



This is a repository copy of *Encoding monomorphic and polymorphic types*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/191510/>

Version: Published Version

Article:

Blanchette, J., Böhme, S., Popescu, A. et al. (1 more author) (2017) Encoding monomorphic and polymorphic types. *Logical Methods in Computer Science*, 12 (4). pp. 1-52. ISSN 1860-5974

[https://doi.org/10.2168/lmcs-12\(4:13\)2016](https://doi.org/10.2168/lmcs-12(4:13)2016)

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

ENCODING MONOMORPHIC AND POLYMORPHIC TYPES

JASMIN CHRISTIAN BLANCHETTE^a, SASCHA BÖHME^b, ANDREI POPESCU^c,
AND NICHOLAS SMALLBONE^d

^a Inria & LORIA, Nancy, France; Max-Planck-Institut für Informatik, Saarbrücken, Germany
e-mail address: jasmin.blanchette@{inria.fr,mpi-inf.mpg.de}

^b Fakultät für Informatik, Technische Universität München, Germany
e-mail address: boehmes@in.tum.de

^c Department of Computer Science, School of Science and Technology, Middlesex University, UK
e-mail address: a.popescu@mdx.ac.uk

^d Dept. of CSE, Chalmers University of Technology, Gothenburg, Sweden
e-mail address: nicsma@chalmers.se

ABSTRACT. Many automatic theorem provers are restricted to untyped logics, and existing translations from typed logics are bulky or unsound. Recent research proposes monotonicity as a means to remove some clutter when translating monomorphic to untyped first-order logic. Here we pursue this approach systematically, analysing formally a variety of encodings that further improve on efficiency while retaining soundness and completeness. We extend the approach to rank-1 polymorphism and present alternative schemes that lighten the translation of polymorphic symbols based on the novel notion of “cover”. The new encodings are implemented in Isabelle/HOL as part of the Sledgehammer tool. We include informal proofs of soundness and correctness, and have formalised the monomorphic part of this work in Isabelle/HOL. Our evaluation finds the new encodings vastly superior to previous schemes.

1. INTRODUCTION

Specification languages, proof assistants, and other theorem proving applications are typically based on polymorphism formalisms, but state-of-the-art automatic provers support only untyped or monomorphic logics. The existing sound (proof-reflecting) and complete (proof-preserving) translation schemes for polymorphic types, whether they revolve around functions (tags) or predicates (guards), produce clutter that severely hampers the proof search [22], and lighter approaches based on type arguments are unsound [22, 29]. As a result, application authors face a difficult choice between soundness and efficiency when interfacing with automatic provers.

The fourth author, together with Claessen and Lillieström [16], designed a pair of sound, complete, and efficient translations from monomorphic many-typed to untyped first-order logic with equality. The key insight is that *monotonic* types—types whose domain can be

2012 ACM CCS: [Theory of computation]: Logic—Automated reasoning.

extended with new elements while preserving satisfiability—can be merged. The remaining types can be made monotonic by introducing suitable protectors.

Example 1.1 (Monkey Village). Imagine a village of monkeys [16] where each monkey owns at least two bananas. The predicate $\text{owns} : \text{monkey} \times \text{banana} \rightarrow o$ (where o denotes truth values) associates monkeys with bananas, and the functions $\text{b}_1, \text{b}_2 : \text{monkey} \rightarrow \text{banana}$ witness the existence of each monkey’s minimum supply of bananas:

$$\begin{aligned} &\forall M : \text{monkey}. \text{owns}(M, \text{b}_1(M)) \wedge \text{owns}(M, \text{b}_2(M)) \\ &\forall M : \text{monkey}. \text{b}_1(M) \not\approx \text{b}_2(M) \\ &\forall M_1, M_2 : \text{monkey}, B : \text{banana}. \text{owns}(M_1, B) \wedge \text{owns}(M_2, B) \rightarrow M_1 \approx M_2 \end{aligned}$$

The axioms are clearly satisfiable.

In the monkey village of Example 1.1, the type *banana* is monotonic, because any model with b bananas can be extended to a model with $b' > b$ bananas. In contrast, *monkey* is nonmonotonic, because there can live at most $\lfloor b/2 \rfloor$ monkeys in a village with a finite supply of b bananas. Syntactically, the monotonicity of *banana* is inferable from the absence of a positive equality $B \approx t$ or $t \approx B$, where B is a variable of type *banana* and t is arbitrary; such a literal would be needed to make the type nonmonotonic.

The example can be encoded as follows, using the predicate $\mathfrak{g}_{\text{monkey}}$ to guard against ill-typed instantiations of M, M_1 , and M_2 :

$$\begin{aligned} &\exists M. \mathfrak{g}_{\text{monkey}}(M) \\ &\forall M. \mathfrak{g}_{\text{monkey}}(M) \rightarrow \text{owns}(M, \text{b}_1(M)) \wedge \text{owns}(M, \text{b}_2(M)) \\ &\forall M. \mathfrak{g}_{\text{monkey}}(M) \rightarrow \text{b}_1(M) \not\approx \text{b}_2(M) \\ &\forall M_1, M_2, B. \mathfrak{g}_{\text{monkey}}(M_1) \wedge \mathfrak{g}_{\text{monkey}}(M_2) \wedge \text{owns}(M_1, B) \wedge \text{owns}(M_2, B) \rightarrow M_1 \approx M_2 \end{aligned}$$

The first axiom states the existence of a monkey; this is necessary for completeness, since model carriers are required to be nonempty. Thanks to monotonicity, it is sound to omit all type information regarding bananas. The intuition behind this is that the $\mathfrak{g}_{\text{monkey}}$ predicate makes the problem fully monotonic, and for such problems it is possible to synchronise the cardinalities of the different types and to merge the types, yielding an equisatisfiable untyped (or singly typed) problem. For example, a model \mathcal{M} of the typed problem with m monkeys and b bananas will give rise to a model \mathcal{M}' of the untyped problem with b “bananamonkeys,” among which m are monkeys according to the interpretation of $\mathfrak{g}_{\text{monkey}}$; conversely, from a model of the untyped problem with b “bananamonkeys” including m values for which $\mathfrak{g}_{\text{monkey}}$ is true (with $m > 0$ thanks to the first axiom), it is easy to construct a model of the typed problem with b bananas and m monkeys.

Monotonicity is not decidable, but it can often be inferred using suitable calculi. In this article, we exploit this idea systematically, analysing a variety of encodings based on monotonicity: some are minor adaptations of existing ones, while others are novel encodings that further improve on the size of the translated formulae.

In addition, we generalise the monotonicity approach to a rank-1 polymorphic logic, as embodied by the typed first-order form TFF1 of the TPTP (Thousand of Problems for Theorem Provers) [7], a de facto standard implemented by a number of reasoning tools. Unfortunately, the presence of a single equality literal $X \approx t$ or $t \approx X$, where X is a polymorphic variable of type α , will lead the analysis to classify all types as possibly nonmonotonic and force the use of protectors everywhere, as in the traditional encodings. A typical example is the list axiom $\forall X : \alpha, Xs : \text{list}(\alpha). \text{hd}(\text{cons}(X, Xs)) \approx X$. We solve this issue through a novel scheme that reduces the clutter associated with nonmonotonic types, based on the

	Traditional (Polymorphic)	Cover-based (Polymorphic)	Monotonicity-based	
			Monomorphic	Polymorphic
Full type erasure	e (§3.1)			
Type arguments	a (§3.2)			
Type tags	t (§3.3)	t@ (§4.2)	$\tilde{t}?, \tilde{t}??$ (§5.4)	t?, t?? (§6.5)
Type guards	g (§3.4)	g@ (§4.1)	$\tilde{g}?, \tilde{g}??$ (§5.5)	g?, g?? (§6.6)

Figure 1: The main encodings

observation that protectors are required only when translating the particular formulae that prevent a type from being inferred monotonic. This contribution improves the monomorphic case as well: for the monkey village example, our scheme detects that the first two axioms are harmless and translates them without the g_{monkey} guards. (In fact, by appealing to a more general notion of monotonicity, it is possible to eliminate all type information in the monkey village problem in a sound fashion.)

Encoding types in an untyped logic is an old problem, and several solutions have been proposed in the literature. We start by reviewing four main traditional approaches (Section 3), which prepare the ground for the more advanced encodings presented in this article. Next, we present improvements of the traditional encodings that aim at reducing the clutter associated with polymorphic symbols, based on the novel notion of “cover” (Section 4). Then we move our attention to monotonicity-based encodings, which try to minimise the number of added tags or guards. We first present known and novel monotonicity-based schemes that handle only ground types (Section 5); these are interesting in their own right and serve as stepping stones for the full-blown polymorphic encodings (Section 6). Proofs of correctness accompany the descriptions of the new encodings. The proofs explicitly relate models of unencoded and encoded problems.

Figure 1 presents a brief overview of the main encodings. The traditional encodings are identified by single letters (e for full type erasure, a for type arguments, t for type tags, g for type guards). The nontraditional encodings append a suffix to the letter: $@$ (= cover-based), $?$ (= monotonicity-based, lightweight), or $??$ (= monotonicity-based, featherweight). The decoration \sim identifies the monomorphic version of an encoding. Among the nontraditional schemes, $\tilde{t}?$ and $\tilde{g}?$ are due to Claessen et al. [16]; the other encodings are novel.

A formalisation [8] of the monomorphic part of our results has been developed in the proof assistant Isabelle/HOL [24, 25]. The encodings have been implemented in Sledgehammer [3, 22], which provides a bridge between Isabelle/HOL and automatic theorem provers (Section 7). They were evaluated with Alt-Ergo, E, SPASS, Vampire, and Z3 on a benchmark suite consisting of proof goals from existing Isabelle formalisations (Section 8). Our comparisons include the traditional encodings as well as the provers’ native support for monomorphic types where it is available. Related work is considered at the end (Section 9).

We presented an earlier version of this work at the TACAS 2013 conference [5]. The current text extends the conference paper with detailed proofs and a discussion of implementational issues (Section 7). It also corrects the side conditions of two encoding definitions, which resulted in an unexpected incompleteness.

Convention 1.2. Given a name, such as t , that ranges over a certain domain, such as the set of terms, its overlined version \bar{t} ranges over lists (tuples) of items in this domain, and $|\bar{t}|$ denotes the length of \bar{t} . We also write $|A|$ for the cardinal of a set A . A set is *countable* if

it is finite or countably infinite. Given an element u , $(u)^n$ or u^n denotes the list consisting of n occurrences of u . Given a nonempty set A , we write $\varepsilon(A)$ for some arbitrary but fixed element of A . We use $\varepsilon(A)$ in definitions where the choice of the element does not matter, as long as it belongs to A .

2. BACKGROUND: LOGICS

This article involves three versions of classical first-order logic with equality: polymorphic, monomorphic, and untyped. They correspond to the TPTP syntaxes TFF1 [7], TFF0 [32], and FOF [30], respectively, excluding interpreted arithmetic.

2.1. Polymorphic First-Order Logic. The source logic is a rank-1 polymorphic logic as specified by TFF1 [7].

We fix \mathcal{A} , a countably infinite set of *type variables* with typical element α , and \mathcal{V} , a countably infinite set of *term variables* with typical element X .

Definition 2.1 (Syntax). A *polymorphic signature* is a triple $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$, where \mathcal{K} is a countable set of type constructors k with arities, \mathcal{F} is a countable set of function symbols f with arities, and \mathcal{P} is a countable set of predicate symbols p with arities.

For type constructors $k \in \mathcal{K}$, the arity is a natural number n . This association is written $k :: n$. Types, forming the set $Type_{\mathcal{K}}$, are then defined inductively starting with type variables and applying type constructors according to their arities:

$$\begin{array}{ll} \text{Types:} & \\ \sigma ::= k(\bar{\sigma}) & \text{where } k :: |\bar{\sigma}| \quad \text{constructor type} \\ \quad \quad \quad | \alpha & \quad \quad \quad \text{type variable} \end{array}$$

For function symbols $f \in \mathcal{F}$, the arity is a triple $(\bar{\alpha}, \bar{\sigma}, \sigma)$, where $\bar{\alpha}$ is a list of distinct type variables, $\bar{\sigma}$ is a list of types, and σ is a type such that all the variables appearing in $\bar{\sigma}$ and σ are among the ones in $\bar{\alpha}$. We write this association as $f : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \sigma$ or $f : \forall \bar{\alpha}. \sigma_1 \times \cdots \times \sigma_n \rightarrow \sigma$. Finally, for predicate symbols $p \in \mathcal{P}$, the arity is a pair $(\bar{\alpha}, \bar{\sigma})$, where $\bar{\alpha}$ is a list of distinct type variables and $\bar{\sigma}$ is a list of types. We write this association as $p : \forall \bar{\alpha}. \bar{\sigma} \rightarrow o$ or $p : \forall \bar{\alpha}. \sigma_1 \times \cdots \times \sigma_n \rightarrow o$.

Overloading of different arities for the same symbol is excluded by definition. The symbols \forall , \times , \rightarrow , and o (omicron) are not type constructors but syntax. The arity declarations for function and predicate symbols can be seen as instances of the general syntax $s : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \zeta$, where $s \in \mathcal{F} \uplus \mathcal{P}$ and ζ is either a type or o . An application of s will require $|\bar{\alpha}|$ type arguments (written in angle brackets) and $|\bar{\sigma}|$ term arguments (written in parentheses). In examples, we may omit type arguments from $\bar{\alpha}$ that are irrelevant or clear from the context.

A *typed (term) variable* is a pair of a variable and a type, written X^σ . We fix $\mathcal{V}_{\text{typed}}$, the set of type variables. The terms and formulae are defined below.

$$\begin{array}{ll} \text{Terms:} & \\ t ::= f\langle \bar{\sigma} \rangle(\bar{t}) & \text{function term} \\ \quad \quad \quad | X^\sigma & \text{typed variable} \end{array}$$

Formulae:

$\varphi ::= p\langle\bar{\sigma}\rangle(\bar{t}) \mid \neg p\langle\bar{\sigma}\rangle(\bar{t})$	predicate literal
$t_1 \approx t_2 \mid t_1 \not\approx t_2$	equality literal
$\varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2$	binary connective
$\forall X : \sigma. \varphi \mid \exists X : \sigma. \varphi$	term quantification
$\forall \alpha. \varphi$	type quantification

All type quantification is universal. We also impose the following syntactic restriction, not captured in the formation rules: type quantifiers may only occur at the top of the formula, i.e. not underneath a term quantifier or a connective, and formulae are in negation normal form (NNF), which is most suitable for defining and reasoning about translations. We sometimes use implication $\varphi_1 \wedge \cdots \wedge \varphi_m \rightarrow \psi_1 \vee \cdots \vee \psi_n$ as an abbreviation for $\neg\varphi_1 \vee \cdots \vee \neg\varphi_m \vee \psi_1 \vee \cdots \vee \psi_n$ to enhance readability. In examples, we freely nest quantifiers and connectives.

Definition 2.2 (Free and Fresh Variables, Groundness, and Sentences). $\text{TVars}(\sigma)$, $\text{TVars}(t)$, and $\text{FTVars}(\varphi)$ denote the sets of type variables occurring in the type σ , occurring in the term t , and occurring freely in the formula φ , respectively. Similarly, $\text{Vars}(t)$ and $\text{FVars}(\varphi)$ denote the sets of typed term variables occurring in the term t and occurring freely in the formula φ , respectively. For example, $\text{TVars}(X^\sigma) = \text{TVars}(\sigma)$, $\text{Vars}(X^\sigma) = \{X^\sigma\}$, $\text{FVars}(\forall X : \sigma. \varphi) = \text{FVars}(\exists X : \sigma. \varphi) = \text{FVars}(\varphi) - \{X^\sigma\}$. A type σ is *ground* if $\text{TVars}(\sigma) = \emptyset$. We let $GType_\Sigma$ denote the set of ground types. A term t is *ground* if $\text{TVars}(t) = \emptyset$ and $\text{Vars}(t) = \emptyset$. A formula φ is a *sentence* if $\text{TVars}(\varphi) = \emptyset$ and $\text{FVars}(\varphi) = \emptyset$.

Convention 2.3. We assume that the set of variables \mathcal{V} is partitioned in two infinite sets: \mathcal{V}^\forall , of *universal* variables, and \mathcal{V}^\exists , of *existential* variables. The former are the only ones allowed to be universally quantified, and the latter are the only ones allowed to be existentially quantified. We let $\mathcal{V}_{\text{typed}}^\forall$ for the set of typed variable X^σ with $X \in \mathcal{V}^\forall$; and similarly for $\mathcal{V}_{\text{typed}}^\exists$. We also assume that each type or term variable is bound only once in a formula, and that if the typed variable X^σ appears in the scope of a quantification $\forall X : \sigma'$ or $\exists X : \sigma'$, then $\sigma = \sigma'$. The last assumption allows us to omit the superscript σ from X^σ in examples.

The typing rules and semantics of the logic are modelled after those of TFF1. Briefly, the type arguments completely determine the types of the term arguments and, for functions, of the result. Polymorphic symbols are interpreted as families of functions or predicates indexed by domains corresponding to ground types. All types are inhabited (nonempty).

Definition 2.4 (Type Substitution). A *type substitution* ρ is a function that maps every type variable α to a type, written $\alpha\rho$. It is extended from type variables to types and type tuples in the standard way: $k(\bar{t})\rho = k(\bar{t}\rho)$ and $(t_1, \dots, t_n)\rho = (t_1\rho, \dots, t_n\rho)$.

Definition 2.5 (Typing Rules). A judgement $t : \sigma$ expresses that the term t is *well typed* and has type σ . A judgement $\varphi : o$ expresses that the formula φ is *well typed*. The *typing rules* of polymorphic first-order logic are given below:

$$\frac{}{X^\sigma : \sigma} \quad \frac{f : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \sigma \quad t_j : \sigma_j \rho \text{ for all } j}{f\langle\bar{\alpha}\rho\rangle(\bar{t}) : \sigma \rho}$$

$$\begin{array}{c}
\frac{p : \forall \bar{\alpha}. \bar{\sigma} \rightarrow o \quad t_j : \sigma_j \rho \text{ for all } j}{p \langle \bar{\alpha} \rho \rangle (\bar{t}) : o} \quad \frac{p : \forall \bar{\alpha}. \bar{\sigma} \rightarrow o \quad t_j : \sigma_j \rho \text{ for all } j}{\neg p \langle \bar{\alpha} \rho \rangle (\bar{t}) : o} \\
\frac{t_1 : \sigma \quad t_2 : \sigma}{t_1 \approx t_2 : o} \quad \frac{t_1 : \sigma \quad t_2 : \sigma}{t_1 \not\approx t_2 : o} \quad \frac{\varphi_1 : o \quad \varphi_2 : o}{\varphi_1 \wedge \varphi_2 : o} \\
\frac{\varphi_1 : o \quad \varphi_2 : o}{\varphi_1 \vee \varphi_2 : o} \quad \frac{\varphi : o}{\forall X : \sigma. \varphi : o} \quad \frac{\varphi : o}{\exists X : \sigma. \varphi : o} \quad \frac{\varphi : o}{\forall \alpha. \varphi : o}
\end{array}$$

Definition 2.6 (Problem). A *problem* is a set of (well-typed) sentences.

Lemma 2.7. *For all terms t , there exists at most one type σ such that $t : \sigma$.*

Proof. Immediate by induction on t . □

Convention 2.8. We write t^σ to indicate that the term t is well typed and its (unique) type is σ . This convention is consistent with the notation X^σ for term variables of type σ .

Definition 2.9 (Semantics). Let $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ be a polymorphic signature. A *structure* \mathcal{M} for Σ is a tuple of families $(\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$, where

- \mathbb{D} is a nonempty collection of nonempty sets called the *domains*;
- if $k :: n$, then $k^{\mathcal{M}} : \mathbb{D}^n \rightarrow \mathbb{D}$. Given a *type variable valuation* $\theta : \mathcal{A} \rightarrow \mathbb{D}$, this induces an *interpretation* of types $\llbracket \cdot \rrbracket_\theta^{\mathcal{M}}$ defined by the equations $\llbracket k(\bar{\sigma}) \rrbracket_\theta^{\mathcal{M}} = k^{\mathcal{M}}(\llbracket \bar{\sigma} \rrbracket_\theta^{\mathcal{M}})$ and $\llbracket \alpha \rrbracket_\theta^{\mathcal{M}} = \theta(\alpha)$;
- if $f : \forall \alpha_1, \dots, \alpha_m. \sigma_1 \times \dots \times \sigma_n \rightarrow \sigma$, then $f^{\mathcal{M}} : \prod_{D \in \mathbb{D}^m} \llbracket \sigma_1 \rrbracket_{\theta_D}^{\mathcal{M}} \times \dots \times \llbracket \sigma_n \rrbracket_{\theta_D}^{\mathcal{M}} \rightarrow \llbracket \sigma \rrbracket_{\theta_D}^{\mathcal{M}}$, where θ_D is a type variable valuation mapping each α_i to D_i ;
- if $p : \forall \alpha_1, \dots, \alpha_m. \sigma_1 \times \dots \times \sigma_n \rightarrow o$, then $p^{\mathcal{M}} \subseteq \prod_{D \in \mathbb{D}^m} \llbracket \sigma_1 \rrbracket_{\theta_D}^{\mathcal{M}} \times \dots \times \llbracket \sigma_n \rrbracket_{\theta_D}^{\mathcal{M}}$, where θ_D is as above.

As expected from the arity of f , the interpretation $f^{\mathcal{M}}$ is a function first taking m type arguments and then n data (element) arguments. Similarly, the interpretation $p^{\mathcal{M}}$ is a predicate respecting the arity of p .

Given a type variable valuation θ and a compatible *term variable valuation* $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} \llbracket \sigma \rrbracket_\theta^{\mathcal{M}}$, the *interpretation* of terms and formulae by the structure \mathcal{M} is as follows:

$$\begin{array}{ll}
\llbracket f \langle \bar{\sigma} \rangle (\bar{t}) \rrbracket_{\theta, \xi}^{\mathcal{M}} = f^{\mathcal{M}}(\llbracket \bar{\sigma} \rrbracket_\theta^{\mathcal{M}}) (\llbracket \bar{t} \rrbracket_{\theta, \xi}^{\mathcal{M}}) & \llbracket X^\sigma \rrbracket_{\theta, \xi}^{\mathcal{M}} = \xi(X)(\sigma) \\
\llbracket p \langle \bar{\sigma} \rangle (\bar{t}) \rrbracket_{\theta, \xi}^{\mathcal{M}} = p^{\mathcal{M}}(\llbracket \bar{\sigma} \rrbracket_\theta^{\mathcal{M}}) (\llbracket \bar{t} \rrbracket_{\theta, \xi}^{\mathcal{M}}) & \llbracket t_1 \approx t_2 \rrbracket_{\theta, \xi}^{\mathcal{M}} = (\llbracket t_1 \rrbracket_{\theta, \xi}^{\mathcal{M}} = \llbracket t_2 \rrbracket_{\theta, \xi}^{\mathcal{M}}) \\
\llbracket \neg p \langle \bar{\sigma} \rangle (\bar{t}) \rrbracket_{\theta, \xi}^{\mathcal{M}} = \neg p^{\mathcal{M}}(\llbracket \bar{\sigma} \rrbracket_\theta^{\mathcal{M}}) (\llbracket \bar{t} \rrbracket_{\theta, \xi}^{\mathcal{M}}) & \llbracket t_1 \not\approx t_2 \rrbracket_{\theta, \xi}^{\mathcal{M}} = (\llbracket t_1 \rrbracket_{\theta, \xi}^{\mathcal{M}} \neq \llbracket t_2 \rrbracket_{\theta, \xi}^{\mathcal{M}}) \\
\llbracket \varphi_1 \wedge \varphi_2 \rrbracket_{\theta, \xi}^{\mathcal{M}} = \llbracket \varphi_1 \rrbracket_{\theta, \xi}^{\mathcal{M}} \wedge \llbracket \varphi_2 \rrbracket_{\theta, \xi}^{\mathcal{M}} & \llbracket \forall X : \sigma. \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}} = \forall a \in \llbracket \sigma \rrbracket_\theta^{\mathcal{M}}. \llbracket \varphi \rrbracket_{\theta, \xi[X \mapsto a]}^{\mathcal{M}} \\
\llbracket \varphi_1 \vee \varphi_2 \rrbracket_{\theta, \xi}^{\mathcal{M}} = \llbracket \varphi_1 \rrbracket_{\theta, \xi}^{\mathcal{M}} \vee \llbracket \varphi_2 \rrbracket_{\theta, \xi}^{\mathcal{M}} & \llbracket \exists X : \sigma. \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}} = \exists a \in \llbracket \sigma \rrbracket_\theta^{\mathcal{M}}. \llbracket \varphi \rrbracket_{\theta, \xi[X \mapsto a]}^{\mathcal{M}} \\
\llbracket \forall \alpha. \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}} = \forall D \in \mathbb{D}. \llbracket \varphi \rrbracket_{\theta[\alpha \mapsto D], \xi}^{\mathcal{M}}
\end{array}$$

We omit irrelevant subscripts to $\llbracket \cdot \rrbracket$, writing $\llbracket \sigma \rrbracket^{\mathcal{M}}$ if σ is ground and $\llbracket \varphi \rrbracket^{\mathcal{M}}$ if φ is a sentence.

A structure \mathcal{M} is a *model* of a problem Φ if $\llbracket \varphi \rrbracket^{\mathcal{M}}$ is true for every $\varphi \in \Phi$. A problem that has a model is *satisfiable*.

Example 2.10 (Algebraic Lists). The following axioms induce a minimalistic first-order theory of algebraic lists that will serve as our main running example:

$$\begin{aligned}
& \forall \alpha. \forall X : \alpha, Xs : list(\alpha). \text{nil} \not\approx \text{cons}(X, Xs) \\
& \forall \alpha. \forall Xs : list(\alpha). Xs \approx \text{nil} \vee (\exists Y : \alpha, Ys : list(\alpha). Xs \approx \text{cons}(Y, Ys)) \\
& \forall \alpha. \forall X : \alpha, Xs : list(\alpha). \text{hd}(\text{cons}(X, Xs)) \approx X \wedge \text{tl}(\text{cons}(X, Xs)) \approx Xs
\end{aligned}$$

We conjecture that cons is injective. The conjecture's negation can be expressed employing an unknown but fixed Skolem type b :

$$\exists X, Y : b, Xs, Ys : list(b). \text{cons}(X, Xs) \approx \text{cons}(Y, Ys) \wedge (X \not\approx Y \vee Xs \not\approx Ys)$$

Because the hd and tl equations force injectivity of cons in both arguments, the problem consisting of the three axioms and the negated conjecture is unsatisfiable. The conjecture is a consequence of the axioms.

We are interested in encoding polymorphic problems Φ in a manner that preserves and reflects their satisfiability. It will be technically convenient to assume that their signatures have at least one nullary type constructor, so that the set of ground types is nonempty. It is obvious that this assumption is harmless: if it is not satisfied, we simply extend the signature with a distinguished nullary type constructor $\iota :: 0$. Since ι does not appear in the formulae of Φ and since in models the set of domains is assumed to be nonempty, this signature extension does not affect its satisfiability: given a model of Φ in the original signature, we obtain one in the extended signature by interpreting ι as an arbitrary domain; given a model in the extended signature, we obtain one in the original signature by omitting the interpretation of ι .

Convention 2.11. For all polymorphic signatures $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ that we consider, we assume that \mathcal{K} contains at least one nullary type constructor.

The following lemma shows that, in structures, the collection of domains can be regarded as a copy of the ground types.

Lemma 2.12. *If a polymorphic Σ -problem Φ has a model, it also has a model $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, _ , _)$ such that the following conditions are met:*

- (1) *each $k^{\mathcal{M}}$ is injective, and $k^{\mathcal{M}}(\bar{D}) \neq k^{\mathcal{M}}(\bar{E})$ whenever $k \neq k'$;*
- (2) $\mathbb{D} = \{\llbracket \tau \rrbracket^{\mathcal{M}} \mid \tau \in GType_{\Sigma}\}$;
- (3) *the type interpretation function $\llbracket _ \rrbracket^{\mathcal{M}}$ is a bijection between $GType$ and \mathbb{D} ;*
- (4) \mathbb{D} *is countable;*
- (5) \mathbb{D} *is disjoint from each $D \in \mathbb{D}$, and any distinct $D_1, D_2 \in \mathbb{D}$ are disjoint.*

Proof. Assume Φ has a model \mathcal{M} . To prove (1), we construct a model \mathcal{M}' from \mathcal{M} by tagging the domains with types, i.e. by defining $\mathbb{D}' = \{D \times \{\sigma\} \mid D \in \mathbb{D}' \wedge \sigma \in GType_{\Sigma}\}$ and maintaining types across the application of type constructors: $k^{\mathcal{M}'}(D_1 \times \{\sigma_1\}, \dots, D_n \times \{\sigma_n\}) = k^{\mathcal{M}}(D_1, \dots, D_n) \times \{k(\sigma_1, \dots, \sigma_n)\}$. The interpretations of the function and predicate symbols are adjusted accordingly. It is easy to prove that \mathcal{M} and \mathcal{M}' satisfy the same polymorphic formulae; in particular, \mathcal{M}' is a model of Φ . For (2), we define a model $\mathcal{M}'' = (\mathbb{D}'', (k^{\mathcal{M}''})_{k \in \mathcal{K}}, _ , _)$ from \mathcal{M}' by taking $\mathbb{D}'' \subseteq \mathbb{D}'$ to be the image of $\llbracket _ \rrbracket^{\mathcal{M}'} : GType_{\Sigma} \rightarrow \mathbb{D}'$ and by taking $k^{\mathcal{M}''}$ and $s^{\mathcal{M}''}$ to be the restrictions of $k^{\mathcal{M}'}$ and $s^{\mathcal{M}'}$. Thanks to Convention 2.11, \mathbb{D}'' is nonempty, and moreover $k^{\mathcal{M}''}$ is well defined on \mathbb{D} . Again, it is easy to prove that \mathcal{M}'' satisfies all the polymorphic formulae that \mathcal{M}' satisfies. In particular, \mathcal{M}'' is a model of Φ . Points (3) and (4) follow from (1) and (2). Finally, we prove (5) by replacing the domains with disjoint copies of them. \square

2.2. Monomorphic First-Order Logic. Monomorphic first-order logic, more commonly known as many-sorted first-order logic and corresponding to TPTP TFF0 [32], has signatures $\Sigma = (\text{Type}, \mathcal{F}, \mathcal{P})$, where Type is a countable set of types (or sorts) ranged over by σ , \mathcal{F} is a countable set of function symbols $f : \bar{\sigma} \rightarrow \sigma$ with arities, and \mathcal{P} is a countable set of predicate symbols $p : \bar{\sigma} \rightarrow o$ with arities. Σ -structures $\mathcal{M} = ((D_\sigma)_{\sigma \in \text{Type}}, (f^\mathcal{M})_{f \in \mathcal{F}}, (p^\mathcal{M})_{p \in \mathcal{P}})$ interpret the types as sets and the function and predicate symbols as functions and predicates of the suitable arities. Given a model \mathcal{M} and a valuation $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D_\sigma$, the interpretations of terms and formulae, $\llbracket t \rrbracket_\xi^\mathcal{M}$ and $\llbracket \varphi \rrbracket_\xi^\mathcal{M}$, are defined as expected.

