

Expressing Polymorphic Types in a Many-Sorted Language

François Bobot^{1,2} and Andrei Paskevich^{1,2}

¹ LRI, Université Paris-Sud 11, CNRS, Orsay F-91405

² INRIA Saclay-Île de France, ProVal, Orsay F-91893

Abstract. In this paper, we study translation from a first-order logic with polymorphic types à la ML (of which we give a formal description) to a many-sorted or one-sorted logic as accepted by mainstream automated theorem provers. We consider a three-stage scheme where the last stage eliminates polymorphic types while adding the necessary “annotations” to preserve soundness, and the first two stages serve to protect certain terms so that they can keep their original unannotated form. This protection allows us to make use of provers’ built-in theories and operations. We present two existing translation procedures as sound and complete instances of this generic scheme. Our formulation generalizes over the previous ones by allowing us to protect terms of arbitrary monomorphic types. In particular, we can benefit from the built-in theory of arrays in SMT solvers such as Z3, CVC3, and Yices. The proposed methods are implemented in the Why3 tool and we compare their performance in combination with several automated provers.

1 Introduction

Polymorphic types are a means of abstraction over families of different types; a polymorphic definition or proposition stands for a potentially infinite number of its type-specific instances. Type systems employing polymorphism arise naturally in programming languages and they are a prominent feature of interactive proof assistants such as Coq [17] or Isabelle/HOL [16].

However, a proof task written in a language with polymorphic types is today a difficult subject for automation. This is not because polymorphism handling in a prover is complicated or inefficient *per se*. As was demonstrated by the Alt-Ergo project [3], this only requires a straightforward extension of the unification procedure and does not impose any significant overhead. The fact is, advanced type systems have not yet become mainstream in automated deduction: SMT solvers use many-sorted languages such as SMT-LIB [1], and TPTP provers are content with one-sorted first-order language. Thus, to apply a mainstream prover to a problem expressed in a polymorphic language, we have to translate it into an equivalent monomorphic or even one-sorted problem.

The challenge is not new and a number of solutions is known, ranging from adding per-variable “type guards” (also known as “relativisation of quantifiers”, see [12] and [11, Sect. 3.0]), to throughout decoration of terms with their types

[9, 6], to various flavours of type erasure [10, 14, 11]. The latter method is logically unsound, though adding type annotations can prevent certain unsound inference steps (see [14, Sect. 2.5,2.6] and [11, Sect. 3.1]).

An important feature of a polymorphism encoding method is special treatment of types and operations that are directly handled by provers’ built-in decision procedures, e.g., for linear arithmetic or bit-vectors. The idea is to prevent the terms that can be interpreted by a prover from being modified by translation, to preserve their original form [6, 11]. In what follows, we call this “type protection” to emphasize that we are interested in terms of particular types.

In this work, we aim to lift (or at least work around) several limitations we perceive in the previous approaches. Firstly, the existing type protection techniques only handle “simple” types, like integers or booleans, but not instances of polymorphic types, like lists of integers or arrays of reals. Yet decision procedures for such “complex” types are implemented in some SMT solvers; for example, Z3 [15], CVC3 [2], and Yices [7] have a built-in support for arrays. Secondly, type protection, as defined in [6], cannot be used to protect finite types such as booleans: given an axiom “every boolean is equal either to ‘true’ or to ‘false’”, one can derive that there are only two values in any encoded type, which can easily lead to a contradiction. Thirdly, while translation by type erasure with addition of type arguments to polymorphic symbols [11, Sect. 3.1] is less intrusive and more efficient than full term decoration [6], the former method is unsound and, according to [11], is only applicable in combination with provers that use trigger-based rather than unification-based instantiation. Such a requirement excludes the superposition-based provers and may be difficult to test when a third-party prover is used.

We begin with a formal presentation of first-order logic with polymorphic types (Section 2). In particular, we show that complete monomorphisation is undecidable, that is, we cannot effectively compute a finite set of monomorphic instances of a polymorphic formula F that is equisatisfiable to F . Then we introduce a generic three-stage scheme of polymorphism encoding (Section 3). In this scheme, we start by replacing interpreted polymorphic symbols (such as operations of access and update in arrays) with selected monomorphic instances. The translation proceeds then to type protection, which we consider as a separate transformation, and concludes with polymorphism elimination proper.

We present a sound and complete method of type instantiation with symbol discrimination for the first stage (Section 3.1). Furthermore, we give a generalized formulation of the type protection method from [6], free from the aforesaid restrictions (Section 3.2). As third-stage transformations, we consider full term decoration from [9, 6] (Section 3.3) and type erasure with added type annotations from [11, Sect. 3.1] (Section 3.4). We show the latter method to be sound on problems that admit models with infinite domains for every non-protected type and we discuss how this condition can be handled in practice.

We conclude by comparing the described techniques in combination with the SMT solvers Z3, CVC3, and Yices [7] on a set of about 4100 proof obligations in the Why3 tool [4] (Section 4).

2 First-Order Logic With Polymorphic Types

The logic \mathbf{FOL}_T , presented below, is an extension of classical first-order many-sorted logic. In \mathbf{FOL}_T , types are built from type constants (such as “integer”), type functions (such as “list of”), and type variables that stand for arbitrary monomorphic types. We do not admit quantifiers over type variables, neither in types, nor in formulas: a polymorphic formula is rather seen as a scheme, a potentially infinite conjunction of its monomorphic type instances. In other words, every type variable that occurs in a formula is bound by an implicit prenex universal quantifier. Basically, we use type polymorphism as a convenient way to write a set of polymorphic axioms — say, for lists or arrays — once, instead of copying them for every particular instance of these types.

The principal purpose of \mathbf{FOL}_T is to help specify and prove programs and its type system can be seen as the first-order fragment of the ML type system. The Why3 verification tool [4] is based upon \mathbf{FOL}_T with some extensions such as algebraic types. The papers [6] and [11] work in a similar setting, though the latter employs explicit quantifiers over type variables in logic formulas.

Syntax. We define types as syntactical expressions built from *type constructors* of fixed arity (denoted with capital sans-serif letters) and *type variables* (denoted α, β, γ). For example, β , l , $F(l, \gamma)$ are well-formed types. Type constructors of arity 0 are called *type constants*. A type that contains no type variables is called *monomorphic type* or *sort*. A vector of types $\langle T_1, \dots, T_n \rangle$ is called *type signature*.

A *type substitution* is a mapping from type variables to types. A *monomorphic* type substitution maps every type variable either to itself or to a sort. A type T is said to *match* another type T' whenever there is a type substitution that instantiates T to T' . This notion is trivially extended to type signatures.

We use letters S and T for types, boldface letters \mathbf{S} , \mathbf{T} for type signatures, and Greek letters τ , θ , and π for type substitutions. We denote the set of available type variables with \mathbb{V}_T , the set of type constructors with \mathbb{F}_T , the set of all types built from \mathbb{V}_T and \mathbb{F}_T with $\mathcal{T}(\mathbb{F}_T, \mathbb{V}_T)$, and the set of all sorts with $\mathcal{T}(\mathbb{F}_T)$. We presume that \mathbb{V}_T is infinite and \mathbb{F}_T contains at least one type constant.

