Semantics for algebraic operations

Gordon Plotkin and John Power *

Division of Informatics, University of Edinburgh, King's Buildings, Edinburgh EH9 3JZ, Scotland

Abstract. Given a complete and cocomplete symmetric monoidal closed category V and a symmetric monoidal V-category C with cotensors and a strong V-monad T on C, we investigate axioms under which an ObCindexed family of operations of the form $\alpha_x : (Tx)^v \longrightarrow (Tx)^w$ provides semantics for algebraic operations, which may be used to extend the usual monadic semantics of the computational λ -calculus uniformly. We recall a definition for which we have elsewhere given adequacy results, and we show that an enrichment of it is equivalent to a range of other possible natural definitions of algebraic operation. We outline examples and non-examples and we show that our definition also enriches one for call-by-name languages with effects.

1 Introduction

Eugenio Moggi, in [8, 10], introduced the idea of giving a unified category theoretic semantics for computational effects such as nondeterminism, probabilistic nondeterminism, side-effects, input/output, and exceptions, by modelling each of them uniformly in the Kleisli category of an appropriate strong monad Ton a base category C with finite products. He supported that construction by developing the computational λ -calculus or λ_c -calculus, for which it provides a sound and complete class of models. The computational λ -calculus is essentially the same as the simply typed λ -calculus except for making a careful systematic distinction between computations and values. However, it does not contain operations, and operations are essential to any programming language. So here, in beginning to address that issue, we provide a unified semantics for algebraic operations, supported by equivalence theorems to indicate definitiveness of the axioms.

We distinguish between algebraic operations and arbitrary operations. The former are, in a sense we shall make precise, a natural generalisation, from Set to an arbitrary symmetric monoidal V-category C with cotensors, of the usual operations of universal algebra, now taking T to be a strong V-monad on C. The key point is that the operations

$$\alpha_x : (Tx)^v \longrightarrow (Tx)^w$$

(where $(-)^v$ denotes cotensor with an object v of V) are parametrically natural in the Kleisli V-category C_T . We could equally formulate this in terms of an

^{*} This work has been done with the support of EPSRC grant GR/M56333.

enriched version of closed Freyd-categories in the spirit of [1]. A preliminary version appears as [15]. The leading class of examples has T being generated by the operations subject to accompanying equations. Examples of such operations are nondeterministic choice operations, probabilistic nondeterministic choice operations, lookup and update operations for side-effects, read and write operations for interactive input/output, and operations for raising exceptions. A non-example is given by an operation for handling exceptions.

In a companion paper [14], we gave an unenriched version of the above definition together with a syntactic counterpart in terms of the computational λ calculus, and proved adequacy results. But such results leave some scope for a precise choice of appropriate semantic axioms. Enrichment allows us to capture examples such as local state [11], where V is the category $\omega - Cpo$ of ω -cpo's, i.e, the category of posets with sups of ω -chains, and C is the $\omega - Cpo$ -category $[W, \omega - Cpo]$, where W is a category of worlds. So, in an enriched setting, we prove a range of equivalence results, which we believe provide strong evidence for the above choice of axioms. Our most interesting result is essentially about the relationship between V-monads and Lawvere V-theories for suitable V [16]: the result, in a more general setting than usual, characterises algebraic operations as generic effects.

Moggi previously gave a semantic formulation of a notion of operation in [9], with an analysis based on his computational metalanguage, but he only required naturality of the operations in C, and we know of no way to provide operational semantics in such generality. Our various characterisation results do not seem to extend to such generality either. He, together with Benton and Hughes [2], also remarked on the construction that, to a generic effect, yields an operation. Evident further work is to consider how other operations such as those for handling exceptions should be modelled. That might involve going beyond monads, as Moggi has suggested to us: one possibility involves the use of distributive laws along the lines of [18]; another involves dyads [17]; possibly, some combination may be appropriate.

The paper is organised as follows. In Section 2, we recall and enrich the definition of algebraic operation given in [14] and we exhibit some simple reformulations of it. In Section 3, we give direct equivalent versions of these statements under the assumption that C is V-closed. In Section 4, we give a more substantial reformulation of the notion in terms of operations on homs, both when Cis closed and more generally when C is not closed. In Section 5, we characterise algebraic operations as generic effects. Finally, in Section 6, we characterise algebraic operations in terms of operations on the V-category T-Alg, as this gives an indication of how to incorporate call-by-name languages with computational effects into the picture. And we give conclusions and an outline of possible future directions in Section 7.

