Model Checking
Continuous-Time Markov Chains

Joost-Pieter Katoen

Software Modeling and Verification Group

RWTH Aachen University

associated to University of Twente, Formal Methods and Tools

Lecture at MOVEP Summerschool, July 1, 2010

Content of this lecture

- Introduction
  - motivation, DTMCs, PCTL model checking
- Negative exponential distribution
  - definition, usage, properties
- Continuous-time Markov chains
  - definition, semantics, examples
- Performance measures
  - transient and steady-state probabilities, uniformization

Content of this lecture

⇒ Introduction

  – motivation, DTMCs, PCTL model checking

- Negative exponential distribution
  – definition, usage, properties
- Continuous-time Markov chains
  – definition, semantics, examples
- Performance measures
  – transient and steady-state probabilities, uniformization

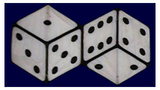

Probabilities help

- 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 problems are undecidable deterministically
  - reachability of channel systems, ...

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)

What is probabilistic model checking?

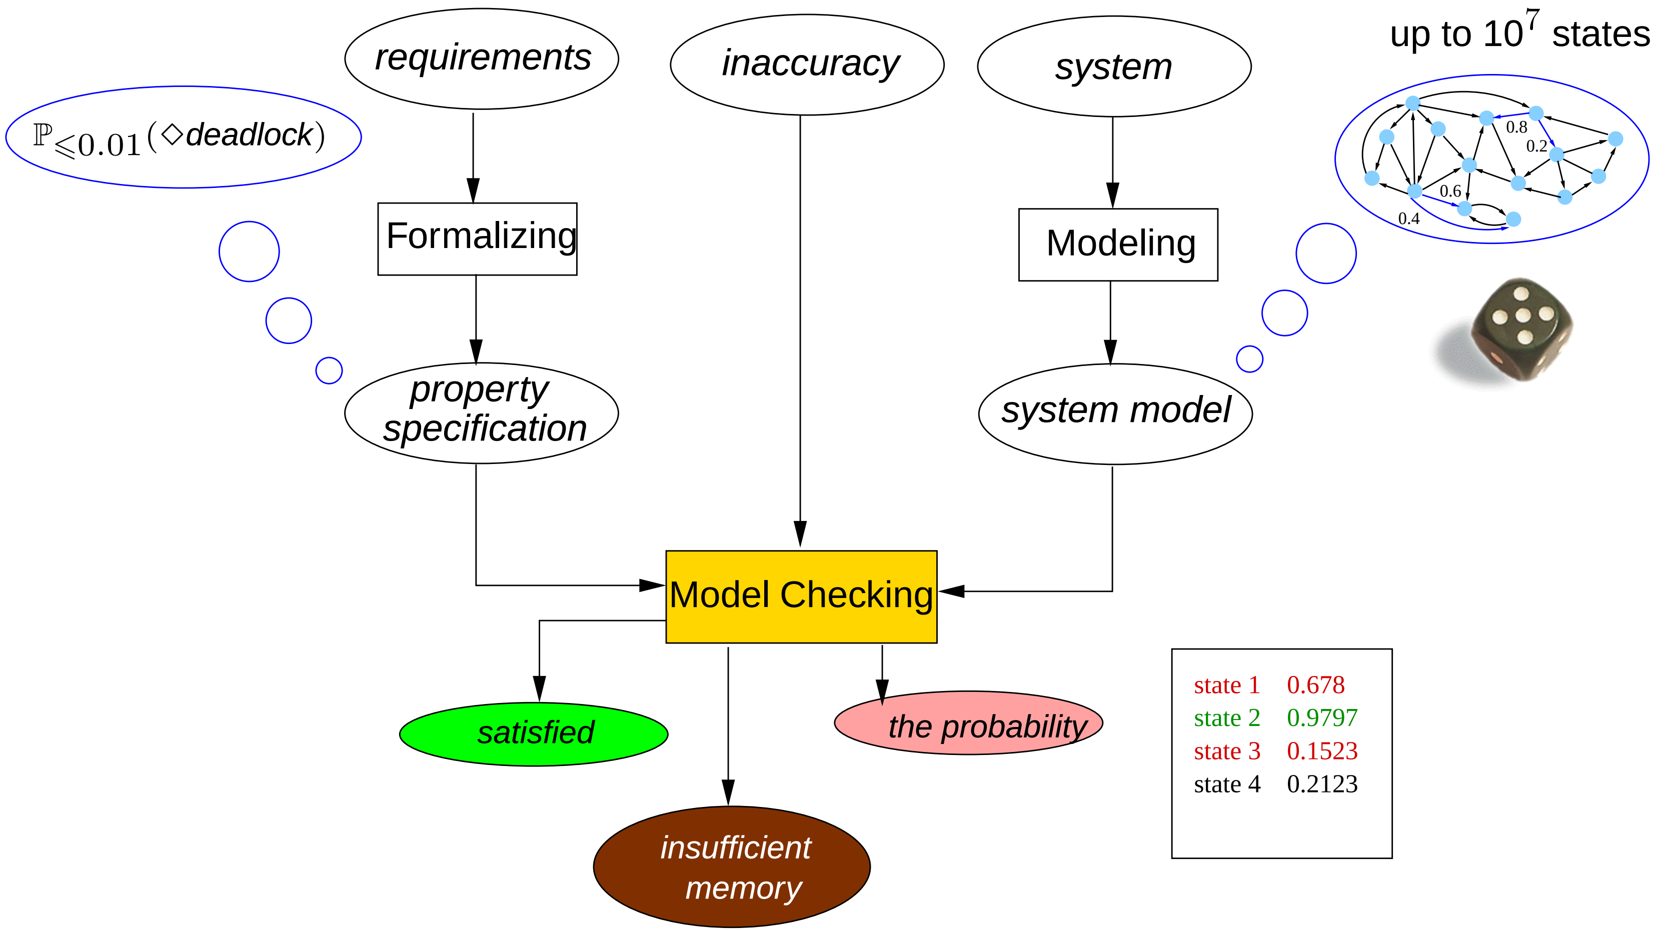

Probabilistic models

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

Other models: probabilistic variants of (priced) timed automata, or hybrid automata

Discrete-time Markov chain

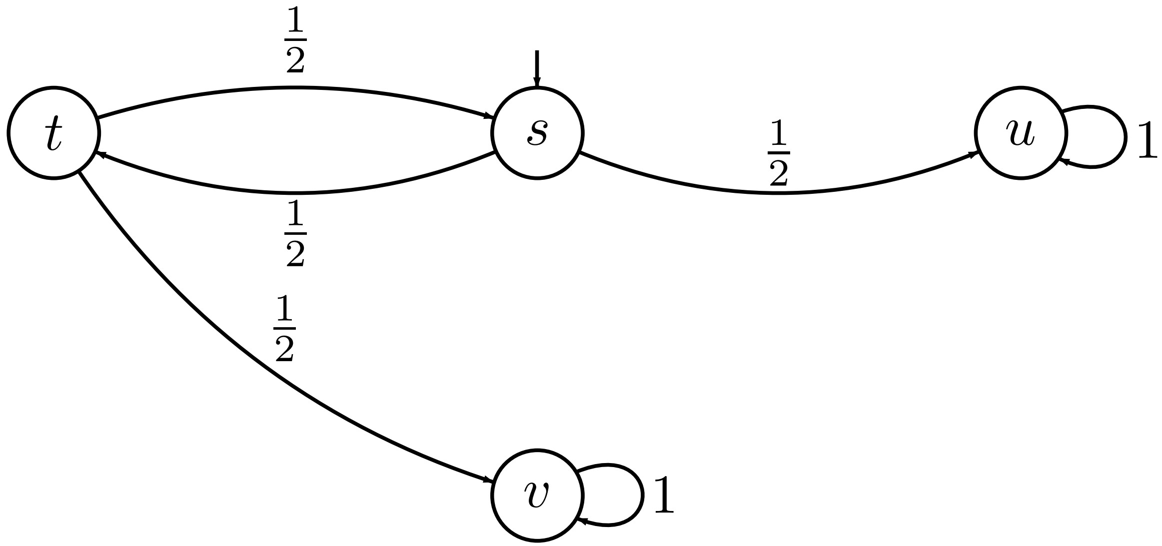

a DTMC $\mathcal{D}$ is a triple $(S, \mathbf{P}, L)$ with state space S and state-labelling L

and $\mathbf{P}$ a stochastic matrix with $\mathbf{P}(s, s') =$ one-step probability to jump from s to s'

