Tradeoffs in Infinite Games

Habilitation thesis submitted to
the Faculty of Mathematics and Computer Science
at Saarland University

by

Dr. rer. nat. Martin Zimmermann.

Saarbrücken, April 26, 2017

Table of Contents

1 Infinite Games in Theoretical Computer Science

Games of infinite duration, infinite games for short, constitute one of the pillars of theoretical computer science with a wide range of applications in logics, automata theory, descriptive set theory, and program synthesis. For example, the model-checking problem, the task of deciding whether a given formula holds true in a given structure, can be interpreted as a game between a player trying to prove that the formula holds true, and a player trying to disprove this [Grä02][1]. For fixed-point logics, such games are of infinite duration. In fact, model-checking the modal μ-calculus is linear-time equivalent to determining the winner of a parity game [EJS01], a certain type of infinite game. Other applications of parity games include the seminal game-based proof of Rabin's theorem by Gurevich and Harrington [GH82], which considerably simplified Rabin's original proof [Rab69] showing the decidability of monadic second-order logic over infinite trees. Furthermore, emptiness problems for tree automata and alternating automata (see [GTW02] for an overview) as well as satisfiability problems for temporal logics (see [DGL16] for an overview) have been expressed as infinite games. Moreover, the existence of a bisimulation between two structures [JM99], the diagonalizability of relations [TL93, TB73], language inclusion between ω-automata [EWS05], the existence of Wadge reductions [Wad72, Wad83], and many other problems have been characterized in terms of infinite games.

In this work, we are most concerned with the reactive synthesis problem, which has its roots in Church's synthesis problem [Chu63]. Church asked for a given specification of the input-output behavior of a circuit, whether one can automatically construct a circuit that satisfies this specification, if such a circuit exists at all. As an example of such a specification, consider the conjunction of the following three requirements on a circuit with a single input bit and a single output bit:

1. Whenever the input bit is 1, then the output bit is 1, too.
2. At least one out of every three consecutive output bits is a 1.
3. If there are infinitely many 0's in the input stream, then there are infinitely many 0's in the output stream.

The following circuit which stores the last three inputs satisfies all three requirements: if the input bit is a 0, answer with a 0, unless the last two output bits were already 0, in this case output a 1. On the other hand, every 1 in the input stream is answered by a 1.

Büchi and Landweber [BL69] solved the problem by framing it as an infinite game between a player representing the environment of the circuit, who produces the stream of input bits, and a player representing the circuit to be synthesized, who is in charge of producing the output bits. The game is played in rounds $n = 0, 1, 2, \ldots$ as follows: in round n, first the environment player picks a bit $\alpha(n)$, then the circuit player picks a bit $\beta(n)$. After ω rounds, the circuit player wins, if the sequence $\alpha(0)\beta(0)\alpha(1)\beta(1)\cdots$ satisfies the specification.

[1]There are two types of references in this introductory part of the thesis: numeric references refer to publications constituting this thesis while all other references are alphanumeric.

Typically, such a game is modeled as a graph-based game. Here, the arena of the game, which describes the rules of the game, is a directed graph whose vertices are partitioned into the positions of the two players. We call them Player 0 and Player 1 and assume for pronomial convenience [McN00] Player 0 to be female and Player 1 to be male. Such a game is played as follows: a token is placed at some initial vertex and always moved to some successor vertex, which is picked by the player who owns the vertex the token is currently at. After ω moves, the players have produced an infinite path through the arena, a so-called play. The winner is determined by the winning condition, a subset of the plays. If the play is in the winning condition, then Player 0 wins, otherwise Player 1 wins.

Figure 1 shows a graph-based game formalizing the instance of Church's synthesis problem given by the specification above. Here, the vertices of Player 0 are depicted as circles and those of Player 1 as squares. As the game models the players picking bits, each vertex has two outgoing edges, one representing picking a 0 (dashed) and the other one representing picking a 1 (solid). The initial vertex is marked by an incoming arrow. Note that the vertices are "colored" by natural numbers, which induce the winning condition: a play is winning for Player 0, if either only vertices of color 0 are seen infinitely often (but none of color 1 or 2) or if color 2 is visited infinitely often. The set of winning plays captures the specification above, the first two conjuncts are encoded in the transition structure and the fact that the lower right vertex is a losing sink for Player 0, and the third condition is expressed by the winning condition. The winning condition of the example above is a parity condition: the vertices are "colored" by natural numbers and a play is winning for Player 0, if the parity of the maximal color visited infinitely often is zero.

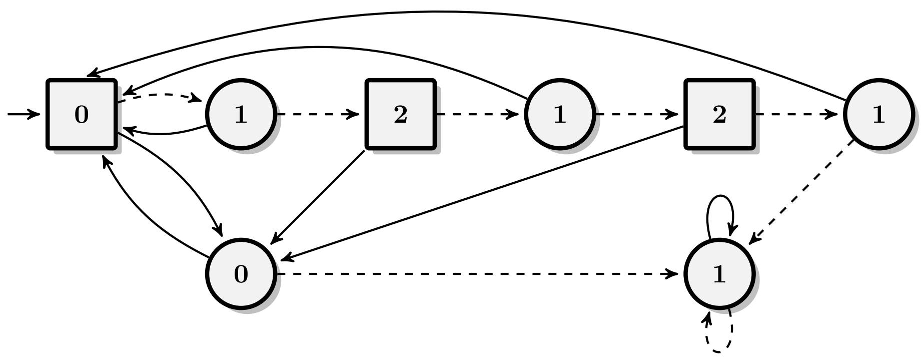

Figure 1: A graph-based infinite game.

The following is a winning strategy for Player 0, who has to pick successors at the circle vertices: at the leftmost two vertices of color 1 move to the right, at the rightmost vertex of color 1 and at the lower vertex of color 0 move back to the initial vertex. It is easy to verify that this strategy is winning for her and corresponds to the strategy described above informally.

In a graph-based game, a strategy for Player i is a function mapping a play prefix ending in a vertex v of Player i to some successor of v. It is said to be winning, if every play that starts in the initial vertex and is played according to the strategy is won by Player i. If some player has a winning strategy, then we say that she wins the game. Typically, one is interested in determining whether Player 0 wins a given game and in computing a winning strategy for her.

In the setting of graph-based infinite games, Büchi and Landweber's seminal theorem solving the synthesis problem reads as follows: "The winner of an infinite game in a finite arena with ω-regular winning condition can be determined effectively and a finite-state winning strategy for the winner can be computed [BL69]." Here, a finite-state strategy is essentially a finite automaton with output that implements a strategy by reading play prefixes and returning a next move to be taken. Finite-state strategies are desirable as they are a finite representation of an a priori infinite object. An important special case of finite-state strategies are positional strategies, which base their decision only on the vertex the token is currently at. The strategy described above for the game in Figure 1 is positional.

A simple proof of the Büchi-Landweber theorem reduces the problem to that of determining the winner (and a winning strategy) of a parity game. To this end, one starts from a deterministic parity automaton $\mathscr{A}$ recognizing the winning condition and takes the product of the arena and the automaton, which is still an infinite game. There is a bijection between the plays ρ in the original arena and the extended plays in the product game, which are products of a play ρ and the unique run of $\mathscr{A}$ on ρ. Thus, declaring those plays of the product to be winning for Player 0 that have an accepting run in the second component, yields an equivalent game. Finally, as winning now only depends on the acceptance of the run, the product game can inherit the acceptance condition of the automaton, i.e., the product is a parity game. Thus, determining the winner of the product game yields also the winner of the original game. Furthermore, positional winning strategies suffice for winning parity games [EJ91, Mos91] and can be computed effectively. Such a strategy and the automaton $\mathscr{A}$ recognizing the winning condition can easily be turned into a finite-state winning strategy for the original game. Finally, if the original game is an instance of Church's synthesis problem, then the finite-state strategy can even be turned into a circuit. This is another example of the importance of parity games.

The setting considered in the original Büchi-Landweber theorem is that of turn-based two-player deterministic zero-sum games of duration ω in finite arenas with ω-regular winning condition and complete information for both players. Since the publication of the original result in 1969 all these characteristics have been generalized to concurrent games [dAH00] and delay games [HL72], multi-player games [GU08], stochastic games [CH12, Sha53], games of ordinal length greater than ω [CH08], non-zero sum games [GU08], games of imperfect information [DR11], games in infinite arenas [Wal01], timed games [BCD$^+$07], and to distributed settings [MW03, PR90]. Finally, the variation receiving the most attention concerns the winning condition: Both special cases with better properties in terms of algorithmic complexity and memory requirements and more expressive winning conditions have been studied. The latter subsumes both qualitative extensions and quantitative extensions. Quantitative extensions are of particular importance for reactive synthesis, as explained later.

In this work, we consider several of these generalizations, often two at a time. The publications described in Section 4.1 and Section 4.2 investigate more expressive (qualitative and quantitative) winning conditions for infinite games and distributed infinite games. The publications described in Section 4.3 and

Section 4.4 are concerned with delay games with (qualitative and quantitative) winning conditions. Furthermore, games in infinite arenas and games of finite duration are used as tools to solve the problems at hand.

2 State of the Art

The two basic problems one is interested in for a class of infinite games are the computational complexity of determining the winner (and a winning strategy), the so-called solution complexity, and the worst-case memory requirements of such a winning strategy. Both of these characteristics depend on the expressiveness and the succinctness of the winning condition. For this reason, it is prudent to distinguish between high-level winning conditions, typically expressive and succinct specification logics, and low-level winning conditions, typically derived from acceptance conditions for automata on infinite objects [GTW02, Tho90], which tend to be less succinct.

For almost all low-level winning conditions, the exact solution complexity and tight upper and lower bounds on the memory requirements for both players are known. The most notable exception are parity games: while it is known that positional winning strategies exist for both players [EJ91, Mos91], solving such games remains the most intriguing problem in the theory of infinite games: There is a plethora of algorithms (e.g., [BDHK06, BCJ+97, EL86, FS13a, HKLN12, Jur00, McN93, Obd03, Sch17, VJ00, Zie98, ZP96]), some with subexponential [BV07, JPZ08, Lud95] and even quasipolynomial [CJK+16, FJS+17, JL17] running time, but it is open whether there is a polynomial time algorithm. In contrast, the solution problem for parity games is known to be in UP ∩ Co-UP [Jur98] and therefore unlikely to be NP-complete.

