The Complexity of Finding Nash Equilibria in Infinite Multiplayer Games

Michael Ummels
ummels@logic.rwth-aachen.de

RWTH AACHEN

ILC 2007

Infinite Games

Let's play!

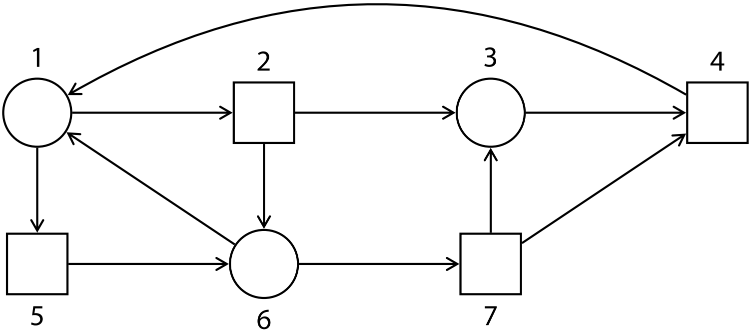

Infinite Games

Let's play!

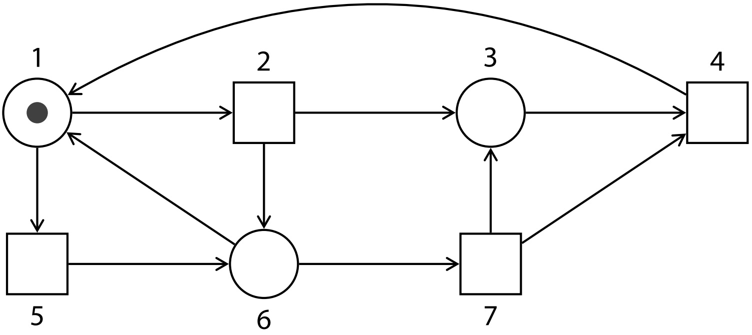

Play: $\pi = 1,$

Infinite Games

Let's play!

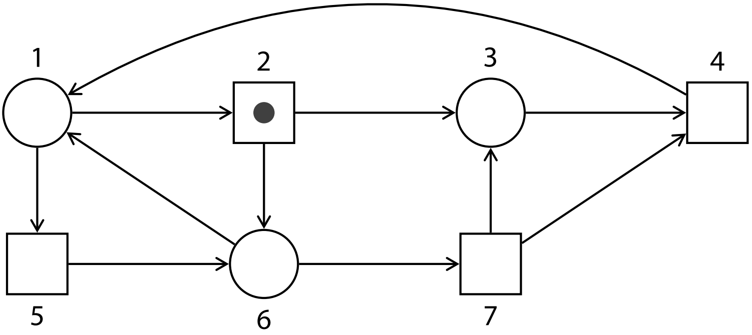

Play: $\pi = 1, 2,$

Let's play!

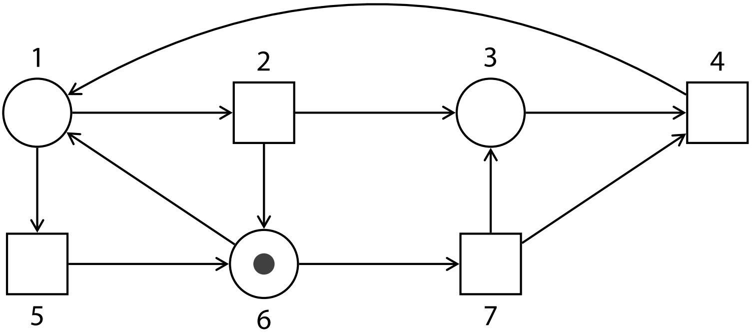

Play: $\pi = 1, 2, 6,$

Infinite Games

Let's play!

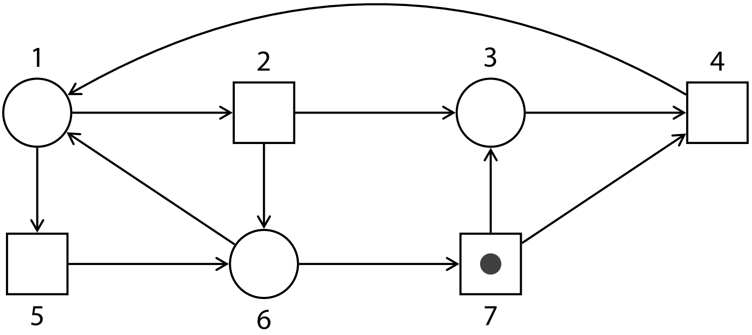

Play: $\pi = 1, 2, 6, 7,$

Infinite Games

Let's play!

