Guarded induction on final coalgebras

D. Pavlović
COGS, University of Sussex, Brighton, and
Kestrel Institute, Palo Alto
e-mail: duskop@cogs.susx.ac.uk

Abstract

We make an initial step towards categorical semantics of guarded induction. While ordinary induction is usually modelled in terms of least fixpoints and initial algebras, guarded induction is based on *unique* fixpoints of certain operations, called guarded, on *final* coalgebras. So far, such operations were treated syntactically [3, 6, 7, 16]. We analyse them categorically. Guarded induction appears as couched in coinduction.

The applications of the presented categorical analysis span across the

"In order to establish that a proposition ϕ follows from other propositions $\phi_1, \ldots, \phi_q$, it is enough to build a proof term e for it, using not only natural deduction, case analysis and already proven lemmas, but also using the proposition we want to prove recursively, provided such a recursive call is guarded by introduction rules. We call this proof principle the 'guarded induction principle'."

— Th. Coquand [6, sec. 2.3]

1 Introduction

Coinduction is usually presented and studied as dual to induction: if induction is interpreted in terms of the universal property of initial algebras, coinduction arises from the couniversal property of final coalgebras [8, 10, 11, 17, 21, 22]. A bit like in the case of monads and comonads, the symmetry, with one side more familiar, opens an easier access to the other side. It provides a very rich source of parallel concepts and techniques [21] — but unfortunately goes only as far as it goes, and not further.

In fact, the most interesting conceptual distinctions often begin to surface only when the symmetry starts breaking down. Going back to monads and comonads, recall, e.g., how the free algebras for a monad form an algebra classifier (the clone), whereas the cofree coalgebras for a comonad do not seem to either classify or "coclassify" anything meaningful. And indeed, the former turns out to be the foundation of a rich mathematical theory, capturing algebraic varieties by *functorial semantics* [14, 15], whereas the latter remains a symptom of the fundamental fact that this theory does not have a dual: coalgebras for comonads on toposes tend to form toposes again, rather than "covarieties".

The present paper is an effort towards analysing an observed asymmetry of induction and coinduction: *coinductively constructed objects conspicuously often come about as domains on which we perform inductive constructions.* Not only models of computation, but even the universes of such models tend to be coinductively constructed — apparently *in order to* accomodate induction [1]. On the other hand, some basic structures of real analysis can be captured in a similar setting, with induction embedded in a coinductively defined domain [20].

1.1 Guarded induction is induction

In basic cases, this interplay of induction and coinduction is easy to understand. Take, e.g., the product functor $\Sigma \times (-) : \mathsf{Set} \longrightarrow \mathsf{Set}$. Its greatest fixpoint is the set Σ^ω of infinite streams in Σ, with the final coalgebra structure

$$\langle \mathrm{head}, \mathrm{tail} \rangle \quad : \quad \Sigma^\omega \longrightarrow \Sigma \times \Sigma^\omega$$

It accomodates the *stream induction*, where head takes care for the base case, and tail for the step. Using the inverse cons : $\Sigma \times \Sigma^\omega \longrightarrow \Sigma^\omega$ of the structure map $\langle \text{head}, \text{tail} \rangle$ — sometimes abbreviated to $a.x = \text{cons}(a, x)$ — the inductive definition

$$\begin{array}{rcl} \text{head}(x) & = & a \\ \text{tail}(x) & = & x \end{array} \tag{1}$$

becomes the equation

$$x = a.x \tag{2}$$

The prefixing $a.(-) : \Sigma^\omega \longrightarrow \Sigma^\omega$ is the simplest guarded operation. Its unique fixpoint is the unique solution of the corresponding inductive system of equations (1).

This surely looks like a very simple example, but it is very typical. For instance, an interesting bit of differential equations can be hidden behind it. Take Σ to be the set $\mathbb{R}$ of real numbers. The final coalgebra Σ^ω then contains the set $\mathbb{A}$ of analytic functions: every $f \in \mathbb{A}$ can indeed be represented as the stream $[f(0), f'(0), f''(0), \ldots]$. As observed by M.H. Escardó[1] [20], the $\langle \text{head}, \text{tail} \rangle$-structure restricts to $\mathbb{A}$ in the form

$$\begin{array}{rcl} \text{head}(f) & = & f(0) \\ \text{tail}(f) & = & f' \end{array}$$

while its inverse becomes

$$\text{cons}(a, g) = a + \int_0^x g\, dt$$

It is not hard to see that the coalgebra $\mathbb{A}$ is final for all $\langle h, t \rangle : A \longrightarrow \mathbb{R} \times A$ such that for every $\alpha \in A$ there is some $x > 0$ with $\sum_{n=0}^{\infty} \frac{ht^n(\alpha)}{n!} x^n < \infty$. An inductive definition in terms of head and tail now becomes an initial value problem, while a guarded equation like (2) becomes the corresponding integral equation.