Monomorphic first-order logic can be viewed as a special case of polymorphic first-order logic, with a polymorphic signature considered monomorphic when all its type constructors are nullary and the arities of its function and predicate symbols contain no type variables.

Example 2.13. A monomorphised version of the algebraic list problem of Example 2.10, with α instantiated by b , follows:

$$\begin{aligned} & \forall X : b, Xs : \text{list_}b. \text{nil}_b \not\approx \text{cons}_b(X, Xs) \\ & \forall Xs : \text{list_}b. Xs \approx \text{nil}_b \vee (\exists Y : b, Ys : \text{list_}b. Xs \approx \text{cons}_b(Y, Ys)) \\ & \forall X : b, Xs : \text{list_}b. \text{hd}_b(\text{cons}_b(X, Xs)) \approx X \wedge \text{tl}_b(\text{cons}_b(X, Xs)) \approx Xs \\ & \exists X, Y : b, Xs, Ys : \text{list_}b. \text{cons}_b(X, Xs) \approx \text{cons}_b(Y, Ys) \wedge (X \not\approx Y \vee Xs \not\approx Ys) \end{aligned}$$

Like the original polymorphic problem, it is unsatisfiable.

2.3. Untyped First-Order Logic. The final target logic for all our encodings, untyped first-order logic, coincides with the TPTP first-order form FOF [30]. This is the logic traditionally implemented in automatic theorem provers and finite model finders. An *untyped signature* is a pair $\Sigma = (\mathcal{F}, \mathcal{P})$, where \mathcal{F} and \mathcal{P} are countable sets of function and predicate symbols with arities, where the arities are natural numbers. The notation s^n indicates that the symbol s has arity n . The untyped syntax is identical to that of the monomorphic logic, except that variable terms do not contain types and quantification is written $\forall X. \varphi$ and $\exists X. \varphi$. The structures for $\Sigma = (\mathcal{F}, \mathcal{P})$ are triples $\mathcal{M} = (D, (f^\mathcal{M})_{f \in \mathcal{F}}, (p^\mathcal{M})_{p \in \mathcal{P}})$, where D is the domain and $f^\mathcal{M}$ and $p^\mathcal{M}$ are n -ary functions and predicates on \mathbb{D} , with n being the symbol's arity.

2.4. Type Encodings. The type encodings discussed in this article are given by functions that take problems Φ in a logic \mathcal{L} to problems Φ' in a logic \mathcal{L}' , where \mathcal{L} and \mathcal{L}' are among the three logics introduced above.

Convention 2.14. Each of the considered encodings will be specified by the following data:

- a function that maps \mathcal{L} -signatures Σ to \mathcal{L}' -signatures Σ' ;
- for all \mathcal{L} -signatures Σ and problems Φ over Σ :
 - a (possibly empty) set Ax_Φ of sentences over Σ , the *axioms*, added by the translation;
 - a function $\llbracket \cdot \rrbracket$ between formulae over Σ and formulae over Σ' , the *formula translation*.

The formula translation is typically based on a *term translation* $\llbracket \cdot \rrbracket$. The encoding $\llbracket \Phi \rrbracket$ of a problem Φ is given by the union between the axioms and the componentwise translations: $\llbracket \Phi \rrbracket = Ax_\Phi \cup \{\llbracket \varphi \rrbracket \mid \varphi \in \Phi\}$.

Central to this article are the notions of soundness and completeness of an encoding:

Definition 2.15 (Correctness). An encoding as above is *sound* for a class of problems \mathcal{C} if satisfiability of $\Phi \in \mathcal{C}$ implies satisfiability of $\langle\langle\Phi\rangle\rangle$; it is *complete* for \mathcal{C} if, given $\Phi \in \mathcal{C}$, satisfiability of $\langle\langle\Phi\rangle\rangle$ implies satisfiability of Φ ; it is *correct* for \mathcal{C} if it is both sound and complete (i.e. Φ and $\langle\langle\Phi\rangle\rangle$ are equisatisfiable). In case \mathcal{C} is the class of all problems, we omit it and simply call the encoding sound, complete, or correct.

3. TRADITIONAL TYPE ENCODINGS

There are four main traditional approaches to encoding polymorphic types: full type erasure, type arguments, type tags, and type guards [18,22,29,36]. Before introducing them, we first establish some conventions that will be useful throughout the article.

Convention 3.1. We will often need to extend signatures Σ with one or more of the following distinguished symbols. Whenever we employ them, we assume they are fresh with respect to Σ and have the indicated arities:

- a nullary type constructor ϑ ;
- a function symbol $t : \forall\alpha. \alpha \rightarrow \alpha$;
- a predicate symbol $g : \forall\alpha. \alpha \rightarrow o$.

Terms of type ϑ will be used to represent the types of Σ , t will be used to tag terms with type information, and g will be used to guard formulae with type information.

Convention 3.2. Since the sets of type and term variables, \mathcal{A} and \mathcal{V} , are countably infinite, we can fix a function from \mathcal{A} to \mathcal{V} , $\alpha \mapsto \mathcal{V}(\alpha)$, such that

- (1) it is injective, i.e. $\mathcal{V}(\alpha_1) = \mathcal{V}(\alpha_2)$ implies $\alpha_1 = \alpha_2$;
- (2) it allows for an infinite supply of term variables that do not correspond to type variables, i.e. the set of term variables not having the form $\mathcal{V}(\alpha)$ is infinite;
- (3) each $\mathcal{V}(\alpha)$ is a universal variable.

This function can be used to encode types as terms. Thanks to (2), we can safely assume that the source problems Φ do not contain variables of the form $\mathcal{V}(\alpha)$.

3.1. Full Type Erasure. The easiest way to translate a typed problem into an untyped logic is to erase all its type information, which means omitting all type arguments, type quantifiers, and types in term quantifiers. We call this encoding \mathbf{e} .

Definition 3.3 (Full Erasure \mathbf{e}). The *full type erasure* encoding \mathbf{e} translates a polymorphic problem over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ into an untyped problem over $\Sigma' = (\mathcal{F}', \mathcal{P}')$, where the symbols in Σ' have the same term arities as in Σ (but without type arguments). Thus, if $f : \forall\bar{\alpha}. \sigma_1 \times \dots \times \sigma_n \rightarrow \sigma$ and $p : \forall\bar{\alpha}. \sigma_1 \times \dots \times \sigma_n \rightarrow o$ are in \mathcal{F} and \mathcal{P} , respectively, then f and p have arities n in \mathcal{F}' and \mathcal{P}' , respectively. The encoding adds no axioms, and the term and formula translations $\langle\langle \cdot \rangle\rangle_{\mathbf{e}}$ are defined as follows:

$$\begin{aligned}
\langle\langle f(\bar{\sigma})(\bar{t}) \rangle\rangle_{\mathbf{e}} &= f(\langle\langle \bar{t} \rangle\rangle_{\mathbf{e}}) & \langle\langle X^\sigma \rangle\rangle_{\mathbf{e}} &= X \\
\langle\langle p(\bar{\sigma})(\bar{t}) \rangle\rangle_{\mathbf{e}} &= p(\langle\langle \bar{t} \rangle\rangle_{\mathbf{e}}) & \langle\langle \forall X : \sigma. \varphi \rangle\rangle_{\mathbf{e}} &= \forall X. \langle\langle \varphi \rangle\rangle_{\mathbf{e}} \\
\langle\langle \neg p(\bar{\sigma})(\bar{t}) \rangle\rangle_{\mathbf{e}} &= \neg p(\langle\langle \bar{t} \rangle\rangle_{\mathbf{e}}) & \langle\langle \exists X : \sigma. \varphi \rangle\rangle_{\mathbf{e}} &= \exists X. \langle\langle \varphi \rangle\rangle_{\mathbf{e}} \\
\langle\langle \forall\alpha. \varphi \rangle\rangle_{\mathbf{e}} &= \langle\langle \varphi \rangle\rangle_{\mathbf{e}}
\end{aligned}$$

Here and elsewhere, we omit the trivial cases where the function is simply applied to its subterms or subformulae, as in $\llbracket \varphi_1 \wedge \varphi_2 \rrbracket_{\mathbf{e}} = \llbracket \varphi_1 \rrbracket_{\mathbf{e}} \wedge \llbracket \varphi_2 \rrbracket_{\mathbf{e}}$. Recall that, according to Section 2.4, the \mathbf{e} translation of a problem Φ is simply the componentwise translation of its formulae: $\llbracket \Phi \rrbracket_{\mathbf{e}} = \{ \llbracket \varphi \rrbracket_{\mathbf{e}} \mid \varphi \in \Phi \}$.

Example 3.4. Encoded using \mathbf{e} , the monkey village axioms of Example 1.1 become

$$\begin{aligned} & \forall M. \text{owns}(M, \mathbf{b}_1(M)) \wedge \text{owns}(M, \mathbf{b}_2(M)) \\ & \forall M. \mathbf{b}_1(M) \not\approx \mathbf{b}_2(M) \\ & \forall M_1, M_2, B. \text{owns}(M_1, B) \wedge \text{owns}(M_2, B) \rightarrow M_1 \approx M_2 \end{aligned}$$

Like the original axioms, the encoded axioms are satisfiable: the requirement that each monkey possesses two bananas of its own can be met by taking an infinite domain (since $2k = k$ for any infinite cardinal k).

However, full type erasure is generally unsound in the presence of equality because equality can be used to encode cardinality constraints on domains. For example, the axiom $\forall U : \text{unit}. U \approx \text{unity}$ forces the domain of *unit* to have only one element. Its erasure, $\forall U. U \approx \text{unity}$, effectively restricts *all* types to one element; a contradiction is derivable from any disequality $t \not\approx u$ or any pair of clauses $p(\bar{t})$ and $\neg p(\bar{u})$. An expedient proposed by Meng and Paulson [22, §2.8], which they implemented in Sledgehammer, is to filter out all axioms of the form $\forall X : \sigma. X \approx \mathbf{a}_1 \vee \dots \vee X \approx \mathbf{a}_n$, but this makes the translation incomplete and generally does not suffice to prevent unsound cardinality reasoning.

An additional issue with full type erasure is that it confuses distinct monomorphic instances of polymorphic symbols. The formula $\mathbf{q}\langle a \rangle(\mathbf{f}\langle a \rangle) \wedge \neg \mathbf{q}\langle b \rangle(\mathbf{f}\langle b \rangle)$ is satisfiable, but its type erasure $\mathbf{q}\langle \mathbf{f} \rangle \wedge \neg \mathbf{q}\langle \mathbf{f} \rangle$ is unsatisfiable. A more intuitive example might be $N \not\approx 0 \rightarrow N > 0$, which we would expect to hold for the natural number versions of 0 and $>$ but not for integers or real numbers.

Nonetheless, full type erasure is complete, and this property will be useful later.

Theorem 3.5 (Completeness of \mathbf{e}). *Full type erasure is complete.*

Proof. From a model $\mathcal{M}' = (D, (f^{\mathcal{M}'})_{f \in \mathcal{F}'}, (p^{\mathcal{M}'})_{p \in \mathcal{P}'})$ of $\llbracket \Phi \rrbracket_{\mathbf{e}}$, we construct a structure $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{X}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ for the signature of Φ by taking the same domain for all types and interpreting all instances of each polymorphic symbol in the same way as \mathcal{M} :

- $\mathbb{D} = \{D\}$, and $k^{\mathcal{M}}$ maps everything to D ;
- if $s \in \mathcal{F} \uplus \mathcal{P}$, then $s^{\mathcal{M}}(\bar{D})(\bar{d}) = s^{\mathcal{M}'}(\bar{d})$.

Given $\theta : \mathcal{A} \rightarrow \mathbb{D}$ and $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}_{\Sigma}} \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$, we define $\xi' : \mathcal{V} \rightarrow D$ by $\xi'(X) = \xi(X)(\alpha)$, where α is any type variable. (The choice of α is irrelevant because θ maps all type variables to D , the only element of \mathbb{D} .) The next facts follow by structural induction on t and φ (for arbitrary θ and ξ):

- $\llbracket t \rrbracket_{\theta, \xi}^{\mathcal{M}} = \llbracket \llbracket t \rrbracket_{\mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'}$;
- $\llbracket \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}} = \llbracket \llbracket \varphi \rrbracket_{\mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'}$.

In particular, for sentences, we have $\llbracket \varphi \rrbracket^{\mathcal{M}} = \llbracket \llbracket \varphi \rrbracket_{\mathbf{e}} \rrbracket^{\mathcal{M}'}$; and since \mathcal{M}' is a model of $\llbracket \Phi \rrbracket_{\mathbf{e}}$, it follows that \mathcal{M} is a model of Φ . \square

By way of composition, the \mathbf{e} encoding lies at the heart of all the encodings presented in this article. Given n encodings $\mathbf{x}_1, \dots, \mathbf{x}_n$, we write $\llbracket \cdot \rrbracket_{\mathbf{x}_1; \dots; \mathbf{x}_n}$ for the composition $\llbracket \cdot \rrbracket_{\mathbf{x}_n} \circ \dots \circ \llbracket \cdot \rrbracket_{\mathbf{x}_1}$. Typically, n will be 2 or 3 and \mathbf{x}_n will be \mathbf{e} . Moreover, \mathbf{x}_i will be correct and will

transform the problem so that it belongs to a fragment for which x_{i+1} is also correct. This will ensure that the whole composition is correct. Finally, because x_2, \dots, x_n will always be fixed for a given x_1 , we will call the entire composition $\langle\langle \rangle\rangle_{x_1; \dots; x_n}$ the “ x_1 encoding”.

3.2. Type Arguments. A natural way to prevent the (unsoundness-causing) confusion arising with full type erasure is to encode types as terms in the untyped logic. Instances of polymorphic symbols can be distinguished using explicit type arguments, encoded as terms: n -ary type constructors k become n -ary function symbols k , and type variables α become term variables A . A polymorphic symbol with m type arguments is passed m additional term arguments. The example given in the previous subsection is translated to $q(a, f(a)) \wedge \neg q(b, f(b))$, and a fully polymorphic instance $f(\alpha)$ would be encoded as $f(A)$ (with A a term variable). We call this encoding \mathbf{a} .

We now proceed with first encoding the types in isolation and then the typed terms.

Definition 3.6 (Term Encoding of Types). Let \mathcal{K} be a finite set of n -ary type constructors. The *term encoding* of a polymorphic type over \mathcal{K} is a term over the signature $(\{\vartheta\}, \mathcal{K}', \emptyset)$, where \mathcal{K}' contains a function symbol $k : \vartheta^n \rightarrow \vartheta$ for each $k \in \mathcal{K}$ with $k :: n$. The encoding is specified by the following equations:

$$\langle\langle k(\bar{\sigma}) \rangle\rangle = k(\langle\langle \bar{\sigma} \rangle\rangle) \qquad \langle\langle \alpha \rangle\rangle = \mathcal{V}(\alpha)$$

Definition 3.7 (Traditional Arguments \mathbf{a}). We first define the encoding function $\langle\langle \rangle\rangle_{\mathbf{a}}$, which translates polymorphic problems over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ to polymorphic problems over $(\mathcal{K} \uplus \{\vartheta\}, \mathcal{F} \uplus \mathcal{K}', \mathcal{P})$, where ϑ and \mathcal{K}' are as in Definition 3.6. It adds no axioms, and its term and formula translations are defined as follows:

$$\begin{aligned} \langle\langle f(\bar{\sigma})(\bar{t}) \rangle\rangle_{\mathbf{a}} &= f(\bar{\sigma})(\langle\langle \bar{\sigma} \rangle\rangle, \langle\langle \bar{t} \rangle\rangle_{\mathbf{a}}) \\ \langle\langle p(\bar{\sigma})(\bar{t}) \rangle\rangle_{\mathbf{a}} &= p(\bar{\sigma})(\langle\langle \bar{\sigma} \rangle\rangle, \langle\langle \bar{t} \rangle\rangle_{\mathbf{a}}) & \langle\langle \forall \alpha. \varphi \rangle\rangle_{\mathbf{a}} &= \forall \alpha. \forall \langle\langle \alpha \rangle\rangle : \vartheta. \langle\langle \varphi \rangle\rangle_{\mathbf{a}} \\ \langle\langle \neg p(\bar{\sigma})(\bar{t}) \rangle\rangle_{\mathbf{a}} &= \neg p(\bar{\sigma})(\langle\langle \bar{\sigma} \rangle\rangle, \langle\langle \bar{t} \rangle\rangle_{\mathbf{a}}) \end{aligned}$$

(Again, we omit the trivial cases, e.g. $\langle\langle \forall X : \sigma. \varphi \rangle\rangle_{\mathbf{a}} = \forall X : \sigma. \langle\langle \varphi \rangle\rangle_{\mathbf{a}}$.) The *traditional type arguments* encoding \mathbf{a} is defined as the composition $\langle\langle \rangle\rangle_{\mathbf{a}; \mathbf{e}}$. It follows from the definition that \mathbf{a} translates a polymorphic problem over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ into an untyped problem over $\Sigma' = (\mathcal{F}' \uplus \mathcal{K}', \mathcal{P}')$, where the symbols in $\mathcal{F}', \mathcal{P}'$ are the same as those in \mathcal{F}, \mathcal{P} ; and for each symbol $s : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \varsigma \in \mathcal{F} \uplus \mathcal{P}$, the arity of s in Σ' is $|\bar{\alpha}| + |\bar{\sigma}|$.

Example 3.8. The \mathbf{a} encoding translates the algebraic list problem of Example 2.10 into the following untyped problem:

$$\begin{aligned} \forall A, X, Xs. \text{nil}(A) &\not\approx \text{cons}(A, X, Xs) \\ \forall A, Xs. Xs &\approx \text{nil}(A) \vee (\exists Y, Ys. Xs \approx \text{cons}(A, Y, Ys)) \\ \forall A, X, Xs. \text{hd}(A, \text{cons}(A, X, Xs)) &\approx X \wedge \text{tl}(A, \text{cons}(A, X, Xs)) \approx Xs \\ \exists X, Y, Xs, Ys. \text{cons}(b, X, Xs) &\approx \text{cons}(b, Y, Ys) \wedge (X \not\approx Y \vee Xs \not\approx Ys) \end{aligned}$$

The \mathbf{a} encoding coincides with \mathbf{e} for monomorphic problems and suffers from the same unsoundness with respect to equality and cardinality constraints. Nonetheless, \mathbf{a} will form the basis of all the sound polymorphic encodings in a slightly generalised version, called \mathbf{a}^x , for suitable instances of x .

Definition 3.9 (Type Argument Filter). Given a signature $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$, a *type argument filter* x maps any $s : \forall \alpha_1, \dots, \alpha_m. \bar{\sigma} \rightarrow \zeta$ to a subset $\mathbf{x}_s = \{i_1, \dots, i_{m'}\} \subseteq \{1, \dots, m\}$ of its type argument indices. Given a list \bar{z} of length m , $\mathbf{x}_s(\bar{z})$ denotes the sublist $z_{i_1}, \dots, z_{i_{m'}}$, where $i_1 < \dots < i_{m'}$. Filters are implicitly extended to $\{1\}$ for the distinguished symbols $t, g \notin \mathcal{F} \uplus \mathcal{P}$ introduced in Convention 3.1.

Definition 3.10 (Generic Arguments \mathbf{a}^x). Given a type argument filter x , we first define the encoding $\llbracket \cdot \rrbracket_{\mathbf{a}^x}$ that translates polymorphic problems over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ to polymorphic problems over $(\mathcal{K} \uplus \{\vartheta\}, \mathcal{F} \uplus \mathcal{K}', \mathcal{P})$, where ϑ and \mathcal{K}' are as in Definition 3.6. It adds no axioms, and its term and formula translations are defined as follows:

$$\begin{aligned} \llbracket f(\bar{\sigma})(\bar{t}) \rrbracket_{\mathbf{a}^x} &= f(\bar{\sigma})(\llbracket \mathbf{x}_f(\bar{\sigma}) \rrbracket, \llbracket \bar{t} \rrbracket_{\mathbf{a}^x}) \\ \llbracket p(\bar{\sigma})(\bar{t}) \rrbracket_{\mathbf{a}^x} &= p(\bar{\sigma})(\llbracket \mathbf{x}_p(\bar{\sigma}) \rrbracket, \llbracket \bar{t} \rrbracket_{\mathbf{a}^x}) & \llbracket \forall \alpha. \varphi \rrbracket_{\mathbf{a}^x} &= \forall \alpha. \forall \llbracket \alpha \rrbracket : \vartheta. \llbracket \varphi \rrbracket_{\mathbf{a}^x} \\ \llbracket \neg p(\bar{\sigma})(\bar{t}) \rrbracket_{\mathbf{a}^x} &= \neg p(\bar{\sigma})(\llbracket \mathbf{x}_p(\bar{\sigma}) \rrbracket, \llbracket \bar{t} \rrbracket_{\mathbf{a}^x}) \end{aligned}$$

The *generic type arguments* encoding \mathbf{a}^x is the composition $\llbracket \cdot \rrbracket_{\mathbf{a}^x, \mathbf{e}}$. It translates a polymorphic problem over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ into an untyped problem over $\Sigma' = (\mathcal{F}' \uplus \mathcal{K}, \mathcal{P}')$, where the symbols in $\mathcal{F}', \mathcal{P}'$ are the same as those in \mathcal{F}, \mathcal{P} ; and for each symbol $s : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \zeta \in \mathcal{F} \uplus \mathcal{P}$, the arity of s in Σ' is $|\mathbf{x}_s| + |\bar{\sigma}|$.

The \mathbf{e} and \mathbf{a} encodings correspond to the special cases of \mathbf{a}^x where x returns none or all of the type arguments, respectively.

Theorem 3.11 (Completeness of \mathbf{a}^x). *The type arguments encoding \mathbf{a}^x is complete.*

Proof. Recall that \mathbf{a}^x is defined as the composition of $\llbracket \cdot \rrbracket_{\mathbf{a}^x}$ and $\llbracket \cdot \rrbracket_{\mathbf{e}}$. Since $\llbracket \cdot \rrbracket_{\mathbf{e}}$ is complete by Theorem 3.5, it suffices to show that $\llbracket \cdot \rrbracket_{\mathbf{a}^x}$ is complete. Let $\mathcal{M}' = (\mathbb{D}', (k^{\mathcal{M}'})_{k \in \mathcal{K}' \uplus \{\vartheta\}}, (f^{\mathcal{M}'})_{f \in \mathcal{F}' \uplus \mathcal{K}'}, (p^{\mathcal{M}'})_{p \in \mathcal{P}'})$ be a model of $\llbracket \Phi \rrbracket_{\mathbf{a}^x}$. We will construct a structure $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ for the signature of Φ by taking the same domains as \mathcal{M}' , interpreting the type constructors other than ϑ in the same way, and interpreting the function and predicate symbols s as in \mathcal{M}' , but supplying, for the extra arguments, a suitable tuple from $\llbracket \vartheta \rrbracket^{\mathcal{M}'}$ that reflects the type arguments.

To this end, we first apply Lemma 2.12 to obtain that every element of \mathbb{D}' is uniquely represented as $\llbracket \tau \rrbracket^{\mathcal{M}'}$ with $\tau \in GType_{\Sigma'}$. We define

- $\mathbb{D} = \mathbb{D}'$ and $k^{\mathcal{M}} = k^{\mathcal{M}'}$;
- if $s \in \mathcal{F} \uplus \mathcal{P}$, then

$$s^{\mathcal{M}}(\llbracket \bar{\tau} \rrbracket^{\mathcal{M}'}) (\bar{d}) = \begin{cases} s^{\mathcal{M}'}(\llbracket \bar{\tau} \rrbracket^{\mathcal{M}'})(\llbracket \mathbf{x}_s(\bar{\tau}) \rrbracket^{\mathcal{M}'}, \bar{d}) & \text{if } \mathbf{x}_s(\bar{\tau}) \in GType_{\Sigma'}^n \\ \text{anything} & \text{otherwise} \end{cases}$$

where n is the length of $\mathbf{x}_s(\bar{\tau})$.

Note how the tuple $\llbracket \mathbf{x}_s(\bar{\tau}) \rrbracket^{\mathcal{M}'}$ reflects the \mathbf{x}_s -selection from the type arguments $\llbracket \bar{\tau} \rrbracket^{\mathcal{M}'}$.

Given $\theta : \mathcal{A} \rightarrow \mathbb{D}$ and $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in Type_{\Sigma}} \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}'}$, we define $\xi' : \mathcal{V} \rightarrow \prod_{\sigma' \in Type_{\Sigma'}} \llbracket \sigma' \rrbracket_{\theta}^{\mathcal{M}'}$ by

$$\xi'(X)(\sigma') = \begin{cases} \xi'(X)(\sigma') & \text{if } \sigma' \in Type_{\Sigma} \\ \llbracket \langle \tau \rangle \rrbracket^{\mathcal{M}'} & \text{if } \sigma' = \vartheta, X = \mathcal{V}(\alpha), \text{ and } \theta(\alpha) = \llbracket \tau \rrbracket^{\mathcal{M}'} \text{ for } \tau \in GType_{\Sigma} \\ \text{anything} & \text{otherwise} \end{cases}$$

The next facts follow by structural induction on t and φ (for arbitrary θ and ξ):

- $\llbracket t \rrbracket_{\theta, \xi}^{\mathcal{M}} = \llbracket \langle t \rangle \mathbf{a}^x \rrbracket_{\theta, \xi'}^{\mathcal{M}'}$;
- $\llbracket \langle \varphi \rangle \mathbf{a}^x \rrbracket_{\theta, \xi'}^{\mathcal{M}'}$ implies $\llbracket \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}}$.

(The reason why for formula interpretation we have only implication and not equality is the $\forall \alpha$ case: when encoding a universally quantified formula $\forall \alpha. \varphi$, the result $\forall \alpha. \forall \langle \alpha \rangle : \vartheta. \langle \varphi \rangle_{\mathbf{a}^x}$ introduces quantification over two variables, α and $\langle \alpha \rangle$, whose interpretations need not be synchronised. As a result, $\forall \alpha. \forall \langle \alpha \rangle : \vartheta. \langle \varphi \rangle_{\mathbf{a}^x}$ could be stronger than $\forall \alpha. \varphi$.) In particular, for a sentence φ , $\llbracket \langle \varphi \rangle \mathbf{a}^x \rrbracket^{\mathcal{M}'}$ implies $\llbracket \varphi \rrbracket^{\mathcal{M}}$, and hence \mathcal{M} is a model of Φ because \mathcal{M}' is a model of $\langle \Phi \rangle_{\mathbf{a}^x}$. \square

3.3. Type Tags. An intuitive approach to encode type information soundly (and completely) is to wrap each term and subterm with its type using type tags. For polymorphic type systems, this scheme relies on a distinguished binary function $t(\langle \sigma \rangle, t)$ that “annotates” each term t with its type σ encoded as a term $\langle \sigma \rangle$. The tags make most type arguments superfluous. We call this scheme t , after the tag function of the same name. It is defined as a two-stage process: the first stage adds tags $t(\sigma)(t)$ while preserving the polymorphism; the second stage encodes t ’s type argument as well as any phantom type arguments.

Definition 3.12 (Phantom Type Argument). Let $s : \forall \alpha_1, \dots, \alpha_m. \bar{\sigma} \rightarrow \varsigma \in \mathcal{F} \uplus \mathcal{P}$. The i th type argument is a *phantom* if α_i does not occur in $\bar{\sigma}$ or ς . Given a list $\bar{z} \equiv z_1, \dots, z_m$, $\text{phn}_s(\bar{z})$ denotes the sublist $z_{i_1}, \dots, z_{i_{m'}}$ corresponding to the positions in $\bar{\alpha}$ of the phantom type arguments.

Definition 3.13 (Traditional Tags t). We first define the encoding $\langle \rangle_t$ that translates polymorphic problems over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ to polymorphic problems over $(\mathcal{K}, \mathcal{F} \uplus \{t : \forall \alpha. \alpha \rightarrow \alpha\}, \mathcal{P})$. It adds no axioms, and its term and formula translations are defined as follows:

$$\llbracket f(\sigma)(\bar{t}) \rrbracket_t = \lfloor f(\sigma)(\langle \bar{t} \rangle_t) \rfloor \quad \llbracket X^\sigma \rrbracket_t = \lfloor \langle X^\sigma \rangle_t \rfloor \quad \text{with } \lfloor t^\sigma \rfloor = t(\sigma)(t)$$

The *traditional type tags* encoding t is the composition $\langle \rangle_{t; \mathbf{a}^{\text{phn}; \mathbf{e}}}$. It translates a polymorphic problem over Σ into an untyped problem over $\Sigma' = (\mathcal{F}' \uplus \mathcal{K}' \uplus \{t^2\}, \mathcal{P}')$, where $\mathcal{K}', \mathcal{F}', \mathcal{P}'$ are as for \mathbf{a}^{phn} (i.e. \mathbf{a}^x with $x = \text{phn}$).

