

Finite-state Strategies in Delay Games*

Martin Zimmermann

Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany

`zimmermann@react.uni-saarland.de`

What is a finite-state strategy in a delay game? We answer this surprisingly non-trivial question and present a very general framework for computing such strategies: they exist for all winning conditions that are recognized by automata with acceptance conditions that satisfy a certain aggregation property. Our framework also yields upper bounds on the complexity of determining the winner of such delay games and upper bounds on the necessary lookahead to win the game. In particular, we cover all previous results of that kind as special cases of our uniform approach.

1 Introduction

What is a finite-state strategy in a delay game? The answer to this question is surprisingly non-trivial due to the nature of delay games in which one player is granted a lookahead on her opponent's moves. This puts her into an advantage when it comes to winning games, i.e., there are games that can only be won with lookahead, but not without. A simple example is a game where one has to predict the third move of the opponent with one's first move. This is impossible when moving in alternation, but possible if one has access to the opponent's first three moves before making the first move. More intriguingly, lookahead also allows Player O to improve the quality of her winning strategies in games with quantitative winning conditions, i.e., there is a tradeoff between quality and amount of lookahead [34].

However, managing (and, if necessary, storing) the additional information gained by the lookahead can be a burden. Consider another game where one just has to copy the opponent's moves. This is obviously possible with or without lookahead (assuming the opponent moves first). In particular, without lookahead one just has to remember the last move of the opponent and copy it. However, when granted lookahead, one has to store the last moves of the opponent in a queue to implement the copying properly. This example shows that lookahead is not necessarily advantageous when it comes to minimizing the memory requirements of a strategy.

In this work, we are concerned with Gale-Stewart games [11], abstract games without an underlying arena.¹ In such a game, both players produce an infinite sequence of letters and the winner is determined by the combination of these sequences. If it is in the winning condition, a set of such combinations, then the second player wins, otherwise the first one wins. In a classical Gale-Stewart game, both players move in alternation while in a delay game, the second player skips moves to obtain a lookahead on the opponent's moves. Which moves are skipped is part of the rules of the game and known to both players.

Delay games have recently received a considerable amount of attention after being introduced by Hosch and Landweber [15] only three years after the seminal Büchi-Landweber theorem [2]. Büchi and Landweber had shown how to solve infinite two-player games with ω -regular winning conditions.

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

¹The models of Gale-Stewart games and arena-based games are interreducible, but delay games are naturally presented as a generalization of Gale-Stewart games. This is the reason we prefer this model here.

Forty years later, delay games were revisited by Holtmann, Kaiser, and Thomas [13] and the first comprehensive study was initiated, which settled many basic problems like the exact complexity of solving ω -regular delay games and the amount of lookahead necessary to win such games [19]. Furthermore, Martin's seminal Borel determinacy theorem [21] for Gale-Stewart games has been lifted to delay games [18] and winning conditions beyond the ω -regular ones have been investigated [10, 20, 32, 34]. Finally, the uniformization problem for relations over infinite words boils down to solving delay games: a relation $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ is uniformized by a continuous function (in the Cantor topology) if, and only if, the delaying player wins the delay game with winning condition L . We refer to [13] for details.

What makes finite-state strategies in infinite games particularly useful and desirable is that a general strategy is an infinite object, as it maps finite play prefixes to next moves. On the other hand, a finite-state strategy is implemented by a transducer, an automaton with output, and therefore finitely represented: the automaton reads a play prefix and outputs the next move to be taken. Thus, the transducer computes a finite abstraction of the play's history using its state space as memory and determines the next move based on the current memory state.

In Gale-Stewart games, finite-state strategies suffice for all ω -regular games [2] and even for deterministic ω -contextfree games, if one allows pushdown transducers [31]. For Gale-Stewart games (and arena-based games), the notion is well-established and one of the most basic questions about a class of winning conditions is that about the existence and size of winning strategies for such games.

While foundational questions for delay games have been answered and many results have been lifted from Gale-Stewart games to those with delay, the issue of computing tractable and implementable strategies has not been addressed before. However, this problem is of great importance, as the existence and computability of finite-state strategies is a major reason for the successful application of infinite games to diverse problems like reactive synthesis, model-checking of fixed-point logics, and automata theory.

In previous work, restricted classes of strategies for delay games have been considered [18]. However, those restrictions are concerned with the amount of information about the lookahead's evolution a strategy has access to, and do not restrict the size of the strategies: In general, they are still infinite objects. On the other hand, it is known that bounded lookahead suffices for many winning conditions of importance, e.g., the ω -regular ones [19], those recognized by parity and Streett automata with costs [34], and those definable in (parameterized) linear temporal logics [20]. Furthermore, for all those winning conditions, the winner of a delay game can be determined effectively. In fact, all these proofs rely on the same basic construction that was already present in the work of Holtmann, Kaiser, and Thomas [13], i.e., a reduction to a Gale-Stewart game using equivalence relations that capture the behavior of the automaton recognizing the winning condition. These reductions and the fact that finite-state strategies suffice for the games obtained in the reductions imply that (some kind of) finite-state strategies exist.

Indeed, in his master's thesis [28], Salzmann recently introduced the first notion of finite-state strategies in delay games and, using these reductions, presented an algorithm computing them for several types of acceptance conditions, e.g., parity conditions and related ω -regular ones. However, the exact nature of finite-state strategies in delay games is not as canonical as for Gale-Stewart games. We discuss this issue in-depth in Sections 3 and 5 by proposing two notions of finite-state strategies, a delay-oblivious one which yields large strategies in the size of the lookahead, and a delay-aware one that follows naturally from the reductions to Gale-Stewart games mentioned earlier. In particular, the number of states of the delay-aware strategies is independent of the size of the lookahead, but often larger in the size of the automaton recognizing the winning condition. However, this is offset by the fact that strategies of the second type are simpler to compute than the delay-oblivious ones and have overall fewer states, if the lookahead is large. In comparison to Salzmann's notion, where strategies syntactically depend on a given automaton representing the winning condition, our strategies are independent of the representation

of the winning condition and therefore more general. Also, our framework is more abstract and therefore applicable to a wider range of acceptance conditions (e.g., qualitative ones) and yields in general smaller strategies, but there are of course some similarities, which we discuss in detail.

To present these notions, we first introduce some definitions in Section 2, e.g., delay games and finite-state strategies for Gale-Stewart games. After introducing the two notions of finite-state strategies for delay games in Section 3, we show how to compute such strategies in Section 4. To this end, we present a generic account of the reduction from delay games to Gale-Stewart games which subsumes, to the best of our knowledge, all decidability results presented in the literature. Furthermore, we show how to obtain the desired strategies from our construction. Then, in Section 5, we compare the two different definitions of finite-state strategies for delay games proposed here and discuss their advantages and disadvantages. Also, we compare our approach to that of Salzmann. We conclude by mentioning some directions for further research in Section 6.

Proofs and constructions omitted due to space restrictions can be found in the full version [33].

Related Work As mentioned earlier, the existence of finite-state strategies is the technical core of many applications of infinite games, e.g., in reactive synthesis one synthesizes a correct-by-construction system from a given specification by casting the problem as an infinite game between a player representing the system and one representing the antagonistic environment. It is a winning strategy for the system player that yields the desired implementation, which is finite if the winning strategy is finite-state. Similarly, Gurevich and Harrington’s game-based proof of Rabin’s decidability theorem for monadic second-order logic over infinite binary trees [25] relies on the existence of finite-state strategies.²

These facts explain the need for studying the existence and properties of finite-state strategies in infinite games [17, 26, 27, 29]. In particular, the seminal work by Dziembowski, Jurdziński, and Walukiewicz [6] addressed the problem of determining upper and lower bounds on the size of finite-state winning strategies in games with Muller winning conditions. Nowadays, one of the most basic questions about a given winning condition is that about such upper and lower bounds. For most conditions in the literature, tight bounds are known, see, e.g., [4, 14, 30]. But there are also surprising exceptions to that rule, e.g., generalized reachability games [9]. More recently, Colcombet, Fijalkow, and Horn presented a very general technique that yields tight upper and lower bounds on memory requirements in safety games, which even hold for games in infinite arenas, provided their degree is finite [5].

