Parity Games with Weights

Sven Schewe[1]*, Alexander Weinert[2]**, and Martin Zimmermann[2]***

[1] University of Liverpool, Liverpool L69 3BX, United Kingdom
sven.schewe@liverpool.ac.uk
[2] Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany
{weinert,zimmermann}@react.uni-saarland.de

Abstract. Quantitative extensions of parity games have recently attracted significant interest. These extensions include parity games with energy and payoff conditions as well as finitary parity games and their generalization to parity games with costs. Finitary parity games enjoy a special status among these extensions, as they offer a native combination of the qualitative and quantitative aspects in infinite games: the quantitative aspect of finitary parity games is a quality measure for the qualitative aspect, as it measures the limit superior of the time it takes to answer an odd color by a larger even one. Finitary parity games have been extended to parity games with costs, where each transition is labelled with a non-negative weight that reflects the costs incurred by taking it. We lift this restriction and consider parity games with costs with arbitrary integer weights. We show that solving such games is in NP∩co-NP, the signature complexity for games of this type. We also show that the protagonist has finite-state winning strategies, and provide tight exponential bounds for the memory he needs to win the game. Naturally, the antagonist may need need infinite memory to win. Finally, we present tight bounds on the quality of winning strategies for the protagonist.

1 Introduction

Finite games of infinite duration offer a wealth of challenges and applications that has garnered to a lot of attention. The traditional class of games under consideration were games with a simple parity [18, 11, 10, 21, 2, 29, 14, 15, 27, 17, 24, 26, 25, 3, 16, 12, 19] or payoff [23, 30, 14, 1, 26] objective. These games form a hierarchy with very simple tractable reductions from parity games through mean payoff games [23, 30, 14, 1, 26] and discounted payoff games [30, 14, 26] to simple stochastic games [9].

More recently, games with a mixture of the qualitative parity condition and further quantitative objectives have been considered, including mean payoff parity games [8] and energy parity games [4]. Finitary parity games [7] take a special role within the class of games with mixed parity and payoff objectives. To win a finitary parity game, Player 0 needs to enforce a play with a bound b such that almost all occurrences of an odd color are followed by a higher even color within at most b steps.

This is interesting, because it provides a natural link between the qualitative and quantitative objective. One aspect that attracted attention is that, as long as one is not interested in optimizing the bound b, these games are the only games of the lot that are known to be tractable [7]. However, the bound b itself is also interesting: It serves as a native quality measure, because it limits the response time [28].

This property calls for a generalization to different cost models, and a first generalization has been made with the introduction of parity games with costs [13]. In parity games with costs, the basic cost function of finitary parity games—where each step incurs the same cost—is replaced with different non-negative costs for different edges. In this paper, we generalize this further to *general* integer costs: We decorate the edges with integer weights. The quantitative aspect in these parity games with weights consists of having to answer almost all odd colors by a higher even color, such that the *absolute* value of the weight of the path to this even color is bounded by a bound b.

In addition to their conceptual charm, we show that parity games with weights are PTime equivalent to energy parity games. This indicates that these games are part of a natural complexity class, whereas

* Supported by the EPSRC projects 'Energy Efficient Control' (EP/M027287/1) and 'Solving Parity Games in Theory and Practice' (EP/P020909/1).
** Supported by the project "TriCS" (ZI 1516/1-1) of the German Research Foundation (DFG) and the Saarbrücken Graduate School of Computer Science.
*** Supported by the project "TriCS" (ZI 1516/1-1) of the German Research Foundation (DFG).

the games with a plain objective appear to form a hierarchy. We use the reduction from parity games with weights to energy parity games to solve them. This reduction goes through intermediate reductions to and from *bounded* parity games with weights. These games have the additional restriction that the limes superior of the absolute weight of initial sequences of unanswered requests in a play is finite. These bounded parity games with weights are then reduced to energy parity games. The other direction of the reduction is through simple gadgets that preserve the main elements of winning strategies in games that are extended in two steps by very simple gadgets. As a result, we obtain the same complexity results for parity games with weights as for energy parity games, i.e., NP∩co-NP, the signature complexity for finite games of infinite duration with parity conditions and their extensions. Thereby, we obtain an argument that these games might be representatives of a natural complexity class, lending a further argument for the relevance of two player games with mixed qualitative and quantitative winning conditions.

Naturally, parity games with weights subsume parity games (as a special case where all weights are zero), finitary parity games (as a special case where all weights are positive), and parity games with costs (as a special case where all weights are non-negative).

Finally, we show that the protagonist has finite-state winning strategies, and provide tight exponential bounds for the memory he needs to win the game. We also present tight bounds on the quality of winning strategies for the protagonist. Naturally, the antagonist may need infinite memory to win.

2 Preliminaries

We denote the non-negative integers by $\mathbb{N}$, the integers by $\mathbb{Z}$, and define $\mathbb{N}_\infty = \mathbb{N} \cup \{\infty\}$. As usual, we have $\infty > n$, $-\infty < n$, $n + \infty = \infty$, and $-\infty - n = -\infty$ for all $n \in \mathbb{Z}$.

An **arena** $\mathcal{A} = (V, V_0, V_1, E)$ consists of a finite, directed graph (V, E) and a partition $\{V_0, V_1\}$ of V into the positions of Player 0 (drawn as ellipses) and Player 1 (drawn as rectangles). The size of $\mathcal{A}$, denoted by $|\mathcal{A}|$, is defined as $|V|$. A **play** in $\mathcal{A}$ is an infinite path $\rho = v_0 v_1 v_2 \cdots$ through (V, E). To rule out finite plays, we require every vertex to be non-terminal. We define $|\rho| = \infty$.

A **game** $\mathcal{G} = (\mathcal{A}, \mathrm{Win})$ consists of an arena $\mathcal{A}$ with vertex set V and a set $\mathrm{Win} \subseteq V^\omega$ of winning plays for Player 0. The set of winning plays for Player 1 is $V^\omega \setminus \mathrm{Win}$. A winning condition Win is 0-extendable if, for all $\rho \in V^\omega$ and all $w \in V^*$, $\rho \in \mathrm{Win}$ implies $w\rho \in \mathrm{Win}$. Dually, Win is 1-extendable if, for all $\rho \in V^\omega$ and all $w \in V^*$, $\rho \notin \mathrm{Win}$ implies $w\rho \notin \mathrm{Win}$.

A **strategy** for Player $i \in \{0, 1\}$ is a mapping $\sigma \colon V^* V_i \to V$ such that $(v, \sigma(wv)) \in E$ holds true for all $wv \in V^* V_i$. We say that σ is **positional** if $\sigma(wv) = \sigma(v)$ holds true for every $wv \in V^* V_i$. A play $v_0 v_1 v_2 \cdots$ is **consistent** with a strategy σ for Player i, if $v_{j+1} = \sigma(v_0 \cdots v_j)$ holds true for every j with $v_j \in V_i$. A strategy σ for Player i is a **winning strategy** for $\mathcal{G}$ from $v \in V$ if every play that starts in v and is consistent with σ is won by Player i. If Player i has a winning strategy from v, then we say Player i wins $\mathcal{G}$ from v. The **winning region** of Player i is the set of vertices, from which Player i wins $\mathcal{G}$; it is denoted by $\mathcal{W}_i(\mathcal{G})$. **Solving** a game amounts to determining its winning regions. If $\mathcal{W}_0(\mathcal{G}) \cup \mathcal{W}_1(\mathcal{G}) = V$, then we say that $\mathcal{G}$ is **determined**.

Let $\mathcal{A} = (V, V_0, V_1, E)$ be an arena and let $X \subseteq V$. The i-attractor of X is defined inductively as $\mathrm{Attr}_i(X) = \mathrm{Attr}_i^{|V|}(X)$, where $\mathrm{Attr}_i^0(X) = X$ and

$$\mathrm{Attr}_i^j(X) = \mathrm{Attr}_i^{j-1}(X) \cup \{v \in V_i \mid \exists v' \in \mathrm{Attr}_i^{j-1}(X).\, (v, v') \in E\} \cup \{v \in V_{1-i} \mid \forall (v, v') \in E.\, v' \in \mathrm{Attr}_i^{j-1}(X)\} \ .$$

Hence, $\mathrm{Attr}_i(X)$ is the set of vertices from which Player i can force the play to enter X: Player i has a positional strategy σ_X such that each play that starts in some vertex in $\mathrm{Attr}_i(X)$ and is consistent with σ_X eventually encounters some vertex from X. We call σ_X an attractor strategy towards X. Moreover, the i-attractor can be computed in time linear in $|E|$ [22]. When we want to stress the arena $\mathcal{A}$ the attractor is computed in, we write $\mathrm{Attr}_i^{\mathcal{A}}(X)$.

A set $X \subseteq V$ is a trap for Player i, if every vertex in $X \cap V_i$ has only successors in X and every vertex in $X \cap V_{1-i}$ has at least one successor in X. In this case, Player $1-i$ has a positional strategy τ_X such that every play starting in some vertex in X and consistent with τ_X never leaves X. We call such a strategy a trap strategy.

Remark 1.

1. The complement of an i-attractor is a trap for Player i.
2. If X is a trap for Player i, then $\mathrm{Attr}_{1-i}(X)$ is also a trap for Player i.
3. If Win is i-extendable and $(\mathcal{A}, \mathrm{Win})$ determined, then $\mathcal{W}_{1-i}(\mathcal{A}, \mathrm{Win})$ is a trap for Player i.

A **memory structure** $\mathcal{M} = (M, \mathrm{init}, \mathrm{upd})$ for an arena (V, V_0, V_1, E) consists of a finite set M of memory states, an initialization function $\mathrm{init}\colon V \to M$, and an update function $\mathrm{upd}\colon M \times E \to M$. The update function can be extended to finite play prefixes in the usual way: $\mathrm{upd}^+(v) = \mathrm{init}(v)$ and $\mathrm{upd}^+(wvv') = \mathrm{upd}(\mathrm{upd}^+(wv), (v, v'))$ for $w \in V^*$ and $(v, v') \in E$. A next-move function $\mathrm{Nxt}\colon V_i \times M \to V$ for Player i has to satisfy $(v, \mathrm{Nxt}(v, m)) \in E$ for all $v \in V_i$ and $m \in M$. It induces a strategy σ for Player i with memory $\mathcal{M}$ via $\sigma(v_0 \cdots v_j) = \mathrm{Nxt}(v_j, \mathrm{upd}^+(v_0 \cdots v_j))$. A strategy is called **finite-state** if it can be implemented by a memory structure. We define $|\mathcal{M}| = |M|$. Slightly abusively, we say that the size of a finite-state strategy is the size of a memory structure implementing it.

3 Parity Games with Weights

Fix an arena $\mathcal{A} = (V, V_0, V_1, E)$. A **weighting** for $\mathcal{A}$ is a function $w\colon E \to \mathbb{Z}$. We define $w(\varepsilon) = w(v) = 0$ for all $v \in V$ and extend w to sequences of vertices of length at least two by summing up the weights of the traversed edges. Given a play (prefix) $\pi = v_0 v_1 v_2 \cdots$, we define the amplitude of π as $\mathrm{Ampl}(\pi) = \sup_{j<|\pi|} |w(v_0 \cdots v_j)| \in \mathbb{N}_\infty$.

A **coloring** of V is a function $\Omega\colon V \to \mathbb{N}$. The classical parity condition requires almost all occurrences of odd colors to be answered by a later occurrence of a larger even color. Hence, let $\mathrm{Ans}(c) = \{c' \in \mathbb{N} \mid c' \geq c \text{ and } c' \text{ is even}\}$ be the set of colors that "answer" a "request" for color c. We denote a vertex v of color c by v/c.

Fijalkow and Zimmermann introduced a generalization of the parity condition and the finitary parity condition [7], the parity condition with costs [13]. There, the edges of the arena are labeled with *non-negative weights* and the winning condition demands that there exists a bound b such that almost all requests are answered with weight at most b, i.e., the weight of the infix between the request and the response has to be bounded by b.

Our aim is to extend the parity condition with costs by allowing for the full spectrum of weights to be used, i.e., by also incorporating negative weights. In this setting, the weight of an infix between a request and a response might be negative. Thus, the extended condition requires the weight of the infix to be bounded from above and from below.[3] To distinguish between the parity condition with costs and the extension introduced here, we call our extension the parity condition with weights.

Formally, let $\rho = v_0 v_1 v_2 \cdots$ be a play. We define the cost-of-response at position $j \in \mathbb{N}$ of ρ by

$$\mathrm{Cor}(\rho, j) = \min\{\mathrm{Ampl}(v_j \cdots v_{j'}) \mid j' \geq j, \Omega(v_{j'}) \in \mathrm{Ans}(\Omega(v_j))\}$$

where we use $\min \emptyset = \infty$. As the amplitude of an infix only increases by extending the infix, $\mathrm{Cor}(\rho, j)$ is the amplitude of the shortest infix that starts at position j and ends at an answer to the request posed at position j. We illustrate this notion in Figure 1.

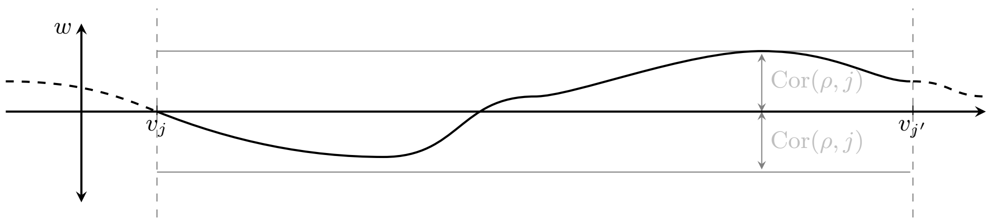

Fig. 1. The cost-of-response of some request posed by visiting vertex v_j, which is answered by visiting vertex $v_{j'}$.

We say that a request at position j is answered with cost b, if $\mathrm{Cor}(\rho, j) = b$. Consequently, a request with an even color is answered with cost zero. The cost-of-response of an unanswered request is infinite,

[3] We discuss other possible interpretations of negative weights in Section 9.

even if the amplitude of the remaining play is bounded. In particular, this means that an unanswered request at position j may be "unanswered with finite cost b" (if the amplitude of the remaining play is $b \in \mathbb{N}$) or "unanswered with infinite cost" (if the amplitude of the remaining play is infinite). In either case, however, we have $\mathrm{Cor}(\rho, j) = \infty$.

We define the parity condition with weights as

$$\mathrm{WeightParity}(\Omega, w) = \{\rho \in V^\omega \mid \limsup_{j\to\infty} \mathrm{Cor}(\rho, j) \in \mathbb{N}\} \ .$$

I.e., ρ satisfies the condition if and only if there exists a bound $b \in \mathbb{N}$ such that almost all requests are answered with cost less than b. In particular, only finitely many requests may be unanswered, even with finite cost. Note that the bound b may depend on the play ρ.

