

Final Coalgebras as Greatest Fixed Points in ZF Set Theory*

Lawrence C. Paulson

lcp@cl.cam.ac.uk

Computer Laboratory, University of Cambridge, England

Received 1 March 1999

A special final coalgebra theorem, in the style of Aczel (1988), is proved within standard Zermelo-Fraenkel set theory. Aczel's Anti-Foundation Axiom is replaced by a variant definition of function that admits non-well-founded constructions. Variant ordered pairs and tuples, of possibly infinite length, are special cases of variant functions. Analogues of Aczel's solution and substitution lemmas are proved in the style of Rutten and Turi (1993). The approach is less general than Aczel's, but the treatment of non-well-founded objects is simple and concrete. The final coalgebra of a functor is its greatest fixedpoint.

Compared with previous work (Paulson, 1995a), iterated substitutions and solutions are considered, as well as final coalgebras defined with respect to parameters. The disjoint sum construction is replaced by a smoother treatment of urelements that simplifies many of the derivations.

The theory facilitates machine implementation of recursive definitions by letting both inductive and coinductive definitions be represented as fixedpoints. It has already been applied to the theorem prover Isabelle (Paulson, 1994).

* Thomas Forster suggested looking into Quine's work. Peter Aczel, Andrew Pitts and Daniele Turi offered considerable advice and help. I have used Paul Taylor's macros for commuting diagrams. K. Mukai commented on the text. Research funded by the ESPRIT Working Group 21900 'Types,' and GR/K57381 'Mechanizing Temporal Reasoning.'

Contents

1	Introduction	1
2	An Alternative Definition of Pairs and Functions	2
2.1	Quine's Ordered Pairs	2
2.2	Basic Definitions	3
2.3	The Role of Atoms	4
2.4	Basic Properties of the Cumulative Hierarchy	4
3	A Final Coalgebra	6
3.1	The Bifunctor Q and the Set U_X	6
3.2	U_X is a Final Q_X -Coalgebra	8
4	Solutions of Equations	9
4.1	Expressing Maps on $Q_X(Y)$	10
4.2	Solution and Substitution Lemmas	10
4.3	Special Final Coalgebra Theorem	13
5	Existence of Functors Uniform on Maps	14
5.1	The Constant Functor	14
5.2	Binary Product	14
5.3	Binary Sum	15
5.4	Sum of a Family of Sets	16
5.5	Product of a Family of Sets	16
5.6	Composition of Functors	17
5.7	The Identity Functor	17
6	Final Coalgebras with Parameters	18
7	Applications to Machine Proof	22
8	Conclusions	22

1. Introduction

A recurring issue in theoretical computer science is the treatment of infinite computations. One important approach is based upon the final coalgebra. This category-theoretic notion relates to the methods of bisimulation and coinduction, which are heavily used in concurrency theory (Milner, 1989), functional programming (Abramsky, 1990) and operational semantics (Milner and Tofte, 1991).

Aczel and Mendler (1989) and also Barr (1993) have proved that final coalgebras exist in set theory for large classes of naturally occurring functors. This might be supposed to satisfy most people's requirements. But Aczel (1988) has argued the case for a non-standard set theory in which infinite computations, and other non-well-founded phenomena, can be modelled directly. He proposes to replace set theory's Foundation Axiom (FA) by an Anti-Foundation Axiom (AFA) that guarantees the existence of solutions to $x = \{x\}$ and more generally of all systems of equations of the form $x_i = \{x_i, x_j, \dots\}$. His general final coalgebra theorem serves as a model construction to justify AFA.

Under AFA, a suitable functor F does not merely have a final coalgebra. That final coalgebra equals F 's greatest fixedpoint. This is the natural dual of the theorem that a functor's initial algebra is its least fixedpoint. These fixedpoints are exact, not up to isomorphism.

The elements of the final coalgebra are easily visualized. For instance, the functor $A \times -$ (the functor F such that $F(Z) = A \times Z$ on objects) yields the set of streams over A . The final coalgebra is also the greatest solution of $S = A \times S$. If $s \in S$ then

$$s = \langle a_1, s_1 \rangle, \quad s_1 = \langle a_2, s_2 \rangle, \quad s_2 = \langle a_3, s_3 \rangle, \dots ;$$

thus s is the infinite stream $\langle a_1, \langle a_2, \langle a_3, \dots \rangle \rangle \rangle$.

In standard set theory, FA outlaws infinite descents under the membership relation. Under the standard definition of ordered pair, we have $b \in \{a, b\} \in \langle a, b \rangle$. Infinitely nested pairs such as s above would create infinite \in -descents, and therefore do not exist: the greatest fixedpoint of $A \times -$ is the empty set. This is not the final coalgebra (which does exist).

The approach proposed in this paper is not to change the axiom system but to adopt new definitions of ordered pairs, functions, and derived concepts such as Cartesian products. Under the new definitions, the stream functor's final coalgebra is indeed its (exact) greatest fixedpoint and each stream is an infinite nest of pairs. Recursion equations are solved up to equality.

The approach handles non-well-founded tuples, and more generally ordered structures. But it does not model true non-well-founded sets, such as solutions of $x = \{x\}$. It does not work for the powerset functor, even with cardinality restrictions. Ironically, the approach requires FA.

Outline. The strategy is to construct a final coalgebra U , which plays the same role as the universe (V) under AFA. Then we can re-play the categorical proofs of Rutten and Turi (1993), generalizing them along the way. Section 2 presents basic motivation—Quine's ordered pairs and their generalization to functions—and proves some lemmas about the cumulative hierarchy, V_α . Section 3 defines the functor \mathcal{Q} and its greatest fixedpoint U and proves that U is a final \mathcal{Q} -coalgebra. Section 4 proves the solution and substitution lemmas for set equations and the special final coalgebra theorem. Section 5 discusses functors that are (or are not!) uniform on

maps. Section 6 considers final coalgebra definitions that take parameters. Section 7 discusses applications of the theory to machine proof. Section 8 presents conclusions.

2. An Alternative Definition of Pairs and Functions

Let us begin with informal motivation based on the work of Quine. The following section will make formal definitions.

2.1. Quine's Ordered Pairs

In ZF set theory, the ordered pair $\langle a, b \rangle$ is usually defined to be $\{\{a\}, \{a, b\}\}$. The rank of $\langle a, b \rangle$ is therefore two levels above those of a and b ; there are no solutions to $b = \langle a, b \rangle$. Quine (1966) has proposed a definition of ordered pair that need not entail an increase of rank. Quine's definition is complicated because (amongst other things) it avoids using standard ordered pairs. Retaining standard pairs lets us define Quine-like ordered pairs easily.

Let $\langle a, b \rangle$ denote the standard ordered pair of a and b . Let tuples of any length consist of ordered pairs nested to the right; thus $\langle a_1, \dots, a_n \rangle$ abbreviates $\langle a_1, \dots, \langle a_{n-1}, a_n \rangle \rangle$ for $n > 2$. Let $A \times B$ denote the standard Cartesian product $\{\langle a, b \rangle \mid a \in A \wedge b \in B\}$.

Define the variant ordered pair, $\langle a; b \rangle$ by

$$\langle a; b \rangle \equiv (\{0\} \times a) \cup (\{1\} \times b).$$

Note that $\langle a; b \rangle$ is just $a + b$, the disjoint sum of a and b (in set theory, everything is a set). The new pairing operator is obviously injective, which is a key requirement. Also, it admits non-well-founded constructions: we have $\langle 0; 0 \rangle = 0$ for a start. (As usual in set theory, the number zero is the empty set.)

The set equation $\langle A; z \rangle = z$ has a unique solution z , consisting of every (standard!) tuple of the form $\langle 1, \dots, 1, 0, x \rangle$ for $x \in A$. The infinite stream

$$\langle A_0; A_1; \dots; A_n; \dots \rangle$$

is the set of all standard tuples of the form

$$\langle \underbrace{1, \dots, 1}_n, 0, x \rangle$$

for $n < \omega$ and $x \in A_n$. Now $\langle a; b \rangle$ is continuous in a and b , in the sense that it preserves arbitrary unions; thus fixedpoint methods can solve recursion equations involving variant tupling.

Variant pairs can be generalized to a variant notion of function:

$$\tilde{\lambda}_{x \in A} b_x \equiv \bigcup_{x \in A} \{x\} \times b_x$$

Note that $\tilde{\lambda}_{x \in A} b_x$ is just $\Sigma_{x \in A} b_x$, the disjoint sum of a family of sets. Also note that $\langle b_0; b_1 \rangle$ is the special case $\tilde{\lambda}_{i \in 2} b_i$, since $2 = \{0, 1\}$. Replacing 2 by larger ordinals such as ω gives us a means of representing infinite sequences. More generally, non-standard functions can represent infinite collections that have non-well-founded elements.

Variant functions are not graphs. Merely replacing $\langle x, b_x \rangle$ by $\langle x; b_x \rangle$ in the usual definition

of function, obtaining $\{(x; b_x) \mid x \in A\}$, would not suffice. It still yields only well-founded constructions because the rank of such a set exceeds the rank of every b_x . For example, if $b = \{(0; b)\}$ then $\{1\} \times b \in b$, violating FA; thus $b = \{(0; b)\}$ has no solution.

Application of variant functions is expressed using the image operator “ $\tilde{\cdot}$ ”. It is easy to check that $(\tilde{\lambda}_{x \in A} b_x)$ “ $\{a\} = b_a$ if $a \in A$ ”. Also if R is a relation with domain A , then $R = \tilde{\lambda}_{x \in A} R$ “ $\{x\}$ ”. Every standard relation is a variant function, and vice versa. The set

$$\{f \subseteq A \times \bigcup B \mid \forall_{x \in A} f \text{ “ } \{x\} \in B\}$$

consists of all variant functions from A to B and will serve as our definition of variant function space, $A \tilde{\rightarrow} B$.

Since $\tilde{\lambda}_{x \in A} b_x$ is not the function’s graph, it does not determine the function’s domain. For instance, $\tilde{\lambda}_{x \in A} 0 = A \times 0 = 0$. Clearly $\tilde{\lambda}_{x \in A} 0 = \tilde{\lambda}_{x \in B} 0$ for all A and B . If $0 \in B$ then $A \tilde{\rightarrow} B$ will contain both total and partial functions: applying a variant function to an argument outside its domain yields 0.

2.2. Basic Definitions

Once we have defined the variant pairs and functions, we can substitute them in the standard definitions of Cartesian product, disjoint sum and function space. The resulting variant operators are decorated by a tilde: $\tilde{\times}$, $\tilde{+}$, $\tilde{\rightarrow}$, etc. Having both standard and variant operators is the simplest way of developing the theory. The standard operators relate the new concepts to standard set theory and they remain useful for defining well-founded constructions. But the duplication of operators may seem inelegant, and it introduces the risk of using the wrong one.

Definition 2.1. The *variant ordered pair* $\langle a; b \rangle$ is defined by

$$\langle a; b \rangle \equiv (\{0\} \times a) \cup (\{1\} \times b).$$

If $\{b_x\}_{x \in A}$ is an A -indexed family of sets then the *variant function* $\tilde{\lambda}_{x \in A} b_x$ is defined by

$$\tilde{\lambda}_{x \in A} b_x \equiv \bigcup_{x \in A} \{x\} \times b_x$$

The *variant Cartesian product*, *disjoint sum* and *partial function space* between two sets A and B are defined by

$$\begin{aligned} A \tilde{\times} B &\equiv \{(x; y) \mid x \in A \wedge y \in B\} \\ A \tilde{+} B &\equiv (\{1\} \tilde{\times} A) \cup (\{1; 1\} \tilde{\times} B) \\ A \tilde{\rightarrow} B &\equiv \{f \subseteq A \times \bigcup B \mid \forall_{x \in A} f \text{ “ } \{x\} \in B\} \end{aligned}$$

The operators $\tilde{\times}$ and $\tilde{\rightarrow}$ can be generalized to a family of sets as usual.

