

Combinatorial Games

Achim Blumensath

last update: 7th March 2021

Contents

1	Positional games	1
2	Reachability games	3
3	Conway games	12
4	Gale-Steward games	14
5	Regular games and parity games	16
	References	37
	Symbol Index	38
	Index	39

1 Positional games

Many games like go, chess, and checkers can be modelled as a directed graph where the vertices represent the different *states*, or *positions*, of the game and

(where the label \diamond denotes positions for Player \diamond and \square positions for Player \square) the shaded part constitutes the winning region for Player \diamond , while the winning region for Player \square is empty.

The winning regions of a reachability game are easy to compute recursively: a player wins from some position v if either v belongs to him and at least one outgoing edge leads to a winning position, or it belongs to his opponent and all the outgoing edges lead to winning positions. To define this formally we introduce the following notation. For $X \subseteq V$, we set

$$\begin{aligned}\diamond X &:= \{v \in V \mid \langle v, w \rangle \in E \text{ for some } w \in X\}, \\ \square X &:= \{v \in V \mid \langle v, w \rangle \in E \text{ implies } w \in X\}.\end{aligned}$$

We denote the opponent of Player σ by $\bar{\sigma}$, i.e., $\overline{\diamond} := \square$ and $\overline{\square} := \diamond$. Then we can define the winning region for Player σ as the least set \mathcal{W}_σ such that

$$\mathcal{W}_\sigma = (V_\sigma \cap \diamond \mathcal{W}_\sigma) \cup (V_{\bar{\sigma}} \cap \square \mathcal{W}_\sigma).$$

Hence, \mathcal{W}_σ is the least fixed point of the following *step function*

$$\text{Step}_\sigma(X) := (V_\sigma \cap \diamond X) \cup (V_{\bar{\sigma}} \cap \square X).$$

It is easy to see that $\text{Step}_\sigma(X)$ contains all the positions from which Player σ can ensure that in the next step the game either reaches some position in X or the game ends with a win for him. Thus, by iterating the step function we obtain the set of all positions from which the player either wins or eventually reaches a position in X . This iteration of Step_σ is called the σ -*attractor* of the set X . The formal definition is

$$\text{Attr}_\sigma(X) := \bigcup_{\alpha} \text{Step}_\sigma^\alpha(X),$$

where α ranges over all ordinals (actually, it is sufficient to take the union for all $\alpha < |V|^+$) and $\text{Step}_\sigma^\alpha$, the α -th iteration of the step function, is defined as

follows

$$\begin{aligned}\text{Step}_\sigma^0(X) &:= X, \\ \text{Step}_\sigma^{\alpha+1}(X) &:= \text{Step}_\sigma(\text{Step}_\sigma^\alpha(X)), \\ \text{Step}_\sigma^\delta(X) &:= \bigcup_{\alpha < \delta} \text{Step}_\sigma^\alpha(X), \quad \text{for limit ordinals } \delta.\end{aligned}$$

Lemma 2.1. *Player σ has a positional strategy s such that, every play p that starts in some position $v \in \text{Attr}_\sigma(X)$ and that conforms to s is winning or contains some position from X .*


Proof. Let $v \in \text{Attr}_\sigma(X)$. Then $v \in \text{Step}_\sigma^\alpha(X)$, for some ordinal α . We prove the claim by induction on α .

If $\alpha = 0$, then $v \in X$ and the claim is trivial. For the successor step, suppose that $\alpha = \beta + 1$. Then $v \in \text{Step}_\sigma(\text{Step}_\sigma^\beta(X))$ and Player σ has a strategy to either win in one step, or to reach some vertex of $\text{Step}_\sigma^\beta(X)$. In the first case we are done. In the second one, we can continue with the strategy from the inductive hypothesis. Finally, if α is a limit ordinal, then $v \in \text{Step}_\sigma^\beta(X)$, for some $\beta < \alpha$, and the claim follows immediately from the inductive hypothesis. \square

Using the notion of an attractor we can define a measure for how long it takes from a given position to win. The *rank* of a position v is the least ordinal α such that

$$v \in \text{Step}_\sigma^{\alpha+1}(\emptyset).$$

If there is no such ordinal, we set the rank to ∞ . In the above example the ranks for Player \diamond are



Next let us take a look at complements of attractors. We call a subset $U \subseteq V$ a σ -trap if $\text{Step}_\sigma(V \setminus U) \subseteq V \setminus U$, that is, if the opponent can ensure that Player σ never leaves the set U once the game has entered it. An easy way to find traps is by computing attractors.

Lemma 2.2. *A set U is a σ -trap if, and only if, it is of the form*

$$U = V \setminus \text{Attr}_\sigma(X), \quad \text{for some } X.$$

Proof. (\Leftarrow) is obvious since $\text{Step}_\sigma(\text{Attr}_\sigma(X)) \subseteq \text{Attr}_\sigma(X)$. For (\Rightarrow), note that $\text{Step}_\sigma(V \setminus U) \subseteq V \setminus U$ implies $U = V \setminus \text{Attr}_\sigma(V \setminus U)$. \square

Lemma 2.3. *for every σ -trap U , Player $\bar{\sigma}$ has a positional strategy s ensuring that, starting from any vertex $v \in U$, the game never leaves U .*

Proof. Setting $A := V \setminus U$, we know that $A = \text{Attr}_\sigma(X)$, for some set X . Hence,

$$\begin{aligned} V_\sigma \setminus A &= V_\sigma \setminus \text{Step}_\sigma(A) = V_\sigma \setminus \diamond A \\ \text{and } V_{\bar{\sigma}} \setminus A &= V_{\bar{\sigma}} \setminus \text{Step}_{\bar{\sigma}}(A) = V_{\bar{\sigma}} \setminus \square A. \end{aligned}$$

In particular, (i) no vertex $v \in V_\sigma \cap U$ has an outgoing edge leading to a vertex in A and (ii) every vertex $v \in V_{\bar{\sigma}} \cap U$ has at least one outgoing edge $\langle v, u \rangle \in E$ with $u \in U$. Consequently, if the game starts in some vertex $v \in U$ Player σ can never move into A , while Player $\bar{\sigma}$ always has the option to stay in U . \square

Theorem 2.4. *Positional reachability games are positionally determined with winning regions $\mathcal{W}_\sigma := \text{Attr}_\sigma(\emptyset)$.*

Proof. Let s be the strategy for Player σ from Lemma 2.1 for $\mathcal{W}_\sigma = \text{Attr}_\sigma(\emptyset)$, and let t be the strategy from Lemma 2.3 for $V \setminus \mathcal{W}_\sigma$. As no play can ever reach a position in \emptyset it follows that s is winning for Player σ from every position $v \in \mathcal{W}_\sigma$. While, t ensures that Player $\bar{\sigma}$ does not lose when starting from a position in $V \setminus \mathcal{W}_\sigma$. This implies that \mathcal{W}_σ is the winning region for Player σ and that, for $V \setminus (\mathcal{W}_\diamond \cup \mathcal{W}_\square)$, each player has a positional strategy ensuring that he does not lose. \square

Having shown that winning regions exist, we next take a look at how to efficiently compute them.

Theorem 2.5. *The winning regions of a finite positional reachability game can be computed in linear time.*

Before presenting the linear time algorithm we start with a simpler, non-linear version that is just a direct translation of the definition of an attractor.

```
function Win( $v, \sigma$ )    // Does Player  $\sigma$  win from position  $v$ ?
  if  $v \in V_\sigma$  then
    if there is an edge  $v \rightarrow u$  with Win( $u, \sigma$ ) then
      return true
    else
      return false
  if  $v \in V_{\bar{\sigma}}$  then
    if for every edge  $v \rightarrow u$  we have Win( $u, \sigma$ ) then
      return true
    else
      return false
end
```

There are several obvious problems with this algorithm. First of all, it might not terminate if the game graph contains a cycle. And secondly, it is very inefficient (exponential time) as it does not remember if it has already computed the winner of a position and recomputes this information every time. There is a rather straightforward fix for both of these issues: we introduce an array where we store whether we have already visited a position and who the winner is. (Thus each entry can have one of four values: (i) not visited yet, (ii) already visited, but we do not know the winner yet, (iii) Player \diamond wins, and (iv) Player \square wins.) With this modification, the algorithm will run in quadratic time.

To improve the runtime to linear, we need to be more clever. We introduce two more array with auxiliary data that helps us to avoid unnecessary work.

