The Complexity of Nash Equilibria in Simple Stochastic Multiplayer Games[*]

Michael Ummels[1] and Dominik Wojtczak[2,3]

[1] RWTH Aachen University, Germany
`ummels@logic.rwth-aachen.de`
[2] CWI, Amsterdam, The Netherlands
`d.k.wojtczak@cwi.nl`
[3] University of Edinburgh, UK

Abstract. We analyse the computational complexity of finding Nash equilibria in simple stochastic multiplayer games. We show that restricting the search space to equilibria whose payoffs fall into a certain interval may lead to undecidability. In particular, we prove that the following problem is undecidable: Given a game $\mathcal{G}$, does there exist a pure-strategy Nash equilibrium of $\mathcal{G}$ where player 0 wins with probability 1. Moreover, this problem remains undecidable if it is restricted to strategies with (unbounded) finite memory. However, if mixed strategies are allowed, decidability remains an open problem. One way to obtain a provably decidable variant of the problem is to restrict the strategies to be positional or stationary. For the complexity of these two problems, we obtain a common lower bound of NP and upper bounds of NP and PSPACE respectively.

1 Introduction

We study *stochastic games* [18] 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. The play ends when it reaches a terminal vertex, in which case each player receives a payoff. In the simplest case, which we discuss here, the possible payoffs of a single play are just 0 and 1 (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.

Stochastic games 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, traditionally, the research in this area

---

has concentrated 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* [17]. In a Nash equilibrium, no player can improve her payoff by unilaterally switching to a different strategy. Chatterjee & al. [6] showed that any simple stochastic multiplayer game has a Nash equilibrium, and they also gave an algorithm for computing one. 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 Condon's *SSG Problem* [8] asking whether in a two-player, zero-sum game one of the two players, say player 0, has a strategy to win the game with probability at least $\frac{1}{2}$.

Our main result is that NE is undecidable if only pure strategies are considered. In fact, even the following, presumably simpler, problem is undecidable: Given a game $\mathcal{G}$, decide whether there exists a pure Nash equilibrium where player 0 wins almost surely. Moreover, the problem remains undecidable if one restricts to pure strategies that use (unbounded) finite memory. However, for the general case of arbitrary mixed strategies, decidability remains an open problem. If one restricts to simpler types of strategies like stationary ones, the problem becomes provably decidable. In particular, for positional (i.e. pure, stationary) strategies the problem becomes NP-complete, and for arbitrary stationary strategies the problem is NP-hard but contained in PSPACE. We also relate the complexity of the latter problem to the complexity of the infamous *Square Root Sum Problem* (SqrtSum) by providing a polynomial-time reduction from SqrtSum to NE with the restriction to stationary strategies.

Let us remark that our game model is rather restrictive: First, players receive a payoff only at terminal vertices. In the literature, a plethora of game models with more complicated modes of winning have been discussed. In particular, the model of a *stochastic parity game* [5,24] has been investigated thoroughly. Second, our model is *turn-based* (i.e. for every non-stochastic vertex there is only one player who controls this vertex) as opposed to *concurrent* [12,11]. The reason that we have chosen to analyse such a restrictive model is that we are focussing on negative results. Indeed, all our lower bounds hold for (multiplayer versions of) the aforementioned models. Moreover, besides Nash equilibria, our negative

results apply to several other solution concepts like subgame perfect equilibria [20,21] and secure equilibria [4].

For games with *rewards* on transitions [15], the situation might be different: While our lower bounds can be applied to games with rewards under the *average reward* or the *total expected reward* criterion, we leave it as an open question whether this remains true in the case of discounted rewards.

Due to space constraints, most proofs are either only sketched or omitted entirely. For the complete proofs, see [23].

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 [10,7]. More in the spirit of our work, Conitzer and Sandholm [9] showed that deciding whether there exists a Nash equilibrium in a two-player game in strategic form where player 0 receives payoff at least x and related decision problems are all NP-hard. For infinite games without stochastic vertices, (a qualitative version of) the problem NE was studied in [22]. In particular, it was shown that the problem is NP-complete for games with parity winning conditions and even in P for games with Büchi winning conditions.

