The Many Faces of Strategy Complexity

James C. A. Main
Thesis realised under the supervision of Mickael Randour

F.R.S.-FNRS and UMONS – Université de Mons, Belgium

Private PhD Defence – September 22, 2025

Synthesis via games

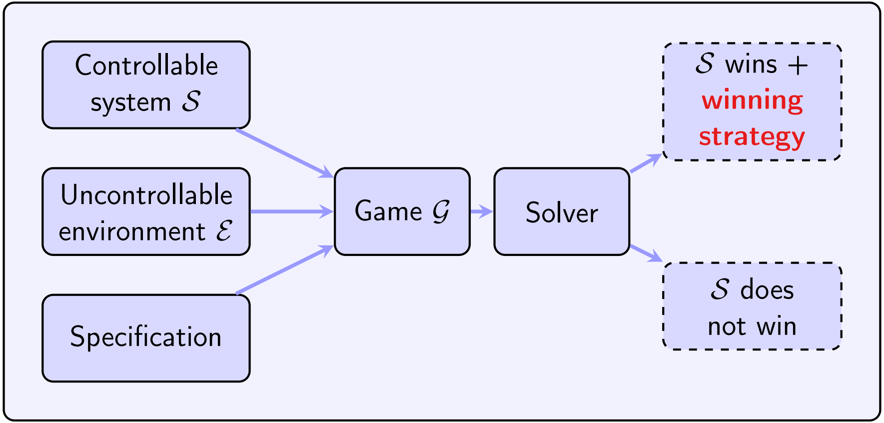

A **strategy** is a formal blueprint for a **controller** of a reactive system.

The simpler, the better

In general, we want **simple strategies**.

Main question

What makes a strategy complex ?

Strategy complexity is **multifaceted**.

Memory

A classical complexity measure

Randomisation

Expressiveness and requirements

Representations

Concise counter-based strategies

Table of contents

Table of contents

Arenas

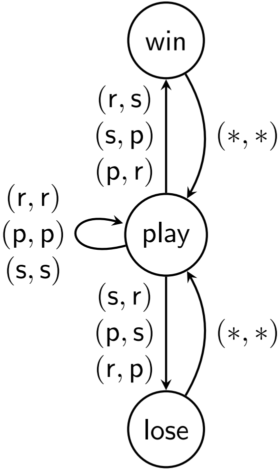

Model: **concurrent game on a graph**

Two-player arena

- Countable **state space** S;
- Countable **action spaces** $A^{(1)}$, $A^{(2)}$;
- **Transition** function $\delta\colon S \times \bar{A} \to \mathcal{D}(S)$.

Play: sequence in $(S\bar{A})^\omega$ coherent with δ.
History: prefix of a play ending in a state.

Objective: measurable subset of $\mathsf{Plays}(\mathcal{A})$.
Payoff: measurable function $f\colon \mathsf{Plays}(\mathcal{A}) \to \bar{\mathbb{R}}$.

Markov decision processes

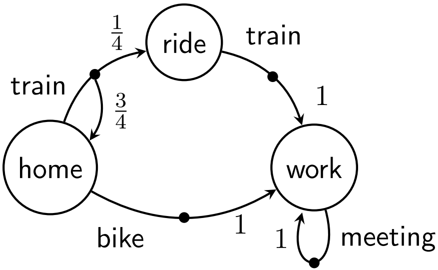

Markov decision process = one-player arena

Markov decision process (MDP)

- Countable **state space** S;
- Countable **action spaces** A;
- **Transition** function $\delta \colon S \times A \to \mathcal{D}(S)$.

Play: sequence in $(SA)^\omega$ coherent with δ.
History: prefix of a play ending in a state.

Objective: measurable subset of Plays$(\mathcal{M})$.
Payoff: measurable function $f \colon \text{Plays}(\mathcal{M}) \to \bar{\mathbb{R}}$.

Strategies

Non-determinism in games is resolved through **strategies**.

Pure strategies

A **pure strategy** is a function $\sigma_i \colon \mathsf{Hist}(\mathcal{A}) \to A^{(i)}$.

A **memoryless strategy** only looks at the current state.

When fixing a **strategy profile** σ and an initial state s, we obtain a **Markov chain** over histories.

- **Probability notation**: $\mathbb{P}^{\sigma}_{s}$.
- **Expectation notation**: $\mathbb{E}^{\sigma}_{s}$.

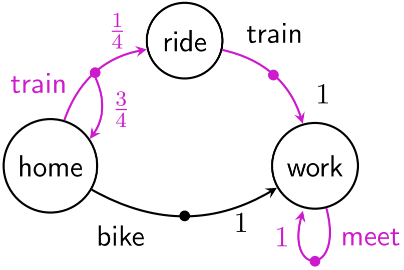

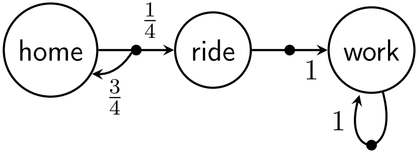

Table of contents

Encoding strategies

Memoryless strategies may not suffice for some specifications.

