Lectures on the modal μ-calculus

Yde Venema*

Abstract

These notes give an introduction to the theory of the modal μ-calculus.

*Institute for Logic, Language and Computation, University of Amsterdam, Science Park 107, NL–1098XG Amsterdam. E-mail: `y.venema@uva.nl`.

Contents

5 Parity games

A large part of the theory of modal fixpoint logic involves nontrivial concepts and results from the theory of infinite games. In this chapter we discuss some of the highlights of this theory in a fair amount of detail. This allows us to be rather informal about game-theoretic concepts in the rest of the notes.

5.1 Board games

The games that we are dealing with here can be classified as *board* or *graph games*. They are played by two agents, here to be called 0 and 1.

Definition 5.1 If $\sigma \in \{0, 1\}$ is a player, then $\overline{\sigma}$ denotes the *opponent* $1 - \sigma$ of σ. $\lhd$

A board game is played on a *board* or *arena*, which is nothing but a directed graph, the nodes of which are usually referred to as *positions*. A *match* or *play* of the game consists of the two players moving a token across the board, following the edges of the graph. To regulate this, the collection of graph nodes, usually referred to as *positions* of the game, is partitioned into two sets, one for each player. Thus with each position we may associate a unique *owner*: the player whose turn it is to move when the token lies on this position.

Definition 5.2 A *board* or *arena* is a structure $\mathbb{B} = \langle B_0, B_1, E\rangle$, such that B_0 and B_1 are disjoint sets, and $E \subseteq B^2$, where $B := B_0 \cup B_1$. We will make use of the notation $E[p]$ for the set of *admissible* or *legitimate moves* from a board position $p \in B$, that is, $E[p] := \{q \in B \mid (p, q) \in E\}$. Positions not in $E[p]$ will sometimes be referred to as *illegitimate moves* with respect to p. A position $p \in B$ is a *dead end* if $E[p] = \varnothing$. If $p \in B$, we let σ_p denote the (unique) player such that $p \in B_{\sigma_p}$, and say that p *belongs to* σ_p, or that it is σ_p's *turn* to move at p. $\lhd$

Remark 5.3 Occasionally it will be convenient to represent a board in an alternative yet equivalent manner, viz., as a triple $\mathbb{B} = \langle B, E, \sigma\rangle$ such that (B, E) is a graph and $\sigma : B \to \{0, 1\}$ is a map assigning a player to each position in B. It is obvious how to switch from one presentation to another. $\lhd$

A match of the game may in fact be identified with the sequence of positions visited during play, and thus corresponds to a *path* through the graph. We refer to the Appendix A for some notation concerning paths.

Definition 5.4 A *path* through a board $\mathbb{B} = \langle B_0, B_1, E\rangle$ is a nonempty (finite or infinite) sequence $\pi \in B^\infty$ such that $E\pi_i\pi_{i+1}$ whenever applicable. A *full* or *complete match* or *play* through $\mathbb{B}$ is either an infinite $\mathbb{B}$-path, or a finite $\mathbb{B}$-path π ending with a dead end (i.e. $E[\mathsf{last}(\pi)] = \varnothing$).

A *partial match* is a finite path through $\mathbb{B}$ that is not a full match; in other words, the last position of a partial match is not a dead end. We let PM denote the set of all partial matches, and PM_σ the set of partial matches such that σ is the player whose turn it is to move at the last position of the match. In the sequel, we will denote this player as σ_π; that is, $\sigma_\pi := \sigma_{\mathsf{last}(\pi)}$. $\lhd$

Each full or completed match is *won* by one of the players, and *lost* by their opponent; that is, there are no draws. A finite match ends if one of the players gets *stuck*, that is, is forced to move the token from a position without successors. Such a finite, completed, match is lost by the player who got stuck.

If neither player ever gets stuck, an infinite match arises. The flavor of a board game is very much determined by the winning conditions of these infinite matches.

Definition 5.5 Given a board $\mathbb{B}$, a *winning condition* is a map $W : B^\omega \to \{0,1\}$. An infinite match π is *won* by $W(\pi)$. A *board game* is a structure $\mathcal{B} = \langle B_0, B_1, E, W\rangle$ such that $\langle B_0, B_1, E\rangle$ is a board, and W is a winning condition on B. $\lhd$

Although the winning condition given above applies to all infinite B-sequences, it will only make sense when applied to matches. We have chosen the above definition because it is usually much easier to formulate maps that are defined on all sequences.

Before players can actually start playing a game, they need a starting position. The following definition introduces some terminology and notation.

Definition 5.6 An *initialized board game* is a pair consisting of a board game $\mathcal{B}$ and a position q on the board of $\mathcal{B}$; such a pair is usually denoted $\mathcal{B}@q$.

Given a (partial) match π, its first element $\mathsf{first}(\pi)$ is called the *starting position* of the match. We let $\mathrm{PM}_\sigma(q)$ denote the set of partial matches for σ that start at position q. $\lhd$

It is sometimes convenient to represent an (initialised) game by its *game tree.*

Remark 5.7 Let a be a position in a board game $\mathcal{B} = (B_0, B_1, E, W)$. The *game tree* of $\mathcal{B}$ at a is given as the tree $\mathbb{T}^{\mathcal{B}}_a := (\mathrm{PM}(b), \vec{E}, a)$, where $\vec{E}$ is the relation

$$\vec{E} := \{(\pi, \pi \cdot b) \mid b \in E[\mathsf{last}(\pi)]\},$$

and a is the one-node match starting and ending at a. That is, $\mathbb{T}^{\mathcal{B}}_a$ is the *unravelling* of the frame (B, E) around a. Clearly, full matches of $\mathcal{B}@a$ correspond to *branches* of $\mathbb{T}^{\mathcal{B}}_a$, with finite matches corresponding to childless nodes, and infinite matches to infinite branches.

One may think of $\mathbb{T}^{\mathcal{B}}_a$ itself as a board game $\mathcal{B}^T_a$ by partitioning its nodes in the obvious way: $\mathrm{PM}(a) = \mathrm{PM}_0(a) \uplus \mathrm{PM}_1(a)$, and formulating the obvious winning conditions. We will usually confine our attention to the version of this game that is initialised at the root a. $\lhd$

Central in the theory of games is the notion of a *strategy.* Roughly, a strategy for a player is a method that the player uses to decide how to continue partial matches when it is their turn to move. More precisely, a strategy is a function mapping partial plays for the player to new positions. It is a matter of definition whether one requires a strategy to always assign moves that are legitimate, or not; here we will not make this requirement. For the following definition, recall that $\sqsubset$ denotes the initial segment relation on sequences (and thus, on matches).

Definition 5.8 Given a board game $\mathcal{B} = \langle B_0, B_1, E, W\rangle$ and a player σ, a *σ-strategy*, or a *strategy for σ*, is a map $f : \mathrm{PM}_\sigma \to B$. In case we are dealing with an initialized game $\mathcal{B}@q$, then we may take a strategy to be a map $f : \mathrm{PM}_\sigma(q) \to B$. A match π is *consistent with* or *guided by* a σ-strategy f if for any partial match $\pi' \sqsubset \pi$ with $\mathsf{last}(\pi') \in B_\sigma$, the next position on π (after π') is indeed the element $f(\pi')$.

A σ-strategy f is *surviving* in $\mathcal{B}@q$ if the moves that it prescribes to f-guided partial matches in $\mathrm{PM}_\sigma@q$ are always admissible to σ, and *winning for σ* in $\mathcal{B}@q$ if in addition all f-guided full matches starting at q are won by σ. A position $q \in B$ is *winning for σ* if σ has a winning strategy for the game $\mathcal{B}@q$; the collection of all winning positions for σ in $\mathcal{B}$ is called the *winning region for σ in $\mathcal{B}$*, and denoted as $\mathrm{Win}_\sigma(\mathcal{B})$. $\lhd$

Intuitively, f being a surviving strategy in $\mathcal{B}@q$ means that σ never gets stuck in an f-guided match of $\mathcal{B}@q$, and so guarantees that σ can stay in the game forever.

Remark 5.9 Where f is a surviving for player σ in $\mathcal{B}@a$, we may *represent* f as the pruned subtree of the game tree $\mathbb{T}^{\mathcal{B}}_a$ that is based on those nodes that correspond to f-guided matches of $\mathcal{B}@a$. In this so-called *strategy tree* $\mathbb{T}^f_a$ we have

$$\vec{E}_f[\pi] := \begin{cases} \vec{E}[\pi] & \text{if } \pi \in \mathrm{PM}_{\overline{\sigma}} \\ \{\pi \cdot f(\pi)\} & \text{if } \pi \in \mathrm{PM}_\sigma. \end{cases}$$

for any node $\pi \in \mathrm{PM}(a)$. $\lhd$

Convention 5.10 Observe that we allow strategies that prescribe illegitimate moves. In practice, it will often be convenient to extend the definition of a strategy even further to include maps f that are *partial* in the sense that they are only defined on a proper subset of PM_σ. We will only permit ourselves such a sloppiness if we can guarantee that $f(\pi)$ is defined for every $\pi \in \mathrm{PM}_\sigma$ that is consistent with the partial σ-strategy f, so that the situation where the partial strategy actually would fail to suggest a move, will never occur.

It is easy to see that a position in a game $\mathcal{B}$ cannot be winning for *both* players. On the other hand, the question whether a position p is always a winning position for *one* of the players, is a rather subtle one. Observe that in such games the two winning regions *partition* the game board.

Definition 5.11 The game $\mathcal{B}$ on the board $\mathbb{B}$ is *determined* if $\mathrm{Win}_0(\mathcal{B}) \cup \mathrm{Win}_1(\mathcal{B}) = B$; that is, each position is winning for one of the players. $\lhd$

It turns out that the axiom of choice implies the existence of infinite games that admit positions from which neither player has a winning strategy.

► `Add some more detail, including a remark on the axiom of determinacy in set theory.`

In principle, when deciding how to move in a match of a board game, players may use information about the entire history of the match played thus far. However, it will turn out to be advantageous to work with strategies that are simple to compute. Particularly nice are so-called *positional* strategies, which only depend on the current position (i.e., the final position of the partial play). Although their importance is sometimes overrated, positional strategies are convenient to work with, and they will be critically needed in the proofs of some of the most fundamental results in the automata-theoretic approach to fixpoint logic.

Definition 5.12 A strategy f is *positional* or *history-free* if $f(\pi) = f(\pi')$ for any π, π' with $\mathsf{last}(\pi) = \mathsf{last}(\pi')$. ◁

Convention 5.13 A positional σ-strategy may be represented as a map $f : B_\sigma \to B$.

As a slight generalisation of positional strategies, *finite-memory strategies* can be computed using only a finite amount of information about the history of the match. More details can be found in Exercise 5.2.

Before finishing this section we note that any set of positions on a board naturally induces a subboard, based on the restricted edge relation, and similarly for games.

Definition 5.14 Given a board $\mathbb{B} = \langle B_0, B_1, E\rangle$, a subset $A \subseteq B$ determines the following *subboard* $\mathbb{B}_A := \langle A \cap B_0, A \cap B_1, E_{\restriction A}\rangle$, where $E_{\restriction A} := E \cap (A \times A)$ is the *restriction* of E to A. Similarly, if $W : B^\omega \to \{0,1\}$ is a winning condition on $\mathbb{B}$, we let $W_{\restriction A} : A^\omega \to \{0,1\}$ denote the restriction of W to the set $A^\omega \subseteq B^\omega$, and with $\mathcal{B} = \langle B_0, B_1, E, W\rangle$, we define $\mathcal{B}_A := \langle A \cap B_0, A \cap B_1, E_{\restriction A}, W_{\restriction A}\rangle$ as the *subgame* induced by A. ◁

Note that, in the setting of this definition, a position $b \in A$ may be a dead end in $\mathbb{B}_A$ but not in $\mathbb{B}$. More in general, a winning strategy for one of the players in $\mathbb{B}_A@b$ is not necessarily a winning strategy for this player in $\mathbb{B}@b$.

5.2 Winning conditions

In case we are dealing with a *finite* board B, then we may nicely formulate winning conditions in terms of the set of positions that occur *infinitely often* in a given match. But in the case of an infinite board, there may be matches in which no position occurs infinitely often (or more than once, for that matter). Nevertheless, we may still define winning conditions in terms of objects that occur infinitely often, if we make use of *finite colorings* of the board. If we assign to each position $b \in B$ a *color*, taken from a finite set C of colors, then we may formulate winning conditions in terms of the infinite *color sequence* induced by an infinite match.

