Once Upon a Time in the West
Determinacy, definability, and complexity of path games*

Dietmar Berwanger[1], Erich Grädel[1], and Stephan Kreutzer[2]

[1] Aachen University
[2] University of Edinburgh

Abstract. We study determinacy, definability and complexity issues of path games on finite and infinite graphs.
Compared to the usual format of infinite games on graphs (such as Gale-Stewart games) we consider here a different variant where the players select, in each move, a path of arbitrary finite length, rather than just an edge. The outcome of a play is an infinite path, the winning condition hence is a set of infinite paths, possibly given by a formula from S1S, LTL, or first-order logic. Such games have a long tradition in descriptive set theory (in the form of Banach-Mazur games) and have recently been shown to have interesting application for planning in nondeterministic domains.
It turns out that path games behave quite differently than classical graph games. For instance, path games with Muller conditions always admit positional winning strategies which are computable in polynomial time. With any logic on infinite paths (defining a winning condition) we can associate a logic on graphs, defining the winning regions of the associated path games. We explore the relationships between these logics. For instance, the winning regions of path games with an S1S-winning condition are definable in the modal mu-calculus. Further, if the winning condition is first-order (on paths), then the winning regions are definable in monadic path logic, or, for a large class of games, even in first-order logic. As a consequence, winning regions of LTL path games are definable in CTL*.

1 Introduction

The story. Once upon a time, two players set out on an infinite ride through the west. More often than not, they had quite different ideas on where to go, but for reasons that have by now been forgotten they were forced to stay together – as long as they were both alive. They agreed on the rule that each player can determine on every second day, where the ride should go. Of course the riders could go only a finite distance every day; but a day's ride might well lead back to the location where it started in the morning, or where it started a day ago.

* This research has been partially supported by the European Community Research Training Network "Games and Automata for Synthesis and Validation" (GAMES)

Hence, one of the players began by choosing the first day's ride: he indicated a finite, non-empty path p_1 from the starting point v; on the second day his opponent selected the next stretch of way, extending p_1 to a finite path p_1q_1; then it was again the turn of the first player to extend the path to $p_1q_1p_2$ and so on. After ω days, an infinite ride is completed and it is time for payoff. There were variants of this game where, after a designated number of days, one of the players would be eliminated (these things happened in the west) and the other was to complete the game by himself for the remaining ω days (a very lonesome ride, indeed).

From descriptive set theory to planning in nondeterministic domains. Path games arise also in other contexts than the wild west. They have been studied in descriptive set theory, in the form of *Banach-Mazur games* (see [3, Chapter 6] or [4, Chapter 8.H]). In their original variant (see [5, pp. 113–117], the winning condition is a set W of real numbers; in the first move, one of the players selects an interval d_1 on the real line, then his opponent chooses an interval $d_2 \subset d_1$, then the first player selects a further refinement $d_3 \subset d_2$ and so on. The first player wins if the intersection $\bigcap_{n\in\omega} d_n$ of all intervals contains a point of W, otherwise his opponent wins. This game is essentially equivalent to a path game on the infinite binary tree $\mathcal{T}^2$ or the ω-branching tree $\mathcal{T}^\omega$. An important issue in descriptive set theory is determinacy: to characterise the winning conditions W such that one of the two players has a winning strategy for the associated game. This is closely related to topological properties of W (see Section 3).

In a quite different setting, Pistore and Vardi [7] have used path games for task planning in nondeterministic domains. In their scenario, the desired infinite behaviour is specified by formulae in linear temporal logic LTL, and it is assumed that the outcome of actions may be nondeterministic; hence a plan does not have only one possible execution path, but an execution tree. Between weak planning (some possible execution path satisfies the specification) and strong planning (all possible outcomes are consistent with the specification) there is a spectrum of intermediate cases such as strong cyclic planning: every possible partial execution of the plan can be extended to an execution reaching the desired goal. In this context, planning can be modelled by a game between a friendly player E and a hostile player A selecting the outcomes of nondeterministic actions. The game is played on the execution tree of the plan, and the question is whether the friendly player E has a strategy to ensure that the outcome (a path through the execution tree) satisfies the given LTL-specification. In contrast to the path games arising in descriptive set theory, the main interest here are path games with finite alternations between players. For instance, strong cyclic planning corresponds to a AE^ω-game where a single move by A is followed by actions of E. Also the relevant questions are quite different: Rather than determinacy (which is clear for winning conditions in LTL) algorithmic issues play the central role. Pistore and Vardi show that the planning problems in this context can be solved by automata-based methods in 2EXPTIME.

Outline of this paper. Here we consider path games in a general, abstract setting, but with emphasis on definability and complexity issues. In Section 2 we describe path games and discuss their basic structure. In Section 3 we review the classical results on determinacy of Banach-Mazur games. We then study in Section 4 path games that are positionally determined. i.e. admit winning strategies that only depend on the current position, not on the history of the play. In Section 5 we investigate definability issues. We are interested in the question how the logical complexity of defining a winning conditions (a property of infinite paths) is related to the logical complexity of defining who wins the associated game (a property of game graphs). In particular, we will see that the winner of path games with LTL winning conditions is definable in CTL*.

2 Path games and their values

Path games are a class of zero-sum infinite two-player games with complete information, where moves of players consist of selecting and extending finite paths through a graph. The players will be called Ego and Alter (in short E and A). All plays are infinite, and there is a payoff function P, defining for each play a real number. The goal of Ego is to maximise the payoff while Alter wants to minimise it.