How can we encode strategies with **memory**?

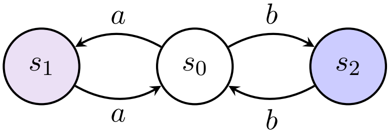

Mealy machines for pure finite-memory strategies

- Finite set of **memory states** M;
- initial **memory state** m_{init};
- **next-move** function $\mathsf{nxt}_{\mathfrak{M}}\colon M \times S \to A^{(i)}$;
- memory **update** function $\mathsf{up}_{\mathfrak{M}}\colon M \times S \times \bar{A} \to M$.

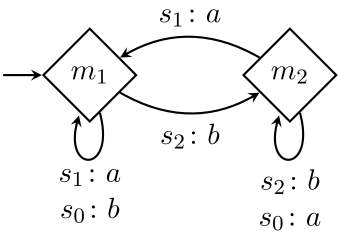

Complexity measure: **size of the memory state space**.

The study of finite memory

Key questions for finite-memory strategies

When does finite memory suffice?

$\rightsquigarrow$ Characterisations of specifications for which finite-memory suffices (e.g., [GZ05; Bou+22]).

How much memory do we need to play optimally?

$\rightsquigarrow$ Computing memory bounds [Bou+23; CO25].

$\rightsquigarrow$ Establishing improved bounds (e.g.,[JLS15; Mai24]).

Can we improve memory requirements by considering more general strategies?

$\rightsquigarrow$ Trading memory for **randomness** (e.g., [CdH04; CRR14]).

Focus on the result of [Mai24]

Memory requirements for Nash equilibria in reachability games

- **Context**: turn-based deterministic **multi-player** arenas.
- **Solution concept**: **Nash equilibria**.

Informal problem statement

How much **memory** do we need to implement a **good enough Nash equilibrium**?

Result for **pure strategies** and **move-independent** Mealy machines.

Theorem (M., STACS 2024)

- *For **reachability** and **shortest-path** games, $n^2 + 2n$ memory states suffice.*
- *For **Büchi** games, **finite memory** suffices.*

Table of contents

The limitations of pure strategies

How should $\mathcal{P}_1$ play to reach win with probability $\geq \frac{1}{3}$ after one move?

$\rightarrow$ **Pure strategies do not suffice!**

Solution: **randomisation**.

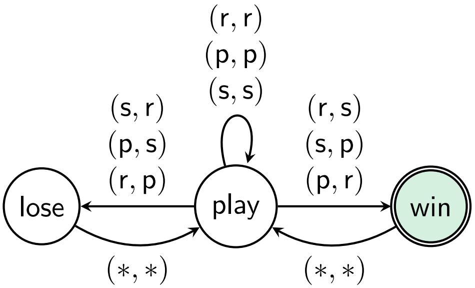

What is a randomised strategy?

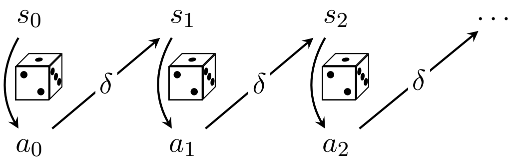

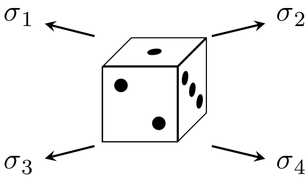

Behavioural strategy

$\sigma_i \colon \mathsf{Hist}(\mathcal{A}) \to \mathcal{D}(A^{(i)})$

Mixed strategy

$\mathcal{D}(\sigma_i \colon \mathsf{Hist}(\mathcal{A}) \to A^{(i)})$

How do these two classes of strategies compare?

Kuhn's theorem: same expressiveness when **perfect recall holds**.

Expressiveness criterion: **outcome-equivalence**.

What about finite-memory strategies?

Components of Mealy machines for **pure** strategies

- Initial **memory state** m_{init};
- **next-move** function $\mathsf{nxt}_{\mathfrak{M}} \colon M \times S \to A^{(i)}$;
- memory **update** function $\mathsf{up}_{\mathfrak{M}} \colon M \times S \times \bar{A} \to M$.

How can we **extend** Mealy machines to model **randomised strategies?**

Stochastic Mealy machines – behavioural version

- Initial **memory state** μ_{init};
- **randomised next-move** function $\mathsf{nxt}_{\mathfrak{M}} \colon M \times S \to \mathcal{D}(A^{(i)})$;
- memory **update** function $\mathsf{up}_{\mathfrak{M}} \colon M \times S \times \bar{A} \to M$.

What about finite-memory strategies?

