Model Checking Markov Chains

Lecture 3: Abstraction

Joost-Pieter Katoen

Software Modeling and Verification Group

RWTH Aachen University

affiliated to University of Twente, Formal Methods and Tools

Lecture at Model Checking Summerschool, October 12, 2010

Probabilistic bisimulation: intuition

- Strong bisimulation is used to compare labeled transition systems
- Strongly bisimilar states exhibit the same step-wise behaviour
- We like to adapt bisimulation to DTMCs
- This yields a probabilistic variant of strong bisimulation
- When do two DTMC states exhibit the same step-wise behaviour?
- Key: if their transition probability for each equivalence class coincides

for simplicity, assume a unique initial state

Probabilistic bisimulation

- Let $\mathcal{M} = (S, \mathbf{P}, AP, L)$ be a DTMC and $R \subseteq S \times S$ an equivalence
- R is a *probabilistic bisimulation* on S if for any $(s, s') \in R$:

$$L(s) = L(s') \text{ and } \mathbf{P}(s, C) = \mathbf{P}(s', C) \quad \text{for all} \quad C \text{ in } S/R$$

where $\mathbf{P}(s, C) = \sum_{s' \in C} \mathbf{P}(s, s')$ [Larsen & Skou, 1989]

- $s \sim s'$ if $\exists$ a probabilistic bisimulation R with $(s, s') \in R$

Example

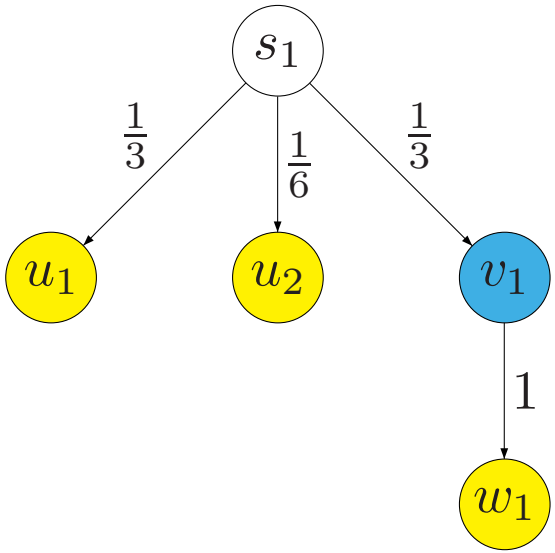
s_1
$\frac{1}{3}$
$\frac{1}{6}$
$\frac{1}{3}$
u_1
u_2
v_1
1
w_1

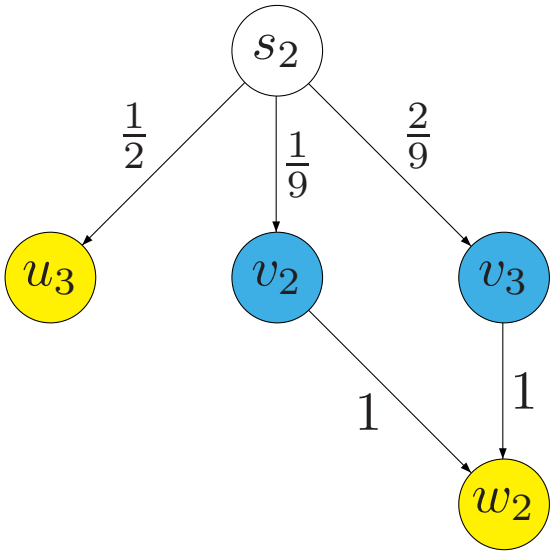
s_2
$\frac{1}{2}$
$\frac{1}{9}$
$\frac{2}{9}$
u_3
v_2
v_3
1
1
w_2

Quotient DTMC under $\sim$

$\mathcal{M}/\!\sim = (S', \mathbf{P}', AP, L')$, the quotient of $\mathcal{M} = (S, \mathbf{P}, AP, L)$ under $\sim$:

- $S' = S/\!\sim = \{ [s]_\sim \mid s \in S \}$
- $\mathbf{P}'([s]_\sim, C) = \mathbf{P}(s, C)$
- $L'([s]_\sim) = L(s)$

get $\mathcal{M}/\!\sim$ by partition-refinement in time $\mathcal{O}(M \cdot \log N + |AP| \cdot N)$ [Derisavi et al., 2001]