For stochastic games, most results concern the classical SSG problem: Condon showed that the problem is in NP ∩ co-NP [8], but it is not known to be in P. We are only aware of two results that are closely related to our problem: First, Etessami & al. [13] investigated Markov decision processes with, e.g., multiple reachability objectives. Such a system can be viewed as a stochastic multiplayer game where all non-stochastic vertices are controlled by one single player. Under this interpretation, one of their results states that NE is decidable in polynomial time for such games. Second, Chatterjee & al. [6] showed that the problem of deciding whether a (concurrent) stochastic game with reachability objectives has a positional-strategy Nash equilibrium with payoff at least $\bar{x}$ is NP-complete. We sharpen their hardness result by showing that the problem remains NP-hard when it is restricted to games with only three players (as opposed to an unbounded number of players) where, additionally, payoffs are assigned at terminal vertices only (cf. Theorem 5).

2   Simple stochastic multiplayer games

The model of a *(two-player, zero-sum) simple stochastic game* [8] easily generalises to the multiplayer case: Formally, a *simple stochastic multiplayer game (SSMG)* is a tuple $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, \Delta, (F_i)_{i \in \Pi})$ such that:

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

– $F_i \subseteq V$ for each $i \in \Pi$.

We call a vertex $v \in V_i$ *controlled by player i* and a vertex that is not contained in any of the sets V_i a *stochastic* vertex. 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 $p = \bot$ if $v \in V_i$ for some $i \in \Pi$. Moreover, 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 $(v, p, w) \in \Delta$. For computational purposes, we require additionally that all these probabilities are rational.

For a given vertex $v \in V$, we denote the set of all $w \in V$ such that there exists $p \in (0, 1] \cup \{\bot\}$ with $(v, p, w) \in \Delta$ by $v\Delta$. For technical reasons, we require that $v\Delta \neq \emptyset$ for all $v \in V$. 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 each vertex v that lies in one of the sets F_i is a *terminal (sink) vertex*: $v\Delta = \{v\}$. So if F is the set of all terminal vertices, then $F_i \subseteq F$ for each $i \in \Pi$.

A *(mixed) strategy of player i in $\mathcal{G}$* is a mapping $\sigma : V^*V_i \to \mathcal{D}(V)$ 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 $(v, \bot, w) \in \Delta$. 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 $\sigma : V^*V_i \to V$. 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 $\sigma : V_i \to \mathcal{D}(V)$. 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 $\sigma : V_i \to V$. 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.[1] 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 SSMG*. 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 SSMG also for initialised SSMGs. It should always be clear from the context if the game is initialised or not.

---

[1] In general, this definition is applicable to mixed strategies as well, but for this paper we will identify finite-state strategies with pure finite-state strategies.

Given an SSMG $(\mathcal{G}, v_0)$ 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 $v_1 \ldots v_k \cdot V^\omega$ is 0 if $v_1 \neq v_0$ and the product of the probabilities $\bar{\sigma}(v_j \mid v_1 \ldots v_{j-1})$ for $j = 2, \ldots, k$ 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 set $U \subseteq V$, let $\mathrm{Reach}(U) := V^* \cdot U \cdot V^\omega$. We are mainly interested in the probabilities $p_i := \Pr_{v_0}^{\bar{\sigma}}(\mathrm{Reach}(F_i))$ of reaching the sets F_i. We call the number p_i the *(expected) payoff of $\bar{\sigma}$ for player i* and the vector $(p_i)_{i \in \Pi}$ the *(expected) payoff of $\bar{\sigma}$*.

Drawing an SSMG. When drawing an SSMG as a graph, we will use the following conventions: The initial vertex is marked by an incoming edge that has no source vertex. Vertices that are controlled by a player are depicted as circles, where the player who controls a vertex is given by the label next to it. Stochastic vertices are depicted as diamonds, where the transition probabilities are given by the labels on its outgoing edges (the default being $\frac{1}{2}$). Finally, terminal vertices are represented by their associated payoff vector. In fact, we allow arbitrary vectors of rational probabilities as payoffs. This does not increase the power of the model since such a payoff vector can easily be realised by an SSMG consisting of stochastic and terminal vertices only.

