Stochastic Games with Branching-Time Winning Objectives

Tomáš Brázdil Václav Brožek Vojtěch Forejt Antonín Kučera
Faculty of Informatics, Masaryk University,
Botanická 68a, 60200 Brno,
Czech Republic.
{brazdil,xbrozek,xforejt,kucera}@fi.muni.cz

Abstract

We consider stochastic turn-based games where the winning objectives are given by formulae of the branching-time logic PCTL. These games are generally not determined and winning strategies may require memory and/or randomization. Our main results concern history-dependent strategies. In particular, we show that the problem whether there exists a history-dependent winning strategy in $1\frac{1}{2}$-player games is highly undecidable*, even for objectives formulated in the $\mathcal{L}(\mathrm{F}^{=5/8}, \mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ fragment of PCTL. On the other hand, we show that the problem becomes decidable (and in fact* ***EXPTIME****-complete) for the $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ fragment of PCTL, where winning strategies require only finite memory. This result is tight in the sense that winning strategies for $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1}, \mathrm{G}^{>0})$ objectives may already require infinite memory.*

1. Introduction

In this paper we consider stochastic turn-based games where the winning objectives are given by formulae of the branching-time logic PCTL. Formally, a $2\frac{1}{2}$-player game G is a finite directed graph where the vertices are partitioned into three subsets $V_\Box$, $V_\Diamond$, $V_\bigcirc$. A *play* is initiated by putting a token on some vertex. The token is then moved from vertex to vertex by two players, $\Box$ and $\Diamond$, who are responsible for selecting outgoing transitions in the vertices of $V_\Box$ and $V_\Diamond$, respectively. In the vertices of $V_\bigcirc$, outgoing transitions are chosen randomly according to a fixed probability distribution. A *strategy* specifies how a player should play. In general, a strategy may or may not depend on the history of a play (we say that a strategy is *history-dependent (H)* or *memoryless (M)*), and the transitions may be chosen deterministically or randomly (*deterministic (D)* and *randomized (R)* strategies). In the case of randomized strategies, a player chooses a probability distribution on the set of outgoing transitions. Note that deterministic strategies can be seen as restricted randomized strategies, where one of the outgoing transitions has probability 1. Each pair of strategies (σ, π) for players $\Box$ and $\Diamond$ determines a unique Markov chain $G(\sigma, \pi)$ where the states are finite paths in G, and $wu \to wuu'$ with probability x iff (u, u') is a transition in the game and x is the probability chosen by player $\Box$ or $\Diamond$ (when $u \in V_\Box$ or $u \in V_\Diamond$, respectively), or the fixed probability of the transition (u, u') when $u \in V_\bigcirc$. A *winning objective* for player $\Box$ is some property of Markov chains that is to be achieved. A *winning strategy* for player $\Box$ is a strategy σ such that for every strategy π of player $\Diamond$ the Markov chain $G(\sigma, \pi)$ has the desired property. Usually, the aim of player $\Diamond$ is to *falsify* this property, which means that his winning objective is dual. A winning strategy for player $\Diamond$ is a strategy π such that $G(\sigma, \pi)$ does *not* have the property for any strategy σ of player $\Box$. A game is *determined* if one of the two players has a winning strategy in every vertex. $1\frac{1}{2}$-player games are "restricted" $2\frac{1}{2}$-player games where $V_\Diamond = \emptyset$. All of the above introduced notions (except for determinacy) are applicable also to $1\frac{1}{2}$-player games.

Infinite games have been studied in various fields of mathematics and computer science (recently written overviews are, e.g., [16, 6]). For example, model-checking problems for certain temporal logics (such as the modal μ-calculus) can be naturally reformulated as the questions to determine the winner in parity games, and a lot of research effort has been invested into this problem. Our work is mainly motivated by applications of games in system design, where systems are modeled as games, player $\Box$ corresponds to a "controller" which determines the system behaviour in a subset of controllable states, player $\Diamond$ models the environment, and the winning objectives for player $\Box$ correspond to the desired property of the system. The task is to find a controller (a strategy σ for player $\Box$) such that the desired property holds no matter what the environment does (i.e., the strategy σ is winning). As for stochastic games, the majority of existing results concern games with *linear time* winning objectives which are specified by some property of runs in Markov chains. Examples include quantitative reachability objectives (the probability of all runs that hit a "good" state is at least ϱ), qualitative Büchi objectives (the probability of all runs along which a "good" vertex appears infinitely often is 1), qualitative/quantitative parity objectives [7, 8], Rabin and Street objectives [5], etc. In this paper we study *branching-time* objectives that are formalized as formulae of the branching-time probabilistic logic PCTL.

Previous and related work. In [1], it is shown that winning strategies for PCTL objectives may require memory and/or randomization in general. Hence, the MD, MR, HD, and HR strategies (see above) need to be considered separately. It is also proven that the problem whether there exists a winning MD strategy in a given $1\frac{1}{2}$-player game for a given PCTL objective is **NP**-complete. MR strategies were considered in [14], where it is shown that the existence of a winning MR strategy in a given $2\frac{1}{2}$-player game for a given PCTL objective is in **EXPTIME**. The construction also yields **PSPACE** upper bound for $1\frac{1}{2}$-player games.

To prevent misunderstanding, we should say that the logic PCTL can also be interpreted directly on games (or Markov decision processes). The decidability of the model-checking problem for stochastic games and PCTL was established in [10] as a simple consequence of the results about quantitative ω-regular games. However, this is a different problem which is not directly related to the subject of this paper (as we shall, the results about stochastic games with branching-time winning objectives are quite different from the results about model-checking).

Main results. We start by observing that stochastic games with branching-time objectives are not determined, even if the objectives are formulae of the $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0})$ fragment of PCTL (in general, $\mathcal{L}(Y_1, \cdots, Y_n)$ denotes the fragment of PCTL containing the connectives $Y_1, \cdots, Y_n$, conjunction, and disjunction (negation can be applied only to atomic propositions)). As a warm-up, we present some simple results about memoryless strategies in Section 3.1. We show that the problem whether player $\Box$ has a winning MD strategy in a given $2\frac{1}{2}$-player game for a given PCTL objective is $\Sigma_2 = \mathbf{NP}^{\mathbf{NP}}$ complete. The Σ_2 lower bound holds even for the $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0})$ fragment of PCTL. Since the existence of a winning MD strategy for player $\Box$ in $1\frac{1}{2}$-player games with PCTL objectives is **NP**-complete [1], we yield a full complexity classification for MD strategies. The lower complexity bounds carry over to MR strategies and hold even for *qualitative* PCTL objectives for which we give the matching upper bounds—we show that the existence of a winning MR strategy for player $\Box$ in $1\frac{1}{2}$-player (or $2\frac{1}{2}$-player) games with *qualitative* PCTL objectives is **NP**-complete (or $\Sigma_2 = \mathbf{NP}^{\mathbf{NP}}$ complete, resp.). Let us note that randomized strategies are strictly more powerful than deterministic ones even for qualitative objectives (a simple example is given in Section 3.1). The existence of a winning MR strategy for player $\Box$ in $1\frac{1}{2}$-player and $2\frac{1}{2}$-player games with general PCTL objectives is known to be in **PSPACE** and **EXPTIME**, respectively [14]. We did not manage to lift the **NP** and Σ_2 lower bounds, and we also failed to improve the mentioned upper bounds. On the other hand, there is some indication that lowering the bounds below **PSPACE** would be quite difficult. We use the same argument as Etessami & Yannakakis in [12], where it is shown that the SQUARE-ROOT-SUM problem is efficiently reducible to the quantitative reachability problem for one-exit recursive Markov chains. An instance of SQUARE-ROOT-SUM is a tuple $(a_1, \cdots, a_n, b)$ of integers. The question is whether $\sum_{i=1}^{n} \sqrt{a_i} \leq b$. This problem is known to be in **PSPACE**, but its exact complexity is a long-standing open problem in computational geometry. Hence, an efficient reduction of SQUARE-ROOT-SUM to another problem $P \in$ **PSPACE** can be seen as an indication that the complexity of P is hard to improve. We show that SQUARE-ROOT-SUM is efficiently reducible to the problem whether player $\Box$ has a winning MR strategy in $1\frac{1}{2}$-player games with PCTL objectives. Let us note that the technique used in the proof is different from the one of [12].

