

Approximating the Value of Energy-Parity Objectives in Simple Stochastic Games

Mohan Dantam

School of Informatics, University of Edinburgh, UK

Richard Mayr

School of Informatics, University of Edinburgh, UK

Abstract

We consider simple stochastic games \mathcal{G} with energy-parity objectives, a combination of quantitative rewards with a qualitative parity condition. The Maximizer tries to avoid running out of energy while simultaneously satisfying a parity condition.

We present an algorithm to approximate the value of a given configuration in 2-NEXPTIME. Moreover, ε -optimal strategies for either player require at most $\mathcal{O}(2\text{-EXP}(|\mathcal{G}|) \cdot \log(\frac{1}{\varepsilon}))$ memory modes.

2012 ACM Subject Classification Computing methodologies → Stochastic games

Keywords and phrases Energy-Parity Games, Simple Stochastic Games, Parity, Energy

Digital Object Identifier [10.4230/LIPIcs.MFCS.2023.36](https://doi.org/10.4230/LIPIcs.MFCS.2023.36)

Related Version *Full version of a paper presented at MFCS 2023.*

1 Introduction

Background. *Simple stochastic games* (SSGs) are 2-player turn-based perfect information stochastic games played on finite graphs. They are also called *competitive Markov decision processes* [20], or *2 $\frac{1}{2}$ -player games* [13, 12]. Introduced by Shapley [36] in 1953, they have since played a central role in the solution of many problems, e.g., synthesis of reactive systems [35, 34] and formal specification and verification [17, 18, 1]. Every state either belongs to one of the players (Maximizer or Minimizer) or is a random state. In each round of the game the player who owns the current state gets to choose the successor state along the game graph. For random states the successor is chosen according to a predefined distribution. Given a start state and strategies of Maximizer and Minimizer, this yields a distribution over induced infinite plays. We consider objectives \mathfrak{O} that are measurable subsets of the set of possible plays, and the players try to maximize (resp. minimize) the probability of \mathfrak{O} .

Many different objectives for SSGs have been studied in the literature. Here we focus on parity, mean-payoff and energy objectives. We assign numeric rewards to transitions and priorities (aka colors), encoded by bounded non-negative numbers, to states. A play satisfies the (min-even) *parity objective* iff the minimal priority that appears infinitely often in a play is even. It subsumes all ω -regular objectives, and in particular safety, liveness, fairness, etc. On finite SSGs, the parity objective can be seen as a special case of the *mean-payoff objective* which requires the limit average reward per transition along a play to be positive (or non-negative). Mean-payoff objectives in SSGs go back to a 1957 paper by Gillette [21] and have been widely studied, due to their relevance for efficient control. The *energy objective* [6] requires that the accumulated reward at any time in a play stays above some finite threshold. The intuition is that a controlled system has some finite initial energy level that must never become depleted. Since the accumulated reward is not bounded a-priori, this essentially turns a finite-state game into an infinite-state one.

Energy-parity. We consider SSGs with *energy-parity* objectives, where plays need to



© Mohan Dantam and Richard Mayr;

licensed under Creative Commons License CC-BY 4.0

48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023).

Editors: Jérôme Leroux, Sylvain Lombardy, and David Peleg; Article No. 36; pp. 36:1–36:18

Leibniz International Proceedings in Informatics



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

satisfy both an energy and a parity objective. The parity objective specifies functional correctness, while the energy condition can encode efficiency or risk considerations, e.g., the system should not run out of energy since manually recharging would be costly or risky.

Previous work. Much work on combined objectives for stochastic systems is restricted to Markov decision processes (MDPs) [8, 9, 4, 28].

For (stochastic) games, the computational complexity of single objectives is often in $\text{NP} \cap \text{coNP}$, e.g., for parity or mean-payoff objectives [25]. Multi-objective games can be harder, e.g., satisfying *two different* parity objectives leads to coNP completeness [11].

Stochastic mean-payoff parity games can be solved in $\text{NP} \cap \text{coNP}$ [10]. However, this does not imply a solution for stochastic energy-parity games, since, unlike in the non-stochastic case [7], there is no known reduction from energy-parity to mean-payoff parity in stochastic games. The reduction in [7] relies on the fact that Maximizer has a winning *finite-memory* strategy for energy-parity, which does not generally hold for stochastic games, or even MDPs [28]. For the same reason, the direct reduction from stochastic energy-parity to ordinary energy games proposed in [8, 9] does not work for general energy-parity but only for energy-Büchi; cf. [28].

Non-stochastic energy-parity games can be solved in $\text{NP} \cap \text{coNP}$ (and even in pseudo-quasi-polynomial time [16]) and Maximizer strategies require only finite (but exponential) memory [7].

Stochastic energy-parity games have been studied in [29], where it was shown that the almost-sure problem is decidable and in $\text{NP} \cap \text{coNP}$. That is, given an initial configuration (control-state plus current energy level), does Maximizer have a strategy to ensure that energy-parity is satisfied with probability 1 against any Minimizer strategy? Unlike in many single-objective games, such an almost-surely winning Maximizer strategy (if it exists) requires infinite memory in general. This holds even in MDPs and for energy-coBüchi objectives [28].

However, [29] did not address quantitative questions about energy-parity objectives, such as computing/approximating the value of a given configuration, or the decidability of exact questions like “Is the value of this configuration $\geq k$?” for some constant k (e.g., $k = 1/2$).

The decidability of the latter type of exact question about the energy-parity value is open, but there are strong indications that it is very hard. In fact, even simpler subproblems are already at least as hard as the *positivity problem for linear recurrence sequences*, which in turn is at least as hard as the *Skolem problem* [19]. (The decidability of these problems has been open for decades; see [30] for an overview.) Given an SSG with an energy-parity objective, suppose we remove the parity condition (assume it is always true) and also suppose that Maximizer is passive (does not get to make any decisions). Then we obtain an MDP where the only active player (the Minimizer in the SSG) has a *termination objective*, i.e., to reach a configuration where the energy level is ≤ 0 . Exact questions about the value of the termination objective in MDPs are already at least as hard as the positivity problem [31, Section 5.2.3] (see also [32, 33]). Thus exact questions about the energy-parity value in SSGs are also at least as hard as the positivity problem.

Our contributions. Since exact questions about the energy-parity value in SSGs are positivity-hard, we consider the problem of computing approximations of the value. We present an algorithm that, given an SSG \mathcal{G} and error ε , computes ε -close approximations of the energy-parity value of any given configuration in 2-NEXPTIME. Moreover, we show that ε -optimal Maximizer (resp. Minimizer) strategies can be chosen as deterministic and using only finite memory with $\mathcal{O}(2\text{-EXP}(|\mathcal{G}|) \cdot \log(\frac{1}{\varepsilon}))$ memory modes. One can understand the idea as a constructive upper bound on the accuracy with which the players need to remember

the current energy level in the game. (This is in contrast to the result in [28] that almost-surely winning Maximizer strategies require infinite memory in general.) Once the upper bound on Maximizer's memory for ε -optimal strategies is established, one might attempt a reduction from energy-parity to mean-payoff parity, along similar lines as for non-stochastic games in [7]. However, instead we use a more direct reduction from energy-parity to parity in a derived SSG for our approximation algorithm.

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.

Games, MDPs and Markov chains. A *Simple Stochastic Game* (SSG) is a finite-state 2-player turn-based perfect-information stochastic game $\mathcal{G} = (S, (S_\square, S_\diamond, S_\circ), E, P)$ where the finite set of states S is partitioned into the states S_\square of the player \square (*Maximizer*), states S_\diamond of player \diamond (*Minimizer*), and chance vertices (aka random states) S_\circ . Let $E \subseteq S \times S$ be the transition relation. We write $s \rightarrow s'$ 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., $P(s) \in \mathcal{D}(\text{Succ}(s))$. For ease of presentation, we extend the domain of P to S^*S_\circ by $P(\rho s) \stackrel{\text{def}}{=} P(s)$ for all $\rho s \in S^+S_\circ$. An *MDP* is a game where one of the two players does not control any states. An MDP is *maximizing* (resp. *minimizing*) iff $S_\diamond = \emptyset$ (resp. $S_\square = \emptyset$). A *Markov chain* is a game with only random states, i.e., $S_\square = S_\diamond = \emptyset$.

Strategies. A *play* is an infinite sequence $s_0s_1 \dots \in S^\omega$ such that $s_i \rightarrow s_{i+1}$ for all $i \geq 0$. A *path* is a finite prefix of a play. Let $\text{Plays}(\mathcal{G}) \stackrel{\text{def}}{=} \{\rho = (q_i)_{i \in \mathbb{N}} \mid q_i \rightarrow q_{i+1}\}$ denote the set of all possible plays. A strategy of the player \square (\diamond) is a function $\sigma : S^*S_\square \rightarrow \mathcal{D}(S)$ ($\pi : S^*S_\diamond \rightarrow \mathcal{D}(S)$) that assigns to every path $ws \in S^*S_\square$ ($\in S^*S_\diamond$) 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 and \diamond in \mathcal{G} is denoted by $\Sigma_{\mathcal{G}}$ and $\Pi_{\mathcal{G}}$, respectively. A play/path $s_0s_1 \dots$ is compatible with a pair of strategies (σ, π) if $s_{i+1} \in \text{supp}(\sigma(s_0 \dots s_i))$ whenever $s_i \in S_\square$ and $s_{i+1} \in \text{supp}(\pi(s_0 \dots s_i))$ whenever $s_i \in S_\diamond$.