Craps

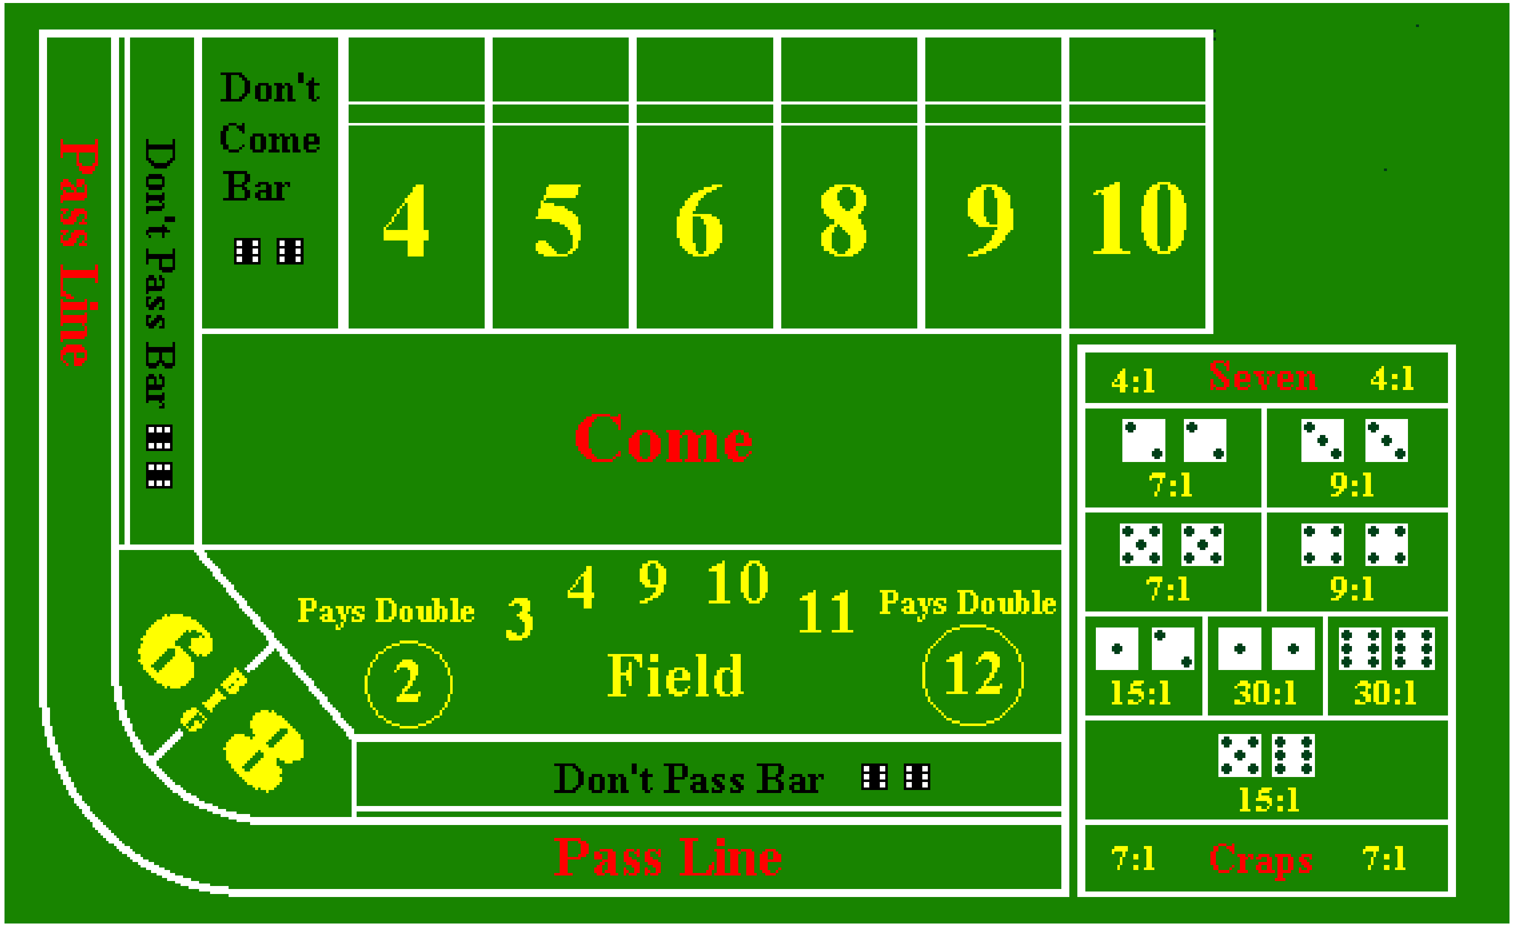

Pass Line
Don't Pass Bar
Don't Come Bar
4
5
6
8
9
10
Come
4:1 Seven 4:1
7:1
9:1
7:1
9:1
15:1
30:1
30:1
15:1
7:1 Craps 7:1
Pays Double
3
4
9
10
11
Pays Double
2
Field
12
6
8
Don't Pass Bar
Pass Line