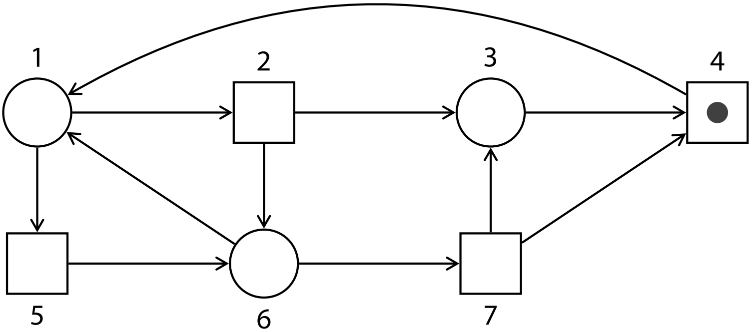

Play: $\pi = 1, 2, 6, 7, 4,$

Let's play!

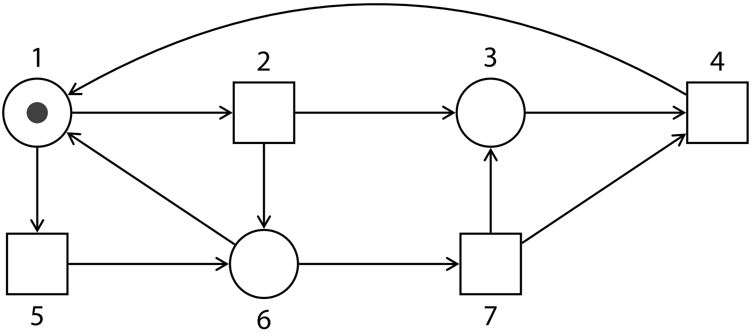

Play: $\pi = 1, 2, 6, 7, 4, 1,$

Infinite Games

Let's play!

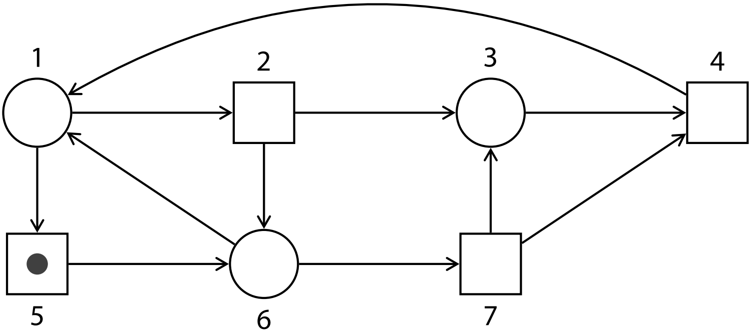

Play: $\pi = 1, 2, 6, 7, 4, 1, 5,$

Infinite Games

Let's play!

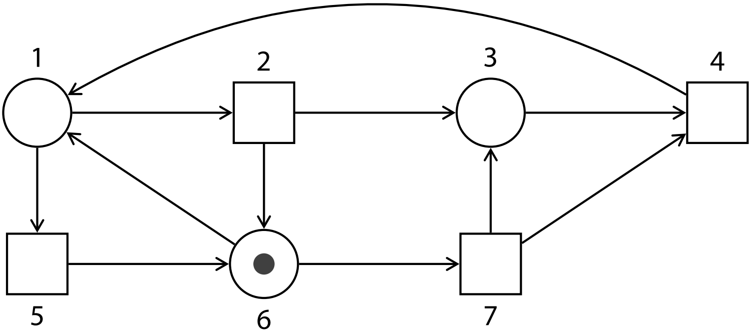

Play: $\pi = 1, 2, 6, 7, 4, 1, 5, 6,$

Infinite Games

Let's play!

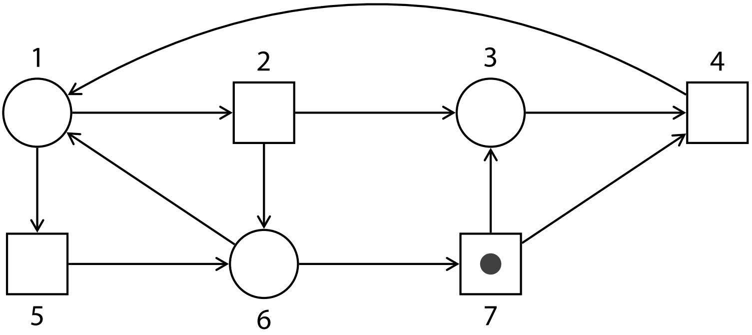

Play: $\pi = 1, 2, 6, 7, 4, 1, 5, 6, 7,$

Infinite Games

Let's play!

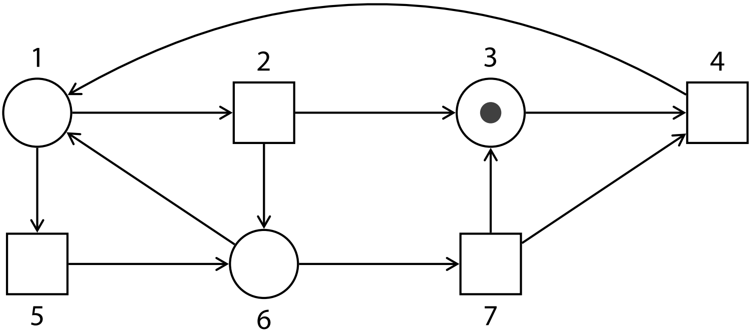

Play: $\pi = 1, 2, 6, 7, 4, 1, 5, 6, 7, 3,$

Let's play!