We use traditional first-order terms and formulas, built from variable symbols (denoted u, v, w), function symbols (denoted f, g, h), and predicate symbols (denoted p, q), with the following additions:

- Every term carries an explicit type, e.g.: $w : C(l)$, $f(u : \alpha, v : L(\alpha)) : L(\alpha)$. We denote terms with letters s and t , and, by abuse of notation, we sometimes write the type of a term to the right of the letter: $s : T_1$, $t : T_2$, and so on.
- A *variable* is a variable symbol with a type, and we treat $w : C(l)$ and $w : C(\alpha)$ as two distinct variables even though they share the same variable symbol.
- To each function symbol of arity n we assign a type signature of length $n + 1$. A term of the form $f(t_1 : T_1, \dots, t_n : T_n) : T$ is well-formed if and only if the type signature of f matches $\langle T_1, \dots, T_n, T \rangle$.
- To each predicate symbol of arity n we give a type signature of length n . An atomic formula of the form $p(t_1 : T_1, \dots, t_n : T_n)$ is well-formed if and only if the type signature of p matches $\langle T_1, \dots, T_n \rangle$.

- An atomic equality formula of the form $t_1 \approx t_2$ is well-formed if and only if the terms t_1, t_2 have the same type.
- Quantifiers bind variables, i.e., typed variable symbols: $\forall(u : \alpha) p(u : \alpha, u : C)$. Here, the first argument of p is bound, but the second one is free.

We treat equality (\approx), negation (\neg), conjunction (\wedge), and the universal quantifier (\forall) as logical symbols and we treat disjunction (\vee), implication (\supset), equivalence (\equiv), disequality ($\not\approx$), and the existential quantifier (\exists) as abbreviations.

We use letters x, y, z for variables, letters F, G, H for formulas, and Greek letters Γ, Δ for sets of formulas. We denote the (infinite) set of variable symbols with \mathbb{V} , the set of function symbols with \mathbb{F} , and the set of predicate symbols with \mathbb{P} . Given a term or a formula e , the set of type variables occurring in e is denoted $\text{FV}_{\mathbb{T}}(e)$ and the set of free variables of e is denoted $\text{FV}(e)$. If $\text{FV}_{\mathbb{T}}(e)$ is empty, we call e *monomorphic*. If $\text{FV}(e)$ is empty, we call e *closed* or *ground*.

Substitutions, denoted with letters σ and δ , apply to a term or a formula e , replacing free variables with terms of the same type (denoted $e\sigma$). The symbol \circ denotes the composition of two (type) substitutions: $x(\sigma \circ \delta) \triangleq x\sigma\delta$.

Type substitutions apply only to closed formulas and ground terms; also, we require type instantiation to rename every bound variable symbol to some fresh variable symbol. In this way, we avoid variable collisions: for example, the type substitution $[l/\alpha]$ would not instantiate the formula $\forall(u : \alpha)\forall(u : l)p(u : \alpha, u : l)$ to $\forall(u : l)\forall(u : l)p(u : l, u : l)$, but to $\forall(u' : l)\forall(u'' : l)p(u' : l, u'' : l)$. In our subsequent examples, we will not use a variable symbol in two different variables in the same formula to avoid confusion.

In what follows, we illustrate our transformations on the following simple polymorphic formula (for the sake of readability, we omit the most obvious type annotations): $\forall(m : \mathbb{M}(\alpha, l))\forall(c : \alpha)\text{get}(\text{set}(m, c, 6) : \mathbb{M}(\alpha, l), c) : l * 7 \approx 42$. Here, the type l represents integers and the type $\mathbb{M}(\alpha, \beta)$ is that of polymorphic α -to- β maps. The function symbol **get** is of type signature $\langle \mathbb{M}(\alpha, \beta), \alpha, \beta \rangle$ and **set** is of type signature $\langle \mathbb{M}(\alpha, \beta), \alpha, \beta, \mathbb{M}(\alpha, \beta) \rangle$.

Satisfiability. Given sets $\mathbb{F}_{\mathbb{T}}, \mathbb{F}, \mathbb{P}$, an interpretation \mathfrak{J} is defined by three maps:

- to each sort $S \in \mathcal{T}(\mathbb{F}_{\mathbb{T}})$, we assign a non-empty *domain* $\mathcal{D}_S^{\mathfrak{J}}$;
- to each symbol $f \in \mathbb{F}$ and each vector of sorts $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$ matched by the type signature of f , we assign a function $f_{\mathbf{S}}^{\mathfrak{J}} : \mathcal{D}_{S_1}^{\mathfrak{J}} \times \dots \times \mathcal{D}_{S_n}^{\mathfrak{J}} \rightarrow \mathcal{D}_S^{\mathfrak{J}}$;
- to each symbol $p \in \mathbb{P}$ and each vector of sorts $\mathbf{S} = \langle S_1, \dots, S_n \rangle$ matched by the type signature of p , we assign a function $p_{\mathbf{S}}^{\mathfrak{J}} : \mathcal{D}_{S_1}^{\mathfrak{J}} \times \dots \times \mathcal{D}_{S_n}^{\mathfrak{J}} \rightarrow \{\top, \perp\}$, where \top and \perp stand for Boolean constants “true” and “false”, respectively.

We call *type valuation* a type substitution that instantiates every type variable in $\mathbb{V}_{\mathbb{T}}$ to a sort. Given a type valuation π , we call *variable valuation under π* a function that maps every variable $u : T$ to some element of $\mathcal{D}_T^{\mathfrak{J}\pi}$. We simply say *variable valuation* when the implied type valuation is known from the context or in a purely monomorphic setting, where every type is already closed.

Let π be a type valuation and ξ be a variable valuation under π . We evaluate terms and formulas according to the following equalities, where $\mathbf{t} : \mathbf{T}$ stands for

a vector of terms $t_1 : T_1, \dots, t_n : T_n$ and $\xi[u : T \mapsto a]$ is a valuation that coincides with ξ everywhere except $u : T$, which is mapped to a .

$$\begin{aligned}
\mathfrak{J}_{\pi, \xi}(u : T) &\triangleq \xi(u : T) & \mathfrak{J}_{\pi, \xi}(t_1 \approx t_2) &\triangleq (\mathfrak{J}_{\pi, \xi}(t_1) = \mathfrak{J}_{\pi, \xi}(t_2)) \\
\mathfrak{J}_{\pi, \xi}(f(\mathbf{t} : \mathbf{T}) : T) &\triangleq f_{\langle \mathbf{T}, T \rangle \pi}^{\mathfrak{J}}(\mathfrak{J}_{\pi, \xi}(\mathbf{t})) & \mathfrak{J}_{\pi, \xi}(\neg F) &\triangleq \neg \mathfrak{J}_{\pi, \xi}(F) \\
\mathfrak{J}_{\pi, \xi}(p(\mathbf{t} : \mathbf{T})) &\triangleq p_{\mathbf{T}\pi}^{\mathfrak{J}}(\mathfrak{J}_{\pi, \xi}(\mathbf{t})) & \mathfrak{J}_{\pi, \xi}(F \wedge G) &\triangleq \mathfrak{J}_{\pi, \xi}(F) \wedge \mathfrak{J}_{\pi, \xi}(G) \\
\mathfrak{J}_{\pi, \xi}(\forall(u : T)F) && &\triangleq \bigwedge_{a \in \mathcal{D}_{T\pi}^{\mathfrak{J}}} \mathfrak{J}_{\pi, \xi}[u : T \mapsto a](F)
\end{aligned}$$

It is easy to see that evaluation of a term or a formula e under $\mathfrak{J}_{\pi, \xi}$ does not depend on the values of π and ξ on (type) variables that do not occur in e . In what follows, when we evaluate closed or monomorphic expressions, we often omit the variable valuation or the type valuation, respectively.

