

Faster Game Solving by Fixpoint Acceleration

Daniel Hausmann*

University of Gothenburg, Sweden

hausmann@chalmers.se


We propose a method for solving parity games with acyclic (DAG) sub-structures by computing nested fixpoints of a DAG attractor function that lives over the non-DAG parts of the game, thereby restricting the domain of the involved fixpoint operators. Intuitively, this corresponds to accelerating fixpoint computation by inlining cycle-free parts during the solution of parity games, leading to earlier convergence. We also present an economic later-appearance-record construction that takes Emerson-Lei games to parity games, and show that it preserves DAG sub-structures; it follows that the proposed method can be used also for the accelerated solution of Emerson-Lei games.

1 Background

The analysis of infinite duration games is of central importance to various problems in theoretical computer science such as formal verification (model checking), logical reasoning (satisfiability checking), or automated program construction (reactive synthesis). Previous work has shown how fixpoint expressions can be used to characterize winning in such games, which in turn has enabled the development of symbolic game solving algorithms that circumvent the state-space explosion to some extent and therefore perform reasonably well in practice in spite of the high complexity of the considered problems.

In this work, we build on the close connection between game solving and fixpoint computation to obtain a method that accelerates the solution of parity games with acyclic sub-structures. Intuitively, we base our method on the observation that cycle-free parts in parity games can be dealt with summarily by using the computation of attractors in place of handling each node individually. In games that are largely cycle-free, but in which attraction along cycle-free parts can be evaluated without exploring most of their nodes, this significantly speeds up the game solution process.

The close relation between games and fixpoint expression has been well researched. On one hand, it has been shown that winning regions in games with various objectives can be characterized by fixpoint expressions [4, 3, 10, 5]. The result for parity games (e.g. [4]) arguably is the best-known case: it has been shown that the winning region in parity games with k priorities can be specified by a modal μ -calculus formula (known as Walukiewicz formulas) with alternation-depth k ; this result connects parity game solving and model checking for the μ -calculus. It corresponds to arrow (1) in the diagram below, which is meant to illustrate the context of our contribution.



A fruitful reduction in the converse direction is more recent. *Fixpoint games*, that is, parity games of a certain structure (defined in detail below), have been used to characterize the semantics of hierarchical

*Supported by the ERC Consolidator grant D-SynMA (No. 772459)

fixpoint equation systems in terms of games [2]. While this reduction (corresponding to arrow (2) in the diagram) incurs blow-up that is exponential in the size of the domain of the equation system, it still proved helpful since it allows for reasoning about nested least and greatest fixpoints in terms of (strategies in) parity games. Fixpoint games have also been used to lift the recent breakthrough quasipolynomial result for parity game solving to the evaluation of general fixpoint equation systems [8].

Later-appearance-record (LAR) constructions have been used to reduce games with richer winning objectives to parity games; this subsumes in particular the LAR-reduction for Emerson-Lei games (arrow (3)), which have recently garnered attention due to their succinctness and favorable compositionality properties.

In this context, the contribution of the current work is two-fold. First, we show how for parity games with cycle-free sub-structures, Walukiewicz’ construction can be adapted to use multi-step attraction along cycle-free parts (we call such parts *DAGs in games*) in place of one-step attraction. In consequence, the domain of the resulting fixpoint expression can be restricted to the parts of the game that are not cycle-free, thereby accelerating convergence in the solution of the fixpoint expression. If multi-step attraction can be evaluated without visiting most cycle-free nodes (which is the case, e.g., when cycle-free parts encode predicates that can be efficiently evaluated), this trick has been shown to significantly speed up game solution. Applications can be found in both model checking [6] and satisfiability checking [7] for generalized μ -calculi in which satisfaction / joint satisfiability of the modalities can be harder to verify than in the standard μ -calculus.

In addition to that, we propose a later-appearance-record construction for the transformation of Emerson-Lei games to parity games (corresponding to arrow (3) in the above diagram) and show that the reduction preserves cycle-free sub-structures, thereby enabling usage of the proposed acceleration method also for games with full Emerson-Lei objectives.

2 Games and Fixpoint Expressions

We start by introducing notions of games and fixpoint expressions pertaining to the acceleration method proposed in Sections 3 and 4 below.

