



# Oink: An Implementation and Evaluation of Modern Parity Game Solvers

Tom van Dijk<sup>(✉)</sup> 

Formal Models and Verification, Johannes Kepler University, Linz, Austria  
tom.vandijk@jku.at

**Abstract.** Parity games have important practical applications in formal verification and synthesis, especially to solve the model-checking problem of the modal  $\mu$ -calculus. They are also interesting from the theory perspective, as they are widely believed to admit a polynomial solution, but so far no such algorithm is known. In recent years, a number of new algorithms and improvements to existing algorithms have been proposed. We implement a new and easy to extend tool Oink, which is a high-performance implementation of modern parity game algorithms. We further present a comprehensive empirical evaluation of modern parity game algorithms and solvers, both on real world benchmarks and randomly generated games. Our experiments show that our new tool Oink outperforms the current state-of-the-art.

## 1 Introduction

Parity games are turn-based games played on a finite graph. Two players *Odd* and *Even* play an infinite game by moving a token along the edges of the graph. Each vertex is labeled with a natural number *priority* and the winner of the game is determined by the parity of the highest priority that is encountered infinitely often. Player Odd wins if this parity is odd; otherwise, player Even wins.

Parity games are interesting both for their practical applications and for complexity theoretic reasons. Their study has been motivated by their relation to many problems in formal verification and synthesis that can be reduced to the problem of solving parity games, as parity games capture the expressive power of nested least and greatest fixpoint operators [17]. In particular, deciding the winner of a parity game is polynomial-time equivalent to checking non-emptiness of non-deterministic parity tree automata [33], and to the explicit model-checking problem of the modal  $\mu$ -calculus [14, 15, 23, 32].

Parity games are interesting in complexity theory, as the problem of determining the winner of a parity game is known to lie in  $UP \cap co-UP$  [26], as well as in  $NP \cap co-NP$  [15]. This problem is therefore unlikely to be NP-complete and it is widely believed that a polynomial solution may exist. Despite much effort, no such algorithm has yet been found.

---

T. van Dijk—Supported by FWF, NFN Grant S11408-N23 (RiSE).

© The Author(s) 2018

D. Beyer and M. Huisman (Eds.): TACAS 2018, LNCS 10805, pp. 291–308, 2018.

[https://doi.org/10.1007/978-3-319-89960-2\\_16](https://doi.org/10.1007/978-3-319-89960-2_16)

Motivated by recent publications with both novel approaches and improvements to known algorithms, we implement a number of modern solvers in our new tool **Oink**, which aims to provide a high-performance implementation of parity game solvers. Oink is designed as a library that integrates with other tools and can easily be extended. We use Oink to provide a modern empirical evaluation of parity game solvers based on both real world benchmarks and randomly generated games.

In past publications new and improved algorithms are often tested against the implementation of Zielonka’s algorithm in the PGSolver tool [19]. However, various recent publications [1, 34, 40] suggest that much better performance can be obtained. We combine a number of improvements from the literature [34, 40, 41] and propose additional optimizations. We show that our implementation of Zielonka’s algorithm outperforms PGSolver by several orders of magnitude.

We describe Oink in Sect. 3 and provide accessible descriptions of the implemented state-of-the-art algorithms in Sects. 4–7. We implement the *strategy improvement* algorithm (Sect. 4), both the *small progress measures* and the recently proposed *quasi-polynomial progress measures* algorithms (Sect. 5), the well-known *Zielonka* algorithm (Sect. 6) as well as a number of related algorithms from the *priority promotion* family (Sect. 7). We also propose an alternative multi-core implementation of strategy improvement.

## 2 Preliminaries

Parity games are two-player turn-based infinite-duration games over a finite directed graph  $G = (V, E)$ , where every vertex belongs to exactly one of two players called player *Even* and player *Odd*, and where every vertex is assigned a natural number called the *priority*. Starting from some initial vertex, a play of both players is an infinite path in  $G$  where the owner of each vertex determines the next move. The winner of such an infinite play is determined by the parity of the highest priority that occurs infinitely often along the play.

More formally, a parity game  $\mathfrak{G}$  is a tuple  $(V_{\diamond}, V_{\square}, E, \text{pr})$  where  $V = V_{\diamond} \cup V_{\square}$  is a set of vertices partitioned into the sets  $V_{\diamond}$  controlled by player *Even* and  $V_{\square}$  controlled by player *Odd*, and  $E \subseteq V \times V$  is a total relation describing all possible moves, i.e., every vertex has at least one successor. We also write  $E(u)$  for all successors of  $u$  and  $u \rightarrow v$  for  $v \in E(u)$ . The function  $\text{pr}: V \rightarrow \{0, 1, \dots, d\}$  assigns to each vertex a *priority*, where  $d$  is the highest priority in the game.

We write  $\text{pr}(v)$  for the priority of a vertex  $v$  and  $\text{pr}(V)$  for the highest priority of a set of vertices  $V$  and  $\text{pr}(\mathfrak{G})$  for the highest priority in the game  $\mathfrak{G}$ . Furthermore, we write  $\text{pr}^{-1}(i)$  for all vertices with the priority  $i$ . A *path*  $\pi = v_0 v_1 \dots$  is a sequence of vertices consistent with  $E$ , i.e.,  $v_i \rightarrow v_{i+1}$  for all successive vertices. A *play* is an infinite path. We write  $\text{pr}(\pi)$  for the highest priority in  $\pi$  if  $\pi$  is finite, or the highest priority that occurs infinitely often if  $\pi$  is infinite. Player Even wins a play  $\pi$  if  $\text{pr}(\pi)$  is even; player Odd wins if  $\text{pr}(\pi)$  is odd.

A *strategy*  $\sigma: V \rightarrow V$  is a partial function that assigns to each vertex in its domain a single successor in  $E$ , i.e.,  $\sigma \subseteq E$ . We typically refer to a *strategy*

of player  $\alpha$  to restrict  $\sigma$  to all vertices controlled by player  $\alpha$ . A player wins a vertex if they have a strategy such that all plays consistent with this strategy are winning for the player. A fundamental result for parity games is that they are memoryless determined [13], i.e., each vertex is winning for exactly one player, and both players have a strategy for each of their winning vertices.

Algorithms for solving parity games frequently employ (variations of) *attractor computation*. Given a set of vertices  $A$ , the attractor of  $A$  for a player  $\alpha$  represents those vertices that  $\alpha$  can force the play toward. We write  $\text{Attr}_\alpha^\triangleright(A)$  to attract vertices in  $\triangleright$  to  $A$  as player  $\alpha$ , i.e.,

$$\mu Z . A \cup \{ v \in V_\alpha \mid E(v) \cap Z \neq \emptyset \} \cup \{ v \in V_{\bar{\alpha}} \mid E(v) \subseteq Z \}$$

Informally, we compute the  $\alpha$ -attractor of  $A$  by iteratively adding vertices to  $A$  of  $\alpha$  that have a successor in  $A$  and of  $\bar{\alpha}$  that have no successors outside  $A$ .