Lemma 1. *For any closed formula F and type valuation π , $\mathfrak{J}_{\pi}(F) = \mathfrak{J}(F\pi)$.*

As we said above, we treat type variables as implicitly universally quantified at the top of a polymorphic formula. Thus, a closed formula F is *satisfied* by \mathfrak{J} if and only if $\mathfrak{J}_{\pi}(F)$ is true for every type valuation π . A closed formula is *satisfiable* if and only if it is satisfied by some interpretation, called a *model* of this formula. These definitions are trivially extended to sets of closed formulas. To prove a polymorphic formula G in a polymorphic context Γ , we take a type substitution τ that replaces all type variables in G with fresh type constants and show that the set $\Gamma, \neg G\tau$ is unsatisfiable. Generally speaking, the semantics of polymorphic formulas in $\mathbf{FOL}_{\mathbf{T}}$ is quite similar to that of first-order clauses, where the free variables are also implicitly universally quantified.

On monomorphic terms and formulas, our definitions correspond to the traditional many-sorted logic with disjoint sorts. Moreover, a trivial corollary of Lemma 1 is that F is satisfiable if and only if the set of all monomorphic type instances of F is satisfiable.

Computing monomorphic instances? A polymorphic formula can have infinitely many monomorphic type instances. But can't we find out, in finite time, all sorts that are potentially relevant to the problem and deal with a finite subset of instances, produced just with these sorts? On one hand, this resembles an attempt to pre-compute the relevant ground instances in a set of first-order formulas — a problem well known to be undecidable. On the other hand, type handling does not need to be as hard as proof search in general, and complete monomorphisation is often possible in programming languages (e.g., C++ templates).

Theorem 1. *There is no computable function that maps an arbitrary closed formula F to an equisatisfiable finite set of monomorphic type instances of F (notice that such a set always exists by compactness).*

Proof. It turns out that our type system is expressive enough to encode an undecidable theory, namely, combinatory logic. Consider the following signature:

$$\mathbb{F}_{\mathbf{T}} = \{ \mathbf{A}(\cdot, \cdot), \mathbf{S}, \mathbf{K} \} \quad \mathbb{F} = \{ \mathbf{A} : \langle \alpha, \beta, \mathbf{A}(\alpha, \beta) \rangle, \mathbf{S} : \langle \mathbf{S} \rangle, \mathbf{K} : \langle \mathbf{K} \rangle \} \quad \mathbb{P} = \{ \mathbf{R} : \langle \alpha, \beta \rangle \}$$

along with five axioms (for brevity, we omit some type annotations):

$$\begin{aligned}
& \forall(u:\alpha)\forall(v:\beta)\forall(w:\gamma) ((\mathbf{R}(u,v) \wedge \mathbf{R}(v,w)) \supset \mathbf{R}(u,w)) \\
& \forall(u:\alpha)\forall(v:\beta)\forall(w:\gamma) (\mathbf{R}(u,v) \supset \mathbf{R}(\mathbf{A}(u,w) : \mathbf{A}(\alpha,\gamma), \mathbf{A}(v,w) : \mathbf{A}(\beta,\gamma))) \\
& \forall(u:\alpha)\forall(v:\beta)\forall(w:\gamma) (\mathbf{R}(u,v) \supset \mathbf{R}(\mathbf{A}(w,u) : \mathbf{A}(\gamma,\alpha), \mathbf{A}(w,v) : \mathbf{A}(\gamma,\beta))) \\
& \forall(u:\alpha)\forall(v:\beta) \mathbf{R}(\mathbf{A}(\mathbf{A}(\mathbf{K},u),v) : \mathbf{A}(\mathbf{A}(\mathbf{K},\alpha),\beta), u : \alpha) \\
& \forall(u:\alpha)\forall(v:\beta)\forall(w:\gamma) \mathbf{R}(\mathbf{A}(\mathbf{A}(\mathbf{S},u),v), w) : \mathbf{A}(\mathbf{A}(\mathbf{S},\alpha),\beta), \gamma), \\
& \qquad \qquad \qquad \mathbf{A}(\mathbf{A}(u,w), \mathbf{A}(v,w)) : \mathbf{A}(\mathbf{A}(\alpha,\gamma), \mathbf{A}(\beta,\gamma))
\end{aligned}$$

Here the binary function symbol \mathbf{A} stands for term application, and the binary predicate symbol \mathbf{R} for CL-reducibility. Notice that every ground combinatory term is reflected in its type.

Now, if we were able to compute a finite set of potentially relevant *closed types* for an arbitrary reducibility problem in this theory, this would readily let us decide the problem itself, as we would thus obtain the set of potentially relevant *ground terms*. Since ground reducibility in CL is undecidable, complete monomorphisation in \mathbf{FOL}_T is undecidable, too. \square

3 Eliminating Polymorphic Types

Being unable to select just a relevant monomorphic subset of a polymorphic problem, we have to resort to some form of encoding, converting the polymorphic problem to an equisatisfiable monomorphic one. Such conversion inevitably implies merging many types into few sorts or just a single sort. This is undesirable if we target an automated prover equipped with special techniques (decision procedures, unification modulo, etc.) for particular types, such as integers, booleans or arrays. These types ought to be separated from the rest, protected against this “type fusion”, expelled from polymorphism in the problem.

To this purpose, we slightly extend our language in order to be able to select the terms that will keep their (monomorphic) type through polymorphism elimination. To every sort S in $\mathcal{T}(\mathbb{F}_T)$ we associate a new *protected sort* \bar{S} . The use of protected sorts is restricted: a protected sort can appear in the type signature of a symbol or as a type of a term, but it cannot occur under a type constructor or in the range of a type substitution. In other words, the only type that matches a protected sort \bar{S} is \bar{S} itself.

For example, $\mathbf{get}(v : \bar{\mathbf{M}}(\bar{\mathbf{l}}, \bar{\mathbf{l}}), c : \bar{\mathbf{l}}) : \bar{\mathbf{l}}$ is a malformed term, since the type signature of \mathbf{get} is $\langle \mathbf{M}(\alpha, \beta), \alpha, \beta \rangle$ and $\mathbf{M}(\alpha, \beta)$ does not match $\bar{\mathbf{M}}(\bar{\mathbf{l}}, \bar{\mathbf{l}})$. Similarly, the term $\mathbf{get}(v : \mathbf{M}(\bar{\mathbf{l}}, \bar{\mathbf{l}}), c : \bar{\mathbf{l}}) : \bar{\mathbf{l}}$ is malformed, because β does not match $\bar{\mathbf{l}}$ and also because $\mathbf{M}(\bar{\mathbf{l}}, \bar{\mathbf{l}})$ is a malformed type expression. However, if we consider a “protected specialization” of \mathbf{get} , denoted $\bar{\mathbf{get}}$, with the type signature $\langle \bar{\mathbf{M}}(\bar{\mathbf{l}}, \bar{\mathbf{l}}), \bar{\mathbf{l}}, \bar{\mathbf{l}} \rangle$, the application $\bar{\mathbf{get}}(v : \bar{\mathbf{M}}(\bar{\mathbf{l}}, \bar{\mathbf{l}}), c : \bar{\mathbf{l}}) : \bar{\mathbf{l}}$ is a well-formed term.

