

Stochastic Reachability Games on Countably Infinite Graphs

Stefan Kiefer, Richard Mayr, Mahsa Shirmohammadi,
Patrick Totzke

University of Oxford, Edinburgh, Paris, Liverpool

Workshop on Current Trends in Graph and Stochastic Games,
7-8 April 2022, in Maastricht, NL.

2-Player Perfect-information Stochastic Games on Countably Infinite Graphs

Concurrent games.

- Both players choose an action from their available action sets.
- For each chosen pair of actions (a, b) , we have a pre-defined distribution over successor states.
- This distribution can have infinite support.
Infinite branching even if the action sets are finite.

Turn-based games (aka switching control).

- Each state is owned by some player. Maximizer, Minimizer or Random.
- The player who owns the state gets to choose among the available successors.
(For a random state, a pre-defined distribution.)

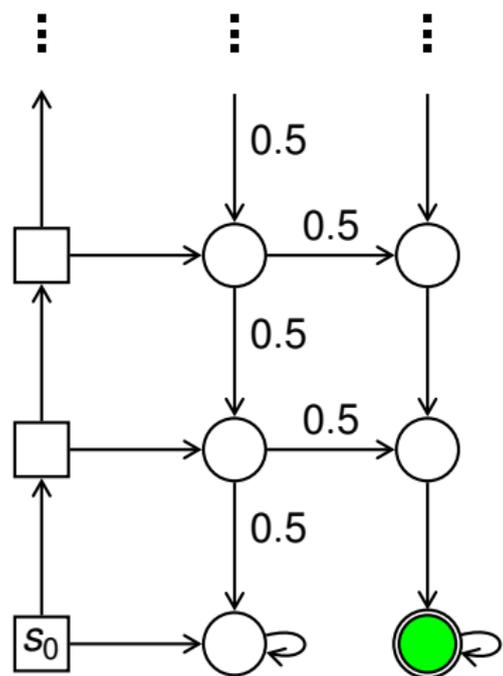
Reachability Objective

Arguably the simplest objective in graph games.

- Start state s_0 .
- Set of target states T .
- Plays from s_0 that reach T are winning for Maximizer.
- Maximizer (resp. Minimizer) wants to maximize (resp. minimize) the probability of reaching T .
- Reachability objective is an open set, thus trivially Borel.
- The **Safety Objective** (avoiding T) is dual to Reachability. (The same, just from Min's point of view.)

No Optimal Max Strategies

Optimal Maximizer strategies might not exist, even in MDPs (Minimizer is passive), even if finitely branching.



$$val_s(E) \stackrel{\text{def}}{=} \sup_{\sigma} \mathcal{P}_s^{\sigma}(E)$$

$$val_s(\text{Reach } \textcircled{\textcircled{\color{red}\bullet}}) = 1$$

Does an ε -optimal Maximizer Strategy need Memory?

Theorem (Kučera 2011)

In finitely-branching turn-based reachability games, Maximizer has ε -optimal MD strategies.

- Reaching the target is approximated by reaching the target in n rounds, for increasing n .
- By finite-branching, only finitely many states can be reached in n rounds.
- ε -close approximation for sufficiently large n .
- Consider the induced finite-state subgame, restricted to states reachable in n steps.
- In finite-state reachability games, Maximizer has ε -optimal MD strategies. (In fact, even optimal MD ones.)

Does an ε -optimal Maximizer Strategy need Memory?

Theorem (Secchi 1997)

*In concurrent reachability games with finite action sets
Maximizer has ε -optimal positional (MR) strategies.*

- Similar ideas as above.
- However, infinitely many states can be reached in n rounds due to the infinite support of the distributions.
- Still, with high probability one stays in a finite subset during the first n rounds.
- Randomization is needed for concurrent games (e.g., Rock-Paper-Scissors).

Infinitely Branching Reachability Games

What about general reachability games, without the restrictions?

Does Maximizer have ε -optimal positional strategies?

- Concurrent reachability games where Minimizer has **infinite action sets**.
(Maximizer still has finite action sets, or else the game would not be determined.)
- Turn-based reachability games with **infinite branching**.
(Only Minimizer's infinite branching is important, since Maximizer's infinite branching can be encoded into finite branching.)

Infinitely Branching Reachability Games

What if one makes the game easier for Maximizer by imposing extra conditions?

- Assume that the game graph is **acyclic**.
- Even stronger, assume that we have a **step-counter encoded** into the states.
- Assume that all states in the game are **almost-sure winning**.

