Banach-Mazur Games on Graphs

Erich Grädel

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

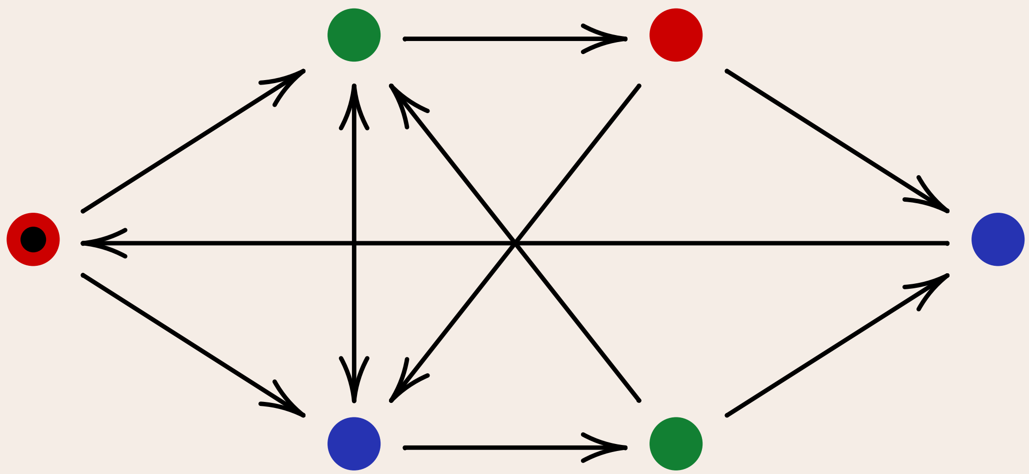

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

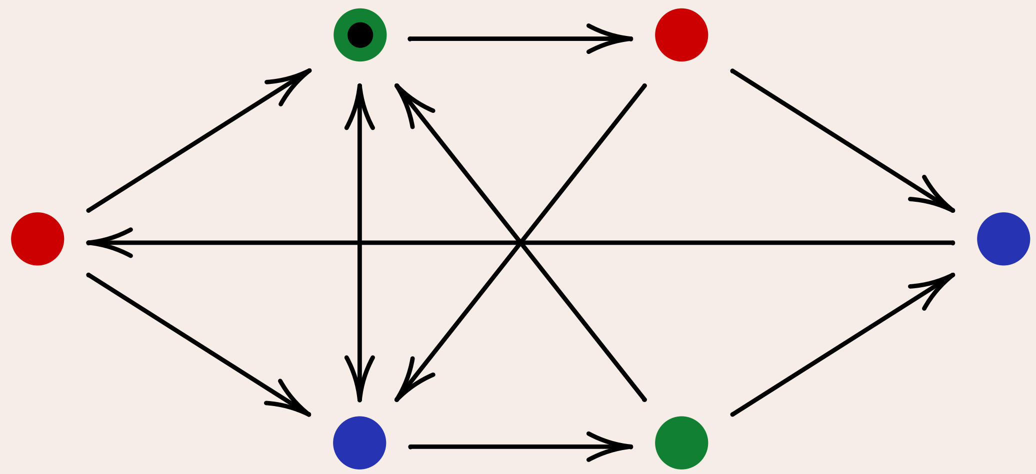

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

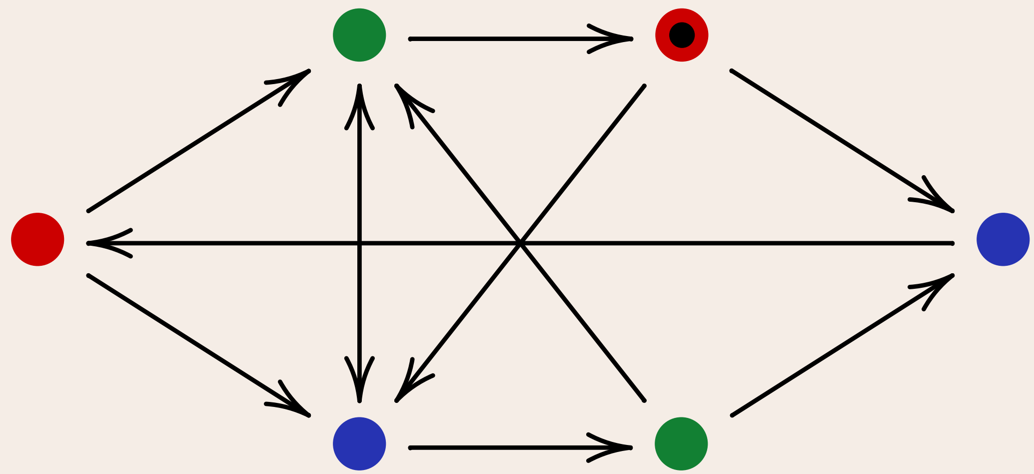

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

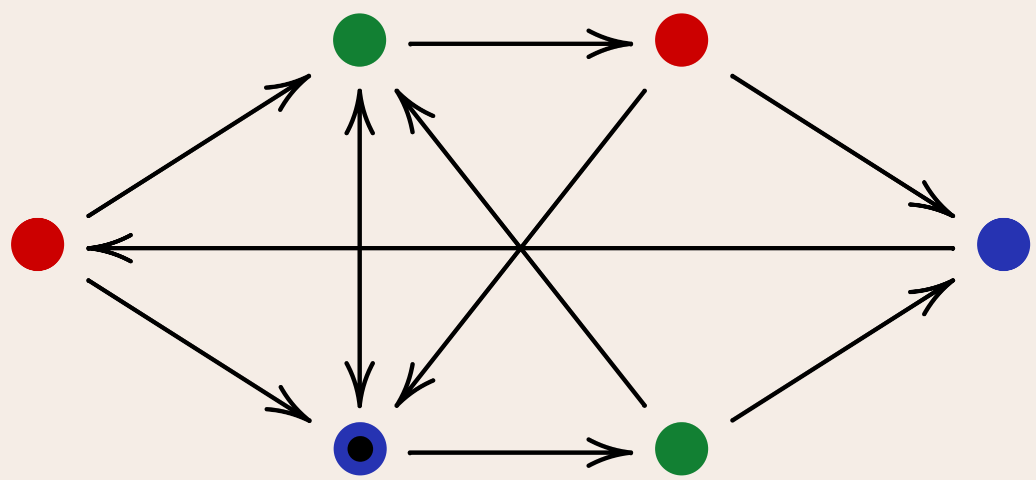

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

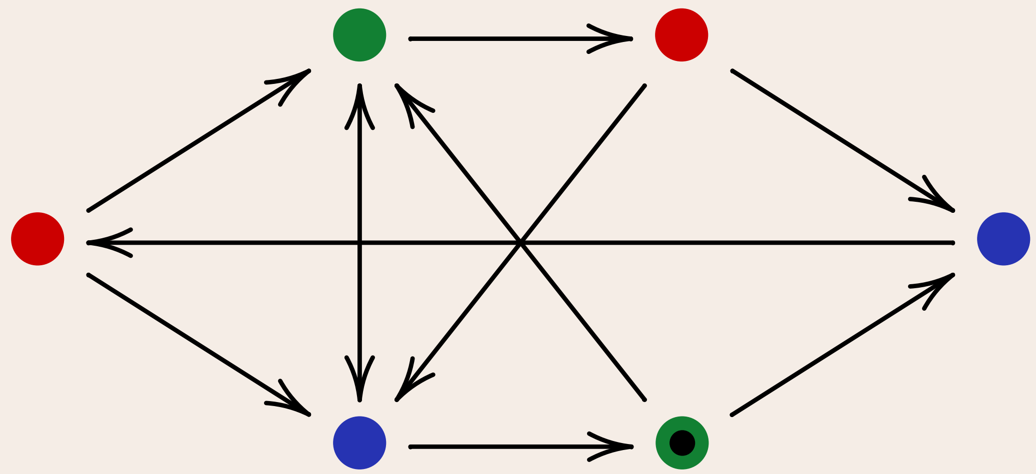

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

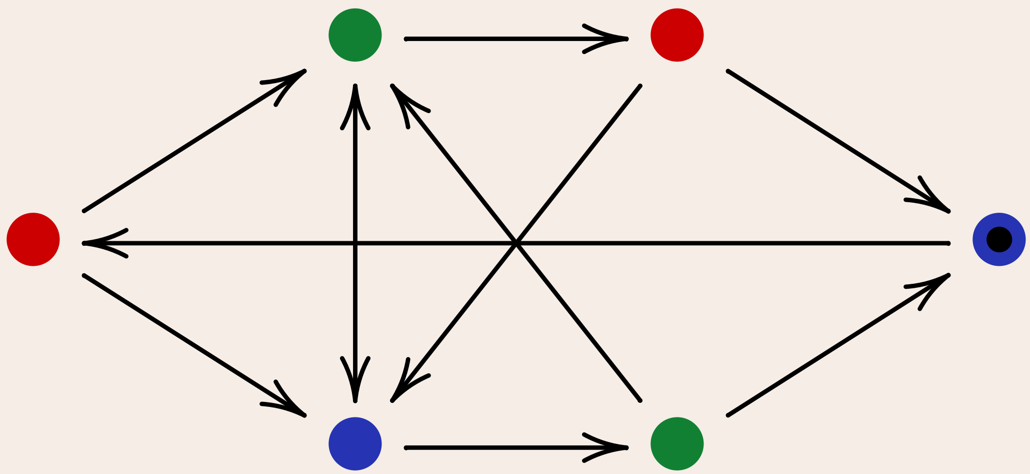

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

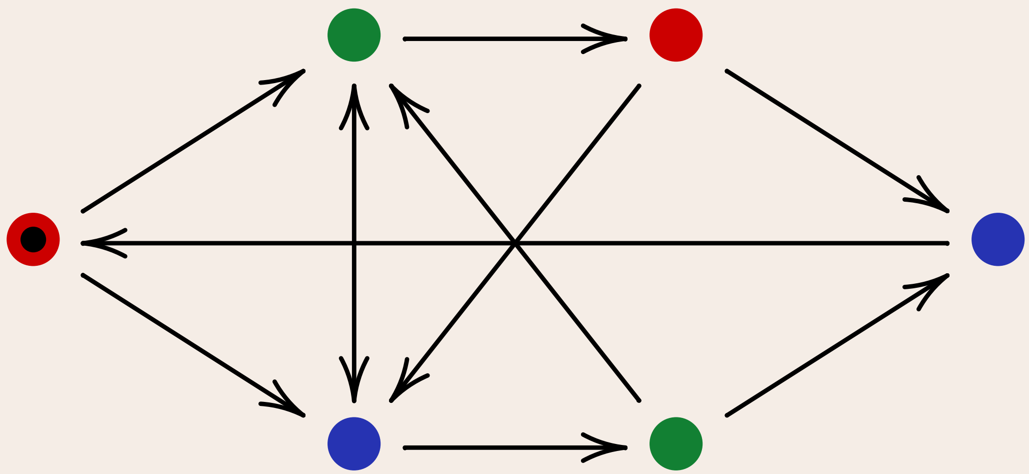

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

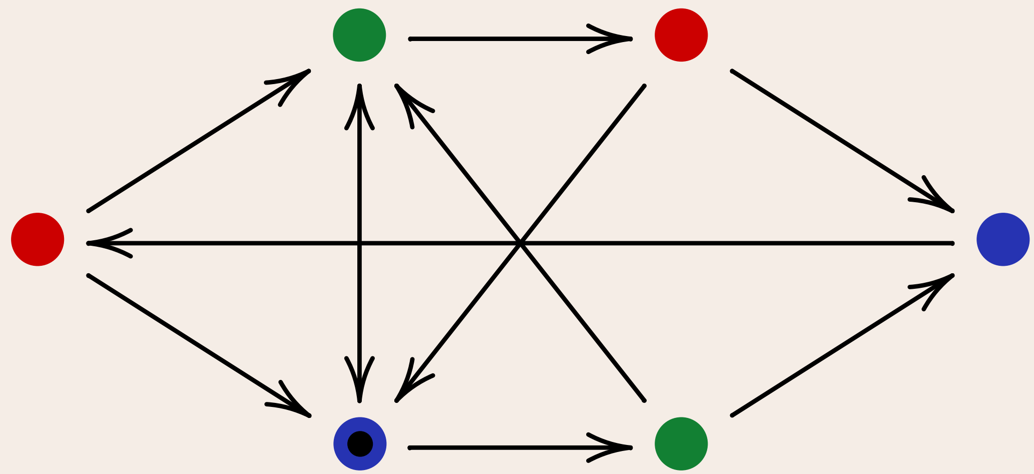

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

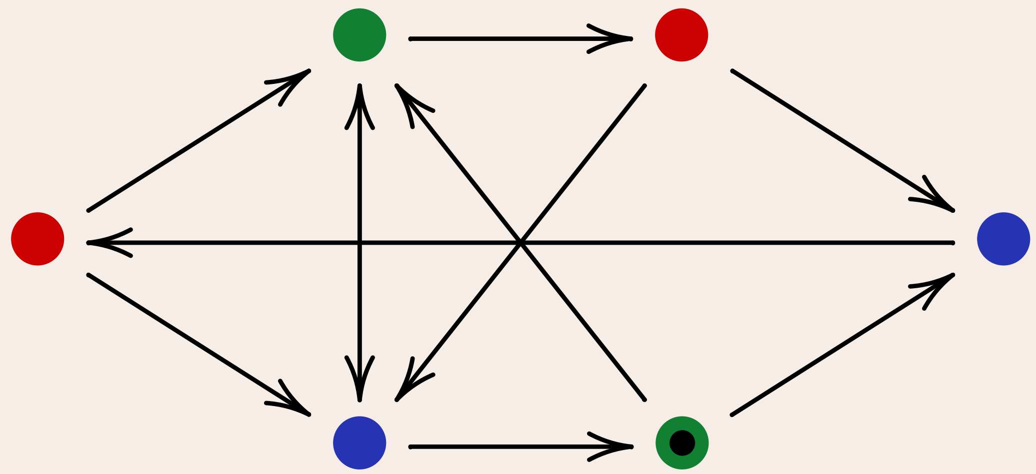

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

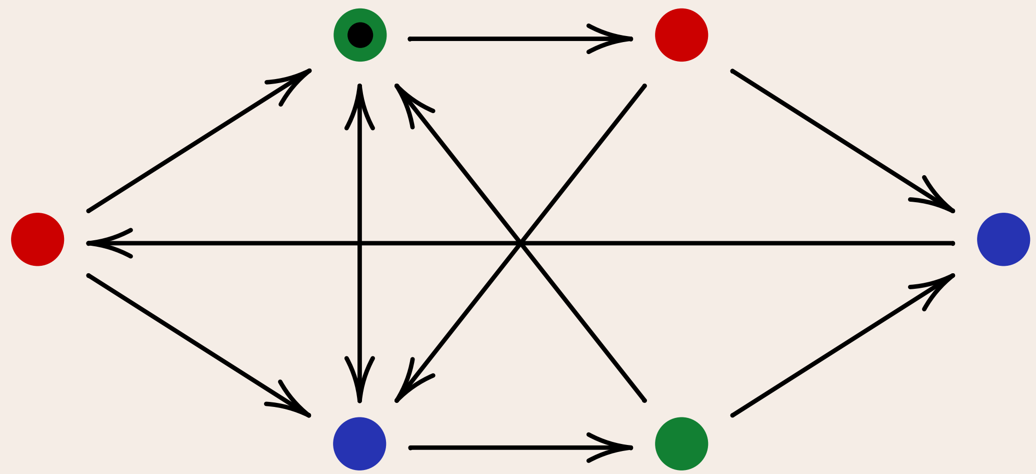

Infinite Games on Graphs

Players take turns (or interact in some other way) to move a token along the edges of a directed graph, tracing out an infinite path.

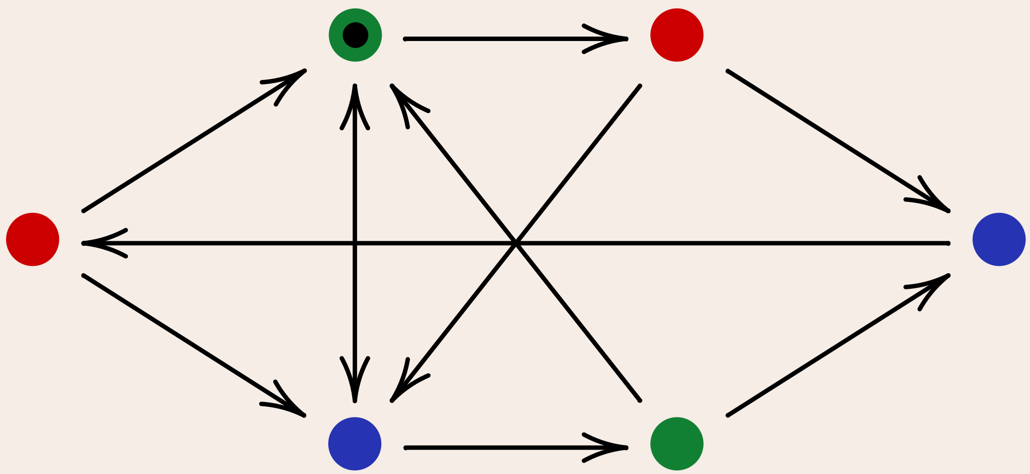

Objectives of the players are given by properties of infinite paths

Mathematical Questions

- What does it mean for a player, to play optimally?
- Compute, for each player, the **optimal value** that she can achieve from a given position.
- Which games are **determined** (from each position, one player has a winning strategy)?
- Are there **optimal strategies**? If so, what is their complexity and how to realise them efficiently?
- How much **knowledge** about the history of a play is necessary to determine an optimal next action?
- ...
- What are the right games to model a particular type of interaction?

The dominant model for infinite graph games

Arena: $\mathcal{G} = (V, V_0, V_1, E, \Omega)$, $V = V_0 \cup V_1$, $E \subseteq V \times V$

$\Omega : V \to C$ assigns to each position a priority or colour.

Move: take token from current position along an edge to a next position

Player 0 moves from positions $v \in V_0$, Player 1 moves from $v \in V_1$.

The dominant model for infinite graph games

Arena: $\mathcal{G} = (V, V_0, V_1, E, \Omega)$, $V = V_0 \cup V_1$, $E \subseteq V \times V$

$\Omega : V \to C$ assigns to each position a priority or colour.

Move: take token from current position along an edge to a next position

