

Simplified Coalgebraic Trace Equivalence

Alexander Kurz¹, Stefan Milius^{3*}, Dirk Pattinson², and Lutz Schröder^{3**}

¹ University of Leicester

² The Australian National University

³ Friedrich-Alexander-Universität Erlangen-Nürnberg

Abstract. The analysis of concurrent and reactive systems is based to a large degree on various notions of process equivalence, ranging, on the so-called linear-time/branching-time spectrum, from fine-grained equivalences such as strong bisimilarity to coarse-grained ones such as trace equivalence. The theory of concurrent systems at large has benefited from developments in coalgebra, which has enabled uniform definitions and results that provide a common umbrella for seemingly disparate system types including non-deterministic, weighted, probabilistic, and game-based systems. In particular, there has been some success in identifying a generic coalgebraic theory of bisimulation that matches known definitions in many concrete cases. The situation is currently somewhat less settled regarding trace equivalence. A number of coalgebraic approaches to trace equivalence have been proposed, none of which however cover all cases of interest; notably, all these approaches depend on explicit termination, which is not always imposed in standard systems, e.g. labelled transition systems. Here, we discuss a joint generalization of these approaches based on embedding functors modelling various aspects of the system, such as transition and branching, into a global monad; this approach appears to cover all cases considered previously and some additional ones, notably standard and probabilistic labelled transition systems.

1 Introduction

It was recognized early on that the initial algebra semantics of Goguen and Thatcher [10] needs to be extended to account for notions of observational or behavioural equivalence, see Giarratana, Gimona and Montanari [9], Reichel [18], and Hennicker and Wirsing [12]. When Aczel [2] discovered that at least one important notion of behavioural equivalence—the bisimilarity of process algebra—is captured by final coalgebra semantics, the study of coalgebras fully entered computer science. Whereas early work emphasized the duality between algebra and coalgebra, it became soon clear that both areas have to be taken together. For example, in the work of Turi and Plotkin [22], monads represent the programs, comonads represent their behaviour (operational semantics), and a distributive law between them ensures that the behaviour of a composed system is given by the behaviours of the components, or, more technically, that bisimilarity is a congruence.

* Supported by the Deutsche Forschungsgemeinschaft (DFG) under project MI 717/5-1

** Supported by the Deutsche Forschungsgemeinschaft (DFG) under project SCHR 1118/11-1

Another example of the interplay of algebraic and coalgebraic structure arises from the desire to make coalgebraic methods available for a larger range of program equivalences, e.g. the ones described in van Glabbeek [23]. To this end, Power and Turi [17] argued that trace equivalence arises from a distributive law $TF \rightarrow FT$ between a monad T describing the non-deterministic part and a functor F describing the deterministic part of a transition system $X \rightarrow TFX$. This was taken up by Hasuo et al. [11] and gave rise to a sequence of papers [16, 14, 20, 7, 8] that discuss coalgebraic aspects of trace equivalence.

We generalize this approach and call a trace semantics for coalgebras $X \rightarrow GX$ simply a natural transformation $G \rightarrow M$ for some monad M . This allows us, for example, and opposed to the work cited in the previous paragraph, to account for non-deterministic transition systems without explicit termination. Moreover, because of the flexibility afforded by choosing M , both trace semantics and bisimilarity can be accounted for in the same setting. We also show that for G being of the specific forms investigated in [11] and in [20, 7, 14] there is a uniform way of constructing a natural transformation of type $G \rightarrow M$ that induces the same generic trace semantics up to canonical forgetting of deadlocks.

2 Preliminaries

We assume that readers are familiar with basic notions of category theory. We work with a base category \mathbf{C} , which we may assume for simplicity to be locally finitely presentable, such as the category \mathbf{Set} of sets and functions (see e.g. Adámek and Rosický [4] for basics on locally finitely presentable categories and finitary functors).

Given a functor $G : \mathbf{C} \rightarrow \mathbf{C}$, a *G-coalgebra* is an arrow $\gamma : X \rightarrow GX$. Given two coalgebras $\gamma : X \rightarrow GX$ and $\gamma' : X' \rightarrow GX'$, a *coalgebra morphism* $f : (X, \gamma) \rightarrow (X', \gamma')$ is an arrow $f : X \rightarrow X'$ in \mathbf{C} such that $\gamma' \circ f = Gf \circ \gamma$.

When \mathbf{C} is a concrete category, i.e. \mathbf{C} comes equipped with a faithful functor $\mathbf{C} \rightarrow \mathbf{Set}$, we say that two states $x \in X$ and $x' \in X'$ in two coalgebras (X, γ) and (X', γ') are *behaviourally equivalent* if there are coalgebra morphisms f, f' with common codomain (Y, δ) such that $f(x) = f'(x')$.

Behavioural equivalence can be computed in a partition-refinement style using the *final coalgebra sequence* $(G^n 1)_{n < \omega}$ where 1 is a final object in \mathbf{C} and G^n is n -fold application of G . The projections $p_n^{n+1} : G^{n+1} 1 \rightarrow G^n 1$ are defined by induction where $p_0^1 : G \rightarrow 1$ is the unique arrow to 1 and $p_{n+1}^{n+2} = G(p_n^{n+1})$.

For any coalgebra (X, γ) , there is a *canonical cone* $\gamma_n : X \rightarrow G^n 1$ defined inductively by $\gamma_0 : X \rightarrow 1$ and $\gamma_{n+1} = G(\gamma_n) \circ \gamma$. We say that two states $x, x' \in X$ in (X, γ) are *finite-depth behaviourally equivalent* if $\gamma_n(x) = \gamma_n(x')$ for all $n < \omega$. (Note that if G is a finitary set functor, then finite-depth behavioural equivalence implies behavioural equivalence; this follows from the results of Worrell [24].)

A *monad* is given by an operation M on the objects of \mathbf{C} and, for each set X , a function $\eta_X : X \rightarrow MX$ and, for each $f : X \rightarrow MY$, a so-called Kleisli star $f^* : MX \rightarrow MY$ satisfying (i) $\eta_X^* = id_{MX}$, (ii) $f^* \circ \eta_X = f$, (iii) $(g^* \circ f)^* = g^* \circ f^*$ for all $g : Y \rightarrow MZ$. It follows that M is a functor, given by $Mf = (\eta \circ f)^*$, and η a natural transformation. Moreover, $\mu = id^* : MM \rightarrow M$ is a natural transformation


and satisfies $\mu \circ M\eta = \mu \circ \eta M = id$ and $\mu \circ M\mu = \mu \circ \mu M$. We obtain the Kleisli star back from μ and M by $f^* = \mu_Y \circ Mf$.

