

On the Complexity of Branching Games with Regular Conditions*

Marcin Przybyłko¹ and Michał Skrzypczak²

- 1 University of Warsaw, University of New Caledonia
M.Przybylko@mimuw.edu.pl
- 2 University of Warsaw
M.Skrzypczak@mimuw.edu.pl

Abstract

Infinite duration games with regular conditions are one of the crucial tools in the areas of verification and synthesis. In this paper we consider a branching variant of such games – the game contains branching vertices that split the play into two independent sub-games. Thus, a play has the form of an infinite tree. The winner of the play is determined by a winning condition specified as a set of infinite trees. Games of this kind were used by Mio to provide a game semantics for the probabilistic μ -calculus. He used winning conditions defined in terms of parity games on trees. In this work we consider a more general class of winning conditions, namely those definable by finite automata on infinite trees. Our games can be seen as a branching-time variant of the stochastic games on graphs.

We address the question of determinacy of a branching game and the problem of computing the optimal game value for each of the players. We consider both the stochastic and non-stochastic variants of the games. The questions under consideration are parametrised by the family of strategies we allow: either mixed, behavioural, or pure.

We prove that in general, branching games are not determined under mixed strategies. This holds even for topologically simple winning conditions (differences of two open sets) and non-stochastic arenas. Nevertheless, we show that the games become determined under mixed strategies if we restrict the winning conditions to open sets of trees. We prove that the problem of comparing the game value to a rational threshold is undecidable for branching games with regular conditions in all non-trivial stochastic cases. In the non-stochastic cases we provide exact bounds on the complexity of the problem. The only case left open is the 0-player stochastic case, i.e. the problem of computing the measure of a given regular language of infinite trees.

1998 ACM Subject Classification F.4.3 Formal Languages

Keywords and phrases stochastic games, meta-parity games, infinite trees, regular languages, parity automata

Digital Object Identifier 10.4230/LIPIcs.MFCS.2016.78

1 Introduction

Since the seminal works of Büchi and Landweber [3], and of McNaughton [17], the infinite duration games are widely used to model interaction between a system and an environment. One of the fundamental questions about such games is the question of determinacy – does always one of the players has a winning strategy? In a more general case of valued zero-sum

* The authors were supported by the Polish National Science Centre grant no. 2014-13/B/ST6/03595.



© Marcin Przybyłko and Michał Skrzypczak;
licensed under Creative Commons License CC-BY

41st International Symposium on Mathematical Foundations of Computer Science (MFCS 2016).

Editors: Piotr Faliszewski, Anca Muscholl, and Rolf Niedermeier; Article No. 78; pp. 78:1–78:14

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

games, determinacy amounts to the equality

$$\sup_{\sigma} \inf_{\pi} \text{val}(\sigma, \pi) = \inf_{\pi} \sup_{\sigma} \text{val}(\sigma, \pi), \quad (1)$$

where σ and π range over strategies of the respective players. It is often crucial to provide a specific information about the strategies that are enough to win a given game. Büchi and Landweber proved that if the winning condition of a game is a regular language of infinite words then the game is determined under finite memory strategies. Further results have established more precise bounds for the amount of memory needed [9, 10]. Also, the stochastic variant of the question was considered [7].

In this work we study a branching variant of stochastic games on graphs – a variant called *branching games* (also known as *tree games*, c.f. [20, Chapter 4]), played on *branching boards*. A play of a branching game consists of a number of threads, each thread develops independently. When a thread reaches a vertex marked as *branching*, it is split into two separate threads. Thus, a play of a branching game has the shape of an infinite tree. The winner of the play is determined by a winning condition specified as a regular set of infinite trees. Since the choices made by players in separate threads are independent, branching games are not games of perfect information, nor of perfect recall in the meaning of [2]. Games of this kind were used by Mio [21] to provide a game semantics for the probabilistic μ -calculus, with the *meta-parity* winning condition defined in terms of parity games on trees. The author called them *meta-parity games* and established their pure determinacy subject to some set-theoretic assumptions, that were later eliminated [13]. This result is interesting, because, as the author notices in conclusion, determinacy results for imperfect information games are not frequent in game theory.

In this article we address the question of determinacy of branching games and the problem of computing the game value for a more general class of winning conditions than the meta-parity conditions studied by Mio, namely those definable by finite automata on infinite trees. We believe that this extension is motivated by the role tree automata play in verification theory. Recall that these automata, introduced by Rabin in his proof of decidability of the Monadic Second-Order theory of the full k -ary tree [23], constitute a general formalism subsuming most of temporal logics of programs. The first step was made by the first author who extended the results of Mio [21] to winning conditions given by game automata [22].

We consider both the stochastic and non-stochastic variants of the games. The questions under consideration are parametrised by the family of strategies we allow, either mixed, behavioural, or pure. The goal of this work is to provide answers to two questions:

- When a branching game is determined in the sense of (1)?
- When the optimal value for a given player can be effectively computed?

Both questions can be asked for stochastic and non-stochastic variants of games, which usually yields different answers, see e.g. [7]; and for different sets of allowed strategies. The distinction between the sets of pure, behavioural, and mixed strategies can significantly alter the techniques and expected outcomes, see e.g. [5].

The answers we provide create an almost complete picture from the point of view of topological complexity of sets:

- non-stochastic branching games are not determined under pure nor behavioural strategies even for winning conditions that are topologically both closed and open,
- non-stochastic branching games are not determined under mixed strategies for winning conditions that are a difference of two open sets,
- non-stochastic branching games are determined under mixed strategies for winning conditions that are open (equivalently, closed) sets,

- the problem of comparing the value of a branching game to a rational number is undecidable in all the non-trivial stochastic cases,
- in the non-stochastic case, when we ask about the existence of a pure winning strategy, the problem is decidable and we provide precise bounds on its complexity.

