

Improved Algorithms for One-Pair and k-Pair Streett Objectives

Krishnendu Chatterjee
IST Austria
krishnendu.chatterjee@ist.ac.at

Monika Henzinger
University of Vienna
Faculty of Computer Science
monika.henzinger@univie.ac.at

Veronika Loitzenbauer
University of Vienna
Faculty of Computer Science
veronika.loitzenbauer@univie.ac.at

Abstract—The computation of the winning set for one-pair Streett objectives and for k-pair Streett objectives in (standard) graphs as well as in game graphs are central problems in computer-aided verification, with application to the verification of closed systems with strong fairness conditions, the verification of open systems, checking interface compatibility, well-formedness of specifications, and the synthesis of reactive systems. We give faster algorithms for the computation of the winning set for (1) one-pair Streett objectives (aka parity-3 problem) in game graphs and (2) for k-pair Streett objectives in graphs. For both problems this represents the first improvement in asymptotic running time in 15 years.

Index Terms—Computer-aided verification; Synthesis; Graph games; Parity games; Streett automata; Graph algorithms.

I. INTRODUCTION

Game graphs and graphs. Consider a directed graph (V, E) with a partition (V_1, V_2) of V , which is called a *game graph*. Let $n = |V|$ and $m = |E|$. Two players play the following *alternating game* on the graph that forms an infinite path. They start by placing a token on an initial vertex and then take turns indefinitely in moving the token: At a vertex $v \in V_1$ player 1 moves the token along one of the outedges of v , at a vertex $u \in V_2$ player 2 moves the token along one of the outedges of u . If $V_2 = \emptyset$, then we simply have a standard graph.

Objectives and winning sets. Objectives are subsets of infinite paths that specify the desired set of paths for player 1, and the objective for player 2 is the complement of the player-1 objective (i.e., we consider *zero-sum* games). Given an objective Φ , an infinite path *satisfies the objective* if it belongs to Φ . Given a starting vertex $x \in V$ and an objective Φ , if player 1 can guarantee that the infinite path starting at x satisfies Φ , *no matter what choices player 2 makes*, then player 1 can *win from x* and x belongs to the *winning set of player 1*. Since the winning sets partition the game graph [33], the complement of the winning set for player 1 is the winning set for player 2. In case the game graph is a standard graph (i.e., $V_2 = \emptyset$), the *winning set* consists of those vertices x such that there exists an infinite path starting at x that satisfies Φ . The winning set computation for game graphs is more involved than for standard graphs due to the presence of the adversarial player 2.

Relevant objectives. The most basic objective is *reachability* where, given a set $U \subseteq V$ of vertices, an infinite path satisfies the objective if the path visits a vertex in U *at least once*. The

next interesting objective is the *Büchi* objective that requires an infinite path to visit some vertex in U *infinitely often*. The next and a very central objective in formal verification and automata theory is the *one-pair Streett objective* that consists of a pair (L_1, U_1) of sets of vertices (i.e., $L_1 \subseteq V$ and $U_1 \subseteq V$), and an infinite path satisfies the objective iff the following condition holds: if some vertex in L_1 is visited infinitely often, then some vertex in U_1 is visited infinitely often (intuitively the objective specifies that if one Büchi objective holds, then another Büchi objective must also hold). A generalization of one-pair Streett objectives is the *k-pair Streett objective* (aka *general Streett objective*) that consists of k -Streett pairs $(L_1, U_1), (L_2, U_2), \dots, (L_k, U_k)$, and an infinite path satisfies the objective iff the condition for every Streett pair is satisfied (in other words the objective is the conjunction of k one-pair Streett objectives).

We study (1) game graphs with one-pair Streett objectives and (2) graphs with general Streett objectives.

Significance in verification. Two-player games on graphs are useful in many problems in computer science, specially in verification and synthesis of systems such as the synthesis of reactive systems [12], [36], [37], verification of open systems [2], checking interface compatibility [15], well-formedness of specifications [16], and many others. General and one-pair Streett objectives are central in verification as most commonly used specifications can be expressed as Streett automata [38], [42].

Game graphs with one-pair Streett objectives arise in many applications in verification. We sketch a few of them. (A) Timed automaton games are a model for real-time systems. The analysis of such games with reachability objectives and safety objectives (which are the dual of reachability objectives) reduces to game graphs with one-pair Streett objectives [9], [11], [13], [14]. (B) The synthesis of Generalized Reactivity(1) (aka GR(1)) specifications exactly require the solution of game graphs with one-pair Streett objectives [4]; GR(1) specifications are standard for hardware synthesis [35] and even used in synthesis of industrial protocols [5], [23]¹. (C) Finally, the

¹A GR(1) specification expresses that if a conjunction of Büchi objectives holds, then another conjunction of Büchi objectives must also hold, and since conjunction of Büchi objectives can be reduced in linear time to a single Büchi objective, a GR(1) specification reduces to implication between two Büchi objectives, which is an one-pair Streett objective.

problem of fair simulation [26] between two systems also reduces to game graphs with one-pair Streett objectives [6].

General Streett objectives in standard graphs arise, for example, in the verification of closed systems with strong fairness conditions [17], [22], [31]. In program verification, a scheduler is *strongly fair* if every event that is enabled infinitely often is scheduled infinitely often. Thus, verification of systems with strong fairness conditions directly corresponds to checking the non-emptiness of Streett automata, which in turn corresponds to determining the winning set in standard graphs with Streett objectives. Note, however, that a Streett objective can either specify desired behaviors of the system or erroneous ones, and for erroneous specifications, it is useful to have a *certificate* (as defined below) to identify an error trace of the system [17], [18], [31].

Note that *standard graphs* are relevant for the verification of *closed* systems or *open* systems with demonic non-determinism (e.g., all inputs are from the environment that are not controllable); while *game graphs* are relevant for the synthesis and verification of *open* systems with both angelic and demonic non-determinism (e.g., certain inputs are controllable, and certain inputs are not controllable).

Previous results. We summarize the previous results for game graphs and graphs with Streett objectives.

Game graphs. We consider the computation of the winning set for player 1 in game graphs. For *reachability* objectives, the problem is PTIME-complete, and the computation can be achieved in time linear in the size of the graph [3], [27]. For *Büchi* objectives, the current best known algorithm requires $O(n^2)$ time [7], [8]. For *general Streett* objectives, the problem is coNP-complete [19], and for *one-pair Streett* objectives the current best known algorithm requires $O(nm)$ time [29]. One-pair Streett objectives also corresponds to the well-known parity games problem with three priorities (the parity games problem in general is in $UP \cap coUP$ [28]; it is one of the rare and intriguing combinatorial problems that lie in $UP \cap coUP$, but not known to be in PTIME). Despite the importance of game graphs with one-pair Streett objectives in numerous applications and several algorithmic ideas to improve the running time for general parity games [30], [39], [43] or Büchi games [7], [8], [10], there has been no algorithmic improvement since 2000 [29] for one-pair Streett games.

Graphs. In standard graphs we study the computation of the winning set for general Streett objectives. If x belongs to the winning set, it is often useful to output a *certificate* for x . Let S be a (not necessarily maximal) strongly connected component (SCC) that is reachable from x such that for all $1 \leq j \leq k$ we have either $S \cap L_j = \emptyset$ or $S \cap U_j \neq \emptyset$ (i.e., if S contains a vertex from L_j then it also contains a vertex from U_j). A certificate is a “lasso-shaped” path that reaches S and then visits all vertices in S infinitely often to satisfy the general Streett objective. The basic algorithm [21], [32] for the winning set problem has an asymptotic running time of $O((m+b) \min(n, k))$ with $b = \sum_{j=1}^k (|L_j| + |U_j|) \leq 2nk$. Within the same time bound Latvala and Heljanko [31] additionally compute a certificate of

size at most $n \min(n, 2k)$. Duret-Lutz et al. [17] presented a space-saving “on-the-fly” algorithm with the same time complexity for the slightly different transition-based Streett automata. The current fastest algorithm for the problem by Henzinger and Telle [25] from 1996 has a running time of $O(m \min(\sqrt{m \log n}, k, n) + b \min(\log n, k))$; however, given a start vertex x , to report the certificate for x adds an additive term of $O(n \min(n, k))$ to the running time bound.

Our contributions. In this work our contributions are two-fold.

Game graphs. We show that the winning set computation for game graphs with one-pair Streett objectives can be achieved in $O(n^{2.5})$ time. Our algorithm is faster for $m > n^{1.5}$, and breaks the long-standing $O(nm)$ barrier for dense graphs. We also discuss the implications of our algorithm for general parity games in Remark 1.

Graphs. We present an algorithm with $O(n^2 + b \log n)$ running time for the winning set computation in graphs with general Streett objectives, which is faster for $m \geq \max(n^{4/3} \log^{-1/3} n, b^{2/3} \log^{1/3} n)$ and $k \geq n^2/m$. We additionally give an algorithm that computes a certificate for a vertex x in the winning set in time $O(m + n \min(n, k))$. We also provide an example where the smallest certificate has size $\Theta(n \min(n, k))$, showing that no algorithm can compute and *output* a certificate faster. In contrast to [25] the running time of our algorithm for the winning set computation does not change with certificate reporting. Thus when certificates need to be reported and $k = \Omega(n)$, our algorithm is *optimal* up to a factor of $\log n$ as the size of the input is at least b and the size of the output is $\Omega(n^2)$.

Technical contributions. Both of our algorithms use a *hierarchical (game) graph decomposition* technique that was developed by Henzinger et al. [24] to handle *edge deletions* in *undirected* graphs. In [7], [8] it was extended to deal with *vertex deletions* in *directed* and *game* graphs. We combine and extend this technique in two ways.