An *Eilenberg-Moore algebra* for the monad M is an arrow $\xi : MX \rightarrow X$ such that $\xi \circ \eta_X = id_X$ and $\xi \circ M\xi = \xi \circ \mu_X$.

Recall (e.g. from Kelly [15]) that an endofunctor G on a category \mathbf{C} is said to generate an *algebraically-free* monad G^* if the category of Eilenberg-Moore algebras of G^* is isomorphic over \mathbf{C} to the category of G -algebras (i.e. morphisms $GX \rightarrow X$). The monad G^* is then also the free monad over G ; conversely, free monads are algebraically-free if the base category \mathbf{C} is complete [5, 15]. For example, when \mathbf{C} is locally finitely presentable, then every finitary functor on \mathbf{C} , representing a type of finitely-branching systems, generates an (algebraically-)free monad (this follows from the free algebra construction in Adámek [3]).

3 A Simple Definition of Coalgebraic Trace Equivalence

Recall the classical distinction between bisimilarity and trace equivalence, the two ends of the *linear-time-branching time spectrum* [23]: to cite a much-belaboured standard example, the two labelled transition systems (over the alphabet $\Sigma = \{a, b, c\}$)



are *trace equivalent* in the usual sense [1], as they both admit exactly the traces ab and ac (and prefixes thereof), but not bisimilar, as bisimilarity is sensitive to the fact that the left hand side decides in the first step whether b or c will be enabled in the second step, while the right hand side leaves the decision between b and c open in the first step. In other words, trace equivalence collapses all future branches, retaining only the branching at the current state. Now observe that we can nevertheless construct the trace semantics by stepwise unfolding; to do this, we need to a) remember the last step reached by a given trace in order to continue the trace correctly, and b) implement the collapsing correctly in each step. E.g. for s_0 above, this takes the following form: let us call a pair (u, x) consisting of a word over Σ and a state x a *pretrace*. Before the first step, we assign, by default, the set $\{(\epsilon, s_0)\}$ of pretraces, where ϵ denotes the empty word. After the first step, we reach, applying both transitions simultaneously, the set $\{(a, s_{10}), (a, s_{11})\}$. After the second step, we reach, again applying two transitions, $\{(ab, s_{20}), (ac, s_{21})\}$. Note that after the third step, the set of pretraces will become empty if we proceed in the same manner, as s_{20} and s_{21} are both deadlocks. Thus, we will in general need to remember all the finite unfoldings of the set of pretraces, as traces ending in deadlocks will be lost on the way. Of course, for purposes of trace equivalence we are no longer interested in the states reached by a given trace, so we

forget the state components of all pretraces that we have accumulated, obtaining the expected prefix-closed trace set $\{\epsilon, a, ab, ac\}$.

Recall that we can understand labelled transition systems as coalgebras $\gamma : X \rightarrow \mathcal{P}(\Sigma \times X)$. What is happening in the unfolding steps is easily recognized as composition with γ in the Kleisli category of a suitable monad, specifically $M = \mathcal{P}(\Sigma^* \times _)$, a monad that contains the functor $\mathcal{P}(\Sigma \times _)$ via an obvious natural transformation α . Defining $\gamma^{(n)}$ as the n -fold iteration of the morphism $\alpha\gamma$ in the Kleisli category of M , we have $\gamma^{(0)}(s_0) = \{(\epsilon, s_0)\}$, $\gamma^{(1)}(s_0) = \{(a, s_{10}), (a, s_{11})\}$, $\gamma^{(2)}(s_0) = \{(ab, s_{20}), (ac, s_{21})\}$, and $\gamma^{(3)}(s_0) = \emptyset$. Forgetting the state component of the pretraces in these sets amounts to postcomposing with $M!$, where $!$ is the unique map into $1 = \{*\}$. These considerations lead to the following definitions.

Definition 1. A *trace semantics* for a functor G is a natural transformation $\alpha : G \rightarrow M$ into a monad M , the *global monad*. Given such an α and a G -coalgebra $\gamma : X \rightarrow GX$, we define the *iterations* $\gamma^{(n)} : X \rightarrow MX$ of γ , for $n \geq 0$, inductively by

$$\gamma^{(0)} = \eta_X \quad \gamma^{(n+1)} = (\alpha\gamma)^* \gamma^{(n)}$$

where the unit η and the Kleisli star $*$ are those of M (in particular $\gamma^{(1)} = \alpha\gamma$). Then the α -*trace sequence* of a state $x \in X$ is the sequence

$$T_\gamma^\alpha(x) = (M!\gamma^{(n)}(x))_{n < \omega},$$

with $!$ denoting the unique map $X \rightarrow 1$ as above. Two states x and y in G -coalgebras $\gamma : X \rightarrow GX$ and $\delta : Y \rightarrow GY$, respectively, are α -*trace equivalent* if

$$T_\gamma^\alpha(x) = T_\delta^\alpha(y).$$

(Although we use an element-based formulation for readability, this definition clearly does make sense over arbitrary complete base categories.)

Of course, one shows by induction over n that

$$\gamma^{(n+1)} = (\gamma^{(n)})^* \alpha\gamma \quad \text{for all } n < \omega. \quad (1)$$

We first note that the trace sequence factors through the initial ω -segment of the terminal sequence. Recall from Section 2 that a G -coalgebra γ induces a cone (γ_n) into the final sequence.

Lemma 2. Let $\alpha : G \rightarrow M$ be a trace semantics for G , and define natural transformations $\alpha_n : G^n \rightarrow M$ for $n < \omega$ recursively by $\alpha_0 = \eta$ and $\alpha_{n+1} = \mu\alpha G\alpha_n$. If γ is a G -coalgebra, then

$$M!\gamma^{(n)} = \alpha_n \gamma_n \quad \text{for all } n < \omega.$$

Proof. Induction on n .

$n = 0$: We have $M!\gamma^{(0)} = M!\eta = \eta! = \alpha_0\gamma_0$.


