Logic and Games
WS 2018/2019

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

Mathematische Grundlagen der Informatik
RWTH Aachen

Contents

1 Reachability Games and First-Order Logic

1.1 Model Checking

One of the fundamental algorithmic tasks in logic is *model checking*. For a logic L and a domain $\mathcal{D}$ of (finite) structures, the *model-checking problem* asks, given a structure $\mathfrak{A} \in \mathcal{D}$ and a formula $\psi \in L$, whether $\mathfrak{A}$ is a model of ψ. Notice that an instance of the model-checking problem has two inputs: a structure and a formula. We can measure the complexity in terms of both inputs, and this is what is commonly refered to as the *combined complexity* of the model-checking problem (for L and $\mathcal{D}$). However, in many cases, one of the two inputs is fixed, and we measure the complexity only in terms of the other. If we fix the structure $\mathfrak{A}$, then the model-checking problem for L on this structure amounts to deciding $\mathrm{Th}_L(\mathfrak{A}) := \{\psi \in L : \mathfrak{A} \models \psi\}$, the *L-theory* of $\mathfrak{A}$. The complexity of this problem is called the *expression complexity* of the model-checking problem (for L on $\mathfrak{A}$). For first-order logic (FO) and for monadic second-order logic (MSO) in particular, such problems have a long tradition in logic and numerous applications in many fields. Of great importance in many areas of logic, in particular for finite model theory or databases, are model-checking problems for a fixed formula ψ, which amounts to deciding the *model class* of ψ inside $\mathcal{D}$, that is $\mathrm{Mod}_{\mathcal{D}}(\psi) := \{\mathfrak{A} \in \mathcal{D} : \mathfrak{A} \models \psi\}$. Its complexity is the *structure complexity* or *data complexity* of the model-checking problem (for ψ on $\mathcal{D}$).

One of the important themes in this course is a game-based approach to model checking. The general idea is to reduce the problem whether $\mathfrak{A} \models \psi$ to a strategy problem for a *model checking game* $\mathcal{G}(\mathfrak{A}, \psi)$ played by two players called *Verifier* (or *Player* 0) and *Falsifier* (or *Player* 1).

We want to have the following relation between these two problems:

$$\mathfrak{A} \models \psi \text{ iff } \text{ Verifier has a winning strategy for } \mathcal{G}(\mathfrak{A}, \psi).$$

We can then do model checking by constructing, or proving the existence of, winning strategies.

To assess the efficiency of games as a solution for model checking problems, we have to consider the complexity of the resulting model checking games based on the following criteria:

- Are all plays necessarily finite?
- If not, what are the winning conditions for infinite plays?
- Do the players always have perfect information?
- What is the structural complexity of the game graphs?
- How does the size of the graph depend on different parameters of the input structure and the formula?

For first-order logic (FO) and modal logic (ML) we have only finite plays with positional winning conditions, and, as we shall see, the winning regions are computable in linear time with respect to the size of the game graph. Model checking games for fixed-point logics however admit infinite plays, and we use so-called *parity conditions* to determine the winner of such plays. It is still an open question whether winning regions and winning strategies in parity games are computable in polynomial time.

1.2 Model Checking Games for Modal Logic

The first logic that we discuss is propositional modal logic (ML). Let us first briefly review its syntax and semantics:

Definition 1.1. Given a set A of actions and a set $\{P_i : i \in I\}$ of atomic propositions, the set of formulae of ML is inductively defined:

- All atomic propositions P_i are formulae of ML.
- If ψ, φ are formulae of ML, then so are $\neg\psi$, $(\psi \wedge \varphi)$ and $(\psi \vee \varphi)$.
- If $\psi \in$ ML and $a \in A$, then $\langle a \rangle \psi \in$ ML and $[a]\psi \in$ ML.

Remark 1.2. If there is only one action $a \in A$, we write $\Diamond\psi$ and $\Box\psi$ instead of $\langle a \rangle \psi$ and $[a]\psi$, respectively.

Definition 1.3. A *transition system* or *Kripke structure* with actions from a set A and atomic properties $\{P_i : i \in I\}$ is a structure

$$\mathcal{K} = (V, (E_a)_{a \in A}, (P_i)_{i \in I})$$

with a universe V of states, binary relations $E_a \subseteq V \times V$ describing transitions between the states, and unary relations $P_i \subseteq V$ describing the atomic properties of states.

A transition system can be seen as a labelled graph where the nodes are the states of $\mathcal{K}$, the unary relations provide labels of the states, and the binary transition relations can be pictured as sets of labelled edges.

Definition 1.4. Let $\mathcal{K} = (V, (E_a)_{a \in A}, (P_i)_{i \in I})$ be a transition system, $\psi \in \mathrm{ML}$ a formula and v a state of $\mathcal{K}$. The *model relationship* $\mathcal{K}, v \models \psi$, i.e., ψ holds at state v of $\mathcal{K}$, is inductively defined:

- $\mathcal{K}, v \models P_i$ if and only if $v \in P_i$.
- $\mathcal{K}, v \models \neg\psi$ if and only if $\mathcal{K}, v \not\models \psi$.
- $\mathcal{K}, v \models \psi \vee \varphi$ if and only if $\mathcal{K}, v \models \psi$ or $\mathcal{K}, v \models \varphi$.
- $\mathcal{K}, v \models \psi \wedge \varphi$ if and only if $\mathcal{K}, v \models \psi$ and $\mathcal{K}, v \models \varphi$.
- $\mathcal{K}, v \models \langle a \rangle \psi$ if and only if there exists w such that $(v, w) \in E_a$ and $\mathcal{K}, w \models \psi$.
- $\mathcal{K}, v \models [a]\, \psi$ if and only if $\mathcal{K}, w \models \psi$ holds for all w with $(v, w) \in E_a$.

For a transition system $\mathcal{K}$ and a formula ψ we define the *extension* $[\![\psi]\!]^{\mathcal{K}} := \{v : \mathcal{K}, v \models \psi\}$ as the set of states of $\mathcal{K}$ where ψ holds.

For the game-based approach to model-checking, it is convenient to assume that modal formulae are written in negation normal form, i.e. negation is applied to atomic propositions only. This does not reduce the expressiveness of modal logic since every formula can be efficiently translated into negation normal form by applying De Morgan's laws and the duality of $\Box$ and $\Diamond$ (i.e. $\neg\langle a \rangle \psi \equiv [a]\neg\psi$ and $\neg[a]\psi \equiv \langle a \rangle \neg\psi$) to push negations to the atomic subformulae.

Syntactically, modal logic is an extension of propositional logic. However, since ML is evaluated over transition systems, i.e. structures, it is often useful to see it as a fragment of first-order logic.

Theorem 1.5. For each formula $\psi \in$ ML there is a first-order formula $\psi^*(x)$ (with only two variables), such that for each transition system $\mathcal{K}$ and all its states v we have that $\mathcal{K}, v \models \psi \iff \mathcal{K} \models \psi^*(v)$.

Proof. The transformation is defined inductively, as follows:

$$\begin{aligned}
P_i &\longmapsto P_i x \\
\neg\psi &\longmapsto \neg\psi^*(x) \\
(\psi \circ \varphi) &\longmapsto (\psi^*(x) \circ \varphi^*(x)), \text{ where } \circ \in \{\wedge, \vee, \rightarrow\} \\
\langle a \rangle \psi &\longmapsto \exists y (E_a xy \wedge \psi^*(y)) \\
[a]\psi &\longmapsto \forall y (E_a xy \rightarrow \psi^*(y))
\end{aligned}$$

where $\psi^*(y)$ is obtained from $\psi^*(x)$ by interchanging x and y everywhere in the formula. Q.E.D.

We are now ready to describe the model checking games for ML. Given a transition system $\mathcal{K}$ and a formula $\psi \in$ ML, we define a game $\mathcal{G}(\mathcal{K}, \psi)$ whose positions are pairs (φ, v) where φ is a subformula of ψ and $v \in V$ is a node of $\mathcal{K}$. From any position (φ, v) in this game, Verifier's goal is to show that $\mathcal{K}, v \models \varphi$, whereas Falsifier tries to establish that $\mathcal{K}, v \not\models \varphi$.

In the game, Verifier moves at positions of the form $(\varphi \vee \vartheta, v)$, with the choice to move either to (φ, v) or to (ϑ, v), and at positions $(\langle a \rangle \varphi, v)$, where she can move to any position (φ, w) with $w \in vE_a$. Analogously, Falsifier moves from positions $(\varphi \wedge \vartheta, v)$ to either (φ, v) or (ϑ, v), and from $([a]\varphi, v)$ to any position (φ, w) with $w \in vE_a$. Finally, at literals, i.e. if $\varphi = P_i$ or $\varphi = \neg P_i$, the position (φ, v) is a terminal position where Verifier has won if $\mathcal{K}, v \models \varphi$, and Falsifier has won if $\mathcal{K}, v \not\models \varphi$.

The correctness of the construction of $\mathcal{G}(\mathcal{K}, \psi)$ follows readily by induction.

Proposition 1.6. For any position (φ, v) of $\mathcal{G}(\mathcal{K}, \psi)$ we have that

$$\mathcal{K}, v \models \varphi \quad \Leftrightarrow \quad \text{Verifier has a winning strategy for } \mathcal{G}(\mathcal{K}, \psi) \text{ from } (\varphi, v).$$

1.3 Reachability and Safety Games

The model-checking games for propositional modal logic, that we have just discussed, are an instance of reachability games played on graphs or, more precisely, two-player games with perfect information and positional winning conditions, played on a *game graph* (or *arena*)

$$\mathcal{G} = (V, V_0, V_1, E)$$

where the set V of positions is partitioned into sets of positions V_0 and V_1 belonging to Player 0 and Player 1, respectively. Player 0 moves from positions $v \in V_0$, while Player 1 moves from positions $v \in V_1$. All moves are along edges, and so the interaction of the players, starting from an initial position v_0, produces a finite or infinite *play* which is a sequence $v_0 v_1 v_2 \dots$ with $(v_i, v_{i+1}) \in E$ for all i.

The winning conditions of the players are based on a simple positional principle: Move or lose! This means that Player σ has won at a position v in the case that position v belongs to his opponent and there are no moves available from that position. Thus the goal of Player σ is to reach a position in $T_\sigma := \{v \in V_{1-\sigma} : vE = \emptyset\}$. We call this a *reachability condition*.

But note that this winning condition applies to finite plays only. If the game graph admits infinite plays (for instance cycles) then we must either consider these as draws, or introduce a winning condition for infinite plays. The dual notion of a reachability condition is a *safety condition* where Player σ just has the objective to avoid a given set of 'bad' positions, which in this case is the set $T_{1-\sigma}$, and to remain inside the safe region $V \setminus T_{1-\sigma}$.

A *(positional) strategy* for Player σ in such a game $\mathcal{G}$ is a (partial) function $f : \{v \in V_\sigma : vE \neq \emptyset\} \to V$ such that $(v, f(v)) \in E$. A finite or infinite play $v_0 v_1 v_2 \dots$ is *consistent with f* if $v_{i+1} = f(v_i)$ for every i

such that $v_i \in V_\sigma$. A strategy f for Player σ is winning from v_0 if every play that starts at initial position v_0 and that is consistent with f is won by Player σ.

We first consider *reachability games* where both players play with the reachability objective to force the play to a position in T_σ. We define *winning regions*

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

If $W_0 \cup W_1 = V$, i.e. for each $v \in V$ one of the players has a winning strategy, the game $\mathcal{G}$ is called *determined*. A play which is not won by any of the players is considered a draw.

Example 1.7. No player can win from one of the middle two nodes:

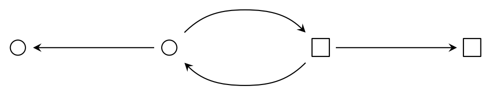

The winning regions of a reachability game $\mathcal{G} = (V, V_0, V_1, E)$ can be constructed inductively as follows:

$$\begin{aligned} W_\sigma^0 &= T_\sigma \text{ and} \\ W_\sigma^{i+1} &= W_\sigma^i \cup \{v \in V_\sigma : vE \cap W_\sigma^i \neq \emptyset\} \cup \{v \in V_{1-\sigma} : vE \subseteq W_\sigma^i\}. \end{aligned}$$

Clearly W_σ^i is the region of those positions from which Player σ has a strategy to win in at most i moves, and for finite game graphs, with $|V| = n$, we have that $W_\sigma = W_\sigma^n$.

Next we consider the case of a *reachability-safety game*, where Player 0, as above, plays with the reachability objective to force the play to a terminal position in T_0, whereas player 1 plays with the safety objective of avoiding T_0, i.e. to keep the play inside the safe region $S_1 := V \setminus T_0$. Notice that there are no draws in such a game.

The winning region W_0 of Player 0 can be defined as in the case above, but the winning region W_1 of Player 1 is now the maximal set $W \subseteq S_1$ such that from all $w \in W$ Player 1 has a strategy to remain inside W, which can be defined as the limit of the descending chain $W_1^0 \supseteq W_1^1 \supseteq W_1^2 \supseteq \dots$ with

$$\begin{aligned} W_1^0 &= S_1 \text{ and} \\ W_1^{i+1} &= W_1^i \cap \{v \in V : (v \in V_0 \text{ and } vE \subseteq W_1^i) \text{ or} \\ &\qquad\qquad (v \in V_1 \text{ and } vE \cap W_1^i \neq \emptyset)\}. \end{aligned}$$

Again on finite game graphs, with $|V| = n$, we have that $W_1 = W_\sigma^n$.

This leads us to two fundamental concepts for the analysis of games on graphs: *attractors* and *traps*. Let $\mathcal{G} = (V, V_0, V_1, E)$ be a game graph and $X \subseteq V$.

Definition 1.8. The *attractor of X for Player σ*, in short $\mathrm{Attr}_\sigma(X)$ is the set of those positions from which Player σ has a strategy to reach X (or to win because the opponent cannot move anymore). We can inductively define $\mathrm{Attr}_\sigma(X) := \bigcup_{n \in \mathbb{N}} X^n$, where

$$\begin{aligned} X^0 &= X \text{ and} \\ X^{i+1} &= X^i \cup \{v \in V_\sigma : vE \cap X^i \neq \emptyset\} \cup \{v \in V_{1-\sigma} : vE \subseteq X^i\}. \end{aligned}$$

For instance, the winning region W_σ in a reachability game is the attractor of the winning positions: $W_\sigma = \mathrm{Attr}_\sigma(T_\sigma)$.

A set $Y \subseteq V \setminus T_{1-\sigma} =: S_\sigma$ is called a *trap* for Player $1 - \sigma$ if Player σ has a strategy to guarantee that from each $v \in Y$ the play will remain inside Y. Note that the complement of an attractor $\mathrm{Attr}_\sigma(X)$ is a trap for player σ. The maximal trap Y of Player $1 - \sigma$ can be defined as $Y = \bigcap_{n \in \mathbb{N}} Y^n$, where

$$\begin{aligned} Y^0 =& S_\sigma \text{ and} \\ Y^{i+1} =& Y^i \cap \{v : (v \in V_\sigma \text{ and } vE \cap Y^i \neq \emptyset) \text{ or} \\ &\qquad (v \in V_{1-\sigma} \text{ and } vE \subseteq Y^i)\}. \end{aligned}$$

The winning region of a Player σ with the safety objective for S_σ is the maximal trap for player $1 - \sigma$.

We consider several algorithmic problems for a given reachability game $\mathcal{G}$: The computation of winning regions W_0 and W_1, the computation of winning strategies, and the associated decision problem

$$\text{Game} := \{(\mathcal{G}, v) : \text{Player 0 has a winning strategy for } \mathcal{G} \text{ from } v\}.$$

Theorem 1.9. Game is P-complete and decidable in time $O(|V| + |E|)$.

Note that this remains true for *strictly alternating games*.

The inductive definition of an attractor shows that winning regions for both players can be computed efficiently. Hence we can also solve Game in polynomial time. To solve Game in linear time, we use the slightly more involved Algorithm 1.1. Procedure Propagate will be called once for every edge in the game graph, so the running time of this algorithm is linear with respect to the number of edges in $\mathcal{G}$.

Furthermore, we can show that the decision problem Game is equivalent to the satisfiability problem for propositional Horn formulae. We recall that propositional Horn formulae are finite conjunctions $\bigwedge_{i \in I} C_i$ of clauses C_i of the form

$$\begin{aligned}
&X_1 \wedge \ldots \wedge X_n \;\rightarrow\; X \quad \text{or} \\
&\underbrace{X_1 \wedge \ldots \wedge X_n}_{\text{body}(C_i)} \;\rightarrow\; \underbrace{0}_{\text{head}(C_i)} \;.
\end{aligned}$$

A clause of the form X or $1 \rightarrow X$ has an empty body.

We will show that Sat-Horn and Game are mutually reducible via logspace and linear-time reductions.

(1) Game $\leq_{\text{log-lin}}$ Sat-Horn
For a game $\mathcal{G} = (V, V_0, V_1, E)$, we construct a Horn formula $\psi_{\mathcal{G}}$ with clauses

$$\begin{aligned}
v \rightarrow u \qquad & \text{for all } u \in V_0 \text{ and } (u, v) \in E, \text{ and} \\
v_1 \wedge \ldots \wedge v_m \rightarrow u \qquad & \text{for all } u \in V_1 \text{ and } uE = \{v_1, \ldots, v_m\}.
\end{aligned}$$

The minimal model of $\psi_{\mathcal{G}}$ is precisely the winning region of Player 0, so

$$(\mathcal{G}, v) \in \text{Game} \iff \psi_{\mathcal{G}} \wedge (v \rightarrow 0) \text{ is unsatisfiable.}$$

Algorithm 1.1. A linear time algorithm for GAME

Input: A game $\mathcal{G} = (V, V_0, V_1, E)$
output: Winning regions W_0 and W_1

for all $v \in V$ **do** (∗ 1: Initialisation ∗)
    $\text{win}[v] := \bot$
    $P[v] := \{u : (u, v) \in E\}$
    $n[v] := |vE|$
end do

for all $v \in V_0$ (∗ 2: Calculate win ∗)
    if $n[v] = 0$ **then** Propagate(v, 1)
for all $v \in V_1$
    if $n[v] = 0$ **then** Propagate(v, 0)
return win

procedure Propagate(v, σ)
    if $\text{win}[v] \neq \bot$ **then return**
    $\text{win}[v] := \sigma$ (∗ 3: Mark v as winning for player σ ∗)
    for all $u \in P[v]$ **do** (∗ 4: Propagate change to predecessors ∗)
        $n[u] := n[u] - 1$
        if $u \in V_\sigma$ or $n[u] = 0$ **then** Propagate(u, σ)
    end do
end

(2) SAT-HORN $\leq_{\text{log-lin}}$ GAME
For a Horn formula $\psi(X_1, \dots, X_n) = \bigwedge_{i \in I} C_i$, we define a game $\mathcal{G}_\psi = (V, V_0, V_1, E)$ as follows:

$$V = \underbrace{\{0\} \cup \{X_1, \dots, X_n\}}_{V_0} \cup \underbrace{\{C_i : i \in I\}}_{V_1} \text{ and}$$

$$E = \{X \to C_i : X = \text{head}(C_i)\} \cup \{C_i \to X_j : X_j \in \text{body}(C_i)\},$$

i.e., Player 0 moves from a variable to some clause containing the variable as its head, and Player 1 moves from a clause to some variable in its body. Player 0 wins a play if, and only if, the play reaches a clause C with $\text{body}(C) = \emptyset$. Furthermore, Player 0 has a winning strategy from position X if, and only if, $\psi \models X$, so we have

$$\text{Player 0 wins from position 0} \quad \Longleftrightarrow \quad \psi \text{ is unsatisfiable.}$$

These reductions show that SAT-HORN is also P-complete and, in particular, also decidable in linear time.

1.4 Games as an Algorithmic Construct: Alternating Algorithms

Alternating algorithms are algorithms whose set of configurations is divided into *accepting*, *rejecting*, *existential* and *universal* configurations. The acceptance condition of an alternating algorithm A is defined by a game played by two players $\exists$ and $\forall$ on the computation graph $\mathcal{G}(A, x)$ (or equivalently, the computation tree $\mathcal{T}(A, x)$) of A on input x. The positions in this game are the configurations of A, and we allow moves $C \to C'$ from a configuration C to any of its successor configurations C'. Player $\exists$ moves at existential configurations and wins at accepting configurations, while Player $\forall$ moves at universal configurations and wins at rejecting configurations. By definition, A accepts some input x if and only if Player $\exists$ has a winning strategy for the game played on $\mathcal{T}_{A,x}$.

We will introduce the concept of alternating algorithms formally, using the model of a Turing machine, and we prove certain relationships between the resulting alternating complexity classes and usual deterministic complexity classes.

1.4.1 Turing Machines

The notion of an alternating Turing machine extends the usual model of a (deterministic) Turing machine which we introduce first. We consider Turing machines with a separate input tape and multiple linear work tapes which are divided into basic units, called cells or fields. Informally, the Turing machine has a reading head on the input tape and a combined reading and writing head on each of its work tapes. Each of the heads is at one particular cell of the corresponding tape during each point of a computation. Moreover, the Turing machine is in a certain state. Depending on this state and the symbols the machine is currently reading on the input and work tapes, it manipulates the current fields of the work tapes, moves its heads and changes to a new state.

Formally, a *(deterministic) Turing machine with separate input tape and k linear work tapes* is given by a tuple $M = (Q, \Gamma, \Sigma, q_0, F_{\text{acc}}, F_{\text{rej}}, \delta)$, where Q is a finite set of *states*, Σ is the *work alphabet* containing a designated symbol $\square$ (*blank*), Γ is the *input alphabet*, $q_0 \in Q$ is the *initial state*, $F := F_{\text{acc}} \cup F_{\text{rej}} \subseteq Q$ is the set of *final states* (with F_{acc} the *accepting states*, F_{rej} the *rejecting states* and $F_{\text{acc}} \cap F_{\text{rej}} = \emptyset$), and $\delta : (Q \setminus F) \times \Gamma \times \Sigma^k \to Q \times \{-1, 0, 1\} \times \Sigma^k \times \{-1, 0, 1\}^k$ is the *transition function*.

A *configuration* of M is a complete description of all relevant facts about the machine at some point during a computation, so it is a tuple $C = (q, w_1, \ldots, w_k, x, p_0, p_1, \ldots, p_k) \in Q \times (\Sigma^*)^k \times \Gamma^* \times \mathbb{N}^{k+1}$ where q is the current state, w_i is the contents of work tape number i, x is the contents of the input tape, p_0 is the position on the input tape and p_i is the position on work tape number i. The contents of each of the tapes is represented as a finite word over the corresponding alphabet[, i.e., a finite sequence of symbols from the alphabet]. The contents of each of the fields with numbers $j > |w_i|$ on work tape number i is the blank symbol (we think of the tape as being infinite). A configuration where x

is omitted is called a *partial configuration*. The configuration C is called *final* if $q \in F$. It is called *accepting* if $q \in F_{\text{acc}}$ and *rejecting* if $q \in F_{\text{rej}}$.

The *successor configuration* of C is determined by the current state and the $k+1$ symbols on the current cells of the tapes, using the transition function: If $\delta(q, x_{p_0}, (w_1)_{p_1}, \ldots, (w_k)_{p_k}) = (q', m_0, a_1, \ldots, a_k, m_1, \ldots, m_k, b)$, then the successor configuration of C is $\Delta(C) = (q', \overline{w}', \overline{p}', x)$, where for any i, w_i' is obtained from w_i by replacing symbol number p_i by a_i and $p_i' = p_i + m_i$. We write $C \vdash_M C'$ if, and only if, $C' = \Delta(C)$.

The *initial configuration* $C_0(x) = C_0(M, x)$ of M on input $x \in \Gamma^*$ is given by the initial state q_0, the blank-padded memory, i.e., $w_i = \varepsilon$ and $p_i = 0$ for any $i \geq 1$, $p_0 = 0$, and the contents x on the input tape.

A *computation* of M on input x is a sequence $C_0, C_1, \ldots$ of configurations of M, such that $C_0 = C_0(x)$ and $C_i \vdash_M C_{i+1}$ for all $i \geq 0$. The computation is called *complete* if it is infinite or ends in some final configuration. A complete finite computation is called *accepting* if the last configuration is accepting, and the computation is called *rejecting* if the last configuration is rejecting. M *accepts* input x if the (unique) complete computation of M on x is finite and accepting. M *rejects* input x if the (unique) complete computation of M on x is finite and rejecting. The machine M *decides* a language $L \subseteq \Gamma^*$ if M accepts all $x \in L$ and rejects all $x \in \Gamma^* \setminus L$.

1.4.2 Alternating Turing Machines

Now we shall extend deterministic Turing machines to nondeterministic Turing machines from which the concept of alternating Turing machines is obtained in a very natural way, given our game theoretical framework.

A *nondeterministic Turing machine* is nondeterministic in the sense that a given configuration C may have several possible successor configurations instead of at most one. Intuitively, this can be described as the ability to *guess*. This is formalised by replacing the transition function $\delta : (Q \setminus F) \times \Gamma \times \Sigma^k \to Q \times \{-1,0,1\} \times \Sigma^k \times \{-1,0,1\}^k$ by a transition relation $\Delta \subseteq ((Q \setminus F) \times \Gamma \times \Sigma^k) \times (Q \times \{-1,0,1\} \times \Sigma^k \times \{-1,0,1\}^k)$. The notion of successor configurations is defined as in the deterministic case, except that the successor configuration of a configuration C may not be uniquely determined. Computations and all related notions carry over from deterministic machines in the obvious way. However, on a fixed input x, a nondeterministic machine now has several possible computations, which form a (possibly infinite) finitely branching computation tree $\mathcal{T}_{M,x}$. A nondeterministic Turing machine M *accepts* an input x if there *exists* a computation of M on x which is accepting, i.e., if there exists a path from the root $C_0(x)$ of $\mathcal{T}_{M,x}$ to some accepting configuration. The language of M is $L(M) = \{x \in \Gamma^* \mid M \text{ accepts } x\}$. Notice that for a nondeterministic machine M to decide a language $L \subseteq \Gamma^*$ it is not necessary that all computations of M are finite. (In a sense, we count infinite computations as rejecting.)

From a game-theoretical perspective, the computation of a nondeterministic machine can be viewed as a solitaire game on the computation tree in which the only player (the machine) chooses a path through the tree starting from the initial configuration. The player wins the game (and hence, the machine accepts its input) if the chosen path finally reaches an accepting configuration.

An obvious generalisation of this game is to turn it into a two-player game by assigning the nodes to the two players who are called $\exists$ and $\forall$, following the intuition that Player $\exists$ tries to show the existence of a *good* path, whereas Player $\forall$ tries to show that all selected paths are *bad*. As before, Player $\exists$ wins a play of the resulting game if, and only if, the play is finite and ends in an accepting leaf of the game tree. Hence, we call a computation tree accepting if, and only if, Player $\exists$ has a winning strategy for this game.

It is important to note that the partition of the nodes in the tree should not depend on the input x but is supposed to be inherent to the machine. Actually, it is even independent of the contents of the work tapes, and thus, whether a configuration belongs to Player $\exists$ or to Player $\forall$ merely depends on the current state.

Formally, an *alternating Turing machine* is a nondeterministic Turing machine $M = (Q, \Gamma, \Sigma, q_0, F_{\text{acc}}, F_{\text{rej}}, \Delta)$ whose set of states $Q = Q_\exists \cup Q_\forall \cup F_{\text{acc}} \cup F_{\text{rej}}$ is partitioned into *existential*, *universal*, *accepting*, and *rejecting* states. The semantics of these machines is given by means of the game described above.

