Solving Parity Games via Priority Promotion

Massimo Benerecetti · Daniele Dell'Erba · Fabio Mogavero

Abstract We consider *parity games*, a special form of two-player infinite-duration games on numerically labeled graphs, whose winning condition requires that the maximal value of a label occurring infinitely often during a play be of some specific parity. The problem of identifying the corresponding winning regions has a rather intriguing status from a complexity theoretic viewpoint, since it belongs to the class UPTIME ∩ COUPTIME, and still open is the question whether it can be solved in polynomial time. Parity games also have great practical interest, as they arise in many fields of theoretical computer science, most notably logic, automata theory, and formal verification. In this paper, we propose a new algorithm for the solution of this decision problem, based on the idea of promoting vertexes to higher priorities during the search for winning regions. The proposed approach has nice computational properties, exhibiting the best space complexity among the currently known solutions. Experimental results on both random games and benchmark families show that the technique is also very effective in practice.

1 Introduction

Parity games [49] are perfect-information two-player turn-based games of infinite duration, usually played on finite directed graphs. Their vertices, labeled by natural numbers called *priorities*, are called positions and assigned to one of two players, named *Even* and *Odd* or, simply, 0 and 1, respectively. The game starts at an arbitrary position and, during its evolution, each player can take a move only at its own positions, which consists in choosing one of the edges outgoing from the current position. The moves selected by the players induce an infinite sequence of positions, called play. If the maximal priority of the positions occurring infinitely often in the play is *even*, then the play is winning for player 0, otherwise, player 1 takes it all.

Parity games have been extensively studied in the attempt to find efficient solutions

---

M. Benerecetti and D. Dell'Erba
Università degli Studi di Napoli Federico I, Napoli, Italy

F. Mogavero
University of Oxford, Oxford, UK

to the problem of determining the winner. From a complexity theoretic perspective, this decision problem lies in NPTIME ∩ CoNPTIME [20,21], since these games are *memoryless determined* [49,19,41,42]. It has been even proved to belong to UPTIME ∩ COUPTIME [34], a status shared with the *factorization problem* [27,28,1]. They are the simplest class of games in a wider family with similar complexities and containing, *e.g.*, *mean payoff games* [18,33], *discounted payoff games* [59], and *simple stochastic games* [17]. In fact, polynomial time reductions exist from parity games to the latter ones. However, despite being the most likely class among those games to admit a polynomial-time solution, the answer to the question whether such a solution exists still eludes the research community.

The effort devoted to provide efficient solutions stems primarily form the fact that many problems in formal verification and synthesis can be reformulated in terms of solving parity games. Emerson, Jutla, and Sistla [20,21] have shown that computing winning strategies for these games is linear-time equivalent to solving the modal μCALCULUS model checking problem [22]. Parity games also play a crucial role in automata theory [48,19,40], where, for instance, they can be applied to solve the complementation problem for alternating automata [32] and the emptiness of the corresponding nondeterministic tree automata [40]. These automata, in turn, can be used to solve the satisfiability and model checking problems for expressive logics, such as the modal [57] and alternating [55,2] μCALCULUS, ATL* [2, 54], Strategy Logic [16,47,44,45], Substructure Temporal Logic [5,6], and fixed-point extensions of guarded first-order logics [8,9].

Previous solutions mainly divide into two families: those that solve a game by first decomposing it into smaller subgames, and those that proceed in a global fashion and approach the game in its entirety. To the first family belongs the *divide et impera* solution originally proposed by McNaughton [43] for Muller games and adapted to parity games by Zielonka [58]. More recent improvements to that recursive algorithm have been proposed by Jurdziński, Paterson, and Zwick [37,38] and by Schewe [52]. Both approaches rely on finding suitably closed dominions, which can then be removed from a game to reduce the size of the subgames to be recursively solved. To the second family belongs the procedure proposed by Jurdziński [35], which exploits the connection between the notions of progress measures [39] and winning strategies. In this approach, an initial measure function on the entire game is iteratively updated until a fixpoint is reached. At that point, a progress measure is obtained that induces a winning strategy for one of the two players. An alternative approach was proposed by Jurdziński and Vöge [56], which directly builds a winning strategy for one of the two players, by iteratively improving an initial non-winning strategy. This technique was later optimized by Schewe [53]. A recent breakthrough [11] by Calude *et al.* proposes a succinct reduction from parity to reachability games based on a clever encoding of the sequences of priorities a player finds along a play. This allows for a mere quasi-polynomial blow up in the size of the underlying graph and sets the basis of the fixed-parameter tractability *w.r.t.* the number of priorities. The approach has been then considerably refined in [24], where these encodings are modeled as progress measures. A similar technique is also used in [36]. Despite the theoretical relevance of this new idea, preliminary experiments conducted in this paper suggest that the practical impact of the result may not match the theoretical one. Indeed, most of the exponential algorithms mentioned above outperform, often by orders of magnitude, the current implementations of the quasi-polynomial ones, which do not scale beyond a few hundred positions. This evaluation is consistent with the fact that the new techniques essentially amount to clever and succinct encodings embedded within a brute force search, which makes matching quasi-polynomial worst cases quite easy to find. As far as space consumption is concerned, we have different and, in some cases, incomparable behaviors. The small progress measure procedure of [35] requires $O(k \cdot n \cdot \log n)$ space,

with n the number of positions in the game and k the number of its priorities. On the other hand, it is $O(n^2)$ for the optimized strategy improvement method of [53]. Due to their inherent recursive nature, the algorithms of the first family require $O(m \cdot n)$ memory, where m denotes the number of edges of the underlying graph. This bound could, in principle, be reduced to $O(n^2)$, by representing subgames implicitly through their sets of positions. The lowest space requirements, however, are those of the more recent quasi-polynomial algorithm described in [36], which only requires $O(n \cdot \log n \cdot \log k)$ additional space. All these bounds do not seem to be amenable to further improvements, as they appear to be intrinsic to the corresponding solution techniques. Polynomial time solutions are only known for restricted versions of the problem, where one among tree-width [50,25,26], dag-width [7], clique-width [51] and entanglement [10] of the underlying graph is bounded.

The main contribution of the paper is a new algorithm for solving parity games, based on the notions of *quasi dominion* and *priority promotion*. A quasi dominion Q for player $\alpha \in \{0, 1\}$, called a *quasi α-dominion*, is a set of positions from each of which player α can enforce a winning play that never leaves the region, unless one of the following two conditions holds: *(i)* the opponent $\overline{\alpha}$ can escape from Q or *(ii)* the only choice for player α itself is to exit from Q (*i.e.*, no move from a position of α remains in Q). Quasi dominions can be ordered by assigning to each of them a priority corresponding to an under-approximation of the best value the opponent can be forced to visit along any play exiting from it. A crucial property is that, under suitable and easy to check assumptions, a higher priority quasi α-dominion Q_1 and a lower priority one Q_2, can be merged into a single quasi α-dominion of the higher priority, thus improving the approximation for Q_2. For this reason we call this merging operation a priority promotion of Q_2 to Q_1. The underlying idea of our approach is to iteratively enlarge quasi α-dominions, by performing sequences of promotions, until an α-dominion is obtained.

We prove soundness and completeness of the algorithm. Moreover, we provide an upper bound $O\left(kn \left(\frac{\mathsf{e}n}{k-1} \right)^{k-1} \right)$ on the time complexity, where e is Euler's number, and a bound $O(n \cdot \log k)$ on the memory requirements. Experimental results, comparing our algorithm with the state of the art solvers, also show that the proposed approach performs very well in practice, most often significantly better than existing ones, on both random games and benchmark families proposed in the literature.

2 Preliminaries

Let us briefly recall the notation and basic definitions concerning parity games that an expert reader can simply skip. We refer to [3,58] for a comprehensive presentation of the subject.

Given a partial function $\mathsf{f} : A \rightharpoonup B$, by $\mathsf{dom}(\mathsf{f}) \subseteq A$ and $\mathsf{rng}(\mathsf{f}) \subseteq B$ we indicate the domain and range of f, respectively. In addition, $\uplus$ denotes the *completion operator* that, taken f and another partial function $\mathsf{g} : A \rightharpoonup B$, returns the partial function $\mathsf{f} \uplus \mathsf{g} \triangleq (\mathsf{f} \setminus \mathsf{dom}(\mathsf{g})) \cup \mathsf{g} : A \rightharpoonup B$, which is equal to g on its domain and assumes the same values of f on the remaining part of A.

A two-player turn-based *arena* is a tuple $\mathcal{A} = \langle \mathrm{Ps}^0, \mathrm{Ps}^1, Mv \rangle$, with $\mathrm{Ps}^0 \cap \mathrm{Ps}^1 = \emptyset$ and $\mathrm{Ps} \triangleq \mathrm{Ps}^0 \cup \mathrm{Ps}^1$, such that $\langle \mathrm{Ps}, Mv \rangle$ is a finite directed graph without sinks. Ps^0 (*resp.*, Ps^1) is the set of positions of player 0 (*resp.*, 1) and $Mv \subseteq \mathrm{Ps} \times \mathrm{Ps}$ is a left-total relation describing all possible moves. A *path* in $\mathrm{V} \subseteq \mathrm{Ps}$ is a finite or infinite sequence $\pi \in \mathrm{Pth}(\mathrm{V})$ of positions in V compatible with the move relation, *i.e.*, $(\pi_i, \pi_{i+1}) \in Mv$, for all $i \in [0, |\pi| - 1[$. For a finite path π, with $\mathsf{lst}(\pi)$ we denote the last position of π. A

positional *strategy* for player $\alpha \in \{0, 1\}$ on $V \subseteq Ps$ is a partial function $\sigma_\alpha \in Str^\alpha(V) \subseteq (V \cap Ps^\alpha) \rightharpoonup V$, mapping each α-position $v \in dom(\sigma_\alpha)$ to position $\sigma_\alpha(v)$ compatible with the move relation, *i.e.*, $(v, \sigma_\alpha(v)) \in Mv$. With $Str^\alpha(V)$ we denote the set of all α-strategies on V. A *play* in $V \subseteq Ps$ from a position $v \in V$ *w.r.t.* a pair of strategies $(\sigma_0, \sigma_1) \in Str^0(V) \times Str^1(V)$, called $((\sigma_0, \sigma_1), v)$-*play*, is a path $\pi \in Pth(V)$ such that $\pi_0 = v$ and, for all $i \in [0, |\pi| - 1[$, if $\pi_i \in Ps^0$ then $\pi_{i+1} = \sigma^0(\pi_i)$ else $\pi_{i+1} = \sigma^1(\pi_i)$. The *play function* $play : (Str^0(V) \times Str^1(V)) \times V \rightarrow Pth(V)$ returns, for each position $v \in V$ and pair of strategies $(\sigma_0, \sigma_1) \in Str^0(V) \times Str^1(V)$, the maximal $((\sigma_0, \sigma_1), v)$-play $play((\sigma^0, \sigma^1), v)$.

A *parity game* is a tuple $\partial = \langle \mathcal{A}, Pr, pr \rangle \in \mathcal{P}$, where $\mathcal{A}$ is an arena, $Pr \subset \mathbb{N}$ is a finite set of priorities, and $pr : Ps \rightarrow Pr$ is a *priority function* assigning a priority to each position. The priority function can be naturally extended to games and paths as follows: $pr(\partial) \triangleq \max_{v \in Ps} pr(v)$; for a path $\pi \in Pth$, we set $pr(\pi) \triangleq \max_{i \in [0, |\pi|[} pr(\pi_i)$, if π is finite, and $pr(\pi) \triangleq \limsup_{i \in \mathbb{N}} pr(\pi_i)$, otherwise. A set of positions $V \subseteq Ps$ is an α-*dominion*, with $\alpha \in \{0, 1\}$, if there exists an α-strategy $\sigma_\alpha \in Str^\alpha(V)$ such that, for all $\overline{\alpha}$-strategies $\sigma_{\overline{\alpha}} \in Str^{\overline{\alpha}}(V)$ and positions $v \in V$, the induced play $\pi = play((\sigma_0, \sigma_1), v)$ is infinite and $pr(\pi) \equiv_2 \alpha$. In other words, σ_α only induces on V infinite plays whose maximal priority visited infinitely often has parity α. By $\partial \setminus V$ we denote the maximal subgame of ∂ with set of positions Ps' contained in $Ps \setminus V$ and move relation Mv' equal to the restriction of Mv to Ps'.

The α-predecessor of V, in symbols $pre^\alpha(V) \triangleq \{v \in Ps^\alpha : Mv(v) \cap V \neq \emptyset\} \cup \{v \in Ps^{\overline{\alpha}} : Mv(v) \subseteq V\}$, collects the positions from which player α can force the game to reach some position in V with a single move. The α-attractor $atr^\alpha(V)$ generalities the notion of α-predecessor $pre^\alpha(V)$ to an arbitrary number of moves, and corresponds to the least fix-point of that operator. When $V = atr^\alpha(V)$, we say that V is α-maximal. Intuitively, V is α-maximal if player α cannot force any position outside V to enter the set. For such a V, the set of positions of the subgame $\partial \setminus V$ is precisely $Ps \setminus V$. Finally, the set $esc^\alpha(V) \triangleq pre^\alpha(Ps \setminus V) \cap V$, called the α-*escape* of V, contains the positions in V from which α can leave V in one move. The dual notion of α-*interior*, defined as $int^\alpha(V) \triangleq (V \cap Ps^\alpha) \setminus esc^\alpha(V)$, contains, instead, the α-positions from which α cannot escape with a single move. Observe that all the operators and sets described above actually depend on the specific game ∂ they are applied in. In the rest of the paper, we shall only add ∂ as subscript of an operator, *e.g.*, $esc^\alpha_\partial(V)$, when the game is not clear from the context.

3 A New Idea

A solution for a parity game $\partial = \langle \mathcal{A}, Pr, pr \rangle \in \mathcal{P}$ over an arena $\mathcal{A} = \langle Ps^0, Ps^1, Mv \rangle$ can trivially be obtained by iteratively computing dominions of some player, namely sets of positions from which that player has a strategy to win the game. Once an α-dominion D for player $\alpha \in \{0, 1\}$ is found, its α-attractor $atr^\alpha_\partial(D)$ gives an α-maximal dominion containing D. In other words, α cannot force any position outside D to enter this set. The subgame $\partial \setminus atr^\alpha_\partial(D)$ can then be solved by iterating the process. This procedure is reported in Algorithm 1. The crucial problem to address, therefore, consists in computing a dominion for some player in the game. The difficulty here is that, in general, no unique priority exists that satisfies the winning condition for a player along all the plays inside the dominion we are looking for. In fact, that value depends on the strategy chosen by the opponent. Our solution to this problem is to proceed in a bottom-up fashion, starting from a weaker notion of α-dominion, called *quasi α-dominion*. Then, we compose quasi α-dominions until we

obtain an α-dominion. Intuitively, a quasi α-dominion is a set of positions on which player α has a strategy whose induced plays either remain in the set forever and are winning for α, or can exit from it passing through a specific set of positions, *i.e.*, the escapes of the set itself. This notion is formalized by the following definition.

Algorithm 1: Parity Game Solver.	**Algorithm 2:** The Searcher.
signature $\mathsf{sol}_\Gamma : \mathcal{P} \to_\partial (2^{\mathrm{Ps}_\partial} \times 2^{\mathrm{Ps}_\partial})$	**signature** $\mathsf{src}_\Gamma : \mathcal{P} \to_\partial \mathrm{QD}_\partial^+$
function $\mathsf{sol}_\Gamma(\partial)$	**function** $\mathsf{src}_\Gamma(\partial)$
1 **if** $\mathrm{Ps}_\partial = \emptyset$ **then**	1 **return** $\mathsf{src}_{\Gamma(\partial)}(\top_{\Gamma(\partial)})$
2 **return** $(\emptyset, \emptyset)$	**signature** $\mathsf{src}_\mathcal{D} : \mathrm{S}_\mathcal{D} \to \mathrm{QD}_{\partial_\mathcal{D}}^+$
else	**function** $\mathsf{src}_\mathcal{D}(s)$
3 $(\mathrm{R}, \alpha) \leftarrow \mathsf{src}_\Gamma(\partial)$	1 $(\mathrm{Q}, \alpha) \leftarrow \Re_\mathcal{D}(s)$
4 $\mathrm{R}^\star \leftarrow \mathsf{atr}_\partial^\alpha(\mathrm{R})$	2 **if** $(\mathrm{Q}, \alpha) \in \mathrm{QD}_{\partial_\mathcal{D}}^+$ **then**
5 $(\mathrm{W_0}', \mathrm{W_1}') \leftarrow \mathsf{sol}_\Gamma(\partial \setminus \mathrm{R}^\star)$	3 **return** (Q, α)
6 $(\mathrm{W}_\alpha, \mathrm{W}_{\overline{\alpha}}) \leftarrow (\mathrm{W}_\alpha' \cup \mathrm{R}^\star, \mathrm{W}_{\overline{\alpha}}')$	**else**
7 **return** $(\mathrm{W_0}, \mathrm{W_1})$	4 **return** $\mathsf{src}_\mathcal{D}(s \downarrow_\mathcal{D} (\mathrm{Q}, \alpha))$

Definition 1 (Quasi Dominion) Let $\partial \in \mathcal{P}$ be a game and $\alpha \in \{0, 1\}$ a player. A non-empty set of positions $\mathrm{Q} \subseteq \mathrm{Ps}_\partial$ is a *quasi α-dominion* in ∂ if there exists an α-strategy $\sigma_\alpha \in \mathrm{Str}_\partial^\alpha(\mathrm{Q})$ such that, for all $\overline{\alpha}$-strategies $\sigma_{\overline{\alpha}} \in \mathrm{Str}_\partial^{\overline{\alpha}}(\mathrm{Q})$, with $\mathsf{int}_\partial^{\overline{\alpha}}(\mathrm{Q}) \subseteq \mathsf{dom}(\sigma_{\overline{\alpha}})$, and positions $v \in \mathrm{Q}$, the induced play $\pi = \mathsf{play}_\partial((\sigma_0, \sigma_1), v)$ satisfies $\mathsf{pr}_\partial(\pi) \equiv_2 \alpha$, if π is infinite, and $\mathsf{lst}(\pi) \in \mathsf{esc}_\partial^{\overline{\alpha}}(\mathrm{Q})$, otherwise.

It is important to observe that the additional requirement that the opponent strategies be defined on all interior positions, formally $\mathsf{int}_\partial^{\overline{\alpha}}(\mathrm{Q}) \subseteq \mathsf{dom}(\sigma_{\overline{\alpha}})$, discards those strategies in which the opponent deliberately chooses to forfeit the play, by declining to take any move at some of its positions.