2 Preliminaries

We denote the non-negative integers by \mathbb{N} . Given two ω -words $\alpha \in (\Sigma_0)^\omega$ and $\beta \in (\Sigma_1)^\omega$, we define $\binom{\alpha}{\beta} = \binom{\alpha^{(0)}}{\beta^{(0)}} \binom{\alpha^{(1)}}{\beta^{(1)}} \binom{\alpha^{(2)}}{\beta^{(2)}} \cdots \in (\Sigma_0 \times \Sigma_1)^\omega$. Similarly, we define $\binom{x}{y}$ for finite words x and y with $|x| = |y|$.

ω -automata A (deterministic and complete) ω -automaton is a tuple $\mathfrak{A} = (Q, \Sigma, q_I, \delta, \text{Acc})$ where Q is a finite set of states, Σ is an alphabet, $q_I \in Q$ is the initial state, $\delta: Q \times \Sigma \rightarrow Q$ is the transition function, and $\text{Acc} \subseteq \delta^\omega$ is the set of accepting runs (here, and whenever convenient, we treat δ as a relation $\delta \subseteq Q \times \Sigma \times Q$). A finite run π of \mathfrak{A} is a sequence $\pi = (q_0, a_0, q_1)(q_1, a_1, q_2) \cdots (q_{i-2}, a_{i-2}, q_{i-1})(q_{i-1}, a_{i-1}, q_i) \in \delta^+$. As usual, we say that π starts in q_0 , ends in q_i , and processes $a_0 \cdots a_{i-1} \in \Sigma^+$. Infinite runs on infinite words are defined analogously. If we speak of *the* run of \mathfrak{A} on $\alpha \in \Sigma^\omega$, then we mean the unique run of

²The proof is actually based on positional strategies, a further restriction of finite-state strategies for arena-based games, because they are simpler to handle. Nevertheless, the same proof also works for finite-state strategies.

\mathfrak{A} starting in q_I processing α . The language $L(\mathfrak{A}) \subseteq \Sigma^\omega$ of \mathfrak{A} contains all those ω -words whose run of \mathfrak{A} is accepting. The size of \mathfrak{A} is defined as $|\mathfrak{A}| = |Q|$.

This definition is very broad, which allows us to formulate our theorems as general as possible. In examples, we consider parity and Muller automata whose set of accepting runs is finitely represented: An ω -automaton $\mathfrak{A} = (Q, \Sigma, q_I, \delta, \text{Acc})$ is a parity automaton, if $\text{Acc} = \{(q_0, a_0, q_1)(q_1, a_1, q_2)(q_2, a_2, q_3) \cdots \in \delta^\omega \mid \limsup_{i \rightarrow \infty} \Omega(q_i) \text{ is even}\}$ for some coloring $\Omega: Q \rightarrow \mathbb{N}$. To simplify our notation, define $\Omega(q, a, q') = \Omega(q)$. Furthermore, \mathfrak{A} is a Muller automaton, if there is a family $\mathcal{F} \subseteq 2^Q$ of sets of states such that $\text{Acc} = \{\rho \in \delta^\omega \mid \text{Inf}(\rho) \in \mathcal{F}\}$, where $\text{Inf}(\rho)$ is the set of states visited infinitely often by ρ .

Delay Games A delay function is a mapping $f: \mathbb{N} \rightarrow \mathbb{N} \setminus \{0\}$, which is said to be constant if $f(i) = 1$ for all $i > 0$. A delay game $\Gamma_f(L)$ consists of a delay function f and a winning condition $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ for some alphabets Σ_I and Σ_O . Such a game is played in rounds $i = 0, 1, 2, \dots$ as follows: in round i , first Player I picks a word $x_i \in \Sigma_I^{f(i)}$, then Player O picks a letter $y_i \in \Sigma_O$. Player O wins a play $(x_0, y_0)(x_1, y_1)(x_2, y_2) \cdots$ if the outcome $\binom{x_0 x_1 x_2 \cdots}{y_0 y_1 y_2 \cdots}$ is in L ; otherwise, Player I wins.

A strategy for Player I in $\Gamma_f(L)$ is a mapping $\tau_I: \Sigma_O^* \rightarrow \Sigma_I^*$ satisfying $|\tau_I(w)| = f(|w|)$ while a strategy for Player O is a mapping $\tau_O: \Sigma_I^+ \rightarrow \Sigma_O$. A play $(x_0, y_0)(x_1, y_1)(x_2, y_2) \cdots$ is consistent with τ_I if $x_i = \tau_I(y_0 \cdots y_{i-1})$ for all i , and it is consistent with τ_O if $y_i = \tau_O(x_0 \cdots x_i)$ for all i . A strategy for Player $P \in \{I, O\}$ is winning, if every play that is consistent with the strategy is won by Player P .

An important special case are delay-free games, i.e., those with respect to the delay function f mapping every i to 1. In this case, we drop the subscript f and write $\Gamma(L)$ for the game with winning condition L . Such games are typically called Gale-Stewart games [11].

Finite-state Strategies in Gale-Stewart Games A strategy for Player O in a Gale-Stewart game is still a mapping $\tau_O: \Sigma_I^+ \rightarrow \Sigma_O$. Such a strategy is said to be finite-state, if there is a deterministic finite transducer \mathfrak{T} that implements τ_O in the following sense: \mathfrak{T} is a tuple $(Q, \Sigma_I, q_I, \delta, \Sigma_O, \lambda)$ where Q is a finite set of states, Σ_I is the input alphabet, $q_I \in Q$ is the initial state, $\delta: Q \times \Sigma_I \rightarrow Q$ is the deterministic transition function, Σ_O is the output alphabet, and $\lambda: Q \rightarrow \Sigma_O$ is the output function. Let $\delta^*(x)$ denote the unique state that is reached by \mathfrak{T} when processing $x \in \Sigma_I^*$ from q_I . Then, the strategy $\tau_{\mathfrak{T}}$ implemented by \mathfrak{T} is defined as $\tau_{\mathfrak{T}}(x) = \lambda(\delta^*(x))$. We say that a strategy is finite-state, if it is implementable by some transducer. Slightly abusively, we identify finite-state strategies with transducers implementing them and talk about finite-state strategies with some number of states. Thus, we focus on the *state complexity* (e.g., the number of memory states necessary to implement a strategy) and ignore the other components of a transducer (which are anyway of polynomial size in $|Q|$, if we assume Σ_I and Σ_O to be fixed).

3 What is a Finite-state Strategy in a Delay Game?

Before we answer this question, we first ask what properties a finite-state strategy should have, i.e., what makes finite-state strategies in Gale-Stewart games useful and desirable? A strategy $\tau_O: \Sigma_I^+ \rightarrow \Sigma_O$ is in general an infinite object and does not necessarily have a finite representation. Furthermore, to execute such a strategy, one needs to store the whole sequence of moves made by Player I thus far: Unbounded memory is needed to execute it.

On the other hand, a finite-state strategy is finitely described by a transducer \mathfrak{T} implementing it. To execute it, one only needs to store a single state of \mathfrak{T} and needs to have access to the transition function δ and the output function λ of \mathfrak{T} . Assume the current state is q at the beginning of some round i (initialized with q_I before round 0). Then, Player I makes his move by picking some $a \in \Sigma_I$, which is processed

by updating the memory state to $q' = \delta(q, a)$. Then, \mathfrak{T} prescribes picking $\lambda(q') \in \Sigma_O$ and round i is completed. Thus, there are two aspects that make finite-state strategies desirable: (1) the next move depends only on a finite amount of information about the history of the play, i.e., a state of the automaton, which is (2) easily updated. In particular, the strategy is completely specified by the transition function and the output function.

