Optimally controlling a random population

Hugo Gimbert
LaBRI, Bordeaux

Corto Mascle
MPI-SWS Kaiserslautern

Patrick Totzke
University of Liverpool

Population MDPs

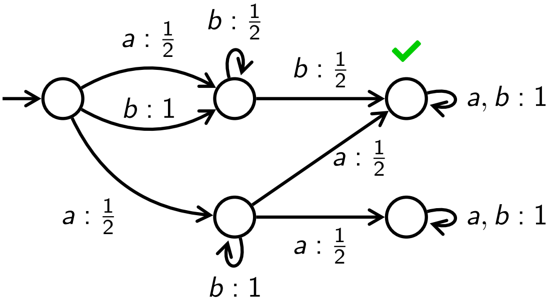

Population MDPs

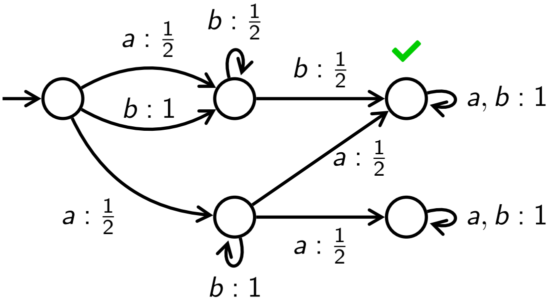

Given a Markov decision process, find a strategy to make it reach ✔ with probability 1.

Population MDPs

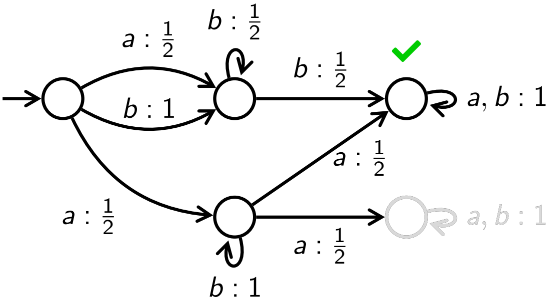

Given a Markov decision process, find a strategy to make it reach ✓ with probability 1.

Population MDPs

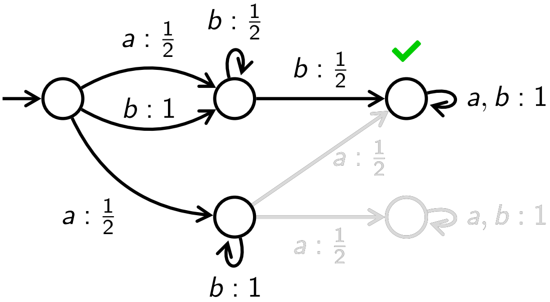

Given a Markov decision process, find a strategy to make it reach ✓ with probability 1.

Population MDPs

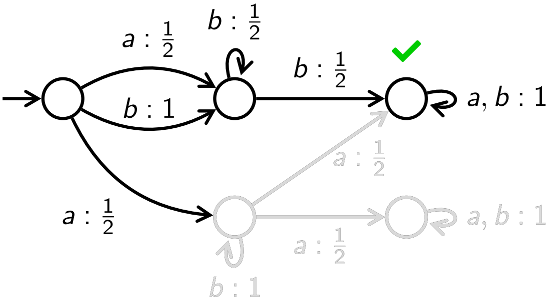

Given a Markov decision process, find a strategy to make it reach ✔ with probability 1.

Population MDPs

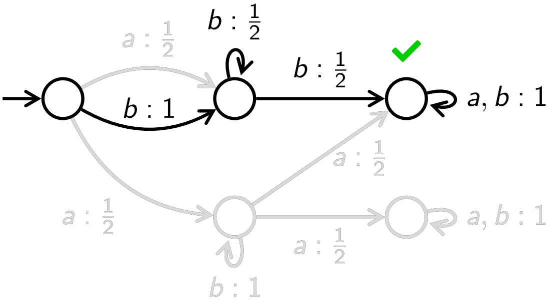

Given a Markov decision process, find a strategy to make it reach ✔ with probability 1.

Population MDPs

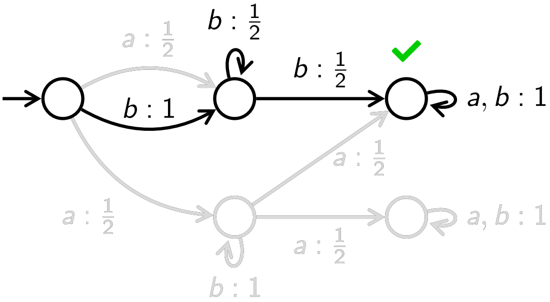

Given a Markov decision process, find a strategy to make it reach ✔ with probability 1.

- Remove dead ends
- Remove risky actions

Population MDPs

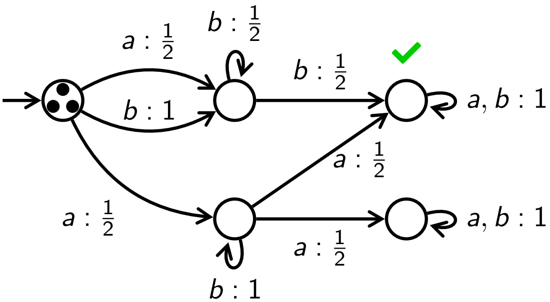

Given N identical Markov decision process, make them all reach ✔ with probability 1.

Population MDPs

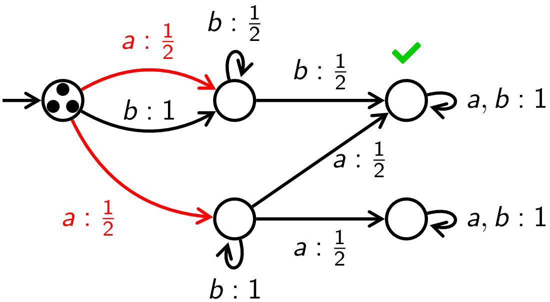

Given N identical Markov decision process, make them all reach ✓ with probability 1.

Population MDPs

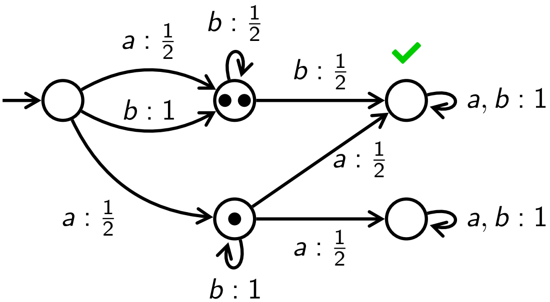

Given N identical Markov decision process, make them all reach ✓ with probability 1.