Player 0 moves from positions $v \in V_0$, Player 1 moves from $v \in V_1$.

Play: finite or infinite sequence $\pi = v_0 v_1 v_2 \cdots \in V^{\leq \omega}$ with $(v_i, v_{i+1}) \in E$

The dominant model for infinite graph games

Arena: $\mathcal{G} = (V, V_0, V_1, E, \Omega)$, $\quad V = V_0 \cup V_1$, $\quad E \subseteq V \times V$

$\Omega : V \to C$ assigns to each position a priority or colour.

Move: take token from current position along an edge to a next position

Player 0 moves from positions $v \in V_0$, Player 1 moves from $v \in V_1$.

Play: finite or infinite sequence $\pi = v_0 v_1 v_2 \cdots \in V^{\leq\omega}$ with $(v_i, v_{i+1}) \in E$

Winning condition: $\quad \text{Win} \subseteq C^\omega$

– finite plays: who cannot move, loses

– infinite plays: Player 0 wins π, if $\Omega(\pi) \in \text{Win}$, otherwise Player 1 wins.

Classical Theory of Gale-Stewart Games

Consider graph games where the arena is the complete bipartite graph K_{22}

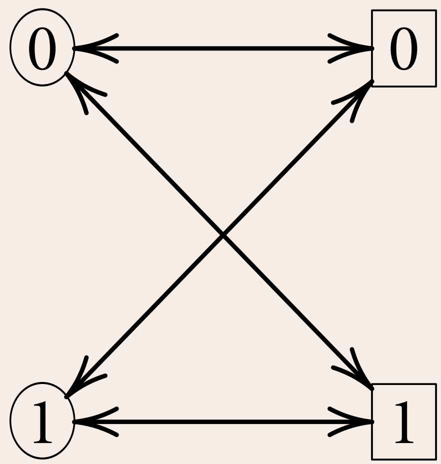

or, equivalently, the infinite binary tree.

Abstract winning condition $\text{Win} \subseteq \{0, 1\}^\omega$

Classical Theory of Gale-Stewart Games

Consider graph games where the arena is the complete bipartite graph K_{22}

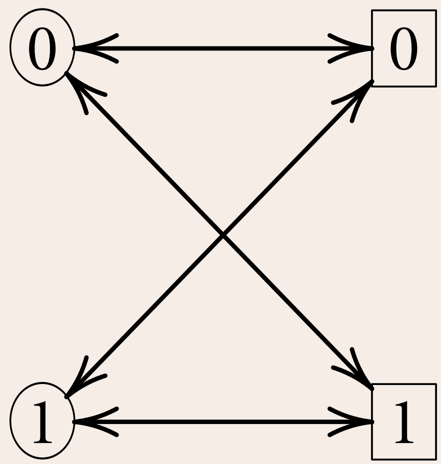

or, equivalently, the infinite binary tree.

Abstract winning condition $\text{Win} \subseteq \{0,1\}^\omega$

This amounts to a game where the players alternatingly select bits $a_i \in \{0,1\}$ in an infinite number of moves and thus produce an infinite word $a_0 a_1 a_2 a_3 \cdots \in \{0,1\}^\omega$.

Classical Theory of Gale-Stewart Games

Consider graph games where the arena is the complete bipartite graph K_{22}

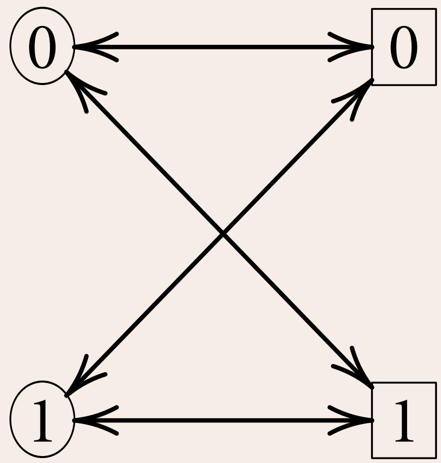

or, equivalently, the infinite binary tree.

Abstract winning condition $\mathrm{Win} \subseteq \{0,1\}^\omega$

This amounts to a game where the players alternatingly select bits $a_i \in \{0,1\}$ in an infinite number of moves and thus produce an infinite word $a_0 a_1 a_2 a_3 \cdots \in \{0,1\}^\omega$.

Similarly for the ω-branching games on the graph $K_{\omega\omega}$ and winning conditions $\mathrm{Win} \subseteq \omega^\omega$

Banach-Mazur games

There is another, in fact older, variant of such infinite games where, in each move, the players select not just a bit, but an arbitrary finite word $w_i \in \{0, 1\}^*$, again producing an infinite word $w_0 w_1 w_2 w_3 \cdots \in \{0, 1\}^\omega$.