Further, there is a generic framework to compute such strategies by reducing them to arena-based games (see, e.g., [12] for an introduction to such games). As an example, consider a game $\Gamma(L(\mathfrak{A}))$ where \mathfrak{A} is a parity automaton with set Q of states and transition function δ . We describe the construction of an arena-based parity game contested between Player I and Player O whose solution allows us to compute the desired strategies (formal details are presented in the full version [33]). The positions of Player I are transitions of \mathfrak{A} while those of Player O are pairs (q, a) where $q \in Q$ and where a is an input letter. From a vertex $(q, \binom{a}{b}, q')$ Player I can move to every vertex (q', a') for $a' \in \Sigma_I$, from which Player O can move to every vertex $(q', \binom{a'}{b'}, \delta(q', \binom{a'}{b'}))$ for $b' \in \Sigma_O$. Finally, Player O wins a play, if the run constructed during the infinite play is accepting. It is easy to see that the resulting game is a parity game with $|\delta| \cdot (|\Sigma_I| + 1)$ vertices, and has the same winner as $\Gamma(L(\mathfrak{A}))$. The winner of the arena-based game has a positional³ winning strategy [7, 23], which can be computed in quasipolynomial time [3, 8, 16]. Such a positional winning strategy can easily be turned into a finite-state winning strategy with $|Q| \cdot |\Sigma_I|$ states for Player O in the game $\Gamma(L(\mathfrak{A}))$, which is implemented by an automaton with state set $Q \times \Sigma_I$. This reduction can be generalized to arbitrary classes of Gale-Stewart games whose winning condition is recognized by an ω -automaton with set Q of states: if Player O has a finite-state strategy with n states in the arena-based game obtained by the construction described above, then Player O has a finite-state winning strategy with $|Q| \cdot |\Sigma_I| \cdot n$ states for the original Gale-Stewart game. Such a strategy is obtained by solving an arena-based game with $|\delta| \cdot (|\Sigma_I| + 1)$ vertices.

So, what is a finite-state strategy in a delay game? In the following, we discuss this question for the case of delay games with respect to constant delay functions, which is the most important case. In particular, constant lookahead suffices for all ω -regular winning conditions [19], i.e., Player O wins with respect to an arbitrary delay function if, and only if, she wins with respect to a constant one. Similarly, constant lookahead suffices for many quantitative conditions like (parameterized) temporal logics [20] and parity conditions with costs [34]. For winning conditions given by parity automata, there is an exponential upper bound on the necessary constant lookahead. On the other hand, there are exponential lower bounds already for winning conditions specified by deterministic automata with reachability or safety acceptance (which are subsumed by parity acceptance).

3.1 Delay-oblivious Finite-state Strategies for Delay Games

Technically, a strategy for Player O in a delay game is still a mapping $\tau_O: \Sigma_I^+ \rightarrow \Sigma_O$. Hence, the definition of finite-state strategies via transducers as given above for Gale-Stewart games is also applicable to delay games. As a (cautionary) example, consider a delay game with winning condition $L_{=} = \{ \binom{\alpha}{\alpha} \mid \alpha \in \{0, 1\}^\omega \}$, i.e., Player O just has to copy Player I 's moves, which she can do with respect to every delay function: Player O wins $\Gamma_f(L_{=})$ for every f . However, a finite-state strategy has to remember the whole lookahead, i.e., those moves that Player I is ahead of Player O , in order to copy his moves. Thus, an automaton implementing a winning strategy for Player O in $\Gamma_f(L_{=})$ needs at least $|\{0, 1\}|^d$ states, if f is a constant delay function with $f(0) = d$. Thus, the memory requirements grow with the size of the

³A strategy in an arena-based games is positional, if its output only depends on the last vertex of the play's history, not on the full history (see, e.g., [12]).

lookahead granted to Player O , i.e., lookahead is a burden, not an advantage. She even needs unbounded memory in the case of unbounded lookahead.

On the other hand, an advantage of this “delay-oblivious” definition is that finite-state strategies can be obtained by a trivial extension of the reduction presented for Gale-Stewart games above: now, states of Player I are from $\delta \times \Sigma_I^{d-1}$ and those of Player O are from $Q \times \Sigma_I^d$. Player I can move from $((q, \binom{a}{b}, q'), w)$ to (q', wa') for $a' \in \Sigma_I$ while Player O can move from (q, aw) to $((q, \binom{a}{b}, \delta(q, \binom{a}{b})), w)$ for $b \in \Sigma_O$. Intuitively, a state now additionally stores a queue of length $d - 1$, which contains the lookahead granted to Player O . Coming back to the parity example, this approach yields a finite-state strategy with $|Q| \cdot |\Sigma_I|^d$ states. To obtain such a strategy, one has to solve a parity game with $|\delta| \cdot (|\Sigma_I| + 1) \cdot |\Sigma_I|^{d-1}$ vertices, which is of doubly-exponential size in $|\mathfrak{A}|$, if d is close to the (tight) exponential upper bound. This can be done in doubly-exponential time, as it still has the same number of colors as the automaton \mathfrak{A} . Again, this reduction can be generalized to arbitrary classes of delay games with constant delay whose winning conditions are recognized by an ω -automaton with set Q of states: if Player O has a finite-state strategy with n states in the arena-based game obtained by the construction, then Player O has a finite-state winning strategy with $|Q| \cdot |\Sigma_I|^d \cdot n$ states for the delay game with constant lookahead of size d . In general, d factors exponentially into the size, as n is the memory size required to win a game with $\mathcal{O}(|\Sigma_I|^d)$ vertices. Also, to obtain the strategy for the delay game, one has to solve an arena-based game with $|\delta| \cdot (|\Sigma_I| + 1) \cdot |\Sigma_I|^{d-1}$ vertices.

3.2 Block Games

We show that one can do better than by decoupling the history tracking and the handling of the lookahead, i.e., by using *delay-aware* finite-state strategies. In the delay-oblivious definition, we hardcode a queue into the arena-based game, which results in a blowup of the arena and therefore also in a blowup in the solution complexity and in the number of memory states for the arena-based game, which is turned into one for the delay game. To overcome this, we introduce a slight variation of delay games with respect to constant delay functions, so-called block games⁴, present a notion of finite-state strategy in block games, and show how to transfer strategies between delay games and block games. Then, we show how to solve block games and how to obtain finite-state strategies for them.

The motivation for introducing block games is to eliminate the queue containing the letters Player I is ahead of Player O , which is cumbersome to maintain, and causes the blowup in the case of games with winning condition $L_{=}$. Instead, in a block game, both players pick blocks of letters of a fixed length with Player I being one block ahead to account for the delay, i.e., Player I has to pick two blocks in round 0 and then one in every round, as does Player O in every round. This variant of delay games lies implicitly or explicitly at the foundation of all arguments establishing upper bounds on the necessary lookahead and at the foundation of all algorithms solving delay games [13, 19, 20, 32, 34]. Furthermore, we show how to transform a (winning) strategy for a delay game into a (winning) strategy for a block game and vice versa, i.e., Player O wins the delay game if, and only if, she wins the corresponding block game.⁵

Formally, the block game $\Gamma^d(L)$, where $d \in \mathbb{N} \setminus \{0\}$ is the block length and where $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$ is the winning condition, is played in rounds as follows: in round 0, Player I picks two blocks $\overline{a_0}, \overline{a_1} \in \Sigma_I^d$, then Player O picks a block $\overline{b_0} \in \Sigma_O^d$. In round $i > 0$, Player I picks a block $\overline{a_{i+1}} \in \Sigma_I^d$, then Player O

⁴Holtmann, Kaiser, and Thomas already introduced a notion of block game in connection to delay games [13]. However, their notion differs from ours in several aspects. Most importantly, in their definition, Player I determines the length of the blocks (within some bounds specified by f) while our block length is fixed.