2 Algebraic operations and simple equivalents

In this section, we give an enriched version of the definition of algebraic operation as we made it in [14]. In that paper, we gave the definition and a syntactic counterpart in terms of the computational λ -calculus, and we proved adequacy results for the latter in terms of the former. Those results did not isolate definitive axioms for the notion of algebraic operation. So in this section, we start with a few straightforward equivalence results on which we shall build later.

We assume that V is a complete and cocomplete symmetric monoidal closed category: those are the conditions on V required for the preponderance of results of Kelly's definitive book [6] on enriched category theory. Implicitly using a larger universe, the category V-CAT of locally small V-categories has a symmetric monoidal structure, with $A \bigotimes B$ having object set $ObA \times ObB$, with

$$(A\bigotimes B)((a,b),(a',b')) = A(a,a') \circ B(b,b')$$

where \circ is the monoidal structure of V, with the evident composition. Using this tensor product on V-CAT, one can routinely define the notion of a symmetric monoidal V-category: it consists of a V-category C together with a V-functor

$$\otimes: C\bigotimes C \longrightarrow C$$

etcetera.

A monoidal V-category C is closed if for every object x of C, the V-functor $- \otimes x : C \longrightarrow C$ has a right V-adjoint. Note that if C is a monoidal Vcategory, then the underlying ordinary category C_0 of C is a monoidal category, similarly for symmetry and closedness. This is ultimately because the functor $(-)_o : V-CAT \longrightarrow CAT$ is monoidal. The notion of strong monad generalises readily from ordinary categories to monoidal V-categories: one asks for the strength to be V-natural. The notion of Kleisli exponential extends routinely too.

We henceforth assume C is a symmetric monoidal V-category with cotensors, and $\langle T, \eta, \mu, st \rangle$ is a strong V-monad on C with Kleisli V-exponentials. It is only for simplicity of exposition that we assume that C has all cotensors: typically, we only need finite cotensors, but occasionally, for instance in modelling state, we need more (see Section 5). To make such a size condition precise requires a corresponding size condition on V, the simplest being that V be locally finitely presentable: we do not want to clutter the paper with details, which may be found, for example, in [16]. We do not take C to be V-closed in general: we shall need to assume it for some later results, but not in general.

Given V and C as we have assumed them, we define parametrised lifting $(-)^{\dagger}$ by

$$C(y\otimes x,Tz) \xrightarrow{T} C(T(y\otimes x),T^2z) \xrightarrow{C(st,\,\mu_z)} C(y\otimes Tx,Tz)$$

This operation can be further extended by parametrisation in V. There is no danger of confusion, so we use the same notation $(-)^{\dagger}$ for the composite, for any

v in V,

$$C(y \otimes x, Tz) \xrightarrow{(-)^{\dagger}} C(y \otimes Tx, Tz) \xrightarrow{(-)^{v}} C((y \otimes Tx)^{v}, (Tz)^{v}) \longrightarrow C(y \otimes (Tx)^{v}, (Tz)^{v})$$

where the unlabelled map is given by composition with the comparison map determined by the universal property of cotensors.

Definition 1. An algebraic operation is an ObC-indexed family of maps

$$\alpha_x: (Tx)^v \longrightarrow (Tx)^w$$

such that the diagram

commutes.

If V = Set, the definition of algebraic operation requires v and w to be sets, typically finite ones n and m. To give the data for an algebraic operation is equivalent to giving $m \ ObC$ -indexed families of maps

$$\alpha_x : (Tx)^n \longrightarrow Tx$$

and the condition is the assertion that for each of these ObC-indexed families, for every map $f: y \otimes x \longrightarrow Tz$ in C, the diagram

commutes.

For some examples of algebraic operations where C = V = Set, let T be the nonempty finite power-set monad with the binary choice operation [12, 1]; alternatively, let T be the monad for probabilistic nondeterminism with a probabilistic choice operation [4,5]; or take T to be the monad for printing with printing operations [13]. Observe the non-commutativity in the latter example. One can, of course, generalise from *Set* to categories such as that of $\omega - Cpo$, for instance considering the various power-domains together with binary choice operators. One can also consider combinations of these, for instance to model internal and external choice operations. Several of these examples are treated in detail in [14].