Banach-Mazur games

There is another, in fact older, variant of such infinite games where, in each move, the players select not just a bit, but an arbitrary finite word $w_i \in \{0,1\}^*$, again producing an infinite word $w_0 w_1 w_2 w_3 \cdots \in \{0,1\}^\omega$.

These games play an important role in topology.

Original version of **Banach-Mazur games** $G^{**}(W)$ (see *"The Scottish Book. Mathematics from the Scottish Café"*, Birkhäuser 1981, Problem 43):

For a given winning condition $W \subseteq \mathbb{R}$, Player 0 first selects an interval $d_1 \subset \mathbb{R}$, then Player 1 chooses a subinterval $d_2 \subset d_1$, then Player 0 selects a further refinement $d_3 \subset d_2$, and so on . . .

Player 0 wins if, and only if, $\bigcap_{n\in\omega} d_n$ contains an element of W.

Banach-Mazur games on graphs

Arena: a directed graph $G = (V, E)$, with initial position v_0

Winning condition: a set W of infinite paths from v_0 through G

Players: Ego (E) and Alter (A)

Banach-Mazur games on graphs

Arena: a directed graph $G = (V, E)$, with initial position v_0

Winning condition: a set W of infinite paths from v_0 through G

Players: Ego (E) and Alter (A)

Playing the game: The first player selects a finite path p_1 from v_0; the opponent extends p_1 to a path $p_1 q_1$; then the first player prolongs this to $p_1 q_1 p_2$; and so on. All moves are non-empty and finite: $1 \leq |p_i|, |q_i| < \omega$.

Banach-Mazur games on graphs

Arena: a directed graph $G = (V, E)$, with initial position v_0

Winning condition: a set W of infinite paths from v_0 through G

Players: Ego (E) and Alter (A)

Playing the game: The first player selects a finite path p_1 from v_0; the opponent extends p_1 to a path $p_1 q_1$; then the first player prolongs this to $p_1 q_1 p_2$; and so on. All moves are non-empty and finite: $1 \leq |p_i|, |q_i| < \omega$.

Winning the game: After ω moves, an infinite path π is completed.
Ego has won if $\pi \in W$, otherwise Alter has won.

Banach-Mazur games on graphs

Arena: a directed graph $G = (V, E)$, with initial position v_0

Winning condition: a set W of infinite paths from v_0 through G

Players: Ego (E) and Alter (A)

Playing the game: The first player selects a finite path p_1 from v_0; the opponent extends p_1 to a path $p_1 q_1$; then the first player prolongs this to $p_1 q_1 p_2$; and so on. All moves are non-empty and finite: $1 \leq |p_i|, |q_i| < \omega$.

Winning the game: After ω moves, an infinite path π is completed. Ego has won if $\pi \in W$, otherwise Alter has won.

Central issue in descriptive set theory: Characterise determined games by topological properties of the winning conditions.

Topology

On B^ω, define a **topology** with **basic open sets** $O(x) := x \cdot B^\omega$, for $x \in B^*$.

- L is **open** $\iff$ $L = W \cdot B^\omega$ (for some $W \subseteq B^*$)
- L is **closed** $\iff$ $B^\omega - L$ is open $\iff$ $L = [T]$
  where $[T]$ is the set of infinite branches of a tree $T \subseteq B^*$.
- L is **nowhere dense** if the closure of L contains no non-empty open set.
- L is **meager** or **topologically small** if it is a countable union of nowhere dense sets.
- L is **co-meager** or **topologically large** if its complement is meager.

Topology

On B^ω, define a **topology** with **basic open sets** $O(x) := x \cdot B^\omega$, for $x \in B^*$.

- L is **open** $\iff$ $L = W \cdot B^\omega$ (for some $W \subseteq B^*$)
- L is **closed** $\iff$ $B^\omega - L$ is open $\iff$ $L = [T]$
  where $[T]$ is the set of infinite branches of a tree $T \subseteq B^*$.
- L is **nowhere dense** if the closure of L contains no non-empty open set.
- L is **meager** or **topologically small** if it is a countable union of nowhere dense sets.
- L is **co-meager** or **topologically large** if its complement is meager.

In precisely the same way, one defines these topological notions on spaces $\text{Paths}(G, v)$, the set of infinite paths through G from v.

Determinacy of Banach-Mazur games

Theorem (Banach-Mazur) In a Banach-Mazur game $\text{BM}(G, v_0, W)$

(1) Alter has a winning strategy $\iff$ W is meager in $\text{Paths}(G, v_0)$.

(2) Ego has a winning strategy $\iff$ there exists a finite path x from v_0 such that $O(x) \setminus W$ is meager (i.e. W co-meager in some basic open set).

Determinacy of Banach-Mazur games

Theorem (Banach-Mazur) In a Banach-Mazur game $\mathrm{BM}(G, v_0, W)$

(1) Alter has a winning strategy $\iff$ W is meager in $\mathrm{Paths}(G, v_0)$.

(2) Ego has a winning strategy $\iff$ there exists a finite path x from v_0 such that $O(x) \setminus W$ is meager (i.e. W co-meager in some basic open set).

Corollary Banach-Mazur games with Borel winning conditions are determined.

(follows from the fact that Borel sets have the Baire property, i.e., their symmetric difference with some open set is meager)

Determinacy of Banach-Mazur games

Theorem (Banach-Mazur) In a Banach-Mazur game $\mathrm{BM}(G, v_0, W)$

(1) Alter has a winning strategy $\iff$ W is meager in Paths(G, v_0).

(2) Ego has a winning strategy $\iff$ there exists a finite path x from v_0 such that $O(x) \setminus W$ is meager (i.e. W co-meager in some basic open set).

Corollary Banach-Mazur games with Borel winning conditions are determined.

(follows from the fact that Borel sets have the Baire property, i.e., their symmetric difference with some open set is meager)

Remark. Standard winning conditions used in applications (e.g. all ω-regular winning conditions) are in low levels of the Borel hierarchy.

Determinacy of Banach-Mazur games

It suffices to prove the claim for Alter.

Let $W = \bigcup_{n\in\omega} X_n$ with X_n nowhere dense. In his n-th move, Player A moves to some x_n such that $O(x_n) \cap X_n = \varnothing$. The play produced in this way is contained in all $O(x_n)$, hence not in W.

For the converse, let f be any strategy of Alter in the game $\mathrm{BM}(G, v, W)$. It suffices to verify that

$$\mathrm{Plays}(f) := \{\pi \in \mathrm{Paths}(G, v) : \pi \text{ is consistent with } f\} \text{ is co-meager.}$$

But for each n, the set

$Y_n(f) := \{\pi \in \mathrm{Paths}(G, v) : \pi \text{ is consistent with } f \text{ for the first } n \text{ moves}\}$

is open and dense (and its complement is nowhere dense). Hence $\mathrm{Plays}(f) = \bigcap_{n\in\omega} Y_n(f)$ is co-meager.

Determinacy of Banach-Mazur games

It suffices to prove the claim for Alter.

Let $W = \bigcup_{n\in\omega} X_n$ with X_n nowhere dense. In his n-th move, Player A moves to some x_n such that $O(x_n) \cap X_n = \varnothing$. The play produced in this way is contained in all $O(x_n)$, hence not in W.

For the converse, let f be any strategy of Alter in the game $\mathrm{BM}(G, v, W)$. It suffices to verify that

$$\mathrm{Plays}(f) := \{\pi \in \mathrm{Paths}(G, v) : \pi \text{ is consistent with } f\} \text{ is co-meager.}$$

But for each n, the set

$Y_n(f) := \{\pi \in \mathrm{Paths}(G, v) : \pi \text{ is consistent with } f \text{ for the first } n \text{ moves}\}$

is open and dense (and its complement is nowhere dense). Hence $\mathrm{Plays}(f) = \bigcap_{n\in\omega} Y_n(f)$ is co-meager.

Consequently, if f is a winning strategy then $\mathrm{Plays}(f) \cap W = \varnothing$ hence W must be meager.

An application: Characterisation of fairness

In verification one sometimes want to exclude **unfair** runs from consideration and verify the specification only for the fair ones.

Fairness: Anything that is infinitely often possible, should indeed happen infinitely often.

Many different ways to make this intuitive notion precise.

An application: Characterisation of fairness

In verification one sometimes want to exclude **unfair** runs from consideration and verify the specification only for the fair ones.

Fairness: Anything that is infinitely often possible, should indeed happen infinitely often.

Many different ways to make this intuitive notion precise.

Definition. (Varacca, Völzer) A set F of infinite runs through a transition system (G, v) is a fairness property if F is topologically large (i.e. co-meager) in the set $\text{Paths}(G, v)$ of all possible runs.

An application: Characterisation of fairness

In verification one sometimes want to exclude **unfair** runs from consideration and verify the specification only for the fair ones.

Fairness: Anything that is infinitely often possible, should indeed happen infinitely often.

Many different ways to make this intuitive notion precise.

Definition. (Varacca, Völzer) A set F of infinite runs through a transition system (G, v) is a fairness property if F is topologically large (i.e. co-meager) in the set $\text{Paths}(G, v)$ of all possible runs.

By the Theorem of Banach-Mazur this means that the first player has a winning strategy in the Banach-Mazur game $\text{BM}(G, v, F)$.

Game-theoretic view of fairness: The scheduler has a strategy to ensure a fair run.

Strategies and positional strategies

A **decomposition invariant** strategy is of the form $f : V^* \to V^*$. It only depends on the finite path constructed so far, not how it has been built up.

Example. "Do the same what your opponent did in his last move" is **not** a decomposition invariant strategy.

Proposition. A player who has a winning strategy for a Banach-Mazur game also has one that is decomposition invariant.

Strategies and positional strategies

A **decomposition invariant** strategy is of the form $f : V^* \to V^*$. It only depends on the finite path constructed so far, not how it has been built up.

Example. "Do the same what your opponent did in his last move" is **not** a decomposition invariant strategy.

Proposition. A player who has a winning strategy for a Banach-Mazur game also has one that is decomposition invariant.

A **positional** strategy is of the form $f : V \to V^*$. It only depends on the current position, not on the history of the play.

A game is **positionally determined** if one of the players has a positional winning strategy. A winning condition **guarantees positional determinacy** if all games with that winning condition are positionally determined.

Positional winning strategy for one player

Proposition. If $W \in \boldsymbol{\Sigma}_2^0$ (countable union of closed sets), and Ego has winning strategy for the game $\mathrm{BM}(G, v, W)$, then he also has a positional winning strategy.

Positional winning strategy for one player

Proposition. If $W \in \mathbf{\Sigma}^0_2$ (countable union of closed sets), and Ego has winning strategy for the game $\mathrm{BM}(G, v, W)$, then he also has a positional winning strategy.

This is not always true for $W \in \mathbf{\Pi}^0_2$:

G_2 :

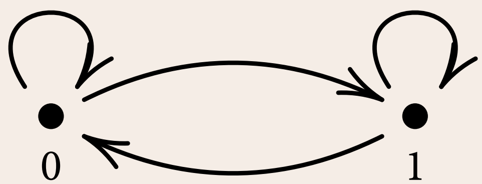

$$W = \{\pi \in \{0,1\}^\omega : (\forall m)(\exists n > m)\, |\{i < n : \pi(i) = 0\}| \geq n/2\}$$

(infinitely many initial segments of π have more zeros than ones)

Ego has a winning strategy for $\mathrm{BM}(G_2, 0, W)$, but no positional one.

ω-regular winning conditions

Game graph $\mathcal{G} = (V, E)$ with colouring $\lambda : V \to \{0, \dots, d-1\}$.

