Logic and Games
SS 2009

Prof. Dr. Erich Grädel
Łukasz Kaiser, Tobias Ganzow

Mathematische Grundlagen der Informatik
RWTH Aachen

Contents

2 Parity Games and Fixed-Point Logics

2.1 Parity Games

In the previous section we presented 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. 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 like temporal logics, dynamic logics, fixed-point logics and others are needed. 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 where to each position a natural number is assigned. This number is called the *priority* of the position, and the winner of an infinite play is determined according to whether the least priority seen infinitely often during the play is even or odd.

Definition 2.1. We describe a *parity game* by a labelled graph $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ where (V, V_0, V_1, E) is a game graph and $\Omega : V \to \mathbb{N}$, with $|\Omega(V)|$ finite, assigns a *priority* to each position. The set V of positions may be finite or infinite, but the number of different priorities, called the *index* of $\mathcal{G}$, must be finite. Recall that a finite play of a game is lost by the player who gets stuck, i.e. cannot move. For infinite plays $v_0 v_1 v_2 \ldots$, we have a special 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.

Definition 2.2. 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 $v_{i+1} = f(v_i)$. 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 winning for Player σ.

In general, a strategy depends on the whole history of the game. However, in this chapter, we are interested in simple strategies that depend only on the current position.

Definition 2.3. 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 game, i.e. $f(hv) = f(h'v)$ for all $h, h' \in V^*, v \in V$. We often view positional strategies simply as functions $f : V \to V$.

We will see that such positional strategies suffice to solve parity games by proving the following theorem.

Theorem 2.4 (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.*

Before proving the theorem, we give two general examples of positional strategies, namely attractor and trap strategies, and show how positional winning strategies on parts of the game graph may be combined to positional winning strategies on larger regions.

Remark 2.5. Let f and f' be positional strategies for Player σ that are winning 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 $W \cup W'$.

Definition 2.6. Let $\mathcal{G} = (V, V_0, V_1, E)$ be a game and $X \subseteq V$. We define the *attractor of X for Player σ* as

$$\begin{aligned} \text{Attr}_\sigma(X) = \{v \in V : \text{Player } \sigma \text{ has a (w.l.o.g. positional) strategy} \\ \text{to reach some position } x \in X \cup T_\sigma \\ \text{in finitely many steps}\} \end{aligned}$$

where $T_\sigma = \{v \in V_{1-\sigma} : vE = \emptyset\}$ denotes the set of terminal positions in which Player σ has won.

A set $X \subseteq V$ is called a *trap* for Player σ if Player $1 - \sigma$ has a (w.l.o.g. positional) strategy that avoids leaving X from every $x \in X$.

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

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 $\Omega(V) = \{0, \dots, m-1\}$ or $\Omega(V) = \{1, \dots, m\}$. We prove the statement by induction over $|\Omega(V)|$.

In the case of $|\Omega(V)| = 1$, i.e., $\Omega(V) = \{0\}$ or $\Omega(V) = \{1\}$, the theorem clearly holds as 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 $\Omega(V) = \{\sigma\}$,

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

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. Concerning W_σ, it can be seen that there is a positional strategy that avoids leaving this $(1-\sigma)$-trap.

Let $|\Omega(v)| = m > 1$. We only consider the case $0 \in \Omega(V)$, i.e., $\Omega(V) = \{0, \dots, m-1\}$ since otherwise we can use the same argumentation with switched roles of the players. We define

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

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 $V \setminus X_1$, so in particular we have $W_1 = X_1$ and $W_0 = V \setminus X_1$.

First of all, observe that $V \setminus X_1$ is a trap for Player 1. Indeed, if Player 1 could move to X_1 from a $v \in V_1 \setminus X_1$, then v would also be in X_1. Thus, there exists a positional *trap strategy* f for Player 0 that guarantees to stay in $V \setminus X_1$.

Let $Y = \Omega^{-1}(0) \setminus X_1$, $Z = \mathrm{Attr}_0(Y)$ and let a be an *attractor strategy* for Player 0 which guarantees that Y (or a terminal winning position $y \in T_0$) can be reached from every $z \in Z \setminus Y$. Moreover, let $V' = V \setminus (X_1 \cup Z)$.