Finite-memory deterministic (FD) strategies are a subclass of strategies described by deterministic transducers $T = (M, m_0, \text{upd}, \text{nxt})$ where M is a finite set of memory modes with initial mode m_0 , $\text{upd} : M \times E \mapsto M$ updates the memory mode upon observing a transition and $\text{nxt} : M \times S_\circ \mapsto S$ chooses the successor state based on the current memory mode and state. FD strategies without memory ($|M| = 1$) are called memoryless deterministic (MD). For deterministic strategies, there is no difference between public memory (observable by the other player) and private memory.

Measure. A game \mathcal{G} with initial state s_0 and strategies (σ, π) yields a probability space $(s_0S^\omega, \mathcal{F}_{s_0}, \mathcal{P}_{\sigma, \pi, s_0}^{\mathcal{G}})$ 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 $\mathcal{P}_{\sigma, \pi, s_0}^{\mathcal{G}}$ is first defined on the cylinder sets. For $\rho = s_0 \dots s_n$, let $\mathcal{P}_{\sigma, \pi, s_0}^{\mathcal{G}}(\rho) \stackrel{\text{def}}{=} 0$ if ρ is not compatible with σ, π and otherwise $\mathcal{P}_{\sigma, \pi, s_0}^{\mathcal{G}}(\rho S^\omega) \stackrel{\text{def}}{=} \prod_{i=0}^{n-1} \tau(s_0 \dots s_i)(s_{i+1})$ where τ is σ or π or P depending on whether $s_i \in S_\square$ or S_\diamond or S_\circ , respectively. By Carathéodory's extension theorem [2], this defines a unique probability measure on the σ -algebra.

Objectives and Payoff functions. General objectives are defined by real-valued measurable functions. However, we only consider indicator functions of measurable sets.

Hence our objectives can be described by measurable subsets $\mathbf{0} \subseteq S^\omega$ of plays. The payoff, under strategies (σ, π) , is the probability that plays belong to $\mathbf{0}$.

We use the syntax and semantics of the LTL operators [14] **F** (eventually), **G** (always) and **X** (next) to specify some conditions on plays.

Reachability & Safety. A reachability objective is defined by a set of target states $T \subseteq S$. A play $\rho = s_0 s_1 \dots$ belongs to **FT** iff $\exists i \in \mathbb{N} s_i \in T$. Similarly, ρ belongs to $\mathbf{F}^{\leq n} T$ (resp. $\mathbf{F}^{\geq n} T$) iff $\exists i \leq n$ (resp. $i \geq n$) such that $s_i \in T$. Dually, the safety objective **GT** consists of all plays which never leave T . We have $\mathbf{GT} = \neg \mathbf{F}\neg T$.

Parity. A parity objective is defined via bounded function $Col : S \rightarrow \mathbb{N}$ that assigns non-negative priorities (aka colors) to states. Given an infinite play $\rho = s_0 s_1 \dots$, let $\text{Inf}(\rho)$ denote the set of numbers that occur infinitely often in the sequence $Col(s_0) Col(s_1) \dots$. A play ρ satisfies *even parity* w.r.t. Col iff the minimum of $\text{Inf}(\rho)$ is even. Otherwise, ρ satisfies *odd parity*. The objective even parity is denoted by **EPAR**(Col) and odd parity is denoted by **OPAR**(Col). Most of the time, we implicitly assume that the coloring function is known and just write **EPAR** and **OPAR**. Observe that, given any coloring Col , we have $\overline{\mathbf{EPAR}} = \mathbf{OPAR}$ and $\mathbf{OPAR}(Col) = \mathbf{EPAR}(Col + 1)$ where $Col + 1$ is the function which adds 1 to the color of every state. This justifies to consider only one of the even/odd parity objectives, but, for the sake of clarity, we distinguish these objectives wherever necessary.

Energy/Reward/Counter based objectives. Let $r : E \rightarrow \{-R, \dots, 0, \dots, R\}$ be a bounded function that assigns weights to transitions. Depending on context, the sum of these weights in a path can be viewed as energy, cost/reward or a counter. If $s \rightarrow s'$ 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 play. We say that ρ satisfies

1. the k -energy objective **EN**(k) iff $(k + \sum_{i=0}^{n-1} c_i) > 0$ for all $n \geq 0$.
2. the l -storage condition 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 play. Let **ST**(k, l) denote the set of plays that satisfy both the k -energy and the l -storage condition. Let $\mathbf{ST}(k) \stackrel{\text{def}}{=} \bigcup_l \mathbf{ST}(k, l)$. Clearly, $\mathbf{ST}(k) \subseteq \mathbf{EN}(k)$.
3. k -Termination **Term**(k) iff there exists $n \geq 0$ such that $(k + \sum_{i=0}^{n-1} c_i) \leq 0$.
4. *Limit* objective **LimInf**($\triangleright z$) iff $(\liminf_{n \rightarrow \infty} \sum_{i=0}^{n-1} c_i) \triangleright z$ for $\triangleright \in \{<, \leq, =, \geq, >\}$ and $z \in \mathbb{R} \cup \{\infty, -\infty\}$ and similarly for **LimSup**($\triangleright z$).
5. *Mean payoff* **MP**($\triangleright c$) for some constant $c \in \mathbb{R}$ iff $(\liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} c_i) \triangleright c$.

Observe that the objectives k -energy and k -termination are mutually exclusive and cover all of the plays. A different way to consider these objectives 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 safety/reachability objective, respectively.

An objective $\mathbf{0}$ is called *shift-invariant* iff for all finite paths ρ and plays $\rho' \in S^\omega$, we have $\rho\rho' \in \mathbf{0} \iff \rho' \in \mathbf{0}$. Parity and mean payoff objectives are shift-invariant, but energy and termination objectives are not. Objective $\mathbf{0}$ is called *submixing* iff for all sequences of finite non-empty words $u_0, v_0, u_1, v_1 \dots$ we have $u_0 v_0 u_1 v_1 \dots \in \mathbf{0} \implies ((u_0 u_1 \dots \in \mathbf{0}) \vee (v_0 v_1 \dots \in \mathbf{0}))$.

Determinacy. Given an objective $\mathbf{0}$ and a game \mathcal{G} , state s has value (w.r.t to $\mathbf{0}$) iff

$$\sup_{\sigma \in \Sigma_{\mathcal{G}}} \inf_{\pi \in \Pi_{\mathcal{G}}} \mathcal{P}_{\sigma, \pi, s}^{\mathcal{G}}(\mathbf{0}) = \inf_{\pi \in \Pi_{\mathcal{G}}} \sup_{\sigma \in \Sigma_{\mathcal{G}}} \mathcal{P}_{\sigma, \pi, s}^{\mathcal{G}}(\mathbf{0}).$$

If s has value then $\text{val}_{\mathbf{0}}^{\mathcal{G}}(s)$ denotes the value of s defined by the above equality. A game with an objective is called *weakly determined* if every state has value. Stochastic games with Borel objectives are weakly determined [26, 27]. Our objectives above are Borel, hence any boolean combination of them is also weakly determined. For $\varepsilon > 0$ and state s , a strategy

1. $\sigma \in \Sigma_{\mathcal{G}}$ is ε -optimal (maximizing) iff $\mathcal{P}_{\sigma, \pi, s}^{\mathcal{G}}(\mathbf{0}) \geq \text{val}_{\mathbf{0}}^{\mathcal{G}}(s) - \varepsilon$ for all $\pi \in \Pi_{\mathcal{G}}$.

2. $\pi \in \Pi_{\mathcal{G}}$ is ε -optimal (minimizing) iff $\mathcal{P}_{\sigma, \pi, s}^{\mathcal{G}}(0) \leq \text{val}_0^{\mathcal{G}}(s) + \varepsilon$ for all $\sigma \in \Sigma_{\mathcal{G}}$.

A 0-optimal strategy is called *optimal*. An MD strategy is called *uniformly ε -optimal* (resp. uniformly optimal) if it is so from every start state. An optimal strategy for player \square from state s is *almost surely winning* if $\text{val}_0^{\mathcal{G}}(s) = 1$. By $\text{AS}(0)$ we denote the set of states that have an almost surely winning strategy for objective 0. For ease of presentation, we drop subscripts and superscripts wherever possible if they are clear from the context.

Energy-parity. We are concerned with approximating the value for the combined energy-parity objective $\text{EN}(k) \cap \text{EPAR}$ and building ε -optimal strategies.

In our constructions we use some auxiliary objectives. Following [29], these are defined as $\text{Gain} \stackrel{\text{def}}{=} \text{LimInf}(\triangleright -\infty) \cap \text{EPAR}$ and $\text{Loss} \stackrel{\text{def}}{=} \overline{\text{Gain}} = \text{LimInf}(=\infty) \cup \text{OPAR}$.

► **Remark 1.** For finite-state SSGs and the following objectives there exist optimal MD strategies for both players. Moreover, if the SSG is just a maximizing MDP then the set of states that are almost surely winning for Maximizer can be computed in polynomial time.

1. FT [15]
2. $\text{LimInf}(\triangleright -\infty)$, $\text{LimInf}(\triangleright \infty)$, $\text{LimSup}(\triangleright -\infty)$, $\text{LimSup}(\triangleright \infty)$, $\text{MP}(\triangleright 0)$ [5, Prop. 1]
3. EPAR [37]

3 The Main Result

The following theorem states our main result.