Definition 2.2. If $\{B_x\}_{x \in A}$ is an A -indexed family of sets then their *variant sum* and *product* are defined by

$$\begin{aligned} \tilde{\sum}_{x \in A} B_x &\equiv \{(x; y) \mid x \in A \wedge y \in B_x\} \\ \tilde{\prod}_{x \in A} B_x &\equiv \{f \subseteq A \times (\bigcup_{x \in A} B_x) \mid \forall_{x \in A} f \text{ “ } \{x\} \in B_x\} \end{aligned}$$

2.3. The Role of Atoms

A first attempt at exploiting these definitions is to fix an index set I and solve the equation $U = I \overset{\sim}{\rightarrow} U$. There is at least one solution, namely $U = \{0\}$, since $\tilde{\lambda}_{i \in I} 0 = 0$. But we cannot build up variant tuples starting from 0 as we can construct the distinct sets $\{0\}$, $\{0, \{0\}\}$, \dots . A variant tuple whose components are all the empty set is itself the empty set.

Since $I \overset{\sim}{\rightarrow} 0 = 0$ if $I \neq 0$, one possible solution to $U = I \overset{\sim}{\rightarrow} U$ is $U = 0$. Also $I \overset{\sim}{\rightarrow} \{0\} = \{0\}$. As it happens, $U = \{0\}$ is the greatest solution.

Proposition 2.3. If $U = I \overset{\sim}{\rightarrow} U$ then $U = 0$ or $U = \{0\}$.

Proof. Suppose not, for contradiction. Then U contains a non-empty element; there exist y_0 and x_0 with $y_0 \in x_0 \in U$. By the definition of $\overset{\sim}{\rightarrow}$ it follows that $y_0 = \langle i, y_1 \rangle$ where $i \in I$ and $y_1 \in x_1 \in U$ for some x_1 . Repeating this argument yields the infinite \in -descent $y_0 = \langle i, y_1 \rangle$, $y_1 = \langle i, y_2 \rangle$, $y_2 = \langle i, y_3 \rangle$, \dots , contradicting FA. \square

If tuples are to get built up, we must start with some atoms. To keep the atoms distinct from the variant tuples, each atom should contain some element that is not a (standard) pair. My earlier work (Paulson, 1995a) regarded one atom as sufficient, choosing 1 since $1 = \{0\}$ and the empty set is not a pair. It presented a final coalgebra theorem based upon the greatest solution of $U = \{1\} \cup (I \overset{\sim}{\rightarrow} U)$. The subsequent development closely followed Rutten and Turi (1993).

Aczel relies on urelements, as do other researchers (Moss and Danner, 1997), to formulate key results such as the solution lemma. He justifies this ‘expanded universe’ by a disjoint sum construction (Aczel, 1988, page 16), which Rutten and Turi (1993) neatly express as the greatest solution of $V_X = \mathcal{P}(X + V_X)$. However, they take this as the definition of V_X , replacing the expanded universe by its disjoint sum model. Abandoning urelements has many drawbacks. Desirable properties such as $V \subseteq V_X$ and $V_X \times V_X \subseteq V_X$ fail, requiring the frequent use of embeddings.

A more streamlined approach is to incorporate an arbitrary set X of atoms into the construction. The final coalgebra U_X is the greatest solution of $U_X = \text{Atoms}(X) \cup \{1\} \cup (I \overset{\sim}{\rightarrow} U_X)$, where $\text{Atoms}(X)$ is a suitable injection. These atoms are analogous to urelements, just as U_X is analogous to V_X , but we always work in standard ZF. The solution and substitution lemmas can be generalized to allow more than one set of indeterminates: we often work with U_X and U_Y , where possibly $Y = 0$, and write U_0 as U .

2.4. Basic Properties of the Cumulative Hierarchy

The following results are needed to prove closure and uniqueness properties in Sect. 3. Let α, β range over ordinals and λ, μ over limit ordinals. The *cumulative hierarchy* of sets is traditionally defined by cases: $V_0 = 0$, $V_{\alpha+1} = \mathcal{P}(V_\alpha)$, and if μ is a limit ordinal, $V_\mu = \bigcup_{\alpha < \mu} V_\alpha$. More convenient is the equivalent definition

$$V_\alpha \equiv \bigcup_{\beta < \alpha} \mathcal{P}(V_\beta).$$

Kunen (1980), Chapter III, is useful background reading; he writes $R(\alpha)$ for V_α . Here are some well-known facts.

Lemma 2.4. If α is an ordinal and μ is a limit ordinal then

$$\begin{aligned}\alpha &\subseteq V_\alpha \\ V_\alpha \times V_\alpha &\subseteq V_{\alpha+2} \\ V_\mu \times V_\mu &\subseteq V_\mu \\ V_\mu + V_\mu &\subseteq V_\mu\end{aligned}$$

The set V_μ is closed under the formation of variant tuples and functions.

Lemma 2.5. If $A \subseteq V_\mu$ and $b_x \subseteq V_\mu$ for all $x \in A$ then $\tilde{\lambda}_{x \in A} b_x \subseteq V_\mu$.

Proof. This follows by the definition of $\tilde{\lambda}$, monotonicity and the facts noted above:

$$\tilde{\lambda}_{x \in A} b_x = \bigcup_{x \in A} \{x\} \times b_x \subseteq \bigcup_{x \in V_\mu} \{x\} \times V_\mu \subseteq V_\mu \times V_\mu \subseteq V_\mu$$

□

Thus $V_{\mu+1}$ has closure properties for variant products and sums analogous to those of V_μ for standard products and sums. It is even closed under variant function space.

Lemma 2.6. Let μ be a limit ordinal.

- (a) If $A \subseteq V_\mu$ then $A \xrightarrow{\sim} V_{\mu+1} \subseteq V_{\mu+1}$.
- (b) $V_{\mu+1} \tilde{\times} V_{\mu+1} \subseteq V_{\mu+1}$.
- (c) $V_{\mu+1} \tilde{+} V_{\mu+1} \subseteq V_{\mu+1}$.

Proof. Obvious by the definitions and the previous lemma. □

These results will allow application of the Knaster-Tarski fixedpoint theorem to construct a final coalgebra. The next group of results will be used in the uniqueness proof.

Lemma 2.7. If $A \cap V_\alpha \subseteq B$ for every ordinal α then $A \subseteq B$.

Proof. By the Foundation Axiom, $V = \bigcup_\alpha V_\alpha$, where V is the universal class. Thus $A = \bigcup_\alpha (A \cap V_\alpha)$. If $A \cap V_\alpha \subseteq B$ for all α then $\bigcup_\alpha (A \cap V_\alpha) \subseteq B$ and the result follows. □

Using this lemma requires some facts about intersection with V_α .

Definition 2.8. A set A is *transitive* if $A \subseteq \mathcal{P}(A)$.

Lemma 2.9. V_α is transitive for every ordinal α .

Proof. See Kunen (1980), page 95. □

Now we can go down the cumulative hierarchy as well as up.

Lemma 2.10. If $\langle a, b \rangle \in V_{\alpha+1}$ then $a \in V_\alpha$ and $b \in V_\alpha$.

Proof. Suppose $\langle a, b \rangle \in V_{\alpha+1}$; this is equivalent to $\{\{a\}, \{a, b\}\} \in \mathcal{P}(V_\alpha)$. Thus $\{a, b\} \in V_\alpha$ and since V_α is transitive $\{a, b\} \subseteq V_\alpha$. □

Lemma 2.11. If $\{b_x\}_{x \in A}$ is an A -indexed family of sets then

- (a) $(\tilde{\lambda}_{x \in A} b_x) \cap V_{\alpha+1} \subseteq \tilde{\lambda}_{x \in A} (b_x \cap V_\alpha)$
- (b) $(\tilde{\lambda}_{x \in A} b_x) \cap V_\alpha \subseteq \bigcup_{\beta < \alpha} \tilde{\lambda}_{x \in A} (b_x \cap V_\beta)$

Proof. For (a) we have, by the previous lemma,

$$\begin{aligned} (\tilde{\lambda}_{x \in A} b_x) \cap V_{\alpha+1} &= \{\langle x, y \rangle \mid x \in A \wedge y \in b_x\} \cap V_{\alpha+1} \\ &\subseteq \{\langle x, y \rangle \mid x \in A \wedge y \in b_x \wedge y \in V_\alpha\} \\ &= \tilde{\lambda}_{x \in A} (b_x \cap V_\alpha). \end{aligned}$$

For (b) we have, by the definition of V_α and properties of unions,

$$\begin{aligned} (\tilde{\lambda}_{x \in A} b_x) \cap V_\alpha &= (\tilde{\lambda}_{x \in A} b_x) \cap \bigcup_{\beta < \alpha} \mathcal{P}(V_\beta) \\ &= \bigcup_{\beta < \alpha} (\tilde{\lambda}_{x \in A} b_x) \cap V_{\beta+1} \\ &\subseteq \bigcup_{\beta < \alpha} \tilde{\lambda}_{x \in A} (b_x \cap V_\beta). \end{aligned}$$

The last step is by (a) above. □

3. A Final Coalgebra

Rutten and Turi (1993), an excellent survey of final semantics, includes a categorical presentation of Aczel's main results. Working in the superlarge category of classes and maps between classes, they note that FA is equivalent to ' V is an initial \mathcal{P} -algebra' while AFA is equivalent to ' V is a final \mathcal{P} -coalgebra.' Put in this way, AFA certainly looks more attractive than the other anti-foundation axioms.

The present treatment of final semantics takes theirs as a starting point. Instead of assuming that V is a final \mathcal{P} -coalgebra, we can define a functor \mathcal{Q}^I , where I is an arbitrary index set, and construct a final \mathcal{Q}^I -coalgebra, called U^I , and obtain generalized forms of the solution and substitution lemmas. We finally arrive at the special final coalgebra theorem.

We shall work not in the category of classes but in the usual category **Set** of sets, which has standard functions as maps. While the former category allows certain statements to be expressed succinctly, it also requires numerous technical lemmas concerning set-based maps, etc. From the standpoint of mechanized proof, one must also bear in mind that classes have no formal existence under the ZF axioms, and class maps are two removes from existence.

3.1. The Bifunctor \mathcal{Q} and the Set U_X

Let I be an index set, which will remain fixed throughout the paper. A typical choice for I would be some limit ordinal such as ω . Note that $\omega \overset{\sim}{\rightarrow} A$ contains all ω -sequences over A ; we shall find that U^ω contains all ω -sequences over itself. Moreover, finite sequences can be represented by ω -sequences containing infinitely many 0s, because $0 \in U^I$ (see remark 3.7 below).

Incorporating atoms (urelements) requires an injection whose range is disjoint from all I -sequences. It suffices to include an element that is not a (standard) pair in its result, since every variant function is a standard relation.

Definition 3.1. The operators atm and Atoms are given by

$$\begin{aligned}\text{atm}(x) &\equiv \{2\} \cup (\{0\} \times x) \\ \text{Atoms}(X) &\equiv \{\text{atm}(x) \mid x \in X\}.\end{aligned}$$

Much is arbitrary in the definition of atm , but it is clearly injective, and $\text{atm}(x)$ is never a standard relation. Moreover, $\text{atm}(x) \neq 1$. The next step is to define the bifunctor $\mathcal{Q}_X^I(Y)$, where I is fixed and X and Y are sets. The intuition is that $\mathcal{Q}_X^I(Y)$ includes a copy of X (the atoms) and also includes I -sequences over Y . It also includes the element 1 to start things off, in case $X = 0$ (recall prop. 2.3). Its effect on a pair of maps is to apply one to the atoms and the other to the sequence elements.

Definition 3.2. The bifunctor $\mathcal{Q}_X^I : \mathbf{Set} \times \mathbf{Set} \rightarrow \mathbf{Set}$ is defined on objects by

$$\mathcal{Q}_X^I(Y) \equiv \text{Atoms}(X) \cup \{1\} \cup (I \rightrightarrows Y)$$

and on maps as follows. If $f : X \rightarrow X'$ and $g : Y \rightarrow Y'$ then $\mathcal{Q}_f^I(g) : \mathcal{Q}_X^I(Y) \rightarrow \mathcal{Q}_{X'}^I(Y')$ satisfies

$$\begin{aligned}\mathcal{Q}_f^I(g)(\text{atm}(x)) &\equiv \text{atm}(f(x)) \quad \text{for } x \in X \\ \mathcal{Q}_f^I(g)(1) &\equiv 1 \\ \mathcal{Q}_f^I(g)(\tilde{\lambda}_{i \in I} y_i) &\equiv \tilde{\lambda}_{i \in I} g(y_i).\end{aligned}$$

Also, $\mathcal{Q}_X(g)$ abbreviates $\mathcal{Q}_{\text{id}_X}(g)$.

