Priority Promotion with Parysian Flair

Massimo Benerecetti[a], Daniele Dell'Erba[b], Fabio Mogavero[a], Sven Schewe[b], Dominik Wojtczak[b]

[a]*Università degli Studi di Napoli Federico II, Naples, Italy*
[b]*University of Liverpool, Liverpool, UK*

Abstract

We develop an algorithm that combines the advantages of Priority Promotion, that is one of the leading approaches to solving large parity games in practice, with the quasi-polynomial time guarantees offered by Parys' algorithm. Hybridising these algorithms sounds both natural and difficult, as they both generalise the classic recursive algorithm in different ways that appear to be irreconcilable: while the promotion transcends the call structure, the guarantees change on each level. We show that an interface that respects both is not only effective, but also efficient.

1. Introduction

Parity games are perfect-information two-player turn-based games of infinite duration, usually played on finite directed graphs. Their vertices, labelled by natural numbers, called *priorities*, are assigned to one of two players, named *Even* and *Odd* (or simply 0 and 1, respectively). A play in the game consists of a sequence of moves between vertices and it is said to be winning for player 0 (*resp.*, 1), if the maximal priority encountered infinitely often along the play is even (*resp.*, odd). Parity games have many applications in the context of formal verification and synthesis of systems. Computing winning strategies for these games is linear-time equivalent to solving the modal μ-calculus model-checking problem [1, 2]. Parity games has been studied in automata theory [3, 4, 5] and can be applied to solve the complementation problem for alternating automata [6] or the emptiness of the corresponding nondeterministic tree automata [7]. These automata, in turn, can be used to solve the satisfiability and model-checking problems for several expressive logics [8, 9, 10, 11, 12, 13, 14], such as μ-calculus [15, 16] and ATL* [17, 18]. On the complexity-theoretic side, determining the winner of a parity game is a problem that lies in NP ∩ co-NP [1], being memoryless determined [19, 20, 21]. It has even been proved to belong to UP ∩ co-UP [22], and later to be solvable in quasi-polynomial time [23]. However, determining whether they belong also to P is still an open challenge.

Research on parity games falls into two categories: on the one hand to develop *fast solvers*; on the other hand to determine the *complexity of parity games* by finding algorithms with a good *worst-case complexity*. With its practical

motivation, one of the most efficient approach for solving parity games is currently the *Priority Promotion* technique [24, 25, 26, 27, 28], a derivation of the classic *recursive algorithm* [29, 30] that follows the iterated fixed-point structure induced by the parity condition. Like all of practically efficient algorithms, it is an exponential approach, however, algorithms that are good in practice do not tend to display their worst-case behaviour, except in carefully designed hostile examples. This holds in particular for strategy improvement algorithms [31, 32, 33, 34, 35, 36, 37], which were considered candidates for tractable algorithms until they were shown to be exponential by Friedman's delicate lower bound constructions [38, 39, 40] and by [41] for the symmetric approach of [37]. But while it is easier to design hard classes for recursive [42, 43, 44] and Priority Promotion algorithms [28], these classes are still not relevant in practice.

One of the most celebrated results in recent years has been the landmark result of [23], which established that parity games can be solved in quasi-polynomial time (QP). This was a major step from former deterministic algorithms, which were (at least) exponential in the number of priorities [29, 2, 45, 46, 30, 47, 34, 35, 48, 27] ($n^{O(c)}$), or in the square-root of the number of game positions [31, 49, 34] (approximately $n^{O(\sqrt{n})}$). The breakthrough of [23] has triggered a new line of research into QP algorithms, including its refinements [50, 51, 52], the register-index algorithm [53, 54], the study on the Strahler Number [55], and the bounded version of the recursive algorithm [56]. Unfortunately, all these quasi-polynomial algorithms can be derived by the separation approach that also provides a lower bound for these techniques [57].

Interestingly, Parys' algorithm [56] and variations thereof [58, 59], which, like Priority Promotion techniques, adjust the classic recursive algorithm [29, 30], are relatively fast among the QP algorithms, where [56] has the edge on benchmarks, while [58] has the edge on theoretical guarantees. On first glance, this seems to invite synthesising one of these algorithms with Priority Promotion. This upgrade would combine the best of both worlds: offer QP lower bounds without undue compromise on efficiency. On second glance, the prospect of this synthesis seems less promising. Priority Promotion techniques [26, 27, 28, 60, 61, 62] achieve their advancement over the classic recursive algorithm [30] by globally bypassing the call structure through temporally increasing the priority of a position. Parys' approach, on the other hand, locally creates sets with guarantees weakened (halved) along the recursive call structure, where subgames are split into areas that contain all 0-dominions with size up to a bound b_0 and all 1-dominions of size up to a bound b_1; one of these bounds is halved in each call until the guarantees are trivial. *Prima facie*, it seems clear that such guarantees are ill suited for a promotion across the call structure. We did, however, find that, when one shifts the view on the essence of a promotion from creating quasi-dominions to creating regions and promoting them to the lowest level where they are no longer dominions, this allows for a concurrent treatment of sets with bounded guarantees (the Parysian flair of our hybrid algorithm) and with unbounded guarantees (the Priority Promotion core of our algorithm). While the integration of these seemingly antagonistic concepts is intricate, it provides an efficient bridge between the behaviour and the data structures of [27] and [56]:

the resulting algorithm guarantees a quasi-polynomial running time, and offers excellent practical behaviour on the benchmarks we have tested it against.

2. Preliminaries

A two-player turn-based *arena* is a tuple $\mathcal{A} = \langle \mathrm{Ps}_0, \mathrm{Ps}_1, \mathsf{Mv} \rangle$, with $\mathrm{Ps}_0 \cap \mathrm{Ps}_1 = \emptyset$ and $\mathrm{Ps} \triangleq \mathrm{Ps}_0 \cup \mathrm{Ps}_1$, such that $\langle \mathrm{Ps}, \mathsf{Mv} \rangle$ is a finite directed graph without sinks. Ps_0 (*resp.*, Ps_1) is the set of positions of the Even (*resp.*, Odd) and $\mathsf{Mv} \subseteq \mathrm{Ps} \times \mathrm{Ps}$ 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.*, $(\pi_i, \pi_{i+1}) \in \mathsf{Mv}$, for all $i \in [0, |\pi| - 1)$. A positional *strategy* for player $\alpha \in \{0, 1\}$ on $\mathrm{V} \subseteq \mathrm{Ps}$ is a function $\sigma_\alpha \in \mathrm{Str}_\alpha(\mathrm{V}) \subseteq (\mathrm{V} \cap \mathrm{Ps}_\alpha) \to \mathrm{V}$, mapping each α-position v in V to position $\sigma_\alpha(v)$ compatible with the move relation, *i.e.*, $(v, \sigma_\alpha(v)) \in \mathsf{Mv}$. With $\mathrm{Str}_\alpha(\mathrm{V})$ we denote the set of all α-strategies on V. When talking about players, we will refer to the opponent player of α as $\overline{\alpha}$. A *play* in $\mathrm{V} \subseteq \mathrm{Ps}$ from a position $v \in \mathrm{V}$ *w.r.t.* a pair of strategies $(\sigma_0, \sigma_1) \in \mathrm{Str}_0(\mathrm{V}) \times \mathrm{Str}_1(\mathrm{V})$, called $((\sigma_0, \sigma_1), v)$-*play*, is a path $\pi \in \mathrm{Pth}(\mathrm{V})$ such that $(\pi)_0 = v$ and, for all $i \in [0, |\pi| - 1)$, if $(\pi)_i \in \mathrm{Ps}_0$ then $(\pi)_{i+1} = \sigma_0((\pi)_i)$ else $(\pi)_{i+1} = \sigma_1((\pi)_i)$. The *play function* $\mathsf{play} : (\mathrm{Str}_0(\mathrm{V}) \times \mathrm{Str}_1(\mathrm{V})) \times \mathrm{V} \to \mathrm{Pth}(\mathrm{V})$ returns, for each position $v \in \mathrm{V}$ and pair of strategies $(\sigma_0, \sigma_1) \in \mathrm{Str}_0(\mathrm{V}) \times \mathrm{Str}_1(\mathrm{V})$, the maximal $((\sigma_0, \sigma_1), v)$-play $\mathsf{play}((\sigma_0, \sigma_1), v)$. Given a partial function $\mathsf{f} : A \rightharpoonup B$, with $\mathsf{dom}(\mathsf{f}) \subseteq A$ and $\mathsf{rng}(\mathsf{f}) \subseteq B$ we denote the domain and the range of f, respectively.

A *parity game* is a tuple $\eth = \langle \mathcal{A}, \mathrm{Pr}, \mathsf{pr} \rangle \in \mathrm{PG}$, where $\mathcal{A}$ is an arena, $\mathrm{Pr} \subset \mathbb{N}$ is a finite set of priorities, and $\mathsf{pr} : \mathrm{Ps} \to \mathrm{Pr}$ is a *priority function* assigning a priority to each position. The priority function can be naturally extended to games and paths as follows: $\mathsf{pr}(\eth) \triangleq \max_{v \in \mathrm{Ps}} \mathsf{pr}(v)$; for a path $\pi \in \mathrm{Pth}$, we define $\mathsf{pr}(\pi) \triangleq \max_{i \in [0, |\pi|)} \mathsf{pr}((\pi)_i)$, if π is finite, and $\mathsf{pr}(\pi) \triangleq \limsup_{i \in \mathbb{N}} \mathsf{pr}((\pi)_i)$, otherwise. A set of positions $\mathrm{V} \subseteq \mathrm{Ps}$ is an α-*dominion*, with $\alpha \in \{0, 1\}$, if there exists an α-strategy $\sigma_\alpha \in \mathrm{Str}_\alpha(\mathrm{V})$ such that, for all $\overline{\alpha}$-strategies $\sigma_{\overline{\alpha}} \in \mathrm{Str}_{\overline{\alpha}}(\mathrm{V})$ and positions $v \in \mathrm{V}$, the induced play $\pi = \mathsf{play}((\sigma_0, \sigma_1), v)$ is infinite and $\mathsf{pr}(\pi) \equiv_2 \alpha$. In other words, σ_α only induces on V infinite plays whose maximal priority visited infinitely often has parity α. The *winning region* for player $\alpha \in \{0, 1\}$ in game $\eth$, denoted by $\mathrm{Wn}_\eth^\alpha$, is the greatest set of positions that is also a α-dominion in $\eth$. Since parity games are memoryless determined [20], meaning that from each position one of the two players wins, the two winning regions of a game $\eth$ form a partition of its positions, *i.e.*, $\mathrm{Wn}_\eth^0 \cup \mathrm{Wn}_\eth^1 = \mathrm{Ps}_\eth$.

The (maximal) subgame $\eth' = \eth \setminus \mathrm{V}$ of $\eth$ is the game with set of positions Ps' contained in $\mathrm{Ps} \setminus \mathrm{V}$ and move relation Mv' equal to the restriction of Mv to Ps'. The α-predecessor of V, in symbols $\mathsf{pre}^\alpha(\mathrm{V}) \triangleq \{v \in \mathrm{Ps}_\alpha \mid \mathsf{Mv}(v) \cap \mathrm{V} \neq \emptyset\} \cup \{v \in \mathrm{Ps}_{\overline{\alpha}} \mid \mathsf{Mv}(v) \subseteq \mathrm{V}\}$, collects the positions from which player α can force the game to reach some position in V with a single move. The α-attractor $\mathsf{atr}^\alpha(\mathrm{V})$ generalises the notion of α-predecessor $\mathsf{pre}^\alpha(\mathrm{V}) \cup \mathrm{V}$ to an arbitrary number of moves. Thus, it corresponds to the least fix-point of that operator. When $\mathrm{V} = \mathsf{atr}^\alpha(\mathrm{V})$, player α cannot force any position outside V to enter this set that is, therefore, called α-maximal. For such a V, the set of positions of the subgame

$\Im \setminus V$ is precisely $Ps \setminus V$. When the computation of the attractor is restricted to a given set of positions X, we will use the notation $\mathsf{atr}^\alpha(V, X)$ which corresponds to the least fix-point of $\mathsf{pre}^\alpha(V) \cap X$. Finally, the set $\mathsf{esc}^\alpha(V) \triangleq \mathsf{pre}^\alpha(Ps \setminus V) \cap V$, called the α-*escape* of V, contains the positions in V from which α can leave V in one move. Observe that all the operators and sets described above actually depend on the specific game $\Im$ they are applied to. In the rest of the paper, we shall only add $\Im$ as subscript of an operator, *e.g.* $\mathsf{atr}^\alpha_\Im(V)$, when the game is not clear from the context.

Outline

We introduce our hybrid algorithm in three steps. In the first step (Section 3), we introduce a variation of the classic Priority Promotion approach, which serves as the backbone of our hybrid algorithm in Section 5. We provide a recap of how Priority Promotion operates and an introduction to its data structure, which we later extend for our hybrid algorithm. In a nutshell, Priority Promotion accelerates the classic recursive algorithm by allowing to merge dominions in subgames that span non-adjacent recursive calls, which is the essence of the promotion operations. In the subsequent section (Section 4), we outline Parys' algorithm, which does not seek to identify *all dominions* on a level, but merely *all small dominions* up to given bounds b_0 and b_1 for the dominions of Player 0 and Player 1, respectively. It truncates the size of the call tree by making all but one call with half the precision for one of the players. Here, we formulate the algorithm with a terminology analogous to Priority Promotion, and present it in a form similar to our hybrid algorithm.

The two concepts of Priority Promotion and truncated tree size through limited guarantees appear to be unlikely allies: not only does the presence of Parys' sets with limited guarantees impede the promotion of dominions, any attempt to promote sets with bounded guarantees is doomed to fail, when the bounds are larger (and thus the required guarantees stronger) along the call tree. In Section 5, we see that, when synthesising the algorithms carefully, sets with the 'region guarantees' from Priority Promotion and with 'bounded guarantees' from Parys' approach can co-exist, so long as they are kept carefully apart and treated differently.

The resulting algorithm can identify dominions in many places, and these dominions can be promoted. This promotion can be to a set with 'region guarantees' at a higher level, but it can also be that the correct target is a set with 'bounded guarantees'; this works across levels because dominions have unbounded guarantees. The identification of the right set to promote to, instead, remains fairly similar to the way such a set is identified in classic Priority Promotion. While sets with bounded guarantees cannot be promoted along the data structure (which follows the call tree), they lose parts of their locality: positions can be promoted into them, and, crucially, they do not prevent promotions to higher levels. This way, we can keep the Priority Promotion part (which usually carries the main burden of solving the parity game) and can play out its practical efficiency in full, while we also retain the quasi-polynomial

4

Figure 1: A game and some of its quasi dominions.

complexity from Parys' algorithm [56], bypassing the known hard cases for recursive algorithms. For practical considerations, it is still computationally attractive to grow the bounded sets more slowly: we found that some of the points where Parys' algorithm applies a closure of sets with bounded guarantees are merely for the convenience of the proof. For efficiency, we have restricted the closure under attractor of these sets to the places where it is necessary for correctness.

3. The Priority Promotion Approach

The *Priority Promotion approaches* [24, 25, 26, 27, 28] attack the problem of solving a parity game $\partial \in \mathrm{PG}$ by iteratively computing, one at a time, a sequence of α-dominions $\mathrm{D}_0^\alpha, \mathrm{D}_1^\alpha, \ldots \subseteq \mathrm{Ps}$, for some player $\alpha \in \mathbb{B} \triangleq \{0, 1\}$ (potentially, for both of them). These are in fact portions of the two winning regions, Wn_0 and Wn_1, that need to be identified (the original notion of dominion has been introduced in [63, 49]). The idea here is to start from a weaker notion, called *quasi dominion*. Similarly to a dominion, a quasi dominion is a set of positions over which one of the two players has a strategy defined on that set, whose induced plays, if infinite, are winning for that player. In contrast to dominions, it's worth noting that some of these plays could be finite. This is because the opponent may have the possibility to escape from those positions towards a different part of the game, hoping for a better outcome. The approach uses quasi dominions by composing them until an actual dominion is obtained. The name of the approach refers to the principle that this composition is computed by applying the following operation called *promotion*. Given two quasi dominions Q_1 and Q_2, to which some priorities $p_1 < p_2$ of the same parity are assigned, Q_1 is combined with Q_2 by promoting the former to the priority of the latter. As a result, together they form a bigger quasi dominion $\mathrm{Q}_1 \cup \mathrm{Q}_2$, whose internal infinite plays are still winning for the same player.

Definition 1 (Quasi Dominion)**.** *A set of positions* $\mathrm{Q} \subseteq \mathrm{Ps}$ *is a* quasi α-dominion, *for some player* $\alpha \in \mathbb{B}$, *if there exists an* α-*strategy* $\sigma_\alpha \in \mathrm{Str}^\alpha(\mathrm{Q})$, *called* α-witness *for* Q, *such that, for all* $\overline{\alpha}$-*strategies* $\sigma_{\overline{\alpha}} \in \mathrm{Str}^{\overline{\alpha}}(\mathrm{Q})$ *and positions* $v \in \mathrm{Q}$, *the induced play* $\pi = \mathsf{play}((\sigma_0, \sigma_1), v)$, *satisfies* $\mathsf{pr}(\pi) \equiv_2 \alpha$, *if infinite*.

5

The usefulness of the above concept, in addition to the property of being suitably composable, resides in the fact that quasi dominions are closed under inclusion. Thus, when a subset of an α-quasi dominion is found for which player $\overline{\alpha}$ has no strategy that leaves the set, a dominion is identified. Note that here no constraints are imposed on the finite plays, unlike in the definition given in the Priority Promotion paper [24], where quasi dominions were not closed under inclusion. However, an analogous constraint is implicitly imposed on the basic data structure leveraged by the algorithm and introduced in Definition 4 below. It is interesting to notice that this version of quasi dominion is similar to the notion of a *snare* [36].

Theorem 1 (Induced Dominion). *Let $\alpha \in \mathbb{B}$ be a player, $Q \subseteq Ps$ a quasi α-dominion, $\sigma \in \mathrm{Str}_\alpha(Q)$ one of its α-witnesses, and $D \subseteq Q$ a subset such that $\sigma(v) \in D$, if $v \in \mathrm{Ps}_\alpha$, and $\mathrm{Mv}(v) \subseteq D$, otherwise, for all positions $v \in D$. Then, D is an α-dominion.*

Example 1. *As an example of quasi dominions, consider the game depicted in Figure 1, where circle-shaped positions belong to Player 0 and square-shaped ones to Player 1. Clearly, g and h are won by Player 0, while the rest of the game is won by Player 1. The dashed blue region, containing the four positions d, f, g, and h, is one of the quasi 0-dominion in the game, where one of its two possible 0-witnesses is highlighted in blue (the other one being the 0-strategy where Player 0 moves from h to g). The solid blue subregion is a 0-dominion induced by the single position g. The dashed red region, containing the remaining half of the game, is instead a quasi 1-dominion with the 1-witnesses highlighted in red. Finally, note that this region is also a quasi 0-dominion with 0-witness mapping position a to b. Indeed, any play compatible with this witness will eventually exit the region. In general, a set of position is both a quasi 0-dominion and a quasi 1-dominion if both players can force exiting the set.*

The solution algorithms following this approach carry on the search for a dominion by exploring a *finite strict partial order* $\langle \mathrm{St}, s_I, \prec \rangle$. Its elements, called *states*, record information about the quasi dominions computed up to that point. In the *initial state* s_I, the quasi dominions are initialised to the sets of positions with the same priority. At each step, a new quasi α-dominion Q, for some player $\alpha \in \mathbb{B}$, is *extracted* from the current state s and used to compute a *successor state w.r.t.* the order $\prec$, if Q is *open*, *i.e.*, if it is not an α-dominion. If, on the other hand, it is *closed*, then the search is over and Q is added to the portion of the winning region Wn_α computed so far.

We describe here a version of the Priority Promotion algorithm that instantiates the above partial order and serves as a basis for the hybrid approach presented later in this section. To do so, we shall need some technical notions, all of which refer to some fixed parity game $\eth \in \mathrm{PG}$. By $\mathrm{Pr}_\perp \triangleq \mathrm{Pr} \cup \{\perp\}$ and $\mathrm{Pr}_\top \triangleq \mathrm{Pr} \cup \{\top_0, \top_1\}$ we denote the set of priorities in $\eth$ extended with the bottom symbol $\perp$ and two top symbols $\top_0$ and $\top_1$, one for each player. The symbols $\perp$, $\top_0$, and $\top_1$ are called *pseudo priorities*. The standard ordering $<$ on Pr is extended to these additional elements in the natural way: $\perp$ is the smallest

element, while both $\top_0$ and $\top_1$ are strictly greater than every other priority; we do not assume any specific order between the two maximal elements, though, we consider $\top_0$ even and $\top_1$ odd.

The first step in the formalisation of the notion of state requires the concept of *priority-lift function*, which is at the basis of all the approaches discussed in this article. Intuitively, it is a partial function from positions to priorities that over-approximates the priority function of the game.

Definition 2 (Priority-Lift Function)**.** *A priority-lift function* $\mathsf{r} \in \mathrm{Pf} \triangleq \mathrm{Ps} \rightharpoonup \mathrm{Pr}_\top$ *is a partial function such that* $\mathsf{r}(v) \geq \mathsf{pr}(v)$*, for every position* $v \in \mathsf{dom}(\mathsf{r})$.