Games. An *arena* is a graph $A = (V_\exists, V_\forall, E)$, consisting of a set $V = V_\exists \cup V_\forall$ of *nodes* and a set $E \subseteq V \times V$ of *moves*; furthermore, the set of nodes is partitioned into the sets V_\exists and V_\forall of nodes *owned* by player \exists and by player \forall , respectively. We write $E(v) = \{w \in V \mid (v, w) \in E\}$ for the set of nodes reachable from node $v \in V$ by a single move. We assume without loss of generality that every node has at least one outgoing edge, that is, that $E(v) \neq \emptyset$ for all $v \in V$. A *play* $\pi = v_0 v_1 \dots$ on A is a (finite or infinite) sequence of nodes such that $v_{i+1} \in E(v_i)$ for all $i \geq 0$. By abuse of notation we denote by A^ω the set of infinite plays on A and by A^+ the set of finite (nonempty) plays on A . A *strategy* for player $i \in \{\exists, \forall\}$ is a function $\sigma : V^* \cdot V_i \rightarrow V$ that assigns a node $\sigma(\pi v) \in V$ to every finite play πv that ends in a node $v \in V_i$. A strategy σ for player i is said to be *positional* if the prescribed moves depend only on the last nodes in plays; formally, this is the case if we have $\sigma(\pi v) = \sigma(\pi' v)$ for all $\pi, \pi' \in V^*$ and $v \in V_i$. A play $\pi = v_0 v_1 \dots$ then is *compatible* with a strategy σ for player $i \in \{\exists, \forall\}$ if for all $j \geq 0$ such that $v_j \in V_i$, we have $v_{j+1} = \sigma(v_0 v_1 \dots v_j)$. A strategy for player $i \in \{\exists, \forall\}$ *with memory* M is a tuple $\sigma = (M, m_0, \text{update} : M \times E \rightarrow M, \text{move} : V_i \times M \rightarrow V)$, where M is some set of *memory values*, $m_0 \in M$ is the initial memory value, the *update function* update assigns the outcome $\text{update}(m, e) \in M$ of updating the memory value $m \in M$ according to the effects of taking the move $e \in E$, and the moving function move prescribes a single move $(v, \text{move}(v, m)) \in E$ to every game node $v \in V_i$ that is owned by player i , depending on the memory value m . We extend update to finite plays π by putting $\text{update}(m, \pi) = m$ in the base case that

π consists of a single node, and by putting $\text{update}(m, \pi) = \text{update}(\text{update}(m, \tau), (v, w))$ if π is of the shape $\tau v w$, that is, contains at least two nodes. Then we obtain a strategy (without explicit memory) $\tau_\sigma : V^* \cdot V_i \rightarrow V$ from a strategy $\sigma = (M, m_0, \text{update} : M \times E \rightarrow M, \text{move} : V_i \times M \rightarrow V)$ with memory by putting, for all $v_0 v_1 \dots v_i \in V^* \cdot V_i$,

$$\tau_\sigma(v_0 v_1 \dots v_i) = \text{move}(v_i, \text{update}(m_0, v_0 v_1 \dots v_i)).$$

In this work we consider two types of objectives: *parity objectives* and, more generally, *Emerson-Lei objectives*.

Emerson-Lei objectives [9] are specified relative to a coloring function $\gamma_C : V \rightarrow 2^C$ (for some set C of colors) that assigns a set $\gamma_C(v) \subseteq C$ of colors to every node $v \in V$. A play $\pi = v_0 v_1 \dots$ then induces a sequence $\gamma_C(\pi) = \gamma_C(v_0) \gamma_C(v_1) \dots$ of sets of colors. Emerson-Lei objectives are given as Boolean formulas $\varphi_C \in \mathbb{B}(\{\text{Inf } c \mid c \in C\})$ over atoms of the shape $\text{Inf}(c)$; throughout, we write $\text{Fin}(c)$ for $\neg(\text{Inf}(c))$. Such formulas are interpreted over infinite sequences $\gamma_0 \gamma_1 \dots$ of sets of colors. We put $\gamma_0 \gamma_1 \dots \models \text{Inf } c$ if and only if there are infinitely many positions i such that $c \in \gamma_i$; satisfaction of Boolean operators is defined in the usual way. Then a play π on A satisfies the objective φ_C if and only if $\gamma_C(\pi) \models \varphi_C$ and we define the Emerson-Lei objective induced by γ_C and φ_C by putting

$$\alpha_{\gamma_C, \varphi_C} = \{\pi \in V^\omega \mid \gamma_C(\pi) \models \varphi_C\}.$$

An *Emerson-Lei game* is a tuple $G = (A, \alpha)$, where $A = (V_\exists, V_\forall, E)$ is an arena and α is an Emerson-Lei objective. A strategy σ is *winning* for player \exists at some node $v \in V$ if all plays that start at v and are compatible with σ satisfy the objective α . A strategy τ for player \forall is defined winning dually. The *winning region* Win_G of player \exists in the game G is the set of nodes for which there is a winning strategy for player \exists .