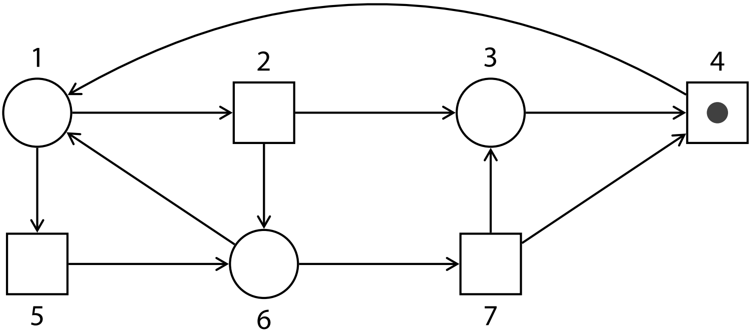

Play: $\pi = 1, 2, 6, 7, 4, 1, 5, 6, 7, 3, 4,$

Infinite Games

Let's play!

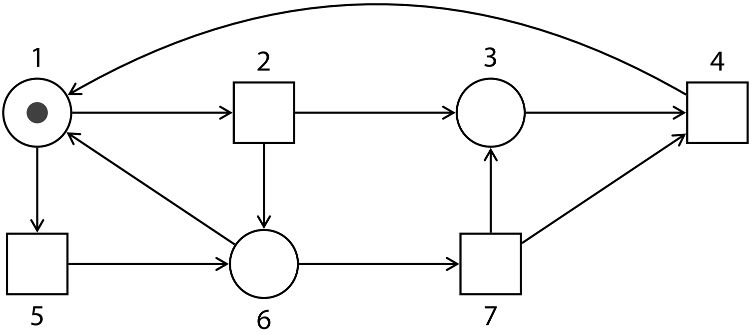

Play: $\pi = 1, 2, 6, 7, 4, 1, 5, 6, 7, 3, 4, 1,$

Infinite Games

Let's play!

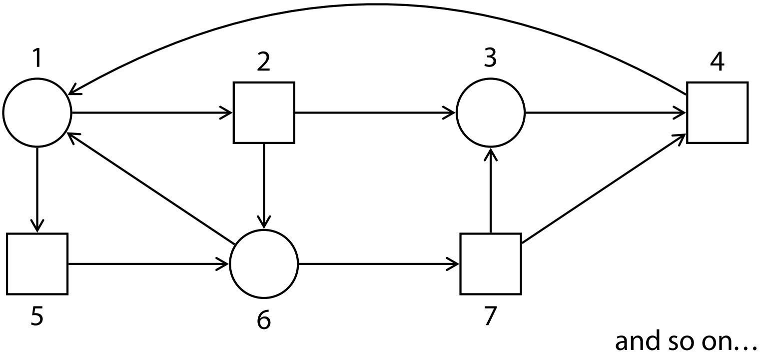

Play: $\pi = 1, 2, 6, 7, 4, 1, 5, 6, 7, 3, 4, 1, \ldots$

Winning conditions

Question: What is the payoff of a play?

Winning conditions

Question: What is the payoff of a play?

Specified by a winning condition for each player:

- Büchi condition: Given a set F of vertices, defines the set of all plays π such that $\pi(k) \in F$ for *infinitely many* positions k.
- Co-Büchi condition: Given a set F of vertices, defines the set of all plays π such that $\pi(k) \in F$ only for *finitely many* positions k.

Player receives payoff 1 if her winning condition is satisfied, otherwise 0.

Winning conditions

Question: What is the payoff of a play?

Specified by a winning condition for each player:

- Büchi condition: Given a set F of vertices, defines the set of all plays π such that $\pi(k) \in F$ for *infinitely many* positions k.
- Co-Büchi condition: Given a set F of vertices, defines the set of all plays π such that $\pi(k) \in F$ only for *finitely many* positions k.

Player receives payoff 1 if her winning condition is satisfied, otherwise 0.

But we are not so much interested in the winner of a certain play, but in the strategic behaviour that can occur.

Classical game problem: Given a game $\mathcal{G}$, a player i and a vertex v, decide whether player i has a *winning strategy* from v.

The Classical Case

Two-player zero-sum Games: Games with two players where winning condition of one player is the complement of the winning condition of the other player.

The Classical Case

Two-player zero-sum Games: Games with two players where winning condition of one player is the complement of the winning condition of the other player.

Two-player zero-sum (co-)Büchi games: If one player has a Büchi winning condition with set F, then the other player has a co-Büchi winning condition with the same set F.

The Classical Case

Two-player zero-sum Games: Games with two players where winning condition of one player is the complement of the winning condition of the other player.

Two-player zero-sum (co-)Büchi games: If one player has a Büchi winning condition with set F, then the other player has a co-Büchi winning condition with the same set F.

