Reachability in Recursive Markov Decision Processes ⋆

Tomáš Brázdil, Václav Brožek, Vojtěch Forejt, Antonín Kučera

Faculty of Informatics, Masaryk University,
Botanická 68a, 60200 Brno,
Czech Republic.

Abstract

We consider a class of infinite-state Markov decision processes generated by stateless pushdown automata. This class corresponds to $1\frac{1}{2}$-player games over graphs generated by BPA systems or (equivalently) 1-exit recursive state machines. An *extended reachability objective* is specified by two sets S and T of *safe* and *terminal* stack configurations, where the membership to S and T depends just on the top-of-the-stack symbol. The question is whether there is a suitable strategy such that the probability of hitting a terminal configuration by a path leading only through safe configurations is equal to (or different from) a given $x \in \{0, 1\}$. We show that the qualitative extended reachability problem is decidable in polynomial time, and that the set of all configurations for which there is a winning strategy is effectively regular. More precisely, this set can be represented by a deterministic finite-state automaton with a fixed number of control states. This result is a generalization of a recent theorem by Etessami & Yannakakis which says that the qualitative termination for 1-exit RMDPs (which exactly correspond to our $1\frac{1}{2}$-player BPA games) is decidable in polynomial time. Interestingly, the properties of winning strategies for the extended reachability objectives are quite different from the ones for termination, and new observations are needed to obtain the result. As an application, we derive the **EXPTIME**-completeness of the model-checking problem for $1\frac{1}{2}$-player BPA games and qualitative PCTL formulae.

Key words: Markov decision processes, temporal logics, stochastic games

⋆ Supported by the research center Institute for Theoretical Computer Science (ITI), project No. 1M0545.

Email addresses: `brazdil@fi.muni.cz` (Tomáš Brázdil), `xbrozek@fi.muni.cz` (Václav Brožek), `forejt@fi.muni.cz` (Vojtěch Forejt), `kucera@fi.muni.cz` (Antonín Kučera).

1 Introduction

$1\frac{1}{2}$-player games (or Markov decision processes) are a fundamental model in the area of system design and control optimization [13,9]. Formally, a $1\frac{1}{2}$-player game G is a directed graph where the vertices are split into two disjoint subsets $V_\square$ and $V_\bigcirc$. For every $v \in V_\bigcirc$, there is a fixed probability distribution over the set of its outgoing transitions. A *play* is initiated by putting a token on some vertex. The token is then moved from vertex to vertex by one "real" player $\square$ (controller) and one "virtual" player $\bigcirc$ (stochastic environment), who are responsible for selecting outgoing transitions in the vertices of $V_\square$ and $V_\bigcirc$, respectively. Player $\bigcirc$ does not make a real choice, but selects his next move randomly according to the fixed probability distribution over the outgoing transitions. A *strategy* specifies how player $\square$ 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, player $\square$ 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 strategy σ determines a unique Markov chain $G(\sigma)$ 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 $\square$, or the fixed probability of the transition (u, u') when $u \in V_\bigcirc$. A *winning objective* for player $\square$ is some property of Markov chains that is to be achieved. A *winning strategy* is a strategy that achieves the objective. In the context of "classical" MDP theory, winning objectives are typically related to long-time characteristics such as the expected total reward, the expected reward per transition, etc. [13,9]. In the context of formal verification, winning objectives are often specified as formulae of suitable temporal logics and their probabilistic variants such as PCTL or PCTL* [11]. For games with finitely many vertices, the corresponding decision algorithms have been designed [11,2,1] and also implemented in verifications tools such as PRISM (see, e.g., [12]). Recently, the scope of this study has been extended to a class of *infinite-state* games generated by *recursive state machines (RSM)* [7,8]. Intuitively, a RSM is a finite collection of finite-state automata which can call each other in a recursive fashion, maintaining the (unbounded) stack of activation records. RSM are semantically equivalent to *pushdown automata (PDA)*, and there are effective linear-time translations between the two models. A given RSM can be encoded in PDA syntax by storing the collection of finite-state automata in the control unit, and the recursive calls/returns are modeled by pushing/popping symbols onto/from the stack. An important subclass of RSM are *1-exit RSM*, where each finite-state automaton in the collection terminates in exactly one state. This means that no information can be returned back to the caller. In PDA terms, this means that whenever a given stack symbol X is popped from the stack, the same control state p_X is entered. Hence, the finite-state control unit can be encoded directly into the stack alphabet and simulated in top-of-the-stack symbol. Thus, 1-exit RSM can effectively be represented as *stateless* PDA, which are also denoted BPA in the context of concurrency theory.

Now we briefly summarize some of the results presented in [7,8]. To be able to give a clear comparison with our work, we reformulate these results in PDA/BPA terminology. A *termination objective* is specified by two control states p, q and one stack symbol X of a given PDA. The task of player $\Box$ is to maximize/minimize the probability of hitting $q\varepsilon$ from pX (each "head" rY in a given PDA is either probabilistic or non-deterministic; transitions from probabilistic heads are chosen randomly according to a fixed distribution, while the transitions from non-deterministic heads can be chosen by player $\Box$). In the case of BPA, there are no control states and the termination objective is specified simply by the stack symbol which is to be removed.

In [7,8], it has been shown that optimal minimizing/maximizing strategies in general $1\frac{1}{2}$-player PDA games with termination objectives do not always exist, and that the problem whether for every $0 < \delta \leq 1$ there is a strategy such that termination is achieved with probability at least $1 - \delta$ is undecidable. The situation is different for $1\frac{1}{2}$-player BPA games, where the optimal minimizing/maximizing strategies do exist, and can be constructed so that they depend only on top-of-the-stack symbol of a given configuration. Hence, the optimal strategies are stackless, memoryless, and deterministic (SMD). Furthermore, the corresponding minimal/maximal termination probabilities are expressible as the least solution of an effectively constructible system of non-linear min/max equations. Since the least solution of this system can effectively be expressed in first-order theory of the reals, this entails the decidability of the *quantitative* termination problem, i.e., the question whether the minimal/maximal achievable termination probability is bounded by a given constant. For the *qualitative* subcase (i.e., the problem whether the minimal/maximal achievable termination probability is equal to one), polynomial-time algorithms have been designed.

Our contribution: In this paper we consider $1\frac{1}{2}$-player BPA games with *extended reachability objectives (EROs).* An ERO is specified by two sets of *safe* and *terminal* stack symbols. A configuration is safe/terminal iff its top-of-the-stack symbol is safe/terminal. A run w satisfies a given ERO iff w visits a terminal configuration and all configurations preceding this visit are safe. The goal of player $\Box$ is to minimize/maximize the probability of all runs satisfying a given ERO. Note that termination objectives can easily be encoded as EROs (this may require a new bottom-of-the-stack symbol). However, the properties of EROs are surprisingly different from the ones of termination objectives (in contrast, methods for termination can easily be extended to EROs for fully probabilistic PDA [5]). We show that optimal maximizing strategies may not exist at all, and even if they do exist, they are not necessarily SMD. The optimal minimizing strategies are guaranteed to exist, but are not necessarily SMD. The method of expressing the minimal/maximal termination probability by a system of non-linear min/max equations used in [7] cannot be easily extended to EROs, and the reasons seem to be fundamental.

At the core of our paper are results about *qualitative* EROs. We show that the sets of all configurations for which there exists a strategy such that the probability of all runs satisfying a given ERO is equal to zero (equal to one, larger than zero, less than one, resp.) are regular and the corresponding finite-state automata can be constructed in polynomial time. In our algorithms, we use the results about qualitative termination as "black boxes" and concentrate on problems that are specific to EROs. We note that the subcase "equal to one", and particularly the subcase "less than one", require non-trivial methods and observations.

As an application, we design an exponential-time model-checking algorithm for $1\frac{1}{2}$-player BPA games and the qualitative fragment of the logic PCTL. More precisely, our algorithm is polynomial in the size of a given BPA and exponential in the size of a given formula (hence, the algorithm becomes polynomial for each fixed formula). Since there is a matching **EXPTIME** lower bound [4], we yield the **EXPTIME**-completeness of the problem. As a consequence we also obtain the **EXPTIME**-completeness of the model-checking problem for fully probabilistic BPA and qualitative PCTL (fully probabilistic BPA correspond to a subclass of $1\frac{1}{2}$-player BPA games where all heads are probabilistic). This problem was studied in [5,4], but the best known upper complexity bound was **EXPSPACE**.

2 Basic Definitions

In this paper, the set of all positive integers, non-negative integers, rational numbers, real numbers, and non-negative real numbers are denoted $\mathbb{N}$, $\mathbb{N}_0$, $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{R}^{\geq 0}$, respectively. For every finite or countably infinite set M, the symbol M^* denotes the set of all finite words over M. The length of a given word w is denoted $|w|$, and the individual letters in w are denoted $w(0), \cdots, w(|w|-1)$. The empty word is denoted by ε, where $|\varepsilon| = 0$. We also use M^+ to denote the set $M^* \smallsetminus \{\varepsilon\}$.

We start by recalling basic notions of probability theory. Let A be a finite or countably infinite set. A *probability distribution* on A is a function $f : A \to \mathbb{R}^{\geq 0}$ such that $\sum_{a \in A} f(a) = 1$. A distribution f is *rational* if $f(a) \in \mathbb{Q}$ for every $a \in A$, *positive* if $f(a) > 0$ 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 a 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 *probability 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 $(X, \mathcal{F})$.

2.1 Markov chains

A *Markov chain* is a triple $\mathcal{M} = (M, \to, \mathit{Prob})$ where M is a finite or countably infinite set of *states*, $\to \subseteq M \times M$ is a set of *transitions* such that for every $s \in M$ there is some transition $(s,t) \in \to$, and Prob is a function which to each $s \in M$ assigns a *positive* probability distribution over the set of its outgoing transitions.

In the rest of this paper we write $s \to t$ instead of $(s,t) \in \to$, and $s \xrightarrow{x} t$ instead of $\mathit{Prob}((s,t)) = x$. A *path* in $\mathcal{M}$ 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. The *length* of a finite path $w = s_0, \cdots, s_i$, denoted $|w|$, is $i+1$. We also use $w(i)$ to denote the state s_i of w, and w_i to denote the path $s_i, s_{i+1}, \cdots$ (by writing $w(i) = s$ or w_i we implicitly impose the condition that $|w| \geq i+1$). A state t is *reachable* from a state s, written $s \to^* t$, if there is a finite path from s to t.

A *run* is an infinite path. The sets of all finite paths and all runs of $\mathcal{M}$ are denoted $\mathit{FPath}(\mathcal{M})$ and $\mathit{Run}(\mathcal{M})$, respectively. Similarly, the sets of all finite paths and runs that start in a given $s \in M$ are denoted $\mathit{FPath}(\mathcal{M}, s)$ and $\mathit{Run}(\mathcal{M}, s)$, respectively.

Each $w \in \mathit{FPath}(\mathcal{M})$ determines a *basic cylinder* $\mathit{Run}(\mathcal{M}, w)$ which consists of all runs that start with w. To every $s \in M$ we associate the probability space $(\mathit{Run}(\mathcal{M}, s), \mathcal{F}, \mathcal{P})$ where $\mathcal{F}$ is the σ-field generated by all basic cylinders $\mathit{Run}(\mathcal{M}, w)$ where w starts with s, and $\mathcal{P} : \mathcal{F} \to \mathbb{R}^{\geq 0}$ is the unique probability measure such that $\mathcal{P}(\mathit{Run}(\mathcal{M}, w)) = \Pi_{i=0}^{m-1} x_i$ 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 $\mathcal{P}(\mathit{Run}(\mathcal{M}, w)) = 1$).

2.2 Games and strategies

A *$1\frac{1}{2}$-player game* (or *Markov decision process*) is a tuple $G = (V, \mapsto, (V_\Box, V_\bigcirc), \mathit{Prob})$ where V is a finite or countably infinite set of *vertices*,

$\mapsto \subseteq V \times V$ is a set of *transitions*, $(V_\Box, V_\bigcirc)$ is a partition of V, and *Prob* is a *probability assignment* which to each $v \in V_\bigcirc$ assigns a positive probability distribution on the set of its outgoing transitions. We write $v \mapsto u$ instead of $(v,u) \in \mapsto$, and $v \overset{x}{\mapsto} u$ when $v \mapsto u$, $v \in V_\bigcirc$, and $Prob((v,u)) = x$. For technical convenience, we assume that each vertex has at least one outgoing transition. We say that G is *finitely-branching* if for each $v \in V$ there are only finitely many $u \in V$ such that $v \mapsto u$.

