

EASY TO WIN, HARD TO MASTER: OPTIMAL STRATEGIES IN PARITY GAMES WITH COSTS

ALEXANDER WEINERT AND MARTIN ZIMMERMANN

Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany
e-mail address: weinert@react.uni-saarland.de

Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany
e-mail address: zimmermann@react.uni-saarland.de

ABSTRACT. The winning condition of a parity game with costs requires an arbitrary, but fixed bound on the cost incurred between occurrences of odd colors and the next occurrence of a larger even one. Such games quantitatively extend parity games while retaining most of their attractive properties, i.e., determining the winner is in NP and co-NP and one player has positional winning strategies.

We show that the characteristics of parity games with costs are vastly different when asking for strategies realizing the minimal such bound: The solution problem becomes PSPACE-complete and exponential memory is both necessary in general and always sufficient. Thus, solving and playing parity games with costs optimally is harder than just winning them. Moreover, we show that the tradeoff between the memory size and the realized bound is gradual in general. All these results hold true for both a unary and binary encoding of costs.

Moreover, we investigate Streett games with costs. Here, playing optimally is as hard as winning, both in terms of complexity and memory.

1. INTRODUCTION

Recently, the focus of research into infinite games for the synthesis of reactive systems moved from studying qualitative winning conditions to quantitative ones. This paradigm shift entails novel research questions, as quantitative conditions induce a (partial) ordering of winning strategies. In particular, there is a notion of semantic optimality for strategies which does not appear in the qualitative setting. Thus, in the quantitative setting, one can ask whether computing optimal strategies is harder than computing arbitrary ones, whether optimal strategies are necessarily larger than arbitrary ones, and whether there are tradeoffs between different quality measures for strategies, e.g., between the size of the strategy and its semantic quality (in terms of satisfaction of the winning condition).

1998 ACM Subject Classification: D.2.4 Software/Program Verification.

Key words and phrases: Parity Games with Costs, Optimal Strategies, Memory Requirements, Tradeoffs. This is a revised and extended version of work first presented at CSL '16 [WZ16].

Supported by the project “TriCS” (ZI 1516/1-1) of the German Research Foundation (DFG).

As an introductory example consider the classical (max)-parity condition, which is defined for an infinite sequence drawn from a finite subset of the natural numbers, so-called colors. The parity condition is satisfied if almost all occurrences of an odd color are *answered* by a later occurrence of a larger even color, e.g., the sequence

$$\pi = 1\ 0\ 2\ 1\ 0\ 0\ 2\ 1\ 0\ 0\ 0\ 2\ 1\ 0\ 0\ 0\ 0\ 2\ 1\ 0\ 0\ 0\ 0\ 0\ 2\ 1\ 0\ 0\ 0\ 0\ 0\ 0\ 2\ \dots$$

satisfies the parity condition, as every 1 is eventually answered by a 2.

The finitary parity condition [CHH09] is obtained by additionally requiring the existence of a bound b such that almost every odd color is answered within at most b steps, i.e., π does not satisfy the finitary parity condition, as the length of the zero-blocks is unbounded. Thus, solving a finitary parity game is a boundedness problem: In order to satisfy the condition, an arbitrary, but fixed bound has to be met. In particular, winning strategies for finitary parity games are naturally ordered by the minimal bound they realize along all consistent plays. Thus, finitary parity games induce an optimization problem: Compute an optimal winning strategy, i.e., one that guarantees the smallest possible bound.

Other examples for such quantitative winning conditions include mean payoff [EM79, ZP95] and energy [BFL⁺08, CdAHS03] conditions and their combinations and extensions, request-response conditions [HTWZ15, Zim09], parity with costs [FZ14b], and parameterized extensions of Linear Temporal Logic (LTL) [AETP01, FZ14a, KPV09, Zim13, Zim15]. Often, these conditions are obtained by interpreting a classical qualitative winning condition quantitatively, as exemplified by the finitary parity condition introduced above.

Often, the best algorithms for solving such boundedness conditions are as fast as the best ones for their respective qualitative variant, while the fastest algorithms for the optimization problem are worse. For example, solving games with winning conditions in Prompt-LTL, a quantitative variant of LTL, is 2EXPTIME-complete [KPV09] (i.e., as hard as solving classical LTL games [PR89]), while computing optimal strategies is only known to be in 3EXPTIME [Zim13]. The same is true for the bounds on the size of winning strategies, which jump from tight doubly-exponential bounds to triply-exponential upper bounds. The situation is similar for other winning conditions as well, e.g., request-response conditions [HTWZ15]. These examples all have in common that there are no known lower bounds on the complexity and the memory requirements in the optimization variant, except for the trivial ones for the qualitative case. A notable exception are finitary parity games, which are solvable in polynomial time [CHH09] and thus simpler than parity games (according to the state-of-the-art).

In this work, we study optimal strategies in parity games with costs, a generalization of finitary parity games. In this setting, we are able to show that computing optimal strategies is indeed harder than computing arbitrary strategies, and that optimal strategies have exponentially larger memory requirements in general. A parity game with costs is played in a finite directed graph whose vertices are partitioned into those of Player 0 and those of Player 1. Starting at an initial vertex, the players move a token through the graph: If it is placed at a vertex of Player i , then this player has to move it to some successor. Thus, after ω rounds, the players have produced an infinite path through the graph, a so-called play. The vertices of the graph are colored by natural numbers and the edges are labeled by (non-negative) costs. These two labelings induce the parity condition with costs: There has to be a bound b such that almost all odd colors are followed by a larger even color such that the cost incurred between these two positions is at most b . Thus, the sequence π from above satisfies the parity condition with costs, if the cost of the zero-blocks is bounded.

Note that the finitary parity condition is the special case where every edge has cost one and the parity condition is the special case where every edge has cost zero.

Thus, to win a parity game with costs, Player 0 has to bound the costs between requests and their responses along all plays. If Player 0 has any such strategy, then she has a positional strategy [FZ14b], i.e., a strategy that determines the next move based only on the vertex the token is currently at, oblivious to the history of the play. Let n denote the number of vertices of the graph the game is played in and let W denote its largest edge cost. Then, a positional winning strategy uniformly bounds the costs to some bound $b \leq nW$, which we refer to as the cost of the strategy. Furthermore, Mogavero et al. showed that the winner of a parity game with costs can be determined in $\text{UP} \cap \text{CO-UP}$ [MMS15]. All previous work on parity games with costs was concerned with the boundedness variant, i.e., the problems ask to find some bound, but not necessarily the best one. Here, in contrast, we study optimal strategies in parity games with costs.

When considering parity games with costs as a boundedness problem, the actual edge costs can be abstracted away: It is only relevant whether an edge has cost zero or not. Thus, it suffices to consider only costs 0 and 1 (typically denoted as ε and \mathbf{i}). We call this setting one of *abstract costs*.

For the optimization variant, abstracting away the actual costs is no longer valid and encoding larger costs by subdividing them into edges of cost \mathbf{i} , i.e., of cost 1, comes at the price of an exponential blowup in the graph's size. Thus, we also consider the case of costs in \mathbb{N} , given in binary encoding. Here, the upper bound nW on the cost of an optimal strategy might be exponential in the size of the game.

Furthermore, we also study Streett conditions, which generalize parity conditions by relaxing the hierarchical structure of the requests and responses: In a parity condition, a large even color answers requests for *all* smaller odd colors. In contrast, in a Streett condition, any two kinds of responses are potentially independent of each other. It is known that solving the boundedness problem for finitary Streett games and for Streett games with costs is EXPTIME -complete [FZ14b], which has to be compared to the CO-NP -completeness of solving classical Streett games [Hor05]. Furthermore, finite-state strategies of exponential size suffice for Player 0 to implement a winning strategy in a Streett game with costs [FZ14b]. As above, one obtains mnW as upper bound on the cost of an optimal strategy in a Streett game with costs with n vertices and largest cost W , where m is the size of a finite-state winning strategy.

1.1. Our Contribution. Our first four results are concerned with the special case of abstract costs, i.e., costs 0 and 1 only; the remaining ones are about games with costs in \mathbb{N} .

The first result shows that determining whether Player 0 has a strategy in a parity game with costs whose cost is smaller than a given bound b is PSPACE -complete. Thus, computing the bound of an optimal strategy is strictly harder than just deciding whether or not some bound exists (unless $\text{PSPACE} \subseteq \text{UP} \cap \text{CO-UP}$). The hardness result is shown by a reduction from QBF and uses the bound b to require Player 0's strategy to implement a satisfying Skolem function for the formula, where picking truth values is encoded by requests of odd colors. The lower bound is complemented by a polynomial space algorithm that is obtained from an alternating polynomial time Turing machine that simulates a finite-duration variant of parity games with costs that is won by Player 0 if and only if she can enforce a cost of at most b in the original game. To obtain the necessary polynomial upper bound on the play

length we rely on the upper bound n on the optimal bound and on a first-cycle variant of parity games (cf. [AR14]) tailored to our setting.

Our second result concerns memory requirements of optimal strategies. A corollary of the correctness of the finite-duration game yields exponential upper bounds: If Player 0 has a strategy of cost b , then she also has one of cost b and of size $(b+2)^d = 2^{d \log(b+2)}$, where d is the number of odd colors in the game. A similar result holds true for Player 1 as well: If he can exceed a given bound b , then he can also do so with a strategy of size $n(b+2)^d$.

Furthermore, we show that the exponential upper bounds are asymptotically tight: We present a family \mathcal{G}_d of parity games with costs such that \mathcal{G}_d has d odd colors and Player 0 requires strategies of size 2^{d-1} to play optimally in each \mathcal{G}_d . This result is based on using the bound b to require Player 0 to store which odd colors have an open request and in which order they were posed. Our result improves a linear bound presented by Chatterjee and Fijalkow [CF13]. Dually, we present an exponential lower bound on the memory size necessary for Player 1 to exceed a given bound b .

Moreover, we study the tradeoff between memory size and cost of a strategy witnessed by the results above: Arbitrary winning strategies are as small as possible, i.e., positional, but in general have cost n . In contrast, optimal strategies realize a smaller bound, but might have exponential size. Hence, one can trade cost for memory and vice versa. We show that this tradeoff is gradual in the games \mathcal{G}_d : There are strategies $\sigma_1, \sigma_2, \dots, \sigma_d$ such that $1 = |\sigma_1| < |\sigma_2| < \dots < |\sigma_d| = 2^{d-1}$ and $b_1 > b_2 > \dots > b_d$, where b_j is the cost of σ_j . Furthermore, we show that the strategy σ_j has minimal size among all strategies of cost b_j . Equivalently, the strategy σ_j has minimal cost among all strategies whose size is not larger than σ_j 's size.

Both lower bounds we prove and the tradeoff result already hold for the special case of finitary parity games, which can even be solved in polynomial time [CHH09]. Hence, in this case, the gap between just winning and playing optimally is even larger.

After the results for the special case of abstract costs (i.e., 0 and 1 only), we consider the general case of arbitrary non-negative costs given in binary encoding. We show that determining whether Player 0 has a strategy in a parity game with costs whose cost is smaller than a given bound b is still PSPACE-complete, i.e., having *larger* costs does not influence the complexity of the problem. The lower bound on the complexity carries over from the special case of abstract costs but the proof of the upper bound is affected by this generalization: The upper bound on the cost of an optimal strategy is now exponential, which implies that the finite-duration variant has exponentially long plays as well. We devise a shortcut criterion to skip parts of a play and prove that this yields the desired alternating polynomial-time algorithm, which places to problem in PSPACE.

As before, this reduction also yields exponential upper bounds on the necessary memory to implement a winning strategy: The memory requirements do not increase asymptotically when considering arbitrary costs, and they still match the lower bounds.

Finally, we consider quantitative Streett conditions. We show that, given Streett game with costs and a bound b , determining whether Player 0 has a strategy with cost at most b is EXPTIME-complete. Thus, playing quantitative Streett games optimally is not harder than just winning them. This is due to the fact that just winning them is already very hard. Furthermore, we present tight exponential bounds on the memory necessary to implement a winning strategy for Player 0 in such a game. All lower bounds already hold for the special case of finitary Streett games while the upper bounds hold for arbitrary costs encoded in binary.

1.2. Related Work. Tradeoffs in infinite games have been studied before, e.g., in stochastic and timed games, one can trade memory for randomness, i.e., randomized strategies are smaller than deterministic ones [CdAH04, CHP08]. A detailed overview of more recent results in this direction and of tradeoffs in multi-dimensional winning conditions is given in the thesis of Randour [Ran14]. The nature of these results is quite different from ours.

Lang investigated optimal strategies in the resource reachability problem on pushdown graphs [Lan14], where there exists a finite number of counters, which may be increased and reset, but not read during a play. He shows that in order to keep the values of the counters minimal during the play, exponential memory in the number of counters is both necessary and sufficient for Player 0. While the author shows the corresponding decision problem to be decidable, he does not provide a complexity analysis of the problem. Furthermore, the setting of the problem is quite different from the model considered in this work: He considers infinite graphs and multiple counters, but only reachability conditions, while we consider finite graphs and implicit counters tied to the acceptance condition, which is a general parity condition.

Also, Fijalkow et al. proved the non-existence of a certain tradeoff between size and quality of strategies in boundedness games [FHKS15], which refuted a conjecture with important implications for automata theory and logics. Such games are similar to those considered by Lang in that they are played in potentially infinite arenas and have multiple counters.

Recently, Bruyère et al. introduced window-parity games [BHR16], another quantitative variant of parity games, and proved tight complexity bounds for the scenario with multiple colorings of the arena. They were also able to show a tight connection between window-parity and finitary parity games.

Finally, the winning conditions considered here have also been studied in the setting of delay games. In such games, Player 0 may delay her moves to obtain a lookahead on her opponent's moves, thereby gaining an advantage that allows her to win games she loses without delay. Now, there are potential tradeoffs between quality, size, and amount of delay. Most importantly, one can trade delay for quality and vice versa [Zim17] which allows Player 0 to improve the quality of her strategies by taking advantage of the delay.

1.3. Organization of the Paper. In Section 2, we introduce basic definitions about infinite games; in Section 3, we introduce parity games with costs. First, we study the variant with abstract costs: we prove the PSPACE-completeness result (Section 4), the exponential bounds on the memory requirements of optimal strategies (Section 5), and the gradual tradeoff between cost and size of winning strategies (Section 6). Then, we turn our attention to the setting of integer-valued costs in Section 7 and to Streett games with costs in Section 8. Finally, we conclude in Section 9 by discussing further research.

2. PRELIMINARIES

We denote the non-negative integers by \mathbb{N} and define $[n] = \{0, 1, \dots, n-1\}$ for every $n \geq 1$.

An *arena* $\mathcal{A} = (V, V_0, V_1, E, v_I)$ consists of a finite, directed graph (V, E) , a partition $\{V_0, V_1\}$ of V into the positions of Player 0 (drawn as circles) and Player 1 (drawn as rectangles), and an initial vertex $v_I \in V$. The size of \mathcal{A} , denoted by $|\mathcal{A}|$, is defined as $|V|$.

A *play* in \mathcal{A} is an infinite path $\rho = v_0 v_1 v_2 \dots$ through (V, E) starting in v_I . To rule out finite plays, we require every vertex to be non-terminal. A *game* $\mathcal{G} = (\mathcal{A}, \text{Win})$ consists of

an arena \mathcal{A} with vertex set V and a set $\text{Win} \subseteq V^\omega$ of winning plays for Player 0. The set of winning plays for Player 1 is $V^\omega \setminus \text{Win}$.

A *strategy* for Player i is a mapping $\sigma: V^*V_i \rightarrow V$ where $(v, \sigma(wv)) \in E$ for all $wv \in V^*V_i$. We say that σ is *positional* if $\sigma(wv) = \sigma(v)$ for every $wv \in V^*V_i$. We often view positional strategies as a mapping $\sigma: V_i \rightarrow V$. A play $v_0v_1v_2 \cdots$ is *consistent* with a strategy σ for Player i , if $v_{j+1} = \sigma(v_0 \cdots v_j)$ for every j with $v_j \in V_i$. A strategy σ for Player i is a *winning strategy* for \mathcal{G} if every play that is consistent with σ is won by Player i . If Player i has a winning strategy, then we say she wins \mathcal{G} . *Solving* a game amounts to determining its winner.

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

An arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and a memory structure $\mathcal{M} = (M, m_I, \text{Upd})$ for \mathcal{A} induce the expanded arena $\mathcal{A} \times \mathcal{M} = (V \times M, V_0 \times M, V_1 \times M, E', (v_I, m_I))$ where E' is defined via $((v, m), (v', m')) \in E'$ if and only if $(v, v') \in E$ and $\text{Upd}(m, (v, v')) = m'$. Every play $\rho = v_0v_1v_2 \cdots$ in \mathcal{A} has a unique extended play $\text{ext}(\rho) = (v_0, m_0)(v_1, m_1)(v_2, m_2) \cdots$ in $\mathcal{A} \times \mathcal{M}$ defined by $m_0 = m_I$ and $m_{j+1} = \text{Upd}(m_j, (v_j, v_{j+1}))$, i.e., $m_j = \text{Upd}^+(m_I, v_0 \cdots v_j)$. The extended play of a finite play prefix in \mathcal{A} is defined similarly.

3. PARITY GAMES WITH COSTS

In this section, we introduce the parity condition with costs [FZ14b]. Fix an arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$. A cost function for \mathcal{A} is an edge-labeling $\text{Cst}: E \rightarrow \{\varepsilon, \mathbf{i}\}$.¹ Edges labeled with \mathbf{i} are called *increment-edges* while edges labeled with ε are called ε -edges. We extend the edge-labeling to a cost function over plays obtained by counting the number of increment-edges traversed during the play, i.e., $\text{Cst}(\rho) \in \mathbb{N} \cup \{\infty\}$ for any play ρ . The cost of a finite play infix is defined analogously. Also, fix a coloring $\Omega: V \rightarrow \mathbb{N}$ of \mathcal{A} 's vertices. 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 $\text{Ans}(c) = \{c' \in \mathbb{N} \mid c' \geq c \text{ and } c' \text{ is even}\}$ be the set of colors that answer a request of color c .

Let $\rho = v_0v_1v_2 \cdots$ be a play. We define the cost-of-response at position $j \in \mathbb{N}$ of ρ by

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

where we use $\min \emptyset = \infty$, i.e., $\text{Cor}(\rho, j)$ is the cost of the infix of ρ from position j to its first answer, and ∞ if there is no answer.

We say that a request at position j is answered with cost b , if $\text{Cor}(\rho, j) = b$. Consequently, a request with an even color is answered with cost zero. The cost-of-response of an unanswered request is infinite, even if it only incurs finite cost during the remainder of the play, i.e., if there are only finitely many increment-edges succeeding the request.

¹Note that using the abstract costs ε and \mathbf{i} essentially entails a unary encoding of costs. We discuss the case of a binary encoding of arbitrary costs in Section 7.




Figure 1: Two parity games with costs. Player 1 only wins the left game.

The parity condition with costs is defined as

$$\text{CostParity}(\Omega, \text{Cst}) = \{\rho \in V^\omega \mid \limsup_{j \rightarrow \infty} \text{Cor}(\rho, j) < \infty\},$$

i.e., ρ satisfies the condition, if there exists a bound $b \in \mathbb{N}$ such that all but finitely many 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 .

A game $\mathcal{G} = (\mathcal{A}, \text{CostParity}(\Omega, \text{Cst}))$ is called a parity game with costs and its size is defined to be $|\mathcal{A}|$. If Cst assigns ε to every edge, then $\text{CostParity}(\Omega, \text{Cst})$ is a classical (max-) parity condition, denoted by $\text{Parity}(\Omega)$. Dually, if Cst assigns \mathbf{i} to every edge, then $\text{CostParity}(\Omega, \text{Cst})$ is equal to the finitary parity condition over Ω , as introduced by Chatterjee et al. [CHH09] and denoted by $\text{FinParity}(\Omega)$. In these cases, we refer to \mathcal{G} as a parity or a finitary parity game, respectively.

