

Coalgebra, lecture 11: Distributive Laws

Jurriaan Rot (and Aleks Kissinger)

November 22, 2016

1 Operations on stream systems

We start with everybody's favorite operation on streams, $\text{alt}: A^\omega \times A^\omega \rightarrow A^\omega$, defined by stream differential equations:

$$\text{alt}(\sigma, \tau)(0) = \sigma(0) \quad \text{alt}(\sigma, \tau)' = \text{alt}(\tau', \sigma') \quad (1)$$

These equations uniquely define alt . One way to see that, is by turning the domain $A^\omega \times A^\omega$ into a stream system, in such a way that alt arises as the unique homomorphism to the final coalgebra A^ω . We're going to look at such definitions a bit more systematically and generally.

First, we rephrase the definition of alt as the following inference rule (or actually rules, one for each $a, b \in A$):

$$\frac{x \xrightarrow{a} x' \quad y \xrightarrow{a} y'}{\text{alt}(x, y) \xrightarrow{a} \text{alt}(y', x')} \quad (2)$$

What do we mean with these rules? The idea is that they define alt as a construction on stream systems. More precisely, let $\langle o, f \rangle: X \rightarrow A \times X$ be a stream system. If we interpret $x \xrightarrow{a} x'$ as $\langle o, f \rangle(x) = (a, x')$, then the above rules give us a new stream system, defined by:

$$\begin{aligned} \text{alt}_f: X \times X &\rightarrow A \times (X \times X) \\ (x, y) &\mapsto (o(x), (f(y), f(x))) \end{aligned}$$

If we apply this construction to the *final* coalgebra $z: A^\omega \rightarrow A \times A^\omega$, we obtain alt , as in (1), as the unique homomorphism to the final coalgebra.

$$\begin{array}{ccc} A^\omega \times A^\omega & \xrightarrow{\text{alt}} & A^\omega \\ \text{alt}_z \downarrow & & \downarrow z \\ A \times (A^\omega \times A^\omega) & \xrightarrow{\text{id}_A \times \text{alt}} & A \times A^\omega \end{array} \quad (3)$$

We make all of this a bit more precise. What we have here is:

- *syntax*: one binary operation alt , modeled by the functor $S: \text{Set} \rightarrow \text{Set}$, $S(X) = X \times X$;
- *behaviour*: stream systems, modeled by (coalgebras for) the well-known functor $B: \text{Set} \rightarrow \text{Set}$, $B(X) = A \times X$;
- the *specification* of alt , as given in (2). We model it by a *distributive law*, which is a natural transformation $\lambda: SB \Rightarrow BS$, in this case defined as follows:

$$\begin{aligned} \lambda_X: (A \times X) \times (A \times X) &\rightarrow A \times (X \times X) \\ ((a, x'), (b, y')) &\mapsto (a, (y', x')) \end{aligned} \quad (4)$$

This allows us to map a stream system $X \xrightarrow{f} A \times X$ to

$$\text{alt}_f: X \times X \xrightarrow{f} (A \times X) \times (A \times X) \xrightarrow{\lambda_X} A \times (X \times X).$$

The point here is that λ here fully captures the specification of alt , in a precise way, which we can immediately apply to any stream system. In particular, applying it to the final coalgebra, we get alt_z and alt as in (3).

2 Distributive laws

The above story about alt is, of course, an instance of a much more general picture. This picture involves a functor $S: \mathcal{C} \rightarrow \mathcal{C}$ on a category \mathcal{C} modelling (for now) the syntax of some operation(s) we would like to define (such as a binary operator, or a whole collection of operators); $B: \mathcal{C} \rightarrow \mathcal{C}$ modelling the behaviour of interest (such as streams, transition systems, or automata); and a distributive law between them, modelling a specification of the operations (such as the specification of alt).

Definition 2.1. Let $S, B: \mathcal{C} \rightarrow \mathcal{C}$ be functors. A *distributive law* of S over B is a natural transformation of the form $\lambda: SB \Rightarrow BS$.