Determinacy: A two-player zero-sum game is determined if from each vertex either one of the two players has a winning strategy.

The Classical Case

Two-player zero-sum Games: Games with two players where winning condition of one player is the complement of the winning condition of the other player.

Two-player zero-sum (co-)Büchi games: If one player has a Büchi winning condition with set F, then the other player has a co-Büchi winning condition with the same set F.

Determinacy: A two-player zero-sum game is determined if from each vertex either one of the two players has a winning strategy.

Theorem (Büchi-Landweber)

Two-player zero-sum (co-)Büchi games are determined.

The Classical Case

Two-player zero-sum Games: Games with two players where winning condition of one player is the complement of the winning condition of the other player.

Two-player zero-sum (co-)Büchi games: If one player has a Büchi winning condition with set F, then the other player has a co-Büchi winning condition with the same set F.

Determinacy: A two-player zero-sum game is determined if from each vertex either one of the two players has a winning strategy.

Theorem (Büchi-Landweber)

Two-player zero-sum (co-)Büchi games are determined.

Solving these games corresponds to computing the winning regions.

Nash Equilibria

In the general case, players may not have a winning strategy, but still they could win by cooperating with other players.

Nash Equilibria

In the general case, players may not have a winning strategy, but still they could win by cooperating with other players.

Definition: A choice of strategies (one for each player) is a Nash equilibrium if no player can gain from unilaterally switching to a different strategy.

Motivation: In a Nash equilibrium, no player has an incentive to switch strategies.

Nash Equilibria

In the general case, players may not have a winning strategy, but still they could win by cooperating with other players.

Definition: A choice of strategies (one for each player) is a Nash equilibrium if no player can gain from unilaterally switching to a different strategy.

Motivation: In a Nash equilibrium, no player has an incentive to switch strategies.

Question: Do Nash equilibria always exist?

Theorem (Chatterjee & al.)

Any infinite game with Borel winning conditions has a Nash equilibrium (independent of the initial vertex).

Nash Equilibria

In the general case, players may not have a winning strategy, but still they could win by cooperating with other players.

Definition: A choice of strategies (one for each player) is a Nash equilibrium if no player can gain from unilaterally switching to a different strategy.

Motivation: In a Nash equilibrium, no player has an incentive to switch strategies.

Question: Do Nash equilibria always exist?

Theorem (Chatterjee & al.)

Any infinite game with Borel winning conditions has a Nash equilibrium (independent of the initial vertex).

But there may be many Nash equilibria (with different payoffs)...

Example

Winning condition for player 0 (round vertices): Visit 3 infinitely often.

Winning condition for player 1 (square vertices): Visit 5 infinitely often.

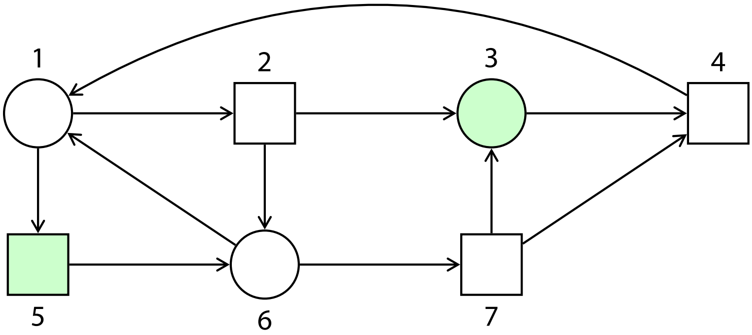

Example

Winning condition for player 0 (round vertices): Visit 3 infinitely often.

Winning condition for player 1 (square vertices): Visit 5 infinitely often.

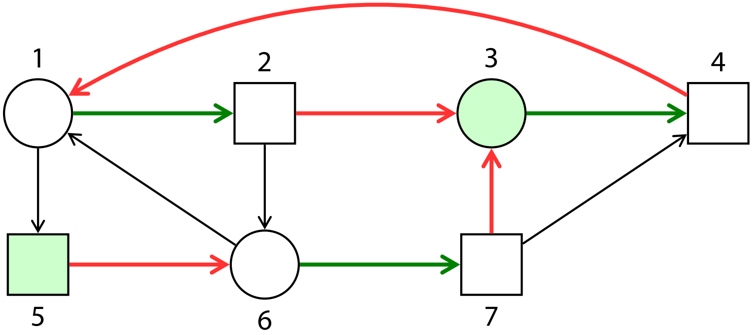

Example

Winning condition for player 0 (round vertices): Visit 3 infinitely often.

Winning condition for player 1 (square vertices): Visit 5 infinitely often.