To model local state, let $V = \omega - Cpo$, with C an $\omega - Cpo$ -category of the form $[W, \omega - Cpo]$, where W is a category of worlds. We study the examples of state and local state in more detail in Section 5.

There are several equivalent formulations of the coherence condition of the definition of algebraic operation. Decomposing it in a maximal way, we have

Proposition 1. An ObC-indexed family of maps

$$\alpha_x: (Tx)^v \longrightarrow (Tx)^w$$

is an algebraic operation if and only if

- 1. α is natural in C
- 2. α respects st in the sense that

$$\begin{array}{cccc} y \otimes (Tx)^v \longrightarrow (y \otimes Tx)^v \xrightarrow{st^v} (T(y \otimes x))^v \\ y \otimes \alpha_x \\ & & & & & \\ y \otimes (Tx)^w \longrightarrow (y \otimes Tx)^w \xrightarrow{st^w} (T(y \otimes x))^w \end{array}$$

commutes, where the unlabelled maps are comparison maps determined by the universal property of cotensors

3. α respects μ in the sense that

commutes.

Proof. It is immediately clear from our formulation of the definition and the proposition that the conditions of the proposition imply the coherence requirement of the definition. For the converse, to prove V-naturality in C, put y = I, the unit of the monoidal structure of C, use composition with η_z applied to C(x, z), and apply the coherence condition of the definition. For coherence with respect to st, take $f : y \otimes x \longrightarrow Tz$ to be $\eta_{y \otimes x}$. And for coherence with respect to μ , put y = I and take f to be id_{Tx} .

There are other interesting decompositions of the coherence condition of the definition too. In the above, we have taken T to be an endo-V-functor on C. But one often also writes T for the right V-adjoint to the canonical V-functor $J: C \longrightarrow C_T$ as the behaviour of the right adjoint on objects is given precisely by the behaviour of T on objects. So with this overloading of notation, we have V-functors $(T-)^v: C_T \longrightarrow C$ and $(T-)^w: C_T \longrightarrow C$, we can speak of V-natural transformations between them, and we have the following proposition.

Proposition 2. An ObC-indexed family of maps

$$\alpha_x: (Tx)^v \longrightarrow (Tx)^u$$

is an algebraic operation if and only if α is V-natural in C_T and α respects st.

In another direction, as we shall investigate further below, it is sometimes convenient to separate the μ part of the coherence condition from the rest of it. We can do that with the following somewhat technical result.

Proposition 3. An ObC-indexed family

$$\alpha_x : (Tx)^v \longrightarrow (Tx)^w$$

forms an algebraic operation if and only if α respects μ and

commutes, where $(-)^*$ is defined, parametrically in V, by the composition of $T: C(y \otimes x, z) \longrightarrow C(T(y \otimes x), Tz)$ with the composite

 $C(T(y \otimes x), Tz) \xrightarrow{C(st, Tz)} C(y \otimes Tx, Tz) \xrightarrow{(-)^v} C((y \otimes Tx)^v, (Tz)^v) \longrightarrow C(y \otimes (Tx)^v, (Tz)^v)$

3 Equivalent formulations if C is V-closed

For our more interesting results, we first assume C is V-closed, explain the results in those terms, and later drop the closedness condition and explain how to reformulate the results without essential change. So for the results in this section, we shall assume C is V-closed.

Let the V-closed structure of C be denoted by [-, -]. Given a V-functor $H : C \longrightarrow C$, an *enrichment* of H is a C-functor $K : C \longrightarrow C$ such that H is the underlying V-functor of K, i.e., H and K agree on objects and the monoidal V-functor $C(I, -) : C \longrightarrow V$ sends [Kx, Ky] to C(Hx, Hy), respecting

composition. Enrichment of a V-natural transformation does not alter the data but requires the stronger property of a commutativity in C rather than one in V. With these definitions, one can speak of the enrichment of a V-monad T to a C-monad.

