A Comparison of BDD-Based Parity Game Solvers

Lisette Sanchez, Wieger Wesselink and Tim A.C. Willemse

Eindhoven University of Technology, The Netherlands

Parity games are two player games with omega-winning conditions, played on finite graphs. Such games play an important role in verification, satisfiability and synthesis. It is therefore important to identify algorithms that can efficiently deal with large games that arise from such applications. In this paper, we describe our experiments with BDD-based implementations of four parity game solving algorithms, *viz.* Zielonka's recursive algorithm, the more recent Priority Promotion algorithm, the Fixpoint-Iteration algorithm and the automata based APT algorithm. We compare their performance on several types of random games and on a number of cases taken from the Keiren benchmark set.

1   Introduction

Parity games [16, 26, 30] are infinite duration games played by two players on a finite directed graph. Each vertex in the game graph is owned by one of the two players and vertices are assigned a colour, or *priority*. The game is played by pushing a token along the edges in the graph; the choice to which vertex the token is to move next is decided by the player owning the vertex currently holding the token. A parity condition determines the winner of this infinite play. A vertex in the game is won by the player that has a strategy for which, no matter how the opponent plays, every play from that vertex is won by her; the winner of each vertex is uniquely determined [26].

The parity game solving problem is to compute the set of vertices won by each player. This problem is known to be in UP ∩ coUP, a class that indicates that the problem is unlikely to be hard for NP. While subexponential solving algorithms have been devised, despite the (deceptive) simplicity of the game and a continuous research effort, no polynomial time algorithm for solving parity games has been found. However, only recently the problem was shown to be solvable in *quasi-polynomial time* [10, 22, 17].

From a practical viewpoint, parity games are interesting since they underlie typical verification, satisfiability and synthesis problems, see [12, 16, 2] and the references therein. For instance, in the mCRL2 toolset [13], parity games, derived from parameterised Boolean equation systems [20], are used to solve modal μ-calculus model checking problems for reactive systems. Games originating from such model checking problems often contain a large number of vertices (typically exceeding 10^6), whereas the number of distinct priorities in the games is typically very small (often at most 3). Consequently, these applications require algorithms and techniques that can efficiently deal with such magnitudes.

Most parity game solving algorithms fall in one of two categories: 'strategy identification' (SI) algorithms and 'dominion identification' (DI) algorithms. Algorithms from the SI category directly compute the winning strategies for both players; *e.g.*, by means of policy iteration or by maintaining some statistics about the plays that can emerge from a vertex. The recent quasi-polynomial algorithms [10, 22, 17] all fall in this category, but also classical algorithms such as Jurdziński's *small progress measures* algorithm [21], and, to some extent, the Fixpoint-Iteration (FI) algorithm [8], which is closely related to the small progress measures algorithm, see *ibid*. The DI category of algorithms proceed by (recursively) decomposing the game graph in *dominions*: small subgraphs that are won by a single player and from which the opponent cannot escape. A classical exponent of the latter category is Zielonka's *recursive*

algorithm [30] but also the recently introduced Priority Promotion (PP) algorithm [6]. While DI algorithms typically have a worst-case running time complexity that is theoretically less attractive than that of SI algorithms, in practice, the SI are significantly outperformed by DI algorithms [14].

Since parity games that originate from practical verification problems can become quite large, it is natural to study whether existing algorithms can be implemented efficiently using symbolic representations of the game graph, *i.e.*, using Binary Decision Diagrams (BDDs). While the size of BDDs may depend very much on the BDD variable ordering, when the compression ratio (the number of BDD nodes versus the number of satisfying assignments) is favourable, BDDs can be used to concisely represent large graphs. DI algorithms typically exploit set-based operations, which can be implemented straightforwardly and efficiently using BDDs. The same applies to SI algorithms such as the FI algorithm [8] and the automata-based APT algorithm [28].

In this paper, we describe and study four set-based algorithms, *viz.* Zielonka's algorithm, the PP algorithm, the FI implementation and the APT algorithm. In particular, we reassess the conclusions from [14] in which Zielonka's algorithm and the PP algorithm are shown to have similar performance when computing with explicit representations of the game graph. Moreover, we revisit the observation of [28] that for games in which the number of vertices is exponentially larger (given a sufficiently large enough base) than the number of distinct priorities in the game, the APT algorithm significantly outperforms Zielonka's algorithm. Like the aforementioned works, we base our observations on random games, which we convert to BDDs using a binary encoding. Such an encoding typically yields BDDs with a poor compression ratio. While this does not reflect scenarios in which BDDs excel, this does yield important insights into which symbolic algorithms are more sensitive to BDDs with poor compression ratios. Finally, we compare the performance of all four algorithms on several of the larger cases taken from the Keiren benchmark set [24] of parity games and demonstrate that our symbolic algorithms can solve games that cannot be solved using explicit techniques.

We find that the symbolic Zielonka and PP algorithms perform mostly similarly; the observations of [14] thus seem to largely carry over to the symbolic setting. However, in our setting we cannot faithfully reproduce the observations of [28] regarding the superior performance of the APT algorithm for favourable vertex/priority ratios. While APT's performance improves with more favourable ratios, it (and also the FI algorithm) suffers more than Zielonka and PP when handling graphs represented by BDDs with a poor compression ratio. Yet, when the BDDs remain small, both APT and FI often outperform Zielonka and PP, as witnessed by the cases taken from the Keiren benchmark set.

Related Work. BDD-based implementations for parity game solving have been studied in the past. In [3], an algorithm that is inspired by Zielonka's algorithm is used to solve planning problems, whereas in [23], a symbolic version of Zielonka's algorithm is implemented in an effort to solve model checking problems represented by parameterised Boolean equation systems [20]. We note that the focus of the latter work is to extract symbolic parity games, represented by *List Decision Diagrams* (LDDs) underlying parameterised Boolean equation systems using a reachability analysis and solving the thus constructed parity games. Since our primary aim is to compare the performance of various symbolic parity game solving algorithms, we have followed a slightly different and less general route in our experiments on model checking, converting all finite data types occurring in such equation systems to Booleans. This allows us to avoid a costly reachability analysis. Obviously, our insights into which algorithms can be used best to solve a given parity game can be combined with the technique of [23] to obtain faster end-to-end verification tooling. Another line of research has concentrated on giving set-based versions of classical algorithms such as the small progress measures algorithm, the bigstep algorithm and the subexponential dominion decomposition algorithm, see *e.g.* [9, 11]. Symbolic algorithms studied for more restricted

types of games include [1, 7].

Outline. In Section 2, we introduce parity games and the relevant concepts. The four algorithms that we compare are then introduced and discussed in Section 3 and we describe how these can be implemented using BDD techniques in Section 4. In Section 5, we describe our experimental evaluation of our implementations and we finish with conclusions in Section 6.

2 Parity Games

A parity game is an infinite duration game, played by players *odd*, denoted by 1 and *even*, denoted by 0, on a directed, finite graph. This game graph is formally defined as follows.

Definition 1 *A parity game is a tuple* $(V, E, p, (V_0, V_1))$, *where:*
- V *is a finite set of vertices, partitioned in a set V_0 of vertices owned by player* 0, *and a set of vertices V_1 owned by player* 1,
- $E \subseteq V \times V$ *is the edge relation; we assume that E is left-total,* i.e., *for all $v \in V$, there is some $w \in V$ such that $(v, w) \in E$,*
- $p : V \to \mathbb{N}$ *is a priority function that assigns priorities to vertices.*

We depict parity games as graphs in which diamond-shaped vertices represent vertices owned by player 0 and box-shaped vertices represent vertices owned by player 1. Priorities, associated with vertices, are typically written inside vertices, see Figure 1.
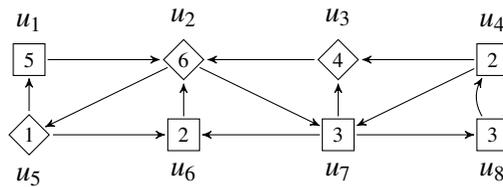
Figure 1: A game with 6 different priorities, 3 vertices owned by player 0 and 5 owned by player 1.

