Probabilistic Model Checking (1)

Lecture #1 of GLOBAN Summerschool

Joost-Pieter Katoen

Software Modeling and Verification Group

affiliated to University of Twente, Formal Methods and Tools

Warsaw University, September 24, 2008

Model checking

- Automated model-based verification and debugging technique
  - model of system = Kripke structure $\approx$ labeled transition system
  - properties expressed in temporal logic like LTL or CTL
  - provides counterexamples in case of property refutation
- Various striking examples
  - Needham-Schroeder protocol, cache coherence, storm surge barrier, C code
- 2008: Pioneers awarded prestigious ACM Turing Award

  

- Today: model checking of probabilistic models

Principles of Model Checking

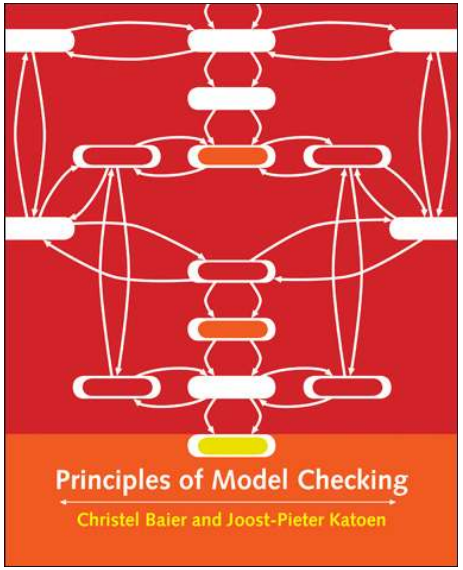

CHRISTEL BAIER

TU Dresden, Germany

JOOST-PIETER KATOEN

RWTH Aachen University, Germany,
and
University of Twente, the Netherlands

"This book offers one of the most comprehensive introductions to logic model checking techniques available today. The authors have found a way to explain both basic concepts and foundational theory thoroughly and in crystal clear prose. Highly recommended for anyone who wants to learn about this important new field, or brush up on their knowledge of the current state of the art."
(**Gerard J. Holzmann, NASA JPL, Pasadena**)

Content of this lecture

- Introduction
  - why probabilities?, history, tools + applications
- Markov chains
  - paths, measurability, reachability probabilities
- Probabilistic CTL
  - syntax, semantics, model checking, PCTL versus CTL
- Abstraction
  - bisimulation, correctness, minimization

Content of this lecture

⇒ Introduction

- why probabilities?, history, tools + applications

- Markov chains
  - paths, measurability, reachability probabilities
- Probabilistic CTL
  - syntax, semantics, model checking, PCTL versus CTL
- Abstraction
  - bisimulation, correctness, minimization

Probabilities help

- When analysing system performance and dependability
  - to quantify arrivals, waiting times, time between failure, QoS, ...
- When modelling uncertainty in the environment
  - to quantify imprecisions in system inputs
  - to quantify unpredictable delays, express soft deadlines, ...
- When building protocols for networked embedded systems
  - randomized algorithms
- When certain problems are undecidable deterministically
  - reachability in communicating finite-state machines

Probabilistic models

	Nondeterminism no	Nondeterminism yes
Discrete time	discrete-time Markov chain (DTMC)	Markov decision process (MDP)
Continuous time	CTMC	CTMDP

Breakthroughs

- Zero-one probabilities for Markov decision processes (Vardi 1985)
  - does an LTL formula hold with probability zero?
- Markov decision processes (Courcoubetis & Yannakakis 1988)
  - does the maximal probability for an LTL formula equal p?
- Discrete-time Markov chains (Hansson & Jonsson 1990)
  - does the probability of a CTL formula equal p?
- Markov decision processes (Bianco & de Alfaro 1995)
  - does the maximal probability for a CTL formula equal p?
- Continuous-time Markov chains (Baier, Katoen & Hermanns 1999)
  - does the probability of a timed CTL formula equal p?

Reachability probabilities

	Nondeterminism no	Nondeterminism yes
Reachability	linear equation system DTMC	linear programming MDP
Timed reachability	transient analysis (+ uniformization) CTMC	greedy backward reachability uniform CTMDP

What is probabilistic model checking?

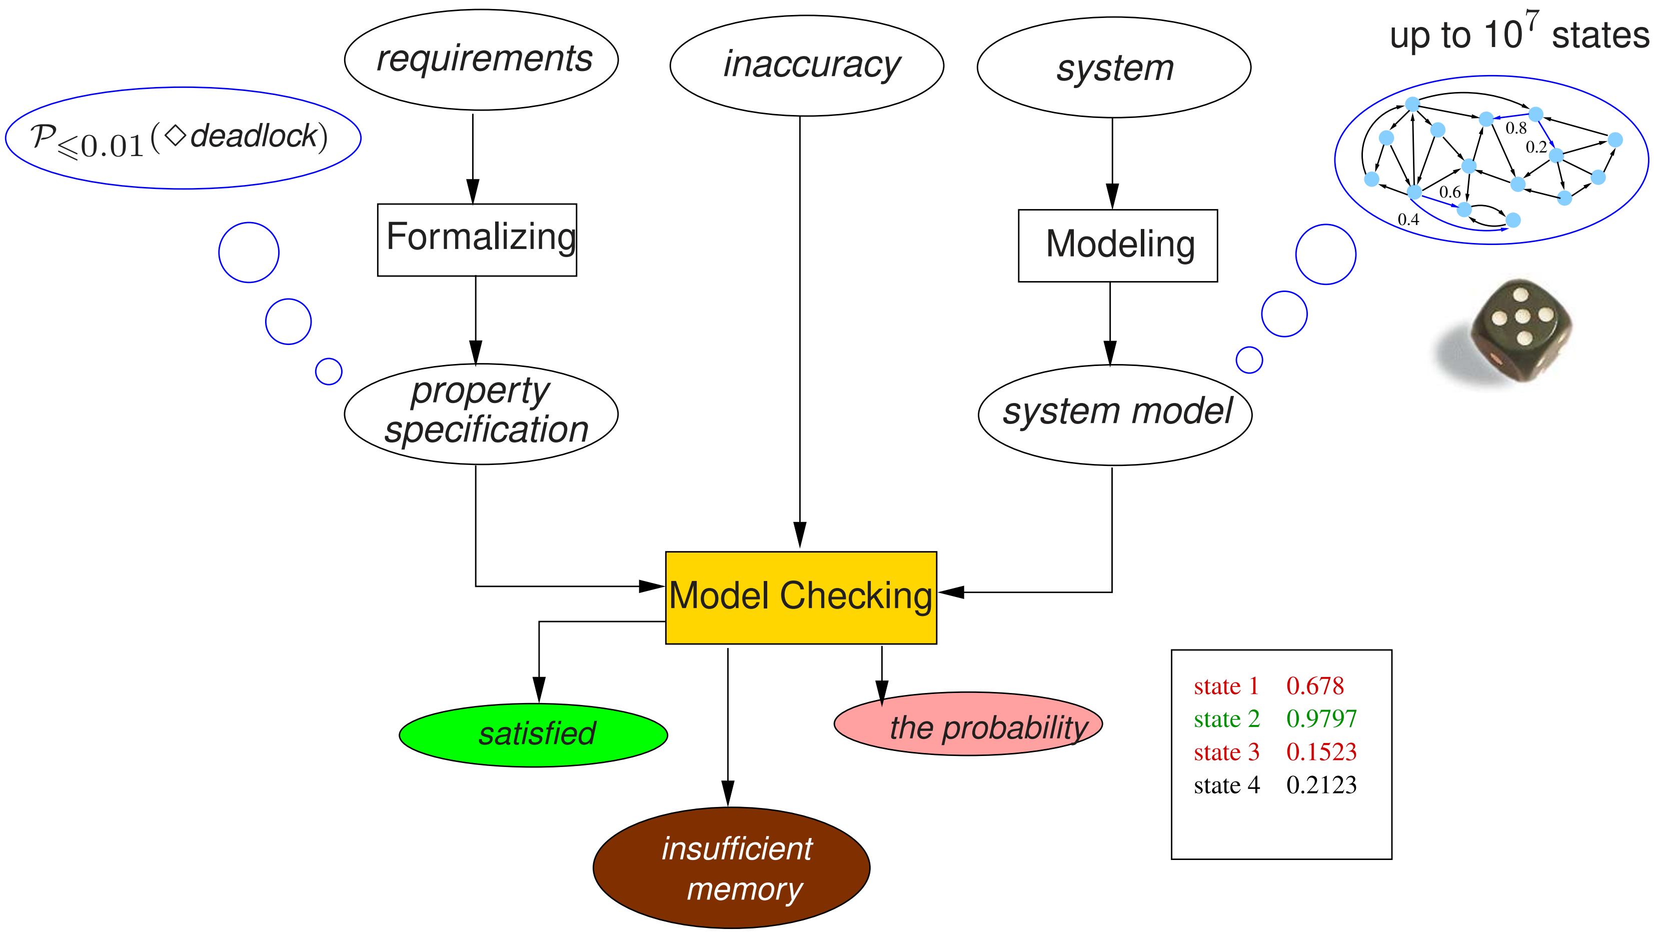

 Characteristics

- What is inside?
  - temporal logics and model checking
  - numerical and optimisation techniques from performance and OR
- What can be checked?
  - time-bounded reachability, long-run averages, safety and liveness
- What is its usage?
  - powerful tools: PRISM (4,000 downloads), MRMC, Petri net tools, Probmela
  - applications: distributed systems, biology, avionics, . . .

Probability elsewhere

- In performance modelling (Erlang, 1907)

  - models: typically continuous-time Markov chains
  - emphasis on steady-state and transient measures

- In stochastic control theory and operations research (Bellman, 1957)

  - models: typically discrete-time Markov decision models
  - emphasis on finding optimal policies for average measures

- Our focus: model checking Markov chains
  - temporal logic $\Rightarrow$ unambiguous and precise *measure-specification*
  - model-checking techniques $\Rightarrow$ no expert algorithmic knowledge needed
  - complex (new) measures are concisely specified and *automatically verified*
  - exchanging techniques with the other two areas

Illustrating examples

- Security: Crowds protocol
  - analysis of probability of anonymity
- IEEE 1394 Firewire protocol
  - proof that biased delay is optimal
- Systems biology
  - probability that enzymes are absent within the deadline
- Software in next generation of satellites
  - mission time probability (ESA project)

A synchronous leader election protocol

[Itai & Rodeh, 1990]

- A round-based protocol in a synchronous ring of $N > 2$ nodes
  - the nodes proceed in a lock-step fashion
  - each slot = 1 message is read + 1 state change + 1 message is sent

  $\Rightarrow$ this synchronous computation yields a Markov chain

- Each round starts by each node choosing a uniform id $\in \{ 1, \ldots, K \}$
- Nodes pass their selected id around the ring
- If there is a unique id, the node with the maximum unique id is leader
- If not, start another round and try again . . .