## 2.1 Solvers

We briefly introduce several approaches to solving parity games. These approaches can be roughly divided into two categories.

*First*, several algorithms iteratively perform local updates to vertices until a fixed point is reached. Each vertex is equipped with some measure which records the best game either player knows that they can play from that vertex so far. By updating measures based on the successors, they play the game backwards. The final measures indicate the winning player of each vertex and typically a winning strategy for one or both players. The *strategy improvement* algorithm (Sect. 4) and the *progress measures* algorithms (Sect. 5) fall into this category.

*Second*, several algorithms employ attractor computation to partition the game into regions that share a certain property. This partition is refined until the winners of some or all vertices can be identified, as well as the strategy for the winning player(s). The *recursive Zielonka* algorithm (Sect. 6) and the recently proposed *priority promotion* algorithms (Sect. 7) fall into this category.

## 2.2 Empirical Evaluation

Our goal in the empirical study is three-fold. *First*, we aim to compare modern algorithms and treat them fairly. We therefore need to establish that our implementation is competitive with existing work. *Second*, we compare the algorithms that are implemented in Oink directly. *Third*, as two algorithms have a parallel implementation, we also study the obtained parallel speedup and the parallel overhead when going from a sequential to a multi-core implementation.

We use the parity game benchmarks from model checking and equivalence checking proposed by Keiren [31] that are publicly available online. This is a total of 313 model checking and 216 equivalence checking games. We also consider different classes of random games, in part because the literature on parity games tends to favor studying the behavior of algorithms on random games. We include three classes of self-loop-free random games generated using PGSolver with a fixed number of vertices:

- low out-degree random games (`randomgame N N 1 2 x`)  
 $N \in \{ 100, 200, 500, 1000, 2000, 5000, 10000, 20000 \}$
- fully random games (`randomgame N N 1 N x`)  
 $N \in \{ 100, 500, 1000, 2000, 4000 \}$
- low-degree steady games (`steadygame N 1 4 1 4`)  
 $N \in \{ 100, 200, 500, 1000, 2000, 5000, 10000, 20000 \}$

We generate 20 games for each parameter  $N$ , in total 420 random games. We include low-degree games, since the solvers may behave differently on games where all vertices have few edges.

We present the evaluation in two ways. We compare runtimes of algorithms and penalize algorithms that do not finish on time (with a timeout of 15 min) by a factor  $2\times$  (PAR2), i.e., we assume that their runtime is  $2\times$  the timeout. This may still be quite optimistic. Compared to a timeout of 10 min, only few more games could be solved in 15 min. We also generate so-called *cactus plots* (often used to compare solvers in the SAT community) that show that a solver solved  $X$  models within  $Y$  seconds individually.

All experimental scripts and log files are available online via <http://www.github.com/trolando/oink-experiments>. The experiments were performed on a cluster of Dell PowerEdge M610 servers with two Xeon E5520 processors and 24 GB internal memory each. The tools were compiled with gcc 5.4.0.

### 3 Oink

We study modern parity game algorithms using our research tool named **Oink**. Oink is written in C++ and is publicly available under a permissive license via <https://www.github.com/trolando/oink>. Oink is easy to extend, as new solvers subclass the `Solver` class and only require a few extra lines in `solvers.cpp`.

Apart from implementing the full solvers described below, Oink also implements several preprocessors similar to other parity game solvers. We base our choices mainly on the practical considerations and observations by Friedmann and Lange [19] and by Verver [41]. We always reorder the vertices by priority and renumber the priorities from 0 to eliminate gaps (not the same as compression). The former is beneficial for the attractor-based algorithms in Sects. 6 and 7. The latter may reduce the amount of memory required for the measures-based algorithms in Sects. 4 and 5.

The following preprocessors are optional. Oink can perform priority inflation and priority compression, as described in [19]. We implement self-loop solving and winner-controlled winning cycle detection, as proposed in [41]. Winner-controlled winning cycle detection is a prerequisite for the strategy improvement algorithm of Sect. 4 but is optional for the other algorithms. We trivially solve games with only a single parity. Finally, we also implement SCC decomposition, which repeatedly solves a bottom SCC of the game until the full game is solved.

The correctness of an algorithm does not imply that implementations are correct. Although a formal proof of the implementations would be preferred, we also implement a fast solution verifier and verify all obtained solutions.

## 4 Strategy Improvement

Strategy improvement is a technique where each player iteratively improves their strategies until they are optimal. Strategy improvement algorithms were first explored for parity games by Jurdziński and Vöge [42] and have been subsequently improved in [7, 16, 20, 35, 38]. Recently, parallel implementations have been studied for the GPU [17, 24, 36]. Fearnley [17] also implements their parallel algorithm for multi-core CPUs. Our treatment in this section is mostly based on [17, 35].

In the strategy improvement algorithm, player Even has a strategy  $\sigma$  and player Odd has a strategy  $\tau$  for all their vertices. They improve their strategies until a fixed point is reached, at which point the game is solved. Instead of choosing a successor, player Even may also end the play. In the specific algorithm here, player Even delays selecting a strategy for a vertex until there is a favorable continuation. As  $\sigma$  and  $\tau$  cover all vertices, they induce a fixed path from each vertex. This path is either infinite (a play) or ends at a vertex of player Even. The strategy is evaluated by computing a valuation for each vertex based on the current paths. Strategies are improved by choosing the most favorable successor.

The valuation used in e.g. [17, 35] assigns to (infinite) plays the value  $\top$  and to (finite) paths a function  $L(p)$  that records for each priority  $p$  how often it occurs in the path. To determine the best move for each player, a total order  $\sqsubset$  compares valuations as follows. For non- $\top$  valuations  $L_1$  and  $L_2$ ,  $L_1 \sqsubset L_2$  iff there exists a highest priority  $z$  that is different in  $L_1$  and  $L_2$ , i.e.,  $z = \max \{ z \mid L_1(z) \neq L_2(z) \}$ , and either  $L_1(z) < L_2(z)$  if  $z$  is even, or  $L_1(z) > L_2(z)$  if  $z$  is odd. Furthermore,  $L \sqsubset \top$  for any  $L \neq \top$ . If  $L_1 \sqsubset L_2$ , then  $L_2$  is more favorable for player Even and  $L_1$  is more favorable for player Odd.

Intuitively, player Even likes plays where higher even priorities occur more often. Furthermore, player Even will end the play unless the highest priority in the continuation is even. Thus infinite paths are won by player Even and the valuation  $\top$  represents this. Player Even will always play to  $\top$  and player Odd will always try to avoid  $\top$ . This assumes that no winner-controlled winning cycles exist where player Odd wins, which can be guaranteed using a preprocessing step that removes these cycles.

