangle$, $\text{list}\langle \text{real} \rangle$, $\text{list}\langle \text{list}\langle \text{int} \rangle \rangle$, etc.), and a possibly infinite set of functions and predicates with simple types ($\text{Nil}\langle \text{int} \rangle$, $\text{Nil}\langle \text{real} \rangle$, $\text{Cons}\langle \text{list}\langle \text{int} \rangle \rangle$, etc.), with suitable extensions for algebraic data types, recursive definitions, and inductive predicates.

2.4 Implementation

The logic we just described is implemented in Why3 at two different levels: an OCaml API and a surface language.

The OCaml API allows the user to build terms, patterns, formulas, declarations, and goals. The API is defensive: only well-typed values can be built, according to the typing rules of Sec. 2.2. Following the presentation above, one first builds symbols and then, later, possible definitions for these symbols. This way, terms, patterns, and formulas can be built as soon as symbols are available, without the need of passing around a context containing definitions.

On top of this API, Why3 provides a surface language. Fig. 11 contains an example of input file. Contrary to the API, when a definition is provided, we do not separate the signature and the definition. For instance, we simply write

```
type list 'a = Nil | Cons 'a (list 'a)
```

to simultaneously introduce the type symbol `list 'a`, the two function symbols `Nil` and `Cons` (its constructors), and the algebraic data type definition.

```

theory Example

  type nat = 0 | S nat

  type list 'a = Nil | Cons 'a (list 'a)

  function length (l: list 'a) : nat =
    match l with
    | Nil → 0
    | Cons _ r → S (length r)
    end

  inductive even (nat) =
    | Even0: even 0
    | EvenS: forall m: nat. even m → even (S (S m))

  goal G: even (length (Cons 0 (Cons 0 Nil)))

end

```

Fig. 11. Example of Why3 input file

In the surface language, there are no angle brackets $\langle . \rangle$ anymore. First, type variables are not explicitly bound in symbol declarations. They are rather gathered, as the set of all type variables appearing in the symbol's type. For instance, the declaration

```
function length (l: list 'a) : nat = ...
```

introduces a function symbol `length` with one type variable. Second, type variables are not explicitly instantiated when a symbol is used. Instantiation is inferred from the arguments whenever possible. For instance, one simply writes

```
length (Cons 1 (Cons 2 Nil))
```

without having to pass $\langle \text{int} \rangle$ to `Nil`, `Cons`, and `length`. When the instantiation cannot be computed from the arguments, it must be provided. A typical example is the following:

```
lemma L: length (Nil: list 'a) = 0
```

Without the type cast, this would result in an **undefined type variable** error. In both the API and the surface language, there is no type inference, only type checking. In particular, quantified variables are given types (such as `forall m: nat` in Fig. 11). This is a deliberate choice: we could implement type inference, but we think that specifications are easier to read when types are given.

Among other features of the surface language of Why3 are type aliases (such as `type t = list int`) and tuple types. Type aliases are inlined systematically.

Tuple types are particular cases of algebraic data types, and are generated on the fly. There is also a syntax for record types (algebraic data types with a single constructor and named projection functions). Finally, logic declarations are organized in units called *theories*, as shown in Fig. 11. Theories can refer to and instantiate other theories. We do not discuss theories here; see [10] for details.

3 Targeting Multiple Provers

We now explain how Why3 translates the logic we just described into the native input format of external theorem provers. Currently, Why3 supports the following theorem provers:

- Proof assistants: Coq 8.4 [41] and PVS 6.0 [36];
- SMT solvers: Alt-Ergo 0.95.1 [7], CVC3 2.4.1 [3], CVC4 1.0 [2], Simplify 1.5.4 [19], Yices 1.0.38 [18] and Yices2 2.0.4, Z3 4.3.1 [17];
- ATP systems: E 1.6 [40], iProver 0.8.1 [27], SPASS 3.7 [42], Vampire 0.6 [38], Zenon 0.7.1 [13];
- Dedicated provers: Gappa 0.15.1 [16], Mathematica 8.0.

We only list the most recent versions that are supported; some older versions may be supported as well.