I.e., from every state s there exists a Maximizer strategy that reaches the target T with probability 1, regardless what Minimizer does.

For many objectives, on MDPs and games, such restrictions make things easier for Maximizer.

Infinitely Branching Reachability Games are hard

Finite memory (even if private) plus a step-counter are still useless for Maximizer.

Theorem (Kiefer, Mayr, Shirmohammadi, Totzke 2022)

There is a turn-based reachability game, infinitely branching for Minimizer, with the following properties:

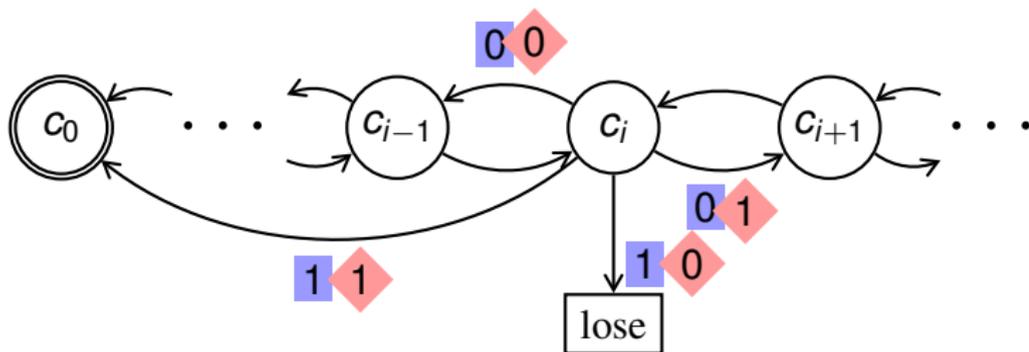
- 1 *for every Maximizer state, Maximizer has at most two successors to choose from;*
- 2 *for every state, Maximizer has a strategy to visit the target state with probability 1, regardless of Minimizer's strategy;*
- 3 *for every Maximizer strategy that uses only a step counter and private finite memory and randomization, for every $\varepsilon > 0$, Minimizer has a strategy so that the target is visited with probability $\leq \varepsilon$.*

This immediately carries over to concurrent games with infinite Minimizer action sets.

Concurrent Big Match on \mathbb{N}

		Min	
		0	1
Max	0	c_{i-1}	c_{i+1}
	1	lose	c_0

actions at c_i ($i \geq 1$)



Concurrent Big Match on \mathbb{N}

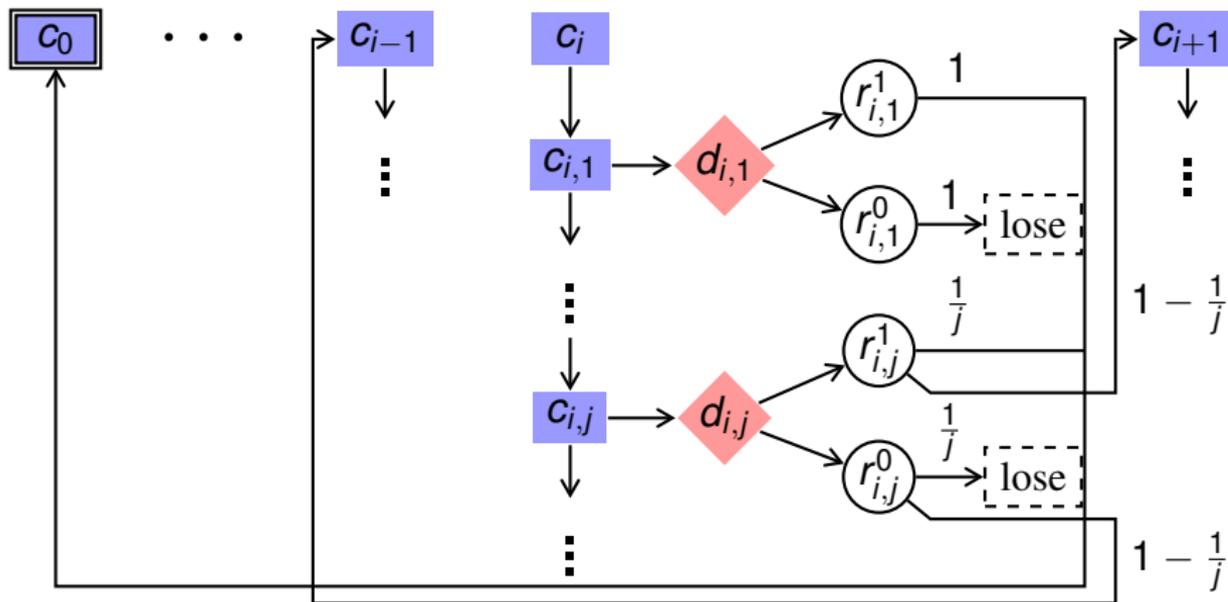
ϵ -optimal positional Maximizer strategies, but **no uniform ones**.