Similarly, for Linear Temporal Logic (LTL) [Pnu77], the de-facto standard linear-time specification language for reactive systems and the foundation of current high-level specification languages [AFF+02, EF06], the solution problem is 2Exptime-complete and doubly-exponential memory is both sufficient and in general necessary for both players [PR89a, PR89b, Ros91]. This result is shown similarly to the proof sketch of the Büchi-Landweber theorem above, starting with a deterministic parity automaton for the winning condition, which is, in general, of doubly-exponential size.

Most of the generalizations discussed in the previous section raise new questions beyond the basic ones about complexity and memory requirements. For example, multi-player games and non-zero sum games require solution concepts beyond winning strategies, e.g., equilibria. Similarly, stochastic strategies, which are necessary in concurrent and imperfect-information games, are typically also not winning in the sense introduced above, but one is interested in bounding the probability of losing. All these extensions have received considerable effort from the community and most of their properties are well understood.

2.1 Delay Games

One notable type of game that has not yet been thoroughly studied are delay games where one of the players is afforded a lookahead on her opponent's moves. These were introduced in 1972 [HL72], only three years after the Büchi-Landweber theorem, but have only recently been revisited [HKT12]. The evolution of the lookahead in a delay game is part of the rules, i.e., a delay game consists of a delay function $f\colon \mathbb{N} \to \mathbb{N}\setminus\{0\}$ and a winning condition $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$, where Σ_I and Σ_O are the input alphabet and the output alphabet, respectively. The game is played for ω rounds. But instead of both players picking one letter in each round as in the game modeling Church's problem, first Player I (the antagonist) has to pick $f(n)$ letters from Σ_I in round n, then Player O (the protagonist) picks a single letter from Σ_O. Thus, the lookahead increases by $f(n) - 1$ letters in round n. After ω rounds, both players have produced an infinite word over their alphabet; Player O wins if the pair of both words is in the winning condition L.

Typically, one is given a winning condition L and is interested in what kind of delay function f is necessary for Player O to win the game induced by L and f. Of particular interest are the *bounded delay functions*, , i.e., those that have finitely many n with $f(n) > 1$. This implies that the size of the lookahead is eventually stable. Even more restricted are the constant delay functions, those with $f(n) = 1$ for all $n > 1$, i.e., the size of the lookahead is constantly equal to $f(0) - 1$.

The asynchronous transmission of data in networks or components with buffers can be modeled via delay games. More theoretically, delay games also solve the diagonalizability problem [TL93, TB73] for relations $L \subseteq (\Sigma_I \times \Sigma_O)^\omega$, which asks, given L, whether there is a continuous function $\sigma\colon \Sigma_I^\omega \to \Sigma_O^\omega$ whose graph is a subset of L. Then, we say σ diagonalizes L. A relation L is diagonalized by a continuous function if, and only if, Player O wins the delay game with winning condition L with respect to some arbitrary delay function [HKT12]. Similarly, L is diagonalized by a Lipschitz-continuous function if, and only if, Player O wins the delay game with winning condition L with respect to some bounded delay function.

Hosch and Landweber [HL72] proved the decidability of the following problem: given an ω-regular L, is there a constant delay function f such that Player O wins the delay game induced by L and f. Almost 40 years later, the problem was revisited by Holtmann, Kaiser, and Thomas [HKT12], who proved that, for an ω-regular winning condition, Player O wins with respect to some arbitrary delay function if, and only if, she wins with respect to some constant delay function, i.e., unbounded lookahead does not provide an additional advantage over constant lookahead if the winning condition is ω-regular. Furthermore, they streamlined the decidability result of Hosch and Landweber, gave an algorithm with doubly-exponential running time, and gave doubly-exponential upper bounds on the constant lookahead necessary to win (in the size of a deterministic parity automaton recognizing the winning condition).

2.2 Quantitative Games: Boundedness vs. Optimization

During the last decade another shortcoming of the classical setting of Büchi and Landweber was addressed by considering quantitative instead of qualitative winning conditions. For example, consider a simple request-response property [WHT03] asking every request to be eventually responded to. This property is ω-regular and is therefore covered by the Büchi-Landweber theorem.

However, it does not rule out that the *waiting time* between requests and responses diverges. Such a behavior is undesirable, but cannot be prevented with ω-regular conditions.

The boundedness variant of the problem asks for the existence of an arbitrary, but fixed bound on the waiting time, i.e., it is a quantitative strengthening. Recently, boundedness problems in logics and automata theory have seen a lot of attention [BtCCV15, BCK+14, Boj04, Boj11, BC06, BGMS14, BPT16, BT09, Col13, CKL10, CL10, FHKS15, HS12] and far-reaching decidability results have been presented. In particular, Rabin's theorem has been extended to quantitative logics [Boj14, BT12] (see also [FHKS15, Van11]). Taking this approach one step further turns infinite games and reactive synthesis from a decision problem into an optimization problem: Instead of determining whether Player 0 wins a game with a qualitative winning condition, one defines semantic quality measures for winning strategies and aims to compute an optimal winning strategy for her. In the request-response example above, one would ask for a strategy that realizes the minimal waiting time among all winning strategies. As in the case of request-response conditions, these quality measures are often obtained by adding quantitative features to a qualitative winning condition, e.g., by bounding the waiting times in request-response games. In the same vein, there are quantitative extensions of LTL by parameterized eventually operators whose scope is bounded in time, e.g., Prompt-LTL and Parametric LTL (PLTL).

For example, the LTL formula $\mathbf{G}\,(q \rightarrow \mathbf{F}\,p)$ expresses the request-response property mentioned above. Formally, every occurrence of the atomic proposition q representing requests has to be answered by a later occurrence of p, representing the response. Again, LTL cannot prevent diverging waiting times. To overcome this limitation, Alur et al. introduced PLTL [AELP01], which adds the unary temporal operators $\mathbf{F}_x$ and $\mathbf{G}_y$, where x and y are variables ranging over natural numbers that are used to bound the scope of the operators. For example, with respect to a valuation $\alpha(x)$ of x, the formula $\mathbf{G}\,(q \rightarrow \mathbf{F}_x\,p)$ expresses that every request has to be answered within $\alpha(x)$ steps. Thus, dually, $\mathbf{G}_y\,\psi$ holds true with respect to $\alpha(y)$, if ψ holds true at least for the next $\alpha(y)$ steps.

In decision problems, the valuation α is typically existentially quantified, i.e., the PLTL synthesis problem asks, given an arena labeled by a set of atomic propositions and given a PLTL formula φ, whether there is a valuation α of φ's variables such that Player 0 has a winning strategy such that every consistent play satisfies φ with respect to α. Alur et al. proved that the PLTL model-checking problem has the same complexity as the classical special case of LTL model-checking, i.e., it is PSPACE-complete [AELP01]. Later, a similar result was shown for the PLTL synthesis problem: it is still 2EXPTIME-complete and doubly-exponential memory is both necessary and sufficient [Zim13]. Thus, one can increase the expressiveness of LTL for free when considering boundedness problems. In contrast, the complexity of the optimization variant of the synthesis problem, asking for a minimal α is still open: there is an algorithm with triply-exponential running time and the problem is trivially 2EXPTIME-complete [Zim13].

Shortly after the introduction of PLTL, Kupferman, Piterman, and Vardi studied Prompt-LTL [KPV09], which can be seen as the fragment of PLTL without parameterized always operators and with a single variable bounding the parameterized eventually operators. They argued that the usefulness of the parameterized always operators is questionable when considering boundedness problems, as they trivialize due to monotonicity: if ψ can be satisfied for some number $\alpha(y)$ of steps, then it can also be satisfied for zero steps, i.e., $\mathbf{G}_y\,\psi$ is equivalent to ψ when existentially quantifying the value for y. Similarly, if ψ holds true within the next $\alpha(x)$ steps at least once, then also at least once within the next k steps for every $k \geq \alpha(x)$. Hence, when asking for the existence of some variable valuation that satisfies the formula, one does not lose generality when assuming all parameterized eventually operators to be parameterized by the same variable. These two observations imply that boundedness problems for PLTL can be reduced to boundedness problems for Prompt-LTL.

The most influential contribution in [KPV09] is the introduction of the so-called *alternating-color technique*. Alur et al.'s decidability proof for PLTL model-checking [AELP01] relies on intricate pumping arguments, which are not generalizable to solving games. In contrast, Kupferman et al. presented a very general approach to dealing with boundedness problems, the alternating-color technique, which essentially consists of a reduction from Prompt-LTL to LTL: a new proposition is introduced that colors infinite traces and a subformula $\mathbf{F}_x\,\psi$ is replaced by an LTL formula that expresses that ψ holds true within at most one color change. If the distance between the color changes is bounded from below and from above, then these two formulas are equivalent. In a game, bounding the distance between color changes can be achieved easily by charging Player 0 with coloring the plays with infinitely many color changes. Then, a simple pumping argument shows that a finite-state strategy for her, which always exists if she wins, bounds the distance. This technique turned out to be applicable in other settings as well [CKL10].

Another important class of quantitative winning conditions are finitary parity and Streett conditions, which strengthen qualitative parity and Streett conditions. Recall that the classical parity condition is defined over sequences of colors (natural numbers) and holds true if the maximal color occurring infinitely often is even, e.g., the sequence

$$\pi = 1\,0\,2\;1\,0\,0\,2\;1\,0\,0\,0\,2\;1\,0\,0\,0\,0\,2\;1\,0\,0\,0\,0\,0\,2\;1\,0\,0\,0\,0\,0\,0\,2\;1\,0\,0\,0\,0\,0\,0\,0\,2\cdots$$

satisfies the parity condition, as the largest color occurring infinitely often is 2. Equivalently, one can require that almost all occurrences of an odd color are followed by a larger even color, e.g., in the example every 1 is eventually followed by a 2. Here, occurrences of odd colors can be understood as requests that are answered by larger even colors. The finitary parity condition introduced by Chatterjee and Henzinger [CH06] now requires the existence of a bound b such that almost all requests are answered within b steps. Hence, the sequence π above does not satisfy the finitary parity condition, as the distance between the occurrences of the request 1 and their answers (the next occurrence of a 2) diverges.