As most of our examples are finitary parity games, we omit the edge-labeling when drawing them for the sake of readability. For the same reason, we sometimes use non-negative integer costs on the edges of finitary parity games. Such games can be transformed into finitary parity games as defined above by subdividing these edges and coloring the newly added vertices with color 0. In all cases in this work, this only incurs a polynomial blowup.

Player 1 has two ways of winning a parity game with costs: Either he violates the classical parity condition, or he delays answers to requests arbitrarily. Consider the two parity games with costs shown in Figure 1. In the game on the left-hand side, Player 1 has a winning strategy, by taking the self-loop of the middle vertex j times upon the j -th visit to it via the increment-edge from the leftmost vertex. Thus, he delays answers to the request for 1 arbitrarily and wins by the second condition. In the game on the right-hand side, however, Player 1 does not have a winning strategy. If he eventually remains in the vertex labeled with 0, then there are only finitely many requests, only one of which is unanswered. Thus, the cost of the play is 0, i.e., it is won by Player 0. If he, on the other hand, always leaves the middle vertex eventually, then each request is answered with cost 2, hence Player 0 wins as well.

Theorem 3.1.

- (1) *Parity games can be solved in quasi-polynomial time and the problem is in $\text{UP} \cap \text{co-UP}$. The winner has a positional winning strategy [CJK⁺17, Jur98, EJ91, Mos91].*
- (2) *Solving finitary parity games is in PTIME. If Player 0 wins, then she has a positional winning strategy, but Player 1 has in general no finite-state winning strategy [CHH09].*
- (3) *Parity games with costs can be solved in quasi-polynomial time and the problem is in $\text{UP} \cap \text{co-UP}$. If Player 0 wins, then she has a positional winning strategy, but Player 1 has in general no finite-state winning strategy [MMS15, CJK⁺17, FZ14b].*

A winning strategy for Player 0 in a parity game with costs does not have to realize a uniform bound b on the value $\limsup_{j \rightarrow \infty} \text{Cor}(\rho, j)$ among all plays ρ that are consistent with σ , but the bound may depend on the play. To capture the cost of a strategy, we first define the cost of a play ρ as $\text{Cst}(\rho) = \limsup_{j \rightarrow \infty} \text{Cor}(\rho, j)$ and the cost of a strategy σ as $\text{Cst}(\sigma) = \sup_{\rho} \text{Cst}(\rho)$, where the supremum ranges over all plays ρ that are consistent with σ . A strategy is optimal for \mathcal{G} if it has minimal cost among all strategies for \mathcal{G} . Analogously, for a strategy τ for Player 1, we define $\text{Cst}(\tau) = \inf_{\rho} \text{Cst}(\rho)$, where ρ again ranges over all plays consistent with τ .

A corollary of Theorem 3.1(3) yields an upper bound on the cost of an optimal strategy: A straightforward pumping argument shows that a positional winning strategy, which always exists if there exists any winning strategy, realizes a uniform bound $b \leq n$ for every play, where n is the number of vertices of the game.

Corollary 3.2. *Let \mathcal{G} be a parity game with costs with n vertices. If Player 0 wins \mathcal{G} , then she has a strategy σ with $\text{Cst}(\sigma) \leq n$, i.e., an optimal strategy has cost at most n .*

This bound is tight, as it is trivial to construct finitary parity games \mathcal{G}_n with $n+1$ vertices and a unique play such that Player 0 wins \mathcal{G}_n with respect to bound $b = n$, but not with respect to any bound $b' < n$.

4. THE COMPLEXITY OF SOLVING PARITY GAMES WITH COSTS OPTIMALLY

In this section we study the complexity of determining the cost of an optimal strategy for a parity game with costs. Recall that solving such games is in $\text{UP} \cap \text{CO-UP}$ (and therefore unlikely to be NP-complete or CO-NP-complete) while solving the special case of finitary parity games is in PTIME. Our main result of this section shows that checking whether a strategy of cost at most b exists is PSPACE-complete, where hardness already holds for finitary parity games. Therefore, this decision problem is harder than just solving the game (unless $\text{PSPACE} \subseteq \text{UP} \cap \text{CO-UP}$, respectively $\text{PSPACE} \subseteq \text{PTIME}$).

Theorem 4.1. *The following problem is PSPACE-complete: “Given a parity game with costs \mathcal{G} and a bound $b \in \mathbb{N}$, does Player 0 have a strategy σ for \mathcal{G} with $\text{Cst}(\sigma) \leq b$?”*

Note that we do not specify how b is encoded: We will argue at the beginning of Section 4.1 that the problem is trivial for bounds $b > n$, i.e., the complexity of the problem is independent of the encoding of b .

The proof of the theorem is split into two lemmas, Lemma 4.6 showing membership and Lemma 4.7 showing hardness, which are presented in Section 4.1 and Section 4.2, respectively.

4.1. Solving Parity Games with Costs Optimally is in Polynomial Space. The remainder of this section is dedicated to showing that parity games with costs with respect to a given bound can be solved in polynomial space. To this end, we fix a parity game with costs $\mathcal{G} = (\mathcal{A}, \text{CostParity}(\Omega, \text{Cst}))$ with $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and a bound b . Let $n = |V|$. First, let us remark that we can assume w.l.o.g. $b < n$: If $b \geq n$, then, due to Corollary 3.2, we just have to check whether Player 0 wins \mathcal{G} . This is possible in polynomial space due to Theorem 3.1(3).

To obtain a polynomial space algorithm, we first turn the quantitative game \mathcal{G} into a qualitative parity game \mathcal{G}' in which the cost of open requests is explicitly tracked up

to the bound b . To this end, we use functions r mapping odd colors to $\{\perp\} \cup [b+1] = \{\perp\} \cup \{0, \dots, b\}$, where \perp denotes that no open request of this color is pending. Additionally, whenever the bound b is exceeded for some request, all open requests are reset and a so-called overflow counter is increased, up to value n . This accounts for a bounded number of unanswered requests, which are allowed by the parity condition with costs. Intuitively, Player 1 wins \mathcal{G}' if he either exceeds the upper bound b at least n times, or if he enforces an infinite play of finite cost with infinitely many unanswered requests. If he wins by the former condition, then he can also enforce infinitely many excesses of b via a pumping argument. The latter condition accounts for plays in which Player 1 wins without violating the bound b repeatedly, but by violating the classical parity condition. We show that Player 0 has a strategy σ in \mathcal{G} with $\text{Cst}(\sigma) \leq b$ if and only if she wins \mathcal{G}' from its initial vertex.

The resulting game \mathcal{G}' is of exponential size in the number of odd colors d and can therefore in general not be solved in polynomial space in n . Thus, in a second step, we construct a finite-duration variant \mathcal{G}'_f of \mathcal{G}' , which is played on the same arena as \mathcal{G}' , but the winner of a play is determined after a polynomial number of moves. We show that Player 0 wins \mathcal{G}' if and only if she wins \mathcal{G}'_f . To conclude, we show how to simulate \mathcal{G}'_f on the fly on an alternating Turing machine in polynomial time in n , which yields a polynomial space algorithm by removing the alternation [CKS81].

We begin by defining \mathcal{G}' . Let $R = (\{\perp\} \cup [b+1])^D$ be the set of request functions, where D is the set of odd colors occurring in \mathcal{G} . Here, $r(c) = \perp$ denotes that there is no open request for the color c , while $r(c) \neq \perp$ encodes that the oldest open request of c has incurred cost $r(c)$. Using these functions, we define the memory structure $\mathcal{M} = ([n+1] \times R, m_I, \text{Upd})$, where the first component $[n+1]$ implements the previously mentioned overflow counter. It suffices to bound this counter by n , since, as we will show, if Player 1 can enforce n overflows in \mathcal{G} , then he can also enforce infinitely many by a pumping argument. If this counter reaches n , we say that it is saturated.

The initial memory state m_I is the pair $(0, r_{v_I})$, where, for an arbitrary $v \in V$, r_v is the function mapping all odd colors to \perp , if $\Omega(v)$ is even. If $\Omega(v)$ is odd, however, r_v maps $\Omega(v)$ to 0, and all other odd colors to \perp . The update function $\text{Upd}(m, e)$ is defined such that traversing an edge $e = (v, v')$ updates the memory state (o, r) to the memory state (o', r') by performing the following steps in order:

- If e is an increment-edge, then for each c with $r(c) \neq \perp$, set $r'(c) = r(c) + 1$. For all other c , set $r'(c) = r(c) = \perp$. If e is an ε -edge, however, then set $r'(c) = r(c)$ for all c .
- Now, if there exists a color c such that $r'(c) > b$, then set $r'(c) = \perp$ for all c and set o' to the minimum of $o + 1$ and n . Otherwise, set o' to o .
- If $\Omega(v')$ is even, set $r(c')$ to \perp for every $c' \leq \Omega(v')$.
- If $\Omega(v')$ is odd, then set $r'(\Omega(v'))$ to the maximum of the previous value of $r'(\Omega(v'))$ and 0, where $\max\{\perp, 0\} = 0$.

The resulting o' is at most n and the resulting function r' is an element of R . We show an example of the evolution of the memory states on a play prefix in Figure 2. In particular, the move from the fifth vertex to the sixth one causes an overflow that resets all requests, but also increments the overflow counter.

We define the parity game $\mathcal{G}' = (\mathcal{A} \times \mathcal{M}, \text{Parity}(\Omega'))$, with $\Omega'(v, o, r) = \Omega(v)$ for $o < n$ and $\Omega'(v, n, r) = 1$. Note that every play that encounters a vertex of the form (v, n, r) at some point, for some $v \in V$ and some $r \in R$, is winning for Player 1, since these vertices form a winning sink component for him. Intuitively, reaching a vertex of this form means




Figure 2: Example of the evolution of the request-functions during a play for $b = 2$.

that Player 1 is able to often open n requests that are not answered with cost at most b . We show that this implies him being able to open infinitely many such requests. However, there is another way of winning the original parity game with costs \mathcal{G} for him, i.e., by violating the underlying parity condition of \mathcal{G} . This is possible even if no request incurs a cost greater than b . Hence, \mathcal{G}' is a parity game.

Let $\mathcal{A}' = \mathcal{A} \times \mathcal{M} = (V', V'_0, V'_1, E', v'_I)$, in particular $V' = V \times M$ and $V'_i = V_i \times M$. Even though \mathcal{G}' has no cost function, we say that an edge $((v, m), (v', m'))$ of \mathcal{G}' is an increment-edge if (v, v') is an increment-edge in \mathcal{G} , otherwise we call it an ε -edge.

It suffices to solve \mathcal{G}' to determine whether Player 0 can bound the cost in \mathcal{G} by b .

Lemma 4.2. *Player 0 has a strategy σ in \mathcal{G} with $\text{Cst}(\sigma) \leq b$ if and only if Player 0 wins \mathcal{G}' .*

Proof. We first introduce some notation. Let $(v_0, o_0, r_0)(v_1, o_1, r_1)(v_2, o_2, r_2) \cdots$ be a play or a play prefix in \mathcal{A}' . An *overflow position* is a j such that either $j = 0$ or $o_j = o_{j-1} + 1$. Note that we have $r_j = r_{v_j}$ for every overflow position, i.e., the request function is reset at each such position.

For the direction from right to left, assume that Player 0 wins \mathcal{G}' and let $\sigma': V'_0 \rightarrow V'$ be a positional winning strategy for her from v'_I in \mathcal{G}' . We define the finite-state strategy σ for Player 0 in \mathcal{G} using the memory structure \mathcal{M} and the next-move function Nxt with $\text{Nxt}(v, m) = v'$, if $\sigma'(v, m) = (v', m')$ for some m' . Let $\rho = v_0 v_1 v_2 \cdots$ be a play that is consistent with σ . A straightforward induction shows that the unique extended play $\text{ext}(\rho) = (v_0, o_0, r_0)(v_1, o_1, r_1)(v_2, o_2, r_2) \cdots$ in \mathcal{G}' is consistent with σ' and therefore winning for Player 0. This in particular implies $o_j < n$ for all $j \in \mathbb{N}$, as the vertices (v, n, r) form a rejecting sink component.

First, assume that ρ traverses infinitely many increment-edges. The overflow-counter in $\text{ext}(\rho)$ stabilizes at some value less than n at some point, i.e., there is a position j such that $o_{j'} = o_j < n$ for every $j' > j$. We claim $\text{Cor}(\rho, j') \leq b$ for every $j' > j$, which finishes this direction of the proof. Assume towards a contradiction that a request at some position after j of ρ is unanswered for $b + 1$ increment-edges. During every traversal of one of these increment-edges, its associated counter in $\text{ext}(\rho)$ is increased by one and not reset until $b + 1$ increment-edges are traversed, which implies encountering an overflow position. This contradicts the choice of the position j . Thus, if ρ traverses infinitely many increment-edges, then almost every request is answered with cost at most b , i.e., $\text{Cst}(\rho) \leq b$.

Now, consider the case where ρ contains only finitely many increment-edges. Such a play satisfies the parity condition with costs if and only if it satisfies the parity condition. Thus, it suffices to note that ρ and $\text{ext}(\rho)$ coincide on their color sequences, due to $o_j < n$ for all $j \in \mathbb{N}$, and that $\text{ext}(\rho)$ satisfies the parity condition, as it is winning for Player 0.

For the other direction, we prove the contrapositive. Assume that Player 0 does not win \mathcal{G}' . Then, due to determinacy of parity games, Player 1 wins \mathcal{G}' , say using the positional

strategy $\tau': V'_1 \rightarrow V'$. This strategy is only useful as long as the overflow counter is not yet saturated, as a play is trivially winning for Player 1 as soon as the sink component is reached. Thus, whenever the overflow counter is increased, we reset it to the smallest possible value for which τ' is still able to enforce a winning play for Player 1.

We introduce the set \mathcal{R} that contains all vertices (v, o, r) that are visited by some play that is consistent with τ' . Then, given a vertex v , let $o_v = \min\{o \mid (v, o, r_v) \in \mathcal{R}\}$ with $\min \emptyset = n$. Note that we have $o_{v_I} = 0$ and that the strategy τ' is winning from (v, o_v, r_v) in \mathcal{G}' for all $v \in V$. Now we define a new memory structure $\mathcal{M}' = (M', m'_I, \text{Upd}')$ with $M' = M = [n+1] \times R$, $m'_I = m_I = (0, r_{v_I})$, and

$$\text{Upd}'((o, r), (v, v')) = \begin{cases} (o, r') & \text{if } \text{Upd}((o, r), (v, v')) = (o, r') \\ (o_v, r') & \text{if } \text{Upd}((o, r), (v, v')) = (o+1, r') \end{cases}.$$

Note that we have $r' = r_{v'}$ in the second case. Now, let τ be the finite-state strategy implemented by \mathcal{M}' and the next-move function Nxt given by $\text{Nxt}(v, m) = v'$, if $\tau(v, m) = (v', m')$ for some $m' \in M$.

Let $\rho = v_0 v_1 v_2 \cdots$ be some play in \mathcal{G} that is consistent with τ and, moreover, let $\rho' = (v_0, o_0, r_0)(v_1, o_1, r_1)(v_2, o_2, r_2) \cdots$ be the extended play of ρ with respect to \mathcal{M}' . We say that j is a *reset position* if $j = 0$ or if $\text{Upd}((o_{j-1}, r_{j-1}), (v_{j-1}, v_j)) = (o_{j-1} + 1, r_j)$, i.e., the second case in the definition of Upd' is applied. Note that ρ' is not necessarily a play in \mathcal{G}' , but every maximal infix of ρ' between two reset positions is a play infix in that game that is consistent with τ' . Furthermore, at every reset position, instead of incrementing the overflow counter, we set it to the minimal value o_v . As a reset position in ρ' is only reached when incurring an overflow, for every reset position but the first one there exists at least one request in ρ that is open for at least $b+1$ increment-edges.

We now prove that the play ρ in \mathcal{G} is winning for Player 1. Recall that the play $\rho' = (v_0, o_0, r_0)(v_1, o_1, r_1)(v_2, o_2, r_2) \cdots$ is the extension of ρ with respect to \mathcal{M}' . In order to show ρ to be winning, we proceed in two steps. First, we show that we have $o_j < n$ for all j , i.e., the strategy τ always uses meaningful moves of τ' for its choice of move. This allows us to argue that τ is indeed winning for Player 1.

First, note that even though ρ' is not a play in \mathcal{G}' , every vertex (v_j, o_j, r_j) of ρ' is in \mathcal{R} . We show this claim by induction over j . For $j = 0$, we obtain $(v_0, o_0, r_0) = v'_I \in \mathcal{R}$. For $j > 0$, we obtain $v'_{j-1} = (v_{j-1}, o_{j-1}, r_{j-1}) \in \mathcal{R}$ by induction hypothesis. Hence, let π be a play prefix ending in v'_{j-1} that is consistent with τ' . If $v'_{j-1} \in V'_0$, then $\pi \cdot (v_j, o_j, r_j)$ is consistent with τ' . Otherwise, i.e., if $v'_{j-1} \in V'_1$, then we obtain $\tau(\pi) = (v_j, o, r_j)$ for some $o \in [n+1]$. In case $o = o_j$, $(v_j, o_j, r_j) \in \mathcal{R}$ follows directly. If $o \neq o_j$, however, then $o_j = o_{v_j}$ and $r_j = r_{v_j}$, i.e., $(v_j, o_j, r_j) \in \mathcal{R}$ by definition of o_{v_j} .

Next, we show $o_{j+1} \leq o_j + 1$ for all $j \in \mathbb{N}$. Assume we have $o_{j+1} > o_j$ for some j . As argued above, the vertex (v_j, o_j, r_j) is in \mathcal{R} . As $o_j \neq o_{j+1}$, there is an edge from (v_j, o_j, r_j) to $(v_{j+1}, o_j + 1, r_{j+1})$ in \mathcal{A}' , where $r_{j+1} = r_{v_{j+1}}$. By construction of τ , we obtain $(v_{j+1}, o_j + 1, r_{j+1}) \in \mathcal{R}$ as argued above. Hence, we indeed have $o_{j+1} = o_{v_{j+1}} \leq o_j + 1$.

Now, we argue $o_j < n$ for all $j \in \mathbb{N}$ by proving the following property by induction over j :

If $o_j = k$, then for every $k' \leq k$ there is a reset position $j_{k'} \leq j$ with $o_{j_{k'}} = k'$.

Let us first argue that this indeed implies $o_j < n$. Towards a contradiction, assume $o_j = n$. Then, there are $n+1$ reset positions, one for each value k in the range $0 \leq k \leq n$ for the overflow counter. Thus, two such positions j', j'' share the same vertex $v_{j'} = v_{j''}$, which

implies that they also share the same overflow counter value $o_{j'} = o_{v_{j'}} = o_{v_{j''}} = o_{j''}$. This yields the desired contradiction to $o_{j'}$ and $o_{j''}$ being distinct.

For the induction start, we have $o_0 = 0$ and pick $j_0 = 0$, which is a reset position. Now, let $j > 0$ and $o_j = k$. If $k \leq o_{j-1}$, then the induction hypothesis yields the necessary positions. Hence, assume we have $k > o_{j-1}$, which implies $k = o_{j-1} + 1$ as shown above. Then, j is a reset position and we can define $j_k = j$ and obtain the remaining $j_{k'}$ for $k' < k$ from the induction hypothesis.

It remains to show that ρ is indeed winning for Player 1. First assume that the overflow counter of ρ' stabilizes, i.e., there exists some $j \in \mathbb{N}$ such that $o_{j'} = o_j$ for all $j' > j$. Then, there exists a suffix of ρ' that is consistent with τ' , which therefore violates the parity condition. Hence, it suffices to note that the colors of ρ' and ρ coincide, i.e., ρ violates the parity condition and thus also the parity condition with costs with respect to any bound.