Population MDPs

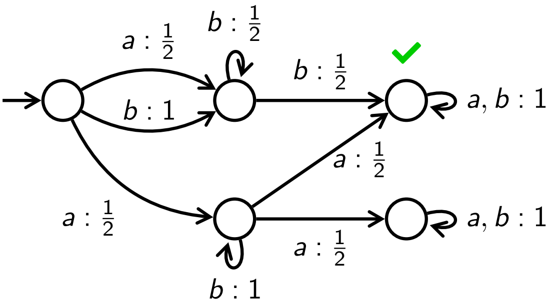

Given N identical Markov decision process, make them all reach ✓ with probability 1.

$\Rightarrow$ Product of N MDPs.

Population MDPs

Random population control problem

Given an MDP $\mathcal{M}$, is there a winning strategy against N tokens, **for all** N?

Population MDPs

Random population control problem

Given an MDP $\mathcal{M}$, is there a winning strategy against N tokens, **for all** N?

- Control of systems made of many simple components
- Parameterised verification
- Probabilistic automata
- Letter games, history-determinism

Population MDPs

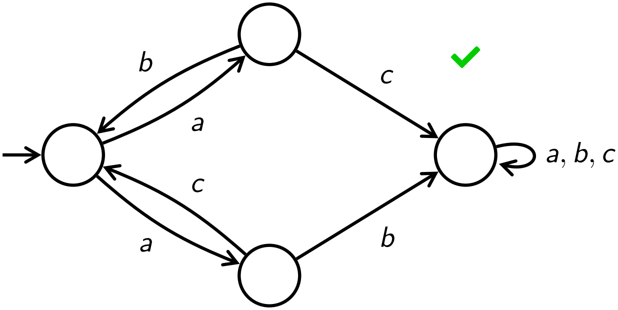

Population MDPs

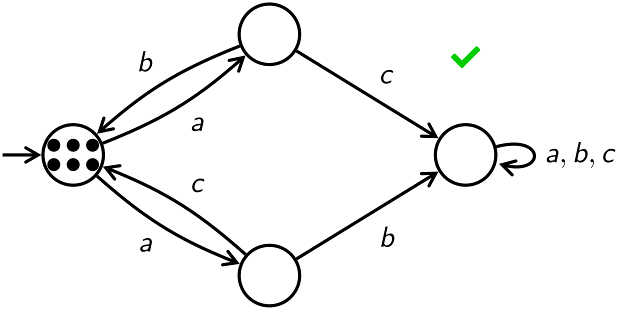

Population MDPs

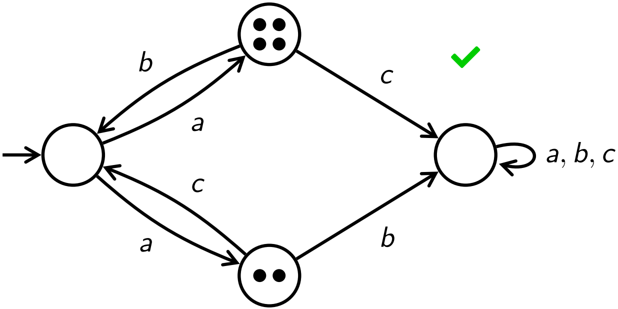

Population MDPs

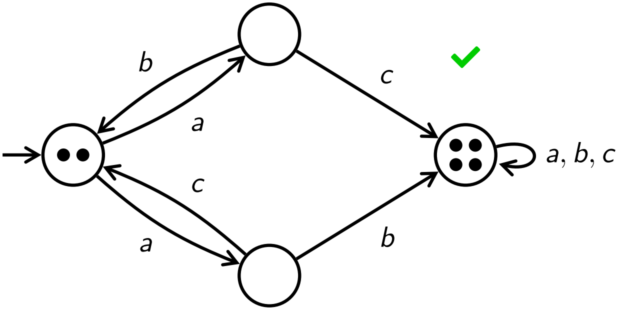