Example 3.14. The t encoding translates the algebraic list problem of Example 2.10 into the following equisatisfiable untyped problem:

$$\begin{aligned} & \forall A, X, Xs. t(\text{list}(A), \text{nil}) \not\approx t(\text{list}(A), \text{cons}(t(A, X), t(\text{list}(A), Xs))) \\ & \forall A, Xs. t(\text{list}(A), Xs) \approx t(\text{list}(A), \text{nil}) \vee \\ & \quad (\exists Y, Ys. t(\text{list}(A), Xs) \approx t(\text{list}(A), \text{cons}(t(A, Y), t(\text{list}(A), Ys)))) \\ & \forall A, X, Xs. t(A, \text{hd}(t(\text{list}(A), \text{cons}(t(A, X), t(\text{list}(A), Xs)))))) \approx t(A, X) \wedge \\ & \quad t(\text{list}(A), \text{tl}(t(\text{list}(A), \text{cons}(t(A, X), t(\text{list}(A), Xs)))))) \approx t(\text{list}(A), Xs) \\ & \exists X, Y, Xs, Ys. t(\text{list}(b), \text{cons}(t(b, X), t(\text{list}(b), Xs))) \approx \\ & \quad t(\text{list}(b), \text{cons}(t(b, Y), t(\text{list}(b), Ys))) \wedge \\ & \quad (t(b, X) \not\approx t(b, Y) \vee t(\text{list}(b), Xs) \not\approx t(\text{list}(b), Ys)) \end{aligned}$$

Since there are no phantoms in this example, \mathbf{a}^{phn} adds no extra arguments. All type information is carried by the t function’s first argument.

Example 3.15. Consider the following formula, with $\text{linorder} : \forall \alpha. o$ and $\text{less_eq} : \forall \alpha. \alpha \times \alpha \rightarrow o$:

$$\forall \alpha. \forall X : \alpha, Y : \alpha. \text{linorder}\langle \alpha \rangle \rightarrow \text{less_eq}(X, Y) \vee \text{less_eq}(Y, X)$$

The α variable in `linorder`'s arity declaration is a phantom. The `t` encoding preserves it as an explicit term argument:

$$\forall A, X, Y. \text{linorder}(A) \rightarrow \text{less_eq}(t(A, X), t(A, Y)) \vee \text{less_eq}(t(A, Y), t(A, X))$$

As the formula suggests, phantom type arguments can be used to encode predicates on types, mimicking type classes [35].

We postpone the proof of the following theorem to Section 4.2, in the context of our improved encodings:

Theorem 3.16 (Correctness of `t`). *The traditional type tags encoding `t` is correct.*

3.4. Type Guards. Type tags heavily burden the terms. An alternative is to introduce type guards, which are predicates that restrict the range of variables. They take the form of a distinguished predicate $\mathbf{g}(\langle \sigma \rangle, t)$ that checks whether t has type σ . The terms are smaller than with tags, but the formulae contain more disjuncts.

With the type tags encoding, only phantom type arguments need to be encoded; here, we must encode any type arguments that cannot be read off the types of the term arguments. Thus, the type argument is encoded for `nil` $\langle \alpha \rangle$ (which has no term arguments) but omitted for `cons` $\langle \alpha \rangle(X, Xs)$, `hd` $\langle \alpha \rangle(Xs)$, and `tl` $\langle \alpha \rangle(Xs)$.

Definition 3.17 (Inferable Type Argument). Let $s : \forall \alpha_1, \dots, \alpha_m. \bar{\sigma} \rightarrow \zeta \in \mathcal{F} \uplus \mathcal{P}$. A type argument α_j is *inferable* if it occurs in some of the term arguments' types, i.e. if there exists an index i such that α_j occurs in σ_i . Given a list $\bar{z} \equiv z_1, \dots, z_m$, let $\text{inf}_s(\bar{z})$ denote the sublist $z_{i_1}, \dots, z_{i_{m'}}$ corresponding to the positions in $\bar{\alpha}$ of the inferable type arguments, and let $\text{ninf}_s(\bar{z})$ denote the sublist for noninferable type arguments.

Observe that a phantom type argument is in particular noninferable, and a noninferable nonphantom type argument is one that appears in ζ but not in $\bar{\sigma}$.

Definition 3.18 (Traditional Guards \mathbf{g}). We first define the encoding $\langle \langle \rangle \rangle_{\mathbf{g}}$, which translates a polymorphic problem over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ into an untyped problem over $(\mathcal{K}, \mathcal{F}, \mathcal{P} \uplus \{\mathbf{g} : \forall \alpha. \alpha \rightarrow o\})$. Its term and formula translations are defined as follows:

$$\langle \langle \forall X : \sigma. \varphi \rangle \rangle_{\mathbf{g}} = \forall X : \sigma. \mathbf{g}\langle \sigma \rangle(X) \rightarrow \langle \langle \varphi \rangle \rangle_{\mathbf{g}} \quad \langle \langle \exists X : \sigma. \varphi \rangle \rangle_{\mathbf{g}} = \exists X : \sigma. \mathbf{g}\langle \sigma \rangle(X) \wedge \langle \langle \varphi \rangle \rangle_{\mathbf{g}}$$

The encoding also adds the following *typing axioms*:

$$\begin{aligned} \forall \bar{\alpha}. \bar{X} : \bar{\sigma}. \left(\bigwedge_j \mathbf{g}\langle \sigma_j \rangle(X_j) \right) &\rightarrow \mathbf{g}\langle \sigma \rangle(f\langle \bar{\alpha} \rangle(\bar{X})) \quad \text{for } f : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \sigma \in \mathcal{F} \\ \forall \alpha. \exists X : \alpha. \mathbf{g}\langle \alpha \rangle(X) & \end{aligned}$$

(Following Convention 2.14, the translation of a problem is given by $\langle \langle \Phi \rangle \rangle_{\mathbf{g}} = \text{Ax} \cup \{\langle \langle \varphi \rangle \rangle_{\mathbf{g}} \mid \varphi \in \Phi\}$, where Ax are the typing axioms.) The *traditional type guards* encoding \mathbf{g} is defined as the composition $\langle \langle \rangle \rangle_{\mathbf{g}; \mathbf{a}^{\text{ninf}}; \mathbf{e}}$. It translates a polymorphic problem over Σ into an untyped problem over $\Sigma' = (\mathcal{F}' \uplus \mathcal{K}', \mathcal{P}' \uplus \{\mathbf{g}^2\})$, where $\mathcal{K}', \mathcal{F}', \mathcal{P}'$ are as for \mathbf{a}^{ninf} .

The last typing axiom in the above definition witnesses inhabitation of every type. It is necessary for completeness, in case some of the types do not appear in the result type of any function symbol.

Example 3.19. The \mathbf{g} encoding translates the algebraic list problem of Example 2.10 into the following:

$$\begin{aligned}
& \forall A. \mathbf{g}(\text{list}(A), \text{nil}(A)) \\
& \forall A, X, Xs. \mathbf{g}(A, X) \wedge \mathbf{g}(\text{list}(A), Xs) \rightarrow \mathbf{g}(\text{list}(A), \text{cons}(X, Xs)) \\
& \forall A, Xs. \mathbf{g}(\text{list}(A), Xs) \rightarrow \mathbf{g}(A, \text{hd}(Xs)) \\
& \forall A, Xs. \mathbf{g}(\text{list}(A), Xs) \rightarrow \mathbf{g}(\text{list}(A), \text{tl}(Xs)) \\
& \forall A. \exists X. \mathbf{g}(A, X) \\
& \forall A, X, Xs. \mathbf{g}(A, X) \wedge \mathbf{g}(\text{list}(A), Xs) \rightarrow \text{nil}(A) \not\approx \text{cons}(X, Xs) \\
& \forall A, Xs. \mathbf{g}(\text{list}(A), Xs) \rightarrow \\
& \quad Xs \approx \text{nil}(A) \vee (\exists Y, Ys. \mathbf{g}(A, Y) \wedge \mathbf{g}(\text{list}(A), Ys) \wedge Xs \approx \text{cons}(Y, Ys)) \\
& \forall A, X, Xs. \mathbf{g}(A, X) \wedge \mathbf{g}(\text{list}(A), Xs) \rightarrow \text{hd}(\text{cons}(X, Xs)) \approx X \wedge \text{tl}(\text{cons}(X, Xs)) \approx Xs \\
& \exists X, Y, Xs, Ys. \mathbf{g}(b, X) \wedge \mathbf{g}(b, Y) \wedge \mathbf{g}(\text{list}(b), Xs) \wedge \mathbf{g}(\text{list}(b), Ys) \wedge \\
& \quad \text{cons}(X, Xs) \approx \text{cons}(Y, Ys) \wedge (X \not\approx Y \vee Xs \not\approx Ys)
\end{aligned}$$

In this and later examples, we push guards inside past quantifiers and group them in a conjunction to enhance readability.

The typing axioms must in general be guarded. Without the guards, any cons , hd , or tl term could be typed with any type, defeating the purpose of the guard predicates.

Example 3.20. Consider the following formula, where $\text{inl} : \forall \alpha, \beta. \alpha \rightarrow \text{sum}(\alpha, \beta)$ and $\text{inr} : \forall \alpha, \beta. \beta \rightarrow \text{sum}(\alpha, \beta)$:

$$\forall \alpha, \beta. \forall X : \alpha, Y : \beta. \text{inl}\langle \alpha, \beta \rangle(X) \not\approx \text{inr}\langle \alpha, \beta \rangle(Y)$$

The β variable in inl 's arity declaration and the α in inr 's are noninferable. The \mathbf{g} encoding preserves them as explicit term arguments:

$$\forall A, B, X, Y. \mathbf{g}(A, X) \wedge \mathbf{g}(B, Y) \rightarrow \text{inl}(B, X) \not\approx \text{inr}(A, Y)$$

Theorem 3.21 (Correctness of \mathbf{g}). *The traditional type guards encoding \mathbf{g} is correct.*

Proof. This will be a consequence of Theorem 4.7 for our parameterised cover-based encoding $\mathbf{g}@$, since \mathbf{g} is a particular case of $\mathbf{g}@$. \square

Remark 3.22. The above encodings, as well as those discussed in the next sections, all lead to an untyped problem. An increasing number of automatic provers support monomorphic types, and it may seem desirable to exploit such support when it is available. With such provers, we can replace the \mathbf{e} encoding with a variant that enforces a basic type discipline by distinguishing two types, ϑ (for encoded types) and ι (for encoded terms). An incomplete (non-proof-preserving) alternative is to perform heuristic monomorphisation (Section 5.6). Hybrid schemes that exploit monomorphic types, including interpreted types, are also possible and have been studied by other researchers (Section 9).

4. COVER-BASED ENCODINGS OF POLYMORPHISM

Type tags and guards considerably increase the size of the problems passed to the automatic provers, with a dramatic impact on their performance. A lot of the type information generated by the traditional encodings \mathbf{t} and \mathbf{g} is redundant. For example, \mathbf{t} translates $\text{cons}\langle \alpha \rangle(X, Xs)$ to $\mathbf{t}(\text{list}(A), \text{cons}(\mathbf{t}(A, X), \mathbf{t}(\text{list}(A), Xs)))$, but intuitively only one of the three tags is necessary to specify the right type for the expression if we know the arity of cons . The cover-based encodings capitalise on this, by supplying only a minimum of protectors and

adding typing axioms that effectively compute the type of function symbols from a selection of their term arguments’ types—the “cover”.

Let us first rigorously define this notion of term arguments “covering” type arguments.

Definition 4.1 (Cover). Let $s : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \zeta \in \mathcal{F} \uplus \mathcal{P}$. A (*type argument*) *cover* $C \subseteq \{1, \dots, |\bar{\sigma}|\}$ for s is a set of term argument indices such that any inferable type argument can be inferred from a term argument whose index belongs to C , i.e. for all j , if α_j appears in $\bar{\sigma}$, it also appears in some σ_i such that $i \in C$. A cover C of s is *minimal* if no proper subset of C is a cover for s ; it is *maximal* if $C = \{1, \dots, |\bar{\sigma}|\}$. We let Cover_s denote an arbitrary but fixed cover for s .

In practice, we would normally take a minimal cover for Cover_s to reduce clutter. Accordingly, $\{1\}$ and $\{2\}$ are minimal covers for $\text{cons} : \forall \alpha. \alpha \times \text{list}(\alpha) \rightarrow \text{list}(\alpha)$, whereas $\{1, 2\}$ is a maximal cover.

Convention 4.2. As canonical cover, we arbitrarily choose $\text{Cover}_{\text{cons}} = \{1\}$.

The cover-based encoding $\mathbf{g}@$ introduced below is a generalisation of the traditional encoding \mathbf{g} . The two encodings coincide if Cover_s is chosen to be maximal for all symbols s . In contrast, the cover-based encoding $\mathbf{t}@$ is not exactly a generalisation of \mathbf{t} , although they share many ideas. For this reason, we momentarily depart from our general policy of considering tags before guards so that we can present the easier case first.

Intuitively, $\mathbf{g}@$ and $\mathbf{t}@$ ensure that each term argument position that is part of its enclosing function or predicate symbol’s cover has a unique type associated with it, from which the omitted type arguments can be inferred. Thus, $\mathbf{t}@$ translates $\text{cons}(\alpha)(X, Xs)$ to $\text{cons}(\mathbf{t}(A, X), Xs)$ with a type tag around X , effectively protecting the term from an ill-typed instantiation of X that would result in the wrong type argument being inferred for cons .

There is no need to protect the second argument, Xs , since it is not part of the cover. We call variables that occur in their enclosing symbol’s cover (and hence that “carry” some type arguments) “undercover variables”. It may seem dangerous to allow ill-typed terms to instantiate Xs , but this is not an issue because such terms cannot contribute meaningfully to a proof. At most, they can act as witnesses for the existence of terms of given types, but even in that capacity they are not necessary.

Definition 4.3 (Undercover Variable). The set of *undercover variables* $\text{UV}(\varphi)$ of a formula φ is defined by the equations

$$\begin{aligned}
 \text{UV}(f(\bar{\sigma})(\bar{t})) &= [\bar{t}]_f \cup \text{UV}(\bar{t}) & \text{UV}(X) &= \emptyset \\
 \text{UV}(p(\bar{\sigma})(\bar{t})) &= [\bar{t}]_p \cup \text{UV}(\bar{t}) & \text{UV}(t_1 \approx t_2) &= (\{t_1, t_2\} \cap \mathcal{V}_{\text{typed}}) \cup \text{UV}(t_1, t_2) \\
 \text{UV}(\neg p(\bar{\sigma})(\bar{t})) &= [\bar{t}]_p \cup \text{UV}(\bar{t}) & \text{UV}(t_1 \not\approx t_2) &= \text{UV}(t_1, t_2) \\
 \text{UV}(\varphi_1 \wedge \varphi_2) &= \text{UV}(\varphi_1, \varphi_2) & \text{UV}(\forall X : \sigma. \varphi) &= \text{UV}(\varphi) \\
 \text{UV}(\varphi_1 \vee \varphi_2) &= \text{UV}(\varphi_1, \varphi_2) & \text{UV}(\exists X : \sigma. \varphi) &= \text{UV}(\varphi) - \{X^\sigma\} \\
 \text{UV}(\forall \alpha. \varphi) &= \text{UV}(\varphi)
 \end{aligned}$$

where $[\bar{t}]_s = \{t_j \mid j \in \text{Cover}_s\} \cap \mathcal{V}_{\text{typed}}$ and $\text{UV}(\bar{t}) = \bigcup_j \text{UV}(t_j)$.

4.1. Cover-Based Type Guards. The cover-based encoding $\mathbf{g}@$ is similar to the traditional encoding \mathbf{g} , except that it guards only undercover occurrences of variables.

Definition 4.4 (Cover Guards $\mathbf{g}@$). The encoding $\llbracket \cdot \rrbracket_{\mathbf{g}@}$ is defined similarly to the encoding $\llbracket \cdot \rrbracket_{\mathbf{g}}$ (used for the traditional \mathbf{g} encoding) except for the \forall case in its formula translation and the typing axioms. Namely, the \forall case adds guards only for universally quantified variables that are undercover:

$$\llbracket \forall X : \sigma. \varphi \rrbracket_{\mathbf{g}@} = \forall X : \sigma. \begin{cases} \llbracket \varphi \rrbracket_{\mathbf{g}@} & \text{if } X \notin \text{UV}(\varphi) \\ \mathbf{g}\langle \sigma \rangle(X) \rightarrow \llbracket \varphi \rrbracket_{\mathbf{g}@} & \text{otherwise} \end{cases}$$

Moreover, the typing axioms take the cover into consideration:

$$\begin{aligned} \forall \bar{\alpha}. \bar{X} : \bar{\sigma}. \left(\bigwedge_{j \in \text{Cover}_f} \mathbf{g}\langle \sigma_j \rangle(X_j) \right) &\rightarrow \mathbf{g}\langle \sigma \rangle(f\langle \bar{\alpha} \rangle(\bar{X})) \quad \text{for } f : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \sigma \in \mathcal{F} \\ \forall \alpha. \exists X : \alpha. \mathbf{g}\langle \alpha \rangle(X) & \end{aligned}$$

The *cover-based type guards* encoding $\mathbf{g}@$ is defined as the composition $\llbracket \cdot \rrbracket_{\mathbf{g}@; \mathbf{a}^{\text{ninf}}; \mathbf{e}}$. It translates a polymorphic problem Φ over Σ into an untyped problem $\llbracket \Phi \rrbracket_{\mathbf{g}@; \mathbf{a}^{\text{ninf}}; \mathbf{e}}$ over $\Sigma' = (\mathcal{F}' \uplus \mathcal{K}', \mathcal{P}' \uplus \{\mathbf{g}^2\})$, where $\mathcal{K}', \mathcal{F}', \mathcal{P}'$ are as for \mathbf{a}^{ninf} .

Example 4.5. If we choose the cover for cons as in Convention 4.2, the $\mathbf{g}@$ encoding of the algebraic list problem is identical to the \mathbf{g} encoding (Example 3.19), except that the guard $\mathbf{g}(\text{list}(A), Xs)$ is omitted in typing axiom and one of the problem axioms:

$$\begin{aligned} \forall A, X, Xs. \mathbf{g}(A, X) &\rightarrow \mathbf{g}(\text{list}(A), \text{cons}(X, Xs)) \\ \forall A, X, Xs. \mathbf{g}(A, X) &\rightarrow \text{nil}(A) \not\approx \text{cons}(X, Xs) \end{aligned}$$

By leaving Xs unconstrained, the typing axiom for cons gives a type to some “ill-typed” terms, such as $\text{cons}(0^{\text{nat}}, 0^{\text{nat}})$. Intuitively, this is safe because such terms cannot be used to prove anything useful that could not be proved with a “well-typed” term. What matters is that “well-typed” terms are associated with their correct type and that “ill-typed” terms are given at most one type.

Lemma 4.6. *Let $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ be such that the domains in \mathbb{D} are mutually disjoint and the type constructors $k^{\mathcal{M}}$ are injective. Assume $\bar{\alpha} = (\alpha_1, \dots, \alpha_m)$, $\bar{\beta} = (\beta_1, \dots, \beta_n)$, $\bar{\sigma} = (\sigma_1, \dots, \sigma_u)$, $\bar{\tau} = (\tau_1, \dots, \tau_v)$, and $s : \forall (\bar{\alpha}, \bar{\beta}). (\bar{\sigma}, \bar{\tau}) \rightarrow \mathcal{S}$ in $\mathcal{F} \uplus \mathcal{P}$, such that the last n type arguments (corresponding to $\bar{\beta}$) are inferable and the first u term arguments (corresponding to $\bar{\sigma}$) constitute a cover for s . Let $(d_1, \dots, d_u) \in (\bigcup_{D \in \mathbb{D}} D)^u$. Then there exists at most one tuple $\bar{E} = (E_1, \dots, E_n) \in \mathbb{D}^n$ such that each d_i is in $\llbracket \sigma_i \rrbracket_{\theta}^{\mathcal{M}}$ for some θ that maps $\bar{\beta}$ to \bar{E} .*

Proof. From the inferability and cover assumptions, we have the condition $\text{TVars}(\bar{\beta}) \subseteq \text{TVars}(\bar{\sigma}, \bar{\tau}) = \text{TVars}(\bar{\sigma})$ on type variables. Assume another such tuple \bar{E}' exists. Let θ' be its corresponding substitution, and let $j \in \{1, \dots, n\}$. By the type variable condition, there exists $i \in \{1, \dots, u\}$ such that $\beta_j \in \text{TVars}(\sigma_i)$; since $d_i \in \llbracket \sigma_i \rrbracket_{\theta}^{\mathcal{M}} \cap \llbracket \sigma_i \rrbracket_{\theta'}^{\mathcal{M}}$ and distinct domains are disjoint, we have $\llbracket \sigma_i \rrbracket_{\theta}^{\mathcal{M}} = \llbracket \sigma_i \rrbracket_{\theta'}^{\mathcal{M}}$, which, together with $\beta_j \in \text{TVars}(\sigma_i)$ and the injectivity of the type constructors, implies $\theta(\beta_j) = \theta'(\beta_j)$, hence $E_j = E'_j$; and since j was arbitrary, we obtain $\bar{E} = \bar{E}'$, as desired. \square

Theorem 4.7 (Correctness of $\mathfrak{g}@$). *The cover-based type guards encoding $\mathfrak{g}@$ is correct.*

Proof. Let $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ be the signature of a polymorphic problem Φ .

SOUND: Let $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ be a model of Φ . By Lemma 2.12, we may assume that \mathbb{D} is disjoint from each of its elements, its elements are mutually disjoint, and the type constructors $k^{\mathcal{M}}$ satisfy distinctness and injectivity.

We define a structure $\mathcal{M}' = (D', (f^{\mathcal{M}'})_{f \in \mathcal{F}' \uplus \mathcal{K}'}, (p^{\mathcal{M}'})_{p \in \mathcal{P}' \uplus \{\mathfrak{g}^2\}})$ for the untyped signature $\Sigma' = (\mathcal{F}' \uplus \mathcal{K}', \mathcal{P}' \uplus \{\mathfrak{g}^2\})$ as follows. D' is the union of the domains of \mathcal{M} and their elements, $\mathfrak{g}^{\mathcal{M}'}$ is the set membership relation, and the symbols of \mathcal{M}' common to those of \mathcal{M} try to emulate the \mathcal{M} interpretation as closely as possible: $k^{\mathcal{M}'}$ acts like $k^{\mathcal{M}}$ on the \mathbb{D} subset of D' , and similarly for function and predicate symbols s whose noninferable type arguments in \mathcal{M} become regular (term) arguments in \mathcal{M}' . The reason why we can omit the inferable type arguments of s is their recoverability from the type of s and the domains of the genuine term arguments. Additionally, when applying $s^{\mathcal{M}'}$ to inputs outside s 's cover that do not belong to their proper domains in \mathcal{M} , we correct these inputs by replacing them with arbitrary inputs from the proper domains. We do this because, in formulae, these inputs will not be guarded by the encoding, but we will nevertheless want to infer the encoded formula holding in \mathcal{M}' from the original formula holding in \mathcal{M} .

Formally, we define the components of \mathcal{M}' as follows. First, $D' = \mathbb{D} \uplus (\bigcup_{D \in \mathbb{D}} D)$ and $\mathfrak{g}^{\mathcal{M}'}(a, b) = (a \in \mathbb{D} \wedge b \in a)$. Assume $k :: n$ is in \mathcal{K} , meaning that k is an n -ary function symbol in \mathcal{K}' . Then

$$k^{\mathcal{M}'}(\bar{D}) = \begin{cases} k^{\mathcal{M}}(\bar{D}) & \text{if } \bar{D} \in \mathbb{D}^n \\ \varepsilon(\mathbb{D}) & \text{otherwise} \end{cases}$$

Assume $\bar{\alpha} = (\alpha_1, \dots, \alpha_m)$, $\bar{\beta} = (\beta_1, \dots, \beta_n)$, $\bar{\sigma} = (\sigma_1, \dots, \sigma_u)$, $\bar{\tau} = (\tau_1, \dots, \tau_v)$, and $s :: \forall \bar{\alpha}, \bar{\beta}. \bar{\sigma} \times \bar{\tau} \rightarrow \zeta$ is in $\mathcal{F} \uplus \mathcal{P}$, such that the first m type arguments are noninferable, the last n type arguments are inferable, and the first u term arguments constitute s 's cover. (The general case, with arbitrary permutations of noninferable and cover arguments, can be handled similarly, albeit with heavier notation.) Then s is an $(m + u + v)$ -ary symbol in $\mathcal{F}' \uplus \mathcal{P}'$. Let $\bar{D} = (D_1, \dots, D_m) \in D'^m$, $\bar{d} = (d_1, \dots, d_u) \in D'^u$ and $\bar{e} = (e_1, \dots, e_v) \in D'^v$, and consider the following condition on domains:

$$(D_1, \dots, D_m) \in \mathbb{D}^m \text{ and there exists } \bar{E} = (E_1, \dots, E_n) \in \mathbb{D}^n \text{ such that each } d_i \text{ is in } \llbracket \sigma_i \rrbracket_{\theta}^{\mathcal{M}}, \\ \text{where } \theta \text{ maps } (\bar{\alpha}, \bar{\beta}) \text{ to } (\bar{D}, \bar{E})$$

Assuming the condition holds, by Lemma 4.6 there exists precisely one tuple \bar{E} satisfying it. We let $\bar{e}' = (e'_1, \dots, e'_v)$, where

$$e'_i = \begin{cases} e_i & \text{if } e_i \in \llbracket \tau_i \rrbracket_{\theta}^{\mathcal{M}} \\ \varepsilon(\llbracket \tau_i \rrbracket_{\theta}^{\mathcal{M}}) & \text{otherwise} \end{cases}$$

Finally, we define

$$s^{\mathcal{M}'}(\bar{D}, \bar{d}, \bar{e}) = \begin{cases} s^{\mathcal{M}}(\bar{D}, \bar{E})(\bar{d}, \bar{e}') & \text{if the domain condition holds} \\ \varepsilon(\mathbb{D}) & \text{if the domain condition fails and } s \in \mathcal{F} \\ \varepsilon(o) & \text{if the domain condition fails and } s \in \mathcal{P} \end{cases}$$

We will show that \mathcal{M}' is a model of $\llbracket \Phi \rrbracket_{\mathfrak{g}@; \mathbf{a}^{\text{int}}; \mathbf{e}}$. Let $\xi' : \mathcal{V} \rightarrow D'$ be a valuation that respects types, in the sense that $\xi'(\mathcal{V}(\alpha)) \in \mathbb{D}$ for all $\alpha \in \mathcal{A}$. We define $\theta : \mathcal{A} \rightarrow \mathbb{D}$ by $\theta(\alpha) = \xi'(\mathcal{V}(\alpha))$ and $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}_{\Sigma}} \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ by $\xi(X)(\sigma) = \xi'(X)$ if $\xi'(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ and

$= \varepsilon(\llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}})$ otherwise. Facts (1)–(3) below follow by induction on σ , t or φ (for arbitrary ξ'), and (1') is a consequence of (1) and the definition of $\mathfrak{g}^{\mathcal{M}'}$:

- (1) $\llbracket \langle \sigma \rangle \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$;
- (1') $\llbracket \mathfrak{g}(\langle \sigma \rangle, X) \rrbracket_{\xi'}^{\mathcal{M}'} = (\xi'(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}})$;
- (2) if $t \notin \mathcal{V}_{\text{typed}}$ and $\xi(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ for all $X^{\sigma} \in \text{UV}(t)$, then $\llbracket \langle t \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket t \rrbracket_{\theta, \xi}^{\mathcal{M}}$;
- (3) if $\xi'(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ for all $X^{\sigma} \in \text{UV}(\varphi)$, then $\llbracket \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}}$ implies $\llbracket \langle \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'}$.

Let us detail a few interesting cases in these proofs:

INDUCTIVE CASE FOR (2): Assume $t = f(\bar{\tau})(t_1, \dots, t_n)$, and let $i \in \{1, \dots, n\}$. If $t_i \notin \mathcal{V}_{\text{typed}}$, then the induction hypothesis applies to it, yielding

$$\llbracket \langle t_i \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket t_i \rrbracket_{\theta, \xi}^{\mathcal{M}}$$

If $t_i = X^{\sigma}$, then $\llbracket \langle t_i \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} = \xi'(X)$, $\llbracket t_i \rrbracket_{\theta, \xi}^{\mathcal{M}} = \xi(X)(\sigma)$, and we have two cases:

- If $\xi'(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$, then $\xi'(X)(\sigma) = \xi(X)$.
- Otherwise, by the assumptions $i \notin \text{Cover}_f$, and hence the definition of $f^{\mathcal{M}'}$ effectively replaces the argument $\xi'(X)$ by $\varepsilon(\llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}})$, which also equals $\xi(X)(\sigma)$.

The above shows $\llbracket \langle t \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket t \rrbracket_{\theta, \xi}^{\mathcal{M}}$.

DISEQUALITY CASE FOR (3): The subcase when one of the terms is a variable and the other is not. Assume φ has the form $X^{\sigma} \not\approx t$. If $\xi'(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$, then $\xi'(X)(\sigma) = \xi(X)$, hence $\llbracket \langle \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}}$. Otherwise, $\llbracket \langle \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'}$ since $\llbracket \langle t \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ (and hence is different from $\xi'(X)$).

UNIVERSAL QUANTIFIER CASE FOR (3): Assume (A) $\llbracket \forall X : \sigma. \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}}$. We must show $\llbracket \langle \forall X : \sigma. \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'}$. Fix $d \in D$ and let $\xi'_1 = \xi'[X \mapsto d]$.

- Assume $X \in \text{UV}(\varphi)$. Then we assume $\llbracket \mathfrak{g}(\langle \sigma \rangle, X) \rrbracket_{\xi'}^{\mathcal{M}'}$, i.e. (B) $d \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$, and need to show $\llbracket \langle \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'_1}^{\mathcal{M}'}$. From (B), we have $\xi(X)(\sigma) = \xi'(X)$, hence $\xi_1 = \xi[X \mapsto d]$; with (A), this implies $\llbracket \varphi \rrbracket_{\theta, \xi_1}^{\mathcal{M}}$, and hence the desired fact follows from the induction hypothesis.
- Assume $X \notin \text{UV}(\varphi)$. We must show $\llbracket \langle \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'_1}^{\mathcal{M}'}$. If $d \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$, then the argument goes the same as above. So assume $d \notin \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$. Then $\xi(X)(\sigma) = \varepsilon(\llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}})$, hence $\xi_1 = \xi[X \mapsto \varepsilon(\llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}})]$; with (A), this implies $\llbracket \varphi \rrbracket_{\theta, \xi_1}^{\mathcal{M}}$, and hence the desired fact follows from the induction hypothesis.