Plays, Strategies and Winning. We write $v \to w$ whenever $(v, w) \in E$; vE denotes the set $\{w \in V \mid v \to w\}$ and Ew denotes the set $\{v \in V \mid v \to w\}$. Henceforth, $\alpha \in \{0, 1\}$ denotes an arbitrary player. We write $\bar{\alpha}$ for α's opponent; *i.e.*, $\bar{0} = 1$ and $\bar{1} = 0$. A sequence of vertices $v_1, \ldots, v_n$ is a *path* if $v_m \to v_{m+1}$ for all $1 \le m < n$. Infinite paths are defined in a similar manner. The n^{th} vertex of a path π is denoted π_n.

A game starts by placing a token on vertex $v \in V$. Players move the token indefinitely according to a single simple rule: if the token is on some vertex $v \in V_\alpha$, player α gets to move the token to an adjacent vertex. We refer to the infinite path through the graph, obtained this way, as a *play*. A play is won by player 0 iff the *maximal* priority that occurs infinitely often along that play is *even*.

The moves of players 0 and 1 are determined by their respective *strategies*. A strategy for a player α determines, for a vertex $\pi_i \in V_\alpha$ the next vertex π_{i+1} that will be visited if the token visits π_i. In general, a strategy is a partial function $\sigma : V^* \times V_\alpha \to V$ which, given a history of vertices visited by the token and a vertex on which the token currently resides, determines the next vertex. However, due to the positional determinacy theorem for parity games [16] it suffices to consider *positional strategies*: partial functions of the form $\sigma : V_\alpha \to V$ that are compatible with E, *i.e.*, for all $v \in V_\alpha$ for which σ is defined, it is the case that $\sigma(v) \in vE$. An infinite path π is compatible with a given strategy σ if for all vertices π_i for which σ

is defined, we have $\pi_{i+1} = \sigma(\pi_i)$. A strategy σ is *closed* on a set of vertices U iff every play compatible with σ remains within U. We say that a set of vertices U is α-*closed* iff α has a strategy that is closed on U.

A strategy σ for player α is *winning* from a vertex v if and only if α is the winner of every play starting in v that is compatible with σ. A set of vertices W_α is won by α if for each vertex $v \in W_\alpha$, player α has some winning strategy from v. Another consequence of the aforementioned positional determinacy theorem is that the set of vertices W_0 won by player 0 and the set of vertices W_1, won by player 1, forms a partition of the set V: every vertex is won by exactly one player.

Example 1 Consider the parity game depicted in Fig. 1. The vertices u_1, u_2, u_3, u_5, u_6 are won by player 0 whereas vertices u_4, u_7 and u_8 are won by player 1. To see why vertices u_4 and u_8 are won by player 1, consider her strategy σ (that is *closed* on $\{u_4, u_8\}$) defined by $\sigma(u_4) = u_8$ and $\sigma(u_8) = u_4$. Any play that is compatible with σ infinitely often visits a dominating odd priority and is thus won by player 1.   □

Subgames, Attractors and Dominions.   For a game $G = (V, E, p, (V_0, V_1))$ and a set $U \subseteq V$, we define the *subgame* of G, denoted $G \cap U$, as the maximal substructure $(V', E', p', (V_0', V_1'))$ that is obtained by restricting the graph (V, E) to U and imposing the restriction on the other elements of G, *i.e.*, $V' = U$, $E' = E \cap (U \times U)$, $p'(v) = p(v)$ for all $v \in U$, and $V_0' = V_0 \cap U$ and $V_1' = V_1 \cap U$. If the edge relation E' of $G \cap U$ is again left-total, then the subgame is again a parity game. We furthermore use the abbreviation $G \setminus A$ to denote the subgame $G \cap (V \setminus A)$.

Parity game solving can essentially be reduced to identifying or computing appropriate subgraphs that are entirely won by a single player and from which the opponent cannot escape. A subgame that has this property for a player α is called an α-*dominion*. Technically, an α-dominion is a non-empty set of vertices $D_\alpha \subseteq V$ such that player α has a winning strategy, closed on D_α, from *all* vertices in the set D_α.

The *pre-image* $pre_\alpha(U)$ (or single-step *force* [28]) for a given non-empty set $U \subseteq V$ and player α is defined as $pre_\alpha(U) = \{v \in V_\alpha \mid vE \cap U \neq \emptyset\} \cup \{v \in V_{\bar{\alpha}} \mid vE \subseteq U\}$. The α-attractor into U, denoted $Attr_\alpha(U)$, is the set of vertices for which player α can force play into U. Formally, $Attr_\alpha(U)$ is the least set $A \supseteq U$ to which α can force play: $Attr_\alpha(U) = \mu A.(U \cup pre_\alpha(A))$. The *confined* α-attractor, denoted $Attr_\alpha(T, U)$, is defined analogously by limiting the set of vertices to those from T. That is, it represents the subset of vertices $T' \subseteq T$ from which α can force play to U while remaining in T. We note that for arbitrary set $U \subseteq V$, the subgame $G \setminus Attr_\alpha(U)$ is again a parity game; that is, the edge relation of the subgame is again total. Furthermore, the α-attractor into an α-dominion D_α is again an α-dominion. Finally, observe that the set $V \setminus Attr_\alpha(U)$ is $\bar{\alpha}$-closed.

Example 2 Reconsider the parity game from Fig. 1. We already identified a closed strategy that is winning for player 1 on the set of vertices $\{u_4, u_8\}$. So this set is a 1-dominion. It is not a maximal 1-dominion: we can add u_7 to extend it. The 0-attractor into the set $\{u_2\}$ contains the vertices u_1 (which is owned by 1 but trivially attracted to u_2 since it has but one successor vertex), u_3, u_5 and u_6. Vertex u_4 does not belong to the 0-attractor into $\{u_2\}$ as player 1 can choose to move to u_7. For similar reasons, vertices u_7 and u_8 do not belong to the 0-attractor into $\{u_2\}$.   □

3   Parity Game Solving Algorithms

We consider four set-based algorithms for solving parity games. These algorithms will be introduced in some detail in the next subsections.

3.1 Zielonka's Recursive Algorithm

The algorithm by Zielonka is a divide and conquer algorithm that searches for dominions in subgames. This algorithm is essentially distilled from a constructive proof of the positional determinacy of parity games by, among others, Zielonka [30]. Despite the fact that the algorithm has a relatively bad theoretical worst-case complexity (it runs in $\mathcal{O}(mn^d)$ where n is the number of vertices, m the number of edges and d is the number of different priorities in the game) and exponential worst-case examples are known for various classes of special games [19], the algorithm remains among the most successful solvers for parity games in practice, see [18] and the recent exploration [14].

Zielonka's algorithm (see Algorithm 1) constructs winning regions for both players out of the solution of subgames with fewer different priorities and vertices. For this, the algorithm relies on the fact that higher priorities in the game dominate lower priorities, and that any forced revisit of these higher priorities is beneficial to the player associated with the parity of the priority. For non-trivial games, the algorithm relies on attractor set computations (see lines 8 and 10) to identify those vertices that can be forced to visit the maximal priority; the remaining vertices are then solved recursively (see line 9).

