

Relational Parametricity for Computational Effects

Rasmus Ejlers Møgelberg* Alex Simpson†
LFCS, School of Informatics
University of Edinburgh, Scotland, UK

Abstract

According to Strachey, a polymorphic program is parametric if it applies a uniform algorithm independently of the type instantiations at which it is applied. The notion of relational parametricity, introduced by Reynolds, is one possible mathematical formulation of this idea. Relational parametricity provides a powerful tool for establishing data abstraction properties, proving equivalences of datatypes, and establishing equalities of programs. Such properties have been well studied in a pure functional setting. Many programs, however, exhibit computational effects. In this paper, we develop a framework for extending the notion of relational parametricity to languages with effects.

1. Introduction

The theory of *relational parametricity*, proposed by Reynolds [21], provides a powerful framework for establishing properties of polymorphic programs and their types. Such properties include the “theorems for free” of Wadler [26], universal properties for datatype encodings, and representation independence properties for abstract datatypes. These results are well established, see e.g. [18], for the pure Girard/Reynolds second-order λ -calculus (a.k.a. system F) which provides a concise yet remarkably powerful calculus of typed total functions.

The generalisation of relational parametricity to richer calculi can be problematic. Even the addition of recursion (hence nontermination) causes difficulties, since the fixed-point property of recursion is incompatible with certain consequences of relational parametricity as usually formulated.¹ This issue led Plotkin [17] to propose using second-order linear type theory as a framework for combining parametricity and recursion. This idea has been developed in an

operational setting in [1] and in a denotational setting in [2]. One of its many good properties is that it supports a rich collection of polymorphic datatype encodings with the desired universal properties following from relational parametricity.

The addition of recursion is just one possible extension of second-order λ -calculus. For example, in [14], Parigot (implicitly) considers an orthogonal extension obtained by adding control operators. Recently, M. Hasegawa [5] has developed a syntactic account of relational parametricity for Parigot’s calculus. An intriguing fact he observes is that, even though the technical frameworks for the two approaches are quite different, there are striking analogies between his “focal” parametricity and Plotkin’s linear parametricity. Accordingly, Hasegawa poses the question of whether it is possible to find a unifying framework for relational parametricity that includes both his work and Plotkin’s linear parametricity as special cases.

In this paper we answer Hasegawa’s question by providing a general theory of relational parametricity for computational effects. Not only does our approach generalise both Plotkin’s and Hasegawa’s, but it also applies across the full range of computational effects (e.g., nondeterminism, probabilistic choice, input/output, side effects, exceptions, etc.).

We build on the work of Moggi [12, 13], who proposed incorporating effects into type theory by adding a new type constructor for typing “computations” rather than values. For every type B , one has a new type $!B$ (our non-standard notation is justified in Section 5) whose elements represent computations that (potentially) return values in B , and which (possibly) perform effects along the way. Semantically, $!$ is interpreted using a *computational monad* that encapsulates the relevant kinds of effect.

In order to obtain an account of relational parametricity for monads, one needs to solve a problem. Basic to relational parametricity is the idea of treating types as relations. Polymorphic functions are required to preserve derived relations under all possible instantiations of relations to type variables. To extend this to computational effects it is necessary to determine how the operation $!$ determines a relation $!R \subseteq !A \times !B$ from any relation $R \subseteq A \times B$. That is one needs a “relational lifting” of the $!$ operation. The literature

*Research supported by EPSRC and the Danish Agency for Science, Technology and Innovation.

†Research supported by an EPSRC Advanced Research Fellowship.

¹Relational parametricity implies types form a cartesian closed category with finite sums, and any such category with fixed points is trivial.

contains two approaches to defining such a relational lifting for $!$ [4, 8] (although neither is presented in the context of polymorphism). In the present paper we instead side-step the issue in a surprising way: we show that, given the right choice of underlying type theory, $!$ is polymorphically definable in terms of more basic primitives.

Our type theory, which we call PE, is presented in Section 2. It is closely related to Levy’s system of *call by push-value* (CBPV) [9], which subsumes call-by-name and call-by-value calculi with effects. Levy emphasises the importance of having two general classes of type: *value types*, which classify “values”, and *computation types*, which classify “computations”. The intuitive difference between the two is that “a value *is*” and “a computation *does*”. Technically, this intuition is supported by a wealth of semantic and operational interpretations of the framework, see [9].

With general computation types at hand, one can give the $!$ constructor the following polymorphic definition:

$$!B \stackrel{\text{def}}{=} \forall \underline{X}. (B \rightarrow \underline{X}) \rightarrow \underline{X} \quad (\underline{X} \text{ not free in } B), \quad (1)$$

where importantly the type variable \underline{X} ranges over computation types only. As we shall see, the type constructors used in the definition all have natural relational interpretations, and hence the defined $!$ operation inherits an induced relational lifting.

In order to reason about parametricity in PE, we build a relationally parametric model of our calculus. Even in the case of ordinary second-order λ -calculus, the construction of parametric models is non-trivial. In our case, the interaction between value and computation types contributes significant additional complexities. To keep things as simple as possible, we work with a set-theoretic model, exploiting the fact that it is consistent to do so if one keeps to intuitionistic reasoning. The details are presented in Sections 3 and 4. As a first application of the model, we prove in Section 5 that the $!$ operator, as defined by (1) above, does indeed enjoy its expected universal property (Theorem 5.2).

In Section 6, we consider how to specialise the generic calculus PE to specific effects of interest. One useful form of specialisation recurs in many examples. It is common for effects to have associated operations that trigger and/or react to “effectful” behaviour. Typically, one would like to give an n -ary such operation the polymorphic type:

$$\forall X. (!X)^n \rightarrow !X . \quad (2)$$

For example, a binary nondeterministic choice operation forms a computation by choosing between two possible continuation computations. Also, the “handle” operation for an exception e , can be viewed as a binary operation where $\text{handle}^e(p, q)$ behaves like p unless p raises exception e , in which case q is executed. Since such operations are computed in a type-independent way, they are “parametric” in the informal sense of Strachey. We show that

such operations are also parametric according to our theory of relational parametricity. This involves two technical developments of independent interest. The first relates to recent work by Plotkin and Power [20], in which they observe that many operations on effects are “algebraic operations” in the sense of universal algebra. As Theorem 6.1, we obtain that n -ary algebraic operations are in one-to-one correspondence with (parametric) elements of type:

$$\forall \underline{X}. \underline{X}^n \rightarrow \underline{X} , \quad (3)$$

where again \underline{X} ranges over computation types. Thus algebraic operations can be incorporated within PE as constants of the above type (which is more informative than (2), since monadic types $!B$ are always computation types).

Not all useful operations on effects arise as algebraic operations; e.g., exception handling is a counterexample. However, exception handling can be added to PE using a different strengthening of (2) for its type:

$$\forall X. (!X)^2 \multimap !X . \quad (4)$$

This correctness of this typing is again based on a general result (Theorem 6.2) which characterises the (parametric) elements of the above type (the nature of the linear arrow is explained in the sequel) in terms of a naturality condition.

Finally, in Section 7, we outline the relationship between PE and other approaches to parametricity and effects. Plotkin’s linear parametricity arises as a specialisation of PE valid in the special case of “commutative” monads. We also briefly discuss how Hasegawa’s account of parametricity and control arises as a specialisation of PE. The details for this appear in a companion paper [10].

Acknowledgements We are indebted to Masahito Hasegawa for first suggesting that (1) should be a general phenomenon within a monadic framework incorporating both linear and continuation-passing settings as special cases. We thank him and Paul Levy for helpful discussions.

