Decision Problems for Nash Equilibria in Stochastic Games

Michael Ummels[1] Dominik Wojtczak[2]

[1]RWTH Aachen University

[2]CWI Amsterdam

CSL '09

Stochastic Games

What is a multiplayer stochastic game (SMG)? Let's play!

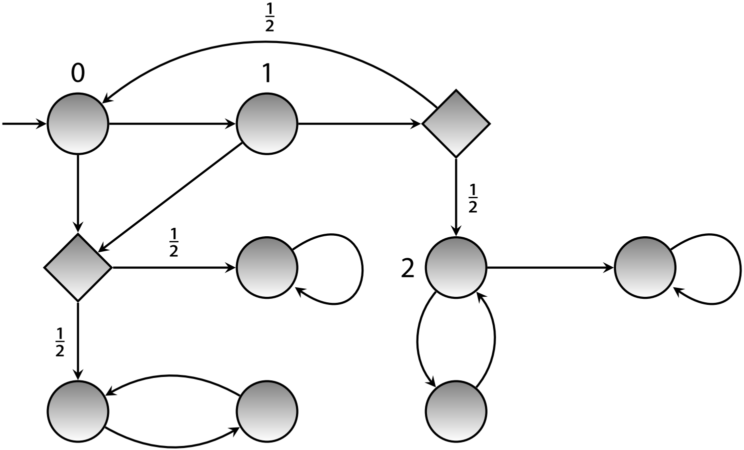

Wining conditions: Büchi, parity, Muller and friends.

Stochastic Games

What is a multiplayer stochastic game (SMG)? Let's play!

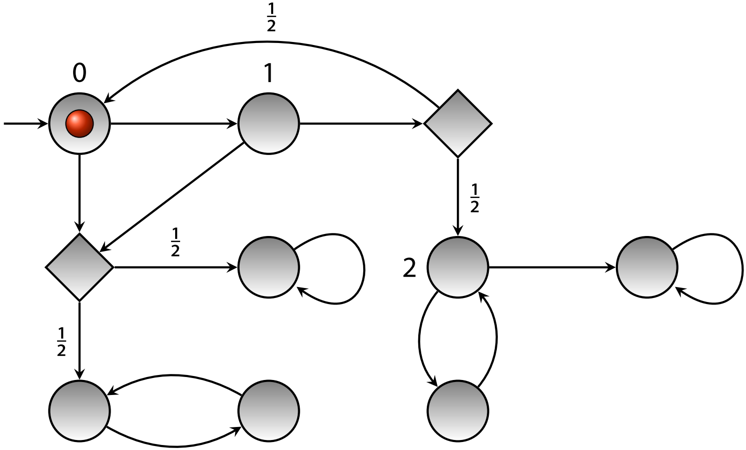

Wining conditions: Büchi, parity, Muller and friends.

What is a multiplayer stochastic game (SMG)? Let's play!

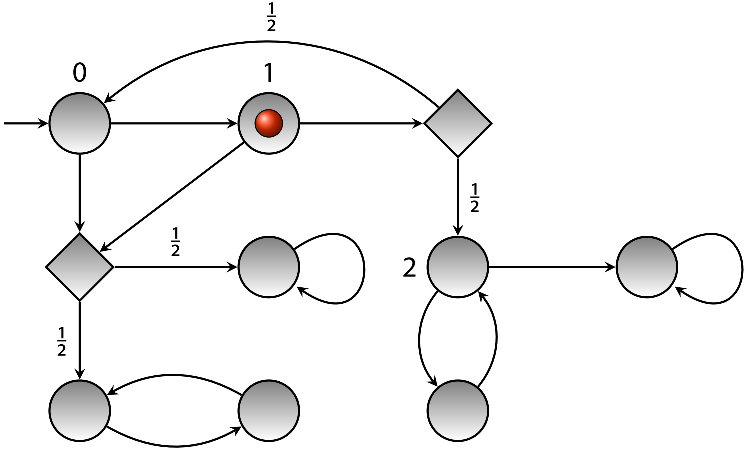

Wining conditions: Büchi, parity, Muller and friends.