Concerning interpretation, every protected sort \bar{S} has its proper non-empty domain $\mathcal{D}_{\bar{S}}^{\bar{J}}$. As with any type substitution, we restrict type valuations to non-protected sorts. Thus, a set $\{\forall(u:\alpha)\forall(v:\alpha)u \approx v, \exists(a:\bar{I})\exists(b:\bar{I})a \not\approx b\}$ is satisfiable. Indeed, the first formula requires the domain of every non-protected sort to be a singleton, but does not constrain the domains of protected sorts.

Using protected sorts, we can define a general three-stage scheme of encoding of polymorphic formulas, explained below from the end to the beginning.

The final, “type-fusing” stage takes a set of polymorphic formulas with protected sorts and converts it into an equisatisfiable set of monomorphic formulas. A common requirement to the methods on this stage is preservation of terms with protected types: monomorphic protected fragments of the problem, e.g., arithmetic expressions, must be sent to a prover as is. We present two “type-fusing” transformations, DEC and EXP, in Sections 3.3 and 3.4. Both methods have been previously described in the literature [9, 10, 14, 6, 11]. Our presentation is more general in that it permits to protect arbitrarily complex monomorphic types, such as “list of integers” or “integer-to-real map”. The ability to preserve such sorts is of more than purely theoretical interest: as we have already mentioned, Z3, CVC3, and Yices provide built-in support for access and update operations on integer-indexed arrays.

The intermediate, “type-protecting” stage takes a set of polymorphic formulas without protected sorts and converts it into an equisatisfiable set of formulas with protection. The methods on this stage take as a parameter the set of sorts that we wish to protect; we expect them to put protection over every occurrence of every sort from this set in the problem. We present a type-protecting transformation called TW in Section 3.2. This method was introduced in [6]; we generalize it to complex sorts.

The first stage can be figuratively called “type-revealing”. Even if our type-protecting and type-fusing transformations are not limited to sort constants and can protect arbitrarily complex sorts, say, arrays of integers, we cannot readily benefit from this capacity. In an initial \mathbf{FOL}_T -problem, arrays are most probably formalized as a polymorphic type, with premises that apply to arrays of any type and with polymorphic function symbols for access and update. In order to produce interpreted monomorphic operations for Z3, CVC3, or Yices in the end, we must start by replacing, wherever possible, these function symbols with their monomorphic specializations. This is the purpose of the DIS transformation, presented in Section 3.1. We show in Section 4 that this “type revealing” brings a considerable improvement to provers’ results.

To fit the page limit, we omit the proofs of our theorems. The reader is referred to the extended technical report [5].

3.1 Symbol Discrimination

The DIS transformation involves producing a sufficient number of type instances of formulas in an initial problem F with subsequent discrimination of function and predicate symbols.

Let f be a function symbol of type signature \mathbf{S} in the initial problem Γ . Let τ be a type substitution in the type variables of \mathbf{S} . A fresh function symbol f_τ with the type signature $\mathbf{S}\tau$ is called a *specialization* of f . We call f_τ a *monomorphic specialization* if $\mathbf{S}\tau$ is monomorphic. Specializations of predicate symbols are defined in the same way.

Let W be a set of monomorphic specializations of function and predicate symbols in Γ . These are the instances that we want to “reveal” in the problem. The set W is fixed for the rest of this section; the DIS transformation is implicitly parametrized by it.

First of all, the DIS transformation modifies the signature of Γ :

1. For every variable symbol u and type T , we add a new variable symbol u_T .
2. We add every function and predicate symbol from W .

Given an arbitrary type substitution θ , the discriminating transformation DIS_θ instantiates and converts terms and formulas into the new signature:

1. Given a variable $u : T$, $\text{DIS}_\theta(u : T) \triangleq u_T : T\theta$.
2. Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$. Let τ be the type substitution that instantiates the type signature of the symbol f to $\langle T_1\theta, \dots, T_n\theta, T\theta \rangle$. If f_τ belongs to W , then $\text{DIS}_\theta(t) \triangleq f_\tau(\text{DIS}_\theta(t_1) : T_1\theta, \dots, \text{DIS}_\theta(t_n) : T_n\theta) : T\theta$. Otherwise, $\text{DIS}_\theta(t) \triangleq f(\text{DIS}_\theta(t_1) : T_1\theta, \dots, \text{DIS}_\theta(t_n) : T_n\theta) : T\theta$.
3. Consider an atomic formula $F = p(t_1 : T_1, \dots, t_n : T_n)$. Let τ be the type substitution that instantiates the type signature of p to $\langle T_1\theta, \dots, T_n\theta \rangle$. If p_τ is in W , then $\text{DIS}_\theta(F) \triangleq p_\tau(\text{DIS}_\theta(t_1) : T_1\theta, \dots, \text{DIS}_\theta(t_n) : T_n\theta)$. Otherwise, $\text{DIS}_\theta(F) \triangleq p(\text{DIS}_\theta(t_1) : T_1\theta, \dots, \text{DIS}_\theta(t_n) : T_n\theta)$.

Equalities and complex formulas are converted in a natural way:

$$\begin{aligned} \text{DIS}_\theta(t_1 \approx t_2) &\triangleq \text{DIS}_\theta(t_1) \approx \text{DIS}_\theta(t_2) & \text{DIS}_\theta(F \wedge G) &\triangleq \text{DIS}_\theta(F) \wedge \text{DIS}_\theta(G) \\ \text{DIS}_\theta(\neg F) &\triangleq \neg \text{DIS}_\theta(F) & \text{DIS}_\theta(\forall x F) &\triangleq \forall (\text{DIS}_\theta(x)) \text{DIS}_\theta(F) \end{aligned}$$

Now, let F be a closed formula. The set of monomorphic type substitutions $\Theta(F)$ is defined as follows:

$$\begin{aligned} \Theta(F) &\triangleq \{ \theta \mid F \text{ contains a term } f(t_1 : T_1, \dots, t_n : T_n) : T \text{ such that} \\ &\quad \theta \text{ only instantiates the variables of } \mathbf{T} = \langle T_1, \dots, T_n, T \rangle \text{ and} \\ &\quad W \text{ contains a specialization of } f \text{ with the type signature } \mathbf{T}\theta \} \\ &\cup \{ \theta \mid F \text{ contains an atomic formula } p(t_1 : T_1, \dots, t_n : T_n) \text{ such that} \\ &\quad \theta \text{ only instantiates the variables of } \mathbf{T} = \langle T_1, \dots, T_n \rangle \text{ and} \\ &\quad W \text{ contains a specialization of } p \text{ with the type signature } \mathbf{T}\theta \} \end{aligned}$$

We call two monomorphic type substitutions *compatible* if they do not substitute two different sorts for the same type variable. The *union* of two compatible monomorphic type substitutions is their composition (the order is irrelevant). We define $\Theta^*(F)$ as the closure of $\Theta(F)$ with respect to finite unions of compatible type substitutions. The empty union, i.e., the identity type substitution, also belongs to $\Theta^*(F)$.

