

Logic and Games

WS 2015/2016

Prof. Dr. Erich Grädel
Notes and Revisions by Matthias Voit

Mathematische Grundlagen der Informatik
RWTH Aachen



This work is licensed under:

<http://creativecommons.org/licenses/by-nc-nd/3.0/de/>

Dieses Werk ist lizenziert unter:

<http://creativecommons.org/licenses/by-nc-nd/3.0/de/>

© 2016 Mathematische Grundlagen der Informatik, RWTH Aachen.

<http://www.logic.rwth-aachen.de>

Contents

1	Reachability Games and First-Order Logic	1
1.1	Model Checking	1
1.2	Model Checking Games for Modal Logic	2
1.3	Reachability and Safety Games	5
1.4	Games as an Algorithmic Construct: Alternating Algorithms	10
1.5	Model Checking Games for First-Order Logic	20
2	Parity Games and Fixed-Point Logics	25
2.1	Parity Games	25
2.2	Algorithms for parity games	30
2.3	Fixed-Point Logics	35
2.4	Model Checking Games for Fixed-Point Logics	37
2.5	Defining Winning Regions in Parity Games	42
3	Infinite Games	45
3.1	Determinacy	45
3.2	Gale-Stewart Games	47
3.3	Topology	53
3.4	Determined Games	59
3.5	Muller Games and Game Reductions	61
3.6	Complexity	74
4	Basic Concepts of Mathematical Game Theory	79
4.1	Games in Strategic Form	79
4.2	Nash equilibria	81
4.3	Two-person zero-sum games	85
4.4	Regret minimization	86
4.5	Iterated Elimination of Dominated Strategies	89
4.6	Beliefs and Rationalisability	95

4.7 Games in Extensive Form	98
4.8 Subgame-perfect equilibria in infinite games	102
Appendix A	111
4.9 Cardinal Numbers	119

3 Infinite Games

After our treatment of reachability (and safety) games in the first, and parity games in the second chapter, we now discuss infinite games in more general setting. More precisely, the games that we study are *two-player, zero-sum games of perfect information, played on game graphs and admitting infinite plays*.

Formally, a *graph game* is a pair $\mathcal{G} = (G, \text{Win})$ where $G = (V, V_0, V_1, E, \Omega)$ is a directed graph with $V = V_0 \cup V_1$ and $\Omega : V \rightarrow C$ for a set C of *colours* (or priorities) and a set $\text{Win} \subseteq C^\omega$ of infinite sequences of colours. We call G the *arena* of \mathcal{G} and Win the *winning condition* of \mathcal{G} .

As before a *play* of \mathcal{G} is a finite or infinite sequence $\pi = v_0v_1v_2 \dots \in V^{\leq\omega}$ such that $(v_i, v_{i+1}) \in E$ for all i . A finite play is lost by the player who cannot move any more, and an infinite play π is won by Player 0 if $\Omega(\pi) = \Omega(v_0)\Omega(v_1) \dots \in \text{Win}$, otherwise Player 1 wins (there are no draws). Let $\text{Plays}(\mathcal{G})$ denote the set of all plays of G and $P_{\text{fin}}(\mathcal{G})$ be set of all initial segments $x \in V^*$ of a play in $\text{Plays}(\mathcal{G})$

3.1 Determinacy

A *strategy* for Player σ in a game $\mathcal{G} = (G, \text{Win})$ is a function $f : V^*V_\sigma \rightarrow V$ such that $(v, f(xv)) \in E$ for all $x \in V^*$ and $v \in V_\sigma$. Thus, a strategy maps prefixes of plays which end in a position in V_σ to legal moves of Player σ . A play $\pi = v_0v_1 \dots$ is *consistent with a strategy f* for Player σ if for all proper prefixes $v_0 \dots v_n$ of π such that $v_n \in V_\sigma$ we have $v_{n+1} = f(v_0 \dots v_n)$. We say that f is a *winning strategy* from position v_0 if every play starting in v_0 that is consistent with f is won by Player σ . The set

$$W_\sigma = \{v \in V : \text{Player } \sigma \text{ has a winning strategy from } v\}$$

is the *winning region* of Player σ . In zero-sum games it always holds that $W_0 \cap W_1 = \emptyset$. We call a game \mathcal{G} *determined* if $W_0 \cup W_1 = V$, i.e. if from each position one player has a winning strategy.

We can generalize the notion of winning regions from initial positions to arbitrary initial segments of plays. Let \tilde{W}_σ be the set of those initial segments $x \in V^*$ of plays for which Player σ has a strategy f to prolong x to a winning play (i.e. every play of form $x\pi \in \text{Plays}(\mathcal{G})$ that is consistent with f is won by Player σ). Clearly if $P_{\text{fin}}(\mathcal{G}) = \tilde{W}_0 \cup \tilde{W}_1$ then \mathcal{G} is determined.

For determinacy questions it suffices to consider games played on trees and forests. Indeed, for an arena G with a node v_0 , let $\mathcal{T}(G, v_0)$ be the tree obtained by unraveling G from v_0 . Obviously, a player has a winning strategy for (G, Win) from v_0 if, and only if she has one for $(\mathcal{T}(G, v_0), \text{Win})$ for the root v_0 . For the forest $\mathcal{F}(G) := \bigcup_{v \in G} \mathcal{T}(G, v)$ we then have that $(\mathcal{F}(G), \text{Win})$ is determined if, and only if, (G, Win) is determined. Notice further that, on trees and forests, all strategies are positional so in this case there is no difference between determinacy and determinacy via positional strategies.

A classical and very old determinacy theorem is due to Zermelo who proved that a game of this kind is always determined if it only admits finite plays. A slightly stronger variant of this result, applying to games with infinite plays, is the following.

Theorem 3.1 (Zermelo). Let \mathcal{G} be a game such that in every play the winner is determined after finitely many moves. Then \mathcal{G} is determined.

Proof. The condition that the winner of every play is determined after finitely many moves means that every infinite play π of G has a finite initial segment $x < \pi$ such that every play of form $x\pi'$ is won by the same player. We claim that this implies that $P_{\text{fin}}(\mathcal{G}) = \tilde{W}_0 \cup \tilde{W}_1$ and hence the determinacy of \mathcal{G} .

Let $X = P_{\text{fin}}(\mathcal{G}) \setminus (\tilde{W}_0 \cup \tilde{W}_1)$, and assume, towards a contradiction, that $X \neq \emptyset$. Take some $x = yv \in X$, with $v \in V_\sigma$.

For all $w \in vE$ it follows that $xw = yvw \notin \tilde{W}_\sigma$ (because otherwise $x \in \tilde{W}_\sigma$). Further, if we had that $xw \in \tilde{W}_{1-\sigma}$ for all $w \in vE$, then also $x \in \tilde{W}_{1-\sigma}$. Thus there exists some prolongation xw of x with $xw \in X$.

By induction, there exists an infinite play $x\pi$ such that $xy \in X$ for all finite y . In particular the winner of $x\pi$ is not determined after any finite initial segment, which contradicts our initial assumption. Q.E.D.

The game that Zermelo originally wanted to study is Chess, which does not quite satisfy our definition of a game given above, since it admits draws. One thus has to slightly modify the determinacy statement for Chess.

Corollary 3.2. For Chess one of the following three possibilities holds:

- White has a winning strategy.
- Black has a winning strategy.
- Both players have a strategy to enforce at least a draw.

In the previous chapter, we proved a strong determinacy theorem for parity games. We now look for general properties of Win that guarantee determinacy. To answer this question we shall need topological arguments. But before we develop them, we introduce the notion of a Gale-Stewart game and prove the existence of non-determined games.

3.2 Gale-Stewart Games

In this chapter we will show that, using the Axiom of Choice, one can construct a non-determined game. Later, we will mention which topological properties guarantee determinacy and how this is related to logic.

Let B be an alphabet (for instance $B = \{0,1\}$ or $B = \omega$). In a Gale-Stewart game the players alternately choose symbols from B in an infinite sequence of moves and thus create an infinite word $\pi \in B^\omega$. Gale-Stewart games can be described as graph games in different ways. For $B = \{0,1\}$, for example, as a game on the infinite binary tree

$$\mathcal{T}^2 = (\{0,1\}^*, V_0, V_1, E, \Omega),$$

where

$$V_0 = \bigcup_{n \in \omega} \{0,1\}^{2n},$$

3 Infinite Games

$$V_1 = \bigcup_{n \in \omega} \{0,1\}^{2n+1},$$

$$E = \{(x, xi) : x \in \{0,1\}^*, i \in \{0,1\}\},$$

and $\Omega : \{0,1\}^* \rightarrow \{0,1,\varepsilon\} : \varepsilon \mapsto \varepsilon, xi \mapsto i$. Alternatively, it can be described as a game on the graph depicted in Figure 3.1. Similar game graphs can be defined for arbitrary B .




Figure 3.1. Game graph for Gale-Stewart game over $B = \{0,1\}$

Theorem 3.3 (Gale-Stewart). There exists a non-determined Gale-Stewart game.

We shall present two proofs. The first one uses enumerations of the strategy spaces of the two player via ordinals (see Appendix A) up to 2^ω . The second uses ultrafilters. Both rely on the Axiom of Choice (AC).

