



HAL
open science

Expressing Polymorphic Types in a Many-Sorted Language

François Bobot, Andrei Paskevich

► **To cite this version:**

François Bobot, Andrei Paskevich. Expressing Polymorphic Types in a Many-Sorted Language. [Research Report] 2011. inria-00591414v1

HAL Id: inria-00591414

<https://inria.hal.science/inria-00591414v1>

Submitted on 9 May 2011 (v1), last revised 26 Jul 2011 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Expressing Polymorphic Types in a Many-Sorted Language (extended version)

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 the 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 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 and CVC3. 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 [15].

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 AltErgo 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 [11] and [10, Sect. 3.0]), to throughout decoration of terms with their types [8, 5], to various flavours of type erasure [9, 12, 10]. The latter method is logically

unsound, though adding type annotations can prevent certain unsound inference steps (see [12, Sect. 2.5, 2.6] and [10, 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 [5, 10]. 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 [13] and CVC3 [2] have a built-in support for arrays. Secondly, type protection, as defined in [5], 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 [10, Sect. 3.1] is less intrusive and more efficient than full term decoration [5], the former method is unsound and, according to [10], 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 that we use (in slightly different variants) both for the first and the second stage (Sections 3.1 and 3.2). Furthermore, we give a generalized formulation of the type protection method from [5], free from the aforesaid restrictions (Section 3.3). As third-stage transformations, we consider full term decoration from [8, 5] (Section 3.4) and type erasure with added type annotations from [10, Sect. 3.1] (Section 3.5). 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 [6] on a set of about 4100 proof obligations in the Why3 tool [4] (Section 4).

2 First-Order Logic With Polymorphic Types

Types. We define types as syntactical expressions constructed from *type variables* (denoted α, β, γ) and *type constructors* of fixed arity (denoted with capital sans-serif letters). For example, $\beta, \mathbf{l}, \mathbf{F}(\mathbf{l}, \gamma)$ are well-formed types. A type constructor of arity 0 is called *type constant*. A type that does not contain type variables is called *monomorphic type*, *closed type* or *sort*. A vector of types $\langle T_1, \dots, T_n \rangle$ is called a *type signature*. We define *type substitution* and *type matching* in the same way as they are usually defined for terms. A *monomorphic* type substitution instantiates every type variable either to itself or to a closed type.

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.

Terms and formulas. 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, be it a variable or a function application, carries a type, e.g.: $w : \mathbf{C}(\mathbf{l}), w : \mathbf{C}(\alpha), f(u : \alpha, v : \mathbf{L}(\alpha)) : \mathbf{L}(\alpha)$. The first two terms are considered as different variables even though they share the variable symbol. 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.
- 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: $\forall(u : \alpha) p(u : \alpha, u : \mathbf{C})$. 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} . We suppose that the sets $\mathbb{V}_T, \mathbb{F}_T, \mathbb{V}, \mathbb{F}, \mathbb{P}$ are pairwise disjoint; these sets form the signature of a given theory.

Given a term or a formula e , the set of type variables occurring in e is denoted $\mathbf{FV}_T(e)$ and the set of free variables of e is denoted $\mathbf{FV}(e)$. If $\mathbf{FV}_T(e)$ is empty, e is called *monomorphic*. If $\mathbf{FV}(e)$ is empty, e is called *closed* or *ground*.

Substitutions apply to terms and formulas, replacing free term variables with terms of the same type. We denote substitutions with letters σ and δ . 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 future examples, we will not use a variable symbol in two different variables in the same formula to avoid confusion.

Running example. Throughout the rest of the paper, we illustrate our transformations on the following simple polymorphic formula (for the sake of readability, we omit the most obvious type annotations):

$$\forall(m : M(\alpha, l)) \forall(c : \alpha) \text{get}(\text{set}(m, c, 6) : M(\alpha, l), c) : l * 7 \approx 42$$

Here, the type l represents integers and the type $M(\alpha, \beta)$ is that of polymorphic α -to- β maps. The function symbol get is of type signature $\langle M(\alpha, \beta), \alpha, \beta \rangle$ and set is of type signature $\langle M(\alpha, \beta), \alpha, \beta, M(\alpha, \beta) \rangle$.

Satisfiability. We are ready to introduce \mathbf{FOL}_T , a classical first-order logic with equality and polymorphic types. Given sets $\mathbb{F}_T, \mathbb{F}, \mathbb{P}$, an interpretation \mathfrak{J} is defined by three mappings:

- to each sort $S \in \mathcal{T}(\mathbb{F}_T)$, we assign a non-empty *domain* $\mathcal{D}_S^{\mathfrak{J}}$;
- to each symbol $f \in \mathbb{F}$ and each vector of sorts $\langle S_1, \dots, S_n, S \rangle$ matched by the type signature of f , we assign a total function $f_S^{\mathfrak{J}} : \mathcal{D}_{S_1}^{\mathfrak{J}} \times \dots \times \mathcal{D}_{S_n}^{\mathfrak{J}} \rightarrow \mathcal{D}_S^{\mathfrak{J}}$;
- to each symbol $p \in \mathbb{P}$ and each vector of sorts $\langle S_1, \dots, S_n \rangle$ matched by the type signature of p , we assign a total function $p_S^{\mathfrak{J}} : \mathcal{D}_{S_1}^{\mathfrak{J}} \times \dots \times \mathcal{D}_{S_n}^{\mathfrak{J}} \rightarrow \{\top, \perp\}$, 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}_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 φ be a variable valuation under π . We evaluate terms and formulas according to the following equalities:

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

where $\mathbf{t} : \mathbf{T}$ stands for a vector of terms $t_1 : T_1, \dots, t_n : T_n$ and $\varphi[u : T \mapsto a]$ is a valuation that coincides with φ everywhere except $u : T$, which is mapped to a .

It is easy to see that evaluation of a term or a formula e under $\mathfrak{J}_{\pi, \varphi}$ does not depend on the values of π and φ on (type) variables that do not occur in e .

Lemma 1. *Let F be a formula and \mathfrak{J} be an interpretation. Let π and τ be type valuations that coincide on every type variable in $\text{FV}_{\mathbf{T}}(F)$. Let φ be a variable valuation under π and ψ be a variable valuation under τ . Let us suppose that $\varphi(u : T) = \psi(u : T)$ for every variable $u : T$ in $\text{FV}(F)$. Then $\mathfrak{J}_{\pi, \varphi}(F) = \mathfrak{J}_{\tau, \psi}(F)$.*

In what follows, when we evaluate closed or monomorphic expressions, we often omit the variable valuation or the type valuation, respectively.

Lemma 2. *Let F be a closed formula and \mathfrak{J} be an interpretation. For any type valuation π , $\mathfrak{J}_{\pi}(F) = \mathfrak{J}(F\pi)$.*

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.

It is easy to see that on monomorphic terms and formulas our definitions correspond to the traditional many-sorted logic with disjoint sorts. Moreover, by Lemma 2, we can see closed polymorphic formulas as mere “schemes” of their monomorphic type instances (of which there may be an infinite number). To put this precisely, we denote the set of all monomorphic type instances of a closed formula F with $\text{MI}(F)$. Similarly, given a set of closed formulas Γ , the set $\text{MI}(\Gamma)$ contains all monomorphic type instances of formulas in Γ . The trivial corollary of Lemma 2 is that Γ is satisfiable in $\mathbf{FOL}_{\mathbf{T}}$ if and only if $\text{MI}(\Gamma)$ is.

Computing monomorphic instances? Of course, as soon as a closed formula F is non-monomorphic and there is at least one non-constant type constructor, F has infinitely many monomorphic type instances. But can’t we find out, in finite time, all sorts that are potentially relevant to the problem in question, 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 (Undecidability of finite monomorphisation). *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; the question is whether it is computable).*

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 equivalent monomorphic one. Such conversion inevitably implies merging many types into few sorts or just a single sort. This may be 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 3.4 and 3.5. Both methods have been previously described in the literature [8, 9, 12, 5, 10]. Our presentation is more general in that it allows 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 and CVC3 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. In this paper, we present two type-protecting transformations, PAR and TW, in Sections 3.2 and 3.3. The latter method was introduced in [5]; 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 or CVC3 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.

3.1 Symbol Discrimination

The DIS transformation consists in producing a sufficient number of type instances of formulas in an initial problem Γ 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 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 $W = \{\mathbf{get}_{[l/\alpha, l/\beta]}, \mathbf{set}_{[l/\alpha, l/\beta]}\}$, Dis produces the following two formulas:

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

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

Theorem 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 $\text{FV}_T(\text{Dis}_\theta(F)) = \text{FV}_T(F) \setminus \text{dom}(\theta)$.*

Proof. The proof goes by induction on the structure of terms and formulas. The only interesting case is the application of a function symbol (predicate symbols are treated in the same way). Let t be a term of the form $f(t_1 : T_1, \dots, t_n : T_n) : T$. Let τ be the type substitution that instantiates the type signature of f to $\langle T_1\theta, \dots, T_n\theta, T\theta \rangle$. For every argument t_i , $\text{Dis}_\theta(t_i)$ is a well-formed term of type $T_i\theta$ by the induction hypothesis. If f_τ belongs to W , then its type signature is exactly $\langle T_1\theta, \dots, T_n\theta, T\theta \rangle$. Otherwise, as τ witnesses, the type signature of f matches $\langle T_1\theta, \dots, T_n\theta, T\theta \rangle$. \square

Theorem 3 (Soundness of Dis). *If a set of closed formulas Γ is satisfiable then $\text{Dis}(\Gamma)$ is satisfiable.*

Proof. Let \mathfrak{M} be a model of Γ . We are going to construct an interpretation \mathfrak{J} of $\text{Dis}(\Gamma)$ such that, for every type substitution θ , type valuation π , and variable valuation φ in \mathfrak{J} , there exist a type valuation π' and a variable valuation φ' in the original model \mathfrak{M} , such that:

$$\mathfrak{J}_{\pi, \varphi}(\text{Dis}_\theta(t)) = \mathfrak{M}_{\pi', \varphi'}(t) \quad \mathfrak{J}_{\pi, \varphi}(\text{Dis}_\theta(F)) = \mathfrak{M}_{\pi', \varphi'}(F)$$

for every term t and formula F occurring in Γ . Since for every closed $F \in \Gamma$ and every $\theta \in \Theta^*(F)$, \mathfrak{J} satisfies $\text{Dis}_\theta(F)$, \mathfrak{J} is a model of $\text{Dis}(\Gamma)$.

For any sort S in the new signature, we define the domain $\mathcal{D}_S^{\mathfrak{J}} \triangleq \mathcal{D}_S^{\mathfrak{M}}$. Moreover, every function and predicate symbol in Γ is interpreted in \mathfrak{J} exactly as in \mathfrak{M} . Every new function symbol f_τ with the type signature \mathbf{S} in $\text{Dis}(\Gamma)$ is interpreted as f on \mathbf{S} in Γ : $(f_\tau)_{\mathbf{S}}^{\mathfrak{J}} \triangleq f_{\mathbf{S}}^{\mathfrak{M}}$. New predicate symbols are interpreted in the same way.

Now, let us consider some type valuation π , a variable valuation φ in \mathfrak{J} , and a type substitution θ . We set π' to be $\theta \circ \pi$ and we define $\varphi'(u : T) \triangleq \varphi(u_T : T\theta)$. Notice that for any type T in Γ , $\mathcal{D}_{T\theta\pi}^{\mathfrak{J}} = \mathcal{D}_{T\theta\pi}^{\mathfrak{M}} = \mathcal{D}_{T\pi}^{\mathfrak{M}}$.

Using this definitions, we proceed to prove the two equalities above. The proof goes by induction on the structure of terms and formulas. We consider three cases: a variable, a function application, and a universally quantified formula (other cases can be handled along the same lines).

1. Consider a variable $u : T$ occurring in Γ . We obtain:

$$\mathfrak{J}_{\pi, \varphi}(\text{Dis}_\theta(u : T)) = \varphi(u_T : T\theta) = \varphi'(u : T) = \mathfrak{M}_{\pi', \varphi'}(u : T)$$

2. Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$ and let \mathbf{T} denote $\langle T_1, \dots, T_n, T \rangle$. Independently of whether W contains a specialization of f with the type signature $\mathbf{T}\theta$, the resulting function application is interpreted in \mathfrak{J} exactly as f on $\mathbf{T}\theta\pi = \mathbf{T}\pi'$ in \mathfrak{M} . Therefore, $\mathfrak{J}_{\pi, \varphi}(\text{DIS}_\theta(t)) = \mathfrak{M}_{\pi', \varphi'}(t)$.
3. Consider a universally quantified formula $\forall(u : T) F$. As we have shown above, $\mathcal{D}_{T\theta\pi}^{\mathfrak{J}} = \mathcal{D}_{T\pi'}^{\mathfrak{M}}$. Also, for any domain element $a \in \mathcal{D}_{T\theta\pi}^{\mathfrak{J}}$, the variable valuation $\varphi[u_T : T\theta \mapsto a]$ in \mathfrak{J} produces the valuation $\varphi'[u : T \mapsto a]$ in \mathfrak{M} . Thus:

$$\begin{aligned} \mathfrak{J}_{\pi, \varphi}(\text{DIS}_\theta(\forall(u : T) F)) &= \bigwedge_{a \in \mathcal{D}_{T\theta\pi}^{\mathfrak{J}}} \mathfrak{J}_{\pi, \varphi[u_T : T\theta \mapsto a]}(\text{DIS}_\theta(F)) \\ &= \bigwedge_{a \in \mathcal{D}_{T\pi'}^{\mathfrak{M}}} \mathfrak{M}_{\pi', \varphi'[u : T \mapsto a]}(F) = \mathfrak{M}_{\pi', \varphi'}(\forall(u : T) F) \end{aligned}$$

Thus, for every formula F in Γ , $\text{DIS}(F)$ is satisfied under \mathfrak{J} . \square

Lemma 3. *For any closed formula F and an arbitrary type valuation π , there exists a unique maximal (with respect to the set of instantiated type variables) type substitution in $\Theta^*(F)$ compatible with (i.e. included in) π .*

Proof. Consider the set S of all type substitutions in $\Theta^*(F)$ included in π . This set is necessarily finite, since F contains a finite number of type variables. The type substitutions in S are pairwise compatible. Therefore, S is a complete upper semi-lattice and its maximal element is included in π . \square

Theorem 4 (Completeness of DIS). *For every set of closed formulas Γ , if $\text{DIS}(\Gamma)$ is satisfiable then Γ is satisfiable.*

Proof. Let \mathfrak{M} be a model of $\text{DIS}(\Gamma)$. We are going to construct an interpretation \mathfrak{J} of Γ such that for every closed formula H in Γ , every type valuation π , and every variable valuation φ in \mathfrak{J} , there exist a type substitution $\theta \in \Theta^*(H)$, a type valuation π' , and a variable valuation φ' in the original model \mathfrak{M} such that:

$$\mathfrak{J}_{\pi, \varphi}(t) = \mathfrak{M}_{\pi', \varphi'}(\text{DIS}_\theta(t)) \quad \mathfrak{J}_{\pi, \varphi}(F) = \mathfrak{M}_{\pi', \varphi'}(\text{DIS}_\theta(F))$$

for every term t and formula F occurring in H . Since for every $\theta \in \Theta^*(H)$, the model \mathfrak{M} satisfies $\text{DIS}_\theta(H)$, the initial formula H is satisfied under \mathfrak{J} .