The restricted game $\mathcal{G}' = \mathcal{G}|_{V'}$ has less 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 $\mathcal{G}'$: $V' = W'_0 \cup W'_1$ and there exist positional winning strategies f' for Player 0 on W'_0 and g' for Player 1 on W'_1 in $\mathcal{G}'$.

We have that $W'_1 = \emptyset$, as 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 $W'_1 = \emptyset$.

Knowing that $W'_1 = \emptyset$, let $f^* = f' + a + f$, i.e.

$$f^*(x) = \begin{cases} f'(x) & \text{if } x \in W'_0 \\ a(x) & \text{if } x \in Z \setminus Y \\ f(x) & \text{if } x \in Y \end{cases}$$

We claim that f^* is a positional winning strategy for Player 0 from $V \setminus X_1$. If π is a play consistent with f^*, then π stays in $V \setminus X_1$.

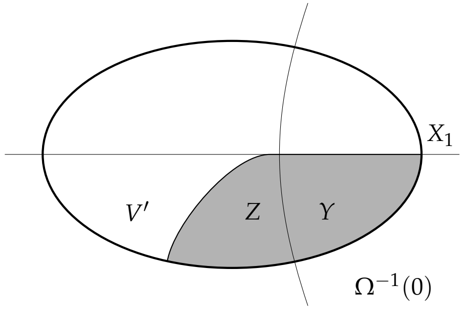

Figure 2.1. Construction of a winning strategy

Case (a): π hits Z only finitely often. Then π eventually stays in W'_0 and is consistent with f' from this point, so 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 2.7. It can be decided in NP ∩ 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.

2.1.1 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. Such algorithms run in polynomial time when the number of priorities in the input parity game is bounded.

One way to intuitively understand an algorithm solving a parity game is to imagine a judge who watches the players playing the game. At some point, the judge is supposed to say "Player 0 wins", and indeed, whenever the judge does so, there should be no question that Player 0 wins. Note that we have no condition in case that Player 1 wins. We will first give a formal definition of a certain kind of judge with bounded memory, and later use this notion to construct algorithms for parity games.

Definition 2.8. A *judge* $\mathcal{M} = (M, m_0, \delta, F)$ 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 judge is thus formally the same as an automaton reading words over the alphabet V. But to be called a judge, two special properties must be fulfilled. Let $v_0 v_1 \dots$ be a play of $\mathcal{G}$ and $m_0 m_1 \dots$ the corresponding sequence of states of $\mathcal{M}$, i.e., m_0 is the initial state of $\mathcal{M}$ and $m_{i+1} = \delta(v_i, m_i)$. Then the following holds:

(1) if $v_0 \dots$ 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 $i < j \le k$ 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. The judge is

indeed truthful, because both players can use a positional strategy in a parity game, so if a cycle with even priority appears, then 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 judge $\mathcal{M}$.

Definition 2.9. 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. The reachability game $\mathcal{G} \times \mathcal{M}$ is defined as follows:

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

where $((v, m), (v', m')) \in E'$ iff $(v, v') \in E$ and $m' = \delta(v, m)$, and the last component $V \times F$ denotes positions which are immediately winning for Player 0 (the goal of Player 0 is to reach such a position).

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 σ in $\mathcal{G}$ we define the strategy $\overline{\sigma}$ in $\mathcal{G} \times \mathcal{M}$ as

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

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

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

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

Having defined $\mathcal{G} \times \mathcal{M}$, we are ready to formally prove that the above definition of a judge indeed makes sense for parity games.

Theorem 2.10. Let $\mathcal{G}$ be a parity game and $\mathcal{M}$ a judge for $\mathcal{G}$. Then Player 0 wins $\mathcal{G}$ from v_0 if and only if he wins $\mathcal{G} \times \mathcal{M}$ from (v_0, m_0).