Given a V-monad $\langle T, \eta, \mu \rangle$ on C, to give a V-strength for T is equivalent to giving an enrichment of T in C: given a strength, one has an enrichment

$$T_{x,y}: [x,y] \longrightarrow [Tx,Ty]$$

given by the transpose of

$$[x,y] \otimes Tx \xrightarrow{st} T([x,y] \otimes x) \xrightarrow{Tev} Ty$$

and given an enrichment of T, one has a V-strength given by the transpose of

$$x \longrightarrow [y, x \otimes y] \xrightarrow{T_{y, x \otimes y}} [Ty, T(x \otimes y)]$$

It is routine to verify that the axioms for a V-strength are equivalent to the axioms for an enrichment. So, given a V-strong V-monad $\langle T, \eta, \mu, st \rangle$ on C, the monad T is enriched in C, and so is the V-functor $(-)^v : C \longrightarrow C$.

The V-category C_T also canonically acquires an enrichment in C, i.e, the homobject $C_T(x, y)$ of C_T in V lifts to a homobject in C: the object [x, Ty] of C acts as a homobject, applying the V-functor $C(I, -) : C \longrightarrow V$ to it giving the V-homobject $C_T(x, y)$; composition

$$C_T(y,z) \circ C_T(x,y) \longrightarrow C_T(x,z)$$

in V lifts to a map in C

$$[y,Tz]\otimes [x,Ty] \longrightarrow [x,Tz]$$

determined by taking a transpose and applying evaluation maps twice and each of the V-strength and the multiplication once; and identities and the axioms for a V-category lift too.

The canonical V-functor $J: C \longrightarrow C_T$ becomes a C-enriched functor with a C-enriched right adjoint. The main advantage of the closedness condition for us is that it allows us to dispense with the parametrisation of the V-naturality, or equivalently with the coherence with respect to the V-strength, as follows.

Proposition 4. If C is V-closed, an ObC-indexed family

$$\alpha_x: (Tx)^v \longrightarrow (Tx)^w$$

forms an algebraic operation if and only if

commutes.

The left-hand vertical map in the diagram here is exactly the behaviour of the *C*-enriched functor $(T-)^w : C_T \longrightarrow C$ on homs, and, correspondingly, the top horizontal map is exactly the behaviour of the *C*-enriched functor $(T-)^v :$ $C_T \longrightarrow C$ on homs. So the coherence condition in the proposition is precisely the statement that α forms a *C*-enriched natural transformation from the *C*-enriched functor $(T-)^v : C_T \longrightarrow C$ to the *C*-enriched functor $(T-)^w : C_T \longrightarrow C$.

Proof. Given an object y of C, applying the V-functor $C(y, -) : C \longrightarrow V$ to the coherence condition here yields the coherence condition of the definition. The converse holds by the (ordinary) Yoneda lemma.

The same argument can be used to give a further characterisation of the notion of algebraic operation if C is V-closed by modifying Proposition 3. This yields

Proposition 5. If C is V-closed, an ObC-indexed family

$$\alpha_x : (Tx)^v \longrightarrow (Tx)^u$$

forms an algebraic operation if and only if α respects μ and

commutes.

This proposition says that if C is V-closed, an algebraic operation is exactly a C-enriched natural transformation from the C-enriched functor $(T-)^v : C \longrightarrow C$ to the C-enriched functor $(T-)^w : C \longrightarrow C$ that is coherent with respect to μ .

4 Algebraic operations as operations on homs

In our various formulations of the notion of algebraic operation so far, we have always had an ObC-indexed family

$$\alpha_x: (Tx)^v \longrightarrow (Tx)^w$$

and considered equivalent conditions on it under which it might be called an algebraic operation. In computing, this amounts to considering an operator on expressions. But there is another approach in which homs of the V-category C_T may be seen as primitive, regarding them as sets or ω -cpo's or the like of

programs. This was the underlying idea of the reformulation [1] of the semantics for finite nondeterminism of [12]. So we should like to reformulate the notion of algebraic operation in these terms. Proposition 4 allows us to do that. In order to explain the reason for the coherence conditions, we shall start by expressing the result assuming C is V-closed; after which we shall drop the closedness assumption and see how the result can be re-expressed using parametrised naturality.

