Non-Parametric Parametricity

Georg Neis
MPI-SWS
neis@mpi-sws.org

Derek Dreyer
MPI-SWS
dreyer@mpi-sws.org

Andreas Rossberg
MPI-SWS
rossberg@mpi-sws.org

Abstract

Type abstraction and intensional type analysis are features seemingly at odds—type abstraction is intended to guarantee parametricity and representation independence, while type analysis is inherently non-parametric. Recently, however, several researchers have proposed and implemented "dynamic type generation" as a way to reconcile these features. The idea is that, when one defines an abstract type, one should also be able to generate at run time a fresh type name, which may be used as a dynamic representative of the abstract type for purposes of type analysis. The question remains: in a language with non-parametric polymorphism, does dynamic type generation provide us with the same kinds of abstraction guarantees that we get from parametric polymorphism?

Our goal is to provide a rigorous answer to this question. We define a step-indexed Kripke logical relation, with a novel form of "possible world", for a language with both non-parametric polymorphism and dynamic type generation. Our logical relation enables us to establish parametricity and representation independence results, even in a non-parametric setting, by attaching arbitrary relational interpretations to dynamically-generated type names. In addition, we explore how programs that are provably equivalent in a more traditional parametric logical relation may be "wrapped" systematically to produce terms that are related by our non-parametric relation, and vice versa. This leads us to a novel polarized form of our logical relation, which distinguishes between positive and negative notions of parametricity.

1. Introduction

When we say that a language supports *parametric polymorphism*, we mean that "abstract" types in that language are really abstract—that is, no client of an abstract type can guess or depend on its underlying implementation [18]. Traditionally, the parametric nature of polymorphism is guaranteed statically by the language's type system, thus enabling the so-called *type-erasure* interpretation of polymorphism by which type abstractions and instantiations are erased during compilation.

However, some modern programming languages include a useful feature that appears to be in direct conflict with parametric polymorphism, namely the ability to perform *intensional type analysis* [11]. Probably the simplest and most common instance of intensional type analysis is found in the implementation of languages supporting a type `Dynamic` [1]. In such languages, any value v may be cast to type `Dynamic`, but the cast *from* type `Dynamic` to any type τ requires a runtime check to ensure that v's actual type equals τ. Other languages such as Acute [23] and Alice ML [21], which are designed to support dynamic loading of modules, require the ability to check dynamically whether a module implements an expected interface, which in turn involves runtime inspection of the module's type components. There have also been a number of more experimental proposals for languages that employ a `typecase` construct to facilitate *polytypic* programming (*e.g.*, [30, 27]).

There is a fundamental tension between type analysis and type abstraction. If one can inspect the identity of an unknown type at run time, then the type is not really abstract, so any invariants concerning values of that type may be broken [30]. Consequently, languages with a type `Dynamic` sometimes prohibit programmers from casting to `Dynamic` any values whose types mention user-defined abstract types. However, this is a rather severe restriction, which effectively penalizes programmers for using type abstraction.

Thus, a number of researchers have proposed that languages with type analysis facilities should also support *dynamic type generation* [22, 19, 27, 20]. That is, when one defines an abstract type, one should also be able to generate at run time a "fresh" type name, which may be used as a dynamic representative of the abstract type for purposes of type analysis. (We will see a concrete example of this in Section 2.) The idea is that the freshness of name generation will ensure that user-defined abstract types are viewed dynamically in the same way that they are viewed statically—*i.e.*, as distinct from all other types.

The question remains: how do we know that dynamic type generation *works*? In a language with intensional type analysis—*i.e.*, *non-parametric* polymorphism—does dynamic type generation provably provide us with the same kinds of abstraction guarantees that we get from traditional parametric polymorphism?

Our goal is to provide a rigorous answer to this question. We study an extension of System F, supporting (1) a type-safe cast operator, which is essentially a variant of Girard's J operator [8], and (2) a facility for dynamic generation of fresh type names. For brevity, we will call this language **G**. As a practical language mechanism, the cast operator is somewhat crude in comparison to the more expressive `typecase`-style constructs proposed in the literature,[1] but it nonetheless renders polymorphism *non-parametric*. Our main technical result is that, in a language with non-parametric polymorphism, parametricity may be provably regained via judicious use of dynamic type generation.

The rest of the paper is structured as follows. In Section 2, we present our language under consideration, G, and also give an example to illustrate how dynamic type generation is useful. In Section 3, we explain informally the approach that we have developed for reasoning about G. Our approach employs a *step-indexed Kripke logical relation* with a novel form of possible world.

[1] That said, the implementation of dynamic modules in Alice ML, for instance, employs a very similar construct [21].

This section is intended to be broadly accessible to readers who are generally familiar with the basic idea of relational parametricity but not with the details of (advanced) logical relations techniques.

In Section 4, we formalize our logical relation for G and show how it may be used to reason about parametricity and representation independence. A particularly appealing feature of our formalization is that the *non*-parametricity of G is encapsulated in the notion of what it means for two *types* to be logically related to each other when viewed as *data*. The definition of this type-level logical relation is a one-liner, which can easily be replaced with an alternative "parametric" version. In Sections 5–7, we explore how terms related by the parametric version of our logical relation may be "wrapped" systematically to produce terms related by the non-parametric version, and vice versa, thus clarifying how dynamic type generation enables parametric reasoning. This leads us to a novel polarized form of our logical relation, which distinguishes between positive and negative notions of parametricity.

Finally, in Section 8, we discuss related work, including recent work on the relevant concepts of dynamic sealing [25] and multi-language interoperation [12].

2. The Language G

Figure 1 defines our non-parametric language G. For the most part, G is a standard call-by-value λ-calculus, consisting of the usual types and terms from System F [8], including pairs and existential types. We also assume an unspecified set of base types b, along with suitable constants c of those types.

Two additional, non-standard constructs isolate the essential aspects of the class of languages we are interested in:

- $\mathsf{cast}\ \tau_1\ \tau_2\ v_1\ v_2$ converts v_1 from type τ_1 to τ_2. It checks that those two types are the same at the time of evaluation. If so, the operator *succeeds* and returns v_1. Otherwise, it *fails* and defaults to v_2, which acts as an else clause of the target type τ_2.
- $\mathsf{new}\ \alpha{\approx}\tau\ \mathsf{in}\ e$ generates a fresh abstract type name α. Values of type α can be formed using its *representation type* τ. Both types are deemed *compatible*, but not equivalent. That is, they are considered equal as *classifiers*, but not as *data*. In particular, $\mathsf{cast}\ \alpha\ \tau\ v\ v'$ will not succeed (*i.e.*, it will return v').

Our cast operator is essentially the same as Harper and Mitchell's *TypeCond* operator [10], which was itself a variant of the non-parametric J operator that Girard studied in his thesis [8]. Our new construct is similar to previously proposed constructs for dynamic type generation [19, 27, 20]. However, we do not require *explicit* term-level type coercions to witness the isomorphism between an abstract type name α and its representation τ. Instead, our type system is simple enough that we perform this conversion *implicitly*.

For convenience, we will occasionally use expressions of the form $\mathsf{let}\ x{=}e_1\ \mathsf{in}\ e_2$, which abbreviate the term $(\lambda x{:}\tau_1.e_2)\ e_1$ (with τ_1 being an appropriate type for e_1). We omit the type annotation for existential packages where clear from context. Moreover, we take the liberty to generalize binary tuples to n-ary ones where necessary and to use pattern matching notation to decompose tuples in the obvious manner.

2.1 Typing Rules

The typing rules for the System F fragment of G are completely standard and thus omitted from Figure 1. We focus on the non-standard rules related to casting and dynamic type generation. Full formal details of the type system appear in the technical appendix.

Typing of casts is straightforward (Rule ECAST): $\mathsf{cast}\ \tau_1\ \tau_2$ is simply treated as a function of type $\tau_1 \to \tau_2 \to \tau_2$. Its first argument is the value to be converted, and its second argument is the default value returned in the case of failure. The rule merely requires that the two types be well-formed.

Types $\tau ::= \alpha \mid b \mid \tau \to \tau \mid \tau \times \tau \mid \forall\alpha.\tau \mid \exists\alpha.\tau$
Values $v ::= x \mid c \mid \lambda x{:}\tau.e \mid \langle v_1, v_2\rangle \mid \lambda\alpha.e \mid \mathsf{pack}\ \langle\tau, v\rangle\ \mathsf{as}\ \tau$
Terms $e ::= v \mid e\ e \mid \langle e_1, e_2\rangle \mid e.1 \mid e.2 \mid e\ \tau \mid$ $\mathsf{pack}\ \langle\tau, e\rangle\ \mathsf{as}\ \tau \mid \mathsf{unpack}\ \langle\alpha, x\rangle{=}e\ \mathsf{in}\ e \mid$ $\mathsf{cast}\ \tau\ \tau \mid \mathsf{new}\ \alpha{\approx}\tau\ \mathsf{in}\ e$
Stores $\sigma ::= \epsilon \mid \sigma, \alpha{\approx}\tau$
Config's $\zeta ::= \sigma; e$

Type Contexts $\Delta ::= \epsilon \mid \Delta, \alpha \mid \Delta, \alpha{\approx}\tau$
Value Contexts $\Gamma ::= \epsilon \mid \Gamma, x{:}\tau$

$$\boxed{\Delta;\Gamma \vdash e : \tau} \qquad \cdots$$

$$(\textsc{Ecast})\,\frac{\Delta \vdash \tau_1 \qquad \Delta \vdash \tau_2}{\Delta;\Gamma \vdash \mathsf{cast}\ \tau_1\ \tau_2 : \tau_1 \to \tau_2 \to \tau_2}$$

$$(\textsc{Enew})\,\frac{\Delta \vdash \tau \qquad \Delta, \alpha{\approx}\tau; \Gamma \vdash e : \tau' \qquad \Delta \vdash \tau'}{\Delta;\Gamma \vdash \mathsf{new}\ \alpha{\approx}\tau\ \mathsf{in}\ e : \tau'}$$

$$(\textsc{Econv})\,\frac{\Delta;\Gamma \vdash e : \tau' \qquad \Delta \vdash \tau \approx \tau'}{\Delta;\Gamma \vdash e : \tau}$$

$$\boxed{\Delta \vdash \tau}$$

$$(\textsc{Tname})\,\frac{\alpha{\approx}\tau \in \Delta}{\Delta \vdash \alpha} \qquad \cdots$$

$$\boxed{\Delta \vdash \tau \approx \tau}$$

$$(\textsc{Qname})\,\frac{\alpha{\approx}\tau \in \Delta}{\Delta \vdash \alpha \approx \tau} \qquad \cdots$$

$$\boxed{\vdash \zeta : \tau}$$

$$(\textsc{Conf})\,\frac{\vdash \sigma \qquad \sigma;\epsilon \vdash e : \tau \qquad \epsilon \vdash \tau}{\vdash \sigma; e : \tau}$$

$$\begin{array}{lrcl}
 & \sigma; (\lambda x{:}\tau.e)\ v & \hookrightarrow & \sigma; e[v/x] \\
 & \sigma; \langle v_1, v_2\rangle.i & \hookrightarrow & \sigma; v_i \\
 & \sigma; (\lambda\alpha.e)\ \tau & \hookrightarrow & \sigma; e[\tau/\alpha] \\
\multicolumn{2}{l}{\sigma; \mathsf{unpack}\ \langle\alpha, x\rangle{=}(\mathsf{pack}\ \langle\tau, v\rangle)\ \mathsf{in}\ e} & \hookrightarrow & \sigma; e[\tau/\alpha][v/x] \\
(\alpha \notin \mathrm{dom}(\sigma)) & \sigma; \mathsf{new}\ \alpha{\approx}\tau\ \mathsf{in}\ e & \hookrightarrow & \sigma, \alpha{\approx}\tau; e \\
(\tau_1 = \tau_2) & \sigma; \mathsf{cast}\ \tau_1\ \tau_2 & \hookrightarrow & \sigma; \lambda x_1{:}\tau_1.\lambda x_2{:}\tau_2.x_1 \\
(\tau_1 \neq \tau_2) & \sigma; \mathsf{cast}\ \tau_1\ \tau_2 & \hookrightarrow & \sigma; \lambda x_1{:}\tau_1.\lambda x_2{:}\tau_2.x_2
\end{array}$$

(... plus standard "search" rules ...)

Figure 1. Syntax and Type System of G (excerpt)

For an expression $\mathsf{new}\ \alpha{\approx}\tau\ \mathsf{in}\ e$, which binds α in e, Rule ENEW checks that the body e is well-typed under the assumption that α is implemented by the representation type τ. For that purpose, we enrich type contexts Δ with entries of the form $\alpha{\approx}\tau$ that keep track of the representation types tied to abstract type names. Note that τ may not mention α.