Definition 5.15 A *coloring* of B is a function $\Gamma : B \to C$ assigning to each position $p \in B$ a *color* $\Gamma(p)$ taken from some finite set C of colors. By putting $\Gamma(p_0 p_1 \cdots) := \Gamma(p_0)\Gamma(p_1)\cdots$ we can naturally extend such a coloring $\Gamma : B \to C$ to a map $\Gamma : B^\omega \to C^\omega$. ◁

Now if $\Gamma : B \to C$ is a coloring, for any infinite sequence $\pi \in B^\omega$, the map $\Gamma \circ \pi \in C^\omega$ forms the associated sequence of colors. Most if not all interesting types of graph games have winning conditions that can be expressed in terms of this induced sequence. An interesting example is given by the so-called *regular games*, that is, board games where the winning condition is given as an ω-regular language over some colouring of the board.

Definition 5.16 An infinite game $\mathcal{B} = \langle B_0, B_1, E, W\rangle$ is called *(ω-)regular* if there exists an ω-regular language L over some finite alphabet C and a colouring $\Gamma : B \to C$, such that player 0 wins a match $\pi = (p_i)_{i<\omega} \in B^\omega$ precisely if the induced sequence $(\Gamma(p_i))_{i<\omega} \in C^\omega$ belongs to L. $\lhd$

In Exercise 5.2 the reader is asked to prove that these regular games are determined, and that we may assume that the winning strategies of each player in a regular game are finite-memory strategies.

A special kind of regular game is given by the Muller game. Here, the definition of the winning condition is based on the observation that if an infinite game has an associated colouring, then for any infinite match there must be a nonempty set of colours that occur infinitely often in this stream.

Definition 5.17 Let $\mathbb{B}$ be a board and $\Gamma : B \to C$ a coloring of B. Given an infinite sequence $\pi \in B^\omega$, we let $\mathit{Inf}_\Gamma(\pi)$ denote the set of colors that occur infinitely often in the sequence $\Gamma \circ \pi$.

A *Muller condition* is a collection $\mathcal{M} \subseteq \wp(C)$ of subsets of C. The corresponding winning condition is defined as the following map $W_\mathcal{M} : B^\omega \to \{0, 1\}$:

$$W_\mathcal{M}(\pi) := \begin{cases} 0 & \text{if } \mathit{Inf}_\Gamma(\pi) \in \mathcal{M} \\ 1 & \text{otherwise.} \end{cases}$$

A *Muller game* is a board game of which the winning conditions are specified by a Muller condition. $\lhd$

In words, player 0 wins an infinite match $\pi = p_0 p_1 \cdots$ if the set of colors one meets infinitely often on this path, belongs to the Muller collection $\mathcal{M}$. It is not difficult to see that Muller games can be presented as regular games, and thus inherit the properties of determinacy and the existence of winning strategies that require only a finite amount of memory.

These results becomes even nicer if the Muller condition allows a formulation in terms of a *priority map*. In this case, as colors we take natural numbers. Note that by definition of a coloring, the range $\Omega[B]$ of the coloring function Ω is finite. This means that every nonempty subset of $\Omega[B]$ has a maximal element. Hence, every match determines a unique natural number, namely, the 'maximal color' that one meets infinitely often during the match. Now the parity winning condition states that the winner of an infinite match is 0 if this number is even, and 1 if it is odd. More succinctly, we formulate the following definition.

Definition 5.18 Let B be some set; a *priority map* on B is a coloring $\Omega : B \to \omega$, that is, a map of finite range. A *parity game* is a board game $\mathcal{B} = \langle B_0, B_1, E, W_\Omega\rangle$ in which the winning condition is given by

$$W_\Omega(\pi) := \max(\mathit{Inf}_\Omega(\pi)) \mod 2.$$

Such a parity game is usually denoted as $\mathcal{B} = \langle B_0, B_1, E, \Omega\rangle$. $\lhd$

▶ `Many examples`

As a variation of the above parity game, we sometimes consider a version where the priority map is only partial.

Definition 5.19 Let $\mathbb{B} = (B, E, \sigma)$ be a board. A *partial priority map* Ω for $\mathbb{B}$ is a partial map $\Omega : B \overset{\circ}{\to} \omega$ such that $\Omega[B]$ is finite and $\mathsf{Dom}(\Omega)$ is cycle-critical in $\mathbb{B}$, that is, every cycle in $\mathbb{B}$ contains at least one position in $\mathsf{Dom}(\Omega)$. To define the winning conditions for these games, we observe that the set $Inf_{\Omega}(\pi)$ is non-empty, for any infinite match π, so that we may define

$$W_{\Omega}(\pi) := \max(Inf_{\Omega}(\pi)) \mod 2,$$

just as in the case of a total priority map, A *ppm-parity game* is a game of the form (B_0, B_1, E, W_ω), where W is induced by a partial priority map Ω for $\mathbb{B}$. ◁

Obviously, any standard parity game corresponds to a ppm-parity game where the priority map is total. Conversely, it is straightforward to transform a ppm-parity game into a standard parity game — we leave the details for the reader.

The key property that makes parity games so interesting and nice to work with is that they enjoy positional determinacy, as we will prove in section 5.4. First, however, we turn to the special case of so-called reachability games.

5.3 Reachability games and attractor sets

Reachability games are a special kind of board games. They are played on a board such as described in section 5.1, but also feature a special set of positions. The aim of the game is for one player to move the token into this special set and for the other to avoid this to happen.

Definition 5.20 Fix a board $\mathbb{B}$ and a subset $A \subseteq B$. The reachability game $\mathcal{R}_\sigma(\mathbb{B}, A)$ is defined as the game over $\mathbb{B}$ in which σ wins as soon as a position in A is reached or if $\overline{\sigma}$ gets stuck. On the other hand, $\overline{\sigma}$ wins if he can manage to keep the token outside of A infinitely long, or if σ gets stuck. ◁

As an example, if $A = \varnothing$, in order to win the game $\mathcal{R}_\sigma(\mathbb{B}, A)$ for player $\overline{\sigma}$ it simply suffices to stay alive forever, while σ can only win by forcing $\overline{\sigma}$ to get stuck.

Remark 5.21 If we want reachability games to fit the format of a board game exactly, we have to modify the board, as follows. Given a reachability game $\mathcal{R}_\sigma(\mathbb{B}, A)$, define the board $\mathbb{B}' := \langle B'_0, B'_1, E' \rangle$ by putting:

$$\begin{array}{lcl} B'_\sigma & := & B_\sigma \setminus A \\ B'_{\overline{\sigma}} & := & B_{\overline{\sigma}} \cup A \\ E' & := & \{(p, q) \in E \mid p \notin A\}. \end{array}$$

In other words, $\mathbb{B}'$ is like $\mathbb{B}$ except that all positions in A are now dead ends, assigned to player $\overline{\sigma}$. This means that $\overline{\sigma}$ gets stuck in a position belonging to A. Furthermore, the winning conditions of such a game are very simple: simply define $W : B^\omega \to \{0, 1\}$ as the constant function mapping all infinite matches to $\overline{\sigma}$. This can easily be formulated as a parity condition. ◁

Since reachability games can thus be formulated as very simple parity games, the following theorem, stating that reachability games enjoy positional determinacy, can be seen as a warming up exercise for the general case. Alternatively, reachability can be seen as an *unfolding game* for some monotone functional over the power set of the board, so that its positional determinacy follows from $\boxed{\text{X}}$. We leave the proof details as an exercise for the reader.

Theorem 5.22 (Positional determinacy of reachability games) *Let $\mathcal{R}$ be a reachability game. Then there are positional strategies f_0 and f_1 for 0 and 1, respectively, such that for every position q there is a player σ such that f_σ is a winning strategy for σ in $\mathcal{R}@q$.*

Definition 5.23 The winning region for σ in $\mathcal{R}_\sigma(\mathbb{B}, A)$ is called the *attractor set* of σ for A in $\mathbb{B}$, notation: $Attr_\sigma^{\mathbb{B}}(A)$. In the sequel we will fix a positional winning strategy for σ in $\mathcal{R}_\sigma(\mathbb{B}, A)$ and denote it as $attr_\sigma^{\mathbb{B}}(A)$. $\lhd$

If no confusion arises, we will generally drop the superscript referring to the board, writing $Attr_\sigma(A)$ and $attr_\sigma(A)$ instead of $Attr_\sigma^{\mathbb{B}}(A)$ and $attr_\sigma^{\mathbb{B}}(A)$. Note that σ-attractor sets always contain all points from which σ can make sure that $\overline{\sigma}$ gets stuck. Furthermore, it is easy to see that in $attr_\sigma(A)$-guided matches the token never leaves $Attr_\sigma(A)$ (at least if the match starts inside $Attr_\sigma(A)$!).

Proposition 5.24 *Let $\mathbb{B}$ be some board. Then $Attr_\sigma^{\mathbb{B}}$ is a closure operation on $\wp(B)$, i.e.*

1. *$A \subseteq A'$ implies $Attr_\sigma^{\mathbb{B}}(A) \subseteq Attr_\sigma^{\mathbb{B}}(A')$,*
2. *$A \subseteq Attr_\sigma^{\mathbb{B}}(A)$,*
3. *$Attr_\sigma^{\mathbb{B}}(Attr_\sigma^{\mathbb{B}}(A)) = Attr_\sigma^{\mathbb{B}}(A)$.*

Proof. For part 1, simply observe that any strategy for σ that forces the token eventually to A, at the same time forces the token into A'. Part 2 is immediate by the definitions.

For part 3, we only need to prove the inclusion $\subseteq$, since the opposite inclusion follows from part 2. Take some arbitrary position $b \in Attr_\sigma^{\mathbb{B}}(Attr_\sigma^{\mathbb{B}}(A))$. That is, σ has some strategy f that is guaranteed to eventually lead the token into $Attr_\sigma^{\mathbb{B}}(A)$. Suppose then that σ uses this strategy, and as soon as the match has arrived to $Attr_\sigma^{\mathbb{B}}(A)$, she switches to the attractor strategy for A. It should be obvious that this is a winning strategy in the reachability game of A, starting from a. QED

A kind of counterpart to attractor sets are *closed sets.* In words, a set A is σ-closed if σ has the power to keep the token inside A while not getting stuck.

Definition 5.25 Given a board $\mathbb{B}$, we call a subset $A \subseteq B$ *σ-closed* (or a *$\overline{\sigma}$-trap*) if $E[b] \subseteq A$ for all $b \in A \cap B_{\overline{\sigma}}$, while $E[b] \cap A \neq \varnothing$ for all $b \in A \cap B_\sigma$. $\lhd$

Note that a σ-closed set does not contain σ-endpoints and that σ will therefore never get stuck in a σ-closed set. We conclude this section with a useful proposition.

Proposition 5.26 *Let $\mathbb{B}$ be a board and $A \subseteq B$ an arbitrary subset of B. Then the following assertions hold.*

1. *If A is σ-closed then so is $Attr_\sigma^{\mathbb{B}}(A)$.*
2. *The complement of $Attr_\sigma^{\mathbb{B}}(A)$ is $\overline{\sigma}$-closed.*
3. *The union $\bigcup\{A_i \mid i \in I\}$ of an arbitrary collection of σ-closed sets is again σ-closed.*
4. *If A is σ-closed in $\mathbb{B}$ then any $C \subseteq A$ is σ-closed in $\mathbb{B}$ iff C is σ-closed in $\mathbb{B}_A$.*
5. *If $C \subseteq B$ is σ-closed, then $Attr_\sigma^{\mathbb{B}_C}(A \cap C) \subseteq Attr_\sigma^{\mathbb{B}}(A) \cap C$.*

Proof. All statements are easily verified and thus the proof is left to the reader. QED

5.4 Positional Determinacy of Parity Games

Probably the most attractive property of parity games is that we may always assume that players have a *positional* winning strategy starting from any of their winning positions.

Definition 5.27 A board game $\mathcal{B} = \langle B_0, B_1, E, W \rangle$ enjoys *positional determinacy* if $B = \text{Win}_0(\mathcal{B}) \cup \text{Win}_1(\mathcal{B})$, and there are positional strategies f_0 and f_1 such that for each player σ and every $b \in \text{Win}_\sigma(\mathcal{B})$, the strategy f_σ is winning for σ in $\mathcal{B}@b$. $\lhd$

Theorem 5.28 (Positional Determinacy of Parity Games) *Let $\mathcal{B}$ be a parity game. Then $\mathcal{B}$ enjoys positional determinacy.*

Related to positional determinacy, and an important tool in the proof of Theorem 5.28 is the notion of a player's *paradise*. In words, a subset $A \subseteq B$ is a σ-paradise if σ has a positional strategy f which guarantees, for any position $a \in A$, both that she wins the game $\mathcal{B}@a$, and that the token stays in A. Thus in particular, a σ-paradise is σ-closed.

Definition 5.29 Given a board game $\mathcal{B}$, we call a σ-closed set A a *σ-paradise* if there exists a positional winning strategy $f : A \cap B_\sigma \to A$. $\lhd$

The importance of this notion is clear from the following proposition. We omit its proof since it is straightforward.

