

Decision Problems for Nash Equilibria in Stochastic Games^{*}

Michael Ummels¹ and Dominik Wojtczak²

¹ RWTH Aachen University, Germany
E-Mail: ummels@logic.rwth-aachen.de

² CWI, Amsterdam, The Netherlands
E-Mail: d.k.wojtczak@cwi.nl

ABSTRACT. *We analyse the computational complexity of finding Nash equilibria in stochastic multiplayer games with ω -regular objectives. While the existence of an equilibrium whose payoff falls into a certain interval may be undecidable, we single out several decidable restrictions of the problem. First, restricting the search space to stationary, or pure stationary, equilibria results in problems that are typically contained in PSPACE and NP, respectively. Second, we show that the existence of an equilibrium with a binary payoff (i.e. an equilibrium where each player either wins or loses with probability 1) is decidable. We also establish that the existence of a Nash equilibrium with a certain binary payoff entails the existence of an equilibrium with the same payoff in pure, finite-state strategies.*

1 Introduction

We study *stochastic games* [21] played by multiple players on a finite, directed graph. Intuitively, a play of such a game evolves by moving a token along edges of the graph: Each vertex of the graph is either controlled by one of the players, or it is *stochastic*. Whenever the token arrives at a non-stochastic vertex, the player who controls this vertex must move the token to a successor vertex; when the token arrives at a stochastic vertex, a fixed probability distribution determines the next vertex. A measurable function maps plays to payoffs. In the simplest case, which we discuss here, the possible payoffs of a single play are binary (i.e. each player either wins or loses a given play). However, due to the presence of stochastic vertices, a player's *expected payoff* (i.e. her probability of winning) can be an arbitrary probability.

^{*}This work was supported by the DFG Research Training Group 1298 (ALGO SYN) and the ESF Research Networking Programme “Games for Design and Verification”.

Stochastic games with ω -regular objectives have been successfully applied in the verification and synthesis of reactive systems under the influence of random events. Such a system is usually modelled as a game between the system and its environment, where the environment's objective is the complement of the system's objective: the environment is considered hostile. Therefore, the research in this area has traditionally focused on two-player games where each play is won by precisely one of the two players, so-called *two-player zero-sum games*. However, the system may comprise of several components with independent objectives, a situation which is naturally modelled by a multiplayer game.

The most common interpretation of rational behaviour in multiplayer games is captured by the notion of a *Nash equilibrium* [20]. In a Nash equilibrium, no player can improve her payoff by unilaterally switching to a different strategy. Chatterjee & al. [7] gave an algorithm for computing a Nash equilibrium in a stochastic multiplayer games with ω -regular winning conditions. We argue that this is not satisfactory. Indeed, it can be shown that their algorithm may compute an equilibrium where all players lose almost surely (i.e. receive expected payoff 0), while there exist other equilibria where all players win almost surely (i.e. receive expected payoff 1).

In applications, one might look for an equilibrium where as many players as possible win almost surely or where it is guaranteed that the expected payoff of the equilibrium falls into a certain interval. Formulated as a decision problem, we want to know, given a k -player game \mathcal{G} with initial vertex v_0 and two thresholds $\bar{x}, \bar{y} \in [0, 1]^k$, whether (\mathcal{G}, v_0) has a Nash equilibrium with expected payoff at least \bar{x} and at most \bar{y} . This problem, which we call NE for short, is a generalisation of the *quantitative decision problem* for two-player zero-sum games, which asks whether in such a game player 0 has a strategy that ensures to win the game with a probability that lies above a given threshold.

In this paper, we analyse the decidability of NE for games with ω -regular objectives. Although the decidability of NE remains open, we can show that several restrictions of NE are decidable: First, we show that NE becomes decidable when one restricts the search space to equilibria in positional (i.e. pure, stationary), or stationary, strategies, and that the resulting decision problems typically lie in NP and PSPACE, respectively (e.g. if the objectives are specified as Muller conditions). Second, we show that the following *qualitative* version of NE is decidable: Given a k -player game \mathcal{G} with initial vertex v_0 and a *binary* payoff $\bar{x} \in \{0, 1\}^k$, decide whether (\mathcal{G}, v_0) has a Nash equilibrium with expected payoff \bar{x} . Moreover, we prove that, depending on the representation of the objective, this problem is typically complete for one of the complexity classes P, NP, coNP and PSPACE, and that the problem is invariant under restricting the search space to equilibria in pure, finite-state strategies.

Our results have to be viewed in light of the (mostly) negative results we

derived in [26]. In particular, it was shown in [26] that NE becomes undecidable if one restricts the search space to equilibria in pure strategies (as opposed to equilibria in possibly mixed strategies), even for *simple stochastic multiplayer games*. These are games with simple reachability objectives. The undecidability result crucially makes use of the fact that the Nash equilibrium one is looking for can have a payoff that is not binary. Hence, this result cannot be applied to the qualitative version of NE, which we show to be decidable in this paper. It was also proven in [26] that the problems that arise from NE when one restricts the search space to equilibria in positional or stationary strategies are both NP-hard. Moreover, we showed that the restriction to stationary strategies is at least as hard as the problem SqrtSum [1], a problem which is not known to lie inside the polynomial hierarchy. This demonstrates that the upper bounds we prove for these problems in this paper will be hard to improve.

Due to lack of space, some of the proofs in this paper are either only sketched or omitted entirely. For the complete proofs, see [27].

RELATED WORK. Determining the complexity of Nash equilibria has attracted much interest in recent years. In particular, a series of papers culminated in the result that computing a Nash equilibrium of a two-player game in strategic form is complete for the complexity class PPAD [12, 8]. However, the work closest to ours is [25], where the decidability of (a variant of) the qualitative version of NE in infinite games without stochastic vertices was proven. Our results complement the results in that paper, and although our decidability proof for the qualitative setting is structurally similar to the one in [25], the presence of stochastic vertices makes the proof substantially more challenging.

Another subject that is related to the study of stochastic multiplayer games are *Markov decision processes with multiple objectives*. These can be viewed as stochastic multiplayer games where all non-stochastic vertices are controlled by one single player. For ω -regular objectives, Etessami & al. [16] proved the decidability of NE for these games. Due to the different nature of the restrictions, this result is incomparable to our results.