Population MDPs

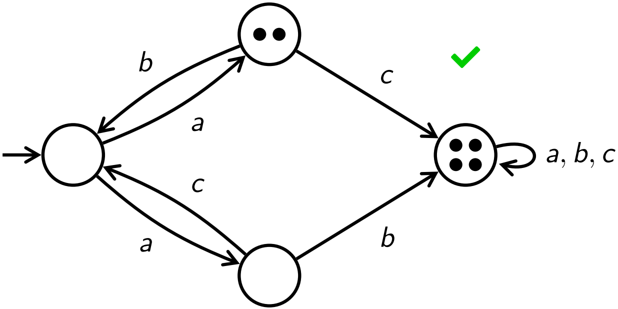

Population MDPs

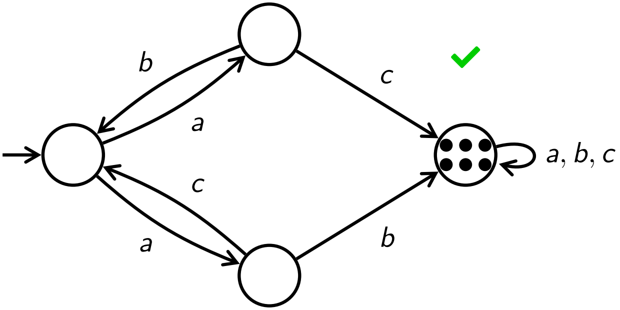

Population MDPs

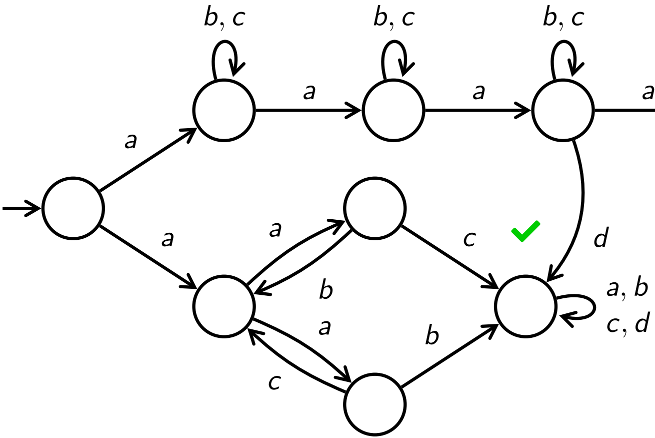

Population MDPs

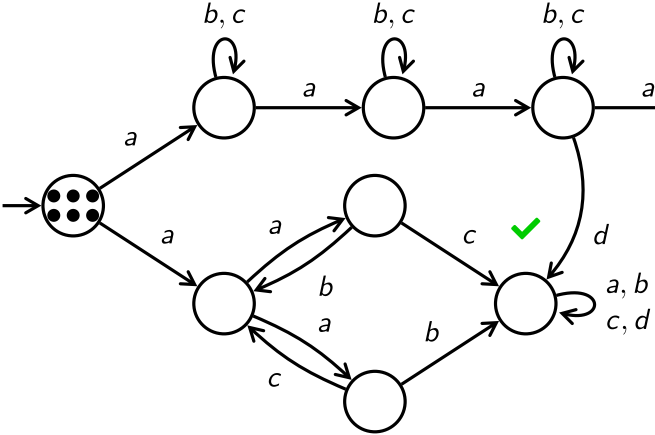

Population MDPs

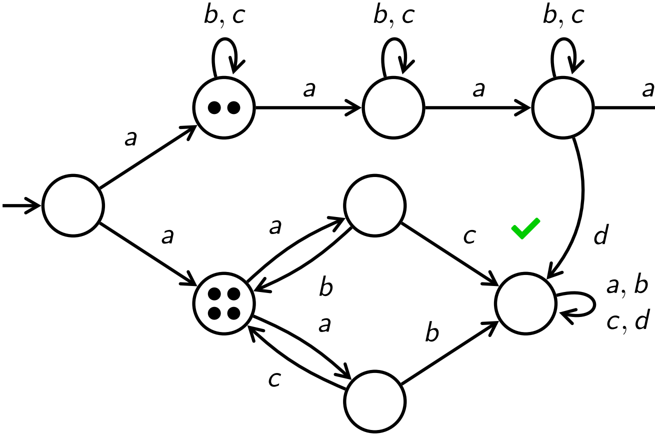

Population MDPs

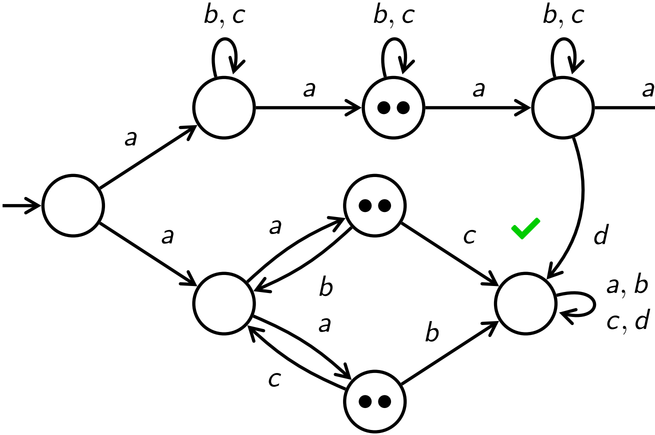

Population MDPs

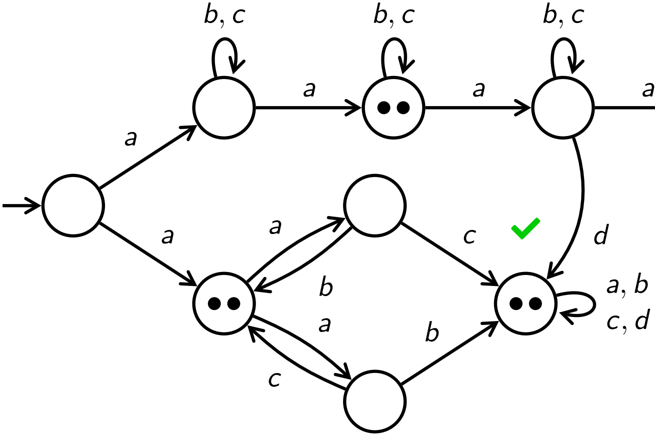

Population MDPs

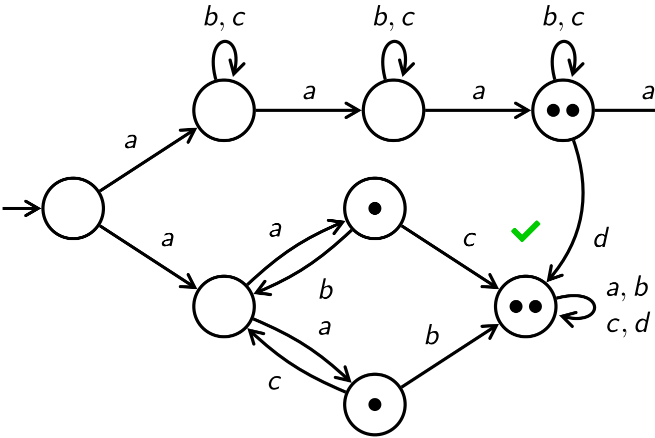

Population MDPs

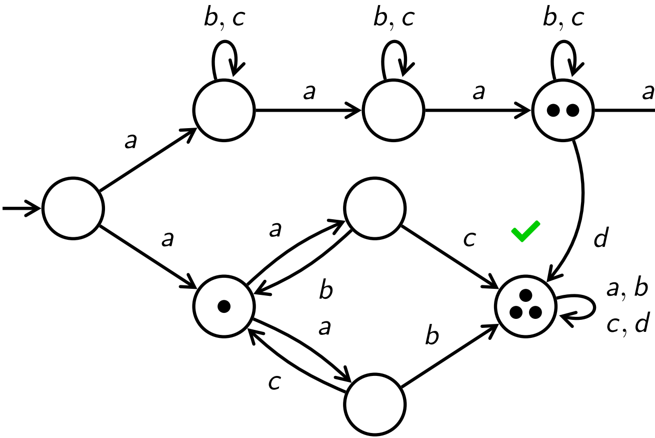

Population MDPs

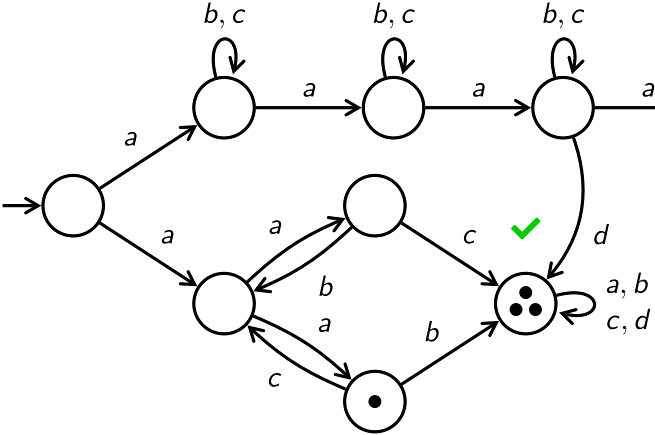

Population MDPs

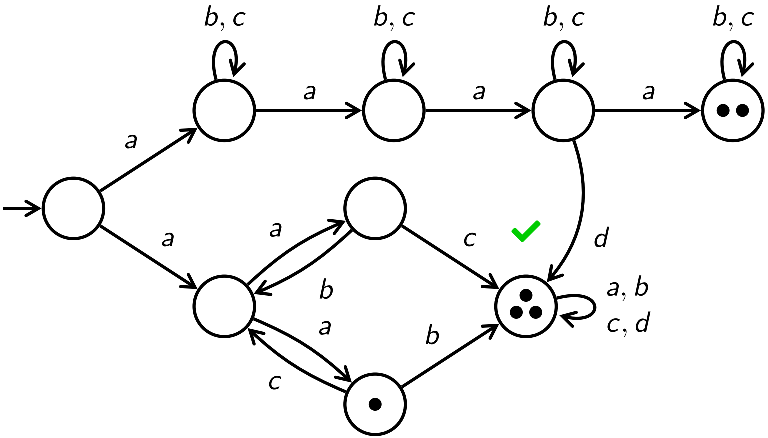

Adapted from [Bertrand Dewaskar Genest Gimbert '15]

There are MDPs of size $\mathcal{O}(k)$ for which we can win against 2^{2^k} tokens and not more.