For any sort S in the signature of Γ , we set $\mathcal{D}_S^{\mathfrak{J}} \triangleq \mathcal{D}_S^{\mathfrak{M}}$. Consider a function symbol f in Γ , a vector of sorts $\mathbf{T} = \langle T_1, \dots, T_n, T \rangle$, and a type substitution τ that instantiates the type signature of f to \mathbf{T} . If the monomorphic specialization f_τ belongs to W , we define $f_{\mathbf{T}}^{\mathfrak{J}} \triangleq (f_\tau)_{\mathbf{T}}^{\mathfrak{M}}$. Otherwise, $f_{\mathbf{T}}^{\mathfrak{J}} \triangleq f_{\mathbf{T}}^{\mathfrak{M}}$. The predicate symbols are interpreted in the same way.

Now, consider a formula H in Γ , a type valuation π , and a variable valuation φ in \mathfrak{J} . By Lemma 3, $\Theta^*(H)$ contains a maximal monomorphic type substitution θ compatible with π (and thus included in π). We set $\pi' \triangleq \pi$, so that $\theta \circ \pi' = \pi$. For every variable $u : T$ occurring in H , we set $\varphi'(u_T : T\theta) \triangleq \varphi(u : T)$. Indeed for every type in F , $\mathcal{D}_{T\pi}^{\mathfrak{J}} = \mathcal{D}_{T\pi'}^{\mathfrak{M}} = \mathcal{D}_{T\theta\pi}^{\mathfrak{M}}$. By Lemma 1, it is sufficient to define φ' on variables that occur in $\text{DIS}(H)$.

Using this definitions, we proceed to prove the two equalities above. The proof goes by induction on the structure of terms and formulas. We consider three cases: a variable, a function application, and a universally quantified formula (other cases can be handled along the same lines).

1. Consider a variable $u:T$ occurring in H . We obtain:

$$\mathfrak{J}_{\pi,\varphi}(u:T) = \varphi(u:T) = \varphi'(u_T:T\theta) = \mathfrak{M}_{\pi',\varphi'}(\text{DIS}_\theta(u:T))$$

2. Consider a term $t = f(t_1:T_1, \dots, t_n:T_n):T$ occurring in H . Let \mathbf{T} denote $\langle T_1, \dots, T_n, T \rangle$ and let τ be the type substitution that instantiates the type signature of f to $\mathbf{T}\theta$. Suppose that the specialization f_τ belongs to W . Thus DIS_θ replaces f with f_τ and $\mathbf{T}\pi = \mathbf{T}\theta$. We obtain:

$$\begin{aligned} \mathfrak{J}_{\pi,\varphi}(t) &= f_{\mathbf{T}\pi}^{\mathfrak{J}}(\mathfrak{J}_{\pi,\varphi}(t_1), \dots, \mathfrak{J}_{\pi,\varphi}(t_n)) \\ &= (f_\tau)_{\mathbf{T}\theta\pi'}^{\mathfrak{M}}(\mathfrak{M}_{\pi',\varphi'}(\text{DIS}_\theta(t_1)), \dots, \mathfrak{M}_{\pi',\varphi'}(\text{DIS}_\theta(t_n))) \\ &= \mathfrak{M}_{\pi',\varphi'}(f_\tau(\text{DIS}_\theta(t_1), \dots, \text{DIS}_\theta(t_n))) \\ &= \mathfrak{M}_{\pi',\varphi'}(\text{DIS}_\theta(t)) \end{aligned}$$

Now, let us assume that f_τ does not belong to W , so that DIS_θ keeps the symbol f in t . Then we can show that there is no monomorphic specialization of f in W with the type signature $\mathbf{T}\pi$. Indeed, if there is one, then the restriction of π to the type variables of \mathbf{T} belongs to $\Theta(H)$. This restriction is compatible with θ , since both of them are included in π . As θ is maximal, this restriction is also included in θ . This gives us $\mathbf{T}\pi = \mathbf{T}\theta$, which leads to a contradiction. We obtain:

$$\begin{aligned} \mathfrak{J}_{\pi,\varphi}(t) &= f_{\mathbf{T}\pi}^{\mathfrak{J}}(\mathfrak{J}_{\pi,\varphi}(t_1), \dots, \mathfrak{J}_{\pi,\varphi}(t_n)) \\ &= f_{\mathbf{T}\theta\pi'}^{\mathfrak{M}}(\mathfrak{M}_{\pi',\varphi'}(\text{DIS}_\theta(t_1)), \dots, \mathfrak{M}_{\pi',\varphi'}(\text{DIS}_\theta(t_n))) \\ &= \mathfrak{M}_{\pi',\varphi'}(f(\text{DIS}_\theta(t_1), \dots, \text{DIS}_\theta(t_n))) \\ &= \mathfrak{M}_{\pi',\varphi'}(\text{DIS}_\theta(t)) \end{aligned}$$

3. Consider a universally quantified formula $\forall(u:T) F$. As we have seen above, $\mathcal{D}_{T\pi}^{\mathfrak{J}} = \mathcal{D}_{T\theta\pi'}^{\mathfrak{M}}$. Also, for any $a \in \mathcal{D}_{T\theta\pi'}^{\mathfrak{M}}$, the variable valuation $\varphi[u:T \mapsto a]$ in \mathfrak{J} produces the valuation $\varphi'[u_T:T\theta \mapsto a]$ in \mathfrak{M} . We obtain:

$$\begin{aligned} \mathfrak{J}_{\pi,\varphi}(\forall(u:T) F) &= \bigwedge_{a \in \mathcal{D}_{T\pi}^{\mathfrak{J}}} \mathfrak{J}_{\pi,\varphi[u:T \mapsto a]}(F) \\ &= \bigwedge_{a \in \mathcal{D}_{T\theta\pi'}^{\mathfrak{M}}} \mathfrak{M}_{\pi',\varphi'[u_T:T\theta \mapsto a]}(\text{DIS}_\theta(F)) \\ &= \mathfrak{M}_{\pi',\varphi'}(\forall(u_T:T\theta) \text{DIS}_\theta(F)) = \mathfrak{M}_{\pi',\varphi'}(\text{DIS}_\theta(\forall(u:T) F)) \end{aligned}$$

Thus, for every formula H in Γ , H is satisfied under \mathfrak{J} . \square

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 $\Theta(H)$ 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 Partial Monomorphisation

It turns out that the above-presented approach of type instantiation with symbol discrimination can be advanced to actually protect types in the sense of the introduction to Section 3.

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 this section and the PAR transformation is parametrized by it. Given a set of polymorphic types M , we define a set of monomorphic substitutions $\Theta(M)$ as follows:

$$\Theta(M) \triangleq \{ \theta \mid M \text{ contains a type } T \text{ such that} \\ \theta \text{ only instantiates the type variables of } T \text{ and} \\ T\theta \text{ belongs to } U \}$$

The closure of $\Theta(M)$ with respect to finite unions, denoted $\Theta^*(M)$, is defined as in the previous section. Recall that the identity substitution always belongs to the closure.

Lemma 4. *Let M be a set of polymorphic types and π a type substitution. Then there exists a unique monomorphic type substitution $\theta \in \Theta^*(M)$ such that:*

- θ is included in π (that is, for every α , either $\alpha\theta = \alpha$ or $\alpha\theta = \alpha\pi$);
- for every type $T \in M$, $T\pi \in U$ if and only if $T\theta \in U$.

Proof. Let M' be the set of types from M that are instantiated “into U ” by π . Let θ be the restriction of π onto the type variables occurring in M' . It is easy to see that θ is a finite union of substitutions in $\Theta(M)$ that instantiate individual types in M' into U . Indeed, all of these individual substitutions are included in π and thus are compatible. Now, let T be an arbitrary type in M . If $T\pi$ belongs to U , then $T \in M'$ and hence $T\theta = T\pi$. If $T\theta$ belongs to U , then $T\pi \in U$, too, since all types in U are closed and θ is included in π .

Now, suppose that there is another monomorphic substitution θ' in $\Theta^*(M)$ that satisfies both conditions. Since both θ and θ' are included in π , they cannot substitute two different sorts for the same type variable. Also, θ' must coincide with θ (and π) on every type variable in M' . Therefore, θ is included in θ' . If θ' differs from θ , then there exists some type variable α outside M' instantiated by θ' . By definition of $\Theta(M)$, there must therefore exist a type $T \in M \setminus M'$ such that $T\theta' \in U$. This contradicts the second condition. \square

Given a type T , the transformed type $[T]$ is defined as follows: if $T \in U$, we set $[T] \triangleq \overline{T}$, otherwise $[T] \triangleq T$. Now, let f be a function symbol of type signature $\langle S_1, \dots, S_n, S \rangle$ in Γ . Let τ be a type substitution in $\Theta^*(\{S_1, \dots, S_n, S\})$. A fresh function symbol \bar{f}_τ with the type signature $\langle [S_1\tau], \dots, [S_n\tau], [S\tau] \rangle$ will be called a *protected specialization* of f . Protected specializations of predicate symbols are defined in the same way.

The idea is to produce every specialization that makes protected sorts appear in the type signature. For example, let M be $\{M(l, l)\}$. Then the symbol \mathbf{get} with the type signature $\langle M(\alpha, \beta), \alpha, \beta \rangle$ induces five protected specializations:

$$\begin{aligned} \mathbf{get}_\square &: \langle M(\alpha, \beta), \alpha, \beta \rangle \\ \mathbf{get}_{[l/\alpha, l/\beta]} &: \langle \overline{M(l, l)}, l, l \rangle \\ \mathbf{get}_{[M(l, l)/\alpha]} &: \langle M(M(l, l), \beta), \overline{M(l, l)}, \beta \rangle \\ \mathbf{get}_{[M(l, l)/\beta]} &: \langle M(\alpha, M(l, l)), \alpha, \overline{M(l, l)} \rangle \\ \mathbf{get}_{[M(l, l)/\alpha, M(l, l)/\beta]} &: \langle M(M(l, l), M(l, l)), \overline{M(l, l)}, \overline{M(l, l)} \rangle \end{aligned}$$

The PAR transformation modifies the signature of a transformed theory:

1. For every variable symbol u and type T , we add a new variable symbol u_T .
2. We replace every function and predicate symbol with the set of all its protected specializations.

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

1. Given a variable $u : T$, $\text{PAR}_\theta(u : T) \triangleq u_T : [T\theta]$.
2. Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$. By Lemma 4, there is a unique protected specialization \bar{f}_τ whose signature matches $\langle [T_1\theta], \dots, [T_n\theta], [T\theta] \rangle$. Indeed, if $T_i\theta$ is in U , then the only type that matches $[T_i\theta]$ is $[T_i\theta]$ itself. Then $\text{PAR}_\theta(t) \triangleq \bar{f}_\tau(\text{PAR}_\theta(t_1), \dots, \text{PAR}_\theta(t_n)) : [T\theta]$.
3. Consider an atomic formula $F = p(t_1 : T_1, \dots, t_n : T_n)$. By Lemma 4, there is one protected specialization \bar{p}_τ whose signature matches $\langle [T_1\theta], \dots, [T_n\theta] \rangle$. Then $\text{PAR}_\theta(F) \triangleq \bar{p}_\tau(\text{PAR}_\theta(t_1), \dots, \text{PAR}_\theta(t_n))$.

Equalities and complex formulas are converted in a natural way:

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

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

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

where M is the set of types of all terms that occur in F .

On our running example, supposing that $U = \{1\}$, PAR produces the following two formulas (via $\text{PAR}_{[\]}$ and $\text{PAR}_{[1/\alpha]}$, respectively):

$$\begin{aligned} & \forall(m : \mathbf{M}(\alpha, 1)) \forall(c : \alpha) \text{get}_{[1/\beta]}(\text{set}_{[1/\beta]}(m, c, 6), c) * 7 \approx 42 \\ & \forall(m : \mathbf{M}(1, 1)) \forall(c : \bar{1}) \text{get}_{[1/\alpha, 1/\beta]}(\text{set}_{[1/\alpha, 1/\beta]}(m, c, 6), c) * 7 \approx 42 \end{aligned}$$

The new symbol $\text{get}_{[1/\beta]}$ has the type signature $\langle \mathbf{M}(\alpha, 1), \alpha, \bar{1} \rangle$ and $\text{get}_{[1/\alpha, 1/\beta]}$ has the monomorphic type signature $\langle \mathbf{M}(1, 1), \bar{1}, \bar{1} \rangle$.