For every strategy  $\sigma$  of player Even, a so-called *best response*  $\tau$  of player Odd minimizes the valuation of each position. Player Even always plays against this best response. In each iteration, after player Odd computes their best response, player Even computes all *switchable edges* based on the current valuation  $L$  and the current strategy  $\sigma$ . An edge  $(u, v)$  is *switchable* if  $L_v \sqsupset L_{\sigma(u)}$ . Not all switchable edges need to be selected for the improved strategy, but as argued in [17], the *greedy all-switches* rule that simply selects the best edge (maximal in  $\sqsubset$ ) for every position performs well in practice.

There are different methods to compute the best response of player Odd. We refer again to [17] for a more in-depth discussion. Player Odd can compute their best response by repeatedly switching all *Odd-switchable* edges, i.e., edges  $(u, v)$  s.t.  $L_v \sqsubset L_{\tau(u)}$  and  $L_v$  is minimal in  $\sqsubset$  of all successors of  $u$ .

```

1 def strategy-improvement:
2    $\sigma \leftarrow (V_0 \mapsto \perp)$ ,  $\tau \leftarrow$  random strategy for Odd
3   repeat
4     repeat
5       compute-valuations( $V$ ,  $\sigma \cup \tau$ ,  $L$ )
6        $\tau \leftarrow \tau[S_{\text{Odd}}]$  where  $S_{\text{Odd}} = \text{All}_{\text{Odd}}(\varnothing, \tau, L)$ 
7     until  $S_{\text{Odd}} = \emptyset$ 
8     mark-won( $\{v \in V : L_v = \top\}$ )
9      $\sigma \leftarrow \sigma[S_{\text{Even}}]$  where  $S_{\text{Even}} = \text{All}_{\text{Even}}(\varnothing, \sigma, L)$ 
10    until  $S_{\text{Even}} = \emptyset$ 
11    return ( $W_0, W_1, \sigma, \tau$ ) where  $W_0 \leftarrow \{v \in V : L_v = \top\}$ ,  $W_1 \leftarrow V \setminus W_0$ 

```

**Algorithm 1.** The strategy improvement algorithm.

```

1 def backwards-update( $v$ , into,  $L$ ):
2   for  $u \in \text{into}(v)$  :
3      $L_u \leftarrow L_v[\text{pr}(u) \mapsto L_v(\text{pr}(u)) + 1]$ 
4     spawn compute-valuation( $u$ ,  $\sigma$ ,  $L$ )
5   sync all
6 def compute-valuations( $V$ ,  $\sigma$ ,  $L$ ):
7   parallel for  $v \in V : L_v \leftarrow \top$  ; into( $v$ )  $\leftarrow \emptyset$ 
8   parallel for  $v \in V \mid \sigma(v) \neq \perp$  : add  $v$  to into( $\sigma(v)$ )
9   parallel for  $v \in V \mid \sigma(v) = \perp$  :
10     $L_v \leftarrow 0[\text{pr}(v) \mapsto 1]$ 
11    backwards-update( $v$ , into,  $L$ )

```

**Algorithm 2.** Computing valuations in parallel.

The players thus improve their strategies as in Algorithm 1, where  $\text{All}_{\text{Odd}}$  and  $\text{All}_{\text{Even}}$  compute all switchable edges as described above. The initial strategy for player Even is to always end the play. Player Odd computes their best response, starting from a random  $\tau$  initially. Player Even then improves their strategy once. They improve their strategies until a fixed point is reached. Then all vertices with valuation  $\top$  are won by player Even with strategy  $\sigma$  and all other vertices are won by player Odd with strategy  $\tau$  [17]. We extend the algorithm given in [17] at line 8 by marking vertices with valuation  $\top$  after player Odd computes their best response as “won”. We no longer need to consider them for  $\text{All}_{\text{Odd}}$  and  $\text{All}_{\text{Even}}$  as player Odd was unable to break the infinite play and thus they are won by Even.

The valuations can be computed in different ways. Fearnley implements a parallel algorithm that uses list ranking in two steps. The first step computes an Euler tour of the game restricted to chosen strategies  $\sigma$  and  $\tau$  resulting in a list. The second step uses a three-step parallel reduction algorithm to sum all values of the list. The list is divided into sublists which are each summed independently in the first sweep, all subresults are then propagated in the second sweep and then the final values are computed in the third sweep. See further [17].

We propose an alternative parallel algorithm to compute the valuation. We start from each Even vertex where the path ends and perform updates along a recursive backwards search, processing predecessors in parallel using task parallelism. Any vertex that is not visited has valuation  $\top$ . See Algorithm 2. This algorithm is implemented in Oink using the high-performance work-stealing framework Lace [12]. When updating the valuations in  $L$ , we first sweep twice over all vertices to initialize  $L$  for each vertex to  $\top$  and to add all vertices to  $\text{into}(v)$  that have their strategy to  $v$ . We also implement computing switchable edges in parallel via a straight-forward binary reduction using Lace.

**Table 1.** Runtimes in sec. (PAR2) and number of timeouts (15 min) of the three solvers PGSolver (**pgsi**), the solver by Fearnley [17] with sequential (**parsi-seq**) and multi-core variants, and Oink with sequential (**psi**) and multi-core variants.

	Model checking		Equiv checking		Random games		Total	
psi-8	<b>694</b>	0	<b>1078</b>	0	<b>315</b>	0	<b>2087</b>	0
psi	860	0	3262	0	480	0	4603	0
psi-1	1190	0	4090	0	487	0	5767	0
parsi-seq	1471	0	4199	0	1534	0	7204	0
parsi-8	2501	1	2908	0	56529	27	61938	28
parsi-1	4200	1	13867	6	71280	39	89347	46
pgsi	167596	88	95407	49	58839	27	321842	164

**Empirical Evaluation.** We compare the performance of Oink with the sequential and parallel solvers (1 or 8 threads) by Fearnley [17] and the “optstratimprov” solver in PGSolver. We disable optional preprocessing in all solvers. We only consider games without winner-controlled winning cycles, which are 289 model checking, 182 equivalence checking and 279 random games, in total 750 games.

See Table 1. We observe that PGSolver is vastly outperformed by Oink and the sequential solver of Fearnley. PGSolver timed out for 160 games, whereas **psi** and **parsi-seq** only timed out for 1 and 5 models, respectively. We observe similar parallel speedup for the parallel solvers, although Fearnley’s solver has more overhead from sequential to parallel with 1 thread. This might be due to the extra work to produce the Euler tour and to perform list ranking. The speedup we obtain with Oink is not very impressive, but the vast majority of the games are solved within 1s already. Furthermore, **psi** and **parsi-seq** are fairly close in performance. This is not a surprise, as their implementations are similar; the main difference is that Fearnley uses a forward search and we use a backward search. Hence, Oink is faster, but not by a large margin. Finally, we remark that Fearnley reports excellent results for list ranking on GPUs, whereas our algorithm is designed for a multi-core architecture.