Now assume that the overflow counter of ρ' does not stabilize. Then, there are infinitely many reset positions in ρ' . Between any two adjacent such positions, by construction, there exists a request that remains unanswered for at least $b + 1$ steps in ρ . Hence, ρ violates the parity condition with costs with respect to bound b and is winning for Player 1. \square

As the parity game \mathcal{G}' is of exponential size and can therefore not be constructed and solved in polynomial space in $|\mathcal{G}|$, we now construct a finite-duration variant \mathcal{G}'_f of \mathcal{G}' . One such variant is obtained by playing the parity game up to the first vertex repetition and declaring the winner according to the maximal color on the induced cycle [AR14]. However, one can show that such a play in \mathcal{G}' is still of exponential length in the worst case. In the following, we exploit the structure of the arena to proclaim a winner after a polynomial number of moves. In particular, we define a preorder on the memory elements and stop a play as soon as a pseudo-cycle is reached, i.e., an infix whose projection to V is a cycle in \mathcal{A} and whose memory states at the start and at the end are in the order relation.

We first introduce the preorder on memory states. To this end, note that not all open requests are “relevant”. In fact, a small request that is opened while a larger one is already open is irrelevant. Answering the larger request is more urgent, as it has already incurred at least as much cost as the newly opened request and answering the larger one answers the smaller one as well. Following this intuition, we define the relevant requests of a request function r as follows: c is relevant in r if and only if $r(c) \neq \perp$ and if there does not exist a color $c' > c$ such that $r(c') \geq r(c)$. Then, for each open request c in r , there exists some $c' \geq c$ such that the request for c' is relevant and $r(c') \geq r(c)$. Also, the largest open request and the one with the highest cost are always relevant. We denote the set of relevant requests of r by $\text{RELREQ}(r)$.

Using the relevant requests, we define a preorder \sqsubseteq on request functions: For two request functions r and r' , we say that r' dominates r , if for each color c that is relevant in r , there exists a relevant request c' in r' that is “more urgent” than that for color c . Formally, we write $r \sqsubseteq r'$ if and only if for each c that is relevant in r , there exists a c' with $c \leq c'$ that is relevant in r' with $r(c) \leq r'(c')$. The relation \sqsubseteq is reflexive and transitive. Moreover, if $r \sqsubseteq r'$ and $r \supseteq r'$ both hold true, we write $r \approx r'$. Finally, we have that $r \approx r'$ implies $\text{RELREQ}(r) = \text{RELREQ}(r')$ and $r(c) = r'(c)$ for all $c \in \text{RELREQ}(r)$.

We extend the preorder \sqsubseteq to memory elements. Since the overflow counter is non-decreasing and every one of its increments brings a play closer towards the winning sink vertices for Player 1, we value this component of the memory state more strongly. Following this intuition, we say that the memory state (o, r) is dominated by the memory state (o', r') ,

written $(o, r) \sqsubseteq (o', r')$, if either $o < o'$, or if both $o = o'$ and $r \sqsubseteq r'$ hold true. Similarly, we extend the notation of \approx such that $(o, r) \approx (o', r')$, if and only if both $(o, r) \sqsubseteq (o', r')$ and $(o, r) \supseteq (o', r')$ hold true.

The preorder on memory elements is preserved under concatenation.

Lemma 4.3. *Let $(v, o_1, r_1), (v, o_2, r_2)$ be vertices in \mathcal{G}' , let $(v, v') \in E$, and, for $j \in \{1, 2\}$, let $\text{Upd}((o_j, r_j), (v, v')) = (o'_j, r'_j)$. If $(o_1, r_1) \sqsubseteq (o_2, r_2)$ and $o_2 < n$, then $(o'_1, r'_1) \sqsubseteq (o'_2, r'_2)$.*

Proof. First assume $o_1 < o_2$. Due to the construction of \mathcal{G}' , this implies $o'_1 \leq o'_2$. If $o'_1 < o'_2$, the given statement holds true. If $o'_1 = o'_2$, however, then $r'_1 = r_{v'}$ and thus $(o'_1, r'_1) \sqsubseteq (o'_2, r'_2)$, since for every (v, o, r) with incoming edges in \mathcal{A}' , we have $r \supseteq r_{v'}$. Thus, assume $o_1 = o_2$ and $r_1 \sqsubseteq r_2$ for the remainder of this proof.

First, assume that $o'_1 = o_1 + 1$. We show the following claim: If the move to v' causes an overflow when starting from (v, o_1, r_1) , then the same move causes an overflow when starting from (v, o_2, r_2) . This then implies $r'_1 = r'_2 = r_{v'}$ and hence, $(o'_1, r'_1) \sqsubseteq (o'_2, r'_2)$. Formally, we claim that if $o'_1 = o_1 + 1$, then $o'_2 = o_2 + 1$. It remains to show $o'_2 = o_2 + 1$. Since the move from v to v' causes an overflow when starting in (v, o_1, r_1) , we have $\text{Cst}(v, v') = \mathbf{i}$. Let c_1 be some color that causes the overflow in r_1 , i.e., $c_1 \in \{c \mid r_1(c) = b\}$. Since $r_2 \supseteq r_1$, there exists a color $c_2 \geq c_1$ with $r_2(c_2) \geq r_1(c_1)$. Moreover, since the range of r_2 is upwards bounded by b , this implies $r_2(c_2) = b$. Hence, the move from v to v' also causes an overflow in r_2 , which implies $o'_2 = o_2 + 1$. This completes the proof in the case $o'_1 = o_1 + 1$.

If, however, $o'_1 = o_1$, we again distinguish two cases: If $o'_2 = o_2 + 1$, then $o'_1 < o'_2$ and hence, $(o'_1, r'_1) \sqsubseteq (o'_2, r'_2)$. On the other hand, if $o'_2 = o_2$, then let c_1 be relevant in r'_1 . We show that there exists a color $c_2 \geq c_1$ with $r'_2(c_2) \geq r'_1(c_1)$. Should c_2 be non-relevant, then there exists a larger one in r'_2 that dominates c_2 .

If a request for c_1 was already open in r_1 , then let $c_2 \geq c_1$ with $r_2(c_2) \geq r_1(c_1)$. Such a color c_2 exists due to $r_2 \supseteq r_1$. Since the request for c_1 was not answered during the move to v' , and since $o_2 = o'_2 < n$, neither was the request for c_2 during the same move. Hence, we have $r'_2(c_2) \geq r'_1(c_1)$. If, however, a request for c_1 was not already open in r_1 , then the request for c_1 must have been opened by moving to v' , i.e., $\Omega(v') = c_1$. Thus, we directly obtain $r'_1(c_1) = 0$ and $r'_2(c_1) \geq 0$. Picking $c_2 = c_1$ concludes the proof in this case. \square

We now define the winning condition in the finite game \mathcal{G}'_f . To this end, let $\pi = (v_0, o_0, r_0) \cdots (v_j, o_j, r_j)$ be a play prefix in \mathcal{G}' and let $\pi' = (v_k, o_k, r_k) \cdots (v_{k'}, o_{k'}, r_{k'})$ be an infix of π . We say π' is a dominating cycle² if $v_k = v_{k'}$, $o_k = o_{k'} < n$, and either

- the maximal color occurring on π' is even and $r_k \supseteq r_{k'}$, or
- the maximal color occurring on π' is odd and $r_k \sqsubseteq r_{k'}$.

We call the former and latter type of dominating cycles even and odd, respectively. Moreover, we say that a play prefix $\pi = (v_0, o_0, r_0) \cdots (v_j, o_j, r_j)$ is *settled* if either $o_j = n$, or if π contains a dominating cycle.³ Fix $\ell = (n + 1)^6$.

Lemma 4.4. *Let π be a play prefix of \mathcal{G}' . If $|\pi| > \ell$, then π is settled.*

²Note that a dominating cycle is only a cycle when projected to a play in \mathcal{A} .

³This definition differs from the one presented in the conference version [WZ16]. The definition here is easily amenable to the case of integer-valued cost functions in Section 7, simplifies the proofs of Lemma 4.4 and Lemma 4.5, and fixes a bug in the conference version, which caused plays to be settled too early in favor of Player 1.

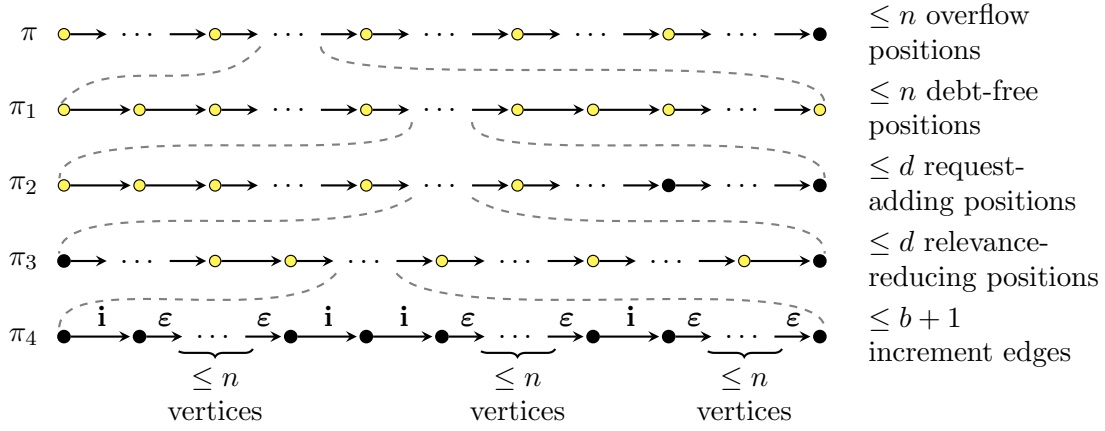


Figure 3: Bounding the length of unsettled play prefixes. The relevant special vertices are marked in yellow.

Proof. Let $\pi = (v_0, o_0, r_0) \cdots (v_j, o_j, r_j)$ be an unsettled play prefix of \mathcal{G}' . We show $|\pi| \leq \ell$, which implies the given statement. Note that, since π is not settled, it does not contain a vertex repetition, since such a repetition induces a dominating cycle.

The structure of our argument is sketched in Figure 3: We recall the definition of overflow positions, define debt-free, request-adding, and relevance-reducing positions, and show

- (1) that there are at most n overflow positions in π ,
- (2) that there are at most n debt-free positions between any two adjacent overflow positions,
- (3) that there are at most d request-adding positions between any two adjacent debt-free positions, where d is the number of odd colors,
- (4) that there are at most d relevance-reducing positions between any two adjacent request-adding positions,
- (5) that there are at most $b + 1$ increment-edges between any two adjacent relevance-reducing positions, and
- (6) that there are at most n vertices between two such increment-edges.

Aggregating these bounds then yields the desired result.

Recall that an overflow position of π is a k with $k = 0$ or with $o_k = o_{k-1} + 1$. As π is unsettled and the o_k are non-decreasing, π has at most n overflow positions, $n - 1$ real increments and the initial position. Hence, by splitting π at the overflow positions we obtain at most n non-empty infixes of π , each without overflow positions. We say such an infix has type 1.

Fix a non-empty type 1 infix π_1 . A *debt-free position* of π is a k with $r_k = r_{v_k}$, i.e., a position that has no other costs than those incurred by visiting v_k . As all vertices of π_1 share the same overflow counter value, there are at most n debt-free positions in π_1 : $n + 1$ such positions would induce a vertex repetition, which we have ruled out above. Hence, by splitting π_1 at the debt-free positions we obtain at most $n + 1$ non-empty infixes of π_1 , each without debt-free and overflow positions. We say such an infix has type 2.

Fix a non-empty type 2 infix π_2 . A *request-adding position* of π is a k with odd $\Omega(v_k)$ such that $r_{k-1}(c) = \perp$ for all $c \geq \Omega(v_k)$. We define d as the number of odd colors assigned

by Ω and claim that there are at most d request-adding positions in π_2 . Assume there are $d + 1$. Then, two request-adding positions $k < k'$ share a color, call it c . As k' is request-adding, only requests strictly smaller than c are open at position $k' - 1$, i.e., c and all larger requests have to be answered in between k and k' . Hence, there is a debt-free position between k and k' , which contradicts π_2 being of type 2. Hence, by splitting π_2 at the request-adding positions we obtain at most $d + 1$ non-empty infixes of π_2 , each without request-adding, debt-free, and overflow positions. We say such an infix has type 3.

Fix a non-empty type 3 infix π_3 . A *relevance-reducing position* of π is a k such that $\text{RELREQ}(r_{k-1}) \supsetneq \text{RELREQ}(r_k)$. We show that π_3 contains at most d relevance-reducing positions. To this end, we first argue that there is at least one request that is open throughout π_3 . First note that some request must be open at the beginning of π_3 , as otherwise the first vertex would be at a debt-free position, which do not occur in π_3 . Let c^* be the maximal open request at the beginning of π_3 . Due to π_3 not containing debt-free nor request-adding positions, all colors c visited during π_4 must satisfy $c \leq c^*$. Hence, the request for c^* remains open throughout π_3 . We now show that the sets of relevant requests along π_3 form a descending chain in the subset-relation. Assume towards a contradiction that an infix $(v, o, r)(v', o', r')$ of π_3 and a color c exist such that $c \notin \text{RELREQ}(r)$, but $c \in \text{RELREQ}(r')$, i.e., $\Omega(v') = c$. Then, $c > c^*$, i.e., there exists a request-adding position in π_3 , a contradiction. Hence, the sets of relevant requests indeed form a descending chain. As at the beginning of π_3 at most d requests are relevant, there are at most d relevance-reducing positions. Thus, by removing relevance-reducing positions, we obtain at most $d + 1$ non-empty infixes, each without relevance-reducing, request-adding, debt-free, and overflow positions. We say such an infix has type 4.

Fix a non-empty type 4 infix π_4 . We show that π_4 contains at most b increment-edges. As argued above, there exists some color c for which a request is open at the beginning of π_4 , which is not answered throughout this infix. Thus, $b + 1$ increment-edges in π_4 would lead to an overflow position. However, π_4 has no overflow positions by construction. Thus, there are at most b increment-edges in π_4 . Hence, by splitting π_4 at the increment-edges, we obtain a decomposition of π_4 into at most $b + 1$ infixes, each without increment-edges and without relevance-reducing, request-adding, debt-free, and overflow positions. We say such an infix has type 5.

Fix a non-empty type 5 infix π_5 . We show that π_5 is of length at most n . Assume towards a contradiction that π_5 contains at least $n + 1$ vertices. Then there exists an infix $\pi' = (v, o, r) \cdots (v, o, r')$ of π_5 , since π_5 does not contain overflow positions. As argued above, we have $\text{RELREQ}(r) = \text{RELREQ}(r')$. Moreover, as π_5 contains no increment-edges, we furthermore obtain $r \approx r'$. Thus, π' is a dominating cycle, which contradicts π being unsettled. Hence, π_5 is of length at most n .

Aggregating all these bounds yields an upper bound of $(n + 1)^6$ on the length of the unsettled play prefix π , as we have $d \leq n$ and $b \leq n$. \square

Using the notion of settled plays, we now define the finite game $\mathcal{G}'_f = (\mathcal{A}', \text{Win}_f)$, in which both players try to settle the play in their favor. Due to Lemma 4.4, every play in \mathcal{A}' is settled. Thus, let ρ be an infinite play in \mathcal{A}' and let π be the minimal settled prefix of ρ , which can be settled due to three mutually exclusive criteria. If π is settled due to containing an even dominating cycle, then ρ is winning for Player 0. Otherwise, i.e., if π is settled due to saturating the overflow counter or due to containing an odd dominating cycle, ρ is winning for Player 1. The game \mathcal{G}'_f is indeed a game of finite duration, as the winner

is certain after ℓ moves, due to Lemma 4.4. Hence, \mathcal{G}'_f is determined [Zer13]. Moreover, in order to solve \mathcal{G}' , it suffices to solve \mathcal{G}'_f .

Lemma 4.5. *Player 0 wins \mathcal{G}' if and only if she wins \mathcal{G}'_f .*

Proof. We first show that Player 0 wins \mathcal{G}' if she wins \mathcal{G}'_f . Let σ'_f be a winning strategy for Player 0 in \mathcal{G}'_f . We construct a winning strategy σ' for Player 0 in \mathcal{G}' by simulating a play in \mathcal{G}'_f that is consistent with σ'_f . As this strategy is only useful as long as the simulating play is not settled, we have to keep the simulating play short by removing settling dominating cycles. We define the simulation $h: (V \times M)^+ \rightarrow (V \times M)^+$ and the strategy σ' simultaneously. The function h satisfies the following invariant:

Let π be consistent with σ' and end in (v, o, r) . Then, $h(\pi)$ is consistent with σ'_f , is unsettled, and ends in (v, o', r') with $(o', r') \sqsupseteq (o, r)$.

Since $h(\pi)$ is consistent with the winning strategy σ'_f for Player 0 and unsettled, this implies that neither the overflow counter of $h(\pi)$ nor that of π reaches the value n .

To begin, let $h(v'_I) = v'_I$, which satisfies the invariant. Now, assume we have a play prefix π consistent with σ' ending in (v, o, r) and let $h(\pi) = (v_0, o_0, r_0) \cdots (v_j, o_j, r_j)$. We consider two cases, depending on whose turn it is at the last vertex (v, o, r) of π .

If $(v, o, r) \in V'_1$, Player 1 moves to some successor of (v, o, r) in \mathcal{G}' , say (v^*, o^*, r^*) . Furthermore, define $(o^*_f, r^*_f) = \text{Upd}((o_j, r_j), (v_j, v^*))$, which is the corresponding memory update in \mathcal{G}'_f . If $(v, o, r) \in V'_0$, let $\sigma'_f(h(\pi)) = (v^*, o^*_f, r^*_f)$ in \mathcal{G}'_f . We mimic the move to (v^*, o^*_f, r^*_f) in \mathcal{G}' by defining $\sigma'(\pi) = (v^*, o^*, r^*)$ with $(o^*, r^*) = \text{Upd}((o, r), (v, v^*))$. This is well-defined due to the invariant $v_j = v$.

In both cases, for a given successor v^* of v in \mathcal{A} , we have $(o^*_f, r^*_f) = \text{Upd}((o_j, r_j), (v, v^*))$ and $(o^*, r^*) = \text{Upd}((o, r), (v, v^*))$. It remains to define $h(\pi \cdot (v^*, o^*, r^*))$. We distinguish two cases: If $\pi^*_f = h(\pi) \cdot (v^*, o^*_f, r^*_f)$ is not settled, we pick $h(\pi \cdot (v^*, o^*, r^*)) = \pi^*_f$. This satisfies the invariant due to Lemma 4.3.

Now assume π^*_f is settled. Since π^*_f is consistent with the winning strategy σ'_f for Player 0, π^*_f is settled due to containing an even dominating cycle. Moreover, since $h(\pi)$ is not settled, the dominating cycle is a suffix of π^*_f . Thus, the cycle starts in a vertex $(v_{j'}, o_{j'}, r_{j'})$ with $v_{j'} = v^*$ and $r_{j'} \sqsupseteq r^*_f$. Removing the settling cycle, we define $h(\pi \cdot (v^*, o^*, r^*)) = (v_0, o_0, r_0) \cdots (v_{j'}, o_{j'}, r_{j'})$, which satisfies the invariant due to transitivity of \sqsubseteq .

Now, consider a play ρ consistent with σ' and let π_j be the prefix of length j of ρ . As argued before, neither the overflow counter of the π_j nor that of the $h(\pi_j)$ reaches n . Hence, the colors of the last vertices of π_j and $h(\pi_j)$ coincide for all j .

Towards a contradiction, assume that the maximal color occurring infinitely often along ρ is odd, call it c . After some finite prefix, c cannot occur on even dominating cycles in the $h(\pi_j)$ anymore, since each occurrence on such a cycle implies at least one occurrence of an even higher even color in ρ . Hence, after this prefix, each time a vertex of color c is visited, say at the end of the prefix π_j , a vertex of the same color is appended to the simulated play $h(\pi_j)$. Moreover, this vertex is never removed from the simulated play, since only vertices occurring on even dominating cycles are removed from the simulated play. Hence, the simulated play becomes longer with each visit to a vertex of color c after a finite prefix. This contradicts the $h(\pi_j)$ being unsettled, as every play of length $\ell + 1$ is settled due to Lemma 4.4. Thus, the maximal color occurring infinitely often in ρ must be even, i.e., σ' is winning for Player 0 in \mathcal{G}' .