From (1') and the definition of $f^{\mathcal{M}'}$, it follows that \mathcal{M}' satisfies the necessary axioms, namely, $\llbracket \varphi \rrbracket_{\text{a}^{\text{ninf}}; \mathbf{e}}$ for each typing axiom φ . It remains to show that \mathcal{M}' satisfies $\llbracket \langle \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket$ for each $\varphi \in \Phi$. So let $\varphi \in \Phi$. Then $\llbracket \varphi \rrbracket_{\theta}^{\mathcal{M}}$. We pick any ξ' that respects types and has the property that $\xi'(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ for all $X^{\sigma} \in \text{UV}(\varphi)$; by (3) and the fact that φ and $\llbracket \langle \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket$ are sentences, it follows that $\llbracket \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}}$, hence $\llbracket \langle \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'}$, hence $\llbracket \langle \varphi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'}$.

COMPLETE: This part is easy. By Theorem 3.11, it suffices to show that $\llbracket \langle \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}} \rrbracket$ is complete. Let $\mathcal{M}' = (\mathbb{D}', (k^{\mathcal{M}'})_{k \in \mathcal{K}}, (f^{\mathcal{M}'})_{f \in \mathcal{F}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}_{\mathbb{W}}(\mathfrak{g})})$ be a model of $\langle \Phi \rangle \mathfrak{g}_{\text{a}^{\text{ninf}}; \mathbf{e}}$, for which again we

may assume that the domains are mutually disjoint. We define $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ by restricting the domains according to the guards and restricting the operations correspondingly. Namely, for each $D \in \mathbb{D}'$, let $D' = \{d \in D \mid g^{\mathcal{M}'}(D)(d)\}$. We take

- $\mathbb{D} = \{D' \mid D \in \mathbb{D}'\}$;
- $k^{\mathcal{M}}(\bar{D}') = k^{\mathcal{M}'}(\bar{D})$;
- $s^{\mathcal{M}}(\bar{D}')(\bar{d}) = s^{\mathcal{M}'}(\bar{D})(\bar{d})$.

Thanks to the typing axioms, each D' is nonempty. Since the domains are disjoint, each D' determines its D uniquely; together with the typing axioms, this also ensures that each $k^{\mathcal{M}}$ and $s^{\mathcal{M}}$ are well defined. Now, $\llbracket U \rrbracket_{\theta, \xi}^{\mathcal{M}} = \llbracket \langle U \rangle \mathfrak{g}_{\theta, \xi} \rrbracket_{\theta, \xi}^{\mathcal{M}'}$, where U is first a term and then a formula, follows by induction on U (for arbitrary θ and ξ). From this, we obtain that \mathcal{M} is a model of $\langle \Phi \rangle \mathfrak{g}_{\theta}$ by the usual route. \square

4.2. Cover-Based Type Tags. The cover-based encoding $\mathfrak{t}_{\@}$ is similar to the traditional encoding \mathfrak{t} , except that it tags only undercover occurrences of variables and requires typing axioms to add or remove tags around function terms.

Definition 4.8 (Cover Tags $\mathfrak{t}_{\@}$). The encoding $\langle \rangle \mathfrak{t}_{\@}$ translates polymorphic problems over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ to polymorphic problems over $(\mathcal{K}, \mathcal{F} \uplus \{\mathfrak{t} : \forall \alpha. \alpha \rightarrow \alpha\}, \mathcal{P})$. Its term and formula translations are the following:

$$\begin{aligned} \langle f(\bar{\sigma})(\bar{t}) \rangle \mathfrak{t}_{\@} &= f\langle \bar{\sigma} \rangle(\llbracket \langle \bar{t} \rangle \mathfrak{t}_{\@} \rrbracket_f) \\ \langle p(\bar{\sigma})(\bar{t}) \rangle \mathfrak{t}_{\@} &= p\langle \bar{\sigma} \rangle(\llbracket \langle \bar{t} \rangle \mathfrak{t}_{\@} \rrbracket_p) & \langle t_1 \approx t_2 \rangle \mathfrak{t}_{\@} &= \llbracket \langle t_1 \rangle \mathfrak{t}_{\@} \rrbracket_{\approx} \approx \llbracket \langle t_2 \rangle \mathfrak{t}_{\@} \rrbracket_{\approx} \\ \langle \neg p(\bar{\sigma})(\bar{t}) \rangle \mathfrak{t}_{\@} &= \neg p\langle \bar{\sigma} \rangle(\llbracket \langle \bar{t} \rangle \mathfrak{t}_{\@} \rrbracket_p) & \langle \exists X : \sigma. \varphi \rangle \mathfrak{t}_{\@} &= \exists X : \sigma. \mathfrak{t}\langle \sigma \rangle(X) \approx X \wedge \langle \varphi \rangle \mathfrak{t}_{\@} \end{aligned}$$

The auxiliary function $\llbracket (t_1^{\sigma_1}, \dots, t_n^{\sigma_n}) \rrbracket_s$ returns a vector (u_1, \dots, u_n) such that

$$u_j = \begin{cases} \mathfrak{t}\langle \sigma_j \rangle(t_j) & \text{if } j \in \text{Cover}_s \text{ and } t_j \in \mathcal{V}_{\text{typed}}^{\forall} \\ t_j & \text{otherwise} \end{cases}$$

taking $\text{Cover}_{\approx} = \{1, 2\}$. The encoding adds the following typing axioms:

$$\begin{aligned} \forall \bar{\alpha}. \forall \bar{X} : \bar{\sigma}. \mathfrak{t}\langle \sigma \rangle(f\langle \bar{\alpha} \rangle(\llbracket \langle \bar{X} \rangle \mathfrak{t}_{\@} \rrbracket_f)) &\approx f\langle \bar{\alpha} \rangle(\llbracket \langle \bar{X} \rangle \mathfrak{t}_{\@} \rrbracket_f) \quad \text{for } f : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \sigma \in \mathcal{F} \\ \forall \alpha. \exists X : \alpha. \mathfrak{t}\langle \alpha \rangle(X) &\approx X \end{aligned}$$

The *cover-based type tags* encoding \mathfrak{t} is the composition $\langle \rangle \mathfrak{t}_{\@}; \mathfrak{a}^{\text{ninf}; \mathfrak{e}}$. It translates a polymorphic problem over Σ into an untyped problem over $\Sigma' = (\mathcal{F}' \uplus \mathcal{K}' \uplus \{\mathfrak{t}^2\}, \mathcal{P}')$, where \mathcal{K}' , \mathcal{F}' , \mathcal{P}' are as for $\mathfrak{a}^{\text{ninf}}$.

Example 4.9. The $\mathfrak{t}_{\@}$ encoding of Example 2.10 is as follows (again, choosing the cover for cons as in Convention 4.2):

$$\begin{aligned} \forall A. \mathfrak{t}(\text{list}(A), \text{nil}(A)) &\approx \text{nil}(A) \\ \forall A, X, Xs. \mathfrak{t}(\text{list}(A), \text{cons}(t(A, X), Xs)) &\approx \text{cons}(\mathfrak{t}(A, X), Xs) \\ \forall A, Xs. \mathfrak{t}(\text{list}(A), \text{hd}(\mathfrak{t}(\text{list}(A), Xs))) &\approx \text{hd}(\mathfrak{t}(\text{list}(A), Xs)) \\ \forall A, Xs. \mathfrak{t}(A, \text{tl}(\mathfrak{t}(\text{list}(A), Xs))) &\approx \text{tl}(\mathfrak{t}(\text{list}(A), Xs)) \\ \forall A, X, Xs. \text{nil}(A) &\not\approx \text{cons}(\mathfrak{t}(A, X), Xs) \\ \forall A, Xs. \mathfrak{t}(\text{list}(A), Xs) &\approx \text{nil}(A) \vee \\ &(\exists Y, Ys. \mathfrak{t}(A, Y) \approx Y \wedge \mathfrak{t}(\text{list}(A), Ys) \approx Ys \wedge \mathfrak{t}(\text{list}(A), Xs) \approx \text{cons}(Y, Ys)) \\ \forall A, X, Xs. \text{hd}(\text{cons}(\mathfrak{t}(A, X), Xs)) &\approx \mathfrak{t}(A, X) \wedge \text{tl}(\text{cons}(\mathfrak{t}(A, X), Xs)) \approx \mathfrak{t}(\text{list}(A), Xs) \end{aligned}$$

$$\begin{aligned} \exists X, Y, Xs, Ys. \mathfrak{t}(b, X) \approx X \wedge \mathfrak{t}(b, Y) \approx Y \wedge \mathfrak{t}(\text{list}(b), Xs) \approx Xs \wedge \mathfrak{t}(\text{list}(b), Ys) \approx Ys \wedge \\ \text{cons}(X, Xs) \approx \text{cons}(Y, Ys) \wedge (X \not\approx Y \vee Xs \not\approx Ys) \end{aligned}$$

Theorem 4.10 (Correctness of $\mathfrak{t}@$). *The cover-based type tags encoding $\mathfrak{t}@$ is correct.*

Proof. Let $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ be the signature of a polymorphic problem Φ .

SOUND: Let $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ be a model of Φ . By Lemma 2.12, we may assume that \mathbb{D} is disjoint from each of its elements, its elements are mutually disjoint, and the type constructors $k^{\mathcal{M}}$ satisfy distinctness and injectivity.

We define a structure $\mathcal{M}' = (\mathbb{D}', (f^{\mathcal{M}'})_{f \in \mathcal{F}' \uplus \mathcal{K}' \uplus \{\mathfrak{t}\}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}'})$ for $\Sigma' = (\mathcal{F}' \uplus \mathcal{K}', \mathcal{P}' \uplus \{\mathfrak{g}^2\})$ intended to be a model of $\langle\langle \Phi \rangle\rangle_{\mathfrak{t}@; \mathbf{a}^{\text{ninf}}; \mathbf{e}}$. The construction of \mathcal{M}' proceeds very similarly to the case of guards from the proof of Theorem 4.7: \mathbb{D}' , $k^{\mathcal{M}'}$ for $k \in \mathcal{K}'$, and $s^{\mathcal{M}'}$ for $s \in \mathcal{F}' \uplus \mathcal{P}'$ are all defined in the same way as in that proof. It remains to define $\mathfrak{t}^{\mathcal{M}'} : \mathbb{D}' \times \mathbb{D}' \rightarrow \mathbb{D}'$:

$$\mathfrak{t}^{\mathcal{M}'}(a, b) = \begin{cases} b & \text{if } a \in \mathbb{D} \text{ and } b \in a \\ \varepsilon(a) & \text{if } a \in \mathbb{D} \text{ and } b \notin a \\ \varepsilon(\mathbb{D}) & \text{if } a \notin \mathbb{D} \end{cases}$$

In the proof of Theorem 4.7, $\mathfrak{g}^{\mathcal{M}'}(a, b)$ captured the notion that a is a domain of \mathcal{M} and b an element in it. The same can now be expressed by $\mathfrak{t}^{\mathcal{M}'}(a, b) = b$. Additionally, here $\mathfrak{t}^{\mathcal{M}'}$ is useful for redirecting any ill-typed element b (one not in the first argument a) to a well-typed element $\varepsilon(a) \in a$.

To show that \mathcal{M}' is a model of $\langle\langle \Phi \rangle\rangle_{\mathfrak{t}@; \mathbf{a}^{\text{ninf}}; \mathbf{e}}$, we proceed almost identically to the proof of Theorem 4.7. Starting with a valuation $\xi' : \mathcal{V} \rightarrow \mathbb{D}'$ that respects types, we define ξ and θ in the same way and state facts (1)–(3) as for guards, but adding a condition about the “well-typedness” of the existential variables interpretation, since only universal variables (i.e. variables in $\mathcal{V}_{\text{typed}}^{\forall}$) are tagged by $\mathfrak{t}@$:

- (1) $\llbracket \langle\langle \sigma \rangle\rangle \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$;
- (2) if $t \notin \mathcal{V}_{\text{typed}}$ and $\xi(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ for all $X^{\sigma} \in \text{UV}(t) \cup \mathcal{V}_{\text{typed}}^{\exists}$, then $\llbracket \langle\langle t \rangle\rangle \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket t \rrbracket_{\theta, \xi}^{\mathcal{M}}$;
- (3) if $\xi'(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ for all $X^{\sigma} \in \text{UV}(\varphi) \cup \mathcal{V}_{\text{typed}}^{\exists}$, then $\llbracket \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}}$ implies $\llbracket \langle\langle \varphi \rangle\rangle \rrbracket_{\xi'}^{\mathcal{M}'}$.

Moreover, (1') from the proof of Theorem 4.7 is replaced with the following two facts:

- (1') $(\llbracket \mathfrak{t}(\langle\langle \sigma \rangle\rangle, X) \rrbracket_{\xi'}^{\mathcal{M}'} = \xi'(X)) = (\xi'(X) \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}})$;
- (1'') $\llbracket \mathfrak{t}(\langle\langle \sigma \rangle\rangle, X) \rrbracket_{\xi'}^{\mathcal{M}'} \in \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$.

Again, (1)–(3) follow by induction on the involved type, term, or formula; (1') and (1'') follow from (1) and the definition of $\mathfrak{t}^{\mathcal{M}'}$ (and are in turn used to establish (2) and (3)). The fact that \mathcal{M}' satisfies $\langle\langle \varphi \rangle\rangle_{\mathfrak{t}@; \mathbf{a}^{\text{ninf}}; \mathbf{e}}$ for each $\varphi \in \Phi$ follows in the same way as in Theorem 4.7. It remains to show that \mathcal{M}' satisfies the necessary axioms, namely, the $\langle\langle \rangle\rangle_{\mathbf{a}^{\text{ninf}}; \mathbf{e}}$ -translations of the $\langle\langle \rangle\rangle_{\mathfrak{t}@}$ typing axioms. This follows from (1') and the definition of $f^{\mathcal{M}'}$.

COMPLETE: This part is again similar to the corresponding one for guards. By Theorem 3.11, it suffices to show $\langle\langle \rangle\rangle_{\mathfrak{t}@}$ complete. Let $\mathcal{M}' = (\mathbb{D}', (k^{\mathcal{M}'})_{k \in \mathcal{K}}, (f^{\mathcal{M}'})_{f \in \mathcal{F} \uplus \{\mathfrak{t}\}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}})$ be a model of $\langle\langle \Phi \rangle\rangle_{\mathfrak{t}@}$, for which we may assume that the domains are mutually disjoint. We define $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ by restricting the domains not according to the guards as in the proof of Theorem 4.7, but according to the property that tags be identity: for each $D \in \mathbb{D}'$, let $D' = \{d \in D \mid \mathfrak{t}^{\mathcal{M}'}(D)(d) = d\}$; we define $\mathbb{D} = \{D' \mid D \in \mathbb{D}'\}$. The

other components of \mathcal{M} are defined as in Theorem 3.11, and the proof is analogous to that for guards. \square

Unlike in the case of guards, the traditional type tag encoding \mathfrak{t} is not a particular case of the cover-based encoding $\mathfrak{t}@$, since only $\mathfrak{t}@$ introduces typing axioms, and also \mathfrak{t} further restricts the type arguments to phantom arguments. Nevertheless, we can provide a somewhat similar argument for \mathfrak{t} 's correctness. Although the encoding is well known, we are not aware of any soundness proof in the literature.

Proof of Theorem 3.16 (Correctness of \mathfrak{t}).

SOUND: Let \mathcal{M} be as in the proof of Theorem 4.10. We define \mathcal{M}' similar to there, but with the following difference: D' contains not only the domains in \mathbb{D} and their elements, but also the set P of “polymorphic values”, i.e. functions that take a domain $D \in \mathbb{D}$ and return an element of D . The reason for the polymorphic values is that, due to the switch from noninferable arguments to phantom arguments, the domain of the result for $f^{\mathcal{M}'}$ with $f : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \sigma$ in \mathcal{F} will no longer be completely inferable from the arguments, but will miss precisely the nonphantom noninferable arguments, which correspond to type variables belonging to the result type σ but not to the argument types $\bar{\sigma}$. Consequently, the result of applying $f^{\mathcal{M}'}$ to “well-typed” arguments will be polymorphic values that wait for the result domain D and, if it has the form $[[\sigma]]_\theta$ for an appropriate θ , return the result from $[[\sigma]]_\theta$ according to the interpretation of $f^{\mathcal{M}'}$. Then the interpretation of the tag applied to polymorphic values will select the desired element by providing the domain. The reason why this approach works with \mathfrak{t} but not with $\mathfrak{t}@$ is that in \mathfrak{t} tags are applied everywhere, thus resolving immediately any polymorphic value emerging from an application of $f^{\mathcal{M}'}$.

Formally, we define the components of \mathcal{M}' as follows. First, $k^{\mathcal{M}'}$ and $p^{\mathcal{M}'}$ are as in the proofs of Theorems 4.7 and 4.10. (For predicate symbols, the notions of phantom and noninferable type arguments coincide.) $D' = \mathbb{D} \cup E \cup P$, where $E = \bigcup_{D \in \mathbb{D}} D$ (the elements) and $P = \prod_{D \in \mathbb{D}} D$ (the polymorphic values). Moreover, $\mathfrak{t}^{\mathcal{M}'} : D' \times D' \rightarrow D'$ is defined as follows:

$$\mathfrak{t}^{\mathcal{M}'}(a, b) = \begin{cases} b & \text{if } a \in \mathbb{D} \text{ and } b \in a \\ \varepsilon(a) & \text{if } a \in \mathbb{D}, b \in E \text{ and } b \notin a \\ b(a) & \text{if } a \in \mathbb{D} \text{ and } b \in P \\ \varepsilon(\mathbb{D}) & \text{otherwise} \end{cases}$$

Assume $\bar{\alpha} = (\alpha_1, \dots, \alpha_m)$, $\bar{\gamma} = (\gamma_1, \dots, \gamma_r)$, $\bar{\beta} = (\beta_1, \dots, \beta_n)$, $\bar{\sigma} = (\sigma_1, \dots, \sigma_u)$, $\bar{\tau} = (\tau_1, \dots, \tau_v)$, and $s : \forall (\bar{\alpha}, \bar{\gamma}, \bar{\beta}). (\bar{\sigma}, \bar{\tau}) \rightarrow \varsigma$ is in $\mathcal{F} \uplus \mathcal{P}$, such that the first m type arguments are phantom, the middle r arguments are noninferable nonphantoms, the last n type arguments are inferable, and the first u term arguments constitute s 's cover. (Again, the general case with arbitrary permutations of the arguments can be handled similarly.) Then s is an $(m + u + v)$ -ary symbol in $\mathcal{F}' \uplus \mathcal{P}'$. Let $\bar{D} = (D_1, \dots, D_m) \in D'^m$, $\bar{d} = (d_1, \dots, d_u) \in D'^u$ and $\bar{e} = (e_1, \dots, e_v) \in D'^v$, and consider the following condition on domains:

$$(D_1, \dots, D_m) \in \mathbb{D}^m \text{ and there exists } \bar{E} = (E_1, \dots, E_n) \in \mathbb{D}^n \text{ such that each } d_i \text{ is in } [[\sigma_i]]_\theta^{\mathcal{M}'}, \\ \text{where } \theta \text{ maps } (\bar{\alpha}, \bar{\beta}) \text{ to } (\bar{D}, \bar{E})$$

Assuming the condition holds, again by Lemma 4.6 (taking $\bar{\alpha}$ from there to be $(\bar{\alpha}, \bar{\gamma})$) there exists precisely one tuple \bar{E} satisfying it, and we define the “correction” vector \bar{e}' from \bar{e} as

in the proof of Theorems 4.7 and 4.10. We now define

$$f^{\mathcal{M}'}(\bar{D}, \bar{d}, \bar{e}) = \begin{cases} \pi_{\bar{D}, \bar{d}, \bar{e}} & \text{if the domain condition holds} \\ \varepsilon(\mathbb{D}) & \text{otherwise} \end{cases}$$

where the polymorphic value $\pi_{\bar{D}, \bar{d}, \bar{e}} \in P$ is defined as follows:

$$\pi_{\bar{D}, \bar{d}, \bar{e}}(D) = \begin{cases} s^{\mathcal{M}}(\bar{D}, \bar{G}, \bar{E})(\bar{d}, \bar{e}') & \text{if } D \text{ has the form } \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}} \text{ and } \theta \text{ assigns } (\bar{\alpha}, \bar{\gamma}, \bar{\beta}) \text{ to } (\bar{D}, \bar{G}, \bar{E}) \\ \varepsilon(\mathbb{D}) & \text{otherwise} \end{cases}$$

To show that \mathcal{M}' is a model of $\langle\langle \Phi \rangle\rangle_{\mathfrak{t}; \mathbf{a}^{\text{phn}}; \mathbf{e}}$, we proceed almost identically to the proof of Theorem 4.10. Starting with a valuation $\xi' : \mathcal{V} \rightarrow D'$ that respects types (as in the proofs of Theorems 4.7 and 4.10). We define $\theta : \mathcal{A} \rightarrow \mathbb{D}$ by $\theta(\alpha) = \xi'(\mathcal{V}(\alpha))$ and $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}_{\Sigma}} \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ by $\xi(X)(\sigma) = \mathfrak{t}^{\mathcal{M}'}(\llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}, \xi'(X))$. Facts (1)–(3) below follow by induction on σ , t , or φ (for arbitrary ξ'):

- (1) $\llbracket \langle\langle \sigma \rangle\rangle \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$;
- (2) $\llbracket \langle\langle t \rangle\rangle_{\mathfrak{t}; \mathbf{a}^{\text{phn}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket t \rrbracket_{\theta, \xi}^{\mathcal{M}}$;
- (3) $\llbracket \langle\langle \varphi \rangle\rangle_{\theta, \xi}^{\mathcal{M}} = \llbracket \langle\langle \varphi \rangle\rangle_{\mathfrak{t}; \mathbf{a}^{\text{phn}}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'}$.

It follows by the usual route that \mathcal{M}' is a model of each $\langle\langle \varphi \rangle\rangle_{\mathfrak{t}; \mathbf{a}^{\text{phn}}; \mathbf{e}}$ with $\varphi \in \Phi$, hence of $\langle\langle \Phi \rangle\rangle_{\mathfrak{t}; \mathbf{a}^{\text{phn}}; \mathbf{e}}$.

COMPLETE: By Theorem 3.11, it suffices to show $\langle\langle \rangle\rangle_{\mathfrak{t}}$ complete. Let $\mathcal{M}' = (\mathbb{D}', (k^{\mathcal{M}'})_{k \in \mathcal{K}}, (f^{\mathcal{M}'})_{f \in \mathcal{F} \cup \{\mathfrak{t}\}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}})$ be a model of $\langle\langle \Phi \rangle\rangle_{\mathfrak{t}}$, for which we may assume that the domains are mutually disjoint. We must construct a model $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ of Φ by somehow removing the tags from \mathcal{M}' . Unlike for $\mathfrak{t}@$, here $\langle\langle \rangle\rangle_{\mathfrak{t}}$ does not contain typing axioms ensuring that a the restriction of \mathcal{M}' to elements for which the tag interpretation is the identity forms a valid submodel; thus \mathcal{M}' cannot be defined that way. On the other hand, we know that the formulae in $\langle\langle \rangle\rangle_{\mathfrak{t}}$ are fully tagged; in particular, all variables are accessed through tags. This means we can take the domains of \mathcal{M} by restricting those of \mathcal{M}' to the images of the tag interpretations. The result will not be a submodel, so cannot take the functions and predicates $s^{\mathcal{M}}$ of \mathcal{M} to be restrictions of those of \mathcal{M}' . Instead, $s^{\mathcal{M}}$ will be defined by applying $s^{\mathcal{M}'}$ through the tag “interface”.

Formally, let, for each $D \in \mathbb{D}'$, let D' be the image of $\mathfrak{t}^{\mathcal{M}'}(D)$ (which is a subset of D). Each D' uniquely determines its D . We define the auxiliary function $T : \prod_{D \in \mathbb{D}'} D \rightarrow D$, which only applies $\mathfrak{t}^{\mathcal{M}'}$ if the element is not already in the image:

$$T_D(d) = \begin{cases} d & \text{if } d \in D' \\ \mathfrak{t}^{\mathcal{M}'}(D)(d) & \text{otherwise} \end{cases}$$

Notice that T_D adds no tags around an existing tag: $T_D(\mathfrak{t}^{\mathcal{M}'}(D)(d)) = \mathfrak{t}^{\mathcal{M}'}(D)(d)$. We define the components of \mathcal{M} as follows: $\mathbb{D} = \{D' \mid D \in \mathbb{D}'\}$; $k^{\mathcal{M}}(\bar{D}') = k^{\mathcal{M}'}(\bar{D})$; if $f : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \sigma$ is in \mathcal{F} , then $f^{\mathcal{M}}(\bar{D}')(d_1, \dots, d_n) = \mathfrak{t}^{\mathcal{M}'}(\llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}'}) (f^{\mathcal{M}'}(\bar{D})(T_{\llbracket \sigma_1 \rrbracket_{\theta}^{\mathcal{M}'}}(d_1), \dots, T_{\llbracket \sigma_n \rrbracket_{\theta}^{\mathcal{M}'}}(d_n)))$, where θ maps each α_i to D_i ; if $p : \forall \bar{\alpha}. \bar{\sigma} \rightarrow o$ is in \mathcal{P} , then $p^{\mathcal{M}}(\bar{D}')(d_1, \dots, d_n) = p^{\mathcal{M}'}(\bar{D})(T_{\llbracket \sigma_1 \rrbracket_{\theta}^{\mathcal{M}'}}(d_1), \dots, T_{\llbracket \sigma_n \rrbracket_{\theta}^{\mathcal{M}'}}(d_n))$, where θ maps each α_i to D_i . The interpretation $f^{\mathcal{M}}$ applies tags $\mathfrak{t}^{\mathcal{M}'}$ at the top, whereas the interpretations of $f^{\mathcal{M}}$ and $p^{\mathcal{M}}$ apply the modified tag T at the bottom. This is to ensure that the interpretations of terms and atomic formulae in \mathcal{M} do not add several consecutive layers of tags. For example, when interpreted in \mathcal{M}

within $f(X^\sigma)$, f should add a tag around X^σ , as well as one around $f(X^\sigma)$; however, when interpreted in \mathcal{M} within $f(g(X^\sigma))$, f should not add a tag around $g(X^\sigma)$, since g already adds one.

Given $\theta : \mathcal{A} \rightarrow \mathbb{D}$ and $\xi' : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}_\Sigma} \llbracket \sigma \rrbracket_{\theta'}^{\mathcal{M}'}$, we define $\theta' : \mathcal{A} \rightarrow \mathbb{D}'$ by letting $\theta'(\alpha)$ be the unique $D \in \mathbb{D}'$ such that $D' = \theta(\alpha)$, and $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}_\Sigma} \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ by $\xi(X)(\sigma) = \mathfrak{t}^{\mathcal{M}'}(\llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}'})(\xi'(X)(\sigma))$. The next facts follow by induction on σ , t , or φ (for arbitrary θ and ξ'):

- (1) $\llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ is in the image of $\mathfrak{t}^{\mathcal{M}'}(\llbracket \sigma \rrbracket_{\theta'}^{\mathcal{M}'})$;
- (2) $\llbracket t \rrbracket_{\theta, \xi}^{\mathcal{M}} = \llbracket \langle t \rangle \rrbracket_{\theta', \xi'}^{\mathcal{M}'}$;
- (3) $\llbracket \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}} = \llbracket \langle \varphi \rangle \rrbracket_{\theta', \xi'}^{\mathcal{M}'}$.

Let $\varphi \in \Phi$. To show that \mathcal{M} is a model of φ , let $\theta : \mathcal{A} \rightarrow \mathbb{D}$ and $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}_\Sigma} \llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$. By (1), there exists $\xi' : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}_\Sigma} \llbracket \sigma \rrbracket_{\theta'}^{\mathcal{M}'}$ such that $\xi(X)(\sigma) = \mathfrak{t}^{\mathcal{M}'}(\llbracket \sigma \rrbracket_{\theta'}^{\mathcal{M}'})(\xi'(X)(\sigma))$ for all X and σ —simply take $\xi'(X)(\sigma)$ to be any element d' in $\llbracket \sigma \rrbracket_{\theta'}^{\mathcal{M}'}$ such that $\mathfrak{t}^{\mathcal{M}'}(\llbracket \sigma \rrbracket_{\theta'}^{\mathcal{M}'})(d') = \xi(X)(\sigma)$. Since $\llbracket \langle \varphi \rangle \rrbracket_{\theta', \xi'}^{\mathcal{M}'}$ holds, it follows from (3) that $\llbracket \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}}$, i.e. $\llbracket \varphi \rrbracket^{\mathcal{M}}$, also holds, as desired. \square

4.3. Polymorphic Löwenheim–Skolem. The previous completeness results can be used to prove Löwenheim–Skolem-like properties for polymorphic first-order logic by appealing to the more manageable (and better known) monomorphic first-order logic.

Lemma 4.11. *If a polymorphic Σ -problem Φ has a model, it also has a model $\mathcal{M} = (\mathbb{D}, (k^{\mathcal{M}})_{k \in \mathcal{K}}, _ , _)$ such that the following conditions are met:¹*

- (1) each $D \in \mathbb{D}$ is countable;
- (2) each $k^{\mathcal{M}}$ is injective, and $k^{\mathcal{M}}(\bar{D}) \neq k^{\mathcal{M}}(\bar{E})$ whenever $k \neq k'$;
- (3) $\mathbb{D} = \{\llbracket \tau \rrbracket^{\mathcal{M}} \mid \tau \in \text{GType}_\Sigma\}$;
- (4) the type interpretation function $\llbracket \cdot \rrbracket^{\mathcal{M}}$ is a bijection between GType and \mathbb{D} ;
- (5) \mathbb{D} is countable;
- (6) \mathbb{D} is disjoint from each $D \in \mathbb{D}$, and any distinct $D_1, D_2 \in \mathbb{D}$ are disjoint.

Proof. Assume Φ has a model, and recall that by definition the polymorphic signatures that we consider are countable. By the soundness of $\mathfrak{g}@$ (Theorem 4.7), $\langle \Phi \rangle_{\mathfrak{g}@}$ also has a model and its signature is countable; hence by classical Löwenheim–Skolem [19], $\langle \Phi \rangle_{\mathfrak{g}@}$ has a countable model \mathcal{M}' . It follows from the proof of completeness of $\mathfrak{g}@$ that Φ has a model $\mathcal{M} = (\mathbb{D}, _ , _ , _)$ constructed from \mathcal{M}' in such a way that each $D \in \mathbb{D}$ is countable—hence \mathcal{M}' satisfies (1). Now, (2)–(6) follow by applying the same reasoning as in Lemma 2.12 and noticing that all the structure modifications from there do not alter the countability of the domains $D \in \mathbb{D}$. \square

¹These correspond to the conditions from Lemma 2.12 plus countability of each $D \in \mathbb{D}$.