Input: game $\langle V_{\diamond}, V_{\square}, E \rangle$
Output: array containing the winner for every position

```

// initialise auxiliary arrays

forall  $v \in V$  do
   $\text{win}[v] := \perp$  // the winner of the position
   $P[v] := \emptyset$  // the set of predecessors of  $v$ 
   $n[v] := 0$  // the number of (not yet processed)
  // successors of  $v$ 
end

forall  $\langle u, v \rangle \in E$  do
   $P[v] := P[v] \cup \{u\}$ 
   $n[u] := n[u] + 1$ 
end

// compute the winning regions

forall  $v \in V_{\diamond}$  do
  if  $n[v] = 0$  then Propagate( $v, \square$ )
forall  $v \in V_{\square}$  do
  if  $n[v] = 0$  then Propagate( $v, \diamond$ )
return win

procedure Propagate( $v, \sigma$ ) =
  if  $\text{win}[v] \neq \perp$  then return
   $\text{win}[v] := \sigma$ 
  forall  $u \in P[v]$  do
     $n[u] := n[u] - 1$ 
    if  $u \in V_{\sigma}$  or  $n[u] = 0$  then Propagate( $u, \sigma$ )
  end
end

```

To see that this algorithm works in linear time (in the number of positions plus the number of edges), note that the body of the procedure Propagate

(except for the first line) is executed exactly once for each vertex v . Furthermore, the loop in Propagate is executed once for each incoming edge which means that, in total, it is executed at most as many times as there are edges in the game. Since the precomputation steps are also linear in the number of vertices or edges, it follows that so is the total runtime of the algorithm.

It remains to show that the algorithm really computes the winning regions \mathcal{W}_{\diamond} and \mathcal{W}_{\square} . To see this it is sufficient to note that, every time the procedure Propagate(v, σ) is called and we still have $w[v] = \perp$, then the vertex v really belongs to the winning region for Player σ . This is clear for the two calls of Propagate in the main part of the algorithm, where only vertices v without successors are considered. For the recursive call inside the body of Propagate we need to distinguish two cases. If $u \in V_{\sigma}$ and v already belongs to \mathcal{W}_{σ} then Player σ can take the edge from u to v to win. Hence, $u \in \mathcal{W}_{\sigma}$. Otherwise, we have $u \in V_{\bar{\sigma}}$ and $n[u] = 0$, which means that we already know for all successors w_0, \dots, w_n of u to which region they belong. If there is some $w_i \in \mathcal{W}_{\bar{\sigma}}$ then we have already called Propagate($u, \bar{\sigma}$) when processing w_i and $\text{win}[u]$ is already set. Otherwise, all successors belong to \mathcal{W}_{σ} , which means that u belongs to it as well.

Horn formulae

As an application of reachability games let us take a look at the satisfiability problem for propositional Horn formulae. Such a formula is an implication of the form

$$A_1 \wedge \dots \wedge A_n \rightarrow B,$$

where we allow both the left-hand side and the right-hand side to be empty, i.e., we allow implications of the form $\top \rightarrow B$ and $A_1 \wedge \dots \wedge A_n \rightarrow \perp$. We are interested in deciding in whether a given set of such formulae is satisfiable. Note that such a set is always satisfiable if there are no implications where the right-hand side is \perp . We call such implications *purely negative*. It is not difficult to prove that every set of Horn formulae with no purely negative implications has a *minimal model*, that is, there exists a unique variable assignment that satisfies all the formulae and that only assigns \top to those

variables that are true in *every* satisfying variable assignment. We will use games to show that this minimal model can be computed in linear time. Then we can check for satisfiability of a given set Φ of Horn formulae by

- (1) removing all the purely negative implications from Φ ,
- (2) computing the minimal model, and
- (3) checking that every of the removed implications is true in this model.

As an example, let us consider the following set of Horn formulae.

$$\begin{array}{lll} \mathbf{1} \rightarrow A & A \wedge C \wedge F \rightarrow D & E \rightarrow G \\ A \wedge D \rightarrow B & B \wedge E \wedge G \rightarrow D & \mathbf{1} \rightarrow E \\ F \rightarrow C & G \rightarrow D & \end{array}$$

The minimal model assigns the value 1 to A, B, D, E , and G , and the value 0 to C and F .

The game corresponding to a set Φ of Horn formulae looks as follows. The positions for Player \diamond are of the form $\langle A \rangle$, where A is a variable appearing in Φ , the positions for Player \square are of the form $[\varphi]$ with $\varphi \in \Phi$. For each formula $A_0 \wedge \dots \wedge A_{n-1} \rightarrow B \in \Phi$, we have edges


$$\langle B \rangle \longrightarrow [A_0 \wedge \dots \wedge A_{n-1} \rightarrow B]$$

and $[A_0 \wedge \dots \wedge A_{n-1} \rightarrow B] \longrightarrow \langle A_i \rangle$, for $i < n$.

Intuitively, in the resulting game Player \diamond tries to prove that a variable A must have value 1 by choosing an implication $B_0 \wedge \dots \wedge B_{n-1} \rightarrow A$ that forces it to be true, while Player \square tries to prove that such an implication is not applicable by finding some condition B_i that is not met. With this intuition it is straightforward to show that our game has the desired properties.

Lemma 2.6. *Let Φ be a set of Horn formulae that does not contain any purely negative implications. A position of the form $\langle A \rangle$ belongs to the winning region of Player \diamond if, and only if, the variable A is true in the minimal model of Φ .*

The game corresponding to the above set of formulae is the following one, whose winning regions we have already computed above.




Modal logic


As a second application let us take a look at the model-checking problem for propositional modal logic. Given a transition system \mathfrak{S} with starting state s and a modal formula φ in *negation normal form*, we can construct a game that is won by Player \diamond if, and only if, $\mathfrak{S}, s \models \varphi$. The positions are of the form $\langle t, \psi \rangle$ where t is a state of \mathfrak{S} and ψ is a subformula of φ . Intuitively, in such a position Player \diamond tries to prove that $\mathfrak{S}, t \models \psi$, while Player \square tries to show that $\mathfrak{S}, t \not\models \psi$. The moves are as follows.

$$\begin{array}{ll} \langle t, \psi_0 \vee \psi_1 \rangle \rightarrow \langle t, \psi_i \rangle, & \text{for } i = 0, 1, \\ \langle t, \psi_0 \wedge \psi_1 \rangle \rightarrow \langle t, \psi_i \rangle, & \text{for } i = 0, 1, \\ \langle t, \diamond \vartheta \rangle \rightarrow \langle u, \vartheta \rangle, & \text{if } t \rightarrow u \text{ is an edge of } \mathfrak{S}, \\ \langle t, \square \vartheta \rangle \rightarrow \langle u, \vartheta \rangle, & \text{if } t \rightarrow u \text{ is an edge of } \mathfrak{S}. \end{array}$$

Finally, positions of the form $\langle t, \psi_0 \vee \psi_1 \rangle$ and $\langle t, \diamond \vartheta \rangle$ belong to Player \diamond ; those of the form $\langle t, \psi_0 \wedge \psi_1 \rangle$ and $\langle t, \square \vartheta \rangle$ belong to Player \square ; and a position of the form $\langle t, P \rangle$ belongs to Player \diamond if $t \notin P$; otherwise, it belongs to Player \square . Similarly, a position of the form $\langle t, \neg P \rangle$ belongs to Player \diamond if $t \in P$. For instance, for the transition system \mathfrak{S} and the formula φ on the left below, the resulting game (or at least its reachable part) is depicted on the right.



$$\varphi := \diamond(P \vee Q) \wedge \diamond\neg P$$



It is straightforward to check by induction on ψ that Player \diamond has a winning strategy from a position $\langle t, \psi \rangle$ if, and only if, $\mathfrak{G}, t \models \psi$.

3 Conway games

For many actual games like chess, go, etc., looking at a given position does not tell you which player's turn it is. To model such games it can be advantageous to assign the players not to the positions but to the moves. A *Conway game* is a graph

$$\mathfrak{G} = \langle V, E_{\diamond}, E_{\square} \rangle$$

with two edge relations $E_{\diamond}, E_{\square} \subseteq V \times V$, one for Player \diamond and one for Player \square . During a play of such a game starting in a given position and with a given starting player, the players alternately choose moves from their respective edge relation, that is, a play of such a game consists of a path through \mathfrak{G} where the edges alternate between E_{\diamond} -edges and E_{\square} -edges. As usual, we consider a player whose turn it is but who cannot make a move to lose the game. For simplicity, we will only consider Conway games where infinite plays are considered draws. Thus, to win a player has to manoeuvre his opponent into a position where he cannot make a move.

The question of determinacy for Conway games becomes more involved as, given a starting position, we must distinguish several more cases: it can be that one of the players has a winning strategy no matter who goes first, or it can be that always the player that make the first move wins, or it is the

player making the second move. We therefore say that a Conway game is *determined* if, for every position v , one of the following cases occurs.

- ◆ When taking the first turn Player \diamond has a winning strategy from v .
- ◆ When taking the first turn Player \square has a winning strategy from v .
- ◆ The player that moves first has a winning strategy from v .
- ◆ The player that moves second has a winning strategy from v .
- ◆ Both players have strategies from v that ensure at least a draw.