Population MDPs

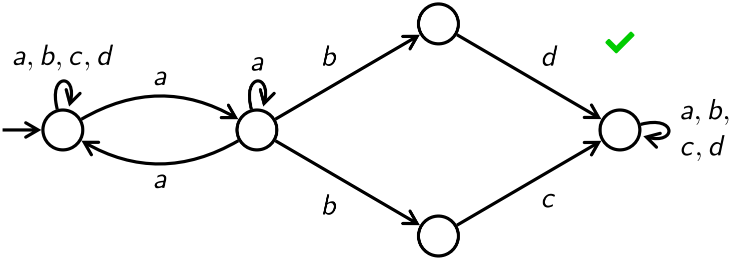

Population MDPs

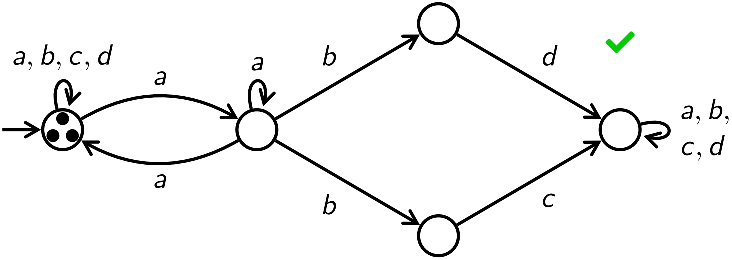

Population MDPs

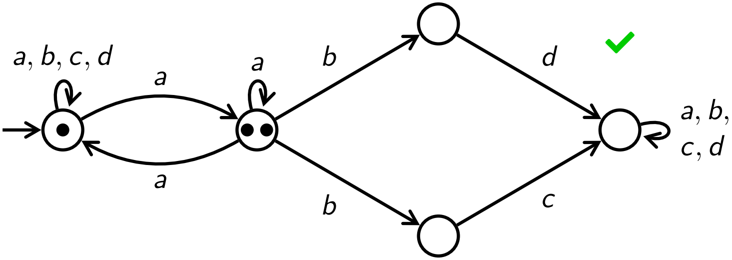

Population MDPs

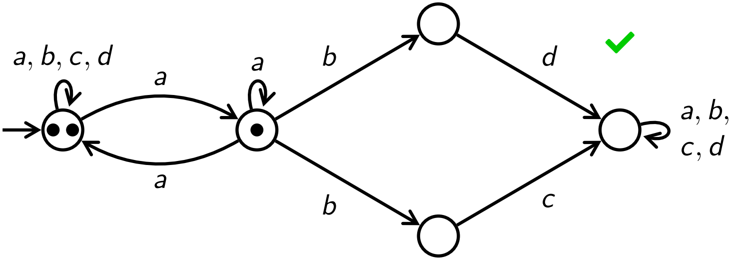

Population MDPs

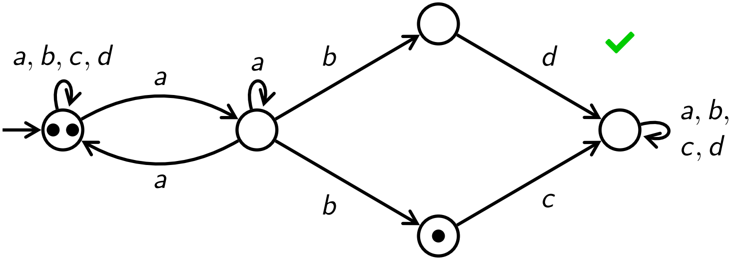

Population MDPs

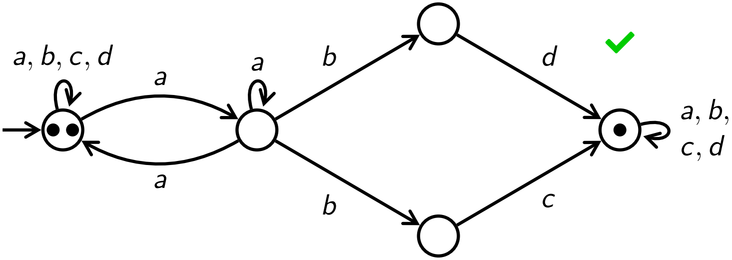

Population MDPs

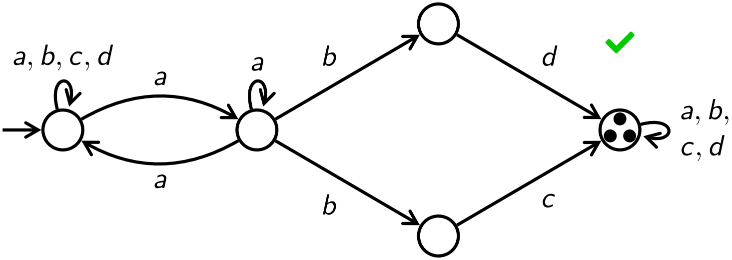

Population MDPs

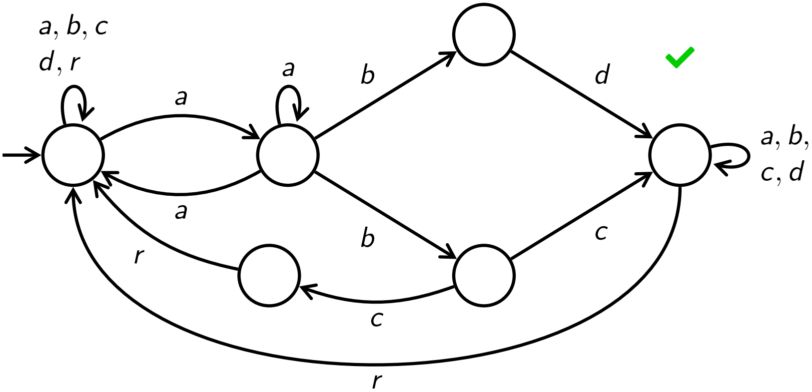

Population MDPs

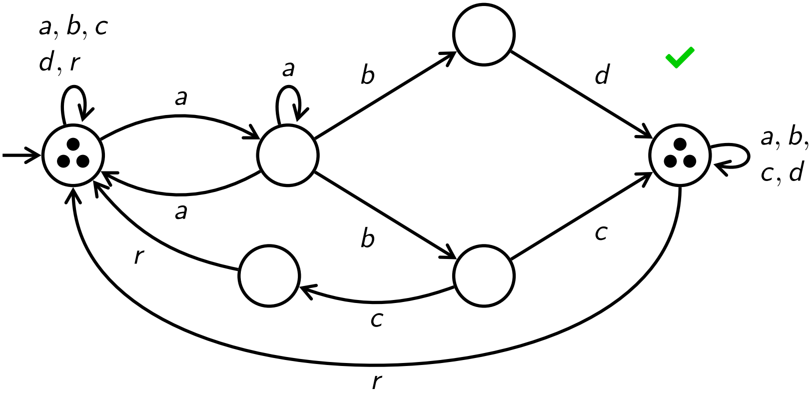

Population MDPs

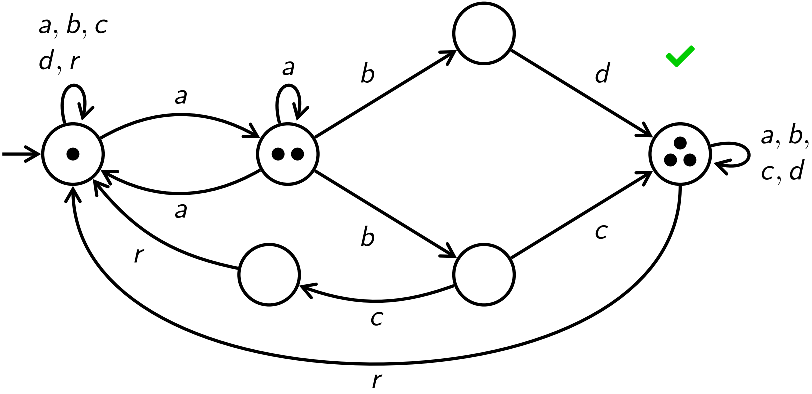

Population MDPs

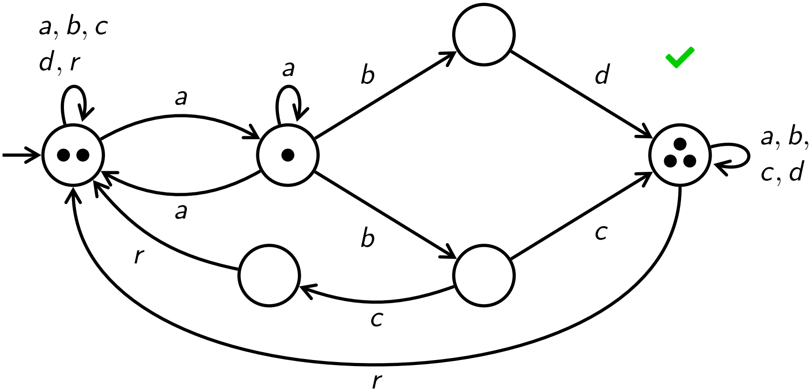

Population MDPs

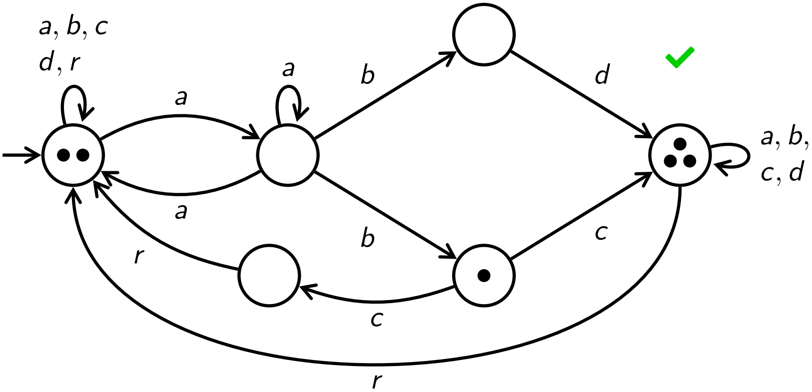

Population MDPs

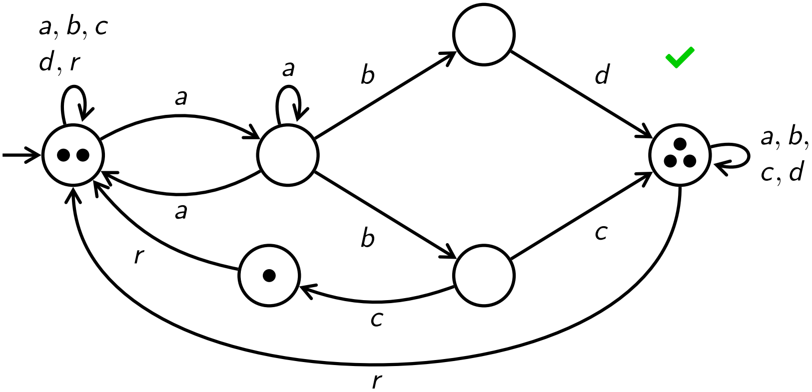

Population MDPs

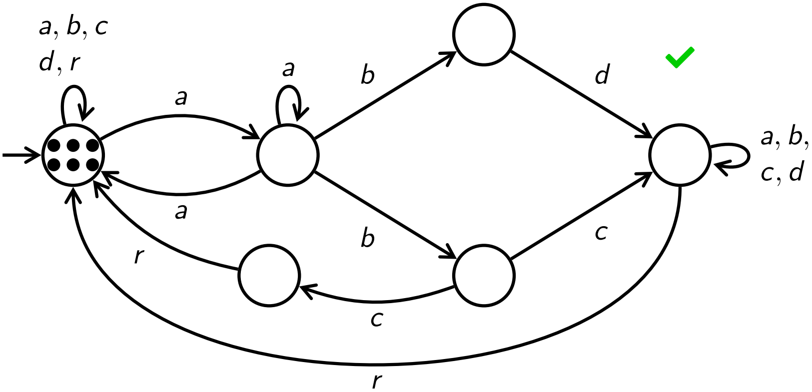

History

Controlling a population [Bertrand Dewaskar Genest Gimbert Godbole '19]

When tokens are controlled by an opponent, existence of a winning strategy is EXPTIME-complete.

History

Controlling a population [Bertrand Dewaskar Genest Gimbert Godbole '19]

When tokens are controlled by an opponent, existence of a winning strategy is EXPTIME-complete.

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

The Randomised Population Control Problem is decidable.

History

Controlling a population [Bertrand Dewaskar Genest Gimbert Godbole '19]

When tokens are controlled by an opponent, existence of a winning strategy is EXPTIME-complete.

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

The Randomised Population Control Problem is decidable.

- Continuous models - [Akshay, Genest, Vyas '18], [Doyen, van den Bogaart '22]