Proposition 5.30 *Let $\mathcal{B}$ be some board game. Then $\mathcal{B}$ enjoys positional determinacy iff its board can be partitioned into a* 0*-paradise and a* 1*-paradise.*

The remainder of this section is organized as follows. We first give a straightforward proof the positional determinacy of parity games in the finite case. In the next subsection we discuss the notion of a paradise in more detail, and in the final subsection we apply our findings to prove Theorem 5.28 in the general setting.

5.4.1 Positional determinacy of parity games: the finite case

In the case where $\mathcal{B}$ is based on a *finite board*, we can give a relatively straightforward proof of positional determinacy.

Proposition 5.31 *Let $\mathcal{B}$ be a finite parity game. Then its board can be partitioned into a* 0*-paradise and a* 1*-paradise.*

Proof. Let $\mathbb{B} = (B_0, B_1, E)$ be the board of $\mathcal{B}$, and let Ω be its priority map. We will prove this proposition by induction on d, the maximal parity in the game (i.e. $d := \max(\Omega[B])$). If $d = 0$ we are dealing with a reachability game (namely $\mathcal{R}_1(\mathbb{B}, \varnothing)$), and so the result follows by Theorem 5.22.

We thus focus on the inductive case, where $d \geq 1$. Without loss of generality we may assume that d is even — in the case where d is odd we proceed in a similar way. The key observation in the finite case is that we may use an inner induction on the size $|\mathcal{B}|$ of $\mathcal{B}$, that is, the number of its positions.

Define $M := \{b \in X \mid \Omega(b) = d\}$ to be the set of all positions in $\mathcal{B}$ that actually reach the maximum priority d and let $M^+ := Attr_0^{\mathbb{B}}(M)$ be its attractor set. Note that M, and hence also M^+, must be nonempty.

We may apply the (outer or inner) inductive hypothesis to the subgame $\mathcal{H}$ of $\mathcal{B}$ that is based on the set $B \setminus M^+$. Let A_0 and A_1, respectively, be the winning regions for the two players of $\mathcal{H}$, with positional winning strategies, respectively, g_0 and g_1. Furthermore, let $A_1^+ := Attr_0^{\mathbb{B}}(M)$ be the attractor set of A_1 (with respect to the full board $\mathbb{B}$). Note that whereas A_1 must be disjoint from M^+, the set A_1^+ may overlap with M^+, and even with M.

We now make a case distinction as to whether A_1^+ is empty or not.

Case $A_1^+ \neq \varnothing$. In this case we may apply the (inner) induction hypothesis to the subgame $\mathcal{B}^-$ of $\mathcal{B}$ that is based on the complement $B \setminus A_1^+$ of A_1^+. This gives us a partition of the board of $\mathcal{B}^-$ into winning regions W_0^- and W_1^-, together with positional winning strategies f_0^- and f_1^- such that for every position $b \in W_\sigma^-$, the strategy f_σ is winning for σ in $\mathcal{B}^-@b$.

Claim 1 1) $W_0 := W_0^-$ is a 0-paradise in $\mathcal{B}$, with positional winning strategy f_0^-;

2) $W_1 := W_1^- \cup A_1^+$ is a 1-paradise in $\mathcal{B}$, with positional winning strategy f_1 given by:

$$f_1(b) := \begin{cases} g_1(b) & \text{if } b \in A_1 \\ attr_1^{\mathbb{B}}(A_1)(b) & \text{if } b \in A_1^+ \setminus A_1 \\ f_1^-(b) & \text{if } b \in W_1^-. \end{cases}$$

Proof of Claim We leave the verification of part (1) as an exercise for the reader. For the second part, consider an arbitrary f_1-guided match π starting from some position $b \in W^1$. If $b \in A_1$, it should be clear that π stays in A_1: player 1 will keep the token there, while player 0 cannot leave A_1 since it is 1-closed. It follows that π is a g_1-guided $\mathcal{H}$-match; it is thus won by 1 in $\mathcal{H}$, and hence, also in $\mathcal{B}$. It is then easy to see that f_1 is also winning when the match starts at some $b \in A_1^+ \setminus A_1$ — we leave the details for the reader.

Finally then, if $b \in W_1^-$, then it is not hard to show that π will stay in W_1^-, unless player 0 manages to move the token out of there. However, player 0 will not be able to move the token to a position in W_0^- — if so, the position from which this were possible would itself belong to W_0^-. The only way out for player 0, then, is to move the token to a position in A_1^+; but if she does so then her opponent simply switches to the positional winning strategy described in the above paragraph. Hence every match that leaves W_1^- is won by player 1. But if the match stays in W_1^-, it is clearly a win for him as well. ◀

Case $A_1^+ = \varnothing$. Clearly this implies that also $A_1 = \varnothing$, and that $E[b] \neq \varnothing$, for any position $b \in B_0$. We may thus pick an element $k(b) \in E[b]$, for any $b \in B_0$. Furthermore, A_1 being empty implies that every position in $B \setminus M^+$ is winning for 0 in $\mathcal{H}$. This means in particular, that $g_0(b)$ is defined for every $b \in B_0 \setminus M^+$. We now claim that every position in $\mathcal{B}$ is winning for player 0, and that she may use the positional strategy f_0, which is defined by putting, for $b \in B_0$:

$$f_0(b) := \begin{cases} k(b) & \text{if } b \in M \\ attr_1^{\mathbb{B}}(M)(b) & \text{if } b \in M^+ \setminus M \\ g_0(b) & \text{if } b \in B \setminus M^+. \end{cases}$$

We leave the proof of this claim as an exercise for the reader.

QED

5.4.2 Paradises

As mentioned, the proof of positional determinacy for parity games in the general case is slightly more involved. We need the following proposition which establishes some basic facts about paradises. In this section we also provide some further observations about paradises that will be of use when we discuss algorithms for solving parity games in section 5.5.

Proposition 5.32 *Let $\mathcal{B}(\mathbb{B}, \Omega)$ be a parity game. Then the following hold:*

1. *The union $\bigcup\{P_i \mid i \in I\}$ of an arbitrary set of σ-paradises is again a σ-paradise.*
2. *There exists a largest σ-paradise.*
3. *If P is a σ-paradise then so is $Attr_\sigma(P)$.*

Proof. The main point of the proof of part (1) is that we somehow have to uniformly choose a strategy on the intersection of paradises, such that we will end up following the strategy of only one paradise. For this purpose, we assume that we have a well-ordering on the index set I (i.e., for the infinite case we assume the Axiom of Choice).

For the details, assume that $\{P_i \mid i \in I\}$ is a family of paradises, and let f_i be the positional winning strategy for P_i. Note that $P := \bigcup\{P_i \mid i \in I\}$ is σ-closed by Proposition 5.26. Assume that $<$ is a well-ordering of I, so that for each $q \in P$ there is a *minimal* index $\min(q)$ such that $q \in P_{\min(q)}$. Define a positional strategy on P by putting

$$f(q) := f_{\min(q)}(q).$$

This strategy ensures at all times that the token either stays in the current paradise, or else it moves to a paradise of lower index, and so, any match where σ plays according to f will proceed through a sequence of σ-paradises of decreasing index. Because of the well-ordering, this decreasing sequence of paradises cannot be strictly decreasing, and thus we know that after finitely many steps the token will remain in the paradise where it is, say, P_j. From that moment on, the match is continued as an f_j-guided match inside P_j, and since f_j is by assumption a winning strategy when played inside P_j, this match is won by σ.

Part (2) of the proposition should now be obvious: clearly the union of all σ-paradises is the greatest σ-paradise.

In order to prove part (3) we need to show that there exists a positional winning strategy for σ starting from any position in $Attr_\sigma(P)$. The principal idea is to first move to P by $attr_\sigma(P)$ and once there to follow the positional winning strategy in P. Let f' be the winning strategy for P, we then define the following strategy f on $Attr_\sigma(P)$ by

$$f(p) := \begin{cases} f'(p) & \text{if } p \in P \\ attr_\sigma(P)(p) & \text{otherwise.} \end{cases}$$

A match consistent with this strategy will stay in $Attr_\sigma(P)$ because this set is σ-closed and $f(p) \in Attr_\sigma(P)$ for all $p \in Attr_\sigma(P)$. It is winning because if ever the match arrives at a point $p \in P$ then play continues as if the match were completely in P; and since f' was supposed to be a winning strategy for σ this play is won by σ. If we start outside P we will at first follow the strategy $attr_\sigma(P)$ which will ensure that σ either wins or that the token ends up in P, in which case σ will also win. QED

Below we list some properties of paradises that will become useful in the next section. The proof of the first proposition is left as a (fairly straightforward) exercise for the reader.

Proposition 5.33 *Let $\mathcal{B} = (\mathbb{B}, \Omega)$ be a parity game, and let $D \subseteq B$ be a σ-paradise, witnessed by a positional winning strategy f. Then the following hold:*

1. *Let $C \subseteq B$ be σ-closed. If $D \subseteq C$ then D is also a σ-paradise in $\mathcal{B}_C$.*
2. *If $A \subseteq B$ is such that $A \cap D = \varnothing$ then D is a σ-paradise in $\mathcal{B}_X$, where $X := B \setminus Attr^{\mathbb{B}}_{\overline{\sigma}}(A)$.*
3. *if $A \subseteq B$ is $\overline{\sigma}$-closed then $D \cap A$ is a σ-paradise in $\mathcal{B}_A$.*

In each case, the statement is witnessed by the positional strategy f (restricted to the domain $D \cap A$, in the third item).

For the following proposition, recall that in a parity game $\mathcal{B} = (B_0, B_1, E, \Omega)$, $\Omega^{-1}[d]$ denotes the set of positions of priority d.

Proposition 5.34 *Let $\mathcal{B} = (B_0, B_1, E, \Omega)$ be a parity game, and let $D \subseteq B$ be a nonempty σ-paradise in $\mathcal{B}$. Write $d := \max\big(\Omega[D]\big)$, and assume that σ is the opponent of $d \bmod 2$. Then D has a nonempty subset $C \subseteq D$ which is a σ-paradise, both in $\mathcal{B}$ and in $\mathcal{B}_D$, and which does not overlap with $Attr^{\mathbb{B}}_{\overline{\sigma}}(\Omega^{-1}[d])$.*

Proof. Let f be player σ's positional winning strategy for D. Every f-guided $\mathcal{B}$-match π starting from some $b \in D$ stays in D, and therefore, the set $Occ(\Omega[\pi])$ of all priorities encountered during π satisfies $\max\big(Occ(\Omega[\pi])\big) \leq d$. Define

$$C := \big\{ b \in D \mid \max\big(Occ(\Omega[\pi])\big) < d \text{ for every } f\text{-guided match } \pi \text{ of } \mathcal{B}@b \big\},$$

that is, C consists of those position in D for which σ can guarantee that f-guided matches *never* pass a position of priority d (or higher).

Clearly then, C is a σ-paradise (both in $\mathcal{B}$ and in $\mathcal{B}_D$), witnessed by the strategy f (restricted to C). We also claim that C is nonempty, for otherwise there would be, for *every*

$a \in D$, a finite f-guided match π, starting at a and ending at a position $a' \in D$ such that $\Omega(a') = d$. But then we may easily construct, starting from an arbitrary position in D, an f-guided match that passes a d-position *infinitely* often; this match would then be won by d mod $2 = \overline{\sigma}$, clearly contradicting our assumptions on f.

It remains to show that C does not overlap with the $\overline{\sigma}$-attractor set of $\Omega^{-1}[d]$, so assume for contradiction that $a \in C \cap Attr^{\mathbb{B}}_{\overline{\sigma}}(\Omega^{-1}[d])$. Now suppose that σ plays her strategy f against her opponent's attractor strategy, starting at this position a. The resulting match must reach a state in $\Omega^{-1}[d]$, but cannot leave D (with f being its paradise strategy). But then a cannot be an element of C, contradicting our assumption. QED

5.4.3 Positional determinacy of parity games: the general case

In this subsection we prove the positional determinacy of arbitrary (that is, not necessarily finite) parity games.

Proposition 5.35 *Let $\mathcal{B}$ be a parity game. Then its board can be partitioned into a 0-paradise and a 1-paradise.*

Proof. We will prove this proposition by induction on d, the maximal parity in the game (i.e. $d = \max(\Omega[B])$). The base case of the induction is the same as in the finite case, so we move to the case where $d \geq 1$. Let $\sigma := d \mod 2$, that is, σ wins an infinite play π if $\max(Inf(\pi)) = \max(\Omega[B]) = d$. Let $P_{\overline{\sigma}}$ be the maximal $\overline{\sigma}$-paradise, with associated positional strategy f. It now suffices to show that $X := B \setminus P_{\overline{\sigma}}$ is a σ-paradise. This proof is depicted in Figure 3.