It is easy to check that the functor preserves the identity map and composition. The next step is to define a set U_X^I to be the greatest solution of $U_X^I = \mathcal{Q}_X^I(U_X^I)$ and prove that U_X^I is a final \mathcal{Q}_X^I -coalgebra. Since $U_X^I = \text{Atoms}(X) \cup \{1\} \cup (I \rightrightarrows U_X^I)$ we may regard the elements of U_X^I as nested I -indexed tuples built up from 1, with further atoms from X .

To solve $U_X^I = \mathcal{Q}_X^I(U_X^I)$ we may apply the Knaster-Tarski fixedpoint theorem. This gives an explicit definition.

Definition 3.3. Let μ be a limit ordinal such that $I \subseteq V_\mu$ and $X \subseteq V_{\mu+1}$. Then

$$U_X^I \equiv \bigcup \{Z \mid Z \subseteq \mathcal{Q}_X^I(Z) \wedge Z \subseteq V_{\mu+1}\}.$$

Henceforth let us regard I as fixed and drop the superscripts. The next two results indicate that U_X really is a fixedpoint of \mathcal{Q}_X , in fact the greatest post-fixedpoint. This justifies proof by co-induction on U_X . The second result also confirms that the choice of the ordinal μ does not matter, provided it is at least the minimum specified.

For the remainder of this section, assume $X \subseteq V_{\mu+1}$.

Lemma 3.4. $\text{Atoms}(X) \subseteq V_{\mu+1}$.

Proof. If $x \in X$ then $x \subseteq V_\mu$, and $\{2\} \cup (\{0\} \times x) \subseteq V_\mu$ by lemma 2.4. So $\text{atm}(x) \in V_{\mu+1}$. \square

Proposition 3.5. $U_X = \mathcal{Q}_X(U_X)$.

Proof. Lemmas 2.6 and 3.4 imply that $\mathcal{Q}_X(V_{\mu+1}) \subseteq V_{\mu+1}$. So \mathcal{Q}_X is an operator over the powerset of V_μ , and it is clearly monotone. The result follows by the Knaster-Tarski theorem. \square

Proposition 3.6. If $Z \subseteq \mathcal{Q}_X(Z)$ then $Z \subseteq U_X$.

Proof. The result follows by the definition of U_X if we can establish $Z \subseteq V_{\mu+1}$. By lemma 2.7 it suffices to prove $\forall z \in Z \ z \cap V_\alpha \subseteq V_\mu$ for all α . Proceed by transfinite induction on the ordinal α .

Let $z \in Z$. Then $z \in \mathcal{Q}_X(Z) = \text{Atoms}(X) \cup \{1\} \cup (I \rightrightarrows Z)$. The case $z = 1$ is trivial, and if $z \in \text{Atoms}(X)$ then $z \subseteq V_\mu$ by lemma 3.4. So we may assume $z = \tilde{\lambda}_{i \in I} z_i$, with $z_i \in Z$ for all $i \in I$. In this case we have

$$\begin{aligned} (\tilde{\lambda}_{i \in I} z_i) \cap V_\alpha &\subseteq \bigcup_{\beta < \alpha} \tilde{\lambda}_{i \in I} (z_i \cap V_\beta) \\ &\subseteq \bigcup_{\beta < \alpha} \tilde{\lambda}_{i \in I} V_\mu \\ &\subseteq V_\mu \end{aligned}$$

by lemma 2.11, the induction hypothesis for z_i and lemma 2.5. Since $z \cap V_\alpha \subseteq V_\mu$ for all α we have $z \subseteq V_\mu$ for all $z \in Z$. This establishes $Z \subseteq V_{\mu+1}$. \square

Remark 3.7. Using this result, we can check that U_X is nontrivial. Clearly $0 \in U_X$ because $\{0\} = I \rightrightarrows \{0\} \subseteq \mathcal{Q}_X(\{0\})$. We also have inclusions such as $\{0, 1\} \cup (I \rightrightarrows \{0, 1\}) \subseteq U_X$.

3.2. U_X is a Final \mathcal{Q}_X -Coalgebra

Proving that U_X is a final \mathcal{Q}_X -coalgebra requires showing that for every map $f : A \rightarrow \mathcal{Q}_X(A)$ there is a unique map $\pi : A \rightarrow U_X$ such that $\pi = \mathcal{Q}_X(\pi) \circ f$:

$$\begin{array}{ccc} A & \overset{\pi}{\dashrightarrow} & U_X \\ \downarrow f & & \parallel \\ \mathcal{Q}_X(A) & \overset{\mathcal{Q}_X(\pi)}{\dashrightarrow} & \mathcal{Q}_X(U_X) \end{array}$$

For the remainder of this section, let the set A and the map $f : A \rightarrow \mathcal{Q}_X(A)$ be fixed.

Lemma 3.8. There exists $\pi : A \rightarrow U_X$ such that $\pi(a) = \mathcal{Q}_X(\pi)(f(a))$ for all $a \in A$.

Proof. The function π is defined by $\pi(a) \equiv \bigcup_{n < \omega} \pi_n(a)$, where $\{\pi_n\}_{n < \omega}$ is a monotonically increasing series of functions:

$$\begin{aligned} \pi_0(a) &\equiv 0 \\ \pi_{n+1}(a) &\equiv \mathcal{Q}_X(\pi_n)(f(a)) \end{aligned}$$

Suppose $a \in A$, and consider $\pi(a) = \mathcal{Q}_X(\pi)(f(a))$ by cases. If $f(a) = 1$ or $f(a) \in \text{Atoms}(X)$ then the equation reduces to $f(a) = f(a)$. If $f(a) = \tilde{\lambda}_{i \in I} a_i$ then simple continuity reasoning establishes the equation:

$$\begin{aligned} \pi(a) &= \bigcup_{n < \omega} \pi_n(a) = \bigcup_{n < \omega} \pi_{n+1}(a) \\ &= \bigcup_{n < \omega} \mathcal{Q}_X(\pi_n)(f(a)) = \bigcup_{n < \omega} \tilde{\lambda}_{i \in I} \pi_n(a_i) = \tilde{\lambda}_{i \in I} \bigcup_{n < \omega} \pi_n(a_i) \\ &= \tilde{\lambda}_{i \in I} \pi(a_i) = \mathcal{Q}_X(\pi)(\tilde{\lambda}_{i \in I} a_i) = \mathcal{Q}_X(\pi)(f(a)) \end{aligned}$$

To show $\pi : A \rightarrow U_X$, use coinduction (prop. 3.6). Let $Z = \{\pi(a) \mid a \in A\}$ and prove $Z \subseteq \mathcal{Q}_X(Z)$. If $z \in Z$ then $z = \pi(a) = \mathcal{Q}_X(\pi)(f(a))$ for some $a \in A$. If $f(a) = 1$ or $f(a) \in \text{Atoms}(X)$ then $f(a) \in \mathcal{Q}_X(Z)$ and $z = f(a)$. If $f(a) = \tilde{\lambda}_{i \in I} a_i$ then $z = \tilde{\lambda}_{i \in I} \pi(a_i) \in \mathcal{Q}_X(Z)$.

Since U_X is the greatest post-fixedpoint of \mathcal{Q}_X , this establishes $Z \subseteq U_X$. And since Z is the range of π , this establishes $\pi : A \rightarrow U_X$. \square

Lemma 3.9. If $\pi = \mathcal{Q}_X(\pi) \circ f$ and $\pi' = \mathcal{Q}_X(\pi') \circ f$ then $\pi = \pi'$.

Proof. Again using lemma 2.7, apply transfinite induction on the ordinal ξ to prove $\forall_{a \in A} \pi(a) \cap V_\xi \subseteq \pi'(a)$.

Let $a \in A$. If $f(a) = 1$ or $f(a) \in \text{Atoms}(X)$ then $\pi(a) = \pi'(a) = f(a)$. If $f(a) = \tilde{\lambda}_{i \in I} a_i$ then

$$\pi(a) \cap V_\xi = (\tilde{\lambda}_{i \in I} \pi(a_i)) \cap V_\xi \subseteq \bigcup_{\eta < \xi} \tilde{\lambda}_{i \in I} (\pi(a_i) \cap V_\eta) \subseteq \bigcup_{\eta < \xi} \tilde{\lambda}_{i \in I} \pi'(a_i) = \pi'(a)$$

using the hypothesis, lemma 2.11, the induction hypothesis for $\eta < \xi$ and monotonicity of $\tilde{\lambda}$.

Since $\pi(a) \cap V_\xi \subseteq \pi'(a)$ for every ordinal ξ , we have $\pi(a) \subseteq \pi'(a)$. By symmetry we have $\pi'(a) \subseteq \pi(a)$ and therefore $\pi(a) = \pi'(a)$ for all $a \in A$. \square

Theorem 1. U_X is a final \mathcal{Q}_X -coalgebra.

Proof. Immediate by the previous two lemmas. \square

Proposition 3.10. If $f : X \rightarrow Y$ then there is a unique map $h : U_X \rightarrow U_Y$ such that $h = \mathcal{Q}_f(h)$. Calling this map U_f makes the operation U_- a functor.

Proof. The map exists by the universal property of U_Y . Routine calculations show that it preserves identities and composition. \square

When $X = 0$ we may omit the subscript, writing $U = \mathcal{Q}(U)$ instead of $U_0 = \mathcal{Q}_0(U_0)$. It is easy to see that U_- is monotone, and in particular that $U \subseteq U_X$.

Lemma 3.11. Let $0[X]$ be the unique map from the empty set into X . Then $U_{0[X]} : U \rightarrow U_X$ equals the inclusion map ι_{U, U_X} .

Proof. Abbreviate ι_{U, U_X} as ι . We find that $\iota(v) = \mathcal{Q}_{0[X]}(\iota)(v)$ for $v \in U$, for if $v = 1$ then $\iota(1) = 1 = \mathcal{Q}_{0[X]}(\iota)(1)$, and if $v = \tilde{\lambda}_{i \in I} v_i$ then

$$\iota(\tilde{\lambda}_{i \in I} v_i) = \tilde{\lambda}_{i \in I} v_i = \tilde{\lambda}_{i \in I} \iota(v_i) = \mathcal{Q}_{0[X]}(\iota)(\tilde{\lambda}_{i \in I} v_i).$$

The result follows by the uniqueness part of prop. 3.10. \square

4. Solutions of Equations

In his development of set theory with AFA, Aczel (1988) defines systems of set-equations and proves the *solution lemma*: each system has a unique solution. Aczel introduces a class X of variables and a class V_X of sets built up from variables (but not themselves variables). His *substitution lemma* says that any assignment $f : X \rightarrow V$ of sets to variables can be extended to a substitution function $\hat{f} : V_X \rightarrow V$. Aczel uses these lemmas to exhibit a unique morphism for his special final coalgebra theorem.

Aczel proves the solution and substitution lemmas using concrete set theory, but in Rutten and Turi's categorical presentation the proofs are much shorter. A key fact in their development is

that V is (assuming AFA) a final \mathcal{P} -coalgebra. My presentation is similar, replacing V by U , V_X by U_X , \mathcal{P} by \mathcal{Q} and AFA by theorem 1. One improvement over Rutten and Turi (1993) is that U is simply U_0 rather than a separate construction. (Sect. 2.3 discusses the advantages at length.) In this setup, the solution and substitution lemmas nicely generalize to relate two sets of variables. Equations in X and Y can be solved with respect to X , and substitutions can be iterated. Also—a matter of taste—I replace the category of classes by the category of sets.

