

Finite-Memory Strategies for Almost-Sure Energy-MeanPayoff Objectives in MDPs

Mohan Dantam

School of Informatics, University of Edinburgh, UK

Richard Mayr

School of Informatics, University of Edinburgh, UK

Abstract

We consider finite-state Markov decision processes with the combined Energy-MeanPayoff objective. The controller tries to avoid running out of energy while simultaneously attaining a strictly positive mean payoff in a second dimension.

We show that *finite memory* suffices for almost surely winning strategies for the Energy-MeanPayoff objective. This is in contrast to the closely related Energy-Parity objective, where almost surely winning strategies require infinite memory in general.

We show that exponential memory is sufficient (even for deterministic strategies) and necessary (even for randomized strategies) for almost surely winning Energy-MeanPayoff. The upper bound holds even if the strictly positive mean payoff part of the objective is generalized to multidimensional strictly positive mean payoff.

Finally, it is decidable in pseudo-polynomial time whether an almost surely winning strategy exists.

2012 ACM Subject Classification Theory of computation → Random walks and Markov chains; Mathematics of computing → Probability and statistics

Keywords and phrases Markov decision processes, energy, mean payoff, parity, strategy complexity

Digital Object Identifier 10.4230/LIPIcs.ICALP.2024.133

Category Track B: Automata, Logic, Semantics, and Theory of Programming

Related Version *Full Version*: <https://arxiv.org/abs/2404.14522> [18]

1 Introduction

Background. Markov decision processes (MDPs) are a standard model for dynamic systems that exhibit both stochastic and controlled behaviour [28]. MDPs play a prominent role in many domains, e.g., artificial intelligence and machine learning [32, 30], control theory [5, 1], operations research and finance [31, 23, 9, 29], and formal verification [2, 31, 20, 14, 3, 11].

An MDP is a directed graph where states are either controlled or random. If the current state is controlled then the controller can choose a distribution over all possible successor states. If the current state is random then the next state is chosen according to a fixed probability distribution. One assigns numeric rewards to transitions (and this can be generalized to multidimensional rewards). Moreover, priorities (aka colours), encoded by bounded non-negative numbers, are assigned to states. By fixing a strategy for the controller and an initial state, one obtains a probability space of runs of the MDP. The goal of the controller is to optimize the expected value of some objective function on the runs.

The *strategy complexity* of a given objective is the amount of memory (and randomization) needed for an optimal (resp. ε -optimal) strategy. Common cases include memoryless strategies, finite-memory strategies, Markov strategies (using a discrete clock, aka step counter), and general infinite-memory strategies.



© Mohan Dantam and Richard Mayr;

licensed under Creative Commons License CC-BY 4.0

51st International Colloquium on Automata, Languages, and Programming (ICALP 2024).

Editors: Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson;

Article No. 133; pp. 133:1–133:17



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Related work. The Parity, MeanPayoff and Energy objectives have been extensively studied in the formal verification community. A run satisfies the (min-even) *Parity objective* iff the minimal priority that appears infinitely often in the run is even. It subsumes all ω -regular objectives, and in particular safety, liveness, fairness, etc. The *MeanPayoff objective* requires that the limit average reward per transition along a run is positive (resp. non-negative in some settings). MeanPayoff objectives go back to a 1957 paper by Gillette [21] and have been widely studied, due to their relevance for efficient control. The *Energy objective* [10] requires that the accumulated reward at any time in a run stays above some finite threshold (typically 0). The intuition is that a controlled system has some finite initial energy level that must never become depleted.

Combinations of these objectives have also been studied, where the runs need to satisfy several of the above conditions simultaneously.

The existence of almost surely winning strategies for *MeanPayoff-Parity* in MDPs is decidable in polynomial time [12]. These strategies require only finite memory for MeanPayoff > 0 [22], but infinite memory for MeanPayoff ≥ 0 [13].

The existence of almost surely winning strategies for *Energy-Parity* in MDPs is decidable in $\text{NP} \cap \text{coNP}$ and in pseudo-polynomial time [25]. (The $\text{NP} \cap \text{coNP}$ upper bound holds even for turn-based stochastic games [26].) Almost surely winning strategies in MDPs require only finite memory in the special case of Energy-Büchi [12], but infinite memory for Energy-co-Büchi and thus for Energy-Parity [25]. However, ε -optimal strategies for Energy-Parity require only finite (at most doubly exponential) memory, and the value can be effectively approximated in doubly exponential time (even for turn-based stochastic games) [17].

The *Energy-MeanPayoff* objective is similar to Energy-Parity, but replaces the Parity part by a MeanPayoff objective for a second reward dimension. I.e., one considers an MDP with 2-dimensional transition rewards, where the Energy condition applies to the first dimension and the MeanPayoff condition applies to the second dimension. (It can be generalized to higher dimensions d , where the MeanPayoff condition applies to all dimensions $2, 3, \dots, d$.) This might look like a direct generalization of the Energy-Parity objective, since Parity games are reducible to MeanPayoff games [27, 24]. However, this reduction does not work in the context of these combined objectives when one considers stochastic systems like MDPs; see below. Non-stochastic Energy-MeanPayoff games have been studied in [8].

A slightly different objective has been studied in [16] who consider MDPs with d -dimensional rewards, where $d = d_1 + d_2$. The objective requires a strictly positive MeanPayoff *surely* in the first d_1 dimensions, and *almost surely* in the remaining d_2 dimensions. This objective is strictly stronger than Energy-MeanPayoff. E.g., a MeanPayoff of zero in the first dimension may or may not satisfy the Energy objective, but it never satisfies the objective in [16].

The objective studied in [6] aims to maximize the expected MeanPayoff (rather than the probability of it being strictly positive) while satisfying the energy constraint. However, unlike in our work, the reward function has a single dimension (i.e., both criteria apply to the same value) and ε -optimal strategies can require infinite memory.

Our contribution. We consider the Energy-MeanPayoff objective in MDPs with d -dimensional rewards. The first dimension needs to satisfy the Energy condition (never drop below 0), while each other dimension needs to have a *strictly* positive MeanPayoff. We show that

almost surely winning strategies for Energy-MeanPayoff require only *finite* memory.¹ This is in contrast to the Energy-Parity objective where almost surely winning strategies require infinite memory in general [25, Page 4] (even for the simpler Energy-co-Büchi objectives). This also shows that Energy-Parity is not reducible to Energy-MeanPayoff in MDPs, unlike the reduction from Parity to MeanPayoff in [27, 24].

We show that almost surely winning strategies for Energy-MeanPayoff, if they exist, can be chosen as deterministic strategies with an exponential number of memory modes. The crucial property is that it suffices to remember the stored energy only up to some exponential upper bound. A small counterexample shows the corresponding exponential lower bound. Even for randomized strategies, an exponential number of memory modes is required, and this holds even for the case of small transition rewards in $\{-1, 0, +1\}$.

Although almost surely winning strategies are “exponentially large” in this sense, their existence is still decidable in pseudo-polynomial time; cf. Section 7.

2 Preliminaries

A *probability distribution* over a countable set S is a function $f: S \rightarrow [0, 1]$ with $\sum_{s \in S} f(s) = 1$. $\text{supp}(f) \stackrel{\text{def}}{=} \{s \mid f(s) > 0\}$ denotes the support of f and $\mathcal{D}(S)$ is the set of all probability distributions over S . Given an alphabet Σ , let Σ^ω and Σ^* (Σ^+) denote the set of infinite and finite (non-empty) sequences over Σ , respectively. Elements of Σ^ω or Σ^* are called words.

MDPs and Markov chains. A *Markov Decision Process* (MDP) is a controlled stochastic directed graph $\mathcal{M} \stackrel{\text{def}}{=} (S, S_\square, S_\circ, E, P, \mathbf{r})$ where the set of vertices S (also called states) is partitioned into the states S_\square of the player \square (*Maximizer*), and chance vertices (aka random states) S_\circ . Let $E \subseteq S \times S$ be the transition relation. We write $s \longrightarrow s'$ if $(s, s') \in E$ and assume that $\text{Succ}(s) \stackrel{\text{def}}{=} \{s' \mid sEs'\} \neq \emptyset$ for every state s . The *probability function* P assigns each random state $s \in S_\circ$ a distribution over its successor states, i.e., $P(s) \in \mathcal{D}(\text{Succ}(s))$. We extend the domain of P to S^*S_\circ by $P(\rho s) \stackrel{\text{def}}{=} P(s)$ for all $\rho s \in S^+S_\circ$. A *Markov chain* is an MDP with only random states, i.e., $S_\square = \emptyset$. In this paper we consider finite-state MDPs, i.e., the set of states S is finite.

Strategies. A *run* is an infinite sequence $s_0s_1 \dots \in S^\omega$ such that $s_i \longrightarrow s_{i+1}$ for all $i \geq 0$. A *path* is a finite prefix of a run. Let $\text{Runs}(\mathcal{M}) \stackrel{\text{def}}{=} \{\rho = (q_i)_{i \in \mathbb{N}} \mid q_i \longrightarrow q_{i+1}\}$ denote the set of all possible runs. A strategy of the player \square is a function $\sigma: S^*S_\square \rightarrow \mathcal{D}(S)$ that assigns to every path $ws \in S^*S_\square$ a probability distribution over the successors of s . If these distributions are always Dirac then the strategy is called *deterministic* (aka pure), otherwise it is called *randomized* (aka mixed). The set of all strategies of player \square in \mathcal{M} is denoted by $\Sigma^\mathcal{M}$. A run/path $s_0s_1 \dots$ is compatible with a strategy σ if $s_{i+1} \in \text{supp}(\sigma(s_0 \dots s_i))$ whenever $s_i \in S_\square$. Finite-memory strategies are a subclass of strategies using a finite set M of memory modes. A function $\text{nxt}: M \times S_\square \mapsto \mathcal{D}(S)$ chooses a (distribution over) successor states based on the current memory mode and state and $\text{upd}: M \times E \mapsto \mathcal{D}(M)$ updates the memory mode upon observing a transition. Let $\sigma[m]$ denote the finite-memory strategy σ starting in

¹ Our results do *not* carry over to Energy-MeanPayoff objectives with *non-strict* inequalities where one just requires a MeanPayoff ≥ 0 almost surely. This needs infinite memory even for the case of $d = 2$, i.e., one energy-dimension and one MeanPayoff-dimension. It suffices to modify the counterexample for Energy-co-Büchi from [25, Page 4] such that a visit to a state with unfavourable colour incurs a reward of -1 in the MeanPayoff-dimension.

memory mode m . The set of all finite-memory strategies in \mathcal{M} is denoted by $\Sigma_f^{\mathcal{M}}$. Strategies with memory $|M| = 1$ are called *memoryless*. Memoryless deterministic (resp. randomized) strategies are called MD (resp. MR). By fixing some finite-memory strategy σ from some initial state in a finite-state MDP \mathcal{M} , we obtain a finite-state Markov chain, denoted by \mathcal{M}^σ .