In the following, we adopt the same notation as in [27]. Given a priority-lift function $\mathsf{r} \in \mathrm{Pf}$ and a priority $p \in \mathrm{Pr}$, we denote with $\mathsf{r}^{(\sim p)}$, for $\sim \in \{<, \leq, \geq, >\}$, the function obtained by restricting the domain of r to the positions $v \in \mathsf{dom}(\mathsf{r})$ whose priority $\mathsf{r}(v)$ satisfies the relation $\mathsf{r}(v) \sim p$, *i.e.*, $\mathsf{r}^{(\sim p)} \triangleq \mathsf{r} \upharpoonright \{v \in \mathsf{dom}(\mathsf{r}) \mid \mathsf{r}(v) \sim p\}$, where $\upharpoonright$ is the standard operation of domain restriction. By $\Delta_\mathsf{r}^\alpha \triangleq \{v \in \mathsf{dom}(\mathsf{r}) \mid \mathsf{r}(v) \equiv_2 \alpha\}$ we denote the set of positions in r with a priority congruent to $\alpha \in \mathbb{B}$, and with $\Delta_\mathsf{r}^{\alpha,p} \triangleq \{v \in \Delta_\mathsf{r}^\alpha \mid \mathsf{r}(v) \geq p\}$ its subset with priorities greater than or equal to p.

Example 2. *Consider again the game reported in Figure 1 and the total function* $\mathsf{r} \triangleq \{\mathsf{a}, \mathsf{c}, \mathsf{e} \mapsto 3; \mathsf{b} \mapsto 7; \mathsf{d}, \mathsf{f}, \mathsf{h} \mapsto 6; \mathsf{g} \mapsto \top_0\}$ *from positions to priorities. It is immediate to see that* r *is a priority-lift function such that* Δ_r^0 *and* Δ_r^1 *denote the dashed blue and red regions, respectively. Moreover,* $\Delta_\mathsf{r}^{0,7} = \{\mathsf{g}\}$*, while* $\Delta_\mathsf{r}^{1,7} = \{\mathsf{b}\}$.

A state encodes information about the quasi dominions computed up to a certain point of the computation. To this end, we require that all the positions in a priority-lift function r with priority of parity $\alpha \in \mathbb{B}$, *i.e.*, the set Δ_r^α, form a quasi α-dominion. Moreover, the idea is to store all α-dominions already identified by associating them with the corresponding pseudo priority $\top_\alpha$.

Definition 3 (Quasi-Dominion Function)**.** *A quasi-dominion function* $\mathsf{r} \in \mathrm{Qs} \subseteq \mathrm{Pf}$ *is a priority-lift function satisfying the following conditions, for every* $\alpha \in \mathbb{B}$*:*

1. *the set* Δ_r^α *is a quasi α-dominion;*

2. *the set* $\mathsf{r}^{-1}(\top_\alpha)$ *is an α-dominion.*

The notion of quasi-dominion function represents the backbone of the algorithm, being the data structure to which the promotion operation is applied.

Example 3. *It is immediate to see that the priority-lift function* r *reported in Example 2 is actually a quasi-dominion function. Indeed, as observed in Example 1, the two sets* $\Delta_\mathsf{r}^0 = \{\mathsf{d}, \mathsf{f}, \mathsf{g}, \mathsf{h}\}$ *and* $\Delta_\mathsf{r}^1 = \{\mathsf{a}, \mathsf{b}, \mathsf{c}, \mathsf{e}\}$ *are quasi dominions for Player* 0 *and* 1*, respectively; in addition,* $\mathsf{r}^{-1}(\top_0) = \{\mathsf{g}\}$ *is a* 0*-dominion. Since, by definition, an empty set is always a dominion,* $\mathsf{r}^{-1}(\top_1) = \emptyset$ *is a* 1*-dominion too, which concludes the verification of the properties required by Definition 3.*

An important property of dominions is that the extension by means of its α-attractor, namely the set of positions obtained by applying the α-attractor operator, of an α-dominion Q is still an α-dominion. This property, however, is not enjoyed by arbitrary quasi dominions. Indeed, there may even be cases where the α-attractor of a quasi α-dominion is a $\overline{\alpha}$-dominion. Moreover, to efficiently verify whether a quasi dominion is actually a dominion, an explicit representation of one of its witnesses is usually required. To overcome these complications, we consider a well-behaving subclass of quasi dominions that meets the following requirements:

1. the set $\mathsf{esc}(Q, \sigma) \triangleq \{v \in \mathrm{Ps}_\alpha \cap Q \,|\, \sigma(v) \notin Q\}$ of α-positions, which leave a quasi α-dominion $Q \subseteq \mathrm{Ps}$ by following one of its α-witnesses $\sigma \in \mathrm{Str}^\alpha(Q)$, is a subset of the $\overline{\alpha}$-escape positions $\mathsf{esc}^{\overline{\alpha}}(Q)$;

2. all these $\overline{\alpha}$-escape positions have priorities congruent to α and greater than the ones of the positions that can be attracted to Q by Player α.

The first requirement states the strategy σ of player α is allowed to leave Q only if it is forced to do so. This ensures that, in order to verify whether Q is an α-dominion, it suffices to check for the emptiness of $\mathsf{esc}^{\overline{\alpha}}(Q)$ (see [26, Definition 1]). The second one, instead, ensures closure under extension by α-attractor. Indeed, since only positions with lower priorities than those of $\mathsf{esc}^{\overline{\alpha}}(Q)$ can be attracted by Q, every cycle induced by the attractor strategy inside Q is necessarily winning for player α (see [24, Proposition 2]). As observed before, a similar structure, which retains weaker properties, is known as snare [36].

Definition 4 (Region Function). *A region function* $\mathsf{r} \in \mathrm{Rg} \subseteq \mathrm{Qs}$ *is a quasi-dominion function satisfying the following conditions, for every* $\alpha \in \mathbb{B}$:

1. *there exists an* α-*witness* $\sigma_\alpha \in \mathrm{Str}^\alpha(\Delta_\mathsf{r}^\alpha)$ *for the quasi* α-*dominion* Δ_r^α *such that* $\mathsf{esc}(\Delta_\mathsf{r}^{\alpha,p}, \sigma_\alpha) \subseteq \mathsf{esc}^{\overline{\alpha}}(\Delta_\mathsf{r}^{\alpha,p})$, *for all priorities* $p \in \mathsf{rng}(\mathsf{r})$, *with* $p \equiv_2 \alpha$;

2. *for all priorities* $p \in \mathsf{rng}(\mathsf{r})$, *with* $p \equiv_2 \alpha$, *and positions* $v \in \mathsf{esc}^{\overline{\alpha}}(\Delta_\mathsf{r}^{\alpha,p})$, *it holds that* $p \leq \mathsf{pr}(v) \equiv_2 \alpha$.

Example 4. *As observed in Example 3, the priority-lift function* r *of Example 2 is a quasi-dominion function. We can now show that it is a region function as well. Indeed, we have (1)* $\Delta_\mathsf{r}^{1,3} = \{\mathsf{a}, \mathsf{b}, \mathsf{c}, \mathsf{e}\}$, *(2)* $\Delta_\mathsf{r}^{0,6} = \{\mathsf{d}, \mathsf{f}, \mathsf{g}, \mathsf{h}\}$, *and (3)* $\Delta_\mathsf{r}^{1,7} = \{\mathsf{b}\}$; *moreover, (1)* $\mathsf{esc}(\Delta_\mathsf{r}^{1,3}, \sigma_1) = \mathsf{esc}^0(\Delta_\mathsf{r}^{1,3}) = \mathsf{esc}(\Delta_\mathsf{r}^{1,7}, \sigma_1) = \mathsf{esc}^0(\Delta_\mathsf{r}^{1,7}) = \{\mathsf{b}\}$ *and (2)* $\emptyset = \mathsf{esc}(\Delta_\mathsf{r}^{0,6}, \sigma_0) \subseteq \mathsf{esc}^1(\Delta_\mathsf{r}^{0,6}) = \{\mathsf{f}\}$, *where* σ_0 *and* σ_1 *are the witnesses of the two quasi dominions highlighted in Figure 1 in blue and red, respectively; moreover,* $\mathsf{pr}(\mathsf{b}) = 7 \equiv_2 1$, *which is greater than or equal to the odd priorities in* r, *namely 3 and 7. Another example of region function is* $\mathsf{r} = \{\mathsf{a} \mapsto 0; \mathsf{c} \mapsto 1; \mathsf{e} \mapsto 3; \mathsf{d}, \mathsf{f}, \mathsf{h} \mapsto 6; \mathsf{b} \mapsto 7; \mathsf{g} \mapsto \top_0\}$, *which is displayed on the left-hand side of Figure 2, and where* $\Delta_\mathsf{r}^{0,0} = \{\mathsf{a}, \mathsf{d}, \mathsf{f}, \mathsf{g}, \mathsf{h}\}$, $\Delta_\mathsf{r}^{1,1} = \{\mathsf{b}, \mathsf{c}, \mathsf{e}\}$, $\Delta_\mathsf{r}^{1,3} = \{\mathsf{b}, \mathsf{e}\}$, $\Delta_\mathsf{r}^{0,6} = \{\mathsf{d}, \mathsf{f}, \mathsf{g}, \mathsf{h}\}$, *and* $\Delta_\mathsf{r}^{1,7} = \{\mathsf{b}\}$. *The easy verification of the properties stated in Definitions 2 and 3 is left to the reader, so we shall only focus on Definition 4. By computing the escape sets for the five priorities* $0, 1, 3, 6, 7 \in \mathsf{rng}(\mathsf{r})$, *we obtain the following, where* σ_0

and σ_1 are two arbitrary witnesses compatible with the partial ones highlighted in the figure: (1) $\{a\} = \mathsf{esc}(\Delta_r^{0,0}, \sigma_0) \subseteq \mathsf{esc}^1(\Delta_r^{0,0}) = \{a, f\}$; (2) $\mathsf{esc}(\Delta_r^{1,1}, \sigma_1) = \mathsf{esc}^0(\Delta_r^{1,1}) = \{b, c\}$; (3) $\{b\} = \mathsf{esc}(\Delta_r^{1,3}, \sigma_1) \subseteq \mathsf{esc}^0(\Delta_r^{1,3}) = \{b, e\}$; (4) $\emptyset = \mathsf{esc}(\Delta_r^{0,6}, \sigma_0) \subseteq \mathsf{esc}^1(\Delta_r^{0,6}) = \{f\}$; (5) $\mathsf{esc}(\Delta_r^{1,7}, \sigma_1) = \mathsf{esc}^0(\Delta_r^{1,7}) = \{b\}$. The first requirement of Definitions 4 is, thus, satisfied. The verification of the second requirement is immediate, once we observe that (1) $\mathsf{pr}(a), \mathsf{pr}(f) \geq 0$, (2) $\mathsf{pr}(b), \mathsf{pr}(c), \mathsf{pr}(e) \geq 1$, (3) $\mathsf{pr}(b), \mathsf{pr}(e) \geq 3$, (4) $\mathsf{pr}(f) \geq 6$, and (5) $\mathsf{pr}(b) \geq 7$.

Note that every set $\Delta_r^{\alpha, p}$, with $p \in \mathsf{rng}(r)$, is a quasi α-dominion, being a subset of the quasi α-dominion Δ_r^α. Also, it is immediate to see that the priority function pr of a given parity game $\eth$ is always a region function. Indeed, it is trivially a priority-lift function. Moreover, the positions in $\Delta_{\mathsf{pr}}^\alpha$, namely those with priorities of parity α, form a quasi α-dominion, whose α-witness can be any strategy that chooses to remain inside the set whenever this is allowed by the move relation. Thus, it is a quasi-dominion function as well. Finally, since $\Delta_{\mathsf{pr}}^{\alpha, p}$ cannot contain positions of parity $\overline{\alpha}$ and thanks to the way the α-witness is chosen, it is clear that pr also satisfies the conditions of Definition 4.

At this point, we have the technical tools to introduce the *search space* that instantiates the (finite) strict partial order mentioned in the intuitive explanation of the approach given above. In particular, to account for the current status of the search of a dominion in a game $\eth$, we define a *state* s as a pair (r, p), comprising a region function r and a priority p, with the idea that

1. all quasi α-dominions computed so far are contained in $\Delta_r^{\alpha, q}$, for some priority q strictly greater than p,

2. the current quasi dominion to focus on is contained in r at priority p, and

3. all positions with priorities smaller than or equal to p correspond to the portion of the game that still has to be processed.

Given a state $s = (r, p)$, we denote by $r_s \triangleq r$ and $p_s \triangleq p$ its two components and by $\alpha_s \triangleq p_s \bmod 2$ the parity of the state, *i.e.*, the player associated to the next quasi dominion to process. Moreover, we set (i) $Q_s \triangleq \Delta_{r_s}^{\alpha_s, p_s}$, (ii) $R_s \triangleq r_s^{-1}(p_s)$, (iii) $L_s \triangleq \{v \in \mathsf{dom}(r_s) \mid r_s(v) \leq p_s\}$, and (iv) $H_s^\alpha \triangleq \Delta_{r_s}^\alpha \setminus L_s$, for all $\alpha \in \mathbb{B}$. We refer to Q_s as *current quasi dominion* and to L_s as the *local area*, *i.e.*, the set of unprocessed positions yet to be analysed (see Item 3 above). Both sets also includes the quasi α_s-dominion R_s on which the next step of the search will focus (see Item 2 above). More precisely, R_s, called α_s-*region* or simply *region* whenever the player is clear from the context, is the intersection of Q_s and L_s. As observed above, the two quasi dominions $\Delta_{r_s}^0$ and $\Delta_{r_s}^1$ partition the entire set of positions in the game, thus, H_s^0 and H_s^1 represent the portions of these quasi dominions already processed, to which the region function has assigned a priority with value strictly greater than p_s (see Item 1 above).

The initial state is composed of the priority function pr of the game and its maximal priority $\mathsf{pr}(\eth)$. Finally, we assume that a state s_1 is smaller than another state s_2 *w.r.t.* the partial order relation $\prec$, if the set of unprocessed positions in s_1 is a subset of those in s_2.

9

Figure 2: A game and a corresponding state representation.

Definition 5 (Search Space). *A search space is a tuple $\mathcal{S} \triangleq \langle \mathrm{St}, s_I, \prec \rangle$, whose three components are defined as follows:*

1. $\mathrm{St} \subseteq \mathrm{Rg} \times \mathrm{Pr}_\perp$ *is the set* states s, *with* $\mathsf{dom}(\mathsf{r}_s) = \mathrm{Ps}$;

2. $s_I \triangleq (\mathsf{pr}, \mathsf{pr}(\ni))$ *is the* initial state;

3. $s_1 \prec s_2$ *if* $\mathrm{L}_{s_1} \subset \mathrm{L}_{s_2}$ *or* $\mathrm{L}_{s_1} = \mathrm{L}_{s_2} = \emptyset$ *and* $p_{s_1} = \perp < p_{s_2}$.

Note that the pseudo priority $\perp$ is used to indicate the situation where all positions have been processed, which corresponds to an empty local area.

Example 5. *To exemplify the notion of state just defined, consider the game of the previous examples, which is reported again on the left-hand side of Figure 2. Let us consider the state $s = (\mathsf{r}, 3)$, where $\mathsf{r} = \{\mathsf{a} \mapsto 0; \mathsf{c} \mapsto 1; \mathsf{e} \mapsto 3; \mathsf{d}, \mathsf{f}, \mathsf{h} \mapsto 6; \mathsf{b} \mapsto 7; \mathsf{g} \mapsto \top_0\}$ is the region function given in Example 4 and $\alpha_s = 1$. The local area L_s contains the positions a, c, and e. Among these positions, only e is part of the current region R_s. This position e, together with b, form the current quasi dominion Q_s. The quasi 0-dominion H_s^0 contains positions d, f, g, and h, while the quasi 1-dominion H_s^1 takes the only remaining position with priority greater than 3, namely b. Position g forms a 0-dominion on its own, represented in the picture by the solid blue line. Apart from this position, all the other ones are contained in open quasi dominions, indicated by dashed lines. For example, the set $\mathsf{r}^{-1}(6) = \{\mathsf{d}, \mathsf{f}, \mathsf{h}\}$ is a quasi 0-dominion. In fact, if Player 1 decides to remain inside, the adversary wins the play. However, Player 1 also has the choice to escape from position f by moving to e. A graphical representation of the state s is given on the right-hand side of Figure 2, where blue-coloured (resp., red-coloured) rectangles correspond to the slots of even (resp., odd) priority in the region function r_s, dashed rectangles indicate that the corresponding set of positions is open and the grey-coloured area denotes the local area L_s.*

During the exploration of the search space, a Priority Promotion algorithm typically traverses several types of states, some of which enjoy important properties that need to be explicitly identified, as they are exploited during the search

10

Algorithm 1: RPP Solver	**Auxiliary Functions & Procedures**
function sol$(s\colon \mathrm{St_M})\colon \mathrm{Rg}$	**function** $\mathsf{NextPr}(s\colon \mathrm{St_S})\colon \mathrm{St_M}$
1 **if** $p_s \neq \bot$ **then**	**1** $q \leftarrow \mathsf{max}\!\left(\mathsf{rng}\!\left(\mathsf{r}_s^{(<p_s)}\right)\right)$
2 $\mathsf{r}_s \leftarrow \mathsf{r}_s[\mathsf{atr}^{\alpha_s}_{\supset_s}(\mathrm{R}_s) \mapsto p_s]$	**2** **return** (r_s, q)
3 **if** s *is closed* **then**	**procedure** $\mathsf{Maximise}(s\colon \mathrm{St})$
4 $\mathsf{Promote}(s)$	**1** **foreach** $\alpha \in \mathbb{B}$ **do**
else	**2** $q \leftarrow \mathsf{min}(\mathsf{rng}(\mathsf{r}_s{\upharpoonright}\mathrm{H}^\alpha_s))$
5 $\mathsf{r}_s \leftarrow \mathbf{sol}(\mathsf{NextPr}(s))$	**3** $\mathrm{X} \leftarrow \mathsf{atr}^\alpha(\mathrm{H}^\alpha_s, \mathrm{L}_s)$
6 **if** s *is closed* **then**	**4** $\mathsf{r}_s \leftarrow \mathsf{r}_s[\mathrm{X} \mapsto q]$
7 $\mathsf{Promote}(s)$	**5** $\mathsf{r}_s \leftarrow \mathsf{r}_s[v \in \mathrm{L}_s \mapsto \mathsf{pr}(v)]$
else	**procedure** $\mathsf{Promote}(s\colon \mathrm{St_P})$
8 $\mathsf{Maximise}(s)$	**1** $q \leftarrow \mathsf{bep}^{\overline{\alpha}_s}(\mathrm{R}_s, \mathsf{r}_s)$
9 $\mathsf{r}_s \leftarrow \mathbf{sol}(s)$	**2** $\mathsf{r}_s \leftarrow \mathsf{r}_s[\mathrm{R}_s \mapsto q]$
10 **return** r_s	

for a dominion. In the following we introduce these types of states and explain what properties they hold. Identify the type of states it is crucial in order to prove the correctness of the approach.

Given a player $\alpha \in \mathbb{B}$, we say that a state $s \in \mathrm{St}$ is α-*maximal* if the quasi α-dominion H^α_s is α-maximal *w.r.t.* L_s. This means that the set of positions α-attracted by H^α_s from the local area L_s is empty, namely, $\mathsf{atr}^\alpha(\mathrm{H}^\alpha_s, \mathrm{L}_s) = \mathrm{H}^\alpha_s$. If s is α-maximal for both players $\alpha \in \mathbb{B}$, we simply say that it is *maximal*. It is useful to denote by $\mathrm{St_M} \subseteq \mathrm{St}$ the corresponding subset of states that are maximal and by $\partial_s \triangleq \partial \setminus (\mathrm{H}^0_s \cup \mathrm{H}^1_s) = \partial \setminus \{v \in \mathsf{dom}(\mathsf{r}_s) \mid \mathsf{r}_s(v) > p_s\}$ the induced subgame. A maximal state s is *strongly maximal* if the current region R_s is α_s-maximal *w.r.t.* L_s as well. We denote by $\mathrm{St_S}$, which is a subset of $\mathrm{St_M}$, the set of strongly maximal states.

Recall that the region R_s of a state s is contained in the current quasi dominion Q_s. We say that s is *open* if the player $\overline{\alpha}_s$ can escape from Q_s starting from R_s using a single move, *i.e.*, if $\mathrm{R}_s \cap \mathsf{esc}^{\overline{\alpha}_s}(\mathrm{Q}_s) \neq \emptyset$. In this case, the opponent may escape from R_s by either moving to the remaining portion of local area $\mathrm{L}_s \setminus \mathrm{R}_s$ or to the quasi $\overline{\alpha}_s$-dominion $\mathrm{H}^{\overline{\alpha}_s}_s$. The state is said to be *closed*, otherwise. For technical convenience, a state with an empty region is always considered open. Finally, a closed state s is *promotable*, if it is $\overline{\alpha}_s$-maximal and R_s is α_s-maximal *w.r.t.* L_s. By $\mathrm{St_P} \subseteq \mathrm{St}$ we denote the set of promotable states. The crucial property of a promotable state s is the following: the only possible moves, if any, that player $\overline{\alpha}_s$ can use to escape from the α_s-region R_s necessarily lead to positions contained in the quasi α_s-dominions in $\mathrm{H}^{\alpha_s}_s$. Indeed, since s is $\overline{\alpha}_s$-maximal, and hence closed, no move from a $\overline{\alpha}_s$-position of $\mathrm{L}_s \supset \mathrm{R}_s$ can lead to $\mathrm{H}^{\overline{\alpha}_s}_s$, which means that all such moves remain inside the current quasi dominion Q_s. This is the property that is leveraged by the promotion operation at the basis of the approach.