Theorem 5. *Let θ be a type substitution, t a term of type T , and F a formula. Then $\text{PAR}_\theta(t)$ is a well-formed term of type $[T\theta]$ and $\text{PAR}_\theta(F)$ is a well-formed formula such that $\text{FV}(\text{PAR}_\theta(F)) = \{\text{PAR}_\theta(x) \mid x \in \text{FV}(F)\}$ and $\text{FV}_T(\text{PAR}_\theta(F)) = \text{FV}_T(F) \setminus \text{dom}(\theta)$.*

Proof. The proof goes by induction on the structure of terms and formulas. The only interesting case is the application of a function symbol (predicate symbols are treated in the same way). Let t be a term of the form $f(t_1 : T_1, \dots, t_n : T_n) : T$. For every t_i , $\text{PAR}_\theta(t_i)$ is a well-formed term of type $[T_i\theta]$ by the induction hypothesis. By definition, the type signature of \bar{f}_τ matches $\langle [T_1\theta], \dots, [T_n\theta], [T\theta] \rangle$. Thus, $\text{PAR}_\theta(t)$ is a well-formed term of type $[T\theta]$. \square

Theorem 6 (Soundness of PAR). *If a set of closed formulas Γ is satisfiable then $\text{PAR}(\Gamma)$ is satisfiable.*

Proof. Let \mathfrak{M} be a model of Γ . We are going to construct an interpretation \mathfrak{J} of $\text{PAR}(\Gamma)$ such that, for every type substitution θ , type valuation π , and variable valuation φ in \mathfrak{J} , there exist a type valuation π' and a variable valuation φ' in the original model \mathfrak{M} , such that:

$$\mathfrak{J}_{\pi, \varphi}(\text{PAR}_\theta(t)) = \mathfrak{M}_{\pi', \varphi'}(t) \quad \mathfrak{J}_{\pi, \varphi}(\text{PAR}_\theta(F)) = \mathfrak{M}_{\pi', \varphi'}(F)$$

for every term t and formula F occurring in Γ . Since for every closed $F \in \Gamma$ and every θ , \mathfrak{J} satisfies $\text{PAR}_\theta(F)$, \mathfrak{J} is a model of $\text{PAR}(\Gamma)$.

For any sort S in Γ , we define $\mathcal{D}_S^{\mathfrak{J}} \triangleq \mathcal{D}_S^{\mathfrak{M}}$ and $\mathcal{D}_{\bar{S}}^{\mathfrak{J}} \triangleq \mathcal{D}_S^{\mathfrak{M}}$. Then we give an interpretation to the function and predicate symbols in $\text{PAR}(\Gamma)$. Consider a function symbol f from Γ and its protected specialization \bar{f}_τ . Denoting the type signature of f with \mathbf{S} , the type signature of \bar{f}_τ is $[\mathbf{S}\tau]$. Let \mathbf{T} be a vector of sorts such that $\mathbf{T} = [\mathbf{S}\tau]\tau'$ for some type substitution τ' . Then we interpret \bar{f}_τ on \mathbf{T} in \mathfrak{J} exactly as f is interpreted on $\mathbf{S}\tau\tau'$ in \mathfrak{M} . This is correct, as $\mathbf{S}\tau\tau'$ is just \mathbf{T} with protection removed. Predicate symbols are interpreted in the same way.

Now, let us consider some type substitution θ , a type valuation π , and a variable valuation φ in \mathfrak{J} . We set $\pi' \triangleq \theta \circ \pi$ and we set $\varphi'(u : T) \triangleq \varphi(u_T : [T\theta])$. Notice that for any type T in Γ , $\mathcal{D}_{[T\theta]\pi}^{\mathfrak{J}} = \mathcal{D}_{T\theta\pi}^{\mathfrak{M}} = \mathcal{D}_{T\pi'}^{\mathfrak{M}}$.

Using this definitions, we proceed to prove the two equalities above. The proof goes by induction on the structure of terms and formulas. We consider three cases: a variable, a function application, and a universally quantified formula (other cases can be handled along the same lines).

1. Consider a variable $u : T$ occurring in Γ . We obtain:

$$\mathfrak{J}_{\pi, \varphi}(\text{PAR}_{\theta}(u : T)) = \varphi(u_T : [T\theta]) = \varphi'(u : T) = \mathfrak{M}_{\pi', \varphi'}(u : T)$$

2. Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$. Let f_{τ} be the protected specialization of f whose type signature matches $[T\theta]$, where $\mathbf{T} = \langle T_1, \dots, T_n, T \rangle$.

$$\begin{aligned} \mathfrak{J}_{\pi, \varphi}(\text{PAR}_{\theta}(t)) &= \mathfrak{J}_{\pi, \varphi}(\bar{f}_{\tau}(\text{PAR}_{\theta}(t_1), \dots, \text{PAR}_{\theta}(t_n)) : [T\theta]) \\ &= (f_{\tau})_{[\mathbf{T}\theta]_{\pi}}^{\mathfrak{J}}(\mathfrak{J}_{\pi, \varphi}(\text{PAR}_{\theta}(t_1)), \dots, \mathfrak{J}_{\pi, \varphi}(\text{PAR}_{\theta}(t_n))) \\ &= f_{\mathbf{T}\theta\pi}^{\mathfrak{M}}(\mathfrak{M}_{\pi', \varphi'}(t_1), \dots, \mathfrak{M}_{\pi', \varphi'}(t_n)) \\ &= f_{\mathbf{T}\pi'}^{\mathfrak{M}}(\mathfrak{M}_{\pi', \varphi'}(t_1), \dots, \mathfrak{M}_{\pi', \varphi'}(t_n)) \\ &= \mathfrak{M}_{\pi', \varphi'}(t) \end{aligned}$$

3. Consider a universally quantified formula $\forall(u : T) F$. As we have shown above, $\mathcal{D}_{[T\theta]_{\pi}}^{\mathfrak{J}} = \mathcal{D}_{T\pi'}^{\mathfrak{M}}$. Also, for any domain element $a \in \mathcal{D}_{[T\theta]_{\pi}}^{\mathfrak{J}}$, the variable valuation $\varphi[u_T : [T\theta] \mapsto a]$ in \mathfrak{J} produces the valuation $\varphi'[u : T \mapsto a]$ in \mathfrak{M} .

$$\begin{aligned} \mathfrak{J}_{\pi, \varphi}(\text{PAR}_{\theta}(\forall(u : T) F)) &= \bigwedge_{a \in \mathcal{D}_{[T\theta]_{\pi}}^{\mathfrak{J}}} \mathfrak{J}_{\pi, \varphi[u_T : [T\theta] \mapsto a]}(\text{PAR}_{\theta}(F)) \\ &= \bigwedge_{a \in \mathcal{D}_{T\pi'}^{\mathfrak{M}}} \mathfrak{M}_{\pi', \varphi'[u : T \mapsto a]}(F) = \mathfrak{M}_{\pi', \varphi'}(\forall(u : T) F) \end{aligned}$$

Thus, for every formula F in Γ , $\text{PAR}(F)$ is satisfied under \mathfrak{J} . \square

Theorem 7 (Completeness of PAR). *For every set of closed formulas Γ , if $\text{PAR}(\Gamma)$ is satisfiable then Γ is satisfiable.*

Proof. Let \mathfrak{M} be a model of $\text{PAR}(\Gamma)$. We are going to construct an interpretation \mathfrak{J} of Γ such that for every formula H in Γ , every type valuation π , and every variable interpretation φ in \mathfrak{J} , there exist a type substitution $\theta \in \Theta^*(\{T \mid T \text{ is the type of a term in } H\})$, a type valuation π' , and a variable valuation φ' in the original model \mathfrak{M} such that:

$$\mathfrak{J}_{\pi, \varphi}(t) = \mathfrak{M}_{\pi', \varphi'}(\text{PAR}_{\theta}(t)) \quad \mathfrak{J}_{\pi, \varphi}(F) = \mathfrak{M}_{\pi', \varphi'}(\text{PAR}_{\theta}(F))$$

for every term t and formula F occurring in H .

For any sort S in the signature of Γ , we set $\mathcal{D}_S^{\mathfrak{J}} \triangleq \mathcal{D}_{[S]}^{\mathfrak{M}}$. Consider a function symbol f in Γ and a vector of sorts $\mathbf{T} = \langle T_1, \dots, T_n, T \rangle$ matched by the type signature of f . By Lemma 4, there is a unique protected specialization \bar{f}_{τ} whose type signature matches the sort vector $[\mathbf{T}]$. Then we define $f_{\mathbf{T}}^{\mathfrak{J}} \triangleq (\bar{f}_{\tau})_{[\mathbf{T}]}^{\mathfrak{M}}$. The predicate symbols are interpreted in the same way.

Now, consider a formula H from Γ , a type valuation π , and a variable valuation φ in \mathfrak{J} . Let M be the set of types of all terms occurring in H . By Lemma 4, there exists a unique monomorphic type substitution $\theta \in \Theta^*(M)$ such that θ is included in π and for every type $T \in M$, $T\pi \in U$ if and only if $T\theta \in U$. We set

$\pi' \triangleq \pi$, so that $\theta \circ \pi' = \pi$. Let a variable $u:T$ occur in H . If it occurs in H as a term, then $T \in M$ and we set $\varphi'(u_T:[T\theta]) \triangleq \varphi(u:T)$. Indeed, for every type T in M , $\mathcal{D}_{T\pi}^{\mathfrak{J}} = \mathcal{D}_{[T\pi]}^{\mathfrak{M}} = \mathcal{D}_{[T\theta]\pi'}^{\mathfrak{M}}$ by definition of θ . Otherwise, if $u:T$ is only bound by a quantifier in H but does not occur under this quantifier, we can safely ignore it in variable valuations by Lemma 1.

Using this definitions, we proceed to prove the two equalities above. The proof goes by induction on the structure of terms and formulas. We consider three cases: a variable, a function application, and a universally quantified formula (other cases can be handled along the same lines).

1. Consider a variable $u:T$ occurring in H as a term.

$$\mathfrak{J}_{\pi,\varphi}(u:T) = \varphi(u:T) = \varphi'(u_T:[T\theta]) = \mathfrak{M}_{\pi',\varphi'}(\text{PAR}_{\theta}(u:T))$$

2. Consider a term $t = f(t_1:T_1, \dots, t_n:T_n):T$ occurring in H and let \mathbf{T} denote $\langle T_1, \dots, T_n, T \rangle$. Let \bar{f}_{τ} be the protected specialization of f whose type signature matches $[\mathbf{T}\theta]$. By definition of θ , $[\mathbf{T}\theta]$ matches $[\mathbf{T}\pi]$. Then \bar{f}_{τ} is also the only protected specialization of f whose type signature matches $[\mathbf{T}\pi]$.

$$\begin{aligned} \mathfrak{J}_{\pi,\varphi}(t) &= \bar{f}_{\tau}^{\mathfrak{J}}(\mathfrak{J}_{\pi,\varphi}(t_1), \dots, \mathfrak{J}_{\pi,\varphi}(t_n)) \\ &= (\bar{f}_{\tau})_{[\mathbf{T}\pi]}^{\mathfrak{M}}(\mathfrak{M}_{\pi',\varphi'}(\text{PAR}_{\theta}(t_1)), \dots, \mathfrak{M}_{\pi',\varphi'}(\text{PAR}_{\theta}(t_n))) \\ &= (\bar{f}_{\tau})_{[\mathbf{T}\theta]\pi'}^{\mathfrak{M}}(\mathfrak{M}_{\pi',\varphi'}(\text{PAR}_{\theta}(t_1)), \dots, \mathfrak{M}_{\pi',\varphi'}(\text{PAR}_{\theta}(t_n))) \\ &= \mathfrak{M}_{\pi',\varphi'}(\bar{f}_{\tau}(\text{PAR}_{\theta}(t_1), \dots, \text{PAR}_{\theta}(t_n))) = \mathfrak{M}_{\pi',\varphi'}(\text{PAR}_{\theta}(t)) \end{aligned}$$

3. Consider a universally quantified formula $\forall(u:T) F$ occurring in H . We can assume that $u:T$ occurs in F and $T \in M$. Then $\mathcal{D}_{T\pi}^{\mathfrak{J}} = \mathcal{D}_{[T\theta]\pi'}^{\mathfrak{M}}$. Also, for any element $a \in \mathcal{D}_{[T\theta]\pi'}^{\mathfrak{M}}$, the variable valuation $\varphi[u:T \mapsto a]$ in \mathfrak{J} produces the valuation $\varphi'[u_T:[T\theta] \mapsto a]$ in \mathfrak{M} . We obtain:

$$\begin{aligned} \mathfrak{J}_{\pi,\varphi}(\forall(u:T) F) &= \bigwedge_{a \in \mathcal{D}_{T\pi}^{\mathfrak{J}}} \mathfrak{J}_{\pi,\varphi[u:T \mapsto a]}(F) \\ &= \bigwedge_{a \in \mathcal{D}_{[T\theta]\pi'}^{\mathfrak{M}}} \mathfrak{M}_{\pi',\varphi'[u_T:[T\theta] \mapsto a]}(\text{PAR}_{\theta}(F)) \\ &= \mathfrak{M}_{\pi',\varphi'}(\forall(u_T:[T\theta]) \text{PAR}_{\theta}(F)) = \mathfrak{M}_{\pi',\varphi'}(\text{PAR}_{\theta}(\forall(u:T) F)) \end{aligned}$$

Thus, for every formula H in Γ , H is satisfied under \mathfrak{J} . \square

In practice, the PAR transformation behaves reasonably well on small sets U , containing two or three sorts. If we try to protect more — say, every sort constant occurring in the problem — the exponential growth due to type instantiation (w.r.t. the number of type variables in a formula) quickly makes the proof search infeasible. Therefore, we do not consider PAR as a practical alternative to the TW method, presented in Section 3.3. Yet, in our view, the existence of an alternative type protection technique justifies the separation of type protection from type fusion and shows that the notion of protected sort gives a good independent characterization of the boundary between the two stages.