Proof. For any countable alphabet B with at least two symbols, let $T_0 = \{x \in B^* : |x| \text{ even}\}$ and $T_1 = \{x \in B^* : |x| \text{ odd}\}$. Then

$$F = \{f : T_0 \rightarrow B\} \text{ and } G = \{g : T_1 \rightarrow B\}$$

are the sets of strategies for Player 0 and for Player 1. Since B is countable, $|F| = |G| = |\mathcal{P}(\omega)| = 2^\omega$. Thus, using the well-ordering principle (which is equivalent to AC) we can enumerate the strategies by ordinals up to 2^ω :

$$F = \{f_\alpha : \alpha < 2^\omega\} \text{ and } G = \{g_\alpha : \alpha < 2^\omega\}.$$

For strategies f and g let $f \hat{\ } g \in B^\omega$ be the uniquely determined play arising from f and g . We shall construct two increasing sequences of sets $X_\alpha, Y_\alpha \subseteq B^\omega$ for $\alpha < 2^\omega$ such that

$$(1) X_\alpha \cap Y_\alpha = \emptyset,$$

$$(2) |X_\alpha|, |Y_\alpha| < 2^\omega,$$

Let $X_0 = Y_0 = \emptyset$. For a successor ordinal $\alpha = \beta + 1$ consider the strategy f_β . The cardinality of X_β and Y_β is smaller than 2^ω but there are 2^ω different strategies $g \in G$ and thus 2^ω different plays that are consistent with f_β . Hence there exists one that is not in X_β . Choose such a play $f_\beta \hat{g}$ (AC again) and add it to Y_β to construct $Y_\alpha := Y_\beta \cup \{f_\beta \hat{g}\}$. Analogously, choose a play $f \hat{g}_\beta$ that is consistent with g_β and which is not in Y_α , and construct $X_\alpha := X_\beta \cup \{f \hat{g}_\beta\}$. For limit ordinals λ let $X_\lambda := \bigcup_{\beta < \lambda} X_\beta$ and $Y_\lambda := \bigcup_{\beta < \lambda} Y_\beta$.

We claim that the Gale-Stewart game with winning condition $\text{Win} := \bigcup_{\alpha < 2^\omega} X_\alpha$ is not determined.

Indeed, assume that $f = f_\alpha$, for some $\alpha < 2^\omega$, is a winning strategy for Player 0. By the construction of Win , there is a strategy $g \in G$ such that $f_\alpha \hat{g} \in Y_{\alpha+1}$ and thus $f_\alpha \hat{g} \notin \text{Win}$, a contradiction.

Now assume that $g = g_\alpha$, for some $\alpha < 2^\omega$, is a winning strategy for Player 1. Analogously, there is a strategy $f \in F$ such that $f \hat{g}_\alpha \in X_{\alpha+1} \subseteq \text{Win}$, a contradiction as well. Q.E.D.

The second proof that we shall present uses the concept of an ultrafilter. We first recall the definition of a filter.

Definition 3.4. Let I be a non-empty set. A non-empty set $F \subseteq \mathcal{P}(I)$ is a *filter* if

- (1) $\emptyset \notin F$,
- (2) $x \in F, y \in F \Rightarrow x \cap y \in F$, and
- (3) $x \in F, y \supseteq x \Rightarrow y \in F$.

The intuition behind a filter is that it is a family of large sets.

Example 3.5. The set $\{x \subseteq \omega : \omega \setminus x \text{ is finite}\}$ is a filter. We call it the *Fréchet filter*.

Definition 3.6. An *ultrafilter* is a filter that satisfies the additional requirement:

- (4) for all $x \subseteq I$ either $x \in F$ or $I \setminus x \in F$.

Example 3.7. Fix $n \in \omega$. Then $U_n = \{a \subseteq \omega : n \in a\}$ is an ultrafilter. Ultrafilters of this form are called *principal ultrafilters*.

Every ultrafilter U that contains a finite set must be principal. Otherwise U would contain a smallest set a which is not a singleton. Pick some $n \in a$. Since $\{n\} \notin U$, the complement $\omega \setminus \{n\}$ is in U , and hence also its intersection with a . But $a \cap (\omega \setminus \{n\}) = a \setminus \{n\} \subsetneq a$ contradicting the minimality of a in U .

On the other side, an ultrafilter that does not contain a finite set must contain all co-finite ones, and thus extend the Fréchet filter. But the Fréchet filter is not an ultrafilter and it is not obvious that it can be extended to one in a consistent way. The proof that this is possible uses Zorn's Lemma or the Compactness Theorem for propositional logic. It holds for every set $F \subseteq \mathcal{P}(\omega)$ such that $a_1 \cap \dots \cap a_m \neq \emptyset$ for all $m \in \mathbb{N}$, $a_1, \dots, a_m \in F$.

Theorem 3.8. The Fréchet filter F can be expanded to an ultrafilter $U \supset F$.

Proof. Let F be the Fréchet filter. We use propositional variables X_a for every $a \in \mathcal{P}(\omega)$. Let $\Phi = \Phi_U \cup \Phi_F$ where

$$\begin{aligned} \Phi_U = & \{ \neg X_\emptyset \} \\ & \cup \{ X_a \wedge X_b \rightarrow X_{a \cap b} : a, b \subseteq \omega \} \\ & \cup \{ X_a \rightarrow X_b : a \subseteq b, a, b \subseteq \omega \} \\ & \cup \{ X_a \leftrightarrow \neg X_{\omega \setminus a} : a \subseteq \omega \} \end{aligned}$$

and

$$\Phi_F = \{ X_a : a \in F \}.$$

Every model \mathcal{I} of Φ defines an ultrafilter U which expands F , namely $U = \{ a \subseteq \omega : \mathcal{I}(X_a) = 1 \}$. It remains to show that Φ is satisfiable.

By the compactness theorem, it suffices to show that every finite subset of Φ is satisfiable. Hence, let Φ_0 be a finite subset of Φ . Then the set $F_0 = \{ a \in F : X_a \in \Phi_0 \}$ is also finite. Now consider the following two cases:

- $F_0 = \emptyset$. Define the interpretation \mathcal{I} by

$$\mathcal{I}(X_a) = \begin{cases} 1 & \text{if } 0 \in a, \\ 0 & \text{otherwise.} \end{cases}$$

Then $\mathcal{I} \models \Phi_0$.

- $F_0 = \{a_1, \dots, a_m\}$. Since F is a filter, there exists $n_0 \in a_1 \cap \dots \cap a_m$. Define the interpretation \mathcal{I} by

$$\mathcal{I}(X_a) = \begin{cases} 1 & \text{if } n_0 \in a \\ 0 & \text{otherwise} \end{cases}$$

Again, we have $\mathcal{I} \models \Phi_0$.

Hence, Φ_0 is satisfiable.

Q.E.D.

We are now able to give an alternative construction for non-determined games. Let U be an ultrafilter that expands the Fréchet filter. We construct a Gale-Stewart game over $B = \omega$ with winning condition Win_U as follows. Player 0 wins a play $x = x_0x_1\dots \in \omega^\omega$ if

- Player 1 has played a number that is not higher than the previously played one, i.e. $\min\{j : x_{j+1} \leq x_j\}$ exists and is even, or
- $x_0 < x_1 < x_2 < \dots$ and

$$A(x) := [0, x_0) \cup \bigcup_{i \in \omega} [x_{2i+1}, x_{2i+2}) \in U$$




Figure 3.2. The winning condition of the ultrafilter game

Proposition 3.9. The Gale-Stewart game with winning condition Win_U is not determined.

Proof. Towards a contradiction, assume that Player 0 has a winning strategy f . We construct two plays x and x' , both of which are consistent with f .

- In the first play the opening move $x_0 = f(\varepsilon)$ of Player 0 is answered by Player 1 with an arbitrary number $x_1 > x_0$. The second move of Player 0 is then $x_2 = f(x_0x_1)$.
- In the second play x' , Player 1 uses x_2 as her answer to the opening move $x_0 = f(\varepsilon)$ by Player 0. The second move of Player 0 in the play x' is then $x_3 = f(x_0x_2)$, and Player 1 uses this in the play x as her answer to $x_0x_1x_2$.
- This is iterated. In play x , Player 1 extends in her $(i + 1)$ st move the sequence $x_0x_1 \dots x_{2i}$ by $x_{2i+1} = f(x_0x_2x_3 \dots x_{2i})$, i.e. she just copies the $(i + 1)$ st move of Player 0 in play x' .
- Similarly, in play x' , Player 1 answers the initial segment $x_0x_2x_3 \dots x_{2i+1}$ by $x_{2i+2} = f(x_0x_1 \dots x_{2i+1})$, i.e she copies the $i + 1$ st move of Player 1 in x .

Thus, in both plays, Player 1 essentially uses the strategy f itself as a counterstrategy against f .




Figure 3.3. Playing the Ultrafilter game

This results in two plays $x = x_0x_1x_2 \dots$ and $x' = x_0x_2x_3x_4 \dots$, where $x_{2i+2} = f(x_0x_1 \dots x_{2i+1})$ but also $x_{2i+1} = f(x_0x_1 \dots x_{2i})$. Both plays are consistent with the winning strategy f for Player 0. Thus we have $A(x) \in U$ and $A(x') \in U$. But

$$A(x) = [0, x_0) \cup \bigcup_{i \in \omega} [x_{2i+1}, x_{2i+2})$$

and

$$A(x') = [0, x_0) \cup \bigcup_{i \in \omega} [x_{2i+2}, x_{2i+3}).$$

Thus $A(x) \cap A(x') = [0, x_0) \in U$. However, since U expands the Fréchet

filter, the co-finite set $\omega \setminus [0, x_0)$ is in U and thus $[0, x_0) \notin U$, a contradiction.

Analogously, one derives a contradiction from the assumption that Player 1 has a winning strategy. Q.E.D.

3.3 Topology

Definition 3.10. A *topology* on a set S is defined by a collection of *open* subsets of S . It is required that

- \emptyset , and S are open;
- if X and Y are open, then $X \cap Y$ is open;
- if $\{X_i : i \in I\}$ is a family of open sets, then $\bigcup_{i \in I} X_i$ is open.

If $\mathcal{O} \subseteq \mathcal{P}(S)$ is a collection of open sets, we call the pair (S, \mathcal{O}) a *topological space*.

Often, a topology is defined by its *base*: A set B of open subsets of S such that every open set can be represented as a union of sets in B .

Example 3.11. The standard topology on \mathbb{R} is defined by the base consisting of all open intervals $(a, b) \subseteq \mathbb{R}$.