The only remaining question is whether the value of a 0-player stochastic branching game can be effectively computed. This is equivalent to asking about computability of the measure of a given regular language of infinite trees.

Although the results of this paper show intractability of branching games, it is still possible that they are determined for a reasonable class of winning conditions. These ideas are discussed in Conclusions.

1.1 Related work

It is known that games with arbitrary pay-off functions are not determined. The celebrated result of Martin [15] states that *Gale-Steward games* with Borel payoffs are determined under pure strategies. His later result establishes an analogous result for the so-called *Blackwell games* and mixed strategies (cf. [16]). In this work we show that branching games are determined only for topologically simplest winning conditions – open or closed sets; even allowing a difference of two open sets leads to indeterminacy.

Since the branching games are not games of *perfect recall* the Kuhn's theorem (cf. e.g. [2]) does not hold; and the behavioural strategies are weaker than the mixed strategies. An example of such a situation was provided by Mio in [20, Chapter 4]. Therefore, there are three variants of the question of determinacy: pure, behavioural, and mixed determinacy.

The concept of branching games is a natural extension of the *meta-parity games* introduced by Mio [21] to provide a game semantics for the probabilistic μ -calculus. The first author proved in [22] that non-stochastic branching games with winning conditions given by the so-called *game automata* are determined under pure strategies. In this work we study determinacy of branching games from the perspective of the topological complexity of the winning condition.

Recently Asarin et al. considered the so-called *entropy games*, cf. [1], which can be easily embedded in our framework. In the authors' own words, an entropy game is played on a finite arena by two-and-a-half players: Despot, Tribune, and the non-deterministic People. The pay-off function is the entropy of the language formed by paths of the resulting tree. The authors of [1] prove that the entropy games are determined under pure strategies and can be solved in $NP \cap coNP$, extending the class of objectives for which branching games are determined.

The question of computing the coin-flipping measure of a given regular languages of infinite trees is one of the crucial open problems about probabilistic logics on infinite trees. Chen, Dräger, and Kiefer proved in [8] that the problem is decidable for regular languages recognisable by deterministic automata. The result was later strengthened by Michalewski and Mio [19] to the so-called game automata. The question of computing the value of a branching game is a natural extension of the above problem obtained by allowing interplay between the players. The first author implicitly provided bounds on the complexity of the problem in the non-stochastic case. In this paper we complete those bounds by proving 2 -EXP-hardness of the problem. Additionally, we prove that the problem becomes undecidable if any form of randomisation is allowed (either by considering randomised or behavioural strategies; or by adding stochastic positions). The only remaining open question is the original one – when there are only stochastic positions and no players.

Branching games fall into a general category of games of *imperfect information*, i.e. games where the full information about the state of the game is not assumed: the definitions assure that the players have no information about the execution of separate threads. The area of imperfect information games is rich and not fully understood, see e.g. [7, 4, 6], see also [5]. In this context, branching games with regular objectives can be seen as a natural extension of *imperfect information games with ω -regular objectives* to the branching-time case.

2 Definitions

In this section we will define the objects studied in the paper. The crucial definitions are those of a *branching game* and *game values*. By ω we denote the set of natural numbers and \mathbb{R} stands for the set of reals.

Words and trees. An *alphabet* Γ is a finite non-empty set. A *word* over Γ is any, possibly infinite, sequence $w = a_0 a_1 \cdots a_n \cdots$ where $a_i \in \Gamma$. By $w[i]$ we denote the i -th letter of w , i.e. a_i . ε stands for the empty sequence. Words are either *finite* (Γ^*) or *infinite* (Γ^ω). $|w|$ is the length of a finite word w . The prefix order on words is denoted \sqsubseteq .

A *tree* over an alphabet Γ is any partial function $t: \{\mathbb{L}, \mathbb{R}\}^* \rightarrow \Gamma$ with a non-empty prefix-closed domain $\text{Dom}(t) \subseteq \{\mathbb{L}, \mathbb{R}\}^*$. The elements $d \in \{\mathbb{L}, \mathbb{R}\}$ are called *directions*, \bar{d} is the direction opposite to d . Elements of the set $\{\mathbb{L}, \mathbb{R}\}^*$ are called *nodes*. We say that a node u of a tree t is *fully branching* if it has two children in the tree, is *uniquely branching* if it has exactly one child in the tree. The set \mathcal{T}_Γ is the set of all trees over an alphabet Γ . This set can naturally be enhanced with a topology in such a way that it becomes a homeomorphic copy of the Cantor set [25]. We say that a tree t_1 is a *prefix* of a tree t_2 if $t_1 \subseteq t_2$, i.e. $\text{Dom}(t_1) \subseteq \text{Dom}(t_2)$ and for every $u \in \text{Dom}(t_1)$ we have $t_1(u) = t_2(u)$.

Regular languages. In this work we use the standard notions of non-deterministic and alternating parity automata over infinite trees. Together with Monadic Second-Order logic, these automata form equivalent formalisms for defining *regular* languages of infinite trees. For an introduction to this area see for instance [24].