⁵Due to their prevalence and importance for solving delay games, one could even argue that the notion of block games is more suitable to model delay in infinite games.

picks a block $\bar{b}_i \in \Sigma_O^d$. Player O wins the resulting play $\overline{a_0 a_1 b_0 a_2 b_1} \dots$, if the outcome $\left(\frac{a_0 a_1 a_2 \dots}{b_0 b_1 b_2 \dots}\right)$ is in L .

A strategy for Player I in $\Gamma^d(L)$ is a map $\tau_I: (\Sigma_O^d)^* \rightarrow (\Sigma_I^d)^2 \cup \Sigma_I^d$ such that $\tau_I(\varepsilon) \in (\Sigma_I^d)^2$ and $\tau_I(\overline{b_0} \dots \overline{b_i}) \in \Sigma_I^d$ for $i \geq 0$. A strategy for Player O is a map $\tau_O: (\Sigma_I^d)^* \rightarrow \Sigma_O^d$. A play $\overline{a_0 a_1 b_0 a_2 b_1} \dots$ is consistent with τ_I , if $(\overline{a_0}, \overline{a_1}) = \tau_I(\varepsilon)$ and $\overline{a_i} = \tau_I(\overline{b_0} \dots \overline{b_{i-2}})$ for every $i \geq 2$; it is consistent with τ_O if $\overline{b_i} = \tau_O(\overline{a_0} \dots \overline{a_{i+1}})$ for every $i \geq 0$. Winning strategies and winning a block game are defined as for delay games.

In the following, we call strategies for block games delay-aware and strategies for delay games delay-oblivious. The next lemma relates delay games with constant lookahead and block games: for a given winning condition, Player O wins a delay game with winning condition L (with respect to some delay function) if, and only if, she wins a block game with winning condition L (for some block size).

Lemma 1. *Let $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$.*

1. *If Player O wins $\Gamma_f(L)$ for some constant delay function f , then she also wins $\Gamma^{f(0)}(L)$.*
2. *If Player O wins $\Gamma^d(L)$, then she also wins $\Gamma_f(L)$ for the constant delay function f with $f(0) = 2d$.*

Proof. 1.) Let $\tau_O: \Sigma_I^+ \rightarrow \Sigma_O$ be a winning strategy for Player O in $\Gamma_f(L)$ and fix $d = f(0)$. Now, define $\tau'_O: (\Sigma_I^d)^* \rightarrow (\Sigma_O^d)^d$ for Player O in $\Gamma^d(L)$ via $\tau'_O(\overline{a_0} \dots \overline{a_i a_{i+1}}) = \beta(0) \dots \beta(d-1)$ with $\beta(j) = \tau_O(\overline{a_0} \dots \overline{a_j} \alpha(0) \alpha(1) \dots \alpha(j-1))$ for $\overline{a_{i+1}} = \alpha(0) \alpha(1) \dots \alpha(d-1)$.

A straightforward induction shows that for every play consistent with τ'_O there is a play consistent with τ_O that has the same outcome. Thus, as τ_O is a winning strategy, so is τ'_O .

2.) Now, let $\tau'_O: (\Sigma_I^d)^* \rightarrow (\Sigma_O^d)^d$ be a winning strategy for Player O in $\Gamma^d(L)$. We define $\tau_O: \Sigma_I^+ \rightarrow \Sigma_O$ for Player O in $\Gamma_f(L)$. To this end, let $x \in \Sigma_I^+$ be a possible input occurring during a play. Hence, by the choice of f , we obtain $|x| \geq f(0) = 2d$. Thus, we can decompose x into $x = \overline{a_0} \dots \overline{a_i} x'$ such that $i \geq 1$, each $\overline{a_i}$ is a block over Σ_I of length d and $|x'| < d$. Now, let $\tau'_O(\overline{a_0} \dots \overline{a_i}) = \beta(0) \dots \beta(d-1)$. Then, we define $\tau_O(x) = \beta(|x'|)$.

Again, a straightforward induction shows that for every play consistent with τ_O there is a play consistent with τ'_O that has the same outcome. Thus, τ_O is a winning strategy. \square

3.3 Delay-aware Finite-state Strategies in Block Games

Now fix a block game $\Gamma^d(L)$ with $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$. A finite-state strategy for Player O in $\Gamma^d(L)$ is implemented by a transducer $\mathfrak{T} = (Q, \Sigma_I, q_I, \delta, \Sigma_O, \lambda)$ where Q , Σ_I , and q_I are defined as in Subsection 2.2. However, the transition function $\delta: Q \times \Sigma_I^d \rightarrow Q$ processes full input blocks and the output function $\lambda: Q \times \Sigma_I^d \times \Sigma_I^d \rightarrow \Sigma_O^d$ maps a state and a pair of input blocks to an output block. The strategy $\tau_{\mathfrak{T}}$ implemented by \mathfrak{T} is defined as $\tau_{\mathfrak{T}}(\overline{a_0} \dots \overline{a_i}) = \lambda(\delta^*(\overline{a_0} \dots \overline{a_{i-2}}, \overline{a_{i-1}}, \overline{a_i}))$ for $i \geq 1$.

Again, we identify delay-aware strategies with transducers implementing them and are interested in the number of states of the transducer. This definition captures the amount of information that is differentiated in order to implement the strategy. Note however, that it ignores the representation of the transition and the output function. These are no longer “small” (in $|Q|$), as it is the case for transducers implementing strategies for Gale-Stewart games. When focussing on executing such strategies, these factors become relevant, but for our purposes they are not: We have decoupled the history tracking from the lookahead-handling. The former is implemented by the automaton as usual while the latter is taken care of by the output function. In particular, the size of the automaton is (a-priori) independent of the block size. In the conclusion, we revisit the issue of presenting the transition and the output function.

In the next section, we present a very general approach to computing finite-state strategies for block games whose winning conditions are specified by automata with acceptance conditions that satisfy a

certain aggregation property. For example, for block games with winning conditions given by deterministic parity automata, we obtain a strategy implemented by a transducer with exponentially many states, which can be obtained by solving a parity game of exponential size. In both aspects, this is an exponential improvement over the delay-oblivious variant for classical delay games.

To conclude the introduction of block games, we strengthen Lemma 1 to transfer finite-state strategies between delay games and block games.

Lemma 2. *Let $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$.*

1. *If Player O has a delay-oblivious finite-state winning strategy for $\Gamma_f(L)$ with n states for some constant delay function f , then she also has a delay-aware finite-state winning strategy for $\Gamma^{f(0)}(L)$ with n states.*
2. *If Player O has a delay-aware finite-state winning strategy for $\Gamma^d(L)$ with n states, then she also has a delay-oblivious finite-state winning strategy for $\Gamma_f(L)$ with $n \cdot |\Sigma_I|^{2d}$ states for the constant delay function f with $f(0) = 2d$.*

Proof. It is straightforward to achieve the strategy transformations described in the proof of Lemma 1 by transforming transducers that implement finite-state strategies. \square

The blowup in the direction from block games to delay games is in general unavoidable, as finite-state winning strategies for the game $\Gamma_f(L_-)$ need at least 2^d states to store the lookahead while winning strategies for the block game need only one state, independently of the block size.

4 Computing Finite-state Strategies for Block Games

The aim of this section is twofold. Our main aim is to compute finite-state strategies for block games (and, by extension, for delay games with constant lookahead). We do so by presenting a general framework for analyzing delay games with winning conditions specified by ω -automata whose acceptance conditions satisfy a certain aggregation property. The technical core is a reduction to a Gale-Stewart game, i.e., we remove the delay from the game. This framework yields upper bounds on the necessary (constant) lookahead to win a given game, but also allows us to determine the winner and a finite-state winning strategy, if the resulting Gale-Stewart game can be effectively solved.

