Coalgebra, lecture 6: Bisimulations

Jurriaan Rot

October 8, 2018

In the lectures so far, we have seen that, given a functor $F\colon \mathsf{Set} \to \mathsf{Set}$, we obtain notions of:

- *F-coalgebras*, which model state based systems;

- *homomorphisms* between F-coalgebras are behaviour-preserving maps;

- *final F-coalgebra* (if it exists), which models the behaviour of all F-coalgebras, and assigns behaviour to states of F-coalgebras.

Final coalgebra thereby gives us a notion of behavioural equivalence: given states $x, y \in X$ of an F-coalgebra (X, f), we say x and y are *behaviourally equivalent* if $\mathsf{beh}(x) = \mathsf{beh}(y)$, where beh is the unique coalgebra homomorphism from (X, f) to the final F-coalgebra. For instance:

- For the functor $F(X) = A \times X$, F-coalgebras are stream systems over A, and the final coalgebra consists of the set of streams A^ω over A; two states of a stream system are behaviourally equivalent if they represent the same stream.

- For $F(X) = 2 \times X^A$, F-coalgebras are deterministic automata, and the final coalgebra consists of the set of languages 2^{A^*}. Two states of a deterministic automaton are behaviourally equivalent if they accept the same language.

- For $F(X) = X + 1$, F-coalgebras, are ..., well, maps $f\colon X \to X + 1$ (deterministic systems with termination?), and the final coalgebra consists of 'extended natural numbers' $\mathbb{N} + \{\omega\}$. Two states are behaviourally equivalent if they terminate in the same number of steps.

It's nice that we can define behavioural equivalence, but how do we actually *prove* that two states are equivalent? This is, of course, an important problem if we want to analyse state-based systems: it's one of the most fundamental decision problems.

This question leads us to the main concept of today's lecture: the notion of *bisimulation* between coalgebras. Bisimulations are one of the most important concepts in the theory: they correspond to *proofs of behavioural equivalence*. What is nice, is that bisimulation is defined abstractly in terms of coalgebras;

by instantiating to a specific functor (such as the one for stream systems, or for automata) one can then derive a concrete proof technique (to prove, for instance, language equivalence of states of an automaton) and actually derive algorithms from this technique!

We'll start with the general definition of bisimulation, show some examples, and then prove that bisimulations are indeed a sound proof technique for behavioural equivalence; this is sometimes called 'coinduction proof principle'. We use automata as an extended example, and derive an algorithm for language equivalence.

1 Bisimulations and bisimilarity

Although it is possible to formulate the notion of bisimulation for coalgebras for functors on arbitrary categories, we'll focus on endofunctors on Set, since this significantly simplifies matters while still conveying the main ideas (and most of the basic examples are on Set anyway). So let $F\colon \mathsf{Set} \to \mathsf{Set}$ be a functor.

Given F-coalgebras (X, f) and (Y, g), a relation $R \subseteq X \times Y$ is a *bisimulation* (between (X, f) and (Y, g)) if there exists a coalgebra structure $c\colon R \to F(R)$ on R such that the following diagram commutes:

$$
\begin{array}{ccccc}
X & \xleftarrow{\;\pi_1\;} & R & \xrightarrow{\;\pi_2\;} & Y \\
{\scriptstyle f}\downarrow & & {\scriptstyle c}\downarrow & & \downarrow{\scriptstyle g} \\
F(X) & \xleftarrow[F(\pi_1)]{} & F(R) & \xrightarrow[F\pi_2]{} & F(Y)
\end{array}
$$

where π_1, π_2 are the projection maps of R. If $(X, f) = (Y, g)$ then we say R is a bisimulation *on* (X, f).

Of particular interest is the *largest* bisimulation (w.r.t subset inclusion) between (X, f) and (Y, g). It is called *bisimilarity*, and is often denoted by $\sim\; \subseteq X \times Y$. So, $x \sim y$ iff there exists a bisimulation $R \subseteq X \times Y$ such that $x \mathrel{R} y$. If $x \sim y$ then we say x and y are *bisimilar*.