As a special case, we define *parity games* to be Emerson-Lei games (A, α) where α is the Emerson-Lei objective induced by a coloring function $\gamma : V \rightarrow 2^{\{p_1, \dots, p_k\}}$ that assigns exactly one of the colors $\{p_1, \dots, p_k\}$ to each game node, and the formula

$$\bigvee_{i \text{ even}} \text{Inf}(p_i) \wedge \bigwedge_{j > i} \text{Fin}(p_j),$$

expressing that for some even i , color p_i is visited infinitely often and no color p_j such that j is larger than i is visited infinitely often. We denote such games by $(V_\exists, V_\forall, E, \Omega : V \rightarrow \mathbb{N})$, where Ω is a priority function, assigning to each node $v \in V$ the number $\Omega(v) = i$ such that $\gamma(v) = \{p_i\}$.

Fixpoint Expressions. Let U be a set, k a natural number (we assume without loss of generality that k is even) and let $f : \mathcal{P}(U)^k \rightarrow \mathcal{P}(U)$ be a monotone function (such that $f(W_1, \dots, W_k) \subseteq f(W'_1, \dots, W'_k)$ whenever $W_i \subseteq W'_i$ for all $1 \leq i \leq k$). This data induces a fixpoint expression

$$e = vX_k. \mu X_{k-1}. \dots \mu X_1. f(X_1, \dots, X_k).$$

The semantics of such fixpoint expressions is defined as usual, following the Knaster-Tarski fixpoint theorem. Formally, we put

$$\text{LFP } g = \bigcap \{W \subseteq U \mid g(W) \subseteq W\} \qquad \text{GFP } g = \bigcup \{W \subseteq U \mid W \subseteq g(W)\}$$

for monotone functions $g : \mathcal{P}(U) \rightarrow \mathcal{P}(U)$ and define

$$\llbracket e \rrbracket_i(X_{i+1}, \dots, X_k) = \begin{cases} \text{LFP}(\lambda X_i. \llbracket e \rrbracket_{i-1}(X_i, X_{i+1}, \dots, X_k)) & \text{if } i \text{ is odd} \\ \text{GFP}(\lambda X_i. \llbracket e \rrbracket_{i-1}(X_i, X_{i+1}, \dots, X_k)) & \text{if } i \text{ is even} \end{cases}$$

for $1 \leq i \leq k$ and $X_{i+1}, \dots, X_k \subseteq V$, where $\llbracket e \rrbracket_0(X_1, \dots, X_k) = f(X_1, \dots, X_k)$. Then we say that $v \in V$ is *contained in e* (denoted $v \in e$) if and only if $v \in \llbracket e \rrbracket_k$; by abuse of notation, we denote $\llbracket e \rrbracket_k$ just by e .

Relation between Games and Fixpoint Expressions. The winning region in parity games can be specified by a fixpoint expression over the underlying game arena. To see how this works, let $G = (V_\exists, V_\forall, E, \Omega)$ be a parity game and define the following notation. Let $\Omega^{-1}(i) = \{v \in V \mid \Omega(v) = i\}$ denote the set of nodes that have priority i . The *controllable predecessor function* $\text{Cpre} : \mathcal{P}(V) \rightarrow \mathcal{P}(V)$ that computes one-step attraction for player \exists in G is defined, for $X \subseteq V$, by putting

$$\text{Cpre}(X) = \{v \in V_\exists \mid E(v) \cap X \neq \emptyset\} \cup \{v \in V_\forall \mid E(v) \subset X\}.$$

The characterization then is given by the fixpoint expression

$$\text{parity}_G = \nu X_k. \mu X_{k-1}. \dots \mu X_1. \bigvee_{1 \leq i \leq k} \Omega^{-1}(i) \wedge \text{Cpre}(X_i). \quad (1)$$

The characterization result (e.g. [4]) then is stated as follows.

Lemma 1. *We have $\text{Win}_G = \text{parity}_G$.*

Similar fixpoint characterizations of winning regions have been given for games with more involved objectives, including GR[1] conditions [3], Rabin and Streett objectives [10], and most recently, also for full Emerson-Lei objectives [5]. While the characterizations of winning according to these more general objectives involve fixpoint expressions of a more general format (given by hierarchical systems of fixpoint equations) than the one we have defined above, the fixpoint for the winning region in parity games will be sufficient for the purposes of this work. All mentioned fixpoint characterizations of winning regions allow for (symbolic) solution of games by computing nested fixpoints, using for instance standard fixpoint iteration [12], or using more recent techniques that employ universal trees to obtain quasipolynomial algorithms to compute nested fixpoints [8, 1].