- Relation to history-determinism - [Hazard, Idir, Kuperberg '25]

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

- Less tokens is always better $\rightsquigarrow$ The set of winning configurations is *downward-closed*.

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

- Less tokens is always better $\rightsquigarrow$ The set of winning configurations is *downward-closed*.
- The set of configurations from which action a is safe is also *downward-closed*.

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

- Less tokens is always better $\rightsquigarrow$ The set of winning configurations is *downward-closed*.
- The set of configurations from which action *a* is safe is also *downward-closed*.
- Downward-closed sets of configurations can be represented as finite unions of *ideals*.

$$(\omega, 5, \omega, 8) \downarrow \cup (4, \omega, \omega, 4) \downarrow \cup (9, 6, 1, 4) \downarrow$$

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

- Less tokens is always better $\rightsquigarrow$ The set of winning configurations is *downward-closed*.
- The set of configurations from which action *a* is safe is also *downward-closed*.
- Downward-closed sets of configurations can be represented as finite unions of *ideals*.

$$(\omega, 5, \omega, 8)\downarrow \cup (4, \omega, \omega, 4)\downarrow \cup (9, 6, 1, 4)\downarrow$$

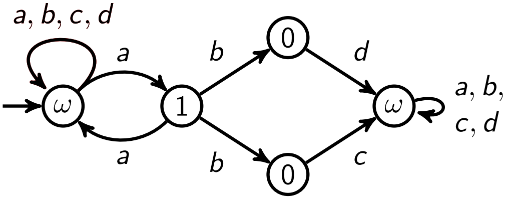

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

- Less tokens is always better $\rightsquigarrow$ The set of winning configurations is *downward-closed*.
- The set of configurations from which action a is safe is also *downward-closed*.
