The Complexity of Games with Randomised Control

Sarvin Bahmani

University of Liverpool

Joint with R Ibsen-Jensen, S Paul, S Schewe, F Slivovsky, Q Tang, D Wojtczak, S Zhu.

Games on Graph

- **Game Arena**

  Directed **finite** graph G = (V, E)

- **Game Setting**

  Two-player **infinite-duration** games.

- **Partition**

  V = Max ∪ Min

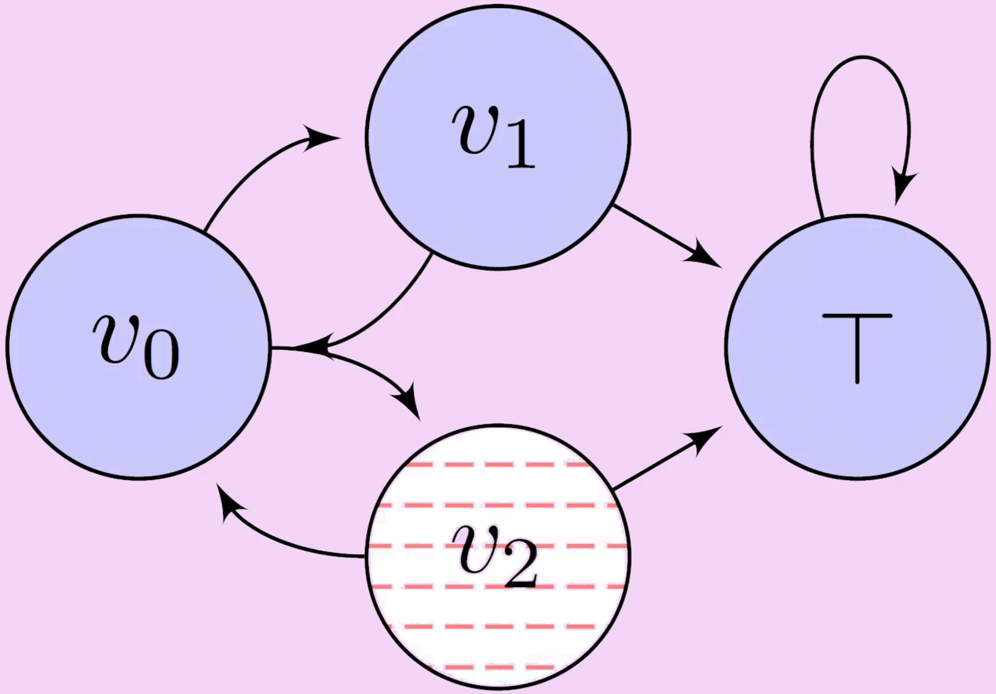

Game Objectives

Reachability

Eventually **reach** a designated **target** set

Parity

Max **priority** seen **infinitely** often is even

Energy

Accumulated weight **never** drops below 0

All three are **memoryless determined**

Reachability games are **P-Complete**

Parity and **Energy** games are in **UP ∩ coUP**

Randomness under Control

Classic Game

The owner of each node is distinguished and **fixed**.

Randomised Control Games

What if **ownership** is randomly decided?

How does it effect the complexity if the ownership **reassigns in every visit** versus **fixed permanently**?

Does it matter **when** to decide the **permanent ownership**?

How to decide the owner of each node?

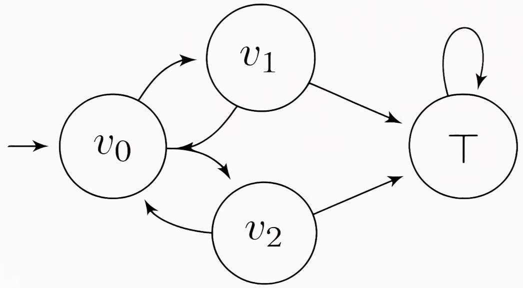

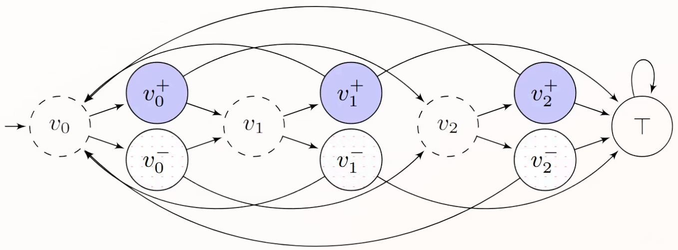

Toss a coin
every time visiting a node
(Random turn)