For the other direction, we show that Player 1 wins \mathcal{G}' if he wins \mathcal{G}'_f . Due to determinacy of \mathcal{G}' and \mathcal{G}'_f , this suffices to show the result. Let τ'_f be a winning strategy for Player 1 in \mathcal{G}'_f . Similarly to the previous case, we simulate play prefixes π in \mathcal{G}' by play prefixes $h(\pi)$ in \mathcal{G}'_f and define a simulation function h and a winning strategy τ' for Player 1 in \mathcal{G}' simultaneously. Again, we need to make sure that the simulating play prefixes remain short and unsettled, as long as the overflow counter of the play in \mathcal{G}' is not saturated. We do so by removing settling dominating cycles from the simulating play prefixes in \mathcal{G}'_f . Formally, h satisfies the following invariant:

Let π be consistent with τ' and end in (v, o, r) with $o < n$. Then, $h(\pi)$ is consistent with τ'_f , is unsettled, and ends in (v, o', r') with $(o', r') \sqsubseteq (o, r)$.

If $o = n$, we can stop the simulation and let Player 1 pick arbitrary successor vertices in \mathcal{G}' , since the play has reached the winning sink component for Player 1 in \mathcal{G}' .

At the beginning, we pick $h(v'_I) = v'_I$, which satisfies the invariant. Now, let π be a play prefix consistent with τ' ending in (v, o, r) and let $h(\pi)$ be defined. We obtain the next vertex similarly to the previous case, i.e., as an arbitrary successor of (v, o, r) if $(v, o, r) \in V'_0$, and by applying τ'_f to $h(\pi)$ if $(v, o, r) \in V'_1$. In the latter case, we moreover define τ' as simulating the move of τ'_f as previously. In both cases, we obtain v^* as the first component of the next vertex in both \mathcal{G}' and \mathcal{G}'_f . Again, we define (o^*, r^*) and (o^*_f, r^*_f) as previously.

Let $\pi^* = \pi \cdot (v^*, o^*, r^*)$. It remains to define $h(\pi^*)$. If $o^* = n$, Player 1 has already won and we can define $h(\pi^*)$ arbitrarily since the invariant contains an empty premise. If $o^* < n$, however, we define $h(\pi^*)$ as in the previous case, i.e., as π^*_f , if π^*_f is not settled, and by removing the settling odd dominating cycle otherwise. No even dominating cycle may occur, since π^*_f is consistent with τ'_f . This maintains the invariant due to the same argument as previously.

Now consider a play ρ that is consistent with τ' . If the overflow counter along ρ reaches the value n , then ρ is winning for Player 1. Thus, we consider the case where the counter is always smaller than n . In this case, however, the maximal color seen infinitely often in ρ is odd, due to the same argument as in the previous case: If it were even, vertices of that color would be appended to the simulation infinitely often without being removed, contradicting the simulated play being unsettled. Hence, τ' is winning for Player 1 in \mathcal{G}' . \square

The combination of Lemmas 4.2 and 4.5 shows that Player 0 wins \mathcal{G} with respect to the bound b if and only if she wins \mathcal{G}'_f . Thus it remains to show that we can simulate \mathcal{G}'_f on an alternating Turing machine in polynomial time.

Lemma 4.6. *The following problem is in PSPACE: “Given a parity game with costs \mathcal{G} and a bound $b \in \mathbb{N}$, does Player 0 have a strategy σ for \mathcal{G} with $\text{Cst}(\sigma) \leq b$?”*

Proof. Given \mathcal{G} and b , we show how to simulate the finite-duration game \mathcal{G}'_f on an alternating polynomial time Turing machine using the game semantics of such machines, i.e., two players construct a single path of a run of the machine. The existential and universal player take the roles of Player 0 and Player 1, respectively. The Turing machine keeps track of the complete prefix of the simulated play of \mathcal{G}'_f . Since every vertex of $\mathcal{A} \times \mathcal{M}$ can be represented in polynomial size and since the length of the play is bounded from above by $(n + 1)^6$ due to Lemma 4.4, the Turing machine can keep track of the history explicitly and check after each step whether a dominating cycle has occurred in polynomial time. If the play is settled due to an even dominating cycle, the machine accepts, if it

is settled otherwise, the machine rejects. Note that this algorithm involves neither the explicit construction of \mathcal{G}' nor that of \mathcal{G}'_f . The Turing machine accepts \mathcal{G} and b if and only if Player 0 wins \mathcal{G}'_f . Due to Lemma 4.4, this machine terminates after polynomially many steps. Hence, $\text{APTIME} = \text{PSPACE}$ [CKS81] completes the proof. \square

4.2. Playing Parity Games with Costs Optimally is PSPACE-hard. Next, we turn our attention to proving a matching lower bound on the complexity, which already holds for finitary parity games, i.e., parity games with costs in which every edge is an increment-edge. This result is proven by a reduction from the canonical PSPACE-hard problem QBF: Given a quantified boolean formula $\varphi = Q_1x_1Q_2x_2 \dots Q_nx_n\psi$ with $Q_i \in \{\exists, \forall\}$ and where ψ is a boolean formula over the variables x_1, x_2, \dots, x_n , determine whether φ evaluates to true. We assume w.l.o.g. that ψ is in conjunctive normal form such that every conjunct has exactly three literals, i.e., $\psi = \bigwedge_{j=1}^m (\ell_{j,1} \vee \ell_{j,2} \vee \ell_{j,3})$, where every $\ell_{j,k}$ is either x or \bar{x} for some $x \in \{x_1, \dots, x_n\}$. We call each $\ell_{j,k}$ for $k \in \{1, 2, 3\}$ a literal and each conjunct of three literals a clause. Furthermore, we assume w.l.o.g. that the quantifiers Q_j are alternating with $Q_1 = Q_n = \exists$.

This proof uses the standard framework for reducing QBF to infinite two-player games of polynomial size: Player 0 implements existential choices, i.e., existential quantification and disjunctions. Dually, Player 1 implements universal quantification and disjunction. Intuitively, the players pick truth values for the variables in the order as they appear in the quantifier prefix. Then, Player 1 picks a clause and then Player 0 picks a literal from that clause. She wins if and only if the literal evaluates to true under the assignment constructed earlier. This requirement has to be encoded by the winning condition, as the polynomial state space of the constructed game is insufficient to encode all possible assignments. If this is the case, then Player 0 wins the game if and only if the formula evaluates to true.

Here, we employ a finitary parity condition with respect to a given bound to achieve this. To this end, we encode assigning a truth value to a variable by opening requests, e.g., setting x_j to false requests color $4j + 1$ and setting it to true requests color $4j + 3$. Crucially, only one color can be requested, but the latter one only after a delay of one step. Now, Player 0 picking a literal ℓ answers the corresponding request, e.g., if $\ell = \bar{x}_j$, then the color $4j + 2$ occurs, and if $\ell = x_j$, then the color $4j + 4$ occurs. Note that $4j + 4$ answers both the request corresponding to setting x_j to false and the one setting it to true. But again, the answer $4j + 4$ comes only after a delay of one step, which allows to invalidate this answer in case x_j has been assigned false, by requiring a well-chosen bound. This is possible with a finitary parity condition with respect to a given bound, but neither with a classical parity condition nor with a finitary parity condition with an existentially quantified bound, which explains the increase in complexity.

Lemma 4.7. *The following problem is PSPACE-hard: “Given a finitary parity game \mathcal{G} and a bound $b \in \mathbb{N}$, does Player 0 have a strategy σ for \mathcal{G} with $\text{Cst}(\sigma) \leq b$?”*

Proof. Let $\varphi = Q_1x_1Q_2x_2 \dots Q_nx_n\psi$ be a quantified boolean formula with $\psi = \bigwedge_{j=1}^m C_j$ and $C_j = (\ell_{j,1} \vee \ell_{j,2} \vee \ell_{j,3})$, where every $\ell_{j,k}$ is either x or \bar{x} for some $x \in \{x_1, \dots, x_n\}$. We construct a finitary parity game \mathcal{G}_φ such that Player 0 has a strategy σ for \mathcal{G}_φ with $\text{Cst}(\sigma) = 3n + 5$ if and only if the formula φ evaluates to true. The arena consists of three parts: In the first part, which begins with the initial vertex v_I , Player 0 and Player 1 determine an assignment for the variables x_1 through x_n , where Player 0 and Player 1 pick

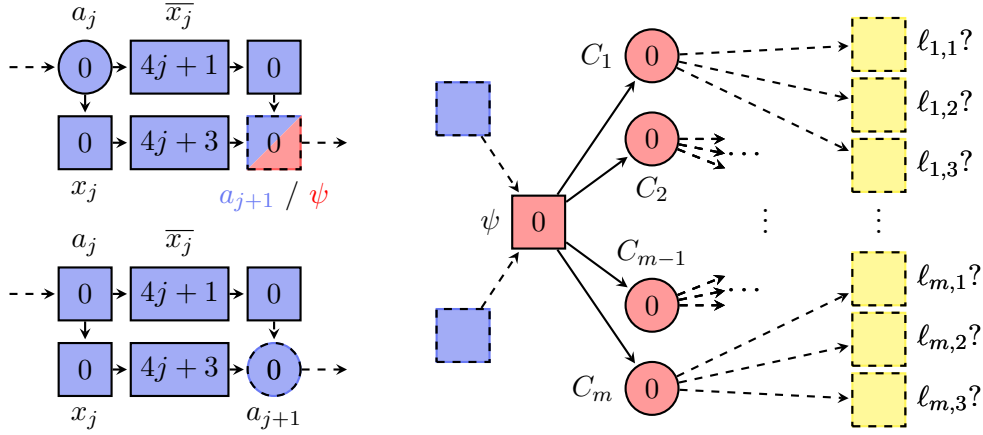


Figure 4: The gadget for existentially and universally quantified variables (left, from top to bottom), and the middle part of the arena (right).

values for the existentially and universally quantified variables, respectively. Each choice of a truth value by either player incurs a request. In the second part, Player 1 first picks a clause, after which Player 0 picks a literal from that clause. In the last part, the play then proceeds without any choice by the players and checks whether or not the chosen literal was set to true in the first part of the arena. If it was set to true, then the corresponding request is answered with cost $3n + 5$. Otherwise, that request is answered with cost $3n + 6$. Furthermore, all other potentially open requests are answered with cost at most $3n + 5$ and the play returns to the initial vertex v_I . Thus, all these gadgets are traversed infinitely often and the traversals are independent of each other. This idea indeed requires the target of the reduction to be a finitary parity game instead of a classical one, as Player 0 is able to answer all requests within at most $3n + 6$ steps independently of the truth value of φ . Only the additional requirement for her to do so within at most $3n + 5$ moves forces her to provide a witness for φ being true.

If φ evaluates to true, then Player 0 can enforce that all requests are answered with cost at most $3n + 5$. Hence, there exists a strategy σ for Player 0 with $\text{Cst}(\sigma) \leq 3n + 5$. If φ evaluates to false, however, then Player 1 can enforce requests that remain unanswered for at least $3n + 6$ steps. Thus, there exists no strategy σ for Player 0 with $\text{Cst}(\sigma) \leq 3n + 5$. We begin by constructing the arena \mathcal{A} together with its coloring Ω .

The left-hand side of Figure 4 shows the gadgets that assign a truth value to variable x_j . The vertex a_j belongs to Player 0 if x_j is existentially quantified, and to Player 1 if x_j is universally quantified. The dashed edges indicate the connections to the pre- and succeeding gadget, respectively. We construct the first part of \mathcal{A} out of n copies of this gadget. Moreover, the vertex a_1 has an incoming edge from the end of \mathcal{A} , in order to allow for infinite plays, and is the initial vertex v_I of the arena. In the remainder of this proof, let $c_{x_j} = 4j + 3$ and $c_{\bar{x}_j} = 4j + 1$ be the colors associated with assigning true or false to x_j , respectively.

The second part of the arena starts with a vertex ψ of Player 1, from which he picks a clause by moving to a vertex C_j of Player 0. Each vertex C_j is connected to three gadgets,




Figure 5: Gadgets checking the assignment of true to x_j (top) or to $\overline{x_j}$ (bottom).

one for each of the three literals contained in C_j . We show this construction in the right-hand side of Figure 4. Note that moving from the vertex of color c_{x_j} or of color $c_{\overline{x_j}}$ to the vertex ψ takes $3(n-j)+1$ or $3(n-j)+2$ steps, respectively..

The last part of the arena consists of one gadget for each literal $x_1, \overline{x_1}$ through $x_n, \overline{x_n}$ occurring in φ . These gadgets check whether or not the literal picked in the middle part was actually set to true in the first part of the arena. We show these gadgets in Figure 5.

In these gadgets, neither player has a non-trivial choice. Thus, the play proceeds by first answering requests for all colors smaller than c_{x_j} and $c_{\overline{x_j}}$. It then either grants the request for color c_{x_j} after $3j+2$ steps, or the request for color $c_{\overline{x_j}}$ after $3j+1$ steps, both counted from the vertices $x_j?$ and $\overline{x_j}?$, respectively. Since traversing the middle part of the arena incurs a constant cost of 2, a request for color c_{x_j} has incurred a cost of $3(n-j)+3$ at $x_j?$ and $\overline{x_j}?$, while a request for color $c_{\overline{x_j}}$ has incurred a cost of $3(n-j)+4$ at these vertices. Hence, the total cost incurred by the request for color c_{x_j} is $(3(n-j)+3)+(3j+2) = 3n+5$ in the gadget corresponding to x_j , and $(3(n-j)+3)+(3j+3) = 3n+6$ in the gadget corresponding to $\overline{x_j}$. The dual reasoning holds true for requests for color $c_{\overline{x_j}}$. Hence, the bound of $3n+5$ is only achieved if the request corresponding to the chosen literal was posed in the initial part of the arena.

After traversing this last gadget, all requests are answered in the vertex $x_{j,3}!$ or $\overline{x_{j,3}}!$ and the play resets to the initial vertex via an edge to a_1 .

The size of \mathcal{A} is polynomial in $|\varphi|$: The first part consists of one constant-size gadget per variable, while the second part is of linear size in the number of clauses in φ . The final part contains a gadget of size $\mathcal{O}(n)$ for each literal occurring in φ . Thus, the size of the arena is in $\mathcal{O}(n^2+m)$. It remains to argue that Player 0 has a strategy σ with $\text{Cst}(\sigma) = 3n+5$ if and only if φ evaluates to true. For any quantifier-free boolean formula ψ that contains variables x_1 through x_n and any partial assignment $\alpha: \{x_1, \dots, x_j\} \rightarrow \{\text{true}, \text{false}\}$ we denote by $\alpha(\psi)$ the formula resulting from replacing the variables in α 's domain with their respective truth values.

It suffices to argue about finite plays that begin and end in a_1 , as all plays start in a_1 , visit a_1 infinitely often, and all requests are answered before moving back to a_1 . Hence, for the remainder of this proof, we only consider a finite play infix π starting and ending in a_1 .

First assume that φ evaluates to true. We construct a strategy σ for Player 0 with the properties described above. Pick j as some index such that x_j is existentially quantified and consider the prefix π' of π up to, but not including a_j . We associate π' with an assignment $\alpha_{j-1}: \{x_1, \dots, x_{j-1}\} \rightarrow \{\text{true}, \text{false}\}$, where $\alpha_{j-1}(x_k) = \text{true}$ if π' visits x_k , and $\alpha_{j-1}(x_k) = \text{false}$ if π' visits $\overline{x_k}$. Due to the structure of the arena, exactly one of these cases holds true, hence α_{j-1} is well-defined.

For $j = 1$, $\exists x_j \cdots Q_n x_n \alpha_{j-1}(\psi)$ evaluates to true by assumption. Let $t \in \{true, false\}$ such that $\forall x_{j+1} \cdots Q_n x_n (\alpha_{j-1}[x_j \mapsto t])(\psi)$ evaluates to true as well, where $\alpha_{j-1}[x_j \mapsto t]$ denotes the mapping α_{j-1} augmented by the mapping $x_j \mapsto t$. Moreover, we define $\sigma(\cdots a_j) = x_j$ if $t = true$, and $\sigma(\cdots a_j) = \bar{x}_j$ otherwise. We proceed inductively, constructing $\sigma(\cdots a_j)$ for all existentially quantified variables x_j according to the Boolean values that satisfy the formulas $\exists x_j Q_{j+1} x_{j+1} \cdots Q_n x_n \alpha_{j-1}(\psi)$, until we reach the vertex ψ .

At this point, the analysis of the play prefix so far yields an assignment of truth values to variables $\alpha_n: \{x_1, \dots, x_n\} \rightarrow \{true, false\}$, such that $\alpha_n(\psi)$ evaluates to true. Let $\alpha = \alpha_n$.

At vertex ψ there exist n open requests. As previously argued, if $\alpha(x_j) = true$, then there is an open request for c_{x_j} with cost $3(n - j) + 1$. Otherwise, there is an open request for $c_{\bar{x}_j}$ with cost $3(n - j) + 2$. At vertex ψ , Player 1 picks a clause C_j by moving to its vertex. Since $\alpha(\psi)$ evaluates to true, there exists a $k \in \{1, 2, 3\}$ with $\alpha(\ell_{j,k}) = true$. We pick $\sigma(\cdots C_j) = \ell_{j,k}$?

If $\ell_{j,k} = x_l$, then $\alpha(x_l) = true$ and hence, there is an open request for c_{x_l} . As argued previously, this request is then answered with cost $3n + 5$, since we picked the gadget corresponding to x_l . Similarly, if $\ell_{j,k} = \bar{x}_l$, then $\alpha(x_l) = false$ and thus there is an open request for $c_{\bar{x}_l}$, which is answered with cost $3n + 5$ as well. All other open requests are answered with cost at most $3n + 5$, as argued previously.

After this traversal of the final gadget, all requests are answered, and the play automatically moves to vertex a_1 to begin anew. The same reasoning then applies ad infinitum. Thus, Player 0 is able to answer all requests with a cost of at most $3n + 5$.

Now assume that φ evaluates to false. Then, irrespective of the choices made by Player 0 when constructing α in the first part of the arena, Player 1 can pick truth values for the universally quantified variables such that $\alpha(\psi)$ evaluates to false and then pick a clause C_j such that $\alpha(C_j)$ evaluates to false. Hence, Player 0 has to pick some $\ell_{j,k}$ with $\alpha(\ell_{j,k}) = false$. If $\ell_{j,k} = x_l$, then there is an open request for $c_{\bar{x}_j}$ at $x_{l,1}$!, which is answered with cost $3n + 6$. Similarly, if $\ell_{j,k} = \bar{x}_l$, then $\alpha(x_l) = true$, hence there is an open request for c_{x_l} , which is also answered with cost $3n + 6$. Thus, in each round Player 1 can open a request that is only answered with cost at least $3n + 6$, i.e., Player 0 has no strategy with cost $3n + 5$. \square

5. MEMORY REQUIREMENTS OF OPTIMAL STRATEGIES IN PARITY GAMES WITH COSTS

Next, we study the memory needed by both players to play optimally in parity games with costs. Recall that Player 0 always has a positional winning strategy if she wins the game, while Player 1 requires infinite memory. In contrast, our main result of this section shows that the memory requirements of optimal strategies are exponential for Player 0, i.e., playing optimally comes at a price in terms of memory, too.

Standard complexity theoretic assumptions already rule out the existence of small strategies: If they existed, guessing and verifying such a strategy would contradict PSPACE-completeness of solving finitary parity games with respect to a given bound. However, here we explicitly present games in which either player requires exponential memory, which we later use to demonstrate gradual tradeoffs between memory and cost in Section 6. We obtain our lower bound by a generalization of a construction of Chatterjee and Fijalkow [CF13] which yielded a linear lower bound.

First, however, let us state a corollary of the construction of the parity game \mathcal{G}' in the proof of Lemma 4.6, which gives an exponential upper bound on the necessary memory states for both players. Recall that the memory structure used in that proof has one counter




Figure 6: The gadgets for Player 1 (left) and Player 0 (right).

with a range of size $b+2$ for each odd color. Furthermore, it has an additional counter that is bounded by n , which counts the number of times the bound b is exceeded. Using similar techniques to [FZ14b], it is possible to remove the overflow counter for Player 0: She can play assuming the largest value for this latter counter that still allows her to win.