Game graphs. The classical algorithm for one-pair Streett objectives repeatedly solves Büchi games such that the union of the winning sets of player 1 in the Büchi games is exactly the winning set for the one-pair Streett objective. Schewe [39] showed that an algorithm for parity games by Jurdziński [29] can be used to compute small subsets of the winning set of player 1, called *dominions*, and thereby improved the running time for general parity games. However his ideas do not improve the running time for one-pair Streett (aka *parity-3*) games. With this algorithm dominions with at most h vertices in Büchi games can be found in time $O(mh)$. We extend this approach by using the hierarchical game graph decomposition technique to find small dominions quickly and call the $O(n^2)$ Büchi game algorithm of [7], [8] for large dominions. This extension is possible as we are able to show that, rather surprisingly, it is sufficient to consider game graphs with $O(nh)$ edges to detect dominions of size h (see Lemma 2).

Graphs. In prior work that used the hierarchical graph decomposition technique the runtime analysis relied on the fact that identified vertex sets that fulfilled a certain desired

condition were *removed* from the (game) graph after their detection. The work for identifying the vertex set was then charged in an amortization argument to the removed vertex set. This is not possible for general Streett objectives on graphs, where SCCs are identified and some but not all of its vertices might be removed. As a consequence a vertex might belong to an identified SCC multiple times. We show how to overcome this difficulty by identifying, when an SCC S splits into multiple SCCs, an SCC $X \subset S$ whose size is at most half of the size of S . We identify X by using Tarjan's SCC algorithm [41] on the graph and its *reverse graph*, thereby finding the smallest *top* (i.e. with no incoming edges) or *bottom* (i.e. with no outgoing edges) SCC contained in S . The smallest such SCC X has size at most $|S|/2$ and the algorithm takes time $O(|X|n)$ to find it, i.e., $O(n)$ time per vertex in X . This will allow us to bound the total running time for this part of the algorithm with $O(n^2)$.

In Section II we present our algorithm for one-pair Streett objectives in game graphs, in Section III the algorithm for general Streett objectives in graphs.

II. ONE-PAIR STREETT OBJECTIVES IN GAME GRAPHS

A. Preliminaries

Parity games. A parity game $P = (G, \alpha)$ consists of a *game graph* $G = ((V, E), (V_{\mathcal{O}}, V_{\mathcal{E}}))$ and a *priority function* $\alpha : V \rightarrow \mathbb{Z}$ that assigns an integer value to each vertex. We denote the two players with \mathcal{O} (for odd) and \mathcal{E} (for even). Player \mathcal{O} (resp. player \mathcal{E}) wins a play if the *lowest* priority occurring in the play infinitely often is odd (resp. even). We say that the vertices in $V_{\mathcal{O}}$ are \mathcal{O} -vertices and the vertices in $V_{\mathcal{E}}$ are \mathcal{E} -vertices. We use p to denote one of the players $\{\mathcal{O}, \mathcal{E}\}$ and \bar{p} to denote his opponent. We will specifically consider *parity-3 games* with $\alpha : V \rightarrow \{-1, 0, 1\}$ and *Büchi games* with $\alpha : V \rightarrow \{0, 1\}$, where the vertices in the set $B = \{v \mid \alpha(v) = 0\}$ are called *Büchi vertices*. Büchi games are denoted as (G, B) .

One-pair Streett and parity-3 games. A one-pair Streett objective with pair (L_1, U_1) is equivalent to a parity game with three priorities. Let the vertices in U_1 have priority -1 , let the vertices in $L_1 \setminus U_1$ have priority 0 and let the remaining vertices have priority 1. Then player 1 wins the game with the one-pair Streett objective if and only if player \mathcal{O} wins the parity-3 game. As the known algorithms for parity-3 games are special cases of algorithms for general parity games, we will use the notion of parity games (i.e., player \mathcal{O} and player \mathcal{E} instead of player 1 and player 2).

Plays. For technical convenience we consider that every vertex in the game graph G has at least one outgoing edge. A game is initialized by placing a token on a vertex. Then the two players form an infinite path called *play* in the game graph by moving the token along the edges. Whenever the token is on a vertex in V_p , player p moves the token along one of the outgoing edges of the vertex. Formally, a *play* is an infinite sequence $\langle v_0, v_1, v_2, \dots \rangle$ of vertices such that $(v_j, v_{j+1}) \in E$ for all $j \geq 0$.

For a vertex $u \in V$, we write $Out(u) = \{v \in V \mid (u, v) \in E\}$ for the set of successor vertices of u and $In(u) = \{v \in V \mid$

$(v, u) \in E\}$ for the set of predecessor vertices of u . We denote by $Outdeg(u) = |Out(u)|$ the number of outgoing edges from u , and by $Indeg(u) = |In(u)|$ the number of incoming edges.

Strategies. A *strategy* of a player $p \in \{\mathcal{O}, \mathcal{E}\}$ is a function that, given a finite prefix of a play ending at $v \in V_p$, selects a vertex from $Out(v)$ to extend the play. *Memoryless strategies* do not depend on the history of a play but only on the current vertex. That is, a memoryless strategy of player p is a function $\tau : V_p \rightarrow V$ such that for all $v \in V_p$ we have $\tau(v) \in Out(v)$. It is well-known that for parity games it is sufficient to consider memoryless strategies (see Theorem 1 below). Therefore we will only consider memoryless strategies from now on. A start vertex v together with a strategy σ for \mathcal{E} and a strategy π for \mathcal{O} defines a unique play $\omega(v, \sigma, \pi) = \langle v_0, v_1, v_2, \dots \rangle$, which is defined as follows: $v_0 = v$ and for all $j \geq 0$, if $v_j \in V_{\mathcal{E}}$, then $\sigma(v_j) = v_{j+1}$, and if $v_j \in V_{\mathcal{O}}$, then $\pi(v_j) = v_{j+1}$.

Winning strategies and sets. A strategy τ is winning for player p at start vertex v if the resulting play is winning for player p irrespective of the strategy of player \bar{p} . A vertex v belongs to the *winning set* W_p of player p if player p has a winning strategy from v . By the following theorem every vertex is winning for exactly one of the two players. When required for explicit reference of a specific game graph G or specific parity game P we use $W_p(G)$ and $W_p(P)$ to refer to the winning sets.

Theorem 1 ([20], [34]): For every parity game the vertices V can be partitioned into the winning set $W_{\mathcal{E}}$ of \mathcal{E} and the winning set $W_{\mathcal{O}}$ of \mathcal{O} . There exists a memoryless winning strategy for \mathcal{E} (resp. \mathcal{O}) for all vertices in $W_{\mathcal{E}}$ (resp. $W_{\mathcal{O}}$).

The algorithmic question for parity games is to compute the set $W_{\mathcal{E}}$. We will use the following algorithm for Büchi games as a subroutine in our algorithm.

Theorem 2 ([7], [8]): Let (G, B) be a Büchi game with game graph G and Büchi vertices B . There is an algorithm $BÜCHI(G, B)$ that computes $W_{\mathcal{E}}$ in time $O(n^2)$.

For the analysis of our algorithm we further introduce the notions of *closed sets*, *attractors*, and *dominions*.

Closed sets and attractors. A set $U \subseteq V$ is *p-closed* if for all p -vertices u in U we have $Out(u) \subseteq U$ and for all \bar{p} -vertices v in U there exists a vertex $w \in Out(v) \cap U$. Note that player \bar{p} can ensure that a play that currently ends in a p -closed set never leaves the p -closed set against any strategy of player p by choosing an edge (v, w) with $w \in Out(v) \cap U$ whenever the current vertex v is in $U \cap V_{\bar{p}}$ (see also [8, Proposition 2.2]). Given a game graph G and a p -closed set X , we will denote by $G[X]$ the game graph induced by the set X of vertices.

In a game graph G , a *p-attractor* $Attr_p(U, G)$ of a set $U \subseteq V$ is the set of vertices from which player p has a strategy to reach U against all strategies of player \bar{p} . We have that $U \subseteq Attr_p(U, G)$. A p -attractor can be constructed inductively as follows: Let $R_0 = U$; and for all $i \geq 0$ let

$$R_{i+1} = R_i \cup \{v \in V_p \mid Out(v) \cap R_i \neq \emptyset\} \cup \{v \in V_{\bar{p}} \mid Out(v) \subseteq R_i\}. \quad (\ddagger)$$

Then $Attr_p(U, G) = \bigcup_{i \geq 0} R_i$. The lemma below summarizes some well-known facts about closed sets, attractors, and winning sets.

Lemma 1: Let $p \in \{\mathcal{O}, \mathcal{E}\}$ and let $U \subseteq V$. The following assertions hold for parity games.

- 1) The set $V \setminus Attr_p(U, G)$ is p -closed in G [44, Lemma 4].
- 2) Let U be p -closed. Then $Attr_{\bar{p}}(U, G)$ is p -closed [44, Lemma 5].
- 3) The attractor $Attr_p(U, G)$ can be computed in $O(\sum_{v \in Attr_p(U, G)} |In(v)|)$ time [3], [27].
- 4) Let U be a subset of the winning set $W_p(G)$ of player p and let A be its p -attractor $Attr_p(U, G)$. Then the winning set $W_p(G)$ of the player p is the union of A and the winning set $W_p(G[V \setminus A])$ in the game graph induced by $V \setminus A$, and the winning set $W_{\bar{p}}(G)$ of the opponent \bar{p} is equal to $W_{\bar{p}}(G[V \setminus A])$ [30, Lemma 4.5].

Dominions. A set of vertices $D \subseteq V$ is a p -dominion if $D \neq \emptyset$, player p has a winning strategy from every vertex in D that also ensures only vertices in D are visited, and D is a \bar{p} -closed set. We will only consider \mathcal{E} -dominions in this paper and therefore usually omit the reference to the player. Dominions of size $|D| \leq h$ can be computed by running the small-progress measure algorithm of Jurdziński [29] with a reduced codomain [39]. A description of the small-progress measure algorithm for Büchi games is given in [8, Section 2.4.1]. We will use the following algorithm as a subroutine.

Theorem 3 ([8], [29], [39]): Let (G, B) be a Büchi game with game graph G and Büchi vertices B . There is an algorithm $BÜCHIPROGRESS(G, B, h)$ that returns the set of all dominions of size at most h in time $O(mh)$.