The first guarded equations, introduced in CCS [16, sec. 3.2], were of a similar kind, e.g.

$$x = a.x + bc.x \tag{3}$$

The operation $+$ can be understood as the union of non-wellfounded sets [2]. Formally, it is the inverse of the structure map

$$\ni \ : \ \mathcal{V} \longrightarrow \wp\mathcal{V}$$

which makes the class $\mathcal{V}$ of non-wellfounded sets into a final coalgebra for the powerset functor $\wp : \mathsf{SET} \longrightarrow \mathsf{SET}$. The map $\ni$ assigns to each element of $\mathcal{V}$ the set of its elements. We write $x \ni y$ instead of $y \in \ni(x)$.

[1] and perhaps also by C.A.R. Hoare [9], who writes respectively α_0 and α' for the head and the tail of a trace α

If non-wellfounded sets are presented as (irredundant) trees [18], it becomes clear that $\ni$ supports the *tree induction.* Equations like (3) are solved by a

So we end up with two methods for constructing unique fixpoints of operations on final coalgebras: one direct, based on their couniversal property, the other inductive, and more general. Can such basic tools lead up to a discipline of *coinductive programming*, where programs, real functions and other infinitary objects would be extracted as fixpoints from specifications written in the form of guarded equations? Section 4 plays with this idea, investigating the compositionality of the prefixing and of the guarded operations.

2 Prefixing

Lemma 2.1 *Let $F : \mathbb{C} \longrightarrow \mathbb{C}$ be a functor and Υ its fixpoint, i.e. an object of $\mathbb{C}$, given together with an isomorphism*

$$\Upsilon \underset{\delta}{\overset{\varrho}{\rightleftarrows}} F\Upsilon \qquad (\cong)$$

Furthermore, let $\eta : \mathrm{id} \longrightarrow F$ be an arbitrary natural transformation, and ∂ the composite

$$\partial \quad : \quad \Upsilon \xrightarrow{\eta\Upsilon} F\Upsilon \xrightarrow{\delta} \Upsilon$$

The following commutativity conditions are then equivalent.

(a) $\forall x \exists! \langle x \rangle.$

$$\begin{array}{ccc} X & \overset{\langle x \rangle}{\dashrightarrow} & \Upsilon \\ {\scriptstyle x}\downarrow & & \\ X & & \downarrow{\scriptstyle \varrho} \\ {\scriptstyle \eta X}\downarrow & & \\ FX & \underset{F\langle x \rangle}{\dashrightarrow} & F\Upsilon \end{array}$$

(b) $\forall x \exists! \langle x \rangle.$

$$\begin{array}{ccc} X & \overset{\langle x \rangle}{\dashrightarrow} & \Upsilon \\ & & \uparrow{\scriptstyle \delta} \\ {\scriptstyle x}\downarrow & & F\Upsilon \quad (\partial) \\ & & \uparrow{\scriptstyle \eta\Upsilon} \\ X & \underset{\langle x \rangle}{\dashrightarrow} & \Upsilon \end{array}$$

(c) $\exists fix \forall x f.$

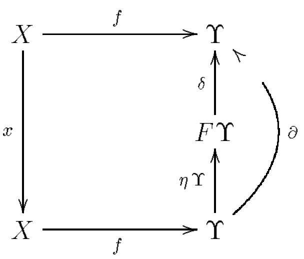

The equivalent condition (c) yields the desired fixpoint $\mathit{fix} : 1 \longrightarrow \Upsilon$. In fact, it is just the coalgebra homomorphism from the prefix $\eta 1 : 1 \longrightarrow F1$ to $\varrho : \Upsilon \longrightarrow F\Upsilon$.

While $\partial : \Upsilon \longrightarrow \Upsilon$ may extend to various natural transformations $\eta : \mathrm{id} \longrightarrow F$

They respectively extend to $\eta^0 : \mathrm{id} \longrightarrow \wp$, the components of which take everything to $\emptyset$, and $\eta^1 : \mathrm{id} \longrightarrow \wp$, where $\eta^1_X : X \longrightarrow \wp X$ takes $x \in X$ to the singleton $\{x\} \in \wp X$.