The game is played by a player $\Box$ who selects the moves in the $V_\Box$ vertices, and a "virtual" player $\bigcirc$ who selects the moves in the $V_\bigcirc$ vertices according to the corresponding probability distribution.

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 $wv \in V^*V_\Box$. Strategies that are not necessarily memoryless are called *history-dependent (H)*, and strategies that are not necessarily deterministic are called *randomized (R)*. 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.

Each strategy σ for player $\Box$ determines a unique *play* of the game G, which is a Markov chain $G(\sigma)$ where V^+ is the set of states, and $wu \xrightarrow{x} wuu'$ iff $u \mapsto u'$ and one of the following conditions holds:

- $u \in V_\bigcirc$ and $u \overset{x}{\mapsto} u'$;
- $u \in V_\Box$ and $\sigma(wu)$ assigns x to $u \mapsto u'$, where $x > 0$.

For every $w \in Run(G(\sigma))$ and every $i \in \mathbb{N}_0$, we define $w[i]$ to be the last vertex of $w(i)$ (realize that $w(i)$ is a finite sequence of vertices of the game G). Further, for all $S, T \subseteq V$ and $u \in V$, we define the sets

- $Run(G(\sigma), u, S\,\mathcal{U}\,T) = \{w \in Run(G(\sigma), u) \mid \exists j \geq 0 : w[j] \in T \wedge \forall i < j : w[i] \in S\}$
- $Run(G(\sigma), u, \mathcal{X}\,S) = \{w \in Run(G(\sigma), u) \mid w[1] \in S\}$
- $Run(G(\sigma), u, \mathcal{F}T) = \{w \in Run(G(\sigma), u) \mid \exists j \geq 0 : w[j] \in T\}$
- $Run(G(\sigma), u, \mathcal{G}T) = \{w \in Run(G(\sigma), u) \mid \forall j \geq 0 : w[j] \in T\}$
- $Run(G(\sigma), u, \neg\mathcal{F}T) = Run(G(\sigma), u) \smallsetminus Run(G(\sigma), u, \mathcal{F}T)$

Now we introduce some notation for MD strategies which is used in proofs of Section 3 and Section 4.

To each MD strategy σ we associate a function $f_\sigma : V_\Box \to V$ where $f_\sigma(v)$ returns the (unique) vertex v' such that $\sigma(v)$ assigns 1 to $v \mapsto v'$. Conversely, each $f : V_\Box \to V$ such that $v \mapsto f(v)$ for all $v \in V_\Box$ determines a unique MD strategy σ_f where $\sigma_f(wv)$ assigns 1 to the transition $v \mapsto f(v)$.

Each MD strategy σ for player $\Box$ determines a Markov chain $G[\sigma]$ where V is the set of states, and $v \xrightarrow{x} u$ in $G[\sigma]$ iff $v \mapsto u$ in G and one of the following conditions holds:

- $v \in V_\bigcirc$ and $v \overset{x}{\mapsto} u$;
- $v \in V_\Box$, $x = 1$, and $f_\sigma(v) = u$.

The chains $G[\sigma]$ and $G(\sigma)$ are "equivalent" in the sense that the unfoldings of $G(\sigma)$ and $G[\sigma]$ are isomorphic for every initial vertex $u \in V$. For our purposes, it suffices to know that if we define the sets $Run(G[\sigma], u, S\,\mathcal{U}\,T)$, $Run(G[\sigma], u, \mathcal{X}\,S)$, $Run(G[\sigma], u, \mathcal{F}T)$, $Run(G[\sigma], u, \mathcal{G}T)$, and $Run(G[\sigma], u, \neg\mathcal{F}T)$ analogously as above, then these sets have the same probability as their respective counterparts in $G(\sigma)$. For example, $\mathcal{P}(Run(G[\sigma], u, S\,\mathcal{U}\,T)) = \mathcal{P}(Run(G(\sigma), u, S\,\mathcal{U}\,T))$. Therefore, if we restrict our attention to MD strategies, we can safely consider the chain $G[\sigma]$ instead of the chain $G(\sigma)$. This becomes particularly convenient in Section 4.

2.3 The logic PCTL

The logic PCTL, the probabilistic extension of CTL, was introduced by Hansson & Jonsson in [11]. Originally, the semantics of PCTL was defined over Markov chains. Here we consider a more general semantics defined over $1\frac{1}{2}$-player games, as proposed by de Alfaro & Bianco in [2].

Let $Ap = \{p, q, \dots\}$ be a countably infinite set of *atomic propositions*. The syntax of PCTL formulae is given by the following abstract syntax equation:

$$\Phi ::= \mathtt{tt} \;\mid\; p \;\mid\; \Phi_1 \wedge \Phi_2 \;\mid\; \neg\Phi \;\mid\; \mathcal{X}^{\bowtie\varrho}\,\Phi \;\mid\; \Phi_1\,\mathcal{U}^{\bowtie\varrho}\,\Phi_2$$

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

Let $G = (V, \mapsto, (V_\Box, V_\bigcirc), Prob)$ be a $1\frac{1}{2}$-player game, and let $\nu : Ap \to 2^V$ be a *valuation*. The semantics of PCTL is defined below.

$$\begin{aligned}
[\![\mathtt{tt}]\!]^\nu &= V\\
[\![p]\!]^\nu &= \nu(p)\\
[\![\Phi_1 \wedge \Phi_2]\!]^\nu &= [\![\Phi_1]\!]^\nu \cap [\![\Phi_2]\!]^\nu\\
[\![\neg\Phi]\!]^\nu &= V \smallsetminus [\![\Phi]\!]^\nu\\
[\![\mathcal{X}^{\bowtie\varrho}\Phi]\!]^\nu &= \{u \in V \mid \forall\sigma \in \mathrm{HR} : \mathcal{P}(Run(G(\sigma), u, \mathcal{X}\,[\![\Phi]\!]^\nu)) \bowtie \varrho\}\\
[\![\Phi_1\,\mathcal{U}^{\bowtie\varrho}\Phi_2]\!]^\nu &= \{u \in V \mid \forall\sigma \in \mathrm{HR} : \mathcal{P}(Run(G(\sigma), u, [\![\Phi_1]\!]^\nu\,\mathcal{U}\,[\![\Phi_2]\!]^\nu)) \bowtie \varrho\}
\end{aligned}$$

The $\mathcal{F}^{\bowtie\varrho}$ and $\mathcal{G}^{\bowtie\varrho}$ operators are defined in the standard way: $\mathcal{F}^{\bowtie\varrho}\Phi$ stands for $\mathtt{tt}\,\mathcal{U}^{\bowtie\varrho}\,\Phi$, and $\mathcal{G}^{\bowtie\varrho}\Phi$ stands for $\mathtt{tt}\,\mathcal{U}^{\widehat{\bowtie}1-\varrho}\,\neg\Phi$, 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. The *qualitative* fragment of PCTL is obtained by restricting the allowed operator/number combinations to '$\bowtie 0$' and '$\bowtie 1$'. Hence, $a\mathcal{U}^{<1}b \vee \mathcal{F}^{>0}c$ is a qualitative PCTL formula.

3 Extended Reachability Objectives

In this section we present several general results about $1\frac{1}{2}$-player games with extended reachability objectives.

Definition 1 *Let $G = (V, \mapsto, (V_\Box, V_\bigcirc), Prob)$ be an (arbitrary) $1\frac{1}{2}$-player game. An* extended reachability objective (ERO) *is a pair (S, T), where $S, T \subseteq V$ are the subsets of* safe *and* terminal *vertices.*

Let (S, T) be an ERO. For every HR strategy σ and every $u \in V$ we define the σ-value of u*, denoted $Val^\sigma(u)$, as follows:*

$$Val^\sigma(u) = \mathcal{P}(Run(G(\sigma), u, S\,\mathcal{U}\,T))$$

We also define the upper *and* lower *value of u, denoted $Val^+(u)$ and $Val^-(u)$, as the* sup *and* inf *of the set $\{Val^\sigma(u) \mid \sigma \in HR\}$, respectively. An* optimal maximizing *and* optimal minimizing *strategy for a vertex u is a strategy σ such that $Val^\sigma(u)$ is equal to $Val^+(u)$ and to $Val^-(u)$, respectively.*

It has been shown in [2] that optimal maximizing/minimizing strategies always exist in $1\frac{1}{2}$-player games with *finitely* many vertices. Moreover, there are efficiently constructible optimal maximizing/minimizing MD strategies. These results are no longer valid for $1\frac{1}{2}$-player games with *infinitely* many vertices. A brief summary of relevant results is given in Proposition 4. The proof is based on standard arguments of Markov decision process theory [13,9], and it is included mainly for the sake of completeness (the only exception is the last claim (5) for which we did not manage to construct a sufficiently simple proof; here we give just a pointer to literature). We start with a preliminary observation.

Definition 2 *Let $G = (V, \mapsto, (V_\Box, V_\bigcirc), Prob)$ be a $1\frac{1}{2}$-player game and (S, T) an ERO. A transition $s \mapsto t$ is* min-optimal *if for every $s \mapsto t'$ we have that $Val^-(t) \leq Val^-(t')$. A given $s \in V_\Box$ is* min-optimizing *if it has at least one min-optimal outgoing transition.*

Lemma 3 *Let $G = (V, \mapsto, (V_\Box, V_\bigcirc), Prob)$ be a $1\frac{1}{2}$-player game and (S,T) an ERO. Let σ_f be (some) MD strategy such that the underlying function $f : V_\Box \to V$ returns a min-optimal transition for every min-optimizing vertex. A vertex $v \in V$ is* covered by σ_f *if one of the following conditions is satisfied:*

- $v \in (V \smallsetminus S) \cup T$;
- $v \in S \smallsetminus T$ *and for each $wt \in (S \smallsetminus T)^*$ that is reachable from v in $G(\sigma_f)$ we have that t is either min-optimizing or belongs to $V_\bigcirc$.*

Then σ_f is an optimal minimizing strategy for every vertex of V which is covered by σ_f.

PROOF. Assume the converse. Then there is $v \in S \smallsetminus T$ covered by σ_f such that $Val^-(v) < Val^{\sigma_f}(v)$. Let $\varepsilon = Val^{\sigma_f}(v) - Val^-(v)$. We show that

(1) For all $k \in \mathbb{N}$, $\delta > 0$, and $u \in V$ such that $v \to^* wu$ in $G(\sigma_f)$ for some $wu \in (S \smallsetminus T)^*$ there is a HR strategy $\pi[k, \delta, u]$ such that
  · $\pi[k, \delta, u](uw) = \sigma_f(uw)$ for all $w \in V^*$ such that $uw \in V^* V_\Box$ and $|uw| < k$;
  · $Val^{\pi[k,\delta,u]}(u) - Val^-(u) < \delta$.

(2) For every $i \in \mathbb{N}$, let Λ_i be the set of all $w \in FPath(G(\sigma_f), v)$ such that $|w| = k \leq i$ and $w(k-1) \in (S \smallsetminus T)^* T$. Then for every $\xi > 0$ there is $m \in \mathbb{N}$ such that $\sum_{w \in \Lambda_m} \mathcal{P}(Run(G(\sigma_f), w)) > Val^{\sigma_f}(v) - \xi$.

Note that (1) and (2) lead to a contradiction—put $\xi = \varepsilon/3$ and consider the m of (2). Since the strategy $\pi[m, \varepsilon/3, v]$ of (1) "agrees" with σ_f on the first m steps of a play, we obtain that $Val^{\pi[m,\varepsilon/3,v]}(v) > Val^{\sigma_f}(v) - \varepsilon/3$ (this is the inequality of (2)). Hence, the inequality $Val^{\pi[m,\varepsilon/3,v]}(v) - Val^-(v) < \varepsilon/3$ of (1) does not hold, which is a contradiction.