Assuming that the maximal priority is of parity α, the outcome of the first recursion yields a dominion for $\bar{\alpha}$ in the subgame. In case this dominion is $\bar{\alpha}$-maximal in the entire game, α wins all vertices outside the dominion of player $\bar{\alpha}$ (lines 11-12). This can be seen as follows: since the game restricted to $V \setminus B$ (where B is computed in line 10) is $\bar{\alpha}$-closed, the opponent can choose to stay within the subgame W_{α}; choosing to do so means she will lose. So the only option she has is to escape. But the only escape she has leads, via A, to the maximal priority which has the parity of α. Note that if α is forced to leave this maximal priority she will again end up in the same subgame; any play that does so *ad infinitum* visits the maximal priority infinitely often and is therefore won by α. In case $\bar{\alpha}$'s dominion is *not* maximal, we can remove its $\bar{\alpha}$-attractor B from the game and recursively solve the remaining subgame and use its solution to construct the solution of the entire game (lines 14 and 15).

Algorithm 1 Zielonka's Algorithm

1: **function** ZIELONKA(G)	11: **if** $B = W'_{\bar{\alpha}}$ **then**
2: **if** $V = \emptyset$ **then**	12: $(W_{\alpha}, W_{\bar{\alpha}}) \leftarrow (A \cup W'_{\alpha}, B)$
3: $(W_0, W_1) \leftarrow (\emptyset, \emptyset)$	13: **else**
4: **else**	14: $(W'_0, W'_1) \leftarrow$ ZIELONKA($G \setminus B$)
5: $m \leftarrow \max\{p(v) \mid v \in V\}$	15: $(W_{\alpha}, W_{\bar{\alpha}}) \leftarrow (W'_{\alpha}, W'_{\bar{\alpha}} \cup B)$
6: $\alpha \leftarrow m \mod 2$	16: **end if**
7: $U \leftarrow \{v \in V \mid p(v) = m\}$	17: **end if**
8: $A \leftarrow Attr_{\alpha}(U)$	18: **return** (W_0, W_1)
9: $(W'_0, W'_1) \leftarrow$ ZIELONKA($G \setminus A$)	19: **end function**
10: $B \leftarrow Attr_{\bar{\alpha}}(W'_{\bar{\alpha}})$	

Note that rather than solving the entire game at once, one can first decompose a game into its strongly connected components (SCCs) and first solve the bottom SCCs. The SCC decomposition can be integrated tightly in the algorithm so that in each recursive call first an SCC decomposition is performed. While this may sound expensive, in [19] it is shown that this actually allows the algorithm to run in polynomial time on many practically relevant classes of special games such as solitaire and dull games for which the original algorithm might require exponential time otherwise.

3.2   Priority Promotion

The recent *Priority Promotion* (PP) algorithm [6] starts, like Zielonka's recursive algorithm, with a game decomposition that aims at identifying dominions for a given player. In contrast to Zielonka's algorithm, however, the PP algorithm does not explicitly solve subgames. Instead, within a fixed game, it uses a dominion searcher that maintains a set of vertices, along with an updated (promoted) priority mapping and zooms in on a dominion within that set of vertices. Once an α-dominion is returned by the dominion searcher, the search for another dominion continues until the entire game is solved, see Algorithm 2.

Algorithm 2 Priority Promotion Solver

1: **function** SOLVEPP(G)
2:       $(W_0, W_1) \leftarrow (\emptyset, \emptyset)$
3:       $U \leftarrow V$
4:       **while** $U \neq \emptyset$ **do**
5:            $m \leftarrow \max\{p(v) \mid v \in U\}$
6:            $(W_0', W_1') \leftarrow$ SEARCHDOMINION($G \cap U, U, p, m$)
7:            $(W_0, W_1) \leftarrow (W_0 \cup W_0', W_1 \cup W_1')$
8:            $U \leftarrow U \setminus (W_0' \cup W_1')$
9:       **end while**
10:      **return** (W_0, W_1)
11: **end function**

The main complexity and novelty of the PP algorithm lies in the way it identifies a dominion. To this end, it relaxes the notion of a dominion to a *quasi α-dominion*. A quasi α-dominion is a set U of vertices for which α has a strategy that guarantees that every play that remains within U is won by α, or that exits U via the $\bar{\alpha}$-*escape* of U. The α-escape of a set U, denoted $esc_G^\alpha(U)$ is the set of vertices from which α can force play out of U in a single move; *i.e.*, $esc_G^\alpha(U) = pre_\alpha(V \setminus U) \cap U$. A quasi α-dominion U in a game G is said to be α-*closed* iff $esc_G^\alpha(U) = \emptyset$; otherwise it is said to be α-open. Note that an α-closed quasi α-dominion is an α-dominion. When the $\bar{\alpha}$-escape set of a quasi α-dominion only contains vertices with the highest priority in the game, the α quasi-dominion is called an α-*region*.

By searching for quasi-dominions, the algorithm avoids the 'hard' problem of identifying the (set of) priorities that lead to a player winning its dominion. Under certain conditions, these quasi-dominions offer just enough guarantees that when composed, the result is a dominion. Regions meet these conditions. The dominion searcher, see Algorithm 3, essentially traverses a partial order of *states* which contain information about the open regions the searcher has computed up to that point. Once the searcher finds a closed region it terminates and returns this dominion to the solver. Conceptually the searcher assigns to each α-region it finds a priority that under-approximates the best value the opponent $\bar{\alpha}$ must visit when escaping from the α-region. A higher α-region U_1 can then be merged with a lower region U_2, yielding a new α-region and improving the under-approximation of U_2 by promoting its best escape value.

The searcher first checks whether the set of vertices U that dominate the current state induces a closed region A within the entire game; if so, the region is returned and the search is finished (lines 2-8). If the region is open, the opponent $\bar{\alpha}$ may escape. When the region is open in the current state, she may try to escape to some inferior priority within the subgraph; if the region is closed in the current state, she can only escape to some region in the larger game and end up in a priority that dominates m_g.

In case the region is open in the current state (line 11-14), the priority function of the current state is updated to set all vertices in A to the currently dominating priority m_g. This is achieved by the update $p_g^* \leftarrow p_g[U \mapsto m_g]$ which is defined as $p_g^*(v) = m_g$ in case $v \in U$ and $p_g^*(v) = p_g(v)$ otherwise. The new priority for the subgraph that will be explored is set to the next largest priority in the graph (line 13) and A is removed from the subgraph (line 14). This new state is then explored recursively (see line 20). In case the region is closed in the current state (line 16-18), the opponent must escape (if she wants to) to vertices in the larger game. The *best escape priority* she can force is the minimal priority among the set

---

Algorithm 3 Priority Promotion Dominion Searcher

1: **function** SEARCHDOMINION(G, V_g, p_g, m_g)
2:     $\alpha \leftarrow m_g \mod 2$
3:     $U \leftarrow \{v \in V_g \mid p_g(v) = m_g\}$
4:     $A \leftarrow Attr_\alpha(V_g, U)$
5:     $X \leftarrow esc_G^{\bar{\alpha}}(A)$
6:     **if** $X = \emptyset$ **then**
7:         $(W_\alpha', W_{\bar{\alpha}}') \leftarrow (Attr_\alpha(A), \emptyset)$
8:         **return** (W_0', W_1')
9:     **else**
10:         $X \leftarrow esc_{G \cap V_g}^{\bar{\alpha}}(A)$
11:         **if** $X \neq \emptyset$ **then**