Logical winning conditions: given by formula φ in some logic on infinite paths, with predicates $\lambda(v) = i$ ($i < d$), such as

- S1S: monadic second-order logic on infinite paths;
- LTL, or equivalently, first-order logic FO($<$).

ω-regular winning conditions

Game graph $\mathcal{G} = (V, E)$ with colouring $\lambda : V \to \{0, \dots, d-1\}$.

Logical winning conditions: given by formula φ in some logic on infinite paths, with predicates $\lambda(v) = i$ ($i < d$), such as

- S1S: monadic second-order logic on infinite paths;
- LTL, or equivalently, first-order logic FO(<).

Muller conditions: given by collection $\mathcal{F} \subseteq \mathcal{P}(\{0, \dots, d-1\})$.

An infinite play $\pi = v_0 v_1 v_2 \dots$ is won by Ego if

$$\mathrm{Inf}(\pi) := \{c : (\forall i)(\exists j > i)\lambda(v_j) = c\} \in \mathcal{F}.$$

ω-regular winning conditions

Game graph $\mathcal{G} = (V, E)$ with colouring $\lambda : V \to \{0, \ldots, d-1\}$.

Logical winning conditions: given by formula φ in some logic on infinite paths, with predicates $\lambda(v) = i$ $(i < d)$, such as

- S1S: monadic second-order logic on infinite paths;
- LTL, or equivalently, first-order logic FO(<).

Muller conditions: given by collection $\mathcal{F} \subseteq \mathcal{P}(\{0, \ldots, d-1\})$.

An infinite play $\pi = v_0 v_1 v_2 \ldots$ is won by Ego if

$$\mathrm{Inf}(\pi) := \{c : (\forall i)(\exists j > i)\lambda(v_j) = c\} \in \mathcal{F}.$$

Parity condition: Ego wins if the least colour seen infinitely often is even.

Positional determinacy of classical graph games

For the more common, single step games on graphs

- **parity games** are positionally determined
- Positional strategies do **not** suffice for **Muller games**.

Example:

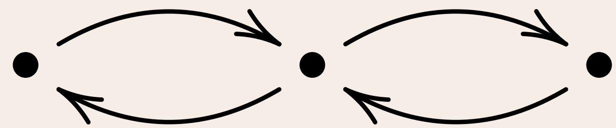

winning condition: all positions must occur infinitely often

Positional Determinacy of Muller Games

Proposition. Muller conditions guarantee positional determinacy for Banach-Mazur games.

Proof. Decompose G into its strongly connected components

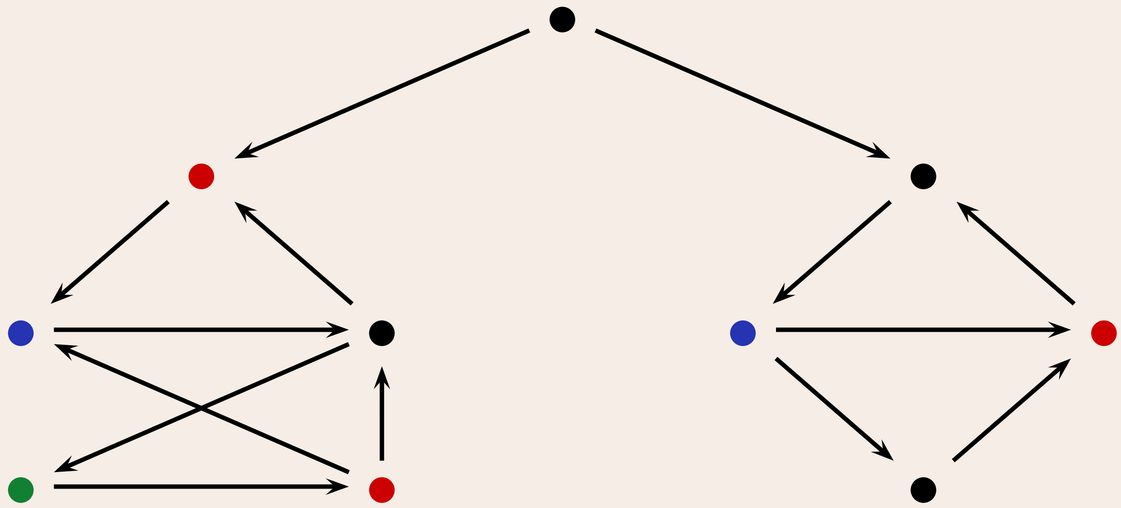

Positional Determinacy of Muller Games

Proposition. Muller conditions guarantee positional determinacy for Banach-Mazur games.

Proof. Decompose G into its strongly connected components

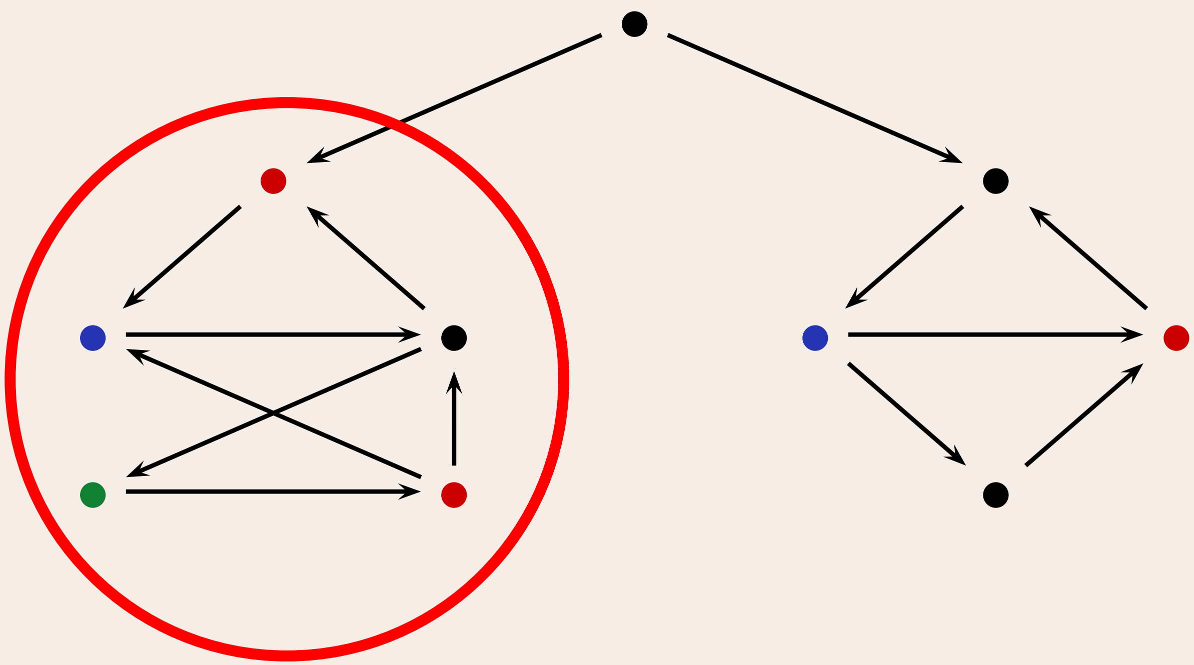

Ego wins iff there is a leaf component C such that $\lambda(C) \in \mathcal{F}$.
But then he also wins with a positional strategy.

Positional Determinacy of Muller Games

Proposition. Muller conditions guarantee positional determinacy for Banach-Mazur games.

Proof. Decompose G into its strongly connected components

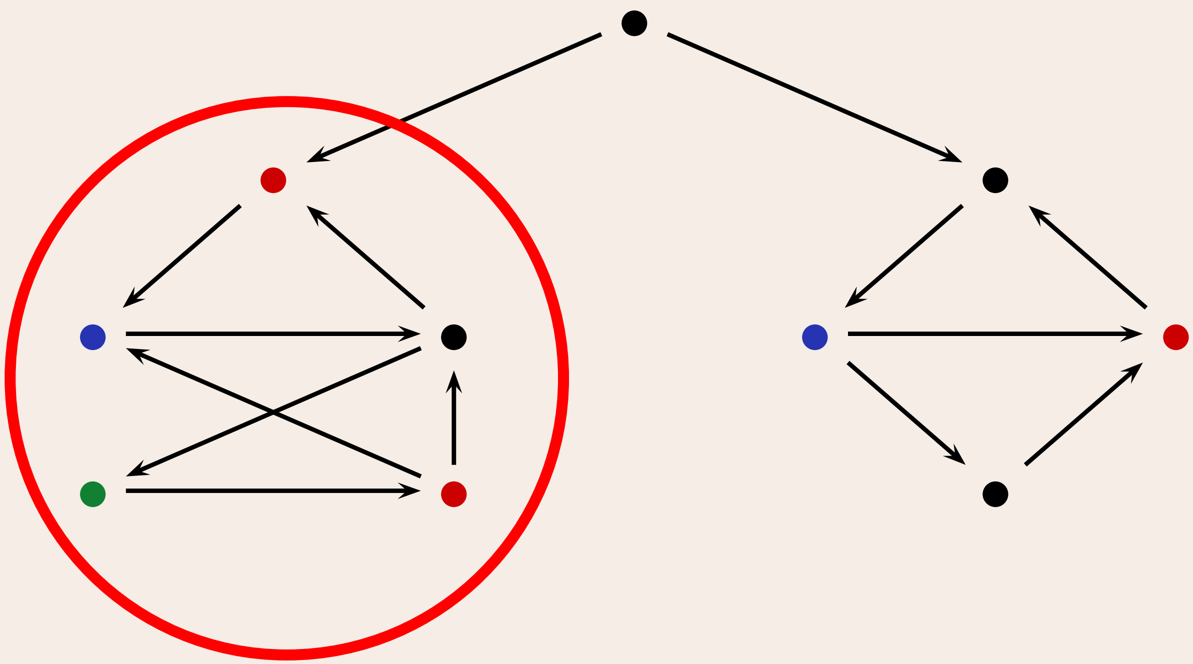

Ego wins iff there is a leaf component C such that $\lambda(C) \in \mathcal{F}$.
But then he also wins with a positional strategy.

Corollary. Muller games can be solved in time $O(|G| \cdot |\mathcal{F}|)$.

Games with ω-regular winning conditions

Theorem. (Büchi-Landweber) Classical graph games with ω-regular winning conditions are determined via finite-memory strategies that can be effectively computed and realised by finite automata.

Games with ω-regular winning conditions

Theorem. (Büchi-Landweber) Classical graph games with ω-regular winning conditions are determined via finite-memory strategies that can be effectively computed and realised by finite automata.

Reduction to parity games. Every game G with ω-regular winning condition W can be transformed into an equivalent game with a game graph $G \times M$ and a parity winning condition.

Simplify the winning condition at the cost of enlarging the game graph

Games with ω-regular winning conditions

Theorem. (Büchi-Landweber) Classical graph games with ω-regular winning conditions are determined via finite-memory strategies that can be effectively computed and realised by finite automata.

Reduction to parity games. Every game G with ω-regular winning condition W can be transformed into an equivalent game with a game graph $G \times M$ and a parity winning condition.

Simplify the winning condition at the cost of enlarging the game graph

These results can be directly carried over to Banach-Mazur games.

Games with ω-regular winning conditions

Theorem. (Büchi-Landweber) Classical graph games with ω-regular winning conditions are determined via finite-memory strategies that can be effectively computed and realised by finite automata.

Reduction to parity games. Every game G with ω-regular winning condition W can be transformed into an equivalent game with a game graph $G \times M$ and a parity winning condition.