(1): By induction on k:

- $\mathbf{k = 1}$. Let $u \in V$ such that $v \to^* wu$ in $G(\sigma_f)$ for some $wu \in (S \smallsetminus T)^*$, and let $\delta > 0$. Let π be a HR strategy such that $Val^\pi(u) - Val^-(u) < \delta$. It suffices to put $\pi[1, \delta, u] = \pi$.
- **Induction step:** Let $u \in V$ such that $v \to^* wu$ in $G(\sigma_f)$ for some $wu \in (S \smallsetminus T)^*$, and let $\delta > 0$. We distinguish two cases:
  · $u \in V_\Box$. Then u is min-optimizing. Let $u' = f(u)$. If $u' \in (V \smallsetminus S) \cup T$, we can simply put $\pi[k{+}1, \delta, u] = \sigma_f$ because $Val^{\sigma_f}(u) = Val^-(u)$ in this case. Otherwise, we use induction hypothesis and conclude that there is a strategy $\pi[k, \delta, u']$ with the respective properties. Now we can define $\pi[k{+}1, \delta, u]$ as follows:
    $\pi[k{+}1, \delta, u](u) = \sigma_f(u)$
    $\pi[k{+}1, \delta, u](uu'w) = \pi[k, \delta, u'](u'w)$ for every $w \in V^*$ such that $u'w \in V^* V_\Box$.

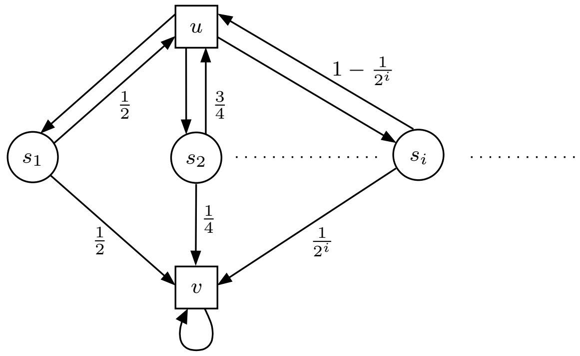

Fig. 1. A $1\frac{1}{2}$-player game without an optimal minimizing strategy.

· $u \in V_{\bigcirc}$. Let $succ(u) = \{t \in V \mid u \mapsto t\}$. For every $t \in succ(u)$, we fix some $\delta_t > 0$ so that $\sum_{t \in succ(u)} \delta_t < \delta$. For every $t \in succ(u) \cap (S \smallsetminus T)$ there is a strategy $\pi[k, \delta_t, t]$ with the respective properties. We define $\pi[k{+}1, \delta, u]$ as follows:

For every $t \in succ(u) \cap (S \smallsetminus T)$ we put $\pi[k{+}1, \delta, u](utw) = \pi[k, \delta_t, t](tw)$ for all $w \in V^*$ such that $tw \in V^* V_{\Box}$;
For every $t \in succ(u) \cap ((V \smallsetminus S) \cup T)$ and all $w \in V^*$ such that $tw \in V^* V_{\Box}$ we define $\pi[k{+}1, \delta, u](utw)$ arbitrarily. One can easily confirm that $Val^{\pi[k+1,\delta,u]}(u) - Val^{-}(u) < \delta$ as needed.

(2): This claim follows directly from the definition of $\mathcal{P}(Run(G(\sigma_f)), v, S\,\mathcal{U}\,T))$. $\Box$

Proposition 4 *Let $G = (V, \mapsto, (V_{\Box}, V_{\bigcirc}), Prob)$ be a $1\frac{1}{2}$-player game, $u \in V$, and (S, T) an ERO (let us note explicitly that V can be infinite and some vertices can have infinitely many successors). The following holds:*

(1) An optimal minimizing strategy for u does not necessarily exist, and the equation $\inf\{Val^{\sigma}(u) \mid \sigma \in MD\} = Val^{-}(u)$ does not necessarily hold.
(2) If there is an optimal minimizing strategy for u, then there is also an optimal minimizing MD strategy for u.
(3) If G is finitely-branching, then there is an optimal minimizing MD strategy for u.
(4) An optimal maximizing strategy for u does not necessarily exist (even if G is finitely-branching), but $\sup\{Val^{\sigma}(u) \mid \sigma \in MD\} = Val^{+}(u)$.
(5) If there is an optimal maximizing strategy for u, then there is also an optimal maximizing MD strategy for u.

PROOF.

(1) Consider the game $G = (V, \mapsto, (V_{\Box}, V_{\bigcirc}), Prob)$ where $V_{\Box} = \{u, v\}$, $V_{\bigcirc} = \{s_i \mid i \in \mathbb{N}\}$, $v \mapsto v$, $u \mapsto s_i$, $s_i \mapsto u$, and $s_i \mapsto v$ for all $i \in \mathbb{N}$, and

$Prob((s_i, v)) = 2^{-i}$, $Prob((s_i, u)) = 1 - 2^{-i}$ for all $i \in \mathbb{N}$. Further, consider an ERO $(V, \{v\})$. The structure of G is shown in Fig. 1. We show that there is no optimal minimizing strategy for u and that

$$1 \quad = \quad \inf\{\, Val^{\sigma}(u) \mid \sigma \in \mathrm{MD}\} \quad > \quad Val^{-}(u) \quad = \quad 0.$$

It is clear that there is no strategy σ such that $Val^{\sigma}(u) = 0$, because from every successor of u there is a transition to v with a positive probability. We show that for every $\delta > 0$ there is a HR strategy $\sigma(\delta)$ such that $Val^{\sigma(\delta)}(u) < \delta$. From this we obtain $\inf\{\, Val^{\sigma}(u) \mid \sigma \in \mathrm{HR}\} = Val^{-}(u) = 0$. For a given $\delta > 0$, we choose a sufficiently large $i \in \mathbb{N}$ such that $\frac{2}{2^i} < \delta$. The strategy $\sigma(\delta)$ is given by $\sigma(\delta)(wu) = s_j$ where $j = \sharp_u(w) + i$ (the symbol $\sharp_u(w)$ denotes the number of occurrences of u in w). It is easy to see that

$$Val^{\sigma(\delta)}(u) \;=\; \mathcal{P}(Run(G(\sigma(\delta)), u, V\,\mathcal{U}\,\{v\})) \;=\; \sum_{j=i}^{\infty} \frac{1}{2^j} \;=\; \frac{2}{2^i} \;<\; \delta.$$

Now, let σ be an MD strategy and let $f_\sigma(u) = s_i$. The only run from u which does not visit v is $u, s_i, u, s_i, \cdots$, hence $\mathcal{P}(Run(G(\sigma), u, V\,\mathcal{U}\,\{v\})) = 1$. This holds for all MD strategies and thus we obtain $\inf\{\, Val^{\sigma}(u) \mid \sigma \in \mathrm{MD}\} = 1$.

(2) Let us fix some optimal minimizing strategy π for u and observe the following:

(I) For every vertex $wv \in (S \setminus T)^*$ of $G(\pi)$ such that $u \to^* wv$ in $G(\pi)$ we have that $\mathcal{P}(Run(G(\pi), wv, S\,\mathcal{U}\,T)) = Val^{-}(v)$. If it was not the case, there would be some strategy π' such that $Val^{\pi'}(v) < \mathcal{P}(Run(G(\pi), wv, S\,\mathcal{U}\,T)$, and hence we could define another strategy π'' behaving identically as π except for vertices of the form wvw' where $\pi''(wvw') = \pi'(vw')$. Since $Val^{\pi''}(u) < Val^{\pi}(u)$, this contradicts the optimality of π.

(II) For every transition $wv \xrightarrow{x} wvv'$ of $G(\pi)$ such that $wv \in (S \setminus T)^*$, $u \to^* wv$ in $G(\pi)$, and $v \in V_\Box$ we have that $v \mapsto v'$ is min-optimal. To see this, realize that

$$\mathcal{P}(Run(G(\pi), wv, S\,\mathcal{U}\,T)) = \sum_{wv \xrightarrow{x} wvv'} x \cdot \mathcal{P}(Run(G(\pi), wvv', S\,\mathcal{U}\,T))$$

By applying (I) and considering the case when $v' \in (V \setminus S) \cup T$ we get

$$Val^{-}(v) = \sum_{wv \xrightarrow{x} wvv'} x \cdot Val^{-}(v')$$

Realize that there cannot be any $v \mapsto v''$ such that $Val^{-}(v'') < Val^{-}(v)$ (it follows directly from Definition 1). This means that all $Val^{-}(v')$ which appear in the above sum are equal to $Val^{-}(v)$. Hence, all of the corresponding $v \mapsto v'$ transitions are min-optimal.

Let $f : V_\Box \to V$ be a function satisfying the following condition: for every $v \in V_\Box \cap S$ such that $u \to^* wv$ for some $w \in (S \setminus T)^*$ we have that $f(v) = v'$, where $wv \xrightarrow{x} wvv'$ is a transition of $G(\pi)$. It follows from (II) that the vertex u is covered by σ_f, and hence we can apply Lemma 3.

(3) If G is finitely branching, all vertices of $V_\Box$ are min-optimizing. Thus, the claim follows from Lemma 3.

(4) In Section 4 we give an example of a $1\frac{1}{2}$-player BPA game and an ERO for which no optimal maximizing strategy exists (see Example 6). Since BPA games are finitely-branching, this example confirms the first part of the claim.

Now we show that for every HR strategy π and every $\delta > 0$ there is a MD strategy σ such that $\mathit{Val}^\pi(u) - \delta \leq \mathit{Val}^\sigma(u)$. Let Λ be the set of all $w \in \mathit{FPath}(G(\pi), u)$ such that $w(k{-}1) \in S^*T$ where $k = |w|$. It follows directly from the definition of $\mathit{Run}(G(\pi), u, S\,\mathcal{U}\,T)$ that there must be a *finite* subset $\Lambda' \subseteq \Lambda$ such that $\sum_{w\in\Lambda'} \mathcal{P}(\mathit{Run}(G(\pi), w)) \geq \mathit{Val}^\pi(u) - \delta$. Let $W \subseteq V$ be the subset of all vertices that are visited by a path in Λ'. We define a *finite-state* $1\frac{1}{2}$-player game F where $W \cup \{*\}$ is the set of vertices (here $* \notin W$ is a fresh symbol), probabilistic vertices are exactly $W \cap V_\bigcirc$, and transitions $\hookrightarrow$ together with their probabilities are determined as follows:

- For all $u, v \in W$ we have that $u \hookrightarrow v$ iff $u \mapsto v$. If, in addition, $u \in V_\bigcirc$, then $u \overset{x}{\hookrightarrow} v$ iff $u \overset{x}{\mapsto} v$.
- For all $u \in W$ such that $u \mapsto v$ for some $v \in V \setminus W$ there is a transition $u \hookrightarrow *$. If $u \in V_\bigcirc$, then $u \overset{x}{\hookrightarrow} *$ where x is the sum of the probabilities of all transition $u \mapsto v$ where $v \in V \setminus W$.
- There is a transition $* \hookrightarrow *$.

The strategy π induces a strategy for F for reaching $T \cap W$ along $S \cap W$ with probability at least $\mathit{Val}^\pi(u) - \delta$. Using the results for finite games [2] we get a MD strategy with the same lower bound on the probability of reaching T along S. This strategy induces at least one MD strategy σ for G such that $\mathit{Val}^\sigma(u) \geq \mathit{Val}^\pi(u) - \delta$.

(5) As we already noted, we did not manage to find a sufficiently simple self-contained proof for this claim. Nevertheless, it follows easily from Theorem 7.2.11 presented in [13] [1]. $\Box$

4 BPA games

A $1\frac{1}{2}$-*player BPA game* is a tuple $\Delta = (\Gamma, \hookrightarrow, (\Gamma_\Box, \Gamma_\bigcirc), \mathit{Prob})$ where Γ is a finite *stack alphabet*, $\hookrightarrow \subseteq \Gamma \times \Gamma^{\leq 2}$ is a set of *rules* (where $\Gamma^{\leq 2} = \{w \in \Gamma^* : |w| \leq 2\}$) such that for each $X \in \Gamma$ there is some $X \hookrightarrow \alpha$, $(\Gamma_\Box, \Gamma_\bigcirc)$ is a partition of Γ, and Prob is a *probability assignment* which to each $X \in \Gamma_\bigcirc$ assigns a rational positive probability distribution on the set of all rules of the form $X \hookrightarrow \alpha$.

[1] We would like to thank Kousha Etessami for providing us with this reference.

Each $1\frac{1}{2}$-player BPA game $\Delta = (\Gamma, \hookrightarrow, (\Gamma_\Box, \Gamma_\bigcirc), \mathit{Prob})$ determines a unique $1\frac{1}{2}$-player game $G_\Delta = (\Gamma^*, \mapsto, (\Gamma_\Box\Gamma^*, \Gamma_\bigcirc\Gamma^* \cup \{\varepsilon\}), \mathit{Prob}_\Delta)$ where the transitions of $\mapsto$ are determined as follows: $\varepsilon \mapsto \varepsilon$, and $X\beta \mapsto \alpha\beta$ iff $X \hookrightarrow \alpha$. The probability assignment Prob_Δ is the natural extension of Prob, i.e., $\varepsilon \stackrel{1}{\mapsto} \varepsilon$ and for all $X \in \Gamma_\bigcirc$ we have that $X\beta \stackrel{x}{\mapsto} \alpha\beta$ iff $X \stackrel{x}{\hookrightarrow} \alpha$. Given a configuration $X\alpha \in \Gamma^*$, we put $\mathit{head}(X\alpha) = X$.

