



FI MU

Faculty of Informatics
Masaryk University Brno

Strategy Synthesis for Markov Decision Processes and Branching-Time Logics

by

Tomáš Brázdil
Vojtěch Forejt

FI MU Report Series

FIMU-RS-2007-03

Copyright © 2007, FI MU

July 2007

**Copyright © 2007, Faculty of Informatics, Masaryk University.
All rights reserved.**

**Reproduction of all or part of this work
is permitted for educational or research use
on condition that this copyright notice is
included in any copy.**

**Publications in the FI MU Report Series are in general accessible
via WWW:**

<http://www.fi.muni.cz/reports/>

Further information can be obtained by contacting:

**Faculty of Informatics
Masaryk University
Botanická 68a
602 00 Brno
Czech Republic**

Strategy Synthesis for Markov Decision Processes and Branching-Time Logics

Tomáš Brázdil* Vojtěch Forejt†

Faculty of Informatics, Masaryk University,

Botanická 68a, 60200 Brno,

Czech Republic.

{brazdil, forejt}@fi.muni.cz

July 23, 2007

Abstract

We consider a class of finite $1\frac{1}{2}$ -player games (Markov decision processes) where the winning objectives are specified in the branching-time temporal logic qPECTL^* (an extension of the qualitative PCTL^*). We study decidability and complexity of existence of a winning strategy in these games. We identify a fragment of qPECTL^* called detPECTL^* for which the existence of a winning strategy is decidable in exponential time, and also the winning strategy can be computed in exponential time (if it exists). Consequently we show that every formula of qPECTL^* can be translated to a formula of detPECTL^* (in exponential time) so that the resulting formula is equivalent to the original one over finite Markov chains. From this we obtain that for the whole qPECTL^* , the existence of a winning finite-memory strategy is decidable in double exponential time. An immediate consequence is that the existence of a winning finite-memory strategy is decidable for the qualitative fragment of PCTL^* in triple exponential time. We also obtain a single exponential upper bound on the same problem for the qualitative PCTL . Finally, we study the power of finite-memory strategies with respect to objectives described in the qualitative PCTL .

*Supported by “Institute for Theoretical Computer Science (ITI)”, project No. 1M0545.

†Supported by the Czech Science Foundation, project No. 102/05/H050.

1 Introduction

We study $1\frac{1}{2}$ -player games (Markov decision processes), which have been applied in various contexts, from computer science and engineering (models of network systems, models of industrial processes, etc.) to biology [13, 9, 12]. A $1\frac{1}{2}$ -player game G is a directed graph whose vertices are partitioned into two disjoint sets V_{\square} and V_{\circ} . For each vertex of V_{\circ} there is a fixed probability distribution on outgoing transitions. A *play* is initiated by putting a token on some vertex. This token is then moved from vertex to vertex by one ‘real’ player \square and one ‘virtual’ player \circ , who choose their moves in vertices of V_{\square} and V_{\circ} , respectively. Player \circ chooses his moves randomly according to the fixed distribution. Player \square chooses his moves according to a *strategy*. Generally, strategies may depend on history of the play and may be either randomized or deterministic (we denote HR and HD the classes of the history-dependent randomized and deterministic strategies, respectively). In this paper we also consider *finite-memory* strategies that depend on a finite-state information about the history of the play¹. The classes of randomized and deterministic finite-memory strategies are denoted FR and FD, respectively.

Once player \square fixes his strategy σ for the game G , we obtain a Markov chain $G(\sigma)$ where the states are finite paths in G , and $ws \xrightarrow{x} wst$ if and only if (s,t) is a transition in G and x is either the fixed probability assigned to (s,t) (if $s \in V_{\circ}$), or the probability of (s,t) assigned by player \square in ws . Now we may ask whether the resulting Markov chain $G(\sigma)$ satisfies a given property. A *winning objective* is a property of Markov chains to be achieved by player \square . A strategy σ is called *winning* if the Markov chain $G(\sigma)$ satisfies the winning objective.

Winning objectives can be expressed using various formalisms. For example, various kinds of linear-time objectives, such as Büchi, parity, and Rabin objectives, were intensively studied in the past (see, e.g., [7, 8, 6]). In this paper we concentrate on a different kind of winning objectives specified by formulae of a branching-time temporal logic.

Let us note that the semantics of branching-time formulae can be defined directly for $1\frac{1}{2}$ -player games (see, e.g., [2]). In that case strategies are chosen separately for each temporal operator occurring in a formula. This approach is different from the one taken in this paper and results on model-checking such games are not related to our results.

The problem of solving games with branching-time winning objectives was for the first time studied in [11], where the existence of a winning memoryless randomized strategy for objectives expressed in PCTL (see, e.g., [10]) was shown to be in **PSPACE**. Results of [11] were substan-

¹More formally, a finite-memory (randomized) strategy is represented by a deterministic finite-state automaton and a function which assigns a distribution on outgoing transitions to the current vertex of the play and the state of the automaton after reading the history of the play.

tially extended in [3] where also history dependent strategies were taken into account. The most relevant results of [3] are the following. First, the existence of a winning HD (and also HR, FR and FD) strategy is undecidable for $1\frac{1}{2}$ -player games with objectives specified in (quantitative) PCTL. Second, the problem of existence of a winning HD strategy is **EXPTIME**-complete for $1\frac{1}{2}$ -player games with objectives specified in the $L(F^{\neq 1}, G^{\neq 1}, F^{>0})$ ² fragment of the qualitative PCTL.

The question is whether the positive result about the fragment $L(F^{\neq 1}, G^{\neq 1}, F^{>0})$ can be extended to more expressive logics at least for finite-memory strategies. In this paper we address this problem and show that the existence of a winning finite-memory strategy is decidable even for a powerful temporal logic $q\text{PECTL}^*$. We also show that the winning finite-memory strategy can always be effectively synthesized. This problem is well motivated because in practice one usually does not only want to know whether a strategy exists but also wants to implement the strategy. Finite-memory strategies have the advantage of being easy to implement.

The logic $q\text{PECTL}^*$ is the qualitative fragment of the logic PECTL^* defined in [5]. PECTL^* is a generalization of the logic PCTL^* (see, e.g., [5, 2]) which is a probabilistic version of the well-known logic CTL^* . Of course, PECTL^* contains the logic PCTL . Hence, our results on $q\text{PECTL}^*$ have immediate consequences for the *qualitative* PCTL^* (denoted $q\text{PCTL}^*$) and the *qualitative* PCTL (denoted $q\text{PCTL}$).

Our contribution: The main results of this paper are summarized below.

- We show that the existence of a winning FR (or FD) strategy for objectives described by $q\text{PCTL}$, $q\text{PECTL}^*$, and $q\text{PCTL}^*$ formulae is decidable in single exponential, double exponential, and triple exponential time, respectively. We also show that the winning strategy can effectively be computed with the same complexity. Moreover, we show that all these problems can be solved in time polynomial in the size of games.
- In the course of the proof of the above results we identify a fragment of $q\text{PECTL}^*$, called detPECTL^* , and show that the existence of a winning HR (or HD) strategy for objectives described in detPECTL^* is decidable in time exponential in the size of formulae and polynomial in the size of games. The fragment detPECTL^* contains the logic $L(F^{\neq 1}, G^{\neq 1}, F^{>0})$, and hence our results improve on the corresponding results of [3] by considering a more general logic, randomized strategies, and providing a polynomial time upper bound in the size of games.

²Formulae of $L(F^{\neq 1}, G^{\neq 1}, F^{>0})$ are built up from literals using conjunction, disjunction, and the temporal operators $F^{\neq 1}, G^{\neq 1}, F^{>0}$ (negation is applied only to atomic propositions).

- Finally, it has been shown in [3] that an infinite-memory strategy is needed for satisfying a formula of the fragment $L(G^{>0}, F^{>0})$ of qPCTL. We extend this result and provide (in a sense) complete classification of the power of finite-memory strategies for various fragments of qPCTL.

Plan of the paper: In Section 2 we review basic definitions for Markov chains and games. We also introduce the logic qPECTL* and its fragments. In Section 3 we consider the problem of existence of a winning history-dependent strategy for objectives described in detPECTL*. In Section 4 we consider the same problem for finite-memory strategies and qPECTL*. Finally, Section 5 deals with the classification of fragments of qPCTL with respect to the power of finite-memory strategies.

2 Basic Definitions

In this section we introduce basic notions of Markov chains, probabilistic temporal logics, and games. Most definitions (except the definition of qPECTL*) are taken from [3].

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

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 $P : \mathcal{F} \rightarrow \mathbb{R}^{\geq 0}$ such that, for each countable collection $\{X_i\}_{i \in I}$ of pairwise disjoint elements of \mathcal{F} , $P(\bigcup_{i \in I} X_i) = \sum_{i \in I} P(X_i)$, and moreover $P(X) = 1$. A *probability space* is a triple (X, \mathcal{F}, P) where (X, \mathcal{F}) is a measurable space and P is a probability measure over (X, \mathcal{F}) .

2.1 Markov Chains

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

In the rest of this paper we also write $s \xrightarrow{x} t$ instead of $Prob(s \rightarrow t) = x$. A *path* in M is a finite or infinite sequence $w = s_0, s_1, \dots$ of states such that $s_i \rightarrow s_{i+1}$ for every i . The *length* of a given path w is the number of transitions in w . 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). The prefix

s_0, s_1, \dots, 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 M 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.

We say that a set $C \subseteq S$ is a bottom strongly connected component (BSCC) of M if for all $s, t \in C$ there is a path from s to t in M , and whenever there is a path from $s \in C$ to $t \in S$, then $t \in C$. Note that if we restrict the set of states of M to a BSCC C , we obtain a Markov chain.

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 probability space $(Run(s), \mathcal{F}, \mathbb{P})$ where \mathcal{F} is the σ -field generated by all basic cylinders $Run(w)$ where w starts with s , and $\mathbb{P} : \mathcal{F} \rightarrow [0, 1]$ is the unique probability measure such that $\mathbb{P}(Run(w)) = \prod_{i=0}^{m-1} x_i$ where $w = s_0, \dots, s_m$ and $s_i \xrightarrow{x_i} s_{i+1}$ for every $0 \leq i < m$ (if $m=0$, we put $\mathbb{P}(Run(w)) = 1$).

2.2 The Logic qPECTL*

A Büchi automaton is a tuple $B = (B, \Sigma, \delta, q_I, F)$, where Σ is a finite *alphabet*, B is a finite set of *states*, $\delta \subseteq B \times \Sigma \times B$ is a *transition relation* (we write $q \xrightarrow{a} q'$ instead of $(q, a, q') \in \delta$), q_I is the *initial state*, and $F \subseteq B$ is a set of *accepting states*. The automaton B is *deterministic* if for each $q \in B$ and each $a \in \Sigma$, there is at most one $q' \in B$ such that $q \xrightarrow{a} q'$.