Note that V_X does not include atoms amongst its elements—they are only allowed in sets—while U_X includes $\text{Atoms}(X)$. This deviation from Aczel will affect many definitions below. The set $\mathcal{Q}(U_X)$ makes a better analogy with V_X : it does not include a copy of the atoms.

4.1. Expressing Maps on $\mathcal{Q}_X(Y)$

Since $\mathcal{Q}_X(Y) = \text{Atoms}(X) \cup \{1\} \cup (I \xrightarrow{\sim} Y)$ and \mathcal{Q} abbreviates \mathcal{Q}_0 , we can write the set $\mathcal{Q}_X(Y)$ as the union of the disjoint sets $\text{Atoms}(X)$ and $\mathcal{Q}(Y)$. Some notation will simplify later calculations.

Definition 4.1. If A and B are sets with B disjoint from $\text{Atoms}(A)$, then

$$A \uplus B \equiv \text{Atoms}(A) \cup B.$$

If moreover $f : A \rightarrow C$ and $g : B \rightarrow C$ are functions, then $\llbracket f, g \rrbracket : A \uplus B \rightarrow C$ is the unique function such that

$$\begin{aligned} \llbracket f, g \rrbracket(\text{atm } x) &= f(x) & (x \in A) \\ \llbracket f, g \rrbracket(y) &= g(y) & (y \in B) \end{aligned}$$

Typically $f : X \rightarrow U_Y$ and $g : \mathcal{Q}(U_X) \rightarrow \mathcal{Q}(U_Y)$. Strictly speaking, the two maps should have the same codomain. Abusing the notation, we can omit the inclusion map $\iota : \mathcal{Q}(U_Y) \rightarrow U_Y$, abbreviating $\llbracket f, \iota \circ g \rrbracket$ as $\llbracket f, g \rrbracket$. Note that $\llbracket f, g \rrbracket : U_X \rightarrow U_Y$ because $X \uplus \mathcal{Q}(U_X) = U_X$. Making ι explicit, a typical calculation is

$$\llbracket j, k \rrbracket \circ \llbracket \text{atm}, \iota \circ g \rrbracket = \llbracket \llbracket j, k \rrbracket \circ \text{atm}, \llbracket j, k \rrbracket \circ \iota \circ g \rrbracket = \llbracket j, k \circ g \rrbracket.$$

The map $\mathcal{Q}_f(g)$ can be written as $\llbracket \text{atm} \circ f, \mathcal{Q}g \rrbracket$, which is sometimes clearer.

4.2. Solution and Substitution Lemmas

Let $f : X \rightarrow U_Y$ be a function. Then the substitution function $\hat{f} : U_X \rightarrow U_Y$ recursively traverses its argument. Given an element of $X \uplus U_X$, it applies f or \hat{f} as appropriate, replacing everything of the form $\text{atm}(x)$ by $f(x)$. We have the case analysis

$$\begin{aligned} \hat{f}(\text{atm}(x)) &= f(x) \\ \hat{f}(1) &= 1 \\ \hat{f}(\tilde{\lambda}_{i \in I} z_i) &= \tilde{\lambda}_{i \in I} \hat{f}(z_i), \end{aligned}$$

which may be put more succinctly as $\hat{f} = \llbracket f, \mathcal{Q}\hat{f} \rrbracket$.

Remark 4.2. In situations where the hat is too short, such as $\widehat{f \circ g}$, the notation $\overline{f \circ g}$ may be used instead.

If X is a set of variables, then a function $\nu : X \rightarrow U_Y \uplus Q(U_X)$ defines a system of equations of the form $x = \nu(x)$ for $x \in X$. Each left-hand side is a variable drawn from X . Each right-hand side is either an expression involving variables from Y or a *guarded* expression involving variables from X . By guarded I mean that the expression must consist of more than just a variable; this restriction excludes degenerate systems of equations such as $\{x = x\}_{x \in X}$, whose solutions are not unique.

A system of equations has a unique solution $f : X \rightarrow U_Y$ that preserves the right-hand sides involving Y while solving for the variables in X . In other words, we require $f(x) = \nu(x)$ if $\nu(x) \in \text{Atoms}(U_Y)$ and $f(x) = Q(\hat{f})(\nu(x))$ otherwise. More concisely, a solution satisfies $f = \llbracket \text{id}_{U_Y}, Q\hat{f} \rrbracket \circ \nu$.

Lemma 4.3. (Solution) Let $\nu : X \rightarrow U_Y \uplus Q(U_X)$ be a function. There exist unique functions $f : X \rightarrow U_Y$ and $\hat{f} : U_X \rightarrow U_Y$ such that $f = \llbracket \text{id}_{U_Y}, Q\hat{f} \rrbracket \circ \nu$ and $\hat{f} = \llbracket f, Q\hat{f} \rrbracket$.

Proof. Let $\iota : Q(U_X) \rightarrow U_Y \uplus Q(U_X)$ be an inclusion and let m be the map

$$\begin{array}{c} U_Y \uplus Q(U_X) = Q_Y(U_Y) \uplus Q(X \uplus Q(U_X)) \\ \downarrow \llbracket Q_Y(\text{atm}), Q(\llbracket \nu, \iota \rrbracket) \rrbracket \\ Q_Y(U_Y \uplus Q(U_X)) \end{array}$$

Now consider the diagram

$$\begin{array}{ccccc} X & \xrightarrow{\nu} & U_Y \uplus Q(U_X) & \overset{\pi}{\dashrightarrow} & U_Y \\ & & \downarrow m & & \parallel \\ & & Q_Y(U_Y \uplus Q(U_X)) & \overset{Q_Y(\pi)}{\dashrightarrow} & Q_Y(U_Y) \end{array}$$

Since $(U_Y \uplus Q(U_X), m)$ is a Q_Y -coalgebra, finality yields a unique coalgebra morphism π into U_Y . The diagram commutes, and we calculate

$$\begin{aligned} \pi &= Q_Y(\pi) \circ m = \llbracket Q_Y(\pi) \circ Q_Y(\text{atm}), Q_Y(\pi) \circ Q(\llbracket \nu, \pi \circ \iota \rrbracket) \rrbracket \\ &= \llbracket Q_Y(\pi \circ \text{atm}), Q(\llbracket \pi \circ \nu, \pi \circ \iota \rrbracket) \rrbracket. \end{aligned}$$

So $\pi \circ \text{atm} = Q_Y(\pi \circ \text{atm}) : U_Y \rightarrow U_Y$, and the uniqueness part of prop. 3.10 yields $\pi \circ \text{atm} = U_{\text{id}_Y} = \text{id}_{U_Y}$. Furthermore, $\pi \circ \iota = Q(\llbracket \pi \circ \nu, \pi \circ \iota \rrbracket)$.

Now put $\hat{f} = \llbracket \pi \circ \nu, \pi \circ \iota \rrbracket$ and $f = \pi \circ \nu$. Then f and \hat{f} satisfy the claimed properties because $\pi = \llbracket \text{id}_{U_Y}, Q\hat{f} \rrbracket$. In particular,

$$\hat{f} = \llbracket f, \llbracket \text{id}_{U_Y}, Q\hat{f} \rrbracket \circ \iota \rrbracket = \llbracket f, Q\hat{f} \rrbracket.$$

As for uniqueness, suppose there are functions $g : X \rightarrow U_Y$ and $\hat{g} : U_X \rightarrow U_Y$ such that $g = \llbracket \text{id}_{U_Y}, Q\hat{g} \rrbracket \circ \nu$ and $\hat{g} = \llbracket g, Q\hat{g} \rrbracket$. Let $\pi' = \llbracket \text{id}_{U_Y}, Q\hat{g} \rrbracket$. Then $g = \pi' \circ \nu$, and π' also

makes the diagram commute:

$$\begin{aligned}
\mathcal{Q}_Y(\pi') \circ m &= \llbracket \mathcal{Q}_Y(\pi' \circ \text{atm}), \mathcal{Q}(\llbracket \pi' \circ \nu, \pi' \circ \iota \rrbracket) \rrbracket \\
&= \llbracket \mathcal{Q}_Y(\text{id}_{U_Y}), \mathcal{Q}(\llbracket g, \mathcal{Q}\hat{g} \rrbracket) \rrbracket \\
&= \llbracket \text{id}_{U_Y}, \mathcal{Q}\hat{g} \rrbracket \\
&= \pi'
\end{aligned}$$

Uniqueness of the final map yields $\pi' = \pi$ and therefore $g = f$ and $\hat{g} = \hat{f}$. \square

The following lemma justifies the \hat{f} notation for substitution by f . The idea is to convert $f : X \rightarrow U_Y$ into a trivial system of equations and then to solve them.

Lemma 4.4. (Substitution) Let $f : X \rightarrow U_Y$ be a function. There exists a unique function $\hat{f} : U_X \rightarrow U_Y$ such that $\hat{f} = \llbracket f, \mathcal{Q}\hat{f} \rrbracket$.

Proof. Let $\nu : X \rightarrow U_Y \uplus \mathcal{Q}(U_X)$ be the map $\text{atm} \circ f$. The solution lemma yields unique maps $g : X \rightarrow U_Y$ and $\hat{g} : U_X \rightarrow U_Y$ such that $g = \llbracket \text{id}_{U_Y}, \mathcal{Q}\hat{g} \rrbracket \circ \nu$ and $\hat{g} = \llbracket g, \mathcal{Q}\hat{g} \rrbracket$. Putting $\hat{f} = \hat{g}$ gives $\hat{f} = \llbracket f, \mathcal{Q}\hat{f} \rrbracket$ because

$$g = \llbracket \text{id}_{U_Y}, \mathcal{Q}\hat{g} \rrbracket \circ \text{atm} \circ f = \text{id}_{U_Y} \circ f = f.$$

As for uniqueness, if $\hat{h} = \llbracket f, \mathcal{Q}\hat{h} \rrbracket$ then $\hat{h} = \llbracket g, \mathcal{Q}\hat{h} \rrbracket$ and so $\hat{h} = \hat{g} = \hat{f}$ by the uniqueness of solutions. \square

Lemma 4.5. (Commutativity) If $f : X \rightarrow U_Y$ and $g : Y \rightarrow U_Z$, then $\widehat{\hat{g} \circ f} = \hat{g} \circ \hat{f}$.

Proof. By uniqueness of substitution, if $h = \llbracket \hat{g} \circ f, \mathcal{Q}h \rrbracket$ then $h = \widehat{\hat{g} \circ f}$. The result follows because

$$\hat{g} \circ \hat{f} = \hat{g} \circ \llbracket f, \mathcal{Q}\hat{f} \rrbracket = \llbracket \hat{g} \circ f, \llbracket g, \mathcal{Q}\hat{g} \rrbracket \circ \mathcal{Q}\hat{f} \rrbracket = \llbracket \hat{g} \circ f, \mathcal{Q}(\hat{g} \circ \hat{f}) \rrbracket.$$

\square

Lemma 4.6. If $f : X \rightarrow Y$ and $g : Y \rightarrow U_Z$, then $\widehat{g \circ f} = \hat{g} \circ U_f$.

Proof. By uniqueness of substitution, if $h = \llbracket g \circ f, \mathcal{Q}h \rrbracket$ then $h = \widehat{g \circ f}$. The result follows because

$$\begin{aligned}
\hat{g} \circ U_f &= \hat{g} \circ \mathcal{Q}_f(U_f) = \hat{g} \circ \llbracket \text{atm} \circ f, \mathcal{Q}(U_f) \rrbracket \\
&= \llbracket \hat{g} \circ \text{atm} \circ f, \llbracket g, \mathcal{Q}\hat{g} \rrbracket \circ \mathcal{Q}(U_f) \rrbracket = \llbracket g \circ f, \mathcal{Q}(\hat{g} \circ U_f) \rrbracket.
\end{aligned}$$

\square

In earlier work (Paulson, 1995a), following previous authors, I defined substitution for a map $f : X \rightarrow U$, with no indeterminates in the codomain. The ability to deal with different sets of variables turns out to be useful. We can recover the original solution and substitution lemmas by applying them with $Y = 0$. The embedding $\sigma_X : U \rightarrow U_X$ becomes the inclusion $U_{0[X]}$ in the present framework.