To compute the winning regions we proceed similar to positional games. We start by defining a step function. Because the game is alternating, it is convenient to use a function that takes *two* steps at a time. We set

$$\text{DStep}_{\sigma}(X) := \langle \sigma \rangle [\bar{\sigma}] X,$$

where, for $X \subseteq V$ and $\sigma \in \{\diamond, \square\}$,

$$\langle \sigma \rangle X := \{ v \in V \mid \langle v, w \rangle \in E_{\sigma} \text{ for some } w \in X \},$$

$$[\sigma] X := \{ v \in V \mid \langle v, w \rangle \in E_{\sigma} \text{ implies } w \in X \}.$$

Then the winning regions are

$$\mathcal{W}_{\sigma,1} := \bigcup_{\alpha} \text{DStep}_{\sigma}^{\alpha}(\emptyset),$$

$$\mathcal{W}_{\sigma,2} := [\bar{\sigma}] \mathcal{W}_{\sigma,1}.$$

Where $\mathcal{W}_{\sigma,1}$ denotes the set of vertices from which Player σ wins if he goes first and $\mathcal{W}_{\sigma,2}$ is the set of vertices from which he wins if the opponent goes first.

Lemma 3.1. *For every $v \in \mathcal{W}_{\sigma,1} \cup \mathcal{W}_{\sigma,2}$, Player σ has a positional winning strategy.*

Proof. It is sufficient to provide winning strategies for every $v \in \mathcal{W}_{\sigma,1}$ since, in a game starting with Player $\bar{\sigma}$ from a position $\mathcal{W}_{\sigma,2}$, the first move necessarily

leads to a position in $\mathcal{W}_{\sigma,1}$. (Or Player $\bar{\sigma}$ has no move and loses immediately.) Hence, suppose that $v \in \mathcal{W}_{\sigma,1}$ and let α be the least ordinal such that $v \in \text{DStep}_{\sigma}^{\alpha+1}(\emptyset)$. Then

$$v \in \langle \sigma \rangle [\bar{\sigma}] \text{DStep}_{\sigma}^{\alpha}(\emptyset)$$

implies that Player σ can make a move to some position

$$w \in [\bar{\sigma}] \text{DStep}_{\sigma}^{\alpha}(\emptyset),$$

and the response of the opponent leads to some position

$$v' \in \text{DStep}_{\sigma}^{\alpha}(\emptyset).$$

Consequently, we can construct the desired strategy by induction on α . \square

Theorem 3.2. *Conway games are positionally determined.*

Proof. Let $\mathcal{W} := \mathcal{W}_{\diamond,1} \cup \mathcal{W}_{\diamond,2} \cup \mathcal{W}_{\square,1} \cup \mathcal{W}_{\square,2}$. By Lemma 3.1, it is sufficient to prove that, for every position $v \in U := V \setminus \mathcal{W}$, both players have strategies that ensure a draw. Similar to the proof of Lemma 2.3, it is straightforward to check that, in every position $v \in V$ both players have the ability to choose an edge $\langle v, u \rangle$ with $u \in U$ and no player σ can choose an edge $\langle v, u \rangle$ with $u \in \mathcal{W}_{\sigma,2}$. \square

4 Gale-Steward games

Reachability games are very simple since we can ignore infinite plays. Let us now take a look at what happens for games where every infinite play is winning for one player or the other. *Gale-Steward games* are simple games of this kind. They are played on a tree with set of vertices $V := A^*$. Every position $v \in V$ has outgoing edges to all vertices of the form va with $a \in A$. The players strictly alternate, that is, positions $v \in (A^2)^*$ of even length belong to Player \diamond and those $v \in A(A^2)^*$ of odd length to Player \square . The winning condition is given by a set $W \subseteq A^\omega$, which contains the infinite plays won by Player \diamond . All other infinite plays are won by Player \square .

Proposition 4.1. *There exist Gale-Steward games that are not determined.*

Proof. We play on the complete binary tree with vertices $V := \{0,1\}^*$. As the game graph is acyclic, every strategy is automatically positional. Thus, a strategy for Player σ is a function $V_{\sigma} \rightarrow \{0,1\}$ where V_{σ} is either the set of all $v \in \{0,1\}^*$ of even length, or the set of all v of odd length. There are $\kappa := 2^{\aleph_0}$ such functions. We fix enumerations $(s_{\alpha})_{\alpha < \kappa}$ and $(t_{\alpha})_{\alpha < \kappa}$ of all strategies for, respectively, Player \diamond and Player \square . To construct a non-determined game we have to find two disjoint sets $W_{\diamond}, W_{\square} \subseteq \{0,1\}^\omega$ of infinite paths such that none of the s_{α} and t_{α} are winning strategies in the game with winning condition $W_{\diamond}, W_{\square}$.

We start with a bit of notation. Given a strategy s , we denote by $[s]$ the set of all infinite plays $p \in A^\omega$ that conform to s . With this notation we can say that a strategy s for Player σ is a winning if, and only if, $[s] \subseteq W_{\sigma}$.

By induction on $i < \kappa$, we construct two sequences $(\xi_i)_{i < \kappa}$ and $(\zeta_i)_{i < \kappa}$ of elements of A^ω as follows. Suppose that we have already defined ξ_i and ζ_i , for all $i < \alpha$. Then we pick some element $\xi_{\alpha} \in [s_{\alpha}]$ that is different from ξ_i and ζ_i , for all $i < \alpha$, and we pick some element $\zeta_{\alpha} \in [t_{\alpha}]$ that is different from ξ_i and ζ_i , for all $i \leq \alpha$, and also from the just chosen ξ_{α} . Note that we can do so since $[s_{\alpha}]$ has size κ , while the set $\{\xi_i \mid i < \alpha\} \cup \{\zeta_i \mid i < \alpha\}$ has size $|\alpha| < \kappa$. The same holds for $[t_{\alpha}]$.

We claim that the game with winning conditions

$$W_{\diamond} := \{\zeta_i \mid i < \kappa\} \quad \text{and} \quad W_{\square} := A^\omega \setminus W_{\diamond}$$

is not determined. For the proof, consider a strategy s for Player \diamond . Then $s = s_{\alpha}$, for some $\alpha < \kappa$. Since $\xi_{\alpha} \in [s_{\alpha}] \cap W_{\square}$, we have $[s] \not\subseteq W_{\diamond}$. Hence, s is not a winning strategy. In the same way it follows that no strategy t for Player \square is winning. \square

As Gale-Steward games can be non-determined in general, we have to put restrictions on the allowed winning conditions to get positive results. One handy way to do so is by equipping the set of all infinite plays with a topology. We call a set O of infinite plays *open* if there exists a set P of finite partial plays such that contains all infinite plays starting with some $p \in P$.

The complement of an open set is called *closed*. Note that the open sets are closed under arbitrary unions and finite intersections. Hence, they form a topology. A set is *Borel* if it is contained in the smallest class of sets that contains the open ones and that is closed under complement and countable unions.

Theorem 4.2 (Martin [8]). *If $W \subseteq A^\omega$ is Borel, then $\langle A^*, W \rangle$ is determined.*

The proof is a bit involved. Instead of proving the result in its full generality, we will only consider the much simpler case of open and closed winning conditions, i.e., where W_\diamond is open and W_\square closed, or vice versa. By symmetry, we may assume that the winning condition for Player \diamond is open. Then W_\diamond is determined by some set $P \subseteq A^*$ of prefixes and every play containing a position from P is winning. Thus, open winning condition corresponds to a reachability game: Player \diamond has to reach a position in P . As we have seen in Section 2, the winning region for Player \diamond in such a game is given by $\text{Attr}_\diamond(P)$. The complement of $\text{Attr}_\diamond(P)$ is a \diamond -trap. Since every infinite play inside this complement cannot contain positions from P , all such plays are winning for Player \square . Consequently, the game is determined with winning regions

$$\text{Attr}_\diamond(P) \quad \text{and} \quad A^* \setminus \text{Attr}_\diamond(P).$$

5 Regular games and parity games

In general, Gale-Steward games have no finite representation as the winning set W can be an arbitrary set of infinite sequences. In this section we take a look at a simple way to represent certain Gale-Steward games in a finite way: if the winning set W is a *regular* set of infinite plays, we can use an automaton to represent it. Or we can use ω -semigroups instead of automata. This leads to the following definition.

A *regular game* over an ω -semigroup $\mathfrak{S} = \langle S, S_\omega \rangle$ is a game


$$\mathfrak{G} = \langle V_\diamond, V_\square, E, \lambda, W \rangle$$

where $\lambda : E \rightarrow S$ is an edge-labelling and the winning set is given by a subset $W \subseteq S_\omega$. Player \diamond wins an infinite play p if the product of the corresponding edge labels evaluates to an element of W .

