Decisiveness for countable MDPs and insights for NPLCSs and POMDPs

Patricia Bouyer

LMF, Université Paris-Saclay,
CNRS, ENS Paris-Saclay
France

Joint work with Nathalie Bertrand (Inria, France), Thomas Brihaye (UMONS, Belgium), Paulin Fournier and Pierre Vandenhove (UMONS, Belgium)

Work partly supported by ANR projects MAVeriQ and BisoUS

Back to 2006/2007…

- Time, probabilities and games join force!
  - My first steps in probabilistic systems
- Trips to Dresden and Oxford

Back to 2006/2007…

- Time, probabilities and games join force!
    - My first steps in probabilistic systems
- Trips to Dresden and Oxford


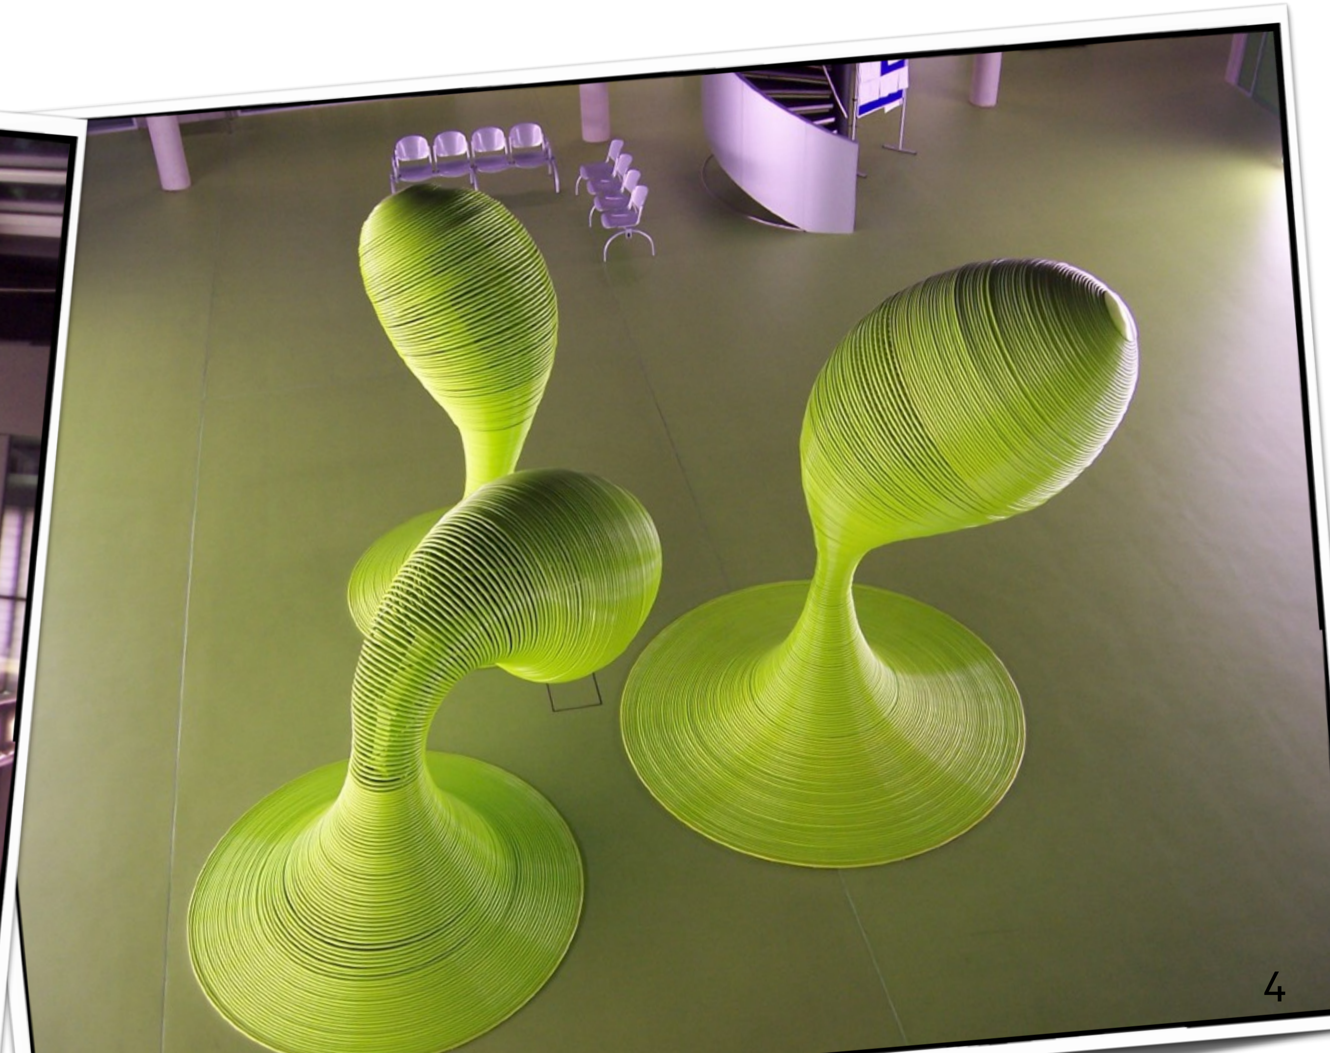




Almost-Sure Model Checking of Infinite Paths in One-Clock Timed Automata

Christel Baier[1] Nathalie Bertrand[2] Patricia Bouyer[3,*] Thomas Brihaye[4] Marcus Größer[1]

[1] *Technische Universität Dresden, Germany*
{baier,groesser}@tcs.inf.tu-dresden.de

[2] *IRISA, INRIA Rennes, France*
nathalie.bertrand@irisa.fr

[3] *LSV, CNRS & ENS Cachan, France*
bouyer@lsv.ens-cachan.fr

[4] *Université de Mons-Hainaut, Belgium*
thomas.brihaye@umh.ac.be

Probabilistic and Topological Semantic[s] for Timed Automata

Christel Baier[1], Nathalie Bertrand[1,*], Patricia Bouyer[2,3,**], Thomas Brihaye[2], and Marcus Größer[1]

[1] Technische Universität Dresden, Germany
[2] LSV - CNRS & ENS Cachan, France
[3] Oxford University, England

When Are Timed Automata Determinizable?

Christel Baier[1], Nathalie Bertrand[2], Patricia Bouyer[3], and Thomas Brihaye[4]

[1] Technische Universität Dresden, Germany
[2] INRIA Rennes Bretagne Atlantique, France
[3] LSV, CNRS & ENS Cachan, France
[4] Université de Mons, Belgium

STOCHASTIC TIMED AUTOMATA

NATHALIE BERTRAND [a], PATRICIA BOUYER [b], THOMAS BRIHAYE [c], QUENTIN ME[NET d], CHRISTEL BAIER [e], MARCUS GRÖSSER [f], AND MARCIN JURDZIŃSKI [g]

[a] Inria Rennes, France
e-mail address: nathalie.bertrand@inria.fr

[b] LSV, CNRS & ENS Cachan, France
e-mail address: bouyer@lsv.ens-cachan.fr

[c,d] Université de Mons, Belgium
e-mail address: {thomas.brihaye,quentin.menet}@umons.ac.be

[e,f] TU Dresden, Germany
e-mail address: {baier,groesser}@tcs.inf.tu-dresden.de

[g] University of Warwick, UK
e-mail address: mju@dcs.warwick.ac.uk

Decisiveness

A successful concept for Markov chains

Markov chains

(Discrete-time) Markov chains

$\mathscr{C} = (S, s_0, \delta)$ with S at most denumerable, $s_0 \in S$ and $\delta : S \rightarrow \mathrm{Dist}(S)$

Markov chains

(Discrete-time) Markov chains

$\mathscr{C} = (S, s_0, \delta)$ with S at most denumerable, $s_0 \in S$ and $\delta : S \to \mathrm{Dist}(S)$

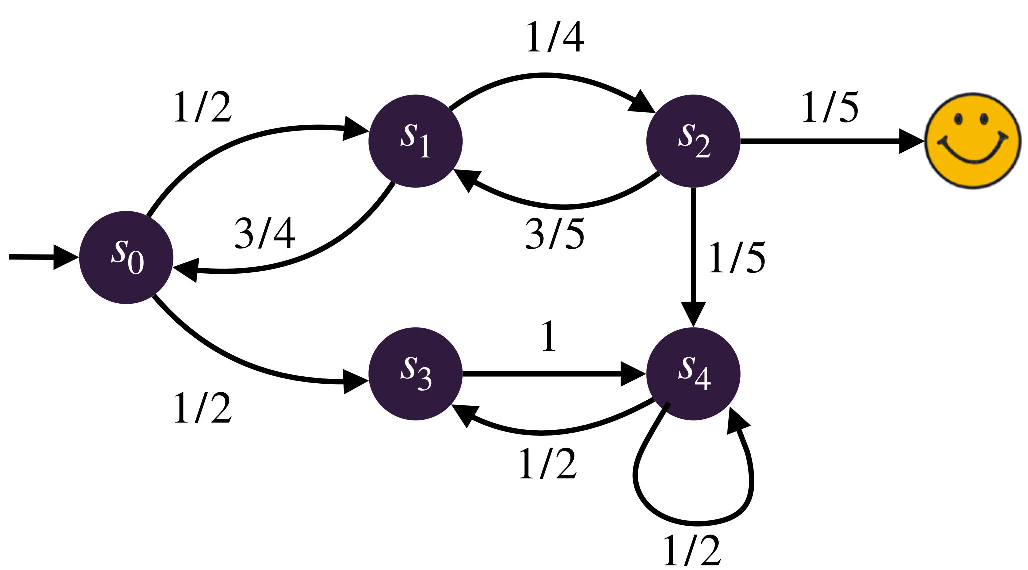

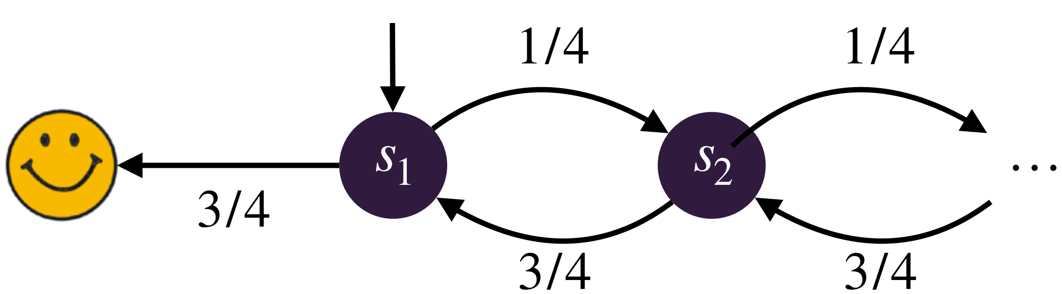

Markov chains

(Discrete-time) Markov chains

$\mathscr{C} = (S, s_0, \delta)$ with S at most denumerable, $s_0 \in S$ and $\delta : S \to \mathrm{Dist}(S)$

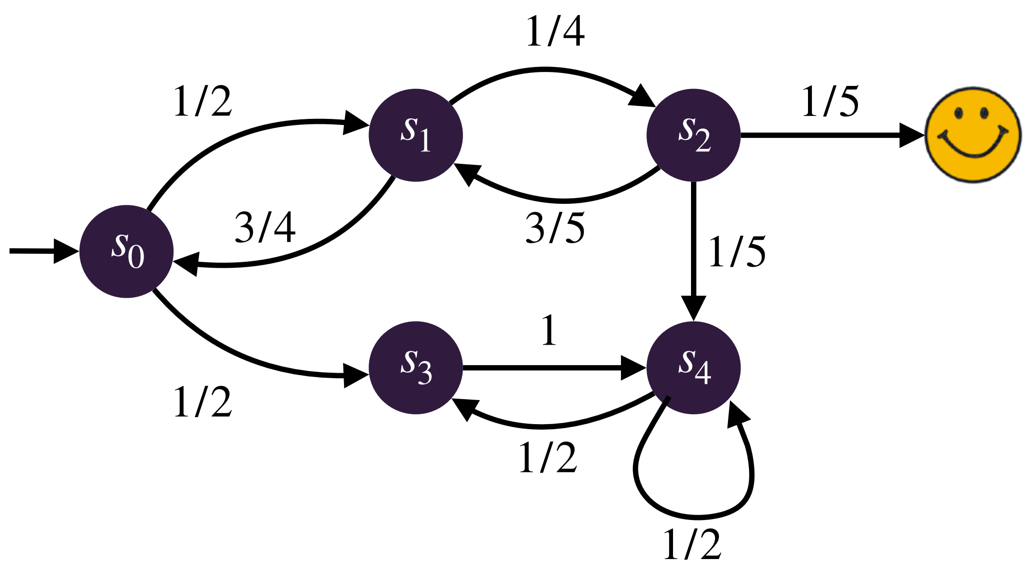

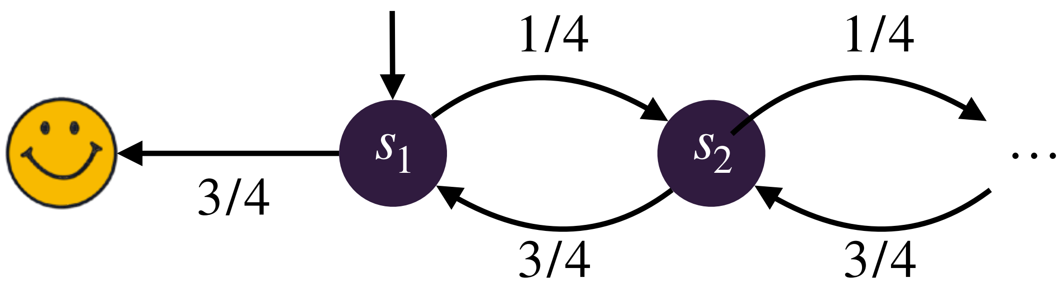

Compute $\mathbb{P}_{\mathscr{C},s_0}\left(\mathbf{F}\ ☺\right)$

Decisiveness

☹ $= \{s \in S \mid s \not\models \exists \mathbf{F}$ ☺ $\}$

Decisiveness

A DTMC $\mathscr{C}$ is decisive from s w.r.t. ☺ if $\mathbb{P}_s(\mathbf{F}$ ☺ $\vee\, \mathbf{F}$ ☹ $) = 1$

[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Decisiveness

☹ $= \{s \in S \mid s \not\models \exists \mathbf{F}$ ☺$\}$

Decisiveness

A DTMC $\mathscr{C}$ is *decisive* from s w.r.t. ☺ if $\mathbb{P}_s(\mathbf{F}$ ☺ $\vee\ \mathbf{F}$ ☹$) = 1$

- Examples of decisive Markov chains: finite Markov chains, probabilistic lossy channel systems, probabilistic VASS, noisy Turing machines, ...

[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Decisiveness

$\frownie = \{ s \in S \mid s \not\models \exists \mathbf{F} \smiley \}$

Decisiveness

A DTMC $\mathscr{C}$ is decisive from s w.r.t. $\smiley$ if $\mathbb{P}_s(\mathbf{F} \smiley \vee \mathbf{F} \frownie) = 1$

- Examples of decisive Markov chains: finite Markov chains, probabilistic lossy channel systems, probabilistic VASS, noisy Turing machines, ...
- Example/counterexample:

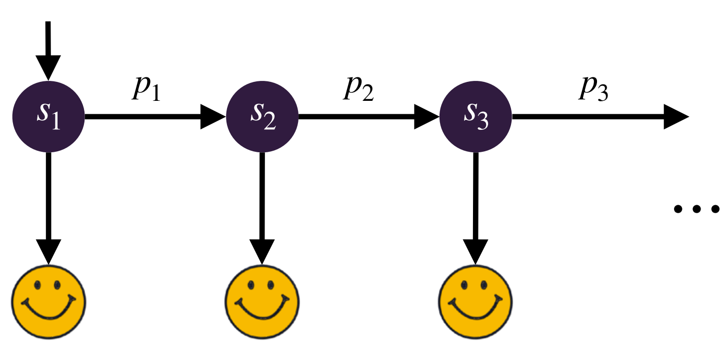

- $\mathbb{P}(\mathbf{G} \neg \smiley) = \prod_{i \geq 1} p_i$
- Decisive iff this product equals 0

[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Decisiveness

🙁 $= \{s \in S \mid s \not\models \exists \mathbf{F}$ 🙂$\}$

Decisiveness

A DTMC $\mathscr{C}$ is decisive from s w.r.t. 🙂 if $\mathbb{P}_s(\mathbf{F}$ 🙂 $\vee \mathbf{F}$ 🙁$) = 1$

- Examples of decisive Markov chains: finite Markov chains, probabilistic lossy channel systems, probabilistic VASS, noisy Turing machines, ...
- Example/counterexample:

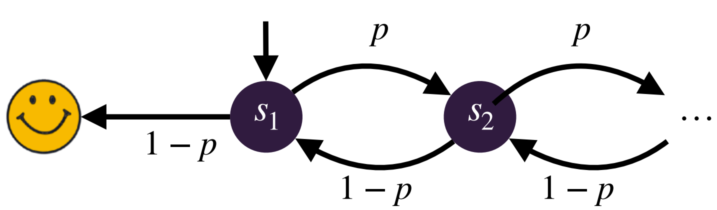

- Recurrent random walk ($p \leq 1/2$): decisive
- Transient random walk ($p > 1/2$): not decisive

[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Approximation scheme

- Aim: compute probability of $\mathbf{F}$ 
-  $= \{ s \in S \mid s \not\models \exists \mathbf{F}$  $\}$

[IN97] P. Iyer, M. Narasimha. Probabilistic lossy channel systems (TAPSOFT'97)
[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Approximation scheme

- Aim: compute probability of $\mathbf{F}$ 😊
- ☹ $= \{s \in S \mid s \not\models \exists \mathbf{F}$ 😊 $\}$

Approximation scheme

Given $\varepsilon > 0$, for every n, compute:

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}(\mathbf{F}_{\leq n}\ \text{😊}) \\ p_n^{\text{no}} = \mathbb{P}(\mathbf{F}_{\leq n}\ \text{☹}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

[IN97] P. Iyer, M. Narasimha. Probabilistic lossy channel systems (TAPSOFT'97)
[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Approximation scheme

- Aim: compute probability of $\mathbf{F}$
- $= \{ s \in S \mid s \not\models \exists \mathbf{F} \}$

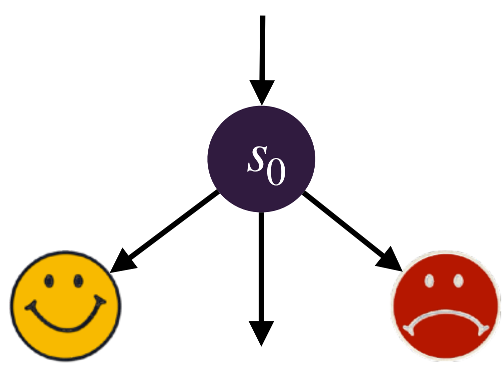

Approximation scheme

Given $\varepsilon > 0$, for every n, compute:

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}(\mathbf{F}_{\leq n}) \\ p_n^{\text{no}} = \mathbb{P}(\mathbf{F}_{\leq n}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$$p_1^{\text{yes}} \leq \mathbb{P}(\mathbf{F}) \leq 1 - p_1^{\text{no}}$$

[IN97] P. Iyer, M. Narasimha. Probabilistic lossy channel systems (TAPSOFT'97)
[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Approximation scheme

- Aim: compute probability of $\mathbf{F}$ ☺
- ☹ $= \{s \in S \mid s \not\models \exists \mathbf{F}$ ☺ $\}$

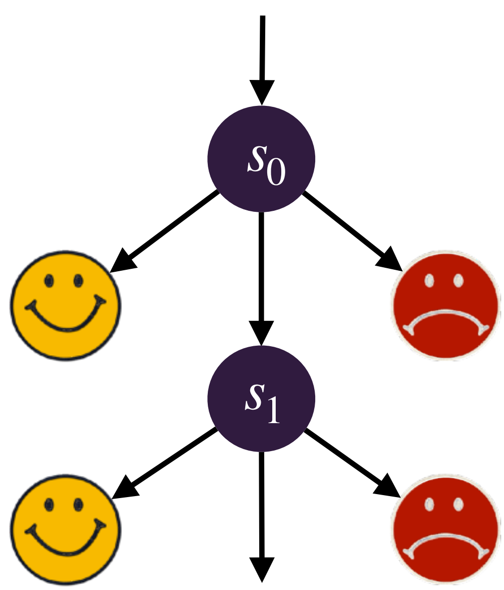

Approximation scheme

Given $\varepsilon > 0$, for every n, compute:

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}(\mathbf{F}_{\leq n} ☺) \\ p_n^{\text{no}} = \mathbb{P}(\mathbf{F}_{\leq n} ☹) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$$p_1^{\text{yes}} \leq \mathbb{P}(\mathbf{F}☺) \leq 1 - p_1^{\text{no}}$$

$$\text{IΛ} \qquad\qquad \text{VI}$$

$$p_2^{\text{yes}} \leq \mathbb{P}(\mathbf{F}☺) \leq 1 - p_2^{\text{no}}$$

[IN97] P. Iyer, M. Narasimha. Probabilistic lossy channel systems (TAPSOFT'97)
[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Approximation scheme

- Aim: compute probability of $\mathbf{F}$ ☺
- ☹ $= \{s \in S \mid s \not\models \exists \mathbf{F}$ ☺ $\}$

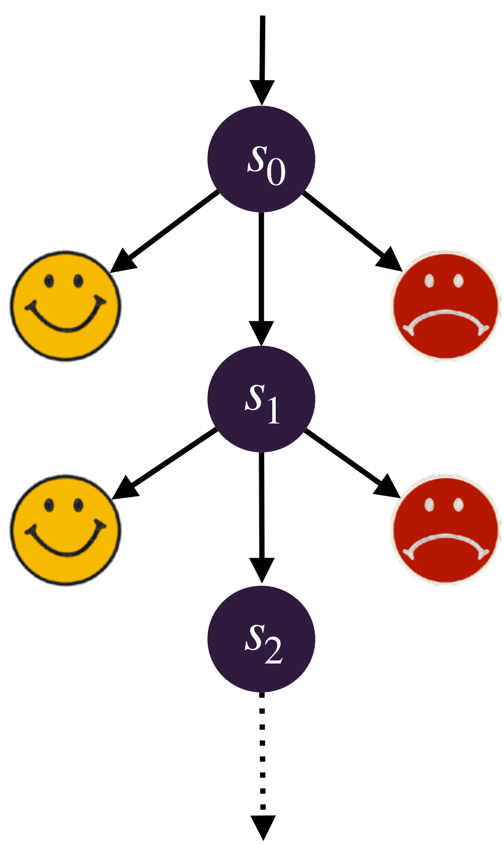

Approximation scheme

Given $\varepsilon > 0$, for every n, compute:

$$\begin{cases} p_n^{\text{yes}} & = & \mathbb{P}(\mathbf{F}_{\leq n} \text{☺}) \\ p_n^{\text{no}} & = & \mathbb{P}(\mathbf{F}_{\leq n} \text{☹}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$$\begin{array}{ccccc} p_1^{\text{yes}} & \leq & \mathbb{P}(\mathbf{F}\text{☺}) & \leq & 1 - p_1^{\text{no}} \\ \rotatebox{90}{$\geq$} & & & & \rotatebox{90}{$\leq$} \\ p_2^{\text{yes}} & \leq & \mathbb{P}(\mathbf{F}\text{☺}) & \leq & 1 - p_2^{\text{no}} \\ \rotatebox{90}{$\geq$} & & \vdots & & \rotatebox{90}{$\leq$} \end{array}$$

[IN97] P. Iyer, M. Narasimha. Probabilistic lossy channel systems (TAPSOFT'97)
[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Approximation scheme

- Aim: compute probability of $\mathbf{F}$ 🙂
- 🙁 $= \{s \in S \mid s \not\models \exists \mathbf{F}$ 🙂 $\}$

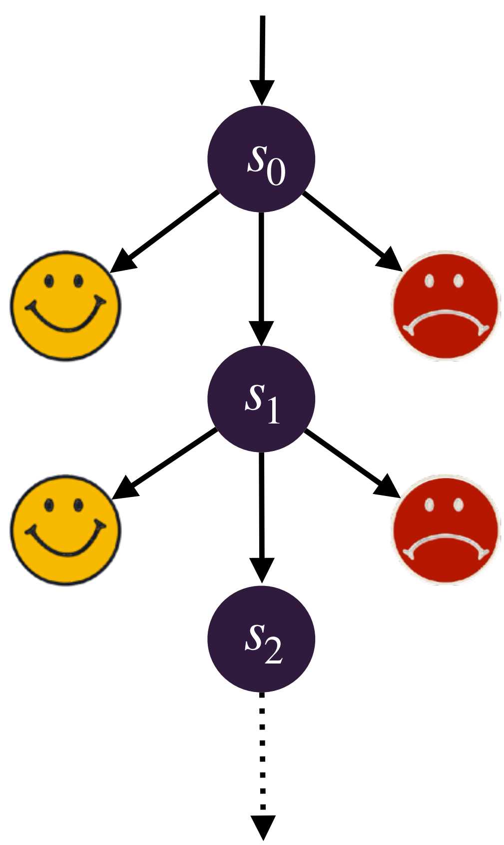

Approximation scheme

Given $\varepsilon > 0$, for every n, compute:

$$\begin{cases} p_n^{\text{yes}} & = \mathbb{P}(\mathbf{F}_{\leq n} \text{🙂}) \\ p_n^{\text{no}} & = \mathbb{P}(\mathbf{F}_{\leq n} \text{🙁}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$$p_1^{\text{yes}} \leq \mathbb{P}(\mathbf{F}\text{🙂}) \leq 1 - p_1^{\text{no}}$$

$$\leq \qquad\qquad \geq$$

$$p_2^{\text{yes}} \leq \mathbb{P}(\mathbf{F}\text{🙂}) \leq 1 - p_2^{\text{no}}$$

$$\leq \qquad \vdots \qquad \geq$$

At the limit: $\mathbb{P}(\mathbf{F}\text{🙂})$ $\qquad 1 - \mathbb{P}(\mathbf{F}\text{🙁})$

[IN97] P. Iyer, M. Narasimha. Probabilistic lossy channel systems (TAPSOFT'97)
[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Approximation scheme

- Aim: compute probability of $\mathbf{F}$ ☺
- ☹ $= \{ s \in S \mid s \not\models \exists \mathbf{F}$ ☺ $\}$

The approximation scheme converges
iff
$\mathscr{C}$ is decisive from s_0 w.r.t. ☺

Approximation scheme

Given $\varepsilon > 0$, for every n, compute:

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}(\mathbf{F}_{\leq n}\,\text{☺}) \\ p_n^{\text{no}} = \mathbb{P}(\mathbf{F}_{\leq n}\,\text{☹}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$$\begin{array}{ccccc} p_1^{\text{yes}} & \leq & \mathbb{P}(\mathbf{F}\,\text{☺}) & \leq & 1 - p_1^{\text{no}} \\ \text{|∧} & & & & \text{∨|} \\ p_2^{\text{yes}} & \leq & \mathbb{P}(\mathbf{F}\,\text{☺}) & \leq & 1 - p_2^{\text{no}} \\ \text{|∧} & & \vdots & & \text{∨|} \end{array}$$

At the limit: $\mathbb{P}(\mathbf{F}\,\text{☺}) \qquad 1 - \mathbb{P}(\mathbf{F}\,\text{☹})$

[IN97] P. Iyer, M. Narasimha. Probabilistic lossy channel systems (TAPSOFT'97)
[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr. Decisive Markov chains (LMCS, 2007)

