

Games with Costs and Delays

Martin Zimmermann

Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany

Email: zimmermann@react.uni-saarland.de

Abstract—We demonstrate the usefulness of adding delay to infinite games with quantitative winning conditions. In a delay game, one of the players may delay her moves to obtain a lookahead on her opponent’s moves. We show that determining the winner of delay games with winning conditions given by parity automata with costs is EXPTIME-complete and that exponential bounded lookahead is both sufficient and in general necessary. Thus, although the parity condition with costs is a quantitative extension of the parity condition, our results show that adding costs does not increase the complexity of delay games with parity conditions.

Furthermore, we study a new phenomenon that appears in quantitative delay games: lookahead can be traded for the quality of winning strategies and vice versa. We determine the extent of this tradeoff. In particular, even the smallest lookahead allows to improve the quality of an optimal strategy from the worst possible value to almost the smallest possible one. Thus, the benefit of introducing lookahead is twofold: not only does it allow the delaying player to win games she would lose without, but lookahead also allows her to improve the quality of her winning strategies in games she wins even without lookahead.

I. INTRODUCTION

Infinite games are one of the pillars of logics and automata theory with a plethora of applications, e.g., as solutions for the reactive synthesis problem and for the model-checking of fixed-point logics, and as the foundation of the game-based proof of Rabin’s theorem to name a few highlights.

The study of infinite games in automata theory was initiated by the seminal Büchi-Landweber theorem [1], which solved Church’s controller synthesis problem [2]: the winner of an infinite-duration two-player zero-sum perfect-information game with ω -regular winning condition can be determined effectively. Furthermore, a finite-state winning strategy, i.e., a strategy that is finitely described by an automaton with output, can be computed effectively. Ever since, this result has been extended along various axes, e.g., the number of players, the type of winning condition, the type of interaction between the players, the informedness of the players, etc. Here, we consider two extensions: first, we allow one player to delay her moves, which gives her a lookahead on her opponent’s moves. Second, we consider quantitative winning conditions, i.e., the finitary parity condition and the parity condition with costs, which both strengthen the classical parity condition.

Delay Games: The addition of delay to the original setting of Büchi and Landweber was already studied by Hosch and Landweber himself [3] shortly after the publication of the

original result. It concerns the ability of one player to delay her moves in order to gain a lookahead on her opponent’s moves. This allows her to win games she would lose without lookahead, e.g., if her first move depends on the third move of the opponent. Delay games are not only useful to model buffers and transmission of data, but the existence of a continuous uniformization function for a given relation is also characterized by winning a delay game (see [4] for details). Recently, delay games have been reconsidered by Holtmann et al. [4] and the first comprehensive study of their properties has been initiated [5], [6], [7], [8], [9].

In particular, determining the winner of delay games with winning conditions given by deterministic parity automata is EXPTIME-complete, where hardness already holds for safety conditions [7]. These results have to be contrasted with those for the delay-free setting: determining the winner of parity games is in $\text{UP} \cap \text{coUP}$ [10], while the special case of safety games is in PTIME. Additionally, exponential constant lookahead is both sufficient and in general necessary [7], i.e., unbounded lookahead does not offer an additional advantage for the delaying player in ω -regular delay games. Either she wins by delaying an exponential number of moves at the beginning of a play and then never again, or not at all.

Parity Games with Costs: Another important generalization of the Büchi-Landweber theorem concerns quantitative winning conditions, as qualitative ones are often too weak to capture the specifications of a system to be synthesized. Prominent examples of quantitative specification languages are parameterized extensions of Linear Temporal Logic (LTL) like Prompt-LTL [11], which adds the prompt-eventually operator whose scope is bounded in time, and the finitary parity condition [12]. The latter strengthens the classical parity condition by requiring a fixed, but arbitrary, bound b on the distance between occurrences of odd colors and the next larger even color. Recently, a further generalization has been introduced in the setting of arenas with costs, i.e., non-negative edge weights. Here, the bound b on the distance is replaced by a bound on the cost incurred between occurrences of odd colors and the next larger even color [13]. This condition, the parity condition with costs, subsumes both the classical parity condition and the finitary parity condition.

Although the finitary parity condition is a strengthening of the classical parity condition, the winner of a finitary parity game can be determined in polynomial time [12], i.e., the solution problem is simpler than that of parity games (unless $\text{UP} \cap \text{coUP} = \text{PTIME}$). In contrast, parity games with costs are not harder than they have to be, i.e., as hard as parity

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

games [14]. Furthermore, such games induce an optimization problem: determine the smallest bound b that allows to satisfy the condition with respect to b . It turns out that determining the optimal bound is much harder [15]: determining whether a given parity game with costs can be won with respect to a given bound b is **PSPACE**-complete, where hardness already holds for the special case of finitary parity games. The same phenomenon manifests itself in the memory requirements of winning strategies for finitary parity games: playing optimally requires exponentially more memory than just winning.

Quantitative Delay Games: Recently, both extensions have been investigated simultaneously, first in the form of delay games with **WMSO+U** winning conditions, i.e., weak monadic second-order logic with the unbounding quantifier [16]. This logic turned out to be too strong: in general, unbounded lookahead is necessary to win such games and the problem of determining the winner has only partially been resolved [9]. Thus, the search for tractable classes of quantitative winning conditions started. A first encouraging candidate was found in Prompt-LTL: determining the winner of Prompt-LTL delay games is **3EXPTIME**-complete and triply-exponential constant lookahead is sufficient and in general necessary [8]. This comparatively high complexity has to be contrasted with that of delay-free Prompt-LTL games, which are already **2EXPTIME**-complete [17]. Thus, adding delay incurs an exponential blowup, which is in line with other results for delay games mentioned above.

Our Contribution: In this work, we investigate delay games with finitary parity conditions and parity conditions with costs and demonstrate the positive effects of adding lookahead to quantitative games. As both the finitary parity and the parity condition with costs subsume the classical safety condition, we immediately obtain **EXPTIME**-hardness and an exponential lower bound on the necessary lookahead.

First, we show that the exponential lower bound on the lookahead is tight, by presenting a matching upper bound via a pumping argument for the opponent of the delaying player: if he wins with a large enough lookahead, then he can pump his moves to win for an arbitrarily large lookahead.

Second, using this result, we construct a *delay-free* parity game with costs that is equivalent to the delay game with fixed exponential lookahead. Furthermore, we show that the resulting game can be solved in exponential time, i.e., we prove **EXPTIME**-completeness of determining the winner of delay games with finitary parity conditions and parity conditions with costs.

These two results show that adding costs to delay games with parity conditions comes for free, i.e., the complexity does not increase, both in terms of necessary lookahead and the computational complexity of determining the winner. These results are similar to the ones for Prompt-LTL delay games, which are just as hard as LTL delay games. However, unlike the latter games, which are **3EXPTIME**-complete, delay games with quantitative extensions of parity conditions are only **EXPTIME**-complete.

Third, we investigate the power of delay in quantitative

games: having a lookahead on the moves of her opponent allows a player not only to win games she would lose without this advantage, but also to improve the (semantic) quality of her winning strategies in games she wins even without lookahead. For example, we present a game induced by an automaton with $\mathcal{O}(n)$ states such that she wins the delay-free game with bound n on every play, which is close to the worst-case [13]. However, with a lookahead of just a single move, she wins the game with bound 1. Thus, lookahead can be traded for quality and vice versa. In further examples, we show that this tradeoff can be gradual, i.e., decrementing the bound requires one additional move lookahead, and that exponential lookahead may be necessary to obtain the optimal bound. Furthermore, we present matching upper bounds on the tradeoff between lookahead and quality.

All examples we present in this work are from a very small fragment, finitary Büchi games, i.e., every edge has cost 1 and the only colors are 1 and 2. Thus, every occurrence of a 1 has to be answered quickly by an occurrence of a 2. Our results show that even these games are as hard as general parity games with costs, i.e., games with arbitrary weights and an arbitrary number of colors.

Fourth, we consider the more general setting of delay games with winning conditions given by Streett automata with costs. Streett conditions are obtained from parity conditions by abandoning the hierarchy between requests (even numbers) and responses (odd numbers) by allowing arbitrary sets of states to represent requests and responses. We prove that the techniques developed for parity conditions can be extended to Streett conditions with costs, at the cost of an exponential blowup: doubly-exponential constant lookahead is sufficient and such games can be solved in doubly-exponential time. Also, we show that an *optimal* winning strategy requires in general doubly-exponential lookahead. Whether general strategies do as well is an open problem. Note that this is already the case for delay games with qualitative Streett conditions.

All proofs omitted due to space restrictions can be found in the full version [18].

II. DEFINITIONS

We denote the non-negative integers by \mathbb{N} . Given two infinite words $\alpha \in \Sigma_I^\omega$ and $\beta \in \Sigma_O^\omega$ we write $\binom{\alpha}{\beta}$ for the word $\binom{\alpha(0)}{\beta(0)} \binom{\alpha(1)}{\beta(1)} \binom{\alpha(2)}{\beta(2)} \dots \in (\Sigma_I \times \Sigma_O)^\omega$. Similarly, we write $\binom{x}{y}$ for finite words x and y , provided they are of equal length.

A. Parity Automata with Costs

A parity automaton with costs $\mathcal{A} = (Q, \Sigma, q_I, \delta, \Omega, \text{Cst})$ consists of a finite set Q of states, an alphabet Σ , an initial state $q_I \in Q$, a (deterministic and complete) transition function $\delta: Q \times \Sigma \rightarrow Q$, a coloring $\Omega: Q \rightarrow \mathbb{N}$, and a cost function $\text{Cst}: \delta \rightarrow \{\varepsilon, \mathbf{i}\}$ (here, and whenever convenient, we treat δ as a relation). If $\text{Cst}(q, a, q') = \varepsilon$, then we speak of an ε -transition, otherwise of an increment-transition. The size of \mathcal{A} is defined as $|\mathcal{A}| = |Q|$.