Slightly more formally, let \mathfrak{A} be the automaton recognizing the winning condition of the block game. Then, the winning condition of the Gale-Stewart game constructed in the reduction is recognized by an automaton \mathfrak{B} that can be derived from \mathfrak{A} . In particular, the acceptance condition of \mathfrak{B} simulates the acceptance condition of \mathfrak{A} . Many types of acceptance conditions are preserved by the simulation, e.g., starting with a parity automaton \mathfrak{A} , we end up with a parity automaton \mathfrak{B} . Thus, the resulting Gale-Stewart game can be effectively solved.

Our second aim is to present a framework as general as possible to obtain upper bounds on the necessary lookahead and on the solution complexity for a wide range of winning conditions. In fact, our framework is a generalization and abstraction of techniques first developed for the case of ω -regular winning conditions [19], which were later generalized to other winning conditions [20, 32, 34]. Here, we cover all these results in a uniform way.

4.1 Aggregations

Let us begin by giving some intuition for the construction. The winning condition of the game is recognized by an automaton \mathfrak{A} . Thus, as usual, the exact input can be abstracted away, only the induced

behavior in \mathfrak{A} is relevant. Such a behavior is characterized by the state transformations induced by processing the input and by the effect on the acceptance condition triggered by processing it. For many acceptance conditions, this effect can be aggregated, e.g., for parity conditions, one can decompose runs into non-empty pieces and then only consider the maximal colors of the pieces. For quantitative winning conditions, one typically needs an additional bound on the lengths of these pieces (cp. [32, 34]).

Thus, we begin by introducing two types of such aggregations of different strength. Fix an ω -automaton $\mathfrak{A} = (Q, \Sigma, q_I, \delta, \text{Acc})$ and let $s: \delta^+ \rightarrow M$ for some finite set M . Given a decomposition $(\pi_i)_{i \in \mathbb{N}}$ of a run $\pi_0 \pi_1 \pi_2 \cdots$ into non-empty pieces $\pi_i \in \delta^+$ we define $s((\pi_i)_{i \in \mathbb{N}}) = s(\pi_0) s(\pi_1) s(\pi_2) \cdots \in M^\omega$.

- We say that s is a strong aggregation (function) for \mathfrak{A} , if for all decompositions $(\pi_i)_{i \in \mathbb{N}}$ and $(\pi'_i)_{i \in \mathbb{N}}$ of runs $\rho = \pi_0 \pi_1 \pi_2 \cdots$ and $\rho' = \pi'_0 \pi'_1 \pi'_2 \cdots$ with $\sup_i |\pi'_i| < \infty$ and $s((\pi_i)_{i \in \mathbb{N}}) = s((\pi'_i)_{i \in \mathbb{N}})$: $\rho \in \text{Acc} \Rightarrow \rho' \in \text{Acc}$.
- We say that s is a weak aggregation (function) for \mathfrak{A} , if for all decompositions $(\pi_i)_{i \in \mathbb{N}}$ and $(\pi'_i)_{i \in \mathbb{N}}$ of runs $\rho = \pi_0 \pi_1 \pi_2 \cdots$ and $\rho' = \pi'_0 \pi'_1 \pi'_2 \cdots$ with $\sup_i |\pi_i| < \infty$, $\sup_i |\pi'_i| < \infty$, and $s((\pi_i)_{i \in \mathbb{N}}) = s((\pi'_i)_{i \in \mathbb{N}})$: $\rho \in \text{Acc} \Rightarrow \rho' \in \text{Acc}$.

Example 1.

- The function $s_{\text{prty}}: \delta^+ \rightarrow \Omega(Q)$ defined as $s_{\text{prty}}(t_0 \cdots t_i) = \max_{0 \leq j \leq i} \Omega(t_j)$ is a strong aggregation for a parity automaton $(Q, \Sigma, q_I, \delta, \text{Acc})$ with coloring Ω (recall that $\Omega(q, a, q') = \Omega(q)$).
- The function $s_{\text{mllr}}: \delta^+ \rightarrow 2^Q$ defined as $s_{\text{mllr}}((q_0, a_0, q_1) \cdots (q_n, a_n, q_{n+1})) = \{q_0, q_1, \dots, q_n\}$ is a strong aggregation for a Muller automaton $(Q, \Sigma, q_I, \delta, \text{Acc})$.
- The exponential time algorithm for delay games with winning conditions given by parity automata with costs, a quantitative generalization of parity automata, is based on a strong aggregation [34].
- The algorithm for delay games with winning conditions given by max automata [1], another quantitative automaton model, is based on a weak aggregation [32].

Due to symmetry, we can replace the implication $\rho \in \text{Acc} \Rightarrow \rho' \in \text{Acc}$ by an equivalence in the definition of a weak aggregation. Also, every strong aggregation is trivially a weak one as well.

Let us briefly comment on the difference between strong and weak aggregations using the examples of parity automata with costs and max-automata: the acceptance condition of the former automata is a boundedness condition on some counters while the acceptance condition of the latter is a boolean combination of boundedness and unboundedness conditions on some counters. The aggregations for these acceptance conditions capture whether a piece of a run induces an increment of a counter or not, but abstract away the actual number of increments if it is non-zero. Now, consider the parity condition with costs, which requires to bound the counters. Assume the counters in some run $\pi_0 \pi_1 \pi_2 \cdots$ are bounded and that we have pieces π'_i of bounded length having the same aggregation. Then, the increments in some piece π'_i have at least one corresponding increment in π_i . Thus, if a counter in $\pi'_0 \pi'_1 \pi'_2 \cdots$ is unbounded, then it is also unbounded in $\pi_0 \pi_1 \pi_2 \cdots$, which yields a contradiction. Hence, the implication $\pi_0 \pi_1 \pi_2 \cdots \in \text{Acc} \Rightarrow \pi'_0 \pi'_1 \pi'_2 \cdots \in \text{Acc}$ holds. For details, see [34]. On the other hand, to preserve boundedness and unboundedness properties, one needs to bound the length of the π'_i and the length of the π_i . Hence, there is only a weak aggregation for max-automata. Again, see [32] for details.

Given a weak aggregation s for \mathfrak{A} with acceptance condition Acc , let

$$s(\text{Acc}) = \{s((\pi_i)_{i \in \mathbb{N}}) \mid \pi_0 \pi_1 \pi_2 \cdots \in \text{Acc} \text{ is an accepting run of } \mathfrak{A} \text{ with } \sup_i |\pi_i| < \infty\}.$$

Next, we consider aggregations that are trackable by automata. A monitor for an automaton \mathfrak{A} with transition function δ is a tuple $\mathfrak{M} = (M, \perp, \text{upd})$ where M is a finite set of memory elements,

$\perp \notin M$ is the empty memory element, and $\text{upd}: M_\perp \times \delta \rightarrow M$ is an update function, where we use $M_\perp = M \cup \{\perp\}$. Note that the empty memory element \perp is only used to initialize the memory, it is not in the image of upd . We say that \mathfrak{M} computes the function $s_{\mathfrak{M}}: \delta^+ \rightarrow M$ defined by $s_{\mathfrak{M}}(t) = \text{upd}(\perp, t)$ and $s_{\mathfrak{M}}(\pi \cdot t) = \text{upd}(s_{\mathfrak{M}}(\pi), t)$ for $\pi \in \delta^+$ and $t \in \delta$.

Example 2. Recall Example 1. The strong aggregation s_{prty} for a parity automaton is computed by the monitor $(\Omega(Q), \perp, (c, t) \mapsto \max\{c, \Omega(t)\})$, where $\perp < c$ for every $c \in \Omega(Q)$.

Next, we take the product of \mathfrak{A} and the monitor \mathfrak{M} for \mathfrak{A} , which simulates \mathfrak{A} and simultaneously aggregates the acceptance condition. Formally, we define the product as $\mathfrak{A} \times \mathfrak{M} = (Q \times M_\perp, (q_I, \perp), \Sigma, \delta', \emptyset)$ where $\delta'((q, m), a) = (q', \text{upd}(m, (q, a, q')))$ for $q' = \delta(q, a)$. Note that $\mathfrak{A} \times \mathfrak{M}$ has an empty set of accepting runs, as these are irrelevant to us.