We call a game $\mathcal{G} = (\mathcal{A}, \mathrm{WeightParity}(\Omega, w))$ a parity game with weights, and we define $|\mathcal{G}| = |\mathcal{A}| + \log(W)$, where W is the largest absolute weight assigned by w; i.e., we assume weights to be encoded in binary. If w assigns zero to every edge, then $\mathrm{WeightParity}(\Omega, w)$ is a classical (max-) parity condition, denoted by $\mathrm{Parity}(\Omega)$. Similarly, if w assigns positive weights to every edge, then $\mathrm{WeightParity}(\Omega, w)$ is equal to the finitary parity condition over Ω, as introduced by Chatterjee and Henzinger [6]. Finally, if w assigns only non-negative weights, then $\mathrm{WeightParity}(\Omega, w)$ is a parity condition with costs, as introduced by Fijalkow and Zimmermann [13]. In these cases, we refer to $\mathcal{G}$ as a parity game, a finitary parity game, or a parity game with costs, respectively. We recall the characteristics of these games in Table 1.

4 Solving Parity Games with Weights

We now show how to solve parity games with weights. Our approach is inspired by the classic work on finitary parity games [7] and parity games with costs [13]: We first define a stricter variant of these games, which we call bounded parity games with weights, and then show two reductions:

- parity games with weights can be solved in polynomial time with oracles that solve bounded parity games with weights (in this section); and
- bounded parity games with weights can be solved in polynomial time with oracles that solve energy parity games (Section 5).

Furthermore, in Section 8 we polynomially reduce solving energy parity games to solving parity games with weights and thereby show that parity games with weights, bounded parity games with weights, and energy parity games belong to the same complexity class.

The energy parity games that we reduce to are known to be efficiently solvable [4]: they are in $\mathrm{NP} \cap \mathrm{CO}\text{-}\mathrm{NP}$ and can be solved in pseudo-polynomial time for a fixed number of colors.

We first introduce the **bounded parity** condition **with weights**, which is a strengthening of the parity condition with weights. Hence, it is also induced by a coloring and a weighting:

$$\begin{aligned}\mathrm{BndWeightParity}(\Omega, w) = {} & \mathrm{WeightParity}(\Omega, w) \\ & \cap \{\rho \in V^\omega \mid \text{no request in } \rho \text{ is unanswered with infinite cost}\} \ .\end{aligned}$$

Note that this condition allows for a finite number of unanswered requests, as long as they are unanswered with finite cost.

We solve parity games with weights by repeatedly solving bounded parity games with weights. To this end, we apply the following two properties of the winning conditions: We have $\mathrm{BndWeightParity}(\Omega, w) \subseteq \mathrm{WeightParity}(\Omega, w)$ as well as that $\mathrm{WeightParity}(\Omega, w)$ is 0-extendable. Hence, if Player 0 has a strategy from a vertex v such that every consistent play has a suffix in $\mathrm{BndWeightParity}(\Omega, w)$, then the strategy is winning for her from v w.r.t. $\mathrm{WeightParity}(\Omega, w)$. Thus, $\mathrm{Attr}_0(\mathcal{W}_0(\mathcal{A}, \mathrm{BndWeightParity}(\Omega w))) \subseteq \mathcal{W}_0(\mathcal{A}, \mathrm{WeightParity}(\Omega, w))$. The algorithm that solves parity games with weights repeatedly removes attractors of winning regions of the bounded parity game with weights until a fixed point is reached. We will later formalize this sketch to show that the removed parts are a subset of Player 0's winning region in the parity game with weights.

To show that the obtained fixed point covers the complete winning region of Player 0, we use the following lemma to show that the remaining vertices are a subset of Player 1's winning region in the parity game with weights. The proof is very similar to the corresponding one for finitary parity games and parity games with costs.

Lemma 1. *Let* $\mathcal{G} = (\mathcal{A}, \mathrm{WeightParity}(\Omega, w))$ *and let* $\mathcal{G}' = (\mathcal{A}, \mathrm{BndWeightParity}(\Omega, w))$*. If* $\mathcal{W}_0(\mathcal{G}') = \emptyset$*, then* $\mathcal{W}_0(\mathcal{G}) = \emptyset$*.*

Proof. As bounded parity conditions with weights are Borel, bounded parity games with weights are determined [20]. Hence, $\mathcal{W}_0(\mathcal{G}') = \emptyset$ implies that, for every vertex v of $\mathcal{A}$, Player 1 has a strategy τ_v that is winning in $\mathcal{G}'$ from v.

We combine these strategies into a single strategy τ for Player 1 that is winning in $\mathcal{G}$ from every vertex of $\mathcal{A}$. This strategy is controlled by a vertex v^* (initialized with the starting vertex of the play) and a counter κ ranging over $\mathbb{N}$ (initialized with zero). The strategy τ mimics the strategy τ_{v^*} from v^* until a request is followed by an infix without an answer and with amplitude κ. This implies that the cost-of-response of this request is at least κ. If such a situation is encountered, then v^* is set to the current vertex and κ is incremented. Furthermore, the history of the play is discarded at this point in the play, and τ behaves henceforth like τ_{v^*} when starting at v^* when this happens.

Consider a play ρ that is consistent with this strategy. If, on the one hand, κ is updated infinitely often along ρ, then ρ contains, for every $b \in \mathbb{N}$, a request that has a cost-of-response that is larger than b. Hence, it violates the parity condition with weights.

If, on the other hand, κ is only updated finitely often, then ρ has a suffix ρ' that starts in some v, which is consistent with τ_v. As τ_v is winning for Player 1 from v in $\mathcal{G}'$, ρ' violates the bounded parity condition with weights. Also, because κ is updated only finitely often during the suffix, there is a bound b such that the amplitude of every suffix of ρ' that starts at a request is bounded by b. Hence, the only way for ρ' to violate the bounded parity condition with weights is to violate the parity condition. Thus, the full play ρ also violates the parity condition, and therefore also the parity condition with weights, which is a strengthening of the parity condition. Therefore, τ is indeed winning for Player 1 from every vertex in $\mathcal{G}$. ◻

Lemma 1 implies that the algorithm for solving parity games with weights by repeatedly solving bounded parity games with weights (see Algorithm 1) is correct. Note that we use an oracle for solving bounded parity games with weights. We provide a suitable algorithm in Section 5.

Algorithm 1 A fixed-point algorithm computing $\mathcal{W}_0(\mathcal{A}, \mathrm{WeightParity}(\Omega, w))$.

```
k = 0; W_0^k = ∅; A_k = A
repeat
    k = k + 1
    X_k = W_0(A_{k-1}, BndWeightParity(Ω, w))
    W_0^k = W_0^{k-1} ∪ Attr_0^{A_{k-1}}(X_k)
    A_k = A_{k-1} \ Attr_0^{A_{k-1}}(X_k)
until X_k = ∅
return W_0^k
```

The loop terminates after at most $|\mathcal{A}|$ iterations (assuming the algorithm solving bounded parity games with weights terminates), as during each iteration at least one vertex is removed from the arena. The correctness proof relies on Lemma 1 and is similar to the one for finitary parity games [7] and for parity games with costs [13].

Lemma 2. *Algorithm 1 returns* $\mathcal{W}_0(\mathcal{A}, \mathrm{WeightParity}(\Omega, w))$

Proof. Let $\mathcal{G} = (\mathcal{A}, \mathrm{WeightParity}(\Omega, w))$ and let k^* be the final iteration when running the algorithm on $\mathcal{G}$, i.e., its output is $W_0^{k^*} = \bigcup_{0<k'<k} \mathrm{Attr}_0^{\mathcal{A}_{k'-1}}(X_{k'})$.

First, we consider Player 0 and show $W_0^{k^*} \subseteq \mathcal{W}_0(\mathcal{G})$. For every vertex v that is in some X_k, Player 0 has a strategy σ_v for $\mathcal{G}_k = (\mathcal{A}_{k-1}, \mathrm{BndWeightParity}(\Omega, w)))$ that is winning from v. Furthermore, for every attractor $\mathrm{Attr}_0^{\mathcal{A}_{k-1}}(X_k)$ he has a positional attractor strategy σ_k. Now, we compose these strategies to a strategy σ for Player 0 in $\mathcal{A}$ via

$$\sigma(v_0 \cdots v_j) = \begin{cases} \sigma_k(v_j) & \text{if } v_j \in \mathrm{Attr}_0^{\mathcal{A}_{k-1}}(X_k) \setminus X_k, \\ \sigma_{v_{j'}}(v_{j'} \cdots v_j) & \text{if } v_j \in X_k. \end{cases}$$

In the second case, $v_{j'} \cdots v_j$ is the longest suffix of $v_0 \cdots v_j$ that only contains vertices from X_k, the set of vertices from which Player 0 has a winning strategy for $\mathcal{G}_k$.

Consider a play $\rho = v_0 v_1 v_2 \cdots$ in $\mathcal{A}$ that starts in $W_0^{k^*}$ and is consistent with σ. For every j there is a unique k_j in the range $0 < k_j < k^*$ such that $v_j \in \mathrm{Attr}_0^{\mathcal{A}_{k_j - 1}}(X_{k_j})$. As BndWeightParity$(\Omega, w)$ is 1-extendable, Items 2 and 3 of Remark 1 imply that each $\mathrm{Attr}_0^{\mathcal{A}_{k-1}}(X_k)$ is a trap for Player 1 in $\mathcal{A}_{k-1}$. Hence, we obtain $k_0 > k_1 > k_2 > \cdots$. As the k_j are always greater than zero, the sequence has to stabilize eventually. This implies that ρ has a suffix $\rho' = v_j v_{j+1} v_{j+2} \cdots$ that is consistent with σ_{v_j}.

Hence, due to σ_{v_j} being a winning strategy for Player 0 in $\mathcal{G}_k$ from v_j, we obtain $\rho' \in$ BndWeightParity(Ω, w). Hence, $\rho \in$ WeightParity(Ω, w) due to 0-extendability of WeightParity(Ω, w). Hence, σ is indeed winning from $W_0^{k^*}$.

Now, consider Player 1. We show $V \setminus W_0^{k^*} \subseteq \mathcal{W}_1(\mathcal{G})$. Then, determinacy of parity games with weights (due to their winning conditions being Borel [20]) yields $W_0^{k^*} = \mathcal{W}_0(\mathcal{G})$ and $V \setminus W_0^{k^*} = \mathcal{W}_1(\mathcal{G})$.

Due to X_{k^*} being empty and bounded parity games with weights being determined (again due to their winning conditions being Borel), Player 1 wins the bounded parity game with weights $\mathcal{G}_{k^*}$ from every vertex. Applying Lemma 1 shows that she also wins the parity game with weights $(\mathcal{A}_{k^*-1}, \text{WeightParity}(\Omega, w))$ from every vertex. Finally, as $V \setminus W_0^{k^*} = \mathcal{W}_1(\mathcal{A}_{k^*-1}, \text{WeightParity}(\Omega, w))$ is a trap for Player 0 in $\mathcal{A}$ by construction, he also wins $\mathcal{G} = (\mathcal{A}, \text{WeightParity}(\Omega, w))$ from every vertex in $V \setminus W_0^{k^*}$. □

The strategy σ defined in the proof of the first item can be implemented by a finite-state strategy of size $\max_{k \leq k^*} s_k$, assuming that the constituent strategies σ_k are finite-state strategies of size s_k. To this end, one uses the fact that the winning regions X_k are disjoint and are never revisited once left. Hence, we can assume the implementations of the σ_k to use the same states.

5 Solving Bounded Parity Games with Weights

After having reduced the problem of solving parity games with weights to that of solving (multiple) bounded parity games with weights, we reduce solving bounded parity games with weights to solving (multiple) energy parity games [4].

Similarly to a parity game with weights, in an energy parity game, the vertices are colored and the edges are equipped with weights. It is the goal of Player 0 to satisfy the parity condition, while, at the same time, ensuring that the weight of every infix, its so-called energy level, is bounded from below. In contrast to a parity game with weights, however, the weights in an energy parity game are not tied to the requests and responses denoted by the coloring.

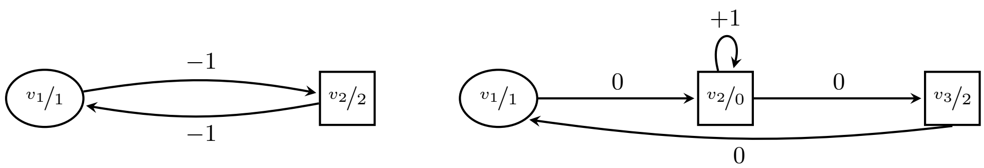

Fig. 2. The difference between energy parity games and parity games with weights.

Consider, for example, the games shown in Figure 2. In the game on the left-hand side, players only have a single, trivial strategy. If we interpret this game as a parity game with weights, Player 0 wins from every vertex, as each request is answered with cost one. If we, however, interpret that game as an energy parity game, Player 1 instead wins from every vertex, since the energy level decreases by one with every move. In the game on the right-hand side, the situation is mirrored: When interpreting this game as a parity game with weights, Player 1 wins from every vertex, as she can easily unbound the costs of the requests for color one by staying in vertex v_2 for an ever-increasing number of cycles. Dually, when interpreting this game as an energy parity game, Player 0 wins from every vertex, since the parity condition is clearly satisfied in every play, and Player 1 is only able to increase the energy level, while it is never decreased.

In Section 5.1, we introduce energy parity games formally and present how to solve bounded parity games with weights via energy games in Section 5.2. Finally, Sections 5.3 and 5.4 are dedicated to the correctness proof of the construction.

5.1 Energy Parity Games

An energy parity game $\mathcal{G} = (\mathcal{A}, \Omega, w)$ consists of an arena $\mathcal{A} = (V, V_0, V_1, E)$, a coloring $\Omega\colon V \to \mathbb{N}$ of V, and an edge weighting $w\colon E \to \mathbb{Z}$ of E. Note that this definition is not compatible with the framework presented in Section 2, as we have not (yet) defined the winner of the plays. This is because they depend on an initial credit, which is existentially quantified in the definition of winning the game $\mathcal{G}$. Formally, the set of winning plays with initial credit $c_0 \in \mathbb{N}$ is defined as

$$\text{EnergyParity}_{c_0}(\Omega, w) = \text{Parity}(\Omega) \cap \{v_0v_1v_2 \cdots \in V^\omega \mid \forall j \in \mathbb{N}.\, c_0 + w(v_0 \cdots v_j) \geq 0\} \ .$$

Now, we say that Player 0 wins $\mathcal{G}$ from v if there exists some initial credit $c_0 \in \mathbb{N}$ such that he wins $\mathcal{G}_{c_0} = (\mathcal{A}, \text{EnergyParity}_{c_0}(\Omega, w))$ from v (in the sense of the definitions in Section 2). If this is not the case, i.e., if Player 1 wins $\mathcal{G}_{c_0}$ from v for every c_0, then we say that Player 1 wins $\mathcal{G}$ from v. Note that the initial credit is uniform for all plays, unlike the bound on the cost-of-response in the definition of the parity condition with weights, which may depend on the play.

Unravelling these definitions shows that Player 0 wins $\mathcal{G}$ from v if there is an initial credit c_0 and a strategy σ, such that every play that starts in v and is consistent with σ satisfies the parity condition *and* the accumulated weight over the play prefixes (the energy level) never drops below $-c_0$. We call such a strategy σ a winning strategy for Player 0 in $\mathcal{G}$ from v. Dually, Player 1 wins $\mathcal{G}$ from v if, for every initial credit c_0, there is a strategy τ_{c_0}, such that every play that starts in v and is consistent with τ_{c_0} violates the parity condition *or* its energy level drops below $-c_0$ at least once. Thus, the strategy τ_{c_0} may, as the notation suggests, depend on c_0. However, Chatterjee and Doyen showed that using different strategies is not necessary: There is a uniform strategy τ that is winning from v for every initial credit c_0.