Realize that all of the previously introduced game-theoretic notions (strategy, upper/lower value, etc.) apply to G_Δ, not directly to Δ. In particular, the vertices of G_Δ are stack configurations of Γ^*, which means that MD strategies generally depend on the whole sequence of symbols which form a given vertex. An MD strategy σ is *stackless (SMD)* if it depends just on the top-of-the-stack symbol of a given vertex. Also note that G_Δ is finitely branching.

In this paper we consider algorithmic issues for EROs in $1\frac{1}{2}$-player BPA games. Since the game G_Δ has infinitely many vertices, we need to restrict ourselves to subclasses of EROs which admit a finite and effective description.

A natural subclass of EROs are *termination objectives*, where $S = \Gamma^*$ and $T = \{\varepsilon\}$. In [7,8], it has been shown that $1\frac{1}{2}$-player BPA games with termination objectives have the following properties:

(a) There are optimal SMD minimizing and maximizing strategies.
(b) For each $X \in \Gamma$, the values $\mathit{Val}^+(X)$ and $\mathit{Val}^-(X)$ are expressible as the least solution of an effectively constructible system of non-linear min/max equations. This allows to express the values $\mathit{Val}^+(X)$ and $\mathit{Val}^-(X)$ in $(\mathbb{R}, +, *, \leq)$, i.e., first-order arithmetic of the reals.
(c) The problems whether $\mathit{Val}^+(\alpha) = x$, where $x \in \{0,1\}$, and whether $\mathit{Val}^-(\alpha) = x$, where $x \in \{0,1\}$, are solvable in polynomial time.

In this paper we consider $1\frac{1}{2}$-player BPA games with a more general subclass of EROs, where the sets S and T are *simple*:

Definition 5 *A set $M \subseteq \Gamma^*$ is* simple *iff there is a* characteristic set $C(M) \subseteq \Gamma$ *such that $M = \bigcup_{Y \in C(M)} \{Y\}\Gamma^*$. An ERO (S,T) is simple if S and T are simple.*

Note that termination objectives can be encoded as simple EROs, but a given $1\frac{1}{2}$-player BPA game must first be modified by introducing a new bottom-of-the-stack symbol.

The properties (a)–(c) stated above do not hold for $1\frac{1}{2}$-player BPA games with simple EROs. In particular, note the following:

(A) An optimal minimizing SMD strategy may not exist, though there must be an optimal minimizing MD strategy by Proposition 4 (3). An optimal

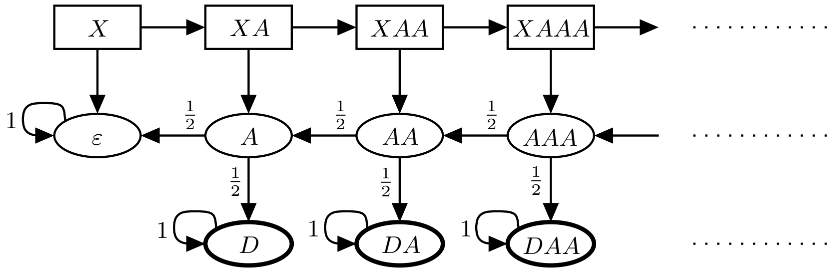

Fig. 2. A $1\frac{1}{2}$-player BPA game without an optimal maximizing strategy.

maximizing strategy may not exist at all (see also [7]). The existence of an optimal maximizing strategy implies the existence of an optimal maximizing MD strategy by Proposition 4 (5), but it does not imply the existence of an optimal maximizing SMD strategy.

(B) The system of non-linear min/max equations which was used in [7] for termination objectives cannot be immediately generalized to simple EROs. Intuitively, the reason is that the optimal minimizing/maximizing strategy in a configuration $X\alpha$ does not depend just on X but also on α, and a small modification of α may lead to a completely different optimal strategy. This is because one has to "balance" between the probability of termination and the probability of hitting a terminal configuration for each stack symbol, depending on what is achievable for the symbols stored below in the stack.

(C) For a given $\alpha \in \Gamma^*$, the problems whether $Val^-(\alpha) = 0$, whether $Val^+(\alpha) = 0$, and whether $Val^-(\alpha) = 1$ are solvable in polynomial time. The decidability of the problem whether $Val^+(\alpha) = 1$ is left open. Nevertheless, we show that the problem whether there is an optimal maximizing strategy σ such that $Val^\sigma(\alpha) = 1$ is decidable in polynomial time (remember that $Val^+(\alpha)$ can be 1 even if no optimal maximizing strategy exists).

The property (A) is demonstrated in the following example.

Example 6

(i) *Let $\Delta = (\{X, A, D\}, \hookrightarrow, (\{X\}, \{A, D\}), Prob)$ be a $1\frac{1}{2}$-player BPA game, where*

$$X \hookrightarrow XA,\ X \hookrightarrow \varepsilon,\ A \stackrel{1/2}{\hookrightarrow} D,\ A \stackrel{1/2}{\hookrightarrow} \varepsilon,\ D \stackrel{1}{\hookrightarrow} D$$

Let us consider a simple ERO (S, T) where $C(S) = \{X, A\}$ and $C(T) = \{D\}$. The structure of G_Δ is shown in Fig. 2. One can easily verify that $Val^+(X) = 1$. However, for every HR strategy σ we have that $Val^\sigma(X) < 1$.

(ii) *Let $\Delta = (\{X, D\}, \hookrightarrow, (\{X\}, \{D\}), Prob)$ be a $1\frac{1}{2}$-player BPA game, where*

$$X \hookrightarrow XD,\ X \hookrightarrow \varepsilon,\ D \stackrel{1}{\hookrightarrow} D$$

Let us consider a simple ERO (S, T) where $C(S) = \{X\}$ and $C(T) = \{D\}$. Then $Val^+(X) = 1$ and there is an optimal maximizing MD strategy, but there is no optimal maximizing SMD strategy.

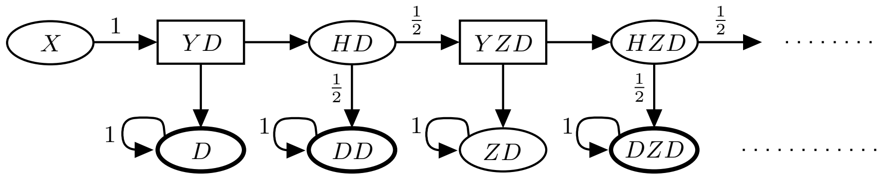

Fig. 3. A $1\frac{1}{2}$-player BPA game without an optimal minimizing SMD strategy.

(iii) Let $\Delta = (\{X,Y,Z,H,D\}, \hookrightarrow, (\{Y\},\{X,D,H,Z\}), Prob)$ be a $1\frac{1}{2}$-player BPA game, where

$$X \stackrel{1}{\hookrightarrow} YD,\ Y \hookrightarrow H,\ Y \hookrightarrow \varepsilon,\ D \stackrel{1}{\hookrightarrow} D,\ H \stackrel{1/2}{\hookrightarrow} YZ,\ H \stackrel{1/2}{\hookrightarrow} D,\ Z \stackrel{1}{\hookrightarrow} Z$$

Let us consider a simple ERO (S,T) where $C(S) = \{X,Y,H,Z\}$ and $C(T) = \{D\}$. The structure of G_Δ is shown in Fig. 3. Observe that $Val^-(X) = 1/2$ and there is an optimal minimizing MD strategy, but for every SMD strategy π we have that $\mathcal{P}(Run(G_\Delta[\pi], X, S\,\mathcal{U}\,T)) = 1$.

Now we present a sequence of results from which Property (C) follows as a simple consequence, and which allow to design the model-checking algorithm for $1\frac{1}{2}$-player BPA games and qualitative PCTL formulae presented in Section 5. Due to Proposition 4, from now on we can safely consider just MD strategies because they are equivalently powerful as HR strategies in the context of $1\frac{1}{2}$-player BPA games and simple qualitative EROs.

For the rest of this section, let us fix a $1\frac{1}{2}$-player BPA game $\Delta = (\Gamma, \hookrightarrow, (\Gamma_\Box, \Gamma_\bigcirc), Prob)$ and a simple ERO (S,T). We use T_ε to denote the set $T \cup \{\varepsilon\}$. For every $\bowtie \in \{<,>,=\}$ and every $\varrho \in \{0,1\}$, let

- $[S\,\mathcal{U}^{\bowtie\varrho}T]$ be the set of all $\alpha \in \Gamma^*$ for which there is a MD strategy σ such that $\mathcal{P}(G_\Delta[\sigma], \alpha, S\,\mathcal{U}\,T) \bowtie \varrho$.
- $[\mathcal{F}^{\bowtie\varrho}T]$ be the set of all $\alpha \in \Gamma^*$ for which there is a MD strategy σ such that $\mathcal{P}(G_\Delta[\sigma], \alpha, \mathcal{F}T) \bowtie \varrho$.

Remark 7 *In general, the sets $[S\,\mathcal{U}^{\bowtie\varrho}T]$ and $[\mathcal{F}^{\bowtie\varrho}T]$ constructed for given $\bowtie$ and ϱ are different. Now consider a modification $\Delta' = (\Gamma, \rightsquigarrow, (\Gamma_\Box, \Gamma_\bigcirc), Prob')$ of the game Δ obtained by replacing every rule of the form $P \hookrightarrow \alpha$, where $P \in C(T) \cup (\Gamma \smallsetminus C(S))$, with a single rule $P \rightsquigarrow P$ (if $P \in \Gamma_\bigcirc$, then $P \rightsquigarrow P$ has probability 1). It is easy to see that the set $[S\,\mathcal{U}^{\bowtie\varrho}T]$ is the* same *in Δ' as in Δ. Moreover, in Δ' we have that $[S\,\mathcal{U}^{\bowtie\varrho}T] = [\mathcal{F}^{\bowtie\varrho}T]$. Thus, $[S\,\mathcal{U}^{\bowtie\varrho}T]$ can be constructed by computing $[\mathcal{F}^{\bowtie\varrho}T]$ in a slightly modified $1\frac{1}{2}$-player BPA game, which leads to simplifications in our proofs.*

Our next four theorems show that the sets $[\mathcal{F}^{>0}T]$, $[\mathcal{F}^{=0}T]$, $[\mathcal{F}^{=1}T]$, and $[\mathcal{F}^{<1}T]$ are regular. The associated finite-state automata have a fixed number of control states and are effectively computable in polynomial time. We also study the relationship between MD and SMD strategies in this specific setting.

Due to Remark 7, all of the presented results can immediately be extended to the sets $[S\,\mathcal{U}^{>0}T]$, $[S\,\mathcal{U}^{=0}T]$, $[S\,\mathcal{U}^{=1}T]$, and $[S\,\mathcal{U}^{<1}T]$.

The difficulty of proofs is increasing. The cases $[\mathcal{F}^{>0}T]$ and $[\mathcal{F}^{=0}T]$ are simple, but the arguments for $[\mathcal{F}^{=1}T]$ and $[\mathcal{F}^{<1}T]$ are more involved.

Theorem 8 *There are $\mathcal{A}, \mathcal{B} \subseteq \Gamma$ computable in polynomial time such that $[\mathcal{F}^{>0}T] = \mathcal{A}^*\mathcal{B}\Gamma^*$. Moreover, there is a fixed SMD strategy π such that $\mathcal{P}(Run(G_\Delta[\pi], \alpha, \mathcal{F}T)) > 0$ for every $\alpha \in [\mathcal{F}^{>0}T]$.*

PROOF. Let $\mathcal{A} = \{X \in \Gamma \mid X \mapsto^* \varepsilon\}$ and $\mathcal{B} = \{X \in \Gamma \mid X \mapsto^* R\beta,$ where $R \in C(T)$ and $\beta \in \Gamma^*\}$. A straightforward induction on the length of α reveals that $\alpha \in [\mathcal{F}^{>0}T]$ iff $\alpha \in \mathcal{A}^*\mathcal{B}\Gamma^*$. The sets $\mathcal{A}$ and $\mathcal{B}$ can be computed as the least fixpoints of monotonic functions $\Theta_\mathcal{A}$ and $\Theta_\mathcal{B}$ defined as follows:

- $\Theta_\mathcal{A}(M) = M \cup \{X \in \Gamma \mid \text{there is } X \hookrightarrow \beta \text{ such that } \beta \in M^*\}$
- $\Theta_\mathcal{B}(M) = M \cup \{X \in \Gamma \mid \text{there is } X \hookrightarrow \beta \text{ such that } \beta \in \mathcal{A}^*(M \cup C(T))\Gamma^*\}$

One can easily verify that $\mathcal{A} = \bigcup_{i=1}^{|\Gamma|} \Theta_\mathcal{A}^i(\emptyset)$ and $\mathcal{B} = \bigcup_{i=1}^{|\Gamma|} \Theta_\mathcal{B}^i(\emptyset)$.