3.3 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 protect as in the previous section. 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 $\mathbf{to}_T : \langle \bar{T}, T \rangle$ and $\mathbf{from}_T : \langle T, \bar{T} \rangle$.

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 .

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

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

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:

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

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

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

Theorem 8. *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.*

Proof. The proof goes by induction on the structure of terms and formulas. The only interesting case is the application of a function symbol (predicate symbols are treated in the same way). The conservation of free (type) variables is trivial.

Let t be a term of the form $f(t_1 : T_1, \dots, t_n : T_n) : T$. Let $\langle S_1, \dots, S_n, S \rangle$ be the type signature of f and τ the type substitution that instantiates $\langle S_1, \dots, S_n, S \rangle$ to $\langle T_1, \dots, T_n, T \rangle$. For every argument t_i , $\text{Tw}(t_i)$ is a well-formed term of type $[T_i]$ by the induction hypothesis, and there are three cases to consider:

1. $T_i \in U$ and $S_i \notin U$. Then $\text{Tw}(t_i)$ has the type \bar{T}_i and $\text{to}_{T_i}(\text{Tw}(t_i)) : T_i$ is a well-formed term. Notice that the i -th element of the type signature of \bar{f} is $[S_i] = S_i$, and $S_i\tau = T_i$.
2. $S_i \in U$. Then S_i is closed, so $T_i = S_i$, and $[T_i] = [S_i] = [S_i]\tau$.
3. $S_i \notin U$ and $T_i \notin U$. Then $[S_i]\tau = S_i\tau = T_i = [T_i]$.

Then we apply \bar{f} to t'_1, \dots, t'_n using τ to instantiate the type signature of \bar{f} . Again, we have three possible cases:

1. $T \in U$ and $S \notin U$. Since $S\tau = T$, we obtain $[S]\tau = S\tau = T$. Then the term $\text{from}_T(\bar{f}(t'_1, \dots, t'_n) : T) : \bar{T}$ is well-formed and has the type $[T] = \bar{T}$.
2. $S \in U$. Then S is closed, so $T = S$ and $[T] = [S] = [S]\tau$.
3. $S \notin U$ and $T \notin U$. Then $[S]\tau = S\tau = T = [T]$.

Thus, $\text{Tw}(t)$ is a well-formed term of type $[T]$. □

Theorem 9 (Soundness of Tw). *If a set of closed formulas Γ is satisfiable then $\text{Tw}(\Gamma)$ is satisfiable.*

Proof. Let \mathfrak{M} be a model of Γ . We are going to construct an interpretation \mathfrak{J} of $\text{Tw}(\Gamma)$ such that, for every type valuation π and every variable valuation φ in \mathfrak{J} , there exist a type valuation π' and a variable valuation φ' in the original model \mathfrak{M} , such that:

$$\mathfrak{J}_{\pi, \varphi}(\text{Tw}(t)) = \mathfrak{M}_{\pi', \varphi'}(t) \quad \mathfrak{J}_{\pi, \varphi}(\text{Tw}(F)) = \mathfrak{M}_{\pi', \varphi'}(F)$$

for every term t and formula F occurring in Γ .

First of all, let us denote with $] [$ the “protection erasure” operation: for any protected sort \bar{S} , $] \bar{S} [\triangleq S$ and for any non-protected type T , $] \bar{T} [\triangleq T$. Obviously, for any type T in the initial signature, $] [T] [= T$. On the other hand, it is not generally true that $] [T] [= T$.

For any sort S in the initial signature, we define $\mathcal{D}_S^{\mathfrak{J}} \triangleq \mathcal{D}_S^{\mathfrak{M}}$ and $\mathcal{D}_{\bar{S}}^{\mathfrak{J}} \triangleq \mathcal{D}_S^{\mathfrak{M}}$. Then we give an interpretation to the function and predicate symbols in $\text{Tw}(\Gamma)$. Apart from the bridge function symbols, every symbol in the new signature originates from some initial symbol, as described above. Consider a function symbol f with the type signature \mathbf{S} in the initial set Γ . In $\text{Tw}(\Gamma)$, it is replaced with the symbol \bar{f} of type signature $[\mathbf{S}]$. Let \mathbf{T} be a sort vector matched by $[\mathbf{S}]$.

It is easy to see that the “unprotected” sort vector \mathbf{T} is matched by \mathbf{S} . Thus we can define $\bar{f}_{\mathbf{T}}^{\mathfrak{J}} \triangleq f_{\mathbf{T}}^{\mathfrak{M}}$. Predicate symbols in $\text{Tw}(\Gamma)$ are interpreted in the same way. Finally, for every sort $T \in U$, the function symbols \mathbf{to}_T and \mathbf{from}_T are interpreted as identities. This finishes the definition of the interpretation \mathfrak{J} .

Now let us take a type valuation π and a variable valuation φ in \mathfrak{J} . We set $\pi' \triangleq \pi$ and $\varphi'(u:T) \triangleq \varphi(u:[T])$. By construction, $\mathcal{D}_{[T]\pi}^{\mathfrak{J}} = \mathcal{D}_{T\pi}^{\mathfrak{J}} = \mathcal{D}_{T\pi'}^{\mathfrak{M}}$.

Using this definitions, we can proceed to prove the two equalities above. The proof goes by induction on the structure of terms and formulas. We will consider three cases: a variable, a function application, and a universally quantified formula (other cases can be handled along the same lines).

1. Let $u:T$ be a variable occurring in Γ .

$$\mathfrak{J}_{\pi,\varphi}(\text{Tw}(u:T)) = \mathfrak{J}_{\pi,\varphi}(u:[T]) = \varphi(u:[T]) = \varphi'(u:T) = \mathfrak{M}_{\pi',\varphi'}(u:T)$$

2. Consider a term $t = f(t_1:T_1, \dots, t_n:T_n):T$. The transformed term $\text{Tw}(t)$ is an application of the function symbol \bar{f} to the transformed arguments $\text{Tw}(t_1), \dots, \text{Tw}(t_n)$ with added, where necessary, functions \mathbf{to}_{\square} and \mathbf{from}_{\square} . By the induction hypothesis and knowing that
 - function symbols \mathbf{to}_{\square} and \mathbf{from}_{\square} are interpreted in \mathfrak{J} as identities, and
 - \bar{f} on a vector sort \mathbf{T} is interpreted in \mathfrak{J} as f on \mathbf{T} in \mathfrak{M} ,
we can conclude that $\mathfrak{J}_{\pi,\varphi}(\text{Tw}(t)) = \mathfrak{M}_{\pi',\varphi'}(t)$.

3. Consider a universally quantified formula $\forall(u:T) F$ from Γ . As we know, $\mathcal{D}_{[T]\pi}^{\mathfrak{J}} = \mathcal{D}_{T\pi'}^{\mathfrak{M}}$. Moreover, for any element $a \in \mathcal{D}_{[T]\pi}^{\mathfrak{J}}$, the variable valuation $\varphi[u:[T] \mapsto a]$ in \mathfrak{J} produces the valuation $\varphi'[u:T \mapsto a]$ in \mathfrak{M} . We obtain:

$$\begin{aligned} \mathfrak{J}_{\pi,\varphi}(\text{Tw}(\forall(u:T) F)) &= \bigwedge_{a \in \mathcal{D}_{[T]\pi}^{\mathfrak{J}}} \mathfrak{J}_{\pi,\varphi[u:[T] \mapsto a]}(\text{Tw}(F)) \\ &= \bigwedge_{a \in \mathcal{D}_{T\pi'}^{\mathfrak{M}}} \mathfrak{M}_{\pi',\varphi'[u:T \mapsto a]}(F) = \mathfrak{M}_{\pi',\varphi'}(\forall(u:T) F) \end{aligned}$$

Thus, for every formula F in Γ , $\text{Tw}(F)$ is satisfied under \mathfrak{J} . The bijection axioms for \mathbf{to}_{\square} and \mathbf{from}_{\square} are trivially satisfied, too, since both function symbols are interpreted as identities. \square

Theorem 10 (Completeness of Tw). *For every set of closed formulas Γ , if $\text{Tw}(\Gamma)$ is satisfiable then Γ is satisfiable.*

Proof. Let \mathfrak{M} be a model of $\text{Tw}(\Gamma)$. Due to the bijection axioms for bridge functions, we can assume without loss of generality that, for every sort $T \in U$, $\mathcal{D}_T^{\mathfrak{M}} = \mathcal{D}_T^{\mathfrak{M}}$ and both \mathbf{to}_T and \mathbf{from}_T are interpreted in \mathfrak{M} as identities.

We are going to construct an interpretation \mathfrak{J} of Γ such that, for every type valuation π and every variable valuation φ in \mathfrak{J} , there exist a valuation π' and a variable valuation φ' in the original model \mathfrak{M} , such that:

$$\mathfrak{J}_{\pi,\varphi}(t) = \mathfrak{M}_{\pi',\varphi'}(\text{Tw}(t)) \quad \mathfrak{J}_{\pi,\varphi}(F) = \mathfrak{M}_{\pi',\varphi'}(\text{Tw}(F))$$

for every term t and formula F occurring in Γ .

For any sort S in the signature of Γ , we define $\mathcal{D}_S^{\mathfrak{J}} \triangleq \mathcal{D}_S^{\mathfrak{M}}$. Then for any type T and its closed type instance $T\tau$, $\mathcal{D}_{T\tau}^{\mathfrak{J}} = \mathcal{D}_{[T]\tau}^{\mathfrak{M}}$. Indeed, if $T \in U$, then $\mathcal{D}_{T\tau}^{\mathfrak{J}} = \mathcal{D}_T^{\mathfrak{J}} = \mathcal{D}_T^{\mathfrak{M}} = \mathcal{D}_{\bar{T}}^{\mathfrak{M}} = \mathcal{D}_{[T]\tau}^{\mathfrak{M}}$, otherwise $[T] = T$.

Every function symbol f (predicate symbol p) of type signature \mathbf{S} in Γ is interpreted on a closed sort vector $\mathbf{S}\tau$ in \mathfrak{J} exactly as the corresponding symbol \bar{f} (respectively, \bar{p}) on $[\mathbf{S}]\tau$ in \mathfrak{M} . This concludes the definition of \mathfrak{J} .

Now, consider a type valuation π and a variable valuation φ in \mathfrak{J} . We set $\pi' \triangleq \pi$ and $\varphi'(u:[T]) \triangleq \varphi(u:T)$ for every variable $u:T$ occurring in Γ . By Lemma 1, it suffices to define φ' on variables that occur in transformed formulas. Also recall that $\mathcal{D}_{T\pi}^{\mathfrak{J}} = \mathcal{D}_{[T]\pi}^{\mathfrak{M}} = \mathcal{D}_{[T]\pi'}^{\mathfrak{M}}$.

Using this definitions, we can proceed to prove the two equalities above. The proof goes by induction on the structure of terms and formulas. We will consider three cases: a variable, a function application, and a universally quantified formula (other cases can be handled along the same lines).

1. Let $u:T$ be a variable occurring in Γ . We obtain:

$$\mathfrak{J}_{\pi,\varphi}(u:T) = \varphi(u:T) = \varphi'(u:[T]) = \mathfrak{M}_{\pi',\varphi'}(u:[T]) = \mathfrak{M}_{\pi',\varphi'}(\text{Tw}(u:T))$$

2. Consider a term $t = f(t_1:T_1, \dots, t_n:T_n):T$. Let $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$ be the type signature of f and τ a type substitution such that $\mathbf{S}\tau = \langle T_1, \dots, T_n, T \rangle$. The term $\text{Tw}(t)$ is of the form $\bar{f}(t'_1:[S_1]\tau, \dots, t'_n:[S_n]\tau):[S]\tau$, possibly under from_T , where every t'_i is $\text{Tw}(t_i)$, possibly under to_{T_i} . Knowing that:
 - bridge function symbols are interpreted in \mathfrak{M} as identities, and
 - f is interpreted on $\mathbf{S}\tau$ in \mathfrak{J} as \bar{f} on $[\mathbf{S}]\tau$ in \mathfrak{M} ,
we can conclude that $\mathfrak{J}_{\pi,\varphi}(t) = \mathfrak{M}_{\pi',\varphi'}(\text{Tw}(t))$.

3. Consider a universally quantified formula $\forall(u:T) F$. For any $a \in \mathcal{D}_{T\pi}^{\mathfrak{J}}$, the valuation $\varphi[u:T \mapsto a]$ in \mathfrak{J} produces the valuation $\varphi'[u:[T] \mapsto a]$ in \mathfrak{M} .

$$\begin{aligned} \mathfrak{J}_{\pi,\varphi}(\forall(u:T) F) &= \bigwedge_{a \in \mathcal{D}_{T\pi}^{\mathfrak{J}}} \mathfrak{J}_{\pi,\varphi[u:T \mapsto a]}(F) \\ &= \bigwedge_{a \in \mathcal{D}_{[T]\pi'}^{\mathfrak{M}}} \mathfrak{M}_{\pi',\varphi'[u:[T] \mapsto a]}(\text{Tw}(F)) \\ &= \mathfrak{M}_{\pi',\varphi'}(\forall(u:[T]) \text{Tw}(F)) = \mathfrak{M}_{\pi',\varphi'}(\text{Tw}(\forall(u:T) F)) \end{aligned}$$

Thus, for every formula F in Γ , F is satisfied under \mathfrak{J} . \square

3.4 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 $\mathbf{deco} : \langle T, U, D \rangle$.