2. A polymorphic calculus

We start by defining the type theory PE for polymorphism and effects. As discussed in the introduction, following [9], PE contains both *value types* A, B, C, \dots and *computation types* $\underline{A}, \underline{B}, \underline{C}, \dots$. A central feature of our type theory is that we allow polymorphic type quantification over both value types and computation types. Accordingly, we use X, Y, Z, \dots to range over a countable set of value-type variables, and $\underline{X}, \underline{Y}, \underline{Z}, \dots$ to range over a disjoint countable set of computation-type variables. Value types and computation types are then mutually defined by:

$$\begin{aligned} B &::= X \mid B \rightarrow C \mid \forall X. B \mid \underline{X} \mid \underline{A} \multimap \underline{B} \mid \forall \underline{X}. B \\ \underline{A} &::= B \rightarrow \underline{A} \mid \forall X. \underline{A} \mid \underline{X} \mid \forall \underline{X}. \underline{A} \end{aligned}$$

$$\begin{array}{c}
\frac{}{\Gamma, x: \underline{B} \mid - \vdash x: \underline{B}} \qquad \frac{\Gamma, x: \underline{B} \mid \Delta \vdash t: \underline{C}}{\Gamma \mid \Delta \vdash \lambda x: \underline{B}. t: \underline{B} \rightarrow \underline{C}} \qquad \frac{\Gamma \mid \Delta \vdash s: \underline{B} \rightarrow \underline{C} \quad \Gamma \mid - \vdash t: \underline{B}}{\Gamma \mid \Delta \vdash s(t): \underline{C}} \\
\\
\frac{\Gamma \mid \Delta \vdash t: \underline{B}}{\Gamma \mid \Delta \vdash \Lambda X. t: \forall X. \underline{B}} \quad X \notin \text{ftv}(\Gamma, \Delta) \qquad \frac{\Gamma \mid \Delta \vdash t: \forall X. \underline{B}}{\Gamma \mid \Delta \vdash t(\underline{A}): \underline{B}[\underline{A}/X]} \\
\\
\frac{}{\Gamma \mid x: \underline{A} \vdash x: \underline{A}} \qquad \frac{\Gamma \mid x: \underline{A} \vdash t: \underline{B}}{\Gamma \mid - \vdash \lambda^\circ x: \underline{A}. t: \underline{A} \multimap \underline{B}} \qquad \frac{\Gamma \mid - \vdash s: \underline{A} \multimap \underline{B} \quad \Gamma \mid \Delta \vdash t: \underline{A}}{\Gamma \mid \Delta \vdash s(t): \underline{B}} \\
\\
\frac{\Gamma \mid \Delta \vdash t: \underline{B}}{\Gamma \mid \Delta \vdash \Lambda \underline{X}. t: \forall \underline{X}. \underline{B}} \quad \underline{X} \notin \text{ftv}(\Gamma, \Delta) \qquad \frac{\Gamma \mid \Delta \vdash t: \forall \underline{X}. \underline{B}}{\Gamma \mid \Delta \vdash t(\underline{A}): \underline{B}[\underline{A}/\underline{X}]}
\end{array}$$

Figure 1. Typing rules.

Note that the computation types form a subcollection of the value types. The intuition here is that any (active) computation has a corresponding (static) value, its “think”. In contrast to [9], we make this passage from computations to values syntactically invisible.

For semantic intuition, one should think of value types as representing sets, and of computation types as representing Eilenberg-Moore algebras for some computational monad on sets. Then $\underline{B} \rightarrow \underline{C}$ is the set of all functions. The special case $\underline{B} \rightarrow \underline{A}$ is a computation type because algebras are closed under powers, with the algebra structure defined pointwise. The type $\underline{A} \multimap \underline{B}$ represents the set of all algebra homomorphisms from \underline{A} to \underline{B} . In general, there is no natural algebra structure on this set, hence the type $\underline{A} \multimap \underline{B}$ is not a computation type. Finally $\forall X. \underline{B}$ and $\forall \underline{X}. \underline{B}$ are polymorphic types, with the polymorphism ranging over value types and computation types respectively. In either case, when \underline{B} is a computation type, the polymorphic type is again a computation type. This is justified by Proposition 4.1 below.

Our types, which are based on function spaces and polymorphism, are not directly comparable with Levy’s [9], which include sums and products. Nonetheless, we shall see in Section 7 that we can encode Levy’s calculus within ours. Given this, our calculus extends Levy’s with polymorphic types (cf. [9, §12.4]) and linear function types. In fact, the latter have a particularly nice explanation in terms of Levy’s stack-based operational framework, within which a value of type $\underline{A} \multimap \underline{B}$ can be understood as a stack turning a computation of type \underline{A} into a computation of type \underline{B} .

Having computation types as special value types allows us to base our type system on a single judgement form:

$$\Gamma \mid \Delta \vdash t: \underline{B} ,$$

where Γ and Δ are disjoint contexts of variable typings subject to the following conditions: either (i) Δ is empty, or

(ii) \underline{B} is a computation type and Δ has the form $x: \underline{A}$, where \underline{A} is also a computation type. Thus the context Δ , which, following [3], we call the *stoup* of the typing judgement, contains at most one typing assertion. When we want to be explicit about which of (i) or (ii) applies, we write:

- (i) $\Gamma \mid - \vdash t: \underline{B}$
- (ii) $\Gamma \mid x: \underline{A} \vdash t: \underline{B}$.

In the first case, the intuitive interpretation of t is as an arbitrary function from the product of all types in Γ to the type \underline{B} . In the second case, the interpretation of t is as a function from $\Gamma \times \underline{A}$ to \underline{B} that is an algebra homomorphism in its right-hand argument (i.e. for every fixed set of values for the Γ variables, the induced function from \underline{A} to \underline{B} is a homomorphism). From this interpretation, one sees why the stoup is restricted to computation types, and also why, when the stoup is nonempty, the result type is required to be a computation type. (Similar considerations are in fact familiar from other stoup-based calculi, e.g., Girard’s LU [3].)

The type system is presented in Figure 1. The side conditions refer to the set $\text{ftv}(\Gamma)$ of free type variables in a context Γ , which is defined in the obvious way. Of course, the type rules are restricted to apply only when the premises satisfy the conditions on judgements imposed above. In such cases, the rule conclusions also satisfy these conditions.

It is immediate that the type system for value types extends the standard second-order λ -calculus of Girard and Reynolds. Indeed, the typing rules for the relevant types (\underline{X} , $\underline{B} \rightarrow \underline{C}$ and $\forall X. \underline{B}$), when restricted to the case with empty stoup, are just the usual ones. It is well-known that the second-order λ -calculus is powerful enough to encode many type constructors including products, sums, inductive and coinductive types. We include those definitions we shall need later in Figure 2.

$$\begin{aligned}
1 &=_{\text{def}} \forall X. X \rightarrow X \\
A \times B &=_{\text{def}} \forall X. (A \rightarrow B \rightarrow X) \rightarrow X \quad (X \notin \text{ftv}(A, B)) \\
0 &=_{\text{def}} \forall X. X \\
A + B &=_{\text{def}} \forall X. (A \rightarrow X) \rightarrow (B \rightarrow X) \rightarrow X \quad (X \notin \text{ftv}(A, B))
\end{aligned}$$

Figure 2. Definable value types

3. Semantic setting

In the previous section, we appealed to semantic intuition by explaining value types as sets and computation types as algebras for a monad on sets. Unfortunately, this intuition runs into the technical problem that there are no set-theoretic models of polymorphism [22]. However, it was shown by Pitts [15] that set-theoretic models of polymorphism are possible if *intuitionistic* set theory is used rather than ordinary *classical* set theory. We shall exploit this by working with such an intuitionistic set-theoretic model. The advantage of this strategy is that the set-theoretic framework allows the development to concentrate entirely on the difficulties inherent in defining a suitable notion of relational parametricity, which are formidable in themselves, rather than on incidental details specific to a particular concrete model. Our approach results in no loss of generality. All denotational models of relational parametricity of which we are aware can be exhibited as full subcategories of models of intuitionistic set theory.

Henceforth in this paper, we use Friedman’s Intuitionistic Zermelo-Fraenkel set theory (IZF) as our meta-theory, see e.g. [25]. IZF is the established intuitionistic counterpart of classical Zermelo-Fraenkel set theory (ZF). Just as in ordinary mathematics one works informally in ZF, we shall work similarly informally within IZF. Readers who are not familiar with IZF and the distinctions between intuitionistic and classical reasoning will anyway be able to follow the development, since IZF is a subtheory of ZF. However, such readers will have to place their trust in the authors that the reasoning principles of IZF are never violated.

Value types will be modelled as sets, but it is known that it is not possible to interpret types in the second-order λ -calculus as arbitrary sets [16]. Thus we require a collection of special sets for interpreting types. Such special sets need to be closed under the set-theoretic operations used in the interpretation. Accordingly, we assume that we have a full subcategory \mathcal{C} of the category **Set** of sets that satisfies:

- (C1) If $A \in \mathcal{C}$ and $A \cong B$ in **Set** then $B \in \mathcal{C}$.
- (C2) For any set-indexed family $\{A_i\}_{i \in I}$ of sets in \mathcal{C} , the set-theoretic product $\prod_{i \in I} A_i$ is again in \mathcal{C} .
- (C3) Given $A, B \in \mathcal{C}$ and functions $f, g: A \rightarrow B$, the equalizer $\{x \in A \mid f(x) = g(x)\}$ is again in \mathcal{C} .

- (C4) There is a set \mathbf{C} of objects of \mathcal{C} such that, for any $A \in \mathcal{C}$, there exists $B \in \mathbf{C}$ with $B \cong A$.