Finally, DIS translates a closed formula F into a set of formulas:

$$\text{DIS}(F) \triangleq \{ \text{DIS}_\theta(F) \mid \theta \in \Theta^*(F) \}$$

On our running example, assuming $W = \{\text{get}_{[l/\alpha, l/\beta]}, \text{set}_{[l/\alpha, l/\beta]}\}$, DIS produces the following two formulas:

$$\begin{aligned} & \forall(m_{\mathbf{M}(\alpha, l)} : \mathbf{M}(\alpha, l)) \forall(c_\alpha : \alpha) \text{get}(\text{set}(m_{\mathbf{M}(\alpha, l)}, c_\alpha, \mathbf{6}), c_\alpha) * 7 \approx 42 \\ & \forall(m_{\mathbf{M}(\alpha, l)} : \mathbf{M}(l, l)) \forall(c_\alpha : l) \text{get}_{[l/\alpha, l/\beta]}(\text{set}_{[l/\alpha, l/\beta]}(m_{\mathbf{M}(\alpha, l)}, c_\alpha, \mathbf{6}), c_\alpha) * 7 \approx 42 \end{aligned}$$

The new symbol $\text{get}_{[l/\alpha, l/\beta]}$ has the monomorphic type signature $\langle \mathbf{M}(l, l), l, l \rangle$.

Lemma 2. *Let θ be a type substitution, t a term of type T , and F a formula. Then $\text{DIS}_\theta(t)$ is a well-formed term of type $T\theta$ and $\text{DIS}_\theta(F)$ is a well-formed formula such that $\text{FV}(\text{DIS}_\theta(F)) = \{\text{DIS}_\theta(x) \mid x \in \text{FV}(F)\}$ and $\text{FV}_T(\text{DIS}_\theta(F)) = \text{FV}_T(F) \setminus \text{dom}(\theta)$.*

Theorem 2. *A set of closed formulas Γ is equisatisfiable to $\text{DIS}(\Gamma)$.*

The definition of DIS can be generalized to a case where W admits polymorphic specializations. This requires W to be closed with respect to unification of type signatures, so that we can always choose the most refined specialization symbol during discrimination. The substitutions in the set $\Theta(F)$ must be considered modulo renaming of type variables in the signatures of specialization symbols. Finally, the union of two substitutions would be their most general common refinement. However, since our transformations target monomorphic theorem provers, we find this generalization of lesser practical interest and do not pursue it in this paper.

3.2 Twin Sorts

The TW transformation converts a set of formulas into an equisatisfiable set with protected sorts. It applies a pair of conversion functions to pass, wherever necessary, from a protected sort to a non-protected one and vice versa.

Let U be a set of sorts that we want to preserve across our type-eliminating transformations. The set U is fixed for the rest of the section and the TW transformation is parametrized by it. Given a type T , the transformed type $[T]$ is \bar{T} if $T \in U$, and T otherwise. Then TW modifies the signature of a transformed theory as follows:

1. We replace every function symbol f of type signature $\langle S_1, \dots, S_n, S \rangle$ with a symbol \bar{f} of type signature $\langle [S_1], \dots, [S_n], [S] \rangle$.
2. We replace every predicate symbol p of type signature $\langle S_1, \dots, S_n \rangle$ with a symbol \bar{p} of type signature $\langle [S_1], \dots, [S_n] \rangle$.
3. For every sort $T \in U$, we add a pair of “bridge” function symbols $\text{to}_T : \langle \bar{T}, T \rangle$ and $\text{from}_T : \langle T, \bar{T} \rangle$.

Then we convert terms and atomic formulas into the new signature. Our aim is to forbid a polymorphic type in a symbol's type signature being instantiated into a type from U . Whenever such instantiation takes place, a bridge function is applied. In more precise terms:

1. Given a variable $u:T$, $\text{Tw}(u:T) \triangleq u:[T]$.
2. Consider a term $t = f(t_1:T_1, \dots, t_n:T_n):T$ and let $\langle S_1, \dots, S_n, S \rangle$ be the type signature of f .

$$\text{For every } t_i, t'_i \triangleq \begin{cases} \text{to}_{T_i}(\text{Tw}(t_i)):T_i & \text{if } T_i \in U \text{ and } S_i \notin U, \\ \text{Tw}(t_i) & \text{if } T_i \notin U \text{ or } S_i \in U. \end{cases}$$

$$\text{Then } \text{Tw}(t) \triangleq \begin{cases} \bar{f}(t'_1, \dots, t'_n):[T] & \text{if } T \notin U \text{ or } S \in U, \\ \text{from}_T(\bar{f}(t'_1, \dots, t'_n):T):\bar{T} & \text{if } T \in U \text{ and } S \notin U. \end{cases}$$

3. Consider a formula $p(t_1:T_1, \dots, t_n:T_n)$ and let $\langle S_1, \dots, S_n \rangle$ be the type signature of p . For every argument t_i , we define t'_i as in the previous case. Then, $\text{Tw}(p(t_1:T_1, \dots, t_n:T_n)) \triangleq \bar{p}(t'_1, \dots, t'_n)$.

Equalities and complex formulas are converted in a natural way:

$$\begin{aligned} \text{Tw}(t_1 \approx t_2) &\triangleq \text{Tw}(t_1) \approx \text{Tw}(t_2) & \text{Tw}(F \wedge G) &\triangleq \text{Tw}(F) \wedge \text{Tw}(G) \\ \text{Tw}(\neg F) &\triangleq \neg \text{Tw}(F) & \text{Tw}(\forall x F) &\triangleq \forall(\text{Tw}(x)) \text{Tw}(F) \end{aligned}$$

Finally, we convert the formulas in Γ and add axioms for the bridge functions:

$$\begin{aligned} \text{Tw}(\Gamma) &\triangleq \{ \text{Tw}(F) \mid F \in \Gamma \} \\ &\cup \{ \forall(v:\bar{T}) \text{from}_T(\text{to}_T(v:\bar{T})) \approx v:\bar{T} \mid T \in U \} \\ &\cup \{ \forall(u:T) \text{to}_T(\text{from}_T(u:T)) \approx u:T \mid T \in U \} \end{aligned}$$

Assuming $U = \{1\}$, the running example is transformed as follows. Notice that **6**, **7**, and **42** have the type $\bar{1}$ and the type signature of $*$ is $\langle \bar{1}, \bar{1}, \bar{1} \rangle$.

$$\forall(m:\mathbb{M}(\alpha, 1)) \forall(c:\alpha) \text{from}_1(\text{get}(\text{set}(m, c, \text{to}_1(6):1), c):1) * 7 \approx 42$$

Lemma 3. *For every term t of type T , $\text{Tw}(t)$ is a well-formed term of type $[T]$, and for every formula F , $\text{Tw}(F)$ is a well-formed formula with the same free variables (modulo conversion of their types) and type variables.*

Theorem 3. *A set of closed formulas Γ is equisatisfiable to $\text{Tw}(\Gamma)$.*

3.3 Decorated Terms

The DEC transformation converts a polymorphic problem with protected sorts into an equisatisfiable monomorphic problem. Roughly speaking, in order to preserve type information, it decorates every term with its type, which itself is transformed to a term of a special sort.