Such a distributive law defines a functor $\bar{S}: \text{Coalg}(B) \rightarrow \text{Coalg}(B)$, defined on objects by:

$$X \xrightarrow{f} BX \quad \mapsto \quad SX \xrightarrow{Sf} SBX \xrightarrow{\lambda_X} BSX.$$

and on arrows by $\bar{S}(h) = S(h)$. It is a nice exercise to prove that this is well-defined, i.e., that $\bar{S}(h)$ is indeed a coalgebra homomorphism. Notice that, if we apply \bar{S} to a coalgebra with carrier X , then the result is a coalgebra with carrier $S(X)$. A fancy way of saying this is that \bar{S} is a *lifting* of S , which means that the following diagram of categories and functors commutes:

$$\begin{array}{ccc} \text{Coalg}(B) & \xrightarrow{\bar{S}} & \text{Coalg}(B) \\ U \downarrow & & \downarrow U \\ \mathcal{C} & \xrightarrow{S} & \mathcal{C} \end{array}$$

where U is the forgetful functor, mapping a coalgebra (X, f) to its carrier X .

Now, if B has a final coalgebra (Z, z) , then we can apply the lifting \bar{S} (the result is on the left-hand side of the diagram below) to it and obtain a homomorphism from SZ to Z , which we'll call beh^λ :

$$\begin{array}{ccc}
 SZ & \xrightarrow{\text{beh}^\lambda} & Z \\
 Sz \downarrow & & \downarrow z \\
 SBZ & & \\
 \lambda_Z \downarrow & & \\
 BSZ & \xrightarrow{B(\text{beh}^\lambda)} & BZ
 \end{array} \tag{5}$$

Notice that beh^λ is an *algebra* for the functor S (which in fact is an algebra for \bar{S}). So a distributive law apparently always defines an *algebra on the final coalgebra*.

If we think of λ as a specification of operations, then we can think of beh^λ as the interpretation of these operations on the final coalgebra. For instance:

- In the setting of the previous section, if we define λ as in (4), $\text{beh}^\lambda: A^\omega \times A^\omega$ becomes the alt operation on A^ω .
- We can define multiple operations at once, simply with a different choice of S . For instance, to combine alt with a constant stream a for every $a \in A$, we take $S(X) = A + (X \times X)$, and define $\lambda: SB \Rightarrow BS$ by

$$\begin{aligned}
 \lambda_X: A + ((A \times X) \times (A \times X)) &\rightarrow A \times (A + (X \times X)) \\
 (a, 1) &\mapsto (a, (a, 1)) \\
 (((a, x'), (b, y')), 2) &\mapsto (a, ((y', x'), 2))
 \end{aligned}$$

- Let $B(X) = 2 \times X^A$. We will define another binary operation, so let $S(X) = X \times X$. Consider $\lambda_X: SB \Rightarrow BX$, defined by

$$\begin{aligned}
 \lambda_X: (2 \times X^A) \times (2 \times X^A) &\rightarrow 2 \times (X \times X)^A \\
 ((b, f), (c, g)) &\mapsto (b \wedge c, a \mapsto (f(a), g(a)))
 \end{aligned}$$

Written in terms of rules, we are defining a binary operation \otimes :

$$\frac{x \xrightarrow{a} x' \quad y \xrightarrow{a} y'}{x \otimes y \xrightarrow{a} x' \otimes y'} \quad \frac{x \downarrow \quad y \downarrow}{(x \otimes y) \downarrow}$$

The (carrier of the) final coalgebra of B is $\mathcal{P}(A^*)$. What will $\text{beh}^\lambda: \mathcal{P}(A^*) \times \mathcal{P}(A^*) \rightarrow \mathcal{P}(A^*)$ be?

2.1 What operations can we define?

The main point so far of using distributive laws, is that they give us an abstract specification format for operations on (final) coalgebras. We have used that to define operations on streams and automata, but you could plug in any B instead. But one can wonder how much the specification formats allow us to do. For instance, consider the rule

$$\frac{x \xrightarrow{a} x'}{f(x) \xrightarrow{a} f(f(x'))}$$

which defines a (silly) operation on stream systems. If we try to model that as $\lambda: SB \Rightarrow BS$, with $S(X) = X$ and $B(X) = A \times X$, we get stuck: for the right-hand side $f(f(x))$ we need “two levels” of S .


The solution is to look instead at different types of natural transformations. In particular, natural transformations of the form $\lambda: SB \Rightarrow BS^*$, where S^* is the free monad for S , do allow to write rules such as the above. They induce (and correspond, in a sense that is beyond the current treatment) distributive laws of the form $\rho: S^*B \Rightarrow BS^*$, essentially by induction. With an even more general version we can define, for instance, the operations of regular expressions as natural transformations. But we won't go into this for now.