B. Algorithm

In this section we present our new algorithm to compute the winning set of player \mathcal{E} in a parity-3 game $P = (G, \alpha)$ in time $O(n^{2.5})$. Its complement is the winning set of player \mathcal{O} .

Initialization (Steps 1–4 of Algorithm 1). First the algorithm constructs the modified game graph $G' = ((V, E'), (V_{\mathcal{O}}, V_{\mathcal{E}}))$ from G . Let Z be the vertices in V with priority -1 . In G' the vertices in Z are made *absorbing*, that is, the outgoing edges of the vertices in Z are replaced with self-loops. Otherwise G' contains the same edges as G . We will consider a Büchi game on G' where the vertices in Z have priority 1, and thus in the Büchi game there are only two priorities (priority 0 and 1). The construction of G' ensures that dominions in the Büchi game are also dominions in the parity-3 game P (see Lemma 5).

Iterated vertex deletions (Steps 5–11 of Algorithm 1). The algorithm will repeatedly remove vertices from the graphs G and G' . Initially the set V is the set of vertices in the input game graph G . During the algorithm, we denote with V the set of remaining vertices after vertex deletions and we denote with $G[V]$ and $G'[V]$ the subgraphs induced by the vertices remaining in V . The set of Büchi vertices B maintains the set of priority-0 vertices in V . The vertex set removal is achieved by identifying dominions and removing their attractors.

Input a game graph $G = ((V, E), (V_{\mathcal{O}}, V_{\mathcal{E}}))$ and a priority function $\alpha : V \rightarrow \{-1, 0, 1\}$

Output the winning set $W_{\mathcal{E}}$ of player \mathcal{E}

```

1:  $Z = \{v \in V \mid \alpha(v) = -1\}$ 
2:  $E' = \{(u, u) \mid u \in Z\} \cup \{(u, v) \in E \mid u \in V \setminus Z\}$ 
3:  $G' = (V, E')$   $\triangleright$  vertices with  $\alpha = -1$  are absorbing in  $G'$ 
4:  $W \leftarrow \emptyset, B \leftarrow \{v \in V \mid \alpha(v) = 0\}$ 
5: repeat
6:    $D \leftarrow \text{BÜCHIDOM}(G'[V], B, \sqrt{n})$ 
7:   if  $D = \emptyset$  then
8:      $D \leftarrow \text{BÜCHI}(G'[V], B)$ 
9:    $A \leftarrow \text{Attr}_{\mathcal{E}}(D, G'[V]); W \leftarrow W \cup A$ 
10:   $V \leftarrow V \setminus A; B \leftarrow B \setminus A$ 
11: until  $D = \emptyset$ 
12: return  $W$ 

13: procedure  $\text{BÜCHIDOM}(G'[V], B, h_{\max})$ 
14:   for  $i \leftarrow 1, \dots, \lceil \log(h_{\max}) \rceil$  do
15:     construct  $G'_i$ 
16:      $Bl_i \leftarrow \{v \in V_{\mathcal{O}} \mid \text{Outdeg}(v) > 2^i\}$ 
17:      $Y_i \leftarrow \text{Attr}_{\mathcal{O}}(Bl_i, G'_i)$ 
18:      $D_i \leftarrow \text{BÜCHIPROGRESS}(G'_i[V \setminus Y_i], B \setminus Y_i, 2^i)$ 
19:     if  $D_i \neq \emptyset$  then
20:       return union of dominions in  $D_i$ 
21:   return  $\emptyset$ 

```

Fig. 1. New Algorithm for Parity-3 aka One-Pair Streett Objective

Dominion find and attractor removal. The algorithm repeatedly finds dominions in the Büchi game $(G'[V], B)$. After a dominion in the Büchi game $G'[V]$ is found, its \mathcal{E} -attractor in $G'[V]$ is removed from V and B . Then the search for dominions is continued on the remaining vertices. If all vertices in the Büchi game are winning for \mathcal{O} , i.e., no dominion exists in the Büchi game, then Algorithm 1 terminates. The winning set of player \mathcal{E} is the union of the \mathcal{E} -attractors of all found dominions. The remaining vertices are winning for player \mathcal{O} . We now describe the steps to find dominions.

Steps of dominion find. For the search for dominions in the Büchi game $(G'[V], B)$ we use two different algorithms, BÜCHI and BÜCHIDOM . We first search for “small” dominions with up to $O(h_{\max})$ vertices with $h_{\max} = \sqrt{n}$ with Procedure BÜCHIDOM . If no dominion is found, we can conclude that either all dominions contain more than \sqrt{n} vertices or the winning set of \mathcal{E} on the current game graph is empty (in this case the algorithm terminates). The former case occurs at most \sqrt{n} times and in such a case we use the $O(n^2)$ algorithm BÜCHI (Theorem 2) to obtain a dominion. Below we describe the details of BÜCHIDOM .

Graph decomposition for BÜCHIDOM . In the Procedure BÜCHIDOM we use the following graph decomposition. For a game graph $G' = ((V, E'), (V_{\mathcal{O}}, V_{\mathcal{E}}))$ we denote its decomposition with $\{G'_i\}$. We consider the incoming edges of each vertex in E' in a fixed order: First the edges from

vertices in $V_{\mathcal{E}}$, then the remaining edges. We define $\lceil \log n \rceil$ graphs $G'_i = (V, E'_i)$, $1 \leq i \leq \lceil \log n \rceil$, where the set of edges E'_i contains for each vertex $v \in V$ with $\text{Outdeg}(v) \leq 2^i$ all its outgoing edges in E' and in addition for each vertex $v \in V$ its first 2^i incoming edges in E' . Note that (1) $E'_i \subseteq E'_{i+1}$, (2) $|E'_i| \leq 2^{i+1}n$, and (3) $G'_{\lceil \log n \rceil} = G'$. We color \mathcal{O} -vertices v with $\text{Outdeg}(v) > 2^i$ blue in G'_i and denote the set of blue vertices with Bl_i . We call vertices with $\text{Outdeg}(v) \leq 2^i$ white. *Procedure BÜCHIDOM (Steps 13–21 of Algorithm 1).* The Procedure BÜCHIDOM searches for dominions in the subgraphs G'_i , starting at $i = 1$. The index i is increased one by one up to at most $i = \lceil \log(h_{\max}) \rceil$ (with $h_{\max} = \sqrt{n}$) as long as no dominion was found. Let Y_i be the \mathcal{O} -attractor of blue vertices in G'_i , i.e., of \mathcal{O} -vertices that are missing outgoing edges in G'_i . To ensure that dominions found in the subgraph G'_i are also dominions in G' , only the vertices in $V \setminus Y_i$ are considered. The BÜCHIPROGRESS algorithm (Theorem 3) is used to find dominions of size at most $O(2^i)$ in $G'_i[V \setminus Y_i]$.

The following key lemma describes the central connection between dominions of a certain size and our graph decomposition. Namely, if a dominion D is found in G'_i but not in G'_{i-1} , then $\text{Attr}_{\mathcal{E}}(D, G')$ contains more than 2^{i-1} vertices. This has the remarkable consequence, detailed in Corollary 1, that every dominion of size h can be found by searching for a dominion in G'_i with $i = \lceil \log(h) \rceil$. This will be crucial for our runtime analysis.

Lemma 2: Let $G' = ((V, E'), (V_{\mathcal{O}}, V_{\mathcal{E}}))$ be a game graph and $\{G'_i\}$ its graph decomposition. For $1 \leq i \leq \lceil \log n \rceil$ we define the following sets: the set of blue vertices $Bl_i = \{v \in V_{\mathcal{O}} \mid \text{Outdeg}(v) > 2^i\}$, the attractor of blue vertices $Y_i = \text{Attr}_{\mathcal{O}}(Bl_i, G'_i)$, and the set of dominions $\mathcal{D}_i = \text{BÜCHIPROGRESS}(G'_i[V \setminus Y_i], B \setminus Y_i, 2^i)$. If a dominion D is contained in \mathcal{D}_i but not in \mathcal{D}_{i-1} , then $\text{Attr}_{\mathcal{E}}(D, G')$ contains more than 2^{i-1} vertices.

Proof: We distinguish three cases:

Case 1: The dominion D contains more than 2^{i-1} vertices. This situation might arise as Procedure BÜCHIPROGRESS($G'_{i-1}[V \setminus Y_{i-1}]$, $B \setminus Y_{i-1}$, 2^{i-1}) only guarantees to detect dominions of size at most 2^{i-1} . In this case the lemma is satisfied trivially.

Case 2: The dominion D contains a vertex $v \in V_{\mathcal{O}}$ that is blue in G'_{i-1} , i.e., an \mathcal{O} -vertex with more than 2^{i-1} outgoing edges. Since D is \mathcal{O} -closed, we have $\text{Out}(v) \subseteq D$. Thus $|\text{Attr}_{\mathcal{E}}(D, G')| \geq |D| > 2^{i-1}$ in this case.

Case 3: All vertices $v \in V_{\mathcal{O}}$ in D are white in G'_{i-1} and thus the outgoing edges of the \mathcal{O} -vertices in D are the same in G'_{i-1} and G'_i . There are two subcases.

Case 3a: All edges (u, v) from vertices $u \in V_{\mathcal{E}} \cap D$ to vertices $v \in D$ that are present in G'_i are also present in G'_{i-1} . Let σ be the winning strategy of \mathcal{E} for the vertices in D found in G'_i . This implies that (i) D is \mathcal{O} -closed in G'_{i-1} and (ii) all edges (u, v) with $u \in D \cap V_{\mathcal{E}}$ and $v = \sigma(u)$ are contained in G'_{i-1} . Thus σ is also a winning strategy of \mathcal{E} for the vertices in D in G'_{i-1} . Hence the set D is a dominion in G'_{i-1} . Thus either Case 1 applies or the dominion would already have been detected in iteration $i - 1$, a contradiction.