We say that a quasi α-dominion Q is α-*open* (resp., α-*closed*) if $\mathsf{esc}_\partial^{\overline{\alpha}}(\mathrm{Q}) \neq \emptyset$ (resp., $\mathsf{esc}_\partial^{\overline{\alpha}}(\mathrm{Q}) = \emptyset$). In other words, in a closed quasi α-dominion, player α has a strategy whose induced plays are all infinite and winning. Hence, when closed, a quasi α-dominion is a dominion for α in ∂. The set of pairs $(\mathrm{Q}, \alpha) \in 2^{\mathrm{Ps}_\partial} \times \{0, 1\}$, where Q is a quasi α-dominion, is denoted by QD_∂, and is partitioned into the sets QD_∂^- and QD_∂^+ of open and closed quasi α-dominion pairs, respectively. As an example, consider Figure 1. It is easy to see that the sets of positions $\{a\}$ and $\{b, c, d\}$ form two open quasi

Fig. 1: A simple game ∂.

0-dominion in the game ∂. Similarly, the sets of positions $\{c, d\}$ and $\{e\}$ are open quasi 1-dominion. Finally, the entire game ∂ is a closed quasi 0-dominion, or simply a 0-dominion, where the 0-strategy correspond to the bold arrows.
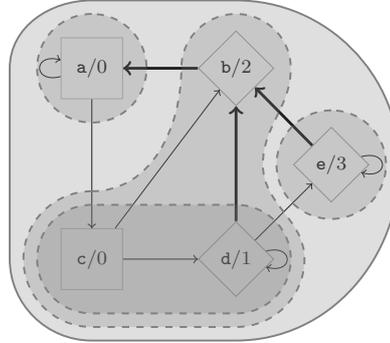
An expert reader might note that quasi α-dominions are loosely related with the concept of *snares*, introduced in [23] and used there for completely different purposes, namely to speed up the convergence of standard strategy improvement algorithms.

During the search for a dominion, we explore a suitable partial order, whose elements, called *states*, record information about the open quasi dominions computed so far. The search starts from the top element, where the quasi dominions are initialized to the sets of nodes with the same priority. At each step, a query is performed on the current state to extract a new quasi dominion, which is then used to compute a successor state, if it is open. If, on the other hand, it is closed, the search is over. Different query and successor operations can in principle be defined, even on the same partial order. However, such operations cannot be completely independent. To account for this intrinsic dependence, we introduce a *compatibility relation* between states and quasi dominions that can be extracted by the query operation. The pairs in such a relation also forms the domain of the successor function. The partial order together with the query and successor operations and the compatibility relation forms what we call a *dominion space*.

Definition 2 (Dominion Space) A *dominion space* for a game $\eth \in \mathcal{P}$ is a tuple $\mathcal{D} \triangleq \langle \eth, \mathcal{S}, \succ, \Re, \downarrow \rangle$, where *(1)* $\mathcal{S} \triangleq \langle S, \top, \prec \rangle$ is a well-founded partial order *w.r.t.* $\prec \subset S \times S$ with distinguished element $\top \in S$, *(2)* $\succ \subseteq \mathcal{S} \times QD_\eth^-$ is the *compatibility relation*, *(3)* $\Re : S \to QD_\eth$ is the *query function* mapping each element $s \in S$ to a quasi dominion pair $(Q, \alpha) \triangleq \Re(s) \in QD_\eth$ such that, if $(Q, \alpha) \in QD_\eth^-$ then $s \succ (Q, \alpha)$, and *(4)* $\downarrow : \succ \to S$ is the *successor function* mapping each pair $(s, (Q, \alpha)) \in \succ$ to the element $s^\star \triangleq s \downarrow (Q, \alpha) \in S$ with $s^\star \prec s$.

The *depth* of a dominion space $\mathcal{D}$ is the length of the longest chain in the underlying partial order $\mathcal{S}$ starting from $\top$. Instead, by *execution depth* of $\mathcal{D}$ we mean the length of the longest chain induced by the successor function $\downarrow$. Obviously, the execution depth is always bound by the depth.

Different dominion spaces can be associated to the same game. Therefore, in the rest of this section, we shall simply assume a function Γ mapping every game $\eth$ to a dominion space $\Gamma(\eth)$. Given the top element of $\mathcal{D} = \Gamma(\eth)$, Algorithm 2 searches for a dominion of either one of the two players by querying the current state s for a region pair (Q, α). If this is closed in $\eth$, it is returned as an α-dominion. Otherwise, a successor state $s \downarrow_\mathcal{D} (Q, \alpha)$ is computed and the search proceeds recursively from it. Clearly, since the partial order is well-founded, termination of the $\mathsf{src}_\mathcal{D}$ procedure is guaranteed. The total number of recursive calls is, therefore, the execution depth $\mathsf{d}_\mathcal{D}(n, m, k)$ of the dominion space $\mathcal{D}$, where n, m, and k are the number of positions, moves, and priorities, respectively. Hence, $\mathsf{src}_\mathcal{D}$ runs in time $O\big(\mathsf{d}_\mathcal{D}(n, m, k) \cdot (\mathsf{T}_\Re(n, m) + \mathsf{T}_\downarrow(n, m))\big)$, where $\mathsf{T}_\Re(n, m)$ and $\mathsf{T}_\downarrow(n, m)$ denote the time needed by the query and successor functions, respectively. Thus, the total time to solve a game is $O\big(m + n \cdot \mathsf{d}_\mathcal{D}(n, m, k) \cdot (\mathsf{T}_\Re(n, m) + \mathsf{T}_\downarrow(n, m))\big)$. Since the query and successor functions of the dominion space considered in the rest of the paper can be computed in linear time *w.r.t.* both n and m, the whole procedure terminates in time $O(n \cdot (n + m) \cdot \mathsf{d}_\mathcal{D}(n, m, k))$. As to the space requirements, observe that $\mathsf{src}_\mathcal{D}$ is a tail recursive algorithm. Hence, the upper bound on memory only depends on the space needed to encode the states of a dominion space, namely $O(\log \|\mathcal{D}\|)$, where $\|\mathcal{D}\|$ is the size of the partial order $\mathcal{S}$ associated with $\mathcal{D}$.

Soundness of the approach follows from the observation that quasi α-dominions closed in the entire game are winning for player α and so are their α-attractors. *Completeness*, instead, is ensured by the nature of dominion spaces. Indeed, algorithm $\mathsf{src}_\mathcal{D}$ always terminates by well-foundedness of the underlying partial order and, when it eventually does, a dominion for some player is returned. Therefore, the correctness of the algorithm reduces to proving the existence of a suitable dominion space, which is the subject of the next section.

4 Priority Promotion

In order to compute dominions, we shall consider a restricted form of quasi dominions that constrains the escape set to have the maximal priority in the game. Such quasi dominions are called *regions*.

Definition 3 (Region) A quasi α-dominion R is an α-region if $\mathsf{pr}(\eth) \equiv_2 \alpha$ and all the positions in $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R)$ have priority $\mathsf{pr}(\eth)$, *i.e.* $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R) \subseteq \mathsf{pr}_{\eth}^{-1}(\mathsf{pr}(\eth))$.

As a consequence of the above definition, if the opponent $\overline{\alpha}$ can escape from an α-region, it must visit a position with the highest priority in the region, which is of parity α. Similarly to the case of quasi dominions, we shall denote with $\mathrm{Rg}_{\eth}$ the set of region pairs in $\eth$ and with $\mathrm{Rg}_{\eth}^-$ and $\mathrm{Rg}_{\eth}^+$ the sets of open and closed region pairs, respectively. A closed α-region is clearly an α-dominion. As an example, consider again Figure 1. The singleton $\{e\}$ is the unique 1-region in the entire game $\eth$, which is also open, while $\{b\}$ is a non-maximal 0-region in the subgame $\eth \setminus \{e\}$ where e is removed. Finally, the set $\{b, c, d\}$ is a maximal open 0-region in the same subgame.

At this point, we have all the tools to explain the crucial steps underlying the search procedure. Open regions are not winning, as the opponent can force plays exiting from them. Therefore, in order to build a dominion starting from open regions, we look for a suitable sequence of regions that can be merged together until a closed one is found. Obviously, the merging operation needs to be applied only to regions belonging to the same player, in such a way that the resulting set of position is still a region of that player. To this end, a mechanism is proposed, where an α-region R in some game $\eth$ and an α-dominion D in a subgame of $\eth$ not containing R itself are merged together, if the only moves exiting from $\overline{\alpha}$-positions of D in the entire game lead to higher priority α-regions and R has the lowest priority among them. As we shall see, this ensures that the new region $R^\star \triangleq R \cup D$ has the same associated priority as R. This merging operation, based on the following proposition, is called *promotion* of the lower region to the higher one.

Proposition 1 (Region Merging) *Let* $\eth \in \mathcal{P}$ *be a game,* $R \subseteq \mathrm{Ps}_{\eth}$ *an* α-region, *and* $D \subseteq \mathrm{Ps}_{\eth \setminus R}$ *an* α-dominion *in the subgame* $\eth \setminus R$. *Then,* $R^\star \triangleq R \cup D$ *is an* α-region *in* $\eth$. *Moreover, if both* R *and* D *are* α-maximal *in* $\eth$ *and* $\eth \setminus R$, *respectively, then* $R^\star$ *is* α-maximal *in* $\eth$ *as well.*

Proof Since R is an α-region, there is an α-strategy σ_R such that, for all $\overline{\alpha}$-strategies $\sigma_{\overline{\alpha}} \in \mathrm{Str}_{\eth}^{\overline{\alpha}}(R)$, with $\mathsf{int}_{\eth}^{\overline{\alpha}}(R) \subseteq \mathsf{dom}(\sigma_{\overline{\alpha}})$, and positions $v \in R$, the play induced by the two strategies is either winning for α or exits from R passing through a position of the escape set $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R)$, which must be one of the position of maximal priority in $\eth$ and of parity α. Set D is, instead, an α-dominion in the game $\eth \setminus R$, therefore an α-strategy $\sigma_D \in \mathrm{Str}_{\eth \setminus R}$ exists that is winning for α from every position in D, regardless of the strategy $\sigma'_{\alpha} \in \mathrm{Str}_{\eth \setminus R}^{\overline{\alpha}}(D)$, with $\mathsf{int}_{\eth \setminus R}^{\overline{\alpha}}(D) \subseteq \mathsf{dom}(\sigma'_{\alpha})$, chosen by the opponent $\overline{\alpha}$. To show that $R^\star$ is an α-region, it suffices to show that the following three conditions hold: *(i)* it is a quasi α-dominion; *(ii)* the maximal priority of $\eth$ is of parity α; *(iii)* the escape set $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R^\star)$ is contained in $\mathsf{pr}_{\eth}^{-1}(\mathsf{pr}(\eth))$.

Condition *(ii)* immediately follows from the assumption that R is an α-region in $\eth$. To show that also Condition *(iii)* holds, we observe that, since D is an α-dominion in $\eth \setminus R$, the only possible moves exiting from $\overline{\alpha}$-positions of D in game $\eth$ must lead to R, *i.e.*, $\mathsf{esc}_{\eth}^{\overline{\alpha}}(D) \subseteq R$. Hence, the only escaping positions of $R^\star$, if any, must belong to R, *i.e.*

$\mathsf{esc}_{\eth}^{\overline{\alpha}}(R^{\star}) \subseteq \mathsf{esc}_{\eth}^{\overline{\alpha}}(R)$. Since R is an α-region in $\eth$, it hods that $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R) \subseteq \mathsf{pr}_{\eth}^{-1}(\mathsf{pr}(\eth))$. By transitivity, we conclude that $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R^{\star}) \subseteq \mathsf{pr}_{\eth}^{-1}(\mathsf{pr}(\eth))$.

Let us now consider Condition *(i)* and let the α-strategy $\sigma_{R^{\star}} \triangleq \sigma_R \cup \sigma_D$ be defined as the union of the two strategies above. Note that, being D and R disjoint sets of positions, $\sigma_{R^{\star}}$ is a well-defined strategy. We have to show that every path π compatible with $\sigma_{R^{\star}}$ and starting from a position in $R^{\star}$ is either winning for α or ends in a position of the escape set $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R^{\star})$.

First, observe that $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R^{\star})$ contains only those positions in the escaping set of R from which α cannot force to move into D, *i.e.* $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R^{\star}) = \mathsf{esc}_{\eth}^{\overline{\alpha}}(R) \setminus \mathsf{pre}_{\eth}^{\alpha}(D)$.

Let now π be a play compatible with $\sigma_{R^{\star}}$. If π is an infinite play, then it remains forever in $R^{\star}$ and we have three possible cases. If π eventually remains forever in D, then it is clearly winning for α, since $\sigma_{R^{\star}}$ coincides with σ_D on all the positions in D. Similarly, if π eventually remains forever in R, then it is also winning for α, as $\sigma_{R^{\star}}$ coincides with σ_R on all the positions in R. If, on the other hand, π passes infinitely often through both R and D, it necessarily visits infinitely often an escaping position in $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R) \subseteq \mathsf{pr}_{\eth}^{-1}(\mathsf{pr}(\eth))$, which has the maximal priority in $\eth$ and is of parity α. Hence, the parity of the maximal priority visited infinitely often along π is α and π is winning for player α. Finally, if π is a finite play, then it must end at some escaping position of R from where α cannot force to move to a position still in $R^{\star}$, *i.e.*, it must end in a position of the set $\mathsf{esc}_{\eth}^{\overline{\alpha}}(R) \setminus \mathsf{pre}_{\eth}^{\alpha}(D) = \mathsf{esc}_{\eth}^{\overline{\alpha}}(R^{\star})$. Therefore, $\mathsf{lst}(\pi) \in \mathsf{esc}_{\eth}^{\overline{\alpha}}(R^{\star})$. We can then conclude that $R^{\star}$ also satisfies Condition *(i)*.

Let us now assume, by contradiction, that $R^{\star}$ is not α-maximal. Then, there must be at least one position v belonging to $\mathsf{atr}_{\eth}^{\alpha}(R^{\star}) \setminus R^{\star}$, from which α can force entering $R^{\star}$ in one move. Assume first that v is an α-position. Then there is a move from v leading either to R or to D. But this means that v belongs to either $\mathsf{atr}_{\eth}^{\alpha}(R) \setminus R$ or $\mathsf{atr}_{\eth \setminus R}^{\alpha}(D) \setminus D$, contradicting α-maximality of those sets. If v is a $\overline{\alpha}$-position, instead, all its outgoing moves must lead to $R \cup D$. If all those moves lead to R, then $v \in \mathsf{atr}_{\eth}^{\alpha}(R) \setminus R$, contradicting α-maximality of R in $\eth$. If not, then in the subgame $\eth \setminus R$, the remaining moves from v must all lead to D. But then, $v \in \mathsf{atr}_{\eth \setminus R}^{\alpha}(D) \setminus D$, contradicting α-maximality of D in $\eth \setminus R$.   □

During the search, we keep track of the computed regions by means of an auxiliary priority function $r \in \mathbb{P}_{\eth} \triangleq \mathrm{Ps}_{\eth} \to \mathrm{Pr}_{\eth}$, called *region function*, which formalizes the intuitive notion of priority of a region described above. Initially, the region function coincides with the priority function $\mathsf{pr}_{\eth}$ of the entire game $\eth$. Priorities are considered starting from the highest one. A region of the same parity $\alpha \in \{0, 1\}$ of the priority p under consideration is extracted from the region function, by collecting the set of positions $r^{-1}(p)$. Then, its attractor $R \triangleq \mathsf{atr}_{\eth^{\star}}^{\alpha}(r^{-1}(p))$ is computed *w.r.t.* the subgame $\eth^{\star}$, which is derived from $\eth$ by removing the regions with priority higher than p. The resulting set forms an α-maximal set of positions from which the corresponding player can force a visit to positions with priority p. This first phase is called *region extension*. If the α-region R is open in $\eth^{\star}$, we proceed and process the next priority. In this case, we set the priority of the newly computed region to p. Otherwise, one of two situations may arise. Either R is closed in the whole game $\eth$ or the only $\overline{\alpha}$-moves exiting from R lead to higher regions of the same parity. In the former case, R is a α-dominion in the entire game and the search stops. In the latter case, R is only an α-dominion in the subgame $\eth^{\star}$, and a promotion of R to a higher region $R^{\natural}$ can be performed, according to Proposition 1. The search, then, restarts from the priority of $R^{\natural}$, after resetting to the original priorities in $\mathsf{pr}_{\eth}$ all the positions of the lower priority regions. The region $R^{\star}$ resulting from the union of $R^{\natural}$ and R will then be reprocessed and, possibly, extended in order to make it α-maximal. If R can be promoted to more than one region, the one with the lowest priority is chosen, so as to ensure the correctness of the merging operation. Due to the
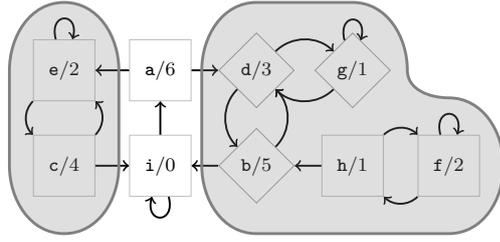
Fig. 2: Running example.

property of maximality, no $\overline{\alpha}$-moves from R to higher priority $\overline{\alpha}$-regions exist. Therefore, only regions of the same parity are considered in the promotion step. The correctness of region extension operation above, the remaining fundamental step in the proposed approach, is formalized by the following proposition.