Syntactically, type names are just type variables. When viewed as data, (*i.e.,* when inspected by the cast operator), types are considered equivalent iff they are syntactically equal. In contrast, when viewed as classifiers for terms, knowledge about the representation of type names may be taken into account. Rule ECONV says that if a term e has a type τ', it may be assigned any other type that is *compatible* with τ'. Type compatibility, in turn, is defined by the judgment $\Delta \vdash \tau_1 \approx \tau_2$. We only show the rule QNAME, which discharges a compatibility assumption $\alpha{\approx}\tau$ from the context; the other rules implement the congruence closure of this axiom. The important point here is that equivalent types are compatible, but compatible types are not necessarily equivalent.

Finally, Rule ENEW also requires that the type τ' of the body e does not contain α (*i.e.,* τ' must be well formed in Δ alone). A type of this form can always be derived by applying ECONV to convert τ' to $\tau'[\tau/\alpha]$.

2.2 Dynamic Semantics

The operational semantics has to deal with generation of fresh type names. To that end, we introduce a *type store* σ to record generated type names. Hence, reduction is defined on *configurations* $(\sigma; e)$ instead of plain terms. Figure 1 shows the main reduction rules. We omit the standard "search" rules for descending into subterms according to call-by-value, left-to-right evaluation order.

The reduction rules for the F fragment are as usual and do not actually touch the store. However, types occurring in F constructs can contain type names bound in the store.

Reducing the expression $\mathsf{new}\,\alpha{\approx}\tau\,\mathsf{in}\,e$ creates a new entry for α in the type store. We rely on the usual hygiene convention for bound variables to ensure that α is fresh with respect to the current store (which can always be achieved by α-renaming).[2]

The two remaining rules are for casts. A cast takes two types and checks that they are equivalent (*i.e.,* syntactically equal). In either case, the expression reduces to a function that will return the appropriate one of the additional value arguments, *i.e.,* the value to be converted in case of success, and the default value otherwise. In the former case, type preservation is ensured because source and target types are known to be equivalent.

Type preservation can be expressed using the typing rule CONF for configurations. We formulate this rule by treating the type store as a type context, which is possible because type stores are a syntactic subclass of type contexts. (In a similar manner, we can write $\vdash\,\sigma$ for well-formedness of store σ, by viewing it as a type context.) It is worth noting that the representation types in the store are actually never inspected by the dynamic semantics. They are only needed for specifying well-formedness of configurations and proving type soundness.

2.3 Motivating Example

Consider the following attempt to write a simple functional "binary semaphore" ADT [15] in G. Following Mitchell and Plotkin [14], we use an existential type, as we would in System F:

$$\begin{aligned}
\tau_{\text{sem}} &:= \exists\alpha.\alpha \times (\alpha \to \alpha) \times (\alpha \to \mathsf{bool})\\
e_{\text{sem}} &:= \mathsf{pack}\,\langle \mathsf{int}, \langle 1, \lambda x{:}\,\mathsf{int}.(1-x), \lambda x{:}\,\mathsf{int}.(x \neq 0)\rangle\rangle\ \mathsf{as}\ \tau_{\text{sem}}
\end{aligned}$$

A semaphore essentially is a flag that can be in two states: either *locked* or *unlocked*. The state can be toggled using the first function of the ADT, and it can be polled using the second. Our little module uses an integer value for representing the state, taking 1 for locked and 0 for unlocked. It is an invariant of the implementation that the integer never takes any other value—otherwise, the toggle function would no longer operate correctly.

In System F, the implementation invariant would be protected by the fact that existential types are parametric: there is no way to inspect the witness of α after opening the package, and hence no client could produce values of type α other than those returned by the module (nor could she apply integer operations to them).

Not so in G. The following program uses cast to forge a value s of the abstract semaphore type α:

$$\begin{aligned}
e_{\text{client}} \quad := \quad & \mathsf{unpack}\,\langle \alpha, \langle s_0, \mathit{toggle}, \mathit{poll}\rangle\rangle = e_{\text{sem}}\ \mathsf{in}\\
& \mathsf{let}\, s = \mathsf{cast}\ \mathsf{int}\ \alpha\ 666\ s_0\ \mathsf{in}\\
& \langle \mathit{poll}\ s, \mathit{poll}\ (\mathit{toggle}\ s)\rangle
\end{aligned}$$

Because reduction of unpack simply substitutes the representation type int for α, the consecutive cast succeeds, and the whole expression evaluates to $\langle \mathsf{true}, \mathsf{true}\rangle$—although the second component should have toggled s and thus be different from the first.

The way to prevent this in G is to create a fresh type name as witness of the abstract type:

$$\begin{aligned}
e_{\text{sem1}} \ := \ & \mathsf{new}\,\alpha' \approx \mathsf{int}\ \mathsf{in}\\
& \mathsf{pack}\,\langle \alpha', \langle 1, \lambda x{:}\,\mathsf{int}.(1-x), \lambda x{:}\,\mathsf{int}.(x \neq 0)\rangle\rangle\ \mathsf{as}\ \tau_{\text{sem}}
\end{aligned}$$

After replacing the initial semaphore implementation with this one, e_{client} will evaluate to $\langle \mathsf{true}, \mathsf{false}\rangle$ as desired—the cast expression will no longer succeed, because α will be substituted by the dynamic type name α', and $\alpha' \neq \mathsf{int}$. (Moreover, since α' is only visible statically in the scope of the new expression, the client has no access to α', and thus cannot *convert* from int to α' either.)

Now, while it is clear that new ensures proper type abstraction in the client program e_{client}, we want to prove that it does so for *any* client program. A standard way of doing so is by showing a more general property, namely *representation independence* [18]: we show that the module e_{sem1} is *contextually equivalent* to another module of the same type, meaning that no G program can observe any difference between the two modules. By choosing that other module to be a suitable reference implementation of the ADT in question, we can conclude that the "real" one behaves properly under all circumstances.

The obvious candidate for a reference implementation of the semaphore ADT is the following:

$$\begin{aligned}
e_{\text{sem2}} \ := \ & \mathsf{new}\,\alpha' \approx \mathsf{bool}\ \mathsf{in}\\
& \mathsf{pack}\,\langle \alpha', \langle \mathsf{true}, \lambda x{:}\,\mathsf{bool}.\neg x, \lambda x{:}\,\mathsf{bool}.x\rangle\rangle\ \mathsf{as}\ \tau_{\text{sem}}
\end{aligned}$$

Here, the semaphore state is represented directly by a Boolean flag and does not rely on any additional invariant. If we can show that e_{sem1} is contextually equivalent to e_{sem2}, then we can conclude that e_{sem1}'s type representation is truly being held abstract.

2.4 Contextual Equivalence

In order to be able to reason about representation independence, we need to make precise the notion of contextual equivalence.

A context C is an expression with a single hole $[\_]$, defined in the usual manner. Typing of contexts is defined by a judgment form $\vdash\, C : (\Delta;\Gamma;\tau) \rightsquigarrow (\Delta';\Gamma';\tau')$, where the triple $(\Delta;\Gamma;\tau)$ indicates the type of the hole. The judgment implies that for any expression e with $\Delta;\Gamma \vdash e : \tau$ we have $\Delta';\Gamma' \vdash C[e] : \tau'$. The rules are straightforward, the key rule being the one for holes:

$$\frac{\Delta \subseteq \Delta' \qquad \Gamma \subseteq \Gamma'}{\vdash [\_] : (\Delta;\Gamma;\tau) \rightsquigarrow (\Delta';\Gamma';\tau)}$$

We can now define contextual approximation and contextual equivalence as follows:

Definition 2.1 (Contextual Approximation and Equivalence) Let $\Delta;\Gamma \vdash e_1 : \tau$ and $\Delta;\Gamma \vdash e_2 : \tau$.

$$\begin{aligned}
\Delta;\Gamma \vdash e_1 \preceq e_2 : \tau \ &\overset{\text{def}}{\Leftrightarrow}\ \forall C, \tau', \sigma.\\
&\qquad \vdash \sigma \wedge \vdash C : (\Delta;\Gamma;\tau) \rightsquigarrow (\sigma;\epsilon;\tau') \ \wedge\\
&\qquad \sigma; C[e_1]\downarrow \ \Longrightarrow \ \sigma; C[e_2]\downarrow\\
\Delta;\Gamma \vdash e_1 \simeq e_2 : \tau \ &\overset{\text{def}}{\Leftrightarrow}\ \Delta;\Gamma \vdash e_1 \preceq e_2 : \tau\ \wedge\\
&\qquad \Delta;\Gamma \vdash e_2 \preceq e_1 : \tau
\end{aligned}$$

That is, contextual approximation $\Delta;\Gamma \vdash e_1 \preceq e_2 : \tau$ means that for any well-typed program context C with a hole of appropriate type, the termination of $C[e_1]$ implies the termination of $C[e_2]$. Contextual equivalence $\Delta;\Gamma \vdash e_1 \simeq e_2 : \tau$ is just approximation in both directions.

Considering that G does not explicitly contain any recursive or looping constructs, the reader may wonder why termination is used

[2] A well-known alternative approach would omit the type store in favour of using scope extrusion rules for new binders, as in Rossberg [19].

as the notion of "distinguishing observation" in our definition of contextual equivalence. The reason is that the cast operator, together with impredicative polymorphism, makes it possible to write well-typed non-terminating programs. (This was Girard's point in studying the J operator in the first place [8].) Moreover, one can encode arbitrary recursive function definitions. Other forms of observation may thus be encoded in terms of (non-)termination.

3. A Logical Relation for G: Main Ideas

Following Reynolds [18] and Mitchell [13], our general approach to reasoning about parametricity and representation independence is to define a *logical relation*. Essentially, logical relations give us a tractable way of proving that two terms are contextually equivalent, which in turn gives us a way of proving that abstract types are really abstract. Of course, since polymorphism in G is non-parametric, the definition of our logical relation in the cases of universal and existential types is somewhat unusual. To place our approach in context, we first review the traditional approach to defining logical relations for languages with parametric polymorphism, such as System F.

3.1 Logical Relations for Parametric Polymorphism

Although the technical meaning of "logical relation" is rather woolly, the basic idea is to define an equivalence (or approximation) relation on programs inductively, following the structure of their types. To take the canonical example of arrow types, we would say that two functions are logically related at the type $\tau_1 \to \tau_2$ if, when passed arguments that are logically related at τ_1, either they both diverge or they both converge to values that are logically related at τ_2. The *fundamental theorem* of logical relations states that the logical relation is a congruence with respect to the constructs of the language. Together with what Pitts [15] calls *adequacy—i.e.*, the fact that logically related terms have equivalent termination behavior—the fundamental theorem implies that logically related terms are contextually equivalent, since contextual equivalence is defined precisely to be the largest adequate congruence.

Traditionally, the parametric nature of polymorphism is made clear by the definition of the logical relation for universal and existential types. Intuitively, two type abstractions, $\lambda\alpha.e_1$ and $\lambda\alpha.e_2$, are logically related at type $\forall\alpha.\tau$ if they map related *type* arguments to related results. But what does it mean for two type arguments to be related? Moreover, once we settle on two related type arguments τ_1' and τ_2', at what type do we relate the results $e_1[\tau_1'/\alpha]$ and $e_2[\tau_2'/\alpha]$?

One approach would be to restrict "related type arguments" to be the *same* type τ'. Thus, $\lambda\alpha.e_1$ and $\lambda\alpha.e_2$ would be logically related at $\forall\alpha.\tau$ iff, for any (closed) type τ', it is the case that $e_1[\tau'/\alpha]$ and $e_2[\tau'/\alpha]$ are logically related at the type $\tau[\tau'/\alpha]$. A key problem with this definition, however, is that, due to the quantification over *any* argument type τ', the type $\tau[\tau'/\alpha]$ may in fact be larger than the type $\forall\alpha.\tau$, and thus the definition of the logical relation is no longer inductive in the structure of the type. Another problem is that this definition does not tell us anything about the parametric nature of polymorphism.

Reynolds' alternative approach is a generalization of Girard's "candidates" method for proving strong normalization for System F [8]. The idea is simple: instead of defining two type arguments to be related only if they are the same, allow *any* two different type arguments to be related by an (almost) arbitrary relational interpretation (subject to certain *admissibility* constraints). That is, we parameterize the logical relation at type τ by an interpretation function ρ, which maps each free type variable of τ to a pair of types τ_1', τ_2' together with some (admissible) relation between values of those types. Then, we say that $\lambda\alpha.e_1$ and $\lambda\alpha.e_2$ are logically related at type $\forall\alpha.\tau$ under interpretation ρ iff, for any closed types τ_1' and τ_2' and any relation R between values of those types, it is the case that $e_1[\tau_1'/\alpha]$ and $e_2[\tau_2'/\alpha]$ are logically related at type τ under interpretation $\rho, \alpha \mapsto (\tau_1', \tau_2', R)$.