Branching games. This paper is about branching games. The two adversaries of our games are called Eve and Adam (or shortly E and A). Since we consider stochastic games, we additionally introduce *Nature* denoted \mathcal{N} . A *branching board* is a tuple $\mathbf{B} = \langle V, \Gamma, s_L, s_R, \rho, \eta, \lambda, v_I \rangle$, where V is the set of *vertices*; Γ is the *alphabet*; $s_L, s_R: V \rightarrow V$ are the *successor functions*; $\lambda: V \rightarrow \Gamma$ is the *labelling* of the vertices; $\rho: V \rightarrow \{A, E, \mathcal{N}, \mathcal{B}\}$ is a partition of the vertices between Adam's, Eve's, *Nature's*, and branching vertices; $\eta: \rho^{-1}(\{\mathcal{N}\}) \rightarrow \text{Dist}(\{\mathbb{L}, \mathbb{R}\})$ maps *Nature's* vertices to random distributions over the successors; $v_I \in V$ is the *initial vertex*. We extend the assignment s to arbitrary sequences of directions in the natural way: $s_\varepsilon(v) = v$ and $s_{u.d}(v) = s_d(s_u(v))$.

For $P \in \{A, E, \mathcal{N}, \mathcal{B}\}$, by V_P we denote the set of vertices belonging to P , i.e. $\rho^{-1}(\{P\})$. We say that \mathbf{B} is *finitary* if the set of vertices V is finite and the values used to define η are rational. For $\mathcal{P} \subseteq \{A, E, \mathcal{N}, \mathcal{B}\}$ we say that \mathbf{B} is \mathcal{P} -*branching* if $\text{Range}(\rho) \subseteq \mathcal{P}$. Every board \mathbf{B} defines the tree $t_{\mathbf{B}}^\lambda: \{\mathbb{L}, \mathbb{R}\}^* \rightarrow \Gamma$ as the unfolding of the adequate labelled sub-graph of the board, i.e. $t_{\mathbf{B}}^\lambda(u) = \lambda(s_u(v_I))$. \mathbf{B} is *non-stochastic* if it is $\{E, A, \mathcal{B}\}$ -branching.

Intuitively, a play over a branching board \mathbf{B} proceeds in threads, each thread has one token located in a vertex of the board. Initially, there is one thread with the token located in v_I . Consider a thread with a token located in a vertex v . If $\rho(v) = \mathcal{B}$ then the thread is

duplicated into two separate threads with tokens located in $s_L(v)$ and $s_R(v)$. If $\rho(v) = \mathcal{N}$ then the token is moved either to $s_L(v)$ or to $s_R(v)$ depending on an independent random event with distribution $\eta(v)$. If $\rho(v) \in \{E, A\}$ then the respective player can make her/his choice depending on the history of the current thread. However, she/he cannot take into account positions of tokens from other threads in the current play. After all the threads moved infinitely many times, a tree-like play has been created. The winning condition of a branching game will indicate which plays are winning for which player. Figure 1 depicts a branching board and a play on this board.

We will now formalise the notions of a play and a pure strategy of a player. Consider a non-empty set $\mathcal{P} \subseteq \{A, E, \mathcal{N}, \mathcal{B}\}$. We say that a tree $t \subseteq t_B^\lambda$ is \mathcal{P} -branching if it is fully branching in the nodes $u \in \{L, R\}^*$ such that $\rho(s_u(v_I)) \in \mathcal{P}$ and uniquely branching in the remaining nodes. A *play* on a board B is a tree $t \subseteq t_B^\lambda$ that is $\{\mathcal{B}\}$ -branching. The set of all plays on a board B is denoted $\text{plays}(B)$. For $P \in \{E, A, \mathcal{N}\}$ we say that a tree $t \subseteq t_B^\lambda$ is a *pure strategy* of P over B if t is $(\{E, A, \mathcal{N}, \mathcal{B}\} \setminus \{P\})$ -branching. The set of pure strategies of P over B is denoted Σ_B^P . Notice that the sets $\text{plays}(B)$ and Σ_B^P for $P \in \{E, A, \mathcal{N}\}$ are closed sets of Γ -labelled trees. If V is finite then all these sets are regular.

Given three pure strategies $\sigma \in \Sigma_B^E$, $\pi \in \Sigma_B^A$, and $\eta \in \Sigma_B^\mathcal{N}$ the *play resulting from* σ , π , and η (denoted $\text{eval}_B(\sigma, \pi, \eta)$) is the tree $\sigma \cap \pi \cap \eta \in \text{plays}(B)$. Thus, $\text{eval}_B: \Sigma_B^E \times \Sigma_B^A \times \Sigma_B^\mathcal{N} \rightarrow \text{plays}(\Gamma)$. Notice that the function eval_B is continuous.

Measure theory. For an introduction to measure theory we refer to [14, Chapter 17]. Measure properties of regular sets of trees are discussed in [13]. Let μ be a Borel measure on a topological space X . We say that μ is a *probability measure* if $\mu(X) = 1$. A function $f: X \rightarrow \mathbb{R}$ is μ -measurable if the pre-image of any measurable set in \mathbb{R} is μ -measurable in X . $f: X \rightarrow \mathbb{R}$ is *universally measurable* if it is μ -measurable for every Borel measure μ on X . If $f: X \rightarrow \mathbb{R}$ is μ -measurable then by $\int_X f(x) \mu(dx)$ we denote the integral of f with respect to the measure μ .


Branching games. A *branching game* is a pair $G = \langle B, \Phi \rangle$ where B is a branching board and Φ is a universally measurable bounded real function $\Phi: \text{plays}(B) \rightarrow \mathbb{R}^+$. The notions of a \mathcal{P} -branching game and a *finitary* game refer to the respective properties of the board.

Mixed strategies. A mixed strategy of a player $P \in \{E, A\}$ is a Borel probability measure over the set Σ_B^P . The set of all mixed strategies of P is denoted by Σ_B^{MP} .

There is a natural way of defining a Borel probability measure η_B^* on the set $\Sigma_B^\mathcal{N}$ of strategies of \mathcal{N} . This measure represents the intuition, that after a sequence of directions $u \in \{L, R\}^*$ corresponding to a vertex $v = s_u(v_I) \in V$ such that $\rho(v) = \mathcal{N}$, *Nature* chooses to move to a direction $d \in \{L, R\}$ with the probability $\eta(v)(d)$ and is called *behavioural*.