First of all, we introduce three fresh sort constants U , D , and T . The first one is assigned to undecorated terms, the second one to decorated terms, and the third one to the terms representing types. To transform the type signatures of function and predicate symbols, we use the following operations on types:

$$[T]^- \triangleq \begin{cases} T & \text{if } T \text{ is protected,} \\ D & \text{otherwise} \end{cases} \quad [T]^+ \triangleq \begin{cases} T & \text{if } T \text{ is protected,} \\ U & \text{otherwise} \end{cases}$$

Now, the signature of the resulting theory is defined as follows:

1. The set of type constructors is extended with U , D , T .
2. We replace every function symbol f of type signature $\langle S_1, \dots, S_n, S \rangle$ with a symbol \hat{f} with the monomorphic type signature $\langle [S_1]^-, \dots, [S_n]^-, [S]^+ \rangle$.
3. We replace every predicate symbol p of type signature $\langle S_1, \dots, S_n \rangle$ with a symbol \hat{p} with the monomorphic type signature $\langle [S_1]^-, \dots, [S_n]^-, \rangle$.
4. For every variable symbol u and type T , we add a new variable symbol u_T .
5. For every type variable $\alpha \in \mathbb{V}_T$, we add a new variable symbol v^α .
6. For every type constructor $F \in \mathbb{F}_T$, we add a new function symbol F of the same arity and with type signature $\langle T, \dots, T, T \rangle$.
7. We add a new “decoration” function symbol $\text{deco} : \langle T, U, D \rangle$.

The DEC transformation applies to non-protected types, translating them to terms of type T :

$$\text{DEC}(\alpha) \triangleq v^\alpha : T \quad \text{DEC}(F(T_1, \dots, T_n)) \triangleq F(\text{DEC}(T_1), \dots, \text{DEC}(T_n)) : T$$

In the next definition, \mathbf{t} stands for a vector of terms, \bar{S} for a protected sort, and T for a non-protected type. The DEC transformation applies to terms:

$$\begin{aligned} \text{DEC}(u : \bar{S}) &\triangleq u_{\bar{S}} : \bar{S} \\ \text{DEC}(u : T) &\triangleq \text{deco}(\text{DEC}(T), u_T : U) : D \\ \text{DEC}(f(\mathbf{t}) : \bar{S}) &\triangleq \hat{f}(\text{DEC}(\mathbf{t})) : \bar{S} \\ \text{DEC}(f(\mathbf{t}) : T) &\triangleq \text{deco}(\text{DEC}(T), \hat{f}(\text{DEC}(\mathbf{t})) : U) : D \end{aligned}$$

and formulas (here, $\{\alpha_1, \dots, \alpha_m\} = \text{FV}_T(H)$):

$$\begin{aligned} \text{DEC}(p(\mathbf{t})) &\triangleq \hat{p}(\text{DEC}(\mathbf{t})) & \text{DEC}(\neg F) &\triangleq \neg \text{DEC}(F) \\ \text{DEC}(t_1 \approx t_2) &\triangleq \text{DEC}(t_1) \approx \text{DEC}(t_2) & \text{DEC}(\forall(u : \bar{S})F) &\triangleq \forall(u_{\bar{S}} : \bar{S}) \text{DEC}(F) \\ \text{DEC}(F \wedge G) &\triangleq \text{DEC}(F) \wedge \text{DEC}(G) & \text{DEC}(\forall(u : T)F) &\triangleq \forall(u_T : U) \text{DEC}(F) \\ \text{DEC}^\circ(H) &\triangleq \forall(v^{\alpha_1} : T) \dots \forall(v^{\alpha_m} : T) \text{DEC}(H) \end{aligned}$$

On the running example, assuming $U = \{1\}$, the transformations TW and DEC° produce the following monomorphic formula:

$$\begin{aligned} \forall(v^\alpha : T) \forall(m_{M(\alpha, 1)} : U) \forall(c_\alpha : U) \text{from}_1(\text{deco}(\mathbf{I}, \text{get}(\text{deco}(\mathbf{M}(v^\alpha), \mathbf{I}), \\ \text{set}(\text{deco}(\mathbf{M}(v^\alpha), \mathbf{I}), m_{M(\alpha, 1)}, \text{deco}(v^\alpha, c_\alpha), \text{deco}(\mathbf{I}, \text{toi}(6))))), \\ \text{deco}(v^\alpha, c_\alpha))) * 7 \approx 42 \end{aligned}$$

The second axiom of bridge functions to_1 and from_1 becomes

$$\forall(u_1 : \mathbf{U}) \text{deco}(\mathbf{I}, \text{to}_1(\text{from}_1(\text{deco}(\mathbf{I}, u_1)))) \approx \text{deco}(\mathbf{I}, u_1)$$

Due to the outer application of deco on the both sides of equality, our translation is sound even when we protect finite types, such as booleans. Without this additional decoration (as in [6, Eq. (8)]), the finiteness of a protected sort implies the finiteness of the whole sort \mathbf{U} .

Lemma 4. *For every term t of type T , $\text{DEC}(t)$ is a well-formed monomorphic term of type $[T]^-$. For every formula F , $\text{DEC}(F)$ is a well-formed monomorphic formula. Also, $\text{FV}(\text{DEC}(t)) = \{u_T : [T]^+ \mid u : T \in \text{FV}(t)\} \cup \{v^\alpha : \mathbf{T} \mid \alpha \in \text{FV}_{\mathbf{T}}(t)\}$ and $\text{FV}(\text{DEC}(F)) = \{u_T : [T]^+ \mid u : T \in \text{FV}(F)\} \cup \{v^\alpha : \mathbf{T} \mid \alpha \in \text{FV}_{\mathbf{T}}(F)\}$. For every closed formula F , $\text{DEC}^\circ(F)$ is a well-formed closed monomorphic formula.*

Theorem 4. *A set of closed formulas with protected sorts Γ is satisfiable if and only if $\text{DEC}^\circ(\Gamma)$ is satisfiable.*

3.4 Explicit Polymorphism

The EXP transformation is similar to DEC except that instead of attaching an explicit type annotation to every term, we add type-representing arguments to polymorphic symbols. This allows for much lighter modifications in the original problem. However, the method is only sound on problems that admit a model where every non-protected sort has an infinite domain.

We introduce fresh sort constants \mathbf{U} and \mathbf{T} . The first one replaces non-protected types and the second one, as in DEC, is the sort of type-representing terms. For any type T , we define $[T]$ to be T if T is protected, and \mathbf{U} otherwise. Then EXP modifies the signature of a transformed theory in the following way:

1. The set of type constructors is extended with \mathbf{U} and \mathbf{T} .
2. Let f be a function symbol of signature $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$ and $\alpha_1, \dots, \alpha_r$ be the free type variables of \mathbf{S} . We replace f with a function symbol \hat{f} of arity $r + n$ with monomorphic type signature $\langle \mathbf{T}, \dots, \mathbf{T}, [S_1], \dots, [S_n], [S] \rangle$.
3. Let p be a predicate symbol of signature $\mathbf{S} = \langle S_1, \dots, S_n \rangle$ and $\alpha_1, \dots, \alpha_r$ be the free type variables of \mathbf{S} . We replace p with a predicate symbol \hat{p} of arity $r + n$ with monomorphic type signature $\langle \mathbf{T}, \dots, \mathbf{T}, [S_1], \dots, [S_n] \rangle$.
4. For every variable symbol u and type T , we add a new variable symbol u_T .
5. For every type variable $\alpha \in \mathbb{V}_{\mathbf{T}}$, we add a new variable symbol v^α .
6. For every type constructor $F \in \mathbb{F}_{\mathbf{T}}$, we add a new function symbol \mathbf{F} of the same arity and with type signature $\langle \mathbf{T}, \dots, \mathbf{T}, \mathbf{T} \rangle$.