The function $\Theta_\mathcal{A}$ determines a SMD strategy $\pi_\mathcal{A}$ where, for each $X\beta \in \Gamma_\Box\Gamma^*$ such that $X \in \Theta_\mathcal{A}^{i+1}(\emptyset) \smallsetminus \Theta_\mathcal{A}^i(\emptyset)$, we have that $f_{\pi_\mathcal{A}}(X\beta) = \xi\beta$ where $X \hookrightarrow \xi$ is a rule witnessing the membership of X to $\Theta_\mathcal{A}^{i+1}(\emptyset)$. If $X \notin \mathcal{A}$, then $f_{\pi_\mathcal{A}}(X\beta)$ is defined arbitrarily. Note that for every $\alpha \in \mathcal{A}^*$ we have that $\mathcal{P}(Run(G_\Delta[\pi_\mathcal{A}], \alpha, \mathcal{F}\{\varepsilon\})) > 0$.

Similarly, the function $\Theta_\mathcal{B}$ determines a SMD strategy $\pi_\mathcal{B}$ where, for each $X\beta \in \Gamma_\Box\Gamma^*$ such that $X \in \Theta_\mathcal{B}^{i+1}(\emptyset) \smallsetminus \Theta_\mathcal{B}^i(\emptyset)$ we have that $f_{\pi_\mathcal{B}}(X\beta) = \xi\beta$ where $X \hookrightarrow \xi$ is a rule witnessing the membership of X to $\Theta_\mathcal{B}^{i+1}(\emptyset)$. If $X \notin \mathcal{B}$, then $f_{\pi_\mathcal{A}}(X\beta)$ is defined arbitrarily.

The SMD strategy π can now be defined as follows: for a given $X\beta \in \Gamma_\Box\Gamma^*$ we put

- $\pi(X\beta) = \pi_\mathcal{B}(X\beta)$ if $X \in \mathcal{B}$;
- $\pi(X\beta) = \pi_\mathcal{A}(X\beta)$ if $X \in \mathcal{A} \smallsetminus \mathcal{B}$;
- for the other arguments, $\pi(X\beta)$ is defined arbitrarily (but consistently with the requirement that π is SMD).

A straightforward induction on the length of α confirms that $\mathcal{P}(Run(G_\Delta[\pi], \alpha, \mathcal{F}T)) > 0$ for every $\alpha \in \mathcal{A}^*\mathcal{B}\Gamma^*$. $\Box$

Theorem 9 *There are $\mathcal{A}, \mathcal{B} \subseteq \Gamma$ computable in polynomial time such that $[\mathcal{F}^{=0}T] = \mathcal{B}^* \cup (\mathcal{B}^*\mathcal{A}\Gamma^*)$. Moreover, there is a fixed SMD strategy π such that*

$\mathcal{P}(Run(G_\Delta[\pi], \alpha, \mathcal{F}T)) = 0$ *for every* $\alpha \in [\mathcal{F}^{=0}T]$.

PROOF. We define the sets $\mathcal{A}$ and $\mathcal{B}$ as follows:

- $X \in \mathcal{A}$ iff there is a MD strategy σ such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}T_\varepsilon)) = 0$
- $X \in \mathcal{B}$ iff there is a MD strategy σ such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}T)) = 0$

It is easy to verify that $\mathcal{A}$ and $\mathcal{B}$ satisfy the property that $[\mathcal{F}^{=0}T] = \mathcal{B}^* \cup \mathcal{B}^*\mathcal{A}\Gamma^*$.

We show that the sets $\mathcal{A}$ and $\mathcal{B}$ can be computed as the greatest fixpoint of a monotonic function $\Theta : 2^\Gamma \times 2^\Gamma \to 2^\Gamma \times 2^\Gamma$, where $\Theta((M, N)) = (M', N')$ is defined as follows:

- $X \in M'$ iff $X \in M \setminus C(T)$ and the following conditions are satisfied:
  - If $X \in \Gamma_\Box$, then there is a rule of one of the following forms: $X \hookrightarrow Y$ where $Y \in M$, or $X \hookrightarrow YZ$ where either $Y \in M$, or $Y \in N$ and $Z \in M$.
  - If $X \in \Gamma_\bigcirc$, then all rules of the form $X \hookrightarrow \alpha$ satisfy either $\alpha = Y$ where $Y \in M$, or $\alpha = YZ$ where either $Y \in M$, or $Y \in N$ and $Z \in M$.
- $X \in N'$ iff $X \in N \setminus C(T)$ and the following conditions are satisfied:
  - If $X \in \Gamma_\Box$, then there is a rule of one of the following forms: $X \hookrightarrow \varepsilon$, or $X \hookrightarrow Y$ where $Y \in N \cup M$, or $X \hookrightarrow YZ$ where either $Y \in M$, or $Y, Z \in N \cup M$.
  - If $X \in \Gamma_\bigcirc$, then all rules of the form $X \hookrightarrow \alpha$ satisfy either $\alpha = \varepsilon$, or $\alpha = Y$ where $Y \in N \cup M$, or $\alpha = YZ$ where either $Y \in M$, or $Y, Z \in N \cup M$.

We prove that

(1) $(\mathcal{A}, \mathcal{B})$ is a fixpoint of Θ.
(2) If (C, D) is a fixpoint of Θ, then $C \subseteq \mathcal{A}$ and $D \subseteq \mathcal{B}$.

(1) Let $\Theta(\mathcal{A}, \mathcal{B}) = (\mathcal{A}', \mathcal{B}')$. It suffices to prove that $\mathcal{A} \subseteq \mathcal{A}'$ and $\mathcal{B} \subseteq \mathcal{B}'$ because the other inclusions follow directly from the definition of Θ. Let $X \in \mathcal{B}$. Then surely $X \notin C(T)$, and let σ be a MD strategy such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}T)) = 0$. If $X \in \Gamma_\bigcirc$, then for all rules of the form $X \hookrightarrow \alpha$ we have that $\alpha \in [\mathcal{F}^{=0}T] = \mathcal{B}^* \cup \mathcal{B}^*\mathcal{A}\Gamma^*$. This means that $X \in \mathcal{B}'$ by definition of Θ. If $X \in \Gamma_\Box$, then the rule $X \hookrightarrow \alpha$ which is selected by the strategy σ at the vertex X must again satisfy $\alpha \in [\mathcal{F}^{=0}T] = \mathcal{B}^* \cup \mathcal{B}^*\mathcal{A}\Gamma^*$. Hence, we obtain $X \in \mathcal{B}'$ by definition of Θ. The inclusion $\mathcal{A} \subseteq \mathcal{A}'$ can be shown similarly.

(2) Let $\Theta((C, D)) = (C', D') = (C, D)$, and let σ be a MD strategy such that

- for all $X \in C \cap \Gamma_\Box$ we have that $f_\sigma(X) = \alpha$ where $X \hookrightarrow \alpha$ is a rule witnessing $X \in C'$;
- for all $X \in (D \setminus C) \cap \Gamma_\Box$ we have that $f_\sigma(X) = \alpha$ where $X \hookrightarrow \alpha$ is a rule witnessing $X \in D'$.

It is easy to show that for all $Y \in \Gamma_\Box$, $w \in Run(G_\Delta[\sigma], Y)$, and $i \in \mathbb{N}_0$ we have that

- $w(i) \in D^* C \Gamma^*$ whenever $Y \in C$;
- $w(i) \in D^* \cup D^* C \Gamma^*$ whenever $Y \in D$.

The rest follows from $(C \cup D) \cap C(T) = \emptyset$.

Hence, the sets $\mathcal{A}$ and $\mathcal{B}$ can be computed in polynomial time by a simple iterative algorithm. The SMD strategy π is easy to define. For each $X\beta \in \Gamma_\Box \Gamma^*$ such that $X \in \mathcal{A}$ (or $X \in \mathcal{B} \smallsetminus \mathcal{A}$) we put $f_\pi(X\beta) = \xi\beta$ where $X \hookrightarrow \xi$ is a rule witnessing that $X \in \mathcal{A}'$ (or $X \in \mathcal{B}'$, resp.), where $(\mathcal{A}', \mathcal{B}') = \Theta((\mathcal{A}, \mathcal{B}))$. For the other arguments, $f_\pi(X\beta)$ is defined arbitrarily. To see that this definition is correct, i.e., $\mathcal{P}(Run(G_\Delta[\pi], \alpha, \mathcal{F}T)) = 0$ for every $\alpha \in \mathcal{B}^* \cup (\mathcal{B}^* \mathcal{A} \Gamma^*)$, consider a game $\widehat{\Delta}$ obtained from Δ just by removing all non-probabilistic rules that are *not* employed in π. Let $(\widehat{A}, \widehat{\mathcal{B}})$ be the greatest fixpoint of Θ computed in $\widehat{\Delta}$. It follows directly from the definition of $\widehat{\Delta}$ and Θ that $(\widehat{\mathcal{A}}, \widehat{\mathcal{B}}) = (\mathcal{A}, \mathcal{B})$. Since $\widehat{\Delta}$ contains exactly one rule for every $X \in \Gamma_\Box$, there is only one strategy $\widehat{\pi}$. By applying the above results, we obtain that $\mathcal{P}(Run(G_{\widehat{\Delta}}[\widehat{\pi}], \alpha, \mathcal{F}T)) = 0$ for every $\alpha \in \mathcal{B}^* \cup (\mathcal{B}^* \mathcal{A} \Gamma^*)$. Since the Markov chains $G_{\widehat{\Delta}}[\widehat{\pi}]$ and $G_\Delta[\pi]$ are isomorphic, we get $\mathcal{P}(Run(G_\Delta[\pi], \alpha, \mathcal{F}T)) = 0$ as needed. $\Box$

Theorem 10 *There are $\mathcal{A}, \mathcal{B}, \mathcal{C} \subseteq \Gamma$ computable in polynomial time such that $[\mathcal{F}^{=1}T] = (\mathcal{B} \cup \mathcal{C})^* \mathcal{A} \Gamma^*$.*

PROOF. We consider the sets $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{C}$ defined as follows:

- $X \in \mathcal{A}$ iff there is a MD strategy σ such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}T)) = 1$
- $X \in \mathcal{B}$ iff there is a MD strategy σ such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}T_\varepsilon)) = 1$ and $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}T)) > 0$
- $X \in \mathcal{C}$ iff there is a MD strategy σ such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}\{\varepsilon\})) = 1$

The set $\mathcal{C}$ can be computed in polynomial time using the algorithm of [8]. Also observe that $[\mathcal{F}^{=1}T] = (\mathcal{B} \cup \mathcal{C})^* \mathcal{A} \Gamma^*$ and $[\mathcal{F}^{=1}T_\varepsilon] = (\mathcal{B} \cup \mathcal{C})^* \cup (\mathcal{B} \cup \mathcal{C})^* \mathcal{A} \Gamma^*$.

We prove that the sets $\mathcal{A}$ and $\mathcal{B}$ are computable in polynomial time. To achieve that, we define a monotonic function $\Theta : 2^\Gamma \times 2^\Gamma \to 2^\Gamma \times 2^\Gamma$ such that $(\mathcal{A}, \mathcal{B})$ is the greatest fixpoint of Θ, and show that Θ (and hence also its greatest fixpoint) is computable in polynomial time.