Corollary 5.1. *Let \mathcal{G} be a parity game with costs containing n vertices and d odd colors.*

- *If Player 0 has a strategy σ for \mathcal{G} with $\text{Cst}(\sigma) = b$, then she also has a strategy σ' with $\text{Cst}(\sigma') \leq b$ and $|\sigma'| = (b+2)^d$.*
- *If Player 1 has a strategy τ for \mathcal{G} with $\text{Cst}(\tau) = b$, then he also has a strategy τ' with $\text{Cst}(\tau') \geq b$ and $|\tau'| = n(b+2)^d$.*

Again, our matching lower bounds already hold for finitary parity games, i.e., parity games with costs in which all edges are increment-edges. These proofs reuse principles underlying the PSPACE-hardness proof presented in Section 4.2, e.g., having a fixed bound requires a player, in the worst case, to store all open requests in order to answer them timely.

We begin by showing the lower bound for Player 0.

Theorem 5.2. *For every $d \geq 1$, there exists a finitary parity game \mathcal{G}_d such that*

- *\mathcal{G}_d has d odd colors and $|\mathcal{G}_d| \in \mathcal{O}(d^2)$,*
- *Player 0 has a strategy σ in \mathcal{G}_d with $\text{Cst}(\sigma) = d^2 + 2d$, but*
- *every strategy σ for Player 0 in \mathcal{G}_d with $\text{Cst}(\sigma) \leq d^2 + 2d$ has size at least 2^{d-1} .*

Proof. Let $d \geq 1$. We construct a finitary parity game \mathcal{G}_d that has the stated properties. To this end, after defining \mathcal{G}_d , we construct a strategy with cost $d^2 + 2d$ for Player 0 and argue that it is optimal, followed by the proof that every optimal strategy has at least size 2^{d-1} .

The game \mathcal{G}_d is played in rounds. In each round, which starts at the initial vertex of \mathcal{G}_d , Player 1 poses d requests for odd colors in the range 1 through $2d-1$. Subsequently, Player 0 gives d answers using colors in the range 2 through $2d$. If she recalls the choices made by Player 1 in the first part of the round, she is able to answer each request optimally. Otherwise, we show that Player 1 can exploit her having not enough memory in order to force requests to go unanswered for more than $d^2 + 2d$ steps. After each round, the play returns to the initial vertex in order to allow for infinite plays.

The arena \mathcal{A} consists of gadgets that each allow exactly one request or response to be made. Moreover, each path through a gadget has the same length $d+2$, including the edge connecting a gadget to its successor. However, low-priority requests and responses are made earlier than high-priority ones when traversing such a gadget, due to its structure. We show both gadgets in Figure 6. The dashed lines denote the edges to the pre- and succeeding gadget and the edge between the final and the initial gadget. As the owner of the succeeding vertex depends on the succeeding gadget's owner, we draw it as a diamond.

More precisely, the arena \mathcal{A} , consists of d repetitions of the gadget for Player 1, followed by d repetitions of the gadget for Player 0. The initial vertex v_I of the arena is the top-left vertex of the first gadget for Player 1. Moreover, the final gadget of Player 0 has a single back-edge to the initial vertex. Clearly, \mathcal{A} satisfies the first statement of the theorem.

Similarly to the proof of Lemma 4.7, it suffices to consider finite plays infixes. Even though the requests are not necessarily all answered after each round, we argue that Player 0 can always do so while playing optimally.

We now construct an optimal strategy from v_I for Player 0. In order to play optimally, Player 0 needs to track the requests made by Player 1 in the first part of each round. Instead of tracking each request precisely, however, it suffices to store the order in which the relevant requests were posed. In order to use memory efficiently, we do not initialize our memory structure with the empty sequence, but rather with the memory element encoding that all requests are relevant. Recall that relevant requests can only be opened by visiting some larger color than all currently open requests. Hence, we define the set of strictly increasing odd sequences

$$\begin{aligned} IncSeq_d = \{ & (c_1, \dots, c_d) \mid 1 \leq c_1 \leq \dots \leq c_d = 2d - 1, \\ & c_j \neq 2d - 1 \text{ implies } c_j < c_{j+1}, \text{ all } c_j \text{ are odd} \} \end{aligned}$$

and use them as the set of memory states $M_d = IncSeq_d$.⁴ Note that $|M_d| = 2^{d-1}$, as each increasing sequence is isomorphic to a subset of $\{1, 3, 5, \dots, 2d - 3\}$. In order to obtain \mathcal{M} , we define $m_I = (1, 3, \dots, 2d - 3, 2d - 1)$. Moreover, we define the update function Upd as follows: $Upd(m, e) = m_I$ if e leads to the initial vertex of \mathcal{G}_d , and as $Upd(m, e) = m$, if e leads to some other vertex of even color. If e leads to the (unique) vertex of odd color c in the j -th gadget of Player 1, however, we differentiate two cases. Let $m = (c_1, \dots, c_d)$. If $c_j \geq c$, then $Upd(m, e) = m$. Otherwise, we define

$$Upd(m, e) = (c_1, \dots, c_{j-1}, c, c + 2, c + 4, \dots, 2d - 1, 2d - 1, \dots, 2d - 1) .$$

Note that this definition of Upd implies that the memory state is fixed once a partial play leaves the gadgets of Player 1, remains unchanged throughout the traversal of the gadgets of Player 0 and is only reset upon moving to the initial vertex of \mathcal{G}_d . Finally, we define the next-move function Nxt such that, if the play leaves the gadgets of Player 1 with memory state $m = (c_1, \dots, c_d)$, then Player 0 moves to color $c_j + 1$ in her j -th gadget.

Now consider a play in which Player 0 plays according to this strategy and consider the request for color c made by Player 1 in his j -th gadget. If the request for color c is the largest request made by Player 1 so far in the current round, then Player 0 answers the j -th request from Player 1 in her j -th gadget. The cost of this request then consists of three components. First, the play has to leave Player 1's j -th gadget, incurring a cost of $d - (c + 1)/2 + 2$. Then, the play passes through $d - 1$ gadgets, incurring a cost of $d + 2$ in each. Finally, moving to color $c + 1$ in Player 0's gadget incurs a cost of $(c + 1)/2$. Hence, answering Player 1's request incurs a cost of $d^2 + 2d$. If there, however, exists some $j' < j$ such that Player 1 has already requested some color $c' > c$ in his j' -th gadget, then the request for color c' will be answered in gadget j' of Player 0. Hence, Player 0 answers the request with cost at most $(d - 1)(d + 2) < d^2 + 2d$.

⁴This definition differs from the one presented in the conference version [WZ16]. There, the padding to length d (here with requests for $2d - 1$) was kept implicit. We show that an optimal strategy for Player 0 can be implemented with $|M_d| = 2^{d-1}$ many memory states. This invalidates the lower bound of $2^d - 2$ claimed in the conference version.

This cost is indeed optimal. Consider the play in which Player 1 always requests $2d-1$. Even if Player 0 answers this request in her first gadget, it still incurs a cost of $d^2 + 2d$.

It remains to show that no optimal strategy of size less than $|M_d|$ exists. To this end we show that Player 1 can exploit a strategy of Player 0 with less than $|M_d|$ memory states and pose requests that will be answered suboptimally.

We associate with each $m \in M_d$ the partial play $req(m)$ which starts in the initial vertex, where Player 1 requests the colors occurring in m in order. Clearly, $m \neq m' \in M_d$ implies $req(m) \neq req(m')$.

Let σ be a strategy for Player 0 that is implemented by (M, m_I, Upd) with $|M| < |M_d|$ and let $m \in M$. Due to the pigeon-hole principle, there exist $m'_1 \neq m'_2 \in M_d$, such that $\text{Upd}^+(m, req(m'_1)) = \text{Upd}^+(m, req(m'_2))$, i.e., Player 0 answers both sequences of requests in the same way. Since $req(m'_1) \neq req(m'_2)$, there exists a gadget of Player 1 in which the requests posed during $req(m'_1)$ and $req(m'_2)$ differ. Pick j as the minimal index of such a gadget and assume that in his j -th gadget, Player 1 requests color c during $req(m'_1)$, and color c' during $req(m'_2)$, where, w.l.o.g., $c < c'$. If Player 0 has already answered the request for c' before her j -th gadget, then some earlier request is not answered optimally when reacting to $req(m'_2)$, as requests are posed in strictly increasing order. Thus, assume that the request for color c' has not been answered upon entering Player 0's j -th gadget. If she visits some color $c'' < c'$ in this gadget, she will only answer c' in some later gadget, thereby incurring a cost of more than $d(d+2) = d^2 + 2d$ when reacting to $req(m'_2)$. If she visits some color $c'' \geq c'$, then she does not answer the request for c optimally, thus incurring a cost of at least $d^2 + 2d + (c' - c)/2 > d^2 + 2d$ when reacting to $req(m'_1)$. Hence, one of the sequences of requests $req(m'_1)$ or $req(m'_2)$ leads to Player 0 answering at least one request non-optimally. As such sequences m'_1 and m'_2 exist for each $m \in M$, Player 1 can force such an “expensive” request in each round. Thus, $\text{Cst}(\sigma) > d^2 + 2d$, i.e., σ is not optimal. \square

After having shown that Player 0 requires exponential memory to keep the cost of the play below a given bound, we now show a similar result for Player 1: In general, he also requires exponential memory to enforce costs above a given bound.



Theorem 5.3. *For every $d \geq 1$, there exists a finitary parity game \mathcal{G}_d such that*

- \mathcal{G}_d has $2d$ odd colors and $|\mathcal{G}_d| \in \mathcal{O}(d)$,
- Player 1 has a strategy τ in \mathcal{G}_d with $\text{Cst}(\tau) = 5(d-1) + 7$, but
- every strategy τ for Player 1 in \mathcal{G}_d with $\text{Cst}(\tau) \geq 5(d-1) + 7$ has size at least 2^d .

Proof. Fix some $d \geq 1$. Similarly to the previous proof, we construct an arena \mathcal{A} using two kinds of gadgets, one for each player, each of which is repeated d times. In \mathcal{A} , first Player 0 opens d requests and subsequently picks one of these requests to be answered. If Player 1 recalls the requests, then he can delay the answer to any chosen request up to $5(d-1) + 7$ steps. Otherwise, Player 0 can find a sequence of requests that is answered optimally.

We show the gadgets in Figure 7 together with their coloring, and call them G_j^0 and G_j^1 for the j -th gadget of Player 0 and Player 1, respectively. The gadget G_j^0 contains the colors $0, 4j-3$, and $4j-1$, while the gadget G_j^1 contains the colors $0, 4j-2$, and $4j$. The complete arena \mathcal{A} is shown in Figure 8. We fix the initial vertex of G_1^0 to be the initial vertex v_I of \mathcal{G}_d . Clearly, \mathcal{A} satisfies the first item of the theorem.

The game is played in rounds. Each such round starts and ends in the initial vertex v_I that answers every request. Thus, it suffices to analyze a single round of the game: In each round, Player 0 starts by posing d requests. With her j -th request, she may either request

Figure 7: The gadgets G_j^0 (above) and G_j^1 (below).Figure 8: The arena \mathcal{A} of \mathcal{G}_d showing necessity of exponential memory for Player 1.

the color $4j - 3$ and take three steps before leaving G_j^0 , or she may request the color $4j - 1$ and take a single step before leaving G_j^0 . After posing d requests, Player 0 then moves to some G_j^1 for $1 \leq j \leq d$ while answering all requests for colors $c \leq 4(j - 1)$. In G_j^1 , Player 1 answers the request posed in G_j^0 . After he has done so, all requests are reset and the next round begins. All requests posed in $G_{j'}^0$ for $j' < j$ are answered before entering G_j^1 due to the structure of the arena. Similarly, all requests posed in $G_{j'}^0$ for $j' > j$ are answered immediately after leaving G_j^1 . We show that by remembering all requests, which takes 2^d memory states, Player 1 can ensure that one request is only answered with cost $5(d - 1) + 7$. If he uses less memory states, we show that Player 0 is able to answer every request with cost $5(d - 1) + 6$, i.e., that strategy has cost at most $5(d - 1) + 6$.

First, consider the strategy τ for Player 1 that is defined as follows: During Player 0's part of the round, Player 1 stores the requests that she makes using 2^d memory states. Assume Player 0 then moves to G_j^1 . If Player 0 requested color $4j - 3$ during her part of the round, Player 1 picks the upper branch shown in Figure 7, while he chooses the lower branch in case Player 0 requested color $4j - 1$.

It remains to argue that this strategy indeed enforces a cost at least $5(d - 1) + 7$. Consider a play prefix that starts in v_I and ends in c_j for some $1 \leq j \leq d$. At this point, the request with the highest cost incurred so far is either a request for color $4j - 3$ with

cost $5(d-1) + 4$, or a request for color $4j-1$ with cost $5(d-1) + 2$. All requests for colors $c < 4j-3$ have already been answered due to the structure of the arena. By coloring the vertices resulting from the subdivision of the edges labeled with cost **5** with the color of the target of these edges, these requests incur cost at most $5(d-1) + 5$. Moreover, all requests for colors $c > 4j-1$ have incurred a cost of at most $5(d-1) - 1$. Now, assume that Player 0 enters G_j^1 . First, consider the case with an open request for color $4j-3$. Then Player 1 moves through the lower branch of his gadget, answering this request with cost $5(d-1) + 7$. If there is, however, an open request for color $4j-1$, then Player 1 moves through the upper branch of his gadget, answering the open request with cost $5(d-1) + 7$ as well. Thus, the strategy τ has cost $5(d-1) + 7$. All other requests are answered immediately after leaving the gadget at vertex v_I with a cost of at most $5(d-1) + 5$.

We now show that Player 1 indeed needs at least 2^d many memory states to enforce a cost of at least $5(d-1) + 7$. Towards a contradiction assume that Player 1 has a finite-state strategy τ with $\text{Cst}(\tau) \geq 5(d-1) + 7$ that is implemented by a memory structure (M, m_I, Upd) , where $|M| < 2^d$. We inductively construct a play ρ consistent with τ such that $\text{Cst}(\rho) \leq 5(d-1) + 6$. Assume we have already defined a prefix π of ρ ending in the initial vertex v_I . We determine a sequence of d requests and a choice of $1 \leq j \leq d$ and prolong π by letting Player 0 first pick the sequence of requests and then move into some gadget G_j^1 . Then, Player 1 applies his strategy, which leads back to the initial vertex.

To this end, let $m = \text{Upd}^+(m_I, \pi)$. Since $|M| < 2^d$, and since there exist 2^d play infixes leading from the unique successor of v_I to c_1 , there exist two such infixes π_1 and π_2 , such that $\text{Upd}^+(m, \pi_1) = \text{Upd}^+(m, \pi_2)$. Let j be minimal such that the choices made in G_j^0 by Player 0 differ in π_1 and π_2 , and w.l.o.g. assume that Player 0 poses a request for color $4j-3$ when playing π_1 , while she poses a request for color $4j-1$ when playing π_2 .

Now consider the response of Player 1 consistent with τ if Player 0 moves to G_j^1 after the play prefix $\pi\pi_1$ and note that this response is the same as the one to the play prefix $\pi\pi_2$ due to $\text{Upd}^+(m, \pi_1) = \text{Upd}^+(m, \pi_2)$. If Player 1 traverses the upper branch of G_j^1 after witnessing $\pi\pi_1$ or $\pi\pi_2$, then he answers the request for $4j-3$ posed during the traversal of π_1 with cost $5(d-1) + 6$. If he, however, traverses the lower branch of G_j^1 after witnessing $\pi\pi_1$ or $\pi\pi_2$, then he answers the request for $4j-1$ posed during π_2 with cost $5(d-1) + 6$. In the former case, we continue π by letting Player 0 play according to π_1 , while in the latter case we continue π by letting her play according to π_2 . In either case, we move to G_j^1 afterwards. In G_j^1 , Player 1 plays consistently with τ .

In both cases all requests posed in gadgets $G_{j'}^0$, for $j' < j$ are answered after at most $5(d-1) + 5$ steps, namely upon reaching the first vertex of the subdivision of the edge leading to the vertex $c_{j'+1}$. Also, the request posed in G_j^0 is answered after $5(d-1) + 6$ steps. Finally, all requests posed in gadgets $G_{j'}^0$, for $j' > j$ are answered after at most $5(d+j-j'-1) + 10 \leq 5(d-1) + 5$ steps upon visiting the vertex v_I .

Since all requests are reset when reaching v_I , and since the reasoning above holds true for any memory state m reached at the end of any π as above, the play ρ resulting from an inductive application of this construction has $\text{Cst}(\rho) \leq 5(d-1) + 6$. Since ρ is consistent with τ , this contradicts $\text{Cst}(\tau) \geq 5(d-1) + 7$ and concludes the proof of the theorem. \square

Note that Player 1 does not win the games \mathcal{G}_d with regards to the classical finitary parity condition, i.e., he cannot unbound the cost of an open request arbitrarily.

6. TRADEOFFS BETWEEN MEMORY AND COST

In the previous section, we have shown that an optimal strategy for either player in a parity game with costs requires exponential memory in general. In contrast, winning strategies of minimal size for Player 0 in parity games with costs are known to be positional [FZ14b], while winning strategies for Player 1 require infinite memory already in the case of finitary parity games in order to violate every bound infinitely often. Here we show that, in general, there exists a gradual tradeoff between the size and the cost of a strategy for both players. For Player 0, this means that she can choose to lower the guaranteed bound b by using a larger winning strategy. Dually, Player 1 can reduce the amount of memory he has to use by not violating every bound, but only a fixed bound b .

Theorem 6.1. *Fix some $d \geq 1$ and let the game \mathcal{G}_d be as defined in the proof of Theorem 5.2. For every j with $1 \leq j \leq d$ there exists a strategy σ_j for Player 0 in \mathcal{G}_d such that*

- $d^2 + 3d - 1 = \text{Cst}(\sigma_1) > \text{Cst}(\sigma_2) > \dots > \text{Cst}(\sigma_d) = d^2 + 2d$, and
- $1 = |\sigma_1| < |\sigma_2| < \dots < |\sigma_d| = 2^{d-1}$.

Also, for every strategy σ' for Player 0 in \mathcal{G}_d with $\text{Cst}(\sigma') \leq \text{Cst}(\sigma_j)$ we have $|\sigma'| \geq |\sigma_j|$.

Proof. Recall that we defined the set of strictly increasing odd sequences IncSeq_d in the proof of Theorem 5.2 and showed that a memory structure using IncSeq_d as memory states implements an optimal strategy with cost $d^2 + 2d$. Intuitively, such a strategy stores up to $d - 1$ requests made by Player 1 in his part of each round, as the final element of each increasing sequence is fixed to be $2d - 1$. The idea behind the construction of the strategies σ_j is to restrict the memory of Player 0 such that she can only store up to $j - 1$ requests. In the extremal cases of $j = 1$ and $j = d$ this implements a positional strategy and the strategy from the proof of Theorem 5.2, respectively.

We implement σ_j by again using strictly increasing odd sequences, where we restrict the maximal number of entries that differ from the maximal value of $2d - 1$. Hence, in strategy σ_j , Player 0 stores at most $j - 1$ requests.

To this end, we define the length-restricted set of strictly increasing odd sequences

$$\text{IncSeq}_d^j = \text{IncSeq}_d \cap \{s = (c_1, \dots, c_{j-1}, 2d - 1, \dots, 2d - 1) \mid s \in \mathbb{N}^d\}$$

and pick $M_d^j = \text{IncSeq}_d^j$. Note that $M_d^d = M_d$ as defined in the proof of Theorem 5.2 and that M_d^1 is a singleton set. Clearly, the second claim of the theorem holds true, since $\text{IncSeq}_d^{j-1} \subsetneq \text{IncSeq}_d^j$ for each $d \geq 1$ and each j with $1 \leq j \leq d$. The initial memory state is $(1, 3, \dots, 2j - 3, 2d - 1, \dots, 2d - 1)$, the update function only stores the first j relevant requests, and the next-move function Nxt_j for Player 0 is the same as that from the proof of Theorem 5.2 in order to obtain the memory structure \mathcal{M}_j implementing σ_j via Nxt_j .