$n \rightarrow n + 1$: We have

$$\begin{aligned}
& \alpha_{n+1}\gamma_{n+1} \\
&= \mu\alpha G(\alpha_n)G\gamma_n\gamma && \text{(Definitions of } \gamma_{n+1}, \alpha_{n+1}\text{)} \\
&= \mu\alpha G(M!\gamma^{(n)})\gamma && \text{(Inductive hypothesis)} \\
&= \mu M(M!\gamma^{(n)})\alpha\gamma && \text{(Naturality of } \alpha\text{)} \\
&= M!\mu M\gamma^{(n)}\alpha\gamma && \text{(Naturality of } \mu\text{)} \\
&= M!(\gamma^{(n)})^*\alpha\gamma \\
&= M!\gamma^{(n+1)} && (1).
\end{aligned}$$

□

Corollary 3. *Finite-depth behaviourally equivalent states are α -trace equivalent.*

Remark 4. In most items of related work, stronger assumptions than we make here allow for identifying an *object* of traces in a suitable category, such as the Kleisli category [11] or the Eilenberg-Moore category [14, 7] of a monad that forms part of the type functor. In our setting, a similar endeavour boils down to characterizing, possibly by means of a limit of a suitable diagram, those α -trace sequences that are *G-realizable*, i.e. induced by a state in some *G*-coalgebra. We do not currently have a general answer for this but point out that in a variant of the special case treated in the beginning of the section where we take *G* to be $\mathcal{P}^*(\Sigma \times _)$, with \mathcal{P}^* denoting nonempty powerset, and $M = \mathcal{P}(\Sigma^* \times _)$, the set of *G*-realizable traces is the limit of the infinite diagram



where R denotes the immediate prefix relation $R = \{(u, ua) \mid u \in \Sigma^*, a \in \Sigma\}$ with projections $\pi_1, \pi_2 : R \rightarrow \Sigma^*$. We expect that this description generalizes to cases where G and M have the form TF and TF^* , respectively, where T is a monad and F^* is the free monad over the functor F , possibly under additional assumptions. In the case at hand, the limit of the diagram is the set of all subsets A of $\Sigma^* \times 1 \cong \Sigma^*$ that are prefix-closed and *extensible* in the sense that for every $u \in A$ there exists $a \in \Sigma$ such that $ua \in A$.

4 Examples

We show that various process equivalences are subsumed under α -trace equivalence.

Finite-depth behavioural equivalence Similarly to the approach via Eilenberg-Moore liftings (explained in more detail in Section 5), α -trace equivalence spans, at least for

finitely branching systems, the entire length of the linear-time-branching-time spectrum, in the sense that even (finite-depth) behavioural equivalence coincides with α -trace equivalence for a suitable α . This is conveniently formulated using the following terminology.

Definition 5. We say that an endofunctor G on a category with a terminal object 1 is *non-empty* if $G1$ has a global element, i.e. there exists a morphism $1 \rightarrow G1$.

Non-emptiness of an endofunctor entails that the component of α_n at 1 are sections where α_n is as in Lemma 2.

Lemma 6. *If G is non-empty and generates an algebraically-free monad G^* with universal arrow α , then $(\alpha_n)_1$ (the component of α_n at the terminal object) is a section for every $n < \omega$.*

Proof. For each set X , G^*X is the initial $G + X$ -algebra, with structure map

$$[\mu\alpha, \eta] : GG^*X + X \rightarrow G^*X$$

where μ and η are the multiplication and unit of G^* [5]. By Lambek's lemma, it follows that $[\mu\alpha, \eta]$ is an isomorphism. Since both summands of the coproduct $GG^*1 + 1$ are nonempty (for GG^*1 , this follows from non-emptiness of G : we obtain a global element of GG^*1 by postcomposing a global element of $G1$ with $G\eta_1 : G1 \rightarrow GG^*1$), the coproduct injections are sections, so we obtain that $\mu\alpha$ and η are sections, each being the composite of a section with an isomorphism. Using (1), it follows by induction that α_n is a section for each $n < \omega$. \square

(Notice that G is non-empty as soon as any GX has a global element; if the base category is Set , then every functor is non-empty except the constant functor for \emptyset .)

Proposition 7. *If G is a non-empty functor and generates an algebraically-free monad via $\alpha : G \rightarrow G^*$, then α -trace equivalence coincides with ω -behavioural equivalence.*

Proof. Immediate from Lemmas 2 and 6. \square

Labelled Transition Systems (LTS) We provide some additional details for our initial example: We have $GX = \mathcal{P}(\Sigma \times X)$ and $MX = \mathcal{P}(\Sigma^* \times X)$, with α the obvious inclusion. The monad M arises from G , as we will see later again in (2), from a distributive law $\delta_X : \Sigma \times \mathcal{P}(X) \rightarrow \mathcal{P}(\Sigma \times X)$ which maps a pair (a, S) to $\{a\} \times S$. Explicitly, the unit of M is given by $\eta(x) = \{(\epsilon, x)\}$, and the multiplication by $\mu(\mathfrak{A}) = \{(uv, x) \mid \exists(u, S) \in \mathfrak{A}. (v, x) \in S\}$ for $\mathfrak{A} \in \mathcal{P}(\Sigma^* \times \mathcal{P}(\Sigma^* \times X))$. For each n and each state x in an LTS $\gamma : X \rightarrow \mathcal{P}(\Sigma \times X)$, $\gamma^{(n)}(x)$ consists of the pretraces of x of length exactly n , i.e.

$$\gamma^{(n)}(x) = \{(u, y) \mid x \xrightarrow{u} y, u \in \Sigma^n\}$$

where \xrightarrow{u} denotes the usual extension of the transition relation to words $u \in \Sigma^*$. Thus, $M!\gamma^{(n)}(x)$ consists of the traces of x of length n , i.e.

$$M!\gamma^{(n)}(x) = \{(u, *) \mid x \xrightarrow{u}, u \in \Sigma^n\}$$

(where, as usual, $x \xrightarrow{u}$ denotes $\exists y. x \xrightarrow{u} y$). Thus, states x and y are α -trace equivalent iff they are trace equivalent in the usual sense, i.e. iff

$$\{u \in \Sigma^* \mid x \xrightarrow{u}\} = \{u \in \Sigma^* \mid y \xrightarrow{u}\}.$$

The entire scenario transfers verbatim to the case of finitely branching LTS, with $G = \mathcal{P}_\omega(\Sigma \times _)$ and $M = \mathcal{P}_{<\omega}(\Sigma^* \times _)$, where $\mathcal{P}_{<\omega}$ denotes finite powerset.

LTS with explicit termination The leading example treated in related work on coalgebraic trace semantics [11, 14, 7] is LTSs with explicit termination, described as coalgebras for the functor

$$\mathcal{P}(1 + \Sigma \times _) \cong 2 \times \mathcal{P}^\Sigma.$$

A state in an LTS with explicit termination can be seen as a non-deterministic automaton; this suggests that one might expect the traces of such a state to be the words accepted by the corresponding automaton, and, in fact, this is the stance taken in previous work [11, 14, 7]; for the sake of distinction, let us call this form of trace semantics *language semantics*. Staring at the problem for a moment reveals that language semantics does not fit directly into our framework: Basically, our definition of trace sequence assembles the traces via successive iteration of the coalgebra structure, and remembers the traces reached in each iteration step. Contrastingly, language semantics will drop a word from the trace set if it turns out that upon complete execution of the word, no accepting state is reached – in α -trace semantics, on the other hand, we will have recorded prefixes of the word on the way, and our incremental approach does not foresee forgetting these prefixes. See Section 5 for a discussion of how α -trace sequences can be further quotiented to obtain language semantics.