Now, if we let accepting configurations belong to player $\forall$ and rejecting configurations belong to player $\exists$, then we have the usual winning condition that a player loses if it is his turn but he cannot move. We can solve such games by determining the winner at leaf nodes and propagating the winner successively to parent nodes. If at some node, the winner at all of its child nodes is determined, the winner at this node can be determined as well. This method is sometimes referred to as backwards induction and it basically coincides with our method for solving GAME on trees (with possibly infinite plays). This gives the following equivalent semantics of alternating Turing machines:

The subtree $\mathcal{T}_C$ of the computation tree of M on x with root C is called *accepting*, if

- C is accepting
- C is existential and there is a successor configuration C' of C such that $\mathcal{T}_{C'}$ is accepting or
- C is universal and $\mathcal{T}_{C'}$ is accepting for all successor configurations C' of C.

M accepts an input x, if $\mathcal{T}_{C_0(x)} = \mathcal{T}_{M,x}$ is accepting.

For functions $T, S : \mathbb{N} \to \mathbb{N}$, an alternating Turing machine M is called *T-time bounded* if, and only if, for any input x, each computation of M on x has length less or equal $T(|x|)$. The machine is called *S-space bounded* if, and only if, for any input x, during any computation of M on x, at most $S(|x|)$ cells of the work tapes are used. Notice that time boundedness implies finiteness of all computations which is not the case for space boundedness. The same definitions apply for deterministic and nondeterministic Turing machines as well since these are just special cases of alternating Turing machines. These notions of resource bounds induce the complexity classes ATIME containing precisely those languages L such that there is an alternating T-time bounded Turing machine deciding L and ASPACE containing precisely those languages L such that there is an alternating S-space bounded

Turing machine deciding L. Similarly, these classes can be defined for nondeterministic and deterministic Turing machines.

We are especially interested in the following alternating complexity classes:

- ALOGSPACE $= \bigcup_{d\in\mathbb{N}} \text{ASPACE}(d \cdot \log n)$,
- APTIME $= \bigcup_{d\in\mathbb{N}} \text{ATIME}(n^d)$,
- APSPACE $= \bigcup_{d\in\mathbb{N}} \text{ASPACE}(n^d)$.

Observe that GAME $\in$ ALOGSPACE. An alternating algorithm which decides GAME with logarithmic space just plays the game. The algorithm only has to store the *current* position in memory, and this can be done with logarithmic space. We shall now consider a slightly more involved example.

Example 1.10. QBF $\in$ ATIME$(\mathrm{O}(n))$. W.l.o.g we assume that negation appears only at literals. We describe an alternating procedure $Eval(\varphi, \mathcal{I})$ which computes, given a quantified Boolean formula ψ and a valuation $\mathcal{I} : \text{free}(\psi) \to \{0,1\}$ of the free variables of ψ, the value $[\![\psi]\!]^{\mathcal{I}}$.

Algorithm 1.2. Alternating algorithm deciding QBF.

Input: $(\psi, \mathcal{I})$ where $\psi \in$ QAL and $\mathcal{I}\colon \text{free}(\psi) \to \{0,1\}$
if $\psi = Y$ **then**
  if $\mathcal{I}(Y) = 1$ **then** accept
    else reject
if $\psi = \varphi_1 \vee \varphi_2$ **then** „$\exists$“ guesses $i \in \{1,2\}$, $Eval(\varphi_i, \mathcal{I})$
if $\psi = \varphi_1 \wedge \varphi_2$ **then** „$\forall$“ chooses $i \in \{1,2\}$, $Eval(\varphi_i, \mathcal{I})$
if $\psi = \exists X \varphi$ **then** „$\exists$“ guesses $j \in \{0,1\}$, $Eval(\varphi, \mathcal{I}[X = j])$
if $\psi = \forall X \varphi$ **then** „$\forall$“ chooses $j \in \{0,1\}$, $Eval(\varphi, \mathcal{I}[X = j])$

1.4.3 Alternating versus Deterministic Complexity Classes

The main results we want to establish in this section concern the relationship between alternating complexity classes and deterministic complexity classes. We will see that alternating time corresponds to deterministic space, while by translating deterministic time into alternating space, we can reduce the complexity by one exponential. Here, we consider the special case of alternating polynomial time and polynomial space. We should mention, however, that these results can be generalised to arbitrary large complexity bounds which are well behaved in a certain sense.

Lemma 1.11. NPSPACE $\subseteq$ APTIME.

Proof. Let $L \in$ NPSPACE and let M be a nondeterministic n^l-space bounded Turing machine which recognises L for some $l \in \mathbb{N}$. The machine M accepts some input x if, and only if, some accepting configuration is reachable from the initial configuration $C_0(x)$ in the configuration tree of M on x in at most $k := 2^{cn^l}$ steps for some $c \in \mathbb{N}$. This is due to the fact that there are most k different configurations of M on input x which use at most n^l cells of the memory which can be seen using a simple combinatorial argument. So if there is some accepting configuration reachable from the initial configuration $C_0(x)$, then there is some accepting configuration reachable from $C_0(x)$ in at most k steps. This is equivalent to the existence of some intermediate configuration C' that is reachable from $C_0(x)$ in at most $k/2$ steps and from which some accepting configuration is reachable in at most $k/2$ steps.

So the alternating algorithm deciding L proceeds as follows. The existential player guesses such a configuration C' and the universal player chooses whether to check that C' is reachable from $C_0(x)$ in at most $k/2$ steps or whether to check that some accepting configuration is reachable from C' in at most $k/2$ steps. Then the algorithm (or equivalently, the game) proceeds with the subproblem chosen by the universal player, and continues in this binary search like fashion. Obviously, the number of steps which have to be performed by this procedure to decide whether x is accepted by M is logarithmic in k. Since k is exponential in n^l, the time bound of M is dn^l for some $d \in \mathbb{N}$, so M decides L in polynomial time. Q.E.D.

Lemma 1.12. APTIME $\subseteq$ PSPACE.

Proof. Let $L \in$ APTIME and let A be an alternating n^l-time bounded Turing machine that decides L for some $l \in \mathbb{N}$. Then there is some $r \in \mathbb{N}$ such that any configuration of A on any input x has at most r successor configurations and w.l.o.g. we can assume that any non-final configuration has precisely r successor configurations. We can think of the successor configurations of some non-final configuration C as being enumerated as $C_1, \ldots, C_r$. Clearly, for given C and i we can compute C_i. The idea for a deterministic Turing machine M to check whether some input x is in L is to perform a depth-first search on the computation tree $\mathcal{T}_{A,x}$ of A on x. The crucial point is that we cannot construct and keep the whole configuration tree $\mathcal{T}_{A,x}$ in memory since its size is exponential in $|x|$ which exceeds our desired space bound. However, since the length of each computation is polynomially bounded, it is possible to keep a single computation path in memory and to construct the successor configurations of the configuration under consideration on the fly.

Roughly, the procedure M can be described as follows. We start with the initial configuration $C_0(x)$. Given any configuration C under consideration, we propagate 0 to the predecessor configuration if C is rejecting and we propagate 1 to the predecessor configuration if C is accepting. If C is neither accepting nor rejecting, then we construct, for $i = 1, \ldots, r$ the successor configuration C_i of C and proceed with checking C_i. If C is existential, then as soon as we receive 1 for some i, we propagate 1 to the predecessor. If we encounter 0 for all i, then we propagate 0. Analogously, if C is universal, then as soon as we receive a 0 for some i, we propagate 0. If we receive only 1 for all i, then we propagate 1. Then x is in L if, and only if, we finally receive 1 at $C_0(x)$. Now, at any point during such a computation we have to store at most one complete computation of A on x. Since A is n^l-time bounded, each such computation has length at most n^l and each configuration has size at most $c \cdot n^l$ for some $c \in \mathbb{N}$. So M needs at most $c \cdot n^{2l}$ memory cells which is polynomial in n. Q.E.D.

So we obtain the following result.

Theorem 1.13. (Parallel time complexity = sequential space complexity)

(1) APTIME = PSPACE.
(2) AEXPTIME = EXPSPACE.

Proposition (2) of this theorem is proved exactly the same way as we have done it for proposition (1). Now we prove that by translating sequential *time* into alternating *space*, we can reduce the complexity by one exponential.

Lemma 1.14. EXPTIME $\subseteq$ APSPACE

Proof. Let $L \in$ EXPTIME. Using a standard argument from complexity theory, there is a deterministic Turing machine $M = (Q, \Sigma, q_0, \delta)$ with time bound $m := 2^{c \cdot n^k}$ for some $c, k \in \mathbb{N}$ with only a single tape (serving as both input and work tape) which decides L. (The time bound of the machine with only a single tape is quadratic in that of the original machine with k work tapes and a separate input tape, which, however, does not matter in the case of an exponential time bound.) Now if $\Gamma = \Sigma \uplus (Q \times \Sigma) \uplus \{\#\}$, then we can describe each configuration C of M by a word

$$\underline{C} = \# w_0 \dots w_{i-1} (q w_i) w_{i+1} \dots w_t \# \in \Gamma^*.$$

Since M has time bound m and only one single tape, it has space bound m. So, w.l.o.g., we can assume that $|\underline{C}| = m + 2$ for all configurations C of M on inputs of length n. (We just use a representation of the tape which has a priori the maximum length that will occur during a computation on an input of length n.) Now the crucial point in the argumentation is the following. If $C \vdash C'$ and $1 \leq i \leq m$, symbol number i of the word $\underline{C'}$ only depends on the symbols number $i - 1$, i and $i + 1$ of $\underline{C}$. This allows us to decide whether $x \in L(M)$ with the following alternating procedure which uses only polynomial space.

Player $\exists$ guesses some number $s \leq m$ of steps of which he claims that it is precisely the length of the computation of M on input x. Furthermore, $\exists$ guesses some state $q \in F_{\text{acc}}$, a Symbol $a \in \Sigma$ and a number $i \in \{0, \dots, s\}$, and he claims that the i-th symbol of the configuration $\underline{C}$ of M after the computation on x is (qa). (So players start inspecting the computation of M on x from the final configuration.)

If M accepts input x, then obviously player $\exists$ has a possibility to choose all these objects such that his claims can be validated. Player $\forall$ wants to disprove the claims of $\exists$. Now, player $\exists$ guesses symbols $a_{-1}, a_0, a_1 \in \Gamma$ of which he claims that these are the symbols number $i-1$, i and $i+1$ of the predecessor configuration of the final configuration $\underline{C}$. Now, $\forall$ can choose any of these symbols and demand that $\exists$ validates his claim for this particular symbol. This symbol is now the symbol under consideration, while i is updated according to the movement of the (unique) head of M. Now, these actions of the players take place for each of the s computation steps of M on x. After s such steps, we check whether the current symbol and the current position are consistent with the initial configuration $C_0(x)$. The only information that has to be stored in the memory is the position i on the tape, the number s which $\exists$ has initially guessed and the current number of steps. Therefore, the algorithm uses space at most $\mathrm{O}(\log(m)) = \mathrm{O}(n^k)$ which is polynomial in n. Moreover, if M accepts input x then obviously player $\exists$ has a winning strategy for the computation game. If, conversely, M rejects input x, then the combination of all claims of player $\exists$ cannot be consistent and player $\forall$ has a strategy to spoil any (cheating) strategy of player $\exists$ by choosing the appropriate symbol at the appropriate computation step. Q.E.D.

Finally, we make the simple observation that it is not possible to gain more than one exponential when translating from sequential time to alternating space. (Notice that Exptime is a proper subclass of 2Exptime.)

Lemma 1.15. APspace $\subseteq$ Exptime

Proof. Let $L \in$ APspace, and let A be an alternating n^k-space bounded Turing machine which decides L for some $k \in \mathbb{N}$. Moreover, for an input x of A, let $\mathrm{Conf}(A, x)$ be the set of all configurations of A on input x. Due to the polynomial space bound of A, this set is finite and its size is at most exponential in $|x|$. So we can construct the graph $G = (\mathrm{Conf}(A, x), \vdash)$ in time exponential in $|x|$. Moreover, a configuration C is reachable from $C_0(x)$ in $\mathcal{T}_{A,x}$ if and only if C is reachable from $C_0(x)$ in G. So to check whether A accepts input x we simply decide whether player $\exists$ has a winning strategy for the game played on G from $C_0(x)$. This can be done in time linear in the size of G, so altogether we can decide whether $x \in L(A)$ in time exponential in $|x|$. Q.E.D.

Theorem 1.16. (Translating sequential time into alternating space)

(1) ALogspace = P.
(2) APspace = Exptime.

Proposition (1) of this theorem is proved using exactly the same arguments as we have used for proving proposition (2). An overview over the relationship between deterministic and alternating complexity classes is given in Figure 1.1.

$$\begin{array}{ccccccccc} \textsc{Logspace} & \subseteq & \textsc{Ptime} & \subseteq & \textsc{Pspace} & \subseteq & \textsc{Exptime} & \subseteq & \textsc{Expspace} \\ & & \| & & \| & & \| & & \| \\ & & \textsc{ALogspace} & \subseteq & \textsc{APtime} & \subseteq & \textsc{APspace} & \subseteq & \textsc{AExptime} \end{array}$$

Figure 1.1. Relation between deterministic and alternating complexity classes

1.5 Model Checking Games for First-Order Logic

Let us first recall the syntax of FO formulae on relational structures. We have that $R_i(\bar{x})$, $\neg R_i(\bar{x})$, $x = y$ and $x \neq y$ are well-formed valid FO formulae, and inductively for FO formulae φ and ψ, we have that $\varphi \vee \psi$, $\varphi \wedge \psi$, $\exists x \varphi$ and $\forall x \varphi$ are well-formed FO formulae. This way, we allow only formulae in *negation normal form* where negations occur only at atomic subformulae and all junctions except $\vee$ and $\wedge$ are eliminated. These constraints do not limit the expressiveness of the logic, but the resulting games are easier to handle.

For a structure $\mathfrak{A} = (A, R_1, \ldots, R_m)$ with $R_i \subseteq A^{r_i}$, we define the evaluation game $\mathcal{G}(\mathfrak{A}, \psi)$ as follows:

We have positions $\varphi(\bar{a})$ for every subformula $\varphi(\bar{x})$ of ψ and every $\bar{a} \in A^k$.

At a position $\varphi \vee \vartheta$, Verifier can choose to move either to φ or to ϑ, while at positions $\exists x \varphi(x, \bar{b})$, he can choose an instantiation $a \in A$ of x and move to $\varphi(a, \bar{b})$. Analogously, Falsifier can move from positions $\varphi \wedge \vartheta$ to either φ or ϑ and from positions $\forall x \varphi(x, \bar{b})$ to $\varphi(a, \bar{b})$ for an $a \in A$.

The winning condition is evaluated at positions with atomic or negated atomic formulae φ, and we define that Verifier wins at $\varphi(\bar{a})$ if, and only if, $\mathfrak{A} \models \varphi(\bar{a})$, and Falsifier wins if, and only if, $\mathfrak{A} \not\models \varphi(\bar{a})$.

In order to determine the complexity of FO model checking, we have to consider the process of determining whether $\mathfrak{A} \models \psi$. To decide this question, we have to construct the game $\mathcal{G}(\mathfrak{A}, \psi)$ and check whether Verifier has a winning strategy from position ψ. The size of the game graph is bound by $|\mathcal{G}(\mathfrak{A}, \psi)| \leq |\psi| \cdot |A|^{\mathrm{width}(\psi)}$, where $\mathrm{width}(\psi)$ is the maximal number of free variables in the subformulae of ψ. So the game graph can be exponential, and therefore we can get only exponential time complexity for Game. In particular, we have the following complexities for the general case:

- alternating time: $O(|\psi| + \mathrm{qd}(\psi) \log |A|)$
  where $\mathrm{qd}(\psi)$ is the quantifier-depth of ψ,
- alternating space: $O(\mathrm{width}(\psi) \cdot \log |A| + \log |\psi|)$,
- deterministic time: $O(|\psi| \cdot |A|^{\mathrm{width}(\psi)})$ and
- deterministic space: $O(|\psi| + \mathrm{qd}(\psi) \log |A|)$.

Efficient implementations of model checking algorithms will construct the game graph on the fly while solving the game.

We obtain that the structural complexity of FO model checking is ALogtime, and both the expression complexity and the combined complexity are PSpace.

Fragments of FO with Efficient Model Checking

We have seen that the size of the model checking games for first-order fomulae is exponential with respect to the width of the formulae, so we do not obtain polynomial-time model-checking algorithms in the general case. We now consider appropriate restrictions of FO, that lead to fragments with small model-checking games and thus to efficient game-based model-checking algorithms.

The *k-variable fragment of* FO is

$$\mathrm{FO}^k := \{\psi \in \mathrm{FO} : \mathrm{width}(\psi) \leq k\}.$$

Clearly $|\mathcal{G}(\mathfrak{A}, \psi)| \leq |\psi| \cdot |A|^k$ for any finite structure $\mathfrak{A}$ and any $\psi \in \mathrm{FO}^k$.

Theorem 1.17. ModCheck(FO^k) is solvable in time $O(|\psi| \cdot |A|^k)$ and P-complete, for every $k \geq 2$.

As shown in Theorem 1.5, modal logic can be embedded (efficiently) into FO^2. Hence, also ML model checking has polynomial time complexity.

It is a general observation that modal logics have many convenient model-theoretic and algorithmic properties. Besides efficient model-checking the following facts are important in many applications of modal logic.

- The satisfiability problem for ML is decidable (in PSPACE),
- ML has the finite model property: each satisfiable formula has a finite model,
- ML has the tree model property: each satisfiable formula has a tree-shaped model,
- algorihmic problems for ML can be solved by automata-based methods.

The embedding of ML into FO^2 has sometimes been proposed as an explanation for the good properties of modal logic, since FO^2 is a first-order fragment that shares some of these properties. However, more recently, it has been seen that this explanation has its limitations and is not really convincing. In particular, there are many extensions of ML to temporal and dynamic logics such as LTL, CTL, CTL*, PDL and the μ-calculus L_μ that are of great importance for applications in computer science, and that preserve many of the good algorithmic properties of ML. Especially the associated satisfiability problems remain decidable. However this is not at all true for the corresponding extension of FO^2.

A better and more recent explanation for the good properties of modal logic is that modal operators correspond to a restricted form of quantification, namely *guarded quantification*. Indeed, in the embedding of ML into FO^2, all quantifiers are *guarded* by atomic formulae. This can be vastly generalised beyond two-variable logic and involving formulae of arbitrary relational vocabularies, leading to the *guarded fragment* of FO.

Definition 1.18. The *guarded fragment* of first-order logic GF is the fragment of first-order logic which allows only guarded quantification

$$\exists \bar{y}(\alpha(\bar{x},\bar{y}) \land \varphi(\bar{x},\bar{y})) \text{ and } \forall \bar{y}(\alpha(\bar{x},\bar{y}) \rightarrow \varphi(\bar{x},\bar{y})),$$

where the *guards* α are atomic formulae containing all free variables of φ.

GF is a generalisation of modal logics: $\mathrm{ML} \subseteq \mathrm{GF} \subseteq \mathrm{FO}$. Indeed, the modal operators $\Diamond$ and $\Box$ can be expressed as

$$\langle a \rangle \varphi \equiv \exists y(E_a xy \land \varphi(y)) \text{ and } [a]\varphi \equiv \forall y(E_a xy \rightarrow \varphi(y)).$$

It has turned out that the guarded fragment preserves (and explains to some extent) essentially all of the good model-theoretic and algorithmic properties of modal logics, in a far more expressive setting. In terms of model-checking games, we can observe that guarded logics have small model checking games of size $\|\mathcal{G}(\mathfrak{A},\psi)\| = O(|\psi| \cdot \|\mathfrak{A}\|)$, and so there exist efficient game-based model-checking algorithms for them.

2 Parity Games and Fixed-Point Logics

In the first chapter we have discussed model checking games for first-order logic and modal logic. These games admit only finite plays and their winning conditions are specified just by sets of positions, that the players want to reach. Winning regions in these games can be computed in linear time with respect to the size of the game graph.

However, in many computer science applications, more expressive logics are needed, such as temporal logics, dynamic logics, fixed-point logics and others. Model checking games for these logics admit infinite plays and their winning conditions must be specified in a more elaborate way. As a consequence, we have to consider the theory of infinite games.

For fixed-point logics, such as LFP or the modal μ-calculus, the appropriate evaluation games are *parity games*. These are games of possibly infinite duration with a function that assigns to each position a natural number, called its *priority*. The winner of an infinite play is determined according to whether the least priority seen infinitely often during the play is even or odd.

2.1 Parity Games

Definition 2.1. A *parity game* is given by a labelled game graph $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ as in Sect. 1.3 with a function $\Omega : V \to \mathbb{N}$ that assigns a *priority* to each position. The set V of positions may be finite or infinite, but $|\Omega(V)|$, the number of different priorities which is called the *index* of $\mathcal{G}$, must be finite. As before, a finite play is lost by the player who gets stuck, i.e. cannot move. For infinite plays $v_0 v_1 v_2 \ldots$, we have the *parity winning condition*: If the least number appearing infinitely often in the sequence $\Omega(v_0)\Omega(v_1)\ldots$ of priorities is even, then Player 0 wins the play, otherwise Player 1 wins.

A strategy (for Player σ) is a function $f : V^*V_\sigma \to V$ such that $f(v_0v_1 \dots v_n) \in v_nE$. We say that a play $\pi = v_0v_1 \dots$ is *consistent* with the strategy f of Player σ if for each $v_i \in V_\sigma$ it holds that $v_{i+1} = f(v_0 \dots v_i)$. The strategy f is *winning* for Player σ from (or on) a set $W \subseteq V$ if each play starting in W that is consistent with f is won by Player σ.

In general, a strategy may depend on the entire history played so far, and can thus be a very complicated object. However, we are interested in simple strategies that depend only on the current position.

Definition 2.2. A strategy (of Player σ) is called *positional* (or *memoryless*) if it only depends on the current position, but not on the history of the play, which means that $f(hv) = f(h'v)$ for all $h, h' \in V^*, v \in V$. We can view positional strategies simply as functions $f : V_\sigma \to V$.

We shall see that positional strategies suffice to solve parity games. Before we formulate and prove this Forgetful Determinacy Theorem, we recall that positional strategies are of course sufficient whenever, as in the previous chapter, the players have purely positional objectives such as reachability or safety. Specifically, for every game $\mathcal{G} = (V, V_0, V_1, E)$ and every $X \subseteq V$ we have defined the attractor

$$\begin{aligned}\mathrm{Attr}_\sigma(X) = \{v \in V : \text{Player } \sigma \text{ has a strategy from } v \text{ to reach some}\\ \text{position } x \in X \cup T_\sigma\}\end{aligned}$$

and such an *attractor strategy* can, without loss of generality, assumed to be positional. Similarly, if $Y \subseteq V$ is a *trap* for Player σ, then Player $(1 - \sigma)$ has a positional *trap strategy* to keep the play inside Y.

Further we note that positional winning strategies on parts of the game graph may be combined to positional winning strategies on larger regions. Indeed, let f and f' be positional strategies for Player σ that are winning exactly on the sets W, W', respectively. Let $(f + f')$ be the positional strategy defined by

$$(f + f')(x) := \begin{cases} f(x) & \text{if } x \in W \\ f'(x) & \text{otherwise.} \end{cases}$$

Then $(f + f')$ is a winning strategy on $W \cup W'$.

We can now turn to the proof of the Forgetful Determinacy Theorem.

Theorem 2.3 (Forgetful Determinacy)**.** In any parity game, the set of positions can be partitioned into two sets W_0 and W_1 such that Player 0 has a positional strategy that is winning on W_0 and Player 1 has a positional strategy that is winning on W_1.

Proof. Let $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ be a parity game with $|\Omega(V)| = m$. Without loss of generality we can assume that $\Omega(V) = \{0, \ldots, m-1\}$ or $\Omega(V) = \{1, \ldots, m\}$. We prove the statement by induction over $|\Omega(V)|$.

In the case that $|\Omega(V)| = 1$, i.e., $\Omega(V) = \{0\}$ or $\Omega(V) = \{1\}$, either Player 0 or Player 1 wins every infinite play. Her opponent can only win by reaching a terminal position that does not belong to him. So we have, for $\Omega(V) = \{\sigma\}$,

$$\begin{aligned} W_{1-\sigma} &= \mathrm{Attr}_{1-\sigma}(T_{1-\sigma}) \text{ and} \\ W_\sigma &= V \setminus W_{1-\sigma}. \end{aligned}$$

Computing $W_{1-\sigma}$ as the attractor of $T_{1-\sigma}$ is a simple reachability problem, and thus it can be solved with a positional strategy. For W_σ there is a positional strategy that avoids leaving this $(1-\sigma)$-trap.

Let now $|\Omega(v)| = m > 1$. We explicitly consider the case that $0 \in \Omega(V)$, i.e., $\Omega(V) = \{0, \ldots, m-1\}$. Otherwise, if the minimal priority is 1, we can use the same argumentation with switched roles of the players. We define

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

and let g be a positional winning strategy for Player 1 on X_1.

Our goal is to provide a positional winning strategy f^* for Player 0 on $X_0 := V \setminus X_1$, so in particular we have $W_1 = X_1$ and $W_0 = V \setminus X_1$.

First of all, observe that X_0 is a trap for Player 1. Indeed, if Player 1 could reach X_1 from some $v \in X_0$, then Player 1 could win with a positional strategy from v, so v would also be in X_1. Thus, there exists a positional *trap strategy* t for Player 0 on X_0 that guarantees that a play remains inside X_0.

Let $Y = \Omega^{-1}(0) \cap X_0$ and $Z = \mathrm{Attr}_0(Y)$. Player 0 has a positional attractor strategy a to ensure, from every position $z \in Z \setminus Y$, that Y (or a terminal winning position in T_0) is reached in finitely many steps.

Let now $V' = V \setminus (X_1 \cup Z)$. The restricted game $\mathcal{G}' = \mathcal{G}|_{V'}$ has strictly fewer priorities than $\mathcal{G}$ (since at least all positions with priority 0 have been removed). Thus, by induction hypothesis, the Forgetful Determinacy Theorem holds for $\mathcal{G}'$. This means that $V' = W_0' \cup W_1'$ and there exist positional winning strategies f' for Player 0 on W_0' and g' for Player 1 on W_1' in $\mathcal{G}'$.

However, it follows that $W_1' = \emptyset$, since the strategy

$$(g + g') : x \mapsto \begin{cases} g(x) & x \in X_1 \\ g'(x) & x \in W_1' \end{cases}$$

is a positional winning strategy for Player 1 on $X_1 \cup W_1'$. Indeed, every play consistent with $(g + g')$ either stays in W_1' and is consistent with g' or reaches X_1 and is from this point on consistent with g. But X_1, by definition, already contains *all* positions from which Player 1 can win with a positional strategy, so $W_1' = \emptyset$.

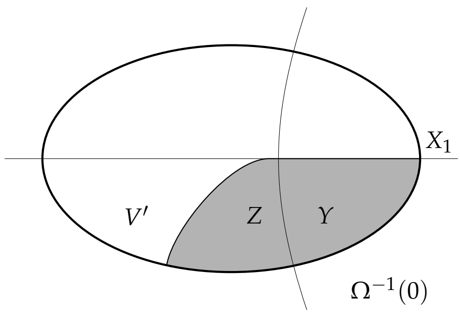

Figure 2.1. Construction of a winning strategy

Knowing that $W_1' = \emptyset$, let $f^* = f' + a + t$, i.e.

$$f^*(x) = \begin{cases} f'(x) & \text{if } x \in W_0' \\ a(x) & \text{if } x \in Z \setminus Y \\ t(x) & \text{if } x \in Y \end{cases}$$

We claim that f^* is a positional winning strategy for Player 0 from X_0. Note that if π is a play that is consistent with f^*, then π remains inside X_0. We distinguish two cases.
Case (a): π hits Z only finitely often. Then π eventually stays in W_0' and is consistent with f' from this point onwards. Hence Player 0 wins π.
Case (b): π hits Z infinitely often. Then π also hits Y infinitely often, which implies that priority 0 is seen infinitely often. Thus, Player 0 wins π. Q.E.D.