Measure. An MDP \mathcal{M} with initial state s_0 and strategy σ yields a probability space $(s_0S^\omega, \mathcal{F}_{s_0}, \mathcal{P}_{\sigma, s_0}^{\mathcal{M}})$ where \mathcal{F}_{s_0} is the σ -algebra generated by the cylinder sets $s_0s_1 \dots s_nS^\omega$ for $n \geq 0$. The probability measure $\mathcal{P}_{\sigma, s_0}^{\mathcal{M}}$ is first defined on the cylinder sets. For $\rho = s_0 \dots s_n$, let $\mathcal{P}_{\sigma, s_0}^{\mathcal{M}}(\rho) \stackrel{\text{def}}{=} 0$ if ρ is not compatible with σ and otherwise $\mathcal{P}_{\sigma, s_0}^{\mathcal{M}}(\rho S^\omega) \stackrel{\text{def}}{=} \prod_{i=0}^{n-1} \tau(s_0 \dots s_i)(s_{i+1})$ where τ is σ or P depending on whether $s_i \in S_\square$ or S_\circ , respectively. If \mathcal{M} is a Markov chain then there is only a single strategy, and we simply write $\mathcal{P}_{s_0}^{\mathcal{M}}$. By Carathéodory's extension theorem [4], this defines a unique probability measure on the σ -algebra. Given some reward function $v : s_0S^\omega \rightarrow \mathbb{R}$, we write $\mathcal{E}(\cdot)$ for the expectation w.r.t. \mathcal{P} and v .

Objectives. General objectives are defined by real-valued measurable functions. However, we mostly consider indicator functions of measurable sets. Hence, our objectives can be described by measurable subsets $\mathbf{0} \subseteq S^\omega$ of runs starting at a given initial state. By $\mathcal{P}_{\sigma, s}^{\mathcal{M}}(\mathbf{0})$ we denote the payoff under σ , i.e., the probability that runs from s belong to $\mathbf{0}$. The value of a state is defined as $\text{val}_0^{\mathcal{M}}(s) \stackrel{\text{def}}{=} \sup_{\sigma \in \Sigma^{\mathcal{M}}} \mathcal{P}_{\sigma, s}^{\mathcal{M}}(\mathbf{0})$. For $\varepsilon > 0$ and state s , a strategy $\sigma \in \Sigma^{\mathcal{M}}$ is ε -optimal iff $\mathcal{P}_{\sigma, s}^{\mathcal{M}}(\mathbf{0}) \geq \text{val}_0^{\mathcal{M}}(s) - \varepsilon$. A 0-optimal strategy is called *optimal*. An MD/MR strategy is called *uniformly ε -optimal* (resp. *uniformly optimal*) if it is so from every start state. An optimal strategy from s is called *almost surely winning* if $\text{val}_0^{\mathcal{M}}(s) = 1$. By $\text{AS}(\mathbf{0})$ (resp. $\text{AS}_f(\mathbf{0})$) we denote the set of states that have an almost surely winning strategy (resp. an almost surely winning finite-memory strategy) for objective $\mathbf{0}$. For ease of presentation, we drop subscripts and superscripts wherever possible if they are clear from the context.

We use the syntax and semantics of the LTL operators [15] **F** (eventually), **G** (always) and **X** (next) to specify some conditions on runs. A reachability objective is defined by a set of target states $T \subseteq S$. A run $\rho = s_0s_1 \dots$ belongs to FT iff $\exists i \in \mathbb{N} s_i \in T$. Similarly, ρ belongs to $\text{F}^{\leq n}T$ (resp. $\text{F}^{\geq n}T$) iff $\exists i \leq n$ (resp. $i \geq n$) such that $s_i \in T$. Dually, the safety objective GT consists of all runs which never leave T . We have $\text{GT} = \neg\text{F}\neg T$.

Energy/Reward/Counter-based objectives. Let $r : E \rightarrow \{-R, \dots, 0, \dots, R\}$ be a bounded function that assigns rewards to transitions. Depending on context, the sum of these rewards in a path can be viewed as energy, cost/profit or a counter. If $s \xrightarrow{c} s'$ and $r((s, s')) = c$, we write $s \xrightarrow{c} s'$. Let $\rho = s_0 \xrightarrow{c_0} s_1 \xrightarrow{c_1} \dots$ be a run. We say that ρ satisfies

1. the *k-energy* objective $\text{EN}(k)$ iff $\left(k + \sum_{i=0}^{n-1} c_i\right) \geq 0$ for all $n \geq 0$.
2. the *l-storage condition* $\text{Infix}(l)$ if $l + \sum_{i=m}^{n-1} c_i \geq 0$ holds for every infix $s_m \xrightarrow{c_m} s_{m+1} \dots s_n$ of the run. Let $\text{ST}(k, l)$ denote the set of runs that satisfy both the *k-energy* and the *l-storage condition*. Let $\text{ST}(k) \stackrel{\text{def}}{=} \bigcup_l \text{ST}(k, l)$. Clearly, $\text{ST}(k) \subseteq \text{EN}(k)$.
3. *Mean payoff* $\text{MP}(\triangleright c)$ for some constant $c \in \mathbb{R}$ iff $\left(\liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} c_i\right) \triangleright c$ for $\triangleright \in \{<, \leq, =, \geq, >\}$.

A different way to consider the energy objective is to encode the energy level (the sum of the transition weights so far) into the state space and then consider the obtained infinite-state game with a safety objective.

An objective $\mathbf{0}$ is called *shift-invariant* iff for all finite paths ρ and plays $\rho' \in S^\omega$, we have $\rho\rho' \in \mathbf{0} \iff \rho' \in \mathbf{0}$. Mean payoff objectives are shift-invariant, but energy and storage/infix objectives are not.

Multidimensional reward-based objectives. Let $\mathbb{N}, \mathbb{Q}, \mathbb{R}$ denote the set of positive integers, rationals and reals respectively. For a d -dimensional real vector $\boldsymbol{\mu}$, let μ_i denote the i^{th} component of $\boldsymbol{\mu}$ for $1 \leq i \leq d$. Given two vectors $\boldsymbol{\mu}, \boldsymbol{\nu} \in \mathbb{R}^d$, $\sim \in \{<, \leq, >, \geq, =\}$ we say $\boldsymbol{\mu} \sim \boldsymbol{\nu}$ if $\mu_i \sim \nu_i$ for every i . In particular, $\boldsymbol{\mu} > \mathbf{0}$ means that *every* component of $\boldsymbol{\mu}$ is strictly greater than 0. For a multidimensional reward function $\mathbf{r} : E \rightarrow [-R, R]^d$, we can consider any boolean combination of reward based objectives using any components of \mathbf{r} . For instance, $\mathbf{0}_1 = \text{EN}_1(k) \cap \text{MP}_2(> 0)$ denotes the objective that contains all runs that satisfy $\text{EN}(k)$ in the 1^{st} dimension and $\text{MP}(> 0)$ in the 2^{nd} one. We denote conjunctions of the same objective across different dimensions in vectorized form, with the dimension information in the subscript. Therefore, $\text{EN}_{[a,b]}(\mathbf{k}) \cap \text{MP}_{[c,d]}(> \mathbf{x})$ denotes the runs where the $\text{EN}_i(k_i)$ objective is satisfied for each $i \in [a, b]$ and the $\text{MP}_j(> x_j)$ objective is satisfied for each $j \in [c, d]$. Given an infinite run $\rho = s_0 \xrightarrow{c_0} s_1 \xrightarrow{c_1} \dots$, let $X_n(\rho) \stackrel{\text{def}}{=} s_n$ denote the n -th state. Let \mathbf{Y}_n be the sum of the rewards in the first n steps, i.e., $\mathbf{Y}_n(\rho) \stackrel{\text{def}}{=} \sum_{i=0}^{n-1} \mathbf{c}_i$. These become random variables once an initial distribution and a strategy are fixed.

Size of an instance. Given an MDP $\mathcal{M} = (S, S_\square, S_\circ, E, P, \mathbf{r})$ with reward function $\mathbf{r} : E \rightarrow [-R, R]^d$, its size $|\mathcal{M}|$ is the number of bits used to describe it. Similarly for $|P|$. Transition probabilities and rewards can thus be stored in binary. We call a size pseudo-polynomial in $|\mathcal{M}|$ if it is polynomial for the case where R is “small”, i.e., if R is given in unary.

3 The Main Result

► **Theorem 1.** *Let $\mathcal{M} = (S, S_\square, S_\circ, E, P, \mathbf{r})$ be an MDP with d -dimensional rewards on the edges $\mathbf{r} : E \rightarrow [-R, R]^d$. For the multidimensional Energy-MeanPayoff objective $\text{EN}_1(k) \cap \text{MP}_{[2,d]}(> \mathbf{0})$ the following properties hold.*

1. *The existence of an almost-surely winning strategy implies the existence of an almost-surely winning finite-memory strategy.*
2. *Moreover, a deterministic strategy with an exponential number of memory modes is sufficient.*
3. *An exponential (in $|P|$) number of memory modes is necessary in general, even for randomized strategies, even for $|S| = 5$, $d = 2$ and $R = 1$.*

In the following three sections we prove items 1.,2.,3. of Theorem 1, respectively.