In our setting, we will only be concerned with the following topology on B^ω , which is due to Cantor. Its base consists of all sets of the form $z \uparrow := z \cdot B^\omega$ for $z \in B^*$. Consequently, a set $X \subseteq B^\omega$ is *open* if it is the union of sets $z \uparrow$, i.e. if there exists a set $W \subseteq B^*$ such that $X = W \cdot B^\omega$. Moreover, a set $X \subseteq B^\omega$ is *closed* if its complement $B^\omega \setminus X$ is open. For $B = \{0, 1\}$, this topology is called the *Cantor space*, and for $B = \omega$ it is called the *Baire space*.




Figure 3.4. Base sets in the Cantor space

Example 3.12.

- The base sets z^\uparrow are both open and closed (*clopen*) since we have $B^\omega \setminus z^\uparrow = W_z \cdot B^\omega$ where $W_z = \{y \in B^* \mid y \not\leq z \text{ and } z \not\leq y\}$. (Here, $u \leq v$ means that u is a prefix of v .)
- $0^*1\{0, 1\}^\omega$ is open. The complement $\{0^\omega\}$ is closed, but not open.
- $L_d = \{x \in \omega^\omega : x \text{ contains } d \text{ infinitely often}\} = \bigcap_{n \in \omega} (\omega^* \cdot d)^n \cdot \omega^\omega$ is a countable intersection of open sets.

Next, we will give another useful characterisation of closed sets. A *tree* $T \subseteq B^*$ is a prefix-closed set of finite words, i.e., $z \in T$ and $y \leq z$ implies $y \in T$. For a tree T let $[T]$ be the set of infinite paths through T (note: $T \subseteq B^*$, but $[T] \subseteq B^\omega$).

Example 3.13. Let $T = 0^* = \{0^n : n \in \omega\}$. Then $[T] = \{0^\omega\}$.

Lemma 3.14. $X \subseteq B^\omega$ is closed if and only if there exists a tree $T \subseteq B^*$ such that $X = [T]$.

Proof.

(\Rightarrow) Let X be closed. Then there is a $W \subseteq B^*$ such that $B^\omega \setminus X = W \cdot B^\omega$. Let $T := \{w \in B^* \mid \forall z(z \leq w \Rightarrow z \notin W)\}$. T is closed under prefixes and $[T] = X$.

(\Leftarrow) Let $X = [T]$. For every $x \notin [T]$ there exists a smallest prefix $w_x \leq x$ such that $w_x \notin T$. Let $W := \{w_x : x \notin X\}$. Then $B^\omega \setminus X = W \cdot B^\omega$ is open, thus X is closed. Q.E.D.

We call a set $W \subseteq B^*$ *prefix-free* if there is no pair $x, y \in W$ such that $x < y$.

Lemma 3.15.

- (1) For every open set $A \subseteq B^\omega$ there is a prefix-free set $W \subseteq B^*$ such that $A = W \cdot B^\omega$.
- (2) Let B be a finite alphabet. $A \subseteq B^\omega$ is clopen if and only if there is a finite set $W \subseteq B^*$ such that $A = W \cdot B^\omega$.

Proof. For (1), let $A = U \cdot B^\omega$ for some open $U \subseteq B^*$. Define

$$W := \{w \in U : U \text{ contains no proper prefix of } w\}.$$

W is prefix-free and $W \cdot B^\omega = U \cdot B^\omega = A$.

For (2) let $A \subseteq B^\omega$ be clopen. Thus there exist prefix-free sets $U, V \subseteq B^*$ such that $A = U \cdot B^\omega$ and $B^\omega \setminus A = V \cdot B^\omega$. We will show that $U \cup V$ is finite. Let $T = \{w \in B^* \mid w \text{ has no prefix in } U \cup V\}$. If T is finite, then $U \cup V$ is also finite. If U (or V) is infinite, then T is also infinite since it contains all proper prefixes of elements of U (respectively V). Hence it suffices to show that T is finite. Notice that T is a finitely branching tree (since B is finite) that contains no infinite path, since otherwise there exists an infinite word $x \in B^\omega$ corresponding to this path with $x \notin U \cdot B^\omega \cup V \cdot B^\omega = A \cup (B^\omega \setminus A) = B^\omega$. By König's Lemma, this implies that T is finite.

For the converse, let $A = W \cdot B^\omega$ where $W \subseteq B^*$ is finite. Let $l = \max\{|w| : w \in W\}$. Then $B^\omega \setminus A = Z \cdot B^\omega$ where

$$Z = \{z \in B^* : |z| = l \text{ and no prefix of } z \text{ is in } W\}.$$

Thus, A is clopen.

Q.E.D.

Notice that (2) does not hold for infinite alphabets B .

Definition 3.16. Let $T = (S, \mathcal{O})$ be a topological space. The class of *Borel sets* is the smallest class $\mathcal{B} \subseteq \mathcal{P}(S)$ that contains all open sets and is closed under countable unions and complementation:

- $\mathcal{O} \subseteq \mathcal{B}$;
- If $X \in \mathcal{B}$ then $S \setminus X \in \mathcal{B}$;
- If $\{X_n : n \in \omega\} \subseteq \mathcal{B}$ then $\bigcup_{n \in \omega} X_n \in \mathcal{B}$.

Most of the ω -languages $L \subseteq B^\omega$ occurring in Computer Science are Borel sets. Borel sets form a natural hierarchy of sets Σ_α^0 and Π_α^0 for $0 \leq \alpha < \omega_1$, where ω_1 is the first uncountable ordinal number.

- $\Sigma_1^0 = \mathcal{O}$;
- $\Pi_\alpha^0 = \text{co}\Sigma_\alpha^0 := \{S \setminus X : X \in \Sigma_\alpha^0\}$ for every α ;
- $\Sigma_\alpha^0 = \{\bigcup_{n \in \omega} X_n : X_n \in \Pi_\beta^0 \text{ for } \beta < \alpha\}$ for $\alpha > 0$.

We are especially interested in the first levels of the Borel hierarchy:

- Σ_1^0 : Open sets
- Π_1^0 : Closed sets

- Σ_2^0 : Countable unions of closed sets
- Π_2^0 : Countable intersections of open sets
- Σ_3^0 : Countable unions of Π_2^0 -sets
- Π_3^0 : Countable intersections of Σ_2^0 -sets

Example 3.17. Let $d \in B$.

$$L_d = \{x \in B^\omega : x \text{ contains } d \text{ infinitely often}\} = \bigcap_{n \in \omega} \underbrace{(B^* \cdot d)^n \cdot B^\omega}_{\in \Sigma_1^0}.$$

Hence, $L_d \in \Pi_2^0$.

To determine the membership of an ω -language in a class Σ_α^0 or Π_α^0 of the Borel hierarchy and to relate the classes, we need a notion of reducibility between ω -languages.

Definition 3.18. A function $f : B^\omega \rightarrow C^\omega$ is called *continuous* if $f^{-1}(Y)$ is open for every open set $Y \subseteq C^\omega$.

Let $X \subseteq B^\omega$, $Y \subseteq C^\omega$. We say that X is *Wadge reducible* to Y , $X \leq Y$, if there exists a continuous function $f : B^\omega \rightarrow C^\omega$ such that $f^{-1}(Y) = X$, i.e. $x \in X$ iff $f(x) \in Y$ for all $x \in B^\omega$. For any such function f , we write $f : X \leq Y$.

Exercise 3.1. Prove that the relation \leq satisfies the following properties:

- $X \leq Y$ and $Y \leq Z$ imply $X \leq Z$;
- $X \leq Y$ implies $B^\omega \setminus X \leq C^\omega \setminus Y$.

Theorem 3.19. Let $X \leq Y$ for $Y \in \Sigma_\alpha^0$ (or $Y \in \Pi_\alpha^0$). Then $X \in \Sigma_\alpha^0$ (respectively $X \in \Pi_\alpha^0$).

Proof. The claim is true by definition for Σ_1^0 (the open sets) and thus also for Π_1^0 .

For $\alpha > 1$, let $f : X \leq Y$ and $Y \in \Sigma_\alpha^0$. We have that $Y = \bigcup_{n \in \omega} Y_n$ where $Y_n \in \bigcup_{\beta < \alpha} \Pi_\beta^0$. Define $X_n := f^{-1}(Y_n)$. Then $X_n \leq Y_n$ for all $n \in \omega$, and thus, by induction hypothesis, $X_n \in \bigcup_{\beta < \alpha} \Pi_\beta^0$. We have:

$$\begin{aligned} x \in X &\Leftrightarrow f(x) \in Y \\ &\Leftrightarrow f(x) \in Y_n \text{ for some } n \in \omega \\ &\Leftrightarrow x \in X_n \text{ for some } n \in \omega. \end{aligned}$$

Hence, $X = \bigcup_{n \in \omega} X_n \in \Sigma_\alpha^0$.

Q.E.D.

In the following we will present a game-theoretic characterisation of the relation \leq in terms of the so-called *Wadge game*.

Definition 3.20. Let $X \subseteq B^\omega$, $Y \subseteq C^\omega$. The *Wadge game* $W(X, Y)$ is an infinite game between two players 0 and 1 who move in alternation. In the i -th round, Player 0 chooses a symbol $x_i \in B$, and afterwards Player 1 chooses a (possibly empty) word $y_i \in C^*$. After ω rounds, Player 0 has produced an ω -word $x = x_0x_1x_2 \cdots \in B^\omega$, and Player 1 has produced a finite or infinite word $y = y_0y_1y_2 \cdots \in C^{\leq\omega}$. Player 1 wins the play (x, y) if, and only if, $y \in C^\omega$ and $x \in X \Leftrightarrow y \in Y$.

Example 3.21. Let $B = C = \{0, 1\}$.

- Player 1 wins $W(0^*1\{0, 1\}^\omega, (0^*1)^\omega)$.
Winning strategy for Player 1: Choose 0 until Player 0 chooses 1 for the first time. Afterwards, always choose 1.
- Player 0 wins $W((0^*1)^\omega, 0^*1\{0, 1\}^\omega)$.
Winning strategy for Player 0: Choose 1 until Player 1 chooses a word containing 1 for the first time. Afterwards, always choose 0.