The following theorem is a consequence of positional determinacy.

Theorem 2.4. It can be decided in NP $\cap$ coNP whether a given position in a parity game is a winning position for Player 0.

Proof. A node v in a parity game $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ is a winning position for Player σ if there exists a positional strategy $f : V_\sigma \to V$ which is winning from position v. It therefore suffices to show that the question whether a given strategy $f : V_\sigma \to V$ is a winning strategy for Player σ from position v can be decided in polynomial time. We prove this for Player 0; the argument for Player 1 is analogous.

Given $\mathcal{G}$ and $f : V_0 \to V$, we obtain a reduced game graph $\mathcal{G}_f = (W, F)$ by retaining only those moves that are consistent with f, i.e.,

$$\begin{aligned} F = \{(v, w) : {} & (v \in W \cap V_\sigma \wedge w = f(v)) \vee \\ & (v \in W \cap V_{1-\sigma} \wedge (v, w) \in E)\}. \end{aligned}$$

In this reduced game, only the opponent, Player 1, makes non-trivial moves. We call a cycle in (W, F) odd if the least priority of its nodes is odd. Clearly, Player 0 wins $\mathcal{G}$ from position v via strategy f if, and only if, in $\mathcal{G}_f$ no odd cycle and no terminal position $w \in V_0$ is reachable from v. Since the reachability problem is solvable in polynomial time, the claim follows. Q.E.D.

2.2 Algorithms for parity games

It is an open question whether winning sets and winning strategies for parity games can be computed in polynomial time. The best algorithms known today are polynomial in the size of the game, but exponential with respect to the number of priorities. On an class of parity games with bounded index, such algorithms run in polynomial time.

One way to intuitively understand an algorithm solving a parity game is to imagine a referee who watches the players playing the game. At some point, the referee is supposed to say "Player 0 wins", and indeed, whenever the referee does so, there should be no question that Player 0 wins. We shall first give a formal definition of a certain kind of referee with bounded memory, and later use this notion to construct algorithms for parity games.

Definition 2.5. A *referee* $\mathcal{M} = (M, m_0, \delta, F)$ for a parity game $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ consists of a set of states M with a distinguished initial state $m_0 \in M$, a set of final states $F \subseteq M$, and a transition function $\delta : V \times M \to M$. Note that a referee is thus formally the same as an automaton reading words over the alphabet V. But to be called a referee, two further conditions must be satisfied, for any play $v_0 v_1 \dots$ of $\mathcal{G}$, and and the corresponding sequence $m_0 m_1 \dots$ of states of $\mathcal{M}$, where m_0 is the initial state of $\mathcal{M}$ and $m_{i+1} = \delta(v_i, m_i)$:

(1) If $v_0 \dots$ is winning for Player 0, then there is a k such that $m_k \in F$,
(2) If $m_k \in F$ for some k, then there exist $i < j \leq k$ such that $v_i = v_j$ and $\min\{\Omega(v_{i+1}), \Omega(v_{i+2}), \dots, \Omega(v_j)\}$ is even.

To illustrate the second condition in the above definition, note that in the play $v_0 v_1 \dots$ the sequence $v_i v_{i+1} \dots v_j$ forms a cycle. Assuming that both players use a positional strategy the decision of the referee is correct. Indeed, if a cycle with even priority appears, then this cycle will be repeated forever, Player 0 can be declared as the winner. To capture this intuition formally, we define the following reachability game, which emerges as the product of the original game $\mathcal{G}$ and the referee $\mathcal{M}$.

Definition 2.6. Let $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ be a parity game and $\mathcal{M} = (M, m_0, \delta, F)$ an automaton reading words over V. We associate with $\mathcal{G}$

and $\mathcal{M}$ a reachability game

$$\mathcal{G} \times \mathcal{M} = (V \times M, V_0 \times M, V_1 \times M, E', V \times F),$$

where $((v,m),\ (v',m')) \in E'$ iff $(v,v') \in E$ and $m' = \delta(v,m)$, and $V \times F$ is the set of positions which are immediately winning for Player 0 (the goal of Player 0 is to reach such a position). Plays that do not reach a position in $V \times F$ are won by Player 1.

Note that $\mathcal{M}$ in the definition above is a deterministic automaton, i.e., δ is a function. Therefore, in $\mathcal{G}$ and in $\mathcal{G} \times \mathcal{M}$ the players have the same choices, and thus it is possible to translate strategies between $\mathcal{G}$ and $\mathcal{G} \times \mathcal{M}$. Formally, for a strategy f in $\mathcal{G}$ we define the strategy $\overline{f}$ in $\mathcal{G} \times \mathcal{M}$ as

$$\overline{f}((v_0, m_0)(v_1, m_1) \dots (v_n, m_n)) = (f(v_0 v_1 \dots v_n), \delta(v_n, m_n)).$$

Conversely, given a strategy f in $\mathcal{G} \times \mathcal{M}$ we define the strategy $\underline{f}$ in $\mathcal{G}$ such that $\underline{f}(v_0 v_1 \dots v_n) = v_{n+1}$ if and only if

$$f((v_0, m_0)(v_1, m_1) \dots (v_n, m_n)) = (v_{n+1}, m_{n+1}),$$

where $m_0 m_1 \dots$ is the unique sequence corresponding to $v_0 v_1 \dots$.

With $\mathcal{G} \times \mathcal{M}$ we are ready to prove that the definition of a referee indeed makes sense for parity games.

Theorem 2.7. Let $\mathcal{G}$ be a parity game and $\mathcal{M}$ a referee for $\mathcal{G}$. Then Player 0 wins $\mathcal{G}$ from v_0 if, and only if, she wins $\mathcal{G} \times \mathcal{M}$ from (v_0, m_0).

Proof. ($\Rightarrow$) Let f be a winning strategy for Player 0 in $\mathcal{G}$ from v_0. Assume that Player 0 does not have a winning strategy for $\mathcal{G} \times \mathcal{M}$ from (v_0, m_0). By determinacy of reachability games, there exists a winning strategy g for Player 1. Consider the unique play $\pi_{\mathcal{G}} = v_0 v_1 \dots$ that is consistent with f and $\underline{g}$ and the unique play $\pi_{\mathcal{G}\times\mathcal{M}} = (v_0, m_0)(v_1, m_1) \dots$ which is consistent with $\overline{f}$ and g. Observe that the positions of $\mathcal{G}$ appearing in both plays are indeed the same due to the way $\overline{f}$ and $\underline{g}$ are defined. Since Player 0 wins $\pi_{\mathcal{G}}$, by Property (1) in the definition of a referee there must be an $m_k \in F$. But this contradicts the fact that Player 1 wins $\pi_{\mathcal{G}\times\mathcal{M}}$.

($\Leftarrow$) Let f be a winning strategy for Player 0 in $\mathcal{G} \times \mathcal{M}$, and assume that Player 1 has a *positional* winning strategy g in $\mathcal{G}$. Again, we consider the unique plays $pi_{\mathcal{G}} = v_0 v_1 \ldots$ and $\pi_{\mathcal{G}\times\mathcal{M}} = (v_0, m_0)(v_1, m_1)\ldots$ such that $\pi_{\mathcal{G}}$ is consistent with $\underline{f}$ and g, and $\pi_{\mathcal{G}\times\mathcal{M}}$ is consistent with f and $\overline{g}$. Since $\pi_{\mathcal{G}\times\mathcal{M}}$ is won by Player 0, there is an $m_k \in F$ appearing in this play.

By Property (2) in the definition of a referee, there exist two indices $i < j$ such that $v_i = v_j$ and the minimum priority appearing between v_i and v_j is even. Let us now consider the following strategy f' for Player 0 in $\mathcal{G}$:

$$f'(w_0 w_1 \ldots w_n) = \begin{cases} \underline{f}(w_0 w_1 \ldots w_n) & \text{if } n < j, \\ \underline{f}(w_0 w_1 \ldots w_m) & \text{otherwise,} \end{cases}$$

where $m = i + [(n-i) \mod (j-i)]$. Intuitively, the strategy f' makes the same choices as $\underline{f}$ up to the $(j-1)$st step, and then repeats the choices of $\underline{f}$ from steps $i, i+1, \ldots, j-1$.

We claim that the unique play π' in $\mathcal{G}$ that is consistent with both f' and g is won by Player 0. Since in the first j steps f' is the same as $\underline{f}$, we have that $\pi[n] = v_n$ for all $n \leq j$. Now observe that $\pi[j+1] = v_{i+1}$. Since g is positional, if v_j is a position of Player 1, then $\pi[j+1] = v_{i+1}$, and if v_j is a position of Player 0, then $\pi[j+1] = v_{i+1}$ because we defined $f'(v_0 \ldots v_j) = f(v_0 \ldots v_i)$. By induction we get that the play π repeats the cycle $v_i v_{i+1} \ldots v_j$ infinitely often, i.e.

$$\pi = v_0 \ldots v_{i-1}(v_i v_{i+1} \ldots v_{j-1})^\omega.$$

Thus, the minimal priority occurring infinitely often in π is the same as $\min\{\Omega(v_i), \Omega(v_{i+1}), \ldots \Omega(v_{j-1})\}$, and thus is even. Therefore Player 0 wins π, which contradicts the fact that g was a winning strategy for Player 1. Q.E.D.

This theorem allows us, if a referee is known, to reduce the problem of solving a parity game to the problem of solving a reachability game, which we already tackled with the GAME algorithm. But to make use of it, we first need to construct a referee for a given parity game.

The most naïve way to build a referee for a parity game is to just remember, for each position v visited during the play, the minimal priority seen since the last occurrence of v. If it happens that a position v is repeated, and the minimal priority seen since the last occurrence of v is even, the referee decides that Player 0 wins the play.

It is easy to check that an automaton defined in this way is indeed a referee for $\mathcal{G}$, but such a referee can be very big. Since for each of the $|V| = n$ positions we need to store one of $|\Omega(V)| = d$ colours, the size of the referee is in the order of $O(d^n)$. We shall present a referee that is much better for small d.

Definition 2.8. A *progress-measuring referee* $\mathcal{M}_P = (M_P, m_0, \delta_P, F_P)$ for a parity game $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ is constructed as follows. If $n_i = |\Omega^{-1}(i)|$ is the number of positions with priority i, then

$$M_P = \{0, 1, \ldots, n_0 + 1\} \times \{0\} \times \{0, 1, \ldots, n_2 + 1\} \times \{0\} \times \ldots$$

and this product ends in $\cdots \times \{0, 1, \ldots, n_m + 1\}$ if the maximal priority m is even, or in $\cdots \times \{0\}$ if it is odd. The initial state is $m_0 = (0, \ldots, 0)$, and the transition function $\delta(v, \bar{c})$ with $\bar{c} = (c_0, 0, c_2, 0, \ldots, c_m)$ is given by

$$\delta(v, \bar{c}) = \begin{cases} (c_0, 0, c_2, 0, \ldots, c_{\Omega(v)} + 1, 0, \ldots, 0) & \text{if } \Omega(v) \text{ is even,} \\ (c_0, 0, c_2, 0, \ldots, c_{\Omega(v)-1}, 0, 0, \ldots, 0) & \text{otherwise.} \end{cases}$$

The set F_P contains all tuples $(c_0, 0, c_2, \ldots, c_m)$ in which some counter $c_j = n_j + 1$ reached the maximum possible value.

The intuition behind $\mathcal{M}_P$ is that it counts, for each even priority p, how many positions with priority p were seen without any lower priority in between. If more than n_p such positions are seen, then at least one must have been repeated, which guarantees that $\mathcal{M}_P$ is a referee.

Lemma 2.9. For each finite parity game $\mathcal{G}$ the automaton $\mathcal{M}_P$ constructed above is a referee for $\mathcal{G}$.

Proof. We need to show that $\mathcal{M}_P$ exhibits the two properties characterising a referee:

(1) if $v_0 \ldots$ is winning for Player 0, then there is a k such that $m_k \in F$,

(2) if, for some k, $m_k \in F$, then there exist $i < j \leq k$ such that $v_i = v_j$ and $\min\{\Omega(v_{i+1}), \Omega(v_{i+2}), \ldots, \Omega(v_j)\}$ is even.

To see (1), assume that $v_0 v_1 \ldots$ is a play winning for Player 0. Let k be such an index that $\Omega(v_k)$ is even, appears infinitely often in $\Omega(v_k)\Omega(v_{k+1})\ldots$, and no priority lower than $\Omega(v_k)$ appears in this play suffix. Then, starting from v_k, the counter $c_{\Omega(v_k)}$ will never be decremented, but it will be incremented infinitely often. Thus, for a finite game $\mathcal{G}$, it will reach $n_{\Omega(v_k)} + 1$ at some point, i.e. a state in F_P.

To prove (2), let $v_0 v_1 \ldots v_k$ be such a prefix of a play that after v_k some counter c_p is set to $n_p + 1$ for an even priority p. Let v_{i_0} be the last position at which this counter was 0, and v_{i_m} the subsequent positions at which it was incremented, up to $i_{n_p} = k$. All positions $v_{i_0}, v_{i_1}, \ldots, v_{i_{n_p}}$ have priority p, but since there are only n_p different positions with priority p, we get that, for some $k < l$, $v_{i_k} = v_{i_l}$. Now i_k and i_l are the positions required to witness (2), because indeed the minimum priority between i_k and i_l is p since c_p was not reset in between. Q.E.D.

For a parity game $\mathcal{G}$ with an even number of priorities d, the above presented referee has size $n_0 \cdot n_2 \cdots n_d$, which is at most $(\frac{n}{d/2})^{d/2}$. We get the following corollary.

Corollary 2.10. Parity games can be solved in time $O((\frac{n}{d/2})^{d/2})$.

Notice that the algorithm using a referee has high space demand: Since the product game $\mathcal{G} \times \mathcal{M}_P$ must be explicitly constructed, the space complexity of this algorithm is the same as its time complexity. There is a method to improve the space complexity by storing the maximal counters the referee $\mathcal{M}_P$ uses in each position and lifting such annotations. This method is called *game progress measures* for parity games. We will not define it here, but the equivalence to modal μ-calculus proven in the next chapter will provide another algorithm for solving parity games with polynomial space complexity.

2.3 Fixed-Point Logics

We will define two fixed-point logics, the modal μ-calculus, L_μ, and the first-order least fixed-point logic, LFP, which extend modal logic and first-order logic, respectively, with the operators for least and greatest fixed-points.

The syntax of L_μ is analogous to modal logic, with two additional rules for building least and greatest fixed-point formulas:

$$\mu X.\varphi(X) \text{ and } \nu X.\varphi(X)$$

are L_μ formulas if $\varphi(X)$ is, where X is a variable that can be used in φ the same way as predicates are used, but must *occur positively in* φ, i.e. under an even number of negations (or, if φ is in negation normal form, simply non-negated).

The syntax of LFP is analogous to first-order logic, again with two additional rules for building fixed-points, which are now syntactically more elaborate. Let $\varphi(T, x_1, x_2, \dots x_n)$ be a LFP formula where T stands for an n-ary relation and occurs only positively in φ. Then both

$$[\mathrm{lfp}\, T\bar{x}.\varphi(T,\bar{x})](\bar{a}) \text{ and } [\mathrm{gfp}\, T\bar{x}.\varphi(T,\bar{x})](\bar{a})$$

are LFP formulas, where $\overline{a} = a_1 \dots a_n$.

To define the semantics of L_μ and LFP, observe that each formula $\varphi(X)$ of L_μ or $\varphi(T,\bar{x})$ of LFP defines an operator $[\![\varphi(X)]\!] : \mathcal{P}(V) \to \mathcal{P}(V)$ on states V of a Kripke structure $\mathcal{K}$ and $[\![\varphi(T,\bar{x})]\!] : \mathcal{P}(A^n) \to \mathcal{P}(A^n)$ on tuples from the universe of a structure $\mathfrak{A}$. The operators are defined in the natural way, mapping a set (or relation) to a set or relation of all these elements, which satisfy φ with the former set taken as argument:

$$[\![\varphi(X)]\!](B) = \{v \in \mathcal{K} : \mathcal{K}, v \models \varphi(B)\}, \text{ and}$$

$$[\![\varphi(T,\bar{x})]\!](R) = \{\bar{a} \in \mathfrak{A} : \mathfrak{A} \models \varphi(R,\bar{a})\}.$$

An argument B is a fixed-point of an operator f if $f(X) = X$, and to complete the definition of the semantics, we say that $\mu X.\varphi(X)$ defines the *smallest* set B that is a fixed-point of $[\![\varphi(X)]\!]$, and $\nu X.\varphi(X)$ defines the *largest* such set. Analogously, $[\text{lfp}\, T\bar{x}.\varphi(T,\bar{x})](\bar{x})$ and $[\text{gfp}\, T\bar{x}.\varphi(T,\bar{x})](\bar{x})$ define the smallest and largest relations being a fixed-point of $[\![\varphi(T,\bar{x})]\!]$, respectively. In a few paragraphs, we will give an alternative characterisation of least and greatest fixed-points, which is better tailored towards an algorithmic computation.

To justify this definition, we have to assure that all notions are well-defined, i.e., in particular, we have to show that the operators actually have fixed-points, and that least and greatest fixed-points always exist. In fact, this relies on the monotonicity of the operators used.

Definition 2.11. An operator F is *monotone* if

$$X \subseteq Y \implies F(X) \subseteq F(Y).$$

The operators $[\![\varphi(X)]\!]$ and $[\![\varphi(T,\bar{x})]\!]$ are monotone because we assumed that X (or T) occurs only positively in φ, and, except for negation, all other logical operators are monotone (the fixed-point operators as well, as we will see). Each monotone operator not only has unique least and greatest fixed-points, but these can be calculated iteratively, as stated in the following theorem.

Remark 2.12. A formal definiton of ordinal numbers can be found in appendix A. For the moment, we think of them as a generalisation of the naturals numbers which allow to count beyond the finite. The first ordinal numbers are the natural numbers $0, 1, 2, \ldots$ itself. The least infinite ordinal number is the set of all natural numbers, written as ω, followed by $\omega+1, \omega+2, \ldots, \omega\cdot 2, \omega\cdot 2+1, \ldots, \omega^2, \ldots, \omega^\omega, \ldots$.

Definition 2.13. Let A be a set, and $F : \mathcal{P}(A^k) \to \mathcal{P}(A^k)$ be a monotone operator. We define the stages X_α of an inductive fixed-point process:

$$\begin{aligned} X_0 &:= \emptyset \\ X_{\alpha+1} &:= f(X_\alpha) \\ X_\lambda &:= \bigcup_{\alpha<\lambda} X_\alpha \qquad \text{for limit ordinals } \lambda. \end{aligned}$$

Due to the monotonicity of F, the sequence of stages is increasing, i.e.

$X_\alpha \subseteq X_\beta$ for $\alpha < \beta$, and hence for some γ, called the *closure ordinal*, we have $X_\gamma = X_{\gamma+1} = F(X_\gamma)$. This fixed-point is called the *inductive fixed-point* and denoted by X_∞.

Analogously, we can define the stages of a similar process:

$$\begin{aligned} X^0 &:= A^k \\ X^{\alpha+1} &:= F(X^\alpha) \\ X^\lambda &:= \bigcap_{\alpha<\lambda} X^\alpha \qquad \text{for limit ordinals } \lambda. \end{aligned}$$

which yields a decreasing sequence of stages X^α that leads to the inductive fixed-point $X^\infty := X^\gamma$ for the smallest γ such that $X^\gamma = X^{\gamma+1}$.

Theorem 2.14 (Knaster, Tarski)**.** Let F be a monotone operator. Then the least fixed-point $\text{lfp}(F)$ and the greatest fixed-point $\text{gfp}(F)$ of F exist, they are unique and correspond to the inductive fixed-points, i.e. $\text{lfp}(F) = X_\infty$, and $\text{gfp}(F) = X^\infty$.

To understand the inductive evaluation let us consider an example. We will evaluate the formula $\mu X.(P \vee \Diamond X)$ on the following Kripke structure:

$$\mathcal{K} = (\{0, \dots, n\}, \{(i, i+1) \mid i < n\}, \{n\}).$$

The structure $\mathcal{K}$ represents a path of length $n+1$ ending in a position marked by the predicate P. The evaluation of this least fixed-point formula starts with $X_0 = \emptyset$ and $X_1 = P = \{n\}$, and in step $i+1$ all nodes having a successor in X_i are added. Therefore, $X_2 = \{n-1, n\}$, $X_3 = \{n-2, n-1, n\}$, and in general $X_k = \{n-k+1, \dots, n\}$. Finally, $X_{n+1} = X_{n+2} = \{0, \dots, n\}$. As you can see, the formula $\mu X.(P \vee \Diamond X)$ describes the set of nodes from which P is reachable. This example shows one motivation for the study of fixed-point logics: It is possible to express transitive closures of various relations in such logics.

2.4 Model Checking Games for Fixed-Point Logics

In this section we will see that parity games are the model checking games for LFP and L_μ.

We will construct a parity game $\mathcal{G}(\mathfrak{A}, \Psi(\bar{a}))$ from a formula $\Psi(\bar{x}) \in$ LFP, a structure $\mathfrak{A}$ and a tuple $\bar{a}$ by extending the FO game with the moves

$$[\mathrm{fp}\, T\bar{x}.\varphi(T, \bar{x})](\bar{a}) \to \varphi(T, \bar{a})$$

and

$$T\bar{b} \to \varphi(T, \bar{b}).$$

We assign priorities $\Omega(\varphi(\bar{a})) \in \mathbb{N}$ to every instantiation of a subformula $\varphi(\bar{x})$. Therefore, we need to make some assumptions on Ψ:

- Ψ is given in negation normal form, i.e. negations occur only in front of atoms.
- Every fixed-point variable T is bound only once in a formula $[\mathrm{fp}\, T\bar{x}.\varphi(T, \bar{x})]$.
- In a formula $[\mathrm{fp}\, T\bar{x}.\varphi(T, \bar{x})]$ there are no other free variables besides $\bar{x}$ in φ.

Then we can assign the priorities using the following schema:

- $\Omega(T\bar{a})$ is even if T is a gfp-variable, and $\Omega(T\bar{a})$ is odd if T is an lfp-variable.
- If T' depends on T (i.e. T occurs freely in $[\mathrm{fp}\, T'\bar{x}.\varphi(T, T', \bar{x})]$), then $\Omega(T\bar{a}) \leq \Omega(T'\bar{b})$ for all $\bar{a}, \bar{b}$.
- $\Omega(\varphi(\bar{a}))$ is maximal if $\varphi(\bar{a})$ is not of the form $T\bar{a}$.

Remark 2.15. The minimal number of different priorities in the game $\mathcal{G}(\mathfrak{A}, \Psi(\bar{a}))$ corresponds to the alternation depth of Ψ.

Before we provide the proof that parity games are in fact the appropriate model checking games for LFP and L_μ, we introduce the notion of an *unfolding* of a parity game.

Let $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ be a parity game. We assume that the minimal priority in $\mathcal{G}$ is 0 and that all positions $v \in V$ with $\Omega(v) = 0$ have a unique successor, i.e., $vE = \{s(v)\}$.

Let $T = \{v \in V : \Omega(v) = 0\}$. We define a modified game $\mathcal{G}^- = (V, V_0, V_1, E^-, \Omega)$ with $E^- = E \setminus (T \times V)$, i.e., positions in T are made terminal positions in $\mathcal{G}^-$. Further, we define a sequence of games $\mathcal{G}^\alpha$ that only differ from $\mathcal{G}^-$ in the assignment of the terminal positions in T to the players. For this purpose, we use a sequence of partitions (T_0^α, T_1^α) of T such that in $\mathcal{G}^\alpha$, Player σ wins at final positions $v \in T_\sigma^\alpha$. The sequence of partitions is inductively defined depending on the winning regions W_σ^α of the players in the games $\mathcal{G}^\alpha$ as follows:

- $T_0^0 := T$,
- $T_0^{\alpha+1} := \{v \in T : s(v) \in W_0^\alpha\}$ for any ordinal α,
- $T_0^\lambda := \bigcup_{\alpha<\lambda} T_0^\alpha$ if λ is a limit ordinal,
- $T_1^\alpha = T \setminus T_0^\alpha$ for any ordinal α.

  We have

- $W_0^0 \supseteq W_0^1 \supseteq W_0^2 \supseteq \dots \supseteq W_0^\alpha \supseteq W_0^{\alpha+1} \dots$
- $W_1^0 \subseteq W_1^1 \subseteq W_1^2 \subseteq \dots \subseteq W_1^\alpha \subseteq W_1^{\alpha+1} \dots$

So there exists an ordinal $\alpha \leq |V|$ such that $W_0^\alpha = W_0^{\alpha+1} = W_0^\infty$ and $W_1^\alpha = W_1^{\alpha+1} = W_1^\infty$.

Lemma 2.16 (Unfolding Lemma)**.**

$$W_0 = W_0^\infty \quad \text{and} \quad W_1 = W_1^\infty.$$

Proof. Let α be such that $W_0^\infty = W_0^\alpha$ and let f^α be a positional winning strategy for Player 0 from W_0^α in $\mathcal{G}$. Define:

$$f \;:\; V_0 \to V \;:\; v \mapsto \begin{cases} f^\alpha(v) & \text{if } v \in V_0 \setminus T, \\ s(v) & \text{if } v \in V_0 \cap T. \end{cases}$$

A play π consistent with f that starts in W_0^∞ never leaves W_0^∞:

- If $\pi(i) \in W_0^\infty \setminus T$, then $\pi(i+1) = f^\alpha(\pi(i)) \in W_0^\alpha = W_0^\infty$ (f_α is a winning strategy in $\mathcal{G}^\alpha$).
- If $\pi(i) \in W_0^\infty \cap T = W_0^\alpha \cap T = W_0^{\alpha+1} \cap T$, then $\pi(i) \in T_0^{\alpha+1}$, i.e. $\pi(i)$ is a terminal position in $\mathcal{G}^\alpha$ from which Player 0 wins, so by the definition of $T_0^{\alpha+1}$ we have $\pi(i+1) = s(v) \in W_0^\alpha = W_0^\infty$.

Thus, we can conclude that Player 0 wins π:

- If π hits T only finitely often, then from some point onwards π is consistent with f^α and stays in W_0^α which results in a winning play for Player 0.

- Otherwise, $\pi(i) \in T$ for infinitely many i. Since we had $\Omega(t) = 0 \leq \Omega(v)$ for all $v \in V, t \in T$, the lowest priority seen infinitely often is 0, so Player 0 wins π.

For $v \in W_1^\infty$, we define $\rho(v) = \min\{\beta : v \in W_1^\beta\}$ and let g^β be a positional winning strategy for Player 1 on W_1^β in $\mathcal{G}^\beta$. We define a positional strategy g of Player 1 in $\mathcal{G}^\infty$ by:

$$g : V_1 \to V, \quad v \mapsto \begin{cases} g^{\rho(v)}(v) & \text{if } v \in W_1^\infty \setminus T \cap V_1 \\ s(v) & \text{if } v \in T \cap V_1 \\ \text{arbitrary} & \text{otherwise} \end{cases}$$

Let $\pi = \pi(0)\pi(1)\ldots$ be a play consistent with g and $\pi(0) \in W_1^\infty$.

Claim 2.17. Let $\pi(i) \in W_1^\infty$. Then

(1) $\pi(i+1) \in W_1^\infty$,
(2) $\rho(\pi(i+1)) \leq \rho(\pi(i))$
(3) $\pi(i) \in T \Rightarrow \rho(\pi(i+1)) < \rho(\pi(i))$.