The main results of this paper concern history-dependent strategies. First, we answer the open question formulated in [1] by showing that the existence of a winning HD (or HR) strategy in $1\frac{1}{2}$-player games is *highly undecidable* even for objectives of the $\mathcal{L}(\mathrm{F}^{=5/8}, \mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ fragment of PCTL. More precisely, we show that the above problem is complete for the Σ_1^1 level of the analytical hierarchy. This is already a deep result relying on specific tricks which were developed to encode and simulate a computation of a given nondeterministic Minsky machine. A slight modification of the proof reveals that the existence of a winning HD (or HR) strategy with *finite memory* in $1\frac{1}{2}$-player games with $\mathcal{L}(\mathrm{F}^{=5/8}, \mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ objectives is also undecidable (and complete for the Σ_1 level of the arithmetical hierarchy). The role of the quantitative $\mathrm{F}^{=5/8}$ operator is very important in these undecidability results[1]. In general, qualitative questions tend to be easier than quantitative ones (this also holds for PCTL and certain classes of infinite-state Markov chains [11, 4, 3]; note that the plays determined by history-dependent strategies are infinite-state Markov chains). Hence, we turn out attention to *qualitative* PCTL objectives. We start by examining the fragments with qualitative forms of reachability and safety connectives, i.e.,

1 Let us note that $5/8$ is not some kind of "magic number", it is just technically convenient. In principle, any operator of the form $\mathrm{F}^{=x}$ where $0 < x < 1$ would suffice for our purposes.

$F^{\bowtie\varrho}$ and $G^{\bowtie\varrho}$, where $\bowtie \in \{=,>,<\}$ and $\varrho \in \{0,1\}$. Even in this simplified setting, the results are not uniform and different combinations of connectives lead to quite different results. First, we show that the role of $F^{=5/8}$ operator in the aforementioned undecidability proof is *provably* crucial in the sense that the existence of a winning HD strategy in $1\frac{1}{2}$-player games with $\mathcal{L}(F^{=1}, F^{>0}, G^{=1})$ objectives is **EXPTIME** complete. Let us note that

- the **EXPTIME** upper bound is proven in two phases. First, we show that the existence of a winning HD strategy in $1\frac{1}{2}$-player games with $\mathcal{L}(F^{=1}, F^{>0}, G^{=1})$ objectives is effectively reducible to the existence of a winning HD strategy in $1\frac{1}{2}$-player games with *mixed linear-time* objectives, which are essentially conjunctions of one *qualitative-Büchi* and one *sure-Büchi* objective. This reduction is exponential. Then, we show that the existence of a winning HD strategy in $1\frac{1}{2}$-player games with mixed linear-time objectives is in **P**. Note that if we had a conjunction of two qualitative-Büchi or two sure-Büchi objectives, we could simply apply known results. To the best of our knowledge, the games where the winning objectives are "mixtures" of stochastic and non-stochastic requirements have not yet been explicitly considered (perhaps due to the lack of motivation). The solution we provide is not trivial.
- The **EXPTIME** lower bound holds even for $\mathcal{L}(F^{=1}, G^{=1})$ objectives and for both HD and HR strategies.

Our construction also reveals that a winning strategy in $1\frac{1}{2}$-player games with $\mathcal{L}(F^{=1}, F^{>0}, G^{=1})$ objectives needs only a finite memory whose size is exponential in the size of a given objective. This result does *not* hold for $\mathcal{L}(F^{=1}, F^{>0}, G^{=1}, G^{>0})$ objectives—we show that even $\mathcal{L}(F^{>0}, G^{>0})$ objectives require infinite memory in general. In this sense, the previous result is tight.

Many interesting questions remain open. For example, it is not clear whether the existence of a winning strategy in $1\frac{1}{2}$-player games with qualitative PCTL objectives is decidable or not (all we know is that these strategies may require infinite memory). Another question is whether some of our positive results can be extended to $2\frac{1}{2}$-player games and/or to *concurrent* stochastic games with branching-time winning objectives. Our knowledge about randomized strategies is also limited, we have not addressed the issue of fairness, and so on. These problems are left for future research. Due to space constrains, some proofs are sketchy or completely omitted. Full proofs can be found in [2].

2. Basic Definitions

We start by recalling basic notions of probability theory. Let A be a finite set. A *probability distribution* on A is a function $f : A \to [0,1]$ such that $\sum_{a\in A} f(a) = 1$. A distribution f is *rational* if $f(a) \in \mathbb{Q}$ for every $a \in A$, and *Dirac* if $f(a) = 1$ for some $a \in A$. The set of all distributions on A is denoted $\mathcal{D}(A)$.

A *σ-field* over a set X is a set $\mathcal{F} \subseteq 2^X$ that includes X and is closed under complement and countable union. A *measurable space* is a pair $(X, \mathcal{F})$ where X is a set called *sample space* and $\mathcal{F}$ is a σ-field over X. A *probability measure* over measurable space $(X, \mathcal{F})$ is a function $\mathcal{P} : \mathcal{F} \to \mathbb{R}^{\geq 0}$ such that, for each countable collection $\{X_i\}_{i\in I}$ of pairwise disjoint elements of $\mathcal{F}$, $\mathcal{P}(\bigcup_{i\in I} X_i) = \sum_{i\in I} \mathcal{P}(X_i)$, and moreover $\mathcal{P}(X) = 1$. A *probabilistic space* is a triple $(X, \mathcal{F}, \mathcal{P})$ where $(X, \mathcal{F})$ is a measurable space and $\mathcal{P}$ is a probability measure over $(X, \mathcal{F})$.

Markov chains. A *Markov chain* is a triple $\mathcal{T} = (S, \to, Prob)$ where S is a finite or countably infinite set of *states*, $\to \subseteq S \times S$ is a *transition relation*, and $Prob$ is a function which to each transition $s \to t$ of $\mathcal{T}$ assigns its probability $Prob(s \to t) \in (0,1]$ so that for every $s \in S$ we have $\sum_{s\to t} Prob(s \to t) = 1$.

In the rest of this paper we also write $s \xrightarrow{x} t$ instead of $Prob(s \to t) = x$. A *path* in $\mathcal{T}$ is a finite or infinite sequence $w = s_0, s_1, \cdots$ of states such that $s_i \to s_{i+1}$ for every i. We also use $w(i)$ to denote the state s_i of w (by writing $w(i) = s$ we implicitly impose the condition that the length of w is at least $i+1$). The prefix $s_0, s_1, \ldots, s_i$ of w is denoted by w^i. A *run* is an infinite path. The sets of all finite paths and all runs of $\mathcal{T}$ are denoted $FPath$ and Run, respectively. Similarly, the sets of all finite paths and runs that start in a given $s \in S$ are denoted $FPath(s)$ and $Run(s)$, respectively.

Each $w \in FPath$ determines a *basic cylinder* $Run(w)$ which consists of all runs that start with w. To every $s \in S$ we associate the probabilistic space $(Run(s), \mathcal{F}, \mathcal{P})$ where $\mathcal{F}$ is the σ-field generated by all basic cylinders $Run(w)$ where w starts with s, and $\mathcal{P} : \mathcal{F} \to [0,1]$ is the unique probability function such that $\mathcal{P}(Run(w)) = \Pi_{i=0}^{m-1} x_i$ where $w = s_0, \cdots, s_m$ and $s_i \xrightarrow{x_i} s_{i+1}$ for every $0 \leq i < m$ (if $m = 0$, we put $\mathcal{P}(Run(w)) = 1$).

The logic PCTL. The logic PCTL, the probabilistic extension of CTL, was introduced by Hansson & Jonsson in [13]. Let $Ap = \{p, q, \ldots\}$ be a countably infinite set of *atomic propositions*. The syntax of PCTL formulae is given by the following abstract syntax equation:

$$\Phi \quad ::= \quad p \mid \neg p \mid \Phi_1 \vee \Phi_2 \mid \Phi_1 \wedge \Phi_2 \mid X^{\bowtie\varrho}\,\Phi \mid \Phi_1 U^{\bowtie\varrho}\,\Phi_2$$

Here $p \in Ap$, $\varrho \in [0,1]$, and $\bowtie \in \{\leq, <, \geq, >, =, \neq\}$.