► **Theorem 2.** *Let $\mathcal{G} = (S, (S_{\square}, S_{\diamond}, S_{\circ}), E, P)$ be an SSG with transition rewards in unary assigned by function r and colors assigned to states by function Col . For every state $s \in S$, initial energy level $i \geq 0$ and error margin $\varepsilon > 0$, one can compute*

1. *a rational number v' such that $0 \leq v' - \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}}(s) \leq \varepsilon$ in 2-NEXPTIME.*¹
2. *ε -optimal FD strategies σ_{ε} and π_{ε} for Maximizer and Minimizer, resp., in 2-NEXPTIME. These strategies use $\mathcal{O}(2\text{-EXP}(|\mathcal{G}|) \cdot \log(\frac{1}{\varepsilon}))$ memory modes.*

For rewards in binary, the bounds above increase by one exponential.

We outline the main steps of the proof; details in the following sections. We begin with the observation that $\text{EN}(i) \subseteq \text{EN}(j)$ for $i \leq j$, and thus for all states s we have $\text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}}(s) \leq \text{val}_{\text{EN}(j) \cap \text{EPAR}}^{\mathcal{G}}(s) \leq 1$. So $\lim_{n \rightarrow \infty} \text{val}_{\text{EN}(n) \cap \text{EPAR}}^{\mathcal{G}}(s)$ exists. We define

$$\text{Lval}^{\mathcal{G}}(s) \stackrel{\text{def}}{=} \lim_{n \rightarrow \infty} \text{val}_{\text{EN}(n) \cap \text{EPAR}}^{\mathcal{G}}(s). \quad (1)$$

We will see that $\text{Lval}^{\mathcal{G}}(s)$ and $\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$ are in fact equal (a consequence of Lemma 9) and $\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$ can be computed in nondeterministic polynomial time (Theorem 5). Intuitively, for high energy levels, the precise energy level does not matter much for the value.

The main steps of the approximation algorithm are as follows.

1. Compute FD strategies $\sigma^*(s)$ that are optimal maximizing for the objective Gain starting from state s in \mathcal{G} . Compute an MD strategy π^* that is uniformly optimal minimizing for the objective Gain . Compute the value $\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$ for every $s \in S$. See Section 4.
2. Compute a natural number N such that for all $s \in S$ and all $i \geq N$ we have

$$0 \leq \text{val}_{\text{Gain}}^{\mathcal{G}}(s) - \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}}(s) \leq \varepsilon.$$

N will be doubly exponential. See Section 5.

¹ We write “computing a number v' in 2-NEXPTIME” as a shorthand for the property that questions like $v' \leq c$ for constants c are decidable in 2-NEXPTIME.

3. Consider the finite-state parity game \mathcal{G}' derived from \mathcal{G} by encoding the energy level up-to N into the states, i.e., the states of \mathcal{G}' are of the form (s, k) for $s \in S$ and $0 \leq k \leq N$, and colors are inherited from s . Moreover, we add gadgets that ensure that states (s, N) at the upper end win with probability $\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$ and states $(s, 0)$ at the lower end lose. By the previous item, $\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$ is ε -close to $\text{val}_{\text{EN}(N) \cap \text{EPAR}}^{\mathcal{G}}(s)$. Thus, for $k < N$ we can ε -approximate the value $v = \text{val}_{\text{EN}(k) \cap \text{EPAR}}^{\mathcal{G}}(s)$ by $v' \stackrel{\text{def}}{=} \text{val}_{\text{EPAR}}^{\mathcal{G}'}((s, k))$. If $k \geq N$ we can ε -approximate v by $v' \stackrel{\text{def}}{=} \text{val}_{\text{Gain}}^{\mathcal{G}}(s)$.

Moreover, we obtain ε -optimal FD strategies σ_ε for Maximizer (resp. π_ε for Minimizer) for $\text{EN}(k) \cap \text{EPAR}$ in \mathcal{G} . Let $\hat{\sigma}$ (resp. $\hat{\pi}$) be optimal MD strategies for Maximizer (resp. Minimizer) for the objective EPAR in \mathcal{G}' . Then σ_ε plays as follows. While the current energy level j (k plus the sum of the rewards so far) stays $< N$, then, at any state s' , play like $\hat{\sigma}$ at state (s', j) in \mathcal{G}' . Once the energy level reaches a value $\geq N$ at some state s' for the first time, then play like $\sigma^*(s')$ forever. Similarly, π_ε plays as follows. While the current energy level j (k plus the sum of the rewards so far) stays $< N$, then, at any state s' , play like $\hat{\pi}$ at state (s', j) in \mathcal{G}' . Once the energy level reaches a value $\geq N$ (at any state) for the first time, then play like π^* forever. See Section 6.

As a technical tool, we sometimes consider the dual of a game \mathcal{G} (resp. the dual maximizing MDP of some minimizing MDP). Consider $\mathcal{G}^d \stackrel{\text{def}}{=} (S', (S'_\square, S'_\diamond, S'_\circ), E', P')$ with the complement objective $\overline{\text{EN}(k) \cap \text{EPAR}} = \text{Term}(k) \cup \text{OPAR}$, where \mathcal{G}^d is simply the game with the roles of Maximizer and Minimizer reversed, i.e.,

$$S' = S \quad S'_\square = S_\diamond \quad S'_\diamond = S_\square \quad S'_\circ = S_\circ \quad E' = E \quad P' = P$$

Hence $\Sigma_{\mathcal{G}^d} = \Pi_{\mathcal{G}}$ and $\Pi_{\mathcal{G}^d} = \Sigma_{\mathcal{G}}$. It is easy to see that for any objective $\mathbb{0}$ and start state s

1. $\text{val}_{\mathbb{0}}^{\mathcal{G}}(s) + \text{val}_{\overline{\mathbb{0}}}^{\mathcal{G}^d}(s) = 1$.
2. σ is ε -optimal maximizing for $\mathbb{0}$ in \mathcal{G} iff it is ε -optimal minimizing for $\overline{\mathbb{0}}$ in \mathcal{G}^d .
3. π is ε -optimal minimizing for $\mathbb{0}$ in \mathcal{G} iff it is ε -optimal maximizing for $\overline{\mathbb{0}}$ in \mathcal{G}^d .

So approximating the value of $\text{EN}(k) \cap \text{EPAR}$ in \mathcal{G} can be reduced in linear time to approximating the value of $\text{Term}(k) \cup \text{OPAR}$ in \mathcal{G}^d .

4 Computing $\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$

Given an SSG $\mathcal{G} = (S, (S_\square, S_\diamond, S_\circ), E, P)$ and a start state s , we will show how to compute $\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$ and the optimal strategies for both players.

We start with the case of maximizing MDPs. The following lemma summarizes some previous results ([29, Lemmas 30,16], [28, Lemma 26], [24, Proposition 4]).

► **Lemma 3.** *Let \mathcal{M} be a maximizing MDP.*

1. $\text{Lval}^{\mathcal{M}}(s) = \text{val}_{\text{Gain}}^{\mathcal{M}}(s)$ for all states $s \in S$.
2. *Optimal strategies for Gain in \mathcal{M} exist and can be chosen FD, with $\mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}))$ memory nodes, and exponential memory is also necessary.*
3. *For any state $s \in S$, $\text{Lval}^{\mathcal{M}}(s)$ is rational and can be computed in $\mathcal{O}(|\mathcal{M}|^8)$ deterministic polynomial time if rewards are in unary, and in NP and coNP if rewards are in binary.*

Proof. Item 1 holds by [29, Lemma 30].

Towards Item 2, we follow the proof of [29, Lemma 16]. Since $\text{Gain} = \text{LimInf}(> -\infty) \cap \text{EPAR}$ is shift-invariant, there exist optimal strategies by [22]. By [28, Theorem 18] and Item 1, an optimal strategy for Gain can be constructed as follows. Let $A \stackrel{\text{def}}{=} \bigcup_{k \in \mathbb{N}} \text{AS}(\text{ST}(k) \cap \text{EPAR})$ and $B \stackrel{\text{def}}{=} \text{AS}(\text{LimInf}(=\infty) \cap \text{EPAR})$ be the subsets of states from which there exist almost

surely winning strategies for the objectives $\text{ST}(k) \cap \text{EPAR}$ and $\text{LimInf}(=\infty) \cap \text{EPAR}$, respectively. By [28, Theorem 8], we can restrict the values k in the definition of A by some $k' = \mathcal{O}(|S| \cdot R)$, i.e., $A = \bigcup_{k \leq k'} \text{AS}(\text{ST}(k) \cap \text{EPAR})$. An optimal strategy σ for **Gain** works in two phases. First it plays an optimal strategy σ_R towards reaching the set $A \cup B$, where σ_R can be chosen MD by Remark 1. Then, upon reaching A (resp. B), it plays an almost surely winning strategy σ_A for the objective $\text{ST}(k) \cap \text{EPAR}$ (resp. σ_B for the objective $\text{LimInf}(=\infty) \cap \text{EPAR}$). By [28, Theorem 8], the strategy σ_A requires $\mathcal{O}(k \cdot |S|)$ memory modes for a given k and thus at most $\mathcal{O}(|S|^2 \cdot R)$, since we can assume that $k \leq k'$. Towards the strategy σ_B , we first observe that in finite MDPs a strategy is almost-surely winning for $\text{LimInf}(=\infty) \cap \text{EPAR}$ iff it is almost-surely winning for $\text{MP}(> 0) \cap \text{EPAR}$. By [24, Proposition 4], there exist optimal deterministic strategies for $\text{MP}(> 0) \cap \text{EPAR}$ that use exponential memory, i.e., $\mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}))$ memory modes. The memory required for σ_B exceeds that of σ_R and σ_A (even when R is given in binary), and the one extra memory mode to record the switch from σ_R to σ_A (resp. σ_B) is negligible in comparison. Thus we can conclude that σ uses $\mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}))$ memory modes. [24, Fig. 1 and Prop. 4] shows that exponential memory is necessary.