Leader election

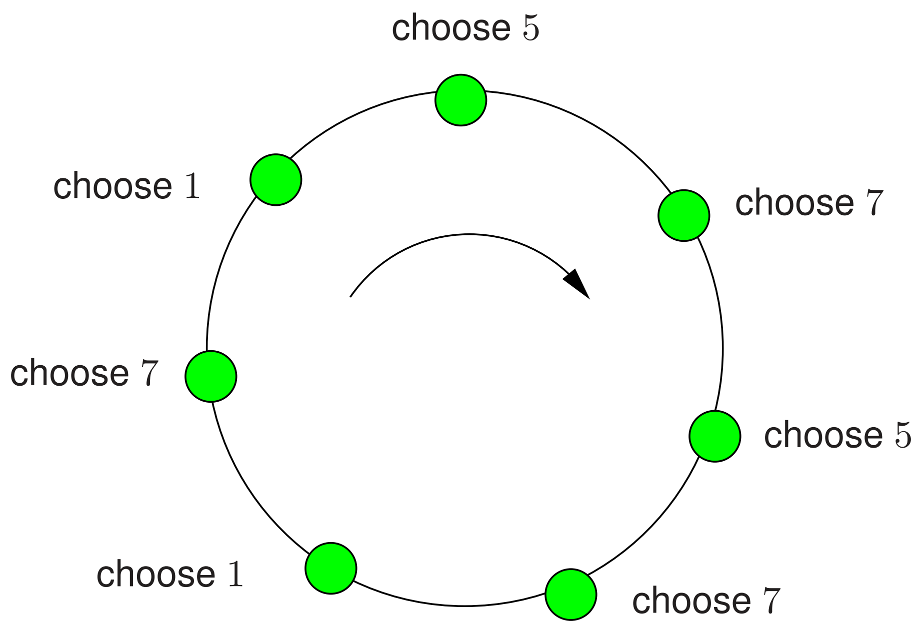

probabilistically choose an id from $[1 \ldots K]$

Leader election

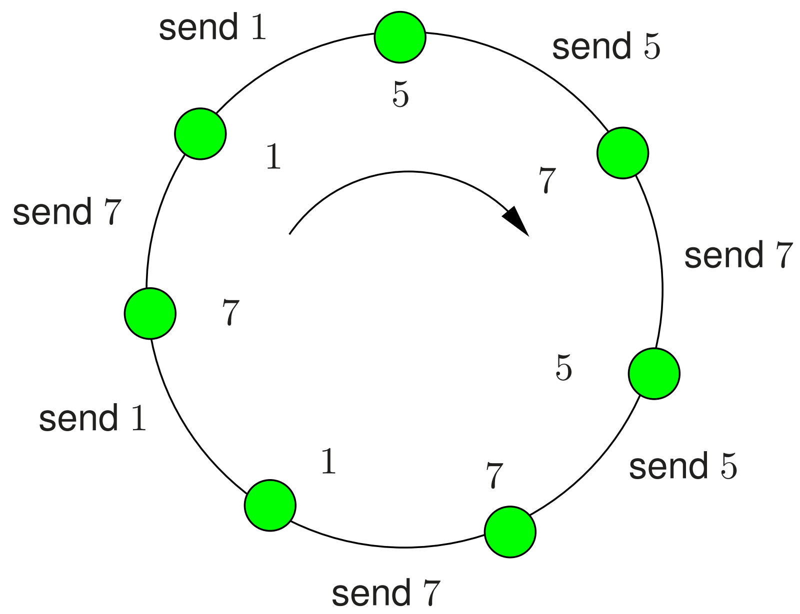

send your selected id to your neighbour

Leader election

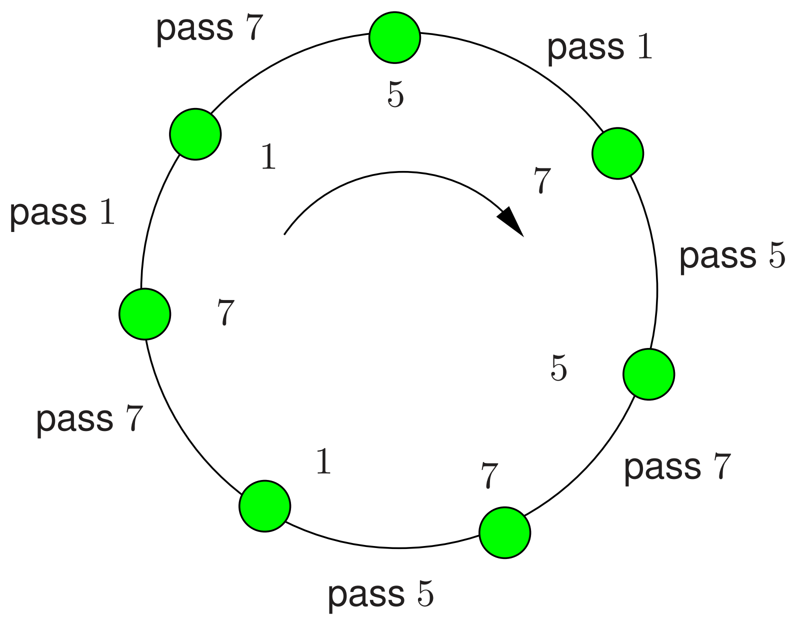

pass the received id, and check uniqueness own id

Leader election

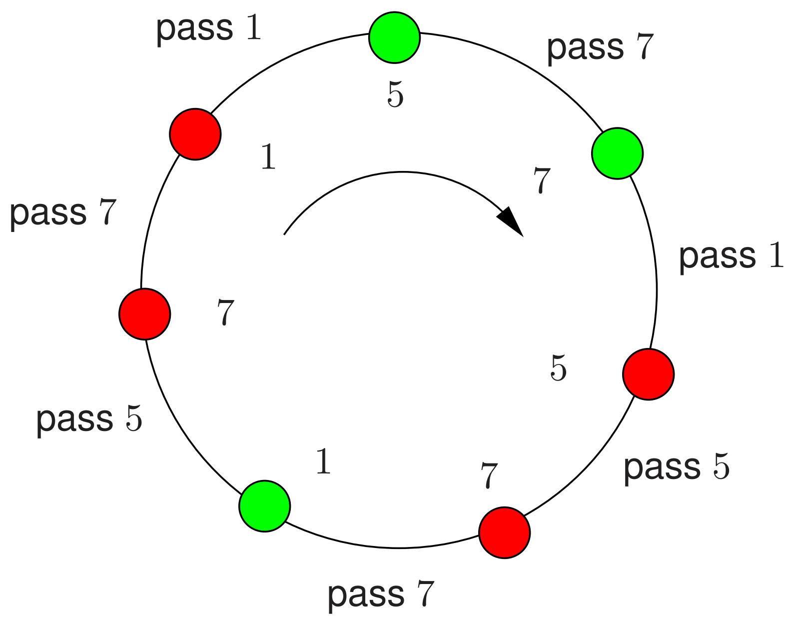

pass the received id, and check uniqueness own id

Leader election

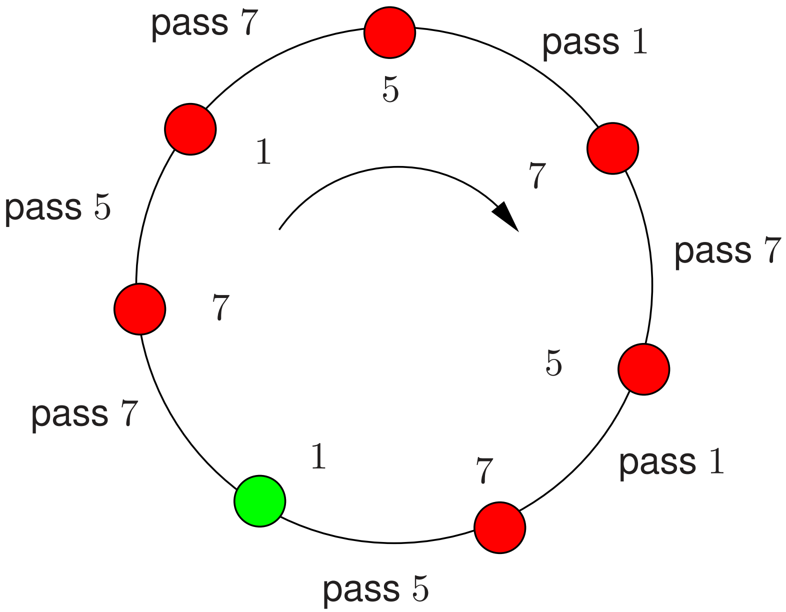

pass the received id, and check uniqueness own id

End of 1st round

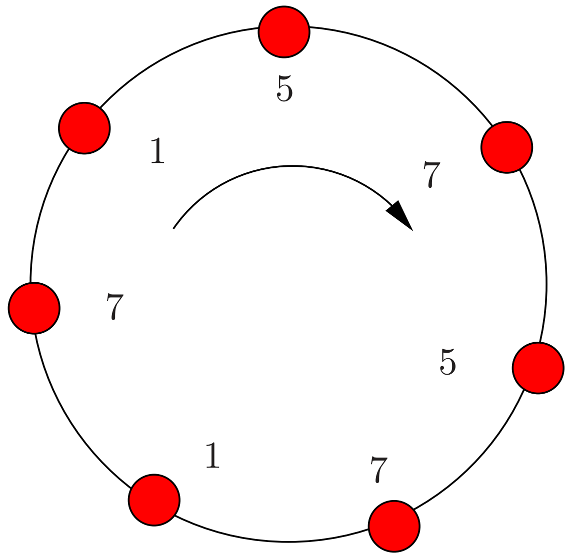

no unique leader has been elected

Start a new round

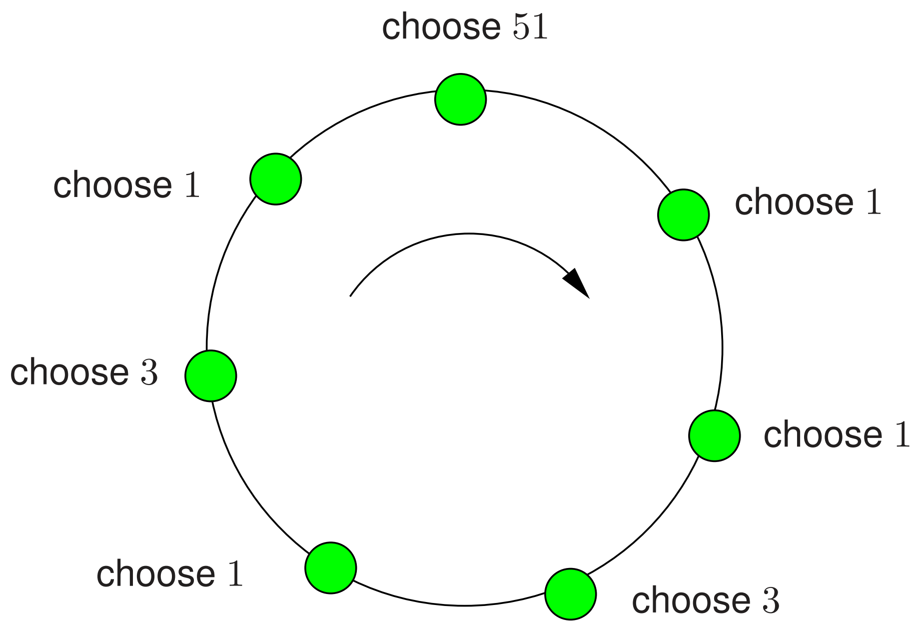

new round and new chances!