Theorem (Nowak-Raghawan'91, Fristedt-Lapic-Sudderth'95)

- 1 $val_G(c_x) = (x + 2)/(2x + 2) \geq 1/2$.
- 2 *For every start state c_x and $N \geq 0$, Maximizer can win with probability $\geq N/(2N + 2)$ by choosing action 1 with probability $1/(n + 1)^2$ whenever the current state is c_i with $i = x + N - n$ for some $n \geq 0$.*
- 3 *Every MR Maximizer strategy σ attains arbitrarily little from c_x as $x \rightarrow \infty$. Formally,*

$$\limsup_{x \rightarrow \infty} \inf_{\pi} \mathcal{P}_{G, c_x, \sigma, \pi}(\text{Reach } \{c_0\}) = 0$$

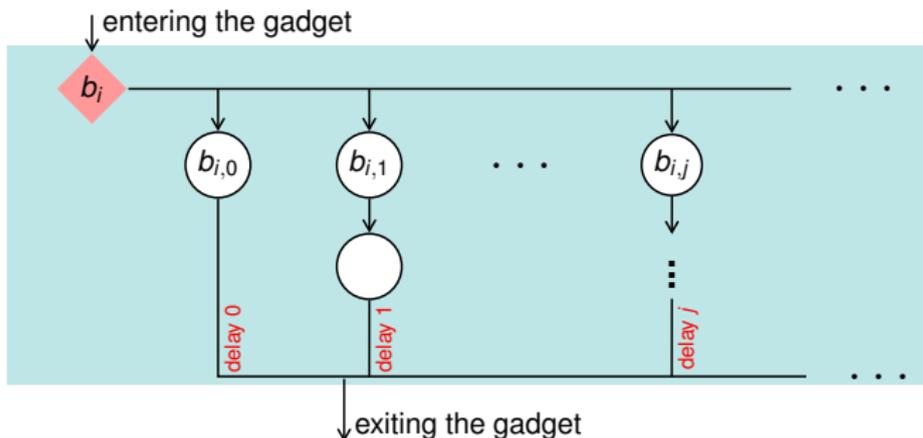
Turn-Based Big Match on \mathbb{N}



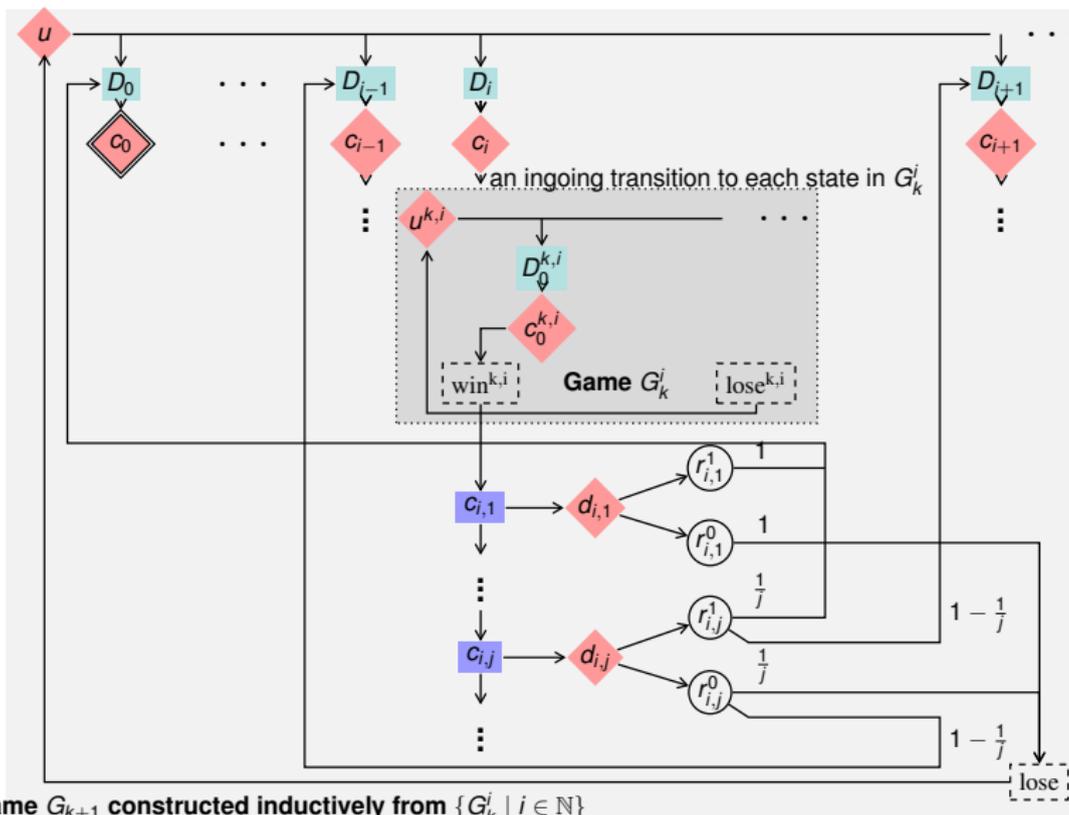
Results on values and Maximizer strategies carry over.