Lemma 4.12. *If Φ has a model $\mathcal{M} = (\mathbb{D}, _, _, _)$ where all $\llbracket \tau \rrbracket^{\mathcal{M}}$ are infinite for all $\tau \in GType$, it also has a model $\mathcal{M}' = (\mathbb{D}', _, _, _)$ where \mathbb{D}' and all $D' \in \mathbb{D}'$ are countably infinite.*

Proof. Let Φ' be Φ extended with axioms $Ax = \{\exists X_1 : \sigma, \dots, X_n : \sigma. \bigwedge_{i < j} X_i \not\approx X_j \mid \sigma \in GType_\Sigma, n \in \mathbb{N}\}$ that state that all ground types are infinite. Then $\mathcal{M} = (\mathbb{D}, _, _, _)$ is a model of Φ' . By Lemma 4.11, Φ' also has a model $\mathcal{M}' = (\mathbb{D}', _, _, _)$ with $\mathbb{D}' = \{\llbracket \tau \rrbracket^{\mathcal{M}'} \mid \tau \in GType_\Sigma\}$ (which is, in particular, countably infinite) and each $D' \in \mathbb{D}'$ at most countably infinite, i.e. due to Ax and the fact that each $D' \in \mathbb{D}$ has the form $\llbracket \tau \rrbracket^{\mathcal{M}'}$, actually countably infinite. \mathcal{M}' is the desired model of Φ . \square

5. MONOTONICITY-BASED TYPE ENCODINGS — THE MONOMORPHIC CASE

The cover-based encodings of Section 4 removed some of the clutter associated with type arguments, which are in general necessary to encode polymorphism soundly. Another family of encodings focuses on the quantifiers and exploit monotonicity. For types that are inferred monotonic, the translation can omit the type information on term variables of these types. Informally, a type is monotonic in a problem when, for any model of that problem, we can increase the size of the domain that interprets the type while preserving satisfiability.

This section focuses on the monomorphic case, where the input problem contains no type variables or polymorphic symbols. This case is interesting in its own right and serves as a stepping stone towards polymorphic monotonicity-based encodings (Section 6).

Before we start, we need to define variants of the traditional \mathbf{e} , \mathbf{t} , and \mathbf{g} encodings that operate on monomorphic problems. Since the monomorphic \mathbf{e} is essentially identical to the polymorphic \mathbf{e} restricted to monomorphic signatures, we use the same notation for both. The monomorphic encodings $\tilde{\mathbf{t}}$ and $\tilde{\mathbf{g}}$ coincide with \mathbf{t} and \mathbf{g} except that the polymorphic function $\mathbf{t}(\sigma)(t)$ and predicate $\mathbf{g}(\sigma)(t)$ are replaced by type-indexed families of unary functions $\mathbf{t}_\sigma(t)$ and predicates $\mathbf{g}_\sigma(t)$, as is customary in the literature [36, §4].

5.1. Monotonicity. The concept of monotonicity used by Claessen et al. [16, §2.2] declares a type τ monotonic for a finite problem Φ if for any model $\mathcal{M} = ((D_\sigma)_{\sigma \in Type}, _, _)$ of Φ such that D_τ is finite, there exists another model $\mathcal{M}' = ((D'_\sigma)_{\sigma \in Type}, _, _)$ of Φ such that $|D'_\tau| = |D_\tau| + 1$ and $|D'_\sigma| = |D_\sigma|$ for all $\sigma \neq \tau$. Their notion, which we call *finite monotonicity*, is designed to ensure that it is possible to produce a model having all types interpreted by countably infinite sets, and finally a model having all types interpreted as the same set, so that type information can be soundly erased. In this article, we take directly the infinite-interpretation property as definition of monotonicity and extend the notion to sets of types:

Definition 5.1 (Monotonicity). Let S be a set of types and Φ be a problem. The set S is *monotonic* in Φ if for all models $\mathcal{M} = ((D_\sigma)_{\sigma \in Type}, _, _)$ of Φ , there exists a model $\mathcal{M}' = ((D'_\sigma)_{\sigma \in Type}, _, _)$ of Φ such that for all types σ , D'_σ is infinite if $\sigma \in S$ and $|D'_\sigma| = |D_\sigma|$ otherwise. A type σ is *monotonic* if $\{\sigma\}$ is monotonic. The problem Φ is *monotonic* if the set $Type$ of all types is monotonic.

Full type erasure is sound for monotonic monomorphic problems. The intuition is that a model of such a problem can be extended into a model where all types are interpreted as sets of the same cardinality, which can be merged to yield an untyped model.

Example 5.2. The monkey village of Example 1.1 is monotonic because any model with finitely many bananas can be extended to a model with infinitely many, and any model with infinitely many bananas and finitely many monkeys can be extended to one where monkeys and bananas have the same infinite cardinality (cf. Example 3.4). However, the type of monkeys is not monotonic on its own, as we argued intuitively in Example 1.1, nor is it finitely monotonic.

Theorem 5.3 (Monotonic Erasure). *Full type erasure is sound for monotonic monomorphic problems.*

Proof. Let Φ be such a problem, let $\Sigma = (\text{Type}, \mathcal{F}, \mathcal{P})$ be its signature, and let $\Sigma' = (\mathcal{F}', \mathcal{P}')$ be the target signature of \mathbf{e} . If Φ is satisfiable, it has a model where all domains are infinite by definition of monotonicity. Since Σ is countable, by the downward Löwenheim–Skolem theorem [19], Φ also has a model where all domains are countably infinite, hence also a model $\mathcal{M} = ((D_\sigma)_{\sigma \in \text{Type}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ where all domains are interpreted as the same set D , i.e. each D_σ is D . We define the Σ' -structure $\mathcal{M}' = (D, (f^{\mathcal{M}'})_{f \in \mathcal{F}'}, (p^{\mathcal{M}'})_{p \in \mathcal{P}'})$ by taking D to be D and each $s^{\mathcal{M}'}$ to be $s^{\mathcal{M}}$. For each $\xi' : \mathcal{V} \rightarrow D$, we define $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D_\sigma$, i.e. $\xi : \mathcal{V} \rightarrow \text{Type} \rightarrow D$, by $\xi(X)(\sigma) = \xi'(X)$. The next facts follow by induction on t and φ (for arbitrary ξ):

- (1) $\llbracket \langle t \rangle \mathbf{e} \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket t \rrbracket_{\xi}^{\mathcal{M}}$;
- (2) $\llbracket \langle \varphi \rangle \mathbf{e} \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket \varphi \rrbracket_{\xi}^{\mathcal{M}}$.

It follows by the usual route that \mathcal{M}' is a model of $\llbracket \Phi \rrbracket_{\mathbf{e}}$. □

Remark 5.4. An alternative to invoking the Löwenheim–Skolem theorem would have been to require countable infinity in the definition of monotonicity. Although it makes no difference in this article, the more general definition would also apply in the presence of uncountable interpreted types (e.g. for the real numbers). The proof of Theorem 5.3 can be adapted to go beyond countable infinity if desired.

It is often convenient to determine the monotonicity of single types separately, viewed as singleton sets. This is enough to make the set of all types, i.e. the problem, monotonic:

Lemma 5.5 (Global Monotonicity from Separate Monotonicity). *If σ is monotonic in Φ for each type σ , then Φ is monotonic.*

Proof. Let $\Sigma = (\text{Type}, \mathcal{F}, \mathcal{P})$ be the signature of Φ , and assume Φ is satisfiable. Let $\sigma_1, \sigma_2, \dots$ be an enumeration of Type , and let φ_k be a formula stating that for all $i \in \{1, \dots, k\}$, σ_i 's domain has at least k elements. Finally, let $\Phi' = \{\varphi_i \mid i \in \mathbb{N}\}$. By the monotonicity of the individual types, it follows by induction on k that each $\Phi \cup \{\varphi_i \mid i \in \{1, \dots, k\}\}$ is satisfiable, and hence that each finite subset of $\Phi \cup \Phi'$ is satisfiable. By the compactness theorem [19], it follows that $\Phi \cup \Phi'$ is itself satisfiable, and hence Φ has a model with only infinite domains. □

5.2. Monotonicity Inference. Claessen et al. introduced a simple calculus to infer finite monotonicity for monomorphic first-order logic [16, §2.3]. The definition below generalises it from clause normal form to negation normal form. The generalisation is straightforward; we present it because we later adapt it to polymorphism. The calculus is based on the observation that a type σ must be monotonic if the problem expressed in NNF contains

no positive literal of the form $X^\sigma \approx t$ or $t \approx X^\sigma$, where X is universal. We call such an occurrence of X a naked occurrence. Naked variables are unavoidable to express upper bounds on the cardinality of types in first-order logic.

Definition 5.6 (Naked Variable). The set of *naked variables* $NV(\varphi)$ of a formula φ is defined as follows:

$$\begin{aligned} NV(p(\bar{t})) &= \emptyset & NV(t_1 \approx t_2) &= \{t_1, t_2\} \cap \mathcal{V}_{\text{typed}} \\ NV(\neg p(\bar{t})) &= \emptyset & NV(t_1 \not\approx t_2) &= \emptyset \\ NV(\varphi_1 \wedge \varphi_2) &= NV(\varphi_1) \cup NV(\varphi_2) & NV(\forall X : \sigma. \varphi) &= NV(\varphi) \\ NV(\varphi_1 \vee \varphi_2) &= NV(\varphi_1) \cup NV(\varphi_2) & NV(\exists X : \sigma. \varphi) &= NV(\varphi) - \{X^\sigma\} \end{aligned}$$

For a problem Φ , we define $NV(\Phi) = \bigcup_{\varphi \in \Phi} NV(\varphi)$.

We see from the \exists case that existential variables never occur naked in sentences.

Variables of types other than σ are irrelevant when inferring whether σ is monotonic; a variable is problematic only if it occurs naked and has type σ . Annoyingly, a single naked variable of type σ , such as X on the right-hand side of the equation $\text{hd}_b(\text{cons}_b(X, Xs)) \approx X$ from Example 2.13, will cause us to classify σ as possibly nonmonotonic. We regain some precision by extending the calculus with an infinity analysis, as suggested by Claessen et al.: trivially, all types with no finite models are monotonic.

Convention 5.7. Abstracting over the specific analysis used to detect infinite types (e.g. Infinox [15]), we fix a set $\text{Inf}(\Phi)$ of types whose interpretations are guaranteed to be infinite in all models of Φ . More precisely, the following property is assumed to hold: if $\tau \in \text{Inf}(\Phi)$ and $\mathcal{M} = ((D_\sigma)_{\sigma \in \text{Type}}, _, _)$ is a model of Φ , then D_τ is infinite.

Our monotonicity calculus takes $\text{Inf}(\Phi)$ into account:

Definition 5.8 (Monotonicity Calculus \triangleright). Let Φ be a monomorphic problem over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$. A judgement $\sigma \triangleright \Phi$ indicates that the ground type σ is inferred monotonic in Φ . The *monotonicity calculus* consists of the following rules:

$$\frac{\sigma \in \text{Inf}(\Phi)}{\sigma \triangleright \Phi} \quad \frac{NV(\Phi) \cap \{X^\sigma \mid X \in \mathcal{V}\} = \emptyset}{\sigma \triangleright \Phi}$$

We write $\sigma \triangleright \Phi$ to indicate that the judgement is derivable and $\sigma \not\triangleright \Phi$ otherwise.

Claessen et al. designed a second, more powerful calculus that extends their first calculus to detect predicates that act as guards for naked variables. Whilst the calculus proved successful on a subset of the TPTP benchmarks [30], we assessed its suitability on about 1000 problems generated by Sledgehammer and found no improvement on the simple calculus. For this reason, we restrict our attention to the first calculus.

Theorem 5.9 (Soundness of \triangleright). *Let Φ be a monomorphic problem. If $\tau \triangleright \Phi$, then τ is monotonic in Φ .*

Proof. Let $\Sigma = (\text{Type}, \mathcal{F}, \mathcal{P})$ be the signature of Φ , and let $\tau \in \text{Type}$ such that $\tau \triangleright \Phi$. Assume Φ has a model $\mathcal{M} = ((D_\sigma)_{\sigma \in \text{Type}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$. We will construct a model $\mathcal{M}' = ((D'_\sigma)_{\sigma \in \text{Type}}, (f^{\mathcal{M}'})_{f \in \mathcal{F}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}})$ with the same domains as \mathcal{M} for $\sigma \neq \tau$ and with D'_τ infinite. If $\tau \in \text{Inf}(\Phi)$, then D_τ is already infinite, so can we take $\mathcal{M}' = \mathcal{M}$. Otherwise,

$\tau \notin \text{Inf}(\Phi)$, which means the second rule of the calculus must have been applied and no variables of type τ occur naked, i.e. $\text{NV}(\Phi) \cap \{X^\tau \mid X \in \mathcal{V}\} = \emptyset$. We define D'_τ by extending D_τ with an infinite number of fresh elements e_1, e_2, \dots and define the functions and predicates of \mathcal{M}' to treat these in the same way as $\varepsilon(D_\tau)$. Intuitively, \mathcal{M}' also satisfies Φ because, since no variables of type τ occur naked, the formulae of Φ cannot tell the difference between a “clone” e_i and the original $\varepsilon(D_\tau)$ except for the case of disequality—and there the formula instantiated with clones is more likely to be true than the one instantiated with $\varepsilon(D_\tau)$ (since clones are mutually disequal and disequal to $\varepsilon(D_\tau)$).

Formally, we let $E = \{e_1, e_2, \dots\}$ be an infinite set disjoint from D_τ , and define

$$D'_\sigma = \begin{cases} D_\sigma \cup E & \text{if } \sigma = \tau \\ D_\sigma & \text{otherwise} \end{cases}$$

Given $s : \bar{\sigma} \rightarrow \varsigma$, we let $s^{\mathcal{M}'}(\bar{d}) = s^{\mathcal{M}}(\bar{d}')$, where

$$d'_i = \begin{cases} \varepsilon(D_\tau) & \text{if } \sigma_i = \tau \text{ and } d_i \in E \\ d_i & \text{otherwise} \end{cases}$$

Given $\xi' : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D'_\sigma$, we define $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D_\sigma$ by

$$\xi(X)(\sigma)(d) = \begin{cases} \varepsilon(D_\tau) & \text{if } \sigma = \tau \text{ and } d \in E \\ \xi'(X)(\sigma)(d) & \text{otherwise} \end{cases}$$

The next facts follow by induction on t or φ (for arbitrary ξ'):

- (1) $t \notin \{X^\tau \mid X \in \mathcal{V}\}$ implies $\llbracket t \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket t \rrbracket_{\xi}^{\mathcal{M}}$;
- (2) $\text{NV}(\varphi) \cap \{X^\sigma \mid X \in \mathcal{V}\} = \emptyset$ and $\llbracket \varphi \rrbracket_{\xi}^{\mathcal{M}}$ imply $\llbracket \varphi \rrbracket_{\xi'}^{\mathcal{M}'}$.

(Notice that (2) is an implication, not an equivalence. An inductive proof of the converse would fail for the case of disequalities.) In particular, thanks to the absence of naked variables of type τ in all $\varphi \in \Phi$, \mathcal{M}' is the desired model of Φ . \square

In the light of the above soundness result, we will allow ourselves to write that σ is monotonic if $\sigma \triangleright \Phi$ and possibly nonmonotonic if $\sigma \not\triangleright \Phi$.

5.3. Encoding Nonmonotonic Types. Monotonic types can be soundly erased when translating to untyped first-order logic, by Theorem 5.3. Nonmonotonic types in general cannot. Claessen et al. [16, §3.2] point out that adding sufficiently many protectors to a nonmonotonic problem will make it monotonic, at which point its types can be erased. Thus the following general two-stage procedure translates monomorphic problems to untyped first-order logic:

1. Selectively introduce protectors (tags or guards) without erasing any types:
 - 1.1. Infer monotonicity to identify the possibly nonmonotonic types in the problem.
 - 1.2. Introduce protectors for the universal variables of possibly nonmonotonic types.
 - 1.3. If necessary (depending on the encoding), generate typing axioms for any function symbol whose result type is possibly nonmonotonic, to make it possible to remove protectors for terms with the right type.
2. Erase all the types.

The purpose of stage 1 is to make the problem monotonic while preserving satisfiability. This paves the way for the sound type erasure of stage 2. The following lemmas will help us prove such two-stage encodings correct.

Lemma 5.10 (Correctness Conditions). *Let Φ be a monomorphic problem, and let x be a monomorphic encoding. The problems Φ and $\llbracket \Phi \rrbracket_{x; \mathbf{e}}$ are equisatisfiable provided that the following conditions hold:*

MONO: $\llbracket \Phi \rrbracket_x$ is monotonic.

SOUND: If Φ is satisfiable, so is $\llbracket \Phi \rrbracket_x$.

COMPLETE: If $\llbracket \Phi \rrbracket_x$ is satisfiable, so is Φ .

Proof. Immediate from Theorems 3.5 and 5.3. \square

5.4. Monotonicity-Based Type Tags. The monotonicity-based encoding $\tilde{\mathfrak{t}}?$ specialises the above procedure for tags. It is similar to the traditional encoding $\tilde{\mathfrak{t}}$ (the monomorphic version of \mathfrak{t}), except that it omits the tags for types that are inferred monotonic. By wrapping all naked variables (in fact, all terms) of possibly nonmonotonic types in a function term, stage 1 yields a monotonic problem.

Definition 5.11 (Lightweight Tags $\tilde{\mathfrak{t}}?$). The encoding $\llbracket \cdot \rrbracket_{\tilde{\mathfrak{t}}?}$ translates monomorphic problems over $\Sigma = (\text{Type}, \mathcal{F}, \mathcal{P})$ to monomorphic problems over $(\text{Type}, \mathcal{F} \uplus \{\mathfrak{t}_\sigma : \sigma \rightarrow \sigma \mid \sigma \in \text{Type}\}, \mathcal{P})$. It adds no axioms, and its term and formula translations are defined as follows:

$$\llbracket f(\bar{t}) \rrbracket_{\tilde{\mathfrak{t}}?} = \llbracket f(\llbracket \bar{t} \rrbracket_{\tilde{\mathfrak{t}}?}) \rrbracket \quad \llbracket X \rrbracket_{\tilde{\mathfrak{t}}?} = \llbracket X \rrbracket \quad \text{with } \llbracket t^\sigma \rrbracket = \begin{cases} t & \text{if } \sigma \triangleright \Phi \\ \mathfrak{t}_\sigma(t) & \text{otherwise} \end{cases}$$

The *monomorphic lightweight type tags* encoding $\tilde{\mathfrak{t}}?$ is the composition $\llbracket \cdot \rrbracket_{\tilde{\mathfrak{t}}?; \mathbf{e}}$. It translates a monomorphic problem over Σ into an untyped problem over $\Sigma' = (\mathcal{F}' \uplus \{\mathfrak{t}_\sigma^1 \mid \sigma \in \text{Type}\}, \mathcal{P}')$, where \mathcal{F}' , \mathcal{P}' are as for \mathbf{e} .

Example 5.12. For the algebraic list problem of Example 2.13, the type *list_b* is monotonic by virtue of being infinite, whereas *b* cannot be inferred monotonic. The $\tilde{\mathfrak{t}}?$ encoding of the problem follows:

$$\begin{aligned} & \forall X, Xs. \text{nil}_b \not\approx \text{cons}_b(\mathfrak{t}_b(X), Xs) \\ & \forall Xs. Xs \approx \text{nil}_b \vee (\exists Y, Ys. Xs \approx \text{cons}_b(\mathfrak{t}_b(Y), Ys)) \\ & \forall X, Xs. \mathfrak{t}_b(\text{hd}_b(\text{cons}_b(\mathfrak{t}_b(X), Xs))) \approx \mathfrak{t}_b(X) \wedge \mathfrak{t}_b(\text{cons}_b(\mathfrak{t}_b(X), Xs)) \approx Xs \\ & \exists X, Y, Xs, Ys. \text{cons}_b(\mathfrak{t}_b(X), Xs) \approx \text{cons}_b(\mathfrak{t}_b(Y), Ys) \wedge (\mathfrak{t}_b(X) \not\approx \mathfrak{t}_b(Y) \vee Xs \not\approx Ys) \end{aligned}$$

The $\tilde{\mathfrak{t}}?$ encoding treats all variables of the same type uniformly. Hundreds of axioms can suffer because of one unhappy formula that uses a type nonmonotonically (or in a way that cannot be inferred monotonic). To address this, we introduce a lighter encoding: if a universal variable does not occur naked in a formula, its tag can safely be omitted.²

This new encoding, called $\tilde{\mathfrak{t}}??$, protects only naked variables and introduces equations $\mathfrak{t}_\sigma(f(\bar{X})^\sigma) \approx f(\bar{X})$ to add or remove tags around each function symbol f whose result type σ is possibly nonmonotonic, and similarly for existential variables.

²This is related to the observation that only paramodulation from or into a variable can cause ill-typed instantiations in a resolution prover [36, §4].

Definition 5.13 (Featherweight Tags $\tilde{t}??$). The encoding $\llbracket \cdot \rrbracket_{\tilde{t}??}$ translates monomorphic problems over $\Sigma = (\text{Type}, \mathcal{F}, \mathcal{P})$ to monomorphic problems over $(\text{Type}, \mathcal{F} \uplus \{\text{t}_\sigma : \sigma \rightarrow \sigma \mid \sigma \in \text{Type}\}, \mathcal{P})$. Its term and formula translations are defined as follows:

$$\begin{aligned} \llbracket t_1 \approx t_2 \rrbracket_{\tilde{t}??} &= \llbracket \llbracket t_1 \rrbracket_{\tilde{t}??} \rrbracket_{\tilde{t}??} \approx \llbracket \llbracket t_2 \rrbracket_{\tilde{t}??} \rrbracket_{\tilde{t}??} \\ \llbracket \exists X : \sigma. \varphi \rrbracket_{\tilde{t}??} &= \exists X : \sigma. \begin{cases} \llbracket \varphi \rrbracket_{\tilde{t}??} & \text{if } \sigma \triangleright \Phi \\ \text{t}_\sigma(X) \approx X \wedge \llbracket \varphi \rrbracket_{\tilde{t}??} & \text{otherwise} \end{cases} \end{aligned}$$

with

$$\llbracket t^\sigma \rrbracket = \begin{cases} t & \text{if } \sigma \triangleright \Phi \text{ or } t \notin \mathcal{V}_{\text{typed}}^\vee \\ \text{t}_\sigma(t) & \text{otherwise} \end{cases}$$

The encoding adds the following typing axioms:

$$\begin{aligned} \forall \bar{X} : \bar{\sigma}. \text{t}_\sigma(f(\bar{X})) \approx f(\bar{X}) & \quad \text{for } f : \bar{\sigma} \rightarrow \sigma \in \mathcal{F} \text{ such that } \sigma \not\triangleright \Phi \\ \exists X : \sigma. \text{t}_\sigma(X) \approx X & \quad \text{for } \sigma \not\triangleright \Phi \text{ that is not the result type of a symbol in } \mathcal{F} \end{aligned}$$

The *monomorphic featherweight type tags* encoding $\tilde{t}??$ is the composition $\llbracket \cdot \rrbracket_{\tilde{t}??; \mathbf{e}}$. The target signature for $\tilde{t}??$ is the same as for \tilde{t} .

The axioms are necessary for the completeness of $\tilde{t}??$. They would have been harmless for \tilde{t} : for soundness, we can think of the t_σ functions as identities. The side condition of the last axiom is a minor optimisation: it avoids asserting that σ is inhabited if the symbols in \mathcal{F} already witness σ 's inhabitation.

Example 5.14. The $\tilde{t}??$ encoding of Example 2.13 requires fewer tags than \tilde{t} , at the cost of a typing axiom for hd and typing equations for the existential variables of type b :

$$\begin{aligned} \forall Xs. \text{t}_b(\text{hd}_b(Xs)) \approx \text{hd}_b(Xs) \\ \forall X, Xs. \text{nil}_b \not\approx \text{cons}_b(X, Xs) \\ \forall Xs. Xs \approx \text{nil}_b \vee (\exists Y, Ys. \text{t}_b(Y) \approx Y \wedge Xs \approx \text{cons}_b(Y, Ys)) \\ \forall X, Xs. \text{hd}_b(\text{cons}_b(X, Xs)) \approx \text{t}_b(X) \wedge \text{tl}_b(\text{cons}_b(X, Xs)) \approx Xs \\ \exists X, Y, Xs, Ys. \text{t}_b(X) \approx X \wedge \text{t}_b(Y) \approx Y \wedge \text{cons}_b(X, Xs) \approx \text{cons}_b(Y, Ys) \wedge \\ (X \not\approx Y \vee Xs \not\approx Ys) \end{aligned}$$

Theorem 5.15 (Correctness of \tilde{t} , $\tilde{t}??$). *The monomorphic type tags encodings \tilde{t} and $\tilde{t}??$ are correct.*

Proof. It suffices to show the three conditions of Lemma 5.10.

MONO: By induction on φ , it follows that

- (1) $X^\sigma \in \text{NV}(\llbracket \varphi \rrbracket_{\tilde{t}}$) implies $\sigma \in \text{Inf}(\Phi)$;
- (2) $X^\sigma \in \text{NV}(\llbracket \varphi \rrbracket_{\tilde{t}??})$ implies $\sigma \in \text{Inf}(\Phi)$.

(Indeed, while transforming φ into $\llbracket \varphi \rrbracket_{\tilde{t}??}$, $\tilde{t}??$ tags precisely the variables that would cause the condition (1) to fail; and \tilde{t} tags even more.) Moreover, the typing axioms of $\tilde{t}??$ have no naked variables, and hence $\sigma \triangleright \llbracket \varphi \rrbracket_{\tilde{t}}$ and $\sigma \triangleright \llbracket \varphi \rrbracket_{\tilde{t}??}$ hold for all types σ . Therefore, by Theorem 5.9 and Lemma 5.5, $\sigma \triangleright \llbracket \varphi \rrbracket_{\tilde{t}}$ and $\sigma \triangleright \llbracket \varphi \rrbracket_{\tilde{t}??}$ are monotonic.

SOUND: This is immediate for both $\llbracket \cdot \rrbracket_{\tilde{t}}$ and $\llbracket \cdot \rrbracket_{\tilde{t}??}$: given a model of Φ , we extend it to a model of the encoded Φ by interpreting all type tags as the identity.

COMPLETE FOR $\llbracket \cdot \rrbracket_{\tilde{t}}$: The proof is analogous to the corresponding case for t (Theorem 3.16), the differences being that here we do not face the complication of interpreting polymorphic

types, and only terms of certain types are tagged. Let $\mathcal{M}' = ((D'_\sigma)_{\sigma \in \text{Type}}, (f^{\mathcal{M}'})_{f \in \mathcal{F} \cup \{\mathfrak{t}_\sigma \mid \sigma \not\triangleright \Phi\}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}})$ be a model of $\langle\langle \Phi \rangle\rangle_{\tilde{t}??}$. For each $\sigma \in \text{Type}$, we let

$$D_\sigma = \begin{cases} D'_\sigma & \text{if } \sigma \triangleright \Phi \\ \{d \in D'_\sigma \mid d \text{ is in the image of } \mathfrak{t}_\sigma^{\mathcal{M}'}\} & \text{otherwise} \end{cases}$$

We define $T_\sigma : D'_\sigma \rightarrow D'_\sigma$ so as to apply $\mathfrak{t}_\sigma^{\mathcal{M}'}$ only if the element is not already in its image:

$$T_\sigma(d) = \begin{cases} d & \text{if } d \in D_\sigma \\ \mathfrak{t}_\sigma^{\mathcal{M}'}(d) & \text{otherwise} \end{cases}$$

Let $\mathcal{M} = ((D_\sigma)_{\sigma \in \text{Type}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$, where

- $f^{\mathcal{M}}(d_1, \dots, d_n) = \mathfrak{t}_\sigma^{\mathcal{M}'}(f^{\mathcal{M}'}(T_{\sigma_1}(d_1), \dots, T_{\sigma_n}(d_n)))$ for $f : \bar{\sigma} \rightarrow \sigma$ is in \mathcal{F} ;
- $p^{\mathcal{M}}(d_1, \dots, d_n) = p^{\mathcal{M}'}(T_{\sigma_1}(d_1), \dots, T_{\sigma_n}(d_n))$ if $p : \bar{\sigma} \rightarrow \sigma$ is in \mathcal{F} .

Given $\xi' : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D'_\sigma$, we define $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D_\sigma$ as follows:

$$\xi(X)(\sigma) = \begin{cases} \xi'(X)(\sigma) & \text{if } \sigma \triangleright \Phi \\ \mathfrak{t}_\sigma(\xi'(X)(\sigma)) & \text{otherwise} \end{cases}$$

The next facts follow by induction on t or φ (for arbitrary ξ'):

- (1) $\llbracket t \rrbracket_\xi^{\mathcal{M}} = \llbracket \langle\langle t \rangle\rangle_{\mathfrak{t}} \rrbracket_{\xi'}^{\mathcal{M}'}$;
- (2) $\llbracket \varphi \rrbracket_\xi^{\mathcal{M}} = \llbracket \langle\langle \varphi \rangle\rangle_{\mathfrak{t}} \rrbracket_{\xi'}^{\mathcal{M}'}$.

Let $\varphi \in \Phi$. To show that \mathcal{M} is a model of φ , let $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D_\sigma$. By the definition of D_σ , there exists $\xi' : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D'_\sigma$ such that the defining property of ξ holds. Now, since $\llbracket \langle\langle \varphi \rangle\rangle_{\tilde{t}??} \rrbracket_{\xi'}^{\mathcal{M}'}$ holds, it follows from (2) that $\llbracket \varphi \rrbracket_\xi^{\mathcal{M}}$, i.e. $\llbracket \varphi \rrbracket^{\mathcal{M}}$ also holds, as desired.