## 5 Progress Measures

Progress measures is a technique that assigns to each vertex a monotonically increasing *measure*. The measure of each vertex is *lifted* based on the measures of its successors. By lifting vertices, players Even and Odd essentially play the game backwards. The measure represents a statistic of the most optimal play so far from the vertex, without storing the plays explicitly.

While progress measures have been used elsewhere, they were introduced for parity games by Jurdziński [27]. Several improvements to the original algorithm are due to Verver [41] and Gazda and Willemse [22]. A number of parallel implementations have been proposed for the Playstation 3 [6], for multi-core architectures [25, 37] and for GPUs [8, 24]. Furthermore, Chatterjee et al. proposed an implementation using BDDs [10]. Different types of progress measures were introduced after the recent breakthrough of a quasi-polynomial time algorithm due to Calude et al. [9], which resulted in the progress measures algorithms by Jurdziński et al. [28] and by Fearnley et al. [18]. This section studies small progress measures [27] and quasi-polynomial progress measures [18].

### 5.1 Small Progress Measures

The original small progress measures algorithm is due to Jurdziński [27]. We rely on the operational interpretation by Gazda and Willemse [22] and propose the *cap-and-carryover* mechanism to further understand the algorithm.

Progress measures record how favorable the game is for one of the players. W.l.o.g. we assume *even* progress measures. Given the highest priority  $d$ , define  $\mathbb{M}^\circ \subseteq \mathbb{N}^d \cup \{\top\}$  to be the largest set containing  $\top$  ( $\top \notin \mathbb{N}^d$ ) and only those  $d$ -tuples with 0 (denoted as  $\_$ ) on *odd* positions. An even progress measure  $m \in \mathbb{N}^d$  essentially records for a vertex  $v$  how often each even priority  $p$  is encountered along the most optimal play (starting at  $v$ ) so far, until a higher priority is encountered, i.e., until  $p$  no longer dominates. Such a prefix of the play is called a  $p$ -dominated stretch. Suppose that the sequence of priorities for a given play  $\pi$  is 00102120232142656201, then  $m = \{2\_3\_1\_2\}$ , since the play starts with a 0-dominated stretch containing two 0s, with a 2-dominated stretch containing three 2s, with a 4-dominated stretch containing one 4, and with a 6-dominated stretch containing two 6s. Furthermore, the special measure  $\top$  represents that the vertex can be won by player Even.

A total order  $\sqsubset$  compares measures as follows. For non- $\top$  measures  $m_1$  and  $m_2$ ,  $m_1 \sqsubset m_2$  iff there exists a highest priority  $z = \max\{z \mid m_1(z) \neq m_2(z)\}$  and  $m_1(z) < m_2(z)$ . Furthermore,  $m \sqsubset \top$  for all  $m \neq \top$ . We define a derived ordering  $\sqsubset_p$  by restricting  $z$  to priorities  $\geq p$ . Examples:

$$\begin{aligned} \{1.1.1\} &\sqsubset_0 \{0.0.2\} \\ \{3.2.1\} &\sqsubset_0 \{0.3.1\} \\ \{1.2.1\} &\sqsubset_1 \{0.2.1\} \\ \{3.3.1\} &\sqsubset_4 \{0.0.1\} \end{aligned}$$



To compute the progress measure for vertex  $v$  when playing to vertex  $w$ , given current measures  $\rho: V \rightarrow \mathbb{M}^\diamond$ , we define  $\text{Prog}(\rho, v, w)$  as follows:

$$\text{Prog}(\rho, v, w) := \begin{cases} \min\{ m \in \mathbb{M}^\diamond \mid m \sqsupset_{\text{pr}(v)} \rho(w) \} & \text{pr}(v) \text{ is even} \\ \min\{ m \in \mathbb{M}^\diamond \mid m \sqsubseteq_{\text{pr}(v)} \rho(w) \} & \text{pr}(v) \text{ is odd} \end{cases}$$

$\text{Prog}$  computes the measure of the play obtained by playing from  $v$  to the play recorded in  $\rho(w)$ . By choosing the lowest measure  $m$  according to  $\sqsupset_{\text{pr}(v)}$ , we ensure that all  $m(p)$  for  $p < \text{pr}(v)$  are set to 0. The inequality is strict for even priorities  $\text{pr}(v)$  to ensure that  $m(\text{pr}(v))$  increases.

Player Even wants to achieve the highest measure, whereas player Odd wants to achieve the lowest measure. We define  $\text{Lift}(\rho, v)$  as follows:

$$\text{Lift}(\rho, v) = \begin{cases} \rho[v \mapsto \max\{ \rho(v), \max\{ \text{Prog}(\rho, v, w) \mid v \rightarrow w \} \}] & \text{if } v \in V_\diamond \\ \rho[v \mapsto \max\{ \rho(v), \min\{ \text{Prog}(\rho, v, w) \mid v \rightarrow w \} \}] & \text{if } v \in V_\square \end{cases}$$

By definition, the  $\text{Lift}$  operation increases measures monotonically. For the specific algorithm described here, we also observe that  $\text{Prog}(\rho, v, w) \sqsupset_{\text{pr}(v)} \rho(w)$  and therefore  $\text{Lift}$  would even monotonically increase  $\rho$  without taking the maximum of the current measure and the best updated successor measure in a lifting procedure that starts with  $\rho = V \mapsto 0$ .

If we iteratively lift vertices from  $\rho = V \mapsto 0$  using  $\text{Lift}$ , eventually some vertex may have a measure  $m$  such that  $m(p)$  for some  $p$  is higher than the number of vertices with priority  $p$ , i.e.,  $m(p) > |V_p|$ . In this case, we know that  $m$  represents a play that visits at least one vertex with priority  $p$  twice and thus contains a cycle dominated by  $p$ . Furthermore, player Odd cannot escape from this cycle unless by playing to a higher losing priority. This follows from the fact that if player Odd could escape from the cycle, then it would not lift to this measure. The option to play to the higher losing priority is not considered because a measure to a higher priority is  $\sqsupset$  a measure that records a cycle.

We need a mechanism to let player Odd play to the next higher priority if it is forced into a cycle. However, we cannot let just any vertex play to a higher priority when its measure records a cycle, since some other vertex may escape to a lower higher priority. Therefore we need a mechanism that finds the lowest escape for player Odd. Small progress measures achieves this using a *cap-and-carryover* mechanism.  $\mathbb{M}^\diamond$  is restricted such that values for each even priority  $p$  may not be higher than  $|V_p|$ . When this cap is reached,  $\text{Prog}$  will naturally find a next higher  $m$  by increasing the value of higher priorities and eventually reach  $\top$ . For example, if we have two vertices of priority 2 and two vertices of priority 4 in a game and there is a self-loop of priority 2, measures increase as follows:  $\{0.2.0\}$ ,  $\{0.0.1\}$ ,  $\{0.1.1\}$ ,  $\{0.2.1\}$ ,  $\{0.0.2\}$ ,  $\{0.1.2\}$ ,  $\{0.2.2\}$ ,  $\top$ .

