Verifying Concisely Represented Strategies in One-Counter Markov Decision Processes

Michal Ajdarów[2] **James C. A. Main**[1]
Petr Novotný[2] Mickael Randour[1]

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

[2]Masaryk University, Brno, Czech Republic

Journées Annuelles du GT Vérif 2024

Verifying one-counter Markov decision processes

- We study **one-counter Markov decision processes** (OC-MDPs).
  - **Markov decision process** (MDP): models systems with **non-determinism** and **randomness**.
  - **Counter**: can be incremented, decremented, left unchanged.
- An OC-MDP induces a **countable-state MDP**.

Verification problem

Given a **strategy**, an **objective** and a **threshold**, is the probability of the objective being satisfied no less than the threshold ?

- We focus on a class of **memoryless strategies** of the infinite MDP that admit a **finite representation**.
- We study variants of **reachability objectives**.

Table of contents

Table of contents

Markov decision processes

Markov decision process (MDP) $\mathcal{M}$

- **Finite or countable** state space S.
- **Finite** action space A.
- **Randomised** transition function $\delta\colon S \times A \to \mathcal{D}(S)$.

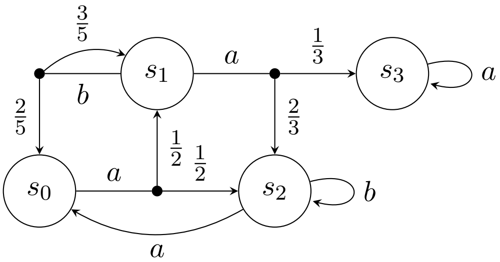

Plays are sequences in $(SA)^\omega$ coherent with transitions.
$\rightsquigarrow$ **Example**: $s_0 a s_1 b s_1 \ldots$

Strategies and induced Markov chains

- A **strategy** is a function $\sigma\colon (SA)^*S \to \mathcal{D}(A)$.
- σ is **memoryless** if its choices depend only on the **current state**.
- We view memoryless strategies as functions $S \to \mathcal{D}(A)$.
- A memoryless strategy σ induces a **Markov chain** over S.

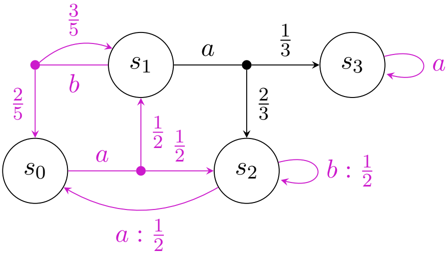

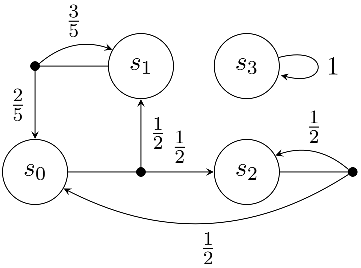

Table of contents

One-counter Markov decision processes

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

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

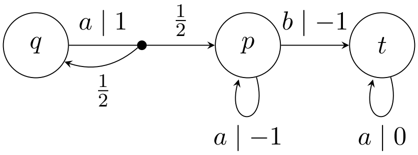

MDP $\mathcal{M}^{\leq\infty}(\mathcal{Q})$ induced by $\mathcal{Q}$

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

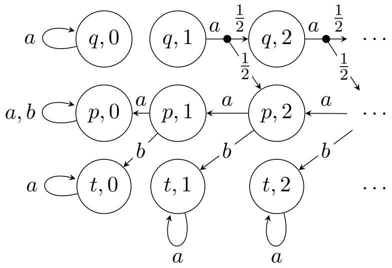

Interval strategies

We study two classes of **memoryless strategies** of $\mathcal{M}^{\leq\infty}(\mathcal{Q})$.

- **Open-ended interval strategies (OEIS)**: σ is an OEIS if there exists $k_0 \in \mathbb{N}$ such that, for all $q \in Q$ and all $k \geq k_0$, $\sigma(q, k) = \sigma(q, k_0)$.

Representing an OEIS

An OEIS is described by a **finite partition** $\mathcal{I}$ of $\mathbb{N}_0$ into intervals and a **function** $Q \times \mathcal{I} \to \mathcal{D}(A)$.

- **Cyclic interval strategies (CIS)**: σ is a CIS if there exists $\rho \in \mathbb{N}_0$ such that, for all $q \in Q$ and all $k \in \mathbb{N}_0$, $\sigma(q, k) = \sigma(q, k + \rho)$.

Representing a CIS

A CIS is described by a **period** ρ, a **partition** $\mathcal{I}$ of $[\![1, \rho]\!]$ into intervals and a **function** $Q \times \mathcal{I} \to \mathcal{D}(A)$.

Conciseness of interval strategies

- Let σ be an interval strategy of $\mathcal{M}^{\leq\infty}(\mathcal{Q})$.
- There exists a **strategy of $\mathcal{Q}$** that induces the **same behaviour** as σ when an initial counter value is fixed $\rightsquigarrow$ **memory = counter value**.

OEISs may require infinite memory

The OEIS σ such that:

- $\sigma(p, 1) = b$
- $\sigma(p, k) = a$ for all $k \geq 2$.

requires **infinite memory** in $\mathcal{Q}$.

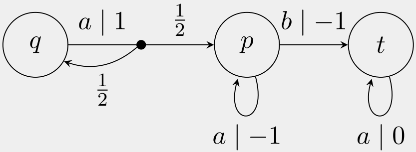

- CISs correspond to **exponential-size finite-memory** strategies of $\mathcal{Q}$.