2 Preliminaries

The model of a (*two-player zero-sum*) *stochastic game* [9] easily generalises to the multiplayer case: Formally, a *stochastic multiplayer game* (SMG) is a tuple $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, \Delta, (\text{Win}_i)_{i \in \Pi})$ where

- ◆ Π is a finite set of *players* (usually $\Pi = \{0, 1, \dots, k-1\}$);
- ◆ V is a finite, non-empty set of *vertices*;
- ◆ $V_i \subseteq V$ and $V_i \cap V_j = \emptyset$ for each $i \neq j \in \Pi$;
- ◆ $\Delta \subseteq V \times ([0, 1] \cup \{\perp\}) \times V$ is the *transition relation*;

- ♦ $\text{Win}_i \subseteq V^\omega$ is a Borel set for each $i \in \Pi$.

The structure $G = (V, (V_i)_{i \in \Pi}, \Delta)$ is called the *arena* of \mathcal{G} , and Win_i is called the *objective*, or the *winning condition*, of player $i \in \Pi$. A vertex $v \in V$ is *controlled by player i* if $v \in V_i$ and a *stochastic vertex* if $v \notin \bigcup_{i \in \Pi} V_i$.

We require that a transition is labelled by a probability iff it originates in a stochastic vertex: If $(v, p, w) \in \Delta$ then $p \in [0, 1]$ if v is a stochastic vertex and $p = \perp$ if $v \in V_i$ for some $i \in \Pi$. Additionally, for each pair of a stochastic vertex v and an arbitrary vertex w , we require that there exists precisely one $p \in [0, 1]$ such that $(v, p, w) \in \Delta$. Moreover, for each stochastic vertex v , the outgoing probabilities must sum up to 1: $\sum_{(p,w):(v,p,w) \in \Delta} p = 1$. Finally, we require that for each vertex the set $v\Delta := \{w \in V : \text{exists } p \in (0, 1] \cup \{\perp\} \text{ with } (v, p, w) \in \Delta\}$ is non-empty, i.e. every vertex has at least one successor.

A special class of SMGs are *two-player zero-sum stochastic games (2SGs)*. These are SMGs played by only two players (player 0 and player 1) and one player's objective is the complement of the other player's objective, i.e. $\text{Win}_0 = V^\omega \setminus \text{Win}_1$. An even more restricted model are *one-player stochastic games*, also known as *Markov decision processes (MDPs)*, where there is only one player (player 0). Finally, *Markov chains* are SMGs with no players at all, i.e. there are only stochastic vertices.

STRATEGIES AND STRATEGY PROFILES. In the following, let \mathcal{G} be an arbitrary SMG. A (*mixed*) *strategy of player i in \mathcal{G}* is a mapping $\sigma : V^* V_i \rightarrow \mathcal{D}(V)$ assigning to each possible *history* $xv \in V^* V_i$ of vertices ending in a vertex controlled by player i a (discrete) probability distribution over V such that $\sigma(xv)(w) > 0$ only if $(v, \perp, w) \in \Delta$. Instead of $\sigma(xv)(w)$, we usually write $\sigma(w \mid xv)$. A (*mixed*) *strategy profile of \mathcal{G}* is a tuple $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ where σ_i is a strategy of player i in \mathcal{G} . Given a strategy profile $\bar{\sigma} = (\sigma_j)_{j \in \Pi}$ and a strategy τ of player i , we denote by $(\bar{\sigma}_{-i}, \tau)$ the strategy profile resulting from $\bar{\sigma}$ by replacing σ_i with τ .

A strategy σ of player i is called *pure* if for each $xv \in V^* V_i$ there exists $w \in v\Delta$ with $\sigma(w \mid xv) = 1$. Note that a pure strategy of player i can be identified with a function $\sigma : V^* V_i \rightarrow V$. A strategy profile $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ is called *pure* if each σ_i is pure.

A strategy σ of player i in \mathcal{G} is called *stationary* if σ depends only on the current vertex: $\sigma(xv) = \sigma(v)$ for all $xv \in V^* V_i$. Hence, a stationary strategy of player i can be identified with a function $\sigma : V_i \rightarrow \mathcal{D}(V)$. A strategy profile $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ of \mathcal{G} is called *stationary* if each σ_i is stationary.

We call a pure, stationary strategy a *positional strategy* and a strategy profile consisting of positional strategies only a *positional strategy profile*. Clearly, a positional strategy of player i can be identified with a function $\sigma : V_i \rightarrow V$. More generally, a pure strategy σ is called *finite-state* if it can be implemented

by a finite automaton with output or, equivalently, if the equivalence relation $\sim \subseteq V^* \times V^*$ defined by $x \sim y$ if $\sigma(xz) = \sigma(yz)$ for all $z \in V^* V_i$ has only finitely many equivalence classes. Finally, a *finite-state strategy profile* is a profile consisting of finite-state strategies only.

It is sometimes convenient to designate an initial vertex $v_0 \in V$ of the game. We call the tuple (\mathcal{G}, v_0) an *initialised SMG*. A strategy (strategy profile) of (\mathcal{G}, v_0) is just a strategy (strategy profile) of \mathcal{G} . In the following, we will use the abbreviation SMG also for initialised SMGs. It should always be clear from the context if the game is initialised or not.

Given an initial vertex v_0 and a strategy profile $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$, the *conditional probability of $w \in V$ given the history $xv \in V^* V$* is the number $\sigma_i(w \mid xv)$ if $v \in V_i$ and the unique $p \in [0, 1]$ such that $(v, p, w) \in \Delta$ if v is a stochastic vertex. We abuse notation and denote this probability by $\bar{\sigma}(w \mid xv)$. The probabilities $\bar{\sigma}(w \mid xv)$ induce a probability measure on the space V^ω in the following way: The probability of a basic open set $v_1 \dots v_k \cdot V^\omega$ is 0 if $v_1 \neq v_0$ and the product of the probabilities $\bar{\sigma}(v_j \mid v_1 \dots v_{j-1})$ for $j = 2, \dots, k$ otherwise. It is a classical result of measure theory that this extends to a unique probability measure assigning a probability to every Borel subset of V^ω , which we denote by $\Pr_{v_0}^{\bar{\sigma}}$.

For a strategy profile $\bar{\sigma}$, we are mainly interested in the probabilities $p_i := \Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i)$ of winning. We call p_i the (*expected*) *payoff of $\bar{\sigma}$ for player i* and the vector $(p_i)_{i \in \Pi}$ the (*expected*) *payoff of $\bar{\sigma}$* .