By items (C2) and (C3), the category \mathcal{C} is small-complete with limits inherited from **Set**. Since function spaces are powers, for any set A and any $B \in \mathcal{C}$, the function space B^A is in \mathcal{C} , i.e. \mathcal{C} is an *exponential ideal* of **Set**. In particular, \mathcal{C} is cartesian closed. By (C1), the category \mathcal{C} is not a small category. However, by (C4) it is weakly equivalent to its small full subcategory on the set of objects \mathbf{C} .

In classical set theory, the above conditions imply that every object in \mathcal{C} is either the empty set or a singleton set. The reason we work in IZF is that this renders it consistent for \mathcal{C} to be an interesting category. Indeed, it is consistent for the natural numbers to be an object of \mathcal{C} . This consistency property derives from the work of Hyland *et. al.* on small-complete small categories [6, 7]. However, our perspective is different. Rather than assuming a small category that is complete only in a restricted technical sense [7, 23], our category \mathcal{C} is genuinely complete, but only weakly equivalent to a small category. This approach, which simplifies the development, is taken from [24].

According to our informal explanation of computation types in Section 2, they should be interpreted as Eilenberg-Moore algebras for a monad T on \mathcal{C} . For any such monad T , the category \mathcal{A} of algebras comes with a forgetful functor $U: \mathcal{A} \rightarrow \mathcal{C}$ and the following properties are satisfied.

- (A1) U “weakly creates limits” in the following sense. For every diagram Δ in \mathcal{A} and limiting cone $\lim U(\Delta)$ in \mathcal{C} , there exists a specified limiting cone $\lim \Delta$ of Δ in \mathcal{A} such that $U(\lim \Delta) = \lim U(\Delta)$.
- (A2) U reflects isomorphisms (i.e. if Uf is an isomorphism in \mathcal{C} then f is an isomorphism in \mathcal{A}).
- (A3) For objects $\underline{A}, \underline{B}$ of \mathcal{A} , the hom-set $\mathcal{A}(\underline{A}, \underline{B})$ is an object of \mathcal{C} .
- (A4) There exists a set \mathbf{A} of objects of \mathcal{A} such that for every $\underline{A} \in \mathcal{A}$, there exists $\underline{B} \in \mathbf{A}$ with \underline{B} isomorphic to \underline{A} .

Indeed, (A1) and (A2) are standard. Property (A3) holds because $\mathcal{A}(\underline{A}, \underline{B})$ arises as an equalizer in \mathcal{C} of two evident functions $(U\underline{B})^{U\underline{A}} \rightarrow (U\underline{B})^{TU\underline{A}}$. Also, (A4) holds because the collection of algebra structures on objects of \mathbf{C} is a set.

The reason for identifying (A1)–(A4) is that, in order to interpret the calculus of Section 2, it is sufficient to work with any category \mathcal{A} and functor $U: \mathcal{A} \rightarrow \mathcal{C}$ satisfying (A1)–(A4) above.² Henceforth, we assume this situation.

It is convenient to maintain algebraic terminology for the category \mathcal{A} . Thus we call the objects of \mathcal{A} *algebras*. By (A1) and (A2), the functor U is faithful, thus we can identify the morphisms $\mathcal{A}(\underline{A}, \underline{B})$ with special functions from $U\underline{A}$ to $U\underline{B}$, which we call *homomorphisms*. We write $\underline{A} \multimap \underline{B}$ for

²In particular, the weakening of limit creation in (A1) is crucial to the application in [10].

the set of homomorphisms from \underline{A} to \underline{B} . (N.B. by (A3) the set $\underline{A} \multimap \underline{B}$ is an object of \mathcal{C} .)

In Section 4 we interpret the type theory of Section 2 using $U: \mathcal{A} \rightarrow \mathcal{C}$. In doing so, we formulate relational parametricity using binary relations in \mathcal{C} and \mathcal{A} . As usual, these are defined as subobjects of products. First, we review the basic properties of subobjects in \mathcal{C} and \mathcal{A} .

We write $\text{Sub}_{\mathcal{C}}(A)$ for the set of subobjects of A in \mathcal{C} . Since the inclusion $\mathcal{C} \hookrightarrow \mathbf{Set}$ preserves limits and hence monomorphisms, this is explicitly defined by:

$$\text{Sub}_{\mathcal{C}}(A) = \{B \in \mathcal{C} \mid B \subseteq A\}.$$

We call the elements of $\text{Sub}_{\mathcal{C}}(A)$ the \mathcal{C} -subsets of A .

Similarly, we write $\text{Sub}_{\mathcal{A}}(\underline{A})$ for the collection of subobjects of an algebra \underline{A} in \mathcal{A} . Because U preserves limits, every mono $\underline{B} \rightarrow \underline{A}$ in \mathcal{A} is mapped by U to a mono $UB \rightarrow UA$ in \mathcal{C} . Thus, for every $\underline{A} \in \mathcal{A}$, the functor U determines a function $\text{Sub}_{\mathcal{A}}(\underline{A}) \rightarrow \text{Sub}_{\mathcal{C}}(UA)$. By (A1) and (A2), this function preserves and reflects the ordering. We say that $A \subseteq UA$ carries a subalgebra if it represents a subobject in the image of the map $\text{Sub}_{\mathcal{A}}(\underline{A}) \rightarrow \text{Sub}_{\mathcal{C}}(UA)$ induced by U . In fact, $\text{Sub}_{\mathcal{A}}(\underline{A})$ is given explicitly by:

$$\text{Sub}_{\mathcal{A}}(\underline{A}) = \{B \in \mathcal{C} \mid B \subseteq UA \text{ and carries a subalg. of } \underline{A}\}.$$

We introduce notation for binary relations. For $A \in \mathcal{C}$, we write Δ_A for the diagonal (identity) relation in $\text{Sub}_{\mathcal{C}}(A \times A)$. Similarly, for $\underline{A} \in \mathcal{A}$, we write $\Delta_{\underline{A}}$ for the diagonal relation on $U\underline{A}$, which is indeed in $\text{Sub}_{\mathcal{A}}(\underline{A} \times \underline{A})$. For $f: A' \rightarrow A, g: B' \rightarrow B$ in \mathcal{C} and $R \in \text{Sub}_{\mathcal{C}}(A \times B)$ we write $(f, g)^{-1}R$ for $\{(x, y) \mid (f(x), g(y)) \in R\}$. Notice that if $f: \underline{A}' \multimap \underline{A}, g: \underline{B}' \multimap \underline{B}$ in \mathcal{A} and $Q \in \text{Sub}_{\mathcal{A}}(\underline{A} \times \underline{B})$ then $(f, g)^{-1}Q \in \text{Sub}_{\mathcal{A}}(\underline{A}' \times \underline{B}')$.

To formulate relational parametricity, we require two specified collections of *admissible* relations, one $\mathcal{R}_{\mathcal{C}}(A, B) \subseteq \text{Sub}_{\mathcal{C}}(A \times B)$ on objects of \mathcal{C} and one $\mathcal{R}_{\mathcal{A}}(\underline{A}, \underline{B}) \subseteq \text{Sub}_{\mathcal{A}}(\underline{A} \times \underline{B})$ on objects of \mathcal{A} . These are required to satisfy:

(R1) For each object A of \mathcal{C} the diagonal relation Δ_A is in $\mathcal{R}_{\mathcal{C}}(A, A)$ and likewise for each object \underline{A} of \mathcal{A} the diagonal $\Delta_{\underline{A}}$ is in $\mathcal{R}_{\mathcal{A}}(\underline{A}, \underline{A})$.

(R2) Admissible relations are closed under reindexing, i.e., if $R \in \mathcal{R}_{\mathcal{C}}(A, B)$ and $f: A' \rightarrow A, g: B' \rightarrow B$, then $(f, g)^{-1}R \in \mathcal{R}_{\mathcal{C}}(A', B')$ and if $Q \in \mathcal{R}_{\mathcal{A}}(\underline{A}, \underline{B})$ and $f: \underline{A}' \multimap \underline{A}, g: \underline{B}' \multimap \underline{B}$, then $(f, g)^{-1}Q \in \mathcal{R}_{\mathcal{A}}(\underline{A}', \underline{B}')$.

(R3) For any set of admissible \mathcal{C} - (respectively \mathcal{A} -)relations on the same pair of objects, the intersection is an admissible \mathcal{C} - (respectively \mathcal{A} -)relation.

(R4) $\mathcal{R}_{\mathcal{A}}(\underline{A}, \underline{B}) \subseteq \mathcal{R}_{\mathcal{C}}(U\underline{A}, U\underline{B})$.