Statistical model-checking

Sample N paths

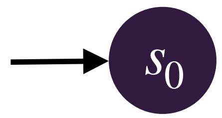

Statistical model-checking

Sample N paths

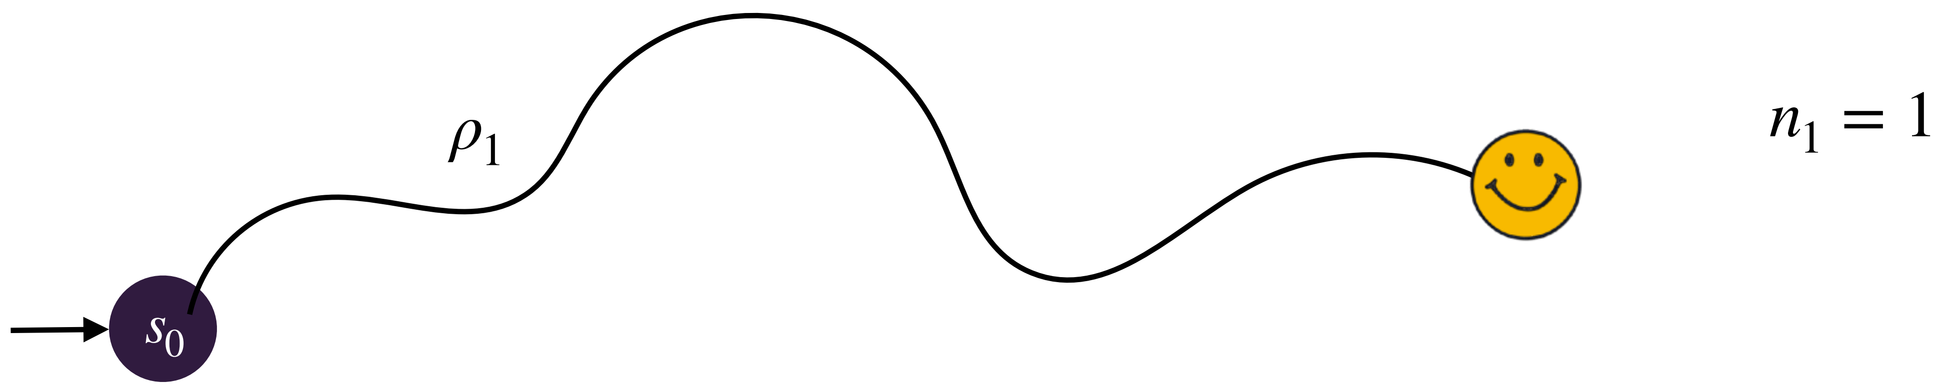

Statistical model-checking

Sample N paths

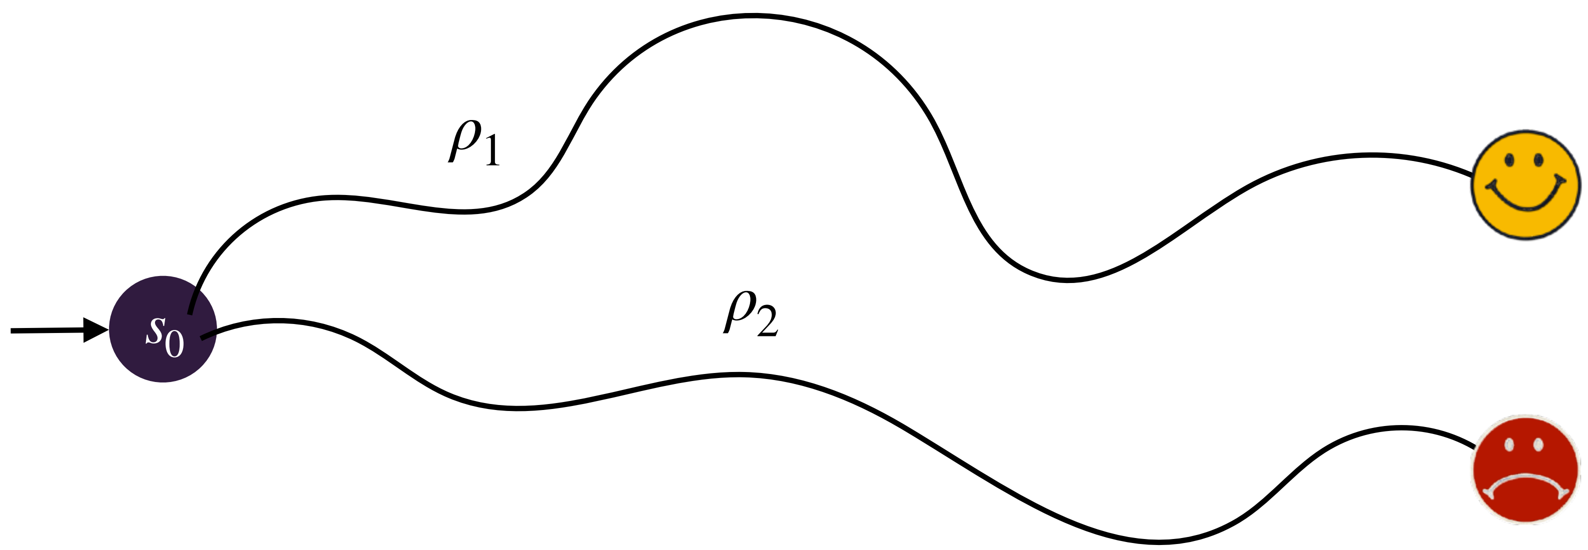

$n_1 = 1$

$n_2 = n_1$

Statistical model-checking

Sample N paths

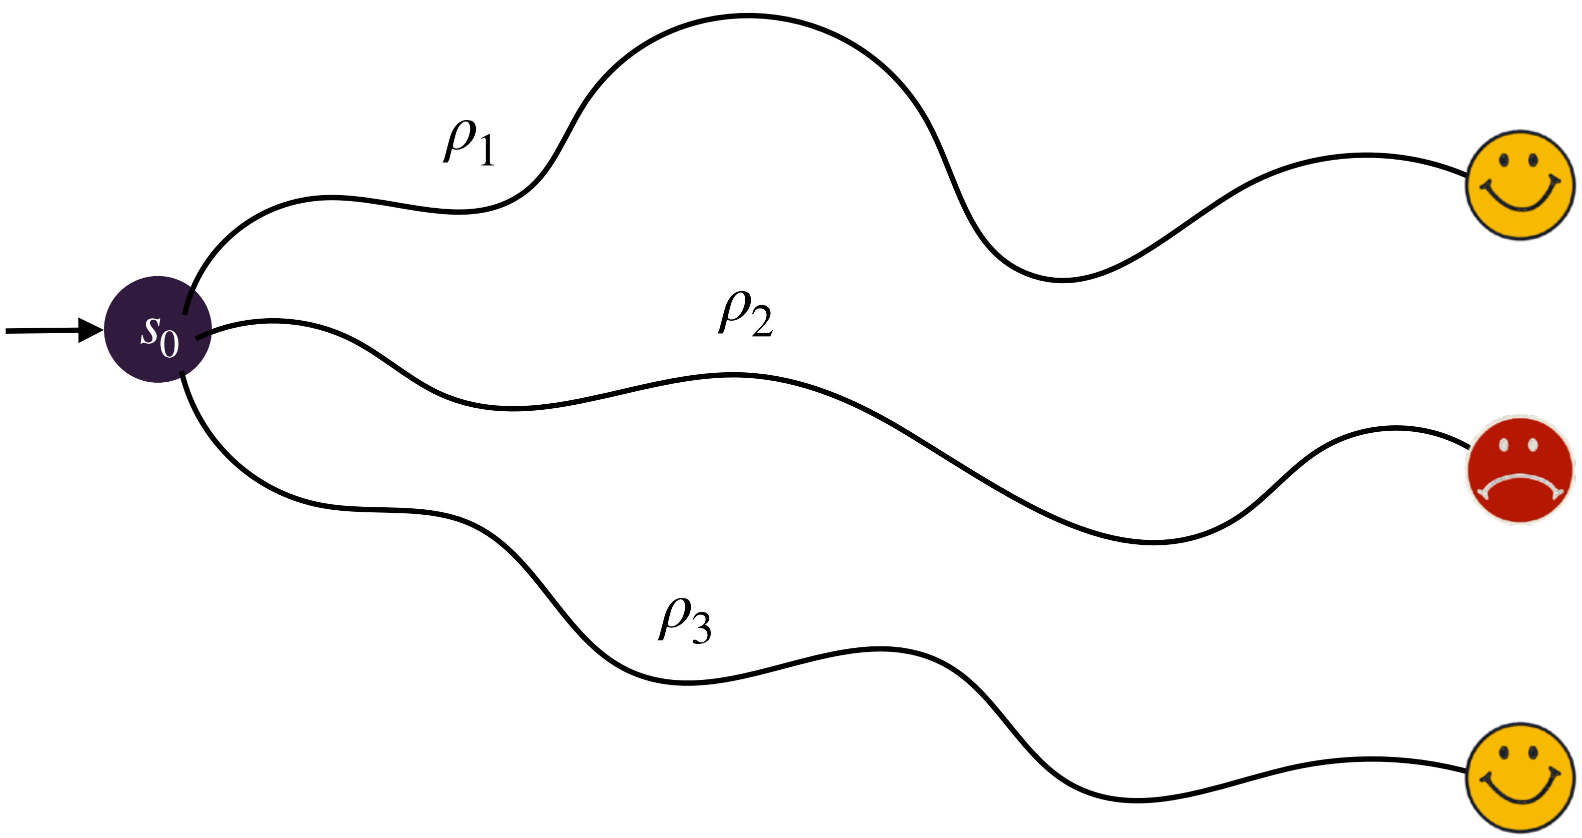

$n_1 = 1$

$n_2 = n_1$

$n_3 = n_2 + 1$

Statistical model-checking

Sample N paths

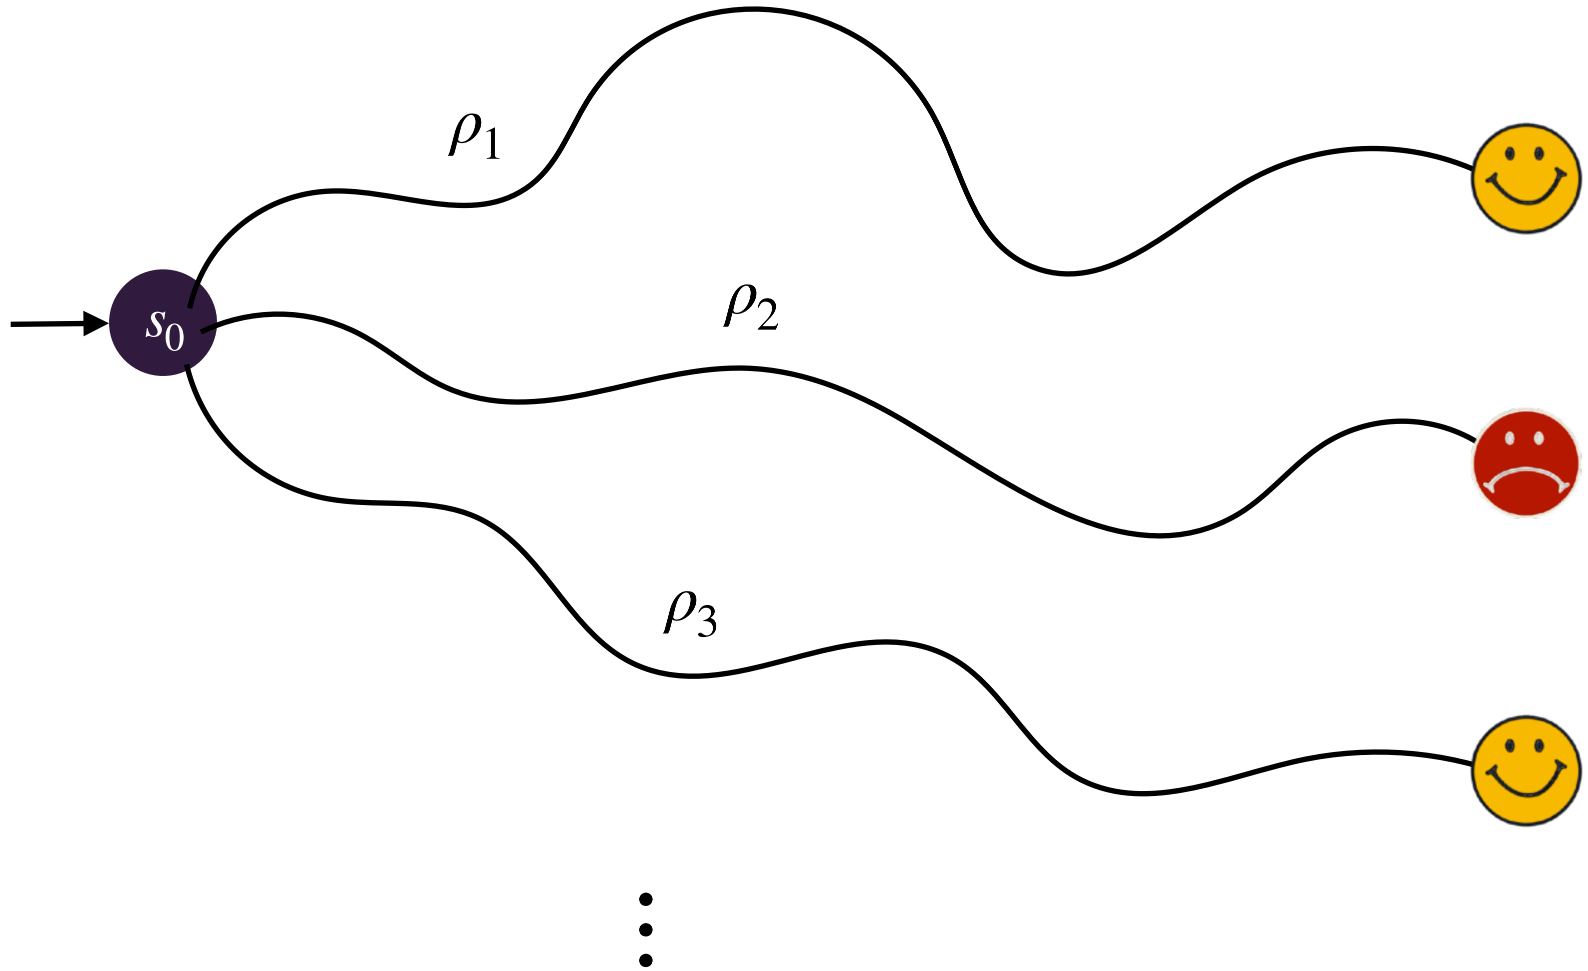

$n_1 = 1$

$n_2 = n_1$

$n_3 = n_2 + 1$

$\vdots$

Statistical model-checking

Sample N paths

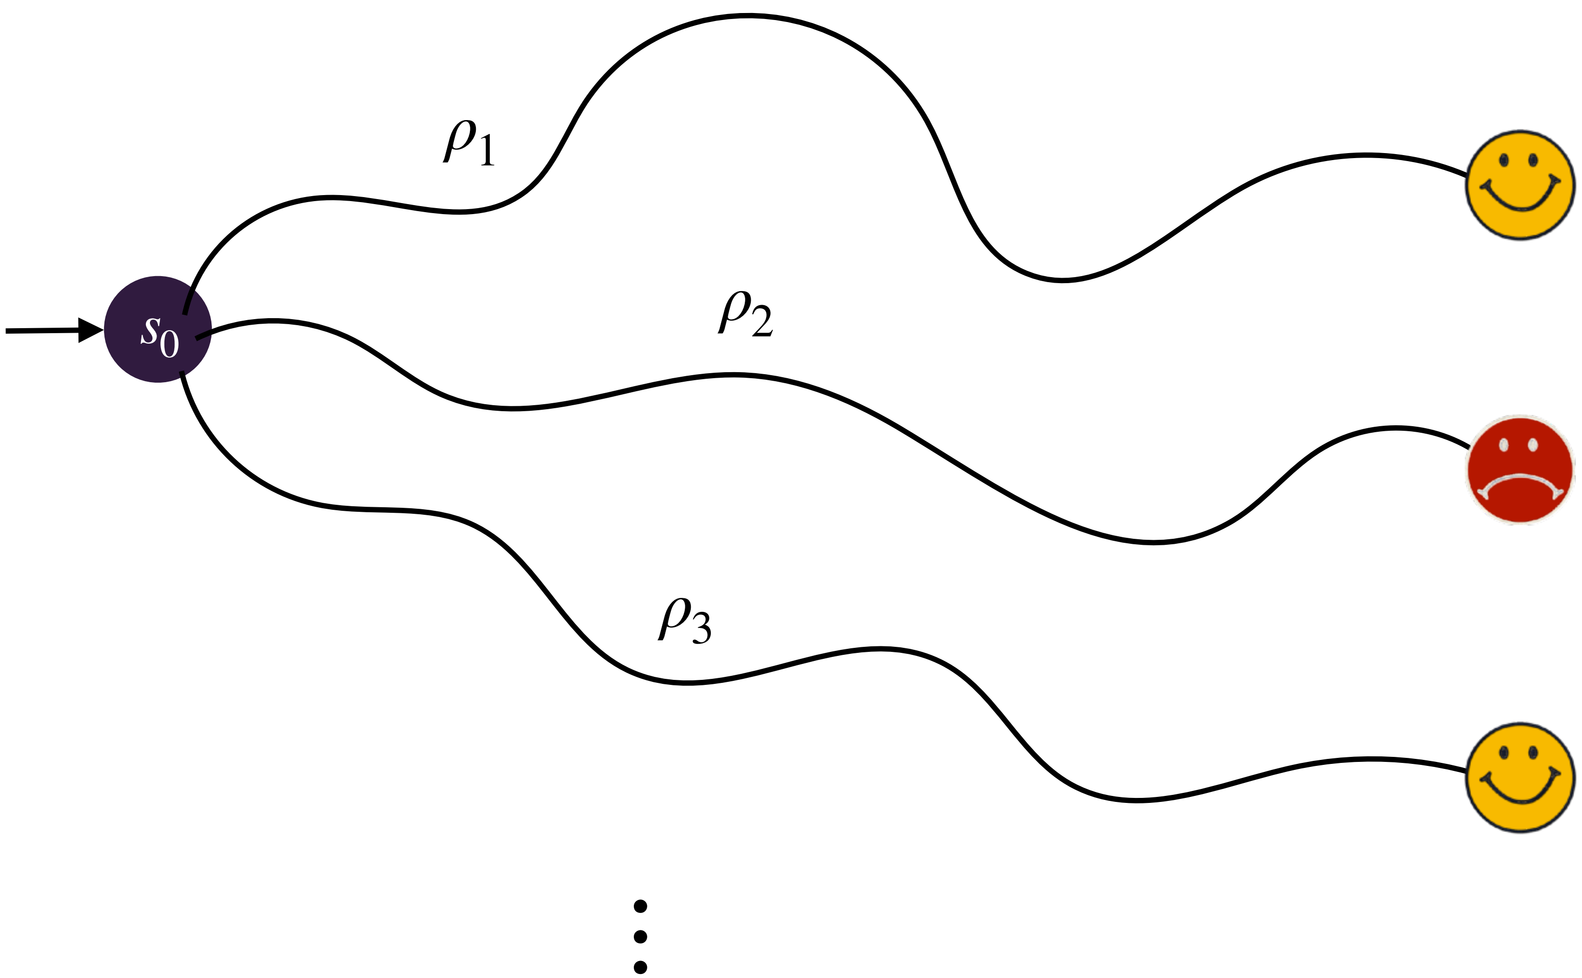

$n_1 = 1$

$n_2 = n_1$

$n_3 = n_2 + 1$

$\vdots$

Return $\dfrac{n_N}{N}$ + some confidence interval

(in the best case)

Statistical model-checking

Sample N paths

Return $\dfrac{n_N}{N}$ + some confidence interval

(in the best case)

Statistical model-checking

Sample N paths

ρ_1

$n_1 = 1$

$n_2 = n_1$

Termination (To our knowledge, never expressed like this)

A sampled path starting at s_0 almost-surely hits 🙂 or ☹️
iff
$\mathscr{C}$ is decisive from s_0 w.r.t. 🙂

⋮

Return $\frac{n_N}{N}$ + some confidence interval

(in the best case)

Where do we stand?

- Decisiveness = crucial concept for Markov chains
  - ensures termination of approximation algorithms
  - ensures feasibility of sampling in algorithms based on SMC

Where do we stand?

- Decisiveness = crucial concept for Markov chains
  - ensures termination of approximation algorithms
  - ensures feasibility of sampling in algorithms based on SMC

Natural question

- Does this concept extend to MDPs and games?

Where do we stand?

- Decisiveness = crucial concept for Markov chains
  - ensures termination of approximation algorithms
  - ensures feasibility of sampling in algorithms based on SMC

Natural question

- Does this concept extend to MDPs and games?

➡ We will focus on MDPs

Markov decision processes

What can we do for them?

Markov decision processes

Markov decision processes (MDPs)

$\mathcal{M} = (S, s_0, \text{Act}, P)$ with S at most denumerable, $s_0 \in S$ and $P : S \times \text{Act} \times S \to [0,1]$ s.t. for each $s \in S$ and $a \in \text{Act}$, $P(s, a, \bullet) \in \text{Dist(S)}$

Markov decision processes

Markov decision processes (MDPs)

$\mathcal{M} = (S, s_0, \text{Act}, P)$ with S at most denumerable, $s_0 \in S$ and
$P : S \times \text{Act} \times S \to [0,1]$ s.t. for each $s \in S$ and $a \in \text{Act}$,
$P(s, a, \bullet) \in \text{Dist(S)}$

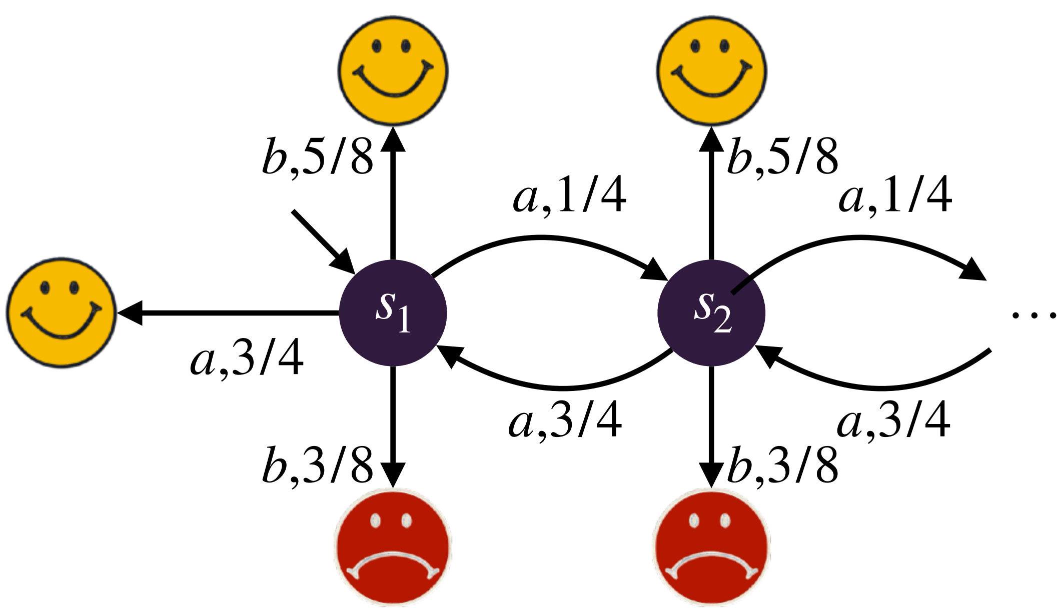

Markov decision processes

Markov decision processes (MDPs)

$\mathcal{M} = (S, s_0, \text{Act}, P)$ with S at most denumerable, $s_0 \in S$ and $P : S \times \text{Act} \times S \to [0,1]$ s.t. for each $s \in S$ and $a \in \text{Act}$, $P(s, a, \bullet) \in \text{Dist(S)}$

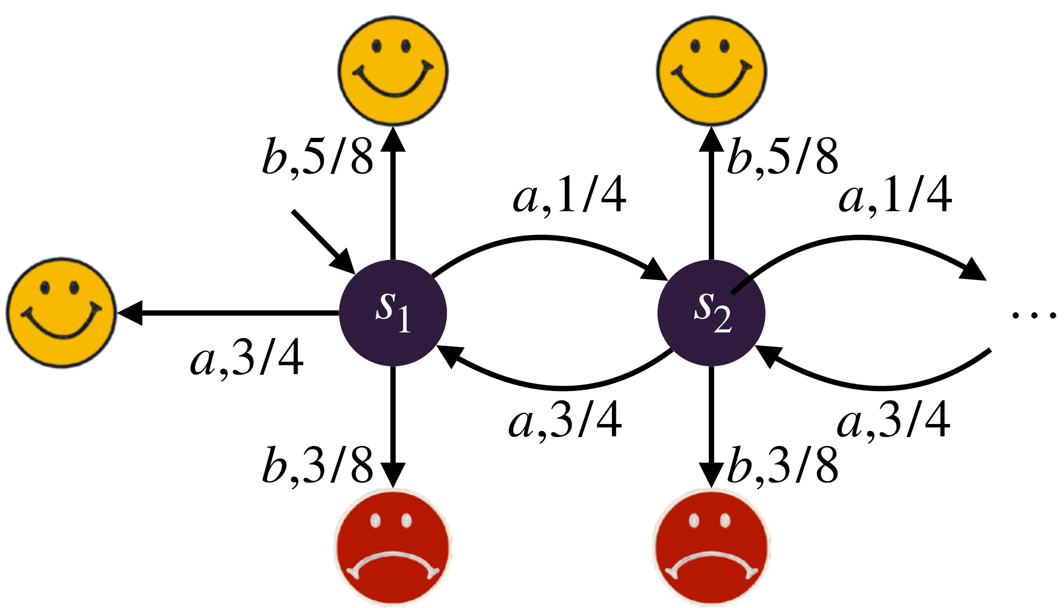

- Requires a scheduler: σ : Hist $\to$ Dist(Act)
- $\mathcal{M}$ under σ is a Markov chain (written $\mathcal{M}^\sigma$)

Markov decision processes

Markov decision processes (MDPs)

$\mathcal{M} = (S, s_0, \text{Act}, P)$ with S at most denumerable, $s_0 \in S$ and
$P : S \times \text{Act} \times S \to [0,1]$ s.t. for each $s \in S$ and $a \in \text{Act}$,
$P(s, a, \bullet) \in \text{Dist(S)}$