While the finitary parity condition strengthens the parity conditions, Horn showed finitary parity games to be solvable in polynomial time [Hor07] (see also [CHH09]), while there is no known algorithm solving parity games in polynomial time. Surprisingly, this situation is reversed for Streett conditions, which generalize parity games by giving up the hierarchical nesting of the requests and responses [Str81]: solving them is Co-NP-complete [Hor05] while solving finitary Streett games, which are defined as expected, is EXPTIME-hard[2], and thus much harder (under standard complexity-theoretic assumptions).

Finally, energy conditions [BFL+08, CRR14, JLR13, JLS15] also have received considerable attention during the last decade modeling systems with a limited resource that can be drained and recharged. Typical requirements are to keep the energy level between some fixed upper and lower bound, often while satisfying some other condition. For example, energy-parity games can be solved in NP $\cap$ CO-NP and Player 0 needs exponential memory while Player 1 still has positional strategies [CD12]. In subsequent work, a multitude of generalizations [CD12, CRR14, JLS15, VCD+15] of the energy condition have been studied, e.g., the average-energy condition which requires the average energy level to be bounded from above by some given threshold. Solving such games with a given lower and upper bound on the energy level and a threshold on the average energy level is EXPTIME-complete [BMR+15].

3 Motivation

The results presented in this thesis are broadly concerned with tradeoffs between different characteristics of infinite games, e.g., the expressiveness of the winning condition, complexity of the solution problem, memory requirements of winning strategies, semantic quality of winning strategies, and amount of lookahead. As an example, consider the work on the boundedness variant of Prompt-LTL and PLTL games. Both logics extend LTL quantitatively, but the solution complexity and the memory requirements are not affected. Thus, one can increase the expressiveness of LTL for free.

In this thesis, we investigate how far this result can be strengthened without affecting the complexity of the solution problem and the memory requirements, i.e., we study the tradeoff between expressiveness and complexity in games with winning conditions specified in linear temporal logics:

How much can the expressiveness of LTL be increased for free?

The second type of tradeoff we investigate is that between the semantic quality of winning strategies and both the solution complexity and the memory requirements. For example, the optimization variant of Prompt-LTL and PLTL games are only known to be solvable in triply-exponential time and only a triply-exponential upper bound on the necessary memory is known. The best corresponding lower bounds are both the trivial doubly-exponential ones for LTL games. This raises the question:

Does playing a quantitative game optimally come at a price?

Then, we shift our attention to delay games. As already mentioned earlier, most of their basic properties are still open, i.e., there are no non-trivial lower bounds on the solution complexity and the necessary lookahead. Also, only ω-regular winning conditions have been considered, but nothing is known about the properties of weaker or stronger winning conditions. Thus, we first study the tradeoff between lookahead and solution complexity, thereby studying the basic properties of delay games:

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

What is the price of adding lookahead to infinite games?

It is straightforward to come up with simple games that can only be won with lookahead, i.e., lookahead allows you to win more games. But does lookahead also allow you to improve the quality of winning strategies in games that can even be won without lookahead:

Can lookahead be traded for quality of winning strategies?

We answer all these questions in this thesis.

4 My Contributions

The publications constituting this thesis are grouped into four lines of work along the four questions motivating this work. The next four subsections summarize our results in terms of the tradeoffs considered.

In Subsection 4.1, we consider the tradeoff between expressiveness of the winning condition and complexity (in terms of the solution problem and the memory requirements of optimal strategies). First, we review several publications on linear temporal logics and then one each on request-response games and on average-energy games.

Then, in Subsection 4.2 we consider the tradeoff between the size of winning strategies and their semantic quality, first for finitary games and games with costs, and then for Prompt-LTL games.

The last two subsections are concerned with delay games: first, in Subsection 4.3 we consider the influence the addition of lookahead has on the complexity of the solution problem. Also, here we are interested in the necessary lookahead to win a delay game with a given winning condition.

Finally, in Subsection 4.4, we consider the tradeoff between lookahead and the semantic quality of strategies.

4.1 Expressiveness vs. Complexity

The overarching theme of the first line of research is the quest for expressiveness: strengthening and generalizing formalisms for specifying winning conditions without increasing the solution complexity and the memory requirements. For LTL, we were able to show that this is indeed possible: in a series of publications, we added both qualitative and quantitative features to LTL while preserving all desirable algorithmic properties of LTL: standing on the shoulders of Kupferman, Piterman, and Vardi [KPV09]

we vastly increased the expressiveness of LTL for free.

Also, we studied an optimization problem for request-response games and solved an open problem concerning a variant of average-energy games. Here, the increase of expressiveness comes at a price in terms of the solution complexity.

4.1.1 Linear Temporal Logics

LTL is the de-facto standard specification language for reactive systems and the basis for industrial logics like PSL [EF06] and ForSpec [AFF$^+$02]. This popularity is based on being expressively equivalent to first-order logic while having intuitive variable-free syntax and semantics. Furthermore, its desirable algorithmic properties are all based on the exponential-compilation property: LTL formulas can be translated into equivalent Büchi automata of exponential size. This property yields the polynomial space model-checking algorithm and the algorithm solving LTL games in doubly-exponential time.

However, there are two shortcomings of LTL that have been identified and tackled in previous work: LTL is unable to express timing constraints and it is less expressive than the ω-regular languages. The first issue was addressed by introducing Prompt-LTL and PLTL while there is a long line of extensions of LTL with the full expressive power of the ω-regular languages [LS07, Var11, VW94, Wol83] and even temporal logics with the full expressive power of the ω-visibly pushdown languages [AEM04, AM04, Boz07, BS14a, BS14b], which are able to express recursive properties that go beyond ω-regularity.

In this line of work, we explore the limits of the exponential-compilation property by strengthening LTL as much as possible without losing the property. My work on this question started during my PhD studies where I studied infinite games with PLTL winning conditions (as a boundedness problem) and was able to show that the alternating-color technique is applicable to them as well [Zim13]: I proved the problem to be 2EXPTIME-complete and gave tight doubly-exponential bounds on the necessary memory. Then, I considered the optimization variant and showed triply-exponential upper bounds on complexity and memory.

Later, in collaboration with Peter Faymonville [1], I introduced and investigated Parametric Linear Dynamic Logic (PLDL), the first parameterized temporal logic with the full expressiveness of the ω-regular languages, which is obtained by adding parameters to the guarded temporal operators of LDL [Var11]. For example, the formula $\mathbf{G}\,(q \rightarrow \langle r\rangle_{\leq x}\, p)$ expresses that every request has to be answered by a response within some fixed, but arbitrary number of steps and the guard r — a regular expression — has to match between the request and the response. Hence, $\langle r\rangle_{\leq x}$ can be understood as a guarded generalization of the parameterized eventually operator $\mathbf{F}_x$.

We developed a generalization of the alternating-color technique for PLDL and proved that PLDL model-checking and solving PLDL games are still PSPACE- and 2EXPTIME-complete, respectively. Also, we gave doubly exponential bounds on the necessary memory in PLDL games. Thus, PLDL retains the same desirable properties of LTL while being more expressive, both qualitatively and and quantitatively.

Afterwards, I generalized the setting even further by considering a weighted framework and interpreted parameterized operators to have a scope of bounded costs instead of bounded time [2], e.g., here, the request-response property $\mathbf{G}\,(q \rightarrow \mathbf{F}_x\, p)$ expresses that the cost between every request and its response has to be bounded by some fixed, but arbitrary value. To this end, I introduced the logics cPLTL and cPLDL and by again generalizing the alternating-color technique to the weighted setting I showed that even these logics have the same desirable properties as LTL. Furthermore, if the weights are given in a unary encoding, then even the complexity of the optimization problems does not change.

All these results are similar in that they are based on generalizations of the alternating-color technique and on proving the exponential compilation property for these logics, i.e., we adapted the alternating-color technique to deal with guarded operators (for PLDL) and with costs (for cPLTL) and devised a novel translation of LDL (with additional operators to capture the color changes) to alternating Büchi automata.

However, there are still desirable properties that cannot be expressed in these logics, e.g., "there are never more responses than requests" ruling out superfluous responses. This property is not ω-regular, but ω-contextfree. Using the full class of ω-contextfree languages for specifying reactive systems is not prudent, as the class is not closed under conjunction and complementation. Instead, the fragment of ω-visibly pushdown languages [AM04] has been the focus of intensive research, as it has much better closure properties. Nevertheless, it still able to express the example property above and many other important properties of recursive systems.

There are several proposals for a temporal logic with the same expressive power as the ω-visibly pushdown languages [AEM04, AM04, Boz07, BS14a, BS14b], but most of them do not come close to LTL in terms of an intuitive syntax and semantics. Furthermore, the algorithmic properties of these logics are no longer as desirable as those of LTL, but not much harder. In particular model checking is often EXPTIME-complete. However, the syntax and semantics of these logics is no longer as intuitive as that of LTL or LDL.

With Alexander Weinert [3], I proposed to equip the intuitive temporal operators of LDL with visibly pushdown automata as guards to obtain a logic with intuitive syntax and semantics and with the expressive power of the ω-visibly pushdown languages, called VLDL (Visibly Linear Dynamic Logic). We presented an exponential translation of VLDL into ω-visibly pushdown automata via alternating jumping automata [Boz07], an automaton model without stack, but with non-local transitions. From this translation, we straightforwardly obtain exponential time algorithms for satisfiability and model-checking as well as an algorithm solving games in triply-exponential running time. Furthermore, we proved all these problems to be complete for the respective complexity classes. Thus, we finally have to pay a price for the vast increase in expressiveness and applicability from LDL to VLDL.