1.1 Stream systems

First, as an example, let's spell out the details in case of our favorite Set endofunctor: $F(X) = A \times X$, where A is a fixed set. And, to slightly simplify notation, we'll just consider bisimulations on a single F-coalgebra, that is, a stream system $\langle o, f \rangle\colon X \to A \times X$. Instantiating the definition: a relation $R \subseteq X \times X$ is a bisimulation if there exists a stream system $\langle o_R, f_R \rangle\colon R \to A \times R$ such that the following diagram commutes:

$$
\begin{array}{ccccc}
X & \xleftarrow{\;\pi_1\;} & R & \xrightarrow{\;\pi_2\;} & X \\
{\scriptstyle \langle o,f \rangle}\downarrow & & {\scriptstyle \langle o_R,f_R \rangle}\downarrow & & \downarrow{\scriptstyle \langle o,f \rangle} \\
A \times X & \xleftarrow[\mathsf{id} \times \pi_1]{} & A \times R & \xrightarrow[\mathsf{id} \times \pi_2]{} & A \times X
\end{array}
$$

So that means we have, for all $(x, y) \in R$,

- $o(x) = o_R((x, y)) = o(y)$

- $f(x) = \pi_1(f_R(x, y))$

- $f(y) = \pi_2(f_R(x, y))$

The second and third point together amount to: $f_R(x, y) = (f(x), f(y))$. For this to be correct, we should of course have $(f(x), f(y)) \in R$.

From the above considerations, we can reformulate the notion of bisimulation on our stream system: R is a bisimulation iff for all $(x, y) \in R$:

- $o(x) = o(y)$, and

- $(f(x), f(y)) \in R$.

This concrete characterisation is equivalent to the instance of coalgebraic bisimulation which we started from.

As an example, consider the stream system $\langle o, f \rangle \colon X \to A \times X$, with $X = \{x, y, z, u, v\}$ and the output and transitions given by:

$$(1)$$

Let's prove that $x \sim u$. We should come up with a bisimulation R such that $x \mathrel{R} u$. The most economical choice would be $R_0 = \{(x, u)\}$. But does it work? Well, $o(x) = a = o(u)$, so that's fine, but $\ldots (f(x), f(u)) = (y, v) \notin R_0$. So R_0 is not a bisimulation. Ok, but we don't give up: lets take $R_1 = \{(x, u), (y, v)\}$. Then (x, u) is fine, but for (y, v) we have $(f(y), f(v)) \notin R_1 \ldots$ and this game continues for a while, until we arrive at

$$R = \{(x, u), (y, v), (z, u), (x, v), (y, u), (z, v)\}$$

and we finally have our desired bisimulation! By the way, which pairs are actually in the *largest* bisimulation $\sim$ in this example?

We conclude our little streams excursion by recovering an old friend from the second lecture: bisimulations between streams. Take the final coalgebra $\langle i, d \rangle \colon A^\omega \to A \times A^\omega$. A relation $R \subseteq A^\omega \times A^\omega$ on it is a bisimulation if for all $(\sigma, \tau) \in R$,

- $i(\sigma) = i(\tau)$, that is, $\sigma(0) = \tau(0)$, and

- $(d(\sigma), d(\tau)) \in R$, that is, $(\sigma', \tau') \in R$.

This is exactly the notion we saw in lecture 2, where we showed for instance that the relation

$$R = \{(even(tail(\sigma)), odd(\sigma)) \mid \sigma \in A^\omega\} \tag{2}$$

is a bisimulation. The point, which we showed in lecture 2, is that, whenever $R \subseteq A^\omega \times A^\omega$ is a bisimulation, then $(\sigma, \tau) \in R$ implies $\sigma = \tau$. Or, phrased differently, the greatest bisimulation $\sim$ is simply the equality relation $\{(\sigma, \sigma) \mid \sigma \in A^\omega\}$. We call this a *coinduction proof principle*. In a moment, we'll see that such a principle holds not only for streams but much more generally, using the final coalgebra.