3 Determinisation

In one of the previous lectures, we've looked at non-deterministic automata. There can be modeled as coalgebras: a non-deterministic automaton is a coalgebra of the form

$$\langle \epsilon, \delta \rangle: X \rightarrow 2 \times (\mathcal{P}(X))^A$$

(we might as well take the finite powerset functor, if we'd like to speak about a final coalgebra, but we don't really need to at this point). The associated coalgebraic notion of bisimilarity does not coincide with language equivalence. The classical example is the automata below:



States u_1 and v_1 are clearly language equivalent, both accepting the empty language, but they are not bisimilar. This is not wrong: the branching aspect is something we can actually observe, and in some case it makes sense not to equate the above systems.

However, in many cases we're actually interested in language equivalence, and of course we'd like to fit that into the coalgebraic picture. To this end, we'll use a classic trick: the powerset construction, which turns a non-deterministic automaton into a deterministic automaton. In fact—like many things in the theory of coalgebras—this turns out to be yet another instance of something general and abstract, involving distributive laws.

First, the concrete definition. From a given non-deterministic automaton $\langle \epsilon, \delta, \rangle : X \rightarrow 2 \times (\mathcal{P}(X))^A$, we define $\langle \epsilon^\#, \delta^\# \rangle : \mathcal{P}(X) \rightarrow 2 \times (\mathcal{P}(X))^A$ by

$$\begin{aligned}\epsilon^\#(S) &= \bigvee_{x \in S} \epsilon(x) \\ \delta^\#(S)(a) &= \bigcup_{x \in S} \delta(x)(a)\end{aligned}$$

for all $S \in \mathcal{P}(X)$ and $a \in A$. This is a *deterministic automaton*, whose carrier is the powerset of states. This process is exactly the classical powerset construction. In a sense, we're hiding the stuff we'd like not to observe inside the state space.

We obtain a language semantics $\text{beh} : \mathcal{P}(X) \rightarrow \mathcal{P}(A^*)$ by finality. This assigns to each set $S \in \mathcal{P}(X)$ a language; we obtain the language of a state $x \in X$ by applying it to the singleton $\{x\}$, i.e., $\text{beh}(\{x\})$. All in all, we get the following picture:

$$\begin{array}{ccccc} X & \xrightarrow{\eta_X} & \mathcal{P}(X) & \xrightarrow{\text{beh}} & \mathcal{P}(A^*) \\ \langle \epsilon, \delta \rangle \downarrow & & \swarrow \langle \epsilon^\#, \delta^\# \rangle & & \downarrow \\ 2 \times (\mathcal{P}(X))^A & \xrightarrow{\text{id} \times (\text{beh})^A} & & & 2 \times \mathcal{P}(A^*)^A \end{array} \quad (7)$$

The arrow on the right is the final coalgebra of languages. The $\eta_X : X \rightarrow \mathcal{P}(X)$ just maps each element of X to the corresponding singleton. If you did the exercises last week, you'll recognise it as the X -component of the unit of the powerset monad—this will be important later on. The *language semantics* now arises as the composite $\text{beh} \circ \eta$. That this is really the intended language semantics is just the correctness of the powerset construction.

It turns out this whole construction is actually concisely captured by a distributive law. First, we decompose the functor of non-deterministic automata as $B\mathcal{P}$, where $BX = 2 \times X^A$ (the functor for deterministic automata) and \mathcal{P} is the powerset functor (giving the branching/non-determinism). Now, the distributive law is $\lambda : \mathcal{P}B \Rightarrow B\mathcal{P}$, given on a component X by:

$$\begin{aligned}\lambda_X : \mathcal{P}(2 \times X^A) &\rightarrow 2 \times (\mathcal{P}(X))^A \\ S &\mapsto \left(\bigvee_{(b,f) \in S} b, a \mapsto \bigcup_{(b,f) \in S} f(a) \right)\end{aligned}$$

This distributive law takes a *set* of “behaviours of a deterministic automaton”, and turns it into a “single” (non-deterministic) behaviour. How does this encode

the powerset construction? Well, a little thought (and/or computation) shows us that the automaton in (8) corresponds to:

$$\mathcal{P}(X) \xrightarrow{\mathcal{P}(\langle \epsilon, \delta \rangle)} \mathcal{P}(2 \times (\mathcal{P}(X))^A) \xrightarrow{\lambda_{\mathcal{P}(X)}} 2 \times (\mathcal{P}(\mathcal{P}(X)))^A \xrightarrow{\text{id} \times (\mu_X)^A} 2 \times (\mathcal{P}(X))^A$$

Here $\mu_X: \mathcal{P}(\mathcal{P}(X)) \rightarrow \mathcal{P}(X)$ is given by union; it is the multiplication of the powerset monad (\mathcal{P}, η, μ) .

