Computing Probabilistic Bisimilarity Distances via Policy Iteration

Franck van Breugel

(joint work with Qiyi Tang)

October 4, 2016

Table of Contents

Fundamental question

Do two states of a systems behave the same?

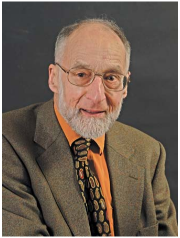

Robin Milner introduced bisimilarity, the most well-known behavioural equivalence, in 1979.

Behavioural Equivalences

Fundamental question

Do two states of a systems behave the same?

Behavioural equivalence is an equivalence relation.

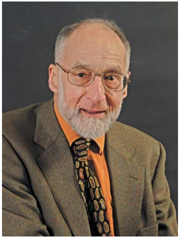

Robin Milner introduced bisimilarity, the most well-known behavioural equivalence, in 1979.

Model of probabilistic system

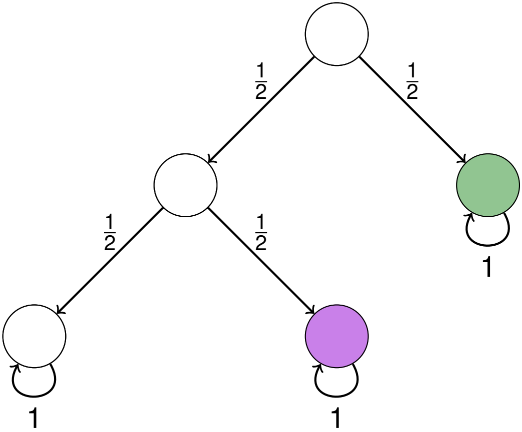

Labelled Markov chain

Andrey Markov produced the first results for Markov chains in 1906.

Transitions

$$\tau \in S \to Dist(S)$$

For each state s, the transitions of s are presented by a probability distribution $\tau(s)$ on S.

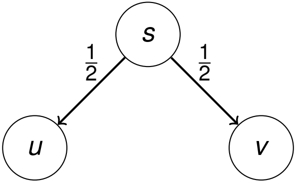

$$\tau(s)(w) = \begin{cases} \frac{1}{2} & \text{if } w = u \\ \frac{1}{2} & \text{if } w = v \\ 0 & \text{otherwise} \end{cases}$$

Probabilistic bisimulation

Definition

An equivalence relation $\mathcal{R}$ is a *probabilistic bisimulation* if for all $(s, t) \in \mathcal{R}$,

- $\ell(s) = \ell(t)$ and
- $(\tau(s), \tau(t)) \in \bar{\mathcal{R}}$.

Definition

Probabilistic bisimilarity is the largest probabilistic bisimulation.

Kim Larsen and Arne Skou introduced probabilistic bisimilarity in 1989.

Lifting

Definition

Let $\mathcal{R} \subseteq S \times S$ be an equivalence relation. The *lifting* of $\mathcal{R}$, $\bar{\mathcal{R}} \subseteq Dist(S) \times Dist(S)$, is defined by

$$(\mu, \nu) \in \bar{\mathcal{R}} \text{ if } \mu([s]) = \nu([s]) \text{ for all } s \in S$$

Next, we will provide an alternative characterization of lifting.

Coupling

Definition

A *coupling* of probability distributions μ and ν on S is a probability distribution ω on $S \times S$ with marginals μ and ν, that is, for all u, $v \in S$,

$$\sum_{v \in S} \omega(u, v) = \mu(u)$$

$$\sum_{u \in S} \omega(u, v) = \nu(v)$$

Wolfgang Doeblin introduced the notion of a coupling in 1936 (published in 1938).

The set of couplings of μ and ν is denoted by $\Omega(\mu, \nu)$.