Properties of leader election

- Almost surely eventually a leader will be elected:

$$\mathbb{P}_{=1}(\Diamond \textit{ leader elected})$$

- With probability $\geqslant \frac{4}{5}$, eventually a leader is elected :

$$\mathbb{P}_{\geqslant 0.8}(\Diamond \textit{ leader elected})$$

- . . . . . . within k steps:

$$\mathbb{P}_{\geqslant 0.8}(\Diamond^{\leqslant k} \textit{ leader elected})$$

Probability to elect a leader within L rounds

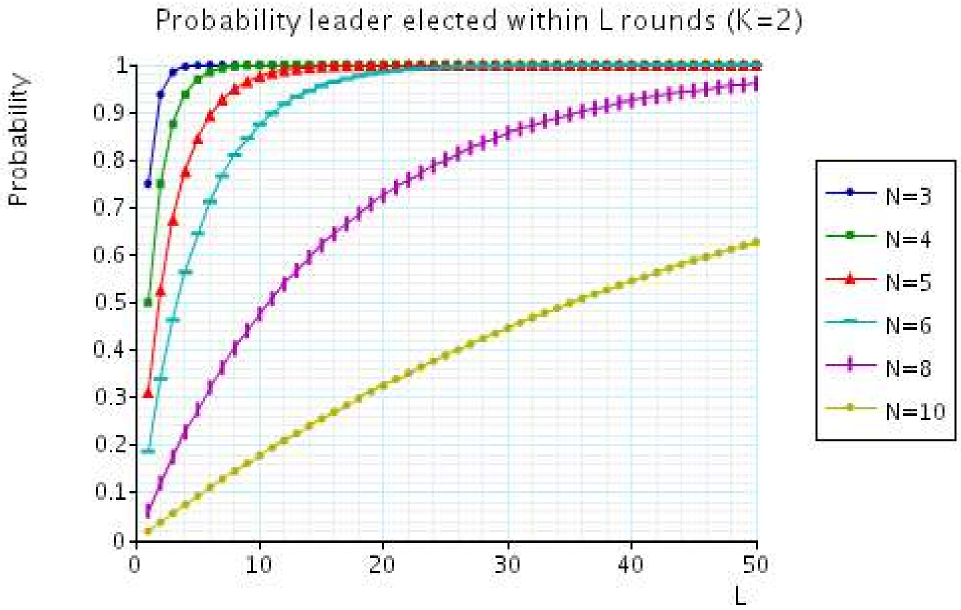

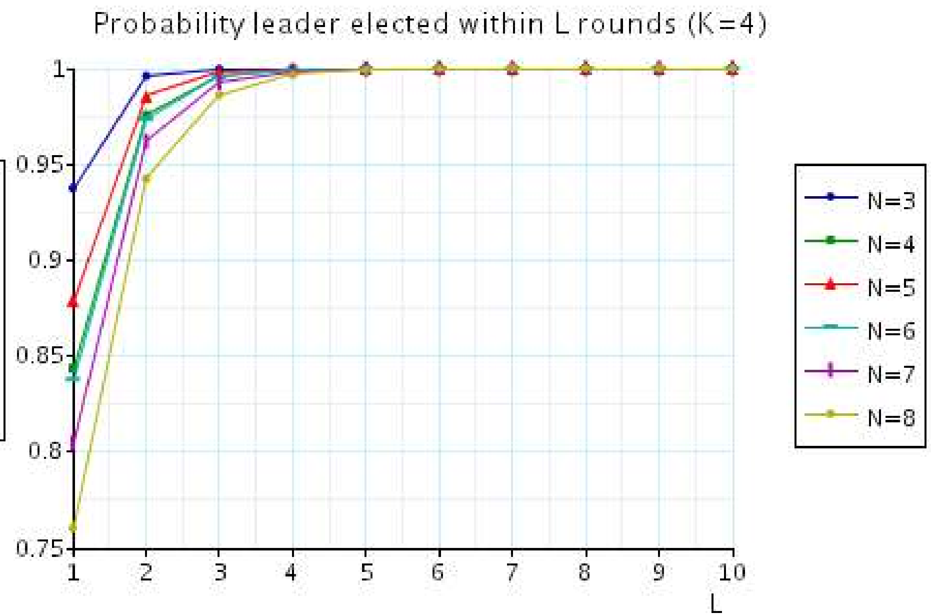

$\mathbb{P}_{\leqslant q}(\Diamond^{\leqslant (N+1)\cdot L}$ *leader elected*$)$ (Itai & Rodeh's algorithm)

Content of this lecture

- Introduction
  - why probabilities?, history, tools + applications

$\Rightarrow$ Markov chains
  - paths, measurability, reachability probabilities

- Probabilistic CTL
  - syntax, semantics, model checking, PCTL versus CTL

- Abstraction
  - bisimulation, correctness, minimization

Discrete-time Markov chains

A DTMC $\mathcal{M}$ is a tuple $(S, \mathbf{P}, \iota_{init}, AP, L)$ with:

- S is a countable nonempty set of states
- $\mathbf{P} : S \times S \rightarrow [0, 1]$, transition probability function s.t. $\sum_{s'} \mathbf{P}(s, s') = 1$
  - $\mathbf{P}(s, s')$ is the probability to jump from s to s' in one step
- $\iota_{init} : S \rightarrow [0, 1]$, the initial distribution with $\sum\limits_{s \in S} \iota_{init}(s) = 1$
  - $\iota_{init}(s)$ is the probability that system starts in state s
  - state s for which $\iota_{init}(s) > 0$ is an initial state
- $L : S \rightarrow 2^{AP}$, the labelling function

$\Rightarrow$ <u>a DTMC is a transition system with probabilistic transitions</u>

Craps

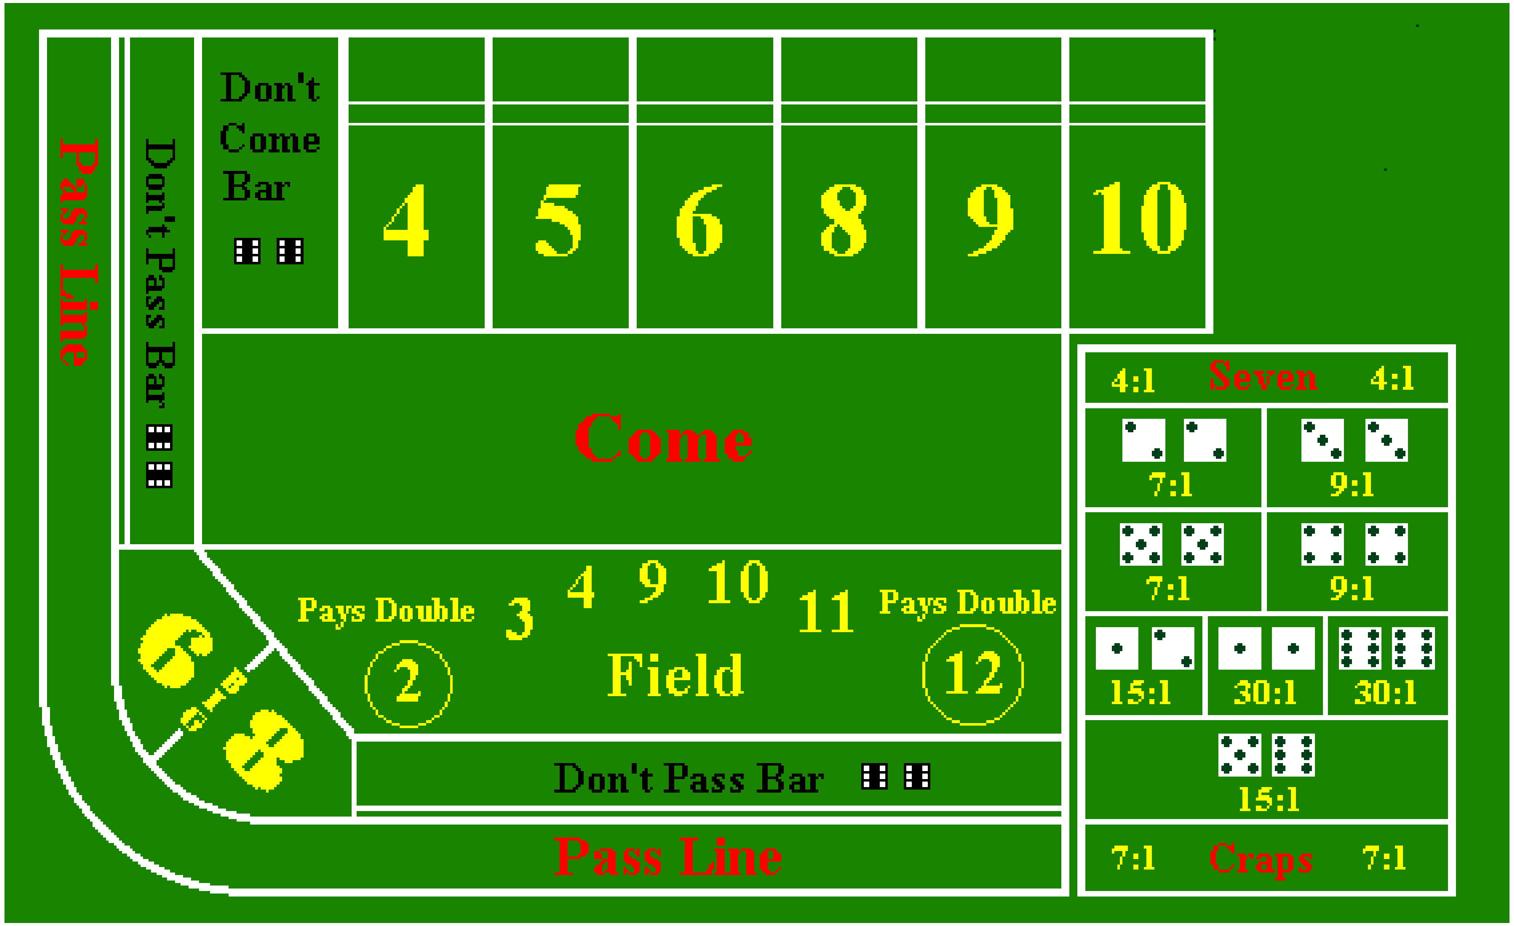

Pass Line
Don't Pass Bar
Don't Come Bar
4
5
6
8
9
10
Come
4:1 Seven 4:1
7:1
9:1
7:1
9:1
15:1
30:1
30:1
15:1
7:1 Craps 7:1
Pays Double
3
4
9
10
11
Pays Double
2
Field
12
6
8
Don't Pass Bar
Pass Line

Craps

- Roll two dice and bet on outcome

- Come-out roll ("pass line" wager):
  - outcome 7 or 11: win
  - outcome 2, 3, or 12: loss ("craps")
  - any other outcome: roll again (outcome is "point")

- Repeat until 7 or the "point" is thrown:
  - outcome 7: loss ("seven-out")
  - outcome the point: win
  - any other outcome: roll again

A DTMC model of Craps

- Come-out roll:
  - 7 or 11: win
  - 2, 3, or 12: loss
  - else: roll again
- Next roll(s):
  - 7: loss
  - point: win
  - else: roll again