A DTMC model of Craps

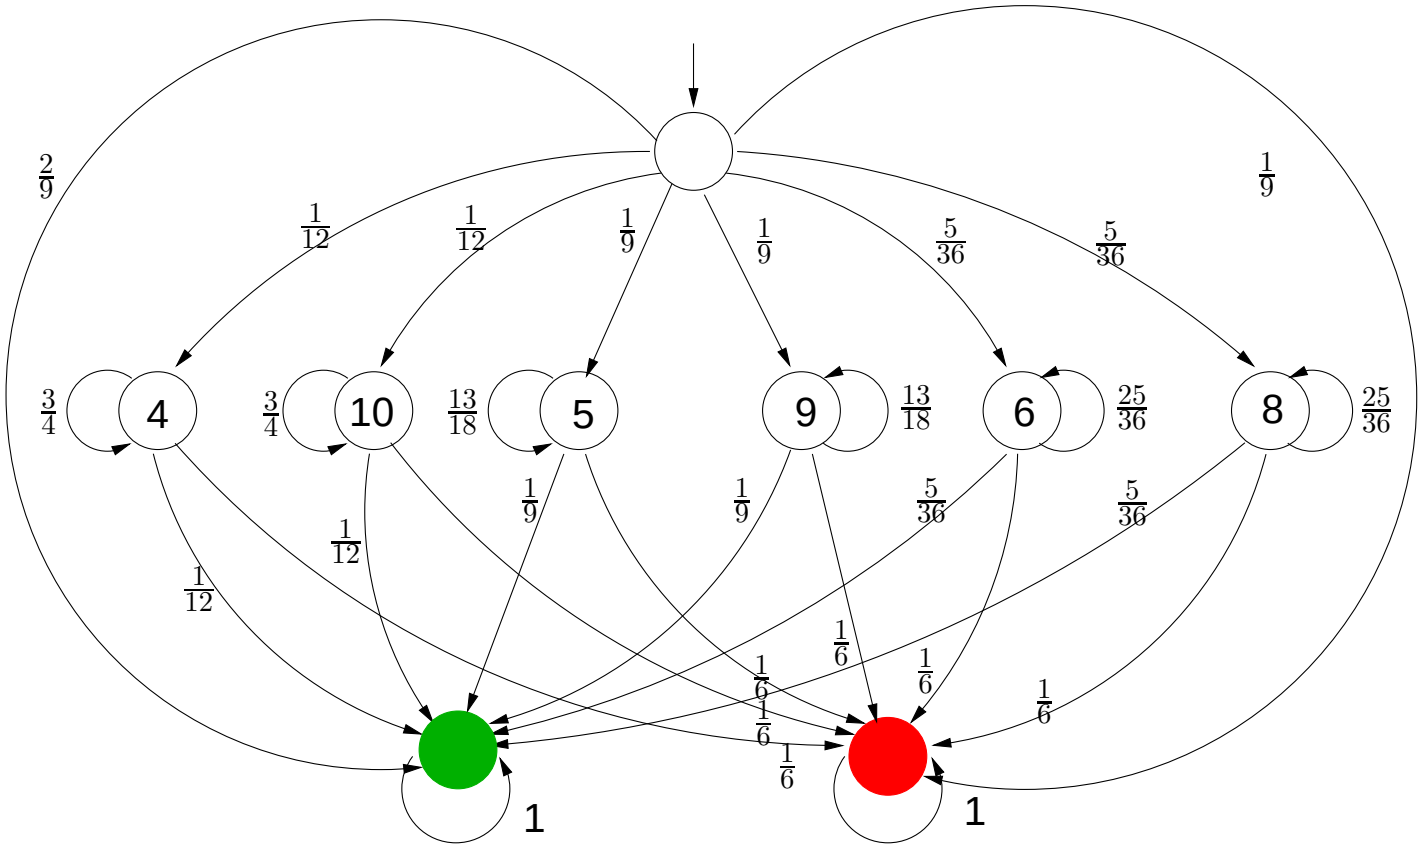

Minimizing Craps

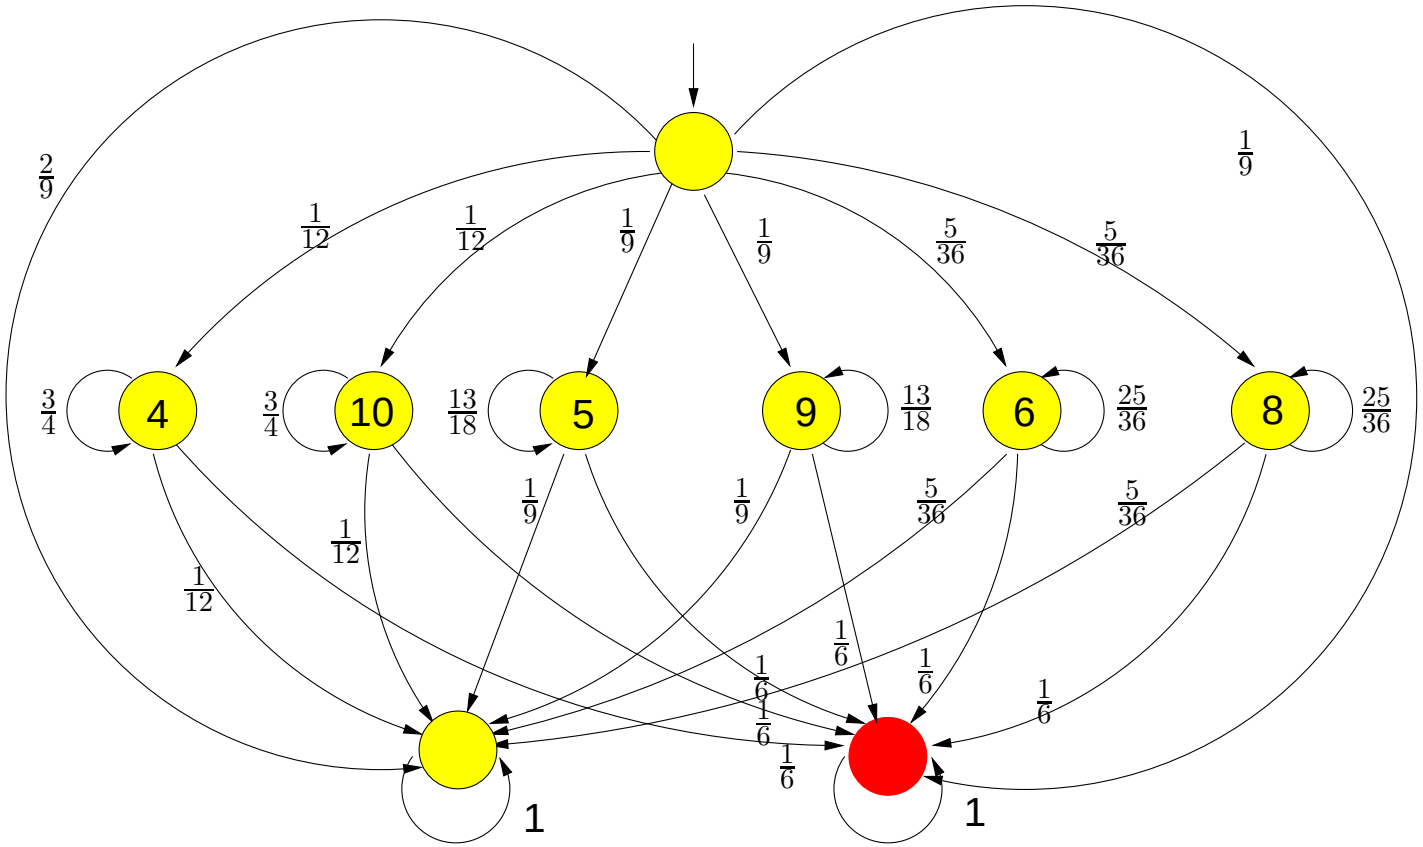

initial partitioning for the atomic propositions $AP = \{ \mathit{loss} \}$

A first refinement

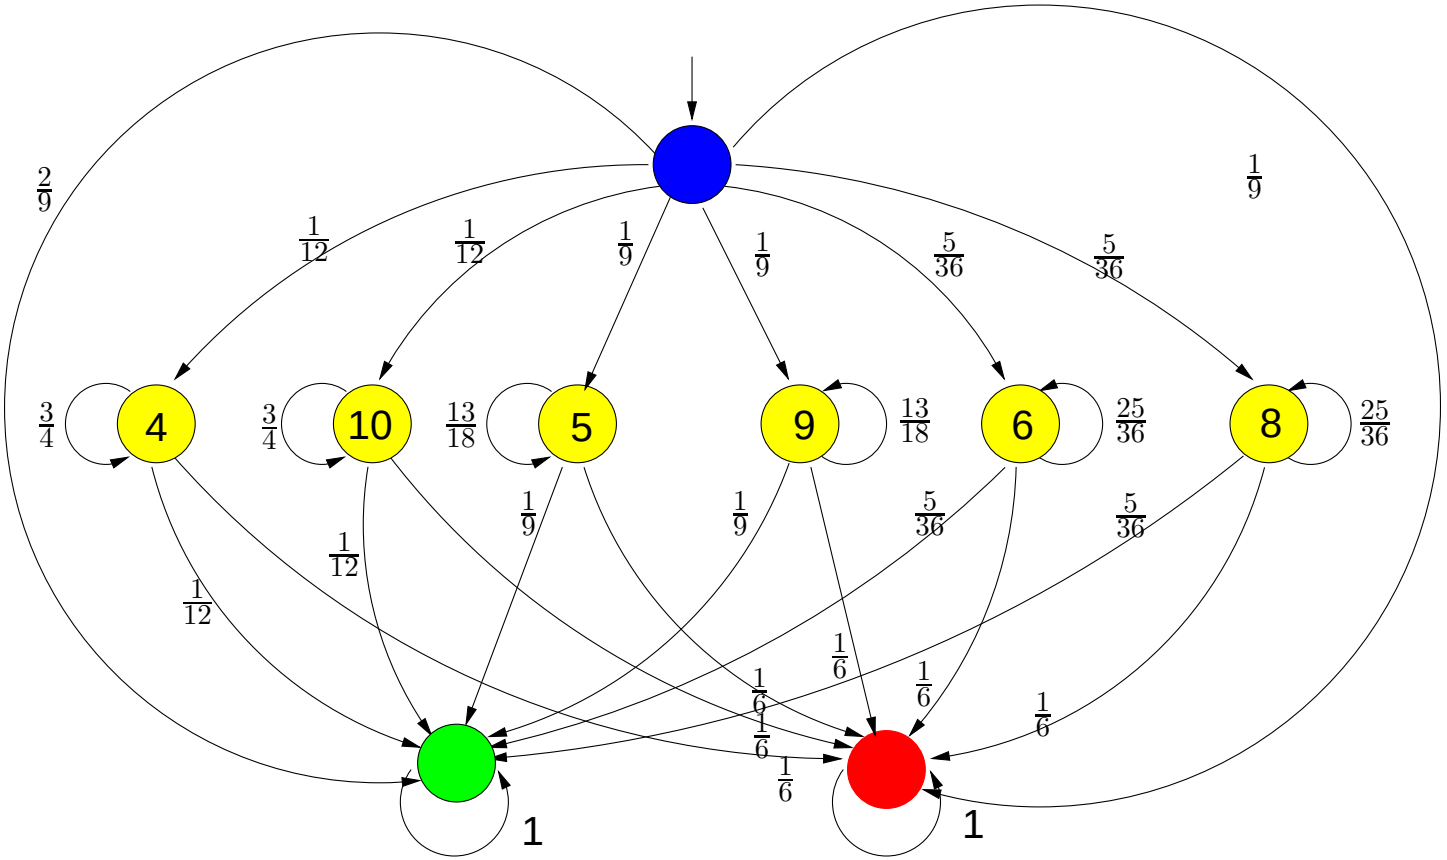

refine (“split”) with respect to the set of red states

A second refinement

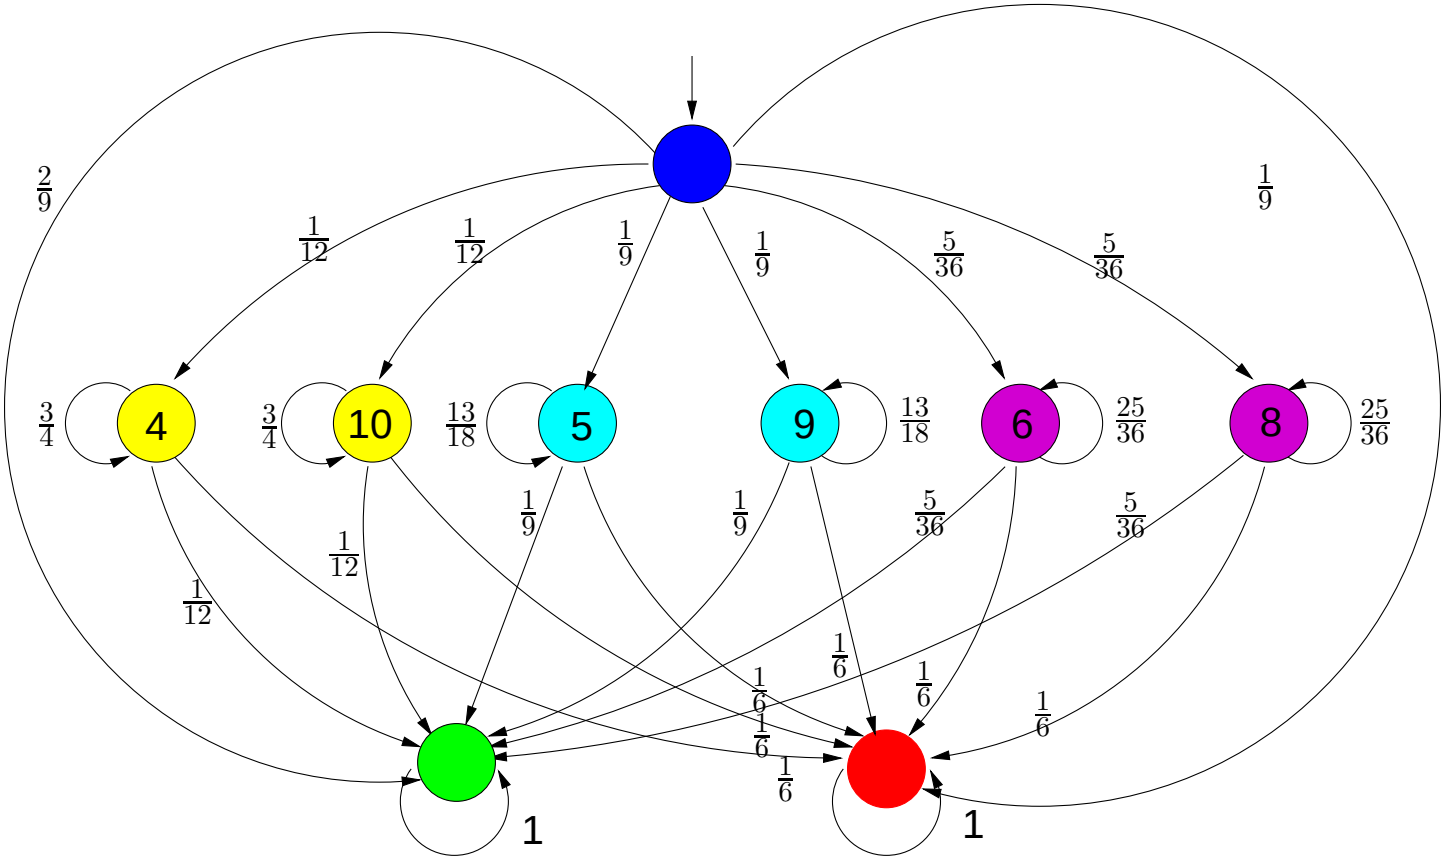

refine (“split”) with respect to the set of green states

Quotient DTMC

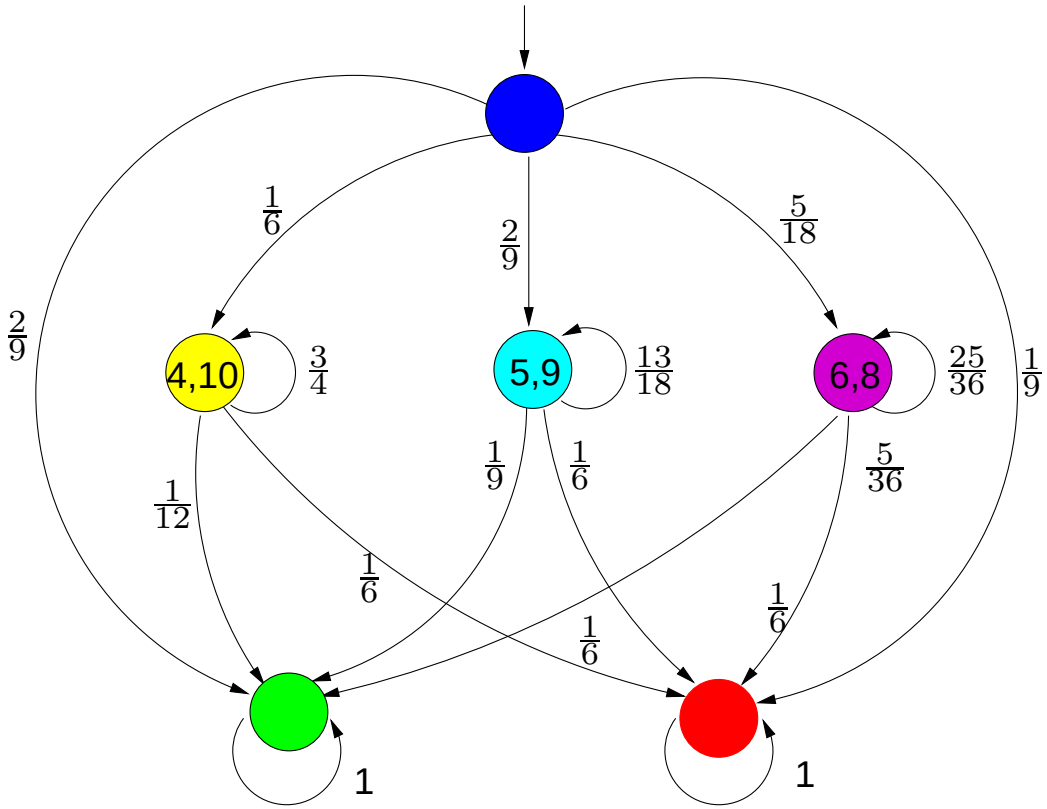

Preservation of PCTL

$$s \sim s' \;\Leftrightarrow\; (\forall \Phi \in \mathit{PCTL} : s \models \Phi \text{ if and only if } s' \models \Phi)$$

IEEE 802.11 group communication protocol

	original CTMC			lumped CTMC		red. factor	
OD	states	transitions	ver. time	blocks	lump + ver. time	states	time
4	1125	5369	121.9	71	13.5	**15.9**	**9.00**
12	37349	236313	7180	1821	642	**20.5**	**11.2**
20	231525	1590329	50133	10627	5431	**21.8**	**9.2**
28	804837	5750873	195086	35961	24716	**22.4**	**7.9**