3   Nash equilibria

To capture rational behaviour of (selfish) players, John Nash [17] introduced the notion of, what is now called, a *Nash equilibrium*. Formally, given a strategy profile $\bar{\sigma}$, a strategy τ of player i is called a *best response to $\bar{\sigma}$* if τ maximises the expected payoff of player i: $\Pr_{v_0}^{(\bar{\sigma}_{-i}, \tau')}(\mathrm{Reach}(F_i)) \leq \Pr_{v_0}^{(\bar{\sigma}_{-i}, \tau)}(\mathrm{Reach}(F_i))$ 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.

Previous research on algorithms for finding Nash equilibria in infinite games has focused on computing *some* Nash equilibrium [6]. 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 us to the following decision problem, which we call NE:[2]

Given an SSMG $(\mathcal{G}, v_0)$ and thresholds $\bar{x}, \bar{y} \in [0, 1]^\Pi$, decide whether there exists a Nash equilibrium of $(\mathcal{G}, v_0)$ with payoff $\geq \bar{x}$ and $\leq \bar{y}$.

---

[2] In the definition of NE, the ordering $\leq$ is applied componentwise.

For computational purposes, we assume that the thresholds $\bar{x}$ and $\bar{y}$ are vectors of rational numbers. A variant of the problem which omits the thresholds just asks about a Nash equilibrium where some distinguished player, say player 0, wins with probability 1:

> Given an SSMG $(\mathcal{G}, v_0)$, decide whether there exists a Nash equilibrium of $(\mathcal{G}, v_0)$ where player 0 wins almost surely.

Clearly, every instance of the threshold-free variant can easily be turned into an instance of NE (by adding the thresholds $\bar{x} = (1, 0, \ldots, 0)$ and $\bar{y} = (1, \ldots, 1)$). Hence, NE is, a priori, more general than its threshold-free variant.

Our main concern in this paper are variants of NE where we restrict the type of strategies that are allowed in the definition of the problem: Let PureNE, FinNE, StatNE and PosNE be the problems that arise from NE by restricting the desired Nash equilibrium to consist of pure strategies, finite-state strategies, stationary strategies and positional strategies, respectively. In the rest of this paper, we are going to prove upper and lower bounds on the complexity of these problems, where all lower bounds hold for the threshold-free variants, too.

Our first observation is that neither stationary nor pure strategies are sufficient to implement any Nash equilibrium, even if we are only interested in whether a player wins or loses almost surely in the Nash equilibrium. Together with a result from Sect. 5 (namely Proposition 8), this demonstrates that the problems NE, PureNE, FinNE, StatNE, and PosNE are pairwise distinct problems, which have to be analysed separately.

Proposition 1. *There exists an SSMG that has a finite-state Nash equilibrium where player 0 wins almost surely but that has no stationary Nash equilibrium where player 0 wins with positive probability.*

Proof. Consider the game $\mathcal{G}$ depicted in Fig. 1 (a) played by three players 0, 1 and 2 (with payoffs in this order). Obviously, the following finite-state strategy profile is a Nash equilibrium where player 0 wins almost surely: Player 1 plays from vertex v_2 to vertex v_3 at the first visit of v_2 but leaves the game immediately (by playing to the neighbouring terminal vertex) at all subsequent visits to v_2; from vertex v_0 player 1 plays to v_1; player 2 plays from vertex v_3 to vertex v_4 at the first visit of v_3 but leaves the game immediately at all subsequent visits to v_3; from vertex v_1 player 2 plays to v_2. It remains to show that there is no stationary Nash equilibrium of $(\mathcal{G}, v_0)$ where player 0 wins with positive probability: Any such equilibrium induces a stationary Nash equilibrium of $(\mathcal{G}, v_2)$ where both players 1 and 2 receive payoff at least $\frac{1}{2}$ since otherwise one of these players could improve her payoff by changing her strategy at v_0 or v_1. However, it is easy to see that in any stationary Nash equilibrium of $(\mathcal{G}, v_2)$ either player 1 or player 2 receives payoff 0. $\qquad\square$

