Decidability of Omega-Regular Objectives for POMDPs with Revelations

[Work in progress]

Pierre Vandenhove

Joint work with Nathanaël Fijalkow, Hugo Gimbert, Guillermo A. Pérez

LaBRI, Université de Bordeaux

Dagstuhl Stochastic Games – June 4, 2024

université de BORDEAUX

LABORATOIRE BORDELAIS DE RECHERCHE EN INFORMATIQUE LaBRI

Outline

Partially observable Markov decision processes (POMDPs):

- stochastic,
- nondeterministic,
- **uncertainty** about the actual state.

Goal

Strategy synthesis for **parity objectives** ($\rightsquigarrow$ ω-regular objectives).
Undecidable in general; **decidable subclasses**?

Means

Two subclasses with probabilistic guarantees about sometimes **knowing the actual state**.

Partially observable MDPs

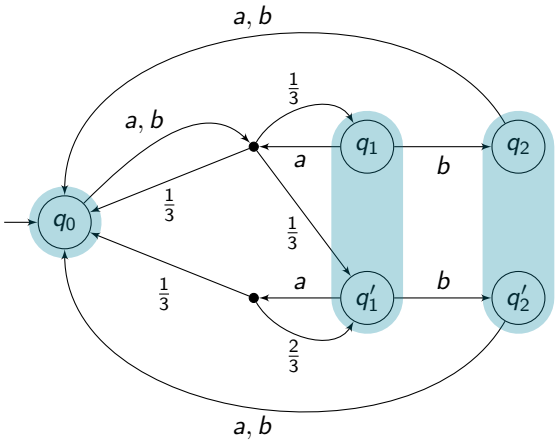

States, **actions**, **observations** (strategies are only based on observations!).

Objective

- Function $p\colon Q \to \{0, \ldots, d\}$ assigning **priorities** to **states**.
- **Parity objective**: the **maximal** priority seen infinitely often is **even**.
- Common subclasses:
  - **Büchi**: $p\colon Q \to \{1, 2\}$: something good (2) occurs infinitely often,
  - **coBüchi**: $p\colon Q \to \{0, 1\}$: something bad (1) occurs finitely often.
- **Almost-sure** strategies; "qualitative".

Theorem[1,2]

- Almost-sure **reachability**, **safety**, and **Büchi** are **EXPTIME-complete**.
- Almost-sure **coBüchi** (and therefore **parity**) is **undecidable**.

[1] Baier, Größer, and Bertrand, "Probabilistic ω-automata", 2012.

[2] Chatterjee, Chmelik, and Tracol, "What is decidable about partially observable Markov decision processes with ω-regular objectives", 2016.

Example

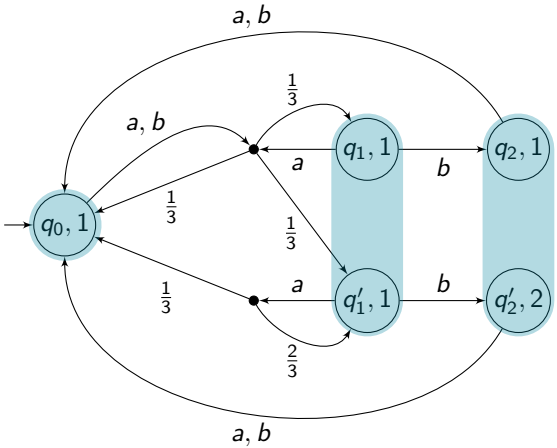

Almost-sure strategy? Yes! Move to q_2/q_2' infinitely often.