3.1 Tasks and Transformations

A central notion in Why3 is that of *task*. A task is a context Γ (symbols, possibly with definitions), a set of axioms Δ , and a formula to be proved. Tasks are massaged using *transformations* until they reach a subset of Why3’s logic that coincides with the input format of a theorem prover. A transformation turns a task into a set of new tasks. Key transformations are the following:

- elimination of recursion, inductive predicates, algebraic data types and pattern matching, **if-then-else** construct, **let** binding;
- encoding of types, to target many-sorted or untyped logic [15,11].

In the process of proving a task, other transformations may be used, such as inlining, splitting, various kinds of simplification, or even induction (following Leino [30]). Why3 can be extended with new transformations via OCaml plugins. The OCaml API allows the user to build tasks and to apply transformations efficiently (thanks to memoization).

3.2 Drivers

The way a particular prover is handled in Why3 is controlled by a text file called a *driver*. Such a file lists the transformations that must be applied to a task prior to its transmission to the prover, defines the pretty-printer that must be

used when we are done with transformations, lists symbols and axioms that are built-in in the prover, and provides regular expressions to interpret the output of calls to the prover.

For instance, the driver for SPASS refers to the printer registered under the name `tptp-fof`, includes transformations such as `eliminate_algebraic` (to get rid of algebraic data types and pattern matching) and `encoding_tptp` (to encode types), and states that equality is built-in with syntax $(x=y)$, among other things.

Users may define new drivers, to add support for a new prover or to experiment with alternative ways of using a prover (*e.g.* not making use of a built-in theory, or using an alternate input format). For that purpose, Why3 can be extended with new pretty-printers via OCaml plug-ins. The OCaml API provides ways to load a driver and then to use it to call the corresponding theorem prover on a task. Drivers are described in more details in our earlier work [10].

Prover Specificities. There are several places where we have to be careful to avoid introducing inconsistencies in the process of translating from Why3 to theorem provers.

An example is integer division. Why3 standard library contains two theories: one for Euclidean division, and another where division rounds towards zero as in most programming languages. When it comes to using a built-in notion of division provided by some prover, we have to identify which one it is, if any, or to provide workarounds otherwise, if possible. For instance, Z3 provides Euclidean division and modulo, but CVC4 provides a division that appears to be none of the two divisions from Why3 standard library.

Another example is the fact that all types in Why3 are inhabited. Automated theorem provers typically make that assumption already. When it comes to proof assistants Coq and PVS, however, it is not granted for free. PVS provides a built-in mechanism for non-empty types, that we use readily. Coq, on the contrary, does not provide any such facility. We resort to type classes to ensure that all types we manipulate are inhabited. For instance, function `length` is translated into

```
Fixpoint length {a:Type} {a_WT:WhyType a} (l:(list a)) ...
```

where `WhyType a` is a type class that states that type `a` is inhabited and has decidable equality.

3.3 Proof Sessions

When a verification tool is built on top of a single automated prover, it is straightforward to replay a proof: one simply reruns the tool on the input files. The same applies for a proof that has been built interactively, say with Coq or KIV, and stored into files: it can be rechecked easily, in batch mode.

With Why3, the situation is somewhat different. Using the graphical user interface `why3ide`, the user interactively applies transformations and calls external theorem provers, until all goals are discharged. Calls to provers record

the maximal amount of CPU time and memory that is allocated to the prover. Calls to proof assistants record user-edited proof scripts. All that information is stored into a set of files called a *proof session*. This way, it can be reloaded on a subsequent call to `why3ide` or replayed in a batch mode with another tool, `why3replayer`. This is of particular interest when the user upgrades the version of one or several provers. A run of `why3replayer` then tells whether the proof session is still valid or must be updated. Proof sessions are the subject of another publication [9].

3.4 Examples

We now give some examples of the way goals are translated to be passed to theorem provers. Let us consider our running example in Fig. 11 and let us assume that we are targeting the SMT solver Alt-Ergo.

Algebraic Data Types. We must get rid of algebraic data types. Thus, `list 'a` is simply turned into some uninterpreted data type, together with uninterpreted function symbols for constructors

```
type 'a list
logic Nil : 'a list
logic Cons : 'a, 'a list → 'a list
```

and axioms (not shown here) to state that `Nil` and `Cons` are distinct and that `Cons` is injective. To get rid of pattern matching on type `list`, we also introduce a function symbol

```
logic match_list : 'a list, 'a1, 'a1 → 'a1
```

together with axioms stating that $\text{match_list}(\text{Nil}, a, b) = a$ and that $\text{match_list}(\text{Cons}(x, y), a, b) = b$. We proceed similarly for type `nat`.