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

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

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:

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

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

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

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

$$\begin{aligned} \forall(v^\alpha : T) \forall(m : U) \forall(c : U) \mathbf{from}_1(\mathbf{deco}(\mathbf{I}, \mathbf{get}(\mathbf{deco}(\mathbf{M}(v^\alpha), \mathbf{I}), \\ \mathbf{set}(\mathbf{deco}(\mathbf{M}(v^\alpha), \mathbf{I}), m), \mathbf{deco}(v^\alpha, c), \mathbf{deco}(\mathbf{I}, \mathbf{to}_1(\mathbf{6}))))), \\ \mathbf{deco}(v^\alpha, c))) * 7 \approx 42 \end{aligned}$$

whereas PAR and DEC^o produce:

$$\begin{aligned} \forall(v^\alpha : \mathbb{T}) \forall(m : \mathbb{U}) \forall(c : \mathbb{U}) \text{get}_{[1/\beta]}(\text{deco}(\mathbb{M}(v^\alpha, \mathbb{I}), \\ \text{set}_{[1/\beta]}(\text{deco}(\mathbb{M}(v^\alpha, \mathbb{I}), m), \text{deco}(v^\alpha, c), 6)), \text{deco}(v^\alpha, c)) * 7 \approx 42 \end{aligned}$$

$$\begin{aligned} \forall(m : \mathbb{U}) \forall(c : \bar{\mathbb{I}}) \text{get}_{[1/\alpha, 1/\beta]}(\text{deco}(\mathbb{M}(\mathbb{I}, \mathbb{I}), \\ \text{set}_{[1/\alpha, 1/\beta]}(\text{deco}(\mathbb{M}(\mathbb{I}, \mathbb{I}), m), c, 6)), c) * 7 \approx 42 \end{aligned}$$

Theorem 11. *For every term t of type T , DEC(t) is a well-formed monomorphic term of type $[T]^-$. For every formula F , DEC(F) is a well-formed monomorphic formula. Moreover,*

$$\begin{aligned} \text{FV}(\text{DEC}(t)) &= \{u_T : [T]^+ \mid u : T \in \text{FV}(t)\} \cup \{v^\alpha : \mathbb{T} \mid \alpha \in \text{FV}_T(t)\} \\ \text{FV}(\text{DEC}(F)) &= \{u_T : [T]^+ \mid u : T \in \text{FV}(F)\} \cup \{v^\alpha : \mathbb{T} \mid \alpha \in \text{FV}_T(F)\} \end{aligned}$$

Finally, for every closed formula F , DEC^o(F) is a well-formed closed monomorphic formula.

Proof. The last claim easily follows from the second one. The proof of the first two statements goes by induction on the structure of terms and formulas. The only interesting case is the application of a function symbol (predicate symbols are treated in the same way).

Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$. Let S_i denote the i -th sort in the type signature of f . We have two possible cases: either S_i and T_i are protected and $S_i = T_i$, or both S_i and T_i are non-protected, by the definition of matching with protected types. Therefore, $[T_i]^- = [S_i]^-$ for every $i \in [1, n]$.

Now, let S be the type of the value in the type signature of f . By the same argument, either S and T are protected and $S = T$, or both S and T are non-protected. Then $[T]^+ = [S]^+$ and the term $t' = \hat{f}(\text{DEC}(t_1), \dots, \text{DEC}(t_n)) : [T]^+$ is well-formed.

Finally, if T is protected, then $\text{DEC}(t) = t'$ and $[T]^- = T = [T]^+$. Notice that the type T does not contain type variables. Otherwise, if T is non-protected, then $[T]^+ = \mathbb{U}$ and $\text{DEC}(t) = \text{deco}(\text{DEC}(T), t') : \mathbb{D}$ is a well-formed term of type $[T]^-$. Furthermore, every type variable α occurring in T appears as a free variable $v^\alpha : \mathbb{T}$ in the first argument of deco in $\text{DEC}(t)$. \square

Theorem 12 (Soundness of DEC). *If a set of closed formulas with protected sorts Γ is satisfiable then DEC^o(Γ) is satisfiable.*

Proof. Let \mathfrak{M} be a model of Γ . We are going to construct an interpretation \mathfrak{J} of DEC^o(Γ) such that, for every variable valuation φ in \mathfrak{J} , there exist a type valuation π' and a variable valuation φ' in the original model \mathfrak{M} , such that:

$$\mathfrak{J}_\varphi(\text{DEC}(t : T)) = \langle T\pi', \mathfrak{M}_{\pi', \varphi'}(t) \rangle \quad \mathfrak{J}_\varphi(\text{DEC}(F)) = \mathfrak{M}_{\pi', \varphi'}(F)$$

for every term t and formula F occurring in Γ . We do not need to consider type valuations in \mathfrak{J} , since the result of DEC° is entirely monomorphic. As soon as these two equalities are proved, it is easy to show that for every closed $H \in \Gamma$, the interpretation \mathfrak{J} satisfies $\text{DEC}^\circ(H)$:

$$\mathfrak{J}(\text{DEC}^\circ(H)) = \bigwedge_{\varphi} \mathfrak{J}_{\varphi}(\text{DEC}(H)) = \bigwedge_{\varphi} \mathfrak{M}_{\pi', \varphi'}(H) = \top$$

For every sort S in the initial signature, protected or not, we keep its domain, pairing its elements with S :

$$\mathcal{D}_S^{\mathfrak{J}} \triangleq \{S\} \times \mathcal{D}_S^{\mathfrak{M}}$$

The domain of \mathbf{D} is defined as a dependent sum over non-protected sorts:

$$\mathcal{D}_{\mathbf{D}}^{\mathfrak{J}} \triangleq \{ \langle S, c \rangle \mid S \in \mathcal{T}(\mathbb{F}_{\mathbf{T}}), c \in \mathcal{D}_S^{\mathfrak{M}} \}$$

The domain of \mathbf{U} is the set of all functions that map every non-protected sort S in $\mathcal{T}(\mathbb{F}_{\mathbf{T}})$ to an element of $\mathcal{D}_S^{\mathfrak{M}}$:

$$\mathcal{D}_{\mathbf{U}}^{\mathfrak{J}} \triangleq \{ S \in \mathcal{T}(\mathbb{F}_{\mathbf{T}}) \mapsto c \in \mathcal{D}_S^{\mathfrak{M}} \}$$

Finally, the domain of \mathbf{T} is the set of all non-protected sorts: $\mathcal{D}_{\mathbf{T}}^{\mathfrak{J}} \triangleq \mathcal{T}(\mathbb{F}_{\mathbf{T}})$.

Then we give an interpretation to the function and predicate symbols in $\text{DEC}^\circ(\Gamma)$. Let f be a function symbol in Γ of type signature $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$. The symbol \hat{f} has the type signature $\mathbf{S}' = \langle [S_1]^{-}, \dots, [S_n]^{-}, [S]^+ \rangle$ and this is the only sort vector to interpret \hat{f} on, since it is monomorphic. Firstly, suppose that S is protected. Then we define:

$$\hat{f}_{\mathbf{S}'}^{\mathfrak{J}}(\langle T_1, c_1 \rangle, \dots, \langle T_n, c_n \rangle) \triangleq \begin{cases} \langle S, f_{\langle T_1, \dots, T_n, S \rangle}^{\mathfrak{M}}(c_1, \dots, c_n) \rangle & \text{if } \mathbf{S} \text{ matches } \langle T_1, \dots, T_n, S \rangle, \\ \langle S, e_S \rangle & \text{otherwise,} \end{cases}$$

where e_S is an arbitrary but fixed element of $\mathcal{D}_S^{\mathfrak{M}}$. Otherwise, if S is non-protected, we set:

$$\hat{f}_{\mathbf{S}'}^{\mathfrak{J}}(\langle T_1, c_1 \rangle, \dots, \langle T_n, c_n \rangle) \triangleq \lambda T \in \mathcal{T}(\mathbb{F}_{\mathbf{T}}). \begin{cases} f_{\langle T_1, \dots, T_n, T \rangle}^{\mathfrak{M}}(c_1, \dots, c_n) & \text{if } \mathbf{S} \text{ matches } \langle T_1, \dots, T_n, T \rangle, \\ e_T & \text{otherwise,} \end{cases}$$

where e_T is an arbitrary but fixed element of $\mathcal{D}_T^{\mathfrak{M}}$. Similarly, let p be a predicate symbol in Γ of type signature $\mathbf{S} = \langle S_1, \dots, S_n \rangle$. The symbol \hat{p} has the type signature $\mathbf{S}' = \langle [S_1]^{-}, \dots, [S_n]^{-} \rangle$, and we set:

$$\hat{p}_{\mathbf{S}'}^{\mathfrak{J}}(\langle T_1, c_1 \rangle, \dots, \langle T_n, c_n \rangle) \triangleq \begin{cases} p_{\langle T_1, \dots, T_n \rangle}^{\mathfrak{M}}(c_1, \dots, c_n) & \text{if } \mathbf{S} \text{ matches } \langle T_1, \dots, T_n \rangle, \\ \top & \text{otherwise.} \end{cases}$$

The symbol deco is interpreted as follows:

$$\text{deco}_{\langle T, \mathbf{U}, \mathbf{D} \rangle}^{\mathfrak{J}}(T, h) \triangleq \langle T, h(T) \rangle$$

Finally, the type-representing function symbols are interpreted exactly as the original type constructors:

$$\mathbf{F}_{\langle T, \dots, T \rangle}^{\mathfrak{J}}(T_1, \dots, T_n) \triangleq \mathbf{F}(T_1, \dots, T_n)$$

Now, consider a variable valuation φ in \mathfrak{J} . For every type variable α , we set $\alpha\pi' \triangleq \varphi(v^\alpha : \mathbf{T})$. Every type in $\mathcal{D}_{\mathfrak{T}}^{\mathfrak{J}}$ is a non-protected sort, so π' is a type valuation. Moreover, for every non-protected type T , we have $\mathfrak{J}_\varphi(\text{DEC}(T)) = T\pi'$.

Then, for any variable $u : T$ in Γ , we set:

$$\varphi'(u : T) \triangleq \begin{cases} c & \text{if } T \text{ is protected and } \varphi(u_T : T) = \langle T, c \rangle, \\ \varphi(u_T : \mathbf{U})(T\pi') & \text{otherwise.} \end{cases}$$

We proceed to prove the two main equalities, and, once again, the proof goes by induction on the structure of terms and formulas. We consider four cases: a variable, a function application, an equality, and a universally quantified formula.

1. Let $u : T$ be a variable in Γ . If T is protected, then $\mathfrak{J}_\varphi(\text{DEC}(u : T)) = \mathfrak{J}_\varphi(u_T : T) = \varphi(u_T : T) = \langle T, \varphi'(u : T) \rangle = \langle T, \mathfrak{M}_{\pi', \varphi'}(u : T) \rangle$. Otherwise, if T is a non-protected type, we have:

$$\begin{aligned} \mathfrak{J}_\varphi(\text{DEC}(u : T)) &= \mathfrak{J}_\varphi(\text{deco}(\text{DEC}(T), u_T : \mathbf{U}) : \mathbf{D}) = \\ &\text{deco}_{\langle T, \mathbf{U}, \mathbf{D} \rangle}^{\mathfrak{J}}(\mathfrak{J}_\varphi(\text{DEC}(T)), \mathfrak{J}_\varphi(u_T : \mathbf{U})) = \langle T\pi', \varphi(u_T : \mathbf{U})(T\pi') \rangle = \\ &\langle T\pi', \varphi'(u : T) \rangle = \langle T\pi', \mathfrak{M}_{\pi', \varphi'}(u : T) \rangle \end{aligned}$$

2. Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$. Let $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$ be the type signature of f and $\mathbf{S}' = \langle [S_1]^- , \dots, [S_n]^- , [S]^+ \rangle$. If T is protected, then

$$\begin{aligned} \mathfrak{J}_\varphi(\text{DEC}(t)) &= \hat{f}_{\mathbf{S}'}^{\mathfrak{J}}(\mathfrak{J}_\varphi(\text{DEC}(t_1)), \dots, \mathfrak{J}_\varphi(\text{DEC}(t_n))) = \\ &\hat{f}_{\mathbf{S}'}^{\mathfrak{J}}(\langle T_1\pi', \mathfrak{M}_{\pi', \varphi'}(t_1) \rangle, \dots, \langle T_n\pi', \mathfrak{M}_{\pi', \varphi'}(t_n) \rangle) = \\ &\langle T, f_{\langle T_1\pi', \dots, T_n\pi', T \rangle}^{\mathfrak{M}}(\mathfrak{M}_{\pi', \varphi'}(t_1), \dots, \mathfrak{M}_{\pi', \varphi'}(t_n)) \rangle = \langle T\pi', \mathfrak{M}_{\pi', \varphi'}(t) \rangle \end{aligned}$$

Now, suppose that T is non-protected. We obtain:

$$\begin{aligned} \mathfrak{J}_\varphi(\text{DEC}(t)) &= \mathfrak{J}_\varphi(\text{deco}(\text{DEC}(T), \hat{f}(\text{DEC}(t_1), \dots, \text{DEC}(t_n)) : \mathbf{U}) : \mathbf{D}) = \\ &\text{deco}_{\langle T, \mathbf{U}, \mathbf{D} \rangle}^{\mathfrak{J}}(T\pi', \hat{f}_{\mathbf{S}'}^{\mathfrak{J}}(\mathfrak{J}_\varphi(\text{DEC}(t_1)), \dots, \mathfrak{J}_\varphi(\text{DEC}(t_n)))) = \\ &\langle T\pi', \hat{f}_{\mathbf{S}'}^{\mathfrak{J}}(\langle T_1\pi', \mathfrak{M}_{\pi', \varphi'}(t_1) \rangle, \dots, \langle T_n\pi', \mathfrak{M}_{\pi', \varphi'}(t_n) \rangle)(T\pi') \rangle = \\ &\langle T\pi', f_{\langle T_1\pi', \dots, T_n\pi', T\pi' \rangle}^{\mathfrak{M}}(\mathfrak{M}_{\pi', \varphi'}(t_1), \dots, \mathfrak{M}_{\pi', \varphi'}(t_n)) \rangle = \langle T\pi', \mathfrak{M}_{\pi', \varphi'}(t) \rangle \end{aligned}$$

3. Consider an equality formula $t_1 \approx t_2$. Let T be the type of t_1 and t_2 . By induction hypothesis, $\mathfrak{J}_\varphi(\text{DEC}(t_1)) = \langle T\pi', \mathfrak{M}_{\pi', \varphi'}(t_1) \rangle$ and $\mathfrak{J}_\varphi(\text{DEC}(t_2)) = \langle T\pi', \mathfrak{M}_{\pi', \varphi'}(t_2) \rangle$. Obviously, these two are equal if and only if $\mathfrak{M}_{\pi', \varphi'}(t_1)$ and $\mathfrak{M}_{\pi', \varphi'}(t_2)$ are equal.