Theorem 3.22 (Wadge). Let $X \subseteq B^\omega$, $Y \subseteq C^\omega$. Then $X \leq Y$ if and only if Player 1 has a winning strategy for $W(X, Y)$.

Proof.

(\Leftarrow) A winning strategy of Player 1 for $W(X, Y)$ induces a mapping $f : B^\omega \rightarrow C^\omega$ such that $x \in X$ iff $y \in Y$. It remains to show that f is continuous. Let $Z = U \cdot C^\omega$ be open. For every $u \in U$ denote by V_u the set of all words $v = x_0x_1 \dots x_n \in B^*$ such that u is the answer of Player 1 to v , i.e. $u = f(x_0)f(x_1) \dots f(x_n)$. Then $f^{-1}(U \cdot C^\omega) = V \cdot B^\omega$ where $V := \bigcup_{u \in U} V_u$.

(\Rightarrow) Let $f : X \leq Y$. We construct a strategy for Player 1 as follows. Player 1 has to answer Player 0's moves $x_0x_1x_2 \dots$ by an ω -word $y_0y_1y_2 \dots$, but Player 1 can delay choosing y_i until he knows $x_0x_1 \dots x_n$ for some appropriate $n \geq i$.

Choice of y_0 : Consider the partition $B^\omega = \bigcup_{c \in C} f^{-1}(c \cdot C^\omega)$. Since $c \cdot C^\omega$ is clopen, $f^{-1}(c \cdot C^\omega)$ is also clopen. For every $x \in B^\omega$ there exists $c \in C$ such that $x \in f^{-1}(c \cdot C^\omega)$, and since $f^{-1}(c \cdot C^\omega)$ is clopen, there is a prefix $w_x \leq x$ such that $w_x \cdot B^\omega \subseteq f^{-1}(c \cdot C^\omega)$. So Player 1 can

wait until Player 0 has chosen a prefix $w \in B^*$ that determines the set $f^{-1}(c \cdot C^\omega)$ the word x will belong to and choose $y_0 = c$.

The subsequent choices are done analogously. Let $y_0 \dots y_i \in C^*$ be Player 1's answer to $x_0 \dots x_n \in B^*$. For the choice of y_{i+1} we consider the partition

$$x_0 \dots x_n \cdot B^\omega = \bigcup_{c \in C} f^{-1}(y_0 \dots y_i \cdot c \cdot C^\omega).$$

Since the sets $f^{-1}(y_0 \dots y_i \cdot c \cdot C^\omega)$ are clopen, after finitely many moves, by choosing a prolongation $x_0 \dots x_n x_{n+1} \dots x_k$, Player 0 has determined in which set $f^{-1}(y_0 \dots y_i \cdot c \cdot C^\omega)$ the word x will be. Player 1 then chooses $y_{i+1} = c$.

By using this strategy, Player 1 constructs the answer $y = f(x)$ for the sequence x chosen by Player 0. Otherwise, there would be a smallest i such that $y_i \neq f(x_i)$. This is impossible since $x \in f^{-1}(y_0 \dots y_i \cdot C^\omega)$. Since $f : X \leq Y$, we have $x \in X$ iff $y \in Y$. Q.E.D.

Definition 3.23. A set $Y \subseteq C^\omega$ is Σ_α^0 -complete if:

- $Y \in \Sigma_\alpha^0$;
- $X \leq Y$ for all $X \in \Sigma_\alpha^0$.

Π_α^0 -completeness is defined analogously.

Note that Y is Σ_α^0 -complete if, and only if, $C^\omega \setminus Y$ is Π_α^0 -complete.

Proposition 3.24. Let $B = \{0, 1\}$. Then:

- $0^*1\{0, 1\}^\omega$ is Σ_1^0 -complete;
- $\{0^\omega\}$ is Π_1^0 -complete;
- $\{0, 1\}^*0^\omega$ is Σ_2^0 -complete;
- $(0^*1)^\omega$ is Π_2^0 -complete.

Proof. By the above remark, it suffices to show that $0^*1\{0, 1\}^\omega$ and $(0^*1)^\omega$ are Σ_1^0 -complete and Π_2^0 -complete, respectively.

- We know that $0^*1\{0, 1\}^\omega \in \Sigma_1^0$. Let $X = W \cdot B^\omega$ be open. We describe a winning strategy for Player 1 in $W(X, 0^*1\{0, 1\}^\omega)$: Pick 0 until Player 0 has completed a word contained in W ; from this point onwards, pick 1. Hence, $X \leq 0^*1\{0, 1\}^\omega$.

- We know that $(0^*1)^\omega \in \Pi_2^0$. Let $X = \bigcap_{n \in \omega} W_n \cdot B^\omega \in \Pi_2^0$. We describe a winning strategy for Player 1 in $W(X, \{0,1\}^*0^\omega)$: Start with $i := 0$; for arbitrary i , answer with 1 and set $i := i + 1$ if the sequence $x_0 \dots x_k$ of symbols chosen by Player 0 so far has a prefix in W_i , otherwise answer with 0 and leave i unaffected. Q.E.D.

3.4 Determined Games

We call a game $\mathcal{G} = (V, V_0, V_1, E, \text{Win})$ clopen, open, closed, etc., or simply a *Borel game*, if the winning condition $\text{Win} \subseteq V^\omega$ has the respective property.

Clopen games are basically finite games: If $A \subseteq B^\omega$ is clopen, then for every $x \in B^\omega$ there exists a finite prefix $w_x \leq x$ such that:

- If $x \in A$ then $w_x \uparrow \subseteq A$;
- If $x \notin A$ then $w_x \uparrow \subseteq B^\omega \setminus A$.

Thus, by Zermelo's Theorem, clopen games are determined.

A stronger result is the following:

Theorem 3.25. Every open game, and thus every closed game, is determined.

Proof. Let $\mathcal{G} = (V, V_0, V_1, E, \text{Win})$ where $\text{Win} = U \cdot V^\omega$ is open. First, we consider finite plays: Let $T_\sigma = \{v \in V_{1-\sigma} : vE = \emptyset\}$ and $A_\sigma = \text{Attr}_\sigma(T_\sigma)$. From every position $v \in A_\sigma$ Player σ wins after finitely many moves with the attractor strategy.

For the infinite plays consider

$$\mathcal{G}' := \mathcal{G} \upharpoonright V \setminus (A_0 \cup A_1)$$

with positions $V' := V \setminus (A_0 \cup A_1)$. In \mathcal{G}' every play is infinite, and Player 0 wins $\pi = v_0v_1v_2\dots$ if and only if $\pi \in U \cdot V^\omega$. Obviously, Player 0 wins in \mathcal{G}' starting from v_0 if she can enforce a sequence $v_0v_1\dots v_n \in U$. Then every infinite prolongation of this sequence is a play in $U \cdot V^\omega$.

Instead of \mathcal{G}' we consider again the equivalent game on the trees $\mathcal{T}(v) = \mathcal{T}_{\mathcal{G}}(v)$, the unfolding of \mathcal{G} from $v \in V$. Positions in $\mathcal{T}(v)$ are

words over V : $\mathcal{T}(v) \subseteq V^*$. Now consider the set

$$B_0 = \{v \in V' : v \in \text{Attr}_0^{\mathcal{T}(v)}(U \cdot V^*)\}$$

of positions from where player 0 can enforce a play prefix in $U \cdot V^*$. From every position in $V' \setminus A_0$, Player 1 has a strategy to guarantee that the play never reaches $U \cdot V^*$ since $V' \setminus A_0$ is a trap for Player 0. But a play that never reaches $U \cdot V^*$ is won by Player 1. It follows that $W_0 = A_0 \cup B_0$ and $W_1 = A_1 \cup (V' \setminus B_0)$, and thus $V = W_0 \cup W_1$. Q.E.D.

A much stronger result was established by Donald Martin in 1975. Its proof is beyond the scope of these lecture notes.

Theorem 3.26 (Martin). All Borel games are determined.

Here are some winning conditions for frequently used games in Computer Science:

- *Muller conditions:* Let B be finite, $\mathcal{F}_0 \subseteq \mathcal{P}(B)$, $\mathcal{F}_1 = \mathcal{P}(B) \setminus \mathcal{F}_0$. Player σ wins $\pi \in B^\omega$ if and only if

$$\text{Inf}(\pi) := \{b \in B : b \text{ appears infinitely often in } \pi\} \in \mathcal{F}_\sigma.$$

Hence, the winning condition is the set

$$\{x \in B^\omega : \text{Inf}(\pi) \in \mathcal{F}_\sigma\} = \bigcup_{X \in \mathcal{F}_0} \left(\bigcap_{d \in X} L_d \cap \bigcup_{d \notin X} (B^\omega \setminus L_d) \right),$$

a finite Boolean combination of Π_2^0 -sets.

- *Parity conditions* (see previous chapter) are special cases of Muller conditions and thus also finite Boolean combinations of Π_2^0 -sets.
- Every ω -regular language is a Boolean combination of Π_2^0 -sets. This follows from the recognisability of ω -regular languages by Muller automata and the fact that Muller conditions are Boolean combinations of Π_2^0 -sets.

In practice, winning conditions are often specified in a suitable logic: ω -words $x \in B^\omega$ are interpreted as structures $\mathfrak{A}_x = (\omega, <, (P_b)_{b \in B})$ with unary predicates $P_b = \{i \in \omega : x_i = b\}$. A sentence ψ (for example

in FO, MSO, etc.) over the signature $\{<\} \cup \{P_b : b \in B\}$ defines the ω -language (winning condition) $L(\psi) = \{x \in B^\omega : \mathfrak{A}_x \models \psi\}$.

Example 3.27. Let $B = \{0, \dots, m\}$. The parity condition is specified by the FO sentence

$$\psi := \bigwedge_{\substack{b \leq m \\ b \text{ odd}}} \left(\exists y \forall z (y < z \rightarrow \neg P_b z) \vee \bigwedge_{c < b} \forall y \exists z (y < z \wedge P_c z) \right).$$

We have:

- FO and LTL define the same ω -languages (winning conditions);
- MSO defines exactly the ω -regular languages;
- There are ω -languages that are definable in MSO but not in FO;
- ω -regular languages are Boolean combinations of Π_2^0 -sets.