Coupling

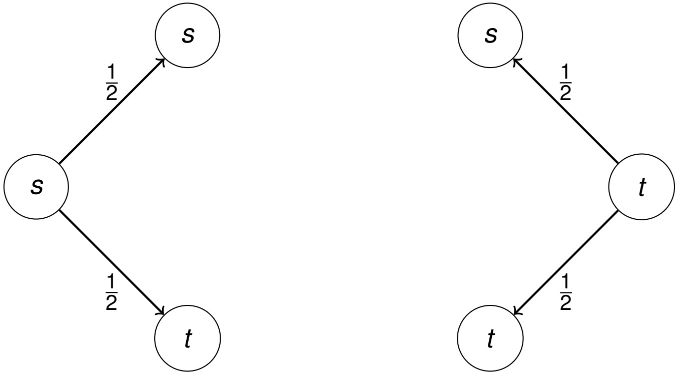

Coupling

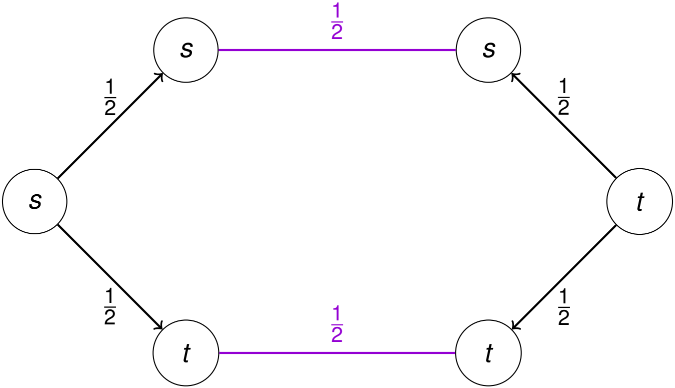

Coupling

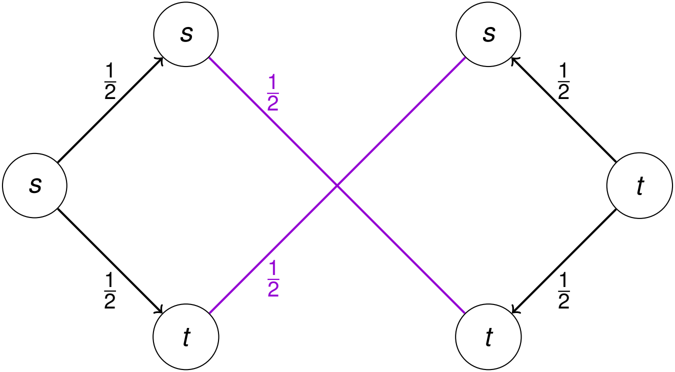

Theorem

Let $\mathcal{R} \subseteq S \times S$ be an equivalence relation.

$$(\mu, \nu) \in \bar{\mathcal{R}} \text{ iff } \exists \omega \in \Omega(\mu, \nu) : \text{support}(\omega) \subseteq \mathcal{R}$$

Bengt Jonsson and Kim Larsen provided the alternative characterization in 1991.

Coupling

There are infinitely many couplings ($r \in [0, \frac{1}{2}]$).

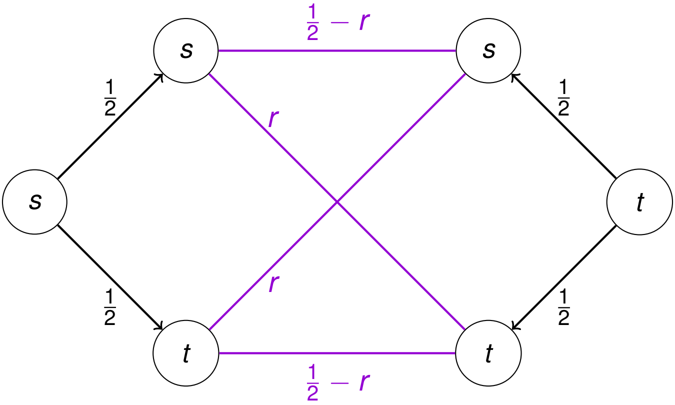

Proposition

$\Omega(\tau(s), \tau(t))$ is a convex polytope.

Proposition

A concave function on a convex polytope attains its minimum at a vertex.

Proposition

The set $V(\Omega(\tau(s), \tau(t)))$ of vertices of $\Omega(\tau(s), \tau(t))$ is finite.

Coupling

There are two vertices ($r \in \{0, \frac{1}{2}\}$).

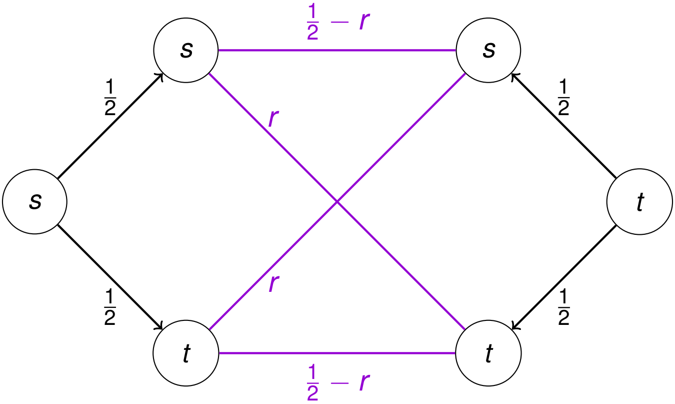

Alternative characterization of lifting

Theorem (TB 2016)

Let $\mathcal{R} \subseteq S \times S$ be an equivalence relation.

$$(\mu, \nu) \in \bar{\mathcal{R}} \text{ iff } \exists \omega \in V(\Omega(\mu, \nu)) : \text{support}(\omega) \subseteq \mathcal{R}$$

Proof sketch

- Order the states $s_1, \ldots, s_n$ such that equivalent states are consecutive.
- Apply the North-West corner method.
- Prove some loop invariants (by means of Dafny).

Fundamental problem

Behavioural equivalences are not robust for systems with real-valued data.

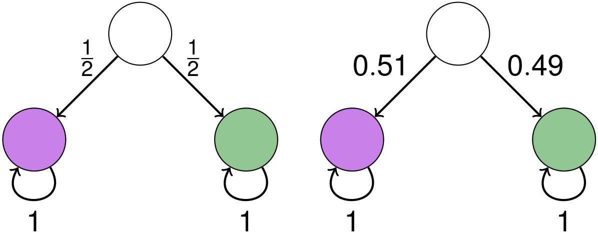

Alessandro Giacalone, Chi-Chang Jou and Scott Smolka observed that probabilistic bisimilarity, the most well-known behavioural equivalence for probabilistic systems, is not robust in 1990.

Behavioural pseudometric

Fundamental problem

Behavioural equivalences are not robust for systems with real-valued data.