A strategy for a player is a function, assigning to every initial segment of a play a next move. Given a strategy f for Ego and a strategy g for Alter in a game $\mathcal{G}$, we write $f\hat{\ }g$ for the unique play defined by f and g, and $P(f\hat{\ }g)$ for its payoff. The values of a game $\mathcal{G}$, from the point of view of Ego and Alter, respectively, are

$$e(\mathcal{G}) := \max_f \min_g P(f\hat{\ }g) \quad \text{and} \quad a(\mathcal{G}) := \min_g \max_f P(f\hat{\ }g).$$

A game is *determined* if $e(\mathcal{G}) = a(\mathcal{G})$. In the case of win-or-lose games, where the payoff of any play is either 0 or 1, this amounts to saying that one of the two players has a winning strategy. For two games $\mathcal{G}$ and $\mathcal{H}$ we write $\mathcal{G} \preceq \mathcal{H}$ if $e(\mathcal{G}) \le e(\mathcal{H})$ and $a(\mathcal{G}) \le a(\mathcal{H})$. Finally, $\mathcal{G} \equiv \mathcal{H}$ if $\mathcal{G} \preceq \mathcal{H}$ and $\mathcal{H} \preceq \mathcal{G}$.

Let $G = (V, F, v)$ be an arena (the west), consisting of a directed graph (V, F) without terminal nodes, a distinguished start node v, and let $P : V^\omega \to \mathbb{R}$ be a payoff function that assigns a real number to each infinite path through the graph.

We denote a move where Ego selects a finite path of length ≥ 1 by E and an ω-sequence of such moves by E^ω; for Alter we use corresponding notation A and A^ω. Hence, for any arena G and payoff function P we have the following games.

- $(EA)^\omega(G, P)$ and $(AE)^\omega(G, P)$ are the path games with infinite alternation of finite path moves.
- $(EA)^k E^\omega(G, P)$ and $A(EA)^k E^\omega(G, P)$, for arbitrary $k \in \mathbb{N}$, are the games ending with an infinite path extension by Ego.
- $(AE)^k A^\omega(G, P)$ and $E(AE)^k A^\omega(G, P)$ are the games where Alter chooses the final infinite lonesome ride.

All these games together form the collection $\mathrm{Path}(G,P)$ of *path games*. (Obviously two consecutive finite path moves by the same players correspond to a single move, so there is no need for prefixes containing EE or AA.)

It turns out that this infinite collection of games collapses to a finite lattice of just eight different games. This has been observed independently by Pistore and Vardi [7].

Theorem 1 *For every arena G and every payoff function P, we have*

$$\begin{array}{ccccc}
E^\omega(G,P) \succeq & EAE^\omega(G,P) & \succeq & AE^\omega(G,P) & \\
 & \rotatebox[origin=c]{-90}{$\succeq$} & & \rotatebox[origin=c]{-90}{$\succeq$} & \\
 & (EA)^\omega(G,P) & \succeq & (AE)^\omega(G,P) & \\
 & \rotatebox[origin=c]{-90}{$\succeq$} & & \rotatebox[origin=c]{-90}{$\succeq$} & \\
 & EA^\omega(G,P) & \succeq & AEA^\omega(G,P) & \succeq A^\omega(G,P)
\end{array}$$

Further, every path game $\mathcal{H} \in \mathrm{Path}(G,P)$ is equivalent to one of these eight games.

Proof. The comparison relations in the diagram follow by trivial arguments. We just illustrate them for one case. To show that $\mathcal{G} \succeq \mathcal{H}$ for $\mathcal{G} = EAE^\omega(G,P)$ and $\mathcal{H} = (EA)^\omega(G,P)$, consider first an optimal strategy f of Ego in $\mathcal{H}$, with $e(\mathcal{H}) = \min_g P(f\hat{\ }g)$. Ego can use this strategy also for $\mathcal{G}$: he just plays as if he would play $\mathcal{G}$, making an arbitrary move whenever it would be A's turn in $\mathcal{H}$. Any play in $\mathcal{G}$ that is consistent with this strategy, is also a play in $\mathcal{H}$ that is consistent with f, and therefore has payoff at least $e(\mathcal{H})$. Hence $e(\mathcal{G}) \geq e(\mathcal{H})$. Second, consider an optimal strategy g of Alter in $\mathcal{G}$, with $a(\mathcal{G}) = \max_f P(f\hat{\ }g)$. In $\mathcal{H} = (EA)^\omega(G,P)$, Alter answers the first move of E as prescribed by g, and moves arbitrarily in all further moves. Again, every play that can be produced against this strategy is also a play of $\mathcal{G}$ that is consistent with g, and therefore has payoff at most $a(\mathcal{G})$. Hence $a(\mathcal{G}) \geq a(\mathcal{H})$. In all other cases the arguments are analogous.

To see that any other path game over G is equivalent to one of those displayed, it suffices to show that

(1) $(EA)^k E^\omega(G,P) \equiv EAE^\omega(G,P)$, for all $k \geq 1$, and
(2) $A(EA)^k E^\omega(G,P) \equiv AE^\omega(G,P)$, for all $k \geq 0$.

By duality, we can then infer the following equivalences: $(AE)^k A^\omega(G,P) \equiv AEA^\omega(G,P)$ for $k \geq 1$ and $E(AE)^k A^\omega(G,P) \equiv EA^\omega(G,P)$ for all $k \geq 0$.