In particular, graph games with winning conditions specified in LTL, FO, MSO, etc. are Borel games and therefore determined.

3.5 Muller Games and Game Reductions

Muller games are infinite games played over an arena $G = (V_0, V_1, E, \Omega : V \rightarrow C)$ with a winning condition depending only on the set of priorities seen infinitely often in a play. It is specified by a partition $\mathcal{P}(C) = \mathcal{F}_0 \cup \mathcal{F}_1$, and a play $\pi = v_0 v_1 v_2 \dots$ is won by Player σ if

$$\text{Inf}(\pi) = \{c : \Omega(v_i) = c \text{ for infinitely many } i \in \omega\} \in \mathcal{F}_\sigma.$$

We will only consider the case that the set C of priorities is finite. Then Muller games are Borel games specified by the FO sentence

$$\bigvee_{X \in \mathcal{F}_\sigma} \left(\bigwedge_{c \in X} \forall x \exists y (x < y \wedge P_c y) \wedge \bigwedge_{c \notin X} \exists x \forall y (x < y \rightarrow \neg P_c y) \right).$$

So Muller games are determined. Parity conditions are special Muller conditions, and we have seen that games with parity winning conditions are even positionally determined. The question arises what kind of strategies are needed to win Muller games. Unfortunately, there are

simple Muller games that are not positionally determined, even solitaire games.

Example 3.28. Consider the game arena depicted in Figure 3.5 with the winning condition $\mathcal{F}_0 = \{\{1, 2, 3\}\}$, i.e. all positions have to be visited infinitely often. Obviously, player 0 has winning a winning strategy, but no positional one: Any positional strategy of player 0 will either visit only positions 1 and 2 or positions 2 and 3.



Figure 3.5. A solitaire Muller game

Although Muller games are, in general, not positionally determined, we will see that Muller games are determined via winning strategies that can be implemented using finite memory. To this end, we introduce the notions of a memory structure and of a memory strategy. Although we will not require that the memory is finite, we will use finite memory in most cases.

Definition 3.29. A *memory structure* for a game \mathcal{G} with positions in V is a triple $\mathfrak{M} = (M, \text{update}, \text{init})$, where M is a set of *memory states*, $\text{update} : M \times V \rightarrow M$ is a *memory update function* and $\text{init} : V \rightarrow M$ is a *memory initialisation function*. The *size* of the memory is the cardinality of the set M .

A *strategy with memory* \mathfrak{M} for Player σ is given by a next-move function $F : V_\sigma \times M \rightarrow V$ such that $F(v, m) \in vE$ for all $v \in V_\sigma, m \in M$. If a play, from starting position v_0 , has gone through positions $v_0 v_1 \dots v_n$, the memory state is $m(v_0 \dots v_n)$, defined inductively by $m(v_0) = \text{init}(v_0)$, and $m(v_0 \dots v_i v_{i+1}) = \text{update}(m(v_0 \dots v_i), v_{i+1})$, and in case $v_n \in V_\sigma$ the strategy leads to position $F(v_n, m(v_0 \dots v_n))$.

Remark 3.30. In case $|M| = 1$, the strategy is positional, and it can be described by a function $F : V_\sigma \rightarrow V$.

Definition 3.31. A game \mathcal{G} is determined via memory \mathfrak{M} if it is determined and both players have winning strategies with memory \mathfrak{M} on their winning regions.

Example 3.32. In the game from Example 3.28, Player 0 wins with a strategy with memory $\mathfrak{M} = (\{1,3\}, \text{update}, \text{init})$ where

- $\text{init}(1) = \text{init}(2) = 1, \text{init}(3) = 3$ and
- $\text{update}(m, v) = \begin{cases} v & \text{if } v \in \{1,3\}, \\ m & \text{if } v = 2. \end{cases}$

The corresponding strategy is defined by

$$F(v, m) = \begin{cases} 2 & \text{if } v \in \{1,3\}, \\ 3 & \text{if } v = 2, m = 1, \\ 1 & \text{if } v = 2, m = 3. \end{cases}$$

Let us consider a more interesting example now.

Example 3.33. Consider the game DJW_2 with its arena depicted in Figure 3.6. Player 0 wins a play π if the maximal number in $\text{Inf}(\pi)$ is equal to the number of letters in $\text{Inf}(\pi)$. Formally:

$$\mathcal{F}_0 = \{X \subseteq \{1,2,a,b\} : |X \cap \{a,b\}| = \max(X \cap \{1,2\})\}.$$




Figure 3.6. Muller game $\mathcal{G} = \text{DJW}_2$

Player 0 has a winning strategy from every position, but no positional one. Assume that $f : \{a,b\} \rightarrow \{1,2\}$ is a positional winning strategy for Player 0. If $f(a) = 2$ (or $f(b) = 2$), then Player 1 always picks a (respectively b) and wins, since this generates a play π with $\text{Inf}(\pi) = \{a,2\}$ (respectively $\text{Inf}(\pi) = \{b,2\}$). If $f(a) = f(b) = 1$, then Player 1 alternates between a and b and wins, since this generates a play π with $\text{Inf}(\pi) = \{a,b,1\}$. However, Player 0 has a winning strategy that uses the memory depicted in Figure 3.7. The corresponding strategy is

defined as follows:

$$F(c, m) = \begin{cases} 1 & \text{if } m = c\#d, \\ 2 & \text{if } m = \#cd. \end{cases}$$




Figure 3.7. Memory for Player 0

Why is this strategy winning? If from some point onwards Player 1 picks only a or only b , then, from this point onwards, the memory state is always $b\#a$ or $a\#b$, respectively, and according to F Player 0 always picks 1 and wins. In the other case, Player 1 picks a and b again and again and the memory state is $\#ab$ or $\#ba$ infinitely often. Thus Player 0 picks 2 infinitely often and wins as well.

The memory structure used in this example is a special case of the LAR memory structure, which we will use for arbitrary Muller games. But first, let us look at a Muller game with infinitely many priorities that allows no winning strategy with finite memory but one with a simple infinite memory structure:

Example 3.34. Consider the game with its arena depicted in Figure 3.8 and with winning condition $\mathcal{F}_0 = \{\{0\}\}$. It is easy to see that every finite-memory strategy of Player 0 (the player who moves at position 0) is losing. A winning strategy with infinite memory is given by the memory structure $\mathfrak{M} = (\omega, \text{init}, \text{update})$ where $\text{init}(v) = v$ and $\text{update}(m, v) = \max(m, v)$ together with the strategy F defined by $F(0, m) = m + 1$.




Figure 3.8. A game where finite-memory strategies do not suffice

Given a game graph $G = (V, V_0, V_1, E)$ and a memory structure $\mathfrak{M} = (M, \text{update}, \text{init})$, we obtain a new game graph

$$G \times \mathfrak{M} = (V \times M, V_0 \times M, V_1 \times M, E_{\text{update}})$$

where

$$E_{\text{update}} = \{((v, m), (v', m')) : (v, v') \in E \text{ and } m' = \text{update}(m, v')\}.$$

Obviously, every play $(v_0, m_0)(v_1, m_1) \dots$ in $G \times \mathfrak{M}$ has a unique projection to the play $v_0 v_1 \dots$ in G . Conversely, every play v_0, v_1, \dots in G has a unique extension to a play $(v_0, m_0)(v_1, m_1) \dots$ in $G \times \mathfrak{M}$ with $m_0 = \text{init}(v_0)$.

Definition 3.35. For games $\mathcal{G} = (G, \Omega, \text{Win})$ and $\mathcal{G}' = (G', \Omega', \text{Win}')$, we say that \mathcal{G} *reduces to* \mathcal{G}' *via memory* \mathfrak{M} , $\mathcal{G} \leq_{\mathfrak{M}} \mathcal{G}'$, if $G' = G \times \mathfrak{M}$ and every play in \mathcal{G}' is won by the same player as the projected play in \mathcal{G} .

Given a memory structure \mathfrak{M} for G and a memory structure \mathfrak{M}' for $G \times \mathfrak{M}$, we obtain a memory structure $\mathfrak{M}^* = \mathfrak{M} \times \mathfrak{M}'$ for G . The set of memory locations is $M \times M'$, and we have memory initialisation

$$\text{init}^*(v) = (\text{init}(v), \text{init}'(v, \text{init}(v)))$$

with the update function

$$\begin{aligned} \text{update}^*((m, m'), v) = \\ (\text{update}(m, v), \text{update}'(m', (v, \text{update}(m, v)))). \end{aligned}$$

Theorem 3.36. Suppose that \mathcal{G} reduces to \mathcal{G}' via memory \mathfrak{M} and that

Player σ has a winning strategy for \mathcal{G}' with memory \mathfrak{M}' from position $(v_0, \text{init}(v_0))$. Then Player σ has a winning strategy for \mathcal{G} with memory $\mathfrak{M} \times \mathfrak{M}'$ from position v_0 .

Proof. Given a strategy $F' : (V_\sigma \times M) \times M' \rightarrow (V \times M)$ for Player σ in \mathcal{G}' , we have to construct a strategy $F : (V_\sigma \times (M \times M')) \rightarrow V$ for Player σ in \mathcal{G} . For any $v \in V_\sigma$ and any pair $(m, m') \in M \times M'$ we have that $F'((v, m), m') = (w, \text{update}(m, w))$ for some $w \in vE$. We put $F(v, (m, m')) = w$. If a play in \mathcal{G} that is consistent with F proceeds from position v with current memory location (m, m') to a new position w , then the memory is updated to (n, n') with $n = \text{update}(m, w)$ and $n' = \text{update}'(m', (w, n))$. In the extended play in \mathcal{G}' , we have an associated move from (v, m) to (w, n) with memory update from m' to n' . Thus, every play in \mathcal{G} from initial position v_0 that is consistent with F is the projection of a play in \mathcal{G}' from $(v_0, \text{init}(v_0))$ that is consistent with F' . Therefore, if F' is a winning strategy from $(v_0, \text{init}(v_0))$, then F is a winning strategy from v_0 . Q.E.D.