Stochastic Games

What is a multiplayer stochastic game (SMG)? Let's play!

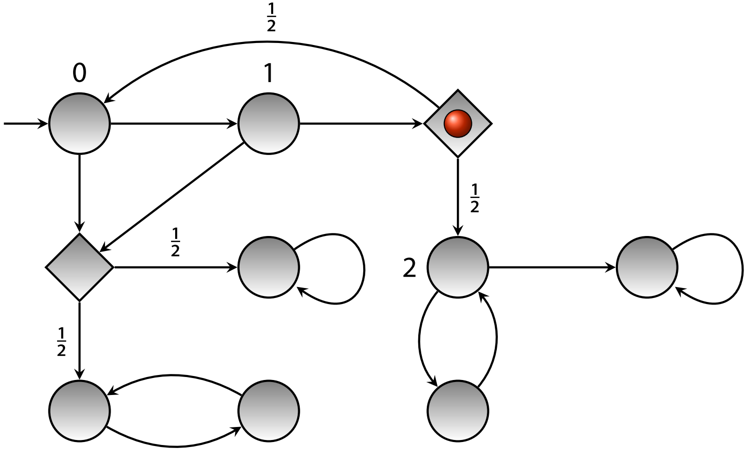

Wining conditions: Büchi, parity, Muller and friends.

Stochastic Games

What is a multiplayer stochastic game (SMG)? Let's play!

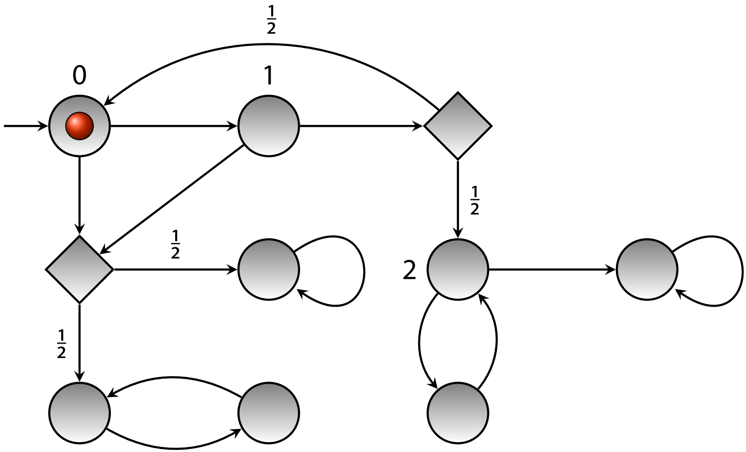

Wining conditions: Büchi, parity, Muller and friends.

Stochastic Games

What is a multiplayer stochastic game (SMG)? Let's play!

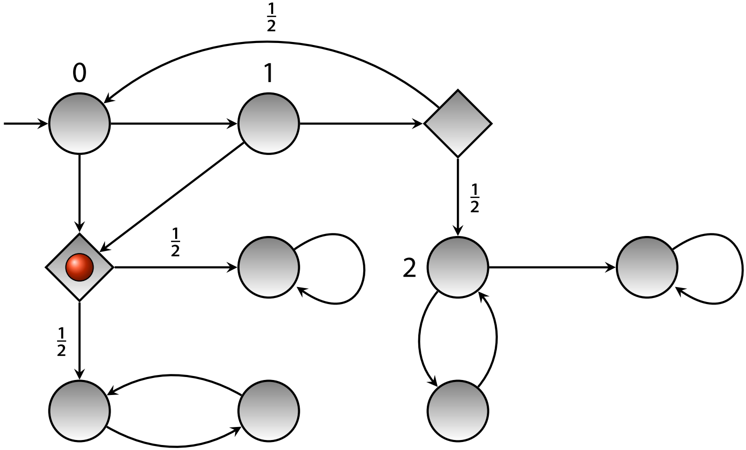

Wining conditions: Büchi, parity, Muller and friends.

Stochastic Games

What is a multiplayer stochastic game (SMG)? Let's play!