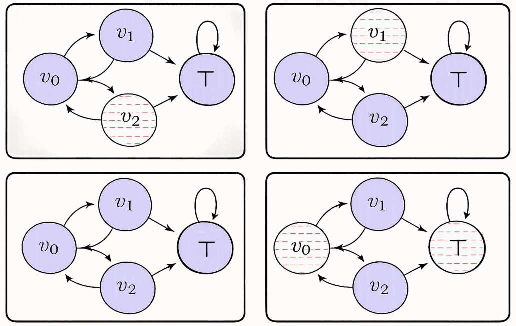

Deciding the ownership of all nodes by coin-toss before starting the game
(Random arena)

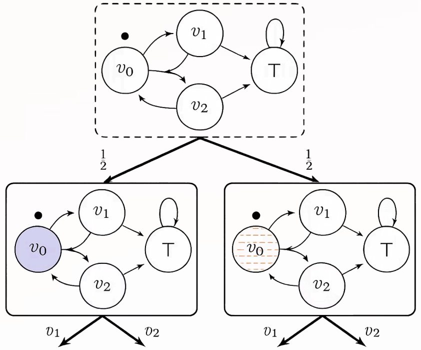

Owner decided the 1st time token reaches a node;
Keep decision permanently
(Unfolding game)

Three Models of Randomised Control

	Model	When is ownership decided?
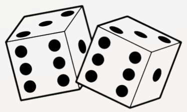	Random turn	Every visit
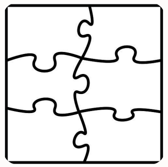	Random arena	Before game
	Unfolding random turn	First visit

Complexity Landscape ?

	Model	Qualitative	Quantitative
	Random turn	?	UP $\cap$ coUP
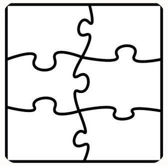	Random arena	?	?
	Unfolding random turn	?	?

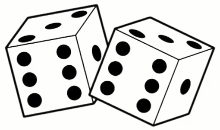

Qualitative Random Turn Game

Infinite plays → eventually repeat

- Any play = prefix + cycle (lasso)
- Max **loses** ⇔ there exists a “**bad**” lasso
- Problem reduces to checking **existence of bad lasso**

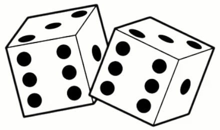

Qualitative
Random Turn Game

Random turn game: **bad lasso path** ⇔ Losing

NL-Membership

Non-deterministically guess a bad lasso path in the game

NL-Hardness

Reduce from st-non-connectivity alternating game

Random-arena game:

bad lasso path ⇔ Losing

Unfolding-random turn:

bad lasso path ⇔ Losing

Qualitative :

NL-Complete

Complexity Landscape

	Model	Qualitative	Quantitative
	Random turn	**NL-Complete**	$\mathrm{UP} \cap \mathrm{coUP}$
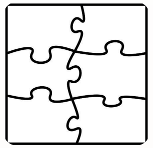	Random arena	**NL-Complete**	**?**
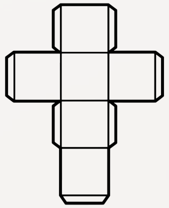	Unfolding random turn	**NL-Complete**	**?**

Two-player reachability game (Classic game) → P-Complete

Random Arena Game

Quantitative: Exact winning Probability

Game Arena

Ownership is decided **once**, **before** game

Number of possible arenas: $2^{|V|}$

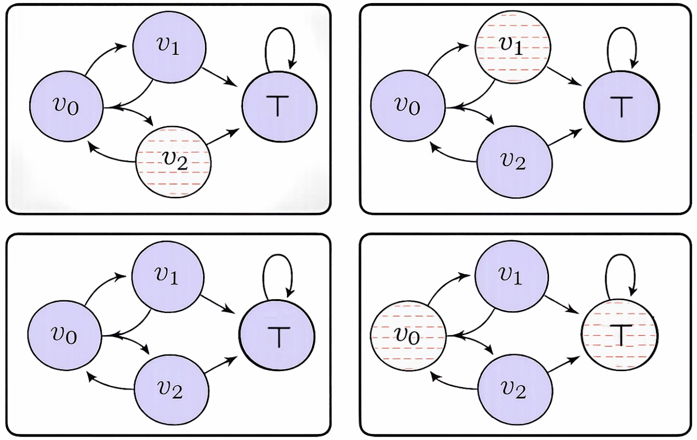

Winning Probability

$$\Pr[\text{Max wins}] = \frac{\#\text{winning arenas}}{\#\text{all arenas}}$$

$$\#\text{all arenas} = 2^{|V|}$$

We must **count** winning arenas →

Counting problem →

$\#P$

$\#\{$winning arenas$\}$ →

For each arena solvable in **polynomial time**

Why Membership in #P?