The miracle of Reynolds/Girard's method is that it simultaneously (1) renders the logical relation inductively well-defined in the structure of the type, and (2) demonstrates the parametricity of polymorphism: logically related type abstractions must behave the same even when passed completely different type arguments, so their behavior may not analyze the type argument and behave in different ways for different arguments. Dually, we can show that two ADTs pack $\langle\tau_1, v_1\rangle$ as $\exists\alpha.\tau$ and pack $\langle\tau_2, v_2\rangle$ as $\exists\alpha.\tau$ are logically related (and thus contextually equivalent) by exhibiting *some* relational interpretation R for the abstract type α, even if the underlying type representations τ_1 and τ_2 are different. This is the essence of what is meant by "representation independence".

Unfortunately, in the setting of G, Reynolds/Girard's method is not directly applicable, precisely because polymorphism in G is not parametric! This essentially forces us back to the first approach suggested above, namely to only consider type arguments to be logically related if they are equal. Moreover, it makes sense: the cast operator views types as data, so types may only be logically related if they are indistinguishable as data.

The natural questions, then, are: (1) what metric do we use to define the logical relation inductively, since the structure of the type no longer suffices, and (2) how do we establish that dynamic type generation regains a form of parametricity? We address these questions in the next two sections, respectively.

3.2 Step-Indexed Logical Relations for Non-Parametricity

First, in order to provide a metric for inductively defining the logical relation, we employ *step-indexing*. Step-indexed logical relations were proposed originally by Appel and McAllester [7] as a way of giving a simple operational-semantics-based model for general recursive types in the context of foundational proof-carrying code. In subsequent work by Ahmed and others [3, 6], the method has been adapted to support relational reasoning in a variety of settings, including untyped and imperative languages.

The key idea of step-indexed logical relations is to index the definition of the logical relation not only by the type of the programs being related, but also by a natural number n representing (intuitively) "the number of steps left in the computation". That is, if two terms e_1 and e_2 are logically related at type τ for n steps, then if we place them in any program context C and run the resulting programs for n steps of computation, we should not be able to produce observably different results (*e.g.*, $C[e_1]$ evaluating to 5 and $C[e_2]$ evaluating to 7). To show that e_1 and e_2 are contextually equivalent, then, it suffices to show that they are logically related for n steps, for any n.

To see how step-indexing helps us, consider how we might define a step-indexed logical relation for G in the case of universal types: two type abstractions $\lambda\alpha.e_1$ and $\lambda\alpha.e_2$ are logically related at $\forall\alpha.\tau$ for n steps iff, for any type argument τ', it is the case that $e_1[\tau'/\alpha]$ and $e_2[\tau'/\alpha]$ are logically related at $\tau[\tau'/\alpha]$ for $n-1$ steps. This reasoning is sound because the only way a program context can distinguish between $\lambda\alpha.e_1$ and $\lambda\alpha.e_2$ in n steps is by first applying them to a type argument τ'—which incurs a step of computation for the β-reduction $(\lambda\alpha.e_i)\,\tau' \hookrightarrow e_i[\tau'/\alpha]$—and then distinguishing between $e_1[\tau'/\alpha]$ and $e_2[\tau'/\alpha]$ within the next $n-1$ steps. Moreover, although the type $\tau[\tau'/\alpha]$ may be larger than $\forall\alpha.\tau$, the step index $n-1$ is smaller, so the logical relation is inductively well-defined.

3.3 Kripke Logical Relations for Dynamic Parametricity

Second, in order to establish the parametricity properties of dynamic type generation, we employ *Kripke logical relations*, *i.e.*, logical relations that are indexed by *possible worlds*.[3] Kripke logical relations are appropriate when reasoning about properties that are true only under certain conditions, such as equivalence of modules with local mutable state. For instance, an imperative ADT might only behave according to its specification if its local data structures obey certain invariants. Possible worlds allow one to codify such *local invariants* on the machine store [15].

In our setting, the local invariant we want to establish is what a dynamically generated type name *means*. That is, we will use possible worlds to assign relational interpretations to dynamically generated type names. For example, consider the programs e_{sem1} and e_{sem2} from Section 2. We want to show they are logically related at $\exists\alpha.\ \alpha \times (\alpha \to \alpha) \times (\alpha \to \mathsf{bool})$ in an empty initial world w_0 (*i.e.*, under empty type stores). The proof proceeds roughly as follows. First, we evaluate the two programs. This will have the effect of generating a fresh type name α', with $\alpha' \approx \mathsf{int}$ extending the type store of the first program and $\alpha' \approx \mathsf{bool}$ extending the type store of the second program. At this point, we correspondingly extend the initial world w_0 with a mapping from α' to the relation $R = \{(1, \mathsf{true}), (0, \mathsf{false})\}$, thus forming a new world w that specifies the semantic meaning of α'.

We now must show that the values

$$\mathsf{pack}\ \langle\alpha', \langle 1, \lambda x{:}\,\mathsf{int}\,.(1-x), \lambda x{:}\,\mathsf{int}\,.(x \neq 0)\rangle\rangle\ \mathsf{as}\ \tau_{\text{sem}}$$

and

$$\mathsf{pack}\ \langle\alpha', \langle \mathsf{true}, \lambda x{:}\,\mathsf{bool}\,.\neg x, \lambda x{:}\,\mathsf{bool}\,.x\rangle\rangle\ \mathsf{as}\ \tau_{\text{sem}}$$

are logically related in the world w. Since G's logical relation for existential types is non-parametric, the two packages must have the *same* type representation, but of course the whole point of using new was to ensure that they do (namely, it is α'). The remainder of the proof is showing that the value components of the packages are related at the type $\alpha' \times (\alpha' \to \alpha') \times (\alpha' \to \mathsf{bool})$ under the interpretation $\rho = \alpha' \mapsto (\mathsf{int}, \mathsf{bool}, R)$ derived from the world w. This last part is completely analogous to what one would show in a standard representation independence proof.

In short, the possible worlds in our Kripke logical relations bring back the ability to assign arbitrary relational interpretations R to abstract types, an ability that was seemingly lost when we moved to a non-parametric logical relation. The only catch is that we can only assign arbitrary interpretations to *dynamic* type names, not to *static*, universally/existentially quantified type variables.

There is one minor technical matter that we glossed over in the above proof sketch but is worth mentioning. Due to nondeterminism of type name allocation, the evaluation of e_{sem1} and e_{sem2} may result in α' being replaced by α'_1 in the former and α'_2 in the latter (for some fresh $\alpha'_1 \neq \alpha'_2$). Moreover, we are also interested in proving equivalence of programs that do not necessarily allocate exactly the same number of type names in the same order.

Consequently, we also include in our possible worlds a partial bijection η between the type names of the first program and the type names of the second program, which specifies how each dynamically generated abstract type is concretely represented in the stores of the two programs. We require them to be in 1-1 correspondence because the cast construct permits the program context to observe equality on type names, as follows:

$$\begin{array}{l}\mathsf{equal?} : \forall\alpha.\forall\beta.\,\mathsf{bool} \ \stackrel{\text{def}}{=} \\ \quad \Lambda\alpha.\Lambda\beta.\ \mathsf{cast}\ \ ((\alpha \to \alpha) \to \mathsf{bool})\ ((\beta \to \beta) \to \mathsf{bool}) \\ \qquad\qquad\qquad (\lambda x{:}(\alpha \to \alpha).\,\mathsf{true})(\lambda x{:}(\beta \to \beta).\,\mathsf{false})(\lambda x{:}\beta.x)\end{array}$$

We then consider types to be logically related if they are the same *up to* this bijection. For instance, in our running example, when extending w_0 to w, we would not only extend its relational interpretation with $\alpha' \mapsto (\mathsf{int}, \mathsf{bool}, R)$ but also extend its η with $\alpha' \mapsto (\alpha'_1, \alpha'_2)$. Thus, the type representations of the two existential packages, α'_1 and α'_2, though syntactically distinct, would still be logically related under w.

4. A Logical Relation for G: Formal Details

Figure 2 displays our step-indexed Kripke logical relation for G in full gory detail. It is easiest to understand this definition by making two passes over it. First, as the step indices have a way of infecting the whole definition in a superficially complex—but really very straightforward—way, we will first walk through the whole definition *ignoring* all occurrences of n's and k's (as well as auxiliary functions like the $\lfloor\cdot\rfloor_n$ operator). Second, we will pinpoint the few places where step indices actually play an important role in ensuring that the logical relation is inductively well-founded.

4.1 Highlights of the Logical Relation

The first section of Figure 2 defines the kinds of semantic objects that are used in the construction of the logical relation. Relations R are sets of *atoms*, which are pairs of terms, e_1 and e_2, indexed by a possible world w. The definition of $\mathrm{Atom}[\tau_1, \tau_2]$ requires that e_1 and e_2 have the types τ_1 and τ_2 under the type stores $w.\sigma_1$ and $w.\sigma_2$, respectively. (We use the dot notation $w.\sigma_i$ to denote the i-th type store component of w, and analogous notation for projecting out the other components of worlds.)

$\mathrm{Rel}[\tau_1, \tau_2]$ defines the set of *admissible* relations, which are permitted to be used as the semantic interpretations of abstract types. For our purposes, admissibility is simply *monotonicity*—*i.e.*, closure under world extension. That is, if a relation in Rel relates two values v_1 and v_2 under a world w, then the relation must relate those values in any future world of w. (We discuss the definition of world extension below.) Monotonicity is needed in order to ensure that we can extend worlds with interpretations of new dynamic type names, without interfering somehow with the interpretations of the old ones.

Worlds w are 4-tuples $(\sigma_1, \sigma_2, \eta, \rho)$, which describe a set of assumptions under which pairs of terms are related. Here, σ_1 and σ_2 are the type stores under which the terms are typechecked and evaluated. The finite mappings η and ρ share a common domain, which can be understood as the set of abstract type names that have been generated dynamically. These "semantic" type names do not exist in either store σ_1 or σ_2.[4] Rather, they provide a way of referring to an abstract type that is represented by *some* type name α_1 in σ_1 and *some* type name α_2 in σ_2. Thus, for each name $\alpha \in \mathrm{dom}(\eta) = \mathrm{dom}(\rho)$, the *concretization* η maps the "semantic" name α to a pair of "concrete" names from the stores σ_1 and σ_2, respectively. (See the end of Section 3.3 for an example of such an η.) As the definition of Conc makes clear, distinct semantic type names must have distinct concretizations; consequently, η represents a *partial bijection* between σ_1 and σ_2.

The last component of the world w is ρ, which assigns relational interpretations to the aforementioned semantic type names. Formally, ρ maps each α to a triple $r = (\tau_1, \tau_2, R)$, where R is a monotone relation between values of types τ_1 and τ_2. (Again, see the end of Section 3.3 for an example of such a ρ.) The final condition in the definition of World stipulates that the closed syntactic types in the range of ρ and the concrete type names in the range of η are compatible. As a matter of notation, we will write η^i and ρ^i to denote the type substitutions $\{\alpha \mapsto \alpha_i \mid \eta(\alpha) = (\alpha_1, \alpha_2)\}$ and $\{\alpha \mapsto \tau_i \mid \rho(\alpha) = (\tau_1, \tau_2, R)\}$, respectively.

[3] In fact, step-indexed logical relations may already be understood as a special case of Kripke logical relations, in which the step index serves as the notion of possible world, and where n is a future world of m iff $n \leq m$.