1.2 Deterministic systems with termination

Of course, if I'm pretending to give a really general definition, we should have at least one more example. So let's consider coalgebras for the functor F given by $F(X) = X + 1$. According to the definition, a relation $R \subseteq X \times X$ on the states of a coalgebra $f \colon X \to X + 1$ is a bisimulation if for all $(x, y) \in R$ (let's assume for simplicity that X is disjoint from $1 = \{*\}$):

- either $f(x) = * = f(y)$, or $(f(x), f(y)) \in R$ (which implicitly means $f(x), f(y) \in X$).

This is equivalent to the coalgebraic definition of bisimulation for this functor; if you're not convinced immediately, working this out is a good exercise. We'll see some further instances in the weekly exercises.

2 Coinduction

So we defined when two states of a coalgebra are bisimilar, denoted $x \sim y$, which means $(x, y) \in R$ for some bisimulation R. But what's the point? Well, first of all, bisimilarity $\sim$ is a really important notion of equivalence in itself. For instance, in concurrency theory, it captures behavioural equivalence of processes; we'll talk about this next week. However, as it turns out, within the theory of coalgebras, it is a *proof technique for behavioural equivalence*. This is formally stated by the following result.

Theorem 2.1 (Coinduction proof principle). *Let $F \colon \mathsf{Set} \to \mathsf{Set}$ be a functor which has a final coalgebra (Z, z), and let (X, f) and (Y, g) be F-coalgebras. If $R \subseteq X \times Y$ is a bisimulation, then for all $x \in X$ and $y \in Y$:*

$$(x, y) \in R \quad \Rightarrow \quad \mathsf{beh}_f(x) = \mathsf{beh}_g(y) \,,$$

where $\mathsf{beh}_f \colon X \to Z$ and $\mathsf{beh}_g \colon Y \to Z$ are the unique coalgebra homomorphisms.

Proof. Suppose $(x, y) \in R$, and R is a bisimulation, witnessed by some coalgebra structure $c\colon R \to F(R)$. Consider the following diagram:

$$
\begin{array}{ccccccccc}
Z & \xleftarrow{\mathsf{beh}_f} & X & \xleftarrow{\pi_1} & R & \xrightarrow{\pi_2} & Y & \xrightarrow{\mathsf{beh}_g} & Z \\
\downarrow{\scriptstyle z} & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle c} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle z} \\
F(Z) & \xleftarrow{F(\mathsf{beh}_f)} & F(X) & \xleftarrow{F(\pi_1)} & F(R) & \xrightarrow{F\pi_2} & F(Y) & \xrightarrow{F(\mathsf{beh}_g)} & F(Z)
\end{array}
$$

The above diagram commutes: the inner squares since R is a bisimulation (with coalgebra structure c) and the outer two squares since beh_f and beh_g are, by definition, coalgebra homomorphisms of this type.

As a consequence, both $\mathsf{beh}_f \circ \pi_1$ and $\mathsf{beh}_g \circ \pi_2$ are coalgebra homomorphisms from (R, c) to the final coalgebra (Z, z), hence

$$
\mathsf{beh}_g \circ \pi_2 = \mathsf{beh}_f \circ \pi_1 .
$$

Thus, for any $(x, y) \in R$, we have $\mathsf{beh}_f(x) = \mathsf{beh}_f(\pi_1(x, y)) = \mathsf{beh}_g(\pi_2(x, y)) = \mathsf{beh}_g(y)$ as needed. $\square$

As a consequence, to show that two states x, y are behaviourally equivalent, it suffices to show that they are related by a bisimulation. For instance, if R is a bisimulation between stream systems, then any $(x, y) \in R$ means that x and y represent the same stream. So in particular, we obtain the rather spectacular fact that x and u in (1) represent the same stream! Ok, that's maybe not the most exciting example; but of course, the technique works for more complicated stream systems as well. One class of examples is the bisimulation proofs on streams which we did in lecture 2, see for instance (2) above. By Theorem 2.1, we recover the correctness of this technique: if $(x, y) \in Z$ are elements of a final coalgebra, then $x = y$ (notice that $\mathsf{beh} = \mathsf{id}$ in that case).