For the converse direction, that is, going from fixpoints to games, it has recently been shown (in a more general framework than the one used in this work, see [2] for details) that the semantics of fixpoint expressions can be equivalently given in terms of parity games. The *fixpoint game* that is associated with a fixpoint expression

$$e = \nu X_k. \mu X_{k-1}. \dots \mu X_1. f(X_1, \dots, X_k)$$

over U is the parity game $G_e = (V_\exists, V_\forall, E, \Omega : V \rightarrow \{1, \dots, k\})$ played over the sets $V_\exists = U$ and $V_\forall = \mathcal{P}(U)^k \cup (U \times [k])$ of nodes; moves and priorities are given by the following table, abbreviating (U_1, \dots, U_k) by \bar{U} .

node	owner	priority	allowed moves
$u \in U$	\exists	0	$\{\bar{U} \in \mathcal{P}(U)^k \mid v \in f(\bar{U})\}$
$\bar{U} \in \mathcal{P}(U)^k$	\forall	0	$\{(u, i) \in U \times [k] \mid u \in V_i\}$
$(u, i) \in U \times [k]$	\forall	i	$\{u\}$

When the game reaches a node $u \in U$, player \exists thus has to provide a tuple $\bar{U} \in \mathcal{P}(U)^k$ such that $f(\bar{U})$. Player \forall in turn can challenge the provided tuple by moving to any (u', i) such that $u' \in U_i$, and then continuing the with $u' \in U$. Crucially, such a sequence triggers the priority i which is even number if X_i belongs to a greatest fixpoint in e , and odd otherwise. We observe that the number of nodes in G_e is exponential in $|U|$; however, player \exists owns only $|U|$ of these nodes.

Lemma 2 ([2]). *We have $e = \text{Win}_{G_e}$.*

While fixpoint games enable reasoning about fixpoint expressions by reasoning about parity games, the exponential size of fixpoint games makes it unfeasible to evaluate fixpoint expressions by solving the associated fixpoint games.

3 Game Arenas with DAG Sub-structures

In this section we introduce notions related to cycle-free sub-structures in (parity) games that will be central to the upcoming acceleration result.

Definition 3. Let $G = (V_\forall, V_\exists, E)$ be a game arena. A set $W \subseteq V$ of nodes is a *DAG* (directed acyclic graph) if it does not contain an E -cycle; then there is no play of G that eventually stays within W forever. A DAG is not necessarily connected, that is, it may consist of several subgames of G . Given a DAG $W \subseteq V$, we write $V' = V \setminus W$ and refer to the set V' as *real* nodes (with respect to W). A DAG is *positional* if for each existential player node $w \in W \cap V_\exists$ in it, there is exactly one real node $v \in V'$ from which w is reachable without visiting any other real node.

We will be interested in parity games over game arenas with DAG structures, and in the computation of attraction (alternating reachability) within such DAGs. To this end, we introduce the following notation of DAG attraction in parity games.

Definition 4. Let $G = (V_\forall, V_\exists, E, \Omega)$ be a parity game with priorities $\{1, \dots, k\}$. Given a DAG W and k sets $\bar{V} = (V_1, \dots, V_k)$ of real nodes and a real node $v \in V'$, we say that player \exists can *attract* to \bar{V} from v within W if there is a positional strategy σ for player \exists such that for all plays π that start at v and adhere to σ , the first node v' in π such that $v' \neq v$ and $v' \in V'$ is contained in V_p , where p is the maximal priority that is visited by the part of π that leads from v to v' .

Given a DAG W , we define the *DAG attractor function* $\text{DAttr}_W : \mathcal{P}(V')^k \rightarrow \mathcal{P}(V')$ by

$$\text{DAttr}_W(\bar{V}) = \{v \in V \mid \text{player } \exists \text{ can attract to } \bar{V} \text{ from } v \text{ within } W\} \quad (2)$$

for $\bar{V} = (V_1, \dots, V_k) \in \mathcal{P}(V')^k$. We assume a bound t_{Attr_W} on the runtime of computing, for any input $\bar{V} \in \mathcal{P}(V')^k$, the dag attractor of \bar{V} through W .

We always have $t_{\text{Attr}_W} \leq |E|$, using a least fixpoint computation to check for alternating reachability, thereby possibly exploring all DAG edges of the game. The method that we propose in the next section can play out its strength best in parity games that have large DAGs that can be efficiently evaluated (more formally, this typically means that we have both $|V'| < \log |V|$ and $t_{\text{Attr}_W} < \log |V|$).

4 Large-step Solving for Parity Games

In this section we show that a parity game with DAG sub-structure W can be solved by computing a fixpoint of the DAG attractor function, using $V' = V \setminus W$ as domain of the fixpoint computation, rather than V .