[4] In fact, technically speaking, we consider $\mathrm{dom}(\eta) = \mathrm{dom}(\rho)$ to be bound variables of the world w.

$$
\begin{aligned}
\mathrm{Atom}_n[\tau_1,\tau_2] &\stackrel{\text{def}}{=} \{(k,w,e_1,e_2) \mid k<n \land w\in \mathrm{World}_k \land \vdash w.\sigma_1; e_1 : \tau_1 \land \vdash w.\sigma_2; e_2 : \tau_2\}\\
\mathrm{Rel}_n[\tau_1,\tau_2] &\stackrel{\text{def}}{=} \{R \subseteq \mathrm{Atom}_n^{\mathsf{val}}[\tau_1,\tau_2] \mid \forall(k,w,v_1,v_2)\in R.\ \forall(k',w')\sqsupseteq(k,w).\ (k',w',v_1,v_2)\in R\}\\
\mathrm{SomeRel}_n &\stackrel{\text{def}}{=} \{r=(\tau_1,\tau_2,R) \mid \mathrm{fv}(\tau_1,\tau_2)=\emptyset \land R\in\mathrm{Rel}_n[\tau_1,\tau_2]\}\\
\mathrm{Interp}_n &\stackrel{\text{def}}{=} \{\rho \in \mathrm{TVar} \xrightarrow{\text{fin}} \mathrm{SomeRel}_n\}\\
\mathrm{Conc} &\stackrel{\text{def}}{=} \{\eta \in \mathrm{TVar} \xrightarrow{\text{fin}} \mathrm{TVar}\times\mathrm{TVar} \mid \forall\alpha,\alpha'\in\mathrm{dom}(\eta).\ \alpha\neq\alpha' \Rightarrow \eta^1(\alpha)\neq\eta^1(\alpha') \land \eta^2(\alpha)\neq\eta^2(\alpha')\}\\
\mathrm{World}_n &\stackrel{\text{def}}{=} \{w=(\sigma_1,\sigma_2,\eta,\rho) \mid\ \vdash\sigma_1 \land \vdash\sigma_2 \land \eta\in\mathrm{Conc} \land \rho\in\mathrm{Interp}_n \land \mathrm{dom}(\eta)=\mathrm{dom}(\rho)\ \land\\
&\qquad \forall\alpha\in\mathrm{dom}(\rho).\ \sigma_1\vdash\rho^1(\alpha)\approx\eta^1(\alpha) \land \sigma_2\vdash\rho^2(\alpha)\approx\eta^2(\alpha)\}
\end{aligned}
$$

$$
\begin{aligned}
\lfloor(\sigma_1,\sigma_2,\eta,\rho)\rfloor_n &\stackrel{\text{def}}{=} (\sigma_1,\sigma_2,\eta,\lfloor\rho\rfloor_n)\\
\lfloor\rho\rfloor_n &\stackrel{\text{def}}{=} \{\alpha\mapsto\lfloor r\rfloor_n \mid \rho(\alpha)=r\}\\
\lfloor(\tau_1,\tau_2,R)\rfloor_n &\stackrel{\text{def}}{=} (\tau_1,\tau_2,\lfloor R\rfloor_n)\\
\lfloor R\rfloor_n &\stackrel{\text{def}}{=} \{(k,w,e_1,e_2)\in R \mid k<n\}\\
\triangleright R &\stackrel{\text{def}}{=} \{(k,w,e_1,e_2) \mid w\in\mathrm{World}_k \land\\
&\qquad (k-1,\lfloor w\rfloor_{k-1},e_1,e_2)\in R\}
\end{aligned}
$$

$$
\begin{aligned}
(k',w')\sqsupseteq(k,w) &\stackrel{\text{def}}{\Leftrightarrow} k'\leq k \land w'\in\mathrm{World}'_k\ \land\\
&\qquad w'.\eta\sqsupseteq w.\eta \land w'.\rho\sqsupseteq\lfloor w.\rho\rfloor_{k'}\ \land\\
&\qquad \forall i\in\{1,2\}.\ w'.\sigma_i\supseteq w.\sigma_i\ \land\\
&\qquad\quad \mathrm{rng}(w'.\eta^i)-\mathrm{rng}(w.\eta^i)\subseteq\\
&\qquad\quad \mathrm{dom}(w'.\sigma_i)-\mathrm{dom}(w.\sigma_i)\\
\eta'\sqsupseteq\eta &\stackrel{\text{def}}{\Leftrightarrow} \forall\alpha\in\mathrm{dom}(\eta).\ \eta'(\alpha)=\eta(\alpha)\\
\rho'\sqsupseteq\rho &\stackrel{\text{def}}{\Leftrightarrow} \forall\alpha\in\mathrm{dom}(\rho).\ \rho'(\alpha)=\rho(\alpha)
\end{aligned}
$$

$$
\begin{aligned}
V_n\llbracket\alpha\rrbracket\rho &\stackrel{\text{def}}{=} \lfloor\rho(\alpha).R\rfloor_n\\
V_n\llbracket b\rrbracket\rho &\stackrel{\text{def}}{=} \{(k,w,c,c)\in\mathrm{Atom}_n[b,b]\}\\
V_n\llbracket\tau\times\tau'\rrbracket\rho &\stackrel{\text{def}}{=} \{(k,w,\langle v_1,v_1'\rangle,\langle v_2,v_2'\rangle)\in\mathrm{Atom}_n[\rho^1(\tau\times\tau'),\rho^2(\tau\times\tau')] \mid\\
&\qquad (k,w,v_1,v_2)\in V_n\llbracket\tau\rrbracket\rho \land (k,w,v_1',v_2')\in V_n\llbracket\tau'\rrbracket\rho\}\\
V_n\llbracket\tau'\to\tau\rrbracket\rho &\stackrel{\text{def}}{=} \{(k,w,\lambda x{:}\tau_1.e_1,\lambda x{:}\tau_2.e_2)\in\mathrm{Atom}_n[\rho^1(\tau'\to\tau),\rho^2(\tau'\to\tau)] \mid\\
&\qquad \forall(k',w',v_1,v_2)\in V_n\llbracket\tau'\rrbracket\rho.\ (k',w')\sqsupseteq(k,w)\Rightarrow\\
&\qquad (k',w',e_1[v_1/x],e_2[v_2/x])\in E_n\llbracket\tau\rrbracket\rho\}\\
V_n\llbracket\forall\alpha.\tau\rrbracket\rho &\stackrel{\text{def}}{=} \{(k,w,\lambda\alpha.e_1,\lambda\alpha.e_2)\in\mathrm{Atom}_n[\rho^1(\forall\alpha.\tau),\rho^2(\forall\alpha.\tau)] \mid\\
&\qquad \forall(k',w')\sqsupseteq(k,w).\ \forall(\tau_1,\tau_2,r)\in T_{k'}\llbracket\Omega\rrbracket w'.\\
&\qquad (k',w',e_1[\tau_1/\alpha],e_2[\tau_2/\alpha])\in\triangleright E_n\llbracket\tau\rrbracket\rho,\alpha\mapsto r\}\\
V_n\llbracket\exists\alpha.\tau\rrbracket\rho &\stackrel{\text{def}}{=} \{(k,w,\mathsf{pack}\,\langle\tau_1,v_1\rangle,\mathsf{pack}\,\langle\tau_2,v_2\rangle)\in\mathrm{Atom}_n[\rho^1(\exists\alpha.\tau),\rho^2(\exists\alpha.\tau)] \mid\\
&\qquad \exists r.\ (\tau_1,\tau_2,r)\in T_k\llbracket\Omega\rrbracket w \land (k,w,v_1,v_2)\in\triangleright V_n\llbracket\tau\rrbracket\rho,\alpha\mapsto r\}\\
E_n\llbracket\tau\rrbracket\rho &\stackrel{\text{def}}{=} \{(k,w,e_1,e_2)\in\mathrm{Atom}_n[\rho^1(\tau),\rho^2(\tau)] \mid\\
&\qquad \forall j<k.\ \forall\sigma_1,v_1.\ (w.\sigma_1;e_1\hookrightarrow^j\sigma_1;v_1)\Rightarrow\\
&\qquad \exists w',v_2.\ (k-j,w')\sqsupseteq(k,w)\land w'.\sigma_1=\sigma_1\land(w.\sigma_2;e_2\hookrightarrow^* w'.\sigma_2;v_2)\land(k-j,w',v_1,v_2)\in V_n\llbracket\tau\rrbracket\rho\}\\
T_n\llbracket\Omega\rrbracket w &\stackrel{\text{def}}{=} \{(w.\eta^1(\tau),w.\eta^2(\tau),(w.\rho^1(\tau),w.\rho^2(\tau),V_n\llbracket\tau\rrbracket w.\rho)) \mid\ \mathrm{fv}(\tau)\subseteq\mathrm{dom}(w.\rho)\}\\
G_n\llbracket\epsilon\rrbracket\rho &\stackrel{\text{def}}{=} \{(k,w,\emptyset,\emptyset) \mid k<n\land w\in\mathrm{World}_k\}\\
G_n\llbracket\Gamma,x{:}\tau\rrbracket\rho &\stackrel{\text{def}}{=} \{(k,w,(\gamma_1,x\mapsto v_1),(\gamma_2,x\mapsto v_2)) \mid\\
&\qquad (k,w,\gamma_1,\gamma_2)\in G_n\llbracket\Gamma\rrbracket\rho\land(k,w,v_1,v_2)\in V_n\llbracket\tau\rrbracket\rho\}\\
D_n\llbracket\epsilon\rrbracket w &\stackrel{\text{def}}{=} \{(\emptyset,\emptyset,\emptyset)\}\\
D_n\llbracket\Delta,\alpha\rrbracket w &\stackrel{\text{def}}{=} \{((\delta_1,\alpha\mapsto\tau_1),(\delta_2,\alpha\mapsto\tau_2),(\rho,\alpha\mapsto r)) \mid\\
&\qquad (\delta_1,\delta_2,\rho)\in D_n\llbracket\Delta\rrbracket w\land(\tau_1,\tau_2,r)\in T_n\llbracket\Omega\rrbracket w\}\\
D_n\llbracket\Delta,\alpha{\approx}\tau\rrbracket w &\stackrel{\text{def}}{=} \{((\delta_1,\alpha\mapsto\beta_1),(\delta_2,\alpha\mapsto\beta_2),(\rho,\alpha\mapsto r)) \mid\\
&\qquad (\delta_1,\delta_2,\rho)\in D_n\llbracket\Delta\rrbracket w\ \land\\
&\qquad \exists\alpha'.\ w.\rho(\alpha')=r\land w.\eta(\alpha')=(\beta_1,\beta_2)\ \land\\
&\qquad w.\sigma_1(\beta_1)=\delta_1(\tau)\land w.\sigma_2(\beta_2)=\delta_2(\tau)\land r.R=V_n\llbracket\tau\rrbracket\rho\}
\end{aligned}
$$

$$
\begin{aligned}
\Delta;\Gamma\vdash e_1\precsim e_2:\tau \stackrel{\text{def}}{\Leftrightarrow}\ & \Delta;\Gamma\vdash e_1:\tau\land\Delta;\Gamma\vdash e_2:\tau\ \land\\
& \forall n\geq 0.\ \forall w_0\in\mathrm{World}_n.\ \forall(\delta_1,\delta_2,\rho)\in D_n\llbracket\Delta\rrbracket w_0.\ \forall(k,w,\gamma_1,\gamma_2)\in G_n\llbracket\Gamma\rrbracket\rho.\\
& (k,w)\sqsupseteq(n,w_0)\Rightarrow(k,w,\delta_1\gamma_1(e_1),\delta_2\gamma_2(e_2))\in E_n\llbracket\tau\rrbracket\rho
\end{aligned}
$$

Figure 2. Logical Relation for G

The second section of Figure 2 displays the definition of world extension. In order for w' to extend w (written $w' \sqsupseteq w$), it must be the case that (1) w' specifies semantic interpretations for a superset of the type names that w interprets, (2) for the names that w interprets, w' must interpret them in the same way, and (3) any new semantic type names that w' interprets may only correspond to *new* concrete type names that did not exist in the stores of w. Although the third condition is not strictly necessary, we have found it to be useful when proving certain examples (*e.g.,* the "order independence" example in Section 4.4).

The last section of Figure 2 defines the logical relation itself. $V\llbracket\tau\rrbracket\rho$ is the logical relation for values, $E\llbracket\tau\rrbracket\rho$ is the one for terms, and $T\llbracket\Omega\rrbracket w$ is the one for *types as data*, as described in Section 3 (here, Ω represents the *kind* of types).

$V\llbracket\tau\rrbracket\rho$ relates values at the type τ, where the free type variables of τ are given relational interpretations by ρ. Ignoring the step indices, $V\llbracket\tau\rrbracket\rho$ is mostly very standard. For instance, at certain points (namely, in the $\rightarrow$ and $\forall$ cases), when we quantify over logically related (value or type) arguments, we must allow them to come from an arbitrary future world w' in order to ensure monotonicity. This kind of quantification over future worlds is commonplace in Kripke logical relations.

The only really interesting bit in the definition of $V\llbracket\tau\rrbracket\rho$ is the use of $T\llbracket\Omega\rrbracket w$ to characterize when the two *type* arguments (resp. components) of a universal (resp. existential) are logically related. As explained in Section 3.3, we consider two types to be logically related in world w iff they are the same up to the partial bijection $w.\eta$. Formally, we define $T\llbracket\Omega\rrbracket w$ as a relation on triples (τ_1, τ_2, r), where τ_1 and τ_2 are the two logically related types and r is a relation telling us how to relate values of those types. To be logically related means that τ_1 and τ_2 are the concretizations (according to $w.\eta$) of some "semantic" type τ'. Correspondingly, r is the logical relation $V\llbracket\tau'\rrbracket w.\rho$ at that semantic type. Thus, when we write $E\llbracket\tau\rrbracket\rho, \alpha \mapsto r$ in the definition of $V\llbracket\forall\alpha.\tau\rrbracket\rho$, this is roughly equivalent to writing $E\llbracket\tau[\tau'/\alpha]\rrbracket\rho$ (which our discussion in Section 3.2 might have led the reader to expect to see here instead). The reason for our present formulation is that $E\llbracket\tau[\tau'/\alpha]\rrbracket\rho$ is not quite right: the free variables of τ are interpreted by ρ, but the free variables of τ' are *dynamic* type names whose interpretations are given by $w.\rho$. It is possible to merge ρ and $w.\rho$ into a unified interpretation ρ', but we feel our present approach is cleaner.