The main function **sol**, the auxiliary function NextPr (with no side-effects), and the two procedures Maximise and Promote (with side-effects on the state, *i.e.* the input state is modified in place) of the new Priority Promotion based approach, called *Recursive Priority Promotion* (RPP, for short), are provided in Algorithm 1.

The function **sol** assumes the input state s to be maximal, *i.e.*, $s \in \text{St}_\text{M}$. At Line 1 it checks whether there are still unprocessed positions in the game, that is, if the priority of the current state is different from the pseudo priority $\bot$. If this is the case, Line 2 maximises the region $\text{R}_s = \mathsf{r}_s^{-1}(p_s)$ of the current state by computing its α_s-attractor $\mathsf{atr}_{\circlearrowright_s}^{\alpha_s}(\text{R}_s)$, so that the resulting set is α_s-maximal. Then, the state s is made strongly maximal, by lifting those positions to priority p_s in the region function r_s, so that $s \in \text{St}_\text{S}$ at the end. If the resulting state s is closed (Line 3) it is also promotable, *i.e.*, $s \in \text{St}_\text{P}$, being maximal by hypothesis, and, therefore, a promotion is applied at Line 4 with a call to procedure Promote.

If, instead, s is open, which means that the opponent can escape from R_s by choosing a move exiting from the current quasi dominion Q_s, the algorithm proceeds to analyse the part of the game still unprocessed. To do this, the algorithm first computes the next state by means of NextPr(s), which simply identifies the next priority to consider, namely the maximum priority among the unprocessed positions. The resulting state is then passed to the recursive call at Line 5.

Once the recursive call completes, the state is updated with the new region function returned by the call. The new state s is such that the local area L_s coincides with R_s, since the recursive call terminates after all the previously unprocessed positions have been analysed. As a consequence, either the opponent cannot escape from R_s anymore or it can only move to its own quasi $\overline{\alpha}_s$-dominion $\text{H}_s^{\overline{\alpha}_s}$.

Line 6 checks which one of the two possibilities occurs. In the first case, the new state s is closed, hence $\overline{\alpha}_s$-maximal. Moreover, since $\text{L}_s = \text{R}_s$, the region R_s cannot attract any other positions and is, therefore, α_s-maximal. This means that s is promotable, *i.e.*, $s \in \text{St}_\text{P}$, and Line 7 promotes the region to the quasi α_s-dominion $\text{H}_s^{\alpha_s}$ inside the current quasi dominion Q_s. If, on the other hand, s is still open, then the opponent can escape to $\text{H}_s^{\overline{\alpha}_s}$ from some positions in R_s. This means that the current state is not $\overline{\alpha}_s$-maximal and Line 8 fixes this by calling the procedure Maximise. The aim of this function is to reestablish maximality of the quasi dominions $\text{H}_s^{\alpha_s}$ and $\text{H}_s^{\overline{\alpha}_s}$ associated with the state s. This is done by allowing the processed quasi dominions recorded in the state, namely the ones with priority greater than p_s, to attract positions from the current region $\text{R}_s = \text{L}_s$. The surviving positions in R_s, if any, may not form a quasi α_s-dominion anymore and are reset to the original priorities recorded in the priority function pr of the game. In any case, when the computation reaches Line 9, whether coming from Line 4, Line 7, or Line 8, the state s is maximal, *i.e.*, $s \in \text{St}_\text{M}$, and a second, and final, recursive call is performed on s to process the remaining positions in L_s, if any (when coming from Line 7, L_s is necessarily empty, so the recursive call could be skipped).

The auxiliary function NextPr generates a new maximal state $\widehat{s} = \text{NextPr}(s) \in$

St_M, starting from a strongly-maximal one $s \in St_S$, that is lower in the partial order *w.r.t.* the one in input, *i.e.*, $\widehat{s} \prec s$. The state $\widehat{s}$ is obtained by changing the current priority p_s to the highest priority q of the positions in $L_s \setminus R_s$. Observe that, when no such position exists, namely when $L_s = R_s$, the new priority coincides with the pseudo-priority $\bot$.

Maximise enforces the maximality property on the state s received as input, so that, in the resulting state obtained by modifying s in-place, no position of the local area L_s can be attracted by the quasi α-dominions H_s^α, with $\alpha \in \mathbb{B}$. To this end, the procedure computes (Line 2) the α-attractor $atr^\alpha(H_s^\alpha, L_s)$, collecting all positions of L_s that player α can force to move into the attracting set H_s^α. These positions are then assigned (Line 4) the priority q corresponding to the minimum one associated with a position in the attracting set (Line 3). Since removing positions from the local area L_s may induce a violation of the two requirements of Definition 4, the positions that remain in L_s at the end of the for-each loop of Lines 1-4 need to be reset to their original priority recorded in pr, as prescribed at Line 5.

To conclude, the procedure Promote requires a promotable state $s \in St_P$ and applies a promotion operation to the region R_s, while preserving any maximality property already enjoyed by the input state. It first computes the opponent best-escape priority q for the set R_s *w.r.t.* r_s (Line 1). Intuitively, this is the smallest priority the opponent can reach with one move when escaping from the region R_s. Formally, it is defined as:

$$\mathsf{bep}^{\overline{\alpha}_s}(R_s, r_s) \triangleq \mathsf{min}(\mathsf{rng}(r_s \restriction \mathsf{rng}(I))),$$

where the binary relation $I \triangleq \mathsf{Mv}_\supseteq \cap (\mathsf{esc}^{\overline{\alpha}_s}(R_s) \times (\mathsf{dom}(r_s) \setminus R_s))$ contains all the moves leading outside R_s that the opponent can use to escape. The main property of promotable states observed above ensures that the best-escape priority q has always the same parity as α_s. The procedure, then, promotes R_s to q, by assigning (Line 2) the priority q to all the positions of R_s in the region function r_s. In particular, when the only possibility for player $\overline{\alpha}_s$ to escape from R_s is to reach $r_s^{-1}(\top_\alpha)$, the value of q is $\top_\alpha$. In this case, we are promoting R_s from the status of quasi α_s-dominion to that of α_s-dominion. The correctness of this is ensured by Theorem 1.

At this point, by defining $\mathbf{sol}(\supseteq) \triangleq (r^{-1}(\top_0), r^{-1}(\top_1))$, where $r \triangleq \mathbf{sol}(s_I)$, we obtain a sound and complete solution algorithm for parity games. In particular, the soundness follows from the fact that RPP always traverses states having as invariant the property that $r^{-1}(\top_0)$ and $r^{-1}(\top_1)$ are dominions (see Item 2 of Definition 3). Completeness, instead, is due to the recursive nature of the algorithm, whose base case ensures that no position is left unprocessed at any given priority.

4. Parys' Algorithm

Before obtaining a quasi-polynomial time algorithm based on the priority-promotion approach, we present here a reformulation of the idea of [56], where

the recursion tree of the McNaughton algorithm is suitably truncated in order to avoid the exponential blowup.

Naturally, cutting some of the recursive calls may prevent us from deciding the winner for some of the positions with certainty. These intermediate results may thus contain positions with an *undetermined* status. The contribution of Parys' idea is the design of the truncation mechanism that offers bounded guarantees. Essentially, the undetermined sets contain all small dominions of one player and do not intersect with small dominions of the other. The *bounds* up to which these limited guarantees hold are shed quickly in the call tree: most of the calls are made with half precision, meaning that one of the bounds is halved, and only one is made with full precision, meaning that both bounds are kept. The result is a quasi-polynomial time solution that we will combine with the Priority Promotion approach in the next section, in order to preserve both the complexity of the first and the efficiency of the second.

A first step in the integration of Parys' approach with Priority Promotion is to formulate it in the same terms by introducing the required notation. The data structures are mostly inherited from the previous section. We need to extend the search space to represent also the precision bounds and the set of undetermined positions. While the two bounds can be simply encoded by natural numbers, the undetermined positions will be stored in a priority-lift function u that assigns to each priority $p > 0$ the set of undetermined positions returned by the recursive calls made at priority $p - 1$. A *Parys' state s* is a tuple $((\mathsf{r}, p), (\mathsf{u}, b_0, b_1))$, where the first pair comes from the the search state define in Section 5 and the second tuple contains the function u, collecting the undetermined positions for each recursion level, and the two bounds b_0 and b_1. Recall that $\Delta_{\mathsf{r}}^{\alpha}$ (*resp.*, $\Delta_{\mathsf{u}}^{\alpha}$) collects the positions assigned to some α-priority by the function r (*resp.*, u), while $\Delta_{\mathsf{r}}^{\alpha,p} \subseteq \Delta_{\mathsf{r}}^{\alpha}$ (*resp.*, $\Delta_{\mathsf{u}}^{\alpha,p} \subseteq \Delta_{\mathsf{u}}^{\alpha}$) collects the positions assigned by r (*resp.*, u) to an α-priority of value at least p. For a state $s = ((\mathsf{r}, p), (\mathsf{u}, b_0, b_1))$, we denote with $\mathsf{u}_s \triangleq \mathsf{u}$ the uncertain component, and with $b_{0s} \triangleq b_0$ and $b_{1s} \triangleq b_1$ the bounds of the state. Once again, denoted $\mathsf{r}_s \triangleq \mathsf{r}$ and $p_s \triangleq p$ the first two state components and by $\alpha_s \triangleq p_s \bmod 2$ the parity of the state, the local area $\mathrm{L}_s \triangleq \{v \in \mathsf{dom}(\mathsf{r}_s) \mid \mathsf{r}_s(v) \leq p_s\}$ of a state s contains the positions still unprocessed in that state, while $\mathrm{R}_s \triangleq \mathsf{r}^{-1}(p_s)$ denotes the region of s. Then, we set $\mathrm{Q}_s \triangleq \Delta_{\mathsf{r}_s}^{\alpha_s, p_s} \cup \Delta_{\mathsf{u}_s}^{\overline{\alpha}_s, p_s}$ and $\mathrm{H}_s^{\alpha} \triangleq (\Delta_{\mathsf{r}_s}^{\alpha} \cup \Delta_{\mathsf{u}_s}^{\overline{\alpha}}) \setminus \mathrm{L}_s$, for all $\alpha \in \mathbb{B}$. Essentially, H_s^0 and H_s^1 represent the portion of the game already processed to which the two functions r_s and u_s have assigned a priority with value strictly greater than p_s.

As opposed to the notions of RPP state, however, H_s^0 and H_s^1 are not necessarily quasi dominions, since they may include undetermined positions, namely those contained in $\Delta_{\mathsf{u}_s}^{0, p_s}$ and $\Delta_{\mathsf{u}_s}^{1, p_s}$. Only the subsets $\Delta_{\mathsf{u}_s}^{0, q}$ and $\Delta_{\mathsf{r}_s}^{1, q}$ are known to be quasi dominions, for all priorities q. The current subgame of a state s is, then, $\partial_s \triangleq \partial \setminus (\mathrm{H}_s^0 \cup \mathrm{H}_s^1)$, containing all the unprocessed positions. In addition, $\mathrm{U}_s \triangleq \mathsf{u}_s^{-1}(p_s)$ collects the currently undetermined set of positions, which is required to satisfy the following property: it must contain all the $\overline{\alpha}_s$-dominions of size no greater than a given bound $b_{\overline{\alpha}_s}$ and it cannot intersect any α_s-dominion of size no greater than a second given bound b_{α_s}.

Definition 6 (Parys' Search Space)**.** *A* Parys' search space *is a tuple* $\mathcal{S} \triangleq \langle \text{St}, s_I, \prec \rangle$, *whose three components are defined as follows:*

1. $\text{St} \subseteq (\text{Rg} \times \text{Pr}_\perp) \times (\text{Pf} \times \mathbb{N} \times \mathbb{N})$ *is the set of* Parys' states s *where:*

   a) $\text{dom}(r_s) \cap \text{dom}(u_s) = \emptyset$, $\text{dom}(r_s) \cup \text{dom}(u_s) = \text{Ps}$;

   b) $\{v \in \text{dom}(u_s) \mid u_s(v) < p_s\} = \emptyset$ *and* $u_s^{-1}(\top_\alpha) = \emptyset$, *where* $\alpha \triangleq \text{pr}(\eth) \bmod 2$;

   c) $\text{dom}(u_s) \cap \text{esc}^{\overline{\alpha}}(\text{H}_s^\alpha) = \emptyset$, *for all* $\alpha \in \mathbb{B}$;

2. $s_I \triangleq ((\text{pr}, \text{pr}(\eth)), (\varnothing, |\eth|, |\eth|))$ *is the* initial state;

3. $s_1 \prec s_2$ *if* $\text{L}_{s_1} \subset \text{L}_{s_2}$ *or* $\text{L}_{s_1} = \text{L}_{s_2} = \emptyset$ *and* $p_{s_1} = \perp < p_{s_2}$.

Item 1a ensures that the set of positions in the game is partitioned into two categories: those contained in the region function r_s, which are considered *determined*, in the sense that they belong to known quasi dominions; and those contained in the priority-lift function u_s, which are *undetermined*, since they form sets that only satisfy the bounded guarantees. Item 1b guarantees that there cannot be undermined positions at priorities lower than p_s and that the pseudo priority $\top_{\alpha_s}$ can never contain any undetermined position. These conditions capture the intuitions that all positions with priorities lower than the current one are still unprocessed, therefore they cannot be undetermined and that positions assigned to a top pseudo-priority must be determined as they are won by the corresponding player. Item 1c, instead, captures a more technical aspect of Parys' approach. Essentially, it requires that the player $\overline{\alpha}$ cannot escape to L_s starting from an undetermined position at any α-priority recursion level. Intuitively, the undetermined positions at any given level are formed by the subgame of recursive sub-calls with bound equal to one (depleted). Thus, they form closed sets, in the terminology of the previous section, from where the opponent cannot escape to lower priorities. For this reason, we shall refer to this property also as *closed under attraction*. The initial state is composed of the pair of the priority function pr and the maximal priority $\text{pr}(\eth)$ of the game, and the tuple with the empty priority-lift function $\varnothing$, and the two (full) bounds that corresponding to number of positions in the game. Finally, the order relation $\prec$, is defined the same way as done for the state space for Priority Promotion.

Example 6. *Figure 3 depicts an example game on the left-side and, on the right-side, the corresponding graphical representation of one of the states computed during the execution of Parys' algohrithm, whose interpretation is similar to that of Figure 2 and where, in addition, we indicate with pattern-filled rectangles the slots of the undetermined function u_s and the undetermined region U_s. Let us consider the state $s = ((r, 1), (u, 2, 2))$, where $r = \{b \mapsto 7; c \mapsto 6; n \mapsto 5; d \mapsto 4; m \mapsto 3; g, i \mapsto 1\}$ is the region function, $u = \{a, e, f \mapsto 2\}$ is the undetermined function, and $\alpha_s = 1$. The current priority p_s is 1 and the two bounds are $b_{0s} = 2$ and $b_{1s} = 2$. The local area L_s contains the positions g, h, i, and l. Among these, g and i are part of the current region R_s. These positions, together with $\Delta_u^1 = \{a, e, f\}$, and b, n, m, form the current set Q_s. The quasi 0-dominion*
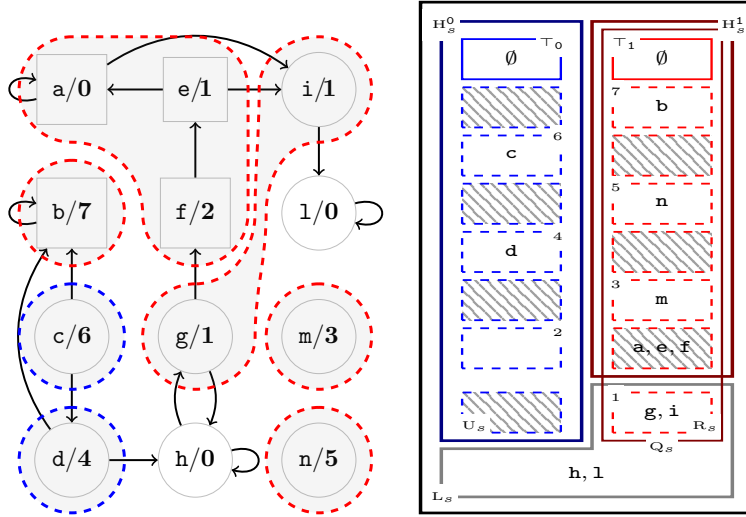
Figure 3: A game and a corresponding state representation.

H_s^0 contains positions d, and c, which are both part of the region function r. The quasi 1-dominion H_s^1, on the other hand, takes the remaining positions with priority greater than 1, namely positions b, n and m, stored in the region function r, and positions a, e and f, which are stored in the uncertain function u.

Just like RPP does not use undetermined positions, Parys' approach does not use promotions, even if it handles regions (by means of the attractor of the positions with highest priority). Consequently, little happens to the region function in our representation of Parys' algorithm: positions that are added to U_s are removed from the region function, and when U_s is destroyed, they are added (with their native priorities) back to r_s.

Since the purpose of this section is to revisit Parys' algorithm, here we only outline the principles, by reporting some of the standard lemmas from [59] that will be used later on also by the hybrid approach.

The first lemma simply states that dominions are closed under opponent attraction from the complement set of its positions.

Lemma 1. *Let* D *be a dominion for player* α *and* S $\subseteq$ Ps *a set of positions. If* D *does not intersect with* S *then it does not intersect with* $\mathsf{atr}^{\overline{\alpha}}(S)$ *either.*

The second specifies that in case a dominion includes positions with a higher opponent priority, then it also has to include a smaller dominion closed under opponent attraction from the set of positions with the higher opponent priority.

Lemma 2. *Let* D *be a dominion for player* α *with highest priority* p *and* R *a* $\overline{\alpha}$*-region of priority* q *with* $q > p$. *If* D *intersects with* R *then it contains a non-empty sub-dominion that does not intersect with* $\mathsf{atr}^{\overline{\alpha}}(R)$.

16

Algorithm 2: Parys Solver	Half-Solver
function sol$(s \colon \mathrm{St}) \colon 2^{\mathrm{Ps}}$	**procedure hsol**$(s \colon \mathrm{St})$

	Algorithm 2: Parys Solver
	function sol$(s\colon \mathrm{St})\colon 2^{\mathrm{Ps}}$
1	**if** $p_s = \bot \vee b_{\alpha_s} = 1$ **then**
2	**return** $(\mathsf{r}_s, \mathsf{u}_s)$
	else
3	$\mathbf{hsol}(s)$
4	$\widehat{s} \leftarrow s$
5	$(\mathsf{r}_s, \mathsf{u}_s) \leftarrow \mathbf{sol}(\mathsf{NextPr}(s))$
6	$\mathsf{Maximise}(s)$
7	**if** $s \prec \widehat{s}$ **then** $\mathbf{hsol}(s)$
8	**return** $\mathsf{Und}(s)$

	Half-Solver
	procedure hsol$(s\colon \mathrm{St})$
1	**repeat**
2	$\mathsf{r}_s \leftarrow \mathsf{r}_s[\mathsf{atr}_{\supseteq_s}^{\alpha_s}(\mathrm{R}_s) \mapsto p_s]$
3	$\widehat{s} \leftarrow s$
4	$(\mathsf{r}_s, \mathsf{u}_s) \leftarrow \mathbf{sol}(\mathsf{Half}(s))$
5	$\mathsf{Maximise}(s)$
	until $s \not\prec \widehat{s}$

	Auxiliary Functions & Procedures
	function $\mathsf{NextPr}(s\colon \mathrm{St})\colon \mathrm{St}$
1	**return** $((\mathsf{r}_s, p_s - 1), (\mathsf{u}_s, p_s, b_{0s}, b_{1s}))$
	function $\mathsf{Half}(s\colon \mathrm{St})\colon \mathrm{St}$
1	$(b_{0s}, b_{1s}) \leftarrow \left(\left\lfloor \frac{b_{0s}}{1 + \alpha_s} \right\rfloor, \left\lfloor \frac{b_{1s}}{2 - \alpha_s} \right\rfloor \right)$
2	**return** $\mathsf{NextPr}(s)$

	Auxiliary Functions & Procedures
	procedure $\mathsf{Maximise}(s\colon \mathrm{St})$
1	$\mathrm{X} \leftarrow \mathsf{atr}^{\overline{\alpha}_s}(\mathrm{H}_s^{\overline{\alpha}_s}, \mathrm{L}_s)$
2	$\mathsf{r}_s \leftarrow \mathsf{r}_s \setminus \mathrm{X}$
3	$\mathsf{u}_s \leftarrow \mathsf{u}_s[\mathrm{X} \mapsto p_s]$
4	**if** $\mathrm{X} \neq \emptyset$ **then**
5	$\mathsf{r}_s \leftarrow \mathsf{r}_s[v \in \mathrm{R}_s \mapsto \mathsf{pr}(v)]$
	function $\mathsf{Und}(s\colon \mathrm{St}) \colon \mathrm{Rg} \times \mathrm{Pf}$
1	$\mathsf{u}_s \leftarrow \mathsf{u}_s[\mathrm{L}_s \mapsto p_s + 1]$
2	$\mathsf{r}_s \leftarrow \mathsf{r}_s^{(>p_s)}[v \in \mathrm{U}_s \mapsto \mathsf{pr}(v)]$
3	**return** $(\mathsf{r}_s, \mathsf{u}_s)$

The main algorithm is composed of a full bound precision function **sol** and an halved precision procedure **hsol** that are reported in the upper part of Algorithm 2. The auxiliary procedure $\mathsf{Maximise}$ and the auxiliary functions NextPr, Half, and Und are reported in the lower part.