So let $G = (V_\exists, V_\forall, E, \Omega)$ be a parity game with (positional) DAG W . Recall (from Lemma 1) that the winning region in G is given by

$$\text{Win}_G = \nu X_k. \mu X_{k-1}. \dots \mu X_1. \bigvee_{1 \leq i \leq k} \Omega^{-1}(i) \wedge \text{Cpre}(X_i),$$

where Cpre operates on subsets X_i of V . We show that

$$\text{Win}_G \cap V' = \nu Y_k. \mu Y_{k-1}. \dots \mu Y_1. \text{DAttr}_W(Y_1, \dots, Y_k), \quad (3)$$

pointing out that $\text{DAttr}_W(Y_1, \dots, Y_k)$ is a function over subsets Y_1, \dots, Y_k of V' . Therefore Equation (3) shows how attraction can be reduced to non-DAG nodes (from V') when solving G ; however, the price for this is that the function DAttr_W of which the fixpoint is computed is complex since it computes multi-step attraction within a DAG in place of one-step attraction as encoded by Cpre .

We make this intuition precise and prove the main result (that is, Equation (3)) as follows.

Lemma 5. *Let G be a parity game with priorities 1 to k and set V of nodes, and let W be a positional DAG in G and let $m := |V \setminus W|$. Then $\text{Win}_G \cap V'$ can be computed with $\mathcal{O}(m^{\log k})$ computations of a DAG attractor; if $k < \log m$, then $\text{Win}_G \cap V'$ can be computed with a number of DAG attractor computations that is polynomial in m .*

Proof. Put $V' = V \setminus W$. First we show that the winning regions in G can be computed as a fixpoint expression (with alternation-depth k) of the DAG attractor function given in Equation (2). Without loss of generality, assume that k is an even number.

We define the fixpoint expression

$$d = \nu X_k. \mu X_{k-1}. \dots \mu X_1. \text{DAttr}_W(X_1, \dots, X_k),$$

noting that the fixpoint variables X_i in d range over subsets of V' rather than over subsets of V . It has been shown in [8] how universal trees can be used to compute the nested fixpoint d with $\mathcal{O}(m^{\log k})$ (that is, quasipolynomially many) iterations of the function DAttr_W . The same paper shows that if $k < \log m$, then d can be computed with a polynomial (in m) number of computations of a dag attractor. It remains to show that d is the winning region (restricted to V') of player \exists in G .

\Rightarrow Let s be a strategy with which player \exists wins from all nodes in their winning region in G_d . We define a positional strategy t for player \exists in G . The definition proceeds in steps that lead from a real node to the next real nodes. Given a real node $v \in V'$ that is contained in the winning region of player \exists in G_d , we have $s(v) = \bar{V} \in \mathcal{P}(V')^{k+1}$ such that player \exists can attract to \bar{V} from v within W . For the dag part of G that starts at v , define the strategy t to simply use the positional strategy that attracts to \bar{V} from v within W .

To see that t is a winning strategy, let π be a play of G that is compatible with t . Then π induces a play τ of G_d that is compatible with s : For any two positions i and i' such that $i < i'$, both $\pi(i)$ and $\pi(i')$ are real nodes and there is no $i < j < i'$ such that $\pi(j)$ is a real node, let p be the maximal priority that is visited by the partial play from $\pi(i)$ to $\pi(i')$. By construction we have $s(\pi(i)) = (V_0, \dots, V_p, \dots, V_k)$ where V_p contains $\pi(i')$. Thus the partial play of G that leads from $\pi(i)$ to $\pi(i')$ induces the partial play $\tau_i = \pi(i)s(\pi(i))(\pi(i'), p)\pi(i')$ of G_d . We define τ to be the concatenation of all partial plays τ_i . By construction, τ is compatible with s ; as s is a winning strategy, player \exists wins τ . Furthermore, the maximal priorities that are visited infinitely often in τ and π agree. Thus player \exists wins π , as required.

\Leftarrow Let t be a positional strategy for player \exists in G with which they win from all nodes in their winning region. We define a positional strategy s for player \exists in G_d as follows. For each node $v \in V'$ that is contained in the winning region of player \exists in G , t yields a positional strategy that attracts from v to some $\bar{V} = (V_0, \dots, V_k) \in \mathcal{P}(V')^{k+1}$ within W : For all v' and all $0 \leq p \leq k$ such that there is a partial play of G that is compatible with t , starts at v , ends at v' , and in between visits only nodes

from W , and in which the maximal priority that is visited is p , add v' to V_p . Put $s(v) = \bar{V}$. As v attracts from v to \bar{V} within W by construction, s is a valid strategy.