Let $\mathcal{T} = (S, \to, Prob)$ be a Markov chain, and let $\nu : Ap \to 2^S$ be a *valuation*. The semantics of PCTL is defined below.

$s \models^\nu p$ iff $s \in \nu(p)$
$s \models^\nu \neg p$ iff $s \notin \nu(p)$

$$
\begin{array}{lll}
s \models^{\nu} \Phi_1 \vee \Phi_2 & \text{iff} & s \models^{\nu} \Phi_1 \text{ or } s \models^{\nu} \Phi_2 \\
s \models^{\nu} \Phi_1 \wedge \Phi_2 & \text{iff} & s \models^{\nu} \Phi_1 \text{ and } s \models^{\nu} \Phi_2 \\
s \models^{\nu} \mathrm{X}^{\bowtie \varrho}\Phi & \text{iff} & \mathcal{P}(\{w \in \mathit{Run}(s) \mid w(1) \models^{\nu} \Phi\}) \bowtie \varrho \\
s \models^{\nu} \Phi_1 \mathrm{U}^{\bowtie \varrho} \Phi_2 & \text{iff} & \mathcal{P}(\{w \in \mathit{Run}(s) \mid \exists j \geq 0 : w(j) \models^{\nu} \Phi_2 \\
 & & \quad \text{and } w(i) \models^{\nu} \Phi_1 \text{ for all } 0 \leq i < j\}) \bowtie \varrho
\end{array}
$$

Note that in our version of PCTL syntax, the negation can be applied only to atomic propositions. This is no restriction because the syntax is closed under dual connectives and relations: For every $\bowtie \in \{\leq, <, \geq, >, =, \neq\}$, let $\overline{\bowtie}$ be the complement of $\bowtie$ (for example, if $\bowtie$ is $\leq$, then $\overline{\bowtie}$ is $>$). The negation of $X^{\bowtie\varrho}\Phi$ and $\Phi_1 \mathrm{U}^{\bowtie\varrho} \Phi_2$ then corresponds to $X^{\overline{\bowtie}\varrho}\Phi$ and $\Phi_1 \mathrm{U}^{\overline{\bowtie}\varrho} \Phi_2$, respectively. The $\mathrm{F}^{\bowtie\varrho}$ and $\mathrm{G}^{\bowtie\varrho}$ operators are defined in the standard way: $\mathrm{F}^{\bowtie\varrho}\Phi$ stands for $\mathtt{tt}\,\mathrm{U}^{\bowtie\varrho}\,\Phi$, and $\mathrm{G}^{\bowtie\varrho}\Phi$ stands for $\mathtt{tt}\,\mathrm{U}^{\widehat{\bowtie} 1-\varrho}\,\neg\Phi$, where $\widehat{\bowtie}$ is $<, >, \leq, \geq, =$, or $\neq$, depending on whether $\bowtie$ is $>, <, \geq, \leq, =$, or $\neq$, respectively.

Various natural fragments of PCTL can be obtained by restricting the PCTL syntax to certain modal connectives and/or certain operator/number combinations. For example, the *qualitative* fragment of PCTL is obtained by restricting the allowed operator/number combinations to '$\bowtie$ 0' and '$\bowtie$ 1'. Hence, $a\mathrm{U}^{<1}b \vee \mathrm{F}^{>0}c$ is a qualitative PCTL formula. In this paper we also consider fragments with unary reachability and safety connectives. Formally, for each tuple $Y_1, \cdots, Y_n$, where each Y_i is of the form $\mathrm{X}^{\bowtie\varrho}$, $\mathrm{F}^{\bowtie\varrho}$, or $\mathrm{G}^{\bowtie\varrho}$, we define the $\mathcal{L}(Y_1, \cdots, Y_n)$ fragment of PCTL:

$$\Phi \quad ::= \quad p \mid \neg p \mid \Phi_1 \vee \Phi_2 \mid \Phi_1 \wedge \Phi_2 \mid Y_1\Phi \mid \cdots \mid Y_n\Phi$$

For example, $\mathrm{F}^{>0}(b \vee \mathrm{G}^{\geq 0.43}(\neg c \wedge \mathrm{F}^{<0.5}d))$ is a formula of $\mathcal{L}(\mathrm{F}^{>0}, \mathrm{G}^{\geq 0.43}, \mathrm{F}^{<0.5})$. Sometimes we also use formulae of the form $p \Rightarrow \Phi$ which stand for $\neg p \vee \Phi$.

Games, strategies, and objectives. A $2\frac{1}{2}$-*player game* is a tuple $G = (V, E, (V_\Box, V_\Diamond, V_\bigcirc), \mathit{Prob})$ where V is a finite set of *vertices*, $E \subseteq V \times V$ is the set of *transitions*, $(V_\Box, V_\Diamond, V_\bigcirc)$ is a partition of V, and Prob is a *probability assignment* which to each $v \in V_\bigcirc$ assigns a rational probability distribution on the set of its outgoing transitions. For technical convenience, we assume that each vertex has at least one outgoing transition. The game is played by two players, $\Box$ and $\Diamond$, who move a single token from vertex to vertex along the transitions of E. Player $\Box$ selects the moves in the $V_\Box$ vertices, and player $\Diamond$ selects the moves in the $V_\Diamond$ vertices. Transitions in the $V_\bigcirc$ vertices are chosen randomly according to the corresponding probability distribution. Game graphs are drawn in the standard way; vertices of $V_\Box$, $V_\Diamond$, and $V_\bigcirc$ are depicted as squares, diamonds, and circles, respectively. Probability distributions are usually uniform, which is indicated by arcs connecting the outgoing transitions of $V_\bigcirc$ vertices. A *strategy* for player $\Box$ is a function σ which to each $wv \in V^*V_\Box$ assigns a probability distribution on the set of outgoing transitions of v. We say that a strategy σ is *memoryless (M)* if $\sigma(wv)$ depends just on the last vertex v, and *deterministic (D)* if $\sigma(wv)$ is a Dirac distribution for each $wv \in V^*V_\Box$. Consistently with [1, 14], strategies that are not necessarily memoryless are called *history-dependent (H)*, and strategies that are not necessarily deterministic are called *randomized (R)*. A special type of history-dependent strategies are strategies with *finite memory*, which are formally defined as pairs $(\mathcal{A}, f)$ where $\mathcal{A} = (Q, V, \delta, q_0)$ is a deterministic finite-state automaton over the alphabet V of vertices and f is a function which to each pair $(q, v) \in Q \times V_\Box$ assigns a probability distribution on the set of outgoing transitions of v. The pair $(\mathcal{A}, f)$ determines a unique strategy $\sigma(\mathcal{A}, f)$ such that $\sigma(\mathcal{A}, f)(wv) = f(q, v)$, where $q = \delta(q_0, wv)$. Intuitively, the states of $\mathcal{A}$ represent a finite memory of size $|Q|$ where selected properties of the history of a play are stored. Hence, we can define the following four classes of strategies: MD, MR, HD, and HR, where MD $\subseteq$ HD $\subseteq$ HR and MD $\subseteq$ MR $\subseteq$ HR, but MR and HD are incomparable. Strategies for player $\Diamond$ are defined analogously. Each pair (σ, π) of strategies for player $\Box$ and $\Diamond$ determines a unique *play* of the game G, which is a Markov chain $G(\sigma, \pi)$ where V^+ is the set of states, and $wu \xrightarrow{x} wuu'$ iff $(u, u') \in E$ and one of the following conditions holds:

- $u \in V_\bigcirc$ and $\mathit{Prob}(u, u') = x$;
- $u \in V_\Box$ and $\sigma(wu)$ assigns x to (u, u');
- $u \in V_\Diamond$ and $\pi(wu)$ assigns x to (u, u').

An *objective* is a pair (ν, φ), where $\nu : Ap \to 2^V$ is a valuation and φ a PCTL formula. Note that each valuation $\nu : Ap \to 2^V$ determines a valuation $\overline{\nu} : Ap \to 2^{V^+}$ defined by $\overline{\nu}(a) = \{wu \in V^+ \mid u \in \nu(a)\}$. For a given objective (ν, φ), each state of $G(\sigma, \pi)$ either does or does not satisfy φ. A (ν, φ)-*winning strategy* for player $\Box$ in a vertex $v \in V$ is a strategy σ such that for every strategy π of player $\Diamond$ we have that $v \models^{\nu} \varphi$. Similarly, a (ν, φ)-winning strategy for player $\Diamond$ in a vertex $v \in V$ is a strategy π such that for every strategy σ of player $\Box$ we have that $v \not\models^{\nu} \varphi$. The game G is (ν, φ)-*determined* if there is a (ν, φ)-winning strategy for one of the two players in every vertex v of G.