Recursive Definitions. We must also get rid of the recursive definition of function `length`. It is turned into some uninterpreted function symbol together with two axioms:

```
logic length : 'a list → nat
axiom length_def : length(Nil : 'a list) = 0
axiom length_def1 : forall x:'a. forall x1:'a list.
  length(Cons(x, x1)) = S(length(x1))
```

Note that we are not making use of function `match_list` here, since it immediately reduces to simpler axioms.

Inductive Definitions. Finally, we must get rid of the inductive definition of predicate `even`. It is axiomatized as follows:

```

logic even : nat → prop
axiom Even0 : even(0)
axiom EvenS : forall m:nat. even(m) → even(S(S(m)))
axiom even_inversion :
  forall z:nat. even(z) →
    (z = 0 or exists m:nat. even(m) and z = S(S(m)))

```

This is incomplete, since we cannot capture the minimality of `even` in first-order logic. For instance, Alt-Ergo cannot prove the goal `forall x: nat. even x → not (even (S x))`, while we could use Coq to prove it. Anyway, Alt-Ergo easily manages to prove the goal

```

goal G: even(length(Cons(0, Cons(0, (Nil : nat list)))))

```

with suitable instantiations of lemmas `length_def`, `length_def1`, `Even0`, and `EvenS`.

Types. As we see on the example above, we did not have to encode types. Indeed, Alt-Ergo supports polymorphic types [8]. If we are instead targeting a prover that only supports simple types (*e.g.* Z3) or some untyped logic (*e.g.* SPASS), we have to encode the polymorphic types of Why3 in some way or another. For instance, on the following Why3 input

```

type list 'a
constant nil: list 'a
function length (list 'a) : int
axiom length_nil: length (nil: list 'a) = 0
goal G: length (nil: list int) = 0

```

the file that is passed to SPASS is the following:

```

fof(length_nil, axiom,
  ![A]: sort(int, length(sort(list(A),nil))) = sort(int,const_0)).

fof(g, conjecture,
  sort(int, length(sort(list(int), nil))) = sort(int,const_0)).

```

Table 1. Comparing eight automated provers using Why3

prover	proved	max. time	avg. time
CVC3 (2.4.1)	2203	21.00	0.17
Alt-Ergo (0.95.1)	2202	29.73	0.16
CVC4 (1.0)	2071	12.00	0.09
Z3 (4.3.1)	1869	45.52	0.11
Yices (1.0.38)	1634	4.30	0.05
Vampire (0.6)	1375	27.72	0.51
E (1.6)	1303	19.73	0.41
Spass (3.7)	1185	23.78	0.52

Here `sort` is a binary function symbol that wraps terms with types. Types themselves are represented as regular terms, such as constant `int` or variable `A` above. The case of an SMT solver is more subtle, as we need to protect built-in types — such as integers, arrays, or reals — if we want to use built-in decision procedures [11].

4 Experimental Results

In this section, we use Why3 to run a small benchmark of 8 different automated provers. The experiment uses 83 proof sessions from our gallery, corresponding to logical theories or programs that were all successfully proved. This includes our solutions to several recent competitions in program verification (VSTTE 2012, FoVeOOS 2011, VSTTE 2012, FM 2012). The total number of subgoals is 2849.

For each subgoal that was discharged by at least one prover, we run the other 7 provers on that subgoal, with the same limit of CPU time that was given to the first one. Provers are run on an 8-cores 3.20GHz Intel Xeon with 24Gb of RAM. Each prover runs on a separate core, with a limit of 1Gb of memory. Results are given in Table 1. For each prover, we give the total number of goals proved, and the maximum/average running time per goal.

The purpose of that table is not really to compare provers, but rather to show the benefits of a collaborative use of several provers: if we were using CVC3 only, we would be left with 646 ($= 2849 - 2203$) unproved subgoals. Besides, it is worth pointing out that most of these goals involve arithmetic; yet provers with no support for arithmetic (E, SPASS, and Vampire, in that case) are able to discharge a large subset of the goals. This was rather unexpected and encourages us to increase our daily use of these provers and to improve the way we use them even further.