Proposition 2 (Region Extension) *Let $\partial \in \mathcal{P}$ be a game and $R^{\star} \subseteq Ps_{\partial}$ an α-region in ∂. Then, $R \triangleq atr_{\partial}^{\alpha}(R^{\star})$ is an α-maximal α-region in ∂.*

Proof Since $R^{\star}$ is an α-region in ∂, then the maximal priority in ∂ is of parity α and $esc_{\partial}^{\overline{\alpha}}(R^{\star}) \subseteq pr_{\partial}^{-1}(pr(\partial))$. Hence, any position v in ∂ must have priority $pr_{\partial}(v) \leq pr(\partial)$. Player α can force entering $R^{\star}$ from every position in $atr_{\partial}^{\alpha}(R^{\star}) \setminus R^{\star}$, with a finite number of moves. Moreover, $R^{\star}$ is a quasi α-dominion and the priorities of the positions in $Ps_{\partial} \setminus R^{\star}$ are lower than or equal to $pr(\partial) \equiv_2 \alpha$. Hence, every play that remains in R forever either eventually remains forever in $R^{\star}$ and is winning for α, or passes infinitely often through $R^{\star}$ and $atr_{\partial}^{\alpha}(R^{\star}) \setminus R^{\star}$. In the latter case, that path must visit infinitely often a position in $esc_{\partial}^{\overline{\alpha}}(R^{\star})$ that has the maximal priority in ∂ and has parity α. Hence, the play is winning for α. If, on the other hand, $\overline{\alpha}$ can force a play to exit from R, it can do so only by visiting some position in $esc_{\partial}^{\overline{\alpha}}(R^{\star})$. In other words, $esc_{\partial}^{\overline{\alpha}}(R) \subseteq esc_{\partial}^{\overline{\alpha}}(R^{\star}) \subseteq pr_{\partial}^{-1}(pr(\partial))$. In either case, we conclude that R is an α-region in ∂. Finally, being R the result of an α-attractor, it is clearly α-maximal. $\square$

Figure 2 and Table 1 illustrate the search procedure on an example game, where diamond shaped positions belong to player 0 and square shaped ones to the opponent 1. Player 0 wins from every position, hence the 0-region containing all the positions is a 0-dominion in this case. Each cell of the table contains a computed region. A downward arrow denotes a region that is open in the subgame where it is computed, while an upward arrow means that the region gets to be promoted to the priority in the subscript. The index of each row corresponds to the priority of the region. Following the idea sketched above, the first region obtained

	1	2	3	4	5	6	7
6	a↓	⋯	⋯	⋯	⋯	a,b,d,g,i↓	⋯
5	b,f,h↓	⋯	⋯	b,d,f,g,h↓	⋯		
4	c↓	c,e↓	⋯	c↓	c,e↓	c↓	c,e,f,h↑$_6$
3	d↓	d↓	d,g↑$_5$				
2	e↑$_4$			e↑$_4$		e, f, h↑$_4$	
1		g↑$_3$					
0					i↑$_6$		

Table 1: PP simulation.

is the single-position 0-region $\{a\}$, which is open because of the two moves leading to d and e. At priority 5, the open 1-region $\{b, f, h\}$ is formed by attracting both f and h to b, which is open in the subgame where $\{a\}$ is removed. Similarly, the 0-region $\{c\}$ at priority 4 and the 1-region $\{d\}$ at priority 3 are open, once removed $\{a, b, f, h\}$ and $\{a, b, c, f, h\}$, respectively, from the game. At priority 2, the 0-region $\{e\}$ is closed in the corresponding subgame. However, it is not closed in the whole game, since it has a move leading to c, *i.e.*, to region 4. A promotion of $\{e\}$ to 4 is then performed, resulting in the new 0-region $\{c, e\}$. The search resumes at the corresponding priority and, after computing the extension of such a region via the attractor, we obtain that it is still open in the corresponding subgame. Consequently, the 1-region of priority 3 is recomputed and, then, priority 1 is processed to build the 1-region $\{g\}$. The latter is closed in the associated subgame, but not in the original game, because of a move leading to position d. Hence, another promotion is performed, leading to closed region in Row 3 and Column 3, which in turn triggers a promotion to 5. Observe that every time a promotion to a higher region is performed, all positions of the regions at lower priorities are reset to their original priorities. The iteration of the region forming and promotion steps proceeds until the configuration in Column 7 is reached. Here only two 0-regions are present: the open region 6 containing $\{a, b, d, g, i\}$ and the closed region 4 containing $\{c, e, f, h\}$. The second one has a move leading to the first one, hence, it is promoted to its priority. This last operation forms a 0-region containing all the positions of the game. It is obviously closed in the whole game and is, therefore, a 0-dominion.

Note that, the positions in 0-region $\{c, e\}$ are reset to their initial priorities, when 1-region $\{d, g\}$ in Column 3 is promoted to 5. Similarly, when 0-region $\{i\}$ in Column 5 is promoted to 6, the priorities of the positions in both regions $\{b, d, f, g, h\}$ and $\{c, e\}$, highlighted by the gray areas, are reset. This is actually necessary for correctness, at least in general. In fact, if region $\{b, d, f, g, h\}$ were not reset, the promotion of $\{i\}$ to 6, which also attracts b, d, and g, would leave $\{f, h\}$ as a 1-region of priority 5. However, according to Definition 3, this is not a 1-region. Even worse, it would also be considered a closed 1-region in the entire game, without being a 1-dominion, since it is actually an open 0-region. This shows that, in principle, promotions to an higher priority require the reset of previously built regions of lower priorities.

In the rest of this section, we shall formalize the intuitive idea described above. The necessary conditions under which promotion operations can be applied are also stated. Finally, query and successor algorithms are provided, which ensure that the necessary conditions are easy to check and always met when promotions are performed.

4.1 The PP Dominion Space.

In order to define the dominion space induced by the *priority-promotion mechanism* (PP, for short), we need to introduce some additional notation. Given a priority function $r \in \mathbb{P}_{\supset}$ and a priority $p \in \Pr$, we denote by $r^{(\geq p)}$ (*resp.*, $r^{(>p)}$ and $r^{(<p)}$) the function obtained by restricting the domain of r to the positions with priority greater than or equal to p (*resp.*, greater than and lower than p). Formally, $r^{(\geq p)} \triangleq r{\restriction}\{v \in \mathsf{dom}(r) : r(v) \geq p\}$ (*resp.*, $r^{(>p)} \triangleq r{\restriction}\{v \in \mathsf{dom}(r) : r(v) > p\}$ and $r^{(<p)} \triangleq r{\restriction}\{v \in \mathsf{dom}(r) : r(v) < p\}$). By $\partial_r^{\leq p}$ we denote the largest subgame contained in the structure $\partial \setminus \mathsf{dom}\left(r^{(>p)}\right)$, which is obtained by removing from ∂ all the positions in the domain of $r^{(>p)}$. A priority function $r \in \mathbb{R}_{\supset} \subseteq \mathbb{P}_{\supset}$ in ∂ is a *region function* iff, for all priorities $q \in \mathsf{rng}(r)$ with $\alpha \triangleq q \bmod 2$, it holds that $r^{-1}(q) \cap \Ps_{\partial_r^{\leq q}}$ is an α-region in the subgame $\partial_r^{\leq q}$, if non-empty. In addition, we say that r

is *maximal above* $p \in \mathrm{Pr}$ iff, for all $q \in \mathsf{rng}(\mathsf{r})$ with $q > p$, we have that $\mathsf{r}^{-1}(q)$ is α-maximal in $\partial_{\mathsf{r}}^{\leq q}$ with $\alpha \triangleq q \bmod 2$.

To account for the current status of the search of a dominion, the states s of the corresponding dominion space need to contain the current region function r and the current priority p reached by the search in ∂. To each of such states $s \triangleq (\mathsf{r}, p)$, we then associate the *subgame at s* defined as $\partial_s \triangleq \partial_{\mathsf{r}}^{\leq p}$, representing the portion of the original game that still has to be processed. For instance, the state s corresponding to Row 4 Column 4 in Table 1 is $(\mathsf{r}, 4)$, where the region function r is such that $\mathsf{r}(\mathtt{a}) = 6, \mathsf{r}(x) = 5$, for $x \in \{\mathtt{b}, \mathtt{d}, \mathtt{f}, \mathtt{g}, \mathtt{h}\}$, and $\mathsf{r}(x) = \mathsf{pr}(x)$, for $x \in \{\mathtt{c}, \mathtt{e}, \mathtt{i}\}$. In addition, the subgame ∂_s of s only contains the positions $\mathtt{c}$, $\mathtt{e}$, and $\mathtt{i}$.

We can now formally define the *Priority Promotion* dominion space, by characterizing the corresponding state space and compatibility relation. Moreover, algorithms for the query and successor functions of that space are provided.

Definition 4 (State Space) A *state space* is a tuple $\mathcal{S}_{\partial} \triangleq \langle \mathrm{S}_{\partial}, \top_{\partial}, \prec_{\partial} \rangle$, where its components are defined as prescribed in the following:

1. $\mathrm{S}_{\partial} \subseteq \mathbb{R}_{\partial} \times \mathrm{Pr}_{\partial}$ is the set of all pairs $s \triangleq (\mathsf{r}, p)$, called *states*, composed of a region function $\mathsf{r} \in \mathbb{R}_{\partial}$ and a priority $p \in \mathrm{Pr}_{\partial}$ such that *(a)* r is maximal above p, *(b)* $p \in \mathsf{rng}(\mathsf{r})$, and *(c)* $\mathsf{r}^{(<p)} \subseteq \mathsf{pr}_{\partial}^{(<p)}$;
2. $\top_{\partial} \triangleq (\mathsf{pr}_{\partial}, \mathsf{pr}(\partial))$;
3. for any two states $s_1 \triangleq (\mathsf{r}_1, p_1), s_2 \triangleq (\mathsf{r}_2, p_2) \in \mathrm{S}_{\partial}$, it holds that $s_1 \prec_{\partial} s_2$ iff either *(a)* there exists a priority $q \in \mathsf{rng}(\mathsf{r}_1)$ with $q \geq p_1$ such that *(a.i)* $\mathsf{r}_1^{(>q)} = \mathsf{r}_2^{(>q)}$ and *(a.ii)* $\mathsf{r}_2^{-1}(q) \subset \mathsf{r}_1^{-1}(q)$, or *(b)* both *(b.i)* $\mathsf{r}_1 = \mathsf{r}_2$ and *(b.ii)* $p_1 < p_2$ hold.

The state space specifies the configurations in which the priority promotion procedure can reside and the relative order that the successor function must satisfy. In particular, for a given state $s \triangleq (\mathsf{r}, p)$, every region $\mathsf{r}^{-1}(q)$, with priority $q > p$, recorded in the region function r has to be α-maximal, where $\alpha = q \bmod 2$. This implies that $\mathsf{r}^{-1}(q) \subseteq \mathrm{Ps}_{\partial_{\mathsf{r}}^{\leq q}}$. Moreover, the current priority p of the state must be the priority of an actual region in r. Finally, all the regions recorded in r at any priority q lower than p must contain positions that have the same priority q in the original priority function pr_{∂} of the game. As far as the order is concerned, a state s_1 is strictly smaller than another state s_2 if either there is a region recorded in s_1 at some higher priority q that strictly contains the corresponding one in s_2 and all regions above q are equal in the two states, or state s_1 is currently processing a lower priority than the one of s_2.

At this point, we can determine the regions that are compatible with a given state. They are the only ones that the query function is allowed to return and that can then be used by the successor function to make the search progress in the dominion space. Intuitively, a region pair (R, α) is compatible with a state $s \triangleq (\mathsf{r}, p)$ if it is an α-region in the current subgame ∂_s. Moreover, if such region is α-open in that game, it has to be α-maximal, and it has to necessarily contain the current region $\mathsf{r}^{-1}(p)$ of priority p in r. These three accessory properties ensure that the successor function is always able to cast R inside the current region function r and obtain a new state.

Definition 5 (Compatibility Relation) An open quasi-dominion pair $(\mathrm{R}, \alpha) \in \mathrm{QD}_{\partial}^{-}$ is *compatible* with a state $s \triangleq (\mathsf{r}, p) \in \mathrm{S}_{\partial}$, in symbols $s \succ_{\partial} (\mathrm{R}, \alpha)$, iff *(1)* $(\mathrm{R}, \alpha) \in \mathrm{Rg}_{\partial_s}$ and *(2)* if R is α-open in ∂_s then *(2.a)* R is α-maximal in ∂_s and *(2.b)* $\mathsf{r}^{-1}(p) \subseteq \mathrm{R}$.

Algorithm 3 provides a possible imple-
mentation for the query function compatible
with the priority-promotion mechanism. Let
$s \triangleq (\mathsf{r}, p)$ be the current state. Line 1 sim-
ply computes the parity α of the priority to
process in that state. Line 2, instead, com-
putes in game $\eth_s$ the attractor *w.r.t.* player
α of the region contained in r at the current
priority p. The resulting set R is, according

Algorithm 3: Query Function.

---

signature $\Re_\eth : \mathrm{S}_\eth \to (2^{\mathrm{Ps}_\eth} \times \{0, 1\})$
function $\Re_\eth(s)$
    let $(\mathsf{r}, p) = s$ **in**
1       $\alpha \leftarrow p \bmod 2$
2       $\mathrm{R} \leftarrow \mathsf{atr}^\alpha_{\eth_s}\big(\mathsf{r}^{-1}(p)\big)$
3   **return** (R, α)

---

to Proposition 2, an α-maximal α-region in $\eth_s$ containing $\mathsf{r}^{-1}(p)$.

Before continuing with the description of the implementation of the successor function,
we need to introduce the notion of *best escape priority* for player $\overline{\alpha}$ *w.r.t.* an α-region R
of the subgame $\eth_s$ and a region function r in the whole game $\eth$. Informally, such a value
represents the best priority associated with an α-region contained in r and reachable by $\overline{\alpha}$
when escaping from R. To formalize this concept, let $I \triangleq Mv_\eth \cap ((\mathrm{R} \cap \mathrm{Ps}^{\overline{\alpha}}_\eth) \times (\mathsf{dom}(\mathsf{r}) \setminus \mathrm{R}))$
be the *interface relation* between R and r, *i.e.*, the set of $\overline{\alpha}$-moves exiting from R and
reaching some position within a region recorded in r. Then, $\mathsf{bep}^{\overline{\alpha}}_\eth(\mathrm{R}, \mathsf{r})$ is set to the minimal
priority among those regions containing positions reachable by a move in I. Formally,
$\mathsf{bep}^{\overline{\alpha}}_\eth(\mathrm{R}, \mathsf{r}) \triangleq \min(\mathsf{rng}(\mathsf{r} {\restriction} \mathsf{rng}(I)))$. Note that, if R is a closed α-region in $\eth_s$, then $\mathsf{bep}^{\overline{\alpha}}_\eth(\mathrm{R}, \mathsf{r})$
is necessarily of parity α and greater than the priority p of R. This property immediately
follows from the maximality of r above p in any state of the dominion space. Indeed, no
move of an $\overline{\alpha}$-position can lead to a $\overline{\alpha}$-maximal $\overline{\alpha}$-region. For instance, in the example of
Figure 2, for 0-region $\mathrm{R} = \{\mathtt{e}, \mathtt{f}, \mathtt{h}\}$ with priority equal to 2 in column 6, we have that
$I = \{(\mathtt{e}, \mathtt{c}), (\mathtt{h}, \mathtt{b})\}$ and $\mathsf{r} {\restriction} \mathsf{rng}(I) = \{(\mathtt{c}, 4), (\mathtt{b}, 6)\}$. Hence, $\mathsf{bep}^1_\eth(\mathrm{R}, \mathsf{r}) = 4$.

In order to perform a reset of the priority of some positions in the game after a promotion,
we use the completing operator $\uplus$. Taken two partial function $\mathsf{f}, \mathsf{g} \colon \mathrm{A} \rightharpoonup \mathrm{B}$, the operator
returns the partial function $\mathsf{f} \uplus \mathsf{g} \colon \mathrm{A} \rightharpoonup \mathrm{B}$, which is equal to g on its domain and assumes the
same values as g on the remaining part of the set A.

Algorithm 4 implements the succes-
sor function informally described at the
beginning of the section. Given the cur-
rent state s and a compatible region pair
(R, α) open in the whole game as inputs,
it produces a successor state $s^\star \triangleq (\mathsf{r}^\star, p^\star)$
in the dominion space. It first checks
whether R is open also in the subgame $\eth_s$
(Line 1). If this is the case, it assigns pri-
ority p to region R and stores it in the new
region function $\mathsf{r}^\star$ (Line 2). The new cur-
rent priority $p^\star$ is, then, computed as the
highest priority lower than p in $\mathsf{r}^\star$ (Line 3).
If, on the other hand, R is closed in $\eth_s$, a

Algorithm 4: Successor Function.

---

signature $\downarrow_\eth \colon \succ_\eth \to \Delta_\eth \times \mathrm{Pr}_\eth$
function $s \downarrow_\eth (\mathrm{R}, \alpha)$
    let $(\mathsf{r}, p) = s$ **in**
1   **if** $(\mathrm{R}, \alpha) \in \mathrm{Rg}^-_{\eth_s}$ **then**
2         $\mathsf{r}^\star \leftarrow \mathsf{r}[\mathrm{R} \mapsto p]$
3         $p^\star \leftarrow \max\left(\mathsf{rng}\left(\mathsf{r}^{\star(<p)}\right)\right)$
      else
4         $p^\star \leftarrow \mathsf{bep}^{\overline{\alpha}}_\eth(\mathrm{R}, \mathsf{r})$
5         $\mathsf{r}^\star \leftarrow \mathsf{pr}_\eth \uplus \mathsf{r}^{(\geq p^\star)}[\mathrm{R} \mapsto p^\star]$
6   **return** $(\mathsf{r}^\star, p^\star)$

---

promotion merging R with some other α-region contained in r is required. The next priority
$p^\star$ is set to the bep of R for player $\overline{\alpha}$ in the entire game $\eth$ *w.r.t.* r (Line 4). Region R is, then,
promoted to priority $p^\star$ and all the priorities below $p^\star$ in the current region function r are
reset (Line 5). The correctness of this last operation follows from Proposition 1.

As already observed in Section 3, a dominion space, together with Algorithm 2, provides
a sound and complete solution procedure. The following theorem states that the priority-

promotion mechanism presented above is indeed a dominion space. For the sake of readability, the proof is provided in the appendix.

Theorem 1 (Dominion Space) *For a game ∂, the structure $\mathcal{D}_{\partial} \triangleq \langle \partial, \mathcal{S}_{\partial}, \succ_{\partial}, \Re_{\partial}, \downarrow_{\partial} \rangle$, where $\mathcal{S}_{\partial}$ is given in Definition 4, $\succ_{\partial}$ is the relation of Definition 5, and $\Re_{\partial}$ and $\downarrow_{\partial}$ are the functions computed by Algorithms 3 and 4 is a dominion space.*

4.2 Complexity of PP Dominion Space

To assess the complexity of the PP approach, we produce an estimate of the size and depth of the dominion space $\mathcal{R}_{\partial}$. This provides upper bounds for both the time and space complexities needed by the search procedure $\mathsf{src}_{\mathcal{R}_{\partial}}$ that computes dominions. By looking at the definition of state space $\mathcal{S}_{\partial}$, it is immediate to see that, for a game ∂ with n positions and k priorities, the number of states is bounded by k^{n+1}. Indeed, there are at most k^n functions $\mathsf{r} : \mathrm{Ps}_{\partial} \to \mathrm{Pr}_{\partial}$ from positions to priorities that can be used as region function of a state. Moreover, every such function can be associated with at most k current priorities.

Measuring the depth is a little trickier. A relatively coarse bound can be obtained by observing that there is an homomorphism from $\mathcal{S}_{\partial}$ to the well-founded partial order, in which the region function r of a state is replaced by a partial function $\mathsf{f} : \mathrm{Pr}_{\partial} \rightharpoonup [1, n]$ with the following properties: it assigns to each priority $p \in \mathsf{rng}(\mathsf{r})$ the size $\mathsf{f}(p)$ of the associated region $\mathsf{r}^{-1}(p)$. The order $(\mathsf{f}_1, p_1) \prec (\mathsf{f}_2, p_2)$ between two pairs is derived from the one on the states, by replacing $\mathsf{r}_2^{-1}(q) \subset \mathsf{r}_1^{-1}(q)$ with $\mathsf{f}_2(q) < \mathsf{f}_1(q)$. This homomorphism ensures that every chain in $\mathcal{S}_{\partial}$ corresponds to a chain in the new partial order. Moreover, there are exactly $\binom{n+k}{k}$ partial functions f such that $\sum_{p \in \mathsf{dom}(\mathsf{f})} \mathsf{f}(p) \leq n$. Consequently, thanks to Stirling's approximation of the factorial function, every chain cannot be longer than $k \binom{n+k}{k} = (n+k) \binom{n+k-1}{k-1} \leq (n+k) \left(\mathsf{e}(n/(k-1) + 1) \right)^{k-1}$, where e is Euler's number.