$1\frac{1}{2}$-*player games* are $2\frac{1}{2}$-player games where the set $V_\Diamond$ is empty. Formally, a $1\frac{1}{2}$-player game is a tuple $G = (V, E, (V_\Box, V_\bigcirc), \mathit{Prob})$ where all elements have the expected meaning.

3. The Results

We start by observing that stochastic games with branching-time objectives are not determined, even if these objectives are taken from the $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0})$ fragment of PCTL. Consider the following game:

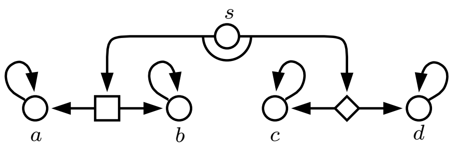

Let ν be a valuation which defines the validity of the propositions a, b, c, d as indicated in the above figure, and let $\varphi \equiv \mathrm{F}^{=1}(a \vee c) \vee \mathrm{F}^{=1}(b \vee d) \vee \big(\mathrm{F}^{>0}c \wedge \mathrm{F}^{>0}d\big)$. Now it is easy to check that none of the two players has a (ν, φ)-winning strategy in the vertex s, regardless whether we consider MD, MR, HD, or HR strategies.

3.1. Memoryless Strategies

In [1], it is shown that the problem whether there exists a winning MD strategy in a given $1\frac{1}{2}$-player game for a given PCTL objective is **NP**-complete. In fact, the **NP** lower bound holds even for the $\mathcal{L}(\mathrm{F}^{=1})$ fragment of PCTL. The following theorem gives a complexity classification for $2\frac{1}{2}$-player games.

Theorem 3.1. *The existence of a winning MD strategy for player $\Box$ in $2\frac{1}{2}$-player games with PCTL objectives is $\Sigma_2 = \boldsymbol{NP}^{\boldsymbol{NP}}$ complete. The Σ_2 lower bound holds even for $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0})$ objectives and for both MD and MR strategies.*

Proof. A Σ_2 *formula* is a formula of the form

$$\exists x_1, \cdots, x_n \, \forall y_1, \cdots, y_m \, B$$

where $n, m \in \mathbb{N}$ and B is a $\wedge, \vee$-expression over the (possibly negated) variables $x_1, \cdots, x_n, y_1, \cdots, y_m$. The problem whether a given Σ_2 formula is valid is Σ_2-complete [15].

Let $\psi \equiv \exists x_1, \cdots, x_n \, \forall y_1, \cdots, y_m \, B$. We construct a $2\frac{1}{2}$-player game $G(\psi)$, a valuation ν, and a formula $\varphi \in \mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0})$ such that player $\Box$ has a (ν, φ)-winning MD (or MR) strategy in a distinguished vertex v of $G(\psi)$ iff ψ is valid. Let us fix two sets $P = \{p_i, \widehat{p_i} \mid 1 \leq i \leq n\}$ and $Q = \{q_j, \widehat{q_j} \mid 1 \leq j \leq m\}$ of fresh atomic propositions, and let $P_i = P \smallsetminus \{p_i\}$, $\widehat{P}_i = P \smallsetminus \{\widehat{p_i}\}$, $Q_j = P \cup \{q_j\}$, $\widehat{Q}_j = P \cup \{\widehat{q_j}\}$ for all $1 \leq i \leq n$ and $1 \leq j \leq m$. The structure of $G(\psi)$ together with the valuation ν are shown in the following figure:

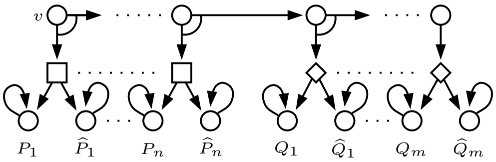

Let

$$\varphi \equiv \left(\bigvee_{j=1}^{m} (\mathrm{F}^{>0} q_j \wedge \mathrm{F}^{>0} \widehat{q_j}) \right) \vee \left(\widehat{B} \wedge \bigwedge_{i=1}^{n} (\mathrm{F}^{=1} p_i \vee \mathrm{F}^{=1} \widehat{p_i}) \right)$$

where $\widehat{B}$ is the formula obtained from B by substituting each occurrence of x_i, $\neg x_i$, y_j, and $\neg y_j$ with $\mathrm{F}^{=1}p_i$, $\mathrm{F}^{=1}\widehat{p_i}$, $\mathrm{F}^{>0}q_j$, and $\mathrm{F}^{>0}\widehat{q_j}$, respectively. Intuitively, player $\Box$ chooses an assignment for the variables $x_1, \cdots, x_n$ (x_i is set to true or false by selecting the transition to a vertex satisfying p_i or $\widehat{p_i}$, resp.). Note that player $\Box$ cannot use randomized moves because then the formula $\mathrm{F}^{=1}p_i \vee \mathrm{F}^{=1}\widehat{p_i}$ would not hold. Similarly, player $\Diamond$ chooses an assignment for $y_1, \cdots, y_m$. Observe that player $\Diamond$ cannot use randomized moves either because this would make some $\mathrm{F}^{>0}q_j \wedge \mathrm{F}^{>0}\widehat{q_j}$ true. Now it is easy to check that ψ is valid iff player $\Box$ has a (ν, φ)-winning MD (or MR) strategy in the vertex v. This establishes the Σ_2 lower bound.

The Σ_2 upper bound holds for all PCTL objectives. First, let us note that the model-checking problem for PCTL formulae and Markov chains is in **P** [13]. Hence, it suffices to "guess" a winning strategy σ for player $\Box$, and then ask the **NP** oracle whether there is a strategy π of player $\Diamond$ such that $G(\sigma, \pi)$ does *not* satisfy a given objective. The answer of the oracle is then simply negated. $\Box$

The complexity classification for MD strategies is thus established. As for MR strategies, the **NP** and Σ_2 lower bounds still hold. However, we managed to provide the matching upper bounds only for the subclass of *qualitative* PCTL objectives. Note that randomized strategies are more powerful than deterministic ones even for qualitative objectives—consider the formula $\mathrm{X}^{>0}p_u \wedge \mathrm{X}^{>0}p_v$ and a simple game G with three vertices $t, u, v \in V_\Box$ where $t \to u$, $t \to v$, $u \to u$, and $v \to v$. The propositions p_u and p_v hold only in u and v, respectively. Obviously, there is no winning (ν, φ)-winning MD strategy, but there are many (ν, φ)-winning MR strategies.

Theorem 3.2. *The existence of a winning MR strategy for player $\Box$ in $1\frac{1}{2}$-player (or $2\frac{1}{2}$-player) games with* qualitative *PCTL objectives is **NP**-complete (or $\Sigma_2 = \boldsymbol{NP}^{\boldsymbol{NP}}$ complete, resp.).*

Proof. A straightforward induction on the structure of a qualitative PCTL formula φ shows that the (in)validity of φ does not depend on the exact values of transition probabilities. It only matters which of the transition have zero/positive probability. Hence, in the case of $1\frac{1}{2}$-player games, it suffices to "guess" the subset of outgoing transitions in each vertex of $V_\Box$ which should have positive probability, and then verify that the guess was correct by a (polynomial time) PCTL model-checking algorithm [13]. The Σ_2 upper bound for $2\frac{1}{2}$-player games is established analogously (see the proof of Theorem 3.1). $\Box$

The existence of a winning MR strategy for player $\Box$ in $1\frac{1}{2}$-player and $2\frac{1}{2}$-player games with general PCTL objectives is known to be in **PSPACE** and **EXPTIME**, respectively [14]. We did not manage to lift the **NP** and Σ_2 lower bounds, and we also failed to improve the mentioned upper bounds. At least, we provide some evidence that lowering these bounds below **PSPACE** is difficult (see the discussion in Section 1). As a byproduct of this construction, we obtain an example of a $1\frac{1}{2}$-player game (where $V_{\bigcirc} = \emptyset$) and an objective (ν, φ) where $\varphi \in \mathcal{L}(\mathrm{X}^{>0}, \mathrm{U}^{=1/2})$ such that the only (ν, φ)-winning MR strategy assigns irrational probabilities to transitions.