Case 3b: There exists a vertex $u \in V_{\mathcal{E}} \cap D$ that has an outgoing edge (u, v) to a vertex $v \in D$ in G'_i but not in G'_{i-1} . This implies $\text{Indeg}(v) > 2^{i-1}$. By the ordering of the incoming edges and the fact that $u \in V_{\mathcal{E}}$, at least 2^{i-1} edges in $\text{In}(v)$ emanate from vertices in $V_{\mathcal{E}}$. By the definition of an attractor, all these vertices are contained in $\text{Attr}_{\mathcal{E}}(D, G')$. Thus we have $|\text{Attr}_{\mathcal{E}}(D, G')| > 2^{i-1}$ as required. ■

Corollary 1: Let G' , $\{G'_i\}$, Bl_i , Y_i , and \mathcal{D}_i be defined as in Lemma 2. Let D be a dominion in G' with $D = \text{Attr}_{\mathcal{E}}(D, G')$ and $h = |D|$. Then for $i = \lceil \log(h) \rceil$ the set of dominions \mathcal{D}_i contains D .

Proof: By the definition of i we have $2^{i-1} < h \leq 2^i$. Assume by contradiction that \mathcal{D}_i does not contain D . Since $G' = G'_{\lceil \log n \rceil}$ and $Y_{\lceil \log n \rceil} = \emptyset$, by Theorem 3 there exists some $i' \in \mathbb{N}$ with $i < i' \leq \lceil \log n \rceil$, such that $\mathcal{D}_{i'}$ contains D . Let i^* be the smallest i' such that $\mathcal{D}_{i'}$ contains D . Note that $i^* \geq i + 1$. We have that $D \notin \mathcal{D}_{i^*-1}$. By Lemma 2 this implies $|\text{Attr}_{\mathcal{E}}(D, G')| > 2^{i^*-1} \geq 2^i$, a contradiction to $h \leq 2^i$. ■

Corollary 2: Either the Procedure BÜCHIDOM($G'[V]$, B , h_{\max}) returns a dominion or every dominion D in $G'[V]$ with $D = \text{Attr}_{\mathcal{E}}(D, G'[V])$ has size greater than h_{\max} .

In the runtime analysis we will additionally use the following lemma, which follows from the inductive construction of attractors.

Lemma 3: Let the game graphs G and G' and the vertex set V be defined as in Algorithm 1. Then for a player $p \in \{\mathcal{E}, \mathcal{O}\}$ and every set $U \subseteq V$ it holds that $\text{Attr}_p(U, G'[V]) \subseteq \text{Attr}_p(U, G[V])$.

Proof: Let us consider the attractor computation $\text{Attr}_p(U, G'[V])$ and $\text{Attr}_p(U, G[V])$ as defined in (§), and let us call the respective sequences as R'_i and R_i respectively. By the definition of G' for every vertex v in V either (1) $\text{Out}'(v) = \text{Out}(v)$ or (2) $\text{Out}'(v) = \{v\}$. It is straightforward to prove by induction that $R'_i \subseteq R_i$ and the desired result follows. ■

Lemma 4 (Runtime): Algorithm 1 can be implemented in $O(n^{2.5})$ time.

Proof: Algorithm 1 can be initialized in $O(m)$ time as the graph G' and the set B can be constructed from G in linear time. Note that the number of edges m' in G' is at most the number of edges m in G .

For the operations in the repeat-until loop we analyze the total running time over all iterations of the loop. The runtime analysis relies on the fact that when a dominion D is identified, the vertices in $\text{Attr}_{\mathcal{E}}(D, G)$ and their incident edges are removed from G and G' . In combination with Corollary 2, this ensures that BÜCHI is called at most $O(n/h_{\max})$ times. By Theorem 2 one call to BÜCHI takes time $O(n^2)$. With $h_{\max} = \sqrt{n}$ we obtain a total time spent in BÜCHI of $O(n^{2.5})$.

To analyze the total time spent in BÜCHIDOM, we first show how to efficiently construct the graph decomposition $\{G'_i\}$ of G' . We maintain the following data structure for G' over all iterations of Algorithm 1. At each vertex v of G' we maintain (a) a sorted list of inedges $\text{In}(v)$, and (b) a list of outedges $\text{Out}(v)$. Additionally we maintain for each edge

(u, v) a pointer to its position in the inlist of v and in the outlist of u . This allows us to update the data structure in time proportional to the degree of v when a vertex v is removed. As each vertex can be deleted at most once, the total time to update this data structure is bounded by $O(m)$. We next analyze the time needed per iteration i of the for-loop in BÜCHIDOM. Given the above data structure, the graph G'_i , the set of blue vertices Bl_i , and the attractor $Y_i = Attr_{\mathcal{O}}(Bl_i, G'_i)$ can be constructed in time $O(n \cdot 2^i)$. By Theorem 3 the time for one call of the subroutine BÜCHI PROGRESS($\cdot, \cdot, 2^i$) on graph G'_i , is $O(n \cdot 2^i \cdot 2^i) = O(n \cdot 2^{2i})$.

Let i^* be the iteration at which Procedure BÜCHIDOM stops after it is called by Algorithm 1. The runtime for this call to Procedure BÜCHIDOM from $i = 1$ to i^* forms a geometric series that is bounded by $O(n \cdot 2^{2i^*})$. By Lemmata 2 and 3 and Corollary 2 either (1) a dominion D with $|Attr_{\mathcal{E}}(D, G)| > 2^{i^*-1}$ vertices was found by BÜCHIDOM or (2) all dominions in G' have more than h_{\max} vertices or there are no more dominions in G' . Thus either (2a) a dominion D with more than h_{\max} vertices is detected in the subsequent call to BÜCHI or (2b) there is no dominion in G' and this is the last iteration of Algorithm 1. Case (2b) can happen at most once and its runtime is bounded by $O(n \cdot 2^{2 \log(h_{\max})}) = O(n^2)$. In the cases (1) and (2a) more than 2^{i^*-1} vertices are removed from the graph in this iteration, as $h_{\max} > 2^{i^*-1}$. We charge each such vertex $O(n \cdot 2^i) = O(n \cdot h_{\max})$ time. Hence the total runtime for these cases is $O(n^2 \cdot h_{\max}) = O(n^{2.5})$.

It remains to consider the total time needed to compute $A = Attr_{\mathcal{E}}(D, G[V])$. By Lemma 1. (3) the attractor A can be computed in time $O(\sum_{v \in A} |In(v)|)$. Since the edges adjacent to vertices in A are removed from G after the iteration in which D was found, this attractor computation can be done in total time $O(m)$. We conclude that the runtime of Algorithm 1 is $O(n^{2.5})$. ■

We will show the correctness of Algorithm 1 by first proving that every dominion found in the Büchi game on G' is indeed a dominion in the parity-3 game on G . Together with Lemma 1. (4) this implies that the computed set W is indeed a part of the winning set of player \mathcal{E} in the parity-3 game. We then provide a winning strategy for player \mathcal{O} for all remaining vertices.

Lemma 5: Let the game graphs G and G' and the vertex sets V and B be defined as in Algorithm 1. If D is a dominion in the Büchi game $(G'[V], B)$, then D is a dominion in the parity-3 game $P = (G[V], \alpha)$.

Proof: Let Z be the vertices in V with priority α equal to -1 . The vertices in Z have priority 1 in the Büchi game, i.e., $Z \cap B = \emptyset$. Whenever a play in $G'[V]$ reaches a vertex u in Z , only u will be visited in the subsequent play since $Out(u) = \{u\}$. Thus no vertex in Z is winning for \mathcal{E} in $(G'[V], B)$, i.e., $D \cap Z = \emptyset$. Hence for all vertices in D the outgoing edges are the same in $G[V]$ and $G'[V]$. Thus D is \mathcal{O} -closed in $G[V]$ and the winning strategy of player \mathcal{E} for D in the Büchi game $(G'[V], B)$ is also winning for player \mathcal{E} for all vertices in D in the parity-3 game P . ■

Lemma 6 (Correctness): Given a parity-3 game P , let W

be the output of Algorithm 1. We have: (1) (*Soundness*). $W \subseteq W_{\mathcal{E}}(P)$; and (2) (*Completeness*). $W_{\mathcal{E}}(P) \subseteq W$.

Proof: The first part on soundness follows from Lemmata 5 and 1. (4). We now prove the completeness result. Given the output W , let \overline{W} denote the complement set. When Algorithm 1 terminates, the winning set of player \mathcal{E} in the Büchi game $(G'[\overline{W}], B)$ is empty (otherwise the algorithm would not have terminated). Also note that since the algorithm removes attractors for \mathcal{E} , the set \overline{W} is closed for \mathcal{E} (by Lemma 1. (1)). Consider the set $Z = \{v \in \overline{W} \mid \alpha(v) = -1\}$, its attractor $X = Attr_{\mathcal{O}}(Z, G'[\overline{W}])$, and the subgame induced by $U = \overline{W} \setminus X$. Note that in U the game graphs G and G' coincide. Thus all vertices in U must be winning for player \mathcal{O} in the Büchi game $(G[U], B)$ as otherwise $W_{\mathcal{E}}$ would have been non-empty for $(G'[\overline{W}], B)$. We prove the lemma by describing a winning strategy for player \mathcal{O} in P for all vertices in \overline{W} . Since \overline{W} is \mathcal{E} -closed, for vertices in $Z \cap V_{\mathcal{O}}$, the winning strategy chooses an edge in \overline{W} . For vertices in X player \mathcal{O} follows his attractor strategy to Z . In the subgame induced by $U = \overline{W} \setminus X$ player \mathcal{O} follows his winning strategy in the Büchi game $(G[U], B)$. Then in a play either (i) X is visited infinitely often; or (ii) from some point on only vertices in U are visited. In the former case, the attractor strategy ensures that then some vertex in Z with priority -1 is visited infinitely often; and in the later case, the subgame winning strategy ensures that only vertices with priority 1 and no vertices with priority 0 are visited infinitely often. It follows that $\overline{W} \subseteq W_{\mathcal{O}}(P)$, i.e., $W_{\mathcal{E}}(P) \subseteq W$, and the desired result follows. ■

Lemmata 4 and 6 yield the following result.