SUBARENAS AND END COMPONENTS. Given an SMG \mathcal{G} , we call a set $U \subseteq V$ a *subarena of \mathcal{G}* if 1. $U \neq \emptyset$; 2. $v\Delta \cap U \neq \emptyset$ for each $v \in U$, and 3. $v\Delta \subseteq U$ for each stochastic vertex $v \in U$.

A set $C \subseteq V$ is called an *end component of \mathcal{G}* if C is a subarena, and additionally C is strongly connected: for every pair of vertices $v, w \in C$ there exists a sequence $v = v_1, v_2, \dots, v_n = w$ with $v_{i+1} \in v_i\Delta$ for each $0 < i < n$. An end component C is *maximal in a set $U \subseteq V$* if there is no end component $C' \subseteq U$ with $C \subsetneq C'$. For any subset $U \subseteq V$, the set of all end components maximal in U can be computed by standard graph algorithms in quadratic time (see e.g. [13]).

The central fact about end components is that, under *any* strategy profile, the set of vertices visited infinitely often is almost surely an end component. For an infinite sequence α , we denote by $\text{Inf}(\alpha)$ the set of elements occurring infinitely often in α .

Lemma 1 ([13, 10]). Let \mathcal{G} be any SMG, and let $\bar{\sigma}$ be any strategy profile of \mathcal{G} . Then $\Pr_v^{\bar{\sigma}}(\{\alpha \in V^\omega : \text{Inf}(\alpha) \text{ is an end component}\}) = 1$ for each vertex $v \in V$.

Moreover, for any end component C , we can construct a stationary strategy profile $\bar{\sigma}$ that, when started in C , guarantees to visit all (and only) vertices in C infinitely often.

Lemma 2 ([13, 11]). Let \mathcal{G} be any SMG, and let C be any end component of \mathcal{G} . There exists a stationary strategy profile $\bar{\sigma}$ with $\Pr_v^{\bar{\sigma}}(\{\alpha \in V^\omega : \text{Inf}(\alpha) = C\}) = 1$ for each vertex $v \in C$.

VALUES, DETERMINACY AND OPTIMAL STRATEGIES. Given a strategy τ of player i in \mathcal{G} and a vertex $v \in V$, the *value of τ from v* is the number $\text{val}^\tau(v) := \inf_{\bar{\sigma}} \Pr_v^{\bar{\sigma}-i, \tau}(\text{Win}_i)$, where $\bar{\sigma}$ ranges over all strategy profiles of \mathcal{G} . Moreover, we define the *value of \mathcal{G} for player i from v* as the supremum of these values, i.e. $\text{val}_i^{\mathcal{G}}(v) = \sup_{\tau} \text{val}^\tau(v)$, where τ ranges over all strategies of player i in \mathcal{G} . Intuitively, $\text{val}_i^{\mathcal{G}}(v)$ is the maximal payoff that player i can ensure when the game starts from v . If \mathcal{G} is a two-player zero-sum game, a celebrated theorem due to Martin [19] states that the game is *determined*, i.e. $\text{val}_0^{\mathcal{G}} = 1 - \text{val}_1^{\mathcal{G}}$ (where the equality holds pointwise). The number $\text{val}^{\mathcal{G}}(v) := \text{val}_0^{\mathcal{G}}(v)$ is consequently called the *value of \mathcal{G} from v* .

Given an initial vertex $v_0 \in V$, a strategy σ of player i in \mathcal{G} is called *optimal* if $\text{val}^\sigma(v_0) = \text{val}_i^{\mathcal{G}}(v_0)$. A *globally optimal* strategy is a strategy that is optimal for every possible initial vertex $v_0 \in V$. Note that optimal strategies do not need to exist since the supremum in the definition of $\text{val}_i^{\mathcal{G}}$ is not necessarily attained. However, if for every possible initial vertex there exists an optimal strategy, then there also exists a globally optimal strategy.

OBJECTIVES. We have introduced objectives as abstract Borel sets of infinite sequences of vertices; to be amenable for algorithmic solutions, all objectives must be finitely representable. In verification, objectives are usually *ω -regular sets* specified by formulae of the logic SIS (monadic second-order logic on infinite words) or LTL (linear-time temporal logic) referring to unary predicates P_c indexed by a finite set C of *colours*. These are interpreted as winning conditions in a game by considering a *colouring* $\chi : V \rightarrow C$ of the vertices in the game. Special cases are the following well-studied conditions:

- ◆ *Büchi* (given by a set $F \subseteq C$): the set of all $\alpha \in C^\omega$ such that $\text{Inf}(\alpha) \cap F \neq \emptyset$.
- ◆ *co-Büchi* (given by set $F \subseteq C$): the set of all $\alpha \in C^\omega$ such that $\text{Inf}(\alpha) \subseteq F$.
- ◆ *Parity* (given by a *priority function* $\Omega : C \rightarrow \mathbb{N}$): the set of all $\alpha \in C^\omega$ such that $\min(\text{Inf}(\Omega(\alpha)))$ is even.
- ◆ *Streett* (given by a set Ω of pairs (F, G) where $F, G \subseteq C$): the set of all $\alpha \in C^\omega$ such that for all pairs $(F, G) \in \Omega$ with $\text{Inf}(\alpha) \cap F \neq \emptyset$ it is the case that $\text{Inf}(\alpha) \cap G \neq \emptyset$.
- ◆ *Rabin* (given by a set Ω of pairs (F, G) where $F, G \subseteq C$): the set of all $\alpha \in C^\omega$ such that there exists a pair $(F, G) \in \Omega$ with $\text{Inf}(\alpha) \cap F \neq \emptyset$ but $\text{Inf}(\alpha) \cap G = \emptyset$.
- ◆ *Muller* (given by a family \mathcal{F} of sets $F \subseteq C$): the set of all $\alpha \in C^\omega$ such that there exists $F \in \mathcal{F}$ with $\text{Inf}(\alpha) = F$.

Note that any Büchi condition is a parity condition with two priorities, that any parity condition is both a Streett and a Rabin condition, and that any Streett or Rabin condition is a Muller condition. (However, the translation from a set of Streett/Rabin pairs to an equivalent family of accepting sets is, in general, exponential.) In fact, the intersection (union) of any two parity conditions is a Streett (Rabin) condition. Moreover, the complement of a Büchi (Streett) condition is a co-Büchi (Rabin) condition and vice versa, whereas the class of parity conditions and the class of Muller conditions are closed under complementation. Finally, note that any of the above condition is *prefix-independent*: for every $\alpha \in C^\omega$ and $x \in C^*$, α satisfies the condition iff $x\alpha$ does.