Behavioural strategies. We say that a mixed strategy of P is *behavioural* if it is a “coin flipping” measure, i.e. a measure induced by supplying some of the nodes of t_B^λ corresponding to vertices of P with a probability distribution over the successors. To produce a pure strategy from a behavioural one, the directions are chosen independently according to the fixed probability distributions.

More formally, a mixed strategy τ of P is *behavioural* if it is, as a measure over Σ_B^P , the measure $\eta_{B'}^*$ for some (possibly not finitary) board B' . The set of all behavioural strategies of P is denoted by Σ_B^{BP} . Clearly we can treat every pure strategy in Σ_B^P as a Dirac delta function in Σ_B^{MP} (in fact in Σ_B^{BP}). Thus, we can assume that $\Sigma_B^P \subseteq \Sigma_B^{BP} \subseteq \Sigma_B^{MP}$.



■ **Figure 1** An example of a branching board and a play on this board. We denote Eve's, Adam's, Nature's, and branching vertices by diamonds, squares, circles, and triangles respectively. Nature's vertices are equipped with a probability distribution over the successors. The successors L and R agree with the directions on the picture, i.e. L moves to the left.

Strategies as functions. There is a different way to define the three types of strategies that may give more intuition to the behaviour and expressive power of the strategies. A pure strategy $\sigma \in \Sigma^P$ can be seen as a function $\sigma: \{\mathsf{L}, \mathsf{R}\}^* \rightarrow \{\mathsf{L}, \mathsf{R}\}$; a behavioural strategy $\sigma_b \in \Sigma^{BP}$ as a function $\sigma_b: \{\mathsf{L}, \mathsf{R}\}^* \rightarrow \mu(\{\mathsf{L}, \mathsf{R}\})$; and a mixed strategy $\sigma_m \in \Sigma^{MP}$ as a measure $\sigma_m \in \mu(\Sigma^P)$, where $\mu(X)$ denotes some Borel probability measure on the set X .

An example. Figure 1 depicts a branching board \mathbf{B} and a play t on this board. We identify the vertices with their labels. A pure strategy of Adam can make different choices in a depending on the history of the thread that lead to this vertex (there are infinitely many such histories). A pure strategy of Eve can make different choices in e_2 depending on the edge taken by Nature in n . A mixed strategy of Eve can synchronise: with probability $\frac{1}{2}$ move to L in both vertices e_1, e_2 ; and with probability $\frac{1}{2}$ move to R in both of them. A behavioural strategy cannot make such a synchronisation: the probability distribution over the successors depends only on the history of the current thread.

Values of strategies. Assume that $\sigma_m \in \Sigma_B^{ME}$ and $\pi_m \in \Sigma_B^{MA}$ are two mixed strategies of the respective players. Our aim is to define the value $\text{val}_G(\sigma_m, \pi_m)$. Intuitively, $\text{val}_G(\sigma_m, \pi_m)$ should be the expected value of $\Phi(\text{eval}_G(\sigma, \pi, \eta))$ where the pure strategies σ, π , and η are chosen according to the probability distributions σ_m, π_m , and η_B^* respectively. This is formalised as follows.

$$\text{val}_G(\sigma_m, \pi_m) \stackrel{\text{def}}{=} \int_{\Sigma_B^E, \Sigma_B^A, \Sigma_B^N} \Phi(\text{eval}_G(\sigma, \pi, \eta)) \sigma_m(d\sigma) \pi_m(d\pi) \eta_B^*(d\eta) \quad (2)$$

If σ and π are pure strategies and the board is non-stochastic then $\text{val}_G(\sigma, \pi) = \Phi(\pi \cap \sigma)$.

Values of a game. The aim of Eve in a branching game is to maximise the value $\text{val}_G(\sigma, \pi)$. Let us define the *partial values* of the game. Consider $X \in \{\varepsilon, B, M\}$ (i.e. X stands for respectively *pure*, *behavioural*, and *mixed* strategies). The X value of G for Eve (resp. Adam) is defined as

$$\begin{aligned} \text{val}_G^{XE} &\stackrel{\text{def}}{=} \sup_{\sigma \in \Sigma_B^{XE}} \text{val}_G(\sigma) \quad \text{where} \quad \text{val}_G(\sigma) \stackrel{\text{def}}{=} \inf_{\pi \in \Sigma_B^A} \text{val}_G(\sigma, \pi), \\ \text{val}_G^{XA} &\stackrel{\text{def}}{=} \inf_{\pi \in \Sigma_B^{XA}} \text{val}_G(\pi) \quad \text{where} \quad \text{val}_G(\pi) \stackrel{\text{def}}{=} \sup_{\sigma \in \Sigma_B^E} \text{val}_G(\sigma, \pi). \end{aligned}$$

Notice, that the second inf/sup is taken over the pure strategies of the opponent. This is explained by the following simple lemma.

► **Lemma 1.** *Let G be a branching game. If σ_m is Eve's mixed strategy then*

$$\inf_{\pi_m \in \Sigma_B^{MA}} \text{val}_G(\sigma_m, \pi_m) = \inf_{\pi_b \in \Sigma_B^{BA}} \text{val}_G(\sigma_m, \pi_b) = \inf_{\pi \in \Sigma_B^A} \text{val}_G(\sigma_m, \pi)$$

Dually, the same holds for mixed strategies of Adam if we replace inf with sup and A with E .