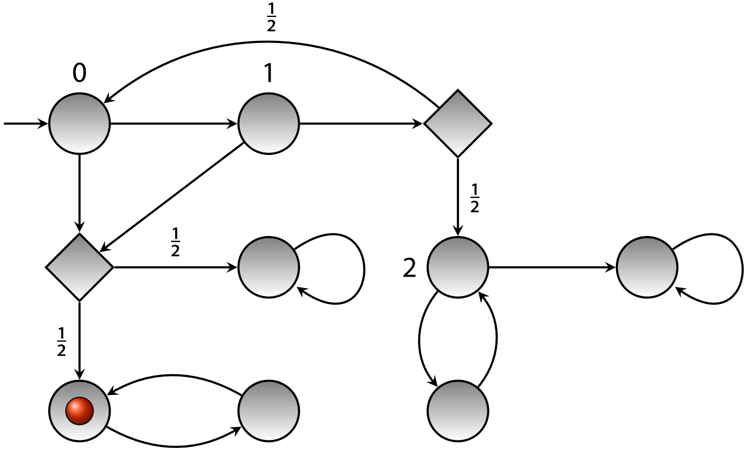

Wining conditions: Büchi, parity, Muller and friends.

Stochastic Games

What is a multiplayer stochastic game (SMG)? Let's play!

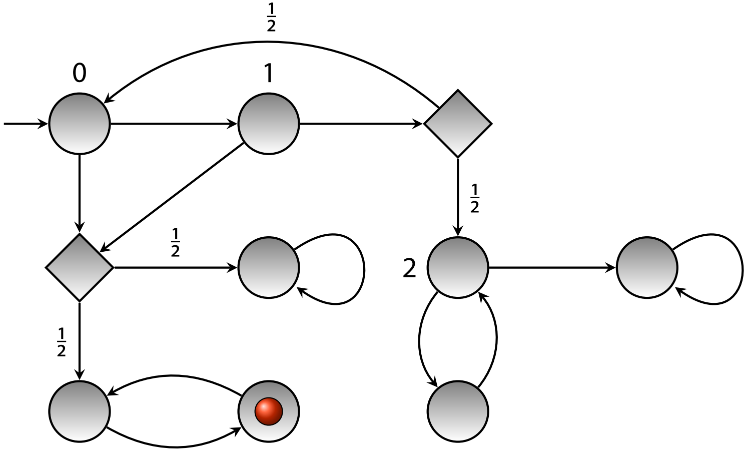

Wining conditions: Büchi, parity, Muller and friends.

Stochastic Games

What is a multiplayer stochastic game (SMG)? Let's play!

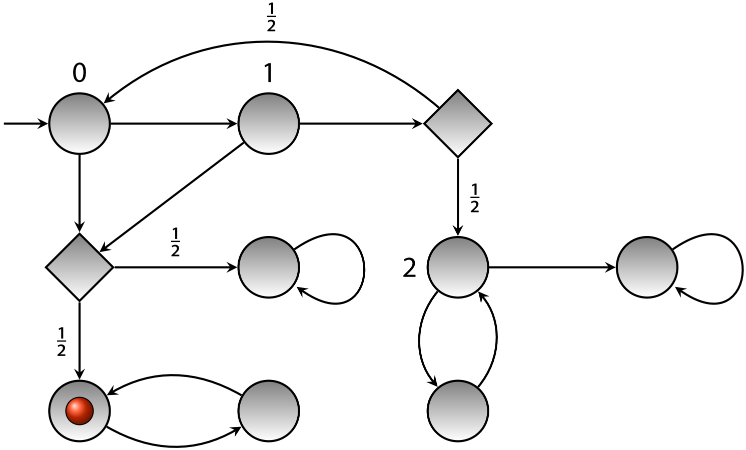

Wining conditions: Büchi, parity, Muller and friends.

Stochastic Games

What is a multiplayer stochastic game (SMG)? Let's play!