Theoretically, parity and Rabin conditions provide the best balance of expressiveness and simplicity: On the one hand, any SMG where player i has a Rabin objective admits a globally optimal positional strategy for this player [4]. On the other hand, any SMG with ω -regular objectives can be reduced to an SMG with parity objectives using *finite memory* (see [24]). An important consequence of this reduction is that there exist globally optimal finite-state strategies in every SMG with ω -regular objectives. In fact, there exist globally optimal pure strategies in every SMG with prefix-independent objectives [17].

In the following, for the sake of simplicity, we will only consider games where each vertex is coloured by itself, i.e. $C = V$ and $\chi = \text{id}$. We would like to point out, however, that all our results remain valid for games with other colourings. For the same reason, we will usually not distinguish between a condition and its finite representation.

DECISION PROBLEMS FOR TWO-PLAYER ZERO-SUM GAMES. The main computational problem for two-player zero-sum games is computing the value (and optimal strategies for either player, if they exist). Rephrased as a decision problem, the problem looks as follows:

Given a 2SG \mathcal{G} , an initial vertex v_0 and a rational probability p , decide whether $\text{val}^{\mathcal{G}}(v_0) \geq p$.

A special case of this problem arises for $p = 1$. Here, we only want to know whether player 0 can win the game almost surely (in the limit). Let us call the former problem the *quantitative* and the latter problem the *qualitative decision problem for 2SGs*.

Table 1 summarises the results about the complexity of the quantitative and the qualitative decision problem for two-player zero-sum stochastic games depending on the type of player 0's objective. For MDPs, both problems are decidable in polynomial time for all of the aforementioned objectives (i.e. up to Muller conditions) [3, 13].

	Quantitative	Qualitative
(co-)Büchi	NP \cap coNP [6]	P-complete [14]
Parity	NP \cap coNP [6]	NP \cap coNP [6]
Streett	coNP-complete [4, 15]	coNP-complete [4, 15]
Rabin	NP-complete [4, 15]	NP-complete [4, 15]
Muller	PSPACE-complete [3, 18]	PSPACE-complete [3, 18]

Table 1. The complexity of deciding the value in 2SGs.

3 Nash equilibria and their decision problems

To capture rational behaviour of (selfish) players, John Nash [20] introduced the notion of, what is now called, a *Nash equilibrium*. Formally, given a strategy profile $\bar{\sigma}$ in an SMG (\mathcal{G}, ν_0) , a strategy τ of player i is called a *best response to $\bar{\sigma}$* if τ maximises the expected payoff of player i : $\Pr_{\nu_0}^{\bar{\sigma}-i, \tau'}(\text{Win}_i) \leq \Pr_{\nu_0}^{\bar{\sigma}-i, \tau}(\text{Win}_i)$ for all strategies τ' of player i . A Nash equilibrium is a strategy profile $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ such that each σ_i is a best response to $\bar{\sigma}$. Hence, in a Nash equilibrium no player can improve her payoff by (unilaterally) switching to a different strategy. For two-player zero-sum games, a Nash equilibrium is nothing else than a pair of optimal strategies.

Proposition 3. Let (\mathcal{G}, ν_0) be a two-player zero-sum game. A strategy profile (σ, τ) of (\mathcal{G}, ν_0) is a Nash equilibrium iff both σ and τ are optimal. In particular, every Nash equilibrium of (\mathcal{G}, ν_0) has payoff $(\text{val}^{\mathcal{G}}(\nu_0), 1 - \text{val}^{\mathcal{G}}(\nu_0))$.

So far, most research on finding Nash equilibria in infinite games has focused on computing *some* Nash equilibrium [7]. However, a game may have several Nash equilibria with different payoffs, and one might not be interested in *any* Nash equilibrium but in one whose payoff fulfils certain requirements. For example, one might look for a Nash equilibrium where certain players win almost surely while certain others lose almost surely. This idea leads to the following decision problem, which we call NE:¹

Given an SMG (\mathcal{G}, ν_0) and thresholds $\bar{x}, \bar{y} \in [0, 1]^{\Pi}$, decide whether there exists a Nash equilibrium of (\mathcal{G}, ν_0) with payoff $\geq \bar{x}$ and $\leq \bar{y}$.

Of course, as a decision problem the problem only makes sense if the game and the thresholds \bar{x} and \bar{y} are represented in a finite way. In the following, we will therefore assume that the thresholds and all transition probabilities are rational, and that all objectives are ω -regular.

Note that NE puts no restriction on the type of strategies that realise the equilibrium. It is natural to restrict the search space to equilibria that are

¹In the definition of NE, the ordering \leq is applied componentwise.

realised in pure, finite-state, stationary, or even positional strategies. Let us call the corresponding decision problems PureNE, FinNE, StatNE and PosNE, respectively.

In a recent paper [26], we studied NE and its variants in the context of *simple* stochastic multiplayer games (SSMGs). These are SMGs where each player's objective is to reach a certain set T of *terminal vertices*: $v\Delta = \{v\}$ for each $v \in T$. In particular, such objectives are both Büchi and co-Büchi conditions. Our main results on SSMGs can be summarised as follows:

- ◆ PureNE and FinNE are undecidable;
- ◆ StatNE is contained in PSPACE, but NP- and SqrtSum-hard;
- ◆ PosNE is NP-complete.

In fact, PureNE and FinNE are undecidable even if one restricts to instances where the thresholds are binary, but distinct, or if one restricts to instances where the thresholds coincide (but are not binary). Hence, the question arises what happens if the thresholds are binary and coincide. This question motivates the following *qualitative* version of NE, a problem which we call QualNE:

Given an SMG (\mathcal{G}, v_0) and $\bar{x} \in \{0,1\}^{\Pi}$, decide whether (\mathcal{G}, v_0) has a Nash equilibrium with payoff \bar{x} .

In this paper, we show that QualNE, StatNE and PosNE are decidable for games with arbitrary ω -regular objectives, and analyse the complexities of these problems depending on the type of the objectives.

4 Stationary equilibria

In this section, we analyse the complexity of the problems PosNE and StatNE. Lower bounds for these problems follow from our results on SSMGs [26].

Theorem 4. PosNE is NP-complete for SMGs with Büchi, co-Büchi, parity, Rabin, Streett, or Muller objectives.