The EXP transformation applies to non-protected types, translating them to terms of type \mathbf{T} , exactly as DEC:

$$\text{EXP}(\alpha) \triangleq v^\alpha : \mathbf{T} \quad \text{EXP}(\mathbf{F}(T_1, \dots, T_n)) \triangleq \mathbf{F}(\text{EXP}(T_1), \dots, \text{EXP}(T_n)) : \mathbf{T}$$

The EXP transformation applies to terms and formulas. In the definition below, \mathbf{t} stands for a list of terms; $\alpha_1, \dots, \alpha_r$ are the type variables of the type signature of f and p ; the type signature of f and p is instantiated with a type substitution τ ; and β_1, \dots, β_m are the type variables of H :

$$\begin{aligned}
\text{EXP}(u : T) &\triangleq u_T : [T] \\
\text{EXP}(f(\mathbf{t}) : T) &\triangleq \hat{f}(\text{EXP}(\alpha_1\tau), \dots, \text{EXP}(\alpha_r\tau), \text{EXP}(\mathbf{t})) : [T] \\
\text{EXP}(p(\mathbf{t})) &\triangleq \hat{p}(\text{EXP}(\alpha_1\tau), \dots, \text{EXP}(\alpha_r\tau), \text{EXP}(\mathbf{t})) \\
\text{EXP}(t_1 \approx t_2) &\triangleq \text{EXP}(t_1) \approx \text{EXP}(t_2) \\
\text{EXP}(\neg F) &\triangleq \neg \text{EXP}(F) \\
\text{EXP}(F \wedge G) &\triangleq \text{EXP}(F) \wedge \text{EXP}(G) \\
\text{EXP}(\forall x F) &\triangleq \forall(\text{EXP}(x)) \text{EXP}(F) \\
\text{EXP}^\circ(H) &\triangleq \forall(v^{\beta_1} : \mathbb{T}) \dots \forall(v^{\beta_m} : \mathbb{T}) \text{EXP}(H)
\end{aligned}$$

On the running example, assuming $U = \{1\}$, the transformations TW and EXP° produce the following formula:

$$\begin{aligned}
&\forall(v^\alpha : \mathbb{T}) \forall(m_{\mathbb{M}(\alpha, 1)} : \mathbb{U}) \forall(c_\alpha : \mathbb{U}) \text{from}_1(\text{get}(v^\alpha, \mathbb{I}, \\
&\quad \text{set}(v^\alpha, \mathbb{I}, m_{\mathbb{M}(\alpha, 1)}, c_\alpha, \text{to}_1(\mathbb{6})), c_\alpha)) * 7 \approx 42
\end{aligned}$$

Lemma 5. *For every term t of type T , $\text{EXP}(t)$ is a well-formed monomorphic term of type $[T]$. For every formula F , $\text{EXP}(F)$ is a well-formed monomorphic formula. Also, $\text{FV}(\text{EXP}(t)) = \{u_T : [T] \mid u : T \in \text{FV}(t)\} \cup \{v^\alpha : \mathbb{T} \mid \alpha \in \text{FV}_T(t)\}$ and $\text{FV}(\text{EXP}(F)) = \{u_T : [T] \mid u : T \in \text{FV}(F)\} \cup \{v^\alpha : \mathbb{T} \mid \alpha \in \text{FV}_T(F)\}$. Finally, for every closed formula F , $\text{EXP}^\circ(F)$ is a well-formed closed monomorphic formula.*

Theorem 5. *Let Γ be a set of closed formulas with protected sorts. If Γ is satisfiable so that every non-protected sort has an infinite domain in the model, then $\text{EXP}^\circ(\Gamma)$ is satisfiable. Conversely, if $\text{EXP}^\circ(\Gamma)$ is satisfiable then Γ is satisfiable.*

From a practical point of view, the soundness part of Theorem 5 is not comforting. Given a \mathbf{FOL}_T -problem Γ , we cannot effectively decide which sorts admit infinite models and which do not (one can postulate a bijection between a given sort and the domain of a partial-recursive function). A practical way out could consist in a small language extension: for every type/sort, we specify explicitly whether it is finite or infinite. We proceed from the assumption that the author of any given problem knows the intended model of every type.

In Why3 [4], every type is declared either as abstract or algebraic (i.e., a sum of products). We postulate that abstract types are all infinite and we analyse the definitions of algebraic types to find out which of their monomorphic instances admit infinite models. For example, given the standard algebraic definitions of booleans (\mathbb{B}), lists ($\mathbb{L}(\alpha)$), and pairs ($\mathbb{P}(\alpha, \beta)$), we can conclude that the sorts \mathbb{B} and $\mathbb{P}(\mathbb{B}, \mathbb{B})$ are finite and \mathbb{L} , $\mathbb{L}(\mathbb{B})$, and $\mathbb{P}(\mathbb{L}, \mathbb{B})$ are infinite.

Once we know the finite sorts, can we transform a problem to eliminate them, so that EXP (or a similar method) can be applied? Meng and Paulson propose

to filter out the premises implying the finiteness of sorts [14, Sect. 2.8]; however, we need an infallible filter to ensure the soundness of type erasure. We have implemented an alternative solution which consists in putting a special “projection” function proj_T over every variable and function symbol of a finite type T . Thus, the premise $\forall(x : B)(x \approx \text{True} \vee x \approx \text{False})$ becomes $\forall(x : B)(\text{proj}_B(x) \approx \text{proj}_B(\text{True}) \vee \text{proj}_B(x) \approx \text{proj}_B(\text{False}))$ and the domain of B does not need to be finite anymore, as we confine ourselves to the range of proj_B . This method is still potentially unsound, as one can state the finiteness of a sort with a polymorphic axiom, where no projection would apply. Precisely, let $\text{isUnit} : \langle \alpha \rangle$ be a unary predicate. Then the formulas $\forall(x : \alpha)(\text{isUnit}(x) \supset \forall(y : \alpha)(y \approx x))$ and $\forall(x : A) \text{isUnit}(x)$ imply that the sort A has a single inhabitant. Today, we know of no way to use EXP soundly on polymorphic problems with finite sorts.

The last remark to make is that EXP provides a path towards one-sorted languages. Indeed, in a monomorphic setting, Theorem 5 comes to: “if every sort admits an infinite domain, then we can safely erase the sort annotations”. Thus, if we want to use a TPTP prover such as Vampire or SPASS, we start by translating a proof task to the many-sorted language, using any of the methods described above. Then we eliminate the protected finite sorts (if any) using projections; in absence of polymorphism, this is a sound and complete transformation. And finally, we apply EXP assuming that all types are non-protected, which amounts to simply erasing all sorts.

4 Experiments and Conclusion

In our experiments, we wanted to compare the impact of different “paths” of polymorphism encoding on the performance of three well-established SMT solvers. We add the classical type encoding technique with per-variable “type guards” [12]. Our implementation of this method (denoted GRD below) closely follows the description given in [11, Sect. 3.0].

