

Rational Behaviour and Strategy Construction in Infinite Multiplayer Games

Michael Ummels
ummels@logic.rwth-aachen.de



FSTTCS 2006

Infinite Multiplayer Games

An infinite multiplayer game \mathcal{G} consists of:

- a finite set Π of **players**,
- a directed graph $G = (V, E)$ (the *arena*),
- a partition $(V_i)_{i \in \Pi}$ assigning each vertex a player,
- a mapping $\chi : V \rightarrow C$ into a finite set C of **colours**,
- for each player i , a **winning condition** $W_i \subseteq C^\omega$.

Notation: (\mathcal{G}, v_0) for a game with **initial vertex** v_0 .

How to play?

Intuitively, a token is moved through the graph forming an infinite path (a **play**). Whenever the token arrives at a vertex $v \in V_i$, player i has to move the token to a successor vertex.

What is a **strategy** in such a game?

- **unrestricted**: depends on the whole sequence of vertices seen so far ($\sigma : V^* V_i \rightarrow V$);
- **finite-state**: depends only on bounded information about the sequence of states seen so far (finite transducer);
- **positional**: depends only on the current vertex ($\sigma : V_i \rightarrow V$).

A **strategy profile** is a collection $(\sigma_i)_{i \in \Pi}$ of strategies, one for each player.

A strategy profile $(\sigma_i)_{i \in \Pi}$ determines a unique play of (\mathcal{G}, v_0) , call it $\langle (\sigma_i)_{i \in \Pi} \rangle_{v_0}$.

Nash Equilibria

Question: What is the solution of a game?

In contrast to so-called **two-player zero-sum games**, winning strategies do, in general, not exist even for very simple winning condition like reachability!


Weaker solutions: Admissibility, **Equilibria**.

Definition: A strategy profile $(\sigma_i)_{i \in \Pi}$ is a **Nash equilibrium** of (\mathcal{G}, v_0) if for all players i and all strategies σ' of player i we have: $\langle \sigma', (\sigma_j)_{j \in \Pi \setminus \{i\}} \rangle_{v_0} \in W_i \Rightarrow \langle (\sigma_j)_{j \in \Pi} \rangle_{v_0} \in W_i$.

Remark: For any two-player zero-sum game, a strategy profile (σ_0, σ_1) is a Nash equilibrium if either σ_0 or σ_1 is a winning strategy.

Nash Equilibria: Example

Game with two players 1 (round) and 2 (boxed):



Fact: If both players want to reach vertex 2, no player has a winning strategy.

However, there are two Nash equilibria:

- 1 moves from 0 to 3, and 2 moves from 3 to 2.
- 1 moves from 0 to 1, and 2 stays in 3.

But why should player 2 stay in vertex 3?

Subgame Perfect Equilibria

Problem: Nash equilibria ignore the sequential structure of an infinite game.

Solution: Require that strategies are optimal after every possible play prefix!


Definition: A strategy profile $(\sigma_1, \dots, \sigma_k)$ is a **subgame perfect equilibrium (SPE)** of (\mathcal{G}, v_0) if, for every non-empty play prefix $h\nu$ (starting with v_0), $(\sigma_1|_h, \dots, \sigma_k|_h)$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$ where

- $\mathcal{G}|_h$ is as \mathcal{G} but with winning conditions $W_i|_h = \{\alpha \in C^\omega : \chi(h) \cdot \alpha \in W_i\}$, and
- $\sigma_i|_h : z\nu \mapsto \sigma_i(hz\nu)$ is the strategy σ_i restricted to $\mathcal{G}|_h$.

In particular: Every SPE is a Nash equilibrium (take $h = \varepsilon$).

Subgame Perfect Equilibria: Example

Game with two players 1 (round) and 2 (boxed):



Recall: This game has two Nash equilibria if both players want to reach vertex 2.

However, the game has a unique SPE:

- 1 moves from 0 to 3, and 2 moves from 3 to 2.
- 1 moves from 0 to 1, and 2 stays in 3.

Note that staying in 3 is not optimal for play prefix 03.

Existence of Nash Equilibria

Recall **Martin's Theorem**: Every two-player zero-sum (quasi-)Borel game is **determined**, i.e. one player has a winning strategy.

Theorem (Chatterjee-Jurdziński-Majumdar, CSL '04)

Every game with (quasi-)Borel winning conditions has a Nash equilibrium.