Proposition 1 ([4]). *Let $\mathcal{G}$ be an energy parity game. If Player 1 wins $\mathcal{G}$ from v, then she has a single positional strategy that is winning from v in $\mathcal{G}_{c_0}$ for every c_0.*

We call such a strategy as in Proposition 1 a winning strategy for Player 1 from v. A play consistent with such a strategy either violates the parity condition, or the energy levels of its prefixes diverge towards $-\infty$.

Furthermore, Chatterjee and Doyen obtained an upper bound on the initial credit necessary for Player 0 to win an energy parity game, as well an upper bound on the size of a corresponding finite-state winning strategy.

Proposition 2 ([4]). *Let $\mathcal{G}$ be an energy parity game with n vertices, d colors, and largest absolute weight W. The following are equivalent for a vertex v of $\mathcal{G}$:*

1. *Player 0 wins $\mathcal{G}$ from v.*
2. *Player 0 wins $\mathcal{G}_{(n-1)W}$ from v with a finite-state strategy with at most ndW states.*

The previous proposition yields that finite-state strategies of bounded size suffice for Player 0 to win.

Such strategies do not admit long expensive descents, which we show by a straightforward pumping argument.

Lemma 3. *Let $\mathcal{G}$ be an energy parity game with n vertices and largest absolute weight W. Further, let σ be a finite-state strategy of size s, and let ρ be a play that starts in some vertex, from which σ is winning, and is consistent with σ. Every infix π of ρ satisfies $w(\pi) > -Wns$.*

Proof. Let σ be implemented by $\mathcal{M} = (M, \text{init}, \text{upd})$ and let $\rho = v_0v_1v_2 \cdots$. We assume towards a contradiction that there is an infix $\pi = v_j \cdots v_{j'}$ with $w(\pi) \leq -Wns$.

We obtain a lower bound of $ns + 1$ on the number of different energy levels attained during π (as W is the maximal absolute value of the occurring weights). Moreover, we obtain $ns + 1$ non-empty prefixes of π with (1) increasing length and (2) strictly decreasing energy levels.

Thus, there are positions j_0, j_1 with $j \leq j_0 < j_1 \leq j'$ with $v_{j_0} = v_{j_1}$, $\text{upd}^+(v_0 \cdots v_{j_0}) = \text{upd}^+(v_0 \cdots v_{j_1})$, and $w(v_{j_0} \cdots v_{j_1}) < 0$. Hence, the play $v_0 \cdots v_{j_0 - 1}(v_{j_0} \cdots v_{j_1 - 1})^\omega$ obtained by repeating the loop between

v_{j_0} and v_{j_1} ad infinitum is consistent with σ and violates the energy condition. However, it starts in a vertex, from which σ is a winning strategy for Player 0 in the energy parity game. This yields the desired contradiction. □

Moreover, Chatterjee and Doyen gave an upper bound on the complexity of solving energy parity games.

Proposition 3 ([4]). *The following decision problem is in* NP∩co-NP*: "Given an energy parity game $\mathcal{G}$ and a vertex v in $\mathcal{G}$, does Player* 0 *win $\mathcal{G}$ from v?"*

5.2 From Bounded Parity Games with Weights to Energy Parity Games

Let $\mathcal{G} = (\mathcal{A}, \mathrm{BndWeightParity}(\Omega, w))$ be a bounded parity game with weights with vertex set V. Without loss of generality, we assume $\Omega(v) \geq 2$ for all $v \in V$. We construct, for each vertex v^* of $\mathcal{A}$, an energy parity game $\mathcal{G}_{v^*}$ with the following property: Player 1 wins $\mathcal{G}_{v^*}$ from some designated vertex induced by v^* if and only if she is able to unbound the amplitude for the request of the initial vertex of the play when starting from v^*. This construction is the technical core of the fixed-point algorithm that solves bounded parity games with weights via solving energy parity games.

The main obstacle towards this is that, in the bounded parity game with weights $\mathcal{G}$, Player 1 may win by unbounding the amplitude for a request from above or from below, while she can only win $\mathcal{G}_{v^*}$ by unbounding the costs from below. We model this in $\mathcal{G}_{v^*}$ by constructing two copies of $\mathcal{A}$. In one of these copies the edge weights are copied from $\mathcal{G}$, while they are inverted in the other copy. We allow Player 1 to switch between these copies arbitrarily. To compensate for Player 1's power to switch, Player 0 can increase the energy level in the resulting energy parity game during each switch.

First, we define the set of polarities $P = \{+, -\}$ as well as $\overline{+} = -$ and $\overline{-} = +$. Given a vertex v^* of $\mathcal{A}$, define the "polarized" arena $\mathcal{A}_{v^*} = (V', V'_0, V'_1, E')$ of $\mathcal{A} = (V, V_0, V_1, E)$ with

- $V' = (V \times P) \cup (E \times P \times \{0, 1\})$,
- $V'_i = (V_i \times P) \cup (E \times P \times \{i\})$ for $i \in \{0, 1\}$, and
- E' contains the following edges for every edge $e = (v, v') \in E$ with $v \notin \mathrm{Ans}(\Omega(v^*))$ and every polarity $p \in P$:
  - $((v, p), (e, p, 1))$: The player whose turn it is at v picks a successor v'. The edge $e = (v, v')$ is stored as well as the polarity p.
  - $((e, p, 1), (v', p))$: Then, Player 1 can either keep the polarity p unchanged and execute the move to v', or
  - $((e, p, 1), (e, p, 0))$: she decides to change the polarity, and another auxiliary vertex is reached.
  - $((e, p, 0), (e, p, 0))$: If the polarity is to be changed, then Player 0 is able to use a self-loop to increase the energy level (see below), before
  - $((e, p, 0), (v', \overline{p}))$: he can eventually complete the polarity switch by moving to v'.
- Furthermore, for every vertex v with $\Omega(v) \in \mathrm{Ans}(\Omega(v^*))$ and every polarity $p \in P$, E' contains the self-loop $((v, p), (v, p))$.[4]

Thus, a play in $\mathcal{A}_{v^*}$ simulates a play in $\mathcal{A}$, unless Player 0 stops the simulation by using the self-loop at a vertex of the form $(e, p, 0)$ ad infinitum, and unless an answer to $\Omega(v^*)$ is reached. We define the coloring and the weighting for $\mathcal{A}_{v^*}$ so that Player 0 loses in the former case and wins in the latter case. Furthermore, the coloring is defined so that all simulating plays that are not stopped have the same color sequence as the simulated play (save for irrelevant colors on the auxiliary vertices in $E \times P \times \{0, 1\}$). Hence, we define

$$\Omega_{v^*}(v) = \begin{cases} \Omega(v') & \text{if } v = (v', p) \text{ with } v' \notin \mathrm{Ans}(\Omega(v^*)) \ , \\ 0 & \text{if } v = (v', p) \text{ with } v' \in \mathrm{Ans}(\Omega(v^*)) \ , \\ 1 & \text{otherwise} \ . \end{cases}$$

[4] Note that this definition introduces some terminal vertices, i.e., those of the form $((v, v'), p, i)$ with $\Omega(v) \in \mathrm{Ans}(\Omega(v^*))$. However, these vertices also have no incoming edges. Hence, to simplify the definition, we just ignore them.

As desired, due to our assumption that $\Omega(v) \geq 2$ for all $v \in V$, the vertices from $E \times P \times \{0,1\}$ do not influence the maximal color visited infinitely often during a play, unless Player 0 opts to remain in some $(e, p, 0)$ ad infinitum (and thereby violating the parity condition) or an answer to the color of v^* is reached (and thereby satisfying the parity condition).

Moreover, recall that our aim is to allow Player 1 to choose the polarity of edges by switching between the two copies of $\mathcal{A}$ occurring in $\mathcal{A}_{v^*}$. Intuitively, Player 1 should opt for positive polarity in order to unbound the costs incurred by the request posed by v^* from above, while she should opt for negative polarity in order to unbound these costs from below. Since in an energy parity game, it is, broadly speaking, beneficial for Player 1 to move along edges of negative weight, we negate the weights of edges in the copy of $\mathcal{A}$ with positive polarity. Thus, we define

$$w_{v^*}(e) = \begin{cases} -w(v, v') & \text{if } e = ((v,+), ((v,v'),+,1)) \ , \\ w(v, v') & \text{if } e = ((v,-), ((v,v'),-,1)) \ , \\ 1 & \text{if } e = ((e,p,0),(e,p,0)) \ , \\ 0 & \text{otherwise } . \end{cases}$$

This definition implies that the self-loops at vertices of the form (v, p) with $\Omega(v) \in \text{Ans}(\Omega(v^*))$ have weight zero. Combined with the fact that these vertices have color zero, this allows Player 0 to win $\mathcal{G}_{v^*}$ by reaching such a vertex. Intuitively, answering the request posed at v^* is beneficial for Player 0. In particular, if $\Omega(v^*)$ is even, then Player 0 wins $\mathcal{G}_{v^*}$ trivially from (v^*, p), as we then have $\Omega(v^*) \in \text{Ans}(\Omega(v^*))$.

Now, we define the energy parity game $\mathcal{G}_{v^*} = (\mathcal{A}_{v^*}, \Omega_{v^*}, w_{v^*})$.

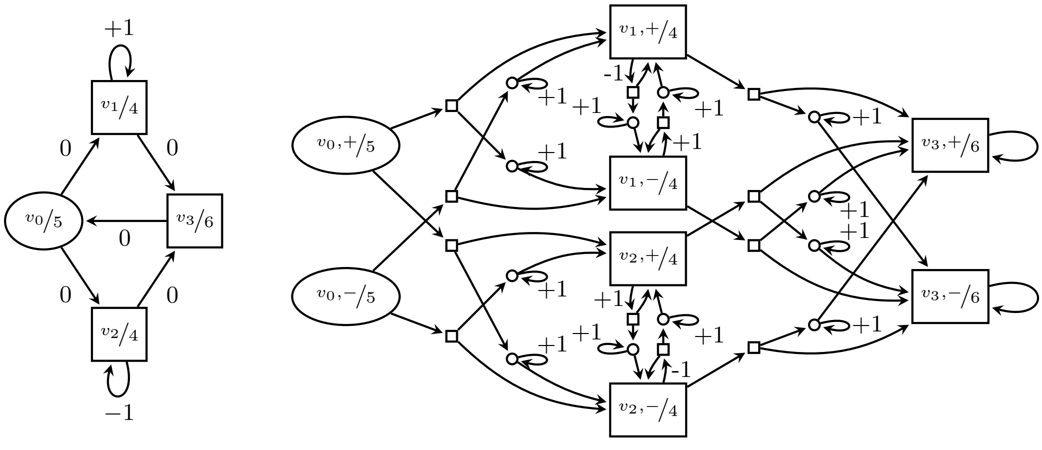

Fig. 3. A bounded parity game with weights $\mathcal{G}$ and the associated energy parity game $\mathcal{G}_{v_0}$. The unnamed vertices of Player 1 (Player 0) are of the form $((v, v'), p, 1)$ (of the form $((v, v'), p, 0)$) when between the vertices (v, p) and v', p'. All missing edge weights in $\mathcal{G}_{v_0}$ are 0.

Example 1. Consider the bounded parity game with weights depicted on the left side of Figure 3 and the associated energy parity game $\mathcal{G}_{v))}$ on the right side. First, let us note that all other $\mathcal{G}_v$ for $v \neq v_0$ are trivial in the following sense: they all consist of a single vertex of even color with a self-loop of weight zero. Hence, Player 0 wins each of these games.

Player 1 wins $\mathcal{G}$ from v_0, where a request for 1 is opened, which is then kept unanswered with infinite cost by using the self-loop at v_1 or v_2 ad infinitum, depending on which successor Player 0 picks.

We show that Player 1 wins $\mathcal{G}_{v_0}$ from $(v_0, +)$): the outgoing edges of $(v_0, +)$ correspond to picking the successor v_1 or v_2 as in $\mathcal{G}$. But before this is executed, Player 1 gets to pick the polarity of the successor: he should pick $+$ for v_1 and $-$ for v_2. Now, Player 0 can either use the self-loop at her "tiny" vertices ad infinitum. These vertices have color one, i.e., Player 1 wins the resulting play. Or, we reach the vertex $(v_1, +)$ or $(v_2, -)$. From both vertices, Player 1 can enforce a loop of negative weight, which allows him to win by violating the parity condition.

Note that the winning strategy for Player 1 for $\mathcal{G}$ from v is very similar to that for him for $\mathcal{G}_{v_0}$ from $(v_0, +)$. We show that one direction holds in general: A winning strategy for Player 0 for $\mathcal{G}_v$ from $(v, +)$ is "essentially" one for him in $\mathcal{G}$ from v.

Note that the other direction does, in general, not hold. This can be seen by adding a vertex v_{-1} of color 3 with a single edge to v_0. Then, vertices of the form (v_i, p) with $i \in \{1, 2\}$ in $\mathcal{G}_{v_{-1}}$ are winning sinks for Player 0. Hence, she wins $\mathcal{G}_{v-1}$ from (v_{-1}) in spite of losing the bounded parity game with weights from v_{-1}.

Hence, the initial request the vertex v inducing $\mathcal{G}_v$ plays a special role in the construction: it is the request Player 1 aims to keep unanswered with infinite cost. To overcome this and to complete our construction, we show a statement reminiscent of Lemma 4: If Player 0 wins $\mathcal{G}_v$ from $(v, +)$ for every v, then she also wins $\mathcal{G}$ from every vertex. With this relation at hand, one can again construct a fixed-point algorithm solving bounded parity games with weights using an oracle for solving energy parity games that is very similar to Algorithm 1.

Lemma 4. *Let $\mathcal{G}$ be a bounded parity game with weights with vertex set V.*

1. *Let $v^* \in V$. If Player 1 wins $\mathcal{G}_{v^*}$ from $(v^*, +)$, then $v^* \in \mathcal{W}_1(\mathcal{G})$.*
2. *If Player 0 wins $\mathcal{G}_{v^*}$ from $(v^*, +)$ for all $v^* \in V$, then $\mathcal{W}_1(\mathcal{G}) = \emptyset$.*

This lemma is the main building block for the algorithm that solves bounded parity games with weights by repeatedly solving energy parity games, which is very similar to Algorithm 1. Indeed, we just swap the roles of the players: We compute 1-attractors instead of 0-attractors and we change the definition of X_k. Hence, we obtain the following algorithm (Algorithm 2).

Algorithm 2 A fixed-point algorithm computing $\mathcal{W}_1(\mathcal{A}, \mathrm{BndWeightParity}(\Omega, w))$.

$k = 0$; $W_1^k = \emptyset$; $\mathcal{A}_k = \mathcal{A}$
repeat
  $k = k + 1$
  $X_k = \{v^* \mid \text{Player 1 wins the energy parity game } ((\mathcal{A}_{k-1})_{v^*}, \Omega_{v^*}, w_{v^*}) \text{ from } (v^*, +)\}$
  $W_0^k = W_0^{k-1} \cup \mathrm{Attr}_1^{\mathcal{A}_{k-1}}(X_k)$
  $\mathcal{A}_k = \mathcal{A}_{k-1} \setminus \mathrm{Attr}_1^{\mathcal{A}_{k-1}}(X_k)$
until $X_k = \emptyset$
return W_1^k

Algorithm 2 terminates after solving at most a quadratic number of energy parity games. Furthermore, the proof of correctness is analogous to the one for Algorithm 1, relying on Lemma 4. We only need two further properties: the 1-extendability of $\mathrm{BndWeightParity}(\Omega, w)$, and an assertion that $\mathrm{Attr}_1^{\mathcal{A}_{k-1}}(X_k)$ is a trap for Player 0 in $\mathcal{A}_{k-1}$. Both are easy to verify.