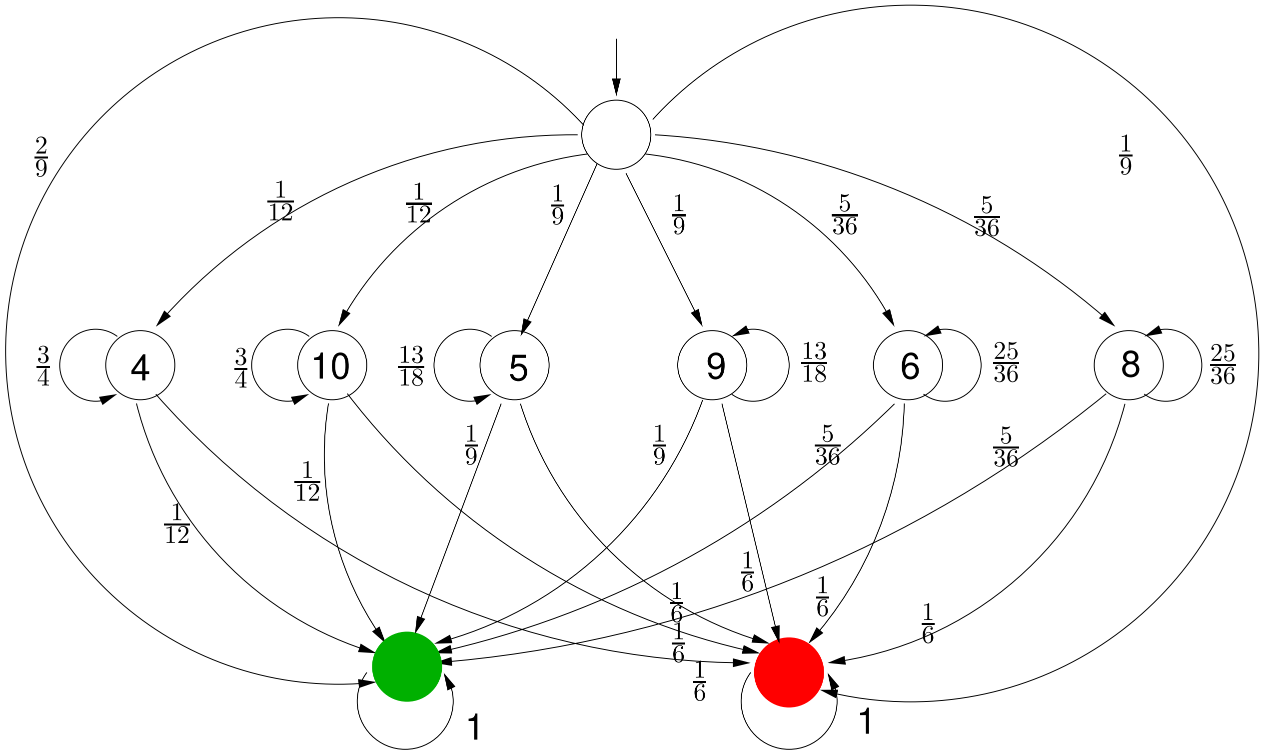

Paths

- State graph of DTMC $\mathcal{M}$
  - vertices are states of $\mathcal{M}$, and (s, s') is an edge iff $\mathbf{P}(s, s') > 0$
- Paths in $\mathcal{M}$ are maximal (i.e., infinite) paths in its state graph
  - *Paths*$(\mathcal{M})$ and *Paths*$_{\mathit{fin}}(\mathcal{M})$ denote the set of (finite) paths in $\mathcal{M}$
- *Post*$(s) = \{s' \in S \mid \mathbf{P}(s, s') > 0\}$ and *Pre*$(s) = \{s' \in S \mid \mathbf{P}(s', s) > 0\}$
  - *Post*$^*(s)$ is the set of states reachable from s via a finite path fragment
  - *Pre*$^*(s) = \{ s' \in S \mid s \in$ *Post*$^*(s') \}$

Probability measure on DTMCs

- Events are *infinite paths* in the DTMC $\mathcal{M}$, i.e., $\Omega = \mathit{Paths}(\mathcal{M})$
- σ-algebra on $\mathcal{M}$ is generated by *cylinder sets* of finite paths $\hat{\pi}$:

$$\mathit{Cyl}(\hat{\pi}) \;=\; \big\{ \pi \in \mathit{Paths}(\mathcal{M}) \;\big|\; \hat{\pi} \text{ is a prefix of } \pi \big\}$$

  – cylinder sets serve as basis events of the smallest σ-algebra on $\mathit{Paths}(\mathcal{M})$

- $\Pr$ is the *probability measure* on the σ-algebra on $\mathit{Paths}(\mathcal{M})$:

$$\Pr\big(\mathit{Cyl}(s_0 \ldots s_n)\big) \;=\; \iota_{init}(s_0) \cdot \mathbf{P}(s_0 \ldots s_n)$$

  – where $\mathbf{P}(s_0\, s_1 \ldots s_n) \;=\; \prod_{0 \leqslant i < n} \mathbf{P}(s_i, s_{i+1})$ and $\mathbf{P}(s_0) = 1$

Reachability probabilities

- What is the probability to reach a set of states $B \subseteq S$ in DTMC $\mathcal{M}$?
- Which event does $\Diamond B$ mean formally?
  - the union of all cylinders $\mathit{Cyl}(s_0 \ldots s_n)$ where
  - $s_0 \ldots s_n$ is an initial path fragment in $\mathcal{M}$ with $s_0, \ldots, s_{n-1} \notin B$ and $s_n \in B$

$$
\begin{aligned}
\Pr(\Diamond B) &= \sum_{s_0 \ldots s_n \in \mathit{Paths}_{\mathit{fin}}(\mathcal{M}) \cap (S \setminus B)^* B} \Pr\big(\mathit{Cyl}(s_0 \ldots s_n)\big) \\
&= \sum_{s_0 \ldots s_n \in \mathit{Paths}_{\mathit{fin}}(\mathcal{M}) \cap (S \setminus B)^* B} \iota_{\mathit{init}}(s_0) \cdot \mathbf{P}(s_0 \ldots s_n)
\end{aligned}
$$

Reachability probabilities in finite DTMCs

- Let $\Pr(s \models \Diamond B) = \Pr_s(\Diamond B) = \Pr_s\{\pi \in \mathit{Paths}(s) \mid \pi \models \Diamond B\}$
  - where $\Pr_s$ is the probability measure in $\mathcal{M}$ with single initial state s
- Let variable $x_s = \Pr(s \models \Diamond B)$ for any state s
  - if B is not reachable from s then $x_s = 0$
  - if $s \in B$ then $x_s = 1$
- For any state $s \in \mathit{Pre}^*(B) \setminus B$:

$$x_s = \underbrace{\sum_{t \in S \setminus B} \mathbf{P}(s,t) \cdot x_t}_{\text{reach } B \text{ via } t} + \underbrace{\sum_{u \in B} \mathbf{P}(s,u)}_{\text{reach } B \text{ in one step}}$$

Remark: expansion law

- Recall in CTL: $\exists(C \,\mathsf{U}\, B)$ is the least solution of expansion law:

  $$\exists(C \,\mathsf{U}\, B) \;\equiv\; B \;\vee\; (C \wedge \exists \bigcirc \exists(C \,\mathsf{U}\, B))$$

- That is: the set $X = \mathit{Sat}(\exists(C \,\mathsf{U}\, B))$ is the smallest set such that:

  $$B \cup \{\, s \in C \setminus B \mid \mathit{Post}(s) \cap X \neq \varnothing \,\} \;\subseteq\; X$$

- Previous slide "replaces" $s \in X$ by values x_s in $[0, 1]$
  - if $s \in B$ then $x_s = 1$ (compare: $s \in B$ implies $s \in X$)
  - if $s \in S \setminus (C \cup B)$ then $x_s = 0$ (compare: $s \notin C \cup B$ implies $s \notin X$)

- If $s \in C \setminus B$ then $x_s = \sum_{t \in C \setminus B} \mathbf{P}(s,t) \cdot x_t + \sum_{t \in B} \mathbf{P}(s,t)$
  - compare: $s \in C \setminus B$ and $\mathit{Post}(s) \cap X \neq \varnothing$ implies $s \in X$

Linear equation system

- These equations can be rewritten into the following form:

$$\mathbf{x} = \mathbf{A}\mathbf{x} + \mathbf{b}$$

  - where vector $\mathbf{x} = (x_s)_{s \in \tilde{S}}$ with $\tilde{S} = \mathit{Pre}^*(B) \setminus B$
  - $\mathbf{A} = \Big(\mathbf{P}(s,t) \Big)_{s,t \in \tilde{S}}$, the transition probabilities in $\tilde{S}$
  - $\mathbf{b} = \Big(b_s \Big)_{s \in \tilde{S}}$ contains the probabilities to reach B within one step

- *Linear equation system*: $(\mathbf{I} - \mathbf{A})\mathbf{x} = \mathbf{b}$

  - note: more than one solution may exist if $\mathbf{I} - \mathbf{A}$ has no inverse (i.e., is singular)

$\Rightarrow$ characterize the desired probability as least fixed point

Unique solution

Let $\mathcal{M}$ be a finite DTMC with state space S partitioned into:

- $S_{=0} \;=\; \mathit{Sat}(\neg\exists(C \,\mathsf{U}\, B))$
- $S_{=1}$ a subset of $\{s \in S \mid \Pr(s \models C \,\mathsf{U}\, B) = 1\}$ that contains B
- $S_{?} \;=\; S \setminus (S_{=0} \cup S_{=1})$

The vector $\quad \big(\Pr(s \models C \,\mathsf{U}\, B)\big)_{s \in S_?}$

is the *unique* solution of the linear equation system:

$$\mathbf{x} \;=\; \mathbf{A}\mathbf{x} + \mathbf{b} \quad \text{where} \quad \mathbf{A} \;=\; \big(\mathbf{P}(s,t)\big)_{s,t \in S_?} \;\text{ and }\; \mathbf{b} \;=\; \big(\mathbf{P}(s, S_{=1})\big)_{s \in S_?}$$

Computing reachability probabilities

- The probabilities of the events $C \,\mathsf{U}^{\leqslant n}\, B$ can be obtained iteratively:

$$\mathbf{x}^{(0)} = \mathbf{0} \quad \text{and} \quad \mathbf{x}^{(i+1)} = \mathbf{A}\mathbf{x}^{(i)} + \mathbf{b} \text{ for } 0 \leqslant i < n$$

- where $\mathbf{A} = \big(\mathbf{P}(s,t)\big)_{s,t \in C \setminus B}$ and $\mathbf{b} = \big(\mathbf{P}(s,B)\big)_{s \in C \setminus B}$

- Then: $\mathbf{x}^{(n)}(s) = \Pr(s \models C \,\mathsf{U}^{\leqslant n} B)$ for $s \in C \setminus B$

Example: Craps game

- $\Pr(start \models C \, \mathsf{U}^{\leqslant n} \, B)$
- $S_{=0} = \{ 8, 9, 10, lost \}$
- $S_{=1} = \{ won \}$
- $S_{?} = \{ start, 4, 5, 6 \}$

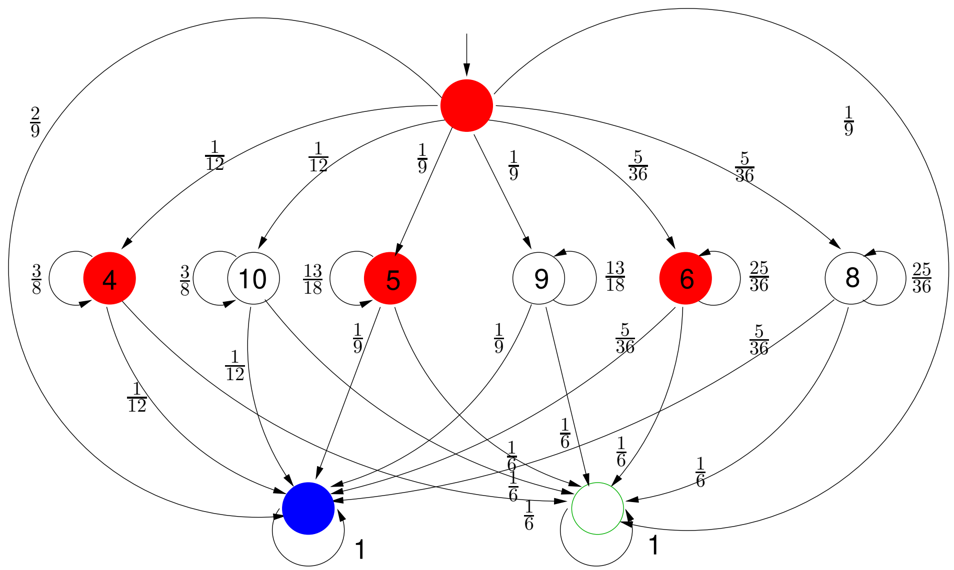

Example: Craps game

- *start* $< 4 < 5 < 6$

- $\mathbf{A} = \frac{1}{36}\begin{pmatrix} 0 & 3 & 4 & 5 \\ 0 & 27 & 0 & 0 \\ 0 & 0 & 26 & 0 \\ 0 & 0 & 0 & 25 \end{pmatrix}$