Determinacy. As a simple consequence of Lemma 1 we obtain the following inequalities

$$\text{val}_G^A \geq \text{val}_G^{BA} \geq \text{val}_G^{MA} \geq \text{val}_G^{ME} \geq \text{val}_G^{BE} \geq \text{val}_G^E. \quad (3)$$

The first two (resp. the last two) inequalities hold by the fact that we take inf (resp. sup) over greater (reps. smaller) sets of strategies. The third inequality holds by Lemma 1 and the fact that $\inf_x \sup_y f(x, y) \geq \sup_y \inf_x f(x, y)$.

We will say that a branching game G is *determined*

- *under pure strategies* if $\text{val}_G^A = \text{val}_G^E$,
- *under behavioural strategies* if $\text{val}_G^{BA} = \text{val}_G^{BE}$,
- *under mixed strategies* if $\text{val}_G^{MA} = \text{val}_G^{ME}$.

Clearly, Equation (3) shows that pure determinacy implies behavioural determinacy and behavioural determinacy implies mixed determinacy. In general, the opposite implications do not hold. The questions of determinacy of branching games are discussed in Section 3.

Regular branching games. The following theorem implies that we can take as Φ an indicator of a regular language of trees $L \subseteq \text{plays}(\mathcal{B})$, i.e. $\Phi(t) = 1$ if $t \in L$ and $\Phi(t) = 0$ otherwise. In that case we say that a game G has L as a winning condition and we write $G = \langle \mathcal{B}, L \rangle$ instead of $G = \langle \mathcal{B}, \Phi \rangle$.

► **Theorem 2** (Michalewski et al. [13]). *Every regular language L of infinite trees is universally measurable, i.e. for every Borel measure μ on the set of trees, we know that L is μ -measurable.*

3 Determinacy


In this section we study determinacy of branching games in the three variants: pure, behavioural, and mixed; see (3). We will show that for general regular winning conditions all three variants fail. However, when we restrict to closed regular winning sets we can recover the mixed determinacy.

Notice that if a branching game is not branching, i.e. it is a $\{E, A, \mathcal{N}\}$ -branching game then the determinacy is well-understood [16, 7]. Similarly, if there are no positions of one of the players then the game is purely determined by Lemma 1. Therefore, the simplest case specific for the branching games are the $\{E, A, \mathcal{B}\}$ -branching games.

3.1 Behavioural indeterminacy

We start by proving the following theorem.

► **Theorem 3.** *There is a $\{E, A, \mathcal{B}\}$ -branching game G with a regular winning condition that is both closed and open such that G is not determined under behavioural strategies.*



■ **Figure 2** A branching board that is not determined under behavioural strategies.

The board B of the game G is depicted in Figure 2. A play $t \in \text{plays}(B)$ over B starts by splitting into four separate threads by the \mathcal{B} -vertices labelled with c . Then, each of the players can perform two separate choices, E in the two vertices labelled x_1 and x_2 , and A in the two vertices labelled x_3 and x_4 . Their choices lead to vertices labelled by either 0 or 1. The rest of the play stays forever in the branching vertex labelled by f . For $i = 1, 2, 3, 4$ let $x_i(t) \in \{0, 1\}$ be label chosen by the respective player in the vertex labelled by x_i , i.e. the label of the unique child of the unique node labelled by x_i in t . Consider a winning set $L \subseteq \text{plays}(B)$ defined as follows

$$L \stackrel{\text{def}}{=} \{t \in \text{plays}(B) \mid x_1(t) = x_2(t) = x_3(t) = x_4(t) \vee x_3(t) \neq x_4(t)\} \quad (4)$$

In other words, Eve wins a play t if either Adam has chosen two different labels in x_3 and x_4 or all the chosen labels are equal. Since the vertices labelled x_i lie at a fixed depth of every play $t \in \text{plays}(B)$, L is a closed and open regular language of infinite trees.

► **Example 4.** The game $G = \langle B, L \rangle$ has the following partial values:

$$\text{val}_G^A = 1; \quad \text{val}_G^{BA} = \frac{3}{4}; \quad \text{val}_G^{MA} = \frac{1}{2} = \text{val}_G^{ME}; \quad \text{val}_G^{BE} = \frac{1}{4}; \quad \text{val}_G^E = 0.$$

We first argue about the pure values – a pure strategy over the board from Figure 2 needs to declare in advance the two values $x_i(t)$ and $x_{i+1}(t)$ for $i = 1, 3$ depending on the player.

If such a strategy is fixed, the opponent can choose his values in such a way to win.


We now consider the mixed value. Let σ_m randomly choose with equal probability between the following two pure strategies σ_i for $i = 0, 1$: the strategy σ_i satisfies $x_1(\sigma_i) = x_2(\sigma_i) = i$. π_m is defined analogously. It is easy to check that these strategies are optimal and witness that the mixed value of the game is $\frac{1}{2}$ for both players.

Consider a behavioural strategy σ_b of Eve (the case of Adam is entirely dual). Such a strategy can be described by two independent random choices:

1. σ_b chooses x_1 to be 0 with probability p_1 ,
2. σ_b chooses x_2 to be 0 with probability p_2 .

Thus, each behavioural strategy of Eve is characterised by a pair of numbers $p_1, p_2 \in [0, 1]$. A simple computation shows that no matter how Eve chooses her values p_1, p_2 , Adam can find a counter-strategy guaranteeing the value of at most $\frac{1}{4}$.

Since $\text{val}_G^{BA} = \frac{3}{4} \neq \frac{1}{4} = \text{val}_G^{BE}$ the proof of Theorem 3 is concluded.



■ **Figure 3** A branching board that is not determined under mixed strategies.

3.2 Mixed indeterminacy

We will now show that the mixed determinacy fails for relatively simple regular sets, as expressed by the following theorem.