Theorem 2 (Size & Depth Upper Bounds) *The size and depth of a PP dominion space $\mathcal{R}$ with $n \in \mathbb{N}_+$ positions and $k \in [1, n]$ priorities are bounded by k^{n+1} and $(n+k) \left(\mathsf{e}(n/(k-1) + 1) \right)^{k-1} = \mathrm{O}\left(kn \left(\frac{\mathsf{e}n}{k-1} \right)^{k-1} \right)$, respectively.*

Unfortunately, due to the reset operations performed after each promotion, an exponential worst-case can actually be built. Indeed, consider the game $\partial_{l,h}$ having all positions ruled by player 0 and containing h chains of length $2l + 1$ that converge into a single position of priority 0 with a self loop. The i-th chain has a head of priority $2(2h - i) + 1$ and a body composed of l blocks of two positions with priority $2i - 1$ and $2i$, respectively. The first position in each block also has a self loop. An instance of this game with $l = 2$ and $h = 4$ is depicted in Figure 3. The labels of the positions, deprived of the possible apexes, correspond to the associated priorities and the highlighted area at the bottom of the figure groups together the last blocks of the chains. Intuitively, the execution depth of the PP dominion

Fig. 3: The $\partial_{2,4}^{\mathrm{PP}}$ game.

space for this game is exponential, since the consecutive promotion operations performed on each chain can simulate the increments of a counter up to l. Also, the priorities are chosen in such a way that, when the i-th counter is incremented, all the j-th counters with $j \in \,]i, h]$
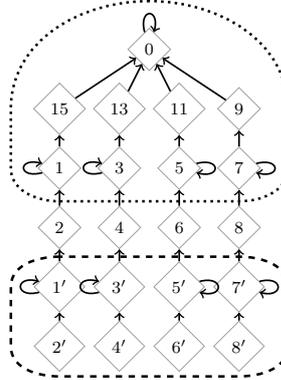
are reset. Therefore, the whole game simulates a counter with h digits taking values from 0 to l. Hence, the overall number of performed promotions is $(l+1)^h$. The search procedure on $\partial_{2,4}$ starts by building the four open 1-regions $\{15\}$, $\{13\}$, $\{11\}$, and $\{9\}$ and the open 0-region $\{8, 7', 8'\}$, where we use apexes to distinguish different positions with the same priority. This state represents the configuration of the counter, where all four digits are set to 0. The closed 1-region $\{7\}$ is then found and promoted to 9. Consequently, the previously computed 0-region with priority 8 is reset and the new region is maximized to obtain the open 1-region $\{9, 7, 8\}$. Now, the counter is set to 0001. After that, the open 0-region $\{8'\}$ and the closed 1-region $\{7'\}$ are computed. The latter one is promoted to 9 and maximized to attract position $8'$. This completes the 1-region containing the entire chain ending in 9. The value of the counter is now 0002. At this point, immediately after the construction of the open 0-region $\{6, 5', 6'\}$, the closed 1-region $\{5\}$ is found, promoted to 11, and maximized to absorb position 6. Due to the promotion, the positions in the 1-region with priority 9 are reset to their original priority and all the work done to build it gets lost. This last operation represents the reset of the least significant digit of the counter, caused by the increment of the second one, *i.e.*, the counter displays 0010. Following similar steps, the process carries on until each chain is grouped in a single region. The corresponding state represents the configuration of the counter in which all digits are set to l. Thus, after an exponential number promotions, the closed 0-region $\{0\}$ is eventually obtained as solution.

Theorem 3 (Execution-Depth Lower Bounds) *For all numbers $h \in \mathbb{N}$, there exists a* PP *dominion space $\mathcal{D}_h^{\mathrm{PP}}$ with $k = 2h + 1$ positions and priorities, whose execution depth is $3 \cdot 2^h - 2 = \Theta\left(2^{k/2}\right)$. Moreover, for all numbers $l \in \mathbb{N}_+$, there exists a* PP *dominion space $\mathcal{D}_{l,h}^{\mathrm{PP}}$ with $n = (2l+1) \cdot h + 1$ positions and $k = 3h + 1$ priorities, whose execution depth is $((3l+1) \cdot (l+1)^h - 1)/l - 2 = \mathrm{O}\left((3n/(2(k-1)))^{k/3}\right)$.*

Proof The single-player game ∂_h^{PP}, with $k = 2h + 1$ positions and priorities, contains a position with priority 0 plus $2h$ positions spanning all the odd priorities from 1 to $4h - 1$. Moreover, the associated moves are those depicted in the highlighted top section of Figure 3. Intuitively, this game encodes a binary counter, where each one of the h chains, composed of two positions of odd priorities connected to position 0, represents a digit. The promotion operation of region $2i - 1$ to region $2(2h - i) + 1$ corresponds to the increment of the i-th digit. As a consequence, the number of promotions for the PP algorithm is equal to the number of configurations of the counter, except for the initial one. Now, to provide a lower bound for the depth of the associated dominion space $\mathcal{D}_h^{\mathrm{PP}}$, consider the recursive function $\mathsf{Q}(h)$ with base case $\mathsf{Q}(0) = 1$ and inductive case $\mathsf{Q}(h) = 2\mathsf{Q}(h-1) + 2$. This function counts the number of queries executed by the the PP algorithm on game ∂_h^{PP}. The correctness of the base case is trivial. Indeed, when h is equal to 0, there are no chains and the only possible region is the one containing position 0. Let us now consider the inductive case $h > 0$. A first query is required to obtain region $\{4h - 1\}$. Then, before reaching the promotion of region $\{1\}$ to region $\{4h - 1\}$, *i.e.*, to set the most-significant digit, all the other digits must be set. To do this, the number of necessary queries is equal to those required on the game instance with $h - 1$ chains minus the query that computes region $\{0\}$, *i.e.*, $\mathsf{Q}(h-1) - 1$. At this point, two additional queries are counted, when region $\{1\}$ is obtained and then merged to region $\{4h - 1\}$. As described above in the execution of the game depicted in Figure 3, this promotion unnecessarily resets all lesser-significant digits. Therefore, $\mathsf{Q}(h-1) - 1$ counts exactly the queries needed to set again to 1 all the digits reset after this promotion. Finally, a query for the 0-closed region $\{0\}$ is required. Summing

up, we have $\mathsf{Q}(h) = 1 + (\mathsf{Q}(h-1) - 1) + 2 + (\mathsf{Q}(h-1) - 1) + 1 = 2\mathsf{Q}(h-1) + 2$, as claimed above. A trivial proof by induction shows that $\mathsf{Q}(h) = 3 \cdot 2^h - 2 = \Theta\left(2^{k/2}\right)$.

We can now turn to the more complex single-player game $\partial^{\mathrm{PP}}_{l,h}$, with $n = (2l+1) \cdot h + 1$ positions and $k = 3h + 1$ priorities, which contains a position with priority 0 plus $(2l+1) \cdot h$ positions, spanning all the priorities from 1 to $2h$ and all odd priorities from $2h+1$ to $4h-1$ as depicted in the entire Figure 3. To provide a lower bound for the depth of the associated dominion space $\mathcal{D}^{\mathrm{PP}}_{l,h}$, consider the recursive function $\mathsf{Q}(h)$ with base case $\mathsf{Q}(l,0) = 1$ and inductive case $\mathsf{Q}(l,h) = (1+l)\mathsf{Q}(l,h-1) + 2l+1$. Also in this case, the correctness of the base cases is trivial. Indeed, when h is equal to 0, there are no chains, independently of the parameter l. Hence, the only possible region is the one containing position 0. For the inductive case $h > 0$, a first query is required to compute region $\{4h-1\}$. Then, before reaching the promotion to region $\{4h-1\}$ of the adjacent region $\{1\}$, $i.e.$, to set the most-significant digit to 1, all the other digits must be set to $l-1$. This requires the same number of queries required to process the game with $h-1$ chains minus the query collecting region $\{0\}$, $i.e.$, $\mathsf{Q}(l,h-1) - 1$. At this point, three additional queries are counted: region $\{2,1',2',\ldots\}$ is computed first, followed by region $\{1\}$; the latter is then merged to region $\{4h-1\}$ absorbing position 2 from the first region. Note that the region of priority 2 contains all positions in the first chain, except for positions 1 and $4h-1$. As in the previous game, such a promotion unnecessarily resets all lesser-significant digits. Therefore, $\mathsf{Q}(l,h-1) - 1$ counts exactly the queries needed to set again to $l-1$ all digits reset after this promotion. After that, and similarly to what is described above, we need to set the most-significant digit to 2, by performing other three queries: $\{2',\ldots\}$ is computed first, followed by region $\{1'\}$; the latter is then merged to region $\{4h-1,1,2\}$, absorbing position $2'$ from the first region. Again, this promotion resets all lesser-significant digits, which requires $\mathsf{Q}(l,h-1) - 1$ steps to be set to $l-1$. This behavior is repeated $l-2$ more times, until the first chain is completely absorbed by the region of highest priority $4h-1$. Finally, a query for the 0-closed region $\{0\}$ is required. Summing up, we have $\mathsf{Q}(n,h) = 1 + (\mathsf{Q}(h-1) - 1) + (3 + (\mathsf{Q}(l,h-1) - 1)) \cdot l + 1 = (1+l)\mathsf{Q}(l,h-1) + 2l+1$, as claimed above. Also in this case, an easy induction shows that $\mathsf{Q}(l,h) = ((3l+1) \cdot (l+1)^h - 1)/l - 2 = \mathrm{O}\left((3n/(2(k-1)))^{k/3}\right)$. □

Observe that, in the above theorem, we provide two different exponential lower bounds. The general one, with $k/3$ as exponent and a parametric base, is the result of the game $\partial_{l,h}$ described in the previous paragraph, where $k = 3h + 1$. The other bound, instead, has a base fixed to 2, but a larger exponent $k/2$. We conjecture that the given upper bound could be improved to match the exponent $k/2$ of this lower bound. In this way, we would obtain an algorithm with an asymptotic behavior comparable with the one exhibited by the small-progress measure procedure [35].

5 Subgame Decomposition

As described in the previous section, the search procedure for the basic PP approach, displayed in Algorithm 2, consists in finding quasi dominions that, whenever closed in the current subgame, are suitably merged with previously computed regions until a dominion in the original game ∂ is found. In particular, when the current quasi dominion is open, the algorithm searches for the next one in a subgame ∂' of ∂. By Algorithm 3 and Lines 2-3 of Algorithm 4, such a subgame corresponds to $\partial_s \setminus \mathrm{R}$, where ∂_s and R are, respectively, the current subgame $w.r.t.$ to state s and the region extracted by the query function on s itself. Due

to this design choice and the fact that the regions are created in decreasing order of priority, the sequence of subgames follows the same order and, in particular, the current priority p contained in the state $s = (\mathsf{r}, p)$ also corresponds to the maximal priority in the part of game yet to be analyzed. This is a very simple and efficient way to compute subgames, but it does not necessarily translate into an efficient way to obtain a solution in terms of the number of promotions required. In fact, the lesser the moves between ∂' and the positions outside ∂', the easier it is to find a dominion in the subsequent iterations, as the probability that a promotion occurs reduces. In this section we propose a generalization of the PP algorithm that does not commit to a specific way of computing the next subgame to process during its dominion search phase. The result is a parametric version of the priority promotion approach that, by exploiting the topology of the underlying game, can be instantiated with different game decomposition techniques.

To this end, we first need to address the crucial issue that the correctness of promotion mechanism relies, according to Proposition 1, on the fact that the positions contained in a dominion to be promoted can only reach some previously computer region. To ensure that this property is preserved, we need to require that, after computing region R in game ∂_s, whatever the subgame ∂' of $\partial_s \setminus R$ we chose to consider next, it cannot contain positions that reach portions of the game yet to be analyzed, namely positions in the set $\mathrm{Ps}_{\partial_s} \setminus (R \cup \mathrm{Ps}_{\partial'})$. Since there is no unique way to choose a suitable subgame, the new algorithm we propose in the following is parametric on a function F that, given a game ∂, computes the positions of the subgame ∂' on which the search for quasi dominions needs to proceed. The standard PP algorithm can be obtained by simply setting $\mathsf{F}(\partial) = \mathrm{Ps}_\partial$. A sharper heuristic, instead, instantiates F with a procedure that collects the positions in some minimal Strongly Connected Component (SCC) of the graph underlying the game ∂. An alternative coarser possibility, which, however, incurs in less overhead, simply computes all the positions reachable from an arbitrary initial position in the game. It is immediate to see that all these instantiations satisfy the above requirement on the subgames.

A second issue concerns the way in which the processed subgames are kept track of. In the original PP algorithm, where $\partial' = \partial_s \setminus R$, the current subgame ∂_s coincides with $\partial_\mathsf{r}^{\leq p}$. Hence, the region function r, together with the priority p, suffices to identify the current subgame ∂_s. On the other hand, if ∂' is chosen as a strict subgame of $\partial_\mathsf{r}^{\leq p}$, we need to introduce an auxiliary function to keep track of the computed subgames and identify the current one. This role is played by a second priority function $\mathsf{g} : \mathrm{Ps} \to \mathrm{Pr}$, called *subgame function*. At the beginning of the procedure, g maps every position to $\mathsf{pr}(\partial)$ in order to ensure that the current subgame is the entire game. Every time a new subgame ∂' of ∂_s is computed, the subgame function g is updated by setting all positions contained in ∂' to the maximal priority p' of ∂' itself. Essentially, g induces a sequence of subgames ordered by game inclusion and indexed by the maximal priorities in each subgame. Consequently, g maps each position to the priority identifying the minimal subgame in the sequence that contains that position. As a result, we have that $\mathrm{Ps}_{\partial'} = \mathsf{g}^{-1}(p')$ and $\mathrm{Ps}_{\partial_s} = \mathrm{dom}\left(\mathsf{g}^{(\leq p)}\right)$. In this new setting, the current priority p of each state can be recovered as the minimal priority in the range of g, while the current subgame corresponds to $\partial_\mathsf{g}^{\leq p} \triangleq \partial \upharpoonright \mathrm{dom}\left(\mathsf{g}^{(\leq p)}\right)$. Formally, a priority function $\mathsf{g} \in \mathbb{G}_\partial \subseteq \mathrm{Ps}_\partial \to \mathrm{Pr}_\partial$ is a *subgame function* iff *(i)* $\mathrm{rng}(\mathsf{g}) \subseteq \mathrm{rng}(\mathsf{pr}_\partial)$ and *(ii)*, for all $p \in \mathrm{rng}(\mathsf{g})$, it holds that *(ii.a)* $\partial_\mathsf{g}^{\leq p} \triangleq \partial \upharpoonright \mathrm{dom}\left(\mathsf{g}^{(\leq p)}\right)$ is a game and *(ii.b)* $\mathsf{g}^{(\leq p)} \subset \mathsf{g}^{(\leq p')}$, for all $p' \in \mathrm{rng}(\mathsf{g})$ with $p < p'$.

A final issue involves the way we keep track of the computed regions. In the PP procedure, we use region function r, which is initialized to pr_∂ and directly stores the stratification of

the regions computed in decreasing order of priority. Therefore, for every priority p, the set $r^{-1}(p)$ exactly corresponds to a region R in $\partial_s = \partial_r^{\leq p}$, where $s \triangleq (r, p)$. If the subgame ∂_s is computed via the subgame function g, instead, a problem with the original definition of the region function r arises: r might contain positions associated with priorities that may not exist in g. Note that this problem is a direct consequence of the strict inclusion $\partial' \subset \partial_s \setminus R$. For this reason, we generalize the notion of region function r to a *quasi-dominion function* q. The regions computed so far can still be recovered by restricting q to the current subgame, identified by g, in symbols $r \triangleq q \cap g$. Formally, given a game ∂, a priority function $q \in \mathbb{Q}_\partial \subseteq \text{Ps}_\partial \to \text{Pr}_\partial$ is a *quasi-dominion function* iff, for all $p \in \text{rng}(q)$, it holds that $q^{-1}(p)$ is an α-quasi dominion in ∂, with $\alpha \triangleq p \bmod 2$. We say that the function $r \triangleq q \cap g$ is a g-*stratified region function* if *(i)* $r^{-1}(p)$ is an α-region in $\partial_g^{\leq p}$, for all $p \in \text{rng}(r)$, and *(ii)* $\text{esc}_\partial^{\overline{\alpha}}(Q) \cap \text{pre}_\partial^{\overline{\alpha}}(X) = \emptyset$, for each $p \in \text{rng}(g)$ and quasi dominion pair $(Q, \alpha) \in \text{QD}_{\partial_g^{\leq p}}$, where $X = \text{Ps}_\partial \setminus \left(\text{dom}\left(g^{(\leq p)} \right) \cup \text{dom}(r) \right)$. Moreover, r is *maximal* w.r.t. g if $r^{-1}(p)$ is α-maximal in $\partial_g^{\leq p}$, for all $p \in \text{rng}(r)$ with $p > \min(\text{rng}(g))$ and $\alpha = p \bmod 2$. Intuitively, Items *(i)* of the g-stratified property and the maximality of r *w.r.t.* g state the same two requirements that a classic PP region function must satisfy. In addition, Item *(ii)* of the former property precisely formalizes the above described requirement on the quasi dominions: every quasi dominion in a subgame identified by the function g cannot have positions with moves that lead outside the domain of the function r.

In the following, we adapt the PP dominion space to support the new subgame decomposition approach. A state is now defined as a pair $s = (q, g)$, where q is a quasi dominion function and g is a subgame function. In addition, the game induced by s is defined as $\partial_s \triangleq \partial_g^{\leq p}$, with $p = \min(\text{rng}(g))$.

Definition 6 (State Space) A *state space* is a tuple $\mathcal{S}_\partial \triangleq \langle \text{S}_\partial, \top_\partial, \prec_\partial \rangle$, where its components are defined as prescribed in the following:

1. $\text{S}_\partial \subseteq \mathbb{Q}_\partial \times \mathbb{G}_\partial$ is the set of all pairs $s \triangleq (q, g)$, called *states*, composed of a quasi dominion function $q \in \mathbb{Q}_\partial$ and a subgame function $g \in \mathbb{G}_\partial$ such that *(a)* $r \triangleq q \cap g$ is maximal *w.r.t.* g, *(b)* r is a g-stratified region function, and *(c)* $q = \text{pr}_\partial \uplus r$;
2. $\top_\partial \triangleq (\text{pr}_\partial, \emptyset[\text{Ps}_\partial \mapsto \text{pr}(\partial)])$;
3. for any two states $s_1 \triangleq (q_1, g_1), s_2 \triangleq (q_2, g_2) \in \text{S}_\partial$, with $p_1 \triangleq \min(\text{rng}(g_1))$ and $p_2 \triangleq \min(\text{rng}(g_2))$, it holds that $s_1 \prec_\partial s_2$ iff either *(a)* there exists a priority $p \in \text{rng}(q_1)$ with $p \geq p_1$ such that *(a.i)* $q_1^{(>p)} = q_2^{(>p)}$ and *(a.ii)* $q_2^{-1}(p) \subset q_1^{-1}(p)$ or *(b)* both *(b.i)* $q_1 = q_2$ and *(b.ii)* $p_1 < p_2$ hold.