Here we sketch the main idea for the upper bound. Except in a special corner case where the energy fluctuates only in a bounded region, almost-surely winning strategies for Energy-MeanPayoff can be chosen among some particular strategies that alternate between two modes, playing two different memoryless strategies. This alternation keeps the balance between the Energy-part and the MeanPayoff-part of the objective. This is similar to almost-surely winning strategies for the Energy-Parity objective in [25]. In one mode, one plays a randomized memoryless strategy that almost surely yields a positive mean payoff in all dimensions (in case of Energy-Parity, instead of mean payoff it satisfies Parity almost surely). This is called the *Gain* phase. Whenever the energy level (the cumulative reward in dimension 1) gets dangerously close to zero, one switches to the other mode and plays a different memoryless strategy that focuses exclusively on getting the energy level up again, while temporarily neglecting the other part of the objective (Parity or Mean payoff, respectively). This is called a *Bailout*. Once the energy level is sufficiently high, one switches back to the Gain phase again. The crucial property is that, except in a null set, only finitely many Bailouts are required, and thus the temporary neglect of the second part of the objective does not matter in the long run. Such a strategy uses infinite memory,

because it needs to remember the unbounded energy level. For Energy-Parity (and even Energy-co-Büchi) this cannot be avoided and finite-memory strategies do not work [25]. However, for Energy-MeanPayoff one can relax the requirements somewhat. Suppose that one records the stored energy only up to a certain bound b , i.e., one forgets about potential excess energy above b . In that case, one might have to do infinitely many Bailouts with high probability, most of which are unnecessary (but one does not know which ones). However, for a sufficiently large bound b , these superfluous Bailouts occur so infrequently that they do not compromise the MeanPayoff-part of the objective. The critical part of the proof is to show this property and an upper bound on b . Once this is established, one obtains a finite-memory strategy, because it suffices to record the energy level only in the range $[0, b]$ (plus one extra bit of memory to record the current phase, Gain or Bailout).

Note that the argument above is different from the one that justifies finite-memory ε -optimal strategies for *Energy-Parity* in [17]. These also record the energy only in a bounded region, but stop doing Bailouts after the upper bound has been visited. I.e., they do too few Bailouts, and thus incur an ε -chance of losing. In contrast, our almost-surely winning strategies for Energy-MeanPayoff rather do too many Bailouts, but sufficiently infrequently such that they don't compromise the objective.

4 Proof of Item 1

W.l.o.g, we assume that every state in \mathcal{M} has an almost surely winning strategy for Energy-MeanPayoff for some initial energy level. (Otherwise, consider a suitably restricted sub-MDP.) For conciseness, we denote the objective by $\mathbf{0}(k) \stackrel{\text{def}}{=} \text{EN}_1(k) \cap \text{MP}_{[2,d]}(> \mathbf{0})$. Let

$$\text{Win}(s) \stackrel{\text{def}}{=} \{k \mid s \in \text{AS}(\mathbf{0}(k))\}, \quad i_s \stackrel{\text{def}}{=} \min(\text{Win}(s))$$

denote the possible initial energy levels and the minimum initial energy level such that one can win almost surely from state s . In particular, i_s is well defined by our assumption on \mathcal{M} .

Towards a contradiction, assume that not all configurations are winnable with a finite-memory strategy. I.e., let $\text{Win}_f(s) \stackrel{\text{def}}{=} \{k \mid s \in \text{AS}_f(\mathbf{0}(k))\}$ denote the energy levels from which one can win almost surely with a *finite-memory* strategy from s , and assume that there is a state s^\dagger such that $i_{s^\dagger} \notin \text{Win}_f(s^\dagger)$. We then construct a finite-memory winning strategy from s^\dagger for $\mathbf{0}(i_{s^\dagger})$, leading to a contradiction. Similar to i_s , let f_s denote the minimal k such that $k \in \text{Win}_f(s)$ and ∞ if there is no such k .

► **Definition 2.** *We construct a new MDP \mathcal{M}^* which abstracts away all the Win_f configurations. At every state s , the player gets the option to enter a winning sink state if the energy level is sufficiently large to win with finite memory, i.e., if the current energy level is at least f_s . The states of the MDP \mathcal{M}^* will have two copies of each state s of \mathcal{M} , namely s and s' . Moreover, we add a new state s_{win} . All states s' are controlled by \square and every step $s_1 \rightarrow s$ in the original MDP \mathcal{M} is now mapped to a step $s_1 \rightarrow s'$ with the same reward (and the same probability if s_1 was a random state). In s' , the player has two choices: he can either go to s with reward $\mathbf{0}$ or go to s_{win} with reward $(-f_s, \mathbf{0})$. The latter choice is only available if $f_s < \infty$. s_{win} is a winning sink where $s_{\text{win}} \rightarrow s_{\text{win}}$ with reward $\mathbf{1}$, i.e., reward $+1$ in all dimensions.*

The following lemma shows that the existence of almost surely winning (finite-memory) strategies coincides in \mathcal{M}^* and \mathcal{M} .

► **Lemma 3.** *Let $s \in S$ and $k \in \mathbb{N}$, and consider the objective $\mathbf{0}(k)$. There exists an almost surely winning strategy σ^* from s in \mathcal{M}^* if and only if there exists an almost surely winning strategy σ from s in \mathcal{M} . Moreover, if σ^* is finite-memory then σ can be chosen as finite-memory, and vice-versa.*

Proof. Towards the “only if” direction, let σ^* be a strategy from s in \mathcal{M}^* that is almost surely winning for $\mathcal{O}(k)$. We define a strategy σ from s in \mathcal{M} that plays as follows. First σ imitates the moves of σ^* until (if ever) σ^* chooses a move $s'_1 \rightarrow s_{\text{win}}$ with non-zero probability at some state s'_1 . This is possible, since any finite path in \mathcal{M}^* that does not contain s_{win} can be bijectively mapped to a path in \mathcal{M} . The only difference is that paths in \mathcal{M}^* contain extra steps via primed states, which are skipped in the paths in \mathcal{M} . Moreover, the transition probabilities at random states coincide in \mathcal{M}^* and \mathcal{M} . If σ^* chooses a move $s'_1 \rightarrow s_{\text{win}}$ with non-zero probability at some state s'_1 then the current energy level must be $\geq f_{s'_1}$, because σ^* satisfies the energy objective almost surely (and thus even surely). Thus, in \mathcal{M} , there exists an almost surely winning finite-memory strategy $\hat{\sigma}$ for $\mathcal{O}(f_{s'_1})$ from s_1 . In this situation σ continues by playing $\hat{\sigma}$ from s_1 . Therefore, σ satisfies the energy objective surely. Moreover, by shift invariance and the properties of $\hat{\sigma}$, it also satisfies the Mean payoff objective almost surely. Thus, σ satisfies $\mathcal{O}(k)$ almost surely. Finally, if σ^* is finite-memory then so is σ , because $\hat{\sigma}$ is also finite-memory.

Towards the “if direction, let σ be a strategy from s in \mathcal{M} that is almost surely winning for $\mathcal{O}(k)$. We define a strategy σ^* from s in \mathcal{M}^* that imitates the moves of σ . Moreover, at primed states q' it always goes to q (and never to s_{win}). Since the probabilities at random states coincide in \mathcal{M}^* and \mathcal{M} , also the probabilities of the induced paths coincide. The only difference is that the runs in \mathcal{M}^* contain extra steps via primed states and these extra steps carry reward zero. Thus, the mean payoff of a run in \mathcal{M}^* is $1/2$ the mean payoff of the corresponding run in \mathcal{M} . However, this does not affect the property that the mean payoff is > 0 almost surely in either MDP. Thus, σ^* satisfies $\mathcal{O}(k)$ almost surely. Finally, if σ is finite-memory then so is σ^* . ◀

The next lemma shows that, in \mathcal{M}^* , it is impossible to satisfy Energy-MeanPayoff from s with arbitrarily high probability, unless one also allows arbitrarily large fluctuations in the energy level, or $f_s = i_s$. (Recall that f_s, i_s are defined relative to \mathcal{M} .)

► **Lemma 4.** *For every state s with $f_s > i_s$ and every $\ell \in \mathbb{N}$, there exists a $\delta_\ell > 0$ such that $\text{val}_{\mathcal{O}(i_s) \cap \text{Inf}_{i_s}(\ell)}^{\mathcal{M}^*}(s) \leq 1 - \delta_\ell$.*

Proof. Towards a contradiction, assume that $\text{val}_{\mathcal{O}(i_s) \cap \text{Inf}_{i_s}(\ell)}^{\mathcal{M}^*}(s) = 1$ for some ℓ .