► **Theorem 5.** *There is a $\{E, A, \mathcal{B}\}$ -branching game with a regular winning set being a difference of two open sets that is not determined under mixed strategies.*

To prove this theorem we will encode the following game as a $\{E, A, \mathcal{B}\}$ -branching game G . Assume that ∞ is an additional symbol such that for every $n \in \omega$ we have $n < \infty$.

► **Example 6 (Folklore).** Consider the following game: Adam and Eve simultaneously and independently choose two numbers: Eve chooses $e \in \omega \cup \{\infty\}$, Adam chooses $a \in \omega \cup \{\infty\}$. Eve wins if $e < \infty$, and either $a = \infty$ or $a \leq e$.

It is easy to see that this game is not determined under mixed strategies. Intuitively, it follows from the fact that both players try to choose a finite number as big as possible.

The board B of the game G is depicted in Figure 3. A play $t \in \text{plays}(B)$ consists of infinitely many independent sub-games that start in the vertices labelled by b . More precisely, the k -th sub-game starts in the node $L^k R$ in the tree t . Such a sub-game is split into two independent choices: Adam chooses a label, either 0 or 1, for the successor of the node labelled by a ; Eve chooses a label, either 2 or 3, for the successor of the node labelled by e .

Let $a(t)$ (resp. $e(t)$) be the smallest number $k \in \omega$ such that Eve (resp. Adam) has chosen an odd label in the k -th sub-game, i.e. $L^k R L R \in \text{Dom}(t)$ (resp. $L^k R R R \in \text{Dom}(t)$). If no such number exists then $a(t)$ (resp. $e(t)$) equals ∞ .

Let the winning condition L of the game G be defined as follows

$$L \stackrel{\text{def}}{=} \{t \in \text{plays}(B) \mid e(t) < \infty \text{ and not } (e(t) < a(t) < \infty)\}. \quad (5)$$

It is easy to see that L is a regular language of infinite trees (to compare $a(t)$ with $e(t)$ it is enough to notice that each of these values corresponds to a node on the left-most branch of the play t). Moreover, both the conditions $e(t) < \infty$ and $e(t) < a(t) < \infty$ are open sets of plays.

Hence, the game $G = \langle B, L \rangle$ is a game as required in Theorem 5. Moreover, there is a clear correspondence between the pure strategies in G and the pure strategies in the game from Example 6. This correspondence extends to the mixed strategies what implies the following claim.

► **Claim 7.** *We have that $val_G^{MA} = 1$ and $val_G^{ME} = 0$.*

This concludes the proof of Theorem 5.

3.3 Mixed determinacy for closed sets

In this section we use Glicksberg's minimax theorem to prove that if a winning condition is a closed set of plays then the game is determined under mixed strategies.

► **Theorem 8.** *If $G = \langle B, L \rangle$ is a $\{E, A, \mathcal{N}, \mathcal{B}\}$ -branching game and L is an arbitrary closed subset of $plays(B)$ then G is determined under mixed strategies.*

Before we recall the statement of Glicksberg's minimax theorem, let us introduce some relevant notions. Assume that X is a metrisable topological space. We say that a function $f: X \rightarrow \mathbb{R}$ is *upper semi-continuous* if for every $x_0 \in X$ we have $\limsup_{x \rightarrow x_0} f(x) \leq f(x_0)$. Clearly, if $C \subseteq X$ is a closed subset of X then the characteristic function of C is upper semi-continuous. Also, a composition of a continuous function and an upper semi-continuous function is upper semi-continuous.

► **Theorem 9** (Glicksberg's minimax theorem [12], see also [18, pages 299–306]). *Let A, B be compact metrisable spaces and $f: A \times B \rightarrow \mathbb{R}$ be an upper semi-continuous function. Then the following holds*

$$\sup_{\mu} \inf_{\nu} \int_{A,B} f(a, b) \mu(da) \nu(db) = \inf_{\nu} \sup_{\mu} \int_{A,B} f(a, b) \mu(da) \nu(db), \quad (6)$$

where μ, ν range over the Borel probability measures on the sets A, B respectively.

It remains to prove that if $G = \langle B, L \rangle$ with $L \subseteq plays(B)$ closed then the function $val_G: \Sigma_B^E \times \Sigma_B^A \rightarrow \mathbb{R}$ is upper semi-continuous. This function can be written as a composition of two functions. The first one maps a pair of pure strategies (σ, π) to a measure on $plays(B)$ defined as $\mu_{(\sigma, \pi)}(T) \stackrel{\text{def}}{=} \nu_B^*(\{t \in \Sigma_B^{\mathcal{N}} \mid \sigma \cap \pi \cap t \in T\})$, i.e. the ν_B^* measure of the pre-image of the set T under the function that intersects the three strategies. This mapping is continuous, as proved by Mio in [20, Lemma 4.1.4]. The second one applies the measure $\mu_{(\sigma, \pi)}$ to the winning set $L \subseteq plays(B)$. For a closed set L this function is upper semi-continuous by [14, Corollary 17.21].

4 Computing game values

In this section we will discuss the computational complexity of determining the partial values of branching games. To be more precise, we consider the following family of problems, parametrised by the set of available positions $\mathcal{P} \subseteq \{A, E, \mathcal{N}, \mathcal{B}\}$ and the type of the value $V \in \{val^A, val^{BA}, val^{MA}, val^{ME}, val^{BE}, val^E\}$.

► **Problem 10** (The value V of a regular \mathcal{P} -branching game).

- **Input:** *A finitary \mathcal{P} -branching game G with the winning condition given by a non-deterministic tree automaton.*
- **Output:** *Does $V > \frac{1}{2}$?*

4.1 The non-stochastic case

If no random choice is involved, i.e. the board has no *Nature's* positions and we consider pure strategies, the values belong to the set $\{0, 1\}$ and we can compute them, as expressed by the following theorem.