As one would expect, the state space defined above is slightly different from the one in Section 4. Indeed, a state (q, g) does not explicitly record the regions anymore. Instead, it records the quasi dominions from which the computed regions can be extracted. In more detail, an α-maximal region is obtained from the derived set $r^{-1}(q) = q^{-1}(q) \cap g^{-1}(q)$, where $q \geq \min(\text{rng}(g))$, $\alpha = q \bmod 2$, and $r \triangleq q \cap g$. Obviously, the current priority $p \triangleq \min(\text{rng}(g))$ of the state must be a priority of an actual region in r. Moreover, the quasi dominion function q, at any priority p' lower than p, must contain positions with the same priority p' in the original priority function pr_∂, *i.e.*, $q^{-1}(p') \subseteq \text{pr}_\partial^{-1}(p')$. At any p' greater that p, instead, q needs to store part of the priority function together with the regions, that is $q^{-1}(p') = (\text{pr}_\partial \uplus r)^{-1}(p')$. All these requirement are ensured by Points *(a)-(c)* of Item 1. Finally, the order follows exactly the same rules as in the original algorithm, where the pair $(q \cap g, \min(\text{rng}(g))$ plays the role of the PP state (r, p).

The compatibility relation is very similar to the one for the PP procedure. Indeed, we only need to specify how to extract the function r from the state s.

Definition 7 (Compatibility Relation) An open quasi dominion pair $(R, \alpha) \in QD_{\partial}^{-}$ is *compatible* with a state $s \triangleq (q, g) \in S_{\partial}$, in symbols $s \succ_{\partial} (R, \alpha)$, iff *(1)* $(R, \alpha) \in Rg_{\partial_s}$ and *(2)* if R is α-open in ∂_s then *(2.a)* R is α-maximal in ∂_s and *(2.b)* $r^{-1}(p) \subseteq R$, where $r \triangleq q \cap g$.

Algorithm 5 provides the implementation for the query function compatible with the generalized priority-promotion mechanism. Let $s \triangleq (q, g)$ be the current state. Line 1 extracts from s the priority p to process, while Line 2 simply computes the associated parity α. Finally, Line 3 computes the attractor *w.r.t.* player α of the quasi dominion contained in q at p that belongs to the subgame ∂_s, *i.e.*, the α-region

---

Algorithm 5: New Query Function.

> **signature** $\Re_{\partial} : S_{\partial} \to (2^{Ps_{\partial}} \times \{0, 1\})$
> **function** $\Re_{\partial}(s)$
>> **let** $(q, g) = s$ **in**
>> 1 $p \leftarrow \min(rng(g))$
>> 2 $\alpha \leftarrow p \bmod 2$
>> 3 $R \leftarrow atr_{\partial_s}^{\alpha}(q^{-1}(p) \cap Ps_{\partial_s})$
>> 4   **return** (R, α)

$q^{-1}(p) \cap Ps_{\partial_s} = q^{-1}(p) \cap g^{-1}(p) = r^{-1}(p)$. The resulting set R is, according to Proposition 2, an α-maximal α-region in ∂_s that contains $r^{-1}(p)$.

Unlike for the classic PP algorithm, the notion of *best escape priority* for player $\overline{\alpha}$ *w.r.t.* an α-region R of the subgame ∂_s is defined *w.r.t.* the quasi dominion function q. Indeed, we do not need to consider the g-stratified region function r, since, due to the property of the subgame function g, the escape positions in every quasi dominion in ∂_s can only reach positions in r. This means that the range of the interface relation I is included in r itself. At this point, the successor function can be defined parametrically on the function F that

---

Algorithm 6: New Successor Function.

> **signature** $\downarrow \, : \, \succ \, \to \Delta \times Pr$
> **function** $s \downarrow (R, \alpha)$
>> **let** $(q, g) = s$ **in**
>> 1   **if** $(R, \alpha) \in Rg_{\partial_s}^{-}$ **then**
>> 2 $q^{\star} \leftarrow q[R \mapsto \min(rng(g))]$
>> 3 $A \leftarrow F(\partial_s \setminus R)$
>> 4 $g^{\star} \leftarrow g[A \mapsto \max(rng(q^{\star} \upharpoonright A))]$
>> **else**
>> 5 $p^{\star} \leftarrow bep_{\partial}^{\overline{\alpha}}(R, q)$
>> 6 $q^{\star} \leftarrow pr_{\partial} \uplus q^{(\geq p^{\star})}[R \mapsto p^{\star}]$
>> 7 $g^{\star} \leftarrow g[dom(g^{(<p^{\star})}) \mapsto p^{\star}]$
>> 8   **return** $(q^{\star}, g^{\star})$

---

identifies a subgame $\partial' \subseteq \partial \setminus R$ such that, for each player α and quasi α-dominion $Q \in QD_{\partial'}$, the set $esc_{\partial'}^{\overline{\alpha}}(Q)$ has no positions that reach $\partial \setminus (R \cup Ps_{\partial'})$. More formally, the set $A \triangleq F(\partial) \subseteq Ps_{\partial}$ must be such that the subgame $\partial \upharpoonright A$ satisfies $esc_{\partial}^{\overline{\alpha}}(Q) \cap pre_{\partial}^{\alpha}(Ps_{\partial} \setminus A) = \emptyset$, for every quasi dominion pair $(Q, \alpha) \in QD_{\partial \upharpoonright A}$. Observe that this requirement exactly mirrors Item *(ii)* in the definition of the g-stratified region function r. As already said before, we propose two instances of such a function. The first one computes the reachability set of an arbitrary position. The second one, instead, simply returns the positions of a minimal SCC. Clearly the escape of every quasi dominion in both the above subgames cannot reach positions outside those subgames.

Algorithm 6 implements the new successor function. Given the current state s and a compatible region pair (R, α) open in the whole input game, it produces a successor state $s^{\star} \triangleq (q^{\star}, g^{\star})$ in the dominion space. In particular, it first checks whether R is open also in the subgame ∂_s (Line 1). If this is the case, it assigns priority $p = \min(rng(g))$ to region R and stores it in the new quasi dominion function $q^{\star}$ (Line 2). The set A of positions of the new subgame is computed via the auxiliary function F (Line 3) and stored in the subgame function $g^{\star}$ at priority $\max(rng(q^{\star} \upharpoonright A))$ (Line 4). If, on the other hand, R is closed in ∂_s,

a promotion is performed. The next priority $p^\star$ is set to the bep of R *w.r.t.* q for player $\overline{\alpha}$ in the entire game $\eth$ (Line 5). Region R is, then, promoted to priority $p^\star$ and all the positions with priorities smaller than $p^\star$ in the current quasi dominion function q are reset (Line 6). Similarly, in the new subgame function $\mathsf{g}^\star$, the priorities of the same positions are set to $p^\star$ (Line 7). Also in this case, the correctness of this last operation follows from Proposition 1.

In conclusion, the priority-promotion mechanism extended with subgame decomposition forms a dominion space, as stated in the following theorem, whose proof is provided in the appendix.

Theorem 4 (Dominion Space) *For a game* $\eth$*, the* structure $\mathcal{D}_\eth \triangleq \langle \eth, \mathcal{S}_\eth, \succ_\eth, \Re_\eth, \downarrow_\eth \rangle$*, where* $\mathcal{S}_\eth$ *is given in Definition 6,* $\succ_\eth$ *is the relation of Definition 7, and* $\Re_\eth$ *and* $\downarrow_\eth$ *are the functions computed by Algorithms 5 and 6 is a dominion space.*

6 Experimental Evaluation

In this section we shall try to assess the effectiveness of the proposed approach. To this end, the priority promotion technique has been implemented in the tool PGSOLVER [31]. The tool, written in OCaml, collects implementations of several parity game solvers proposed in the literature and provides benchmarking tools, which can generate different forms of parity games. The available benchmarks divide into concrete and synthetic problems. The *concrete benchmarks* encode validity and verification problems for temporal logics. They consist in parity games resulting from encodings of the language inclusion problem between automata, specifically a non-deterministic Büchi automaton and a deterministic one, simple reachability problems, the Tower of Hanoi problems, and a fairness verification problem, the Elevator problem (see [31] for more details on this benchmarks). The *synthetic benchmarks* divide into randomly generated games and various families, corresponding to difficult cases, such as clique and ladder-like games, and worst cases for the solvers implemented in PGSOLVER. To fairly compare the different solution techniques used by the underlying algorithms, the solvers involved in the experiments have been isolated from the generic solver implemented in PGSOLVER, which exploits few game transformation and decomposition techniques in the attempt to speed up the solution process. Indeed, those optimizations can, in some cases, solve the game without even calling the selected algorithm, and, in other cases, the resulting overhead can even outweigh the solver time, making the comparison among solvers virtually worthless [31]. Experiments were also conducted with different optimizations enabled and the results exhibit patterns similar to the ones emerging in the following experimental evaluation, even though solution times and the gap among solvers may reduce considerably in some cases, depending on the benchmarks considered.

The algorithms considered in the experimentation are the Zielonka algorithm *Rec* [58], its two dominion decomposition variants, *Dom* [37,38] and *Big* [52], the strategy improvement algorithm *Str* [56][1], the small progress measure algorithm *SmPr* [35], and the one of this article, PP (available at `https://github.com/tcsprojects/pgsolver`).[2]

---

[1] Note that the version of small-progress measure used in the experiments, which is included in the official release of PGSolver, is not the original one proposed by Jurdzinski, as it performs a progress interleaving for both players at once, instead of focusing on a single player like the original algorithm. According to the authors of PGSolver, this optimization allows for a non-trivial performance improvement.

[2] Experiments were carried out on a 64-bit 3.1GHz INTEL® quad-core machine, with i5-2400 processor and 8GB of RAM, running UBUNTU 12.04 with LINUX kernel version 3.2.0. PGSOLVER was compiled with OCaml version 2.12.1.

Benchmark	Positions	Dom	Big	Str	Rec	SmPr	PP
Hanoi	$7 \cdot 10^5$	1.4	1.4	†	1.8	†	**0.7**
Hanoi	$2.1 \cdot 10^6$	4.4	4.4	†	5.7	†	**2.3**
Hanoi	$6.3 \cdot 10^6$	21.4	21.4	‡	17.4	†	**7.0**
Elevator	$1 \cdot 10^5$	1.0	1.0	†	0.8	12.8	**0.2**
Elevator	$8.6 \cdot 10^5$	10.86	10.86	†	7.0	155.5	**2.0**
Elevator	$7.7 \cdot 10^6$	†	‡	‡	‡	†	**19.8**
Lang. Incl.	$3.7 \cdot 10^5$	6.6	6.6	†	3.0	78.4	**0.6**
Lang. Incl.	$1.6 \cdot 10^6$	51.0	51.3	†	26.3	†	**3.6**
Lang. Incl.	$5 \cdot 10^6$	†	‡	‡	145.5	‡	**16.5**
Ladder	$4 \cdot 10^6$	†	‡	‡	35.0	130.5	**7.9**
Str. Imp.	$4.5 \cdot 10^6$	81.0	82.8	†	71.0	‡	**57.0**
Clique	$8 \cdot 10^3$	†	‡	†	†	†	**10.8**
MC. Lad.	$7.5 \cdot 10^6$	†	‡	‡	**4.3**	‡	**4.3**
Rec. Lad.	$5 \cdot 10^4$	†	‡	†	‡	‡	**62.8**
Jurdziński	$4 \cdot 10^4$	†	†	188.2	†	93.2	**69.6**

Table 2: Execution times (in seconds) on several benchmark families. Time out (†) is set to 600 seconds and memory out (‡) to 7.5Gb.

6.1 Special Families.

Table 2 displays the results of all the solvers involved on the benchmark families available in PGSOLVER. We only report on the biggest instances we could deal with, given the available computational resources[3]. The parameter Positions refers to the number of positions in the games and the best results are emphasized in bold. The first three sections of the table, each comprising three instances of increasing size for the same benchmark, consider the concrete verification problems mentioned above. On all the instances of the Tower of Hanoi problem most of the solvers perform reasonably well, except for *SmPr* and for *Str*. The Elevator problem, instead, proved to be very demanding, both in terms of time and memory, for all the solvers, except for our new algorithm and for *Dom*, which, however, could not solve the biggest instance within the time limit of 10 minutes. Our solver performs extremely well on both this benchmarks and on the instances of the Language Inclusion problem, whose biggest instance could be solved only by *Rec* among the other solvers. On the worst case benchmarks, PP performs quite well also on Ladder, Strategy Improvement, Clique, and Jurdziński games, all of which proved to be considerably difficult for all the other solvers. The Modelchecker game is a tie with *Rec*. On the Recursive Ladder game PP is the only algorithm that could solve the instance used in the experiments. However, the even instances of that game, unlike the odd ones, turn out to be very easy to solve by *Str*, which requires less than a second even for the instance with $5 \cdot 10^5$ positions.

The reason is that those instances are completely won by player odd and, in addition, since *Str* looks for a winning strategy of player even, it immediately discovers that the initial even strategy cannot be improved further. On the other hand, the odd instances are completely won by player even and the algorithm requires a linear number of iterations to find the winning strategy. The new solver exhibits the most consistent behavior overall on these benchmarks, significantly outperforming the other solvers considered in the experiments. Moreover, for all of them the priority promotion algorithm requires no promotions regardless

---

[3] The biggest instances in the table are generated by the following PGSOLVERcommands:
`towersofhanoi 13`, `elevatorgame 8`, `langincl 500 100`, `laddergame 4000000`, `stratimprgen -pg friedmannsubexp 1000`, `modelcheckerladder 2500000`, `cliquegame 8000`, `recursiveladder 10001`, and `jurdzinskigame 100 100`.

of the input parameters, except for the Elevator problem, on which it performs only two. As a consequence, PP requires polynomial time on all the benchmarks reported in the table.

Table 3, instead, reports a comparison with the quasi-polynomial time algorithms proposed in [24] and [36]. We considered a preliminary C++ implementation of the former, *QPT*, made available by the authors themselves (the iOS executable can be obtained at `https://cgi.csc.liv.ac.uk/~dominik/parity`), and an OCaml implementation of the latter, *SPM*, recently included in PG-SOLVER[4]. Despite the better worst case upper bound and the theoretical relevance of the result, neither algorithm could solve the benchmark instances reported in Table 2. Hence, we had to resort to much smaller instances. The comparison may not be particularly meaningful, given the differences in implementation, neither fair towards the solvers implemented in OCaml. It suggests, nonetheless, that the practical effectiveness of the two quasi-polynomial

Benchmark	Positions	QPT	SPM	PP
Hanoi	$1.9 \cdot 10^4$	7.7	†	**0.0**
Hanoi	$5.9 \cdot 10^4$	72.5	†	**0.0**
Elevator	$2.7 \cdot 10^3$	13.0	†	**0.0**
Elevator	$1.5 \cdot 10^4$	342.3	†	**0.0**
Elevator	$1 \cdot 10^5$	†	†	**0.2**
Lang. Incl.	$4.4 \cdot 10^4$	25.7	†	**0.0**
Lang. Incl.	$7.8 \cdot 10^4$	71.8	†	**0.0**
Lang. Incl.	$1.2 \cdot 10^5$	149.3	†	**0.1**
Ladder	$1 \cdot 10^4$	121.1	†	**0.0**
Ladder	$2 \cdot 10^4$	501.5	†	**0.0**
Str. Imp.	$1.2 \cdot 10^4$	3.7	†	**0.0**
Str. Imp.	$4.6 \cdot 10^4$	42.6	†	**0.0**
Str. Imp.	$1 \cdot 10^5$	227.9	†	**0.1**
Clique	$2 \cdot 10^2$	0.8	†	**0.0**
Clique	$1 \cdot 10^3$	102.1	†	**0.1**
Clique	$1.5 \cdot 10^3$	518.6	†	**0.4**
MC. Lad.	$3 \cdot 10^4$	9.2	†	**0.0**
MC. Lad.	$1.5 \cdot 10^5$	254.5	†	**0.0**
Rec. Lad.	$5 \cdot 10^3$	13.1	†	**0.5**
Rec. Lad.	$2.5 \cdot 10^4$	288.2	†	**13.8**
Jurdziński	$7.5 \cdot 10^3$	54.6	†	**3.4**
Jurdziński	$1.2 \cdot 10^4$	89.6	†	**9.6**
Jurdziński	$1.9 \cdot 10^4$	320.4	†	**25.41**

Table 3: Execution times (in seconds) on several benchmark families. Time out (†) is set to 600 seconds.

solvers, at least of the versions available at the time of writing, is very far from that of the exponential algorithms and, in particular, of PP. The table shows indeed a considerable gap, often of many orders of magnitude, which may only to some extent be explained by a possible lack of optimization.
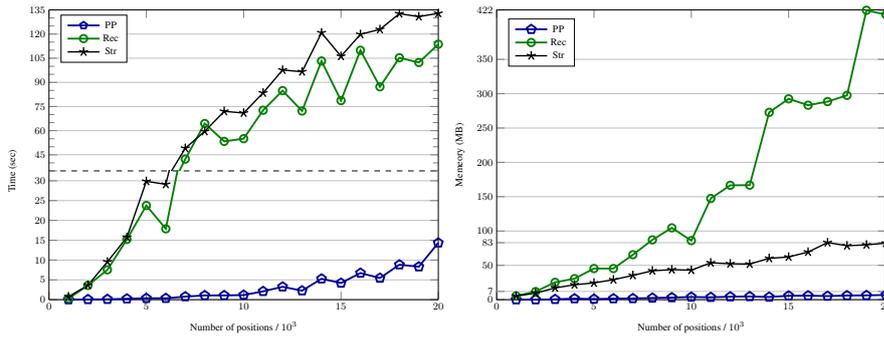
6.2 Random Games.

Figure 4 compares the running times (left-hand side) and memory requirements (right-end side) of the new algorithm PP against *Rec* and *Str* on 2000 random games of size ranging from 5000 to 20000 positions and 2 outgoing moves per position. Interestingly, these random games proved to be quite challenging for all the considered solvers. We set a time-out to 180 seconds (3 minutes). Both *Dom* and *Big* perform quite poorly on those games, hitting the time-out already for very small instances, and we decided to leave them out of the picture. The behavior of the solvers is typically highly variable even on games of the same size and priorities. To summaries the results, the average running time on clusters of games seemed the most appropriate choice in this case. Therefore, each point in the graph shows the average time over a cluster of 100 different games of the same size: for each size value n, we chose a number $k = n \cdot i / 10$ of priorities, with $i \in [1, 10]$, and 10 random games were generated

---

[4] At the time of writing, the *SPM* implementation is only provided in the development version of PGSOLVER, currently available at `https://github.com/tcsprojects/pgsolver/tree/develop`.