5 Related Work

There is actually very little in the literature regarding the extension of first-order logic with polymorphism. For instance, a classical textbook such as Manzano’s [32] does not contain a single occurrence of the word ‘polymorphism’. On the other side of the logical spectrum, rich logics such as the Calculus of Inductive Constructions or HOL do have polymorphism (even beyond rank-1) but are seldom interested in identifying their first-order fragments. Among the recent work on this subject, we can mention the work by Leino and Rümmer on Boogie 2’s type system [28] and the definition of the TPTP TFF1 format by Blanchette and Paskevich [6].

There is more related work regarding the second part of this paper, as a lot has been done in the context of Isabelle’s Sledgehammer tactic [34,12]. It translates Isabelle’s logic to several external provers, using type encodings different from ours, ranging from mere type erasure (which is unsound, but Sledgehammer uses proof reconstruction²) to partial monomorphisation [5] (which is proved

² A significant difference between Sledgehammer and our work is that we do not perform any kind of proof reconstruction. Thus we have to keep to sound encodings.

incomplete [11] but seems efficient in practice anyway). Earlier, Hurd had already investigated encodings from higher-order to first-order logic [24].

6 Conclusion and Perspectives

We have presented a logic whose purpose is to provide a unified front-end to many existing theorem provers, either interactive or automated, and its implementation in Why3. We have designed this logic with the idea that specification must be as natural as possible and that tools should adapt themselves to proof obligations (and not the opposite). Using a wide range of theorem provers encourages this attitude.

One of the key features of our logic is polymorphism. It is defined and handled roughly the same way it is in tools Boogie, Sledgehammer, and Alt-Ergo. There is no competition, but rather a nice convergence. A contribution of this paper is to add the formalization of algebraic data types, recursive definitions, and inductive predicates on top of that.

Currently, Why3 is successfully used as a sub-component in various projects [23,33,4,14], as well as the vehicle for many non-trivial case studies in program verification (see for instance [21]). We also envision that verification environments that are currently built on top of a dedicated prover (*e.g.* B, KIV, SmallFoot) could also benefit from additional, external theorem provers. For instance, we are currently experimenting with the use of Why3 to discharge goals obtained from Atelier B [35], and the first results are promising.

We could still improve the way we are using theorem provers. For instance, we can observe that some tools that are based on a single automated prover — *e.g.* VCC, Dafny, or VeriFast — are able to carry out impressive case studies. It is clear that these tools are achieving a level of intimacy with the prover that is beyond that of Why3. We should learn from these tools and transpose relevant techniques to Why3. In particular, we our support of built-in theories.

Of course, it would be much simpler if we had native support for polymorphic types in provers. Alt-Ergo demonstrates that this is possible, and even simple to implement [8], yet this is currently the only automated prover with such a feature. We hope that TFF1 [6] will encourage some ATP developers to take the plunge.

Acknowledgments. I would like to thank Andrei Paskevich and Sylvain Conchon for fruitful discussions during the preparation of this paper. Some ideas behind the formalization in Sec. 2 already appear in Bobot and Paskevich’s work [11]. The development of Why3 is joint work with François Bobot, Claude Marché, Guillaume Melquiond, and Andrei Paskevich.