Proof. Case (1): $\pi(i) \in W_1^\infty \setminus T$, $\rho(\pi(i)) = \beta$ (so $\pi(i) \in W_1^\beta$). We have $\pi(i+1) = g(\pi(i)) = g^\beta(\pi(i))$, so $\pi(i+1) \in W_1^\beta \subseteq W_1^\infty$ and $\rho(\pi(i+1)) \leq \beta = \rho(\pi(i))$.
Case (2): $\pi(i) \in W_1^\infty \cap T$, $\rho(\pi(i)) = \beta$. Then we have $\pi(i) \in W_1^\infty$, $\beta = \gamma + 1$ for some ordinal γ, and $\pi(i+1) = s(\pi(i)) \in W_1^\gamma$, so $\pi(i+1) \in W_1^\infty$ and $\rho(\pi(i+1)) \leq \gamma < \beta = \rho(\pi(i))$. Q.E.D.

As there is no infinite descending chain of ordinals, there exists an ordinal β such that $\rho(\pi(i)) = \rho(\pi(k)) = \beta$ for all $i \geq k$, which means that $\pi(i) \notin T$ for all $i \geq k$. As $\pi(k)\pi(k+1)\ldots$ is consistent with g^β and $\pi(k) \in W_1^\beta$, so π is won by Player 1.

Therefore we have shown that Player 0 has a winning strategy from all vertices in W_0^∞ and Player 1 has a winning strategy from all vertices in W_1^∞. As $V = W_0^\infty \cup W_1^\infty$, this shows that $W_0 = W_0^\infty$ and $W_1 = W_1^\infty$. Q.E.D.

We can now give the proof that parity games are indeed appropriate model checking games for LFP and L_μ.

Theorem 2.18. If $\mathfrak{A} \models \Psi(\bar{a})$, then Player 0 has a winning strategy in the game $\mathcal{G}(\mathfrak{A}, \Psi(\bar{a}))$ starting at position $\Psi(\bar{a})$.

Proof. By structural induction over $\Psi(\bar{a})$. We will only consider the interesting cases $\Psi(\bar{a}) = [\mathrm{gfp}\, T\bar{x}.\varphi(T, \bar{x})](\bar{a})$ and $\Psi(\bar{a}) = [\mathrm{lfp}\, T\bar{x}.\varphi(T, \bar{x})](\bar{a})$.

Let $\Psi(\bar{a}) = [\mathrm{gfp}\, T\bar{x}.\varphi(T, \bar{x})](\bar{a})$. In the game $\mathcal{G}(\mathfrak{A}, \Psi(\bar{a}))$, the positions $T\bar{b}$ have priority 0. Every such position has a unique successor $\varphi(T, \bar{b})$, so the unfoldings $\mathcal{G}^\alpha(\mathfrak{A}, \Psi(\bar{a}))$ are well defined.

Let us take the chain of steps of the gfp-induction of $\varphi(\bar{x})$ on $\mathfrak{A}$.

$$X^0 \supseteq X^1 \supseteq \ldots \supseteq X^\alpha \supseteq X^{\alpha+1} \supseteq \ldots$$

We have

$$\begin{aligned}
\mathfrak{A} \models \Psi(\bar{a}) \quad &\Leftrightarrow \quad \bar{a} \in \mathrm{gfp}(\varphi^{\mathfrak{A}}) \\
&\Leftrightarrow \quad \bar{a} \in X^\alpha \text{ for all ordinals } \alpha \\
&\Leftrightarrow \quad \bar{a} \in X^{\alpha+1} \text{ for all ordinals } \alpha \\
&\Leftrightarrow \quad (\mathfrak{A}, X^\alpha) \models \varphi(\bar{a}) \text{ for all ordinals } \alpha.
\end{aligned}$$

Induction hypothesis: For every $X \subset A^k$

$$(\mathfrak{A}, X) \models \varphi(\bar{b}) \text{ iff Player 0 has a winning strategy in } \mathcal{G}((\mathfrak{A}, X), \varphi(\bar{a})) \text{ from } \varphi(\bar{a}).$$

We show: If Player 0 has a winning strategy in $\mathcal{G}((\mathfrak{A}, X^\alpha), \varphi(\bar{a}))$ starting at position $\varphi(\bar{a})$, then Player 0 has a winning strategy in $\mathcal{G}^\alpha(\mathfrak{A}, \Psi(\bar{a}))$ starting at position $\varphi(\bar{a})$.

By the unfolding lemma, the second statement is true for all ordinals α if and only if Player 0 has a winning strategy in $\mathcal{G}(\mathfrak{A}, \Psi(\bar{a}))$ starting at $\varphi(\bar{a})$.

As $\varphi(\bar{a})$ is the only successor of $\Psi(\bar{a}) = [\mathrm{gfp}\, T\bar{x}.\varphi(T, \bar{x})](\bar{a})$, this holds exactly if Player 0 has a winning strategy in $\mathcal{G}(\mathfrak{A}, \Psi(\bar{a}))$ starting at $\Psi(\bar{a})$.

It remains to show that Player 0 has indeed a winning strategy in the game $\mathcal{G}((\mathfrak{A}, X^\alpha), \varphi(\bar{a}))$ starting at the position $\varphi(\bar{a})$.

There are few differences between $\mathcal{G}((\mathfrak{A}, X^\alpha), \varphi(\bar{a}))$ and the unfolding $\mathcal{G}^\alpha(\mathfrak{A}, \Psi(\bar{a}))$:

- In $\mathcal{G}^\alpha(\mathfrak{A}, \Psi(\bar{a}))$, there is an additional position $\Psi(\bar{a})$, but this position is not reachable.
- The assignment of the atomic propositions $T\bar{b}$:
  - Player 0 wins at position $T\bar{b}$ in $\mathcal{G}((\mathfrak{A}, X^\alpha), \varphi(\bar{a}))$ if and only if $\bar{b} \in X^\alpha$.
  - Player 0 wins at position $T\bar{b}$ in $\mathcal{G}^\alpha(\mathfrak{A}, \Psi(\bar{a}))$ if and only if $T\bar{b} \in T_0^\alpha$.

So we need to show using an induction over α that

$$\bar{b} \in X^\alpha \text{ iff } T\bar{b} \in T_0^\alpha.$$

Base case $\alpha = 0$: $X^0 = A^k$ and $T_\alpha^0 = T = \{T\bar{b} : \bar{b} \in A^k\}$.
Induction step $\alpha = \gamma + 1$: Then $\bar{b} \in X^\alpha = X^{\alpha+1}$ if and only if $(\mathfrak{A}, X^\gamma) \models \varphi(\bar{b})$, which in turn holds if Player 0 wins $\mathcal{G}((\mathfrak{A}, X^\gamma), \varphi(\bar{b}))$ starting at $\varphi(\bar{b})$. By induction hypothesis, this holds if and only if Player 0 wins the unfolding $\mathcal{G}^\gamma(\mathfrak{A}, \Psi(\bar{a}))$ starting at $\varphi(\bar{b}) = s(T\bar{b})$ if and only if $T\bar{b} \in T_0^{\gamma+1} = T_0^\alpha$.
Induction step with α being a limit ordinal: We have that $\bar{b} \in X^\alpha$ if $\bar{b} \in X^\gamma$ for all ordinals $\gamma < \alpha$, which holds, by induction hypothesis, if and only if $T\bar{b} \in T_0^\gamma$ for all $\gamma < \alpha$, which is equivalent to $T\bar{b} \in T_0^\alpha$.

The proof for $\Psi(\bar{a}) = [\text{lfp}\, T\bar{x}.\varphi(T, \bar{x})](\bar{a})$ is analogous. Q.E.D.

2.5 Defining Winning Regions in Parity Games

To conclude this chapter, we consider the converse question—whether winning regions in a parity game can be defined in fixed-point logic—and show that, given an appropriate representation of parity games as structures, winning regions are definable in the μ-calculus.

A parity game $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ with priorities $\Omega(V) = \{0, 1, \ldots, d-1\}$, can be described by the Kripke structure $\mathcal{K}_\mathcal{G} = (V, V_0, V_1, E, P_0, \ldots, P_{d-1})$ with atomic propositions $P_j = \{v \in V : \Omega(v) = j\}$.

Given the above representation, the μ-calculus formula

$$\varphi_d^{\mathrm{Win}} = \nu X_0.\mu X_1.\nu X_2.\dots\lambda X_{d-1} \bigvee_{j=0}^{d-1} \big((V_0 \wedge P_j \wedge \Diamond X_j) \vee (V_1 \wedge P_j \wedge \Box X_j)\big),$$

where $\lambda = \nu$ if d is odd, and $\lambda = \mu$ otherwise, defines the winning region of Player 0 in the sense of the following theorem.

Theorem 2.19. $\mathcal{K}_{\mathcal{G}}, v \models \varphi_d^{\mathrm{Win}}$ if and only if Player 0 has a winning strategy from v in $\mathcal{G}$.

Proof. The model checking game for φ_d^{Win} on $\mathcal{K}_{\mathcal{G}}$ is essentially the same as the game $\mathcal{G}$ itself, up to the elimination of 'stupid moves':

- Eliminate moves after which the opponent wins in at most two steps. For instance, Player 0 would never move to a position $(V_0 \wedge P_j \wedge \Diamond X_j, v)$ if v was not a vertex of Player 0 or did not have priority j. Similarly, Player 1 would not move to a position (P_j, v) or (V_σ, v) if $v \in P_j$ or $v \in V_\sigma$.
- Contract sequences of trivial moves and remove the intermediate positions.

A schematic view of a model checking game for φ_d^{Win} is sketched in Figure 2.2. Q.E.D.

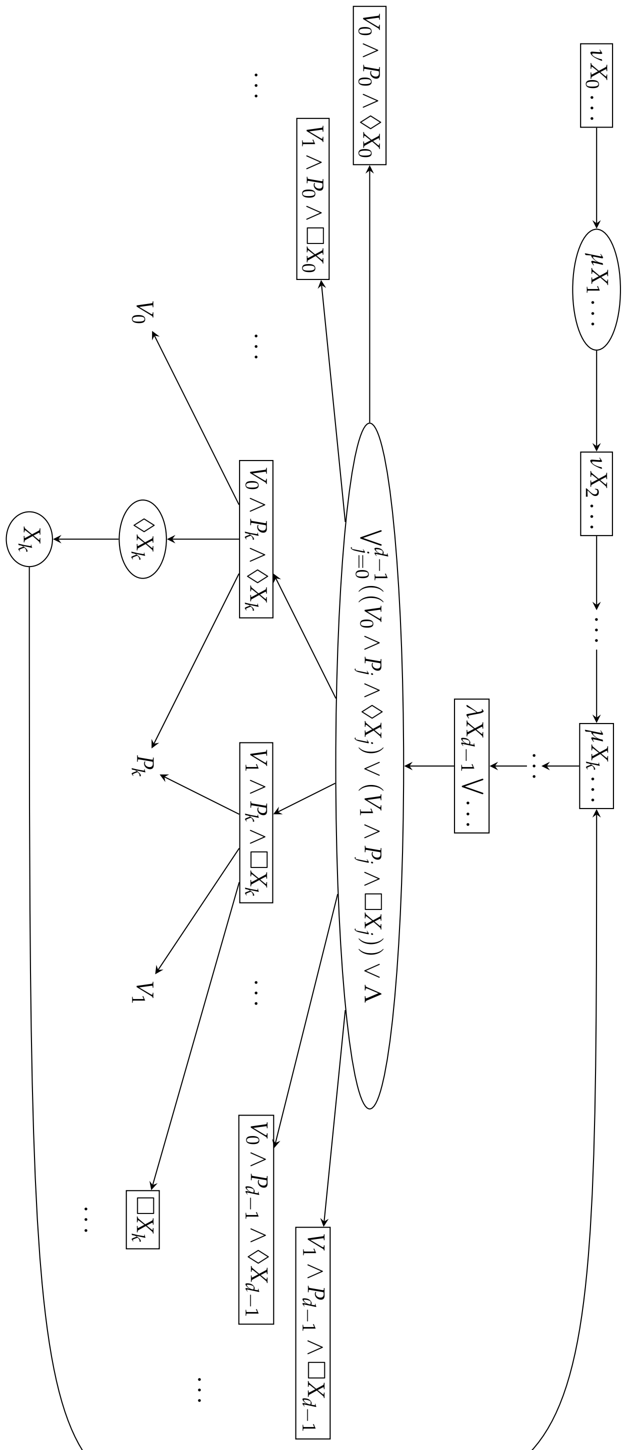

Figure 2.2. Part of a model checking game for φ_d^{Win}.

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 $\mathcal{G} = (G, \mathrm{Win})$ where $G = (V, V_0, V_1, E, \Omega)$ is a directed graph with $V = V_0 \cupdot V_1$ and $\Omega : V \to C$ for a set C of *colours* (or priorities) and a set $\mathrm{Win} \subseteq C^\omega$ 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 $\pi = v_0 v_1 v_2 \ldots \in V^{\leq\omega}$ 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 $\Omega(\pi) = \Omega(v_0)\Omega(v_1)\ldots \in \mathrm{Win}$, otherwise Player 1 wins (there are no draws). Let $\mathrm{Plays}(\mathcal{G})$ denote the set of all plays of G and $P_{\mathrm{fin}}(\mathcal{G})$ be set of all initial segments $x \in V^*$ of a play in $\mathrm{Plays}(\mathcal{G})$

3.1 Determinacy

A *strategy* for Player σ in a game $\mathcal{G} = (G, \mathrm{Win})$ is a function $f : V^* V_\sigma \to V$ 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 $\pi = v_0 v_1 \ldots$ is *consistent with a strategy* f for Player σ if for all proper prefixes $v_0 \ldots v_n$ of π such that $v_n \in V_\sigma$ we have $v_{n+1} = f(v_0 \ldots v_n)$. 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 $\tilde{W}_\sigma$ 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 wining 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 at the roots 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 prefixes y of π. 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

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

$$V_1 = \bigcup_{n\in\omega} \{0,1\}^{2n+1},$$
$$E = \{(x, xi) : x \in \{0,1\}^*, i \in \{0,1\}\},$$

and $\Omega : \{0,1\}^* \to \{0,1,\varepsilon\} : \varepsilon \mapsto \varepsilon, xi \mapsto i$. 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.

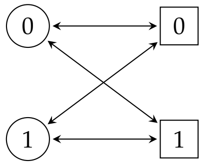

Figure 3.1. Game graph for Gale-Stewart game over $B = \{0,1\}$

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

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

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 $\mathrm{Win} := \bigcup_{\alpha<2^\omega} X_\alpha$ 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 $f_\alpha\hat{}g \notin \mathrm{Win}$, 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 $f\hat{}g_\alpha \in X_{\alpha+1} \subseteq \mathrm{Win}$, 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 $a_1 \cap \cdots \cap a_m \neq \emptyset$ for all $m \in \mathbb{N}$, $a_1, \ldots, a_m \in F$.

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 $x = x_0 x_1 \ldots \in \omega^\omega$ if

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

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

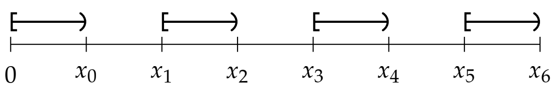

Figure 3.2. The winning condition of the ultrafilter game

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 $x_2 = f(x_0 x_1)$.
- 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 $x_3 = f(x_0 x_2)$, and Player 1 uses this in the play x as her answer to $x_0 x_1 x_2$.
- This is the iterated. In play x, Player 1 extends in her $(i+1)$st move the sequence $x_0 x_1 \dots x_{2i}$ by $x_{2i+1} = f(x_0 x_2 x_3 \dots x_{2i})$, 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 $x_0 x_2 x_3 \dots x_{2i+1}$ by $x_{2i+2} = f(x_0 x_1 \dots x_{2i+1})$, 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.

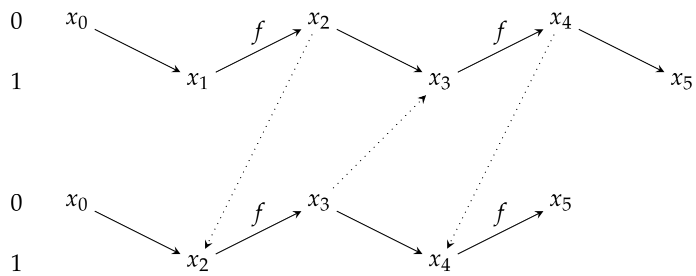

Figure 3.3. Playing the Ultrafilter game

This results in two plays $x = x_0 x_1 x_2 \dots$ and $x' = x_0 x_2 x_3 x_4 \dots$, where $x_{2i+2} = f(x_0 x_1 \dots x_{2i+1})$ but also $x_{2i+1} = f(x_0 x_1 \dots x_{2i})$. 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*.

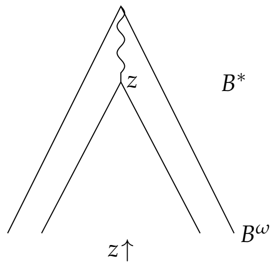

Figure 3.4. Base sets in the Cantor space

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^0_α and Π^0_α for $0 < \alpha < \omega_1$, where ω_1 is the first uncountable ordinal number.

- $\Sigma^0_1 = \mathcal{O}$;
- $\Pi^0_\alpha = \mathrm{co}\Sigma^0_\alpha := \{S \setminus X : X \in \Sigma^0_\alpha\}$ for every α;
- $\Sigma^0_\alpha = \{\bigcup_{n\in\omega} X_n : X_n \in \Pi^0_\beta$ for $\beta < \alpha\}$ for $\alpha > 0$.

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

- Σ^0_1: Open sets
- Π^0_1: Closed sets
- Σ^0_2: Countable unions of closed sets
- Π^0_2: Countable intersections of open sets
- Σ^0_3: Countable unions of Π^0_2-sets
- Π^0_3: Countable intersections of Σ^0_2-sets

Example 3.17. Let $d \in B$.

$$L_d = \{x \in B^\omega : x \text{ contains } d \text{ infinitely often}\} = \bigcap_{n\in\omega} \underbrace{(B^* \cdot d)^n \cdot B^\omega}_{\in\Sigma^0_1}.$$

Hence, $L_d \in \Pi^0_2$.

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

Definition 3.18. A function $f : B^\omega \to C^\omega$ 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 $f : B^\omega \to C^\omega$ 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:

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

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 $x = x_0 x_1 x_2 \cdots \in B^\omega$, and Player 1 has produced a finite or infinite word $y = y_0 y_1 y_2 \cdots \in C^{\leq \omega}$. 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 $f : B^\omega \to C^\omega$ 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 $v = x_0 x_1 \dots x_n \in B^*$ 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 $x_0 x_1 x_2 \dots$ by an ω-word $y_0 y_1 y_2 \dots$, but Player 1 can delay choosing y_i until he knows $x_0 x_1 \dots x_n$ for some appropriate $n \geq i$.

Choice of y_0: Consider the partition $B^\omega = \dot{\bigcup}_{c \in C} f^{-1}(c \cdot C^\omega)$. 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

$$x_0 \cdots x_n \cdot B^\omega = \dot{\bigcup}_{c \in C} f^{-1}(y_0 \cdots y_i \cdot c \cdot C^\omega).$$

Since the sets $f^{-1}(y_0 \cdots y_i \cdot c \cdot C^\omega)$ are clopen, after finitely many moves, by choosing a prolongation $x_0 \cdots x_n x_{n+1} \cdots x_k$, Player 0 has determined in which set $f^{-1}(y_0 \cdots y_i \cdot c \cdot C^\omega)$ 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 $x \in f^{-1}(y_0 \cdots y_i \cdot C^\omega)$. 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^0_α-*complete* if:

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

Π^0_α-completeness is defined analogously.

Note that Y is Σ^0_α-complete if, and only if, $C^\omega \setminus Y$ is Π^0_α-complete.

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

- $0^*1\{0,1\}^\omega$ is Σ^0_1-complete;
- $\{0^\omega\}$ is Π^0_1-complete;
- $\{0,1\}^*0^\omega$ is Σ^0_2-complete;
- $(0^*1)^\omega$ is Π^0_2-complete.

Proof. By the above remark, it suffices to show that $0^*1\{0,1\}^\omega$ and $(0^*1)^\omega$ are Σ^0_1-complete and Π^0_2-complete, respectively.

- We know that $0^*1\{0,1\}^\omega \in \Sigma^0_1$. 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 $(0^*1)^\omega \in \Pi^0_2$. Let $X = \bigcap_{n\in\omega} W_n \cdot B^\omega \in \Pi^0_2$. 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 $x_0 \ldots x_k$ 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 $\mathcal{G} = (V, V_0, V_1, E, \mathrm{Win})$ clopen, open, closed, etc., or simply a *Borel game*, if the winning condition $\mathrm{Win} \subseteq V^\omega$ 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 $\mathcal{G} = (V, V_0, V_1, E, \mathrm{Win})$ where $\mathrm{Win} = U \cdot V^\omega$ is open. First, we consider finite plays: Let $T_\sigma = \{v \in V_{1-\sigma} : vE = \emptyset\}$ and $A_\sigma = \mathrm{Attr}_\sigma(T_\sigma)$. From every position $v \in A_\sigma$ Player σ wins after finitely many moves with the attractor strategy.

For the infinite plays consider

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

with positions $V' := V \setminus (A_0 \cup A_1)$. In $\mathcal{G}'$ every play is infinite, and Player 0 wins $\pi = v_0 v_1 v_2 \dots$ 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 $v_0 v_1 \dots v_n \in U$. 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

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

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

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

Hence, the winning condition is the set

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

a finite Boolean combination of Π^0_2-sets.

- *Parity conditions* (see previous chapter) are special cases of Muller conditions and thus also finite Boolean combinations of Π^0_2-sets.
- Every ω-regular language is a Boolean combination of Π^0_2-sets. This follows from the recognisability of ω-regular languages by Muller automata and the fact that Muller conditions are Boolean combinations of Π^0_2-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

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

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^0_2-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 $G = (V_0, V_1, E, \Omega : V \to C)$ with a winning condition depending only on the set of priorities seen infinitely often in a play. It is specified by a partition $\mathcal{P}(C) = \mathcal{F}_0 \uplus \mathcal{F}_1$, and a play $\pi = v_0 v_1 v_2 \dots$ is won by Player σ if

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

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

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

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 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.

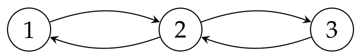

Figure 3.5. A solitaire Muller game

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 $\mathfrak{M} = (M, \mathrm{update}, \mathrm{init})$, where M is a set of *memory states*, update : $M \times V \to M$ is a *memory update function* and init : $V \to M$ 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 $F : V_\sigma \times M \to V$ 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 $F : V_\sigma \to V$.

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:

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

Player 0 has a winning strategy from every position, but no positional one. Assume that $f : \{a, b\} \to \{1, 2\}$ is a positional winning

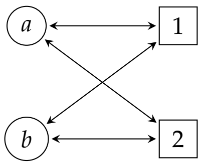

Figure 3.6. Muller game $\mathcal{G} = DJW_2$

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 $\mathrm{Inf}(\pi) = \{a, 2\}$ (respectively $\mathrm{Inf}(\pi) = \{b, 2\}$). If $f(a) = f(b) = 1$, then Player 1 alternates between a and b and wins, since this generates a play π with $\mathrm{Inf}(\pi) = \{a, b, 1\}$. However, Player 0 has a winning strategy that uses the memory depicted in Figure 3.7. The corresponding strategy is defined as follows:

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

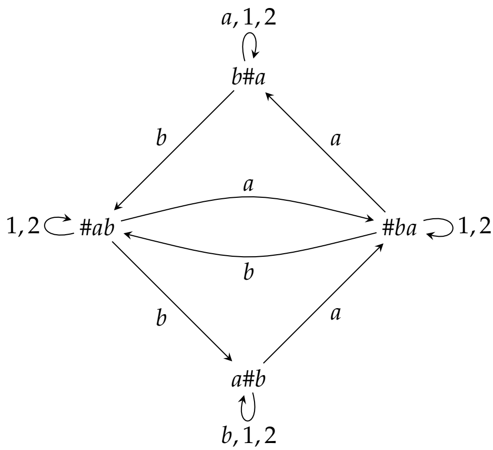

Figure 3.7. Memory for Player 0

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 $F(0, m) = m + 1$.

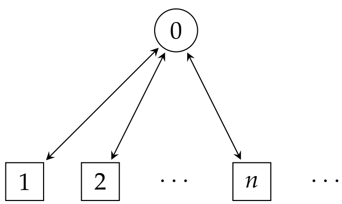

Figure 3.8. A game where finite-memory strategies do not suffice

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

Definition 3.35. For games $\mathcal{G} = (G, \Omega, \mathrm{Win})$ and $\mathcal{G}' = (G', \Omega', \mathrm{Win}')$, 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

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

with the update function

$$\mathrm{update}^*((m, m'), v) = \\ (\mathrm{update}(m, v), \mathrm{update}'(m', (v, \mathrm{update}(m, v)))).$$

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 $(v_0, \mathrm{init}(v_0)))$. 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 $F' : (V_\sigma \times M) \times M' \to (V \times M)$ for Player σ in $\mathcal{G}'$, we have to construct a strategy $F : (V_\sigma \times (M \times M')) \to V$ for Player σ in $\mathcal{G}$. For any $v \in V_\sigma$ and any pair $(m, m') \in M \times M'$ we have that $F'((v, m), m') = (w, \mathrm{update}(m, w))$ 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 $n = \mathrm{update}(m, w)$ and $n' = \mathrm{update}'(m', (w, n))$. 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 $(v_0, \mathrm{init}(v_0))$ that is consistent with F'.

Therefore, if F' is a winning strategy from $(v_0, \mathrm{init}(v_0))$, 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 $\mathfrak{M} = (M, \mathrm{update}, \mathrm{init})$ and $\mathcal{G}'$ reduces to $\mathcal{G}''$ with memory $\mathfrak{M}' = (M', \mathrm{update}', \mathrm{init}')$ then $\mathcal{G}$ reduces to $\mathcal{G}''$ with memory $(M \times M', \mathrm{update}'', \mathrm{init}'')$ where

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

and

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

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, $\mathrm{LAR}(C)$ is the set of sequences $c_1 \dots c_k \# c_{k+1} \dots c_l$ of elements from $C \mathbin{\dot\cup} \{\#\}$ 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 $c_1 c_2 c_3 \# c_4 c_5$ is updated to $c_1 \# c_3 c_4 c_5 c_2$. (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 $\Omega : V \to C$ is the triple $(\mathrm{LAR}(C), \mathrm{update}, \mathrm{init})$ with $\mathrm{init}(v) = \#\Omega(v)$ and

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

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 $\pi =$

$v_0v_1\dots$ 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 $\mathrm{Inf}(\pi)$ 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 $\mathrm{Inf}(\pi)$ and infinitely often coincides with $\mathrm{Inf}(\pi)$.