As an example consider the ω -semigroup $\mathfrak{S} = \langle S, S_\omega \rangle$ where $S = \wp\{0, 1\}$ and $S_\omega = \wp\{0, 1\}$. We define the product by

$$\begin{aligned} a \cdot b &:= a \cup b, & \text{for } a, b \in S, \\ a \cdot u &:= u, & \text{for } a \in S \text{ and } u \in S_\omega, \\ \pi(a_0, a_1, \dots) &:= \bigcap_{i < \omega} \bigcup_{i \leq k < \omega} a_i, & \text{for } a_0, a_1, \dots \in S. \end{aligned}$$

In the game



with winning set $W = \{\{0, 1\}\}$ Player \diamond has a winning strategy by alternating between the two edges. But note that he does not have a positional winning strategy as using only one of the edges will result in a loss.

Computing the winning regions of a regular game is more complicated than for open games. The difference is that, instead of reaching a certain set once, we have to be able to reach it over and over again. To simplify our task, let us start by considering a special case of regular games of the following form. A *parity game* is a game of the form

$$\mathfrak{G} = \langle V_\diamond, V_\square, E, \Omega \rangle$$

where $\Omega : V \rightarrow D$ for some finite set $D \subseteq \omega$. We call Ω the *priority function* and $\Omega(v)$ the *priority* of the position v . Player \diamond wins an infinite play $p = (v_i)_{i < \omega}$ if it satisfies the *parity condition*:


$$\liminf_{i < \omega} \Omega(v_i) \text{ is even.}$$

For instance, in the following parity game Player \diamond wins from every position except for the one in the lower right. (The numbers denote the priorities.)

$W_\sigma = V \setminus \text{Attr}_{\bar{\sigma}}(W_\sigma)$ is a $\bar{\sigma}$ -trap. For a set $X \subseteq V$, we denote by $\mathfrak{G}[X]$ the subgame of \mathfrak{G} consisting of all positions in X . Define

$$K := W_\sigma \cap \Omega^{-1}(k) \quad \text{and} \quad U := W_\sigma \setminus \text{Attr}_\sigma(K/W_\sigma),$$

where $\text{Attr}_\sigma(K/W_\sigma)$ is the σ -attractor of K computed in $\mathfrak{G}[W_\sigma]$. Since no position in U uses the priority k , we can apply the inductive hypothesis to the game $\mathfrak{G}[U]$ and obtain a partition $U = U_\sigma \cup U_{\bar{\sigma}}$ of U into a σ -domain U_σ and a $\bar{\sigma}$ -domain $U_{\bar{\sigma}}$.



We will show next that $W_{\bar{\sigma}} \cup U_{\bar{\sigma}}$ is a $\bar{\sigma}$ -domain. By definition of $W_{\bar{\sigma}}$, it then follows that $U_{\bar{\sigma}} = \emptyset$. To see that $W_{\bar{\sigma}} \cup U_{\bar{\sigma}}$ is a σ -trap, we distinguish four cases.

- (i) Let $v \in V_{\bar{\sigma}} \cap W_{\bar{\sigma}}$. As $W_{\bar{\sigma}}$ is a $\bar{\sigma}$ -domain, we can find a successor u of v that belongs to $W_{\bar{\sigma}} \subseteq W_{\bar{\sigma}} \cup U_{\bar{\sigma}}$.
- (ii) Let $v \in V_\sigma \cap W_{\bar{\sigma}}$. As $W_{\bar{\sigma}}$ is a $\bar{\sigma}$ -domain, every successor u of v belongs to $W_{\bar{\sigma}} \subseteq W_{\bar{\sigma}} \cup U_{\bar{\sigma}}$.
- (iii) Let $v \in V_{\bar{\sigma}} \cap U_{\bar{\sigma}}$. As $U_{\bar{\sigma}}$ is a $\bar{\sigma}$ -domain in $\mathfrak{G}[U]$, we can find a successor u of v that belongs to $U_{\bar{\sigma}} \subseteq W_{\bar{\sigma}} \cup U_{\bar{\sigma}}$.
- (iv) Let $v \in V_\sigma \cap U_{\bar{\sigma}}$. As $U_{\bar{\sigma}}$ is a $\bar{\sigma}$ -domain in $\mathfrak{G}[U]$, every successor u of v either belongs to $U_{\bar{\sigma}}$ or to $V \setminus U$. Since U is a σ -trap in the subgame $\mathfrak{G}[W_\sigma]$, the latter is only possible if $u \notin W_\sigma$, i.e., $u \in W_{\bar{\sigma}}$. Consequently, all successors belong to $W_{\bar{\sigma}} \cup U_{\bar{\sigma}}$.

It remains to find a positional strategy for Player $\bar{\sigma}$ on $W_{\bar{\sigma}} \cup U_{\bar{\sigma}}$. As $W_{\bar{\sigma}}$ and $U_{\bar{\sigma}}$ are $\bar{\sigma}$ -domains in, respectively, \mathfrak{G} and $\mathfrak{G}[U]$, there are positional

strategies t_W and t_U for Player $\bar{\sigma}$ on these to sets. We define a strategy s by

$$s(v) := \begin{cases} t_U(v) & \text{if } v \in U, \\ t_W(v) & \text{otherwise.} \end{cases}$$

To show that s is winning, consider a play p conforming to s and starting in some position in $W_{\bar{\sigma}} \cup U_{\bar{\sigma}}$. If p enters $W_{\bar{\sigma}}$, it never leaves this set. Consequently, the rest of p conforms to t_W and is therefore winning. Otherwise, the play stays the whole time in $U_{\bar{\sigma}}$ and conforms to t_U . Thus, it is also winning.

We have shown that $U = U_\sigma$ is a σ -domain in $\mathfrak{G}[U]$. Let t be the corresponding strategy. To conclude the proof it is sufficient to show that W_σ is a σ -domain in \mathfrak{G} . We have already seen above that it is a $\bar{\sigma}$ -trap. Hence, it remains to construct a positional strategy for Player σ on W_σ . Note that $W_\sigma = U_\sigma \cup \text{Attr}_\sigma(K/W_\sigma)$. For positions $v \in U$, we use the strategy t for $\mathfrak{G}[U]$. For positions $v \in K$, we choose an arbitrary successor in W_σ . For the remaining positions, we use the attractor strategy that ensures that we visit K . Let s be the resulting strategy. To show that it is winning, consider a play p conforming to s and starting in some position in W_σ . If p enters U_σ it will stay in this set and the remainder of the play conforms to t . Hence, p is winning. Otherwise, p is entirely contained in $\text{Attr}_\sigma(K/W_\sigma)$. This implies that p either ends in a terminal vertex belonging to Player $\bar{\sigma}$, or it visits the set K infinitely often. In both cases Player σ wins. \square


Finite-memory strategies

One can show that all regular games are Borel. Hence, determinacy follows from the Theorem of Martin. But we can prove a stronger statement: regular games admit what is called *finite-memory* strategies. A *finite-memory strategy* for Player σ is given by a finite set M (the *memory*) and two functions $s : M \times V_\sigma \rightarrow E$ and $\alpha : M \times E \rightarrow M$ such that, given a state $m \in M$ and a vertex $v \in V_\sigma$, $s(m, v)$ returns the outgoing edge e to be chosen by Player σ after which $\alpha(m, e)$ will be the new memory state. (If there is no outgoing edge, we let s remain undefined.) Formally, we say that a play $p = (e_i)_i$

(which, for regular games where the edge labelling matters, we consider as a sequence of edges) *conforms* to such a strategy if there exists a sequence $(m_i)_i$ of memory states such that, for every step i ,

- ◆ if $e_i = \langle v_i, v_{i+1} \rangle$ with $v_i \in V_\sigma$, then $e_i = s(m_i, v_i)$, and
- ◆ $m_{i+1} := \alpha(m_i, e_i)$.

In the game



from above, Player \diamond has a finite-memory strategy with $M = \{0, 1\}$. If $m = 0$, he takes the left edge and sets the memory state to 1. Otherwise, he takes the right edge and sets the state to 0. This results in alternately taking the two edges, which is winning for him.

We can reformulate the definition of a finite-memory strategy as follows. For a regular game $\mathfrak{G} = \langle V_\diamond, V_\square, E, \lambda, W \rangle$ and a function $\alpha : M \times E \rightarrow M$, we define the *product game*

$$\mathfrak{G} \times_\alpha M := \langle V'_\diamond, V'_\square, E', \lambda', W' \rangle$$

with positions


$$V'_\diamond := V_\diamond \times M \quad \text{and} \quad V'_\square := V_\square \times M,$$

edge relations

$$E' := \{ \langle \langle u, m \rangle, \langle v, n \rangle \rangle \mid \langle u, v \rangle \in E \text{ and } n = \alpha(u, \langle u, v \rangle) \},$$

edge labelling $\lambda'(\langle \langle u, m \rangle, \langle v, n \rangle \rangle) := \lambda(\langle u, v \rangle)$, and the same winning condition $W' := W$. Then a strategy is finite-memory for \mathfrak{G} if, and only if, there exists a finite set M and a function $\alpha : M \times E \rightarrow M$ such that s is a positional strategy in the game $\mathfrak{G} \times_\alpha M$.