4. Consider a universally quantified formula $\forall(u : S) F$, where S is protected. By construction, $\mathcal{D}_S^{\mathfrak{J}} = \{S\} \times \mathcal{D}_S^{\mathfrak{M}}$. Let us take an arbitrary element $\langle S, c \rangle \in \mathcal{D}_S^{\mathfrak{J}}$. According to the definition above, the variable valuation $\varphi[u_S : S \mapsto \langle S, c \rangle]$ will produce in \mathfrak{M} the same type valuation π' and the variable valuation $\varphi'[u : S \mapsto c]$. We obtain:

$$\begin{aligned} \mathfrak{J}_\varphi(\text{DEC}(\forall(u : S) F)) &= \bigwedge_{\langle S, c \rangle \in \mathcal{D}_S^{\mathfrak{J}}} \mathfrak{J}_{\varphi[u_S : S \mapsto \langle S, c \rangle]}(\text{DEC}(F)) \\ &= \bigwedge_{c \in \mathcal{D}_S^{\mathfrak{M}}} \mathfrak{M}_{\pi', \varphi'[u : S \mapsto c]}(F) = \mathfrak{M}_{\pi', \varphi'}(\forall(u : S) F) \end{aligned}$$

5. Consider a formula $\forall(u : T) F$, where T is a non-protected type. After the transformation, the variable symbol u_T acquires the type \mathbf{U} . Any element of $\mathcal{D}_{\mathbf{U}}^{\mathfrak{J}}$ is a function that maps any non-protected sort S to some element of $\mathcal{D}_S^{\mathfrak{M}}$. However, we are only interested in the value of this function on $T\pi'$, since for any two functions $h, h' \in \mathcal{D}_{\mathbf{U}}^{\mathfrak{J}}$ that coincide on $T\pi'$, $\mathfrak{J}_{\varphi[u_T : \mathbf{U} \mapsto h]}(\text{DEC}(F)) = \mathfrak{J}_{\varphi[u_T : \mathbf{U} \mapsto h']}(\text{DEC}(F))$. Indeed, $u_T : \mathbf{U}$ can only occur in $\text{DEC}(F)$ inside a term $\text{deco}(\text{DEC}(T), u_T : \mathbf{U}) : \mathbf{D}$. Since $\text{DEC}(F)$ does not contain quantifiers over \mathbf{T} , the term $\text{DEC}(T)$ will always evaluate to $T\pi'$. Therefore, the term $\text{deco}(\text{DEC}(T), u_T : \mathbf{U}) : \mathbf{D}$ will evaluate to $\langle T\pi', h(T\pi') \rangle$ in both cases. For any $c \in \mathcal{D}_{T\pi'}^{\mathfrak{M}}$, let h_c denote an arbitrary but fixed function in $\mathcal{D}_{\mathbf{U}}^{\mathfrak{J}}$ that maps $T\pi'$ to c . The valuation $\varphi[u_T : \mathbf{U} \mapsto h_c]$ produces in \mathfrak{M} the same type valuation π' and the variable valuation $\varphi'[u : T \mapsto c]$. We obtain:

$$\begin{aligned} \mathfrak{J}_\varphi(\text{DEC}(\forall(u : T) F)) &= \bigwedge_{h \in \mathcal{D}_{\mathbf{U}}^{\mathfrak{J}}} \mathfrak{J}_{\varphi[u_T : \mathbf{U} \mapsto h]}(\text{DEC}(F)) \\ &= \bigwedge_{c \in \mathcal{D}_{T\pi'}^{\mathfrak{M}}} \mathfrak{J}_{\varphi[u_T : \mathbf{U} \mapsto h_c]}(\text{DEC}(F)) \\ &= \bigwedge_{c \in \mathcal{D}_{T\pi'}^{\mathfrak{M}}} \mathfrak{M}_{\pi', \varphi'[u : T \mapsto c]}(F) = \mathfrak{M}_{\pi', \varphi'}(\forall(u : T) F) \end{aligned}$$

Thus, for every formula F in Γ , $\text{DEC}^\circ(F)$ is satisfied under \mathfrak{J} . \square

Theorem 13 (Completeness of DEC). *For every set of closed formulas with protected sorts Γ , if $\text{DEC}^\circ(\Gamma)$ is satisfiable then Γ is satisfiable.*

Proof. Let \mathfrak{M} be a model of $\text{DEC}^\circ(\Gamma)$. We are going to construct an interpretation \mathfrak{J} of Γ such that, for every type valuation π and every variable valuation φ in \mathfrak{J} , there exists a variable valuation φ' in the original model \mathfrak{M} such that:

$$\mathfrak{J}_{\pi, \varphi}(t) = \mathfrak{M}_{\varphi'}(\text{DEC}(t)) \quad \mathfrak{J}_{\pi, \varphi}(F) = \mathfrak{M}_{\varphi'}(\text{DEC}(F))$$

for every term t and formula F occurring in Γ . Since $\text{DEC}^\circ(F)$ is the universal closure of $\text{DEC}(F)$, the second equality suffices to show $\mathfrak{J}_\pi(F)$ for every $F \in \Gamma$.

For any sort S in the signature of Γ , we set:

$$\mathcal{D}_S^{\mathfrak{J}} \triangleq \begin{cases} \mathcal{D}_S^{\mathfrak{M}} & \text{if } S \text{ is protected,} \\ \text{deco}_{(\mathfrak{T}, \mathfrak{U}, \mathfrak{D})}^{\mathfrak{M}}(\mathfrak{M}(\text{DEC}(S)), \mathcal{D}_{\mathfrak{U}}^{\mathfrak{M}}) & \text{otherwise,} \end{cases}$$

In the second case, we speak about the image of $\text{deco}^{\mathfrak{M}}$ on the domain of \mathfrak{U} . We omit the variable valuation when evaluating $\text{DEC}(S)$, as this term does not contain variables. Notice that for every non-protected sort S , $\mathcal{D}_S^{\mathfrak{J}} \subseteq \mathcal{D}_{\mathfrak{D}}^{\mathfrak{M}}$.

Consider a function symbol f in Γ of type signature $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$. Let \mathbf{S}' stand for the type signature of \hat{f} , $\langle [S_1]^{-}, \dots, [S_n]^{-}, [S]^+ \rangle$. The symbol f is then interpreted in \mathfrak{J} on a sort vector $\mathbf{S}\tau$ as follows:

$$f_{\mathbf{S}\tau}^{\mathfrak{J}}(\mathbf{c}) \triangleq \begin{cases} \hat{f}_{\mathbf{S}'}^{\mathfrak{M}}(\mathbf{c}) & \text{if } S \text{ is protected,} \\ \text{deco}_{(\mathfrak{T}, \mathfrak{U}, \mathfrak{D})}^{\mathfrak{M}}(\mathfrak{M}(\text{DEC}(S\tau)), \hat{f}_{\mathbf{S}'}^{\mathfrak{M}}(\mathbf{c})) & \text{otherwise.} \end{cases}$$

Notice that every argument c_i from \mathbf{c} belongs either to $\mathcal{D}_{S_i}^{\mathfrak{M}}$ (if the sort S_i is protected) or to $\mathcal{D}_{\mathfrak{D}}^{\mathfrak{M}}$ (if it is not), and thus the interpretation is well-defined. Furthermore, every predicate symbol p of type signature $\mathbf{S} = \langle S_1, \dots, S_n \rangle$ is interpreted in \mathfrak{J} on $\mathbf{S}\tau$ as $\hat{p}_{([S_1]^{-}, \dots, [S_n]^{-})}^{\mathfrak{M}}$ on the same arguments. This concludes the definition of the interpretation \mathfrak{J} .

Now, consider some type valuation π and a variable valuation φ in \mathfrak{J} . For every variable occurring in $\text{DEC}^{\circ}(\Gamma)$, we define (assuming that S is a protected sort and T is a non-protected type):

$$\begin{aligned} \varphi'(v^{\alpha} : \mathfrak{T}) &\triangleq \mathfrak{M}(\text{DEC}(\alpha\pi)) \\ \varphi'(u_S : S) &\triangleq \varphi(u : S) \\ \varphi'(u_T : \mathfrak{U}) &\triangleq b_{\varphi(u : T)} \end{aligned}$$

where, for any $a \in \mathcal{D}_{T\pi}^{\mathfrak{J}}$, b_a is an arbitrary but fixed element of $\mathcal{D}_{\mathfrak{U}}^{\mathfrak{M}}$ such that

$$\text{deco}_{(\mathfrak{T}, \mathfrak{U}, \mathfrak{D})}^{\mathfrak{M}}(\mathfrak{M}(\text{DEC}(T\pi)), b_a) = a$$

By definition of $\mathcal{D}_{T\pi}^{\mathfrak{J}}$, there exists at least one such element. Notice that as T is non-protected, $T\pi$ is non-protected, too. Also, $\mathfrak{M}_{\varphi'}(\text{DEC}(T)) = \mathfrak{M}(\text{DEC}(T\pi))$.

The proof of the equalities above goes by induction on the structure of terms and formulas. We will consider three cases: a variable, a function application, and a universally quantified formula.

1. Let $u : T$ be a variable occurring in Γ . If T is a protected sort, we have $\mathfrak{J}_{\pi, \varphi}(u : T) = \varphi(u : T) = \varphi'(u_T : T) = \mathfrak{M}_{\varphi'}(\text{DEC}(u : T))$. If T is a non-protected type, we obtain:

$$\begin{aligned} \mathfrak{J}_{\pi, \varphi}(u : T) &= \varphi(u : T) = \text{deco}_{(\mathfrak{T}, \mathfrak{U}, \mathfrak{D})}^{\mathfrak{M}}(\mathfrak{M}(\text{DEC}(T\pi)), b_{\varphi(u : T)}) \\ &= \text{deco}_{(\mathfrak{T}, \mathfrak{U}, \mathfrak{D})}^{\mathfrak{M}}(\mathfrak{M}_{\varphi'}(\text{DEC}(T)), \varphi'(u_T : \mathfrak{U})) \\ &= \mathfrak{M}_{\varphi'}(\text{deco}(\text{DEC}(T), u_T : \mathfrak{U}) : \mathfrak{D}) = \mathfrak{M}_{\varphi'}(\text{DEC}(u : T)) \end{aligned}$$

- Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$ in Γ . Let $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$ be the type signature of f and \mathbf{S}' be the type signature of \hat{f} . If S is a protected sort, then f is interpreted on $\langle T_1, \dots, T_n, T \rangle \pi$ in \mathfrak{J} exactly as \hat{f} on \mathbf{S}' in \mathfrak{M} , thus $\mathfrak{J}_{\pi, \varphi}(t) = \mathfrak{M}_{\varphi'}(\text{DEC}(t))$. Assume that S is non-protected. Denoting the list of arguments with \mathbf{t} , we get:

$$\begin{aligned} \mathfrak{J}_{\pi, \varphi}(t) &= f_{\mathbf{T}\pi}^{\mathfrak{J}}(\mathfrak{J}_{\pi, \varphi}(\mathbf{t})) = \text{deco}_{\langle \mathbf{T}, \mathbf{U}, \mathbf{D} \rangle}^{\mathfrak{M}}(\mathfrak{M}(\text{DEC}(T\pi)), \hat{f}_{\mathbf{S}'}^{\mathfrak{M}}(\mathfrak{J}_{\pi, \varphi}(\mathbf{t}))) = \\ &= \text{deco}_{\langle \mathbf{T}, \mathbf{U}, \mathbf{D} \rangle}^{\mathfrak{M}}(\mathfrak{M}_{\varphi'}(\text{DEC}(T)), \hat{f}_{\mathbf{S}'}^{\mathfrak{M}}(\mathfrak{M}_{\varphi'}(\text{DEC}(\mathbf{t})))) = \mathfrak{M}_{\varphi'}(\text{DEC}(t)) \end{aligned}$$