Lemma 4.7. $\overline{0[U_X]} = U_{0[X]}$.

Proof. The result follows by the uniqueness aspect of prop. 3.10, since

$$\overline{0[U_X]} = \llbracket 0[U_X], \mathcal{Q}(\overline{0[U_X]}) \rrbracket = \mathcal{Q}_{0[U_X]}(\overline{0[U_X]}).$$

\square

Lemma 4.8. (Inclusion) If $f : X \rightarrow U_Y$ then $\hat{f} \circ U_{0[X]} = U_{0[Y]}$, and thus $\hat{f}(v) = v$ for $v \in U$.

Proof. By the previous lemmas, $\hat{f} \circ U_{0[X]} = \overline{f \circ 0[X]} = \overline{0[U_Y]} = U_{0[Y]}$. If $v \in U$ then $\hat{f}(v) = \hat{f}(U_{0[X]}(v)) = U_{0[Y]}(v) = v$ by lemma 3.11. \square

4.3. Special Final Coalgebra Theorem

We shall no longer work in the category **Set** of sets but rather in the full subcategory **Set_U** whose objects are the subsets of U . Recall that U , in turn, depends upon the choice of index set I ; we can make U as large as necessary.

For a suitable functor, our goal is to show that its final coalgebra coincides with its greatest fixed point. Let us only consider functors that preserve inclusion maps. This is a natural restriction since all functors preserve identity maps, and inclusion maps are identity maps when regarded as sets. All such functors have a greatest fixedpoint.

Lemma 4.9. If the functor $F : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ preserves inclusions then there exists an object $J[F] : \mathbf{Set}_U$ such that $J[F]$ is the greatest fixedpoint and greatest post-fixedpoint of F .

Proof. Apply the Knaster-Tarski fixedpoint theorem to the lattice of subsets of U . The functor F is necessarily monotone because it preserves inclusions: if $A \subseteq B$ then $F(\iota_{A,B}) = \iota_{FA,FB}$, giving $FA \subseteq FB$. \square

Definition 4.10. A functor $F : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is *uniform on maps* if it preserves inclusions and for all A such that $A \subseteq U$ there exists a mapping $\phi_A : FA \rightarrow \mathcal{Q}(U_A)$ satisfying $Fh(w) = (\mathcal{Q}\hat{h} \circ \phi_A)(w)$ for all $h : A \rightarrow U$ and $w \in FA$. The mapping ϕ_A is called the U_A translation.

Remark 4.11. The condition above can be abbreviated as $\iota_{F(U),U} \circ Fh = \mathcal{Q}\hat{h} \circ \phi_A$, where $\iota_{F(U),U}$ is the inclusion map from $F(U)$ into U . And since the domain of \hat{h} includes that of $\mathcal{Q}\hat{h}$, we have

$$\hat{h}(\phi_A(w)) = \llbracket h, \mathcal{Q}\hat{h} \rrbracket(\phi_A(w)) = \mathcal{Q}\hat{h}(\phi_A(w)) = Fh(w).$$

The main theorem applies to functors that are uniform on maps. This notion is due to Aczel (1988), but the presentation owes much to Rutten and Turi (1993).

Theorem 2. (Special Final Coalgebra) If the functor $F : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is uniform on maps, then $J[F]$ is a final F -coalgebra.

Proof. Let (A, f) be an F -coalgebra. We must exhibit a unique map $h : A \rightarrow J[F]$ such that $h = Fh \circ f$:

$$\begin{array}{ccc} A & \overset{h}{\dashrightarrow} & J[F] \\ \downarrow f & & \parallel \\ FA & \overset{Fh}{\dashrightarrow} & F(J[F]) \end{array}$$

Since F is uniform on maps, there is a U_A -translation $\phi_A : FA \rightarrow \mathcal{Q}(U_A)$. Let $\iota : \mathcal{Q}(U_A) \rightarrow U \uplus \mathcal{Q}(U_A)$ be an embedding and apply the solution lemma with $\nu = \iota \circ \phi_A \circ f$. We obtain a unique map $h : A \rightarrow U$ such that $h = \llbracket \text{id}_U, \mathcal{Q}\hat{h} \rrbracket \circ \iota \circ \phi_A \circ f = \mathcal{Q}\hat{h} \circ \phi_A \circ f$. So $h(a) = (\mathcal{Q}\hat{h} \circ \phi_A)(f(a)) = Fh(f(a))$ for $a \in A$.

Regarding the maps as set-theoretic functions, a standard coinduction argument proves $h \in A \rightarrow J[F]$. Writing $h \text{ `` } A$ for the image of A under h , we have

$$h \text{ `` } A = (Fh \circ f) \text{ `` } A = Fh \text{ `` } (f \text{ `` } A) \subseteq Fh \text{ `` } FA \subseteq F(h \text{ `` } A)$$

since $h \in A \rightarrow h \text{ `` } A$ and $Fh \in FA \rightarrow F(h \text{ `` } A)$.

The range of h is thus a post-fixedpoint of F and is included in the greatest post-fixedpoint, namely $J[F]$. \square

5. Existence of Functors Uniform on Maps

If F is uniform on maps then, in essence, its effect upon a map $h : A \rightarrow U$ can be expressed as the substitution of h over a pattern derived from the argument; if $w \in FA$ then $Fh(w) = \mathcal{Q}\hat{h}(\phi_A(w))$. Most natural functors are uniform on maps, but there is one glaring exception. Let us examine some typical cases, starting with a trivial one.

This section illustrates the advantages of constructing U_X to include atoms, or (assuming AFA) having urelements. If instead we used $V_X = \mathcal{P}(X + V_X)$, then the failure of $V_X \times V_X \subseteq V_X$ and $V_X + V_X \subseteq V_X$ would complicate the translations and the proofs.

5.1. The Constant Functor

If $C \subseteq U$ then let K_C be the constant functor such that $K_C(A) = C$ for all $A : \mathbf{Set}_U$ and such that $K_C(f) = \text{id}_C$ for all maps $f : A \rightarrow A'$.

Proposition 5.1. If $C : \mathbf{Set}_U$ then the constant functor $K_C : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is uniform on maps.

Proof. Let $A \subseteq U$. Now $C \subseteq U = \mathcal{Q}U \subseteq \mathcal{Q}(U_A)$, so we can define $\phi_A : C \rightarrow \mathcal{Q}(U_A)$ by $\phi_A(c) = c$ for $c \in C$. Now $(\mathcal{Q}\hat{h} \circ \phi_A)(c) = \mathcal{Q}\hat{h}(c) = c = K_C(h)(c)$ for all $c \in C$ by lemma 4.8. \square

5.2. Binary Product

The set U satisfies the inclusion $U \tilde{\times} U \subseteq U$. So it is easy to see that $\tilde{\times} : \mathbf{Set}_U \times \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is a functor when extended to maps in the standard way. If $f : A \rightarrow A'$ and $g : B \rightarrow B'$ are maps then $f \tilde{\times} g : A \tilde{\times} B \rightarrow A' \tilde{\times} B'$ is the map that takes $\langle a; b \rangle$ to $\langle f(a); g(b) \rangle$.

Proposition 5.2. If $F, G : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ are uniform on maps, then the functor $F(-) \tilde{\times} G(-) : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is uniform on maps.

Proof. Let A be a set such that $A : \mathbf{Set}_U$, or equivalently $A \subseteq U$. Clearly we have $FA \tilde{\times} GA : \mathbf{Set}_U$. Since F and G are uniform on maps there exist U_A translations

$$\begin{aligned} \phi_A : FA &\rightarrow \mathcal{Q}(U_A) \text{ such that } Fh(u) = (\mathcal{Q}\hat{h} \circ \phi_A)(u) \quad \text{and} \\ \psi_A : GA &\rightarrow \mathcal{Q}(U_A) \text{ such that } Gh(v) = (\mathcal{Q}\hat{h} \circ \psi_A)(v) \end{aligned}$$

for all $u \in FA, v \in GA$ and $h : A \rightarrow U$.

To define the U_A translation for $F(-) \tilde{\times} G(-)$, put $\theta_A = \phi_A \tilde{\times} \psi_A$. Thus $\theta_A(\langle b; c \rangle) =$

$\langle \phi_A(b); \psi_A(b) \rangle$ for all $b \in FA$ and $c \in GA$. Now for $\langle u; v \rangle \in FA \tilde{\times} FB$,

$$\begin{aligned} (F(-) \tilde{\times} G(-))(h)(\langle u; v \rangle) &= \langle Fh(u); Gh(v) \rangle \\ &= \langle \hat{h}(\phi_A(u)); \hat{h}(\psi_A(v)) \rangle \\ &= \mathcal{Q}\hat{h}(\langle \phi_A(u); \psi_A(v) \rangle) \\ &= \mathcal{Q}\hat{h}(\theta_A(\langle u; v \rangle)) \end{aligned}$$

by remark 4.11. So θ_A is the desired U_A translation. \square

5.3. Binary Sum

Recall that $\tilde{+}$ is the variant form of disjoint sum, defined by $A \tilde{+} B \equiv (\{1\} \tilde{\times} A) \cup (\{1; 1\} \tilde{\times} B)$. We have $U \tilde{+} U \subseteq U$ because U is closed under $\tilde{\times}$ and contains 1 and $\langle 1; 1 \rangle$ as elements.

Variant sum is a coproduct in both **Set** and **Set_U**. The injections $\tilde{\text{Inl}} : A \rightarrow A \tilde{+} B$ and $\tilde{\text{Inr}} : B \rightarrow A \tilde{+} B$ are defined in the obvious way. For every pair of maps $f : A \rightarrow C$ and $g : B \rightarrow C$, there exists a unique map $[f, g] : A \tilde{+} B \rightarrow C$ such that $[f, g] \circ \tilde{\text{Inl}} = f$ and $[f, g] \circ \tilde{\text{Inr}} = g$. Another useful law is $h \circ [f, g] = [h \circ f, h \circ g]$.

To make $\tilde{+}$ into a functor, we must define its action on maps. If $j : A \rightarrow A'$ and $k : B \rightarrow B'$ then $j \tilde{+} k : A \tilde{+} A' \rightarrow B \tilde{+} B'$ is defined (as usual) by $j \tilde{+} k \equiv [\tilde{\text{Inl}} \circ j, \tilde{\text{Inr}} \circ k]$.

Lemma 5.3. If $f : A \rightarrow U$ and $z \in U_A$ then $\mathcal{Q}\hat{h}(\tilde{\text{Inl}}(z)) = \tilde{\text{Inl}}(\hat{h}(z))$ and $\mathcal{Q}\hat{h}(\tilde{\text{Inr}}(z)) = \tilde{\text{Inr}}(\hat{h}(z))$.

Proof. Calculate

$$\mathcal{Q}\hat{h}(\tilde{\text{Inl}}(z)) = \mathcal{Q}\hat{h}(\langle 1; z \rangle) = \langle \hat{h}1; \hat{h}(z) \rangle = \langle 1; \hat{h}(z) \rangle = \tilde{\text{Inl}}(\hat{h}(z)).$$

The treatment of $\tilde{\text{Inr}}(z)$ is similar, since $\hat{h}(\langle 1; 1 \rangle) = \mathcal{Q}\hat{h}(\langle 1; 1 \rangle) = \langle \hat{h}1; \hat{h}1 \rangle = \langle 1; 1 \rangle$. \square

The tags of the disjoint sum are arbitrary distinct sets, usually 0 and 1. However, 0 has complicated properties in our framework because $0 = \tilde{\lambda}_{i \in I} 0$. Using $\langle 1; 1 \rangle$ simplifies the proof above. In fact, 0 would work too. From $\hat{h}(0) = \mathcal{Q}\hat{h}(0) = \tilde{\lambda}_{i \in I} \hat{h}(0)$ we get $\hat{h}(0) = 0$, but making this argument rigorous requires establishing a coinduction principle for equations. That would be a distraction.

Proposition 5.4. If $F, G : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ are uniform on maps, then the functor $F(-) \tilde{+} G(-) : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is uniform on maps.