Proof. For each play $\pi = v_0v_1v_2\dots$ there is a position v_m such that $\Omega(v_n) \in \mathrm{Inf}(\pi)$ for all $n \geq m$. Since no priority outside $\mathrm{Inf}(\pi)$ is seen after position v_m, the hit set will, from that position onwards, always be contained in $\mathrm{Inf}(\pi)$, 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 $\mathrm{Inf}(\pi) = \{c_j,\dots,c_l\}$. Since all priorities in $\mathrm{Inf}(\pi)$ 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 $c_1,\dots,c_{j-1}\#c_{j+1}\dots c_lc_j$, and the hit set coincides with $\mathrm{Inf}(\pi)$. 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 $\Omega : V \to C$ and winning condition $(\mathcal{F}_0, \mathcal{F}_1)$. We have to prove that $\mathcal{G} \leq_{\mathrm{LAR}} \mathcal{G}'$ for a parity game $\mathcal{G}'$ with game graph $G \times \mathrm{LAR}(C)$ and an appropriate priority labelling Ω' on $V \times \mathrm{LAR}(C)$, which is defined as follows:

$$\Omega'(v, c_1c_2\dots c_k\#c_{k+1}\dots c_l) = \begin{cases} 2k & \text{if } \{c_{k+1},\dots,c_l\} \in \mathcal{F}_0, \\ 2k+1 & \text{if } \{c_{k+1},\dots,c_l\} \in \mathcal{F}_1. \end{cases}$$

Let $\pi = v_0v_1v_2\dots$ be a play on $\mathcal{G}$ and fix a number m such that, for all $n \geq m$, $\Omega(v_n) \in \mathrm{Inf}(\pi)$ and the LAR at position v_n has the form $c_1\dots c_jc_{j+1}\dots c_k\#c_{k+1}\dots c_l$ where $\mathrm{Inf}(\pi) = \{c_{j+1},\dots c_l\}$ 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., $\mathrm{Inf}(\pi) \in \mathcal{F}_\sigma$. Since the hit set of the LAR coincides with $\mathrm{Inf}(\pi)$ 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.

A set $Y \subseteq C$ belongs to $\mathcal{F}_\sigma$ 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

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

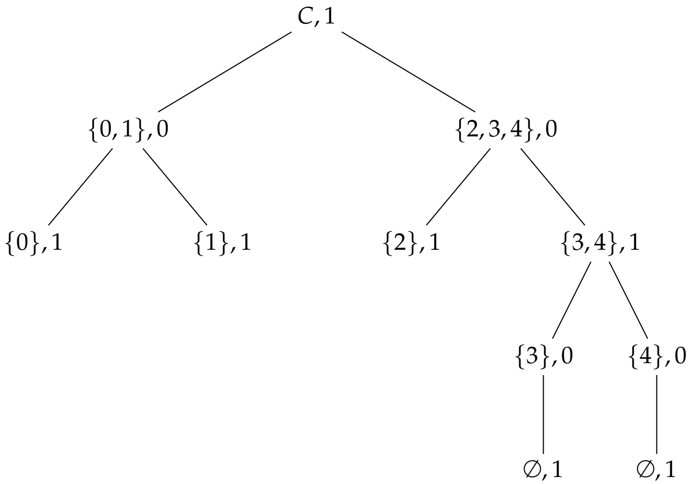

Figure 3.9. A Zielonka tree

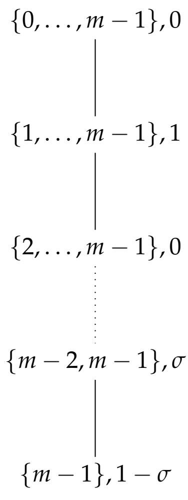

Figure 3.10. The Zielonka tree of a parity-condition with m priorities

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.

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 $Z = \mathrm{Attr}_1(Y) \setminus Y$. 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 $X_0 \cup W'_0 \supsetneq X_0$, a contradiction to the definition of X_0. Hence $W'_1 = V'$.

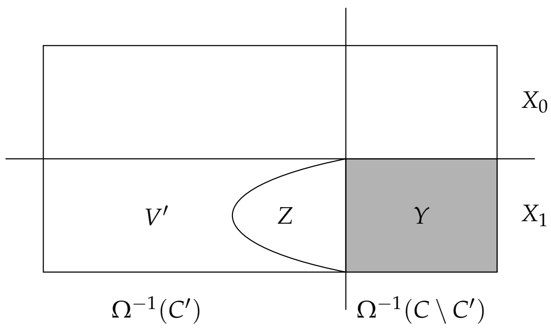

Figure 3.11. Constructing a winning strategy for Player 1

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, $\mathrm{Inf}(\pi) \cap (C \setminus C') \neq \emptyset$ and $\mathrm{Inf}(\pi) \in \mathcal{F}_1$. So Player 1 has a positional winning strategy on X_1.

Case 2: $C \in \mathcal{F}_0$. There exist maximal subsets $C_0, \ldots, C_{k-1} \subseteq C$ 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 $f_0, \ldots, f_{k-1}$ 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 $M_0 \cup \cdots \cup M_{k-1}$.

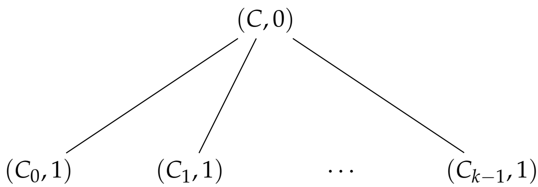

Figure 3.12. The top of the Zielonka tree $Z(\mathcal{F}_0, \mathcal{F}_1)$

For $i = 0, \ldots, k-1$, let $Y_i = X_0 \cap \Omega^{-1}(C \setminus C_i)$, and $Z_i = \mathrm{Attr}_0(Y_i) \setminus Y_i$, 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 $\widetilde{Y} := \bigcap_{i<k} Y_i$, Player 0 just plays with a (positional) trap strategy t ensuring that the play remains in X_0. At the first position $v \notin \widetilde{Y}$, 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 $\widetilde{Y}$.

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 $M = M_0 \cup \cdots \cup M_{k-1}$. 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 $Y_0, \ldots, Y_{k-1}$ infinitely often. But this means that $\mathrm{Inf}(\pi) \cap (C \setminus C_i) \neq \emptyset$ for all $i \leq k$; as observed above this implies that $\mathrm{Inf}(\pi) \in \mathcal{F}_0$. 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 NP $\cap$ 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 $\mathcal{G}_\Psi$ as follows: Positions for Player 0 are the literals $X_1, \ldots, X_k, \neg X_1, \ldots, \neg X_k$ occurring in Ψ; positions for Player 1 are the clauses $C_1, \ldots, C_n$. 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 $\mathcal{G}_\Psi$ (from any initial position).

($\Rightarrow$) Assume that Ψ is satisfiable. There exists a satisfying interpretation $\mathcal{I} : \{X_1, \ldots, X_k\} \to \{0, 1\}$. Player 1 moves from a clause C to a literal $Y \in C$ such that $[\![Y]\!]^\mathcal{I} = 1$. In the resulting play only literals with $[\![Y]\!]^\mathcal{I} = 1$ 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 $\{(L_1, R_1), \ldots, (L_k, R_k)\}$ 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 $\{(L_1, R_1), \ldots, (L_{2m}, R_{2m})\}$ 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 $\mathcal{P} = \{(L_1, R_1), \ldots, (L_k, R_k)\}$ 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, $\mathrm{Attr}_0(U) \subseteq W_0$. 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

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

is visited again and again. Thus we will reduce the game from U to $U \setminus \widetilde{U}$ 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.

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 $W_0 = \mathrm{WinReg}(G, \mathcal{P})$:

($\subseteq$) 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 $v \in \mathrm{WinReg}(G, \mathcal{P})$. Otherwise,

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**
  $\mathcal{P}' := \{(L_i, R_i) : \Omega(U) \cap L_i \neq \emptyset\}$;
  $\widetilde{U} := \{u \in U : \Omega(u) \in L_i$ for some i such that $\Omega(U) \cap R_i = \emptyset\}$;
  if $\widetilde{U} = \emptyset$ **then** $W := W \cup U$;
  else $W := W \cup \text{WinReg}(G \restriction_{U \setminus \widetilde{(U)}}, \mathcal{P}')$;
enddo;
$W_0 := \text{Attr}_0(W)$;
Output W_0;

$S \subseteq U \setminus \widetilde{U}$, and S is contained in an SCC of $G \restriction_{U \setminus \widetilde{U}}$. 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.

4 Basic Concepts of Mathematical Game Theory

Up to now we considered finite or infinite games

- with two players,
- played on finite or infinite graphs,
- with perfect information (the players know the whole game, the history of the play and the actual position),
- with qualitative (win or loss) winning conditions (zero-sum games),
- with ω-regular winning conditions (or Borel winning conditions) specified in a suitable logic or by automata, and
- with asynchronous interaction (turn-based games).

Those games are used for verification or to evaluate logic formulae.

In this section we move to concurrent multi-player games in which players get real-valued *payoffs*. The games will still have perfect information and additionally throughout this chapter we assume that the set of possible plays is *finite*, so there exist only finitely many strategies for each of the players.

4.1 Games in Strategic Form

Definition 4.1. A *game in strategic form* is described by a tuple $\Gamma = (N, (S_i)_{i \in N}, (p_i)_{i \in N})$ where

- $N = \{1, \dots, n\}$ is a finite set of players
- S_i is a set of *strategies* for Player i
- $p_i : S \to \mathbb{R}$ is a *payoff function* for Player i

and $S := S_1 \times \cdots \times S_n$ is the set of *strategy profiles*. Γ is called a *zero-sum game* if $\sum_{i \in N} p_i(s) = 0$ for all $s \in S$.

The number $p_i(s_1, \ldots, s_n)$ is called the *value* or *utility* of the strategy profile $(s_1, \ldots, s_n)$ for Player i. The intuition for zero-sum games is that the game is a closed system.

Many important notions can best be explained by two-player games, but are defined for arbitrary multi-player games.

In the sequel, we will use the following notation: Let Γ be a game. Then $S_{-i} := S_1 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_n$ is the set of all strategy profiles for the players except i. For $s \in S_i$ and $s_{-i} \in S_{-i}$, (s, s_{-i}) is the strategy profile where Player i chooses the strategy s and the other players choose s_{-i}.

Definition 4.2. Let $s, s' \in S_i$. Then s *dominates* s' if

- for all $s_{-i} \in S_{-i}$ we have $p_i(s, s_{-i}) \geq p_i(s', s_{-i})$, and
- there exists $s_{-i} \in S_{-i}$ such that $p_i(s, s_{-i}) > p_i(s', s_{-i})$.

A strategy s is *dominant* if it dominates every other strategy of the player.

Definition 4.3. An *equilibrium in dominant strategies* is a strategy profile $(s_1, \ldots, s_n) \in S$ such that all s_i are dominant strategies.

Definition 4.4. A strategy $s \in S_i$ is a *best response* to $s_{-i} \in S_{-i}$ if $p_i(s, s_{-i}) \geq p_i(s', s_{-i})$ for all $s' \in S_i$.

Obviously, a dominant strategy is a best response to all strategy profiles of the other players.

Example 4.5. **The Prisoner's Dilemma.**

Two suspects are arrested, but there is insufficient evidence for a conviction. Both prisoners are questioned separately, and are offered the same deal: if one testifies for the prosecution against the other and the other remains silent, the betrayer goes free and the silent accomplice receives the full 10-year sentence. If both stay silent, both prisoners are sentenced to only one year in jail for a minor charge. If both betray each other, each receives a five-year sentence. So this dilemma poses the question: How should the prisoners act?

	stay silent	betray
stay silent	$(-1, -1)$	$(-10, 0)$
betray	$(0, -10)$	$(-5, -5)$

An entry (a, b) at position i, j of the matrix means that if profile (i, j) is chosen, Player 1 (who chooses the rows) receives payoff a and Player 2 (who chooses the columns) receives payoff b.

Betraying is a dominant strategy for every player, call this strategy b. Therefore, (b, b) is an equilibrium in dominant strategies. Problem: The payoff $(-5, -5)$ of the dominant equilibrium is not optimal.

The Prisoner's Dilemma is an important metaphor for many decision situations, and there exists extensive literature concerned with the problem. Especially interesting is the situation, where the Prisoner's Dilemma is played repeatedly, possibly infinitely often.

Example 4.6. **Battle of the sexes.**

	meat	fish
red wine	$(2,1)$	$(0,0)$
white wine	$(0,0)$	$(1,2)$

There are no dominant strategies, and thus there is no dominant equilibrium. The pairs (red wine, meat) and (white wine, fish) are distinguished since every player plays with a best response against the strategy of the other player: No player would change his or her strategy unilaterally.

4.2 Nash equilibria

Definition 4.7. A strategy profile $s = (s_1, \ldots, s_n) \in S$ is a *Nash equilibrium* in Γ if for all $i \in N$ and all strategies $s'_i \in S_i$

$$p_i(\underbrace{s_i, s_{-i}}_{s}) \geq p_i(s', s_{-i}).$$

Thus, in a Nash equilibrium, every player plays with a best response to the profile of his opponents, and thus has no incentive to deviate unilaterally to a different strategy. Is there a Nash equilibrium in every game? The following example shows that this is not always the case, at least not in pure strategies.

Example 4.8. **Rock, paper, scissors.**

	rock	scissors	paper
rock	$(0,0)$	$(1,-1)$	$(-1,1)$
scissors	$(-1,1)$	$(0,0)$	$(1,-1)$
paper	$(1,-1)$	$(-1,1)$	$(0,0)$

There are no dominant strategies and no Nash equilibria: For every pair (f,g) of strategies one of the players can change to a better strategy. Note that this game is a zero-sum game.

Although there are no Nash equilibria in pure strategies in rock, paper, scissors, there is of course an obvious good method to play this game: Randomly pick one of the three actions with equal probability. This observation leads us to the notion of mixed strategies, where the players are allowed to randomise over strategies.

Definition 4.9. A *mixed strategy* of Player i in Γ is a probability distribution $\mu_i : S_i \to [0,1]$ on S_i (so that $\sum_{s\in S_i} \mu(s) = 1$).
$\Delta(S_i)$ denotes the set of probability distributions on S_i. $\Delta(S) := \Delta(S_1) \times \cdots \times \Delta(S_n)$ is the set of all strategy profiles in mixed strategies. The expected payoff is $\widehat{p}_i : \Delta(S) \to \mathbb{R}$,

$$\widehat{p}_i(\mu_1, \ldots, \mu_n) = \sum_{(s_1,\ldots,s_n)\in S} \left(\prod_{j\in N} \mu_j(s_j) \right) \cdot p_i(s_1, \ldots, s_n)$$

For every game $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$ we define the *mixed expansion* $\widehat{\Gamma} = (N, (\Delta(S_i))_{i\in N}, (\widehat{p}_i)_{i\in N})$.

Definition 4.10. A *Nash equilibrium of* Γ *in mixed strategies* is a Nash equilibrium in $\widehat{\Gamma}$, i.e. a Nash equilibrium in Γ in mixed strategies is a mixed strategy profile $\overline{\mu} = (\mu_1, \ldots, \mu_n) \in \Delta(S)$ such that, for every player i and every $\mu'_i \in \Delta(S)$, $\widehat{p}_i(\mu_i, \mu_{-i}) \geq \widehat{p}_i(\mu'_i, \mu_{-i})$.

Nash equilibria (in mixed strategies) provide the arguably most important solution concept in classical game theory (although, as we shall point out later, this concept is not without problems). An important reason for the success of Nash equilibrium as a solution concept is the fact that every finite game has one. To prove this, we shall use a well-known classical fixed-point theorem.

Theorem 4.11 (Brouwer's Fixed-Point Theorem)**.** Let $X \subseteq \mathbb{R}^n$ be compact (i.e., closed and bounded) and convex. Then every continuous function $f : X \to X$ has a fixed point.

We do not prove this here but remark that, interestingly, the Brouwer Fixed-Point Theorem can itself be proved via a game-theoretic result, namely the determinacy of HEX.

Theorem 4.12 (Nash)**.** Every finite game Γ in strategic form has at least one Nash equilibrium in mixed strategies.

Proof. Let $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$. Every mixed strategy of Player i is a tuple $\mu_i = (\mu_{i,s})_{s\in S_i} \in [0,1]^{|S_i|}$ such that $\sum_{s\in S_i} \mu_{i,s} = 1$. Thus, $\Delta(S_i) \subseteq [0,1]^{|S_i|}$ is a compact and convex set, and the same applies to $\Delta(S) = \Delta(S_1) \times \cdots \times \Delta(S_n)$ for $N = \{1, \ldots, n\}$. For every $i \in N$, every pure strategy $s \in S_i$ and every mixed strategy profile $\overline{\mu} \in \Delta(S)$ let

$$g_{i,s}(\overline{\mu}) := \max\left(\widehat{p}_i(s, \overline{\mu}_{-i}) - \widehat{p}_i(\overline{\mu}), 0\right)$$

be the gain of Player i if he unilaterally changes from the mixed profile $\overline{\mu}$ to the pure strategy s (only if this is reasonable).

Note that if $g_{i,s}(\overline{\mu}) = 0$ for all i and all $s \in S_i$, then $\overline{\mu}$ is a Nash equilibrium. We define the function

$$\begin{aligned} f : \Delta(S) &\to \Delta(S) \\ \overline{\mu} &\mapsto f(\overline{\mu}) = (\nu_1, \ldots, \nu_n) \end{aligned}$$

where $\nu_i : S_i \to [0,1]$ is a mixed strategy defined by

$$\nu_{i,s} = \frac{\mu_{i,s} + g_{i,s}(\overline{\mu})}{1 + \sum_{s\in S_i} g_{i,s}(\overline{\mu})}.$$

For every Player i and all $s \in S_i$, $\overline{\mu} \mapsto \nu_{i,s}$ is continuous since $\widehat{p}_i$ is continuous and thus $g_{i,s}$, too. $f(\overline{\mu}) = (\nu_1, \ldots, \nu_n)$ is in $\Delta(S)$: Every

$\nu_i = (\nu_{i,s})_{s\in S_i}$ is in $\Delta(S_i)$ since

$$\sum_{s\in S_i} \nu_{i,s} = \frac{\sum_{s\in S_i} \mu_{i,s} + \sum_{s\in S_i} g_{i,s}(\overline{\mu})}{1+\sum_{s\in S_i} g_{i,s}(\overline{\mu})} = \frac{1+\sum_{s\in S_i} g_{i,s}(\overline{\mu})}{1+\sum_{s\in S_i} g_{i,s}(\overline{\mu})} = 1.$$

By the Brouwer fixed point theorem f has a fixed point. Thus, there is a $\overline{\mu} \in \Delta(S)$ such that

$$\mu_{i,s} = \frac{\mu_{i,s} + g_{i,s}(\overline{\mu})}{1+\sum_{s\in S_i} g_{i,s}(\overline{\mu})}$$

for all i and all s.

Case 1: There is a Player i such that $\sum_{s\in S_i} g_{i,s}(\overline{\mu}) > 0$.
Multiplying both sides of the fraction above by the denominator, we get $\mu_{i,s} \cdot \sum_{s\in S_i} g_{i,s}(\mu) = g_{i,s}(\overline{\mu})$. This implies $\mu_{i,s} = 0 \;\Leftrightarrow\; g_{i,s}(\overline{\mu}) = 0$, and thus $g_{i,s}(\overline{\mu}) > 0$ for all $s \in S_i$ where $\mu_{i,s} > 0$.

But this leads to a contradiction: $g_{i,s}(\overline{\mu}) > 0$ means that it is profitable for Player i to switch from (μ_i, μ_{-i}) to (s, μ_{-i}). This cannot be true for all s where $\mu_{i,s} > 0$ since the payoff for (μ_i, μ_{-i}) is the mean of the payoffs (s, μ_{-i}) with arbitrary $\mu_{i,s}$. However, the mean cannot be smaller than all components:

$$\begin{aligned}
\widehat{p}_i(\mu_i, \mu_{-i}) &= \sum_{s\in S_i} \mu_{i,s} \cdot \widehat{p}_i(s, \mu_{-i}) \\
&= \sum_{\substack{s\in S_i \\ \mu_{i,s}>0}} \mu_{i,s} \cdot \widehat{p}_i(s, \mu_{-i}) \\
&> \sum_{\substack{s\in S_i \\ \mu_{i,s}>0}} \mu_{i,s} \cdot \widehat{p}_i(\mu_i, \mu_{-i}) \\
&= \widehat{p}_i(\mu_i, \mu_{-i})
\end{aligned}$$

which is a contradiction.

Case 2: $g_{i,s}(\overline{\mu}) = 0$ for all i and all $s \in S_i$, but this already means that $\overline{\mu}$ is a Nash equilibrium as stated before. Q.E.D.

The *support* of a mixed strategy $\mu_i \in \Delta(S_i)$ is $\operatorname{supp}(\mu_i) = \{s \in S_i : \mu_i(s) > 0\}$.

Theorem 4.13. Let $\mu^* = (\mu_1, \ldots, \mu_n)$ be a Nash equilibrium in mixed strategies of a game Γ. Then for every Player i and every pure strategy $s, s' \in \operatorname{supp}(\mu_i)$

$$\widehat{p}_i(s, \mu_{-i}) = \widehat{p}_i(s', \mu_{-i}).$$

Proof. Assume $\widehat{p}_i(s, \mu_{-i}) > \widehat{p}_i(s', \mu_{-i})$. Then Player i could achieve a higher payoff against μ_{-i} if she played s instead of s': Define $\tilde{\mu}_i \in \Delta(S_i)$ as follows:

- $\tilde{\mu}_i(s) = \mu_i(s) + \mu_i(s')$,
- $\tilde{\mu}_i(s') = 0$,
- $\tilde{\mu}_i(t) = \mu_i(t)$ for all $t \in S_i - \{s, s'\}$.

Then

$$\begin{aligned}\widehat{p}_i(\tilde{\mu}_i, \mu_{-i}) &= \widehat{p}_i(\mu_i, \mu_{-i}) + \underbrace{\mu_i(s')}_{>0} \cdot \underbrace{\left(\widehat{p}_i(s, \mu_{-i}) - \widehat{p}_i(s', \mu_{-i})\right)}_{>0} \\ &> \widehat{p}_i(\mu_i, \mu_{-i})\end{aligned}$$

which contradicts the fact that μ is a Nash equilibrium. Q.E.D.

4.3 Two-person zero-sum games

We want to apply Nash's Theorem to two-person games. First, we note that in every game $\Gamma = (\{0,1\}, (S_0, S_1), (p_0, p_1))$

$$\max_{f \in \Delta(S_0)} \min_{g \in \Delta(S_1)} p_0(f, g) \leq \min_{g \in \Delta(S_1)} \max_{f \in \Delta(S_0)} p_0(f, g).$$

The maximal payoff which one player can enforce cannot exceed the minimal payoff the other player has to cede. This is a special case of the general observation that for every function $h : X \times Y \to \mathbb{R}$

$$\sup_x \inf_y h(x, y) \leq \inf_y \sup_x h(x, y).$$

(For all x', y: $h(x', y) \leq \sup_x h(x, y)$. Thus $\inf_y h(x', y) \leq \inf_y \sup_x h(x, y)$ and $\sup_x \inf_y h(x, y) \leq \inf_y \sup_x h(x, y)$.)

Remark 4.14. Another well-known special case from mathematical logic is that $\exists x \, \forall y \, Rxy \models \forall y \, \exists x \, Rxy$.

Theorem 4.15 (v. Neumann, Morgenstern)**.**
Let $\Gamma = (\{0,1\}, (S_0, S_1), (p, -p))$ be a two-person zero-sum game. For every Nash equilibrium (f^*, g^*) in mixed strategies

$$\max_{f \in \Delta(S_0)} \min_{g \in \Delta(S_1)} p(f,g) = p(f^*, g^*) = \min_{g \in \Delta(S_1)} \max_{f \in \Delta(S_0)} p(f,g).$$

In particular, all Nash equilibria have the same payoff which is called the *value* of the game. Furthermore, both players have optimal strategies to realise this value.

Proof. Since (f^*, g^*) is a Nash equilibrium, for all $f \in \Delta(S_0)$, $g \in \Delta(S_1)$

$$p(f^*, g) \geq p(f^*, g^*) \geq p(f, g^*).$$

Thus

$$\min_{g \in \Delta(S_1)} p(f^*, g) = p(f^*, g^*) = \max_{f \in \Delta(S_1)} p(f, g^*).$$

So

$$\max_{f \in \Delta(S_0)} \min_{g \in \Delta(S_1)} p(f,g) \geq p(f^*, g^*) \geq \min_{g \in \Delta(S_1)} \max_{f \in \Delta(S_0)} p(f,g)$$

and

$$\max_{f \in \Delta(S_0)} \min_{g \in \Delta(S_1)} p(f,g) \leq \min_{g \in \Delta(S_1)} \max_{f \in \Delta(S_0)} p(f,g)$$

imply the claim. Q.E.D.

4.4 Regret minimization

To motivate the concept of regret minimization we consider

Example 4.16. **Traveller's Dilemma.** This is a symmetric two-player game $\Gamma = (\{1,2\}, (S_1, S_2), (p_1, p_2))$ with $S_1 = S_2 = \{2, \ldots, 100\}$ and

$$p_1(x,y) = \begin{cases} x+2 & \text{if } x < y, \\ y-2 & \text{if } y < x, \\ x & \text{if } x = y, \end{cases} \qquad p_2(x,y) = p_1(y,x)$$

The only Nash equilibrium in pure strategies is $(2,2)$ since for each (i,j) with $i \neq j$ the player that has chosen the greater number, say i, can do better by switching to $j-1$, and also, for every (i,i) with $i > 2$ each player can do better by playing $i-1$ (and getting the payoff $i+1$ then). Also most other solution concepts from game theory (such as the iterated elimination of dominated strategies discussed in the next section) suggest that the players should choose 2.

However, experiments show that people (even game theorists!) tend to select large numbers, in the range between 90 and 100; moreover they seem right to do so, since they perform much better in these experiments than those who follow what game theory proposes and select the strategy 2.

The question arises whether there are alternative solution concepts that justify the choice of large strategies in the Traveller's Dilemma, and if yes, which one. A relatively recent proposal that seems to achieve this is *regret minimization*. When a player uses this concept, he wants to minimize the lost payoff (which he would "regret") due to not playing with the best response to the strategies of the other players.

This idea was formulated in the context of decision theory, concerned with the choices of individual agents rather than the interaction of different agents as in game theory. Accordingly, the payoff is determined by a binary function $p : S \times Z \to \mathbb{R}$, where S is the set of strategies of the player we are considering, and Z is an abstract set of possible *states*.

Before we can introduce regret minimization, we need several definitions. In state $z \in Z$, the maximal payoff for our player is

$$p^*(z) := \max_{s \in S} p(s,z),$$

and if the player chooses the strategy $s \in S$, he will miss the following payoff:

$$\text{regret}_p(s,z) := p^*(z) - p(s,z).$$

The overall maximal regret for the strategy s is

$$\text{maxreg}_p(s) := \max_{z \in Z} \text{regret}_p(s,z).$$

Now, the decision with respect to regret minimization would be: *Choose $s \in S$ such that* $\text{maxreg}_p(s)$ *is minimal.*

Let us reconsider Example 4.16. Since it belongs to game theory, Z is the set of strategy profiles of the other players. We claim that exactly the strategies $s \in \{96, \ldots, 100\}$ minimize the maximal regret. To see this, note that for those s, we have that $\text{maxreg}_p(s) = 3$, since

- if $t \leq s$, then $p(s,t) \geq t-2$ and $p^*(t) \leq t+1$, thus $\text{regret}_p(s,t) = p^*(t) - p(s,t) \leq t+1-(t-2) = 3$,
- if $t > s$, then $p(s,t) = s+2$ and $p^*(t) \leq 101$, thus $\text{regret}_p(s,t) \leq 101-(s+2) = 99-s \leq 3$,

and on the other hand,

- $\text{regret}_p(96, 100) = 101 - 98 = 3$,
- for $s \in \{97, \ldots, 100\}$, $\text{regret}_p(s, 96) = 97 - 94 = 3$.

Also, for $s \leq 95$, we have that $\text{maxreg}_p(s) \geq 4$, as $\text{maxreg}_p(s) \geq \text{regret}_p(s, 100) = 101 - (s+2) = 99 - s \geq 4$.