Proof. ($\Rightarrow$) By contradiction, let σ be the winning strategy for Player 0 in $\mathcal{G}$ from v_0, and assume that there exists a winning strategy ρ for

Player 1 in $\mathcal{G} \times \mathcal{M}$ from (v_0, m_0). (Note that we just used determinacy of reachability games.) Consider the unique plays

$$\pi_{\mathcal{G}} = v_0 v_1 \dots \quad \text{and} \quad \pi_{\mathcal{G}\times\mathcal{M}} = (v_0, m_0)(v_1, m_1) \dots$$

in $\mathcal{G}$ and $\mathcal{G} \times \mathcal{M}$, respectively, which are consistent with both σ and $\underline{\rho}$ (the play $\pi_{\mathcal{G}}$) and with $\overline{\sigma}$ and ρ ($\pi_{\mathcal{G}\times\mathcal{M}}$). Observe that the positions of $\mathcal{G}$ appearing in both plays are indeed the same due to the way $\overline{\sigma}$ and $\underline{\rho}$ are defined. Since Player 0 wins $\pi_{\mathcal{G}}$, by Property (1) in the definition of a judge 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 σ be a winning strategy for Player 0 in $\mathcal{G} \times \mathcal{M}$, and let ρ be a *positional* winning strategy for Player 1 in $\mathcal{G}$. Again, we consider the unique plays

$$\pi_{\mathcal{G}} = v_0 v_1 \dots \quad \pi_{\mathcal{G}\times\mathcal{M}} = (v_0, m_0)(v_1, m_1) \dots$$

such that $\pi_{\mathcal{G}}$ is consistent with $\underline{\sigma}$ and ρ, and $\pi_{\mathcal{G}\times\mathcal{M}}$ is consistent with σ and $\overline{\rho}$. 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 judge, 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 σ' for Player 0 in $\mathcal{G}$:

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

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

We will now show that the unique play π' in $\mathcal{G}$ that is consistent with both σ' and ρ is won by Player 0. Since in the first j steps σ' is the same as $\underline{\sigma}$, we have that $\pi[n] = v_n$ for all $n \leq j$. Now observe that $\pi[j+1] = v_{i+1}$. Since ρ 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 $\sigma'(v_0 \dots v_j) = \sigma(v_0 \dots v_i)$. Inductively repeating this reasoning, we get that the play π repeats the cycle $v_i v_{i+1} \dots v_j$ infinitely often, i.e.

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

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

The above theorem allows us, if only a judge 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 judge for an input parity game.

The most naïve way to build a judge for a *finite* parity game $\mathcal{G}$ is to just remember, for each position v visited during the play, what is the minimal priority seen in the play since the last occurrence of v. If it happens that a position v is repeated and the minimal priority since v last occurred is even, then the judge decides that Player 0 won the play.

It is easy to check that an automaton defined in this way indeed is a judge for any finite parity game $\mathcal{G}$, but such judge 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 judge is in the order of $O(d^n)$. We will present a judge that is much better for small d.

Definition 2.11. A *progress-measuring judge* $\mathcal{M}_P = (M_P, m_0, \delta_P, F_P)$ 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

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

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

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

The set F_P contains all tuples $(c_0, 0, c_2, \dots, c_m)$ 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 judge.

Lemma 2.12. For each finite parity game $\mathcal{G}$ the automaton $\mathcal{M}_P$ constructed above is a judge for $\mathcal{G}$.

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

(1) if $v_0 \dots$ 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 $\min\{\Omega(v_{i+1}), \Omega(v_{i+2}), \dots, \Omega(v_j)\}$ is even.

To see (1), assume that $v_0 v_1 \dots$ is a play winning for Player 0. Let k be such an index that $\Omega(v_k)$ is even, appears infinitely often in $\Omega(v_k)\Omega(v_{k+1})\dots$, and no priority higher 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 $v_0 v_1 \dots v_k$ 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 $v_{i_0}, v_{i_1}, \dots, v_{i_{n_p}}$ 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 judge 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 2.13. Parity games can be solved in time $O((\frac{n}{d/2})^{d/2})$.

Notice that the algorithm using a judge 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 judge $\mathcal{M}_P$ 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.

2.2 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 2.14. 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.

Definition 2.15. 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 2.16 (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.

2.3 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 $\Omega(T\bar{a}) \le \Omega(T'\bar{b})$ 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 2.17. 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 lowest priority $m = \min_{v \in V} \Omega(v)$ is even and that for all positions $v \in V$ with minimal priority $\Omega(v) = m$ we have a unique successor $vE = \{s(v)\}$. This assumption can be easily satisfied by changing the game slightly.