The equivalences (1) and (2) follow with similar reasoning as above. Ego can modify a strategy f for $EAE^\omega(G,P)$ to a strategy for $(EA)^k E^\omega(G,P)$. He chooses the first move according to f and makes arbitrary moves the next $k-1$ times; he then considers the entire $A(EA)^{k-1}$-sequence of moves, which were played after his first move, as one single move of A in $EAE^\omega(G,P)$ and completes the play again according to f. The resulting play of $(EA)^k E^\omega(G,P)$ is a consistent play with f in $EAE^\omega(G,P)$. Conversely a strategy of Ego for $(EA)^k E^\omega$ also works if his opponent lets Ego move for him in all moves after the first one, i.e., in the game $EAE^\omega(G,P)$. This proves that the e-values of the two games coincide. All other equalities are treated in a similar way. $\square$

The question arises whether the eight games displayed in the diagram are really different or whether they can be collapsed further. The answer depends on the game graph and the payoff function, but for each comparison $\succeq$ in the diagram we find simple cases where it is strict. Indeed, standard winning conditions $W \subseteq \{0,1\}^\omega$ (defining the payoff function $P(\pi) = 1$ if $\pi \in W$, and $P(\pi) = 0$ otherwise) show that the eight games in the diagram are distinct on appropriate game graphs. Let us consider here the completely connected graph with two nodes 0 and 1.

If the winning condition requires some initial segment ("our journey will start with a ride through the desert") then Ego wins the path games where he moves first and loses those where Alter moves first. Thus, starting conditions separate the left half of the diagram from the right one.

Games with reachability conditions ("some day, there will be the showdown") and safety conditions ("I will always wait for you") separate games in which only one player moves, i.e. with prefix E^ω or A^ω respectively, from the other ones.

A game with a Büchi condition ("again and again someone will play the harmonica") is won by Ego if he has infinite control and lost if he only has a finite number of finite moves (prefix ending with A^ω). Similarly, Co-Büchi conditions ("some day, he will ride alone toward the sunset and never come back") separate the games which are controlled by Ego from some time onwards (with prefix ending in E^ω) from the others.

3 Determinacy

From now on we consider win-or-lose games, with a winning condition given by a set of plays W. Player E wins the path game if the resulting infinite path belongs to W, otherwise Player A wins.

The topological properties of winning conditions W implying that the associated path games are determined are known from descriptive set theory. We just recall the basic topological notions and the results. In the following section, we will proceed to the issue of *positional determinacy*, i.e. to the question which path games admit winning strategies, that only depend on the current position, not on the history of the play.

Note that path games with only finite alternations between the two players are trivially determined, for whatever winning condition; hence we restrict attention to path games with prefix $(EA)^\omega$ or $(AE)^\omega$, and by duality, it suffices to consider $(EA)^\omega$. By unravelling the game graph to a tree, we can embed any game $(EA)^\omega(G,W)$ in a Banach-Mazur game, say over T^ω. The determinacy of Banach-Mazur games is closely related to the Baire property, a notion that arose from topological classifications due to René Baire.

Topology. We consider the space B^ω of infinite sequences over a set B, endowed with the topology whose basic open sets are $O(x) := x \cdot B^\omega$, for $x \in B^*$. A set $L \subseteq B^\omega$ is *open* if it is a union of set $O(x)$, i.e., if $L = W \cdot B^\omega$ for some $W \subseteq B^*$. A tree $T \subseteq B^*$ is a set of finite words that is closed under prefixes. It is easily seen that $L \subseteq B^\omega$ is closed (i.e. the complement of an open set) if L is the set of infinite branches of some tree $T \subseteq B^*$, denoted $L = [T]$. This topological space is called *Cantor space* in case $B = \{0, 1\}$, and *Baire space* in case $B = \omega$.

The class of *Borel sets* is the closure of the open sets under countable union and complementation. Borel sets form a natural hierarchy of classes $\boldsymbol{\Sigma}^0_\eta$ for $1 \leq \eta < \omega_1$, whose first levels are

$$
\begin{array}{lll}
\boldsymbol{\Sigma}^0_1 & (\text{or } G): & \text{the open sets} \\
\boldsymbol{\Pi}^0_1 & (\text{or } F): & \text{the closed sets} \\
\boldsymbol{\Sigma}^0_2 & (\text{or } F_\sigma): & \text{countable unions of closed sets} \\
\boldsymbol{\Pi}^0_2 & (\text{or } G_\delta): & \text{countable intersections of open sets}
\end{array}
$$

In general, $\boldsymbol{\Pi}^0_\eta$ contains the complements of the $\boldsymbol{\Sigma}^0_\eta$-sets, $\boldsymbol{\Sigma}^0_{\eta+1}$ is the class of countable unions of $\boldsymbol{\Pi}^0_\eta$-sets, and $\boldsymbol{\Sigma}^0_\lambda = \bigcup_{\eta<\lambda} \boldsymbol{\Sigma}^0_\eta$ for limit ordinals λ.

We recall that a set X in a topological space is *nowhere dense* if its closure does not contain a non-empty open set. A set is *meager* if it is a union of countably many nowhere dense sets and it has the *Baire property* if its symmetric difference with some open set is meager. In particular, every Borel set has the Baire property.

We are now ready to formulate the Theorem of Banach and Mazur (see e.g. [3, 4]). To keep in line with our general notation for path games we write $(EA)^\omega(T^\omega, W)$ for the Banach-Mazur game on the ω-branching tree with winning condition W.

Theorem 2 (Banach-Mazur) *(1) Player A has a winning strategy for the game $(EA)^\omega(T^\omega, W)$ if, and only if, W is meager.*

(2) Player E has a winning strategy for $(EA)^\omega(T^\omega, W)$ if, and only if, there exists a finite word $x \in \omega^*$ such that $x \cdot \omega^\omega \setminus W$ is meager (i.e., W is co-meager in some basic open set).