We first need to explain an enriched version of the Yoneda lemma as in [6]. If D is a small C-enriched category, then D^{op} may also be seen as a C-enriched category. We will not assume C is complete. However, assuming for the moment that it was, we would have a C-enriched functor category $[D^{op}, C]$ and a C-enriched Yoneda embedding

$$Y_D: D \longrightarrow [D^{op}, C]$$

The C-enriched Yoneda embedding Y_D would be a C-enriched functor and it would be fully faithful in the strong sense that the map

$$D(x,y) \longrightarrow [D^{op}, C](D(-,x), D(-,y))$$

would be an isomorphism in the category C: see [6] for all the details. It follows by applying the V-functor $C(I, -) : C \longrightarrow V$ that this would induce an isomorphism from the homobject of V underlying D(x, y) to the object of V underlying the homobject from the C-enriched functor $D(-, x) : D^{op} \longrightarrow C$ to the C-enriched functor $D(-, y) : D^{op} \longrightarrow C$: if V = Set, the former object is the set of maps from x to y, and the latter is the set of C-enriched natural transformations from D(-, x) to D(-, y).

This is the result we need, except that, as we wrote, we do not want to assume that C is complete, and the C-enriched categories of interest to us are of the form C_T , so in general are not small. These are not major problems although they go a little beyond the scope of the standard formulation of enriched category theory in [6]: one can embed C into a larger universe C' just as one can embed Set into a larger universe Set' when necessary, and the required mathematics for the enriched analysis appears in [6]. We still have what can reasonably be called a Yoneda embedding of D into $[D^{op}, C]$, with both categories regarded as C'-enriched rather than C-enriched, and it is still fully faithful as a C'-enriched functor. However, we can formulate our result even without reference to C' by stating a restricted form of the enriched Yoneda lemma: letting $Fun_C(D^{op}, C)$ denote the V'-category (for a suitable extension V' of V) of C-enriched functors from D^{op} to C, the underlying V-functor

$$D \longrightarrow Fun_C(D^{op}, C)$$

of the Yoneda embedding is fully faithful.

We use this latter statement both here and in the following section. Now for our main result of this section under the assumption that C is V-closed.

Theorem 1. If C is V-closed, to give an algebraic operation is equivalent to giving an $ObC^{op} \times ObC$ family of maps

$$a_{y,x}: [y,Tx]^v \longrightarrow [y,Tx]^w$$

that is C-natural in y as an object of C^{op} and C-natural in x as an object of C_T , i.e., such that

$$\begin{array}{cccc} [y,Tx]^{v} \otimes [y',y] \xrightarrow{\cong} [y,(Tx)^{v}] \otimes [y',y] \xrightarrow{comp} [y',Tx]^{v} \\ a_{y,x} \otimes [y',y] & & & & & \\ [y,Tx]^{w} \otimes [y',y] \xrightarrow{\cong} [y,(Tx)^{w}] \otimes [y',y] \xrightarrow{comp} [y',Tx]^{w} \end{array}$$

and

$$\begin{array}{cccc} [x,Tz] \otimes [y,Tx]^v \longrightarrow ([x,Tz] \otimes [y,Tx])^v & \xrightarrow{comp_K^v} & [y,Tz]^v \\ [x,Tz] \otimes a_{y,x} & & & \\ [x,Tz] \otimes [y,Tx]^w \longrightarrow ([x,Tz] \otimes [y,Tx])^w & \xrightarrow{comp_K^v} & [y,Tz]^w \end{array}$$

commute, where comp is the C-enriched composition of C, the unlabelled isomorphisms of the first diagram are determined by the fact that $[y, -] : C \longrightarrow C$ is a right adjoint, so preserves cotensors, $comp_K$ is C-enriched Kleisli composition, and the unlabelled maps of the second diagram are determined by the universal property of cotensors.

Proof. It follows from our C-enriched version of the Yoneda lemma that to give the data together with the first axiom of the theorem is equivalent to giving an ObC-indexed family

$$\alpha: (Tx)^v \longrightarrow (Tx)^w$$

By a further application of our C-enriched version of the Yoneda lemma, it follows that the second condition of the theorem is equivalent to the coherence condition of Proposition 4.

As mentioned earlier, we can still state essentially this result even without the condition that C be closed. There are two reasons for this. First, for the paper, we have assumed the existence of Kleisli exponentials, as are essential in order to model λ -terms. But most of the examples of the closed structure of Cwe have used above are of the form [y, Tx], which can equally be expressed as the Kleisli exponential $y \Rightarrow x$. The Kleisli exponential routinely extends to a V-functor

$$- \Rightarrow -: C_T^{op} \times C_T \longrightarrow C$$

Second, in the above, we made one use of a construct of the form [y', y] with no T protecting the second object. But we can replace that by using the V'enriched Yoneda lemma to express the first condition of the theorem in terms of homobjects of V of the form $C(w \otimes y', y)$.