Proof. Let $A \subseteq U$. Then $FA \tilde{+} GA : \mathbf{Set}_U$ and for all $h : A \rightarrow U$ there exist U_A translations ϕ_A and ψ_A as in the proof of prop. 5.2. Let $\theta_A = \phi_A \tilde{+} \psi_A$ and $w \in FA \tilde{+} GA$. If $w = \tilde{\text{Inl}}(u)$ for $u \in FA$, then

$$\begin{aligned} (F(-) \tilde{+} G(-))(h)(\tilde{\text{Inl}}(u)) &= \tilde{\text{Inl}}(Fh(u)) \\ &= \tilde{\text{Inl}}(\hat{h}(\phi_A(u))) \\ &= \mathcal{Q}\hat{h}(\tilde{\text{Inl}}(\phi_A(u))) \\ &= \mathcal{Q}\hat{h}(\theta_A(\tilde{\text{Inl}}(u))) \end{aligned}$$

by lemma 5.3 and remark 4.11. So θ_A is the desired translation. The case $w = \tilde{\text{Inr}}(v)$ follows by symmetry. \square

5.4. Sum of a Family of Sets

Let $\{B_x\}_{x \in C}$ be a C -indexed family of sets. If $C \subseteq U$ and $B_x \subseteq U$ for all $x \in C$ then we have $\tilde{\sum}_{x \in C} B_x \subseteq U$. Note that $\tilde{\sum}_{x \in C} B_x$ is the usual generalization of $C \tilde{\times} B$ to allow B to depend upon $x \in C$; the two functors have a similar effect upon maps. But $\tilde{\sum}$ is less general than $\tilde{\times}$ in one key respect: the index set is not given by a functor but is constant. Neither $\tilde{\sum}$ nor $\tilde{\times}$ supersedes the other.

Proposition 5.5. If $C : \mathbf{Set}_U$ and if $\{F_x : \mathbf{Set}_U \rightarrow \mathbf{Set}_U\}_{x \in C}$ is a C -indexed family of functors that are uniform on maps, then the functor

$$\tilde{\sum}_{x \in C} F_x(-) : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$$

is uniform on maps.

Proof. Let $A \subseteq U$. For each $x \in C$ there exists a U_A translation $\phi_{x,A} : F_x(A) \rightarrow U_A$ such that $F_x(h)(y) = (\hat{h} \circ \phi_{x,A})(y)$ for all $h : A \rightarrow U$ and $y \in F_x(A)$. The U_A translation for $\tilde{\sum}_{x \in C} F_x(-)$, called θ_A , is defined by

$$\theta_A(\langle x; y \rangle) = \langle x; \phi_{x,A}(y) \rangle$$

for all $x \in C$ and $y \in F_x(A)$. Now, we have

$$\begin{aligned} \left(\tilde{\sum}_{x \in C} F_x(-) \right)(h)(\langle x; y \rangle) &= \langle x; F_x(h)(y) \rangle \\ &= \langle \hat{h}(x); \hat{h}(\phi_{x,A}(y)) \rangle = \mathcal{Q}\hat{h}(\theta_A(\langle x; y \rangle)) \end{aligned}$$

by lemma 4.8. □

5.5. Product of a Family of Sets

Again let $\{B_x\}_{x \in C}$ be a C -indexed family of sets. If $C \subseteq I$ (not $C \subseteq U$ as above) and $B_x \subseteq U$ for all $x \in C$ then $\tilde{\prod}_{x \in C} B_x \subseteq I \tilde{\rightarrow} U \subseteq U$.

Thus $\tilde{\prod} : \mathbf{Set}_U^I \rightarrow \mathbf{Set}_U$ is a functor whose effect on maps is as follows. If $\{f_x : B_x \rightarrow B'_x\}_{x \in C}$ is a C -indexed family of maps then

$$\tilde{\prod}_{x \in C} f_x : \tilde{\prod}_{x \in C} B_x \rightarrow \tilde{\prod}_{x \in C} B'_x$$

is the usual pointwise map that takes $\tilde{\lambda}_{x \in C} b_x$ to $\tilde{\lambda}_{x \in C} f_x(b_x)$.

Proposition 5.6. If $C \subseteq I$ and $\{F_x : \mathbf{Set}_U \rightarrow \mathbf{Set}_U\}_{x \in C}$ is a C -indexed family of functors that are uniform on maps, then the functor

$$\tilde{\prod}_{x \in C} F_x(-) : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$$

is uniform on maps.

Proof. Let $A \subseteq U$. For each $x \in C$ there exists a U_A translation $\phi_{x,A}$ as in the proof of prop. 5.5.

Let $\theta_A = \prod_{x \in C} \phi_{x,A}$. If $w \in (\prod_{x \in C} F_x(-))(A)$ then $w = \tilde{\lambda}_{x \in C} w_x$, and

$$\begin{aligned} (\prod_{x \in C} F_x(-))(h)(\tilde{\lambda}_{x \in C} w_x) &= \tilde{\lambda}_{x \in C} F_x(h)(w_x) \\ &= \tilde{\lambda}_{x \in C} (\hat{h} \circ \phi_{x,A})(w_x) \\ &= \mathcal{Q}\hat{h}(\tilde{\lambda}_{x \in C} \phi_{x,A}(w_x)) \\ &= (\mathcal{Q}\hat{h} \circ \theta_A)(\tilde{\lambda}_{x \in C} w_x) \end{aligned}$$

and θ_A is the desired U_A translation. \square

5.6. Composition of Functors

That the composition $F \circ G$ should preserve uniformity on maps seems obvious, but the proof requires the notion of iterated substitution of lemma 4.4.

Proposition 5.7. If $F, G : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ are uniform on maps, then so is the functor $F \circ G : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$.

Proof. Let $A \subseteq U$. Since F and G are uniform on maps, there exist U_A and U_{GA} translations

$$\begin{aligned} \psi_A : GA &\rightarrow \mathcal{Q}(U_A) \text{ such that } \iota_{GU,U} \circ Gh = \mathcal{Q}\hat{h} \circ \psi_A \quad \text{and} \\ \phi_{GA} : FGA &\rightarrow \mathcal{Q}(U_{GA}) \text{ such that } \iota_{FU,U} \circ Fj = \mathcal{Q}\hat{j} \circ \phi_{GA} \end{aligned}$$

for $h : A \rightarrow U$ and $j : GA \rightarrow U$.

Let $\iota : \mathcal{Q}(U_A) \rightarrow U_A$ be an inclusion map and put $\theta_A = \mathcal{Q}(\overline{\iota \circ \psi_A}) \circ \phi_{GA}$. If $h : A \rightarrow U$ and $u \in FGA$ then

$$\begin{aligned} F(Gh)(u) &= F(\mathcal{Q}\hat{h} \circ \psi_A)(u) \\ &= F(\hat{h} \circ \iota \circ \psi_A)(u) \\ &= (\mathcal{Q}(\overline{\hat{h} \circ \iota \circ \psi_A}) \circ \phi_{GA})(u) \\ &= (\mathcal{Q}(\hat{h} \circ \overline{\iota \circ \psi_A}) \circ \phi_{GA})(u) \\ &= (\mathcal{Q}\hat{h} \circ \mathcal{Q}(\overline{\iota \circ \psi_A}) \circ \phi_{GA})(u) \\ &= (\mathcal{Q}\hat{h} \circ \theta_A)(u) \end{aligned}$$

by commutativity of substitution (lemma 4.5). The first equality, in which Gh is replaced by $\mathcal{Q}\hat{h} \circ \psi_A$, holds because F preserves inclusions. \square

5.7. The Identity Functor

These results suggest that any functor that operates on constructions in a pointwise fashion is uniform on maps. But there is one glaring exception.

Proposition 5.8. The identity functor $\text{Id} : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is not uniform on maps.

Proof. Suppose $\text{Id} : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is uniform on maps. Then if $A \subseteq U$ then there is a mapping $\phi_A : A \rightarrow \mathcal{Q}(U_A)$ such that $h(w) = \mathcal{Q}\hat{h}(\phi_A(w))$ for all $h : A \rightarrow U$ and $w \in A$.

Let $A = \{1\}$ and define $h_1, h_2 : \{1\} \rightarrow U$ by $h_1(1) = 1$ and $h_2(1) = \langle 1; 1 \rangle$. Then $1 =$

$\mathcal{Q}(\widehat{h}_1)(\phi_A(1))$, so $\phi_A(1) = 1$ by the definition of \mathcal{Q} . Also $\langle 1; 1 \rangle = \mathcal{Q}(\widehat{h}_2)(\phi_A(1))$, which implies $\phi_A(1) = \langle a; b \rangle$ for some $a, b \in U_A$. But then $1 = \langle a; b \rangle$, which is absurd. \square

An alternative proof uses the special final coalgebra theorem. If Id is uniform on maps then $J[\text{Id}]$ is a final Id -coalgebra. But a final Id -coalgebra must be a singleton set, while $J[\text{Id}] = U$ and U contains 1 and $\langle 1; 1 \rangle$ as elements.

This circumstance is awkward. The natural way of constructing suitable functors is to combine constant and identity functors by products, sums, etc. Since the identity functor is not uniform on maps, this approach fails. Various similar functors are uniform on maps, such as $-\tilde{\times} K_{\{0\}}$ and $-\tilde{\times} -$; both have the singleton set $\{0\}$ as their greatest fixedpoint. One can prove variants of the lemmas above, for example that if F is uniform on maps then so is $F(-) \tilde{\times} -$. Assuming AFA does not help; the identity functor is not uniform on maps in Aczel's system either.

6. Final Coalgebras with Parameters

Section 1 discussed the set S of streams over A , which satisfies $S = A \times S$. But ‘streams over A ’ should be a construction taking A as a parameter. Can we define it as a functor that can itself be used in further constructions?

Suppose F is a bifunctor. If A is an object then $F(-, A)$ is a functor, which we abbreviate to F_A . If F_A has a final coalgebra $J[F_A]$ for every A , then the map $A \mapsto J[F_A]$ determines a functor. The idea is to show that this functor is uniform on maps and to express other functors in terms of it. For example, the functor of streams over A , $\text{stream}(A)$, is uniform on maps. It can express the functor of ω -branching trees as the final coalgebra of the bifunctor $F(A', A) = A \times \text{stream}(A')$, etc.

Our existing machinery already suffices to handle mutually recursive coinductive definitions, finding greatest fixedpoints in the product category $\mathbf{Set}_U \times \mathbf{Set}_U$. The idea is to generalize the special final coalgebra theorem, applying the solution lemma to a set of indeterminates of the form $A_1 \tilde{+} A_2$. But it is more general to handle definitions that have parameters. This topic appears to be little discussed in the final coalgebra literature, but see Hensel and Jacobs (1997), who work in total categories of fibrations. The approach outlined below is simple and applies (making the obvious changes) to approaches based on AFA.

Definition 6.1. A bifunctor $F : \mathbf{Set}_U \times \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is *uniform on maps* if it preserves inclusions and for all subsets A, B of U there exists a mapping $\phi_{A,B} : F(A, B) \rightarrow U_{A \tilde{+} B}$ such that $F(f, g)(w) = (\overline{[f, g]} \circ \phi_{A \tilde{+} B})(w)$ for all $f : A \rightarrow U, g : B \rightarrow U$ and $w \in F(A, B)$.

In this section, A and B range over subsets of U . If the bifunctor F is uniform on maps then so are the functors $F(-, B)$ and $F(A, -)$ for objects A and B . To prove this, we need a few more results.