Fig. 1. The automaton \mathcal{A}_n (on the left), which contains gadgets G_1, \dots, G_n (on the right). Transitions not depicted are defined as follows: those of the form $(\frac{\#}{b})$ for arbitray $b \in \Sigma_O$ lead to an (accepting) sink of color 2, those of the form $(\frac{\neq \#}{\#})$ lead to a (rejecting) sink of color 1. Both sinks are not drawn.

The run of \mathcal{A} starting in $q_0 \in Q$ on a finite word $w = a_0 \dots a_{n-1}$ is the unique sequence

$$(q_0, a_0, q_1) (q_1, a_1, q_2) \dots (q_{n-1}, a_{n-1}, q_n) \in \delta^*, \quad (1)$$

i.e., $q_{j+1} = \delta(q_j, a_j)$ for every $j < n$. We say the run ends in q_n and define its cost as the number of traversed increment-transitions. The run of \mathcal{A} starting in q_0 on an infinite word and its cost is defined analogously. If we speak of *the* run of \mathcal{A} on an infinite word, then we mean the one starting in q_I .

As usual in parity games, we interpret the occurrence of an odd color c as a request, which has to be answered by an even $c' > c$. Hence, we define $\text{Ans}(c) = \{c' \in \Omega(Q) \mid c' > c \text{ and } c' \text{ is even}\}$. An infinite run $\rho = (q_0, a_0, q_1) (q_1, a_1, q_2) (q_2, a_2, q_3) \dots$ is accepting, if it satisfies the parity condition with costs: there is a bound $b \in \mathbb{N}$ such that for almost all n with odd $\Omega(q_n)$, the cost between the positions n and n' is at most b , where n' is minimal such that $\Omega(q_{n'})$ answers the request $\Omega(q_n)$. Formally, we require the cost $\limsup_{n \rightarrow \infty} \text{Cor}(\rho, n)$ of ρ to be finite, where the cost-of-response $\text{Cor}(\rho, n)$ is defined to be 0 if $\Omega(q_n)$ is even and defined to be

$$\min\{\text{Cst}((q_n, a_n, q_{n+1}) \dots (q_{n'-1}, a_{n'-1}, q_{n'})) \mid n' > n \text{ and } \Omega(q_{n'}) \in \text{Ans}(\Omega(q_n))\}$$

if $\Omega(q_n)$ is odd, where $\min \emptyset = \infty$.

The language of \mathcal{A} , denoted by $L(\mathcal{A})$, contains all infinite words whose run of \mathcal{A} is accepting. If every transition in \mathcal{A} is an ε -transition, then the parity condition with costs boils down to the parity condition, i.e., \mathcal{A} is a classical parity automaton and $L(\mathcal{A})$ is ω -regular. In contrast, if every transition is an increment-transition, then \mathcal{A} is a finitary parity automaton, which have been studied by Chatterjee and Fijalkow [19].

Remark 1. In the following, and without mentioning it again, we only consider parity automata with both even and odd colors. If this is not the case, then the automaton is trivial, i.e., its language is either universal or empty.

In our examples, we often use finitary Büchi automata, i.e., automata having only increment-transitions and colors 1 and 2. Hence, every state of color 2 answers all pending requests

and all other states open a request (of color 1). Hence, a run of such an automaton is accepting, if there is a bound b such that every infix of length b contains a transition whose target state has color 2. Furthermore, the cost of the run is the smallest b' such that almost every infix of length b' contains a transition whose target state has color 2. In figures, we denote states of color 2 by doubly-lined states, all other states have color 1, and we do not depict the cost function, as it is trivial.

Example 1. Consider the finitary Büchi automaton \mathcal{A}_n , for $n > 1$, on the left side of Figure 1 over the alphabet $\Sigma_I \times \Sigma_O = \{1, \dots, n, \#\}^2$. A word $(\alpha^{(0)} \dots \alpha^{(m)})_{\beta^{(0)} \dots \beta^{(m)}}$ leads \mathcal{A}_n from the initial state, which opens a request, to the rightmost state, which answers all requests, without visiting the latter in the meantime, if, and only if, it has the following form: $\alpha(i) \neq \#$ for all i , $\beta(i) = \#$ if, and only if $i = m$, and $\alpha(1) \dots \alpha(m)$ has two occurrences of $\beta(0)$ with no larger letter in between (a so-called bad j -pair for $j = \beta(0)$), where the second occurrence is $\alpha(m)$. Call such a word productive.

Then, the language of \mathcal{A}_n contains two types of words: first all those that can be decomposed into $w_0 x_0 (\frac{\#}{\neq \#}) w_1 x_1 (\frac{\#}{\neq \#}) w_2 x_2 (\frac{\#}{\neq \#}) \dots$ where each w_i is productive, where $\sup_i |w_i| < \infty$, and where each x_i is in $(\{1, \dots, n\}^2)^*$. The other way \mathcal{A}_n accepts a word is by reaching the accepting sink state, which it does for all words of the form $w_0 x_0 (\frac{\#}{\neq \#}) \dots w_{m-1} x_{m-1} (\frac{\#}{\neq \#}) w'_m (\frac{\#}{*}) (\{1, \dots, n\}^2)^\omega$ where each w_i is productive, each x_i is in $(\{1, \dots, n\}^2)^*$, w'_m is a strict prefix of a productive word, and $*$ denotes an arbitrary letter.

To conclude this introductory subsection, we state a replacement lemma for accepting runs: fix a parity automaton with costs $\mathcal{A} = (Q, \Sigma, q_I, \delta, \Omega, \text{Cst})$ and let $(q_0, a_0, q_1) \dots (q_{n-1}, a_{n-1}, q_n)$ be a non-empty finite run of \mathcal{A} . Its type is the tuple $(q_0, q_n, c_0, c_1, \ell)$ where $c_0, c_1 \in \Omega(Q) \cup \{\perp\}$ and $\ell \in \{\varepsilon, i\}$ are defined (using $\max \emptyset = \perp$) as

- $c_0 = \max\{\Omega(q_j) \mid 0 \leq j \leq n \text{ and } \Omega(q_j) \text{ is even}\},$
- $c_1 = \max\{\Omega(q_j) \mid 0 \leq j \leq n, \Omega(q_j) \text{ is odd, and } \Omega(q_{j'}) \notin \text{Ans}(\Omega(q_j)) \text{ for all } j' < j' \leq n\},$ i.e., c_1 is the maximal unanswered request, and

- $\ell = \mathbf{i}$ if, and only if, the run contains an increment-transition.

The following lemma shows that we can replace infixes of accepting runs by infixes of the same type, provided the replacements are of bounded length.

Lemma 1. *Let $(\rho_j^1)_{j \in \mathbb{N}}$ and $(\rho_j^2)_{j \in \mathbb{N}}$ be sequences of finite runs of a parity automaton with costs such that ρ_j^1 and ρ_j^2 have the same type for every j , and such that $\sup_j |\rho_j^2| = d < \infty$. Let $\rho' = \rho_0^1 \rho_1^1 \rho_2^1 \dots$ and $\rho'' = \rho_0^2 \rho_1^2 \rho_2^2 \dots$.*

If $\limsup_{n \rightarrow \infty} \text{Cor}(\rho', n) \leq b$ for some b , then $\limsup_{n \rightarrow \infty} \text{Cor}(\rho'', n) \leq (b+2) \cdot d$, i.e., if ρ' is an accepting run, then ρ'' is an accepting run as well.

The special case where both the lengths of the ρ_j^1 and the lengths of the ρ_j^2 are bounded is useful later on, too.

Corollary 1. *Let $(\rho_j^1)_{j \in \mathbb{N}}$ and $(\rho_j^2)_{j \in \mathbb{N}}$ be sequences of finite runs of a parity automaton with costs such that ρ_j^1 and ρ_j^2 have the same type for every j , and such that $\sup_j |\rho_j^1| < \infty$ and $\sup_j |\rho_j^2| < \infty$.*

Then, $\rho_0^1 \rho_1^1 \rho_2^1 \dots$ is an accepting run if, and only if, $\rho_0^2 \rho_1^2 \rho_2^2 \dots$ is an accepting run.

Finally, let us remark that the type of a run can be computed on the fly. Let $T_{\mathcal{A}} = Q^2 \times (\Omega(Q) \cup \{\perp\})^2 \times \{\varepsilon, \mathbf{i}\}$ be the set of types of \mathcal{A} . Define $\text{Init}_{\mathcal{A}}(q) = (q, q, c_0, c_1, \varepsilon)$ with

- $c_0 = \Omega(q)$, if $\Omega(q)$ is even, and $c_0 = \perp$ otherwise, and
- $c_1 = \Omega(q)$, if $\Omega(q)$ is odd, and $c_1 = \perp$ otherwise.

Thus, $\text{Init}_{\mathcal{A}}(q)$ can be understood as the type of the empty prefix of a run starting in state q . Furthermore, define $\text{Upd}_{\mathcal{A}}((q_0, q_1, c_0, c_1, \ell), a) = (q_0, q'_1, c'_0, c'_1, \ell')$ where

- $q'_1 = \delta(q_1, a)$,
- $c'_0 = \begin{cases} \max\{c_0, \Omega(q'_1)\} & \text{if } \Omega(q_1) \text{ is even,} \\ c_0 & \text{if } \Omega(q_1) \text{ is odd,} \end{cases}$
- $c'_1 = \begin{cases} \max\{c_1, \Omega(q'_1)\} & \text{if } \Omega(q'_1) \text{ is odd,} \\ \perp & \text{if } \Omega(q'_1) \in \text{Ans}(c_1), \text{ and} \\ c_1 & \text{else,} \end{cases}$
- $\ell' = \mathbf{i}$ if $\ell = \mathbf{i}$ or if (q_1, a, q'_1) is an increment transition.