Theorem 4: Algorithm 1 correctly computes the winning sets in parity-3 games in $O(n^{2.5})$ time.

Computation of winning strategies. In parity-3 games the previous results for computing winning strategies for the players in their respective winning sets are as follows: The small-progress measure algorithm of [29] requires $O(nm)$ time to compute the winning strategy of the player whose parity is equal to the parity of the lowest priority and $O(n^2m)$ time to compute the respective winning strategies for both players; Schewe [40] shows how to modify the small-progress measure algorithm to compute the respective winning strategies of both players in $O(nm)$ time. We show that our algorithm also computes the respective winning strategies in $O(n^{2.5})$ time. We first observe that the algorithm of [7], [8] that solves Büchi games in $O(n^2)$ time also computes the respective winning strategies of both players (the algorithm is based on identifying traps and attractors, and the corresponding winning strategies are identified immediately with the computation). In Lemma 6 we describe the strategy computation for a winning strategy for player \mathcal{O} which involves an attractor strategy and the sub-game strategy for Büchi games, each of which can be computed in $O(n^2)$ time. A winning strategy for player \mathcal{E} is obtained in the iterations of the algorithm, i.e., whenever we obtain a dominion by solving Büchi games we also obtain a corresponding winning strategy, and similarly for the attractor computation. Thus the winning strategy for player \mathcal{E} can be

computed in $O(n^{2.5})$ time.

Corollary 3: Winning strategies for player \mathcal{E} and player \mathcal{O} in parity-3 games in their respective winning sets can be computed in $O(n^{2.5})$ time.

Remark 1: (DISCUSSION ON GENERAL PARITY GAMES). We now discuss the implication of our result for general parity games (we do not discuss general Streett games where the problem is coNP-complete [19]). The current best known algorithm for parity games with dependence on the number of priorities d is from [39], and the algorithm is referred as the *Big-step* algorithm. The (simplified) running time of the Big-step algorithm for $d = o(\sqrt{n})$ is $O(n^{\gamma(d)} \cdot m)$, where $\gamma(d)$ is approximately $d/3$ for large d . More precisely, $\gamma(d) = d/3 + 1/2 - 1/(\lceil 0.5d \rceil \lfloor 0.5d \rfloor)$ for odd d , and $\gamma(d) = d/3 + 1/2 - 1/(3d) - 1/(\lceil 0.5d \rceil \lfloor 0.5d \rfloor)$ for even d . Our algorithm for parity-3 games also extends to parity games as a recursive algorithm as follows: we apply our initialization step and iterated vertex deletions, and to find dominions we replace BÜCHI by our recursive algorithm that handles games that have one less priority and replace BÜCHIDOM by a procedure to find dominions with small-progress measure of [29] with our graph decomposition and codomain bounded by h_{\max} (where h_{\max} is chosen to balance the running time of the two dominion find procedures). Note that, in contrast to the Big-step algorithm, we apply the small-progress measure algorithm on a game with one less priority; this does not influence the correctness as we find (possibly different) dominions for the same player as in the recursion of the Big-step algorithm but slightly changes the running time analysis. For the sake of simplicity of presentation we present the analysis for the case of constantly many priorities, extending the analysis of [39] in combination with our approach (the extension for the general case follows from a similar extension of the analysis of the general case of [39]). Similar to the analysis and notation of [39], let $\beta(d) = \gamma(d)/(\lceil 0.5d \rceil + 1)$. With $h_{\max} = n^{\beta(d)}$ we obtain that the running time of our algorithm is $O(n^{1+\gamma(d+1)}) = O(n^{2+\gamma(d)-\beta(d)})$ for parity games with d priorities, i.e., it replaces m of [39] by $n^{2-\beta(d)}$. We present the details of the calculation. We show by induction that our algorithm solves parity games with $d-1$ priorities in $O(n^{1+\gamma(d)})$ time. The base case of $d-1 = 3$ follows from our algorithm for parity-3 games. The inductive case is as follows: To solve a parity game with d priorities, our algorithm calls the progress measure algorithm (on the graph decomposition) for $d-1$ priorities and recursively calls the algorithm for $d-1$ priorities at most $O(n^{1-\beta(d)})$ times. The total time for the progress measure algorithm is bounded by $O(n^2 \cdot (h_{\max})^{\lfloor 0.5d \rfloor}) = O(n^2 \cdot n^{\beta(d)\lfloor 0.5d \rfloor})$, and the total time for all calls to the algorithm for $d-1$ priorities is bounded by $O((n/h_{\max}) \cdot n^{1+\gamma(d)}) = O(n^{1-\beta(d)} \cdot n^{1+\gamma(d)})$. We obtain the recurrence $\gamma(d+1) = 1 + \gamma(d) - \beta(d)$, which yields $\gamma(d)$ as defined above and a running time of $O(n^{1+\gamma(d+1)})$ for d priorities. In the limit $\beta(d)$ approaches $2/3$. For small d we compare our running times with the Big-step algorithm in Table I. We have presented the details for parity-3 games for the following reasons: (1) All the key ideas and conceptual

details are easily demonstrated for the simpler case of parity-3 games and (2) while all previous ideas for general parity games (such as [30], [39]) and for Büchi games (such as [7], [8], [10]) fail to improve the running time for parity-3 games, our approach succeeds to break the long-standing $O(nm)$ barrier for dense graphs.

III. K-PAIR STREETT OBJECTIVES IN GRAPHS

A. Preliminaries

Let $G[S]$ denote the subgraph of a graph $G = (V, E)$ induced by the set of vertices $S \subseteq V$. $RevG$ denotes the graph with vertices V and all edges of G reversed. Let $Reach(S, G)$ be the set of vertices in G that can reach a vertex in $S \subseteq V$. A strongly connected component (SCC) of a directed graph $G = (V, E)$ is a subgraph $G[S]$ induced by a subset of vertices $S \subseteq V$ such that there is a path in $G[S]$ between every pair of vertices in S . We call an SCC *trivial* if it only contains a single vertex and no edges. All other SCCs are *non-trivial*. The set $Reach(S, G)$ and the maximal SCCs of a graph G can be found in linear time [3], [27], [41].

Algorithm STREETT and good component detection. The input is a directed graph $G = (V, E)$ and k Streett pairs (L_j, U_j) , $j = 1, \dots, k$. The size of the input is measured in terms of $m = |E|$, $n = |V|$, k , and $b = \sum_{j=1}^k (|L_j| + |U_j|) \leq nk$. Consider a maximal SCC C ; the *good component detection problem* asks to (a) output a non-trivial SCC $G[X] \subseteq C$ induced by some set of vertices X such that for all $1 \leq j \leq k$ either no vertex in L_j or at least one vertex in U_j is contained in the SCC (i.e., $L_j \cap X = \emptyset$ or $U_j \cap X \neq \emptyset$), or (b) detect that no such SCC exists. In the former case, there exists an infinite path that visits X infinitely often and satisfies the Streett objective, while in the later case there exists no infinite path that visits vertices in C infinitely often and satisfies the Streett objective. It follows from the results of [1] that the following algorithm, called Algorithm STREETT, suffices for the winning set computation: (1) Compute the maximal SCC decomposition of the graph; (2) for each maximal SCC C for which the good component detection returns an SCC, label the maximal SCC C as *satisfying*; (3) output the set of vertices that can reach a satisfying maximal SCC as the winning set. Since the first and last step are linear time, the runtime of Algorithm STREETT is dominated by the detection of good components in maximal SCCs. In the following we assume that the input graph is strongly connected and focus on good component detection.

B. Certificate computation

In this section we present our results for the certificate computation.

Given a start vertex x that belongs to the winning set, a certificate is an example of an *accepting run*, i.e., an infinite path from x that satisfies the objective. The output of Algorithm STREETT can be used to construct such an accepting run. Given a start vertex x and a good component $G[X]$ induced by some set of vertices X that is reachable from x , we generate the accepting run as follows. A path from x to X can be found

TABLE I
COMPARISON OF RUNNING TIMES FOR FEW PRIORITIES.

Algorithm	# priorities				
	3	4	5	6	7
Big-step [39]	$O(mn)$	$O(mn^{3/2})$	$O(mn^2)$	$O(mn^{7/3})$	$O(mn^{11/4})$
Big-step [39] with $m = \Theta(n^2)$	$O(n^3)$	$O(n^{7/2})$	$O(n^4)$	$O(n^{13/3})$	$O(n^{19/4})$
Our algorithm	$O(n^{2.5})$	$O(n^3)$	$O(n^{10/3})$	$O(n^{15/4})$	$O(n^{65/16})$

in linear time by a depth-first search. Let v be the vertex in X where this path ends. We call v the *root* of the SCC $G[X]$. We show next how to obtain, in $O(m + n \min(n, k))$ time, from the SCC $G[X]$ a cycle starting and ending at the root v such that the resulting certificate is indeed an accepting run. For this it is sufficient that the cycle in $G[X]$ contains for each L_j with $L_j \cap X \neq \emptyset$ a vertex of $U_j \cap X$, i.e., we do not have to include *all* vertices in X .

We can use Tarjan’s depth-first search based SCC algorithm [41] to traverse the subgraph $G[X]$ in linear $O(m)$ time, starting from root v . Tarjan’s algorithm constructs a graph called *jungle* with $O(|X|)$ edges that for an SCC $G[X]$ consists of a spanning tree and at most one *backedge* per vertex in X . The vertices are assigned pre-order numbers in the order they are traversed. We say an edge of $G[X]$ is a *backedge* if it leads from a vertex with a higher number to a vertex with a lower number. Spanning tree edges always lead from lower numbered vertices to higher numbered vertices. In Tarjan’s algorithm a *lowlink* is determined for each vertex u which refers to the lowest numbered vertex w that u can reach by a sequence of tree edges followed by at most one backedge. We additionally store at each vertex $u \neq v$ a *backlink* that is the first edge on the path from u to its lowlink. The backlinks can be determined and stored during the depth-first search without increasing its running time.