4.2 Removing Delay via Aggregation

Consider a play prefix in a delay game $\Gamma_f(L(\mathfrak{A}))$: Player I has produced a sequence $\alpha(0) \cdots \alpha(i)$ of letters while Player O has produced $\beta(0) \cdots \beta(i')$ with, in general, $i' < i$. Now, she has to determine $\beta(i' + 1)$. The automaton $\mathfrak{A} \times \mathfrak{M}$ can process the joint sequence $\binom{\alpha(0) \cdots \alpha(i')}{\beta(0) \cdots \beta(i')}$, but not the sequence $\alpha(i' + 1) \cdots \alpha(i)$, as Player O has not yet picked the letters $\beta(i' + 1) \cdots \beta(i)$. However, one can determine which states are reachable by some completion $\binom{\alpha(i'+1) \cdots \alpha(i)}{\beta(i'+1) \cdots \beta(i)}$ by projecting away Σ_O from $\mathfrak{A} \times \mathfrak{M}$.

Thus, from now on assume $\Sigma = \Sigma_I \times \Sigma_O$ and define $\delta_P: 2^{Q \times M_\perp} \times \Sigma_I \rightarrow 2^{Q \times M}$ via

$$\delta_P(S, a) = \left\{ \delta' \left((q, m), \binom{a}{b} \right) \mid (q, m) \in S \text{ and } b \in \Sigma_O \right\}.$$

Intuitively, δ_P is obtained as follows: take $\mathfrak{A} \times \mathfrak{M}$, project away Σ_O , and apply the power set construction (while discarding the anyway empty acceptance condition). Then, δ_P is the transition function of the resulting deterministic automaton. As usual, we extend δ_P to $\delta_P^+: 2^{Q \times M_\perp} \times \Sigma_I^+ \rightarrow 2^{Q \times M}$ via $\delta_P^+(S, a) = \delta_P(S, a)$ and $\delta_P^+(S, wa) = \delta_P(\delta_P^+(S, w), a)$.

Remark 1. The following are equivalent for $q \in Q$ and $w \in \Sigma_I^+$:

1. $(q', m') \in \delta_P^+(\{(q, \perp)\}, w)$.
2. There is a $w' \in (\Sigma_I \times \Sigma_O)^+$ whose projection to Σ_I is w such that the run π of \mathfrak{A} processing w' starting from q ends in q' and satisfies $s_{\mathfrak{M}}(\pi) = m'$.

We use this property to define an equivalence relation formalizing the idea that words having the same behavior in $\mathfrak{A} \times \mathfrak{M}$ do not need to be distinguished. To this end, to every $w \in \Sigma_I^+$ we assign the transition summary $r_w: Q \rightarrow 2^{Q \times M}$ defined via $r_w(q) = \delta_P^+(\{(q, \perp)\}, w)$. Having the same transition summary is a finite equivalence relation \equiv over Σ_I^+ whose index is bounded by $2^{|Q|^2 \cdot |M|}$. For an \equiv -class $S = [w]_{\equiv}$ define $r_S = r_w$, which is independent of representatives. Let R be the set of infinite \equiv -classes.

Now, we define a Gale-Stewart game in which Player I determines an infinite sequence of equivalence classes from R . By picking representatives, this induces a word $\alpha \in \Sigma_I^\omega$. Player O picks states (q_i, m_i) such that the m_i aggregate a run of \mathfrak{A} on some completion $\binom{\alpha}{\beta}$ of α . Player O wins if the m_i imply that the run of \mathfrak{A} on $\binom{\alpha}{\beta}$ is accepting. To account for the delay, Player I is always one move ahead, which is achieved by adding a dummy move for Player O in round 0.

Formally, in round 0, Player I picks an \equiv -class $S_0 \in R$ and Player O has to pick $(q_0, m_0) = (q_I, \perp)$. In round $i > 0$, first Player I picks an \equiv -class $S_i \in R$, then Player O picks a state $(q_i, m_i) \in r_{S_{i-1}}(q_{i-1})$ of the product automaton. Player O wins the resulting play $S_0(q_0, m_0)S_1(q_1, m_1)S_2(q_2, m_2) \cdots$ if $m_1 m_2 m_3 \cdots \in$

$s_{\mathfrak{M}}(\text{Acc})$ (note that m_0 is ignored). The notions of (finite-state and winning) strategies are inherited from Gale-Stewart games, as this game is indeed such a game $\Gamma(L(\mathfrak{B}))$ for some automaton \mathfrak{B} of size $|R| \cdot |Q| \cdot |M|$ which can be derived from \mathfrak{A} and \mathfrak{M} .

Formally, we define $\mathfrak{B} = (R \times Q \times M_{\perp}, R \times (Q \times M), (S_I, q_I, m_I), \delta', \text{Acc}')$ for some arbitrary $S_I \in R$, some arbitrary $m_I \in M$, $\delta'((S, q, m), \binom{S'}{q', m'}) = (S', q', m')$, and $(S_0, q_0, m_0)(S_1, q_1, m_1)(S_2, q_2, m_2) \cdots \in \text{Acc}'$ if, and only if,

- $(q_0, m_0) = (q_I, \perp)$,
- $(q_i, m_i) \in r_{S_{i-1}}(q_{i-1})$ for all $i > 0$, and
- $m_1 m_2 m_3 \cdots \in s_{\mathfrak{M}}(\text{Acc})$.

It is straightforward to prove that \mathfrak{B} has the desired properties.

Note that, due to our very general definition of acceptance conditions, we are able to express the local consistency requirement “ $(q_i, m_i) \in r_{S_{i-1}}(q_{i-1})$ ” using the acceptance condition. For less general acceptance modes, e.g., parity, one has to check this property using the state space of the automaton, which leads to a polynomial blowup, as one has to store each S_{i-1} for one transition.

Theorem 2. *Let \mathfrak{A} be an ω -automaton and let \mathfrak{M} be a monitor for \mathfrak{A} such that $s_{\mathfrak{M}}$ is a strong aggregation for \mathfrak{A} , let \mathfrak{B} be constructed as above, and define $d = 2^{|\Omega|^2 \cdot |M_{\perp}|}$.*

1. *If Player O wins $\Gamma_f(L(\mathfrak{A}))$ for some delay function f , then she also wins $\Gamma(L(\mathfrak{B}))$.*
2. *If Player O wins $\Gamma(L(\mathfrak{B}))$, then she also wins the block-game $\Gamma^d(L(\mathfrak{A}))$. Moreover, if she has a finite-state winning strategy for $\Gamma(L(\mathfrak{B}))$ with n states, then she has a delay-aware finite-state winning strategy for $\Gamma^d(L(\mathfrak{A}))$ with n states.*

By applying both implications and Item 2 of Lemma 1, we obtain upper bounds on the complexity of determining for a given \mathfrak{A} whether Player O wins $\Gamma_f(L(\mathfrak{A}))$ for some f and on the necessary constant lookahead necessary to do so.

Corollary 1. *Let \mathfrak{A} , \mathfrak{B} , and d be as in Theorem 2. Then, the following are equivalent:*

1. *Player O wins $\Gamma_f(L(\mathfrak{A}))$ for some delay function f .*
2. *Player O wins $\Gamma_f(L(\mathfrak{A}))$ for the constant delay function f with $f(0) = 2d$.*
3. *Player O wins $\Gamma(L(\mathfrak{B}))$.*

Thus, determining whether, given \mathfrak{A} , Player O wins $\Gamma_f(L(\mathfrak{A}))$ for some f is achieved by determining the winner of the Gale-Stewart game $\Gamma(L(\mathfrak{B}))$ and, independently, we obtain an exponential (in $|Q| \cdot |M|$) upper bound on the necessary constant lookahead.