In the above example, the product $\mathfrak{G} \times_\alpha M$ is the game



which clearly has a positional winning strategy.

Note that this operation of equipping a game with memory does not change the game much. There exist one-to-one correspondences between

- ◆ plays of \mathfrak{G} and of $\mathfrak{G} \times_\alpha M$;
- ◆ strategies of \mathfrak{G} and of $\mathfrak{G} \times_\alpha M$;
- ◆ winning strategies of \mathfrak{G} and of $\mathfrak{G} \times_\alpha M$.

The only difference between these two games is the amount of memory a strategy needs. In particular, some positional strategies of $\mathfrak{G} \times_\alpha M$ might correspond to strategies of \mathfrak{G} which are not positional. This is exactly what we need to prove the following result.

Theorem 5.3 (Büchi, Landweber). *In every regular game both players have a finite-memory winning strategy on their respective winning regions.*

Proof. Let $\mathfrak{G} = \langle V_\diamond, V_\square, E, \lambda, W \rangle$ be a regular game. We fix a deterministic parity automaton $\mathcal{A} = \langle Q, S, \delta, q_0, \Omega \rangle$ recognising the set W of winning plays and construct the product game $\mathfrak{G} \times_\delta \mathcal{Q}$. In this game a play p is winning for Player \diamond if, and only if, its projection to the second component produces an accepting run of \mathcal{A} . Consequently, we can turn $\mathfrak{G} \times_\delta \mathcal{Q}$ into a parity game by using as priority function the function Ω from \mathcal{A} applied to the second component. Since parity games are positionally determined, we obtain two positional winning strategies s_\diamond and s_\square for the two players in their respective winning regions. As we have seen in the remark before the theorem, these two strategies induce finite-memory strategies in the original game. \square

Positionally determined games

We have seen that parity games admit positional strategies while arbitrary regular games only admit finite-memory ones. One might wonder whether

there exists a larger class of games with positional strategies. It turns out that is not the case. We will prove below that (nearly) every regular, positionally determined game is equivalent to a parity game – with one notable caveat: we will only be able to establish this statement for games with a winning condition of the following form. A winning condition $W \subseteq S_\omega$ is called *prefix-invariant* if

$$w \in W \Leftrightarrow aw \in W, \quad \text{for all } a \in S \text{ and } w \in S_\omega.$$

Note that this is not much of a restriction since most of the common winning conditions used in game theory or automata theory are of this form.

To analyse such conditions we start with an observation from semigroup theory. In a *finite* ω -semigroup $\mathfrak{S} = \langle S, S_\omega \rangle$ every infinite product $a_0 a_1 a_2 \dots$ has a factorisation of the form be^ω . If W is prefix-invariant, we have $be^\omega \in W \Leftrightarrow e^\omega \in W$. Thus, the set W is completely determined by the powers e^ω it contains.

Given a regular game $\mathfrak{G} = \langle V_\diamond, V_\square, E, \lambda, W \rangle$ over an ω -semigroup \mathfrak{S} , let us write $W_\diamond := W$ and $W_\square := S_\omega \setminus W$. The *winning condition* of \mathfrak{G} is the pair $\langle \mathfrak{S}, W \rangle$. We call the set

$$P_\sigma := \{ e \in S \mid e^\omega \in W_\sigma \}$$

the *period set* for Player σ . Finally, let us say that $\langle \mathfrak{S}, W \rangle$ is *equivalent* to a parity condition if there exists a function $\Omega : S \rightarrow \omega$ such that

$$\pi(a_0, a_1, \dots) \in W \quad \text{iff} \quad \liminf_i \Omega(a_i) \text{ is even.}$$

In this case, we can turn every game \mathfrak{G} with winning condition $\langle \mathfrak{S}, W \rangle$ into a parity game as follows. We first replace all edge labels by their image under Ω . In this way we obtain a kind of parity game where the priorities are attached to the edges instead of the vertices. We can turn the resulting game into an ordinary parity game by adding intermediate vertices to the edges where we can put the priorities.

Theorem 5.4 (Colcombet, Niwiński). *Let \mathfrak{S} be an ω -semigroup (not necessarily finite) and $W \subseteq S_\omega$ prefix-invariant. If all games with winning condition*

$\langle \mathfrak{S}, W \rangle$ *are positionally determined, then $\langle \mathfrak{S}, W \rangle$ is equivalent to a parity condition.*

We split the proof of this theorem into two lemmas. We start with listing some basic properties of the period sets.

Lemma 5.5. *Let \mathfrak{S} be an ω -semigroup and $W \subseteq S_\omega$ prefix-invariant. If all games with winning condition $\langle \mathfrak{S}, W \rangle$ are positionally determined, then the following condition holds:*

- (a) $a, b \in P_\sigma$ implies $ab \in P_\sigma$.
- (b) $ab \in P_\sigma$ implies $ba \in P_\sigma$.
- (c) Every element $w \in S_\omega$ that can be written as an infinite product of elements of P_σ belongs to W_σ .
- (d) For all $A, B \subseteq S$,

$$(\exists a \in A)(\forall b \in B)[ab \in P_\sigma] \Leftrightarrow (\forall b \in B)(\exists a \in A)[ab \in P_\sigma].$$


- (e) For every $a \in S$, there exists some $n > 0$ such that

$$ab_0, \dots, ab_k \in P_\sigma \Rightarrow a^n b_0 \dots b_k \in P_\sigma, \quad \text{for all } b_0, \dots, b_k \in S.$$

- (f) For $a \in P_\sigma$ and $B \subseteq S$,

$$aB \subseteq P_\sigma \quad \text{implies} \quad aB^* \subseteq P_\sigma.$$

Proof. (c) Suppose that $w = \pi(a_0, a_1, \dots) \notin W_\sigma$. We have to show that there is some index k with $a_k^\omega \notin W_\sigma$. Consider the game \mathfrak{G} with a single position v belonging to Player $\bar{\sigma}$ and one a_i -labelled edge $v \rightarrow v$, for every $i < \omega$.




Since $w \notin W$, Player $\bar{\sigma}$ has a winning strategy in \mathfrak{G} by choosing in turn i the edge with label a_i . By assumption, he also has a positional winning strategy s . Let a_k be the label of the edge chosen by s . As the resulting play is winning, it follows that $a_k^\omega \notin W_\sigma$. Hence, $a_k \notin P_\sigma$.

(a) If $a, b \in P_\sigma$, then $(ab)^\omega \in W_\sigma$ by (c).

(b) Let $ab \in P_\sigma$. Then $a(ba)^\omega = (ab)^\omega \in W_\sigma$ implies, by prefix-invariance that $(ba)^\omega \in W_\sigma$.

(d) (\Rightarrow) is trivial. For (\Leftarrow), consider the game \mathfrak{G} with positions $V_\sigma = \{u\}$ and $V_{\bar{\sigma}} = \{v\}$. For every $a \in A$, we add an a -labelled edge $u \rightarrow v$, and for every $b \in B$, a b -labelled edge $v \rightarrow u$.



Then Player σ can win this game by playing as follows. Every time Player $\bar{\sigma}$ chooses a b -labelled edge, Player σ responds with an a -labelled edge for some $ab \in P$. By assumption, Player σ also has a positional winning strategy s . Let a be the label of the edge chosen by s . For every $b \in B$, there exists a play with labelling $(ab)^\omega$ conforming to s . Since these plays must be winning, it follows that $(ab)^\omega \in W_\sigma$, i.e., $ab \in P_\sigma$, for all b .

(e) Fix $a \in S$ and set $B_\circ := \{b \in S \mid ab \in P_\sigma\}$. Applying (d) to the sets

$$A := \{a^n \mid n > \circ\} \quad \text{and} \quad B := \{b_\circ \cdots b_k \mid b_\circ, \dots, b_k \in B_\circ\},$$

we see that it is sufficient to prove that

$$a^{k+1}b_\circ \cdots b_k \in P_\sigma, \quad \text{for } b_\circ, \dots, b_k \in B_\circ.$$

We do so by induction on k . If $k = \circ$, then $b_\circ \in B_\circ$ implies that $a^1b_\circ \in P_\sigma$. Hence, suppose that $k > \circ$. By inductive hypothesis and the fact that $b_k \in B_\circ$, we have $a^k b_\circ \cdots b_{k-1} \in P_\sigma$ and $ab_k \in P_\sigma$. Hence, $(ab_k)^\omega \in W_\sigma$. Since $a(b_k a)^\omega = (ab_k)^\omega \in W_\sigma$, prefix-invariance implies that $(b_k a)^\omega \in W_\sigma$, i.e., $b_k a \in P_\sigma$. Consequently, it follows by (c) that $(a^k b_\circ \cdots b_{k-1} b_k a)^\omega \in W_\sigma$ and

we can again use prefix-invariance to show that

$$(a^{k+1}b_\circ \cdots b_k)^\omega = a(a^k b_\circ \cdots b_{k-1} b_k a)^\omega \in W_\sigma.$$

Thus, $a^{k+1}b_\circ \cdots b_k \in P_\sigma$.