12:             $p_g^* \leftarrow p_g[A \mapsto m_g]$
13:             $m_g^* \leftarrow \max\{p_g^*(v) \mid v \in V_g \wedge p_g^*(v) < m_g\}$
14:             $V_g^* \leftarrow V_g \setminus A$
15:         **else**
16:             $m_g^* \leftarrow bep_{p_g}^{\bar{\alpha}}(A)$
17:             $p_g^* \leftarrow (p \uplus p_g^{\geq m_g^*})[A \mapsto m_g^*]$
18:             $V_g^* \leftarrow \{v \in V \mid p_g^*(v) \leq m_g^*\}$
19:         **end if**
20:         **return** SEARCHDOMINION(G, V_g^*, p_g^*, m_g^*)
21:     **end if**
22: **end function**

---

of vertices that she can force to reach in a single step. This is given by the function $bep_{p_g}^{\alpha}(A)$ which yields the minimum priority (according to p_g) of the set $\{w \in V \setminus A \mid \exists v \in A \cap V_\alpha : w \in vE\}$. The dominating priority m_g is updated to m_g^* to reflect this best escape priority (see line 16) after which p_g is updated to p_g^*. In this update, all vertices in A are set to priority m_g^* while all vertices with priorities exceeding m_g^* remain unchanged. All vertices with priorities dominated by m_g^* are reset to their original value. This is achieved by the update in line 17. Here, $p_g^{\geq m}$ yields the partial function that coincides with p_g on the (maximal) domain $V' \subseteq V$ for which $p_g(v) \geq m$ for all $v \in V'$, and is undefined elsewhere. The *update* $p \uplus p_g$, for a partial function p_g is then defined as $(p \uplus p_g)(v) = p(v)$ in case $v \notin \mathrm{dom}(p_g)$ and $p_g(v)$ otherwise. The subgraph of G that is explored next is set to all vertices with (promoted) priority no larger than m_g^*. This newly constructed state is then again recursively explored (see line 20).

The PP algorithm resets all information in lower regions when promoting vertices. As observed in [4, 5], this can be improved by only resetting regions of the opponent, or even a subset of that. While this affects the performance for some artificial worst-case examples, these improvements do not lead to improvements on games stemming from practical applications, nor on random games [14].

3.3   The Fixpoint-Iteration Algorithm

The third algorithm we consider is the so-called *Fixpoint-Iteration* (FI) algorithm. It is essentially based on the reduction of the parity game solving problem to the modal μ-calculus model checking problem. This transformation views a parity game as a graph with state labels indicating the owner (by associating a proposition Even or Odd to a state) and priority of a vertex (by associating a proposition $prio_i$ to a state representing a vertex with priority i), while the winning regions W_0 and W_1 of the game are characterised by a so-called Walukiewicz μ-calculus formula.

Rather than using a general-purpose modal μ-calculus model checker for model checking a parity game, we here recall the algorithm given by Bruse *et al.* [8], see Algorithm 4. This algorithm works directly on a (pre-processed) parity game, essentially checking the following formula[1] for games with least priority 0, and a similar formula for games with least priority 1:

$$\sigma X_{d-1} \ldots \mu X_1. \nu X_0. \left(\left(\text{Even} \wedge \bigvee_{i=0}^{d-1} \diamond (\text{prio}_i \wedge X_i) \right) \vee \left(\text{Odd} \wedge \bigwedge_{i=0}^{d-1} \square (\overline{\text{prio}_i} \vee X_i) \right) \right)$$

---

[1]Here, μ and ν are the least and greatest fixpoint, respectively, $\diamond$ is the *may* modality, $\square$ is the *must* modality, Even, Odd and $prio_i$ are propositions and $\overline{prio_i}$ is the negation of $prio_i$.

The pre-processing guarantees that the parity game is *compressed*: *i.e.*, it ensures that the set of priorities occurring in the game constitute an integer interval with least priority 0 or 1. Note that converting a parity game to a compressed parity game can be done without changing the solution to the game. The

---

Algorithm 4 Fixpoint-Iteration algorithm for compressed parity games with least priority 0 and maximal priority $d-1$. The algorithm for compressed parity games with least priority 1 is analogous.

```
 1: function FI(G)                                          11:        X'_0, X_0 ← X_0, DIAMOND_BOX()
 2:     function DIAMOND_BOX                                 12:        i ← 0
 3:         U_0 ← {v ∈ V_0 | ∃w ∈ vE : w ∈ X_{p(w)}}        13:        while X_i = X'_i and i < d−1 do
 4:         U_1 ← {v ∈ V_1 | ∀w ∈ vE : w ∈ X_{p(w)}}        14:            i ← i+1
 5:         return U_0 ∪ U_1                                 15:            X'_i, X_i ← X_i, X_{i−1}
 6:     end function                                         16:            X_{i−1} ← if i is odd then V else ∅
 7:     for i ← d−1,…,0 do                                   17:        end while
 8:         X_i ← if i is even then V else ∅                 18:     until i = d−1 and X_{d−1} = X'_{d−1}
 9:     end for                                              19:     return (X_{d−1}, V \ X_{d−1})
10:     repeat                                               20: end function
```

---

main loop of the algorithm evaluates the fixpoint-free part of the Walukiewicz formula using the current values for the variables $X_0 \ldots X_{d-1}$. The evaluation gives rise to new values for these variables until a fixpoint is reached for all variables; at that moment, variable X_{d-1} contains all vertices that satisfy the formula, and, by construction, are won by player 0.

Note that the re-initialisation of all variables $X_0 \ldots X_{d-1}$ can be optimised by exploiting monotonicity of the formula, reducing the overall complexity from exponential in d to exponential in $\lceil d/2 \rceil$, see *e.g.* [27]. This can be done by first computing the smallest index i for which X_i is not yet stable, and only resetting $X_{i-1}, X_{i-3}, \ldots$ afterwards, thus yielding an algorithm that is closely related to the small progress measures algorithm, see [8]. We have included this optimisation in our implementation.

3.4 The APT Algorithm

The APT[2] algorithm, introduced by Kupferman and Vardi [25], is studied and implemented in [28]. This algorithm solves parity games by solving the emptiness problem of a corresponding alternating parity word automaton, which, in turn, is solved by checking the emptiness of an equivalent weak alternating word automaton. Rather than explicitly constructing this weak alternating word automaton, the APT algorithm, see Algorithm 5, constructs this automaton implicitly, checking emptiness on-the-fly.

APT uses two disjoint sets of vertices V_V (*Visiting* vertices) and V_A (*Avoiding* vertices) for solving a parity game. It declares a vertex *Visiting* for a player just whenever by reaching that vertex she can force a winning play; dually, a vertex is *Avoiding* for a player whenever by reaching that vertex she cannot force any winning play. These sets derive their meaning from a generalisation of the parity game winning condition, in which player 0 wins a play as soon as V_V is visited, player 1 wins as soon as V_A is visited, and the traditional parity game condition applies if neither V_V nor V_A are ever visited. Furthermore, the algorithm keeps track of an index i, representing the most significant priority in the game for which the winners of the vertices with that priority have not yet been decided for the given sets V_V and V_A.

---

[2]The description of the algorithm in [28] seems to contain several small mistakes and implementing it that way yields incorrect solutions; our exposition is based on their actual OCaml implementation in PGSolver.

Algorithm 5 The APT algorithm for compressed parity games with minimal priority m and maximal priority d.