$\mathcal{O}(i_s) \cap \text{Inf}_{i_s}(\ell) = \text{EN}_1(i_s) \cap \text{MP}_{[2,d]}(> \mathbf{0}) \cap \text{Inf}_{i_s}(\ell) = \text{ST}_1(i_s, \ell) \cap \text{MP}_{[2,d]}(> \mathbf{0})$. Therefore, we have $\text{val}_s^{\mathcal{M}^*}(\text{ST}_1(i_s, \ell) \cap \text{MP}_{[2,d]}(> \mathbf{0})) = 1$. Below we prove that this objective has a finite-memory almost-surely winning strategy σ in \mathcal{M}^* . Consider a modified MDP \mathcal{M}_1^* that encodes the energy level up to $i_s + \ell$ in the states. A step exceeding the upper energy bound $i_s + \ell$ results in a truncation to $i_s + \ell$, while a step leading to a negative energy leads to a losing sink. There exists a memoryless randomized (MR) strategy σ_1 in \mathcal{M}_1^* from state (s, i_s) that wins $\text{MP}_{[2,d]}(> \mathbf{0})$ almost surely, by Lemma 6. We can then carry σ_1 back to \mathcal{M}^* as a finite-memory strategy σ with $i_s + \ell + 1$ memory modes such that $\mathcal{P}_{\sigma, s}^{\mathcal{M}^*}(\text{ST}_1(i_s, \ell) \cap \text{MP}_{[2,d]}(> \mathbf{0})) = 1$. By set inclusion, $\mathcal{P}_{\sigma, s}^{\mathcal{M}^*}(\mathcal{O}(i_s)) = 1$. By Lemma 3, there also exists a finite-memory strategy from s in \mathcal{M} that is almost surely winning for $\mathcal{O}(i_s)$. This implies $f_s = i_s$, a contradiction to our assumption $f_s > i_s$. Hence, we obtain $\delta_\ell \stackrel{\text{def}}{=} 1 - \text{val}_{\mathcal{O}(i_s) \cap \text{Inf}_{i_s}(\ell)}^{\mathcal{M}^*}(s) > 0$. ◀

The following three lemmas show that almost surely winning strategies for Energy-MeanPayoff can be found by combining two different memoryless strategies for the simpler **Bailout** and **Gain** objectives.

First, we define the objective $\text{Bailout}(k) \stackrel{\text{def}}{=} \text{EN}_1(k) \cap \text{MP}_1(> 0)$. Let i_s^{Bailout} denote the minimal energy value k with which one can almost surely satisfy $\text{Bailout}(k)$ when starting from state s (or ∞ if it does not exist).

► **Lemma 5** ([6, Lemma 3]). *Let \mathcal{M} be an MDP. If $s \in \text{AS}(\text{Bailout}(k))$ for some $k \in \mathbb{N}$ then $i_s^{\text{Bailout}} \leq 3 \cdot |\mathcal{M}| \cdot R$. Moreover, there exists a uniform MD strategy $\sigma_{\text{Bailout}}^*$ which is almost surely winning $\text{Bailout}(k)$ from every state $s \in \text{AS}(\text{Bailout}(k))$.*

We define the **Gain** objective as $\text{MP}_{[1,d]}(> 0)$. The following lemma shows that an almost surely winning strategy σ_{Gain}^* for this objective can be chosen as memoryless randomized.

► **Lemma 6** ([7, Proposition 5.1]). *There is a uniform MR strategy σ_{Gain}^* which is almost surely winning for **Gain** (or any subset of dimensions) from all states $s \in \text{AS}(\text{Gain})$.*

A difference between \mathcal{M}^* and \mathcal{M} is that if one can almost surely win Energy-MeanPayoff in \mathcal{M}^* then one can also push the energy level arbitrarily high. This does not always hold in \mathcal{M} . (Consider, e.g., a single-state Markov chain with a single loop with reward 0 in the 1st dimension and +1 in all other dimensions.) The difference comes from the loop at state s_{win} in \mathcal{M}^* which has a strictly positive reward in all dimensions. Thus, the following lemma only holds for \mathcal{M}^* .

► **Lemma 7.** *In \mathcal{M}^* , there are two uniform memoryless strategies $\sigma_{\text{Bailout}}^*$ and σ_{Gain}^* which, starting from any state $s \in \text{AS}(\mathbf{0}(k))$, almost surely satisfy $\text{Bailout}(k)$ and **Gain**, respectively.*

Proof. Let $s \in \text{AS}(\mathbf{0}(k))$. We show that $s \in \text{AS}(\text{Bailout}(k))$ and $s \in \text{AS}(\text{Gain})$. The existence of the memoryless strategies $\sigma_{\text{Bailout}}^*$ and σ_{Gain}^* then follows from Lemma 5 and Lemma 6, respectively.

We assumed that all states s in \mathcal{M} admit an almost surely winning strategy for Energy-MeanPayoff. By Lemma 3, this also holds for all states q in \mathcal{M}^* . Let σ_q^\sharp denote an almost surely winning strategy from q for $\mathbf{0}(i_q)$ in \mathcal{M}^* (without restrictions on memory).

Recall from Section 2 that the random variable X_t denotes the state at time t , and Y_t denotes the (d -dimensional) sum of the rewards until time t .

▷ **Claim 8.** For every state $q \in \mathcal{M}^*$ there exists some number of steps $n_q \in \mathbb{N}$ and a probability $p_q > 0$ such that

$$\mathcal{P}_{\sigma_q^\sharp, q}^{\mathcal{M}^*} \left(\bigcup_{j=0}^{n_q} ((Y_j)_1 > i_{X_j} - i_q) \cup ((Y_j)_1 \geq f_{X_j} - i_q) \right) \geq p_q.$$

Proof. Towards a contradiction, assume that for all m

$$\mathcal{P}_{\sigma_q^\sharp, q}^{\mathcal{M}^*} \left(\bigcup_{j=0}^m ((Y_j)_1 > i_{X_j} - i_q) \cup ((Y_j)_1 \geq f_{X_j} - i_q) \right) = 0.$$

Due to the second part of the union, this implies that never $(Y_j)_1 + i_q \geq f_{X_j}$. Since σ_q^\sharp satisfies $\text{EN}_1(i_q)$ almost surely, it can never choose the step to s_{win} . This implies $\mathcal{P}_{\sigma_q^\sharp, q}^{\mathcal{M}^*}(\mathbf{F}_{s_{\text{win}}}) = 0$, i.e., X_j is always different from s_{win} . (The values f_s were initially defined with respect to states s of the original MDP \mathcal{M} , but the definition is naturally extended to the MDP \mathcal{M}^* , by giving the primed states the same value, i.e., $f_{s'} = f_s$. The state s_{win} does not appear in \mathcal{M} , but only in \mathcal{M}^* . We can extend the definition by having $f_{s_{\text{win}}} = 0$. However, this is not strictly required. The f_{X_j} is already defined, since X_j is always different from s_{win} .)

Since σ_q^\sharp satisfies $\text{EN}_1(i_q)$ almost surely, all runs always satisfy $(Y_j)_1 \geq i_{X_j} - i_q$ for all j . On the other hand, our assumption yields $\mathcal{P}_{\sigma_q^\sharp, q}^{\mathcal{M}^*} \left(\bigcup_{j=0}^m (Y_j)_1 > i_{X_j} - i_q \right) = 0$. This implies that $(Y_j)_1 = i_{X_j} - i_q$ for all j . Hence, in all runs the energy fluctuates by at most $\ell \stackrel{\text{def}}{=} 2 \max_q i_q$.

Thus, $\mathcal{P}_{\sigma_q^\#}^{\mathcal{M}^*}(0(i_q) \cap \text{Infix}_1(\ell)) = 1$. Then Lemma 4 implies that $f_q = i_q$. Since $X_0 = q$ we have $f_{X_0} = f_q$ and thus $(Y_0)_1 \geq f_{X_0} - i_q = 0$. This contradicts our assumption, since the second part of the union is surely satisfied. \triangleleft

For any state q , let n_q, p_q denote the values from Claim 8.

Now we show that $s \in \text{AS}(\text{Bailout}(k))$. Define a strategy σ_{Bailout} which plays in phases, separated by resets. It remembers the number of steps $t \geq 0$ since last reset, the (under-approximated) sum of rewards Q_t and the current state X_t . The first phase starts at state s and σ_{Bailout} plays like $\sigma_s^\#$ until one of the following events occur.

1. There is enough energy such that it is safe to move to s_{win} , i.e., $(Q_t \geq f_{X_t} - i_s)$, or
2. The current energy level is strictly greater than the minimal required energy level of the current state, i.e., $(Q_t > i_{X_t} - i_s)$, or
3. n_s steps have elapsed, i.e., $(t = n_s)$.

If at any point Item 1 happens, then the strategy simply goes to s_{win} . If it is the case that Item 2 occurs before $t = n_s$, let's say at some time t' , then the phase ends at t' . The sum of the rewards in the phase, between the last reset (where $t = 0$) and the current time is $\geq i_{X_{t'}} - i_s + 1$. If neither Item 1 nor Item 2 occurs before $t = n_s$, then the phase ends and we let $t' \stackrel{\text{def}}{=} t = n_s$. The sum of the rewards in this phase is then exactly $i_{X_{t'}} - i_s$. At the end of the phase σ_{Bailout} resets the number of steps ($t = 0$), and Q_t to 0. In the following phase it moves according to $\sigma_{X_{t'}}^\#$ until the next reset.

σ_{Bailout} clearly satisfies $\text{EN}_1(k)$ as it is a mix of energy safe strategies $(\sigma_q^\#)_{q \in S^*}$ and since we are starting from a safe energy level. By Claim 8, there is a positive probability (lower-bounded by $\min_q p_q > 0$) that either Item 1 or Item 2 happens in each phase.

Hence, unless event Item 1 occurs, Item 2 occurs infinitely often almost surely. Moreover, since the length of phases is upper bounded by $\max_q n_q$, it occurs frequently. We obtain $\mathcal{P}_{\sigma_{\text{Bailout},s}^{\mathcal{M}^*}}(\text{MP}_1 \geq \min_q \left(\frac{p_q}{n_q}\right) > 0 \mid \neg \text{FS}_{s_{\text{win}}}) = 1$. On the other hand, if s_{win} is reached, then MP_1 holds by shift invariance and the definition of the positive rewards in the loop at s_{win} . Therefore, $\mathcal{P}_{\sigma_{\text{Bailout},s}^{\mathcal{M}^*}}(\text{EN}_1(i_s) \cap \text{MP}_1(> 0)) = 1$.

Now we show that $s \in \text{AS}(\text{Gain})$. We make use of the following strategies.