Components of Mealy machines for **pure** strategies

- Initial **memory state** m_{init};
- **next-move** function $\mathsf{nxt}_{\mathfrak{M}} \colon M \times S \to A^{(i)}$;
- memory **update** function $\mathsf{up}_{\mathfrak{M}} \colon M \times S \times \bar{A} \to M$.

How can we **extend** Mealy machines to model **randomised strategies?**

Stochastic Mealy machines – mixed version

- Initial **memory distribution** $\mu_{\mathsf{init}} \in \mathcal{D}(M)$;
- **next-move** function $\mathsf{nxt}_{\mathfrak{M}} \colon M \times S \to A^{(i)}$;
- memory **update** function $\mathsf{up}_{\mathfrak{M}} \colon M \times S \times \bar{A} \to M$.

What about finite-memory strategies?

Components of Mealy machines for **pure** strategies

- Initial **memory state** m_{init};
- **next-move** function $\mathsf{nxt}_{\mathfrak{M}} \colon M \times S \to A^{(i)}$;
- memory **update** function $\mathsf{up}_{\mathfrak{M}} \colon M \times S \times \bar{A} \to M$.

How can we **extend** Mealy machines to model **randomised strategies?**

Stochastic Mealy machines – full randomisation

- Initial **memory distribution** $\mu_{\mathsf{init}} \in \mathcal{D}(M)$;
- **randomised next-move** function $\mathsf{nxt}_{\mathfrak{M}} \colon M \times S \to \mathcal{D}(A^{(i)})$;
- **randomised** memory **update** function $\mathsf{up}_{\mathfrak{M}} \colon M \times S \times \bar{A} \to \mathcal{D}(M)$.

Kuhn's theorem crumbles

Are all stochastic Mealy machine models equivalent?

Example: can $\mathcal{P}_1$ ensure Büchi(win) almost-surely with a

- **behavioural**-like Mealy machine? **Yes**.
- **mixed**-like Mealy machine? **No**.

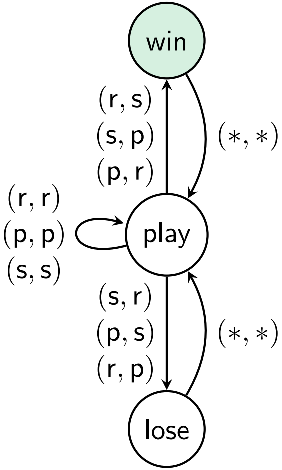

Main question

How do these different models compare?

Randomisation and finite memory

Acronyms **XYZ** where X, Y, Z $\in \{D, R\}$ and D = deterministic and R = random, and

- X is for initialisation;
- Y is for the next-move function,
- Z is for updates.

Classification of Mealy machine expressiveness (M., Randour, Inf. Comp., 2024)

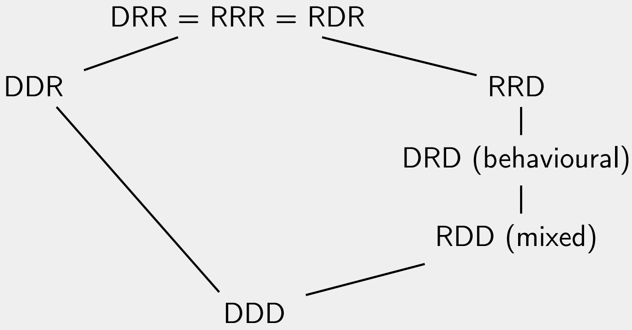

Table of contents

Another use for randomisation

Randomisation can be used to **balance multiple goals**. For instance:

- reaching work under 40 minutes with **high probability**;
- minimising the **expected** time to reach work.

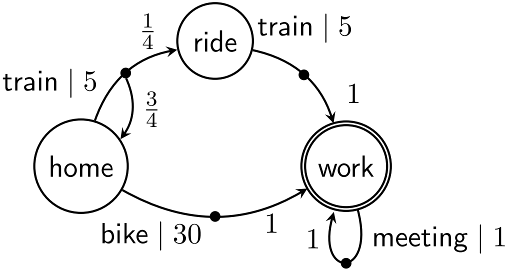

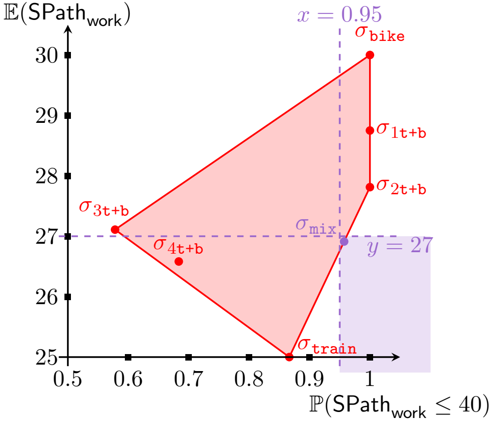

Randomisation requirements in multi-objective MDPs

Setting: MDPs with **multi-dimensional payoffs**.