Here, we use $\max\{\perp, c\} = c$ for every $c \in \Omega(Q)$.

Remark 2. *The run $(q, a, \delta(q, a))$ has type $\text{Upd}_{\mathcal{A}}(\text{Init}_{\mathcal{A}}(q), a)$, and for every non-empty finite run $\rho = (q_0, a_0, q_1) \dots (q_{n-1}, a_{n-1}, q_n)$ and every $a \in \Sigma$: if ρ has type t , then $\rho \cdot (q_n, a, \delta(q_n, a))$ has type $\text{Upd}_{\mathcal{A}}(t, a)$.*

B. Delay Games

A delay function is a map $f: \mathbb{N} \rightarrow \mathbb{N} \setminus \{0\}$, which is said to be constant, if $f(i) = 1$ for every $i > 0$. A delay game $\Gamma_f(L)$ consists of a delay function f and a winning condition $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$. It is played in rounds $i = 0, 1, 2, \dots$ between Player I and Player O . In round i , first Player I picks $u_i \in \Sigma_I^{f(i)}$, then Player O answers by picking $v_i \in \Sigma_O$. Player O wins the resulting play $(u_0, v_0)(u_1, v_1)(u_2, v_2) \dots$ if the outcome $(\frac{u_0 u_1 u_2 \dots}{v_0 v_1 v_2 \dots})$ is in L , else Player I wins.

A strategy for Player I in $\Gamma_f(L)$ is a mapping $\tau_I: \Sigma_O^* \rightarrow \Sigma_I^*$ such that $|\tau_I(w)| = f(|w|)$; a strategy for Player O is a

mapping $\tau_O: \Sigma_I^* \rightarrow \Sigma_O^*$. A play $(u_0, v_0)(u_1, v_1)(u_2, v_2) \dots$ is consistent with τ_I , if $u_i = \tau_I(v_0 \dots v_{i-1})$ for every i and it is consistent with τ_O , if $v_i = \tau_O(u_0 \dots u_i)$ for every i . A strategy for Player $P \in \{I, O\}$ is winning, if the outcome of every consistent play is winning for Player P . In this case we say that Player P wins $\Gamma_f(L)$. A game is determined if one of the players wins it.

Parity automata with costs recognize Borel languages [13].¹ Thus, the Borel determinacy result for delay games [6] is applicable: delay games with winning conditions given by parity automata with costs are determined.

In this work, we study delay games whose winning conditions are given by a parity automaton with costs \mathcal{A} . In particular, we determine the complexity of determining whether Player O wins $\Gamma_f(L(\mathcal{A}))$ for some f and what kind of f is necessary in general. Furthermore, we study the tradeoff between the quality and the necessary lookahead of winning strategies.

To this end, given a winning strategy τ_O for Player O in a delay game $\Gamma_f(L(\mathcal{A}))$, let

$$\text{Cst}_{\mathcal{A}}(\tau_O) = \sup_w (\limsup_{n \rightarrow \infty} \text{Cor}(\rho(w), n)),$$

where w ranges over the outcomes of τ_O and where $\rho(w)$ denotes the unique run of \mathcal{A} on w . A winning strategy for $\Gamma_f(L(\mathcal{A}))$ is optimal, if it has minimal cost among the winning strategies for $\Gamma_f(L(\mathcal{A}))$. Note that the delay function f is fixed here, i.e., the cost of optimal strategies might depend on f . Furthermore, the cost of a winning strategy might be infinite, if it has outcomes with arbitrarily large cost. This is possible, as the parity condition with costs only requires a bound for every run, but not a uniform one among all runs.

Example 2. *Consider delay games with winning conditions $L_n = L(\mathcal{A}_n)$, where \mathcal{A}_n is the parity automaton with costs from Example 1.*

We claim that Player O wins $\Gamma_f(L_n)$ if $f(0) \geq 2^n + 1$, based on the following fact: every word over $\{1, \dots, n\}$ of length 2^n contains a bad j -pair for some j [7].

Now, consider the situation after the first move of Player I : \mathcal{A}_n is in the initial state, as Player O has not yet produced any letters, and Player I has picked at least $2^n + 1$ letters $\alpha(0) \dots \alpha(f(0) - 1)$ lookahead. Let $w = \alpha(1) \dots \alpha(2^n)$. We consider several cases:

If w contains a $\#$ without a bad j -pair in the prefix before the first $\#$, then Player O can pick 1's until the automaton has reached the accepting sink state. The resulting run is accepting with cost 0.

Otherwise, w contains a bad j -pair without an earlier occurrence of $\#$, for some j . Then, Player O picks $\beta(0) = j$ and afterwards arbitrary letters (but $\#$) until the second j constituting the bad j -pair, which she answers by a $\#$. The resulting word is productive, i.e., it leads the automaton to the rightmost state. Then, either the run stays in this state forever,

¹The result proven there is about winning conditions of games, which covers languages of automata as a special case.

which implies it is accepting, or the initial state is reached again. Then, Player O can iterate her strategy, as the size of the lookahead is non-decreasing.

Hence, if the initial state is visited infinitely often, then the rightmost state of color 2 is visited at least once in every run infix of length $2^n + 1$, which implies that the resulting outcome is winning for Player O . Every other outcome has only finitely many requests and is therefore winning with cost 0. In particular, the cost of this strategy is bounded from above by $2^n + 1$.

On the other hand, Player I wins $\Gamma_f(L_n)$ if $f(0) < 2^n + 1$. Let $w_n \in \Sigma_I$ be inductively defined by $w_1 = 1$ and $w_j = w_{j-1} \cdot j \cdot w_{j-1}$ for $j > 0$. Then, w_n has length $2^n - 1$ and contains no bad j -pair for every j , and trivially no $\#$.

Let Player I pick the prefix of length $f(0)$ of $1w_n$ in the first round and let Player O answer by picking $\beta(0)$. If $\beta(0) = \#$, then \mathcal{A}_n reaches the rejecting sink-state of color 1. If $\beta(0) = j$, then Player I picks j' ad infinitum for some $j' \notin \{j, \#\}$. Then, \mathcal{A}_n neither reaches the accepting sink state (since he never plays a $\#$) nor the rightmost state (as he has not produced a bad j -pair). Hence, the resulting run is rejecting.

One can even show an exponential lower bound on the cost of an optimal winning strategy for Player O : using the word w_n to begin each round and then restart the play by picking $\#$ after Player O has reached the rightmost state by picking $\#$, he can enforce a cost-of-response of $2^n + 1$ for infinitely many requests, i.e., a winning strategy for Player O has at least cost $2^n + 1$ as well.

Proposition 1. *For every n , there is a finitary parity automaton \mathcal{A}_n of size $\mathcal{O}(n)$ such that Player O wins $\Gamma_f(L(\mathcal{A}_n))$ if, and only if, $f(0) \geq 2^n + 1$. Furthermore, the cost of an optimal winning strategy for $\Gamma_f(L(\mathcal{A}_n))$ is $2^n + 1$, for every such f .*

III. CONSTANT LOOKAHEAD SUFFICES

In this section, we prove that constant delay functions suffice for Player O to win delay games with winning conditions specified by parity automata with costs and give an exponential upper bound on the necessary initial lookahead. To this end, we generalize a similar result proven for delay games with winning conditions given by parity automata [7]. This is the first step towards proving that the winner of such a game can be determined effectively.

Theorem 1. *Let L be recognized by a parity automaton with costs \mathcal{A} with n states and k colors, and let f be the constant delay function with $f(0) = 2^{2n^4k^2+1}$. The following are equivalent:*

- 1) Player I wins $\Gamma_f(L(\mathcal{A}))$.
- 2) Player I wins $\Gamma_{f'}(L(\mathcal{A}))$ for every delay function f' .

Proof. We only prove the non-trivial direction by taking a winning strategy for $\Gamma_f(L(\mathcal{A}))$ and pumping its moves to obtain a winning strategy for $\Gamma_{f'}(L(\mathcal{A}))$, which might require Player I to provide more lookahead.

To this end, let $\mathcal{A} = (Q, \Sigma_I \times \Sigma_O, q_I, \delta, \Omega, \text{Cst})$, let $T_{\mathcal{A}}$ be the set of types of \mathcal{A} , and recall that $\text{Init}_{\mathcal{A}}$ and $\text{Upd}_{\mathcal{A}}$ compute the type of a run as described in Remark 2. We extend \mathcal{A} so that it tracks the type of its runs using the state set $Q \times T_{\mathcal{A}}$. This information does not change the language of the automaton, but is useful when pumping the moves in $\Gamma_f(L(\mathcal{A}))$.

Formally, we define $\mathcal{A}' = (Q', \Sigma_I \times \Sigma_O, q'_I, \delta', \Omega', \text{Cst}')$ with

- $Q' = Q \times T_{\mathcal{A}}$,
- $q'_I = (q_I, \text{Init}_{\mathcal{A}}(q_I))$,
- $\delta'((q, t), \binom{a}{b}) = (\delta(q, \binom{a}{b}), \text{Upd}(t, \binom{a}{b}))$,
- $\Omega'(q, t) = \Omega(q)$, and
- $\text{Cst}'((q, t), \binom{a}{b}, (q', t')) = \Omega(q, \binom{a}{b}, q')$.

Now, define $\delta_P: 2^{Q'} \times \Sigma_I \rightarrow 2^{Q'}$ via

$$\delta_P(S, a) = \left\{ \delta' \left((q, t), \binom{a}{b} \right) \mid (q, t) \in S \text{ and } b \in \Sigma_O \right\}.$$

Intuitively, δ_P is obtained as follows: take \mathcal{A}' , project away Σ_O , and apply the power set construction (ignoring the coloring and the costs). Then, δ_P is the transition function of the resulting deterministic automaton. As usual, we extend δ_P to $\delta_P^+: 2^{Q'} \times \Sigma_I^+ \rightarrow 2^{Q'}$ via $\delta_P^+(S, a) = \delta_P(S, a)$ and $\delta_P^+(S, wa) = \delta_P(\delta_P^+(S, w), a)$.