Per Ownership Assignment

- Reachability → solve in **polynomial time**
- Parity / Energy → solvable in **UP ∩ coUP**

Counting Argument

Each **winning arena** contributes exactly one accepting branch in the nondeterministic computation.

Random Arena Game

Quantitative: #P-Hardness

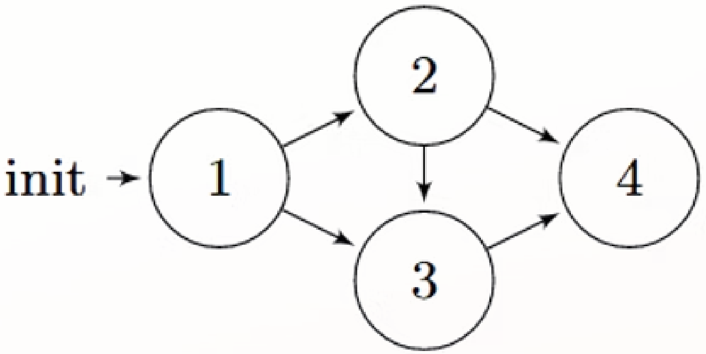

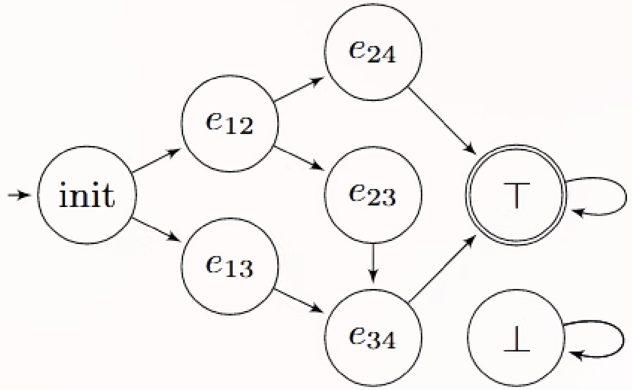

Source-Sink Reliability	Random Arena
Edge exists (or not)	Node owned by Max (or Min)
Subgraph chosen	Arena chosen
s→t path exists	Max has winning strategy

Complexity Landscape

	Model	Qualitative	Quantitative
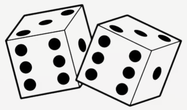	**Random turn**	NL-Complete	$\mathrm{UP} \cap \mathrm{coUP}$
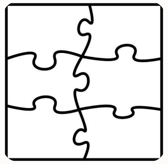	**Random arena**	NL-Complete	#P-complete
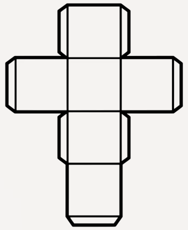	**Unfolding random turn**	NL-Complete	?

Unfolding random-turn Game

Quantitative: Is Pr[Max wins] ≥ θ ?

Game Structure

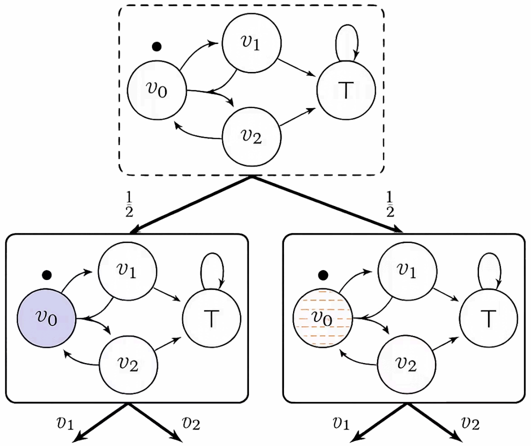

Unfolding at each step depends on:

1. current position of token
2. which nodes have already been assigned
3. to whom the nodes were assigned

$\implies$ Unfolding Game is **Exponentially** big

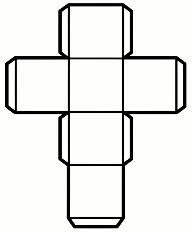

Unfolding Game

Quantitative: Is Pr[Max wins] ≥ θ ?