Alter plugging Algorithm 2 into Algorithm 1, Proposition 3 yields our main theorem, settling the complexity of solving parity games with weights.

Theorem 1. *The following problem is in* NP $\cap$ CO-NP*: "Given a parity game with weights $\mathcal{G}$ and a vertex v in $\mathcal{G}$, does Player 0 win $\mathcal{G}$ from v?"*

The following two subsections are dedicated to showing the two assertions of Lemma 4. In order to prepare for this, we first introduce some notation. Let $v^* \in V$ and consider $\mathcal{G}_{v^*}$. We distinguish three types of plays in $\mathcal{G}_{v^*}$:

Type -1: Plays that have a suffix $(e, p, 0)^\omega$ for some $e \in E$ and some $p \in P$.
Type 0: Plays that visit infinitely many vertices from both $V \times P$ and $E \times P \times \{0, 1\}$.
Type 1: Plays that have a suffix $(v, p)^\omega$. Note that this implies $\Omega(v) \in \mathrm{Ans}(\Omega(v^*))$.

Remark 2. Let ρ' be a play in $\mathcal{G}_{v^*}$ that starts in (v^*, p).

1. If ρ' is consistent with a winning strategy for Player 0 from (v^*, p), then ρ' is not a play of type -1.
2. If ρ' is consistent with a winning strategy for Player 1 from (v^*, p), then ρ' is not a play of type 1.

We define the mapping unpol: $V' \to V \cup \{\varepsilon\}$ as $\text{unpol}(v,p) = v$ and $\text{unpol}(e,p,i) = \varepsilon$ for $v \in V$, $e \in E$, $p \in P$, and $i \in \{0,1\}$, which we extend to sequences of vertices in the straightforward way. Let $\rho' \in (V')^* \cup (V')^\omega$. We call $\text{unpol}(\rho')$ the **unpolarization** of ρ'.

Remark 3. Let ρ' be a play of type 0 in some $\mathcal{G}_{v^*}$. We have $\rho' \in \text{Parity}(\Omega_{v^*})$ if and only if $\text{unpol}(\rho') \in \text{Parity}(\Omega)$.

5.3 Proof of Lemma 4.1

Let τ_{v^*} be a winning strategy for Player 1 from $(v^*,+)$ in $\mathcal{G}_{v^*}$. We define a winning strategy τ for Player 1 from v^* in $\mathcal{G}$ such that τ mimicks the moves made by τ_{v^*}. To this end, τ keeps track of a play prefix $\mathcal{G}_{v^*}$. Formally, we define τ together with a simulation function h that satisfies the following invariant:

> If π is a non-empty play prefix in $\mathcal{A}$ that starts in v^*, is consistent with τ, and ends in some v, then $h(\pi)$ is a play prefix in $\mathcal{A}_{v^*}$ that starts in $(v^*,+)$, is consistent with τ_{v^*}, and ends in some (v,p). Furthermore, $\text{unpol}(h(\pi)) = \pi$.

Recall that, if h has the properties described above, then, due to the structure of $\mathcal{A}_{v^*}$, for each π, given $h(\pi)$, the strategy τ_{v^*} prescribes a move to some vertex $((v,v'),p,1)$, where $(v,v') \in E$. We can mimic this choice by moving to v' in $\mathcal{G}$.

We now define h and τ formally and begin with $h(v^*) = (v^*,+)$, which clearly satisfies the invariant. Now let $\pi = v_0 \cdots v_j$ be some non-empty play prefix in $\mathcal{A}$ beginning in v^* and consistent with τ such that $h(\pi)$ is defined. Due to the invariant, $h(\pi)$ ends in (v_j,p_j) for some $p_j \in P$.

If $v_j \in V_1$, let v_{j+1} be such that $h(\pi) \cdot ((v_j,v_{j+1}),p_j,1)$ is consistent with τ_{v^*} and define $\tau(\pi) = v_{j+1}$. Such a v_{j+1} exists, because (v_j,p_j), the last vertex of $h(\pi)$, satisfies $\Omega(v_j) \notin \text{Ans}(\Omega(v^*))$ due to the invariant, Remark 2.2, and because the answering vertices are sinks. If, however, $v_j \in V_0$, then let v_{j+1} be an arbitrary successor of v_j in $\mathcal{A}$. In either case, it remains to define $h(\pi \cdot v_{j+1})$.

Since we want to simulate the move from v_j to v_{j+1} in $h(\pi \cdot v_{j+1})$, we first move from (v_j,p_j) to $((v_j,v_{j+1}),p_j,1)$. Moreover, in order to satisfy the invariant, we aim to simulate the play prefix $\pi \cdot v_{j+1}$ such that $h(\pi \cdot v_{j+1})$ is consistent with τ_{v^*}. This strategy may prescribe for Player 1 to either preserve the polarity p_j, or to switch it during the simulated move from v_j to v_{j+1}.

In the former case, i.e., if $\tau_{v^*}(h(\pi) \cdot ((v_j,v_{j+1}),p_j,1)) = (v_{j+1},p_j)$, we define $h(\pi \cdot v_{j+1}) = h(\pi) \cdot ((v_j,v_{j+1}),p_j,1) \cdot (v_{j+1},p_j)$. In the latter case, Player 0 gets an opportunity to recharge the energy by taking the self-loop of the vertex $((v_j,v_{j+1}),p_j,0)$ finitely often. We opt to let her recover the energy lost so far in the play prefix, i.e., we pick $c_j = \max\{0, -w(h(\pi) \cdot ((v_j,v_{j+1}),p_j,1)\}$ and define $h(\pi \cdot v_{j+1}) = h(\pi) \cdot ((v_j,v_{j+1}),p_j,1) \cdot ((v_j,v_{j+1}),p_j,0)^{c_j} \cdot (v_{j+1},\overline{p_j})$ in this case. Since $h(\pi \cdot v_{j+1})$ is consistent with τ_{v^*} in either case, we satisfy the invariant in either case. This completes the definition of τ and h.

It remains to show that τ is indeed winning from v^* in $\mathcal{G}$. To this end, let $\rho = v_0 v_1 v_2 \cdots$ be a play in $\mathcal{A}$ that starts in v^* and is consistent with τ. We show $\rho \notin \text{BndWeightParity}(\Omega, w)$ by examining the play ρ' in $\mathcal{A}_{v^*}$, which is limit of the $h(\pi)$ for increasing prefixes π of ρ. Due to the invariant, ρ' starts in $(v^*,+)$ and is consistent with τ_{v^*}. Moreover, due to the construction of h, we obtain $\text{unpol}(\rho') = \rho$. Finally, we have that ρ' is a play of type 0 in $\mathcal{G}_{v^*}$. Hence, due to Remark 3, ρ satisfies the parity condition if and only if ρ' satisfies the parity condition.

As ρ' is consistent with the winning strategy τ_{v^*}, we have $\rho' \notin \text{EnergyParity}(\Omega_{v^*}, w_{v^*})$, i.e., ρ' either violates the parity condition or the energy condition. Hence, as argued above, if ρ' violates the parity condition, then so does ρ, i.e., ρ is indeed winning for Player 1.

Now assume that ρ' violates the energy condition. Due to the structure of $\mathcal{A}_{v^*}$ and the construction of h we have

$$\rho' = \mathit{\Pi}_{j=0,1,2,\ldots}(v_j,p_j) \cdot ((v_j,v_{j+1}),p_j,1) \cdot ((v_j,v_{j+1}),p_j,0)^{m_j}$$

for some $m_j \in \mathbb{N}$. Since ρ' violates the energy condition, we have $\inf_{j \in \mathbb{N}} w((v_0,p_0) \cdots (v_j,p_j) \cdot ((v_j,v_{j+1}),p_j,1)) = -\infty$. The restriction to play prefixes of this form suffices due to the structure of $\mathcal{A}_{v^*}$ and, in turn, the structure of ρ'. Moreover, since Player 1 wins $\mathcal{G}_{v^*}$ from $(v^*,+)$, the initial vertices v^* and $(v^*,+)$ of ρ and ρ', respectively, have the same odd color. Also, as ρ' is a play of type 0, the request for the color $\Omega(v^*)$ is never answered in ρ or ρ'. We show that the request for $\Omega(v^*)$ in ρ is unanswered with infinite cost, which concludes the proof.

To this end, we split ρ' into infixes of constant polarity. Given a vertex $v = (v', p)$ or $v = ((v', v''), p, i)$, we call p the **polarity** of v. Let $\rho' = \nu'_0 \nu'_1 \nu'_2 \cdots$, where each ν'_j is a maximal finite (or infinite) infix (or suffix) of ρ', such that all vertices in ν'_j have the same polarity. We call an infix ν'_j of ρ' an **equi-polarity infix** (EPI) of ρ'.

Since the polarity remains constant throughout each ν'_j, Player 0 only resets the energy via repeatedly traversing a self-loop of a vertex in $\mathcal{A}_{v^*}$ at the last vertex visited in ν'_j, if at all. Hence, the energy levels attained during ν'_j and $\mathrm{unpol}(\nu'_j)$ are closely related.

Remark 4. Let ν' be an EPI beginning in (v_j, p_j) and let $\nu = \mathrm{unpol}(\nu') = v_j v_{j+1} v_{j+2} \cdots$. For each j' with $j \leq j' < j + |\nu|$, we have $|w(v_j \cdots v_{j'})| = |w((v_j, p_j) \cdots (v_{j'}, p_{j'}))|$.

In particular, Remark 4, the structure of $\mathcal{A}_{v^*}$, and the definition of h imply $\mathrm{Ampl}(\mathrm{unpol}(\nu')) = \mathrm{Ampl}(\nu')$ for all EPIs ν' of ρ'. Thus, if there exist only finitely many EPIs of ρ', let ν' be the infinite final EPI of ρ', let $\nu = \mathrm{unpol}(\nu')$, and note that, due to $\mathrm{Ampl}(\rho') = \infty$, we have $\mathrm{Ampl}(\nu') = \infty$. Due to Remark 4, we obtain $\mathrm{Ampl}(\nu) = \infty$, which implies that the request posed at the initial position of ρ is unanswered with infinite cost due to the reasoning above, as ν is a suffix of ρ.

If, however, there exist infinitely many EPIs of ρ', assume towards a contradiction that the cost of answering the request posed at the initial position of ρ is finite. By construction of ρ', the energy level is non-negative at the end of each EPI. Since ρ' violates the energy condition, for each bound $b \in \mathbb{N}$ there exists an EPI ν' of ρ' with a prefix of weight strictly smaller than $-b$. We obtain $\mathrm{Ampl}(\mathrm{unpol}(\nu')) > b$ via Remark 4. This contradicts the cost of answering the request posed at the initial position of ρ being bounded and concludes the proof of Lemma 4.1.

5.4 Proof of Lemma 4.2

To prove Lemma 4.2, we construct a strategy σ for Player 0 in $\mathcal{G}$ that is winning for her from each vertex $v \in V$. As winning regions are disjoint, this implies the desired result.

For each energy parity game $\mathcal{G}_v = (\mathcal{A}_v, \Omega_v, w_v)$ we have $n' = |\mathcal{A}_v| \in \mathcal{O}(|\mathcal{A}|^2)$, we have $d' = |\Omega_v(V')| = |\Omega(V)| + 2$, and we have $W' = \max(w(E')) = \max(w(E) \cup \{1\})$, where E and E' are the sets of edges in $\mathcal{A}$ and the $\mathcal{A}_v$, respectively. Note that the values n', d', and W' of $\mathcal{G}_v$ are independent of the vertex v, which explains our notation. Due to the assumption of the statement and Proposition 2, for each $v \in V$, there exists a finite-state strategy σ_v with at most $n'd'W'$ states that is winning for Player 0 from $(v, +)$ in $\mathcal{G}_v$.

We construct the winning strategy σ for Player 0 in $\mathcal{G}$ by "stitching together" the individual σ_v. To this end, given a play prefix, we identify the request which should be answered most urgently. Say this request was opened by visiting vertex v. The strategy σ then mimics the moves made by σ_v when starting in $(v, +)$. Once the request for $\Omega(v)$ is answered, σ makes arbitrary moves until a new request is opened.

Formally, given a play prefix $\pi = v_0 \cdots v_j$, we say that a request for color c is open in π if there exists a position j' with $0 \leq j' \leq j$ such that $\Omega(v_{j'}) = c$ and, for all positions j'' with $j' \leq j'' \leq j$, we have $\Omega(v_{j''}) \notin \mathrm{Ans}(\Omega(v'_j))$. Note that there is never an open request for an even color.

If there is no open request in π, the position of the most relevant request is undefined and we write $\mathrm{posMRR}(\pi) = \bot$. Otherwise, let c be the maximal color, for which there is an open request in π. We define $\mathrm{posMRR}(\pi)$ as the smallest position j', such that the request for color c is open in all prefixes of π of length greater than j'.

As an example, consider the play prefix shown in Figure 4 using the notation $v/\Omega(v)$. We mark a position j in red with solid background if $\mathrm{posMRR}(v_0 \cdots v_j) = j$ and in green with dashed background if $\mathrm{posMRR}(v_0 \cdots v_j) = \bot$. Otherwise, i.e., if $\bot \neq \mathrm{posMRR}(v_0 \cdots v_j) < j$, we leave j unmarked. For those positions, $\mathrm{posMRR}(v_0 \cdots v_j)$ is equal to the largest (i.e., last visited) earlier position marked in red.

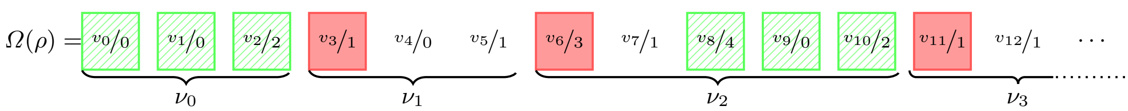

Fig. 4. A play ρ, its induced color sequence, its most relevant requests, and the ESIs of ρ.

In order to leverage the moves made by the strategies for the $\mathcal{G}_v$ in $\mathcal{G}$, we need to simulate play prefixes in the latter game in the former ones. To this end, we again define σ together with a simulation function h. This function h maps a play prefix consistent with σ to a sequence of vertices from V' (not necessarily a play prefix) such that we are able to leverage the choices made by the σ_v in order to define σ. Our aim is to define h such that it satisfies the following invariant:

> Let $\pi = v_0 \cdots v_j$ be a play in $\mathcal{A}$ consistent with σ. Then $h(\pi)$ ends in some (v_j, p_j). Moreover, if $\text{posMRR}(\pi) = j' \neq \bot$, then $h(\pi)$ has a (unique) suffix $(v_{j'}, +) \cdots (v_j, p_j)$ that is consistent with $\sigma_{v_{j'}}$ and satisfies $\text{unpol}((v_{j'}, +) \cdots (v_j, p_j)) = v_{j'} \cdots v_j$.