(R1) and (R2) imply that graphs of functions are admissible, i.e., if $f: A \rightarrow B$ then $\langle f \rangle =_{\text{def}} \{(x, y) \mid f(x) = y\} \in \mathcal{R}_{\mathcal{C}}(A, B)$ and if $g: \underline{A} \multimap \underline{B}$ then $\langle g \rangle \in \mathcal{R}_{\mathcal{A}}(\underline{A}, \underline{B})$.

By a *parametric model of PE* we shall mean any category \mathcal{C} satisfying (C1)–(C4), together with a category \mathcal{A} and functor $U: \mathcal{A} \rightarrow \mathcal{C}$ satisfying (A1)–(A4) and collections $\mathcal{R}_{\mathcal{C}}$ and $\mathcal{R}_{\mathcal{A}}$ satisfying (R1)–(R4) above. The proposition below shows that every monad gives rise to a parametric model of PE. Thus the theory of relational parametricity for PE that we shall develop over such models is applicable to arbitrary computational monads.

Proposition 3.1 *Given \mathcal{C} satisfying (C1)–(C4) and a monad T on \mathcal{C} , let \mathcal{A} be the category of algebras for the monad, U the forgetful functor and define $\mathcal{R}_{\mathcal{C}}(A, B) = \text{Sub}_{\mathcal{C}}(A \times B)$ and $\mathcal{R}_{\mathcal{A}}(\underline{A}, \underline{B}) = \text{Sub}_{\mathcal{A}}(\underline{A} \times \underline{B})$. This data defines a parametric model of PE.*

Although the above is a useful general result, we comment that some applications of PE require a different choice of model. For example, the application of PE to control in [10] makes crucial use of the permitted flexibility in the choice of \mathcal{A} , U and $\mathcal{R}_{\mathcal{A}}$. (The situation is analogous to that for Levy's CBPV [9], where the natural adjunction model of control does not involve the Eilenberg-Moore category.)

4. Interpreting the calculus

In this section we interpret PE in any parametric model as defined in Section 3. As adumbrated there, a value type B will be interpreted as a set $\mathcal{C}[\![B]\!]$ in \mathcal{C} , and a computation type \underline{A} will be interpreted as an algebra $\mathcal{A}[\![\underline{A}]\!]$. In order to incorporate relational parametricity, we shall also give a second interpretation of a value type B as an admissible \mathcal{C} -relation $\mathcal{R}[\![B]\!]$. For computation types \underline{A} , it will hold automatically that $\mathcal{R}[\![\underline{A}]\!]$ is also an admissible \mathcal{A} -relation.

Given a set of type variables Θ , a Θ -environment is a function γ mapping every value-type variable $X \in \Theta$ to an object $\gamma(X)$ of \mathcal{C} , and every computation-type variable $\underline{X} \in \Theta$ to an object $\gamma(\underline{X})$ of \mathcal{A} . A *relational Θ -environment* is a tuple $\rho = (\rho_1, \rho_2, \rho_{\mathcal{R}})$, where: ρ_1, ρ_2 are Θ -environments; for every value-type variable $X \in \Theta$,

$$\rho_{\mathcal{R}}(X) \in \mathcal{R}_{\mathcal{C}}(\rho_1(X), \rho_2(X)) ;$$

and, for every computation-type variable $\underline{X} \in \Theta$,

$$\rho_{\mathcal{R}}(\underline{X}) \in \mathcal{R}_{\mathcal{A}}(\rho_1(\underline{X}), \rho_2(\underline{X})) .$$

For each value type $B(\Theta)$ (i.e. type B with $\text{ftv}(B) \subseteq \Theta$) and Θ -environment γ , we define an object $\mathcal{C}[\![B]\!]_{\gamma}$ of \mathcal{C} ; and, for each computation type $\underline{A}(\Theta)$ and Θ -environment γ , we define an object $\mathcal{A}[\![\underline{A}]\!]_{\gamma}$ of \mathcal{A} . Interdependently with the above, for each value type $B(\Theta)$ and relational Θ -environment ρ , we define an admissible \mathcal{C} -relation $\mathcal{R}[\![B]\!]_{\rho} \in \mathcal{R}_{\mathcal{C}}(\mathcal{C}[\![B]\!]_{\rho_1}, \mathcal{C}[\![B]\!]_{\rho_2})$. The definitions are given in Figure 3 (although see below for an important caveat). In these definitions, the products and powers used in the definition of

$$\begin{aligned}
\mathcal{C}[\mathbf{X}]_\gamma &= \gamma(\mathbf{X}) \\
\mathcal{C}[\mathbf{B} \rightarrow \mathbf{C}]_\gamma &= \mathcal{C}[\mathbf{C}]_\gamma^{C[\mathbf{B}]_\gamma} \\
\mathcal{C}[\forall \mathbf{X}. \mathbf{B}]_\gamma &= \left\{ \pi \in \prod_{A \in \mathcal{C}} \mathcal{C}[\mathbf{B}]_{\gamma[A/\mathbf{X}]} \mid \forall A, B \in \mathcal{C}, \forall R \in \mathcal{R}_C(A, B). \mathcal{R}[\mathbf{B}]_{\Delta_\gamma[R/\mathbf{X}]}(\pi_A, \pi_B) \right\} \\
\mathcal{C}[\underline{\mathbf{X}}]_\gamma &= U(\gamma(\underline{\mathbf{X}})) \\
\mathcal{C}[\underline{\mathbf{A}} \multimap \underline{\mathbf{B}}]_\gamma &= \mathcal{A}[\underline{\mathbf{A}}]_\gamma \multimap \mathcal{A}[\underline{\mathbf{B}}]_\gamma \\
\mathcal{C}[\forall \underline{\mathbf{X}}. \underline{\mathbf{B}}]_\gamma &= \left\{ \kappa \in \prod_{A \in \mathcal{A}} \mathcal{C}[\underline{\mathbf{B}}]_{\gamma[\underline{A}/\underline{\mathbf{X}}]} \mid \forall \underline{A}, \underline{B} \in \mathcal{A}, \forall Q \in \mathcal{R}_A(\underline{A}, \underline{B}). \mathcal{R}[\underline{\mathbf{B}}]_{\Delta_\gamma[Q/\underline{\mathbf{X}}]}(\kappa_{\underline{A}}, \kappa_{\underline{B}}) \right\} . \\
\mathcal{A}[\mathbf{B} \rightarrow \underline{\mathbf{A}}]_\gamma &= \mathcal{A}[\underline{\mathbf{A}}]_\gamma^{C[\mathbf{B}]_\gamma} \\
\mathcal{A}[\forall \mathbf{X}. \underline{\mathbf{A}}]_\gamma &= \left\{ \pi \in \prod_{A \in \mathcal{C}} \mathcal{A}[\underline{\mathbf{A}}]_{\gamma[A/\mathbf{X}]} \mid \forall A, B \in \mathcal{C}, \forall R \in \mathcal{R}_C(A, B). \mathcal{R}[\underline{\mathbf{A}}]_{\Delta_\gamma[R/\mathbf{X}]}(\pi_A, \pi_B) \right\} \\
\mathcal{A}[\underline{\mathbf{X}}]_\gamma &= \gamma(\underline{\mathbf{X}}) \\
\mathcal{A}[\forall \underline{\mathbf{X}}. \underline{\mathbf{A}}]_\gamma &= \left\{ \kappa \in \prod_{A \in \mathcal{A}} \mathcal{A}[\underline{\mathbf{A}}]_{\gamma[\underline{A}/\underline{\mathbf{X}}]} \mid \forall \underline{A}, \underline{B} \in \mathcal{A}, \forall Q \in \mathcal{R}_A(\underline{A}, \underline{B}). \mathcal{R}[\underline{\mathbf{A}}]_{\Delta_\gamma[Q/\underline{\mathbf{X}}]}(\kappa_{\underline{A}}, \kappa_{\underline{B}}) \right\} . \\
\mathcal{R}[\mathbf{X}]_\rho(x_1, x_2) &\Leftrightarrow \rho_{\mathcal{R}}(\mathbf{X})(x_1, x_2) \\
\mathcal{R}[\mathbf{B} \rightarrow \mathbf{C}]_\rho(f_1, f_2) &\Leftrightarrow \forall x_1 \in \mathcal{C}[\mathbf{B}]_{\rho_1}, x_2 \in \mathcal{C}[\mathbf{B}]_{\rho_2}. \mathcal{R}[\mathbf{B}]_\rho(x_1, x_2) \implies \mathcal{R}[\mathbf{C}]_\rho(f_1(x_1), f_2(x_2)) \\
\mathcal{R}[\forall \mathbf{X}. \mathbf{B}]_\rho(\pi_1, \pi_2) &\Leftrightarrow \forall A_1, A_2 \in \mathcal{C}, \forall R \in \mathcal{R}_C(A_1, A_2). \mathcal{R}[\mathbf{B}]_{\rho[R/\mathbf{X}]}((\pi_1)_{A_1}, (\pi_2)_{A_2}) \\
\mathcal{R}[\underline{\mathbf{X}}]_\rho(x_1, x_2) &\Leftrightarrow \rho_{\mathcal{R}}(\underline{\mathbf{X}})(x_1, x_2) \\
\mathcal{R}[\underline{\mathbf{A}} \multimap \underline{\mathbf{B}}]_\rho(h_1, h_2) &\Leftrightarrow \forall x_1 \in \mathcal{C}[\underline{\mathbf{A}}]_{\rho_1}, x_2 \in \mathcal{C}[\underline{\mathbf{A}}]_{\rho_2}. \mathcal{R}[\underline{\mathbf{A}}]_\rho(x_1, x_2) \implies \mathcal{R}[\underline{\mathbf{B}}]_\rho(h_1(x_1), h_2(x_2)) \\
\mathcal{R}[\forall \underline{\mathbf{X}}. \underline{\mathbf{B}}]_\rho(\kappa_1, \kappa_2) &\Leftrightarrow \forall \underline{A}_1, \underline{A}_2 \in \mathcal{A}, \forall Q \in \mathcal{R}_A(\underline{A}_1, \underline{A}_2). \mathcal{R}[\underline{\mathbf{B}}]_{\rho[Q/\underline{\mathbf{X}}]}((\kappa_1)_{\underline{A}_1}, (\kappa_2)_{\underline{A}_2}) .
\end{aligned}$$