We formally state the above observation in a little theorem, which also includes a converse. In fact, the coinduction proof principle is sometimes stated in the following form: bisimilarity on the final coalgebra coincides with equality.

Theorem 2.2. *Let $F\colon \mathsf{Set} \to \mathsf{Set}$ be a functor which has a final coalgebra (Z, z). Then for all $(x, y) \in Z$: $x \sim y$ iff $x = y$.*

Proof. From left to right, this follows from Theorem 2.1. From right to left, one shows that the relation $\{(x, x) \mid x \in Z\}$ is a bisimulation. In fact, that holds on any coalgebra: a nice exercise. $\square$

The above theorem generalises what we showed in lecture 2: that bisimilarity of streams coincides with equality. Of course, the above theorem is much more general, applying to any kind of coalgebra; for instance, it also allows us to capture equality of languages through bisimilarity on the final deterministic automaton. We'll use this in another lecture; and in a moment, we'll concider bisimulations on (arbitrary) determinstic automata, which, by Theorem 2.1 suffice to prove language equivalence of states.

We finish our discussion of coinduction with a little observation on homomorphisms and bisimilarity. Notice that we obtain (one side of) Theorem 2.2 as a special case Theorem 2.1. As it turns out, one can also go the other way, and obtain Theorem 2.1 from Theorem 2.2. That argument uses that, if h, k are homomorphisms, then $x \sim y$ implies $h(x) \sim k(y)$ — we won't go into details at this point.

However, one may also wonder whether there is a converse: is it the case that $h(x) \sim k(y)$ implies $x \sim y$? And specifically, do we have a converse for Theorem 2.1, that is, does $\mathrm{beh}(x) = \mathrm{beh}(y)$ imply $x \sim y$? This is a reasonable question, as it is a kind of completeness: if it holds, this means that whenever two states are behaviourally equivalent, we can prove this using bisimulations. However, as it turns out, the answer is no: there are some functors for which this is not the case. Fortunately, it holds under mild conditions, and in particular for most functors that we are interested in: for instance, all (Kripke) polynomial functors, that is, functors built using products, coproducts, exponents, constant functors, the identity functor and the (finite) powerset functor. The condition typically used is that the functor preserves weak pullbacks, but this is a bit beyond the scope of the course right now; if you're interested, have a look, for instance, at [4].

3 Bisimulations on automata

In the previous lecture, we considered the functor $F(X) = 2 \times X^A$. Recall that it's coalgebras are deterministic automata, of the form:

$$\langle o, \delta \rangle \colon X \to 2 \times X^A$$

where o is the output function and δ the transition function. Further, recall that the language semantics $\mathrm{beh} \colon X \to 2^{A^*}$ arises as the unique map into the final coalgebra $(2^{A^*}, \langle e, d \rangle)$ with $e(L) = L(\varepsilon)$ and $d(L)(a) = L_a$ (language derivative, that is, $L_a(w) = L(aw)$). Thus, beh is the unique map making the following diagram commute:

$$
\begin{array}{ccc}
X & \xrightarrow{\ \mathrm{beh}\ } & 2^{A^*} \\
{\scriptstyle \langle o, \delta \rangle} \downarrow & & \downarrow {\scriptstyle \langle e, d \rangle} \\
2 \times X^A & \xrightarrow[\mathrm{id} \times \mathrm{beh}^A]{} & 2 \times (2^{A^*})^A
\end{array}
$$

Concretely, we have

- $\mathrm{beh}(x)(\varepsilon) = o(x)$, and

- $\mathrm{beh}(x)(aw) = \mathrm{beh}(\delta(x)(a))(w)$

for all $x \in X$, $a \in A$, $w \in A^*$. This corresponds to the usual language semantics of deterministic automata.