- $\sigma_q^\#$ which satisfies $\text{EN}_1(k) \cap \text{MP}_{[2,d]}(> 0)$ almost surely from q for every $k \geq i_q$.
- a uniform MD strategy $\sigma_{\text{MP}_1}^*$ which satisfies $\text{MP}_1(> 0)$ almost surely from every state. It exists since $\text{AS}(\text{MP}_1(> 0)) = S^*$ (where S^* is the set of states of \mathcal{M}^*), because $\mathcal{P}_{\sigma_{\text{Bailout},s}^{\mathcal{M}^*}}(\text{EN}_1(i_s) \cap \text{MP}_1(> 0)) = 1$.

From the former, we get probabilistic bounds on the achievable mean payoff in all the dimensions, i.e., for all states s , and $0 \leq \varepsilon < 1$, there is a $d - 1$ dimensional vector $\nu_\varepsilon > \mathbf{0}$ such that $\mathcal{P}_{\sigma_s^\#}^{\mathcal{M}^*}(\text{MP}_{[2,d]} \geq \nu_\varepsilon) \geq 1 - \frac{\varepsilon}{2}$. This follows from the fact that for any sequence of decreasing vectors $\nu_n \rightarrow \mathbf{0}$ in \mathbb{R}^{d-1} , $\text{MP}_{[2,d]}(> \mathbf{0}) = \bigcup_n \text{MP}_{[2,d]}(\geq \nu_n)$ and continuity of measures. Furthermore, denoting by \mathbf{Y}_t the sum of rewards in all dimensions until time t , there exists a sufficiently large bound $n_\varepsilon \in \mathbb{N}$ such that $\mathcal{P}_{\sigma_s^\#}^{\mathcal{M}^*}\left(\frac{(Y_t)_j}{t} \geq \frac{(\nu_\varepsilon)_j}{2}\right) \geq 1 - \varepsilon$ in each of the dimensions $j \in [2, d]$ for all $t \geq n_\varepsilon$ steps. This can be shown by observing that $\text{MP}_j(\geq (\nu_\varepsilon)_j) = \bigcap_{k=1}^\infty \bigcup_{n=1}^\infty \bigcap_{t=n}^\infty \left(\frac{(Y_t)_j}{t} \geq (\nu_\varepsilon)_j \cdot \left(1 - \frac{1}{2^k}\right)\right)$ and using continuity of measures.

Similarly, there exists a bound $n_\varepsilon^* \in \mathbb{N}$ and value $\nu_\varepsilon^* > 0$ such that $\mathcal{P}_{\sigma_{\text{MP}_1}^*,s}^{\mathcal{M}^*}\left(\frac{(Y_t)_1}{t} \geq \frac{\nu_\varepsilon^*}{2}\right) \geq 1 - \varepsilon$ after $t \geq n_\varepsilon^*$ steps for every state s .

Now consider the following strategy σ_{Gain} , which switches between two phases.

Phase 1: If the current state is q , it moves according to $\sigma_q^\#$ for some number $\alpha > n_\varepsilon$ of steps.

Then it switches to phase 2.

133:10 Finite-Memory Strategies for Almost-Sure Energy-MeanPayoff in MDPs

Phase 2: It moves according to $\sigma_{\text{MP}_1}^*$ for some number $\beta > n_\varepsilon^*$ of steps, and then switches back to phase 1.

The strategy σ_{Gain} is a finite-memory strategy, since the lengths of the alternating phases are bounded by α and β , respectively. (Even if σ_q^\sharp is an infinite-memory strategy, it can only use bounded memory in each phase.)

We fix σ_{Gain} from the start state s and obtain a finite-state Markov chain. In every BSCC of this Markov chain, the expected mean payoff in the 1st dimension will be

$$\geq \frac{-i^\sharp + \beta \cdot (1 - \varepsilon) \cdot \left(\frac{\nu_\varepsilon^*}{2}\right) - \beta \cdot \varepsilon \cdot R}{\alpha + \beta}.$$

where $i^\sharp = \max_s i_s$ denotes the maximum (over all states) minimal safe energy.

Similarly, in every BSCC, the expected mean payoff in the j^{th} dimension for $j \geq 2$ can be lower-bounded by

$$\geq \frac{\alpha \cdot \left((1 - \varepsilon) \cdot \left(\frac{\nu_\varepsilon}{2}\right)_j - \varepsilon \cdot R\right) - \beta \cdot R}{\alpha + \beta}.$$

By choosing ε sufficiently small, β sufficiently large to make the first term positive and $\alpha \gg \beta$ sufficiently large to make the second term positive, we can get positive expected mean payoff in all dimensions. Since this holds in every BSCC of the induced finite Markov chain, the objective **Gain** is satisfied almost surely. ◀

The following lemma shows the converse of Lemma 7. In \mathcal{M}^* , it is always possible to win $\mathbf{0}(i_s)$ almost surely from s by playing a particular strategy $\sigma_{\text{alt}, Z_b, Z_g}^*$ which combines the two uniform memoryless strategies $\sigma_{\text{Bailout}}^*$ and σ_{Gain}^* . Let Z_b denote the minimal universally safe energy level for **Bailout**, i.e., $Z_b \stackrel{\text{def}}{=} \max_s \min\{k \mid s \in \text{AS}(\text{Bailout}(k))\}$. Moreover, let $Z_g > Z_b$ be a larger energy level at which our strategy switches from $\sigma_{\text{Bailout}}^*$ to σ_{Gain}^* .

Similarly to [25], we define an infinite-memory strategy $\sigma_{\text{alt}, Z_b, Z_g}^*$ that always records the current energy level and operates by switching between two phases. It starts by playing σ_{Gain}^* (**Gain**-phase) if our starting energy level is sufficiently high ($\geq Z_b + R$), and otherwise starts by playing $\sigma_{\text{Bailout}}^*$ (**Bailout**-phase). In the **Bailout**-phase, the primary goal is to pump the energy level up until it is $\geq Z_g$, and then it switches to the **Gain**-phase. It enters the **Bailout**-phase again if the energy level drops below $Z_b + R$ (in which case it will still be $\geq Z_b$).

► **Lemma 9.** *There exists a $Z_g \in \mathbb{N}$ such that for every s in \mathcal{M}^* the strategy $\sigma_{\text{alt}, Z_b, Z_g}^*$ is almost surely winning for $\mathbf{0}(i_s)$ from s .*

Proof. The parameter Z_g is chosen sufficiently large such that there is a fixed non-zero probability that after every **Bailout**-phase one never needs another **Bailout**. (Thus, except in a null set there are only finitely many **Bailouts**.) The existence of such a finite Z_g is guaranteed by the fact that $\lim_{k \rightarrow \infty} \mathcal{P}_{\sigma_{\text{Gain}}^*, s}(\mathbf{0}(k)) = 1$. ([18, Lemma 22]). Eventually, except in a null set, $\sigma_{\text{alt}, Z_b, Z_g}^*$ plays **Gain** forever, thus satisfying $\mathbf{0}(i_s)$ almost surely from s . ◀

Some combined objectives like **Energy-Parity** really require infinite memory for almost surely winning strategies [25]. However, we show that a sufficiently large *finite* memory is enough to win **Energy-MeanPayoff** almost surely. The idea is to modify the strategy $\sigma_{\text{alt}, Z_b, Z_g}^*$ such that it remembers the current energy only in the interval $[0, b]$, for some sufficiently large $b > Z_g$, and ignores any possible excess energy above b . This modified strategy is denoted by $\sigma_{\text{alt}, Z_b, Z_g, b}^*$, and it has a finite set of memory modes $[0, b] \times \{0, 1\}$. The $\{0, 1\}$ part is used to remember the current phase (**Gain** = 0 or **Bailout** = 1). Then $\sigma_{\text{alt}, Z_b, Z_g, b}^*[(u, x)]$ denotes the strategy $\sigma_{\text{alt}, Z_b, Z_g, b}^*$ with current memory mode $(u, x) \in [0, b] \times \{0, 1\}$.

The finite bound b on the remembered energy has the effect that $\sigma_{\text{alt},Z_b,Z_g,b}^*$ can no longer guarantee a fixed positive probability of not needing another Bailout after each Bailout-phase. Thus, one might have infinitely many Bailouts with positive probability. (Most of these are unnecessary, but one cannot be sure which ones). Unlike for Energy-Parity, where using infinitely many Bailout phases can compromise the objective, the nature of the $\text{MP}_{[2,d]}(> \mathbf{0})$ objective allows us to use infinitely many Bailouts with non-zero probability, provided that they happen sufficiently infrequently.

By its construction, the strategy $\sigma_{\text{alt},Z_b,Z_g,b}^*[(i_s, x)]$ is energy-safe from every state s , every initial energy $\geq i_s$ and $x \in \{0, 1\}$. It remains to show that it also satisfies $\text{MP}_{[2,d]}(> \mathbf{0})$ almost surely. Since $\sigma_{\text{alt},Z_b,Z_g,b}^*$ is finite-memory, it suffices to consider the induced finite Markov chain \mathcal{A} and show that the expected mean payoff is strictly positive in every BSCC. I.e., we prove that $\mathcal{E}_{\sigma_{\text{alt},Z_b,Z_g,b}^*}(\text{MP}_{[2,d]}) > \mathbf{0}$ for a sufficiently large b . To this end, we consider the finite Markov chains $\mathcal{A}^{\text{Gain}}$ and $\mathcal{A}^{\text{Bailout}}$ obtained by fixing the memoryless strategies σ_{Gain}^* and $\sigma_{\text{Bailout}}^*$ in \mathcal{M}^* , respectively. The application of $\sigma_{\text{alt},Z_b,Z_g,b}^*$ can then be seen as alternating between these two Markov chains based on hitting certain energy levels.

Let T^{Gain} denote the random variable that measures the length of a Gain-phase, when starting at energy level Z_g and assuming that the energy is truncated at b . Similarly, T^{Bailout} is the random variable that measures the length of a Bailout-phase when starting at energy level Z_b . (Here it does not matter that the energy is truncated at b , since the Bailout-phase ends when the energy reaches $Z_g < b$.) Since R can be > 1 , the Bailout-phase might actually start at a slightly higher energy level $u \in [Z_b, Z_b + R - 1]$, and thus T^{Bailout} over-approximates the actual length of the Bailout-phase, which is conservative for our analysis. Similarly, the Gain phase might start with an energy slightly higher than Z_g , and T^{Gain} under-approximates the length of the Gain-phase, which is again conservative. The random variables $(Y_{T^{\text{Gain}}})_i$ and $(Y_{T^{\text{Bailout}}})_i$ then measure the sum of the rewards the i^{th} dimension obtained during the Gain and Bailout phases, respectively.