Fig. 4: Time and auxiliary memory on random games with 2 moves per position.

for each pair of n and k. The new algorithm perform significantly better than the others on those games. The right-hand side graph also shows that the theoretical improvement on the auxiliary memory requirements of the new algorithm has a considerable practical impact on memory consumption compared to the other solvers. In these particular benchmarks, the additional memory required by PP amounts to 7 Mb on average for the biggest instances, while *Str* requires up to 83 Mb and *Rec* up to 422 Mb.

| *Positions*	Dom	Big	Str	Rec	SmPr	PP
$1 \cdot 10^4$	0.66	60.00	24.57	0.43	51.35	**0.12**
$3 \cdot 10^4$	3.29	60.00	55.06	1.54	60.00	**0.46**
$5 \cdot 10^4$	7.43	60.00	59.70	3.00	60.00	**0.97**
$7 \cdot 10^4$	12.85	60.00	60.00	4.54	60.00	**1.50**
$1 \cdot 10^5$	19.56	60.00	60.00	7.62	60.00	**2.28**
Time/Mem-out	1%	100%	74%	**0%**	92%	**0%**

Table 4: Average execution time (in seconds) on random games with 10 and 100 moves per position. Time out is set to 60 seconds.

Table 4 reports on some experiments on random games with higher number of moves per position. The resulting games turn out to be much easier to solve for most the solvers, in particular for *Rec*, *Dom*, and PP. The benchmarks here are divided into 5 cluster with increasing number of positions from 10^4 to 10^5. Each cluster contains 120 games each, with number of priorities linear on the number of positions and 10 to 100 moves per position. On each row, the table reports the average time required for the solution over the 120 games of each cluster. The last row, instead, details the percentage of games on which the corresponding solver could not terminate before the timeout set to 60 seconds. The higher number of moves significantly increases the dimension of the regions computed by PP and, consequently, also the chances for it to find a closed one. Indeed, the number of promotions required by the algorithm on all those games is typically below half a dozen. The whole solution time is almost exclusively due to a very limited number of attractors needed to compute the few regions contained in the games. The only other solvers that can easily solve all these games within 60 seconds are *Rec*, whose performance is only slightly worse than that of PP, and *Dom*, which could solve almost all of the games before the timeout.

Table 5 shows an experimental comparison on random games where the number of priorities grows logarithmically *w.r.t.* the number of positions. The relevance of these benchmarks stems from the fact that in practical applications, such as verification problems, the resulting encodings into parity games produce games with low number of priorities *w.r.t.* the number

$Positions$	Dom	Big	Str	Rec	SmPr	PP
$1 \cdot 10^4$	60.00	15.22	55.24	0.65	60.00	**0.14**
$3 \cdot 10^4$	60.00	35.31	59.37	2.94	60.00	**0.40**
$5 \cdot 10^4$	60.00	45.04	60.00	5.79	60.00	**1.39**
$7 \cdot 10^4$	60.00	50.67	60.00	5.12	59.70	**0.79**
$1 \cdot 10^5$	60.00	58.43	60.00	8.89	60.00	**1.83**
Timeout	100%	49%	95.4%	1%	99.67%	**0%**

Table 5: Average execution time (in seconds) on random games with logarithmic number of priorities. Time out is set to 60 seconds.
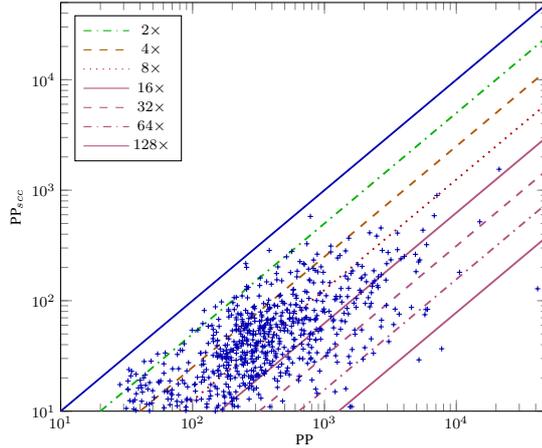
Fig. 5: Comparison between PP and PP_{scc} on random games with 50000 positions.

of positions. This number is usually connected to a measure of complexity of the temporal formula to verify, *e.g.*, the alternation depth the fixpoint operators of the μ-calculus. In many cases, this number is bounded from above by a logarithmic function of the total number of positions in the game. As in the previous table, these benchmarks are divided into five clusters, according the number n of positions. Each cluster, in turn, contains 100 games with two moves per position and with number of priorities varying from $5 \log_2 n$ to $20 \log_2 n$. The results show that PP perform much better than the other solvers, being able to solve all of the benchmarks within the timeout, with *Rec* coming a close second.

6.3 PP with subgame decomposition.

Finally, Figure 5 compares the performance of the subgame decomposition version of PP with the original algorithm. Here, the decomposition schema has been realized by instantiating function F in Algorithm 6 with an SCC decomposition procedure that computes some minimal SCC in the current subgame. We refer to resulting procedure with PP_{scc}. In order to assess its effectiveness, we used a new pool of benchmarks that contains 740 games, each with 50000 positions, 2 moves per positions and priorities varying from 8000 to 12000. These games are much harder than the ones reported in Figure 4 and have been specifically selected among random games whose solution requires PP more 30 seconds and up to more than 10 hours to be solved. The results, drawn on a logarithmic scale, show that the decomposition technique, described in Section 6, does pay off significantly. It can run up to 128 times faster

than the original version, reducing the solution time from several hours to a few seconds in some cases.

Very similar results can be obtained by instantiating function F in Algorithm 6 with a simpler reachability computation procedure that collects all the positions reachable from some arbitrary position in the current game. The resulting technique, called PP_{rch}, cannot, in general, provide as deep a decomposition as does the SCC-based decomposition. The additional overhead of the latter one is, however, usually higher. As shown in Figure 6, the two algorithm are not comparable, in the sense that there is no clear winner between the two: on games with more structure the SCC decomposition usually performs better, being able to compute smaller subgames; on games with less structure, the reachability decomposition technique obtains sufficiently small subgames with less overhead.

7 Discussion

This article considers the problem of solving *Parity Games*, a special form of infinite-duration games over graphs having relevant applications in various branches of Theoretical Computer Science. We propose a novel solution technique, based on a *priority-promotion mechanism*. Following on this approach, a new solution algorithm is presented and studied. We gave proofs of its correctness and provided an accurate analysis of its time and space complexities.

As far as time complexity is concerned, an exponential upper bound in the number of priorities is given. A lower bound for the worst-case is also presented in the form of a family of parity games on which the new technique exhibits an exponential behavior. On the bright side, the new solution exhibits the best space complexity among the currently known algorithms for parity games. In fact, we showed that the maximal additional space needed to solve a parity game is linear in the number of positions, logarithmic in the number of priorities, and independent of the number of moves in the game. This is an important result, in particular considering that in practical applications we often need to deal with games having
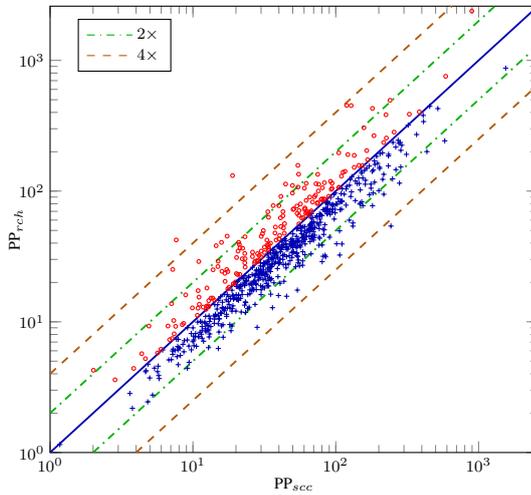
Fig. 6: Comparison between PP_{scc} and PP_{rch} on random games with 50000 positions.

a very high number of positions, moves, and, in some cases, priorities. Therefore, low space requirements are essential for practical scalability.

To assess the effectiveness of the new approach, experiments are conducted against concrete and synthetic problems. We compare the new algorithm with the state-of-the-art solvers implemented in PGSOLVER. The results are very promising, showing that the proposed approach is extremely effective in practice, often substantially better than existing ones, including the quasi-polynomial algorithms recently proposed in the literature. We also devise a quite general priority-promotion-based schema that can be instantiated with various game decomposition techniques, such as SCC-decomposition. Experiments comparing the original PP algorithm with two such instantiations reveal that a deep coupling of PP with decomposition techniques can be extremely helpful in cutting down solution time, often of several orders of magnitude.

These results suggest that the new approach is worth pursuing further. Therefore, we are currently investigating new and finer priority-promotion policies that try to minimize the number of region resets after a priority promotion.

It would also be interesting to investigate the applicability of the priority promotion approach to related problems, such as *prompt-parity games* [46] and similar conditions [14, 29, 30], and even in wider contexts like *mean-payoff games* [18, 15] and *energy games* [12, 13].

A Proof of Theorem 1

To prove Theorem 1, we have to show that the three components $\mathcal{S}_\supset$, $\Re_\supset$, and $\downarrow_\supset$ of the structure $\mathcal{D}_\supset$ satisfy the properties required by Definition 2 of dominion space. We do this through the Lemmas 1, 2, and 3.

Lemma 1 (State Space) *The* PP *state space* $\mathcal{S}_\supset = \langle S_\supset, \top_\supset, \prec_\supset \rangle$ *for a game* $\supset \in \mathcal{P}$ *is a well-founded partial order* w.r.t. $\prec_\supset$ *with designated element* $\top_\supset \in S_\supset$.

Proof Since $S_\supset$ is a finite set, to show that $\prec_\supset$ is a well-founded partial order on $S_\supset$, it is enough to prove that it is simply a strict partial order on the same set, *i.e.*, an irreflexive and transitive relation.

For the irreflexive property, by Item 3 of Definition 4, it is immediate to see that $s \nprec_\supset s$, for all states $s \triangleq (r, p) \in S_\supset$, since neither there exists a priority $q \in \mathsf{rng}(r)$ such that $r^{-1}(q) \subset r^{-1}(q)$ nor $p < p$.

For the transitive property, instead, consider three states $s_1 \triangleq (r_1, p_1), s_2 \triangleq (r_2, p_2), s_3 \triangleq (r_3, p_3) \in S_\supset$ for which $s_1 \prec_\supset s_2$ and $s_2 \prec_\supset s_3$ hold. Due to Items 3.a and 3.b of the same definition, four cases may arise.

- Item 3.a for both $s_1 \prec_\supset s_2$ and $s_2 \prec_\supset s_3$: there exist two priorities $q_1 \in \mathsf{rng}(r_1)$ and $q_2 \in \mathsf{rng}(r_2)$ with $q_1 \geq p_1$ such that $r_1{}^{(>q_1)} = r_2{}^{(>q_1)}$, $r_2{}^{(>q_2)} = r_3{}^{(>q_2)}$, $r_2^{-1}(q_1) \subset r_1^{-1}(q_1)$, and $r_3^{-1}(q_2) \subset r_2^{-1}(q_2)$. Let $q \triangleq \max\{q_1, q_2\} \geq p_1$. If $q = q_1 = q_2$ then $r_1{}^{(>q)} = r_2{}^{(>q)} = r_3{}^{(>q)}$ and $r_3^{-1}(q) \subset r_2^{-1}(q) \subset r_1^{-1}(q)$. If $q = q_1 > q_2$ then $r_1{}^{(>q)} = r_2{}^{(>q)} = (r_2{}^{(>q_2)})^{(>q)} = (r_3{}^{(>q_2)})^{(>q)} = r_3{}^{(>q)}$ and $r_3^{-1}(q) = r_2^{-1}(q) \subset r_1^{-1}(q)$. Finally, if $q = q_2 > q_1$ then $r_1{}^{(>q)} = (r_1{}^{(>q_1)})^{(>q)} = (r_2{}^{(>q_1)})^{(>q)} = r_2{}^{(>q)} = r_3{}^{(>q)}$ and $r_3^{-1}(q) \subset r_2^{-1}(q) = r_1^{-1}(q)$. Moreover, $q \in \mathsf{rng}(r_2{}^{(>q_1)}) = \mathsf{rng}(r_1{}^{(>q_1)}) \subseteq \mathsf{rng}(r_1)$. Summing up, it holds that $s_1 \prec_\supset s_3$.
- Item 3.a for $s_1 \prec_\supset s_2$ and Item 3.b for $s_2 \prec_\supset s_3$: there exists a priority $q \in \mathsf{rng}(r_1)$ with $q \geq p_1$ such that $r_1{}^{(>q)} = r_2{}^{(>q)}$ and $r_2^{-1}(q) \subset r_1^{-1}(q)$; moreover, $r_2 = r_3$. Thus, $r_1{}^{(>q)} = r_2{}^{(>q)} = r_3{}^{(>q)}$ and $r_3^{-1}(q) = r_2^{-1}(q) \subset r_1^{-1}(q)$. Consequently, $s_1 \prec_\supset s_3$.
- Item 3.a for $s_2 \prec_\supset s_3$ and Item 3.b for $s_1 \prec_\supset s_2$: there exists a priority $q \in \mathsf{rng}(r_2)$ with $q \geq p_2$ such that $r_2{}^{(>q)} = r_3{}^{(>q)}$ and $r_3^{-1}(q) \subset r_2^{-1}(q)$; moreover, $r_1 = r_2$ and $p_1 < p_2$. Thus, $r_1{}^{(>q)} = r_2{}^{(>q)} = r_3{}^{(>q)}$, $r_3^{-1}(q) \subset r_2^{-1}(q) = r_1^{-1}(q)$, $q \in \mathsf{rng}(r_1)$, and $q > p_1$. Consequently, $s_1 \prec_\supset s_3$.
- Item 3.b for both $s_1 \prec_\supset s_2$ and $s_2 \prec_\supset s_3$: $r_1 = r_2$, $r_2 = r_3$, $p_1 < p_2$, and $p_2 < p_3$. Hence, $r_1 = r_3$ and $p_1 < p_3$, which implies that $s_1 \prec_\supset s_3$ also in this case.

To complete the proof, we need to show that $\top_\supset \triangleq (\mathsf{pr}_\supset, \mathsf{pr}(\supset))$ belongs to $S_\supset$. Indeed, r is vacuously maximal above p, so Item 1.a holds. Item 1.c is trivially verified, since $r = \mathsf{pr}_\supset$. Obviously, Item 1.b also follows from the fact that $p = \mathsf{pr}(\supset) = \max(\mathsf{rng}(\mathsf{pr}_\supset)) \in \mathsf{rng}(\mathsf{pr}_\supset) = \mathsf{rng}(r)$. Finally, $\mathsf{pr}_\supset$ is a region

function as, for all priorities $q \in \mathsf{rng}(\mathsf{pr}_\partial)$ with $\alpha \triangleq q \bmod 2$, it holds that $\mathsf{pr}_\partial^{-1}(q) \cap \mathrm{Ps}_{\partial_{\mathsf{pr}_\partial}^{\leq q}}$ is an α-region in the subgame $\partial_{\mathsf{pr}_\partial}^{\leq q}$, if non-empty. Indeed, the set $\mathsf{pr}_\partial^{-1}(q) \cap \mathrm{Ps}_{\partial_{\mathsf{pr}_\partial}^{\leq q}}$ can only contain positions of priority p, which is the maximal one in the corresponding subgame. Therefore, player α has an obvious strategy that forces every infinite play inside this set to be winning for it. $\qquad\square$

Lemma 2 (Query Function) *The function $\Re_\partial$ is a query function,* i.e., *for all states $s \in \mathrm{S}_\partial$, it holds that* (1) $\Re_\partial(s) \in \mathrm{QD}_\partial$ *and* (2) *if* $\Re_\partial(s) \in \mathrm{QD}_\partial^-$ *then* $s \succ_\partial \Re_\partial(s)$.

Proof Let $s \triangleq (\mathsf{r}, p) \in \mathrm{S}_\partial$ be a state and $(\mathrm{R}, \alpha) \triangleq \Re_\partial(s)$ the pair of a set of positions $\mathrm{R} \subseteq \mathrm{Ps}_\partial$ and a player $\alpha \in \{0, 1\}$ obtained by computing the function $\Re_\partial$ on s. Due to Line 1 of Algorithm 3, it follows that $\alpha \equiv_2 p$. By Item 1.b, it holds that $p \in \mathsf{rng}(\mathsf{r})$. Thus, by Item 1.a and the definition of region function, we have that $\mathsf{r}^{-1}(p)$ is an α-region in $\partial_\mathsf{r}^{\leq p}$, the latter being equal to ∂_s. Now, by Line 2 of Algorithm 3 and Proposition 2, where we assume $\mathrm{R}^\star \triangleq \mathsf{r}^{-1}(p)$ and $\partial \triangleq \partial_s$, it follows that $\mathrm{R} \supseteq \mathsf{r}^{-1}(p)$ is an α-region in ∂_s, i.e., $(\mathrm{R}, \alpha) \in \mathrm{Rg}_{\partial_s}$, and, so a quasi dominion in ∂, i.e., $(\mathrm{R}, \alpha) \in \mathrm{QD}_\partial$. In addition, R is α-maximal in ∂_s. Consequently, $s \succ_\partial (\mathrm{R}, \alpha)$, since all requirements of Definition 5 are satisfied. $\qquad\square$

Lemma 3 (Successor Function) *The function $\downarrow_\partial$ is an successor function,* i.e., *for all states $s \in \mathrm{S}_\partial$ and region pairs $(\mathrm{R}, \alpha) \in \mathrm{Rg}_\partial^-$ with $s \succ_\partial (\mathrm{R}, \alpha)$, it holds that* (1) $s \downarrow_\partial (\mathrm{R}, \alpha) \in \mathrm{S}_\partial$ *and* (2) $s \downarrow_\partial (\mathrm{R}, \alpha) \prec_\partial s$.

Proof Let $s \triangleq (\mathsf{r}, p) \in \mathrm{S}_\partial$ be a state, $(\mathrm{R}, \alpha) \in \mathrm{Rg}_\partial^-$ an open region pair in ∂ compatible with s, and $s^\star = (\mathsf{r}^\star, p^\star) \triangleq s \downarrow_\partial (\mathrm{R}, \alpha)$ the result obtained by computing the function $\downarrow_\partial$ on s and (R, α). Due to Item 1 of Definition 4, we have that *(1)* r is maximal above p and *(2)* $\mathsf{r}^{(<p)} \subseteq \mathsf{pr}_\partial^{(<p)}$. Moreover, by Item 1 of Definition 5, it holds that $\mathrm{R} \subseteq \mathrm{Ps}_{\partial_s}$, which implies *(3)* $\mathsf{dom}(\mathsf{r}^{(>p)}) \cap \mathrm{R} = \emptyset$.