The notion of bisimulation will provide us with the means for checking *language equivalence*: whether two states $x, y \in X$ of a deterministic automaton accept the same language, i.e., whether $\mathsf{beh}(x) = \mathsf{beh}(y)$. Consider, for instance, the following automaton (for simplicity of notation, we view the entire thing as a single automaton):

(3)

In this example, x_1, x_3 and y_2 are all behaviourally equivalent; so are x_2, y_1 and y_3. Of course, this is a simple example, just for the sake of argument; for bigger automata, language equivalence becomes much less obvious.

First of all, we'll spell out the notion of bisimulation for deterministic automata. As always, we start by instantiating the definition: a relation $R \subseteq X \times X$ on the states of an automaton $(X, \langle o, \delta \rangle)$ is a bisimulation if there exists a coalgebra structure $\langle o_R, \delta_R \rangle \colon X \to 2 \times X^A$ such that the following diagram commutes:

So, equivalently, we must have, for all $(x, y) \in R$:

- $o(x) = o_R(x, y) = o(y)$,

- $\forall a. \delta(x)(a) = \pi_1(\delta_R(x, y)(a))$ and

- $\forall a. \delta(y)(a) = \pi_2(\delta_R(x, y)(a))$.

A coalgebra structure $\langle o_R, \delta_R \rangle$ satisfying the above conditions exists if and only if the following holds, for all $(x, y) \in R$:

- $o(x) = o(y)$, and

- $\forall a \in A. \; \delta(x)(a) \; R \; \delta(y)(a)$.

We have obtained a concrete notion of bisimulation for deterministic automata: a relation $R \subseteq X \times X$ satisfying the above two requirements (which don't anymore mention a coalgebra structure!): if two states x, y are related, then they should either both be accepting or both be non-accepting; and for every alphabet letter, the states reached after taking a transition (labelled by that letter) from both states should be again related.

It follows from Theorem 2.1, that if $(x, y) \in R$ for some bisimulation R, then x and y are language equivalent: $\mathsf{beh}(x) = \mathsf{beh}(y)$. Thus, to prove language equivalence, it suffices to construct a suitable bisimulation.

To start with, let's show that x_1, y_2 in (3) are equivalent, using bisimulations. The first attempt could be

$$R_0 = \{(x_1, y_2)\}\,.$$

Both x_1, y_2 are non-accepting. However, $\delta(x_1)(a) = x_3$, whereas $\delta(y_2)(a) = y_2$, so this is not a bisimulation, since $(x_3, y_2) \notin R_0$; also, $(\delta(x_1)(b), \delta(y_2)(b)) = (x_2, y_1) \notin R_0$. So we add both pairs, and let

$$R_1 = \{(x_1, y_2), (x_3, y_2), (x_2, y_1)\}$$

This is still not a bisimulation ... however, after one more step, we reach

$$R_2 = \{(x_1, y_2), (x_3, y_2), (x_2, y_1), (x_2, y_3)\}$$

which is actually a bisimulation. Hence, $\mathsf{beh}(x_1) = \mathsf{beh}(y_2)$, that is, x_1 and y_2 are language equivalent.

Note that we could have (somehow) guessed the entire relation R_2 at once; it's not at all necessary to construct in the above, incremental way. However, what is nice about this, is that it looks pretty algorithmic: one just adds successor states and keeps checking whether both are accepting/non-accepting, until either the relation is a bisimulation, or one has to add some pair (x, y) which violates the first condition, that is, where one is accepting and the other is not. In fact, in the latter case, we have found a counterexample: a reason that the two states are *not* language equivalent.