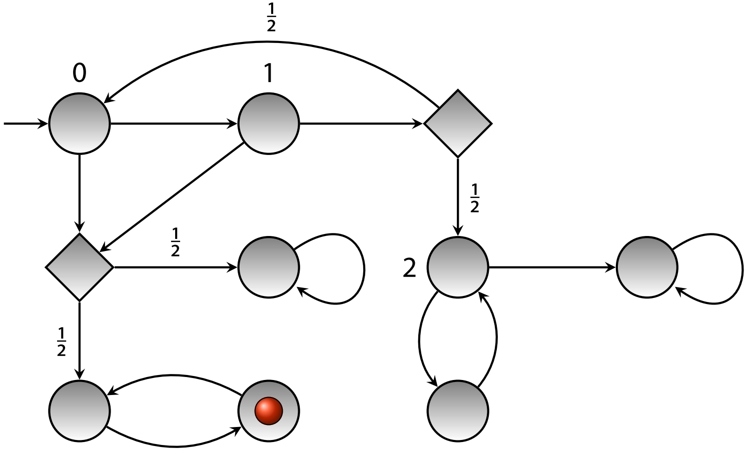

Wining conditions: Büchi, parity, Muller and friends.

Strategies and Probabilities

A strategy profile induces a probability distribution on sets of plays.

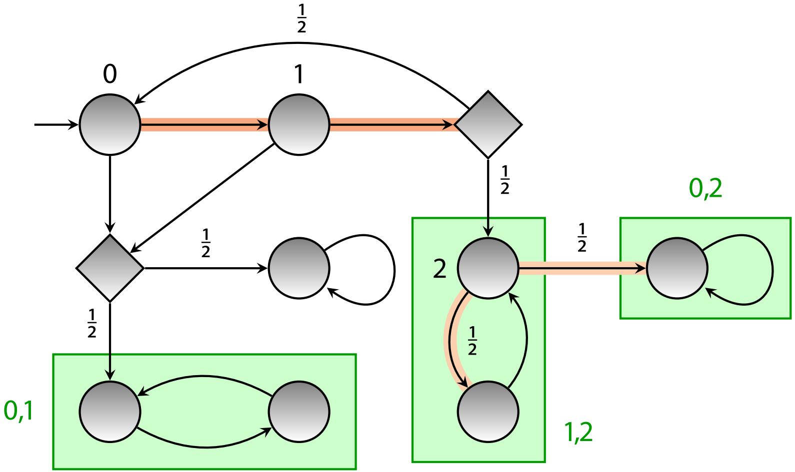

Payoff of this strategy profile: $(1, 0, 1)$

Two-Player Zero-Sum Games

Classical setting: Games played by two players Max and Min with opposing objectives (Two-Player Zero-sum Games).

Theorem (Martin, 1998)

Stochastic two-player zero-sum games with Borel winning conditions are determined, i.e. $\sup_\sigma \inf_\tau \Pr^{\sigma,\tau}(\text{Win}) = \inf_\tau \sup_\sigma \Pr^{\sigma,\tau}(\text{Win}) =: \text{val}(\mathcal{G})$.
Both players have ε-optimal pure strategies.

Theorem (Chatterjee & al., 2003)

In any stochastic two-player zero-sum parity game, both players have *optimal positional* strategies.

Corollary

Deciding the value of a stochastic two-player zero-sum parity game *qualitatively* or *quantitatively* can be done in NP $\cap$ co-NP.

Nash Equilibria

Definition: A strategy profile is a Nash equilibrium if no player can gain from unilaterally switching to a different strategy.

Question: Do Nash equilibria always exist?

Theorem (Chatterjee & al., 2004)

Any SMG with ω-regular winning conditions a Nash equilibrium in pure finite-state strategies.

Next Question: Can we compute one?

Proposition

The problem of computing a Nash equilibrium of a parity SMG is in FNP.

Open Problem: Can one find a Nash equilibrium in polynomial time?