36	2076773	15187833	5103900	91391	77694	**22.7**	**6.6**
40	3101445	22871849	7725041	135752	127489	**22.9**	**6.1**

BitTorrent-like P2P protocol

			symmetry reduction				
	original CTMC		reduced CTMC			red. factor	
N	states	ver. time	states	red. time	ver. time	states	time
2	1024	5.6	528	12	2.9	**1.93**	0.38
3	32768	410	5984	100	59	**5.48**	2.58
4	1048576	22000	52360	360	820	**20.0**	18.3

			bisimulation minimisation				
	original CTMC		lumped CTMC			red. factor	
N	states	ver. time	blocks	lump time	ver. time	states	time
2	1024	5.6	56	1.4	0.3	**18.3**	3.3
3	32768	410	252	170	1.3	**130**	2.4
4	1048576	22000	792	10200	4.8	**1324**	2.2

bisimulation may reduce a factor 66 after (manual) symmetry reduction

Weak probabilistic bisimulation

- Let $\mathcal{M} = (S, \mathbf{P}, AP, L)$ be a DTMC and $R \subseteq S \times S$ an equivalence
- R is a *weak* probabilistic bisimulation on S if for any $(s_1, s_2) \in R$:
  - $L(s_1) = L(s_2)$
  - s_1 can reach a state outside $[s_1]_R$ iff s_2 can do so
  - if $\mathbf{P}(s_i, [s_i]_R) < 1$ for $i{=}1, 2$ then:

$$\frac{\mathbf{P}(s_1, C)}{1 - \mathbf{P}(s_1, [s_1]_R)} = \frac{\mathbf{P}(s_2, C)}{1 - \mathbf{P}(s_2, [s_2]_R)} \quad \text{for all } C \in S/R, C \neq [s_1]_R$$

- $s \approx s'$ if $\exists$ a weak probabilistic bisimulation R with $(s, s') \in R$

Logical characterization

$$s \approx s' \;\Leftrightarrow\; \big(\forall \Phi \in \mathit{PCTL}_{\setminus \bigcirc} : s \models \Phi \text{ if and only if } s' \models \Phi\big)$$

Probabilistic simulation

- For transition systems, state s' simulates state s if
  - for each successor t of s there is a one-step successor t' of s' that simulates t

  $\Rightarrow$ simulation of two states is defined in terms of simulation of successor *states*

- What are successor states in the probabilistic setting?
  - the target of a transition is in fact a probability distribution

  $\Rightarrow$ the simulation relation $\sqsubseteq$ needs to be lifted from states to distributions

Weight function Δ

- Δ *"distributes"* a distribution μ over set X to one μ' over set Y
  - such that the total probability assigned by Δ to $y \in Y$ . . . equals the original probability $\mu'(y)$ on Y
  - and symmetrically for the total probability mass of $x \in X$ assigned by Δ

- Δ is *a distribution on* $R \subseteq X \times Y$ such that:
  - the probability to select (x, y) with $(x, y) \in R$ is one, and
  - the probability to select $(x, \cdot) \in R$ equals $\mu(x)$, and
  - the probability to select $(\cdot, y) \in R$ equals $\mu'(y)$

Weight function

- Let $R \subseteq S \times S$, and $\mu, \mu' \in \mathit{Distr}(S)$
- $\Delta \in \mathit{Distr}(S \times S)$ is a *weight function* for (μ, μ') and R whenever:

$$\Delta(s, s') > 0 \text{ implies } (s, s') \in R \quad \text{and}$$

$$\mu(s) = \sum_{s' \in S} \Delta(s, s') \quad \text{and} \quad \mu'(s') = \sum_{s \in S} \Delta(s, s') \text{ for any } s, s' \in S$$