Simplify the winning condition at the cost of enlarging the game graph

These results can be directly carried over to Banach-Mazur games.

But we can do better!

Eliminating the finite memory

Theorem. Banach-Mazur games that are determined via finite-memory strategies are in fact positionally determined.

Intuition. A finite-memory winning strategy for G leads to a positional winning strategy in an expanded game $G \times M$. Combine many segments of plays consistent with this strategy to a single move of the opponent, and combine answers to a fixed answer $f(v, m)$ that does not depend on m. Prove that each play consistent with this strategy is also consistent with the original strategy. Since this strategy does not depend on the memory component, it induces a positional winning strategy in the original game.

Corollary. For Banach-Mazur games, ω-regular winning conditions guarantee positional determinacy.

Positional determinacy and fairness

From the positional determinacy of the Banach-Mazur games with ω-regular winning conditions, it follows that, on finite graphs, Ego has a winning strategy if, and only if, W is probabilistically large (under positive Markov measures).

Corollary (Varacca, Völzer) Any ω-regular fairness property has probability 1 under randomised scheduling.

As a further consequence, one can use results about finite Markov chains for checking whether a finite system is fairly correct with respect to LTL or ω-regular specification.

Banach-Mazur games with few alternations

We are also interested in Banach-Mazur games where, after a fixed number of moves, only one player remains.

A **game prefix** $\gamma \in \{E, A\}^\omega$ indicates who begins and how many alternations are played. Obviously, $EE \equiv E$ and $AA \equiv A$. Hence, for any arena (G, W) we have the following games:

- $(EA)^\omega(G, W)$ and $(AE)^\omega(G, W)$:
  usual Banach-Mazur games, with infinite alternations
- $(EA)^k E^\omega(G, W)$ and $A(EA)^k E^\omega(G, W)$:
  games ending with infinite path extension by Ego
- $(AE)^k A^\omega(G, W)$ and $E(AE)^k A^\omega(G, W)$:
  games where Alter chooses the final infinite lonesome ride

Planning in nondeterministic domains

Planning domain: transition system $G := (V, (E_a)_{a \in A}, (P_b)_{b \in B})$

Planning goal: property of execution paths, specified by $\varphi \in \mathrm{LTL}$

Plan: $\pi : V^* \longrightarrow A$, assigns to each history an action

Execution tree: if the planning domain G is deterministic, then π defines a unique execution path. However, if G is nondeterministic, an action may have several outcomes, and a plan π then has not only one execution path, but an execution tree $\mathcal{T}_{G,\pi}$.

Planning in nondeterministic domains

Planning domain: transition system $G := (V, (E_a)_{a \in A}, (P_b)_{b \in B})$

Planning goal: property of execution paths, specified by $\varphi \in \mathrm{LTL}$

Plan: $\pi : V^* \longrightarrow A$, assigns to each history an action

Execution tree: if the planning domain G is deterministic, then π defines a unique execution path. However, if G is nondeterministic, an action may have several outcomes, and a plan π then has not only one execution path, but an execution tree $\mathcal{T}_{G,\pi}$.

It may be unrealistic to expect that all execution paths of a plan π satisfy the goal φ. On the other side, it is too optimistic to assume that a plan is good if just one execution path is consistent with φ.

Banach-Mazur games for planning

Pistore/Vardi study nondeterministic planning by means of games: a plan π is good for the goal φ on domain G if Ego wins an associated game on the execution tree $\mathcal{T}_{G,\pi}$.

- **Weak planning**
  There is a path in $\mathcal{T}_{G,\pi}$ that satisfies φ.
- **Strong planning**
  Every path in $\mathcal{T}_{G,\pi}$ satisfies φ.
- **Strong cyclic planning**
  Every partial execution of π can be extended to a successful path.

Banach-Mazur games for planning

Pistore/Vardi study nondeterministic planning by means of games: a plan π is good for the goal φ on domain G if Ego wins an associated game on the execution tree $\mathcal{T}_{G,\pi}$.

- **Weak planning** $E^\omega(\mathcal{T}_{G,\pi}, \varphi)$

  There is a path in $\mathcal{T}_{G,\pi}$ that satisfies φ.

- **Strong planning** $A^\omega(\mathcal{T}_{G,\pi}, \varphi)$

  Every path in $\mathcal{T}_{G,\pi}$ satisfies φ.

- **Strong cyclic planning** $AE^\omega(\mathcal{T}_{G,\pi}, \varphi)$

  Every partial execution of π can be extended to a successful path.

Banach-Mazur games for planning

Pistore/Vardi study nondeterministic planning by means of games: a plan π is good for the goal φ on domain G if Ego wins an associated game on the execution tree $\mathcal{T}_{G,\pi}$.

- **Weak planning** $E^\omega(\mathcal{T}_{G,\pi}, \varphi)$
  There is a path in $\mathcal{T}_{G,\pi}$ that satisfies φ.
- **Strong planning** $A^\omega(\mathcal{T}_{G,\pi}, \varphi)$
  Every path in $\mathcal{T}_{G,\pi}$ satisfies φ.
- **Strong cyclic planning** $AE^\omega(\mathcal{T}_{G,\pi}, \varphi)$
  Every partial execution of π can be extended to a successful path.

Theorem. The planning problem for LTL-goals, described by Banach-Mazur games, can be solved by automata-based methods and is 2EXPTIME-complete.

Comparing games

$\mathcal{G} \succeq \mathcal{H}$ means that $\mathcal{G}$ is better for Ego than $\mathcal{H}$:

$$\text{Ego wins } \mathcal{H} \quad \Longrightarrow \quad \text{Ego wins } \mathcal{G}$$

$$\text{Alter wins } \mathcal{G} \quad \Longrightarrow \quad \text{Alter wins } \mathcal{H}$$

$\mathcal{G} \equiv \mathcal{H}$ if $\mathcal{G} \succeq \mathcal{H}$ and $\mathcal{H} \succeq \mathcal{G}$

Comparing games

$\mathcal{G} \succeq \mathcal{H}$ means that $\mathcal{G}$ is better for Ego than $\mathcal{H}$:

$$\text{Ego wins } \mathcal{H} \quad \Longrightarrow \quad \text{Ego wins } \mathcal{G}$$

$$\text{Alter wins } \mathcal{G} \quad \Longrightarrow \quad \text{Alter wins } \mathcal{H}$$

$\mathcal{G} \equiv \mathcal{H}$ if $\mathcal{G} \succeq \mathcal{H}$ and $\mathcal{H} \succeq \mathcal{G}$

Problem: Do the game prefixes γ induce an infinite hierarchy (wrt. $\succeq$)?

Comparing games

$\mathcal{G} \succeq \mathcal{H}$ means that $\mathcal{G}$ is better for Ego than $\mathcal{H}$:

$$\text{Ego wins } \mathcal{H} \implies \text{Ego wins } \mathcal{G}$$
$$\text{Alter wins } \mathcal{G} \implies \text{Alter wins } \mathcal{H}$$

$\mathcal{G} \equiv \mathcal{H}$ if $\mathcal{G} \succeq \mathcal{H}$ and $\mathcal{H} \succeq \mathcal{G}$

Problem: Do the game prefixes γ induce an infinite hierarchy (wrt. $\succeq$)?

Obviously, $EAE^\omega(G, W) \succeq (EA)^k E^\omega(G, W)$.

Comparing games

$\mathcal{G} \succeq \mathcal{H}$ means that $\mathcal{G}$ is better for Ego than $\mathcal{H}$:

$$\text{Ego wins } \mathcal{H} \quad \Longrightarrow \quad \text{Ego wins } \mathcal{G}$$
$$\text{Alter wins } \mathcal{G} \quad \Longrightarrow \quad \text{Alter wins } \mathcal{H}$$

$\mathcal{G} \equiv \mathcal{H}$ if $\mathcal{G} \succeq \mathcal{H}$ and $\mathcal{H} \succeq \mathcal{G}$

Problem: Do the game prefixes γ induce an infinite hierarchy (wrt. $\succeq$)?

Obviously, $EAE^\omega(G, W) \succeq (EA)^k E^\omega(G, W)$.

But is $EAE^\omega(G, W)$ strictly better for Ego?

Game equivalence

Proposition. $EAE^\omega(G, W) \equiv (EA)^k E^\omega(G, W)$

Proof. Ego wins $EAE^\omega(G, W) \implies$ Ego wins $(EA)^k E^\omega(G, W)$.

Let $k = 2$.

$EAE^\omega(G, W)$

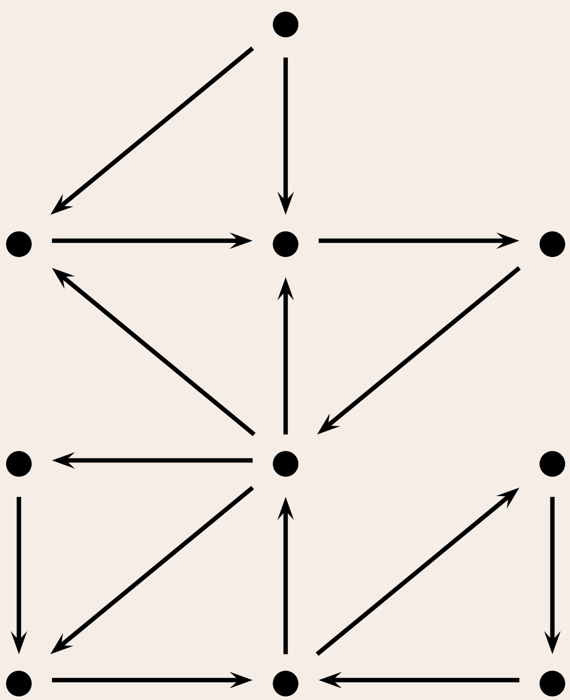

$(EA)^2 E^\omega(G, W)$

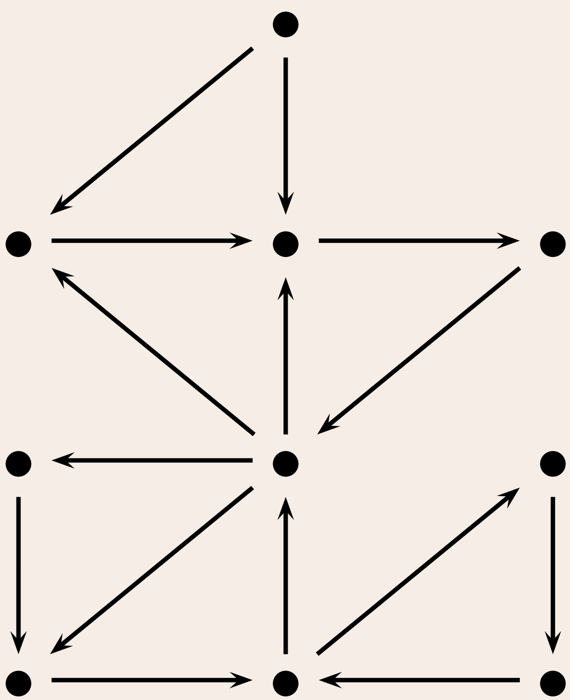

Game equivalence

Proposition. $EAE^\omega(G, W) \equiv (EA)^k E^\omega(G, W)$

Proof. Ego wins $EAE^\omega(G, W)$ $\implies$ Ego wins $(EA)^k E^\omega(G, W)$.

Let $k = 2$.