The symbol Σ^ω denotes the set of all infinite words over the alphabet Σ . A *computation* of B on a word $w = w(0)w(1)\dots \in \Sigma^\omega$ is a sequence $\omega = q_0, q_1, \dots$ of states of B such that $q_0 = q_I$ and for all $i \geq 0$ we have $q_i \xrightarrow{w(i)} q_{i+1}$. A computation ω of B is *accepting* if a state of F occurs infinitely many times in ω . The automaton B accepts a word $w \in \Sigma^\omega$ if there exists an accepting computation of B on w . The set of all $w \in \Sigma^\omega$ accepted by B is denoted $L(B)$.

The logic qPECTL* has the following syntax:

$$\Phi ::= a \mid \neg a \mid B^{\sim\rho}(\Phi_1, \dots, \Phi_n)$$

Here a ranges over the set Ap of atomic propositions, $\sim\rho \in \{=1, <1, >0, =0\}$, $n \geq 1$, B is a Büchi automaton over an alphabet $\Sigma \subseteq 2^{\{1, \dots, n\}}$, and each Φ_i is a qPECTL* formula.

The semantics of qPECTL* formulae is defined below. Let $M = (S, \rightarrow, Prob)$ be a Markov chain and let $v : Ap \rightarrow 2^S$ be a valuation. We define $s \models^v a$ iff $s \in v(a)$, and $s \models^v \neg a$ iff $s \notin v(a)$. The semantics of a qPECTL* formula $\Phi = B^{\sim\rho}(\Phi_1, \dots, \Phi_n)$, where B is a Büchi automaton with the alphabet $\Sigma \subseteq 2^{\{1, \dots, n\}}$, is defined as follows: First, we can assume that the semantics of the qPECTL* formulae Φ_1, \dots, Φ_n has already been defined. For every state s of M , let $Run(s, \Phi)$ be the set of all runs $w \in Run(s)$ satisfying the following condition: There is a word $v \in L(B)$

such that for all $i \geq 0$ and all $k \in v(i)$ holds $w(i) \models^v \Phi_k$. We stipulate that $s \models^v \Phi$ if and only if $P(\text{Run}(s, \Phi)) \sim \rho$.

We say that formulae Φ and Ψ are *equivalent* ($\Phi \equiv \Psi$) iff for each state s of an arbitrary Markov chain M and for arbitrary valuation v holds: $s \models^v \Phi$ iff $s \models^v \Psi$.

Remark 2.1. Note that once a formula $B \sim^p(\Phi_1, \dots, \Phi_n)$, a Markov chain M , and a valuation v are fixed, we can say that B (or any automaton with the alphabet $\Sigma \subseteq 2^{\{1, \dots, n\}}$) accepts a run w of M if there is a word $v \in L(B)$ such that for all $i \geq 0$ we have $\bigwedge_{k \in v(i)} w(i) \models^v \Phi_k$. Then, e.g., $P(\text{Run}(s, \Phi))$ is the probability that B accepts a run of $\text{Run}(s)$. We can also say that the automaton B goes from a state q_0 to q_{i+1} after reading a finite path s_0, \dots, s_i in M if there is a sequence q_0, \dots, q_{i+1} of states of B and a word X_0, \dots, X_i such that $q_j \xrightarrow{X_j} q_{j+1}$ and $\bigwedge_{k \in X_j} w(j) \models^v \Phi_k$ for all $0 \leq j \leq i$.


For computational purposes we assume that each formula is represented as a directed acyclic multigraph obtained from the parse tree of the formula by merging similar subtrees. For example, the formula $B_1 \sim^p(B_1 \sim^p(a, a, a), B_1 \sim^p(a, a, a), B_2 \sim^p(B_1 \sim^p(a, a, a)))$ is represented by a multigraph with four nodes n_1, n_2, n_3, n_4 labeled with $B_1 \sim^p, B_2 \sim^p, B_1 \sim^p, a$, respectively, and transitions: $n_1 \xrightarrow{1,2} n_3$ (here the numbers 1, 2 stand for the first and the second argument), $n_1 \xrightarrow{3} n_2$, $n_2 \xrightarrow{1} n_3$, $n_3 \xrightarrow{1,2,3} n_4$. Here n_1 corresponds to the whole formula.

Expressing other operators in qPECTL*. The logic qPECTL*, as defined above, is very powerful and succinct, and hence ideal for theoretical considerations. However, it is easier to express complex properties when we have some additional operators. We show that all operators of qPCTL can be expressed in qPECTL*. We define automata B_\vee, B_\wedge as follows:



It is easy to see that formulae $B_{\vee}^{-1}(\Phi_1, \Phi_2)$ and $B_{\wedge}^{-1}(\Phi_1, \Phi_2)$ are equivalent to logical disjunction and conjunction, respectively, of Φ_1 and Φ_2 . Hence, in what follows we write $\Phi_1 \vee \Phi_2$ and $\Phi_1 \wedge \Phi_2$ instead of $B_{\vee}^{-1}(\Phi_1, \Phi_2)$ and $B_{\wedge}^{-1}(\Phi_1, \Phi_2)$, respectively.

We also define Büchi automata representing ‘next’, ‘until’ and ‘release’ (the dual of ‘until’) operators:



We write $X \sim^p \Phi_1$, $\Phi_1 U \sim^p \Phi_2$ and $\Phi_1 R \sim^p \Phi_2$ instead of $B_X \sim^p(\Phi_1)$, $B_U \sim^p(\Phi_1, \Phi_2)$ and $B_R \sim^p(\Phi_1, \Phi_2)$, respectively. We also define ‘future’ and ‘globally’ operators as follows: Let tt and ff stand for

$a \vee \neg a$ and $a \wedge \neg a$, respectively, for some $a \in Ap$. Let $F^{\sim\rho}\Phi$ stands for $\text{tt}U^{\sim\rho}\Phi$, and let $G^{\sim\rho}\Phi$ stands for $\text{ff}R^{\sim\rho}\Phi$.

Given a formula of the form $B^{\sim\rho}(\Phi_1, \dots, \Phi_n)$, we write $\neg B^{\sim\rho}(\Phi_1, \dots, \Phi_n)$ to stand for $B^{\bowtie\rho}(\Phi_1, \dots, \Phi_n)$, where ' $\bowtie\rho$ ' is ' $=1$ ', ' <1 ', ' >0 ', or ' $=0$ ', depending on whether ' $\sim\rho$ ' is ' <1 ', ' $=1$ ', ' $=0$ ', or ' >0 ', respectively. This clearly corresponds to the logical operation of negation. Note that $\Phi_1 R^{\sim\rho} \Phi_2$ is equivalent to $\neg(\neg\Phi_1 U^{\sim\rho} \neg\Phi_2)$.

Now qPCTL is the fragment of qPECTL* consisting of all formulae of the following form:

$$\Phi ::= a \mid \neg a \mid \Phi_1 \vee \Phi_2 \mid \Phi_1 \wedge \Phi_2 \mid X^{\sim\rho}\Phi_1 \mid \Phi_1 U^{\sim\rho}\Phi_2 \mid \Phi_1 R^{\sim\rho}\Phi_2$$

Here $\sim\rho$ ranges over $\{=1, =0, <1, >0\}$.

One can also show that all formulae of qPCTL* (for definition see, e.g., [5]) can be translated to equivalent qPECTL* formulae. This translation employs the algorithm for translating LTL formulae to Büchi automata (see, e.g., [15]) which results in a single exponential blow-up in the size of formulae.

The logic detPECTL*. Now we define the *deterministic* fragment of qPECTL* (called detPECTL*), which generalizes the fragment $L(F^=1, F^{>0}, G^=1)$ defined in [3] (see also Section 5). This fragment (together with Theorem 3.4) plays the crucial role in the proof of Theorem 4.2.

Given an automaton B with an alphabet $\Sigma \subseteq 2^{\{1, \dots, n\}}$, we say that a state q of B is *terminal* iff there is a transition $q \xrightarrow{\emptyset} q$ and no transition of the form $q \xrightarrow{X} q'$ where $q \neq q'$ in B . A formula Φ of qPECTL* is a detPECTL* formula if all subformulae of Φ of the form $B^{\sim\rho}(\Phi_1, \dots, \Phi_n)$ satisfy the following conditions:

1. ' $\sim\rho$ ' is either ' $=1$ ' or ' >0 ';
2. if ' $\sim\rho$ ' is ' >0 ', then all accepting states of B are terminal;
3. All states q of B satisfy the following condition: For *distinct nonterminal* states q' and q'' such that $q \xrightarrow{A} q'$ and $q \xrightarrow{B} q''$, we have that $\bigwedge_{i \in A \cup B} \Phi_i$ is *not* satisfied in any state of any Markov chain for any valuation (this must hold even for $A = B$).

Observe that $X^=1, X^{>0}, U^=1, U^{>0}$, and $R^=1$ are operators of detPECTL*.

2.3 Games and Strategies

A $1\frac{1}{2}$ -player game (or *Markov decision process*) is a tuple $G = (V, E, (V_{\square}, V_{\circ}), Prob)$ where V is a finite set of *vertices*, $E \subseteq V \times V$ is a set of *transitions*, (V_{\square}, V_{\circ}) is a partition of V , and *Prob*

is a *probability assignment* which to each $v \in V_{\square}$ assigns a positive 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 a player \square who selects the moves in the V_{\square} vertices, and a “virtual” player \circlearrowleft who selects the moves in the V_{\circlearrowleft} vertices according to the corresponding probability distribution.

A *strategy* for player \square is a function σ which to each $vs \in V^*V_{\square}$ assigns a probability distribution on the set of outgoing transitions of s . We say that a strategy σ is *deterministic* if $\sigma(vs)$ is a Dirac distribution for each $vs \in V^*V_{\square}$. Consistently with [1, 11, 3], we use HR and HD to denote the classes of all (history-dependent randomized) strategies and (history-dependent) deterministic strategies, respectively. A special type of strategies are strategies with *finite-memory*, which are formally defined as pairs (A, f) where $A = (Q, V, \delta, q_0)$ is a deterministic finite-state automaton over the alphabet V of vertices and f is a function which to each pair $(q, s) \in Q \times V_{\square}$ assigns a probability distribution on the set of outgoing transitions of s . The pair (A, f) determines a unique strategy $\sigma(A, f)$ such that $\sigma(A, f)(vs) = f(q, s)$, where $q = \delta(q_0, vs)$. Intuitively, the states of A represent a finite memory of size $|Q|$ where selected properties of the history of a play are stored. We denote FR and FD the classes of all finite-memory strategies and finite-memory deterministic strategies, respectively.

Each strategy σ for player \square determines a unique *play* of the game G , which is a Markov chain $G(\sigma)$ where V^+ is the set of states, and $vs \xrightarrow{x} vst$ iff $(s, t) \in E$ and one of the following conditions holds:

- $s \in V_{\circlearrowleft}$ and $Prob(s, t) = x$;
- $s \in V_{\square}$ and $\sigma(vs)$ assigns x to (s, t) .