Thus all vertices involved in a cycle will find their measures slowly rising until the measure of some vertex controlled by Odd becomes equal to the measure when playing to a vertex that is not rising. This is the lowest escape. If no such escape is found, then the measures rise until  $\top$  and these vertices are won by player Even. The slowly increasing measures no longer follow the operational

interpretation described above, but can be understood as player Odd looking for the lowest escape.

We refer to [27] for the proof that the fixed point of applying the above lifting operation solves the parity game, such that vertices with measure  $\top$  are won by player Even and all other vertices are won by player Odd with a strategy that chooses the successor for which  $\text{Prog}(\rho, v, w)$  is the lowest.

We implement three known improvements. Improvements 2 and 3 are also implemented by PGSolver [19].

1. When a vertex with some even priority  $p$  is raised to  $\top$ , the cap of  $p$  may be lowered. The reason is that if a play records priority  $p$   $|V_p|$  times, it either contains a vertex now won by player Even or a cycle of priority  $p$  [41].
2. Small progress measures only computes the strategy for player Odd according to measures for player Even. We compute both *even* and *odd* measures simultaneously to compute the strategy for both players.
3. In addition, we occasionally halt the lifting procedure to perform an attractor computation for player Even to the set of *even*-liftable vertices. Any vertices not in this set are won by player Odd. We can immediately lift these vertices to  $\top$  in the *odd* measures. We perform this analysis also for *odd*-liftable measures to potentially lift vertices to  $\top$  in the *even* measures.

## 5.2 Quasi-polynomial Progress Measures

Different types of progress measures were introduced after the recent breakthrough of a quasi-polynomial time algorithm due to Calude et al. [9], which resulted in the progress measures algorithms by Jurdziński et al. [28] and by Fearnley et al. [18]. We only briefly and informally describe the idea of [18]. (Even) measures are  $k$ -tuples  $M: (\mathbb{N} \cup \{\perp\})^k \cup \{\top\}$ , which record that the optimal play consists of consecutive stretches that are dominated by vertices with even priority. For example, in the path  $\underline{1}2\underline{13}1\underline{4}23\underline{2}15\underline{6}3\underline{2}1\underline{2}$ , all vertices are dominated by each pair of underlined vertices of even priority.  $k$  is such that there are fewer than  $2^k$  vertices with even priority in the game. An 8-tuple  $\{2\ 2\ 4\ \perp\ 5\ \perp\ 6\ \perp\}$  denotes a game with consecutive stretches of 1, 2, 4, 16 and 64 even vertices, where the first dominating vertex has priority  $M(i)$  and may actually be *odd* instead of even. If the first dominating vertex has an odd priority, then player Even must reach a higher priority before continuing to build a play where they have more dominating even vertices than are in the game. If player Even can visit more dominating even vertices than are in the game, then at least one of these is visited twice and therefore player Even knows that they can win and lifts to  $\top$ .

## 5.3 Empirical Evaluation

We compare our implementation of small progress measures and quasi-polynomial progress measures to the small progress measures implementation of

**Table 2.** Runtimes in sec. (PAR2) and number of timeouts (15 min) of PGSolver (pgspm), pbesspsolve (pbesspm) and the implementations spm and qpt in Oink.

	Model checking		Equiv checking		Random games		Total	
spm	<b>3637</b>	1	<b>7035</b>	0	168271	93	<b>178944</b>	94
qpt	122549	64	65310	31	<b>66303</b>	35	254162	130
pbesspm	38397	20	52422	27	183742	101	274561	148
pgspm	88800	45	59885	30	320666	171	469351	246

pbesspsolve that comes with the mCRL2 model checker [11,41] and the implementation of small progress measures in PGSolver [19]. Unfortunately, the solver used in [18] contains proprietary source code and cannot be compiled and compared. For this comparison, we disabled optional preprocessing, i.e., removing self-loops, winner-controlled winning cycles and solving single-parity games.

See Table 2. Although Fearnley et al. [18] say that the QPT solver is mainly interesting for the theoretical result rather than practical performance, we observe that qpt outperforms the other solvers for random games. Oink is faster than PGSolver, especially for model checking and equivalence checking.

## 6 Zielonka’s Recursive Algorithm

The algorithm by Zielonka [43] is a recursive solver that despite its relatively bad theoretical complexity is known to outperform other algorithms in practice [19]. Furthermore, tight bounds are known for various classes of games [21].

Zielonka’s recursive algorithm is based on attractor computation. At each step, given current subgame  $\mathcal{D}$ , the algorithm removes the attractor  $A := \text{Attr}_\alpha^\mathcal{D}(\text{pr}^{-1}(\text{pr}(\mathcal{D})))$ , i.e., all vertices attracted to the current highest vertices of priority  $p := \text{pr}(\mathcal{D})$  for player  $\alpha = p \bmod 2$ , and recursively computes the winning regions  $(W_\diamond, W_\square)$  of the remaining subgame  $\mathcal{D} \setminus A$ . If the opponent  $\bar{\alpha}$  can attract vertices in  $A$  to  $W_{\bar{\alpha}}$ , then  $\bar{\alpha}$  wins  $W'_{\bar{\alpha}} := \text{Attr}_{\bar{\alpha}}^\mathcal{D}(W_{\bar{\alpha}})$  and the solution for the remainder  $\mathcal{D} \setminus W'_{\bar{\alpha}}$  is computed recursively. Otherwise,  $\alpha$  wins  $A$  and no further recursion is necessary. The strategies for both players are trivially obtained during attractor computation and by assigning to winning  $p$ -vertices in  $A$  any strategy to vertices in  $W_\alpha \cup A$ .