In order to outline the structure of function **sol**, it is helpful to first describe the halved precision procedure **hsol**. The half-solver **hsol** maximises, at Line 2, the region $\mathrm{R}_s = \mathsf{r}_s^{-1}(p_s)$ of the current state by computing its α_s-attractor $\mathsf{atr}_{\supseteq_s}^{\alpha_s}(\mathrm{R}_s)$. The remaining subgame from which R_s has been removed, is recursively solved by a call of **sol** with half precision by means of the function Half that halves the bound b_{α_s} for the current player α_s. As long as the maximisation of the state returned by the call to **sol** at Line 4 differs from the copy of the local state saved at Line 3, the procedure will repeat the loop of Lines 2-5. Note that in all these iterations the priority p_s remains the same; therefore, the positions in current region R_s have to change at every iteration (and they can only do that by shrinking), otherwise the repeat-until loop will stop. Then this loop ensures that every small $\overline{\alpha}_s$-dominion, those of size at most $\left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor$ in the subgame that are not included in small α_s-regions of size at most b_{α_s}, are found and returned in the uncertain function u_s. The observation above is a consequence of Lemmas 1 and 2. If such a small $\overline{\alpha}_s$-dominion does not intersect with the current α_s-region R_s, then by Lemma 1, it is included in the remaining subgame, and identified by the call of **sol**. If, instead, it intersects with R_s, then, by Lemma 2, there

must be a sub-dominion that cannot be attracted by R_s and, therefore, will be identified by the next call to **sol**.

The function **sol** has a base case at Line 1, where it checks if there are unprocessed positions in the current state (which implies that $p_s > \perp$) or the bound for player α_s is not depleted. The function can perform three recursive calls, two of the half-solver, and one of **sol** itself. The first call to **hsol**, at Line 3, serves to find small dominions in the subgame, while the call to **sol**, at Line 5, finds big dominions that the previous call might have been unable to find due to the bound constraint. If the full precision call has not found a $\overline{\alpha}_s$-dominion, then the subgame does not contain any dominion for the bound available. In this case a second call to **hsol** would be pointless as **hsol** has a smaller bound than **sol**. Therefore, the function ends by calling Und, which records the entire current subgame L_s as the undetermined set for the caller at level $p_s + 1$. Otherwise, when the full precision call returns a new $\overline{\alpha}_s$-dominion D, there might be other small $\overline{\alpha}_s$-dominions in the subgame to find. This follows for the fact that D is bigger than $\left\lfloor \frac{\mathsf{b}_{\overline{\alpha}_s}}{2} \right\rfloor$ and it has been closed under $\overline{\alpha}_s$-attractor by the call to Maximise at Line 6. This operation can reset R_s that has been identified by **hsol** and possibly release a small $\overline{\alpha}_s$-dominion previously included in R_s. For this reason a second and final call to **hsol** is performed to identify such small $\overline{\alpha}_s$-dominions. This case occur only when the updated state returned at Line 5 differs from the copy of the local state saved at Line 4. At the end of the function, when Und is called at Line 8, all $\overline{\alpha}_s$-dominions of size at most $\mathsf{b}_{\overline{\alpha}_s}$ have been identified and processed in U_s. Therefore, none of them intersects with L_s that instead contains the bounded winning region of player α_s. At this point, since the guarantees for the winning α_s-region are bounded, L_s is moved to the uncertain set U of the caller priority, that is $p_s + 1$.

The auxiliary function NextPr generates a new maximal state by decreasing the current priority by 1. Also Half generates a new maximal state, by first halving the precision bound for player α_s and then calling NextPr.

As in the case of the RPP solver, Maximise enforces the maximality property on the current state s. As a result, in the modifies state s, no positions of the local area L_s can be attracted by the quasi $\overline{\alpha}_s$-dominion $H_s^{\overline{\alpha}_s}$. To obtain maximality, the procedure computes at Line 1 the $\overline{\alpha}_s$-attractor $\mathsf{atr}^{\overline{\alpha}_s}\!\left(H_s^{\overline{\alpha}_s}, L_s\right)$, which collects all the position in L_s that player $\overline{\alpha}_s$ can force to move into $H_s^{\overline{\alpha}_s}$. The attracted set is then removed from r_s at Line 2 and assigned at Line 3 to u_s at the current priority p_s. Attracting positions from L_s may result in the violation of the requirements of Definition 4 for the positions that remain in R_s. For this reason, if $H_s^{\overline{\alpha}_s}$ does attract some positions from L_s, the rest of R_s is reset, *i.e.* the priority of its positions is set to their original priority in pr, as prescribed by Line 5. Note that the remaining positions of L_s need not be reset, since their priorities are already those recorded in pr.

Finally, the function Und takes care of handling the undetermined positions computed at the end of a call to the solver. It receives in input a state where all $\overline{\alpha}_s$-dominions of size at most $\mathsf{b}_{\overline{\alpha}_s}$ (that are not included in α_s-regions of size at most b_{α_s}) have been moved into U_s. It follows that L_s is a bounded winning

18

region for player α_s.

To explain how Und works, let s be the current state and s' the one of the caller priority level $p_s + 1$, where, in case $p_s = \mathsf{pr}(\partial)$, then $p_s + 1$ denotes the top pseudo-priority $\top_{\overline{\alpha}}$, with $\mathsf{pr}(\partial) \equiv_2 \alpha$. Since Und is the last operation executed before the return to the call corresponding to state s', the bounded winning α_s-region L_s at the current priority level p_s is added to the uncertain winning set of positions for player $\overline{\alpha}_s$ in s', that is $\mathrm{U}_{s'}$ (Line 1). The positions of U_s are instead reset at Line 2 and moved into r_s at their original priority. Therefore, these positions will be part of the unprocessed partition of the game of s', that is $\mathrm{L}_{s'}$. Clearly, no new position added to the undetermined function can be forced to move to $\mathrm{L}_{s'}$ by means of an opponent attractor operation in accordance with Item 1c of Definition 6 applied to the new state s'.

When $\mathbf{sol}(s_I)$ terminates by calling Und, it returns the pair composed of r_{s_I} and u_{s_I}, and the two winning regions correspond to the sets $\mathrm{Wn}_{\partial}^{\alpha} = \mathsf{u}_{s_I}^{-1}(\top_{\overline{\alpha}})$ and $\mathrm{Wn}_{\partial}^{\overline{\alpha}} = \partial \setminus \mathrm{Wn}_{\partial}^{\alpha}$.

Corollary 1. *Let* (r, u) *be the pair of region and promotion functions computed by Algorithm 2 on a game* ∂. *The winning region for player* α *corresponds to* $\mathsf{u}^{-1}(\top_{\overline{\alpha}})$. $\qquad\qquad\square$

We provided a different formalisation of Parys' algorithm by means of the pair of region and promotion functions. The correctness corollary can be found, in a different formulation, in the original paper [56].

5. A Hybrid Algorithm

In our hybrid approach, we have to synthesise the use of regions employed in the RPP solver and the use of undetermined sets introduced to implement Parys' approach. To formalise this intuition into a hybrid state, we need to further modify the simple state of RPP *w.r.t.* the changes already introduced in the Parys' search space of Definition 6. The good news is that the data structure of a Parys' state stores almost all the information required by the hybrid algorithm. It follows that, syntactically, the hybrid state corresponds to the Parys' one with an additional priority and, semantically, it satisfy an additional property.

The most relevant difference between the two approaches is that, unlike Parys' algorithm, the hybrid algorithm does not explore sequences of states with consecutive priorities. Instead, like the RPP approach, it can skip priorities that do not occur in the current subgame. To account for the gap between the priority of the parent recursive call and the current one associated with the state, we introduce in the state an additional component c that records the caller priority. Clearly, since no recursion levels are present between the caller and the current priorities, no position can be recorded between those two levels in either of the two functions r and u of that state. This follows form the fact that the priority selected by the caller for the next sub-call is the maximum one existing in the subgame. Therefore, both r and u will not contain positions with priority lower that the caller one and higher than the sub-caller one.

A *hybrid state* s is a tuple $((\mathsf{r}, p), (\mathsf{u}, b_0, b_1), c)$, which contains all the components of the Parys' state, namely, the region function r, the current priority p, the undetermined function u, and the two bounds b_0 and b_1, one for each player, as well as the additional priority c of the caller. We set $c_s \triangleq c$, while the elements α_s, p_s, b_{α_s}, r_s, u_s, Q_s, H_s^α, U_s, R_s, and L_s are defined as in the previous section. The caller priority of the initial state is set to the pseudo priority $\top_{\overline{\alpha_s}}$.

Definition 7 (Hybrid Search Space). *A* hybrid search space *is a tuple* $\mathcal{S} \triangleq \langle \mathrm{St}, s_I, \prec \rangle$, *whose three components are defined as follows:*

1. $\mathrm{St} \subseteq (\mathrm{Rg} \times \mathrm{Pr}_\perp) \times (\mathrm{Pf} \times \mathbb{N} \times \mathbb{N}) \times \mathrm{Pr}_\top$ *is the set of* hybrid states s *where:*

    a) $\mathsf{dom}(\mathsf{r}_s) \cap \mathsf{dom}(\mathsf{u}_s) = \emptyset$, $\mathsf{dom}(\mathsf{r}_s) \cup \mathsf{dom}(\mathsf{u}_s) = \mathrm{Ps}$;

    b) $\{v \in \mathsf{dom}(\mathsf{u}_s) \mid \mathsf{u}_s(v) < p_s\} = \emptyset$ *and* $\mathsf{u}_s^{-1}(\top_\alpha) = \emptyset$, *where* $\alpha \triangleq \mathsf{pr}(\eth) \bmod 2$;

    c) $\mathsf{dom}(\mathsf{u}_s) \cap \mathsf{esc}^{\overline{\alpha}}(\mathrm{H}_s^\alpha) = \emptyset$, *for all* $\alpha \in \mathbb{B}$;

    d) $\{v \in \mathsf{dom}(\mathsf{f}) \mid p_s < \mathsf{f}(v) < c_s\} = \emptyset$, *for* $\mathsf{f} \in \{\mathsf{r}_s, \mathsf{u}_s\}$, *and* $p_s < c_s$;

2. $s_I \triangleq ((\mathsf{pr}, \mathsf{pr}(\eth)), (\varnothing, |\eth|, |\eth|), \top_{\overline{\alpha}})$ *is the* initial state, *where* $\alpha \triangleq \mathsf{pr}(\eth) \bmod 2$;

3. $s_1 \prec s_2$ *if* $\mathrm{L}_{s_1} \subset \mathrm{L}_{s_2}$ *or* $\mathrm{L}_{s_1} = \mathrm{L}_{s_2} = \emptyset$ *and* $p_{s_1} = \perp < p_{s_2}$.

The only new condition *w.r.t.* Definition 6 is Item 1d, which requires that no position be contained in either r_s or u_s at priorities between p_s, the current priority, and c_s, the caller one.

Most of the concepts and notation introduced in a Parys' state have a similar meaning and play a similar role in a hybrid state. In particular, given a hybrid state $s \in \mathrm{St}$, the set L_s identifies the *local area*, *i.e.*, the set of positions yet to analyse, while R_s is the quasi α_s-dominion, called *region*, included in L_s, which the algorithm is currently focusing on. Moreover, the two sets H_s^0 and H_s^1 partition the portion of game already processed.

Example 7. *Figure 4 depicts a snapshot of the execution of the hybrid algorithm on the example game described in Example 6. Again, the left-side depicts the game, while the right-side provides a graphical representation of the hybrid state in the considered snapshot with the same graphical conventions adopted in Figure 3. The hybrid state s corresponds to the tuple $((\mathsf{r}, 1), (\mathsf{u}, 2, 1), 1)$, where the first two components are region function $\mathsf{r} = \{\mathsf{i}, \mathsf{l} \mapsto \top_0; \mathsf{b} \mapsto \top_1; \mathsf{c} \mapsto 6; \mathsf{n} \mapsto 5; \mathsf{d} \mapsto 4; \mathsf{n} \mapsto 3; \mathsf{g} \mapsto 1\}$ and the current priority $p = 1$. It follows that the parity of the state is $\alpha_s = 1$. In the second part of the state we have $\mathsf{u} = \{\mathsf{a}, \mathsf{e}, \mathsf{f} \mapsto 2\}$ that is the undetermined function, the two bounds b_0 and b_1, that are both 1, and the caller priority c set to 2. The local area L_s, represented in grey at the bottom of the picture, contains the positions g, and h. Among these, only g is part of the current region R_s. This position, together with $\Delta_\mathsf{u}^1 = \{\mathsf{a}, \mathsf{e}, \mathsf{f}\}$ and $\Delta_\mathsf{r}^1 = \{\mathsf{b}, \mathsf{n}, \mathsf{m}\}$ forms the current quasi dominion Q_s represented by the right column in red of the picture. The quasi 0-dominion H_s^0, represented in blue, contains positions $\mathsf{i}, \mathsf{l}, \mathsf{c}$, and d, that are all stored in the region function r, while the quasi 1-dominion H_s^1 takes the remaining positions with priority greater than 1. The positions stored*
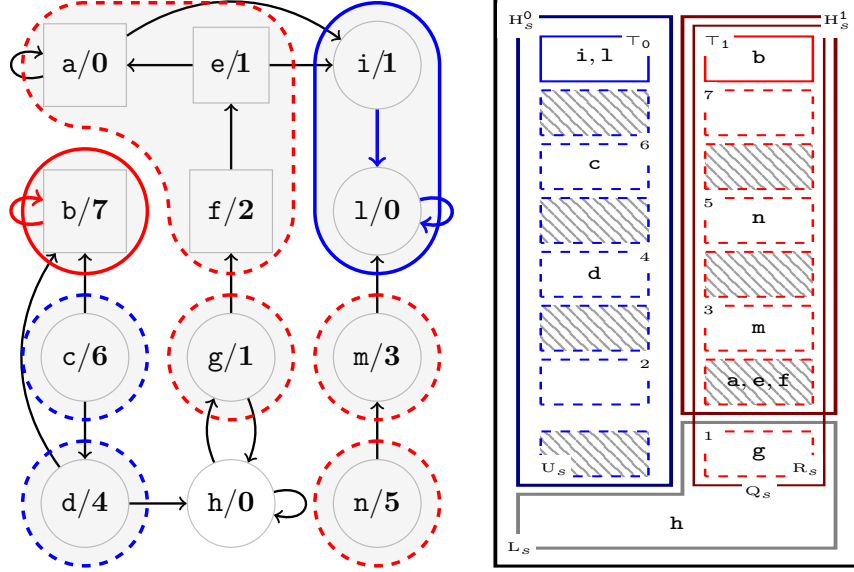
Figure 4: A game and a corresponding state representation.

in r *with value* $\top_0$ *or* $\top_1$ *are dominion already identified, and therefore solved by the algorithm. For this reason, in the picture, these positions are represented with non-dashed lines and their winning strategy is highlighted with coloured edges.*

It is worth noticing that Parys' algorithm reaches the state depicted in Example 6 *regardless the two positions* m, *and* n*. On the contrary, for the hybrid algorithm, these two positions are required to deplete the precision bound. Without* m *and* n*, the hybrid algorithm can easily solve the game thanks to the skip of the priorities that do not occur in the subgame. It follows that whenever a game has priorities that are not consecutive, the hybrid algorithm solves it easier than Parys' approach.*

Given a player $\alpha \in \mathbb{B}$, we say that a hybrid state $s \in \mathrm{St}$ is α-*maximal*, if the quasi α-dominion H_s^α is α-maximal *w.r.t.* L_s. If s is α-maximal *w.r.t.* both players $\alpha \in \mathbb{B}$, we say that it is *maximal*. We denote with $\mathrm{St}_{\mathtt{M}} \subseteq \mathrm{St}$ the set of maximal hybrid states and with $\supset_s \triangleq \supset \setminus (\mathrm{H}_s^0 \cup \mathrm{H}_s^1)$ the induced subgame of s. A maximal hybrid state s is *strongly maximal*, if the current region R_s is α_s-maximal *w.r.t.* L_s. With $\mathrm{St}_{\mathtt{S}} \subseteq \mathrm{St}_{\mathtt{M}}$ we denote the set of strongly maximal hybrid states. Again, we say that s is *open* if $\mathrm{R}_s \cap \mathsf{esc}^{\overline{\alpha}_s}(\mathrm{Q}_s) \neq \emptyset$, and that it is *closed*, otherwise. For technical convenience, we consider open a hybrid state whose region R_s is empty. Finally, a closed hybrid state s is *promotable*, if it is $\overline{\alpha}_s$-maximal and R_s is α_s-maximal *w.r.t.* L_s. With $\mathrm{St}_{\mathtt{P}} \subseteq \mathrm{St}$ we denote the set of promotable hybrid states.

The main functions and the auxiliaries procedures of the *Hybrid Priority Promotion algorithm* (HPP, for short), which are provided in Algorithm 5, synthesise the recursive Priority Promotion technique of Algorithm 1 and the

recursion-tree truncation idea of Algorithm 2 into a single approach.

Algorithm 5: HPP Solver	**Algorithm 6:** Half-Solver

function sol(s: $\mathrm{St_M}$): $\mathrm{Rg} \times \mathrm{Pf}$

1   **if** $p_s = \perp \vee b_\alpha = 0$ **then**
2   **return** $(\mathsf{r}_s, \mathsf{u}_s)$
   else
3   **hsol**(s)
4     $\widehat{s} \leftarrow s$
5     $(\mathsf{r}_s, \mathsf{u}_s) \leftarrow \mathbf{sol}(\mathsf{NextPr}(s))$
6   **if** s *is open* **then**
7      $\mathsf{Maximise}(s)$
    else
8      $\mathsf{Promote}(s)$
9   **if** $s \prec \widehat{s}$ **then** **hsol**(s)
10   **return** $\mathsf{Und}(s)$

procedure hsol(s: $\mathrm{St_M}$)

1   **repeat**
2     $\mathsf{r}_s \leftarrow \mathsf{r}_s[\mathsf{atr}_{\supseteq_s}^{\alpha_s}(\mathsf{R}_s) \mapsto p_s]$
3     $\widehat{s} \leftarrow s$
4   **if** s *is open* **then**
5      $(\mathsf{r}_s, \mathsf{u}_s) \leftarrow \mathbf{sol}(\mathsf{Half}(s))$
6   **if** s *is open* **then**
7       $\mathsf{Maximise}(s)$
     else
8       $\mathsf{Promote}(s)$
    else
9      $\mathsf{Promote}(s)$
  until $s \not\prec \widehat{s}$

Auxiliary Functions	Auxiliary Procedures

procedure Promote(s: $\mathrm{St_P}$)

1   $(p_\mathsf{r}, p_\mathsf{u}) \leftarrow (\mathsf{bep}^{\overline{\alpha}_s}(\mathsf{R}_s, \mathsf{r}_s), \mathsf{bep}^{\overline{\alpha}_s}(\mathsf{R}_s, \mathsf{u}_s))$
2   **if** $p_\mathsf{r} \leq p_\mathsf{u}$ **then**
3     $\mathsf{r}_s \leftarrow \mathsf{r}_s[\mathsf{R}_s \mapsto p_\mathsf{r}]$
   else
4     $(\mathsf{r}_s, \mathsf{u}_s) \leftarrow (\mathsf{r}_s \backslash \mathsf{R}_s, \mathsf{u}_s[\mathsf{R}_s \mapsto p_\mathsf{u}])$

function NextPr(s: $\mathrm{St_S}$): $\mathrm{St_M}$

1   $q \leftarrow \mathsf{max}\left(\mathsf{rng}\left(\mathsf{r}_s^{(<p_s)}\right)\right)$
2   **return** $((\mathsf{r}_s, q), (\mathsf{u}_s, p_s, b_{0s}, b_{1s}))$

function Half(s: $\mathrm{St_S}$): $\mathrm{St_M}$

1   $(b_{0s}, b_{1s}) \leftarrow \left(\left\lfloor \frac{b_{0s}}{1+\alpha_s} \right\rfloor, \left\lfloor \frac{b_{1s}}{2-\alpha_s} \right\rfloor\right)$
2   **return** $\mathsf{NextPr}(s)$

procedure Maximise(s: St)

1   $\mathrm{Z} \leftarrow \mathsf{R}_s$
2   **foreach** $\alpha \in \mathbb{B}$ **do**
3     $\mathrm{X} \leftarrow \mathsf{atr}^\alpha(\mathrm{H}_s^\alpha, \mathrm{L}_s)$
4     $q \leftarrow \mathsf{min}(\mathsf{rng}((\mathsf{r}_s \cup \mathsf{u}_s) \upharpoonright \mathrm{H}_s^\alpha))$
5   **if** $q \equiv_2 \alpha$ **then**
6      $\mathsf{r}_s \leftarrow \mathsf{r}_s[\mathrm{X} \mapsto q]$
7      $\mathsf{u}_s \leftarrow \mathsf{u}_s \backslash \mathrm{X}$
    else
8      $\mathsf{r}_s \leftarrow \mathsf{r}_s \backslash \mathrm{X}$
9      $\mathsf{u}_s \leftarrow \mathsf{u}_s[\mathrm{X} \mapsto q]$
10   **if** $\mathrm{Z} \neq \mathsf{R}_s$ **then**
11     $\mathsf{r}_s \leftarrow \mathsf{r}_s[v \in \mathrm{R}_s \mapsto \mathsf{pr}(v)]$

function Und(s: $\mathrm{St_S}$): $\mathrm{Rg} \times \mathrm{Pf}$

1   **if** $c_s \equiv_2 \alpha_s$ **then**
2     $\mathsf{u}_s \leftarrow \mathsf{u}_s[\mathrm{U}_s \mapsto c_s]$
   else
3     $\mathsf{u}_s \leftarrow \mathsf{u}_s^{(\geq c_s)}[\mathrm{L}_s \mapsto c_s]$
4     $\mathsf{r}_s \leftarrow \mathsf{r}_s^{(\geq c_s)}[v \in \mathrm{U}_s \mapsto \mathsf{pr}(v)]$
5   **return** $(\mathsf{r}_s, \mathsf{u}_s)$

Like RPP, the main function **sol** assumes the input state s to be maximal, *i.e.*, $s \in \mathrm{St_M}$. Line 1 checks whether (i) there are no unprocessed positions in the game or (ii) the bound for player $\alpha[s]$ for the guarantees on the undetermined positions has reached the threshold zero. If one of these conditions is satisfied,

the current region function r_s and the undetermined function u_s are returned unmodified at Line 2, as no further progress can be obtained in the current recursive call. Otherwise, similar to Parys' approach, the search for a dominion is split into three phases: (i) a first search with halved precision made by calling the auxiliary mutually-recursive procedure **hsol** (Line 3); (ii) a second search with full precision with a recursive call to **sol** itself (Lines 4 to 8); (iii) a final search by means of **hsol**, again with halved precision, conditioned to the actual progress obtained during the previous phase (Line 9). On termination of these three phases, the information about the undetermined positions contained in the local area L_s or in the undetermined set U_s are handled by the function Und at Line 10.