Consequently, regret minimization suggests a strategy s with $96 \leq s \leq 100$. We will now iterate this idea. If both players eliminate strategies which do not minimize the regret, we obtain a subgame with strategies $\{96, \ldots, 100\}$. In this game, we have that

- $\text{maxreg}_p(97) = 2$, since
  - $\text{regret}_p(97, 100) = 101 - 99 = 2$,
  - $\text{regret}_p(97, 99) = 100 - 99 = 1$,
  - $\text{regret}_p(97, 98) = 99 - 99 = 0$,

  - $\text{regret}_p(97, 97) = 98 - 97 = 1$,
  - $\text{regret}_p(97, 96) = 96 - 95 = 2$.
- $\text{maxreg}_p(100) \geq \text{regret}_p(100, 99) = 100 - 97 = 3$.
- $\text{maxreg}_p(99) \geq \text{regret}_p(99, 98) = 99 - 96 = 3$.
- $\text{maxreg}_p(98) \geq \text{regret}_p(98, 97) = 98 - 95 = 3$.
- $\text{maxreg}_p(96) \geq \text{regret}_p(96, 100) = 101 - 98 = 3$.

Hence, 97 is the unique strategy which minimizes the regret in this subgame and thus is the choice of a player who assumes that his opponent wants to minimize his regret as well.

4.5 Iterated Elimination of Dominated Strategies

Besides Nash equilibria and (iterated) regret minimization, the iterated elimination of dominated strategies is a promising solution concept for strategic games which is inspired by the following ideas. Assuming that each player behaves rational in the sense that he will not play a strategy that is dominated by another one, dominated strategies may be eliminated. Assuming further that it is common knowledge among the players that each player behaves rational, and thus discards some of her strategies, such elimination steps may be iterated as it is possible that some other strategies become dominated due to the elimination of previously dominated strategies. Iterating these elimination steps eventually yields a fixed point where no strategies are dominated.

Example 4.17.

$$
\underbrace{\begin{array}{ccc} & L & R \\ T & (1,0,1) & \boxed{(1,1,0)} \\ B & (1,1,1) & (0,0,1) \end{array}}_{X}
\quad\Bigg|\quad
\underbrace{\begin{array}{cc} L & R \\ (1,0,1) & (0,1,0) \\ (1,1,1) & (1,0,0) \end{array}}_{Y}
$$

Player 1 picks rows, Player 2 picks columns, and Player 3 picks matrices.

- No row dominates the other (for Player 1);
- no column dominates the other (for Player 2);

- matrix X dominates matrix Y (for Player 3).

Thus, matrix Y is eliminated.

- In the remaining game, the upper row dominates the lower one (for Player 1).

Thus, the lower row is eliminated.

- Of the remaining two possibilities, Player 2 picks the better one.

The only remaining profile is (T, R, X).

There are different variants of strategy elimination that have to be considered:

- dominance by *pure* or *mixed* strategies;
- (*weak*) dominance or *strict* dominance;
- dominance by strategies in the *local* subgame or by strategies in the *global* game.

The possible combinations of these parameters give rise to eight different operators for strategy elimination that will be defined more formally in the following.

Let $\Gamma = (N, (S_i)_{i \in N}, (p_i)_{i \in N})$ such that S_i is finite for every Player i. A subgame is defined by $T = (T_1, \ldots, T_n)$ with $T_i \subseteq S_i$ for all i. Let $\mu_i \in \Delta(S_i)$, and $s_i \in S_i$. We define two notions of dominance:

(1) Dominance with respect to T:
$\mu_i >_T s_i$ if and only if
  - $p_i(\mu_i, t_{-i}) \geq p_i(s_i, t_{-i})$ for all $t_{-i} \in T_{-i}$
  - $p_i(\mu_i, t_{-i}) > p_i(s_i, t_{-i})$ for some $t_{-i} \in T_{-i}$.

(2) Strict dominance with respect to T:
$\mu_i \gg_T s_i$ if and only if $p_i(\mu_i, t_{-i}) > p_i(s_i, t_{-i})$ for all $t_{-i} \in T_{-i}$.

We obtain the following operators on $T = (T_1, \ldots, T_n)$, $T_i \subseteq S_i$, that are defined component-wise:

$$
\begin{aligned}
\mathrm{ML}(T)_i &:= \{t_i \in T_i : \neg\exists \mu_i \in \Delta(T_i)\ \mu_i >_T t_i\}, \\
\mathrm{MG}(T)_i &:= \{t_i \in T_i : \neg\exists \mu_i \in \Delta(S_i)\ \mu_i >_T t_i\}, \\
\mathrm{PL}(T)_i &:= \{t_i \in T_i : \neg\exists t_i' \in T_i\ t_i' >_T t_i\}, \text{and}
\end{aligned}
$$

$$\mathrm{PG}(T)_i := \{t_i \in T_i : \neg\exists s_i \in S_i \; s_i >_T t_i\}.$$

MLS, MGS, PLS, PGS are defined analogously with $\gg_T$ instead of $>_T$. For all T we have the following obvious inclusions:

- Every M-operator eliminates more strategies than the corresponding P-operator.
- Every operator considering (weak) dominance eliminates more strategies than the corresponding operator considering strict dominance.
- With dominance in global games more strategies are eliminated than with dominance in local games.

$$\begin{tikzcd}
& \mathrm{PL}(T) \arrow[rr] & & \mathrm{PLS}(T) \\
\mathrm{ML}(T) \arrow[ur] \arrow[rr] & & \mathrm{MLS}(T) \arrow[ur] & \\
& \mathrm{PG}(T) \arrow[uu] \arrow[rr] & & \mathrm{PGS}(T) \arrow[uu] \\
\mathrm{MG}(T) \arrow[uu] \arrow[ur] \arrow[rr] & & \mathrm{MGS}(T) \arrow[uu] \arrow[ur] &
\end{tikzcd}$$

Figure 4.1. Inclusions between the eight strategy elimination operators

Each of these operators is deflationary, i.e. $F(T) \subseteq T$ for every T and every operator F. We iterate an operator beginning with $T = S$, i.e. $F^0 := S$ and $F^{\alpha+1} := F(F^\alpha)$. Obviously, $F^0 \supseteq F^1 \supseteq \cdots \supseteq F^\alpha \supseteq F^{\alpha+1}$. Since S is finite, we will reach a fixed point F^α such that $F^\alpha = F^{\alpha+1} =: F^\infty$. We expect that for the eight fixed points MG^∞, ML^∞, etc. the same inclusions hold as for the operators $\mathrm{MG}(T)$, $\mathrm{ML}(T)$, etc. But this is not the case: For the following game $\Gamma = (\{0,1\}, (S_0, S_1), (p_0, p_1))$ we have $\mathrm{ML}^\infty \not\subseteq \mathrm{PL}^\infty$.

	X	Y	Z
A	$(2,1)$	$(0,1)$	$(1,0)$
B	$(0,1)$	$(2,1)$	$(1,0)$
C	$(1,1)$	$(1,0)$	$(0,0)$
D	$(1,0)$	$(0,1)$	$(0,0)$

We have:

- Z is dominated by X and Y.
- D is dominated by A.
- C is dominated by $\frac{1}{2}A + \frac{1}{2}B$.

Thus:

$$\mathrm{ML}(S) = \mathrm{ML}^1 = (\{A, B\}, \{X, Y\}) \subset \mathrm{PL}(S) = \mathrm{PL}^1$$
$$= (\{A, B, C\}, \{X, Y\}).$$

$\mathrm{ML}(\mathrm{ML}^1) = \mathrm{ML}^1$ since in the following game there are no dominated strategies:

	X	Y
A	$(2,1)$	$(0,1)$
B	$(0,1)$	$(2,1)$

$\mathrm{PL}(\mathrm{PL}^1) = (\{A, B, C\}, \{X\}) = \mathrm{PL}^2 \subsetneq \mathrm{PL}^1$ since Y is dominated by X (here we need the presence of C). Since B and C are now dominated by A, we have $\mathrm{PL}^3 = (\{A\}, \{X\}) = \mathrm{PL}^\infty$. Thus, $\mathrm{PL}^\infty \subsetneq \mathrm{ML}^\infty$ although ML is the stronger operator.

We are interested in the inclusions of the fixed points of the different operators. But we only know the inclusions for the operators. So the question arises under which assumptions can we prove, for two deflationary operators F and G on S, the following claim:

If $F(T) \subseteq G(T)$ for all T, then $F^\infty \subseteq G^\infty$?

The obvious proof strategy is induction over α: We have $F^0 = G^0 = S$, and if $F^\alpha \subseteq G^\alpha$, then

$$F^{\alpha+1} = F(F^\alpha) \subseteq G(F^\alpha)$$

$$F(G^\alpha) \subseteq G(G^\alpha) = G^{\alpha+1}$$

If we can show one of the inclusions $F(F^\alpha) \subseteq F(G^\alpha)$ or $G(F^\alpha) \subseteq G(G^\alpha)$, then we have proven the claim. These inclusions hold if the operators are monotone: $H : S \to S$ is monotone if $T \subseteq T'$ implies $H(T) \subseteq H(T')$. Thus, we have shown:

Lemma 4.18. Let $F, G : \mathcal{P}(S) \to \mathcal{P}(S)$ be two deflationary operators such that $F(T) \subseteq G(T)$ for all $T \subseteq S$. If either F or G is monotone, then $F^\infty \subseteq G^\infty$.

Corollary 4.19. PL and ML are not monotone on every game.

Which operators are monotone? Obviously, MGS and PGS are monotone: If $\mu_i \gg_T s_i$ and $T' \subseteq T$, then also $\mu_i \gg_{T'} s_i$. Let $T' \subseteq T$ and $s_i \in \text{PGS}(T')_i$. Thus, there is no $\mu_i \in S_i$ such that $\mu_i \gg_{T'} s_i$, and there is also no $\mu_i \in S_i$ such that $\mu_i \gg_T s_i$ and we have $s_i \in \text{PGS}(T)_i$. The reasoning for MGS is analogous if we replace S_i by $\Delta(S_i)$.

MLS and PLS are not monotone. Consider the following simple game:

	X
A	$(1,0)$
B	$(0,0)$

$$\begin{aligned} \text{MLS}(\{A,B\},\{X\}) = \text{PLS}(\{A,B\},\{X\}) = (\{A\},\{X\}) \text{ and} \\ \text{MLS}(\{B\},\{X\}) = \text{PLS}(\{B\},\{X\}) = (\{B\},\{X\}), \end{aligned}$$

but $(\{B\},\{X\}) \not\subseteq (\{A\},\{X\})$.

Thus, none of the local operators (those which only consider dominant strategies in the current subgame) is monotone. We will see that also MG and PG are not monotone in general. The monotonicity of the global operators MGS and PGS will allow us to prove the expected inclusions $\text{ML}^\infty \subseteq \text{MLS}^\infty \subseteq \text{PLS}^\infty$ and $\text{PL}^\infty \subseteq \text{PLS}^\infty$ between the local operators. To this end, we will show that the fixed points of the local and corresponding global operators coincide (although the operators are different).

Lemma 4.20. $\mathrm{MGS}^\infty = \mathrm{MLS}^\infty$ and $\mathrm{PGS}^\infty = \mathrm{PLS}^\infty$.

Proof. We will only prove $\mathrm{PGS}^\infty = \mathrm{PLS}^\infty$. Since $\mathrm{PGS}(T) \subseteq \mathrm{PLS}(T)$ for all T and PGS is monotone, we have $\mathrm{PGS}^\infty \subseteq \mathrm{PLS}^\infty$. Now we will prove by induction that $\mathrm{PLS}^\alpha \subseteq \mathrm{PGS}^\alpha$ for all α. Only the induction step $\alpha \mapsto \alpha + 1$ has to be considered: Let $s_i \in \mathrm{PLS}_i^{\alpha+1}$. Therefore, $s_i \in \mathrm{PLS}_i^\alpha$ and there is no $s_i' \in \mathrm{PLS}_i^\alpha$ such that $s_i' \gg_{\mathrm{PLS}^\alpha} s_i$. Assume $s_i \notin \mathrm{PGS}_i^{\alpha+1}$, i.e.

$$A = \{s_i' \in S_i : s_i' \gg_{\mathrm{PGS}^\alpha} s_i\} \neq \emptyset$$

(note: By induction hypothesis $\mathrm{PGS}^\alpha = \mathrm{PLS}^\alpha$). Pick an $s_i^* \in A$ which is maximal with respect to $\gg_{\mathrm{PLS}^\alpha}$. Claim: $s_i^* \in \mathrm{PLS}^\alpha$. Otherwise, there exists a $\beta \leq \alpha$ and an $s_{i'} \in S_i$ with $s_i' \gg_{\mathrm{PLS}^\beta} s_{i^*}$. Since $\mathrm{PLS}^\beta \supseteq \mathrm{PLS}^\alpha$, it follows that $s_i' \gg_{\mathrm{PLS}^\alpha} s_i^* \gg_{\mathrm{PLS}^\alpha} s_i$. Therefore, $s_i' \in A$ and s_i^* is not maximal with respect to $\gg_{\mathrm{PLS}^\alpha}$ in A. Contradiction.

But if $s_i^* \in \mathrm{PLS}^\alpha$ and $s_i^* \gg_{\mathrm{PLS}^\alpha} s_i$, then $s_i \notin \mathrm{PLS}^{\alpha+1}$ which again constitutes a contradiction.

The reasoning for MGS^∞ and MLS^∞ is analogous. Q.E.D.

Corollary 4.21. $\mathrm{MLS}^\infty \subseteq \mathrm{PLS}^\infty$.

Lemma 4.22. $\mathrm{MG}^\infty = \mathrm{ML}^\infty$ and $\mathrm{PG}^\infty = \mathrm{PL}^\infty$.

Proof. We will only prove $\mathrm{PG}^\infty = \mathrm{PL}^\infty$ by proving $\mathrm{PG}^\alpha = \mathrm{PL}^\alpha$ for all α by induction. Let $\mathrm{PG}^\alpha = \mathrm{PL}^\alpha$ and $s_i \in \mathrm{PG}_i^{\alpha+1}$. Then $s_i \in \mathrm{PG}_i^\alpha = \mathrm{PL}_i^\alpha$ and hence there is no $s_i' \in S_i$ such that $s_i' >_{\mathrm{PG}^\alpha} s_i$. Thus, there is no $s_i' \in \mathrm{PL}_i^\alpha$ such that $s_i' >_{\mathrm{PL}^\alpha} s_i$ and $s_i \in \mathrm{PL}^{\alpha+1}$. So, $\mathrm{PG}^{\alpha+1} \subseteq \mathrm{PL}^{\alpha+1}$.

Now, let $s_i \in \mathrm{PL}_i^{\alpha+1}$. Again we have $s_i \in \mathrm{PL}_i^\alpha = \mathrm{PG}_i^\alpha$. Assume $s_i \notin \mathrm{PG}_i^{\alpha+1}$. Then

$$A = \{s_i' \in S_i : s_i' >_{\mathrm{PL}^\alpha} s_i\} \neq \emptyset.$$

For every $\beta \leq \alpha$ let $A^\beta = A \cap \mathrm{PL}_i^\beta$. Pick the maximal β such that $A^\beta \neq \emptyset$ and a $s_i^* \in A^\beta$ which is maximal with respect to $>_{\mathrm{PL}^\beta}$.

Claim: $\beta = \alpha$. Otherwise, $s_i^* \notin \mathrm{PL}_i^{\beta+1}$. Then there exists an $s_i' \in \mathrm{PL}_i^\beta$ with $s_i' >_{\mathrm{PL}^\beta} s_i^*$. Since $\mathrm{PL}^\beta \supseteq \mathrm{PL}^\alpha$ and $s_i^* >_{\mathrm{PL}^\alpha} s_i$, we have

$s_i' >_{\mathrm{PL}^\alpha} s_i$, i.e. $s_i' \in A^\beta$ which contradicts the choice of s_i^*. Therefore, $s_i^* \in \mathrm{PL}_i^\alpha$. Since $s_i^* >_{\mathrm{PL}^\alpha} s_i$, we have $s_i \notin \mathrm{PL}_i^{\alpha+1}$. Contradiction, hence the assumption is wrong, and we have $s_i \in \mathrm{PG}^{\alpha+1}$. Altogether $\mathrm{PG}^\alpha = \mathrm{PL}^\alpha$. Again, the reasoning for $\mathrm{MG}^\infty = \mathrm{ML}^\infty$ is analogous. Q.E.D.

Corollary 4.23. $\mathrm{PL}^\infty \subseteq \mathrm{PLS}^\infty$ and $\mathrm{ML}^\infty \subseteq \mathrm{MLS}^\infty$.

Proof. We have $\mathrm{PL}^\infty = \mathrm{PG}^\infty \subseteq \mathrm{PGS}^\infty = \mathrm{PLS}^\infty$ where the inclusion $\mathrm{PG}^\infty \subseteq \mathrm{PGS}^\infty$ holds because $\mathrm{PG}(T) \subseteq \mathrm{PGS}(T)$ for any T and PGS is monotone. Analogously, we have $\mathrm{ML}^\infty = \mathrm{MG}^\infty \subseteq \mathrm{MGS}^\infty = \mathrm{MLS}^\infty$. Q.E.D.

This implies that MG and PG cannot be monotone. Otherwise, we would have $\mathrm{ML}^\infty = \mathrm{PL}^\infty$. But we know that this is wrong.

4.6 Beliefs and Rationalisability

Let $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$ be a game. A *belief* of Player i is a probability distribution over S_{-i}.

Remark 4.24. A belief is not necessarily a product of independent probability distributions over the individual S_j ($j \neq i$). A player may believe that the other players play correlated.

A strategy $s_i \in S_i$ is called a *best response to a belief* $\gamma \in \Delta(S_{-i})$ if $\widehat{p}_i(s_i, \gamma) \geq \widehat{p}_i(s_i', \gamma)$ for all $s_i' \in S_i$. Conversely, $s_i \in S_i$ is *never a best response* if s_i is not a best response for any $\gamma \in \Delta(S_{-i})$.

Lemma 4.25. For every game $\Gamma = (N, (S_i)_{i\in N}, (p_i)_{i\in N})$ and every $s_i \in S_i$, s_i is never a best response if and only if there exists a mixed strategy $\mu_i \in \Delta(S_i)$ such that $\mu_i \gg_S s_i$.

Proof. If $\mu_i \gg_S s_i$, then $\widehat{p}_i(\mu_i, s_{-i}) > \widehat{p}_i(s_i, s_{-i})$ for all $s_{-i} \in S_{-i}$. Thus, $\widehat{p}_i(\mu_i, \gamma) > \widehat{p}_i(s_i, \gamma)$ for all $\gamma \in \Delta(S_{-i})$. Then, for every belief $\gamma \in \Delta(S_{-i})$, there exists an $s_i' \in \mathrm{supp}(\mu_i)$ such that $\widehat{p}_i(s_i', \gamma) > \widehat{p}_i(s_i, \gamma)$. Therefore, s_i is never a best response.

Conversely, let $s_i^* \in S_i$ be never a best response in Γ. We define a two-person zero-sum game $\Gamma' = (\{0,1\}, (T_0, T_1), (p, -p))$ where $T_0 = S_i - \{s_i^*\}$, $T_1 = S_{-i}$ and $p(s_i, s_{-i}) = p_i(s_i, s_{-i}) - p_i(s_i^*, s_{-i})$. Since s_i^* is never a best response, for every mixed strategy $\mu_1 \in \Delta(T_1) = \Delta(S_{-i})$ there is a strategy $s_0 \in T_0 = S_i - \{s_i^*\}$ such that $\widehat{p}_i(s_0, \mu_1) > \widehat{p}_i(s_i^*, \mu_1)$ (in Γ), i.e. $p(s_0, \mu_1) > 0$ (in Γ'). So, in Γ'

$$\min_{\mu_1 \in \Delta(T_1)} \max_{s_0 \in T_0} p(s_0, \mu_1) > 0,$$

and therefore

$$\min_{\mu_1 \in \Delta(T_1)} \max_{\mu_0 \in \Delta(T_0)} p(\mu_0, \mu_1) > 0.$$

By Nash's Theorem, there is a Nash equilibrium (μ_0^*, μ_1^*) in Γ'. By von Neumann and Morgenstern we have

$$\begin{aligned} \min_{\mu_1 \in \Delta(T_1)} \max_{s_0 \in \Delta(T_0)} p(\mu_0, \mu_1) &= p(\mu_0^*, \mu_1^*) \\ &= \max_{s_0 \in \Delta(T_0)} \min_{\mu_1 \in \Delta(T_1)} p(\mu_0, \mu_1) > 0. \end{aligned}$$

Thus, $0 < p(\mu_0^*, \mu_1^*) \leq p(\mu_0^*, \mu_1)$ for all $\mu_1 \in \Delta(T_1) = \Delta(S_{-i})$. So, we have in Γ $\widehat{p}_i(\mu_0^*, s_{-i}) > p_i(s_i^*, s_{-i})$ for all $s_{-i} \in S_{-i}$ which means $\mu_0^* \gg_S s_i^*$. Q.E.D.

Definition 4.26. Let $\Gamma = (N, (S_i)_{i \in N}, (p_i)_{i \in N})$ be a game. A strategy $s_i \in S_i$ is *rationalisable* in Γ if for any Player j there exists a set $T_j \subseteq S_j$ such that

- $s_i \in T_i$, and
- every $s_j \in T_j$ (for all j) is a best response to a belief $\gamma_j \in \Delta(S_{-j})$ where $\text{supp}(\gamma_j) \subseteq T_{-j}$.

Theorem 4.27. For every finite game Γ we have: s_i is rationalisable if and only if $s_i \in \text{MLS}_i^\infty$. This means, the rationalisable strategies are exactly those surviving iterated elimination of strategies that are strictly dominated by mixed strategies.

Proof. Let $s_i \in S_i$ be rationalisable by $T = (T_1, \ldots, T_n)$. We show $T \subseteq \text{MLS}^\infty$. We will use the monotonicity of MGS and the fact that $\text{MLS}^\infty = \text{MGS}^\infty$. This implies $\text{MGS}^\infty = \text{gfp}(\text{MGS})$ and hence, MGS^∞

contains all other fixed points. It remains to show that $\mathrm{MGS}(T) = T$. Every $s_j \in T_j$ is a best response (among the strategies in S_j) to a belief γ with $\mathrm{supp}(\gamma) \subseteq T_{-j}$. This means that there exists no mixed strategy $\mu_j \in \Delta(S_j)$ such that $\mu_j \gg_T s_j$. Therefore, s_j is not eliminated by MGS: $\mathrm{MGS}(T) = T$.

Conversely, we have to show that every strategy $s_i \in \mathrm{MLS}_i^\infty$ is rationalisable by MLS^∞. Since $\mathrm{MLS}^\infty = \mathrm{MGS}^\infty$, we have $\mathrm{MGS}(\mathrm{MLS}^\infty) = \mathrm{MLS}^\infty$. Thus, for every $s_i \in \mathrm{MLS}_i^\infty$ there is no mixed strategy $\mu_i \in \Delta(S_i)$ such that $\mu_i \gg_{\mathrm{MLS}^\infty} s_i$. So, s_i is a best response to a belief in MLS_i^∞. Q.E.D.

Intuitively, the concept of rationalisability is based on the idea that every player keeps those strategies that are a best response to a possible combined rational action of his opponents. As the following example shows, it is essential to also consider correlated actions of the players.

Example 4.28. Consider the following cooperative game in which every player receives the same payoff:

$$\begin{array}{c|cc|cc|cc|cc}
 & L & R & L & R & L & R & L & R \\
T & 8 & 0 & 4 & 0 & 0 & 0 & 3 & 3 \\
B & 0 & 0 & 0 & 4 & 0 & 8 & 3 & 3 \\
\hline
 & \multicolumn{2}{c}{1} & \multicolumn{2}{c}{2} & \multicolumn{2}{c}{3} & \multicolumn{2}{c}{4}
\end{array}$$

Matrix 2 is not strictly dominated. Otherwise there were $p, q \in [0,1]$ with $p + q \leq 1$ and

$$8 \cdot p + 3 \cdot (1 - p - q) > 4 \text{ and}$$
$$8 \cdot q + 3 \cdot (1 - p - q) > 4.$$

This implies $2 \cdot (p+q) + 6 > 8$, i.e. $2 \cdot (p+q) > 2$, which is impossible.

So, matrix 2 must be a best response to a belief $\gamma \in \Delta(\{T, B\} \times \{L, R\})$. Indeed, the best responses to $\gamma = \frac{1}{2} \cdot ((T, L) + (B, R))$ are matrices 1, 2 or 3.

On the other hand, matrix 2 is not a best response to a belief of independent actions $\gamma \in \Delta(\{T, B\}) \times \Delta(\{L, R\})$. Otherwise, if matrix 2 were a best response to $\gamma = (p \cdot T + (1-p) \cdot B, q \cdot L + (1-q) \cdot R)$, we would have that

$$4pq + 4 \cdot (1-p) \cdot (1-q) \geq \max\{8pq, 8 \cdot (1-p) \cdot (1-q), 3\}.$$

We can simplify the left side: $4pq + 4 \cdot (1-p) \cdot (1-q) = 8pq - 4p - 4q + 4$. Obviously, this term has to be greater than each of the terms from which we chose the maximum:

$$8pq - 4p - 4q + 4 \geq 8pq \Rightarrow p + q \geq 1$$

and

$$8pq - 4p - 4q + 4 \geq 8 \cdot (1-p) \cdot (1-q) \Rightarrow p + q \leq 1.$$

So we have $p + q = 1$, or $q = 1 - p$. But this allows us to substitute q by $1 - p$, and we get

$$8pq - 4p - 4q + 4 = 8p \cdot (1-p).$$

However, this term must still be greater or equal than 3, so we get

$$\begin{aligned} & 8p \cdot (1-p) \geq 3 \\ \Leftrightarrow\ & p \cdot (1-p) \geq \frac{3}{8}, \end{aligned}$$

which is impossible since $\max(p \cdot (1-p)) = \frac{1}{4}$ (see Figure 4.2).

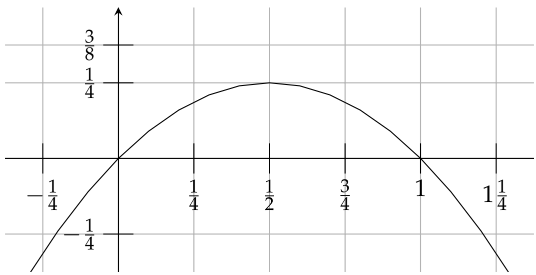

Figure 4.2. Graph of the function $p \mapsto p \cdot (1-p)$

4.7 Games in Extensive Form

A *game in extensive form* (with perfect information) is described by a game tree. For two-person games this is a special case of the games on graphs which we considered in the earlier chapters. The generalisation to n-person games is obvious: $\mathcal{G} = (V, V_1, \dots, V_n, E, p_1, \dots, p_n)$ where (V, E) is a directed tree (with root node w), $V = V_1 \uplus \cdots \uplus V_n$, and the payoff function $p_i : \text{Plays}(\mathcal{G}) \to \mathbb{R}$ for Player i, where $\text{Plays}(\mathcal{G})$ is the set of paths through (V, E) beginning in the root node, which are either infinite or end in a terminal node.

A strategy for Player i in $\mathcal{G}$ is a function $f : \{v \in V_i : vE \neq \emptyset\} \to V$ such that $f(v) \in vE$. S_i is the set of all strategies for Player i. If all players $1, \dots, n$ each fix a strategy $f_i \in S_i$, then this defines a unique play $f_1\hat{\ } \cdots \hat{\ } f_n \in \text{Plays}(\mathcal{G})$.

We say that $\mathcal{G}$ has *finite horizon* if the depth of the game tree (the length of the plays) is finite.