First we shall show that X is σ-closed. By proposition 5.32(3) it follows that $Attr_{\overline{\sigma}}(P_{\overline{\sigma}})$ is itself also a $\overline{\sigma}$-paradise. By maximality of $P_{\overline{\sigma}}$ and the fact that $Attr_{\overline{\sigma}}$ is a closure operation, it follows that $P_{\overline{\sigma}} = Attr_{\overline{\sigma}}(P_{\overline{\sigma}})$. Thus by Proposition 5.26(2) we see that X, being the complement of a $\overline{\sigma}$-attractor set is σ-closed indeed.

Consider $\mathcal{B}_X$, the subgame of $\mathcal{B}$ induced by X. Define $M := \{b \in X \mid \Omega(b) = d\}$ to be the set of all points in X with priority d and let $Z := X \setminus Attr^{\mathbb{B}_X}_{\sigma}(M)$. Since Z is the complement of a σ-attractor set in $\mathbb{B}_X$ it is $\overline{\sigma}$-closed in $\mathbb{B}_X$.

By the induction hypothesis we can split the subgame $\mathcal{B}_Z$ into a 0-paradise Z_0 and a 1-paradise Z_1, see the picture. The positional winning strategies in these paradises we call f_0 and f_1 respectively. (All notions are with regard to the game $\mathcal{B}_Z$.) We claim that

$$Z_{\overline{\sigma}} = \varnothing. \tag{51}$$

To prove this, we first show that $P_{\overline{\sigma}} \cup Z_{\overline{\sigma}}$ is a $\overline{\sigma}$-paradise in $\mathcal{B}$. Consider the following strategy g for $\overline{\sigma}$:

$$g(b) := \begin{cases} f(b) & \text{if } b \in P_{\overline{\sigma}} \\ f_{\overline{\sigma}}(b) & \text{if } b \in Z_{\overline{\sigma}}. \end{cases}$$

It is left as an exercise for the reader to show that this is indeed a positional winning strategy for $\overline{\sigma}$ in $\mathcal{B}$, which in addition keeps the token inside $P_{\overline{\sigma}} \cup Z_{\overline{\sigma}}$. By the definition of $P_{\overline{\sigma}}$ as the maximal $\overline{\sigma}$-paradise, we see that $P_{\overline{\sigma}} = P_{\overline{\sigma}} \cup Z_{\overline{\sigma}}$ and since $P_{\overline{\sigma}}$ and $Z_{\overline{\sigma}}$ are disjoint we conclude that $Z_{\overline{\sigma}}$ must be empty indeed. This proves (51).

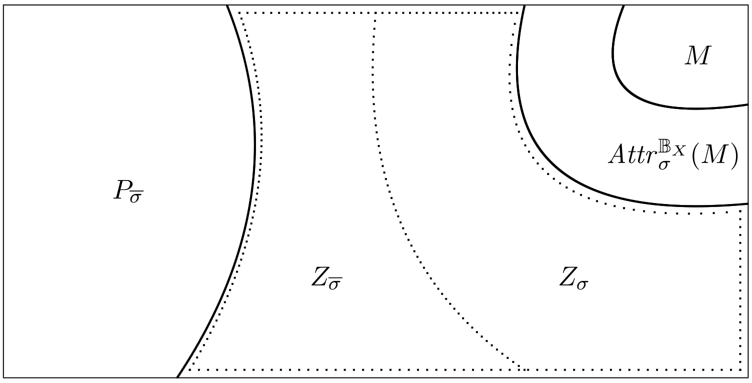

Figure 3: The proof of Proposition 5.35 in a picture

But from $Z_{\overline{\sigma}} = \varnothing$ it immediately follows that $Z = Z_\sigma$, so that we find

$$X = Z_\sigma \cup Attr_{\sigma}^{\mathbb{B}_X}(M).$$

Recall that X is σ-closed, so that for each $b \in X \cap B_\sigma$, we may pick an arbitrary element $k(b) \in E[b] \cap X$. Now define the following strategy h in $\mathcal{B}$ for σ on X.

$$h(b) := \begin{cases} k(b) & \text{if } b \in M \\ attr_\sigma(M)(b) & \text{if } b \in Attr_{\sigma}^{\mathbb{B}_X}(M) \setminus M \\ f_\sigma(b) & \text{if } b \in Z_\sigma = Z. \end{cases}$$

It is left as an exercise for the reader to show that h is indeed a winning strategy for σ in $\mathcal{B}$ and that it keeps the token in X. QED

5.5 Algorithms for solving parity games

Solving a parity game means to determine the players' winning regions in the game; by determinacy it suffices to determine the winning region of one player. In the case of a *finite* game, there is an obvious interest in finding *efficient algorithms* for this task — one of the reasons for this interest lies in the connection with the *model checking* problem for the modal μ-calculus that we will discuss in the next Chapter. The algorithmic angle also brings up the theoretical issue of the *computational complexity* of solving parity games, one of the most important open problems in the area. These algorithmic and complexity theoretic aspects of parity games form the topic of this section. Before going into the details we mention some notational and terminological conventions that apply in this section.

Convention 5.36 In this section it will be convenient to admit the *empty* game: this allows for much simpler formulation of the recursion base of our algorithms. Furthermore, in the context of a fixed ambient game $\mathcal{B}$, we will write $Attr_\sigma(A, D)$ rather than $Attr^{\mathcal{B}_D}(A)$ for the attractor set of A in the subgame $\mathcal{B}_D$, and similarly we use $Attr(\mathbb{B}, A)$ as the notation for the attractor set $Attr^{\mathbb{B}}(A)$. We will also frequently be sloppy about the distinction between a subset and the game it induces.

5.5.1 A quadratic algorithm for reachability games

We first consider the special case where the parity game is actually a reachability game; that is, we are interested in algorithms for computing attractor sets. The following theorem states that these games can be solved in *linear* time if we measure the board by its nunber of edges; obviously this algorithm is then *quadratic* in terms of the number of positions.

Theorem 5.37 *Let $\mathbb{B} = (B_0, B_1, E)$ be some board $\mathbb{B}$ and let $A \subseteq B$ be a set of positions. Then there is an algorithm* **R** *that computes the attractor set $Attr(\mathbb{B}, A)$ and terminates in time $\mathcal{O}(|E|)$.*

▶ `Provide the algorithm and prove the theorem`

5.5.2 An algorithm for solving parity games

As our starting point we take the algorithm **Z** of Figure 4, which is basically the algorithmic presentation of our proof for positional determinacy in the finite case (Proposition 5.31). Based on some ambient parity game $\mathcal{B}$, **Z** computes, given a subgame G of $\mathcal{B}$ with priorities up to some even number d, a subset $\texttt{solve}_0(G, d)$ of G which constitutes exactly player 0's winning region in $\mathbb{B}_G$. Dually, in case d is odd **Z** computes a subset $\texttt{solve}_1(G, d)$ consisting of player 1's winning region .

The correctness of **Z** is secured by the proof of Proposition 5.31, so we refrain from giving details here. For a rough complexity analysis of **Z**, we focus on its recursive calls. We will think of the first of these, in line 4, as the main recursion, and of the second, in line 7, as an iteration — this perspective can be made precise by replacing the second recursive call with a while-loop. The depth of the main recursion is bounded by d, while the number of iterations is bounded by n. This means, basically, that the algorithm is polynomial in n, but *exponential* in d.

Algorithm Z: $\texttt{solve}_0(G, d)$
1. $\texttt{if } G = \varnothing \texttt{ then return } G$
2. $M_d := \{b \in G \mid \Omega(b) = d\}$
3. $H := G \setminus \mathit{Attr}_0(M_d, G)$
4. $A_1 := \texttt{solve}_1(H, d-1)$
5. $G_1 := G \setminus \mathit{Attr}_1(A_1, G)$
6. $W_0 := \texttt{solve}_0(G_1, d)$
7. $\texttt{return } W_0$

Figure 4: Basic algorithm **Z** for solving parity games

5.5.3 A quasi-polynomial algorithm for solving parity games

Interestingly, the algorithm **Z** for solving parity games that we discussed in the previous subsection can be upgraded to a so-called *quasi-polynomial* algorithm.

An algorithm is said to take *quasi-polynomial time* if its time complexity is quasi-polynomially bounded. This means that there is some constant c such that the algorithm, when given an input of size n, is guaranteed to stop after at most

$$2^{\mathcal{O}((\log n)^c)} \tag{52}$$

computation steps. Note that if we take $c = 1$ in (52) we obtain the class of all polynomial functions — the *degree* of the polynomial is given by the constant that is hidden in the $\mathcal{O}$-notation. If $c > 1$ the function will not be polynomial, but its growth rate will still be substantially smaller than exponential.

The idea underlying the upgrade of Z is that, instead of each recursive call of the procedure $\texttt{solve}_\sigma$ *completely* solving a subgame, the algorithm will be based on a weaker assumption, viz., that each call of $\texttt{solve}_\sigma$ returns a partition of the board that separates *small* paradises of the two players, up to a size specified by a pair or explicit parameters. The key observation is then that, in order to find all paradises for player σ of size, say, up to p, it suffices to search for all paradises of size up to $\lfloor p/2 \rfloor$, and *at most one* paradise of size $\geq \lfloor p/2 \rfloor$. (Recall that $\lfloor q \rfloor$ denotes the largest natural number smaller than or equal to q.)

The quasi-polynomical upgrade **L** of **Z** is presented in Figure 5. Based on some ambient parity game $\mathcal{B}$, **L** computes a function $\texttt{solve}_0$ as follows. Given a subset G of B, again with priorities up to some even number d, **L** returns a subset $\texttt{solve}_0(G, d, p_0, p_1) \subseteq G$ which contains all 0-paradises in $\mathbb{B}_G$ of size up to p_0, and which does not intersect with any 1-paradise in $\mathbb{B}_G$ of size up to p_1. Dually, **L** computes a map $\texttt{solve}_1(G, d, p_1, p_0)$ for player 1 and odd d.

The main observation about **L** is that it makes *three* recursive calls, which compute three 'successive' subsets of G: $G \supseteq G_1 \supseteq G_2 \supseteq G_3$. The goal is then that G_3 contains all 0-paradises of size up to p_0; and that the three recursive calls of **L** collect all 1-paradises of size $\leq p_1$ in $G \setminus G_3$, by successively fishing in $G \setminus G_1$, $G_1 \setminus G_2$, and $G_2 \setminus G_3$. Crucially, the first and third recursive call use only *half* of the precision parameter for 1-paradises: only the second call asks for full precision. The key step in the correctness proof of **L** is that the first two

Algorithm L: $\texttt{solve}_0(G, d, p_0, p_1)$
$\texttt{if}\ p_1 = 0\ \texttt{then return}\ G$
$\texttt{if}\ p_0 = 0\ \texttt{then return}\ \varnothing$
$\texttt{if}\ d = 0\ \texttt{then return}\ Attr_0(G, \varnothing)$
$G_1 := \texttt{solve}_0(G, d, p_0, \lfloor p_1/2 \rfloor)$
$M_d := \{b \in G_1 \mid \Omega(b) = d\}$
$H := G_1 \setminus Attr_0(M_d, G_1)$
$A_1 := \texttt{solve}_1(H, d{-}1, p_1, p_0)$
$G_2 := G_1 \setminus Attr_1(A_1, G_1)$
$G_3 := \texttt{solve}_0(G_2, d, p_0, \lfloor p_1/2 \rfloor)$
$\texttt{return}\ G_3$

Figure 5: Quasi-polynomial algorithm **L** for solving parity games

recursive calls have already collected over the half of any 1-paradise of size p_0, so that the final call takes care of what is left.

Theorem 5.38 *Let $\mathcal{B} = (B_0, B_1, E, \Omega)$ be a finite parity game, let $d \geq \max\big(\mathsf{Ran}[\Omega]\big)$ and let $\sigma := d \bmod 2$. Then for any $G \subseteq B$, the algorithm* **L** *returns a set $\texttt{solve}_\sigma(G, d, p_\sigma, p_{\overline{\sigma}})$ of nodes such that*

(i) $\texttt{solve}_\sigma(G, d, p_\sigma, p_{\overline{\sigma}})$ contains each σ-paradise in $\mathcal{B}_G$ of size up to p_σ; and

(ii) $\texttt{solve}_\sigma(G, d, p_\sigma, p_{\overline{\sigma}})$ does not overlap with any $\overline{\sigma}$-paradise in $\mathcal{B}_G$ of size up to $p_{\overline{\sigma}}$.

Proof. By induction on the sum $d + p_0 + p_1$ we will prove the statements (i) and (ii), as well as the following statement:

(iii) the set $\texttt{solve}_\sigma(B, d, p_\sigma, p_{\overline{\sigma}})$ is σ-closed.

As the basis of the induction we take the cases where at least one of d, p_0 or p_1 is zero. Since in any of these cases the proof is trivial, we may focus on the inductive case, where we may assume that d, p_0 and p_1 are all bigger than zero. Without loss of generality we assume that d is even, so that $\sigma = 0$. In the case where d is odd we proceed in a completely analogous way.