- $\mathbf{b} = \frac{1}{36}\begin{pmatrix} 8 \\ 3 \\ 4 \\ 5 \end{pmatrix}$

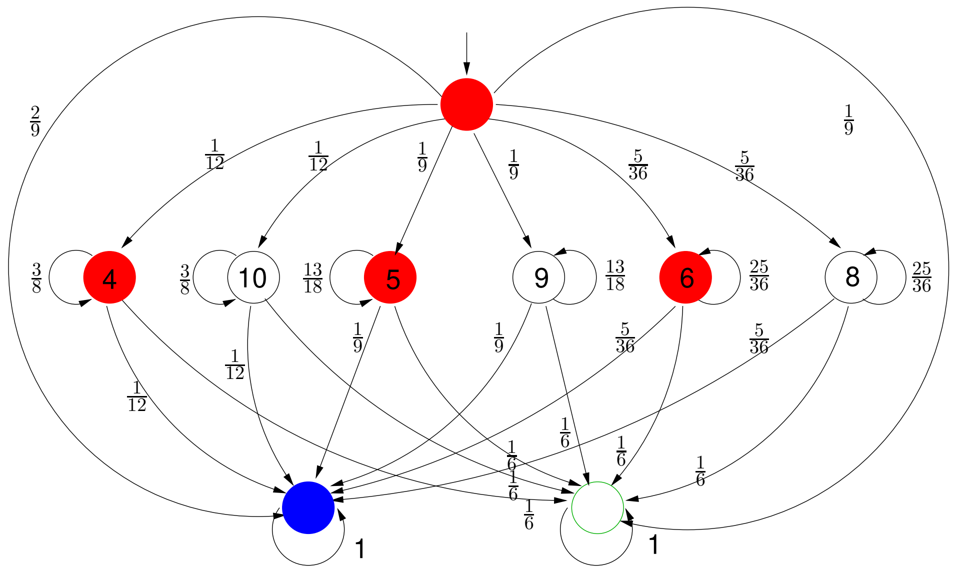

$$\mathbf{x}^{(0)} = \mathbf{0} \quad \text{and} \quad \mathbf{x}^{(i+1)} = \mathbf{A}\mathbf{x}^{(i)} + \mathbf{b} \text{ for } 0 \leqslant i < n.$$

Example: Craps game

$$\mathbf{x}^{(2)} = \underbrace{\frac{1}{36}\begin{pmatrix} 0 & 3 & 4 & 5 \\ 0 & 27 & 0 & 0 \\ 0 & 0 & 26 & 0 \\ 0 & 0 & 0 & 25 \end{pmatrix}}_{\mathbf{A}} \cdot \underbrace{\frac{1}{36}\begin{pmatrix} 8 \\ 3 \\ 4 \\ 5 \end{pmatrix}}_{\mathbf{x}^{(1)}} + \underbrace{\frac{1}{36}\begin{pmatrix} 8 \\ 3 \\ 4 \\ 5 \end{pmatrix}}_{\mathbf{b}} = \left(\frac{1}{36}\right)^2 \begin{pmatrix} 338 \\ 189 \\ 248 \\ 305 \end{pmatrix}$$

Content of this lecture

- Introduction
  - why probabilities?, history, tools + applications
- Markov chains
  - paths, measurability, reachability probabilities

$\Rightarrow$ Probabilistic CTL
  - syntax, semantics, model checking, PCTL versus CTL
- Abstraction
  - bisimulation, correctness, minimization

PCTL Syntax

- For $a \in AP$, $J \subseteq [0,1]$ an interval with rational bounds, and natural n:

$$\boxed{\begin{array}{rcl} \Phi & ::= & \text{true} \;\big|\; a \;\big|\; \Phi \wedge \Phi \;\big|\; \neg\Phi \;\big|\; \mathbb{P}_J(\varphi) \\[2ex] \varphi & ::= & \bigcirc\Phi \;\big|\; \Phi_1 \,\mathsf{U}\, \Phi_2 \;\big|\; \Phi_1 \,\mathsf{U}^{\leqslant n}\, \Phi_2 \end{array}}$$

- $s_0 s_1 s_2 \ldots \models \Phi \,\mathsf{U}^{\leqslant n}\, \Psi$ if Φ holds until Ψ holds within n steps
- $s \models \mathbb{P}_J(\varphi)$ if probability that paths starting in s fulfill φ lies in J

abbreviate $\mathbb{P}_{[0,0.5]}(\varphi)$ by $\mathbb{P}_{\leqslant 0.5}(\varphi)$ and $\mathbb{P}_{]0,1]}(\varphi)$ by $\mathbb{P}_{>0}(\varphi)$ and so on

Derived operators

$$\Diamond\Phi \;=\; \mathsf{true}\,\mathsf{U}\,\Phi$$

$$\Diamond^{\leqslant n}\Phi \;=\; \mathsf{true}\,\mathsf{U}^{\leqslant n}\,\Phi$$

$$\mathbb{P}_{\leqslant p}(\Box\Phi) \;=\; \mathbb{P}_{\geqslant 1-p}(\Diamond\neg\Phi)$$

$$\mathbb{P}_{]p,q]}(\Box^{\leqslant n}\,\Phi) \;=\; \mathbb{P}_{[1-q,1-p[}(\Diamond^{\leqslant n}\,\neg\Phi)$$

operators like weak until W or release R can be derived analogously

Example properties

- With probability $\geqslant 0.92$, a goal state is reached via legal ones:

$$\mathbb{P}_{\geqslant 0.92}\,(\neg\, \mathit{illegal}\ \mathsf{U}\ \ \mathit{goal})$$

- ... in maximally 137 steps: $\mathbb{P}_{\geqslant 0.92}\,(\neg\, \mathit{illegal}\ \mathsf{U}^{\leqslant 137}\ \mathit{goal})$

- ... once there, remain there almost surely for the next 31 steps:

$$\mathbb{P}_{\geqslant 0.92}\left(\neg\, \mathit{illegal}\ \mathsf{U}^{\leqslant 137}\ \mathbb{P}_{=1}(\Box^{[0,31]}\ \mathit{goal})\right)$$

PCTL semantics (1)

$\mathcal{M}, s \models \Phi$ if and only if formula Φ holds in state s of DTMC $\mathcal{M}$

Relation $\models$ is defined by:

$$
\begin{array}{lll}
s \models a & \text{iff} & a \in L(s) \\
s \models \neg \Phi & \text{iff} & \text{not } (s \models \Phi) \\
s \models \Phi \vee \Psi & \text{iff} & (s \models \Phi) \text{ or } (s \models \Psi) \\
s \models \mathbb{P}_J(\varphi) & \text{iff} & \Pr(s \models \varphi) \in J
\end{array}
$$

where $\Pr(s \models \varphi) = \Pr_s\{\pi \in \mathit{Paths}(s) \mid \pi \models \varphi\}$

PCTL semantics (2)

A *path* in $\mathcal{M}$ is an infinite sequence $s_0\, s_1\, s_2 \ldots$ with $\mathbf{P}(s_i, s_{i+1}) > 0$

Semantics of path-formulas is defined as in CTL:

$$
\begin{array}{lll}
\pi \models \bigcirc\Phi & \text{iff} & s_1 \models \Phi \\
\pi \models \Phi \,\mathsf{U}\, \Psi & \text{iff} & \exists n \geqslant 0.(\, s_n \models \Psi \,\wedge \forall 0 \leqslant i < n.\, s_i \models \Phi\,) \\
\pi \models \Phi \,\mathsf{U}^{\leqslant n}\, \Psi & \text{iff} & \exists k \geqslant 0.(\, k \leqslant n \,\wedge\, s_k \models \Psi \wedge \\
 & & \qquad \forall 0 \leqslant i < k.\, s_i \models \Phi\,)
\end{array}
$$

Measurability

For any PCTL path formula φ and state s of DTMC $\mathcal{M}$

the set $\{ \pi \in \mathit{Paths}(s) \mid \pi \models \varphi \}$ is measurable

PCTL model checking

- Given a finite DTMC $\mathcal{M}$ and PCTL formula Φ, how to check $\mathcal{M} \models \Phi$?

- Check whether state s in a DTMC satisfies a PCTL formula:
  - compute recursively the set *Sat*(Φ) of states that satisfy Φ
  - check whether state s belongs to *Sat*(Φ)

  $\Rightarrow$ bottom-up traversal of the parse tree of Φ (like for CTL)

- For the propositional fragment: as for CTL

- How to compute *Sat*(Φ) for the probabilistic operators?

PCTL model checking

- Alternative formulation: $s \models \mathbb{P}_J(\bigcirc \Phi)$ if and only if $\Pr(s \models \bigcirc \Phi) \in J$
- Next: $\Pr(s \models \bigcirc \Phi)$ equals $\displaystyle\sum_{s' \in Sat(\Phi)} \mathbf{P}(s, s')$
- Matrix-vector multiplication:

$$\big(\Pr(s \models \bigcirc \Phi)\big)_{s \in S} = \mathbf{P} \cdot \iota_\Phi$$

where ι_Φ is the characteristic vector of *Sat*(Φ), i.e., $\iota_\Phi(s) = 1$ iff $s \in Sat(\Phi)$

Checking probabilistic reachability

- $s \models \mathbb{P}_J(\Phi \, \mathsf{U}^{\leqslant h} \, \Psi)$ if and only if $\Pr(s \models \Phi \, \mathsf{U}^{\leqslant h} \, \Psi) \in J$
- $\Pr(s \models \Phi \, \mathsf{U}^{\leqslant h} \, \Psi)$ is the least solution of: (Hansson & Jonsson, 1990)
  - 1 if $s \models \Psi$
  - for $h > 0$ and $s \models \Phi \wedge \neg \Psi$:

$$\sum_{s' \in S} \mathbf{P}(s, s') \cdot \Pr(s' \models \Phi \, \mathsf{U}^{\leqslant h-1} \, \Psi)$$

  - 0 otherwise
- Standard reachability for $\mathbb{P}_{>0}(\Phi \, \mathsf{U}^{\leqslant h} \, \Psi)$ and $\mathbb{P}_{\geqslant 1}(\Phi \, \mathsf{U}^{\leqslant h} \, \Psi)$
  - for efficiency reasons (avoiding solving system of linear equations)

Reduction to transient analysis

- Make all Ψ- and all $\neg(\Phi \vee \Psi)$-states absorbing in $\mathcal{M}$
- Check $\Diamond^{=h} \Psi$ in the obtained DTMC $\mathcal{M}'$
- This is a standard transient analysis in $\mathcal{M}'$:

$$\sum_{s' \models \Psi} \Pr_s\{\pi \in \textit{Paths}(s) \mid \sigma[h] = s'\}$$

  - compute by $(\mathbf{P}')^h \cdot \iota_\Psi$ where ι_Ψ is the characteristic vector of $\textit{Sat}(\Psi)$

$\Rightarrow$ Matrix-vector multiplication

Time complexity

For finite DTMC $\mathcal{M}$ and PCTL formula Φ, $\mathcal{M} \models \Phi$ can be solved in time

$$\mathcal{O}\big(\, poly(|\mathcal{M}|) \cdot n_{\max} \cdot |\Phi| \,\big)$$

where $n_{\max} \;=\; \max\{\, n \mid \Psi_1 \,\mathsf{U}^{\leqslant n}\, \Psi_2$ occurs in $\Phi \,\}$ with $\max \varnothing = 1$