On the one hand, suppose that R is α-open in ∂_s, i.e., $(\mathrm{R}, \alpha) \in \mathrm{Rg}_{\partial_s}^-$. By Lines 1-3 of Algorithm 4, we have that *(4)* $\mathsf{r}^\star = \mathsf{r}[\mathrm{R} \mapsto p]$ and *(5)* $p^\star = \max(\mathsf{rng}(\mathsf{r}^{\star(<p)}))$. In addition, by Item 2 of Definition 5, it holds that *(6)* R is α-maximal in ∂_s and *(7)* $\mathsf{r}^{-1}(p) \subseteq \mathrm{R}$. Now, by Point *(5)*, it immediately follows that *(8)* $p^\star \in \mathsf{rng}(\mathsf{r}^\star)$, *(9)* $p^\star < p$, and *(10)* there is no priority $q \in \mathsf{rng}(\mathsf{r}^\star)$ such that $p^\star < q < p$. Also, by Points *(4)*, *(3)* and *(7)*, we have that $\mathsf{r}^{\star(>p)} = \mathsf{r}^{(>p)}$, $\mathsf{r}^{\star-1}(p) = \mathrm{R}$, and $\mathsf{r}^{\star(<p)} \subseteq \mathsf{r}^{(<p)}$. Thus, by Points *(1)*, *(6)*, *(9)*, and *(10)*, it holds that *(11)* $\mathsf{r}^\star$ is maximal above $p^\star$. Moreover, by Points *(2)* and *(9)*, we derive that *(12)* $\mathsf{r}^{\star(<p^\star)} \subseteq \mathsf{pr}_\partial^{(<p^\star)}$. Finally, to prove that *(13)* $\mathsf{r}^\star$ is a region function, we need to show that $\mathsf{r}^{\star-1}(q) \cap \mathrm{Ps}_{\partial_{\mathsf{r}^\star}^{\leq q}}$, if non-empty, is a β-region in the subgame $\partial_{\mathsf{r}^\star}^{\leq q}$, for all priorities $q \in \mathsf{rng}(\mathsf{r}^\star)$ with $\beta \triangleq q \bmod 2$. Indeed, if $q > p$, by Point *(1)*, it holds that $\mathsf{r}^{\star-1}(q) \cap \mathrm{Ps}_{\partial_{\mathsf{r}^\star}^{\leq q}} = \mathsf{r}^{-1}(q) \neq \emptyset$. Hence, the property follows directly from the fact that r is a region function. If $q = p$, again by Point *(1)*, we have that $\mathsf{r}^{\star-1}(q) \cap \mathrm{Ps}_{\partial_{\mathsf{r}^\star}^{\leq q}} = \mathrm{R} \neq \emptyset$. So, the property is ensured by the hypothesis that R is an α-region. For the last case $q < p$, the property follows by Point *(12)* and the observation that, if non-empty, the set $\mathsf{r}^{\star-1}(q) \cap \mathrm{Ps}_{\partial_{\mathsf{r}^\star}^{\leq q}} \subseteq \mathsf{pr}_\partial^{-1}(q)$ only contains positions of priority q that necessarily form a region of the corresponding parity. Summing up, Points *(13)*, *(11)*, *(8)*, and *(12)* ensure that $(\mathsf{r}^\star, p^\star) \in \mathrm{S}_\partial$. At this point, it remains just to show that $(\mathsf{r}^\star, p^\star) \prec_\partial (\mathsf{r}, p)$. By Point *(7)*, two cases may arise. If $\mathsf{r}^{-1}(p) \subset \mathrm{R}$, the thesis follows from Item 3.a of Definition 4, where the priority q is set to p. On the contrary, if $\mathrm{R} = \mathsf{r}^{-1}(p)$, by Point *(4)*, we have that $\mathsf{r}^\star = \mathsf{r}$. Therefore, due to Point *(9)*, the thesis is derived from Item 3.b of the same definition.

On the other hand, suppose that R is α-closed in ∂_s, i.e., $(\mathrm{R}, \alpha) \notin \mathrm{Rg}_{\partial_s}^-$. By Lines 1, 4, and 5 of Algorithm 4, we have that *(14)* $p^\star = \mathsf{bep}_\partial^{\overline{\alpha}}(\mathrm{R}, \mathsf{r})$ and *(15)* $\mathsf{r}^\star = \mathsf{pr}_\partial \uplus \mathsf{r}^{(\geq p^\star)}[\mathrm{R} \mapsto p^\star]$. It is not hard to see that, by Points *(14)* and *(7)*, the definition of bep and the fact that R is closed, we have that *(16)* $p^\star > p$. Now, by Points *(15)* and *(3)*, it follows that $\mathsf{r}^{\star(>p^\star)} = \mathsf{r}^{(>p^\star)}$, $\mathsf{r}^{\star-1}(p^\star) = \mathsf{r}^{-1}(p^\star) \cup \mathrm{R}$, and *(18)* $\mathsf{r}^{\star(<p^\star)} \subseteq \mathsf{pr}_\partial^{(<p^\star)}$. Consequently, *(19)* $p^\star \in \mathsf{rng}(\mathsf{r}^\star)$. Moreover, due to Points *(1)* and *(16)*, we have that *(20)* $\mathsf{r}^\star$ is maximal above $p^\star$. Finally, the proof that *(21)* $\mathsf{r}^\star$ is a region function easily follows by observing that $\mathsf{r}^{-1}(p^\star) \cup \mathrm{R}$ is an α-region, due to Proposition 1. Summing up, Points *(21)*, *(20)*, *(19)*, and *(18)* ensure that $(\mathsf{r}^\star, p^\star) \in \mathrm{S}_\partial$. At this point, as for the previous case, it remains to show that $(\mathsf{r}^\star, p^\star) \prec_\partial (\mathsf{r}, p)$. This fact easily follows from Item 3.a of Definition 4, where the priority q is set to $p^\star$, since, by Points *(3)* and *(16)*, we have that $\mathsf{r}^{-1}(p^\star) \cap \mathrm{R} = \emptyset$, so $\mathsf{r}^{-1}(p^\star) \subset \mathsf{r}^{-1}(p^\star) \cup \mathrm{R} = \mathsf{r}^{\star-1}(p^\star)$. $\qquad\square$

B Proof of Theorem 4

Similarly to the previous appendix, to prove Theorem 4, we have to show that the three components $\mathcal{S}_\partial$, $\Re_\partial$, and $\downarrow_\partial$ of the structure $\mathcal{D}_\partial$ satisfy the properties required by Definition 6 of dominion space. We do this through the Lemmas 4, 5, and 6.

Lemma 4 (State Space) *The generalized* PP *state space* $\mathcal{S}_\partial = \langle S_\partial, \top_\partial, \prec_\partial \rangle$ *for a game* $\partial \in \mathcal{P}$ *is a well-founded partial order* w.r.t. $\prec_\partial$ *with designated element* $\top_\partial \in S_\partial$.

Proof To prove that $\prec_\partial$ is a strict partial order on S_∂, we can follow the same approach used in the proof of Theorem 1, where the pair $(\mathsf{q} \cap \mathsf{g}, \min(\mathsf{rng}(\mathsf{g})))$ plays the same role of the state (r, p) in the PP algorithm. Therefore, to complete the proof, we only need to show that $\top_\partial \triangleq (\mathsf{pr}_\partial, \varnothing[\mathrm{Ps}_\partial \mapsto \mathsf{pr}(\partial)])$ belongs to S_∂. The function $\mathsf{q} \triangleq \mathsf{pr}_\partial$ is trivially a quasi dominion function, since $\mathsf{pr}_\partial^{-1}(p)$ is an α-quasi dominion in the original game ∂ as it only contains positions of the same priority p of parity α, for all $p \in \mathsf{rng}(\mathsf{pr}_\partial)$ with $\alpha \triangleq p \bmod 2$. To prove that $\mathsf{g} \triangleq \varnothing[\mathrm{Ps}_\partial \mapsto \mathsf{pr}(\partial)]$ is a subgame function, first observe that Item *(i)* of the corresponding definition holds due to the fact that $\mathsf{rng}(\mathsf{g}) = \{\mathsf{pr}(\partial)\}$ and $\mathsf{pr}(\partial) \in \mathsf{rng}(\mathsf{pr}_\partial)$. Moreover, $\mathsf{dom}\big(\mathsf{g}^{(\leq \mathsf{pr}(\partial))}\big) = \mathrm{Ps}_\partial$, so, $\partial_\mathsf{g}^{\leq \mathsf{pr}(\partial)} = \partial$ is clearly a game. Consequently, Item *(ii.a)* holds as well. Finally, Item *(ii.b)* is vacuously verified, since there is no priority $p \in \mathsf{rng}(\mathsf{g})$ with $p > \mathsf{pr}(\partial)$. We can now consider the three points in Item 1 of Definition 6. Point *(a)* vacuously holds, since, as already observed, there is no priority in g greater than $\mathsf{pr}(\partial)$. Now, let $\mathsf{r} \triangleq \mathsf{q} \cap \mathsf{g}$. First notice that $\mathsf{r}^{-1}(\mathsf{pr}(\partial)) = \mathsf{pr}_\partial^{-1}(\mathsf{pr}(\partial))$ is a region in $\partial_{\top_\partial} = \partial_\mathsf{g}^{\leq \mathsf{pr}(\partial)} = \partial$, since it is a quasi dominion whose escape positions have the maximal priority in the game. Moreover, $\mathrm{X} = \mathrm{Ps}_\partial \setminus (\mathsf{dom}\big(\mathsf{g}^{(\leq \mathsf{pr}(\partial))}\big) \cup \mathsf{dom}(\mathsf{r})) = \mathrm{Ps}_\partial \setminus \mathrm{Ps}_\partial = \emptyset$. Consequence, Point *(b)* also holds. Finally, Point *(c)* is implied by the fact that $\mathsf{pr}_\partial \uplus \mathsf{r} = \mathsf{pr}_\partial \uplus (\mathsf{q} \cap \mathsf{g}) = \mathsf{pr}_\partial \uplus \mathsf{q} = \mathsf{pr}_\partial \uplus \mathsf{pr}_\partial = \mathsf{pr}_\partial = \mathsf{q}$. $\square$

Lemma 5 (Query Function) *The function* $\Re_\partial$ *is a query function,* i.e., *for all states* $s \in S_\partial$, *it holds that* (1) $\Re_\partial(s) \in \mathrm{QD}_\partial$ *and* (2) *if* $\Re_\partial(s) \in \mathrm{QD}_\partial^-$ *then* $s \succ_\partial \Re_\partial(s)$.

Proof Let $s \triangleq (\mathsf{q}, \mathsf{g}) \in S_\partial$ be a state and $(\mathrm{R}, \alpha) \triangleq \Re_\partial(s)$ the pair of a set of positions $\mathrm{R} \subseteq \mathrm{Ps}_\partial$ and a player $\alpha \in \{0, 1\}$ obtained by computing the function $\Re_\partial$ on s. Moreover, assume $\mathsf{r} \triangleq \mathsf{q} \cap \mathsf{g}$. First observe that Item 1.c implies $\mathsf{rng}(\mathsf{g}) \subseteq \mathsf{rng}(\mathsf{q})$. Indeed, suppose by contradiction that there exists a priority $p \in \mathsf{rng}(\mathsf{g}) \setminus \mathsf{rng}(\mathsf{q})$. Clearly, $p \notin \mathsf{rng}(\mathsf{r})$. Moreover, $p \in \mathsf{rng}(\mathsf{g}) \subseteq \mathsf{rng}(\mathsf{pr}_\partial)$, due to Item *(i)* of the definition of subgame function applied to g. Therefore, $p \in \mathsf{rng}(\mathsf{pr}_\partial \uplus \mathsf{r}) = \mathsf{rng}(\mathsf{q})$, but this is obviously impossible. Now, by Line 1 of Algorithm 5 and the definition of the subgame ∂_s, we have that $\mathsf{q}^{-1}(p) \cap \mathrm{Ps}_{\partial_s} = \mathsf{q}^{-1}(p) \cap \mathsf{dom}\big(\mathsf{g}^{(\leq p)}\big) = \mathsf{q}^{-1}(p) \cap \mathsf{g}^{-1}(p) = \mathsf{r}^{-1}(p) \neq \emptyset$. Thus, due to Line 2 of the same algorithm and Item 1.b, it holds that $\mathsf{r}^{-1}(p)$ is an α-region in ∂_s. Finally, by Line 3 and Proposition 2, R is α-maximal in ∂_s. Consequently, both Point *(1)* and the requirements of Definition 5 are satisfied. $\square$

Lemma 6 (Successor Function) *The function* $\downarrow_\partial$ *is an successor function,* i.e., *for all states* $s \in S_\partial$ *and quasi dominion pairs* $(\mathrm{R}, \alpha) \in \mathrm{QD}_\partial^-$ *with* $s \succ_\partial (\mathrm{R}, \alpha)$, *it holds that* (1) $s \downarrow_\partial (\mathrm{R}, \alpha) \in S_\partial$ *and* (2) $s \downarrow_\partial (\mathrm{R}, \alpha) \prec_\partial s$.

Proof Let $s \triangleq (\mathsf{q}, \mathsf{g}) \in S_\partial$ be a state, $(\mathrm{R}, \alpha) \in \mathrm{QD}_\partial^-$ an open quasi dominion pair in ∂ compatible with s, and $s^\star = (\mathsf{q}^\star, \mathsf{g}^\star) \triangleq s \downarrow_\partial (\mathrm{R}, \alpha)$ the result obtained by computing the function $\downarrow_\partial$ on s and (R, α). Moreover, assume $p = \min(\mathsf{rng}(\mathsf{g}))$ and $\mathsf{r} = \mathsf{q} \cap \mathsf{g}$. Two cases may be arises: R is either α-open or α-closed in ∂_s, i.e., $(\mathrm{R}, \alpha) \in \mathrm{Rg}_{\partial_s}^-$ or $(\mathrm{R}, \alpha) \notin \mathrm{Rg}_{\partial_s}^-$, respectively.

In the first case, by Lines 1-4 of Algorithm 6, we have that *(1)* $\mathsf{q}^\star = \mathsf{q}[\mathrm{R} \mapsto p]$, *(2)* $\mathrm{A} = \mathsf{F}(\partial_s \setminus \mathrm{R})$, and *(3)* $\mathsf{g}^\star = \mathsf{g}[\mathrm{A} \mapsto \max(\mathsf{rng}(\mathsf{q}^\star \restriction \mathrm{A}))]$. By Definition 7, it holds that $\mathrm{R} \subseteq \mathrm{Ps}_{\partial_s}$. Thus, *(4)* $\mathrm{R} \subseteq \mathsf{g}^{-1}(p)$, i.e. R does not contain positions with priority higher than p. As consequence of Points *(1)* and *(4)*, we have that *(5)* $\mathsf{q}^{\star(>p)} = \mathsf{q}^{\star(>p)}$, *(6)* $\mathsf{q}^{\star(<p)} \subseteq \mathsf{q}^{(<p)}$, and *(7)* $\mathsf{q}^{\star-1}(p) = \mathrm{R} \cup \mathsf{q}^{-1}(p)$.

By Point *(5)* and the assumption on q, we have that $\mathsf{q}^{\star-1}(p')$ is an α'-quasi dominion in ∂, for all priorities $p \in \mathsf{rng}(\mathsf{q})$ with $p' > p$ and $\alpha' \triangleq p' \bmod 2$. Moreover, by Item 1.c of Definition 6, for every priority $p \in \mathsf{rng}(\mathsf{q})$ with $p' < p$ and $\alpha' \triangleq p' \bmod 2$, it holds that *(8)* $\mathsf{q}^{-1}(p') \subseteq \mathsf{pr}_\partial^{-1}(p')$, so, by Point *(6)*, $\mathsf{q}^{\star-1}(p')$ is an α'-quasi dominion in ∂ as well. Again by Item 1.c, we have that $\mathsf{q}^{-1}(p)$ is the union of $\mathsf{r}^{-1}(p)$ with a set $\mathrm{P} \subseteq \mathsf{pr}_\partial^{-1}(p) \setminus \mathsf{g}^{-1}(p)$. Moreover, *(9)* $\mathsf{r}^{-1}(p) \subseteq \mathrm{R}$, due to Definition 7. Thus, by Point *(7)*, it holds that *(10)* $\mathsf{q}^{\star-1}(p) = \mathrm{R} \cup \mathsf{q}^{-1}(p) = \mathrm{R} \cup \mathrm{P} \cup \mathsf{r}^{-1}(p) = \mathrm{R} \cup \mathrm{P}$, i.e., $\mathsf{q}^{\star-1}(p)$ is the union of an α-quasi dominion in ∂ with some set of positions of priority p having the same parity. Consequently, also $\mathsf{q}^{\star-1}(p)$ is a α-quasi dominion in ∂. In conclusion, we have that *(11)* $\mathsf{q}^\star$ is a quasi-dominion function.

By Point *(2)* and the assumption on the function F, we know that *(12)* $\emptyset \neq \mathrm{A} \subseteq \partial_s \setminus \mathrm{R} = \mathrm{dom}\big(\mathsf{g}^{\leq p}\big) \setminus \mathrm{R}$ induces a subgame $\partial\!\restriction\!\mathrm{A}$ of $\partial_s \setminus \mathrm{R}$ such that $\mathrm{esc}_{\partial}^{\overline{\alpha}}(\mathrm{Q}) \cap \mathrm{pre}_{\partial}^{\overline{\alpha}}(\mathrm{Ps}_{\partial} \setminus \mathrm{A}) = \emptyset$, for every quasi dominion pair $(\mathrm{Q}, \alpha) \in \mathrm{QD}_{\partial\restriction\mathrm{A}}$. In addition, it is easy to see that $\mathsf{pr}_{\partial}^{-1}(p) \cap \mathrm{Ps}_{\partial_s} \subseteq \mathrm{R}$. Indeed, if by contradiction this was not the case, by definition of the subgame ∂_s and Point *(9)*, there would be a position $v \in \mathsf{pr}_{\partial}^{-1}(p) \cap \mathsf{g}^{-1}(p)$ such that $v \notin \mathsf{r}^{-1}(p)$. Therefore, due to Item 1.c of Definition 6, $v \in \mathsf{q}^{-1}(p)$. However, $\mathsf{r} = \mathsf{q} \cap \mathsf{g}$, so, $v \in \mathsf{q}^{-1}(p) \cap \mathsf{q}^{-1}(p) = \mathsf{r}^{-1}(p)$, which is impossible. As a consequence of this inclusion together with Point *(12)*, we have that *(13)* $p^\star = \max(\mathrm{rng}(\mathsf{q}^\star \restriction \mathrm{A})) < p$. Now, due to Points *(3)* and *(13)*, it holds that *(14)* $\mathsf{g}^{\star(>p)} = \mathsf{g}^{(>p)}$, *(15)* $\mathsf{g}^{\star-1}(p) = \mathsf{g}^{-1}(p) \setminus \mathrm{A}$, *(16)* $\mathrm{dom}\big(\mathsf{g}^{\star(<p)}\big) = \mathsf{g}^{\star-1}(p^\star) = \mathrm{A}$, and *(17)* $\mathrm{dom}\big(\mathsf{g}^{\star(<p^\star)}\big) = \emptyset$. Observe that, by Point *(14)*, g and $\mathsf{g}^\star$ only differ on the priorities p and $p^\star$. In particular, $\mathrm{rng}(\mathsf{g}^\star) = \mathrm{rng}(\mathsf{g}) \cup \{p^\star\}$. Thus, by Item *(i)* of the definition of subgame function and Points *(8)* and *(13)*, we have that $\mathrm{rng}(\mathsf{g}^\star) \subseteq \mathrm{rng}(\mathsf{pr}_{\partial})$. By Points *(15)* and *(16)*, *(18)* $\partial_{\mathsf{g}^\star}^{\leq p} = \partial_{\mathsf{g}}^{\leq p}$ is obviously a game. Moreover, due to Point *(4)*, it holds that *(19)* $\mathrm{dom}\big(\mathsf{g}^{\star(\leq p^\star)}\big) = \mathrm{A} \subset \mathrm{dom}\big(\mathsf{g}^{\star(\leq p)}\big)$. Finally, $\partial_{\mathsf{g}^\star}^{\leq p^\star}$ is a game, by Point *(12)*. Summing up, it follows that *(20)* $\mathsf{g}^\star$ is a subgame function.