Example

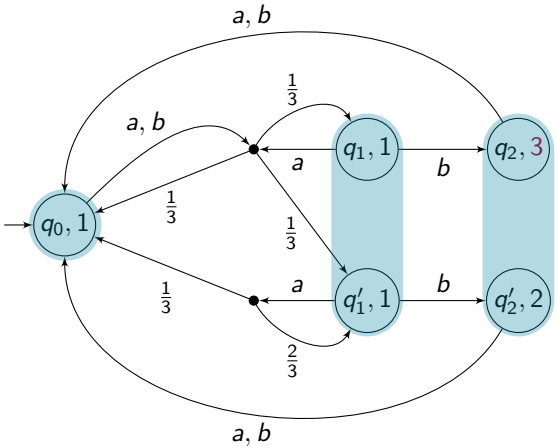

Almost-sure strategy? Yes! Move to q_2/q_2' when increasingly high probability to be in q_1'.

Belief (support) MDP

POMDPs induce **infinite belief MDPs**:

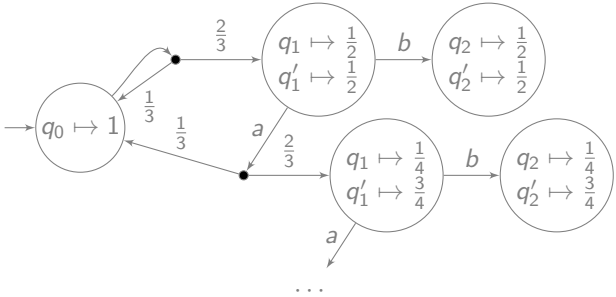

Finite: only keep belief **supports**:

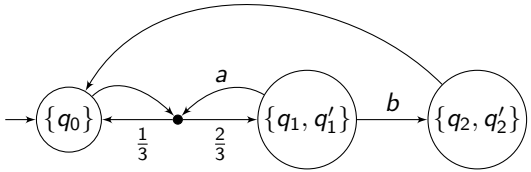

When does the analysis of the belief **support** MDP suffice?

Non-soundness of the belief support MDP

No almost-sure strategy in the POMDP, but OK in the belief support MDP.

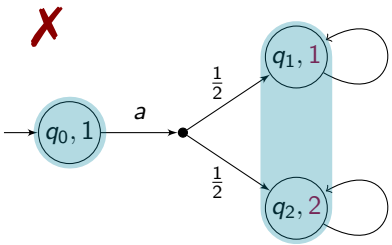

POMDP

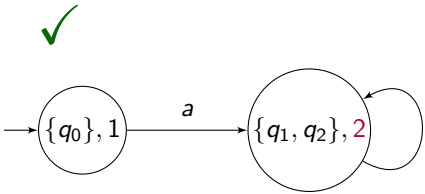

Belief support MDP

(Technical detail: how to lift the priority function? Take the **max**.)

Incompleteness of the belief support MDP

Almost-sure strategy in the POMDP, **not** in the belief support MDP.

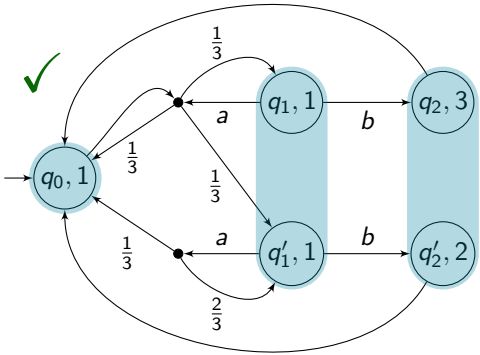

POMDP

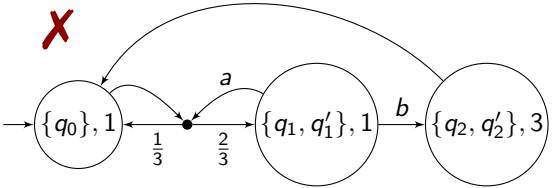

Belief support MDP

First revealing property

Property 1

A POMDP is **weakly revealing** if for all strategies,
almost surely, the **current state is known** infinitely often.