Proof. Hardness was already proven in [26]. To prove membership in NP, we give a nondeterministic polynomial-time algorithm for deciding PosNE. On input $\mathcal{G}, v_0, \bar{x}, \bar{y}$, the algorithm simply guesses a positional strategy profile $\bar{\sigma}$ (which is basically a mapping $\bigcup_{i \in \Pi} V_i \rightarrow V$). Next, the algorithm computes the payoff z_i of $\bar{\sigma}$ for each player i by computing the probability of the event Win_i in the Markov chain $(\mathcal{G}^{\bar{\sigma}}, v_0)$, which arises from \mathcal{G} by fixing all transitions according to $\bar{\sigma}$. Once each z_i is computed, the algorithm can easily check whether $x_i \leq z_i \leq y_i$. To check whether $\bar{\sigma}$ is a Nash equilibrium, the algorithm needs to compute, for each player i , the value r_i of the MDP $(\mathcal{G}^{\bar{\sigma}-i}, v_0)$, which arises from \mathcal{G} by fixing all transitions but the ones leaving vertices controlled by player i according to $\bar{\sigma}$ (and imposing the objective Win_i). Clearly, $\bar{\sigma}$ is a

Nash equilibrium iff $r_i \leq z_i$ for each player i . Since we can compute the value of any MDP (and thus any Markov chain) with one of the above objectives in polynomial time [3, 13], all these checks can be carried out in polynomial time. Q.E.D.

To prove the decidability of StatNE, we appeal to results established for the *Existential Theory of the Reals*, $\text{ExTh}(\mathfrak{R})$, the set of all existential first-order sentences (over the appropriate signature) that hold in $\mathfrak{R} := (\mathbb{R}, +, \cdot, 0, 1, \leq)$. The best known upper bound for the complexity of the associated decision problem is PSPACE [2], which leads to the following theorem.

Theorem 5. StatNE is in PSPACE for SMGs with Büchi, co-Büchi, parity, Rabin, Streett, or Muller objective.

Proof (Sketch). Since $\text{PSPACE} = \text{NPSPACE}$, it suffices to provide a nondeterministic algorithm with polynomial space requirements for deciding StatNE. On input $\mathcal{G}, v_0, \bar{x}, \bar{y}$, where w.l.o.g. \mathcal{G} is an SMG with Muller objectives $\mathcal{F}_i \in 2^V$, the algorithm starts by guessing the *support* $S \subseteq V \times V$ of a stationary strategy profile $\bar{\sigma}$ of \mathcal{G} , i.e. $S = \{(v, w) \in V \times V : \bar{\sigma}(w | v) > 0\}$. From the set S alone, by standard graph algorithms (see [3, 13]), one can compute (in polynomial time) for each player i the following sets:

1. the union F_i of all end components (i.e. bottom SCCs) C of the Markov chain $\mathcal{G}^{\bar{\sigma}}$ that are winning for player i , i.e. $C \in \mathcal{F}_i$;
2. the set R_i of vertices v such that $\Pr_v^{\bar{\sigma}}(\text{Reach}(F_i)) > 0$;
3. the union T_i of all end components of the MDP $\mathcal{G}^{\bar{\sigma}-i}$ that are winning for player i .

After computing all these sets, the algorithm evaluates a suitable existential first-order sentence ψ , which can be computed in polynomial time from $\mathcal{G}, v_0, \bar{x}, \bar{y}, (R_i)_{i \in \Pi}, (F_i)_{i \in \Pi}$ and $(T_i)_{i \in \Pi}$ over \mathfrak{R} and returns the answer to this query. The sentence ψ states that there exists a stationary Nash equilibrium of (\mathcal{G}, v_0) with payoff $\geq \bar{x}$ and $\leq \bar{y}$ whose support is S . Q.E.D.

5 Equilibria with a binary payoff

In this section, we prove that QualNE is decidable. We start by characterising the existence of a Nash equilibrium with a binary payoff in any game with prefix-independent objectives.

5.1 Characterisation of existence

For a subset $U \subseteq V$, we denote by $\text{Reach}(U)$ the set $V^* \cdot U \cdot V^\omega$; if $U = \{v\}$, we just write $\text{Reach}(v)$ for $\text{Reach}(U)$. Finally, given an SMG \mathcal{G} and a player i , we denote by $V_i^{>0}$ the set of all vertices $v \in V$ such that $\text{val}_i^{\mathcal{G}}(v) > 0$. The following

lemma allows to infer the existence of a Nash equilibrium from the existence of a certain strategy profile. The proof uses so-called *threat strategies* (also known as *trigger strategies*), which are the basis of the *folk theorems* in the theory of repeated games (cf. [22, Chapter 8]).

Lemma 6. Let $\bar{\sigma}$ be a pure strategy profile of \mathcal{G} such that, for each player i , $\Pr_{v_0}^{\bar{\sigma}}(\text{Win}_i) = 1$ or $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V_i^{>0})) = 0$. Then there exists a pure Nash equilibrium $\bar{\sigma}^*$ with $\Pr_{v_0}^{\bar{\sigma}^*} = \Pr_{v_0}^{\bar{\sigma}}$. If, additionally, all winning conditions are ω -regular and $\bar{\sigma}$ is finite-state, then there exists a finite-state Nash equilibrium $\bar{\sigma}^*$ with $\Pr_{v_0}^{\bar{\sigma}^*} = \Pr_{v_0}^{\bar{\sigma}}$.

Proof (Sketch). Let $\mathcal{G}_i = (\{i, \Pi \setminus \{i\}\}, V, V_i, \bigcup_{j \neq i} V_j, \Delta, \text{Win}_i, V^\omega \setminus \text{Win}_i)$ be the 2SG where player i plays against the coalition $\Pi \setminus \{i\}$ of all other players. Since the set Win_i is prefix-independent, there exists a globally optimal pure strategy τ_i for the coalition in this game [17]. For each player $j \neq i$, this strategy induces a pure strategy $\tau_{j,i}$ in \mathcal{G} . To simplify notation, we also define $\tau_{i,i}$ to be an arbitrary finite-state strategy of player i in \mathcal{G} . Player i 's equilibrium strategy σ_i^* is defined as follows:

$$\sigma_i^*(xv) = \begin{cases} \sigma_i(xv) & \text{if } \Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) > 0, \\ \tau_{i,j}(x_2v) & \text{otherwise,} \end{cases}$$

where, in the latter case, $x = x_1x_2$ with x_1 being the longest prefix of xv such that $\Pr_{v_0}^{\bar{\sigma}}(x_1 \cdot V^\omega) > 0$ and $j \in \Pi$ being the player that has deviated from $\bar{\sigma}$, i.e. x_1 ends in V_j ; if x_1 is empty or ends in a stochastic vertex, we set $j = i$. Intuitively, σ_i^* behaves like σ_i as long as no other player j deviates from playing σ_j , in which case σ_i^* starts to behave like $\tau_{i,j}$.