$EAE^\omega(G, W)$

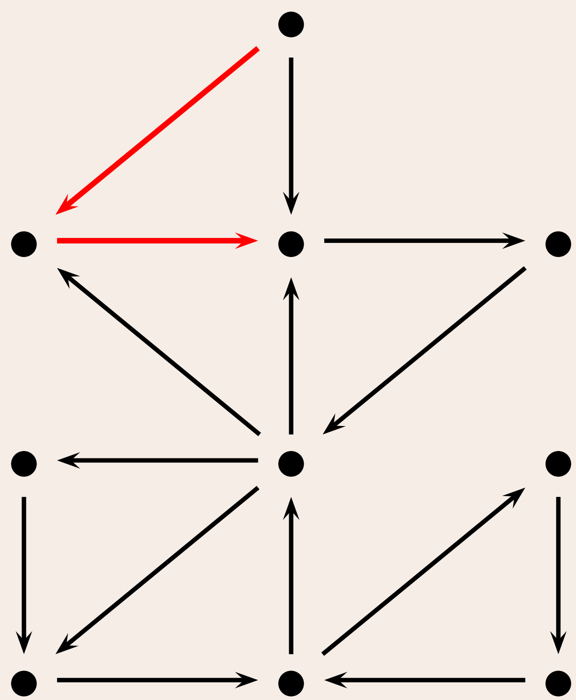

$(EA)^2 E^\omega(G, W)$

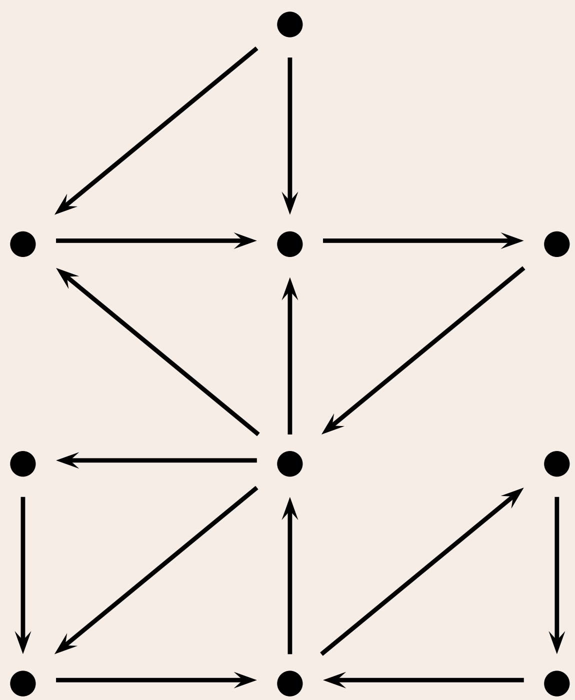

Game equivalence

Proposition. $EAE^{\omega}(G, W) \equiv (EA)^{k}E^{\omega}(G, W)$

Proof. Ego wins $EAE^{\omega}(G, W)$ $\Longrightarrow$ Ego wins $(EA)^{k}E^{\omega}(G, W)$.

Let $k = 2$.

$EAE^{\omega}(G, W)$

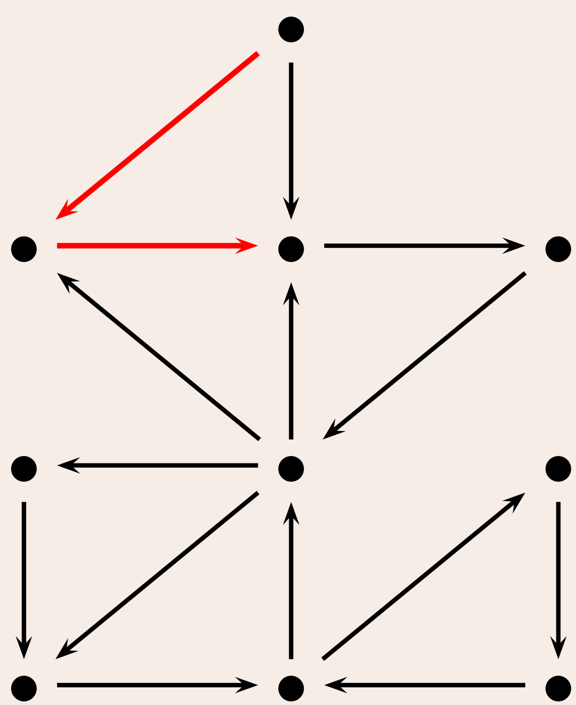

$(EA)^{2}E^{\omega}(G, W)$

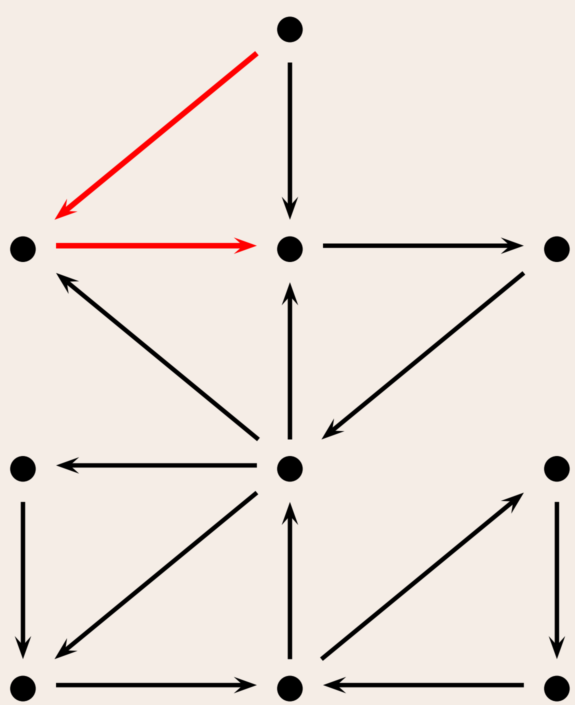

Game equivalence

Proposition. $EAE^\omega(G, W) \equiv (EA)^k E^\omega(G, W)$

Proof. Ego wins $EAE^\omega(G, W) \implies$ Ego wins $(EA)^k E^\omega(G, W)$.

Let $k = 2$.

$EAE^\omega(G, W)$

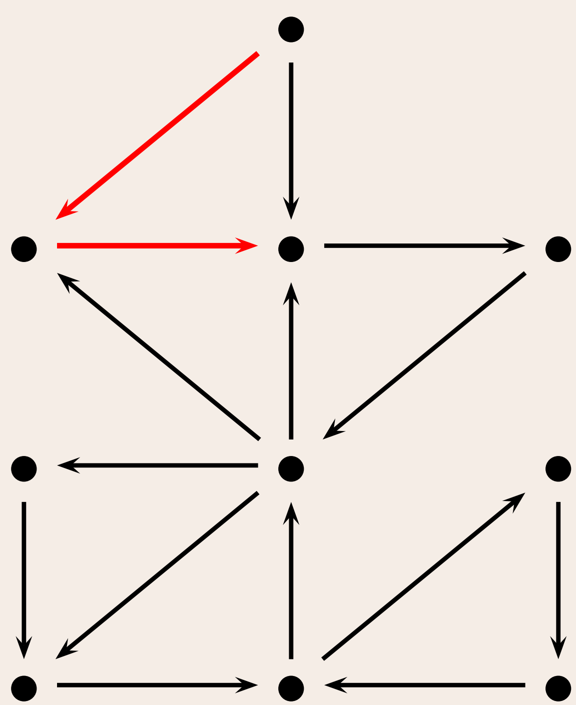

$(EA)^2 E^\omega(G, W)$

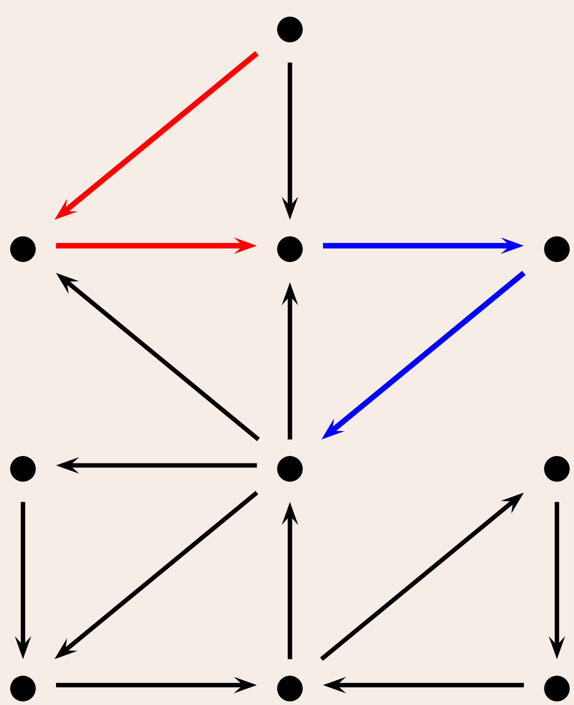

Game equivalence

Proposition. $EAE^{\omega}(G, W) \equiv (EA)^k E^{\omega}(G, W)$

Proof. Ego wins $EAE^{\omega}(G, W) \implies$ Ego wins $(EA)^k E^{\omega}(G, W)$.

Let $k = 2$.

$EAE^{\omega}(G, W)$

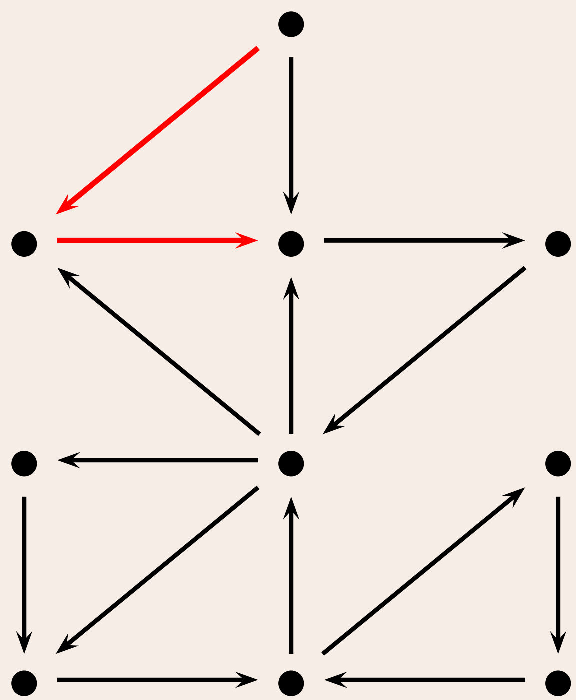

$(EA)^2 E^{\omega}(G, W)$

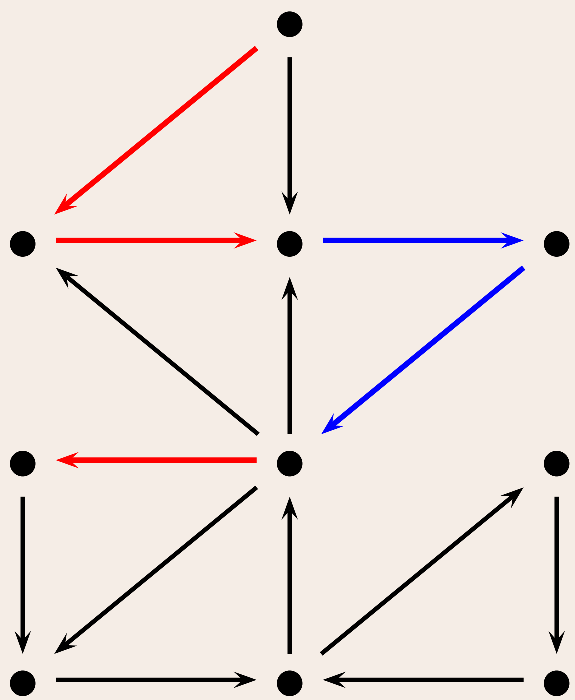

Proposition. $EAE^{\omega}(G, W) \equiv (EA)^k E^{\omega}(G, W)$

Proof. Ego wins $EAE^{\omega}(G, W) \implies$ Ego wins $(EA)^k E^{\omega}(G, W)$.