For every $vs \in V^*V$ we denote $last(vs) = s$. For every run w of $G(\sigma)$ and every $i \geq 0$, we define $w[i] = last(w(i))$ (note that $w(i)$ is a finite sequence of vertices of the game G).

An *objective* is a pair (ν, Φ) , where $\nu : Ap \rightarrow 2^V$ is a valuation and Φ a qPECTL* formula. Note that each valuation $\nu : Ap \rightarrow 2^V$ determines a valuation $\bar{\nu} : Ap \rightarrow 2^{V^+}$ defined by $\bar{\nu}(a) = \{vs \in V^*V \mid s \in \nu(a)\}$. For a given objective (ν, Φ) , each state of $G(\sigma)$ either does or does not satisfy Φ . A (ν, Φ) -*winning strategy* for player \square in a vertex $s \in V$ is a strategy σ such that $s \models^{\nu} \Phi$. We are interested in the following problem:

Synthesis problem: Given a vertex s and an objective (ν, Φ) ,

is there a (ν, Φ) -winning strategy for player \square in s ?

Moreover, if the winning strategy exists, then return its finite representation.

Remark 2.2. Let σ be a finite-memory strategy determined by (A, f) where $A = (Q, V, \delta, q_0)$. Observe, that the chain $G(\sigma)$ can be seen as an ‘unfolding’ of a finite Markov chain. Formally, let $\approx \subseteq V^+ \times V^+$ be an equivalence defined as follows: $u \approx v$ if and only if $\delta(q_0, u) = \delta(q_0, v)$ and $\text{last}(u) = \text{last}(v)$. Given $v \in V^+$ we denote $[v] = \{u \mid u \approx v\}$, the equivalence class of v , and we denote $V^+ / \approx = \{[v] \mid v \in V^+\}$. Let us define a finite Markov chain $\bar{G}(\sigma)$ where V^+ / \approx is a set of states, and $[v] \xrightarrow{x} [vs]$ in $\bar{G}(\sigma)$ if and only if $v \xrightarrow{x} vs$ in $G(\sigma)$. Each valuation $\nu : Ap \rightarrow 2^V$ determines a valuation $\bar{\nu} : Ap \rightarrow 2^{V^+ / \approx}$ defined by $\bar{\nu}(a) = \{[vs] \mid s \in \nu(a)\}$. Now, it is easy to verify that for every qPECTL* formula Φ , every valuation ν , and every state v of $G(\sigma)$, we have that $v \models^\nu \Phi$ iff $[v] \models^{\bar{\nu}} \Phi$.

3 The Synthesis Problem for detPECTL*

In this section we prove that the synthesis problem is decidable in exponential time for both HR and HD strategies and detPECTL* objectives. The proof follows similar lines as the analogous proof for HD strategies and the $L(F^{=1}, G^{=1}, F^{>0})$ fragment of qPCTL (see [3]). However, new technical difficulties arise from the use of automata connectives and randomization.

Similarly to [3], we reduce the synthesis problem for detPECTL* to the problem of solving $1\frac{1}{2}$ -player games for a different type of objectives (and strategies) defined as follows. Let $G = (V, E, (V_\square, V_\circ), Prob)$ be a $1\frac{1}{2}$ -player game. A *mixed* Büchi objective is a pair (P, O) where $P, O \subseteq V$. A strategy σ is (P, O) -winning in a vertex s iff *all* runs in $G(\sigma)$ initiated in s visit a vertex of P infinitely often, and *almost all* runs initiated in s visit a vertex of O infinitely often. To be able to control randomization in games we introduce a new restriction on strategies defined as follows: Given a set $\mathfrak{X} \subseteq V_\square$, we say that a (HR) strategy σ is \mathfrak{X} -*must* iff for every $\nu \in V^*$, each $s \in \mathfrak{X}$ and each $(s, t) \in E$ we have $\sigma(\nu s)(s, t) > 0$, and for each $t \in V_\square \setminus \mathfrak{X}$ we have that $\sigma(\nu t)$ is a Dirac distribution. Intuitively, a strategy is \mathfrak{X} -must if it always assigns non-zero probability to all successors of vertices of \mathfrak{X} and behaves deterministically in vertices of $V_\square \setminus \mathfrak{X}$.

Let us fix a game $G = (V, E, (V_\square, V_\circ), Prob)$, a vertex s_{in} of G , a detPECTL* formula Φ , and a valuation ν . The following lemma allows us to assume that the branching degree of all vertices of G is at most two (we sketch the proof in Appendix).

Lemma 3.1. *There is a $1\frac{1}{2}$ -player game \bar{G} , a vertex \bar{s}_{in} , a formula $\bar{\Phi}$, and a valuation \bar{v} (computable in polynomial time), such that each vertex of \bar{G} has at most two successors, and there is a (v, Φ) -winning strategy in s_{in} iff there is a $(\bar{v}, \bar{\Phi})$ -winning strategy in \bar{s}_{in} . Moreover, each (v, Φ) -winning FR (FD) strategy in s_{in} can be polynomially translated to a $(\bar{v}, \bar{\Phi})$ -winning FR (FD) strategy in \bar{s}_{in} , and vice versa.*

We construct a game G' , a vertex s'_{in} of G' , a mixed Büchi objective (P, O) , and a set \mathfrak{R} , such that there is a (P, O) -winning \mathfrak{R} -must HR strategy in s'_{in} iff there is (v, Φ) -winning HR strategy in s_{in} . The size of G' will be single exponential in the size of Φ and polynomial in the size of G .

To simplify our presentation we introduce some additional notation. We say that a Büchi automaton B corresponds to a formula Ψ if and only if Ψ is of the form $B \sim^p(\Phi_1, \dots, \Phi_n)$ (for some \sim^p and Φ_1, \dots, Φ_n). For technical convenience, we assume that each Büchi automaton corresponds to at most one subformula of Φ and that all automata occurring in Φ have pairwise disjoint sets of states. Let $States$ denote the set of all states of all automata occurring in Φ , and let $States^{>0}$ and $States^{=1}$ denote sets of states of all automata that correspond to subformulae of Φ of the form $B^{>0}(\Phi_1, \dots, \Phi_n)$ and $B^{=1}(\Phi_1, \dots, \Phi_n)$, respectively. By $L(s)$ we denote the set of all literals (i.e., atomic propositions and negated atomic propositions) satisfied in the vertex s .

Now we present a formal definition of the game G' . An intuition behind the definition is given below.

Formal definition of G' . We define $G' = (V', E', (V'_\square, V'_\circ), Prob')$ where the set V' consists of vertices of the following three forms:

- f -vertices are of the form $(s, A)^f$, where $s \in V$ and $A \subseteq States \times \{\circ, \star\}$.
- g -vertices are of the form $(s, D)^g$, where

$$D \subseteq \{(t, B) \mid (s, t) \in E, B \subseteq States \times \{\circ, \star\}\}$$

is a non-empty set, for each t there is at most one pair of the form (t, B) in D , and if $s \in V_\circ$ and $(s, t) \in E$, then D contains a pair of the form (t, B) .

- distinguished vertices s'_{in} and $dead$.

To V'_\circ we put all g -vertices whose first component belongs to V_\circ , and we put $V'_\square = V' \setminus V'_\circ$. To formally define E' we need some additional notation. Given a formula of the form $\Psi = B \sim^p(\Phi_1, \dots, \Phi_n)$, we denote $Rep(\Psi)$ the tuple (q_0, \star) where q_0 is the initial state of B . Given a literal Ψ , we define $Rep(\Psi) = \Psi$. Given a transition $q \xrightarrow{X} q'$ of an automaton corresponding to

a formula of the form $B \sim^p(\Phi_1, \dots, \Phi_n)$, we denote $Start(q \xrightarrow{X} q')$ the set of all $Rep(\Phi_i)$ where $i \in X$. The set of transitions E' is defined as follows:

- $((s,A)^f, (s,D)^g) \in E'$ for all $(s,A)^f$ and $(s,D)^g$ satisfying the following: for every $(q,x) \in A$ there exists (q',x') satisfying $q \xrightarrow{X} q'$ and $Start(q \xrightarrow{X} q') \subseteq A \cup L(s)$ and

$$x' = \begin{cases} \star & \text{if } q' \text{ is accepting;} \\ \circ & \text{if } q' \in States^{\equiv 1} \text{ is not accepting and } A \cap (States^{\equiv 1} \times \{\circ\}) = \emptyset; \\ \circ & \text{if } q' \in States^{>0} \text{ is not accepting and } A \cap (States^{>0} \times \{\circ\}) = \emptyset; \\ x & \text{otherwise.} \end{cases}$$

and either $(q',x') \in \bigcap_{(t,B) \in D} B$ or $(q',x') \in \bigcup_{(t,B) \in D} B$, depending on whether $q \in States^{\equiv 1}$ or $q \in States^{>0}$, respectively.

- $((s,A)^f, dead) \in E'$ for all f -vertices, and $(dead, dead) \in E'$;
- $((s,D)^g, (t,B)^f) \in E'$ for all $(t,B) \in D$;
- $(s'_{in}, (s_{in}, A)^f) \in E'$ for all $A \subseteq States \times \{\circ, \star\}$ satisfying $Rep(\Phi) \in A$ (here we assume that Φ is not a literal, because otherwise the synthesis problem is trivially solved)

We define $Prob'((s,D)^g, (t,B)^f) = Prob(s,t)$ whenever $s \in V_{\square}$ and $(t,B) \in D$.

Let P be the set of all vertices of the form $(s,A)^f$ such that A does not contain any pair of the form (q,\circ) where $q \in States^{>0}$. Let O be the set of all vertices of the form $(s,A)^f$ such that A does not contain any pair of the form (q,\circ) where $q \in States^{\equiv 1}$.

Let \mathfrak{R} be the set of all g -vertices of the form $(s,D)^g$ where $s \in V_{\square}$.

Intuition behind G' . The game G' simulates the game G (in the first component of vertices) and at the same time maintains some information about subformulae of Φ (the second component of vertices). Each step of the simulated play of G corresponds to two steps in G' . The first step, going from the current f -vertex to a g -vertex, updates the information about Φ . The next one, going from the g -vertex to an f -vertex, simulates a move in G .

While playing the game G' , player \square simulates a play of the game G and at the same time simulates computations of Büchi automata corresponding to subformulae of Φ , verifying that these subformulae are satisfied in appropriate places in the simulated play. Given an f -vertex $(s,A)^f$, every pair $(q,x) \in A$ represents a running instance of an automaton which is in the state q . (Here x maintains the information whether this particular instance recently entered an accepting state.)