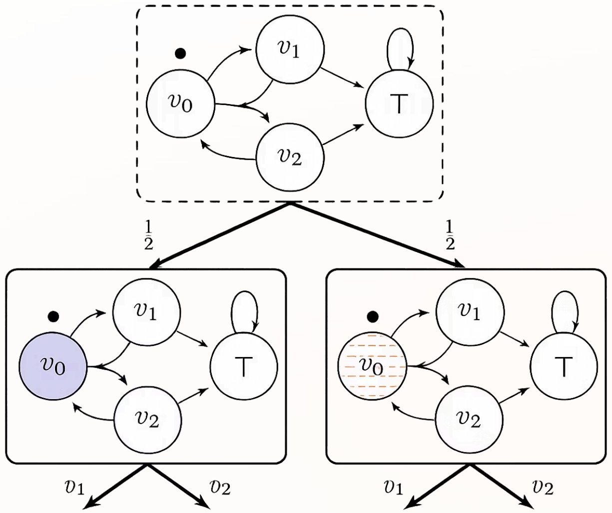

Structural properties:

1. Ownership never reverts.
2. Number of unassigned nodes strictly decreases.
3. These gives us a **layered DAG structure.**

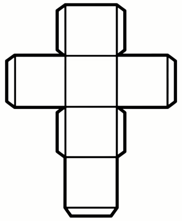

Unfolding Game

Quantitative: Is Pr[Max wins] $\geq \theta$?

PSPACE-Membership

$$\text{Structural properties} \implies \begin{cases} \text{Exponential Tree} \\ \text{Polynomial Depth} \end{cases}$$

$$\left.\begin{array}{r} \text{Exponential Tree} \\ \text{Polynomial Depth} \end{array}\right\} \implies \text{APTIME}$$

PSPACE = APTIME

$\implies$

PSPACE-Membership

Unfolding Game

Quantitative: Is Pr[Max wins] ≥ θ ?

PSPACE-Hardness

$$\Phi = \forall x_1 \exists y_1 \forall x_2 \exists y_2\ \varphi$$

$$\varphi = \underbrace{(x_1 \vee y_1 \vee \neg y_2)}_{C_1} \wedge \underbrace{(\neg x_2 \vee y_1)}_{C_2} \wedge \underbrace{(x_2 \vee y_2)}_{C_3}$$

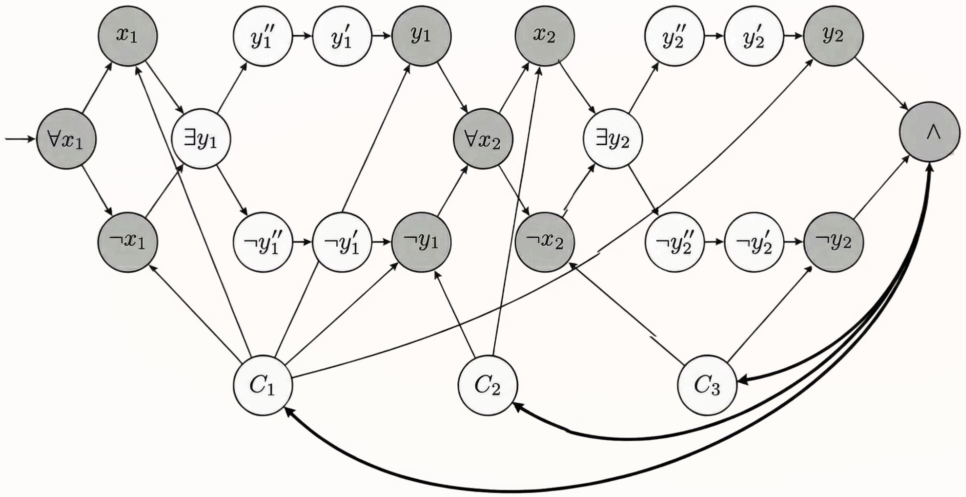

Gap Construction:

1. If formula is true → Pr[Max wins] ≥ P + ½ Q
2. If formula is false → Pr[Max wins] ≤ P + ¼ Q

Unfolding Game
Quantitative: Is Pr[Max wins] ≥ θ ?
PSPACE-Hardness

$$\varphi = \underbrace{(x_1 \vee y_1)}_{C_1} \wedge \underbrace{(x_2 \vee y_2)}_{C_2} \wedge \underbrace{(\neg y_1 \vee \neg y_2)}_{C_3}.$$