Verification times

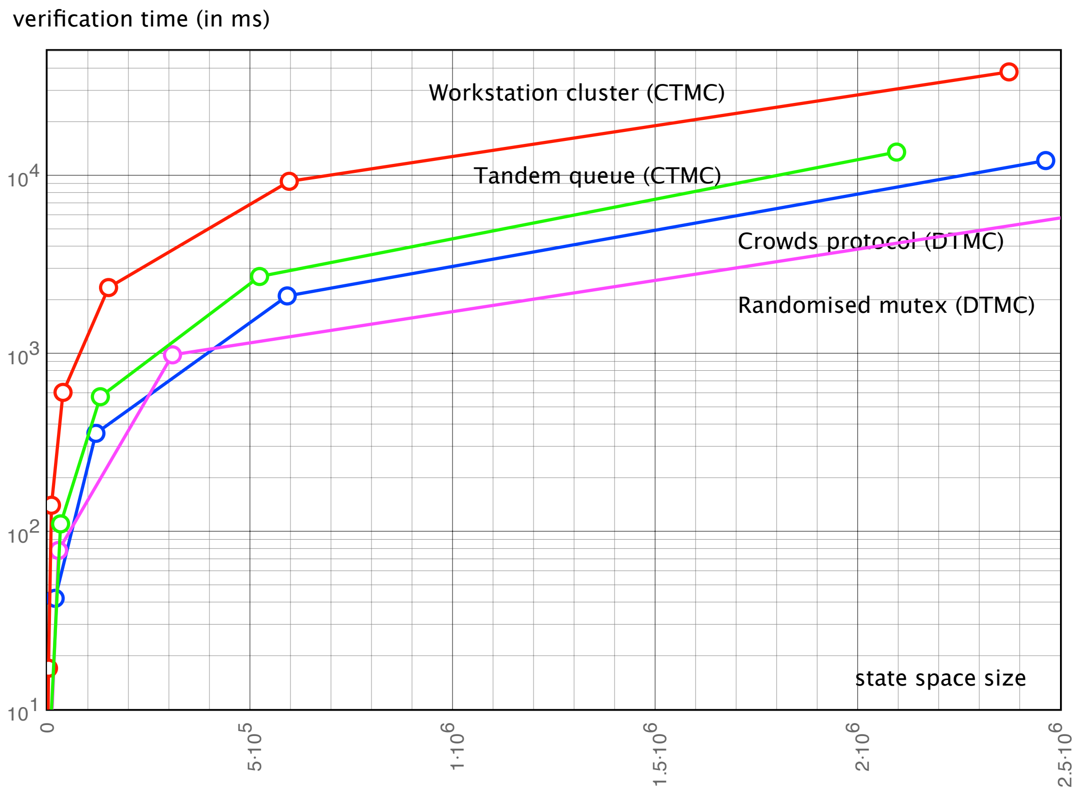

command-line tool MRMC ran on a Pentium 4, 2.66 GHz, 1 GB RAM laptop

The qualitative fragment of PCTL

- For $a \in AP$:

$$\Phi \; ::= \; \text{true} \;\big|\; a \;\big|\; \Phi \wedge \Phi \;\big|\; \neg\Phi \;\big|\; \mathbb{P}_{>0}(\varphi) \;\big|\; \mathbb{P}_{=1}(\varphi)$$

$$\varphi \; ::= \; \bigcirc\Phi \;\big|\; \Phi_1 \,\mathsf{U}\, \Phi_2$$

- The probability bounds $= 0$ and < 1 can be derived:

$$\mathbb{P}_{=0}(\varphi) \equiv \neg\mathbb{P}_{>0}(\varphi) \quad \text{and} \quad \mathbb{P}_{<1}(\varphi) \equiv \neg\mathbb{P}_{=1}(\varphi)$$

- No bounded until, and only > 0, $= 0$, > 1 and $= 1$ intervals

so: $\mathbb{P}_{=1}(\Diamond\mathbb{P}_{>0}(\bigcirc a))$ and $\mathbb{P}_{<1}(\mathbb{P}_{>0}(\Diamond a)\,\mathsf{U}\, b)$ are qualitative PCTL formulas

Qualitative PCTL = CTL?

- PCTL-formula Φ is *equivalent* to CTL-formula Ψ:
  - $\Phi \equiv \Psi$ if and only if $Sat_{\mathcal{M}}(\Phi) = Sat_{TS(\mathcal{M})}(\Psi)$ for each DTMC $\mathcal{M}$
- $\exists\varphi$ requires φ on some paths, $\mathbb{P}_{>0}(\varphi)$ with positive probability
  - $\mathbb{P}_{>0}(\bigcirc a) \equiv \exists \bigcirc a$ and $\mathbb{P}_{>0}(\Diamond a) \equiv \exists\Diamond a$ and $\mathbb{P}_{>0}(a \,\mathsf{U}\, b) \equiv \exists a \,\mathsf{U}\, b$
- $\forall\varphi$ requires φ to hold for all paths, $\mathbb{P}_{=1}(\varphi)$ for almost all
  - $\mathbb{P}_{=1}(\bigcirc a) \equiv \forall \bigcirc a$ and $\mathbb{P}_{=1}(\Box a) \equiv \forall\Box a$
- But: $\mathbb{P}_{>0}(\varphi) \equiv \exists\varphi$ and $\mathbb{P}_{=1}(\varphi) \equiv \forall\varphi$ do not hold in general!

Qualitative PCTL versus CTL

- There is no CTL-formula that is equivalent to $\mathbb{P}_{=1}(\Diamond a)$
- There is no CTL-formula that is equivalent to $\mathbb{P}_{>0}(\Box a)$
- There is no qualitative PCTL-formula that is equivalent to $\forall \Diamond a$
- There is no qualitative PCTL-formula that is equivalent to $\exists \Box a$

$\Rightarrow$ PCTL with $\forall \varphi$ and $\exists \varphi$ is more expressive than PCTL

Content of this lecture

- Introduction
  - why probabilities?, history, tools + applications
- Markov chains
  - paths, measurability, reachability probabilities
- Probabilistic CTL
  - syntax, semantics, model checking, PCTL versus CTL

$\Rightarrow$ Abstraction
  - bisimulation, correctness, minimization

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 $\mathcal{M} = (S, \mathbf{P}, \mathit{AP}, L)$ 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

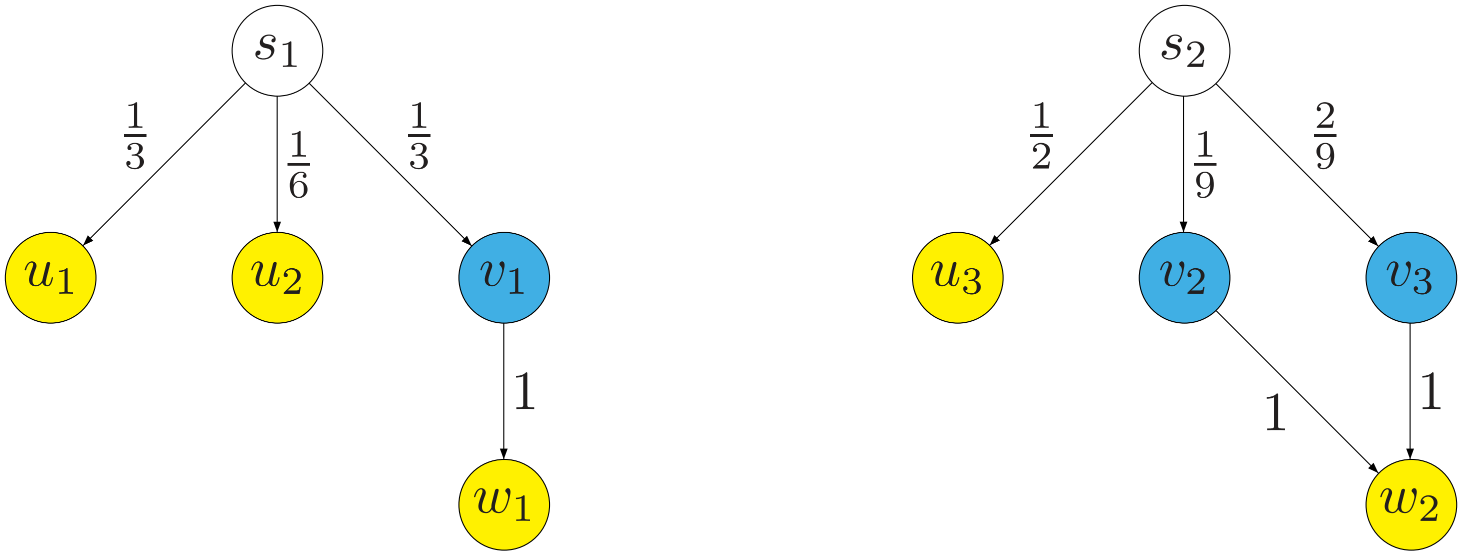

s_1
$\frac{1}{3}$
$\frac{1}{6}$
$\frac{1}{3}$
u_1
u_2
v_1
1
w_1
s_2
$\frac{1}{2}$
$\frac{1}{9}$
$\frac{2}{9}$
u_3
v_2
v_3
1
1
w_2