Example 3. *Continuing our example for the parity acceptance condition, we obtain the exponential upper bound $2^{|\Omega|^2 \cdot |\Omega(Q)| + 2}$ on the constant lookahead necessary to win the delay game and an exponential-time algorithm for determining the winner, as \mathfrak{B} has exponentially many states, but the same number of colors as \mathfrak{A} . Both upper bounds are tight [19].*

In case there is no strong aggregation for \mathfrak{A} , but only a weak one, one can show that finite-state strategies exist, if Player O wins with respect to some constant delay function at all.

Theorem 3. *Let \mathfrak{A} be an ω -automaton and let \mathfrak{M} be a monitor for \mathfrak{A} such that $s_{\mathfrak{M}}$ is a weak aggregation for \mathfrak{A} , let \mathfrak{B} be constructed as above, and define $d = 2^{|\Omega|^2 \cdot |M_{\perp}|}$.*

1. *If Player O wins $\Gamma_f(L(\mathfrak{A}))$ for some constant delay function f , then she also wins $\Gamma(L(\mathfrak{B}))$.*

2. If Player O wins $\Gamma(L(\mathfrak{B}))$, then she also wins the block-game $\Gamma^d(L(\mathfrak{A}))$. Moreover, if she has a finite-state winning strategy for $\Gamma(L(\mathfrak{B}))$ with n states, then she has a delay-aware finite-state winning strategy for $\Gamma^d(L(\mathfrak{A}))$ with n states.

Again, we obtain upper bounds on the solution complexity (here, with respect to constant delay functions) and on the necessary constant lookahead.

Corollary 2. *Let \mathfrak{A} , \mathfrak{B} , and d be as in Theorem 3. Then, the following are equivalent:*

1. Player O wins $\Gamma_f(L(\mathfrak{A}))$ for some constant delay function f .
2. Player O wins $\Gamma_f(L(\mathfrak{A}))$ for the constant delay function f with $f(0) = 2d$.
3. Player O wins $\Gamma(L(\mathfrak{B}))$.

5 Discussion

Let us compare the two approaches presented in the previous section with three use cases: delay games whose winning conditions are given by deterministic parity automata, by deterministic Muller automata, and by LTL formulas. All formalisms only define ω -regular languages, but vary in their succinctness.

The following facts about arena-based games will be useful for the comparison:

- The winner of arena-based parity games has positional winning strategies [7, 23], i.e., finite-state strategies with a single state.
- The winner of an arena-based Muller game has a finite-state strategy with $n!$ states [22], where n is the number of vertices of the arena.
- The winner of an arena-based LTL game has a finite-state strategy with $2^{2^{|\varphi|}}$ states [24], where φ is the formula specifying the winning condition.

Also, we need the following bounds on the necessary lookahead in delay games:

- In delay games whose winning conditions are given by deterministic parity automata, exponential (in the size of the automata) constant lookahead is both sufficient and in general necessary [19].
- In delay games whose winning conditions are given by deterministic Muller automata, doubly-exponential (in the size of the automata) constant lookahead is sufficient. This follows from the transformation of deterministic Muller automata into deterministic parity automata of exponential size (see, e.g., [12]). However, the best lower bound is the exponential one for parity automata, which are also Muller automata.
- In delay games whose winning conditions are given by LTL formulas, triply-exponential (in the size of the formula) constant lookahead is both sufficient and in general necessary [20].

Using these facts, we obtain the following complexity results for finite-state strategies: Figure 1 shows the upper bounds on the number of states of delay-oblivious finite-state strategies for delay games and on the number of states of delay-aware finite-state strategies for block games. In all three cases, the former strategies are at least exponentially larger. This illustrates the advantage of decoupling tracking the history from managing the lookahead.

Finally, let us compare our approach to that of Salzmann. Fix a delay game $\Gamma_f(L(\mathfrak{A}))$ and assume Player I has picked $\alpha(0) \cdots \alpha(i)$ while Player O has picked $\beta(0) \cdots \beta(i')$ with $i' < i$. His strategies are similar to our delay-aware ones for block games. The main technical difference is that his strategies have access to the state reached by \mathfrak{A} when processing $\binom{\alpha(0) \cdots \alpha(i')}{\beta(0) \cdots \beta(i')}$. Thus, his strategies explicitly depend on

	parity	Muller	LTL
delay-oblivious	doubly-exp.	quadruply-exp.	quadruply-exp.
delay-aware	exp.	doubly-exp.	triply-exp.

Figure 1: Memory size for delay-oblivious strategies (for delay games) and delay-aware finite-state strategies (for block games), measured in the size of the representation of the winning condition. For the sake of readability, we only present the orders of magnitude, but not exact values.

the specification automaton \mathcal{A} while ours are independent of it. In general, his strategies are therefore smaller than ours, as our transducers have to simulate \mathcal{A} if they need access to the current state. On the other hand, our aggregation-based framework is more general and readily applicable to quantitative winning conditions as well, while he only presents results for selected qualitative conditions like parity, weak parity, and Muller.

6 Conclusion

We have presented a very general framework for analyzing delay games. If the automaton recognizing the winning condition satisfies a certain aggregation property, our framework yields upper bounds on the necessary lookahead to win the game, an algorithm for determining the winner (under some additional assumptions on the acceptance condition), and finite-state winning strategies for Player O , if she wins the game at all. These results cover all previous results on the first two aspects (although not necessarily with optimal complexity of determining the winner).

Thereby, we have lifted another important aspect of the theory of infinite games to the setting with delay. However, many challenging open questions remain, e.g., a systematic study of memory requirements in delay games is now possible. For delay-free games, tight upper and lower bounds on these requirements are known for almost all winning conditions.

Another exciting question concerns the tradeoff between memory and amount of lookahead: can one trade memory for lookahead? In other settings, such tradeoffs exist, e.g., lookahead allows Player O to improve the quality of her strategies [34]. Salzman has presented some tradeoffs between memory and lookahead, e.g., linear lookahead allows exponential reductions in memory size in comparison to delay-free strategies [28]. In current work, we investigate whether these results are inherent to his setting, which differs subtly from the one proposed here, or whether they exist in our setting as well.

Finite-state strategies in arena-based games are typically computed by game reductions, which turn a game with a complex winning condition into one in a larger arena with a simpler winning condition. In future work, we plan to lift this approach to delay games. Note that the algorithm for computing finite-state strategies presented here can already be understood as a reduction, as we turn a delay game into a Gale-Stewart game. This removes the delay, but preserves the type of winning condition. However, it is also conceivable that staying in the realm of delay games yields better results, i.e., by keeping the delay while simplifying the winning condition. In future work, we address this question.

In our study here we focussed on the state complexity of the automata implementing the strategies, i.e., we measure the quality of a strategy in the number of states of a transducer implementing it. However, this is not the true size of such a machine, as we have ignored the need to represent the transition function and the output function, which have an exponential domain (in the block size) in the case of delay-aware strategies. Thus, when represented as lookup tables, they are prohibitively large. However,

our delay-removing reduction hints at these functions also being implementable by transducers. For the transition function this is straightforward; in current work, we investigate whether this is also possible for the output function.

Finally, in future work we will determine the complexity of computing finite-state strategies in delay games and investigate notions of finite-state strategies for Player I , which should be much simpler since he does not have to deal with the lookahead.

Acknowledgements The author is very grateful to the anonymous reviewers whose feedback significantly improved the exposition.