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

Example

Nash Equilibrium where Player 2 wins almost surely:

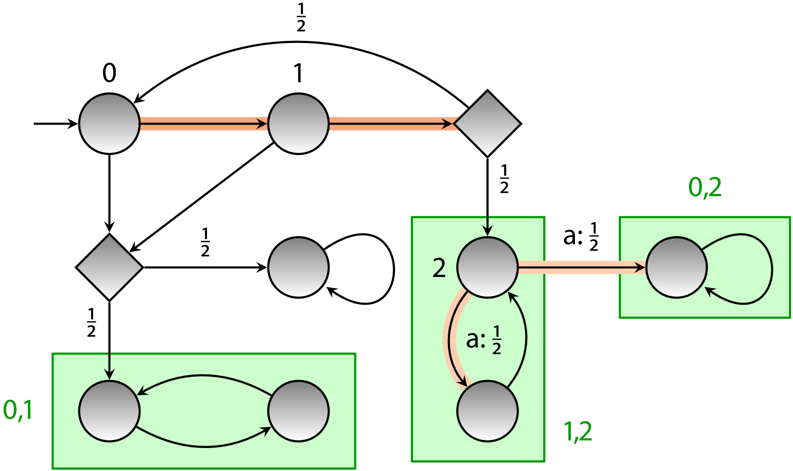

Observation: Memory and randomisation is useful.

Example

Nash Equilibrium where Player 2 wins almost surely:

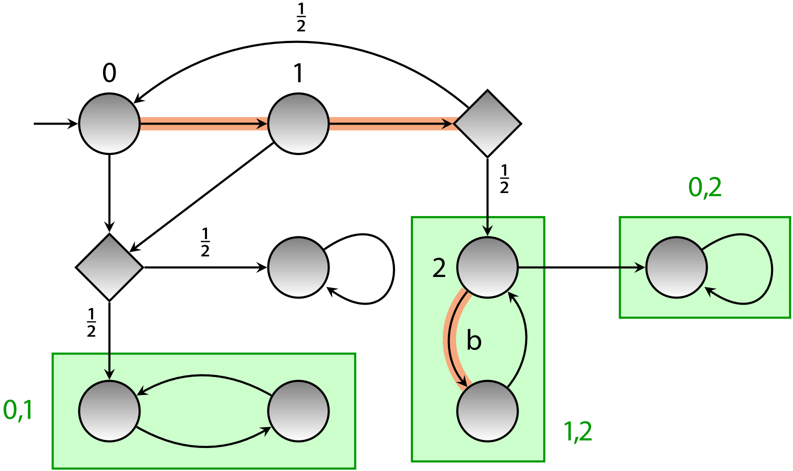

Observation: Memory and randomisation is useful.

The Problem NE

Goal: Compute a Nash equilibrium that meets certain requirements on the payoff.

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

Note: This is a generalisation of the quantitative decision problem for stochastic two-player zero-sum games.

Variants of the Problem:

- Arbitrary strategies
- Pure strategies
- Stationary strategies
- Positional strategies

Results

What is known about the complexity of NE and its variants?

- NP-complete for *simple* stochastic games wrt. positional strategies.
- NP-hard and contained in PSPACE for *simple* stochastic games wrt. stationary strategies.
- Undecidable wrt. pure strategies, even if $\overline{x} = \overline{y}$ *or* if $\overline{x}$ and $\overline{y}$ are binary.
- Complexity of NE wrt. arbitrary strategies wide open.

In this paper:

- Decidability of NE wrt. positional/stationary strategies.
- Decidability of NE wrt. arbitrary/pure strategies if $\overline{x} = \overline{y}$ is binary.

Results hold for SMGs with parity/Muller objectives.

Positional Nash equilibria

Theorem

NE wrt. positional strategies is NP-complete for parity or Muller SMGs.

The algorithm:

- Guess a positional strategy profile $\overline{\sigma}$.
- For each player i:
  1. Compute the payoff r_i player i receives with $\overline{\sigma}$.
  2. Compute the maximal payoff z_i player i can achieve if all other players stick to $\overline{\sigma}$.
  3. Check whether $x_i \leq r_i \leq y_i$ and $z_i \leq r_i$.

1. and 2. are doable in polynomial time (via linear programming).

Remark: NP-hardness holds even for games with only two players.