- $\mu \sqsubseteq_R \mu'$ iff there exists a weight function for (μ, μ') and R

Weight function example

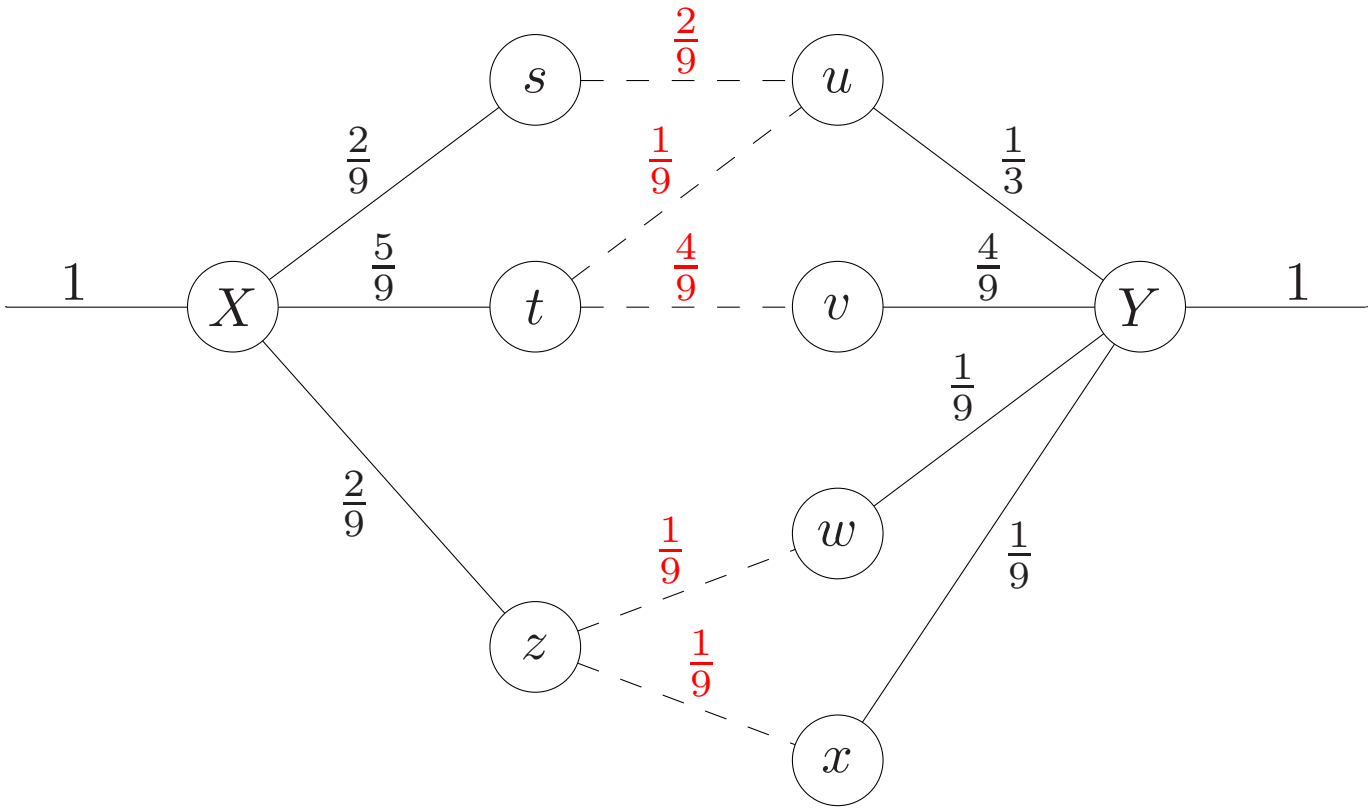

Probabilistic simulation

- Let $\mathcal{M} = (S, \mathbf{P}, AP, L)$ be a DTMC and $R \subseteq S \times S$
- R is a *probabilistic simulation* on S if for all $(s, s') \in R$:

$$L(s) \;=\; L(s') \quad \text{and} \quad \mathbf{P}(s, \cdot) \;\sqsubseteq_R\; \mathbf{P}(s', \cdot)$$

- $s \sqsubseteq_p s'$ if there exists a probabilistic simulation R with $(s, s') \in R$