Figure 3. Interpretation of Types

$\mathcal{C}[\mathbf{B}]_\gamma$ are the ones in \mathcal{C} , and those used in the definition of $\mathcal{A}[\underline{\mathbf{A}}]_\gamma$ are those in \mathcal{A} , as (weakly) created by U . We write Δ_γ for the relational Θ -environment that maps X (resp. \underline{X}) to $\Delta_{\gamma(X)}$ (resp. $\Delta_{\gamma(\underline{X})}$). We also use an obvious notation for update of environments. The algebras defined by $\mathcal{A}[\forall Y. \underline{\mathbf{A}}]_\gamma$ and $\mathcal{A}[\forall \underline{X}. \underline{\mathbf{A}}]_\gamma$ are the canonical algebras carried by the subsets of the product algebras.

The caveat referred to above is that the interpretations of the polymorphic types in Figure 3 do not, on the face of it, make sense, since they involve products over proper classes of objects. Remarkably, the definitions are rescued by the fact that, in them, the fictitious products are cut down to their parametric elements. For example, using condition (C4), one can show that each tuple π satisfying the *parametricity property*:

$$\forall A, B \in \mathcal{C}, \forall R \in \mathcal{R}_C(A, B). \mathcal{R}[\mathbf{B}]_{\Delta_\gamma[R/\mathbf{X}]}(\pi_A, \pi_B)$$

is determined by its elements $\{\pi_A\}_{A \in \mathcal{C}}$. Similarly, by (A4), each parametric κ is determined by $\{\kappa_{\underline{A}}\}_{\underline{A} \in \mathcal{A}}$. Thus, in either case, there are only set-many parametric tuples. For space reasons, we omit the proof of this claim (which essentially goes back to [24]), preferring to focus on applications of the model rather than on technicalities in its construction.

Proposition 4.1 $\mathcal{C}[\mathbf{B}]_\gamma$, $\mathcal{A}[\underline{\mathbf{A}}]_\gamma$ and $\mathcal{R}[\mathbf{B}]_\rho$ are well defined by Figure 3. Further, for every computation type $\underline{\mathbf{A}}$, it holds that $\mathcal{C}[\underline{\mathbf{A}}]_\gamma = U(\mathcal{A}[\underline{\mathbf{A}}]_\gamma)$ and $\mathcal{R}[\underline{\mathbf{A}}]_\rho \in \mathcal{R}_A(\mathcal{A}[\underline{\mathbf{A}}]_{\rho_1}, \mathcal{A}[\underline{\mathbf{A}}]_{\rho_2})$.

Lemma 4.2 (Identity extension) For any type $\mathbf{B}(\Theta)$ and Θ -environment γ , it holds that $\mathcal{R}[\mathbf{B}]_{\Delta_\gamma} = \Delta_{\mathcal{C}[\mathbf{B}]_\gamma}$.

Next, we define the interpretation of terms. Given a context Γ with all free type variables in Θ , a Θ - Γ -environment is a function defined on both the type variables in Θ and the term variables in Γ , such that the restriction of γ to Θ is a Θ -environment, and, for every type assignment $x : \mathbf{B}$ in Γ , it holds that $\gamma(x) \in \mathcal{C}[\mathbf{B}]_\gamma$. A term $\Gamma \mid \Delta \vdash_\Theta t : \mathbf{B}$ (i.e. such that $\text{ftv}(\Gamma, \Delta, t, \mathbf{B}) \subseteq \Theta$) is interpreted as an element $\llbracket t \rrbracket_\gamma \in \mathcal{C}[\mathbf{B}]_\gamma$, relative to any Θ - (Γ, Δ) -environment γ . The definition of $\llbracket t \rrbracket_\gamma$ is given in Figure 4. In the two clauses that apply to $t(\underline{\mathbf{A}})$, we distinguish between the cases for t of type $\forall \mathbf{X}. \mathbf{B}$ and $\forall \underline{\mathbf{X}}. \underline{\mathbf{B}}$. Note that the definition of $\llbracket s(t) \rrbracket_\gamma$ applies uniformly, whether s has type $\mathbf{B} \rightarrow \mathbf{C}$ or $\underline{\mathbf{A}} \multimap \underline{\mathbf{B}}$.

Proposition 4.3 If $\Gamma \mid \Delta \vdash_\Theta t : \mathbf{B}$ then:

1. (Well-definedness) For any Θ - (Γ, Δ) -environment γ , the value $\llbracket t \rrbracket_\gamma \in \mathcal{C}[\mathbf{B}]_\gamma$ is well defined.

$$\begin{aligned}
\llbracket x \rrbracket_\gamma &= \gamma(x) \\
\llbracket \lambda x : B. t \rrbracket_\gamma &= \llbracket \lambda^\circ x : \underline{A}. t \rrbracket_\gamma = (d : \mathcal{C}[\underline{B}]_\gamma \mapsto \llbracket t \rrbracket_{\gamma[d/x]}) \\
\llbracket s(t) \rrbracket_\gamma &= \llbracket s \rrbracket_\gamma(\llbracket t \rrbracket_\gamma) \\
\llbracket \Lambda X. t \rrbracket_\gamma &= \{ \llbracket t \rrbracket_{\gamma[A/X]} \}_{A \in \mathcal{C}} \\
\llbracket t : \forall X. B \rrbracket_\gamma &= (\llbracket t \rrbracket_\gamma)(\mathcal{C}[\underline{A}]_\gamma) \\
\llbracket \Lambda \underline{X}. t \rrbracket_\gamma &= \{ \llbracket t \rrbracket_{\gamma[\underline{A}/\underline{X}]} \}_{\underline{A} \in \mathcal{A}} \\
\llbracket t : \forall \underline{X}. B \rrbracket_\gamma &= (\llbracket t \rrbracket_\gamma)(\mathcal{A}[\underline{A}]_\gamma)
\end{aligned}$$

Figure 4. Interpretation of Terms

2. (Relational invariance) For any relational Θ -environment ρ , and Θ - (Γ, Δ) -environments γ_1, γ_2 extending ρ_1, ρ_2 respectively, define

$$\mathcal{R}[\Gamma]_\rho(\gamma_1, \gamma_2) \Leftrightarrow \forall x : A \in (\Gamma, \Delta). \mathcal{R}[\underline{A}]_\rho(\gamma_1(x), \gamma_2(x)).$$

Then $\mathcal{R}[\Gamma]_\rho(\gamma_1, \gamma_2)$ implies $\mathcal{R}[\underline{B}]_\rho(\llbracket t \rrbracket_{\gamma_1}, \llbracket t \rrbracket_{\gamma_2})$.

If $\Gamma \mid x : \underline{A} \vdash_\Theta t : \underline{B}$ then:

3. (Homomorphism property) For any Θ - Γ -environment γ , the function $d \in \mathcal{C}[\underline{A}]_\gamma \mapsto \llbracket t \rrbracket_{\gamma[d/x]}$ is a homomorphism from $\mathcal{A}[\underline{A}]_\gamma$ to $\mathcal{A}[\underline{B}]_\gamma$.