Delay Gadget of Minimizer

Minimizer can make arbitrary, but finite, delays.



k-Nested Turn-Based Big Match on \mathbb{N} with Restarts



k-Nested Turn-Based Big Match on \mathbb{N} with Restarts

- Let \mathcal{G}_1 be the Turn-Based Big Match on \mathbb{N} with restarts.
- Construct \mathcal{G}_k as the k-nested Turn-Based Big Match on \mathbb{N} with restarts.
Uses \mathcal{G}_{k-1} as building blocks.
- Prove that every Max strategy with k memory modes (in addition to the step-counter) is useless in \mathcal{G}_k .
By induction on k .
- Proof is very technical, because one cannot assume that Max neatly separates his memory between \mathcal{G}_k and the \mathcal{G}_{k-1} subgames.
- Finally, combine all the games \mathcal{G}_k into a single game where Minimizer gets to choose k . (Infinitely branching again.)
If Maximizer's strategy has some k memory modes then Minimizer goes to \mathcal{G}_k and wins there.

Uniform Strategies

- In MDPs there exist **uniform** ε -optimal MD strategies for reachability. [Ornstein'69]
- This does not carry over to concurrent games with finite action sets. ε -optimal strategies can be positional [Secchi'97], but not uniform (Concurrent Big Match on \mathbb{N}).
- The same holds for finitely branching turn-based games. (Turn-based Big Match on \mathbb{N}).

Uniform Strategies

- In MDPs there exist **uniform** ε -optimal MD strategies for reachability. [Ornstein'69]
- This does not carry over to concurrent games with finite action sets. ε -optimal strategies can be positional [Secchi'97], but not uniform (Concurrent Big Match on \mathbb{N}).
- The same holds for finitely branching turn-based games. (Turn-based Big Match on \mathbb{N}).

Theorem (Kiefer, Mayr, Shirmohammadi, Totzke 2022)

*For any concurrent game with finite action sets and reachability objective, for any $\varepsilon > 0$, Maximizer has a uniformly ε -optimal public **1-bit** strategy.*

If the game is turn-based and finitely branching, Maximizer has a deterministic such strategy.

Infinitely branching games:

- Reachability games with infinite Minimizer action sets are hard. Even in the special case of turn-based games.
- In such games, good Maximizer strategies for reachability need **infinite memory** (beyond a step counter).

Infinitely branching games:

- Reachability games with infinite Minimizer action sets are hard. Even in the special case of turn-based games.
- In such games, good Maximizer strategies for reachability need **infinite memory** (beyond a step counter).

Uniformity:

- In concurrent reachability games with finite action sets there are ε -optimal positional Max strategies, but not uniform ones.
- Can be made **uniform with 1 bit** of public memory.

Infinitely branching games:

- Reachability games with infinite Minimizer action sets are hard. Even in the special case of turn-based games.
- In such games, good Maximizer strategies for reachability need **infinite memory** (beyond a step counter).

Uniformity:

- In concurrent reachability games with finite action sets there are ε -optimal positional Max strategies, but not uniform ones.
- Can be made **uniform with 1 bit** of public memory.

Other problems:

- How much memory do optimal strategies need, if they exist?
- How much memory do good Minimizer strategies need?
- \implies Strategy Complexity of Reachability in Countable Stochastic 2-Player Games. **arXiv:2203.12024**