To see that s is a winning strategy, let τ be a play of G_d that starts a real node from the winning region of player \exists in G and that is compatible with s . For each partial subplay $v_i s(v_i)(v_{i+1}, p_i)v_{i+1}$ of τ , we define π_i to be some partial play of G that is compatible with t , starts at v_i , ends at v_{i+1} , and in between visits only nodes from W , and in which the maximal priority that is visited is p_i . Such a play exists by construction of s . Let π denote the concatenation of all π_i . Again, the maximal priorities that are visited infinitely often in τ and π agree. As t is a winning strategy, player \exists wins both τ and π . □

In case that $m < \log n$ and $t_{\text{Attr}_W} \in \mathcal{O}(\log n)$ (that is, DAG attractability can be decided without exploring most of the DAG nodes), this trick leads to exponentially faster game solving.

5 Large-step Solving for Emerson-Lei Games

In this section, we show how games with an Emerson-Lei objective can be transformed to equivalent parity games, using an economic variant of the later-appearance-record (LAR) construction; we also show that this transformation preserves dag sub-structures. Thereby we enable usage of the acceleration method for parity games from the previous section also for Emerson-Lei games.

The LAR reduction annotates game nodes with permutations that act as a memory and record the order in which colors have been recently visited; this allows to identify the set of colors that is visited infinitely often. A parity condition is used to check whether this set satisfies the Emerson-Lei objective. A slight difference to previous LAR constructions for Emerson-Lei automata and games [11, 9] is that our reduction moves at most one color within the LAR in each game step, thereby allowing to bound the branching of memory updates along DAG parts of games.

Before we present the formal reduction, we first fix a set C of colors and introduce notation for permutations over C . We let $\Pi(C)$ denote the set of permutations over C , and for a permutation $\pi \in \Pi(C)$ and a position $1 \leq i \leq |C|$, we let $\pi(i) \in C$ denote the element at the i -th position of π . Let π_0 be some fixed element of $\Pi(C)$. For $D \subseteq C$ and $\pi \in \Pi(C)$, we let $\pi @ D$ denote the permutation that is obtained from π by moving the element of D that occurs at the right-most position in π to the front of π ; for instance, for $C = \{a, b, c, d\}$ and $\pi = (a, d, c, b) \in \Pi(C)$, we have $\pi @ \{a, d\} = \pi @ \{d\} = (d, a, c, b)$ and $(d, a, c, b) @ \{a, d\} = \pi$. Crucially, restricting the reordering to single colors, rather than sets of colors, ensures that for each $\pi \in \Pi(C)$, there are only $|C|$ many π' such that $\pi @ D = \pi'$ for some $D \subseteq C$. Given a permutation $\pi \in \Pi(C)$ and an index $1 \leq i \leq |C|$, we furthermore let $\pi[i]$ denote the set of colors that occur in one of the first i positions in π .

Definition 6. Let $G = (V_\exists, V_\forall, E, \gamma_C, \varphi_C)$ be an Emerson-Lei game with set C of colors. We define the parity game

$$P(G) = (V_\exists \times \Pi(C), V_\forall \times \Pi(C), E', \Omega)$$

by putting $E'(v, \pi) = \{(w, \pi @ \gamma_C(v, w)) \mid (v, w) \in E\}$ for $(v, \pi) \in V \times \Pi(C)$, and

$$\Omega(v, \pi) = \begin{cases} 2p & \pi[p] \models \varphi_C \\ 2p+1 & \pi[p] \not\models \varphi_C \end{cases}$$

for $(v, \pi) \in V \times \Pi(C)$. Here, p denotes the right-most position in π that contains some color from $\gamma_C(v)$.

Lemma 7. *For all $v \in V$, we have $v \in \text{Win}_G$ if and only if $(v, \pi_0) \in \text{Win}_{P(G)}$.*

Proof. For a finite play $\tau = (v_0, \pi_0)(v_1, \pi_1) \dots (v_n, \pi_n)$ of $P(G)$, we let $p(\tau)$ denote the right-most position in π_0 such that $\pi_0(p(\tau)) \in \gamma_C(v_0 v_1 \dots v_n)$. By slight abuse of notation, we let $\Omega(\tau)$ denote the maximal Ω -priority that is visited in τ , having $\Omega(\tau) = 2(p(\tau))$ if $\pi_0[p(\tau)] \models \varphi_C$ and $\Omega(\tau) = 2(p(\tau)) + 1$ if $\pi_0[p(\tau)] \not\models \varphi_C$.