All algorithmic results in these papers were concerned with model-checking and solving graph-based games. In collaboration with Swen Jacobs and Leander Tentrup, I also considered distributed synthesis for Prompt-LTL specifications [4], the first work on distributed synthesis for parameterized specifications. Finkbeiner and Schewe [FS05] proved that decidability of distributed LTL synthesis depends on the imperfect information induced by the communication structure of the components and gave a necessary and sufficient criterion for decidability. By a reduction to the LTL case via the alternating-color technique we were able to show that the synchronous distributed Prompt-LTL synthesis is decidable for a fixed communication structure if, and only if, the synchronous distributed LTL synthesis problem is decidable for this structure. Furthermore, both problems have asymptotically the same complexity. For the asynchronous case, we gave a semi-algorithm, as the problem is already undecidable for two processes and LTL specifications. Whether the one-process asynchronous case is decidable for Prompt-LTL, as it is for LTL, is an open problem. In unpublished work, we have generalized these results to stronger logics like PLTL and PLDL [JTZ17].

Finally, with Bernd Finkbeiner [5], I considered an orthogonal extension of LTL: instead of adding generalized temporal operators, HyperLTL adds trace quantifiers to LTL. HyperLTL is able to express information flow properties to specify the secrecy and integrity of security-critical systems, e.g., non-interference and observational determinism [CFK+14]. For example, the formula

$$\forall \pi \forall \pi' (\mathbf{G}\,(i_\pi \leftrightarrow i_{\pi'})) \rightarrow (\mathbf{G}\,(o_\pi \leftrightarrow o_{\pi'}))$$

expresses that the system is input-deterministic, i.e., for all traces π and π', if they coincide on their truth values of the input bit i, then they also have to coincide on their output bit o.

It is known that HyperLTL model-checking is decidable [CFK+14, FRS15] while HyperLTL satisfiability is undecidable, even for a single quantifier alternation [FH16]. However, the foundations of HyperLTL had been unexplored: We proved that HyperLTL differs significantly from LTL, where every satisfiable formula has a model that is an ultimately periodic trace. In contrast, HyperLTL formulas have in general only infinite models, have in general no ω-regular models (not even ω-contextfree ones), and one can express in HyperLTL that every trace in a model is aperiodic. In a second part of this work, we presented a first-order logic for hyperproperties that is expressively equivalent to LTL, thereby lifting the seminal Theorem of Kamp [GPSS80, Kam68] to HyperLTL.

4.1.2 Request-response Conditions

In collaboration with Florian Horn, Wolfgang Thomas, and Nico Wallmeier [6], I investigated request-response games as introduced earlier. These games are very versatile, as conjunctions of request-response conditions cover conjunctions of safety conditions and Büchi conditions. One way to turn the qualitative request-response condition into a quantitative one is to require every request to be answered within an arbitrary, but fixed number of steps, as it is done for Prompt-LTL. In this work, however, we considered a more sophisticated approach by measuring the quality of a play by the limit superior of the mean accumulated waiting times between requests and their responses, i.e., the penalty grows quadratically.

By a reduction to mean-payoff games we proved that if Player 0 has a winning strategy for a request-response game, then she also has an optimal winning strategy, which can be determined in doubly-exponential time and implemented with doubly-exponential memory. Whether these bounds are tight is an open problem.

4.1.3 Average-energy Conditions

Furthermore, with multiple co-authors I studied another quantitative winning condition for infinite games, the average-energy condition. Here, the edges of the arena are labeled by integers that model the discharge or recharge of a resource, e.g., a battery. While the energy condition requires the energy level to be always between a given lower and/or a given upper bound, the average-energy condition requires the average energy level to be below a given threshold. This winning condition was first rigorously studied by Bouyer, Markey, Randour, Larsen, and Laursen [BMR+15]. One problem left open by them was the decidability of the average-energy condition in conjunction with a non-negativity constraint on the energy level, i.e., there is a lower bound on the energy level, but no upper bound. In collaboration with the authors of the original work on average-energy games, I was able to solve this open problem.

First, with Kim G. Larsen and Simon Laursen [LLZ16], I considered the variation where the bound on the average energy level is not part of the input, but existentially quantified. For applications this means that one does not check whether a given battery size suffices to satisfy the specification, but whether there is a battery size that allows to satisfy it. We proved this problem to be decidable by showing it to be equivalent to asking for the existence of an upper bound on the energy level. However, these reductions only proved the existence of some bound, but they were not sharp enough the solve the original problem with a given fixed bound on the average-energy.

Then, together with Patricia Bouyer, Piotr Hofman, Nicolas Markey, and Mickael Randour [7], I was able to show decidability of the original problem by proving a sufficient upper bound on the energy level, which yielded a reduction to the known case of average-energy games with both a given lower and upper bound on the energy level. The resulting algorithm has doubly-exponential running time and doubly-exponential memory is sufficient. Furthermore, we improved the lower bound on the complexity of the problem from EXPTIME-hardness to EXPSPACE-hardness.

4.2 Size vs. Quality

My results on playing games with quantitative winning conditions optimally mentioned above often show that the best (known) algorithms for the optimization variant are exponentially worse than those for the boundedness variant, e.g., for request-response games and for games with winning conditions in parameterized temporal logics. The same is true in these cases for the memory requirements. This naturally raised the second question motivating my work: Does playing a quantitative game optimally come at a price?

There are some examples of such tradeoffs, e.g., for the resource reachability problem in pushdown graphs [Lan14], where adding memory allows to keep counters smaller than the bounds that are obtainable with positional strategies. In contrast, there are also examples where such tradeoffs do not exist, e.g., a certain tradeoff between size and quality of strategies in boundedness games [FHKS15] has been ruled out, which refuted a conjecture with important implications for automata theory and logics.

For several other types of winning conditions, I was able to show that there is indeed a tradeoff:

> **Computing optimal strategies can be harder than computing arbitrary winning strategies and playing optimally may require more memory than just winning.**

These results have profound implications for the quest turning synthesis into an optimization problem. Computing the optimal strategy may be infeasible while computing an arbitrary, i.e., non-optimal, strategy might still be feasible. This motivates the study of approximation algorithms for quantitative games and the study of tradeoff-aware algorithms.

4.2.1 Winning Conditions with Costs

This line of work is about tradeoffs in finitary parity and Streett games as introduced above and in parity and Streett games with costs, an extension of the finitary variants introduced by Nathanaël Fijalkow and me [8]. Recall that the finitary parity condition requires a bound on the number of steps between a request and its response. The conditions with costs are evaluated in a setting with edge-costs and require a bound on the cost between a request and its response. Solving these games is at least as hard as solving parity games, as they are subsumed as the special case of cost zero on every edge. Similarly, finitary parity games are obtained as the special case of cost one on every edge. We showed that parity games with costs are also not harder than they have to be, i.e., as hard as parity games, and that Player 0 always has positional winning strategies for the boundedness variant.

Then, with Alexander Weinert [9], I considered the optimization variant of parity games with costs: determining the optimal bound is PSPACE-complete and an optimal strategy might need exponential memory to implement. Both lower bounds already hold for finitary parity games. Thus, unless P $=$ PSPACE, playing them optimally is harder than winning them. On the other hand, the upper bounds even hold for games whose costs are encoded in binary. In unpublished work, we showed the situation to be different for Streett conditions: determining the optimal bound is still EXPTIME-complete and exponential memory is necessary and sufficient [WZ17]. Both results already hold for the boundedness variant. Thus, playing Streett games with costs optimally is not harder than just winning them.

4.2.2 Prompt-LTL

Together with Leander Tentrup and Alexander Weinert [10], I was also able to exhibit tradeoffs between size and quality of winning strategies for Prompt-LTL games. Here, playing optimally may increase the memory requirements exponentially. However, the exact complexity of the optimization problem remains open.

Nevertheless, we took a step towards closing the gap by presenting an approximation algorithm with doubly-exponential running time that computes a strategy with respect to the bound $2k_{\text{opt}}$, where k_{opt} is the optimal bound. Thus, the algorithm matches the 2EXPTIME lower bound inherited from solving LTL games. Furthermore, being based on the principle of bounded synthesis [FS13b], the algorithm allows to specify upper bounds on the size of a strategy and on the quality of the strategy and then check whether there is a winning strategy with these characteristics, i.e., one can precisely prioritize size over quality and vice versa. Finally, we empirically evaluated a prototype implementation which was able to solve small examples.

4.3 Lookahead vs. Complexity

I initiated the first in-depth study of delay games, both with quantitative winning conditions (ω-regular and ω-contextfree ones), as well as with quantitative ones. With my collaborators, I resolved the problems left open by Holtmann, Kaiser, and Thomas, and we

determined the exact price of adding lookahead

for a wide range of winning conditions by developing a very general proof technique that is applicable to all these conditions.

4.3.1 Delay Games with Qualitative Winning Conditions

My work on delay games started shortly after Holtmann, Kaiser, and Thomas revisited the problem and proved decidability of ω-regular delay games with respect to arbitrary delay functions and that doubly-exponential constant lookahead is sufficient to win such games [HKT12]. A natural question back then was the extension of these results to ω-contextfree winning conditions, as delay-free games with such winning conditions remain decidable [Wal01].

Together with Wladimir Fridman and Christoph Löding [11], I proved that delay games with ω-contextfree winning conditions behave quite differently: determining the winner with respect to arbitrary delay functions or with respect to bounded delay functions is undecidable. Furthermore, unbounded lookahead might be necessary to win. Finally, even the growth rate of the necessary lookahead might be non-elementary. All our results hold for very small subclasses of ω-pushdown automata, which completed the quest to extend the positive results for ω-regular winning conditions to (subclasses of) ω-contextfree winning conditions.

Instead, I turned my attention back to the ω-regular case whose exact complexity had been left open by Holtmann, Kaiser, and Thomas. Together with Felix Klein [12][3], I was able to prove the problem to be EXPTIME-complete by presenting a faster algorithm and by proving the first lower bound on the problem, which already holds for very weak fragments. Furthermore, we were able to settle the necessary lookahead to win ω-regular games, again by improving the upper bound and by exhibiting the first non-trivial lower bound. All these results assume the winning condition to be given by a deterministic parity automaton while both lower bounds already hold for very weak subclasses.