It remains to show that each strategy σ_j realizes a cost of $d^2 + 3d - j$ and that it is minimal for its respective cost. To this end, we fix some j with $1 \leq j \leq d$ for the remainder of this proof. First, we show that Player 1 can enforce a cost of $d^2 + 3d - j$ if Player 0 plays consistently with σ_j . Intuitively, Player 1 fills the memory of Player 0 as quickly as possible, and requests the minimal color that has not yet been requested afterwards. Thus, he maximizes the gap between the smallest unstored request and the “default” answer of $2d$.

More precisely, in each turn Player 1 requests the colors $1, 3, \dots, 2j - 3, 2j - 1, 2j - 1, \dots, 2j - 1$. Playing consistently with σ_j , Player 0 answers these requests with $2, 4, \dots, 2j - 2, 2d, 2d, \dots, 2d$. Hence, the cost of the resulting play is that incurred by answering a request for $2j - 1$ in the j -th gadget of Player 1 with $2d$ in the j -th gadget of Player 0. As argued

in the proof of Theorem 5.2, the cost incurred by such a request-response-pair amounts to

$$[d - j + 2] + [(d - 1)(d + 2)] + d = d^2 + 3d - j.$$

As the game restarts after Player 0's turn, Player 1 can enforce this cost infinitely often. Hence, $\text{Cst}(\sigma_j) \geq d^2 + 3d - j$.

This sequence of requests is indeed optimal for Player 1, i.e., he cannot enforce a higher cost. Assume that Player 1 does not pose requests as specified above, but poses the requests c_1, \dots, c_d . Then either there exist some k and k' with $k < k' \leq j$, such that $c_k \geq c_{k'}$, or there exists a $k \leq j$ with $2j - 1 < c_k \leq 2d - 1$.

In the former case, let k' be minimal such that such a k exists. Player 0 answers the first $k' - 1$ requests optimally before answering all remaining requests with costs at most $(d - 1)(d + 2)$, as she ignores the request for $c_{k'}$. In the latter case, Player 0 again answers all requests up to the first request as described above optimally. Afterwards, she answers all succeeding requests with cost at most $d^2 + 2d + (2d - 1 - c_k)/2 \leq d^2 + 3d - j$. Hence, there exists no play ρ consistent with σ_j and $\text{Cst}(\rho) > d^2 + 3d - j$.

To conclude the proof, we observe that there exists no strategy σ' with $|\sigma'| < |\sigma_j|$ and $\text{Cst}(\sigma') \leq \text{Cst}(\sigma_j)$. The argument is nearly identical to the argument of minimality of the strategy constructed in the proof of Theorem 5.2 and can in fact be obtained by replacing all occurrences of 2^{d-1} and $d^2 + 2d$ by $|\sigma_j|$ and $d^2 + 3d - j$, respectively. Hence, the strategies σ_j are minimal for their respective cost. \square

The similar result for Player 1 has a much simpler proof. Again, it suffices to reuse the games from the lower bound.

Theorem 6.2. *For each $d \geq 1$ there exists a finitary parity game \mathcal{G}_d with $|\mathcal{G}_d| \in \mathcal{O}(d^2)$, such that for every j with $1 \leq j \leq d$, there exists a strategy τ_j for Player 1 in \mathcal{G}_d such that*

- $7 = \text{Cst}(\tau_1) < \text{Cst}(\tau_2) < \dots < \text{Cst}(\tau_d) = 5(d - 1) + 7$, and
- $2 = |\tau_1| < |\tau_2| < \dots < |\tau_d| = 2^d$.

Moreover, for every strategy τ' with $\text{Cst}(\tau') \geq \text{Cst}(\tau_j)$, we have $|\tau'| \geq |\tau_j|$.

Proof. For each j with $1 \leq j \leq d$, let \mathcal{G}'_j be the game with $2j$ odd colors defined in Theorem 5.3. We construct \mathcal{G}_d such that it contains a dedicated initial vertex v_I of color 0, from which Player 1 may choose to move to the initial vertex of any of the \mathcal{G}'_j . Once the play ρ has moved into some \mathcal{G}'_j , it never leaves that part of the arena, i.e., the suffix starting at the second position of ρ is a play of \mathcal{G}'_j .

For each j , the strategy τ_j defined in Theorem 5.3, augmented by a single move from v_I to the sub-game \mathcal{G}'_j , satisfies the properties above. Moreover, as the sub-games \mathcal{G}'_j are isolated from each other, each strategy τ' for Player 1 in \mathcal{G}_d can be trivially transformed into a strategy for him in the subgame \mathcal{G}'_j that τ' chooses at the beginning of \mathcal{G}_d . Hence, every strategy τ_j is of minimal size for the cost that it realizes. \square

7. PARITY GAMES WITH INTEGER COSTS

In this section, we extend the model of parity games with costs to non-negative integer costs. Recall that the cost function $\text{Cst}: E \rightarrow \{\varepsilon, \mathbf{i}\}$ assigns to each edge of \mathcal{G} either no cost, or a unit cost. We now consider cost functions that assign to each edge a natural number. Formally, let \mathcal{A} and Ω be an arena and a coloring as defined previously and let $\text{Cst}: E \rightarrow \mathbb{N}$ be a cost function. We call $\mathcal{G} = (\mathcal{A}, \text{CostParity}(\Omega, \text{Cst}))$ a parity game with

(non-negative) integer-valued costs. Now the cost of a play infix is the sum of the costs of its edges and the winning condition $\text{CostParity}(\Omega, \text{Cst})$ is defined analogously to the case of abstract costs. We denote the highest cost assigned by Cst by W if Cst is clear from the context. Since we assume Cst to be encoded in binary encoding, W may be of exponential size in the number of bits required to describe Cst . To distinguish these games from those we considered previously, we call a game $\mathcal{G} = (\mathcal{A}, \text{CostParity}(\Omega, \text{Cst}))$ with $\text{Cst}: E \rightarrow \{\varepsilon, \mathbf{i}\}$ a parity game with abstract costs and omit this qualifier if it is clear which game is meant from the context.

Formally, the size of a parity game with integer-valued costs is defined as $|\mathcal{G}| = |\mathcal{A}| + \log W$. Similarly to the abstract case, we call an edge e an ε - or an increment-edge if $\text{Cst}(e) = 0$ or $\text{Cst}(e) > 0$, respectively.

We first show that determining the winner in \mathcal{G} with respect to a given bound b is as easy as determining the winner in a parity game with abstract costs. Since the former problem subsumes the latter one, this yields PSPACE-completeness of the former problem via Lemma 4.7. Afterwards, we briefly discuss the memory requirements of both players as well as the potential tradeoffs present in this setting. In particular, we show that the tradeoffs between the cost and the size of a strategy are, in general, more pronounced in the case of integer-valued costs.

7.1. The Complexity of Solving Parity Games with Integer-Valued Costs Optimally. We first examine the complexity of playing optimally in parity games with integer-valued costs and show that this problem is not harder than the special case of abstract costs. Afterwards, we argue that this implies an exponential upper bound on the size of optimal strategies for both players.

Theorem 7.1. *The following problem is PSPACE-complete: “Given a parity game with integer-valued costs \mathcal{G} and a bound $b \in \mathbb{N}$ (in binary encoding), does Player 0 have a strategy σ for \mathcal{G} with $\text{Cst}(\sigma) \leq b$?”*

PSPACE-hardness of this problem follows directly from Lemma 4.7, as every finitary parity game is a parity game with integer-valued costs. Thus, the remainder of this section is dedicated to showing PSPACE-membership of the given problem. We follow the same approach that solved the case of abstract costs, i.e., we first extend \mathcal{G} with the same memory structure yielding the qualitative parity game \mathcal{G}' , which is still only of exponential size. In contrast to the abstract case, however, plays of \mathcal{G}' are not settled after polynomially many moves, as the upper bound on b is now exponential. Thus, we construct a novel finite-duration game $\mathcal{G}'_{\curvearrowright}$, which extends \mathcal{G}'_f by a shortcut mechanism. This reduces unsettled play prefixes to at most polynomial length. By showing equivalence of \mathcal{G}' and $\mathcal{G}'_{\curvearrowright}$ and by the fact that $\mathcal{G}'_{\curvearrowright}$ can be simulated on a polynomially time-bounded alternating Turing machine, we obtain PSPACE-membership of the given problem. Due to the similarities between the two proofs, we reuse notation from Section 4.1.

We begin by bounding the parameter b from above. While we could assume $b < n$ in the abstract case, this is not the case with integer-valued costs anymore: Here, a positional winning strategy for Player 0, which still exists, has cost at most nW .

Corollary 7.2. *Let \mathcal{G} be a parity game with integer-valued costs with n vertices and largest cost W . If Player 0 wins \mathcal{G} , then she has a strategy σ with $\text{Cst}(\sigma) \leq nW$, i.e., an optimal strategy has cost at most nW .*

Due to this corollary, if $b \geq nW$, we check for the existence of a winning strategy for Player 0 in $(\mathcal{A}, \text{CostParity}(\text{Cst}', \Omega))$, where $\text{Cst}'(e) = \varepsilon$, if $\text{Cst}(e) = 0$, and $\text{Cst}'(e) = \mathbf{i}$ otherwise. This can be decided in polynomial space due to Theorem 3.1(3) and is equivalent to deciding the given problem. Hence, we assume $b < nW$ for the remainder of this proof.

Given \mathcal{G} and b , we define $\mathcal{G}' = (\mathcal{A} \times \mathcal{M}, \text{Parity}(\Omega'))$ as before, with $M = [n+1] \times R$, where $[n+1]$ implements the overflow counter and $R = (\{\perp\} \cup [b+1])^D$ is the set of request functions. Since the proof of Lemma 4.2 does not rely on $b < n$, Player 0 wins \mathcal{G}' if and only if she has a strategy of cost at most b for \mathcal{G} . Recall that v'_I is the initial vertex of \mathcal{A}' and that we call a play settled if it either contains a dominating cycle or if the overflow counter is saturated.

In contrast to the abstract case, plays of \mathcal{G}' are not settled after polynomially many steps, as up to nW many increment-edges may have to be traversed before an overflow occurs, which is exponential in the size of the game. In order to be able to declare the winner of a play after polynomially many steps, we define $\mathcal{G}'_{\curvearrowright}$. In this game, we skip infixes that form cycles when projected to \mathcal{A} , during which the costs incurred by the relevant requests increase, but the set of these requests is stable. We say that such infixes satisfy the shortcut criterion. When such an infix π is traversed, Player 1 has demonstrated that he can increase the cost of the currently open relevant requests by $\text{Cst}(\pi)$ without answering any such request, i.e., he can strictly improve his situation. In $\mathcal{G}'_{\curvearrowright}$, Player 1 does not have to traverse this infix (or others) until an overflow occurs, but the play continues at a position that is as “close” to an overflow as Player 1 can get by repeatedly traversing π without actually causing an overflow.

Formally, we say that a play infix $\pi = (v_0, o_0, r_0) \cdots (v_j, o_j, r_j)$ in \mathcal{A}' satisfies the shortcut criterion if $v_0 = v_j$, if $o_0 = o_j$, if $\text{RELREQ}(r_0) = \text{RELREQ}(r_{j'}) \neq \emptyset$ for all $0 \leq j' \leq j$, if $\text{Cst}(\pi) > 0$, and if $r_j(c^*) + \text{Cst}(\pi) \leq b$ for $c^* = \arg \max_c r_j(c)$. Note that the condition $r_j(c^*) + \text{Cst}(\pi) \leq b$ is equivalent to demanding $\text{Cst}(\pi) \leq \frac{b - r_0(c^*)}{2}$.

For the sake of readability, we refrain from defining the arena underlying $\mathcal{G}'_{\curvearrowright}$ formally. We rather define the set of play prefixes of $\mathcal{G}'_{\curvearrowright}$ inductively, which are subsequences of plays in \mathcal{A}' . In particular, the vertices in $\mathcal{G}'_{\curvearrowright}$ inherit the coloring from \mathcal{G}' . First, v'_I is a play prefix of $\mathcal{G}'_{\curvearrowright}$. Now, let $\pi = (v_0, o_0, r_0) \cdots (v_j, o_j, r_j)$ be a play prefix of $\mathcal{G}'_{\curvearrowright}$ and let (v', o', r') be a successor of (v_j, o_j, r_j) in \mathcal{A}' . If there exists no j' such that the infix $\pi' = (v_{j'}, o_{j'}, r_{j'}) \cdots (v_j, o_j, r_j)(v', o', r')$ of π satisfies the shortcut criterion, then $\pi(v', o', r')$ is a play prefix of $\mathcal{G}'_{\curvearrowright}$. If, however, such a j' exists, let it be the maximal one, let $c^* = \arg \max_c r'(c)$, let $s = \text{Cst}(\pi')$, and let $t = \max\{t' > 0 \mid r'(c^*) + s \cdot t' \leq b\}$. Moreover, define r^* as $r^*(c) = r'(c) + s \cdot t$ if $r'(c) \neq \perp$ and $r^*(c) = \perp$, otherwise. Then, $(v_0, o_0, r_0) \cdots (v_j, o_j, r_j)(v', o', r^*)$ is a play prefix of $\mathcal{G}'_{\curvearrowright}$. We define the cost of the transition from (v_j, o_j, r_j) to (v', o', r^*) as $\text{Cst}(v_j, v') + s \cdot t$ in $\mathcal{G}'_{\curvearrowright}$ and redefine the notions of the cost of a play accordingly in order to obtain uniform notation.

Moreover, we use the following notions:

- The transition from (v_j, o_j, r_j) to (v', o', r^*) is a *shortcut*.
- The infix $(v_{j'}, o_{j'}, r_{j'}) \cdots (v_j, o_j, r_j)(v', o', r^*)$ is a *shortcut cycle*, where we call the vertex (v', o', r^*) its *destination*.⁵
- The infix $(v_{j'}, o_{j'}, r_{j'}) \cdots (v_j, o_j, r_j)(v', o', r')$ is the *detour* corresponding to the shortcut cycle, with *destination* (v', o', r') .

⁵Note that, similarly to the case of dominating cycles, a shortcut cycle is only a cycle when projected to its first component.

Let $(v_j, o_j, r_j)(v', o', r^*)$ be a shortcut with (v', o', r') as the target of its associated detour as defined above. We obtain $\text{RELREQ}(r') = \text{RELREQ}(r^*)$ and $r' \sqsubseteq r^*$. Moreover, if c^* is the open request that has incurred the highest cost in $r_{j'}$, this shortcut closes at least half the distance between $r_{j'}(c^*)$ and b , i.e., $r^*(c^*) \geq r_{j'}(c^*) + \frac{b-r_{j'}(c^*)}{2}$. Hence, no infix π containing a shortcut π' satisfies the shortcut criterion, as the cost of π' is already larger than half the cost that would cause an overflow. Thus, π violates the condition that it must be able to be traversed at least twice without causing an overflow.

However, a shortcut may be part of a dominating cycle. If the maximal color along the detour associated with some shortcut is odd, then the shortcut cycle is an odd dominating cycle, i.e., a play may be settled due to a dominating cycle that contains a shortcut.

The introduction of shortcuts in $\mathcal{G}'_{\curvearrowright}$ ensures that plays in $\mathcal{G}'_{\curvearrowright}$ are settled in polynomial time. Fix $\ell_{\curvearrowright} = (\log(nW) + 1)(n + 1)^6$, which is polynomial in the size of \mathcal{G} .

Lemma 7.3. *Let π be a play prefix of $\mathcal{G}'_{\curvearrowright}$. If $|\pi| > \ell_{\curvearrowright}$, then π is settled.*

Proof. Let $\pi = (v_0, o_0, r_0) \cdots (v_j, o_j, r_j)$. First, due to the same argument as in the proof of Lemma 4.4, π may not contain a vertex repetition, as it is unsettled. Also note that π only contains at most n overflow positions, each type 1 infix contains at most n debt-free positions, each type 2 infix contains at most d request-adding positions and each type 3 infix contains at most d relevance-reducing positions, again due to the same arguments as in Lemma 4.4.

Fix a non-empty type 4 infix π_4 and recall that there is at least one request continuously open throughout π_4 . Let c^* be the request that has incurred the greatest cost s at the beginning of π_4 and note that the request for c^* is continuously open throughout π_4 , as π_4 has no relevance-reducing positions. We define the first *halving position* as the minimal position k such that $r_k(c^*) \geq s + \frac{b-s}{2}$ (if it exists). Inductively, if k is a halving position, then the minimal $k' > k$ with $r_{k'}(c^*) \geq r_k(c^*) + \frac{b-r_k(c^*)}{2}$ is a halving position as well (if it exists). Since at each halving position the difference between the currently incurred cost of c^* and the bound b has been halved when compared to the previous halving position, there exist at most $\log(b) \leq \log(nW)$ many such positions. Hence, splitting π_4 at its halving positions yields at most $\log(nW) + 1$ many infixes without overflow, debt-free, request-adding, relevance-reducing or halving positions. We say such an infix has type 4'.

Fix a non-empty type 4' infix $\pi_{4'}$. We show that $\pi_{4'}$ contains at most n increment-edges. Towards a contradiction assume that it contains $n + 1$ increment-edges and let c^* be the request that has incurred the highest cost s at the beginning of $\pi_{4'}$. Since there exist more than n increment-edges, there exist two such edges leading to the vertices (v_k, o_k, r_k) and $(v_{k'}, o_{k'}, r_{k'})$ with $k < k'$ and $v_k = v_{k'}$. If the cost of the infix $(v_k, o_k, r_k) \cdots (v_{k'}, o_{k'}, r_{k'})$ is larger than $\frac{b-s}{2}$, then $\pi_{4'}$ contains a halving position, which yields the desired contradiction. If the cost is lower, however, then this infix satisfies the shortcut condition, since $\pi_{4'}$ does not contain overflow positions and the relevant requests are stable throughout $\pi_{4'}$. Hence, the vertex $(v_{k'}, o_{k'}, r_{k'})$ is the destination of a shortcut cycle, which contradicts the infix from (v_k, o_k, r_k) to $(v_{k'}, o_{k'}, r_{k'})$ having a cost of less than $\frac{b-s}{2}$. Thus, $\pi_{4'}$ contains at most n increment-edges and, by splitting $\pi_{4'}$ at the increment-edges, we obtain a decomposition of $\pi_{4'}$ into at most $b + 1$ infixes, each without increment-edges and without request-adding, debt-free, and overflow positions.

Moreover, as none of the infixes contain increment-edges, they also contain no shortcuts. Hence, each infix is of type 5, i.e., at most of length n as argued in the proof of Lemma 4.4.

Aggregating all these bounds yields an upper bound of $(\log(nW) + 1)(n + 1)^6$ on the length of an unsettled play prefix π . \square

We again define the winner of a play ρ of $\mathcal{G}'_{\curvearrowright}$ such that Player 0 wins ρ if the minimal settled prefix of ρ is settled due to an even dominating cycle. Otherwise, Player 1 wins ρ . Due to Lemma 7.3, the winner of a play of $\mathcal{G}'_{\curvearrowright}$ is determined after finitely many moves, hence $\mathcal{G}'_{\curvearrowright}$ is determined [Zer13]. It remains to show that solving $\mathcal{G}'_{\curvearrowright}$ actually solves \mathcal{G}' .

Lemma 7.4. *Player 0 wins \mathcal{G}' if and only if she wins $\mathcal{G}'_{\curvearrowright}$.*

Proof. We first show that, if Player 0 wins $\mathcal{G}'_{\curvearrowright}$, then she also wins \mathcal{G}' . To this end, let $\sigma_{\curvearrowright}$ be a winning strategy for Player 0 in $\mathcal{G}'_{\curvearrowright}$. We construct a winning strategy σ' for her in \mathcal{G}' by mimicking the moves made in \mathcal{G}' in $\mathcal{G}'_{\curvearrowright}$ using a simulation function h mapping play prefixes in \mathcal{G}' to play prefixes in $\mathcal{G}'_{\curvearrowright}$. This simulation function satisfies the same invariant as in the proof of Lemma 4.5, i.e.:

Let π be consistent with σ' and end in (v, o, r) . Then, $h(\pi)$ is consistent with $\sigma_{\curvearrowright}$, is unsettled, and ends in (v, o', r') with $(o', r') \supseteq (o, r)$.