For every game $\mathcal{G}$ in extensive form, we can construct a game $S(\mathcal{G}) = (N, (S_i)_{i \in N}, (p_i)_{i \in N})$ with $N = \{1, \dots, n\}$ and $p_i(f_1, \dots, f_n) = p_i(f_1\hat{\ } \cdots \hat{\ } f_n)$. Hence, we can apply all solution concepts for strategic games (Nash equilibria, iterated elimination of dominated strategies, etc.) to games in extensive form. First, we will discuss Nash equilibria in extensive games.

Example 4.29. Consider the game $\mathcal{G}$ (of finite horizon) depicted in Figure 4.3 presented as (a) an extensive-form game and as (b) a strategic-form game. The game has two Nash equilibria:

- The natural solution (b, d) where both players win.
- The second solution (a, c) which seems to be irrational since both players pick an action with which they lose.

What seems irrational about the second solution is the following observation. If Player 0 picks a, it does not matter which strategy her opponent chooses since the position v is never reached. Certainly, if Player 0 switches from a to b, and Player 1 still responds with c, the payoff of Player 0 does not increase. But this threat is not credible since if v is reached after action a, then action d is better for Player 1 than c. Hence, Player 0 has an incentive to switch from a to b.

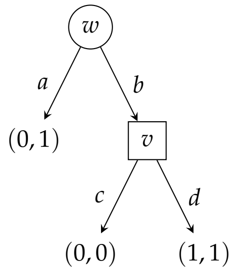

(a) extensive form

	c	d
a	$(0,1)$	$(0,1)$
b	$(0,0)$	$(1,1)$

(b) strategic form

Figure 4.3. A game of finite horizon

This example shows that the solution concept of Nash equilibria is not sufficient for games in extensive form since they do not take the sequential structure into account. Before we introduce a stronger notion of equilibrium, we will need some more notation: Let $\mathcal{G}$ be a game in extensive form and v a position of $\mathcal{G}$. $\mathcal{G}\restriction_v$ denotes the *subgame* of $\mathcal{G}$ beginning in v (defined by the subtree of $\mathcal{G}$ rooted at v). Payoffs: Let h_v be the unique path from w to v in $\mathcal{G}$. Then $p_i^{\mathcal{G}\restriction_v}(\pi) = p_i^{\mathcal{G}}(h_v \cdot \pi)$. For every strategy f of Player i in $\mathcal{G}$ let $f\restriction_v$ be the restriction of f to $\mathcal{G}\restriction_v$.

Definition 4.30. A *subgame perfect equilibrium* of $\mathcal{G}$ is a strategy profile $(f_1, \dots, f_n)$ such that, for every position v, $(f_1\restriction_v, \dots, f_n\restriction_v)$ is a Nash equilibrium of $\mathcal{G}\restriction_v$. In particular, $(f_1, \dots, f_n)$ itself is a Nash equilibrium.

In the example above, only the natural solution (b, d) is a subgame perfect equilibrium. The second Nash equilibrium (a, c) is not a subgame perfect equilibrium since $(a\restriction_v, c\restriction_v)$ is not a Nash equilibrium in $\mathcal{G}\restriction_v$.

Let $\mathcal{G}$ be a game in extensive form, $f = (f_1, \dots, f_n)$ be a strategy profile, and v a position in $\mathcal{G}$. We denote by $\widetilde{f}(v)$ the play in $\mathcal{G}\restriction_v$ that is uniquely determined by $f_1 \dots, f_n$.

Lemma 4.31. Let $\mathcal{G}$ be a game in extensive form with finite horizon. A strategy profile $f = (f_1, \ldots, f_n)$ is a subgame perfect equilibrium of $\mathcal{G}$ if and only if for every Player i, every $v \in V_i$, and every $w \in vE$: $p_i(\widetilde{f}(v)) \geq p_i(\widetilde{f}(w))$.

Proof. Let f be a subgame perfect equilibrium. If $p_i(\widetilde{f}(w)) > p_i(\widetilde{f}(v))$ for some $v \in V_i$, $w \in vE$, then it would be better for Player i in $\mathcal{G}\restriction_v$ to change her strategy in v from f_i to f_i' with

$$f_i'(u) = \begin{cases} f_i(u) & \text{if } u \neq v \\ w & \text{if } u = w\,. \end{cases}$$

This is a contradiction.

Conversely, if f is not a subgame perfect equilibrium, then there is a Player i, a position $v_0 \in V_i$ and a strategy $f_i' \neq f_i$ such that it is better for Player i in $\mathcal{G}\restriction_{v_0}$ to switch from f_i to f_i' against f_{-i}. Let $g := (f_i', f_{-i})$. We have $q := p_i(\widetilde{g}(v_0)) > p_i(\widetilde{f}(v_0))$. We consider the path $\widetilde{g}(v_0) = v_0 \ldots v_t$ and pick a maximal $m < t$ with $p_i(\widetilde{g}(v_0)) > p_i(\widetilde{f}(v_m))$. Choose $v = v_m$ and $w = v_{m+1} \in vE$. Claim: $p_i(\widetilde{f}(v)) < p_i(\widetilde{f}(w))$ (see Figure 4.4):

$$p_i(\widetilde{f}(v)) = p_i(\widetilde{f}(v_m)) < p_i(\widetilde{g}(v_m)) = q$$

$$p_i(\widetilde{f}(w)) = p_i(\widetilde{f}(v_{m+1})) \geq p_i(\widetilde{g}(v_{m+1})) = q$$

Q.E.D.

If f is not a subgame perfect equilibrium, then we find a subgame $\mathcal{G}\restriction_v$ such that there is a profitable deviation from f_i in $\mathcal{G}\restriction_v$, which only differs from f_i in the first move.

In extensive games with finite horizon we can directly define the payoff at the terminal nodes (the leaves of the game tree). We obtain a payoff function $p_i : T \to \mathbb{R}$ for $i = 1, \ldots, n$ where $T = \{v \in V : vE = \emptyset\}$.

Backwards induction: For finite games in extensive form we define a strategy profile $f = (f_1, \ldots, f_n)$ and values $u_i(v)$ for all positions v and every Player i by backwards induction:

- For terminal nodes $t \in T$ we do not need to define f, and $u_i(t) := p_i(t)$.

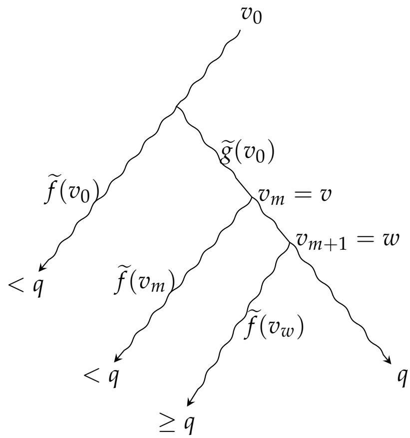

Figure 4.4. $p_i(\widetilde{f}(v)) < p_i(\widetilde{f}(w))$

- Let $v \in V \setminus T$ such that all $u_i(w)$ for all i and all $w \in vE$ are already defined. For i with $v \in V_i$ define $f_i(v) = w$ for some w with $u_i(w) = \max\{u_i(w') : w' \in vE\}$ and $u_j(v) := u_j(f_i(v))$ for all j.

We have $p_i(\widetilde{f}(v)) = u_i(v)$ for every i and every v.

Theorem 4.32. The strategy profile defined by backwards induction is a subgame perfect equilibrium.

Proof. Let $f_i' \neq f_i$. Then there is a node $v_0 \in V_i$ with minimal height in the game tree such that $f_i'(v) \neq f_i(v)$. Especially, for every $w \in vE$, $\widetilde{(f_i', f_{-i})}(w) = \widetilde{f}(w)$. For $w = f_i'(v)$ we have

$$\begin{aligned} p_i(\widetilde{(f_i', f_{-i})}(v)) &= p_i(\widetilde{(f_i', f_{-i})}(w)) \\ &= p_i(\widetilde{f}(w)) \\ &= u_i(w) \leq \max_{w' \in vE}\{u_i(w')\} \\ &= u_i(v) \\ &= p_i(\widetilde{f}(v)). \end{aligned}$$

Therefore, $f\restriction_v$ is a Nash equilibrium in $\mathcal{G}\restriction_v$. Q.E.D.

Corollary 4.33. Every finite game in extensive form has a subgame perfect equilibrium (and thus a Nash equilibrium) in pure strategies.

4.8 Subgame-perfect equilibria in infinite games

We now consider cases of infinite games in extensive form, for which we can establish the existence of subgame-perfect equilibria. Generalizing the model of infinite two-person zero-sum games on graphs, we consider multi-player, turn-based games on graphs with arbitrary (not necesssarily antagonistic) qualitative objectives.

Definition 4.34. An *infinite (turn-based, qualitative) multiplayer game* is a tuple $\mathcal{G} = (N, V, (V_i)_{i \in N}, E, \Omega, (\mathrm{Win}_i)_{i \in N})$ where N is a finite set of *players*, (V, E) is a (finite or infinite) directed graph $(V_i)_{i \in N}$ is a partition of V into the position sets for each player, $\Omega : V \to C$ is a colouring of the positions by some finite set C of colours, and $\mathrm{Win}_i \subseteq C^\omega$ is the winning condition for Player i.

For the sake of simplicity, we assume that $uE := \{v \in V : (u, v) \in E\} \neq \emptyset$ for all $u \in V$, i.e. each vertex of G has at least one outgoing edge. We call $\mathcal{G}$ a *zero-sum game* if the sets Win_i define a partition of C^ω.

A *play* of $\mathcal{G}$ is an infinite path through the graph (V, E), and a *history* is a finite initial segment of a play. We say that a play π is *won* by Player i if $\Omega(\pi) \in \mathrm{Win}_i$. A *(pure) strategy of Player i in $\mathcal{G}$* is a function $f : V^* V_i \to V$ assigning to each sequence xv ending in a position v of Player i a next position $f(xv) \in vE$. We say that a play $\pi = \pi(0)\pi(1)\ldots$ of $\mathcal{G}$ is *consistent* with a strategy f of Player i if $\pi(k+1) = f(\pi(0)\ldots\pi(k))$ for all $k < \omega$ with $\pi(k) \in V_i$. A *strategy profile* of $\mathcal{G}$ is a tuple $(f_i)_{i \in N}$ where f_i is a strategy of Player i.

It is sometimes convenient to designate an initial vertex $v_0 \in V$ of the game. We call the tuple $(\mathcal{G}, v_0)$ an *initialized infinite multiplayer game*. A *play (history) of* $(\mathcal{G}, v_0)$ is a play (history) of $\mathcal{G}$ starting with v_0. A strategy (strategy profile) of $(\mathcal{G}, v_0)$ is just a strategy (strategy profile) of $\mathcal{G}$. A strategy f of some player i in $(\mathcal{G}, v_0)$ is *winning* if every play of $(\mathcal{G}, v_0)$ consistent with σ is won by player i. A strategy profile $(f_i)_{i \in N}$ of $(\mathcal{G}, v_0)$ determines a unique play of $(\mathcal{G}, v_0)$ consistent with each f_i, called the *outcome of* $(f_i)_{i \in N}$ and denoted by $\langle (f_i)_{i \in N} \rangle$ or, in the case that the initial vertex is not understood from the context, $\langle (f_i)_{i \in N} \rangle_{v_0}$. In the following we will often use the term *game* to denote an *(initialized) infinite multiplayer game* according to Definition 4.34.

For turn-based (non-stochastic) games with qualitative winning conditions, mixed strategies play no relevant role. Nash equilibria in pure strategies take the following form:

A strategy profile $(f_i)_{i \in N}$ of a game $(\mathcal{G}, v_0)$ is a *Nash equilibrium* if for every player i and all her possible strategies f'_i in $(\mathcal{G}, v_0)$ the play $\langle f'_i, (f_j)_{j \in N \setminus \{i\}} \rangle$ is won by player i only if the play $\langle (f_j)_{j \in N} \rangle$ is also won by her.

Despite the importance and popularity of Nash equilibria, there are several problems with this solution concept, in particular for games that extend over time. This is due to the fact that Nash equilibria do not take into account the sequential nature of games and all the consequences of this. After any initial segment of a play, the players face a new situation and may change their strategies. Choices made because of a threat by the other players may no longer be rational, because the opponents have lost their power of retaliation in the remaining play.

Example 4.35. Consider a two-player Büchi game with its arena depicted in Figure 4.5; round vertices are controlled by player 1; boxed vertices are controlled by player 2; both players win if and only if vertex 3 is visited (infinitely often); the initial vertex is 1. Intuitively, the only rational outcome of this game should be the play 123^ω. However, the game has two Nash equilibria:

(1) Player 1 moves from vertex 1 to vertex 2, and player 2 moves from vertex 2 to vertex 3. Hence, both players win.
(2) Player 1 moves from vertex 1 to vertex 4, and player 2 moves from vertex 2 to vertex 5. Both players lose.

The second equilibrium certainly does not describe a rational behaviour. Indeed both players move according to a strategy that is always losing (whatever the other player does), and once player 1 has moved from vertex 1 to vertex 2, then the rational behaviour of player 2 would be to change her strategy and move to vertex 3 instead of vertex 5 as this is then the only way for her to win.

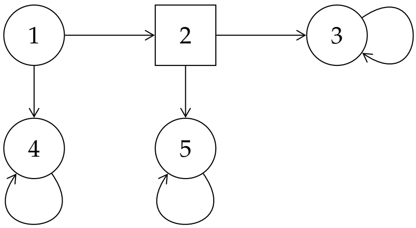

Figure 4.5. A two-player Büchi game.

This example can be modified in many ways. Indeed we can construct games with Nash equilibria in which every players moves infinitely often according to a losing strategy, and only has a chance to win if she deviates from the equilibrium strategy. The following is an instructive example with quantitative objectives.

Example 4.36. Let $\mathcal{G}_n$ be an n-player game with positions $0, \dots, n-1$. Position n is the initial position, and position 0 is the terminal position. Player i moves at position i and has two options. Either she loops at position i (and stays in control) or moves to position $i-1$ (handing control to the next player). For each player, the value of a play π is $n/|\pi|$. Hence, for all players, the shortest possible play has value 1, and all infinite plays have value 0. Obviously, the rational behaviour for each player i is to move from i to $i-1$. This strategy profile, which is of course a Nash equilibrium, gives value 1 to all players. However, the 'most stupid' strategy profile, where each player loops forever at his position, i.e. moves forever according to a losing strategy, is also a Nash equilibrium.

For a game $\mathcal{G} = (N, V, (V_i)_{i\in N}, E, \Omega, (\mathrm{Win}_i)_{i\in\Pi})$ and a history h of $\mathcal{G}$, let the game $\mathcal{G}|_h = (N, V, (V_i)_{i\in N}, E, \Omega, (\mathrm{Win}_i|_h)_{i\in N})$ be defined by $\mathrm{Win}_i|_h = \{\alpha \in C^\omega : \Omega(h) \cdot \alpha \in \mathrm{Win}_i\}$. For an initialized game $(\mathcal{G}, v_0)$ and a history hv of $(\mathcal{G}, v_0)$, we call the initialized game $(\mathcal{G}|_h, v)$ the *subgame of* $(\mathcal{G}, v_0)$ *with history* hv. For a strategy f of Player i in $\mathcal{G}$, let $f|_h : V^*V_i \to V$ be defined by $f|_h(xv) := f(hxv)$. Obviously, $f|_h$ is a strategy of Player i in $\mathcal{G}|_h$.

Recall that a strategy profile $(f_i)_{i \in N}$ is a *subgame perfect equilibrium (SPE)* if $(f_i|_h)_{i \in N}$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$ for every history hv of $(\mathcal{G}, v_0)$.

Example 4.37. Consider again the game described in Example 4.35. The Nash equilibrium where Player 1 moves from vertex 1 to vertex 4 and Player 2 moves from vertex 2 to vertex 5 is not a subgame perfect equilibrium since moving from vertex 2 to vertex 5 is not optimal for Player 2 after the play has reached vertex 2. On the other hand, the Nash equilibrium where Player 1 moves from vertex 1 to vertex 2 and Player 2 moves from vertex 2 to vertex 3 is also a subgame perfect equilibrium.

The first step in the analysis of subgame perfect equilibria for *infinite* duration games is the notion of subgame-perfect determinacy. While the notion of subgame perfect equilibrium makes sense for more general classes of extensive games, the notion of subgame-perfect determinacy applies only to games with qualitative winning conditions.

Definition 4.38. A game $(\mathcal{G}, v_0)$ is *subgame-perfect determined* if there exists a strategy profile $(f_i)_{i \in N}$ such that for each history hv of the game one of the strategies $f_i|_h$ is a winning strategy in $(\mathcal{G}|_h, v)$.

Proposition 4.39. Let $(\mathcal{G}, v_0)$ be a qualitative zero-sum game such that every subgame is determined. Then $(\mathcal{G}, v_0)$ is subgame-perfect determined.

Proof. Let $(\mathcal{G}, v_0)$ be a multiplayer game such that, for every history hv, there exists a strategy f_i^h for some player i, which is winning in $(\mathcal{G}|_h, v)$. We have to combine these strategies in an appropriate way to strategies f_i. (Let us point out that the trivial combination, namely $f_i(hv) := f_i^h(v)$ does not work in general.) We say that a decomposition $h = h_1 \cdot h_2$ is *good* for player i w.r.t. vertex v if $f_i^{h_1}|_{h_2}$ is winning in $(\mathcal{G}|_h, v)$. If the strategy f_i^h is winning in $(\mathcal{G}|_h, v)$, then the decomposition $h = h \cdot \varepsilon$ is good w.r.t. v, so a good decomposition exists.

For each history hv, if f_i^h is winning in $(\mathcal{G}|_h, v)$, we choose the good (w.r.t. vertex v) decomposition $h = h_1 h_2$ with minimal h_1, and put

$$f_i(hv) := f_i^{h_1}(h_2 v)\,.$$

Otherwise, we set $f_i(hv) := f_i^h(v)$.

It remains to show that for each history hv of $(\mathcal{G}, v_0)$ the strategy $f_i|_h$ is winning in $(\mathcal{G}|_h, v)$ whenever the strategy f_i^h is. Hence, assume that f_i^h is winning in $(\mathcal{G}|_h, v)$, and let $\pi = \pi(0)\pi(1)\ldots$ be a play starting in $\pi(0) = v$ and consistent with $f_i|_h$. We need to show that π is won by player i in $(\mathcal{G}|_h, v)$.

First, we claim that for each $k < \omega$ there exists a decomposition of the form $h\pi(0)\ldots\pi(k-1) = h_1 \cdot (h_2\pi(0)\ldots\pi(k-1))$ that is good for player i w.r.t. $\pi(k)$. This is obviously true for $k = 0$. Now, for $k > 0$, assume that there exists a decomposition $h\pi(0)\ldots\pi(k-2) = h_1 \cdot (h_2\pi(0)\ldots\pi(k-2))$ that is good for player i w.r.t. $\pi(k-1)$ and with h_1 being minimal. Then $\pi(k) = f_i(h\pi(0)\ldots\pi(k-1)) = f^{h_1}(h_2\pi(0)\ldots\pi(k-1)$, and $h\pi(0)\ldots\pi(k-1) = h_1(h_2\pi(0)\ldots\pi(k-1))$ is a decomposition that is good w.r.t. $\pi(k)$.

Now consider the sequence $h_1^0, h_1^1, \ldots$ of prefixes of the good decompositions $h\pi(0)\ldots\pi(k-1) = h_1^k h_2^k \pi(0)\ldots\pi(k-1)$ (w.r.t. $\pi(k)$) with each h_1^k being minimal. Then we have $h_1^0 \succeq h_1^1 \succeq \ldots$, since for each $k > 0$ the decomposition $h\pi(0)\ldots\pi(k-1) = h_1^{k-1} h_2^{k-1}\pi(0)\ldots\pi(k-1)$ is also good for player i w.r.t. $\pi(k)$. As $\prec$ is well-founded, there must exist $k < \omega$ such that $h_1 := h_1^k = h_1^l$ for each $k \leq l < \omega$. Hence, we have that the play $\pi(k)\pi(k+1)\ldots$ is consistent with $f_i^{h_1}|_{h_2\pi(0)\ldots\pi(k-1)}$, which is a winning strategy in $(\mathcal{G}|_{h\pi(0)\ldots\pi(k-1)}, \pi(k))$. So the play $h\pi$ is won by player i in $(\mathcal{G}, v_0)$, which implies that the play π is won by player i in $(\mathcal{G}|_h, v)$. Q.E.D.

We say that a class of winning conditions is closed under taking subgames, if for every condition $X \subseteq C^\omega$ in the class, and every $h \in C^*$, also $X|_h := \{x \in C^\omega : hx \in X\}$ belongs to the class. Since Borel winning conditions are closed under taking subgames, it follows that any two-player zero-sum game with Borel winning condition is subgame-perfect determined.

Corollary 4.40. Let $(\mathcal{G}, v_0)$ be a two-player zero-sum Borel game. Then $(\mathcal{G}, v_0)$ is subgame-perfect determined.

Multiplayer games are usually not zero-sum games. Indeed when we have many players the assumption that the winning conditions of the players form a partition of the set of plays is very restrictive and unnatural. We now drop this assumption and establish general conditions under which a multiplayer game admits a subgame perfect equilibrium. In fact we will relate the existence of subgame perfect equilibria with the determinacy of associated two-player games. In particular, it will follow that every multiplayer game with Borel winning conditions has a subgame perfect equilibrium.

In the rest of this subsection, we are only concerned with the *existence* of equilibria, not with their complexity. Thus, without loss of generality, we tacitly assume that the arena of the game under consideration is a tree or a forest with the initial vertex as one of its root. The justification for this assumption is that we can always replace the arena of an arbitrary game by its unravelling from the initial vertex, ending up in an equivalent game.

Definition 4.41. Let $\mathcal{G} = (N, V, (V_i)_{i \in N}, E, \Omega, (\mathrm{Win}_i)_{i \in N})$ be a multiplayer game (played on a forest), with winning conditions $\mathrm{Win}_i \subseteq C^\omega$. The associated class $\mathrm{Two}(\mathcal{G})$ of two-player zero-sum games is obtained as follows:

(1) For each player i, $\mathrm{Two}(\mathcal{G})$ contains the game $\mathcal{G}_i$ where player i plays $\mathcal{G}$, with his winning condition Win_i, against the coalition of all other players, with winning condition $C^\omega \setminus \mathrm{Win}_i$.

(2) Close the class under taking subgames (i.e. consider plays after initial histories).

(3) Close the class under taking subgraphs (i.e. admit deletion of positions and moves).

Note that the order in which the operations (1), (2), and (3) are applied has no effect on the class $\mathrm{Two}(\mathcal{G})$.

Theorem 4.42. Let $(\mathcal{G}, v_0)$ be a multiplayer game such that every game in $\mathrm{Two}(\mathcal{G})$ is determined. Then $(\mathcal{G}, v_0)$ has a subgame perfect equilibrium.

Proof. Let $\mathcal{G} = (N, V, (V_i)_{i\in N}, E, \Omega, (\mathrm{Win}_i)_{i\in N})$ be a multiplayer game such that every game in $\mathrm{Two}(\mathcal{G})$ is determined. For each ordinal α we define a set $E^\alpha \subseteq E$ beginning with $E^0 = E$ and

$$E^\lambda = \bigcap_{\alpha<\lambda} E^\alpha$$

for limit ordinals λ. To define $E^{\alpha+1}$ from E^α, we consider for each player $i \in N$ the two-player zero-sum game $\mathcal{G}_i^\alpha = (V, V_i, E^\alpha, \Omega, \mathrm{Win}_i)$ where player i plays, with his winning condition Win_i against the coalition of all other players (with winning condition $C^\omega \setminus \mathrm{Win}_i$). Every subgame of $\mathcal{G}_i^\alpha$ belongs to $\mathrm{Two}(\mathcal{G})$ and is therefore determined. Hence we can use Proposition 4.39 to fix a subgame perfect equilibrium $(f_i^\alpha, f_{-i}^\alpha)$ of $(\mathcal{G}, v_0)$ where f_i^α is a strategy of player i and f_{-i}^α is a strategy of the coalition. Moreover, as the arena of $\mathcal{G}^\alpha$ is a forest, these strategies can be assumed to be positional. Let X_i^α be the set of all $v \in V$ such that f_i^α is winning in $(\mathcal{G}_i^\alpha|_h, v)$ for the unique maximal history h of $\mathcal{G}$ leading to v. For vertices $v \in V_i \cap X_i^\alpha$ we delete all outgoing edges except the one taken by the strategy f_i^α, i.e. we define

$$E^{\alpha+1} = E^\alpha \setminus \bigcup_{i\in N} \{(u, v) \in E : u \in V_i \cap X_i^\alpha \text{ and } v \neq f_i^\alpha(u)\}.$$

Obviously, the sequence $(E^\alpha)_{\alpha\in\mathrm{On}}$ is non-increasing. Thus we can fix the least ordinal γ with $E^\gamma = E^{\gamma+1}$ and define $f_i = f_i^\gamma$ and $f_{-i} = f_{-i}^\gamma$. Moreover, for each player $j \neq i$ let $f_{j,i}$ be the positional strategy of player j in $\mathcal{G}$ that is induced by f_{-i}.

Intuitively, Player i's equilibrium strategy g_i is as follows: Player i plays f_i as long as no other player deviates. Whenever some player $j \neq i$ deviates from her equilibrium strategy f_j, player i switches to $f_{i,j}$. Formally, define for each vertex $v \in V$ the player $p(v)$ who has to be "punished" at vertex v where $p(v) = \perp$ if nobody has to be punished.

If the game has just started, no player should be punished. Thus we let

$$p(v) = \bot \text{ if } v \text{ is a root.}$$

At vertex v with predecessor u, the same player has to be punished as at vertex u as long as the player whose turn it was at vertex u did not deviate from her prescribed strategy. Thus for $u \in V_i$ and $v \in uE$ we let

$$p(v) = \begin{cases} \bot & \text{if } p(u) = \bot \text{ and } v = f_i(u), \\ p(u) & \text{if } p(u) \neq i,\ p(u) \neq \bot \text{ and } v = f_{i,p(u)}(u), \\ i & \text{otherwise.} \end{cases}$$

Now, for each player $i \in N$ we can define the equilibrium strategy g_i by setting

$$g_i(v) = \begin{cases} f_i(v) & \text{if } p(v) = \bot \text{ or } p(v) = i, \\ f_{i,p(v)}(v) & \text{otherwise} \end{cases}$$

for each $v \in V$.

It remains to show that $(g_i)_{i \in N}$ is a subgame perfect equilibrium of $(\mathcal{G}, v_0)$. First note that f_i is winning in $(\mathcal{G}_i^\gamma|_h, v)$ if f_i^α is winning in $(\mathcal{G}_i^\alpha|_h, v)$ for some ordinal α because if f_i^α is winning in $(\mathcal{G}_i^\alpha|_h, v)$ every play of $(\mathcal{G}_i^{\alpha+1}|_h, v)$ is consistent with f_i^α and therefore won by player i. As $E^\gamma \subseteq E^{\alpha+1}$, this also holds for every play of $(\mathcal{G}_i^\gamma|_h, v)$. Now let v be any vertex of $\mathcal{G}$ with h the unique maximal history of $\mathcal{G}$ leading to v. We claim that $(g_j)_{j \in N}$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$. Towards this, let g' be any strategy of any player $i \in N$ in $\mathcal{G}$; let $\pi = \langle (g_j)_{j \in N} \rangle_v$, and let $\pi' = \langle g', (g_j)_{j \in N \setminus \{i\}} \rangle_v$. We show that $h\pi$ is won by player i or that $h\pi'$ is not won by player i. The claim is trivial if $\pi = \pi'$. Thus assume that $\pi \neq \pi'$ and fix the least $k < \omega$ such that $\pi(k+1) \neq \pi'(k+1)$. Clearly, $\pi(k) \in V_i$ and $g'(\pi(k)) \neq g_i(\pi(k))$. Without loss of generality, let $k = 0$. We distinguish the following two cases:

- f_i is winning in $(\mathcal{G}_i^\gamma|_h, v)$. By the definition of each g_j, π is a play of $(\mathcal{G}_i^\gamma|_h, v)$. We claim that π is consistent with f_i, which implies that $h\pi$ is won by player i. Otherwise fix the least $l < \omega$

such that $\pi(l) \in V_i$ and $f_i(\pi(l)) \neq \pi(l+1)$. As f_i is winning in $(\mathcal{G}_i^\gamma|_h, v)$, f_i is also winning in $(\mathcal{G}_i^\gamma|_{h\pi(0)\dots\pi(l-1)}, \pi(l))$. But then $(\pi(l), \pi(l+1)) \in E^\gamma \setminus E^{\gamma+1}$, a contradiction to $E^\gamma = E^{\gamma+1}$.