Proposition 2. *There exists an SSMG that has a stationary Nash equilibrium where player 0 wins almost surely but that has no pure Nash equilibrium where player 0 wins with positive probability.*
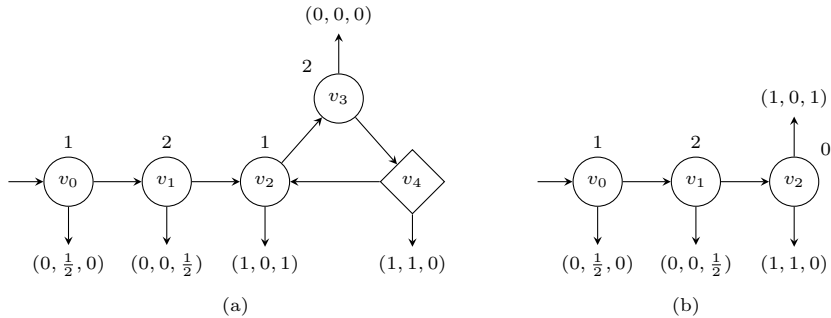
Fig. 1. Two SSMGs with three players

Proof. Consider the game depicted in Fig. 1 (b) played by three players 0, 1 and 2 (with payoffs given in this order). Clearly, the stationary strategy profile where from vertex v_2 player 0 selects both outgoing edges with probability $\frac{1}{2}$ each, player 1 plays from v_0 to v_1 and player 2 plays from v_1 to v_2 is a Nash equilibrium where player 0 wins almost surely. However, for any pure strategy profile where player 0 wins almost surely, either player 1 or player 2 receives payoff 0 and could improve her payoff by switching her strategy at v_0 or v_1 respectively. $\qquad\square$

4 Decidable variants of NE

In this section, we show that the problems PosNE and StatNE are decidable and analyse their complexity.

Theorem 3. *PosNE is in* NP.

Proof (Sketch). Let $(\mathcal{G}, v_0)$ be an SSMG. Any positional strategy profile of $\mathcal{G}$ can be identified with a mapping $\bar{\sigma} : \bigcup_{i \in \Pi} V_i \to V$ such that $(v, \bot, \bar{\sigma}(v)) \in \Delta$ for each non-stochastic vertex v, an object whose size is linear in the size of $\mathcal{G}$. To prove that PosNE is in NP, it suffices to show that we can check in polynomial time whether such a mapping $\bar{\sigma}$ constitutes a Nash equilibrium whose payoff lies in between the given thresholds $\bar{x}$ and $\bar{y}$. This can be done by computing, for each player i, 1. the payoff z_i of $\bar{\sigma}$ for player i and 2. the maximal payoff $r_i = \sup_\tau \Pr_{v_0}^{(\bar{\sigma}_{-i}, \tau)}(\text{Reach}(F_i))$ that player i can achieve when playing against $\bar{\sigma}_{-i}$, and then to check whether $x_i \le z_i \le y_i$ and $r_i \le z_i$. It follows from results on *Markov chains* and *Markov decision processes* that both these numbers can be computed in polynomial time (via linear programming). $\qquad\square$

To prove the decidability of StatNE, we appeal to results established for the *Existential Theory of the Reals*, ExTh($\mathfrak{R}$), the set of all existential first-order sentences (over the appropriate signature) that hold in $\mathfrak{R} := (\mathbb{R}, +, \cdot, 0, 1, \le)$. The best known upper bound for the complexity of the associated decision problem is PSPACE [3,19], which leads to the following theorem.

Theorem 4. *StatNE is in* PSPACE.