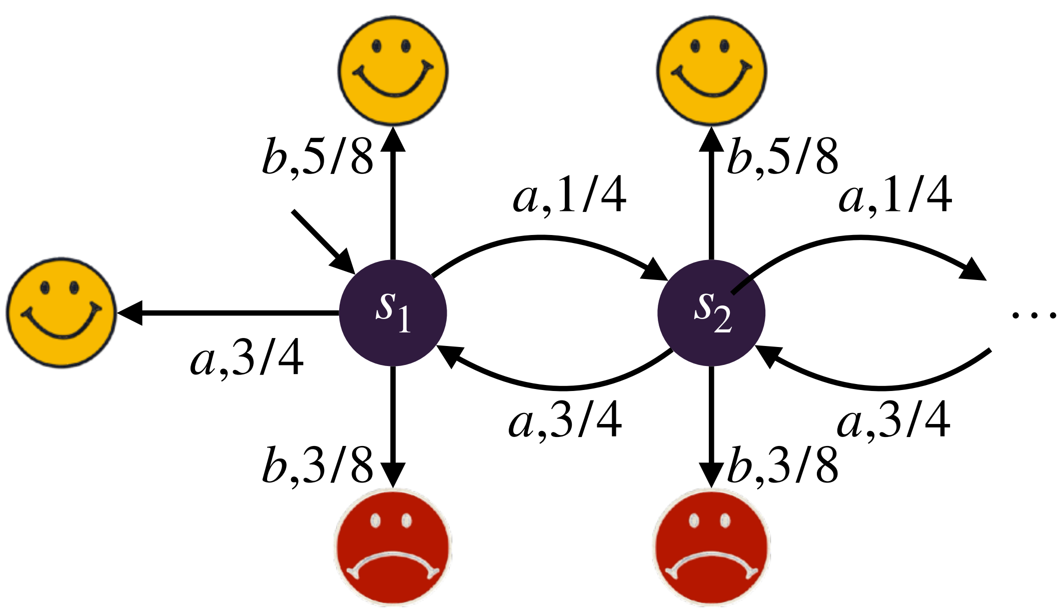

Compute $\mathbb{P}^{\text{opt}}_{\mathcal{M},s_0}\left(\mathbf{F}\ ☺\right)$

$=$

opt_σ scheduler $\mathbb{P}_{\mathcal{M}^\sigma,s_0}\left(\mathbf{F}\ ☺\right)$

- Requires a scheduler: $\sigma : \text{Hist} \to \text{Dist(Act)}$
- $\mathcal{M}$ under σ is a Markov chain (written $\mathcal{M}^\sigma$)

Markov decision processes

Markov decision processes (MDPs)

$\mathscr{M} = (S, s_0, \text{Act}, P)$ with S at most denumerable, $s_0 \in S$ and
$P : S \times \text{Act} \times S \to [0,1]$ s.t. for each $s \in S$ and $a \in \text{Act}$,
$P(s, a, \bullet) \in \text{Dist}(S)$

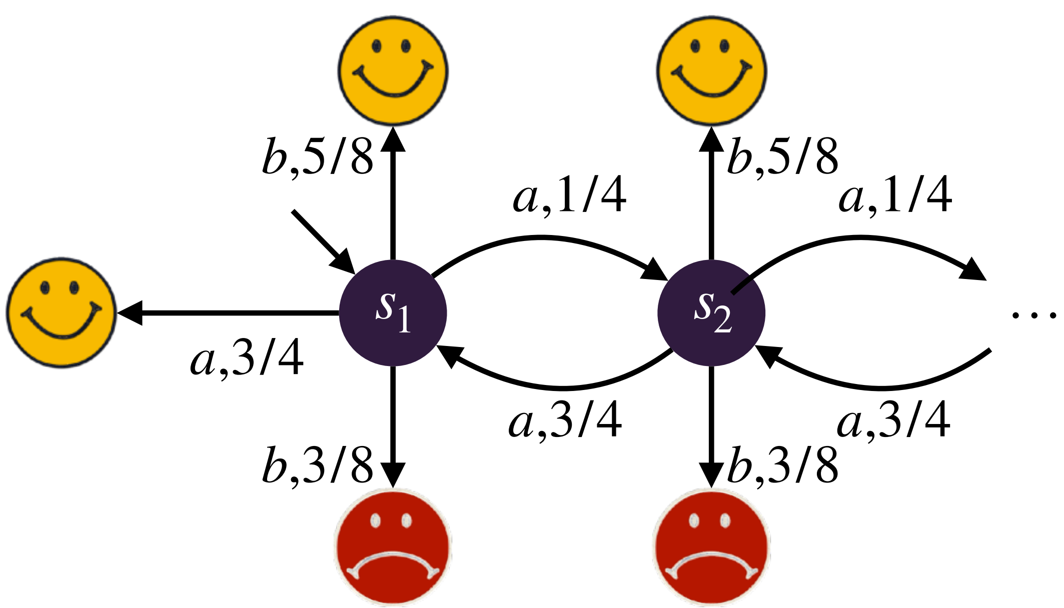

Compute $\mathbb{P}^{\text{opt}}_{\mathscr{M},s_0}(\mathbf{F}\ ☺)$

$=$

$\text{opt}_{\sigma \text{ scheduler}}\ \mathbb{P}_{\mathscr{M}^\sigma,s_0}(\mathbf{F}\ ☺)$

Focus on reachability properties

- Requires a scheduler: $\sigma : \text{Hist} \to \text{Dist}(\text{Act})$
- $\mathscr{M}$ under σ is a Markov chain (written $\mathscr{M}^\sigma$)

Where do we go?

Where do we go?

Our goal

We want to design (approximation) algorithms to compute inf/sup-values of MDPs

Where do we go?

Our goal

We want to design (approximation) algorithms to compute inf/sup-values of MDPs

Some remarks

- The values can be characterized via fixpoints, but no general stopping criterion. The values are not known to be computable in general.
- Remarks: inf and sup are very different (in terms of (existence and nature of) optimal schedulers for example).
  - Optimal positional schedulers exist for inf
  - Optimal schedulers may not exist for sup

Where do we go?

Our goal

We want to design (approximation) algorithms to compute inf/sup-values of MDPs

Some remarks

- The values can be characterized via fixpoints, but no general stopping criterion. The values are not known to be computable in general.
- Remarks: inf and sup are very different (in terms of (existence and nature of) optimal schedulers for example).
  - Optimal positional schedulers exist for inf
  - Optimal schedulers may not exist for sup

➡ We call $\text{Approx}^{\text{opt}}_{\text{unfold}}$ the approximation scheme which extends in a natural way the approximation scheme for Markov chains

Several « natural » decisiveness concepts for MDPs

Several « natural » decisiveness concepts for MDPs

$☹️^{\text{opt}} = \{s \in S \mid \mathbb{P}_s^{\text{opt}}(\mathbf{F}\,🙂) = 0\}$

opt-Decisiveness

An MDP $\mathscr{M}$ is opt-decisive from s w.r.t. 🙂 if $\mathbb{P}^{\text{opt}}_{\mathscr{M},s}(\mathbf{F}\,🙂 \vee \mathbf{F}\,☹️^{\text{opt}}) = 1$

Several « natural » decisiveness concepts for MDPs

$\frownie^{\text{opt}} = \{s \in S \mid \mathbb{P}_s^{\text{opt}}(\mathbf{F}\smiley) = 0\}$

opt-Decisiveness

An MDP $\mathscr{M}$ is opt-decisive from s w.r.t. $\smiley$ if $\mathbb{P}^{\text{opt}}_{\mathscr{M},s}(\mathbf{F}\smiley \vee \mathbf{F}\frownie^{\text{opt}}) = 1$

$\frownie^{\sigma} = \{s \in S \mid \mathbb{P}_{\sigma,s}(\mathbf{F}\smiley) = 0\}$

σ-Decisiveness

An MDP $\mathscr{M}$ is σ-decisive from s w.r.t. $\smiley$ if $\mathbb{P}_{\mathscr{M}^\sigma,s}(\mathbf{F}\smiley \vee \mathbf{F}\frownie^{\sigma}) = 1$

Several « natural » decisiveness concepts for MDPs

$\frownie^{\mathrm{opt}} = \{s \in S \mid \mathbb{P}^{\mathrm{opt}}_{s}(\mathbf{F}\,\smiley) = 0\}$

opt-Decisiveness

An MDP $\mathscr{M}$ is opt-decisive from s w.r.t. $\smiley$ if $\mathbb{P}^{\mathrm{opt}}_{\mathscr{M},s}(\mathbf{F}\,\smiley \vee \mathbf{F}\,\frownie^{\mathrm{opt}}) = 1$

$\frownie^{\sigma} = \{s \in S \mid \mathbb{P}_{\sigma,s}(\mathbf{F}\,\smiley) = 0\}$

σ-Decisiveness

An MDP $\mathscr{M}$ is σ-decisive from s w.r.t. $\smiley$ if $\mathbb{P}_{\mathscr{M}^{\sigma},s}(\mathbf{F}\,\smiley \vee \mathbf{F}\,\frownie^{\sigma}) = 1$

univ-Decisiveness

pure positional

An MDP $\mathscr{M}$ is univ-decisive from s w.r.t. $\smiley$ if it is σ-decisive for every scheduler pp σ

A small example

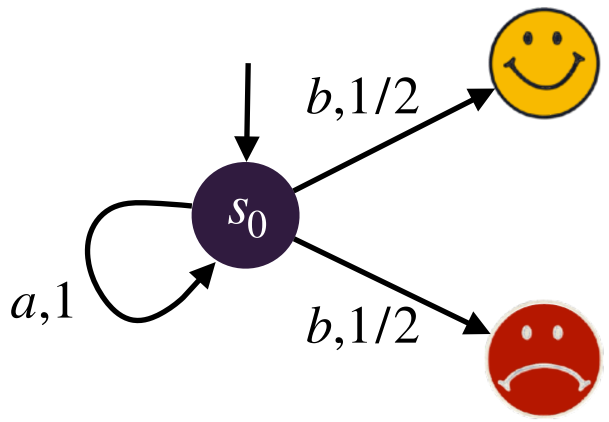

A small example

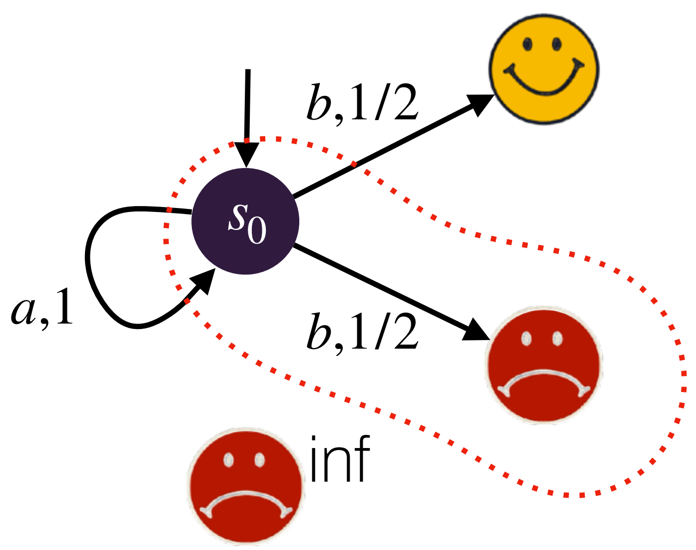

- It is inf-decisive

A small example

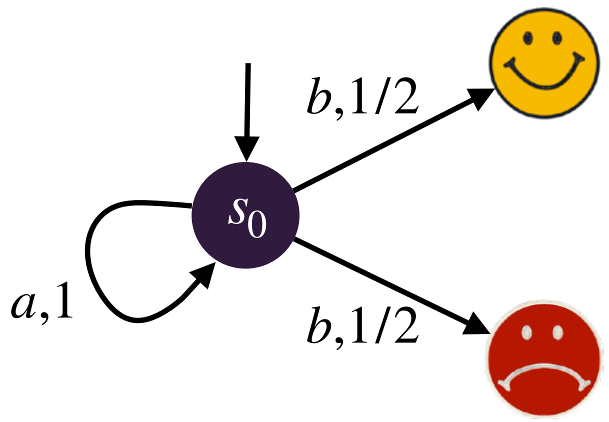

- It is inf-decisive
- It is univ-decisive

A small example

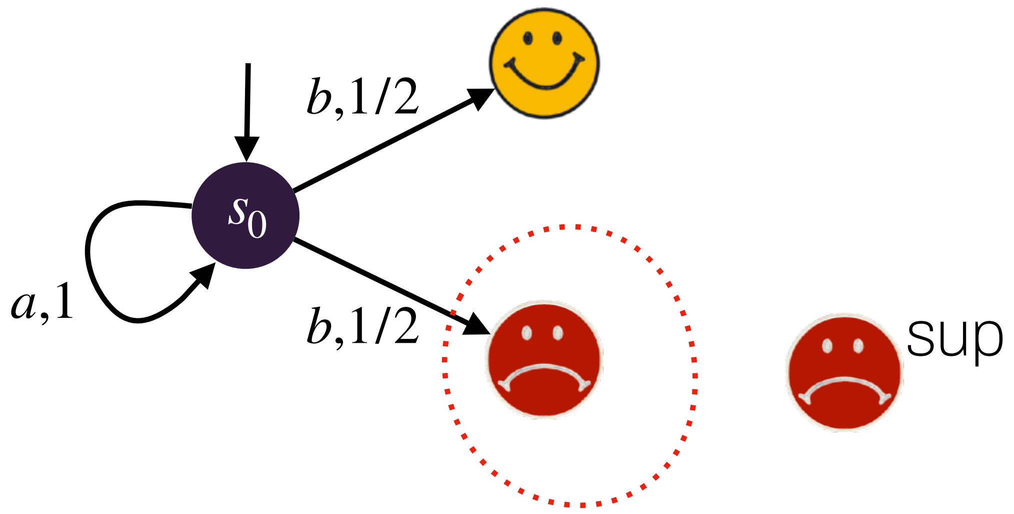

- It is inf-decisive
- It is univ-decisive
- It is **not** sup-decisive

A small example

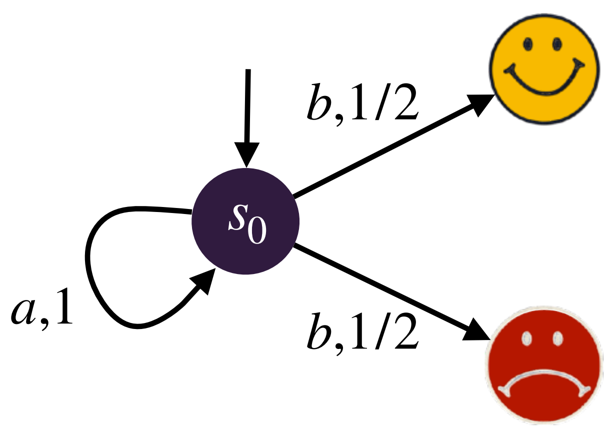

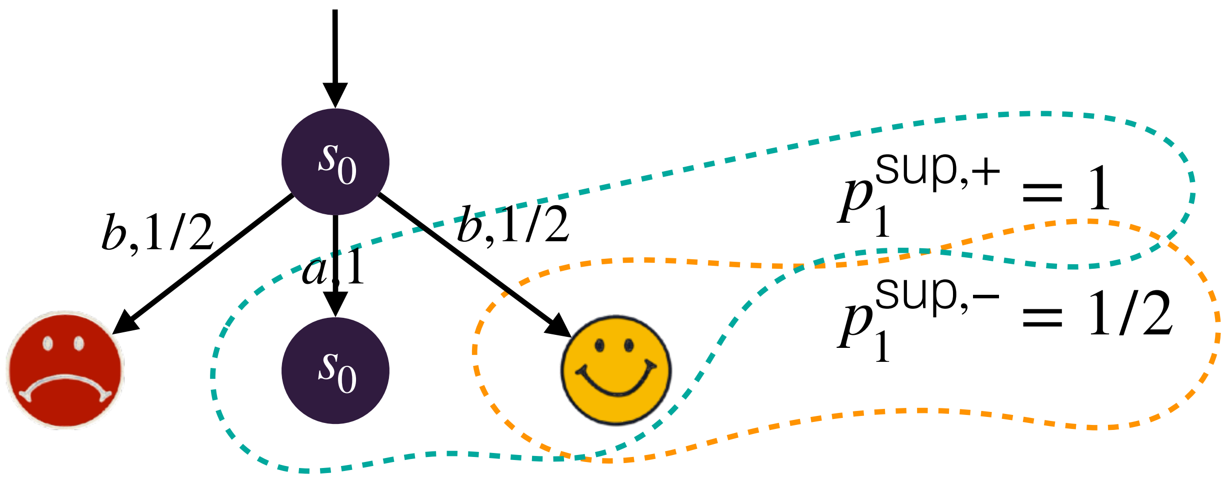

- It is inf-decisive
- It is univ-decisive
- It is **not** sup-decisive

A small example

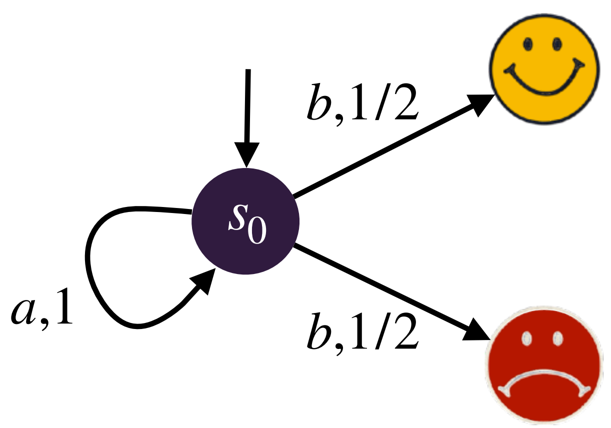

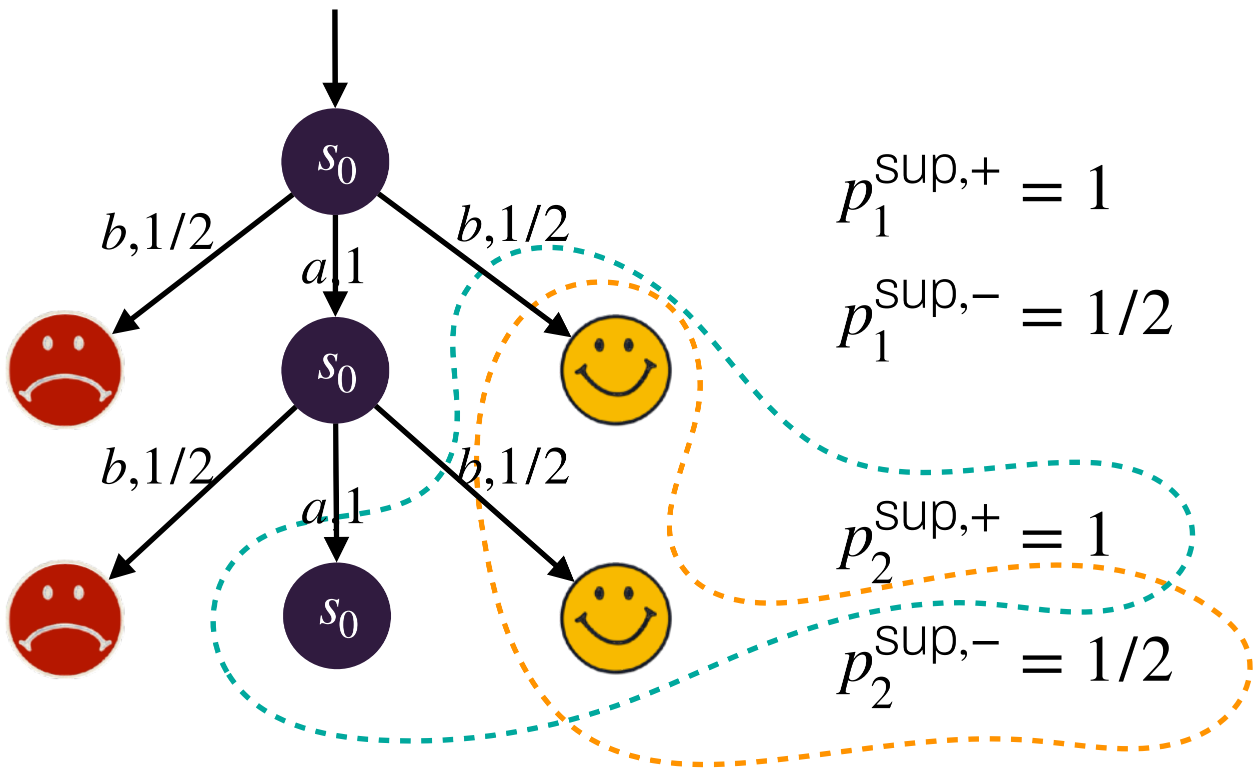

- It is inf-decisive
- It is univ-decisive
- It is **not** sup-decisive

A small example

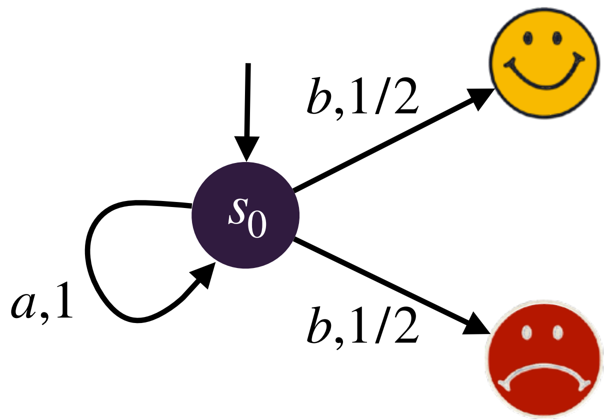

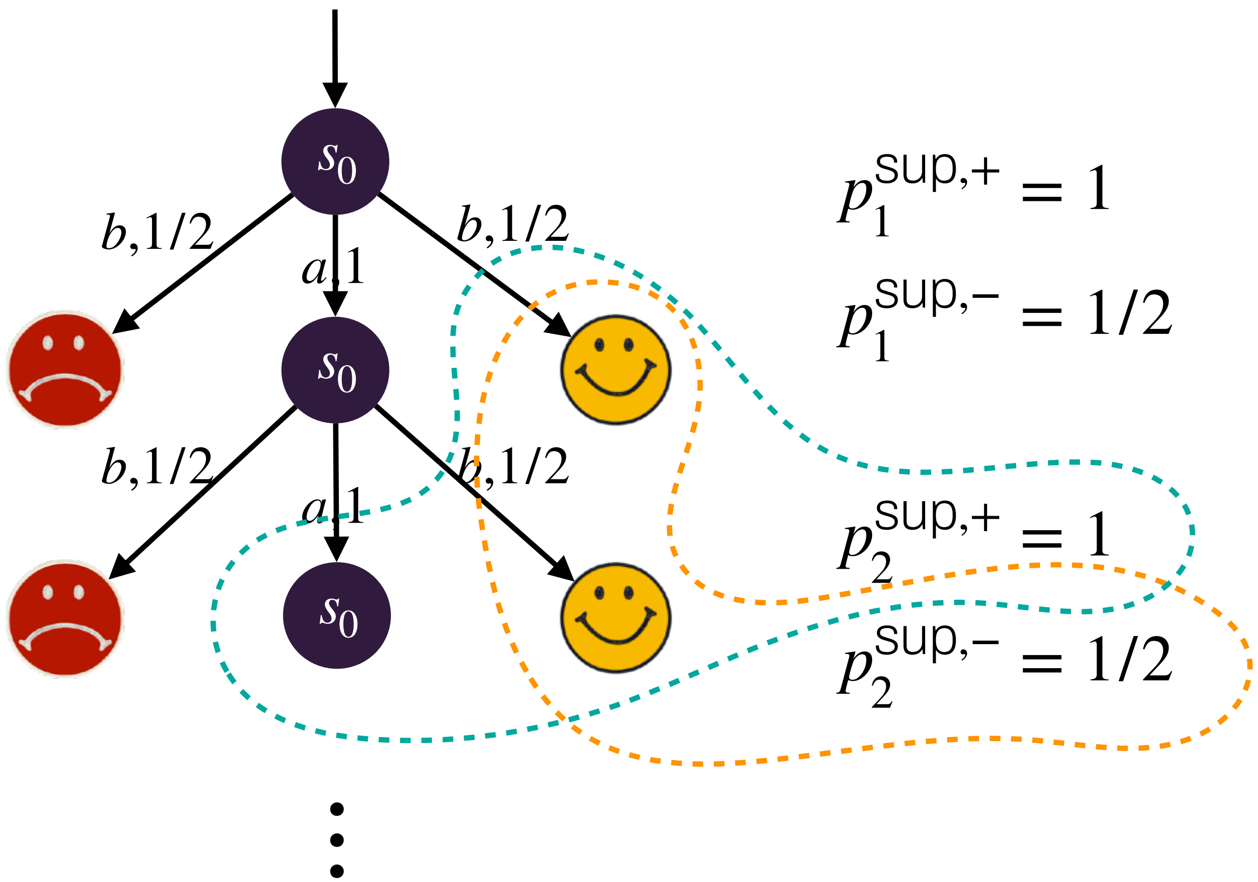

$p_1^{\text{sup},+} = 1$

$p_1^{\text{sup},-} = 1/2$

$p_2^{\text{sup},+} = 1$

$p_2^{\text{sup},-} = 1/2$

- It is inf-decisive
- It is univ-decisive
- It is **not** sup-decisive

The two sequences $(p_n^{\text{sup},-})_n$ and $(p_n^{\text{sup},+})_n$ are not adjacent

A small example

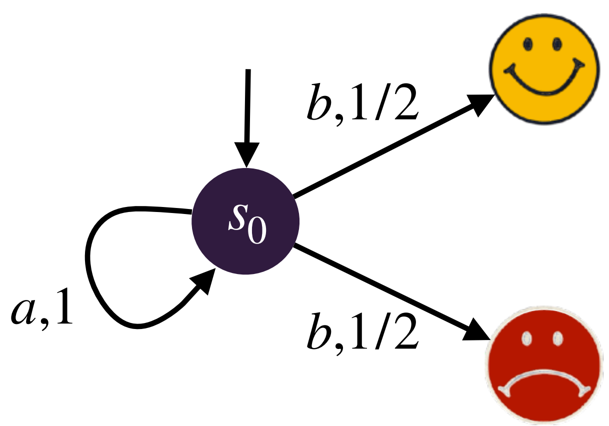

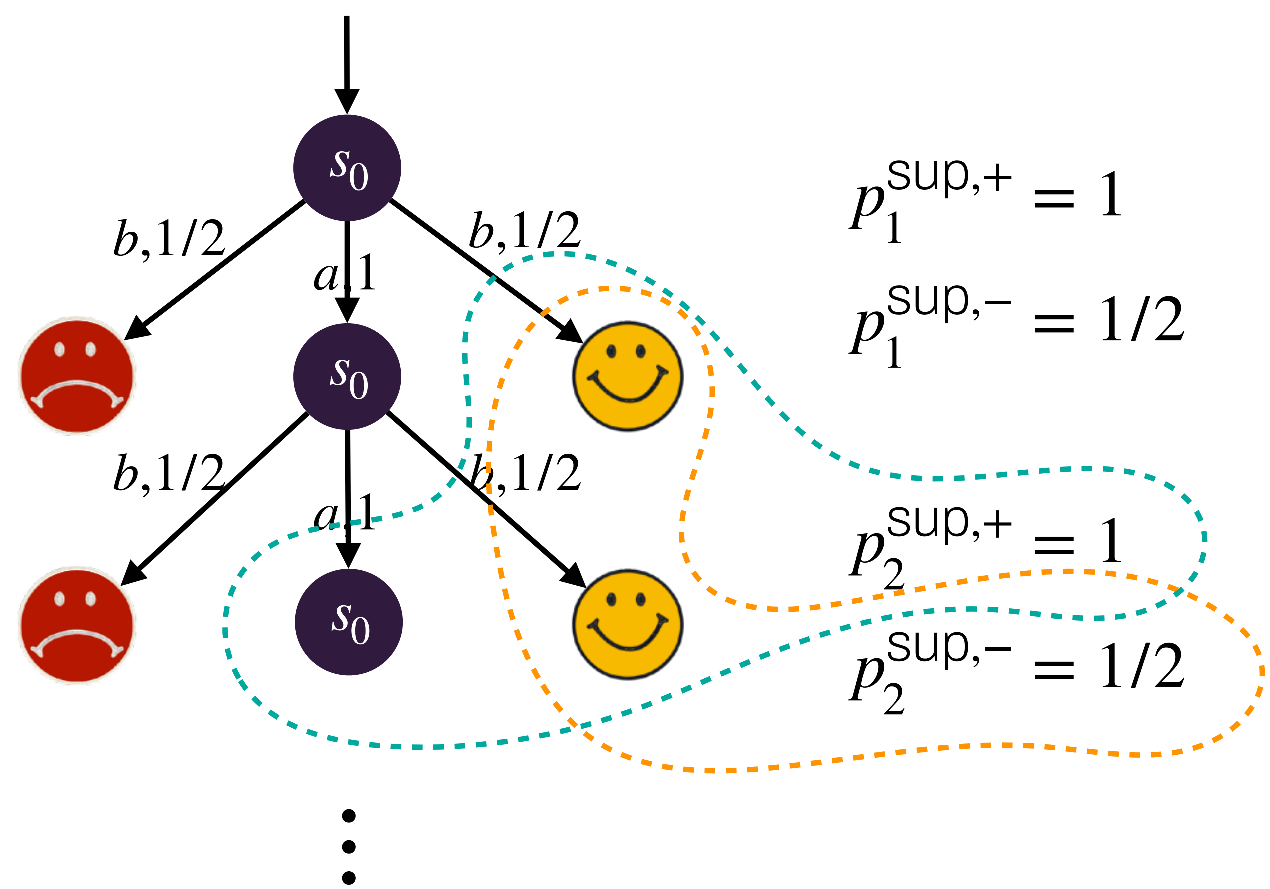

- It is inf-decisive
- It is univ-decisive
- It is **not** sup-decisive

The two sequences $(p_n^{\text{sup},-})_n$ and $(p_n^{\text{sup},+})_n$ are not adjacent

$\text{Approx}^{\text{sup}}_{\text{unfold}}$ is not converging, even for finite MDPs!