Indeed one might contend that a more natural trace semantics of an LTS with explicit termination will distinguish two types of traces: those induced by the plain LTS structure, disregarding acceptance, and those that additionally end up in accepting states; this is related to the trace semantics of CSP [13], which distinguishes deadlock from successful termination \checkmark . Such a semantics is generated by our framework as follows. As the global monad, we take $MX = \mathcal{P}(\Sigma^* \times (X + 1))$ (where we regard X and $1 = \{\checkmark\}$ as subsets of $X + 1$), with $\eta(x) = \{(\epsilon, x)\}$ and

$$f^*(S) = \{(uv, b) \mid \exists (u, x) \in S \cap (\Sigma^* \times X). (v, b) \in f(x)\} \cup (S \cap (\Sigma^* \times 1))$$

for $f : X \rightarrow MY$ and $S \in MY$. This is exactly the monad induced by the distributive law $\lambda_X : 1 + \Sigma \times \mathcal{P}(X) \rightarrow \mathcal{P}(1 + \Sigma \times X)$ with $\lambda_X(\checkmark) = \{\checkmark\}$ and $\lambda_X(a, S) = a \times S$ as used by Hasuo et al. [11]. We embed $\mathcal{P}(1 + \Sigma \times _)$ into M by the natural transformation α given by

$$\alpha_X(S) = \{(\epsilon, \checkmark) \mid \checkmark \in S\} \cup \{(a, x) \mid (a, x) \in S\}$$

(implicitly converting letters into words in the second part). Then $M1 \cong \mathcal{P}(\Sigma^*)^2$ where the first component records accepted words and the second component non-blocked words; in α -trace sequences, the first component is always contained in the second one, and increases monotonically over the sequence as the Kleisli star as defined above always keeps traces that are already accepted. Two states are α -trace equivalent iff they generate the same traces and the same accepted traces, in the sense discussed above.

Our framework also covers the language semantics of non-deterministic automata. Note that we can impose w.l.o.g. that a non-deterministic automaton never blocks an input letter – if a state fails to have an a -successor, just add an a -transition into a non-accepting state that loops on all input letters and has no transitions into other states; this clearly leaves the language of the automaton unchanged. This restriction amounts to considering coalgebras for the subfunctor

$$G = 2 \times (\mathcal{P}^*)^\Sigma$$

of the functor $\mathcal{P}(1 + \Sigma \times _)$ modelling LTS with explicit termination, where \mathcal{P}^* denotes non-empty powerset. We embed this functor into the same monad M as above, by restricting $\alpha : \mathcal{P}(1 + \Sigma \times _) \rightarrow M$ to G . Calling G -coalgebras *non-blocking non-deterministic automata*, we now have that *two states in a non-blocking non-deterministic automaton are α -trace equivalent iff they accept the same language*. For a coalgebra $\gamma : X \rightarrow GX$, the maps $\gamma^{(n)} : X \rightarrow M1$, of course, still record accepted traces as well as plain traces, but the plain traces no longer carry any information: all α -trace sequences have the form $(L_n, \Sigma^n)_{n < \omega}$ (with $L_n \subseteq \Sigma^*$ recording the accepted words of length at most n).

Probabilistic Transition Systems Recall that *generative probabilistic (transition) systems* (for simplicity without the possibility of deadlock, not to be confused with explicit termination) are modelled as coalgebras for the functor $\mathcal{D}(\Sigma \times _)$ where \mathcal{D} denotes the discrete distribution functor (i.e. $\mathcal{D}(X)$ is the set of discrete probability distributions on X , and $\mathcal{D}(f)$ takes image measures under f). That is, each state has a probability distribution over pairs of actions and successor states. We embed $\mathcal{D}(\Sigma \times _)$ into the global monad $MX = \mathcal{D}(\Sigma^* \times _)$ via the natural transformation α that takes a discrete distribution μ on $\Sigma \times X$ to the discrete distribution on $\Sigma^* \times X$ that behaves like μ on $\Sigma \times X$ (where we see Σ as a subset of Σ^*) and is 0 outside $\Sigma \times X$. The unit η of M maps $x \in X$ to the Dirac distribution at (ϵ, x) , and for $f : X \rightarrow MY$,

$$f^*(\mu)(u, y) = \sum_{u=vw, x \in X} \mu(v, x) f(x)(w, y)$$

for all $\mu \in MX$, $(u, y) \in \Sigma^* \times Y$. This is the monad induced by the canonical distributive law $\lambda : \Sigma \times \mathcal{D} \rightarrow \mathcal{D}(\Sigma \times _)$ given by $\lambda_X(a, \mu) = \delta(a) * \mu$ where δ forms Dirac measures and $*$ is product measure [11]. We identify $M1$ with $\mathcal{D}(\Sigma^*)$. Given these data, observe that for $\gamma : X \rightarrow \mathcal{D}(\Sigma \times X)$ and $x \in X$, each distribution $M!\gamma^{(n)}(x)$ is concentrated at traces of length n .

Assume from now on that Σ is finite. Recall that the usual σ -algebra on the set Σ^ω of infinite words over Σ is generated by the *cones*, i.e. the sets $v\uparrow = \{vw \mid w \in \Sigma^\omega\}$, $v \in \Sigma^*$, which (by finiteness of Σ) form a semiring of sets. We let states x in a coalgebra $\gamma : X \rightarrow \mathcal{D}(\Sigma \times X)$ induce distributions μ_x on Σ^ω via the Hahn-Kolmogorov theorem [21], defining a content $\mu(v\uparrow)$ inductively by

$$\begin{aligned} \mu_x(\epsilon\uparrow) &= 1 \\ \mu_x(av\uparrow) &= \sum_{x' \in X} \gamma(a, x') \mu_{x'}(v\uparrow) \end{aligned}$$

– a compactness argument, again hinging on the finiteness of Σ , shows that no cone can be written as a countably infinite disjoint union of cones, so μ is in fact a pre-measure, i.e. σ -additive.

We note explicitly

Proposition 8. *States in generative probabilistic systems over a finite alphabet Σ are α -trace equivalent iff they induce the same distribution on Σ^ω .*

Proof. For v a word of length n and x a state in a generative probabilistic system, we have $\mu_x(v\uparrow) = (M!\gamma^{(n)}(x))(v)$. \square

5 Relation to Other Frameworks

Kleisli Liftings Hasuo et al. [11] treat the case where the type functor G has the form TF for a monad T and a finitary endofunctor F on sets. They require that F lifts to a functor \bar{F} on the Kleisli category of T , which is equivalent to having a (functor-over-monad) distributive law

$$\lambda : FT \rightarrow TF.$$

They impose further conditions that include a cppo structure on the hom-sets¹ of the Kleisli category $\text{Kl}(T)$ of T and ensure that

- $T\emptyset$ is a singleton, so that \emptyset is a terminal object in $\text{Kl}(T)$ (unique Kleisli morphisms into \emptyset of course being \perp); and
- the final sequence of \bar{F} coincides on objects with the initial sequence of F , and converges to the final \bar{F} -coalgebra in ω steps.