Going from the f -vertex $(s, A)^f$ to a g -vertex $(s, \{(t_1, B_1), (t_2, B_2)\})^g$, player \square simulates one computational step for each running instance $(q, x) \in A$. More concretely, let $(q, x) \in A$ where q is a state of an automaton corresponding to $B^{\sim\rho}(\Phi_1, \dots, \Phi_n)$. Player \square chooses a transition $q \xrightarrow{X} q'$ such that $Start(q \xrightarrow{X} q') \subseteq A \cup L(s)$, which intuitively means that the running instances of automata corresponding to formulae of $\{\Phi_i \mid i \in X\}$ have already been initiated in $(s, A)^f$. Then (q', x') (here x' is an appropriate update of x) is put either to both B_1, B_2 or to at least one of them, depending on whether ' $\sim\rho$ ' is either ' $=1$ ' or ' >0 ', respectively. Note that in the case ' $=1$ ', the simulated computation goes to the same state q' for both successors. To ensure correctness of the simulation, we need Φ to be a detPECTL^* formula (intuitively this means that there is at most one 'correct' non-terminal successor q' of q after reading the current state).

Note that the definition of x' and the winning objective (P, O) ensure that *all* running instances of automata corresponding to formulae of the form $B^=1(\dots)$ are almost surely accepting (i.e., enter an accepting state infinitely many times), and that *all* running instances of automata corresponding to formulae of the form $B^{>0}(\dots)$ are surely accepting (i.e., enter a terminal accepting state). Note that the sets B_1 and B_2 may, in addition to the obligatory contents, contain arbitrary pairs from $States \times \{o, \star\}$. This may be used by player \square to initiate new running instances needed in the next simulation step to perform transitions of Büchi automata (see above).

Finally, going from the g -vertex $(s, \{(t_1, B_1), (t_2, B_2)\})^g$ to an f -vertex, player \square randomly chooses one of the successors $(t_1, B_1)^f, (t_2, B_2)^f$, by which he chooses a successor in the simulated play. The \mathfrak{R} -must restriction ensures that each of the successors is chosen with non-zero probability, which prevents player \square from erasing pairs of $States^{>0} \times \{o, \star\}$.

The following lemma is proved in Appendix.

Lemma 3.2. *There is a (v, Φ) -winning HR strategy in s_{in} if and only if there is a (P, O) -winning \mathfrak{R} -must HR strategy in s'_{in} . Moreover, each (P, O) -winning \mathfrak{R} -must FR strategy (A, f) in s'_{in} induces a (v, Φ) -winning FR strategy in s_{in} computable in time polynomial in the size of (A, f) .*

It has been shown in [3] that the existence of a winning HD strategy in $1\frac{1}{2}$ -player games with mixed Büchi objectives is decidable in polynomial time, and moreover, that the existence of a winning HD strategy in such games implies the existence of a winning FD strategy computable in polynomial time. By a slight modification of the proof from [3] we obtain the following analogy for \mathfrak{R} -must HR strategies:

Lemma 3.3. *The existence of a winning \mathfrak{R} -must HR strategy in $1\frac{1}{2}$ -player games with mixed Büchi objectives is decidable in polynomial time. Moreover, in these games, the existence of a winning \mathfrak{R} -must HR strategy implies the existence of a winning \mathfrak{R} -must FR strategy computable in polynomial time.*

Applying Lemma 3.2 and Lemma 3.3 we obtain the following theorem.

Theorem 3.4. *The existence of a winning HR (HD) strategy in $1\frac{1}{2}$ -player games with detPECTL^* objectives is decidable in time exponential in the size of formulae and polynomial in the size of games. Moreover, in these games, the existence of a winning HR (HD) strategy implies the existence of a winning FR (FD) strategy computable in time exponential in the size of formulae and polynomial in the size of games.*

Sketch. For HR strategies, the result follows immediately from Lemma 3.2 and Lemma 3.3. For HD strategies, it suffices to slightly modify the construction of the game G' by erasing all g -vertices $(s, D)^g$ such that $s \in V_\square$ and $|D| > 1$. Now for $\mathfrak{R} = \emptyset$, each \mathfrak{R} -must strategy in s'_{in} is deterministic. An inspection of the proof of Lemma 3.2 reveals that the lemma remains valid even for deterministic strategies. Now using Lemma 3.3 we obtain the desired result. \square

4 The Synthesis Problem for $q\text{PECTL}^*$ and Finite-Memory Strategies

In this section we show how to solve the synthesis problem for $q\text{PECTL}^*$ and finite-memory strategies. We show that, in fact, the logic $q\text{PECTL}^*$ and its fragment detPECTL^* are expressively equivalent over finite Markov chains, and then obtain the solution to the synthesis problem as an immediate corollary of our previous results. To formally capture this equivalence, we write $\Phi \equiv_{fin} \Psi$ whenever for arbitrary state s of arbitrary *finite* Markov chain and arbitrary valuation v , holds $s \models^v \Phi$ iff $s \models^v \Psi$. The main aim of this section is formalized by the following theorem.

Theorem 4.1. *For every $q\text{PECTL}^*$ formula Φ there is a detPECTL^* formula Ψ , computable in exponential time, such that $\Phi \equiv_{fin} \Psi$. Moreover, if Φ is a $q\text{PCTL}$ formula, then Ψ is computable in polynomial time.*

An immediate corollary of Theorem 4.1, Theorem 3.4, and Remark 2.2 is the following

Theorem 4.2. *For both FR and FD strategies, the synthesis problem for $q\text{PCTL}$, $q\text{PECTL}^*$, and $q\text{PCTL}^*$ objectives can be solved in single exponential, double exponential, and triple exponential time, respectively. Moreover, in all these cases, the synthesis problem can be solved in time polynomial in the size of games.*

The rest of this section is devoted to the proof of Theorem 4.1. Let us fix a qPECTL* formula Φ . First, observe that we may assume (w.l.o.g.) that for each subformula of Φ of the form $B \sim^p(\Phi_1, \dots, \Phi_n)$, every state s of an arbitrary Markov chain and an arbitrary valuation v , there is *exactly one* letter A in the alphabet of B such that $s \models^v \bigwedge_{i \in A} \Phi_i$. Indeed, if Φ does not have this property, it suffices to substitute each subformula of Φ of the form $B \sim^p(\Phi_1, \dots, \Phi_n)$ with a formula of the form $\bar{B} \sim^p(\Phi_1, \dots, \Phi_n, \neg\Phi_1, \dots, \neg\Phi_n)$, where \bar{B} is obtained from B by substituting each transition of the form $q \xrightarrow{A} q'$ with all transitions of the form $q \xrightarrow{A' \cup A''} q'$ where $A \subseteq A' \subseteq \{1, \dots, n\}$ and $A'' = \{m + n \mid m \in \{1, \dots, n\} \setminus A'\}$, and consequently making the transition function total. Observe that although the alphabet of \bar{B} may be exponentially larger than the alphabet of B , the number of states increases only by 1 due to making the transition function total. Observe also that the size of the (graph representing) resulting formula is polynomial in the size of Φ .

Now, to determinize the formula Φ , it suffices to determinize (syntactically) transition relations of all Büchi automata in Φ . However, the problem is that deterministic Büchi automata are strictly weaker than non-deterministic Büchi automata. We solve this problem by first translating the Büchi automata to equivalent deterministic Rabin automata (using results of [14]) and then encoding the deterministic Rabin automata to detPECTL* formulae.

First, let us formally define the notion of deterministic Rabin automata. A *deterministic Rabin automaton* R is a tuple $(Q, \Sigma, \gamma, q_0, Acc)$, where Q is a finite set of states, Σ is an input alphabet, $\gamma: Q \times \Sigma \rightarrow Q$ is a transition function, q_0 is an initial state, and $Acc = \{(C_1, D_1), \dots, (C_k, D_k)\}$, where $C_1, \dots, C_k, D_1, \dots, D_k \subseteq Q$, specifies the acceptance condition. A word $w \in \Sigma^\omega$ is accepted by R iff there is a sequence $\omega = q_0, q_1, \dots$ of states of R such that for all $i \geq 0$ we have $\gamma(q_i, w(i)) = q_{i+1}$, and there is $j \in \{1, \dots, k\}$ such that some state of C_j occurs infinitely often in ω and no state of D_j occurs infinitely often in ω . We denote $L(R)$ the set of all words accepted by R . The following proposition was proved in [14].

Theorem 4.3 ([14]). *Given a Büchi automaton $B = (B, \Sigma, \delta, q_I, F)$ there is an effectively computable deterministic Rabin automaton $R = (Q, \Sigma, \gamma, q_0, Acc)$ such that $L(R) = L(B)$, $|R| = 2^{O(|B| \log |B|)}$ and $|Acc| = O(|B|)$.*

Now let $B \sim^p(\Phi_1, \dots, \Phi_n)$ be a subformula of Φ and let us assume that Φ_1, \dots, Φ_n are already detPECTL* formulae. Let us denote Σ the alphabet of B , and let us fix a Rabin automaton $R = (Q, \Sigma, \gamma, q_0, Acc)$, where $Acc = \{(C_1, D_1), \dots, (C_k, D_k)\}$, such that $L(R) = L(B)$. Let us assume (w.l.o.g.) that the transition function of R is total.

Let us first assume that ‘ $\sim\rho$ ’ is either of the form ‘ $>_0$ ’ or ‘ $=_1$ ’. Based on R , we define deterministic Büchi automata B_{fin} and $B_{q,i}$, for all $q \in Q$ and $1 \leq i \leq k$, such that

$$B^{\sim\rho}(\Phi_1, \dots, \Phi_n) \equiv_{fin} B_{fin}^{\sim\rho}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$$

where each Ψ_j is of the form $B_{q,i}^{\bar{1}}(\Phi_1, \dots, \Phi_n)$ for one of the automata $B_{q,i}$ (i.e., $\ell = |Q| \cdot k$). In what follows we denote $index(q, i)$ the number j such that Ψ_j is the formula $B_{q,i}^{\bar{1}}(\Phi_1, \dots, \Phi_n)$.

Before we formally define B_{fin} and $B_{q,i}$, let us explain the intuition behind the definition. Let us fix a state s_0 of a finite Markov chain M and a valuation v , and let us assume that B (and hence also R) accepts a run of $Run(s_0)$ (see Remark 2.1) with a probability greater than 0 (the explanation is analogous for the probability $=_1$). Because M is finite, there is a finite path $v \in FPath(s_0)$ such that R accepts almost all runs of $Run(v)$. Here, however, using basic results of the theory of finite Markov chains, one can say even more. There is a finite path $v \in FPath(s_0)$ such that almost all $w \in Run(v)$ satisfy the following condition: the automaton R , after reading the prefix v of w , enters a state of C_j infinitely often and no state of D_j at all, for a suitable j . We define B_{fin} and $B_{q,i}$ so that $B_{fin}^{>_0}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$ expresses precisely this property.