Theorem 3.3. *The* SQUARE-ROOT-SUM *problem is efficiently reducible to the problem whether player $\Box$ has a winning MR strategy in $1\frac{1}{2}$-player games with PCTL objectives.*

Proof. Let $a_1, \cdots, a_n, b$ be an instance of SQUARE-ROOT-SUM (see Section 1). Let G be a game where

- the set $V = V_{\Box}$ contains the vertices v, u, s, $c(i)$, $d(i)$, $e(i)$, $f(i)$, $g(i)$, and $h(i)$ for all $1 \leq i \leq n$;
- the set of transitions contains $v \to u$, $u \to u$, $s \to u$, $v \to c(i)$, $c(i) \to d(i) \to f(i)$, $c(i) \to e(i) \to f(i)$, $f(i) \to g(i) \to s$, and $f(i) \to h(i) \to s$ for all $1 \leq i \leq n$.

The structure of G is shown in the following figure:

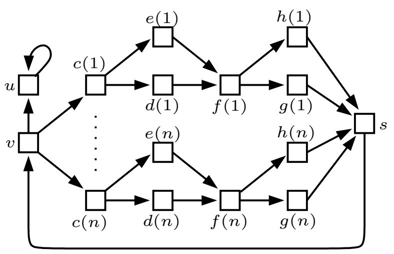

We assume that for each vertex $t \in V$ there is an atomic proposition p_t which is valid only in t (thus we obtain our valuation ν). Slightly abusing notation, we write t instead of p_t in our formulae.

Every strategy σ for player $\Box$ assigns (some) probabilities $p(c_i)$, $p(e_i)$, and $p(h_i)$ to transitions $v \to c(i)$, $c(i) \to e(i)$, and $f(i) \to h(i)$, respectively, where $1 \leq i \leq n$. Let $q = b + \sum_{i=1}^{n} a_i$. We construct a PCTL formula φ such that every (ν, φ)-winning MR strategy in v has to assign $p(c_i) = p(e_i) = p(h_i) = \sqrt{a_i}/q$ for every $1 \leq i \leq n$. Then the probability of $v \to u$ must be $1 - \sum_{i=1}^{n} \sqrt{a_i}/q$. The formula φ contains the clause $\mathrm{X}^{\geq 1-b/q} u$. Hence, player $\Box$ has a (ν, φ)-winning MR strategy in v iff $1 - \sum_{i=1}^{n} \sqrt{a_i}/q \geq 1 - b/q$, i.e., iff $\sum_{i=1}^{n} \sqrt{a_i} \leq b$.

Now we describe the formula φ in greater detail. For every $1 \leq i \leq n$, let $\Phi_i \equiv (v \vee c(i))\ \mathrm{U}^{a_i/q^2} e(i)$. Note that $v \models^{\nu} \Phi_i$ iff $p(c_i) \cdot p(e_i) = a_i/q^2$. Similarly, we construct the formulae Ψ_i and Ξ_i such that $v \models^{\nu} \Psi_i$ and $v \models^{\nu} \Xi_i$ iff $p(e_i) \cdot p(h_i) = a_i/q^2$ and $p(h_i) \cdot p(c_i) = a_i/q^2$, respectively:

$$\begin{aligned}\Psi_i &\equiv \mathrm{X}^{>0}\big(c(i) \vee (e(i) \vee f(i))\ \mathrm{U}^{a_i/q^2} h(i)\big)\\ \Xi_i &\equiv \mathrm{X}^{>0}\mathrm{X}^{>0}\mathrm{X}^{>0}\big((f(i) \vee h(i) \vee s \vee v)\ \mathrm{U}^{a_i/q^2} c(i)\big)\end{aligned}$$

Observe that if $p(c_i) \cdot p(e_i) = p(e_i) \cdot p(h_i) = p(h_i) \cdot p(c_i) = a_1/q^2$, then necessarily $p(c_i) = p(e_i) = p(h_i) = \sqrt{a_i}/q$. We put $\varphi \equiv \mathrm{X}^{\geq 1-b/q} u \wedge \bigwedge_{i=1}^{n} (\Phi_i \wedge \Psi_i \wedge \Xi_i)$.

Let us consider the game obtained for $n = 1$, $a_1 = 2$, and $b = 0$. Then $\Phi_1 \wedge \Psi_1 \wedge \Xi_1 \in \mathcal{L}(X^{>0}, \mathrm{U}^{1/2})$ and the only $(\nu, \Phi_1 \wedge \Psi_1 \wedge \Xi_1)$-winning MR strategy in v assigns irrational probabilities to certain transitions. Thus, we obtain the example promised above. $\Box$

3.2. History-Dependent Strategies

The results presented in this section constitute the main contribution of our paper. We start with the negative ones.

Theorem 3.4. *The existence of a winning HD (or HR) strategy in $1\frac{1}{2}$-player games with $\mathcal{L}(\mathrm{F}^{=5/8}, \mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ objectives is undecidable (and Σ_1^1-hard).*

Proof (sketch). The result is obtained by reduction of the problem whether a given nondeterministic Minsky machine with two counters initialized to zero has an infinite computation such that the initial instruction is executed infinitely often (this problem is known to be Σ_1^1-complete [**?**]). Formally, a nondeterministic Minsky machine with two counters c_1, c_2 is a finite sequence $\mathcal{M}$ of numbered instructions $1{:}ins_1, \cdots, n{:}ins_n$, where each ins_i is of one of the following forms (where $j \in \{1, 2\}$):

- $c_j := c_j{+}1;$ *goto* k
- *if* $c_j{=}0$ *then goto* k *else* $c_j := c_j{-}1;$ *goto* m
- *goto* $\{k$ *or* $m\}$

Here the indexes k, m range over $\{1, \cdots, n\}$. A *configuration* of $\mathcal{M}$ is a triple $[ins_i, v_1, v_2]$, where ins_i is the instruction to be executed, and $v_1, v_2 \in \mathbb{N}_0$ are the current values of c_1, c_2. A *computational step* $\hookrightarrow$ between configurations is defined in the expected way. A *recurrent computation* of $\mathcal{M}$ is an infinite computation initiated in $[ins_1, 0, 0]$ along which ins_1 is executed infinitely often. As we already mentioned, the problem whether a given $\mathcal{M}$ has a recurrent computation is Σ_1^1-complete.

Let $\mathcal{M} \equiv 1{:}ins_1, \cdots, n{:}ins_n$ be a nondeterministic Minsky machine. We construct a $1\frac{1}{2}$-player game $G(\mathcal{M})$ and a formula $\varphi \in \mathcal{L}(\mathrm{F}^{=5/8}, \mathrm{F}^{>0}, \mathrm{F}^{=1}, \mathrm{G}^{=1})$ such that player $\Box$ has a winning HD or HR strategy in a distinguished vertex v of $G(\mathcal{M})$ iff $\mathcal{M}$ has a recurrent computation.

Intuitively, the game $G(\mathcal{M})$ is constructed so that *every* play of $G(\mathcal{M})$ corresponds to an infinite sequence

$$[ins_1, 0, 0], \cdots, [ins_i, V_1, V_2], [ins_k, U_1, U_2], \cdots$$

of *extended* configurations of $\mathcal{M}$, where the counters can also take the ω (i.e., “infinite”) value. Player $\Box$ can (to some extent) determine the sequence. In particular, he is responsible for “guessing” the counter values in each extended configuration (see below). Of course, this sequence does not necessarily correspond to a valid computation of $\mathcal{M}$. The definition of $G(\mathcal{M})$ guarantees that the above sequence *does* correspond to a recurrent computation of $\mathcal{M}$ iff the following conditions are satisfied:

(a) Counter values in all extended configurations of the sequence are finite.

(b) The sequence contains infinitely many configurations of the form $[ins_1, \cdots]$.

(c) For each pair $[ins_i, V_1, V_2], [ins_k, U_1, U_2]$ of successive configurations we have that $[ins_i, V_1, V_2] \hookrightarrow [ins_k, U_1, U_2]$.