We define h and σ inductively and begin with $h(v) = (v, +)$ for each $v \in V$, which clearly satisfies the invariant. Now let $\pi = v_0 \cdots v_j$ be a non-empty play prefix in $\mathcal{A}$ starting in v^* and consistent with σ such that $h(\pi)$ is defined. If $v_j \in V_1$, let v_{j+1} be an arbitrary successor of v_j in $\mathcal{A}$. If, however, $v_j \in V_0$, we distinguish two cases based on $\text{posMRR}(\pi)$. If $\text{posMRR}(\pi) = \bot$, again let v_{j+1} be an arbitrary successor of v_j. If, however, $\text{posMRR}(\pi) = j' \neq \bot$, then the invariant of h yields a suffix $(v_{j'}, +) \cdots (v_j, p_j)$ of $h(\pi)$ that is consistent with $\sigma_{v_{j'}}$. Let v_{j+1} such that $(v_{j'}, +) \cdots (v_j, p_j) \cdot ((v_j, v_{j+1}), p_j, 1)$ is consistent with $\sigma_{v_{j'}}$. Such a vertex v_{j+1} exists, because the request posed by visiting $v_{j'}$ is open in π due to $\text{posMRR}(\pi) = j'$. Since $v_j \in V_0$, the vertex v_{j+1} is unique.

It remains to define $h(\pi \cdot v_{j+1})$ in such a way, that it satisfies the above invariant. To this end, we use one of two operations. Firstly, we define the **discontinuous extension of $h(\pi)$ with v_{j+1}** as $h(\pi) \cdot (v_{j+1}, +)$. Secondly, we define a simulated extension of $h(\pi)$ such that we obtain $h(\pi \cdot v_{j+1})$ by simulating the move from v_j to v_{j+1} in some $\mathcal{G}_v$. Formally, we define the **simulated extension of $h(\pi)$ with v_{j+1} and charge** m as $h(\pi) \cdot ((v_j, v_{j+1}), p_j, 1) \cdot ((v_j, v_{j+1}), p_j, 0)^m \cdot (v_{j+1}, p_{j+1})$, where $p_{j+1} = p_j$ if $m = 0$ and $p_{j+1} = \overline{p_j}$ otherwise. This ensures that the extension is a play infix in some $\mathcal{G}_v$.

We first distinguish two cases. If $\text{posMRR}(\pi) = \bot$, we define $h(\pi \cdot v_{j+1})$ to be the discontinuous extension of $h(\pi)$ with v_{j+1}. This clearly satisfies the first condition of the invariant. Moreover, the second condition of the invariant is satisfied as well: If $\text{posMRR}(\pi \cdot v_{j+1}) = \bot$, this condition holds true vacuously. Otherwise, we have $\text{posMRR}(\pi \cdot v_{j+1}) = j+1$ and observe that the suffix $(v_{j+1}, +)$ of $h(\pi \cdot v_{j+1})$ satisfies the second condition of the invariant.

If, however, $\text{posMRR}(\pi) \neq \bot$, let $\text{posMRR}(\pi) = j'$. We distinguish three sub-cases. First, assume that the move to v_{j+1} neither opens a new most relevant request, nor answers the existing one, i.e., $\text{posMRR}(\pi \cdot v_{j+1}) = j'$. In this case, we extend the suffix of $h(\pi)$ that is consistent with $\sigma_{v_{j'}}$ by simulating the move from v_j to v_{j+1}. Recall that we picked v_{j+1} such that $(v_{j'}, +) \cdots (v_j, p_j) \cdot ((v_j, v_{j+1}), p_j, 1)$ is consistent with $\sigma_{v_{j'}}$. As we can freely define the choice of Player 1 in the simulation, we follow the intuition stated during the construction of the polarized arena. Recall that both players are currently playing "with respect to" the request for $\Omega(v_{j'})$ opened by visiting $v_{j'}$. Hence, we opt to let Player 1 move to positive polarity if the cost of the request for $\Omega(v_{j'})$ so far is nonnegative, and let her move to negative polarity otherwise.

We use $\text{Sgn}(n) = +$ for $n \geq 0$ and $\text{Sgn}(n) = -$ for all other n. If $\text{Sgn}(w(v_{j'} \cdots v_{j+1})) = p_j$, we define $h(\pi \cdot v_{j+1})$ to be the simulated extension of $h(\pi)$ with v_{j+1} and charge 0. Otherwise, i.e., if $\text{Sgn}(w(v_{j'} \cdots v_{j+1})) = \overline{p_j}$, let $m \in \mathbb{N}$ such that $\sigma_{v_{j'}}((v_{j'}, +) \cdots (v_j, p_j) \cdot ((v_j, v_{j+1}), p_j, 1) \cdot ((v_j, v_{j+1}), p_j, 1)^m) = (v_{j+1}, \overline{p_j})$. Such an m exists, as otherwise the play $(v_{j'}, +) \cdots (v_j, p_j) \cdot ((v_j, v_{j+1}), p_j, 1) \cdot ((v_j, v_{j+1}), p_j, 0)^\omega$ of type -1 that starts in $(v_{j'}, +)$ would be consistent with the winning strategy $\sigma_{v_{j'}}$ from $(v_{j'}, +)$ for Player 0, a contradiction to Remark 2.1. In this case, we define $h(\pi \cdot v_{j+1})$ to be the simulated extension of $h(\pi)$ with v_{j+1} and charge m. Since we have $\text{posMRR}(v_0 \cdots v_{j+1}) = j'$ by assumption, either definition of $h(\pi \cdot v_{j+1})$ satisfies the invariant.

It remains to define $h(\pi \cdot v_{j+1})$ for the two remaining cases, i.e., for the case that the move to v_{j+1} answers the most relevant request in π, and for the case that the request posed by visiting v_{j+1} is the most relevant request of $\pi \cdot v_{j+1}$. Formally, we have $\text{posMRR}(\pi \cdot v_{j+1}) = \bot$ in the former case and $\text{posMRR}(\pi \cdot v_{j+1}) = j+1$ in the latter one. In either case, we define $h(\pi \cdot v_{j+1})$ to be the discontinuous extension of $h(\pi)$ with v_{j+1} and observe that the first condition of the invariant holds. If $\text{posMRR}(\pi \cdot v_{j+1}) = \bot$, the second condition of the invariant vacuously holds. If, however, $\text{posMRR}(\pi \cdot v_{j+1}) = j+1$, then the suffix $(v_{j+1}, +)$ witnesses that the second condition of the invariant holds. This completes the definition of σ and h.

It remains to show that the strategy σ is indeed winning for Player 0 from v^*. To this end, fix some play $\rho = v_0 v_1 v_2 \cdots$ consistent with σ starting in v^*, and let ρ' be the limit of the $h(\pi)$ for increasing prefixes π of ρ. By the construction of h, we obtain $\mathrm{unpol}(\rho') = \rho$. Hence, the play ρ' is of the form $(v_0, p_0) \cdots (v_1, p_1) \cdots (v_2, p_2) \cdots$. We call a position j of ρ a **discontinuity** of ρ if either $j = 0$ or if $h(v_0 \cdots v_j)$ is the discontinuous extension of $h(v_0 \cdots v_{j-1})$ by v_j.

Let j and j' be adjacent discontinuities of ρ. We call the infix $v_j \cdots v_{j'-1}$ of ρ an **equi-strategic infix** (ESI) of ρ. Moreover, if there only exist finitely many discontinuities of ρ, let j^* be its final discontinuity. We call the suffix $v_{j^*} v_{j^*+1} v_{j^*+2} \cdots$ of ρ the **terminal** ESI of ρ.

Remark 5. Let $\nu = v_j v_{j+1} v_{j+2} \cdots$ be an ESI.

1. If ν is finite, then the infix ν' of ρ' starting at position $|h(v_0 \cdots v_j)|$ and ending at position $|h(v_0 \cdots v_{j+|\nu|-1})|$ starts in $(v_j, +)$, ends in some $(v_{j+|\nu|-1}, p)$, and is consistent with σ_{v_j}.
2. If ν is infinite, then the suffix ν' of ρ' starting at position $|h(v_0 \cdots v_j)|$ starts in $(v_j, +)$ and is consistent with σ_{v_j}.

For each position j of ρ we define $\mathrm{ESI}(j) = k$ if the k-th ESI of ρ contains v_j. Moreover, if $\nu = v_j v_{j+1} v_{j+2} \cdots$ is an ESI of ρ, then we call $\Omega(v_j)$ the **characteristic color** of ν. By the construction of h, if the characteristic color of an ESI ν is even, then ν consists only of a single vertex. If, however, the characteristic color c of an ESI ν is odd, then we have $\Omega(v) \leq c$ for all vertices v in ν. Moreover, let c' be the characteristic color of the ESI succeeding ν, if ν is not the terminal ESI of ρ. Due to the construction of h, we have $c' > c$. If c' is even, this observation implies $c' \in \mathrm{Ans}(\Omega(v))$ for all vertices v in ν. As the number of colors in $\mathcal{G}$ is finite, this in turn implies that the number of ESIs between a request and its response (if a response exists at all) is bounded.

Remark 6. Let j be some position in ρ and let $k = \mathrm{ESI}(j)$. Moreover, let d be the number of colors in $\mathcal{G}$.

1. If the request at position j is first answered at position j', then $\mathrm{ESI}(j') < k + d$
2. If the request at position j is unanswered in ρ, then ρ contains less than $k + d$ many ESIs.

Recall that the bounded parity condition with weights requires the play ρ to not only satisfy the parity condition, but also that the cost of almost all requests is bounded and that there exists no unanswered request with infinite cost in ρ. We first show that ρ satisfies the classical parity condition. In a second step, we then show that there exists a bound on the cost of each (answered or unanswered) request in ρ. The former condition, i.e., that ρ satisfies the parity condition, is in large parts implied by Remark 6.

Lemma 5. *The play ρ satisfies the parity condition.*

Proof. If ρ contains no unanswered requests, then it vacuously satisfies the parity condition. Hence, let j be the position of such an unanswered request in ρ. Due to Remark 6, we obtain that there exist only finitely many ESIs in ρ. Let $\nu = v_{j^*} v_{j^*+1} v_{j^*+2} \cdots$ be the terminal ESI of ρ. By the construction of h, there exists a suffix ν' of ρ' with $\mathrm{unpol}(\nu') = \nu$. Due to Remark rem:esi:equivalence:infinite.5, the suffix ν' begins in $(v_{j^*}, +)$ and is consistent with the winning strategy $\sigma_{v_{j^*}}$ for Player 0 from $(v_{j^*}, +)$ in $\mathcal{G}_{v_{j^*}}$. Moreover, ν' is a play of type 0 due to ν' being the terminal ESI of ρ and due to being consistent with $\sigma_{v_{j^*}}$. Hence, we obtain that ν satisfies the parity condition via Remark 3, which in turn implies that ρ satisfies the parity condition. □

It remains to show that the costs of requests in ρ are bounded. Recall that we defined $n' = |\mathcal{A}_v|$, d' as the largest color of a vertex in the $\mathcal{G}_v$, and W' as the largest absolute weight of an edge. We claim that the costs of the most relevant requests in ρ are bounded by $(n'd'W')^2$. This implies that the cost of all requests is bounded: Due to Remark 6 we obtain that the number of ESIs between a request and its response, if one exists, is bounded by d. Hence, it suffices to show that each ESI contributes at most a bounded amount to the cost of answering a request.

Lemma 6. *Let $\nu = v_j v_{j+1} v_{j+2} \cdots$ be an ESI of ρ. For each j' with $j \leq j' < j + |\nu|$ we have $|w(v_j \cdots v_{j'})| \leq d'(n'W')^2$.*

Proof. Towards a contradiction let j' be a position with $j \leq j' < j + |\nu|$, such that we have $|w(v_j \cdots v_{j'})| > d'(n'W')^2$. We assume $w(v_j \cdots v_{j'}) > d'(n'W')^2$, i.e., that the infix $v_j \cdots v_{j'}$ violates the claimed bound from above. The other case is dual. Since each traversed edge adds a cost of at most W', there exists a minimal position j'' such that, for all k with $j'' \leq k \leq j'$, we have $w(v_j \cdots v_k) > 0$. Let $\pi' =$

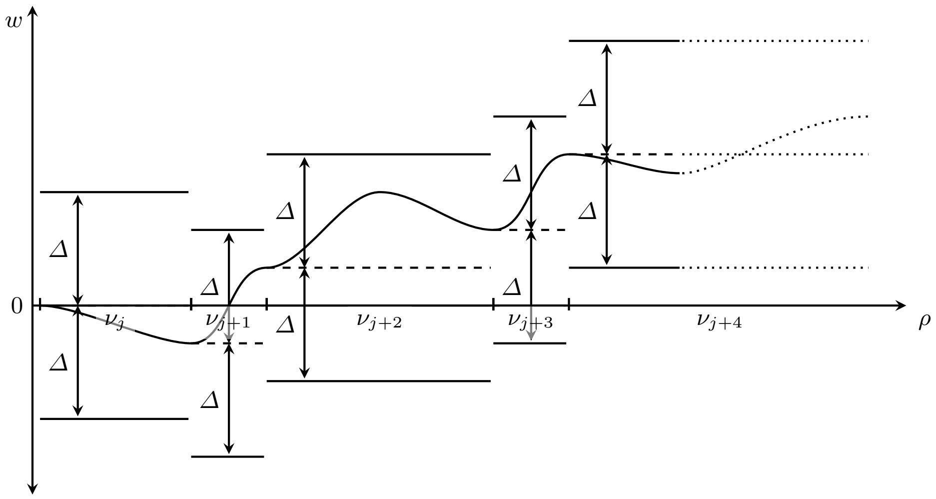

Fig. 5. Bounds on the cost of a request over time given by Lemma 6. We write $\varDelta = d'(n'W')^2$.

$(v_{j''}, p_{j''}) \cdots (v_{j'}, p_{j'})$ be an infix of ρ' such that $\mathrm{unpol}(\pi') = v_{j''} \cdots v_{j'}$. The polarity remains positive throughout π' due to the construction of h, while we have $w_{v_j}(\pi') < -d'(n'W')^2$. Moreover, by definition of ESIs, the infix π' is an infix of a play that starts in $(v_j, +)$ and is consistent with σ_{v_j}. This, however, contradicts Lemma 3. □

Due to Lemma 6, each ESI strictly in-between a request and its response contributes at most $d'(n'W')^2$ to the cost incurred by the request. Similarly, the ESI containing the request and its response also contribute at most $d'(n'W')^2$ each to the cost of answering the given request. Hence, via Remark 6, we obtain that each (answered or unanswered) request in ρ incurs a cost of at most $(d'n'W')^2$. We illustrate this argument in Figure 5. Hence, σ is a winning strategy for Player 0 from v^* in $\mathcal{G}$, as each play that starts in v^* and is consistent with σ satisfies the parity condition due to Lemma 5 and because no such play contains a request that is unanswered with infinite cost, which concludes the proof of Lemma 4.2.

Before we conclude this section, we formalize the above observation about the winning strategy for Player 0 uniformly bounding the costs of requests in the following corollary.