Not only does this look algorithmic, we can just turn it into an algorithm! To this end, we'll focus now on *finite* automata: we assume that the set of states X is finite. We call the algorithm Naive, as we'll see in a moment how to improve it. This, as well as the adapated algorithm in the next section, is more or less literally copied from [2] and from notes written by Filippo Bonchi and myself (see [1]). Given a deterministic automaton $(X, \langle o, \delta \rangle)$, and states $x_0, y_0 \in X$, the algorithm is as follows:

<u>$\mathtt{Naive}(x_0, y_0)$</u>

```
(1)  R is empty; todo is empty;
(2)  insert (x₀, y₀) in todo;
(3)  while todo is not empty do
 (3.1)   extract (x, y) from todo;
 (3.2)   if (x, y) ∈ R then continue;
 (3.3)   if o(x) ≠ o(y) then return false;
 (3.4)   for all a ∈ A,
             insert (δ(x)(a), δ(y)(a)) in todo;
 (3.5)   insert (x, y) in R;
(4)  return true;
```

First of all, observe that the algorithm is correct: $\mathtt{Naive}(x_0, y_0)$ returns true iff $x_0 \sim y_0$ (iff $\mathsf{beh}(x_0) = \mathsf{beh}(y_0)$). To see this, notice that the following is a loop invariant: for all $(x, y) \in R$, we have

- $o(x) = o(y)$, and

- $\forall a \in A.\ (\delta(x)(a), \delta(y)(a)) \in R \cup todo$.

Since in the end $todo$ is empty, it follows that, if the algorithm terminates returning true, the relation R is a bisimulation. Conversely, suppose $\mathsf{beh}(x_0) \neq \mathsf{beh}(y_0)$. Then eventually a pair (x, y) will be added such that $o(x) \neq o(y)$, hence the algorithm returns false, in (3.3).

If the state space X is finite, then the algorithm terminates. To see this, notice first that every iteration of the loop, an element is taken from $todo$. Further, new elements are added to $todo$ only if some new element is also added to R. Hence, the number of times new elements are added to $todo$ is bounded by the number of pairs of elements of X.

3.1 Bisimulation up to equivalence

The previous algorithm was derived more or less directly from the instance of coalgebraic bisimulation for deterministic automata. One can think of similar algorithms for other types of coalgebras, just based on the associated notion of bisimulation.

This is all nice, but … it turns out that, in many cases, one can do better: it is possible to prove behavioural equivalence in an 'easier' way, by slightly relaxing the requirements of a bisimulation. In case of deterministic automata, this leads to a classical algorithm for language equivalence, as was observed in [2].

We first look at a motivating example. Consider the following automaton, over a singleton alphabet.


For simplicity, no state is accepting. So it should be easy to prove that x_0 and y_0 are bisimilar! However … the bisimulation we compute by starting with $\{(x_0, y_0)\}$ and proceeding with our algorithm is slightly unpleasant. Depicting it with dashed arrows, we obtain a very readable picture.

This is in fact the smallest bisimulation that contains (x_0, y_0).

Luckily we can do better, by using a so-called *bisimulation up-to technique*. In the course, we'll hear much more about up-to techniques, but this is our first aquaintance. Bisimulations up-to are variants of bisimulations which are easier to construct than bisimulations. However, if they're any good (they come in different kinds) than they are *sound*, which means that, even though they'ren not a bisimulation themselves, they *represent* a good old bisimulation. Therefore, they still suffice to prove behavioural equivalence.