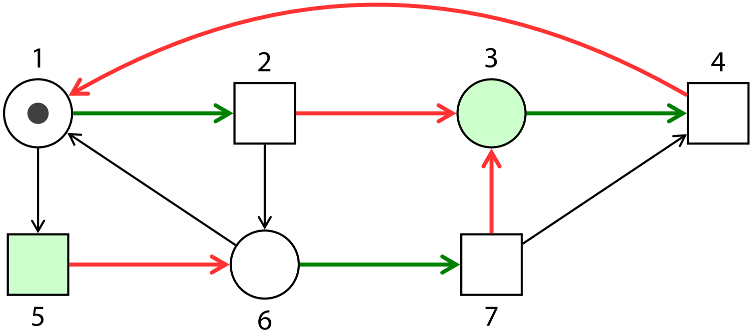

Example

Winning condition for player 0 (round vertices): Visit 3 infinitely often.

Winning condition for player 1 (square vertices): Visit 5 infinitely often.

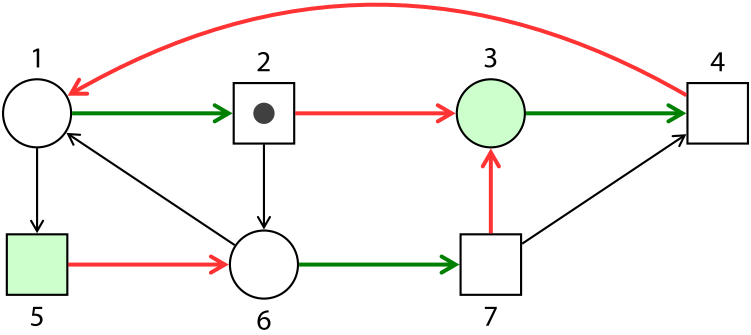

Example

Winning condition for player 0 (round vertices): Visit 3 infinitely often.

Winning condition for player 1 (square vertices): Visit 5 infinitely often.

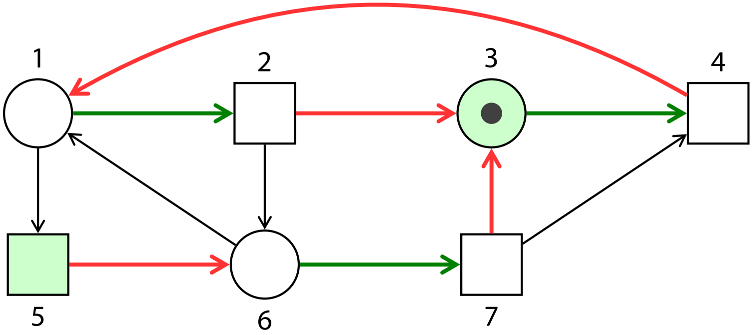

Example

Winning condition for player 0 (round vertices): Visit 3 infinitely often.

Winning condition for player 1 (square vertices): Visit 5 infinitely often.

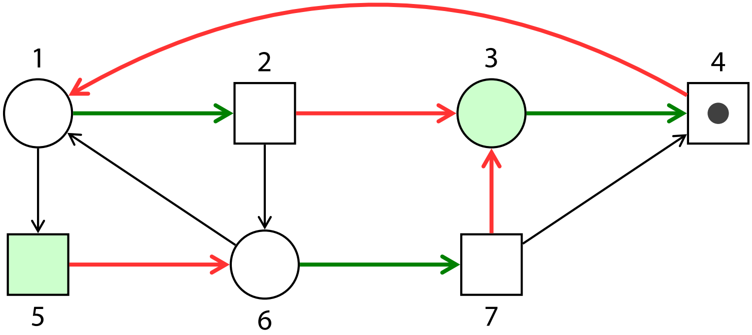

Example

Winning condition for player 0 (round vertices): Visit 3 infinitely often.

Winning condition for player 1 (square vertices): Visit 5 infinitely often.

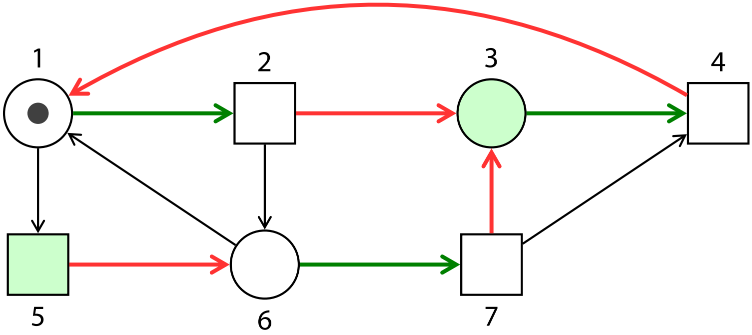

Example

Winning condition for player 0 (round vertices): Visit 3 infinitely often.

Winning condition for player 1 (square vertices): Visit 5 infinitely often.

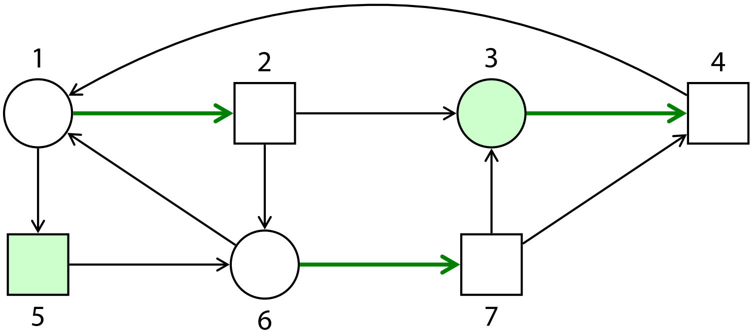

Example

Winning condition for player 0 (round vertices): Visit 3 infinitely often.