- Downward-closed sets of configurations can be represented as finite unions of *ideals*.

$$(\omega, 5, \omega, 8)\downarrow \cup\ (4, \omega, \omega, 4)\downarrow \cup\ (9, 6, 1, 4)\downarrow$$

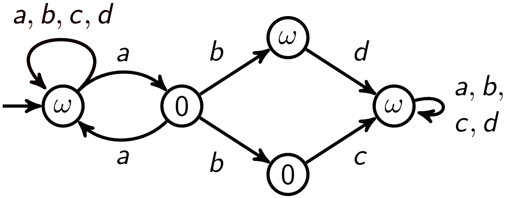

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

Given N identical Markov decision processes, make them all reach ✓ with probability 1.

$S \leftarrow \mathbb{N}^Q$ $\qquad\rightarrow$ configurations

$C \leftarrow \mathbb{N}^Q \times \Sigma$ $\qquad\rightarrow$ commits

while not fixpoint **do**

  if $\exists (s, a) \in C,\ s \xrightarrow{a} s',\ s' \notin S$ **then**

    $C \leftarrow C \setminus \{(s, a)\} \uparrow$

  if $\exists s \in S$, no path from s to F in C **then**

    $S \leftarrow S \setminus \{s\} \uparrow,\ C \leftarrow C \cap S \times \Sigma$

return $I \subseteq S$

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

Given N identical Markov decision processes, make them all reach ✓ with probability 1.

Terminates by well quasi-order argument

$S \leftarrow \mathbb{N}^Q$ $\rightarrow$ configurations

$C \leftarrow \mathbb{N}^Q \times \Sigma$ $\rightarrow$ commits

while not fixpoint **do**

  if $\exists (s, a) \in C,\ s \xrightarrow{a} s',\ s' \notin S$ **then**

    $C \leftarrow C \setminus \{(s, a)\} \uparrow$

  if $\exists s \in S$, no path from s to F in C **then**

    $S \leftarrow S \setminus \{s\} \uparrow,\ C \leftarrow C \cap S \times \Sigma$

return $I \subseteq S$

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

Given N identical Markov decision processes, make them all reach ✔ with probability 1.