Remark 3. *The following are equivalent for $q \in Q$ and $w \in \Sigma_I^+$:*

- 1) $(q', t') \in \delta_P^+(\{(q, \text{Init}_{\mathcal{A}}(q))\}, w)$.
- 2) *There is a $w' \in (\Sigma_I \times \Sigma_O)^+$ whose projection to Σ_I is w such that the run of \mathcal{A} processing w' from q ends in q' and has type t' .*

Let τ_I be a winning strategy for Player I in $\Gamma_f(L(\mathcal{A}))$ and let f' be an arbitrary delay function. We construct a winning strategy τ'_I for Player I in $\Gamma_{f'}(L(\mathcal{A}))$ by simulating a play of $\Gamma_{f'}(L(\mathcal{A}))$ by a play in $\Gamma_f(L(\mathcal{A}))$. For the sake of brevity, we denote $\Gamma_f(L(\mathcal{A}))$ by Γ and $\Gamma_{f'}(L(\mathcal{A}))$ by Γ' from now on.

Recall that we only consider automata \mathcal{A} with non-trivial colorings. Thus, we can bound the number of \mathcal{A} 's types by $2n^2k^2$, where n is the number of states and k the number of colors. We define $d = |(2^{Q'})^Q| = 2^{2n^4k^2}$. In the simulating play in Γ , the players make their moves in blocks of length d : Player I 's are denoted by \bar{a}_i and Player O 's by \bar{b}_i , i.e., in the following, every \bar{a}_i is in Σ_I^d and every \bar{b}_i is in Σ_O^d . Furthermore, we say that a decomposition $\bar{a}_i = xyz$ is *pumpable*, if y is non-empty and if

$$\delta_P^+(\{(q, \text{Init}_{\mathcal{A}}(q))\}, x) = \delta_P^+(\{(q, \text{Init}_{\mathcal{A}}(q))\}, xy)$$

for every $q \in Q$. As \mathcal{A} is complete, $\delta_P^+(\{(q, \text{Init}_{\mathcal{A}}(q))\}, w)$ is always non-empty, which implies the following remark.

Remark 4. *Every $\bar{a} \in \Sigma_I^d$ has a pumpable decomposition.*

Now, we begin the construction of τ'_I . Note that we have $f(0) = 2d$. Thus, let $\tau_I(\varepsilon) = \bar{a}_0\bar{a}_1$ be the first move of Player I in Γ according to τ_I . Remark 4 yields pumpable decompositions $\bar{a}_0 = x_0y_0z_0$ and $\bar{a}_1 = x_1y_1z_1$. We pick $h_0 > 0$ such that $|x_0(y_0)^{h_0}z_0| \geq f'(0)$ and define $\alpha(0) \cdots \alpha(\ell_0 -$




Fig. 2. *The situation* (in solid lines): in Γ , Player I has picked $\overline{a_0}, \dots, \overline{a_i}$, Player O has picked $\overline{b_0}, \dots, \overline{b_{i-2}}$ (hidden in the thin part at the beginning), and q_{i-1} is the state of \mathcal{A} reached when processing $\left(\frac{\overline{a_0}}{b_0}\right) \dots \left(\frac{\overline{a_{i-2}}}{b_{i-2}}\right)$ (denoted by the curly line). In Γ' , Player I has repeated y_i sufficiently often so that Player O has provided an answer to $x_{i-1}y_{i-1} \dots y_{i-1}z_{i-1}$, i.e., up to position $\ell_{i-1} - 1$, and q_{i-1}^* is the state reached when processing x_{i-1} and all but the last copy of y_{i-1} (and the corresponding answers of Player O) starting in q_{i-1} . By construction, there is an answer x'_{i-1} to x_{i-1} such that processing $\left(\frac{x_{i-1}}{x'_{i-1}}\right)$ from q_{i-1} brings \mathcal{A} to q_{i-1}^* as well. The block b_{i-1} (dotted) is the concatenation of x'_{i-1} , and $y'_{i-1}z'_{i-1}$ from Γ' .

1) $= x_0(y_0)^{h_0}z_0$. Similarly, we pick $h_1 > 0$ such that $|x_1(y_1)^{h_1}z_1| \geq \sum_{j=1}^{\ell_0-1} f'(j)$ and define $\alpha(\ell_0) \dots \alpha(\ell_1 - 1) = x_1(y_1)^{h_1}z_1$. Now, we define the strategy τ'_I for Player I in Γ' to pick the prefix of length $\sum_{j=0}^{\ell_0-1} f'(j)$ of $\alpha(0) \dots \alpha(\ell_1 - 1)$ during the first ℓ_0 rounds, independently of the choices of Player O . This prefix is well-defined by the choices of h_0 and h_1 . The remaining letters of $\alpha(0) \dots \alpha(\ell_1 - 1)$ are stored in a buffer γ_1 . During these first ℓ_0 rounds, Player O answers by producing some $\beta(0) \dots \beta(\ell_0 - 1)$.

Thus, we are in the following situation for $i = 1$ (see the solid part of Figure 2).

- In Γ , Player I has picked $\overline{a_0}, \dots, \overline{a_i}$ such that for every $j \leq i$: $x_j y_j z_j$ is a pumpable decomposition of $\overline{a_j}$. Furthermore, Player O has picked $\overline{b_0}, \dots, \overline{b_{i-2}}$.
- In Γ' , Player I has picked the prefix of length $\sum_{j=0}^{\ell_{i-1}-1} f'(j)$ of

$$\alpha(0) \dots \alpha(\ell_i - 1) = x_0(y_0)^{h_0}z_0 \dots x_i(y_i)^{h_i}z_i$$

while the remaining suffix is the buffer γ_i . Player O has picked $\beta(0) \dots \beta(\ell_{i-1} - 1)$.

Now, let $i > 0$ be arbitrary and let q_{i-1} be the state reached by \mathcal{A} when processing $\left(\frac{\overline{a_0}}{b_0}\right) \dots \left(\frac{\overline{a_{i-2}}}{b_{i-2}}\right)$ from q_I . Furthermore, let q_{i-1}^* be the state reached by \mathcal{A} when processing $x_{i-1}(y_{i-1})^{h_{i-1}-1}$ and the corresponding part of $\beta(0) \dots \beta(\ell_{i-1} - 1)$ starting in q_{i-1} , and let t_{i-1}^* be the type of the run. Then, due to Remark 3 and the decomposition $x_{i-1}y_{i-1}z_{i-1}$ being pumpable, there is a word $x'_{i-1} \in \Sigma_O^{|x_{i-1}|}$ such that \mathcal{A} reaches the same state q_{i-1}^* when processing $\left(\frac{x_{i-1}}{x'_{i-1}}\right)$ starting in q_{i-1} , and the run has type t_{i-1}^* as well. Now, we define $\overline{b_{i-1}} = x'_{i-1}y'_{i-1}z'_{i-1}$ where y'_{i-1} and z'_{i-1} are the letters picked by Player O at the positions of the last repetition of y_{i-1} and at the positions of z_{i-1} , respectively (see Figure 2).

Using $\overline{b_{i-1}}$ we continue the simulation in Γ by letting Player O pick the letters of $\overline{b_{i-1}}$ during the next d rounds, which yields d moves for Player I by applying τ_I . Call this sequence of letters $\overline{a_{i+1}}$, which again has a pumpable decomposition $x_{i+1}y_{i+1}z_{i+1}$. We again pick $h_{i+1} > 0$ such that $|x_{i+1}(y_{i+1})^{h_{i+1}}z_{i+1}| \geq \sum_{j=\ell_{i-1}}^{\ell_i-1} f'(j)$ and define $\alpha(\ell_i) \dots \alpha(\ell_{i+1} - 1) = x_{i+1}(y_{i+1})^{h_{i+1}}z_{i+1}$. The strategy τ'_I for Player I in Γ' is defined so that it picks the prefix of length $\sum_{j=\ell_{i-1}}^{\ell_i-1} f'(j)$ of $\gamma_i \alpha(\ell_i) \dots \alpha(\ell_{i+1} - 1)$ during the next ℓ_i rounds, independently of the choices of Player O , and the remaining letters are stored in the buffer γ_{i+1} . The prefix is again well-defined by the choice of h_{i+1} . Hence, during these rounds, Player O answers by producing $\beta(\ell_{i-1}) \dots \beta(\ell_i - 1)$. Then, we are again in the situation described above for $i + 1$.

We conclude by showing that τ'_I is indeed a winning strategy for Player I in Γ' . To this end, let $w' = \left(\frac{\alpha}{\beta}\right)$ be an outcome of a play that is consistent with τ'_I and let $w = \left(\frac{\overline{a_0}}{b_0}\right) \left(\frac{\overline{a_1}}{b_1}\right) \left(\frac{\overline{a_2}}{b_2}\right) \dots$ be the play in Γ constructed during the simulation as described above, which is consistent with τ_I and therefore winning for Player I . Hence, the run of \mathcal{A} on w' is rejecting.