With this distributive law, its not so difficult to prove that we obtain a functor

$$\bar{\mathcal{P}}: \text{Coalg}(B\mathcal{P}) \rightarrow \text{Coalg}(\mathcal{P})$$

from the category of non-deterministic automata to deterministic automata, defined on objects by the powerset construction. Actually, this is another example of a *lifting* defined by a distributive law, although it is different than the one we've seen before.

$$\begin{array}{ccc} \text{Coalg}(B\mathcal{P}) & \xrightarrow{\bar{\mathcal{P}}} & \text{Coalg}(B) \\ U \downarrow & & \downarrow U \\ \text{Set} & \xrightarrow{\mathcal{P}} & \text{Set} \end{array}$$

The fact that it is a lifting, means that it acts as the powerset on the state space of automata (and on homomorphisms).

Let's see a different example. A *partial automaton* is a coalgebra of the form $\langle \epsilon, \delta \rangle: X \rightarrow 2 \times (X + 1)^A$, where $1 = \{*\}$. For every state and alphabet symbol, we have either a single next state, or no next state. The latter should just mean that no more words with that letter in front should be accepted. Again, this is not captured by coalgebraic bisimilarity, which sees the 1 part as explicit termination (a nice exercise to work out).

So we'll use another determinisation procedure. Now, we construct a deterministic automaton $\langle \epsilon^\#, \delta^\# \rangle: X + 1 \rightarrow 2 \times (X + 1)^A$ (notice that it's state space is $X + 1$) defined as follows: ... (left as an exercise!).

The idea is that this construction canonically adds a "sink" or "trap" state, which is not accepting and has transitions to itself for every alphabet letter. Since the result is a deterministic automaton, we get a unique coalgebra morphism $\text{beh}: X + 1 \rightarrow \mathcal{P}(A^*)$. And the language of a state $x \in X$ is just the language $\text{beh}(x, 1)$. We get another nice picture.

$$\begin{array}{ccc} X & \xrightarrow{\eta_X} & X + 1 & \xrightarrow{\text{beh}} & \mathcal{P}(A^*) & (8) \\ \langle \epsilon, \delta \rangle \downarrow & & \swarrow \langle \epsilon^\#, \delta^\# \rangle & & \downarrow & \\ 2 \times (X + 1)^A & & & \xrightarrow{\text{id} \times (\text{beh})^A} & 2 \times \mathcal{P}(A^*)^A & \end{array}$$

The η_X here maps x to $(x, 1)$. The *language semantics* arises as the composite $\text{beh} \circ \eta$.

Last week, we've seen the Option monad, which we'll call (O, η, μ) here to save some typing work for the author. Partial automata are coalgebras for the composite functor BO . Similar to the case of non-deterministic automata, the determinisation construction for partial automata is captured concisely as a distributive law $\lambda: OB \Rightarrow BO$. And again, the determinisation construction can be rephrased in terms of λ . All of this is left as an exercise.

4 Determinisation: the general picture

We have seen two determinisation constructions, for non-deterministic automata and partial automata. In both cases, the main ingredients were:

1. a functor B , whose coalgebras model deterministic automata;
2. a monad (T, η, μ) , which captures the branching aspect that we would like to hide;
3. a distributive law $\lambda: TB \Rightarrow BT$, which captures the actual determinisation construction.

So both seem to be part of a general picture. Indeed, with the above ingredients, for an arbitrary coalgebra $f: X \rightarrow BTX$, we can define a B -coalgebra

$$TX \xrightarrow{Tf} TBTX \xrightarrow{\lambda_X} BTTX \xrightarrow{B\mu_X} BTX \quad (9)$$

which is the “determinisation” of f . This construction defines a lifting of T :

$$\begin{array}{ccc} \text{Coalg}(BT) & \xrightarrow{\bar{T}} & \text{Coalg}(B) \\ U \downarrow & & \downarrow U \\ \mathcal{C} & \xrightarrow{T} & \mathcal{C} \end{array}$$

giving an abstract determinisation procedure. This picture generalizes both the one for non-deterministic automata and for partial automata; and there are many other examples.

There is a lot more to say about this, in particular on the interplay between the distributive law and the monad structure. If that sounds exciting to you, have a look at the paper below.

References

- [1] Alexandra Silva, Filippo Bonchi, Marcello M. Bonsangue, Jan J. M. M. Rutten: Generalizing determinization from automata to coalgebras. *Logical Methods in Computer Science* 9(1) (2013)