Back to the example: consider the following relation R depicted by the dashed lines:

This is not a bisimulation: the pair (x_4, y_3) violates it, since (x_1, y_4) is not related. However, (x_1, y_4) is related by the least equivalence relation $\mathsf{eqv}(R)$ containing R. More precisely, $\mathsf{eqv}(R)$ is the least relation satisfying the following rules:

$$\frac{x \; R \; y}{x \; \mathsf{eqv}(R) \; y} \qquad \frac{}{x \; \mathsf{eqv}(R) \; x} \qquad \frac{x \; \mathsf{eqv}(R) \; y}{y \; \mathsf{eqv}(R) \; x} \qquad \frac{x \; \mathsf{eqv}(R) \; y \quad y \; \mathsf{eqv}(R) \; z}{x \; \mathsf{eqv}(R) \; z}$$

With that in place, for the above relation given by the dashed lines, we have, for all $(x, y) \in R$:

- $o(x) = o(y)$, and

- $\forall a \in A. \; (\delta(x)(a), \delta(y)(a)) \in \mathsf{eqv}(R)$.

Such a relation is called a *bisimulation up to equivalence*.

And now what? Well, it turns out that bisimulations up to equivalence suffice to prove language equivalence. This is called *soundness*, and we may express it nicely as another coinduction principle:

Theorem 3.1. *If R is a bisimulation up to equivalence on a deterministic automaton $(X, \langle o, \delta \rangle)$, then for any $(x, y) \in R$: $\mathsf{beh}(x) = \mathsf{beh}(y)$.*

Proof. One can show that, if R is a bisimulation up to equivalence, then $\mathsf{eqv}(R)$ is a bisimulation (exercise!). $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad\square$

In the second half of the course, we'll treat up-to techniques in general, and show really nice ways of proving such soundness results.

Anyway, the point here is that bisimulations up to equivalence suffice to prove language equivalence. And they're easier to construct than plain bisimulations! We can now update our Naive algorithm to something we call HK:

$$\underline{\mathtt{HK}(x_0,\, y_0)}$$

```
(1) R is empty; todo is empty;
(2) insert (x₀, y₀) in todo;
(3) while todo is not empty do
 (3.1)   extract (x, y) from todo;
 (3.2)   if (x, y) ∈ eqv(R) then continue;
 (3.3)   if o(x) ≠ o(y) then return false;
 (3.4)   for all a ∈ A,
            insert (δ(x)(a), δ(y)(a)) in todo;
 (3.5)   insert (x, y) in R;
(4) return true;
```

The only difference with `Naive` is in (3.2); this algorithm constructs a bisimulation up to equivalence rather than a bisimulation. Hence, it is still correct, but it terminates earlier in general. In fact, the above algorithm was devised by Hopcroft and Karp (not knowing about bisimulations, as far as we know); it was observed in [2] that this amounts to the computation of a bisimulation up to equivalence.

4   Final remarks

The main thing of this lecture is the notion of bisimulation, which provides a general notion and proof technique for behavioural equivalence of coalgebras. We've mainly seen our favorite instances, streams and automata; however, bisimulation is defined rather generally, in terms of a (Set) endofunctor. For instance, next week, we'll consider the case of labelled transition systems, recovering the classical case of bisimulations introduced first, historically (and which is still the most famous instance of bisimulation).

A little remark on terminology: we've been using the word *coinduction* sometimes to *define* maps into a final coalgebra ("coinductive" definition of an operation such as *even* or *odd* on streams), and, mainly today, to talk about a *proof* technique. The common theme here is that we refer to coinduction as the "use of final coalgebras"; as, for instance, in [3]. In computer science, there's also a more classical use of the word 'coinduction', in the context of partial orders; we'll see this in the last part of the course.

References

[1] Filippo Bonchi, Marcello Bonsangue, and Jurriaan Rot. Lecture notes: coalgebraic methods for automata, 2018. `http://esslli2018.folli.info/wp-content/uploads/coma.pdf`.

[2] Filippo Bonchi and Damien Pous. Checking NFA equivalence with bisimulations up to congruence. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2013, proceedings*, pages 457–468. ACM, 2013.

[3] Bart Jacobs and Jan Rutten. An introduction to (co)algebras and (co)induction. In *Advanced Topics in Bisimulation and Coinduction*, pages 38–99. Cambridge University Press, 2012.

[4] Jan J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theor. Comput. Sci.*, 249(1):3–80, 2000.