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

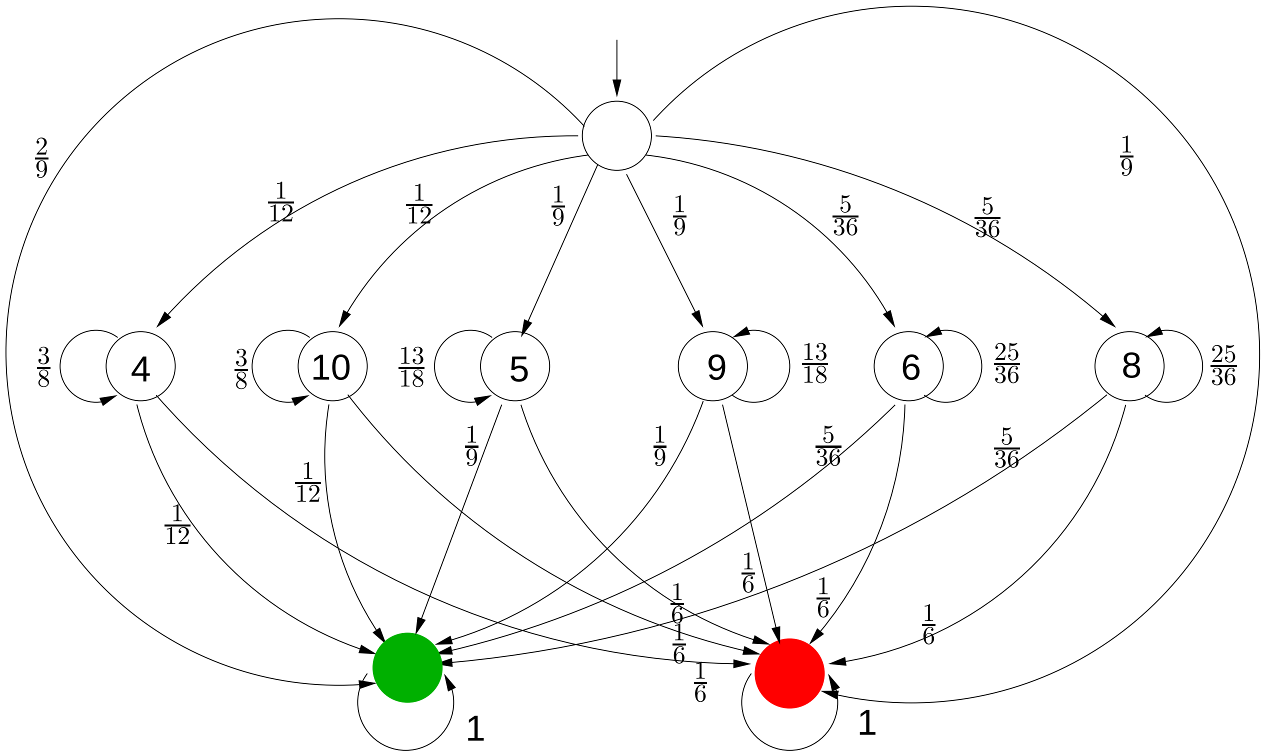

Probability measure on DTMCs

- Events are *infinite paths* in the DTMC $\mathcal{D}$, i.e., $\Omega = \textit{Paths}(\mathcal{D})$
  - a path in a DTMC is just a sequence of states
- A σ-algebra on $\mathcal{D}$ is generated by *cylinder sets* of finite paths $\hat{\pi}$:

$$\textit{Cyl}(\hat{\pi}) \; = \; \left\{ \pi \in \textit{Paths}(\mathcal{D}) \mid \hat{\pi} \text{ is a prefix of } \pi \right\}$$

  - cylinder sets serve as basis events of the smallest σ-algebra on $\textit{Paths}(\mathcal{D})$
- $\Pr$ is the *probability measure* on the σ-algebra on $\textit{Paths}(\mathcal{D})$:

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

- where $\mathbf{P}(s_0\, s_1 \ldots s_n) \;=\; \prod\limits_{0 \leqslant i < n} \mathbf{P}(s_i, s_{i+1})$ and $\mathbf{P}(s_0) = 1$, and
- $\iota_{init}(s_0)$ is the initial probability to start in state s_0

Reachability probabilities

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

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

Reachability probabilities in finite DTMCs

- Let $\Pr(s \models \Diamond B) = \Pr_s(\Diamond B) = \Pr_s\{\pi \in \textit{Paths}(s) \mid \pi \models \Diamond B\}$
  - where $\Pr_s$ is the probability measure in $\mathcal{D}$ 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 $s \in \textit{Pre}^*(B) \setminus B$:

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

Unique solution

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

- $S_{=0} \ = \ \textit{Sat}(\neg\exists(C \,\mathsf{U}\, B))$
- $S_{=1}$ 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 $\big(\Pr(s \models C \,\mathsf{U}\, B) \big)_{s \in S_?}$

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

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

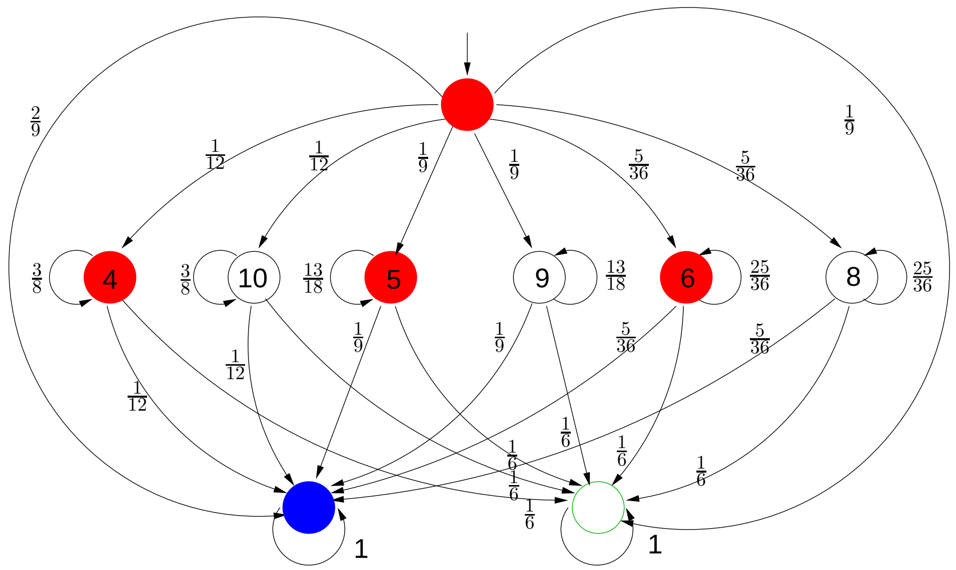

Example: Craps game

- *start* $< 4 < 5 < 6$
- $\mathbf{A} = \frac{1}{36} \begin{pmatrix} 0 & 3 & 4 & 5 \\ 0 & 27 & 0 & 0 \\ 0 & 0 & 26 & 0 \\ 0 & 0 & 0 & 25 \end{pmatrix}$
- $\mathbf{b} = \frac{1}{36} \begin{pmatrix} 8 \\ 3 \\ 4 \\ 5 \end{pmatrix}$