Another point of note: since r is uniquely determined from τ_1 and τ_2, it is not really necessary to include it in the $T\llbracket\Omega\rrbracket w$ relation. However, as we shall see in Section 6, formulating the logical relation in this way has the benefit of isolating all of the non-parametricity of our logical relation in the definition of $T\llbracket\Omega\rrbracket w$.

The term relation $E\llbracket\tau\rrbracket\rho$ is very similar to that in previous step-indexed Kripke logical relations [6]. Briefly, it says that two terms are related in an initial world w if whenever the first evaluates to a value under $w.\sigma_1$, the second evaluates to a value under $w.\sigma_2$, and the resulting stores and values are related in some future world w'.

The remainder of the definitions in Figure 2 serve to formalize a logical relation for *open* terms. $G\llbracket\Gamma\rrbracket\rho$ is the logical relation on value substitutions γ, which asserts that related γ's must map variables in $\mathrm{dom}(\Gamma)$ to related values. $D\llbracket\Delta\rrbracket w$ is the logical relation on type substitutions. It asserts that related δ's must map variables in $\mathrm{dom}(\Delta)$ to types that are related in w. For type variables α bound as $\alpha \approx \tau$, the δ's must map α to a type name whose semantic interpretation in w is precisely the logical relation at τ. Analogously to $T\llbracket\Omega\rrbracket w$, the relation $D\llbracket\Delta\rrbracket w$ also includes a relational interpretation ρ, which may be uniquely determined from the δ's.

Finally, the open logical relation $\Delta; \Gamma \vdash e_1 \precsim e_2 : \tau$ is defined in a fairly standard way. It says that for any starting world w_0, and any type substitutions δ_1 and δ_2 related in that world, if we are given related value substitutions γ_1 and γ_2 in any future world w, then $\delta_1\gamma_1 e_1$ and $\delta_2\gamma_2 e_2$ are related in w as well.

4.2 Why and Where the Steps Matter

As we explained in Section 3.2, step indices play a critical role in making the logical relation well-founded. Essentially, whenever we run into an apparent circularity, we "go down a step" by defining an n-level property in terms of an $(n-1)$-level one. Of course, this trick only works if, at all such "stepping points", the only way that an adversarial program context could possibly tell whether the n-level property holds or not is by taking one step of computation and then checking whether the underlying $(n-1)$-level property holds. Fortunately, this is the case.

Since worlds contain relations, and relations contain sets of tuples that include worlds, a naïve construction of these objects would have an inconsistent cardinality. We thus stratify both worlds and relations by a step index: n-level worlds $w \in \mathrm{World}_n$ contain n-level interpretations $\rho \in \mathrm{Interp}_n$, which map type variables to n-level relations; n-level relations $R \in \mathrm{Rel}_n[\tau_1, \tau_2]$ only contain atoms indexed by a step level $k < n$ and a world $w \in \mathrm{World}_k$. Although our possible worlds have a different structure than in previous work, the technique of mutual world and relation stratification is similar to that used in Ahmed's thesis [2], as well as recent work by Ahmed, Dreyer and Rossberg [6].

Intuitively, the reason this works in our setting is as follows. Viewed as a judgment, our logical relation asserts that two terms e_1 and e_2 are logically related for k steps in a world w at a type τ under an interpretation ρ (whose domain contains the free type variables of τ). Clearly, in order to handle the case where τ is just a type variable α, the relations r in the range of ρ must include atoms at step index k (*i.e.*, the r's must be in $\mathrm{SomeRel}_{k+1}$).

But what about the relations in the range of $w.\rho$? Those relations only come into play in the universal and existential cases of the logical relation for values. Consider the existential case (the universal one is analogous). There, $w.\rho$ pops up in the definition of the relation r that comes from $T_k\llbracket\Omega\rrbracket w$. However, that r is only needed in defining the relatedness of the values v_1 and v_2 at step level $k-1$ (note the definition of $\triangleright R$ in the second section of Figure 2). Consequently, we only need r to include atoms at step $k-1$ and lower (*i.e.*, r must be in $\mathrm{SomeRel}_k$), so the world w from which r is derived need only be in World_k.

As this discussion suggests, it is *imperative* that we "go down a step" in the universal and existential cases of the logical relation. For the other cases, it is not necessary to go down a step, although we have the option of doing so. For example, we could define k-level relatedness at pair type $\tau_1 \times \tau_2$ in terms of $(k-1)$-level relatedness at τ_1 and τ_2. But since the type gets smaller, there is no need to. For clarity, we have only gone down a step in the logical relation at the points where it is absolutely necessary, and we have used the $\triangleright$ notation to underscore those points.

4.3 Key Properties

The main results concerning our logical relation are as follows:

Theorem 4.1 (Fundamental Property for $\precsim$)
If $\Delta; \Gamma \vdash e : \tau$, then $\Delta; \Gamma \vdash e \precsim e : \tau$.

Theorem 4.2 (Soundness of $\precsim$ wrt. Contextual Approximation)
If $\Delta; \Gamma \vdash e_1 \precsim e_2 : \tau$, then $\Delta; \Gamma \vdash e_1 \preceq e_2 : \tau$.

These theorems establish that our logical relation provides a sound technique for proving contextual equivalence of G programs. The proofs of these theorems rely on many technical lemmas, most of which are standard and straightforward to prove. We highlight a few of them here, and refer the reader to the technical appendix for full details of the proofs.

One key lemma we have mentioned already is the *monotonicity* lemma, which states that the logical relation for values is closed under world extension, and therefore belongs to the Rel class of relations. Another key lemma is *transitivity of world extension*.

There are also a group of lemmas—Pitts terms them *compatibility* lemmas [15]—which show that the logical relation is a pre-congruence with respect to the constructs of the G language. Of particular note among these are the ones for `cast` and `new`.

For `cast`, we must show that `cast` $\tau_1\,\tau_2$ is logically related to itself under a type context Δ assuming that τ_1 and τ_2 are well-formed in Δ. This boils down to showing that, for logically related type substitutions δ_1 and δ_2, it is the case that $\delta_1\tau_1 = \delta_1\tau_2$ if and only if $\delta_2\tau_1 = \delta_2\tau_2$. This follows easily from the fact that δ_1 and δ_2, by virtue of being logically related, map the variables in $\mathrm{dom}(\Delta)$ to types that are syntactically identical up to some bijection on type names.

For new, we must show that, if $\Delta, \alpha{\approx}\tau'; \Gamma \vdash e_1 \precsim e_2 : \tau$, then $\Delta; \Gamma \vdash \mathsf{new}\,\alpha{\approx}\tau'\ \mathsf{in}\ e_1 \precsim \mathsf{new}\,\alpha{\approx}\tau'\ \mathsf{in}\ e_2 : \tau$ (assuming $\Delta \vdash \Gamma$ and $\Delta \vdash \tau$). The proof involves extending the η and ρ components of some given initial world w_0 with bindings for the fresh dynamically-generated type name α. The η is extended with $\alpha \mapsto (\alpha_1, \alpha_2)$, where α_1 and α_2 are the concrete fresh names that are chosen when evaluating the left and right new expressions. The ρ is extended so that the relational interpretation of α is simply the logical relation at type τ'. The proof of this lemma is highly reminiscent of the proof of compatibility for ref (reference allocation) in a language with mutable references [6].

Finally, another important compatibility property is *type compatibility*, *i.e.*, that if $\Delta \vdash \tau_1 \approx \tau_2$ and $(\delta_1, \delta_2, \rho) \in D_n\llbracket\Delta\rrbracket w$, then $V_n\llbracket\tau_1\rrbracket\rho = V_n\llbracket\tau_2\rrbracket\rho$ and $E_n\llbracket\tau_1\rrbracket\rho = E_n\llbracket\tau_2\rrbracket\rho$. The interesting case is when τ_1 is a variable α bound in Δ as $\alpha \approx \tau_2$, and the result in this case follows easily from the definition of $D\llbracket\Delta, \alpha \approx \tau\rrbracket w$.

4.4 Examples

Semaphore. We now return to our semaphore example from Section 2 and show how to prove representation independence for the two different implementations e_{sem1} and e_{sem2}. Recall that the former uses int, the latter bool. To show that they are contextually equivalent, it suffices by Soundness to show that each logically approximates the other. We prove only one direction, namely $\vdash e_{\mathsf{sem1}} \precsim e_{\mathsf{sem2}} : \tau_{\mathsf{sem}}$; the other is proven analogously.

Expanding the definitions, we need to show $(k, w, e_{\mathsf{sem1}}, e_{\mathsf{sem2}}) \in E_n\llbracket\tau_{\mathsf{sem}}\rrbracket\emptyset$. Note how each term generates a fresh type name α_i in one step, resulting in a package value. Hence all we need to do is come up with a world w' satisfying

- $(k-1, w') \sqsupseteq (k, w)$,
- $w'.\sigma_1 = w.\sigma_1, \alpha_1{\approx}\mathsf{int}$ and $w'.\sigma_2 = w.\sigma_2, \alpha_2{\approx}\mathsf{bool}$,
- $(k-1, w', \mathsf{pack}\langle\alpha_1, v_1\rangle, \mathsf{pack}\langle\alpha_2, v_2\rangle) \in V_n\llbracket\tau_{\mathsf{sem}}\rrbracket\emptyset$.

where v_i is the term component of $e_{\mathsf{sem}i}$'s implementation. We construct w' by extending w with mappings that establish the relation between the new type names:

$$\begin{aligned} R &:= \{(k'', w'', v_{\mathsf{int}}, v_{\mathsf{bool}}) \in \mathrm{Atom}^{\mathsf{val}}_{k-1}[\mathsf{int}, \mathsf{bool}] \mid \\ &\qquad (v_{\mathsf{int}}, v_{\mathsf{bool}}) = (1, \mathsf{true}) \vee (v_{\mathsf{int}}, v_{\mathsf{bool}}) = (0, \mathsf{false})\} \\ r &:= (\mathsf{int}, \mathsf{bool}, R) \\ w' &:= \lfloor w \rfloor_{k-1} \uplus (\alpha_1{\approx}\mathsf{int}, \alpha_2{\approx}\mathsf{bool}, \alpha{\mapsto}(\alpha_1, \alpha_2), \alpha{\mapsto}r) \end{aligned}$$

The first two conditions above are satisfied by construction. To show that the packages are related we need to show the existence of an r' with $(\alpha_1, \alpha_2, r') \in T_{k-1}\llbracket\Omega\rrbracket w'$ such that $(k-2, \lfloor w' \rfloor_{k-2}, v_1, v_2) \in V_n\llbracket\tau'_{\mathsf{sem}}\rrbracket\rho, \alpha{\mapsto}r'$, where $\tau'_{\mathsf{sem}} = \alpha \times (\alpha \to \alpha) \times (\alpha \to \mathsf{bool})$. Since $\alpha_i = w'.\eta^i(\alpha)$, r' must be $(\mathsf{int}, \mathsf{bool}, V_{k-1}\llbracket\alpha\rrbracket w'.\rho)$ by definition of $T\llbracket\Omega\rrbracket$. Of course, we defined w' the way we did so that this r' is exactly r.

The proof of $(k-2, \lfloor w' \rfloor_{k-2}, v_1, v_2) \in V_n\llbracket\tau'_{\mathsf{sem}}\rrbracket\rho, \alpha{\mapsto}r$ decomposes into three parts, following the structure of τ'_{sem}:

1. $(k-2, \lfloor w' \rfloor_{k-2}, 1, \mathsf{true}) \in V_n\llbracket\alpha\rrbracket\rho, \alpha{\mapsto}r$
   This holds because $V_n\llbracket\alpha\rrbracket\rho, \alpha{\mapsto}r = R$.
2. $(k-2, \lfloor w' \rfloor_{k-2}, \lambda x{:}\mathsf{int}.(1-x), \lambda x{:}\mathsf{bool}.\neg x)$
   $\in V_n\llbracket\alpha \to \alpha\rrbracket\rho, \alpha{\mapsto}r$
   - Suppose we are given related arguments in a future world: $(k'', w'', v'_1, v'_2) \in V_n\llbracket\alpha\rrbracket\rho, \alpha{\mapsto}r = R$.
   - Hence either $(v'_1, v'_2) = (1, \mathsf{true})$ or $(v'_1, v'_2) = (0, \mathsf{false})$.
   - Consequently, $1 - v'_1$ and $\neg v'_2$ will evaluate in one step, without effects, to values again related by R.
   - In other words, $(k'', w'', 1 - v'_1, \neg v'_2) \in E_n\llbracket\alpha\rrbracket\rho, \alpha{\mapsto}r$.
3. $(k-2, \lfloor w' \rfloor_{k-2}, \lambda x.(x \neq 0), \lambda x.x) \in V_n\llbracket\alpha \to \mathsf{bool}\rrbracket\rho, \alpha{\mapsto}r$
   Like in the previous part, the arguments v'_1 and v'_2 will be related by R in some future (k'', w''). Therefore $v'_1 \neq 0$ will reduce in one step without effects to v'_2, which already is a value. Because of the definition of the logical relation at type bool, this implies $(k'', w'', v'_1 \neq 0, v'_2) \in E_n\llbracket\mathsf{bool}\rrbracket\rho, \alpha{\mapsto}r$.