The trace semantics of a TF -coalgebra is then defined as the unique Kleisli morphism into the final \bar{F} -coalgebra; in keeping with distinguishing terminology used in Section 4, we refer to this as language semantics. Thus, two states in a TF -coalgebra are *language equivalent*, i.e. trace equivalent in the sense of Hasuo et al., iff they map to the same values in the final sequence of \bar{F} under the cones induced by the respective coalgebras. Explicitly: the underlying sets of the final sequence of \bar{F} have the form $TF^n\emptyset$, $n < \omega$, and given a coalgebra $\gamma : X \rightarrow TF^2X$, the canonical cone $(\bar{\gamma}_n : X \rightarrow TF^n\emptyset)_{n < \omega}$ is defined recursively by $\gamma_0 = \perp$ and

$$\bar{\gamma}_{n+1} = X \xrightarrow{\gamma} TF^2X \xrightarrow{TF\bar{\gamma}_n} TF^2TF^n\emptyset \xrightarrow{T\lambda} TTF^{n+1}\emptyset \xrightarrow{\mu} TF^{n+1}\emptyset.$$

Now the distributive law λ induces a monad structure on the functor

$$M = TF^*, \tag{2}$$

where F^* denotes the (algebraically-)free monad on F (cf. Section 4), and we have a natural transformation $\alpha : TF \rightarrow M$, so that the situation fits our current framework. The sets TF^nX embed into MX , so that the objects in the final sequence of \bar{F} can be

¹ More precisely, hom-sets are partial orders with a least element \perp , joins of ω -chains and composition is continuous in both arguments and *left-strict*, i.e. $\perp \circ f = \perp$.

seen as living in $M0$. The definition of $\tilde{\gamma}_{n+1}$ is then seen to be just an explicit form of Kleisli composition in M ; that is, we can, for purposes of language equivalence, replace the $\tilde{\gamma}_n$ with maps $\tilde{\gamma}_n : X \rightarrow M0$ defined recursively by

$$\tilde{\gamma}_0 = \perp \quad \tilde{\gamma}_{n+1} = \tilde{\gamma}_n^* \alpha \gamma$$

where the Kleisli star is that of M . Comparing with (1), we see that the only difference with the definition of $\gamma^{(n)}$ is in the base of the recursion: $\gamma^{(0)} = \eta_X$. Noting moreover that

$$\perp^* M! \eta_X = \perp^* \eta! = \perp! = \perp,$$

we obtain

$$\tilde{\gamma}_n = \perp^* M! \gamma^{(n)}.$$

(Kissig and Kurz [16] use a very similar definition in a more general setting that in particular, for non-commutative T , does not restrict $T\emptyset$ to be a singleton, and instead assume some distinguished element $e \in T\emptyset$. They then put $\tilde{\gamma}_0 = \lambda x. e$; the comparison with our framework is then entirely analogous.)

Summing up, *language equivalence is induced from α -trace equivalence by postcomposing α -trace sequences with $\perp^* : M1 \rightarrow M0$* . Intuitively, this means that any information tied to poststates in a pretrace is erased in language equivalence, as opposed to just forgetting the poststate itself in α -trace equivalence. An example of this phenomenon are LTS with explicit termination as discussed in Section 4. Moreover, this observation elucidates why language equivalence becomes trivial in cases without explicit termination, such as standard LTS: here, all traces are tied to poststates and hence are erased when postcomposing with \perp^* . (This is also easily seen directly [11]: without explicit termination, e.g. $F = \Sigma \times _$, one typically has $F\emptyset = \emptyset$ so that the final \bar{F} -coalgebra is trivial in the Kleisli category of M .)

Eilenberg-Moore Liftings An alternative route to final objects for trace semantics was first suggested by the generalized powerset construction of Silva et al. [19] and explicitly formulated in [7] (see also Jacobs et al. [14] where this is compared to the semantics given by Kleisli liftings). In this approach one considers liftings of functors to Eilenberg-Moore categories in lieu of Kleisli categories. The setup applies to functors of the form $G = FT$ where F is an endofunctor and T is a monad on a base category \mathbf{C} . It is based on assuming a final F -coalgebra Z and a (functor-over-monad) distributive law

$$\rho : TF \rightarrow FT.$$

Under these assumptions, F lifts to an endofunctor \hat{F} on the Eilenberg-Moore category \mathbf{C}^T of T , and the free-algebra functor $\mathbf{C} \rightarrow \mathbf{C}^T$ lifts to a functor D from FT -coalgebras to \hat{F} -coalgebras, which can be seen as a generalized powerset construction (the standard powerset construction is obtained by taking $FX = 2 \times X^\Sigma$, $T = \mathcal{P}$). Explicitly, $D(\gamma) = F\mu_X^T \rho_{TX} T\gamma$ for $\gamma : X \rightarrow FTX$, where μ^T denotes the multiplication of T . In other words, $D(\gamma) : TX \rightarrow FTX$ is the unique T -algebra morphism with $D(\gamma) \cdot \eta_X^T = \gamma$. Moreover, \hat{F} has a final coalgebra with carrier Z . The *extension semantics* (i.e. trace semantics obtained via the powerset extension) of an FT -coalgebra

$\gamma : X \rightarrow FTX$ is then obtained by first applying D to γ , obtaining a \hat{F} -coalgebra with carrier TX and hence a \hat{F} -coalgebra map $TX \rightarrow Z$, and finally precomposing with $\eta_X^T : X \rightarrow TX$ where η^T denotes the unit of T . Similarly to our approach, one obtains standard bisimilarity as a special case of trace equivalence by taking $T = Id$.