Let $k = 2$.

$EAE^{\omega}(G, W)$

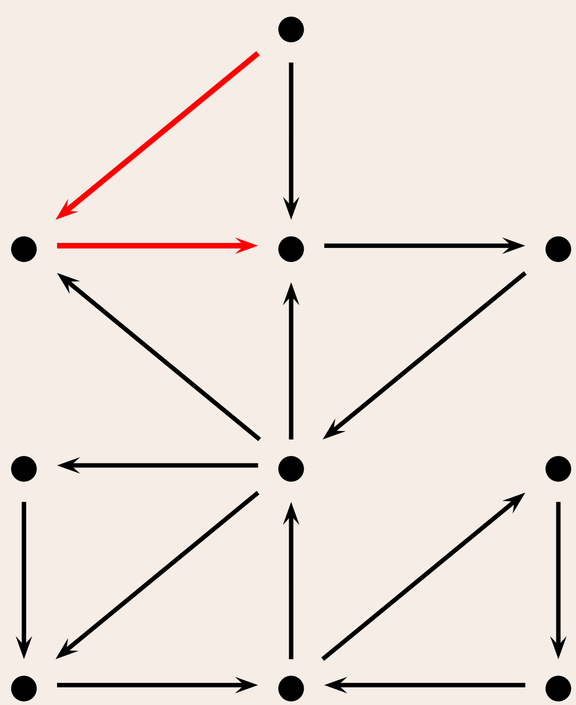

$(EA)^2 E^{\omega}(G, W)$

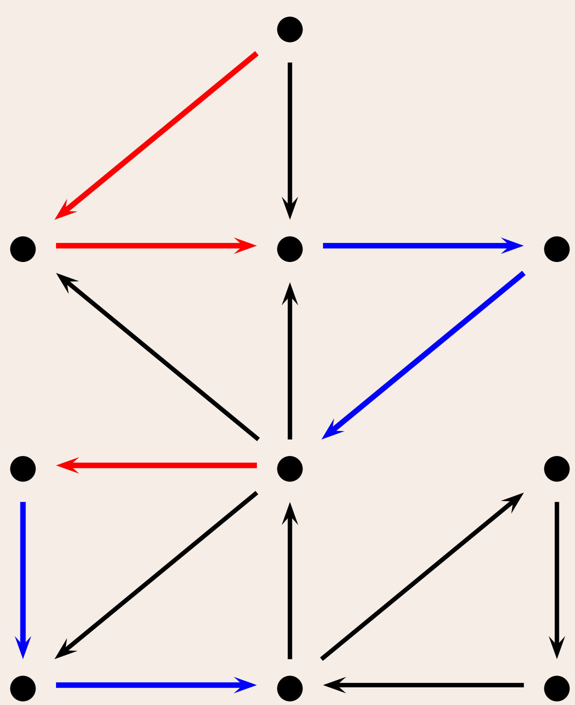

Game equivalence

Proposition. $EAE^\omega(G, W) \equiv (EA)^k E^\omega(G, W)$

Proof. Ego wins $EAE^\omega(G, W)$ $\implies$ Ego wins $(EA)^k E^\omega(G, W)$.

Let $k = 2$.

$EAE^\omega(G, W)$

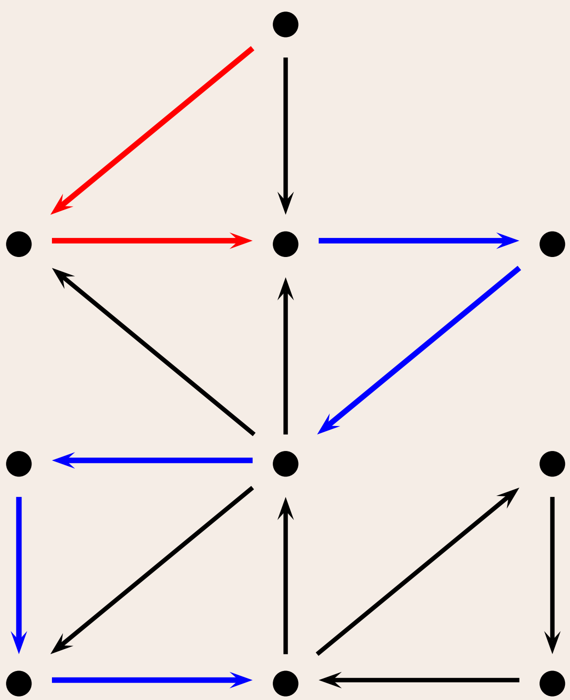

$(EA)^2 E^\omega(G, W)$

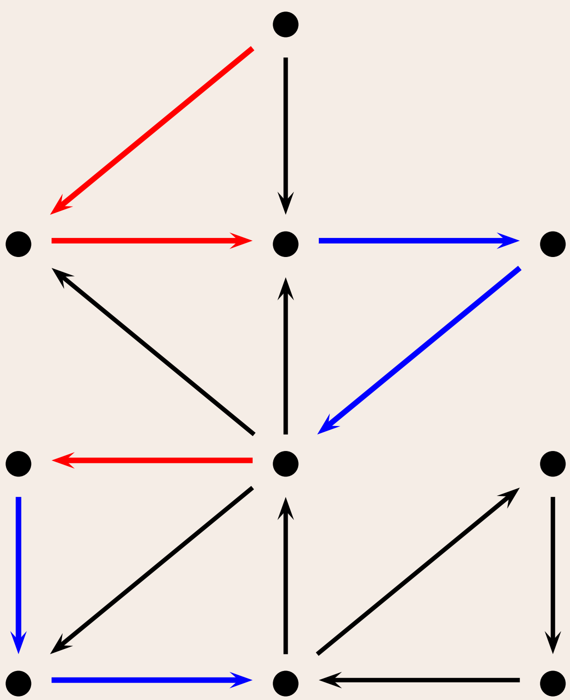

Game equivalence

Proposition. $EAE^\omega(G, W) \equiv (EA)^k E^\omega(G, W)$

Proof. Ego wins $EAE^\omega(G, W) \implies$ Ego wins $(EA)^k E^\omega(G, W)$.

Let $k = 2$.

$EAE^\omega(G, W)$

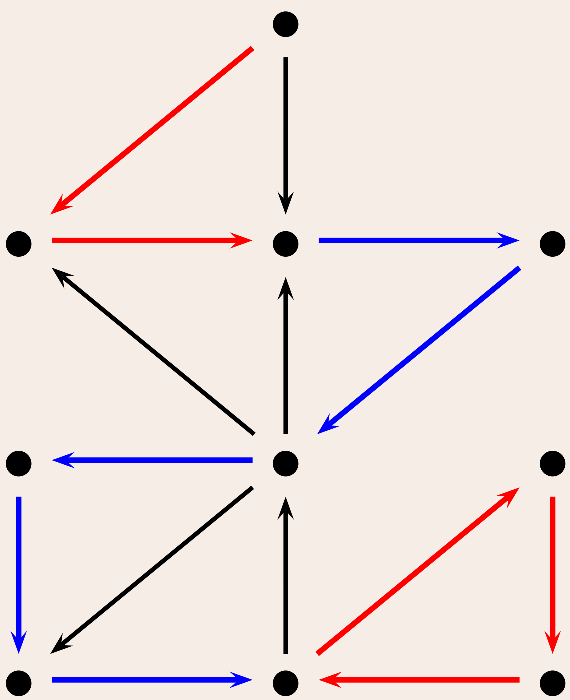

$(EA)^2 E^\omega(G, W)$

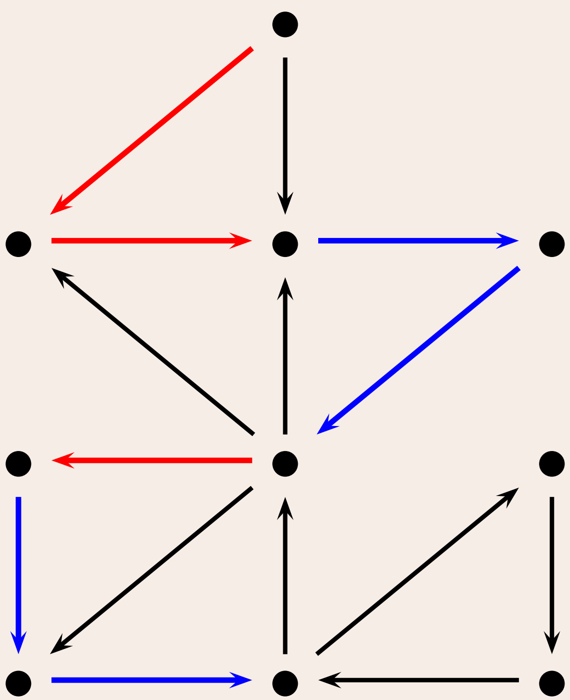

Game equivalence

Proposition. $EAE^{\omega}(G, W) \equiv (EA)^{k}E^{\omega}(G, W)$

Proof. Ego wins $EAE^{\omega}(G, W) \implies$ Ego wins $(EA)^{k}E^{\omega}(G, W)$.

Let $k = 2$.

$EAE^{\omega}(G, W)$

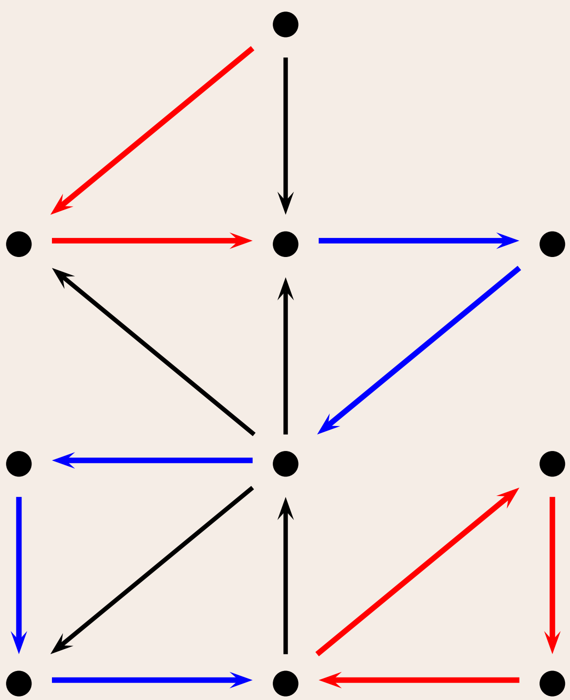

$(EA)^{2}E^{\omega}(G, W)$

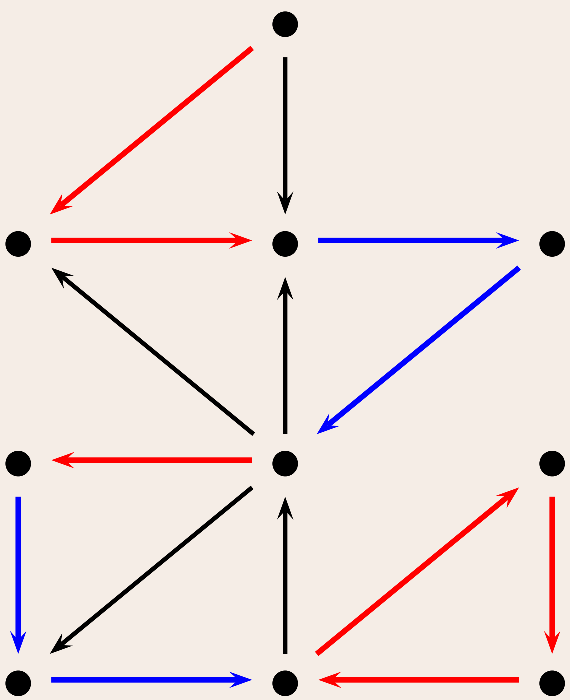

Alternations: one, two, three, infinity

The hierarchy defined by the game prefixes collapses

Theorem. For any game graph G and any winning condition W

$$\begin{array}{ccccc} E^\omega(G,W) & \succeq & EAE^\omega(G,W) & \succeq & AE^\omega(G,W) \\ & & \text{\rotatebox{90}{$\preceq$}} & & \text{\rotatebox{90}{$\preceq$}} \\ & & (EA)^\omega(G,W) & \succeq & (AE)^\omega(G,W) \\ & & \text{\rotatebox{90}{$\preceq$}} & & \text{\rotatebox{90}{$\preceq$}} \\ & & EA^\omega(G,W) & \succeq & AEA^\omega(G,W) & \succeq & A^\omega(G,W) \end{array}$$

Every Banach-Mazur game over (G, W) is equivalent to one of these eight games.