In general, **randomised strategies are necessary** in multi-objective MDPs.

Main questions

- What is the relationship between expected payoffs of **pure strategies** and expected payoffs of **general strategies**?
- What **type of randomisation** do we need for multi-objective queries?

Applicability of our results

We want results that apply to a **broad class of payoffs**.

Which payoffs f do we consider?

- A payoff f is **good** (universally unambiguously integrable) if it has a **well-defined expectation** under all strategies from all initial states.
- A payoff f is **universally integrable** if its expectation is finite under all strategies from all initial states.

For a **multi-dimensional payoff** $\bar{f} = (f_1, \ldots, f_d)$ and $s \in S$, we study:

- $\mathsf{Pay}_s(\bar{f}) = \{\mathbb{E}_s^\sigma(\bar{f}) \mid \sigma \text{ strategy}\}$;
- $\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f}) = \{\mathbb{E}_s^\sigma(\bar{f}) \mid \sigma \text{ pure strategy}\}$.

Universally integrable payoffs

Theorem (M., Randour, 2025)

*Let $\bar{f}$ be **universally integrable**. Then for all $s \in S$,*

$$\mathsf{Pay}_s(\bar{f}) = \mathsf{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f})).$$

Proof idea: we reason on **lexicographic multi-objective MDPs**.

Lemma (M., Randour, 2025)

*If $\bar{f}$ is **universally integrable**, then for all strategies σ, there exists a **pure strategy** τ such that $\mathbb{E}_s^{\sigma}(\bar{f}) \leq_{\mathsf{lex}} \mathbb{E}_s^{\tau}(\bar{f})$.*

Mixing for universally integrable payoffs
Proof

Let $\bar{f}$ be universally integrable and $s \in S$.

Goal: show that $\mathsf{Pay}_s(\bar{f}) \subseteq \mathsf{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f}))$.

Fix a strategy σ and $\mathbf{q} = \mathbb{E}_s^{\sigma}(\bar{f})$.

Step 1: **isolate** $\mathbf{q}$ as much as possible with an intersection of **supporting hyperplanes**.

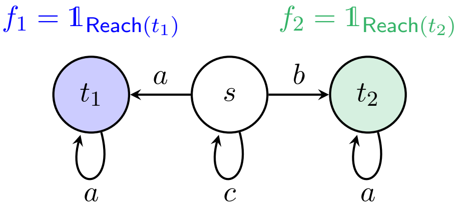

Example 1: $\mathbf{q} = (0, 1)$.

- First hyperplane: $x = 0 \rightsquigarrow x_1^*(x, y) = -x$.
- Second hyperplane: $y = 1 \rightsquigarrow x_2^*(x, y) = y$

σ is **lexicographically optimal** for $(x_1^*, x_2^*) \circ \bar{f}$

$\implies \mathbf{q} \in \mathsf{Pay}_s^{\mathsf{pure}}(\bar{f})$.

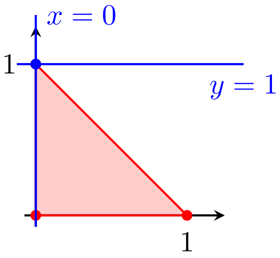

Mixing for universally integrable payoffs

Proof

Let $\bar{f}$ be universally integrable and $s \in S$.

Goal: show that $\mathsf{Pay}_s(\bar{f}) \subseteq \mathsf{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f}))$.

Fix a strategy σ and $\mathbf{q} = \mathbb{E}_s^{\sigma}(\bar{f})$.

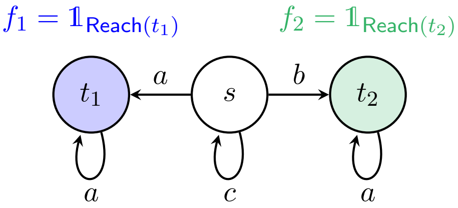

Step 1: **isolate** $\mathbf{q}$ as much as possible with an intersection of **supporting hyperplanes**.

Example 2: $\mathbf{q} = (\frac{1}{2}, \frac{1}{2})$.

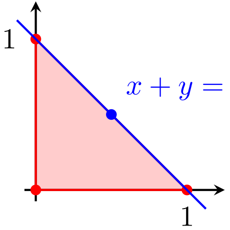

We construct $L_{\mathbf{q}}$ **linear** such that:

- σ **lexicographically optimal** from s for $L_{\mathbf{q}} \circ \bar{f}$;
- $\mathbf{q} \in \mathsf{ri}(\mathsf{Pay}_s(\bar{f}) \cap V))$ for $V = L_{\mathbf{q}}^{-1}(L_{\mathbf{q}}(\mathbf{q}))$

Mixing for universally integrable payoffs

Proof – continued

Goal: $\mathbf{q} \in \mathsf{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f}))$.