The sliced MDPs

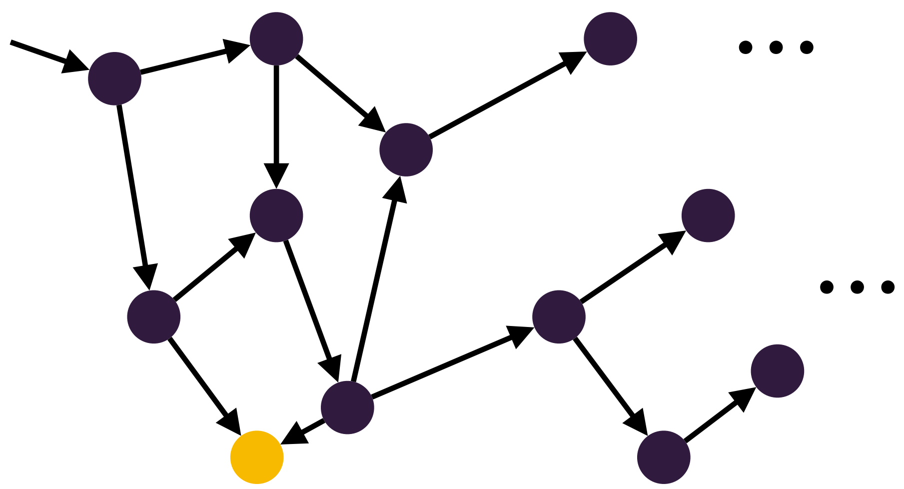

$\mathscr{M}$

Some technical details are ignored

The sliced MDPs

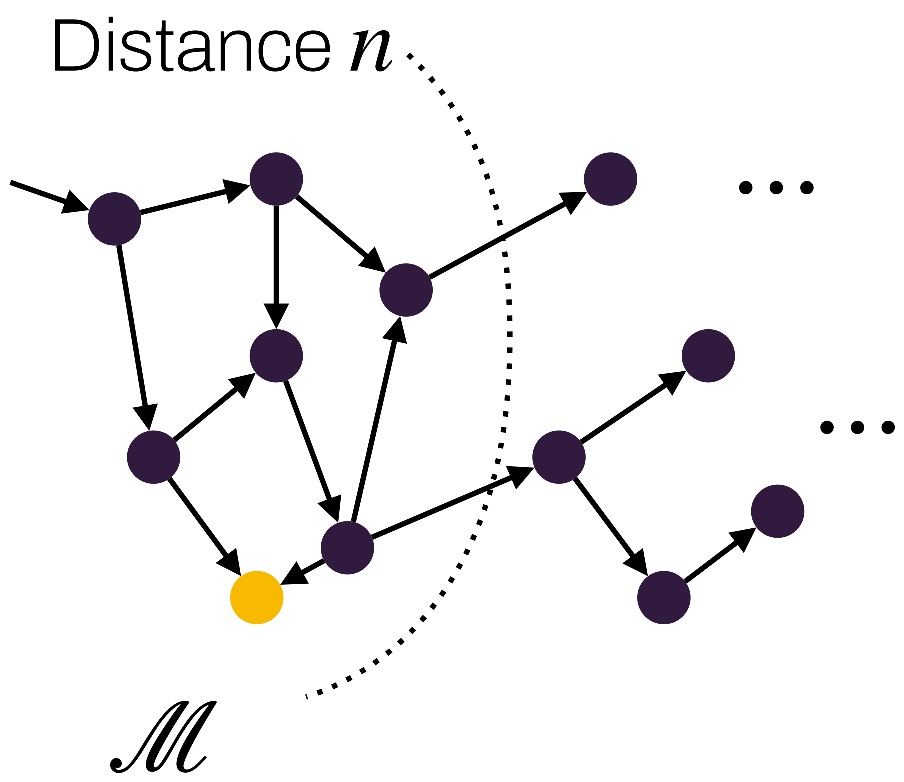

Some technical details are ignored

The sliced MDPs

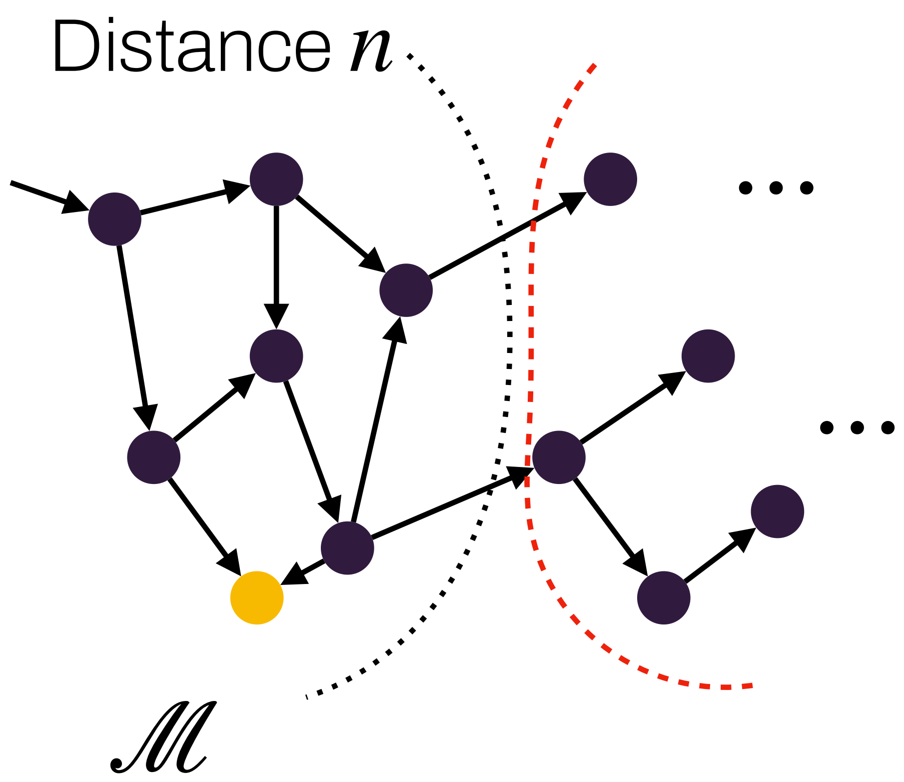

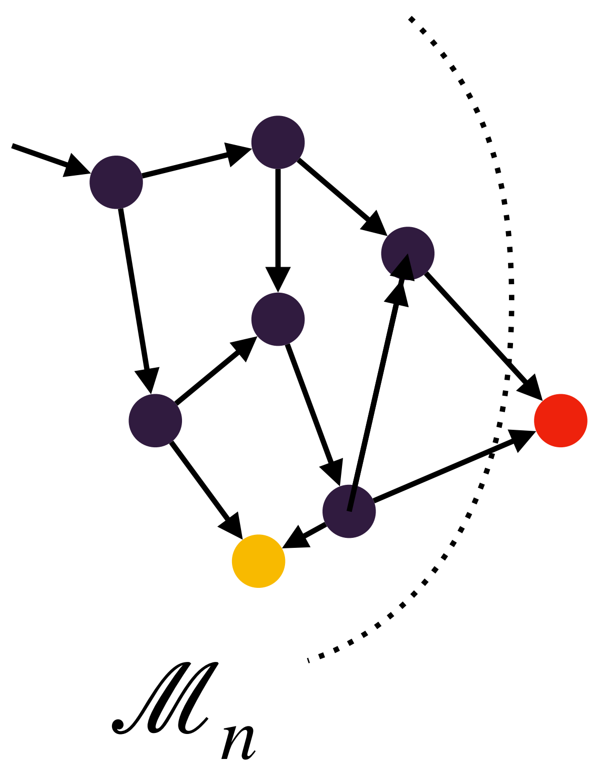

Some technical details are ignored

The sliced MDPs

Distance n

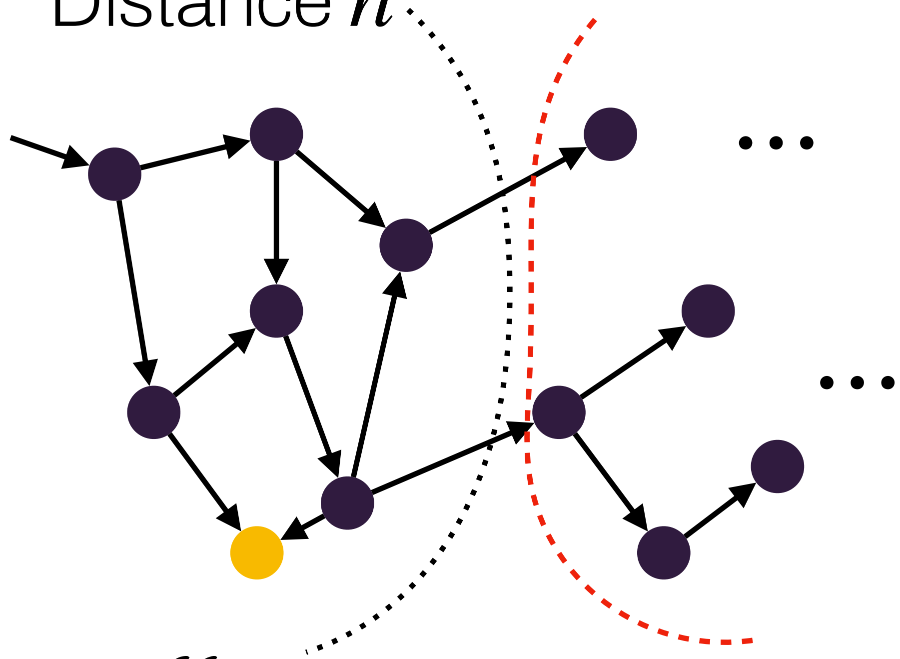

$\mathscr{M}$

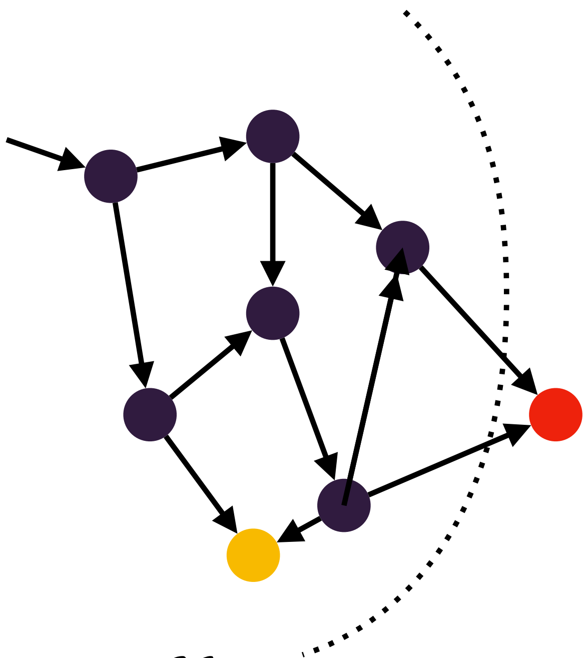

$\mathscr{M}_n$

- Scheme $\text{Approx}^{\text{opt}}_{\text{sliced}}$

Some technical details are ignored

The sliced MDPs

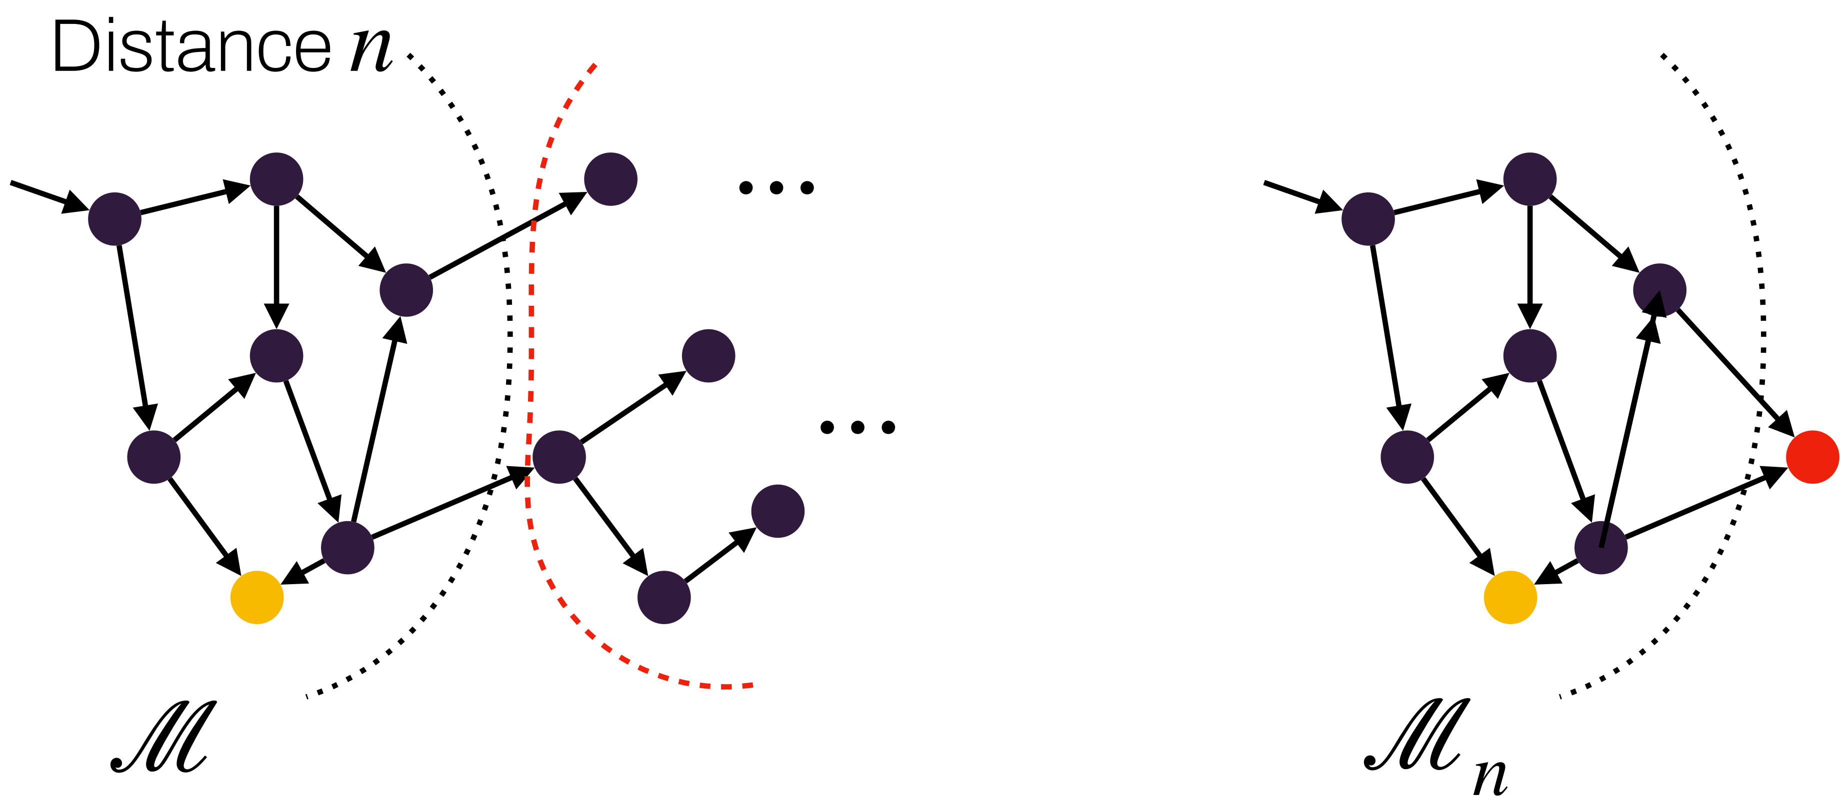

- Scheme $\text{Approx}^{\text{opt}}_{\text{sliced}}$

Converges over all finite MDPs…

Some technical details are ignored

Results — Summary

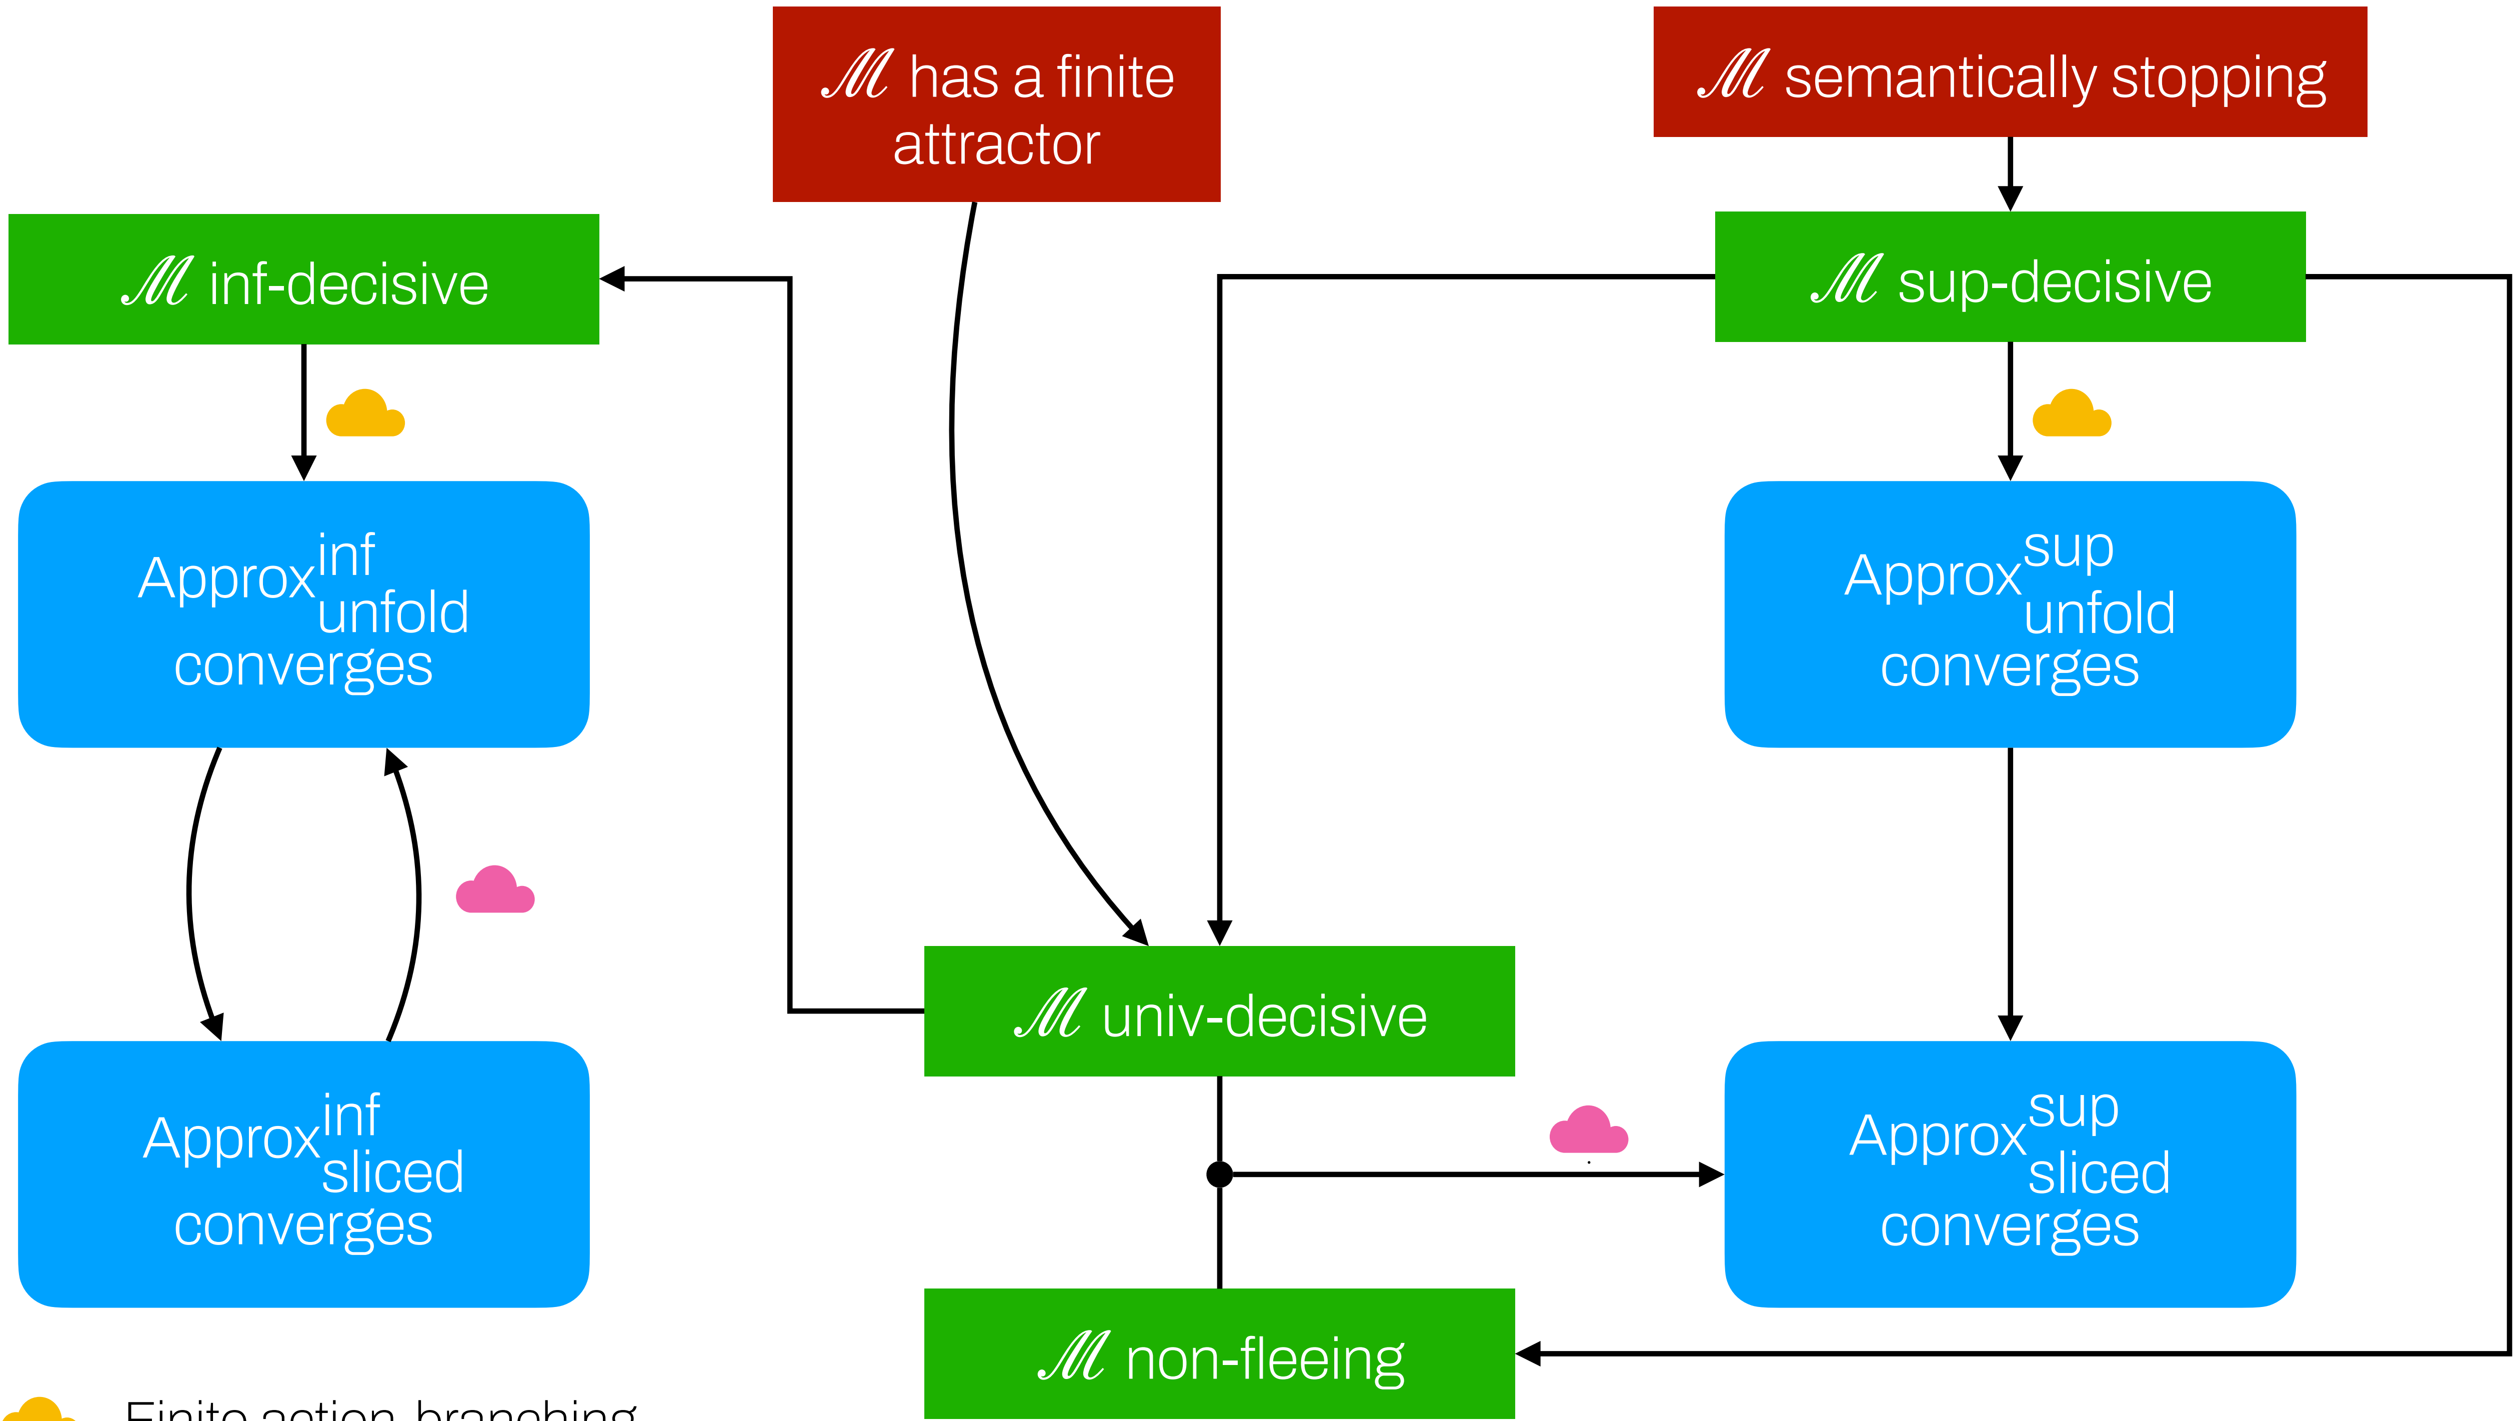