Then we show how to express these conditions in $\mathcal{L}(\mathrm{F}^{=5/8}, \mathrm{F}^{>0}, \mathrm{F}^{=1}, \mathrm{G}^{=1})$. Conditions (a) and (b) are relatively easy. Condition (c) requires more effort, and this is the (only) place where we need the $\mathrm{F}^{=5/8}$ operator. The main problem is to verify the “compatibility” of counter values V_1, U_1 and V_2, U_2. For example, if $ins_i \equiv c_1{:=}c_1{+}1;\, goto\ k$, then we must verify that $U_1 = V_1 + 1$ and $U_2 = V_2$. Here we illustrate just the basic idea of this construction (technical details can be found in [2]). We show how player $\Box$ can “guess” two numbers $I, J \in \mathbb{N}_0 \cup \{\omega\}$, and how to design a temporal formula which says that $I = J < \omega$. Let us consider the following game:

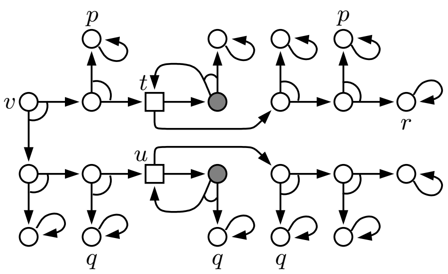

The p, q, and r are atomic propositions which are valid exactly in the indicated vertices. A play of this game (initiated in v) looks as follows:

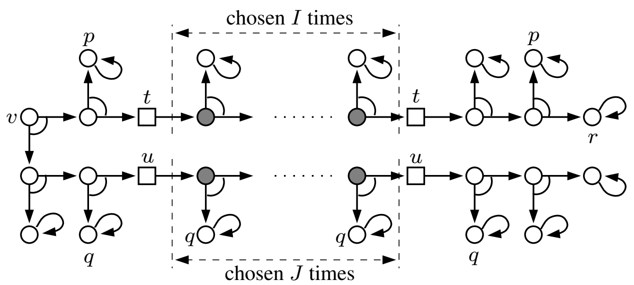

Observe that when a t vertex of the play is visited, player $\Box$ can choose between transitions leading to a “gray” or “white” vertex. If he chooses a gray vertex, then with probability $1/2$ he will make another choice after performing the next transition. Thus, player $\Box$ may decide to visit a gray vertex I-times, where I ranges from 0 to infinity, and the number of such choices represents the value of I. Similarly, the value of J is represented by the number of choices leading to a gray vertex at u vertices. The condition that $I < \omega$ is easy to express—we simply say that v satisfies $\mathrm{F}^{>0} r$. We claim that $I = J$ iff v satisfies $\mathrm{F}^{=1/2}(p \vee q)$. A closer look reveals that the probability of all $w \in Run(v)$ satisfying the formula $\mathrm{F}(p \vee q)$ is equal to the following sum of two binary numbers:

$$0.01\underbrace{0\cdots0}_{I}01 + 0.001\underbrace{1\cdots1}_{J}1$$

Obviously, this sum is equal to $1/2$ iff $I = J$, and we are done.

Note that in the above construction we used the $\mathrm{F}^{=1/2}$ operator, and not $\mathrm{F}^{=5/8}$. The exact value of the index does not really matter, any operator of the form $\mathrm{F}^{=x}$ where $0 < x < 1$ would suffice for our purposes. In the “full” proof, the operator $\mathrm{F}^{=5/8}$ is technically convenient, because then we can keep all transition probabilities in $G(\mathcal{M})$ equal to $1/2$. See [2] for the details. $\Box$

On the other hand, the existence of a winning HD strategy in $1\frac{1}{2}$ games with general PCTL objectives can be encoded by a Σ^1_1 formula in a straightforward way. Hence, the problem is Σ^1_1-complete.

A slight modification of the construction presented in Theorem 3.4 reveals the following:

Theorem 3.5. *The existence of a winning HD (or HR) strategy with finite memory in $1\frac{1}{2}$ games with $\mathcal{L}(\mathrm{F}^{=5/8}, \mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ objectives is undecidable.*

Proof (sketch). First, let us realize that the problem is semidecidable (i.e., belongs to the Σ_1 level of the arithmetical hierarchy). Obviously, one can effectively enumerate all $(\mathcal{A}, f)$ and for each such $(\mathcal{A}, f)$ decide whether $\sigma(\mathcal{A}, f)$ is winning, because the corresponding play has only finitely many states (more precisely, the play is obtained as unfolding of an effectively constructible finite-state Markov chain). The undecidability result is obtained by a slight modification of the construction presented in Theorem 3.4. In this case, we reduce the halting problem for “ordinary” deterministic Minsky machines (i.e., there is no *goto* $\{k\ or\ m\}$ instruction, and the last instruction is *halt*). Note that if a given Minsky machine halts, then it halts after finitely many steps and the corresponding winning strategy needs only finite memory (of course, there is no bound on its size).

If the machine does not halt, there is no winning strategy at all. $\square$

Now we show that the previous undecidability results are tight in the sense that the existence of a winning HD strategy in $1\frac{1}{2}$-player games with $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ objectives is decidable, and in fact **EXPTIME**-complete.

Let G be a $1\frac{1}{2}$-player game where V is the set of vertices. A *mixed* objective is a pair (P, Q) where $P, Q \subseteq V$. A strategy σ for player $\Box$ is (P, Q)-*winning* in a vertex $v \in V$ iff all runs in $G(\sigma)$ initiated in v visit some state of P infinitely often, and the probability of all runs which visit some state of Q infinitely often is 1. Hence, a mixed objective is essentially a conjunction of a *sure-Büchi* objective specified by P and a *qualitative-Büchi* objective specified by Q. The first step towards the promised **EXPTIME** upper bound is the following:

Lemma 3.6. *Let G be a $1\frac{1}{2}$-player game, s_{in} a vertex of G, and (ν, ψ) an objective where $\psi \in \mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$. Then there effectively exists a $1\frac{1}{2}$-player game G', a vertex s'_{in} of G', and a mixed objective (P, Q) such that player $\Box$ has a (ν, ψ)-winning HD strategy in the vertex s_{in} iff player $\Box$ has a (P, Q)-winning HD strategy in the vertex s'_{in}. Moreover, the G', s'_{in}, and (P, Q) are computable in time which is linear in the size of G and exponential in the size of ψ.*

Proof (sketch). For the rest of this proof, let us fix a $1\frac{1}{2}$-player game $G = (V, E, (V_\Box, V_\bigcirc), \mathit{Prob})$, a vertex $s_{in} \in V$, and an objective (ν, ψ) where $\psi \in \mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$. For technical convenience, we assume that all subformulae of ψ are pairwise distinct (this can be achieved by replacing atomic propositions in φ with fresh propositions so that each proposition has a unique occurrence in ψ; the valuation ν is extended accordingly). Our aim is to define another $1\frac{1}{2}$-player game G', a vertex s'_{in} of G', and a mixed objective (P, Q) such that player $\Box$ has a (ν, ψ)-winning HD strategy in s_{in} iff player $\Box$ has a (P, Q)-winning HD strategy in s'_{in}.

Let L be the set of all *literals*, i.e., atomic propositions and their negations. Let $\mathcal{S}$ be the set of all subformulae of ψ, where negation is not considered as a connective (for example, if $\psi \equiv F^{=1}\neg q$, then $\mathcal{S} = \{\neg q, F^{=1}\neg q\}$). For each connective $\alpha \in \{\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1}, \vee, \wedge\}$, we use $\mathcal{S}_\alpha$ to denote the subset of $\mathcal{S}$ consisting of all formulae where the topmost connective is α. We also use $\mathcal{S}_{Ap}, \mathcal{S}_{\mathrm{F}}, \mathcal{S}_{Temp}, \mathcal{S}_{Bool}$, and $\mathcal{S}_{\bar{\mathrm{F}}}$ to denote the sets $\mathcal{S} \cap L$, $\mathcal{S}_{\mathrm{F}^{=1}} \cup \mathcal{S}_{\mathrm{F}^{>0}}$, $\mathcal{S}_{\mathrm{F}^{=1}} \cup \mathcal{S}_{\mathrm{F}^{>0}} \cup \mathcal{S}_{\mathrm{G}^{=1}}$, $\mathcal{S}_\vee \cup \mathcal{S}_\wedge$, and $\{\bar{\mathrm{F}}^{\bowtie\varrho}\varphi \mid \mathrm{F}^{\bowtie\varrho}\varphi \in \mathcal{S}_{\mathrm{F}}\}$, respectively. The purpose of "barred" formulae of $\mathcal{S}_{\bar{\mathrm{F}}}$ becomes clear later.