We put $\Theta((R, H)) = (R', H')$, where R' (or H') is the set of all $X \in R$ (or all $X \in H$, resp.) for which there is a sequence $\mathcal{S}_X \equiv \alpha_0, \cdots, \alpha_n$ such that

- $\alpha_0 = X$, $\alpha_n \in T$;
- $\alpha_i \mapsto \alpha_{i+1}$ for all $0 \leq i < n$;

- for every $\beta \in \Gamma^*$ such that β either appears in the sequence $\mathcal{S}_X$ or $\alpha_i \mapsto \beta$ for some $\alpha_i \in \Gamma_{\bigcirc}\Gamma^*$ and $0 \leq i < n$ we have that $\beta \in (H \cup \mathcal{C})^* R \Gamma^*$ (or $\beta \in (H \cup \mathcal{C})^* R \Gamma^* \cup (H \cup \mathcal{C})^*$, resp.).

It follows directly from the definition that Θ is monotonic. It remains to show that $(\mathcal{A}, \mathcal{B})$ is the greatest fixpoint of Θ. First, we prove that $(\mathcal{A}, \mathcal{B})$ is a fixpoint. Let $\Theta((\mathcal{A}, \mathcal{B})) = (\mathcal{A}', \mathcal{B}')$. Since $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{B}' \subseteq \mathcal{B}$ by definition of Θ, it suffices to show the opposite inclusions. Let $X \in \mathcal{A}$ and let σ be a MD strategy which witnesses that $X \in \mathcal{A}$. Let us consider a path of minimal length in $G_\Delta[\sigma]$ from X to a configuration of T. Since every configuration reachable from X along a path which does not visit T belongs to $[\mathcal{F}^{=1}T] = (\mathcal{B} \cup \mathcal{C})^* \mathcal{A} \Gamma^*$, we can conclude $X \in \mathcal{A}'$. Similarly, we can show that $\mathcal{B} \subseteq \mathcal{B}'$.

Now suppose that $\Theta((R, H)) = (R', H') = (R, H)$. We prove that $R \subseteq \mathcal{A}$ and $H \subseteq \mathcal{B}$. For every $Y \in R$ (or $Y \in H$), let us fix a sequence $\mathcal{S}_Y$ which witnesses that $Y \in R'$ (or $Y \in H'$, resp.). It follows from the definition of Θ that if $Y \in R'$ (or $Y \in H'$) then all immediate successors of all stochastic configurations that appear in $\mathcal{S}_Y$ are of the form $(H \cup \mathcal{C})^* R \Gamma^*$ (or $(H \cup \mathcal{C})^* R \Gamma^* \cup (H \cup \mathcal{C})^*$, resp.).

Due to [8], there is a (SMD) strategy ϱ such that $\mathcal{P}(Run(G_\Delta[\varrho], X, \mathcal{F}\{\varepsilon\})) = 1$ for every $X \in \mathcal{C}$. Now we define a (MD) strategy π such that $\mathcal{P}(Run(G_\Delta[\pi], Y, \mathcal{F}^{=1}T)) = 1$ for all $Y \in R$. Let $X\xi \in \Gamma_\Box\Gamma^*$. If $X \in C(T)$, we define $f_\pi(X\xi)$ arbitrarily. Otherwise, let $\gamma \in \Gamma^*$ be the maximal prefix of $X\xi$ satisfying one of the following conditions:

(1) $\gamma = \mathcal{S}_Y(k)$ for some $Y \in R$ and $k \in \mathbb{N}_0$;
(2) $\gamma = \mathcal{S}_Y(k)$ for some $Y \in H$ and $k \in \mathbb{N}_0$ such that $\eta \in (H \cup \mathcal{C})^* R \Gamma^*$, where $X\xi = \gamma\eta$.

If there is no such γ, we either put $f_\pi(X\xi) = \varrho(X)\xi$ or define $f_\pi(X\xi)$ arbitrarily, depending on whether $X \in \mathcal{C}$ or not, respectively. Otherwise, we fix some $Y \in R \cup H$ and $k \in \mathbb{N}_0$ such that (1) or (2) is satisfied and $|\mathcal{S}_Y| - k$ is *minimal*. Note that γ cannot be the last configuration of $\mathcal{S}_Y$ because $X \notin C(T)$. Now we put $f_\pi(X\xi) = \mathcal{S}_Y(k{+}1)\eta$ where $X\xi = \gamma\eta$.

For all $Y \in R \cup H$ and $1 \leq i \leq |\mathcal{S}_Y|$, we put x_i to be either 1 or the probability of $\mathcal{S}_Y(i) \mapsto \mathcal{S}_Y(i{+}1)$, depending on whether $\mathcal{S}_Y(i) \in \Gamma_\Box\Gamma^*$ or not, respectively. Further, we define $\delta_Y = \prod_{i=1}^{|\mathcal{S}_Y|} x_i$ and $\delta = \min\{\delta_Y \mid Y \in R \cup H\}$. Clearly $\delta > 0$.

Let $Y \in R$. It follows directly from the definition of π that every $w \in Run(G_\Delta[\pi], Y)$ belongs either to $Run(G_\Delta[\pi], Y, \mathcal{F}T)$ or to $Run(G_\Delta[\pi], Y, \mathcal{G}((H \cup \mathcal{C})^* R \Gamma^* \smallsetminus T))$. However, almost all $w \in Run(G_\Delta[\pi], Y, \mathcal{G}((H \cup \mathcal{C})^* R \Gamma^* \smallsetminus T))$ visit infinitely often a configuration of the form $X\xi$ where $X \in H \cup R$, and from each such $X\xi$ a configuration of T can be visited with probability at least δ. This

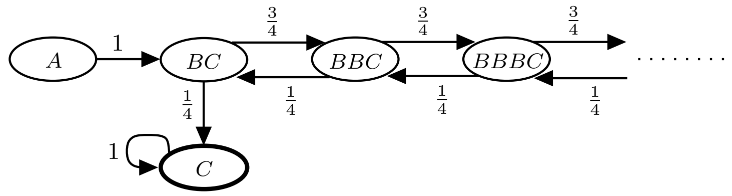

Fig. 4. The structure of $\bar{\Delta}$.

means that $\mathcal{P}(Run(G_\Delta[\pi], Y, \mathcal{G}((H \cup \mathcal{C})^* R\Gamma^* \smallsetminus T))) = 0$, and hence $\mathcal{P}(Run(G_\Delta[\pi], Y, \mathcal{F}T)) = 1$. Thus, we obtain $Y \in \mathcal{A}$.

The inclusion $H \subseteq \mathcal{B}$ can be shown similarly. We define a MD strategy π' such that $\mathcal{P}(Run(G_\Delta[\pi'], Y, \mathcal{F}T_\varepsilon)) = 1$ for all $Y \in H$. The strategy π' is defined in the same way as π. The only difference is that the condition (2) is relaxed to

(2') $\gamma = \mathcal{S}_Y(k)$ for some $Y \in H$ and $k \in \mathbb{N}_0$.

Then, for every $Y \in H$ we have that each $w \in Run(G_\Delta[\pi'], Y)$ either hits a configuration of T_ε, or visits only configurations of $((H \cup \mathcal{C})^* R\Gamma^* \cup (H \cup \mathcal{C})^*) \smallsetminus T_\varepsilon$. Again, one can easily show that the total probability of all runs of the second type is zero, hence $\mathcal{P}(Run(G_\Delta[\pi'], Y, \mathcal{F}T_\varepsilon)) = 1$ and $Y \in \mathcal{B}$ as needed.

We proved that $(\mathcal{A}, \mathcal{B})$ is the greatest fixpoint of Θ. Observe that $\Theta((R, H)) = (R', H')$ is computable in polynomial time, because the membership conditions for R' and H' are variants of simple reachability problems for non-probabilistic BPA that are solvable in polynomial time by standard techniques (for example, one can efficiently reduce these problems to the model-checking problem for BPA and a fixed CTL formula, which is decidable in polynomial time [14]). □

Let us note that, for a given $\alpha \in [\mathcal{F}^{=1}T]$, a SMD strategy π such that $\mathcal{P}(Run(G_\Delta[\pi], \alpha, \mathcal{F}T)) = 1$ does not necessarily exist. A counterexample is given in Example 6 (ii).

Theorem 11 *There are $\mathcal{A}, \mathcal{B} \subseteq \Gamma$ computable in polynomial time such that $[\mathcal{F}^{<1}T] = \mathcal{A}^* \cup (\mathcal{A}^* \mathcal{B} \Gamma^*)$.*

PROOF. Let us define the sets $\mathcal{A}$ and $\mathcal{B}$ as follows:

- $X \in \mathcal{A}$ iff there is a MD strategy σ such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}\{\varepsilon\})) > 0$
- $X \in \mathcal{B}$ iff there is a MD strategy σ such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \neg\mathcal{F}T_\varepsilon)) > 0$

It is easy to show that $[\mathcal{F}^{<1}T] = \mathcal{A}^* \cup (\mathcal{A}^* \mathcal{B} \Gamma^*)$. The set $\mathcal{A}$ is constructible in polynomial time by employing the algorithm of Theorem 8. For the rest of this proof we fix some $X \in \Gamma$ and examine the conditions under which $X \in \mathcal{B}$.

We say that $Y \in \Gamma \smallsetminus C(T)$ is a *type-I witness* if there are two MD strategies σ, π such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}(\{Y\}\Gamma^*))) > 0$ and $\mathcal{P}(Run(G_\Delta[\pi], Y, \mathcal{F}T_\varepsilon)) = 0$. The existence of such Y, σ, and π is obviously a sufficient condition for $X \in \mathcal{B}$, because the strategies σ and π can be combined into a single MD strategy σ' which behaves like σ until a configuration α such that $\mathcal{P}(Run(G_\Delta[\pi], \alpha, \mathcal{F}T_\varepsilon)) = 0$ is reached, and then it behaves like π. Obviously, $\mathcal{P}(Run(G_\Delta[\sigma'], X, \neg\mathcal{F}T_\varepsilon)) > 0$, because any configuration of the form $Y\gamma$ can play the role of α. Also note that the existence of a type-I witness can be decided in polynomial time using Theorem 8 and Theorem 9.

Unfortunately, $X \in \mathcal{B}$ does not necessarily imply the existence of a type-I witness, as illustrated in the following example: Let us consider a $1\frac{1}{2}$-player BPA game $\bar{\Delta} = (\{A, B, C\}, \hookrightarrow, (\emptyset, \{A, B, C\}), Prob)$ where

$$A \overset{1}{\hookrightarrow} BC,\ B \overset{3/4}{\hookrightarrow} BB,\ B \overset{1/4}{\hookrightarrow} \varepsilon,\ C \overset{1}{\hookrightarrow} C.$$

Note that $G_{\bar{\Delta}}$ closely resembles a one-dimensional asymmetric random walk (see Fig. 4). Let $T = \{C\}\Gamma^*$. Since $\Gamma_\Box = \emptyset$ in $\bar{\Delta}$, there is only one strategy (the empty strategy $\emptyset$). By applying standard result for one-dimensional random walks (see, e.g., [10]) we obtain that $\mathcal{P}(Run(G_{\bar{\Delta}}[\emptyset], A, \neg\mathcal{F}T_\varepsilon)) > 0$. However, for every $Y \in \{A, B\}$ we have that $\mathcal{P}(Run(G_{\bar{\Delta}}[\emptyset], Y, \mathcal{F}T_\varepsilon)) > 0$.

We say that $Y \in \Gamma \smallsetminus C(T)$ is a *type-II witness* if there is a $1\frac{1}{2}$-player BPA game $\Delta' = (\mathcal{D}, (\mathcal{D} \cap \Gamma_\Box, \mathcal{D} \cap \Gamma_\bigcirc), \twoheadrightarrow, Prob')$ and two MD strategies σ, π such that

- $Y \in \mathcal{D} \subseteq \Gamma$, $\mathcal{D} \cap C(T) = \emptyset$;
- $\twoheadrightarrow \subseteq \hookrightarrow$, where $Z \overset{x}{\twoheadrightarrow} \gamma$ iff $Z \in \mathcal{D} \cap \Gamma_\bigcirc$ and $Z \overset{x}{\hookrightarrow} \gamma$;
- $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}(\{Y\}\Gamma^*))) > 0$ and $\mathcal{P}(Run(G_{\Delta'}[\pi], Y, \neg\mathcal{F}\{\varepsilon\})) > 0$.

In the example above, B is a type-II witness for $A \in \mathcal{B}$ where $(\{B\}, (\emptyset, \{B\}), \twoheadrightarrow, Prob')$ is the associated $1\frac{1}{2}$-player BPA game.