Finite action-branching

Finite branching

Results — Summary

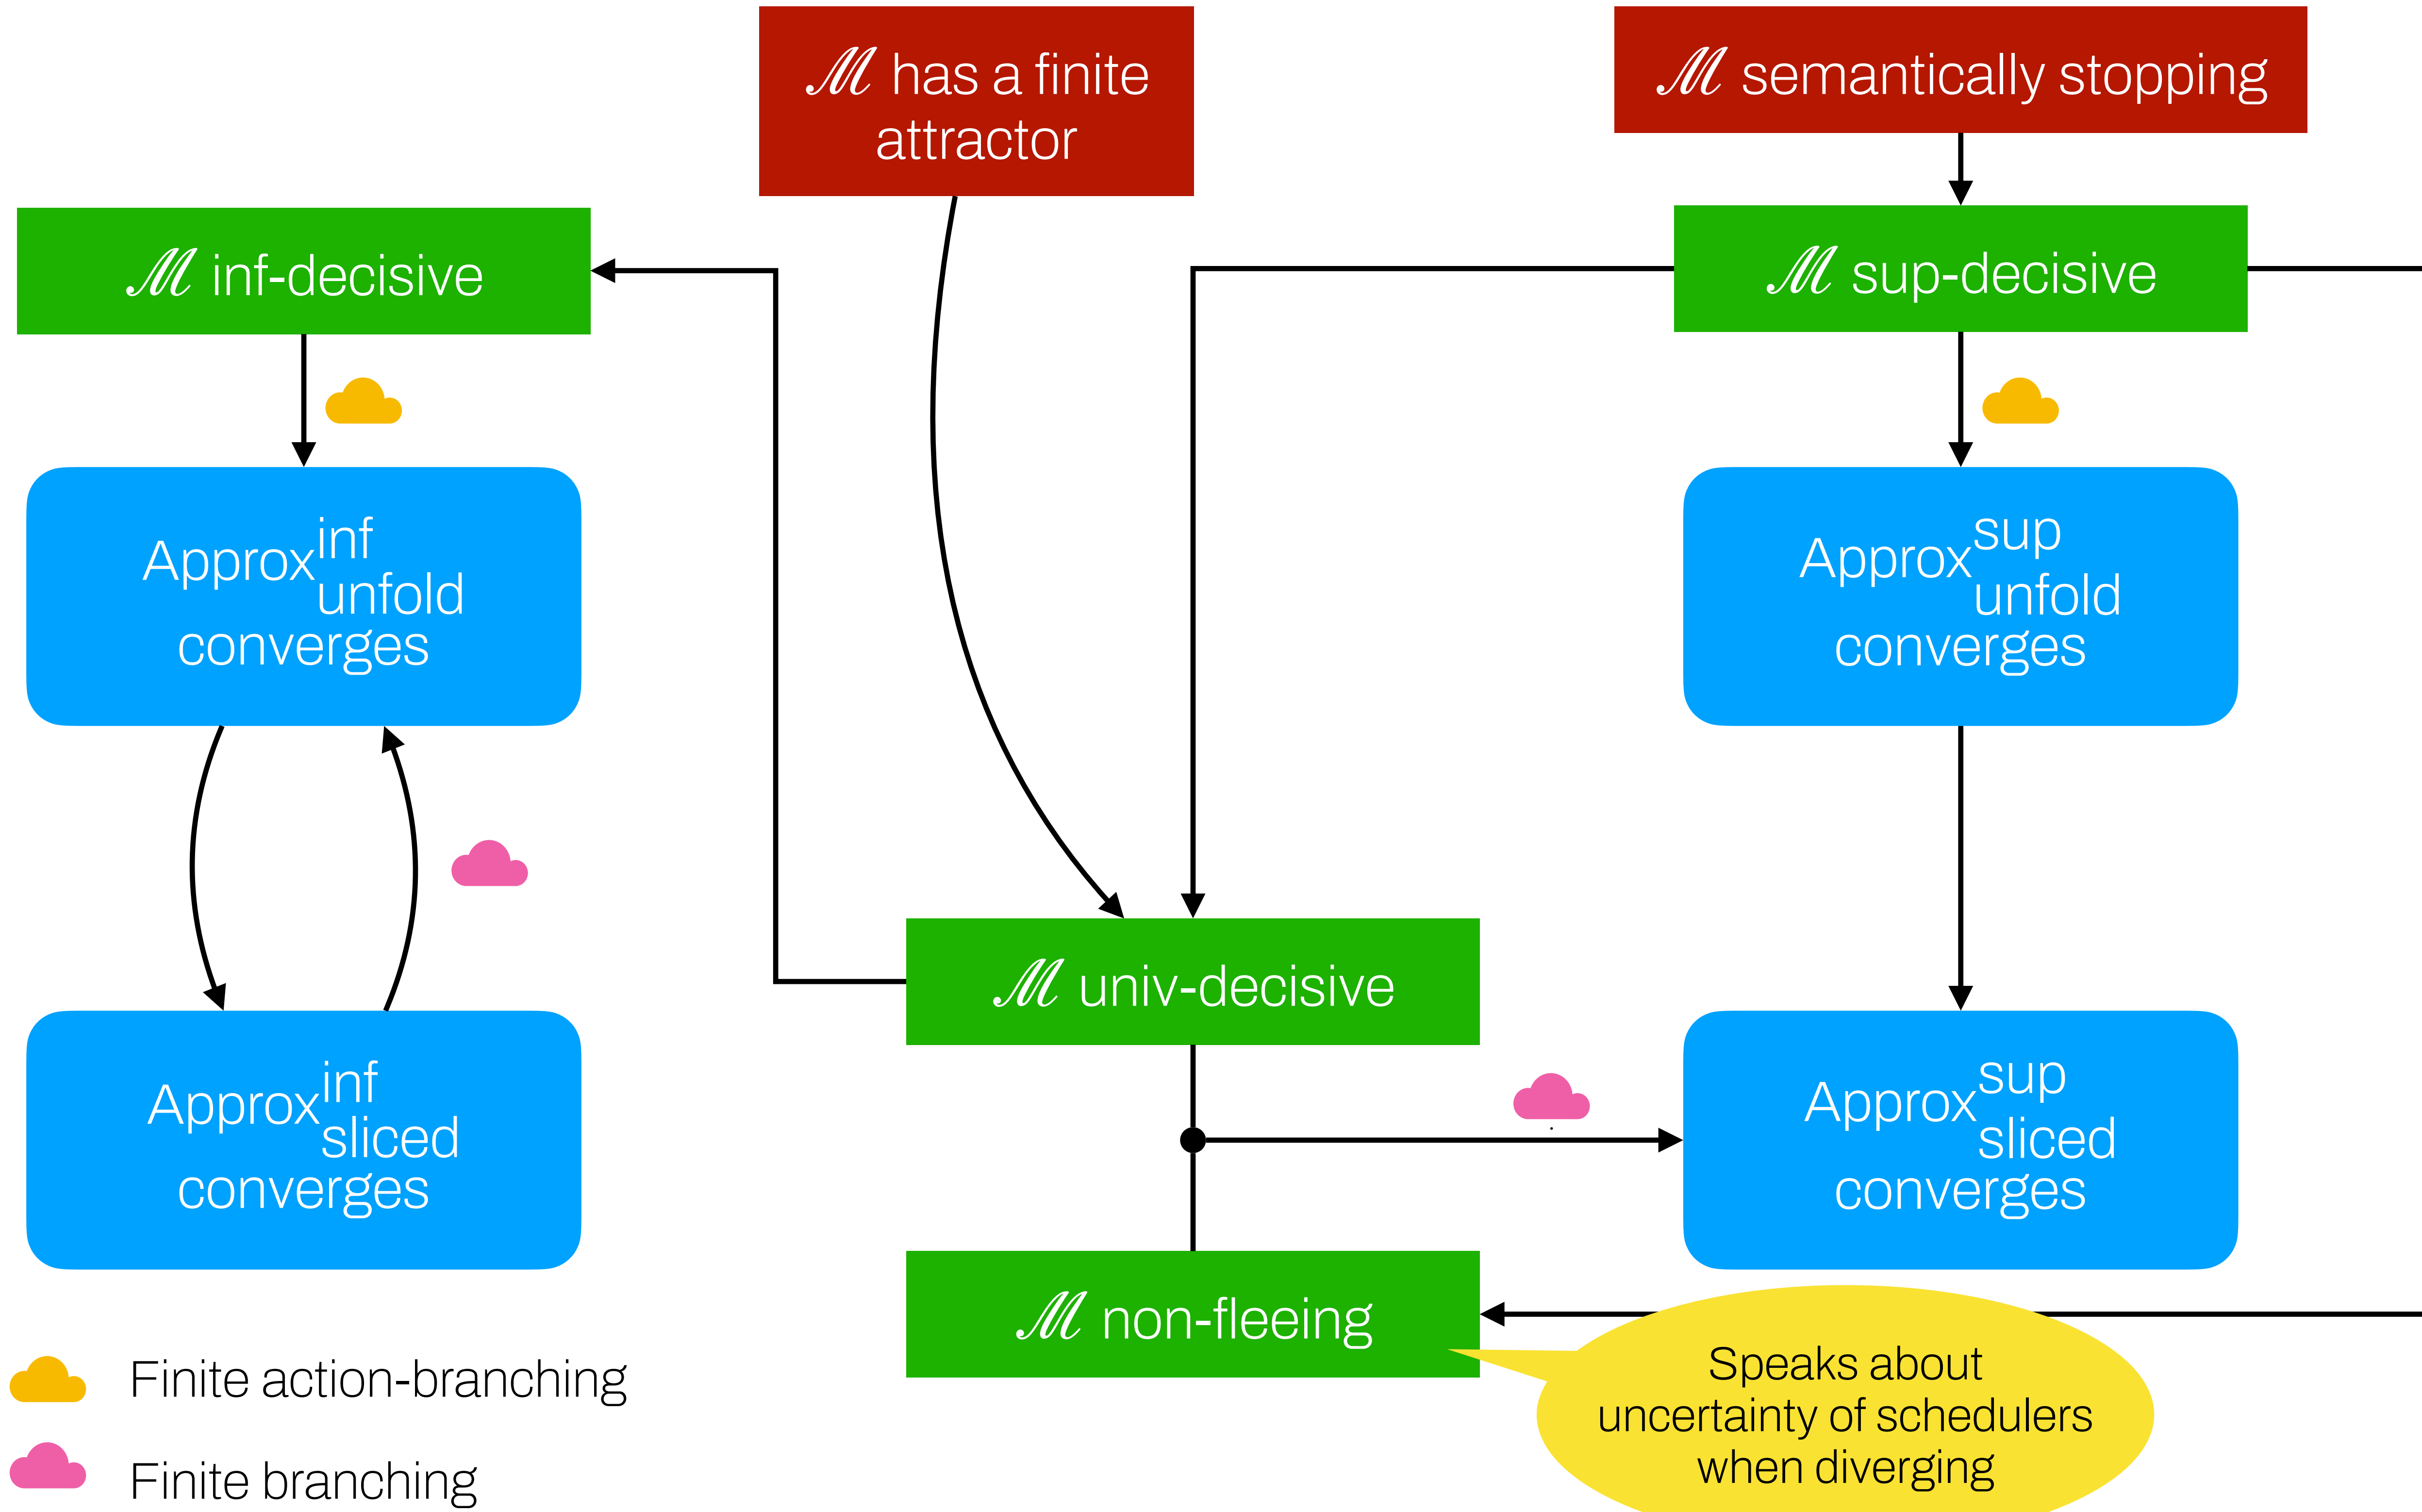

Results — Summary

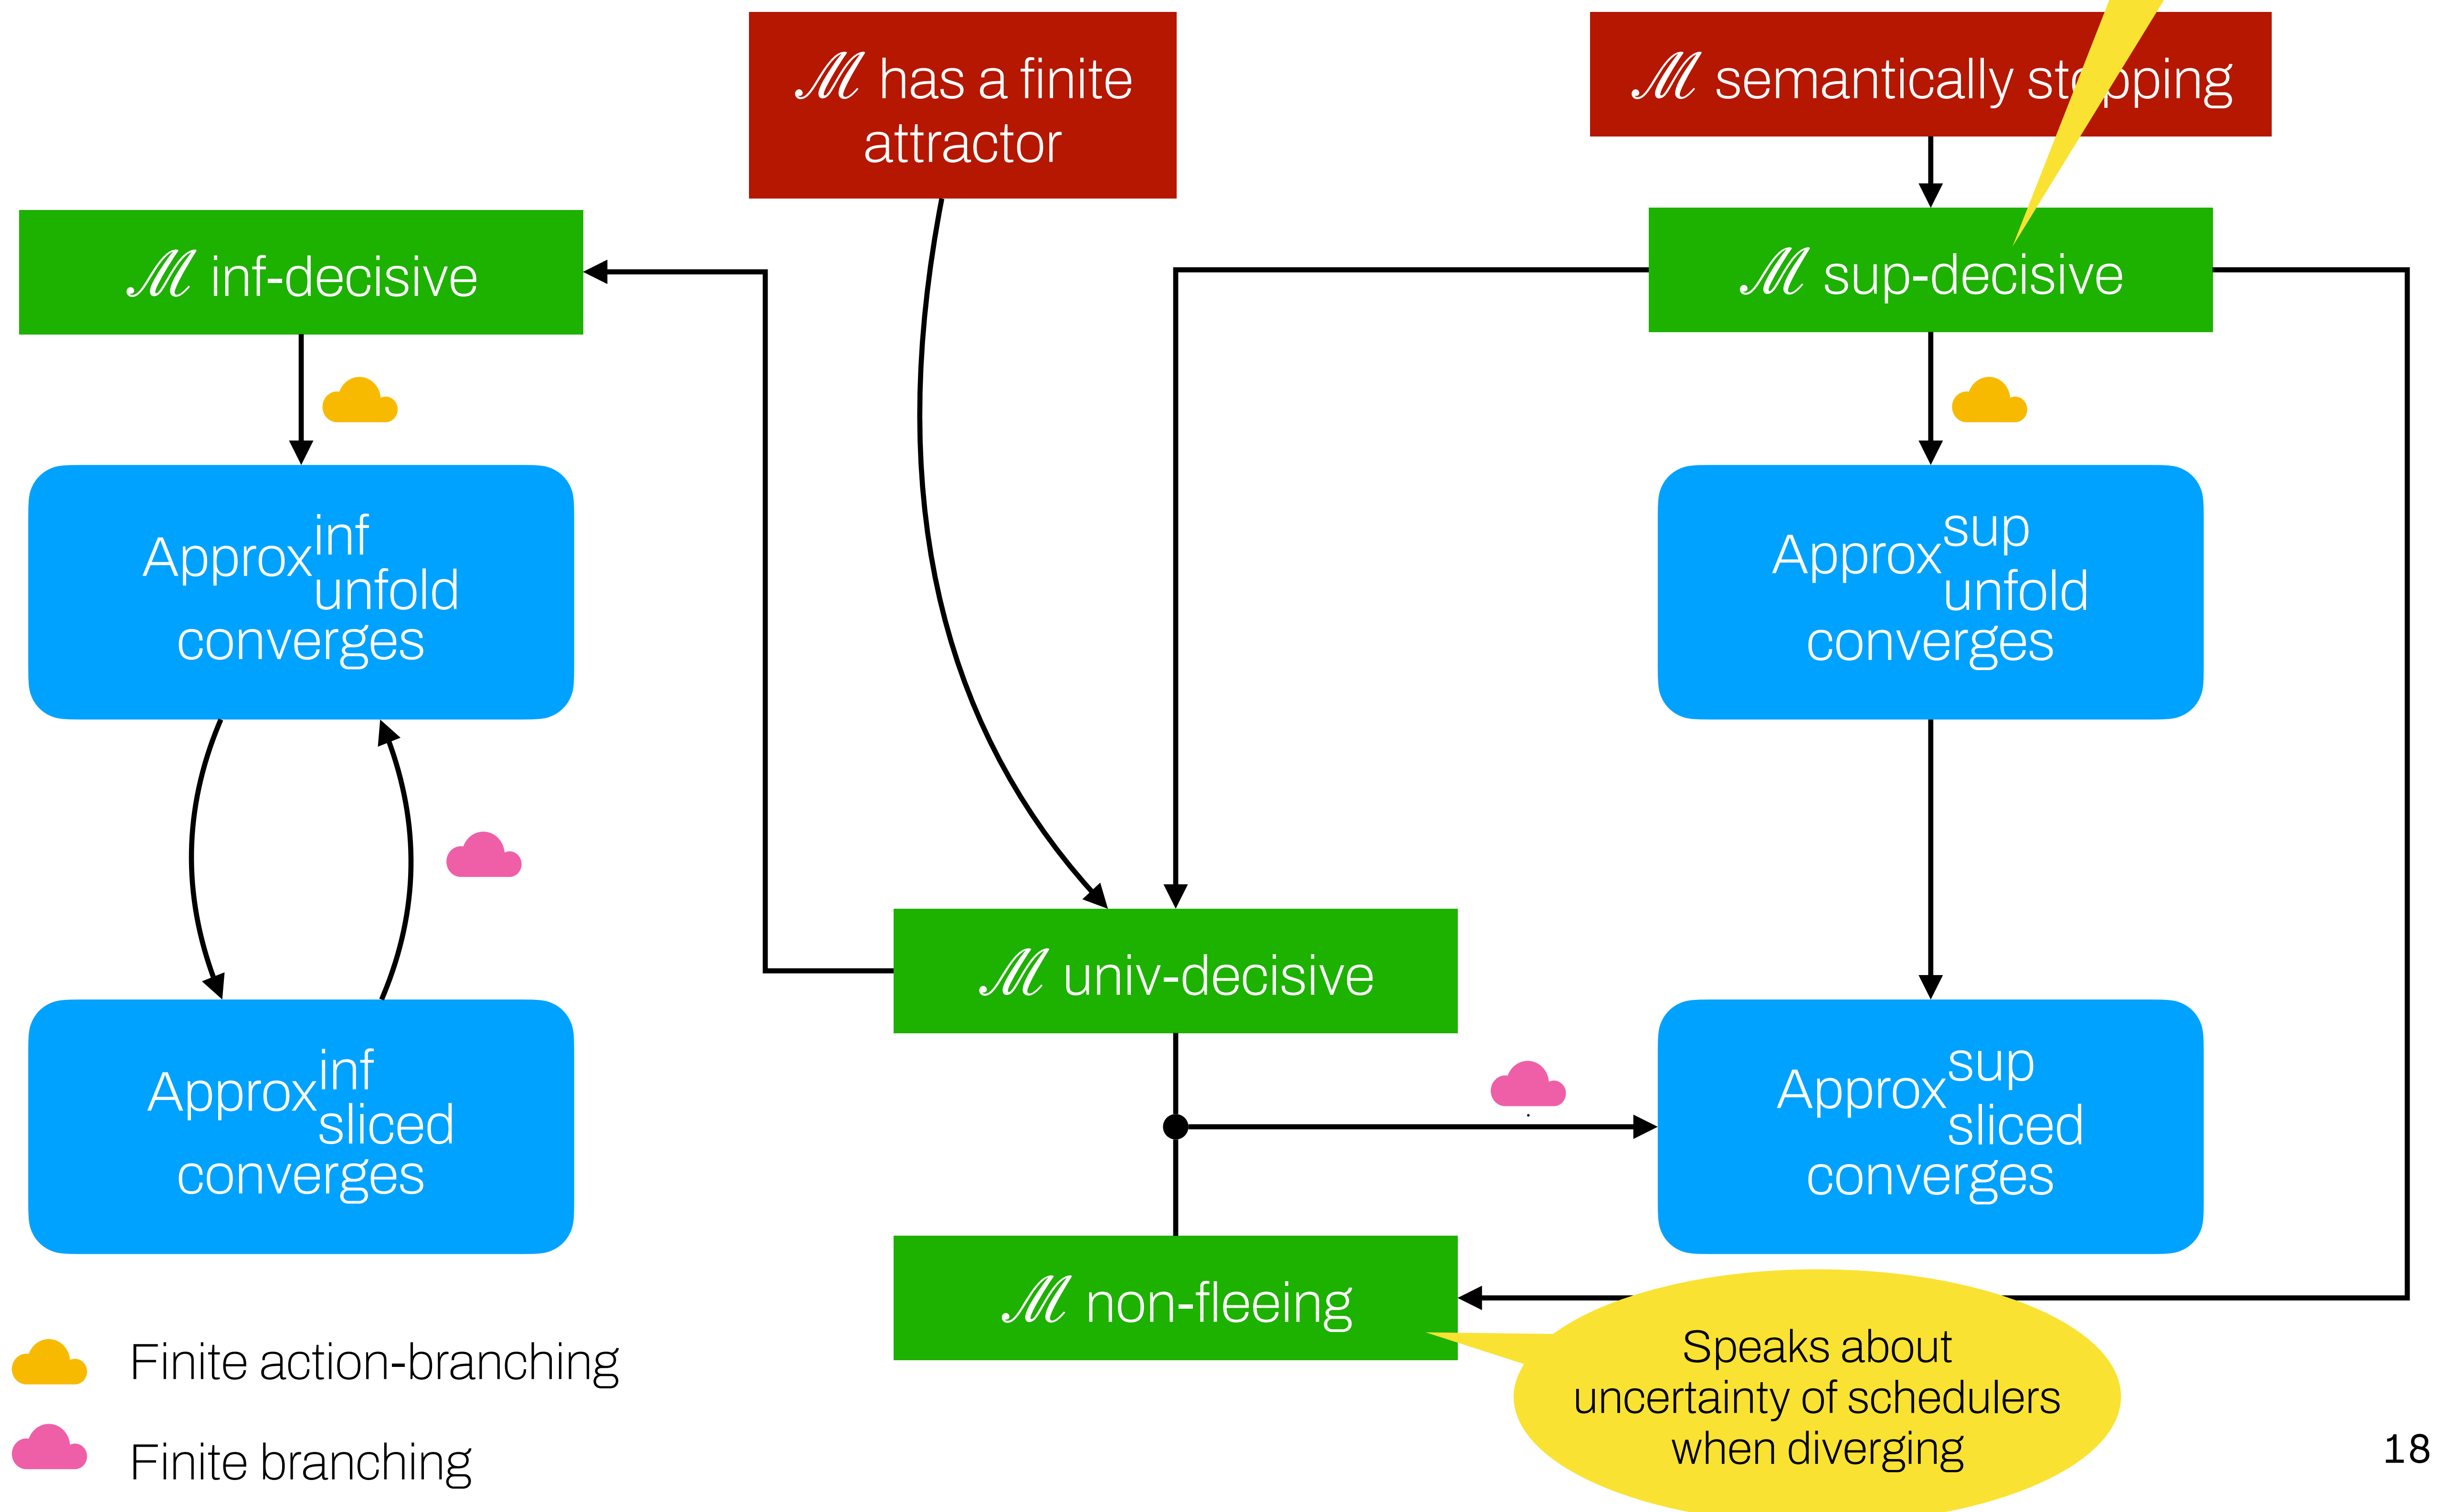

Laboratoire Méthodes Formelles

Applications

Two families of models Christel has worked on

NPLCS

- NPLCS = non-deterministic and probabilistic lossy channel systems

NPLCS

Verifying nondeterministic probabilistic channel systems against ω-regular linear-time properties

CHRISTEL BAIER
Universität Bonn, Institut für Informatik I
and
NATHALIE BERTRAND and PHILIPPE SCHNOEBELEN
Lab. Specification & Verification, CNRS & ENS de Cachan

- NPLCS = non-deterministic and probabilistic lossy channel systems

NPLCS

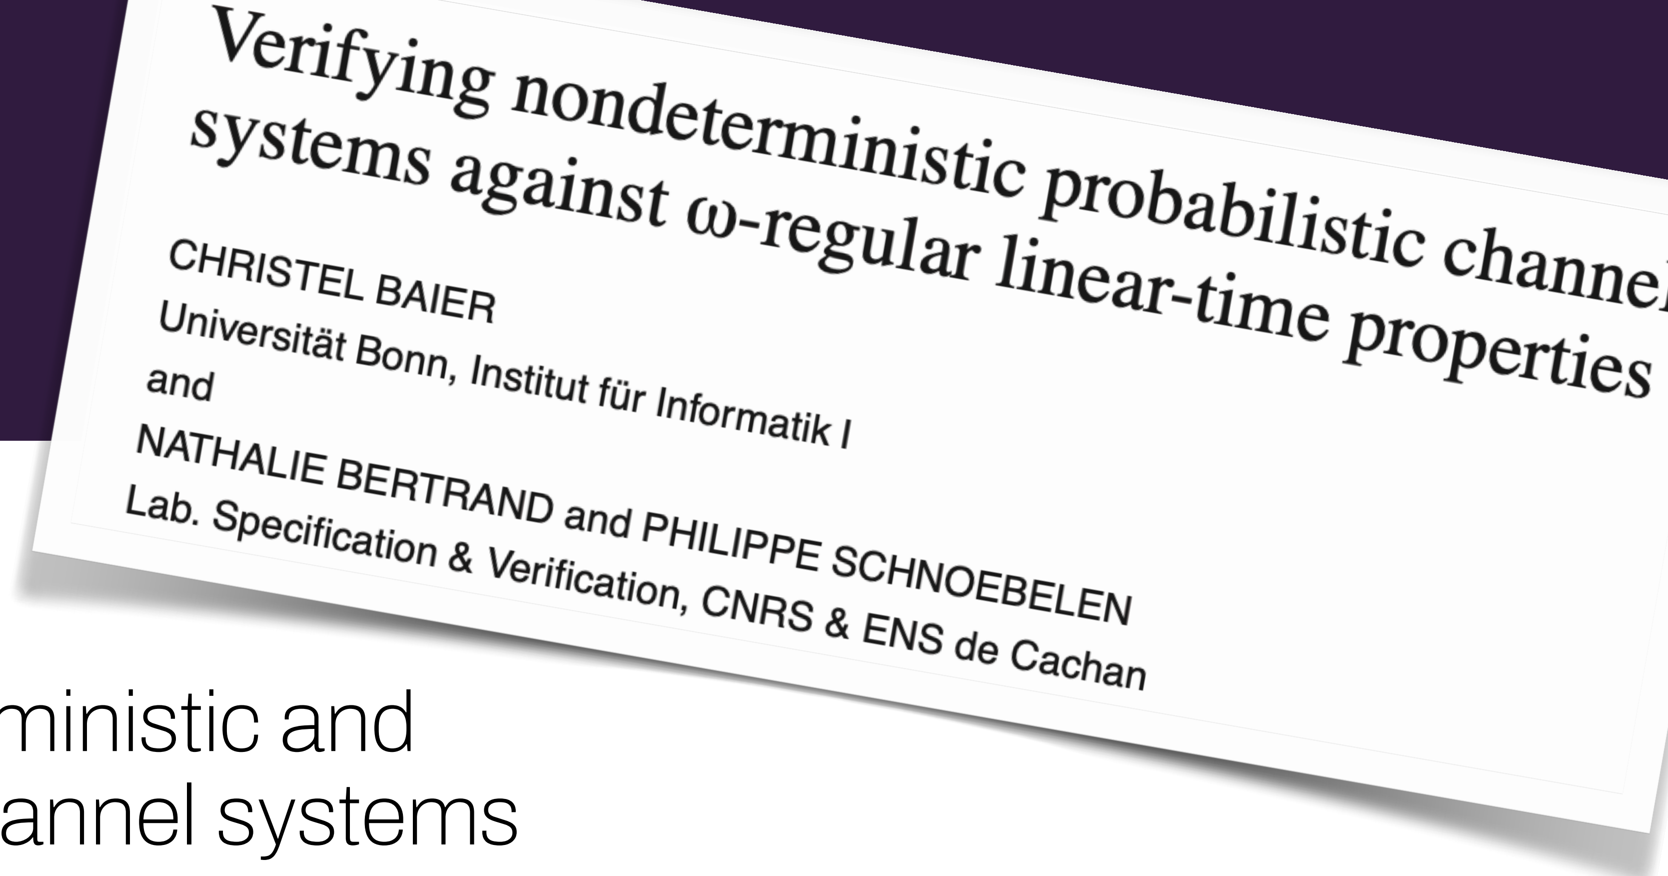

Verifying nondeterministic probabilistic channel systems against ω-regular linear-time properties