In the following, we assume that $\mathcal{S} = \{\varphi_1, \ldots, \varphi_n\}$ where $i < j$ implies that φ_i is *not* a subformula of φ_j. The first step towards the definition of G' is the function $\Theta : 2^{\mathcal{S}} \to 2^{2^{\mathcal{S} \cup \mathcal{S}_{\bar{\mathrm{F}}}}}$ which decomposes subformulae of ψ into "subgoals". Let $A \subseteq \mathcal{S}$. If $A \subseteq \mathcal{S}_{Ap}$, then $\Theta(A) = \{A\}$. Otherwise, let i be the least index such that $\varphi_i \in A \smallsetminus \mathcal{S}_{Ap}$. We distinguish among the following possibilities:

- If $\varphi_i \equiv \varphi_k \vee \varphi_\ell$, then
$$\Theta(A) = \Theta((A \smallsetminus \{\varphi_i\}) \cup \{\varphi_k\}) \cup \Theta((A \smallsetminus \{\varphi_i\}) \cup \{\varphi_\ell\})$$
- If $\varphi_i \equiv \varphi_k \wedge \varphi_\ell$, then $\Theta(A) = \Theta((A \smallsetminus \{\varphi_i\}) \cup \{\varphi_k, \varphi_\ell\})$
- If $\varphi_i \equiv \mathrm{G}^{=1}\varphi_j$, then
$$\Theta(A) = \{D \cup \{\mathrm{G}^{=1}(\varphi_j)\} \mid D \in \Theta((A \smallsetminus \{\varphi_i\}) \cup \{\varphi_j\})\}$$
- $\varphi_i \equiv \mathrm{F}^{\bowtie\varrho}\varphi_j$, then
$$\begin{array}{rcl}\Theta(A) & = & \{D \cup \{\mathrm{F}^{\bowtie\varrho}(\varphi_j)\} \mid D \in \Theta(A \smallsetminus \{\varphi_i\})\} \\ & \cup & \{D \cup \{\bar{\mathrm{F}}^{\bowtie\varrho}(\varphi_j)\} \mid D \in \Theta((A \smallsetminus \{\varphi_i\}) \cup \{\varphi_j\})\}\end{array}$$

The intuition behind the function Θ is the following: to find out whether there is a (ν, ψ)-winning HD strategy in s_{in}, we extend each vertex of G (and hence each state of an arbitrary play of G) with a set A of subformulae of ψ that should be valid when the play is in the state. Some of these formulae represent temporal "goals" which can be achieved either in the current state or in its successors. The function Θ "offers" all admissible possibilities how to distribute the goals among the current state and its successors so that all formulae in A are valid. Selecting the right alternative becomes the responsibility of player $\Box$. For example, $\Theta(\{\mathrm{F}^{=1}p\}) = \{\{\mathrm{F}^{=1}p\}, \{\bar{\mathrm{F}}^{=1}p, p\}\}$, because the "current" state satisfies $\mathrm{F}^{=1}p$ iff either all of its successors satisfy $\mathrm{F}^{=1}p$ (the goal is "postponed"), or the proposition p is satisfied in the current state (the goal is "achieved"). In the latter case, the function Θ also "marks" the current state with $\bar{\mathrm{F}}^{=1}(p)$, which means that the goal $\mathrm{F}^{=1}(p)$ has been achieved. The exact purpose of these marks will be clarified later.

The game $G' = (V', E', (V'_\Box, V'_\bigcirc), \mathit{Prob}')$ is defined as follows. The set of vertices V' consists of vertices of the following two forms (f-vertices and g-vertices):

- f-vertices are of the form $(s, A, B, C)^f$ where $s \in V$, $A \subseteq \mathcal{S}$, $B \subseteq \mathcal{S}_{\mathrm{F}^{=1}} \cup \{\bullet\}$, and $C \subseteq \mathcal{S}_{\mathrm{F}^{>0}}$. Intuitively, the set A consists of formulas that should be satisfied in the current state (see the intuitive description of Θ above). The sets B and C assure that all subgoals of the form $F^{=1}\varphi$ and $F^{>0}\varphi$ are eventually fullfilled (see the mixed winning objective defined below).
- g-vertices are of the form $(s, A, B, C, \vec{D})^g$ where $s \in V$, $A \subseteq \mathcal{S} \cup \mathcal{S}_{\bar{\mathrm{F}}}$, $B \subseteq \mathcal{S}_{\mathrm{F}^{=1}} \cup \{\bullet\}$, $C \subseteq \mathcal{S}_{\mathrm{F}^{>0}}$, and $\vec{D} \in \prod_{t \in V} 2^{\mathcal{S}_{\mathrm{F}^{>0}}}$. The purpose of B and C is similar as in the case of f-vertices. The set A consists of subgoals that should be satisfied in successors of the current state. The vector $\vec{D}$ is used to distribute the subgoals of the form $\mathrm{F}^{>0}\varphi$ to the successors of the current state.

The set $V'_\Box$ consists of all f-vertices and of all g-vertices of the form $(s, A, B, C, \vec{D})^g$ where $s \in V_\Box$. The set $V'_\bigcirc$ consists of all g-vertices of the form $(s, A, B, C, \vec{D})^g$ where $s \in V_{\bigcirc}$. The set E' of transitions of G' is defined as follows:

1. $(s, A, B, C)^f \to (s, A', B', C', \vec{D})^g$ iff the following conditions are satisfied:
   - $A' \in \Theta(A)$
   - B' is equal to
     - $\{\bullet\}$, if $A' \cap Ap \not\subseteq \nu(s)$ or there is $\neg p \in A'$ such that $p \in \nu(s)$;
     - $A' \cap \mathcal{S}_{\mathrm{F}^{=1}}$, if $B = \emptyset$;
     - $B \smallsetminus \{\mathrm{F}^{=1}\xi \mid \bar{\mathrm{F}}^{=1}\xi \in A'\}$ otherwise.
   - if $C = \emptyset$ then $C' = A' \cap \mathcal{S}_{\mathrm{F}^{>0}}$; otherwise $C' = C \smallsetminus \{\mathrm{F}^{>0}\xi \mid \bar{\mathrm{F}}^{>0}\xi \in A'\}$.
   - $\bigcup_{(s,t)\in E} \vec{D}_t = A' \cap \mathcal{S}_{\mathrm{F}^{>0}}$
   - if $s \in V_{\Box}$ then for each $t \in V$ such that $(s, t) \in E$ we have that $D_t = A' \cap \mathcal{S}_{\mathrm{F}^{>0}}$.

   Intuitively, the f-vertices are controlled by player $\Box$ who chooses a set of subgoals $A' \in \Theta(A)$. The atomic propositions in A' are immediately verified (if there is some inconsistency then $\bullet$ is put into B') while the other formulae in A' are passed to successors. The sets B' and C' are updated depending on which subgoals (subformulae) are chosen by player $\Box$ as "achieved" in the current state (cf. the intuitive description of Θ above). Note that the vertex s is not changed in the successors of f-vertices. The transitions of G are simulated in g-vertices (see below).

2. $(s, A, B, C, \vec{D})^g \to (t, A', B', C')^f$ if $(s, t) \in E$, $A' = (A \smallsetminus (\mathcal{S}_{\mathrm{F}^{>0}} \cup \mathcal{S}_{\bar{\mathrm{F}}} \cup \mathcal{S}_{Ap})) \cup \vec{D}_t$, $B' = B$, and $C' = C \cap \vec{D}_t$

3. There are no other transitions in E' than those given by the rules 1. and 2.

$Prob'$ is defined as follows: For all $s \in V_{\bigcirc}$, the probability of $(s, A, B, C, \vec{D})^g \to (t, A', B', C')^f$ is the same as the probability of $s \to t$ in G. We put $s'_{in} = (s_{in}, \{\psi\}, \emptyset, \emptyset)^f$. Finally, we define the mixed (P, Q) objective as follows:

- the set P consists of all vertices of the form $(s, A, B, \emptyset, \vec{D})^g$;
- the set Q consists of all vertices of the form $(s, A, \emptyset, C, \vec{D})^g$.