Robust alternative

Instead of an equivalence relation

$$\sim : S \times S \to \{\text{true}, \text{false}\}$$

use a pseudometric

$$d : S \times S \to [0, 1].$$

Probabilistic bisimilarity

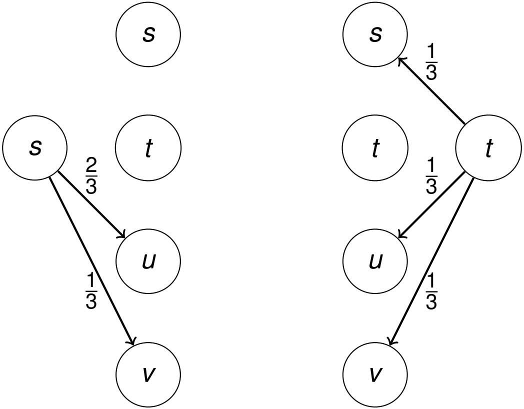

$$\mathrm{support}(\omega) \subseteq \mathcal{R}$$

Probabilistic bisimilarity

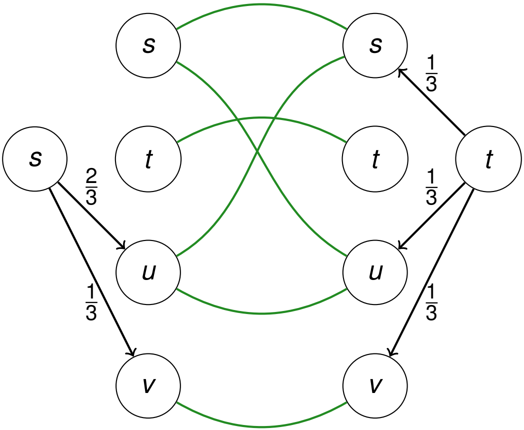

$$\text{support}(\omega) \subseteq \mathcal{R}$$

Probabilistic bisimilarity

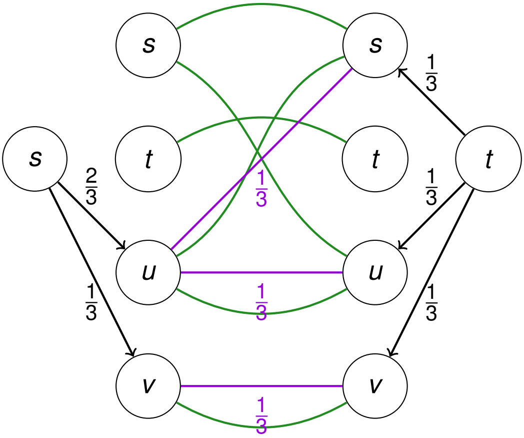

$$\text{support}(\omega) \subseteq \mathcal{R}$$

Let us represent the equivalence relation $\mathcal{R}$ with the following distance function.

$$r(s,t) = \begin{cases} 0 & \text{if } (s,t) \in \mathcal{R} \\ 1 & \text{otherwise} \end{cases}$$

Then the condition

$$\text{support}(\omega) \subseteq \mathcal{R}$$

is equivalent to

$$\sum_{u,v \in S} \omega(u,v)\, r(u,v) = 0$$

Quantitative generalization of probabilistic bisimilarity

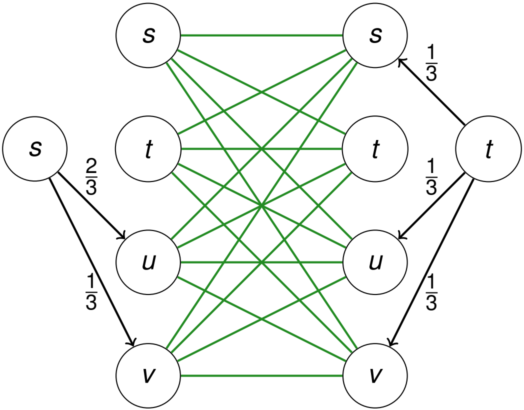

$$\text{minimize} \sum_{u,v \in S} \omega(u, v)\ d(u, v)$$

Quantitative generalization of probabilistic bisimilarity

Definition

Probabilistic bisimilarity is the largest equivalence relation $\sim$ such that $s \sim t$ implies

- $\ell(s) = \ell(t)$ and
- $\exists \omega \in V(\Omega(\tau(s), \tau(t))) : \mathrm{support}(\omega) \subseteq \sim$.

Definition

The *probabilistic bisimilarity pseudometric* is the smallest $d : S \times S \to [0, 1]$ such that

$$d(s, t) = \begin{cases} 1 & \text{if } \ell(s) \neq \ell(t) \\ \min\limits_{\omega \in V(\Omega(\tau(s), \tau(t)))} \sum\limits_{u,v \in S} \omega(u, v)\, d(u, v) & \text{otherwise} \end{cases}$$

Probabilistic bisimilarity pseudometric

Josee Desharnais, Vineet Gupta, Radha Jagadeesan and Prakash Panangaden. Metrics for Labeled Markov Systems. CONCUR 1999.




Theorem (DGJP 1999)

$s \sim t$ if and only if $d(s, t) = 0$.