Proof (Sketch). Since PSPACE = NPSPACE, it suffices to give a nondeterministic polynomial-space algorithm for StatNE. On input $\mathcal{G}, v_0, \bar{x}, \bar{y}$, the algorithm starts by guessing a set $S \subseteq V \times V$ and proceeds by computing (in polynomial time), for each player i, the set R_i of vertices from where the set F_i is reachable in the graph $G = (V, S)$. Finally, the algorithm evaluates a certain existential first-order sentence ψ, which can be computed in polynomial time from $(\mathcal{G}, v_0)$, $\bar{x}$, $\bar{y}$, S and $(R_i)_{i \in \Pi}$, over $\mathfrak{R}$ and returns the answer to this query. The sentence ψ states that there exists a stationary Nash equilibrium $\bar{\sigma}$ of $(\mathcal{G}, v_0)$ with payoff $\geq \bar{x}$ and $\leq \bar{y}$ whose *support* is S, i.e. $S = \{(v, w) \in V \times V : \bar{\sigma}(w \mid v) > 0\}$. $\qquad\square$

Having shown that PosNE and StatNE are in NP and PSPACE respectively, the natural question arises whether there is a polynomial-time algorithm for PosNE or StatNE. The following theorem implies that this is not the case (unless, of course, P = NP) since both problems are NP-hard. Moreover, both problems are already NP-hard for games with only two players.

Theorem 5. *PosNE and StatNE are* NP-*hard, even for games with only two players (three players for the threshold-free variants).*

It follows from Theorems 3 and 5 that PosNE is NP-complete. For StatNE, we have provided an NP lower bound and a PSPACE upper bound, but the exact complexity of the problem remains unclear. Towards gaining more insight into the problem StatNE, we relate its complexity to the complexity of the *Square Root Sum Problem* (SqrtSum), the problem of deciding, given numbers $d_1, \ldots, d_n, k \in \mathbb{N}$, whether $\sum_{i=1}^{n} \sqrt{d_i} \geq k$. Recently, it was shown that SqrtSum belongs to the 4th level of the *counting hierarchy* [1], which is a slight improvement over the previously known PSPACE upper bound. However, it is an open question since the 1970s whether SqrtSum falls into the polynomial hierarchy [16,14]. We identify a polynomial-time reduction from SqrtSum to StatNE.[3] Hence, StatNE is at least as hard as SqrtSum, and showing that StatNE resides inside the polynomial hierarchy would imply a major breakthrough in understanding the complexity of numerical computation.

Theorem 6. *SqrtSum is polynomial-time reducible to StatNE.*

5   Undecidable variants of NE

In this section, we argue that the problems PureNE and FinNE are undecidable by exhibiting reductions from two undecidable problems about *two-counter machines*. Our construction is inspired by a construction used by Brázdil & al. [2] to prove the undecidability of stochastic games with branching-time winning conditions.

A two-counter machine $\mathcal{M}$ is given by a list of instructions $\iota_1, \ldots, \iota_m$ where each instruction is one of the following:

---

[3] Some authors define SqrtSum with $\leq$ instead of $\geq$. With this definition, we would reduce from the complement of SqrtSum instead.

- "inc(j); goto k"   (increment counter j by 1 and go to instruction number k);
- "zero(j) ? goto k : dec(j); goto l"   (if the value of counter j is zero, go to instruction number k; otherwise, decrement counter j by one and go to instruction number l);
- "halt"   (stop the computation).

Here j ranges over $1, 2$ (the two counters), and $k \neq l$ range over $1, \ldots, m$. A configuration of $\mathcal{M}$ is a triple $C = (i, c_1, c_2) \in \{1, \ldots, m\} \times \mathbb{N} \times \mathbb{N}$, where i denotes the number of the current instruction and c_j denotes the current value of counter j. A configuration C' is the *successor* of configuration C, denoted by $C \vdash C'$, if it results from C by executing instruction ι_i; a configuration $C = (i, c_1, c_2)$ with $\iota_i = $ "halt" has no successor configuration. Finally, the *computation of* $\mathcal{M}$ is the unique maximal sequence $\rho = \rho(0)\rho(1)\ldots$ such that $\rho(0) \vdash \rho(1) \vdash \ldots$ and $\rho(0) = (1, 0, 0)$ (the *initial configuration*). Note that ρ is either infinite, or it ends in a configuration $C = (i, c_1, c_2)$ such that $\iota_i = $ "halt".