Winning condition for player 1 (square vertices): Visit 5 infinitely often.

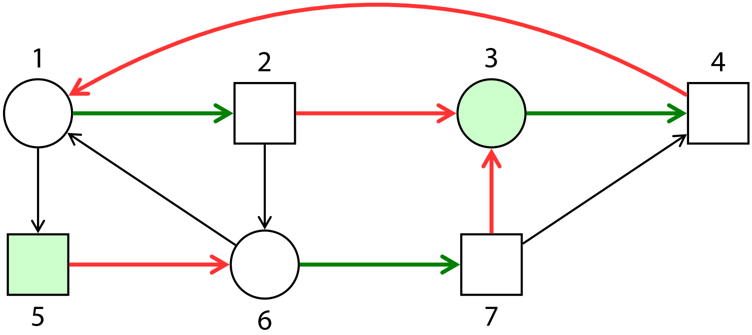

Fact: Player 0's winning region is empty, but there is a Nash equilibrium where she wins (for any initial vertex).

The Problem NE

Two computational problems:

- Compute any Nash equilibrium.
- Compute one with certain requirements on the payoff.

We focus on the second (more involved) task.

The Problem NE

Two computational problems:

- Compute any Nash equilibrium.
- Compute one with certain requirements on the payoff.

We focus on the second (more involved) task.

The problem NE: Given a game $\mathcal{G}$, two payoff thresholds $\overline{x}, \overline{y} \in \{0,1\}^k$ and an initial vertex v, decide whether the game has a Nash equilibrium with payoff $\overline{x} \leq \overline{z} \leq \overline{y}$.

Thresholds determine which players may win and which players shall lose.

The Problem NE

Two computational problems:

- Compute any Nash equilibrium.
- Compute one with certain requirements on the payoff.

We focus on the second (more involved) task.

The problem NE: Given a game $\mathcal{G}$, two payoff thresholds $\overline{x}, \overline{y} \in \{0,1\}^k$ and an initial vertex v, decide whether the game has a Nash equilibrium with payoff $\overline{x} \leq \overline{z} \leq \overline{y}$.

Thresholds determine which players may win and which players shall lose.

Remark: The game problem is decidable in polynomial time for games with (co-)Büchi winning conditions.

Question: What is the complexity of NE?

Characterising Nash Equilibria

Necessary condition for having a Nash equilibrium with payoff $\overline{z} \in \{0,1\}^k$:

- There must be a play with payoff $\overline{z}$.
- This play must avoid the winning region of each player i with $z_i = 0$.

Characterising Nash Equilibria

Necessary condition for having a Nash equilibrium with payoff $\overline{z} \in \{0,1\}^k$:

- There must be a play with payoff $\overline{z}$.
- This play must avoid the winning region of each player i with $z_i = 0$.

Key observation: This condition is also sufficient!

Characterising Nash Equilibria

Necessary condition for having a Nash equilibrium with payoff $\overline{z} \in \{0, 1\}^k$:

- There must be a play with payoff $\overline{z}$.
- This play must avoid the winning region of each player i with $z_i = 0$.

Key observation: This condition is also sufficient!

Proof idea:

- Let all players play according to the given play.
- If some player deviates, let all the other players switch to a joint counter-strategy (threat strategy).

Characterising Nash Equilibria

Necessary condition for having a Nash equilibrium with payoff $\overline{z} \in \{0,1\}^k$:

- There must be a play with payoff $\overline{z}$.
- This play must avoid the winning region of each player i with $z_i = 0$.

Key observation: This condition is also sufficient!

Proof idea:

- Let all players play according to the given play.
- If some player deviates, let all the other players switch to a joint counter-strategy (threat strategy).

Remark: Proof relies on winning conditions to be determined.

Characterising Nash Equilibria

Necessary condition for having a Nash equilibrium with payoff $\overline{z} \in \{0,1\}^k$:

- There must be a play with payoff $\overline{z}$.
- This play must avoid the winning region of each player i with $z_i = 0$.

Key observation: This condition is also sufficient!

Proof idea:

- Let all players play according to the given play.
- If some player deviates, let all the other players switch to a joint counter-strategy (threat strategy).

Remark: Proof relies on winning conditions to be determined.

Let's turn this characterisation into an algorithm!

Deciding NE

Theorem

NE is in NP for games with (co-)Büchi winning conditions.

Deciding NE

Theorem

NE is in NP for games with (co-)Büchi winning conditions.

To prove the theorem, we describe an NP algorithm for NE:

1. Compute the winning region W_i for each player i.
2. Guess a payoff $\overline{x} \leq \overline{z} \leq \overline{y}$.
3. For each player i with $z_i = 0$, remove all vertices in W_i.
4. Check whether there exists a play with payoff $\overline{z}$ in the induced subgraph.