► **Theorem 11.** *The value val^E problem of a regular $\{A, E, \mathcal{B}\}$ -branching game is in 2-EXP, the value val^A problem of a regular $\{A, E, \mathcal{B}\}$ -branching game is EXP-complete.*

Moreover, the value val^E problem of a regular $\{A, E, \mathcal{B}\}$ -branching game is 2-EXP-complete if the winning condition is given by an alternating tree automaton.

This theorem follows from the constructions in [22], performed in a bit different language. The asymmetry in this theorem comes from the fact that in Problem 10 we assume that the winning condition of a game is given as a non-deterministic automaton. In this work we strengthen the second part of the above theorem by proving that the value val^E problem of a regular $\{A, E, \mathcal{B}\}$ -branching game is 2-EXP-hard also for non-deterministic automata. This is achieved by using the completeness result from [22] together with the following reduction. It is somehow surprising to notice that in the context of branching games one can de-alternate an automaton in polynomial time.

► **Theorem 12.** *There exists a polynomial time reduction that inputs a $\{A, E, \mathcal{B}\}$ -branching game G with the winning condition given as an alternating tree automaton and constructs a $\{A, E, \mathcal{B}\}$ -branching game G' with the winning condition given by a non-deterministic tree automaton, such that $val_G^E = val_{G'}^E$.*

The proof is straightforward, its main idea is to split the alternation of the given automaton into two parts: the choices of Adam and the choices of Eve. In the game G' the former choices will be done explicitly on the board while the latter choices will be performed by the non-deterministic automaton that recognises the winning condition of G' .

► **Corollary 13.** *val^E problem of a regular $\{A, E, \mathcal{B}\}$ -branching game is 2-EXP-complete.*

4.2 The stochastic cases

The above decidability results hold for non-stochastic games and pure strategies. Restoring any of those features yields undecidability, as expressed by the two theorems of this section.


► **Theorem 14.** *For every $V \in \{val^A, val^{BA}, val^{MA}, val^{ME}, val^{BE}, val^E\}$ and $P \in \{E, A\}$, the value V problem of a regular $\{P, \mathcal{N}, \mathcal{B}\}$ -branching game is undecidable.*

Observe that by Lemma 1 a $\{P, \mathcal{N}, \mathcal{B}\}$ -branching game is determined under pure strategies. It means that all the six partial values are the same for such games. Thus, by the symmetry we can assume that $P = E$ and $V = val^E$.

To prove Theorem 14 we reduce the following undecidable problem, cf. [11]. It can be shown that the word problem is undecidable even if we restrict our attention to a two-letters alphabet and the so-called *very simple* automata: a non-deterministic automaton is *very simple* if from every state and every letter there are exactly two possible transitions leading to two distinct states.

► **Problem 15** (Word problem for VSNA).

- **Input:** *A very simple non-deterministic automaton \mathcal{A} on finite words over $\{a, b\}$.*
- **Output:** *Does there exist a finite word such that more than half of the runs of \mathcal{A} on this word is accepting?*



(a) A branching board used in the proof of Theorem 14. (b) A gadget used in the proof of Theorem 16. to replace *Nature's* vertex in the board

■ **Figure 4** Boards used in undecidability proofs.

We will now sketch the proof of Theorem 14. Let us take a very simple non-deterministic automaton \mathcal{A} and assume that the two transitions over a letter $l \in \{a, b\}$ from a state $q \in Q^{\mathcal{A}}$ lead to the states $\delta_0(q, l)$ and $\delta_1(q, l)$. Assume that l_0, l_1, \dots, l_k is a sequence of letters $l_i \in \{a, b\}$ and n_0, n_1, \dots, n_k is a sequence of numbers $n_i \in \{0, 1\}$. These two sequences allow us to naturally define a run $\rho = \text{run}(\vec{l}, \vec{n})$ of \mathcal{A} over the word l_0, \dots, l_k that follows the respective transitions of \mathcal{A} : $\rho[0] = q_I$ and $\rho[i + 1] = \delta_{n_i}(\rho[i], l_i)$.

Consider the board B depicted on Figure 4a. A play on this board consists of a sequence of decisions made by Eve, whether to move from the vertex labelled l to a or to b . At every moment Eve can stop this sequence by choosing the right successor of the vertex labelled p . For every choice of a or b by Eve, the *Nature* simultaneously chooses a number 0 or 1. Thus, a play t results in two finite or infinite sequences of the same length: l_0, l_1, \dots with $l_i \in \{a, b\}$ and n_0, n_1, \dots with $n_i \in \{0, 1\}$. Consider the following winning condition

$$L \stackrel{\text{def}}{=} \{t \in \text{plays}(B) \mid \text{the sequences } \vec{l} \text{ and } \vec{n} \text{ are finite and } \text{run}(\vec{l}, \vec{n}) \text{ is accepting}\}. \quad (7)$$

Now let $G = \langle B, L \rangle$. It is easy to see that the winning condition L can be represented as a regular language of infinite trees. A pure strategy of Eve in G either never moves from the vertex labelled p to the vertex labelled f (in that case its value is 0) or in the opposite case it corresponds to a finite word l_0, l_1, \dots, l_k . The value of such a strategy is the probability that the choices of *Nature* will represent an accepting run of \mathcal{A} over the word \vec{l} . Thus, Eve has a pure strategy σ with $\text{val}_G^E(\sigma) > \frac{1}{2}$ if and only if more than half of the runs of \mathcal{A} over the word \vec{l} produced by σ is accepting.