The *halting problem* is to decide, given a machine $\mathcal{M}$, whether the computation of $\mathcal{M}$ is finite. It is well-known that two-counter machines are Turing powerful, which makes the halting problem and its dual, the *non-halting problem*, undecidable.

Theorem 7. *PureNE is undecidable.*

Proof (Sketch). The proof is by a reduction from the non-halting problem to PureNE: we show that one can compute from a two-counter machine $\mathcal{M}$ an SSMG $(\mathcal{G}, v_0)$ with nine players such that the computation of $\mathcal{M}$ is infinite iff $(\mathcal{G}, v_0)$ has a pure Nash equilibrium where player 0 wins almost surely.

Any pure strategy profile $\bar{\sigma}$ of $(\mathcal{G}, v_0)$ where player 0 wins almost surely determines an infinite sequence ρ of *pseudo configurations* of $\mathcal{M}$ (where the counters may take the value ω). Of course, in general, ρ is not the computation of $\mathcal{M}$. However, the game $\mathcal{G}$ is constructed in such a way that $\bar{\sigma}$ is a Nash equilibrium iff ρ is the computation of $\mathcal{M}$. Since ρ is infinite, this equivalence implies that $\mathcal{G}$ has a pure Nash equilibrium where player 0 wins almost surely iff $\mathcal{M}$ does not halt.

To get a flavour of the full proof, let us consider a machine $\mathcal{M}$ that contains the instruction "inc(1); goto k". An abstraction of the corresponding part of the game $\mathcal{G}$ is depicted in Fig. 2: the game is restricted to three players 0, A and B (with payoffs given in this order), and some irrelevant vertices have been removed.

In the following, let $\bar{\sigma}$ be a pure strategy profile of $(\mathcal{G}, v)$ where player 0 wins almost surely. At both vertices u and u', player 0 can play to either a grey or a white vertex; if she plays to a grey vertex, then with probability $\frac{1}{2}$ the play returns to u or u' respectively; if she plays to the white vertex, the play never returns to u or u'. Let c and c' denote the maximal number of visits to the grey vertex connected to u and u' respectively (the number being ω if player 0 always plays to the grey vertex): these two ordinal numbers represent the counter value before and after executing the instruction.
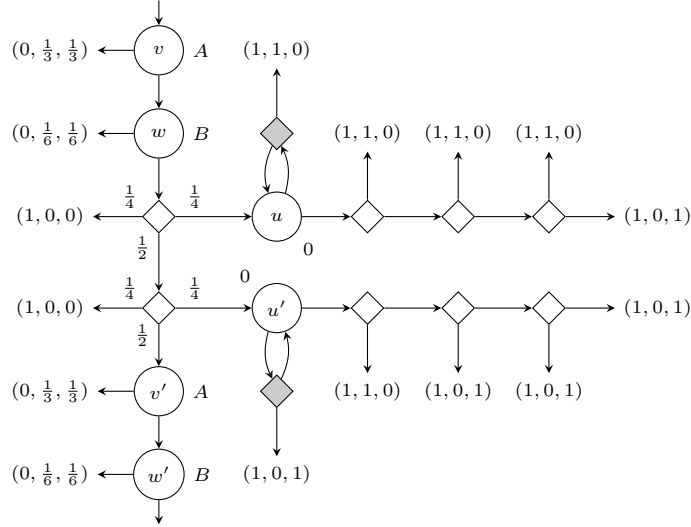
Fig. 2. Incrementing a counter.