1: **function** APT(G)
2:     **function** WIN(α, i, V_V, V_A)
3:         **return** $V \setminus \text{FP}(\bar{\alpha}, i, V_V, V_A)$ **if** $i \geq m$ **else** $pre_\alpha(V_V)$
4:     **end function**
5:     **function** FP(α, i, V_V, V_A)
6:         $X, F \leftarrow \emptyset, \{v \in V \mid p(v) = i\}$
7:         $Y \leftarrow \text{WIN}(\alpha, i-1, V_A \cup (F \setminus X), V_V \cup (F \cap X))$
8:         **while** $Y \neq X$ **do**
9:             $X \leftarrow Y$
10:             $Y \leftarrow \text{WIN}(\alpha, i-1, V_A \cup (F \cap X), V_V \cup (F \setminus X))$
11:         **end while**
12:     **end function**
13:     $\alpha \leftarrow (d+1) \mod 2$
14:     $X \leftarrow \text{WIN}(\alpha, d, \emptyset, \emptyset)$
15:     **return** $(X, V \setminus X)$ **if** $\alpha = 0$ **else** $(V \setminus X, X)$
16: **end function**

---

4 Implementing Parity Game Solvers Using BDDs

Using an efficient data structure to explicitly represent a graph, one may be able to represent parity games up to 2^{25} vertices on a 16Gb main memory machine. Parity games encoding verification problems of (software or hardware) systems easily require more memory. A symbolic representation of the game graph may then help to sidestep this problem. Binary Decision Diagrams (BDDs) provide just such a representation. We assume that the reader is familiar with BDDs and the algorithms manipulating these; for a comprehensive treatment we refer to [29, 15].

BDDs can be used to canonically (and often concisely) represent propositional formulae, which in turn can be used to characterise sets. In our setting we have to represent the following sets: the set of vertices V of a game, the edges E, the partition (V_0, V_1) of V and the priority function p. For most parity games stemming from practical verification problems, the number of distinct priorities in the game is small. For that reason, it suffices to represent p as a mapping from priorities to sets of vertices. For the computations involved in the parity game solving algorithms of the previous section we provide the key ingredients below; each operation is described using an expression for the BDD-based implementation.

Henceforward we assume that $\vec{x}$ is the vector of Boolean variables used to span the set of vertices V of a given parity game; we assume ι is a total injective mapping from V to truth-assignments for $\vec{x}$. The i-th Boolean variable of $\vec{x}$ is denoted $\vec{x}_i$; this notation extends to other vectors of Boolean variables. We represent a set $U \subseteq V$ by a propositional formula encoded as a BDD $\mathscr{U}(\vec{x})$ that ranges over $\vec{x}$; as is standard, for a truth-assignment $\vec{u}$ for $\vec{x}$ we have $\mathscr{U}(\vec{u}) = 1$ iff $\iota(v) = \vec{u}$ for some $v \in U$. This way, we can represent, *e.g.*, the sets V_0 and V_1 by BDDs $\mathscr{V}_0(\vec{x})$ and $\mathscr{V}_1(\vec{x})$. The set of edges is given by a BDD $\mathscr{E}(\vec{x}, \vec{x}')$ that ranges over $\vec{x}$ and $\vec{x}'$, where $\vec{x}'$ is the vector of Boolean variables used to represent a successor vertex. Finally, the collection of BDDs $\mathscr{P}_i(\vec{x})$ represents the set of vertices with priority i.

Set comparison, intersection, union and complement are easily expressed as BDD operations using equivalence, conjunction, disjunction and negation. The pre-image $pre_\alpha(U)$, which is the basis for the APT algorithm, can then be encoded using existential quantification and the above operations:

$$pre_\alpha(\mathscr{U}(\vec{x})) = (\mathscr{V}_\alpha(\vec{x}) \wedge \exists \vec{x}'.(\mathscr{E}(\vec{x}, \vec{x}') \wedge \mathscr{U}(\vec{x}'))) \vee (\mathscr{V}_{\bar{\alpha}}(\vec{x}) \wedge \neg \exists \vec{x}'.(\mathscr{E}(\vec{x}, \vec{x}') \wedge \neg \mathscr{U}(\vec{x}')))$$

Using the pre-image, the α-attractor $Attr_\alpha(\mathscr{U}(\vec{x}))$ and the confined α-attractor $Attr_\alpha(\mathscr{T}(\vec{x}), \mathscr{U}(\vec{x}))$ can then be implemented effectively as a least fixpoint computation. Note also that the operation $esc_G^\alpha(A)$, used in lines 5 and 10 of the PP algorithm, can be implemented using the pre-image operation. The DIAMONDBOX() function of the FI algorithm can be implemented in a similar manner:

$$(\mathscr{V}_0(\vec{x}) \wedge \bigvee_i (\exists \vec{x}'.(\mathscr{E}(\vec{x}, \vec{x}') \wedge \mathscr{P}_i(\vec{x}') \wedge \mathscr{X}_i(\vec{x}')))) \vee (\mathscr{V}_1(\vec{x}) \wedge \neg \bigvee_i (\exists \vec{x}'.(\mathscr{E}(\vec{x}, \vec{x}') \wedge \mathscr{P}_i(\vec{x}') \wedge \neg \mathscr{X}_i(\vec{x}'))))$$

where, for each i, $\mathscr{X}_i(\vec{x})$ is the set of vertices represented by approximation X_i in Algorithm 4.

The operations in the PP algorithm that remain to be encoded involve the computation of the best escape priority and the computation of new priority mappings. For the best escape priority we are required to compute the minimal priority among a set of successor vertices of a given set of vertices U. We split this computation in two parts. First we identify the set of successor vertices using the operation $post(U)$ defined as the union of uE, for all $u \in U$; its BDD encoding is standard: $post(\mathscr{U}(\vec{x})) = (\exists \vec{x}. \mathscr{E}(\vec{x}, \vec{x}') \wedge \mathscr{U}(\vec{x}))[\vec{x}' \mapsto \vec{x}]$, where $[\vec{x}' \mapsto \vec{x}]$ denotes a variable renaming. Computing $bep^{\alpha}_{p_g}(A)$, done in line 16 of Alg. 3, then can be implemented by searching for the least value m in the domain of mapping p_g (assuming that we have BDDs $\mathscr{P}_{g,i}(\vec{x})$ representing the set of vertices $\{v \in V \mid p_g = i\}$) for which the BDD $\neg \mathscr{A}(\vec{x}) \wedge \mathscr{P}_{g,m}(\vec{x}) \wedge post(\mathscr{A}(\vec{x}) \wedge \mathscr{V}_{\alpha}(\vec{x}))$ is satisfiable. This can be done using a simple iteration.