Towards Item 3, let $d \stackrel{\text{def}}{=} |\text{Col}(S)|$ be the number of priorities in the parity condition. By [28, Lemma 26], for each $s \in S$, $\text{Lval}^{\mathcal{M}}(s)$ is rational and can be computed in deterministic time $\tilde{\mathcal{O}}(|E| \cdot d \cdot |S|^4 \cdot R + d \cdot |S|^{3.5} \cdot (|P| + |r|)^2)$ (and still in NP and coNP if R is given in binary). So $\text{Lval}^{\mathcal{M}}(s)$ can be computed in $\mathcal{O}(|\mathcal{M}|^8)$ deterministic polynomial time if weights are given in unary, and in NP and coNP if weights are given in binary. ◀

In order to extend Lemma 3 from MDPs to games, we need the notion of derived MDPs, obtained by fixing the choices of one player according to some FD strategy. Given an SSG $\mathcal{G} = (S, (S_{\square}, S_{\diamond}, S_{\circ}), E, P)$ and a finite memory deterministic (FD) strategy π for Minimizer (resp. σ for Maximizer) from a state s , described by $(M, m_0, \text{upd}, \text{nxt})$, let \mathcal{G}_{π} (resp. \mathcal{G}^{σ}) be the maximizing (resp. minimizing) MDP with state space $M \times S$ obtained by fixing Minimizer's (resp. Maximizer's) choices according to π (resp. σ).

► **Lemma 4.** *For every SSG \mathcal{G} , objective 0 and Minimizer (resp. Maximizer) FD strategy $\pi = (M, m_0, \text{upd}, \text{nxt})$ (resp. σ), from state s we get $\text{val}_0^{\mathcal{G}^{\sigma}}((m_0, s)) \leq \text{val}_0^{\mathcal{G}}(s) \leq \text{val}_0^{\mathcal{G}_{\pi}}((m_0, s))$ and equality holds if π (resp. σ) is optimal from state s .*

► **Theorem 5.** *Consider an SSG $\mathcal{G} = (S, (S_{\square}, S_{\diamond}, S_{\circ}), E, P)$ with the **Gain** objective.*

1. *Optimal Minimizer strategies exist and can be chosen uniform MD.*
2. *$\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$ is rational and questions about it, i.e., $\text{val}_{\text{Gain}}^{\mathcal{G}}(s) \leq c$ for constants c , are decidable in NP.*
3. *Optimal Maximizer strategies exist and can be chosen FD, with $\mathcal{O}(\exp(|\mathcal{G}|^{\mathcal{O}(1)}))$ memory modes. Moreover, exponential memory is also necessary.*

Proof. Towards Item 1, observe that since both the objectives $\text{LimInf}(=-\infty)$ and **OPAR** are shift-invariant and submixing, so is their union, i.e., $\overline{\text{Gain}}$ is shift-invariant and submixing. Hence, by [23, Theorem 1.1], an optimal MD strategy π_s^* for Minimizer exists from any state $s \in S$. Since S is finite and **Gain** is shift-invariant, we can also obtain a uniformly optimal MD strategy π^* , i.e., π^* is optimal from every state.

Towards Item 2, consider the maximizing MDP \mathcal{G}_{π^*} obtained from \mathcal{G} by fixing π^* (cf. Definition 11). Since π^* is MD, the states of \mathcal{G}_{π^*} are the same as the states as \mathcal{G} . Since π^* is optimal for Minimizer from every state s , we obtain that $\text{val}_{\text{Gain}}^{\mathcal{G}}(s) = \text{val}_{\text{Gain}}^{\mathcal{G}_{\pi^*}}(s)$ for every state s by Lemma 4. By Lemma 3, the latter is rational and can be computed in polynomial time for weights in unary (resp. in NP and coNP for weights in binary). Thus, by guessing π^* , we can decide questions $\text{val}_{\text{Gain}}^{\mathcal{G}}(s) \leq c$ in NP.

Towards Item 3, we again use the property that $\overline{\text{Gain}}$ is shift-invariant and submixing (see above). By [29, Theorem 6, Def. 24], optimal FD Maximizer strategies for Gain in an SSG require only $|S_\diamond| \cdot \lceil \log(|E|) \rceil$ many extra bits of memory above the memory required for optimal Maximizer strategies in any derived MDP \mathcal{M} where Minimizer's choices are fixed. Hence, by Lemma 3, one can obtain optimal FD Maximizer strategies in \mathcal{G} that use at most $2^{|S_\diamond| \cdot \lceil \log(|E|) \rceil} \cdot \mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)})) = \mathcal{O}(\exp(|\mathcal{G}|^{\mathcal{O}(1)}))$ memory modes. The corresponding exponential lower bound on the memory holds already for MDPs by Lemma 3. \blacktriangleleft

5 Computing the Upper Bound N

We show how to compute the upper bound N , up-to which Maximizer needs to remember the energy level, for any given error margin $\varepsilon > 0$. Similarly as in Section 4, we first solve the problem for maximizing MDPs and then extend the solution to SSGs.

5.1 Computing N for maximizing MDPs

Given a maximizing MDP $\mathcal{M} = (S, S_\square, S_\circ, E, P)$ and $\varepsilon > 0$, we will compute an $N \in \mathbb{N}$ such that for all $s \in S$ and all $j \geq N$

$$0 \leq \text{val}_{\text{Term}(j) \cup \text{OPAR}}^{\mathcal{M}}(s) - \text{val}_{\text{Loss}}^{\mathcal{M}}(s) \leq \varepsilon.$$

Recall that $\text{Loss} = \text{LimInf}(= -\infty) \cup \text{OPAR}$. We now define the sets of states $W_0 \stackrel{\text{def}}{=} \text{AS}(\text{Loss})$, $W_1 \stackrel{\text{def}}{=} \text{AS}(\text{LimInf}(= -\infty))$ and $W_2 \stackrel{\text{def}}{=} \text{AS}(\text{OPAR})$. By Remark 1, there exist optimal MD strategies for $\text{LimInf}(= -\infty)$ and OPAR . Since Loss is shift-invariant and submixing, there exists an optimal MD strategy for it by [23, Theorem 1.1].

► **Lemma 6.** *For every state s in the MDP \mathcal{M} we have*

1. $W_1 \cup W_2 \subseteq W_0$
2. $\text{val}_{\text{FW}_0}(s) \leq \text{val}_{\text{Loss}}(s)$
3. $\text{val}_{\text{OPAR} \cap \overline{\text{FW}_2}}(s) = 0$
4. for every initial energy level $j \geq 0$

$$\text{val}_{(\text{Term}(j) \cup \text{OPAR}) \cap \text{FW}_0}(s) = \text{val}_{\text{FW}_0}(s) \tag{2}$$

$$\text{val}_{\text{Loss}}(s) \leq \text{val}_{\text{Term}(j) \cup \text{OPAR}}(s) \leq \text{val}_{\text{Loss}}(s) + \sup_{\sigma} \mathcal{P}_{\sigma, s}(\text{Term}(j) \cap \overline{\text{FW}_1}) \tag{3}$$

Proof.

1. This follows directly from the definitions of W_0, W_1, W_2 .
2. Let σ' be an optimal MD strategy for FW_0 from s and σ'' be an almost surely winning MD strategy for Loss from any state in W_0 . Let σ be the strategy that plays σ' until reaching W_0 and then switches to σ'' . We have $\text{val}_{\text{Loss}}(s) \geq \mathcal{P}_{\sigma, s}(\text{Loss}) \geq \mathcal{P}_{\sigma', s}(\text{FW}_0) = \text{val}_{\text{FW}_0}(s)$.
3. For $s \in W_2$ the statement is obvious. So let $s \notin W_2$ and consider the modified MDP $\mathcal{M}' = (S', S'_\square, S'_\circ, E', P')$ where all states in W_2 are collapsed into a losing sink. I.e., $S' \stackrel{\text{def}}{=} (S \setminus W_2) \uplus \{\text{trap}\}$, with trap a new random sink state having color 0 (thus losing for objective OPAR), E' contains all of $(E \cap \{(S \setminus W_2) \times (S \setminus W_2)\} \cup (\text{trap}, \text{trap}))$ and all transitions to W_2 are redirected to trap and P' is derived accordingly from P . Then $\text{val}_{\text{OPAR}}^{\mathcal{M}'}(\hat{s}) = \text{val}_{\text{OPAR} \cap \overline{\text{FW}_2}}^{\mathcal{M}}(\hat{s})$ for all states $\hat{s} \in S \setminus W_2$. Towards a contradiction, assume that $\text{val}_{\text{OPAR} \cap \overline{\text{FW}_2}}^{\mathcal{M}}(s) > 0$. Hence $\text{val}_{\text{OPAR}}^{\mathcal{M}'}(s) > 0$. Then, by [22, Theorem 3.2], there exists a state $s' \in S'$ such that $\text{val}_{\text{OPAR}}^{\mathcal{M}'}(s') = 1$, and it is easy to see that $s' \neq \text{trap}$ and thus $s' \in S \setminus W_2$. But this implies that $\text{val}_{\text{OPAR}}^{\mathcal{M}'}(s') = 1$ and thus $s' \in W_2$, a contradiction.