We define h and σ' inductively and simultaneously, starting with $h(v'_I) = v'_I$, which obviously satisfies the invariant. Now let π be a play prefix of \mathcal{G}' consistent with σ' , ending in (v, o, r) , and assume $h(\pi)$ is defined. Due to the invariant, $h(\pi)$ ends in some $(v, o_{\curvearrowright}, r_{\curvearrowright})$ with $(o, r) \sqsubseteq (o_{\curvearrowright}, r_{\curvearrowright})$. If $(v, o, r) \in V'_0$, let $\sigma_{\curvearrowright}(h(\pi)) = (v^*, o^*_{\curvearrowright}, r^*_{\curvearrowright})$ and define $\sigma'(\pi) = (v^*, o^*, r^*)$, where $(o^*, r^*) = \text{Upd}((o, r), (v, v^*))$. Otherwise, if $(v, o, r) \in V'_1$, let (v^*, o^*, r^*) be an arbitrary successor of (v, o, r) in \mathcal{A}' . In either case, let $\pi^* = \pi \cdot (v^*, o^*, r^*)$.

It remains to define $h(\pi^*)$. To this end, let $(o^*_{\curvearrowright}, r^*_{\curvearrowright})$ be the unique memory state such that $\pi^*_{\curvearrowright} = h(\pi) \cdot (v^*, o^*_{\curvearrowright}, r^*_{\curvearrowright})$ is a play prefix of $\mathcal{G}'_{\curvearrowright}$. If π^*_{\curvearrowright} is unsettled, we define $h(\pi^*) = \pi^*_{\curvearrowright}$. This choice satisfies the invariant: If the vertex $(v^*, o^*_{\curvearrowright}, r^*_{\curvearrowright})$ is the destination of a shortcut, then let $(v^*, o^*_{\rightarrow}, r^*_{\rightarrow})$ be the destination of its corresponding detour. We obtain $(o^*_{\curvearrowright}, r^*_{\curvearrowright}) \supseteq (o^*_{\rightarrow}, r^*_{\rightarrow}) \supseteq (o^*, r^*)$ due to Lemma 4.3. Otherwise, i.e., if $(v^*, o^*_{\curvearrowright}, r^*_{\curvearrowright})$ is not the destination of a shortcut, then Lemma 4.3 yields the invariant directly. Now consider the case that π^*_{\curvearrowright} is settled. Then it is settled due to containing an even dominating cycle as a suffix, due to the invariant and due to π^*_{\curvearrowright} being consistent with the winning strategy $\sigma_{\curvearrowright}$ for Player 0. We define $h(\pi^*)$ by removing the settling dominating cycle similarly to the proof of Lemma 4.5. Using the same argument as in the proof of that lemma, we obtain that $h(\pi^*)$ satisfies the invariant.

It remains to show that σ' is winning for Player 0. To this end, consider a play ρ consistent with σ' and let π_j be the prefix of length j of ρ . As all π_j are consistent with σ' , due to the invariant, neither the overflow counter of the $h(\pi_j)$, nor that of the π_j reaches n . Hence, the colors of the last vertices of π_j and $h(\pi_j)$ coincide. Recall the argument in the proof of Lemma 4.5: If the largest color c appearing infinitely often in ρ is odd, then it can only occur finitely often on even dominating cycles. Hence, after some prefix, every time a vertex of color c is visited, this vertex is added to the simulated play prefix and never removed. This unbounded growth contradicts the $h(\pi_j)$ being unsettled, as every play prefix of $\mathcal{G}'_{\curvearrowright}$ of length $\ell_{\curvearrowright} + 1$ is settled due to Lemma 7.3. Thus, ρ satisfies the parity condition, i.e., σ' is indeed winning for Player 0.

For the other direction, we show that if Player 1 wins $\mathcal{G}'_{\curvearrowright}$, he wins \mathcal{G}' , which suffices due to determinacy. To this end, let τ_{\curvearrowright} be a winning strategy for Player 1 in $\mathcal{G}'_{\curvearrowright}$. We construct a winning strategy τ' for him in \mathcal{G}' by simulating play prefixes in \mathcal{G}' by such prefixes in $\mathcal{G}'_{\curvearrowright}$, from which we remove shortcut- and dominating cycles. We again define a simulation function h that maintains the same invariant as in the proof of Lemma 4.5, i.e.:

Let π be consistent with τ' and end in (v, o, r) with $o < n$. Then, $h(\pi)$ is consistent with τ_{\frown} , is unsettled, and ends in (v, o', r') with $(o', r') \sqsubseteq (o, r)$.

We define h and τ' inductively and simultaneously, starting with $h(v'_I) = v'_I$, which clearly satisfies the invariant. Now let π be a play prefix of \mathcal{G}' consistent with τ' and ending in (v, o, r) . If $(v, o, r) \in V'_0$, then let (v^*, o^*, r^*) be an arbitrary successor of (v, o, r) in \mathcal{A}' . Otherwise, if $(v, o, r) \in V'_1$, let $\tau_{\frown}(h(\pi)) = (v^*, o^*, r^*)$ and define $\tau'(\pi) = (v^*, o^*, r^*)$, where $(o^*, r^*) = \text{Upd}((o, r), (v, v^*))$. In either case, let $\pi^* = \pi \cdot (v^*, o^*, r^*)$.

It remains to define $h(\pi^*)$ in the case that $o^* < n$. To this end, let (o^*, r^*) be the unique memory state such that $\pi^*_{\frown} = h(\pi) \cdot (v^*, o^*, r^*)$ is a play prefix of \mathcal{G}'_{\frown} . If π^*_{\frown} is unsettled and if (v^*, o^*, r^*) is not the destination of a shortcut, we define $h(\pi^*) = \pi^*_{\frown}$, which satisfies the invariant due to Lemma 4.3. If π^*_{\frown} is settled due to $o^* = n$, then, due to the invariant and Lemma 4.3, we obtain $o^* = n$, i.e., we can define $h(\pi^*)$ arbitrarily. If π^*_{\frown} is settled due to reaching a dominating cycle, we remove this cycle from π^*_{\frown} similarly to the construction from the proof of Lemma 4.5.

Finally, consider the case that π^*_{\frown} is unsettled and (v^*, o^*, r^*) is the destination of a shortcut, with $(v^*, o^*_{\rightarrow}, r^*_{\rightarrow})$ as the destination of the corresponding detour. We differentiate whether the destination $(v^*, o^*_{\rightarrow}, r^*_{\rightarrow})$ of the shortcut merely allows Player 1 to catch up to the play prefix constructed in \mathcal{G}' , or whether it is more advantageous for him than the position (v^*, o^*, r^*) actually reached in \mathcal{G}' . In the former case, i.e., if $(o^*, r^*) \sqsupseteq (o^*_{\rightarrow}, r^*_{\rightarrow})$, we define $h(\pi^*) = \pi^*_{\frown}$, which satisfies the invariant by assumption. In the latter case, however, i.e., if $(o^*, r^*) \sqsupseteq (o^*_{\rightarrow}, r^*_{\rightarrow})$ does not hold true, we remove the shortcut cycle similarly to the removal of a settling dominating cycle, obtaining π^*_{\frown} , and define $h(\pi^*) = \pi^*_{\frown}$. This satisfies the invariant due to $(o^*, r^*) \sqsupseteq (o^*_{\rightarrow}, r^*_{\rightarrow})$, which we obtain via Lemma 4.3, and the definition of the shortcut condition.

It remains to show that τ' is indeed winning for Player 1 in \mathcal{G}' . To this end, consider a play ρ consistent with τ' and let $\pi_j = (v_0, o_0, r_0) \cdots (v_j, o_j, r_j)$ be the prefix of length $j + 1$ of ρ . If the overflow counter along ρ eventually saturates, ρ is clearly winning for Player 1. Hence, assume the opposite, and note that, due to the invariant of h , the colors of the last vertices of π_j and $h(\pi_j)$ coincide for all $j \in \mathbb{N}$. Let c be the largest color occurring infinitely often along ρ and assume towards a contradiction that c is even. Similarly to the previous argument, c must either occur on odd dominating cycles or on removed shortcuts after some finite prefix, as these are the only play infixes that are removed from the simulation. Similarly to the proof of Lemma 4.5, c can only occur finitely often on odd dominating cycles, as each such occurrence implies one occurrence of some larger, odd color.

Now assume that c occurs infinitely often on removed shortcut cycles. Since the overflow counter along ρ never saturates, none of the $h(\pi_j)$ contains a saturated overflow counter either. Moreover, as both the removal of an odd dominating cycle and that of a shortcut retain the value of the overflow counter, the values of the overflow counter of the $h(\pi_j)$ eventually stabilize. Let $h(\pi_j) = (v_0^j, o_0^j, r_0^j) \cdots (v_{k_j}^j, o_{k_j}^j, r_{k_j}^j)$. Pick p such that $o_p = o_j$ and $o_{k_p}^p = o_{k_j}^j$ for all $j > p$, and such that c is the largest color occurring on the suffix of ρ starting at position p .

If $o_{k_p}^p < o_p$, then $h(\pi_j)$ results from $h(\pi_{j-1})$ by removing a shortcut cycle only finitely often. In fact, after reaching π_p , no shortcut cycle is removed anymore: If a shortcut is used in the move from $h(\pi_{j-1})$ to $h(\pi_j)$, then $(o_j, r_j) \sqsupseteq (o_{k_j}^j, r_{k_j}^j)$, i.e., the shortcut cycle is not removed. Hence, only finitely many shortcut cycles are removed, which contradicts c occurring on infinitely many of these. Thus, we obtain $o_{k_p}^p = o_p$, which implies $r_{k_p}^p \sqsubseteq r_p$

due to the invariant of h . In particular, for each relevant request that is open in $r_{k_j}^j$, some larger one is open in r_j for each $j > p$.

If c occurs on a removed shortcut cycle, then c must be smaller than the smallest relevant request that is open during the witnessing infix: Otherwise it would answer that relevant request, due to c being even. Hence the detour corresponding to the infix would violate the shortcut condition. While there may be some open requests for colors $c' < c$ in the corresponding infix in ρ , visiting c does not answer all relevant requests in that corresponding infix in ρ , as argued before. This implies traversing the shortcut cycle increases the cost of some request in ρ . Furthermore, since c is the maximal color visited in the considered suffix, this request eventually overflows after traversing at most $b+1$ many increment-edges. This contradicts the choice of p such that no overflows occur after π_p . If less than $b+1$ increment-edges occur during the remainder of the play, then also at most $b+1$ shortcuts occur, since each shortcut requires the traversal of at least one increment-edge. This in turn contradicts c occurring on infinitely many removed shortcuts.

Hence, since vertices of color c only occur finitely often on odd dominating cycles and on removed shortcut cycles, after some finite prefix, each visited vertex of color c is added to the simulated play and never removed. Thus, the $h(\pi_j)$ grow increasingly longer. Such unbounded growth contradicts them being unsettled, as required by the invariant. This is due to every play prefix of length at least l_{\sim} being settled, due to Lemma 7.3. Hence, c is odd, i.e., ρ is winning for Player 1. \square

Having shown \mathcal{G}'_{\sim} to be equivalent to \mathcal{G}' , which is in turn equivalent to \mathcal{G} , we can use the same construction as in the proof of Lemma 4.6, i.e., simulate \mathcal{G}'_{\sim} on an alternating Turing machine, in order to decide the winner of \mathcal{G}'_{\sim} . Due to Lemma 7.3, and due to $\log(nW)$ being polynomial in the size of the description of \mathcal{G} , this Turing machine is polynomially time-bounded. Thus, a similar proof to that of Lemma 4.6 yields PSPACE-membership of the given problem. Together with the previously stated PSPACE-hardness of the given problem, this concludes the proof of Theorem 7.1

Due to the same reasoning as for the results of Section 5, we obtain asymptotically tight exponential bounds for the memory required by both players in order to win with respect to a given bound b . The upper bounds are obtained as a corollary of the equivalence of \mathcal{G} and \mathcal{G}' .

Corollary 7.5. *Let \mathcal{G} be a parity game with integer-valued costs containing n vertices and d odd colors.*

- *If Player 0 has a strategy σ for \mathcal{G} with $\text{Cst}(\sigma) = b$, then she also has a strategy σ' with $\text{Cst}(\sigma') \leq b$ and $|\sigma'| = (b+2)^d$.*
- *If Player 1 has a strategy τ for \mathcal{G} with $\text{Cst}(\tau) = b$, then he also has a strategy τ' with $\text{Cst}(\tau') \geq b$ and $|\tau'| = n(b+2)^d$.*

We obtain matching lower bounds from Theorem 5.2 and Theorem 5.3.

7.2. Tradeoffs Between Memory and Cost in Parity Games with Integer-Valued Costs. As every finitary parity game is a parity game with integer-valued costs, Theorem 6.1 holds true for the latter kind of games as well. In this result, the tradeoff in terms of costs ranges over an interval of size d , while the size of strategies ranges from positional ones to strategies of exponential size. Hence, even a small improvement in the cost realized by a strategy comes at the price of an exponential increase in memory. In the case

of a binary encoding of integer-valued costs, however, an exponential increase in the size of memory used by a strategy may yield an exponential improvement in terms of the cost realized by the strategy.

Theorem 7.6. *For each $d \geq 1$, there exists a parity game with integer-valued costs \mathcal{G}_d with $|\mathcal{G}_d| \in \mathcal{O}(d^2)$, such that for every j with $1 \leq j \leq d$ there exists a strategy σ_j for Player 0 in \mathcal{G}_d such that*

- $(d+1)2^d + 2^{d-1} = \text{Cst}(\sigma_1) > \text{Cst}(\sigma_2) > \dots > \text{Cst}(\sigma_d) = (d+1)2^d$, and
- $1 = |\sigma_1| < |\sigma_2| < \dots < |\sigma_d| = 2^{d-1}$.

Also, for every strategy σ' for Player 0 in \mathcal{G}_d with $\text{Cst}(\sigma') \leq \text{Cst}(\sigma_j)$ we have $|\sigma'| \geq |\sigma_j|$.

Proof. We reuse the arena from the game \mathcal{G}_d constructed in the proof of Theorem 5.2 and redefine the cost function Cst . We do so by assigning a cost of 2^{c-1} to each edge leading to some vertex colored with either the odd color $2c-1$ or the even color $2c$. Moreover, we assign a cost of $2^d - 2^{c-1}$ to each edge leading away from some such vertex. All other edges are edges of cost 0. The traversal of such a gadget with costs incurs a uniform cost of 2^d , regardless of the path taken through it. As the cost of each edge can be encoded using d bits, \mathcal{G}_d is indeed of size $\mathcal{O}(d^2)$.

For any j with $1 \leq j \leq d$ let σ_j be a strategy as defined in Theorem 6.1. Due to similar arguments as in the proof of that theorem, we obtain $\text{Cst}(\sigma_j) = (d+1)2^d + 2^{d-1} - 2^{j-1}$, which satisfies the first property stated in the given theorem. As the strategies remain unchanged from the proof of Theorem 6.1, they also satisfy $|\sigma_j| < |\sigma_{j+1}|$. Finally, similar reasoning to the proof of Theorem 6.1 yields that each σ_j is minimal for its cost. \square

8. STREETT GAMES WITH COSTS

In this section, we consider the Streett condition with costs [FZ14b], which generalizes both the parity condition with costs as well as the classical Streett condition. We show that, given some Streett game with costs \mathcal{G} and bound b , the problem of deciding whether there exists a strategy for Player 0 in \mathcal{G} with cost at most b , is EXPTIME-complete. Thus, this problem is harder to solve than that of solving classical Streett games (unless $\text{CO-NP} = \text{EXPTIME}$), and as hard as solving both finitary Streett games and Streett games with costs [FZ14b]. As a corollary of this result and of those previously obtained in this work we furthermore obtain tight exponential bounds on the memory required by both players in such a game. We begin by formally defining the Streett condition with costs, before examining its complexity and the memory required by both players.

Let $\mathcal{A} = (V, V_0, V_1, E, v_I)$ be an arena and let $\Gamma = (Q_c, P_c)_{c \in [d]}$ for some $d \geq 1$ be a non-empty, finite family of so-called Streett pairs of subsets of V . Intuitively, for each $c \in [d]$, the set Q_c denotes vertices requesting condition c , which are answered by visiting some vertex in P_c . Finally, let $\text{Cst} = (\text{Cst}_c)_{c \in [d]}$ be a family of cost functions, where $\text{Cst}_c: E \rightarrow \mathbb{N}$ for each $c \in [d]$, which we extend to cost functions over plays as usual. We denote the highest cost assigned by any $\text{Cst}_c \in \text{Cst}$ by W .

Let $\rho = v_0 v_1 v_2 \dots$ be a play in \mathcal{A} and let $j \in \mathbb{N}$ be a position. We first define the cost-of-response for a single Streett pair (Q_c, P_c) as

$$\text{StCor}_c(\rho, j) = \begin{cases} \min\{\text{Cst}(v_j \dots v_{j'}) \mid j' \geq j \text{ and } v_{j'} \in P_c\} & \text{if } v_j \in Q_c \\ 0 & \text{otherwise} \end{cases}$$

with $\min \emptyset = \infty$. Note that, in contrast to parity games, the visit to v_j may open multiple requests, as there may exist multiple c such that $v_j \in Q_c$. Thus, we define the (total) cost-of-response at position $j \in \mathbb{N}$ of ρ by

$$\text{StCor}(\rho, j) = \max\{\text{StCor}_c(\rho, j) \mid c \in [d]\} .$$

Thus, $\text{StCor}(\rho, j)$ is the cost of the infix of ρ from position j to the earliest position where all requests opened at position j are answered, and ∞ , if at least one such request is not answered. Moreover, $\text{StCor}(\rho, j)$ is zero if no requests are opened at position j .

The Streett condition with costs is then defined as

$$\text{CostStreett}(\Gamma, \text{Cst}) = \{\rho \in V^\omega \mid \limsup_{j \rightarrow \infty} \text{StCor}(\rho, j) < \infty\} ,$$

i.e., ρ satisfies the condition if there exists a bound $b \in \mathbb{N}$ such that all but finitely many requests are answered with cost less than b . In particular, only finitely many requests may be unanswered, even if they only incur finite cost. Similarly to the case of the parity condition with costs, the bound b may depend on the play ρ .

A game $\mathcal{G} = (\mathcal{A}, \text{CostStreett}(\Gamma, \text{Cst}))$ is called a Streett game with costs. If all Cst_c assign 0 to every edge, then $\text{CostStreett}(\Gamma, \text{Cst})$ is a classical Streett condition [Str81], denoted by $\text{Streett}(\Gamma)$. Dually, if all Cst_c assign 1 to every edge, then $\text{CostStreett}(\Omega, \text{Cst})$ is equal to the finitary Streett condition over Γ , as introduced by Chatterjee et al. [CHH09] and denoted by $\text{FinStreett}(\Gamma)$. In these cases, we refer to \mathcal{G} as a Streett or a finitary Streett game, respectively.

We assume the cost functions to be given in binary encoding.⁶ Hence, in general, the largest cost W is exponential in $2^{|\text{Cst}|}$, where $|\text{Cst}|$ is the length of the encoding of Cst . Thus, we define $|\mathcal{G}| = |\mathcal{A}| + d + \log W$.

Theorem 8.1.

- (1) *Solving Streett games is co-NP-complete. If Player 0 wins, then she has a winning strategy of size $d!$, while Player 1 has uniform positional winning strategies [Hor05].*
- (2) *Solving finitary Streett games is EXPTIME-complete.⁷ If Player 0 wins, then she has a winning strategy of size $d2^d$, but Player 1 has in general no finite-state winning strategy [CHH09].*
- (3) *Solving Streett games with costs is EXPTIME-complete. If Player 0 wins, then she has a winning strategy of size $2^d((2d)!)$, but Player 1 has in general no finite-state winning strategy [FZ14b].*


We define the cost of strategies for Streett games with costs analogously to the parity case. In contrast to that case, however, we obtain an exponential upper bound on the cost of an optimal strategy for a Streett game with costs by using Theorem 8.1(3) and applying the same pumping argument as for Corollary 3.2.