Finally, the computations of p_g^* in lines 12 and 17 of Alg. 3 are simple updates of the function p_g which can be done by iterating over the domain of the mapping p (resp. p_g) and computing new BDDs using conjunction and disjunction of the BDDs given by p and p_g. More specifically, for updates of the form $(p \uplus p_g^{\geq m})[A \mapsto m]$, we first compute the set of vertices in the domain of p_g that are above m:

$$\mathscr{A}bove_{\geq m}(\vec{x}) = \bigvee \{\mathscr{P}_{g,i}(\vec{x}) \mid i \geq m\}$$

The updated priority function p_g^* coinciding with $(p \uplus p_g^{\geq m})[A \mapsto m]$, where A is encoded by BDD $\mathscr{A}(\vec{x})$ and the set of vertices $\{v \in V \mid p(v) = i\}$ is encoded by $\mathscr{P}_i(\vec{x})$ is then given by:

$$\mathscr{P}_{g,i}^*(\vec{x}) = \begin{cases} \neg \mathscr{A}(\vec{x}) \wedge ((\mathscr{P}_i(\vec{x}) \wedge \neg \mathscr{A}bove_{\geq m}(\vec{x})) \vee (\mathscr{P}_{g,i}(\vec{x}) \wedge \mathscr{A}bove_{\geq m}(\vec{x}))) & \text{if } i \neq m \\ \mathscr{A}(\vec{x}) \vee (\mathscr{P}_i(\vec{x}) \wedge \neg \mathscr{A}bove_{\geq m}(\vec{x})) \vee (\mathscr{P}_{g,i}(\vec{x}) \wedge \mathscr{A}bove_{\geq m}(\vec{x})) & \text{else} \end{cases}$$

5 Experimental Evaluation

We have implemented all four algorithms in Python, utilising the BDD package dd[3] which, next to a native Python BDD implementation, offers wrappers to C-based BDD implementations such as CUDD, Sylvan and BuDDy. For the timings we report on for our experiments we rely on the native Python implementation as our independent experiments suggest that its performance is comparable to that of CUDD. Note that the choice of programming language is secondary given that all essential time-consuming computations involve BDD creation and manipulation. We conducted all experiments on a Macbook Pro, 3.5 GHz Intel Core i7 (13-inch, 2017 model, *i.e.*, dual core), with 16 Gb 2133 MHz LPDDR3 main memory, running macOS 10.13.4.

We compare the performance of the implementation of all four algorithms as described in Section 4, and, for reference, we also compare these to explicit-state state-of-the-art C++ implementations of both the Zielonka and priority promotion algorithms offered by the `pbespgsolve` tool of the mCRL2 toolset [13]. First, we assess their performance on random games with different characteristic. Such games, however, may be poor predictors for the performance of the algorithms on practical problems. Second, since one of our aims is to assess the scalability and the speed at which BDD-based solvers can solve parity games that encode practical verification problems, we also compare their performance on some of the larger cases of the Keiren benchmark set [24].

---

[3]See `https://github.com/johnyf/dd`, version 0.5.4, by Ioannis Filippidis

5.1 Random Parity Games

PGSolver [18] is a collection of parity game solving algorithms programmed in OCaml, which includes a tool `randomgame` to generate (possibly self-loop-free) random parity games with a fixed number of vertices. Running `randomgame N P L H x` generates a game with N vertices; each vertex v is assigned a priority that is chosen uniformly from the interval $[0, P-1]$, and ownership 0 with probability $\frac{1}{2}$ and ownership 1 with probability $\frac{1}{2}$. Each vertex is connected to d distinct neighbours, where d is chosen uniformly from the interval $[L, H]$. Since the tool generates explicit graphs in the PGSolver format, we automatically converted these games to BDDs using a binary encoding for the vertices.

For our first batch of experiments, we considered both dense and low out-degree, self-loop-free random games with a high number of distinct priorities, varying $N \in \{250, 500, 1\,000, 2\,000, 4\,000\}$:

- low out-degree random games generated using `randomgame N N 1 2 x`;

- dense random games generated using `randomgame N N 1 N x`;

For each class of random games and each value of N we generated 20 parity games. Table 1 lists the cumulative runtime (in seconds) for solving these parity games. We set a timeout of 12 000 (cumulative) seconds for each experiment; a timeout is denoted by $\dagger_T$. Experiments requiring over 15Gb memory are aborted; we indicate this by $\dagger_M$.

	low out-degree games					*dense games*				
Solver	250	500	1 000	2 000	4 000	250	500	1 000	2 000	4 000
Zielonka	9	53	302	1 130	3 463	43	228	1 004	$\dagger_M$	$\dagger_M$
PP	10	63	316	815	3 048	50	227	1 039	$\dagger_M$	$\dagger_M$
FI	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_M$	$\dagger_M$
APT	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_T$	$\dagger_M$	$\dagger_M$
Zielonka (explicit)	0	0	0	0	1	0	0	0	0	5
PP (explicit)	0	0	0	0	0	0	0	0	0	3

Table 1: Cumulative time (in seconds) to solve the random games for a given value of N. We set a timeout of 12 000 seconds (indicated by $\dagger_T$) for the cumulative solving time; $\dagger_M$ indicates out-of-memory.

We observe that for both types of games, the explicit Zielonka and PP solvers substantially outperform the symbolic solvers, see Table 1. Of note is that both APT and FI fail on all accounts, which we believe is due to the large number of distinct priorities in the random games, causing both algorithms to require an excessively large number of iterations. The symbolic PP implementation outperforms the symbolic Zielonka implementation for low out-degree games that have large number of vertices and different priorities, but for smaller number of vertices and small number of different priorities, Zielonka performs better, but only marginally so. For games that consist of more than 2 000 vertices, the BDD approach fails dramatically, especially in the case of dense games.

We can trace the poor performance of the symbolic solvers to the overly complex and unstructured BDDs underlying the random graphs. These unwieldy BDD representations are caused by the lack of structure provided by the binary encoding of vertices and edges, leading to a poor compression ratio: typically, BDDs remain concise if two adjacent vertices (*i.e.*, vertices related by *E*) only differ in their representation by a relatively small number of bits. The binary encoding, however, does not guarantee this. Variable-reordering heuristics may help in some cases but are not guaranteed to do so. As a result, the BDDs themselves require an excessive amount of memory (as witnessed by the out-of-memory issues in modest-sized dense random graphs), and operations on these BDDs lead to excessive running times. As such, these experiments may not be good predictors for running times of the symbolic algorithms on games encoding practical verification problems. However, they do suggest that FI and APT both suffer

significantly from a poor compression ratio. Zielonka and PP, on the other hand, are (relatively) less susceptible to a poor compression ratio.

For our next batch of random experiments, we repeat the experiments conducted in [28]. In that work, the authors observe that the APT algorithm outperforms Zielonka's algorithm for games in which the number of vertices is exponential in the number of priorities in a game, as long as the base is sufficiently large: they show that APT outperforms Zielonka for $N \approx c^P$, for $c = e$ or $c = 10$, but for $c = 2$, Zielonka outperforms APT. Games in which the number of vertices is exponential in the number of distinct priorities are typical of games encoding practical verification and synthesis problems. We generate dense random games with a priority/vertex ratio given by the relation $N \approx c^P$, for $c \in \{2, e, 10, 13\}$, using `randomgame N P 1 N x`; *i.e.*, we extend the experiments of [28] by also considering $c = 13$. We again generate 20 random games per experiment, and set a timeout at 12 000 seconds (cumulative), see Table 2.

| | $N = 2^P$ *games* | | | | | $N \approx e^P$ *games* | | | | | $N = 10^P$ *games* | | | $N = 13^P$ *games* | | |
Solver	128	256	512	1 024	2 048	21	55	149	404	1 097	10	100	1 000	13	169	2197
Zielonka	9	33	127	559	2 203	1	2	14	88	646	0	7	514	0	18	3099
PP	10	37	132	543	2 355	1	2	15	91	717	0	8	544	0	19	4299
FI	348	2 212	$†_T$	$†_T$	$†_T$	1	6	89	955	10 695	0	6	530	0	12	2912
APT	173	1 092	8 195	$†_T$	$†_T$	1	4	61	659	7 526	0	5	545	0	13	3982
Zielonka (explicit)	0	0	0	0	1	0	0	0	1	4	0	0	3	0	0	17
PP (explicit)	0	0	0	0	1	0	0	0	1	4	0	0	3	0	0	17

Table 2: Cumulative time (in seconds) to solve the random games for a given value of N. A timeout (a cumulative time of 12 000 seconds or more) is indicated by $†_T$.