We start with proving the statement (iii). By the induction hypothesis we find that G_1 is 0-closed in G, and that G_3 is 0-closed in G_2. Hence by Proposition 5.26(4) it suffices to show that G_2 is 0-closed in G_1. But this follows by Proposition 5.26(2), since G_2 is the complement in G_1 of a 1-attractor set.

In order to prove statement (i), let D be a 0-paradise in G of size at most p_0; we need to show that $D \subseteq G_3$. This proof is depicted in Figure 6.

By the induction hypothesis we have $D \subseteq G_1$. Moreover, D is a 0-paradise in G_1: this follows from Proposition 5.33(3) and the fact that $G_1 = G \setminus Attr_1(G \setminus G_1, G)$, being the complement of a 0-attractor, is 1-closed in G. We now claim that

$$D \text{ is a 0-paradise in } G_2. \tag{53}$$

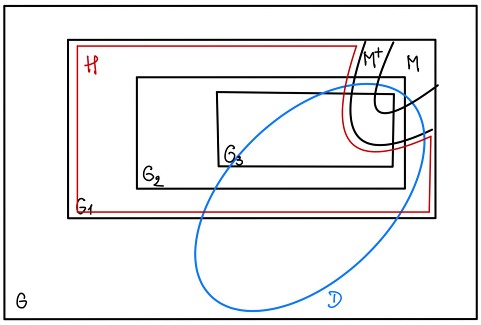

Figure 6: The proof of Theorem 5.38, part (i)

To prove this, we first observe that since H, being the complement of a 0-attractor set, is 1-closed in G_1, the intersection $D' := D \cap H$ is a 0-paradise in H by Proposition 5.33(3). But since $|D'| \leq |D| \leq p_0$, by the induction hypothesis we find $D' \cap A_1 = \varnothing$; but then by $A_1 \subseteq H$ clearly we also have $D \cap A_1 = D \cap H \cap A_1 = D' \cap A_1 = \varnothing$. It follows that D is a 0-paradise in G_2 by Proposition 5.33(2). This proves (53).

But this means that we are done since by the induction hypothesis, D being a 0-paradise in G_2 implies that $D \subseteq G_3$ as required.

Finally, we address statement (ii). Let D be a 1-paradise in G of size at most p_1; we need to show that D has no overlap with G_3. This proof is depicted in Figure 7.

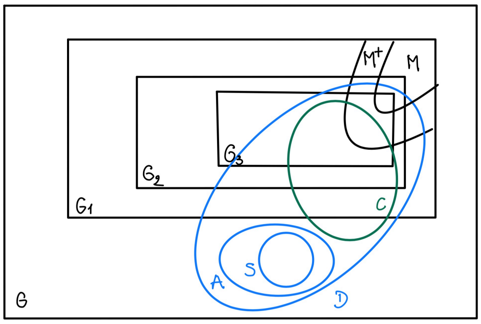

Figure 7: The proof of Theorem 5.38, part (ii)

Let S be the union of all 1-paradises in D of size at most $\lfloor p_1/2 \rfloor$, and let $A := Attr_1(S, D)$ be its 1-attractor set in D. We first show that

$$A \cap G_1 = \varnothing. \tag{54}$$

But since D is 0-closed in G, by Proposition 5.26(5) we find that $A \subseteq Attr_1(S, G)$ and that every 1-paradise in D is also a 1-paradise in G. Hence by the induction hypothesis we get

$S \cap G_1 = \varnothing$. But since G_1 is 0-closed, the set $A \subseteq Attr_1(S, G)$ does not overlap with G_1. This proves (54).

Now if $A = D$ we are done since then we have $D \cap G_3 = A \cap G_3 \subseteq A \cap G_1 = \varnothing$. In the sequel we may therefor assume that $A \subset D$. This implies that $D \setminus A$, being 0-closed in D, is a (non-empty) 1-paradise in $D \setminus A$ by Proposition 5.33(3). Then by Proposition 5.34, the set $D \setminus A$ contains a nonempty 1-paradise C such that $C \cap M = \varnothing$. We now claim that

$$C \cup A \text{ is a 1-paradise in } G\text{, of size greater than } \lfloor p_1/2 \rfloor. \tag{55}$$

We first show that $C \cup A$ is a 1-paradise in D. To witness this, player 1 may use her positional winning strategy for positions in $D \setminus A$, the attractor strategy in D towards A for positions in $A \setminus S$, and the assumed positional winning strategy for any position in the 1-paradise S in D.

But D is 0-closed in G, and so this strategy works for 1 in G as well, for all nodes in $C \cup A$. It remains to show that $C \cup A$ is of size greater than $\lfloor p_1/2 \rfloor$. Suppose for contradiction that $|C \cup A| \leq \lfloor p_1/2 \rfloor$; then by definition of S we would find $C \cup A \subseteq S$, and since $S \subseteq A$ this would contradict the fact that C is a nonempty subset of $D \setminus A$. This finishes the proof of (55).

Our next claim is that

$$(C \cup A) \cap G_2 = \varnothing. \tag{56}$$

Since G_1 is 1-closed in G, by Proposition 5.33(3) the set $(C \cup A) \cap G_1 = C \cap G_1$ is a 1-paradise in G_1, and, containing no vertices in M, also in H. Hence by the induction hypothesis, the set $(C \cup A) \cap G_1$ does not intersect G_2 (here we use that $(C \cup A) \cap G_1 \subseteq D$, so that $|(C \cup A) \cap G_1| \leq |D| \leq p_1$). This proves (56).

But since G_2 is 0-closed in G (as we saw in the proof of item (iii)), by Proposition 5.33(3) the set $D \cap G_2$ is a 1-paradise in G_2. The key observation is now that because $D \cap G_2$ has no overlap with $C \cup A$ and $|C \cup A| > \lfloor p_1/2 \rfloor$, the size of $D \cap G_2$ can be at most $\lfloor p_1/2 \rfloor$. But then a simple application of the induction hypothesis yields $(D \cap G_2) \cap G_3 = \varnothing$, which obviously implies $D \cap G_3 = \varnothing$. QED

With the correctness of **L** established, we turn to its complexity analysis. The theorem below states that **L** terminates in quasi-polynomial time, using quadratic space. In its proof we use the following arithmetical fact.

Proposition 5.39 *For any pair k, ℓ of natural numbers we have*

$$\binom{k}{\ell} \leq \left(\frac{ek}{\ell}\right)^{\ell}.$$

Proof. Starting from the well-known fact that $e^\ell = \sum_{j=0}^{\infty} \frac{\ell^j}{j!}$, we find that $\frac{\ell^\ell}{\ell!} < e^\ell$, and from this we immediately infer that $\frac{1}{\ell!} < \frac{e^\ell}{\ell^\ell}$. This gives $\binom{k}{\ell} \leq \frac{k^\ell}{\ell!} \leq \left(\frac{ek}{\ell}\right)^\ell$. QED

Theorem 5.40 *Let $\mathcal{B}$ be a parity game with n positions and maximal priority d. Then the algorithm **L** computes the winning regions of $\mathcal{B}$ in time*

$$2^{\mathcal{O}\left(\log n \cdot \log\left(1 + \frac{d}{\log n}\right)\right)},$$

and space $\mathcal{O}(n^2)$.

Proof. Let $R(d,\ell)$ be the maximum number of calls to $\texttt{solve}_0$ and $\texttt{solve}_1$ during an execution of $\texttt{solve}_\sigma$, where $\sigma = d \mod 2$ and $\ell := \lfloor \log p_0 \rfloor + \lfloor \log p_1 \rfloor$. We only consider the cases where $p_0, p_1 \geq 1$, so that we have $\ell \geq 0$.

By induction on $d + \ell$ we show that

$$R(d,\ell) \leq 2^{\ell+1} \cdot \binom{d+\ell}{\ell} - 1. \tag{57}$$

We take the basis of this induction to be the cases where $d = 0$; here we may easily verify (57) by calculating $R(d,\ell) = 0 \leq 2^{\ell+1} - 1 = 2^{\ell+1} \cdot \binom{0+\ell}{\ell} - 1$.

In the inductive step of the proof we may then assume that $d \geq 1$, so that a straightforward inspection of the algorithm shows that $R(d,\ell) \leq 1 + 2 \cdot R(d, \ell-1) + R(d-1,\ell)$. This leads to the following calculation:

$$\begin{aligned}
R(d,\ell) &\leq 1 + 2 \cdot R(d,\ell-1) + R(d-1,\ell) \\
&\leq 2 \cdot 2^{(\ell-1)+1} \cdot \left(\binom{d+(\ell-1)}{\ell-1} - 1 \right) + 2^{\ell+1} \cdot \binom{(d-1)+\ell}{\ell} - 1 \qquad \text{(induction)} \\
&= 2^{\ell+1} \cdot \left(\binom{d+(\ell-1)}{\ell-1} + \binom{d+(\ell-1)}{\ell} \right) - 1 \\
&= 2^{\ell+1} \cdot \binom{d+\ell}{\ell} - 1
\end{aligned}$$

as required for proving (57).

We now apply the inequality $\binom{k}{\ell} \leq \left(\frac{ek}{\ell}\right)^\ell$ from Proposition 5.39 for $k = d + \ell$ and obtain

$$R(d,\ell) \leq 2^{\ell+1} \cdot \left(\frac{e \cdot (d+\ell)}{\ell}\right)^\ell \leq 2 \cdot \left(\frac{2e \cdot (d+\ell)}{\ell}\right)^\ell \leq 2 \cdot 2^{\ell \cdot \left(1 + \log e + \log(1 + \frac{d}{\ell})\right)}$$

In a parity game with n vertices, we start with $\ell = 2 \cdot \lfloor \log n \rfloor$, finding

$$R(d, 2 \cdot \lfloor \log n \rfloor) \leq 2 \cdot 2^{2 \cdot \lfloor \log n \rfloor \cdot \left(1 + \log e + \log(1 + \frac{d}{2 \cdot \lfloor \log n \rfloor})\right)}$$

Finally, the cost of each call, apart from the recursive calls, is dominated by the construction of the attractor sets in the lines 6 and 8. Since we obtain by Theorem 5.37 that there are algorithms for computing attractor set that runs in time quadratic in the size of the vertices, this means that the run time of $\mathbf{L}$ is bounded by

$$n^2 \cdot 2 \cdot 2^{2 \cdot \lfloor \log n \rfloor \cdot \left(1 + \log e + \log(1 + \frac{d}{2 \cdot \lfloor \log n \rfloor})\right)}$$

and hence in $2^{\mathcal{O}\left(\log n \cdot \log\left(1 + \frac{d}{\log n}\right)\right)}$ as stated by the theorem.

▶ `say something about space complexity`

QED

Remark 5.41 Generally the index of a parity game is small when compared to its size. If we have $d = \mathcal{O}(\log n)$ the complexity of $\mathbf{L}$ is bounded by $2^{\mathcal{O}(\log n)}$. This means that it runs in *polynomial* time, where the degree of the polynomial that provides an upper bound is 'hidden' in the $\mathcal{O}$-notation. ◁

► `Add some exercises:`
`o` $\mathbf{Z}$ `in EXPTIME`
`o also compute positional winning strategies`

5.6 Game equivalence

In this section we explore some notions of *equivalence* for board games. In this setting we will frequently represent a board as a triple $\mathbb{B} = \langle B, E, \sigma \rangle$, where $\sigma : B \to \{0, 1\}$ is a map assigning a player to each position in B.

5.6.1 Simple shadow games

In this book we often use the technique of 'shadow matching' to supply a player σ with a winning strategy in some initialised game $\mathcal{B}@b$, using a winning strategy for σ in some similar, related initialised game $\mathcal{B}'@b'$. These proofs follow a common pattern that we bring out here, starting with the notion of a *(simple) shadow simulation*. For its definition we recall the definition of the *lifted* version $\widetilde{Z} \subseteq A^\infty \times B^\infty$ of a relation $Z \subseteq A \times B$:

$$\widetilde{Z} := \big\{ \big((a_i)_{0 \leq i < \kappa}, (b_i)_{0 \leq i < \lambda} \big) \mid \kappa = \lambda \text{ and } (a_i, b_i) \in Z, \text{ for all } i \big\}.$$

Definition 5.42 Let $\mathcal{B} = \langle B, E, \sigma, W \rangle$ and $\mathcal{B}' = \langle B', E', \sigma', W' \rangle$ be board games. A relation $Z \subseteq B' \times B$ is called a *(simple) shadow simulation* if it meets the following conditions:

1. Z repects players: if $(b', b) \in Z$ then $\sigma'(b') = \sigma(b)$;
2. Z respects winners: for finite matches we require that if $(b', b) \in Z$ then we have $E'[b] = \varnothing$ iff $E[b] = \varnothing$; for infinite matches we require that $(\pi', \pi) \in \widetilde{Z}$ implies $W'(\pi') = W(\pi)$.