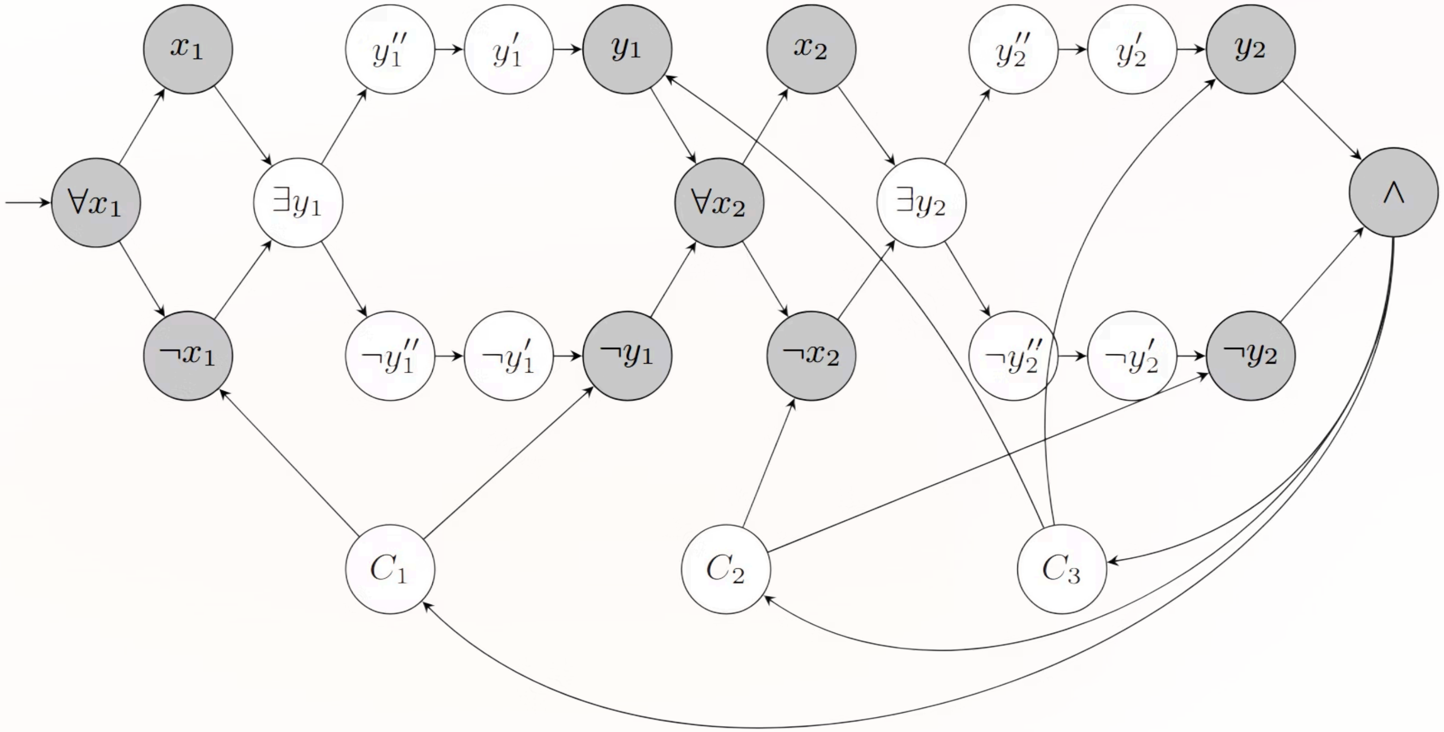

Counting vs Evaluation

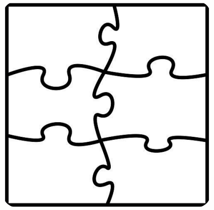

Random **arena** game

Count the number of **winning deterministic arenas** across $2^{|V|}$ assignments.

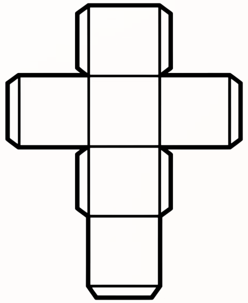

Unfolding random turn game

Evaluate the value of **one stochastic game** symbolically

Complexity Landscape

	Model	Qualitative	Quantitative
	Random turn	NL-complete	UP $\cap$ coUP
	Random arena	NL-complete	#P-complete
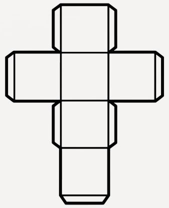	Unfolding random turn	NL-complete	PSPACE-complete

Thank you!

Approximation: FPRAAS

Exact counting is **#P-hard**. Instead, use a randomised approximation scheme:

01

Sample Random Ownership

Assign each node to Max or Min independently at random.

02

Solve Deterministic Game

Solve the resulting two-player game in polynomial time.

03

Repeat n Times

Estimate frequency of Max winning across all samples.

Hoeffding Inequality

$\Pr(|\hat{p} - p| \geq \varepsilon) \leq 2e{-}2\varepsilon^{2}n$

Set $\mathbf{n = O((1/\varepsilon^2)\log(1/\delta))}$ to guarantee:

$|\hat{p} - p| \leq \varepsilon$ with probability $\geq 1 - \delta$

This gives an **additive (ε, δ)-approximation**.