To discuss the guarantees and their effects in more detail, let us fix a small $\overline{\alpha}$-dominion D, with $|D| \le b_{\overline{\alpha}_s}$. Then, the call to **hsol** at Line 3 modifies the maximal state s given as input in-place, turning it into a strongly-maximal one so that L_s no longer contains any tiny dominions of player $\overline{\alpha}_s$ of size less than or equal to $\lfloor \frac{b_{\overline{\alpha}_s}}{2} \rfloor$. It follows that **hsol** might have not enough precision bound to guarantee that D is found. Therefore, what eventually remains is a set $D' = D \cap L_s$ that is a dominion in L_s by Lemma 1. The set $D \setminus D'$, has been instead processed and added either to u_s at some priority $q \ge p$ or to r_s at some priority $q > p$. Moreover, by Lemma 2, D' has a non-empty sub-dominion D'' that does not intersect with the current region R_s. Since L_s contains no tiny $\overline{\alpha}_s$-dominions, it holds that $|D''| > \lfloor \frac{b_{\overline{\alpha}_s}}{2} \rfloor$, and $D' \setminus D'' < \lfloor \frac{b_{\overline{\alpha}_s}}{2} \rfloor$. After this first call to **hsol**, the resulting state is locally stored, at Line 4, in order to determine, later on, whether the second phase achieves any progress.

The algorithm then proceeds to analyse the remaining part of the game still unprocessed. To do so, the next state computed by NextPr(s) is given as input to the recursive call **sol** at Line 5. Once the call completes, the state is updated with the two new functions returned by the call. At this point, also D'' has been processed, since **sol** has enough precision bound to find D''. What eventually remains unsolved is the set $D' \setminus D''$. This set, however can only contain tiny $\overline{\alpha}_s$-dominions that can also be found by **hsol**.

At this point, depending on whether or not the state is closed, either a promotion or a maximisation operation is performed (Lines 6 to 8) to ensure that the new state is maximal. If the middle phase with full precision made some progress in the search, a last call to **hsol** at Line 9 is performed, which again modifies the current state in-place turning it into a strongly-maximal one. Due to the small size of the residual unsolved dominions, this call suffices to solve the corresponding subgame ∂_s. If no progress occurred, instead, the current state s is equal to the one stored in and previously returned by the first call to **hsol**. Thus, s is strongly-maximal, D' was empty in the first place, and D has been completely processed in that first call to **hsol**. In either case, the state is fed to the function Und to handle possible undetermined positions, after which the current call terminates.

The procedure **hsol** simply executes the main body of the RPP algorithm by making mutually-recursive calls to the function **sol** (Line 5) with halved precision until no progress on the search for a dominion can be made. Similarly to

RPP, the auxiliary function NextPr identifies the next priority to consider, and the promotion and maximisation procedures, Promote and Maximise, generalise the corresponding procedures defined for RPP. Specifically, Promote applies a promotion, but this time the best escape priority can belong either to r or to u. Therefore, in base of the escaping priority, the current region R can be promoted to r, as for the classic Priority Promotion approach, or removed for it and promoted to u. Finally, Maximise turns a state s into a $\overline{\alpha_s}$-maximal one, by maximising both r and u. Similar to Parys' algorithm, the Half function halves the bound of the player $\overline{\alpha}_s$, leaving the bound of player α_s unchanged. Finally, the Und function starts from a strongly maximal state $s \in \mathrm{St_S}$ where all small dominions of player $\overline{\alpha}_s$ of size $\leq b_{\overline{\alpha}_s}$ *at the time of the current call* are processed. Thus, neither do any of the small dominions ($\leq b_{\overline{\alpha}_s}$) of player $\overline{\alpha}_s$ intersect with U_s, nor do any of the small dominions ($\leq b_{\overline{\alpha}_s}$) of player $\overline{\alpha}_s$ intersect with L_s. Depending on the parity of the calling priority c_s, we can return the respective set (U_s or L_s) and, where the parity is different, reset the positions of U_s in r to their original priority.

At this point, by defining the winning regions of the players as $\mathrm{Wn}_\supset^\alpha = \mathrm{r}^{-1}(\top_\alpha) \cup \mathrm{u}^{-1}(\top_{\overline{\alpha}})$ and $\mathrm{Wn}_\supset^{\overline{\alpha}} = \supset \setminus \mathrm{Wn}_\supset^\alpha$, where $(\mathrm{r}, \mathrm{u}) \triangleq \mathbf{sol}(s_I)$, we obtain a sound and complete solution algorithm, whose time-complexity is quasi-polynomial, as we shall show in the next section.

6. Correctness and Complexity

We now discuss how we can entangle the concepts of Priority Promotion, *i.e.* the transfer of information across the call structure that makes it more efficient in practice, with the concept of relative guarantees that provides favourable complexity bounds to Parys' algorithm. Before turning to the principle guarantees provided by the algorithm, we note that the two algorithms from the previous sections (Parys' and the selected variation of Priority Promotion) can be viewed as variations of our hybrid algorithm. This is particularly easy to see for the RPP algorithm from Section 3: when we set the bounds to infinity (or to 2^c, where c is the number of different priorities of the game) then the algorithm never runs out of bounds. In this case, the function u is never used, and the algorithm behaves exactly as Algorithm 1. In more detail, the whole game is recursively solved by the first call to **hsol** at Line 3 of Algorithm 5 that act like **sol** of Algorithm 1. Half has no effect, because the bounds always remain greater than 0, and Und does nothing since U_s is constantly empty and, at the end of **sol**, L_s is empty as well. Finally, Promote and Maximise only use the function r as a consequence of the fact that u is never filled by Und.

The connection to Parys' algorithm is slightly looser, but essentially it works one priority at a time solving the game purely with attractors. Hence it gives up promoting positions to higher priorities. This change does not impact on the partial correctness argument as it guarantees progress, but significantly slow down the solution procedure. In more detail, neither **sol** nor **hsol** check for open/closed regions as all regions are always treated as open. Half is the same, NextPr does not skip any priority and, as a consequence, u always applies the

case in which the parity of the current priority is different from the parity of the caller priority, assigning the uncertain positions to the previous (caller) priority. Finally, Maximise closes only u under attractor.

As the algorithm is a hybrid one, its correctness proof has both local and global aspects. The global guarantees are that the regions stored in $\mathsf{r}^{(>p_s)}$ and the bounded dominions stored in $\mathsf{u}^{(\geq p_s)}$ retain their properties (that are described in the last lemma) in all function calls. The local guarantees refer instead to the local area L_s of the current state s introduced in Definition 7. To conveniently reason about the guarantees we define the set P_s, which stores the local area L_s at the beginning of the call of **sol**; P_s is then available also in the **hsol** at the level where they are called.

This additional set P_s of initial positions is relevant as the guarantees of finding small dominions is formulated relative to it and not relative to the local area L_s at the end of **sol**. The correctness proof falls into lemmas that refer to the guarantees maintained by the auxiliary functions, and an inductive proof of the main lemma. We introduce in the following the lemmas for the auxiliary functions, while their proofs are reported in Appendix A. A final one, instead, outline the inductive proof for the main functions making use of the previous lemmas. After them, we report a theorem for the correctness of the algorithm that immediately follows from the invariants that have been proved before.

Lemma 3. *Given a strongly-maximal hybrid state* $s \in \mathrm{St}_\mathsf{S}$*, the functions* NextPr(s) *and* Half(s) *return a maximal hybrid state* $\widehat{s} \in \mathrm{St}_\mathsf{M}$ *such that* $\widehat{s} \prec s$.

This lemma states that both NextPr and Half take as input a strongly-maximal hybrid state and return a new maximal hybrid state which is smaller $w.r.t. \prec$ (the ordering among states from Definition 7).

Lemma 4. *Given a hybrid state* $s \in \mathrm{St}$*, the procedure* Maximise(s) *modifies in-place* s *into a maximal hybrid state* $\widehat{s} \in \mathrm{St}_\mathsf{M}$ *such that* $\widehat{s} \preccurlyeq s$.

The procedure Maximise takes as input a hybrid state that modifies in-place into a maximal hybrid state which is no greater than the original one $w.r.t. \prec$.

Lemma 5. *Given a promotable hybrid state* $s \in \mathrm{St}_\mathsf{P}$*, the procedure* Promote(s) *modifies in-place* s *into a maximal hybrid state* $\widehat{s} \in \mathrm{St}_\mathsf{M}$ *such that* $\widehat{s} \prec s$.

The procedure Promote takes as input a promotable hybrid state that modifies in-place into a maximal hybrid state which is smaller than the original one $w.r.t. \prec$.

Lemma 6. *Given a strongly-maximal hybrid state* $s \in \mathrm{St}_\mathsf{S}$*, the function* Und($s$) *returns a pair* $(\widehat{\mathsf{r}}, \widehat{\mathsf{u}}) \in \mathrm{Rg} \times \mathrm{Pf}$ *of a region function* $\widehat{\mathsf{r}}$ *and a promotion function* $\widehat{\mathsf{u}}$ *such that: 1)* $\mathsf{r}_s^{(>p_s)} \subseteq \widehat{\mathsf{r}}^{(>p_s)}$*; 2)* $\mathsf{u}_s^{(>p_s)} \subseteq \widehat{\mathsf{u}}^{(>p_s)}$*; 3)* $\widehat{s} = ((\widehat{\mathsf{r}}, c_s), (\widehat{\mathsf{u}}, \mathsf{b}_0, \mathsf{b}_1), q))$ *is a promotable hybrid state if* $\widehat{s}$ *is* closed *and, an hybrid state otherwise, for all priority* $q \in \mathsf{pr}$ *such that* $c_s < q \leq \min(\mathsf{rng}(\mathsf{r}_s \cup \mathsf{u}_s^{>c_s}))$*, and* $\mathsf{b}_0, \mathsf{b}_1 \in \mathbb{N}$.

The function Und takes as input a strongly-maximal hybrid state and returns a pair of region and promotion functions $(\widehat{\mathsf{r}}, \widehat{\mathsf{u}})$ that can be used to compose

a new hybrid state. In particular, both $\widehat{r}$ and $\widehat{u}$ include the positions mapped above the priority p_s in the input state components r and u, respectively. The new hybrid state can be composed by swapping, in the input state, the old pair (r, u) with $(\widehat{r}, \widehat{u})$ and the current priority p with the caller one c. Moreover, in case such new state is closed, then it is also promotable.

Lemma 7. *Given a maximal hybrid state $s \in \mathrm{St_M}$, the function $\mathbf{sol}(s)$ terminates and returns a pair $(\widehat{r}, \widehat{u}) \in \mathrm{Rg} \times \mathrm{Pf}$ of a region function $\widehat{r}$ and a promotion function $\widehat{u}$ such that: 1) $r_s^{(>p_s)} \subseteq \widehat{r}^{(>p_s)}$; 2) $u_s^{(>p_s)} \subseteq \widehat{u}^{(>p_s)}$; 3) $\widehat{s} = ((\widehat{r}, c_s), (\widehat{u}, b_0, b_1), q)$ is a promotable hybrid state if $\widehat{s}$ is* closed *and, an hybrid state otherwise, for all priorities $q \in \mathrm{pr}$ such that $c_s < q \leq \min(\mathrm{rng}(r_s \cup u_s^{>c_s}))$, and $b_0, b_1 \in \mathbb{N}$. Moreover, the procedure $\mathbf{hsol}(s)$ modifies in-place s into a strongly-maximal hybrid state $\widehat{s} \in \mathrm{St_S}$ such that $\widehat{s} \preccurlyeq s$.*

This lemma states that the function $\mathbf{sol}$ takes as input a maximal hybrid state and, when it terminates, it returns a pair of region and promotion functions $(\widehat{r}, \widehat{u})$. In particular, both $\widehat{r}$ and $\widehat{u}$ include the positions mapped above the priority p_s in the input state components r and u, respectively. The new hybrid state can be composed by swapping, in the input state, the old pair (r, u) with $(\widehat{r}, \widehat{u})$ and the current priority p with the caller one c. Moreover, in case such new state is closed, then it is also promotable. The procedure $\mathbf{hsol}$, instead, modifies in-place a maximal hybrid state into a strongly-maximal hybrid state which is no greater than the original one *w.r.t.* $\prec$.

We can now proceed with the main lemma.

Lemma 8. *main Let $s \in \mathrm{St_M}$ be a maximal hybrid state for a parity game ∂, where b_{0s} and b_{1s} are positive. Assume $\mathbf{sol}$ is called on s and let $s' \triangleq ((\widehat{r}, p'), (\widehat{u}, b_{0s}, b_{1s}), c')$ be the hybrid state, for $(\widehat{r}, \widehat{u}) \triangleq \mathbf{sol}(s)$, $p' \triangleq c_s$, and some $c' > c_s$. For $\alpha = \alpha_s$, the following holds:*

- *at the end of the call, H_s^α and $\mathrm{H}_s^{\overline{\alpha}}$ contain all small dominions of player α and size at most b_α, and small dominions of player $\overline{\alpha}$ and size at most $b_{\overline{\alpha}}$, respectively, in P_s;*

- *if $c_s \equiv_2 \alpha$ then $\mathrm{H}_{s'}^{\overline{\alpha}} = \mathrm{H}_{s'}^{\overline{\alpha}}$; hence, $\mathrm{H}_{s'}^{\overline{\alpha}}$ preserves all global guarantees of $\mathrm{H}_s^{\overline{\alpha}}$;*

- *if $c_s \not\equiv_2 \alpha$ then $\mathrm{H}_{s'}^{\overline{\alpha}} = \mathrm{H}_s^\alpha$; hence, $\mathrm{H}_{s'}^{\overline{\alpha}}$ preserves the dual of all global guarantees of H_s^α (since the caller player c_s has dual parity w.r.t. α).*

Moreover, if $\mathbf{hsol}$ is called on $s \in \mathrm{St_M}$, then it modifies s into a strongly-maximal hybrid state with the following properties:

- *at the end of the call, H_s^α and $\mathrm{H}_s^{\overline{\alpha}}$ contain all small dominions of player α and size at most b_α, and tiny dominions of player $\overline{\alpha}$ and size at most $\left\lfloor \frac{b_{\overline{\alpha}}}{2} \right\rfloor$, respectively, in P_s;*

- L_s *does not contain a small dominion of player $\overline{\alpha}$ of size at most $\left\lfloor \frac{b_{\overline{\alpha}}}{2} \right\rfloor$, and $\mathrm{H}_s^{\overline{\alpha}}$ is closed under attractor in P_s.*

Proof. We provide an inductive proof over the highest priority.

For the **induction basis** we consider the four cases in this order: (1) **sol** with highest priority $\perp$ (*i.e.* on the empty game); (2) **hsol** being called with highest priority 0; (3) **hsol** with highest priority $\perp$; and (4) **sol** being called with highest priority 0.

Case 1: When **sol** is called on the empty game the highest priority is $\perp$ and there is nothing to do (and due to Lines 1-2 **sol** does nothing).

Case 2: For games with maximal (and thus only) priority 0 in **hsol**, all states are in R_s, hence s is closed, and consequently, due to Lines 4 and 9, all states are promoted. The game is then empty (s is open), and by Lines 4 and 5, **sol** is called on the empty game. After that, Maximise is called (again on an empty set of positions). All global guarantees are retained due to Lemma 5, case 1, and Lemma 4. Moreover, as $\mathsf{dom}\big(\mathsf{r}_s^{\leq 0}\big) = \mathsf{u}_s^{-1}(0) = \emptyset$, the additional requirement for **hsol** holds.

Case 3: The case in which **hsol** is called on the empty game, is a sub-case of case 2 since s is open, and **sol** is called on the empty game.

Case 4: For games with maximal (and thus only) priority 0 in **sol**, all positions are promoted in the first call of **hsol** at Line 3 of Algorithm 5, similarly as shown for case 2. Then, s is empty and the successive call of **sol** at Line 5, of Maximise at Line 7, and of Und at Line 10 do nothing. The second call of **hsol** at Line 9 is skipped as s has not changed since Line 4.

Thus, the global guarantees are retained by the other base cases 1 and 2, and due to Lemmas 4 and 6.

We now implement the **induction step**, again first for **hsol**, and then for **sol**, using the results from **hsol**.

When executed, the call of **sol** at Line 5 of **hsol** provides, by induction hypothesis, an increase of $H_s^{\overline{\alpha}}$, which now contains every small dominion of player $\overline{\alpha}$ of size at most $\left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor$ in $L_s \setminus R_s$ from before the call, while retaining all global guarantees of H_s^{α} for which the bound is still b_{α_s}. When $\left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor = 0$, instead **sol** does nothing, trivially retaining the guarantees.

Thus, the global guarantees are preserved in every step either by induction hypothesis and Lemma 3, by the base case $\left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor = 0$, or by Lemmas 4 and 5 when functions Promote and Maximise are called.

Finally, to see that $H_s^{\overline{\alpha}}$ is closed, we observe that it is closed after Maximise is called, due to Lemma 4. Moreover, in a generic iteration of **hsol**, if s is closed, it follows a promotion and then $s \prec \widehat{s}$ triggers a new iteration, otherwise when s is open the last operation executed is a call of Maximise. Hence, in both cases we have that $H_s^{\overline{\alpha}}$ is closed when **hsol** ends.

So far we have shown that $H_s^{\overline{\alpha}}$ contains all small $\overline{\alpha}$-dominions of size at most $\left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor$ in $L_s \setminus R_s$. And since s has not changed during the call, it also contains the small dominions after the call, and on the time of return. To conclude this induction step we have to show that $H_s^{\overline{\alpha}}$ contains also the small dominions in L_s. Do to this, we observe that R_s is an α-region with highest priority p_s, and that, by Lemma 2, for any non-empty $\overline{\alpha}$-dominion D' in L_s, there must be a

non-empty sub-dominion $D'' \subseteq D'$, which is also a $\overline{\alpha}$-dominion in $L_s \setminus R_s$. As a consequence, since $L_s \setminus R_s$ does not contain a $\overline{\alpha}$-dominion D with size at most $\left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor$ when **hsol** returns, we can conclude that neither does L_s.

We can now turn to the induction step for **sol**. We first check that the global guarantees are retained up to the point where Und is called, then we make the argument that **sol** will, for a $\overline{\alpha}$-dominion D in P_s of size $|D| \leq b_{\overline{\alpha}_s}$, guarantee that $D \subseteq H_s^{\overline{\alpha}}$ holds. We will finally show that, under these conditions, Und will retain the global guarantees.

The check that the global guarantees are retained up to the point where Und is called is straightforward as we have seen that **hsol** (called at Lines 3 and 9) does by the induction step and Lemma 7, and both Maximise and Promote do retain the global guarantees by Lemmas 4 and 5, respectively.