Corollary 3.37. Every game that reduces via memory \mathfrak{M} to a positionally determined game is determined via memory \mathfrak{M} .

Obviously, memory reductions between games can be composed. If \mathcal{G} reduces to \mathcal{G}' with memory $\mathfrak{M} = (M, \text{update}, \text{init})$ and \mathcal{G}' reduces to \mathcal{G}'' with memory $\mathfrak{M}' = (M', \text{update}', \text{init}')$ then \mathcal{G} reduces to \mathcal{G}'' with memory $(M \times M', \text{update}'', \text{init}'')$ where

$$\text{init}''(v) = (\text{init}(v), \text{init}'(v, \text{init}(v)))$$

and

$$\text{update}''((m, m'), v) = (\text{update}(m, v), \text{update}'(m', (v, \text{update}(m, v)))).$$

The classical example of a game reduction with finite memory is the reduction of Muller games to parity games via latest appearance records. Intuitively, a *latest appearance record* (LAR) is a list of priorities ordered by their latest occurrence. More formally, for a finite set C of priorities, $\text{LAR}(C)$ is the set of sequences $c_1 \dots c_k \# c_{k+1} \dots c_l$ of elements from $C \cup \{\#\}$ in which each priority $c \in C$ occurs at most once and $\#$ occurs precisely once. At a position v , the LAR $c_1 \dots c_k \# c_{k+1} \dots c_l$ is

updated by moving the priority $\Omega(v)$ to the end, and moving $\#$ to the previous position of $\Omega(v)$ in the sequence. For instance, at a position with priority c_2 , the LAR $c_1c_2c_3\#c_4c_5$ is updated to $c_1\#c_3c_4c_5c_2$. (If $\Omega(v)$ did not occur in the LAR, we simply append $\Omega(v)$ at the end). Thus, the LAR memory for an arena with priority labelling $\Omega : V \rightarrow C$ is the triple $(\text{LAR}(C), \text{update}, \text{init})$ with $\text{init}(v) = \#\Omega(v)$ and

$$\text{update}(c_1 \dots c_k \# c_{k+1} \dots c_l, v) = \begin{cases} c_1 \dots c_k \# c_{k+1} \dots c_l \Omega(v) & \text{if } \Omega(v) \notin \{c_1, \dots, c_l\}, \\ c_1 \dots c_{m-1} \# c_{m+1} \dots c_l c_m & \text{if } \Omega(v) = c_m. \end{cases}$$

The *hit set* of an LAR $c_1 \dots c_k \# c_{k+1} \dots c_l$ is the set $\{c_{k+1} \dots c_l\}$ of priorities occurring after the symbol $\#$. Note that if in a play $\pi = v_0 v_1 \dots$ the LAR at position v_n is $c_1 \dots c_k \# c_{k+1} \dots c_l$, then $\Omega(v_n) = c_l$ and the hit set $\{c_{k+1} \dots c_l\}$ is the set of priorities that have been visited since the latest previous occurrence of c_l in the play.

Lemma 3.38. Let π be a play of a Muller game \mathcal{G} with finitely many priorities, and let $\text{Inf}(\pi)$ be the set of priorities occurring infinitely often in π . Then the hit set of the latest appearance record is, from some point onwards, always a subset of $\text{Inf}(\pi)$ and infinitely often coincides with $\text{Inf}(\pi)$.

Proof. For each play $\pi = v_0 v_1 v_2 \dots$ there is a position v_m such that $\Omega(v_n) \in \text{Inf}(\pi)$ for all $n \geq m$. Since no priority outside $\text{Inf}(\pi)$ is seen after position v_m , the hit set will, from that position onwards, always be contained in $\text{Inf}(\pi)$, and the LAR will always have the form $c_1 \dots c_{j-1} c_j \dots c_k \# c_{k+1} \dots c_l$ where c_1, \dots, c_{j-1} remains fixed and $\text{Inf}(\pi) = \{c_j, \dots, c_l\}$. Since all priorities in $\text{Inf}(\pi)$ are seen again and again, it happens infinitely often that, among these, the one occurring leftmost in the LAR is hit. At such positions, the LAR is updated to $c_1, \dots, c_{j-1} \# c_{j+1} \dots c_l c_j$, and the hit set coincides with $\text{Inf}(\pi)$. Q.E.D.

Theorem 3.39. Every Muller game with finitely many priorities reduces via LAR memory to a parity game.

Proof. Let \mathcal{G} be a Muller game with game graph G , priority labelling $\Omega : V \rightarrow C$ and winning condition $(\mathcal{F}_0, \mathcal{F}_1)$. We have to prove that

$\mathcal{G} \leq_{\text{LAR}} \mathcal{G}'$ for a parity game \mathcal{G}' with game graph $G \times \text{LAR}(C)$ and an appropriate priority labelling Ω' on $V \times \text{LAR}(C)$, which is defined as follows:

$$\Omega'(v, c_1 c_2 \dots c_k \# c_{k+1} \dots c_l) = \begin{cases} 2k & \text{if } \{c_{k+1}, \dots, c_l\} \in \mathcal{F}_0, \\ 2k+1 & \text{if } \{c_{k+1}, \dots, c_l\} \in \mathcal{F}_1. \end{cases}$$

Let $\pi = v_0 v_1 v_2 \dots$ be a play on \mathcal{G} and fix a number m such that, for all $n \geq m$, $\Omega(v_n) \in \text{Inf}(\pi)$ and the LAR at position v_n has the form $c_1 \dots c_j c_{j+1} \dots c_k \# c_{k+1} \dots c_l$ where $\text{Inf}(\pi) = \{c_{j+1}, \dots, c_l\}$ and the prefix $c_1 \dots c_j$ remains fixed. In the corresponding play $\pi' = (v_0, r_0)(v_1, r_1) \dots$ in \mathcal{G}' , all nodes (v_n, r_n) for $n \geq m$ have a priority $2k + \rho$ with $k \geq j$ and $\rho \in \{0, 1\}$. Assume that the play π is won by Player σ , i.e., $\text{Inf}(\pi) \in \mathcal{F}_\sigma$. Since the hit set of the LAR coincides with $\text{Inf}(\pi)$ infinitely often, the minimal priority seen infinitely often on the extended play is $2j + \sigma$. Thus the extended play in the parity game \mathcal{G}' is won by the same player as the original play in \mathcal{G} . Q.E.D.

Corollary 3.40. Muller games are determined via finite memory strategies. The size of the memory is bounded by $(|C| + 1)!$.

The question arises which Muller conditions $(\mathcal{F}_0, \mathcal{F}_1)$ guarantee positional winning strategies for arbitrary game graphs? One obvious answer are parity conditions. But there are others:

Example 3.41. Let $C = \{0, 1\}$, $\mathcal{F}_0 = \{C\}$ and $\mathcal{F}_1 = \mathcal{P}(C) \setminus \{C\} = \{\{0\}, \{1\}, \emptyset\}$. $(\mathcal{F}_0, \mathcal{F}_1)$ is not a parity condition, but every Muller game with winning condition $(\mathcal{F}_0, \mathcal{F}_1)$ is positionally determined.

Definition 3.42. The *Zielonka tree* for a Muller condition $(\mathcal{F}_0, \mathcal{F}_1)$ over C is a tree $Z(\mathcal{F}_0, \mathcal{F}_1)$ whose nodes are labelled with pairs (X, σ) such that $X \in \mathcal{F}_\sigma$. We define $Z(\mathcal{F}_0, \mathcal{F}_1)$ inductively as follows. Let $C \in \mathcal{F}_\sigma$ and C_0, \dots, C_{k-1} be the maximal sets in $\{X \subseteq C : X \in \mathcal{F}_{1-\sigma}\}$. Then $Z(\mathcal{F}_0, \mathcal{F}_1)$ consists of a root, labelled with (C, σ) , to which we attach as subtrees the Zielonka trees $Z(\mathcal{F}_0 \cap \mathcal{P}(C_i), \mathcal{F}_1 \cap \mathcal{P}(C_i))$, $i = 0, \dots, k-1$.

Example 3.43. Let $C = \{0, 1, 2, 3, 4\}$ and $\mathcal{F}_0 = \{\{0, 1\}, \{2, 3, 4\}, \{2, 3\}, \{2, 4\}, \{3\}, \{4\}\}$, $\mathcal{F}_1 = \mathcal{P}(C) \setminus \mathcal{F}_0$. The Zielonka tree $Z(\mathcal{F}_0, \mathcal{F}_1)$ is depicted in Figure 3.9.




Figure 3.9. A Zielonka tree

A set $Y \subseteq C$ belongs to \mathcal{F}_σ if there is a node of $Z(\mathcal{F}_0, \mathcal{F}_1)$ that is labelled with (X, σ) for some $X \supseteq Y$ and for all children $(Z, 1 - \sigma)$ of (X, σ) we have $Y \not\subseteq Z$.

Example 3.44. Consider again the tree $Z(\mathcal{F}_0, \mathcal{F}_1)$ from Example 3.43. It is the case that $\{2, 3\} \in \mathcal{F}_0$, since $(\{2, 3, 4\}, 0)$ is a node of $Z(\mathcal{F}_0, \mathcal{F}_1)$ and

- $\{2, 3\} \subseteq \{2, 3, 4\}$;
- $\{2, 3\} \not\subseteq \{2\}$;
- $\{2, 3\} \not\subseteq \{3, 4\}$.

The Zielonka tree of a parity-condition is especially simple, as Figure 3.10 shows.

Besides parity games there are other important special cases of Muller games. Of special relevance are games with Rabin and Streett conditions because these admit positional winning strategies for one player.

Definition 3.45. A *Streett-Rabin condition* is a Muller condition $(\mathcal{F}_0, \mathcal{F}_1)$ such that \mathcal{F}_0 is closed under union.