The automata B_{fin} and $B_{q,i}$ are formally defined as follows. Let $B_{fin} = (Q \cup \{q_a\}, \Sigma \cup T, \delta_{fin}, q_0, \{q_a\})$, where $T = \{\{n+1\}, \dots, \{n+\ell\}\}$, and transitions of B are defined as follows:

- $q \xrightarrow{A} q'$ for all $A \in \Sigma$ and $q, q' \in Q$ such that $\gamma(q, A) = q'$;
- $q \xrightarrow{\{n+index(q,i)\}} q_a$ for all $q \in Q$ and all $1 \leq i \leq k$;
- $q_a \xrightarrow{\emptyset} q_a$;
- nothing else is a transition.

We define $B_{q,i} = (Q, \Sigma, \delta_{q,i}, q, C_i)$ where transitions are defined as follows: for all $q \in Q$ and all $A \in \Sigma$, we define $q \xrightarrow{A} q'$ if and only if $q, q' \notin D_i$ and $\gamma(q, A) = q'$ (i.e., there are no transitions leaving or entering states of D_i).

Lemma 4.4. *If ‘ $\sim\rho$ ’ is either of the form ‘ $>_0$ ’ or ‘ $=_1$ ’, then*

$$B^{\sim\rho}(\Phi_1, \dots, \Phi_n) \equiv_{fin} B_{fin}^{\sim\rho}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$$

Moreover, the right hand side formula is in $detPECTL^$.*

Sketch. The fact that $B_{fin}^{\sim\rho}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$ is a $detPECTL^*$ formula follows immediately from our assumption about Φ and from the fact that the automata B_{fin} and $B_{q,i}$ are obtained from

the deterministic automaton R either by deleting transitions or by adding transitions to the newly added terminal state q_a .

It remains to prove the equivalence. Let us fix a finite Markov chain M with the set of states S , a state s_0 of M , and a valuation v . Observe that for every state s of M there is exactly one $A_s \in \Sigma$ such that $s \models \bigwedge_{i \in A_s} \Phi_i$. Now an automaton with the alphabet Σ accepts a run s_0, s_1, \dots of M if it accepts the word $A_{s_0} A_{s_1} \dots$.

First, let us assume that $s_0 \models B_{fin}^{\sim \rho}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$. It follows immediately from the definitions that, with probability $\sim \rho$, there is a path $s_0, s_1, \dots, s_i, s_{i+1}$ in M satisfying the following conditions:

- the automaton B_{fin} (and hence also the automaton R) moves from its initial state to a state r after reading the word $A_{s_0} \dots A_{s_i}$;
- $s_{i+1} \models B_{r,j}^{\leq 1}(\Phi_1, \dots, \Phi_n)$ for some $1 \leq j \leq k$ (and hence almost all runs of $Run(s_{i+1})$ are accepted by the Rabin automaton R initiated in r).

However, this immediately implies that, with probability $\sim \rho$, the automaton R (and hence the automaton B) accepts a run of $Run(s_0)$.

For the opposite direction, let $M \times R$ be a Markov chain (the synchronous product of M and R) whose set of states is $S \times Q$ and transitions are defined as follows: $(s, q) \xrightarrow{x} (t, r)$ iff $s \xrightarrow{x} t$ and $q \xrightarrow{A_s} r$. Given $1 \leq j \leq k$, we say that a BSCC C of M is j -accepting iff a state of C_j occurs in (the second component of a state of) C and no state of D_j occurs in C . Basic results of the theory of Markov chains imply that the probability measure of all runs of $Run(s)$ accepted by R (hence also by B) is equal to the probability of reaching some j -accepting BSCC of $M \times R$. Thus, if $s_0 \models B^{\sim \rho}(\Phi_1, \dots, \Phi_n)$, then, with probability $\sim \rho$, there is a path $(s_0, q_0), \dots, (s_i, q_i)$ in $M \times R$ such that (s_i, q_i) belongs to a j -accepting BSCC. It is easy to show that $s_i \models B_{q_i, j}^{\leq 1}(\Phi_1, \dots, \Phi_n)$ and B_{fin} moves from q_0 to q_a after reading the word $A_{s_0}, \dots, A_{s_{i-1}}, \{n + index(q_i, j)\}$. This implies that $s_0 \models B_{fin}^{\sim \rho}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$. \square

Now, let us consider the case where ' $\sim \rho$ ' is either of the form ' $=0$ ' or ' <1 '. We denote $\widehat{\sim \rho}$ the 'dual' of ' $\sim \rho$ ', i.e., $\widehat{=0}$ is ' $=1$ ', and $\widehat{<1}$ is ' >0 '. Using very similar arguments, we prove the following analogy of Lemma 4.4 (a proof is in Appendix).

Lemma 4.5. *If ‘ $\sim\rho$ ’ is either of the form ‘ $=0$ ’ or ‘ <1 ’, then there are Büchi automata B_{fin} and $B_{q,i}$, for all $q \in Q$ and $1 \leq i \leq k$, such that*

$$B^{\sim\rho}(\Phi_1, \dots, \Phi_n) \equiv_{fin} B_{fin}^{\widehat{\rho}}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$$

where each Ψ_j is of the form $B_{q,i}^{\overline{1}}(\Phi_1, \dots, \Phi_n)$ for one of the automata $B_{q,i}$ (i.e., $\ell = |Q| \cdot k$). Moreover, the right hand side formula is in $detPECTL^*$.

Using Lemma 4.4 and Lemma 4.5 one can easily design an algorithm which transforms the formula Φ to a $detPECTL^*$ formula Ψ using appropriate substitutions in a ‘bottom-up’ manner.

For general $qPECTL^*$ formulae, the time complexity of the algorithm is single exponential in the size of Φ . Indeed, by [14] the Rabin automaton R can be computed in time exponential in the number of states of B and polynomial in the size of the alphabet, and consequently the automata B_{fin} and $B_{q,i}$ are computable in single exponential time. It follows that the formula Ψ is computable in exponential time (here we make use of the representation of Ψ as a directed acyclic graph, i.e., we assume that several occurrences of the same subformula are represented by one vertex).

Now observe that in the case of PCTL formulae, there are only five distinct Büchi automata used to define Boolean connectives and operators X , U , and R . It follows that the corresponding Rabin automata have bounded size, and thus each PCTL formula can be translated to a $detPECTL^*$ formula in polynomial time. This finishes the proof of Theorem 4.1.

5 qPCTL and Finite-Memory Strategies

In this section we study the power of finite-memory strategies w.r.t. the synthesis problem for $1\frac{1}{2}$ -player games and $qPCTL$ objectives.

Given $Y \subseteq \{X^{\sim\rho}, F^{\sim\rho}, G^{\sim\rho} \mid \sim\rho \in \{=0, >0, <1, =1\}\}$, we denote $L(Y)$ the fragment of $qPCTL$ which consists of formulae of the following form:


$$\Phi ::= a \mid \neg a \mid \Phi_1 \vee \Phi_2 \mid \Phi_1 \wedge \Phi_2 \mid Y^{\sim\rho} \Phi_1$$

where $Y^{\sim\rho} \in Y$. For example, the fragment $L(\{F^=1, G^=1, F^{>0}\})$ (we usually omit the set brackets and write $L(F^=1, G^=1, F^{>0})$) is the fragment of $qPCTL$ whose formulae are built up from literals using conjunction, disjunction, and three temporal operators $F^=1, G^=1, F^{>0}$ (there is no operation of complement). Note that we work with the operators $F^{\sim\rho}$ and $G^{\sim\rho}$ only for simplicity: All results of this section remain valid even if one replaces $F^{\sim\rho}$ with $U^{\sim\rho}$, and $G^{\sim\rho}$ with $R^{\sim\rho}$.

Definition 5.1. A fragment $L(Y)$ is finitely determined if for every formula Φ of $L(Y)$, arbitrary vertex s of a $1\frac{1}{2}$ -player game, and arbitrary valuation v , the following holds: If there is a (v, Φ) -winning strategy in s , then there is a (v, Φ) -winning finite-memory strategy in s .

First, we show which fragments are *not* finitely determined. Then we prove that no finitely determined fragment is more expressive than the fragment $L(X^{=1}, X^{>0}, G^{=1}, F^{=1}, F^{>0})$.

It has been shown in [3] that $L(G^{>0}, F^{>0})$ is not finitely determined. We extend this result and show that also $L(G^{>0})$ is not finitely determined. Let us consider a game G depicted in the following figure (it is very similar to the corresponding game from [3]):



We use names of vertices as atomic propositions with an obvious semantics. Let us denote $\Phi = G^{>0}(\neg stop \wedge (\neg left \vee G^{>0} \neg right_2))$. The proof of the following lemma is presented in Appendix.

Lemma 5.2. *There is a (v, Φ) -winning HD strategy in start. There is no (v, Φ) -winning FR strategy in start.*

We continue by proving that also $L(G^{=0})$ is not finitely determined. Let us consider a formula $\Psi = G^{>0}(\neg stop \wedge (\neg left \vee F^{=1} G^{>0} \neg right_2))$ and the game G defined above. It is easy to show, using arguments similar to the proof of Lemma 5.2, that there is a (v, Ψ) -winning HD strategy in $start$, and no (v, Ψ) -winning FR strategy in $start$. Now we transform Ψ to a $L(G^{=0})$ formula Ψ' such that for an arbitrary strategy we have $start \models^v \Psi$ iff $start \models^v \Psi'$. Using obvious equivalences $\neg G^{>0} \phi \equiv G^{=0} \phi$ and $F^{=1} \phi \equiv G^{=0} \neg \phi$, one can easily show that Ψ is equivalent to $\neg \chi$ where $\chi = G^{=0}(\neg stop \wedge (\neg left \vee G^{=0} G^{=0} \neg right_2))$. Now it is easy to verify that $start \models^v \neg \chi$ if and only if $start \models^v G^{=0}(\neg start \vee \chi)$. It follows that $L(G^{=0})$ is not finitely determined.

Next, let us consider the fragment $L(F^{<1})$. Using the equivalence $G^{>0} \phi \equiv F^{<1} \neg \phi$, one can easily show that Ψ is equivalent to $F^{<1}(stop \vee (left \wedge F^{<1} F^{<1} right_2))$. It follows that $L(F^{<1})$ is not finitely determined.

The last fragments we analyze are $L(X^{=0}, F^{=1})$, $L(X^{<1}, F^{=1})$, $L(F^{=0}, F^{=1})$ and $L(G^{<1}, F^{=1})$. Using the equivalences $F^{=1} \phi \equiv G^{=0} \neg \phi$ and $\neg F^{=1} \phi \equiv G^{>0} \neg \phi$ one can show that $\Phi = G^{>0}(\neg stop \wedge (\neg left \vee G^{>0} \neg right_2))$ is equivalent to $\neg \chi$ where $\chi = F^{=1}(stop \vee (left \wedge F^{=1} right_2))$. Now, using a similar trick as above, we obtain $start \models^v \neg \chi$ iff $start \models^v X^{=0}(\chi)$ iff $start \models^v X^{<1}(\chi)$ iff $start \models^v F^{=0}(start \wedge \chi)$ iff $start \models^v G^{<1}(\neg start \vee \chi)$.