Our main application of the model will be to establish equalities between terms. Henceforth, for $\Gamma \mid \Delta \vdash s : B$ and $\Gamma \mid \Delta \vdash t : B$, we write $\Gamma \mid \Delta \vdash s = t : B$ to mean that $\llbracket s \rrbracket_\gamma = \llbracket t \rrbracket_\gamma$ for all appropriate γ .

5. Monadic types

In this section, we study the encoding of monadic types $!B$ in our calculus, as defined by equation (1) of Section 1. One sees immediately that $!B$ is always a computation type. We show that it enjoys the following derived introduction and elimination rules.

$$\frac{\Gamma \mid - \vdash t : B}{\Gamma \mid - \vdash !t : !B} \quad \frac{\Gamma \mid \Delta \vdash t : !B \quad \Gamma, x : B \mid - \vdash u : \underline{A}}{\Gamma \mid \Delta \vdash \text{let } !x \text{ be } t \text{ in } u : \underline{A}}$$

Indeed, for this simply define:

$$\begin{aligned}
!t &=_{\text{def}} \Lambda \underline{X}. \lambda p : B \rightarrow \underline{X}. p(t) \\
\text{let } !x \text{ be } t \text{ in } u &=_{\text{def}} t(\underline{A})(\lambda x : B. u) .
\end{aligned}$$

It is the above rules that motivate our notation for the $!$ type constructor, since these are simply restrictions of the usual rules for the exponential $!$ of intuitionistic linear logic.

As a first application of relational parametricity for our system, we show that $!B$ has the correct universal property for Moggi's monadic type. To keep the notation bearable,

we frequently omit semantic brackets, treating syntactic objects as the semantic elements they define, and we freely mix syntactic expressions with semantic values. For example, given any set A in \mathcal{C} , we simply write $!A$ rather than $\mathcal{C}[\!X]_{[A/X]}$ or $\mathcal{A}[\!X]_{[A/X]}$, referring to $!A$ as a set or as an algebra respectively when disambiguation is needed.

Lemma 5.1 1. If $\Gamma \mid - \vdash t : B$ and $\Gamma, x : B \mid - \vdash u : \underline{A}$ then $\Gamma \mid - \vdash \text{let } !x \text{ be } t \text{ in } u = u[t/x] : \underline{A}$.

2. $\Gamma \mid y : !A \vdash y = \text{let } !x \text{ be } y \text{ in } !x : !A$.

3. Suppose that $\Gamma \mid \Delta \vdash s : !A$, $\Gamma, x : A \mid - \vdash t : \underline{B}$ and $\Gamma \mid y : \underline{B} \vdash u : \underline{C}$, then $\Gamma \mid \Delta \vdash \text{let } !x \text{ be } s \text{ in } u[t/y] = u[\text{let } !x \text{ be } s \text{ in } t / y] : \underline{C}$.

Proof. Item 1 is a straightforward consequence of the semantic validity of beta equality.

For 2, we must show that $y = y(!A)(\lambda x : A. !x)$ at type $\forall \underline{X}. (A \rightarrow \underline{X}) \rightarrow \underline{X}$. By evident extensionality properties of the model, it suffices to show that, for any algebra \underline{B} and $f : A \rightarrow U\underline{B}$ in \mathcal{C} , we have $y(\underline{B})(f) = y(!A)(\lambda x : A. !x)(\underline{B})(f)$.

Consider the homomorphism $g : !A \multimap \underline{B}$ defined by $g(z) = z(\underline{B})(f)$. Then $\langle g \rangle \in \mathcal{R}_{\mathcal{A}}(!A, \underline{B})$. So, by parametricity,

$$((\Delta_A \rightarrow \langle g \rangle) \rightarrow \langle g \rangle) (y(!A), y(\underline{B})) . \quad (5)$$

For any $x \in A$, we have $g(!x) = (\Lambda \underline{X}. \lambda p. p(x))(\underline{B})(f) = f(x)$, i.e.

$$(\Delta_A \rightarrow \langle g \rangle) (\lambda x : A. !x, f) . \quad (6)$$

Combining (5) and (6), we obtain that

$$\langle g \rangle (y(!A)(\lambda x : A. !x), y(\underline{B})(f)) ,$$

i.e. $g(y(!A)(\lambda x : A. !x)) = y(\underline{B})(f)$. Thus it indeed holds that $y(!A)(\lambda x : A. !x)(\underline{B})(f) = y(\underline{B})(f)$.

For 3, $h = \lambda^\circ y : \underline{B}. u : \underline{C}$ is a homomorphism, so $\langle h \rangle \in \mathcal{R}_{\mathcal{A}}(\underline{B}, \underline{C})$. By parametricity, we have that

$$((\Delta_A \rightarrow \langle h \rangle) \rightarrow \langle h \rangle) (s(\underline{B}), s(\underline{C})) . \quad (7)$$

Consider $\lambda x : A. t : A \rightarrow \underline{B}$ and $\lambda x : A. u[t/y] : A \rightarrow \underline{C}$. Then, for $x \in A$, it holds that $h((\lambda x : A. t)(x)) = u[t/y] = (\lambda x : A. u[t/y])(x)$, i.e.

$$(\Delta_A \rightarrow \langle h \rangle) (\lambda x : A. t, \lambda x : A. u[t/y]) . \quad (8)$$

Combining (7) and (8), we obtain that

$$\langle h \rangle (s(\underline{B})(\lambda x : A. t), s(\underline{C})(\lambda x : A. u[t/y])) ,$$

i.e. $h(s(\underline{B})(\lambda x : A. t)) = s(\underline{C})(\lambda x : A. u[t/y])$. So indeed we have $u[\text{let } !x \text{ be } s \text{ in } t / y] = h(s(\underline{B})(\lambda x : A. t)) = s(\underline{C})(\lambda x : A. u[t/y]) = \text{let } !x \text{ be } s \text{ in } u[t/y]$. \square

For any set A in \mathcal{C} define $\eta_A : A \rightarrow !A$ by $\eta_A = \lambda x. !x$.

Theorem 5.2 *The function $\eta_A: A \rightarrow !A$ presents $!A$ as the free algebra over A , i.e. for any algebra \underline{A} and function $f: A \rightarrow U\underline{A}$, there exists a unique homomorphism $h: !A \rightarrow \underline{A}$ such that $h \circ \eta_A = f$. Indeed, h is given by $\lambda^{\circ}y. \text{let } !x \text{ be } y \text{ in } f(x)$.*

Proof. Clearly $\lambda^{\circ}y. \text{let } !x \text{ be } y \text{ in } f(x)$ is a homomorphism, and $(\lambda^{\circ}y. \text{let } !x \text{ be } y \text{ in } f(x)) \circ \eta_A = f$ because $\text{let } !x \text{ be } !x \text{ in } f(x) = f(x)$ by Lemma 5.1.1. For uniqueness, suppose h is such that $h \circ \eta_A = f$. Then

$$h(y) = h(\text{let } !x \text{ be } y \text{ in } !x) \quad (\text{Lemma 5.1.2})$$

$$= \text{let } !x \text{ be } y \text{ in } h(!x) \quad (\text{Lemma 5.1.3})$$

$$= \text{let } !x \text{ be } y \text{ in } f(x) \quad (h \circ \eta_A = f) ,$$

as required. \square

It follows from the above theorem that the operation mapping A to the algebra $!A$ is the object part of a functor $F: \mathcal{C} \rightarrow \mathcal{A}$ left adjoint to U . We write T for the associated monad UF on \mathcal{C} .

The bijective correspondence of Theorem 5.2 can be expressed in the type theory PE as an isomorphism of (value) types between $!A \rightarrow \underline{B}$ and $A \rightarrow \underline{B}$. Thus we have a Girard decomposition of function spaces with computation type codomains, further motivating the $!$ notation.

We end this section with a characterisation of the induced relational lifting of the $!$ type constructor.

Proposition 5.3 *Suppose A, B are objects of \mathcal{C} and $R: \mathcal{R}_{\mathcal{C}}(A, B)$ is a relation. Then $!R: \mathcal{R}_A(!A, !B)$ is the smallest admissible A -relation containing all pairs of the form $(\eta(x), \eta(y))$ for $(x, y) \in R$.*

6. Specialising the calculus to specific effects

The type theory PE is a generic calculus for effects since the type $!B$ can be interpreted as an arbitrary monad, and no further effect-specific features are included. In this regard, PE is analogous to Moggi’s computational λ -calculus [12], computational metalanguage [13] and Levy’s call-by-push-value [9]. As with those calculi, specific effects can be incorporated by specialising the calculus appropriately. In this section we consider various such specialisations, emphasising, in particular, the interaction with parametricity.