Lemma 6.2. For every map $\phi : B \rightarrow \mathcal{Q}(U_{A \tilde{+} B})$ there exists a unique map $\text{Outl}[\phi] : U_{A \tilde{+} B} \rightarrow U_A$ such that

$$\text{Outl}[\phi] = \overline{[\text{atm}, \mathcal{Q}(\text{Outl}[\phi]) \circ \phi]}.$$

Proof. Let m be the map

$$\begin{array}{ccc}
 U_{A\tilde{+}B} & \xlongequal{\quad} & (A \tilde{+} B) \uplus \mathcal{Q}(U_{A\tilde{+}B}) \\
 & & \downarrow \\
 & & \llbracket [\text{atm}, \phi], \mathcal{Q}(\text{id}_{U_{A\tilde{+}B}}) \rrbracket \\
 & & A \uplus \mathcal{Q}(U_{A\tilde{+}B})
 \end{array}$$

Now consider the diagram

$$\begin{array}{ccc}
 U_{A\tilde{+}B} & \overset{\pi}{\dashrightarrow} & U_A \\
 m \downarrow & & \parallel \\
 \mathcal{Q}_A(U_{A\tilde{+}B}) & \overset{\mathcal{Q}_A(\pi)}{\dashrightarrow} & \mathcal{Q}_A(U_A)
 \end{array}$$

Since $(U_{A\tilde{+}B}, m)$ is a \mathcal{Q}_A -coalgebra, there is a unique map π into the final coalgebra U_A making the diagram commute. Now

$$\begin{aligned}
 \pi &= \mathcal{Q}_A(\pi) \circ m \\
 &= \llbracket \mathcal{Q}_A(\pi) \circ [\text{atm}, \phi], \mathcal{Q}_A(\pi) \circ \mathcal{Q}(\text{id}_{U_{A\tilde{+}B}}) \rrbracket \\
 &= \llbracket [\text{atm}, \mathcal{Q}\pi \circ \phi], \mathcal{Q}\pi \rrbracket.
 \end{aligned}$$

By the substitution lemma (4.4), the desired map $\text{Outl}[\phi]$ is π . (Note: $\mathcal{Q}_A(\pi)$ becomes $\mathcal{Q}\pi$ after composition with the implicit inclusion map for $\mathcal{Q}(U_{A\tilde{+}B}) \subseteq A \uplus \mathcal{Q}(U_{A\tilde{+}B}) = \mathcal{Q}_A(U_{A\tilde{+}B})$.) \square

Lemma 6.3. For every map $\phi : A \rightarrow \mathcal{Q}(U_{A\tilde{+}B})$ there exists a unique map $\text{Outr}[\phi] : U_{A\tilde{+}B} \rightarrow U_B$ such that

$$\text{Outr}[\phi] = \overline{[\mathcal{Q}(\text{Outr}[\phi]) \circ \phi, \text{atm}]}.$$

Proof. As above, by symmetry. \square

Lemma 6.4. If $f : A \rightarrow U$ then $\hat{f} \circ \text{Outl}[\iota_B, \mathcal{Q}(U_{A\tilde{+}B})] = \overline{[f, \iota_B, U]}$.

Proof. Abbreviate $\iota_B, \mathcal{Q}(U_{A\tilde{+}B})$ as ι . Since $B \subseteq U = \mathcal{Q}U \subseteq \mathcal{Q}(U_{A\tilde{+}B})$, we have $\iota = \mathcal{Q}(U_{0[A\tilde{+}B]}) \circ \iota_{B,U}$. Lemmas 6.2 and 4.8 give

$$\text{Outl}[\iota] \circ U_{0[A\tilde{+}B]} = \overline{[\text{atm}, \mathcal{Q}(\text{Outl}[\iota]) \circ \iota]} \circ U_{0[A\tilde{+}B]} = U_{0[A]}. \quad (1)$$

Apply (1) and lemma 4.8 in a preliminary derivation:

$$\begin{aligned}
 \hat{f} \circ [\text{atm}, \mathcal{Q}(\text{Outl}[\iota]) \circ \iota] &= [\hat{f} \circ \text{atm}, \hat{f} \circ \mathcal{Q}(\text{Outl}[\iota]) \circ \mathcal{Q}(U_{0[A\tilde{+}B]}) \circ \iota_{B,U}] \\
 &= [f, \llbracket f, \mathcal{Q}\hat{f} \rrbracket \circ \mathcal{Q}(\text{Outl}[\iota] \circ U_{0[A\tilde{+}B]}) \circ \iota_{B,U}] \\
 &= [f, \mathcal{Q}(\hat{f} \circ U_{0[A]}) \circ \iota_{B,U}] \\
 &= [f, \mathcal{Q}(\text{id}_U) \circ \iota_{B,U}] \\
 &= [f, \iota_{B,U}]
 \end{aligned}$$

And so, applying lemmas 4.4 and 4.5, we obtain

$$\begin{aligned}\hat{f} \circ \text{Outl}[\iota] &= \hat{f} \circ \overline{[\text{atm}, \mathcal{Q}(\text{Outl}[\iota]) \circ \iota]} \\ &= \overline{\hat{f} \circ [\text{atm}, \mathcal{Q}(\text{Outl}[\iota]) \circ \iota]} \\ &= \overline{[f, \iota_{B,U}]}. \end{aligned}$$

□

Lemma 6.5. If $g : B \rightarrow U$ then $\hat{g} \circ \text{Outl}[\iota_{A, \mathcal{Q}(U_{A \dagger B})}] = \overline{[\iota_{A,U}, g]}$.

Proof. By symmetry in the previous proof. □

Proposition 6.6. If the bifunctor $F : \mathbf{Set}_U \times \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is uniform on maps, then so are the functors $F(-, B)$ and $F(A, -)$ for $A, B : \mathbf{Set}_U$.

Proof. Since F is uniform on maps, it has a $U_{A \dagger B}$ translation $\phi_{A,B} : F(A, B) \rightarrow \mathcal{Q}(U_{A \dagger B})$.

Let $B \subseteq U$ be fixed and consider the functor $F(-, B)$. Then, for $A \subseteq U$, we shall see that the U_A translation for $F(-, B)$ is $\mathcal{Q}(\text{Outl}[\iota_{B, \mathcal{Q}(U_{A \dagger B})}]) \circ \phi_{A,B}$. For $h : A \rightarrow U$ and $w \in F(A, B)$, we have by lemma 6.4

$$\begin{aligned}(\mathcal{Q}\hat{h} \circ \mathcal{Q}(\text{Outl}[\iota_{B, \mathcal{Q}(U_{A \dagger B})}])) \circ \phi_{A,B}(w) &= (\mathcal{Q}(\overline{[\hat{h}, \iota_{B,U}]}) \circ \mathcal{Q}(\text{Outl}[\iota_{B, \mathcal{Q}(U_{A \dagger B})}])) \circ \phi_{A,B}(w) \\ &= (\mathcal{Q}(\overline{[h, \iota_{B,U}]}) \circ \phi_{A,B})(w) \\ &= F(h, \iota_{B,U})(w) \\ &= (F(\text{id}_U, \iota_{B,U}) \circ F(h, \text{id}_B))(w) \\ &= F(h, \text{id}_B)(w), \end{aligned}$$

since F preserves inclusions.

If $A \subseteq U$ is fixed, the U_A translation for $F(A, -)$ is $\mathcal{Q}(\text{Outl}[\iota_{A, \mathcal{Q}(U_{A \dagger B})}]) \circ \phi_{A,B}$, by symmetry. □

I do not know whether the converse of this proposition holds. This question might be considered in future research.

Theorem 3. Let $F : \mathbf{Set}_U \times \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ be a bifunctor that is uniform on maps. For $A : \mathbf{Set}_U$, let F_A abbreviate the functor $F(-, A)$. Then $J[F_A]$ is a final F_A -coalgebra, and the map $A \mapsto J[F_A]$ determines a functor that is uniform on maps.

Proof. If $A \subseteq U$ then the functor F_A is uniform on maps by prop. 6.6. By the special final coalgebra theorem, $J[F_A]$ is a final F_A -coalgebra. The fixedpoint property yields $J[F_A] = F(J[F_A], A)$.

As is well known, the map $A \mapsto J[F_A]$ determines a functor. Given $h : A \rightarrow B$, finality of $J[F_B]$ yields a unique map $J[F_h]$ such that $J[F_h] = F(J[F_h], h)$. By uniqueness, it is easy to check that $J[\text{id}_A] = \text{id}_{J[F_A]}$ and $J[f \circ g] = J[f] \circ J[g]$.

The functor also preserves inclusions. If $A \subseteq B$ then $J[F_A] \subseteq J[F_B]$ by monotonicity of the greatest fixedpoint operator. Since F preserves inclusions, $F(\iota_{J[F_A], J[F_B]}, \iota_{A,B}) = \iota_{J[F_A], J[F_B]}$. By uniqueness, $J[F_{\iota_{A,B}}] = \iota_{J[F_A], J[F_B]}$.

Let $A \subseteq U$ be given. To show that the functor $J[F_-]$ is uniform on maps, it remains to exhibit a U_A translation θ_A such that $\mathcal{Q}\hat{h} \circ \theta_A = \iota_{F_U, U} \circ F_h$ for $h : A \rightarrow U$. Abbreviate $J[F_A]$ as J .

Let $\phi_{J\tilde{\dagger}A} : F(J, A) \rightarrow \mathcal{Q}(U_{J\tilde{\dagger}A})$ be the translation for the bifunctor F ; since $J = F(J, A)$, we have $\phi_{J\tilde{\dagger}A} : J \rightarrow \mathcal{Q}(U_{J\tilde{\dagger}A})$.

The required translation is $\theta_A = \mathcal{Q}(\text{Outr}[\phi_{J\tilde{\dagger}A}]) \circ \phi_{J\tilde{\dagger}A}$. If $h : A \rightarrow U$ then by lemmas 6.3 and 4.5,

$$\begin{aligned} \mathcal{Q}\hat{h} \circ \theta_A &= \mathcal{Q}(\hat{h} \circ \text{Outr}[\phi_{J\tilde{\dagger}A}]) \circ \phi_{J\tilde{\dagger}A} \\ &= \mathcal{Q}(\hat{h} \circ [\overline{\mathcal{Q}(\text{Outr}[\phi_{J\tilde{\dagger}A}]) \circ \phi_{J\tilde{\dagger}A}, \text{atm}}]) \circ \phi_{J\tilde{\dagger}A} \\ &= \mathcal{Q}(\hat{h} \circ [\overline{\mathcal{Q}(\text{Outr}[\phi_{J\tilde{\dagger}A}]) \circ \phi_{J\tilde{\dagger}A}, \text{atm}}]) \circ \phi_{J\tilde{\dagger}A} \\ &= \mathcal{Q}([\overline{\mathcal{Q}(\hat{h} \circ \text{Outr}[\phi_{J\tilde{\dagger}A}]) \circ \phi_{J\tilde{\dagger}A}, h}]) \circ \phi_{J\tilde{\dagger}A} \\ &= \iota_{F(U,U),U} \circ F(\mathcal{Q}(\hat{h} \circ \text{Outr}[\phi_{J\tilde{\dagger}A}]) \circ \phi_{J\tilde{\dagger}A}, h) \\ &= \iota_{F(U,U),U} \circ F(\mathcal{Q}\hat{h} \circ \theta_A, h) \end{aligned}$$

because F is uniform on maps.

It remains to eliminate the inclusion map. Considering $\mathcal{Q}\hat{h} \circ \theta_A$ as a set theoretic function, its range R satisfies $R = F(R, U)$, but the greatest solution to that equation is $J[F_U]$. So $\mathcal{Q}\hat{h} \circ \theta_A = \iota_{J[F_U],U} \circ j$ for some $j : J \rightarrow J[F_U]$. Since F preserves inclusions, we find

$$\iota_{J[F_U],U} \circ j = \mathcal{Q}\hat{h} \circ \theta_A = \iota_{F(U,U),U} \circ F(\iota_{J[F_U],U} \circ j, h) = \iota_{J[F_U],U} \circ F(j, h)$$

and so $j = F(j, h)$. By uniqueness, $j = J[F_h]$. Summarizing, we have

$$\mathcal{Q}\hat{h} \circ \theta_A = \mathcal{Q}\hat{h} \circ \theta_A = \iota_{J[F_U],U} \circ J[F_h],$$

and θ_A is the required U_A translation. \square

How do we create bifunctors that are uniform on maps? It would not do to rehearse the proofs of Sect. 5, but they clearly apply with obvious changes, replacing $h : A \rightarrow U$ by $[f, g] : A\tilde{\dagger}B \rightarrow U$. Ordinary functors give us material to start with.