As a consequence, it can be shown that for any class $\Gamma \subseteq \mathcal{P}(\omega^\omega)$ that is closed under complement and under union with open sets, all games $(EA)^\omega(T^\omega, W)$ with $W \in \Gamma$ are determined if, and only if, all sets in Γ have the Baire property. Since Borel sets have the Baire property, it follows that Banach-Mazur games are determined for Borel winning conditions. (Via a coding argument, this can also been easily derived form Martin's Theorem, saying that Gale-Stewart games with Borel winning conditions are determined.)

Standard winning conditions used in applications (in particular the winning conditions that can be described in S1S) are contained in very low levels of the Borel hierarchy. Hence all path games of this form are determined.

4 Positional determinacy

In general, winning strategies may be rather complicated. However, there are interesting classes of games that are determined via relatively simple winning strategies. Of particular interest are *positional* (also called *memoryless*) strategies, that only depend on the current position, not on the history of the play. On a game graph $G = (V, F)$ a positional strategy has the form $f : V \to V^*$ assigning to every move v a finite path from v through G.

To start, we present a simple example of a path game, that is determined, but does not admit a positional strategy.

Example 3 Let G_2 be the completely connected graph with nodes 0 and 1, and let the winning condition for Ego be the set of infinite sequences with infinitely many initial segments that contain more ones than zeros. Clearly, Ego has a winning strategy for $(EA)^\omega(G, W)$, but not a positional one.

Note that this winning condition is on the Π_2-level of the Borel hierarchy. It has been pointed out by Jacques Duparc, that this is the lowest level with such an example.

Proposition 4 *If Ego has a winning strategy for a path game $(EA)^\omega(G, W)$ with $W \in \boldsymbol{\Sigma}^0_2$, then he also has a positional winning strategy.*

Proof. Let $G = (V, F)$ be the game graph. Since W is a countable union of closed sets, we have $W = \bigcup_{n<\omega}[T_n]$ where each $T_n \subseteq V^*$ is a tree. Further, let f be any (non-positional) winning strategy for Ego. We claim that, in fact, Ego can win with one move.

We construct this move by induction. Let x_1 be the initial path chosen by Ego according to f. Let $i \geq 1$ and suppose that we have already constructed a finite path $x_i \notin \bigcup_{n<i} T_n$. If $x_i y \in T_i$ for all finite y, then all infinite plays extending x_i remain in W, hence Ego wins with the initial move $w = x_i$. Otherwise choose some y_i such that $x_i y_i \notin T_i$, and suppose that Alter prolongs the play from x_i to $x_i y_i$. Let $x_{i+1} := f(x_i y_i)$ the result of the next move of Ego, according to his winning strategy f.

If this process did not terminate, then it would produce an infinite play that is consistent with f and won by Alter. Since f is a winning strategy for Ego, this is impossible. Hence there exists some $m < \omega$ such that $x_m y \in T_m$ for all y. Thus, if Ego moves to x_m in his opening move, then he wins, no matter how the play proceeds afterwards. In particular, Ego wins with a positional strategy. □

While many important winning conditions are outside $\boldsymbol{\Sigma}^0_2$, they may well be Boolean combinations of $\boldsymbol{\Sigma}^0_2$-sets.

For instance, this is the case for parity conditions, Muller conditions, and more generally, S1S-definable winning conditions. In the classical framework of infinite games on graphs (where moves are along edges rather than paths) it is well-known that parity games admit positional winning strategies, whereas there are simple games with Muller conditions that require strategies with some memory. We will see that for path games, the class of winning conditions admitting positional winning strategies is much larger than for classical graph games.

Let $G = (V, F)$ be a game graph with a colouring $\lambda : V \to C$ of the nodes with a finite number of colours. The winning condition is given by an ω-regular set $W \subseteq C^\omega$ which is defined by a formula in some appropriate logic over infinite paths. In the most general case, we have S1S-formulae (i.e. MSO-formulae on infinite paths with vocabulary $\{<\} \cup \{P_c : c \in C\}$) but we will also consider weaker formalisms like first-order logic or, equivalently, LTL.

Muller and parity conditions. As we mentioned before, typical examples of winning conditions for which strategies require memory on single-step games are Muller conditions. Such a condition is specified by a family $\mathcal{F} \subseteq 2^C$ of winning sets; a play is winning if the set of colours seen infinitely often belongs to $\mathcal{F}$.

Proposition 5 *All Muller path games $(EA)^\omega(G, \mathcal{F})$ and $(AE)^\omega(G, \mathcal{F})$ admit positional winning strategies.*

Proof. We will write $w \geq v$ to denote that position w is reachable from position v. For every position $v \in V$, let $C(v)$ be the set of colours reachable from v, that is, $C(v) := \{\lambda(w) : w \geq v\}$. Obviously, $C(w) \subseteq C(v)$ whenever $w \geq v$. In case $C(w) = C(v)$ for all $w \geq v$, we call v a *stable* position. Note that from every $u \in V$ some stable position is reachable. Further, if v is stable, then every reachable position $w \geq v$ is stable as well.

We claim that Ego has a winning strategy in $(EA)^\omega(G, \mathcal{F})$ iff there is a stable position v that is reachable from the initial position v_0, so that $C(v) \in \mathcal{F}$.