To complete the landscape of decidability we state.

► **Theorem 16.** *For every $V \in \{\text{val}^{BA}, \text{val}^{MA}, \text{val}^{ME}, \text{val}^{BE}\}$ the value V problem of a regular $\{E, A, \mathcal{B}\}$ -branching game is undecidable.*

The theorem follows from the fact that the game used in the proof of Theorem 14 can be simulated on the board with *Nature's* position replaced by the gadget depicted in Figure 4b.

5 Conclusions

In this work we have studied questions of determinacy and decidability of regular branching games. We have shown that the games are not determined even for topologically simple regular conditions. In the case of mixed determinacy, the frontier lies in the first level of the difference hierarchy of closed sets. Additionally, we have shown that the question whether the value of a given game is greater than a fixed threshold is undecidable in all non-trivial stochastic cases. In the non-stochastic cases (i.e. when the board is non-stochastic and we ask about pure strategies) we have given exact bounds on the complexity of the problem. The only remaining case is the 0-player stochastic case, i.e. the problem of computing the measure of a regular language of infinite trees.

Further work. It seems interesting to understand for which classes of regular winning conditions, the branching games are determined. It was proved by the first author in [22] that the non-stochastic branching games with winning conditions given by *game automata* are determined under pure strategies. We believe that the proof can be naturally extended to the stochastic case. However, there are regular languages of infinite trees L that are not recognisable by game automata, but still all the branching games with the winning condition L are purely determined. The characterisation of such objectives poses an interesting research direction as it could give a broader class of games with decidable value problem.

On the frontier of mixed determinacy, it seems that allowing the objective to check local consistency at arbitrary depths of the tree is the cause of both the indeterminacy and the undecidability. This intuition suggests the following conjecture. We say that L is a *path language* if L is a Boolean combination of languages of the form

$$\{t \mid \text{there exists a branch of } t \text{ belonging to a regular language of infinite words } K \subseteq \Gamma^\omega\}.$$

► **Conjecture 17.** *If $G = \langle B, L \rangle$ is a branching game and L is a path language then the game G is determined under mixed strategies.*

Acknowledgement. The authors would like to thank Damian Niwiński for insightful discussions on the subject.

References

- 1 Eugene Asarin, Julien Cervelle, Aldric Degorre, Catalin Dima, Florian Horn, and Victor Kozyakin. Entropy games. *CoRR*, abs/1506.04885, 2015. URL: <http://arxiv.org/abs/1506.04885>.
- 2 Robert J. Aumann. Mixed and behavior strategies in infinite extensive games. Research Memorandum 32, Economic Research Program, Princeton University, 1961.
- 3 Julius Richard Büchi and Lawrence H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.
- 4 Krishnendu Chatterjee and Laurent Doyen. The complexity of partial observation parity games. In *In Proc. LPAR'10, LNCS 6397*. Springer, 2010.
- 5 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. A survey of partial-observation stochastic parity games. *Formal Methods in System Design*, 43(2):268–284, 2012.
- 6 Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. Algorithms for omega-regular games with imperfect information. *CoRR*, abs/0706.2619, 2007. URL: <http://arxiv.org/abs/0706.2619>.

- 7 Krishnendu Chatterjee and Thomas A. Henzinger. A survey of stochastic ω -regular games. *J. Comput. Syst. Sci.*, 78(2):394–413, 2012.
- 8 Taolue Chen, Klaus Dräger, and Stefan Kiefer. Model checking stochastic branching processes. In *MFCS*, pages 271–282, 2012.
- 9 Thomas Colcombet and Damian Niwiński. On the positional determinacy of edge-labeled games. *Theor. Comput. Sci.*, 352(1–3):190–196, 2006.
- 10 Stefan Dziembowski, Marcin Jurdziński, and Igor Walukiewicz. How much memory is needed to win infinite games? In *LICS*, pages 99–110, 1997.
- 11 Hugo Gimbert and Youssef Oualhaj. Probabilistic automata on finite words: Decidable and undecidable problems. In *ICALP (2)*, pages 527–538, 2010.
- 12 Irving Leonard Glicksberg. *Minimax Theorem for Upper and Lower Semi-continuous Pay-offs*. Memorandum (Rand Corporation). Rand Corporation, 1950.
- 13 Tomasz Gogacz, Henryk Michalewski, Matteo Mio, and Michał Skrzypczak. Measure properties of game tree languages. In *MFCS*, pages 303–314, 2014.
- 14 Alexander Kechris. *Classical descriptive set theory*. Springer-Verlag, New York, 1995.
- 15 Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
- 16 Donald A. Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
- 17 Robert McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65(2):149–184, 1993.
- 18 P. Wolfe Melvin Dresher, A. W. Tucker. *Contributions to the Theory of Games*. Number 3 in Annals of mathematics studies. Princeton University Press, 1957.
- 19 Henryk Michalewski and Matteo Mio. On the problem of computing the probability of regular sets of trees. In *FSTTCS*, pages 489–502, 2015.
- 20 Matteo Mio. *Game Semantics for Probabilistic μ -Calculus*. PhD thesis, University of Edinburgh, 2012.
- 21 Matteo Mio. On the equivalence of game and denotational semantics for the probabilistic mu-calculus. *Logical Methods in Computer Science*, 8(2), 2012.
- 22 Marcin Przybyłko. Tree games with regular objectives. In *GandALF*, pages 231–244, 2014.
- 23 Michael Oser Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. of the American Math. Soc.*, 141:1–35, 1969.
- 24 Wolfgang Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, pages 389–455. Springer, 1996.
- 25 Wolfgang Thomas and Helmut Lescow. Logical specifications of infinite computations. In *REX School/Symposium*, pages 583–621, 1993.