Corollary 8.2. *Let \mathcal{G} be a Streett game with costs with n vertices and d Streett pairs. Moreover, let W be the largest cost in \mathcal{G} . If Player 0 wins \mathcal{G} , then she has a strategy σ with $\text{Cst}(\sigma) \leq nW \cdot 2^d((2d)!)$.*

Similarly to the case of parity games with costs, this bound is tight. The games demonstrating the lower bound are, however, no longer trivial. We adapt the games used

⁶All lower bounds shown for this setting already hold for that of finitary Streett games.

⁷Shown in unpublished work by Chatterjee, Henzinger, and Horn, obtained by a minor modification to the proof of EXPTIME-hardness of solving request-response games [CHH11].

Figure 9: The game \mathcal{G}_3 .

in [CHH11] for demonstrating the necessity for exponential memory for Player 0 in request-response games in order to show this bound.

Theorem 8.3. *For each $d \geq 0$, there exists a finitary Streett game \mathcal{G}_d with $\mathcal{O}(d)$ vertices and $d + 1$ Streett pairs, such that Player 0 has a strategy with cost $b = 3(2^d - 1) + 2$, but no strategy with cost less than b .*

Proof. Figure 9 shows the game \mathcal{G}_3 . In general, the game \mathcal{G}_d consists of an initializing prefix, i.e., the vertices P and Q , a central vertex m , and one branch for each of the $d + 1$ Streett pairs, which can be entered from m . Each branch consists of one path leading back to the central vertex as well as one path restarting the game via moving to the initializing prefix.

The vertices P and Q answer and open all $d + 1$ requests, respectively, while a visit to m neither opens nor answers any requests. When moving into the branch associated with condition c , the initial visit to P_c answers the request for condition c . Afterwards, visiting c or \bar{c} opens requests for all conditions $c' < c$ and answers all requests for conditions $c' > c$, respectively.

Intuitively, by moving to P_c , Player 0 claims that c is the smallest index for which a request is open and answers the request for that condition. If this claim holds true, the best choice for Player 1 is to move to c , where requests for all conditions $c' < c$ are opened and the play returns to m . If, on the other hand, this claim does not hold true, then Player 1 can move to \bar{c} , where all requests for conditions $c' > c$ are answered and Player 1 can increase the cost of the remaining requests arbitrarily before moving to P and thereby starting the next round. Note that staying in \bar{c} infinitely long is losing for Player 1, as no requests are opened in that vertex and thus, the cost of the resulting play is 0.

The optimal strategy for Player 0 implements a binary counter: Whenever the play reaches m , she has to recall the smallest c for which there is an open request and move to P_c . Since moving to P_c and returning to m via c reopens requests for all conditions $c' < c$, Player 0 has to repeat answering these requests from smallest to largest before she can answer the outstanding request for condition $c + 1$. Figure 10 shows the cost of requests during a play consistent with this strategy in \mathcal{G}_3 .

Note that, in this example, after answering the request for 3, Player 0 continues implementing a binary counter until no more requests are open. At this point, Player 0 may choose an arbitrary branch. If Player 1 then moves to the trap-vertex and subsequently

Branch	0	1	0	2	0	1	0	3	0	1	0	2	0	1	0	...	
$r(0)$	1	⊥	1	⊥	1	⊥	1	⊥	1	⊥	1	⊥	1	⊥	1	⊥	...
$r(1)$	1	4	⊥	⊥	1	4	⊥	⊥	1	4	⊥	⊥	1	4	⊥	⊥	...
$r(2)$	1	4	7	10	⊥	⊥	⊥	⊥	1	4	7	10	⊥	⊥	⊥	⊥	...
$r(3)$	1	4	7	10	13	16	19	22	⊥	⊥	⊥	⊥	⊥	⊥	⊥	⊥	...

Figure 10: An optimal play for Player 0 in \mathcal{G}_3 . We only consider the positions at which the play is at vertex m . The topmost row denotes the branch visited in-between two visits to that vertex.

to P , he starts the next round of the game. If he, however, returns to m , Player 0 obtains new open requests which she has to answer as she did before. Since no request for 3 can be opened without moving to P , not moving on to the next round will not yield higher costs than doing so.

As illustrated in Figure 10, answering the request for condition d posed at the beginning of each round requires $2^d - 1$ many visits to branches, as well as an additional step into the first vertex of the branch of condition d . As each of the visits to the branches implies the traversal of three edges, this request is answered b steps after it is posed. Hence, the strategy implementing a binary counter has a cost of $b = 3(2^d - 1) + 2$.

Moreover, as argued above, in each round in which Player 0 deviates from this strategy, Player 1 can move the play into the “trap”-vertex \bar{c} of the current branch upon the first deviation, where he can loop until the cost of an open request increases beyond b , before he moves to the next round. Hence, Player 0 has to adhere to this strategy after finitely many rounds, i.e., each strategy for her has cost at least b . \square

8.1. The Complexity of Solving Streett Games with Costs Optimally. We now show that solving Streett games with costs with respect to a given bound b is EXPTIME-complete. Since solving finitary Streett games is complete for the same complexity class, and since an exponential b suffices for Player 0 to win in such games due to Corollary 8.2, we can encode the problem of solving a finitary Streett game as the given problem with only a linear blowup, due to binary encoding of b . Hence, the latter problem is EXPTIME-hard.

In order to show membership of the given problem in EXPTIME, we reduce it to that of solving a classical Streett game with exponentially many vertices, but with only a single additional Streett pair. Since Streett games with n vertices and d Streett pairs can be solved in time $\mathcal{O}(nd(d!))$ using the algorithm from Piterman and Pnueli [PP06], this construction yields EXPTIME-membership of the given problem.

Theorem 8.4. *The following problem is EXPTIME-complete: “Given a Streett game with costs \mathcal{G} and a bound $b \in \mathbb{N}$ in binary encoding, does Player 0 have a strategy σ for \mathcal{G} with $\text{Cst}(\sigma) \leq b$?”*

Proof. EXPTIME-hardness of this problem follows from the EXPTIME-hardness of deciding the winner in a finitary Streett game [FZ14b] as argued above.

We now show the given problem to be in EXPTIME. The idea for this proof is the same as for the proof of Lemma 4.6, i.e., we reduce the problem to solving a classical Streett game \mathcal{G}' . Instead of simulating \mathcal{G}' on the fly, however, we construct and solve it explicitly.

Let $\mathcal{G} = (\mathcal{A}, \text{CostStreett}(\Gamma, \text{Cst}))$ be a Streett game with costs with n vertices and d Streett pairs. Moreover, let W be the largest cost assigned by any Cst_c in Cst . If $b \geq nW \cdot 2^d((2d)!)$, then we construct $\mathcal{G}' = (\mathcal{A}, \text{CostStreett}(\Gamma, \text{Cst}'))$ with $\text{Cst}'(e) = \varepsilon$, if $\text{Cst}(e) = 0$ and $\text{Cst}'(e) = \mathbf{i}$ otherwise and solve \mathcal{G}' using an exponential-time algorithm [FZ14b]. If Player 0 can ensure some upper bound on the cost incurred in \mathcal{G}' , then she can also do so in \mathcal{G} using the same strategy. Thus, by Corollary 8.2, she can bound the cost from above by b . Similarly, if Player 1 wins \mathcal{G}' , then he can still do so in \mathcal{G} using the same strategy. Hence, solving \mathcal{G}' solves \mathcal{G} with respect to b .

Thus, assume $b < nW \cdot 2^d((2d)!)$. For the reduction of \mathcal{G} to a Streett game we again use a memory structure that keeps track of the costs of responses accumulated so far, up to the bound b , while allowing this counter to overflow at most n times. Hence, let $\mathcal{M} = (M, m_I, \text{Upd})$ be the memory structure with memory states $M = [n+1] \times (\{\perp\} \cup [b+1])^{[d]}$, where we define m_I and Upd analogously to the parity case. Note that \mathcal{M} is of size $|\mathcal{M}| = (n+1) \cdot (b+2)^d$.

Also, given $\Gamma = (Q_c, P_c)_{c \in [d]}$, we construct $\Gamma' = (Q'_c, P'_c)_{c \in [d+1]}$ as $(Q'_c, P'_c) = (Q_c \times M, P_c \times M)$ for $c \in [d]$, and $(Q'_d, P'_d) = (V \times (\{\perp\} \cup [b+1])^{[d]} \times \{n\}, \emptyset)$. Thus, Γ' is the conjunction of the extension of Γ to the game $\mathcal{G} \times \mathcal{M}$ and one additional Streett pair which causes Player 0 to lose once the overflow counter reaches the value n . Hence $|\Gamma'| = |\Gamma| + 1$.

We define $\mathcal{G}' = (\mathcal{A} \times \mathcal{M}, \text{Streett}(\Gamma'))$. Player 0 wins \mathcal{G}' if and only if she has a strategy with cost at most b in \mathcal{G} , due to the same arguments as in the proof of Lemma 4.2.

Moreover, using the algorithm presented in [PP06], we can solve \mathcal{G}' in exponential time in $|\mathcal{G}'|$. The Streett game \mathcal{G}' has $n' \in \mathcal{O}(n^2 b^d)$ vertices. Since we assume $b \leq nW \cdot 2^d((2d)!)$, we obtain $n' \in \mathcal{O}((nW \cdot 2^d((2d)!))^d)$. Moreover, \mathcal{G}' has $d' = d + 1$ many Streett pairs.

As discussed above, by using the algorithm for solving Streett games by Piterman and Pnueli [PP06], we obtain an algorithm that is polynomial in the number of vertices n of \mathcal{G} , while it is exponential in the number of Streett pairs d and $\log W$. Hence, the given problem is in EXPTIME. \square

This result also holds true if the bound b is given in unary encoding. As every number in unary encoding can be rewritten in binary encoding in polynomial time, membership in EXPTIME follows directly. Moreover, recall that EXPTIME-hardness of the problem of Theorem 8.4 follows from EXPTIME-hardness of solving finitary Streett games. This problem is in turn shown to be EXPTIME-hard via a reduction from the word problem for polynomially time-bounded alternating Turing machines. A minor modification of that proof yields that a polynomial bound suffices for Player 0 in order to win the resulting finitary parity game. Hence, the problem is still EXPTIME-hard when considering a unary encoding of the bound b .

8.2. Memory Requirements of Optimal Strategies in Streett Games with Costs.

As we have shown in the previous section, Streett games with costs can be solved by reducing them to classical Streett games of exponential size, but with only linearly many Streett pairs. Similarly to Corollary 5.1, we obtain an exponential upper bound on the memory necessary for both players to win a Streett game with costs with respect to a given bound as a corollary of Theorems 8.1(1) and 8.4.

Also, similarly to Corollary 5.1, we are able to remove the overflow counter from our memory structure for Player 0 by letting her play according to the largest value of this counter for which she still has a winning strategy. This yields an improved upper bound

on her required memory. Moreover, we obtain matching lower bounds for both players as a corollary of Theorems 5.2 and 5.3, as every finitary parity game with costs with d colors can be turned into a finitary Streett game of the same size with d Streett pairs.

Corollary 8.5. *Let \mathcal{G} be a Streett game with costs with n vertices and d Streett pairs. Moreover, let $b \in \mathbb{N}$ be some bound.*

- (1) *If Player 0 has a strategy σ in \mathcal{G} with $\text{Cst}(\sigma) \leq b$, then she also has a strategy σ' with $\text{Cst}(\sigma') \leq b$ and $|\sigma'| = (d+1)! \cdot (b+2)^d$.*
- (2) *If Player 1 has a strategy τ in \mathcal{G} with $\text{Cst}(\tau) \geq b$, then he also has a strategy τ' with $\text{Cst}(\tau') \geq b$ and $|\tau'| = n(b+2)^d$.*

These bounds are asymptotically tight already for finitary Streett conditions.

9. CONCLUSION

In this work we have shown that playing parity games with costs optimally is harder than just winning them, both in terms of computational complexity as well as in terms of memory requirements of strategies. We proved checking an upper bound on the value of an optimal strategy to be complete for polynomial space, while just solving such games is in $\text{UP} \cap \text{CO-UP}$, respectively in PTIME for the special case of finitary parity games. Moreover, we have shown that optimal strategies in general require exponential memory, but also that exponential memory is always sufficient to implement optimal strategies. In contrast, winning strategies in these games are positional. Finally, we have shown that, in general, there exists a gradual tradeoff between the size and the cost of strategies.

Also, we considered Streett games with costs: checking an upper bound on the cost of an optimal strategy is EXPTIME -complete and exponential memory is sufficient to implement optimal strategies. Thus, playing optimally is as hard as just winning.

All our proofs can be adapted for the case of bounded parity (Streett) games and bounded parity (Streett) games with costs [CHH09, FZ14b]. While the parity condition with costs only restricts the cost-of-response in the limit, the bounded parity condition prohibits any unanswered request with cost ∞ (but still allows finitely many unanswered requests with finite cost). The other conditions are defined similarly.

In further research, we are considering two additional directions in which to extend the cost function: By allowing negative integers as costs and by allowing multiple cost functions, i.e., by extending the parity winning condition with a family of cost functions similar to the case of Streett conditions. Moreover, we investigate tradeoffs in delay games with quantitative winning conditions. Preliminary results exhibited a tradeoff between costs of strategies and delay [Zim17], but there are no results involving the size of strategies.

REFERENCES

- [AETP01] Rajeev Alur, Kousha Etessami, Salvatore La Torre, and Doron Peled. Parametric temporal logic for “model measuring”. *ACM Trans. Comput. Log.*, 2(3):388–407, 2001.
- [AR14] Benjamin Aminof and Sasha Rubin. First cycle games. In Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi, editors, *SR 2014*, volume 146 of *EPTCS*, pages 83–90, 2014.
- [BFL⁺08] Patricia Bouyer, Ulrich Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jiri Srba. Infinite runs in weighted timed automata with energy constraints. In Franck Cassez and Claude Jard, editors, *FORMATS 2008*, volume 5215 of *LNCS*, pages 33–47. Springer, 2008.

- [BHR16] Véronique Bruyère, Quentin Hautem, and Mickael Randour. Window parity games: an alternative approach toward parity games with time bounds. In Domenico Cantone and Giorgio Delzanno, editors, *GandALF 2016*, volume 226 of *EPTCS*, pages 135–148, 2016.
- [CdAH04] Krishnendu Chatterjee, Luca de Alfaro, and Thomas A. Henzinger. Trading memory for randomness. In Gethin Norman and William Sanders, editors, *QEST 2004*, pages 206–217. IEEE, 2004.
- [CdAHS03] Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource interfaces. In Rajeev Alur and Insup Lee, editors, *EMSOFT 2003*, volume 2855 of *LNCS*, pages 117–133. Springer, 2003.
- [CF13] Krishnendu Chatterjee and Nathanaël Fijalkow. Infinite-state games with finitary conditions. *arXiv*, 1301.2661, 2013.
- [CHH09] Krishnendu Chatterjee, Thomas A. Henzinger, and Florian Horn. Finitary winning in ω -regular games. *ACM Trans. Comput. Log.*, 11(1), 2009.
- [CHH11] Krishnendu Chatterjee, Thomas A. Henzinger, and Florian Horn. The complexity of request-response games. In Adrian-Horia Dediu, Shunsuke Inenaga, and Carlos Martín-Vide, editors, *LATA 2011*, pages 227–237. Springer, 2011.
- [CHP08] Krishnendu Chatterjee, Thomas A. Henzinger, and Vinayak S. Prabhu. Trading infinite memory for uniform randomness in timed games. In Magnus Egerstedt and Bud Mishra, editors, *HSCC 2008*, volume 4981 of *LNCS*, pages 87–100. Springer, 2008.
- [CJK⁺17] Cristian S. Calude, Sanjay Jain, Bakhadyr Khossainov, Wei Li, and Frank Stephan. Deciding parity games in quasipolynomial time. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *STOC 2017*, pages 252–263. ACM, 2017.
- [CKS81] Ashok K. Chandra, Dexter Kozen, and Larry J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981.
- [EJ91] E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *FOCS 1991*, pages 368–377. IEEE, 1991.
- [EM79] Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *Int. J. Game Theory*, 8:109–113, 1979.
- [FHKS15] Nathanaël Fijalkow, Florian Horn, Denis Kuperberg, and Michal Skrzypczak. Trading bounds for memory in games with counters. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *ICALP 2015, Part II*, volume 9135 of *LNCS*, pages 197–208. Springer, 2015.
- [FZ14a] Peter Faymonville and Martin Zimmermann. Parametric linear dynamic logic. In Adriano Peron and Carla Piazza, editors, *GandALF 2014*, volume 161 of *EPTCS*, pages 60–73, 2014.
- [FZ14b] Nathanaël Fijalkow and Martin Zimmermann. Parity and Streett Games with Costs. *LMCS*, 10(2), 2014.
- [Hor05] Florian Horn. Streett games on finite graphs. In *GDV 2005*, 2005. available at <https://www.irif.fr/~horn/publications.html>.
- [HTWZ15] Florian Horn, Wolfgang Thomas, Nico Wallmeier, and Martin Zimmermann. Optimal strategy synthesis for request-response games. *RAIRO - Theor. Inf. and Applic.*, 49(3):179–203, 2015.
- [Jur98] Marcin Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Inf. Process. Lett.*, 68(3):119–124, 1998.
- [KPV09] Orna Kupferman, Nir Piterman, and Moshe Y. Vardi. From liveness to promptness. *Form. Met. in Sys. Des.*, 34(2):83–103, 2009.
- [Lan14] Martin Lang. Resource reachability games on pushdown graphs. In Anca Muscholl, editor, *FOS-SACS 14*, volume 8412 of *LNCS*, pages 195–209. Springer, 2014.
- [MMS15] Fabio Mogavero, Aniello Murano, and Loredana Sorrentino. On promptness in parity games. *Fundam. Inform.*, 139(3):277–305, 2015.
- [Mos91] Andrzej Mostowski. Games with forbidden positions. Technical Report 78, University of Gdańsk, 1991.
- [PP06] Nir Piterman and Amir Pnueli. Faster solutions of Rabin and Streett games. In *LICS 2006*, pages 275 – 284. IEEE, 2006.
- [PR89] Amir Pnueli and Roni Rosner. On the synthesis of an asynchronous reactive module. In Giorgio Ausiello, Mariangiola Dezani-Ciancaglini, and Simona Ronchi Della Rocca, editors, *ICALP 1989*, volume 372 of *LNCS*, pages 652–671. Springer, 1989.

- [Ran14] Mickael Randour. *Synthesis in Multi-Criteria Quantitative Games*. PhD thesis, University of Mons, 2014.
- [Str81] Robert S. Streett. Propositional dynamic logic of looping and converse. In *STOC 1981*, pages 375 – 383. ACM, 1981.
- [WZ16] Alexander Weinert and Martin Zimmermann. Easy to win, hard to master: Optimal strategies in parity games with costs. In Jean-Marc Talbot and Laurent Regnier, editors, *CSL 2016*, volume 62 of *LIPICs*, pages 31:1–31:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [Zer13] Ernst Zermelo. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In *Proc. Fifth Congress of Mathematicians, Vol. 2*, pages 501–504. Cambridge Press, 1913.
- [Zim09] Martin Zimmermann. Time-optimal winning strategies for poset games. In Sebastian Maneth, editor, *CIAA 2009*, volume 5642 of *LNCS*, pages 217–226. Springer, 2009.
- [Zim13] Martin Zimmermann. Optimal bounds in parametric LTL games. *Theoret. Comput. Sci.*, 493(0):30 – 45, 2013.
- [Zim15] Martin Zimmermann. Parameterized linear temporal logics meet costs: Still not costlier than LTL. In Javier Esparza and Enrico Tronci, editors, *GandALF 2015*, volume 193 of *EPTCS*, pages 144–157, 2015.
- [Zim17] Martin Zimmermann. Games with costs and delays. *arXiv*, 1701.02168, 2017.
- [ZP95] Uri Zwick and Mike Paterson. The complexity of mean payoff games. In Ding-Zhu Du and Ming Li, editors, *COCOON 1995*, volume 959 of *LNCS*, pages 1–10. Springer, 1995.