To see this, let us assume that there is such a stable position v with $C(v) \in \mathcal{F}$. Then, for every $u \geq v$, we choose a path p from u so that, when moving along p, each colour of $C(u) = C(v)$ is visited at least once, and set $f(u) := p$. In case v_0 is not reachable from v, we assign $f(v_0)$ to some path that leads from v_0 to v. Now f is a positional winning strategy for Ego in $(EA)^\omega(G, \mathcal{F})$, because, after the first move, no colours other then those in $C(v)$ are seen. Moreover, every colour in $C(v)$ is visited at each move of Ego, hence, infinitely often.

Conversely, if for every stable position v reachable from v_0 we have $C(v) \notin \mathcal{F}$, we can construct a winning strategy for Alter in a similar way. $\square$

Note that in a finite arena all positions of a strongly connected component that is terminal, i.e., with no outgoing edges, are stable. Thus, the above characterisation translates as follows: Ego wins the game iff there is a terminal component whose set of colours belongs to $\mathcal{F}$. Obviously this can be established in linear time w.r.t. the size of the arena and the description of $\mathcal{F}$.

Corollary 6 *On a finite arena G, path games with a Muller winning condition $\mathcal{F}$ can be solved in time $O(|G| \cdot |\mathcal{F}|)$.*

We remark that solving single-step graph games with Muller wining condition is PSPACE-complete.

A special case of the Muller condition is the parity condition. Given an arena $G = (V, F)$ with positions coloured by a priority function $\Omega : V \to \mathbb{N}$ of finite range, this condition requires that the least priority seen infinitely often on a play is even. It turns out that path games with parity conditions are positionally determined for any game prefix. (By Theorem 1 we can restrict attention to the eight prefixes $E^\omega, A^\omega, AE^\omega, EA^\omega$, EAE^ω, AEA^ω, $(EA)^\omega$, and $(AE)^\omega$.)

Proposition 7 *Every parity path game $\gamma(G, \text{parity})$ is determined via a positional winning strategy.*

General S1S-winning conditions. In the following, we will use parity games as an instrument to investigate path games with winning conditions specified in the monadic second-order logic of paths, S1S. It is well known that every S1S-definable class of infinite words can be recognised by a deterministic parity automaton. For words over the set of colours C, such an automaton has the form $\mathcal{A} = (Q, C, q_0, \delta, \Omega)$, where Q is a finite set of states, q_0 the initial state, $\delta : Q \times C \to Q$ a deterministic transition function, and $\Omega : Q \to \mathbb{N}$ a priority function. Given an input word, a run of $\mathcal{A}$ starts at the first word position in state q_0; if, at the current position v the automaton is in state q, it proceeds to the next position assuming the state $\delta(q, \lambda(v))$. The input is accepted if the least priority of a state occurring infinitely often in the run is even.

Via a reduction to parity games, we will first show that S1S-games admit finite-memory (or, automatic) strategies. By refining these, we will then establish strategies that are independent of the memory state, that is, positional.

Proposition 8 *For any winning condition $\psi \in$ S1S and any game prefix γ, the path games $\gamma(G, \psi)$ admit finite-memory winning strategies.*

Proof. Let $\mathcal{A} = (Q, C, q_0, \delta, \Omega)$ be an automaton that recognises the set of words defined by ψ. Given an arena $G = (V, E)$ with starting position v_0, we define the *synchronised product* $G \times \mathcal{A}$ to be the arena with positions $V \times Q$, edges from (v, q) to (v', q') whenever $(v, v') \in E$ and $\delta(q, \lambda(v)) = q'$, and designated starting position (v_0, q_0). We will use two sets of colours for $G \times \mathcal{A}$: one inherited from G, $\lambda(v, q) := \lambda(v)$, and the other one inherited from $\mathcal{A}$, $\Omega(v, q) := \Omega(q)$. When referring to a specific colouring we write, respectively, $G \times \mathcal{A}{\restriction_\lambda}$ and $G \times \mathcal{A}{\restriction_\Omega}$. Between the games on G and $G \times \mathcal{A}$ we can observe a strong relationship.

(1) For every prefix γ, a play starting from position (v_0, q_0) is winning in $\gamma(G \times \mathcal{A}{\restriction_\lambda}, \psi)$ if, and only if, it is winning in $\gamma(G \times \mathcal{A}{\restriction_\Omega}, \text{parity})$.
(2) The arenas G, v_0 and $G \times \mathcal{A}{\restriction_\lambda}, (v_0, q_0)$ are bisimilar.

The first assertion follows from the meaning of the automaton $\mathcal{A}$, and entails a strategical equivalence between the two games: Any winning strategy in $\gamma(G \times \mathcal{A}{\restriction_\Omega}, \text{parity})$ is also a winning strategy in $\gamma(G \times \mathcal{A}{\restriction_\lambda}, \psi)$ and vice versa. By Proposition 7, there always exists a positional winning strategy for the former game and, hence, for the latter one as well.

The second statement holds because $\mathcal{A}$ is deterministic. It implies that every winning strategy for a path game $\gamma(G, \psi)$ starting from position v_0 is also a winning strategy for the game $\gamma(G \times \mathcal{A}, \psi)$ starting from (v_0, q_0). Conversely, every winning strategy f for the latter game induces a winning strategy f' for the former one, namely $f'(v,s) := f\big((v,q'),s\big)$ where $q' := \delta(q_0,s)$ is the state reached by the automaton after processing the word s. Since f can be chosen to be positional, we obtain a winning strategy f' on $\gamma(G,\psi)$ that does not depend on the entire history, but only on a finite memory, namely the set of states Q. $\square$

Note that the finite-memory strategy f' constructed above does not yet need to be positional, since a position v in G has several copies (v,q) in $G \times \mathcal{A}$ at which the prescriptions of f may differ. In order to obtain a state-independent winning strategy for $\gamma(G,\psi)$ we will unify, for each node $v \in V$, the prescriptions $f(v,q)$ for those position (v,q) which are reachable in a play of according to f.

