Infinite Strategies in Two-Player Games

Benjamin Plummer, Corina Cîrstea

{bjp1g19@,cc2@ecs.}soton.ac.uk

The coalgebraic study of infinite linear-time semantics was initiated by Jacobs in [4], where it was shown that linear-time semantics in labelled transition systems cannot be captured be finality in a Kleisli category, but instead by the *greatest* coalgebra morphism. This contribution describes some recent work [8], about infinite linear-time semantics in the richer structure of a *two-player game*: a transition system with two flavours of non-determinism. The semantics at a state in the game is described by the set of strategies starting from that state; we find that the infinite trace map, assigning states to semantics, can be computed as a greatest, and least, fixed point.

Our motivation for studying two-player games stems from their use in reactive synthesis. The problem of controller synthesis can be phrased as constructing a *strategy*, a function which prescribes a move in the game based on the history of the play, subject to some *correctness* and *optimality* requirements. Hence, understanding the linear-time semantics (i.e. strategies) is vital in this setting. In the future, we see that coalgebra should play a role in extending these techniques from qualitative settings to quantitative settings (e.g. when the environment in probabilistic).

To obtain a linear-time semantics, the key aspect of the coalgebraic approach is to identify a *monad* to model the branching type. The classical case of labelled transition systems uses the *powerset monad* $P : \mathbf{Set} \to \mathbf{Set}$, where the *monad multiplication* $\bigcup : PP \to P$ lets us flatten multiple steps into one. What enables us to model two-player games is identifying a suitable monad $\mathcal{G} : \mathbf{Set} \to \mathbf{Set}$ to model the branching structure of games, provided by work on alternating automata in [2] (and implicitly in [5]). The set $\mathcal{G}(X)$ consists of sets of non-empty subsets which are closed under under arbitrary non-empty union, and the monad multiplication $\mu^{\mathcal{G}} : \mathcal{G}\mathcal{G}(X) \to \mathcal{G}(X)$ maps $\alpha \in \mathcal{G}\mathcal{G}(X) \mapsto \{\bigcup_{\mathcal{U}\in\beta} U_{\mathcal{U}} \mid \exists\beta \in \alpha, \forall\mathcal{U} \in \beta : U_{\mathcal{U}} \in \mathcal{U}\}$. We see $\mu^{\mathcal{G}}$ as flattening two sequential moves in the game into a single move, as depicted on the right. We have not explicitly drawn choices which are unions of others, we see these choices as letting the $\square$ player leave choices undetermined.

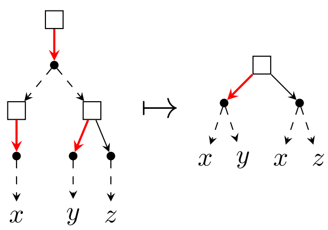

One-player Models We briefly describe the situation for one-player models (essentially from [4]), before describing our results for games. We can model a labelled transition system as a function $\alpha : X \to PF(X)$, where for simplicity we take the functor $F : \mathbf{Set} \to \mathbf{Set}$ to be $A \times (-)$. The terminal coalgebra for F is A^ω equipped with $\zeta : A^\omega \to A \times A^\omega$, where A^ω denotes the set of infinite words over A. As with every monad, there is a functor $\overline{(-)} : \mathbf{Set} \to \mathbf{Kl}(P)$ which composes a function with the unit η^P of P. We will denote $\mathbf{Kl}(P)$ maps with $\nrightarrow$, it is well-known that these are *relations* (i.e. that $\mathbf{Kl}(P) \cong \mathbf{Rel}$), and that F lifts to a functor $\overline{F}$ on $\mathbf{Rel}$. Also recall that each homset in $\mathbf{Rel}$ naturally carries a complete lattice structure.

Viewing (X, α) as an $\overline{F}$-coalgebra, we can study the complete lattice of relations $X \nrightarrow A^\omega$, equipped with a monotone function Φ which maps $(X \xrightarrow{f} A^\omega) \mapsto (X \xrightarrow{\alpha} F(X) \xrightarrow{\overline{F}(f)} F(A^\omega) \xrightarrow{\overline{\zeta^{-1}}} A^\omega)$. It is known that the *greatest fixed point* of Φ gives us a good notion of *infinite trace map* $X \nrightarrow A^\omega$, which maps a state to all of the infinite traces which can arise from said state.

Another view is presented in [6], which phrases infinite traces as a categorical limit of the n-step unfoldings. A biproduct of the investigation in [8] is showing that these approaches are equivalent.