\Rightarrow Let $\sigma = (\Pi(C), \text{update}_\sigma, \text{move}_\sigma)$ be a strategy with memory $\Pi(C)$ for player \exists in G with which they win every node from their winning region. It has been shown in a previous LAR reduction for Emerson-Lei games [9] that winning strategies with this amount of memory always exist. We define a strategy $\rho = (\Pi(C), \text{update}_\rho, \text{move}_\rho)$ with memory $\Pi(C)$ for player \exists in $P(G)$ by putting $\text{update}_\rho(\pi, ((v, \pi'), (w, \pi''))) = \text{update}_\sigma(\pi, (v, w))$ and $\text{move}_\rho((v, \pi'), \pi) = (w, \pi @ \gamma_C(v, w))$ where $w = \text{move}_\sigma(v, \pi)$. Thus ρ updates the memory and picks moves just as σ does, but also updates the permutation component in $P(G)$ according to the taken moves; hence ρ is a valid strategy.

We show that ρ wins a node (v, π) in $P(G)$ whenever v is in the winning region of player \exists in G . To this end, let $\tau = (v_0, \pi_0)(v_1, \pi_1) \dots$ be a play of $P(G)$ that starts at $(v_0, \pi_0) = (v, \pi)$ and is compatible with ρ . By construction, $\pi = v_0 v_1 \dots$ is a play that is compatible with σ . Since σ is a winning strategy for player \exists , we have $\gamma_C(\pi) \models \varphi_C$. There is a number i such that all colors that appear in π from position i on occur infinitely often. Let p be the number of colors that appear infinitely often in π . It follows by definition of $\pi @ D$ for $D \subseteq C$ (which moves the single right-most element of π that is contained in D to the very front of π), that there is a position $j \geq i$ such that the left-most p elements of π_j are exactly the colors occurring infinitely often in π (and all colors to the right of $\pi_j(p)$ are never visited from position j on). It follows that from position j on, τ never visits a priority larger than $2p$. To see that τ infinitely often visits priority $2p$ we note that $\pi_{j'}[p] \models \varphi_C$ for any $j' > j$, so it suffices to show that p infinitely often is the rightmost position in the permutation component of τ that is visited. This is the case since, from position j on, the p -th element in the permutation component of τ cycles fairly through all colors that are visited infinitely often by π .

\Leftarrow Let ρ be a positional strategy for player \exists in $P(G)$ with which they win every node from their winning region. We define a strategy $\sigma = (\Pi(C), \text{update}_\sigma, \text{move}_\sigma)$ with memory $\Pi(C)$ for player \exists in G by putting $\text{update}_\sigma(\pi, (v, w)) = \pi @ \gamma_C(v, w)$ and $\text{move}_\sigma(v, \pi) = w$ where w is such that $\rho(v, \pi) = (w, \pi @ \gamma_C(v, w))$. Thus σ updates the memory and picks moves just as plays that follow ρ do.

We show that σ wins a node v in G whenever (v, π) is in the winning region of player \exists in $P(G)$. To this end, let $\pi = v_0 v_1 \dots$ be a play of G that starts at $v_0 = v$ and is compatible with σ . By construction, π induces a play $\tau = (v_0, \pi_0)(v_1, \pi_1) \dots$ with $(v_0, \pi_0) = (v, \pi)$ and $\pi_{i+1} = \pi_i @ \gamma_C(v_i, v_{i+1})$ for $i \geq 0$ that is compatible with ρ . Since ρ is a winning strategy for player \exists , the maximal priority in it is even (say $2p$). Again, p is the position such that the left-most p elements in the permutation component of ρ from some point on contain exactly the colors that are visited infinitely often by π . It follows that $\gamma_C(\pi) \models \varphi_C$.

□

The LAR reduction from Definition 6 preserves DAG sub-structures:

Lemma 8. *Let G be an Emerson-Lei games with set C of colors. If W is a (positional) DAG in G , then $W \times \Pi(C)$ is a (positional) DAG in $P(G)$.*

Proof. The claim follows immediately since the LAR reduction just annotates game nodes with additional memory, meaning that each game node v in G is replaced with copies (v, π) in $P(G)$ so that all cycles in $P(G)$ have corresponding cycles in G . \square

6 Conclusion

We have presented a method for solving parity games with DAG sub-structures by computing nested fixpoints over the non-DAG nodes only, thereby intuitively accelerating fixpoint computation by summarizing cycle-free parts during the solution of parity games. This can significantly reduce the domain of the game solving process, and in cases where DAG sub-structures in addition can be evaluated without exploring all nodes in them, it improves the time complexity of parity game solving. Furthermore, we have proposed a later-appearance-record construction with linear branching on the memory values that transforms Emerson-Lei games to parity games, and have shown that this transformation preserves DAG sub-structures, enabling usage of the proposed method also for the solution of games with general Emerson-Lei objectives.