Probabilistic simulation example

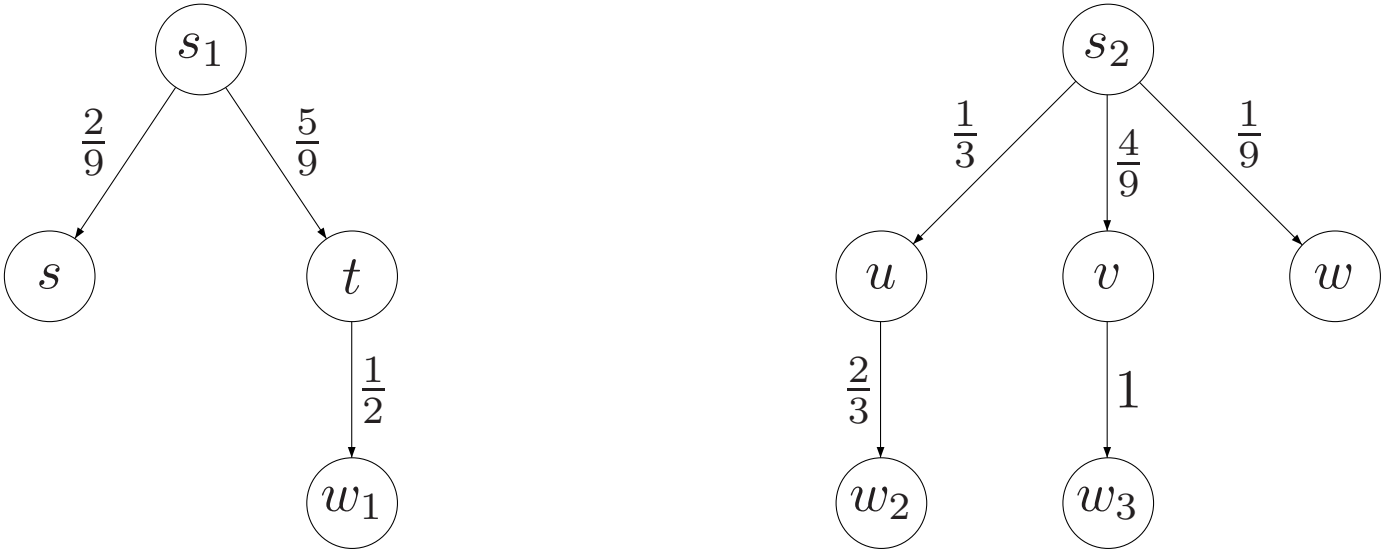

$$R = \{ (s_1, s_2), (s, u), (t, u), (t, v), (w_1, w_2), (w_1, w_3) \}$$

is a probabilistic simulation (cf. weight function before)

Simulation equivalence = bisimulation

For any DTMC:

probabilistic simulation equivalence

coincides with

probabilistic bisimulation

this does only hold for deterministic labeled transition systems

Logical characterization

$$s \sqsubseteq s' \Leftrightarrow (\forall \Phi \in \textit{safePCTL} : s' \models \Phi \text{ implies } s \models \Phi)$$

The syntax of the safe fragment of PCTL is given by:

$$\Phi ::= \textit{true} \mid a \mid \neg a \mid \Phi \wedge \Phi \mid \Phi \vee \Phi \mid \mathbb{P}_{\geqslant p}(\Phi \,\mathsf{W}\, \Phi) \mid \mathbb{P}_{\geqslant p}(\Phi \,\mathsf{W}^{\leqslant n}\, \Phi)$$

A typical safe PCTL formula: $\mathbb{P}_{\geqslant 0.99}(\Box^{\leqslant 100} \neg \textit{error})$

Overview

	strong bisimulation $\sim$	weak bisimulation $\approx$	strong simulation $\sqsubseteq$	weak simulation $\precsim$
logical preservation	PCTL	PCTL$_{\setminus\bigcirc}$	safePCTL	safePCTL$_{\setminus\bigcirc}$
checking equivalence	partition refinement $\mathcal{O}(m \log n)$	partition refinement $\mathcal{O}(n^3)$	parametric maximal flow problem $\mathcal{O}(m^2 \cdot n)$	parametric maximal flow problem $\mathcal{O}(m^2 \cdot n^3)$
graph minimization	$\mathcal{O}(m \log n)$	$\mathcal{O}(n^3)$	–	–

Can we abstract more?

- Partition the state space into groups of concrete states
  - allow any partitioning, not just grouping of bisimilar states
- Use a three-valued semantics
  - abstraction is conservative for *both* negative and positive verification results
  - if verification yields *don't know*, validity in concrete model is unknown
- Challenges:
  - what are abstract probabilistic models?
  - how to interpret PCTL on these abstract models?
  - how to verify abstractions?
  - how accurate are abstractions in practice?

The discrete-time setting

An abstract MC (AMC) is a quintuple $\mathcal{D} = (S, \mathbf{P}^l, \mathbf{P}^u, L)$ with:

- $\mathbf{P}^l, \mathbf{P}^u : S \times S \mapsto [0, 1]$, transition probability bounds where

$$\mathbf{P}^l(s, S) \leqslant 1 \leqslant \mathbf{P}^u(s, S) < \infty \quad \text{for all } s \in S$$

- $L : S \times AP \mapsto \{ \top, \bot, ? \}$, the labeling function

This is also known as interval Markov chains (Kozine & Utkin, 2002)

Abstraction

For $\mathcal{A} = \{ A_1, \ldots, A_n \}$ let AMC $\alpha(\mathcal{A}, \mathcal{D}) := (\mathcal{A}, \tilde{\mathbf{P}}^l, \tilde{\mathbf{P}}^u, \tilde{L})$ with:

$$\tilde{\mathbf{P}}^l(A_i, A_j) = \inf_{s \in A_i} \mathbf{P}^l(s, A_j) \quad \text{and} \quad \tilde{\mathbf{P}}^u(A_i, A_j) = \min\{ 1, \sup_{s \in A_i} \mathbf{P}^u(s, A_j) \}$$

and $\tilde{L}(A_i, a) = \begin{cases} \top & \text{if } L(s,a) = \top \text{ for all } s \in A_i \\ \bot & \text{if } L(s,a) = \bot \text{ for all } s \in A_i \\ ? & \text{otherwise} \end{cases}$