Kantorovich metric

Let $\mu, \nu \in \mathit{Dist}(S)$ and $d : S \times S \to [0, 1]$.

$$
\begin{aligned}
& \max_{f \in (S,d) \rightarrow [0,1]} \sum_{s \in S} f(s)\, (\mu(s) - \nu(s)) \\
= \quad & \min_{\omega \in \Omega(\mu,\nu)} \sum_{u.v \in S} \omega(u, v)\, d(u, v) \\
= \quad & \min_{\omega \in V(\Omega(\mu,\nu))} \sum_{u.v \in S} \omega(u, v)\, d(u, v)
\end{aligned}
$$

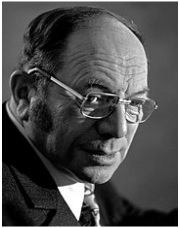

Leonid Kantorovich first published this metric in 1942.

Table of Contents

Algorithm

Question

How to compute the probabilistic bisimilarity distances for a labelled Markov chain?

Algorithm to compute the bisimilarity distances

- Express $d(s, t) < q$ in the first order theory over the reals.
- Use the binary search method to approximate $d(s, t)$.

Babita Sharma, Franck van Breugel and James Worrell. Approximating a Behavioural Pseudometric without Discount for Probabilistic Systems. FoSSaCS 2007.

Alfred Tarski showed that the first order theory over the reals is decidable in 1948.

Algorithm to compute the bisimilarity distances

- Express $d(s, t)$ as a linear program.
- Use the ellipsoid method to compute $d(s, t)$.
  - As separation algorithm, to solve a minimum cost flow problem, use the network simplex algorithm.

Di Chen, Franck van Breugel and James Worrell. On the Complexity of Computing Probabilistic Bisimilarity. FoSSaCS 2012.

Leonid Khachiyan proved the polynomial-time solvability of linear programs in 1979.

Algorithm to compute the bisimilarity distances

Giorgio Bacci, Giovanni Bacci, Kim Larsen and Radu Mardare. On-the-Fly Exact Computation of Bisimilarity Distances. TACAS 2013.




$\text{B}^2\text{LM algorithm} = \underbrace{\text{basic algorithm}}_{\text{this talk}} + \underbrace{\text{optimization}}_{\text{on-the-fly}}$

Simple stochastic game (SSG)

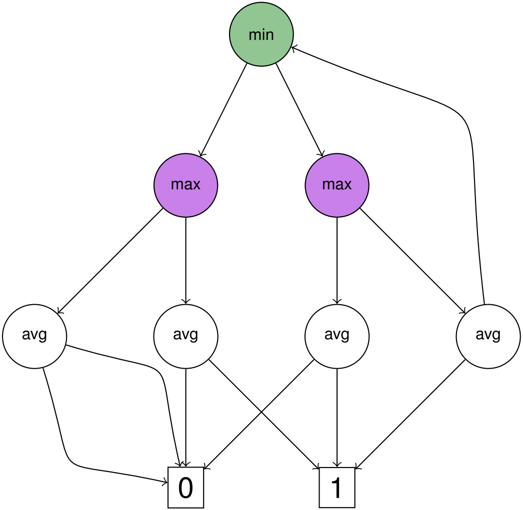

Anne Condon was the first to study simple stochastic games from a computational point of view in 1992.

Definition

The *value* of a vertex is the probability that the max player wins the game (reaches 1) provided that both players use optimal strategies (the min player tries not to reach 1).

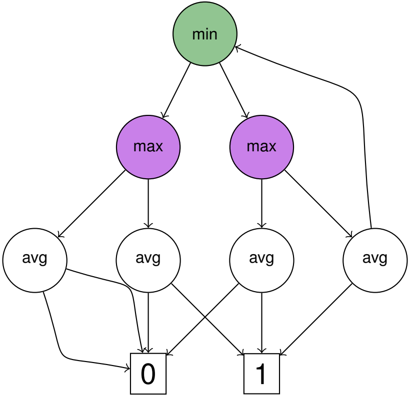

From LMCs to SSGs

For each labelled Markov chain we construct a corresponding simple stochastic game such that

LMC	SSG
distance	value
algorithm	simple policy iteration

Ronald Howard introduced policy iteration in 1958.

From LMCs to SSGs

With every pair of states (s, t) of the LMC we associate a vertex of the SSG.

- If $\ell(s) \neq \ell(t)$ then $d(s, t) = 1$.

$$\boxed{1}$$

- If $s \sim t$ then $d(s, t) = 0$.

$$\boxed{0}$$

- Otherwise,

$$d(s,t) = \min_{\omega \in V(\Omega(\tau(s),\tau(t)))} \sum_{u,v \in S} \omega(u,v)\, d(u,v)$$

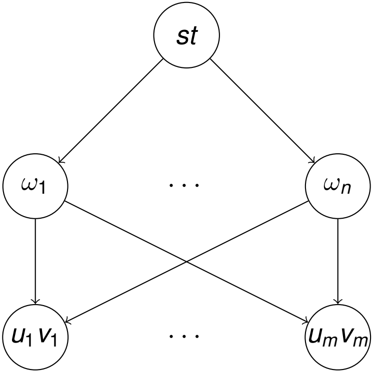

From LMCs to SSGs

- Otherwise,

$$d(s,t) = \min_{\omega \in V(\Omega(\tau(s),\tau(t)))} \sum_{u,v \in S} \omega(u,v)\, d(u,v)$$

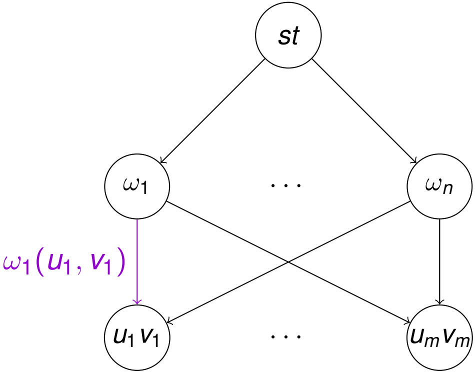

Correctness of simple policy iteration

Simple policy iteration

```
choose a random initial policy
while exists a vertex which is not locally optimal
  adjust the policy at that vertex
```

Theorem (Condon 1992)

Simple policy iteration computes the value function if the simple stochastic game terminates with probability one (no matter which strategy the players use).