Furthermore, these results also yielded trivial upper bounds for non-deterministic, universal, and alternating automata via determinization, which involves an exponential respectively doubly-exponential blowup, both in complexity and necessary lookahead. In later work, we showed these bounds to be tight.

Finally, together with Felix Klein [13] I studied the determinacy problem for delay games. A game is determined if it has a winner. This foundational property allows to infer the existence of a winning strategy for Player $1 - i$ from the non-existence of a winning strategy for Player i. This quantifier change

[3]Note that there is a bug in the proof of Theorem 4.8 of that paper. A corrected version is currently under review at LMCS. This thesis contains the paper as it is published right now, i.e., with the bug. The revised proof technique is also used to prove Theorem 1 in [16], which generalizes Theorem 4.8 of [12].

lies at the heart of Rabin's theorem. While the games one typically encounters, in particular those in this thesis, are determined, there are also undetermined games [GS53]. Martin's determinacy theorem [Mar75] posits that a wide class of winning conditions, those in the Borel hierarchy, all induce determined games. We lifted this far-reaching result to games with lookahead: delay games with Borel winning conditions are determined.

4.3.2 Delay Games with Quantitative Winning Conditions

In later work, I showed that the techniques developed for settling the qualitative ω-regular case are applicable to quantitative winning conditions as well.

For example, weak monadic second-order with the unbounding quantifier (WMSO+U) is a quantitative extension of monadic second-order logic (which captures the ω-regular languages) able to express (un)boundedness properties. Extending the approach to solving ω-regular delay games in this direction instead of ω-contextfree winning conditions turned out to be more promising [14]: delay games with WMSO+U winning conditions restricted to bounded delay functions are decidable and doubly-exponential lookahead is always sufficient in that case. However, I also found a WMSO+U winning condition that requires unbounded lookahead. Nevertheless, I also proved that WMSO+U cannot enforce a minimal growth rate of the lookahead, i.e., any unbounded delay function is sufficient. The problem of deciding delay games with WMSO+U winning conditions with respect to arbitrary delay functions remains open. Also, the need for unbounded lookahead makes these games undesirable: WMSO+U is just too expressive. Hence, in subsequent work I investigated fragments of WMSO+U for which bounded lookahead suffices.

With Felix Klein [15], I showed that the most natural candidate, Prompt-LTL, is much better behaved: again by generalizing the techniques developed for the ω-regular case and by combining it with the alternating-color technique, we proved solving Prompt-LTL delay games to be 3EXPTIME-complete and triply-exponential lookahead to be sufficient and in general necessary to win. On the one hand, all lower bounds already hold for LTL; on the other hand, the upper bounds also hold for the stronger parameterized logics discussed earlier. Finally, the techniques developed in this work also yielded the matching lower bounds on solution complexity and necessary lookahead for delay games with winning conditions given by non-deterministic, universal, and alternating automata.

Throughout my work on delay games, a pattern became apparent: adding delay causes an exponential increase in the complexity, e.g., from P-completeness to EXPTIME-completeness for safety conditions and from 2EXPTIME-completeness to 3EXPTIME-completeness for LTL and Prompt-LTL. Also, the necessary lookahead is of the same order of magnitude, i.e., exponential, respectively triply-exponential.

4.4 Lookahead vs. Quality

Finally, I studied the tradeoff between the necessary lookahead and the quality of strategies in delay games with finitary parity conditions and parity conditions with costs [16]. First of all, these games are still only EXPTIME-complete and therefore not harder than delay games with parity conditions. Furthermore, exponential lookahead is again always sufficient. In that respect, delay games behave similarly to delay-free games: adding costs comes for free.

Then, I constructed examples where one can improve the quality of a winning strategy by employing more lookahead and vice versa. Thus, there is a tradeoff:

Lookahead can indeed be traded for quality of strategies.

This result demonstrates the usefulness of adding delay to infinite games with quantitative winning conditions.

5 A Note on Notation

Several publications constituting this thesis share the same notation.

The publications [6, 8, 9] are all concerned with graph-based games and share most basic definitions, e.g., the ones of arenas, games, strategies, memory structures, and reductions. The introductory part of Section 2 in [8] contains all necessary notation and definitions. However, the paper on average-energy games [7] uses different notation and slightly different definitions due to technical necessities.

Similarly, the publications [11, 12, 13, 14, 15, 16] are all concerned with delay games and also share their basic definitions, which can be found, for example, in Section 2.2 of [12]. This section also contains some introductory examples of delay games.

Finally, the logics Prompt-LTL, PLTL, and PLDL appear in the publications [4, 10, 15], sometimes without giving the full syntax and semantics. Prompt-LTL is introduced in, for example, [4], we use the definition of PLTL as in [Zim13], and PLDL is introduced in [1].

6 Conclusion and Outlook

With the work presented in this thesis, my co-authors and I have advanced the state of the art in infinite games on multiple fronts, e.g., we laid the foundations for employing much stronger specification languages, exhibited the need for tradeoff-aware algorithms for quantitative winning conditions, presented the first in-depth study of delay games, and illustrated the benefits of lookahead in infinite games. In future work, I plan to continue our study of all these tradeoffs.

Of particular theoretical and practical interest is the Prompt-LTL optimization problem, whose exact complexity is still open. Also, we plan to revisit the prototype implementation of our tradeoff-aware algorithm for Prompt-LTL games in order to make it competitive with current algorithms solving LTL games.

Furthermore, we continue our investigation of tradeoffs in delay games. We have already proved that lookahead can be traded for quality and vice versa. The natural dimension to study next are the memory requirements of winning strategies. However, it is surprisingly non-trivial to come up with the "right" notion of finite-state strategy in such games. Based on a proposal put forth by Salzmann [Sal15] we study whether lookahead can be traded for memory and vice versa. Note that in this setting, it is not even clear that adding lookahead decreases the memory requirements in comparison to the delay-free game with the same winning condition.

Also, we are currently developing a very general framework for studying delay games, which would explain the similarities between all our results: adding delay incurs an exponential blowup in the solution complexity. Also, it will allow to obtain upper bounds on the necessary lookahead.

Finally, we plan to intensify our work on logics for hyperproperties: the discovery of a first-order logic that is equivalent to HyperLTL raised multiple questions, e.g., about similar results for extensions with knowledge operators, quantification of atomic propositions, etc.

For more avenues for further research, we refer to the respective sections of the papers constituting this thesis.

Acknowledgments

Many people had a direct or indirect role in my journey leading to this thesis.

First, I am very grateful to Bernd Finkbeiner for his constant guidance and support during my time at Saarland University. Also, he encouraged me to apply for my own funding, which allowed me to pursue my own research interests. The results of this work are presented in this thesis.

Although there is a single name on the title page, obtaining the results presented in this thesis has been a true team effort: thirteen papers are outcomes of collaborations. I am grateful to all my coauthors. I am especially proud to have collaborated with five PhD students of the Reactive Systems Group on nine different projects. It is a privilege to be a member of this group.

Also, I am thankful to Mikołaj Bojańczyk, Patricia Bouyer, Kim G. Larsen, and Nicolas Markey for hosting me during research visits in Warsaw, Paris, and Aalborg, which all left their mark on this work.

Most of the research presented in this thesis was financially supported by the German Research Foundation, first via the Transregional Collaborative Research Center "AVACS" and then by the project "TriCS". Earlier research was supported by the ESF project "Gasics" and the ERC Starting Grant "SOSNA". In particular, several of the results presented here were obtained in collaboration with other members of Gasics. I am very grateful for the opportunities that arose (and still arise) from being a member of this project.

Also, let me thank all the members of the habilitation committee, in particular Bernd Finkbeiner (mentor and reviewer), Christoph Weidenbach (scientific tutor), and Holger Herrmanns (committee chair).

Also, I am thankful to Swen Jacobs and Alexander Weinert for proof-reading the introductory part of this thesis and to Christa Schäfer for the administrative support, which allowed me to focus on my research.

Finally, I am most grateful for my family, who took this journey with me. To my parents, who supported me during my diploma studies, which started this journey. And to Nadine and Paulina, who joined later and ever since gave me the reason to continue it, but also, very crucially, provided the necessary relief.

Pau, this one is dedicated to you. I am excited to find out where your journey will take you.

List of Publications Constituting this Thesis

[1] Peter Faymonville and Martin Zimmermann. Parametric linear dynamic logic. *Information and Computation*, 253(2):237–256, 2017. Journal version of [FZ14].

[2] Martin Zimmermann. Parameterized linear temporal logics meet costs: Still not costlier than LTL. *Acta Informatica*, 2017. Article in Press. Journal version of [Zim15b].

[3] Alexander Weinert and Martin Zimmermann. Visibly linear dynamic logic. In Akash Lal, S. Akshay, Saket Saurabh, and Sandeep Sen, editors, *FSTTCS 2016*, volume 65 of *LIPIcs*, pages 28:1–28:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[4] Swen Jacobs, Leander Tentrup, and Martin Zimmermann. Distributed Prompt-LTL synthesis. In Domenico Cantone and Giorgio Delzanno, editors, *GandALF 2016*, volume 226 of *EPTCS*, pages 228–241, 2016.

[5] Bernd Finkbeiner and Martin Zimmermann. The First-Order Logic of Hyperproperties. In Heribert Vollmer and Brigitte Vallée, editors, *STACS 2017*, volume 66 of *LIPIcs*, pages 30:1–30:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[6] Florian Horn, Wolfgang Thomas, Nico Wallmeier, and Martin Zimmermann. Optimal strategy synthesis for request-response games. *RAIRO - Theor. Inf. and Applic.*, 49(3):179–203, 2015.

[7] Patricia Bouyer, Piotr Hofman, Nicolas Markey, Mickael Randour, and Martin Zimmermann. Bounding average-energy games. In Javier Esparza and Andrzej S. Murawski, editors, *FOSSACS 2017*, volume 10203 of *LNCS*, pages 179–195. Springer, 2017.

[8] Nathanaël Fijalkow and Martin Zimmermann. Parity and Streett games with costs. *LMCS*, 10(2), 2014. Journal version of [FZ12].