In order to compare this with our framework, in which we currently consider only finite iterates of the given coalgebra, we need to assume that F -behavioural equivalence coincides with finite-depth behavioural equivalence; this is ensured e.g. by assuming that F is a finitary endofunctor on Set . In this case, two states have the same extension semantics iff they induce the same values in the first ω steps of the final sequence of \hat{F} , whose carriers coincide with the final sequence of F . Combining the definition of $D\gamma$ for a coalgebra $\gamma : X \rightarrow FTX$ with the usual construction of the canonical cone for $D\gamma$, which we denote by $\bar{\gamma}_n : TX \rightarrow F^n 1$ for distinction from the canonical cone of γ in the final sequence of FT , we obtain that $\bar{\gamma}_n$ is recursively defined by

$$\begin{aligned}\bar{\gamma}_0 &= !_{TX} : TX \rightarrow 1 \\ \bar{\gamma}_{n+1} &= F\bar{\gamma}_n T\gamma\rho F\mu^T.\end{aligned}$$

Now let us also assume that T is a finitary monad on Set . Then Set^T is a locally finitely presentable category, and since the forgetful functor to Set creates filtered colimits, we see that the lifting \hat{F} is finitary on Set^T . Hence free \hat{F} -algebras exists, which implies that we have the adjunction on the right below

$$\text{Set} \xrightarrow{\perp} \text{Set}^T \xleftarrow{\perp} \text{Alg } \hat{F},$$

and the adjunction on the left is the canonical one. We define M to be the monad of the composed adjunction; it assigns to a set X the underlying set \hat{F}^*TX of a free \hat{F} -algebra on the free T -algebra TX ; here \hat{F}^* denotes the free monad on \hat{F} (notice that this is not in general a lifting of the free monad on F to Set^T). Intuitively, M is defined by forming the disjoint union of the algebraic theories associated to T and F , respectively, and then imposing the distributive law between the operations of T and F embodied by ρ . In the following we shall denote the unit and multiplication of \hat{F}^* by $\hat{\eta}$ and $\hat{\mu}$, respectively. We also write $\hat{\varphi}_X : \hat{F}\hat{F}^*X \rightarrow \hat{F}^*X$ for the structures of the free \hat{F} -algebras and note that these yield a natural transformation $\hat{\varphi}$.

Now denote by $\hat{\kappa} : \hat{F} \rightarrow \hat{F}^*$ the universal natural transformation into the free monad; it is easy to see that $\hat{\kappa} = \hat{\varphi} \cdot \hat{F}\hat{\eta}$. Then it follows that $\alpha = \hat{\kappa}T$ yields a natural transformation from FT to M (on Set). Let us further recall that there exist canonical natural transformations $\hat{\beta}^n : \hat{F}^n \rightarrow \hat{F}^*$ defined inductively by

$$\hat{\beta}^0 = (Id \xrightarrow{\hat{\eta}} \hat{F}^*) \quad \text{and} \quad \hat{\beta}^{n+1} = (\hat{F}^{n+1} = \hat{F}\hat{F}^n \xrightarrow{\hat{F}\hat{\beta}^n} \hat{F}\hat{F}^* \xrightarrow{\hat{\varphi}} \hat{F}^*).$$

We can assume w.l.o.g. that F preserves monos (hence, so does \hat{F} since monos in Set^T are precisely injective T -algebra homomorphisms) and that coproduct injections are monic in Set^T . Then an easy induction shows that the $\hat{\beta}^n$ are monic, too. (One uses that $[\hat{\eta}, \hat{\varphi}] : Id + \hat{F}\hat{F}^* \cong \hat{F}^*$.) This implies that for testing equivalence in the extension semantics we can replace $\bar{\gamma}_n$ with

$$\hat{\gamma}_n = \hat{\beta}_1^n \cdot \bar{\gamma}_n : TX \rightarrow \hat{F}^* 1.$$

We are now ready to state the semantic comparison result:

Theorem 9. *Let F be a finitary endofunctor, and let T be a finitary monad, both on Set . Further let $\rho : TF \rightarrow FT$ be a functor-over-monad distributive law. Then two states in FT -coalgebras are equivalent under the extension semantics iff for $\alpha : FT \rightarrow M$ as given above, their α -trace sequences are identified under componentwise postcomposition with $\hat{F}^*!_{T1}$. That is, in the above notation,*

$$\hat{\gamma}_n \cdot \eta_X^T = \hat{F}^*!_{T1} \cdot M!_X \cdot \gamma^{(n)}. \quad (3)$$

Proof. We first recall how the Kleisli extension $f \mapsto f^*$ for the monad M is obtained. Given $f : X \rightarrow MY$ one first extends this to the unique T -algebra morphism $f^\sharp : TX \rightarrow MY$ with $f^\sharp \cdot \eta_X^T = f$ (i. e. one applies the Kleisli extension of T). Then one obtains $f^* : MX = \hat{F}^*TX \rightarrow \hat{F}^*TY = MY$ as the unique \hat{F} -algebra morphism with $f^* \cdot \hat{\eta}_{TX} = f^\sharp$. Notice that in this notation we have $D(\gamma) = \gamma^\sharp$ and that the inductive step of the definition on $\bar{\gamma}_n$ can be written as $\bar{\gamma}_{n+1} = \hat{F}\bar{\gamma}_n \cdot \gamma^\sharp : TX \rightarrow \hat{F}^n1$. Observe further that, since $\hat{\gamma}_n$, $\hat{F}^*!_{T1}$ and $M!$ are T -algebra homomorphisms, (3) is equivalent to

$$\hat{\gamma}_n = \hat{F}^*!_{T1} \cdot M!_X \cdot (\gamma^{(n)})^\sharp. \quad (4)$$

We now prove (3) by induction on n . For the base case $n = 0$ we have:

$$\begin{aligned} \hat{F}^*!_{T1} \cdot M!_X \cdot \gamma^{(0)} &= \hat{F}^*!_{T1} \cdot \hat{F}^*T!_X \cdot \eta_X^M && M = \hat{F}^*T \text{ and def. of } \gamma^{(0)} \\ &= \hat{F}^*!_{T1} \cdot \hat{F}^*T!_X \cdot \hat{\eta}_{TX} \cdot \eta_X^T && \text{since } \eta^M = \hat{\eta}T \cdot \eta^T \\ &= \hat{\eta}_1 \cdot !_{T1} \cdot T!_X \cdot \eta_X^T && \text{naturality of } \hat{\eta} \\ &= \hat{\eta}_1 \cdot !_{TX} \cdot \eta_X^T && \text{uniqueness of } !_{TX} \\ &= \hat{\beta}_1^0 \cdot \bar{\gamma}_0 \cdot \eta_X^T && \text{def. of } \hat{\beta}^0 \text{ and } \bar{\gamma}_0 \\ &= \hat{\gamma}_0 \cdot \eta_X^T && \text{def. of } \hat{\gamma}_0. \end{aligned}$$