References

- [1] André Arnold, Damian Niwinski, and Pawel Parys. A quasi-polynomial black-box algorithm for fixed point evaluation. In Christel Baier and Jean Goubault-Larrecq, editors, *29th EACSL Annual Conference on Computer Science Logic, CSL 2021, January 25-28, 2021, Ljubljana, Slovenia (Virtual Conference)*, volume 183 of *LIPIcs*, pages 9:1–9:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. URL: <https://doi.org/10.4230/LIPIcs.CSL.2021.9>, doi:10.4230/LIPIcs.CSL.2021.9.
- [2] Paolo Baldan, Barbara König, Christina Mika-Michalski, and Tommaso Padoan. Fixpoint games on continuous lattices. *Proc. ACM Program. Lang.*, 3(POPL):26:1–26:29, 2019. doi:10.1145/3290339.
- [3] Roderick Bloem, Barbara Jobstmann, Nir Piterman, Amir Pnueli, and Yaniv Sa’ar. Synthesis of reactive(1) designs. *J. Comput. Syst. Sci.*, 78(3):911–938, 2012. doi:10.1016/j.jcss.2011.08.007.
- [4] Florian Bruse, Michael Falk, and Martin Lange. The fixpoint-iteration algorithm for parity games. In *International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2014*, volume 161 of *EPTCS*, pages 116–130, 2014. doi:10.4204/EPTCS.161.12.
- [5] Daniel Hausmann, Mathieu Lehaut, and Nir Piterman. Symbolic reactive synthesis for the safety and el-fragment of LTL. *CoRR*, abs/2305.02793, 2023. URL: <https://doi.org/10.48550/arXiv.2305.02793>, arXiv:2305.02793, doi:10.48550/ARXIV.2305.02793.
- [6] Daniel Hausmann and Lutz Schröder. Game-based local model checking for the coalgebraic mu-calculus. In Wan J. Fokkink and Rob van Glabbeek, editors, *30th International Conference on Concurrency Theory, CONCUR 2019, August 27-30, 2019, Amsterdam, the Netherlands*, volume 140 of *LIPIcs*, pages 35:1–35:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. URL: <https://doi.org/10.4230/LIPIcs.CONCUR.2019.35>, doi:10.4230/LIPIcs.CONCUR.2019.35.
- [7] Daniel Hausmann and Lutz Schröder. Optimal satisfiability checking for arithmetic μ -calculi. In Mikolaj Bojanczyk and Alex Simpson, editors, *Foundations of Software Science and Computation Structures - 22nd International Conference, FOSSACS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings*, volume 11425 of *Lecture Notes in Computer Science*, pages 277–294. Springer, 2019. doi:10.1007/978-3-030-17127-8_16.
- [8] Daniel Hausmann and Lutz Schröder. Quasipolynomial computation of nested fixpoints. In Jan Friso Groote and Kim Guldstrand Larsen, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 27th International Conference, TACAS 2021, Held as Part of the European Joint Conferences on Theory and*

- Practice of Software, ETAPS 2021, Luxembourg City, Luxembourg, March 27 - April 1, 2021, Proceedings, Part I*, volume 12651 of *Lecture Notes in Computer Science*, pages 38–56. Springer, 2021. doi:10.1007/978-3-030-72016-2_3.
- [9] Paul Hunter and Anuj Dawar. Complexity bounds for regular games. In *30th International Symposium on Mathematical Foundations of Computer Science*, volume 3618 of *Lecture Notes in Computer Science*, pages 495–506. Springer, 2005. doi:10.1007/11549345_43.
 - [10] Nir Piterman and Amir Pnueli. Faster solutions of rabin and streett games. In *21th IEEE Symposium on Logic in Computer Science (LICS 2006), 12-15 August 2006, Seattle, WA, USA, Proceedings*, pages 275–284. IEEE Computer Society, 2006. doi:10.1109/LICS.2006.23.
 - [11] Florian Renkin, Alexandre Duret-Lutz, and Adrien Pommellet. Practical “paritizing” of emerson-lei automata. In Dang Van Hung and Oleg Sokolsky, editors, *Automated Technology for Verification and Analysis - 18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19-23, 2020, Proceedings*, volume 12302 of *Lecture Notes in Computer Science*, pages 127–143. Springer, 2020. doi:10.1007/978-3-030-59152-6_7.
 - [12] Helmut Seidl. Fast and simple nested fixpoints. *Universität Trier, Mathematik/Informatik, Forschungsbericht*, 96-05, 1996.