Proposition

If we do not map a pair of probabilistic bisimilar states to a zero sink, then the resulting simple stochastic game may not terminate with probability one.

Correctness of simple policy iteration

The labelled Markov chain

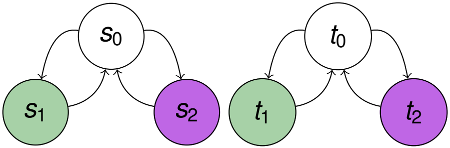

is mapped to the simple stochastic game

$$3\times \boxed{0} + 6\times \boxed{1}$$

Proposition

This simple stochastic game terminates with probability one.

Correctness of simple policy iteration

If we do not map a pair of probabilistic bisimilar states to a zero sink, the labelled Markov chain

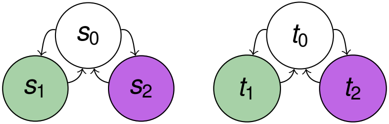

is mapped to the simple stochastic game

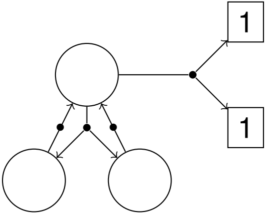

Proposition

This simple stochastic game does not terminate with probability one.

B^2LM algorithm is exponential

Simple policy iteration

```
choose a random initial policy
while exists a vertex which is not locally optimal
  adjust the policy at that vertex
```

Theorem

For each $n \in \mathbb{N}$, there exists a labelled Markov chain of size $O(n)$ such that simple policy iteration takes $\Omega(2^n)$ iterations.

Proof idea: Implement an "*n*-bit counter."

B^{2}LM algorithm is exponential

We start with the following labelled Markov chain.

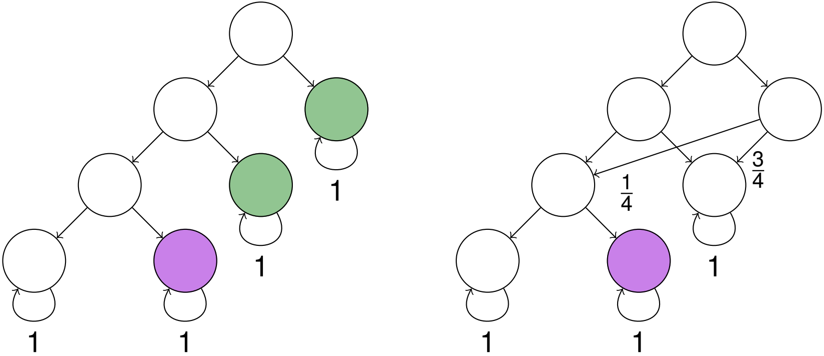

B^{2}LM algorithm is exponential

The labelled Markov chain corresponds to the following simple stochastic game.