Again, the existence of a type-II witness is obviously a sufficient condition for $X \in \mathcal{B}$, because the strategies σ and π can be combined into a single MD strategy σ' such that $\mathcal{P}(Run(G_\Delta[\sigma'], X, \neg\mathcal{F}T_\varepsilon)) > 0$. Also note that the existence of a type-II witness can be decided in polynomial time as follows:

- We compute the largest candidate for $\mathcal{D}$, which is the greatest fixpoint of $\Theta : 2^\Gamma \to 2^\Gamma$, where $Z \in \Theta(M)$ iff either $Z \in \Gamma_\bigcirc \smallsetminus C(T)$ and $\alpha \in M^*$ for all $Z \hookrightarrow \alpha$, or $Z \in \Gamma_\Box \smallsetminus C(T)$ and $\alpha \in M^*$ for some $Z \hookrightarrow \alpha$.
- If $\mathcal{D} = \emptyset$, there is no type-II witness for $X \in \mathcal{B}$. Otherwise, we put $\Delta' = (\mathcal{D}, (\mathcal{D} \cap \Gamma_\Box, \mathcal{D} \cap \Gamma_\bigcirc), \twoheadrightarrow, Prob')$ where
  · $Z \twoheadrightarrow \gamma$ iff $Z \hookrightarrow \gamma$, $Z \in \Gamma$, and $\gamma \in \mathcal{D}^*$
  · $Z \overset{x}{\twoheadrightarrow} \gamma$ iff $Z \in \mathcal{D} \cap \Gamma_\bigcirc$ and $Z \overset{x}{\hookrightarrow} \gamma$

It follows directly from the definition of Θ that Δ' is indeed $1\frac{1}{2}$-player BPA game.

- Now we decide if there are $Y \in \mathcal{D}$ and MD strategies σ, π with the required properties. The existence of σ can be easily checked in polynomial time (cf. Theorem 8). The existence of π can be decided in polynomial time by the algorithm presented in [8].

Now we prove that if $X \in \mathcal{B}$, then there is a type-I or type-II witness.

Let $X \in \mathcal{B}$ and let us fix a MD strategy σ such that $\mathcal{P}(Run(G_\Delta[\sigma], X, \neg\mathcal{F}T_\varepsilon)) > 0$. For every $\twoheadrightarrow \subseteq \hookrightarrow$, we denote $\mathcal{L}^{\twoheadrightarrow}$ the set of all $w \in Run(G_\Delta[\sigma], X, \neg\mathcal{F}T_\varepsilon)$ such that the set of all rules that induce infinitely many transitions of w is exactly $\twoheadrightarrow$. Since there are only finitely many subsets $\twoheadrightarrow \subseteq \hookrightarrow$ and $\bigcup_{\twoheadrightarrow \subseteq \hookrightarrow} \mathcal{L}^{\twoheadrightarrow} = Run(G_\Delta[\sigma], X, \neg\mathcal{F}T_\varepsilon)$, there must be some $\twoheadrightarrow \subseteq \hookrightarrow$ such that $\mathcal{P}(\mathcal{L}^{\twoheadrightarrow}) > 0$. For the rest of this proof, let us fix such a $\twoheadrightarrow$.

Let $\mathcal{D}$ be the set of all $Z \in \Gamma$ such that $Z \twoheadrightarrow \gamma$ for some $\gamma \in \Gamma^*$. Observe that $\mathcal{D} \cap C(T) = \emptyset$. Now we distinguish two possibilities:

(i) *There is a rule $Z \twoheadrightarrow \gamma$ such that $\gamma \notin \mathcal{D}^*$.* Then $\gamma = PQ$ where $P \in \mathcal{D}$ and $Q \notin \mathcal{D}$. We show that P is a type-I witness. Since $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}(\{P\}\Gamma^*))) > 0$, it suffices to prove that $\mathcal{P}(Run(G_\Delta[\pi], P, \mathcal{F}T_\varepsilon)) = 0$ for a suitable MD strategy π.

First we show that for every $\delta > 0$ there is a configuration of the form $PQ\beta$ such that $\mathcal{P}(Run(G_\Delta[\sigma], PQ\beta, \mathcal{F}\{Q\beta\})) < \delta$. Assume the converse. Then the probability of hitting a configuration starting with Q from a configuration starting with PQ is at least δ for some fixed $\delta > 0$. Since all $w \in \mathcal{L}^{\twoheadrightarrow}$ visit a configuration starting with PQ infinitely often, each $w \in \mathcal{L}^{\twoheadrightarrow}$ must visit a configuration starting with Q infinitely often, because $\mathcal{P}(\mathcal{L}^{\twoheadrightarrow})$ would be zero otherwise. However, this is a contradiction with $Q \notin \mathcal{D}$.

For every configuration of the form $PQ\beta$ we define a MD strategy $\pi_{Q\beta}$ by putting $f_{\pi_{Q\beta}}(\gamma) = f_\sigma(\gamma Q\beta)$ for every $\gamma \in \Gamma_\Box\Gamma^*$. It is easy to see that $\mathcal{P}(Run(G_\Delta[\pi_{Q\beta}], P, \mathcal{F}\{\varepsilon\})) = \mathcal{P}(Run(G_\Delta[\sigma], PQ\beta, \mathcal{F}\{Q\beta\}))$. Since $\inf\{\mathcal{P}(Run(G_\Delta[\pi_{Q\beta}], P, \mathcal{F}\{\varepsilon\})) \mid PQ\beta \in \Gamma^*\} = 0$, we can apply Proposition 4 (3) and conclude that there is MD strategy π such that $\mathcal{P}(Run(G_\Delta[\pi], P, \mathcal{F}\{\varepsilon\})) = 0$.

(ii) *For every rule $Z \twoheadrightarrow \gamma$ we have that $\gamma \in \mathcal{D}^*$.* First, observe that for every $V \in \mathcal{D} \cap \Gamma_\bigcirc$ and every $V \hookrightarrow \eta$ we have that $V \twoheadrightarrow \eta$. This is because V appears infinitely many times along every run of $\mathcal{L}^{\twoheadrightarrow}$, and hence $\mathcal{P}(\mathcal{L}^{\twoheadrightarrow})$ would be zero if some $V \hookrightarrow \eta$ was used only finitely many times in every run of $\mathcal{L}^{\twoheadrightarrow}$. This means that $\Delta' = (\mathcal{D}, (\mathcal{D} \cap \Gamma_\Box, \mathcal{D} \cap \Gamma_\bigcirc), \twoheadrightarrow, Prob')$, where $V \stackrel{x}{\twoheadrightarrow} \eta$ iff $V \in \mathcal{D} \cap \Gamma_\bigcirc$ and $V \stackrel{x}{\hookrightarrow} \eta$, is a $1\frac{1}{2}$-player BPA game. We show that there is a type-II witness $Y \in \mathcal{D}$ where Δ' is the associated $1\frac{1}{2}$-player BPA game.

For every $w \in \mathcal{L}^{\twoheadrightarrow}$, let $v_w = w(0), \cdots, w(i)$ be the finite prefix of w where $i \in \mathbb{N}_0$ is the least index such that for every $j \geq i$ we have that $w(j) \rightarrow w(j+1)$ is induced by a rule of $\twoheadrightarrow$ and $|w(i)| \leq |w(j)|$. For a given $v \in FPath(G_\Delta[\sigma], X)$, let $\mathcal{L}_v^{\twoheadrightarrow}$ be the set of all $w \in \mathcal{L}^{\twoheadrightarrow}$ such that $v_w = v$. Since $\mathcal{P}(\mathcal{L}^{\twoheadrightarrow}) > 0$ and there are only countably many $v \in FPath(G_\Delta[\sigma], X)$, there is some $v \in FPath(G_\Delta[\sigma], X)$ such that $\mathcal{P}(\mathcal{L}_v^{\twoheadrightarrow}) > 0$. For the rest of this proof, we fix such a v.

Let $v = v(0), \cdots, v(i)$, where $v(i) = Y\beta$. Obviously $\mathcal{P}(Run(G_\Delta[\sigma], X, \mathcal{F}(\{Y\}\Gamma^*))) > 0$. It remains to find a suitable MD strategy π such that $\mathcal{P}(Run(G_{\Delta'}[\pi], Y, \neg\mathcal{F}\{\varepsilon\})) > 0$. For every $\gamma \in \Gamma_\Box\Gamma^*$ we put $f_\pi(\gamma) = f_\sigma(\gamma\beta)$. Now it suffices to realize that every $w \in \mathcal{L}_v^{\twoheadrightarrow}$ is of the form $w \equiv v(0), \cdots, v(i{-}1), \alpha_0\beta, \alpha_1\beta, \cdots$, where $\alpha_0 = Y$. Thus, each $w \in \mathcal{L}_v^{\twoheadrightarrow}$ determines a unique run $\mathcal{R}(w) = \alpha_0, \alpha_1, \cdots$ of $Run(G_{\Delta'}[\pi], Y)$. Since w does not hit $T \cup \{\beta\}$, the run $\mathcal{R}(w)$ does not hit T_ε. Let $\mathcal{R}(\mathcal{L}_v^{\twoheadrightarrow}) = \{\mathcal{R}(w) \mid w \in \mathcal{L}_v^{\twoheadrightarrow}\}$. It is easy to show that

$$\mathcal{P}(\mathcal{R}(\mathcal{L}_v^{\twoheadrightarrow})) = \frac{\mathcal{P}(\mathcal{L}_v^{\twoheadrightarrow})}{\mathcal{P}(Run(G_\Delta[\sigma], v))}$$

Hence, $\mathcal{P}(\mathcal{R}(\mathcal{L}_v^{\twoheadrightarrow})) > 0$. Since $\mathcal{R}(\mathcal{L}_v^{\twoheadrightarrow}) \subseteq Run(G_{\Delta'}[\pi], Y, \neg\mathcal{F}\{\varepsilon\})$, we obtain that $\mathcal{P}(Run(G_{\Delta'}[\pi], Y, \neg\mathcal{F}\{\varepsilon\})) > 0$ as needed. $\Box$

Example 6 (iii) demonstrates that, for a given $\alpha \in [\mathcal{F}^{<1}T]$, a SMD strategy π such that $\mathcal{P}(Run(G_\Delta[\pi], \alpha, \mathcal{F}T)) < 1$ does not necessarily exist.

5 Model-checking Qualitative PCTL for $1\frac{1}{2}$-player BPA Games

In this section we show that the results about $1\frac{1}{2}$-player BPA games with extended reachability objectives (see Section 3) can be used to design an essentially optimal model-checking algorithm for the qualitative fragment of PCTL and $1\frac{1}{2}$-player BPA games. For technical convenience, we restrict ourselves to *simple* valuations, where $\nu(p)$ is a simple set for each $p \in Ap$ (see Definition 5).

Infinite sets of stack configurations will be represented by deterministic finite-state automata (DFA) which read the stack bottom-up. Formally, a DFA is a tuple $\mathcal{F} = (Q, \Sigma, \delta, \hat{q}, F)$ where Q is a finite set of *control states*, Σ is a finite *input alphabet*, $\delta : (Q \times \Sigma) \rightarrow Q$ is a total *transition function*, $\hat{q} \in Q$ is the *initial state*, and $F \subseteq Q$ is a subset of *final* states. The function δ is extended to the elements of $Q \times \Sigma^*$ in the natural way. A word $w \in \Sigma^*$ is *accepted* by $\mathcal{F}$ iff $\delta(q_0, w) \in F$.

Let Δ be a $1\frac{1}{2}$-player BPA game with the stack alphabet Γ, and let $\mathcal{F}$ be a

DFA with the input alphabet Γ. We say that a stack configuration $\alpha \in \Gamma^*$ is *recognized* by $\mathcal{F}$ iff the *reverse* of α is accepted by $\mathcal{F}$. Note that stack configurations are traditionally written as words starting with the top-of-the-stack symbol, but for technical reasons we prefer to read them in the bottom-up (i.e., right to left) direction.

In the proof of our next theorem we use the standard technique of simulating DFA in the stack alphabet (see, e.g., [6]).

Theorem 12 *Let $\Delta = (\Gamma, \hookrightarrow, (\Gamma_\Box, \Gamma_\bigcirc), Prob)$ by a $1\frac{1}{2}$-player BPA game. Let ν be a simple valuation and Φ a qualitative PCTL formula. Then there is a DFA $\mathcal{F}_\Phi$ of size $|\Delta| \cdot 2^{\mathcal{O}(|\Phi|)}$ constructible in time which is polynomial in $|\Delta|$ and exponential in $|\Phi|$ such that for all $\alpha \in \Gamma^*$ we have that $\alpha \models^\nu \Phi$ iff α is recognized by $\mathcal{F}_\Phi$.*

PROOF. We proceed by induction on the structure of Φ. The cases when $\Phi \equiv p$, $\Phi \equiv \Phi_1 \wedge \Phi_2$, and $\Phi \equiv \neg\Phi_1$ follow immediately.