COMPLETE FOR $\langle\langle \rangle\rangle_{\tilde{t}??}$: The proof is analogous to the corresponding case for $\mathfrak{t}@$ (Theorem 4.10). Starting with a model $\mathcal{M}' = ((D'_\sigma)_{\sigma \in \text{Type}}, (f^{\mathcal{M}'})_{f \in \mathcal{F} \cup \{\mathfrak{t}_\sigma \mid \sigma \not\triangleright \Phi\}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}})$ of $\langle\langle \Phi \rangle\rangle_{\tilde{t}??}$, we define a model \mathcal{M} of Φ as follows, where the typing axioms ensure that each D_σ is nonempty and each $s^{\mathcal{M}}$ is well defined:

- $D_\sigma = \begin{cases} D'_\sigma & \text{if } \sigma \triangleright \Phi \\ \{d \in D'_\sigma \mid \mathfrak{t}_\sigma^{\mathcal{M}'}(d) = d\} & \text{otherwise} \end{cases}$
- $s^{\mathcal{M}}$ is the restriction of $s^{\mathcal{M}'}$.

The next facts follow by induction on t or φ (for arbitrary $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D_\sigma$):

- (1) $\llbracket t \rrbracket_\xi^{\mathcal{M}} = \llbracket \langle\langle t \rangle\rangle_{\mathfrak{t}} \rrbracket_{\xi'}^{\mathcal{M}'}$;
- (2) $\llbracket \varphi \rrbracket_\xi^{\mathcal{M}} = \llbracket \langle\langle \varphi \rangle\rangle_{\mathfrak{t}} \rrbracket_{\xi'}^{\mathcal{M}'}$.

Then, by the usual route, it follows that \mathcal{M} is a model of Φ . □

5.5. Monotonicity-Based Type Guards. The $\tilde{g}?$ and $\tilde{g}??$ encodings are defined analogously to $\tilde{t}?$ and $\tilde{t}??$ but using type guards. The $\tilde{g}?$ encoding omits the guards for types that are inferred monotonic, whereas $\tilde{g}??$ omits more guards that are not needed to make the intermediate problem monotonic.

Definition 5.16 (Lightweight Guards $\tilde{\mathfrak{g}}?$). The encoding $\llbracket \cdot \rrbracket_{\tilde{\mathfrak{g}}?}$ translates monomorphic problems over $\Sigma = (\text{Type}, \mathcal{F}, \mathcal{P})$ to monomorphic problems over $(\text{Type}, \mathcal{F}, \mathcal{P} \uplus \{\mathfrak{g}_\sigma : \sigma \rightarrow o \mid \sigma \in \text{Type}\})$. Its term and formula translations are defined as follows:

$$\begin{aligned} \llbracket \forall X : \sigma. \varphi \rrbracket_{\tilde{\mathfrak{g}}?} &= \forall X : \sigma. \begin{cases} \llbracket \varphi \rrbracket_{\tilde{\mathfrak{g}}?} & \text{if } \sigma \triangleright \Phi \\ \mathfrak{g}_\sigma(X^\sigma) \rightarrow \llbracket \varphi \rrbracket_{\tilde{\mathfrak{g}}?} & \text{otherwise} \end{cases} \\ \llbracket \exists X : \sigma. \varphi \rrbracket_{\tilde{\mathfrak{g}}?} &= \exists X : \sigma. \begin{cases} \llbracket \varphi \rrbracket_{\tilde{\mathfrak{g}}?} & \text{if } \sigma \triangleright \Phi \\ \mathfrak{g}_\sigma(X^\sigma) \wedge \llbracket \varphi \rrbracket_{\tilde{\mathfrak{g}}?} & \text{otherwise} \end{cases} \end{aligned}$$

The encoding adds the following typing axioms:

$$\begin{aligned} \forall \bar{X} : \bar{\sigma}. \mathfrak{g}_\sigma(f(\bar{X}^{\bar{\sigma}})) \quad &\text{for } f : \bar{\sigma} \rightarrow \sigma \in \mathcal{F} \text{ such that } \sigma \not\triangleright \Phi \\ \exists X : \sigma. \mathfrak{g}_\sigma(X^\sigma) \quad &\text{for } \sigma \not\triangleright \Phi \text{ that is not the result type of a symbol in } \mathcal{F} \end{aligned}$$

The *monomorphic lightweight type guards* encoding $\tilde{\mathfrak{g}}?$ is the composition $\llbracket \cdot \rrbracket_{\tilde{\mathfrak{g}}?; \mathbf{e}}$. It translates a monomorphic problem over Σ into an untyped problem over $\Sigma' = (\mathcal{F}', \mathcal{P}' \uplus \{\mathfrak{g}_\sigma^1 \mid \sigma \in \text{Type}\})$, where $\mathcal{F}', \mathcal{P}'$ are as for \mathbf{e} .

Example 5.17. The $\tilde{\mathfrak{g}}?$ encoding of Example 2.13 is as follows:

$$\begin{aligned} &\forall Xs. \mathfrak{g}_b(\text{hd}_b(Xs)) \\ &\forall X, Xs. \mathfrak{g}_b(X) \rightarrow \text{nil}_b \not\approx \text{cons}_b(X, Xs) \\ &\forall Xs. Xs \approx \text{nil}_b \vee (\exists Y, Ys. \mathfrak{g}_b(Y) \wedge Xs \approx \text{cons}_b(Y, Ys)) \\ &\forall X, Xs. \mathfrak{g}_b(X) \rightarrow \text{hd}_b(\text{cons}_b(X, Xs)) \approx X \wedge \text{tl}_b(\text{cons}_b(X, Xs)) \approx Xs \\ &\exists X, Y, Xs, Ys. \mathfrak{g}_b(X) \wedge \mathfrak{g}_b(Y) \wedge \text{cons}_b(X, Xs) \approx \text{cons}_b(Y, Ys) \wedge (X \not\approx Y \vee Xs \not\approx Ys) \end{aligned}$$

Notice that the tl_b equation is needlessly in the scope of the guard. The encoding is more precise if the problem is clausified.

Our novel encoding $\tilde{\mathfrak{g}}??$ omits the guards for variables that do not occur naked, regardless of whether they are of a monotonic type.

Definition 5.18 (Featherweight Guards $\tilde{\mathfrak{g}}??$). The *monomorphic featherweight type guards* encoding $\tilde{\mathfrak{g}}??$ is identical to the lightweight encoding $\tilde{\mathfrak{g}}?$ except that the condition “if $\sigma \triangleright \Phi$ ” in the \forall case is weakened to “if $\sigma \triangleright \Phi$ or $X \notin \text{NV}(\varphi)$ ”.

Example 5.19. The $\tilde{\mathfrak{g}}??$ encoding of the algebraic list problem is identical to $\tilde{\mathfrak{g}}?$ except that the $\text{nil}_b \not\approx \text{cons}_b$ axiom does not have any guard.

Theorem 5.20 (Correctness of $\tilde{\mathfrak{g}}?$, $\tilde{\mathfrak{g}}??$). *The monomorphic type guards encodings $\tilde{\mathfrak{g}}?$ and $\tilde{\mathfrak{g}}??$ are correct.*

Proof. It suffices to show the three conditions of Lemma 5.10.

MONO: By Lemma 5.5, it suffices to show that each type τ such that $\sigma \not\triangleright \Phi$ is monotonic in the encoded problem. By Theorem 5.9, types are monotonic unless they are possibly finite and variables of their types occur naked in the original problem. Both encodings guard all such variables— $\tilde{\mathfrak{g}}??$ guards exactly those variables, while $\tilde{\mathfrak{g}}?$ guards more. The typing axioms contain no naked variables. We cannot use Theorem 5.9 directly, because guarding a naked variable does not make it less naked—but we can generalise the proof slightly to exploit the guards. Given a model \mathcal{M} of the encoded problem $\mathcal{M} = ((D_\sigma)_{\sigma \in \text{Type}}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P} \uplus \{\mathfrak{g}_\sigma \mid \sigma \not\triangleright \Phi\}})$, we construct a model $\mathcal{M}' = ((D'_\sigma)_{\sigma \in \text{Type}},$

$(f^{\mathcal{M}'})_{f \in \mathcal{F}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}}$) with all types infinite as in the proof of Theorem 5.9, but defining the τ -guard interpretation to be false on newly added elements (recall that $D'_\tau = D_\tau \uplus E$ for some countably infinite set E):

$$\mathfrak{g}_\sigma^{\mathcal{M}'}(d) = \begin{cases} \mathfrak{g}_\sigma^{\mathcal{M}}(d) & \text{if } \sigma \neq \tau \text{ or } \sigma = \tau \text{ and } d \in D_\sigma \\ \text{false} & \text{if } \sigma = \tau \text{ and } d \in E \end{cases}$$

Given $\xi' : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D'_\sigma$, we define $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}} D_\sigma$ by

$$\xi(\sigma)(d) = \begin{cases} \varepsilon(D_\tau) & \text{if } \sigma = \tau \text{ and } d \in E \\ \xi'(\sigma)(d) & \text{otherwise} \end{cases}$$

The encoded problem has the form $\Phi_1 \cup \Phi_2$, where Φ_1 are the added axioms and Φ_2 are the formula translations from Φ . It is easy to check that \mathcal{M}' satisfies Φ_1 . We say that a formula is τ -*guarded* if all its subformulae of the form $\forall X : \tau. \varphi$ have φ of the form $\neg \mathfrak{g}_\tau(X^\tau) \vee \chi$ and all its subformulae of the form $\exists X : \tau. \varphi$ have φ of the form $\mathfrak{g}_\tau(X^\tau) \wedge \chi$. All the formulae in Φ_2 are clearly guarded. The next facts follow by induction on t or φ (for arbitrary ξ'):

- (1) $\llbracket t \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket t \rrbracket_{\xi}^{\mathcal{M}}$;
- (2) If φ is guarded and $\mathfrak{g}_\tau(\xi(\tau)(X))$ is false whenever $X^\tau \in \text{NV}(\varphi) \cap \text{FVars}(\varphi)$, then $\llbracket \varphi \rrbracket_{\xi}^{\mathcal{M}}$ implies $\llbracket \varphi \rrbracket_{\xi'}^{\mathcal{M}'}$.

It follows that \mathcal{M}' satisfies Φ_2 as well. Hence it is the desired model of the encoded problem.

SOUND: Just like for tags, this is immediate for both $\llbracket \tilde{\mathfrak{g}}? \rrbracket$ and $\llbracket \tilde{\mathfrak{g}}?? \rrbracket$: given a model of Φ , we extend it to a model of the encoded Φ by interpreting all type guards as everywhere-true predicates.

COMPLETE: The proofs for both $\llbracket \tilde{\mathfrak{g}}? \rrbracket$ and $\llbracket \tilde{\mathfrak{g}}?? \rrbracket$ are very similar to that for $t??$ (from Theorem 5.15)—the only change is the replacement of the condition $t_\sigma^{\mathcal{M}'}(d) = d$ by $\mathfrak{g}_\sigma^{\mathcal{M}'}(d)$ in the definition of the model \mathcal{M} . (Just like for $t??$, the typing axioms ensure that the structure is well defined.) \square

A simpler but less instructive way to prove MONO is to observe that the second monotonicity calculus by Claessen et al. [16, §2.4] can infer monotonicity of all problems generated by $\tilde{\mathfrak{g}}, \tilde{\mathfrak{g}}?,$ and $\tilde{\mathfrak{g}}??$.

Remark 5.21. The proofs of Theorems 5.15 and 5.20 remain valid as they are even if the generated problems contain more tags or guards than inferred as unnecessary by the monotonicity calculus, i.e. if in the definitions of $\llbracket \tilde{\mathfrak{t}}? \rrbracket, \llbracket \tilde{\mathfrak{t}}?? \rrbracket, \llbracket \tilde{\mathfrak{g}}? \rrbracket,$ and $\llbracket \tilde{\mathfrak{g}}?? \rrbracket$ we replace the condition “ $\sigma \triangleright \Phi$ ” by “ $\sigma \in J$ ”, where $J \subseteq \{\sigma \in \text{Type}_\Sigma \mid \sigma \triangleright \Phi\}$. (The replacement should be done everywhere, including in the axioms.)

5.6. Heuristic Monomorphisation. Section 6 will show how to translate polymorphic types soundly and completely. If we are willing to sacrifice completeness, an easy way to extend $\tilde{\mathfrak{t}}?, \tilde{\mathfrak{t}}??, \tilde{\mathfrak{g}}?,$ and $\tilde{\mathfrak{g}}??$ to polymorphism is to perform *heuristic monomorphisation* on the polymorphic problem:

1. Heuristically instantiate all type variables with suitable ground types, taking finitely many copies of each formula if desired.

2. Map each ground occurrence $s(\bar{\alpha}\rho)$ of a polymorphic symbol $s : \forall\bar{\alpha}. \bar{\sigma} \rightarrow \sigma$ to a fresh monomorphic symbol $s_{\bar{\alpha}\rho} : [\bar{\sigma}\rho] \rightarrow [\sigma\rho]$, where ρ is a ground type substitution (a function from type variables to ground types) and $[\]$ is an injection from ground types to nullary type constructors (e.g. $\{b \mapsto b, \text{list}(b) \mapsto \text{list}_b\}$).

Heuristic monomorphisation is generally incomplete [12, §2] and often overlooked in the literature, but by eliminating type variables it considerably simplifies the generated formulae, leading to very efficient encodings. It also provides a simple and effective way to exploit the native support for monomorphic types in some automatic provers.

6. COMPLETE MONOTONICITY-BASED ENCODINGS OF POLYMORPHISM

Heuristic monomorphisation is simple and effective, but its incompleteness can be a cause for worry, and its nonmodular nature makes it unsuitable for some applications that need to export an entire polymorphic theory independently of any conjecture. Here we adapt the type encodings to a polymorphic setting.

We start by defining a correct but infinitary translation of a polymorphic into a monomorphic problem, called complete monomorphisation. Then we address the genuinely polymorphic issue of encoding type arguments, proving conditions under which composition of an encoding x with the type arguments encoding a is correct: x must be complete and produce monotonic problems. Finally, we define polymorphic counterparts of the guard and tag encodings and show that they satisfy the required conditions by reducing (most of) the problem to the monomorphic case via complete monomorphisation.

6.1. Complete Monomorphisation. The main insight behind complete monomorphisation is that a polymorphic formula having all its type quantification at the top is equisatisfiable to the (generally infinite) set of its monomorphic instances. Complete monomorphisation does not obey our convention about encodings, since it translates each polymorphic formula not into a single monomorphic formula but into a set of formulae.

Definition 6.1 (Instance Type and Most General Instance). A type τ is an *instance* of a type σ if there exists a type substitution ρ such that $\tau = \sigma\rho$. If this is the case, we also say that τ is *less general* than σ and that σ is *more general* than τ , and we write $\tau \leq \sigma$. Given two types σ and τ , if they have a common instance (i.e. a unifier), then they also have a most general common instance, which we denote by $\text{mgi}(\sigma, \tau)$.

Definition 6.2 (Complete Monomorphisation ∞). We define the encoding $\langle\langle \ \rangle\rangle_\infty$ that translates polymorphic problems over $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ to monomorphic problems over $\Sigma' = (\text{Type}', \mathcal{F}', \mathcal{P}')$, where

- $\text{Type}' = \text{GType}_\Sigma$;
- for each $s : \forall\bar{\alpha}. \bar{\sigma} \rightarrow \varsigma$ in $\mathcal{F} \uplus \mathcal{P}$ and each ρ such that each α_i is mapped to a ground type $\tau_i \in \text{GType}_\Sigma$, $\mathcal{F}' \uplus \mathcal{P}'$ contains a symbol $s_{\bar{\tau}} : \bar{\sigma}\rho \rightarrow \sigma\rho$.

We first define the translation of terms and formulae that contain no type variables, i.e. that have no type quantifiers and such that all the occurring types in applications of the function

or predicate symbols are ground:

$$\begin{aligned} \langle\langle f(\bar{\sigma})(\bar{t}) \rangle\rangle_\infty &= f_{\bar{\sigma}}(\langle\langle \bar{t} \rangle\rangle_\infty) & \langle\langle X^\sigma \rangle\rangle_\infty &= X^\sigma \\ \langle\langle p(\bar{\sigma})(\bar{t}) \rangle\rangle_\infty &= p_{\bar{\sigma}}(\langle\langle \bar{t} \rangle\rangle_\infty) & \langle\langle \forall X : \sigma. \varphi \rangle\rangle_\infty &= \forall X : \sigma. \langle\langle \varphi \rangle\rangle_\infty \\ \langle\langle \neg p(\bar{\sigma})(\bar{t}) \rangle\rangle_\infty &= \neg p_{\bar{\sigma}}(\langle\langle \bar{t} \rangle\rangle_\infty) & \langle\langle \exists X : \sigma. \varphi \rangle\rangle_\infty &= \exists X : \sigma. \langle\langle \varphi \rangle\rangle_\infty \end{aligned}$$

Now, given a sentence $\forall \bar{\alpha}. \varphi$ where $\bar{\alpha}$ indicates all its universally quantified types, we encode it as the set of encodings of its monomorphic instances (via ground type substitutions ρ):

$$\langle\langle \forall \bar{\alpha}. \varphi \rangle\rangle_\infty = \{ \langle\langle \varphi \rho \rangle\rangle_\infty \mid \rho : \mathcal{A} \rightarrow GType \}$$

Finally, the encoding of a problem is the union of the encoding of its formulae:

$$\langle\langle \Phi \rangle\rangle_\infty = \bigcup_{\varphi \in \Phi} \langle\langle \varphi \rangle\rangle_\infty$$

Convention 6.3. Whenever we write a polymorphic formula as $\forall \bar{\alpha}. \varphi$, we implicitly assume that $\bar{\alpha}$ indicates all its universally quantified types, so that φ has no type quantifiers.

Lemma 6.4 (Correctness of ∞). *The complete monomorphisation encoding ∞ is correct.*

Proof. First, observe that in any model \mathcal{M} of a polymorphic signature, the interpretation $\llbracket \cdot \rrbracket_{\theta, \xi}^{\mathcal{M}}$ of a term or formula that does not contain type variables does not depend on θ and, from ξ , it only depends on the restriction of ξ to ground types, $\xi' : \mathcal{V} \rightarrow \prod_{\sigma \in GType_\Sigma} \llbracket \sigma \rrbracket^{\mathcal{M}}$. We can therefore write $\llbracket \cdot \rrbracket_{\xi'}^{\mathcal{M}}$ instead of $\llbracket \cdot \rrbracket_{\theta, \xi'}^{\mathcal{M}}$.

SOUND: Assume Φ is satisfiable and let $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ be its polymorphic signature. By Lemma 4.11, Φ also has a model $\mathcal{M} = (\mathbb{D}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ for which $\llbracket \cdot \rrbracket^{\mathcal{M}} : GType_\Sigma \rightarrow \mathbb{D}$ is a bijection—let $\nu : \mathbb{D} \rightarrow GType_\Sigma$ be its inverse. We define a structure $\mathcal{M}' = ((D'_\tau)_{\tau \in Type'_{\Sigma}}, (f^{\mathcal{M}'})_{f \in \mathcal{F}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}'})$ for Σ' as follows:

- $D'_\tau = \llbracket \tau \rrbracket^{\mathcal{M}}$;
- if $s : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \zeta$ is in $\mathcal{F} \uplus \mathcal{P}$ and $\bar{\tau} = (\tau_1, \dots, \tau_m)$ are ground instances of $\bar{\alpha} = (\alpha_1, \dots, \alpha_m)$ via ρ , we define by $s_{\bar{\tau}}^{\mathcal{M}'}(\bar{d}) = s^{\mathcal{M}}(\nu(\tau_1), \dots, \nu(\tau_m))(\bar{d})$.

The next fact follows by induction on the term or formula δ (for arbitrary $\xi : \mathcal{V} \rightarrow \prod_{\tau \in Type_{\Sigma'}} D'_\tau$):

(1) If δ contains no type variables, then $\llbracket \langle\langle \delta \rangle\rangle_\infty \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket \delta \rrbracket_{\xi}^{\mathcal{M}}$.

In particular, for a sentence φ , we have $\llbracket \langle\langle \varphi \rangle\rangle_\infty \rrbracket^{\mathcal{M}'} = \llbracket \varphi \rrbracket^{\mathcal{M}}$.

Now let $\forall \bar{\alpha}. \varphi \in \Phi$. Since \mathcal{M} is a model of $\forall \bar{\alpha}. \varphi$, \mathcal{M} is also a model of all its ground-type instances $\varphi \rho$, and hence, since $\llbracket \langle\langle \varphi \rangle\rangle_\infty \rrbracket^{\mathcal{M}'} = \llbracket \varphi \rrbracket^{\mathcal{M}}$, \mathcal{M}' is a model of each formula in $\langle\langle \forall \bar{\alpha}. \varphi \rangle\rangle_\infty$. Thus, \mathcal{M}' is a model of $\langle\langle \Phi \rangle\rangle_\infty$, as desired.

COMPLETE: Let $\mathcal{M}' = ((D'_\tau)_{\tau \in Type'_{\Sigma}}, (f^{\mathcal{M}'})_{f \in \mathcal{F}}, (p^{\mathcal{M}'})_{p \in \mathcal{P}'})$ be a model of $\langle\langle \Phi \rangle\rangle_\infty$, for which we can assume without loss of generality that $D'_\tau \cap D'_{\tau'} = \emptyset$ if $\tau \neq \tau'$. We define a structure $\mathcal{M} = (\mathbb{D}, (f^{\mathcal{M}})_{f \in \mathcal{F}}, (p^{\mathcal{M}})_{p \in \mathcal{P}})$ for Σ as follows:

- $\mathbb{D} = \{D'_\tau \mid \tau \in Type'_{\Sigma}\}$ —for each $D \in \mathbb{D}$, we let $\nu(D)$ be the unique $\tau \in Type'_{\Sigma} = GType_{\Sigma}$ such that $D = D'_\tau$;
- if $\bar{\alpha} = (\alpha_1, \dots, \alpha_m)$, $s : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \zeta$ in $\mathcal{F} \uplus \mathcal{P}$, and $\bar{D} \in \mathbb{D}^m$, we define $s^{\mathcal{M}}(\bar{D}) = s^{\mathcal{M}'}_{(\nu(D_1), \dots, \nu(D_m))}$.

The next fact follows by induction on τ :

(1) $\tau \in GType_{\Sigma}$ implies $\llbracket \tau \rrbracket^{\mathcal{M}} = D'_\tau$.

The next fact follows by induction on the term or formula δ (for arbitrary $\xi : \mathcal{V} \rightarrow \prod_{\tau \in GType_\Sigma} \llbracket \tau \rrbracket^{\mathcal{M}}$):

(2) If δ contains no type variables, then $\llbracket \delta \rrbracket_\xi^{\mathcal{M}} = \llbracket \langle \delta \rangle_\infty \rrbracket_{\xi'}^{\mathcal{M}'}$.

In particular, for a sentence φ , we have $\llbracket \varphi \rrbracket^{\mathcal{M}} = \llbracket \langle \varphi \rangle_\infty \rrbracket^{\mathcal{M}'}$. Now let $\forall \bar{\alpha}. \varphi \in \Phi$, and assume by absurdity that \mathcal{M} is not a model of $\forall \bar{\alpha}. \varphi$. Then, since by (1) and the definition of \mathbb{D} we have that $\llbracket \cdot \rrbracket^{\mathcal{M}} : GType \rightarrow \mathbb{D}$ is surjective, we obtain a ground-type instance $\varphi\rho$ of φ such that \mathcal{M} is not a model of $\varphi\rho$. By $\llbracket \varphi \rrbracket^{\mathcal{M}} = \llbracket \langle \varphi \rangle_\infty \rrbracket^{\mathcal{M}'}$, \mathcal{M}' is not a model of $\llbracket \langle \varphi\rho \rangle_\infty \rrbracket^{\mathcal{M}'}$, hence not a model of $\langle \forall \bar{\alpha}. \varphi \rangle_\infty$, which is a contradiction. Therefore \mathcal{M} is a model of $\forall \bar{\alpha}. \varphi$. We obtain that \mathcal{M} is a model Φ , as desired. \square

6.2. Monotonicity. The definition of monotonicity from Section 5.1 (Definition 5.1) must be adapted to the polymorphic case.

Definition 6.5 (Monotonicity). Let S be a set of types and Φ be a polymorphic problem. The set S is *monotonic* in Φ if for all models \mathcal{M} of Φ , there exists a model \mathcal{M}' of Φ such that for all ground types σ , $\llbracket \sigma \rrbracket^{\mathcal{M}'}$ is infinite if σ is an instance of a type in S and $|\llbracket \sigma \rrbracket^{\mathcal{M}'}| = |\llbracket \sigma \rrbracket^{\mathcal{M}}|$ otherwise. A type σ is *monotonic* if $\{\sigma\}$ is monotonic. The problem Φ is *monotonic* if the set *Type* of all types is monotonic.

Example 6.6. For the algebraic list problem of Example 2.10, $S = \{list(\alpha)\}$ is monotonic because all models \mathcal{M} of the problem necessarily interpret all of the ground instances of $list(\alpha)$ (e.g. $list(b)$, $list(list(b))$) by infinite domains. Monotonicity is trivially witnessed by taking $\mathcal{M}' = \mathcal{M}$.

Theorem 6.7 (Monotonic Erasure). *The traditional type arguments encoding \mathbf{a} is sound for monotonic polymorphic problems.*

Proof. Let Φ be such a problem and $\Sigma = (\mathcal{K}, \mathcal{F}, \mathcal{P})$ be its signature, and assume Φ is satisfiable. By monotonicity and Lemma 4.12, it also has a model \mathcal{M} where all $D \in \mathbb{D}$ are countably infinite. From this model, we construct a model \mathcal{M}' of $\langle \Phi \rangle_{\mathbf{a}; \mathbf{e}}$ with a countably infinite domain E that interprets the encoded types as distinct elements of E . The function and predicate tables for \mathcal{M}' are based on those from \mathcal{M} , with encoded type arguments corresponding to actual type arguments.

More precisely, let E be an countably infinite set and consider the following functions:

- the mutually inverse bijections $u : \mathbb{D} \rightarrow E$ and $v : E \rightarrow \mathbb{D}$;
- for each $D \in \mathbb{D}$, the mutually inverse bijections $u_D : D \rightarrow E$ and $v_D : E \rightarrow D$.

Let $\Sigma' = (\mathcal{F}' \uplus \mathcal{K}', \mathcal{P}')$ be the target untyped signature of \mathbf{a} . We define the structure \mathcal{M}' for Σ' as follows:

- $D = E$.
- Assume k is an n -ary function in \mathcal{K}' , meaning $k :: n$ is in \mathcal{K} . Then $k^{\mathcal{M}'}(\bar{e}) = u(k^{\mathcal{M}}(v(\bar{e})))$.
- Assume $\bar{\alpha} = \alpha_1, \dots, \alpha_m$, $\bar{\sigma} = \sigma_1, \dots, \sigma_n$, and $s : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \varsigma \in \mathcal{F} \uplus \mathcal{P}$, meaning that s is an $(m+n)$ -ary symbol in $\mathcal{F}' \uplus \mathcal{P}'$. Given $\bar{c} = (c_1, \dots, c_m) \in E^m$ and $\bar{e} = (e_1, \dots, e_n) \in E^n$, we define $s^{\mathcal{M}'}(\bar{c}, \bar{e}) = u_{\llbracket \sigma \rrbracket_\theta^{\mathcal{M}'}}(s^{\mathcal{M}}(v(\bar{c}))(v_{\llbracket \sigma_1 \rrbracket_\theta^{\mathcal{M}'}}(e_1), \dots, v_{\llbracket \sigma_n \rrbracket_\theta^{\mathcal{M}'}}(e_n)))$, where θ maps each α_i to $v(c_i)$.

Given $\xi' : \mathcal{V} \rightarrow E$, we define $\theta : \mathcal{A} \rightarrow \mathbb{D}$ by $\theta(\alpha) = v(\xi'(\mathcal{V}(\alpha)))$ and $\xi : \mathcal{V} \rightarrow \prod_{\sigma \in \text{Type}_\Sigma} \llbracket \sigma \rrbracket_\theta^{\mathcal{M}}$ by $\xi(X)(\sigma) = v_{\llbracket \sigma \rrbracket_{\xi'}^{\mathcal{M}'}}(\xi'(X))$, where $\mathcal{V}(\alpha)$ and $\llbracket \sigma \rrbracket$ are as in Definition 3.6.

The next facts follow by structural induction on σ , t , and φ (for arbitrary ξ'):

- $\llbracket \llbracket \sigma \rrbracket \rrbracket_{\xi'}^{\mathcal{M}'} = u(\llbracket \sigma \rrbracket_\theta^{\mathcal{M}})$;
- $t : \sigma$ implies $\llbracket \llbracket t \rrbracket_{\mathbf{a}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} = u_{\llbracket \sigma \rrbracket_\theta^{\mathcal{M}}}(\llbracket t \rrbracket_{\theta, \xi}^{\mathcal{M}})$;
- $\llbracket \llbracket \varphi \rrbracket_{\mathbf{a}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket \varphi \rrbracket_{\theta, \xi}^{\mathcal{M}}$.

In particular, for sentences, we have $\llbracket \llbracket \varphi \rrbracket_{\mathbf{a}; \mathbf{e}} \rrbracket_{\xi'}^{\mathcal{M}'} = \llbracket \varphi \rrbracket^{\mathcal{M}}$; and since \mathcal{M} is a model of Φ , it follows that \mathcal{M}' is a model of $\llbracket \Phi \rrbracket_{\mathbf{a}; \mathbf{e}}$. \square

Lemma 6.8 (Correctness Conditions). *Let Φ be a polymorphic problem, and let x be a polymorphic encoding. The problems Φ and $\llbracket \Phi \rrbracket_{x; \mathbf{a}; \mathbf{e}}$ are equisatisfiable provided that the following conditions hold:*

MONO: $\llbracket \Phi \rrbracket_x$ is monotonic.

SOUND: If Φ is satisfiable, so is $\llbracket \Phi \rrbracket_x$.

COMPLETE: If $\llbracket \Phi \rrbracket_x$ is satisfiable, so is Φ .

Proof. Immediate from Theorems 3.11 and 6.7. \square

Lemma 6.9 (Monotonicity Preservation and Reflection by ∞). *A polymorphic problem Φ is monotonic iff its monomorphic encoding $\llbracket \Phi \rrbracket_\infty$ is monotonic.*

Proof. Let φ_k be the polymorphic formula stating that all type domains have at least k elements, and let $\Phi' = \{\varphi_k \mid k \in \mathbb{N}\}$. We note that the following four statements are equivalent:

- (1) Φ has a model with all domains infinite;
- (2) $\Phi \cup \Phi'$ has a model;
- (3) $\llbracket \Phi \rrbracket_\infty \cup \llbracket \Phi' \rrbracket_\infty$ has a model;
- (4) $\llbracket \Phi \rrbracket_\infty$ has a model with all domains infinite.