Remark. This holds for games with arbitrary payoff functions (that may take other values than 0 and 1). The games need not be determined.

Start conditions: *Let's have a drink first.*

$$
\begin{array}{ccccccc}
E^{\omega}(G,W) & \succeq & EAE^{\omega}(G,W) & \succeq & AE^{\omega}(G,W) & & \\
 & & \rotatebox{90}{$\preceq$} & & \rotatebox{90}{$\preceq$} & & \\
 & & (EA)^{\omega}(G,W) & \succeq & (AE)^{\omega}(G,W) & & \\
 & & \rotatebox{90}{$\preceq$} & & \rotatebox{90}{$\preceq$} & & \\
 & & EA^{\omega}(G,W) & \succeq & AEA^{\omega}(G,W) & \succeq & A^{\omega}(G,W)
\end{array}
$$

Won by Ego **Won by Alter**

Alternations: one, two, three, infinity

Reachability conditions: *Some day, we'll have a drink.*

Guarantee conditions: *Every day, we have a drink.*

$$
\begin{array}{ccccccc}
E^{\omega}(G,W) & \succeq & EAE^{\omega}(G,W) & \succeq & AE^{\omega}(G,W) & & \\
 & & \rotatebox{-90}{$\succeq$} & & \rotatebox{-90}{$\succeq$} & & \\
 & & (EA)^{\omega}(G,W) & \succeq & (AE)^{\omega}(G,W) & & \\
 & & \rotatebox{-90}{$\succeq$} & & \rotatebox{-90}{$\succeq$} & & \\
 & & EA^{\omega}(G,W) & \succeq & AEA^{\omega}(G,W) & \succeq & A^{\omega}(G,W)
\end{array}
$$

Co-Büchi conditions:

X occurs only finitely often

$$
\begin{array}{ccccccc}
E^\omega(G,W) & \succeq & EAE^\omega(G,W) & \succeq & AE^\omega(G,W) & & \\
 & & \rotatebox{-90}{$\succeq$} & & \rotatebox{-90}{$\succeq$} & & \\
 & & (EA)^\omega(G,W) & \succeq & (AE)^\omega(G,W) & & \\
 & & \rotatebox{-90}{$\succeq$} & & \rotatebox{-90}{$\succeq$} & & \\
 & & EA^\omega(G,W) & \succeq & AEA^\omega(G,W) & \succeq & A^\omega(G,W)
\end{array}
$$

Won by Ego

Won by Alter

Büchi conditions: *Y occurs infinitely often*

$$\begin{array}{ccccc}
E^\omega(G,W) & \succeq & EAE^\omega(G,W) & \succeq & AE^\omega(G,W) \\
& & \rotatebox{-90}{$\succeq$} & & \rotatebox{-90}{$\succeq$} \\
& & (EA)^\omega(G,W) & \succeq & (AE)^\omega(G,W) \\
& & \rotatebox{-90}{$\succeq$} & & \rotatebox{-90}{$\succeq$} \\
& & EA^\omega(G,W) & \succeq & AEA^\omega(G,W) & \succeq & A^\omega(G,W)
\end{array}$$

Won by Ego

Won by Alter

Alternations: one, two, three, infinity

$$
\begin{array}{ccccccccc}
E^\omega(G,W) & \succeq & EAE^\omega(G,W) & \succeq & AE^\omega(G,W) & & \\
 & & \rotatebox{-90}{$\succeq$} & & \rotatebox{-90}{$\succeq$} & & \\
 & & (EA)^\omega(G,W) & \succeq & (AE)^\omega(G,W) & & \\
 & & \rotatebox{-90}{$\succeq$} & & \rotatebox{-90}{$\succeq$} & & \\
 & & EA^\omega(G,W) & \succeq & AEA^\omega(G,W) & \succeq & A^\omega(G,W)
\end{array}
$$

Definability of Banach-Mazur games

Question: How relates the logical complexity of defining the winner of a game to the logical complexity of the winning condition ?

Definability of Banach-Mazur games

Question: How relates the logical complexity of defining the winner of a game to the logical complexity of the winning condition ?

From a game prefix γ and a formula φ (e.g. from LTL or S1S) defining a winning condition, we obtain a game formula $\gamma.\varphi$, with

$$G \models \gamma.\varphi \quad \Longleftrightarrow \quad \text{Ego wins the game } \gamma(G, \varphi).$$

Definability of Banach-Mazur games

Question: How relates the logical complexity of defining the winner of a game to the logical complexity of the winning condition ?

From a game prefix γ and a formula φ (e.g. from LTL or S1S) defining a winning condition, we obtain a game formula $\gamma.\varphi$, with

$$G \models \gamma.\varphi \quad \Longleftrightarrow \quad \text{Ego wins the game } \gamma(G, \varphi).$$

This maps a logic L on paths to a logic $\gamma.L := \{\gamma.\varphi : \varphi \in L\}$ on graphs.

Definability of Banach-Mazur games

Question: How relates the logical complexity of defining the winner of a game to the logical complexity of the winning condition ?

From a game prefix γ and a formula φ (e.g. from LTL or S1S) defining a winning condition, we obtain a game formula $\gamma.\varphi$, with

$$G \models \gamma.\varphi \quad \Longleftrightarrow \quad \text{Ego wins the game } \gamma(G, \varphi).$$

This maps a logic L on paths to a logic $\gamma.L := \{\gamma.\varphi : \varphi \in L\}$ on graphs.

Question: Expressive power of γ.LTL and γ.S1S, compared to common logics on (game) graphs, like μ-calculus, CTL*, FO, and MSO ?

Definability theorem for Banach-Mazur games

Theorem. For any game prefix γ

(1) $\gamma\,.\,\mathrm{S1S} \leq L_\mu$

(2) $\gamma\,.\,\mathrm{LTL} \equiv \gamma\,.\,\mathrm{FO} \leq \mathrm{CTL}^*$

That is, the winner of any Banach-Mazur with S1S resp. LTL winning condition is definable in the modal μ-calculus resp. CTL^*.

Simplification via bisimulation invariance

It suffices to prove, that on trees,

(1) $\gamma\,.\,\text{S1S} \leq \text{MSO}$

(2) $\gamma\,.\,\text{FO} \leq \text{MPL}$

MPL is monadic path logic: MSO on trees with second-order quantification restricted to (finite or infinite) paths.

Simplification via bisimulation invariance

It suffices to prove, that on trees,

(1) $\gamma\,.\,\mathrm{S1S} \leq \mathrm{MSO}$

(2) $\gamma\,.\,\mathrm{FO} \leq \mathrm{MPL}$

MPL is monadic path logic: MSO on trees with second-order quantification restricted to (finite or infinite) paths.

- winning a (Banach-Mazur) game is invariant under bisimulation
- bisimulation-invariant MSO $\equiv$ L_μ (Janin/Walukiewicz)
- bisimulation-invariant MPL $\equiv$ CTL^* (Hafer/Thomas) (Moller/Rabinovitch)

The simple case: games with finite alternations

Ego wins $EAE^{\omega}(\mathcal{T}, \varphi) \iff \mathcal{T} \models \psi$, where

$$\psi := (\exists X \,.\, X \text{ finite path})(\forall Y \,.\, X \subseteq Y \wedge Y \text{ finite path})$$
$$(\exists Z \,.\, Y \subseteq Z \wedge Z \text{ infinite path} \wedge \varphi|_Z)$$

$\varphi|_Z := \varphi$ relativized to the path Z

- $\varphi \in \mathrm{S1S} \implies \psi \in \mathrm{MSO}$
- $\varphi \in \mathrm{LTL} \implies \psi \in \mathrm{MPL}$

The slightly harder case: games with infinite alternations

Winning strategy of Ego for $(EA)^\omega(\mathcal{T}, \varphi)$ on tree $\mathcal{T} = (V, E)$: described by set $X \subseteq V$ with

- X is non-empty
- for all $x \in X$ and $y > x$ there is a $z > y$ with $z \in X$
- every path through $\mathcal{T}$ hitting X infinitely often satisfies φ

The slightly harder case: games with infinite alternations

Winning strategy of Ego for $(EA)^\omega(\mathcal{T}, \varphi)$ on tree $\mathcal{T} = (V, E)$:
described by set $X \subseteq V$ with

- X is non-empty
- for all $x \in X$ and $y > x$ there is a $z > y$ with $z \in X$
- every path through $\mathcal{T}$ hitting X infinitely often satisfies φ

Obviously, this can be formalised in MSO, if $\varphi \in$ S1S.

The slightly harder case: games with infinite alternations

Winning strategy of Ego for $(EA)^\omega(\mathcal{T}, \varphi)$ on tree $\mathcal{T} = (V, E)$:
described by set $X \subseteq V$ with

- X is non-empty
- for all $x \in X$ and $y > x$ there is a $z > y$ with $z \in X$
- every path through $\mathcal{T}$ hitting X infinitely often satisfies φ

Obviously, this can be formalised in MSO, if $\varphi \in$ S1S.

For $\varphi \in$ FO, we have to formalise in MPL.
In fact we can even formalise in FO !

Normal form for FO:

on infinite paths, every first-order formula is equivalent to

$$\bigvee_i \Big(\exists x (\forall y \geq x) \varphi_i \wedge \forall x (\exists y \geq x) \vartheta_i \Big)$$

where φ_i and ϑ_i contain only bounded quantifiers $(Qz \leq y)$.

In terms of LTL: Every LTL-formula is equivalent to a disjunction of formulae $(\mathrm{FG}\varphi \wedge \mathrm{GF}\vartheta)$, where φ and ϑ are past-formulae.

Use this to show that on trees,

$$(EA)^\omega.\,\mathrm{FO} \leq \mathrm{FO} \quad \text{and} \quad (AE)^\omega.\,\mathrm{FO} \leq \mathrm{FO}.$$

Conclusion

- Banach-Mazur games are a natural kind of games on graphs.
- Mathmatically and algorithmically, Banach-Mazur games tend to be simpler than the common single-step graph games.
- Finite memory determinacy collapses to positional determinacy.
- All Banach-Mazur games with ω-regular winning conditions are positionally determined. Muller games can be solved efficiently.
- For games with bounded alternations, the hierarchy defined by game prefixes collapses to eight different games.
- The winner of ω-regular Banch-Mazur games can be defined in the modal μ-calculus, and the winner of LTL Banach-Mazur games can be defined in CTL*.

The End