References

1. Abrial, J.-R.: *The B-Book, assigning programs to meaning*. Cambridge University Press (1996)
2. Barrett, C., Conway, C.L., Deters, M., Hadarean, L., Jovanović, D., King, T., Reynolds, A., Tinelli, C.: CVC4. In: Gopalakrishnan, G., Qadeer, S. (eds.) *CAV 2011*. LNCS, vol. 6806, pp. 171–177. Springer, Heidelberg (2011)
3. Barrett, C.W., Tinelli, C.: CVC3. In: Damm, W., Hermanns, H. (eds.) *CAV 2007*. LNCS, vol. 4590, pp. 298–302. Springer, Heidelberg (2007)
4. Barthe, G., Grégoire, B., Heraud, S., Béguelin, S.Z.: Computer-aided security proofs for the working cryptographer. In: Rogaway, P. (ed.) *CRYPTO 2011*. LNCS, vol. 6841, pp. 71–90. Springer, Heidelberg (2011)
5. Blanchette, J.C., Böhme, S., Popescu, A., Smallbone, N.: Encoding monomorphic and polymorphic types. In: Piterman, N., Smolka, S.A. (eds.) *TACAS 2013*. LNCS, vol. 7795, pp. 493–507. Springer, Heidelberg (2013)
6. Blanchette, J.C., Paskevich, A.: TFF1: The TPTP typed first-order form with rank-1 polymorphism. In: Bonacina, M.P. (ed.) *CADE 2013*. LNCS (LNAI), vol. 7898, pp. 414–420. Springer, Heidelberg (2013)
7. Bobot, F., Conchon, S., Contejean, E., Iguernelala, M., Lescuyer, S., Mebsout, A.: *The Alt-Ergo automated theorem prover* (2008), <http://alt-ergo.lri.fr/>
8. Bobot, F., Conchon, S., Contejean, E., Lescuyer, S.: Implementing Polymorphism in SMT solvers. In: Barrett, C., de Moura, L. (eds.) *SMT 2008: 6th International Workshop on Satisfiability Modulo*. ACM International Conference Proceedings Series, vol. 367, pp. 1–5 (2008)
9. Bobot, F., Filliâtre, J.-C., Marché, C., Melquiond, G., Paskevich, A.: Preserving user proofs across specification changes. In: Cohen, E., Rybalchenko, A. (eds.) *Verified Software: Theories, Tools, Experiments (5th International Conference VSTTE)*, Atherton, USA, May 2013. LNCS, Springer (2013)
10. Bobot, F., Filliâtre, J.-C., Marché, C., Paskevich, A.: Why3: Shepherd your herd of provers. In: *Boogie 2011: First International Workshop on Intermediate Verification Languages*, Wrocław, Poland, pp. 53–64 (August 2011)
11. Bobot, F., Paskevich, A.: Expressing Polymorphic Types in a Many-Sorted Language. In: Tinelli, C., Sofronie-Stokkermans, V. (eds.) *FroCoS 2011*. LNCS, vol. 6989, pp. 87–102. Springer, Heidelberg (2011)
12. Böhme, S., Nipkow, T.: Sledgehammer: Judgement day. In: Giesl, J., Hähnle, R. (eds.) *IJCAR 2010*. LNCS, vol. 6173, pp. 107–121. Springer, Heidelberg (2010)
13. Bonichon, R., Delahaye, D., Doligez, D.: Zenon: An extensible automated theorem prover producing checkable proofs. In: Dershowitz, N., Voronkov, A. (eds.) *LPAR 2007*. LNCS (LNAI), vol. 4790, pp. 151–165. Springer, Heidelberg (2007)
14. Comar, C., Kanig, J., Moy, Y.: Integrating formal program verification with testing. In: *Proceedings of the Embedded Real Time Software and Systems Conference, ERTS² 2012* (February 2012)
15. Couchot, J.-F., Lescuyer, S.: Handling polymorphism in automated deduction. In: Pfenning, F. (ed.) *CADE 2007*. LNCS (LNAI), vol. 4603, pp. 263–278. Springer, Heidelberg (2007)
16. Daumas, M., Melquiond, G.: Certification of bounds on expressions involving rounded operators. *Transactions on Mathematical Software* 37(1), 1–20 (2010)
17. de Moura, L., Bjørner, N.S.: Z3, an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) *TACAS 2008*. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)