In the Zielonka tree for a Streett-Rabin condition, the nodes labelled with $(X, 1)$ have only one successor. It follows that if both \mathcal{F}_0 and \mathcal{F}_1 are closed under union, then the Zielonka tree $Z(\mathcal{F}_0, \mathcal{F}_1)$ is a path, which implies that $(\mathcal{F}_0, \mathcal{F}_1)$ is equivalent to a parity condition.




Figure 3.10. The Zielonka tree of a parity-condition with m priorities

In a Streett-Rabin game, Player 1 has a positional winning strategy on his winning region. On the other hand, Player 0 can win on his winning region via a finite-memory strategy, and the size of the memory can be directly read off from the Zielonka tree. We present an elementary proof of this result.

Theorem 3.46. Let $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ be a game with a Streett-Rabin winning condition $(\mathcal{F}_0, \mathcal{F}_1)$. Then \mathcal{G} is determined, i.e. $V = W_0 \cup W_1$, with a finite memory winning strategy for Player 0 on W_0 , and a positional winning strategy for Player 1 on W_1 . The size of the memory required by the winning strategy for Player 0 is bounded by the number of leaves of the Zielonka tree $Z(\mathcal{F}_0, \mathcal{F}_1)$.

Proof. We proceed by induction on the number of priorities in C or, equivalently, the depth of the Zielonka tree $Z(\mathcal{F}_0, \mathcal{F}_1)$. Let l be the number of leaves of $Z(\mathcal{F}_0, \mathcal{F}_1)$. We distinguish two cases.

Case 1: $C \in \mathcal{F}_1$. Let

$$X_0 := \left\{ v : \begin{array}{l} \text{Player 0 has a winning strategy with memory} \\ \text{of size } \leq l \text{ from } v \end{array} \right\},$$

and $X_1 = V \setminus X_0$. It suffices to prove that Player 1 has a positional winning strategy on X_1 . To construct this strategy, we combine three

positional strategies of Player 1: A trap strategy, an attractor strategy, and a winning strategy on a subgame with fewer priorities.

At first, we observe that X_1 is a trap for Player 0. This means that Player 1 has a positional trap strategy t on X_1 to enforce that the play stays within X_1 .

Since \mathcal{F}_0 is closed under union, there is a unique maximal subset $C' \subseteq C$ with $C' \in \mathcal{F}_0$. Let $Y := X_1 \cap \Omega^{-1}(C \setminus C')$, and let $Z = \text{Attr}_1(Y) \setminus Y$. Observe that Player 1 has a positional attractor strategy a , by which he can force, from any position $z \in Z$, that the play reaches Y .

Finally, let $V' = X_1 \setminus (Y \cup Z)$ and let \mathcal{G}' be the subgame of \mathcal{G} induced by V' , with winning condition $(\mathcal{F}_0 \cap \mathcal{P}(C'), \mathcal{F}_1 \cap \mathcal{P}(C'))$ (see Figure 3.11). Since this game has fewer priorities, the induction hypothesis applies, i.e. we have $V' = W'_0 \cup W'_1$, and Player 0 has a winning strategy with memory $\leq l$ on W'_0 , whereas Player 1 has a positional winning strategy g' on W'_1 . However, $W'_0 = \emptyset$: Otherwise we could combine the strategies of Player 0 to obtain a winning strategy with memory $\leq l$ on $X_0 \cup W'_0 \supseteq X_0$, a contradiction to the definition of X_0 . Hence $W'_1 = V'$.




Figure 3.11. Constructing a winning strategy for Player 1

We can now define a positional strategy g for Player 1 on X_1 by

$$g(x) = \begin{cases} g'(x) & \text{if } x \in V', \\ a(x) & \text{if } x \in Z, \\ t(x) & \text{if } x \in Y. \end{cases}$$

Consider any play π that starts at a position $v \in X_1$ and is consistent with g . We have to show that π is won by Player 1. Obviously, π stays

within X_1 . If it hits $Y \cup Z$ only finitely often, then from some point onwards it stays within V' and coincides with a play consistent with g' . It is therefore won by Player 1. Otherwise, π hits $Y \cup Z$, and hence also Y , infinitely often. Thus, $\text{Inf}(\pi) \cap (C \setminus C') \neq \emptyset$ and $\text{Inf}(\pi) \in \mathcal{F}_1$. So Player 1 has a positional winning strategy on X_1 .

Case 2: $C \in \mathcal{F}_0$. There exist maximal subsets $C_0, \dots, C_{k-1} \subseteq C$ with $C_i \in \mathcal{F}_1$ (see Figure 3.12). Observe that if $D \cap (C \setminus C_i) \neq \emptyset$ for all $i < k$ then $D \in \mathcal{F}_0$. Now let

$$X_1 := \{v \in V : \text{Player 1 has a positional winning strategy from } v\},$$

and $X_0 = V \setminus X_1$. We claim that Player 0 has a finite memory winning strategy of size $\leq l$ on X_0 . To construct this strategy, we proceed in a similar way as above, for each of the sets $C \setminus C_i$. We will obtain strategies f_0, \dots, f_{k-1} for Player 0 such that each f_i has finite memory M_i , and we will use these strategies to build a winning strategy f on X_0 with memory $M_0 \cup \dots \cup M_{k-1}$.




Figure 3.12. The top of the Zielonka tree $Z(\mathcal{F}_0, \mathcal{F}_1)$

For $i = 0, \dots, k-1$, let $Y_i = X_0 \cap \Omega^{-1}(C \setminus C_i)$, and $Z_i = \text{Attr}_0(Y_i) \setminus Y_i$, and let a_i be a positional attractor strategy by which Player 0 can force a play from any position in Z_i to reach Y_i . Furthermore, let $U_i = X_0 \setminus (Y_i \cup Z_i)$, and let \mathcal{G}_i be the subgame of \mathcal{G} induced by U_i with winning condition $(\mathcal{F}_0 \cap \mathcal{P}(C_i), \mathcal{F}_1 \cap \mathcal{P}(C_i))$. The winning region of Player 1 in \mathcal{G}_i is empty: Indeed, if Player 1 could win \mathcal{G}_i from v , then, by the induction hypothesis, he could win with a positional winning strategy. By combining this strategy with the positional winning strategy of Player 1 on X_1 , this would imply that $v \in X_1$, but $v \in U_i \subseteq V \setminus X_1$.

Hence, by the induction hypothesis, Player 0 has a winning strategy f_i with finite memory M_i on U_i . Let $(f_i + a_i)$ be the combination of f_i

with the attractor strategy a_i , defined by

$$(f_i + a_i)(v) := \begin{cases} f_i(v) & \text{if } v \in U_i, \\ a_i(v) & \text{if } v \in Z_i. \end{cases}$$

From any position $v \in U_i \cup Z_i$ this strategy ensures that the play either remains inside U_i and is winning for Player 1, or that it eventually reaches a position in Y_i .

We now combine the strategies $(f_0 + a_0), \dots, (f_{k-1} + a_{k-1})$ to a winning strategy f on X_0 , which ensures that either the play ultimately remains within one of the regions U_i and coincides with a play according to f_i , or that it cycles infinitely often through all the regions Y_0, \dots, Y_{k-1} .

At positions in $\tilde{Y} := \bigcap_{i < k} Y_i$, Player 0 just plays with a (positional) trap strategy t ensuring that the play remains in X_0 . At the first position $v \notin \tilde{Y}$, Player 0 takes the minimal i such that $v \notin Y_i$, i.e. $v \in U_i \cup Z_i$, and uses the strategy $(f_i + a_i)$ until a position $w \in Y_i$ is reached. At this point, Player 0 switches from i to $j = i + l \pmod{k}$ for the minimal l such that $w \notin Y_j$. Hence $w \in U_j \cup Z_j$; Player 0 now plays with strategy $(f_j + a_j)$ until a position in Y_j is reached. There Player 0 again switches to the appropriate next strategy, as he does every time he reaches \tilde{Y} .

Assuming that $M_i \cap M_j = \emptyset$ for $i \neq j$, it is not difficult to see that f can be implemented with memory $M = M_0 \cup \dots \cup M_{k-1}$. We leave the formal definition of f as an exercise.

Note that, by the induction hypothesis, the size of the memory M_i is bounded by the number of leaves of the Zielonka subtrees $Z(\mathcal{F}_0 \cap \mathcal{P}(C_i), \mathcal{F}_1 \cap \mathcal{P}(C_i))$. Consequently, the size of M is bounded by the number of leaves of $Z(\mathcal{F}_0, \mathcal{F}_1)$.

It remains to prove that f is winning on X_0 . Let π be a play that starts in X_0 and is consistent with f . If π eventually remains inside some U_i , then from some point onwards it coincides with a play that is consistent with f_i and is therefore won by Player 0. Otherwise, it is easy to see that π hits each of the sets Y_0, \dots, Y_{k-1} infinitely often. But this means that $\text{Inf}(\pi) \cap (C \setminus C_i) \neq \emptyset$ for all $i \leq k$; as observed above this implies that $\text{Inf}(\pi) \in \mathcal{F}_0$. Q.E.D.

An immediate consequence of Theorem 3.46 is that parity games are positionally determined.

3.6 Complexity

We will now determine the complexity of computing the winning regions for games over finite game graphs. The associated decision problem is

Given: Game graph \mathcal{G} , winning condition $(\mathcal{F}_0, \mathcal{F}_1)$, $v \in V$,
 $\sigma \in \{0, 1\}$.
Question: $v \in W_\sigma$?

For parity games, we already know that this problem is in $\text{NP} \cap \text{coNP}$, and it is conjectured to be in P. Moreover, for many special cases, we know that it is indeed in P. Now we will examine the complexity of Streett-Rabin games and games with arbitrary Muller conditions.

Theorem 3.47. Deciding whether Player σ wins from a given position in a Streett-Rabin game is

- coNP-hard for $\sigma = 0$,
- NP-hard for $\sigma = 1$.