We now argue that, when Und is called, $H_s^{\overline{\alpha}}$ contains all $\overline{\alpha}$-dominions in P_s of size at most $b_{\overline{\alpha}_s}$. Let D be such a dominion.

We first observe that by the global guarantees $H_s^{\alpha} \setminus R_s$ and its α-attractor $\mathsf{atr}^{\alpha}(H_s^{\alpha} \setminus R_s)$, cannot intersect with D at any point. Thus, D is a dominion in $P_s \setminus \mathsf{atr}^{\alpha}(H_s^{\alpha} \setminus R_s)$ at any point by Lemma 1.

After the first call of **hsol** (at Line 3), we have established by induction step that L_s does not contain a $\overline{\alpha}$-dominion of size at most $\left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor$. Now two cases may arise: L_s does not intersect with D or it does. The first case entails that D must be contained in $H_s^{\overline{\alpha}}$. In the second case, there is a set $D' = D \cap L_s$ that, by Lemma 1, must be a $\overline{\alpha}$-dominion in L_s, as well as in $P_s \setminus \mathsf{atr}^{\alpha}(H_s^{\alpha} \setminus R_s)$. By Lemma 2, we also know that D' has a sub-dominion $D'' \subseteq D'$ that does not intersect with R_s. Moreover, since L_s contains no smaller dominions than $\left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor$, it follows that $|D''| > \left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor$. Thus, by induction hypothesis, after the call of **sol** from Line 5, $H_s^{\overline{\alpha}}$ also contains D''. Note that in this case $s \prec \widehat{s}$ holds, and therefore **hsol** is called at Line 9. Again, by induction step, the global guarantees are retained, therefore also $D' \setminus D''$ is included in $H_s^{\overline{\alpha}}$.

Let us assume for contradiction it is not. Then, after the return from **hsol**, there is a non empty set $E = L_s \cap (D' \setminus D'')$ that still belongs to the current subgame. Clearly $D \supseteq D' \supseteq D' \setminus D''$, and since $|D| \leq b_{\overline{\alpha}_s}$ and $|D''| > \left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor$, it easily follows that $|E| \leq \left\lfloor \frac{b_{\overline{\alpha}_s}}{2} \right\rfloor$. Moreover, by the global guarantees, E cannot intersect with $H_s^{\alpha} \setminus R_s$ or with its attractor $\mathsf{atr}^{\alpha}(H_s^{\alpha} \setminus R_s)$. Hence, E is contained in $P_s \setminus \mathsf{atr}^{\alpha}(H_s^{\alpha} \setminus R_s)$, and $H_s^{\overline{\alpha}}$ (that now includes D'') is closed under $\overline{\alpha}$-attractor in P_s. Thus, by Lemma 2, E must be a $\overline{\alpha}$-dominion in L_s, as well as in $P_s \setminus \mathsf{atr}^{\alpha}(H_s^{\alpha} \setminus R_s)$. However, due to the inductive hypothesis, after the call of **hsol** at Line 9 we have that L_s cannot contain such a small dominion for player $\overline{\alpha}$, contradicting the assumption that E is not empty.

We have now established that the global guarantees hold when Und is called, as $H_s^{\overline{\alpha}}$ contains all $\overline{\alpha}$-dominions of size at most $b_{\overline{\alpha}_s}$ in P_s. This entails that $H_s^{\alpha} \cup L_s$ contains all α-dominions of size at most b_{α_s} in P_s, when Und is called.

Und then moves the respective set, U_s (or L_s), to $\mathsf{u}_s^{-1}(c_s)$ of the caller priority

c_s, such that $\mathrm{H}_{s'}^{\overline{\alpha}}$ is equal to $\mathrm{H}_s^{\overline{\alpha}}$ (or to $\mathrm{H}_s^{\alpha} \cup \mathrm{L}_s$) when $c_s \equiv_2 p_s$ (or $c_s \not\equiv_2 p_s$, respectively). Thus, due to Lemma 6, Und retains the global guarantees.

This completes the induction step for **sol** and the proof. □

Correctness is then the special case when **sol** is called with $c_s = \top_{\overline{\alpha}}$, where $\alpha \triangleq \mathsf{pr}(\eth) \bmod 2$, and full precision, *i.e.* $b_\alpha = |\eth|$, for all $\alpha \in \mathbb{B}$.

Theorem 2. *When* **sol** *is called with the initial state* s_I, *i.e. with* $c_s = \top_{\overline{\alpha_I}}$, *where* $\alpha_I \triangleq \mathsf{pr}(\eth) \bmod 2$, *and full precision* $b_\alpha = |\eth|$, *for all* $\alpha \in \mathbb{B}$, *then, after* **sol** *returns*, $\mathsf{r}^{-1}(\top_{\alpha_I}) \cup \mathsf{u}^{-1}(\top_{\overline{\alpha_I}})$ *is the winning region of player* α_I. □

The theorem immediately follows from the previous lemma since the initial state has full precision and, therefore, the guarantees are not bounded. The assignment of the winning regions is due to the fact that $\mathsf{r}^{-1}(\top_0)$ and $\mathsf{r}^{-1}(\top_1)$ contain dominions of the respective players and are closed under attractor by our global guarantees. It is also clear that $\mathsf{u}^{-1}(\top_{\overline{\alpha_I}})$ can only be filled in the final call of Und, such that $\mathsf{r}^{-1}(\top_{\alpha_I}) \cup \mathsf{u}^{-1}(\top_{\overline{\alpha_I}})$ contains all dominions (as the bound does not exclude any) of Player α_I, but does not intersect with any dominion (as the bound again does not include any) dominion of Player $\overline{\alpha_I}$. Note that the wining region of Player 0 is simply the complement of the winning region of Player 1.

It is interesting to note an algorithmic difference between the parts of the winning regions in the dominions $\mathsf{r}^{-1}(\top_0)$ and $\mathsf{r}^{-1}(\top_1)$, and the rest of the winning regions of both players. The two sets are computed constructively, and winning strategies are simply the contribution of attractor strategies. This is not the case for the remainder of the winning regions, as their calculation is not constructive.

As a hybrid between Priority Promotion and Parys' approach, the algorithm retains the quasi-polynomial bound of Parys and the practical efficiency of Priority Promotion from [26, 27, 28]. Our argument is exactly the same as Parys' (end of Section 4 of [59]): we use 2 parameters, h for the number of priorities, and $l = 2\lfloor \log_2(n) \rfloor + 1$ for the bounds, where n is the number of positions. We estimate the number of times **sol** is called, excluding the trivial calls that return immediately (at Line 2) because we run out of priorities or bounds, by $R(h, l)$. If $h = 0$, then we have run out of priorities ($p = \bot$), and $R(h, l) = 0$. If $l = 0$, then we have run out of bounds ($b_0 = 0$ or $b_1 = 0$, with the other value being 1). As argued in Section 4 of [59], we can estimate $R(h, l) \leq 2n^l \binom{h+l}{l} - 1$. Therefore, since the cost of all operations is linear in the size of the game, and since l is logarithmic in the number of positions, this provides a quasi-polynomial running time.

7. Experimental Evaluation

The practical effectiveness of the algorithms presented here, namely RPP and HPP, has been assessed by means of an extensive experimentation on both concrete and synthetic benchmarks. The algorithms have been incorporated in Oink [64], a tool written in C++ that collects implementations of several parity

game solvers proposed in the literature, including some of the known quasi-polynomial ones. We shall compare solution times against the quasi-polynomial solvers SSPM [51], QPT [50], and the improved version of Parys' algorithm PAR [58], as well as the original version of the best exponential solver classes, namely the recursive algorithm ZLK from [30] and the original priority promotion PP [27], whose superiority in practical contexts is widely acknowledged[1] (see *e.g.*, [68, 64]).

The results give a simple argument for why and when to use this algorithm[2]. The first question is why one should use a QP algorithm. The answer to that question is quite simple: it is not hard to produce pathological cases for exponential time solvers. For complex games, it is well known that even the most efficient solvers in practice, *i.e.* ZLK and PP [42, 24, 69, 44], take exponential time, while HPP has a quasi-polynomial worst case complexity. To show this behaviour, we have evaluated these two solvers, PP and ZLK, with HPP and the improved version of Parys' algorithm PAR, a quasi polynomial version of ZLK, on the worst case family developed for the approaches based on quasi-dominions [24]. The results are reported in Table 1. Clearly the complex infrastructure required by the HPP can pay off in terms of running time, while PAR does not outperform ZLK on these small examples.

Worst Case		PP		PAR		ZLK		HPP	
Index	Nodes	Time	Iterations	Time	Iterations	Time	Iterations	Time	Iterations
10	88	0.00	11267	0.00	12750	0.00	1295	0.00	69
15	170	0.18	524294	0.21	502439	0.00	10175	0.00	177
20	278	7.70	22020104	4.54	9561994	0.03	133631	0.00	339
25	410	-	-	51.13	82520213	0.17	796671	0.00	879
30	568	-	-	-	-	1.99	8863743	0.00	1217
35	750	-	-	-	-	11.25	47120383	0.01	2125
40	958	-	-	-	-	-	-	0.02	2952

Table 1: Solution times in seconds on the worst case family [24].

The reason for why most QP algorithms should not be used in practice becomes evident when the algorithms are run on Keiren's family of (for current solvers) simple benchmarks [70]. This set of benchmarks comprises a number of concrete verification problems, ranging from model-checking to equivalence-checking and decision problems for different temporal logics. They can be divided in the following four categories.

---

[1]Variations of Zielonka's algorithm [65] as well as of Priority Promotion approaches [26, 27, 28] (including Tangle Learning [66] and Justification [67] based approaches), who share the same basic data structure and promotion principles, are available. Their performance on benchmarks does not vary significantly. We therefore went with the original Parity Promotion approach to compare the principle performance.

[2]Experiments were carried out on a 64-bit 3.9 GHz INTEL quad-core machine, with i5-6600K processor and 16GB of RAM, running UBUNTU 18.04 with LINUX kernel version 4.15.0. Oink was compiled with gcc 7.4.

Model-checking benchmarks. The first group contains 313 games, with size up to in the order of magnitude of 10^7 positions. It includes a number of different verification problems. A first set contains encodings of a variety of communication protocols from [71, 72, 73, 74]: the alternating bit protocol, the positive acknowledgement with retransmission protocol, the bounded retransmission protocol, and the sliding window protocols. The protocols are parameterised with the number of messages to send and, when applicable, the window size. The set also contains verification problems for the cache coherence protocol of [75] and the wait-free handshake register of [76], as well as the classic elevator and towers-of-Hanoi benchmarks from [77]. The verification tasks under analysis cover fairness, liveness and safety properties. A second set, instead, contains encodings of two-player board games, such as Clobber, Domineering, Hex, Othello, and Snake, all parameterised by their board size. Here, the existence of a winning strategy for the game is the property considered. The encoding into parity games results in games with very few priorities: up to 4 in some cases.

Equivalence checking benchmarks. This group contains 216 games encoding equivalence tests between processes. The verification problems test various forms of process equivalences, such as strong, weak and branching bisimulation, as well as branching simulation. Most of the processes are the ones already considered in the model-checking benchmarks. The encoding into parity games results in games with at most two priorities, hence the only relevant measure of difficulty is the size, again reaching $O(10^7)$ positions for the bigger instances.

Decision problem benchmarks. The third group contains encodings of satisfiability and validity problems for formulae of various temporal logics: LTL, CTL, CTL*, PDL and the μ-CALCULUS, and comprises 192 games. The maximal size of a benchmark is around $3 \cdot 10^6$ positions. The parity games encoding have been obtained with the tool MLSolver [78]. The situation here is more interesting, since these concrete problems feature a higher number of priority, up to 20 in few cases. Hence, unlike the previous two groups, these benchmarks allow us to stress a bit more the scalability of the solution algorithms *w.r.t.* the increase in priorities.

PGSolver. This group contains 291 synthetic benchmarks, corresponding to known families of hard cases for specific solvers and randomly generated ones. The sizes and number of priorities vary significantly, depending on the specific class of games.

Table 2 reports the results of the experiments for all the solvers considered in the analysis, divided by class of benchmarks[3]. For each solver, the total completion time, the average time per benchmark and the percentage of timed-out executions are given. We set a timeout of 10 seconds for all the benchmarks, except for the equivalence-check class, for which 40 seconds is used instead. As expected, the exponential solvers perform better on all the classes, with PP

---

[3]The benchmarks were run by issuing the following OINK commands: oink –no-single –no-loops –no-wcwc *GameName SolverName*; where solvers are: ZLK, NPP, RPP, SSPM, QPT, ZLKQ.

Benchmarks		Exponential			Quasi Polynomial			
		ZLK	PP	RPP	SSPM	QPT	PAR	**HPP**
Model-Ch.	Tot. Time	27.16	20.12	53.66	>849.01	>1259.84	42.38	**64.21**
	Avg. Time	0.08	0.06	0.17	>2.71	>4.02	0.13	**0.2**
	TimeOut	0%	0%	0%	22%	36.4%	0%	**0%**
Equiv. Ch.	Tot. Time	202.95	137.92	242.32	>2681.33	>3139.85	208.3	**280.81**
	Avg. Time	0.94	0.63	1.12	>12.41	>14.53	0.93	**1.3**
	TimeOut	0%	0%	0%	28.2%	27.3%	0%	**0%**
Decision Prb.	Tot. Time	13.20	11.75	13.27	>360.8	>853.64	66.85	**14.27**
	Avg. Time	0.07	0.06	0.07	>1.87	>4.44	0.35	**0.07**
	TimeOut	0%	0%	0%	11%	26.5%	0%	**0%**
PGSolver	Tot. Time	1.54	2.21	2.89	>3615.22	>4069.12	8.62	**4.04**
	Avg. Time	0.005	0.007	0.009	>12.42	>13.98	0.03	**0.01**
	TimeOut	0%	0%	0%	78%	92.4%	0%	**0%**

Table 2: Solution times in seconds on Keiren's benchmarks (1012 games).

taking the lead most of the time. SSPM and QPT both perform quite poorly, between two and three orders of magnitude worse than the other solvers, and do not seem to scale beyond the simplest instances, as also evidenced by the high number of timeouts. Both PAR and HPP, instead, perform relatively well in all the benchmarks, being able to solve all the instances without incurring any timeouts and maintaining a short distance from the exponential solvers performance-wise. PAR has a slight edge over HPP on the model-checking and equivalence checking problems, both of which feature a very low number of priorities, though the time advantage on average is typically negligible. On the other hand, when the number of priorities increases, like in the decision problems, the situation reverses and HPP takes the lead over PAR and practically matches the performance of the exponential solvers. We can observe the same behaviour also for the non-bounded recursive version RPP, whose performance relies less on the number of priorities, so that its advantage increases with a rising number of priorities. This seems to suggest that HPP may scale better *w.r.t.* the number of priorities in the games. To further investigate this behaviour we decided to perform additional experiments.

We then tried to assess scalability *w.r.t.* the number of positions and priorities, so as to evaluate how sensitive the solvers are to variations of those two parameters. To this end, we set up two types of synthetic benchmarks. The first kind of benchmarks keeps the number of priorities fixed and only increases the size of the underlying graph, while the second one maintains a linear relation between positions and priorities. Here we drop both SSPM and QPT, since they could not solve any of these benchmarks.

Figure 5 reports the solution times of the quasi polynomial solvers on 10 clusters, each composed of 100 randomly generated games, of increasing size varying from 10^4 to 10^5. Each point corresponds to the total time to solve all
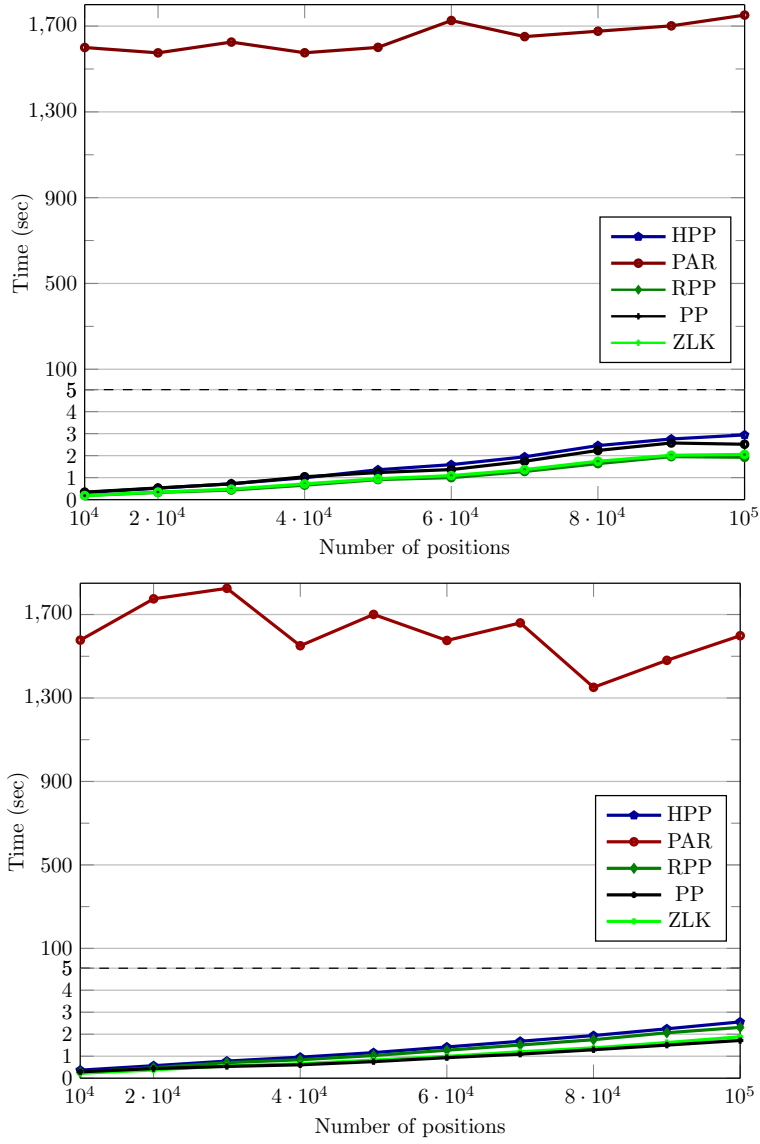
Figure 5: Results on random games with linear (on the left) and fixed (on the right) number priorities.

the games in the cluster[4]. On the top graph the number of priorities grows linearly with the positions, *i.e.*, equal to $\frac{n}{4}$, with n the number of positions. On the bottom graph, instead, all the games have 500 priorities. In both cases, the timeout was set to 25 seconds. In these experiments, HPP and PAR are tested together with the exponential solvers. The results seem to confirm what we already observed with the decision procedure benchmarks of the previous subsection. HPP definitely scales very well *w.r.t.* the number of priorities, as opposed to PAR, which is very sensitive to this parameter and starts hitting the timeout already on the smallest instances. HPP, indeed, behaves very much like the exponential solvers, none of which seems to be particularly sensitive to this parameter in practice, despite requiring time exponential in the number of priorities in the worst case.

What seems to emerge from the experimental analysis is that HPP behaves quite nicely in practice, competing with the leading exponential solution algorithms. The algorithmic overhead that guarantees its quasi-polynomial upper bound does not seem to impact the performance in any meaningful way. This is in contrast to what happens for all the other quasi-polynomial solvers, which do *not* scale with the number of positions and/or the number of priorities. This bodes well for the applicability of the hybrid approach in more challenging practical contexts, such as deciding temporal logic properties or solving reactive synthesis problems, where the number of priorities is typically higher.

Thus, HPP should be used when the game is hard; for this, it should have *some* structure (as opposed to be essentially random) as well as a high number of priorities, as the used advanced data structure only kicks in when the number of priorities is higher than $\log(n)$.

The excellent performance of the basic interplay between the two parts of the data structure invites exploring the limits of this approach. For example, one could refine the interplay between these parts in our algorithm, by extracting the guarantees on the bounds provided upon return instead of the guarantees required when a call is made.

As a final remark regarding the computation of the strategy, a very relevant question in reactive synthesis applications, we have observed that, in all the experiments we conducted, the final undetermined region is always empty. This means that, in practice, the winning strategies of the two players are always complete.

8. Discussion

We have combined two generalisations of the classic recursive algorithm: the quasi-polynomial recursion scheme of Parys, which relies on the local spread of imperfect guarantees, and a Priority Promotion scheme, which relies on identifying and realising the global potential of perfect guarantees. That these

---

[4]The instances were generated by issuing the following OINK command for top graph games: rngame n n/4 2 10; bottom graph games: rngame n 500 2 10

improvements can be synthesised bodes well, as it promises to perfectly join the advantages of both schemes, and the first experimental data collected suggests that the algorithm lives up to this promise.

Acknowledgments

The authors would like to thank the anonymous reviewers for the insightful comments and suggestions (in particular regarding the experimental analysis on the strategy coverage) that helped to considerably improve the earlier version of the manuscript.

Appendix A. Properties of the auxiliary functions for Section 6

Lemma 9. *Given a strongly-maximal hybrid state $s \in \mathrm{St_S}$, the functions* $\mathsf{NextPr}(s)$ *and* $\mathsf{Half}(s)$ *return a maximal hybrid state* $\widehat{s} \in \mathrm{St_M}$ *such that* $\widehat{s} \prec s$.