Step 2: it suffices to prove:
$\mathsf{cl}(\mathsf{Pay}_s(\bar{f}) \cap V) = \mathsf{cl}(\mathsf{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f})) \cap V)$.

Proof by contradiction.
Let $\mathbf{p} \in \mathsf{Pay}_s(\bar{f}) \cap V \setminus \mathsf{cl}(\mathsf{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f})) \cap V)$.
Separate $\mathbf{p}$ and $\mathsf{cl}(\mathsf{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f})) \cap V)$ with x^*.

There is a **pure strategy** τ such that

$$\mathbb{E}_s^{\tau}((L_{\mathbf{q}}, x^*) \circ \bar{f}) \geq_{\mathsf{lex}} (L_{\mathbf{q}}(\mathbf{p}), x^*(\mathbf{p})).$$

$$\implies x^*(\mathbb{E}_s^{\tau}(\bar{f})) \geq x^*(\mathbf{p}) \quad (\textbf{contradiction}).$$

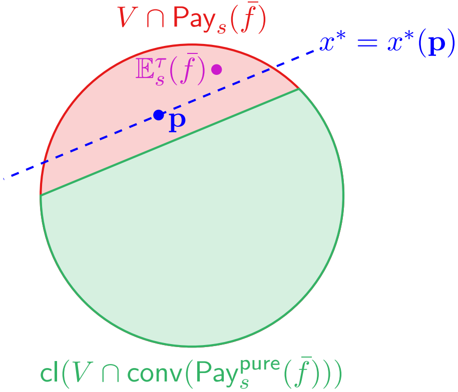

Beyond universally integrable payoffs

What if $\bar{f}$ is not universally integrable?

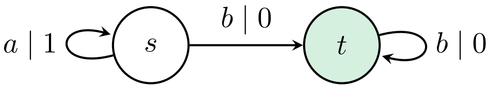

Non-universally-integrable example

1. reaching $t \rightsquigarrow f_1 = \mathbb{1}_{\mathsf{Reach}(t)}$;
2. sum of weights $\rightsquigarrow f_2 = \sum_{\ell=0}^{\infty} w(c_\ell)$.

The theorem does not generalise:

- $\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f}) = \{(0, +\infty)\} \cup \{(1, \ell) \mid \ell \in \mathbb{N}\}$
  $\implies \mathsf{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f})) = (\{1\} \times \mathbb{R}_{\geq 0}) \cup ([0, 1[\times \{+\infty\}),$
- $(1, +\infty) \in \mathsf{Pay}_s(\bar{f})$.

Theorem (M., Randour, 2025)

Let $\bar{f} = (f_1, \dots, f_d)$ be a good payoff and $s \in S$. Then

$$\mathsf{cl}(\mathsf{Pay}_s(\bar{f})) = \mathsf{cl}(\mathsf{conv}(\mathsf{Pay}_s^{\mathsf{pure}}(\bar{f}))).$$

Other results

How many strategies do we have to mix?

Theorem (M., Randour, 2025)

- *Payoffs of* ***finite-support mixed strategies*** *can be* ***obtained*** *by* ***mixing*** $d+1$ ***strategies***.
- *Payoffs of* ***finite-support mixed strategies*** *can be* ***dominated*** *by* ***mixing*** d ***strategies***.

When is a payoff set closed?

Theorem (M., Randour, 2025)

If $\bar{f}$ *is* ***continuous*** *and* ***universally square integrable***, *then* $\mathsf{Pay}_s(\bar{f})$ ***is compact*** *for all* $s \in S$.

Table of contents

Memory does not tell the whole story (1/2)

Action choices influence simplicity

Memory and randomisation do **not fully reflect** the complexity of a strategy.

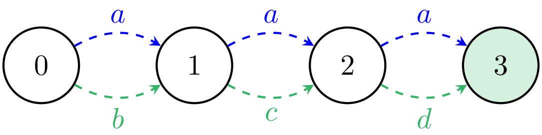

$\rightarrow$ Strategy σ_1 is **simpler to represent** than σ_2

- The **action choices** can impact how concise the strategy can be made.

Memory does not tell the whole story (2/2)

Counter-based strategies

Memory and **randomisation** do **not fully reflect** the complexity of a strategy.

- We consider a game with an **energy-Büchi** objective [CD12], where $W \in \mathbb{N}$.

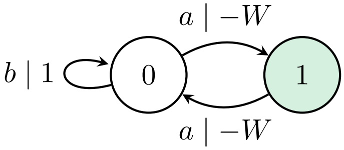

- Need memory **exponential** in the binary encoding of W to satisfy the objective.
- **Polynomial** representation with a **counter**-based approach.

Related challenge

How to represent and analyse **memoryless strategies** when the state space is **infinite**?

Memoryless strategies in one-counter MDPs

- We study **one-counter Markov decision processes**.
- We consider **interval strategies**: **counter-based** strategies with a **compact representation**.