Deciding NE

Theorem

NE is in NP for games with (co-)Büchi winning conditions.

To prove the theorem, we describe an NP algorithm for NE:

1. Compute the winning region W_i for each player i.
2. Guess a payoff $\overline{x} \leq \overline{z} \leq \overline{y}$.
3. For each player i with $z_i = 0$, remove all vertices in W_i.
4. Check whether there exists a play with payoff $\overline{z}$ in the induced subgraph.

The check corresponds to deciding the existence of a winning strategy in a one-player Strett game (decidable in polynomial time).

Deciding NE

Theorem

NE is in NP for games with (co-)Büchi winning conditions.

To prove the theorem, we describe an NP algorithm for NE:

1. Compute the winning region W_i for each player i.
2. Guess a payoff $\overline{x} \leq \overline{z} \leq \overline{y}$.
3. For each player i with $z_i = 0$, remove all vertices in W_i.
4. Check whether there exists a play with payoff $\overline{z}$ in the induced subgraph.

The check corresponds to deciding the existence of a winning strategy in a one-player Strett game (decidable in polynomial time).

Remark: Extends to games with parity or Streett winning conditions.

Games with a Small Number of Players

Question: Can we do better if the number of players is small?

Games with a Small Number of Players

Question: Can we do better if the number of players is small?

Theorem

NE is decidable in polynomial time for games with (co-)Büchi winning conditions where the number of players is bounded.

Games with a Small Number of Players

Question: Can we do better if the number of players is small?

Theorem

NE is decidable in polynomial time for games with (co-)Büchi winning conditions where the number of players is bounded.

Here is a polynomial-time algorithm for the problem:

1. Compute the winning region W_i for each player i.
2. For each payoff $\overline{x} \leq \overline{z} \leq \overline{y}$ do:
   - For each player i with $z_i = 0$, temporarily remove all vertices in W_i.
   - Check whether there exists a play with payoff $\overline{z}$ in the induced subgraph.

Games with a Small Number of Players

Question: Can we do better if the number of players is small?

Theorem

NE is decidable in polynomial time for games with (co-)Büchi winning conditions where the number of players is bounded.

Here is a polynomial-time algorithm for the problem:

1. Compute the winning region W_i for each player i.
2. For each payoff $\overline{x} \le \overline{z} \le \overline{y}$ do:
   - For each player i with $z_i = 0$, temporarily remove all vertices in W_i.
   - Check whether there exists a play with payoff $\overline{z}$ in the induced subgraph.

Remark: NE is also decidable in polynomial time for games with Büchi winning conditions only.

NP-Hardness

Theorem

NE is NP-hard for multiplayer co-Büchi games.

NP-Hardness

Theorem

NE is NP-hard for multiplayer co-Büchi games.

The proof is by a reduction from SAT.

NP-Hardness

Theorem

NE is NP-hard for multiplayer co-Büchi games.

The proof is by a reduction from SAT.

Given a formula φ in CNF over variables $X_1, \ldots, X_k$, we construct (in polynomial time) a co-Büchi game $\mathcal{G}_\varphi$ with players $0, 1, \ldots, k$ and initial position v such that:

$$
\begin{aligned}
& \varphi \text{ is satisfiable} \\
\Leftrightarrow\ & (\mathcal{G}_\varphi, v) \text{ has a Nash equilibrium where player 0 wins.} \\
\Leftrightarrow\ & (\mathcal{G}_\varphi, v) \text{ has a Nash equilibrium with payoff} \geq (1, 0, \ldots, 0) \\
& \text{and} \leq (1, 1, \ldots, 1).
\end{aligned}
$$

The Construction

The game $\mathcal{G}_\varphi$:

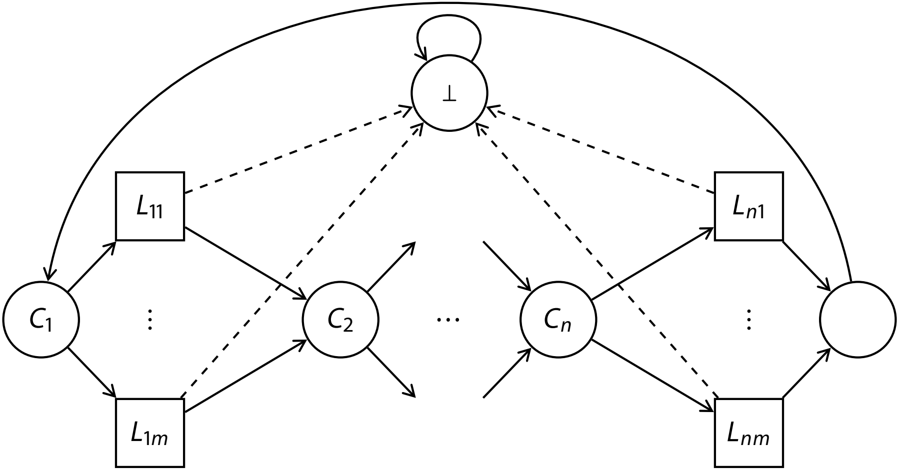

The Construction

The game $\mathcal{G}_\varphi$:

Edges to $\bot$ only from negated variables!