Proof. Let $s \triangleq ((\mathsf{r}, p), (\mathsf{u}, b_0, b_1), c)$ be a strongly-maximal hybrid state. For function NextPr, let $\widehat{s} = \mathsf{NextPr}(s)$ be the result obtained by calling the function NextPr on s. Due to Line 2 of Algorithm NextPr, $\widehat{s} = ((\mathsf{r}_s, q), (\mathsf{u}_s, b_{0s}, b_{1s}), p_s)$ where $q = \mathsf{max}(\mathsf{rng}(\mathsf{r}_s^{(<p_s)}))$ by Line 1. Consequently, trivially holds Condition 1a and 1d of Definition 7 and that r_s is a region function, u_s is a promotion function, p_s is a priority, and b_α for $\alpha \in \mathbb{B}$ are two integers. Moreover, clearly $q \in [\bot, p_s[$ and, therefore, also q is a priority and both $\{v \in \mathsf{dom}(\mathsf{u}_s) \mid \mathsf{u}_s(v) < p_s\}$ and $\mathsf{u}_s^{-1}(\top_\alpha)$, where $\alpha \triangleq \mathsf{pr}(\circlearrowright) \bmod 2$, are equal to $\emptyset$ as required by Condition 1b of the same definition. In addition, since $\mathrm{H}_s^\beta = \mathrm{H}_{\widehat{s}}^\beta \cup \mathsf{r}^{-1}(q)$ and $\mathrm{H}_{\widehat{s}}^{\overline{\beta}} = \mathrm{H}_s^{\overline{\beta}}$, with $\beta = q \bmod 2$, it also follows Condition 1c. Finally, to prove that $\widehat{s}$ is maximal it suffices to observe that $\mathrm{H}_{\widehat{s}}^\alpha = \mathrm{H}_s^\alpha \cup \mathrm{R}_s$ and $\mathrm{H}_{\widehat{s}}^{\overline{\alpha}} = \mathrm{H}_s^{\overline{\alpha}}$. Consequently, since s is strongly-maximal it holds that R_s is α-maximal *w.r.t.* L_s and, therefore, $\mathrm{H}_{\widehat{s}}^\alpha$ is α-maximal as well. To conclude, if $\mathrm{R}_s \neq \emptyset$ trivially we have that $\mathrm{L}_{\widehat{s}} \subset \mathrm{L}_s$, otherwise we have that $\mathrm{L}_{\widehat{s}} = \mathrm{L}_s$ and $p_{\widehat{s}} = q < p_s$. In both cases we can have that $\widehat{s} \prec s$.

We can now focus on Half function. Let us consider $\widehat{s}$ as the argument of function NextPr at Line 2 of Algorithm Half. Due to Line 1, we have that $\widehat{s} = ((\mathsf{r}, p), (\mathsf{u}, \widehat{b}_0, \widehat{b}_1), c)$, where $\widehat{b}_0 = \lfloor \frac{b_{0s}}{1+\alpha} \rfloor$ and $\widehat{b}_1 = \lfloor \frac{b_{1s}}{2-\alpha} \rfloor$. It is easy to observe, by the fact that s is a strongly-maximal hybrid state and that both $\widehat{b}_0$ and $\widehat{b}_1$ are integers, that $\widehat{s}$ is a strongly-maximal hybrid state as well. Hence, Half returns $\mathsf{NextPr}(\widehat{s})$ that in turn returns a new maximal hybrid state $s' \in \mathrm{St_M}$ for which $s' \prec \widehat{s}$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad\square$

Lemma 10. *Given a hybrid state* $s \in \mathrm{St}$, *the procedure* $\mathsf{Maximise}(s)$ *modifies in-place s into a maximal hybrid state* $\widehat{s} \in \mathrm{St_M}$ *such that* $\widehat{s} \preccurlyeq s$.

Proof. Let $s \triangleq ((\mathsf{r}, p), (\mathsf{u}, b_0, b_1), c)$ be an hybrid state and $\widehat{s} \triangleq ((\widehat{\mathsf{r}}, \widehat{p}), (\widehat{\mathsf{u}}, \widehat{b}_0, \widehat{b}_1), \widehat{c}) = \mathsf{Maximise}(s)$ the modified state obtained by computing the function $\mathsf{Maximise}$ on s. It is easy to see that $\widehat{p} = p$, $\widehat{c} = c$, and $\widehat{b}_\alpha = b_\alpha$ for $\alpha \in \mathbb{B}$. Moreover, due to Line 3 of Algorithm $\mathsf{Maximise}$, it holds that $\mathrm{X} \subseteq \mathrm{L}_s \cup \mathrm{U}_s \subseteq \{v \in \mathsf{dom}(\mathsf{pr}) \mid \mathsf{pr}(v) \leq p_s\}$ while, due to Line 4, we have that $q \geq p$. As a consequence, by Lines 6 and 8, $\widehat{\mathsf{r}}$ is a promotion function. Furthermore, from the observation that in case $q = \top_\beta$, with $\beta \in \mathbb{B}$, it necessarily holds that $\alpha \equiv_2 \beta$, we can also conclude that, by Lines 7 and 9, $\widehat{\mathsf{u}}$ is a promotion function as well. To prove that $\widehat{\mathsf{r}}$ is a region function, let us consider again the case in which $q = \top_\beta$ with $\beta \in \mathbb{B}$. Now, since $\alpha \equiv_2 \beta$, the set X corresponds to the β-attractor of $\mathrm{L}_s \cup \mathrm{U}_s$ by Line 3. Once X is merged to $\mathsf{r}^{-1}(\top_\beta)$, by Line 6, it is obvious that $\mathsf{r}^{-1}(\top_\beta)$ is still a β-dominion due to the fact that the β-attractor of a β-dominion is a β-dominion as well. We can now consider $q < \top_\beta$ for any $\beta \in \mathbb{B}$. Let us first assume that at Line 2 $\alpha = \alpha_s$. Two cases may arise: $q \equiv_2 \alpha$ or $q \not\equiv_2 \alpha$. In the

first case, due to Lines 5-6, we have that $\Delta_{\hat{r}}^{\alpha} = \Delta_{r}^{\alpha} \cup X$ which is clearly a quasi α-dominion since X corresponds to the α-attractor of L_s, while $\Delta_{\hat{r}}^{\overline{\alpha}} = \Delta_{r}^{\overline{\alpha}} \setminus X$ is a quasi $\overline{\alpha}$-dominion. Hence, $\hat{r}$ is a quasi dominion function. Finally, clearly X does not belong to $\mathsf{esc}^{\overline{\alpha}}(H^{\alpha})$, since Player α has a strategy to attract every position $v \in X$ to $r^{-1}(q)$, from which the opponent can only escape through positions having priority $q' \geq q$ and congruent to the parity of q. Summing up, the resulting $\hat{r}$ is a region function and due to the fact that the priority function pr is always a region function, we can conclude that also after Line 11 $\hat{r}$ is a region function. On the other hand, if $q \not\equiv_2 \alpha$, we can easily conclude that $\hat{r}$ is a region function due to Lines 5, 8, and 11. When, instead, at Line 2 $\alpha \neq \alpha_s$ and the priority selected at Line 4 has the same parity as α, the proof that $\hat{r}$ is a region function is equivalent to the one provided above for the case in which $q \not\equiv_2 \alpha$ and $\alpha = \alpha_s$. Similarly, when $q \not\equiv_2 \alpha$ the proof corresponds to the one provided for the case in which $q \equiv_2 \alpha$ and $\alpha = \alpha_s$. To prove Condition 1a of Definition 7 it suffices to observe that whenever X is merged to r (resp. u), it is removed from u (resp. u) due to Lines 6-7 and 8-9, while the operation at Line 11 does not change $\mathsf{dom}(r)$, since R_s is already stored in r. Moreover, due to Line 4, q refers to a priority in the domain of r or u, hence also Condition 1d holds. Condition 1b, instead, follows form the fact that $q \geq p$ and $q = \top_{\beta}$ with $\beta \in \mathbb{B}$ only if $\alpha \equiv_2 \beta$. Therefore, it holds that $\hat{u}^{-1}(\top_{\alpha}) = u^{-1}(\top_{\alpha})$, where $\alpha \triangleq \mathsf{pr}(\supset) \bmod 2$, and $\{v \in \mathsf{dom}(\hat{u}) \mid \hat{u}(v) < p\} = \{v \in \mathsf{dom}(u) \mid u(v) < p\}$ which are all empty. It remains to prove Condition 1c. Since s is a hybrid state, it holds that U_s and every $v \in \Delta_{u}^{\overline{\alpha}}$ cannot be directly attracted by R_s and more in general by L_s, due to Condition 1c. In addition, by Lines 5 and 9, the domain of u only increases when the priority selected at Line 4 has parity $\overline{\alpha}$. However, the positions of X merged in this case are not part of $\mathsf{esc}^{\overline{\alpha}}(H_{\hat{s}}^{\alpha})$ since X has been attracted to $\mathsf{dom}(u_s) \cap H_{s}^{\overline{\alpha}}$ and, therefore, Player $\overline{\alpha}$ has a strategy to reach this set from X. The maximality of $\hat{s}$ is a trivial consequence of Line 3, together with the observation that the operation at Line 11 cannot affect the maximality of the state. Finally, since $\mathsf{dom}(\hat{u}) \supseteq \mathsf{dom}(u)$ and $\mathsf{dom}(\hat{r}) \supseteq \mathsf{dom}(r)$, while $\hat{p} = p$, we can also conclude that $\hat{s} \preccurlyeq s$. $\qquad\square$

Lemma 11. *Given a promotable hybrid state $s \in \mathrm{St_P}$, the procedure $\mathsf{Promote}(s)$ modifies in-place s into a maximal hybrid state $\hat{s} \in \mathrm{St_M}$ such that $\hat{s} \prec s$.*

Proof. Let $s \triangleq ((r, p), (u, b_0, b_1), c)$ be a promotable hybrid state and $\hat{s} \triangleq ((\hat{r}, \hat{p}), (\hat{u}, \hat{b_0}, \hat{b_1}), \hat{c}) = \mathsf{Promote}(s)$ the modified state obtained by computing the function $\mathsf{Promote}$ on s. It is easy to see that $\hat{p} = p$, $\hat{c} = c$, and $\hat{b}_{\alpha} = b_{\alpha}$ for $\alpha \in \mathbb{B}$. At Line 1 of Algorithm $\mathsf{Promote}$ the two best escape priorities (p_r, p_u) from R_s to r and u are computed by the function bep. Due to the definition of bep and the fact that R_s is an α-dominion in $\supset_s$, it follows that both p_r and p_u are priorities higher than p. Moreover, by the same definition of bep and the fact that s is promotable, it also holds that $p_r \equiv_2 p$, while $p_u \not\equiv_2 p$ if $p_u < \top_{\alpha_s}$. Now, two cases may arise: $p_r \leq p_u$ or $p_r > p_u$. In the first case, by Line 2 and 3, we have that $\hat{u} = u$ and $\hat{r} = r[R_s \mapsto p_r]$. Hence, $\hat{u}$ is a promotion function and Condition 1b of Definition 7 is satisfied. Moreover, since $R_s \subseteq r^{(\leq p)}$, we have

37

that $\mathsf{dom}(\widehat{r}) = \mathsf{dom}(r)$ and, therefore, also the two requirements of Condition 1a are guaranteed. By the definition of bep, p_r belongs to the domain of r or is equal to the pseudo priority $\top$. In both cases Condition 1d still holds. Finally, since $p_r \equiv_2 p$, it follows that $H_{\widehat{s}}^{\beta} = H_s^{\beta}$ for any $\beta \in \mathbb{B}$. Hence, Condition 1c is satisfied. At this point, it remains to prove that $\widehat{r}$ is a region function. Now, since $R_s \subseteq r^{(\leq p)}$ and $p_r > p$, $\widehat{r}$ is clearly a promotion function. In addition, it is also a quasi dominion function as a consequence of the fact that $H_{\widehat{s}}^{\beta} = H_s^{\beta}$ for any $\beta \in \mathbb{B}$. Moreover, being r a region function, it holds that $\mathsf{pr}(v) \geq p$ and $\mathsf{pr}(v) \equiv_2 \alpha$ for each position v in $\mathsf{esc}^{\overline{\alpha}}(\Delta_r^{\alpha,p})$. However, due to Line 1 we have that $\mathsf{bep}_{\supset}^{\overline{\alpha}}(R_s, r_s) = p_r$, which implies that $\mathsf{pr}(v) \geq p_r$ and $\mathsf{pr}(v) \equiv_2 \alpha$ for each position v in $\mathsf{esc}^{\overline{\alpha}}(\Delta_r^{\alpha,p})$. Hence, $\widehat{r}$ is a region function. Finally, since s is promotable, it follows that $H_{\widehat{s}}^{\alpha} = H_s^{\alpha} \cup R_s$ is α-maximal, while $H_{\widehat{s}}^{\overline{\alpha}} = H_s^{\overline{\alpha}}$ is $\overline{\alpha}$-maximal and, therefore, $\widehat{s}$ is maximal. To conclude, being s promotable, it necessarily holds that $R_s \neq \emptyset$, which implies that $L_{\widehat{s}} \subset L_s$. Hence, $\widehat{s} \prec s$. Let us now consider the case in which $p_r > p_u$. Due to Line 2 and 4 of Algorithm Promote we have that $\widehat{r} = r \setminus R_s$ and $\widehat{u} = u[R_s \mapsto p_u]$. Conditions 1a and 1b follow from the fact that $\top_{\alpha_s} > p_u > p$ and $R_s \subseteq r^{(\leq p)}$, which also imply that $\widehat{u}$ is a promotion function. In addition, since $R_s \triangleq r^{-1}(p)$, it holds that $\Delta_{\widehat{r}}^{\alpha,p} = \Delta_r^{\alpha,p} \setminus R_s$, hence $\widehat{r}$ is a region function. By the definition of bep, p_r belongs to the domain of u, hence Condition 1d still holds. To prove Condition 1c, instead, let us observe that R_s is an α-dominion in $\supset_s$ and, consequently, $\mathsf{esc}^{\alpha}(R_s) \subseteq H_s^{\alpha}$. Moreover, we have that $p_u \not\equiv_2 p$, hence, as a consequence of the fact that $\widehat{u} = u[R_s \mapsto p_u]$, it holds that $R_s \in \Delta_{\widehat{u}}^{\overline{\alpha},p} \subseteq H_{\widehat{s}}^{\overline{\alpha}}$. In addition, it suffice to observe that $\{v \in \mathsf{dom}(\widehat{r}) \mid \widehat{r}(v) > p\} = \{v \in \mathsf{dom}(r) \mid r(v) > p\}$ and $\{v \in \mathsf{dom}(\widehat{u}) \mid \widehat{u}(v) \geq p\} = \{v \in \mathsf{dom}(u) \mid u(v) \geq p\} \cup R_s$, with $R_s \neq \emptyset$, to conclude that $\widehat{s} \prec s$. Finally, similarly to the previous case, it can be proved that both $H_{\widehat{s}}^{\alpha} = H_s^{\alpha} \cup R_s$ is α-maximal, while $H_{\widehat{s}}^{\overline{\alpha}} = H_s^{\overline{\alpha}}$ is $\overline{\alpha}$-maximal, as a direct consequence of the fact that s is promotable. □

Lemma 12. *Given a strongly-maximal hybrid state $s \in \mathrm{St_S}$, the function $\mathsf{Und}(s)$ returns a pair $(\widehat{r}, \widehat{u}) \in \mathrm{Rg} \times \mathrm{Pf}$ of a region function $\widehat{r}$ and a promotion function $\widehat{u}$ such that: 1) $r_s^{(>p_s)} \subseteq \widehat{r}^{(>p_s)}$; 2) $u_s^{(>p_s)} \subseteq \widehat{u}^{(>p_s)}$; 3) $\widehat{s} = ((\widehat{r}, c_s), (\widehat{u}, b_0, b_1), q))$ is a promotable hybrid state if $\widehat{s}$ is* closed *and, an hybrid state otherwise, for all priority $q \in \mathrm{pr}$ such that $c_s < q \leq \min(\mathsf{rng}(r_s \cup u_s^{>c_s}))$, and $b_0, b_1 \in \mathbb{N}$.*

Proof. Let $s \triangleq ((r,p), (u, b_0, b_1), c)$ be a strongly-maximal hybrid state and $(\widehat{r}, \widehat{u})$ the result obtained by computing the function Und on s. Moreover, let $\widehat{s} = ((\widehat{r}, c), (\widehat{u}, b_0, b_1), q)$ the composition of state s and the pair $(\widehat{r}, \widehat{u})$. Two cases may arise: $c \equiv_2 \alpha$ or $c \equiv_2 \overline{\alpha}_s$. In the first case, by Lines 1-2 of Algorithm Und, we have that $\widehat{r} = r$ and $\widehat{u} = u[U_s \mapsto c]$, where $U \triangleq u^{-1}(p)$. By the fact that the priorities of the states are strictly decreasing we have that $c > p$. Moreover, it holds that $\mathsf{dom}(\widehat{u}) = \mathsf{dom}(u)$ and, therefore, $\widehat{r}$ is a region function, $\widehat{u}$ is a promotion function, both Points 1) and 2) of the Lemma and Conditions 1a, 1b, and 1d of Definition 7 are satisfied for the composed state $\widehat{s}$. To prove Condition 1c it suffices to observe that due to Line 3 and the fact that $c \equiv_2 \alpha_s$, it holds that $H_{\widehat{s}}^{\beta} = H_s^{\beta}$ for any $\beta \in \mathbb{B}$. In particular, we have that $\Delta_{\widehat{r}}^{\alpha_s,q} = \Delta_r^{\alpha_s,q}$ for

all priorities $q \geq p$, $\Delta_{\widehat{\mathsf{u}}}^{\overline{\alpha}_s, q} = \Delta_{\mathsf{u}}^{\overline{\alpha}_s, q}$ for all priorities $q > c$ and $\Delta_{\widehat{\mathsf{u}}}^{\overline{\alpha}_s, c} = \Delta_{\mathsf{u}}^{\overline{\alpha}_s, c} \cup \mathrm{U}_s$. Finally, if $\widehat{s}$ is *closed*, it necessarily holds that $\widehat{s}$ is $\overline{\alpha}_s$-maximal and that $\mathrm{R}_{\widehat{s}}$ is α-maximal as a direct consequence of the fact that s is strongly-maximal and $\mathsf{r}^{(>p) \wedge (<c_s)} = \emptyset$. Hence, $\widehat{s}$ is also promotable. We can now focus on the case in which $c \equiv_2 \overline{\alpha}_s$. By Lines 1 and 3-4 of Algorithm Und we have that $\widehat{\mathsf{u}} = \mathsf{u}^{(\geq c)}[\mathrm{L}_s \mapsto c]$ and $\widehat{\mathsf{r}} = \mathsf{r}^{(\geq c)}[v \in \mathrm{U}_s \mapsto \mathsf{pr}(v)]$. Now, it is easy to see that also in this case both Points 1) and 2) of the Lemma and Conditions 1a, 1b, and 1d are satisfied for the composed state $\widehat{s}$, since $\mathrm{L} \triangleq \mathsf{dom}(\mathsf{r}^{\leq p})$ and $\mathrm{U} \triangleq \mathsf{u}^{-1}(p)$ and, it also holds that $\widehat{\mathsf{u}}$ is a promotion function and $\widehat{\mathsf{r}}$ is a region function due to the fact that the priority function pr is always a region function. Condition 1c, instead, follows from the fact that s is a strongly-maximal hybrid state and, therefore, it holds that $\mathsf{esc}^{\alpha}(\mathrm{L}_s) \subseteq \mathrm{H}_s^{\overline{\alpha}_s}$. By the same Condition 1c, we have that $\mathsf{esc}^{\alpha_s}(\Delta_{\mathsf{u}}^{\alpha_s, p}) \subseteq \mathsf{pre}^{\overline{\alpha}_s}(\Delta_{\mathsf{r}}^{\alpha_s, p} \setminus \mathrm{L}_s)$ and, consequently, for each priority $q > p$ the set $\mathsf{u}^{-1}(q)$ cannot reach U_s that, due to Line 5, is reset to the values of the priority function pr. The latter observation also implies that, if $\widehat{s}$ is *closed*, $\widehat{s}$ is also $\overline{\alpha}_s$-maximal. Indeed, when $\widehat{s}$ is *closed*, it holds that player $\overline{\alpha}_s$ cannot attracts $\mathrm{R}_{\widehat{s}}$ and since $\mathrm{L}_{\widehat{s}} \supseteq \mathrm{R}_{\widehat{s}}$ we can conclude that $\mathrm{H}_{\widehat{s}}^{\alpha_s} \setminus \mathrm{L}_{\widehat{s}}$ is $\overline{\alpha}_s$-maximal *w.r.t.* $\mathrm{L}_{\widehat{s}}$. Finally, similarly to the previous case, $\mathrm{R}_{\widehat{s}}$ turns out to be α-maximal as a direct consequence of the fact that s is strongly-maximal and that $\mathsf{r}^{(>p) \wedge (<c_s)} = \emptyset$. Hence, $\widehat{s}$ is also promotable. $\qquad \square$