$S \leftarrow \mathbb{N}^Q$ $\rightarrow$ configurations

$C \leftarrow \mathbb{N}^Q \times \Sigma$ $\rightarrow$ commits

Terminates by well quasi-order argument

$$\left\{\begin{array}{l}
\textbf{while } \text{not fixpoint } \textbf{do} \\
\quad \textbf{if } \exists (s,a) \in C,\ s \xrightarrow{a} s',\ s' \notin S \textbf{ then} \\
\qquad C \leftarrow C \setminus \{(s,a)\}\uparrow \\
\quad \textbf{if } \exists s \in S, \text{ no path from } s \text{ to } F \text{ in } C \textbf{ then} \\
\qquad S \leftarrow S \setminus \{s\}\uparrow,\ C \leftarrow C \cap S \times \Sigma \\
\textbf{return } I \subseteq S
\end{array}\right.$$

Easy to check

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

Given N identical Markov decision processes, make them all reach ✔ with probability 1.

$S \leftarrow \mathbb{N}^Q$ $\rightarrow$ configurations
$C \leftarrow \mathbb{N}^Q \times \Sigma$ $\rightarrow$ commits

Terminates by well quasi-order argument

while not fixpoint **do**
  if $\exists (s,a) \in C,\ s \xrightarrow{a} s',\ s' \notin S$ **then** — Easy to check
    $C \leftarrow C \setminus \{(s,a)\}\uparrow$
  if $\exists s \in S$, no path from s to F in C **then** — Unclear
    $S \leftarrow S \setminus \{s\}\uparrow,\ C \leftarrow C \cap S \times \Sigma$
return $I \subseteq S$

Controlling a random population [Colcombet Fijalkow Ohlmann '20]

Given N identical Markov decision processes, make them all reach ✔ with probability 1.

$S \leftarrow \mathbb{N}^Q$ $\rightarrow$ configurations

$C \leftarrow \mathbb{N}^Q \times \Sigma$ $\rightarrow$ commits

Terminates by well quasi-order argument

while not fixpoint **do**

  if $\exists (s,a) \in C,\ s \xrightarrow{a} s',\ s' \notin S$ **then** Easy to check

    $C \leftarrow C \setminus \{(s,a)\}\uparrow$

  if $\exists s \in S$, no path from s to F in C **then** Unclear

    $S \leftarrow S \setminus \{s\}\uparrow,\ C \leftarrow C \cap S \times \Sigma$

return $I \subseteq S$

The remaining problem

$$C = ((\omega, 5, \omega, 8), a)\downarrow \cup ((4, \omega, \omega, 4), a)\downarrow \cup ((9, 6, 1, 4), b)\downarrow$$

Is there a path in C from every configuration in $(\omega, 5, \omega, 8)\downarrow$ to $(0, 0, 0, \omega)\downarrow$?

The remaining problem

$$C = ((\omega, 5, \omega, 8), a)\downarrow \cup ((4, \omega, \omega, 4), a)\downarrow \cup ((9, 6, 1, 4), b)\downarrow$$

Is there a path in C from every configuration in $(\omega, 5, \omega, 8)\downarrow$ to $(0, 0, 0, \omega)\downarrow$?

Reachability relation in some infinite-state system

The remaining problem

$$C = ((\omega, 5, \omega, 8), a)\downarrow \cup ((4, \omega, \omega, 4), a)\downarrow \cup ((9, 6, 1, 4), b)\downarrow$$

Is there a path in C from every configuration in $(\omega, 5, \omega, 8)\downarrow$ to $(0, 0, 0, \omega)\downarrow$?

Reachability relation in some infinite-state system

- In EXPSPACE, PSPACE-hard [Colcombet, Fijalkow, Ohlmann '20]
- PSPACE-complete [Gimbert, Mascle, Totzke '26a]

Results

[Colcombet Fijalkow Ohlmann '20]

The Randomised Population Control Problem is decidable.

Results

[Colcombet Fijalkow Ohlmann '20]

The Randomised Population Control Problem is decidable.

This paper

The Randomised Population Control Problem is EXPTIME-complete.

Results

[Colcombet Fijalkow Ohlmann '20]

The Randomised Population Control Problem is decidable.

This paper

The Randomised Population Control Problem is EXPTIME-complete.

- Lower bound by reduction from *countdown games*
- Upper bound by showing that if there is a winning strategy then there is a union of ideals with small constants that is a winning region.

Proof plan

Theorem

If there is a winning strategy then there is a collection $W = (W_a)_{a \in \Sigma}$ such that

- Each W_a is a union of ideals of the form $(b_1, ..., b_{|S|}) \downarrow$ with $b_k \in \{0, \ldots, |S|, \omega\}$
- From everywhere in W there is a path to a final configuration
- We cannot leave W by playing a within W_a