Let $\Phi \equiv \mathcal{X}^{=1}\Phi_1$, and let $\mathcal{F}_1 = (Q_1, \Gamma, \delta_1, \hat{q}, F_1)$ be the DFA associated with Φ_1. The automaton $\mathcal{F}$ associated with Φ should then recognize exactly all $\alpha \in \Gamma^*$ such that for every transition $\alpha \to \beta$ we have that β is recognized by $\mathcal{F}_1$. Hence, we put $\mathcal{F} = (Q_1 \cup Q_1', \Gamma, \delta, \hat{r}, Q_1')$, where $Q_1' = \{q' \mid q \in Q_1\}$ and the transition function δ is constructed as follows: Let $q \in Q_1$, $A \in \Gamma$, and let $t = \delta_1(q, A)$. If for all rules $A \hookrightarrow \gamma$ we have that $\delta_1(q, \gamma^r) \in F_1$ (where γ^r denotes the reverse of γ), then $\delta(q, A) = \delta(q', A) = t'$. Otherwise, $\delta(q, A) = \delta(q', A) = t$. The initial state $\hat{r}$ of $\mathcal{F}$ is either $\hat{q}'$ or $\hat{q}$, depending on whether ε is recognized by $\mathcal{F}_1$ or not, respectively.

The cases when $\Phi \equiv \mathcal{X}^{<1}\Phi_1$, $\Phi \equiv \mathcal{X}^{=0}\Phi_1$, and $\Phi \equiv \mathcal{X}^{>0}\Phi_1$ are handled similarly.

Now, let us consider the case when $\Phi \equiv \Phi_1 \mathcal{U}^{=1}\Phi_2$. Note that $\alpha \models^\nu \Phi_1 \mathcal{U}^{=1}\Phi_2$ iff there is no strategy σ such that $\mathcal{P}(Run(G_\Delta(\sigma), \alpha, [\![\Phi_1]\!]\,\mathcal{U}\,[\![\Phi_2]\!])) < 1$. Let $\mathcal{F}_1 = (Q_1, \Gamma, \delta_1, \hat{q}, F_1)$ and $\mathcal{F}_2 = (Q_2, \Gamma, \delta_2, \hat{r}, F_2)$ be the DFA associated with Φ_1 and Φ_2. We construct another DFA $\mathcal{F}$ which accepts exactly those $\alpha \in \Gamma^*$ for which there *exists* a strategy σ such that $\mathcal{P}(Run(G_\Delta(\sigma), \alpha, [\![\Phi_1]\!]\,\mathcal{U}\,[\![\Phi_2]\!])) < 1$. The desired DFA is then obtained simply by complementing the automaton $\mathcal{F}$. First we construct a $1\frac{1}{2}$-player BPA game $\bar{\Delta}$ which is obtained from Δ by encoding the automata $\mathcal{F}_1$, $\mathcal{F}_2$ into the stack alphabet and simulating them "on-the-fly". Formally, $\bar{\Delta} = (\bar{\Gamma}, \rightsquigarrow, (\bar{\Gamma}_\Box, \bar{\Gamma}_\bigcirc), Pr)$ where $\bar{\Gamma} = (\Gamma \cup \{\varepsilon\}) \times Q_1 \times Q_2$, $\bar{\Gamma}_\Box = \Gamma_\Box \times Q_1 \times Q_2$, $\bar{\Gamma}_\bigcirc = (\Gamma_\bigcirc \cup \{\varepsilon\}) \times Q_1 \times Q_2$ and the transition relation $\rightsquigarrow$ together with Pr are defined as follows (A, q, and r range over Γ, Q_1, and Q_2, respectively):

- $(A, q, r) \overset{x}{\rightsquigarrow} \varepsilon$ iff $A \overset{x}{\hookrightarrow} \varepsilon$;

- $(A,q,r) \stackrel{x}{\rightsquigarrow} (B,q,r)$ iff $A \stackrel{x}{\hookrightarrow} B$
- $(A,q,r) \stackrel{x}{\rightsquigarrow} (B,q',r')(C,q,r)$ iff $A \stackrel{x}{\hookrightarrow} BC$, $\delta_1(q,C) = q'$ and $\delta_2(r,C) = r'$;
- $(\varepsilon,q,r) \stackrel{1}{\rightsquigarrow} (\varepsilon,q,r)$.

For every configuration $\alpha \in \Gamma^*$ of the form $A_n \cdots A_1$ there is a unique configuration $[\alpha] \in \bar{\Gamma}^*$ of the form $(A_n,q_n,r_n)\cdots(A_1,q_1,r_1)(\varepsilon,\hat{q},\hat{r})$ where $q_1 = \hat{q}$, $r_1 = \hat{r}$, and for all $0 \leq i < n$ we have that $\delta_1(q_i,A_i) = q_{i+1}$ and $\delta_2(r_i,A_i) = r_{i+1}$. Note that for every $\alpha \in \Gamma^*$, the subgraphs of G_Δ and $G_{\bar{\Delta}}$ which consist of all vertices reachable from α and $[\alpha]$ are isomorphic. Let $S, T \subseteq \bar{\Gamma}^*$ be the simple sets where

- $C(S) = \{(x,q,r) \mid x \in \Gamma \cup \{\varepsilon\}, \delta_1(q,x) \in F_1, r \in Q_2\}$
- $C(T) = \{(x,q,r) \mid x \in \Gamma \cup \{\varepsilon\}, q \in Q_1, \delta_2(r,x) \in F_2\}$.

Now it is easy to see that $\{\alpha \in \Gamma^* \mid \exists\sigma : \mathcal{P}(\mathit{Run}(G_\Delta(\sigma),\alpha,[\![\Phi_1]\!]\,\mathcal{U}\,[\![\Phi_2]\!])) < 1\}$ is equal to the set $K = \{\alpha \in \Gamma^* \mid \exists\sigma : \mathcal{P}(\mathit{Run}(G_{\bar{\Delta}}(\sigma),[\alpha],S\,\mathcal{U}\,T)) < 1\}$. By Theorem 11 (see also Remark 7), there effectively exist the sets $\mathcal{A},\mathcal{B} \subseteq \bar{\Gamma}$ such that $K = \{\alpha \in \Gamma^* \mid [\alpha] \in \mathcal{A}^* \cup (\mathcal{A}^*\mathcal{B}\bar{\Gamma}^*)\}$. Hence, the automaton $\mathcal{F}$ recognizing the set K can now be constructed as follows: we put $\mathcal{F} = (Q,\Gamma,\delta,\hat{t},F)$ where

- $Q = Q_1 \times Q_2 \times \{0,1\}$.
- For all $A \in \Gamma$, $q \in Q_1$, $r \in Q_2$, and $i \in \{0,1\}$ we put $\delta((q,r,i),A) = (\delta_1(q,A),\delta_2(r,A),j)$, where
  - $j = 0$ iff either $i = 0$ and $(q,r,A) \in \mathcal{A} \cup \mathcal{B}$, or $i = 1$ and $(q,r,A) \in \mathcal{B}$;
  - $j = 1$ iff either $i = 0$ and $(q,r,A) \in \Gamma \smallsetminus (\mathcal{A} \cup \mathcal{B})$, or $i = 1$ and $(q,r,A) \in \Gamma \smallsetminus \mathcal{B}$.
- The initial state $\hat{t}$ is either $(\hat{q},\hat{r},0)$ or $(\hat{q},\hat{r},1)$, depending on whether $(\varepsilon,\hat{q},\hat{r}) \in \mathcal{A} \cup \mathcal{B}$ or not, respectively.
- $F = Q_1 \times Q_2 \times \{0\}$.

The cases when $\Phi \equiv \Phi_1 \mathcal{U}^{=0} \Phi_2$, $\Phi \equiv \Phi_1 \mathcal{U}^{>0} \Phi_2$, and $\Phi \equiv \Phi_1 \mathcal{U}^{<1} \Phi_2$ are handled similarly, using Theorem 8, 9, and 10, respectively.

The complexity of the whole algorithm is easy to evaluate (it suffices to consider the worst subcase $\Phi \equiv \Phi_1 \mathcal{U}^{\bowtie\varrho} \Phi_2$). $\square$

Since the model-checking problem for qualitative PCTL and fully probabilistic BPA (i.e., the subclass of $1\frac{1}{2}$-player BPA games where $\Gamma_\Box = \emptyset$) is known to be **EXPTIME**-hard [4], we obtain the following:

Corollary 13 *The model-checking problem for qualitative PCTL and $1\frac{1}{2}$-player BPA games is **EXPTIME**-complete. For each fixed formula, the problem becomes solvable in polynomial time.*

6 Conclusions

We have shown that the sets of all configurations of a given $1\frac{1}{2}$-player BPA game for which there exists a strategy such that the probability of all runs satisfying a given simple ERO is greater than zero, equal to zero, equal to one, and less than one, are regular. Moreover, the corresponding finite-state automata have a fixed number of control states and are effectively constructible in polynomial time. With the help of this result, we derived **EXPTIME**-completeness of the model-checking problem for $1\frac{1}{2}$-player BPA games and qualitative PCTL formulae.

One natural question we left open is the decidability of the problem whether $\mathit{Val}^+(\alpha) = 1$ for a given configuration α of a given $1\frac{1}{2}$-player BPA game. Note that Theorem 10 entails the decidability of a slightly different problem—we ask whether there is a strategy such that the probability of all runs satisfying a given simple ERO is equal to one. If this is the case, then obviously $\mathit{Val}^+(\alpha) = 1$. However, it can happen that $\mathit{Val}^+(\alpha) = 1$ even if no optimal maximizing strategy exists (see Example 6 (i)).

Note that the proofs of Theorem 10 and Theorem 11 are not fully constructive in the sense that the associated MD strategies are not explicitly designed (we are interested just in their existence). Although the transitions chosen by these strategies for a given $X\alpha$ generally depend both on X and α, we conjecture that the required information is actually finite and can effectively be encoded by a finite-state automaton.

References

[1] C. Baier and M. Kwiatkowska. Model checking for a probabilistic branching time logic with fairness. *Distributed Computing*, 11(3):125–155, 1998.

[2] A. Bianco and L. de Alfaro. Model checking of probabalistic and nondeterministic systems. In *Proceedings of FST&TCS'95*, volume 1026 of *Lecture Notes in Computer Science*, pages 499–513. Springer, 1995.

[3] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: application to model checking. In *Proceedings of CONCUR'97*, volume 1243 of *Lecture Notes in Computer Science*, pages 135–150. Springer, 1997.

[4] T. Brázdil, A. Kučera, and O. Stražovský. On the decidability of temporal properties of probabilistic pushdown automata. In *Proceedings of STACS'2005*, volume 3404 of *Lecture Notes in Computer Science*, pages 145–157. Springer, 2005.

[5] J. Esparza, A. Kučera, and R. Mayr. Model-checking probabilistic pushdown automata. In *Proceedings of LICS 2004*, pages 12–21. IEEE Computer Society Press, 2004.

[6] J. Esparza, A. Kučera, and S. Schwoon. Model-checking LTL with regular valuations for pushdown systems. *Information and Computation*, 186(2):355–376, 2003.

[7] K. Etessami and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. In *Proceedings of ICALP 2005*, volume 3580 of *Lecture Notes in Computer Science*, pages 891–903. Springer, 2005.

[8] K. Etessami and M. Yannakakis. Efficient qualitative analysis of classes of recursive Markov decision processes and simple stochastic games. In *Proceedings of STACS'2006*, volume 3884 of *Lecture Notes in Computer Science*, pages 634–645. Springer, 2006.

[9] E. Feinberg and A. Shwartz, editors. *Handbook of Markov Decision Processes.* Kluwer, 2002.

[10] W. Feller. *An Introduction to Probability Theory and Its Applications, Vol. 1.* Wiley, 1968.

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

[12] A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker. PRISM: a tool for automatic verification of probabilistic systems. In *Proceedings of TACAS 2006*, Lecture Notes in Computer Science. Springer, 2006. To appear.

[13] M.L. Puterman. *Markov Decision Processes.* Wiley, 1994.

[14] I. Walukiewicz. Pushdown processes: Games and model-checking. *Information and Computation*, 164(2):234–263, 2001.