Now we can give a complete classification of finitely determined fragments.

Lemma 5.3. *The fragment $L_1 = L(X^{=1}, X^{<1}, X^{>0}, X^{=0}, G^{=1}, F^{>0}, F^{=0}, G^{<1})$ and the fragment $L_2 = L(X^{=1}, X^{>0}, G^{=1}, F^{>0}, F^{=1})$ are maximal (w.r.t. inclusion) finitely determined fragments.*

Proof. We have proved above that the fragments $L(G^{>0})$, $L(G^{=0})$, $L(F^{<1})$, $L(X^{=0}, F^{=1})$, $L(X^{<1}, F^{=1})$, $L(F^{=0}, F^{=1})$ and $L(G^{<1}, F^{=1})$ are not finitely determined. Clearly, any fragment which contains one of these fragments is not finitely determined. On the other hand, it follows from Theorem 3.4 that the fragment L_2 is finitely determined. By a close inspection of various possibilities, we obtain that the only fragment we have not yet classified is L_1 (and some of its subsets).

However, we show that each formula of L_1 can efficiently be translated to L_2 , which implies that L_1 is finitely determined. Let Ψ be a formula of L_1 . First, using the equivalences $X^{=0}\Phi \equiv X^{=1}\neg\Phi$, $X^{<1}\Phi \equiv X^{>0}\neg\Phi$, $F^{=0}\Phi \equiv G^{=1}\neg\Phi$, $G^{<1}\Phi \equiv F^{>0}\neg\Phi$, one can remove operators $X^{=0}$, $X^{<1}$, $F^{=0}$, $G^{<1}$ (introducing, however, some negations to the formula). The negations introduced in the previous step can be pushed to atomic propositions using equivalences $\neg X^{=1}\phi \equiv X^{>0}\neg\phi$, $\neg X^{>0}\phi \equiv X^{=1}\neg\phi$, $\neg G^{=1}\phi \equiv F^{>0}\neg\phi$, $\neg G^{>0}\phi \equiv F^{=1}\neg\phi$. \square

Corollary 5.4. *A fragment $L(Y)$, where $Y \subseteq \{X^{\sim\rho}, F^{\sim\rho}, G^{\sim\rho} \mid \sim\rho \in \{=0, >0, <1, =1\}\}$, is finitely determined if and only if each formula of $L(Y)$ can be efficiently translated to an equivalent formula of $L(X^{=1}, X^{>0}, G^{=1}, F^{=1}, F^{>0})$.*

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] A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Proceedings of FST&TCS'95*, vol. 1026 of *LNCS*, pp. 499–513. Springer, 1995.
- [3] T. Brázdil, V. Brožek, V. Forejt, and A. Kučera. Stochastic games with branching-time winning objectives. In *Proceedings of LICS 2006*. IEEE, 2006.
- [4] T. Brázdil and V. Forejt. Strategy synthesis for Markov decision processes and branching-time logics. Technical report FIMU-RS-2007-03, 2007.
- [5] 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.

- [6] 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.
- [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] E. Feinberg and A. Shwartz, editors. *Handbook of Markov Decision Processes*. Kluwer, 2002.
- [10] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6:512–535, 1994.
- [11] 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.
- [12] S. Mahadevan. Partially observable semi-Markov decision processes: Theory and applications in engineering and cognitive science. In *AAAI: Fall Symposium on Planning with Partially Observable Markov Decision Processes*, 1998.
- [13] M.L. Puterman. *Markov Decision Processes*. Wiley, 1994.
- [14] S. Safra. *Complexity of automata on infinite objects*. PhD thesis, 1989.
- [15] M. Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Banff Higher Order Workshop*, pp. 238–266, 1995.

A Appendix

A.1 Sketch of the proof of Lemma 3.1

We proceed as follows. Each vertex $s \in V$ together with its n successors t_1, \dots, t_n can be viewed as n -ary tree of depth 1 with s as a root and t_1, \dots, t_n as leaves. We change this tree to a binary tree while keeping s as a root and t_1, \dots, t_n as leaves. This requires adding new nodes to the tree and altering the probabilities of transitions.

Formally, we define the game $\bar{G} = (\bar{V}, \bar{E}, (\bar{V}_\square, \bar{V}_\circ), \bar{Prob})$ as follows: For each $s \in V$ having successors t_1, \dots, t_n , we put to \bar{V} the vertices s, t_1, \dots, t_n and new vertices t_2^s, \dots, t_{n-1}^s (i.e., if $n < 2$, then only the vertices s, t_1, \dots, t_n are put to \bar{V}). The transition relation \bar{E} is defined as follows. For each $s \in V$ with successors t_1, \dots, t_n where $n > 2$, we put the following transitions to \bar{E} : (s, t_1) , (s, t_2^s) , and $(t_2^s, t_3^s), \dots, (t_{n-2}^s, t_{n-1}^s)$, and $(t_2^s, t_2), (t_3^s, t_3), \dots, (t_{n-1}^s, t_{n-1})$, and (t_{n-1}^s, t_n) . For s with two successors t_1, t_2 , we put $(s, t_1), (s, t_2)$ to \bar{E} , and for s with one successor t , we put (s, t) to \bar{E} . There are no other vertices and transitions in \bar{G} than those defined above.

We define $V_\square = \{s, t_i^s \mid s \in V_\square\}$ and $V_\circ = \{s, t_i^s \mid s \in V_\circ\}$. Let $s \in V_\circ$ with successors t_1, \dots, t_n where $n > 2$. Let \bar{Prob} be a (unique) function such that $\bar{Prob}(s, t_1) = Prob(s, t_1)$ and $\bar{Prob}(t_i^s, t_i) = \frac{Prob(s, t_i)}{1 - \sum_{j=1}^{i-1} Prob(s, t_j)}$ for $1 < i < n$.

The formula $\bar{\Phi}$ is obtained from Φ by substituting (in a bottom up manner) each subformula of the form $B \sim^p(\Phi_1, \dots, \Phi_n)$ with a formula $B_o \sim^p(\Psi_1, \dots, \Psi_n, o, \neg o)$ where B_o is obtained from B by substituting each transition of the form $q \xrightarrow{A} q'$ with a transition $q \xrightarrow{A \cup \{n+2\}} q'$, and by adding self-loops $q \xrightarrow{\{n+1\}} q$ for each state q . Finally, the valuation \bar{v} is the same as v except that $\bar{v}(o) = \bar{V} \setminus V$. Intuitively, the automaton B_o simulates B on vertices of V and remains idle on the newly added vertices.

Given a sequence v of vertices of \bar{G} , we define $\Theta(v)$ the sequence of vertices of G obtained from v by erasing all vertices of $\bar{V} \setminus V$. Now each strategy σ for G determines a strategy $\bar{\sigma}$ for \bar{G} as follows: Let $v \in \bar{V}$ such that $last(v) = s \in V_\square$ and t_1, \dots, t_n are all successors of s in G . Then we define $\bar{\sigma}(v)(s, t_1) = \sigma(\Theta(v))(s, t_1)$, and $\bar{\sigma}(vt_2^s \dots t_i^s)(t_i^s, t_i) = \frac{\sigma(\Theta(v))(s, t_i)}{1 - \sum_{j=1}^{i-1} \sigma(\Theta(v))(s, t_j)}$ for $1 < i < n$. One can prove, using a straightforward induction on the structure of Φ , that $s_{in} \models^v \Phi$ in $G(\sigma)$ if and only if $s_{in} \models^{\bar{v}} \bar{\Phi}$ in $\bar{G}(\bar{\sigma})$.

Similarly, a strategy $\bar{\sigma}$ in \bar{G} determines a strategy σ in G as follows: Let $v \in \bar{V}$ such that $last(v) = s \in V_\square$ and t_1, \dots, t_n are all successors of s in G . We define $\sigma(\Theta(v))(s, t_1) = \bar{\sigma}(v)(s, t_1)$, and

$$\sigma(\Theta(v))(s, t_i) = \bar{\sigma}(v)(s, t_2^s) \cdot \prod_{j=2}^{i-1} \bar{\sigma}(vt_2^s \dots t_j^s)(t_j^s, t_{j+1}^s) \cdot \bar{\sigma}(vt_2^s \dots t_i^s)(t_i^s, t_i)$$

In other words, $\sigma(\Theta(v))(s, t_i)$ the probability of reaching $vt_2^s \cdots t_i^s t_i$ from v in $\bar{G}(\bar{\sigma})$. Similarly as above, $s_{in} \models^v \Phi$ in $G(\sigma)$ if and only if $s_{in} \models^{\bar{v}} \bar{\Phi}$ in $\bar{G}(\bar{\sigma})$. □

A.2 Proof of Lemma 3.2

A.2.1 From G' to G .

Let σ' be a \mathfrak{R} -must (P, O) -winning strategy in s'_{in} . By R we denote the set of all f -states of $G'(\sigma')$ that are reachable from s'_{in} (a state v of $G'(\sigma')$ is an f -state if $last(v)$ is an f -vertex). We define a function $\Lambda : R \rightarrow V^*$ inductively as follows:

- $\Lambda(s'_{in} \cdot \sigma'(s'_{in})) = s_{in}$
- $\Lambda(v \cdot (s, D)^g \cdot (t, A)^f) = \Lambda(v) \cdot t$

It is easy to show that Λ is injective.

We define a strategy σ in the game G as follows: For each $v \in \Lambda(R) \cap V^*V_{\square}$ the strategy $\sigma(v)$ assigns to (s, t) the probability that σ' assigns to $((s, D)^g, (t, A)^f)$ in $\Lambda^{-1}(v) \cdot (s, D)^g$ where $\Lambda^{-1}(v) \cdot (s, D)^g$ is the unique successor of $\Lambda^{-1}(v)$ in $G'(\sigma')$ (remember that σ' is \mathfrak{R} -must). For all other states we define σ arbitrarily. We show that such states are not reachable in $G(\sigma)$ from s'_{in} .

Lemma A.1. $\Lambda(R)$ is precisely the set of states reachable from s_{in} in $G(\sigma)$. Moreover, for all $v, v' \in \Lambda(R)$ we have $P(v \rightarrow^* v') = P(\Lambda^{-1}(v) \rightarrow^* \Lambda^{-1}(v'))$ ³.

Proof. Let $v \in \Lambda(R)$. We show that $v \xrightarrow{x} v'$ iff $v' \in \Lambda(R)$ and $\Lambda^{-1}(v) \xrightarrow{1} u \xrightarrow{x} \Lambda^{-1}(v')$, where u is the unique predecessor of $\Lambda^{-1}(v')$ in $G'(\sigma')$.