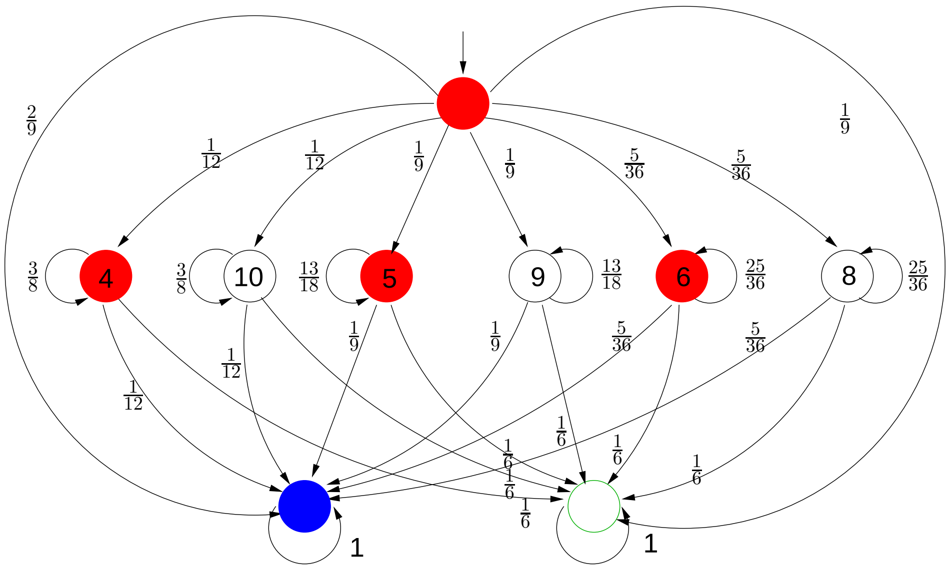

$$\mathbf{x}^{(0)} = \mathbf{0} \quad \text{and} \quad \mathbf{x}^{(i+1)} = \mathbf{A}\mathbf{x}^{(i)} + \mathbf{b} \text{ for } 0 \leqslant i < n.$$

Example: Craps game

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

PCTL Syntax

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

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

- $s_0 s_1 s_2 \ldots \models \Phi \, \mathsf{U}^{\leqslant n} \, \Psi$ 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

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

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

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

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

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:

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

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

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

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

PCTL semantics (1)

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

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 $\mathcal{D}$ 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 $\mathcal{D}$

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

PCTL model checking

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

- 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?

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 $\mathcal{D}$
- Check $\Diamond^{=h}\,\Psi$ in the obtained DTMC $\mathcal{D}'$
- This is a standard transient analysis in $\mathcal{D}'$:

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

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

$\Rightarrow$ Matrix-vector multiplication

Time complexity

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

$$\mathcal{O}\Big(\, poly(|\mathcal{D}|) \, \cdot \, n_{\max} \cdot |\Phi| \, \Big)$$

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

The qualitative fragment of PCTL

- For $a \in AP$:

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

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

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

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
  - motivation, DTMCs, PCTL model checking

$\Rightarrow$ Negative exponential distribution
  - definition, usage, properties

- Continuous-time Markov chains
  - definition, semantics, examples

- Performance measures
  - transient and steady-state probabilities, uniformization

Time in DTMCs

- Time in a DTMC proceeds in discrete steps
- Two possible interpretations
  - accurate model of (discrete) time units
    - e.g., clock ticks in model of an embedded device
  - time-abstract
    - no information assumed about the time transitions take
- Continuous-time Markov chains (CTMCs)
  - dense model of time
  - transitions can occur at any (real-valued) time instant
  - modelled using negative exponential distributions

Continuous random variables

- X is a random variable (r.v., for short)
  - on a sample space with probability measure Pr
  - assume the set of possible values that X may take is dense

- X is *continuously distributed* if there exists a function $f(x)$ such that:

$$\mathrm{Pr}\{X \leqslant d\} \; = \; \int_{-\infty}^{d} f(x)\,dx \qquad \text{for each real number } d$$

where f satisfies: $f(x) \geqslant 0$ for all x and $\int_{-\infty}^{\infty} f(x)\,dx = 1$

  - $F_X(d) = \mathrm{Pr}\{X \leqslant d\}$ is the *(cumulative) probability distribution function*
  - $f(x)$ is the *probability density function*

Negative exponential distribution

The density of an *exponentially distributed* r.v. Y with rate $\lambda \in \mathbb{R}_{>0}$ is:

$$f_Y(x) = \lambda{\cdot}e^{-\lambda \cdot x} \quad \text{for } x > 0 \quad \text{and } f_Y(x) = 0 \text{ otherwise}$$

The cumulative distribution of Y:

$$F_Y(d) \;=\; \int_0^d \lambda{\cdot}e^{-\lambda \cdot x}\, dx \;=\; \left[-e^{-\lambda \cdot x}\right]_0^d \;=\; 1 - e^{-\lambda \cdot d}$$

- expectation $E[Y] = \int_0^\infty x{\cdot}\lambda{\cdot}e^{-\lambda \cdot x}\, dx = \frac{1}{\lambda}$
- variance *Var*$[Y] = \frac{1}{\lambda^2}$

the rate $\lambda \in \mathbb{R}_{>0}$ uniquely determines an exponential distribution.