- Consider a universal formula $\forall(u : T) F$. The case of protected T is trivial, so we assume that T is non-protected. Let b and b' be two elements of $\mathcal{D}_{\mathbf{U}}^{\mathfrak{M}}$ such that $\text{deco}_{\langle \mathbf{T}, \mathbf{U}, \mathbf{D} \rangle}^{\mathfrak{M}}(\mathfrak{M}(\text{DEC}(T\pi)), b) = \text{deco}_{\langle \mathbf{T}, \mathbf{U}, \mathbf{D} \rangle}^{\mathfrak{M}}(\mathfrak{M}(\text{DEC}(T\pi)), b')$. We can show that $\mathfrak{M}_{\varphi'[u_T : \mathbf{U} \mapsto b]}(\text{DEC}(F)) = \mathfrak{M}_{\varphi'[u_T : \mathbf{U} \mapsto b']}(\text{DEC}(F))$. Indeed, the variable $u_T : \mathbf{U}$ can only occur in $\text{DEC}(F)$ in terms $\text{deco}(\text{DEC}(T), u_T : \mathbf{U}) : \mathbf{D}$. Since $\text{DEC}(F)$ does not contain quantifiers over \mathbf{T} , the term $\text{DEC}(T)$ will always evaluate to $\mathfrak{M}(\text{DEC}(T\pi))$. Let a be an arbitrary element of $\mathcal{D}_{T\pi}^{\mathfrak{J}}$. The type valuation π and the variable valuation $\varphi[u : T \mapsto a]$ produce in \mathfrak{M} the variable valuation $\varphi'[u_T : \mathbf{U} \mapsto b_a]$.

$$\begin{aligned} \mathfrak{J}_{\pi, \varphi}(\forall(u : T) F) &= \bigwedge_{a \in \mathcal{D}_{T\pi}^{\mathfrak{J}}} \mathfrak{J}_{\pi, \varphi[u : T \mapsto a]}(F) \\ &= \bigwedge_{a \in \mathcal{D}_{T\pi}^{\mathfrak{J}}} \mathfrak{M}_{\varphi'[u_T : \mathbf{U} \mapsto b_a]}(\text{DEC}(F)) \\ &= \bigwedge_{b \in \mathcal{D}_{\mathbf{U}}^{\mathfrak{M}}} \mathfrak{M}_{\varphi'[u_T : \mathbf{U} \mapsto b]}(\text{DEC}(F)) \\ &= \mathfrak{M}_{\varphi'}(\forall(u_T : \mathbf{U}) \text{DEC}(F)) = \mathfrak{M}_{\varphi'}(\text{DEC}(\forall(u : T) F)) \end{aligned}$$

Thus, for every formula F in Γ , F is satisfied under \mathfrak{J} . \square

3.5 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:

- The set of type constructors is extended with \mathbf{U} and \mathbf{T} .
- 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 $\langle \mathbb{T}, \dots, \mathbb{T}, [S_1], \dots, [S_n] \rangle$.
4. For every variable symbol u and type T , we add a new variable symbol u_T .
5. For every type variable $\alpha \in \mathbb{V}_{\mathbb{T}}$, we add a new variable symbol v^α .
6. For every type constructor $F \in \mathbb{F}_{\mathbb{T}}$, we add a new function symbol F of the same arity and with type signature $\langle \mathbb{T}, \dots, \mathbb{T}, \mathbb{T} \rangle$.

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

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

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 $U = \{\mathbb{1}\}$, the transformations TW and EXP° produce the following formula:

$$\forall(v^\alpha : \mathbb{T}) \forall(m : \mathbb{U}) \forall(c : \mathbb{U}) \text{from}_1(\text{get}(v^\alpha, \mathbb{I}, \text{set}(v^\alpha, \mathbb{I}, m, c, \text{to}_1(6)), c)) * 7 \approx 42$$

and PAR and EXP° give:

$$\begin{aligned} \forall(v^\alpha : \mathbb{T}) \forall(m : \mathbb{U}) \forall(c : \mathbb{U}) \text{get}_{[1/\beta]}(v^\alpha, \text{set}_{[1/\beta]}(v^\alpha, m, c, 6), c) * 7 \approx 42 \\ \forall(m : \mathbb{U}) \forall(c : \bar{\mathbb{1}}) \text{get}_{[1/\alpha, 1/\beta]}(\text{set}_{[1/\alpha, 1/\beta]}(m, c, 6), c) * 7 \approx 42 \end{aligned}$$

Theorem 14. *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. Moreover,*

$$\begin{aligned} \text{FV}(\text{EXP}(t)) &= \{u_T : [T] \mid u : T \in \text{FV}(t)\} \cup \{v^\alpha : \mathbb{T} \mid \alpha \in \text{FV}_{\mathbb{T}}(t)\} \\ \text{FV}(\text{EXP}(F)) &= \{u_T : [T] \mid u : T \in \text{FV}(F)\} \cup \{v^\alpha : \mathbb{T} \mid \alpha \in \text{FV}_{\mathbb{T}}(F)\} \end{aligned}$$

Finally, for every closed formula F , $\text{EXP}^\circ(F)$ is a well-formed closed monomorphic formula.

Proof. The last claim easily follows from the second one. The proof of the first two statements goes by induction on the structure of terms and formulas. The only interesting case is the application of a function symbol (predicate symbols are treated in the same way).

Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$. Let S_i denote the i -th sort in the type signature of f . We have two possible cases: either S_i and T_i are protected and $S_i = T_i$, or both S_i and T_i are non-protected, by the definition of matching with protected types. Therefore, $[T_i] = [S_i]$ for every $i \in [1, n]$.

Now, let S be the type of the value in the type signature of f . By the same argument, either S and T are protected and $S = T$, or both S and T are non-protected. Therefore, $[T] = [S]$ and the application is well-formed.

Finally, every type variable α occurring in $\langle T_1, \dots, T_n, T \rangle$ appears as a free variable $v^\alpha : T$ in the supplementary “type” arguments of \hat{f} in $\text{EXP}(t)$. \square

Theorem 15 (Soundness of EXP). *If a set of closed formulas with protected sorts Γ is satisfiable so that every non-protected sort has an infinite domain in the model, then $\text{EXP}^\circ(\Gamma)$ is satisfiable.*

Proof. Every model of Γ is a model of the set of monomorphic type instances $\text{MI}(\Gamma)$ and vice versa. In this set, protected sorts require no special treatment and the whole problem is reduced to many-sorted logic. Let \mathfrak{M} be a model of Γ such that, for every non-protected sort S , the domain $\mathcal{D}_S^{\mathfrak{M}}$ is infinite. By Löwenheim-Skolem theorem for many-sorted logic [11, Sect. 7.6], we can assume that the domain of every non-protected sort is simply \mathbb{N} .

We are going to construct an interpretation \mathfrak{J} of $\text{EXP}^\circ(\Gamma)$ such that, for every variable valuation φ in \mathfrak{J} , there exist a type valuation π' and a variable valuation φ' in the original model \mathfrak{M} , such that:

$$\mathfrak{J}_\varphi(\text{EXP}(t)) = \mathfrak{M}_{\pi', \varphi'}(t) \quad \mathfrak{J}_\varphi(\text{EXP}(F)) = \mathfrak{M}_{\pi', \varphi'}(F)$$

for every term t and formula F occurring in Γ . We do not need to consider type valuations in \mathfrak{J} , since the result of EXP° is entirely monomorphic. As soon as these two equalities are proved, it is easy to show that for every closed formula $H \in \Gamma$, the interpretation \mathfrak{J} satisfies $\text{EXP}^\circ(H)$. Indeed,

$$\mathfrak{J}(\text{EXP}^\circ(H)) = \bigwedge_{\varphi} \mathfrak{J}_\varphi(\text{EXP}(H)) = \bigwedge_{\varphi} \mathfrak{M}_{\pi', \varphi'}(H) = \top$$

For every sort S in the initial signature, protected or not, we keep its domain: $\mathcal{D}_S^{\mathfrak{J}} \triangleq \mathcal{D}_S^{\mathfrak{M}}$. The domain of \mathbf{U} is the set of natural numbers: $\mathcal{D}_{\mathbf{U}}^{\mathfrak{J}} \triangleq \mathbb{N}$. The domain of \mathbf{T} is the set of non-protected sorts: $\mathcal{D}_{\mathbf{T}}^{\mathfrak{J}} \triangleq \mathcal{T}(\mathbb{F}_{\mathbf{T}})$.

Then we give an interpretation to the function and predicate symbols in $\text{EXP}^\circ(\Gamma)$. Let f be a function symbol in Γ of type signature $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$ and $\alpha_1, \dots, \alpha_r$ be the type variables in \mathbf{S} . The symbol \hat{f} has the arity $r + n$ and the type signature $\mathbf{S}' = \langle \mathbf{T}, \dots, \mathbf{T}, [S_1], \dots, [S_n], [S] \rangle$ and this is the only sort vector to interpret \hat{f} on, since it is monomorphic. We define:

$$\hat{f}_{\mathbf{S}'}^{\mathfrak{J}}(T_1, \dots, T_r, c_1, \dots, c_n) \triangleq f_{\mathbf{S}'}^{\mathfrak{M}}(c_1, \dots, c_n)$$

where τ is the type substitution $[T_1/\alpha_1, \dots, T_r/\alpha_r]$. Notice that every type S in \mathbf{S} verifies $\mathcal{D}_{[S]}^{\mathfrak{J}} = \mathcal{D}_{S\tau}^{\mathfrak{M}}$. Indeed, either S is protected and $[S] = S = S\tau$, or both S and $S\tau$ are non-protected and $\mathcal{D}_{[S]}^{\mathfrak{J}} = \mathcal{D}_{\mathbb{U}}^{\mathfrak{J}} = \mathbb{N} = \mathcal{D}_{S\tau}^{\mathfrak{M}}$.

Similarly, let p be a predicate symbol in Γ of type signature $\mathbf{S} = \langle S_1, \dots, S_n \rangle$ and $\alpha_1, \dots, \alpha_r$ be the type variables in \mathbf{S} . The symbol \hat{p} has the arity $r + n$ and the type signature $\mathbf{S}' = \langle \mathbb{T}, \dots, \mathbb{T}, [S_1], \dots, [S_n] \rangle$. Then we set:

$$\hat{p}_{\mathbf{S}'}^{\mathfrak{J}}(T_1, \dots, T_r, c_1, \dots, c_n) \triangleq p_{\mathbf{S}[T_1/\alpha_1, \dots, T_r/\alpha_r]}^{\mathfrak{M}}(c_1, \dots, c_n)$$

Finally, the type-representing function symbols are interpreted exactly as the original type constructors:

$$\mathbf{F}_{\langle \mathbb{T}, \dots, \mathbb{T} \rangle}^{\mathfrak{J}}(T_1, \dots, T_n) \triangleq \mathbf{F}(T_1, \dots, T_n)$$

Now, let us consider a variable valuation φ in \mathfrak{J} . For any type variable α , we set $\alpha\pi' \triangleq \varphi(v^\alpha : \mathbb{T})$. For any variable $u : T$ in Γ , we set $\varphi'(u : T) \triangleq \varphi(u_T : [T])$. As we have shown above, $\mathcal{D}_{T\pi'}^{\mathfrak{M}} = \mathcal{D}_{[T]}^{\mathfrak{J}}$. Also, for every non-protected type T , we have $\mathfrak{J}_\varphi(\text{EXP}(T)) = T\pi'$.

The proof of the two main equalities goes by induction on the structure of terms and formulas. We consider three cases: a variable, a function application, and a universally quantified formula.

1. Let $u : T$ be a variable in Γ . We obtain:

$$\mathfrak{J}_\varphi(\text{EXP}(u : T)) = \mathfrak{J}_\varphi(u_T : [T]) = \varphi(u_T : [T]) = \varphi'(u : T) = \mathfrak{M}_{\pi', \varphi'}(u : T)$$

2. Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$. Let $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$ be the type signature of f and $\alpha_1, \dots, \alpha_r$ be the type variables of \mathbf{S} . Let \mathbf{S}' be the type signature of \hat{f} . Finally, let τ be the type substitution that instantiates \mathbf{S} to $\langle T_1, \dots, T_n, T \rangle$. We obtain:

$$\begin{aligned} \mathfrak{J}_\varphi(\text{EXP}(t)) &= \mathfrak{J}_\varphi(\hat{f}(\text{EXP}(\alpha_1\tau), \dots, \text{EXP}(\alpha_r\tau), \text{EXP}(t_1), \dots, \text{EXP}(t_n)) : [S]) \\ &= \hat{f}_{\mathbf{S}'}^{\mathfrak{J}}(\alpha_1\tau\pi', \dots, \alpha_r\tau\pi', \mathfrak{M}_{\pi', \varphi'}(t_1), \dots, \mathfrak{M}_{\pi', \varphi'}(t_n)) \\ &= f_{\mathbf{S}\tau\pi'}^{\mathfrak{M}}(\mathfrak{M}_{\pi', \varphi'}(t_1), \dots, \mathfrak{M}_{\pi', \varphi'}(t_n)) = \mathfrak{M}_{\pi', \varphi'}(t) \end{aligned}$$

3. Consider a universally quantified formula $\forall(u : T) F$. Let us take an arbitrary element $c \in \mathcal{D}_{[T]}^{\mathfrak{J}} = \mathcal{D}_{T\pi'}^{\mathfrak{M}}$. The variable valuation $\varphi[u_T : [T] \mapsto c]$ produces in \mathfrak{M} the same type valuation π' and the variable valuation $\varphi'[u : T \mapsto c]$.

$$\begin{aligned} \mathfrak{J}_\varphi(\text{EXP}(\forall(u : T) F)) &= \bigwedge_{c \in \mathcal{D}_{[T]}^{\mathfrak{J}}} \mathfrak{J}_{\varphi[u_T : [T] \mapsto c]}(\text{EXP}(F)) \\ &= \bigwedge_{c \in \mathcal{D}_{T\pi'}^{\mathfrak{M}}} \mathfrak{M}_{\pi', \varphi'[u : T \mapsto c]}(F) = \mathfrak{M}_{\pi', \varphi'}(\forall(u : T) F) \end{aligned}$$

Thus, for every formula F in Γ , $\text{EXP}^\circ(F)$ is satisfied under \mathfrak{J} . \square

Theorem 16 (Completeness of EXP). *For every set of closed formulas with protected sorts Γ , if $\text{EXP}^\circ(\Gamma)$ is satisfiable then Γ is satisfiable.*

Proof. Let \mathfrak{M} be a model of $\text{EXP}^\circ(\Gamma)$. We are going to construct an interpretation \mathfrak{J} of Γ such that, for every type valuation π and every variable valuation φ in \mathfrak{J} , there exists a variable valuation φ' in the original model \mathfrak{M} , such that:

$$\mathfrak{J}_{\pi,\varphi}(t) = \mathfrak{M}_{\varphi'}(\text{EXP}(t)) \quad \mathfrak{J}_{\pi,\varphi}(F) = \mathfrak{M}_{\varphi'}(\text{EXP}(F))$$

for every term t and formula F occurring in Γ . Since $\text{EXP}^\circ(F)$ is the universal closure of $\text{EXP}(F)$, the second equality suffices to show $\mathfrak{J}_\pi(F)$ for every $F \in \Gamma$.