Proof idea: Let each player play an optimal strategy w.r.t. her winning condition. Upon deviation of one player, react with an optimal joint counter-strategy against this player.

Drawback: Assume that player j has deviated after history hv . It may be the case that player i loses the resulting play, although she has a winning strategy in $(\mathcal{G}|_h, v)$. No SPE!

Existence of Subgame Perfect Equilibria

Theorem

Every game with (quasi-)Borel winning conditions has a SPE.

Proof sketch: Consider the unravelling of \mathcal{G} from the initial vertex v_0 (the **game tree**). We prune the tree as follows:

- If some player has a winning strategy from some of her positions, delete all outgoing edges other than the one taken by this strategy (do this in a uniform way!).
- Repeat this construction, until we reach a fixed point.

Let each player play any (fixed) strategy in the pruned game. Upon deviation, react with an optimal counter strategy in the pruned game. The resulting strategy profile is a subgame perfect equilibrium of the original game. □

Winning Conditions

Abstract winning conditions are nice, but to be of algorithmic use, we have to use finitely representable conditions:

- **Reachability:** Given $F_i \subseteq C$, player i wins if F_i is visited.
- **Büchi:** Given $F_i \subseteq C$, player i wins if F_i is visited infinitely often.
- **parity:** Given a **priority function** $\Omega_i : C \rightarrow \omega$, player i wins if the least priority seen infinitely often is even.
- **ω -regular:** Given an MSO (or LTL) formula φ_i using $<$ and predicates P_c for each $c \in C$, player i wins if $(\omega, <, (P_c)_{c \in C}) \models \varphi_i$.

Note: Any game with ω -regular winning conditions can be *reduced* to a game with parity winning conditions.

From now on: Let us only consider games over finite arenas!

Computing Equilibria

Recall that two-player zero-sum parity games are **positionally determined**: One of the two players not only has a winning strategy but a positional one.

This allows us to perform the construction used in the existence proof directly on the game arena.

As positional winning strategies in two-player zero-sum parity games can be computed in exponential time, we get the following theorem.

Theorem

For any multiplayer parity game (played on a finite arena), we can compute a finite-state subgame perfect equilibrium in exponential time.

The Decision Problem SPE

Problem: Subgame perfect equilibria are not unique. Indeed, it can be shown that the algorithm may compute an arbitrary “poor” equilibrium.

Definition: The **payoff** of a strategy profile is the vector $x \in \{0, 1\}^\Pi$ def. by $x(i) = 1$ if player i wins the induced play.

Goal: Find a SPE with an optimal payoff.

The decision problem SPE is defined as follows:

Given a multiplayer parity game (\mathcal{G}, v_0) (played on a finite arena) and thresholds $x, y \in \{0, 1\}^\Pi$, decide whether (\mathcal{G}, v_0) has a SPE with payoff $\geq x$ and $\leq y$.

If we can solve the problem, we can compute an optimal SPE payoff by using “binary search”.

Solving The Problem

Proviso: Choose a suitable encoding of strategy profiles by infinite binary trees.

We use tree automata to solve the problem:

- Construct a **nondeterministic Büchi tree automaton** accepting precisely those trees that are encodings of strategy profiles.
- Construct an **alternating parity tree automaton** accepting those strategy profiles that are SPEs with a desired payoff.
- Make the latter tree automaton nondeterministic and check the product automaton for nonemptiness.

This gives an exponential time algorithm for the problem.

Theorem

The decision problem SPE is decidable in exponential time. For games with a bounded number of players and priorities, the problem is in P.

Open question: What is the exact complexity of the problem?

The proof also reveals that, for any payoff thresholds x, y , the set of SPEs with a payoff $\geq x$ and $\leq y$ is regular.

Theorem

For any multiplayer parity game (\mathcal{G}, v_0) (played on a finite arena) and $x \in \{0, 1\}^\Pi$, if (\mathcal{G}, v_0) has a SPE with payoff x , then (\mathcal{G}, v_0) has a finite-state SPE with payoff x .

Remark: Positional strategies do not suffice!

Conclusion

Summary:

- Every infinite game with (quasi-)Borel winning conditions has a SPE.
- Deciding whether a parity game played on a finite arena has a SPE with a payoff between two given thresholds can be done in exponential time.
- Finite-state strategies suffice to implement SPEs in parity games played on finite arenas.

Future work:

- Generalisation to stochastic and concurrent games (in progress).
- Generalisation to infinite arenas (e.g. pushdown graphs).