Benign Effects. When side effects are introduced into a pure language, they often falsify various equational laws concerning repeatability and order independence of computations. In this section, we offer some evidence that the effect of dynamic type generation is fairly *benign* in that it does not falsify such equational laws.

First, consider the following functions:

$$\begin{aligned} v_1 &:= \lambda x{:}(\mathsf{unit} \to \tau).\ \mathsf{let}\ x' = x\,()\ \mathsf{in}\ x\,() \\ v_2 &:= \lambda x{:}(\mathsf{unit} \to \tau).\ x\,() \end{aligned}$$

The only difference between v_1 and v_2 is whether the argument x is applied once or twice. Intuitively, either $x\,()$ diverges, in which case both programs diverge, or else the first application of x terminates, in which case so should the second.

Second, consider the following functions:

$$\begin{aligned} v'_1 &:= \lambda x{:}(\mathsf{unit} \to \tau).\lambda y{:}(\mathsf{unit} \to \tau').\ \mathsf{let}\ y' = y\,()\ \mathsf{in}\ \langle x\,(), y'\rangle \\ v'_2 &:= \lambda x{:}(\mathsf{unit} \to \tau).\lambda y{:}(\mathsf{unit} \to \tau').\langle x\,(), y\,()\rangle \end{aligned}$$

The only difference between v'_1 and v'_2 is the order in which they call their argument callbacks x and y. Those calls may both result in the generation of fresh type names, but the order in which the names are generated should not matter.

Using our logical relation, we can prove that v_1 and v_2 are contextually equivalent, and so are v'_1 and v'_2. Due to space considerations, we refer the interested reader to the technical appendix for full proof details.

5. Wrapping

We have seen that parametricity can be re-established in G by introducing name generation in the right place. But what is the "right place" in general? That is, given an arbitrary expression e with polymorphic type τ_e, how can we *systematically* transform it into an expression e' of the same type τ_e that is parametric?

One obvious—but unfortunately bogus—idea is the following: transform e such that every existential *introduction* and every universal *elimination* creates a fresh name for the respective witness or instance type. Formally, apply the following rewrite rules to e:

$$\begin{array}{lcl} \mathsf{pack}\,\langle\tau, e\rangle\ \mathsf{as}\ \tau' & \rightsquigarrow & \mathsf{new}\,\alpha{\approx}\tau\ \mathsf{in}\ \mathsf{pack}\,\langle\alpha, e\rangle\ \mathsf{as}\ \tau' \\ e\ \tau & \rightsquigarrow & \mathsf{new}\,\alpha{\approx}\tau\ \mathsf{in}\ e\ \alpha \end{array}$$

Obviously, this would make every quantified type abstract, so that any cast that tries to inspect it would fail.

Or would it? Perhaps surprisingly, the answer is no. To see why, consider the following expressions of type $(\exists\alpha.\tau') \times (\exists\alpha.\tau')$:

$$\begin{array}{lcl} e_1 & := & \mathsf{let}\ x = \mathsf{pack}\,\langle\tau, v\rangle\ \mathsf{in}\ \langle x, x\rangle \\ e_2 & := & \langle\mathsf{pack}\,\langle\tau, v\rangle, \mathsf{pack}\,\langle\tau, v\rangle\rangle \end{array}$$

They are clearly equivalent in a parametric language (and in fact they are even equivalent in G). Yet rewriting yields:

$$\begin{array}{lcl} e'_1 & := & \mathsf{let}\ x = \mathsf{new}\,\alpha{\approx}\tau\ \mathsf{in}\ \mathsf{pack}\,\langle\alpha, v\rangle\ \mathsf{in}\ \langle x, x\rangle \\ e'_2 & := & \langle\mathsf{new}\,\alpha{\approx}\tau\ \mathsf{in}\ \mathsf{pack}\,\langle\alpha, v\rangle, \mathsf{new}\,\alpha{\approx}\tau\ \mathsf{in}\ \mathsf{pack}\,\langle\alpha, v\rangle\rangle \end{array}$$

$$
\begin{array}{lcl}
\mathrm{Wr}^{\pm}_{\tau}(e) & \overset{\text{def}}{=} & \mathsf{let}\, x{=}e\ \mathsf{in}\ \mathrm{Wr}^{\pm}_{\tau}(x) \\
\mathrm{Wr}^{\pm}_{\alpha}(v) & \overset{\text{def}}{=} & v \\
\mathrm{Wr}^{\pm}_{b}(v) & \overset{\text{def}}{=} & v \\
\mathrm{Wr}^{\pm}_{\tau_1\times\tau_2}(v) & \overset{\text{def}}{=} & \langle \mathrm{Wr}^{\pm}_{\tau_1}(v.1), \mathrm{Wr}^{\pm}_{\tau_2}(v.2)\rangle \\
\mathrm{Wr}^{\pm}_{\tau_1\to\tau_2}(v) & \overset{\text{def}}{=} & \lambda x_1{:}\tau_1.\, \mathrm{Wr}^{\pm}_{\tau_2}(v\ \ \mathrm{Wr}^{\mp}_{\tau_1}(x_1)) \\
\mathrm{Wr}^{\pm}_{\forall\alpha.\tau}(v) & \overset{\text{def}}{=} & \lambda\alpha.\, \mathsf{new}^{\mp}\ \alpha\ \mathsf{in}\ \mathrm{Wr}^{\pm}_{\tau}(v\ \alpha) \\
\mathrm{Wr}^{\pm}_{\exists\alpha.\tau}(v) & \overset{\text{def}}{=} & \mathsf{unpack}\,\langle\alpha, x\rangle{=}v\ \mathsf{in} \\
 & & \quad \mathsf{new}^{\pm}\ \alpha\ \mathsf{in}\ \mathsf{pack}\,\langle\alpha, \mathrm{Wr}^{\pm}_{\tau}(x)\rangle\ \mathsf{as}\ \exists\alpha.\tau \\
\mathsf{new}^{+}\ \alpha\ \mathsf{in}\ e & \overset{\text{def}}{=} & \mathsf{new}\,\alpha'{\approx}\alpha\ \mathsf{in}\ e[\alpha'/\alpha] \\
\mathsf{new}^{-}\ \alpha\ \mathsf{in}\ e & \overset{\text{def}}{=} & e
\end{array}
$$

Figure 3. Wrapping

The resulting expressions are *not* equivalent anymore, because they perform different effects. Here is one distinguishing context:

$$
\begin{array}{l}
\mathsf{let}\, p = [\_]\ \mathsf{in}\ \mathsf{unpack}\,\langle\alpha_1, x_1\rangle = p.1\ \mathsf{in} \\
\qquad\qquad\qquad \mathsf{unpack}\,\langle\alpha_2, x_2\rangle = p.2\ \mathsf{in}\ \mathsf{equal?}\ \alpha_1\ \alpha_2
\end{array}
$$

Although the representation type τ is not disclosed as such, *sharing* between the two abstract types in e_1' is. In a parametric language, that would not be possible.

In order to introduce effects uniformly, and to hide internal sharing, the transformation we are looking for needs to be defined on the structure of types, not terms. Roughly, for each quantifier occurring in τ_e we need to generate one fresh type name. That is, instead of transforming e itself, we simply *wrap* it with some expression that introduces the necessary names at the boundary, by induction on the type τ_e.

In fact, we can refine the problem further. When looking at a G expression e, what do we actually mean by “making it parametric”? We can mean two different things: either ensuring that e *behaves* parametrically, or dually, that any context *treats* e parametrically. In the former case, we are protecting the *context* against e, in the latter we protect e against malicious contexts. The latter is what is sometimes referred to as *abstraction safety*.

Figure 3 defines a pair of wrapping operators that correspond to these two dual requirements: Wr^{+} protects an expression $e : \tau_e$ from being *used* in a non-parametric way, by inserting fresh names for each existential quantifier. Dually, Wr^{-} forces e to *behave* parametrically by creating a fresh name for each polymorphic instantiation. The definitions extend to other types in the usual functorial manner. Both definitions are interdependent, because roles switch for function arguments.

Given these operators, we can go back to our semaphore example: e_{sem1} can now be obtained as $\mathrm{Wr}^{+}_{\tau_{\text{sem}}}(e_{\text{sem}})$ (modulo some harmless η-expansions). This generalises to any ADT: wrapping its implementation positively will guarantee abstraction by making it parametric. We prove that in the next section.

Positive wrapping is reminiscent of *module sealing* (or opaque signature ascription) in ML-style module languages. If we view e as a module and its type τ_e as a signature, then $\mathrm{Wr}^{+}_{\tau_e}(e)$ corresponds to the sealing operation $e :> \tau_e$. While module sealing typically only performs static abstraction, wrapping describes the dynamic equivalent [20]. In fact, positive wrapping is precisely how sealing is implemented in Alice ML [21], where the module language is non-parametric otherwise.

6. Parametric Reasoning

The logical relation developed in Section 4 enables us to do *non-parametric* reasoning about equivalence of G programs. It also enables us to do *parametric* reasoning, but only indirectly: we have to explicitly deal with the effects of new and to define worlds containing relations between type names. It would be preferable if we were able to do parametric reasoning directly. For example, given two expressions e_1, e_2 that do not use casts, and assuming that the context does not do so either, we should be able to reason about equivalence of e_1 and e_2 in a manner similar to what we do when reasoning about System F.

6.1 A Parametric Logical Relation

Thanks to the modular formulation of our logical relation in Figure 2, it is easy to modify it so that it becomes parametric. All we need to do is swap out the definition of $T[\![\Omega]\!]w$, which relates types as data. Figure 4 gives an alternative definition that allows choosing an arbitrary relation between arbitrary types. Everything else stays exactly the same. We decorate the set of *parametric logical relations* thus obtained with $^\circ$ (*i.e.*, V°, E°, etc.) to distinguish them from the original ones. Likewise, we write $\precsim^\circ$ for the notion of *parametric logical approximation* defined as in Figure 2 but in terms of the parametric relations. For clarity, we will refer to the original definition as the *proper* logical relation.

This modification gives us a seemingly parametric definition of logical approximation for G terms. But what does that actually *mean*? What is the relation between parametric and proper logical approximation and, ultimately, *contextual* approximation? Since the language is not parametric, clearly, parametrically equivalent terms generally are not contextually equivalent.

The answer is given by the wrapping functions we defined in the previous section. The following theorem connects the two notions of logical relation and approximation that we have introduced:

Theorem 6.1 (Wrapping for $\precsim^\circ$)

1. If $\vdash e_1 \precsim^\circ e_2 : \tau$, then $\vdash \mathrm{Wr}^{+}_{\tau}(e_1) \precsim \mathrm{Wr}^{+}_{\tau}(e_2) : \tau$.
2. If $\vdash e_1 \precsim e_2 : \tau$, then $\vdash \mathrm{Wr}^{-}_{\tau}(e_1) \precsim^\circ \mathrm{Wr}^{-}_{\tau}(e_2) : \tau$.

This theorem justifies the definition of the parametric logical relation. At the same time it can be read as a correctness result for the wrapping operators: it says that whenever we can relate two terms using parametric reasoning, then the positive wrappings of the first term contextually approximates the positive wrapping of the second. Dually, once any properly related terms are wrapped negatively, they can safely be passed to any term that depends on its context behaving parametrically.

What can we say about the content of the parametric relation? Obviously, it cannot contain arbitrary G terms—*e.g.,* cast $\tau_1\,\tau_2$ will generally not be related to anything (including itself) in E°. However, for all constructs besides cast, we obtain the following:

Theorem 6.2 (Fundamental Property for $\precsim^\circ$)
If $\Delta; \Gamma \vdash e : \tau$ and e does not use cast, then $\Delta; \Gamma \vdash e \precsim^\circ e : \tau$.

In particular, this implies that any well-typed System F term is parametrically related to itself. The relation will also contain terms with cast, but only if the cast is “harmless”, *e.g.,* does not inspect types received from the context.

6.2 Examples

Semaphore. Consider our running example of the semaphore module again. Using the parametric relation, we can prove that the two implementations are related without actually reasoning about type generation. That aspect is covered once and for all by the Wrapping Theorem.