(f) Fix $a \in P_\sigma$ and $B \subseteq S$ with $aB \subseteq P_\sigma$. By (e) the set $N := \{n \geq 1 \mid a^n B^+ \subseteq W_\diamond\}$ is non-empty. Note that, if $n \in N$ then $a^n u \in P_\sigma$, for all $u \in B^+$. By (a) this implies that $aa^n u \in P_\sigma$. Consequently, $n+1 \in N$. Thus, $N = \{k, k+1, k+2, \dots\}$ for some $k < \omega$.

We claim that $k = 1$. Then $aB^+ \subseteq P_\sigma$. Since also $a \in P_\sigma$, it follows that $aB^* \subseteq P_\sigma$, as desired. Hence, it remains to prove the claim.

For a contradiction, suppose that $k > 1$ and set $m := k-1$. Then $m \notin N$, but $2m \in N$. Hence, there is some $u \in B^+$ with $a^m u \in P_{\bar{\sigma}}$. By (a) and (b), it follows that $ua^m \in P_{\bar{\sigma}}$, $a^m uua^m \in P_{\bar{\sigma}}$, and $a^{2m} uu \in P_{\bar{\sigma}}$. Hence, $2m \notin N$. A contradiction. \square

The next lemma now concludes the proof of Theorem 5.4.

Lemma 5.6. *Let \mathfrak{S} be an ω -semigroup and $W \subseteq S_\omega$ a prefix-invariant set such that all games with winning condition (\mathfrak{S}, W) are positionally determined. There exists a function $\Omega : S \rightarrow [2n]$ such that*

- (1) Ω maps P_\diamond to even numbers and P_\square to odd ones,
- (2) $\Omega(a) \leq \Omega(b)$ implies $\Omega(ab) \equiv \Omega(a) \pmod{2}$.
- (3) $\pi(a_\circ, a_1, \dots) \in W_\diamond$ iff $\liminf_i \Omega(a_i)$ is even.

Proof. Consider the relation $\sqsubseteq \subseteq P_\square \times P_\square$ defined by

$$a \sqsubseteq b \quad \text{iff} \quad ac \in P_\square \Rightarrow bc \in P_\square, \quad \text{for all } c \in S.$$

We start by proving that it is a linear preorder of finite index.

Reflexivity and transitivity of \sqsubseteq follows immediately from the definition. For linearity, suppose that a and b are non-comparable. Then there are elements c and d such that

$$ac \in P_\square, \quad bc \notin P_\square, \quad ad \notin P_\square, \quad bd \in P_\square.$$

By Lemma 5.5 (a) and (b), it follows that

$$acbd \in P_{\square}, \quad da \in P_{\diamond}, \quad cb \in P_{\diamond}, \quad dacb \in P_{\diamond}, \quad acbd \in P_{\diamond}.$$

A contradiction.

It remains to prove that \sqsubseteq has finite index. For a contradiction, suppose otherwise. We distinguish two cases. If there exists an infinite strictly increasing chain $a_0 \sqsubset a_1 \sqsubset a_2 \sqsubset \dots$, we can fix elements $c_i \in S$ with $a_i c_i \notin P_{\square}$ and $a_{i+1} c_i \in P_{\square}$. Then Lemma 5.5 (a) and (c) implies that $c_i a_{i+1} \in P_{\square}$ and

$$c_0 a_1 c_1 a_2 c_2 a_3 \dots \in W_{\square}.$$

But by the same argument as above, $a_i c_i \in P_{\diamond}$ implies that

$$a_0 c_0 a_1 c_1 a_2 c_2 a_3 \dots \in W_{\diamond}.$$

A contradiction to prefix-invariance.

Similarly, if there exists an infinite strictly decreasing chain $a_0 \supset a_1 \supset a_2 \supset \dots$, we can fix elements $c_i \in S$ with $a_i c_i \in P_{\square}$ and $a_{i+1} c_i \notin P_{\square}$. In the same way as above it follows that

$$c_0 a_1 c_1 a_2 c_2 a_3 \dots \in W_{\diamond} \quad \text{and} \quad a_0 c_0 a_1 c_1 a_2 c_2 a_3 \dots \in W_{\square}.$$

Again a contradiction.

To conclude the proof, let $B_0 \supset \dots \supset B_{m-1}$ be a decreasing enumeration of all \sqsubseteq -classes and set

$$A_i := \{c \in P_{\diamond} \mid ac \in P_{\square} \text{ for some/all } a \in B_i\}, \quad \text{for } 0 \leq i < m.$$

In addition, we set $A_{-1} := P_{\diamond}$ and $A_m := \emptyset$. Note that, by definition of \sqsubseteq , we have $P_{\diamond} = A_{-1} \supseteq A_0 \supseteq \dots \supseteq A_{m-1} = A_m = \emptyset$. We claim that the function $\Omega : S \rightarrow [2m+1]$ defined by

$$\Omega(a) := \begin{cases} 2k & \text{if } a \in A_{k-1} \setminus A_k, \\ 2k+1 & \text{if } a \in B_k, \end{cases}$$

has the desired properties.

(1) Clearly, Ω maps each $A_k \subseteq P_{\diamond}$ to an even number and each $B_k \subseteq P_{\square}$ to an odd one.

(2) Suppose that $\Omega(a) \leq \Omega(b)$. We distinguish four cases. If both $\Omega(a)$ and $\Omega(b)$ are even, then $a, b \in P_{\diamond}$, which implies by Lemma 5.5 (a) that $ab \in P_{\diamond}$. Hence, $\Omega(ab)$ is also even. In the same way it follows that, if $\Omega(a)$ and $\Omega(b)$ are odd, then so is $\Omega(ab)$.

Suppose that $\Omega(a) = 2k+1$ and $\Omega(b) = 2i$. Then $k < i$, $a \in B_k$ and $b \in A_{i-1} \setminus A_i \subseteq A_k$. By definition of A_k it follows that $ab \in P_{\square}$. Hence, $\Omega(ab)$ is odd.

Finally, suppose that $\Omega(a) = 2k$ and $\Omega(b) = 2i+1$. Then $k \leq i$, $a \in A_{k-1} \setminus A_k$ and $b \in B_i \subseteq B_k$. By definition of A_k and the fact that $a \notin A_k$, it follows that $ba \in P_{\diamond}$. Hence, Lemma 5.5 (a) implies that $ab \in P_{\diamond}$ and $\Omega(ab)$ is even.

(3) Fix $a_0, a_1, \dots \in S$ and set $k := \liminf_i \Omega(a_i)$. Since W is prefix-invariant, we may assume w.l.o.g. that there is no i with $\Omega(a_i) < k$. Set

$$B := \Omega^{-1}(k) \quad \text{and} \quad C := \Omega^{-1}[\{k+1, \dots, 2m\}].$$

Then we can factorise the sequence $(a_i)_i$ into words $u_0, u_1, \dots \in BC^*$. Let c_i be the product of u_i . If k is even, it follows by (b) that $bc \in P_{\diamond}$, for all $b \in B$ and $c \in C$. Consequently, we can use Lemma 5.5 (f) and (c) to show that $c_i \in P_{\diamond}$ and $\pi(a_0, a_1, \dots) = \pi(c_0, c_1, \dots) \in W_{\diamond}$. If k is odd, it follows in the same way that $\pi(a_0, a_1, \dots) \in W_{\square}$. \square

Solving parity games

To solve reachability games, we introduced the notion of a rank which, intuitively counts how far away from the goal we are. For parity games the situation is more complicated since we have to reach the goal not only once but repeatedly. It is possible to define ranks also for parity games if, instead of a single ordinal, one uses tuples of them, one for each priority k that counts how far away we are from a position of that priority. As this turns out to be a bit technical and not very enlightening, we will not do so.

Instead, we will present an algorithm for computing the winning regions that is similar to the construction in the Theorem of Büchi and Landweber. We will prove that, for every parity game \mathfrak{G} , there exists an action $\alpha : M \times E \rightarrow M$ that turns the product game $\mathfrak{G} \times_\alpha M$ into a reachability game. Then we can use the linear time algorithm from Section 2 to compute the winning regions. The memory M we will construct below has size $n^{\mathcal{O}(\log d)}$. As we can solve reachability games in linear time, we therefore obtain the following complexity bound.

Theorem 5.7. *The winning regions of a parity game \mathfrak{G} with n positions and d priorities can be computed in time $n^{\mathcal{O}(\log d)}$.*

The precise complexity of computing the winning regions of a parity game are still unknown. One can show that the problem belongs to the complexity class $U \cap \text{co-}U$, which means that it is probably not NP-complete. It might even belong to P, but no one has found a polynomial time algorithm so far.