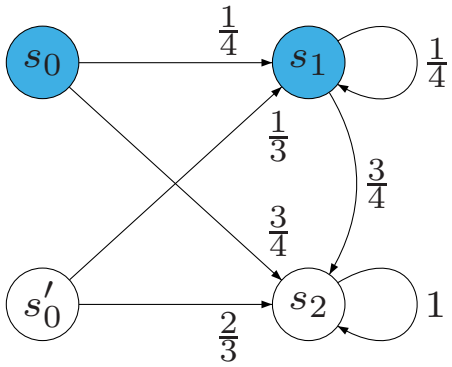

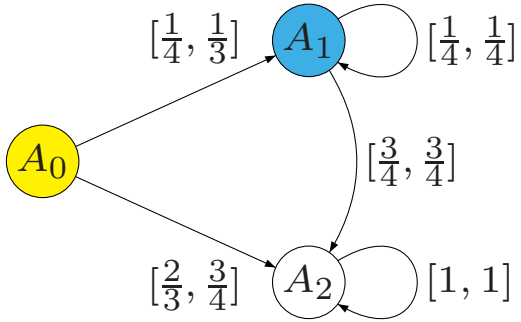

Normalization

removes illegal probability combinations

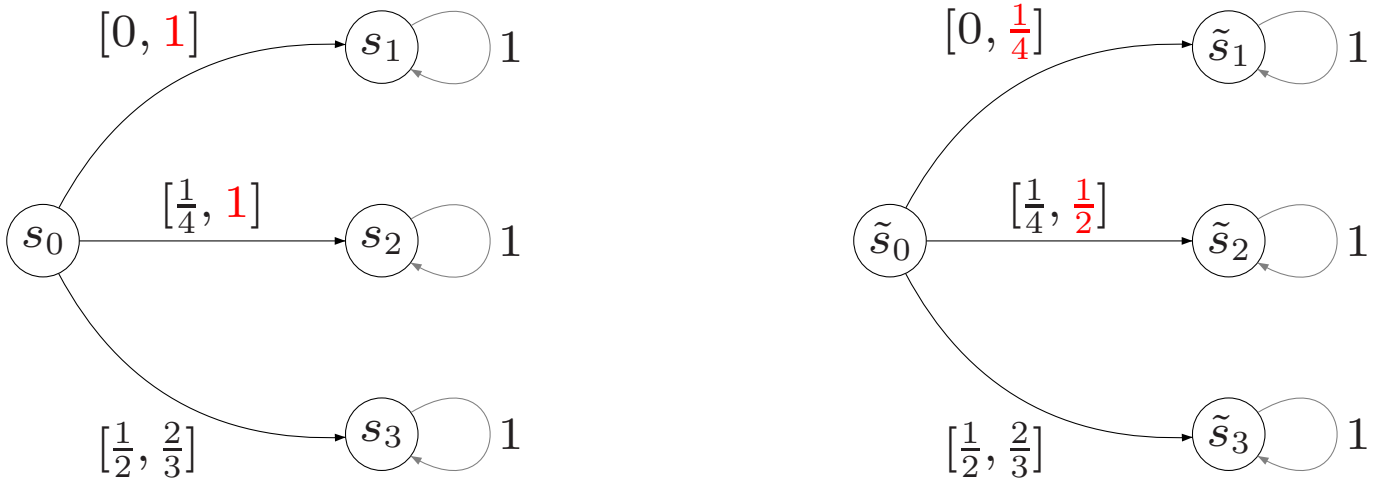

an AMC is normalized if for each pair (s, s') and $p \in [\mathbf{P}^l(s, s'), \mathbf{P}^u(s, s')]$
there exists a distribution μ with $\mu(s') = p$

Correctness

For AMC $\mathcal{D}$ with state space S, and partitioning $\mathcal{A}$ of S:

$$\mathcal{D} \sqsubseteq \alpha(\mathcal{A}, \mathcal{D})$$

For states s and s' of AMC $\mathcal{D}$ with $s \sqsubseteq s'$:

$\forall \Phi \in$ PCTL : $[\![\Phi]\!](s') \neq ?$ implies $[\![\Phi]\!](s) = [\![\Phi]\!](s')$

Policies

- A policy resolves the nondeterminism as given by the intervals
  - consider history-dependent deterministic policies
  - there are infinitely many of such policies
  - on an AMC, such policies induce an (infinite-state) Markov chain
- Extreme policies only select bounds of intervals
  - there are finitely (possibly exponentially) many extreme policies

For any measurable event E (in the σ-algebra on infinite paths):

$$\inf_{\text{extreme } \mathfrak{S}} \overset{\mathfrak{S}}{\Pr}(E) = \inf_{\text{any } \mathfrak{S}} \overset{\mathfrak{S}}{\Pr}(E) \quad \text{and} \quad \sup_{\text{extreme } \mathfrak{S}} \overset{\mathfrak{S}}{\Pr}(E) = \sup_{\text{any } \mathfrak{S}} \overset{\mathfrak{S}}{\Pr}(E)$$

Reachability probabilities

For $\mathcal{D} \subseteq \mathcal{D}'$ and compatible sets $G \subseteq S$, $G' \subseteq S'$
there exists for any policy $\mathfrak{S}$ on $\mathcal{D}$ a policy $\mathfrak{S}'$ on $\mathcal{D}'$ such that:

$$\Pr^{\mathfrak{S}}(\Diamond^{\leqslant k}\, G) \;=\; \Pr^{\mathfrak{S}'}(\Diamond^{\leqslant k}\, G') \quad \text{for any } k \in \mathbb{N}$$

$$\Pr^{\mathfrak{S}}(\Diamond G) \;=\; \Pr^{\mathfrak{S}'}(\Diamond G')$$

computing (step-)bounded probabilities is as in MDPs

谢谢大家！