For the induction step we compute:

$$\begin{aligned} \hat{F}^*!_{T1} \cdot M!_X \cdot \gamma^{(n+1)} &= \hat{F}^*!_{T1} \cdot \hat{F}^*T!_X \cdot (\gamma^{(n)})^* \cdot \alpha_X \cdot \gamma && M = \hat{F}^*T \text{ and def. of } \gamma^{(n+1)} \\ &= \hat{F}^*!_{T1} \cdot \hat{F}^*T!_X \cdot (\gamma^{(n)})^* \cdot \hat{\varphi}_X \cdot \hat{F}\hat{\eta}_{TX} \cdot \gamma && \text{def. of } \alpha \\ &= \hat{\varphi}_1 \cdot \hat{F}\hat{F}^*!_{T1} \cdot \hat{F}\hat{F}^*T!_X \cdot \hat{F}(\gamma^{(n)})^* \cdot \hat{F}\hat{\eta}_{TX} \cdot \gamma && \hat{F}\text{-algebra morphisms} \\ &= \hat{\varphi}_1 \cdot \hat{F}\hat{F}^*!_{T1} \cdot \hat{F}\hat{F}^*T!_X \cdot \hat{F}(\gamma^{(n)})^\sharp \cdot \gamma && \text{def. of } (-)^* \\ &= \hat{\varphi}_1 \cdot \hat{F}\hat{\gamma}_n \cdot \gamma && \text{induction hypothesis (4)} \\ &= \hat{\varphi}_1 \cdot \hat{F}\hat{\beta}_1^n \cdot \hat{F}\bar{\gamma}_n \cdot \gamma && \text{def. of } \hat{\gamma}^n \\ &= \hat{\beta}_1^{n+1} \cdot F\bar{\gamma}_n \cdot \gamma && \text{def. of } \beta^{n+1} \\ &= \hat{\beta}_1^{n+1} \cdot F\bar{\gamma}_n \cdot \gamma^\sharp \cdot \eta_X^T && (-)^\sharp \text{ Kleisli extension} \\ &= \hat{\beta}_1^{n+1} \cdot \bar{\gamma}_{n+1} \cdot \eta_X^T && \text{def. of } \bar{\gamma}_{n+1} \\ &= \hat{\gamma}_{n+1} \cdot \eta_X^T && \text{def. of } \hat{\gamma}_{n+1}. \quad \square \end{aligned}$$

In the base example in work on extension semantics [14, 7], the case of non-deterministic automata understood as coalgebras of the form $\gamma : X \rightarrow 2 \times \mathcal{P}(X)^\Sigma$, the situation is as follows. The extension semantics of γ [14, Section 5.1] yields a map

$tr : X \rightarrow \mathcal{P}(\Sigma^*)$ that maps each state $x \in X$ to the language accepted by the automaton with starting state x .

To understand the above theorem in terms of this concrete example, we fix $FX = 2 \times X^\Sigma$ and $TX = \mathcal{P}_{<\omega}(X)$ (to ensure finitariness). Understood as an algebraic signature, F can be represented by two Σ -ary function symbols y and n . The monad $M = \hat{F}^*T$ has these operations and those of $\mathcal{P}_{<\omega}$, i.e. the join semilattice operations, which we write using set notation; the distributive law ρ allows us to distribute joins over y and n , favouring y over n to reflect the acceptance condition of (existential) non-deterministic automata. The trace semantics $\alpha_X : FTX \rightarrow MX$ embeds flat terms, i.e. terms of the form $y((U_a)_{a \in \Sigma})$ or $n((U_a)_{a \in \Sigma}) \in FTX$ (with $U_a \in \mathcal{P}(X)$), into general (non-flat) terms. Every step in the construction of $\gamma^n(c)$ puts a flat term on top of terms constructed in the previous step, and then distributes T -operations (joins) over their arguments as indicated. Therefore, the terms $\gamma^n(c)$ are terms of uniform depth in the F -operations over sets of variables, i.e. they are elements of F^nTC . For the alphabet $\Sigma = \{0, 1\}$, a typical component of the trace sequence $T_\gamma^\alpha(c)$, i.e. $M!_X \gamma^n(c)$ for some n can be visualised as a tree like the one on the left:



This tree conveys the information that the empty word ϵ and the word 1 lead to final states (i.e. are accepted in the sense of language semantics), and additionally that 00 and 11 are not blocked; generally, the α -trace sequence records at each stage which words are accepted and additionally which words can be executed without deadlock. The tree on the right is then obtained by applying $\hat{F}^*!_{T1}$. This erases the information on non-blocked words, so that only the information that ϵ and 1 are accepted remains; this yields the extension semantics [14, 7], i.e. language semantics of the automaton, as formally stated in Theorem 9. As noted already in Section 4, if we move to non-blocking non-deterministic automata, then α -trace equivalence coincides directly with language equivalence – note that in this case, T is the non-empty powerset monad, so that $!_{T1}$ is a bijection, i.e. postcomposing the α -trace sequence with $\hat{F}^*!_{T1}$ does not lose information. Informally, this is clear as non-acceptance of words due to deadlock never happens in a non-blocking nondeterministic automaton.

Fixpoint Definitions Trace semantics, and associated linear-time logics, are also considered by Cirstea [8]. The setup is similar to that of Hasuo et al. [11] in that it applies to systems of type $X \rightarrow TFX$ where T is a monad (that describes the branching) and F a polynomial endofunctor (modelling the traces). The monad T is required to be commutative and partially additive, thus inducing a partial additive semiring structure on $T1$. In the examples of interest, one recovers the monad T as induced by this semiring structure.

Given a system $(X, f : X \rightarrow TFX)$, trace semantics then arises as a $T1$ -valued relation $R : X \times Z \rightarrow T1$ where $Z = \nu F$ is the final coalgebra of the functor F defining traces. For this to be well-defined, one additionally requires that the semiring $T1$ has suprema of chains, with order defined in the standard way.

The crucial difference to our approach is that trace semantics is defined *coinductively* on the *infinite unfolding* of the functor F defining the shape of traces, whereas our definition is *inductive* and based on *finite unfoldings*.

The difference becomes apparent when looking at examples. For labelled transition systems $X \rightarrow \mathcal{P}(A \times X)$, the trace semantics of [8] is a function $X \rightarrow \mathcal{P}(A^\omega)$ that maps x to the set of maximal traces, and two states are trace equivalent if they have the same set of *infinite* traces. This contrasts with our treatment where equivalent states have the same *finite* traces. Similarly, for generative probabilistic systems, i.e. systems of shape $X \rightarrow \mathcal{D}(A \times X)$ where \mathcal{D} is the discrete distributions functor, the trace semantics obtained in [8] associates probabilities to maximal (infinite) traces whereas our treatment is centered around probabilities of finite prefixes. In summary, the main conceptual difference between [8] and our approach is that between infinite and finite traces. Technically, this difference is manifest in the fact that our approach defines traces inductively instead of coinductively.

6 Conclusions

One of the main important aspects of the general theory of coalgebra is a uniform theory of strong bisimulation. In coalgebraic terms, strong bisimulation is a simple concept, readily defined, supports a rich theory and instantiates to the natural and known notions for concretely given transition types. Instead of re-establishing facts about strong bisimulation on a case-by-case basis, separately for each type of transition system, the coalgebraic approach provides a general theory of which specific results for concretely given systems are mere instances: a coalgebraic success story.