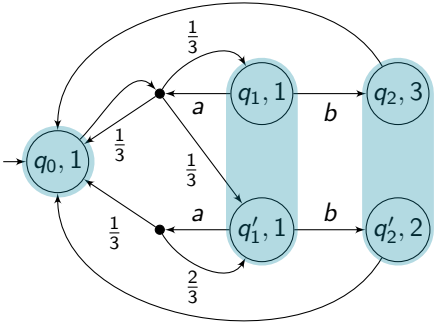

Weakly revealing

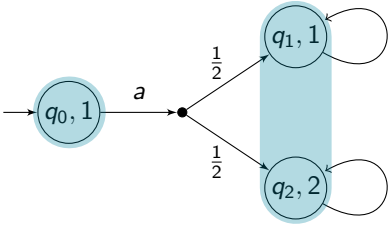

Not weakly revealing

First revealing property

Property 1

A POMDP is **weakly revealing** if for all strategies,
almost surely, the **current state is known** infinitely often.

As much information in the belief **support** MDP as in the infinite belief MDP.

Deciding the property

Deciding whether a POMDP is **weakly revealing** is EXPTIME-complete.

Weakly revealing POMDPs

Let $\mathcal{P}$ be a **weakly revealing** POMDP with a parity objective.

Soundness for parity

Almost-sure winning strategy in the **belief support MDP** of $\mathcal{P} \implies$ also in **POMDP** $\mathcal{P}$.

Reciprocal OK for Büchi and coBüchi:

Completeness for Büchi and coBüchi

Almost-sure winning strategy in **POMDP** $\mathcal{P} \implies$ also in the **belief support MDP** of $\mathcal{P}$.

Analysing the belief support MDP is **sound** and **complete** for Büchi and coBüchi.

Decidability of weakly revealing POMDPs

Decidability

Almost-sure **coBüchi** (and **Büchi**) is decidable for **weakly revealing** POMDPs.

Algorithm: solve the **belief support MDP** $\rightsquigarrow$ in EXPTIME.

Compared to general POMDPs:
$\rightsquigarrow$ makes **coBüchi decidable**,
$\rightsquigarrow$ **simpler algorithm** for Büchi.

Parity still not decidable

Belief support MDP is "incomplete" for this weakly revealing POMDP:

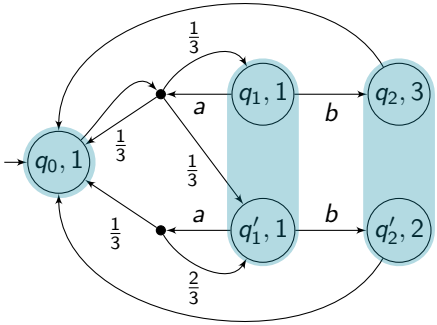

Undecidability

Almost-sure **parity** $\{1, 2, 3\}$ is **undecidable** for **weakly revealing** POMDPs.

Reduction from the value-1 problem for probabilistic automata.[3]

[3] Gimbert and Oualhadj, "Probabilistic Automata on Finite Words: Decidable and Undecidable Problems", 2010.

Second revealing property

Property 2

A POMDP is **strongly revealing** if after every action,
there is **non-zero probability to know the current state**.

- Strongly revealing $\implies$ weakly revealing.
- Easy syntactic transformation $\rightsquigarrow$ "optimistic POMDP".

Completeness for **parity**

Almost-sure winning strategy in **strongly revealing POMDP** $\mathcal{P} \implies$ also in the **belief support MDP** of $\mathcal{P}$.

Theorem

Almost-sure **parity** is decidable for **strongly revealing** POMDPs.

Summary

Almost-sure strategies:

	Büchi	coBüchi	Parity	coBüchi **games**
In general	EXPTIME	Undec.	Undec. $(0, 1)$	Undec.
Weakly revealing	EXPTIME	EXPTIME	Undec. $(1, 2, 3)$	Undec.
Strongly revealing	EXPTIME	EXPTIME	EXPTIME	Undec.

- CoBüchi is **undecidable** for **strongly revealing** partial-observation **games**!
- Planning for **tool support**.
- **Open problems**:
  - Larger class where the **belief support MDP** is sound and complete?
  - Larger **decidable classes** for coBüchi/parity?

Thanks!