Corollary 1. *Let $\mathcal{G}$ be a bounded parity game with weights with n vertices, d colors, and largest absolute weight W. There exists a strategy σ for Player 0 that is winning from $\mathcal{W}_0(\mathcal{G})$, such that in each play ρ consistent with σ, each request is answered or unanswered with cost at most $((d+2)(2n+4n^2)(W+1))^2$.*

Using arguments from Section 7, this bound can be improved to $((d+2)(6n)(W+1))^2$. However, as we only use Corollary 1 later to obtain some upper bound on the quality of such strategies, we refrain from repeating these arguments here.

6 Memory Requirements

We now discuss the upper and lower bounds on the memory required to implement winning strategies for either player. Recall that we use binary encoding to denote weights, i.e., weights may be exponential in the size of the game. In this section we show polynomial (in n, d, and W) upper and lower bounds on the necessary and sufficient memory for Player 0 to win parity games with weights. Due to the binary encoding of weights, these bounds are exponential in the size of the game. In contrast, Player 1 requires infinite memory.

Theorem 2. *Let $\mathcal{G}$ be a parity game with weights with n vertices, d colors, and largest absolute weight W assigned to any edge in $\mathcal{G}$. Moreover, let v be a vertex of $\mathcal{G}$.*

1. *Player 0 has a winning strategy σ from $\mathcal{W}_0(\mathcal{G})$ with $|\sigma| \in \mathcal{O}(nd^2W)$. This bound is tight.*
2. *Player 1 requires, in general, infinite memory to win from $\mathcal{W}_1(\mathcal{G})$.*

The proof of the second item of Theorem 2 is straightforward, since Player 1 already requires infinite memory to implement winning strategies in finitary parity games [7]. Since parity games with weights subsume finitary parity games, this result carries over to our setting. We show the game witnessing this lower bound on the right-hand side of Figure 2.

Proposition 4 ([7]). *There exists a parity game with weights $\mathcal{G}$, such that Player 1 has a winning strategy from each vertex v in $\mathcal{G}$, but she has no finite-state winning strategy from any v in $\mathcal{G}$.*

Having argued that no finite upper bound on the space requirements of winning strategies for Player 1 exists, we now show that, in contrast, exponential memory is sufficient, but also necessary, for Player 0. To this end, we first prove that the winning strategy for him constructed in the proof of Lemma 4.2 suffers at most a linear blowup in comparison to his winning strategies in the underlying energy parity games. This is sufficient as we have argued in Section 4 that the construction of a winning strategy for Player 0 in a parity game with weights suffers no blowup in comparison to the underlying bounded parity games with weights.

Lemma 7. *Let $\mathcal{G}$, n, d, and W be as in Theorem 2. Player 0 has a finite-state winning strategy of size at most $d(6n)(d+2)(W+1)$ from $\mathcal{W}_0(\mathcal{G})$.*

Proof. We only show the above result for the case where $\mathcal{G}$ is a bounded parity game with weights. The result for the case where $\mathcal{G}$ is a parity game with weights then follows directly as argued above.

Let V and E be the vertex set of $\mathcal{G}$ and recall that we have defined $P = \{+, -\}$. In the proof of Lemma 4.2, we have constructed an energy parity game $\mathcal{G}_v$ with vertices $(V \times P) \cup (E \times P \times \{0, 1\})$ for each vertex v of $\mathcal{G}$. We have then constructed a winning strategy σ for Player 0 for $\mathcal{G}$ out of winning strategies for her in the $\mathcal{G}_v$. As it is straightforward to implement σ via the disjoint union of memory structures implementing the constituent strategies, this approach yields an upper bound of $n(2n+4n^2)(d+2)(W+1)$ on the size of σ due to the upper bound on the size of winning strategies for Player 0 in energy parity games from Proposition 2.

In the construction of the $\mathcal{G}_v$, however, we only store the edges chosen by the players in the vertices of the form $E \times P \times \{0, 1\}$ for didactic purposes. In fact, it suffices to store the target vertex of an edge instead, resulting in a vertex set of size $6n$ of the $\mathcal{G}_v$. Moreover, recall that the definition of the $\mathcal{G}_v$ only takes the color of v into account: If the vertices v and v' have the same color, then the games $\mathcal{G}_v$ and $\mathcal{G}_{v'}$ are isomorphic. Further, Chatterjee and Doyen have shown that, if Player 0 wins an energy parity game $\mathcal{G}'$ with n' vertices, d' colors, and largest absolute weight W', then he has a uniform strategy of size $n'd'W'$ that is winning from all vertices, from which he wins $\mathcal{G}'$ [4]. Hence, it suffices to combine at most d strategies, each of size $(6n)(d+2)(W+1)$, in order to obtain a winning strategy for Player 0 in $\mathcal{G}$. □

Having established an upper bound on the memory required by Player 0, we now proceed to show that this exponential bound is indeed tight, which is witnessed by the games $\mathcal{G}_n$ depicted in Figure 6.

Lemma 8. *Let $n \in \mathbb{N}$. There exists a parity game with weights $\mathcal{G}_n$ with $|\mathcal{G}_n| \in \mathcal{O}(n)$ such that Player 0 wins $\mathcal{G}_n$ from every vertex, but each winning strategy for her is of size at least $n2^n + 1$.*

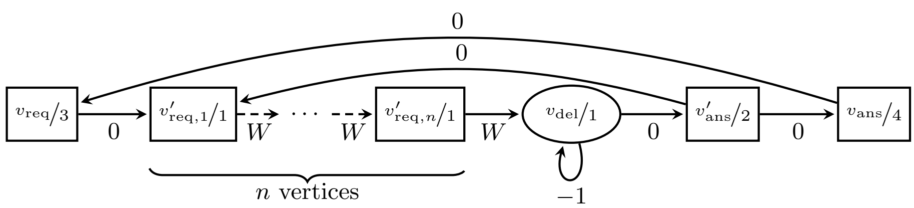

Fig. 6. A game of size $\mathcal{O}(n)$ in which Player 0 only wins with strategies of size at least $n2^n + 1$.

Proof. We show the game $\mathcal{G}_n$ in Figure 6. This game has $n+4$ vertices and the largest absolute weight of an edge is $W = 2^n$. Hence, we have $|\mathcal{G}_n| = n + 4 + \log(W) \in \mathcal{O}(n)$. Note that the only choice of Player 0 in $\mathcal{G}_n$ consists in determining how often to take the self-loop of vertex v_{del} upon each visit. Dually, the only choice of Player 1 consists of deciding whether or not to move from v'_{ans} to $v'_{\text{req},1}$, or to continue to v_{ans}.

Player 0 wins $\mathcal{G}_n$ from each vertex by remaining in v_{del} for nW moves after each visit to v_{del} and by subsequently moving to v'_{ans}. Each request in each play that is consistent with this strategy is answered or unanswered with cost $n2^n$, independent of the choices of Player 1 in v_{ans}. Moreover, as the only way to visit v_{req} is to move there from v_{ans}, the play visits v_{ans} infinitely often if and only if it visits v_{req} infinitely often. Further, the play visits $v'_{\text{req},1}$ and v'_{ans} infinitely often. Hence, all requests are answered, i.e., this strategy is winning for Player 0 from all vertices. This strategy can be implemented by a counter that counts the number of self-loops of v_{del} taken so far, which is reset to nW after leaving v_{del}. As this counter is bounded by nW, the strategy is of size $nW + 1 = n2^n + 1$.

It remains to show that each winning strategy for Player 0 has at least $n2^n+1$ memory states. Towards a contradiction, let σ be a winning strategy for Player 0 from some vertex v with less than $n2^n + 1$ many memory states and let ρ be a play that starts in v and is consistent with σ. We implement a strategy for Player 1 using a counter κ that is initialized with one if $v = v_{\text{req}}$ and with zero otherwise. Moreover, we increment κ upon each visit to v_{req}. After each visit to v_{req}, the strategy τ prescribes moving from v'_{ans} to $v'_{\text{req},1}$ for the first κ visits to v'_{ans}, and it prescribes moving from v'_{ans} to v_{ans} at the $\kappa + 1$-th visit to v'_{ans}. Hence, after the $\kappa + 1$-th visit to v'_{ans}, the vertex v_{req} is visited again, κ is incremented and the behavior of τ described above repeats with incremented κ.

Let ρ be the unique play consistent with σ and τ. Since σ is winning for Player 0, the play ρ does not have a suffix of the form $(v_{\text{del}})^\omega$. Hence, playing consistently with τ, Player 1 enforces a play that starts with a (possibly empty) finite prefix that ends before the first visit to v_{req} and infinitely many rounds, where each round starts in v_{req}. The j-th round is of the form $v_{\text{req}} \cdot \mathit{\Pi}_{k=0,\dots,j}(v'_{\text{req},1} \cdots v'_{\text{req},n} \cdot (v_{\text{del}})^{\ell_{j,k}} \cdot v'_{\text{ans}}) \cdot v_{\text{ans}}$.

We first show $\ell_{j,k} < nW + 1$ for all j, k. Towards a contradiction, assume $\ell_{j,k} \geq nW + 1$ for some j, k. Since σ is of size less than $nW + 1$, a straightforward pumping argument shows that the play ρ' consisting of the finite prefix of ρ concatenated with the first $j - 1$ rounds of ρ, but ending in the infinite suffix

$$v_{\text{req}} \cdot \mathit{\Pi}_{k'=0,\dots,k-1}(v'_{\text{req},1} \cdots v'_{\text{req},n} \cdot (v_{\text{del}})^{\ell_{j,k'}} \cdot v'_{\text{ans}}) \cdot v'_{\text{req},1} \cdots v'_{\text{req},n} \cdot (v_{\text{del}})^\omega$$

is consistent with σ. This, however, contradicts that σ is a winning strategy for Player 0, as v_{del} has odd color. Hence, $\ell_{j,k} < nW + 1$ for all j, k.

Since each edge $(v'_{\text{req},n'}, v'_{\text{req},n'+1})$ for $1 \leq n' \leq n$ has weight W, we obtain $w(v'_{\text{req},1} \cdot (v_{\text{del}})^{\ell_{j,k}} \cdot v'_{\text{ans}}) > 0$ for all j, k. This, in turn, implies $w(v_{\text{req}} \cdot \mathit{\Pi}_{k=0,\dots,j}(v'_{\text{req},1} \cdots v'_{\text{req},n} \cdot (v_{\text{del}})^{\ell_{j,k'}} \cdot v'_{\text{ans}}) \geq j + 1$ for each j. Since, as argued above, the play ρ consistent with σ consists of infinitely many rounds, we obtain that for each $b \in \mathbb{N}$ there exist infinitely many requests in ρ that are answered with cost at least b. This contradicts σ being a winning strategy for Player 0. ◻

7 Quality of Strategies

We have shown in the previous section that finite-state strategies of bounded size suffice for Player 0 to win in parity games with weights, while Player 1 clearly requires infinite memory. However, as we are dealing with a quantitative winning condition, we are not only interested in the size of winning strategies, but also in their quality. More precisely, we are interested in an upper bound on the cost of requests that Player 0 can ensure. In this section, we show that he can guarantee an exponential upper bound on such costs. Dually, Player 1 is required to unbound the cost of responses.

Theorem 3. *Let $\mathcal{G}$ be a parity game with weights with n vertices, d colors, and largest absolute weight W.*

There exists a $b \in \mathcal{O}((ndW)^2)$ and a strategy σ for Player 0 such that, for all plays ρ beginning in $\mathcal{W}_0(\mathcal{G})$ and consistent with σ, we have $\limsup_{j\to\infty} \mathrm{Cor}(\rho, j) \leq b$. This bound is tight.

We first show that Player 0 can indeed ensure an upper bound as stated in Theorem 3. We obtain this bound via a straightforward pumping argument leveraging the upper bound on the size of winning strategies obtained in Lemma 7 .

Lemma 9. *Let $\mathcal{G}$, n, d, and W, and v be as in the statement of Theorem 3 and let $s = d(6n)(d+2)(W+1)$. Player 0 has a winning strategy σ such that, for each play ρ that starts in $\mathcal{W}_0(\mathcal{G})$ and is consistent with σ, we have $\limsup_{j\to\infty} \mathrm{Cor}(\rho, j) \leq nsW$.*

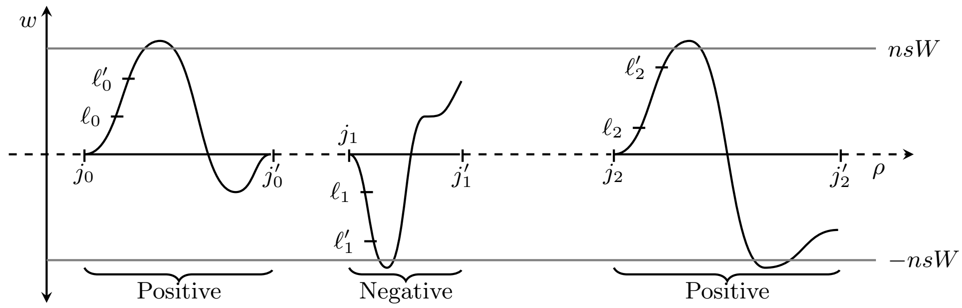

Fig. 7. Illustration of the approach to the proof of Lemma 9.

Proof. We illustrate our approach as well in Figure 7 Let σ be a winning strategy for Player 0 in $\mathcal{G}$ from $\mathcal{W}_0(\mathcal{G})$ of size at most s. Due to Lemma 7, such a strategy exists. Let $\rho = v_0 v_1 v_2 \cdots$ be a play that starts in $\mathcal{W}_0(\mathcal{G})$ and is consistent with σ. We call a position $j \in \mathbb{N}$ **sumptuous** if $nsW < \mathrm{Cor}(\rho, j) < \infty$. Each sumptuous position j has some odd color c, and the request for c posed by visiting v_j is eventually answered.

Assume towards a contradiction that there exist infinitely many sumptuous positions. We define a sequence of positions that begins with the first sumptuous position j_0. For each j_k, we define j'_k as the minimal position with $\Omega(v_{j'_k}) \in \mathrm{Ans}(\Omega(v_{j_k}))$. Finally, each j_k is the smallest sumptuous position greater than j'_{k-1}. The sequence $p = j_0 < j'_0 < j_1 < j'_1 < j_2 < j'_2 < \cdots$ is infinite: There exist infinitely many sumptuous positions by assumption. Moreover, the requests posed at almost all sumptuous positions are eventually answered. Otherwise the play ρ would violate the parity condition, which would contradict that ρ is consistent with the winning strategy σ for Player 0.

Due to the definition of sumptuous positions and the j'_k, we have $\mathrm{Ampl}(v_{j_k} \cdots v_{j'_k}) > nsW$ for each $k \in \mathbb{N}$. Since ρ is consistent with the finite-state strategy σ of size s, in each such $v_{j_k} \cdots v_{j'_k}$ there exists an infix that can be repeated arbitrarily often while retaining consistency with σ. To identify such infixes, we separate the sumptuous positions j_k into two groups: We call a position j_k **positively sumptuous** if there exists a j' with $j_k \leq j' \leq j'_k$ such that $w(v_{j_k} \cdots v_{j'}) > nsW$ and **negatively sumptuous** otherwise.