We define the set

$$T = \{v \in V : \Omega(v) = m\}$$

of positions with minimal priority. For any such set T we get 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 rendered terminal positions.

Additionally, we define a sequence of games $\mathcal{G}^\alpha = (V, V_0^\alpha, V_1^\alpha, E^-, \Omega)$ that only differ in the assignment of the terminal positions in T to the players. For this purpose, we use a sequence of disjoint pairs of sets T_0^α and T_1^α such that each pair partitions the set T, and let $V_\sigma^\alpha = (V_\sigma \setminus T) \cup T_{1-\sigma}^\alpha$, i.e., Player σ wins at final positions $v \in T_\sigma^\alpha$. The sequence of partitions is inductively defined depending on the winning regions 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

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

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

Lemma 2.18 (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 $\Omega(t) = m \le \Omega(v)$ for all $v \in V, t \in T$, the lowest priority seen infinitely often is m, which we have assumed to be even, 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 $\pi = \pi(0)\pi(1)\dots$ be a play consistent with g and $\pi(0) \in W_1^\infty$.

Claim 2.19. Let $\pi(i) \in W_1^\infty$. Then

(1) $\pi(i+1) \in W_1^\infty$,
(2) $\rho(\pi(i+1)) \le \rho(\pi(i))$
(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 $\rho(\pi(i+1)) \le \beta = \rho(\pi(i))$.
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 $\rho(\pi(i+1)) \le \gamma < \beta = \rho(\pi(i))$. 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 $i \ge k$, which means that $\pi(i) \notin T$ for all $i \ge k$. As $\pi(k)\pi(k+1)\dots$ 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 2.20. 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 $\mathfrak{A}$.

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

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 $X \subset A^k$

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

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

2.3.1 Defining Winning Regions in Parity Games

To conclude, 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.

To represent a parity game $\mathcal{G} = (V, V_0, V_1, E, \Omega)$ with priorities $\Omega(V) = \{0, 1, \ldots, d-1\}$, we use the Kripke structure $\mathcal{K}_\mathcal{G} = (V, V_0, V_1, E, P_0, \ldots, P_{d-1})$. The universe and edge relation of this Kripke structure are the same as in the parity game, and so are the predicates V_0 and V_1 assigning positions to players. The only difference is in the predicates P_j, which are used to explicitly represent positions with priority j, i.e. we define $P_j = \{v \in V : \Omega(v) = j\}$.

Given the above representation, the μ-calculus formula

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

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 2.21. $\mathcal{K}_\mathcal{G}, v \models \varphi_d^{\text{Win}}$ if and only if Player 0 has a winning strategy from v_0 in $\mathcal{G}$.

Proof (Idea). 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 some negligible modifications:

- eliminate moves after which the opponent wins in at most two steps (e.g. Verifier 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),
- 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.

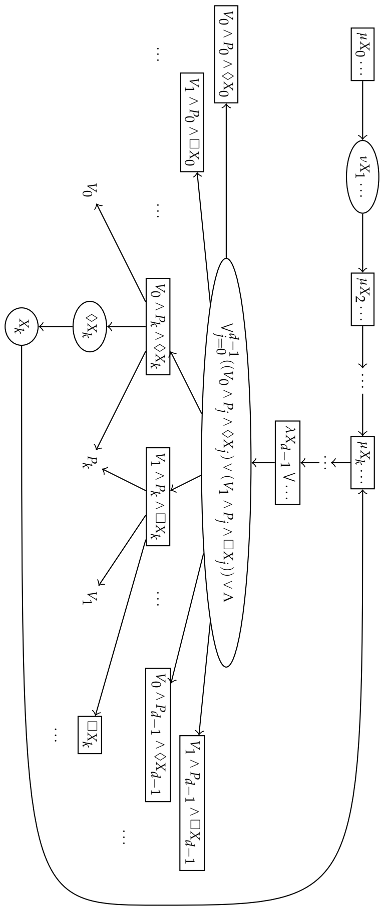

Figure 2.2. Part of a model checking game for φ_d^{Win}.