Proposition 6.7. If $G : \mathbf{Set}_U \rightarrow \mathbf{Set}_U$ is uniform on maps, then so are the degenerate bifunctors F and F' defined by $F(A, B) = GA$ and $F'(B, A) = GA$.

Proof. Let A and B be objects of \mathbf{Set}_U , and let $\phi_A : GA \rightarrow U_A$ be the U_A translation for G . We shall see that the $U_{A\tilde{\dagger}B}$ translation for F is $\mathcal{Q}(U_{\tilde{\text{Inl}}}) \circ \phi_A$.

Suppose $h : A \rightarrow U$, $j : B \rightarrow U$ and $w \in F(A, B)$. Of course $w \in GA$ and, by lemma 4.6,

$$\begin{aligned} (\mathcal{Q}(\overline{[h, j]}) \circ \mathcal{Q}(U_{\tilde{\text{Inl}}}) \circ \phi_A)(w) &= (\mathcal{Q}(\overline{[h, j]} \circ U_{\tilde{\text{Inl}}}) \circ \phi_A)(w) \\ &= (\mathcal{Q}([h, j] \circ \tilde{\text{Inl}}) \circ \phi_A)(w) = (\mathcal{Q}\hat{h} \circ \phi_A)(w) = Gh(w) = F(h, j)(w). \end{aligned}$$

The translation for F' is $\mathcal{Q}(U_{\tilde{\text{Inr}}}) \circ \phi_A$, and the proof follows by symmetry. \square

If another example is needed, the bifunctor $\tilde{\times}$ is uniform on maps with translation $\phi : A\tilde{\times}B \rightarrow \mathcal{Q}(U_{A\tilde{\dagger}B})$ defined by

$$\phi(\langle x; y \rangle) = \langle \text{atm}(\tilde{\text{Inl}}(x)); \text{atm}(\tilde{\text{Inr}}(y)) \rangle.$$

It seems clear that uniformity on maps could be defined for functors in $\mathbf{Set}_U^n \rightarrow \mathbf{Set}_U$, generalizing the proofs of Sect. 5 to an arbitrary positive integer n .

7. Applications to Machine Proof

The context for this work is my mechanization of ZF set theory, using the theorem prover Isabelle (Paulson, 1993). Proof tools should allow users to define sets inductively. Adding induction principles to the formalism is popular (Paulin-Mohring, 1993), but is not suitable for ZF set theory, where strong induction principles can be derived from the axioms. I have put much effort into supporting inductive definitions in Isabelle/ZF, basing the representation on least fixedpoints (Paulson, 1995b).

Coinductive definitions should also be supported. The simplest approach is to base the representation on greatest fixedpoints. If the bulk of the implementation works for any fixedpoint, admitting coinductive definitions will cost almost nothing.

AFA could be the basis for a greatest fixedpoint approach in Isabelle/ZF. It would be straightforward to separate FA from the other ZF axioms and to move most of the formalization into the resulting theory of ZF^- . Isabelle can support parallel developments in ZF and $ZF^- + AFA$. However, implementation of AFA would require much further work. The axiom and its consequences, such as the solution lemma, would have to be mechanized in a form suitable for constructing particular coalgebras (as opposed to developing metatheory).

My approach to final coalgebras is easy to mechanize. Most of the facts required of greatest fixed points are obtained by dualizing facts already proved about least fixed points. The definitions of variant pairs, products, sums, etc., are elementary. Their properties are easily established; many proofs can be adapted from those for the standard operators. A set (analogous to U) closed under the most important constructors can be defined in terms of V_ω , whose theory is already needed for the inductive case.

This fixedpoint approach has been implemented as an Isabelle package (Paulson, 1994). In order to admit both inductive and coinductive definitions, the package takes the relevant notions of products, sums, etc., as parameters. The package does not prove that particular coinductively defined sets are final coalgebras, but the script needed to generate such a proof is fairly short. It was by developing this script that I obtained the ideas underlying lemma 3.9.

Frost (1995) has used the package to mechanize a substantial example taken from a tutorial on coinduction (Milner and Tofte, 1991). The semantics of a simple functional programming language is defined an unusual way: recursive functions are modelled as non-well-founded expressions. The theorem relates the dynamic and static semantics—values and types—via a correspondence relation that is defined coinductively. The chief difficulty in the mechanization is to justify the basic definitions, which involve mutual recursion and variant functions; fortunately, the package does most of the work. The proofs themselves are routine. The full development takes just over a minute to run.

Recall that the identity functor is not uniform on maps. The corresponding declaration in Isabelle/ZF turns out to have the wrong properties: the greatest fixedpoint is U when it should be a singleton.

8. Conclusions

Researchers in semantics seldom worry about how an object is constructed, provided it has the right abstract properties. From this point of view, the general theorems of Aczel and Mendler

(1989) and Barr (1993) yield final coalgebras for a great many functors, using techniques such as inverse limits and quotienting.

But there is an undoubted interest in the special final coalgebra theorem of Aczel (1988), proved using AFA. This theorem is weaker but concrete. The set of streams over A is simply the greatest fixedpoint of the functor $A \times -$, which is also that functor's final coalgebra. Its elements are easily visualized objects of the form $\langle a_0, a_1, a_2, \dots \rangle$.

The original motivation for my work was to treat streams and other infinite data structures. I wished to use the standard ZF axiom system as it was automated using Isabelle. Thomas Forster suggested that Quine's treatment of ordered pairs might help. Generalizing this treatment led to the new definition of functions (and thus infinite streams), in order to compare the approach with AFA. This part of the work closely follows Aczel (1988) and Rutten and Turi (1993), from the substitution lemma onwards. As Aczel has pointed out to me, this reuse of the development suggests general conditions under which a category possessing final coalgebras analogous to U and U_X satisfies a special final coalgebra theorem.

Compared with my early paper (Paulson, 1995a), the present development is more streamlined and goes further. Its treatment of urelements eliminates most embeddings, simplifying the derivations. New laws govern iterated substitution and maps of the form U_f . Final coalgebras may be defined with respect to parameters. Much of the new material is relevant to systems based upon AFA.

My version of the theorem is less general than the version using AFA, especially for modelling concurrency. Here is a typical example. Let \mathcal{P}_f be the finite powerset operator, which returns the set of all finite subsets of its argument. Let A be a set of actions, and consider the set P of processes defined as the final coalgebra of $\mathcal{P}_f(A \times -)$. With AFA the final coalgebra is the greatest solution of $P = \mathcal{P}_f(A \times P)$, and if $p \in P$ then

$$p = \{\langle a_1, p_1 \rangle, \dots, \langle a_n, p_n \rangle\}$$

with $n < \omega$, $a_1, \dots, a_n \in A$ and $p_1, \dots, p_n \in P$. Here p represents a process that can execute action a_i and become process p_i , with no restriction that a_1, \dots, a_n are distinct. In this way, Aczel (1988) modelled the transition systems of SCCS, and other process algebras require at least as much generality.

My approach does not handle general set constructions, only variant tuples and functions; I do not know how to model \mathcal{P}_f respecting set equalities such as $\{x, y\} = \{y, x\} = \{x, y, x\}$. However, it is not entirely useless for modelling concurrency. In the UNITY formalism Chandy and Misra (1988), nondeterminism lies only in the choice of action, the actions themselves being deterministic. We could model UNITY by the set of the non-well-founded A -branching trees, but not by the greatest solution of $P = A \overset{\sim}{\rightarrow} P$, which is trivial (prop. 2.3). Instead we should use the greatest solution of $P = \{1\} \cup (A \overset{\sim}{\rightarrow} P)$, which is of course U^A , taking A as the index set.

The approach works best in its original application, infinite data structures. We can model the main constructions in U^ω . Since $U^\omega \subseteq V_{\omega+1}$, each infinite data structure is a subset of V_ω and thus is a set of hereditarily finite sets.* Section 2.1 discussed infinite streams. The set S of streams

* An *hereditarily finite set* is one built in finitely many stages from the empty set. There are countably many of them.

over A is the greatest solution of $S = A \tilde{\times} S$, and is the final coalgebra of the functor $A \tilde{\times} -$. The construction is parametric in A , yielding the functor $\text{stream}(A)$ that can be used in further definitions. Another possible application is the modelling of object-oriented languages (Hensel *et al.*, 1998).

Thus we have an account of non-well-founded phenomena that is concrete enough to be understood directly, and simple enough to use in machine proof. One can argue about the constructive validity of the cumulative hierarchy, but V_ω is uncontroversial even from an intuitionistic viewpoint. An infinite data structure is represented by a countable set of elementary objects.

Aczel has shown that by adopting AFA we can obtain final coalgebras as greatest fixedpoints, dualizing a standard result about initial algebras. My approach is another way of doing the same thing, though for fewer functors. Whether or not one chooses to adopt AFA hinges on a number of issues: philosophical, theoretical, practical. Variant tuples and functions are a simple alternative.

References

- Abramsky, S. (1990). The lazy lambda calculus. In D. A. Turner, editor, *Research Topics in Functional Programming*, pages 65–116. Addison-Wesley.
- Aczel, P. (1988). *Non-Well-Founded Sets*. CSLI.
- Aczel, P. and Mendler, N. (1989). A final coalgebra theorem. In D. H. Pitt, D. E. Rydeheard, P. Dybjer, A. M. Pitts, and A. Poigné, editors, *Category Theory and Computer Science*, pages 356–365. Springer. LNCS 389.
- Barr, M. (1993). Terminal coalgebras in well-founded set theory. *Theoretical Computer Science*, **114**(2), 299–315.
- Chandy, K. M. and Misra, J. (1988). *Parallel Program Design: A Foundation*. Addison-Wesley.
- Frost, J. (1995). A case study of co-induction in Isabelle. Technical Report 359, Comp. Lab., Univ. Cambridge.
- Hensel, U. and Jacobs, B. (1997). Proof principles for datatypes with iterated recursion. In E. Moggi and G. Rosolini, editors, *Category Theory and Computer Science*, pages 220–241. Springer. LNCS 1290.
- Hensel, U., Huisman, M., Jacobs, B., and Tews, H. (1998). Reasoning about classes in object-oriented languages: Logical models and tools. In C. Hankin, editor, *European Symposium on Programming*, LNCS 1381, pages 105–121.
- Kunen, K. (1980). *Set Theory: An Introduction to Independence Proofs*. North-Holland.
- Milner, R. (1989). *Communication and Concurrency*. Prentice-Hall.
- Milner, R. and Tofte, M. (1991). Co-induction in relational semantics. *Theoretical Computer Science*, **87**, 209–220.
- Moss, L. S. and Danner, N. (1997). On the foundations of corecursion. *Logic Journal of the IGPL*, **5**(2), 231–257.
- Paulin-Mohring, C. (1993). Inductive definitions in the system Coq: Rules and properties. In M. Bezem and J. Groote, editors, *Typed Lambda Calculi and Applications*, LNCS 664, pages 328–345. Springer.
- Paulson, L. C. (1993). Set theory for verification: I. From foundations to functions. *J. Auto. Reas.*, **11**(3), 353–389.

- Paulson, L. C. (1994). A fixedpoint approach to implementing (co)inductive definitions. In A. Bundy, editor, *Automated Deduction — CADE-12 International Conference*, LNAI 814, pages 148–161. Springer.
- Paulson, L. C. (1995a). A concrete final coalgebra theorem for ZF set theory. In P. Dybjer, B. Nordström, and J. Smith, editors, *Types for Proofs and Programs: International Workshop TYPES '94*, LNCS 996, pages 120–139. Springer.
- Paulson, L. C. (1995b). Set theory for verification: II. Induction and recursion. *J. Auto. Reas.*, **15**(2), 167–215.
- Quine, W. V. (1966). On ordered pairs and relations. In *Selected Logic Papers*, chapter VIII, pages 110–113. Random House. Originally published 1945–6.
- Rutten, J. J. M. M. and Turi, D. (1993). On the foundations of final semantics: Non-standard sets, metric spaces, partial orders. In J. de Bakker, W.-P. de Roever, and G. Rozenberg, editors, *Semantics: Foundations and Applications*, pages 477–530. Springer. LNCS 666.