Let us present the algorithm the above theorem is based on. Consider a parity game \mathfrak{G} and let p be a play of $\mathfrak{G} \times_\alpha M$. We say that p contains an *even cycle* if $p = xyz$ where y is a non-empty path starting and ending at the same position of \mathfrak{G} (the memory contents may differ) and such that the least priority seen along y is even. Note that, whether or not a given play p contains an even cycle is a reachability property: once we have found the end of the cycle, we do not need to look at the rest of p . We will design our memory M in such a way that detection of such cycles becomes easy. But first, let us show that the existence of even cycles is equivalent to winning.

Lemma 5.8. *Let \mathfrak{G} be a finite parity game and $\alpha : M \times E \rightarrow M$ an action. The following statements are equivalent.*

- (1) Player \diamond has a winning strategy s in \mathfrak{G} .
- (2) Player \diamond has a strategy s' for $\mathfrak{G} \times_\alpha M$ such that all cycles in every play conforming to s' are even.
- (3) Player \diamond has a strategy s' for $\mathfrak{G} \times_\alpha M$ such that every play conforming to s' has an even cycle.

Proof. (2) \Rightarrow (3) is trivial.

(3) \Rightarrow (1) Suppose that Player \diamond does not have a winning strategy for \mathfrak{G} . By determinacy, it then follows that Player \square has a positional winning strategy s in that game. This strategy induces a strategy s' for Player \square in the game $\mathfrak{G} \times_\alpha M$. Let p' be a play conforming to s' and let p be the corresponding play of \mathfrak{G} . Then p conforms to s . If p' contained an even cycle then, s being positional, it would follow that p contained infinite repetitions of this cycle. In particular, the least priority seen infinitely often in p would be even and p would be winning for Player \diamond . A contradiction to our choice of s .

(1) \Rightarrow (2) Let s be a winning strategy for Player \diamond in \mathfrak{G} . W.l.o.g. we may assume that s is positional. Let s' be the strategy in $\mathfrak{G} \times_\alpha M$ induced by s . To show that it has the desired property, consider a play p' conforming to s' and let p be the corresponding play in \mathfrak{G} . Then p conforms to s and is, therefore, winning. Since s is positional, p consists of a path leading to a cycle which is repeated infinitely often. Let k be the minimal priority along this cycle. As p is winning, it follows that k is even. Hence, so is (every copy of) the cycle. \square

How can we detect an even cycle? The easiest way would be to store all the positions of \mathfrak{G} we have already seen. Once we see one of them for the second time, we have found a cycle. Unfortunately, storing that many positions requires too much memory. So instead, we resort to a counting trick.

Let \mathfrak{G} be a parity game with n positions and priorities $\{0, \dots, d-1\}$. We use the ω -semigroup \mathfrak{S} with domains $S := [d]$ and $S_\omega := [d]$ and product


$$\begin{aligned} k \cdot k' &:= \min\{k, k'\}, & \text{for } k, k' \in S, \\ k \cdot l &:= l, & \text{for } k \in S \text{ and } l \in S_\omega, \\ \pi(k_0, k_1, k_2, \dots) &:= \liminf_{i < \omega} k_i, & \text{for } k_i \in S, \end{aligned}$$

and we label an edge $u \rightarrow v$ of \mathfrak{G} by the semigroup element $\Omega(u) \in S$.

Given a finite word $w = k_0 \cdots k_{n-1} \in [d]^*$, we call a sequence $z_0 < \dots < z_{m-1} < n$ an *even factorisation* of w if

- k_{z_i} is even, for all $i < m$,
- $k_j \geq \min \{k_{z_i}, k_{z_{i+1}}\}$, for all $z_i \leq j \leq z_{i+1}$, $i < m-1$,
- $k_j \geq k_{z_0}$, for $j \leq z_0$,
- $k_j \geq k_{z_{m-1}}$, for $j \geq z_{m-1}$.

We call m the *length* of the factorisation and the number $\min_i k_{z_i}$ its *value*.



Note that, if w has an even factorisation of length m and value k and w' has one of length m' and value k' , the ww' has an even factorisation of length $m + m'$ and value $\min \{k, k'\}$.

For each $l < \omega$, we will define a deterministic automaton \mathcal{A}_l over the alphabet $[d]$ that computes the length and value of an even factorisation of (some suffix of) its input. The precise definition is as follows. The set of states M_l consists of a special *accepting state* $*$ plus all triples

$$\langle s, \bar{k}, \bar{n} \rangle \in [l+1] \times [d]^* \times [l]^*,$$

where $s \leq l$ is the *size* of the state, \bar{k} is a non-decreasing sequence of *priorities* of length s , and \bar{n} a strictly decreasing sequence of *counters*, also of length s . To compute the cardinality of M_l , note that we can encode each state $\langle s, \bar{k}, \bar{n} \rangle$ as a word $c_0 \dots c_{l-1} \in ([d] + \square)^l$ where

$$c_i := \begin{cases} k_j & \text{if } n_j = i, \\ \square & \text{otherwise.} \end{cases} \quad \langle 3, \langle 0, 3, 4 \rangle, \langle 5, 2, 1 \rangle \rangle$$

4	3			0
---	---	--	--	---

Below we will choose l such that $2^{l-1} \leq n < 2^l$. Then it follows

$$\begin{aligned} |M_l| &\leq (d+1)^l + 1 \leq (d+1)^{\log(n+1)} + 1 \\ &= (n+1)^{\log(d+1)} + 1 \in n^{\mathcal{O}(\log d)}, \end{aligned}$$

which is the right size for the theorem.

Before defining the transition relation of \mathcal{A}_l , let us state the intended behaviour of the automaton.

Lemma 5.9. *Suppose that after having read a word $w \in [d]^*$ the automaton \mathcal{A}_l enters the state $\langle s, \bar{k}, \bar{n} \rangle$. Then*

$$w = w' w_0 \dots w_{s-1},$$

where each w_i has an even factorisation of length at least 2^{n_i} with value k_i . In particular, the word w has a suffix with an even factorisation of length at least

$$\sum_{i < s} 2^{n_i}.$$

Before giving the proof, we have to finish the definition of \mathcal{A}_l . The initial state is the pair $\langle 0, \langle \rangle, \langle \rangle \rangle$ consisting of two empty sequences. In a state $\langle s, \bar{k}, \bar{n} \rangle$ when reading the letter c , the automaton can enter the state $\langle s', \bar{k}', \bar{n}' \rangle$ if one of the following three conditions is met.

(i) $s = 0$ or $k_{s-1} \leq c$,

$$s' = s, \quad \bar{k}' = \bar{k}, \quad \text{and} \quad \bar{n}' = \bar{n}.$$

(ii) There is some $0 < i < s$ such that $k_{i-1} < c < k_i$,

$$s' = i+1, \quad \bar{k}' = \langle k_0, \dots, k_{i-1}, c \rangle, \quad \bar{n}' = \langle n_0, \dots, n_{i-1}, n_i \rangle.$$

(iii) There is some $0 < i < s$ such that $k_{i-1} \leq c$, the priorities k_i, \dots, k_{s-1} are even, $\langle n_i, \dots, n_{s-1} \rangle = \langle s-i-1, s-i-2, \dots, 1, 0 \rangle$,

$$s' = i+1, \quad \bar{k}' = \langle k_0, \dots, k_{i-1}, c \rangle, \quad \bar{n}' = \langle n_0, \dots, n_{i-1}, n_i+1 \rangle.$$

(iv) There is some $o < i < s$ such that $k_{i-1} \leq c$, the priority k_i is odd, while k_{i+1}, \dots, k_{s-1} are even,

$$\langle n_i, \dots, n_{s-1} \rangle = \langle s-i-1, s-i-2, \dots, 1, o \rangle,$$

$$s' = i+1, \quad \bar{k}' = \langle k_o, \dots, k_{i-1}, c \rangle, \quad \bar{n}' = \langle n_o, \dots, n_{i-1}, n_i+1 \rangle.$$

(v) $s = l$, each priority k_i is even, c is even, and the next state is $*$.

(vi) Once \mathcal{A}_l has reached the state $*$, it remains there.

If there are several possible transitions, we choose the one that leads to a state of minimal length.

Proof of Lemma 5.9. We prove the claim by induction on the length of w . For $w = \langle \rangle$, \mathcal{A}_l is in the initial state $\langle o, \langle \rangle, \langle \rangle \rangle$ and the claim is trivial.

For the inductive step, suppose that the input is wc with $w \in [d]^*$ and $c \in [d]$, and let $\langle s, \bar{k}, \bar{c} \rangle$ be the state after reading w . By inductive hypothesis, w has a suffix of the form $w_o \cdots w_{s-1}$ where each w_i has an even factorisation of length at least 2^{n_i} with value k_i . We distinguish several cases, depending on which transition the automaton takes while reading the last letter c .