The question about whether a similar success story for trace equivalence can also be told in a coalgebraic setting has been the subject of numerous papers (discussed in the previous section in detail) but has so far not received a satisfactory answer.

One of the reasons why trace semantics has so far been a more elusive concept is the fact that – even for concretely given systems such as labelled transition systems with explicit termination – there are many, equally natural, formulations of trace equivalence. This suggests that trace equivalence, by its very nature, cannot be captured by one general definition, but needs an additional parameter that defines the precise nature of traces one wants to capture.

We account for this fact by parametrising trace semantics by an embedding of a functor (that defines the coalgebraic type of system under consideration) into a monad (that allows us to sequence transitions). We have shown that our definition subsumes existing notions of generic trace semantics [11, 14, 7]; it is more flexible than these in that it avoids various additional technical conditions like order enrichment or partial additivity of a monad. This has allowed us to deal with new examples not covered by previous approaches, notably systems not featuring explicit termination, such as standard labelled transition systems.

Proposition 7 shows that finite-depth strong bisimilarity is a specific instance of our parametrized definition of trace equivalence, similarly as strong bisimilarity arises as a special case of Eilenberg-Moore-style extension semantics [14, 7]. One important point to explore in future research is how further equivalences from the linear-time-branching time spectrum [23] fit into our framework, extending previous work on extension semantics [6]. Further open issues include a generalisation to behavioural preorders, as well as appropriate logics that characterise these preorders and ensuing equivalences.

On a more personal note, and reflecting on the last 15 years of development in universal coalgebra, we have learned that neither coalgebraic nor algebraic approaches by themselves are powerful enough to describe *all* aspects of state based systems, and that only the combination of algebraic theories (here in the form of monads) and universal coalgebra provide a satisfactory understanding.

Apart from providing a puzzle piece, we also hope to provide some physical motivation for the person honoured by this collection of papers to occasionally stand on his head, and while in this only seemingly anomalous position, to apply algebraic techniques from a coalgebraic viewpoint. We hope that this exercise will contribute to continued good health and many more insights into the algebraic foundations of our discipline. Happy Birthday, Martin!

Acknowledgements The authors wish to thank Bartek Klin for useful discussions, and Erwin R. Catesbeiana for repeated hints regarding inconsistent monads.

References

1. L. Aceto, A. Ingólfssdóttir, K. Larsen, and J. Srba. *Reactive systems: modelling, specification and verification*. Cambridge University Press, 2007.
2. P. Aczel. *Non-Well-Founded Sets*. CSLI, Stanford, 1988.
3. J. Adámek. Free algebras and automata realizations in the language of categories. *Comment. Math. Univ. Carolin.*, 15:589–602, 1974.
4. J. Adámek and J. Rosický. *Locally presentable and accessible categories*. Cambridge University Press, 1994.
5. M. Barr. Coequalizers and free triples. *Math. Zeitschr.*, 116:307–322, 1970.
6. F. Bonchi, M. Bonsangue, G. Caltais, J. Rutten, and A. Silva. Final semantics for decorated traces. In *Mathematical Foundations of Programming Semantics, MFPS 2012*, vol. 286 of *ENTCS*, pp. 73–86. Elsevier, 2012.
7. M. M. Bonsangue, S. Milius, and A. Silva. Sound and complete axiomatizations of coalgebraic language equivalence. *ACM Trans. Comput. Log.*, 14(1:7), 2013.
8. C. Cirstea. A coalgebraic approach to linear-time logics. In *Foundations of Software Science and Computation Structures, FoSSaCS 2014*, vol. 8412 of *LNCS*, pp. 426–440. Springer, 2014.
9. V. Giarratana, F. Gimona, and U. Montanari. Observability concepts in abstract data type specifications. In *Mathematical Foundations of Computer Science, MFCS 1976*, vol. 45 of *LNCS*, pp. 576–587. Springer, 1976.
10. J. Goguen and J. Thatcher. Initial algebra semantics. In *Switching and Automata Theory, SWAT (FOCS) 1974*, pp. 63–77. IEEE Computer Society, 1974.
11. I. Hasuo, B. Jacobs, and A. Sokolova. Generic trace semantics via coinduction. *Log. Methods Comput. Sci.*, 3, 2007.

12. R. Hennicker and M. Wirsing. Observational Specification: A Birkhoff Theorem. In *Workshop on Theory and Applications of Abstract Data Types, WADT 1985, Selected Papers*, pp. 119–135. Springer, 1985.
13. A. Hoare. *Communicating sequential processes*. Prentice Hall, 1985.
14. B. Jacobs, A. Silva, and A. Sokolova. Trace semantics via determinization. In *Coalgebraic Methods in Computer Science, CMCS 2012*, vol. 7399 of *LNCS*, pp. 109–129. Springer, 2012.
15. M. Kelly. A unified treatment of transfinite constructions for free algebras, free monoids, colimits, associated sheaves, and so on. *Bull. Austral. Math. Soc.*, 22:1–83, 1980.
16. C. Kissig and A. Kurz. Generic trace logics. arXiv preprint 1103.3239, 2011.
17. J. Power and D. Turi. A coalgebraic foundation for linear time semantics. In *Coalgebraic Methods in Computer Science, CMCS 1999*, vol. 29 of *ENTCS*, pp. 259–274. Elsevier, 1999.
18. H. Reichel. Behavioural equivalence a unifying concept for initial and final specification methods. In *Math. Models in Comp. Systems, Proc. 3rd Hungarian Comp. Sci. Conference*, pp. 27–39, 1981.
19. A. Silva, F. Bonchi, M. Bonsangue, and J. Rutten. Generalizing the powerset construction, coalgebraically. In K. Lodaya and M. Mahajan, eds., *Proc. IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, vol. 8 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 272–283, 2010.
20. A. Silva, F. Bonchi, M. M. Bonsangue, and J. J. M. M. Rutten. Generalizing determinization from automata to coalgebras. *Log. Methods Comput. Sci.*, 9(1:9), 2013.
21. T. Tao. *An Introduction to Measure Theory*. AMS, 2011.
22. D. Turi and G. Plotkin. Towards a mathematical operational semantics. In *Logic in Computer Science, LICS 1997*, pp. 280–291, 1997.
23. R. van Glabbeek. The linear time-branching time spectrum (extended abstract). In *Theories of Concurrency: Unification and Extension, CONCUR 1990*, vol. 458 of *LNCS*, pp. 278–297. Springer, 1990.
24. J. Worrell. On the final sequence of a finitary set functor. *Theoret. Comput. Sci.*, 338:184–199, 2005.