Summarising, we have

Corollary 1. To give an algebraic operation is equivalent to giving an $ObC^{op} \times ObC$ family of maps

$$a_{y,x}: (y \Rightarrow x)^v \longrightarrow (y \Rightarrow x)^v$$

in C, such that for all objects z' and y' of C, the diagram

commutes, and for every object z of C, the diagram

$$\begin{array}{c|c} (x \Rightarrow z) \otimes (y \Rightarrow x)^v \longrightarrow ((x \Rightarrow z) \otimes (y \Rightarrow x))^v & \xrightarrow{comp_K^v} (y \Rightarrow z)^v \\ (x \Rightarrow z) \otimes a_{y,x} & & & & \\ (x \Rightarrow z) \times (y \Rightarrow x)^w \longrightarrow ((x \Rightarrow z) \otimes (y \Rightarrow x))^w & \xrightarrow{comp_K^w} (y \Rightarrow z)^w \end{array}$$

commutes, where $comp_K$ is the canonical internalisation of Kleisli composition.

5 Algebraic operations as generic effects

In this section, we apply our formulation of the *C*-enriched Yoneda lemma to characterise algebraic operations in entirely different terms again as maps in C_T , i.e., in terms of generic effects. Observe that if *C* has a tensor \mathbf{v} of v with *I*, the *V*-functor $(T-)^v : C_T \longrightarrow C$ is isomorphic to the *V*-functor $\mathbf{v} \Rightarrow -: C_T \longrightarrow C$. If *C* is *V*-closed, the *V*-functor $\mathbf{v} \Rightarrow -$ enriches canonically to a *C*-enriched functor, namely the representable *C*-functor $C_T(\mathbf{v}, -) : C_T \longrightarrow C$, where C_T is regarded as a *C*-enriched category. So by Proposition 4 together with our *C*-enriched version of the Yoneda lemma, we immediately have

Theorem 2. If C is V-closed, the C-enriched Yoneda embedding induces a bijection between maps $\mathbf{w} \longrightarrow \mathbf{v}$ in C_T and algebraic operations

$$\alpha_x: (Tx)^v \longrightarrow (Tx)^w$$

This result is essentially an enriched version of the identification of maps in a Lawvere theory with operations of the Lawvere theory [16]. If C is locally finitely presentable as a closed category, one can define a notion of Lawvere C-theory, and prove it is equivalent to the notion of finitary C-monad on C, generalising the usual equivalence in the case that C = Set. Given a finitary C-monad T,

the corresponding Lawvere C-theory is given by the full sub-C-category of C_T determined by the finitely presentable objects. So the above result generalises the relationship between maps in the Lawvere C-theory with algebraic operations in two ways: the above result does not use finitariness and it allows V and C to differ.

In studying Lawvere theories, where V = Set, one typically restricts to natural numbers, but there are occasions when we want to drop the finitariness. For instance, this allows us to include an account of infinitary operations as one might use to model state as detailed below. For specific choices of C such as $\omega - Cpo$, one can also consider more exotic arities such as that given by Sierpinski space. For an enriched version of Lawvere's idea without the finitariness but with the restriction to C = V, see [3].

Once again, by use of parametrisation, we can avoid the closedness assumption on C here, yielding the stronger statement

Theorem 3. Functoriality of $- \Rightarrow -: C_T^{op} \times C_T \longrightarrow C$ in its first variable induces a bijection from the set of maps $\mathbf{w} \longrightarrow \mathbf{v}$ in C_T to the set of algebraic operations

$$\alpha_x: (Tx)^v \longrightarrow (Tx)^w$$

We could of course extend this to an isomorphism between the homobject $C_T(\mathbf{w}, \mathbf{v})$ and an object of appropriate algebraic operations of corresponding arity. We regard the theorem as the most interesting result of the paper. This result shows that to give an algebraic operation is equivalent to giving a generic effect, i.e., a constant of type the arity of the operation. Moreover, it follows from inspection of the category theoretic formulation of the notion of equation that equations to accompany the operator correspond to equations to be satisfied by the constant.