If each Win_i is ω -regular, then each of the strategies τ_i can be chosen to be a finite-state strategy. Consequently, each $\tau_{j,i}$ can be assumed to be finite-state. If additionally $\bar{\sigma}$ is finite-state, it is easy to see that the strategy profile $\bar{\sigma}^*$, as defined above, is also finite-state.

Note that $\Pr_{v_0}^{\bar{\sigma}^*} = \Pr_{v_0}^{\bar{\sigma}}$. We claim that $\bar{\sigma}^*$ is a Nash equilibrium of (\mathcal{G}, v_0) .

Q.E.D.

Finally, we can state the main result of this section.

Proposition 7. Let (\mathcal{G}, v_0) be any SMG with prefix-independent winning conditions, and let $\bar{x} \in \{0, 1\}^{\Pi}$. Then the following statements are equivalent:

1. There exists a Nash equilibrium with payoff \bar{x} ;
2. There exists a strategy profile $\bar{\sigma}$ with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V_i^{>0})) = 0$ for each player i with $x_i = 0$;
3. There exists a pure strategy profile $\bar{\sigma}$ with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V_i^{>0})) = 0$ for each player i with $x_i = 0$;

4. There exists a pure Nash equilibrium with payoff \bar{x} .

If additionally all winning conditions are ω -regular, then any of the above statements is equivalent to each of the following statements:

5. There exists a finite-state strategy profile $\bar{\sigma}$ with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V_i^{>0})) = 0$ for each player i with $x_i = 0$;
6. There exists a finite-state Nash equilibrium with payoff \bar{x} .

Proof. (1. \Rightarrow 2.) Let $\bar{\sigma}$ be a Nash equilibrium with payoff \bar{x} . We claim that $\bar{\sigma}$ is already the strategy profile we are looking for: $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V_i^{>0})) = 0$ for each player i with $x_i = 0$. Towards a contradiction, assume that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V_i^{>0})) > 0$ for some player i with $x_i = 0$. Since V is finite, there exists a vertex $v \in V_i^{>0}$ and a history x such that $\Pr_{v_0}^{\bar{\sigma}}(xv \cdot V^\omega) > 0$. Let τ be an optimal strategy for player i in the game (\mathcal{G}, v) , and consider her strategy σ' defined by

$$\sigma'(yw) = \begin{cases} \sigma(yw) & \text{if } xv \not\leq yw, \\ \tau(y'w) & \text{otherwise,} \end{cases}$$

where, in the latter case, $y = xy'$. A straightforward calculation yields that $\Pr_{v_0}^{(\bar{\sigma}_{-i}, \sigma')}(\text{Win}_i) > 0$. Hence, player i can improve her payoff by playing σ' instead of σ_i , a contradiction to the fact that $\bar{\sigma}$ is a Nash equilibrium.

(2. \Rightarrow 3.) Let $\bar{\sigma}$ be a strategy profile of (\mathcal{G}, v_0) with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V_i^{>0})) = 0$ for each player i with $x_i = 0$. Consider the MDP \mathcal{M} that is obtained from \mathcal{G} by removing all vertices $v \in V$ such that $v \in V_i^{>0}$ for some player i with $x_i = 0$, merging all players into one, and imposing the objective

$$\text{Win} = \bigcap_{\substack{i \in \Pi \\ x_i = 1}} \text{Win}_i \cap \bigcap_{\substack{i \in \Pi \\ x_i = 0}} V^\omega \setminus \text{Win}_i.$$

The MDP \mathcal{M} is well-defined since its domain is a subarena of \mathcal{G} . Moreover, the value $\text{val}^{\mathcal{M}}(v_0)$ of \mathcal{M} is equal to 1 because the strategy profile $\bar{\sigma}$ induces a strategy σ in \mathcal{M} satisfying $\Pr_{v_0}^\sigma(\text{Win}) = 1$. Since each Win_i is prefix-independent, so is the set Win . Hence, there exists a pure, optimal strategy τ in (\mathcal{M}, v_0) . Since the value is 1, we have $\Pr_{v_0}^\tau(\text{Win}) = 1$, and τ induces a pure strategy profile of \mathcal{G} with the desired properties.

(3. \Rightarrow 4.) Let $\bar{\sigma}$ be a pure strategy profile of (\mathcal{G}, v_0) with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V_i^{>0})) = 0$ for each player i with $x_i = 0$. By Lemma 6, there exists a pure Nash equilibrium $\bar{\sigma}^*$ of (\mathcal{G}, v_0) with $\Pr_{v_0}^{\bar{\sigma}} = \Pr_{v_0}^{\bar{\sigma}^*}$. In particular, $\bar{\sigma}^*$ has payoff \bar{x} .

(4. \Rightarrow 1.) Trivial.

Under the additional assumption that all winning conditions are ω -regular,

the implications (2. \Rightarrow 5.) and (5. \Rightarrow 6.) are proven analogously; the implication (6. \Rightarrow 1.) is trivial. Q.E.D.

As an immediate consequence of Proposition 7, we can conclude that finite-state strategies are as powerful as arbitrary mixed strategies as far as the existence of a Nash equilibrium with a binary payoff in SMGs with ω -regular objectives is concerned. (This is not true for Nash equilibria with a non-binary payoff [25].)

Corollary 8. Let (\mathcal{G}, ν_0) be any SMG with ω -regular objectives, and let $x \in \{0, 1\}^\Pi$. There exists a Nash equilibrium of (\mathcal{G}, ν_0) with payoff \bar{x} iff there exists a finite-state Nash equilibrium of (\mathcal{G}, ν_0) with payoff \bar{x} .

Proof. The claim follows from Proposition 7 and the fact that every SMG with ω -regular objectives can be reduced to one with prefix-independent ω -regular (e.g. parity) objectives. Q.E.D.

5.2 Computational complexity

We can now describe an algorithm for deciding QualNE for games with Muller objectives. The algorithm relies on the characterisation we gave in Proposition 7, which allows to reduce the problem to a problem about a certain MDP.