4. Let $0 \stackrel{\text{def}}{=} \text{Term}(j) \cup \text{OPAR}$. For Equation (2), the first inequality $\text{val}_{0 \cap \text{FW}_0}(s) \leq \text{val}_{\text{FW}_0}(s)$ is trivial, since $0 \cap \text{FW}_0 \subseteq \text{FW}_0$. To show the reverse inequality, consider the strategy σ that first plays like an optimal MD strategy σ' for the objective FW_0 and after reaching W_0 switches to an almost surely winning MD strategy σ'' for the objective Loss . Then $\text{val}_{0 \cap \text{FW}_0}(s) \geq \mathcal{P}_{\sigma,s}(0 \cap \text{FW}_0) \geq \mathcal{P}_{\sigma,s}(\text{Loss} \cap \text{FW}_0) = \mathcal{P}_{\sigma',s}(\text{FW}_0) = \text{val}_{\text{FW}_0}(s)$, where the second inequality is due to $\text{LimInf}(= -\infty) \subseteq \text{Term}(j)$.

For Equation (3), the first inequality is again due to the fact that $\text{LimInf}(= -\infty) \subseteq \text{Term}(j)$ for all $j \geq 0$. Towards the second inequality of Equation (3) we have

$$\begin{aligned}
& \text{val}_0(s) \\
&= \sup_{\sigma} \mathcal{P}_{\sigma,s}(0) \\
&= \sup_{\sigma} (\mathcal{P}_{\sigma,s}(0 \cap \text{FW}_0) + \mathcal{P}_{\sigma,s}(0 \cap \overline{\text{FW}_0})) && \text{(Law of total probability)} \\
&\leq \sup_{\sigma} \mathcal{P}_{\sigma,s}(0 \cap \text{FW}_0) + \sup_{\sigma} \mathcal{P}_{\sigma,s}(0 \cap \overline{\text{FW}_0}) && (\sup(f+g) \leq \sup f + \sup g) \\
&= \sup_{\sigma} \mathcal{P}_{\sigma,s}(\text{FW}_0) + \sup_{\sigma} \mathcal{P}_{\sigma,s}(0 \cap \overline{\text{FW}_0}) && \text{(Equation (2))} \\
&\leq \text{val}_{\text{Loss}}(s) + \sup_{\sigma} \mathcal{P}_{\sigma,s}(0 \cap \overline{\text{FW}_0}) && \text{(Item 2)}
\end{aligned}$$

We can upper-bound the second summand above as follows.

$$\begin{aligned}
& \sup_{\sigma} \mathcal{P}_{\sigma,s}(0 \cap \overline{\text{FW}_0}) \\
&= \sup_{\sigma} \mathcal{P}_{\sigma,s}((\text{Term}(j) \cup \text{OPAR}) \cap \overline{\text{FW}_0}) \\
&\leq \sup_{\sigma} \mathcal{P}_{\sigma,s}(\text{Term}(j) \cap \overline{\text{FW}_0}) + \sup_{\sigma} \mathcal{P}_{\sigma,s}(\text{OPAR} \cap \overline{\text{FW}_0}) && \text{(Union bound)} \\
&\leq \sup_{\sigma} \mathcal{P}_{\sigma,s}(\text{Term}(j) \cap \overline{\text{FW}_1}) + \sup_{\sigma} \mathcal{P}_{\sigma,s}(\text{OPAR} \cap \overline{\text{FW}_2}) && \text{(Item 1)} \\
&= \sup_{\sigma} \mathcal{P}_{\sigma,s}(\text{Term}(j) \cap \overline{\text{FW}_1}) && \text{(Item 3)} \quad \blacktriangleleft
\end{aligned}$$

We show that the term $\sup_{\sigma} \mathcal{P}_{\sigma,s}(\text{Term}(j) \cap \overline{\text{FW}_1})$ in Equation (3) can be made arbitrarily small for large j . To this end, we use [3, Lemma 3.9] (adapted to our notation).

► **Lemma 7.** [3, Lemma 3.9 and Claim 6] *Let $\mathcal{M} = (S, S_{\square}, S_{\circ}, E, P)$ be a maximizing finite MDP with rewards in unary and $W_1 \stackrel{\text{def}}{=} \text{AS}(\text{LimInf}(= -\infty))$. One can compute, in polynomial time, a rational constant $c < 1$, and an integer $h \geq 0$ such that for all $j \geq h$ and $s \in S$*

$$\sup_{\sigma} \mathcal{P}_{\sigma,s}(\text{Term}(j) \cap \overline{\text{FW}_1}) \leq \frac{c^j}{1-c}.$$

Moreover, $1/(1-c) \in \mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}))$ and $h \in \mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}))$.

► **Lemma 8.** *Consider a maximizing MDP $\mathcal{M} = (S, S_{\square}, S_{\circ}, E, P)$, $\varepsilon > 0$ and the constants c, h from Lemma 7. For rewards in unary and $i \geq N$ we have $\text{val}_{\text{Term}(i) \cup \text{OPAR}}^{\mathcal{M}}(s) - \text{val}_{\text{Loss}}^{\mathcal{M}}(s) \leq \varepsilon$ where $N \stackrel{\text{def}}{=} \max(h, \lceil \log_c(\varepsilon \cdot (1-c)) \rceil) \in \mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}) \cdot \log(1/\varepsilon))$.*

For rewards in binary we have $N \in \mathcal{O}(\exp(\exp(|\mathcal{M}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$, i.e., the size of N increases by one exponential.

Proof sketch. For rewards in unary, the result follows from Lemma 6 (Equation (3)) and Lemma 7. For rewards in binary, the constants increase by one exponential via encoding binary rewards into unary rewards in a modified MDP. ◀

5.2 Computing N for SSGs

In order to compute the bound N for an SSG \mathcal{G} , we first consider bounds $N(s)$ for individual states s and then take their maximum. Given a state s , we can use Theorem 5(Item 3) to obtain an optimal FD strategy (with $\mathcal{O}(\exp(|\mathcal{G}|^{\mathcal{O}(1)}))$ memory modes) $\sigma^*(s) = (\mathbf{M}, \mathbf{m}_0, \text{upd}, \text{nxt})$ for Maximizer from state s w.r.t. the **Gain** objective. Theorem 5(Item 1) yields a uniform MD strategy π^* that is optimal for Minimizer from all states s w.r.t. the **Gain** objective.

► **Lemma 9.** *Given an SSG $\mathcal{G} = (S, (S_{\square}, S_{\diamond}, S_{\circ}), E, P)$ and $\varepsilon > 0$, we can compute a number $N \in \mathbb{N}$ such that for all $i \geq N$ and states $s \in S$ we have*

$$\text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}}(s) - \varepsilon \leq \text{val}_{\text{Gain}}^{\mathcal{G}}(s) - \varepsilon \leq \inf_{\pi} \mathcal{P}_{\sigma^*(s), \pi, s}^{\mathcal{G}}(\text{EN}(i) \cap \text{EPAR}) \leq \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}}(s) \quad (4)$$

i.e., $\sigma^(s)$ is ε -optimal for Maximizer for $\text{EN}(i) \cap \text{EPAR}$ for all $i \geq N$. In particular, $0 \leq \text{val}_{\text{Gain}}^{\mathcal{G}}(s) - \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}}(s) \leq \varepsilon$.*

Moreover, π^ is ε -optimal for Minimizer from any state s for $i \geq N$.*

$$\sup_{\sigma} \mathcal{P}_{\sigma, \pi^*, s}^{\mathcal{G}}(\text{EN}(i) \cap \text{EPAR}) \leq \sup_{\sigma} \mathcal{P}_{\sigma, \pi^*, s}^{\mathcal{G}}(\text{Gain}) = \text{val}_{\text{Gain}}^{\mathcal{G}}(s) \leq \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}}(s) + \varepsilon \quad (5)$$

For rewards in unary, N is doubly exponential, i.e., $N \in \mathcal{O}(\exp(\exp(|\mathcal{G}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$ and it can be computed in exponential time. For rewards in binary, the size of N and its computation time increase by one exponential, respectively.

Proof. Assume that rewards are in unary. The first inequality of (4) holds because $\text{EN}(i) \cap \text{EPAR} \subseteq \text{Gain}$ for any i . The third inequality of (4) follows from the definition of the value. Towards the second inequality of (4), we consider the minimizing MDP $\mathcal{M}(s) \stackrel{\text{def}}{=} \mathcal{G}^{\sigma^*(s)}$ obtained by fixing the Maximizer strategy $\sigma^*(s)$. Since $\sigma^*(s)$ is optimal for Maximizer from state s wrt. the objective **Gain**, Lemma 4 yields that

$$\text{val}_{\text{Gain}}^{\mathcal{G}}(s) = \text{val}_{\text{Gain}}^{\mathcal{M}(s)}((\mathbf{m}_0, s)). \quad (6)$$

Since $\sigma^*(s)$ has $\mathcal{O}(\exp(|\mathcal{G}|^{\mathcal{O}(1)}))$ memory modes, the size of $\mathcal{M}(s)$ is exponential in $|\mathcal{G}|$ and $\mathcal{M}(s)$ can be computed in exponential time.