Combining the above, one gets on the class of synchronisation trees $\mathcal{V}_\Sigma$, as the greatest fixpoint of $F = \wp_\Sigma$, a prefixing operation

$$\partial^a(x) \quad {}^{a}$$

3 Guarded operations

3.1 Cones and coalgebras

In a category $\mathbb{C}$ with a final object 1, every functor $F : \mathbb{C} \longrightarrow \mathbb{C}$ induces a tower νF, like on

$$\begin{array}{ccccccccccc} \nu F & = & 1 & \xleftarrow{!} & F1 & \xleftarrow{F!} & F^2 1 & \xleftarrow{F^2!} & F^3 1 & \xleftarrow{F^3!} & \cdots \\ & & \uparrow{\scriptstyle !} & & \uparrow{\scriptstyle F!} & & \uparrow{\scriptstyle F^2!} & & \uparrow{\scriptstyle F^3!} & & \\ \Xi & = & X & \xrightarrow[\xi]{} & FX & \xrightarrow[F\xi]{} & F^2 X & \xrightarrow[F^2\xi]{} & F^3 X & \xrightarrow[F^3\xi]{} & \cdots \end{array} \tag{5}$$

while every coalgebra $\xi : X \longrightarrow FX$ induces a tower Ξ. Hence the cone $p = p^\xi : X \longrightarrow \nu F$, with the components

$$\begin{aligned} p_0 &: X \xrightarrow{!} 1 \\ p_{i+1} &: X \xrightarrow{\xi} FX \xrightarrow{Fp_i} F^{i+1}1 \end{aligned} \tag{6}$$

If $F^\omega 1$ is defined as the limit of νF, the cone p factorizes through $p_\omega : X \longrightarrow F^\omega 1$. On the other hand, $F^{\omega+1}1 = FF^\omega 1$ comes with an obvious cone to νF as well, which induces $F^\omega ! : F^{\omega+1}1 \longrightarrow F^\omega 1$. Proceeding in this way, the tower νF and the cone p can both be extended transfinitely.

If νF ever becomes stationary, in the sense that for some ordinal α, the arrow $\delta = F^\alpha ! : F^{\alpha+1}1 \longrightarrow F^\alpha 1$ is an isomorphism, then $\Upsilon = F^\alpha 1$ will be the greatest fixpoint of F: the inverse $\varrho : \Upsilon \longrightarrow F\Upsilon$ of δ will yield the final F-coalgebra structure [13, 23].

Of course, νF will surely become stationary at α if F preserves limits of the towers of length α. In fact, if $F : \mathbb{C} \longrightarrow \mathbb{C}$ does not preserve such limits, but $\mathbb{C}$ is a concrete category with objects bounded by some inaccessible cardinal κ, then F can usually be extended to a larger category $\widehat{\mathbb{C}}$, containing $\mathbb{C}$ as a full subcategory, *and* having the limits of κ-towers. The extension of F to $\widehat{\mathbb{C}}$ is then defined as to preserve such limits — and hence to have the greatest fixpoint. The familiar construction [2] of the universe of non-wellfounded sets as the greatest fixpoint of (the extension of) the powerset functor $\wp : \mathsf{Set} \longrightarrow \mathsf{Set}$ (to the category SET of classes) can be viewed as an example of this method [4, prop. 1.3].

Alternatively, if the F-images of finite objects are finite, and $\mathbb{C}$ has the limits of countable towers, one can take the finitary restriction $F_{\mathsf{fin}} : \mathbb{C}_{\mathsf{fin}} \longrightarrow \mathbb{C}_{\mathsf{fin}}$ of F and then extend it to $F_{\mathsf{fin}} : \mathbb{C} \longrightarrow \mathbb{C}$, but in such a way that the limits of the countable towers are preserved. Applied to the powersets $\wp : \mathsf{Set} \longrightarrow \mathsf{Set}$, this method 3].

functor as to preserve the limits of κ-towers: here, indeed, F_{fin} gets extended as to preserve the limits of the $\aleph_0$-towers[3].

In any case, the preceding duscussion shows that the following assumption causes no significant loss of generality, as it can usually be enforced with enough inaccessible cardinals (or Grothendieck universes), and often even without them.

Assumption. *In the sequel, the functor F will always preserve the limits of κ-towers, for some fixed κ, so that its greatest fixpoint Υ comes about as the limit $F^\kappa 1$, where the κ-tower νF stabilizes.*

As pointed out before, the coalgebra structure $\varrho : \Upsilon \longrightarrow F\Upsilon$ is obtained as the inverse of the stabilizing isomorphism $\delta : F^\kappa 1 \longrightarrow FF^\kappa 1$. The cone $p : \Upsilon \longrightarrow \nu F$, induced as in (6) by $\xi = \varrho$, will in this case be a limit cone.