Objectives

Let $\mathcal{Q} = (Q, A, \delta, w)$ be an OC-MDP. We consider two **objectives** for a **target** $T \subseteq Q$.

- **State reachability**: Reach(T) is the set of plays **visiting** T.
- **Selective termination**: Term(T) is the set of plays for which **counter value** 0 **is reached in** T.

Interval strategy verification problem

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

Goal: explain how to solve the interval strategy verification problem in **polynomial space**.

Table of contents

Verification and interval strategies

- We develop techniques to analyse the **infinite** Markov chain induced by an **interval strategy**.
- For OEISs, we reduce to the analysis of a **finite** Markov chain.
- For CISs, we reduce to the analysis of a **one-counter** Markov chain.

$\rightarrow$ **We focus on an OEIS σ based on a partition $\mathcal{I}$ from here.**

Main idea: compressing the configuration space

For each interval $I \in \mathcal{I}$:

- we only keep a **subset of the configurations** in $Q \times I$ and
- we **aggregate several transitions** of $\mathcal{M}^{\leq\infty}(\mathcal{Q})$ into one.

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

Compressing the unbounded interval

- Let $I \in \mathcal{I}$ be the unbounded interval, i.e., $Q \times I$ is **infinite**.
- We keep only **one configuration per state**.

Example: σ choosing a **in all states** for the interval $\mathbb{N}_0 = [\![1, \infty]\!]$.

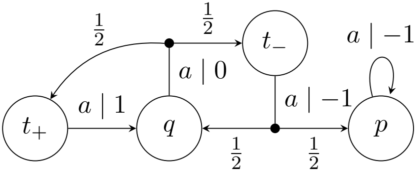

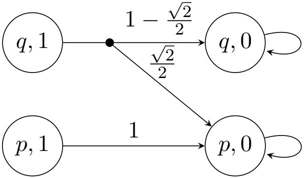

$\rightsquigarrow$ Transition probabilities can be **irrational**.

Theorem ([KEM06][1])

The transition probabilities of $\mathcal{C}^{\sigma}_{\mathcal{I}}$ with respect to $Q \times I$ are the least non-negative solution of a quadratic system of equations.

[1] Kucera et al., "Model Checking Probabilistic Pushdown Automata", LMCS 2006.

Compressing bounded intervals

Motivation

- Let $I \in \mathcal{I}$ be bounded.
- The set of configurations $Q \times I$ is **finite**.

Why do we want to compress bounded intervals ?

- The bounds of I are encoded in **binary**.
- Thus $Q \times I$ is of **exponential size**.
- **Goal**: **polynomial-size** compressed Markov chain.

Compressing bounded intervals

State space and transition structure

Main idea: retain configurations by considering counter changes by **powers of two**.

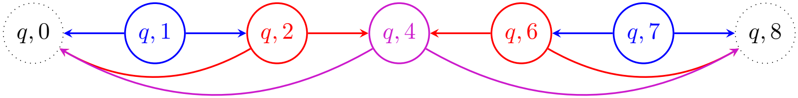

- The construction requires that $|I| = 2^x - 1$ for some $x \in \mathbb{N}_0$.
- We retain at most $2x - 1$ counter values.

Compressing bounded intervals

Transition probabilities

Example: σ playing **uniformly at random** for the interval $[\![1, 15]\!]$.

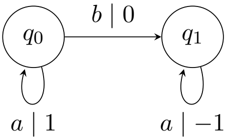

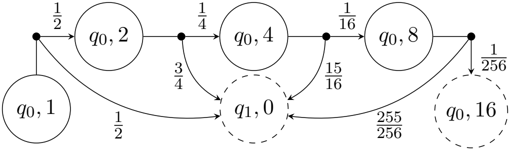

$\rightsquigarrow$ Transition probabilities can require **exponential-size** representations.

Theorem

*The transition probabilities of $\mathcal{C}^{\sigma}_{\mathcal{I}}$ with respect to $Q \times I$ are the **least non-negative solution** of a **quadratic system of equations**.*

Verification via compressed Markov chains

Summary: compressed Markov chain $\mathcal{C}_{\mathcal{I}}^{\sigma}$

- **Polynomial-size** state space.
- Transition probabilities given by **polynomial-size equations systems**.
- **Preserves termination probabilities**.

For CISs, we can use the same approach to derive a **compressed one-counter Markov chain**.

Unbounded counter		Bounded counter
OEIS	**CIS**	**OEIS**
co-ETR	co-ETR	$\mathsf{P}^{\mathsf{PosSLP}}$
Square-root sum-hard [EWY10][2]		Square-root sum-hard

[2] Etessami et al., "Quasi-Birth-Death Processes, Tree-Like QBDs, Probabilistic 1-Counter Automata, and Pushdown Systems", Perform. Evaluation 2010.

References I

Kousha Etessami, Dominik Wojtczak, and Mihalis Yannakakis. "Quasi-Birth-Death Processes, Tree-Like QBDs, Probabilistic 1-Counter Automata, and Pushdown Systems". In: *Perform. 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.

Antonín Kucera, Javier Esparza, and Richard Mayr. "Model Checking Probabilistic Pushdown Automata". In: *Log. Methods Comput. Sci.* 2.1 (2006). DOI: 10.2168/LMCS-2(1:2)2006. URL: https://doi.org/10.2168/LMCS-2(1:2)2006.