CHRISTEL BAIER
Universität Bonn, Institut für Informatik I
and
NATHALIE BERTRAND and PHILIPPE SCHNOEBELEN
Lab. Specification & Verification, CNRS & ENS de Cachan

- NPLCS = non-deterministic and probabilistic lossy channel systems

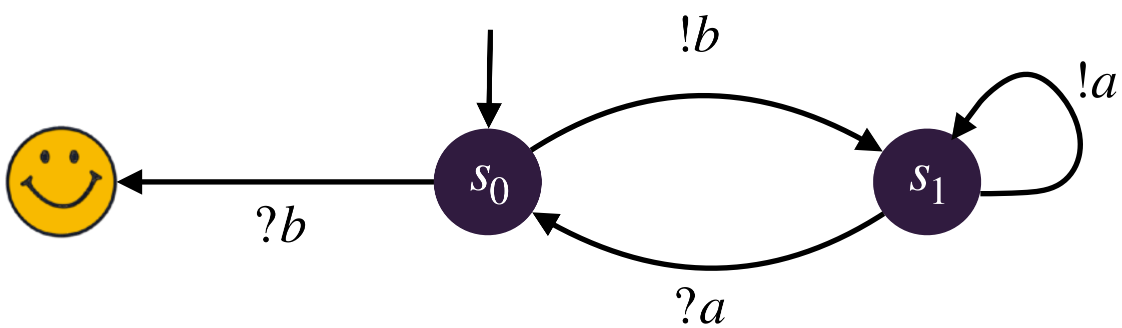

NPLCS

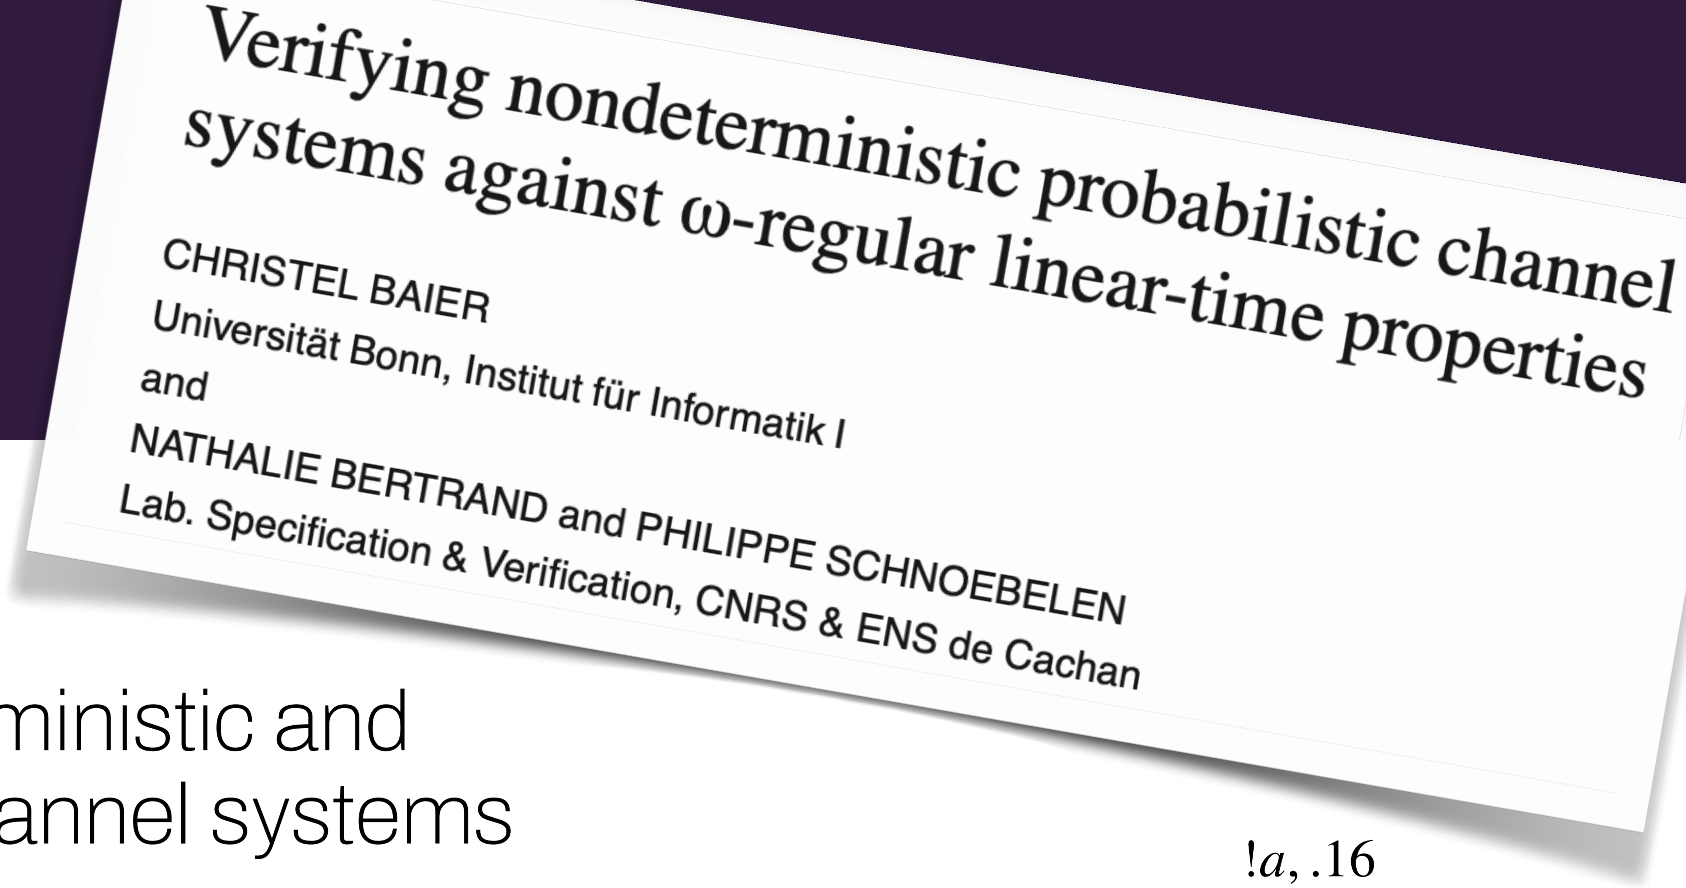

Verifying nondeterministic probabilistic channel systems against ω-regular linear-time properties

CHRISTEL BAIER
Universität Bonn, Institut für Informatik I
and
NATHALIE BERTRAND and PHILIPPE SCHNOEBELEN
Lab. Specification & Verification, CNRS & ENS de Cachan

- NPLCS = non-deterministic and probabilistic lossy channel systems

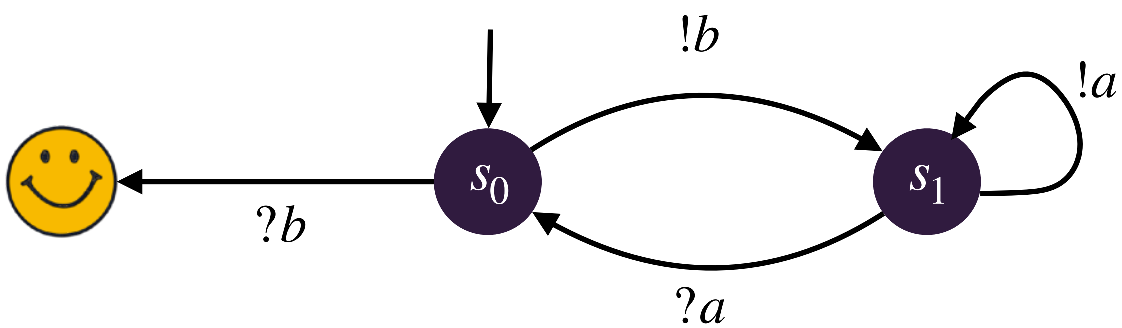

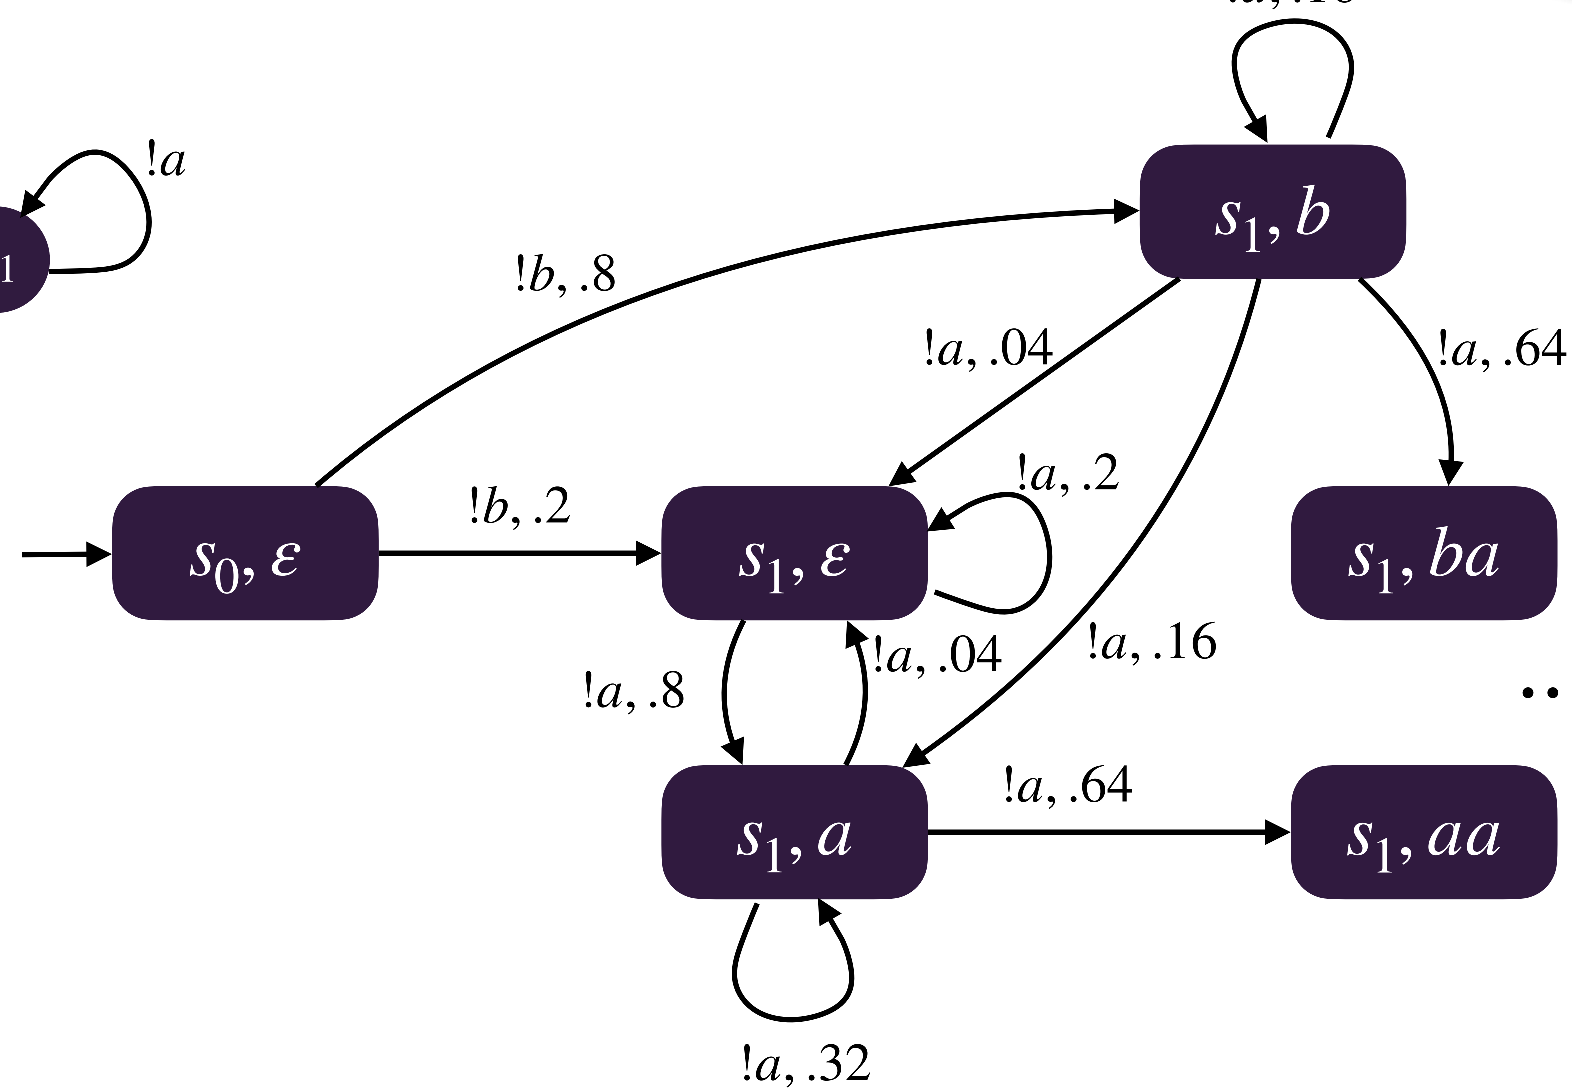

NPLCS

- The set of configurations with an empty channel is an attractor for an NPLCS [BBS07]
- Any NPLCS is univ-decisive:
  - The « inf schemes » can be applied
  - Under *nonfleeingness*, $\text{Approx}^{\text{sup}}_{\text{sliced}}$ will converge as well
- Checking $\mathbb{P}^{\text{sup}}_{\mathcal{M},s_0}\left(\mathbf{F}\ ☺\right) = 1$ is undecidable

[BBS07] C. Baier, N. Bertrand, P. Schnoebelen. Verifying nondeterministic probabilistic channel systems against ω-regular linear-time properties (ACM ToCL, 2007)

POMDPs

- POMDPs = partially observable MDPs

POMDPs

- POMDPs = partially observable MDPs

Probabilistic ω-Automata

CHRISTEL BAIER and MARCUS GRÖSSER, Technisc
NATHALIE BERTRAND, INRIA Rennes Bretagne Atlantique

POMDPs

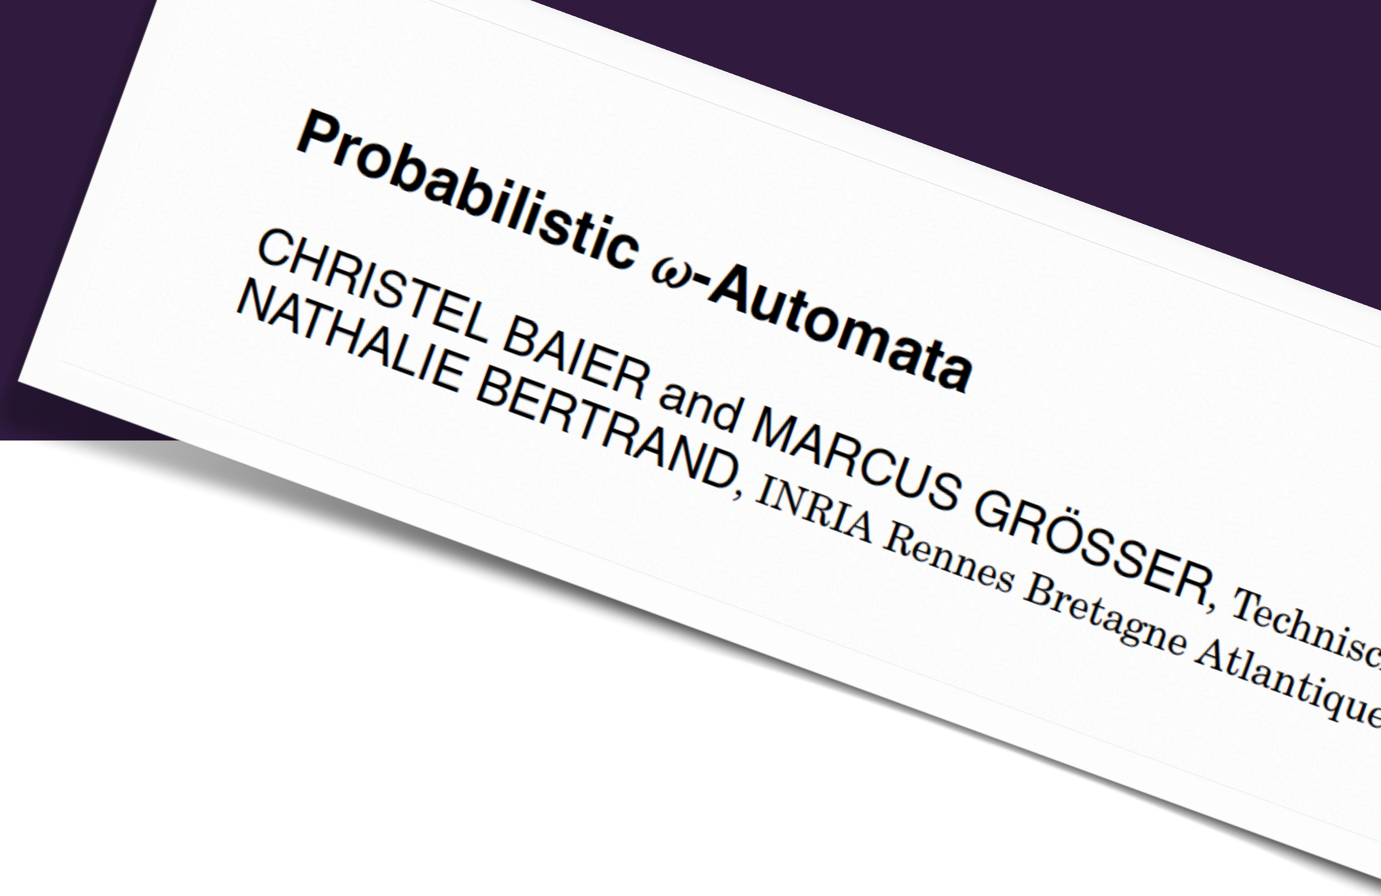

- POMDPs = partially observable MDPs

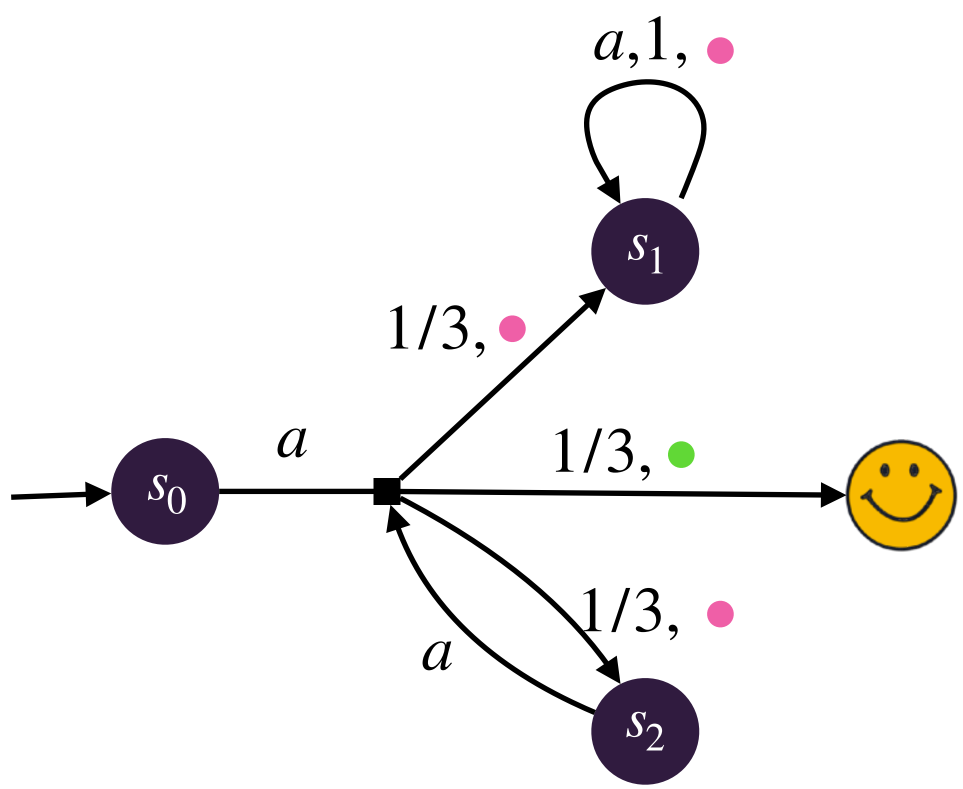

POMDPs

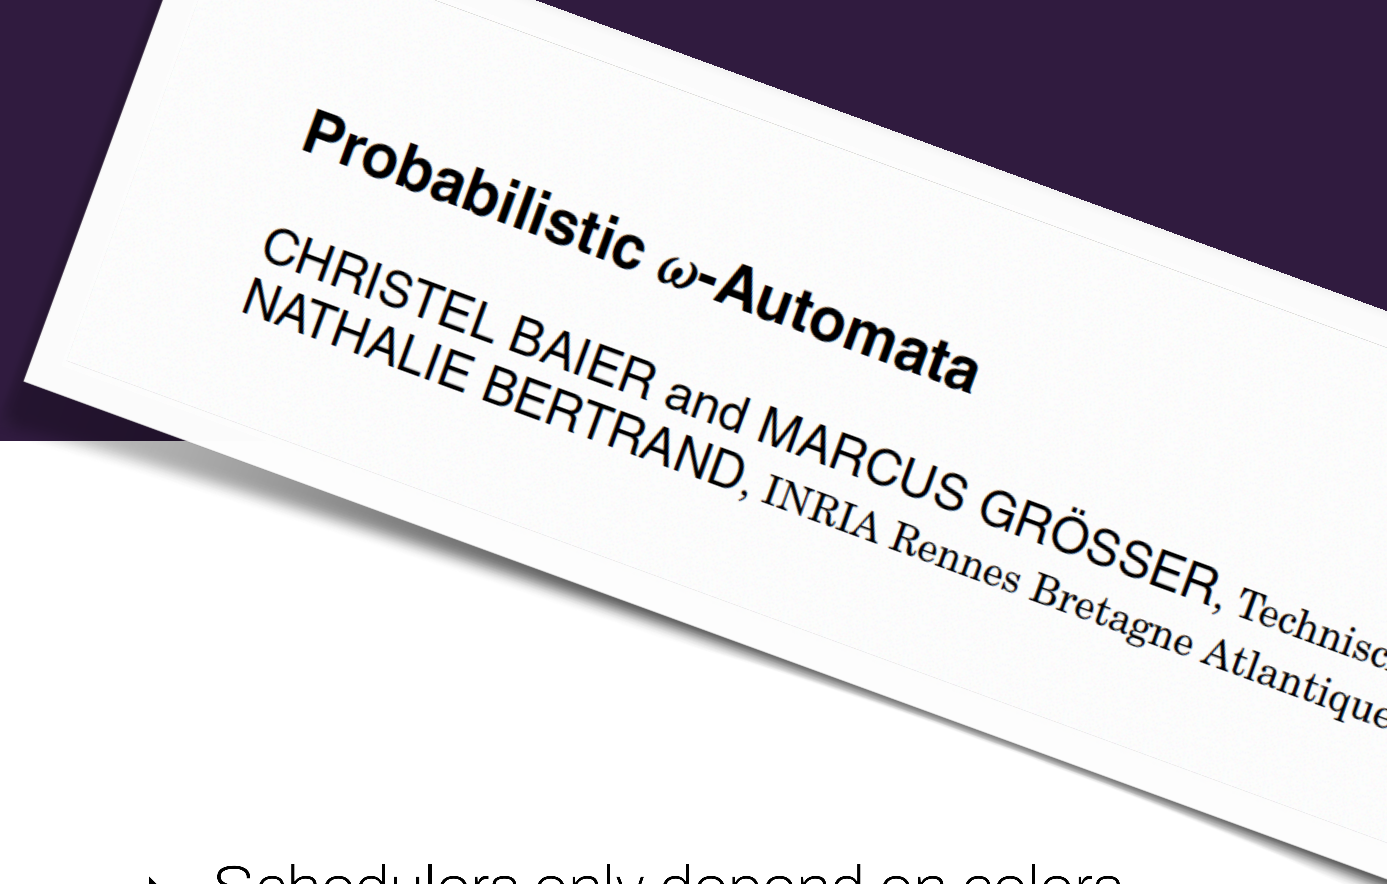
Probabilistic ω-Automata
CHRISTEL BAIER and MARCUS GRÖSSER, Technisc
NATHALIE BERTRAND, INRIA Rennes Bretagne Atlantique