With this data structure we can find within $G[X]$ a path from root v to a vertex $u \in X$, $u \neq v$, and back by first searching for u in the spanning tree and then following the backlinks back to v . Since no vertex will appear more than twice on this path, its size and the time to compute it is $O(|X|)$. As it suffices to find such paths for one vertex per non-empty set $U_j \cap X$, we can generate a certificate from $G[X]$ in $O(m + |X| \min(|X|, |\{j \mid U_j \cap X \neq \emptyset\}|))$ time, which can be bounded with $O(m + n \min(n, k))$. This certificate has a size of $O(n \min(n, k))$.

Example 1 (Illustration of Tarjan’s jungle graph.): Figure 2 shows the types of edges and the values at the vertices as assigned by Tarjan’s SCC algorithm for a small example graph.

Example 2 (Lower bound.): Figure 3 shows that the smallest existing certificate can be as large as $\Theta(n \min(n, k))$.

C. Algorithm

In this section we present the algorithm for good component detection. First we introduce the different concepts used in the algorithm for good component detection. We start with




Fig. 2. An example for a “jungle” constructed by Tarjan’s SCC algorithm for an SCC. Backedges are dotted, spanning tree edges are solid. Backlinks are marked with a dot. The numbers of the vertices represent the order in which the vertices are visited, the numbers in brackets are the lowlinks.




Fig. 3. Let the only path between s and t be of length $\Theta(n/2) = \Theta(n)$, not containing any of the vertices v_j for $1 \leq j \leq k$. Let the Streett pairs (L_j, U_j) be given by $L_j = \{s\}$ and $U_j = \{v_j\}$ for $1 \leq j \leq k$. Then the size of the smallest certificate is $\Theta(nk)$, where k can be of order $\Theta(n)$.

describing the hierarchical graph decomposition technique for this setting, which will be crucial for the runtime analysis.

Graph decomposition. In our algorithm we decompose a graph G in the following way. For $i \in \{1, \dots, \lceil \log n \rceil\}$, let $G_i = (V, E_i)$ be a subgraph of G with $E_i = \{(u, v) \mid \text{Outdeg}(u) \leq 2^i\}$, i.e., the edges of G_i are the outedges of the vertices with outdegree at most 2^i . Note that for $i = \lceil \log n \rceil$ we have that $G_i = G$. We say vertices in G with $\text{Outdeg}(v) > 2^i$ are *colored blue* in G_i and denote the set of blue vertices in G_i by Bl_i . All other vertices are *white*. Note that all vertices in $G = G_{\lceil \log n \rceil}$ are white and that all vertices in Bl_i have outdegree zero in G_i .

Top and bottom strongly connected components. The algorithm will repeatedly find a top or a bottom SCC in the remaining graph G . A bottom SCC $G[S]$ in a directed graph G , induced by some set of vertices S , is an SCC with no edges from vertices in S to vertices in $V \setminus S$, i.e., no *outgoing* edges. A top SCC is a bottom SCC of $RevG$, i.e., an SCC without *incoming* edges. Top and bottom SCCs are by definition maximal SCCs. Note that every graph has at least one bottom and at least one top SCC. If they are not the same, then they are disjoint and thus one of them contains at most half of the vertices of G .

Bad vertices. In contrast to *good components* we also define

bad vertices. The basic idea behind the algorithms for good component detection, described for example in [25], is to repeatedly delete *bad* vertices until either a good component is found or it can be concluded that no such component exists. A vertex is *bad* if for some index j with $1 \leq j \leq k$ the vertex is in L_j but it is not strongly connected to any vertex in U_j . All other vertices are *good*. Note that good vertices can become bad if some vertex deletion disconnects an SCC or a vertex of a set U_j is deleted. A good component is a non-trivial SCC that only contains good vertices.

Data structure. The algorithm maintains for the current graph $G = (V, E)$ (some vertices of the input graph might have been deleted) a decomposition into vertex sets $S \subseteq V$ such that every SCC of G is completely contained in $G[S]$ for one of the sets S . For all the sets S a data structure $D(S)$ is saved in a list Q . The data structure $D(S)$ supports the following operations: (1) $Construct(S)$ initializes the data structure for the set S , (2) $Remove(S, D(S), B)$ removes a set $B \subseteq V$ from S and updates the data structure of S accordingly, and (3) $Bad(D(S))$ returns the set $\{v \in S \mid \exists j \text{ with } v \in L_j \text{ and } U_j \cap S = \emptyset\}$. In [25] an implementation of this data structure was given that achieves the following running times. For a set of vertices $S \subseteq V$ let $bits(S)$ be defined as $\sum_{j=1}^k (|S \cap L_j| + |S \cap U_j|)$.

Lemma 7 (Lemma 2.1 in [25]): After a one-time preprocessing of time $O(k)$, the data structure $D(S)$ can be implemented in time $O(bits(S) + |S|)$ for $Construct(S)$, time $O(bits(B) + |B|)$ for $Remove(S, D(S), B)$, and constant running time for $Bad(D(S))$.

By abuse of notation we denote by G the *current* graph maintained by the algorithm where some edges and vertices might have been deleted and use *input graph* to denote the unmodified, strongly connected graph for which a good component is searched. Our algorithm for good component detection is given in Algorithm 4. It maintains in a list Q a partition of the vertices in G into sets such that every SCC of G is contained in the subgraph induced by one of the vertex sets. The list is initialized with the set of all vertices in the strongly connected input graph. We will show that if a good component exists, its vertices must be fully contained in one of the vertex sets in the partition. The algorithm repeatedly removes a set S from Q and identifies and deletes bad vertices from $G[S]$. If no edge is contained in $G[S]$, the set S is removed as it can only induce trivial components. Otherwise the subgraph $G[S]$ is either determined to be strongly connected and output as a good component or a “small” maximal SCC in $G[S]$ is identified. To find a small maximal SCC the algorithm searches alternately in $G[S]$ and in $RevG[S]$ for a bottom SCC and stops as soon as one of the searches stops. (A bottom SCC in $RevG[S]$ is a top SCC in $G[S]$.) We use $Outdeg_H(v)$ to denote the out-degree of a vertex v in $H \in \{G, RevG\}$. We only describe the search in $G[S]$ here, the search in $RevG[S]$ is analogous. The algorithm uses the hierarchical graph decomposition of $G[S]$. The subgraph $G_i[S]$ for any i contains only the outedges of vertices with an outdegree of at most 2^i . The search for a bottom SCC is started at $i = 1$, then i is increased one

Input strongly connected graph $G = (V, E)$, Streett pairs (L_j, U_j) for $j = 1, \dots, k$

Output a good component in G if one exists

```

1: add  $Construct(V)$  to  $Q$ 
2: while  $Q \neq \emptyset$  do
3:   pull  $D(S)$  from  $Q$ 
4:   while  $Bad(D(S)) \neq \emptyset$  do
5:      $D(S) \leftarrow Remove(S, D(S), Bad(D(S)))$ 
6:   if  $G[S]$  contains at least one edge then
7:     for  $i \leftarrow 1, \dots, \lceil \log(|S|) \rceil$  do
8:       for all  $H \in \{G, RevG\}$  do
9:         construct  $H_i[S]$ 
10:         $Bl_i \leftarrow \{v \in S \mid Outdeg_H(v) > 2^i\}$ 
11:         $Z \leftarrow S \setminus Reach(Bl_i, H_i[S])$ 
12:                                      $\triangleright Z$  cannot reach  $Bl_i$ 
13:        if  $Z \neq \emptyset$  then
14:           $X \leftarrow SmallestBSCC(H_i[Z])$ 
15:          if  $X = S$  then  $\triangleright$  good component
16:            return  $G[S]$ 
17:          if  $|X| \leq |S|/2$  then
18:            add  $Remove(S, D(S), X)$  to  $Q$ 
19:            add  $Construct(X)$  to  $Q$ 
20:            go to Line 2
21: return no good component exists

```

Fig. 4. Detection of good components for the winning set computation in graphs with k -pair Streett objectives

by one if necessary, up to at most $\lceil \log(|S|) \rceil$. If for some i we can identify a bottom SCC that does not contain any blue vertex (i.e., a vertex for which some edges are missing in G_i), then the found SCC in $G_i[S]$ must also be a bottom SCC in $G[S]$. If multiple bottom SCCs (without blue vertices) are found in $G_i[S]$, we only consider the smallest one. The Procedure $SmallestBSCC(H')$ returns the set of vertices that induces the smallest bottom SCC in the graph H' . We then put the newly detected SCC and the “rest” of S back into Q .

The idea of the running time analysis is as follows. We can show that a bottom SCC of $G[S]$ identified in iteration i of the outer for-loop must contain $\Omega(2^i)$ vertices. In time $O(n2^i)$ a standard SCC algorithm can compute all SCCs of $G_i[S]$ and thus also the smallest bottom SCC. The time needed for the search in all graphs $G_{i'}[S]$ for $1 \leq i' < i$ can be bounded with an additional factor of two. Thus the work for the search is $O(n)$ per vertex in the identified SCC.

Given that the subgraph $G[S]$ was split into at least one top and one bottom SCC, the smallest top or bottom SCC contains at most half of the vertices of the subgraph. By searching for a smallest bottom SCC (without blue vertices) in $G_i[S]$ and $RevG_i[S]$, we find one top or bottom SCC with at most half of the vertices of the subgraph. We charge the work for finding such an SCC to the vertices in this SCC. We will show that this yields a total running time of $O(n^2)$ for computing SCCs.

We additionally have to take the time for the maintenance

of the data structures into account. Here we use the properties of the data structure $D(S)$ described in Lemma 7 to obtain a running time of $O((n+b)\log n)$ for the maintenance of the data structures and the identification of bad vertices over the whole algorithm. Combined these ideas lead to a total running time of $O(n^2 + b\log n)$.