In a recent programme of research [20], Plotkin and Power have shown that many monads of computational interest can be profitably viewed as free algebra constructions for equational theories. This approach arises naturally from a computational viewpoint: the “algebraic operations” used to specify the theory correspond to programming primitives that cause effects, and the equational theory simply expresses natural behavioural equivalences between such

primitives. We begin this section with an analysis of how to specialise PE to the case of such “algebraic effects”.

Our approach is justified by a general theorem, which we now present. As one of their central results about algebraic effects, Plotkin and Power establish a one-to-one correspondence between “algebraic operations” and (what they call) “generic effects” [19]. The theorem below reformulates this correspondence in our setting, and adds a third equivalent induced by our polymorphic description of monadic types. We shall apply this third equivalent to obtain the correct polymorphic typing for algebraic operations in effect-specific specialisations of PE.

Theorem 6.1 *For any set A in \mathcal{C} , there are one-to-one correspondences between:*

1. “algebraic operations of arity A ”, i.e. natural transformations from the functor $(U(-))^A: \mathcal{A} \rightarrow \mathcal{C}$ to U ,
2. “generic effects over A ”, i.e. elements of TA , and
3. “polymorphic computation type operations of arity A ”, i.e. elements of $\forall \underline{X}. (A \rightarrow \underline{X}) \rightarrow \underline{X}$.

The simplifications in the formulation of statement 1 above, compared with [19], are due to our set-theoretic setting, which renders it unnecessary to consider issues relating to enrichment or tensorial strength. Also note that, by statement 2, the other two statements, in spite of appearances, depend only on the monad T on \mathcal{C} , not on how it is resolved into an adjunction $F \dashv U: \mathcal{A} \rightarrow \mathcal{C}$.

To illustrate how Theorem 6.1 informs the specialisation of PE to algebraic effects, we consider nondeterminism as a typical example. As in [20], nondeterministic choice is naturally formulated using a binary operation “or” satisfying the semilattice equations:

$$x \text{ or } x = x, \quad x \text{ or } y = y \text{ or } x, \quad x \text{ or } (y \text{ or } z) = (x \text{ or } y) \text{ or } z .$$

Define the category \mathcal{A}_{nd} of “nondeterministic algebras” to have, as objects, structures (A, or_A) where A is a set in \mathcal{C} and $\text{or}_A: A \times A \rightarrow A$ satisfies the semilattice equations, and, as morphisms from (A, or_A) to (B, or_B) , functions from A to B that are homomorphisms with respect to the “or” operations. It is easily verified that the obvious forgetful functor $U: \mathcal{A}_{\text{nd}} \rightarrow \mathcal{C}$ satisfies conditions (A1)–(A4).

Since the morphisms in \mathcal{A}_{nd} are homomorphisms, the operation mapping any nondeterministic algebra (A, or_A) to the function $\text{or}_A: A^2 \rightarrow A$ is an algebraic operation of arity 2 in the sense of statement 1 of Theorem 6.1. Thus, applying Theorem 6.1 and currying, one obtains a corresponding polymorphic operation:

$$\text{or} : \forall \underline{X}. \underline{X} \rightarrow \underline{X} \rightarrow \underline{X} .$$

Accordingly, nondeterministic choice can be incorporated in PE by adding a constant or , typed as above, to the type

theory. This example illustrates the general pattern for adding algebraic operations as polymorphic constants to our type theory, and readily adapts to the algebraic operations associated with other algebraic effects.

A limitation of the notion of algebraic operation is that there exist effect-specific programming primitives that are not algebraic operations. One well-known example of such a primitive is exception handling. Below, we show how exception handling may also be incorporated within our approach as a suitably typed polymorphic constant. The approach is justified by a general theorem, giving another instance of a coincidence between natural transformations and elements of polymorphic type.

Theorem 6.2 *For any $n \in \mathbb{N}$, there are one-to-one correspondences between:*

1. *Natural transformations from $(F(-))^n: \mathcal{C} \rightarrow \mathcal{A}$ to $F: \mathcal{C} \rightarrow \mathcal{A}$, and*
2. *elements of $\forall X. (n \rightarrow !X) \multimap !X$,*

where, in statement 2, we write n for the n -fold coproduct type $1 + \dots + 1$, as defined in Figure 2.

We now consider exception handling in detail. We assume we have a set E of exceptions with decidable equality (i.e. for all $e, e' \in E$ either $e = e'$ or $e \neq e'$). We also assume (for simplicity) that \mathcal{C} is closed under binary coproduct in \mathbf{Set} . We define the category \mathcal{A}_{exc} of “exception algebras” to have, as objects, structures $(A, \{\text{raise}_A^e\}_{e \in E})$ where $\text{raise}_A^e \in A$, and, as morphisms from $(A, \{\text{raise}_A^e\}_{e \in E})$ to $(B, \{\text{raise}_B^e\}_{e \in E})$, functions from A to B that map each raise_A^e to raise_B^e . Since the raise^e elements are algebraic constants (operations of arity 0), they can be added to PE as constants:

$$\text{raise}^e : \forall X. \underline{X} .$$

As is standard, the forgetful functor from \mathcal{A}_{exc} to \mathcal{C} , has as its left adjoint the functor F mapping A to the exception algebra $(A + E, \{\text{inr}(e)\}_{e \in E})$. For an exception $e \in E$, the handling operation over A is the function $\text{handle}_A^e: (F(A))^2 \rightarrow F(A)$ defined by

$$\text{handle}_A^e(p, q) = \begin{cases} p & \text{if } p \neq \text{inr}(e) \\ q & \text{if } p = \text{inr}(e) \end{cases} .$$

It is easily shown that this specifies a natural transformation from $(F(-))^2: \mathcal{C} \rightarrow \mathcal{A}_{\text{exc}}$ to $F: \mathcal{C} \rightarrow \mathcal{A}_{\text{exc}}$. In particular, the component handle_A^e of the natural transformation does lie in \mathcal{A}_{exc} because the interpretation of raise^e in the exception algebra $F(A)^2$ is the pair $(\text{inr}(e), \text{inr}(e))$. Thus, by Theorem 6.2, exception handling can be incorporated in PE by adding typed constants:

$$\text{handle}^e : \forall X. (2 \rightarrow !X) \multimap !X .$$

The main surprise with this typing is that exception handling is given a “linear” type. From this typing, one of course obtains an associated term of the expected (but less informative) type $\forall X. (2 \rightarrow !X) \rightarrow !X$.

Both Theorems 6.1 and 6.2 relate elements of certain polymorphic types with natural transformations between associated functors. In fact, more generally, for types that determine functors, parametricity implies naturality (cf. [18]). However, the exact correspondences between natural transformations and parametric elements established above depend heavily on the precise forms of types considered there.

The forms of n -ary operation considered in this section by no means exhaust the collection of operations of interest from an effects perspective. Control operators provide a particularly interesting class of examples, since their associated continuations monads do not naturally fit into the Plotkin-Power framework for algebraic effects. One way of specialising PE to the case of control is discussed briefly in the next section.

7. Relation to other systems

Several computational effects of interest, including non-termination, nondeterminism, and probabilistic choice, give rise to monads on \mathcal{C} that are *commutative*, cf. [13]. The collection of models of PE in which \mathcal{A} is the category of algebras for a commutative monad T is of special interest since, for such monads, the set of homomorphisms $\underline{A} \multimap \underline{B}$ between algebras $\underline{A}, \underline{B}$ carries a natural algebra structure which provides a closed structure on the category \mathcal{A} . For such models, it is thus natural to modify our type system by including $\underline{A} \multimap \underline{B}$ as a computation type. Making this adjustment, one obtains second-order intuitionistic linear type theory as the fragment of computation types:

$$\underline{X} \mid \underline{A} \multimap \underline{B} \mid \underline{A} \rightarrow \underline{B} \mid \forall X. \underline{A} . \quad (9)$$

Thus we obtain a rich collection of models for the type theory proposed by Plotkin as a foundation for combining polymorphism and recursion [17].

By the above, PE is naturally viewed as a generalisation of second-order intuitionistic linear type theory valid in a wider collection of models. A remarkable feature of second-order intuitionistic linear type theory, due to Plotkin, is that a rich collection of type constructors can be defined in terms of the three primitives in (9) above, cf. [17, 1, 2]. In fact, using well chosen variants of Plotkin’s definitions, a similar richness of definability is available in PE, see Figure 5 (which makes use of the definitions in Figure 2). We briefly discuss these encodings.