Given a shadow simulation Z, any pair $(\pi', \pi) \in \widetilde{Z}$ is called a *shadow pair*, where we refer to π' as a *shadow match* of π. ◁

Shadow simulations can be used to transfer winning strategies from one game to another. The key idea here is that a player may win a match of a $\mathcal{B}'$-game by maintaining a safe connection with a shadow match of $\mathcal{B}$ that is guided by her winning strategy.

Definition 5.43 Let $\mathcal{B} = \langle B, E, \sigma, W \rangle$ and $\mathcal{B}' = \langle B', E', \sigma', W' \rangle$ be board games, linked by a shadow simulation $Z \subseteq B' \times B$. Furthermore, let f' be a strategy for σ which is winning for each of her winning positions in $\mathcal{B}'$. Then a shadow pair (π', π) is called *safe for* σ (relative to f) if $\mathsf{first}(\pi') \in \mathrm{Win}_\sigma(\mathcal{B}')$ and π' is f'-guided.

We say that Z satisfies the *safe continuation condition* (scc) with respect to f, if for every safe shadow pair (π', π), we have, writing $a := \mathsf{first}(\pi)$ and $a' := \mathsf{first}(\pi')$:
1) if $a \in B_\sigma$ then $(f'(\pi'), b) \in Z$, for some $b \in E[a]$;
2) if $a \in B_{\overline{\sigma}}$ then for all $b \in E[a]$ there is some $b' \in Z$ such that $(b', b) \in Z$. ◁

The key result about shadow simulations is the following; we leave its proof as an exercise for the reader.

Theorem 5.44 (Shadow Play Theorem) *Let $\mathcal{B} = \langle B, E, \sigma, W \rangle$ and $\mathcal{B}' = \langle B', E', \sigma', W' \rangle$ be board games, linked by a shadow simulation $Z \subseteq B' \times B$, and let f' be a strategy for σ which is winning for each of her winning positions in $\mathcal{B}'$. Furthermore, assume that Z satisfies the safe continuation conditions with respect to f. Then for every pair $(b, b') \in Z$ we find that $b' \in \mathrm{Win}_\sigma(\mathcal{B}')$ implies $b \in \mathrm{Win}_\sigma(\mathcal{B})$.*

Note that the safe continuation condition refers to *both* matches of the shadow pair. The following example may clarify why this is needed.

Example 5.45 Let $\mathcal{B}, \mathcal{B}', f'$ and Z be as in Definition 5.43. Call a $\mathcal{B}$-match π *safe* if it is part of a safe shadow pair (π', π). Furthermore, say that Z satisfies the *weak safe continuation condition* (wscc) if every safe $\mathcal{B}$-match π is either full, and won by σ, or else it has a safe extension. The following example reveals that this does not suffice to transfer f' to a winning strategy in $\mathcal{B}$.

We compare two models $\mathbb{S} = (S, R, V)$ and $\mathbb{S}' = (S', R', V')$, where $S = S' := \omega \cup \{\infty\}$, $R' = >$, $R = > \cup \{(\infty, \infty)\}$, and $V = V'$ is arbitrary. These models can be used to show that bisimilarity does not coincide with equivalence in the basic modal language, but they are also of interest here. Clearly we have, with $\varphi := \mu x.\Box x$ that $\mathbb{S}, \infty \Vdash \varphi$ while $\mathbb{S}', \infty \not\Vdash \varphi$. In game-theoretic terms, writing $\mathcal{B} := \mathcal{E}(\varphi, \mathbb{S})$ and $\mathcal{B}' := \mathcal{E}(\varphi, \mathbb{S}')$, we find $(\varphi, \infty) \in \text{Win}_\exists(\mathcal{B})$ but $(\varphi, \infty) \notin \text{Win}_\exists(\mathcal{B}')$. Now consider the relation

$$Z := \Big\{ \big((\psi, s), (\psi, s)\big) \mid s \in S, \psi \in Cl(\varphi) \Big\} \cup \Big\{ \big((\psi, n), (\psi, \infty)\big) \mid n \in \omega, \psi \in Cl(\varphi) \Big\}.$$

The point is that Z satisfies the *weak* safe continuation condition, while obviously, the position (φ, ∞) is linked to itself by Z. This provides a counterexample to the weaker version of Theorem 5.44 where we merely require Z to satisfies the wscc. ◁

▶ `discuss generalisations: may link` σ `to their opponent as well`

5.6.2 Covers

A very tight link between two games arises if one is a *cover* of the other. Intuitively, $\mathcal{B}$ is a cover of $\mathcal{B}'$ if it is some kind of (finitary) unravelling of $\mathcal{B}'$, witnessed by some *cover map* f. In order to define the conditions on these cover maps, we remind the reader of the following notation: where $f : B \to B'$ is some map, and $\pi \in B^\infty$ is some (finite or infinite) B-sequence, say, $\pi = (b_i)_{0 \le i < \kappa}$, we let $f \circ \pi$ denote the B'-sequence $f \circ \pi := (fb_i)_{0 \le i < \kappa}$.

Definition 5.46 Let $\mathbb{B} = \langle B, E, \sigma \rangle$ and $\mathbb{B}' = \langle B', E', \sigma' \rangle$ be two boards. Then we call a function $f : B \to B'$ a *cover map* for $\mathbb{B}$ and $\mathbb{B}'$ if f is surjective and satisfies the following conditions:

1) f restricts to a bijection between the sets $E[b]$ and $E'[fb]$, for every $b \in B$;
2) f respects ownership: $\sigma'(fb) = \sigma(b)$, for every $b \in B$.

For two board games $\mathcal{B} = \langle B, E, \sigma, W \rangle$ and $\mathcal{B}' = \langle B', E', \sigma', W' \rangle$, we call a function $f : B \to B'$ a *cover map* if f is a cover map for the underlying boards, and in addition we have

3) f respects winners: $W(\pi) = W(f \circ \pi)$, for every infinite $\mathcal{B}$-match π.

In the case of initialised games, we call $\mathcal{B}@b$ a cover of $\mathcal{B}'@b'$ if there is a cover map f with

4) $f(b) = b'$.

If f is a cover map from $\mathcal{B}$ to $\mathcal{B}'$, we write $f : \mathcal{B} \twoheadrightarrow \mathcal{B}'$, say that $\mathcal{B}$ *covers* $\mathcal{B}'$ *through* f, and we call $\mathcal{B}$ a *cover* of $\mathcal{B}'$. We use similar terminology and notation for initialised games. ◁

Cover maps resemble bounded morphisms between Kripke models, but the bijection condition on successor sets is much stronger than the back- and forth condition of bounded morphisms.

We gather some basic facts about covers and cover maps. Our first observation is that, if f is a cover map from $\mathbb{B}$ to $\mathbb{B}'$, then for any position $b \in B$ there is a 1-1 correspondence between $\mathbb{B}$-paths starting at b and $\mathbb{B}'$-paths starting at fb.

Proposition 5.47 *Let $\mathbb{B} = \langle B, E, \sigma\rangle$ and $\mathbb{B}' = \langle B', E', \sigma'\rangle$ be two boards, and let $f : B \twoheadrightarrow B'$ be a cover map. Then for any path π' in $\mathbb{B}'$ and every $b \in B$ such that $f(b) = \mathsf{first}(\pi')$, there is a unique path π in $\mathbb{B}$ such that $\pi' = \pi \circ f$.*

At the level of games, Proposition 5.47 means that the cover map induces a bijection between the matches of an initialised game and any of its covers. Formulated in terms of game trees, we even find an *isomorphism*, which is sometimes taken as a reason to *identify* a game with its covers.

Proposition 5.48 *Let $a \in B$ be some position in the board game $\mathcal{B} = \langle B, E, \sigma, W\rangle$. Then*

1. *the tree game $\mathcal{B}^T_a@a$ covers $\mathcal{B}@a$, via the cover map $\pi \mapsto \mathsf{last}(\pi)$;*
2. *if $f : \mathcal{B} \twoheadrightarrow \mathcal{H}$, then $\mathcal{B}^T_a \cong \mathcal{H}^T_a$ for any a in $\mathcal{B}$.*

The following Proposition, the proof of which is left for the reader, states that cover maps preserve winning regions.

Proposition 5.49 *Let $f : \mathcal{B} \twoheadrightarrow \mathcal{B}'$ be a cover map. Then for each player $\sigma \in \{0, 1\}$ we have* $\mathrm{Win}_\sigma(\mathcal{B}') = f[\mathrm{Win}_\sigma(\mathcal{B})]$.

As an important example of a cover map, the following proposition states that every ω-regular game is covered by a parity game.

Proposition 5.50 (Parity Cover Lemma) *Let $\mathcal{B} = \langle B, E, \sigma, W\rangle$ be an ω-regular game, and let M be any deterministic parity automaton recognizing the ω-regular language used to define W. Then $\mathcal{B}$ is covered by a parity game which is based on the set $B \times M$.*

Proof. See Exercise 5.2. QED

The following observation is an immediate corollary of the Cover Lemma and Proposition 5.49.

Corollary 5.51 *Let $\mathcal{B}$ be an ω-regular game. Then $\mathcal{B}$ is determined.*

The concept of cover game allows us to manipulate some structural properties of board games. For instance, the following Proposition shows how covers may be used to confine the domain of a partial priority map to the pre-image of any cycle-critical subset of the board of the covered game.

Proposition 5.52 (Strengthened Parity Cover Lemma) *Let $\mathcal{B} = \langle B, E, \sigma, W\rangle$ be an ω-regular game. Then $\mathcal{B}$ is covered by a ppm-parity game $\mathcal{B}' = \langle B', E', \sigma', \Omega'\rangle$ through a cover map f. Furthermore,*

1. *if $D \subseteq B$ is cycle-critical, then we may assume that $\mathsf{Dom}(\Omega') \subseteq f^{-1}[D]$;*
2. *if $D \subseteq B$ is cycle-critical and $E \cap (D \times D) = \varnothing$, then we may assume injectivity of E' on $f^{-1}[D]$: if $f(u'_i) \in D$ and $(u'_i, v') \in E'$, for $i = 0, 1$, then $u'_0 = u'_1$.*

Proof. By the standard cover lemma we may without loss of generality assume that $\mathcal{B}$ is itself already a (standard) parity game, say, with (total) priority map $\Omega : B \to \omega$. Based on this we will take care of the two parts of the proposition.

For part 1, assume that $D \subseteq B$ is cycle-critical, and define $\mathcal{B}' := \langle V', E', \sigma', \Omega'\rangle$ as follows.

$$
\begin{array}{lcl}
V' & := & V \times \big(\mathsf{Ran}(\Omega) \cup \{-1\}\big) \\
E'[(u,k)] & := & \left\{ \begin{array}{ll} \{(v, \Omega(v)) \mid v \in E[u]\} & \text{if } u \in D \\ \big\{\big(v, \max(k, \Omega(v))\big) \mid v \in E[u]\big\} & \text{otherwise} \end{array} \right. \\
\sigma'(u,k) & := & \sigma(u) \\
\Omega(u,k) & := & \left\{ \begin{array}{ll} k & \text{if } u \in D \\ \text{undefined} & \text{otherwise} \end{array} \right.
\end{array}
$$

Roughly, the intuition underlying this construction is as follows. A position (u, k) represents a path π in $\mathcal{B}$ with $\mathsf{last}(\pi) = u$, $\mathsf{first}(\pi)$ is the only position on π that belongs to D, and k is a counter that records the highest priority (if any) encountered on π (after $\mathsf{first}(\pi)$). The value k is *reset* following a position $u \in D$; that is, if $(v, m) \in E(u, k)$ and $u \in D$ then the value of m is no longer dependent on any value encountered on the path before v, but simply defined as the priority of v.

We claim that the projection map $(u, k) \mapsto u$ is the required cover map from $\mathcal{B}'$ to $\mathcal{B}$. The details of verifying this claims are left as an exercise for the reader.

For part 2, assume that $D \subseteq B$ is cycle-critical and that $E \cap (D \times D) = \varnothing$. Now we define

$$
\begin{array}{lcl}
V' & := & ((V \setminus D) \times V) \cup (D \times \{*\}) \\
E'[(u,x)] & := & \big\{(v, u) \mid v \in E[u] \setminus D\big\} \cap \big\{(v, *) \mid v \in E[u] \cap D\big\} \\
\sigma'(u,x) & := & \sigma(u) \\
\Omega'(u,x) & := & \left\{ \begin{array}{ll} \Omega(u) & \text{if } u \in \mathsf{Dom}(\Omega) \\ \text{undefined} & \text{otherwise} \end{array} \right.
\end{array}
$$

Here the intuition underlying the definition of V' is also that a position (u, x) is basically a version of $u \in V$, with x providing some record of u's 'provenance', so to say. More precisely, if $(v, y) \in E'[(u, x)]$ then we always have $v \in E[u]$; in addition we 'remember' u if $v \notin D$ (in the sense that $y = u$), while we 'forget' u if $v \in D$ (in the sense that $y = *$).