Now we consider the dual maximizing MDP $\mathcal{M}(s)^d$ and the objectives **Term**(i) \cup **OPAR** and **Loss**. (Note that $\mathcal{M}(s)^d$ has the same size as $\mathcal{M}(s)$.) From Lemma 8, we obtain a bound $N(s) \in \mathbb{N}$ such that for all $i \geq N(s)$

$$0 \leq \text{val}_{\text{Term}(i) \cup \text{OPAR}}^{\mathcal{M}(s)^d}((\mathbf{m}_0, s)) - \text{val}_{\text{Loss}}^{\mathcal{M}(s)^d}((\mathbf{m}_0, s)) \leq \varepsilon. \quad (7)$$

By Lemma 8 and Lemma 7, $N(s)$ is exponential in $|\mathcal{M}(s)^d|$ and thus doubly exponential in $|\mathcal{G}|$, i.e., $N(s) \in \mathcal{O}(\exp(\exp(|\mathcal{G}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$. Moreover, $N(s)$ can be computed in time polynomial in $|\mathcal{M}(s)^d|$ and thus in time exponential in $|\mathcal{G}|$. By duality, we can rewrite Equation (7) for $\mathcal{M}(s)$ as follows. For all $i \geq N(s)$

$$0 \leq \text{val}_{\text{Gain}}^{\mathcal{M}(s)}((\mathbf{m}_0, s)) - \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{M}(s)}((\mathbf{m}_0, s)) \leq \varepsilon. \quad (8)$$

In order to get a uniform upper bound that holds for all states, let $N \stackrel{\text{def}}{=} \max_{s \in S} N(s)$. Since $|S|$ is linear, we still have $N \in \mathcal{O}(\exp(\exp(|\mathcal{G}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$ and it can be computed in

exponential time in $|\mathcal{G}|$. Finally, we can show the second inequality of (4).

$$\begin{aligned}
& \inf_{\pi} \mathcal{P}_{\sigma^*(s), \pi, s}^{\mathcal{G}} (\text{EN}(i) \cap \text{EPAR}) \\
&= \inf_{\pi} \mathcal{P}_{\pi, (m_0, s)}^{\mathcal{M}(s)} (\text{EN}(i) \cap \text{EPAR}) \\
&= \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{M}(s)} ((m_0, s)) \\
&\geq \text{val}_{\text{Gain}}^{\mathcal{M}(s)} ((m_0, s)) - \varepsilon && \text{by } i \geq N \geq N(s) \text{ and Equation (8)} \\
&= \text{val}_{\text{Gain}}^{\mathcal{G}} (s) - \varepsilon && \text{by (6)}
\end{aligned}$$

The first inequality of (5) holds because $\text{EN}(i) \cap \text{EPAR} \subseteq \text{Gain}$ for any i . The equality in (5) holds by the optimality of π^* . The second inequality of (5) follows from the previously stated consequence of (4).

For rewards in binary, the sizes of the numbers $N(s)$ (and hence N) and the time to compute it increase by one exponential by Lemma 8. ◀

6 Unfolding the Game to Energy Level N

Given an SSG $\mathcal{G} = (S, (S_{\square}, S_{\diamond}, S_{\circ}), E, P)$ and error tolerance $\varepsilon > 0$, for each state $s \in S$ and energy level $i \geq 0$, we want to compute a rational number v' which satisfies $0 \leq v' - \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}} (s) \leq \varepsilon$, and ε -optimal FD strategies σ_{ε} and π_{ε} for Maximizer and Minimizer, resp. We achieve this by constructing a finite-state parity game \mathcal{G}' that closely approximates the original game \mathcal{G} , as described in Section 3 (Item 3).

For clarity, we explain the construction in two steps. In the first step, we consider a finite-state parity game $\mathcal{G}[N]$. (Unlike \mathcal{G}' , the game $\mathcal{G}[N]$ is not actually constructed. It just serves as a part of the correctness proof.) $\mathcal{G}[N]$ encodes the energy level up-to $N + R$ (where R is the maximal transition reward) into the states, i.e., it has states of the form (s, k) with $k \leq N + R$. It imitates the original game \mathcal{G} till energy level $N + R$, but at any state (s, i) with energy level $i \geq N$ it jumps to a winning state with probability $\text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}} (s)$ and to a losing state with probability $1 - \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}} (s)$. (We need the margin up-to $N + R$, because transitions can have rewards > 1 , so the level N might not be hit exactly.) Similarly, at states $(s, 0)$ with energy level 0, we jump to a losing state. The coloring function in the new game $\mathcal{G}[N]$ derives its colors from the colors in the original game \mathcal{G} , i.e., all states (s, i) have the same color as s in \mathcal{G} .

By construction of $\mathcal{G}[N]$, for $i \leq N$, the EPAR value of (s, i) in $\mathcal{G}[N]$ coincides with $\text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}} (s)$.

In the second step, since we do not know the exact values $\text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}} (s)$ for $N + R \geq i > N$, we approximate these by the slightly larger $\text{val}_{\text{Gain}}^{\mathcal{G}} (s)$. I.e., we modify $\mathcal{G}[N]$ by replacing the probability values $\text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}} (s)$ for the jumps to the winning state by $\text{val}_{\text{Gain}}^{\mathcal{G}} (s)$. Let \mathcal{G}' be the resulting finite-state parity game. It follows from Lemma 9 that $0 \leq \text{val}_{\text{Gain}}^{\mathcal{G}} (s) - \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}} (s) \leq \varepsilon$ for $i \geq N$ and $\text{Lval}_{\text{EN} \cap \text{EPAR}}^{\mathcal{G}} (s) = \text{val}_{\text{Gain}}^{\mathcal{G}} (s)$. Thus \mathcal{G}' ε -over-approximates $\mathcal{G}[N]$ and \mathcal{G} , and we obtain the following lemma.

► **Lemma 10.** *For all states s and all $0 \leq i \leq N$*

$$\begin{aligned}
& \text{val}_{\text{EPAR}}^{\mathcal{G}[N]} ((s, i)) = \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}} (s), \text{ and} \\
& 0 \leq \text{val}_{\text{EPAR}}^{\mathcal{G}'} ((s, i)) - \text{val}_{\text{EPAR}}^{\mathcal{G}[N]} ((s, i)) \leq \varepsilon.
\end{aligned}$$

Now we are ready to prove the main theorem.

► **Theorem 2.** Let $\mathcal{G} = (S, (S_{\square}, S_{\diamond}, S_{\circ}), E, P)$ be an SSG with transition rewards in unary assigned by function r and colors assigned to states by function Col . For every state $s \in S$, initial energy level $i \geq 0$ and error margin $\varepsilon > 0$, one can compute

1. a rational number v' such that $0 \leq v' - \text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}}(s) \leq \varepsilon$ in 2-NEXPTIME.²
2. ε -optimal FD strategies σ_{ε} and π_{ε} for Maximizer and Minimizer, resp., in 2-NEXPTIME. These strategies use $\mathcal{O}(2\text{-EXP}(|\mathcal{G}|) \cdot \log(\frac{1}{\varepsilon}))$ memory modes.

For rewards in binary, the bounds above increase by one exponential.

Proof. For $i > N$ we output $v' = \text{val}_{\text{Gain}}^{\mathcal{G}}(s)$, which satisfies the condition by Lemma 9. For $i \leq N$ we output $v' = \text{val}_{\text{EPAR}}^{\mathcal{G}'}((s, i))$, which satisfies the condition by Lemma 10. By Theorem 5, the values $\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$ are rational for all states s . Therefore all probability values in \mathcal{G}' are rational and thus the EPAR values of all states in \mathcal{G}' are rational. Hence our numbers v' are always rational.

By Theorem 5, the values $\text{val}_{\text{Gain}}^{\mathcal{G}}(s)$ for all states $s \in S$ can be computed in exponential time. By Lemma 9, $N \in \mathcal{O}(\exp(\exp(|\mathcal{G}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$ is doubly exponential. Therefore, we can construct \mathcal{G}' in $\mathcal{O}(\exp(\exp(|\mathcal{G}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$ time and space. Questions about the parity values of states in \mathcal{G}' can be decided in nondeterministic time polynomial in $|\mathcal{G}'|$. Thus the numbers v' are computed in 2-NEXPTIME.

Towards Item 2, we construct ε -optimal FD strategies σ_{ε} for Maximizer (resp. π_{ε} for Minimizer) for $\text{EN}(i) \cap \text{EPAR}$ in \mathcal{G} . Let $\hat{\sigma}$ (resp. $\hat{\pi}$) be optimal MD strategies for Maximizer (resp. Minimizer) for the objective EPAR in \mathcal{G}' , which exist by Remark 1. Since these strategies are MD, they can be guessed in nondeterministic time polynomial in the size $|\mathcal{G}'|$, and thus in $\mathcal{O}(\exp(\exp(|\mathcal{G}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$ nondeterministic time.

Then σ_{ε} plays as follows. While the current energy level j (i plus the sum of the rewards so far) stays $< N$, then, at any state s' , play like $\hat{\sigma}$ at state (s', j) in \mathcal{G}' . Once the energy level reaches a value $\geq N$ at some state s' for the first time, then play like $\sigma^*(s')$ forever. (Recall that $\sigma^*(s')$ is the optimal FD Maximizer strategy for Gain from state s' from Section 5.2.) σ_{ε} is ε -optimal by Lemma 10 and Lemma 9. It needs to remember the energy level up-to N while simulating $\hat{\sigma}$. Moreover, $\sigma^*(s')$ needs $\mathcal{O}(\exp(|\mathcal{G}|^{\mathcal{O}(1)}))$ memory modes by Theorem 5. Finally, it needs to remember the switch from $\hat{\sigma}$ to $\sigma^*(s')$. Since $N \in \mathcal{O}(\exp(\exp(|\mathcal{G}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$ dominates the rest, σ_{ε} uses $\mathcal{O}(\exp(\exp(|\mathcal{G}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$ memory modes.

Similarly, π_{ε} plays as follows. While the current energy level j stays $< N$, at any state s' , play like $\hat{\pi}$ at state (s', j) in \mathcal{G}' . Once the energy level reaches a value $\geq N$ (at any state) for the first time, then play like π^* forever (where π^* is the uniform optimal MD Minimizer strategy for Gain from Section 5.2.) π_{ε} is ε -optimal by Lemma 10 and Lemma 9. While π^* is MD and does not use any memory, π_{ε} still needs to remember the energy level up-to N while simulating $\hat{\pi}$, and thus it uses $\mathcal{O}(\exp(\exp(|\mathcal{G}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$ memory modes.