A simple induction shows that \mathcal{A} reaches the same state when processing

$$\left(\frac{\alpha(0) \dots \alpha(\ell_i - 1)}{\beta(0) \dots \beta(\ell_i - 1)}\right) \quad \text{and} \quad \left(\frac{\overline{a_0} \dots \overline{a_i}}{\overline{a_b} \dots \overline{b_i}}\right),$$

call it q_i . By construction, the run of \mathcal{A} starting in q_i processing $\left(\frac{\alpha(\ell_{i-1}) \dots \alpha(\ell_i - 1)}{\beta(\ell_{i-1}) \dots \beta(\ell_i - 1)}\right)$ (using $\ell_{-1} = 0$) and the one starting in q_i processing $\left(\frac{\overline{a_i}}{b_i}\right)$ have the same type. The runs on the $\left(\frac{\overline{a_i}}{b_i}\right)$ all have length d , hence applying Lemma 1 to the rejecting run of \mathcal{A} on w shows that the run of \mathcal{A} on w' is rejecting as well. Thus, the outcome is winning for Player I and τ_I is indeed a winning strategy for Player I in Γ' . \square

Applying both directions of the equivalence proved in Theorem 1 and determinacy yields an upper bound on the

necessary lookahead for Player O .

Corollary 2. *Let L be recognized by a parity automaton with costs \mathcal{A} with n states and k colors. If Player O wins $\Gamma_f(L)$ for some delay function f , then also for the constant delay function f with $f(0) = 2^{2n^4k^2+1}$.*

Finally, the upper bound of Corollary 2 is asymptotically tight due to Proposition 1, which is a generalization of the corresponding lower bound for delay games with winning conditions given by deterministic safety automata [7].

IV. DETERMINING THE WINNER

The main result of this section is that the following problem is **EXPTIME**-complete: given a parity automaton with costs \mathcal{A} , does Player O win $\Gamma_f(L(\mathcal{A}))$ for some f ? Hardness already holds for the special case of safety automata, thus we focus our attention on membership. To this end, we revisit the analogous result for delay games with classical parity conditions (i.e., without costs) [7]: such games are reduced to equivalent delay-free parity games of exponential size, which can be solved in exponential time (in the size of the original parity automaton). Here, we extend this proof to automata with costs while simplifying its structure. Furthermore, we obtain an exponential upper bound on the cost of a winning strategy for Player O .

For the remainder of this section, fix $\mathcal{A} = (Q, \Sigma_I \times \Sigma_O, q_I, \delta, \Omega, \text{Cst})$, let $\mathcal{A}' = (Q', \Sigma_I \times \Sigma_O, q'_I, \delta', \Omega', \text{Cst}')$ and $\delta_P^+ : 2^{Q'} \times \Sigma_I \rightarrow 2^{Q'}$ be defined as in Section III, and recall that $Q' = Q \times T_{\mathcal{A}}$, where $T_{\mathcal{A}}$ is the set of types of runs of \mathcal{A} .

Given $x \in \Sigma_I^+$, we define the function $r_x : Q \rightarrow 2^{Q'}$ via

$$r_x(q) = \delta_P^+(\{(q, \text{Init}_{\mathcal{A}}(q))\}, x).$$

Now, we define $x \equiv_{\mathcal{A}} x'$ if, and only if, $r_x = r_{x'}$, which is a finite equivalence relation. Furthermore, we can assign to every $\equiv_{\mathcal{A}}$ equivalence class S a function r_S from Q to $2^{Q'}$, i.e., $r_S = r_x$ for all $x \in S$, which is independent of representatives. Finally, let \mathfrak{R} denote the set of $\equiv_{\mathcal{A}}$ equivalence classes of words in Σ_I^{2d} , where $d = 2^{2n^4k^2}$ as before.

Next, we construct a delay-free game $\mathcal{G}(\mathcal{A})$ between Player I and Player O that is won by Player O if, and only if, she wins $\Gamma_f(L(\mathcal{A}))$ for some delay function f . The game $\mathcal{G}(\mathcal{A})$ is a zero-sum infinite-duration two-player game of perfect information played in rounds $i = 0, 1, 2, \dots$. Intuitively, Player I picks a sequence $S_0 S_1 S_2 \dots$ of equivalence classes from \mathfrak{R} , which induces an infinite word α over Σ_I by picking representatives. Player O implicitly picks an infinite word over Σ_O by constructing a run of \mathcal{A} on a word over $\Sigma_I \times \Sigma_O$ whose projection to Σ_I is α . She wins, if the run is accepting. To account for the delay, she is always one move behind.

Formally, in round 0, Player I picks an equivalence class $S_0 \in \mathfrak{R}$ and then Player O has to pick $(q_0, t_0) = q'_I$, the initial state of \mathcal{A}' .² In round $i > 0$, Player I picks an equivalence class $S_i \in \mathfrak{R}$ and then Player O picks

a state $(q_i, t_i) \in r_{S_{i-1}}(q_{i-1})$ (due to completeness of \mathcal{A} , Player O always has an available move).

For every $t \in T_{\mathcal{A}}$ such that there is some run of \mathcal{A} of type t , fix one such run ρ_t . Now, consider a play $\pi = S_0(q_0, t_0)S_1(q_1, t_1)S_2(q_2, t_2) \dots$ of $\mathcal{G}(\mathcal{A})$. By construction, $\rho_{t_{i+1}}$ starts in q_i and ends in q_{i+1} , for every $i \geq 0$. The play π is winning for Player O if, and only if, the run $\rho_{t_1}\rho_{t_2}\rho_{t_3} \dots$ of \mathcal{A} is accepting (note that t_0 is disregarded). As the length of the representatives ρ_{t_i} is bounded (there are only finitely many), Corollary 1 implies that the winner is independent of the choice of the representatives.

A strategy for Player I in $\mathcal{G}(\mathcal{A})$ is a mapping $\tau_I : (\mathfrak{R} \cdot Q')^* \rightarrow \mathfrak{R}$ while a strategy for Player O is a mapping $\tau_O : (\mathfrak{R} \cdot Q')^* \cdot \mathfrak{R} \rightarrow Q'$ that has to satisfy $\tau_O(S_0) = q'_I$ for every $S_0 \in \mathfrak{R}$ and $\tau_O(S_0 \dots (q_i, t_i)S_{i+1}) \in r_{S_i}(q_i)$ for all $S_0 \dots (q_i, t_i)S_{i+1} \in (\mathfrak{R} \cdot Q')^+ \cdot \mathfrak{R}$. A play $S_0(q_0, t_0)S_1(q_1, t_1)S_2(q_2, t_2) \dots$ is consistent with τ_I if $S_i = \tau_I(S_0 \dots (q_{i-1}, t_{i-1}))$ for every i , and it is consistent with τ_O , if $(q_i, t_i) = \tau_O(S_0 \dots S_i)$ for every i . A strategy is winning for a Player $P \in \{I, O\}$, if every play that is consistent with the strategy is won by Player P . As usual, we say that Player P wins $\mathcal{G}(\mathcal{A})$, if she has a winning strategy.

Lemma 2. *Player O wins $\mathcal{G}(\mathcal{A})$ if, and only if, she wins $\Gamma_f(L(\mathcal{A}))$ for some f .*

Now, we are able to state and prove our main theorem of this section.

Theorem 2. *The following problem is **EXPTIME**-complete: given a parity automaton with costs \mathcal{A} , does Player O win $\Gamma_f(L(\mathcal{A}))$ for some f ?*

Proof. We focus on membership as **EXPTIME**-hardness already holds for safety automata [7]. To this end, we show how to model the abstract game $\mathcal{G}(\mathcal{A})$ as an arena-based parity game with costs [13] of exponential size with at most one more color than \mathcal{A} . This game can be constructed (argued below) and solved in exponential time (in the size of \mathcal{A}) [13]. Lemma 2 shows that solving this game yields the correct answer. Hence, the problem is in **EXPTIME**.

Intuitively, the arena encodes the rules of $\mathcal{G}(\mathcal{A})$: the players pick equivalence classes from \mathfrak{R} and states from Q' in alternation. The restrictions on the states that may be picked are enforced by storing the last equivalence class and the last state in the vertices of the arena. Finally, to encode the winning condition of $\mathcal{G}(\mathcal{A})$, we simulate the effect of a run of type $t = (q, q', c, c', \ell)$ every time Player O picks a state (q, t) . Recall that c encodes the largest answer, c' the largest unanswered request, and ℓ whether the overall cost is zero or greater than zero. The effect is simulated by first visiting a state of color c , then one of color c' , and equipping the edge between these vertices with cost ℓ . Afterwards, Player I again picks another equivalence class.

Formally, we define the parity game with costs $(\mathcal{A}, \text{CostParity}(\Omega))$ with arena $\mathcal{A} = (V, V_I, V_O, E, \text{Cst})$ where

²This move is trivial, but we add it to keep the definition consistent.




Fig. 3. The construction of \mathcal{A} : a play prefix $S_0, (q_0, t_0)S_1, (q_1, t_1), S_2, (q_2, t_2)$ of $\mathcal{G}(\mathcal{A})$ with $t_i = (q_i, q'_i, c_i^0, c_i^1, \ell_i)$ is simulated by the depicted sequence of vertices. Colors are depicted above vertices, edge weights above edges.