On the other hand, taking (5) with $X = 1$, any $\xi : 1 \longrightarrow F1$ induces a corresponding tower Ξ as a "splitting" of νF. For each $i < \kappa$, (5) now gives a cone $\upsilon_i : F^i 1 \longrightarrow \nu F$, with $\upsilon_{i+1} \circ F^i \xi = \upsilon_i$. Since Υ is the limit of νF, these cones induce $u_i : F^i 1 \longrightarrow \Upsilon$, satisfying $u_{i+1} \circ F^i \xi = u_i$.

Since each u_i is defined as the factorisation of the cone $\upsilon_i : F^i 1 \longrightarrow \nu F$ through the limit cone $p : \Upsilon \longrightarrow \nu F$, the arrow $p_m \circ u_n : F^n 1 \longrightarrow F^m 1$ must be the m-th component of υ^n, that is

$$p_m \circ u_n \quad = \quad \begin{cases} F^{m-1}\xi \circ \cdots \circ F^n \xi & \text{if } m > n \\ \mathrm{id} & \text{if } m = n \\ F^m ! \circ \cdots \circ F^{n-1} ! & \text{if } m < n \end{cases} \tag{7}$$

In particular,

Lemma 3.1 *For a final F-coalgebra Υ, all limit cone components $p_i : \Upsilon \longrightarrow F^i 1$ are split epi, as soon as there is some arrow $1 \longrightarrow F1$.*

3.2 Guards

Definition 3.2 *A* guard *of an operation $\partial : X \longrightarrow X$ with respect to a coalgebra $\xi : X \longrightarrow FX$ is a family $\eta = \langle \eta_0, \eta_1, \eta_2 \ldots \rangle$, such that the squares*

$$\begin{array}{ccc} F^i 1 & \xleftarrow{\;p_i\;} & X \\ & & \big\downarrow \partial \\ \eta_i \big\downarrow & & X \\ & & \big\downarrow \xi \\ F^{i+1} 1 & \xleftarrow[\;Fp_i\;]{} & FX \end{array} \tag{8}$$

[3] Although $\aleph_0$ is often explicitly, by definition, excluded from the class of inaccessible cardinals, it actually possesses both of the relevant closure properties: for all $\zeta < \aleph_0$ holds $2^\zeta < \aleph_0$ and $|\cup \zeta| < \aleph_0$.

commute for all $i \geq 0$, with p_i constructed as in (6). An operation ∂ is said to be guarded *if there is some guard η for it.*

Remark. By definition (6) of p_i, square (8) commutes if and only if

$$\begin{array}{ccc} F^i 1 & \xleftarrow{\;p_i\;} & X \\ \Big\downarrow{\scriptstyle \eta_i} & & \Big\downarrow{\scriptstyle \partial} \\ F^{i+1} 1 & \xleftarrow[\;p_{i+1}\;]{} & X \end{array} \tag{9}$$

commutes.

Proposition 3.3 *Every prefixing operation is guarded.*

Proof. If the composite $\varrho \circ \partial : \Upsilon \longrightarrow F\Upsilon$ extends to a natural transformation $\eta : \mathrm{id} \longrightarrow F$, then the family consisting of $\eta_i = \eta F^i 1$ constitutes a guard of ∂ with respect to ϱ. □

Examples. On the coalgebra $\mathbb{A}$ of analytic functions, a guard η actually approximates the action of the corresponding operation ∂ polynomially. The commutativity of (9) means that the approximation of $\partial(f)$ of order $i+1$ is completely determined by the approximation of f of order i. The component η_i of the guard expresses that determination. In fact, any initial value problem that can be solved by the method of power series — i.e. inductively — induces a guarded operator. The details are in [19, 20].

The constant $\partial^{\infty}(x) = \infty$ on $\mathcal{V}$ is guarded by the maps $\eta_i : \wp^i 1 \longrightarrow \wp^{i+1} 1$, defined

$$\begin{array}{rcl} \eta_0 & = & 1 \\ \eta_{i+1}(x) & = & \{\eta_i\} \end{array}$$

The operation $\partial^{a,b}(x) = a.x + b.x$ on $\mathcal{V}_\Sigma$ is guarded by $\eta_i : \wp^i_\Sigma 1 \longrightarrow \wp^{i+1}_\Sigma 1 = \wp(\Sigma \times$

3.3 Guarded o erations on final coalgebras and their fix oints

As explained in 3.1, when Υ is the final coalgebra for F, it is natural to assume that $p : \Upsilon \longrightarrow \nu F$ is a limit cone. This means, of course, that the arrows $p_i : \Upsilon \longrightarrow F^i 1$ are jointly monic.