For rewards in binary, all bounds increase by one exponential via an encoding of \mathcal{G} into an exponentially larger but equivalent game with rewards in unary. ◀

No nontrivial lower bounds are known on the computational complexity of approximating $\text{val}_{\text{EN}(i) \cap \text{EPAR}}^{\mathcal{G}}(s)$. However, even without the parity part, the problem appears to be hard. The best known algorithm for approximating the value of the energy objective (resp. the dual termination objective) runs in NEXPTIME for SSGs with rewards in unary [3].

As for lower bounds on the strategy complexity, ε -optimal Maximizer strategies need at least an exponential number of memory modes (for any $0 < \varepsilon < 1$) even in maximizing

² We write “computing a number v' in 2-NEXPTIME” as a shorthand for the property that questions like $v' \leq c$ for constants c are decidable in 2-NEXPTIME.

MDPs. This can easily be shown by extending the example in Lemma 3(Item 2) and [24, Fig. 1 and Prop. 4] that shows the lower bound for the **Gain** objective. First loop in a state with an unfavorable color to accumulate a sufficiently large reward (depending on ε) and then switch to the MDP in [24, Fig. 1 and Prop. 4] to play for **Gain** (since $\text{EN}(i) \cap \text{EPAR}$ will be very close to **Gain** then). Even the latter part requires exponentially many memory modes.

7 Conclusion & Extensions

We gave a procedure to compute ε -approximations of the value of combined energy-parity objectives in SSGs. The decidability of questions about the exact values is open, but the problem is at least as hard as the positivity problem for linear recurrence sequences [31, Section 5.2.3]. Unlike almost surely winning Maximizer strategies which require infinite memory in general [28, 29], ε -optimal strategies for either player require only finite memory with at most doubly exponentially many memory modes.

An interesting topic for further study is whether these results can be extended to other combined objectives where the parity part is replaced by something else, i.e., energy-X for some objective X (e.g., some other color-based condition like Rabin/Streett, or a quantitative objective about multi-dimensional transition rewards). While our proofs are not completely specific to parity, they do use many strong properties that parity satisfies.

- Shift-invariance of **EPAR** is used in several places, e.g. in Lemma 6 (and thus its consequences) and for the correctness of the constructions in Section 6.
- We use the fact that **EPAR** goes well together with $\text{LimInf}(> -\infty)$, i.e., the objective $\text{Gain} = \text{LimInf}(> -\infty) \cap \text{EPAR}$ allows optimal FD strategies for Maximizer in MDPs; cf. Lemma 3.
- The submixing property of $\text{OPAR} = \overline{\text{EPAR}}$ is used in Theorem 5 to lift Lemma 3 from MDPs to SSGs.

References

- 1 Rajeev Alur, Thomas A Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
- 2 Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2008.
- 3 T. Brázdil, V. Brožek, K. Etessami, and A. Kučera. Approximating the Termination Value of One-Counter MDPs and Stochastic Games. *Information and Computation*, 222:121–138, 2013. URL: <https://arxiv.org/abs/1104.4978>.
- 4 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.
- 5 Tomáš Brázdil, Václav Brozek, and Kousha Etessami. One-Counter Stochastic Games. In Kamal Lodaya and Meena Mahajan, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, volume 8 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 108–119, Dagstuhl, Germany, 2010. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. Full version at <http://arxiv.org/abs/1009.5636>. URL: <http://drops.dagstuhl.de/opus/volltexte/2010/2857>, doi: [10.4230/LIPIcs.FSTTCS.2010.108](https://doi.org/10.4230/LIPIcs.FSTTCS.2010.108).
- 6 Arindam Chakrabarti, Luca De Alfaro, Thomas A Henzinger, and Mariëlle Stoelinga. Resource interfaces. In *International Workshop on Embedded Software*, pages 117–133, 2003.
- 7 Krishnendu Chatterjee and Laurent Doyen. Energy parity games. In *International Colloquium on Automata, Languages and Programming (ICALP)*, volume 6199 of *LNCS*, pages 599–610, 2010.
- 8 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.
- 9 Krishnendu Chatterjee and Laurent Doyen. Games and Markov decision processes with mean-payoff parity and energy parity objectives. In *Mathematical and Engineering Methods in Computer Science (MEMICS)*, volume 7119 of *LNCS*, pages 37–46. Springer, 2011.
- 10 Krishnendu Chatterjee, Laurent Doyen, Hugo Gimbert, and Youssouf Oualhadj. Perfect-information stochastic mean-payoff parity games. In *International Conference on Foundations of Software Science and Computational Structures (FoSSaCS)*, volume 8412 of *LNCS*, 2014.
- 11 Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Generalized parity games. In Helmut Seidl, editor, *International Conference on Foundations of Software Science and Computational Structures (FoSSaCS)*, volume 4423 of *LNCS*, pages 153–167. Springer, 2007. URL: https://doi.org/10.1007/978-3-540-71389-0_12, doi: [10.1007/978-3-540-71389-0_12](https://doi.org/10.1007/978-3-540-71389-0_12).
- 12 Krishnendu Chatterjee, Marcin Jurdziński, and Thomas A. Henzinger. Simple stochastic parity games. In *Computer Science Logic (CSL)*, volume 2803 of *LNCS*, pages 100–113. Springer, 2003.
- 13 Krishnendu Chatterjee, Marcin Jurdziński, and Thomas A. Henzinger. Quantitative stochastic parity games. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 121–130. SIAM, 2004.
- 14 E.M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, Dec. 1999.
- 15 Anne Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992. doi: [https://doi.org/10.1016/0890-5401\(92\)90048-K](https://doi.org/10.1016/0890-5401(92)90048-K).
- 16 Laure Daviaud, Martin Jurdziński, and Ranko Lazić. A pseudo-quasi-polynomial algorithm for mean-payoff parity games. In *Logic in Computer Science (LICS)*, pages 325–334, 2018.
- 17 Luca De Alfaro and Thomas A Henzinger. Interface automata. *ACM SIGSOFT Software Engineering Notes*, 26(5):109–120, 2001.
- 18 David L. Dill. *Trace theory for automatic hierarchical verification of speed-independent circuits*, volume 24. MIT press Cambridge, 1989.

- 19 Graham Everest, Alfred Jacobus van der Poorten, Igor Shparlinski, and Thomas Ward. *Recurrence sequences*. ACM, 2003.
- 20 J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.
- 21 Dean Gillette. Stochastic games with zero stop probabilities. *Contributions to the Theory of Games*, 3:179–187, 1957.
- 22 Hugo Gimbert and Florian Horn. Solving simple stochastic tail games. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 847–862, 2010. URL: <http://epubs.siam.org/doi/abs/10.1137/1.9781611973075.69>.
- 23 Hugo Gimbert and Edon Kelmendi. Submixing and shift-invariant stochastic games. 2022. URL: <https://arxiv.org/abs/1401.6575>.
- 24 Hugo Gimbert, Youssouf Oualhadj, and Soumya Paul. Computing optimal strategies for Markov decision processes with parity and positive-average conditions. 2011. URL: <https://hal.science/hal-00559173/en/>.
- 25 Marcin Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Information Processing Letters*, 68(3):119–124, 1998.
- 26 A. Maitra and W. Sudderth. Stochastic games with Borel payoffs. In *Stochastic Games and Applications*, pages 367–373. Kluwer, Dordrecht, 2003.
- 27 Donald A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
- 28 Richard Mayr, Sven Schewe, Patrick Totzke, and Dominik Wojtczak. MDPs with Energy-Parity Objectives. In *Logic in Computer Science (LICS)*. IEEE, 2017. URL: <https://arxiv.org/abs/1701.02546>.
- 29 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. URL: <https://arxiv.org/abs/2101.06989>.
- 30 Joël Ouaknine and James Worrell. On linear recurrence sequences and loop termination. *ACM SIGLOG News*, 2(2):4–13, 2015.
- 31 Jakob Piribauer. *On non-classical stochastic shortest path problems*. PhD thesis, Technische Universität Dresden, Germany, 2021. URL: <https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa2-762812>.
- 32 Jakob Piribauer and Christel Baier. On Skolem-Hardness and Saturation Points in Markov Decision Processes. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *Proc. of ICALP*, volume 168 of *LIPICs*, pages 138:1–138:17, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/opus/volltexte/2020/12545>, doi:10.4230/LIPICs.ICALP.2020.138.
- 33 Jakob Piribauer and Christel Baier. Positivity-hardness results on Markov decision processes. 2023. URL: <https://arxiv.org/abs/2302.13675v1>.
- 34 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *Annual Symposium on Principles of Programming Languages (POPL)*, pages 179–190, 1989.
- 35 Peter J. Ramadge and W. Murray Wonham. Supervisory control of a class of discrete event processes. *SIAM journal on control and optimization*, 25(1):206–230, 1987.
- 36 Lloyd S. Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953.
- 37 W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1-2):135–183, 1998.

A Appendix for Section 4

► **Definition 11.** Given an SSG $\mathcal{G} = (S, (S_{\square}, S_{\diamond}, S_{\circ}), E, P)$ and a finite memory deterministic (FD) strategy $\pi = (M, m_0, \text{upd}, \text{nxt})$ for Minimizer let \mathcal{G}_{π} be the maximizing MDP with state space $M \times S$ obtained by fixing Minimizer's choices according to π . The transition rules \longrightarrow' in the derived MDP \mathcal{G}_{π} are given as follows.