Zielonka’s original algorithm has been extended and improved over the years. In his thesis, Verver [41] improves the partitioning of the game after computing  $A$  by extending  $A$  with the attractors of the next highest vertices if they are of the same parity. The original algorithm always recomputes the solution of  $\mathcal{D} \setminus W'_{\bar{\alpha}}$  if  $W_{\bar{\alpha}}$  is nonempty, even if no vertices are attracted to  $W_{\bar{\alpha}}$ . Liu et al. propose that this is not necessary [34]. See Algorithm 3 for the recursive algorithm with these modifications. Other extensions that we do not consider here are the subexponential algorithm [29] and the big steps algorithm [39] that have been reported to perform slower than ordinary Zielonka [19]. Also, variations using BDDs have been proposed [2,30].

```

1 def zielonka( $\mathcal{D}$ ):
2   if  $\mathcal{D} = \emptyset$  : return  $\emptyset, \emptyset$ 
3    $\alpha \leftarrow \text{pr}(\mathcal{D}) \bmod 2$ 
4    $A \leftarrow \text{attr}(\mathcal{D}, \alpha)$ 
5    $W_{\diamond}, W_{\square} \leftarrow \text{zielonka}(\mathcal{D} \setminus A)$ 
6    $W'_{\bar{\alpha}} \leftarrow \text{Attr}_{\bar{\alpha}}^{\mathcal{D}}(W_{\bar{\alpha}})$ 
7   if  $W'_{\bar{\alpha}} = W_{\bar{\alpha}}$  :
8      $W_{\alpha} \leftarrow W_{\alpha} \cup A$ 
9   else:
10     $W_{\diamond}, W_{\square} \leftarrow \text{zielonka}(\mathcal{D} \setminus W'_{\bar{\alpha}})$ 
11     $W_{\bar{\alpha}} \leftarrow W_{\bar{\alpha}} \cup W'_{\bar{\alpha}}$ 
12   return  $W_{\diamond}, W_{\square}$ 

13 def attr( $\mathcal{D}, \alpha$ ):
14    $A \leftarrow \emptyset$ 
15   while  $\text{pr}(\mathcal{D} \setminus A) =_2 \alpha$  :  $A \leftarrow A \cup \text{Attr}_{\alpha}^{\mathcal{D} \setminus A}(\text{pr}^{-1}(\text{pr}(\mathcal{D} \setminus A)))$ 
16   return  $A$ 

```

**Algorithm 3.** The recursive Zielonka algorithm.

Although the implementation of the recursive algorithm in PGsolver [19] is typically used for comparisons in the literature, improved implementations have been proposed by Verver [41], Di Stasio et al. [40], Liu et al. [34], and Arcucci et al. [1]. Verver suggests to record the number of remaining “escaping” edges for each vertex during attractor computation, to reduce the complexity of attractor computation at the cost of an extra integer per vertex. Di Stasio et al. avoid creating copies of the game for recursive operations by recording which vertices are removed in a special array. Recently, Arcucci et al. extended the implementation in [40] with a multi-core implementation of attractor computation [1].

The implementation in Oink is based upon the ideas described above. Furthermore, we improve the implementation using the following techniques.

- Instead of creating copies of the “removed” array [40] for each recursive step, we use a single “region” array that stores for each vertex that it is attracted by the  $r$ th call to `attr`. This value is initially  $\perp$  for all vertices and is reset to  $\perp$  for vertices in  $\mathcal{D} \setminus W'_{\bar{\alpha}}$  (line 10). We record the initial  $r$  at each depth and thus derive that all vertices with a value  $\geq r$  or  $\perp$  are part of the subgame.
- As a preprocessing step, we order all vertices by priority. We can then quickly obtain the highest vertex of each subgame.
- We eliminate the recursion using a stack.
- We implement an alternative lock-free attractor, relying on the work-stealing library Lace [12] that provides fine-grained load balancing.

In the interest of space, we cannot describe the multi-core attractor in-depth. This implementation is fairly straightforward. We implement the attractor recursively where the work-stealing framework runs the recursive operations in parallel. Like typical lock-free algorithms, we rely on the `compare-and-swap` operation

**Table 3.** Runtimes in sec. (PAR2) and number of timeouts (15 min) of the four solvers PGSolver (`pgzlk`), SPGSolver (`spg`), pbespgsolve (`pbeszlk`) and Oink (sequential `zlk`, multi-core `zlk-1` and `zlk-8`, unoptimized `uzlk`).

	Model checking		Equiv checking		Random games		Total	
<code>zlk-8</code>	94	0	415	0	11	0	<b>521</b>	0
<code>zlk</code>	88	0	472	0	<b>6</b>	0	566	0
<code>zlk-1</code>	97	0	512	0	7	0	616	0
<code>uzlk</code>	89	0	472	0	69	0	630	0
<code>pbeszlk</code>	64	0	513	0	338	0	915	0
<code>spg-seq</code>	<b>58</b>	0	<b>198</b>	0	694	0	950	0
<code>spg-mc</code>	389	0	1451	0	72608	37	74447	37
<code>pgzlk</code>	65905	33	68013	36	41629	14	175547	83

to implement safe communication between threads. The attractor uses this operation when manipulating the number of escaping edges and to “claim” a vertex by setting its value in the region array from  $\perp$  to  $r$ .

**Empirical Evaluation.** We compare our implementation of Zielonka’s recursive algorithm with and without the optimizations of Algorithm 3 to PGSolver, to Verver’s implementation pbespgsolve [11, 41] and to SPGSolver [1, 40]. Unfortunately, the Java version of SPGSolver (all three variations) suffers from severe performance degradation for unknown reasons. They also provide a C++ implementation in their online repository, which we used instead. The multi-core version of the SPGSolver tool relies on `async` tasks provided by C++11. Similar to the previous sections, we disable the optional preprocessors that solve single parity games, remove self-loops and solve winner-controlled winning cycles.

See Table 3. The results show that the implementation in Oink outperforms PGSolver by several orders of magnitude on all benchmark types. PGSolver timed out for 83 of all 949 games. The solvers `spg-seq` and `pbeszlk` are faster than Oink on the model checking and equivalence checking games, but are significantly outperformed on random games. We also observe severe performance degradation for `spg-mc` on random games. It appears that our parallel implementation of Zielonka’s algorithm also does not scale well. Finally, there seems to be no significant difference between the optimized and unoptimized versions of Zielonka’s algorithm, except for random games.

## 7 Priority Promotion

In recent work, a new family of algorithms has been proposed based on *priority promotion* [5]. Priority promotion starts with a similar decomposition of the game as Zielonka’s recursive algorithm. Priority promotion is based on the insight that a recursive decomposition based on attractor computation leads to regions

**Table 4.** Runtimes in sec. (PAR2) and number of timeouts (15 min) of the five priority promotion solvers in Oink.

	Model checking		Equiv checking		Random games		Total	
ppp	<b>81</b>	0	<b>382</b>	0	<b>12</b>	0	<b>475</b>	0
pp	82	0	<b>382</b>	0	<b>12</b>	0	476	0
rr	<b>81</b>	0	385	0	<b>12</b>	0	477	0
dp	84	0	389	0	15	0	488	0
rrdp	83	0	394	0	14	0	491	0

with a specific property related to the highest priority in the region, called its *measure*  $p$ . This property is that all plays that stay in the region are won by the player who wins the highest priority  $p$ , denoted by player  $\alpha$ . The other player  $\bar{\alpha}$  has three options. They either lose the game by staying in the region, or they can leave the region by playing to an  $\alpha$ -region of higher measure, or they can leave the region to a lower region of either player via a vertex with priority  $p$ . The goal of  $\alpha$  is to find “closed”  $\alpha$ -regions, where  $\bar{\alpha}$  cannot escape to lower regions. The result is a region where player  $\bar{\alpha}$  either loses, or leaves the region to a higher  $\alpha$ -region which may or may not be closed. The measure of the closed  $\alpha$ -region is then “promoted” to the measure of the lowest higher region to which  $\bar{\alpha}$  can escape and the attractor-based decomposition is recomputed for all lower regions. The promoted region may now attract from regions with a measure between its original measure and its promoted measure, thus requiring recomputing the decomposition. When player  $\bar{\alpha}$  cannot escape from an  $\alpha$ -region to a higher  $\alpha$ -region, player  $\alpha$  is the winner of all vertices in the region.