For example, to give a binary nondeterministic operator for a strong monad T is equivalent to giving a constant of type 2. For instance, let T be the nonempty finite powerset monad. Given a nondeterministic operator \lor , the constant is given by $true \lor false$, and given a constant c, the operator is given by $M \lor N =$ if c then M else N.

For another example, let L be a set of locations and let V be a set of values. We denote the exponential [L, V] by S, representing a set of states. (We ignore partiality here for simplicity of exposition.) Let C be the category $\omega - Cpo$ of ω -cpo's, and let T be the monad $(S \times -)^S$. Then, one naturally considers operations

$$lookup: (TX)^V \longrightarrow (TX)^L$$

and

$$update: TX \longrightarrow (TX)^{L \times V}$$

It is easiest to understand these operations in terms of the corresponding generic elements, which are of the form

$$lookup_q: L \longrightarrow TV$$

$$update_{a}: L \times V \longrightarrow T1$$

respectively. These are defined, using Currying freely, by

$$lookup_g(l,s) = (s,s(l))$$

and

$$update_q(l, v, s) = s(l \mapsto v)$$

respectively, which is exactly how *lookup* and *update* are supposed to behave.

One can readily extend this analysis to operations on an ω – *Cpo*-monad for local state on the ω – *Cpo*-enriched category $[N, \omega - Cpo]$ to model *lookup*, *update*, and *block*.

For a further example, one can consider a monad for interactive input and output and give generic operations $read : 1 \longrightarrow TI$ and $write : O \longrightarrow T1$, behaving as expected: see [9] for details.

6 Algebraic operations and the category of algebras

Finally, in this section, we characterise the notion of algebraic operation in terms of the V-category of algebras T-Alg. The co-Kleisli category of the comonad on T-Alg induced by the monad T is used to model call-by-name languages with effects, so this formulation gives us an indication of how to generalise our analysis to call-by-name computation or perhaps to some combination of call-by-value and call-by-name, cf [7].

If C is V-closed and has equalisers, generalising Lawvere, the results of the previous section can equally be formulated as equivalences between algebraic operations and operations

$$\alpha_{(A,a)}: U(A,a)^v \longrightarrow U(A,a)^w$$

natural in (A, a), where $U : T - Alg \longrightarrow C$ is the *C*-enriched forgetful functor: equalisers are needed in *C* in order to give an enrichment of T - Alg in *C*. We prove the result by use of our *C*-enriched version of the Yoneda lemma again, together with the observation that the canonical *C*-enriched functor $I : C_T \longrightarrow T - Alg$ is fully faithful. Formally, the result is

Theorem 4. If C is V-closed and has equalisers, the C-enriched Yoneda embedding induces a bijection between maps $\mathbf{w} \longrightarrow \mathbf{v}$ in C_T and C-enriched natural transformations

$$\alpha: (U-)^v \longrightarrow (U-)^w.$$

Combining this with Theorem 2, we have

Corollary 2. If C is closed and has equalisers, to give an algebraic operation

$$\alpha_x: (Tx)^v \longrightarrow (Tx)^u$$

is equivalent to giving a C-enriched natural transformation

$$\alpha: (U-)^v \longrightarrow (U-)^w.$$

and

One can also give a parametrised version of this result if C is neither closed nor complete along the lines for C_T as in the previous section. It yields

Theorem 5. To give an algebraic operation

$$\alpha_x: (Tx)^v \longrightarrow (Tx)^u$$

is equivalent to giving an Ob(T-Alg)-indexed family of maps

$$\alpha_{(A,a)}: U(A,a)^v \longrightarrow U(A,a)^w$$

such that commutativity of

implies commutativity of

7 Conclusions and Further Work

For some final comments, we note that little attention has been paid in the literature to the parametrised naturality condition on the notion of algebraic operation that we have used heavily here. And none of the main results of [14] used it, although they did require naturality in C_T . So it is natural to ask why that is the case.

For the latter point, in [14], we addressed ourselves almost exclusively to closed terms, and that meant that parametrised naturality of algebraic operations was not emphasised as we did not need a parameter for our main results. Had we given an equational theory in that paper, parametrised naturality would have been essential.