References

- [1] Mikołaj Bojańczyk (2011): *Weak MSO with the Unbounding Quantifier*. *Theory Comput. Syst.* 48(3), pp. 554–576, doi:10.1007/s00224-010-9279-2.
- [2] J. Richard Büchi & Lawrence H. Landweber (1969): *Solving Sequential Conditions by Finite-State Strategies*. *Transactions of the American Mathematical Society* 138, pp. 295–311, doi:10.2307/1994916.
- [3] Cristian S. Calude, Sanjay Jain, Bakhadyr Khoussainov, Wei Li & Frank Stephan (2017): *Deciding Parity Games in Quasipolynomial Time*. In Hamed Hatami, Pierre McKenzie & Valerie King, editors: *STOC 2017*, ACM, pp. 252–263, doi:10.1145/3055399.3055409.
- [4] Krishnendu Chatterjee, Thomas A. Henzinger & Florian Horn (2011): *The Complexity of Request-Response Games*. In Adrian Horia Dediu, Shunsuke Inenaga & Carlos Martín-Vide, editors: *LATA 2011, LNCS 6638*, Springer, pp. 227–237, doi:10.1007/978-3-642-21254-3_17.
- [5] Thomas Colcombet, Nathanaël Fijalkow & Florian Horn (2014): *Playing Safe*. In Venkatesh Raman & S. P. Suresh, editors: *FSTTCS 2014, LIPIcs 29*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 379–390, doi:10.4230/LIPIcs.FSTTCS.2014.379.
- [6] Stefan Dziembowski, Marcin Jurdziński & Igor Walukiewicz (1997): *How Much Memory is Needed to Win Infinite Games?* In: *LICS 1997*, IEEE Computer Society, pp. 99–110, doi:10.1109/LICS.1997.614939.
- [7] E. Allen Emerson & Charanjit S. Jutla (1991): *Tree Automata, Mu-Calculus and Determinacy (Extended Abstract)*. In: *FOCS 1991*, IEEE Computer Society, pp. 368–377, doi:10.1109/SFCS.1991.185392.
- [8] John Fearnley, Sanjay Jain, Sven Schewe, Frank Stephan & Dominik Wojtczak (2017): *An Ordered Approach to Solving Parity Games in Quasi Polynomial Time and Quasi Linear Space*. In Hakan Erdogmus & Klaus Havelund, editors: *SPIN 2017*, ACM, pp. 112–121, doi:10.1145/3092282.3092286.
- [9] Nathanaël Fijalkow & Florian Horn (2013): *Les jeux d’accessibilité généralisée*. *Technique et Science Informatiques* 32(9-10), pp. 931–949, doi:10.3166/tsi.32.931-949. See also arXiv:1010.2420.
- [10] Wladimir Fridman, Christof Löding & Martin Zimmermann (2011): *Degrees of Lookahead in Context-free Infinite Games*. In Marc Bezem, editor: *CSL 2011, LIPIcs 12*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 264–276, doi:10.4230/LIPIcs.CSL.2011.264.
- [11] David Gale & Frank M. Stewart (1953): *Infinite Games with Perfect Information*. *Annals of Mathematics* 28, pp. 245–266, doi:10.1515/9781400881970-014.
- [12] Erich Grädel, Wolfgang Thomas & Thomas Wilke, editors (2002): *Automata, Logics, and Infinite Games: A Guide to Current Research*. LNCS 2500, Springer, doi:10.1007/3-540-36387-4.
- [13] Michael Holtmann, Łukasz Kaiser & Wolfgang Thomas (2012): *Degrees of Lookahead in Regular Infinite Games*. *LMCS* 8(3), doi:10.2168/LMCS-8(3:24)2012.
- [14] Florian Horn (2005): *Streit Games on Finite Graphs*. In: *GDV 2005*. Available at <https://www.irif.fr/~horn/publications.html>.

- [15] Frederick A. Hosch & Lawrence H. Landweber (1972): *Finite Delay Solutions for Sequential Conditions*. In: *ICALP 1972*, pp. 45–60.
- [16] Marcin Jurdzinski & Ranko Lazic (2017): *Succinct Progress Measures for Solving Parity Games*. In: *LICS 2017*, IEEE Computer Society, pp. 1–9, doi:10.1109/LICS.2017.8005092.
- [17] Bakhadyr Khoussainov (2003): *Finite State Strategies in One Player McNaughton Games*. In Cristian Calude, Michael J. Dinneen & Vincent Vajnovszki, editors: *DMTCS 2003, Dijon, LNCS 2731*, Springer, pp. 203–214, doi:10.1007/3-540-45066-1_16.
- [18] Felix Klein & Martin Zimmermann (2015): *What are Strategies in Delay Games? Borel Determinacy for Games with Lookahead*. In Stephan Kreutzer, editor: *CSL 2015, LIPIcs 41*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 519–533, doi:10.4230/LIPIcs.CSL.2015.519.
- [19] Felix Klein & Martin Zimmermann (2016): *How Much Lookahead is Needed to Win Infinite Games?* *LMCS* 12(3), doi:10.2168/LMCS-12(3:4)2016.
- [20] Felix Klein & Martin Zimmermann (2016): *Prompt Delay*. In Akash Lal, S. Akshay, Saket Saurabh & Sandeep Sen, editors: *FSTTCS 2016, LIPIcs 65*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 43:1–43:14, doi:10.4230/LIPIcs.FSTTCS.2016.43.
- [21] Donald A. Martin (1975): *Borel determinacy*. *Annals of Mathematics* 102, pp. 363–371, doi:10.2307/1971035.
- [22] Robert McNaughton (1993): *Infinite Games Played on Finite Graphs*. *Ann. Pure Appl. Logic* 65(2), pp. 149–184, doi:10.1016/0168-0072(93)90036-D.
- [23] Andrzej Mostowski (1991): *Games with Forbidden Positions*. Technical Report 78, University of Gdańsk.
- [24] Amir Pnueli & Roni Rosner (1989): *On the Synthesis of an Asynchronous Reactive Module*. In Giorgio Ausiello, Mariangiola Dezani-Ciancaglini & Simona Ronchi Della Rocca, editors: *ICALP 1989, LNCS 372*, Springer, pp. 652–671, doi:10.1007/BFb0035790.
- [25] Michael O. Rabin (1969): *Decidability of Second-Order Theories And Automata on Infinite Trees*. *Transactions of the American Mathematical Society* 141(1-35), p. 4, doi:10.2307/1995086.
- [26] Alexander Rabinovich (2009): *Synthesis of Finite-state and Definable Winning Strategies*. In Ravi Kannan & K. Narayan Kumar, editors: *FSTTCS 2009, LIPIcs 4*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 359–370, doi:10.4230/LIPIcs.FSTTCS.2009.2332.
- [27] Stéphane Le Roux & Arno Pauly (2016): *Extending Finite Memory Determinacy to Multiplayer Games*. In Alessio Lomuscio & Moshe Y. Vardi, editors: *SR 2016, EPTCS 218*, pp. 27–40, doi:10.4204/EPTCS.218.3.
- [28] Tobias Salzmann (2015): *How Much Memory is Needed to Win Regular Delay Games?* Master’s thesis, Saarland University.
- [29] Wolfgang Thomas (1994): *Finite-State Strategies in Regular Infinite Games*. In P. S. Thiagarajan, editor: *FSTTCS 1994, LNCS 880*, Springer, pp. 149–158, doi:10.1007/3-540-58715-2_121.
- [30] Nico Wallmeier, Patrick Hütten & Wolfgang Thomas (2003): *Symbolic Synthesis of Finite-State Controllers for Request-Response Specifications*. In Oscar H. Ibarra & Zhe Dang, editors: *CIAA 2003, LNCS 2759*, Springer, pp. 11–22, doi:10.1007/3-540-45089-0_3.
- [31] Igor Walukiewicz (2001): *Pushdown Processes: Games and Model-Checking*. *Inf. and Comput.* 164(2), pp. 234–263, doi:10.1006/inco.2000.2894.
- [32] Martin Zimmermann (2016): *Delay Games with WMSO+U Winning Conditions*. *RAIRO - Theor. Inf. and Applic.* 50(2), pp. 145–165, doi:10.1051/ita/2016018.
- [33] Martin Zimmermann (2017): *Finite-state Strategies in Delay Games (full version)*. arXiv:1704.08887. Available at <https://arxiv.org/abs/1704.08887>.
- [34] Martin Zimmermann (2017): *Games with Costs and Delays*. In: *LICS 2017*, IEEE Computer Society, pp. 1–12, doi:10.1109/LICS.2017.8005125.