[9] Alexander Weinert and Martin Zimmermann. Easy to win, hard to master: Optimal strategies in parity games with costs. In Jean-Marc Talbot and Laurent Regnier, editors, *CSL 2016*, volume 62 of *LIPIcs*, pages 31:1–31:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[10] Leander Tentrup, Alexander Weinert, and Martin Zimmermann. Approximating optimal bounds in Prompt-LTL realizability in doubly-exponential time. In Domenico Cantone and Giorgio Delzanno, editors, *GandALF 2016*, volume 226 of *EPTCS*, pages 302–315, 2016.

[11] Wladimir Fridman, Christof Löding, and Martin Zimmermann. Degrees of lookahead in context-free infinite games. In Marc Bezem, editor, *CSL 2011*, volume 12 of *LIPIcs*, pages 264–276. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2011.

[12] Felix Klein and Martin Zimmermann. How much lookahead is needed to win infinite games? *LMCS*, 12(3), 2016. Journal version of [KZ15].

[13] Felix Klein and Martin Zimmermann. What are strategies in delay games? Borel determinacy for games with lookahead. In Stephan Kreutzer, editor, *CSL 2015*, volume 41 of *LIPIcs*, pages 519–533. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.

[14] Martin Zimmermann. Delay games with WMSO+U winning conditions. *RAIRO - Theor. Inf. and Applic.*, 50(2):145–165, 2016. Journal version of [Zim15a].

[15] Felix Klein and Martin Zimmermann. Prompt delay. In Akash Lal, S. Akshay, Saket Saurabh, and Sandeep Sen, editors, *FSTTCS 2016*, volume 65 of *LIPIcs*, pages 43:1–43:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[16] Martin Zimmermann. Games with costs and delays. *arXiv*, 1701.02168, 2017. Accepted for publication at LICS 2017.

References

[AELP01] Rajeev Alur, Kousha Etessami, Salvatore La Torre, and Doron Peled. Parametric temporal logic for "model measuring". *ACM Trans. Comput. Log.*, 2(3):388–407, 2001.

[AEM04] Rajeev Alur, Kousha Ettesami, and Parthasarathy Madhusudan. A temporal logic of nested calls and returns. In *TACAS 2004*, volume 2988 of *LNCS*, pages 467–481. Springer, 2004.

[AFF+02] Roy Armoni, Limor Fix, Alon Flaisher, Rob Gerth, Boris Ginsburg, Tomer Kanza, Avner Landver, Sela Mador-Haim, Eli Singerman, Andreas Tiemeyer, Moshe Y. Vardi, and Yael Zbar. The ForSpec temporal logic: A new temporal property-specification language. In Joost-Pieter Katoen and Perdita Stevens, editors, *TACAS 2002*, volume 2280 of *LNCS*, pages 296–311. Springer, 2002.

[AM04] Rajeev Alur and Parthasarathy Madhusudan. Visibly Pushdown Languages. In *STOC 2004*, pages 202–211. ACM, 2004.

[BC06] Mikołaj Bojańczyk and Thomas Colcombet. Bounds in ω-regularity. In *LICS 2006*, pages 285–296. IEEE Computer Society, 2006.

[BCD+07] Gerd Behrmann, Agnès Cougnard, Alexandre David, Emmanuel Fleury, Kim Guldstrand Larsen, and Didier Lime. UPPAAL-Tiga: Time for playing games! In Werner Damm and Holger Hermanns, editors, *CAV 2007*, volume 4590 of *LNCS*, pages 121–125. Springer, 2007.

[BCJ+97] Anca Browne, Edmund M. Clarke, Somesh Jha, David E. Long, and Wilfredo R. Marrero. An improved algorithm for the evaluation of fixpoint expressions. *Theor. Comput. Sci.*, 178(1-2):237–255, 1997.

[BCK+14] Achim Blumensath, Thomas Colcombet, Denis Kuperberg, Pawel Parys, and Michael Vanden Boom. Two-way cost automata and cost logics over infinite trees. In Thomas A. Henzinger and Dale Miller, editors, *CSL-LICS 2014*, pages 16:1–16:9. ACM, 2014.

[BDHK06] Dietmar Berwanger, Anuj Dawar, Paul Hunter, and Stephan Kreutzer. DAG-width and parity games. In Bruno Durand and Wolfgang Thomas, editors, *STACS 2006*, volume 3884 of *LNCS*, pages 524–536. Springer, 2006.

[BFL+08] Patricia Bouyer, Uli Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jiří Srba. Infinite runs in weighted timed automata with energy constraints. In Franck Cassez and Claude Jard, editors, *FORMATS 2008*, volume 5215 of *LNCS*, pages 33–47. Springer, 2008.

[BGMS14] Mikołaj Bojańczyk, Tomasz Gogacz, Henryk Michalewski, and Michał Skrzypczak. On the decidability of MSO+U on infinite trees. In Esparza et al. [EFHK14], pages 50–61.

[BL69] J. Richard Büchi and Lawrence H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.

[BMR+15] Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim Guldstrand Larsen, and Simon Laursen. Average-energy games. In Javier Esparza and Enrico Tronci, editors, *GandALF 2015*, volume 193 of *EPTCS*, pages 1–15, 2015.

[Boj04] Mikołaj Bojańczyk. A bounding quantifier. In Jerzy Marcinkowski and Andrzej Tarlecki, editors, *CSL 2004*, volume 3210 of *LNCS*, pages 41–55. Springer, 2004.

[Boj11] Mikołaj Bojańczyk. Weak MSO with the unbounding quantifier. *Theory Comput. Syst.*, 48(3):554–576, 2011.

[Boj14] Mikołaj Bojańczyk. Weak MSO+U with path quantifiers over infinite trees. In Esparza et al. [EFHK14], pages 38–49.

[Boz07] Laura Bozzelli. Alternating Automata and a Temporal Fixpoint Calculus for Visibly Pushdown Languages. In L. Caires and V. T. Vasconcelos, editors, *CONCUR 2007*, volume 4703 of *LNCS*, pages 476–491. Springer, 2007.

[BPT16] Mikołaj Bojańczyk, Pawel Parys, and Szymon Toruńczyk. The MSO+U theory of $(\mathbb{N}, <)$ is undecidable. In Nicolas Ollinger and Heribert Vollmer, editors, *STACS 2016*, volume 47 of *LIPIcs*, pages 21:1–21:8. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[BS14a] Laura Bozzelli and César Sánchez. Visibly linear temporal logic. In *IJCAR 2014*, volume 8562 of *LNCS*, pages 418–483, 2014.

[BS14b] Laura Bozzelli and César Sánchez. Visibly rational expressions. *Acta Informatica*, 51(1):25–49, 2014.

[BT09] Mikołaj Bojańczyk and Szymon Toruńczyk. Deterministic automata and extensions of weak MSO. In Ravi Kannan and K. Narayan Kumar, editors, *FSTTCS 2009*, volume 4 of *LIPIcs*, pages 73–84. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2009.

[BT12] Mikołaj Bojańczyk and Szymon Toruńczyk. Weak MSO+U over infinite trees. In Christoph Dürr and Thomas Wilke, editors, *STACS 2012*, volume 14 of *LIPIcs*, pages 648–660. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012.

[BtCCV15] Michael Benedikt, Balder ten Cate, Thomas Colcombet, and Michael Vanden Boom. The complexity of boundedness for guarded logics. In *LICS 2015*. IEEE Computer Society, 2015.

[BV07] Henrik Björklund and Sergei G. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discrete Applied Mathematics*, 155(2):210–229, 2007.

[CD12] Krishnendu Chatterjee and Laurent Doyen. Energy parity games. *Theor. Comput. Sci.*, 458:49–60, 2012.

[CFK+14] Michael R. Clarkson, Bernd Finkbeiner, Masoud Koleini, Kristopher K. Micinski, Markus N. Rabe, and César Sánchez. Temporal logics for hyperproperties. In Martín Abadi and Steve Kremer, editors, *POST 2014*, volume 8414 of *LNCS*, pages 265–284. Springer, 2014.

[CH06] Krishnendu Chatterjee and Thomas A. Henzinger. Finitary winning in omega-regular games. In Holger Hermanns and Jens Palsberg, editors, *TACAS 2006*, volume 3920 of *LNCS*, pages 257–271. Springer, 2006.

[CH08] Julien Cristau and Florian Horn. On reachability games of ordinal length. In Viliam Geffert, Juhani Karhumäki, Alberto Bertoni, Bart Preneel, Pavol Návrat, and Mária Bieliková, editors, *SOFSEM 2008*, volume 4910 of *LNCS*, pages 211–221. Springer, 2008.

[CH12] Krishnendu Chatterjee and Thomas A. Henzinger. A survey of stochastic ω-regular games. *J. Comput. Syst. Sci.*, 78(2):394–413, 2012.

[CHH09] Krishnendu Chatterjee, Thomas A. Henzinger, and Florian Horn. Finitary winning in ω-regular games. *ACM Trans. Comput. Log.*, 11(1), 2009.

[CHH11] Krishnendu Chatterjee, Thomas A. Henzinger, and Florian Horn. The complexity of request-response games. In Adrian Horia Dediu, Shunsuke Inenaga, and Carlos Martín-Vide, editors, *LATA 2011*, volume 6638 of *LNCS*, pages 227–237. Springer, 2011.

[Chu63] Alonzo Church. Logic, arithmetic, and automata. In *Proc. Int. Congr. Math. 1962*, pages 23–35. Inst. Mittag-Lefler, Djursholm, Sweden, 1963.

[CJK+16] Cristian S. Calude, Sanjay Jain, Bakhadyr Khoussainov, Wei Li, and Frank Stephan. Deciding parity games in quasipolynomial time. Technical report, Centre for Discrete Mathematics and Theoretical Computer Science, 2016. To appear at FOCS 2017.

[CKL10] Thomas Colcombet, Denis Kuperberg, and Sylvain Lombardy. Regular temporal cost functions. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *ICALP 2010, Part II*, volume 6199 of *LNCS*, pages 563–574. Springer, 2010.