Regarding why parametrised naturality does not seem to have been addressed much in the past, observe that for C = Set, every monad has a unique strength, so parametrised naturality of α is equivalent to ordinary naturality of α . More generally, if the functor $C(1, -) : C \longrightarrow Set$ is faithful, i.e., if 1 is a generator in C, then parametrised naturality is again equivalent to ordinary naturality of α . That is true for categories such as *Poset* and that of ω -cpo's, which have been the leading examples of categories studied in this regard. The reason we have a distinction is because we have not assumed that 1 is a generator, allowing us to include examples such as toposes or *Cat*.

In future, we hope to address other operations that are not algebraic, such as one for handling exceptions. It seems unlikely that the approach of this paper extends directly. Eugenio Moggi has recommended we look beyond monads. In a sense suggested by Andrzej Filinski, operations such as those for handling exceptions seem to be destructors, whereas those we have considered here are constructors; and it seems that additional data and axioms are required to model the former, possibly along the lines of a distributive law as in [18]. We should also like to extend and integrate this work with work addressing other aspects of giving a unified account of computational effects. We note here especially Paul Levy's work [7] which can be used to give accounts of both call-by-value and call-by-name in the same setting, and work on modularity [17], which might also help with other computational effects.

References

- S. O. Anderson and A. J. Power, A Representable Approach to Finite Nondeterminism, *Theoret. Comput. Sci.*, Vol. 177, No. 1, pp. 3–25, 1997.
- N. Benton, J. Hughes, and E. Moggi, *Monads and Effects* APPSEM '00 Summer School, 2000.
- E. Dubuc, Kan extensions in enriched category theory, LNM, Vol. 145, Berlin: Springer-Verlag, 1970.
- C. Jones, *Probabilistic Non-Determinism*, Ph.D. Thesis, University of Edinburgh, Report ECS-LFCS-90-105, 1990.
- C. Jones and G. D. Plotkin, A Probabilistic Powerdomain of Evaluations, in *Proc.* 4th LICS, Asilomar, pp. 186–195, Washington: IEEE Press, 1989.
- G. M. Kelly, Basic Concepts of Enriched Category Theory, Cambridge: Cambridge University Press, 1982.
- P. B. Levy, Call-by-Push-Value: A Subsuming Paradigm, in *Proc. TLCA '99* (ed. J.-Y. Girard), LNCS, Vol. 1581, pp. 228-242, Berlin: Springer-Verlag, 1999.
- E. Moggi, Computational lambda-calculus and monads, in *Proc. LICS '89*, pp. 14–23, Washington: IEEE Press, 1989.
- E. Moggi, An abstract view of programming languages, University of Edinburgh, Report ECS-LFCS-90-113, 1989.
- E. Moggi, Notions of computation and monads, in *Inf. and Comp.* Vol. 93, No. 1, pp. 55–92, 1991.
- P. W. O'Hearn and R. D. Tennent, *Algol-like Languages*, Progress in Theoretical Computer Science, Boston: Birkhauser, 1997.
- G. D. Plotkin, A Powerdomain Construction, in SIAM J. Comput. Vol. 5, No. 3, pp. 452–487, 1976.
- 13. G. D. Plotkin, Domains, (http://www.dcs.ed.ac.uk/home/gdp/), 1983.
- G. D. Plotkin and A. J. Power, Adequacy for Algebraic Effects, in *Proc. FOSSACS* '01 (eds. F. Honsell and M. Miculan), LNCS, Vol. 2030, pp. 1–24, Berlin: Springer-Verlag, 2001.

- 15. G. D. Plotkin and A. J. Power, Semantics for Algebraic Operations (extended abstract), in *Proc. MFPS XVII* (eds. S. Brookes and M. Mislove), ENTCS, Amsterdam: Elsevier, 2001 (to appear).
- A. J. Power, Enriched Lawvere Theories, in *Theory and Applications of Categories*, pp. 83–93, 2000.
- A. J. Power and E. P. Robinson, Modularity and Dyads, in *Proc. MFPS XV* (eds. S. Brookes, A. Jung, M. Mislove and A. Scedrov), ENTCS Vol. 20, Amsterdam: Elsevier, 1999.
- D. Turi and G. D. Plotkin, Towards a mathematical operational semantics, in Proc. LICS '97, pp. 280–291, Washington: IEEE Press, 1997.