Our contribution (Ajdarów, M., Novotný, Randour, ICALP 2025)

- PSPACE **verification** algorithms for interval strategies.
- PSPACE **realisability** algorithms for **structurally-constrained** interval strategies.

- Our algorithms are based on a **finite abstraction** of an **infinite system**.

One-counter Markov decision processes

One-counter MDP (OC-MDP) $\mathcal{Q}$

- **Finite** MDP (Q, A, δ).
- **Weight function** $w\colon Q \times A \to \{-1, 0, 1\}$.

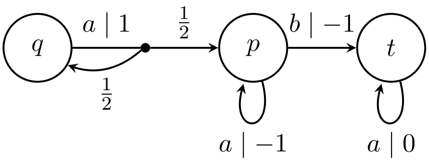

MDP $\mathcal{M}^{\leq\infty}(\mathcal{Q})$ induced by $\mathcal{Q}$

- **Countable** MDP over $S = Q \times \mathbb{N}$.
- State transitions via δ.
- Counter updates via w.

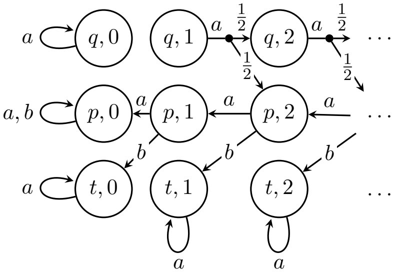

Objectives

- We study variants of **reachability objectives**.
- Let $T \subseteq Q$ **be a target**.

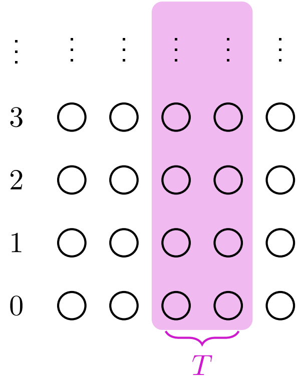

State reachability $\mathsf{Reach}(T)$

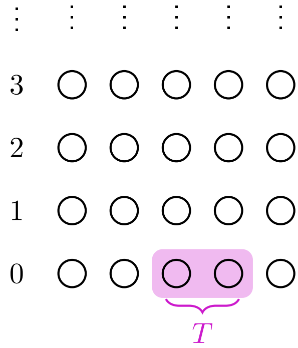

Selective termination $\mathsf{Term}(T)$

Interval strategies

We study a restricted class of **memoryless strategies** of $\mathcal{M}^{\leq\infty}(\mathcal{Q})$.

An **open-ended interval strategy (OEIS)** is a strategy over $Q \times \mathbb{N}_{>0}$ of the form:

$\mathbb{N}_0$	1	2	$\cdots$	$k_0 - 1$	k_0	$k_0 + 1$	$\cdots$
Q	σ_1	σ_2	$\cdots$	σ_{k_0-1}	σ_{k_0}	σ_{k_0}	$\cdots$

constant

Group counter values in intervals

Inter.	I_1	I_2	$\cdots$	$I_d = [\![k_0, \infty]\!]$
Q	τ_1	τ_2	$\cdots$	$\tau_d = \sigma_{k_0}$

$=$ **Finite partition** of $\mathbb{N}_{>0}$ into **intervals**

Verification

Interval strategy verification problem

Given an **interval strategy** σ, **an objective** $\Omega \in \{\mathsf{Reach}(T), \mathsf{Term}(T)\}$, a **threshold** $\theta \in \mathbb{Q} \cap [0,1]$ and an **initial configuration** $s_{\mathsf{init}} \in Q \times \mathbb{N}$, decide whether $\mathbb{P}^{\sigma}_{\mathcal{M}^{\leq\infty}(\mathcal{Q}), s_{\mathsf{init}}}(\Omega) \geq \theta$

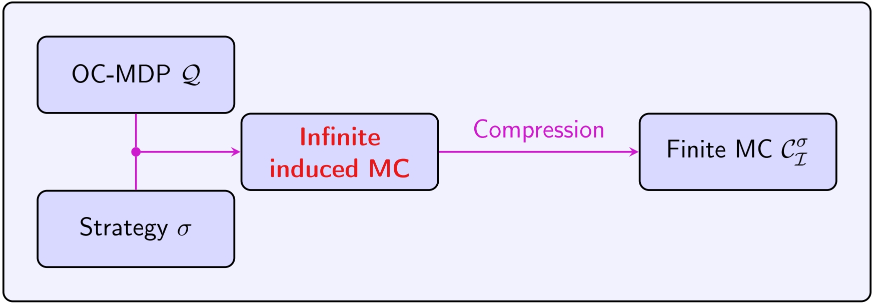

Interval strategy verification problem

We construct a finite **compressed Markov chain** $\mathcal{C}^{\sigma}_{\mathcal{I}}$.