- $V = V_I \cup V_O$ with $V_I = \{v_I\} \cup \mathfrak{R} \times Q' \times \{0, 1\}$ and $V_O = \mathfrak{R} \times Q' \times \mathfrak{R}$, and
- E is the union of the following sets of edges:
 - $\{(v_I, (S_0, q'_I, S_1)) \mid S_0, S_1 \in \mathfrak{R}\}$: the initial moves of Player I (which subsume the first (trivial) move of Player O),
 - $\{((S_0, (q_0, t_0), 0), (S_0, (q_0, t_0), 1)) \mid S_0 \in \mathfrak{R}, (q_0, t_0) \in Q'\}$: deterministic moves of Player I used to simulate the effect of a run of type t_0 ,
 - $\{((S_0, (q_0, t_0), 1), (S_0, (q_0, t_0), S_1)) \mid S_0, S_1 \in \mathfrak{R}, (q_0, t_0) \in Q'\}$: regular moves of Player I picking the next equivalence class S_1 , and
 - $\{((S_0, (q_0, t_0), S_1), (S_1, (q_1, t_1), 0)) \mid S_0, S_1 \in \mathfrak{R}, (q_0, t_0), (q_1, t_1) \in Q', (q_1, t_1) \in r_{S_0}(q_0)\}$: moves of Player O picking the next state (q_1, t_1) .
- We define $\text{Cst}(e) = i$, if, and only if, $e = ((S_0, (q_0, t_0), 0), (S_0, (q_0, t_0), 1))$ with $t_0 = (q, q', c, c', i)$ for some $q, q' \in Q$ and $c, c' \in \Omega(Q)$, i.e., we simulate the cost encoded in the type t_0 .
- Finally, for $t = (q, q', c_0, c_1, \ell)$ we define $\Omega(S, (q, t), 0) = c_0$ and $\Omega(S, (q, t), 1) = c_1$, i.e., we simulate the largest response and afterwards the largest unanswered request encoded in t . Every other state has color 0, which has no effect on the satisfaction of the parity condition with costs, as it is too small to answer requests.

As an illustration of the construction, consider Figure 3, which depicts the vertices reached while simulating a play prefix of $\mathcal{G}(\mathcal{A})$. Note that the infix $(S_1, (q_1, t_1), 0)(S_1, (q_1, t_1), 1)$ has largest response c_0^1 and largest unanswered request c_1^1 , just as encoded by the type t_1 . Similarly, the cost of this infix is the one encoded by t_1 . All other vertices and edges are neutral. Furthermore, the type t_0 encoded in the initial state $q'_I = (q_I, \text{Init}_{\mathcal{A}}(q_I)) = (q_0, t_0)$ is not simulated, just as it is ignored when it comes to determining the winner of a play in $\mathcal{G}(\mathcal{A})$.

Thus, Corollary 1, implies that Player O wins $\mathcal{G}(\mathcal{A})$ if, and only if, Player O has a winning strategy for the parity game with costs $(\mathcal{A}, \text{CostParity}(\Omega))$ from v_I . The construction of $(\mathcal{A}, \text{CostParity}(\Omega))$ is possible in exponential time using the same automata construction to determine the elements of \mathfrak{R} as in the case of plain parity conditions [7]. \square

If Player O wins an arena-based parity game with costs, then there is also a winning strategy for her whose cost is bounded

by the number of vertices of the arena [13], [15]. Hence, an application of Lemma 1 yields an exponential upper bound on the cost of a winning strategy in a delay game with such a winning condition.

Corollary 3. *Let \mathcal{A} be a parity automaton with costs with n states and k colors. If Player O wins $\Gamma_f(L(\mathcal{A}))$ for some f , then she also wins $\Gamma_f(L(\mathcal{A}))$ for the constant delay function f given by $f(0) = 2^{2n^4k^2+2}$ with a winning strategy τ_O satisfying $\text{Cst}_{\mathcal{A}}(\tau_O) \leq n^3k^22^{2n^7k^4+3}$.*

Due to Proposition 1, these bounds are asymptotically tight.

V. STREETT CONDITIONS

In this section, we consider the more general case of delay games with winning conditions given by automata with finitary Streett acceptance or with Streett conditions with costs. In a parity condition, the requests and responses are hierarchically ordered. Streett conditions generalize parity conditions by giving up this hierarchy.

Formally, a Streett automaton with costs is a tuple $\mathcal{A} = (Q, \Sigma, q_I, \delta, (Q_j, P_j)_{j \in J}, (\text{Cst}_j)_{j \in J})$ where Q , Σ , q_I , and δ are defined as for parity automata with costs. Furthermore, the acceptance condition $(Q_j, P_j)_{j \in J}$ consists of a finite collection of Streett pairs (Q_j, P_j) of subsets $Q_j, P_j \subseteq Q$. Here, states in Q_j are requests of condition j which are answered by visiting a response in P_j . Finally, $(\text{Cst}_j)_{j \in J}$ is a collection of cost functions for \mathcal{A} , one for each Streett pair. The size of \mathcal{A} is defined as $|Q| + |J|$.

For a run $(q_0, a_0, q_1)(q_1, a_1, q_2)(q_2, a_2, q_3) \dots$ and a position n , we define the cost-of-response $\text{StCor}_j(\rho, n)$ of pair j to be 0, if $q_n \notin Q_j$, and to be

$$\min\{\text{Cst}_j((q_n, a_n, q_{n+1}) \dots (q_{n'-1}, a_{n'-1}, q_{n'})) \mid n' \geq n \text{ and } q_{n'} \in P_j\},$$

if $q_n \in Q_j$. Furthermore, we aggregate these costs by defining

$$\text{StCor}(\rho, n) = \max_{j \in J} \text{StCor}_j(\rho, n).$$

Finally, we say that ρ is accepting if it satisfies the Streett condition with costs, i.e., if $\limsup_{n \rightarrow \infty} \text{StCor}(\rho, n) < \infty$. We recover classical Streett acceptance as the special case where every edge is an ε -edge w.r.t. every cost function. Similarly, finitary Streett [12] acceptance is the special case where every edge is an increment-edge w.r.t. every cost function.

In the following, we consider delay games with winning conditions specified by such automata. The notion of the cost of a strategy for Player O and that of optimality is defined as in the case of parity automata. Also, note that every parity condition is a Streett condition, i.e., all lower bounds already proven hold for Streett conditions as well. In particular, exponential lookahead is necessary to win delay games with finitary Streett conditions and solving such games is **EXPTIME**-hard. We complement these lower bounds by doubly-exponential upper bounds, both on the necessary lookahead and on the solution complexity. It is open whether this gap can be closed. Nevertheless, we show in Section VI that doubly-exponential lookahead is necessary for *optimal* strategies!

Our first step towards these results is the generalization of the replacement lemma for parity conditions with costs. To this end, we have to generalize the notion of types. To this end, fix a Streett automaton with costs $\mathcal{A} = (Q, \Sigma, q_I, \delta, (Q_j, P_j)_{j \in J}, (\text{Cst}_j)_{j \in J})$. The set of types of \mathcal{A} is defined as $T_{\mathcal{A}} = Q^2 \times \{\perp, p, q, pq\}^J \times \{\varepsilon, i\}^J$. The type of a non-empty finite run $(q_0, a_0, q_1) \cdots (q_{n-1}, a_{n-1}, q_n)$ is defined as (q_0, q_n, g, ℓ) where

- $g(j) = pq$ if the run contains a response of condition j as well as an unanswered request of condition j ,
- $g(j) = p$ if the run contains a response of condition j , but no unanswered request of condition j ,
- $g(j) = q$ if the run contains no response of condition j , but an unanswered request of condition j , and
- $g(j) = \perp$ if the run contains neither a request nor a response of condition j .

Furthermore, $\ell(j)$ is equal to i if, and only if, the run contains an increment-transition with respect to Cst_j .

With this definition, the replacement property formalized in Lemma 1 also holds for runs of Streett automata with costs, which is proven using essentially the same argument as for parity automata with costs. Similarly, Corollary 1 holds for Streett automata with costs as well. Also, as for parity automata with costs, the type of a run can be computed on the fly using functions $\text{Init}_{\mathcal{A}}$ and $\text{Upd}_{\mathcal{A}}$ with the same properties as their analogues in Remark 2.

Using these results, we determine upper bounds on the necessary lookahead and the complexity of solving delay games induced by Streett automata with costs. Here, the exponential increase in complexity in comparison to games induced by parity automata with costs stems from the fact that there are exponentially many types for Street automata, but only polynomially many for parity automata.

Theorem 3. *Let L be recognized by a Streett automaton with costs \mathcal{A} with n states and k Streett pairs, and let f be the constant delay function with $f(0) = 2^{n \cdot 2^{3k} + 1}$. The following are equivalent:*

- 1) *Player I wins $\Gamma_f(L(\mathcal{A}))$.*
- 2) *Player I wins $\Gamma_{f'}(L(\mathcal{A}))$ for every delay function f' .*

Proof. Similar to the one of Theorem 1 using $d = |(2^{Q \times T_{\mathcal{A}}})^Q| = 2^{n \cdot 2^{3k}}$. \square

Thus, we obtain $f(0) = 2^{n \cdot 2^{3k} + 1}$ as an upper bound on the necessary constant lookahead for Player O to win a delay game with winning condition L .

Also, the decidability proof for parity conditions is applicable to Streett conditions, again with an exponential blowup.

Theorem 4. *The following problem is in **2EXPTIME**: given a Streett automaton with costs \mathcal{A} , does Player O win $\Gamma_f(L(\mathcal{A}))$ for some f ?*

Proof. Given \mathcal{A} , one constructs an abstract game $\mathcal{G}(\mathcal{A})$ as for the parity case and proves the analogue of Lemma 2. Then, one models $\mathcal{G}(\mathcal{A})$ as an arena-based Streett game with costs of doubly-exponential size with the same number of Streett pairs as \mathcal{A} , which can be solved in doubly-exponential time [13]. \square

Again, modeling the abstract game as an arena-based game yields an upper bound on the cost of a winning strategy for a delay game with winning condition given by a Streett automaton with costs: in an arena-based Streett game with costs, there is a tight exponential bound on the cost of an optimal strategy [13], [20]. This implies a triply-exponential upper bound for the original delay game.

VI. TRADING LOOKAHEAD FOR COSTS

Introducing lookahead allows Player O to win games she loses in a delay-free setting. In this section, we study another positive effect of lookahead: it allows to reduce the cost of optimal strategies, i.e., one can trade lookahead for quality and vice versa. To simplify our notation, let f_k for $k \geq 0$ denote the unique constant delay function with $f_k(0) = k + 1$. Thus, k denotes the size of the lookahead. In particular, a delay game $\Gamma_{f_0}(L)$ is a delay-free game.