Exponential pdf and cdf

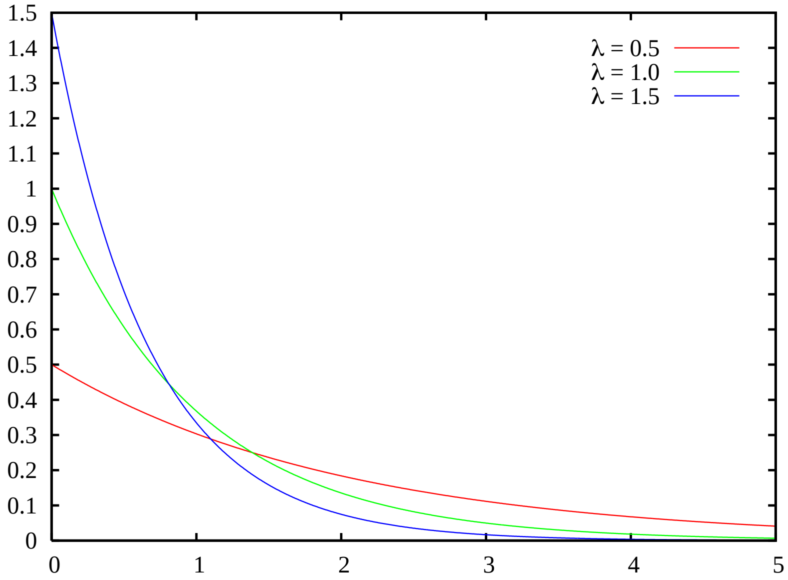

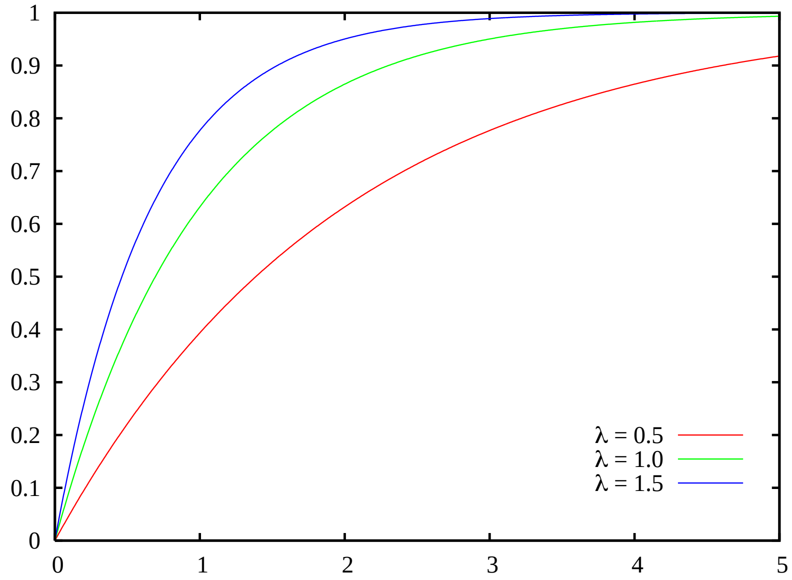

the higher λ, the faster the cdf approaches 1

Why exponential distributions?

- Are *adequate* for many real-life phenomena
  - the time until a radioactive particle decays
  - the time between successive car accidents
  - inter-arrival times of jobs, telephone calls in a fixed interval
- Are the continuous counterpart of *geometric* distribution
- Heavily used in physics, performance, and reliability analysis
- Can *approximate* general distributions arbitrarily closely
- Yield a *maximal entropy* if only the mean is known

Memoryless property

1. For any random variable X with an exponential distribution:

   $\Pr\{X > t + d \mid X > t\} \;=\; \Pr\{X > d\}$ for any $t, d \in \mathbb{R}_{\geqslant 0}$.

2. Any continuous distribution which is memoryless is an exponential one.

Proof of 1. : Let λ be the rate of X's distribution. Then we derive:

$$\Pr\{X > t + d \mid X > t\} \;=\; \frac{\Pr\{X > t{+}d \,\cap\, X > t\}}{\Pr\{X > t\}} \;=\; \frac{\Pr\{X > t{+}d\}}{\Pr\{X > t\}}$$

$$= \frac{e^{-\lambda\cdot(t+d)}}{e^{-\lambda\cdot t}} \;=\; e^{-\lambda\cdot d} \;=\; \Pr\{X > d\}.$$

Proof of 2. : by contradiction, using the total law of probability.