- f_i is not winning in $(\mathcal{G}_i^\gamma|_h, v)$. Hence f_{-i} is winning in $(\mathcal{G}_i^\gamma|_h, v)$. As $g'(v) \neq g_i(v)$, player i has deviated, and it is the case that $\pi' = \langle g', (f_{j,i})_{j\in N\setminus\{i\}}\rangle_v$. We claim that π' is a play of $(\mathcal{G}_i^\gamma|_h, v)$. As f_{-i} is winning in $(\mathcal{G}_i^\gamma|_h, v)$, this implies that $h\pi'$ is not won by player i. Otherwise fix the least $l < \omega$ such that $(\pi'(l), \pi'(l+1)) \notin E^\gamma$ together with the ordinal α such that $(\pi'(l), \pi'(l+1)) \in E^\alpha \setminus E^{\alpha+1}$. Clearly, $\pi'(l) \in V_i$. Thus f_i^α is winning in $(\mathcal{G}_i^\alpha|_{h\pi'(0)\dots\pi'(l-1)}, \pi'(l))$, which implies that f_i is winning in $(\mathcal{G}_i^\gamma|_{h\pi'(0)\dots\pi'(l-1)}, \pi'(l))$. As π' is consistent with f_{-i}, this means that f_{-i} is not winning in $(\mathcal{G}_i^\gamma|_h, v)$, a contradiction.

It follows that $(g_j)_{j\in N} = (g_j|_h)_{j\in N}$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$ for every history hv of $(\mathcal{G}, v_0)$, hence $(g_j)_{j\in N}$ is a subgame perfect equilibrium of $(\mathcal{G}, v_0)$. Q.E.D.

Corollary 4.43. Every multiplayer game with Borel winning conditions has a subgame perfect equilibrium.

O course this also implies that every multiplayer game with Borel winning conditions has a Nash equilibrium. Indeed, for the existence of Nash equilibria, a slightly weaker condition suffices. Let $\mathrm{Two}(\mathcal{G})_{\mathrm{Nash}}$ be defined in the same way as $\mathrm{Two}(\mathcal{G})$ but without closure under subgraphs.

Corollary 4.44. If every game in $\mathrm{Two}(\mathcal{G})_{\mathrm{Nash}}$ is determined, then $\mathcal{G}$ has a Nash equilibrium.

Appendix A - Ordinal Numbers

The standard basic notion used in mathematics is the notion of a set, and all mathematical theorems follow from *the axioms of set theory*. The standard set of axioms, which (among others) guarantee the existence of an empty set, an infinite set, and the powerset of any set, and that no set is a member of itself (i.e. $\forall x\ \neg x \in x$) is called the *Zermelo-Fränkel Set Theory ZF*. Furthermore, it is consequence of ZF that every set a contains an $\in$-minimal element $b \in a$, i.e. $b \cap a = \emptyset$. This implies that there are no infinite $\ni$-sequences $x_1 \ni x_2 \ni x_3 \ni \dots$, because otherwise the set $\{x_1, x_2, x_3, \dots\}$ would not contain an $\in$-minimal element. It is standard in mathematics to use ZF extended by *the axiom of choice AC*, which together are called ZFC.

Since everything is a set in mathematics, there is a need to represent numbers as sets. The standard way to do this is to start with the empty set, let $0 = \emptyset$, and proceed by induction, defining $n + 1 = n \cup \{n\}$. Here are the first few numbers in this coding:

- $0 = \emptyset$,
- $1 = \{\emptyset\}$,
- $2 = \{\emptyset, \{\emptyset\}\}$,
- $3 = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\} = 2 \cup \{2\}$,
- $4 = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}\} = 3 \cup \{3\}$.

Observe that for each number n (as a set) it holds that

$$m \in n \implies m \subseteq n \text{ for every set } m. \tag{4.1}$$

Sets satisfying property (4.1) are called *transitive* sets, because (4.1) is equivalent to

$$x \in y \in n \implies x \in n \text{ for every set } x, y.$$

For example, the set $a := \{\emptyset, \{\emptyset\}, \{\{\emptyset\}\}\} \neq 3$ is a transitive set, but a does not occur on our list of natural numbers. Intuitively, the problem is that $\{\{\emptyset\}\} \notin \emptyset$ and $\emptyset \notin \{\{\emptyset\}\}$, so $\in$ is not trichotomous on a. This is why, $\in$ does not constitute a linear order on a. Now, we define a more general class of numbers, the so-called *von Neumann ordinal* numbers.

Definition 4.45. A set α is an *ordinal* if

(1) α is transitive, i.e. $x \in y \in \alpha \implies x \in \alpha$ for every x, y, and

(2) $\in$ is trichotomous on α, i.e. for every $x, y \in \alpha$ either $x = y$ or $x \in y$ or $y \in x$.

$\mathrm{On} := \{\alpha \ : \ \alpha \text{ is an ordinal}\}$ is the class of all ordinals.

We are going to prove in Theorem 4.47 that for all ordinal α, β it holds that either

$$a = b \text{ or } a \in b \text{ or } a \ni b.$$

It is even the cases, that the class of ordinal numbers forms a *well-founded* order (w.r.t. $\in$). This means, that $\in$ is a linear order on the class of ordinals and that every non-empty class X of ordinal number contains an $\in$-minimal ordinal $\alpha \in X$, i.e. $\alpha \in \beta$ for every $\beta \in X \setminus \{\alpha\}$. Note, that this also implies that the class On is a proper class, which means that On is not a set itself (otherwise On would satisfy Definition 4.45 and, hence, $\mathrm{On} \in \mathrm{On}$ in contradiction to the ZFC axioms).

It is easy to check that the natural numbers we defined above are ordinal numbers: Indeed, if n is a natural number, then we have that $n = \{0, \ldots, n-1\}$ and, consequently, for every $i \in n$ follows that $i = \{0, \ldots, i-1\} \subseteq \{0, \ldots, i-1, i, \ldots n-1\} = n$. Similarly, it is easy to see that for every $m, k \in n$ that either $m = k$ or $m \in k$ or $k \in m$ holds. It is worth mentioning that the relation $\in$ coincides with the usual order $<$ on natural numbers.

Except for natural numbers, are there any other ordinal numbers? In fact, we shall see that there are infinite many ordinals which are infinitely large. For example, consider ω which is defined by

$$\omega = \bigcup_n n = \bigcup_n \{0, \dots, n-1\} = \{0, 1, 2, 3, \dots\}.$$

ω is the set of all natural numbers, but it is easy to verify that it satisfies Definition 4.45 and, hence, ω is also an ordinal number. But it does not stop here! It is always possible to apply the $+1$ operation, which is defined as

$$\alpha + 1 := \alpha \cup \{\alpha\}.$$

Lemma 4.46. Let α be an ordinal and $\beta \in \alpha$. Then β and $\alpha + 1$ are ordinals as well.

Proof. First, we prove that β is an ordinal. To do this, we need to prove that β satisfies (1) and (2) of Definition 4.45.

(1) For this, let $d \in c \in \beta$. We need to show that $d \in \beta$. Due to $b \in c \in \beta \in \alpha$ and the transitivity of α (Definition 4.45 (1)), it follows that $b, c \in \alpha$. Thus, $\beta, c, d \in \alpha$. By Definition 4.45 (2), we can conclude that $\beta = d$ or $\beta \in d$ or $d \in \beta$ holds.
    $\beta = d$ is impossible, because $\beta = d$ would implies that $d \in c \in \beta = d$ and, thus, $c \ni d \ni c \ni \dots$ but due to the ZFC axioms there are no infinite $\ni$-sequences. Similarly, $\beta \in d$ is also wrong since otherwise $d \in c \in \beta \in d$. Therefore, $d \in \beta$ has to be true.

(2) It remains to show that Definition 4.45 (2) is true for β. But this is trivial because, due to Definition 4.45 (1), it is the case that $\beta \subseteq \alpha$ and condition (2) is assumed to be true for α.

    Now we demonstrate that $\alpha + 1$ is an ordinal number.

(1) Transitivity of $\alpha + 1$: Let $c \in b \in \alpha + 1$. Our goal is to prove that $c \in \alpha + 1$. Since $\alpha + 1 = \alpha \cup \{\alpha\}$, we can distinguish the following two cases. If $b = \alpha$, then $c \in b \in \alpha$ and, by using the transitivity of α, we can deduce that

$$c \in \alpha \subseteq \alpha \cup \{\alpha\} = \alpha + 1.$$

Otherwise, $b \neq \alpha$ and then $b \in \alpha$ (because $b \in \alpha + 1$). By transitivity of α, we obtain $c \in \alpha \subseteq \alpha + 1$.

(2) Trichotomy: Let $x, y \in \alpha + 1$. We need to prove that $x = y$ or $x \in y$ or $y \in x$ holds. If both $x, y \in \alpha$, then there is nothing to prove. Hence, $x \notin \alpha$ or $y \notin \alpha$. W.l.o.g. we assume that $x \notin \alpha$. Since $x \in \alpha + 1$, it follows that $x = \alpha$. If $y \in \alpha$, then we are done. If, on the other hand, $y \notin \alpha$, then $x = \alpha = y$. Thus, we obtain $x \in y$ or $x = y$.

Q.E.D.

But does it make sense to say that $\omega + 1$ is the *next* ordinal, is there an order on ordinals?

In fact, the ordinal numbers are linearly ordered by $\in$.

Theorem 4.47. For every ordinal α, β either $\alpha = \beta$ or $\alpha \in \beta$ or $\beta \in \alpha$. Furthermore, $\alpha \subseteq \beta$ holds, if and only if $\alpha \in \beta$ or $\alpha = \beta$.

Before can prove this theorem, we need some lemmas first.

Lemma 4.48. If X is a non-empty class of ordinals, then

$$\bigcap X := \{x \ : \ x \in a \text{ for every } a \in X\}$$

is an ordinal.

Proof. Since X is non-empty, there is an ordinal $\alpha \in X$ and, then, $\bigcap X \subseteq \alpha$. Because α is a set, it is possible to prove (by using the ZFC axioms) that $\bigcap X$ is a set. Now it suffices to prove that $\bigcap X$ satisfies the two conditions from Definition 4.45:

(1) Transitivity: Let $a \in b \in \bigcap X$. Then $a \in b \in \gamma$ for all $\gamma \in X$. Since X is a class of ordinals, it follows that $a \in \gamma$ for all $\gamma \in X$ and, finally, $a \in \bigcap X$.

(2) Trichotomy: Let $a, b \in \bigcap X$. Then $a, b \in \alpha$ and, because α is an ordinal, $a \in b$ or $a = b$ or $b \in a$.

Q.E.D.

The transitivity of ordinals allows us to prove that elements of ordinals are subsets. Of course, the converse is not true in general, because not every subset of an ordinal is an element. However, proper subsets that are ordinals turn out to be elements. As usual we write $\alpha \subset \beta$ as a shorthand for $\alpha \subseteq \beta$ and $\alpha \neq \beta$.

Lemma 4.49. Let α, β be ordinals and $\alpha \subset \beta$. Then $\alpha \in \beta$.

Proof. Towards a contradiction, we assume there are some ordinals $\alpha \subset \beta$ with $\alpha \notin \beta$.

In order to obtain a contradiction, we prove that there is an infinite $\ni$-sequence $\beta_0 \ni \beta_1 \ni \beta_2 \dots$ of ordinals starting at β such that $\alpha \subset \beta_i$ but $\alpha \notin \beta_i$ for all $i \in \{0, 1, 2, \dots\}$.

We start with $\beta_0 := \beta$. Now, consider the set

$$\beta_0 \setminus \alpha := \{y \in \beta_0 \ : \ y \notin \alpha\} .$$

We define $\gamma := \bigcap(\beta_0 \setminus \alpha)$. Due to $\alpha \subset \beta = \beta_0$, there is a $\mu \in \beta_0 \setminus \alpha$. As a result, $\beta_0 \setminus \alpha \neq \emptyset$ and γ is an ordinal (by Lemma 4.48).

Claim 4.50. $\alpha \subseteq \gamma$.

Proof. Let $\delta \in \alpha$. We are going to prove that $\delta \in \gamma$.

Since $\alpha \subset \beta_0$ we have $\delta \in \beta_0$. Let $\mu' \in \beta_0 \setminus \alpha$ be picked arbitrarily. As a result $\mu', \delta \in \beta_0$ and, by Definition 4.45 (2), it follows that

$$\mu' = \delta \text{ or } \mu' \in \delta \text{ or } \delta \in \mu'.$$

We observe that $\mu' \neq \delta$, because $\mu' \notin \alpha$ but $\delta \in \alpha$. Furthermore, $\mu' \notin \delta$, because otherwise $\mu' \in \delta \in \alpha$ and (since α is an ordinal) $\mu' \in \alpha$ but $\mu' \notin \alpha$.

Therefore, it must be the case that $\delta \in \mu'$. $\mu' \in \beta_0 \setminus \alpha$ was chosen arbitrarily, so $\delta \in \bigcap(\beta_0 \setminus \alpha) = \gamma$. Q.E.D.

Now we have

$$\alpha \subseteq \gamma = \bigcap \beta_0 \setminus \alpha.$$

Recall that $\mu \in \beta_0 \setminus \alpha$ and, therefore, $\gamma = \bigcap(\beta_0 \setminus \alpha) \subseteq \mu$. Together with $\alpha \subseteq \gamma$ this leads to $\alpha \subseteq \mu$. Since $\mu \in \beta_0$ and $\alpha \notin \beta_0$, it follows that $\alpha \subset \mu$. Furthermore, $\alpha \notin \mu$, because otherwise $\alpha \in \mu \in \beta_0$ and then $\alpha \in \beta_0$ (because β_0 is an ordinal) in contradiction to $\alpha \notin \beta_0$.

All in all, we managed to prove that $\mu \in \beta_0$ is an ordinal (due to Lemma 4.46) with $\alpha \subset \mu$ but $\alpha \notin \mu$. Hence, we can set $\beta_1 := \mu$.

By repetition, we can construct the desired sequence $\beta_0 \ni \beta_1 \ni \beta_2 \ni \ldots$, but this contradicts the ZFC axioms!

Q.E.D.

Now we have all the tools we need to finally prove Theorem 4.47.

Proof (of Theorem 4.47). First we prove that $\alpha \subseteq \beta \Longleftrightarrow \alpha = \beta \vee \alpha \in \beta$.

The direction "$\Longleftarrow$" follows intermediately from Definition 4.45 (1), while "$\Longrightarrow$" is Lemma 4.49.

Now we demonstrate that $\in$ is a linear order on the class of ordinal numbers. Towards a contradiction, assume that there are ordinals α, β that are incomparable w.r.t. $\in$, i.e., we have

$$\alpha \neq \beta \text{ and } \alpha \notin \beta \text{ and } \beta \notin \alpha. \tag{4.2}$$

Consider $\alpha \cap \beta$. By Lemma 4.48, $\alpha \cap \beta$ is an ordinal. Furthermore, $\alpha \cap \beta \subseteq \alpha$ and $\alpha \cap \beta \subseteq \beta$. If $\alpha = \alpha \cap \beta$, then $\alpha \subseteq \beta$ and by Lemma 4.49 either $\alpha = \beta$ or $\alpha \in \beta$ in contradiction to (4.2). Thus, $\alpha \neq \alpha \cap \beta$ and, similarly, $\beta \neq \alpha \cap \beta$.

But then, $\alpha \cap \beta \subset \alpha$ and $\alpha \cap \beta \subset \beta$, which implies that $\alpha \cap \beta \in \alpha$ and $\alpha \cap \beta \in \beta$, which leads to $\alpha \cap \beta \in \alpha \cap \beta$, but due to the ZFC axioms this is not possible! Contradiction!

So, $\in$ is in fact a linear order on the class of ordinal numbers.

Q.E.D.

Recall that On is the class of all ordinals. Theorem 4.47 tells us that $\in$ is a linear order on On. More general, $\in$ is a *well-founded order* on On. An order $(A, <)$ is a *well-founded order*, if

(1) $(A, <)$ is a linear order and

(2) for every non-empty set $X \subseteq A$ there is a $<$-minimal element $x \in X$, i.e., $x < y$ for every $y \in X$.

For example, $(\mathbb{N}, <)$ is a well-founded order but $(\mathbb{Z}, <)$ or $(\mathbb{Q}_{\geq 0}, <)$ are not well-founded orders.

It is not difficult to see that ordinal numbers are well-founded orders (w.r.t. $\in$). Indeed, if $X \subseteq \mathrm{On}$ is a non-empty class of ordinals, then $\gamma := \bigcap X$ is an ordinal (by Lemma 4.48) and $\gamma \subseteq x$ for all $x \in X$. It remains to prove that $\gamma \in X$: Otherwise $\gamma \notin X$ and then $\gamma \subset x$ for all $x \in X$. This leads to $\gamma \in x$ for all $x \in X$ (by Theorem 4.47). Thus, $\{\gamma\} \subseteq x$ and, as a consequence, $\gamma + 1 \subseteq x$ for all $x \in X$. But then $\gamma + 1 \subseteq \bigcap X = \gamma \implies \gamma \in \gamma$ which violates the ZFC axioms! Hence, $\bigcap X = \gamma \in X$.

Now we turn our attention towards the construction of bigger ordinals. For this, we need the following lemma which states that ordinal numbers are closed under unions.

Lemma 4.51. Let x be a set of ordinals, i.e., every $\alpha \in x$ is an ordinal. Then

$$\bigcup x := \{\beta \ : \ \beta \in \alpha \text{ for some } \alpha \in x\}$$

is an ordinal number.

Proof. Using the ZFC axioms, it is possible to prove that $\bigcup x$ is a set. Hence, it remains to show that (1) and (2) of Definition 4.45 are satisfied.

(1) Transitivity of $\bigcup x$: If $a \in b \in \bigcup x$, then there is a $c \in x$ such that $a \in b \in c$ and, by transitivity, $a \in c$ which implies that $a \in \bigcup x$.
(2) Trichotomy: If $a, b \in \bigcup x$. Then there are some $c, d \in x$ such that $a \in c$ and $b \in d$. Applying Lemma 4.46 yields that a, b are ordinals and, by Theorem 4.47, either $a = b$ or $a \in b$ or $b \in a$.

Q.E.D.

$\omega := \bigcup_n n$, the union of all natural numbers, is again an ordinal number. To prove this, we observe that $\omega = \bigcup \omega$ and use Lemma 4.51 (that ω is a set is a consequence of the axiom of infinity).

What is the next ordinal number after ω? We can again apply the $+1$ operation in the same way as for natural numbers, so

$$\omega + 1 = \omega \cup \{\omega\} = \{0, 1, 2, \ldots, \{0, 1, 2, \ldots\}\}.$$

Of course it is now possible to construction ordinals like $\omega + 2 := (\omega + 1) + 1, \omega + 3, \ldots$ and then we can build the union

$$\omega + \omega = \bigcup_{i \in \omega} \omega + i = \{0, 1, 2, \ldots, \omega, \omega + 1, \omega + 2, \ldots\},$$

which is an ordinal because of Lemma 4.51. The fact that $\omega \cup \{\omega + i \ : \ i \in \omega\}$ is a set can be proven by using the axiom of replacement.

To get an intuition on how ordinals look like, consider the following examples of infinite ordinals: $\omega + 1,\ \omega + \omega = 2\omega,\ 3\omega, \ldots, \omega \cdot \omega = \omega^2,\ \omega^3, \ldots, \omega^\omega$.

For some ordinals α it is the case that $\alpha = \beta + 1$ for some β. However, it is not possible to find an ordinal γ such that $\gamma + 1 = \omega$ holds (Why?).

Definition 4.52. Let $\alpha \neq 0$ be an ordinal. If $\beta + 1 \in \alpha$ for every $\beta \in \alpha$, then we call α an *limes* ordinal.

It is easy to see that λ is an limes ordinal, if and only if $\lambda \neq 0$ and $\bigcup \lambda = \lambda$.

Ordinals that are not limes ordinals are called successor ordinals because of the following theorem.

Theorem 4.53. Let $\alpha \neq 0$ be an ordinal that is not an limes ordinal. Then there is an ordinal β such that $\beta + 1 = \alpha$.

Proof. By Definition 4.52 there is a $\beta \in \alpha$ such that $\beta + 1 \notin \alpha$. By Theorem 4.47, either $\beta + 1 = \alpha$ or $\beta + 1 \ni \alpha$.

So, we only need to show that $\beta + 1 \not\ni \alpha$ holds. Otherwise $\alpha \in \beta + 1 = \beta \cup \{\beta\}$. Clearly, $\alpha \notin \{\beta\}$ because $\alpha = \beta \in \alpha$ would violate the ZFC axioms. But then $\alpha \in \beta \in \alpha$ which contradicts the ZFC axioms as well. Hence $\beta + 1 \ni \alpha$ is impossible which leads to $\beta + 1 = \alpha$. Q.E.D.

Ordinals are intimately connected to well-orders. In fact any well-ordering $(A, <)$ is isomorphic to some $(\alpha, \in)$ where α is an ordinal. For example, $(\mathbb{N}, <)$ is isomorphic to $(\omega, \in)$ and $\omega + \omega$ represents $(\{0,1\} \times \mathbb{N}, <_{\text{lex}})$ where $<_{\text{lex}}$ is the lexicographical order.

The well-ordering of ordinals allows to define and prove the principle of *transfinite induction*. This principle states that On, the class of *all ordinals*, is generated from $\emptyset$ by taking the successor $(+1)$ and the union on limit steps, as shown on the examples before.

The principle of *transfinite induction* allows us to define sets X_α where α is an ordinal number. Since On is a well-order, we only need to describe how X_α is constructed under the assumption that X_β is already defined for every $\beta \in \alpha$.

For example, it is possible to define (via transfinite induction) the winning region of player 0 in a reachability game (V, V_0, V_1, E). To do this, we define sets W^0_α for every ordinal number α:

$$\begin{aligned} W^0_0 &:= \emptyset, \\ W^0_{\alpha+1} &:= \left\{ x \in V_0 \ : \ xE \cap W^0_\alpha \neq \emptyset \right\} \cup \left\{ x \in V_1 \ : \ xE \subseteq W^0_\alpha \right\}, \\ W^0_\lambda &:= \bigcup_{\beta \in \lambda} W^0_\beta \text{ for limes ordinals } \lambda. \end{aligned}$$

Now it is easy to verify that $\bigcup_{\alpha \in \mathrm{On}} W^0_\alpha$ is the winning region of Player 0.

4.9 Cardinal Numbers

Besides ordinals, we sometimes need cardinal numbers which are special ordinal number that can be used to measure the size of sets. We say that two sets x, y have the same cardinality, if there is a bijection between x and y.

Definition 4.54. An ordinal κ is a *cardinal number*, if for every $\alpha \in \kappa$ there is no bijection between κ and α. Furthermore, we say that a cardinal number κ is *the* cardinality of a set x, if there is a bijection between x and κ. In this case we let $|x| := \kappa$.

$\mathrm{Cn} := \{\kappa \in \mathrm{On} \ : \ \kappa \text{ is a cardinal number}\}$ is the class of all cardinal numbers.

But is it guaranteed that we really find a cardinal number for every possible set out there? The next theorem answers this question.

Theorem 4.55. For every set x there is a cardinal number $|x|$.

Proof. Consider the class Y of ordinals, which is given by

$$\begin{aligned} Y &:= \{\alpha \in \mathrm{On} \ : \ \text{there is a bijection } f : x \to \alpha\} \\ &= \{\alpha \in \mathrm{On} \ : \ \text{there is a bijection } f : \alpha \to x\}\,. \end{aligned}$$

If Y is non-empty, then $|x| := \bigcap Y \in Y$ is the desired cardinal number.

Now we prove that $X \neq \emptyset$ is indeed the case. By the axiom of choice, there is a choice function g for x, i.e., for every $y \subseteq x$ with $y \neq \emptyset$ we have $g(y) \in y$.

Using transfinite induction, we define for every ordinal α an object x_α by

$$x_\alpha := \begin{cases} g(y_\alpha) & \text{if } y_\alpha := x \setminus \left\{x_\beta \ : \ \beta \in \alpha\right\} \neq \emptyset \\ x & \text{if } y_\alpha = \emptyset \end{cases}$$

It is easy to see that for every $x_\alpha \neq x$ we have that x_α is an element of x but $x_\alpha \neq x_\beta$ for every $\beta \in \alpha$.

If $x_\alpha = x$ holds for some ordinal α, then there is a minimal ordinal $\alpha' \subseteq \alpha$ such that $x_{\alpha'} = x$ and, by definition of $x_{\alpha'}$, this means that $x = \left\{x_\beta \ : \ \beta \in \alpha'\right\}$. Furthermore, the function $f : \alpha' \to x$, $\beta \mapsto x_\beta$ is a bijection between x and α'. This implies that $\alpha' \in Y$.

So, it only remains to prove that $x_\alpha = x$ for some ordinal α. Towards a contradiction, we assume that $x_\alpha \neq x$ for every ordinal α. Then every $x_\alpha \in x$ and, therefore, the mapping $f : \mathrm{On} \to x' := \{x_\alpha \ : \ \alpha \in \mathrm{On}\}\,,\ \alpha \mapsto x_\alpha$ is a bijection between On and x'. Since x is a set, $x' \subseteq x$ is a set as well. Therefore, by the axiom of replacement,

$$f^{-1}[x'] := \left\{f^{-1}(y) \ : \ y \in x'\right\} = \mathrm{On}$$

is a set. As a result, On satisfies Definition 4.45 and, consequently, On $\in$ On which violates the ZFC axioms! Contradiction!

Q.E.D.

It is worth mentioning that the enumeration $(x_\alpha)_{\alpha \in |x|}$ induces a well-ordering $<$ on x by

$$x_\alpha < x_\beta \Longleftrightarrow \alpha \in \beta.$$

Corollary 4.56 (Well-ordering theorem)**.** Every set x can be well-ordered, i.e., there is a well-order $<$ on x.

Every finite ordinal number is a cardinal number but there are also infinite cardinal numbers. For example, $\beth_0 := \omega$ is the smallest infinite cardinal number and, by using the power set, we can construct strictly larger cardinal numbers:

$$\begin{aligned} \beth_{\alpha+1} &:= 2^{\beth_\alpha} := |\mathcal{P}(\beth_\alpha)|\,, \\ \beth_\lambda &:= \bigcup_{\beta \in \lambda} \beth_\beta \text{ for limes ordinals } \lambda. \end{aligned}$$

Please observe that $\beth_1 = |\mathcal{P}(\omega)| = |\mathbb{R}|$.

Whether there exists cardinal numbers between $\beth_0$ and $\beth_1$ is called the continuum hypothesis (CH) which has turned out to be independent of ZFC, i.e., neither (CH) nor $\neg$(CH) are consequences of ZFC.