Assume that $last(v) = s$ and $last(\Lambda^{-1}(v)) = (s, A)^f$. There are two possibilities:

- If $s \in V_{\circ}$, then we have $v \xrightarrow{x} v'$ iff $v' = v \cdot t$ for some $t \in V$ such that $s \xrightarrow{x} t$ iff $\Lambda^{-1}(v) \xrightarrow{1} \Lambda^{-1}(v) \cdot (s, D)^g \xrightarrow{x} \Lambda^{-1}(v')$ where $\Lambda^{-1}(v') = \Lambda^{-1}(v) \cdot (s, D)^g \cdot (t, A)^f$
- If $s \in V_{\square}$, then we have $v \xrightarrow{x} v'$ iff $v' = v \cdot t$ for some $t \in V$ such that $\sigma(v)(s, t) = x$ iff $\Lambda^{-1}(v) \xrightarrow{1} \Lambda^{-1}(v) \cdot (s, D)^g \xrightarrow{x} \Lambda^{-1}(v')$ where $\Lambda^{-1}(v') = \Lambda^{-1}(v) \cdot (s, D)^g \cdot (t, A)^f$

The rest of the proof can be done by straightforward induction. □

³ $P(v \rightarrow^* v')$ is the probability of reaching v' from v in $G(\sigma)$, and $P(\Lambda^{-1}(v) \rightarrow^* \Lambda^{-1}(v'))$ is the probability of reaching $\Lambda^{-1}(v')$ from $\Lambda^{-1}(v)$ in $G'(\sigma')$.

In the rest of this proof we use the following notation: Given a Büchi automaton B and a state q of B , we denote B_q the Büchi automaton obtained from B by changing the initial state to q .

Lemma A.2. $s_{in} \models^v \Phi$

Proof. We prove, by induction on the structure of Φ , that for all subformulae of the form $B \sim^p(\Phi_1, \dots, \Phi_n)$, for all states q_0 of B and for all $v = v' \cdot (s_0, A_0)^f \in R$, we have $\Lambda(v) \models B \sim^p(\Phi_1, \dots, \Phi_n)$ whenever $(q_0, x) \in A_0$ for some $x \in \{o, \star\}$.

To simplify our notation, we say that a word $u \in (\mathcal{Z}^{\{1, \dots, n\}})^\omega$ is *consistent* with a run w in $G(\sigma)$ iff for every $i \geq 0$ we have $\bigwedge_{j \in u(i)} w(i) \models^v \Phi_j$.

First, let us assume that ' \sim^p ' is ' $>o$ '. Because σ' is (P, O) -winning, there is a finite path $w \in FPath(v)$ of length $2i$ such that

- $w[0] = (s_0, A_0)^f, w[2] = (s_1, A_1)^f, \dots, w[2i] = (s_i, A_i)^f$;
- there is a sequence $q_0, q_1, q_2, \dots, q_i$ of states of B such that q_i is accepting (and terminal) and for all $0 \leq j < i$ holds $(q_j, x_j) \in A_j$ and $q_j \xrightarrow{X_j} q_{j+1}$ where $Start(X_j) \subseteq A_j \cup L(s_j)$.

However, then, by induction, each run w' of $Run(\Lambda(v))$ satisfying $w^i = (w')^i$ is consistent with a word $u \in (\mathcal{Z}^{\{1, \dots, n\}})^\omega$ with the prefix X_0, X_1, \dots, X_n . It is clear that $u \in L(B_{q_0})$, and hence that $\Lambda(v) \models B_{q_0}^{>o}(\Phi_1, \dots, \Phi_n)$

Now, let us assume that ' \sim^p ' is ' $=1$ '. For almost all runs $w \in Run(v)$ we have that $w[i] \in O$ for infinitely many indexes i and hence

- $w[0] = (s_0, A_0)^f, w[2] = (s_1, A_1)^f, \dots, w[2j] = (s_j, A_j)^f, \dots$;
- there is a sequence q_0, q_1, q_2, \dots of states of B such that for all $j \geq 0$ holds $(q_j, x_j) \in A_j$ and $q_j \xrightarrow{X_j} q_{j+1}$ and $Start(X_j) \in A_j$;
- q_l is accepting for infinitely many $l \geq 0$.

However, then, by induction, the run s_0, s_1, \dots is consistent with the the word X_0, X_1, \dots . Finally, by Lemma A.1, we obtain that $\Lambda(v) \models B_{q_0}^{=1}(\Phi_1, \dots, \Phi_n)$. \square

Finally, let us assume that σ' is induced by a finite-memory strategy (A', f') where $A' = (Q', V', \delta', q'_0)$. We define a finite-memory strategy (A, f) , which induces σ , as follows. Let $A = (Q' \times V', V, \delta, (q'_0, s'_{in}))$ where the transition function δ is defined as follows: Let $q \in Q'$ and $(s, A)^f \in V'$, and let us assume that $f'(q, (s, A)^f)$ assigns $\mathbf{1}$ to a vertex $(s, D)^g$.

We define $\delta((q, (s, A)^f), t) = (q', (t, A')^f)$ where $\delta'(\delta'(q, (s, A)^f), (s, D)^g) = q'$. Other transitions of A are defined arbitrarily. The function f is defined as follows: $f(q, (s, A)^f)(s, t) = f'(q, (s, D)^g)((s, D)^g, (t, A')^f)$. It is easy to show that the strategy σ (see above) is induced by (A, f) .

A.2.2 From G to G' .

Let σ be a (v, Φ) -winning HR strategy in s_{in} . We construct \mathfrak{R} -must (P, O) -winning strategy in s'_{in} . By R we denote the set of all states of $G(\sigma)$ reachable from s_{in} . Given a state q of a Büchi automaton B which corresponds to a subformula $B \sim^P(\Phi_1, \dots, \Phi_n)$, we denote $Aut(q) = B_q$ and $Form(q) = B_q \sim^P(\Phi_1, \dots, \Phi_n)$. Observe that due to our assumptions, both Aut and $Form$ are well-defined functions.

We define a function $Dist : R \times States \rightarrow \mathbb{N}_0$ as follows: We define $Dist(v, q)$, where $Form(q) = B_q \sim^P(\Phi_1, \dots, \Phi_n)$, to be the length of a minimal path w of $FPath(v)$ such that the automaton $Aut(q)$ moves from q to an accepting state after reading w (see Remark 2.1). If there is no such a path, then $Dist(v, q)$ remains undefined. Given $A \subseteq States \times \{\circ, \star\}$, we define $Dist(v, A) = \sum_{(q, \circ) \in A \cap (States^{>0} \times \{\circ\})} Dist(v, q)$.

We define $\Lambda : R \rightarrow (V')^*$ and σ' inductively as follows: $\sigma'(s'_{in}) = (s_{in}, A)^f$ and $\Lambda(s_{in}) = s'_{in} \cdot (s_{in}, A)^f$ where A consists of all pairs (q, \star) such that $s_{in} \models Form(q)$. Given $v \in R$, such that $last(\Lambda(v)) = (s, A)^f$, we define $\sigma'(\Lambda(v))((s, A)^f, (s, D)^g) = 1$ where D is a minimal set⁴ satisfying the following conditions:

1. $((s, A)^f, (s, D)^g) \in E'$, always giving a priority to pairs of the form (q', x') (see the definition of E') where q' is a terminal accepting state, and avoiding pairs of the form (q', x') where q' is a non-accepting terminal state;
2. if $s \in V_{\square}$ and $\sigma(v)(s, t) > 0$, then $(t, B) \in D$ for some B ;
3. if $Dist(v, A) > 0$, then $Dist(v, A) > Dist(v, \bigcup_{(t, B) \in D} B)$;
4. for every $q \in States$ and $(t, B) \in D$, we have that $v \cdot t \models^v Form(q)$ implies $(q, \star) \in B$.

Furthermore, if $s \in V_{\square}$, then we define $\sigma'(\Lambda(v) \cdot (s, D)^g)((s, D)^g, (t, B)^f) = x$ for $(t, B) \in D$ and $x = \sigma(v)(s, t)$, and we define $\Lambda(v \cdot t) = \Lambda(v) \cdot (s, D)^g \cdot (t, B)^f$ where $(t, B) \in D$ and $\Lambda(v) \cdot (s, D)^g$ is the unique successor of $\Lambda(v)$ in $G'(\sigma')$.

⁴By minimal we mean that whenever D' satisfies the conditions 1–4 and for each $(t, B) \in D$ there is $(t, B') \in D'$ satisfying $B' \subseteq B$, then $D = D'$. (Note that due to the conditions 1 and 2, D' contains exactly one pair of the form (t, B) whenever $\sigma(v)(s, t) > 0$, and no pair of the form (t, B) whenever $\sigma(v)(s, t) = 0$.)

Observe that σ' is indeed \mathfrak{R} -must, because for f -states and s'_{in} it chooses a successor deterministically, while for all other states it assigns non-zero probability to all successors.

Lemma A.3. $\Lambda(R)$ is precisely the set of f -states reachable from s'_{in} in $G'(\sigma')$. For all $v, v' \in \Lambda(R)$ we have $P(v \rightarrow^* v') = P(\Lambda^{-1}(v) \rightarrow^* \Lambda^{-1}(v'))$. Moreover, given $\Lambda(v) = v' \cdot (s, A)^f$ and $(q, x) \in A$, we have that $v \models^v \text{Form}(q)$.

Proof. Let $v \in R$ such that $\text{last}(\Lambda(v)) = (s, A)^f$ where for all pairs $(q, x) \in A$ we have $v \models \text{Form}(q)$. We show that for every $v \cdot t$, such that $v \xrightarrow{x} v \cdot t$ and $\text{last}(\Lambda(v \cdot t)) = (t, A')^f$, we have that $\Lambda(v) \xrightarrow{1} \Lambda(v) \cdot (s, D)^g \xrightarrow{x} \Lambda(v \cdot t)$ and that for all pairs $(q, x) \in A'$ holds $v \cdot t \models \text{Form}(q)$. The rest follows from the definition of $\Lambda(s_{in})$ by simple induction.

We have to prove that $\sigma'(v)$ is defined (i.e., that there is $(s, D)^g$ satisfying the conditions 1–3). We show that for every $(q, x) \in A$ there exists q' and X such that $q \xrightarrow{X} q'$ and $\text{Start}(q \xrightarrow{X} q') \subseteq A \cup L(s)$. Because $v \models \text{Form}(q)$, where $\text{Form}(q)$ has the form $B \sim^p(\Phi_1, \dots, \Phi_n)$, we obtain that there is q' such that $q \xrightarrow{X} q'$ and $v \models \bigwedge_{i \in X} \Phi_i$. However, then $\text{Start}(q \xrightarrow{X} q') \subseteq A \cup L(s)$ because A contains all pairs of the form (r, \star) where $v \models \text{Form}(r)$ (these pairs were put into A in the previous step).