If the last transition is of the form (i), we obtain the desired suffix $w'_o \cdots w'_{s-1}$ of wc by setting $w'_i := w_i$, for $i < s-1$, and $w'_{s-1} := w_{s-1}c$.

If the last transition is of the form (ii), let i be the index such that $k_{i-1} < c < k_i$. We set $w'_j := w_j$, for $j < i$, and $w'_i := w_i \cdots w_{s-1}c$. Then w'_i has an even factorisation of length at least $2^{n_i} + \cdots + 2^{n_{s-1}} + 1 \geq 2^{n_i}$.

If the last transition is of the form (iii) or (iv), let i be the index from the above definition. We set $w'_j := w_j$, for $j < i$, and $w'_i := w_i \cdots w_{m-1}c$. Then w'_i has an even factorisation of length at least

$$2^{n_i} + \cdots + 2^{n_{s-1}} + 1 = 2^{s-i-1} + 2^{s-i-2} + \cdots + 2^1 + 2^0 + 1 = 2^{s-i}. \quad \square$$

With the help of Lemma 5.9, we are able to show that the information contained in the states of \mathcal{A}_l is sufficient to detect whether the input contains an even cycle.

Lemma 5.10. *Let p be an infinite play in a parity game \mathfrak{G} with n positions, and let $w = (c_i)_{i < \omega}$ be the sequence of priorities along p . If $l > \log n$ and \mathcal{A}_l accepts w , then p contains an even cycle.*

Proof. Let $\langle l, \bar{k}, \bar{n} \rangle \in M_l$ be the last state in the run of \mathcal{A}_l before it enters the state $*$, let w' be the prefix of w leading to this state, and let c be the next input letter. By Lemma 5.9, the word w' has a suffix with an even factorisation of length at least

$$2^{l-1} + \cdots + 2^0 = 2^l - 1 \geq n,$$

while $w'c$ has a suffix with an even factorisation $z_o < \cdots < z_{m-1}$ of length $m \geq 2^l - 1 + 1 > n$. By the Pigeon Hole Principle it follows that, after reading the additional letter c , there must be two positions $z_i < z_j$ in the factorisation that correspond to the same vertex v of \mathfrak{G} . Let p_o be the part of p corresponding to the path between these two positions. To see that p_o is an even cycle, note that the minimal priority seen along the closed path is the minimal value of c_{z_i}, \dots, c_{z_j} . In particular, it is even. Hence, the play is winning. \square

Lemma 5.11. *The automaton \mathcal{A}_l accepts every word w satisfying the parity condition.*

Proof. By induction on l we will prove that, starting from an arbitrary state $\langle s, \bar{k}, \bar{n} \rangle$ the automaton \mathcal{A}_l accepts every infinite word w that satisfies the parity condition. Hence, fix l and a run ρ of \mathcal{A}_l on w . If ρ contains the accepting state $*$, we are done. Hence, suppose otherwise.

If every state of size l appears only finitely often in ρ , some suffix of ρ is a run of \mathcal{A}_{l-1} . By inductive hypothesis, this suffix is accepting. Hence, so is ρ .

Consequently, we may assume that some state $\langle l, \bar{k}', \bar{n}' \rangle$ of size l appears infinitely often in ρ . The way the transitions are defined it follows that, after the first appearance of $\langle l, \bar{k}', \bar{n}' \rangle$ every state $\langle s, \bar{k}'', \bar{n}'' \rangle$ in ρ satisfies $k''_o = k'_o$. Let ρ' be the sequence of states obtained from ρ by (i) removing the part before the first appearance of $\langle l, \bar{k}', \bar{n}' \rangle$ and (ii) removing the first components of all remaining states, i.e., replacing $\langle s, \bar{k}'', \bar{n}'' \rangle$ by

$$\langle s-1, \langle k''_1, \dots, k''_{s-1} \rangle, \langle n''_1, \dots, n''_{s-1} \rangle \rangle.$$

Then ρ' is a run of \mathcal{A}_{l-1} on the corresponding suffix of w . Again it follows by inductive hypothesis that this run is accepting. Hence, so is ρ . \square

We are finally able to prove Theorem 5.7. Set $l := \lceil \log n \rceil + 1$ and let $\alpha : M_l \times E \rightarrow M_l$ be the action induced by \mathcal{A}_l . We claim that a position v of \mathfrak{G} belongs to the winning region of Player \diamond in \mathfrak{G} if, and only if, Player \diamond has a strategy in the game $\mathfrak{G} \times_{\alpha} M_l$ from the position $\langle v, \langle \circ, \langle \rangle, \langle \rangle \rangle$ to reach some position of the form $\langle u, * \rangle$. Since the latter game is a reachability game of size $n \times \mathcal{O}(n^{\log n}) = \mathcal{O}(n^{\log n})$ and we can compute its winning regions in linear time, the theorem follows. Hence, it remains to prove the claim.

(\Leftarrow) follows immediately by Lemma 5.10 and the implication (2) \Rightarrow (1) in Lemma 5.8, while (\Rightarrow) follows by Lemma 5.11 and the implication (1) \Rightarrow (3) in Lemma 5.8.

Notes

One of the first articles on games is by Zermelo [9], who proved the determinacy of chess. Conway games were subsequently fully developed by Conway [5].

Gale and Steward [6] proved the existence of indetermined games and the determinacy for open games. The full proof of Borel determinacy is by Martin [8].

The section on parity games follows [10] and [7]. The Theorem of Büchi and Landweber was originally proved in [2], and Theorem 5.4 is taken from [4]. The algorithm to solve parity games in Theorem 5.7 is from [3]. Our presentation owes much to a set of lecture notes by Bojańczyk and Czerwiński [1].

References

- [1] M. BOJAŃCZYK AND W. CZERWIŃSKI, *An Automata Toolbox*. lecture notes, 2018.
- [2] J. R. BÜCHI AND L. H. LANDWEBER, *Solving sequential conditions by finite-state strategies*, Trans. Amer. Math. Soc., 139 (1969), pp. 367–378.
- [3] C. S. CALUDE, S. JAIN, B. KHOUSSAINOV, W. LI, AND F. STEPHAN, *Deciding parity games in quasipolynomial time*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19–23, 2017, 2017, pp. 252–263.
- [4] T. COLCOMBET AND D. NIWINSKI, *On the positional determinacy of edge-labeled games*, Theoretical Computer Science, 352 (2006), pp. 190–196.
- [5] J. H. CONWAY, *On Numbers and Games*, Academic Press, 1976.
- [6] D. GALE AND F. M. STEWARD, *Infinite Games with Perfect Information*, in Contributions to the Theory of Games, vol. 28 of Annals of Mathematical Studies, Princeton University Press, 1953, pp. 245–266.
- [7] E. GRÄDEL AND I. WALUKIEWICZ, *Positional Determinacy of Games with Infinitely Many Priorities*, Logical Methods in Computer Science, 2 (2006).
- [8] D. A. MARTIN, *Borel Determinacy*, Annals of Mathematics, 102 (1975), pp. 363–371.
- [9] E. ZERMELO, *Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels*, in Proc. 5th Congress of Mathematicians, Cambridge University Press, 1913, pp. 501–504.
- [10] W. ZIELONKA, *Infinite Games on Finitely Coloured Graphs with Applications to Automata on Infinite Trees*, Theoretical Computer Science, 200 (1998), pp. 135–183.

Symbol Index

Chapter 1

V_{\diamond}	positions for \diamond , 2
V_{\square}	positions for \square , 2
\mathcal{W}_{σ}	winning region, 3
$\diamond X$	some successor in X , 4
$\square X$	all successor in X , 4
$\bar{\sigma}$	opponent, 4
Step_{σ}	step function, 4
Attr_{σ}	attractor, 4
DStep_{σ}	double step function, 13
$\langle \sigma \rangle X$	some σ -successor in X , 13
$[\sigma]X$	all σ -successors in X , 13
$\mathfrak{S} \times_{\alpha} M$	product game, 22
P_{σ}	period set, 24

Index

alternating game, 2	period set, 24
attractor, 4	play, 2
	play of a Conway game, 12
Borel determinacy, 16	player, 2
Borel set, 16	position, 1
	positional determinacy, 3
conforming to a finite-memory strategy, 22	positional game, 2
conforming to a strategy, 2	positional strategy, 2
Conway game, 12	prefix-invariant, 24
	priority function, 17
determined Conway games, 13	product game, 22
determined games, 3	
σ -domain, 18	rank for reachability games, 5
draw, 2	reachability game, 3
	regular game, 16
even cycle, 30	regular winning condition, 24
finite-memory strategy, 21	simultaneous game, 2
	strategy, 2
Gale-Steward game, 14	
game with memory, 22	trap, 6
Horn formula, 9	win, 2
Horn satisfiability, 9	winning a game, 3
	winning region, 3
memory-free strategy, 2	winning strategy, 3
modal logic, 11	
move, 2	
outcome, 2	
parity condition, 17	
parity game, 17	
partial play, 2	