Priority promotion was proposed in [5] and improved in [3, 4]. The original PP algorithm [5] forgets all progress (“resets”) in lower regions after promotion. The PP+ algorithm [3] only resets lower regions of player  $\bar{\alpha}$ . The RR algorithm [4] only resets *some* lower regions of player  $\bar{\alpha}$ . The DP algorithm [3] uses a heuristic to delay certain promotions to avoid resets. We implement all four algorithms and also combine the DP algorithm, which is based on PP+, with the RR algorithm.

**Empirical Evaluation.** We compare our implementation of five variations of priority promotion in Oink. As we do not compare with other solvers, we enable the optional preprocessors that solve single parity games, remove self-loops and solve winner-controlled winning cycles.

See Table 4. Overall, we see that the simplest solver pp performs just as good as the more complex solvers. The motivation for the variations is based on crafted families that require an exponential number of promotions. The pp solver may be most vulnerable to these crafted families, but on practical games and random games there is no significant difference.

## 8 Conclusions


See Table 5 for a comparison of the five main sequential algorithms in Oink, including the preprocessing that removes winner-controlled winning cycles, self-loops and solves single parity games. The results show that the `zlk` and the `pp` solvers have similar performance and outperform the other solvers. See also Fig. 1 for a cactus plot of these five solvers.

Priority promotion is a powerful and attractive idea, as promoting closed  $\alpha$ -regions is similar to *cap-and-carryover* in small progress measures. Attractor computation finds such regions directly whereas value iteration algorithms may require many iterations. We confirm the observations in [5] that the algorithm has a good performance but it is not faster than Zielonka’s algorithm.

In this work, we studied modern parity game algorithms using a new tool named **Oink**. Oink is publicly available via <https://www.github.com/trolando/oink>. We implemented a number of modern algorithms and provided a comprehensive description of these algorithms, introducing *cap-and-carryover* to understand small progress measures. We proposed improvements to strategy improvement and to Zielonka’s algorithm. We presented an empirical evaluation of Oink, comparing its performance with state-of-the-art solvers, especially the popular PGSolver tool. The results demonstrate that Oink is competitive with other

**Table 5.** Runtimes in sec. (PAR2) and number of timeouts (15 min) of the sequential implementations of the five solvers in Oink described in this paper.

	Model checking		Equiv checking		Random games		Total	
pp	82	0	<b>382</b>	0	12	0	<b>476</b>	0
zlk	<b>78</b>	0	393	0	<b>10</b>	0	481	0
psi	231	0	2440	0	689	0	3359	0
spm	1007	0	3079	0	156885	87	160971	87
qpt	59559	31	60728	31	62104	33	182391	95



**Fig. 1.** A cactus plot of five sequential solvers implemented in Oink. The plot shows how many games are (individually) solved within a certain amount of time.

implementations and in fact outperforms PGSolver for all algorithms, especially Zielonka's recursive algorithm. This result is particularly interesting considering that many publications compare the performance of novel ideas to Zielonka's algorithm in PGSolver.

**Acknowledgements.** We thank Tim Willemse and John Fearnley for their helpful comments and Jaco van de Pol for the use of their computer cluster.

## References

1. Arcucci, R., Marotta, U., Murano, A., Sorrentino, L.: Parallel parity games: a multicore attractor for the Zielonka recursive algorithm. In: ICCS, *Procedia Computer Science*, vol. 108, pp. 525–534. Elsevier (2017)
2. Bakera, M., Edelkamp, S., Kissmann, P., Renner, C.D.: Solving  $\mu$ -calculus parity games by symbolic planning. In: Peled, D.A., Wooldridge, M.J. (eds.) *MoChArt 2008. LNCS (LNAI)*, vol. 5348, pp. 15–33. Springer, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-00431-5\\_2](https://doi.org/10.1007/978-3-642-00431-5_2)
3. Benerecetti, M., Dell'Erba, D., Mogavero, F.: A delayed promotion policy for parity games. In: GandALF 2016. *EPTCS*, vol. 226, pp. 30–45 (2016)
4. Benerecetti, M., Dell'Erba, D., Mogavero, F.: Improving priority promotion for parity games. In: Bloem, R., Arbel, E. (eds.) *HVC 2016. LNCS*, vol. 10028, pp. 117–133. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-49052-6\\_8](https://doi.org/10.1007/978-3-319-49052-6_8)
5. Benerecetti, M., Dell'Erba, D., Mogavero, F.: Solving parity games via priority promotion. In: Chaudhuri, S., Farzan, A. (eds.) *CAV 2016. LNCS*, vol. 9780, pp. 270–290. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-41540-6\\_15](https://doi.org/10.1007/978-3-319-41540-6_15)
6. van der Berg, F.: Solving parity games on the playstation 3. In: *Twente Student Conference* (2010)
7. Björklund, H., Vorobyov, S.G.: A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discret. Appl. Math.* **155**(2), 210–229 (2007)
8. Bootsma, P.: Speeding up the small progress measures algorithm for parity games using the GPU. Master's thesis, Eindhoven University of Technology (2013)
9. Calude, C.S., Jain, S., Khoussainov, B., Li, W., Stephan, F.: Deciding parity games in quasipolynomial time. In: *STOC*, pp. 252–263. ACM (2017)
10. Chatterjee, K., Dvořák, W., Henzinger, M., Loitzenbauer, V.: Improved set-based symbolic algorithms for parity games. In: *CSL, LIPIcs*, vol. 82, pp. 18:1–18:21. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2017)
11. Cranen, S., Groote, J.F., Keiren, J.J.A., Stappers, F.P.M., de Vink, E.P., Wesselink, W., Willemse, T.A.C.: An overview of the mCRL2 toolset and its recent advances. In: Piterman, N., Smolka, S.A. (eds.) *TACAS 2013. LNCS*, vol. 7795, pp. 199–213. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-36742-7\\_15](https://doi.org/10.1007/978-3-642-36742-7_15)
12. van Dijk, T., van de Pol, J.C.: Lace: non-blocking split deque for work-stealing. In: Lopes, L., et al. (eds.) *Euro-Par 2014. LNCS*, vol. 8806, pp. 206–217. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-14313-2\\_18](https://doi.org/10.1007/978-3-319-14313-2_18)
13. Emerson, E.A., Jutla, C.S.: Tree automata,  $\mu$ -calculus and determinacy (extended abstract). In: *FOCS*, pp. 368–377. IEEE Computer Society (1991)
14. Emerson, E.A., Jutla, C.S., Sistla, A.P.: On model-checking for fragments of  $\mu$ -calculus. In: Courcoubetis, C. (ed.) *CAV 1993. LNCS*, vol. 697, pp. 385–396. Springer, Heidelberg (1993). [https://doi.org/10.1007/3-540-56922-7\\_32](https://doi.org/10.1007/3-540-56922-7_32)