The equivalences “(1) iff (2)” and “(3) iff (4)” are obvious, and “(2) iff (3)” follows from the correctness of ∞ (Lemma 6.4). \square

6.3. Monotonicity Inference. The monotonicity inference of Section 5.2 must be adapted to the polymorphic setting. We start by generalising the notion of naked variable.

Definition 6.10 (Polymorphic Naked Variable). The set of *naked variables* $\text{NV}(\varphi)$ of a polymorphic formula φ is defined similarly to the monomorphic case (Definition 5.6). The following equations are new or slightly different from there:

$$\text{NV}(p(\bar{\sigma})(\bar{t})) = \emptyset \quad \text{NV}(\neg p(\bar{\sigma})(\bar{t})) = \emptyset \quad \text{NV}(\forall \alpha. \varphi) = \text{NV}(\varphi)$$

Again, we take $\text{NV}(\Phi) = \bigcup_{\varphi \in \Phi} \text{NV}(\varphi)$.

Note from Definitions 5.6 and 6.10 that the naked variables of a polymorphic formula occur at the same positions as in all its monomorphic instances. Moreover, the types of the naked variables in the completely monomorphised problem are simply the ground instances of the types of the naked variables in the original problem:

Lemma 6.11. $\text{NV}(\llbracket \Phi \rrbracket_\infty) = \{X^\tau \mid \tau \in G\text{Type and there exists } \sigma \geq \tau \text{ such that } X^\sigma \in \text{NV}(\Phi)\}$.

The calculus presented below captures the insight that a polymorphic type is monotonic if each of its common instances with the type of any naked variable is an instance of an infinite type. Similarly to its monomorphic counterpart, it is parameterised by a fixed set $\text{Inf}(\Phi)$ of (not necessarily ground) types for which any interpretation $\llbracket \sigma \rrbracket_{\theta}^{\mathcal{M}}$ (for any type valuation θ) in any model \mathcal{M} of Φ is known to be infinite.

Convention 6.12. It is easy to see that if $\sigma \in \text{Inf}(\Phi)$ and $\rho : \mathcal{A} \rightarrow GType_{\Sigma}$, then $\sigma\rho$ is infinite in all models of $\langle\langle \Phi \rangle\rangle_{\infty}$. We fix the “known as infinite” types of $\langle\langle \Phi \rangle\rangle_{\infty}$, $\text{Inf}(\langle\langle \Phi \rangle\rangle_{\infty})$, to be $\{\sigma\rho \mid \sigma \in \text{Inf}(\Phi) \text{ and } \rho : \mathcal{A} \rightarrow GType_{\Sigma}\}$.

Definition 6.13 (Polymorphic Monotonicity Calculus \triangleright). A judgement $\sigma \triangleright \Phi$ indicates that the type σ is inferred monotonic in Φ . The *monotonicity calculus* consists of the single rule

$$\frac{\text{for all } X^{\sigma'} \in \text{NV}(\Phi), \text{ if } \sigma \text{ and } \sigma' \text{ have a common instance, then } \text{mgi}(\sigma, \sigma') \in \text{Inf}^*(\Phi)}{\sigma \triangleright \Phi}$$

where $\text{Inf}^*(\Phi) = \{\sigma \in Type_{\Sigma} \mid \text{there exists } \sigma' \in \text{Inf}(\Phi) \text{ such that } \sigma \leq \sigma'\}$ consists of all instances of all types in $\text{Inf}(\Phi)$.

Example 6.14. For the algebraic list problem of Example 2.10, the only naked variables are X^{α} and $Xs^{list(\alpha)}$, i.e. $\text{NV}(\Phi) = \{X^{\alpha}, Xs^{list(\alpha)}\}$. Assume $\text{Inf}(\Phi) = \{list(\alpha)\}$. Then $\text{Inf}^*(\Phi) = \{list(\sigma) \mid \sigma \in Type_{\Sigma}\}$. The type $list(\alpha)$ is inferred monotonic by \triangleright because its most general common instance with α (the type of X) and $list(\alpha)$ (the type of Xs) is in both cases $list(\alpha)$, which is known to be infinite. In contrast, α and b cannot be inferred monotonic: each of them is its most general common instance with α , and neither of them is among the types that are known to be infinite.

Lemma 6.15. *The following properties hold for any polymorphic signature Σ and any Σ -problem Φ :*

- (1) *If $\sigma \triangleright \Phi$ and $\tau \in GType_{\Sigma}$ such that $\tau \leq \sigma$, then $\tau \triangleright \langle\langle \Phi \rangle\rangle_{\infty}$.*
- (2) *If $\tau \in GType_{\Sigma}$, then $\tau \triangleright \langle\langle \Phi \rangle\rangle_{\infty}$ iff $\tau \triangleright \Phi$.*
- (3) *If $\sigma \triangleright \Phi$ and $\sigma' \leq \sigma$, then $\sigma' \triangleright \Phi$.*

Proof. (1): Assume $\sigma \triangleright \Phi$ and $\sigma \geq \tau \in GType_{\Sigma}$. To show $\tau \triangleright \langle\langle \Phi \rangle\rangle_{\infty}$, it suffices to let $X^{\tau} \in \text{NV}(\langle\langle \Phi \rangle\rangle_{\infty})$ and prove $\tau \in \text{Inf}(\langle\langle \Phi \rangle\rangle_{\infty})$. By Lemma 6.12, we obtain $\sigma' \geq \tau$ such that $X^{\sigma'} \in \text{NV}(\Phi)$. Since σ and σ' have a common instance, by the first assumption we have $\text{mgi}(\sigma, \sigma') \in \text{Inf}^*(\Phi)$, and hence, since $\tau \leq \text{mgi}(\sigma, \sigma')$, we have $\tau \in \text{Inf}^*(\Phi)$ —and since τ is ground, we obtain $\tau \in \text{Inf}(\langle\langle \Phi \rangle\rangle_{\infty})$, as desired.

(2): Let (A) $\tau \in GType_{\Sigma}$. One implication follows immediately from (1). For the second, assume (B) $\tau \triangleright \langle\langle \Phi \rangle\rangle_{\infty}$. To prove $\tau \triangleright \Phi$, we fix (C) $X^{\sigma} \in \text{NV}(\Phi)$ such that σ and τ have a common instance (i.e. by (A), $\tau \leq \sigma$) and prove $\text{mgi}(\sigma, \tau) \in \text{Inf}^*(\Phi)$, i.e. by (A), $\tau \in \text{Inf}^*(\Phi)$, i.e. again by (A), $\tau \in \text{Inf}(\langle\langle \Phi \rangle\rangle_{\infty})$. From (C), $\tau \leq \sigma$ and Lemma 6.12, we obtain $X^{\tau} \in \text{NV}(\langle\langle \Phi \rangle\rangle_{\infty})$; hence, by (A), we have $\tau \in \text{Inf}(\langle\langle \Phi \rangle\rangle_{\infty})$, as desired.

(3): Immediate from the definition. □

Note that property (1) of the above lemma states that if a type is inferred monotonic in the polymorphic calculus, all its ground instances can be inferred monotonic in the monomorphic calculus associated to the completely monomorphised problem. This allows us to prove soundness of the former from soundness of the latter:

Theorem 6.16 (Soundness of \triangleright). *Let Φ be a polymorphic problem. If $\sigma \triangleright \Phi$ for all $\sigma \in \text{Type}_\Sigma$, then Φ is monotonic.*

Proof. Assume $\sigma \triangleright \Phi$ for all $\sigma \in \text{Type}_\Sigma$. By Lemma 6.15, $\tau \triangleright \langle\langle \Phi \rangle\rangle_\infty$ for all ground instances τ of types in Type_Σ , i.e. for all τ in the signature of $\langle\langle \Phi \rangle\rangle_\infty$. By Theorem 5.9 and Lemma 5.5, this makes $\langle\langle \Phi \rangle\rangle_\infty$ monotonic. Then, by Lemma 6.4, Φ is also monotonic. \square

6.4. General Strategy. We will define polymorphic versions of the featherweight and lightweight tags and guard encodings, altogether four encodings: $t?$, $t??$, $g?$, and $g??$. If x ranges over the polymorphic encodings, $\tilde{x} \in \{\tilde{t?}, \tilde{t??}, \tilde{g?}, \tilde{g??}\}$ denotes its monomorphic counterpart. We base the correctness of each x on the correctness of \tilde{x} , by proving the problems $\langle\langle \Phi \rangle\rangle_{\infty; \tilde{x}}$ and $\langle\langle \Phi \rangle\rangle_{x; \infty}$ equisatisfiable in a monotonicity-preserving way. The following lemma implicitly assumes that the signatures of $\langle\langle \Phi \rangle\rangle_{\infty; \tilde{x}}$ and $\langle\langle \Phi \rangle\rangle_{x; \infty}$ coincide, which is easy to check for each encoding x .

Lemma 6.17 (Correctness Conditions via Complete Monomorphisation). *Let Φ be a polymorphic problem and let $x \in \{t?, t??, g?, g??\}$. The problems Φ and $\langle\langle \Phi \rangle\rangle_{x; \mathbf{a}; \mathbf{e}}$ are equisatisfiable provided that the following conditions hold:*

SOUND: *If $\langle\langle \Phi \rangle\rangle_{\infty; \tilde{x}}$ has a model, $\langle\langle \Phi \rangle\rangle_{x; \infty}$ has a model with the same interpretation of the type constructors.*

COMPLETE: *If $\langle\langle \Phi \rangle\rangle_{x; \infty}$ has a model, $\langle\langle \Phi \rangle\rangle_{\infty; \tilde{x}}$ has a model with the same interpretation of the type constructors.*

Proof. By the correctness of ∞ (Lemma 6.4) and the corresponding correctness theorems for \tilde{x} (more precisely, from the SOUND and COMPLETE statements from the proofs of Theorems 5.15 and 5.20), we have that Φ and $\langle\langle \Phi \rangle\rangle_{\infty; \tilde{x}}$ are equisatisfiable. Together with the assumptions, this implies that (A) Φ and $\langle\langle \Phi \rangle\rangle_{x; \infty}$ are equisatisfiable. Moreover, from Lemma 6.9 and monotonicity of \tilde{x} (more precisely, from the MONO statements from the proofs of Theorems 5.15 and 5.20), we know that $\langle\langle \Phi \rangle\rangle_{\infty; \tilde{x}}$ is monotonic. Hence, by the assumptions, (B) $\langle\langle \Phi \rangle\rangle_{x; \infty}$ is also monotonic. The desired fact follows from (A), (B), and Lemma 6.8 (with x instantiated with $x; \infty$). \square

To apply the above lemma, we need to provide an x such that the equation $\langle\langle \Phi \rangle\rangle_{x; \infty} = \langle\langle \Phi \rangle\rangle_{\infty; \tilde{x}}$ almost holds, in the sense that the two problems are equisatisfiable without changing the type constructor interpretation, but possibly changing some of the function or predicate symbol interpretation.

Given \tilde{x} , we will come up with an encoding x in a systematic way. But first we need to define some relevant sets of types for a polymorphic Σ -problem Φ :

Definition 6.18. Let

- $\text{Type}_\Phi = \{\sigma \in \text{Type}_\Sigma \mid \sigma \text{ is the type of a subterm occurring in } \Phi\};$
- $S_\Phi = \{\tau \in \text{GType}_\Sigma \mid \tau \not\triangleright \langle\langle \Phi \rangle\rangle_\infty \text{ and there exists } \sigma \in \text{Type}_\Phi \text{ such that } \tau \leq \sigma\};$
- $T_\Phi = \{\tau \in \text{GType}_\Sigma \mid \text{there exists } \sigma \in \text{Type}_\Phi \text{ such that } \sigma \not\triangleright \Phi \text{ and } \tau \leq \sigma\}.$

Here is how we proceed with the definition of x :

- (1) We emulate the definition of \tilde{x} , using the polymorphic monotonicity calculus instead of the monomorphic one.

- (2) Step (1) will cause $\langle\langle\Phi\rangle\rangle_{x;\infty}$ to introduce more protectors than $\langle\langle\Phi\rangle\rangle_{\infty;\bar{x}}$. This is because the former protects all ground types τ that are instances of types $\sigma \in \text{Type}_{\Phi}$ such that $\sigma \not\triangleright \Phi$ (namely, T_{Φ}), whereas the latter protects all ground types τ that are instances of types in Type_{Φ} and satisfy $\tau \not\triangleright \langle\langle\Phi\rangle\rangle_{\infty}$ (namely, S_{Φ}). We have $S_{\Phi} \subseteq T_{\Phi}$ but generally not vice versa. To repair this mismatch, we add axioms that semantically eliminate the protectors for the types in T_{Φ} but not in S_{Φ} . To achieve this, we must characterise $T_{\Phi} - S_{\Phi}$ (which is an infinite set even for finite problems Φ) as the set of ground instances of a set U_{Φ} of types so that U_{Φ} is finite whenever Φ is.

We define U_{Φ} next:

Definition 6.19 (Cap). Given a set of types S , a *cap* for it is a set $S' \subseteq S$ such that all types in S are instances of types in S' . A cap is *minimal* if it contains no two distinct types σ and σ' such that $\sigma \leq \sigma'$. For each σ , let U_{σ} be a minimal cap of the set

$$U'_{\sigma} = \{\sigma' \leq \sigma \mid \sigma' \in \text{Inf}^*(\Phi) \text{ or there exists no } X^{\sigma''} \in \text{NV}(\Phi) \\ \text{such that } \sigma' \text{ and } \sigma'' \text{ have a common instance}\}$$

We let U_{Φ} be an arbitrary cap of $\bigcup\{U_{\sigma} \mid \sigma \in \text{Type}_{\Phi} \text{ and } \sigma \not\triangleright \Phi\}$.

The set U_{Φ} is both a subset of the monotonic types and a precise characterisation of the sets of types whose ground instances give our difference of interest, $T_{\Phi} - S_{\Phi}$.

Lemma 6.20. *The following properties hold:*

- (1) $S_{\Phi} \subseteq T_{\Phi}$;
- (2) If $\sigma' \in U'_{\sigma}$, then $\sigma' \triangleright \Phi$;
- (3) If $\sigma \in U_{\Phi}$, then $\sigma \triangleright \Phi$;
- (4) If $\tau \in T_{\Phi} - S_{\Phi}$, then $\tau \triangleright \langle\langle\Phi\rangle\rangle_{\infty}$;
- (5) $T_{\Phi} - S_{\Phi} = \{\tau \in \text{GType}_{\Sigma} \mid \text{there exists } \sigma \in U_{\Phi} \text{ such that } \tau \leq \sigma\}$.

Proof. (1): Assume $\tau \in S_{\Phi}$ and let $\sigma \geq \tau$ as in the definition of S_{Φ} . By Lemma 6.15(2), we have $\tau \not\triangleright \Phi$, and hence by Lemma 6.15(3) we have $\sigma \not\triangleright \Phi$, ensuring that $\tau \in T_{\Phi}$.

(2): It is clear that the condition defining U'_{σ} is a strengthening of that defining $\sigma \triangleright \Phi$.

(3): Immediate from (2).

(4): Immediate from the definitions of S_{Φ} and T_{Φ} .

(5): Let $\tau \in T_{\Phi}$ such that (A) $\tau \notin S_{\Phi}$. There exists σ such that $\tau \leq \sigma \in \text{Type}_{\Phi}$ and $\sigma \not\triangleright \Phi$. From (4), we have $\tau \triangleright \langle\langle\Phi\rangle\rangle_{\infty}$; hence, by Lemma 6.15(2), $\tau \triangleright \Phi$. Since τ is ground, either $\tau \in \text{Inf}^*(\Phi)$ or there exists no $X^{\sigma''} \in \text{NV}(\Phi)$ such that $\tau \leq \sigma''$, implying $\tau \in U'_{\sigma}$. It follows that τ is an instance of an element of U_{σ} , hence of an element of U_{Φ} , as desired. Now, assume τ is ground such that $\tau \leq \sigma' \in U_{\Phi}$. We obtain $\sigma \in \text{Type}_{\Phi}$ such that $\sigma \not\triangleright \Phi$ and $\sigma' \in U_{\sigma}$. In particular, we have $\sigma' \leq \sigma$, hence $\tau \in T_{\Phi}$. Moreover, by $\sigma' \in U_{\sigma}$ and (2), we have $\sigma' \triangleright \Phi$; hence by Lemma 6.15(1), we have $\tau \triangleright \langle\langle\Phi\rangle\rangle_{\infty}$, implying $\tau \notin S_{\Phi}$, as desired. \square

The following results will hold for any choice of set V_{Φ} between U_{Φ} and $\{\sigma \in \text{Type}_{\Sigma} \mid \sigma \triangleright \Phi\}$. The smaller the chosen set, the lighter the encoding.

Convention 6.21. We fix V_{Φ} such that $U_{\Phi} \subseteq V_{\Phi} \subseteq \{\sigma \in \text{Type}_{\Sigma} \mid \sigma \triangleright \Phi\}$.

6.5. Monotonicity-Based Type Tags. We are now equipped to present the definitions of the polymorphic monotonicity-based encodings and prove their correctness. The polymorphic $t?$ encoding can be seen as a hybrid between traditional tags (t) and monomorphic lightweight tags (\tilde{t}): as in t , tags take the form of a function $t\langle\sigma\rangle(t)$; as in \tilde{t} , tags are omitted for types that are inferred monotonic.

The main novelty concerns the typing axioms. The \tilde{t} encoding omits all typing axioms for monotonic types. In the polymorphic case, the monotonic type σ might be an instance of a more general, potentially nonmonotonic type for which tags are generated. For example, if α is tagged (because it is possibly nonmonotonic) but its instance $list(\alpha)$ is not (because it is infinite and hence monotonic), there will be mismatches between tagged and untagged terms. Our solution is to add the typing axiom $t\langle list(\alpha)\rangle(Xs) \approx Xs$, which allows the prover to add or remove a tag for the infinite type $list(\alpha)$. Such an axiom is sound for any monotonic type.

Definition 6.22 (Lightweight Tags $t?$). The encoding $t?$ translates polymorphic problems over $\Sigma = (Type, \mathcal{F}, \mathcal{P})$ to polymorphic problems over $(Type, \mathcal{F} \uplus \{t : \forall\alpha. \alpha \rightarrow \alpha\}, \mathcal{P})$. Its term and formula translations are defined as follows:

$$\llbracket f\langle\sigma\rangle(\bar{t})^\sigma \rrbracket_{t?} = \llbracket f\langle\sigma\rangle(\llbracket \bar{t} \rrbracket_{t?}) \rrbracket \quad \llbracket X^\sigma \rrbracket_{t?} = \llbracket X \rrbracket \quad \text{with } \llbracket t^\sigma \rrbracket = \begin{cases} t & \text{if } \sigma \triangleright \Phi \\ t\langle\sigma\rangle(t) & \text{otherwise} \end{cases}$$

The encoding adds the following typing axioms:

$$\forall TVars(\sigma). \forall X : \sigma. t\langle\sigma\rangle(X^\sigma) \approx X^\sigma \quad \text{for } \sigma \in V_\Phi$$

The *polymorphic lightweight type tags* encoding $t?$ is the composition $\llbracket \cdot \rrbracket_{t?; \mathbf{a}; \mathbf{e}}$. It translates a polymorphic problem over Σ into an untyped problem over $\Sigma' = (\mathcal{F}' \uplus \{t^2\}, \mathcal{P}')$, where \mathcal{F}' , \mathcal{P}' are as for \mathbf{a} .

Example 6.23. The $t?$ encoding of of Example 2.10 follows:

$$\begin{aligned} & \forall A, Xs. t\langle list(A)\rangle(Xs) \approx Xs \\ & \forall A, X, Xs. nil(A) \not\approx cons(A, t(A, X), Xs) \\ & \forall A, Xs. Xs \approx nil(A) \vee (\exists Y, Ys. Xs \approx cons(A, t(A, Y), Ys)) \\ & \forall A, X, Xs. t(A, hd(A, cons(A, t(A, X), Xs))) \approx t(A, X) \wedge \\ & \quad t(A, cons(A, t(A, X), Xs)) \approx Xs \\ & \exists X, Y, Xs, Ys. cons(b, t(b, X), Xs) \approx cons(b, t(b, Y), Ys) \wedge \\ & \quad (t(b, X) \not\approx t(b, Y) \vee Xs \not\approx Ys) \end{aligned}$$

The typing axiom allows any term to be typed as $list(\alpha)$, which is sound because $list(\alpha)$ is infinite. It would have been equally correct to provide separate axioms for nil , $cons$, and tl . Either way, the axioms are needed to remove the $t(A, X)$ tags in case the proof requires reasoning about $list(list(\alpha))$.

The lighter encoding $t??$ protects only naked variables and introduces equations of the form $t\langle\sigma\rangle(f\langle\bar{\alpha}\rangle(\bar{X})) \approx f\langle\bar{\alpha}\rangle(\bar{X})$ to add or remove tags around each function symbol f of a possibly nonmonotonic type σ , and similarly for existential variables.

Definition 6.24 (Featherweight Tags $t??$). The encoding $t??$ translates polymorphic problems over $\Sigma = (Type, \mathcal{F}, \mathcal{P})$ to polymorphic problems over $(Type, \mathcal{F} \uplus \{t : \forall\alpha. \alpha \rightarrow \alpha\}, \mathcal{P})$.

Its term and formula translations are defined as follows:

$$\begin{aligned} \langle\langle t_1 \approx t_2 \rangle\rangle_{t??} &= [\langle\langle t_1 \rangle\rangle_{t??}] \approx [\langle\langle t_2 \rangle\rangle_{t??}] \\ \langle\langle \exists X : \sigma. \varphi \rangle\rangle_{t??} &= \exists X : \sigma. \begin{cases} \langle\langle \varphi \rangle\rangle_{t??} & \text{if } \sigma \triangleright \Phi \\ t(\sigma)(X) \approx X \wedge \langle\langle \varphi \rangle\rangle_{t??} & \text{otherwise} \end{cases} \end{aligned}$$

with

$$[t^\sigma] = \begin{cases} t & \text{if } \sigma \triangleright \Phi \text{ or } t \notin \mathcal{V}_{\text{typed}}^\vee \\ t(\sigma)(t) & \text{otherwise} \end{cases}$$

The encoding adds the typing axioms of $t?$ (from Definition 6.22) and the following:

$$\begin{aligned} \forall \bar{\alpha}. \forall \bar{X} : \bar{\sigma}. t(\sigma)(f(\bar{\alpha})(\bar{X}^{\bar{\sigma}})) \approx f(\bar{\alpha})(\bar{X}^{\bar{\sigma}}) & \text{ for } f : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \sigma \in \mathcal{F} \text{ such that } \sigma \not\triangleright \Phi \\ \forall \text{TVars}(\sigma). \exists X : \sigma. t(\sigma)(X^\sigma) \approx X^\sigma & \text{ for } \sigma \in \text{Type}_\Phi \text{ such that } \sigma \not\triangleright \Phi \text{ and } \sigma \text{ is not} \\ & \text{an instance of the result type of some } f \in \mathcal{F} \end{aligned}$$

The *polymorphic featherweight type tags* encoding $t??$ is the composition $\langle\langle \rangle\rangle_{t??; \mathbf{a}; \mathbf{e}}$. The target signature of $t??$ is the same as that of $t?$.

Example 6.25. The $t??$ encoding of Example 2.10 requires fewer tags than $\tilde{t}?$, at the cost of two additional typing axioms and two typing equations for the existential variables of type b :

$$\begin{aligned} \forall A, Xs. t(A, \text{hd}(A, Xs)) \approx \text{hd}(A, Xs) \\ \forall A, Xs. t(\text{list}(A), Xs) \approx Xs \\ \forall A. \exists X. t(A, X) \approx X \\ \forall A, X, Xs. \text{nil}(A) \not\approx \text{cons}(A, X, Xs) \\ \forall A, Xs. Xs \approx \text{nil}(A) \vee (\exists Y, Ys. t(A, Y) \approx Y \wedge Xs \approx \text{cons}(A, Y, Ys)) \\ \forall A, X, Xs. \text{hd}(A, \text{cons}(A, X, Xs)) \approx t(A, X) \wedge \text{tl}(A, \text{cons}(A, X, Xs)) \approx Xs \\ \exists X, Y, Xs, Ys. t(b, X) \approx X \wedge t(b, Y) \approx Y \wedge \text{cons}(b, X, Xs) \approx \text{cons}(b, Y, Ys) \wedge \\ (X \not\approx Y \vee Xs \not\approx Ys) \end{aligned}$$

Theorem 6.26 (Correctness of $t?$, $t??$). *The polymorphic type tags encodings $t?$ and $t??$ are correct.*

Proof. First we discuss the case of $t?$. The following property can be routinely checked:

$$(A) \ \langle\langle \Phi \rangle\rangle_{t?; \infty} = \{\varphi^\# \mid \varphi \in \langle\langle \Phi \rangle\rangle_{\infty; \tilde{t}??}\} \cup \{\forall X^\tau. t_\tau(X^\tau) \approx X^\tau \mid \tau \in \text{GInst}(V_\Phi)\}$$

where $\text{GInst}(V_\Phi)$ denotes the set of ground instances of types in V_Φ and $\varphi^\#$ denotes the modification of φ obtained by adding, for each $\tau \in T_\Phi - S_\Phi$, a tag t_τ around every term of type τ . Recall that $U_\Phi \subseteq V_\Phi \subseteq \{\sigma \in \text{Type}_\Sigma \mid \sigma \triangleright \Phi\}$. We have the following:

- (1) $T_\Phi - S_\Phi \subseteq \text{GInst}(V_\Phi)$;
- (2) $\text{GInst}(V_\Phi) \subseteq \{\sigma \in \text{GType}_\Sigma \mid \sigma \triangleright \Phi\}$;
- (3) if $\tau \in V_\Phi$, the symbol t_τ does not occur in $\langle\langle \rangle\rangle_{\infty; \tilde{t}??}$.

Property (1) follows from Lemma 6.20(5), (2) from Lemma 6.15(3), and (3) from (2) and the definition of $\langle\langle \rangle\rangle_{\tilde{t}??}$. Now we are ready to check the conditions of Lemma 6.17.

SOUND: Given a model \mathcal{M} of $\langle\langle \Phi \rangle\rangle_{\infty; \tilde{t}??}$, we modify it into a structure \mathcal{M}' by reinterpreting, for all $\tau \in \text{GInst}(V_\Phi)$, the symbols t_τ as identity. Thanks to (3), \mathcal{M}' is still a model of $\langle\langle \Phi \rangle\rangle_{\infty; \tilde{t}??}$ and, thanks to (A) and (1), it is also a model of $\langle\langle \Phi \rangle\rangle_{t?; \infty}$.

COMPLETE: Any model of $\langle\langle\Phi\rangle\rangle_{t?;\infty}$ is also a model of $\langle\langle\Phi\rangle\rangle_{\infty;\tilde{t}?$, since, thanks to (1), in the presence of $\{\forall X^\tau. t_\tau(X^\tau) \approx X^\tau \mid \tau \in \text{GInst}(V_\Phi)\}$ each $\varphi^\#$ is equivalent to φ .

The case of $t??$ is similar, with the following modifications:

$$(A) \quad \langle\langle\Phi\rangle\rangle_{t?;\infty} = \{\varphi^\# \mid \varphi \in \langle\langle\Phi\rangle\rangle_{\infty;\tilde{t}?\} \cup \{\forall X^\tau. t_\tau(X^\tau) \approx X^\tau \mid \tau \in \text{GInst}(V_\Phi)\} \cup Ax$$

where Ax consists of all ground instances $\varphi\rho$ of the additional axioms $\forall\bar{\alpha}. \varphi$ from Definition 6.24 where the occurring t_τ is such that $\tau \in T_\Phi - S_\Phi$. What we retain about Ax is that Ax is satisfied by any structure that interprets as identity each t_τ with $\tau \in T_\Phi - S_\Phi$, and hence, by (1), we have

$$(4) \quad Ax \text{ is satisfied by any structure that interprets as identity each } t_\tau \text{ with } \tau \in \text{GInst}(V_\Phi).$$

Moreover, $\varphi^\#$ is modified to add tags to φ in fewer places—not to arbitrary terms, but only to naked universal variables. Then the proof for $t?$ works here too, additionally invoking (4) in the proof of soundness. \square

6.6. Monotonicity-Based Type Guards. Analogously to $t?$, the $g?$ encoding is best understood as a hybrid between traditional guards (g) and monomorphic lightweight guards ($\tilde{g}?$): as in g , guards take the form of a predicate $g\langle\sigma\rangle(t)$; as in $\tilde{g}?$, guards are omitted for types that are inferred monotonic.

Once again, the main novelty concerns the typing axioms. The $\tilde{g}?$ encoding omits all typing axioms for monotonic types. In the polymorphic case, the monotonic type σ might be an instance of a more general, potentially nonmonotonic type for which guards are generated. Our solution is to add the typing axiom $g\langle\sigma\rangle(X)$, which allows the prover to discharge any guard for the monotonic type σ .

Definition 6.27 (Lightweight Guards $g?$). The encoding $g?$ translates polymorphic problems over $\Sigma = (\text{Type}, \mathcal{F}, \mathcal{P})$ to polymorphic problems over $(\text{Type}, \mathcal{F}, \mathcal{P} \uplus \{g : \forall\alpha. \alpha \rightarrow o\})$. Its term and formula translations $\langle\langle \rangle\rangle_{g?;a;e}$ are defined as follows:

$$\begin{aligned} \langle\langle \forall X : \sigma. \varphi \rangle\rangle_{g?} &= \forall X : \sigma. \begin{cases} \langle\langle \varphi \rangle\rangle_{g?} & \text{if } \sigma \triangleright \Phi \\ g\langle\sigma\rangle(X^\sigma) \rightarrow \langle\langle \varphi \rangle\rangle_{g?} & \text{otherwise} \end{cases} \\ \langle\langle \exists X : \sigma. \varphi \rangle\rangle_{g?} &= \exists X : \sigma. \begin{cases} \langle\langle \varphi \rangle\rangle_{g?} & \text{if } \sigma \triangleright \Phi \\ g\langle\sigma\rangle(X^\sigma) \wedge \langle\langle \varphi \rangle\rangle_{g?} & \text{otherwise} \end{cases} \end{aligned}$$

The encoding adds the following typing axioms:

$$\begin{aligned} \forall \text{TVars}(\sigma). \forall X : \bar{\sigma}. g\langle\sigma\rangle(X^\sigma) & \quad \text{for } \sigma \in V_\Phi \\ \forall \bar{\alpha}. \forall \bar{X} : \bar{\sigma}. g\langle\sigma\rangle(f(\bar{\alpha})(\bar{X}^{\bar{\sigma}})) & \quad \text{for } f : \forall \bar{\alpha}. \bar{\sigma} \rightarrow \sigma \in \mathcal{F} \text{ such that } \sigma \not\triangleright \Phi \\ \forall \text{TVars}(\sigma). \exists X : \sigma. g\langle\sigma\rangle(X^\sigma) & \quad \text{for } \sigma \in \text{Type}_\Phi \text{ such that } \sigma \not\triangleright \Phi \text{ and } \sigma \text{ is not} \\ & \quad \text{an instance of the result type of some } f \in \mathcal{F} \end{aligned}$$

The *polymorphic lightweight type guards* encoding $g?$ is the composition $\langle\langle \rangle\rangle_{g?;a;e}$. It translates a polymorphic problem over Σ into an untyped problem over $\Sigma' = (\mathcal{F}', \mathcal{P}' \uplus \{g^2\})$, where \mathcal{F}' , \mathcal{P}' are as for a .