Recall the two implementations, here given in unwrapped form:

$$
\begin{array}{l}
e'_{\text{sem1}} := \mathsf{pack}\,\langle \mathsf{int}, \langle 1, \lambda x{:}\,\mathsf{int}\,.(1-x), \lambda x{:}\,\mathsf{int}\,.(x \neq 0)\rangle\rangle\ \mathsf{as}\ \tau_{\text{sem}} \\
e'_{\text{sem2}} := \mathsf{pack}\,\langle \alpha', \langle \mathsf{true}, \lambda x{:}\,\mathsf{bool}\,.\neg x, \lambda x{:}\,\mathsf{bool}\,.x\rangle\rangle\ \mathsf{as}\ \tau_{\text{sem}}
\end{array}
$$

$$T^{\circ}_n[\![\Omega]\!]w \quad \stackrel{\text{def}}{=} \quad \{(\tau_1, \tau_2, (\tau'_1, \tau'_2, R)) \mid \;\vdash \tau'_i \;\wedge\; w.\sigma_i \vdash \tau_i \approx \tau'_i \;\wedge\; R \in \mathrm{Rel}_n[\tau'_1, \tau'_2]\}$$

(everything else as in Figure 2)

Figure 4. Parametric Logical Relation

We can prove $\vdash e'_{\text{sem1}} \precsim^{\circ} e'_{\text{sem2}} : \tau_{\text{sem}}$ using conventional parametric reasoning about polymorphic terms. Now define $e_{\text{sem1}} = \mathrm{Wr}^{+}_{\tau_{\text{sem}}}(e'_{\text{sem1}})$ and $e_{\text{sem2}} = \mathrm{Wr}^{+}_{\tau_{\text{sem}}}(e'_{\text{sem2}})$, which is morally equivalent to the original definitions. The Wrapping Theorem then immediately tells us that $\vdash e_{\text{sem1}} \precsim e_{\text{sem2}} : \tau_{\text{sem}}$.

A Free Theorem. We can use the parametric relation for proving free theorems [28] in G. For example, for any $\vdash g : \forall\alpha.\alpha \to \alpha$ in G it holds that $\mathrm{Wr}^{-}(g)$ either diverges for all possible arguments τ and $\vdash v : \tau$, or it returns v in all cases. We first apply the Fundamental Property for $\precsim$ to relate g to itself in E, then transfer this to E° for $\mathrm{Wr}^{-}(g)$ using the Wrapping Theorem. From there the proof proceeds in the usual way.

6.3 Full Abstraction

The definition of the parametric relation E° is largely very similar to that of a typical (step-indexed) logical relation E_F for System F plus nontermination (see *e.g.* [3]). The main difference is the presence of worlds, but they are not actually used in a particularly interesting way in E°. Consequently, one might expect that any two terms related by the hypothetical E_F would also be related by E° and vice versa, *i.e.*, all logical equivalence results carry over from F to G. However, this is not obvious: G is more expressive than F, *i.e.*, even terms in the parametric relation can contain non-trivial uses of casts, and there is no evident way to back-translate these terms into F, as would be needed in the function case. That invalidates a proof approach like the one taken by Ahmed and Blume [5].

Ultimately, the property we would like to be able to show is *full abstraction* for the translation of F terms into G by wrapping:

$$\vdash e_1 \simeq_F e_2 : \tau \quad\Longleftrightarrow\quad \vdash \mathrm{Wr}^{+}_{\tau}(e_1) \simeq_G \mathrm{Wr}^{+}_{\tau}(e_2) : \tau$$

We conjecture that this holds, but are not aware of any suitable technique to prove it. This equivalence is even stronger than the one about logical relatedness in E_F and E° because our logical relation is only sound w.r.t. contextual equivalence, not complete.

7. Polarized Logical Relation

The parametric relation is useful for proving parametricity properties about (the positive wrappings of) G terms. However, it is all-or-nothing: it can only be used to prove parametricity for terms that are *treated* parametrically *and* also *behave* parametrically—cf. the two dual aspects of parametricity described in Section 5. This is adequate for proving representation independence for F terms, where we want the context to behave parametrically and the term itself is parametric by construction (because it has no casts). But we might also be interested in proving representation independence for terms that do not behave parametrically themselves. One situation where this might show up is if we want to show representation independence for generic modules—*e.g.*, in the form of ML functors—that have specialized behavior for particular argument types.

Here is a somewhat contrived example to illustrate the point. Consider the following two polymorphic functions of type $\forall\alpha.\tau_\alpha$:

$$\begin{array}{ll}
\tau_\alpha := \exists\beta.\,(\alpha \to \beta) \times (\beta \to \alpha) & \\
f_1 := \lambda\alpha.\,\mathsf{cast}\,\tau_{\mathsf{int}}\,\tau_\alpha & (\mathsf{pack}\,\langle \mathsf{int}, \langle \lambda x{:}\mathsf{int}.x{+}1, \lambda x{:}\mathsf{int}.x\rangle\rangle\ \mathsf{as}\ \tau_{\mathsf{int}}) \\
 & (\mathsf{pack}\,\langle \alpha, \langle \lambda x{:}\alpha.x, \lambda x{:}\alpha.x\rangle\rangle\ \mathsf{as}\ \tau_\alpha) \\
f_2 := \lambda\alpha.\,\mathsf{cast}\,\tau_{\mathsf{int}}\,\tau_\alpha & (\mathsf{pack}\,\langle \mathsf{int}, \langle \lambda x{:}\mathsf{int}.x, \lambda x{:}\mathsf{int}.x{+}1\rangle\rangle\ \mathsf{as}\ \tau_{\mathsf{int}}) \\
 & (\mathsf{pack}\,\langle \alpha, \langle \lambda x{:}\alpha.x, \lambda x{:}\alpha.x\rangle\rangle\ \mathsf{as}\ \tau_\alpha)
\end{array}$$

These functions can be understood as simplistic functors with a type argument α. Both functors implement a simple ADT β. Values of type α can be injected into β, and projected out again. However, both functors specialize the behavior of this ADT for type int—for integers, injecting n and projecting again will not give back n, but $n + 1$. This is true for both functors, but they implement it in a different way.

We want to prove that both implementations are equivalent under wrapping using a form of parametric reasoning. However, we cannot do that using the parametric relation from the previous section—since the functors do not *behave* parametrically (*i.e.*, they return different packages for different types), they will not be related in E°.

To support that kind of reasoning, we need a more refined treatment of parametricity in the logical relation. The idea is to separate the two aforementioned aspects of parametricity. Consequently, we are going to have a pair of separate relations, E^{+} and E^{-}. The former enforces parametric usage, the latter parametric behavior.

Figure 5 gives the definition of these relations. We call them *polarized*, because they are mutually dependent and the polarity ($+$ or $-$) switches for contravariant positions, *i.e.*, for function arguments and for universal quantifiers. Intuitively, in these places, term and context switch roles.

Except for the consistent addition of polarities, the definition of the polarized relations again only represents a minor modification of the original one.[5] We merely refine the definition of the type relation $T[\![\Omega]\!]w$ to distinguish polarity: in the positive case it behaves parametrically (*i.e.*, allowing an arbitrary relation) and in the negative case non-parametrically (*i.e.*, demanding r be the *logical* relation at some type). Thus, existential types behave parametrically in E^{+} but non-parametrically in E^{-}, and vice versa for universals.

7.1 Key Properties

The way in which polarities switch in the polarized relations mirrors what is going on in the definition of wrapping. That of course is no accident, and we can show the following theorem that relates the polarized relations with the proper and the parametric ones through uses of wrapping:

Theorem 7.1 (Wrapping for $\precsim^{\pm}$)

1. If $\vdash e_1 \precsim^{+} e_2 : \tau$, then $\vdash \mathrm{Wr}^{+}_{\tau}(e_1) \precsim \mathrm{Wr}^{+}_{\tau}(e_2) : \tau$.
2. If $\vdash e_1 \precsim e_2 : \tau$, then $\vdash \mathrm{Wr}^{-}_{\tau}(e_1) \precsim^{-} \mathrm{Wr}^{-}_{\tau}(e_2) : \tau$.
3. If $\vdash e_1 \precsim^{+} e_2 : \tau$, then $\vdash \mathrm{Wr}^{-}_{\tau}(e_1) \precsim^{\circ} \mathrm{Wr}^{-}_{\tau}(e_2) : \tau$.
4. If $\vdash e_1 \precsim^{\circ} e_2 : \tau$, then $\vdash \mathrm{Wr}^{+}_{\tau}(e_1) \precsim^{-} \mathrm{Wr}^{+}_{\tau}(e_2) : \tau$.

Moreover, we can show that the inverse directions of these implications require no wrapping at all:

Theorem 7.2 (Inclusion for $\precsim^{\pm}$)

1. If $\vdash e_1 \precsim e_2 : \tau$ or $\vdash e_1 \precsim^{\circ} e_2 : \tau$, then $\vdash e_1 \precsim^{+} e_2 : \tau$.
2. If $\vdash e_1 \precsim^{-} e_2 : \tau$, then $\vdash e_1 \precsim e_2 : \tau$ and $\vdash e_1 \precsim^{\circ} e_2 : \tau$.

This theorem can equivalently be stated as the chains $E^{-} \subseteq E \subseteq E^{+}$ and $E^{-} \subseteq E^{\circ} \subseteq E^{+}$.

Note that Theorem 6.1 follows directly from Theorems 7.1 and 7.2. Similarly, the following property follows from Theorem 7.2 together with Theorem 4.1:

[5] In fact, all four relations can easily be formulated in a single unified definition indexed by $\iota ::= \epsilon \mid \circ \mid + \mid -$. We refrained from doing so here for the sake of clarity; see the technical appendix for details.

$$
\begin{array}{lcl}
V_n^{\pm}\llbracket\alpha\rrbracket\rho & \stackrel{\text{def}}{=} & \lfloor\rho(\alpha).R\rfloor_n \\
V_n^{\pm}\llbracket b\rrbracket\rho & \stackrel{\text{def}}{=} & \{(k,w,c,c)\in \mathrm{Atom}_n[b,b]\} \\
V_n^{\pm}\llbracket\tau\times\tau'\rrbracket\rho & \stackrel{\text{def}}{=} & \{(k,w,\langle v_1,v_1'\rangle,\langle v_2,v_2'\rangle)\in \mathrm{Atom}_n[\rho^1(\tau\times\tau'),\rho^2(\tau\times\tau')]\mid \\
 & & \quad (k,w,v_1,v_2)\in V_n^{\pm}\llbracket\tau\rrbracket\rho \wedge (k,w,v_1',v_2')\in V_n^{\pm}\llbracket\tau'\rrbracket\rho\} \\
V_n^{\pm}\llbracket\tau'\to\tau\rrbracket\rho & \stackrel{\text{def}}{=} & \{(k,w,\lambda x{:}\tau_1.e_1,\lambda x{:}\tau_2.e_2)\in \mathrm{Atom}_n[\rho^1(\tau'\to\tau),\rho^2(\tau'\to\tau)]\mid \\
 & & \quad \forall(k',w',v_1,v_2)\in V_n^{\mp}\llbracket\tau'\rrbracket\rho.\,(k',w')\sqsupseteq(k,w)\Rightarrow \\
 & & \quad (k',w',e_1[v_1/x],e_2[v_2/x])\in E_n^{\pm}\llbracket\tau\rrbracket\rho\} \\
V_n^{\pm}\llbracket\forall\alpha.\tau\rrbracket\rho & \stackrel{\text{def}}{=} & \{(k,w,\lambda\alpha.e_1,\lambda\alpha.e_2)\in \mathrm{Atom}_n[\rho^1(\forall\alpha.\tau),\rho^2(\forall\alpha.\tau)]\mid \\
 & & \quad \forall(k',w')\sqsupseteq(k,w).\,\forall(\tau_1,\tau_2,r)\in T_{k'}^{\mp}\llbracket\Omega\rrbracket w'. \\
 & & \quad (k',w',e_1[\tau_1/\alpha],e_2[\tau_2/\alpha])\in \triangleright E_n^{\pm}\llbracket\tau\rrbracket\rho,\alpha\mapsto r\} \\
V_n^{\pm}\llbracket\exists\alpha.\tau\rrbracket\rho & \stackrel{\text{def}}{=} & \{(k,w,\mathsf{pack}\,\langle\tau_1,v_1\rangle,\mathsf{pack}\,\langle\tau_2,v_2\rangle)\in \mathrm{Atom}_n[\rho^1(\exists\alpha.\tau),\rho^2(\exists\alpha.\tau)]\mid \\
 & & \quad \exists r.\,(\tau_1,\tau_2,r)\in T_k^{\pm}\llbracket\Omega\rrbracket w\wedge(k,w,v_1,v_2)\in\triangleright V_n^{\pm}\llbracket\tau\rrbracket\rho,\alpha\mapsto r\} \\
E_n^{\pm}\llbracket\tau\rrbracket\rho & \stackrel{\text{def}}{=} & \{(k,w,e_1,e_2)\in \mathrm{Atom}_n[\rho^1(\tau),\rho^2(\tau)]\mid \\
 & & \quad \forall j<k.\,\forall\sigma_1,v_1.\,(w.\sigma_1;e_1\hookrightarrow^j\sigma_1;v_1)\Rightarrow \\
 & & \quad \exists w',v_2.\,(k-j,w')\sqsupseteq(k,w)\wedge w'.\sigma_1=\sigma_1\wedge(w.\sigma_2;e_2\hookrightarrow^* w'.\sigma_2;v_2)\wedge(k-j,w',v_1,v_2)\in V_n^{\pm}\llbracket\tau\rrbracket\rho\}
\end{array}
$$

$$
\begin{array}{lcllcl}
T_n^{+}\llbracket\Omega\rrbracket w & \stackrel{\text{def}}{=} & T_n^{\circ}\llbracket\Omega\rrbracket w & D_n^{+}\llbracket\Delta\rrbracket w & \stackrel{\text{def}}{=} & D_n^{\circ}\llbracket\Delta\rrbracket w \\
T_n^{-}\llbracket\Omega\rrbracket w & \stackrel{\text{def}}{=} & T_n\llbracket\Omega\rrbracket w & D_n^{-}\llbracket\Delta\rrbracket w & \stackrel{\text{def}}{=} & D_n\llbracket\Delta\rrbracket w
\end{array}
$$

$$
\begin{array}{lcl}
\Delta;\Gamma\vdash e_1\precsim^{\pm}e_2:\tau & \stackrel{\text{def}}{\Leftrightarrow} & \Delta;\Gamma\vdash e_1:\tau\wedge\Delta;\Gamma\vdash e_2:\tau\wedge \\
 & & \forall n\geq 0,\forall w_0\in\mathrm{World}_n\,.\,\forall(\delta_1,\delta_2,\rho)\in D_n^{\mp}\llbracket\Delta\rrbracket w_0.\,\forall(k,w,\gamma_1,\gamma_2)\in G_n^{\mp}\llbracket\Gamma\rrbracket\rho. \\
 & & (k,w)\sqsupseteq(n,w_0)\Rightarrow(k,w,\delta_1\gamma_1(e_1),\delta_2\gamma_2(e_2))\in E_n^{\pm}\llbracket\tau\rrbracket\rho
\end{array}
$$