Lemma 13. *Given a maximal hybrid state $s \in \mathrm{St}_{\mathsf{M}}$, the function $\mathbf{sol}(s)$ terminates and returns a pair $(\widehat{\mathsf{r}}, \widehat{\mathsf{u}}) \in \mathrm{Rg} \times \mathrm{Pf}$ of a region function $\widehat{\mathsf{r}}$ and a promotion function $\widehat{\mathsf{u}}$ such that: 1) $\mathsf{r}_s^{(>p_s)} \subseteq \widehat{\mathsf{r}}^{(>p_s)}$; 2) $\mathsf{u}_s^{(>p_s)} \subseteq \widehat{\mathsf{u}}^{(>p_s)}$; 3) $\widehat{s} = ((\widehat{\mathsf{r}}, c_s), (\widehat{\mathsf{u}}, \mathsf{b}_0, \mathsf{b}_1), q)$ is a promotable hybrid state if $\widehat{s}$ is closed and, an hybrid state otherwise, for all priorities $q \in \mathsf{pr}$ such that $c_s < q \leq \min(\mathsf{rng}(\mathsf{r}_s \cup \mathsf{u}_s^{>c_s}))$, and $\mathsf{b}_0, \mathsf{b}_1 \in \mathbb{N}$. Moreover, the procedure $\mathbf{hsol}(s)$ modifies in-place s into a strongly-maximal hybrid state $\widehat{s} \in \mathrm{St}_{\mathsf{S}}$ such that $\widehat{s} \preccurlyeq s$.*

Proof. Let $\widehat{s} \triangleq \mathbf{sol}(s)$ be the result obtained by computing $\mathbf{sol}$ on a maximal hybrid state $s \triangleq ((\mathsf{r}, p), (\mathsf{u}, \mathsf{b}_0, \mathsf{b}_1), c)$. The proof proceeds by induction, where in the base case we have that $p = \bot$ or $b_{\alpha_s} = 0$. In this case, by Line 1 and 2 of Algorithm 5 we have that $\widehat{s} = ((\mathsf{r}, c_s), (\mathsf{u}, \mathsf{b}_0, \mathsf{b}_1), q)$, which trivially satisfies all the requirements. Therefore, let us now consider an arbitrary recursive call of $\mathbf{sol}$ where $p > \bot$ and $b_{\alpha_s} > 0$. At Line 3 of $\mathbf{sol}$, the procedure $\mathbf{hsol}$ is called on the state s. Consequently, at Line 1 of Algorithm 9, the set R_s is maximised by computing its α_s-attractor $\mathsf{atr}_{\supseteq_s}^{\alpha_s}$. It is not hard to show that the new state s is still a strongly-maximal hybrid state greater or equal to the old one, *w.r.t.* the ordering of Definition 7. Then, two cases may arise at Line 4 of Half: s is *open* or *closed*. In the first case, at Line 5 a recursive call of $\mathbf{sol}$ is performed on the state generated by Half. Now, due to the fact that s is a strongly maximal hybrid state and by Lemma 3, it follows that the resulting state of $\mathsf{Half}(s)$ is a maximal hybrid state greater or equal to the input s, *w.r.t.* the ordering of Definition 7. Hence, by external induction, we have that $s = ((\widehat{\mathsf{r}}, p), (\widehat{\mathsf{u}}, b_0, b_1), c)$ is a hybrid state and, by Point 1) and 2) it is also greater or equal to the previous s, *w.r.t.* the ordering of Definition 7. At this point (Line 6 of Half), s can be *open* or *closed*.

Let us first consider again the case in which it is *open*. By Lemma 4 and Line 7, the function Maximise modifies s into a maximal hybrid state that is greater or equal to the previous s, *w.r.t.* the ordering of Definition 7. On the other hand, when s is *closed*, by external induction it holds that s is promotable and, by Lemma 5 the function Promote modifies s into a maximal hybrid state greater than the previous one *w.r.t.* the ordering of Definition 7. This last case also apply when s is *closed* at Line 4. Now, since the modified state by Promote(s) is always greater than s, *w.r.t.* the ordering of Definition 7, we can conclude that function **hsol** only ends when the function Maximise modify the input hybrid state such that there is no progress *w.r.t.* the ordering of Definition 7, otherwise **hsol** starts a new iteration due to the Loop at Line 1. By the fact that the state space is finite and that it starts a new iteration only if there is a progress in the ordering it is easy to prove that **hsol** always terminates. Observe that, when $s' \not\prec s$ where s' is the state modified by Maximise(s), it holds that s' is also strongly-maximal and, therefore, **hsol** returns a strongly maximal hybrid state. We can now go back to Line 3 of Algorithm 5 where the modified state s is strongly maximal as observed above. Now, by Lemma 3 the input state of **sol** at Line 5 is maximal. Hence, by external induction, the modified state composed of the results of Line 5 is $s = ((\widehat{r}, p), (\widehat{u}, b_0, b_1), c)$ that is a hybrid state and, by Point 1) and 2) it is also greater or equal to the previous state stored as $\widehat{s}$, *w.r.t.* the ordering of Definition 7. Similarly to the proof for **hsol**, the state s can be *open* or *closed*. In the first case, two more cases may arise, indeed, if s is *open*, then by Lemma 4 and Line 7, the function Maximise modifies s into a maximal hybrid state that is greater or equal to the previous s, *w.r.t.* the ordering of Definition 7. On the other hand, if s is *closed*, then by external induction it holds that s is promotable and, by Lemma 5 the function Promote modifies s into a maximal hybrid state greater than the previous one *w.r.t.* the ordering of Definition 7. Thus, at Line 9, in case $s \prec \widehat{s}$, where $\widehat{s}$ is the state stored before the recursive calls at Line 4, **hsol** is called on a maximal state and, as proved above, the modified resulting state is strongly maximal. Otherwise, as observed in the proof for **hsol**, s is already strongly-maximal. As a consequence, by Lemma 6 and Line 10, we can conclude that **sol** returns a pair of region function and priority function $(\widehat{r}, \widehat{u})$ such that $\widehat{s} = ((\widehat{r}, c_s), (\widehat{u}, b_0, b_1), q)$ is a promotable hybrid state if $\widehat{s}$ is *closed* and, an hybrid state otherwise, for all priorities $q \in \mathsf{pr}$ such that $c_s < q \leq \min(\mathsf{rng}(r_s \cup u_s^{>c_s}))$, and $b_0, b_1 \in \mathbb{N}$. $\qquad\square$

References

[1] E. Emerson, C. Jutla, A. Sistla, On Model Checking for the muCalculus and its Fragments, Theoretical Computer Science 258 (1-2) (2001) 491–522.

[2] E. Emerson, C. Lei, Efficient Model Checking in Fragments of the Propositional muCalculus, in: Logic in Computer Science'86, IEEE Computer Society, 1986, pp. 267–278.

[3] M. Vardi, Reasoning about The Past with Two-Way Automata, in: International Colloquium on Automata, Languages, and Programming'98, LNCS 1443, Springer, 1998, pp. 628–641.

[4] N. Piterman, From Nondeterministic Büchi and Streett Automata to Deterministic Parity Automata, in: Logic in Computer Science'06, IEEE Computer Society, 2006, pp. 255–264.

[5] S. Schewe, Tighter Bounds for the Determinisation of Büchi Automata, in: Foundations of Software Science and Computational Structures'09, LNCS 5504, Springer, 2009, pp. 167–181.

[6] E. Grädel, W. Thomas, T. Wilke, Automata, Logics, and Infinite Games: A Guide to Current Research, LNCS 2500, Springer, 2002.

[7] O. Kupferman, M. Vardi, Weak Alternating Automata and Tree Automata Emptiness, in: Symposium on Theory of Computing'98, Association for Computing Machinery, 1998, pp. 224–233.

[8] K. Chatterjee, T. Henzinger, N. Piterman, Strategy Logic, Information and Computation 208 (6) (2010) 677–693.

[9] F. Mogavero, A. Murano, M. Vardi, Reasoning About Strategies, in: Foundations of Software Technology and Theoretical Computer Science'10, LIPIcs 8, Leibniz-Zentrum fuer Informatik, 2010, pp. 133–144.

[10] F. Mogavero, A. Murano, G. Perelli, M. Vardi, What Makes ATL* Decidable? A Decidable Fragment of Strategy Logic, in: Concurrency Theory'12, LNCS 7454, Springer, 2012, pp. 193–208.

[11] M. Benerecetti, F. Mogavero, A. Murano, Substructure Temporal Logic, in: Logic in Computer Science'13, IEEE Computer Society, 2013, pp. 368–377.

[12] D. Berwanger, E. Grädel, Fixed-Point Logics and Solitaire Games, Theoretical Computer Science 37 (6) (2004) 675–694.

[13] D. Kozen, Results on the Propositional muCalculus, Theoretical Computer Science 27 (3) (1983) 333–354.

[14] L. de Alfaro, T. Henzinger, R. Majumdar, From Verification to Control: Dynamic Programs for Omega-Regular Objectives, in: Logic in Computer Science'01, IEEE Computer Society, 2001, pp. 279–290.

[15] T. Wilke, Alternating Tree Automata, Parity Games, and Modal muCalculus, Bulletin of the Belgian Mathematical Society 8 (2) (2001) 359–391.

[16] S. Schewe, B. Finkbeiner, Satisfiability and Finite Model Property for the Alternating-Time muCalculus, in: Computer Science Logic'06, LNCS 4207, Springer, 2006, pp. 591–605.

[17] R. Alur, T. Henzinger, O. Kupferman, Alternating-Time Temporal Logic, Journal of the ACM 49 (5) (2002) 672–713.

[18] S. Schewe, ATL* Satisfiability is 2ExpTime-Complete, in: International Colloquium on Automata, Languages, and Programming'08, LNCS 5126, Springer, 2008, pp. 373–385.

[19] A. Mostowski, Games with Forbidden Positions, Tech. rep., University of Gdańsk, Gdańsk, Poland (1991).

[20] E. Emerson, C. Jutla, Tree Automata, muCalculus, and Determinacy, in: Foundation of Computer Science'91, IEEE Computer Society, 1991, pp. 368–377.

[21] A. Martin, Borel Determinacy, Annals of Mathematics 102 (2) (1975) 363–371.

[22] M. Jurdziński, Deciding the Winner in Parity Games is in UP ∩ co-UP, Information Processing Letters 68 (3) (1998) 119–124.

[23] C. Calude, S. Jain, B. Khoussainov, W. Li, F. Stephan, Deciding Parity Games in Quasipolynomial Time, in: Symposium on Theory of Computing'17, Association for Computing Machinery, 2017, pp. 252–263.

[24] M. Benerecetti, D. Dell'Erba, F. Mogavero, Solving Parity Games via Priority Promotion, in: Computer Aided Verification'16, LNCS 9780 (Part II), Springer, 2016, pp. 270–290.

[25] M. Benerecetti, D. Dell'Erba, F. Mogavero, A Delayed Promotion Policy for Parity Games, in: Games, Automata, Logics, and Formal Verification'16, EPTCS 226, 2016, pp. 30–45.

[26] M. Benerecetti, D. Dell'Erba, F. Mogavero, Improving Priority Promotion for Parity Games, in: Haifa Verification Conference'16, LNCS 10028, Springer, 2016, pp. 1–17.

[27] M. Benerecetti, D. Dell'Erba, F. Mogavero, Solving Parity Games via Priority Promotion, Formal Methods in System Design 52 (2) (2018) 193–226.

[28] M. Benerecetti, D. Dell'Erba, F. Mogavero, A Delayed Promotion Policy for Parity Games, Information and Computation 262 (2) (2018) 221–240.

[29] R. McNaughton, Infinite Games Played on Finite Graphs, Annals of Pure and Applied Logic 65 (1993) 149–184.

[30] W. Zielonka, Infinite Games on Finitely Coloured Graphs with Applications to Automata on Infinite Trees, Theoretical Computer Science 200 (1-2) (1998) 135–183.

[31] W. Ludwig, A Subexponential Randomized Algorithm for the Simple Stochastic Game Problem, Information and Computation 117 (1) (1995) 151–155.

[32] A. Puri, Theory of Hybrid Systems and Discrete Event Systems, Ph.D. thesis, University of California, Berkeley, CA, USA (1995).

[33] J. Vöge, M. Jurdziński, A Discrete Strategy Improvement Algorithm for Solving Parity Games, in: Computer Aided Verification'00, LNCS 1855, Springer, 2000, pp. 202–215.

[34] H. Björklund, S. Vorobyov, A Combinatorial Strongly Subexponential Strategy Improvement Algorithm for Mean-Payoff Games, Discrete Applied Mathematics 155 (2) (2007) 210–229.

[35] S. Schewe, An Optimal Strategy Improvement Algorithm for Solving Parity and Payoff Games, in: Computer Science Logic'08, LNCS 5213, Springer, 2008, pp. 369–384.

[36] J. Fearnley, Non-Oblivious Strategy Improvement, in: Logic for Programming Artificial Intelligence and Reasoning'10, LNCS 6355, Springer, 2010, pp. 212–230.

[37] S. Schewe, A. Trivedi, T. Varghese, Symmetric Strategy Improvement, in: International Colloquium on Automata, Languages, and Programming'15, LNCS 9135, Springer, 2015, pp. 388–400.

[38] O. Friedmann, An Exponential Lower Bound for the Latest Deterministic Strategy Iteration Algorithms, Logical Methods in Computer Science 7 (3) (2011) 1–42.

[39] O. Friedmann, A Subexponential Lower Bound for Zadeh's Pivoting Rule for Solving Linear Programs and Games, in: Integer Programming and Combinatorial Optimization'11, LNCS 6655, Springer, 2011, pp. 192–206.

[40] O. Friedmann, A Superpolynomial Lower Bound for Strategy Iteration Based on Snare Memorization, Discrete Applied Mathematics 161 (10-11) (2013) 1317–1337.

[41] T. van Dijk, G. Loho, M. Maat, The Worst-Case Complexity of Symmetric Strategy Improvement, in: Computer Science Logic'24, LIPIcs 288, Leibniz-Zentrum fuer Informatik, 2024, pp. 24:1–19.

[42] O. Friedmann, Recursive Algorithm for Parity Games Requires Exponential Time, RAIRO - Theoretical Informatics and Applications 45 (4) (2011) 449–457.

[43] M. Benerecetti, D. Dell'Erba, F. Mogavero, Robust Exponential Worst Cases for Divide-et-Impera Algorithms for Parity Games, in: Games, Automata, Logics, and Formal Verification'17, EPTCS 256, 2017, pp. 121–135.

[44] M. Benerecetti, D. Dell'Erba, F. Mogavero, Robust Worst Cases for Parity Games Algorithms, Information and Computation 272 (2020) 104501:1–31.

[45] U. Zwick, M. Paterson, The Complexity of Mean Payoff Games on Graphs, Theoretical Computer Science 158 (1-2) (1996) 343–359.

[46] A. Browne, E. Clarke, S. Jha, D. Long, W. Marrero, An Improved Algorithm for the Evaluation of Fixpoint Expressions, Theoretical Computer Science 178 (1-2) (1997) 237–255.

[47] M. Jurdziński, Small Progress Measures for Solving Parity Games, in: Symposium on Theoretical Aspects of Computer Science'00, LNCS 1770, Springer, 2000, pp. 290–301.

[48] S. Schewe, Solving Parity Games in Big Steps, Journal of Computer and System Sciences 84 (2017) 243–262.

[49] M. Jurdziński, M. Paterson, U. Zwick, A Deterministic Subexponential Algorithm for Solving Parity Games, SIAM Journal on Computing 38 (4) (2008) 1519–1532.

[50] J. Fearnley, S. Jain, S. Schewe, F. Stephan, 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'17, Association for Computing Machinery, 2017, pp. 112–121.

[51] M. Jurdziński, R. Lazic, Succinct Progress Measures for Solving Parity Games, in: Logic in Computer Science'17, Association for Computing Machinery, 2017, pp. 1–9.

[52] D. Dell'Erba, S. Schewe, Smaller Progress Measures and Separating Automata for Parity Games, Frontiers in Computer Science 4 (2022) 8:1–18.

[53] K. Lehtinen, A Modal mu Perspective on Solving Parity Games in Quasi-Polynomial Time, in: Logic in Computer Science'18, Association for Computing Machinery & IEEE Computer Society, 2018, pp. 639–648.

[54] P. Parys, Parity Games: Another View on Lehtinen's Algorithm., in: Computer Science Logic'20, LIPIcs 152, Leibniz-Zentrum fuer Informatik, 2020, pp. 32:1–15.

[55] L. Daviaud, M. Jurdzinski, K. Thejaswini, The Strahler Number of a Parity Game, in: International Colloquium on Automata, Languages, and Programming'20, LIPIcs 168, Leibniz-Zentrum fuer Informatik, 2020, p. 123:1–19.

[56] P. Parys, Parity Games: Zielonka's Algorithm in Quasi-Polynomial Time, in: Mathematical Foundations of Computer Science'19, LIPIcs 138, Leibniz-Zentrum fuer Informatik, 2019, pp. 10:1–13.

[57] W. Czerwinski, L. Daviaud, N. Fijalkow, M. Jurdzinski, R. Lazic, P. Parys, Universal Trees Grow Inside Separating Automata: Quasi-Polynomial Lower Bounds for Parity Games, in: Symposium on Discrete Algorithms'19, SIAM, 2019, p. 2333–2349.

[58] K. Lehtinen, S. Schewe, D. Wojtczak, Improving the Complexity of Parys' Recursive Algorithm, Tech. rep., arXiv (2019).

[59] K. Lehtinen, P. Parys, S. Schewe, D. Wojtczak, A Recursive Approach to Solving Parity Games in Quasipolynomial Time, Logical Methods in Computer Science 18 (1) (2022).

[60] M. Benerecetti, D. Dell'Erba, M. Faella, F. Mogavero, From Quasi-Dominions to Progress Measures, in: Automata Theory and Applications: Games, Learning and Structures, World Scientific, 2023, pp. 159–199.

[61] M. Benerecetti, D. Dell'Erba, F. Mogavero, Solving Mean-Payoff Games via Quasi Dominions, in: Tools and Algorithms for the Construction and Analysis of Systems'20, LNCS 12079, Springer, 2020, pp. 289–306.

[62] M. Benerecetti, D. Dell'Erba, F. Mogavero, Solving Mean-Payoff Games via Quasi Dominions, Information and Computation 297 (2024) 105151:1–25.

[63] M. Jurdziński, M. Paterson, U. Zwick, A Deterministic Subexponential Algorithm for Solving Parity Games, in: Symposium on Discrete Algorithms'06, SIAM, 2006, pp. 117–123.

[64] T. van Dijk, Oink: an Implementation and Evaluation of Modern Parity Game Solvers, in: Tools and Algorithms for the Construction and Analysis of Systems'18, LNCS 10805, Springer, 2018, pp. 291–308.

[65] Y. Liu, Z. D., C. Tian, An Improved Recursive Algorithm for Parity Games, in: Theoretical Aspects of Software Engineering'14, IEEE Computer Society, 2014, pp. 154–161.

[66] T. van Dijk, Attracting Tangles to Solve Parity Games, in: Computer Aided Verification'18, LNCS 10982, Springer, 2018, pp. 198–215.

[67] R. Lapauw, M. Bruynooghe, M. Denecker, Improving Parity Game Solvers with Justifications, in: Verification, Model Checking, and Abstract Interpretation'20, LNCS 11990, Springer, 2020, pp. 449–470.

[68] L. Sanchez, W. Wesselink, T. A. C. Willemse, A Comparison of BDD-Based Parity Game Solvers, in: Games, Automata, Logics, and Formal Verification'18, EPTCS 277, 2018, pp. 103–117.

[69] T. van Dijk, A Parity Game Tale of Two Counters, in: Games, Automata, Logics, and Formal Verification'19, EPTCS 305, 2019, pp. 107–122.

[70] J. Keiren, Benchmarks for Parity Games, in: Fundamentals of Software Engineering'15, LNCS 9392, Springer, 2015, pp. 127–142.

[71] C. Koymans, J. Mulder, A Modular Approach to Protocol Verification Using Process Algebra, in: Applications of Process Algebra, Cambridge University Press, 1990, pp. 261–306.

[72] V. Cerf, R. Kahn, A Protocol for Packet Network Intercommunication, Transactions on Communications 22 (5) (1974) 637–648.

[73] J. Groote, J. van de Pol, A Bounded Retransmission Protocol for Large Data Packets, in: American Mathematical Society Translations96, LNCS 1101, Springer, 1996, pp. 536–550.

[74] K. Bartlett, R. Scantlebury, P. Wilkinson, A Note on Reliable Full-Duplex Transmission over Half-Duplex Links, Communication of the ACM 12 (5) (1969) 260–261.

[75] R. Veldema, R. Hofman, R. Bhoedjang, C. Jacobs, H. Bal, Source-Level Global Optimizations for Fine-Grain Distributed Shared Memory Systems, in: Principles and Practice of Parallel Programming'01, Association for Computing Machinery, 2001, pp. 83–92.

[76] W. Hesselink, Invariants for the Construction of a Handshake Register, Information Processing Letters 68 (4) (1998) 173–177.

[77] O. Friedmann, M. Lange, Solving Parity Games in Practice, in: Automated Technology for Verification and Analysis'09, LNCS 5799, Springer, 2009, pp. 182–196.

[78] O. Friedmann, M. Lange, A Solver for Modal Fixpoint Logics, Electronic Notes in Theoretical Computer Science 262 (2010) 99–111.