Formally, given an SMG $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, \Delta, (\mathcal{F}_i)_{i \in \Pi})$ with Muller objectives $\mathcal{F}_i \subseteq 2^V$, and a binary payoff $\bar{x} \in \{0, 1\}^\Pi$, we define the Markov decision process $\mathcal{G}(\bar{x})$ as follows: Let $Z \subseteq V$ be the set of all v such that $\text{val}_i^{\mathcal{G}}(v) = 0$ for each player i with $x_i = 0$; the set of vertices of $\mathcal{G}(\bar{x})$ is precisely the set Z , with the set of vertices controlled by player 0 being $Z_0 := Z \cap \bigcup_{i \in \Pi} V_i$. (If $Z = \emptyset$, we define $\mathcal{G}(\bar{x})$ to be a trivial MDP with the empty set as its objective.) The transition relation of $\mathcal{G}(\bar{x})$ is the restriction of Δ to transitions between Z -states. Note that the transition relation of $\mathcal{G}(\bar{x})$ is well-defined since Z is a subarena of \mathcal{G} . We say that a subset $U \subseteq V$ has payoff \bar{x} if $U \in \mathcal{F}_i$ for each player i with $x_i = 1$ and $U \notin \mathcal{F}_i$ for each player i with $x_i = 0$. The objective of $\mathcal{G}(\bar{x})$ is $\text{Reach}(T)$ where $T \subseteq Z$ is the union of all end components $U \subseteq Z$ that have payoff \bar{x} .

Lemma 9. Let (\mathcal{G}, ν_0) be any SMG with Muller objectives, and let $\bar{x} \in \{0, 1\}^\Pi$. Then (\mathcal{G}, ν_0) has a Nash equilibrium with payoff \bar{x} iff $\text{val}^{\mathcal{G}(\bar{x})}(\nu_0) = 1$.

Proof. (\Rightarrow) Assume that (\mathcal{G}, ν_0) has a Nash equilibrium with payoff \bar{x} . By Proposition 7, this implies that there exists a strategy profile $\bar{\sigma}$ of (\mathcal{G}, ν_0) with payoff \bar{x} such that $\Pr_{\nu_0}^{\bar{\sigma}}(\text{Reach}(V \setminus Z)) = 0$. We claim that $\Pr_{\nu_0}^{\bar{\sigma}}(\text{Reach}(T)) = 1$. Otherwise, by Lemma 1, there would exist an end component $C \subseteq Z$ such that $C \notin \mathcal{F}_i$ for some player i with $x_i = 1$ or $C \in \mathcal{F}_i$ for some player i with $x_i = 0$, and $\Pr_{\nu_0}^{\bar{\sigma}}(\{\alpha \in V^\omega : \text{Inf}(\alpha) = C\}) > 0$. But then, $\bar{\sigma}$ cannot have payoff \bar{x} , a contradiction. Now, since $\Pr_{\nu_0}^{\bar{\sigma}}(\text{Reach}(V \setminus Z)) = 0$, $\bar{\sigma}$ induces a strategy σ in

$\mathcal{G}(\bar{x})$ such that $\Pr_{v_0}^\sigma(B) = \Pr_{v_0}^{\bar{\sigma}}(B)$ for every Borel set $B \subseteq Z^\omega$. In particular, $\Pr_{v_0}^\sigma(\text{Reach}(T)) = 1$ and hence $\text{val}^{\mathcal{G}(\bar{x})}(v_0) = 1$.

(\Leftarrow) Assume that $\text{val}^{\mathcal{G}(\bar{x})}(v_0) = 1$ (in particular, $v_0 \in Z$), and let σ be an optimal strategy in $(\mathcal{G}(\bar{x}), v_0)$. From σ , using Lemma 2, we can devise a strategy σ' such that $\Pr_{v_0}^{\sigma'}(\{\alpha \in V^\omega : \text{Inf}(\alpha) \text{ has payoff } \bar{x}\}) = 1$. Finally, σ' can be extended to a strategy profile $\bar{\sigma}$ of \mathcal{G} with payoff \bar{x} such that $\Pr_{v_0}^{\bar{\sigma}}(\text{Reach}(V \setminus Z)) = 0$. By Proposition 7, this implies that (\mathcal{G}, v_0) has a Nash equilibrium with payoff \bar{x} . Q.E.D.

Since the value of an MDP with reachability objectives can be computed in polynomial time (via linear programming, cf. [23]), the difficult part lies in computing the MDP $\mathcal{G}(\bar{x})$ from \mathcal{G} and \bar{x} (i.e. its domain Z and the target set T).

Theorem 10. QualNE is in PSPACE for games with Muller objectives.

Proof. Since PSPACE = NPSpace, it suffices to give a nondeterministic algorithm with polynomial space requirements. On input $\mathcal{G}, v_0, \bar{x}$, the algorithm starts by computing for each player i with $x_i = 0$ the set of vertices v with $\text{val}_i^{\mathcal{G}}(v) = 0$, which can be done in polynomial space (see Table 1). The intersection of these sets is the domain Z of the Markov decision process $\mathcal{G}(\bar{x})$. If v_0 is not contained in this intersection, the algorithm immediately rejects. Otherwise, the algorithm proceeds by guessing a set $T' \subseteq Z$ and for each $v \in T'$ a set $U_v \subseteq Z$ with $v \in U_v$. If, for each $v \in T'$, the set U_v is an end component with payoff \bar{x} , the algorithm proceeds by computing (in polynomial time) the value $\text{val}^{\mathcal{G}(\bar{x})}(v_0)$ of the MDP $\mathcal{G}(\bar{x})$ with T' substituted for T and accepts if the value is 1. In all other cases, the algorithm rejects.

The correctness of the algorithm follows from Lemma 9 and the fact that $\Pr_{v_0}^\sigma(\text{Reach}(T')) \leq \Pr_{v_0}^\sigma(\text{Reach}(T))$ for any strategy σ in $\mathcal{G}(\bar{x})$ and any subset $T' \subseteq T$. Q.E.D.

Since any SMG with ω -regular can effectively be reduced to one with Muller objectives, Theorem 10 implies the decidability of QualNE for games with arbitrary ω -regular objectives (e.g. given by SIS formulae). Regarding games with Muller objectives, a matching PSPACE-hardness result appeared in [18], where it was shown that the qualitative decision problem for 2SGs with Muller objectives is PSPACE-hard, even for games without stochastic vertices. However, this result relies on the use of arbitrary colourings.