For any sort S in the signature of Γ , we set $\mathcal{D}_S^{\mathfrak{J}} \triangleq \mathcal{D}_{[S]}^{\mathfrak{M}}$.

Consider a function symbol f in Γ of type signature $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$ and let $\alpha_1, \dots, \alpha_r$ be the type variables of \mathbf{S} . Let \mathbf{S}' stand for the type signature of \hat{f} , $\langle \mathbb{T}, \dots, \mathbb{T}, [S_1], \dots, [S_n], [S] \rangle$. The symbol f is then interpreted in \mathfrak{J} on a sort vector $\mathbf{S}\tau$ as follows:

$$f_{\mathbf{S}\tau}^{\mathfrak{J}}(\mathbf{c}) \triangleq \hat{f}_{\mathbf{S}'\tau}^{\mathfrak{M}}(\mathfrak{M}(\text{EXP}(\alpha_1\tau)), \dots, \mathfrak{M}(\text{EXP}(\alpha_r\tau)), \mathbf{c}).$$

We omit the variable valuation when evaluating $\text{EXP}(\alpha_i\tau)$, as these terms do not contain variables. Notice that every argument c_i from \mathbf{c} belongs either to $\mathcal{D}_{S_i}^{\mathfrak{M}}$ (if the sort S_i is protected) or to $\mathcal{D}_{\mathbb{U}}^{\mathfrak{M}}$ (if it is not), and thus the interpretation is well-defined.

Similarly, let p be a predicate symbol in Γ of type signature $\mathbf{S} = \langle S_1, \dots, S_n \rangle$. Let $\alpha_1, \dots, \alpha_r$ be the type variables of \mathbf{S} and \mathbf{S}' be the type signature of \hat{p} . The symbol p is then interpreted in \mathfrak{J} on a sort vector $\mathbf{S}\tau$ as follows:

$$p_{\mathbf{S}\tau}^{\mathfrak{J}}(\mathbf{c}) \triangleq \hat{p}_{\mathbf{S}'\tau}^{\mathfrak{M}}(\mathfrak{M}(\text{EXP}(\alpha_1\tau)), \dots, \mathfrak{M}(\text{EXP}(\alpha_r\tau)), \mathbf{c}).$$

This concludes the definition of interpretation \mathfrak{J} .

Now, let us consider some type valuation π and a variable valuation φ in \mathfrak{J} . For every variable occurring in $\text{EXP}^\circ(\Gamma)$, we define:

$$\varphi'(v^\alpha : \mathbb{T}) \triangleq \mathfrak{M}(\text{EXP}(\alpha\pi)) \quad \varphi'(u_T : [T]) \triangleq \varphi(u : T)$$

Notice that $\mathfrak{M}_{\varphi'}(\text{EXP}(T)) = \mathfrak{M}(\text{EXP}(T\pi))$ and $\mathcal{D}_{T\pi}^{\mathfrak{J}} = \mathcal{D}_{[T]}^{\mathfrak{M}}$.

The proof of the equalities above goes by induction on the structure of terms and formulas. We consider three cases: a variable, a function application, and a universally quantified formula.

1. Let $u : T$ be a variable occurring in Γ . We obtain:

$$\mathfrak{J}_{\pi,\varphi}(u : T) = \varphi(u : T) = \varphi'(u_T : [T]) = \mathfrak{M}_{\varphi'}(u_T : [T]) = \mathfrak{M}_{\varphi'}(\text{EXP}(u : T))$$

2. Consider a term $t = f(t_1 : T_1, \dots, t_n : T_n) : T$ in Γ . Let $\mathbf{S} = \langle S_1, \dots, S_n, S \rangle$ be the type signature of f , let $\alpha_1, \dots, \alpha_r$ be the type variables of \mathbf{S} , let \mathbf{S}'

be the type signature of \hat{f} , and τ the type substitution that instantiates \mathbf{S} to $\langle T_1, \dots, T_n, T \rangle$. Denoting the list of arguments with \mathbf{t} , we get:

$$\begin{aligned} \mathfrak{J}_{\pi, \varphi}(t) &= f_{\mathbf{S}\tau\pi}^{\mathfrak{J}}(\mathfrak{J}_{\pi, \varphi}(\mathbf{t})) \\ &= \hat{f}_{\mathbf{S}'}^{\mathfrak{M}}(\mathfrak{M}(\text{EXP}(\alpha_1\tau\pi)), \dots, \mathfrak{M}(\text{EXP}(\alpha_r\tau\pi)), \mathfrak{J}_{\pi, \varphi}(\mathbf{t})) \\ &= \hat{f}_{\mathbf{S}'}^{\mathfrak{M}}(\mathfrak{M}_{\varphi'}(\text{EXP}(\alpha_1\tau)), \dots, \mathfrak{M}_{\varphi'}(\text{EXP}(\alpha_r\tau)), \mathfrak{M}_{\varphi'}(\text{EXP}(\mathbf{t}))) \\ &= \mathfrak{M}_{\varphi'}(\hat{f}(\text{EXP}(\alpha_1\tau), \dots, \text{EXP}(\alpha_r\tau), \text{EXP}(\mathbf{t})) : [S]) = \mathfrak{M}_{\varphi'}(\text{EXP}(t)) \end{aligned}$$

3. Consider a universal formula $\forall(u:T) F$. Let a be some element of $\mathcal{D}_{T\pi}^{\mathfrak{J}}$. The type valuation π and the variable valuation $\varphi[u:T \mapsto a]$ produces in \mathfrak{M} the variable valuation $\varphi'[u_T:[T] \mapsto a]$. We obtain:

$$\begin{aligned} \mathfrak{J}_{\pi, \varphi}(\forall(u:T) F) &= \bigwedge_{a \in \mathcal{D}_{T\pi}^{\mathfrak{J}}} \mathfrak{J}_{\pi, \varphi[u:T \mapsto a]}(F) \\ &= \bigwedge_{a \in \mathcal{D}_{[T]}^{\mathfrak{M}}} \mathfrak{M}_{\varphi'[u_T:[T] \mapsto a]}(\text{EXP}(F)) \\ &= \mathfrak{M}_{\varphi'}(\forall(u_T:[T]) \text{EXP}(F)) = \mathfrak{M}_{\varphi'}(\text{EXP}(\forall(u:T) F)) \end{aligned}$$

Thus, for every formula F in Γ , F is satisfied under \mathfrak{J} . \square

From a practical point of view, Theorem 15 is not comforting. Given a **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 closed instances admit infinite models. For example, given the standard algebraic definitions of booleans (\mathbf{B}), lists ($\mathbf{L}(\alpha)$), and pairs ($\mathbf{P}(\alpha, \beta)$), we can conclude that the sorts \mathbf{B} and $\mathbf{P}(\mathbf{B}, \mathbf{B})$ are finite and \mathbf{L} , $\mathbf{L}(\mathbf{B})$, and $\mathbf{P}(\mathbf{L}, \mathbf{B})$ 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 [12, 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, we can prove that there exist only two (projected) booleans. This 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 $\forall(x:\mathbf{A}) \text{isUnit}(x)$ imply that the sort \mathbf{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 purely monomorphic (many-sorted) setting, Theorem 15 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 [16] or SPASS [18], we can first translate a proof task to many-sorted language, eliminate the protected finite sorts (if any) using projections³, and then apply EXP assuming all types are non-protected.

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 omit the PAR transformation (see the remark in the end of Section 3.2) and we add the classical type encoding technique with per-variable “type guards” [11]. Our implementation of this method (denoted GRD below) closely follows the description given in [10, Sect. 3.0].

We tested the encodings TW+DEC, TW+EXP, TW+GRD both with and without DIS on 4123 verification conditions⁴ generated by the Why platform from 166 programs, which originate from Caduceus [7], Jessie [14], or directly from Why. Translated tasks are sent to Z3, CVC3, and Yices with a time limit of 60 seconds. We use the built-in theory of arrays of Z3 and CVC3.

The choice of monomorphic specializations for DIS and of sorts to protect for TW is, certainly, of critical importance. After a series of tests, we decided in favor of the following settings. In DIS, the set W contains every monomorphic specialization that occurs in the goal formula. For Z3 and CVC3, we add the specializations of access and update operations on every closed array type that occurs in the goal. In TW, the set U contains the sorts of integers and reals. In DIS+TW, we add all sorts that occur in the goal. The results of our experiments are given in Table 1.

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 to prove 87 goals that were not proved by TW+DEC. On the other hand, with TW+DEC, CVC3 proves 20 goals that were not proved with DIS+TW+DEC.

We observe that the addition of the DIS phase allows to prove more goals in every case except for CVC3 with TW+EXP. In particular, the GRD transformation is remarkably helped by symbol discrimination. Notice that the results of Yices also improve with DIS, even though Yices does not have built-in theories for complex sorts like arrays. Just as the PAR transformation, DIS can also lead to an exponential growth of the problem. In practice, this growth is more limited and better controlled than in PAR. Still, it is to this introduction of mostly irrelevant type instances that we impute the loss of some proofs (up to 59 in the case of Z3 with TW+DEC). We believe that better heuristics to choose the sets W and U can be devised.

³ In absence of polymorphism, this is a sound and complete transformation.

⁴ Available in the source repository of Why3 [4] in the directory `bench/encoding`.

Z3	TW+GRD	TW+EXP	TW+DEC	DIS+TW+GRD	DIS+TW+EXP
DIS+TW+DEC	+272 -49	+30 -30	+89 -59	+37 -10	+41 -46
DIS+TW+EXP	+256 -28	+33 -28	+84 -49	+52 -20	
DIS+TW+GRD	+250 -54	+11 -38	+72 -69		
TW+DEC	+204 -11	+58 -88			
TW+EXP	+275 -52				

CVC3	TW+GRD	TW+EXP	TW+DEC	DIS+TW+GRD	DIS+TW+EXP
DIS+TW+DEC	+268 -21	+2 -26	+87 -20	+66 -4	+1 -7
DIS+TW+EXP	+271 -18	+3 -21	+92 -19	+69 -1	
DIS+TW+GRD	+202 -17	+3 -89	+49 -44		
TW+DEC	+185 -5	+1 -92			
TW+EXP	+272 -1				

Yices	TW+GRD	TW+EXP	TW+DEC	DIS+TW+GRD	DIS+TW+EXP
DIS+TW+DEC	+700 -4	+12 -284	+317 -38	+283 -0	+1 -285
DIS+TW+EXP	+980 -0	+26 -14	+563 -0	+567 -0	
DIS+TW+GRD	+441 -28	+3 -558	+141 -145		
TW+DEC	+429 -12	+0 -551			
TW+EXP	+968 -0				

Table 1. Experimental comparison of polymorphism encoding methods

The comparison between EXP, DEC, and GRD shows that EXP is generally more efficient than DEC which in its turn is much more efficient than GRD. This is quite different from the experimental results given in [10], 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 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.

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., Filliâtre, J.C., Marché, C., Paskevich, A.: Why3, a theorem proving framework (2010), <http://why3.gforge.inria.fr/>
5. Couchot, J.F., Lescuyer, S.: Handling polymorphism in automated deduction. In: CADE-21. LNAI, vol. 4603, pp. 263–278. Springer (2007)
6. Dutertre, B., de Moura, L.: The YICES SMT solver. Tech. rep., SRI International (2006)
7. Filliâtre, J.C., Marché, C.: Multi-prover verification of C programs. In: ICFEM'04. LNCS, vol. 3308, pp. 15–29. Springer (2004)
8. Hurd, J.: An LCF-style interface between HOL and first-order logic. In: CADE-18. LNAI, vol. 2392, pp. 134–138. Springer (2002)
9. 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)
10. 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)
11. Manzano, M.: Extensions of First-Order Logic, Cambridge Tracts in Theoretical Computer Science, vol. 19. Cambridge University Press (1996)
12. Meng, J., Paulson, L.C.: Translating higher-order clauses to first-order clauses. *Journal of Automated Reasoning* 40(1), 35–60 (2008)
13. de Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: TACAS'08. LNCS, vol. 4963, pp. 337–340. Springer (2008)
14. Moy, Y., Marché, C.: Jessie Plugin Tutorial, Lithium version. INRIA (2008), <http://www.frama-c.cea.fr/jessie.html>
15. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL — A Proof Assistant for Higher-Order Logic, LNCS, vol. 2283. Springer (2002)
16. Riazanov, A., Voronkov, A.: Vampire 1.1. In: IJCAR'01. LNCS, vol. 2083, pp. 376–380. Springer (2001)
17. The Coq Development Team: The Coq Proof Assistant Reference Manual – Version V8.0 (2004), <http://coq.inria.fr>
18. Weidenbach, C., Dimova, D., Fietzke, A., Kumar, R., Suda, M., Wischniewski, P.: SPASS version 3.5. In: CADE-22. LNCS, vol. 5663, pp. 140–145. Springer (2009)