On the other hand, if there is a guarded operation on Υ, each $p_i : \Upsilon \longrightarrow F^i 1$ will be a split epi. Indeed, a guard

On the other hand, the $(i+1)$-st component of the cone corresponding to $\partial \circ \mathit{fix} : 1 \longrightarrow \Upsilon$ is

$$\begin{aligned} p_{i+1} \circ \partial \circ \mathit{fix} &= \eta_i \circ p_i \circ \mathit{fix} \\ &= \eta_i \circ \mathit{fix}_i \\ &= \mathit{fix}_{i+1} \end{aligned}$$

Hence $\partial \circ \mathit{fix} = \mathit{fix}$.

Towards the uniqueness, suppose $\partial \circ f = f : X \longrightarrow \Upsilon$. Writing $p_i \circ f$ as f_i, we have

$$\begin{aligned} f_{i+1} &= p_{i+1} \circ f \\ &= p_{i+1} \circ \partial \circ f \\ &= \eta_i \circ p_i \circ f \\ &= \eta_i \circ f_i \end{aligned}$$

Since f_0 is obviously $! : X \longrightarrow 1$,

$$f_i = \mathit{fix}_i \circ !$$

follows by induction over i. $\square$

Remark. If a coalgebra is not final, a guarded operation may not have a fixpoint, or may have many. E.g., the universe V of *wellfounded* setsof

For any $n \geq 1$ and the n-tuple composite F^n of $F : \mathbb{C} \longrightarrow \mathbb{C}$, each F-coalgebra $\xi : X \longrightarrow FX$ gives rise to an F^n-coalgebra

$$\xi^n \quad : \quad X \xrightarrow{\xi} FX \xrightarrow{F\xi} F^2X \xrightarrow{F^2\xi} \cdots \xrightarrow{F^{n-1}\xi} F^nX \tag{10}$$

Clearly, if ξ

Corollary 4.3 *If $\varrho : \Upsilon \longrightarrow F\Upsilon$ is a final coalgebra as above, then any composite of prefixing operations with respect to it has a unique fixpoint.*

Proof. By lemma 4.1, a composite of n prefixing operations with respect to F will be a prefixing operation with respect to F^n. By lemma 4.2, the final F-coalgebra $\varrho : \Upsilon \longrightarrow F\Upsilon$ yields the final F^n-coalgebra $\varrho^n : \Upsilon \longrightarrow F^n\Upsilon$. Applying corollary 2.3 (i.e. the constructions preceding it), we get the unique fixpoint of the composite prefixing as the unique coalgebra homomorphism to ϱ^n. □

4.2 Com osite guards

Similarly as above, a composite of n operations guarded with respect to ξ is guarded with respect to ξ^n. The point is now that it is also guarded with respect to ξ itself.

Proposition 4.4 *An operation $\partial : X \longrightarrow X$ is guarded with respect to $\xi : X \longrightarrow FX$ as soon as it is guarded with respect to any of $\xi^n : X \longrightarrow F^nX$, for $n \geq 1$.*

Proof. Given a guard $\eta^n = \langle \eta_0^n, \eta_1^n, \eta_2^n, \ldots \rangle$ of $\partial : X \longrightarrow X$ with respect to $\xi^n : X \longrightarrow F^nX$, a guard $\eta = \langle \eta_0, \eta_1, \eta_2, \ldots \rangle$ with respect to $\xi : X \longrightarrow FX$ will

The arrow p^n_k is a component of the cone $p^n : X \longrightarrow \nu F^n$, induced by ξ^n and (5–6). Clearly, p^n is a subcone of $p : X \longrightarrow \nu F$, and in particular

$$p^n_k \quad = \quad p_{nk}$$

References

[1] S. Abramsky and D. Pavlovic, Process Categories as Fixpoints *(in preparation)*

[2] P. Aczel, *Non-Well-Founded Sets*

[21] J. Rutten, Universal coalgebra: a theory of systems. CWI Report CS-R9652

[22] J. Rutten and D. Turi, Initial algebra and final coalgebra semantics for concurrency. In: *A Decade of Concurrency* (J.W. de Bakker et al., eds.) Lecture Notes in Computer Science 803 (Springer 1994) 530–582

[23] M.B. Smyth and G.D. Plotkin, The category theoretic solution of recursive domain equations. *SIAM J. Comp.* 11(1982) 761–783