18. de Moura, L., Dutertre, B.: Yices: An SMT Solver, <http://yices.csl.sri.com/>
19. Detlefs, D., Nelson, G., Saxe, J.B.: Simplify: a theorem prover for program checking. *J. ACM* 52, 365–473 (2005)
20. Filliâtre, J.-C.: Deductive software verification. *International Journal on Software Tools for Technology Transfer (STTT)* 13(5), 397–403 (2011)
21. Filliâtre, J.-C.: Verifying two lines of C with Why3: an exercise in program verification. In: Joshi, R., Müller, P., Podelski, A. (eds.) *VSTTE 2012. LNCS*, vol. 7152, pp. 83–97. Springer, Heidelberg (2012)
22. Filliâtre, J.-C., Paskevich, A.: Why3 — where programs meet provers. In: Felleisen, M., Gardner, P. (eds.) *Programming Languages and Systems. LNCS*, vol. 7792, pp. 125–128. Springer, Heidelberg (2013)
23. The Frama-C platform for static analysis of C programs (2008), <http://www.frama-c.cea.fr/>
24. Hurd, J.: An lcf-style interface between hol and first-order logic. In: Voronkov, A. (ed.) *CADE 2002. LNCS (LNAI)*, vol. 2392, pp. 134–138. Springer, Heidelberg (2002)
25. Jacobs, B., Piessens, F.: The VeriFast program verifier. *CW Reports CW520*, Department of Computer Science, K.U.Leuven (August 2008)
26. Klein, G., Andronick, J., Elphinstone, K., Heiser, G., Cock, D., Derrin, P., Elkaduwe, D., Engelhardt, K., Kolanski, R., Norrish, M., Sewell, T., Tuch, H., Winwood, S.: seL4: Formal verification of an OS kernel. *Communications of the ACM* 53(6), 107–115 (2010)
27. Korovin, K.: iProver – an instantiation-based theorem prover for first-order logic (system description). In: Armando, A., Baumgartner, P., Dowek, G. (eds.) *IJCAR 2008. LNCS (LNAI)*, vol. 5195, pp. 292–298. Springer, Heidelberg (2008)
28. Leino, K.R.M., Rümmer, P.: A polymorphic intermediate verification language: Design and logical encoding. In: Esparza, J., Majumdar, R. (eds.) *TACAS 2010. LNCS*, vol. 6015, pp. 312–327. Springer, Heidelberg (2010)
29. Leino, K.R.M.: Dafny: An Automatic Program Verifier for Functional Correctness. In: Clarke, E.M., Voronkov, A. (eds.) *LPAR-16 2010. LNCS*, vol. 6355, pp. 348–370. Springer, Heidelberg (2010)
30. Leino, K.R.M.: Automating induction with an SMT solver. In: Kuncak, V., Rybalchenko, A. (eds.) *VMCAI 2012. LNCS*, vol. 7148, pp. 315–331. Springer, Heidelberg (2012)
31. Leroy, X.: A formally verified compiler back-end. *Journal of Automated Reasoning* 43(4), 363–446 (2009)
32. Manzano, M.: *Extensions of first order logic*. Cambridge University Press, New York (1996)
33. Marché, C., Paulin-Mohring, C., Urbain, X.: The KRAKATOA tool for certification of JAVA/JAVACARD programs annotated in JML. *Journal of Logic and Algebraic Programming* 58(1–2), 89–106 (2004), <http://krakatoa.lri.fr>
34. Meng, J., Paulson, L.C.: Translating higher-order clauses to first-order clauses. *Journal of Automated Reasoning* 40, 35–60 (2008)
35. Mentré, D., Marché, C., Filliâtre, J.-C., Asuka, M.: Discharging proof obligations from Atelier B using multiple automated provers. In: Derrick, J., Fitzgerald, J., Gnesi, S., Khurshid, S., Leuschel, M., Reeves, S., Riccobene, E. (eds.) *ABZ 2012. LNCS*, vol. 7316, pp. 238–251. Springer, Heidelberg (2012), <http://hal.inria.fr/hal-00681781/en/>
36. Owre, S., Rushby, J.M., Shankar, N.: PVS: A prototype verification system. In: Kapur, D. (ed.) *CADE 1992. LNCS*, vol. 607, pp. 748–752. Springer, Heidelberg (1992)

37. Reif, W., Schnellhorn, G., Stenzel, K.: Proving system correctness with KIV 3.0. In: McCune, W. (ed.) CADE 1997. LNCS, vol. 1249, pp. 69–72. Springer, Heidelberg (1997)
38. Riazanov, A., Voronkov, A.: Vampire. In: Ganzinger, H. (ed.) CADE 1999. LNCS (LNAI), vol. 1632, pp. 292–296. Springer, Heidelberg (1999)
39. Schulte, W., Xia, S., Smans, J., Piessens, F.: A glimpse of a verifying C compiler, <http://www.cs.ru.nl/~tews/cv07/cv07-smans.pdf>
40. Schulz, S.: System description: E 0.81. In: Basin, D., Rusinowitch, M. (eds.) IJCAR 2004. LNCS (LNAI), vol. 3097, pp. 223–228. Springer, Heidelberg (2004)
41. The Coq Development Team. The Coq Proof Assistant Reference Manual – Version V8.3 (2010), <http://coq.inria.fr>
42. Weidenbach, C., Dimova, D., Fietzke, A., Kumar, R., Suda, M., Wischniewski, P.: SPASS version 3.5. In: Schmidt, R.A. (ed.) CADE 2009. LNCS, vol. 5663, pp. 140–145. Springer, Heidelberg (2009)