Table 2 illustrates that Zielonka has a slight edge over PP. Comparing the results in Table 2 to the experimental outcomes in [28], we find some interesting differences. First, whereas in [28], the APT algorithm is found to significantly outperform Zielonka's algorithm for $c = e$ and $c = 10$, our results do not provide such clear indications. We do see that for $c = 10$ and $c = 13$, the running times of the APT algorithm becomes more in line with that of Zielonka and PP. The behaviour of FI illustrates a more pronounced correlation between the number of vertices and the number of priorities in a game, suggesting that in a symbolic setting, the FI algorithm may actually perform quite well on games that encode verification problems that typically result in only a few distinct priorities. We test this hypothesis by comparing the performance on a number of (large) parity games originating from model checking problems, mostly taken from the Keiren Benchmark set.

5.2 The Keiren Parity Game Solving Benchmark Set

In [24], Keiren describes and provides a set of parity games that originate from over 300 model checking problems, originating from 21 specifications. The main obstacle in reusing the data set is that the games are encoded as explicit graphs in the PGSolver format. This means that most of the structure that a BDD solver can typically exploit for compactly representing the game graph is lost. As we concluded from the random games generated by PGSolver, a binary encoding of the sets involved leads to severe performance degradation of the solvers, and given the size of the graphs, there is little hope the running times of our algorithms on such encoded BDDs provide meaningful information. We have coped with this by generating BDDs from the original specifications for several of the model checking games included in the Keiren benchmark set.[4] Note that the conversion of the original specifications to BDDs

---

[4]To this end, we developed a special-purpose `pbes2bdd` tool for converting (a fragment of) parameterised Boolean equation systems in Standard Recursive Form, containing only Boolean data parameters, to BDDs. The tool currently does not guarantee

is not straightforward: most games stem from model checking problems for system models that employ (unbounded) lists, natural numbers, *etcetera*, for which no trivial bounds can be established. The results can be found in Table 3; here, $\dagger_M$ indicates that the computation was aborted because it needed more than 15Gb of memory.

Model	Property	#prio	Zielonka (sym/exp)	PP (sym/exp)	FI	APT	#$\vec{x}$
Chatbox	Absence of Deadlock	2	0.04/0.02	0.05/0.01	**0.00**	0.01	18
	It is possible to JOIN infinitely often	2	0.02/0.02	0.02/0.01	**0.00**	0.01	18
	Invariantly LEAVE can follow JOIN	3	0.04/0.07	0.05/0.02	**0.01**	0.02	18
Onebit$_2$	Infinitely often send and receive	3	0.59/0.14	0.58/**0.04**	0.64	0.69	26
	Absence of Deadlock	2	0.16/0.11	0.19/**0.03**	0.04	0.09	26
Onebit$_8$	Infinitely often send and receive	3	3.30/24.17	3.28/4.49	3.36	**3.19**	38
	Absence of Deadlock	2	0.86/22.70	1.02/6.82	**0.19**	0.47	38
Onebit$_{32}$	Infinitely often send and receive	3	10.93/$\dagger_M$	10.68/$\dagger_M$	10.22	**9.81**	50
	Absence of Deadlock	2	2.42/$\dagger_M$	3.16/$\dagger_M$	**0.52**	1.65	50
Hesselink$_4$	Absence of Deadlock	2	1.97/$\dagger_M$	1.87/$\dagger_M$	**1.21**	1.31	60
	Cache consistency	2	3.47/$\dagger_M$	3.55/$\dagger_M$	**1.55**	1.93	61
Dining$_{13}$	Absence of Deadlock	2	1.45/8.40	1.46/7.68	0.20	**0.19**	89
	Aristotle can eat infinitely often	2	7.58/3.44	8.00/**1.18**	4.24	4.03	89
	Aristotle can, and thus will eat infinitely often	3	1.69/10.35	1.56/4.74	0.43	**0.37**	90
	Cumulative:		34.52/$\dagger_M$	35.47/$\dagger_M$	**22.61**	23.77	

Table 3: Running times (in seconds) for selected model checking problems; $\dagger_M$ indicates an out-of-memory error; #$\vec{x}$ indicates the number of Boolean variables needed to encode the set of vertices.

We again compared our BDD-based implementations of the four parity game solving algorithms to the explicit Zielonka and PP solvers of the `pbespgsolve` tool[5] (indicated by the 'exp' fields in Table 3).

First, observe that all four symbolic solvers perform quite well on all games, sometimes significantly outperforming the two explicit solvers. In particular, none of the symbolic solvers run out-of-memory; in fact, memory usage for all games stays well below 500Mb. This sharply contrasts the explicit solvers, which, for the larger games that can be solved consume 8Gb or more. We furthermore note that the runtime performance of the symbolic solvers appears to be more robust than that of the explicit solvers.

Second, we observe that both FI and APT clearly outperform Zielonka and PP. The PP and Zielonka algorithms perform comparably on these games, with Zielonka having an insignificant edge over PP; in this regard, the outcomes of these experiments are largely similar to the collective outcome of our experiments using random games.

6 Discussion and Conclusions

We studied four algorithms for solving parity games, *viz.* Zielonka's recursive algorithm, the recently introduced Priority Promotion (PP) algorithm, the Fixpoint-Iteration (FI) algorithm and the APT algorithm, and compared the performance of their symbolic implementations.

---

totality of the generated parity game; instead, it relies on the user to guarantee that this is the case.

[5]Of note is the fact that for our explicit solvers, we only compute the solution to the vertices reachable from a designated vertex (typically the vertex that indicates whether the modal μ-calculus formula holds for the initial state of a system); we have excluded the (typically significant) time required for constructing this reachable subgame and only list the time required to solve the resulting (sub)game. The game that is solved by the explicit solvers can thus be orders of magnitude smaller than the games solved by our symbolic solvers, as these do not conduct a reachability analysis.

Our results indicate that, overall, the symbolic implementations of Zielonka's algorithm and the PP algorithm have very similar performance characteristics. The experiments using random games indicate that Zielonka's algorithm performs slightly better than PP for games with a small number of priorities, and for games in which the number of vertices is exponentially larger than the number of distinct priorities, but the differences are not significant in most cases. The same observation applies to the cases taken from Keiren's benchmark set.

The symbolic implementations of the APT and FI algorithms both perform poorly on most of our random games. There are indications that also in the symbolic case, these algorithms are more suited for games in which the number of vertices is exponentially larger than the number of distinct priorities, but the results are not as pronounced as reported in [28] for the explicit case. We believe that part of this difference is due to the susceptibility both algorithms demonstrate on the size of the BDDs. This is further supported by the cases taken from Keiren's benchmark set, in which the BDDs remain rather small and both FI and APT outperform Zielonka and PP, but further research is needed to confirm this hypothesis and to see if there are ways to sidestep such issues.

While our experimental results on the Keiren benchmark set show that both FI and APT beat Zielonka and PP on games stemming from verification problems, their unpredictability in the presence of large BDDs makes them less reliable for solving such problems as one cannot upfront predict the size of BDDs in practice. Given that all solvers require at most seconds for even the largest models, containing up to 2^{90} vertices, any of the four algorithms would work in practice, and one may therefore favour the more robust PP or Zielonka algorithms in practice.

Finally, comparing the performance of our symbolic Zielonka and PP solvers to the explicit solvers, we find that there are only few cases in which the explicit solvers have an edge over the symbolic solvers. More importantly, the explicit solvers are unable to solve games which are readily solved by the symbolic solvers, indicating the practical significance of BDD-based solvers.