We can now prove the tree parts in which Item 1 of Definition 6 applied to $(\mathsf{q}^\star, \mathsf{g}^\star)$ is split. First notice that, since $\mathsf{r} = \mathsf{q} \cap \mathsf{g}$, we have that *(21)* $\mathsf{r}^{\star(>p)} = \mathsf{r}^{(>p)}$, due to Points *(5)* and *(14)*. Moreover, *(22)* $\mathsf{r}^{\star-1}(p^\star) \subseteq \mathsf{pr}_{\partial}^{-1}(p^\star)$, by Point *(8)*, and *(23)* $\mathsf{r}^{\star(<p^\star)} = \varnothing$, by Point *(17)*. Finally, by Points *(10)*, *(13)*, and *(15)* and the observation that $\mathrm{P} \cap \mathsf{g}^{-1}(p) = \emptyset$, it holds that *(24)* $\mathsf{r}^{\star-1}(p) = \mathrm{R}$. At this point, Item 1.a of the definition of state space follows from Points *(21)* and *(24)*, since R is α-maximal in ∂_s due to Definition 7. By putting together Points *(18)*-*(23)* with Point *(12)*, also Item 1.b can be derived. To prove that Item 1.c, *i.e.*, $\mathsf{q}^\star = \mathsf{pr}_{\partial} \uplus \mathsf{r}^\star$, one can use Points *(5)* and *(21)* for the priorities greater than p, Points *(10)* and *(24)* for the priority p, and Point *(8)* for the priorities smaller than p.

In the end, by Points *(11)* and *(20)*, and the fact that Item 1 of Definition 6 holds on $(\mathsf{q}^\star, \mathsf{g}^\star)$, it follows that $(\mathsf{q}^\star, \mathsf{g}^\star) \in \mathrm{S}_{\partial}$.

To conclude this case of the proof, we need just to show that $(\mathsf{q}^\star, \mathsf{g}^\star) \prec_{\partial} (\mathsf{q}, \mathsf{g})$. If $\mathrm{R} \subseteq \mathsf{q}^{-1}(p)$, by Point *(1)* and *(13)*, we have that $\mathsf{q}^\star = \mathsf{q}$ and $p^\star < p$. Thus, the thesis is derived from Item 3.b of Definition 6. If, on the other hand, $\mathrm{R} \setminus \mathsf{q}^{-1}(p) \neq \emptyset$, due to Point *(5)* and *(7)*, the thesis follows from Item 3.a of the same definition.

Consider now the case in which the set R is α-closed in ∂_s. By Lines 1, 5-7 of Algorithm 6, we have that *(25)* $p^\star = \mathrm{bep}_{\partial}^{\overline{\alpha}}(\mathrm{R}, \mathsf{q})$, *(26)* $\mathsf{q}^\star = \mathsf{pr}_{\partial} \uplus \mathsf{q}^{(\geq p^\star)}[\mathrm{R} \mapsto p^\star]$, and *(27)* $\mathsf{g}^\star = \mathsf{g}[\mathrm{dom}\big(\mathsf{g}^{(<p^\star)}\big) \mapsto p^\star]$. Due to Point *(25)*, the definition of best escape priority, and the fact that R is α-closed in ∂_s, it is not hard to see that *(28)* $p^\star > p$. Moreover, due to Item *(ii)* of the definition of g stratified region function, it holds that *(29)* $p^\star \in \mathrm{rng}(\mathsf{r}^\star)$, since the moves escaping from R can only reach positions in $\mathrm{dom}\big(\mathsf{r}^{(>p)}\big)$.

Now, by Point *(27)*, it follows that *(30)* $\mathsf{g}^{\star(>p^\star)} = \mathsf{g}^{(>p^\star)}$, *(31)* $\mathsf{g}^{\star(<p^\star)} = \emptyset$, and *(32)* $\mathrm{dom}\big(\mathsf{g}^{\star(\leq p^\star)}\big) = \mathrm{dom}\big(\mathsf{g}^{(\leq p^\star)}\big)$. Therefore, it is immediate to see that *(33)* $\mathsf{g}^\star$ is a subgame function.

Moreover, by Point *(26)*, it holds that *(34)* $\mathsf{q}^{\star(>p^\star)} = \mathsf{q}^{(>p^\star)}$, *(35)* $\mathsf{q}^{\star-1}(p^\star) = \mathsf{q}^{-1}(p^\star) \cup \mathrm{R}$, and *(36)* $\mathsf{q}^{\star(<p^\star)} \subseteq \mathsf{pr}_{\partial}^{(<p^\star)}$. Consequently, *(37)* $\mathsf{r}^{\star(>p^\star)} = \mathsf{r}^{(>p^\star)}$, by Points *(30)* and *(34)*, and *(38)* $\mathsf{r}^{\star(<p^\star)} = \emptyset$, by Point *(31)*. In addition, due to Points *(32)* and *(36)*, together with Item *(ii.b)* of the definition of subgame function and Proposition 1, we have that *(39)* $\mathsf{r}^{\star-1}(p^\star) = \mathsf{q}^{\star-1}(p^\star) \cap \mathsf{g}^{\star-1}(p^\star) = (\mathsf{q}^{-1}(p^\star) \cup \mathrm{R}) \cap \mathsf{g}^{\star-1}(p^\star) = \mathsf{r}^{-1}(p^\star) \cup \mathrm{R}$ is an α-region in $\partial_{s^\star}$. At this point, Item 1 of Definition 6 follows by simply verifying that all requirements are satisfied.

Because of Points *(34)*-*(36)*, to prove that *(40)* $\mathsf{q}^\star$ is a quasi-dominion function, it is enough to observe that $\mathsf{q}^{\star-1}(p^\star)$ is the union of $\mathsf{r}^{\star-1}(p^\star))$ with a set $\mathrm{P} \subseteq \mathsf{pr}_{\partial}^{-1}(p^\star)) \setminus \mathsf{g}^{\star-1}(p^\star))$. Indeed, $\mathsf{q}^{\star-1}(p^\star)$ is an α-quasi dominion in ∂, being the union of an α-region in $\partial_{s^\star}$ and a set of positions having priority $p^\star$).

In the end, by Points *(33)* and *(40)*, and the fact that Item 1 of Definition 6 holds on $(\mathsf{q}^\star, \mathsf{g}^\star)$, it follows that $(\mathsf{q}^\star, \mathsf{g}^\star) \in \mathrm{S}_{\partial}$.

To conclude, as for the previous case, it remains to show that $(\mathsf{q}^\star, \mathsf{g}^\star) \prec_{\partial} (\mathsf{q}, \mathsf{g})$. This fact easily follows from Item 3.a of Definition 6, where the priority p is set to $p^\star$. Indeed, by Definition 7, we have that $\emptyset \neq \mathrm{R} \subseteq \mathrm{Ps}_{\partial_s}$, so, $\mathsf{q}^{-1}(p^\star) \cap \mathrm{R} = \emptyset$, since $p^\star > p$ as stated in Point *(28)*. Therefore, by Point *(31)*, it follows that $\mathsf{q}^{-1}(p^\star) \subset \mathsf{q}^{-1}(p^\star) \cup \mathrm{R} = \mathsf{q}^{\star-1}(p^\star)$. $\qquad\square$

References

1. M. Agrawal, N. Kayal, and N. Saxena. PRIMES is in P. *Annals of Mathematics*, 160(2):781–793, 2004.

2. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. *Journal of the ACM*, 49(5):672–713, 2002.
3. K. Apt and E. Grädel. *Lectures in Game Theory for Computer Scientists.* Cambridge University Press, 2011.
4. M. Benerecetti, D. Dell'Erba, and F. Mogavero. Solving Parity Games via Priority Promotion. In *Computer Aided Verification'16*, LNCS 9780 (Part II), pages 270–290. Springer, 2016.
5. M. Benerecetti, F. Mogavero, and A. Murano. Substructure Temporal Logic. In *Logic in Computer Science'13*, pages 368–377. IEEE Computer Society, 2013.
6. M. Benerecetti, F. Mogavero, and A. Murano. Reasoning About Substructures and Games. *Transactions On Computational Logic*, 16(3):25:1–46, 2015.
7. D. Berwanger, A. Dawar, P. Hunter, and S. Kreutzer. DAG-Width and Parity Games. In *Symposium on Theoretical Aspects of Computer Science'06*, LNCS 3884, pages 524–536. Springer, 2006.
8. D. Berwanger and E. Grädel. Games and Model Checking for Guarded Logics. In *Logic for Programming Artificial Intelligence and Reasoning'01*, LNCS 2250, pages 70–84. Springer, 2001.
9. D. Berwanger and E. Grädel. Fixed-Point Logics and Solitaire Games. *Theoretical Computer Science*, 37(6):675–694, 2004.
10. D. Berwanger, E. Grädel, L. Kaiser, and R. Rabinovich. Entanglement and the Complexity of Directed Graphs. *Theoretical Computer Science*, 463:2–25, 2012.
11. C.S. Calude, S. Jain, B. Khoussainov, W. Li, and F. Stephan. Deciding Parity Games in Quasipolynomial Time. In *Symposium on Theory of Computing'17*, pages 252–263. Association for Computing Machinery, 2017.
12. K. Chatterjee and L. Doyen. Energy Parity Games. *Theoretical Computer Science*, 458:49–60, 2012.
13. K. Chatterjee, L. Doyen, T.A. Henzinger, and J.-F. Raskin. Generalized Mean-Payoff and Energy Games. In *Foundations of Software Technology and Theoretical Computer Science'10*, LIPIcs 8, pages 505–516. Leibniz-Zentrum fuer Informatik, 2010.
14. K. Chatterjee, T.A. Henzinger, and F. Horn. Finitary Winning in omega-Regular Games. *Transactions On Computational Logic*, 11(1):1:1–26, 2010.
15. K. Chatterjee, T.A. Henzinger, and M. Jurdziński. Mean-Payoff Parity Games. In *Logic in Computer Science'05*, pages 178–187. IEEE Computer Society, 2005.
16. K. Chatterjee, T.A. Henzinger, and N. Piterman. Strategy Logic. *Information and Computation*, 208(6):677–693, 2010.
17. A. Condon. The Complexity of Stochastic Games. *Information and Computation*, 96(2):203–224, 1992.
18. A. Ehrenfeucht and J. Mycielski. Positional Strategies for Mean Payoff Games. *International Journal of Game Theory*, 8(2), 1979.
19. E.A. Emerson and C.S. Jutla. Tree Automata, muCalculus, and Determinacy. In *Foundation of Computer Science'91*, pages 368–377. IEEE Computer Society, 1991.
20. E.A. Emerson, C.S. Jutla, and A.P. Sistla. On Model Checking for the muCalculus and its Fragments. In *Computer Aided Verification'93*, LNCS 697, pages 385–396. Springer, 1993.
21. E.A. Emerson, C.S. Jutla, and A.P. Sistla. On Model Checking for the muCalculus and its Fragments. *Theoretical Computer Science*, 258(1-2):491–522, 2001.
22. E.A. Emerson and C.-L. Lei. Temporal Reasoning Under Generalized Fairness Constraints. In *Symposium on Theoretical Aspects of Computer Science'86*, LNCS 210, pages 267–278. Springer, 1986.
23. J. Fearnley. Non-Oblivious Strategy Improvement. In *Logic for Programming Artificial Intelligence and Reasoning'10*, LNCS 6355, pages 212–230. Springer, 2010.
24. J. Fearnley, S. Jain, S. Schewe, F. Stephan, and D. Wojtczak. An Ordered Approach to Solving Parity Games in Quasi Polynomial Time and Quasi Linear Space. In *SPIN Symposium on Model Checking of Software'2017*, pages 112–121. Association for Computing Machinery, 2017.
25. J. Fearnley and O. Lachish. Parity Games on Graphs with Medium Tree-Width. In *Mathematical Foundations of Computer Science'11*, LNCS 6907, pages 303–314. Springer, 2011.
26. J. Fearnley and S. Schewe. Time and Parallelizability Results for Parity Games with Bounded Treewidth. In *International Colloquium on Automata, Languages, and Programming'12*, LNCS 7392, pages 189–200. Springer, 2012.
27. M.R. Fellows and N. Koblitz. Self-Witnessing Polynomial-Time Complexity and Prime Factorization. In *Conference on Structure in Complexity Theory'92*, pages 107–110. IEEE Computer Society, 1992.
28. M.R. Fellows and N. Koblitz. Self-Witnessing Polynomial-Time Complexity and Prime Factorization. *Designs, Codes, and Cryptography*, 2(3):231–235, 1992.
29. N. Fijalkow and M. Zimmermann. Cost-Parity and Cost-Streett Games. In *Foundations of Software Technology and Theoretical Computer Science'12*, LIPIcs 18, pages 124–135. Leibniz-Zentrum fuer Informatik, 2012.
30. N. Fijalkow and M. Zimmermann. Cost-Parity and Cost-Streett Games. *Logical Methods in Computer Science*, 10(2):1–29, 2014.

31. O. Friedmann and M. Lange. Solving Parity Games in Practice. In *Automated Technology for Verification and Analysis'09*, LNCS 5799, pages 182–196. Springer, 2009.
32. E. Grädel, W. Thomas, and T. Wilke. *Automata, Logics, and Infinite Games: A Guide to Current Research*. LNCS 2500. Springer, 2002.
33. V.A. Gurvich, A.V. Karzanov, and L.G. Khachivan. Cyclic Games and an Algorithm to Find Minimax Cycle Means in Directed Graphs. *USSR Computational Mathematics and Mathematical Physics*, 28(5):85–91, 1990.
34. M. Jurdziński. Deciding the Winner in Parity Games is in UP ∩ co-UP. *Information Processing Letters*, 68(3):119–124, 1998.
35. M. Jurdziński. Small Progress Measures for Solving Parity Games. In *Symposium on Theoretical Aspects of Computer Science'00*, LNCS 1770, pages 290–301. Springer, 2000.
36. M. Jurdziński and R. Lazic. Succinct Progress Measures for Solving Parity Games. In *Logic in Computer Science'17*, pages 1–9. Association for Computing Machinery, 2017. Accepted for publication.
37. M. Jurdziński, M. Paterson, and U. Zwick. A Deterministic Subexponential Algorithm for Solving Parity Games. In *Symposium on Discrete Algorithms'06*, pages 117–123. Society for Industrial and Applied Mathematics, 2006.
38. M. Jurdziński, M. Paterson, and U. Zwick. A Deterministic Subexponential Algorithm for Solving Parity Games. *SIAM Journal on Computing*, 38(4):1519–1532, 2008.
39. N. Klarlund and D. Kozen. Rabin Measures and Their Applications to Fairness and Automata Theory. In *Logic in Computer Science'91*, pages 256–265. IEEE Computer Society, 1991.
40. O. Kupferman and M.Y. Vardi. Weak Alternating Automata and Tree Automata Emptiness. In *Symposium on Theory of Computing'98*, pages 224–233. Association for Computing Machinery, 1998.
41. A.D. Martin. Borel Determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
42. A.D. Martin. A Purely Inductive Proof of Borel Determinacy. In *Symposia in Pure Mathematics'82*, Recursion Theory., pages 303–308. American Mathematical Society and Association for Symbolic Logic, 1985.
43. R. McNaughton. Infinite Games Played on Finite Graphs. *Annals of Pure and Applied Logic*, 65:149–184, 1993.
44. F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. What Makes ATL* Decidable? A Decidable Fragment of Strategy Logic. In *Concurrency Theory'12*, LNCS 7454, pages 193–208. Springer, 2012.
45. F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. Reasoning About Strategies: On the Model-Checking Problem. *Transactions On Computational Logic*, 15(4):34:1–42, 2014.
46. F. Mogavero, A. Murano, and L. Sorrentino. On Promptness in Parity Games. In *Logic for Programming Artificial Intelligence and Reasoning'13*, LNCS 8312, pages 601–618. Springer, 2013.
47. F. Mogavero, A. Murano, and M.Y. Vardi. Reasoning About Strategies. In *Foundations of Software Technology and Theoretical Computer Science'10*, LIPIcs 8, pages 133–144. Leibniz-Zentrum fuer Informatik, 2010.
48. A.W. Mostowski. Regular Expressions for Infinite Trees and a Standard Form of Automata. In *Symposium on Computation Theory'84*, LNCS 208, pages 157–168. Springer, 1984.
49. A.W. Mostowski. Games with Forbidden Positions. Technical report, University of Gdańsk, Gdańsk, Poland, 1991.
50. J. Obdržálek. Fast Mu-Calculus Model Checking when Tree-Width Is Bounded. In *Computer Aided Verification'03*, LNCS 2725, pages 80–92. Springer, 2003.
51. J. Obdržálek. Clique-Width and Parity Games. In *Computer Science Logic'07*, LNCS 4646, pages 54–68. Springer, 2007.
52. S. Schewe. Solving Parity Games in Big Steps. In *Foundations of Software Technology and Theoretical Computer Science'07*, LNCS 4855, pages 449–460. Springer, 2007.
53. S. Schewe. An Optimal Strategy Improvement Algorithm for Solving Parity and Payoff Games. In *Computer Science Logic'08*, LNCS 5213, pages 369–384. Springer, 2008.
54. S. Schewe. ATL* Satisfiability is 2ExpTime-Complete. In *International Colloquium on Automata, Languages, and Programming'08*, LNCS 5126, pages 373–385. Springer, 2008.
55. S. Schewe and B. Finkbeiner. Satisfiability and Finite Model Property for the Alternating-Time muCalculus. In *Computer Science Logic'06*, LNCS 6247, pages 591–605. Springer, 2006.
56. J. Vöge and M. Jurdziński. A Discrete Strategy Improvement Algorithm for Solving Parity Games. In *Computer Aided Verification'00*, LNCS 1855, pages 202–215. Springer, 2000.
57. T. Wilke. Alternating Tree Automata, Parity Games, and Modal muCalculus. *Bulletin of the Belgian Mathematical Society*, 8(2):359–391, 2001.
58. W. Zielonka. Infinite Games on Finitely Coloured Graphs with Applications to Automata on Infinite Trees. *Theoretical Computer Science*, 200(1-2):135–183, 1998.
59. U. Zwick and M. Paterson. The Complexity of Mean Payoff Games on Graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996.