Theorem 9 *For any winning condition $\psi \in$ S1S, the games $(EA)^\omega(G,\psi)$ and $(AE)^\omega(G,\psi)$ admit positional winning strategies.*

Proof. Let us assume that Ego wins the game $(EA)^\omega(G,\psi)$ starting from position v_0. We will base our argumentation on the game $(EA)^\omega(G \times \mathcal{A},\psi)$, where Ego has a positional winning strategy f.

For any $v \in V$, we denote by $Q_f(v)$ the set of states q so that the position (v,q) can be reached from position (v_0,q_0) in a play according to f:

$$Q_f(v) := \{\delta(q_0,s) : s \text{ prolongs } f(v_0,q_0) \text{ and leads to } v\}.$$

Let $\{q_1,q_2,\dots,q_n\}$ be an enumeration of $Q_f(v)$, in which the initial state q_0 is taken first, in case it belongs to $Q_f(v)$. We construct a path associated to v along the following steps. First, set $p_1 := f(v,q_1)$; for $1 < i \le n$, let (v',q') be the node reached after playing the path $p_1 \cdot p_2 \cdot \dots \cdot p_{i-1}$ from position (v,q_i) and set $p_i := f(v',q')$. Finally, let $f'(v)$ be the concatenation of $p_1,p_2,\dots,p_n$.

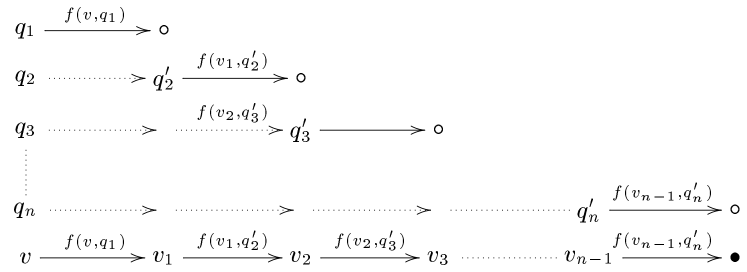

Fig. 1. Merging strategies at node v

Now, consider a play on $(EA)^\omega(G \times \mathcal{A},\psi)$ in which Ego chooses the path $f'(v)$ at any node $(v,q) \in V \times Q$. This way, the play will start with $f(q_0,v_0)$. Further, at any position (v,q) at which Ego moves, the prescription $f'(v)$ contains some segment of the form $(v', q') \cdot f(v', q')$. In other words, every move of Ego has some "good part" which would also have been produced by f at the position (v', q'). But this means that the play cannot be distinguished, post-hoc, from a play where Ego always moved according to the strategy f while all the "bad parts" were produced by Alter. Accordingly, Ego wins every play of $(EA)^\omega(G \times \mathcal{A}, \psi)$ starting from (q_0, v_0).

This proves that f' is a positional strategy for Ego in the game $(EA)^\omega(A \times G, \psi)$. Since the values do not depend on the first component, f' induces a positional strategy for Ego in $(EA)^\omega(G, \psi)$.

The same construction works for the case $(AE)^\omega(G, \psi)$, if we take instead of $Q_f(v)$ the set $Q(v) := \{\delta(q_0, s) : s \text{ is a path from } v_0 \text{ to } v\}$. □

The above proof relies upon the fact that the players always take turns. If we consider games where the players alternate only finitely many times, the situation changes. Intuitively, a winning strategy of the solitaire player eventually forms an infinite path which may not be broken apart into finite pieces to serve as a positional strategy.

Proposition 10 *For any prefix γ with finitely many alternations between the players, there are arenas G and winning conditions $\psi \in$ S1S so that no positional strategy is winning in the game $\gamma(G, \psi)$.*

Proof. Consider, for instance, the arena G_2 from Example 3 and a winning condition $\psi \in$ S1S that requires the number of zeroes occurring in a play to be odd. When starting from position 1, Ego obviously has winning strategies for each of the games $E^\omega(G, \psi)$, $AE^\omega(G, \psi)$, and $EAE^\omega(G, \psi)$, but no positional ones. □

Nevertheless, these games may be positionally determined for one of the players. For instance, if Alter wins a game $\gamma(G, \psi)$ finally controlled by Ego, he always has a positional winning strategy. This is trivial when $\gamma \in \{E^\omega, AE^\omega\}$; for $\gamma = EAE^\omega$, a positional strategy can be constructed as in the proof of Theorem 9. The reason for this asymmetry is that the lonely ride of a solitaire player may depend on an initial segment in which the opponent may interfere (or which is fixed by the winning condition, as in the case of E^ω).

Finally we consider winning conditions that are invariant to initial segments. We say that ψ is a *future*-formula, if, for any ω-word π and any finite words x and y, we have $x\pi \models \psi$ if, and only if, $y\pi \models \psi$.