Proof

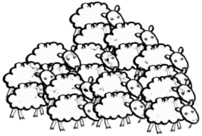

Proof

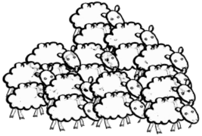

Proof

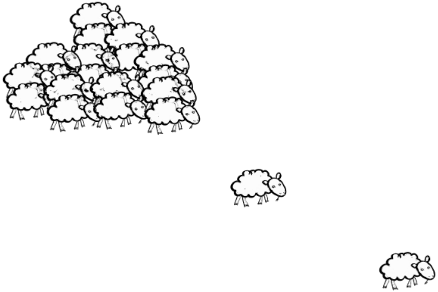

Proof

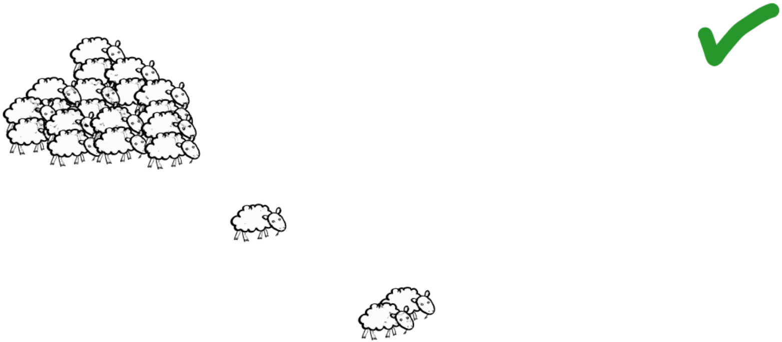

Proof

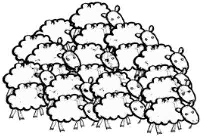

Proof

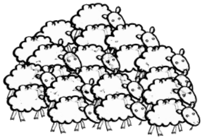

Proof

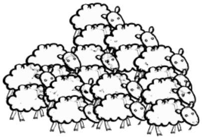

Proof

Proof

Proof

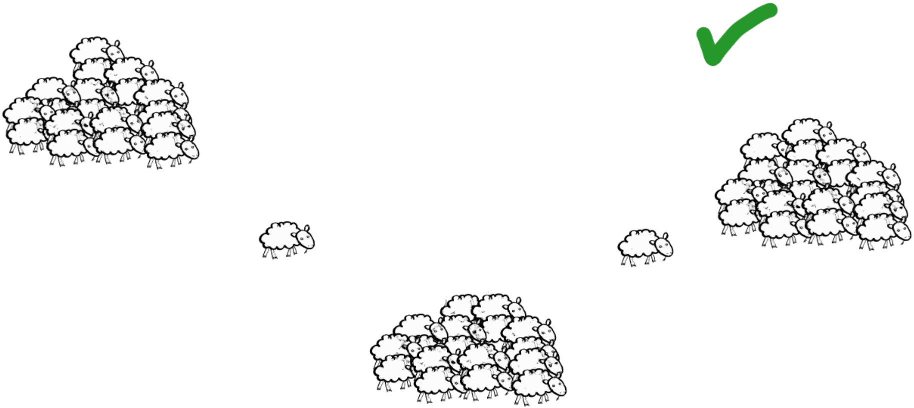

Proof

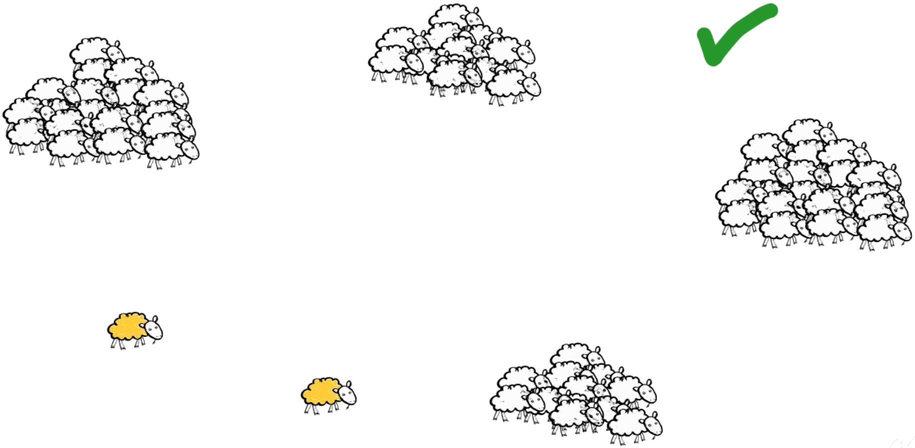

Proof

Proof

Proof

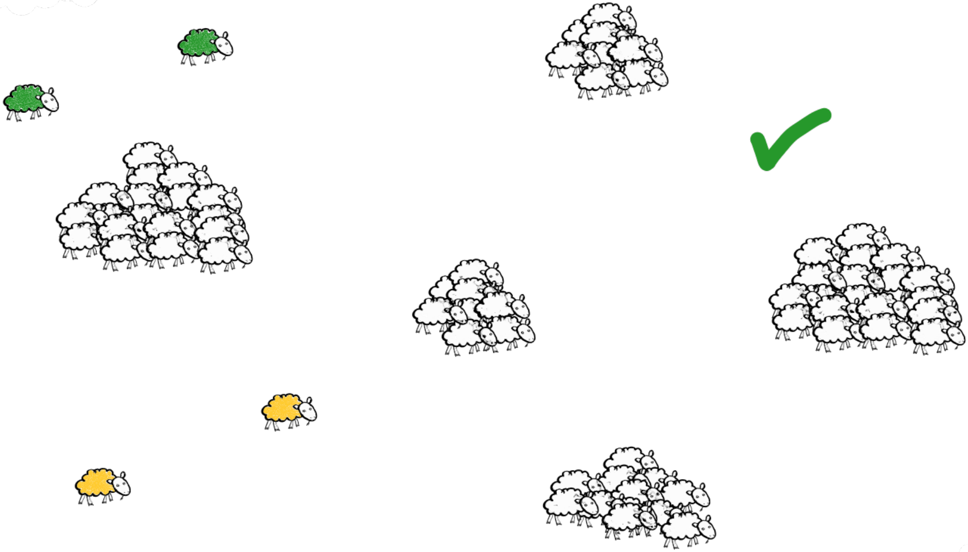

Proof

Proof

Proof

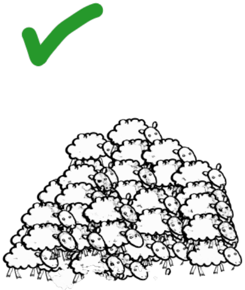

Future work

- Quantitative objectives
- Infinite-horizon specifications

Future work

- Quantitative objectives
- Infinite-horizon specifications
- **Fast strategies**

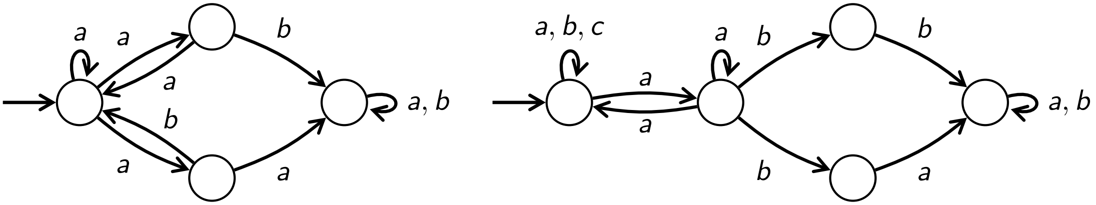

Future work

- Quantitative objectives
- Infinite-horizon specifications
- **Fast strategies**

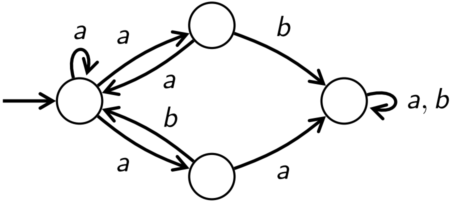

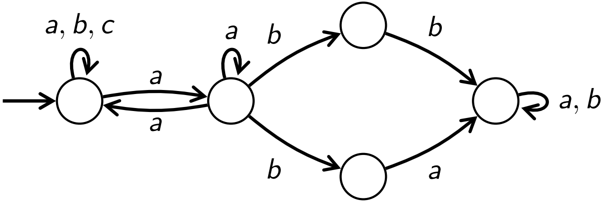

Thank you for your attention!