Figure 5. Polarized Logical Relation

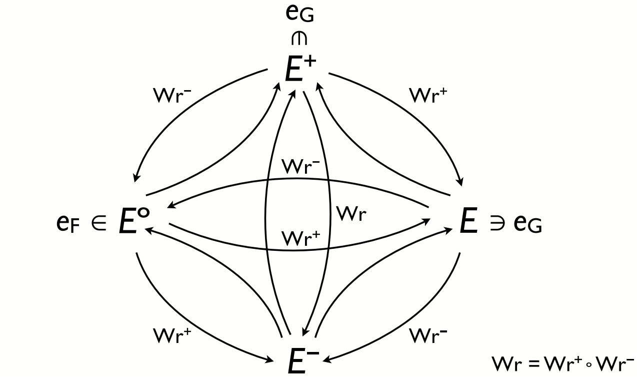

Figure 6. Relating the Relations

Theorem 7.3 (Fundamental Property for $\precsim^+$)
If $\Delta;\Gamma\vdash e:\tau$, then $\Delta;\Gamma\vdash e\precsim^+ e:\tau$.

Interestingly, compatibility does not hold for $\precsim^{\pm}$ (consider the polarities in the rule for application), which has the consequence that we cannot show Theorem 7.3 directly. For a similar reason, we cannot show any such property for $\precsim^-$ at all.

Figure 6 depicts all of the above properties in a single diagram. Plain arrows denote simple inclusion, while annotated arrows require respective wrapping to go from one relation to the other. The $\in$-operators show the fundamental properties for the respective relations, *i.e.*, which class of terms are included (G terms or F terms).

7.2 Example

Getting back to our motivating example from the beginning of the section, it is essentially straightforward to prove that $\vdash f_1 \precsim^+ f_2 : \forall\alpha.\tau_\alpha$. The proof proceeds as usual, except that we have to make a case distinction where we want to show that the functor bodies are related in E^+. At that point, we are given a triple $(\tau_1,\tau_2,r)\in T^-\llbracket\Omega\rrbracket w$.

If $\tau_1 = \mathsf{int}$, then we know from the definition of T^- that $\tau_2 = \mathsf{int}$, too. We hence know that both sides will evaluate to the specialized version of the ADT. Since we are in E^+, we get to pick some $(\tau_1',\tau_2',r')\in T^+\llbracket\Omega\rrbracket w$ as the interpretation of β, where the choice of r' is up to us. The natural choice is to use $\tau_1'=\tau_2'=\mathsf{int}$ with the relation $r'=(\mathsf{int},\mathsf{int},\{(k,w,n+1,n)\mid n\in\mathbb{N}\})$. The rest of the proof is then straightforward.

If $\tau_1\neq\mathsf{int}$ we similarly know that $\tau_2\neq\mathsf{int}$ from the definition of T^-. Hence, both sides use the default implementations, which are trivially related in E^+, thanks to Theorem 7.3.

Finally, applying the Wrapping Theorem 7.1, we can conclude that $\vdash \mathrm{Wr}^+(f_1)\precsim\mathrm{Wr}^+(f_2):\forall\alpha.\tau_\alpha$, and hence by Soundness, $\vdash \mathrm{Wr}^+(f_1)\preceq\mathrm{Wr}^+(f_2):\forall\alpha.\tau_\alpha$.

Note how we relied on the knowledge that τ_1 and τ_2 can only be int at the same time. This holds for types related in T^- but not in T^+ or T°. If we had tried to do this proof in E°, the types would have been related by T° only, which would give us too little information to proceed with the necessary case distinction.

8. Related Work

Type Generation vs. Other Forms of Data Abstraction. Traditionally, authors have distinguished between two complementary forms of data abstraction, sometimes dubbed the *static* and the *dynamic* approach [12]. The former is tied to the type system and relies on parametricity (especially for existential types) to hide an ADT's representation from clients [14]. The latter approach is typically employed in untyped languages, which do not have the ability to place static restrictions on clients. Consequently, data hiding has to be enforced on the level of individual values. For that, languages provide means for generating unique names and using them as *keys* for *dynamically sealing* values. A value sealed by a given key can only be inspected by principals that have access to the key [25].

Dynamic type generation as we employ it [19, 27, 20] can be seen as a middle ground, because it bears resemblance to both approaches. As in the dynamic approach, we cannot rely on parametricity and instead generate dynamic names to protect abstractions. However, these are type-level names, not term-level names, and they only 'seal' type information. In particular, individual values of abstract type are still directly represented by the underlying representation type, so that crossing abstraction boundaries has no runtime cost. In that sense, we are closer to the static approach.

Another approach to reconciling type abstraction and type analysis has been proposed by Washburn and Weirich [29]. They introduce a type system that tracks information flow for terms and types-as-data. By distinguishing security levels, the type system can statically prevent unauthorized inspection of types by clients.

Multi-Language Interoperation. The closest work to ours is that of Matthews and Ahmed [12]. They describe a pair of mutually recursive logical relations that deal with the interoperation between a typed language ("ML") and an untyped language ("Scheme"). Unlike in G, parametric behavior is hard-wired into their ML side: polymorphic instantiation unconditionally performs a form of dynamic sealing to protect against the non-parametric Scheme side. (In contrast, we treat new as its own language construct, orthogonal to universal types.) Dynamic sealing can then be defined in terms of the primitive coercion operators that bridge between the ML and Scheme sides. These coercions are similar to our (meta-level) wrapping operators, but ours perform type-level sealing, not term-level sealing.

The logical relations in Matthews and Ahmed's formalism are somewhat reminiscent of E° and E, although theirs are distinct logical relations for two languages, while ours are for a single language and differ only in the definition of $T[\![\Omega]\!]w$. In order to prove the fundamental property for their relations, they prove a "bridge lemma" transferring relatedness in one language to the other via coercions. This is analogous to our Wrapping Theorem for $\precsim^\circ$, but the latter is an independent theorem, not a lemma. Also, they do not propose anything like our polarized logical relations.

A fundamental technical difference is that their formulation of the logical relations does not use possible worlds to capture the type store (which is also left implicit in their operational semantics). Unfortunately, this resulted in a somewhat significant flaw in their technical development [4]. We believe that a reformulation of their relation along the lines of ours would fix this problem.

Proof Methods. Logical relations in various forms are routinely used to reason about program equivalence and type abstraction [18, 13, 15, 3]. In particular, Ahmed, Dreyer and Rossberg recently applied step-indexed logical relations with possible worlds to reason about type abstraction for a language with higher-order state [6]. State in G is comparatively benign, but still requires a circular definition of worlds that we stratify using steps.

Pitts and Stark used logical relations to reason about program equivalence in a language with (term-level) name generation [16] and subsequently generalized their technique to handle mutable references [17]. Sumii and Pierce use them for proving secrecy results for a language with dynamic sealing [24], where generated names are used as keys. In another line of work, Sumii and Pierce have used *bisimulations* to establish abstraction results for both untyped and polymorphic languages [25, 26]. However, none of the languages they investigate mixes the two paradigms.

Grossman, Morrisett and Zdancewic have proposed the use of *abstraction brackets* for syntactically tracing abstraction boundaries [9] during program execution. However, this is a comparatively weak method that does not seem to help in proving parametricity or representation independence results.

References

[1] Martín Abadi, Luca Cardelli, Benjamin Pierce, and Didier Rémy. Dynamic typing in polymorphic languages. *JFP*, 5(1), 1995.

[2] Amal Ahmed. *Semantics of Types for Mutable State*. PhD thesis, Princeton University, 2004.

[3] Amal Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *ESOP*, 2006.

[4] Amal Ahmed. Personal communication, 2009.

[5] Amal Ahmed and Matthias Blume. Typed closure conversion preserves observational equivalence. In *ICFP*, 2008.

[6] Amal Ahmed, Derek Dreyer, and Andreas Rossberg. State-dependent representation independence. In *POPL*, 2009.

[7] Andrew W. Appel and David McAllester. An indexed model of recursive types for foundational proof-carrying code. *Transactions on Programming Languages and Systems*, 23(5):657–683, 2001.

[8] Jean-Yves Girard. *Interprétation fonctionelle et élimination des coupures de l'arithmétique d'ordre supérieur*. PhD thesis, Université Paris VII, 1972.

[9] Dan Grossman, Greg Morrisett, and Steve Zdancewic. Syntactic type abstraction. *TOPLAS*, 22(6), 2000.

[10] Robert Harper and John Mitchell. Parametricity and variants of Girard's J operator. *Information Processing Letters*, 1999.

[11] Robert Harper and Greg Morrisett. Compiling polymorphism using intensional type analysis. In *POPL*, 1995.

[12] Jacob Matthews and Amal Ahmed. Parametric polymorphism through run-time sealing, or, theorems for low, low prices! In *ESOP*, 2008.

[13] John C. Mitchell. Representation independence and data abstraction. In *POPL*, 1986.

[14] John C. Mitchell and Gordon D. Plotkin. Abstract types have existential type. *TOPLAS*, 10(3), 1988.

[15] Andrew Pitts. Typed operational reasoning. In Benjamin C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 7. MIT Press, 2005.

[16] Andrew Pitts and Ian Stark. Observable properties of higher order functions that dynamically create local names, or: What's new? In *MFCS*, volume 711 of *LNCS*, 1993.

[17] Andrew Pitts and Ian Stark. Operational reasoning for functions with local state. In *HOOTS*, 1998.

[18] John C. Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing*, 1983.

[19] Andreas Rossberg. Generativity and dynamic opacity for abstract types. In *PPDP*, 2003.

[20] Andreas Rossberg. Dynamic translucency with abstraction kinds and higher-order coercions. In *MFPS*, 2008.

[21] Andreas Rossberg, Didier Le Botlan, Guido Tack, Thorsten Brunklaus, and Gert Smolka. Alice ML through the looking glass. In *TFP*, volume 5, 2004.

[22] Peter Sewell. Modules, abstract types, and distributed versioning. In *POPL*, 2001.

[23] Peter Sewell, James Leifer, Keith Wansbrough, Francesco Zappa Nardelli, Mair Allen-Williams, Pierre Habouzit, and Viktor Vafeiadis. Acute: high-level programming language design for distributed computation. *JFP*, 17(4&5):547–612, 2007.

[24] Eijiro Sumii and Benjamin Pierce. Logical relations for encryption. *JCS*, 11(4):521–554, 2003.

[25] Eijiro Sumii and Benjamin Pierce. A bisimulation for dynamic sealing. *TCS*, 375(1–3), 2007.

[26] Eijiro Sumii and Benjamin Pierce. A bisimulation for type abstraction and recursion. *JACM*, 54(5), 2007.

[27] Dimitrios Vytiniotis, Geoffrey Washburn, and Stephanie Weirich. An open and shut typecase. In *TLDI*, 2005.

[28] Philip Wadler. Theorems for free! In *FPCA*, 1989.

[29] Geoffrey Washburn and Stephanie Weirich. Generalizing parametricity using information flow. In *LICS*, 2005.

[30] Stephanie Weirich. Type-safe cast. *JFP*, 14(6), 2004.