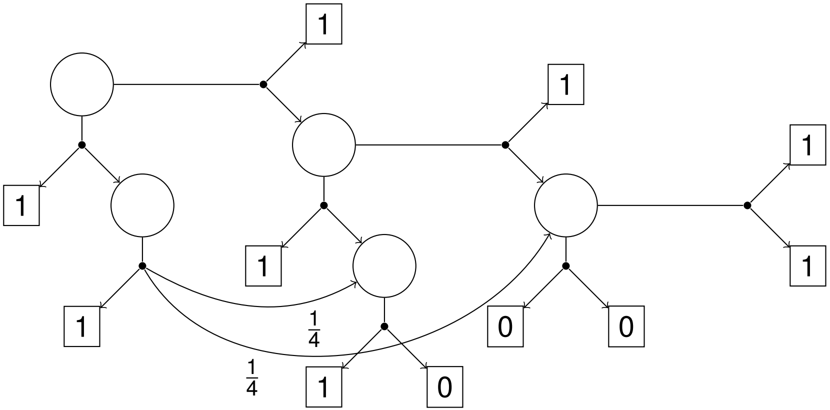

B^2LM algorithm is exponential

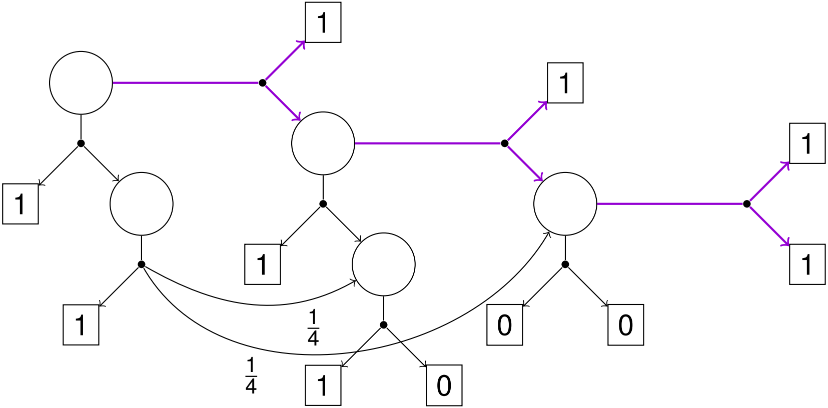

000

B^2LM algorithm is exponential

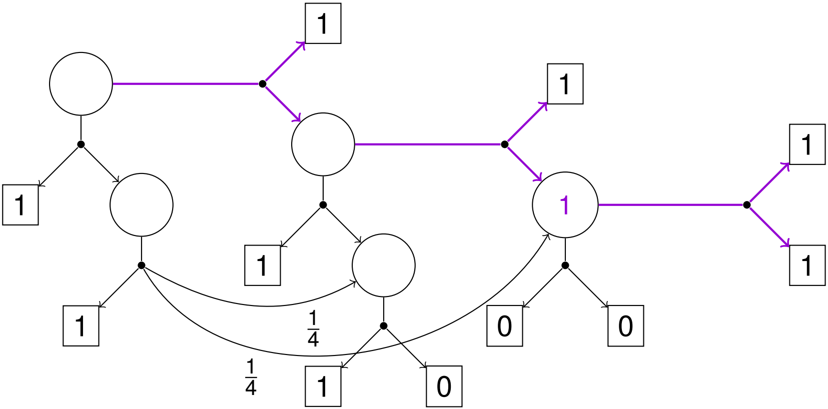

000

B^2LM algorithm is exponential

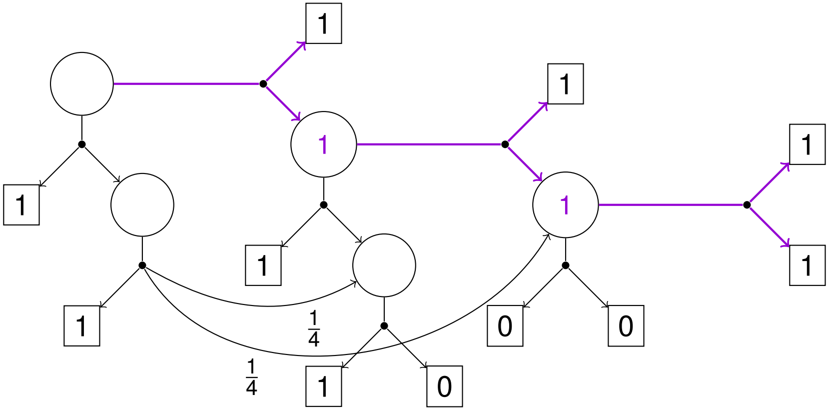

000

B^2LM algorithm is exponential

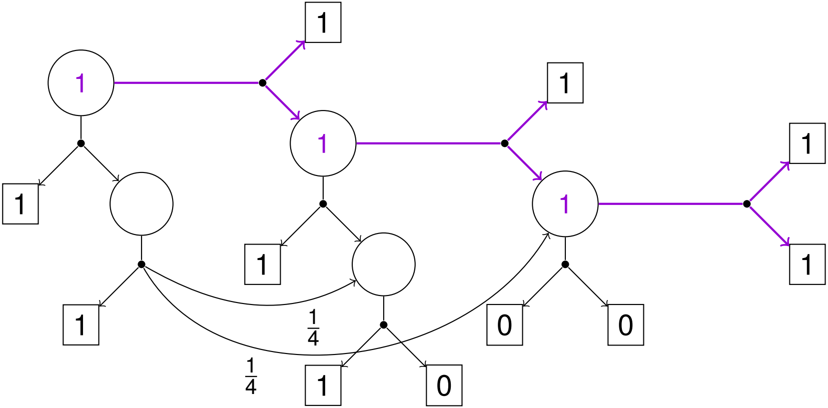

000

B^{2}LM algorithm is exponential

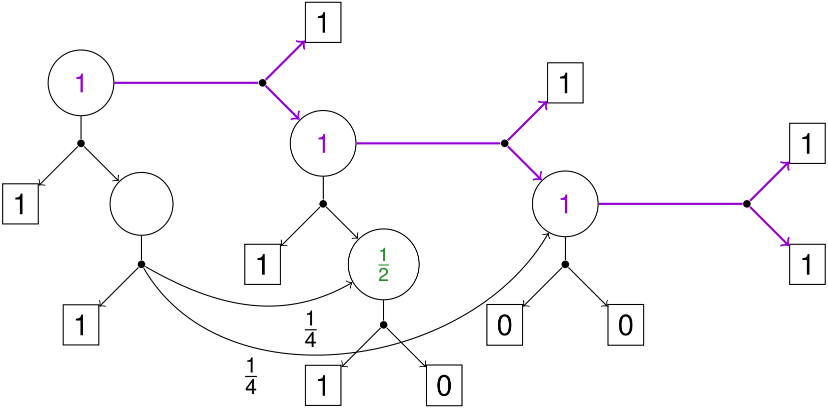

000

B^2LM algorithm is exponential

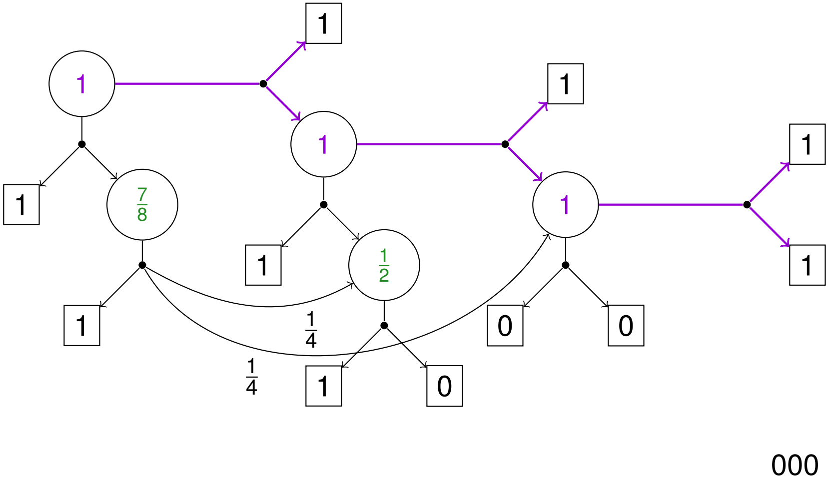

B^2LM algorithm is exponential

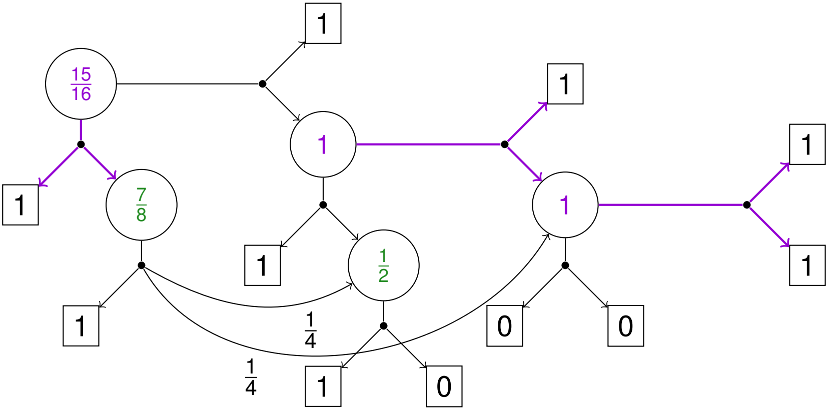

100

B^2LM algorithm is exponential

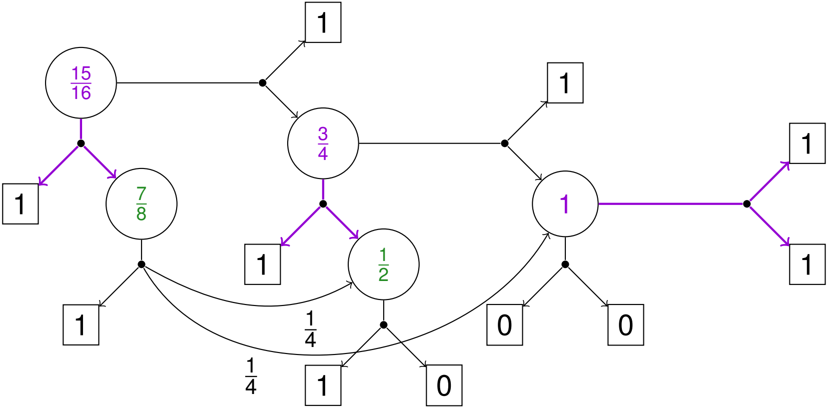

110