Let σ be implemented by $(M, \mathrm{init}, \mathrm{upd})$. As each edge contributes cost at most W to $\mathrm{Ampl}(v_{j_k} \cdots v_{j'_k})$, this implies that there exist positions ℓ_k and ℓ'_k with $j_k < \ell_k < \ell'_k < j'_k$ such that

- $v_{\ell_k} = v_{\ell'_k}$,
- $\mathrm{upd}^+(v_0 \cdots v_{\ell_k}) = \mathrm{upd}^+(v_0 \cdots v_{\ell'_k})$,
- $w(v_{\ell_k} \cdots v_{\ell'_k - 1}) > 0$, if j_k is positively sumptuous, and such that
- $w(v_{\ell_k} \cdots v_{\ell'_k - 1}) < 0$, if j_k is negatively sumptuous.

The positions j_k, ℓ_k, ℓ'_k, and j'_k split ρ into infixes $\rho = \Pi_{k=0,1,2,\dots} \pi_{k,I} \cdot \pi_{k,II} \cdot \pi_{k,III} \cdot \pi_{k,IV}$, where $\pi_{k,I}$, $\pi_{k,II}$, $\pi_{k,III}$, and $\pi_{0,IV}$ start at j_k, ℓ_k, ℓ'_k, and j'_k, respectively. Due to the definition of ℓ_k and ℓ'_k, the play $\rho' = \Pi_{k=0,1,2,\dots} \pi_{k,I} \cdot (\pi_{k,II})^k \cdot \pi_{k,III} \cdot \pi_{k,IV}$ is consistent with σ. The costs-of-response of the requests opened by visiting the v_{j_k}, however, diverge due to $|w(\pi_{k,II})| = |w(v_{\ell_k} \cdots v_{\ell'_k - 1})| > 0$. Hence, ρ' violates the parity condition with weights, which contradicts that σ is a winning strategy of Player 0. □

Having thus shown that Player 0 can indeed ensure an exponential upper bound on the incurred cost, we now proceed to show that this bound is tight. A simple example shows that there exists a series of parity games with weights, in which Player 0 wins from every vertex, but in which he cannot enforce a sub-exponential cost of any request.

Lemma 10. *Let $n \in \mathbb{N}$. There exists a parity game with weights $\mathcal{G}_n$ with $|\mathcal{G}_n| \in \mathcal{O}(n)$ and a vertex $v \in \mathcal{W}_0(\mathcal{G})$, such that for each winning strategy for Player 0 from v there exists a play ρ starting in v and consistent with σ with* $\limsup_{j \to \infty} \mathrm{Cor}(\rho, j) \geq (n-1)2^n$.

Proof. We show the game $\mathcal{G}_n$ in Figure 8. The arena of $\mathcal{G}_n$ is a cycle with n vertices of Player 1, where each edge has weight 2^n. Moreover, one vertex is labeled with color two, its directly succeeding vertex is labeled with color one. All remaining vertices have color zero.

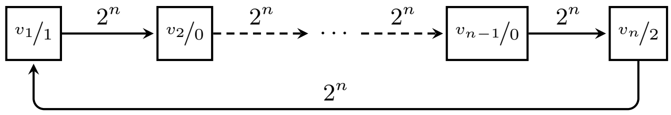

Fig. 8. The game $\mathcal{G}_n$ witnessing an exponential lower bound on the cost that Player 0 can ensure.

Player 0 only has a single strategy in this game and there exist only n plays in $\mathcal{G}_n$, each starting in a different vertex of $\mathcal{G}_n$. In each play, each request for color one is only answered after $n-1$ steps, each contributing a cost of 2^n. Hence, this request incurs a cost of $(n-1)2^n$. Moreover, as this request is posed and answered infinitely often in each play, we obtain the desired result. □

8 From Energy Parity Games to (Bounded) Parity Games with Weights

We have discussed in Sections 4 and 5 how to solve parity games with weights via solving bounded parity games with weights and how to solve the latter games by solving energy parity games, both steps with a polynomial overhead. An obvious question is whether one can also solve energy parity game by solving (bounded) parity games with weights. In this section, we answer this question affirmatively. We show how to transform an energy parity game into a bounded parity game with weights so that solving the latter also solves the former. Then, we show how to transform a bounded parity game with weights into a parity game with weights with the same relation: Solving the latter also solves the former. Both constructions here are gadget based and increase the size of the arenas only linearly. Hence, all three types of games considered here are interreducible with at most polynomial overhead.

8.1 From Energy Parity Games to Bounded Parity Games with Weights

Note that, in an energy parity game, Player 0 wins if the energy increases without a bound, as long as there is a lower bound. However, in a bounded parity game, he has to ensure an upper and a lower bound. Thus, we show in a first step how to modify an energy parity game so that Player 0 still has to ensure a lower bound on the energy, but can also *throw away* unnecessary energy during each transition, thereby also ensuring an upper bound. The most interesting part of this construction is to determine when energy becomes unnecessary to ensure a lower bound. Here, we rely on Lemma 3.

Formally, let $\mathcal{G} = (\mathcal{A}, \Omega, w)$ be an energy parity game with $\mathcal{A} = (V, V_0, V_1, E)$ where we assume w.l.o.g. that the minimal color in $\Omega(V)$ is strictly greater than 1. Now, we define $\mathcal{G}' = (\mathcal{A}', \Omega', w')$ with $\mathcal{A} = (V, V_0, V_1, E)$ where

- $V' = V \cup E$, $V'_0 = V_0 \cup E$, and $V'_1 = V_1$,
- $E' = \{(v, e), (e, e), (e, v') \mid e = (v, v') \in E\}$,
- $\Omega'(v) = \Omega(v)$ and $\Omega'(e) = 1$, and
- $w'(v, e) = w(e)$, $w'(e, e) = -1$, and $w(e, v') = 0$ for every $e = (v, v') \in E$.

Intuitively, every edge of $\mathcal{A}$ is subdivided and a new vertex for Player 0 is added, where he can decrease the energy level. The negative weight ensures that he eventually leaves this vertex in order to satisfy an energy condition.

We say that a strategy σ for Player 0 in $\mathcal{A}'$ is corridor-winning for him from some $v \in V$, if there is a $b \in \mathbb{N}$ such that every play ρ that starts in v and is consistent with σ satisfies the parity condition and $\mathrm{Ampl}(\rho) \leq b$. Hence, instead of just requiring a lower bound on the energy level as in the energy parity condition, we also require a uniform upper bound on the energy level (where we w.l.o.g. assume these bounds to coincide).

Lemma 11. *Let $\mathcal{G}$ and $\mathcal{G}'$ be as above and let $v \in V$. The following are equivalent:*

1. *Player 0 has a winning strategy for $\mathcal{G}$ from v.*
2. *Player 0 has a corridor-winning strategy for $\mathcal{G}'$ from v.*

Proof. 1 ⇒ 2: Due to Proposition 2, Player 0 has a winning finite-state strategy σ for $\mathcal{G}$ from v, say of size s. Furthermore, there is an initial credit c_0 such that every play prefix π that starts in v and is consistent with σ satisfies $w(\pi) \geq -c_0$. Finally, define $b = Wns$, where n and W are defined as in Lemma 3.

We define a strategy σ' for Player 0 in $\mathcal{G}'$ such that it mimics the behavior of σ and additionally ensures that the energy level of a play prefix never exceeds b by at most W. Formally, consider a play prefix π' in $\mathcal{G}'$ starting in v. If π' ends in some $v' \in V$, then we define $\sigma'(\pi) = (v', \sigma(f(\pi)))$ where f is the homomorphism from $V' = V \cup E$ to V induced by $f(v) = v$ and $f(e) = \varepsilon$. On the other hand, assume π' ends in some $e = (v', v'') \in E$. If $w(\pi') \leq b$, then we define $\sigma'(\pi') = v''$, otherwise, we define $\sigma'(\pi') = e$, i.e., the self-loop is used until the energy level of the play prefix is exactly b. This completes the definition of σ'.

Now, consider a play ρ' in $\mathcal{G}'$ that starts in v and is consistent with σ'. By definition of σ', ρ' visits infinitely many vertices in V. Hence, by construction of $\mathcal{A}$, $\rho = f(\rho')$ is a play in $\mathcal{G}$ that starts in v. Further, ρ is consistent with σ, as σ' mimics σ. Hence, $f(\rho')$ satisfies the parity condition. As the vertices removed from ρ' all have color one, and as all colors in $\Omega(V)$ are greater than one, we conclude that ρ' satisfies the parity condition as well.

To conclude, we show that every prefix π' of ρ' satisfies $-c_0 \leq w'(\pi') \leq b + W$. This implies that σ' is indeed a corridor-winning strategy from v. The upper bound $b + W$ is satisfied by construction of σ': As soon as the weight exceeds b, it is decreased to b by the strategy. As this correction happens after each transition, the bound b can be exceeded by at most W, the largest absolute weight of an edge.

To conclude, we consider two cases: first, assume ρ' has no prefix whose energy level exceeds b, then we have $w'(\pi') = w(f(\pi)) \geq -c_0$ for every prefix π' of ρ'. Second, assume ρ' has at least one prefix whose weight exceeds b. We show that every longer prefix has non-negative weight, which concludes the proof.

Towards a contradiction, assume there is a longer suffix with negative weight. Then, there is an infix of ρ' of weight strictly smaller than $-b$, such that Player 0 never uses a self-loop in $\mathcal{A}'$ to throw away energy. Hence, $f(\rho)$ also has an infix with weight strictly smaller than $-b$. This, however, contradicts Lemma 3.

2 ⇒ 1: Let σ' be a corridor-winning strategy for Player 0 in $\mathcal{G}'$ from v. Further, let f be defined as above.

We define a strategy σ for Player 0 from v in $\mathcal{G}$ that is obtained by simulating play prefixes in $\mathcal{G}'$. To this end, we again use a simulation function h that maps a play prefix $v_0 \cdots v_j$ in $\mathcal{G}$ that starts in v and is consistent with σ to a play prefix $h(v_0 \cdots v_j)$ in $\mathcal{G}'$ that starts in v, is consistent with σ', and ends in v_j.

Hence, we define $h(v) = v$. Now, assume we have a play prefix $v_0 \cdots v_j$ in $\mathcal{G}$ that starts in v and is consistent with σ. From our construction, we obtain a play prefix $h(v_0 \cdots v_j)$ in $\mathcal{G}'$ that starts in v, is consistent with σ', and ends in v_j. If $v_j \in V_0 \subseteq V_0'$, then let $\sigma'(h(v_0 \cdots v_j)) = (v_j, v_{j+1})$. We define $\sigma(v_0 \cdots v_j) = v_{j+1}$, which is a legal move due to the construction of $\mathcal{A}'$. If $v_j \in V_1$, then let v_{j+1} be an arbitrary successor of v_j in $\mathcal{A}$.

In both cases, we have to define $h(v_0 \cdots v_j v_{j+1})$. As σ' is a corridor-winning strategy for Player 0 from v in $\mathcal{G}'$, there is a unique play of the form $h(v_0 \cdots v_j)(v_j, v_{j+1})^m v_{j+1}$ that is consistent with σ'. We define $h(v_0 \cdots v_j v_{j+1})$ to be equal to this play, which satisfies the properties required above.

Let b the uniform bound on the amplitude of plays in $\mathcal{G}'$ consistent with σ' starting in v. Now, fix a play ρ in $\mathcal{G}$ starting in v and consistent with σ. Furthermore, let ρ' be the limit of the $h(\pi)$ for increasing prefixes of ρ. By construction, ρ' starts in v as well and is consistent with σ'. Hence, ρ' visits infinitely many vertices from V and never gets stuck in a self-loop throwing away energy. This implies $f(\rho') = \rho$. Furthermore, as ρ' satisfies the parity condition, ρ does as well: the colors removed by applying f are inconsequential in this situation.

Let π_j be the prefix of length j of ρ. A straightforward induction proves that the energy level of π_j is greater or equal to that of $h(\pi_j)$. As the latter is bounded from below by b, we conclude that σ is winning for Player 0 in $\mathcal{G}$ from v with initial credit b. ☐

Now, we turn $\mathcal{G}'$ into a bounded parity game with weights. In such a game, the cost-of-response of every request has to be bounded, but the overall energy level of the play may still diverge to $-\infty$. To rule this out, we open one unanswerable request at the beginning of each play, which has to be unanswered with finite cost in order to satisfy the bounded parity condition with weights. If this is the case, then the energy level of the play is always in a bounded corridor, i.e., we obtain a corridor-winning strategy.

Formally, for every vertex $v \in V$, we add a vertex $\overline{v}$ to $\mathcal{A}'$ of an odd color c^* that is larger than every color in $\Omega(V)$, i.e., the request can never be answered. Furthermore, $\overline{v}$ has a single outgoing edge to v of weight 0, i.e., it is irrelevant whose turn it is. Call the resulting arena $\mathcal{A}''$, the resulting coloring Ω'', and the resulting weighting w'', and let $\mathcal{G}'' = (\mathcal{A}'', \text{BndWeightParity}(\Omega'', w''))$.

Lemma 12. *Let $\mathcal{G}'$ and $\mathcal{G}''$ be as above and let $v \in V$. The following are equivalent:*

1. *Player 0 has a corridor-winning strategy for $\mathcal{G}'$ from v.*
2. *$\overline{v} \in \mathcal{W}_0(\mathcal{G}'')$.*

Proof. 1 ⇒ 2: Let σ' be a corridor-winning strategy for Player 0 in $\mathcal{G}'$ from v. Further, let b be the corresponding uniform bound on the amplitude of plays that start in v and are consistent with σ'. We define a strategy σ'' for Player 0 from $\overline{v}$ via $\sigma''(\overline{v}\pi) = \sigma'(\pi)$.

Let $\overline{v}\rho$ be a play that is consistent with σ''. By construction, ρ starts in v and is consistent with σ'. Hence, it satisfies the parity condition and its amplitude is bounded by b. Thus, almost all requests in ρ are answered with cost at most b and there is no unanswered request of infinite cost. This implies that $\overline{v}\rho$ satisfies the bounded parity condition with weights. Hence, $\overline{v} \in \mathcal{W}_0(\mathcal{G}'')$.

2 ⇒ 1: Let σ'' and b be a winning strategy for Player 0 in $\mathcal{G}''$ from $\overline{v}$ and a bound such that every request in a play starting in $\overline{v}$ and consistent with σ'' is answered or unanswered with cost at most b. Due to Corollary 1, such a strategy σ'' exists. We define a strategy σ' for Player 0 from v in $\mathcal{G}'$ via $\sigma'(\pi) = \sigma''(\overline{v}\pi)$.

Let ρ be a play starting in v that is consistent with σ'. By construction, $\overline{v}\rho$ is consistent with σ''. Hence, $\overline{v}\rho$ satisfies the parity condition and every request is answered or unanswered with cost at most b. In particular, this holds true for the unanswered request posed by visiting $\overline{v}$. Hence, the amplitude of $\overline{v}\rho$ (and thus also that of ρ) is bounded by b.

Thus, ρ satisfies the parity condition and the energy level of all its prefixes is between $-b$ and b. As ρ is picked arbitrarily, we have that σ' is corridor-winning from v. ☐

8.2 From Bounded Parity Games with Weights to Parity Games with Weights

Next, we show how to turn a bounded parity game with weights into a parity game with weights so that solving the latter also solves the former. The construction here uses the same restarting mechanism that underlies the proof of Lemma 1: as soon as a request has incurred a cost of b, restart the play and enforce a request of cost $b+1$, and so on. Unlike the proof of Lemma 1, where Player 1 could restart the play at any vertex, here we always have to return to a fixed initial vertex we are interested in. While resetting, we have to answer all requests in order to prevent Player 1 to use the reset to prevent requests from being answered. Assume $v^* \in V$ is the initial vertex we are interested in. Then, we subdivide every edge in $\mathcal{A}''$ to allow Player 1 to restart the play by answering all open requests and then moving back to v^*.