The following lemma shows that the strategy $\sigma_{\text{alt},Z_b,Z_g,b}^*$ can attain a strictly positive mean payoff in all dimensions $i \in [2, d]$, provided that the expected reward during the Gain-phase is sufficiently large (positive) and the expected reward during the Bailout-phase (though possibly negative) is not too small.

► **Lemma 10.** *If there are constants $v_i^1 > 0$ and v_i^2 such that, for all $i \in [2, d]$ and states q*

$$\begin{aligned} \mathcal{E}_{\sigma_{\text{alt},Z_b,Z_g,b}^*}^{\mathcal{M}^*}((Y_{T^{\text{Gain}}})_i) &\geq v_i^1 \\ \mathcal{E}_{\sigma_{\text{alt},Z_b,Z_g,b}^*}^{\mathcal{M}^*}((Y_{T^{\text{Bailout}}})_i) &\geq v_i^2 \\ v_i^1 + v_i^2 &> 0 \end{aligned}$$

then $\mathcal{E}_{\sigma_{\text{alt},Z_b,Z_g,b}^*}^{\mathcal{M}^*}(\text{MP}_i) > 0$ for all s and $m \in [i_s, b] \times \{0, 1\}$.

Proof. By fixing the finite-memory strategy $\sigma_{\text{alt},Z_b,Z_g,b}^*$, we obtain a finite Markov chain. Consider any BSCC in this Markov chain. In this BSCC, except for a null set of runs, either no Bailouts happen or infinitely many. In the former case, this BSCC behaves like playing σ_{Gain}^* forever, which attains a strictly positive mean payoff in all dimensions almost surely, and thus a strictly positive expected mean payoff in each dimension i . In the second case, almost surely there happen infinitely many Bailouts, each starting at an every level $\geq Z_b$. Then, by the finiteness of the BSCC, we obtain that $\mathcal{E}(T^{\text{Gain}}) < \infty$. Moreover, by the definition of $\sigma_{\text{Bailout}}^*$, the expected duration of the Bailout-phase is always finite, i.e., $\mathcal{E}(T^{\text{Bailout}}) < \infty$. Thus, by linearity of expectations, $\mathcal{E}_{\sigma_{\text{alt},Z_b,Z_g,b}^*}^{\mathcal{M}^*}(\text{MP}_i) \geq (v_i^1 + v_i^2)/(\mathcal{E}(T^{\text{Gain}}) + \mathcal{E}(T^{\text{Bailout}})) > 0$. ◀

The following technical Lemma 11 (proof in [18, Appendix B]) shows that the constants v_i^1, v_i^2 from Lemma 10 exist. Recall that the finite Markov chains $\mathcal{A}^{\text{Gain}}$ and $\mathcal{A}^{\text{Bailout}}$ are obtained by fixing the memoryless strategies σ_{Gain}^* and $\sigma_{\text{Bailout}}^*$ in \mathcal{M}^* , respectively. Let $x_{\min,1}$ and $x_{\min,2}$ denote the minimal occurring non-zero probabilities in these two Markov chains, respectively. (They come from solutions of linear programs and can be chosen as only exponentially small, i.e., described by a polynomial number of bits; cf. [18, Appendix B]). The proof works by applying general results about expected first passage times in truncated Markov chains to the induced Markov chains $\mathcal{A}^{\text{Gain}}$ and $\mathcal{A}^{\text{Bailout}}$. The general idea is that in the Gain-phase one has a general up drift in all dimensions, and in particular in the first (energy) dimension. It is thus unlikely to go down very far in the energy dimension, even if the energy is truncated at b . Thus, for a sufficiently large truncation point b (actually $b = Z_g + 1$ suffices), the expected time spent in the Gain-phase is very large relative to the expected time spent in the Bailout phase. More exactly, the former increases exponentially in b , while the latter is polynomial in b . For a sufficiently large b (exponential in $|\mathcal{M}^*|$), the condition $v_i^1 + v_i^2 > 0$ is met.

► **Lemma 11.** *Let $\mu_i > 0$ denote the lower bound on the mean payoff in dimension i in any BSCC in the Markov chain $\mathcal{A}^{\text{Gain}}$ with corresponding computable constants $c_i, g_{\text{Gain}}, h_{\text{Gain}}$, and let μ denote the lower bound on the mean payoff in the 1st dimension in any BSCC of $\mathcal{A}^{\text{Bailout}}$ with the corresponding constants $g_{\text{Bailout}}, h_{\text{Bailout}}$. All the above constants, except c_i , can be chosen as at most exponential in $|\mathcal{M}^*|$ and $1/(1 - c_i) \in \mathcal{O}(\exp(\exp(|\mathcal{M}^*|^{\mathcal{O}(1)})))$.*

Then there are constants $0 < C_1 < 1, C_2 > 0, C_3 > 0, C_4 > 0, C_5 > 0$, all exponential in $|\mathcal{M}^|$ and dependent only on \mathcal{M} , such that for $k \stackrel{\text{def}}{=} \frac{2 \cdot |\mathcal{S}^*|}{x_{\min,1}^{\mathcal{O}(1)}} \in \mathcal{O}(\exp(|\mathcal{M}^*|^{\mathcal{O}(1)}))$, any*

$\delta \in (0, 1)$ sufficiently small such that $(|\mathcal{S}^| + 1) \cdot (\frac{1}{\delta} - 1) + \lceil (\log_{c_i}(\delta \cdot (1 - c_i))) \rceil \geq \frac{h_{\text{Gain}}}{\mu_i}$ for all $2 \leq i \leq d$, one can choose*

$Z_g \stackrel{\text{def}}{=} Z_b + R + k \cdot R + \max_i (R \cdot \lceil \log_{c_i}(\delta \cdot (1 - c_i)) \rceil - R + 1, h_{\text{Gain}}) \in \mathcal{O}(e^{|\mathcal{M}^|^{\mathcal{O}(1)}} \cdot \log(1/\delta))$*

and $b \stackrel{\text{def}}{=} Z_g + 1$ so that

$$\mathcal{E}_{\sigma_{\text{alt}, Z_b, Z_g, b}^*}^{\mathcal{M}^*}((Y_{T^{\text{Gain}}})_i) \geq C_1 \cdot \frac{1}{\delta} - C_2 \log_2\left(\frac{1}{\delta}\right) - C_3 \stackrel{\text{def}}{=} v_i^1$$

$$\mathcal{E}_{\sigma_{\text{alt}, Z_b, Z_g, b}^*}^{\mathcal{M}^*}((Y_{T^{\text{Bailout}}})_i) \geq -C_4 \log_2\left(\frac{1}{\delta}\right) - C_5 \stackrel{\text{def}}{=} v_i^2$$

In particular, in order to satisfy the condition $v_i^1 + v_i^2 > 0$, it suffices to choose $1/\delta \in \mathcal{O}(\max(1/C_1, \max_{2 \leq j \leq 5} C_j)^{\mathcal{O}(1)})$. Since the constants C_j are exponential in $|\mathcal{M}^|$, and by the conditions on the other constants above, the value Z_g , and hence the overall bound $b = Z_g + 1$, can be chosen such that $b \in \mathcal{O}(\exp(|\mathcal{M}^*|^{\mathcal{O}(1)}))$.*

Now we can prove the first item of our main result.

Proof of Theorem 1 (Item 1). Towards a contradiction, we assume that there exists a state s^\dagger such that there is no finite-memory almost surely winning strategy from s^\dagger for $\mathcal{O}(i_{s^\dagger})$ in the MDP \mathcal{M} .

First we consider the MDP \mathcal{M}^* . The finite-memory strategy $\sigma_{\text{alt}, Z_b, Z_g, b}^*[(i_{s^\dagger}, 1)]$ from s^\dagger is energy-safe by construction and satisfies $\text{EN}_1(i_{s^\dagger})$ surely. Now consider the finite Markov chain induced by fixing this finite-memory strategy in \mathcal{M}^* . By Lemma 10 and Lemma 11, for a sufficiently large (exponential) b it yields a strictly positive expected mean payoff $v_i^1 + v_i^2 > 0$ in every dimension $i \in [2, d]$ in every BSCC of this Markov chain. Since the Markov chain is finite, this implies that the mean payoff in every dimension $i \in [2, d]$ is strictly positive almost

surely. Hence, $\mathcal{P}_{\sigma_{\text{alt},z_b,z_g,b}^*}^{\mathcal{M}^*}(\text{MP}_{[2,d]}(> 0)) = 1$ and thus $\mathcal{P}_{\sigma_{\text{alt},z_b,z_g,b}^*}^{\mathcal{M}^*}(\mathcal{O}(i_{s^\dagger})) = 1$. So there exists an almost surely winning finite-memory strategy from s^\dagger for $\mathcal{O}(i_{s^\dagger})$ in \mathcal{M}^* . However, Lemma 3 then implies that there also exists an almost surely winning finite-memory strategy from s^\dagger for $\mathcal{O}(i_{s^\dagger})$ in \mathcal{M} . Contradiction. \blacktriangleleft

► **Remark 12.** If $\sigma_{\text{alt},z_b,z_g,b}^*$ satisfies $\mathcal{O}(i_s)$ almost surely from some state s then it also satisfies the stronger objective $\mathcal{O}(i_s) \cap \text{Infix}(b)$ almost surely. Consider a winning run induced by $\sigma_{\text{alt},z_b,z_g,b}^*$. While the true energy might sometimes be higher than b , the energy remembered by $\sigma_{\text{alt},z_b,z_g,b}^*$ is always $\leq b$. Even with this conservative under-approximation of the energy, the run still satisfies the energy objective. Therefore, in any winning run induced by $\sigma_{\text{alt},z_b,z_g,b}^*$, the energy can never *decrease* by more than b . Thus, also $\text{Infix}(b)$ is satisfied almost surely.