B^2LM algorithm is exponential

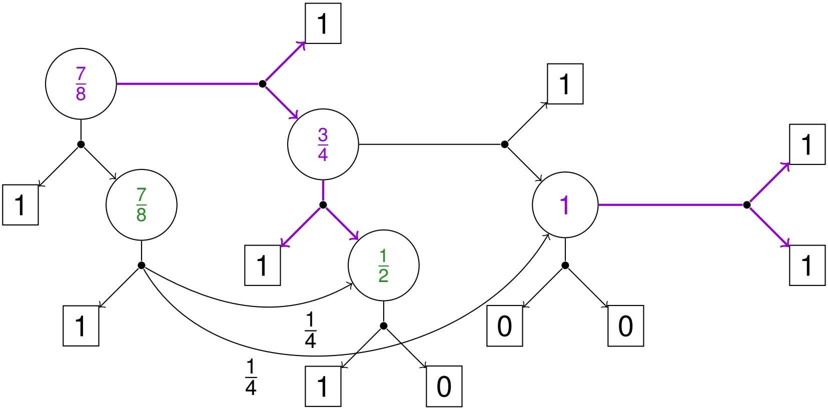

010

B^2LM algorithm is exponential

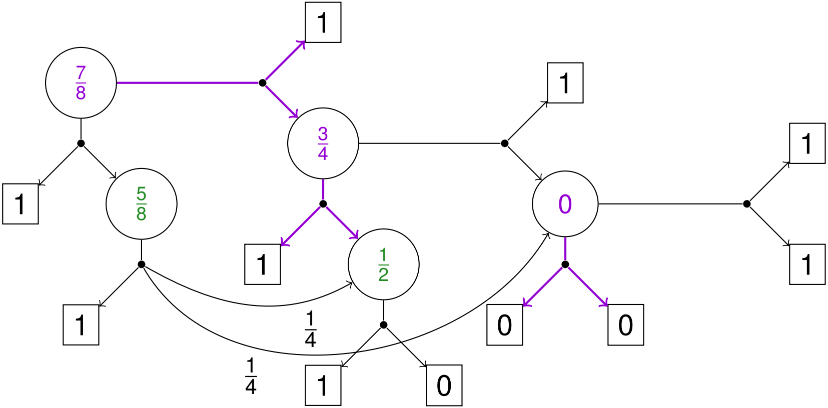

011

B^2LM algorithm is exponential

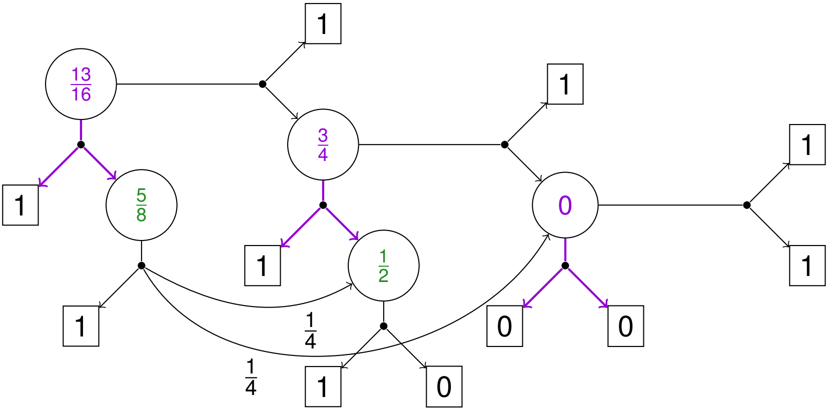

111

B^2LM algorithm is exponential

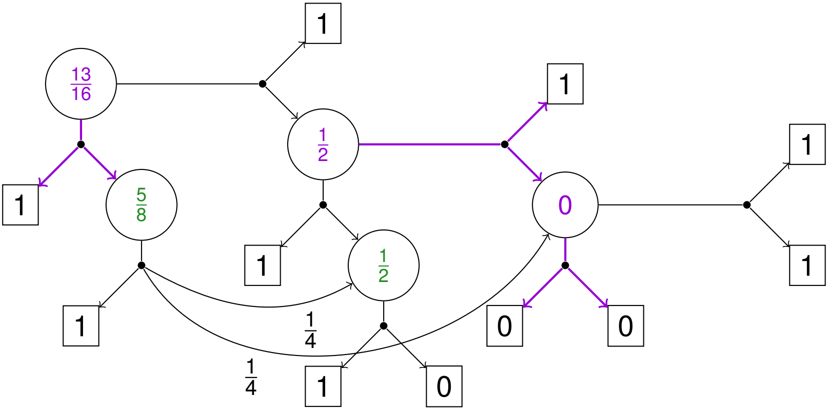

101

B^{2}LM algorithm is exponential

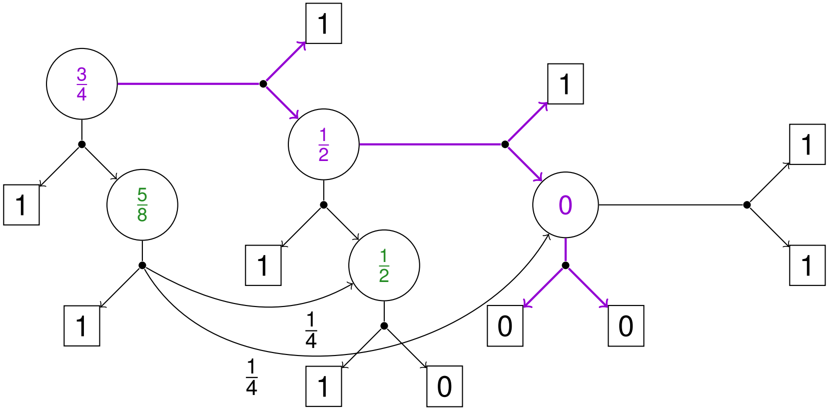

001

Performance comparison

Aron Itai and Michael Rodeh. Symmetry breaking in distributed networks. Information and Computation, 88(1):60–87, 1990.

- SBW: –
- CBW: more than 10 hours for $N = 3$ and $K = 2$
- B^2LM:

		without bisimilarity		with bisimilarity	
N	K	μ	σ	μ	σ
3	2	4.02	0.15	2.70	0.08
4	2	478.67	1.49	399.833	0.82
3	3	1151.10	0.36	753.73	0.58
5	2	62126.78	1264.50	58862.35	1512.58

Conclusion

- The basic B^2LM algorithm is simple policy iteration.
  - To define the simple stochastic game we need to decide probabilistic bisimilarity.
- In the worst case, the (basic) B^2LM algorithm is exponential.
- In practice, the (basic) B^2LM algorithm performs much better than all other algorithms.

Related and future work

- Use general policy iteration to compute the probabilistic bisimilarity distances for labelled Markov chains.
- Use (simple/general) policy iteration to compute the probabilistic bisimilarity distances for probabilistic automata.