Games Many of the same facts carry over to the two-player setting. We denote maps in $\mathbf{Kl}(\mathcal{G})$ with $\mathrel{\mapstochar\twoheadrightarrow}$, and again it is known that $F = A \times (-)$ lifts to a functor $\overline{\overline{F}} : \mathbf{Kl}(\mathcal{G}) \to \mathbf{Kl}(\mathcal{G})$. The space of maps $X \mathrel{\mapstochar\twoheadrightarrow} Y$ form a complete lattice. We model a game with an $\overline{\overline{F}}$-coalgebra $\gamma : X \mathrel{\mapstochar\twoheadrightarrow} F(X)$, and have a monotone function Φ on $\mathbf{Kl}(\mathcal{G})(X, A^\omega)$ given by $(X \xrightarrow{f} A^\omega) \mapsto (X \xrightarrow{\gamma} F(X) \xrightarrow{\overline{\overline{F}}(f)} F(A^\omega) \xrightarrow{\overline{\overline{\zeta^{-1}}}} A^\omega)$. Notice that our games are fundamentally from $\square$'s perspective, when we speak of "a strategy" we mean a strategy for $\square$.

While in the one-player case, it is clear what the infinite trace map $X \nrightarrow A^\omega$ should be, the question is more subtle for the infinite trace map $X \twoheadrightarrow A^\omega$ in games. This was tackled conclusively for finite traces in [7], by instantiating the coalgebraic approach to finite traces in [3] to games. It was shown in loc. cit. that the *finite* trace map assigns each state the set of subsets of *finite* traces which can be forced by a *finitely completing* strategy in the game. This insight leads us to define the *infinite trace semantics* at a state in the game, as the set consisting of subsets of A^ω which can be forced by some strategy from said state. One way to view this definition is that the correct generalisation of a *path* through a single-player model, is a *strategy* in a game, because both arise from of a sequence of choices. The infinite trace map will assign each state to its infinite trace semantics, in [8], we find that this map is *not* the greatest, or least, fixed point of Φ.

Despite this negative result, we can still recover the infinite trace map as a greatest fixed point, just on a smaller lattice. For the restriction, we require the notion of *limit-closure* used in the semantics of temporal logic [1]. A subset $U \subseteq A^\omega$ is *limit-closed* when for any $p \in A^\omega$, if we can the find every prefix of p in U (followed by some suffix), then we find that $p \in U$. This is a natural condition when we consider that the subset of plays which conform to a strategy is always limit-closed. We also technically work with *executions* rather than traces, by using $F = X \times A \times (-)$ and assuming that the game records state information.

Theorem 1 ([8])**.** *We can restrict* $\mathbf{Kl}(\mathcal{G})(X, (XA)^\omega)$ *and* Φ *to the set of maps* $f : X \to \mathcal{G}((XA)^\omega)$ *where each* $U \in f(x)$ *is limit-closed. The greatest fixed point in this lattice recovers the infinite trace map.*

We can even see the infinite trace map as being approximated by the least fixed point.

Theorem 2 ([8])**.** *Let* (X, γ) *be a game with no deadlocks (i.e.* $\gamma(x) \neq \emptyset$*). We can restrict* $\mathbf{Kl}(\mathcal{G})(X, (XA)^\omega)$ *and* Φ *to the collection of maps* $f : X \to \mathcal{G}((XA)^\omega)$ *where there is some* $U \in f(x)$ *which corresponds to* $\Box$ *leaving all choices undetermined on* (X, γ)*. Take the least fixed point* $f : X \to \mathcal{G}((XA)^\omega)$ *in this lattice, if we close each* $f(x)$ *under intersections of descending chains, we obtain the infinite trace map.*

Theorem 2 can be viewed as *strategy refinement*: where one starts with the maximally permissive strategy at each state, and then refines strategies in a inductive fashion (with Φ). This explains the assumption that there are no deadlocks, as this is required for there to exist a maximally permissive strategy from every state. A consequence is that we can approximate any infinite memory strategy with a sequence of finite memory strategies.

References

[1] E. Emerson. Alternative semantics for temporal logics. *Theoretical Computer Science*, 1983.

[2] A. Goy. *On the compositionality of monads via weak distributive laws.* PhD thesis, Université Paris-Saclay, Gif-sur-Yvette, France, Oct. 2021.

[3] I. Hasuo, B. Jacobs, and A. Sokolova. Generic trace semantics via coinduction. *Logical Methods in Computer Science*, Volume 3, Issue 4, Nov 2007.

[4] B. Jacobs. Trace semantics for coalgebras. *Electronic Notes in Theoretical Computer Science*, 2004. Proceedings of the Workshop on Coalgebraic Methods in Computer Science (CMCS).

[5] B. Klin and J. Rot. Coalgebraic trace semantics via forgetful logics. *Logical Methods in Computer Science*, Volume 12, Issue 4, Apr 2017.

[6] A. Kurz, S. Milius, D. Pattinson, and L. Schröder. *Simplified Coalgebraic Trace Equivalence*, pages 75–90. Springer International Publishing, Cham, 2015.

[7] B. Plummer and C. Cirstea. Traces via strategies in two-player games. *Electronic Notes in Theoretical Informatics and Computer Science*, Dec 2025.

[8] B. Plummer and C. Cîrstea. A coalgebraic approach to infinite games. In *Foundations of Software Science and Computation Structures, FoSSaCS 2026*, 2026.