Lemma 8: Let $H \in \{G, RevG\}$ be the graph and let i^* be the iteration for which in Algorithm 4 the outer for-loop stops. Let Z be the non-empty set $S \setminus Reach(Bl_{i^*}, H_{i^*}[S])$ and let X be the set of vertices that induces the smallest bottom SCC $H[X]$ in $H_{i^*}[Z]$ returned by *SmallestBSCC*($H_{i^*}[Z]$). Assume we have $|X| \leq |S|/2$. Then $H[X]$ contains at least 2^{i^*-1} vertices.

Proof: As Bl_{i^*-1} is the set of vertices in $H_{i^*-1}[S]$ with outdegree larger than 2^{i^*-1} , any bottom SCC $H[Y]$ that contains a vertex of Bl_{i^*-1} , has $|Y| \geq 2^{i^*-1}$. Hence it suffices to show that $X \cap Bl_{i^*-1} \neq \emptyset$. Assume by contradiction that $X \cap Bl_{i^*-1} = \emptyset$. Since $H[X]$ is a bottom SCC, no vertex in X can reach any vertex in Bl_{i^*-1} , i.e., $X \subseteq S \setminus Reach(Bl_{i^*-1}, H_{i^*}[S])$. As all edges in $H_{i^*-1}[S]$ are contained in $H_{i^*}[S]$, this implies $X \subseteq S \setminus Reach(Bl_{i^*-1}, H_{i^*-1}[S])$. Since *SmallestBSCC* finds the smallest bottom SCC in graph H_i for each i , the outer for-loop would thus have terminated in an iteration $i \leq i^* - 1$. Contradiction. ■

Lemma 9 (Runtime): Algorithm 4 can be implemented in time $O(n^2 + b\log n)$.

Proof: The preprocessing and initialization of the data structure and the removal of bad vertices in the whole algorithm take time $O(m+k+b)$ using Lemma 7. Additionally we maintain at each vertex a list of its incoming and a list of its outgoing edges including pointers to the lists of its neighbors, which we use to update the lists of its neighbors. Since each vertex is deleted at most once, this data structure can be constructed and maintained in total time $O(m)$.

Consider the while loop where a set S is removed from Q . The search in $G[S]$ and $RevG[S]$ only increases the running time by a factor of two, thus we restrict the analysis of the running time to $G[S]$. Let $n' \leq n$ be the number of vertices in S . The construction of $G_i[S]$, Z , and $G[X]$ can all be done in time $O(n' \cdot 2^i)$ for each i , i.e., in total time $O(n' \cdot 2^{i^*})$ up to level i^* . If $X = S$, then the algorithm terminates and the time for processing S can be bounded by $O(n' \cdot 2^{\log n'}) = O(n'^2)$. If the processing of S ends when some bottom SCC $G[X] \subset G[S]$ induced by some set of vertices X is found, let i^* be the value of i when $G[X]$ is detected and inserted into Q , and let c be some constant such that the time spent in this search for X is bounded by $c \cdot n' \cdot 2^{i^*-1}$. By Lemma 8 the set X contains at least 2^{i^*-1} vertices. Let $|X| = n_1$. The algorithm ensures $n_1 \leq n'/2$. We claim that the total running time for processing all sets S , except for the work in *Remove* and *Construct*, can be bounded by $f(n) = 2cn^2$. Whenever the algorithm does not terminate, we have by induction, and in particular for $n' = n$,

$$\begin{aligned} f(n') &\leq f(n_1) + f(n' - n_1) + cn'n_1, \\ &\leq 2cn_1^2 + 2c(n' - n_1)^2 + cn'n_1, \end{aligned}$$

$$\begin{aligned} &= 2cn_1^2 + 2cn'^2 - 4cn'n_1 + 2cn_1^2 + cn'n_1, \\ &= 2cn'^2 + 4cn_1^2 - 3cn'n_1, \\ &\leq 2cn'^2, \end{aligned}$$

where the last inequality follows from $n_1 \leq n'/2$.

The operations *Remove* and *Construct* are called once per found bottom SCC $G[X]$ with $X \neq S$ and take by Lemma 7 $O(|X| + bits(X))$ time. By $n_1 \leq n'/2$ we have that whenever a vertex v is in X , the size of the set in Q containing v is halved; this can happen at most $\lceil \log n \rceil$ times. Hence, by charging $O(1)$ to the vertices in X and, respectively, to $bits(X)$, the total running time for this part can be bounded by $O((n+b)\log n)$, as each vertex and bit will only be charged $O(\log n)$ times. Combining all parts yields the claimed running time bound of $O(n^2 + b\log n)$. ■

To prove the correctness of Algorithm 4 we first show that all candidates for good components are in Q before each iteration of the algorithm.

Lemma 10: Before each iteration of the outer while-loop every good component of the input graph is contained in one of the subgraphs $G[S]$ for which the data structure $D(S)$ is maintained in the list Q .

Proof: We will show that the algorithm never removes edges or vertices that belong to a good component, which together with a correct initialization of the list Q will imply the lemma. At the beginning of the algorithm one data structure for the whole strongly connected input graph is added to Q . Thus every good component is contained in this data structure in Q after the initialization. At the beginning of each iteration of the outer while-loop the data structure of one of the subgraphs $G[S]$ is pulled from the list Q . In Lines 4–5 we remove vertices from the subgraph that are in some set L_j but not strongly connected to any vertex in U_j , i.e., bad vertices. In Line 6 we remove trivial SCCs. Observe that a good component is non-trivial and does not contain any bad vertices. Thus the removal of bad vertices and trivial SCCs does not remove any vertices of a good component, i.e., after the removal of these vertices the updated subgraph $G[S]$ still contains the good components it contained before. If no good component is identified in this iteration, i.e., the algorithm does not terminate, we find a bottom or top SCC $G[X]$, induced by some set of vertices X , which is by definition a maximal SCC. Since a good component is strongly connected, every good component in $G[S]$ either is a subgraph of the newly identified SCC $G[X]$ or does not contain any vertex of X . Thus the removed edges between $G[X]$ and the remaining subgraph cannot belong to a good component. Finally, we add the data structures for $G[X]$ as well as for $G[S \setminus X]$ to Q . Thus no vertex or edge of a good component was removed and every good component continues to be completely contained in a subgraph in Q . ■

As all candidates for good components are maintained in the list Q , it remains to show that the algorithm makes progress in each iteration and correctly outputs a good component if and only if one exists. We use the following lemma to show the first part and then conclude the correctness proof with Lemma 12.

Lemma 11: In each iteration of the outer while-loop either (a) the number of elements in Q decreases, (b) the algorithm terminates, or (c) one of the subgraphs maintained in Q is replaced by two smaller subgraphs.

Proof: Let S be the set of vertices for which $D(S)$ is pulled from Q at the beginning of the outer while-loop. Let $S' \subseteq S$ be the set of vertices from the set S that are not removed as bad vertices in Lines 4–5.

Case (a): If $G[S']$ does not contain any edge, then S' is either empty or induces only trivial SCCs. In this case S' is not considered further after the if-statement in Line 6, i.e., no element is added to Q and thus the number of elements in Q decreases by one.

For Cases (a) and (b) note that whenever in the algorithm the set Z is non-empty (for some $i \in \{1, \dots, \lceil \log n \rceil\}$ and some set of vertices S') and a bottom SCC in $H_i[Z]$ for $H \in \{G, RevG\}$ induced by a set of vertices X is identified, then $H_i[X] = H[X]$ is a bottom SCC in $H[S']$ because $X \cap Bl_i = \emptyset$.

Case (b): If $G[S']$ is a non-trivial SCC, then the only bottom SCC in $G[S']$ or $RevG[S']$ is induced by S' . Since we have $H_{\lceil \log n \rceil} = H$ and S' is not empty, a set of vertices X that induces a bottom SCC in $H_i[Z]$ will be identified for some i . As $H[X]$ is a bottom SCC in $H[S']$, we have $X = S'$ and the algorithm terminates.

Case (c): If $G[S']$ contains an edge but is not strongly connected, then it contains a top SCC and a bottom SCC that are disjoint and at least one of them contains at most half of the vertices in S' . Thus for some i and one of G and $RevG$ a set of vertices X with $|X| \leq |S'|/2$ that induces a bottom SCC in $H_i[Z]$ will be identified and the data structures of two non-empty subsets of S' , namely X and $V \setminus X$, are added to Q in Lines 18–19. ■

Lemma 12 (Correctness): Algorithm 4 outputs a good component if one exists, otherwise the algorithm reports that no such component exists.

Proof: First we show that whenever Algorithm 4 outputs a subgraph $G[S]$ induced by some set of vertices S , then $G[S]$ is a good component. Line 6 ensures only non-trivial SCCs are considered. After the removal of bad vertices from S in Lines 4–5, we know that for all $1 \leq j \leq k$ and all vertices in $S \cap L_j$ there exists a vertex in $S \cap U_j$. Thus if $G[S]$ is strongly connected, then $G[S]$ is a good SCC, and the only maximal SCC in $G[S]$ is $G[S]$ itself. Only in this case $G[S]$ is output (in Line 16).