Formally, fix a bounded parity game with weights $\mathcal{G} = (\mathcal{A}, \text{BndWeightParity}(\Omega, w))$ with $\mathcal{A} = (V, V_0, V_1, E)$ and a vertex $v^* \in V$. We define the parity game with weights $\mathcal{G}_{v^*} = (\mathcal{A}_{v^*}, \text{WeightParity}(\Omega_{v^*}, w_{v^*}))$ with $\mathcal{A}_{v^*} = (V', V_0', V_1', E')$ where

- $V' = V \cup E \cup \{\top\}$, $V_0' = V_0$, and $V_1' = V_1 \cup E \cup \{\top\}$,
- $E' = \{(v, e), (e, \top), (e, v') \mid e = (v, v') \in E\} \cup \{(\top, v^*)\}$,
- $\Omega_{v^*}(v) = \Omega(v)$, $\Omega_{v^*}(e) = 0$ for every $e \in E$, and $\Omega_{v^*}(\top) = 2\max(\Omega(V))$, and
- $w_{v^*}(v, e) = w(e)$ for $(v, e) \in V \times E$ and $w_{v^*}(e) = 0$ for every other edge $e \in E'$.

Lemma 13. *Let $\mathcal{G}$ and $\mathcal{G}_{v^*}$ as above. The following are equivalent:*

1. *$v^* \in \mathcal{W}_0(\mathcal{G})$.*
2. *$v^* \in \mathcal{W}_0(\mathcal{G}_{v^*})$.*

Proof. 1 ⇒ 2: Let σ be a winning strategy for Player 0 for $\mathcal{G}$ from v^*. Due to Corollary 1 we can assume that there is a b, such that every request in a play that starts in v^* and is consistent with σ is answered or unanswered with cost at most b.

We define a strategy for Player 0 in $\mathcal{G}_{v^*}$. Given a play prefix π' in $\mathcal{A}_{v^*}$ let $\text{sfx}_\top(\pi')$ be the longest suffix of π' that does not contain the vertex $\top$. Hence, if π' starts in v^*, then $\text{sfx}_\top(\pi')$ starts in v^* as well, as v^* is the only successor of $\top$. Further, let f be the homomorphism from $V \cup E$ to V induced by $f(v) = v$ for $v \in V$ and $f(e) = \varepsilon$ for $e \in E$. Note that, if π' is a play (prefix) in $\mathcal{A}_{v^*}$ that does not visit $\top$, then $f(\pi')$ is a play prefix in $\mathcal{A}$ of the same weight and with the same sequence of colors (save for the occurrences of color zero at the deleted vertices, which is inconsequential for our condition considered here, since it is the smallest color).

Now, let π' be a play prefix in $\mathcal{A}_{v^*}$ that ends in a vertex of Player 0, say in $v \in V'_0 = V_0$. Then, we define $\sigma'(\pi') = (v, \sigma(f(\mathrm{sfx}_\top(\pi'))))$. It remains to prove that σ' is winning for Player 0 in $\mathcal{G}_{v^*}$ from v^*. To this end, let ρ' be a play in $\mathcal{A}_{v^*}$ starting in v^* and consistent with σ'. We consider two cases, depending on how often the vertex $\top$ is visited by ρ'.

If $\top$ is visited only finitely often, then let ρ be the suffix of ρ' starting after the last occurrence of $\top$. By construction of σ', $f(\rho)$ is a play in $\mathcal{A}$ starting in v^* and consistent with σ. Hence, it satisfies the bounded parity condition with weights. Hence, due to $\mathrm{BndWeightParity}(\Omega, w) \subseteq \mathrm{WeightParity}(\Omega, w)$, and since the bounded parity condition with weights is 0-extendable, we conclude that $f(\rho')$ satisfies the parity condition with weights as well. Finally, by construction of $\mathcal{A}_{v^*}$, ρ' satisfies the parity condition with weights, too.

Now, assume $\top$ is visited infinitely often. Then, we can decompose ρ' into $\pi'_0 \top \pi'_1 \top \pi'_2 \top \cdots$, so that each π'_j does not visit $\top$. Hence, by definition of σ', each of the $f(\pi'_j)$ starts in v^* and is consistent with σ. Furthermore, every request in some π'_j is answered by the next visit to $\top$. Thus, it suffices to show that the cost-of-response of all requests in ρ' is bounded. This follows immediately from the fact that σ allows answered or unanswered requests of cost at most b in $f(\pi'_j)$. This property is inherited by the π'_j by construction of $\mathcal{A}_{v^*}$.

So, in both cases $\rho' \in \mathrm{WeightParity}(\Omega, w)$, i.e., σ' has the desired properties.

$2 \Rightarrow 1$: We proceed by contraposition. Due to the determinacy of both games, it suffices to show that $v^* \in \mathcal{W}_1(\mathcal{G})$ implies $v^* \in \mathcal{W}_1(\mathcal{G}_{v^*})$. Hence, let τ be a winning strategy for Player 1 in $\mathcal{G}$ from v. Further, let $\mathrm{sfx}_\top$ and f be defined as above.

Now, we define a strategy τ' for Player 1 from v^* in $\mathcal{G}_{v^*}$ that is controlled by a counter κ, which is initialized with zero, and which is incremented during a play every time the costs of some request exceed κ. We construct our strategy such that each time κ is updated, Player 1 restarts the play by moving to $\top$ and then to v^*.

Assume we have a play prefix π' in $\mathcal{A}_{v^*}$ that ends in a vertex of Player 1 and have to define $\tau'(\pi')$. We consider several cases depending on the last vertex of π'. If π' ends in $\top$, then we define $\tau'(\pi') = v^*$, which is the only successor of $\top$.

If π' ends in $v \in V_1 \subseteq V'_1$, then we define $\tau'(\pi') = (v, \tau(f(\mathrm{sfx}_\top(\pi'))))$, i.e., we discard everything up to and including the last occurrence of $\top$. Finally, if π' ends in $e = (v, v') \in E \subseteq V'_1$, then we consider two cases. Let κ be the current counter value. If $\mathrm{sfx}_\top(\pi')$ contains a request such that the remaining part of π' that starts at this request has amplitude greater than κ, then we define $\tau'(\pi') = \top$ and increment κ. Otherwise, we define $\tau'(\pi') = v'$ and leave κ unchanged.

It remains to show that τ' is winning in $\mathcal{G}_{v^*}$ from v^*. To this end, let ρ' be a play in $\mathcal{G}_{v^*}$ that starts in v^* and is consistent with τ'. If ρ' visits $\top$ infinitely often, then ρ' contains, for every $b \in \mathbb{N}$, a (different) request that is answered or unanswered with cost at least b. Hence, ρ' violates the parity condition with costs.

Finally, if ρ' visits $\top$ only finitely often, then there is a $b \in \mathbb{N}$ (the final value of κ, which is only finitely often incremented in this case) such that every request in ρ' is answered or unanswered with cost at most b. Furthermore, let ρ be the suffix of ρ' that starts after the last occurrence of $\top$. As in the previous case, $f(\rho)$ is a play in $\mathcal{A}$ that starts in v^* and is consistent with τ. As ρ and $f(\rho)$ have essentially the same evolution of the weights (save for the removed edges of weight zero) and the same color sequence (save for the removed vertices of color zero), every request in $f(\rho)$ is answered or unanswered with cost at most b. However, as ρ is consistent with τ, it violates the bounded parity condition with weights. This is, in this situation, only possible by violating the parity condition. Hence ρ, and thus also ρ', violates the parity condition as well. Therefore, ρ' in particular violates the parity condition with weights.

In both cases, ρ' is winning for Player 1, i.e., τ' has the desired properties. □

9 Conclusions and Future Work

We have established that parity games with weights and bounded parity games fall into the same complexity class as energy parity games. This is interesting, because, while the complexity of solving such games have the signature complexity class NP ∩ co-NP, they are not yet considered a class in their own right. It is also interesting because the properties appear to be inherently different: While they both combine the qualitative parity condition with quantified costs, parity games with weights *combine* these aspects on the property level, whereas energy parity games simply look at the combined—and totally unrelated—properties. We show the characteristic properties of parity games and of games with combinations of a parity condition with quantitative conditions relevant for this work in Table 1.

	Complexity	Memory Player 0	Memory Player 1	Bounds
Parity Games [3]	quasi-polynomial	pos.	pos.	–
Energy Parity Games [4]	NP ∩ co-NP	$\mathcal{O}(ndW)$	pos.	$\mathcal{O}(nW)$
Finitary Parity Games [7]	polynomial	pos.	inf.	$\mathcal{O}(nW)$
Parity Games with Costs [13]	NP ∩ co-NP	pos.	inf.	$\mathcal{O}(nW)$
Parity Games with Weights	NP ∩ co-NP	$\mathcal{O}(nd^2W)$	inf.	$\mathcal{O}((ndW)^2)$

Table 1. Characteristic properties of variants of parity games.

As future work, we are looking into the natural extensions of parity games with weights to Streett games with weights [7, 13], and at the complexity of determining optimal bounds and strategies that obtain them [28]. We are also looking at variations of the problem. The two natural variations are

– to use a one-sided definition (instead of the absolute value) for the amplitude of a play, i.e., using $\text{Ampl}(\pi) = \sup_{j<|\pi|} w(v_0 \cdots v_j) \in \mathbb{N}_\infty$ (instead of $\text{Ampl}(\pi) = \sup_{j<|\pi|} |w(v_0 \cdots v_j)| \in \mathbb{N}_\infty$), and
– to use an arbitrary consecutive subsequence of a play, i.e., $\text{Ampl}(\pi) = \sup_{j\le k<|\pi|} |w(v_j \cdots v_k)| \in \mathbb{N}_\infty$.

There are good arguments in favor and against using these individual variations—and their combination to $\text{Ampl}(\pi) = \sup_{j\le k<|\pi|} w(v_j \cdots v_k) \in \mathbb{N}_\infty$—but we feel that the introduction of parity games with weights benefit from choosing one of the four combinations as *the* parity games with weights.

We expect the complexity to rise when changing from maximizing over the absolute value to maximizing over the value, as this appears to be close to pushdown boundedness games [5], and we conjecture this problem to be PSPACE complete.

References

1. Björklund, H., Vorobyov, S.: A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. Discrete Appl. Math. 155(2), 210–229 (2007)
2. Browne, A., Clarke, E.M., Jha, S., Long, D.E., Marrero, W.R.: An improved algorithm for the evaluation of fixpoint expressions. Theor. Comput. Sci. 178(1–2), 237–255 (1997)
3. Calude, C.S., Jain, S., Khoussainov, B., Li, W., Stephan, F.: Deciding parity games in quasipolynomial time. In: Proc. of STOC 2017. pp. 252–263. ACM Press (2017)
4. Chatterjee, K., Doyen, L.: Energy Parity Games. TCS 458, 49–60 (2012)
5. Chatterjee, K., Fijalkow, N.: Infinite-state games with finitary conditions. In: Computer Science Logic 2013 (CSL 2013). Leibniz International Proceedings in Informatics (LIPIcs), vol. 23, pp. 181–196. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2013)
6. Chatterjee, K., Henzinger, T.A.: Finitary winning in omega-regular games. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 257–271 (2006)
7. Chatterjee, K., Henzinger, T.A., Horn, F.: Finitary winning in omega-regular games. Trans. Comput. Log. 11(1), 1:1–1:27 (2009)
8. Chatterjee, K., Henzinger, T.A., Jurdzinski, M.: Mean-payoff parity games. In: Proc. of LICS. pp. 178–187. IEEE Computer Society (2005)
9. Condon, A.: On algorithms for simple stochastic games. In: Advances in Computational Complexity Theory. pp. 51–73. American Mathematical Society (1993)
10. Emerson, E.A., Jutla, C.S.: Tree automata, μ-calculus and determinacy. In: FOCS. pp. 368–377 (1991)
11. Emerson, E.A., Lei, C.L.: Efficient model checking in fragments of the propositional μ-calculus. In: LICS. pp. 267–278 (1986)

12. Fearnley, J., Jain, S., Schewe, S., Stephan, F., Wojtczak, D.: An ordered approach to solving parity games in quasi polynomial time and quasi linear space. In: Proc. of SPIN. pp. 112–121. ACM (2017)
13. Fijalkow, N., Zimmermann, M.: Parity and streett games with costs. LMCS 10(2) (2014), https://doi.org/10.2168/LMCS-10(2:14)2014
14. Jurdziński, M.: Deciding the winner in parity games is in UP ∩ co-UP. Information Processing Letters 68(3), 119–124 (November 1998)
15. Jurdziński, M.: Small progress measures for solving parity games. In: STACS. LNCS, vol. 1770, pp. 290–301 (2000)
16. Jurdziński, M., Lazić, R.: Succinct progress measures for solving parity games. In: Proc. of LICS 2017. pp. 1–9. IEEE Computer Society (2017)
17. Jurdziński, M., Paterson, M., Zwick, U.: A deterministic subexponential algorithm for solving parity games. SIAM Journal on Computing 38(4), 1519–1532 (2008)
18. Kozen, D.: Results on the propositional μ-calculus. TCS 27, 333–354 (1983)
19. Lehtinen, K.: A modal μ perspective on solving parity games in quasipolynomial time. In: Proc. of LICS 2018. p. (to appear) (2018)
20. Martin, D.A.: Borel determinacy. Annals of Mathematics 102(2), 363–371 (1975), http://www.jstor.org/stable/1971035
21. McNaughton, R.: Infinite games played on finite graphs. Ann. Pure Appl. Logic 65(2), 149–184 (1993)
22. Nerode, A., Remmel, J.B., Yakhnis, A.: Mcnaughton games and extracting strategies for concurrent programs. Annals of Pure and Applied Logic 78(1-3), 203–242 (1996), https://doi.org/10.1016/0168-0072(95)00032-1
23. Puri, A.: Theory of hybrid systems and discrete event systems. Ph.D. thesis, Computer Science Department, University of California, Berkeley (1995)
24. Schewe, S.: An optimal strategy improvement algorithm for solving parity and payoff games. In: CSL. LNCS, vol. 5213, pp. 368–383 (2008)
25. Schewe, S.: Solving parity games in big steps. Journal of Computer and System Sciences 84, 243–262 (2017)
26. Schewe, S., Trivedi, A., Varghese, T.: Symmetric strategy improvement. In: ICALP. LNCS, vol. 9135, pp. 388–400 (2015)
27. Vöge, J., Jurdziński, M.: A discrete strategy improvement algorithm for solving parity games. In: Proceedings of the CAV. pp. 202–215. Springer-Verlag (July 2000)
28. Weinert, A., Zimmermann, M.: Easy to win, hard to master: Optimal strategies in parity games with costs. Logical Methods in Computer Science 13(3) (2017), https://doi.org/10.23638/LMCS-13(3:29)2017
29. Zielonka, W.: Infinite games on finitely coloured graphs with applications to automata on infinite trees. Theor. Comput. Sci. 200(1-2), 135–183 (1998)
30. Zwick, U., Paterson, M.S.: The complexity of mean payoff games on graphs. Theoretical Computer Science 158(1–2), 343–359 (1996)