Solving the verification problem through compressed Markov chains

- To compress, we **keep few configurations** and adjust transitions.
- We have formulae (in the signature $\{0, 1, +, -, \cdot, \leq\}$):
  - $\Phi^{\mathcal{I}}_{\delta}(\mathbf{x}, \mathbf{z}^{\sigma})$ for **transition probabilities** of $\mathcal{C}^{\sigma}_{\mathcal{I}}$;
  - $\Phi^{\mathcal{I}}_{\Omega}(\mathbf{x}, \mathbf{y})$ for **termination probabilities** from configurations of $\mathcal{C}^{\sigma}_{\mathcal{I}}$.

We can solve the verification problem by checking if

$$\mathbb{R} \models \forall\, \mathbf{x} \forall\, \mathbf{y}\, (\Phi^{\mathcal{I}}_{\delta}(\mathbf{x}, \mathbf{z}^{\sigma}) \wedge \Phi^{\mathcal{I}}_{\Omega}(\mathbf{x}, \mathbf{y})) \implies y_{s_{\mathsf{init}}} \geq \theta.$$

	Unbounded counter	**Bounded counter**
Upper bound	co-ETR	$\mathsf{P}^{\mathsf{PosSLP}}$
Lower bound	Square-root-sum-hard [EWY10]	Square-root-sum-hard

Synthesis of interval strategies

We have also studied the **synthesis** of **structurally-constrained interval strategies**.

Parameterised interval strategy synthesis problem

Given **parameters** d **and** $n \in \mathbb{N}_0$, does there exists an interval partition $\mathcal{I}$ of $\mathbb{N}$ and an OEIS σ such that

1. $|\mathcal{I}| \leq d$ and all bounded $I \in \mathcal{I}$ satisfy $|I| \leq n$;
2. σ is **based on** $\mathcal{I}$ and
3. $\mathbb{P}^{\sigma}_{s_{\mathsf{init}}}(\Omega) \geq \theta$.

	Unbounded counter	**Bounded counter**
Upper bound	PSPACE	$\mathsf{NP}^{\mathsf{ETR}}$
Lower bound	Square-root-sum-hard and NP-hard	

Table of contents

Summary

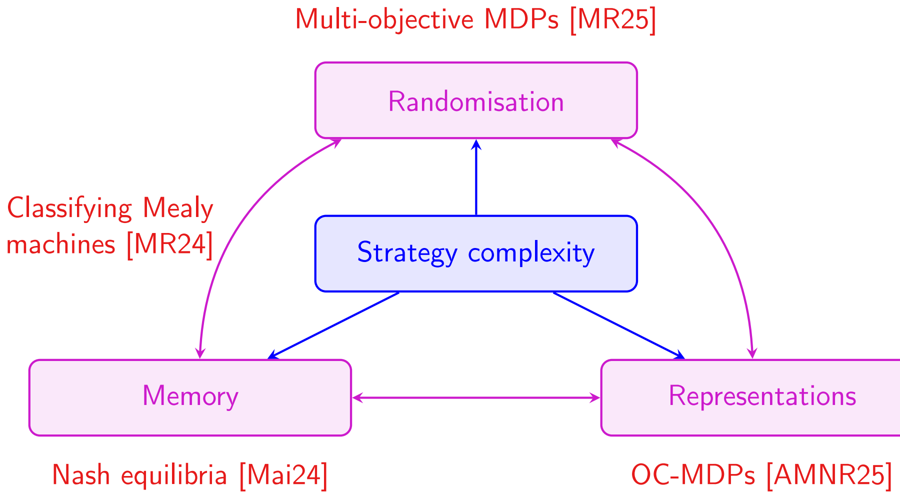

Multi-objective MDPs [MR25]
Randomisation
Classifying Mealy machines [MR24]
Strategy complexity
Memory
Representations
Nash equilibria [Mai24]
OC-MDPs [AMNR25]

Future work

Long-term goal

Developing an extensive and comprehensive framework of strategy complexity.

Other lines of work:

- Understanding **memory** requirements for **equilibria in multiplayer games**.
- Studying the power of **(finite-memory) randomised strategies** with respect to **given classes of payoffs**.
- Extending our results on **multi-objective MDPs** to also refer to **memory**.
- Finding whether there exist **well-structured** optimal strategies in **finite-horizon MDPs**.

Thanks!

References I

[Ajd+25] Michal Ajdarów et al. "Taming Infinity one Chunk at a Time: Concisely Represented Strategies in One-Counter MDPs". In: *Proceedings of the 52th International Colloquium on Automata, Languages, and Programming, ICALP 2025, July 8–11, 2025, Aarhus, Denmark*. Ed. by Keren Censor-Hillel et al. Vol. 334. LIPIcs. Schloss Dagstuhl –Leibniz-Zentrum für Informatik, 2025.