[CL10] Thomas Colcombet and Christof Löding. Regular cost functions over finite trees. In *LICS 2010*, pages 70–79. IEEE Computer Society, 2010.

[Col13] Thomas Colcombet. Regular cost functions, part I: logic and algebra over words. *LMCS*, 9(3), 2013.

[CRR14] Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. Strategy synthesis for multi-dimensional quantitative objectives. *Acta Informatica*, 51(3-4):129–163, 2014.

[dAH00] Luca de Alfaro and Thomas A. Henzinger. Concurrent omega-regular games. In *LICS 2000*, pages 141–154. IEEE Computer Society, 2000.

[DGL16] Stéphane Demri, Valentin Goranko, and Martin Lange. *Temporal Logics in Computer Science*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016.

[DR11] L. Doyen and J.-F. Raskin. Games with imperfect information: Theory and algorithms. In *Lectures in Game Theory for Computer Scientists*, pages 185–212. Cambridge University Press, 2011.

[EF06] Cindy Eisner and Dana Fisman. *A Practical Introduction to PSL*. Integrated Circuits and Systems. Springer, 2006.

[EFHK14] Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors. *ICALP 2014, Part II*, volume 8573 of *LNCS*. Springer, 2014.

[EJ91] E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *FOCS 1991*, pages 368–377. IEEE Computer Society, 1991.

[EJS01] E. Allen Emerson, Charanjit S. Jutla, and A. Prasad Sistla. On model checking for the μ-calculus and its fragments. *Theor. Comput. Sci.*, 258(1-2):491–522, 2001.

[EL86] E. Allen Emerson and Chin-Laung Lei. Efficient model checking in fragments of the propositional mu-calculus (extended abstract). In *LICS 1986*, pages 267–278. IEEE Computer Society, 1986.

[EWS05] Kousha Etessami, Thomas Wilke, and Rebecca A. Schuller. Fair simulation relations, parity games, and state space reduction for Büchi automata. *SIAM J. Comput.*, 34(5):1159–1175, 2005.

[FH16] Bernd Finkbeiner and Christopher Hahn. Deciding hyperproperties. In Josée Desharnais and Radha Jagadeesan, editors, *CONCUR 2016*, volume 59 of *LIPIcs*, pages 13:1–13:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.

[FHKS15] Nathanaël Fijalkow, Florian Horn, Denis Kuperberg, and Michał Skrzypczak. Trading bounds for memory in games with counters. In Halldórsson et al. [HIKS15], pages 197–208.

[FJS$^+$17] John Fearnley, Sanjay Jain, Sven Schewe, Frank Stephan, and Dominik Wojtczak. An ordered approach to solving parity games in quasi polynomial time and quasi linear space. Technical report, arXiv:1703.01296, 2017.

[FRS15] Bernd Finkbeiner, Markus N. Rabe, and César Sánchez. Algorithms for model checking HyperLTL and HyperCTL*. In Daniel Kroening and Corina S. Pasareanu, editors, *CAV 2015, Part I*, volume 9206 of *LNCS*, pages 30–48. Springer, 2015.

[FS05] Bernd Finkbeiner and Sven Schewe. Uniform distributed synthesis. In *LICS 2005*, pages 321–330. IEEE Computer Society, 2005.

[FS13a] John Fearnley and Sven Schewe. Time and space results for parity games with bounded treewidth. *LMCS*, 9(2), 2013.

[FS13b] Bernd Finkbeiner and Sven Schewe. Bounded synthesis. *STTT*, 15(5-6):519–539, 2013.

[FZ12] Nathanaël Fijalkow and Martin Zimmermann. Cost-parity and cost-Streett games. In Deepak D'Souza, Telikepalli Kavitha, and Jaikumar Radhakrishnan, editors, *FSTTCS 2012*, volume 18 of *LIPIcs*, pages 124–135. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012.

[FZ14] Peter Faymonville and Martin Zimmermann. Parametric linear dynamic logic. In Adriano Peron and Carla Piazza, editors, *GandALF 2014*, volume 161 of *EPTCS*, pages 60–73, 2014.

[GH82] Yuri Gurevich and Leo Harrington. Trees, automata, and games. In *STOC*, pages 60–65. ACM, 1982.

[GPSS80] Dov M. Gabbay, Amir Pnueli, Saharon Shelah, and Jonathan Stavi. On the temporal basis of fairness. In Paul W. Abrahams, Richard J. Lipton, and Stephen R. Bourne, editors, *POPL 1980*, pages 163–173. ACM Press, 1980.

[Grä02] Erich Grädel. Model checking games. *Electr. Notes Theor. Comput. Sci.*, 67:15–34, 2002.

[GS53] David Gale and Frank M. Stewart. Infinite games with perfect information. *Annals of Mathematics*, 28:245–266, 1953.

[GTW02] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *LNCS*. Springer, 2002.

[GU08] Erich Grädel and Michael Ummels. Solution concepts and algorithms for infinite multiplayer games. In Krzysztof Apt and Robert van Rooij, editors, *New Perspectives on Games and Interaction*, volume 4 of *Texts in Logic and Games*, pages 151–178. Amsterdam University Press, 2008.

[HIKS15] Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors. *ICALP 2015*, volume 9135 of *LNCS*. Springer, 2015.

[HKLN12] Keijo Heljanko, Misa Keinänen, Martin Lange, and Ilkka Niemelä. Solving parity games by a reduction to SAT. *J. Comput. Syst. Sci.*, 78(2):430–440, 2012.

[HKT12] Michael Holtmann, Łukasz Kaiser, and Wolfgang Thomas. Degrees of lookahead in regular infinite games. *LMCS*, 8(3), 2012.

[HL72] Frederick A. Hosch and Lawrence H. Landweber. Finite delay solutions for sequential conditions. In *ICALP 1972*, pages 45–60, 1972.

[Hor05] Florian Horn. Streett games on finite graphs. In *GDV 2005*, 2005. Available at `https://www.irif.fr/~horn/publications.html`.

[Hor07] Florian Horn. Faster algorithms for finitary games. In Orna Grumberg and Michael Huth, editors, *TACAS 2007*, volume 4424 of *LNCS*, pages 472–484. Springer, 2007.

[HS12] Szczepan Hummel and Michał Skrzypczak. The topological complexity of MSO+U and related automata models. *Fundam. Inform.*, 119(1):87–111, 2012.

[HTW08] Florian Horn, Wolfgang Thomas, and Nico Wallmeier. Optimal strategy synthesis in request-response games. In Sung Deok Cha, Jin-Young Choi, Moonzoo Kim, Insup Lee, and Mahesh Viswanathan, editors, *ATVA 2008*, volume 5311 of *LNCS*, pages 361–373. Springer, 2008.

[JL17] Marcin Jurdziński and Ranko Lazic. Succinct progress measures for solving parity games. Technical report, arXiv:1702.05051, 2017. To appear at LICS 2017.

[JLR13] Line Juhl, Kim Guldstrand Larsen, and Jean-François Raskin. Optimal bounds for multiweighted and parametrised energy games. In Zhiming Liu, Jim Woodcock, and Yunshan Zhu, editors, *Theories of Programming and Formal Methods – Essays Dedicated to Jifeng He on the Occasion of His 70th Birthday*, volume 8051 of *LNCS*, pages 244–255. Springer, 2013.

[JLS15] Marcin Jurdziński, Ranko Lazic, and Sylvain Schmitz. Fixed-dimensional energy games are in pseudo-polynomial time. In Halldórsson et al. [HIKS15], pages 260–272.

[JM99] Petr Jancar and Faron Moller. Techniques for decidability and undecidability of bisimilarity. In Jos C. M. Baeten and Sjouke Mauw, editors, *CONCUR 1999*, volume 1664 of *LNCS*, pages 30–45. Springer, 1999.

[JPZ08] Marcin Jurdziński, Mike Paterson, and Uri Zwick. A deterministic subexponential algorithm for solving parity games. *SIAM J. Comput.*, 38(4):1519–1532, 2008.

[JTZ17] Swen Jacobs, Leander Tentrup, and Martin Zimmermann. Distributed Prompt-LTL synthesis. 2017. Extended version of [4], invited to special issue of *Information and Computation* dedicated to *GandALF 2016*.

[Jur98] Marcin Jurdziński. Deciding the winner in parity games is in UP $\cap$ CO-UP. *Inf. Process. Lett.*, 68(3):119–124, 1998.

[Jur00] Marcin Jurdziński. Small progress measures for solving parity games. In Horst Reichel and Sophie Tison, editors, *STACS 2000*, volume 1770 of *LNCS*, pages 290–301. Springer, 2000.

[Kam68] Hans W. Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, Computer Science Department, University of California at Los Angeles, USA, 1968.

[KPV09] Orna Kupferman, Nir Piterman, and Moshe Y. Vardi. From liveness to promptness. *Form. Method. Syst. Des.*, 34(2):83–103, 2009.

[KZ15] Felix Klein and Martin Zimmermann. How much lookahead is needed to win infinite games? In Halldórsson et al. [HIKS15], pages 452–463.

[Lan14] Martin Lang. Resource reachability games on pushdown graphs. In Anca Muscholl, editor, *FOSSACS 2014*, volume 8412 of *LNCS*, pages 195–209. Springer, 2014.

[LLZ16] Kim G. Larsen, Simon Laursen, and Martin Zimmermann. Limit your consumption! Finding bounds in average-energy games. In Mirco Tribastone and Herbert Wiklicky, editors, *QAPL 2016*, volume 227 of *EPTCS*, pages 1–14, 2016.

[LS07] Martin Leucker and César Sanchéz. Regular linear temporal logic. In C. B. Jones, Z. Liu, and J. Woodcock, editors, *ICTAC 2007*, number 4711 in LNCS, pages 291 – 305, 2007.

[Lud95] Walter Ludwig. A subexponential randomized algorithm for the simple stochastic game problem. *Inf. and Comput.*, 117(1):151–155, 1995.

[Mar75] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.

[McN93] Robert McNaughton. Infinite games played on finite graphs. *Ann. Pure Appl. Logic*, 65(2):149–184, 1993.