References

[1] L. de Alfaro & M. Faella (2007): *An Accelerated Algorithm for 3-Color Parity Games with an Application to Timed Games*. In: *CAV*, *LNCS* 4590, Springer, pp. 108–120, doi:10.1007/978-3-540-73368-3_13.

[2] A. Arnold, A. Vincent & I. Walukiewicz (2003): *Games for Synthesis of Controllers with Partial Observation*. *TCS* 303(1), pp. 7–34, doi:10.1016/S0304-3975(02)00442-5.

[3] M. Bakera, S. Edelkamp, P. Kissmann & C.D. Renner (2008): *Solving μ-Calculus Parity Games by Symbolic Planning*. In: *MoChArt*, *LNCS* 5348, Springer, pp. 15–33, doi:10.1007/978-3-642-00431-5_2.

[4] M. Benerecetti, D. Dell'Erba & F. Mogavero (2016): *A Delayed Promotion Policy for Parity Games*. In: *GandALF*, *EPTCS* 226, pp. 30–45, doi:10.4204/EPTCS.226.3.

[5] M. Benerecetti, D. Dell'Erba & F. Mogavero (2016): *Improving Priority Promotion for Parity Games*. In: *Haifa Verification Conference*, *LNCS* 10028, Springer, pp. 117–133, doi:10.1007/978-3-319-49052-6_8.

[6] M. Benerecetti, D. Dell'Erba & F. Mogavero (2016): *Solving Parity Games via Priority Promotion*. In: *CAV (2)*, *LNCS* 9780, Springer, pp. 270–290, doi:10.1007/978-3-319-41540-6_15.

[7] D. Berwanger, K. Chatterjee, M. De Wulf, L. Doyen & T.A. Henzinger (2009): *Alpaga: A Tool for Solving Parity Games with Imperfect Information*. In: *TACAS*, *LNCS* 5505, Springer, pp. 58–61, doi:10.1007/978-3-642-00768-2_7.

[8] F. Bruse, M. Falk & M. Lange (2014): *The Fixpoint-Iteration Algorithm for Parity Games*. In: *GandALF*, *EPTCS* 161, pp. 116–130. Available at https://doi.org/10.4204/EPTCS.161.12.

[9] D. Bustan, O. Kupferman & M.Y. Vardi (2004): *A Measured Collapse of the Modal μ-Calculus Alternation Hierarchy*. In: *STACS*, *LNCS* 2996, Springer, pp. 522–533, doi:10.1007/978-3-540-24749-4_46.

[10] C.S. Calude, S. Jain, B. Khoussainov, W. Li & F. Stephan (2017): *Deciding Parity Games in Quasipolynomial Time*. In: *STOC*, ACM, pp. 252–263, doi:10.1145/3055399.3055409.

[11] K. Chatterjee, W. Dvorák, M. Henzinger & V. Loitzenbauer (2017): *Improved Set-Based Symbolic Algorithms for Parity Games*. In: *CSL*, *LIPIcs* 82, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 18:1–18:21, doi:10.4230/LIPIcs.CSL.2017.18.

[12] T. Chen, B. Ploeger, J. van de Pol & T.A.C. Willemse (2007): *Equivalence Checking for Infinite Systems Using Parameterized Boolean Equation Systems*. In: *CONCUR'07*, pp. 120–135, doi:10.1007/978-3-540-74407-8_9.

[13] S. Cranen, J.F. Groote, J.J.A. Keiren, F. Stappers, E.P. de Vink, W. Wesselink & T.A.C. Willemse (2013): *An Overview of the mCRL2 Toolset and Its Recent Advances*. In: *TACAS*, *LNCS* 7795, Springer, pp. 199–213, doi:10.1007/978-3-642-36742-7_15.

[14] T. van Dijk (2018): *Oink: An Implementation and Evaluation of Modern Parity Game Solvers*. In: *TACAS (1)*, *LNCS* 10805, Springer, pp. 291–308, doi:10.1007/978-3-319-89960-2_16.

[15] R. Drechsler & B. Becker (1998): *Binary Decision Diagrams - Theory and Implementation*. Springer, doi:10.1007/978-1-4757-2892-7.

[16] E.A. Emerson & C.S. Jutla (1991): *Tree automata, Mu-Calculus and determinacy*. In: *FOCS'91*, IEEE Computer Society, Washington, DC, USA, pp. 368–377, doi:10.1109/SFCS.1991.185392.

[17] J. Fearnley, S. Jain, S. Schewe, F. Stephan & D. Wojtczak (2017): *An Ordered Approach to Solving Parity Games in Quasi Polynomial Time and Quasi Linear Space*. In: *SPIN*, ACM, pp. 112–121, doi:10.1145/3092282.3092286.

[18] O. Friedmann & M. Lange (2009): *Solving Parity Games in Practice*. In: *ATVA*, *LNCS* 5799, Springer, pp. 182–196, doi:10.1007/978-3-642-04761-9_15.

[19] M. Gazda & T.A.C. Willemse (2013): *Zielonka's Recursive Algorithm: dull, weak and solitaire games and tighter bounds*. In: *GandALF*, *EPTCS* 119, pp. 7–20, doi:10.4204/EPTCS.119.4.

[20] J.F. Groote & T.A.C. Willemse (2005): *Parameterised Boolean Equation Systems*. Theor. Comput. Sci. 343(3), pp. 332–369, doi:10.1016/j.tcs.2005.06.016.

[21] M. Jurdziński (2000): *Small Progress Measures for Solving Parity Games*. In: *STACS'00*, *LNCS* 1770, Springer, pp. 290–301, doi:10.1007/3-540-46541-3_24.

[22] M. Jurdziński & R. Lazic (2017): *Succinct Progress Measures for Solving Parity Games*. In: *LICS*, IEEE Computer Society, pp. 1–9, doi:10.1109/LICS.2017.8005092.

[23] G. Kant & J. van de Pol (2014): *Generating and Solving Symbolic Parity Games*. In: *GRAPHITE*, *EPTCS* 159, pp. 2–14, doi:10.4204/EPTCS.159.2.

[24] J.J.A. Keiren (2015): *Benchmarks for Parity Games*. In: *FSEN*, *LNCS* 9392, Springer, pp. 127–142, doi:10.1007/978-3-319-24644-4_9.

[25] O. Kupferman & M.Y. Vardi (1998): *Weak Alternating Automata and Tree Automata Emptiness*. In: *STOC*, ACM, pp. 224–233, doi:10.1145/276698.276748.

[26] R. McNaughton (1993): *Infinite Games Played on Finite Graphs*. APAL 65(2), pp. 149–184, doi:10.1016/0168-0072(93)90036-D.

[27] H. Seidl (1996): *Fast and Simple Nested Fixpoints*. Inf. Process. Lett. 59(6), pp. 303–308, doi:10.1016/0020-0190(96)00130-5.

[28] A. Di Stasio, A. Murano, G. Perelli & M.Y. Vardi (2016): *Solving Parity Games Using an Automata-Based Algorithm*. In: *CIAA*, *LNCS* 9705, Springer, pp. 64–76, doi:10.1007/978-3-319-40946-7_6.

[29] I. Wegener (2000): *Branching programs and binary decision diagrams*. SIAM, doi:10.1137/1.9780898719789.

[30] W. Zielonka (1998): *Infinite Games on Finitely Coloured Graphs with Applications to Automata on Infinite Trees*. TCS 200(1-2), pp. 135 – 183, doi:10.1016/S0304-3975(98)00009-7.