The featherweight cousin is a straightforward generalisation of $g?$ along the lines of the generalisation of $\tilde{g}?$ into $\tilde{g}??$.

Definition 6.28 (Featherweight Guards $g??$). The *polymorphic featherweight type guards* encoding $g??$ is identical to the lightweight encoding $g?$ except that the condition “if $\sigma \triangleright \Phi$ ” in the \forall case is weakened to “if $\sigma \triangleright \Phi$ or $X^\sigma \notin \text{NV}(\varphi)$ ”.

Example 6.29. The $g??$ encoding of Example 2.10 follows:

$$\begin{aligned} & \forall A, Xs. g(A, \text{hd}(A, Xs)) \\ & \forall A, Xs. g(\text{list}(A), Xs) \\ & \forall A, X, Xs. \text{nil}(A) \not\approx \text{cons}(A, X, Xs) \\ & \forall A, Xs. Xs \approx \text{nil}(A) \vee (\exists Y, Ys. g(A, Y) \wedge Xs \approx \text{cons}(A, Y, Ys)) \\ & \forall A, X, Xs. g(A, X) \rightarrow \text{hd}(A, \text{cons}(A, X, Xs)) \approx X \wedge \text{tl}(A, \text{cons}(A, X, Xs)) \approx Xs \\ & \exists X, Y, Xs, Ys. g(b, X) \wedge g(b, Y) \wedge \text{cons}(b, X, Xs) \approx \text{cons}(b, Y, Ys) \wedge (X \not\approx Y \vee Xs \not\approx Ys) \end{aligned}$$

Theorem 6.30 (Correctness of $g?$, $g??$). *The polymorphic type guards encodings $g?$ and $g??$ are correct.*

Proof. The argument is very similar to that for tags (Theorem 6.26), with the following modifications:

$$(A) \llbracket \Phi \rrbracket_{t?, \infty} = \{\varphi^\# \mid \varphi \in \llbracket \Phi \rrbracket_{\infty, \tilde{t}?, \infty}\} \cup \{\forall X^\tau. g_\tau(X^\tau) \mid \tau \in \text{GInst}(V_\Phi)\} \cup Ax$$

where Ax consists of a set of formulae which is satisfied by each structure that interprets as everywhere-true each g_τ with $\tau \in \text{GInst}(V_\Phi)$ and $\varphi^\#$ denotes the modification of φ obtained by adding, for each $\tau \in T_\Phi - S_\Phi$, a guard g_τ on several negative positions in φ .

Apart from this, the proofs for $g?$ and $g??$ are identical to that for $t??$, except that instead of interpreting tags as identity we interpret guards as everywhere-true. \square

7. IMPLEMENTATION

Our research on polymorphic type encodings was driven by Sledgehammer, a component of Isabelle/HOL that harnesses first-order automatic theorem provers to discharge interactive proof obligations. The tool heuristically selects hundreds of background facts, translates them to untyped or monomorphic first-order logic, invokes the external provers in parallel, and reconstructs machine-generated proofs in Isabelle.

All the encodings presented in this article except for the complete monomorphisation translation from Section 6.1 (which plays only an auxiliary theoretical role in our development) enjoy the desirable property that if the source signature and problem are finite, then the target signature and problem are also finite. All of these encodings, including the traditional ones, are implemented in Sledgehammer and can be used to target external first-order provers. The rest of this section considers implementation issues in more detail.

7.1. Heuristic Monomorphisation Algorithm. The monomorphisation algorithm implemented in Sledgehammer translates a polymorphic problem into a monomorphic problem by heuristically instantiating type variables. It involves three stages:

1. Separate the monomorphic and the polymorphic formulae, and collect all symbols occurring in the monomorphic formulae (the “mono-symbols”).
2. For each polymorphic axiom, stepwise refine a set of substitutions, starting from the singleton set containing only the empty substitution, by matching known mono-symbols against their polymorphic counterparts in the axiom. So long as new mono-symbols emerge, collect them and repeat this stage.

3. Apply the substitutions to the corresponding polymorphic formulae. Only keep fully monomorphic formulae.

To ensure termination, the iterations performed in stage 2 are limited to a configurable number K . To curb the exponential growth, the algorithm also enforces an upper bound Δ on the number of new formulae. Sledgehammer operates with $K = 3$ and $\Delta = 200$ by default, so that a problem with 500 formulae comprises at most 700 formulae after monomorphisation. Experiments found these values suitable. Given formulae about b and $list(\alpha)$, the third iteration already generates $list(list(list(b)))$ instances; adding yet another layer of $list$ is unlikely to help. Increasing Δ sometimes helps solve more problems, but its potential for clutter is real.

7.2. Extension to Higher-Order Logic. Isabelle/HOL’s logic, polymorphic higher-order logic with axiomatic type classes [35], is not the same as the polymorphic first-order logic considered in this article. Sledgehammer’s translation is a three-step process, where the first and last step may be omitted, depending on whether monomorphisation is desired and whether the target prover supports monomorphic types:

1. *Optionally monomorphise the problem.*
2. *Eliminate the higher-order constructs* [22, §2.1]. λ -abstractions are rewritten to SK combinators or to supercombinators (λ -lifting). Functions are passed varying numbers of arguments via an apply operator $\mathbf{hAPP} : \forall\alpha, \beta. fun(\alpha, \beta) \times \alpha \rightarrow \beta$ (where fun is uninterpreted). Boolean terms are converted to formulae using a unary predicate $\mathbf{hBOOL} : bool \rightarrow o$ (where $bool$ is uninterpreted).
3. *Encode the type information.* Polymorphic types are encoded using the techniques described in this article. Type classes are essentially sets of types; they are encoded as polymorphic predicates $\forall\alpha. o$ (where α is a phantom type variable, Definition 3.12). For example, a predicate $\mathbf{linorder} : \forall\alpha. o$ could be used to restrict the axioms specifying that $\mathbf{less_eq} : \forall\alpha. \alpha \times \alpha \rightarrow o$ is a linear order to those types that satisfy the $\mathbf{linorder}$ predicate (cf. Example 3.15). The type class hierarchy is expressible as Horn clauses [22, §2.3].

The symbol \mathbf{hAPP} would hugely burden problems if it were introduced systematically for all arguments to functions. To reduce clutter, Sledgehammer computes the minimum arity n needed for each symbol and passes the first n arguments directly, falling back on \mathbf{hAPP} for additional arguments. In general, more arguments can be passed directly if monomorphisation is performed before \mathbf{hAPP} is introduced, because each monomorphic instance of a polymorphic symbol is considered individually. Similar observations can be made for \mathbf{hBOOL} .

7.3. Infinite Types and Constructors. The monotonicity calculus \triangleright is parameterised by a set $\text{Inf}(\Phi)$ of infinite types. One could employ an approach similar to that implemented in `Infinox` [15] to automatically infer finite unsatisfiability of types. This tool relies on various proof principles to show that a set of untyped first-order formulae only has models with infinite domains. For example, it can infer that $list_b$ is infinite in Example 2.13 because \mathbf{cons}_b is injective in its second argument but not surjective. However, in a proof assistant such as Isabelle, it is simpler to exploit meta-information available through introspection. Isabelle’s datatypes are registered with their constructors; if some of them are recursive, or take an argument of an infinite type, the datatype must be infinite and hence monotonic.

More specifically, the monotonicity inference is run on the entire problem and maintains two finite sets of polymorphic types: the surely infinite types J and the possibly nonmonotonic types N . Every type of a naked variable in the problem is tested for infinity. If the test succeeds, the type is inserted into J ; otherwise, it is inserted into N . Simplifications are performed: there no need to insert σ to J or N if it is an instance of a type already in the set; when inserting σ to a set, it is safe to remove any type in the set that is an instance of σ . The monotonicity check then becomes

$$\sigma \triangleright \Phi \iff (\exists \tau \in J. \exists \rho. \sigma = \tau\rho) \vee (\forall \tau \in N. \nexists \rho. \sigma\rho = \tau\rho)$$

7.4. Proof Reconstruction. To guard against bugs in the external provers, Sledgehammer reconstructs machine-generated proofs in Isabelle. This is usually accomplished by the *metis* proof method [26], supplying it with the short list of facts referenced in the proof found by the prover. The proof method is based on the Metis prover [20], a complete resolution prover for untyped first-order logic. The *metis* call is all that remains from the Sledgehammer invocation in the Isabelle theory, which can then be replayed without external provers. Given only a handful of facts, *metis* usually succeeds within milliseconds.

Prior to our work, a large share of the reconstruction failures were caused by type-unsound proofs found by the external provers, due to the use of the unsound encoding a [13, §4.1]. We now replaced the internals of Sledgehammer and *metis* so that they use a translation module supporting all the type encodings described in this article.

Nonetheless, despite the typing information, individual inferences in Metis can be ill-typed when types are reintroduced, causing the *metis* proof method to fail. There are two main failure scenarios.

First, the prover may in principle instantiate variables with “ill-typed” terms at any point in the proof. Fortunately, this hardly ever arises in practice, because like other resolution provers Metis heavily restricts paramodulation from and into variables [1].

An issue that is more likely to plague users concerns the infinite types $\text{Inf}(\Phi)$. In the theoretical part of the paper, we required infinity to be a consequence of the problem Φ . The implementation is less rigorous; it will happily treat types that are known to be infinite in the Isabelle background theories even if Φ itself does not imply infinity of the types. For example, assuming *nat* is known to be infinite, the implementation of the monotonicity-based encodings will not introduce any protectors around the naked variables M and N when translating the problem

$$\text{on} \not\approx \text{off}^{\text{state}} \wedge (\forall X, Y : \text{nat}. X \approx Y)$$

(where the second conjunct is presumably the negation of a conjecture stating that there exist two distinct natural numbers). That problem is satisfiable on its own but unsatisfiable with respect to the background theory. Untyped provers will instantiate X and Y with *on* and *off* to derive a contradiction; and no “type-sound” proof is possible unless we also provide characteristic theorems for *nat*. In general, we would need to provide infinity axioms for all types in $\text{Inf}(\Phi)$ to make the encoding sound irrespective of the background theory; for example:

$$\forall N : \text{nat}. \text{zero} \not\approx \text{suc}(N) \qquad \forall M, N : \text{nat}. \text{suc}(M) \approx \text{suc}(N) \rightarrow M \approx N$$

Although this now makes a sound proof possible (by instantiating X and Y with `zero` and `suc(zero)`), it does not prevent the prover from discovering the spurious proof with `on` and `off`, which cannot be reconstructed by *metis*.

Although the above scenarios rarely occur in practice, it would be more satisfactory if proof reconstruction were always possible. A solution would be to connect our formalised soundness proofs with a verified checker for untyped first-order proofs. This remains for future work.

8. EVALUATION

To evaluate the type encodings described in this article, we put together two sets of 1000 polymorphic first-order problems originating from 10 existing Isabelle theories, translated with Sledgehammer’s help (100 problems per theory).³ Nine of the theories are the same as in a previous evaluation [3]; the tenth one is an optimality proof for Huffman’s algorithm. Our test data are publicly available [4].

The problems in the first benchmark set include about 50 heuristically selected facts (before monomorphisation); that number is increased to 500 for the second set, to reveal how well the encodings scale with the problem size.

We evaluated each type encoding with five modern automatic theorem provers: the resolution provers E 1.8 [28], SPASS 3.8ds [9], and Vampire 3.0 (revision 1803) [27] and the SMT solvers Alt-Ergo 0.95.2 [10] and Z3 4.3.2 (revision 944dfee008) [23]. To make the evaluation more informative, we also tested the provers’ native support for types where it is available; it is referred to as \tilde{n} (monomorphic) and n (polymorphic). Only Alt-Ergo supports polymorphic types natively.

Each prover was invoked with the set of options we had previously determined worked best for Sledgehammer.⁴ The provers were granted 15 seconds of CPU time per problem on one core of a 3.06 GHz Dual-Core Intel Xeon processor. Most proofs were found within a few seconds; a higher time limit would have had a restricted impact on the success rate [13, §4]. To avoid giving the unsound encodings (`e` and `a`) an unfair advantage, for these proof search was followed by a certification phase that attempted to re-find the proof using a combination of sound encodings, based on its referenced facts. This phase slightly penalises the unsound encodings by rejecting a few sound proofs, but such is the price of unsoundness in practice.

Figures 2 and 3 give, for each combination of prover and encoding, the number of solved problems from each problem set. Rows marked with \sim concern the monomorphic encodings. The encodings \tilde{a} , $\tilde{t}@$, and $\tilde{g}@$ are omitted; the first two coincide with \tilde{e} , whereas $\tilde{t}@$ and $\tilde{g}@$ are identical to degenerate versions of $\tilde{t}??$ and $\tilde{g}??$ that treat all types as possibly nonmonotonic. We observe the following:

- Among the encodings to untyped first-order logic, the monomorphic monotonicity-based encodings (especially $\tilde{g}??$ but also $\tilde{t}??$, $\tilde{g}?$, and $\tilde{t}?$) performed best overall. Their performance is close to that of the provers’ native support for simple types (\tilde{n}). Polymorphic encodings lag behind; this is likely due in part to the synergy between the monomorphiser and the translation of higher-order constructs (cf. Section 7.2).

³The TPTP benchmark suite [30], which is customarily used for evaluating theorem provers, has just begun collecting polymorphic (TFF1) problems.

⁴The setup for E was suggested by Stephan Schulz and includes the little known “symbol offset” weight function. We ran Alt-Ergo with the default setup, SPASS in Isabelle mode, Vampire in CASC mode, and Z3 through the `z3_tptp` wrapper but otherwise with the default setup.

	e	a	t	t?	t??	t@	g	g?	g??	g@	n
Alt-Ergo	268	275	243	218	296	272	233	291	295	259	301
~	293	–	245	293	301	–	296	303	304	–	302
E	319	330	322	311	325	269	268	325	336	288	–
~	338	–	337	353	343	–	335	351	352	–	–
SPASS	289	316	290	275	319	196	223	302	314	244	–
~	322	–	322	328	332	–	317	330	337	–	349
Vampire	328	335	291	304	320	243	281	320	328	301	–
~	340	–	331	355	351	–	329	356	354	–	370
Z3	295	323	299	245	317	311	289	301	317	295	–
~	326	–	307	320	341	–	342	344	344	–	343

Figure 2: Number of solved problems with 50 facts

	e	a	t	t?	t??	t@	g	g?	g??	g@	n
Alt-Ergo	268	218	183	223	273	216	157	240	273	190	260
~	285	–	209	302	311	–	251	311	313	–	318
E	170	355	261	271	353	193	205	345	355	237	–
~	376	–	331	388	393	–	316	388	390	–	–
SPASS	136	321	289	259	299	154	163	278	296	194	–
~	333	–	287	323	333	–	229	327	334	–	343
Vampire	318	396	164	286	365	230	219	312	349	259	–
~	406	–	231	384	406	–	314	406	405	–	440
Z3	247	358	253	241	362	273	213	305	363	270	–
~	363	–	260	358	369	–	349	370	369	–	370

Figure 3: Number of solved problems with 500 facts

	e	a	t	t?	t??	t@	g	g?	g??	g@
Clauses	89	99	100	108	140	167	166	139	139	166
~	125	–	127	127	141	–	242	141	141	–
Literals per clause	2.3	2.4	2.4	2.3	2.2	2.5	4.3	3.2	2.6	3.8
~	2.3	–	2.3	2.3	2.2	–	4.4	2.5	2.4	–
Symbols per atom	6.3	8.0	18.3	16.0	10.3	9.9	5.7	8.0	8.6	5.7
~	6.2	–	10.6	7.0	6.2	–	4.3	5.5	5.7	–
Symbols	1276	1870	4339	3924	3235	4070	4051	3609	3103	3610
~	1757	–	3060	2040	1935	–	4548	1951	1904	–

Figure 4: Average size of classified problems with 50 facts

- Among the polymorphic encodings, our novel cover-based and monotonicity-based encodings (t@, t?, t??, g@, g?, and g??), with the exception of t@, constitute a substantial improvement over the traditional sound schemes (t and g).

- As suggested in the literature, there is no clear winner between tags and guards. We expected monomorphic guards to be especially effective with SPASS, since they are internally mapped to soft sorts (an optimised representation of unary predicates [34]), but this is not corroborated by the data.
- Despite the proof reconstruction phase, the unsound encoding \mathbf{a} achieved similar results to the sound polymorphic encodings. In contrast, its monomorphic cousin $\tilde{\mathbf{e}}$ is generally no match for the sound monomorphic schemes.
- Oddly, the polymorphic prover Alt-Ergo performs significantly better on monomorphised problems than on the corresponding polymorphic ones. This raises serious doubts about the quality of the prover’s heuristics for instantiating type variables.

For the first benchmark set, Figure 4 presents the average number of clauses, literals per clause, symbols per atom, and symbols for clausified problems (using E’s clausifier), to give an idea of each encoding’s overhead. The striking point is the lightness of the monomorphic encodings, as witnessed by the number of symbols. Because monomorphisation generates several copies of the same formulae, we could have expected it to lead to larger problems, but this underestimates the cost of encoding types as terms in the polymorphic encodings.⁵ The strong correlation between the success rates in Figure 3 and the average number of symbols in Figure 4 confirms the expectation that clutter (whether type arguments, tags, or guards) slows down automatic provers.

Independently of these empirical results, the new type encodings made an impact at the 2012 edition of CASC, the annual automatic prover competition [31]. Isabelle’s automatic proof tools, including Sledgehammer, compete against the automatic provers LEO-II, Satalax, and TPS in the higher-order division. Largely thanks to the new schemes (but also to improvements in the underlying first-order provers), Isabelle moved from the third place it had occupied since 2009 to the first place.

9. RELATED WORK

The earliest descriptions of type tags and type guards we are aware of are due to Enderton [18, §4.3] and Stickel [29, p. 99]. Wick and McCune [36, §4] compare type arguments, tags, and guards in a monomorphic setting. Type arguments are described by Meng and Paulson [22], who also consider full type erasure and polymorphic type tags and present a translation of axiomatic type classes. As part of the MPTP project, Urban [33] extended the untyped TPTP FOF syntax with dependent types to accommodate Mizar and designed translations to plain FOF.

The intermediate verification language and tool Boogie 2 [21] supports a restricted form of higher-rank polymorphism (with polymorphic maps), and Why3 [11] provides rank-1 polymorphism. Both define translations to a monomorphic logic and rely on proxies to handle interpreted types [12, 17, 21]. One of the Boogie translations [21, §3.1] uses SMT triggers to prevent ill-typed instantiations in conjunction with type arguments; however, this approach is risky in the absence of a semantics for triggers. Bouillaguet et al. [14, §4]

⁵The increase visible in the \mathbf{e} column, from 89 to 125 clauses and from 1276 to 1757 symbols, is due to the multiple copies arising from monomorphisation. Even though these become identical after type erasure, they are counted as separate in the statistics, and it is up to the automatic provers to notice that they are the same.

showed that full type erasure is sound if all types can be assumed to have the same cardinality and exploit this in the verification system Jahob.

An alternative to encoding polymorphic types or monomorphising them away is to support them natively in the prover. This is ubiquitous in interactive theorem provers, but perhaps the only production-quality automatic prover that supports polymorphism is Alt-Ergo [10].

Blanchette and Krauss [6] studied monotonicity inferences for higher-order logic without polymorphism. Claessen et al. [16] were first to apply them to type erasure.

10. CONCLUSION

This article introduced a family of translations from polymorphic into untyped first-order logic, with a focus on efficiency. Our monotonicity-based encodings soundly erase all types that are inferred monotonic, as well as most occurrences of the remaining types. The best translations outperform the traditional encoding schemes.

We implemented the new translations in the Sledgehammer tool for Isabelle/HOL, thereby addressing a recurring user complaint. Although Isabelle certifies external proofs, unsound proofs are annoying and often conceal sound proofs. The same translation module forms the core of Isabelle’s TPTP exporter tool, which makes entire theorem libraries available to first-order reasoners. Our refinements to the monomorphic case have made their way into the Monotonox translator [16]. Applications such as Boogie [21], LEO-II [2], and Why3 [11] also stand to gain from lighter encodings.

The TPTP family recently welcomed the addition of TFF1 [7], an extension of the monomorphic TFF0 logic with rank-1 polymorphism. Equipped with a concrete syntax and translation tools, we can turn any popular automatic theorem prover into an efficient polymorphic prover. Translating the untyped proof back into a typed proof is usually straightforward, but there are important corner cases that call for more research. It should also be possible to extend our approach to interpreted arithmetic.

From both a conceptual and an implementation point of view, the encodings are all instances of a general framework, in which mostly orthogonal features can be combined in various ways. Defining such a large number of encodings makes it possible to select the most appropriate scheme for each automatic prover, based on empirical evidence. In fact, using strategy scheduling or parallelism, it is advantageous to have each prover employ a combination of encodings with complementary strengths.

ACKNOWLEDGEMENT

Koen Claessen and Tobias Nipkow made this collaboration possible. Lukas Bulwahn, Peter Lammich, Rustan Leino, Tobias Nipkow, Nir Piterman, Alexander Steen, Mark Summerfield, Tjark Weber, and several anonymous reviewers suggested dozens of textual improvements; in particular, the three reviewers of this extended, journal version accomplished a remarkable work, going well beyond the call of duty. Blanchette was partly supported by the Deutsche Forschungsgemeinschaft (DFG) projects Quis Custodiet and Hardening the Hammer (grants NI 491/11-2 and NI 491/14-1). Popescu was partly supported by the DFG project Security Type Systems and Deduction (grant NI 491/13-2) as part of the program Reliably Secure Software Systems (RS³, priority program 1496). The authors are listed alphabetically.

REFERENCES

- [1] Bachmair, L., Ganzinger, H., Lynch, C., Snyder, W.: Basic paramodulation. *Inf. Comput.* **121**(2), 172–192 (1995)
- [2] Benzmüller, C., Paulson, L.C., Theiss, F., Fietzke, A.: LEO-II—A cooperative automatic theorem prover for higher-order logic. In: A. Armando, P. Baumgartner, G. Dowek (eds.) *IJCAR 2008, LNAI*, vol. 5195, pp. 162–170. Springer (2008)
- [3] Blanchette, J.C., Böhme, S., Paulson, L.C.: Extending Sledgehammer with SMT solvers. *J. Autom. Reasoning* **51**(1), 109–128 (2013)
- [4] Blanchette, J.C., Böhme, S., Popescu, A., Smallbone, N.: Empirical data associated with this article. http://www21.in.tum.de/~blanchet/enc_types_data_lmcs.tar.gz (2012)
- [5] Blanchette, J.C., Böhme, S., Popescu, A., Smallbone, N.: Encoding monomorphic and polymorphic types. In: N. Piterman, S. Smolka (eds.) *TACAS 2013, LNCS*, vol. 7795, pp. 493–507. Springer (2013)
- [6] Blanchette, J.C., Krauss, A.: Monotonicity inference for higher-order formulas. *J. Autom. Reasoning* **47**(4), 369–398 (2011)
- [7] Blanchette, J.C., Paskevich, A.: TFF1: The TPTP typed first-order form with rank-1 polymorphism. In: M.P. Bonacina (ed.) *CADE-24, LNAI*, vol. 7898, pp. 414–420. Springer (2013)
- [8] Blanchette, J.C., Popescu, A.: Mechanizing the metatheory of Sledgehammer. In: P. Fontaine, C. Ringeissen, R.A. Schmidt (eds.) *FroCoS 2013, LNCS*, vol. 8152, pp. 245–260. Springer (2013)
- [9] Blanchette, J.C., Popescu, A., Wand, D., Weidenbach, C.: More SPASS with Isabelle—Superposition with hard sorts and configurable simplification. In: L. Beringer, A. Felty (eds.) *ITP 2012, LNCS*, vol. 7406, pp. 345–360. Springer (2012)
- [10] Bobot, F., Conchon, S., Contejean, E., Lescuyer, S.: Implementing polymorphism in SMT solvers. In: C. Barrett, L. de Moura (eds.) *SMT 2008* (2008)
- [11] Bobot, F., Filliâtre, J.C., Marché, C., Paskevich, A.: Why3: Shepherd your herd of provers. In: K.R.M. Leino, M. Moskal (eds.) *Boogie 2011*, pp. 53–64 (2011)
- [12] Bobot, F., Paskevich, A.: Expressing polymorphic types in a many-sorted language. In: C. Tinelli, V. Sofronie-Stokkermans (eds.) *FroCoS 2011, LNCS*, vol. 6989, pp. 87–102. Springer (2011)
- [13] Böhme, S., Nipkow, T.: Sledgehammer: Judgement Day. In: J. Giesl, R. Hähnle (eds.) *IJCAR 2010, LNAI*, vol. 6173, pp. 107–121. Springer (2010)
- [14] Bouillaguet, C., Kuncak, V., Wies, T., Zee, K., Rinard, M.: Using first-order theorem provers in the Jahob data structure verification system. In: B. Cook, A. Podelski (eds.) *VMCAI 2007, LNCS*, vol. 4349, pp. 74–88. Springer (2007)
- [15] Claessen, K., Lillieström, A.: Automated inference of finite unsatisfiability. *J. Autom. Reasoning* **47**(2), 111–132 (2011)
- [16] Claessen, K., Lillieström, A., Smallbone, N.: Sort it out with monotonicity—Translating between many-sorted and unsorted first-order logic. In: N. Bjørner, V. Sofronie-Stokkermans (eds.) *CADE-23, LNAI*, vol. 6803, pp. 207–221. Springer (2011)
- [17] Couchot, J.F., Lescuyer, S.: Handling polymorphism in automated deduction. In: F. Pfenning (ed.) *CADE-21, LNAI*, vol. 4603, pp. 263–278. Springer (2007)
- [18] Enderton, H.B.: *A Mathematical Introduction to Logic*. Academic Press (1972)
- [19] Hodges, W.: *Model Theory*. Cambridge University Press (1993)
- [20] Hurd, J.: First-order proof tactics in higher-order logic theorem provers. In: M. Archer, B. Di Vito, C. Muñoz (eds.) *Design and Application of Strategies/Tactics in Higher Order Logics*, no. CP-2003-212448 in NASA Tech. Reports, pp. 56–68 (2003)
- [21] Leino, K.R.M., Rümmer, P.: A polymorphic intermediate verification language: Design and logical encoding. In: J. Esparza, R. Majumdar (eds.) *TACAS 2010, LNCS*, vol. 6015, pp. 312–327. Springer (2010)
- [22] Meng, J., Paulson, L.C.: Translating higher-order clauses to first-order clauses. *J. Autom. Reasoning* **40**(1), 35–60 (2008)
- [23] de Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: C.R. Ramakrishnan, J. Rehof (eds.) *TACAS 2008, LNCS*, vol. 4963, pp. 337–340. Springer (2008)
- [24] Nipkow, T., Klein, G.: *Concrete Semantics—with Isabelle/HOL*. Springer (2014)
- [25] Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL: A Proof Assistant for Higher-Order Logic, *LNCS*, vol. 2283. Springer (2002)

- [26] Paulson, L.C., Susanto, K.W.: Source-level proof reconstruction for interactive theorem proving. In: K. Schneider, J. Brandt (eds.) TPHOLs 2007, *LNCS*, vol. 4732, pp. 232–245. Springer (2007)
- [27] Riazanov, A., Voronkov, A.: The design and implementation of Vampire. *AI Comm.* **15**(2-3), 91–110 (2002)
- [28] Schulz, S.: System description: E 0.81. In: D. Basin, M. Rusinowitch (eds.) IJCAR 2004, *LNAI*, vol. 3097, pp. 223–228. Springer (2004)
- [29] Stickel, M.E.: Schubert’s steamroller problem: Formulations and solutions. *J. Autom. Reasoning* **2**(1), 89–101 (1986)
- [30] Sutcliffe, G.: The TPTP problem library and associated infrastructure—The FOF and CNF parts, v3.5.0. *J. Autom. Reasoning* **43**(4), 337–362 (2009)
- [31] Sutcliffe, G.: Proceedings of the 6th IJCAR ATP system competition (CASC-J6). In: G. Sutcliffe (ed.) CASC-J6, *EPiC*, vol. 11, pp. 1–50. EasyChair (2012)
- [32] Sutcliffe, G., Schulz, S., Claessen, K., Baumgartner, P.: The TPTP typed first-order form with arithmetic. In: N. Bjørner, A. Voronkov (eds.) LPAR-18, *LNCS*, vol. 7180, pp. 406–419. Springer (2012)
- [33] Urban, J.: MPTP 0.2: Design, implementation, and initial experiments. *J. Autom. Reasoning* **37**(1-2), 21–43 (2006)
- [34] Weidenbach, C.: Combining superposition, sorts and splitting. In: A. Robinson, A. Voronkov (eds.) Handbook of Automated Reasoning, vol. 2, pp. 1965–2013. Elsevier (2001)
- [35] Wenzel, M.: Type classes and overloading in higher-order logic. In: E.L. Gunter, A. Felty (eds.) TPHOLs 1997, *LNCS*, vol. 1275, pp. 307–322. Springer (1997)
- [36] Wick, C.A., McCune, W.W.: Automated reasoning about elementary point-set topology. *J. Autom. Reasoning* **5**(2), 239–255 (1989)