1. If $s \in S_{\square}$ for every $(s, s') \in E$, $m \in M$, we have $(m, s) \longrightarrow' (\text{upd}(m, (s, s')), s')$, i.e., Maximizer determines the successor state and Minimizer updates its memory according to the observed transition.
2. Similarly if $s \in S_{\circ}$ for every $(s, s') \in E$, $m \in M$ we have $(m, s) \longrightarrow' (\text{upd}(m, (s, s')), s')$ and $P((m, s))(\text{upd}(m, (s, s')), s') = P(s)(s')$, i.e., transition probabilities are inherited and Minimizer's memory is updated according to the observed transition.
3. If $s \in S_{\diamond}$ then $(m, s) \longrightarrow' (\text{upd}(m, (s, s')), s')$ where $s' = \text{nxt}(m, s)$, i.e., Minimizer chooses the successor state according to the strategy π and updates its memory accordingly.

The reward of each transition is the same as the reward of the transition in \mathcal{G} from which it is derived. Similarly for the priorities (aka coloring) of the states. The ownership of the vertices (m, s) in \mathcal{G}_{π} is as follows. If $s \in S_{\square}$ then (m, s) belongs to Maximizer. If $s \in S_{\circ}$ then (m, s) is also a chance vertex. If $s \in S_{\diamond}$ then (m, s) also becomes a chance vertex (with exactly one successor), since Minimizer's choice has been fixed.

In the dual case where a FD strategy σ for Maximizer is fixed, we obtain a minimizing MDP \mathcal{G}^{σ} . The construction is the same as above, with the roles of Minimizer and Maximizer swapped.

B Appendix for Section 5

► **Lemma 8.** Consider a maximizing MDP $\mathcal{M} = (S, S_{\square}, S_{\circ}, E, P)$, $\varepsilon > 0$ and the constants c, h from Lemma 7. For rewards in unary and $i \geq N$ we have $\text{val}_{\text{Term}(i) \cup \text{OPAR}}^{\mathcal{M}}(s) - \text{val}_{\text{Loss}}^{\mathcal{M}}(s) \leq \varepsilon$ where $N \stackrel{\text{def}}{=} \max(h, \lceil \log_c(\varepsilon \cdot (1 - c)) \rceil) \in \mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}) \cdot \log(1/\varepsilon))$.

For rewards in binary we have $N \in \mathcal{O}(\exp(\exp(|\mathcal{M}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon))$, i.e., the size of N increases by one exponential.

Proof. By Lemma 6 (Equation (3)) and Lemma 7, we have

$$\text{val}_{\text{Term}(i) \cup \text{OPAR}}^{\mathcal{M}}(s) - \text{val}_{\text{Loss}}^{\mathcal{M}}(s) \leq \sup_{\sigma} \mathcal{P}_{\sigma, s}(\text{Term}(i) \cap \overline{\text{FW}}_1) \leq \frac{c^i}{1 - c}$$

for all $i \geq h$ and $s \in S$. To obtain a bound $N \geq h$ with $\frac{c^N}{1 - c} \leq \varepsilon$, it suffices to choose

$$N \stackrel{\text{def}}{=} \max(h, \lceil \log_c(\varepsilon \cdot (1 - c)) \rceil).$$

We observe that $\log_c(\varepsilon \cdot (1 - c)) = -\ln(\varepsilon \cdot (1 - c)) \cdot (-\ln(c))^{-1}$.

However, $-\ln(c) = -\ln(1 - (1 - c)) \geq (1 - c)$. Thus $\log_c(\varepsilon \cdot (1 - c)) \leq \ln\left(\frac{1}{\varepsilon} \cdot \frac{1}{1 - c}\right) \cdot \frac{1}{1 - c}$.

For rewards in unary, by Lemma 7, we have $1/(1 - c) \in \mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}))$ and h is only $\mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}))$. Thus $N \in \mathcal{O}(\exp(|\mathcal{M}|^{\mathcal{O}(1)}) \cdot \log(1/\varepsilon))$.

Now consider the case where rewards are given in binary. Following the proof of [3, Lemma 3.9], the bounds are derived from the size of solutions of the constructed linear program. While the MDPs in [3] only consider unary rewards from $\{-1, 0, 1\}$, one can extend it to the case where the rewards come from the set $\{-R, \dots, 0, \dots, R\}$ in a natural way. This affects the complexity of the above computed constants and thereby size of N . More

precisely, the proof of Lemma 7 can be split into three steps. Firstly, given an MDP \mathcal{M} construct a new “rising” MDP \mathcal{M}' . Then from this derived \mathcal{M}' , construct a linear program. From the solutions of constructed LP, compute the required c and h . We evaluate the effect of having non-unary rewards in each of these steps.

When rewards are given in unary, the resulting \mathcal{M}' has overall size $|\mathcal{M}'| \leq 10|\mathcal{M}|^4$. More exactly, $|S'| \leq 10 * |S|^3 * (|S| + |E|)$ and similarly for $|E'|$. When the rewards are given in binary, the construction results in an additional R^2 factor. So the resulting \mathcal{M}' is pseudo-polynomially big when compared to \mathcal{M} in our case.

The constructed LP (cf. [3, Fig.1]) has $S' + 2$ variables (z_s for each state, x for the mean payoff and ξ for converting the constraint $x > 0$ to $x \geq \xi$). Moreover all variables can be assumed non-negative. The number of constraints is bounded by $E' + 1$. Furthermore all the constants appearing in the constraints are either constants in the original MDP \mathcal{M} or 1 or 0.

Finally, from an optimal solution of the LP (z_s, x, ξ) one can compute $\exp\left(\frac{-x^2}{2 \cdot (z_{\max} + x + R)^2}\right)$ and to get c , then take a rational over-approximation and also take h as $\lceil z_{\max} \rceil$ where $z_{\max} \stackrel{\text{def}}{=} \max_{s \in S'} z_s - \min_{s \in S'} z_s$. The only difference compared to the unary rewards case here is that the one step change of the submartingale is bounded by $z_{\max} + x + R$ instead of $z_{\max} + x + 1$.

From the complexity point of view, both the construction of the LP and the computation from its optimal solutions aren't affected by changes in the rewards, i.e., the previous bounds for c , h and N in terms of $|\mathcal{M}'|$ still hold. In particular, $c \in \mathcal{O}\left(\exp\left(1/2^{|\mathcal{M}'|^{\mathcal{O}(1)}}\right)\right)$, $h \in \mathcal{O}\left(\exp\left(|\mathcal{M}'|^{\mathcal{O}(1)}\right)\right)$ and thus $N \in \mathcal{O}\left(\exp\left(|\mathcal{M}'|^{\mathcal{O}(1)}\right) \cdot \log(1/\varepsilon)\right)$ by [3, Claim 6].

While previously, \mathcal{M}' is only polynomially larger than \mathcal{M} , introducing binary rewards blows up the construction (cf. [3, Appendix A.2]). As a result we have that $|\mathcal{M}'| \in \mathcal{O}\left(2^{|\mathcal{M}|^{\mathcal{O}(1)}}\right)$. Therefore N can be doubly exponential in the size of the original MDP \mathcal{M} , i.e., $N \in \mathcal{O}\left(\exp(\exp(|\mathcal{M}|^{\mathcal{O}(1)})) \cdot \log(1/\varepsilon)\right)$. ◀

C Appendix for Section 6

► **Definition 12** (Definition of $\mathcal{G}[N]$). *We present formally the definition of the game $\mathcal{G}[N]$, which unfolds the energy level in \mathcal{G} till N*

$$\mathcal{G}[N] = (S[N], (S_{\square}[N], S_{\diamond}[N], S_{\circ}[N]), E[N], P[N])$$

where

1. $S[N] \stackrel{\text{def}}{=} S \times \{0, \dots, N + R\} \uplus \{s_{\text{win}}, s_{\text{lose}}\}$, the set of states is the tuple with the game state and energy level until $N + R$ as the maximum change in a single step is R and since we are only interested in energy levels $\leq N$, it suffices to consider till $N + R$.
2. $S_{\circ}[N] \stackrel{\text{def}}{=} S_{\circ} \times \{1, \dots, N\}$, both players control their respective states until energy level N . Every state with energy $> N$ becomes a chance node. Consequently,
3. $S_{\circ}[N] \stackrel{\text{def}}{=} S_{\circ} \times \{1, \dots, N\} \cup S \times \{0, N + 1, \dots, N + R\} \cup \{s_{\text{win}}, s_{\text{lose}}\}$, since the Maximizer loses when the energy level becomes ≤ 0 , we make these states as a chance vertex which go to a losing loop.
4. $E[N], P[N]$
 - a. For $0 < i \leq N$, $(s, i) \longrightarrow (s', \max(0, j))$ iff $s \xrightarrow{j-i} s' \in E$, this is just simulating the transitions of the game until energy level N and taking care of border cases. When energy drops below 0, we move to level 0 as there is no difference. When it shoots above N , it cannot go beyond $N + R$ and thus the transition is well defined.
 - b. If $s \in S_{\circ}$ above, then the probability is carried over.

36:18 Approximating the Value of Energy-Parity Games

- c. $(s, 0) \longrightarrow s_{\text{lose}}$ with probability 1.
- d. $(s, N + k) \longrightarrow s_{\text{win}}$ with probability $\text{val}_{\text{EN}(N+k) \cap \text{EPAR}}^G(s)$ and with remaining probability moves to s_{lose} for $1 \leq k \leq R$
- e. $s_{\text{lose}} \longrightarrow s_{\text{lose}}$ with probability 1. Similarly for s_{win} .