To see that E' is injective on $f^{-1}[D]$, suppose that we have $(v, y) \in E'[(u_0, x_0)] \cap E'[(u_1, x_1)]$ for $u_0, u_1 \in D$. But since $u_0, u_1 \in D$ we can only have $x_0 = x_1 = *$. Furthermore, $u_i \in D$ implies $v \notin D$ by the assumption on E, so that by definition of E' we find $y = u_0 = u_1$. But then we have $(u_0, v_0) = (u_1, v_1)$ indeed.

As in the proof of the previous claim, we leave it for the reader to verify that the projection map $(u, k) \mapsto u$ is the required cover map from $\mathcal{B}'$ to $\mathcal{B}$. QED

5.6.3 Game bisimulations

In the previous subsections we saw examples of rather *tight* links between two games. Here we introduce a notion of game bisimulation which is more general, in at least two ways.

First, game bisimulations do not necessarily link a player in one game to the *same* player in the oppositie game. This is for instance useful if we want to give a game-theoretic proof that the formula $\overline{\varphi}$ of Definition 3.22 indeed corresponds to the negation of the formulas φ.

Second, game bisimulations do not need to link *every* position in one of the two games to some position in the other; rather we select some key positions in both games, and we only require that the bisimulation links key positions in one game to key positions in the other. This is for instance useful in a situation where we compare evaluation games for formulas that are equivalent.

▶ `Details to be supplied.`

Definition 5.53 Given an graph game $\mathcal{B} = (B_0, B_1, E, W)$, call a subset $A \subseteq B_0 \cup B_1$ of positions *prime* if it satisfies the following two conditions:

1. any infinite full match starting at some $b \in A$ passes through another position in A;
2. the winner of an infinite match is completely determined by the sequence of positions in A induced by the match.

By a slight abuse of language we refer to the elements of a fixed, prime set as *prime positions*. $\lhd$

Remark 5.54 To spell out the second condition, in the appendix we formally define what we mean with 'the sequence of positions in A, induced by π'. Where $\pi_{\upharpoonright A}$ denotes this induced sequence, condition 2 states that that two matches π and π' have the same winner if $\pi_{\upharpoonright A} = \pi'_{\upharpoonright A}$. Furthermore, note that, given the first condition, π_A is an infinite sequence if π is infinite. $\lhd$

▶ `Examples to be added.`

The point about the introduction of prime positions is that, just as in the case of acceptance games of modal automata, they allow us to think of the game proceeding in *rounds* that start and finish at a prime position (unless one of the players gets stuck during the round). Formally we define these rounds, that correspond to subgraphs of the board graph, via their unravellings as follows.

Definition 5.55 Let A be some prime set of positions in a graph game $\mathcal{B} = (B_0, B_1, E, W)$, and let $a \in A$ be a prime position. The *neighbourhood game tree* associated with a is defined as the following tree $\mathbb{T}^a := (T^a, E^a, a)$; here T^a is the set of all paths β in (B, E), starting in a, and of which the only prime positions are a and $\mathsf{last}(\beta)$ (in case it is prime). We let T^a inherit the partition of B, by putting $T^a_\sigma := \{\beta \in T^a \mid \mathsf{last}(\beta) \in B_\sigma\}$; $E^a := \{(\beta, \beta b) \mid b \in E[\mathsf{last}(\beta)]\}$ is the one-item extension relation on paths, and the one-state path a is the root of $\mathbb{T}^a$. $\lhd$

It follows from the definitions that all branches of $\mathbb{T}^a$ are finite. Clearly $\mathbb{T}^a$ is a subtree of the game tree of the initialized game $\mathcal{B}@b$, where we only take those finite matches that start at the prime position a and that end either because one the players gets stuck, or because a new prime position is reached.

Definition 5.56 Let A be some prime set of positions in a graph game $\mathcal{B} = (B_0, B_1, E, W)$, and let $a \in A$ be a prime position. A node $\beta \in T^a$ is a *prime leaf* of $\mathbb{T}^a$ if $|\beta| > 1$ and $\mathsf{last}(\beta) \in A$; we denote the set of these leaves as $\mathsf{PrL}(b)$, and write $N(a) := \{\mathsf{last}(\beta) \mid \beta \in \mathsf{PrL}\}$. $\lhd$

Prime leaves of $\mathbb{T}^a$ represent full rounds of $\mathcal{B}$ starting at a, that is, partial matches that start at a and end because a second prime position has been reached. Intuitively then, the set $N(a)$ denotes the set of all positions that one may encounter in such a round of $\mathcal{B}$. Non-prime leafs correspond to full matches in which one of the two players gets stuck before a next prime position is met. These intuitions will be made precise in Proposition 5.60 below.

Once we have established a disection of a graph game using prime positions, we may think of the game being in an 'iterative strategic normal form', in the sense that in each round, each player only makes *one choice*, announcing their complete strategy for that round right at the beginning. This idea will be formalized below in two ways, using, respectively, the notion of a player's *power* and that of a *neighbourhood strategy*. Underlying both approaches is the notion of a *neighbourhood game*; different from the games that we have seen so far this game can end in a *draw*.

Definition 5.57 Let $\mathcal{B} = (B_0, B_1, E, W)$ be a board game, with a prime set $A \subseteq B$, and let $a \in A$ be a prime position. The *neighbourhood game* $\mathcal{B}^a$ induced by a is played on the board $(B^a_0, B^a_1, E^a))$, where $B^a_\sigma := \{\mathsf{last}(\beta) \mid \beta \in T^a \cap B_\sigma\}$, and $E^a := \{(b, c) \in E \mid b = a \text{ or } b \in B \setminus A\}$. Noting that all matches played on this board are finite, we declare a match π of $\mathcal{B}^a$ to be a *draw* if $|\pi| > 1$ and $\mathsf{last}(\pi)$ is a prime position; in all other cases one of the players got stuck, and we declare their opponent to be the winner. $\lhd$

It is easy to see that $\mathbb{T}^a$ is the game tree of $\mathcal{B}^a$. The game-theoretic notion of *power* describes the terminal positions which a player can force in this game by deploying a strategy.

Definition 5.58 Let A be a prime set relative to some board game $\mathcal{B} = (B_0, B_1, E, W)$, and let a be a prime position. By induction on the height of a node $\beta \in T^a$ we define, for each player σ, the *power* of σ at a as a collection $P_\sigma(\beta) \subseteq N(a)$ of subsets of $N(a)$.

- if $\beta \in \mathsf{PrL}$, we define for each player σ:

$$P_\sigma(\beta) := \big\{\{\mathsf{last}(\beta)\}\big\};$$

- if $\beta \notin \mathsf{PrL}$, we put

$$P_\sigma(\beta) := \begin{cases} \bigcup \Big\{ P_\sigma(\beta c) \mid c \in E^a[\mathsf{last}(\beta)] \Big\} & \text{if } \beta \in T^a_\sigma \\ \Big\{ \bigcup_{c \in E^a[\mathsf{last}(\beta)]} Y_c \mid Y_c \in P_\sigma(\beta c), \text{ all } c \Big\} & \text{if } \beta \in T^a_{\bar{\sigma}}. \end{cases}$$

Finally, we define the power of σ at a as the set $P_\sigma(a)$. $\lhd$

The first clause of the definition should be clear: at the end of a round of $\mathcal{B}$, the players have no influence on the outcome any more. For some intuitive explanation of the second clause, we distinguish the two cases. If $\beta \in T^a_\sigma$, then in $\mathcal{B}^a$, when play has arrived at the match β, σ may choose a successor c of $\mathsf{last}(\beta)$, and hence $P_\sigma(\beta)$ consists of those sets that σ can pick at any of the continuationss βc of β. On the other hand, if $\beta \in T^a_{\overline{\sigma}}$, then σ has no influence on the choice of a successor c of $\mathsf{last}(\beta)$, so the best σ can do is the following. If σ picks, for each $c \in E^a[\mathsf{last}(\beta)]$, a set $Y_c \in P_\sigma(\beta c)$, then the union of these sets does belong to her power at β.

Perhaps some special attention is due to the paths $\beta \in T^a$ such that $E^a[\beta] = \varnothing$. In case such a β is a prime leaf, we follow the first clause of Definition 5.58; intuitively, this means that $E[\mathsf{last}(\beta)]$ being empty will not be noticed in this round of $\mathcal{B}$, but only at the beginning of the next one. In case $E^a[\beta] = \varnothing$ for some non-prime leaf, we obtain, by the second clause of the definition, that

$$P_\sigma(\beta) := \begin{cases} \varnothing & \text{if } \beta \in T^a_\sigma \\ \{\varnothing\} & \text{if } \beta \in T^a_{\overline{\sigma}}. \end{cases}$$

This confirms our intuition that at such a position, the player who is to move gets stuck and looses the match.

As we will see now, the powers of the two players can also be formulated in terms of strategies in the neighbourhood game.

Definition 5.59 Let A be a prime set relative to some board game $\mathcal{B} = (B_0, B_1, E, W)$, and let a be a prime position. A *neighbourhood strategy* of a player σ is simply a strategy for σ in the neighbourhood game. Such a neighbourhood strategy is *surviving* for σ if it guarantees that σ will not get stuck (in $\mathcal{B}^a$), and *winning* if it guarantees that $\overline{\sigma}$ gets stuck (in $\mathcal{B}^a$).

If players 0 and 1 play $\mathcal{B}^a$ using, respectively, neighbourhood strategies f_0 and f_1, then the unique resulting match will be denoted as $\mathsf{res}(f_0, f_1)$. Given a neighfourhood strategy f_0 for player 0, we define

$$X_{f_0} := \{\mathsf{last}(\mathsf{res}(f_0, f_1)) \in A \mid f_1 \text{ is a neighbourhood strategy for player 1}\},$$

and similarly we define X_{f_1} for a neighbourhood strategy for player 1. $\lhd$

Obviously, neighbourhood strategies for σ are linked to (fragments of) strategies for σ in $\mathcal{B}$. Since these links are generally obvious, but rather tedious to spell out, we refrain from introducing detailed notation and terminology here.

The next proposition spells out the link between neighbourhood strategies and the power of players. We leave its proof as an exercise to the reader.

Proposition 5.60 *Let A be a prime set relative to some board game $\mathcal{B} = (B_0, B_1, E, W)$, and let a be some prime position. Then the following hold, for either player σ:*

1) $P_\sigma(a) = \varnothing$ iff $\overline{\sigma}$ has a winning neighbourhood strategy in $\mathcal{B}^a$;

2) $\varnothing \in P_\sigma(a)$ iff σ has a winning neighbourhood strategy in $\mathcal{B}^a$;

3) for any nonempty $W \subseteq N(a)$ it holds that

$$W \in P_\sigma(a) \text{ iff } W = X_f \text{ for some surviving neighbourhood strategy } f \text{ for } \sigma \text{ in } \mathcal{B}^a.$$

The partitioning of matches of a board game $\mathcal{B}$ into rounds, moving form one prime position to another, and the analysis of the players' powers in one round of such a match, lay the foundations to the introduction of a structural equivalence relation between board games that we refer to as *prime game bisimulation*.

For its definition we recall how to *lift* a binary relation Z between objects to a relation $\widetilde{Z}$ between sequences of these objects. Where $Z \subseteq S \times S'$, $\pi = (s_i)_{0 \leq i < \kappa} \in S^\infty$ and $\pi' = (s'_i)_{0 \leq i < \lambda} \in S'^\infty$, we put $(\pi, \pi') \in \widetilde{Z}$ iff $\kappa = \lambda$ and $(s_i, s'_i) \in Z$, for all $i < \kappa$.

Definition 5.61 Let $\mathcal{B} = (B_0, B_1, E, W)$ and $\mathcal{B}' = (B'_0, B'_1, E', W')$ be two board games, with prime sets A and A', respectively, and let σ and σ' be (not necessarily distinct) players in $\mathcal{B}$ and $\mathcal{B}'$, respectively.