Now, observe also that if q' is non-terminal then it is uniquely determined. This follows from the fact that Φ is a formula of detPECTL^* . Let us assume that q' is non-terminal. It follows from the uniqueness of q' that if $s \in V_{\bigcirc}$ and $q' \in \text{States}^{-1}$, then $v \cdot t \models^v \text{Form}(q')$ for all t such that $(s, t) \in E$. If $s \in V_{\bigcirc}$ and $q' \in \text{States}^{>0}$ is not accepting, then there is t such that $\text{Dist}(v, q) > \text{Dist}(v \cdot t, q')$ and thus $v \cdot t \models^v \text{Form}(q')$. If $s \in V_{\square}$ and $q' \in \text{States}^{>0}$, then there is t such that $\sigma(v)(s, t) > 0$, $v \cdot t \models^v \text{Form}(q')$ and $\text{Dist}(v, q) > \text{Dist}(v \cdot t, q')$. If $s \in V_{\square}$ and $q' \in \text{States}^{-1}$, then for all t such that $\sigma(v)(s, t) > 0$ we have $v \cdot t \models \text{Form}(q')$.

This shows that there is $(s, D)^g$ satisfying the condition 3 and also that all states of $G'(\sigma')$ of the form $\Lambda(v \cdot t) = \Lambda(v) \cdot (s, D)^g \cdot (t, A')^f$ satisfy that for all pairs $(q, x) \in A'$ holds $v \cdot t \models \text{Form}(q)$ (here we used the minimality condition from the definition of σ'). \square

Lemma A.4. All runs of $\text{Run}(\Lambda(s_{in}))$ enter a vertex from P infinitely many times.

Proof. Let $w \in \text{Run}(\Lambda(s_{in}))$. For each $j \geq 0$ we put $w[2j] = (s_j, A_j)^f$. Note that $w[2j]$ is in P whenever $\text{Dist}(\Lambda^{-1}(w[2j]), A_j) = 0$. Now observe that due to the definition of σ' , we have that $\text{Dist}(w[2j+2], A_{j+1}) \geq \text{Dist}(w[2j], A_j)$ implies $\text{Dist}(w[2j]) = 0$. However, then, clearly, the run w enters a vertex from P infinitely many times. \square

Lemma A.5. Almost all runs of $\text{Run}(\Lambda(s_{in}))$ enter a vertex of O infinitely many times.

Proof. Let $v \in R$, and let $last(\Lambda(v)) = (s, A)^f$. Let $(q_0, \circ) \in A$ and let us assume that $Form(q_0) = B_{q_0}^{-1}(\Phi_1, \dots, \Phi_n)$. Let $w \in Run(v)$ be a run for which there is a sequence q_0, q_1, \dots, q_i of states of B such that q_i is accepting, $q_j \xrightarrow{X_j} q_{j+1}$ and $w(j) \models \bigwedge_{\ell \in X_j} \Phi_\ell$ for all $0 \leq j < i$. Let us assume that the sequence q_0, q_1, \dots, q_i is the shortest one with such properties. Now, let w' be the only run of $Run(\Lambda(v))$ such that for all $j \geq 0$ holds $w'(2j) = \Lambda(w(j))$. It follows from the definition of σ' that for every j such that $0 \leq j < i$ we have $w'[2j] = (w[j], A_j)$ where $(q_j, \circ) \in A_j$, and $w'[2i] = (w[i], A_i)$ where $(q_i, \star) \in A_i$. It is easy to observe that the number of pairs of the form $(q, \circ) \in States^{-1} \times \{\circ\}$ in A_i is strictly less than in $A_0 = A$.

We have shown that all $v \in \Lambda(R)$, such that $last(v) = (s, A)^f$, satisfy the following: With probability 1, a run initiated in v reaches a state v' such that $last(v') = (s', A')^f$ where $|A' \cap (States^{-1} \times \{\circ\})| < |A \cap (States^{-1} \times \{\circ\})|$. Now, using basic results of theory of Markov chains we obtain that with probability 1, a run initiated in v reaches a state v' such that $last(v) \in O$. However, this implies that almost all runs of $Run(s'_{in})$ reach vertices of O infinitely many times. \square

A.3 Proof of Lemma 4.5

We define the Büchi automata B_{fin} and $B_{q,i}$, for all $q \in Q$ and $1 \leq i \leq k$, as follows. Let $R = (Q, \Sigma, \gamma, q_0, Acc)$, where $Acc = \{(C_1, D_1), \dots, (C_k, D_k)\}$, be a Rabin automaton such that $L(R) = L(B)$. We make use of the fact that the Rabin automaton R is easily complemented by dualizing the accepting condition (i.e., considering R to be a Streett automaton). Let $B_{fin} = (Q \cup \{q_a\}, \Sigma \cup T, \delta_{fin}, q_0, \{q_a\})$, where $T = \{A_q \mid q \in Q\}$ where $A_q = \{n + index(q, i) \mid 1 \leq i \leq k\}$, and transitions of B_{fin} are defined as follows:

- $q \xrightarrow{A} q'$ for all $A \in \Sigma$ and $q, q' \in Q$ such that $\gamma(q, a) = q'$;
- $q \xrightarrow{A_q} q_a$ for all $q \in Q$ and all $1 \leq j \leq k$;
- $q_a \xrightarrow{\emptyset} q_a$;
- nothing else is a transition.

Let us define $B_{q,i} = (Q \times \{1, 2\}, \Sigma, \delta_{q,i}, (q, 1), F)$, where $F = (Q \times \{1\}) \cup (D_i \times \{2\})$, and transitions are defined as follows:

- $(p, 1) \xrightarrow{A} (p', 1)$ if and only if $p \in Q \setminus C_i$ and $\gamma(p, A) = p'$;
- $(p, 1) \xrightarrow{A} (p', 2)$ if and only if $p \in C_i$ and $\gamma(p, A) = p'$;
- $(p, 2) \xrightarrow{A} (p', 2)$ if and only if $\gamma(p, A) = p'$;

- nothing else is a transition.

We prove that $B^{\sim\rho}(\Phi_1, \dots, \Phi_n) \equiv_{fin} B_{fin}^{\widehat{\rho}}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$. The proof is very similar to the proof of Lemma 4.4 and uses similar notation. Let us prove the equivalence. Let us fix a finite Markov chain M with a set of states S , a state s_0 of M , and a valuation v . First, let $s_0 \models B_{fin}^{\widehat{\rho}}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$. Immediately from definitions follows that, with probability $\widehat{\rho}$, there is a path $s_0, s_1, \dots, s_i, s_{i+1}$ in M satisfying the following conditions:

- the automaton B_{fin} (and hence also the automaton R) moves from its initial state to a state r after reading the word A_{s_0}, \dots, A_{s_i} ;
- $s_{i+1} \models B_{r,j}^{\equiv 1}(\Phi_1, \dots, \Phi_n)$ for all $1 \leq j \leq k$ (and hence for almost all runs $w \in Run(s_{i+1})$ and all $1 \leq j \leq k$ holds that either no state of C_j occurs in the computation of R on w initiated in r , or some state of D_j occurs infinitely many times in the computation of R on w initiated in r).


However, this immediately implies that, with probability $\widehat{\rho}$, the automaton R *does not* accept a run of $Run(s_0)$. It follows that, with probability $\sim\rho$, the automaton R accepts a run of $Run(s_0)$.

For the opposite direction, let $M \times R$ be the product Markov chain defined in the proof of Lemma 4.4. We say that a BSCC C of $M \times R$ is *rejecting* iff for all $1 \leq j \leq k$ we have that either no state of C_j occurs in C , or some state of D_j occurs in C . Basic results of the theory of Markov chains imply that the probability measure of all runs of $Run(s)$ accepted by R (hence not accepted by B) is equal to one minus the probability of reaching a rejecting BSCC of $M \times R$. Thus, if $s_0 \models B^{\sim\rho}(\Phi_1, \dots, \Phi_n)$, then, with probability $\widehat{\rho}$, there is a path $(s_0, q_0), \dots, (s_i, q_i)$ in $M \times R$ such that (s_i, q_i) lies in a rejecting BSCC. Now, it follows immediately from the definition of $B_{q_i,j}$ that $s_i \models B_{q_i,j}^{\equiv 1}(\Phi_1, \dots, \Phi_n)$, and also that B_{fin} moves from q_0 to q_a after reading the word $A_{s_0}, \dots, A_{s_i}, A_{q_i}$. This implies that $s_0 \models B_{fin}^{\widehat{\rho}}(\Phi_1, \dots, \Phi_n, \Psi_1, \dots, \Psi_\ell)$. \square

A.4 Proof of Lemma 5.2

First we show that there is a (v, Φ) -winning HD strategy σ in *start*. The proof is very similar to the analogous proof for the $L(G^{>0}, F^{>0})$ fragment given in [3]. We define $\sigma(ws)$ to be the Dirac distribution which assigns 1 to the transitions leading to s_1 or s_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 $right_1$ or $right_2$ in w and $\#_{left}(w)$ denotes the number of occurrences of a state satisfying the proposition $left$ in w . We claim that the state *start* in the play $G(\sigma)$ satisfies the formula $G^{\equiv 2/3}(\neg stop \wedge (\neg left \vee G^{>0} \neg right_2))$ and hence also the

formula Φ . To see this, realize that the play $G(\sigma)$ corresponds to the following infinite Markov chain:



A standard calculation reveals that the probability of hitting the *stop* state from *start* is equal to $1/3$. Hence, the probability of all runs initiated in *start* which do *not* hit the *stop* state is $2/3$. All states in all these runs can reach the *stop* on the path that does not go through *right₂* state. And because *right₂* is not reachable from *stop* state, we have that *start* satisfies the formula $G^{=2/3}(\neg stop \wedge (\neg left \vee G^{>0} \neg right_2))$.

Now suppose that there is a (v, Φ) -winning FR strategy σ . We claim that the state *start* of $\bar{G}(\sigma)$ satisfies the formula $G^{=0}(\neg stop \wedge (\neg left \vee G^{>0} \neg right_2))$. Let us consider the finite Markov chain $\bar{G}(\sigma)$ (see Remark 2.2 in Section 2.3). Observe, that each BSCC of $\bar{G}(\sigma)$ contains either a *stop* state (i.e., a state of the form $[v \cdot stop]$, remember that states of $\bar{G}(\sigma)$ are equivalence classes of strings of vertices), or contains both a *right₂* state and a *left* state. Now because almost all runs of an arbitrary finite Markov chain enter each state of some BSCC infinitely many times, we obtain that no *left* state contained in a BSCC of $\bar{G}(\sigma)$ satisfies $\neg left \vee (G^{>0} \neg right_2)$, and hence that the state *start* satisfies $G^{=0}(\neg stop \wedge (\neg left \vee G^{>0} \neg right_2))$. \square