[McN00] Robert McNaughton. Playing infinite games in finite time. In Arto Salomaa, Derick Wood, and Sheng Yu, editors, *A Half-Century of Automata Theory*, pages 73–91. World Scientific, 2000.

[Mos91] Andrzej Mostowski. Games with forbidden positions. Technical Report 78, University of Gdańsk, 1991.

[MW03] Swarup Mohalik and Igor Walukiewicz. Distributed games. In Paritosh K. Pandya and Jaikumar Radhakrishnan, editors, *FSTTCS 2003*, volume 2914 of *LNCS*, pages 338–351. Springer, 2003.

[Obd03] Jan Obdrzálek. Fast mu-calculus model checking when tree-width is bounded. In Warren A. Hunt Jr. and Fabio Somenzi, editors, *CAV 2003*, volume 2725 of *LNCS*, pages 80–92. Springer, 2003.

[Pnu77] Amir Pnueli. The temporal logic of programs. In *FOCS 1977*, pages 46–57. IEEE Computer Society, 1977.

[PR89a] Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *POPL 1989*, pages 179–190. ACM Press, 1989.

[PR89b] Amir Pnueli and Roni Rosner. On the synthesis of an asynchronous reactive module. In Giorgio Ausiello, Mariangiola Dezani-Ciancaglini, and Simona Ronchi Della Rocca, editors, *ICALP 1989*, volume 372 of *LNCS*, pages 652–671. Springer, 1989.

[PR90] Amir Pnueli and Roni Rosner. Distributed reactive systems are hard to synthesize. In *FOCS 1990, Volume II*, pages 746–757. IEEE Computer Society, 1990.

[Rab69] Michael O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141(1-35):4, 1969.

[Ros91] Roni Rosner. *Modular Synthesis of Reactive Systems*. PhD thesis, Weizmann Institute of Science, February 1991.

[Sal15] Tobias Salzmann. How much memory is needed to win regular delay games? Master's thesis, Saarland University, 2015.

[Sch17] Sven Schewe. Solving parity games in big steps. *J. Comput. Syst. Sci.*, 84:243–262, 2017.

[Sha53] L. S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100, 1953.

[Str81] Robert S. Streett. Propositional dynamic logic of looping and converse. In *STOC 1981*, pages 375 – 383. ACM, 1981.

[TB73] Boris A. Trakhtenbrot and Ya. M. Barzdin. *Finite Automata; Behavior and Synthesis*. Fundamental Studies in Computer Science, V. 1. North-Holland Publishing Company; New York: American Elsevier, 1973.

[Tho90] Wolfgang Thomas. Automata on infinite objects. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (B)*, pages 133–192. MIT Press, 1990.

[TL93] Wolfgang Thomas and Helmut Lescow. Logical specifications of infinite computations. In J. W. de Bakker, Willem P. de Roever, and Grzegorz Rozenberg, editors, *A Decade of Concurrency, Reflections and Perspectives, REX School/Symposium*, volume 803 of *LNCS*, pages 583–621. Springer, 1993.

[Van11] Michael Vanden Boom. Weak cost monadic logic over infinite trees. In Filip Murlak and Piotr Sankowski, editors, *MFCS*, volume 6907 of *LNCS*, pages 580–591. Springer, 2011.

[Var11] Moshe Vardi. The Rise and Fall of LTL. In Giovanna D'Agostino and Salvatore La Torre, editors, *GandALF 2011*, volume 54 of *EPTCS*, 2011.

[VCD+15] Yaron Velner, Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Alexander Rabinovich, and Jean-François Raskin. The complexity of multi-mean-payoff and multi-energy games. *Information and Computation*, 241:177–196, 2015.

[VJ00] Jens Vöge and Marcin Jurdziński. A discrete strategy improvement algorithm for solving parity games. In E. Allen Emerson and A. Prasad Sistla, editors, *CAV 2000*, volume 1855 of *LNCS*, pages 202–215. Springer, 2000.

[VW94] Moshe Vardi and Pierre Wolper. Reasoning about infinite computations. *Inf. and Comp.*, 115:1–37, 1994.

[Wad72] William W. Wadge. Degrees of complexity of subsets of the Baire space. *Notices A.M.S.*, pages A714–A715, 1972.

[Wad83] William W. Wadge. *Reducibility and determinateness on the Baire space*. PhD thesis, University of California, Berkeley, 1983.

[Wal01] Igor Walukiewicz. Pushdown processes: Games and model-checking. *Inf. and Comput.*, 164(2):234–263, 2001.

[Wal08] Nico Wallmeier. *Strategien in unendlichen Spielen mit Liveness-Gewinnbedingungen: Syntheseverfahren, Optimierung und Implementierung*. PhD thesis, RWTH Aachen University, 2008.

[WHT03] Nico Wallmeier, Patrick Hütten, and Wolfgang Thomas. Symbolic synthesis of finite-state controllers for request-response specifications. In Oscar H. Ibarra and Zhe Dang, editors, *CIAA 2003*, volume 2759 of *LNCS*, pages 11–22. Springer, 2003.

[Wol83] Pierre Wolper. Temporal logic can be more expressive. *Inf. and Cont.*, 56:72–99, 1983.

[WZ17] Alexander Weinert and Martin Zimmermann. Easy to win, hard to master: Optimal strategies in parity games with costs. 2017. Extended version of [9], submitted to *LMCS*. Available as *arXiv:1604.05543*.

[Zie98] Wieslaw Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theor. Comput. Sci.*, 200(1-2):135–183, 1998.

[Zim09] Martin Zimmermann. Time-optimal winning strategies in infinite games. Diploma thesis, RWTH Aachen University, 2009.

[Zim13] Martin Zimmermann. Optimal bounds in parametric LTL games. *Theor. Comput. Sci.*, 493:30–45, 2013.

[Zim15a] Martin Zimmermann. Delay games with WMSO+U winning conditions. In Lev D. Beklemishev and Daniil V. Musatov, editors, *CSR 2015*, volume 9139 of *LNCS*, pages 412–425. Springer, 2015.

[Zim15b] Martin Zimmermann. Parameterized linear temporal logics meet costs: Still not costlier than LTL. In Javier Esparza and Enrico Tronci, editors, *GandALF 2015*, volume 193 of *EPTCS*, pages 144–157, 2015.

[ZP96] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theor. Comput. Sci.*, 158(1&2):343–359, 1996.

Proposals for Scientific Talks According to § 11

Tradeoffs in Infinite Games

Games of infinite duration, infinite games for short, constitute one of the pillars of theoretical computer science with a wide range of applications in logics, automata theory, descriptive set theory, and synthesis. Recently, a lot of effort has been invested in turning the synthesis problem for reactive systems, which can be formulated as a game, into an optimization problem. As a simple example consider a request-response condition: the qualitative variant asks every request to be answered, the boundedness variant asks for some bound on the response time between request and responses, and the optimization variant asks for the minimal such bound.

Throughout this work a pattern became apparent: the best algorithms solving the optimization problem have worse running times than those solving the qualitative and the boundedness variant. Also, the optimal strategies computed by these algorithms are larger than those for the qualitative and the boundedness variant. However, there were no lower bounds, i.e., it was open whether playing optimally is indeed harder than just winning.

We answer this question affirmatively for several winning conditions: computing an optimal strategy can be harder than computing an arbitrary strategy and optimal strategies are necessarily larger than arbitrary ones. These results have profound consequences for the quest of turning reactive synthesis into an optimization problem. To overcome this, we report on a prototypical implementation of a tradeoff-aware synthesis algorithm.

Based on joint work with Nathanaël Fijalkow, Leander Tentrup, and Alexander Weinert presented in the habilitation thesis.

The complexity of Counting Models of Linear Temporal Logic

We determine the complexity of counting models of bounded size of specifications expressed in linear-time temporal logic.

Counting word-models is #P-complete, if the bound is given in unary, and as hard as counting accepting runs of nondeterministic polynomial space Turing machines, if the bound is given in binary.

Counting tree-models is as hard as counting accepting runs of nondeterministic exponential time Turing machines, if the bound is given in unary. For a binary encoding of the bound, the problem is at least as hard as counting accepting runs of nondeterministic exponential space Turing machines, and not harder than counting accepting runs of nondeterministic doubly-exponential time Turing machines.

Finally, counting arbitrary transition systems satisfying a formula is #P-hard and not harder than counting accepting runs of nondeterministic polynomial time Turing machines with a PSPACE oracle, if the bound is given in unary. If the bound is given in binary, then counting arbitrary models is as hard as counting accepting runs of nondeterministic exponential time Turing machines.

Based on joint work with Hazem Torfah not presented in the habilitation thesis.

Some Recent Results on Pattern Matching

Every sufficiently long word over an arbitrary finite alphabet contains a pattern of the form xyx, i.e., as soon as a letter has two non-adjacent occurrences. Similarly, every sufficiently long word over a *binary* alphabet contains a square, an infix of the form xx for some non-empty x. In fact, every word of length four has such a pattern. On the other hand, there is an infinitely long square-free word over a ternary alphabet and an infinitely long cube-free word over a binary alphabet.

Thus, the pattern xyx is unavoidable while squares and cubes are avoidable, as there are alphabets for which they are avoidable.

There is a beautiful characterization of the unavoidable and avoidable patterns via so-called Zimin patterns Z_n. These patterns are defined inductively over variables x_j as $Z_1 = x_1$ and $Z_n = Z_{n-1}x_nZ_{n-1}$ for n greater than 1. Now, Zimin and independently Bean, Ehrenfeucht and McNulty proved that a pattern over n variables is unavoidable if, and only if, it appears in the Zimin pattern Z_n. For example, xyx appears in the pattern Z_2.

A natural question concerns bounds on the maximal length $f(n,k)$ of words over an alphabet of size k that avoid the pattern Z_n. Also, this question has important connections to the equivalence problem of deterministic pushdown automata, which was famously solved by Sénizergues.

Cooper and Rorabaugh proved that $f(n,k)$ is upper-bounded by a tower of exponentials and Carayol and Göller gave a matching lower bound using Stockmeyer's yardstick construction.

We give an overview of these results, focussing on the recently presented lower bound.