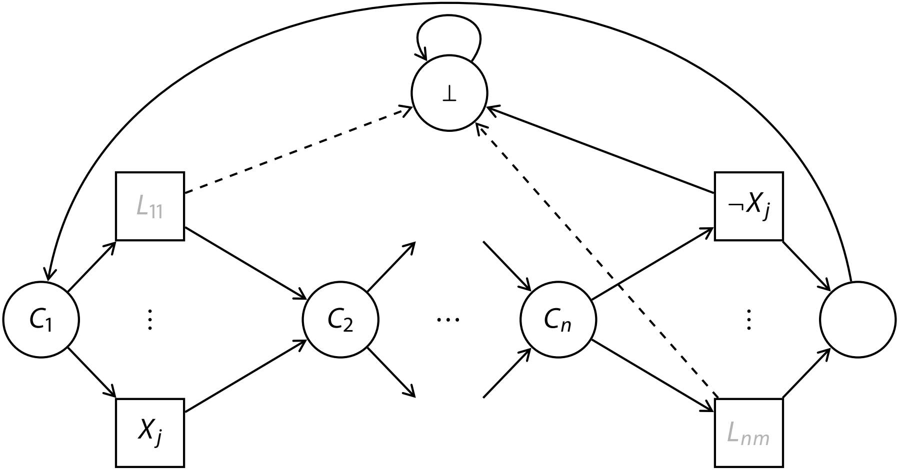

The Construction

Winning condition of player 0 (controls clauses):

Visit $\perp$ only finitely often (i.e. never).

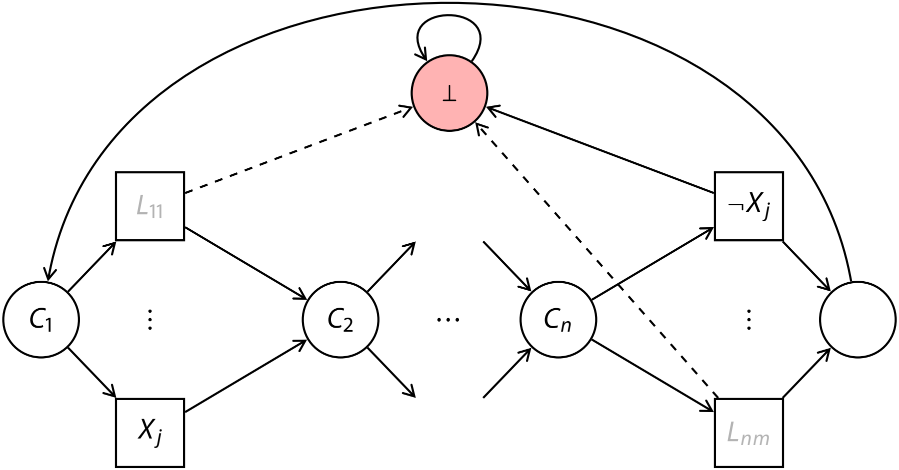

The Construction

Winning condition of player j (controls literals X_j and $\neg X_j$):
Visit X_j only finitely often.

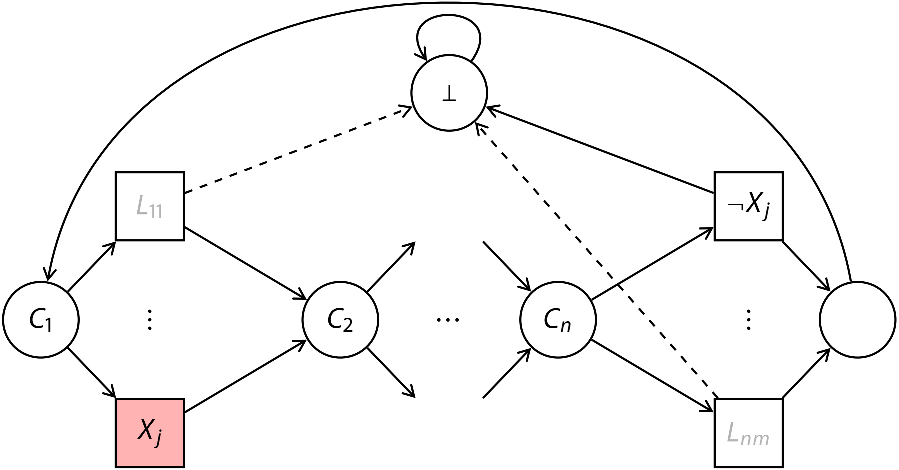

Conclusion

What you might want to remember from this talk:

Finding Nash equilibria in infinite games is hard!

Conclusion

What you might want to remember from this talk:

Finding Nash equilibria in infinite games is hard!

However, it can be done efficiently if…

- there is only a bounded number of players.
- there are only Büchi winning conditions.