Proof. It is sufficient to prove the claim for $\sigma = 1$ since Streett-Rabin games are determined. We will reduce the satisfiability problem for Boolean formulae in CNF to the given problem. For every formula

$$\Psi = \bigwedge_i C_i, \quad C_i = \bigvee_j Y_{ij}$$

in CNF, we define the game \mathcal{G}_Ψ as follows: Positions for Player 0 are the literals $X_1, \dots, X_k, \neg X_1, \dots, \neg X_k$ occurring in Ψ ; positions for Player 1 are the clauses C_1, \dots, C_n . Player 1 can move from a clause C to a literal $Y \in C$; Player 0 can move from Y to any clause. The winning condition is given by

$$\mathcal{F}_0 = \{Z : \{X, \neg X\} \subseteq Z \text{ for at least one variable } X\}.$$

Obviously, $(\mathcal{F}_0, \mathcal{F}_1)$ is a Streett-Rabin condition.

We claim that Ψ is satisfiable if and only if Player 1 wins \mathcal{G}_Ψ (from any initial position).

(\Rightarrow) Assume that Ψ is satisfiable. There exists a satisfying interpretation $\mathcal{I} : \{X_1, \dots, X_k\} \rightarrow \{0, 1\}$. Player 1 moves from a clause C to a literal $Y \in C$ such that $\llbracket Y \rrbracket^{\mathcal{I}} = 1$. In the resulting play only literals with $\llbracket Y \rrbracket^{\mathcal{I}} = 1$ are seen, and thus Player 1 wins.

(\Leftarrow) Assume that Ψ is unsatisfiable. It is sufficient to show that Player 1 has no positional winning strategy. Every positional strategy f for Player 1 chooses a literal $Y = f(C) \in C$ for every clause C . Since Ψ is unsatisfiable, there exist clauses C, C' and a variable X such that $f(C) = X$, $f(C') = \neg X$. Otherwise, f would define a satisfying interpretation for Ψ . Player 0's winning strategy is to move from $\neg X$ to C and from any other literal to C' . Then X and $\neg X$ are seen infinitely often, and Player 0 wins. Thus, f is not a winning strategy for Player 1. If Player 1 has no positional winning strategy, he has no winning strategy at all.

Is $\Psi \mapsto \mathcal{G}_\Psi$ a polynomial reduction? The problem that arises is the winning condition: Both \mathcal{F}_0 and \mathcal{F}_1 contain exponentially many sets. Moreover, the Zielonka tree $Z(\mathcal{F}_0, \mathcal{F}_1)$ has exponential size. On the other hand, \mathcal{F}_0 and \mathcal{F}_1 can be represented in a very compact way using a Boolean formula in the following sense: Let $(\mathcal{F}_0, \mathcal{F}_1)$ be a Muller condition over C . A Boolean formula Ψ with variables in C defines the set $\mathcal{F}_\Psi = \{Y \subseteq C : \mathcal{I}_Y \models \Psi\}$ where

$$\mathcal{I}_Y(c) = \begin{cases} 1 & \text{if } c \in Y \\ 0 & \text{if } c \notin Y. \end{cases}$$

Ψ defines $(\mathcal{F}_0, \mathcal{F}_1)$ if $\mathcal{F}_\Psi = \mathcal{F}_0$ (and thus $\mathcal{F}_{\neg\Psi} = \mathcal{F}_1$). Representing the winning condition by a Boolean formula makes the reduction a polynomial reduction. Q.E.D.

Another way of defining Streett-Rabin games is by a collection of pairs (L, R) with $L, R \subseteq C$. The collection $\{(L_1, R_1), \dots, (L_k, R_k)\}$ defines the Muller condition $(\mathcal{F}_0, \mathcal{F}_1)$ given by:

$$\mathcal{F}_0 = \{X \subseteq C : X \cap L_i \neq \emptyset \Rightarrow X \cap R_i \neq \emptyset \text{ for all } i \leq k\}.$$

We have:

- Every Muller condition defined by a collection of pairs is a Streett-Rabin condition.
- Every Streett-Rabin condition is definable by a collection of pairs.
- Representing a Streett-Rabin condition by a collection of pairs can be exponentially more succinct than a representation by its Zielonka tree or an explicit enumeration of \mathcal{F}_0 or \mathcal{F}_1 : There are Streett-Rabin conditions definable with k pairs such that the corresponding Zielonka tree has $k!$ leaves.

The reduction $\Psi \mapsto \mathcal{G}_\Psi$ can be modified such that the winning condition is given by $2m$ pairs, where m is the number of variables in Ψ :

$$L_{2i} = \{X_i\}, \quad R_{2i} = \{\neg X_i\}, \quad L_{2i-1} = \{\neg X_i\}, \quad R_{2i-1} = \{X_i\}.$$

For the Streett-Rabin condition defined by $\{(L_1, R_1), \dots, (L_{2m}, R_{2m})\}$ we have that

$$\mathcal{F}_1 = \left\{ Z : \begin{array}{l} Z \text{ contains a Literal } X_i \text{ (or } \neg X_i) \text{ such that the} \\ \text{complementary literal } \neg X_i \text{ (respectively } X_i) \text{ is} \\ \text{not contained in } Z \end{array} \right\}.$$

The winning strategies used in the proof remain winning for the modified winning condition.

To prove the upper bounds for the complexity of Streett-Rabin games we will consider solitaire games first.

Theorem 3.48. Let \mathcal{G} be a Streett-Rabin game such that only Player 0 can do non-trivial moves. Then the winning regions W_0 and W_1 can be computed in polynomial time.

Proof. Let us assume that the winning condition is given by the collection $\mathcal{P} = \{(L_1, R_1), \dots, (L_k, R_k)\}$ of pairs. It is sufficient to prove the claim for W_0 since Streett-Rabin games are determined. Every play π will ultimately stay in a strongly connected set $U \subseteq V$ such that all positions in U are seen infinitely often. Therefore, we call a strongly connected set

U good for Player 0 if for all $i \leq k$

$$\Omega(U) \cap L_i \neq \emptyset \Rightarrow \Omega(U) \cap R_i \neq \emptyset.$$

For every such U , $\text{Attr}_0(U) \subseteq W_0$. If U is not good for Player 0 then there is a node in U which violates a pair (L_i, R_i) . In this case Player 0 wants to find a (strongly connected) subset of U where she can win nevertheless. We can eliminate the pairs (L_i, R_i) where $\Omega(U) \cap L_i = \emptyset$ since they never violate the winning condition. On the other hand, Player 0 loses if a node of

$$\tilde{U} = \{u \in U \mid \Omega(u) \in L_i \text{ for some } i \text{ such that } \Omega(U) \cap R_i = \emptyset\}$$

is visited again and again. Thus we will reduce the game from U to $U \setminus \tilde{U}$ with the modified winning condition $\mathcal{P}' = \{(L_i, R_i) \in \mathcal{P} : \Omega(U) \cap L_i \neq \emptyset\}$. This yields Algorithm 3.1.

Algorithm 3.1. A polynomial time algorithm solving solitaire Streett-Rabin games

Algorithm WinReg(G, \mathcal{P})

Input: Streett-Rabin game with game graph G and pairs condition \mathcal{P} .

Output: W_0 , the winning region for Player 0.

$W_0 := \emptyset;$

Decompose G into its SCCs;

For every SCC U **do**

$\mathcal{P}' := \{(L_i, R_i) : \Omega(U) \cap L_i \neq \emptyset\};$

$\tilde{U} := \{u \in U : \Omega(u) \in L_i \text{ for some } i \text{ such that } \Omega(U) \cap R_i = \emptyset\};$

if $\tilde{U} = \emptyset$ **then** $W := W \cup U;$

else $W := W \cup \text{WinReg}(G \upharpoonright_{U \setminus \tilde{U}}, \mathcal{P}');$

enddo;

$W_0 := \text{Attr}_0(W);$

Output $W_0;$

The SCC decomposition can be computed in linear time. The decomposition algorithm will be called less than $|V|$ times, the rest are elementary steps. Therefore, the algorithm runs in polynomial time.

It remains to show that $W_0 = \text{WinReg}(G, \mathcal{P})$:

(\subseteq) Let $v \in W_0$. Player 0 can reach from v a strongly connected set S that satisfies the winning condition. S is a subset of an SCC U of G . If U satisfies the winning condition, then $v \in \text{WinReg}(G, \mathcal{P})$. Otherwise, $S \subseteq U \setminus \tilde{U}$, and S is contained in an SCC of $G \upharpoonright_{U \setminus \tilde{U}}$. The repetition of the argument leads to $S \subseteq W$ and therefore $v \in \text{WinReg}(G, \mathcal{P})$

(\supseteq) Let $v \in \text{WinReg}(G, \mathcal{P})$. The algorithm finds a strongly connected set U (an SCC of a subgraph) that is reachable from v and that satisfies the winning condition. By moving from v into U , staying there, and visiting all positions in U infinitely often, Player 0 wins. Thus $v \in W_0$. Q.E.D.

Theorem 3.49. Deciding whether Player σ wins from a given position in a Streett-Rabin game is

- coNP-complete for $\sigma = 0$,
- NP-complete for $\sigma = 1$.

Proof. It suffices to prove the claim for Player 1 since W_0 is the complement of W_1 . Hardness follows from Theorem 3.47. To decide whether $v \in W_1$, guess a positional strategy for Player 1 and construct the induced solitaire game, in which only Player 0 has non-trivial moves. Then decide in polynomial time whether v is in the winning region of Player 1 in the solitaire game (according to Theorem 3.48), i.e. whether the strategy is winning from v . If this is the case, accept; otherwise reject. Q.E.D.

Remark 3.50. The complexity of computing the winning regions in arbitrary Muller games depends to a great amount on the representation of the winning condition. For any reasonable representation, the problem is in PSPACE, and many representations are so succinct as to render the problem PSPACE-hard. Only recently, it was shown that, given an explicit representation of the winning condition, the problem of deciding the winner is in P.