Theorem 11 *For any winning condition $\psi \in$ S1S specified by a future-formula and every prefix γ, the games $\gamma(G, \psi)$ admit a positional winning strategies.*

5 Definability

We now study the question in what logics (MSO, μ-calculus, FO, CTL*, ...) winning positions of path games with ω-regular winning conditions can be defined. Given any formula φ from a logic on infinite paths (like S1S or LTL) and a quantifier prefix γ for path games, we define the game formula $\gamma.\varphi$, to be evaluated over game graphs, with the meaning that

$$G \models \gamma\,.\,\varphi \iff \text{Player } E \text{ wins the path game } \gamma(G,\varphi).$$

Note that the operation $\varphi \mapsto \gamma.\varphi$ maps a formula over infinite paths to a formula over graphs. Given a logic L over infinite paths, and a prefix γ, let $\gamma.L := \{\gamma.\varphi : \varphi \in L\}$. As usual we write $L \leq L'$ to denote that every formula in the logic L is equivalent to some formula from the logic L'.

Our main definability result can be stated as follows.

Theorem 12 *For any game prefix γ,*

(1) $\gamma\,.\,\mathrm{S1S} \leq L_\mu$
(2) $\gamma\,.\,\mathrm{LTL} \equiv \gamma\,.\,\mathrm{FO} \leq \mathrm{CTL}^*$.

Obviously, the properties expressed by formulae $\gamma.\varphi$ are invariant under bisimulation. This has two relevant consequences:

(a) We can restrict attention to trees (obtained for instance by unravelling the given game graph from the start node).
(b) It suffices to show that, on trees, $\gamma\,.\,\mathrm{S1S} \leq \mathrm{MSO}$, and $\gamma\,.\,\mathrm{FO} \leq \mathrm{MPL}$ where MPL is *monadic path logic*, i.e., monadic second-order logic where second-order quantification is restricted to infinite paths.

Indeed, it has been proved by Janin and Walukiewicz [2] that every bisimulation-invariant class of trees that is MSO-formula is also definable in the modal μ-calculus. Similarly, it is known from results by Hafer and Thomas [1] and by Moller and Rabinovitch [6], that every bisimulation invariant property of trees expressible in MPL is also expressible in CTL^*.

Proposition 13 *On trees,* $(EA)^\omega.\,\mathrm{S1S} \leq \mathrm{MSO}$ *and* $(AE)^\omega.\,\mathrm{S1S} \leq \mathrm{MSO}$.

Proof. Let $x \leq y$ denote that y is reachable from x. A strategy for Player E in a game $(EA)^\omega(T,W)$ on a tree $T = (V,F)$ is a partial function $f : V \to V$, such that $w < f(w)$ for every w; it is winning if every infinite path through T containing $f(\varepsilon), y_1, f(y_1), y_2, f(y_2)\ldots$, where $f(y_i) < y_{i+1}$ for all i, satisfies W. An equivalent description can be given in terms of the set $X = f(V)$. A set $X \subseteq V$ defines a winning strategy for Player E in the game $(EA)^\omega(T,W)$ if

(1) $(\forall x \in X)\forall y(x < y \to (\exists z \in X)(y < z))$
(2) every path hitting X infinitely often is in W (i.e. is winning for Player E)
(3) X is non-empty.

Clearly these conditions are expressible in MSO. For the game $(AE)^\omega(G,W)$ we only have to replace (3) by the condition that the start node v is contained in X. □

Proposition 14 *Let γ be a game prefix with a bounded number of alternations between E and A. Then* $\gamma\,.\,\mathrm{S1S} \leq \mathrm{MSO}$ *and* $\gamma\,.\,\mathrm{FO} \leq \mathrm{MPL}$.

Proof. Every move is represented by a path quantification; by relativizing the formula φ that defines the winning condition to the infinite path produced by the players, we obtain an MSO-formula expressing that Player E has a winning strategy for the game given by γ and φ. If φ a first-order formula over paths, then the entire formula remains in MPL. $\square$

The most interesting case concerns winning conditions defined in first-order logic (or equivalently, LTL). In our proof, we will use a normal form for first-order logic on infinite paths (with $<$) that has been established by Thomas [10]. Recall that a first-order formula $\varphi(\overline{x})$ is *bounded* if it only contains bounded quantifiers of form $(\exists y \leq x_i)$ or $(\forall y \leq x_i)$.

Proposition 15 *On infinite paths, every first-order formula is equivalent to a formula of the form*

$$\bigvee_i \Big(\exists x (\forall y \geq x) \varphi_i \wedge \forall y (\exists z \geq y) \vartheta_i \Big)$$

where φ_i and ϑ_i are bounded.

Theorem 16 *On trees, $(EA)^\omega.\mathrm{FO} \leq \mathrm{FO}$ and $(AE)^\omega.\mathrm{FO} \leq \mathrm{FO}$.*

Proof. Let $\psi = \bigvee_i \Big(\exists x (\forall y \geq x) \varphi_i \wedge \forall y (\exists z \geq y) \vartheta_i \Big)$ be a first-order formula on infinite paths describing a winning condition. We claim that, on trees, $(EA)^\omega \psi$ is equivalent to the first-order formula

$$\psi^* := (\exists p_1)(\forall p_2 \geq p_1)(\exists p_3 \geq p_2) \bigvee_{i \in I} \psi_i^{(b)} \quad \text{where}$$
$$\psi_i^{(b)} := (\exists x \leq p_1)(\forall y \,.\, x \leq y \leq p_2)\varphi_i \wedge (\forall y \leq p_2)(\exists z \,.\, y \leq z \leq p_3)\vartheta_i.$$