- POMDPs = partially observable MDPs

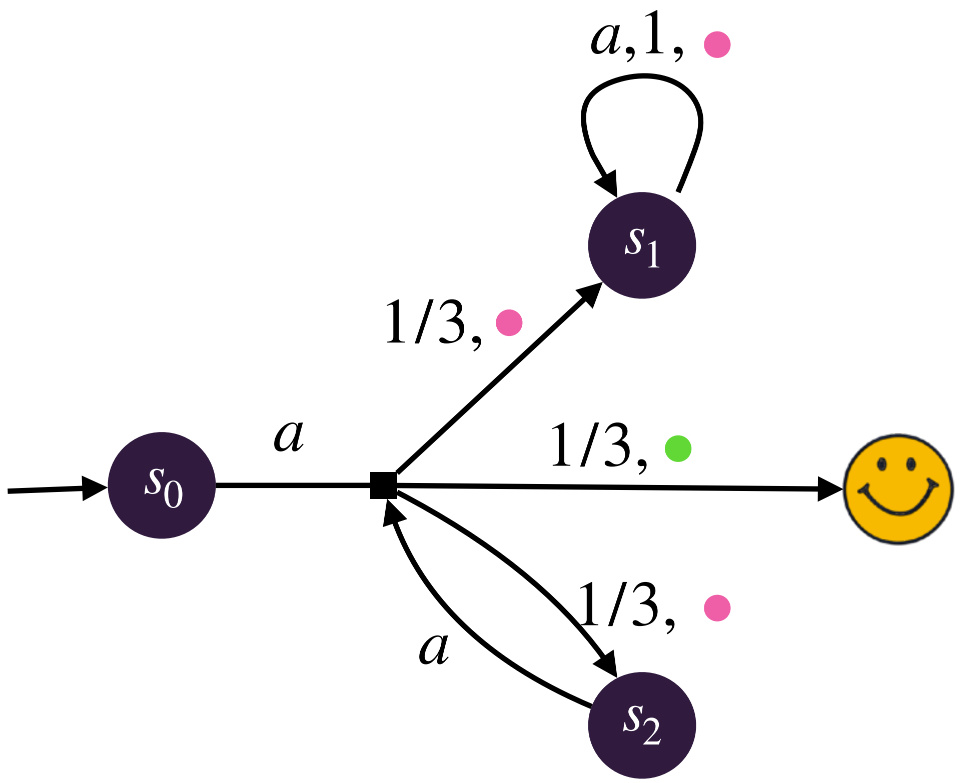

- Schedulers only depend on colors

POMDPs

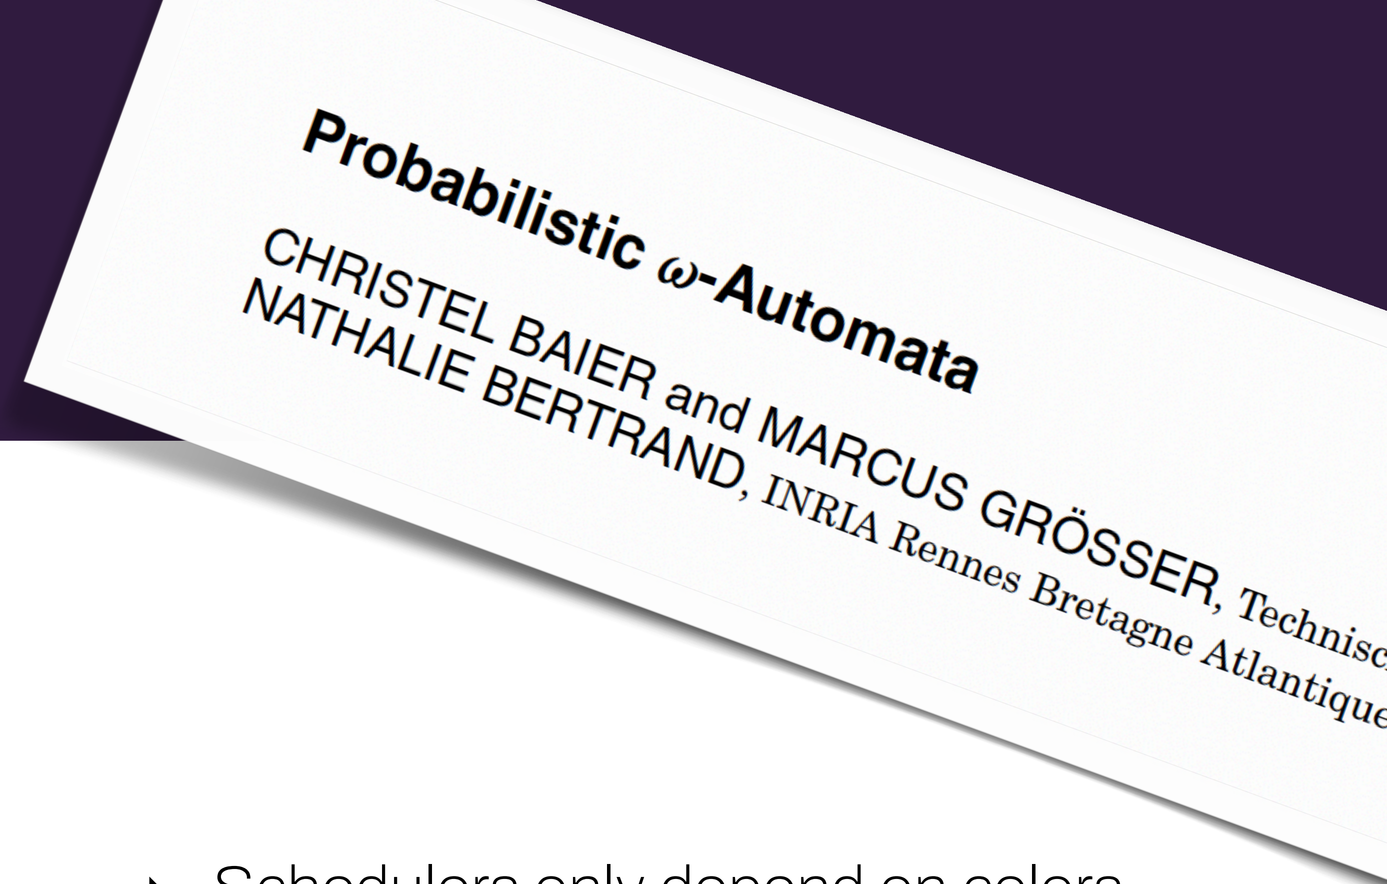

- POMDPs = partially observable MDPs

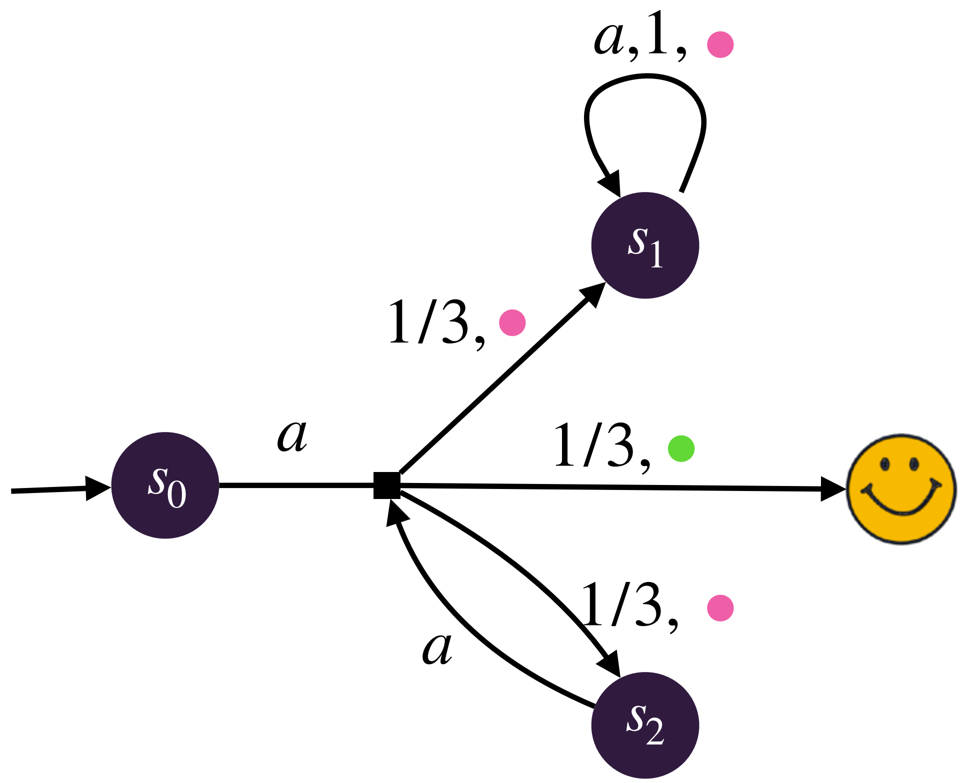

- Schedulers only depend on colors

- Checking $\mathbb{P}^{\text{sup}}_{\mathscr{M},s_0}\left(\mathbf{F}\ \right) = 1$ is undecidable [GO10]
  There is no approximation algorithm for $\mathbb{P}^{\text{sup}}(\text{reach})$ in POMDPs [MHCS99]

- What can we do for $\mathbb{P}^{\text{inf}}(\text{reach})$?

[GO10] H. Gimbert, Y. Oualhadj. Probabilistic automats on finite words: Decidable and undecidable problems (ICALP'10)
[MHC99] O. Madani, S. Hanks, A. Condon. On the undecidability of probabilistic planning and infinite-horizon partially observable Markov decision processes (1999)

POMDPs

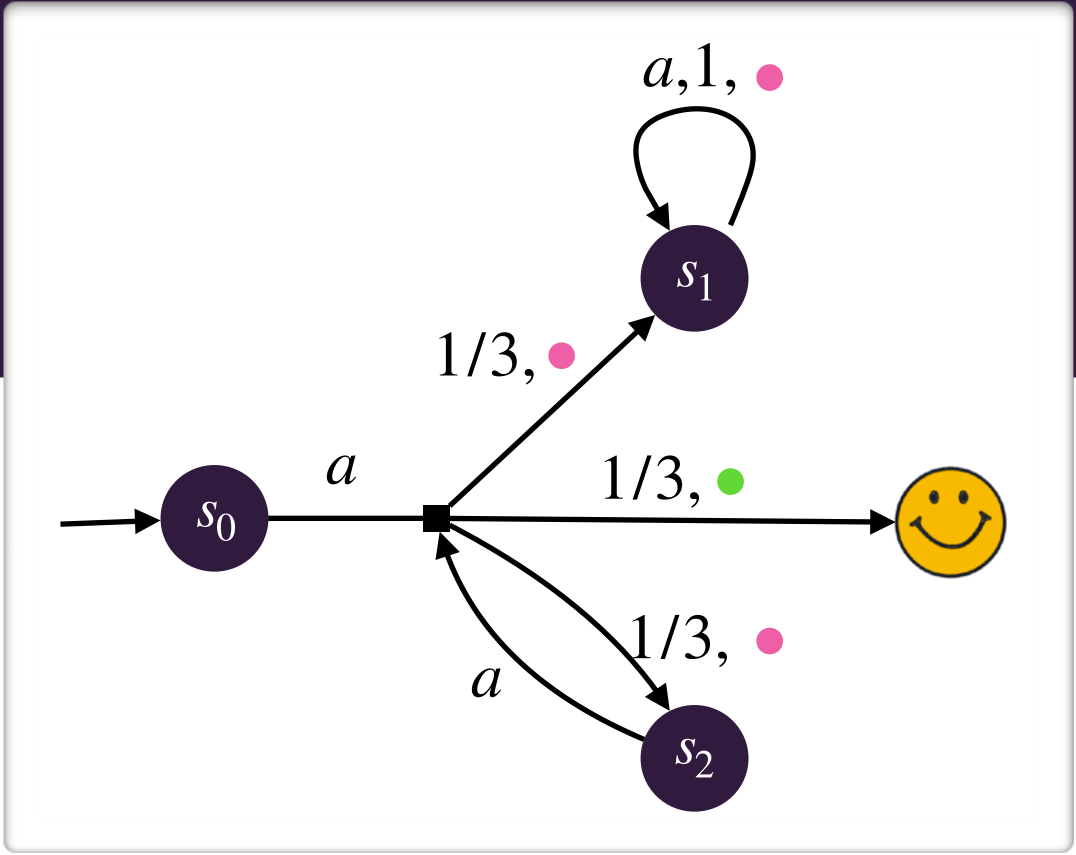

[KLC98] L.P. Kaelbing, M.L. Littman, A.R. Cassandra. Planning and acting in partially observable stochastic domains (1998)

POMDPs

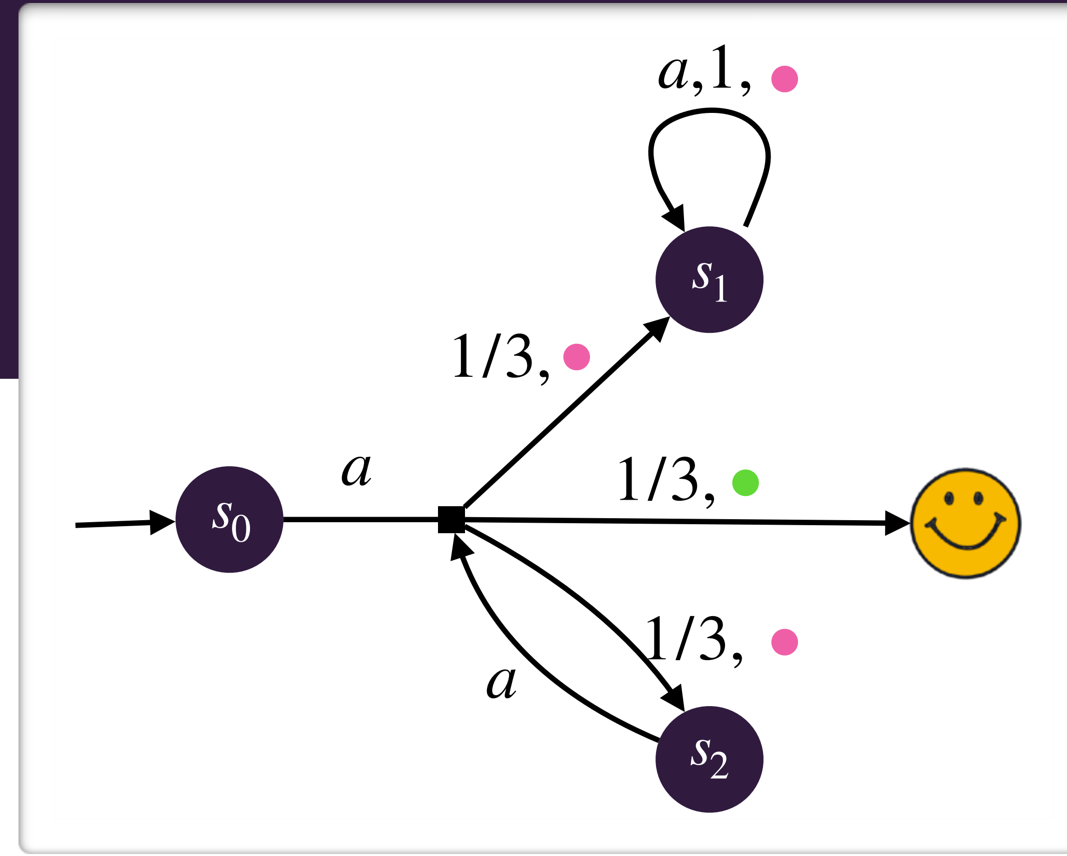

The belief MDP [KLC98]

[KLC98] L.P. Kaelbing, M.L. Littman, A.R. Cassandra. Planning and acting in partially observable stochastic domains (1998)

POMDPs

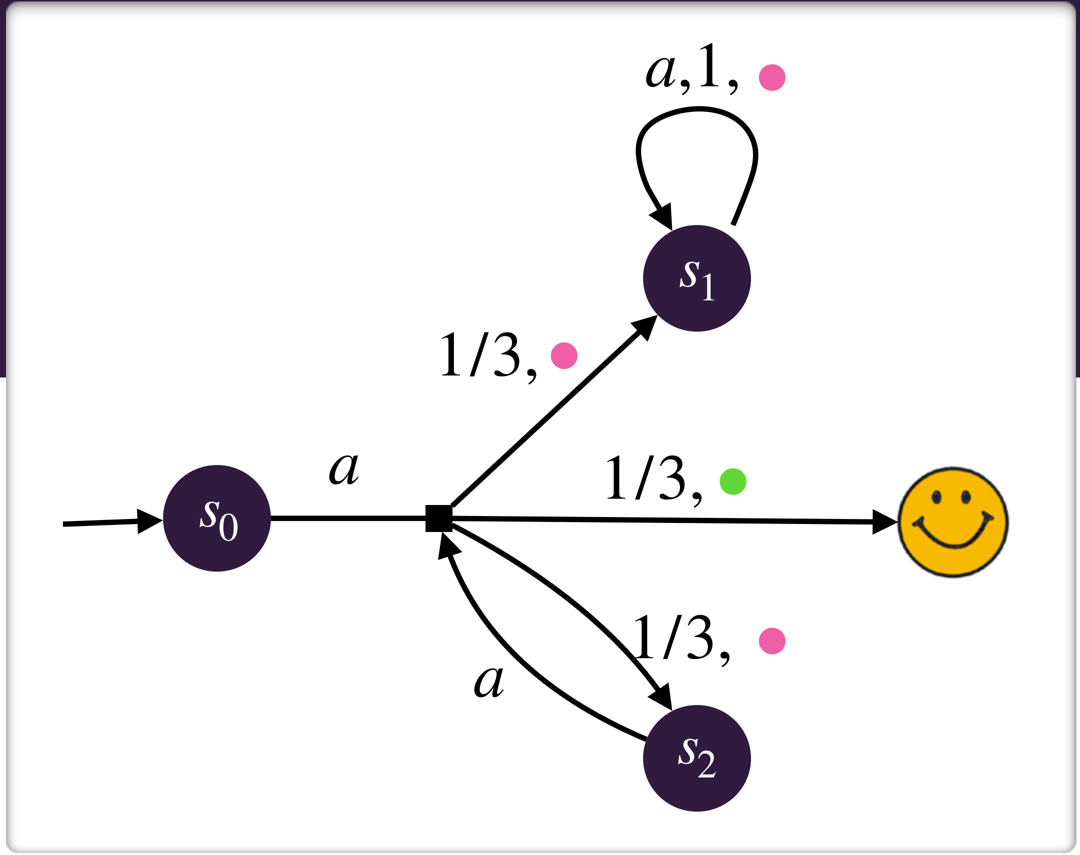

The belief MDP [KLC98]

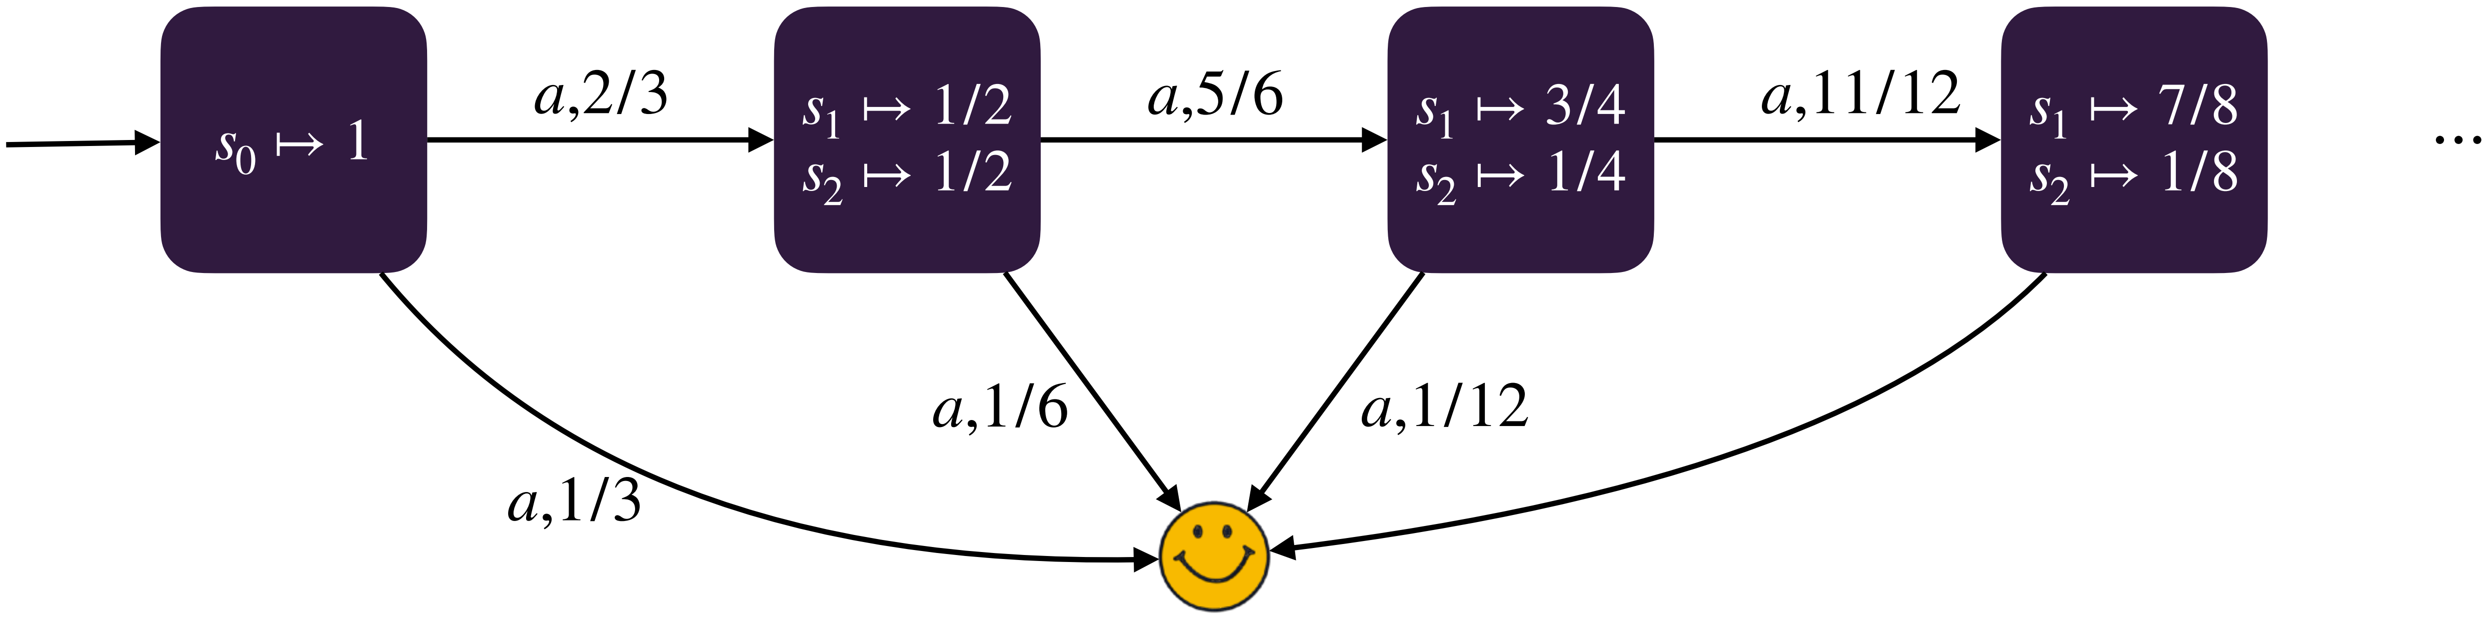

[KLC98] L.P. Kaelbing, M.L. Littman, A.R. Cassandra. Planning and acting in partially observable stochastic domains (1998)

POMDPs

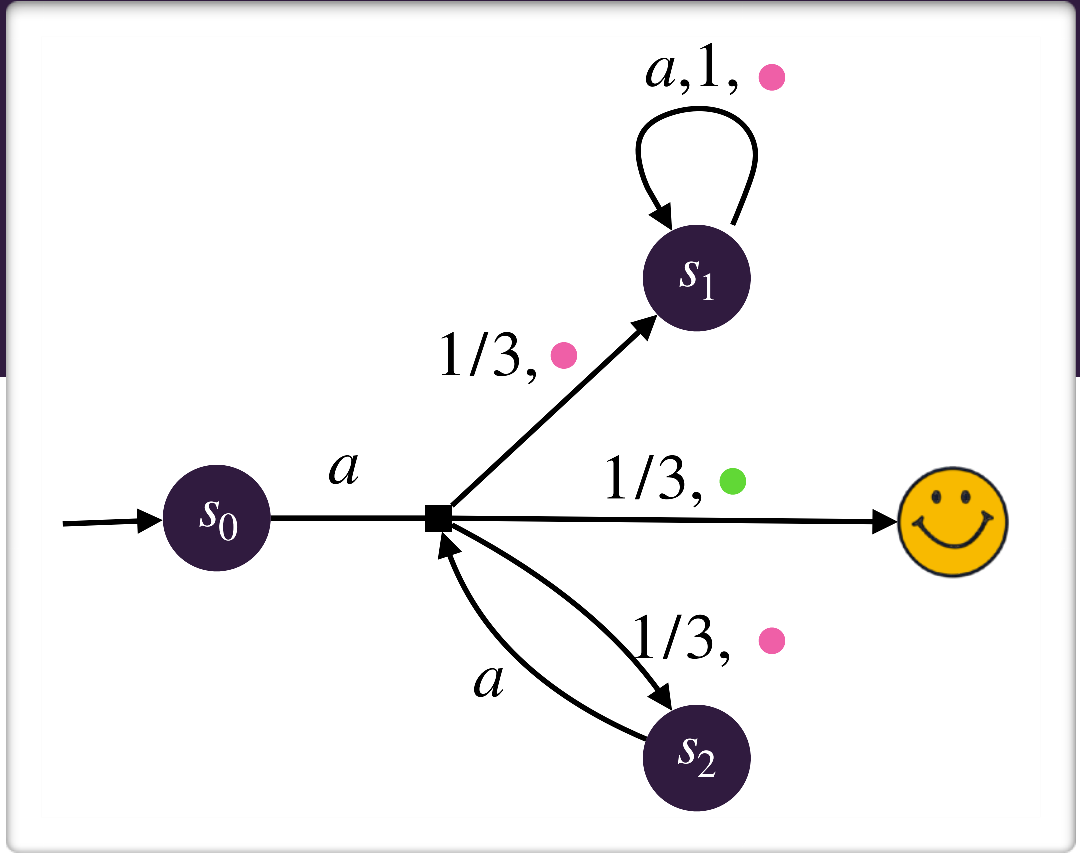

The belief MDP [KLC98]

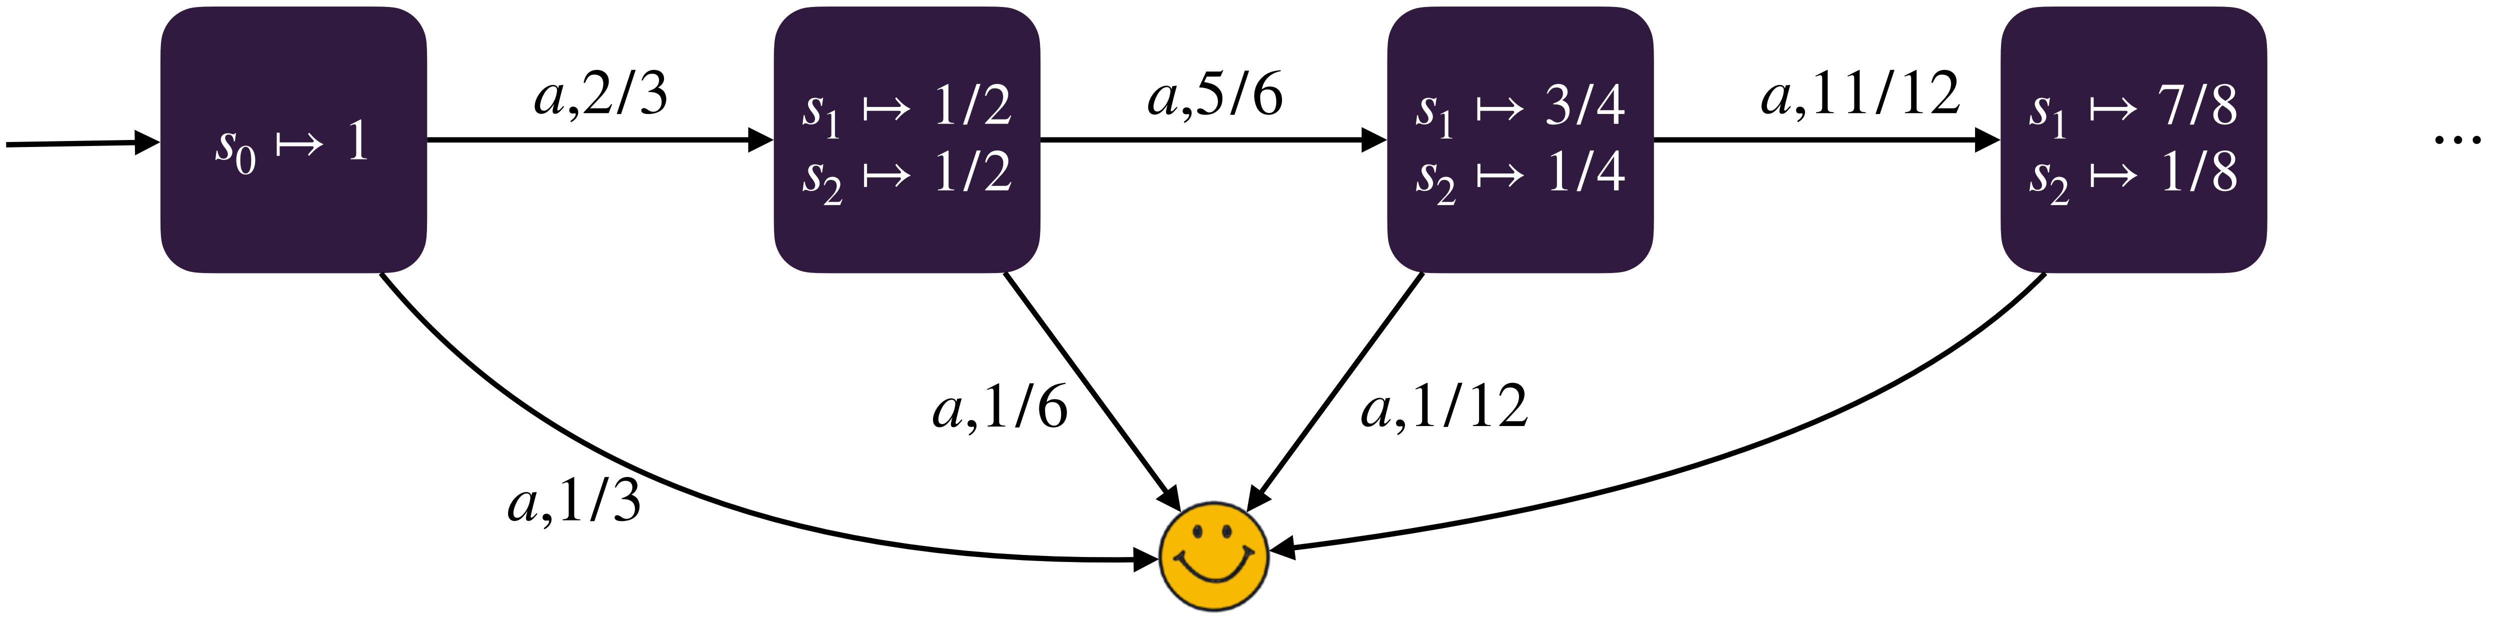

This MDP is not inf-decisive!