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

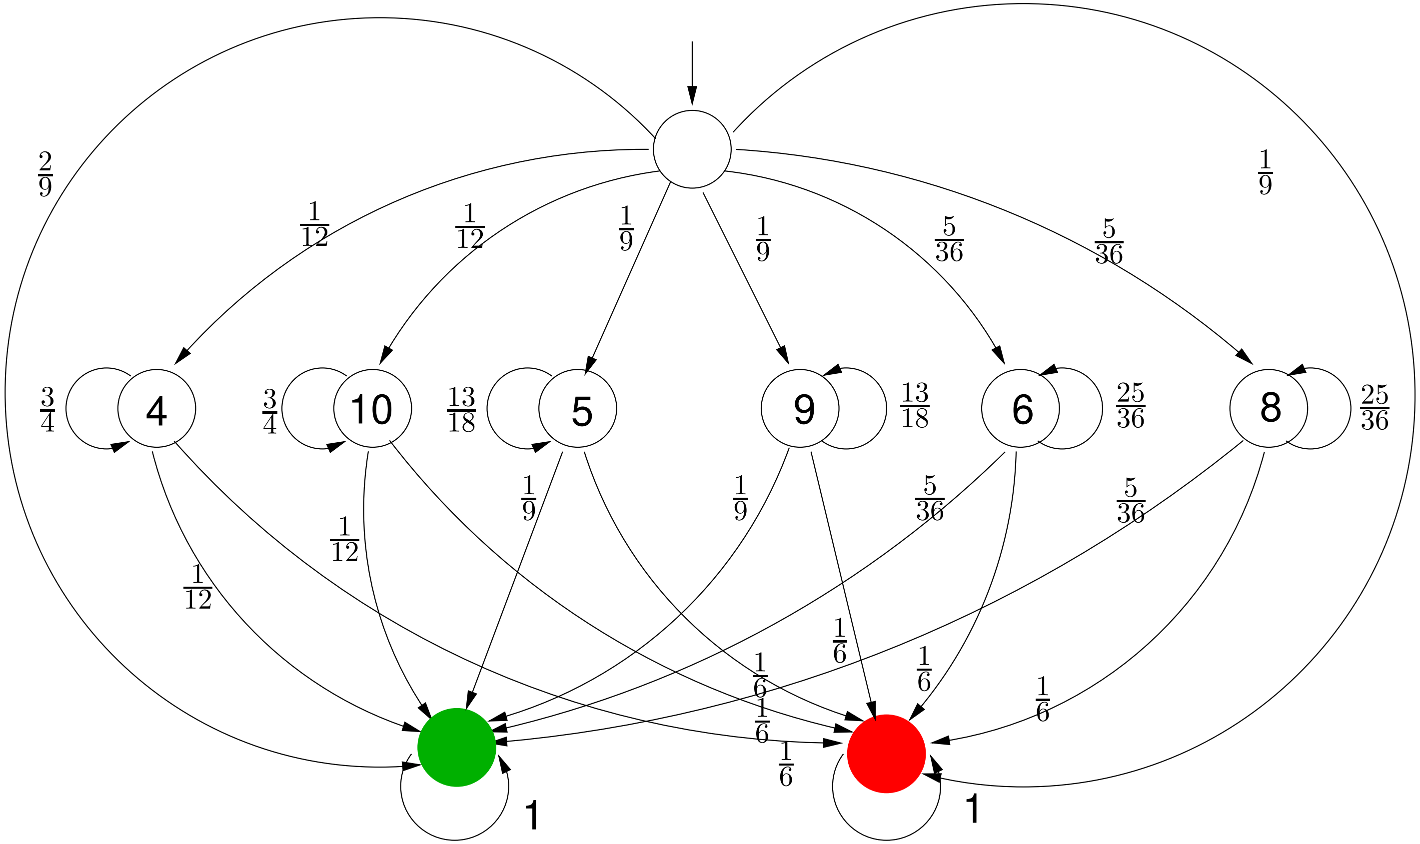

Minimizing Craps

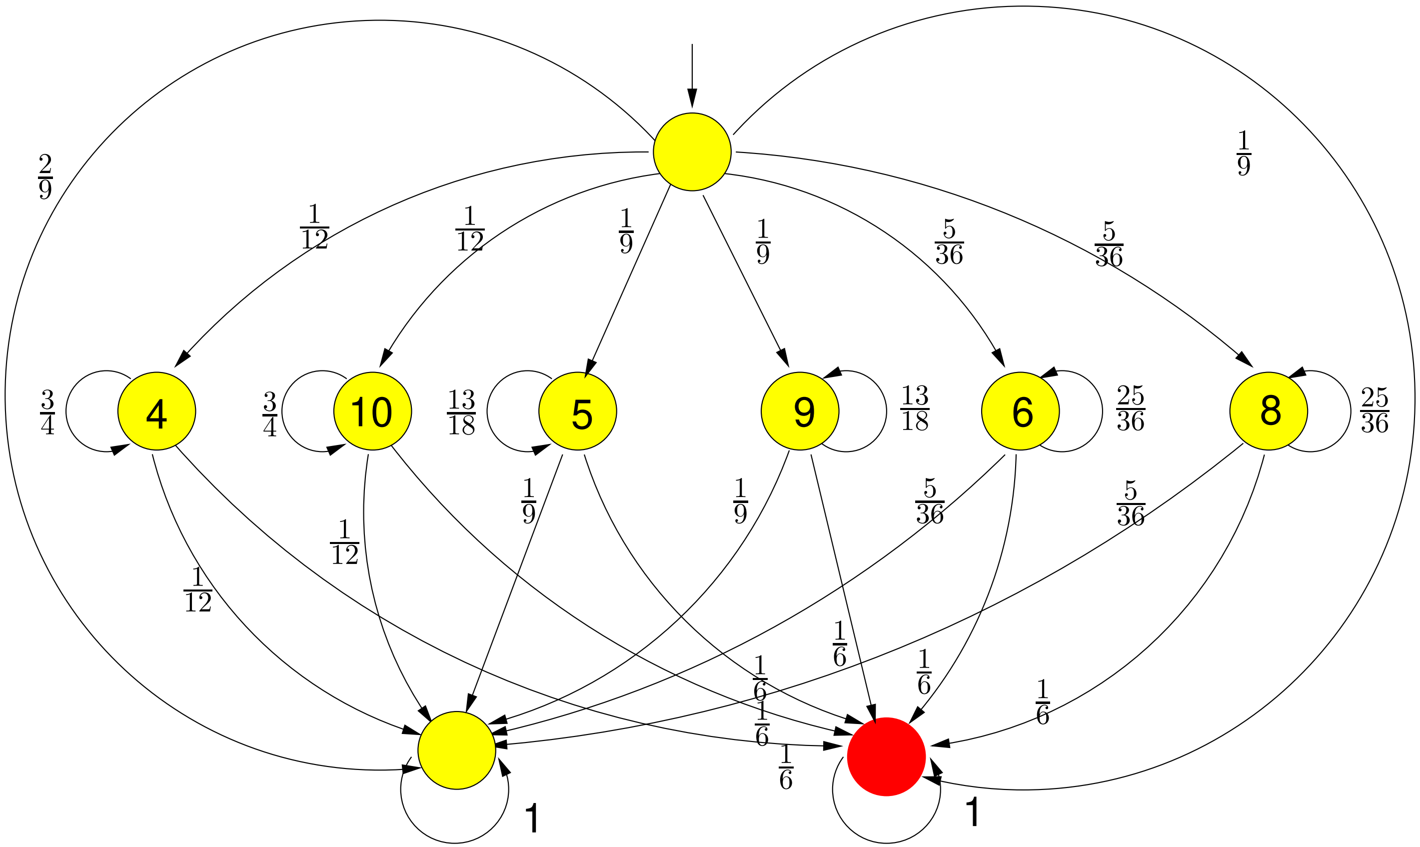

initial partitioning for the atomic propositions $AP = \{ \textit{loss} \}$

A first refinement

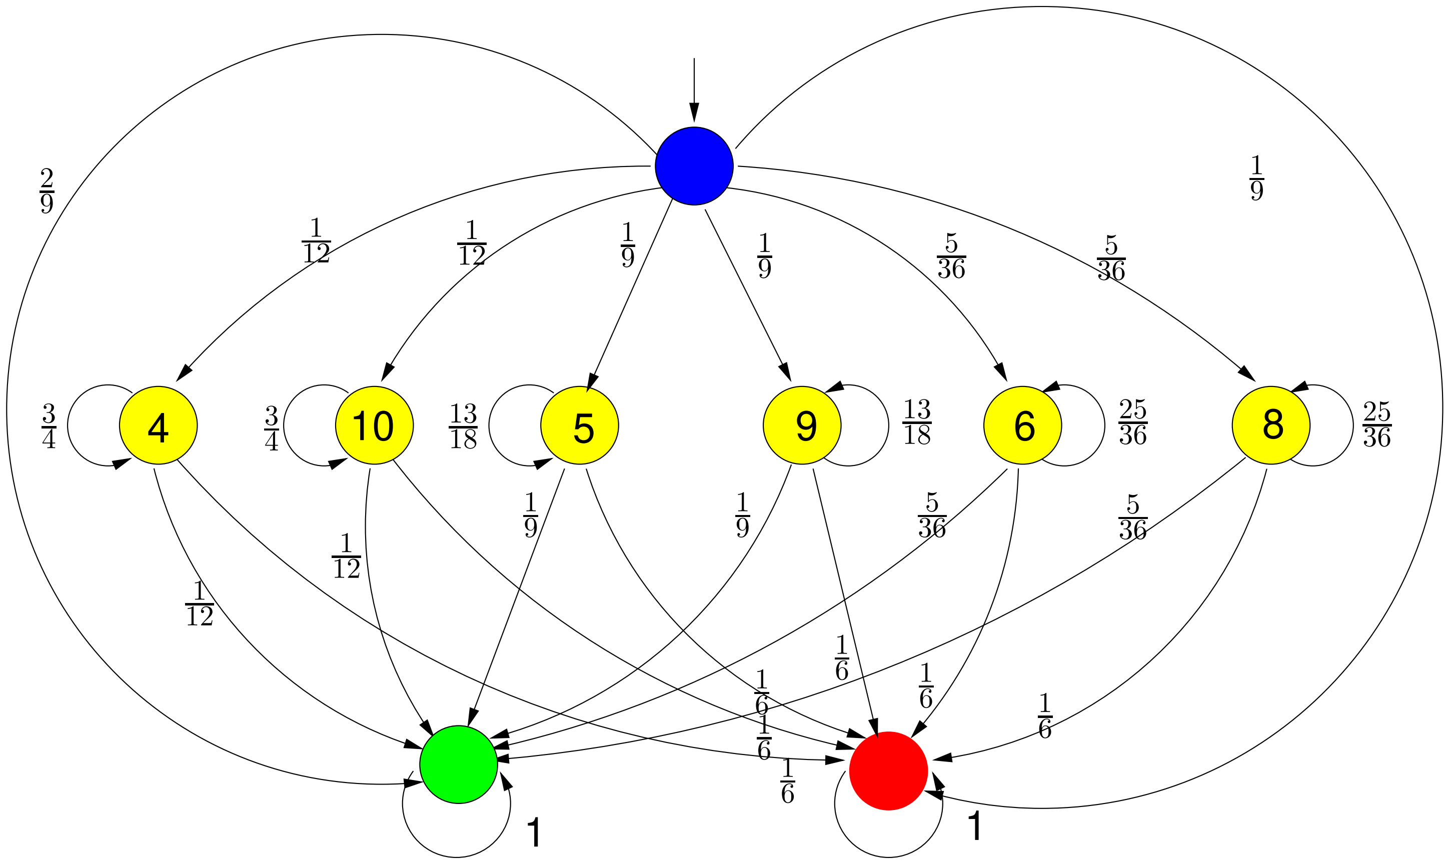

refine (“split”) with respect to the set of red states

A second refinement

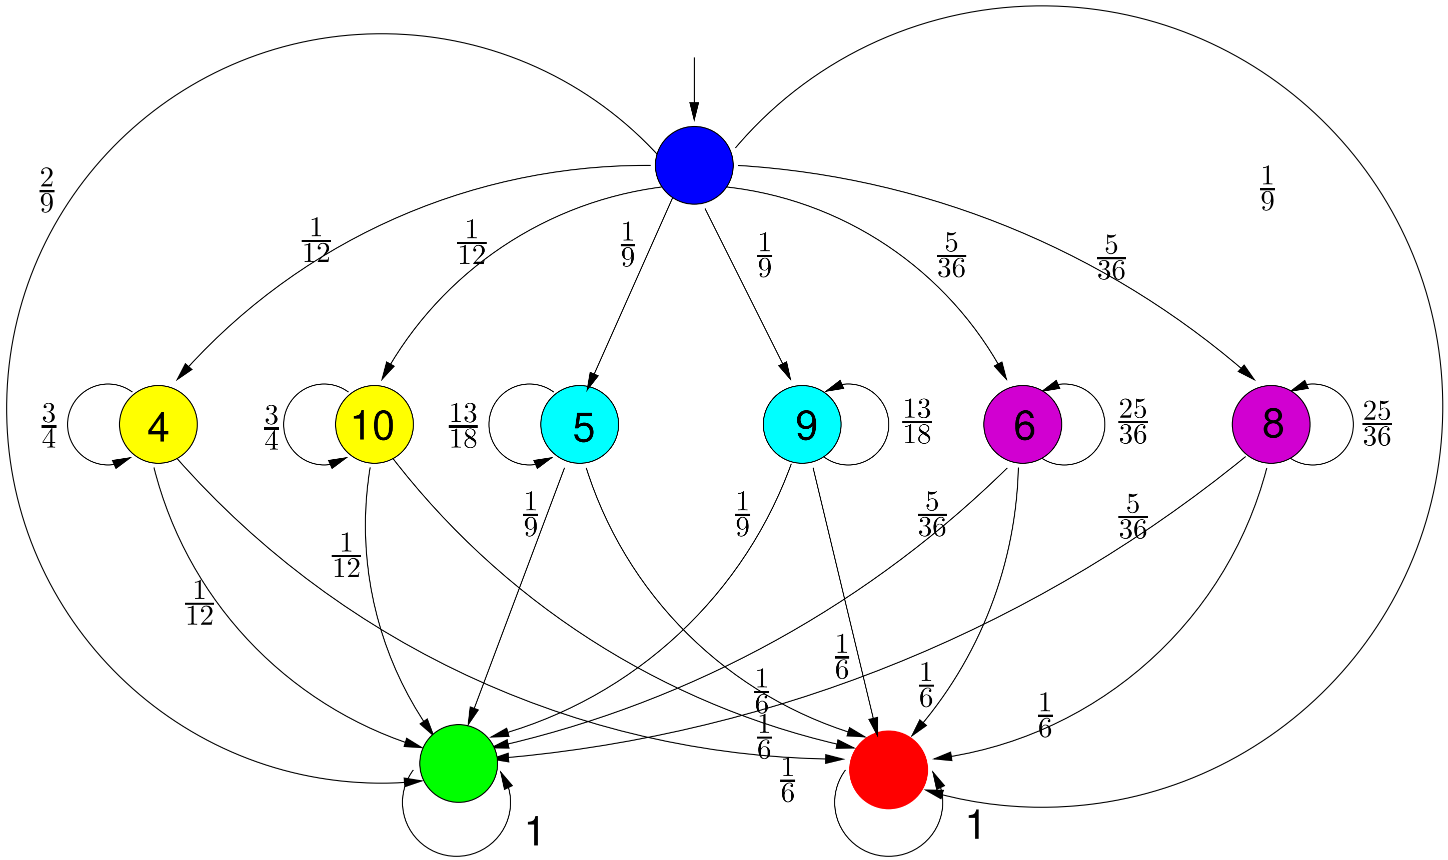

refine (“split”) with respect to the set of green states

Quotient DTMC

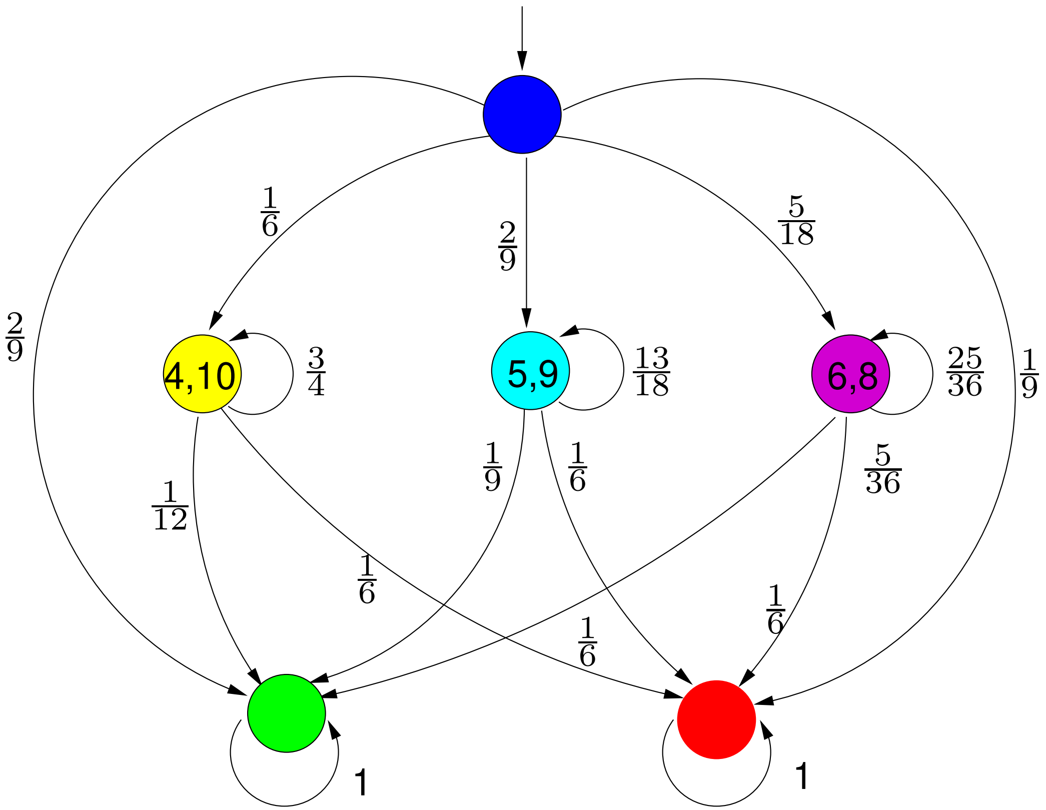

Preservation of PCTL

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

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**

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

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

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

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

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

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