Closure under minimum

For independent, exponentially distributed random variables X and Y with rates $\lambda, \mu \in \mathbb{R}_{>0}$, r.v. $\min(X, Y)$ is exponentially distributed with rate $\lambda + \mu$, i.e.,:

$$\Pr\{\min(X, Y) \leqslant t\} \;=\; 1 - e^{-(\lambda+\mu)\cdot t} \quad \text{for all } t \in \mathbb{R}_{\geqslant 0}$$

Proof

Let λ (μ) be the rate of X's (Y's) distribution. Then we derive:

$$
\begin{aligned}
\Pr\{\min(X, Y) \leqslant t\} &= \Pr_{X,Y}\{(x, y) \in \mathbb{R}^2_{\geqslant 0} \mid \min(x, y) \leqslant t\} \\
&= \int_0^\infty \left(\int_0^\infty \mathbf{I}_{\min(x,y)\leqslant t}(x, y) \cdot \lambda e^{-\lambda x} \cdot \mu e^{-\mu y}\, dy \right)\, dx \\
&= \int_0^t \int_x^\infty \lambda e^{-\lambda x} \cdot \mu e^{-\mu y}\, dy\, dx + \int_0^t \int_y^\infty \lambda e^{-\lambda x} \cdot \mu e^{-\mu y}\, dx\, dy \\
&= \int_0^t \lambda e^{-\lambda x} \cdot e^{-\mu x}\, dx + \int_0^t e^{-\lambda y} \cdot \mu e^{-\mu y}\, dy \\
&= \int_0^t \lambda e^{-(\lambda+\mu)x}\, dx + \int_0^t \mu e^{-(\lambda+\mu)y}\, dy \\
&= \int_0^t (\lambda+\mu) \cdot e^{-(\lambda+\mu)z}\, dz \;=\; 1 - e^{-(\lambda+\mu)t}
\end{aligned}
$$

Winning the race with two competitors

For independent, exponentially distributed random variables

X and Y with rates $\lambda, \mu \in \mathbb{R}_{>0}$, it holds:

$$\Pr\{X \leqslant Y\} = \frac{\lambda}{\lambda + \mu}$$

Proof

Let λ (μ) be the rate of X's (Y's) distribution. Then we derive:

$$
\begin{aligned}
\Pr\{X \leqslant Y\} &= \Pr_{X,Y}\{(x,y) \in \mathbb{R}^2_{\geqslant 0} \mid x \leqslant y\} \\
&= \int_0^\infty \mu e^{-\mu y} \left(\int_0^y \lambda e^{-\lambda x}\, dx \right) dy \\
&= \int_0^\infty \mu e^{-\mu y} \left(1 - e^{-\lambda y} \right) dy \\
&= 1 - \int_0^\infty \mu e^{-\mu y} \cdot e^{-\lambda y}\, dy \;=\; 1 - \int_0^\infty \mu e^{-(\mu+\lambda)y}\, dy \\
&= 1 - \frac{\mu}{\mu+\lambda} \cdot \underbrace{\int_0^\infty (\mu+\lambda) e^{-(\mu+\lambda)y}\, dy}_{=1} \\
&= 1 - \frac{\mu}{\mu+\lambda} \;=\; \frac{\lambda}{\mu+\lambda}
\end{aligned}
$$

Winning the race with many competitors

For independent, exponentially distributed random variables

$X_1, X_2, \ldots, X_n$ with rates $\lambda_1, \ldots, \lambda_n \in \mathbb{R}_{>0}$, it holds:

$$\Pr\{X_i = \min(X_1, \ldots, X_n)\} = \frac{\lambda_i}{\sum_{j=1}^{n} \lambda_j}$$

Content of this lecture

- Introduction
  - motivation, DTMCs, PCTL model checking
- Negative exponential distribution
  - definition, usage, properties

$\Rightarrow$ Continuous-time Markov chains
  - definition, semantics, examples

- Performance measures
  - transient and steady-state probabilities, uniformization

Continuous-time Markov chain

A *continuous-time Markov chain* (CTMC) is a tuple $(S, \mathbf{P}, r, L)$ where:

- S is a countable (today: finite) set of *states*
- $\mathbf{P} : S \times S \rightarrow [0, 1]$, a *stochastic matrix*
  - $\mathbf{P}(s, s')$ is one-step probability of going from state s to state s'
  - s is called *absorbing* iff $\mathbf{P}(s, s) = 1$
- $r : S \rightarrow \mathbb{R}_{>0}$, the *exit-rate function*
  - $r(s)$ is the rate of exponential distribution of residence time in state s

$\Rightarrow$ a CTMC is a Kripke structure with random state residence times

Continuous-time Markov chain

a CTMC $(S, \mathbf{P}, r, L)$ is a DTMC plus an exit-rate function $r : S \rightarrow \mathbb{R}_{>0}$