With similar arguments as for games with Muller objectives, we can show that QualNE is in NP for games with Streett objectives, and in coNP for games with Rabin objectives. A matching NP-hardness result for games with Streett objectives was proven in [25], and the proof of this result can easily be modified to prove coNP-hardness for games with Rabin objectives; both hardness results hold for games with only two players and without stochastic vertices.

Theorem 11. QualNE is NP-complete for games with Streett objectives, and coNP-complete for games with Rabin objectives.

Since any parity condition can be turned into both a Streett and a Rabin condition where the number of pairs is linear in the number of priorities, we can immediately infer from Theorem 11 that QualNE is in $\text{NP} \cap \text{coNP}$ for games with parity objectives.

Corollary 12. QualNE is in $\text{NP} \cap \text{coNP}$ for games with parity objectives.

It is a major open problem whether the qualitative decision problem for 2SGs with parity objectives is in P. This would imply that QualNE is decidable in polynomial time for games with parity objectives since this would allow us to compute the domain of the MDP $\mathcal{G}(\bar{x})$ in polynomial time. For each $d \in \mathbb{N}$, a class of games where the qualitative decision problem is provably in P is the class of all 2SGs with parity objectives that uses at most d priorities [5]. For $d = 2$, this class includes all 2SGs with a Büchi or a co-Büchi objective (for player 0). Hence, we have the following theorem.

Theorem 13. For each $d \in \mathbb{N}$, QualNE is in P for games with parity winning conditions that use at most d priorities. In particular, QualNE is in P for games with (co-)Büchi objectives.

6 Conclusion

We have analysed the complexity of deciding whether a stochastic multiplayer game with ω -regular objectives has a Nash equilibrium whose payoff falls into a certain interval. Specifically, we have isolated several decidable restrictions of the general problem that have a manageable complexity (PSPACE at most). For instance, the complexity of the qualitative variant of NE is usually not higher than for the corresponding problem for two-player zero-sum games.

Apart from settling the complexity of NE (where arbitrary mixed strategies are allowed), two directions for future work come to mind: First, one could study other restrictions of NE that might be decidable. For example, it seems plausible that the restriction of NE to games with two players is decidable. Second, it seems interesting to see whether our decidability results can be extended to more general models of games, e.g. *concurrent games* or games with infinitely many states like *pushdown games*.

References

1. E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen & P. B. Miltersen. On the complexity of numerical analysis. In *CCC '06*, pp. 331–339. IEEE Computer Society Press, 2006.

2. J. Canny. Some algebraic and geometric computations in PSPACE. In *STOC '88*, pp. 460–469. ACM Press, 1988.
3. K. Chatterjee. *Stochastic ω -regular games*. PhD thesis, U.C. Berkeley, 2007.
4. K. Chatterjee, L. de Alfaro & T. A. Henzinger. The complexity of stochastic Rabin and Streett games. In *ICALP 2005*, vol. 3580 of *LNCS*, pages 878–890. Springer-Verlag, 2005.
5. K. Chatterjee, M. Jurdziński & T. A. Henzinger. Simple stochastic parity games. In *CSL 2003*, vol. 2803 of *LNCS*, pp. 100–113. Springer-Verlag, 2003.
6. K. Chatterjee, M. Jurdziński & T. A. Henzinger. Quantitative stochastic parity games. In *SODA 2004*, pp. 121–130. ACM Press, 2004.
7. K. Chatterjee, R. Majumdar & M. Jurdziński. On Nash equilibria in stochastic games. In *CSL 2004*, vol. 3210 of *LNCS*, pp. 26–40. Springer-Verlag, 2004.
8. X. Chen & X. Deng. Settling the complexity of two-player Nash equilibrium. In *FOCS 2006*, pp. 261–272. IEEE Computer Society Press, 2006.
9. A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.
10. C. A. Courcoubetis & M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.
11. C. A. Courcoubetis & M. Yannakakis. Markov decision processes and regular events. *IEEE Transactions on Automatic Control*, 43(10):1399–1418, 1998.
12. C. Daskalakis, P. W. Goldberg & C. H. Papadimitriou. The complexity of computing a Nash equilibrium. In *STOC 2006*, pp. 71–78. ACM Press, 2006.
13. L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
14. L. de Alfaro & T. A. Henzinger. Concurrent omega-regular games. In *LICS 2000*, pp. 141–154. IEEE Computer Society Press, 2000.
15. E. A. Emerson & C. S. Jutla. The complexity of tree automata and logics of programs (extended abstract). In *FoCS '88*, pp. 328–337. IEEE Computer Society Press, 1988.
16. K. Etessami, M. Z. Kwiatkowska, M. Y. Vardi & M. Yannakakis. Multi-objective model checking of Markov decision processes. *Logical Methods in Computer Science*, 4(4), 2008.
17. F. Horn & H. Gimbert. Optimal strategies in perfect-information stochastic games with tail winning conditions. *CoRR*, abs/0811.3978, 2008.
18. P. Hunter & A. Dawar. Complexity bounds for regular games. In *MFCSS 2005*, vol. 3618 of *LNCS*, pp. 495–506. Springer-Verlag, 2005.
19. D. A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
20. J. F. Nash Jr. Equilibrium points in N -person games. *Proceedings of the National Academy of Sciences of the USA*, 36:48–49, 1950.
21. A. Neyman & S. Sorin, editors. *Stochastic Games and Applications*, vol. 570 of *NATO Science Series C*. Springer-Verlag, 2003.
22. M. J. Osborne & A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
23. M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, 1994.
24. W. Thomas. On the synthesis of strategies in infinite games. In *STACS '95*, vol. 900 of *LNCS*, pages 1–13. Springer-Verlag, 1995.

25. M. Ummels. The complexity of Nash equilibria in infinite multiplayer games. In *FOSSACS 2008*, vol. 4962 of *LNCS*, pp. 20–34. Springer-Verlag, 2008.
26. M. Ummels & D. Wojtczak. The complexity of Nash equilibria in simple stochastic multiplayer games. In *ICALP 2009(II)*, vol. 5556 of *LNCS*, pp. 297–308. Springer-Verlag, 2009.
27. M. Ummels & D. Wojtczak. Decision problems for Nash equilibria in stochastic games. *CoRR*, abs/0904.3325, 2009.