Algorithm 4 terminates if a good component is identified or Q is empty. Lemma 10 shows that before every iteration of the outer while-loop every good component is contained in one of the subgraphs $G[S]$ in Q . That is, if a good component exists in G , the algorithm will not terminate until a good component is identified. Lemma 11 shows that whenever the algorithm does not terminate in an iteration of the outer while-loop, either a subgraph is removed from Q or one of the subgraphs from Q is split into two smaller subgraphs. Each case can happen at most n times. This implies that the algorithm terminates after a finite number of steps if no good component exists. Next we show that if there exists a good component in G ,

then the algorithm will output a good component. Let Y be a maximal good component in G and let S_Y be the vertex set maintained in Q that currently contains the vertices in Y . By the arguments above after a finite number of steps either (1) another good component is detected or (2) $D(S_Y)$ is pulled from Q . By Lemma 10 the component Y is never split by the algorithm thus after Case (2) happened at most n times, one of the following two cases occurs: either (2a) $D(S_Y)$ is pulled from Q with $G[S_Y] \supset Y$ and after the removal of bad vertices from S_Y , $G[S_Y]$ without the bad vertices is equal to Y or (2b) $G[S_Y] = Y$ is pulled from Q . In both cases the good component Y is output and the algorithm terminates: Since Y is non-trivial, the condition in Line 6 is satisfied. The algorithm searches for a top or bottom SCC in Y . Since Y is strongly connected, the only top or bottom SCC in Y is Y itself. Hence the algorithm outputs Y in Line 16. ■

Recall Algorithm STRETT that calls Algorithm 4 for each maximal SCC in the input graph and then computes reachability to the union of the identified good components. Lemmata 9 and 12 yield the following result.

Theorem 5: Algorithm STRETT correctly computes the winning set in graphs with k -pair Streett objectives in $O(n^2 + b \log n)$ time. Given a vertex x in the winning set, a certificate for x can be output in time $O(m + n \min(n, k))$.

Remark 2 (Optimality.): We showed that in a graph with k -pair Streett objectives the winning set and a certificate can be computed and output in time $O(n^2 + b \log n)$. Example 2 shows a lower bound of $\Omega(n \min(n, k))$ for outputting a certificate. Note that the size of the input is at least b . Hence the presented algorithm is optimal up to a log factor when $k = \Omega(n)$ and a certificate is required.

ACKNOWLEDGEMENT

We would like to thank Tom Henzinger for his useful comments. V. L. would like to thank Lara Turner for her feedback on an early version of the manuscript.

K. C. is supported by the Austrian Science Fund (FWF): P23499-N23 and S11407-N23 (RiSE/SHiNE), an ERC Start Grant (279307: Graph Games), and a Microsoft Faculty Fellows Award. M. H. is supported by the Austrian Science Fund (FWF): P23499-N23 and the Vienna Science and Technology Fund (WWTF) grant ICT10-002. V. L. is supported by the Vienna Science and Technology Fund (WWTF) grant ICT10-002. Additionally, the research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement no. 340506.

REFERENCES

- [1] R. Alur and T. A. Henzinger, “Computer-aided verification,” 2004, unpublished, available at <http://www.cis.upenn.edu/cis673/>.
- [2] R. Alur, T. Henzinger, and O. Kupferman, “Alternating-time temporal logic,” *Journal of the ACM*, vol. 49, pp. 672–713, 2002.
- [3] C. Beeri, “On the membership problem for functional and multivalued dependencies in relational databases,” *ACM Transactions on Database Systems*, pp. 241–259, 1980.
- [4] R. Bloem, K. Chatterjee, K. Greimel, T. A. Henzinger, and B. Jobstmann, “Robustness in the presence of liveness,” in *CAV*, 2010, pp. 410–424.

- [5] R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Sa'ar, "Synthesis of reactive(1) designs," *J. Comput. Syst. Sci.*, vol. 78, no. 3, pp. 911–938, 2012.
- [6] K. Chatterjee, S. Chaudhary, and P. Kamath, "Faster algorithms for alternating refinement relations," in *CSL*, ser. LIPIcs 16. Schloss Dagstuhl, 2012, pp. 167–182.
- [7] K. Chatterjee and M. Henzinger, "An $O(n^2)$ Time Algorithm for Alternating Büchi Games," in *SODA*, 2012, pp. 1386–1399.
- [8] —, "Efficient and Dynamic Algorithms for Alternating Büchi Games and Maximal End-component Decomposition," *Journal of the ACM*, vol. 61, no. 3, p. 15, 2014.
- [9] K. Chatterjee, T. A. Henzinger, and V. S. Prabhu, "Timed parity games: Complexity and robustness," *Logical Methods in Computer Science*, vol. 7, no. 4, 2011.
- [10] K. Chatterjee, M. Jurdziński, and T. Henzinger, "Simple stochastic parity games," in *CSL'03*, ser. LNCS, vol. 2803. Springer, 2003, pp. 100–113.
- [11] K. Chatterjee and V. S. Prabhu, "Synthesis of memory-efficient, clock-memory free, and non-zeno safety controllers for timed systems," *Inf. Comput.*, vol. 228, pp. 83–119, 2013.
- [12] A. Church, "Logic, arithmetic, and automata," in *Proceedings of the International Congress of Mathematicians*. Institut Mittag-Leffler, 1962, pp. 23–35.
- [13] L. de Alfaro and M. Faella, "An accelerated algorithm for 3-color parity games with an application to timed games," in *CAV*, 2007, pp. 108–120.
- [14] L. de Alfaro, M. Faella, T. A. Henzinger, R. Majumdar, and M. Stoelinga, "The element of surprise in timed games," in *CONCUR*, 2003, pp. 142–156.
- [15] L. de Alfaro and T. Henzinger, "Interface automata," in *FSE'01*. ACM Press, 2001, pp. 109–120.
- [16] D. Dill, *Trace Theory for Automatic Hierarchical Verification of Speed-independent Circuits*. The MIT Press, 1989.
- [17] A. Duret-Lutz, D. Poitrenaud, and J.-M. Couvreur, "On-the-fly Emptiness Check of Transition-Based Streett Automata," in *ATVA*, 2009, pp. 213–227.
- [18] R. Ehlers, "Short Witnesses and Accepting Lassos in ω -Automata," in *LATA*, 2010, pp. 261–272.
- [19] E. Emerson and C. Jutla, "The complexity of tree automata and logics of programs," in *FOCS'88*. IEEE, 1988, pp. 328–337.
- [20] —, "Tree automata, mu-calculus and determinacy," in *FOCS*, 1991, pp. 368–377.
- [21] E. Emerson and C.-L. Lei, "Modalities for Model Checking: Branching Time Logic Strikes Back," *Science of Computer Programming*, vol. 8, no. 3, pp. 275–306, 1987.
- [22] N. Francez, *Fairness*. New York: Springer, 1986.
- [23] Y. Godhal, K. Chatterjee, and T. A. Henzinger, "Synthesis of AMBA AHB from formal specification: a case study," *STTT*, vol. 15, no. 5-6, pp. 585–601, 2013.
- [24] M. Henzinger, V. King, and T. Warnow, "Constructing a Tree from Homeomorphic Subtrees, with Applications to Computational Evolutionary Biology," *Algorithmica*, vol. 24, pp. 1–13, 1999.
- [25] M. Henzinger and J. Telle, "Faster Algorithms for the Nonemptiness of Streett Automata and for Communication Protocol Pruning," in *SWAT*, 1996, pp. 16–27.
- [26] T. A. Henzinger, O. Kupferman, and S. K. Rajamani, "Fair simulation," *Inf. Comput.*, vol. 173, no. 1, pp. 64–81, 2002.
- [27] N. Immerman, "Number of quantifiers is better than number of tape cells," *Journal of Computer and System Sciences*, pp. 384–406, 1981.
- [28] M. Jurdziński, "Deciding the winner in parity games is in $UP \cap co-UP$," *Information Processing Letters*, vol. 68, no. 3, pp. 119–124, 1998.
- [29] —, "Small Progress Measures for Solving Parity Games," in *STACS*, 2000, pp. 290–301.
- [30] M. Jurdziński, M. Paterson, and U. Zwick, "A Deterministic Subexponential Algorithm for Solving Parity Games," *SIAM Journal on Computing*, vol. 38, no. 4, pp. 1519–1532, 2008.
- [31] T. Latvala and K. Heljanko, "Coping With Strong Fairness," *Fundamenta Informaticae*, vol. 43, no. 1-4, pp. 175–193, 2000.
- [32] O. Lichtenstein and A. Pnueli, "Checking That Finite State Concurrent Programs Satisfy Their Linear Specification," in *POPL*, 1985, pp. 97–107.
- [33] D. Martin, "Borel determinacy," *Annals of Mathematics*, vol. 102(2), pp. 363–371, 1975.
- [34] R. McNaughton, "Infinite games played on finite graphs," *Annals of Pure and Applied Logic*, vol. 65, no. 2, pp. 149–184, 1993.
- [35] N. Piterman, A. Pnueli, and Y. Sa'ar, "Synthesis of reactive(1) designs," in *VMCAI*, ser. LNCS 3855, Springer, 2006, pp. 364–380.
- [36] A. Pnueli and R. Rosner, "On the synthesis of a reactive module," in *POPL'89*. ACM Press, 1989, pp. 179–190.
- [37] P. Ramadge and W. Wonham, "Supervisory control of a class of discrete-event processes," *SIAM Journal of Control and Optimization*, vol. 25, no. 1, pp. 206–230, 1987.
- [38] S. Safra, "On the complexity of ω -automata," in *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, 1988, pp. 319–327.
- [39] S. Schewe, "Solving Parity Games in Big Steps," in *FSTTCS*, 2007, pp. 449–460.
- [40] —, "Synthesis of Distributed Systems," Ph.D. dissertation, Universität des Saarlandes, 2008.
- [41] R. E. Tarjan, "Depth First Search and Linear Graph Algorithms," *SIAM Journal of Computing*, vol. 1, no. 2, pp. 146–160, 1972.
- [42] W. Thomas, "Languages, automata, and logic," in *Handbook of Formal Languages*, G. Rozenberg and A. Salomaa, Eds. Springer, 1997, vol. 3, Beyond Words, ch. 7, pp. 389–455.
- [43] J. Vöge and M. Jurdziński, "A discrete strategy improvement algorithm for solving parity games," in *CAV'00*. LNCS 1855, Springer, 2000, pp. 202–215.
- [44] W. Zielonka, "Infinite games on finitely coloured graphs with applications to automata on infinite trees," *Theoretical Computer Science*, vol. 200, no. 1–2, pp. 135–183, 1998.