To see why $c' = 1 + c$ if $\bar{\sigma}$ is a Nash equilibrium, consider the probabilities $a := \Pr_v^{\bar{\sigma}}(\mathrm{Reach}(F_A))$ and $b := \Pr_v^{\bar{\sigma}}(\mathrm{Reach}(F_B))$. We have $a \geq \frac{1}{3}$ and $b \geq \frac{1}{6}$ since otherwise player A or B could improve by changing her strategy at vertex v or w respectively. In fact, the construction of $\mathcal{G}$ ensures that $a + b = \frac{1}{2}$; hence, $a = \frac{1}{3}$. Moreover, the same argumentation proves that $a' := \Pr_{v'}^{\bar{\sigma}}(\mathrm{Reach}(F_A)) = \frac{1}{3}$.

Let $p := \Pr_v^{\bar{\sigma}}(\mathrm{Reach}(F_A) \mid V^\omega \backslash v \ldots v' \cdot V^\omega)$ be the conditional probability that player A wins given that v' is not reached; then $a = p + \frac{1}{4} \cdot a'$ and consequently $p = \frac{1}{4}$. But p can also be written as the following sum of two binary numbers:

$$0.00 \underbrace{1 \ldots 1}_{c \text{ times}} 111 + 0.000 \underbrace{0 \ldots 0}_{c' \text{ times}} 100 \,.$$

Obviously, this sum is equal to $\frac{1}{4}$ iff $c' = 1 + c$. $\qquad\qquad\square$

It follows from the proof of Theorem 7 that Nash equilibria may require infinite memory (even if we are only interested in whether a player wins with probability 0 or 1). More precisely, we have the following proposition.

Proposition 8. *There exists an SSMG that has a pure Nash equilibrium where player 0 wins almost surely but that has no finite-state Nash equilibrium where player 0 wins with positive probability.*

Proof. Consider the game $(\mathcal{G}, v_0)$ constructed in the proof of Theorem 7 for the machine $\mathcal{M}$ consisting of the single instruction "inc(1); goto 1". We modify this game by adding a new initial vertex v_1 which is controlled by a new player, player 1, and from where she can either move to v_0 or to a new terminal vertex where she receives payoff 1 and every other player receives payoff 0. Additionally,

player 1 wins at every terminal vertex of the game $\mathcal{G}$ that is winning for player 0. Let us denote the modified game by $\mathcal{G}'$.

Since the computation of $\mathcal{M}$ is infinite, the game $(\mathcal{G}, v_0)$ has a pure Nash equilibrium where player 0 wins almost surely. This equilibrium induces a pure Nash equilibrium of $(\mathcal{G}', v_1)$ where player 0 wins almost surely. However, it is easy to see that there is no finite-state Nash equilibrium of $(\mathcal{G}, v_0)$ where player 0 (or player 1) wins almost surely. Consequently, in any finite-state Nash equilibrium of $(\mathcal{G}', v_1)$ player 1 will play from v_1 to the new terminal vertex, giving player 0 a payoff of 0. □

It follows from Proposition 8 that the problems FinNE and PureNE are distinct. Nevertheless, we can show that FinNE is undecidable, too. Note however that FinNE is recursively enumerable: To decide whether an SSMG $(\mathcal{G}, v_0)$ has a finite-state Nash equilibrium with payoff $\geq \bar{x}$ and $\leq \bar{y}$, one can just enumerate all possible finite-state profiles and check for each of them whether the profile is a Nash equilibrium with the desired properties (by analysing the finite Markov chain that is generated by this profile).

Theorem 9. *FinNE is undecidable.*

Proof (Sketch). The proof is by a reduction from the halting problem for two-counter machines and similar to the proof of Theorem 7. □

6 Conclusion

We have analysed the complexity of deciding whether a simple stochastic multiplayer game has a Nash equilibrium whose payoff falls into a certain interval. Our results demonstrate that the presence of both stochastic vertices and more than two players makes the problem much more complicated than when one of these factors is absent. In particular, the problem of deciding the existence of a pure-strategy Nash equilibrium where player 0 wins almost surely is undecidable for simple stochastic multiplayer games, whereas it is contained in NP∩co-NP for two-player, zero-sum simple stochastic games [8] and even in P for non-stochastic infinite multiplayer games with, e.g., Büchi winning conditions [22].