Stationary Nash Equilibria

Theorem

NE wrt. stationary strategies is in PSPACE for parity or Muller SMGs.

The algorithm:

- Guess the support S of a stationary strategy profile $\overline{\sigma}$.
- For each player i, compute the set R_i of vertices from where player i wins with positive probability when playing $\overline{\sigma}$.
- For each player i, compute the set T_i of vertices from where player i has a strategy win with probability 1 if all other players stick to $\overline{\sigma}$.
- Evaluate an existential first-order sentence ψ (which is polynomial-time computable from $\mathcal{G}, \overline{x}, \overline{y}, S, \overline{R}$ and $\overline{T}$) over $\mathfrak{R} = (\mathbb{R}, +, \cdot, 0, 1)$.

ψ states that there exists a stationary Nash equilibrium $\overline{\sigma}$ with payoff $\geq \overline{x}$ and $\leq \overline{y}$ whose support is precisely S.

The Strict Qualitative Fragment

Goal: Deciding the existence of a Nash equilibrium with payoff $\overline{x} \in \{0,1\}^k$.

Denote by P_i the set of vertices from where player i has a strategy to win with probability > 0.

Observation: $\overline{\sigma}$ can only be a Nash equilibrium of $\mathcal{G}$ with payoff $\overline{x} \in \{0,1\}^k$ if $\Pr^{\overline{\sigma}}(\mathrm{Reach}(P_i)) = 0$ for each player i with $x_i = 0$.

In fact, this condition is sufficient for having a pure equilibrium with payoff $\overline{x}$.

Lemma

If $\overline{\sigma}$ is a *pure* strategy profile with payoff $\overline{x} \in \{0,1\}^k$ such that $\Pr^{\overline{\sigma}}(\mathrm{Reach}(P_i)) = 0$ for each player i with $x_i = 0$, then there exists a pure Nash equilibrium $\overline{\sigma}^*$ such that $\Pr^{\overline{\sigma}} = \Pr^{\overline{\sigma}^*}$.

Proof idea: Threat strategies.

The Strict Qualitative Fragment

We can now characterize the existence of a Nash equilibrium with payoff $\overline{x}$.

Proposition

Let $\mathcal{G}$ be a parity/Muller SMG, and $\overline{x} \in \{0,1\}^k$. Then the following statements are equivalent:

1. There exists a Nash equilibrium with payoff $\overline{x}$;
2. There exists a strategy profile $\overline{\sigma}$ with payoff $\overline{x}$ such that $\Pr^{\overline{\sigma}}(\text{Reach}(P_i)) = 0$ for every player i with $x_i = 0$;
3. There exists a pure strategy profile $\overline{\sigma}$ with payoff $\overline{x}$ such that $\Pr^{\overline{\sigma}}(\text{Reach}(P_i)) = 0$ for every player i with $x_i = 0$;
4. There exists a pure Nash equilibrium with payoff $\overline{x}$.

Corollary: Randomisation does not help.

But deciding 2. is essentially an MDP problem...

The Strict Qualitative Fragment

The (high-level) algorithm:

- Compute for each player i with $x_i = 0$ the set P_i of vertices from where player i has a strategy to win with probability > 0.
- Construct the MDP $\mathcal{M}$ that arises from $\mathcal{G}$ by merging all players into one, removing all sets P_i and imposing the winning condition $\bigwedge_{i:x_i=1} \mathrm{Win}_i \wedge \bigwedge_{i:x_i=0} \neg\mathrm{Win}_i$.
- Check whether $\mathrm{val}(\mathcal{M}) = 1$.

To compute the sets P_i, we have to solve the qualitative value problem.

Theorem

Deciding whether there exists a Nash equilibrium with payoff $\overline{x} \in \{0,1\}^k$ can be done in P for Büchi SMGs, in NP ∩ co-NP for parity SMGs and PSPACE-complete for Muller SMGs.

Conclusion

Open Problems:

- Is NE wrt. arbitrary strategies decidable?
- Is NE decidable for two-player games?
- Do our results carry over to infinite-state games?
- …