A. Tradeoffs for Parity Conditions with Costs

First, we consider parity conditions and show that already the smallest possible lookahead allows to improve the cost of an optimal strategy from $|\mathcal{A}|$ to 1.

Theorem 5. *For every $n > 0$, there is a language L_n recognized by a finitary Büchi automaton with costs \mathcal{A}_n with $n + 2$ states such that*

- *an optimal strategy for $\Gamma_{f_0}(L_n)$ has cost n , but*
- *an optimal strategy for $\Gamma_{f_1}(L_n)$ has cost 1.*

Proof. Consider the finitary Büchi automaton \mathcal{A}_n depicted in Figure 4 over $\Sigma_I \times \Sigma_O = \{0, 1\}^2$. Every run of \mathcal{A}_n visits the initial state infinitely often, which answers all requests. Thus, consider a run starting and ending in the initial state, but not visiting it in between. There are two types of such runs, those of length two and those of length $n + 1$ (visiting the gray state). Runs of the first type process a word of the form $\binom{*}{b} \binom{b}{*}$ for some $b \in \{0, 1\}$ and an arbitrary letter $*$, runs of the second type a word of the form $\binom{*}{b} \binom{1-b}{*} \binom{*}{*}^{n-1}$. A run having infinitely many infixes of the second type has cost n , otherwise it has cost 1. Thus, to achieve cost 1, Player O has

to predict the next move of Player I . This is possible with constant lookahead 1, but not without lookahead. \square

Another simple example shows that even exponential lookahead might be necessary to achieve the smallest cost possible, relying on the exponential lower bound shown in Example 2.

Theorem 6. *For every $n > 0$, there is a language L'_n recognized by a finitary Büchi automaton with costs \mathcal{A}'_n with $\mathcal{O}(n)$ states such that*

- Player O wins $\Gamma_f(L'_n)$ for every delay function f , but
- an optimal strategy for $\Gamma_{f_{2^n}}(L'_n)$ has cost 0, and
- an optimal strategy for $\Gamma_{f_k}(L'_n)$ for $k < 2^n$ has cost n .

Finally, we generalize Theorem 5 to a gradual tradeoff, i.e., with every additional increase of the lookahead decreases the cost of an optimal strategy, up to some upper bound.

Theorem 7. *For every $n > 0$, there is a language L''_n recognized by a finitary Büchi automaton with costs \mathcal{A}''_n with $\mathcal{O}(n^2)$ states such that for every $j \in \{0, 1, \dots, n\}$: an optimal strategy for $\Gamma_{f_j}(L''_n)$ exists, but has cost $2(n+1) - j$.*

Proof. The finitary Büchi automaton \mathcal{A}''_n has alphabet $\{0, 1\} \times \{(i, j) \mid i \in \{1, \dots, n\}, j \in \{0, 1\}\}$ and is depicted in Figure 5 for $n = 3$.

The automaton generalizes the idea from Theorem 5. For every $j \in \{1, 2, 3\}$, with lookahead j at the initial state, Player O can use a transition of the form (j, b) , where b is the letter picked by Player I j positions ahead. The resulting run infix from the initial state back to it has a request that is answered with cost $2(n+1) - j$, none with larger cost, and ends with all requests being answered. With less lookahead, Player I can falsify the prediction by moving from the gray states to the sink state and thereby win. Thus, an optimal strategy for $\Gamma_{f_j}(L''_n)$ with $j > 0$ has cost $2(n+1) - j$.

Finally, for $j = 0$, Player O has to always pick a letter of the form $(1, b)$ when at the initial state. The prediction can be immediately falsified by Player I by picking $1 - b$ in the next round, leading to a request that is answered with cost $8 = 2(n+1) - j$.

The automaton \mathcal{A}''_3 can easily be generalized to an arbitrary n by allowing Player O for every $j \in \{1, \dots, n\}$ to predict the letter picked by Player I j positions ahead with a cost of $2(n+1) - j$. \square

After exhibiting these tradeoffs, a natural question concerns upper bounds on the tradeoff between quality and lookahead. The results on (delay-free, arena-based) parity games with costs imply that an optimal strategy for $\Gamma_{f_0}(L(\mathcal{A}))$ has cost at most $2|\mathcal{A}|$ (with a little more effort, the factor 2 can be eliminated): every such game can be modeled as an arena-based parity game with costs by splitting the transitions of the automaton into two moves. For such games, it is known that the cost of an optimal strategy is at most the number of states of the arena [13], [15]. On the other hand, we have shown in Corollary 3 that exponential lookahead and exponential cost is (simultaneously) achievable, if Player O wins at all.

These results constrain the type of lookahead exhibited in the previous theorems.

B. Tradeoffs for Streett Conditions with Costs

To conclude this section, we consider Streett conditions with costs. Recall that there is a trivial exponential lower bound on the necessary lookahead in delay games with Streett conditions with costs obtained from the same lower bound for parity conditions with costs. However, we only proved a doubly-exponential upper bound. Next, we show that this upper bound is tight, when considering strategies realizing the smallest possible cost with respect to *all* delay functions.

To this end, we consider a modification of the *bad j-pair game* described in Examples 1 and 2 showing an exponential lower bound on the necessary lookahead for parity conditions with costs. Recall that Player O needs lookahead $2^n + 1$ when picking numbers from $\{1, \dots, n\}$. Thus, to prove a doubly-exponential lower bound, it suffices to implement this game with numbers from the range $\{0, \dots, 2^n - 1\}$ (encoded in binary to keep the alphabet small) by an automaton of polynomial size in n . However, we have to modify the rules of the game, as such a small automaton cannot recognize the winning condition, which requires to distinguish 2^n different choices for y_0 . Instead, we would like to require Player O to pick $y_i = y_0$ for all $i > 0$ and to mark the two positions i and i' inducing the bad y_0 -pair by some special markers \rightarrow and \leftarrow . Then, the automaton just has to check that x_j and y_j are equal at the marked positions and that x_j is strictly smaller than y_j in between these positions. Due to the binary encoding, both checks are easily implemented using the transition structure.

It remains to explain how to require Player O to copy her choice y_0 . As before, using the state space requires too many states. Instead, we employ the finitary Streett condition with respect to a small bound to enforce the copying. The j -th bit of a binary encoding of a number opens a request that can only be answered by encountering the same bit at the same position of a later encoding. Thus, to answer these requests with cost n , the numbers have to be copied.

To simplify our notation, we say that a delay function f eventually grants a lookahead of size m , if there is an i such that $\sum_{0 \leq i' \leq i} (f(i') - 1) \geq m$. Furthermore, we say a Streett pair in a Streett automaton is qualitative (finitary), if the associated cost function assigns $\varepsilon(i)$ to every transition.

Theorem 8. *For every $n > 0$, there is a language L_n recognized by a Streett automaton with costs \mathcal{A}_n of polynomial size in n such that*

- Player O has a winning strategy τ_O for $\Gamma_f(L_n)$ for some f with $\text{Cst}(\tau_O) = n$, but
- if Player O has a winning strategy τ_O for $\Gamma_f(L_n)$ with $\text{Cst}(\tau_O) = n$, then f eventually grants a lookahead of size $n \cdot (2^{2^n} - 1)$.

Proof. We start by describing the language L_n and by arguing that it can be recognized by a Streett automaton with costs \mathcal{A}_n of polynomial size in n . Fix $\Sigma_I = \{0, 1\} \cup \{0, 1\} \times \{\#\}$ and $\Sigma_O = \{0, 1\} \cup \{0, 1\} \times \{\rightarrow, \leftarrow\}$ and consider a




Fig. 4. The automaton \mathcal{A}_n for the proof of Theorem 5. The path from the gray state to the doubly-lined state has $n-1$ edges and $*$ denotes an arbitrary letter.




Fig. 5. The automaton \mathcal{A}_3 for the proof of Theorem 7. A transition labeled by $(*)^2$ represents a path of two transitions, each labeled with $(*)$, where $*$ denotes an arbitrary letter. The missing transitions of the gray states lead to a sink state of color 1.

word $(\alpha_\beta) \in (\Sigma_I \times \Sigma_O)^\omega$. By grouping the bits of α (ignoring the mark $\#$) into blocks of length n , α can be interpreted as a sequence $\bar{\alpha} = x_0x_1x_2 \dots \in \{0, \dots, 2^n - 1\}^\omega$ of natural numbers. Analogously, the bits of β can be interpreted as a sequence $\bar{\beta} = y_0y_1y_2 \dots \in \{0, \dots, 2^n - 1\}^\omega$ when ignoring the marks \rightarrow and \leftarrow . We say that x_i is marked, if $\#$ holds at the first position of the block encoding x_i , that y_i is marked by \rightarrow , if \rightarrow holds at the first position of the block encoding y_i , and y_i being marked by \leftarrow is defined analogously. Player I uses his mark to start a new round while Player O uses her marks to pick bad j -pairs.

Fix a word $w = (\alpha_\beta)$ with $\bar{\alpha} = x_0x_1x_2 \dots$ and $\bar{\beta} = y_0y_1y_2 \dots$. If x_0 is not marked, then $w \in L_n$. Thus, assume from now on that x_0 is marked and let $i_\#$ be arbitrary with $x_{i_\#}$ being marked. To be in L_n , w has to satisfy the following condition (amongst others): there have to be exactly two marked y_i with $i \geq i_\#$ before the next $x_{i'_\#}$ is marked. The first one after $i_\#$ has to be marked by \rightarrow , the second one by \leftarrow . However, if there is another marked $x_{i'_\#}$ before the y_i marked by \leftarrow then w is in L_n , i.e., Player I may only start a new round after Player O has picked a bad j -pair in the current round. These properties can be implemented using the transition structure of the automaton and a classical Streett pair to require Player O to use her marks eventually. Also, whenever a y_i is marked, it has to be equal to x_i . Furthermore, let y_i be marked by \rightarrow and let $y_{i'}$ be the next marked number. Then, we require $x_j < y_j$ for every j in the range $i < j < i'$. These requirements can be enforced by the transition structure.