5 Proof of Item 2

Given some state s , let $\sigma = (\mathcal{M}, \mathbf{m}_0, \text{upd}, \text{nxt})$ be a finite-memory strategy that is almost surely winning for $\mathcal{O}(i_s)$ (which exists by Item 1). We show there exists an almost surely winning strategy σ' for $\mathcal{O}(i_s)$ such that the energy fluctuations are bounded by some constant which is exponential in $|\mathcal{M}|$.

First, inside any BSCC B of \mathcal{M}^σ , we construct an almost surely winning strategy σ_B and upper bound the minimal safe energy levels and energy fluctuation while following σ_B . Using this, we upper bound the energy fluctuations in paths before reaching a BSCC. We use the fact that the set of states and transitions that occur in any BSCC of a Markov chain induced by fixing some finite-memory strategy in an MDP is an end component of this MDP ([19, Theorem 3.2]).

► **Lemma 13.** *Let B be a BSCC of \mathcal{M}^σ and let $\mathcal{M}(B)$ be the corresponding end component in \mathcal{M} with states S_B and transitions E_B . Then there is a strategy σ_B , a bound $b_B \in \mathcal{O}(\exp(|\mathcal{M}(B)|^{\mathcal{O}(1)}))$ such that for any state $q \in S_B$, there is a minimal safe energy level $j_q \stackrel{\text{def}}{=} i_q^{\mathcal{M}(B)} \leq 3 \cdot |S_B| \cdot R$ such that $\mathcal{P}_{\sigma_B,q}^{\mathcal{M}(B)}(\mathcal{O}(j_q) \cap \text{Infix}(b_B)) = 1$.*

Proof Sketch. (Full proof in [18, Appendix C].) The idea is that for $\mathcal{M}(B)$ there are two cases. In the first case it behaves similar to \mathcal{M}^* from Section 4, in the sense that it is possible to win **Gain** and **Bailout** almost surely, and thus **Energy-MeanPayoff** can be won almost surely by switching between the two strategies for **Gain** and **Bailout** like in the strategy $\sigma_{\text{alt},z_b,z_g}^*$. Then one can invoke Lemma 11 and Remark 12 on $\mathcal{M}(B)$ to get an exponential bound b_B such that $\mathcal{P}_{\sigma_B,q}^{\mathcal{M}(B)}(\mathcal{O}(j_q) \cap \text{Infix}(b_B)) = 1$.

If the first case does not hold then $\mathcal{M}(B)$ is very restrictive, and one can show that the energy level fluctuations are bounded by a constant in $\mathcal{O}(|S_B| \cdot R)$. \blacktriangleleft

Since the minimal safe energy levels inside these end components are not too large, one can then bound the energy fluctuations in paths before they reach any such end component $\mathcal{M}(B)$. The following lemma is shown in [18, Appendix C].

► **Lemma 14.** *Let T denote the union of all S_B of every BSCC B of \mathcal{M}^σ , as in Lemma 13. Then one can almost surely reach any state in T with the corresponding minimal safe energy level with energy fluctuations of at most $5 \cdot |S| \cdot R$.*

Proof of Theorem 1 (Item 2). By Lemmas 13 and 14, for each state s , one can choose a strategy σ and some constant $b \in \mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}))$ such that $\mathcal{P}_{\sigma,s}^{\mathcal{M}}(\mathcal{O}(i_s) \cap \text{Infix}(b)) = 1$. This means if one encodes the energy levels between $[0, b]$ into the state space by discarding

any excess energy above b and redirecting all the transitions which result in a negative energy to a losing sink (for $\text{MP}_{[2,d]}(> 0)$) and constructs this larger MDP $\mathcal{M}[0, b]$, then there is a strategy σ' such that $\mathcal{P}_{\sigma', (s,k)}^{\mathcal{M}[0,b]}(\text{MP}_{[2,d]}(> 0)) = 1$ for every $k \in [i_s, b]$. Then, by Lemma 6, there also exists a memoryless (MR) strategy σ^* in $\mathcal{M}[0, b]$ which is almost surely winning $\text{MP}_{[2,d]}(> 0)$ from (s, k) .

We can carry the memoryless strategy σ^* in $\mathcal{M}[0, b]$ back to \mathcal{M} as a finite-memory strategy $\sigma_{\mathcal{M}}^*$ with memory $[0, b]$. It stores the encoded under-approximated energy level from $\mathcal{M}[0, b]$ in its finite memory instead. Thus, $\sigma_{\mathcal{M}}^*$ is a finite-memory strategy from s that satisfies $\mathfrak{O}(i_s)$ almost surely, and the size of its memory is bounded by $b \in \mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}))$.

The strategy $\sigma_{\mathcal{M}}^*$ uses randomization, because σ^* from Lemma 6 is MR. However, the MR strategy σ^* for the mean payoff objective could be replaced by a deterministic strategy with an exponential number of memory modes. Hence, the overall number of memory modes in the obtained deterministic version of $\sigma_{\mathcal{M}}^*$ is still only exponential. ◀

6 The Lower Bound (Proof of Item 3)

In the previous sections we have shown that finite memory suffices for almost surely winning strategies for the Energy-MeanPayoff objective. However, the required memory depends on the given MDP. We show that no fixed finite amount of memory is sufficient for all MDPs. In fact, the required memory is exponential in the transition probabilities even for an otherwise fixed 5-state MDP with just one controlled state, $R = 1$ and $d = 2$.

► **Definition 15.** Let $1 > \delta > 0$ and $\mathcal{M}_{\delta} = (S, S_{\square}, S_{\circ}, E, P, \mathbf{r})$ be an MDP with 2-dimensional rewards. It has just one controlled state s with transitions $s \rightarrow s_l$ and $s \rightarrow s_r$. From s_l there are two transitions $e_1 = (s_l \rightarrow s_l^1)$ and $e_2 = (s_l \rightarrow s_l^2)$. Let $\mathcal{P}(e_1) = (1 + \delta)/2$ and $\mathcal{P}(e_2) = (1 - \delta)/2$ and $\mathbf{r}(e_1) = (+1, +1)$ and $\mathbf{r}(e_2) = (-1, -1)$. s_l^1 and s_l^2 are random states which each have just one transition back to s with probability 1 and reward $\mathbf{0}$. From s_r there is only one transition e_3 back to s with probability 1 and $\mathbf{r}(e_3) = (+1, -1)$.

The following lemma directly implies the exponential lower bound on the number of memory modes in Theorem 1(Item 3).

► **Lemma 16.** Consider the Energy-MeanPayoff objective. For every finite bound $m \in \mathbb{N}$ on the number of memory modes there exists a $\delta \stackrel{\text{def}}{=} 1/(6m) > 0$ such that the finite MDP $\mathcal{M}_{\delta} = (S, S_{\square}, S_{\circ}, E, P, \mathbf{r})$ from Definition 15 satisfies the following properties.

1. $\exists \sigma' \mathcal{P}_{\sigma', s}^{\mathcal{M}_{\delta}}(\text{EN}_1(0) \cap \text{MP}_2(> 0)) = 1$, i.e., it is possible to win almost surely from s in \mathcal{M}_{δ} , even with initial energy 0.
2. For every finite-memory strategy σ with $\leq m$ memory modes we have $\mathcal{P}_{\sigma, s}^{\mathcal{M}_{\delta}}(\text{EN}_1(k) \cap \text{MP}_2(> 0)) = 0$ for every $k \in \mathbb{N}$, i.e., σ attains nothing in \mathcal{M}_{δ} , regardless of the initial energy k .
3. For \mathcal{M}_{δ} we have $|S| = 5$, $d = 1$ and $R = 1$. The number of memory modes required for an almost-surely winning strategy in \mathcal{M}_{δ} is exponential in $|P|$ (and in $|\mathcal{M}_{\delta}|$).

Proof. Towards item 1, consider a strategy σ' that plays as follows. It keeps a counter that records the current energy, which is initially 0. Whenever the current energy is 0, it plays $s \rightarrow s_r$, otherwise it plays $s \rightarrow s_l$. Thus, σ' satisfies $\text{EN}_1(0)$ surely from s . Since $\delta > 0$ it follows from the classic Gambler's ruin problem (with strictly positive expected gain, here in the first reward dimension) that σ' plays $s \rightarrow s_r$ only finitely often, except in a null set of the runs. Therefore, the expected mean payoff (in the second dimension) under σ' is $(1 + \delta)/2 - (1 - \delta)/2 = \delta > 0$. Hence, $\mathcal{P}_{\sigma', s}^{\mathcal{M}_{\delta}}(\text{MP}_2(> 0)) = 1$. Since the energy objective is satisfied surely, we obtain $\mathcal{P}_{\sigma', s}^{\mathcal{M}_{\delta}}(\text{EN}_1(0) \cap \text{MP}_2(> 0)) = 1$.

Towards item 2, let $\delta \stackrel{\text{def}}{=} 1/(6m) > 0$ and let σ be a finite-memory strategy with $\leq m$ memory modes. Consider the finite-state Markov chain \mathcal{C} that is induced by playing σ from s in \mathcal{M}_δ . This Markov chain has $\leq 5m$ states, since \mathcal{M} has 5 states and σ has $\leq m$ memory modes. Let B be any BSCC of \mathcal{C} that is reachable from s and the initial memory mode of σ . In particular, $|B| \leq 5m$. In B there must not exist any loop that does not contain s_r , because otherwise the energy objective cannot be satisfied almost surely. Thus, every path in B of length $\geq 5m$ must contain s_r (and hence a reward $(+1, -1)$) at least once. Therefore, the expected mean payoff in B (in the second reward dimension) is $\leq 5m\delta - 1 = -1/6 < 0$. Since this holds in every reachable BSCC, we obtain $\mathcal{P}_{\sigma,s}^{\mathcal{M}_\delta}(\text{MP}_2(> 0)) = 0$ and thus $\mathcal{P}_{\sigma,s}^{\mathcal{M}_\delta}(\text{EN}_1(k) \cap \text{MP}_2(> 0)) = 0$.