A (σ, σ')-*game bisimulation* is a binary relation $Z \subseteq A \times A'$ which satisfies the following compatibility condition:

$$\text{if } \pi \in B^\omega \text{ and } \pi' \in B'^\omega \text{ are such that } (\pi{\upharpoonright_A}, \pi{\upharpoonright_{A'}}) \in \widetilde{Z}, \text{ then } W(\pi) = \sigma \text{ iff } W'(\pi') = \sigma', \tag{58}$$

as well as the following back-and-forth condition:

$$(P^{\mathcal{B}}_\sigma(a), P^{\mathcal{B}'}_{\sigma'}(a')) \in \overrightarrow{\wp}\overleftarrow{\wp}(Z) \cap \overleftarrow{\wp}\overrightarrow{\wp}(Z), \tag{59}$$

for every $a \in A$ and $a' \in A'$ with $(a, a') \in Z$. ◁

Unravelling the rather concise statement (59), we find that the back-and-forth condition boils down to the following requirements:

$$\begin{array}{ll}
(\sigma, \text{forth}) & \forall W \in P^{\mathcal{B}}_\sigma(a)\, \exists W' \in P^{\mathcal{B}'}_{\sigma'}(a')\, \forall b' \in W'\, \exists b \in W.\, (b, b') \in Z \\
(\overline{\sigma}, \text{forth}) & \forall W \in P^{\mathcal{B}}_{\overline{\sigma}}(a)\, \exists W' \in P^{\mathcal{B}'}_{\overline{\sigma'}}(a')\, \forall b' \in W'\, \exists b \in W.\, (b, b') \in Z \\
(\sigma', \text{back}) & \forall W' \in P^{\mathcal{B}'}_{\sigma'}(a')\, \exists W \in P^{\mathcal{B}}_\sigma(a)\, \forall b \in W\, \exists b' \in W'.\, (b, b') \in Z \\
(\overline{\sigma'}, \text{back}) & \forall W' \in P^{\mathcal{B}'}_{\overline{\sigma'}}(a')\, \exists W \in P^{\mathcal{B}}_{\overline{\sigma}}(a)\, \forall b \in W\, \exists b' \in W'.\, (b, b') \in Z.
\end{array}$$

Given the symmetries in this definition, it should be clear that there are only two kinds of game bisimulations: the $(0, 0)$-bisimulations coincide with the $(1, 1)$-bisimulations, and the $(0, 1)$-bisimulations with the $(1, 0)$-bisimulations.

Example 5.62 Clearly our definition of the compatibility condition (58) does not follow the idea that bisimulation-like conditions should be 'local' in some sense. But then again, this is hardly possible since at this level of generality the winning conditions themselves are non-local. In many cases where the winning conditions are given by a more local definition we can significantly improve on this.

In particular, in case $\mathcal{B} = (B_0, B_1, E, \Omega)$ and $\mathcal{B}' = (B'_0, B'_1, E', \Omega')$ are *parity* game, we may require the following two conditions:

(matching parities) if $(v, v') \in Z$ then $\Omega(v) \bmod 2 = \sigma$ iff $\Omega'(v') \bmod 2 = \sigma'$;

(contraction) if $(v, v'), (w, w') \in Z$ then $\Omega(v) \leq \Omega(w)$ iff $\Omega'(v') \leq \Omega'(w')$.

We leave it for the reader to verify that any relation satisfying these two conditions also meets the compatibility constraint (58). ◁

The following theorem bears witness to the fact that game bisimulation is indeed a useful notion of structural equivalence between board games.

Theorem 5.63 *Let $\mathcal{B} = (B_0, B_1, E, W)$ and $\mathcal{B}' = (B'_0, B'_1, E', W')$ be two board games, with prime sets A and A', respectively, and let σ and σ' be (not necessarily distinct) players in $\mathcal{B}$ and $\mathcal{B}'$, respectively. Whenever $a \in A$ and $a' \in A'$ are related via some (σ, σ')-game bisimulation Z, we have*

$$a \in \mathrm{Win}_\sigma(\mathcal{B}) \text{ iff } a' \in \mathrm{Win}_{\sigma'}(\mathcal{B}').$$

Proof. (Sketch) Without loss of generality we may assume that $\sigma = \sigma' = 0$. Let $a \in A$ and $a' \in A'$ be such that $(a, a') \in Z$, and assume that $a \in \mathrm{Win}_0(\mathcal{B})$. That is, 0 has a winning strategy f in the initialized game $\mathcal{B}@a$. We will use this strategy to provide her with a winning strategy f' in the game $\mathcal{B}'@a'$.

The key idea here is that 0 will make sure that any f'-guided $\mathcal{B}'$-match is linked to an f-guided $\mathcal{B}$-match π through the game bisimulation Z. This is completely analogous to the shadow matches discussed earlier on, the main difference being that now, 0 will only be able to maintain the right link *round by round*, not necessarily move by move.

Call a finite $\mathcal{B}'$-match π' *safe* if $\mathsf{last}(\pi') \in A'$ and there is an f-guided $\mathcal{B}$-match π such that $\mathsf{last}(\pi) \in A$ and $(\pi_A, \pi'_{A'}) \in \widetilde{Z}$. In this case we say that π *witnesses* the safety of π'. Note that this means in particular that $(\mathsf{last}(\pi), \mathsf{last}(\pi')) \in Z$.
The main claim in the proof is the following.

Claim 1 Let π' be a safe $\mathcal{B}'$-match, as witnessed by the f-guided $\mathcal{B}$-match π. Write $b := \mathsf{last}(\pi)$, $b' := \mathsf{last}(\pi')$. Then player 0 has a neigbourhood strategy g'_0 at b' such that for every neighbourhood strategy g'_1 for 1 at b' either 1 gets stuck, or else the resulting extension $\pi' \diamond \mathsf{res}(g'_0, g'_1)$ of π' is safe, witnessed by an f-guided extension ρ of π.

Proof of Claim Note that the continuation of player 0's winning strategy f at the match π induces a neighbourhood strategy $\widehat{f}$ for 0, given by $\widehat{f}(\rho) := f(\pi \diamond \rho)$. Since f is winning for 0 in $\mathcal{B}@a$, $\widehat{f}$ must be at least surviving for her in $\mathcal{B}^a$.

Hence Proposition 5.60 $\widehat{f}$ induces a set $X_{\widehat{f}} \in P_0(b)$. But then by definition of a game bisimulation there is a set $W' \in P_0^{\mathcal{B}'}(b')$ such that $(W, W') \in \overleftarrow{\wp}(Z)$. Using Proposition 5.60 once more, we find a surviving neighbourhood strategy g'_0 in $\mathcal{B}'^{b'}$ such that $W' = X_{g'}$.

Now let 0 use this strategy against player 1 in a match of the neighbourhood match $\mathcal{B}'^{b'}$. Since g' is surviving, 0 cannot get stuck. In case player 1 gets stuck we are done, so assume otherwise. Where g'_1 is the neighbourhood strategy employed by player 1, we let $\tau' := \mathsf{res}(g'_0, g'_1)$ be the resulting neighborhood match. It should then be clear that $c' := \mathsf{last}(\tau) \in W' \subseteq A'$.

Since $(X_{\widehat{f}}, W') \in \overleftarrow{\wp}(Z)$ we may find a position $c \in X_{\widehat{f}}$ such that $(c, c') \in Z$; but $c \in X_{\widehat{f}}$ implies the existence of some neighbourhood strategy f_1 for 1 in $\mathcal{B}^b$ such that $c = \mathsf{last}(\mathsf{res}(f_0, f_1))$.

Let $\tau := \mathsf{res}(f_0, f_1)$ be the corresponding neighbourhood match, then $c = \mathsf{last}(\tau)$ and so $\rho := \pi \diamond \tau$ is the required f-guided match extending π. ◀

Finally, we leave it for the reader to make explicit how the claim is used to prove the theorem — this part of the proof is similar to many earlier ones. QED

Notes

The application of game-theoretic methods in the area of logic and automata theory goes back to work of Büchi. The positional determinacy of parity games was proved independently by Emerson & Jutla [6] and by Mostowski in an unpublished technical report. Our proof of this result is based on Zielonka [24].

Exercises

Exercise 5.1 (positional determinacy of reachability games) Give a direct proof of the positional determinacy of reachability games, that is: prove Theorem 5.22.

Exercise 5.2 (regular games & finite memory strategies) A strategy α for player σ in an infinite game $\mathcal{B} = \langle B_0, B_1, E, W\rangle$ is a *finite memory strategy* if there exists a finite set M, called the *memory set*, an element $m_I \in M$ and a map $(\alpha_1, \alpha_2) : B \times M \to B \times M$ such that for all pairs of sequences $p_0 \cdots p_k \in B^*$ and $m_0 \cdots m_k \in M^*$: if $m_0 = m_I$, $p_0 \cdots p_k \in \mathrm{PM}_\sigma$ and $m_{i+1} = \alpha_2(p_i, m_i)$ (for all $i < k$), then $\alpha(p_0 \cdots p_k) = \alpha_1(p_k, m_k)$.

Now let $\mathcal{B}$ be a regular game.

(a) Define an parity game which covers $\mathcal{B}$, with positions $B \times M$, where M is the carrier of a deterministic parity automaton $\mathbb{M}$ recognizing L.

(a) Show that each player i has a finite memory strategy which is winning for them in $\mathcal{B}@p$ for every $p \in \mathrm{Win}_i$.

Exercise 5.3 (closed sets) Prove the properties of closed sets as listed in Proposition 5.26.

Exercise 5.4 (strengthened cover lemma) Supply the missing details in the proof of the Strengthed Cover Lemma, Proposition 5.52.

Exercise 5.5 (game bisimulation)

- ▶ `Prove the equivalence of the slow and the fast acceptance game for modal automata using game bisimulations.`

Exercise 5.6 (players' powers) Let A be a prime set relative to some board game $\mathcal{B} = (B_0, B_1, E, W)$, let a be some prime position, and let $\sigma \in \{0, 1\}$ be a player. Then for any subset $U \subseteq N(a)$ we have that either $V \subseteq U$ for some $V \in P_\sigma(a)$ or $U \cap V = \varnothing$ for some $V \in P_{\overline{\sigma}}(a)$. Argue that this can be called a *determinacy* property.

References

[1] P. Aczel. An introduction to inductive definitions. In J. Barwise, editor, *Handbook of Mathematical Logic*, volume 90 of *Studies in Logic and the Foundations of Mathematics*, chapter C.5, pages 739–782. North-Holland Publishing Co., Amsterdam, 1977.

[2] J. van Benthem. *Modal Correspondence Theory.* PhD thesis, Mathematisch Instituut & Instituut voor Grondslagenonderzoek, University of Amsterdam, 1976.

[3] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic.* Number 53 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.

[4] J.R. Büchi. On a decision method in restricted second order arithmetic. In E. Nagel, editor, *Proceedings of the International Congress on Logic, Methodology and the Philosophy of Science*, pages 1–11. Stanford University Press, 1962.

[5] A. Chagrov and M. Zakharyaschev. *Modal Logic*, volume 35 of *Oxford Logic Guides.* Oxford University Press, 1997.

[6] E.A. Emerson and C.S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proceedings of the 32nd Symposium on the Foundations of Computer Science*, pages 368–377. IEEE Computer Society Press, 1991.

[7] D. Janin and I. Walukiewicz. Automata for the modal μ-calculus and related results. In *Proceedings of the Twentieth International Symposium on Mathematical Foundations of Computer Science, MFCS'95*, volume 969 of *LNCS*, pages 552–562. Springer, 1995.

[8] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ-calculus w.r.t. monadic second-order logic. In *Proceedings of the Seventh International Conference on Concurrency Theory, CONCUR '96*, volume 1119 of *LNCS*, pages 263–277, 1996.

[9] B. Knaster. Un théorème sur les fonctions des ensembles. *Annales de la Societé Polonaise de Mathematique*, 6:133–134, 1928.

[10] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[11] D. Kozen and R. Parikh. A decision procedure for the propositional μ-calculus. In *Proceedings of the Workshop on Logics of Programs 1983*, LNCS, pages 313–325, 1983.

[12] R. McNaughton. Testing and generating infinite sequences by a finite automaton. *Information and Control*, 9:521–530, 1966.

[13] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *Ann.P.Appl.Log.* 99:241–259, 1999).

[14] A.W. Mostowski. Regular expressions for infinite trees and a standard form of automata. In A. Skowron, editor, *Computation Theory*, LNCS, pages 157–168. Springer-Verlag, 1984.

[15] D.E. Muller. Infinite sequences and finite machines. In *Proceedings of the 4th IEEE Symposium on Switching Circuit Theory and Logical Design*, pages 3–16, 1963.

[16] D. Niwiński. On fixed point clones. In L. Kott, editor, *Proceedings of the 13th International Colloquium on Automata, Languages and Programming (ICALP 13)*, volume 226 of *LNCS*, pages 464–473, 1986.

[17] D. Park. Concurrency and automata on infinite sequences. In *Proceedings 5th GI Conference*, pages 167–183. Springer, 1981.

[18] A. Pnueli. The temporal logic of programs. In *Proc. 18th Symp. Foundations of Computer Science*, pages 46–57, 1977.

[19] V.R. Pratt. Semantical considerations on Floyd-Hoare logic. In *Proc. 17th IEEE Symposium on Computer Science*, pages 109–121, 1976.

[20] S. Safra. On the complexity of ω-automata. In *Proceedings of the 29th Symposium on the Foundations of Computer Science*, pages 319–327. IEEE Computer Society Press, 1988.

[21] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.

[22] I. Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ-calculus. In *Proceedings of the 10th Annual IEEE Symposium on Logic in Computer Science (LICS'95)*, pages 14–24. IEEE Computer Society Press, 1995.

[23] I. Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ-calculus. *Information and Computation*, 157:142–182, 2000.

[24] W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200:135–183, 1998.