Finally, we employ finitary Streett pairs and a small bound on the cost to enforce the copying. Every occurrence of a bit b (ignoring the marks) in β at a position k opens a request that is only answered by a later b (again ignoring the marks) in β at a

position k' with $k \bmod n = k' \bmod n$ or by a later block that is marked by \leftarrow . This property is enforced by finitary Streett pairs. Thus, to answer these requests with cost n , all the y_j between a marked x_i (the start of a round) and the next $y_{i'}$ marked by \leftarrow (the end of the round) have to coincide.

It is straightforward to construct a Streett automaton with costs \mathcal{A}_n of polynomial size that recognizes the language L_n described above.

Now, analogously to the arguments in Example 2, one can show that Player O wins $\Gamma_f(L_n)$ for the constant f with $f(0) = n \cdot 2^{2^n}$: at the start of each round, she has enough lookahead to pick a y_i such that Player I has already produced a bad y_i -pair in the current round. Then, she copies the y_i and marks the pair correctly and waits for the start of the next round. This satisfies the qualitative Streett pairs as well as the finitary ones with cost n , i.e., she wins.

Now fix some delay function f such that Player O has a winning strategy τ_O for $\Gamma_f(L_n)$ with $\text{Cst}(\tau_O) = n$. Assume towards a contradiction that f does not eventually grant a lookahead of size $n \cdot (2^{2^n} - 1)$. As mentioned in Example 2, there is a sequence \bar{x} of length $2^{2^n} - 1$ without a bad j -pair for every $j \in \{0, \dots, 2^n - 1\}$. Player I 's strategy against τ_O is to first play the binary encoding of \bar{x} , with the first number marked. Due to the small lookahead, Player O has to specify y_0 during these moves. Then, Player I just plays some $x \neq y_0$ until Player O has played both her marks. If she never does, then Player I wins the play. Otherwise, he just starts a new round and proceeds as previously described.

Consider a round of an outcome of this strategy and say Player O picks y_i as first number in this round. Then, the sequence of numbers picked by Player I in this round contains no bad y_i -pair. If Player O does not mark any numbers in this round, or they do not constitute a bad j -pair for some j , then

she loses the play, which contradicts our assumption. Thus, assume she does mark a bad j -pair correctly. Then, we have $j \neq y_0$. This means she does not copy y_i throughout the round until she plays her second mark. Thus, there is a request that is not answered with cost n . As Player I is able to enforce such a request in each round, the resulting play has at least cost $n + 1$, again a contradiction. \square

On the other hand, the previous example does not yield a doubly-exponential lower bound on the necessary lookahead for arbitrary bounds on the costs, as Player O can satisfy the acceptance condition of the automaton by starting each round by playing 0^n and then 1^n , which answers all finitary Streett pairs in the game and then correctly marking a bad j -pair for some j . This strategy has much larger cost than n , but is still winning. Whether there is a tradeoff (and, if yes, its extent) remains an open problem.

Also note that the automaton \mathcal{A}_n has both classical and finitary Streett pairs. It is an open problem to show the same result for finitary Streett automata. The problem one encounters is that the acceptance condition has to force Player O to mark a bad j -pair in each round. Implementing this with a finitary Streett pair increases the cost of a winning strategy, as it may take doubly-exponentially long before such a pair appears. This large bound allows Player O to cheat in the copying process, as described above.

VII. CONCLUSION

We have demonstrated the usefulness of adding delay to games with quantitative winning conditions, here finitary parity and Streett conditions as well as parity and Streett conditions with costs.

We have shown that delay games with parity conditions with costs are just as hard as delay games with parity conditions, both in terms of the necessary lookahead and in terms of the computational complexity of determining the winner. Thus, adding quantitative features to such games comes for free, which is in line with similar results for both delay-free games [13], [11], [21] and delay games [8]. Furthermore, we exhibited the usefulness of delay by showing that lookahead can be traded for quality of strategies. This phenomenon goes beyond the advantages in qualitative delay games, where lookahead only allows to win more games.

Another interesting property of delay-free finitary parity games is that playing them optimally is much harder than just winning them: the bounding player always has a positional winning strategy [13], which is winning with respect to some uniform bound b , but satisfying the optimal uniform bound might require exponential memory [15]. Similarly, determining the optimal bound is **PSPACE**-complete [15] while just determining the winner of a delay-free finitary parity game is in **PTIME** [12].

In current work, we study the tradeoffs between quality, memory requirements, lookahead, and solution complexity. In particular, this requires to develop a theory of finite-state strategies for delay games, which is, due to the presence of

lookahead, non-trivial (see [22] for a proposal of finite-state strategies in a setting that is similar to the definition of $\mathcal{G}(\mathcal{A})$).

Furthermore, we gave a doubly-exponential upper bound on the necessary lookahead for delay games with Streett conditions with costs and showed that such games can be solved in doubly-exponential time. The best lower bounds are those for parity conditions with costs, i.e., there is an exponential gap in both cases. Note that these gaps already exists in the case of qualitative Streett conditions: doubly-exponential constant lookahead is sufficient and solving such games is in **2EXPTIME** (this follows via determinization from the results for parity conditions), but the best lower bounds are exponential for the lookahead and **EXPTIME**-completeness. In future work, we aim to close these gaps.

Acknowledgments: We thank Alexander Weinert for numerous fruitful discussions.

REFERENCES

- [1] J. R. Büchi and L. H. Landweber, "Solving sequential conditions by finite-state strategies," *Transactions of the American Mathematical Society*, vol. 138, pp. 295–311, 1969.
- [2] A. Church, "Logic, arithmetic, and automata," in *Proc. Int. Congr. Math. 1962*. Inst. Mittag-Leffler, Djursholm, Sweden, 1963, pp. 23–35.
- [3] F. A. Hosch and L. H. Landweber, "Finite delay solutions for sequential conditions," in *ICALP 1972*, 1972, pp. 45–60.
- [4] M. Holtmann, L. Kaiser, and W. Thomas, "Degrees of lookahead in regular infinite games," *LMCS*, vol. 8, no. 3, 2012.
- [5] W. Fridman, C. Löding, and M. Zimmermann, "Degrees of lookahead in context-free infinite games," in *CSL 2011*, ser. LIPIcs, M. Bezem, Ed., vol. 12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2011, pp. 264–276.
- [6] F. Klein and M. Zimmermann, "What are strategies in delay games? Borel determinacy for games with lookahead," in *CSL 2015*, ser. LIPIcs, S. Kreutzer, Ed., vol. 41. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015, pp. 519–533.
- [7] —, "How much lookahead is needed to win infinite games?" *LMCS*, vol. 12, no. 3, 2016.
- [8] —, "Prompt delay," in *FSTTCS 2016*, ser. LIPIcs, A. Lal, S. Akshay, S. Saurabh, and S. Sen, Eds., vol. 65. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2016, pp. 43:1–43:14.
- [9] M. Zimmermann, "Delay games with WMSO+U winning conditions," *RAIRO-Theor. Inf. Appl.*, vol. 50, no. 2, pp. 145–165, 2016.
- [10] M. Jurdziński, "Deciding the winner in parity games is in **UP** \cap **coUP**," *Inf. Process. Lett.*, vol. 68, no. 3, pp. 119–124, 1998.
- [11] O. Kupferman, N. Piterman, and M. Y. Vardi, "From liveness to promptness," *Form. Method. Syst. Des.*, vol. 34, no. 2, pp. 83–103, 2009.
- [12] K. Chatterjee, T. A. Henzinger, and F. Horn, "Finitary winning in ω -regular games," *ACM Trans. Comput. Log.*, vol. 11, no. 1, 2009.
- [13] N. Fijalkow and M. Zimmermann, "Parity and Streett games with costs," *LMCS*, vol. 10, no. 2, 2014.
- [14] F. Mogavero, A. Murano, and L. Sorrentino, "On promptness in parity games," *Fundam. Inform.*, vol. 139, no. 3, pp. 277–305, 2015.
- [15] A. Weinert and M. Zimmermann, "Easy to win, hard to master: Optimal strategies in parity games with costs," in *CSL 2016*, ser. LIPIcs, J. Talbot and L. Regnier, Eds., vol. 62. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016, pp. 31:1–31:17.
- [16] M. Bojańczyk, "Weak MSO with the unbounding quantifier," *Theory Comput. Syst.*, vol. 48, no. 3, pp. 554–576, 2011.
- [17] A. Pnueli and R. Rosner, "On the synthesis of an asynchronous reactive module," in *ICALP 1989*, ser. LNCS, G. Ausiello, M. Dezaniciancagnini, and S. R. D. Rocca, Eds., vol. 372. Springer, 1989, pp. 652–671.
- [18] M. Zimmermann, "Games with costs and delays," *arXiv*, vol. 1701.02168, 2017.
- [19] K. Chatterjee and N. Fijalkow, "Finitary languages," in *LATA 2011*, ser. LNCS, A. H. Dediu, S. Inenaga, and C. Martín-Vide, Eds., vol. 6638. Springer, 2011, pp. 216–226.

- [20] A. Weinert and M. Zimmermann, “Easy to win, hard to master: Optimal strategies in parity games with costs,” 2017, Journal version of [15]. Under submission.
- [21] M. Zimmermann, “Optimal bounds in parametric LTL games,” *Theoret. Comput. Sci.*, vol. 493, no. 0, pp. 30 – 45, 2013.
- [22] T. Salzmänn, “How much memory is needed to win regular delay games?” Master’s thesis, Saarland University, 2015.