15. Emerson, E.A., Jutla, C.S., Sistla, A.P.: On model checking for the  $\mu$ -calculus and its fragments. *Theor. Comput. Sci.* **258**(1–2), 491–522 (2001)
16. Fearnley, J.: Non-oblivious strategy improvement. In: Clarke, E.M., Voronkov, A. (eds.) LPAR 2010. LNCS (LNAI), vol. 6355, pp. 212–230. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-17511-4\\_13](https://doi.org/10.1007/978-3-642-17511-4_13)
17. Fearnley, J.: Efficient parallel strategy improvement for parity games. In: Majumdar, R., Kunčák, V. (eds.) CAV 2017. LNCS, vol. 10427, pp. 137–154. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-63390-9\\_8](https://doi.org/10.1007/978-3-319-63390-9_8)
18. Fearnley, J., Jain, S., Schewe, S., Stephan, F., Wojtczak, D.: An ordered approach to solving parity games in quasi polynomial time and quasi linear space. In: SPIN, pp. 112–121. ACM (2017)
19. Friedmann, O., Lange, M.: Solving parity games in practice. In: Liu, Z., Ravn, A.P. (eds.) ATVA 2009. LNCS, vol. 5799, pp. 182–196. Springer, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-04761-9\\_15](https://doi.org/10.1007/978-3-642-04761-9_15)
20. Friedmann, O., Lange, M.: Local strategy improvement for parity game solving. In: GandALF. EPTCS, vol. 25, pp. 118–131 (2010)
21. Gazda, M., Willemse, T.A.C.: Zielonka’s recursive algorithm: dull, weak and solitaire games and tighter bounds. In: GandALF. EPTCS, vol. 119, pp. 7–20 (2013)
22. Gazda, M., Willemse, T.A.C.: Improvement in small progress measures. In: GandALF. EPTCS, vol. 193, pp. 158–171 (2015)
23. Grädel, E., Thomas, W., Wilke, T. (eds.): Automata Logics, and Infinite Games: A Guide to Current Research. LNCS, vol. 2500. Springer, Heidelberg (2002). <https://doi.org/10.1007/3-540-36387-4>
24. Hoffmann, P., Luttenberger, M.: Solving parity games on the GPU. In: Van Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 455–459. Springer, Cham (2013). [https://doi.org/10.1007/978-3-319-02444-8\\_34](https://doi.org/10.1007/978-3-319-02444-8_34)
25. Huth, M., Kuo, J.H.-P., Piterman, N.: Concurrent small progress measures. In: Eder, K., Lourenço, J., Shehory, O. (eds.) HVC 2011. LNCS, vol. 7261, pp. 130–144. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-34188-5\\_13](https://doi.org/10.1007/978-3-642-34188-5_13)
26. Jurdzinski, M.: Deciding the winner in parity games is in  $UP \cap co-UP$ . *Inf. Process. Lett.* **68**(3), 119–124 (1998)
27. Jurdziński, M.: Small progress measures for solving parity games. In: Reichel, H., Tison, S. (eds.) STACS 2000. LNCS, vol. 1770, pp. 290–301. Springer, Heidelberg (2000). [https://doi.org/10.1007/3-540-46541-3\\_24](https://doi.org/10.1007/3-540-46541-3_24)
28. Jurdzinski, M., Lazic, R.: Succinct progress measures for solving parity games. In: LICS, pp. 1–9. IEEE Computer Society (2017)
29. Jurdzinski, M., Paterson, M., Zwick, U.: A deterministic subexponential algorithm for solving parity games. *SIAM J. Comput.* **38**(4), 1519–1532 (2008)
30. Kant, G., van de Pol, J.: Generating and solving symbolic parity games. In: GRAPHITE. EPTCS, vol. 159, pp. 2–14 (2014)
31. Keiren, J.J.A.: Benchmarks for parity games. In: Dastani, M., Sirjani, M. (eds.) FSEN 2015. LNCS, vol. 9392, pp. 127–142. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-24644-4\\_9](https://doi.org/10.1007/978-3-319-24644-4_9)
32. Kozen, D.: Results on the propositional mu-calculus. *Theor. Comput. Sci.* **27**, 333–354 (1983)
33. Kupferman, O., Vardi, M.Y.: Weak alternating automata and tree automata emptiness. In: STOC, pp. 224–233. ACM (1998)
34. Liu, Y., Duan, Z., Tian, C.: An improved recursive algorithm for parity games. In: TASE, pp. 154–161. IEEE Computer Society (2014)
35. Luttenberger, M.: Strategy iteration using non-deterministic strategies for solving parity games. *CoRR abs/0806.2923* (2008)

36. Meyer, P.J., Luttenberger, M.: Solving mean-payoff games on the GPU. In: Artho, C., Legay, A., Peled, D. (eds.) ATVA 2016. LNCS, vol. 9938, pp. 262–267. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46520-3\\_17](https://doi.org/10.1007/978-3-319-46520-3_17)
37. van de Pol, J., Weber, M.: A multi-core solver for parity games. *Electr. Notes Theor. Comput. Sci.* **220**(2), 19–34 (2008)
38. Schewe, S.: An optimal strategy improvement algorithm for solving parity and payoff games. In: Kaminski, M., Martini, S. (eds.) CSL 2008. LNCS, vol. 5213, pp. 369–384. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-87531-4\\_27](https://doi.org/10.1007/978-3-540-87531-4_27)
39. Schewe, S.: Solving parity games in big steps. *J. Comput. Syst. Sci.* **84**, 243–262 (2017)
40. Di Stasio, A., Murano, A., Prignano, V., Sorrentino, L.: Solving parity games in scala. In: Lanese, I., Madelaine, E. (eds.) FACS 2014. LNCS, vol. 8997, pp. 145–161. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-15317-9\\_9](https://doi.org/10.1007/978-3-319-15317-9_9)
41. Verver, M.: Practical improvements to parity game solving. Master’s thesis, University of Twente (2013)
42. Vöge, J., Jurdziński, M.: A discrete strategy improvement algorithm for solving parity games. In: Emerson, E.A., Sistla, A.P. (eds.) CAV 2000. LNCS, vol. 1855, pp. 202–215. Springer, Heidelberg (2000). [https://doi.org/10.1007/10722167\\_18](https://doi.org/10.1007/10722167_18)
43. Zielonka, W.: Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theor. Comput. Sci.* **200**(1–2), 135–183 (1998)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