[Bou+22] Patricia Bouyer et al. "Games Where You Can Play Optimally with Arena-Independent Finite Memory". In: *Logical Methods in Computer Science* 18.1 (2022). DOI: 10.46298/lmcs-18(1:11)2022.

References II

[Bou+23] Patricia Bouyer et al. "How to Play Optimally for Regular Objectives?" In: *Proceedings of the 50th International Colloquium on Automata, Languages, and Programming, ICALP 2023, July 10–14, 2023, Paderborn, Germany*. Ed. by Kousha Etessami, Uriel Feige, and Gabriele Puppis. Vol. 261. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 118:1–118:18. DOI: 10.4230/LIPIcs.ICALP.2023.118.

[CD12] Krishnendu Chatterjee and Laurent Doyen. "Energy parity games". In: *Theoretical Computer Science* 458 (2012), pp. 49–60. DOI: 10.1016/j.tcs.2012.07.038. URL: https://doi.org/10.1016/j.tcs.2012.07.038.

References III

[CdH04] Krishnendu Chatterjee, Luca de Alfaro, and Thomas A. Henzinger. "Trading Memory for Randomness". In: *Proceedings of the 1st International Conference on Quantitative Evaluation of Systems, QEST 2004, Enschede, The Netherlands, 27–30 September 2004*. IEEE Computer Society, 2004, pp. 206–217. DOI: 10.1109/QEST.2004.1348035.

[CO25] Antonio Casares and Pierre Ohlmann. "The Memory of ω-Regular and $BC({\Sigma_2}^0)$ Objectives". In: *Proceedings of the 52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, July 8–11, 2025, Aarhus, Denmark*. Ed. by Keren Censor-Hillel et al. Vol. 334. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025, 149:1–149:18. DOI: 10.4230/LIPICS.ICALP.2025.149. URL: https://doi.org/10.4230/LIPIcs.ICALP.2025.149.

References IV

[CRR14] Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. "Strategy synthesis for multi-dimensional quantitative objectives". In: *Acta Informatica* 51.3-4 (2014), pp. 129–163. DOI: 10.1007/s00236-013-0182-6.

[EWY10] Kousha Etessami, Dominik Wojtczak, and Mihalis Yannakakis. "Quasi-Birth-Death Processes, Tree-Like QBDs, Probabilistic 1-Counter Automata, and Pushdown Systems". In: *Performance Evaluation* 67.9 (2010), pp. 837–857. DOI: 10.1016/J.PEVA.2009.12.009. URL: https://doi.org/10.1016/j.peva.2009.12.009.

[GZ05] Hugo Gimbert and Wiesław Zielonka. "Games Where You Can Play Optimally Without Any Memory". In: *Proceedings of the 16th International Conference on Concurrency Theory, CONCUR 2005, San Francisco, CA, USA, August 23–26, 2005*. Ed. by Martín Abadi and Luca de Alfaro. Vol. 3653. Lecture Notes in Computer Science. Springer, 2005, pp. 428–442. ISBN: 3-540-28309-9. DOI: 10.1007/11539452_33.

References V

[JLS15] Marcin Jurdzinski, Ranko Lazic, and Sylvain Schmitz. "Fixed-Dimensional Energy Games are in Pseudo-Polynomial Time". In: *Proceedings (Part II) of the 42nd International Colloquium on Automata, Languages, and Programming, ICALP 2015, Kyoto, Japan, July 6–10, 2015*. Ed. by Magnús M. Halldórsson et al. Vol. 9135. Lecture Notes in Computer Science. Springer, 2015, pp. 260–272. DOI: 10.1007/978-3-662-47666-6_21. URL: https://doi.org/10.1007/978-3-662-47666-6_21.

[Mai24] James C. A. Main. "Arena-Independent Memory Bounds for Nash Equilibria in Reachability Games". In: *41st International Symposium on Theoretical Aspects of Computer Science, STACS 2024, March 12-14, 2024, Clermont-Ferrand, France*. Ed. by Olaf Beyersdorff et al. Vol. 289. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024, 50:1–50:18. DOI: 10.4230/LIPICS.STACS.2024.50. URL: https://doi.org/10.4230/LIPIcs.STACS.2024.50.

References VI

[MR24] James C. A. Main and Mickael Randour. "Different strokes in randomised strategies: Revisiting Kuhn's theorem under finite-memory assumptions". In: *Information and Computation* 301 (2024), p. 105229. DOI: 10.1016/J.IC.2024.105229. URL: https://doi.org/10.1016/j.ic.2024.105229.

[MR25] James C. A. Main and Mickael Randour. "Mixing Any Cocktail with Limited Ingredients: On the Structure of Payoff Sets in Multi-Objective MDPs and its Impact on Randomised Strategies". In: *CoRR* abs/2502.18296 (2025). DOI: 10.48550/arXiv.2502.18296.