[KLC98] L.P. Kaelbing, M.L. Littman, A.R. Cassandra. Planning and acting in partially observable stochastic domains (1998)

POMDPs

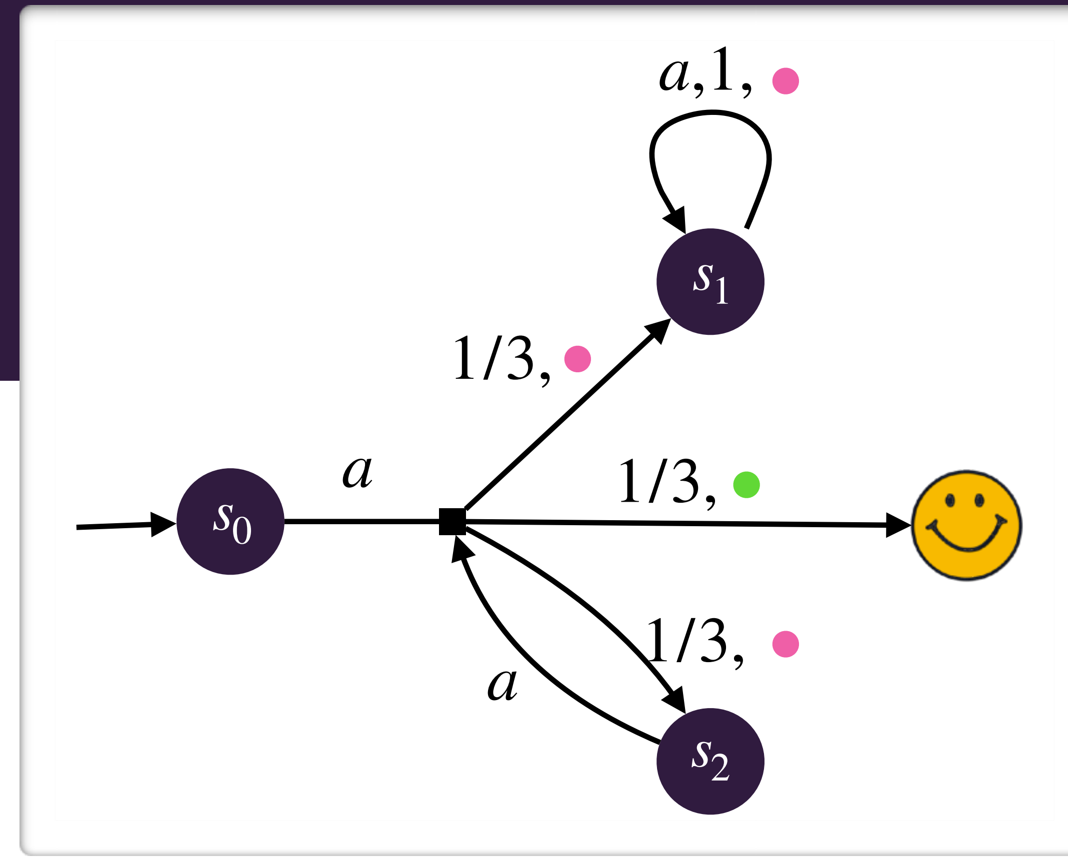

The belief MDP [KLC98]

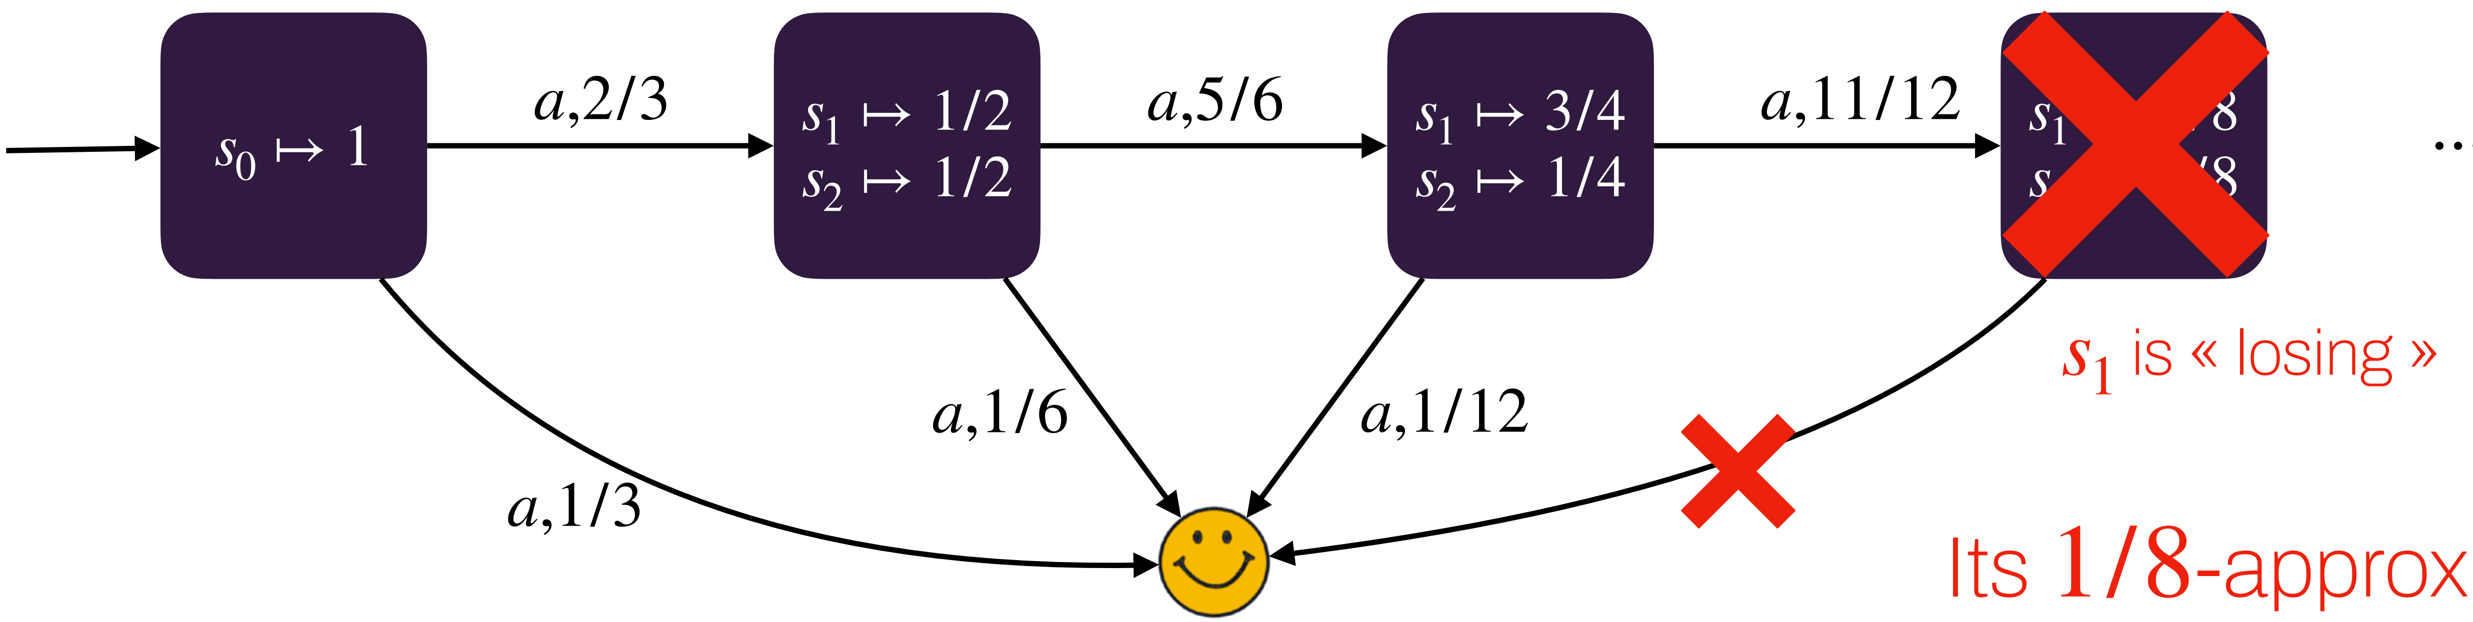

This MDP is not inf-decisive!

[KLC98] L.P. Kaelbing, M.L. Littman, A.R. Cassandra. Planning and acting in partially observable stochastic domains (1998)

POMDPs

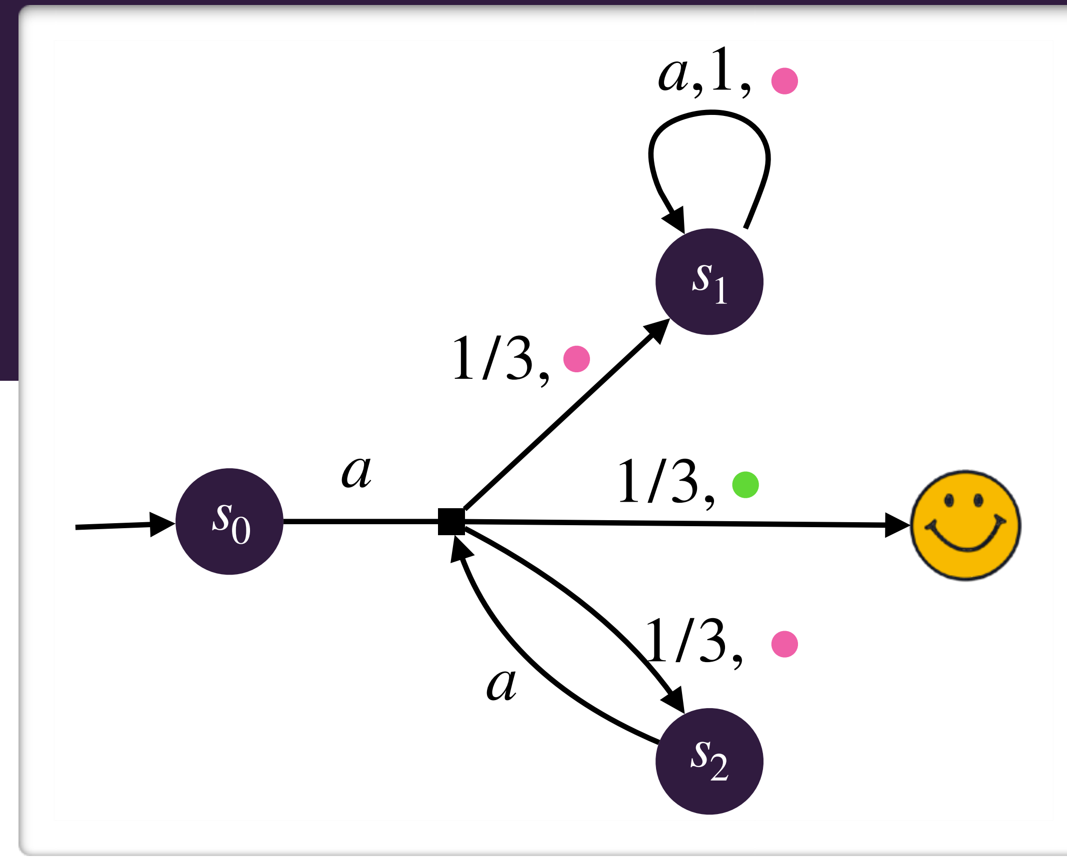

The belief MDP [KLC98]

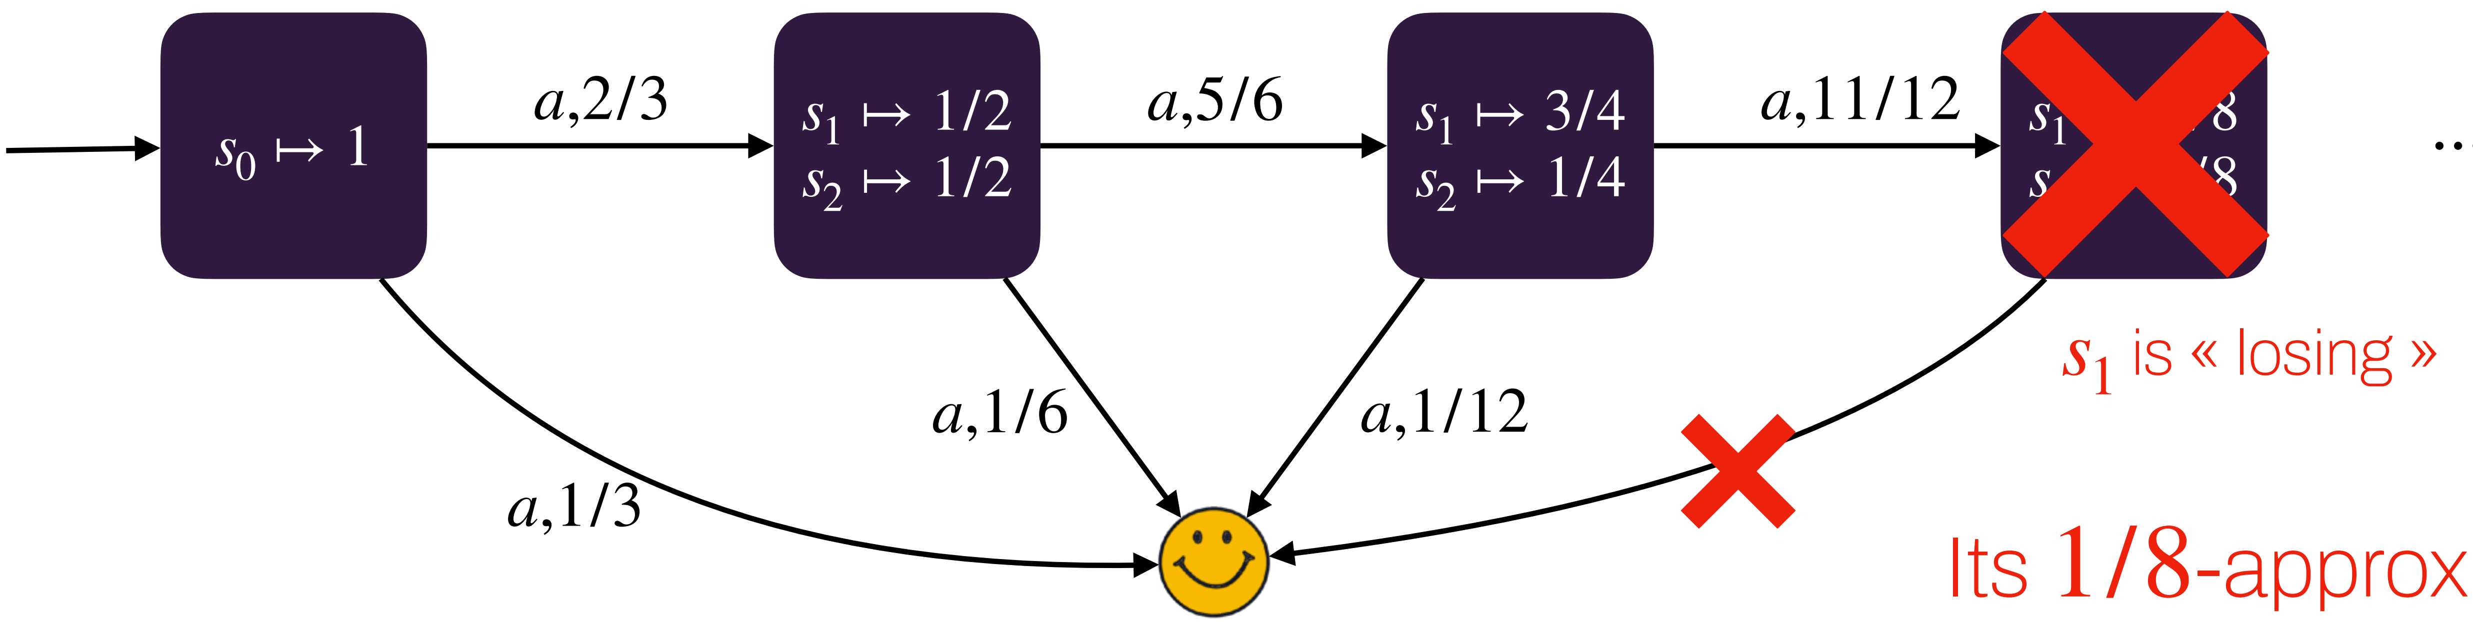

s_1 is « losing »

Its $1/8$-approx

It is inf-decisive!

This MDP is not inf-decisive!

[KLC98] L.P. Kaelbing, M.L. Littman, A.R. Cassandra. Planning and acting in partially observable stochastic domains (1998)

POMDPs

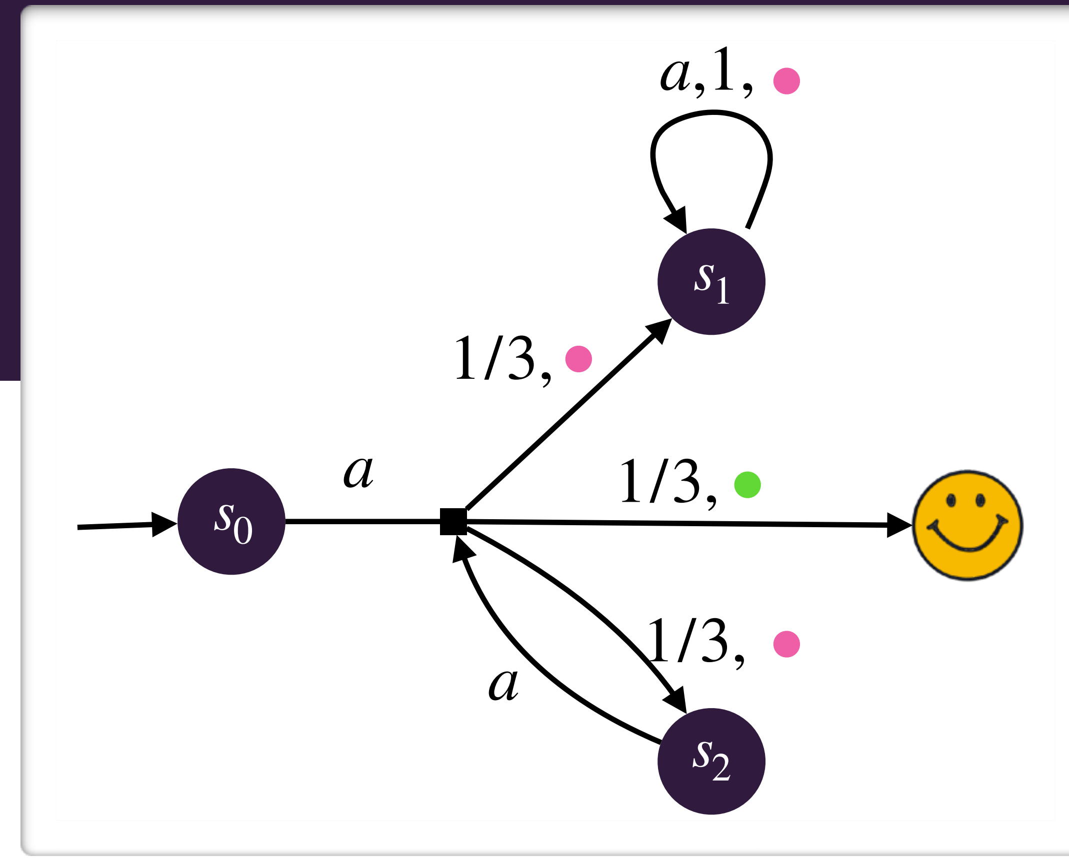

The belief MDP [KLC98]

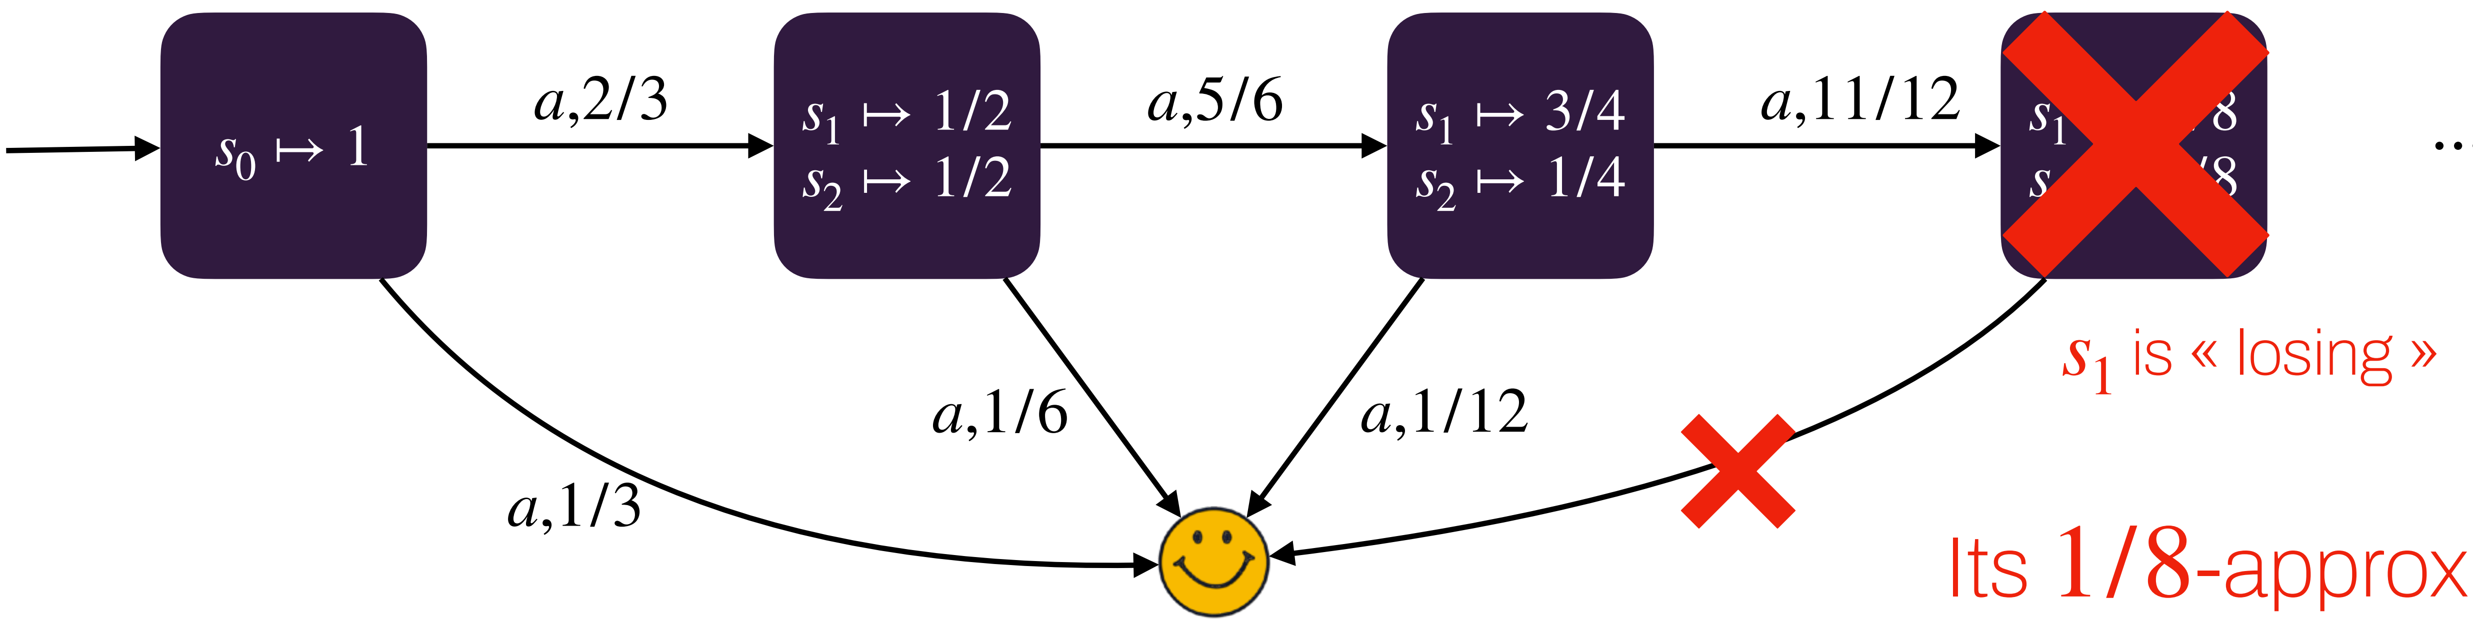

Its $1/8$-approx

It is inf-decisive!

This MDP is not inf-decisive!

- The « inf schemes » can be applied

[KLC98] L.P. Kaelbing, M.L. Littman, A.R. Cassandra. Planning and acting in partially observable stochastic domains (1998)

Conclusion

- Notions of decisiveness lifted to the setting of MDPs
- Inf-decisiveness very relevant: two classes of systems satisfy this property
- Sup-decisiveness less interesting, alternatives need to be worked out
- Further work: more expressive properties, more classes, refined notions for specific classes, ...