Weight function example

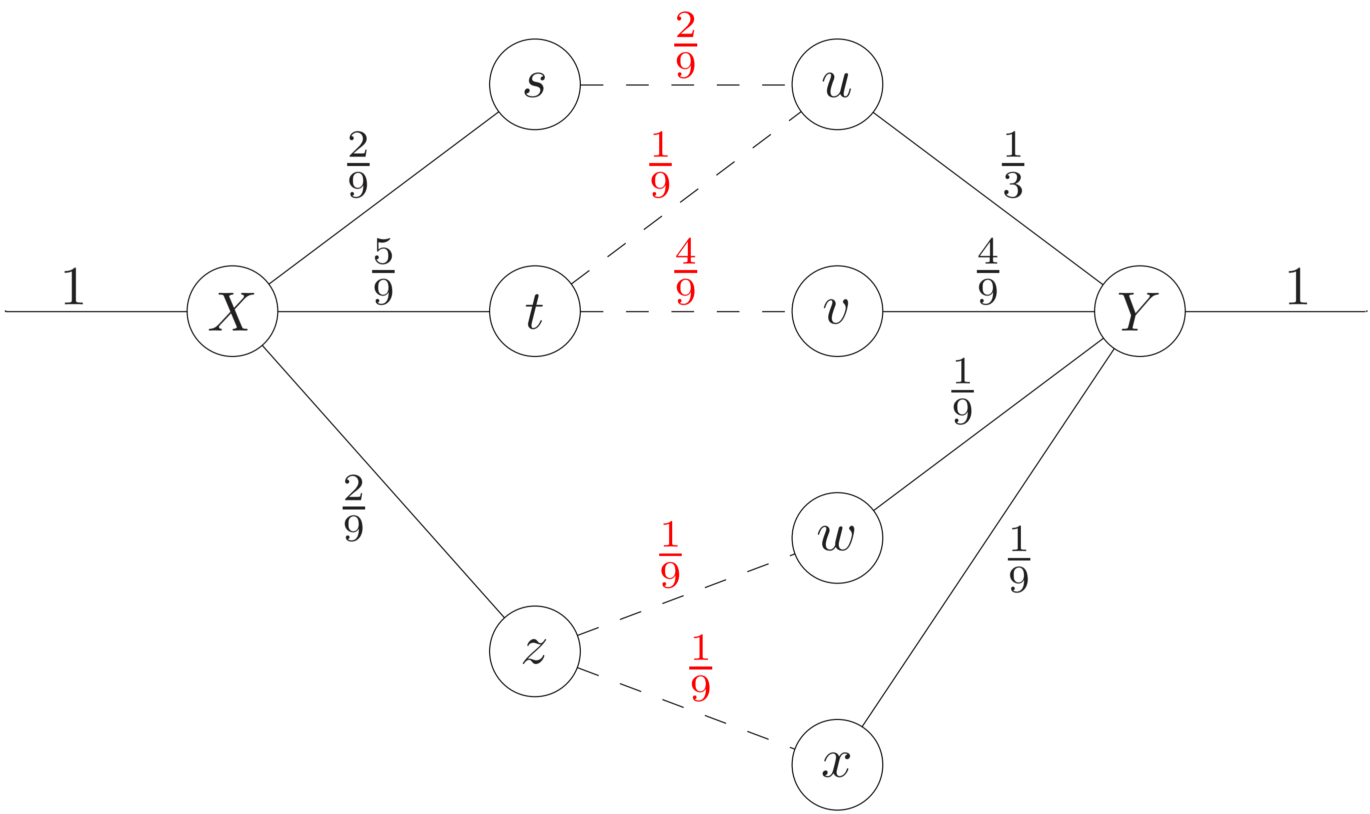

Probabilistic simulation

- 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

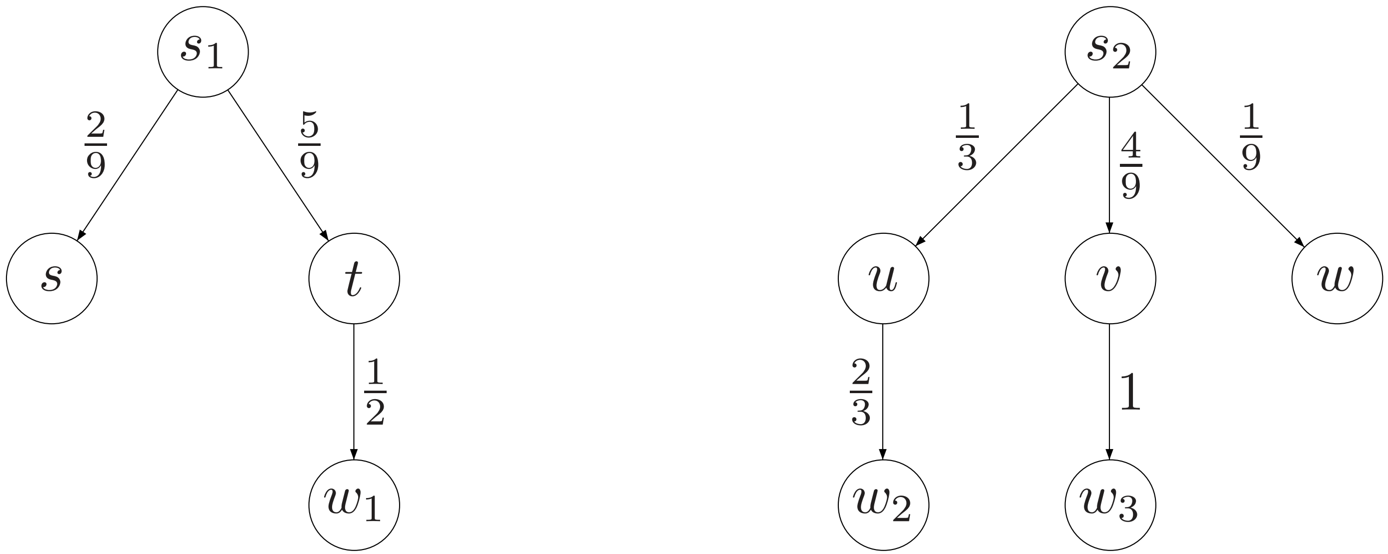

$$R = \{ (s_1, s_2), (s, u), (t, u), (t, v), (w_1, w_2), (w_1, w_3) \}$$

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

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

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

$$\boxed{\Phi \;::=\; \text{true} \;\big|\; a \;\big|\; \neg a \;\big|\; \Phi \wedge \Phi \;\big|\; \Phi \vee \Phi \;\big|\; \mathbb{P}_{\geqslant p}(\Phi \,\mathsf{W}\, \Phi) \;\big|\; \mathbb{P}_{\geqslant p}(\Phi \,\mathsf{W}^{\leqslant n} \Phi)}$$

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 $\precapprox$
logical preservation	PCTL	$\text{PCTL}_{\setminus \bigcirc}$	safePCTL	$\text{safePCTL}_{\setminus \bigcirc}$
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	$\mathcal{O}(m \log n)$	$\mathcal{O}(n^3)$	–	–

Thank you for the attendance

谢谢大家！