We run our tests on 4123 verification conditions generated by the Why platform from 166 programs, which originate from Caduceus [8], Jessie [13], or directly from Why. Translated tasks were sent to Z3, CVC3, and Yices with a time limit of 60 seconds. On the whole, 3993 proof obligations were proved by at least one prover. The initial Why3 files and our results are available at http://why3.lri.fr/download/polyfol_encoding.tar.gz.

We have tested the encodings TW+DEC, TW+EXP, and TW+GRD, both with and without DIS. In the latter case, these methods correspond to what is described in [6] and [11]; the set U of sorts to protect in TW is set to contain only integers and reals, which are natively supported by the three provers. In presence of DIS, we put in the set W every monomorphic specialization that occurs in the goal formula along with the specializations of access and update operations on every monomorphic array type in the goal; we also protect every sort in the goal (as well as integers and reals). This configuration of DIS and TW gives better results comparing to other configurations that we tried, e.g., collect the specializations and the sorts to protect from the whole proof task.

Z3 (3809)	TW+GRD	TW+EXP	TW+DEC	DIS+TW+GRD	DIS+TW+EXP
DIS+TW+DEC	+203 -36	+20 -49	+66 -37	+18 -5	+26 -30
DIS+TW+EXP	+191 -20	+13 -38	+63 -30	+35 -18	
DIS+TW+GRD	+195 -41	+11 -53	+59 -43		
TW+DEC	+157 -19	+15 -73			
TW+EXP	+211 -15				

CVC3 (3756)	TW+GRD	TW+EXP	TW+DEC	DIS+TW+GRD	DIS+TW+EXP
DIS+TW+DEC	+269 -20	+0 -26	+84 -19	+66 -4	+0 -6
DIS+TW+EXP	+272 -17	+0 -20	+88 -17	+69 -1	
DIS+TW+GRD	+204 -17	+1 -89	+46 -43		
TW+DEC	+188 -4	+0 -91			
TW+EXP	+275 -0				

Yices (3717)	TW+GRD	TW+EXP	TW+DEC	DIS+TW+GRD	DIS+TW+EXP
DIS+TW+DEC	+882 -6	+13 -276	+379 -79	+204 -2	+3 -272
DIS+TW+EXP	+1149 -4	+39 -33	+574 -5	+472 -1	
DIS+TW+GRD	+684 -10	+6 -471	+241 -143		
TW+DEC	+577 -1	+5 -568			
TW+EXP	+1140 -1				

Our results are given in the table above. To the right of the prover's name, we put the number of goals proved by at least one encoding method. In every cell we specify the number of goals proved by one encoding but not by the other one. For example, with CVC3, the encoding by DIS+TW+DEC allows us to prove 84 goals that were not proved by TW+DEC. On the other hand, with TW+DEC, CVC3 proves 19 goals that were not proved with DIS+TW+DEC.

On the average, symbol discrimination increases the number of premises by a factor of 1.8 (ranging from 1 to 10 on some examples). Nevertheless, adding the DIS phase allows us to prove more goals in every case except for Z3 and CVC3 with EXP. In particular, the GRD transformation is remarkably helped by DIS. Apart from the possibility to use the built-in support for arrays, the effectiveness of DIS is also explained by the fact that we protect the sorts that occur in the selected monomorphic specializations. Thus, the new premises generated by DIS are not only instantiated to the relevant sorts, they are also liberated from decorations imposed by the third, type-fusing, stage. This effect is less important in the case of EXP, because this transformation, unlike DEC and GRD, adds very little clutter to the encoded formulas in the first place.

Also notice that type protection, Tw, is crucially important: if we protect no types at all, we prevent provers from using their built-in theories, and the total number of goals proved (using only EXP, DEC, or GRD) drops to 1861.

The comparison between EXP, DEC, and GRD shows that EXP is generally more efficient than DEC which in its turn is more efficient than GRD. This is quite different from the results given in [11], where EXP and GRD have roughly the same performance. We have not yet identified whether this discrepancy comes from the difference in our test cases or in our implementations.

Conclusion. In the present paper, we described first-order logic with polymorphic types and introduced generic notions to define and reason about practical methods of polymorphism elimination. Using these notions, we generalized and proved two translation techniques known from literature. We also proposed to combine type protection with symbol discrimination. As our experiments show, this improves the performance of automated proof search and allows us to use built-in theories of complex types, such as arrays, in SMT solvers. One interesting problem we would like to resolve in the future is protection of polymorphic types, allowing to merge all monomorphic instances of a given complex type in a single protected sort. We also believe that better heuristics to choose the sets W and U can be devised, and further experiments are in order.

References

1. Barrett, C., Stump, A., Tinelli, C.: The SMT-LIB Standard: Version 2.0. Tech. rep., Department of Computer Science, The University of Iowa (2010)
2. Barrett, C., Tinelli, C.: CVC3. In: CAV'07. LNCS, vol. 4590, pp. 298–302. Springer (2007)
3. Bobot, F., Conchon, S., Contejean, E., Lescuyer, S.: Implementing polymorphism in SMT solvers. In: SMT'08. ACM ICPS, vol. 367, pp. 1–5 (2008)
4. Bobot, F., Filiâtre, J.C., Marché, C., Paskevich, A.: Why3: Shepherd your herd of provers. In: Boogie 2011 (co-loc. with CADE-23). Wrocław, Poland (August 2011)
5. Bobot, F., Paskevich, A.: Expressing polymorphic types in a many-sorted language (extended version) (Jul 2011), <http://hal.inria.fr/inria-00591414/en/>
6. Couchot, J.F., Lescuyer, S.: Handling polymorphism in automated deduction. In: CADE-21. LNAI, vol. 4603, pp. 263–278. Springer (2007)
7. Dutertre, B., de Moura, L.: The YICES SMT solver. Tech. rep., SRI International (2006)
8. Filiâtre, J.C., Marché, C.: Multi-prover verification of C programs. In: ICFEM'04. LNCS, vol. 3308, pp. 15–29. Springer (2004)
9. Hurd, J.: An LCF-style interface between HOL and first-order logic. In: CADE-18. LNAI, vol. 2392, pp. 134–138. Springer (2002)
10. Hurd, J.: First-order proof tactics in higher-order logic theorem provers. In: Design and Application of Strategies/Tactics in Higher Order Logics. pp. 56–68. NASA Technical Report NASA/CP-2003-212448 (2003)
11. Leino, K.R.M., Rümmer, P.: A polymorphic intermediate verification language: Design and logical encoding. In: TACAS'10. LNCS, vol. 6015, pp. 312–327. Springer (2010)
12. Manzano, M.: Extensions of First-Order Logic, Cambridge Tracts in Theoretical Computer Science, vol. 19. Cambridge University Press (1996)
13. Marché, C., Moy, Y.: Jessie plug-in (2010), <http://frama-c.com/jessie.html>
14. Meng, J., Paulson, L.C.: Translating higher-order clauses to first-order clauses. Journal of Automated Reasoning 40(1), 35–60 (2008)
15. de Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: TACAS'08. LNCS, vol. 4963, pp. 337–340. Springer (2008)
16. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL — A Proof Assistant for Higher-Order Logic, LNCS, vol. 2283. Springer (2002)
17. The Coq Development Team: The Coq Proof Assistant Reference Manual – Version V8.0 (2004), <http://coq.inria.fr>