Semantically, because $U: \mathcal{A} \rightarrow \mathcal{C}$ weakly creates limits, algebras are closed under products in \mathcal{C} . Syntactically, however, the types 1 and $\underline{A} \times \underline{B}$ from Figure 2 are *not* computation types. Thus the alternative encodings 1° and $\underline{A} \times^\circ \underline{B}$ are

needed to obtain products of computation types as computation types. The types 0° and $\underline{A} \oplus \underline{B}$ from Figure 5 define respectively an initial object and binary coproduct in the category \mathcal{A} . This structure is *not* preserved by U , and coproducts of algebras behave very differently from coproducts of sets in \mathcal{C} . (The latter are implemented by the sum types in Figure 2.) The type $\underline{B} \cdot \underline{A}$ defines a $\mathcal{C}[\underline{B}]$ -fold copower of $\mathcal{A}[\underline{A}]$ in \mathcal{A} . Figure 5 also contains: existential types, $\exists^\circ X. \underline{A}$ and $\exists^\circ \underline{X}. \underline{A}$, packaged up as computation types; inductive computation types, $\mu^\circ \underline{X}. \underline{A}$; and coinductive computation types, $\nu^\circ \underline{X}. \underline{A}$. As is standard, the (co)inductive types rely on the functoriality of type expressions in their positive arguments. It is a consequence of relational parametricity that the above types all enjoy the correct universal properties. The arguments are carried out most naturally using a suitable logic for relational parametricity in PE, and will appear in a forthcoming paper [11].

A simple application of Figures 2 and 5 is to translate Levy's CBPV calculus [9] into PE. For this, coproducts and products of value types are translated using $+$ and \times from Figure 2, products of computation types are translated using \times° from Figure 5, Levy's F constructor is translated using $!$, and U is simply ignored.

Finally, we mention the case of control, which was one of the motivations for this work. Control primitives can be modelled naturally within PE by adding a polymorphic constant of type (using 0° from Figure 5):

$$\forall \underline{X}. ((\underline{X} \multimap 0^\circ) \multimap 0^\circ) \multimap \underline{X} .$$

The resulting theory is studied in a companion article [10], where it is shown that Hasegawa's [5] results on polymorphic definability in the second-order $\lambda\mu$ -calculus fall out as special cases of constructions from Figure 5.

In general, the approach of this paper should be applicable whenever there is interaction between polymorphism and effects. Such applications are a topic for future work.

References

- [1] G. Bierman, A. Pitts, and C. Russo. Operational properties of Lily, a polymorphic linear lambda calculus with recursion. *ENTCS*, 41:70–88, 2000.
- [2] L. Birkedal, R. E. Møgelberg, and R. L. Petersen. Linear Abadi & Plotkin logic. *Log. Meth. in Comp. Sci.*, 2, 2006.
- [3] J.-Y. Girard. On the unity of logic. *Annals of Pure and Applied Logic*, 59:201–217, 1993.
- [4] J. Goubault-Larrecq, S. Lasota, and D. Nowak. Logical relations for monadic types. In *Computer Science Logic*, pages 553–568. Springer LNCS 2571, 2002.
- [5] M. Hasegawa. Relational parametricity and control. *Logical Methods in Computer Science*, 2, 2006. Special issue for selected papers from LICS 2005.
- [6] J.M.E. Hyland. A small complete category. *Annals of Pure and Applied Logic*, 40:135 – 165, 1988.

$$\begin{aligned}
1^\circ &=_{\text{def}} \forall \underline{X}. 0 \multimap \underline{X} \\
\underline{A} \times^\circ \underline{B} &=_{\text{def}} \forall \underline{X}. ((\underline{A} \multimap \underline{X}) + (\underline{B} \multimap \underline{X})) \multimap \underline{X} \quad (\underline{X} \notin \text{ftv}(\underline{A}, \underline{B})) \\
0^\circ &=_{\text{def}} \forall \underline{X}. \underline{X} \\
\underline{A} \oplus \underline{B} &=_{\text{def}} \forall \underline{X}. (\underline{A} \multimap \underline{X}) \multimap (\underline{B} \multimap \underline{X}) \multimap \underline{X} \quad (\underline{X} \notin \text{ftv}(\underline{A}, \underline{B})) \\
\underline{B} \cdot \underline{A} &=_{\text{def}} \forall \underline{X}. (\underline{B} \multimap \underline{A} \multimap \underline{X}) \multimap \underline{X} \quad (\underline{X} \notin \text{ftv}(\underline{B}, \underline{A})) \\
\exists^\circ X. \underline{A} &=_{\text{def}} \forall \underline{Y}. (\forall X. (\underline{A} \multimap \underline{Y})) \multimap \underline{Y} \quad (\underline{Y} \notin \text{ftv}(\underline{A})) \\
\exists^\circ \underline{X}. \underline{A} &=_{\text{def}} \forall \underline{Y}. (\forall \underline{X}. (\underline{A} \multimap \underline{Y})) \multimap \underline{Y} \quad (\underline{Y} \notin \text{ftv}(\underline{A})) \\
\mu^\circ \underline{X}. \underline{A} &=_{\text{def}} \forall \underline{X}. (\underline{A} \multimap \underline{X}) \multimap \underline{X} \quad (\underline{X} \text{ +ve in } \underline{A}) \\
\nu^\circ \underline{X}. \underline{A} &=_{\text{def}} \exists^\circ \underline{X}. (\underline{X} \multimap \underline{A}) \cdot \underline{X} \quad (\underline{X} \text{ +ve in } \underline{A})
\end{aligned}$$

Figure 5. Definable computation types

- [7] J.M.E. Hyland, E. Robinson, and G. Rosolini. The discrete objects in the effective topos. *Proc. LMS.*, 3(60), 1990.
- [8] S. Katsumata. A semantic formulation of $\top\top$ -lifting and logical predicates for computational metalanguage. In *Computer Science Logic*, Springer LNCS 3634, 2005.
- [9] P. Levy. *Call-By-Push-Value*. Springer, 2004.
- [10] R.E. Møgelberg and A. Simpson. Relational Parametricity for Control Considered as a Computational Effect. In *Proc. MFPS XXIII, ENTCS 173:295–312*, 2007.
- [11] R.E. Møgelberg and A. Simpson. A logic for parametric polymorphism with effects. In preparation, 2007.
- [12] E. Moggi. Computational lambda-calculus and monads. In *LICS'89, Proc. 4th LICS Symposium*, pages 14–23, 1989.
- [13] E. Moggi. Notions of computation and monads. *Information and Computation*, 93(1), 1991.
- [14] M. Parigot. Strong normalization for second order classical natural deduction. *J. Symb. Logic*, 62:1461–1479, 1997.
- [15] A. Pitts. Polymorphism is set theoretic, constructively. In *Proc. CTCS*, pages 12–39. Springer LNCS 283, 1987.
- [16] A. Pitts. Non-trivial power types can't be subtypes of polymorphic types. In *Proc. 4th LICS Symp.*, pages 6–13, 1989.
- [17] G. Plotkin. Type theory and recursion (extended abstract). In *Proc. 8th LICS Symposium*, page 374, 1993.
- [18] G. Plotkin and M. Abadi. A logic for parametric polymorphism. *Proc. TLCA*, pp.361–375. Springer LNCS 664, 1993.
- [19] G. Plotkin and A.J. Power. Algebraic operations and generic effects. *Applied categorical Structures*, 11:69–94, 2003.
- [20] G. Plotkin and A.J. Power. Computational effects and operations: an overview. *ENTCS*, 73:149–163, 2004.
- [21] J. Reynolds. Types, abstraction and parametric polymorphism. In *Inf. Processing*, pp.513–523. N. Holland, 1983.
- [22] J. Reynolds. Polymorphism is not set-theoretic. In *Semantics of Data Types*. Springer LNCS 173, 1984.
- [23] E. Robinson. How complete is PER? In *Proc. 4th LICS Symposium*, pages 106–111, 1989.
- [24] G. Rosolini and A. Simpson. *Using Synthetic Domain Theory to Prove Operational Properties of a Polymorphic Programming Language Based on Strictness*. Preprint, 2004.
- [25] A. Ščedrov. Intuitionistic set theory. In *Harvey Friedman's Research on The Foundations of Mathematics*, pages 257–284. Elsevier Science Publishers, 1985.
- [26] P. Wadler. Theorems for free! In *Proc. 4th Int. Conf. on Funct. Prog. Languages and Computer Arch.* London, 1989.