Let $T = (V, E)$ and suppose first that Alter has a winning strategy for the game $(EA)^\omega(T, \psi)$. We prove that $T \models \neg\psi^*$. To see this we have to define an appropriate Skolem function $g : p_1 \mapsto p_2$ such that for all $p_3 \geq p_2$ and all $i \in I$

$$T \models \neg\psi_i^{(b)}(p_1, p_2, p_3).$$

Fix any p_1 which we can consider as the first move of Ego in the game $(EA)^\omega(T, \psi)$ and any play P (i.e., any infinite path through T) that prolongs this move and that is consistent with Alter's winning strategy. Since Alter wins, we have that $P \models \neg\psi$. Hence there exists some $J \subseteq I$ such that

$$P \models \bigwedge_{i \in J} \forall x (\exists y \geq x) \neg\varphi_i \wedge \bigwedge_{i \in I - J} \exists y (\forall z \geq y) \neg\vartheta_i.$$

To put it differently, there exist

- for every $i \in J$ and every $a \in P$ a witness $h_i(a) \in P$ such that $P \models \neg\varphi_i(a, h_i(a))$, and

– for every $i \in I - J$ an element b_i such that $P \models (\forall z \geq b_i)\neg\vartheta_i(b_i, z)$.

Now set

$$p_2 := \max(\{h_i(a) : a \leq p_1, i \in J\} \cup \{b_i : i \in I - J\}).$$

For any p_3 we now obviously have that $T \models \neg\psi_i^{(b)}(p_1, p_2, p_3)$.

For the converse, let $f : V \to V$ be a winning strategy for Ego in the game $(EA)^\omega(T, \psi)$. We claim that $T \models \psi^*$. Toward a contradiction, suppose that $T \models \neg\psi^*$. Hence there exists a Skolem function $g : V \to V$ assigning to each p_1 an appropriate $p_2 \geq p_1$ such that $T \models \neg\psi_i^{(b)}(p_1, p_2, p_3)$ for all $p_3 \geq p_2$ and all $i \in I$. We can view g as a strategy for Alter in the game $(EA)^\omega(T, \psi)$. If Ego plays according to f and Alter plays according to g, then the resulting infinite play $f\hat{\ }g = q_1q_2q_3\ldots$ satisfies ψ (because f is a winning strategy). Hence there exists some $i \in I$ such that

$$f\hat{\ }g \models \exists x(\forall y \geq x)\varphi_i \wedge \forall y(\exists z \geq y)\vartheta_i.$$

Let a be a witness for x so that $f\hat{\ }g \models (\forall y \geq a)\varphi_i(a, y)$. Choose the minimal odd k, such that $a \leq q_k$, and set $p_1 := q_k$. Then $q_{k+1} = g(q_k) = g(p_1) = p_2$. Since $f\hat{\ }g \models \forall y(\exists z \geq y)\vartheta_i(y, z)$, we have, in particular, for every $b \leq p_2$ a witness $h(b) \geq b$ on $f\hat{\ }g$ such that $f\hat{\ }g \models \vartheta_i(b, h(b))$. Choose $p_3 = \max\{h(b) : b \leq p_2\}$ It follows that $f\hat{\ }g \models \psi_i^{(b)}(p_1, p_2, p_3)$. Since $\psi_i^{(b)}$ is bounded, its evaluation on T is equivalent to its evaluation on $f\hat{\ }g$. Hence we have shown that there exists p_1 such that for $p_2 = g(p_1)$, given by the Skolem function g, we can find a p_3 with $T \models \psi_i^{(b)}(p_1, p_2, p_3)$. But this contradicts the assumption that g is an appropriate Skolem function for $\neg\psi^*$.

We have shown that whenever Ego has a winning strategy for $(EA)^\omega(T, \psi)$ then $T \models \psi^*$ and whenever Alter has a winning strategy, then $T \models \neg\psi^*$. By contraposition and determinacy, the reverse implications also hold. For games of form $(AE)^\omega(T, \psi)$ the arguments are analogous. $\square$

Theorem 12 is implied by Propositions 13, Proposition 14, and Theorem 16.

Acknowledgement. We are grateful to Jacques Duparc for information on descriptive set theory and for pointing out Proposition 4 to us.

References

[1] T. Hafer and W. Thomas, *Computation tree logic CTL* and path quantifiers in the monadic theory of the binary tree*, in Automata, Languages and Programming, 14th International Colloquium, ICALP87, Lecture Notes in Computer Science Nr. 267, Springer, 1987, pp. 269–279.

[2] D. Janin and I. Walukiewicz, *On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic*, in Proceedings of 7th International Conference on Concurrency Theory CONCUR '96, no. 1119 in Lecture Notes in Computer Science, Springer-Verlag, 1996, pp. 263–277.

[3] A. KANAMORI, *The Higher Infinite*, Springer, 1991.
[4] A. KECHRIS, *Classical Descriptive Set Theory*, Springer, 1995.
[5] R. MAULDIN, ed., *The Scottish Book. Mathematics from the Scottish Café*, Birkhäuser, 1981.
[6] F. MOLLER AND A. RABINOVITCH, *Counting on CTL*: On the expressive power of monadic path logic*, Information and Computation, (2003). To appear.
[7] M. PISTORE AND M. VARDI, *The planning spectrum — one, two, three, infinity*, in Proc. 18th IEEE Symp. on Logic in Computer Science, 2003.
[8] A. PNUELI AND R. ROSNER, *On the synthesis of a reactive module*, in Proceeding of The Sixteenth Annual ACM Symposium on The Principles of Programming Languages, 1988, pp. 179–190.
[9] A. SISTLA AND E. CLARKE, *The complexity of propositional linear temporal logics*, Journal of the ACM, 32 (1985), pp. 733–749.
[10] W. THOMAS, *A combinatorial approach to the theory of omega-automata*, Information and Control, 48 (1981), pp. 261–283.