Apart from settling the complexity of NE when arbitrary mixed strategies are considered, future research may, for example, investigate restrictions of NE to games with a small number of players. In particular, we conjecture that the problem is decidable for two-player games, even if these are not zero-sum.

References

1. E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen & P. B. Miltersen. On the complexity of numerical analysis. In *Proc. CCC '06*, pp. 331–339. IEEE Computer Society Press, 2006.
2. T. Brázdil, V. Brožek, V. Forejt & A. Kučera. Stochastic games with branching-time winning objectives. In *Proc. LICS 2006*, pp. 349–358. IEEE Computer Society Press, 2006.

3. J. Canny. Some algebraic and geometric computations in PSPACE. In *Proc. STOC 1988*, pp. 460–469. ACM Press, 1988.

4. K. Chatterjee, T. A. Henzinger & M. Jurdziński. Games with secure equilibria. *Theoretical Computer Science*, 365(1–2):67–82, 2006.

5. K. Chatterjee, M. Jurdziński & T. A. Henzinger. Quantitative stochastic parity games. In *Proc. SODA 2004*, pp. 121–130. ACM Press, 2004.

6. K. Chatterjee, R. Majumdar & M. Jurdziński. On Nash equilibria in stochastic games. In *Proc. CSL 2004*, vol. 3210 of *LNCS*, pp. 26–40. Springer-Verlag, 2004.

7. X. Chen & X. Deng. Settling the complexity of two-player Nash equilibrium. In *Proc. FOCS 2006*, pp. 261–272. IEEE Computer Society Press, 2006.

8. A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.

9. V. Conitzer & T. Sandholm. Complexity results about Nash equilibria. In *Proc. IJCAI 2003*, pp. 765–771. Morgan Kaufmann, 2003.

10. C. Daskalakis, P. W. Goldberg & C. H. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proc. STOC 2006*, pp. 71–78. ACM Press, 2006.

11. L. de Alfaro & T. A. Henzinger. Concurrent omega-regular games. In *Proc. LICS 2000*, pp. 141–154. IEEE Computer Society Press, 2000.

12. L. de Alfaro, T. A. Henzinger & O. Kupferman. Concurrent reachability games. In *Proc. FOCS '98*, pp. 564–575. IEEE Computer Society Press, 1998.

13. 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.

14. K. Etessami & M. Yannakakis. On the complexity of Nash equilibria and other fixed points. In *Proc. FOCS 2007*, pp. 113–123. IEEE Computer Society Press, 2007.

15. J. Filar & K. Vrieze. *Competitive Markov decision processes*. Springer-Verlag, 1997.

16. M. R. Garey, R. L. Graham & D. S. Johnson. Some NP-complete geometric problems. In *Proc. STOC '76*, pp. 10–22. ACM Press, 1976.

17. J. F. Nash Jr. Equilibrium points in N-person games. *Proc. National Academy of Sciences of the USA*, 36:48–49, 1950.

18. A. Neyman & S. Sorin, eds. *Stochastic Games and Applications*, vol. 570 of *NATO Science Series C*. Springer-Verlag, 1999.

19. J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. *Journal of Symbolic Computation*, 13(3):255–352, 1992.

20. R. Selten. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*, 121:301–324 and 667–689, 1965.

21. M. Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In *Proc. FSTTCS 2006*, vol. 4337 of *LNCS*, pp. 212–223. Springer-Verlag, 2006.

22. M. Ummels. The complexity of Nash equilibria in infinite multiplayer games. In *Proc. FOSSACS 2008*, vol. 4962 of *LNCS*, pp. 20–34. Springer-Verlag, 2008.

23. M. Ummels & D. Wojtczak. The complexity of Nash equilibria in simple stochastic multiplayer games. Technical report, University of Edinburgh, 2009.

24. W. Zielonka. Perfect-information stochastic parity games. In *Proc. FOSSACS 2004*, vol. 2987 of *LNCS*, pp. 499–513. Springer-Verlag, 2004.