Towards item 3, the size of \mathcal{M}_δ follows from Definition 15. By items 1 and 2, the required number of memory modes m for an almost-surely winning strategy satisfies $m > 1/(6\delta)$. Since $|P| = \Theta(\log(1/\delta))$ and $|\mathcal{M}_\delta| = \Theta(|P|)$, we obtain $m = \Omega(\exp(|P|))$ and $m = \Omega(\exp(|\mathcal{M}_\delta|))$. ◀

The exponential lower bound on the required memory does not require probabilities encoded in binary like in Lemma 16. One can construct an equivalent example with polynomially many states where all transition probabilities are $1/2$. This is because one can encode exponentially small probabilities 2^{-k} with a chain of k extra states and transition probabilities $1/2$.

7 Computational Complexity

We have shown that the existence of an almost surely winning strategy for the Energy-MeanPayoff objective for a given state and initial energy level in an MDP implies the existence of a deterministic such strategy with exponentially many memory modes (unlike for Energy-Parity which requires infinite memory in general [25]).

A related problem is the decidability of the question whether a given state in an MDP and a given initial energy level admit an almost surely winning strategy for Energy-MeanPayoff. This problem is decidable in *pseudo-polynomial* time, using an algorithm very similar to the one for Energy-Parity presented in [25]. I.e., the time is polynomial, provided that the bound R on the rewards is given in unary. Transition probabilities in the MDP can still be represented in binary. The crucial point is that it suffices to witness the mere *existence* of an almost surely winning strategy, regardless of its memory. Basically, it suffices that the algorithm proves that the infinite-memory strategy $\sigma_{\text{alt},z_b,z_g}^*$ wins almost surely (plus a small extra argument about a corner case where the energy fluctuates only in a bounded region). The algorithm does not need to compute the bound b or to explicitly construct the finite-memory strategy $\sigma_{\text{alt},z_b,z_g,b}^*$.

▶ **Proposition 17.** *Let $\mathcal{M} = (S, S_\square, S_\circ, E, P, \mathbf{r})$ be an MDP with d -dimensional rewards on the edges $\mathbf{r} : E \rightarrow [-R, R]^d$. For any state s and $k \in \mathbb{N}$, the existence of an almost surely winning strategy from s for the multidimensional Energy-MeanPayoff objective $\text{EN}_1(k) \cap \text{MP}_{[2,d]}(> \mathbf{0})$ is decidable in pseudo-polynomial time (i.e., polynomial for R in unary).*

Proof. The proof is similar to the one for Energy-Parity presented in [25]. We outline the differences below. First, in the corner case where it is impossible to pump the energy up arbitrarily high almost surely from some state q , the only possible way to win Energy-MeanPayoff (resp. Energy-Parity) almost surely (if at all) is by using a non-null set of runs where the energy only ever fluctuates in a bounded region. In that case, the size of the energy fluctuations in these runs can safely be restricted to a region that is polynomial in

$|S| \cdot R$, and thus pseudo-polynomial in $|\mathcal{M}|$ [25]. It thus suffices to win multi-dimensional $\text{MP}_{[2,d]}(> 0)$ almost surely in a derived MDP \mathcal{M}' where the bounded energy is encoded into the states. Deciding this requires time polynomial in $|\mathcal{M}'|$ [12, 22] and thus pseudo-polynomial in $|\mathcal{M}|$. The winning situations of the corner case can then be encoded into \mathcal{M} , yielding a derived MDP \mathcal{M}' of pseudo-polynomial size, where Energy-MeanPayoff can be won almost surely if and only if it can be won almost surely by a combination of **Gain** and **Bailout** strategies, i.e., by strategy $\sigma_{\text{alt}, z_b, z_g}^*$. Therefore, it suffices to compute the states (and minimal initial energy levels k) where **Gain** and **Bailout**(k) can be won almost surely. The objective **Bailout**(k) $\stackrel{\text{def}}{=} \text{EN}_1(k) \cap \text{MP}_1(> 0)$ is exactly the same as the **Bailout** objective analysed in [25], and winning it almost surely is decidable in pseudo-polynomial time. Our objective **Gain** $\stackrel{\text{def}}{=} \text{MP}_{[1,d]}(> 0)$ differs from the **Gain** objective considered in [25] (which was $\text{MP}_1(> 0) \cap \text{Parity}$), but winning it almost surely is still decidable in polynomial time [12, 22] by solving a linear program. So overall the algorithm runs in pseudo-polynomial time. ◀

References

- 1 Pieter Abbeel and Andrew Y. Ng. Learning first-order Markov models for control. In *Advances in Neural Information Processing Systems 17*, pages 1–8. MIT Press, 2004.
- 2 Galit Ashkenazi-Golan, János Flesch, Arkadi Predtetchinski, and Eilon Solan. Reachability and safety objectives in Markov decision processes on long but finite horizons. *Journal of Optimization Theory and Applications*, 185:945–965, 2020.
- 3 Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.
- 4 Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2008.
- 5 Vincent D. Blondel and John N. Tsitsiklis. A survey of computational complexity results in systems and control. *Automatica*, 36(9):1249–1274, 2000.
- 6 T. Brázdil, A. Kučera, and P. Novotný. Optimizing the expected mean payoff in energy Markov decision processes. In *International Symposium on Automated Technology for Verification and Analysis (ATVA)*, volume 9938 of *LNCS*, pages 32–49, 2016.
- 7 Tomáš Brázdil, Václav Brožek, Krishnendu Chatterjee, Vojtěch Forejt, and Antonín Kučera. Markov decision processes with multiple long-run average objectives. *Logical Methods in Computer Science*, 10, 2014.
- 8 Véronique Bruyère, Quentin Hautem, Mickael Randour, and Jean-François Raskin. Energy Mean-Payoff Games. In Wan Fokkink and Rob van Glabbeek, editors, *30th International Conference on Concurrency Theory (CONCUR 2019)*, volume 140 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 21:1–21:17, Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CONCUR.2019.21.
- 9 Nicole Bäuerle and Ulrich Rieder. *Markov Decision Processes with Applications to Finance*. Springer-Verlag Berlin Heidelberg, 2011.
- 10 Arindam Chakrabarti, Luca De Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource interfaces. In *International Workshop on Embedded Software*, pages 117–133, 2003.
- 11 K. Chatterjee and T. Henzinger. A survey of stochastic ω -regular games. *Journal of Computer and System Sciences*, 78(2):394–413, 2012.
- 12 Krishnendu Chatterjee and Laurent Doyen. Energy and mean-payoff parity Markov decision processes. In *International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 6907, pages 206–218, 2011.
- 13 Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdziński. Mean-payoff parity games. In *Logic in Computer Science (LICS)*, pages 178–187, 2005.
- 14 Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors. *Handbook of Model Checking*. Springer, 2018. doi:10.1007/978-3-319-10575-8.
- 15 E.M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, December 1999.

- 16 Lorenzo Clemente and Jean-Francois Raskin. Multidimensional beyond worst-case and almost-sure problems for mean-payoff objectives. In *Proceedings of the 2015 30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 257–268, 2015.
- 17 Mohan Dantam and Richard Mayr. Approximating the value of energy-parity objectives in simple stochastic games. In *48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023)*, pages 38:1–38:15, 2023. doi:10.4230/LIPIcs.MFCS.2023.38.
- 18 Mohan Dantam and Richard Mayr. Finite-memory Strategies for Almost-sure Energy-MeanPayoff Objectives in MDPs, 2024. arXiv:2404.14522.
- 19 Luca De Alfaro. *Formal verification of probabilistic systems*. PhD thesis, Stanford University, 1997.
- 20 János Flesch, Arkadi Predtetchinski, and William Sudderth. Simplifying optimal strategies in limsup and liminf stochastic games. *Discrete Applied Mathematics*, 251:40–56, 2018.
- 21 Dean Gillette. Stochastic games with zero stop probabilities. *Contributions to the Theory of Games*, 3:179–187, 1957.
- 22 Hugo Gimbert, Youssouf Oualhadj, and Soumya Paul. Computing optimal strategies for Markov decision processes with parity and positive-average conditions. Working paper or preprint, 2011.
- 23 T.P. Hill and V.C. Pestien. The existence of good Markov strategies for decision processes with general payoffs. *Stoch. Processes and Appl.*, 24:61–76, 1987.
- 24 M. Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Information Processing Letters*, 68:119–124, 1998.
- 25 Richard Mayr, Sven Schewe, Patrick Totzke, and Dominik Wojtczak. MDPs with Energy-Parity Objectives. In *Logic in Computer Science (LICS)*. IEEE, 2017.
- 26 Richard Mayr, Sven Schewe, Patrick Totzke, and Dominik Wojtczak. Simple stochastic games with almost-sure energy-parity objectives are in NP and coNP. In *Proc. of Fossacs*, volume 12650 of *LNCS*, 2021. Extended version on arXiv. arXiv:2101.06989.
- 27 A. Puri. *Theory of hybrid systems and discrete event structures*. PhD thesis, University of California, Berkeley, 1995.
- 28 Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1994.
- 29 Manfred Schäl. Markov decision processes in finance and dynamic options. In *Handbook of Markov Decision Processes*, pages 461–487. Springer, 2002.
- 30 Olivier Sigaud and Olivier Buffet. *Markov Decision Processes in Artificial Intelligence*. John Wiley & Sons, 2013.
- 31 William D. Sudderth. Optimal Markov strategies. *Decisions in Economics and Finance*, 43:43–54, 2020.
- 32 R.S. Sutton and A.G Barto. *Reinforcement Learning: An Introduction*. Adaptive Computation and Machine Learning. MIT Press, 2018.