It remains to show that player $\Box$ has a (ν, ψ)-winning HD strategy in s_{in} iff player $\Box$ has a (P, Q)-winning HD strategy in s'_{in}. A full proof of this assertion can be found in [2]. $\Box$

Hence, the problem of our interest is reducible to another game-theoretic problem, whose complexity is analyzed in our next lemma.

Lemma 3.7. *The existence of a winning strategy in* $1\frac{1}{2}$-*player games with mixed objectives is decidable in polynomial time.*

A direct consequence of Lemma 3.6 and Lemma 3.7 is that the existence of a winning HD strategy in $1\frac{1}{2}$-player games with $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ objectives is in **EXPTIME**. It remains to establish the matching lower bound.

Lemma 3.8. *The existence of a winning HD (or HR) strategy in* $1\frac{1}{2}$-*player games with* $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{G}^{=1})$ *objectives is* ***EXPTIME****-hard.*

A simple corollary of Lemma 3.6, Lemma 3.7, and Lemma 3.8 is the following:

Theorem 3.9. *The existence of a winning HD strategy in* $1\frac{1}{2}$-*player games with* $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ *objectives is* ***EXPTIME****-complete. The* ***EXPTIME*** *lower bound holds even for* $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{G}^{=1})$ *objectives.*

It follows from the proofs of Lemma 3.6 and Lemma 3.7 that a winning HD strategy in $1\frac{1}{2}$-player games with $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1})$ objectives actually requires only *finite memory* whose size is linear in the size of a given game and exponential in the size of a given objective. A natural question is whether Theorem 3.9 can be generalized to a larger class of qualitative PCTL objectives. One natural possibility is to add the $\mathrm{G}^{>0}$ operator, which yields the $\mathcal{L}(\mathrm{F}^{=1}, \mathrm{F}^{>0}, \mathrm{G}^{=1}, \mathrm{G}^{>0})$ fragment. However, there is a strong evidence that the method of Lemma 3.6 cannot be generalized to this class of objectives. This is because these objectives may already require *infinite memory*, which is demonstrated in our last theorem:

Theorem 3.10. *A winning HD strategy in* $1\frac{1}{2}$ *games with* $\mathcal{L}(\mathrm{F}^{>0}, \mathrm{G}^{>0})$ *objectives may require infinite memory.*

Proof. Let $\varphi \equiv \mathrm{G}^{>0}\big(\neg stop \wedge \mathrm{F}^{>0} stop\big)$ and let G be the following game (the valuation ν for atomic propositions *stop*, *left*, and *right* is also indicated in the figure):

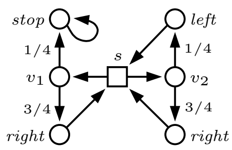

First we show that there is a (ν, φ)-winning HD strategy σ for player $\Box$ in the vertex v_1. We define $\sigma(ws)$ to be the Dirac distribution which assigns 1 to the transition leading to v_1 or v_2, depending on whether $\#_{right}(w) - \#_{left}(w) \leq 0$ or $\#_{right}(w) - \#_{left}(w) > 0$, respectively. Here $\#_{right}(w)$ denotes the number of occurrences of a state satisfying the proposition *right* in w. We claim that the state v_1 in the play $G(\sigma)$ satisfies the formula $\mathrm{G}^{=2/3}\big(\neg stop \wedge \mathrm{F}^{>0} stop\big)$ and hence also the formula φ. To see this, realize that the play $G(\sigma)$ corresponds to the unfolding of the following infinite Markov chain:

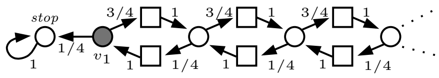

A standard calculation reveals that the probability of hitting the $stop$ state from v_1 is equal to $1/3$. Hence, the probability of all runs initiated in v_1 which do *not* hit the $stop$ state is $2/3$. All states in all these runs can reach the $stop$ state with positive probability. Hence, v_1 satisfies the formula $\mathrm{G}^{=2/3}(\neg stop \wedge \mathrm{F}^{>0} stop)$.

Now we show that there is no (ν, φ)-winning HD strategy with finite memory. Suppose the converse. Let $(\mathcal{A}, f)$ be such a strategy where the automaton $\mathcal{A} = (Q, V, \delta, q_0)$ has n states. We show that the state v_1 in the corresponding play satisfies the formula $\mathrm{G}^{=0}(\neg stop \wedge \mathrm{F}^{>0} stop)$, which means that v_1 does not satisfy φ. We say that a state w in the play $G(\sigma(\mathcal{A}, f))$ is *live* if there is a state $ww's$ such that $w \to^* ww's$ and $f(\delta(q_0, ww's), s)$ assigns 1 to the transition leading to v_1. A state which is not live is *dead*. We claim that there is a fixed $\varepsilon > 0$ such that the probability of hitting a $stop$ state from a given live state w is at least ε. To see this, it suffices to observe that whenever w is a live state, then there is a path from w to a $stop$ state of length at most $3n+1$. Note that a state w is dead iff w is a $stop$ state or w cannot reach a $stop$ state at all. By applying standard arguments of Markov chain theory, we can now conclude that the probability of hitting a dead state from v_1 is equal to one. Since a dead state does not satisfy $\neg stop \wedge \mathrm{F}^{>0} stop$, we obtain that v_1 satisfies $\mathrm{G}^{=0}(\neg stop \wedge \mathrm{F}^{>0} stop)$ and we are done. $\square$

References

[1] C. Baier, M. Größer, M. Leucker, B. Bollig, and F. Ciesinski. Controller synthesis for probabilistic systems. In *Proceedings of IFIP TCS'2004*. Kluwer, 2004.

[2] T. Brázdil, V. Brožek, V. Forejt, and A. Kučera. Stochastic games with branching-time winning objectives. Technical report FIMU-RS-2006-02, Faculty of Informatics, Masaryk University, 2006.

[3] T. Brázdil, J. Esparza, and A. Kučera. Analysis and prediction of the long-run behavior of probabilistic sequential programs with recursion. In *Proceedings of FOCS 2005*, pp. 521–530. IEEE, 2005.

[4] T. Brázdil, A. Kučera, and O. Stražovský. On the decidability of temporal properties of probabilistic pushdown automata. In *Proceedings of STACS'2005*, vol. 3404 of *LNCS*, pp. 145–157. Springer, 2005.

[5] K. Chatterjee, L. de Alfaro, and T. Henzinger. The complexity of stochastic Rabin and Streett games. In *Proceedings of ICALP 2005*, vol. 3580 of *LNCS*, pp. 878–890. Springer, 2005.

[6] K. Chatterjee and T. Henzinger. Semiperfect-information games. In *Proceedings of FST&TCS 2005*, vol. 3821 of *LNCS*, pp. 1–18. Springer, 2005.

[7] K. Chatterjee, M. Jurdzinski, and T. Henzinger. Simple stochastic parity games. In *Proceedings of CSL'93*, vol. 832 of *LNCS*, pp. 100–113. Springer, 1994.

[8] K. Chatterjee, M. Jurdzinski, and T. Henzinger. Quantitative stochastic parity games. In *Proceedings of SODA 2004*, pp. 121–130. SIAM, 2004.

[9] L. de Alfaro. Computing minimum and maximum reachability times in probabilistic systems. In *Proceedings of CONCUR'99*, vol. 1664 of *LNCS*, pp. 66–81. Springer, 1999.

[10] L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. *JCSS*, 68:374–397, 2004.

[11] J. Esparza, A. Kučera, and R. Mayr. Model-checking probabilistic pushdown automata. In *Proceedings of LICS 2004*, pp. 12–21. IEEE, 2004.

[12] K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of non-linear equations. In *Proceedings of STACS'2005*, vol. 3404 of *LNCS*, pp. 340–352. Springer, 2005.

[13] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6:512–535, 1994.

[14] A. Kučera and O. Stražovský. On the controller synthesis for finite-state Markov decision processes. In *Proceedings of FST&TCS 2005*, vol. 3821 of *LNCS*, pp. 541–552. Springer, 2005.

[15] Ch. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[16] W. Thomas. Infinite games and verification. In *Proceedings of CAV 2003*, vol. 2725 of *LNCS*, pp. 58–64. Springer, 2003.