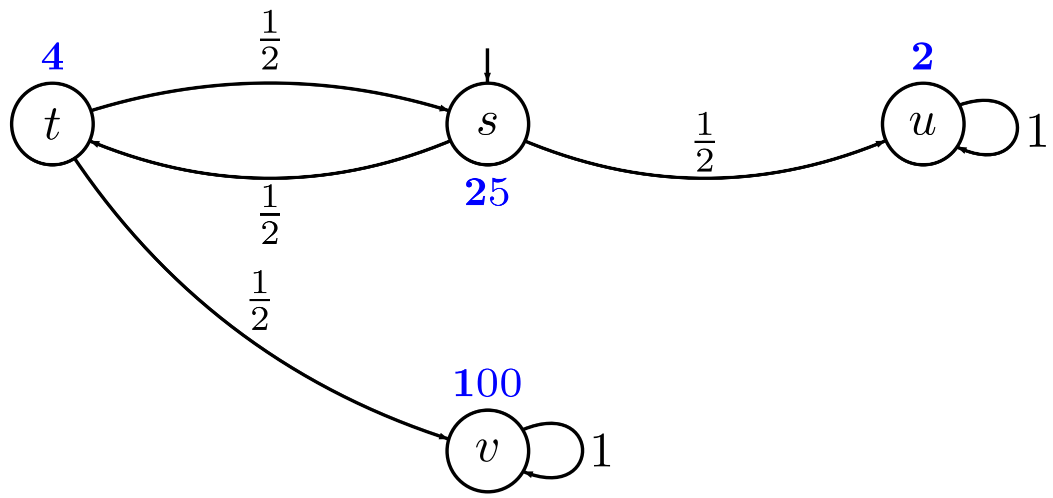

the average residence time in state s is $\dfrac{1}{r(s)}$

A classical (though equivalent) perspective

a CTMC is a triple $(S, \mathbf{R}, L)$ with $\mathbf{R}(s, s') = \mathbf{P}(s, s') \cdot r(s)$

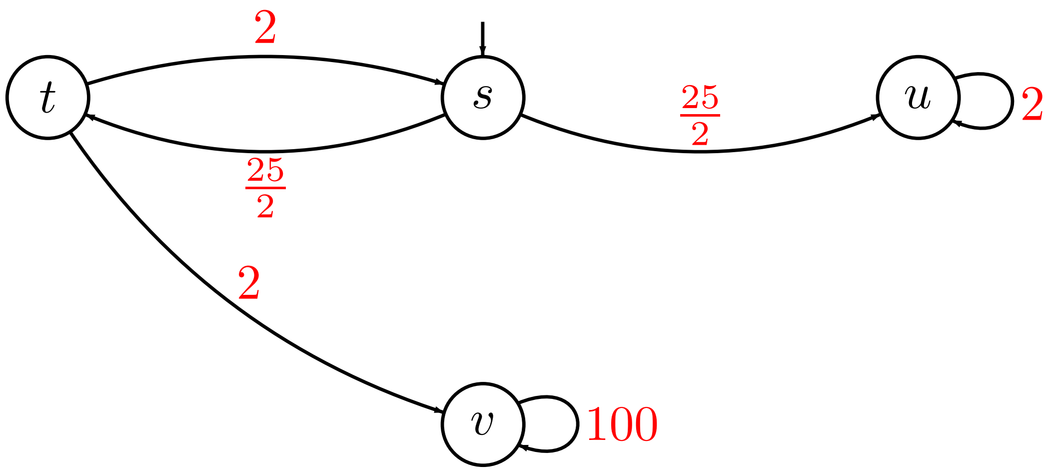

CTMC semantics: example

- Transition $s \longrightarrow s'$:= r.v. $X_{s,s'}$ with rate $\mathbf{R}(s, s')$

- Probability to go from state s_0 to, say, state s_2 is:

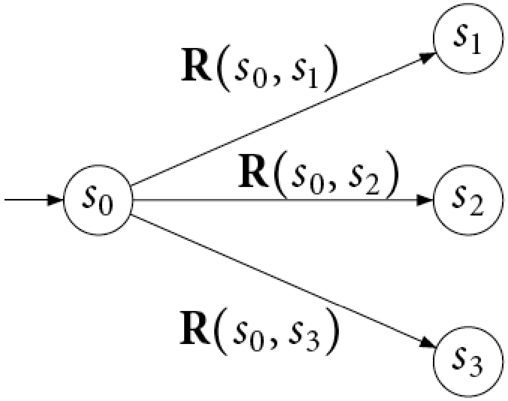

$$\Pr\{X_{s_0,s_2} \leqslant X_{s_0,s_1} \cap X_{s_0,s_2} \leqslant X_{s_0,s_3}\}$$
$$=$$
$$\frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0, s_1) + \mathbf{R}(s_0, s_2) + \mathbf{R}(s_0, s_3)} = \frac{\mathbf{R}(s_0, s_2)}{r(s_0)}$$

- Probability of staying at most t time in s_0 is:

$$\Pr\{\min(X_{s_0,s_1}, X_{s_0,s_2}, X_{s_0,s_3}) \leqslant t\}$$
$$=$$
$$1 - e^{-(\mathbf{R}(s_0,s_1)+\mathbf{R}(s_0,s_2)+\mathbf{R}(s_0,s_3))\cdot t} = 1 - e^{-r(s_0)\cdot t}$$

CTMC semantics

- The probability that transition $s \longrightarrow s'$ is *enabled* in $[0, t]$:

$$1 - e^{-\mathbf{R}(s,s')\cdot t}$$

- The probability to *move* from non-absorbing s to s' in $[0, t]$ is:

$$\frac{\mathbf{R}(s, s')}{r(s)} \cdot \left(1 - e^{-r(s)\cdot t}\right)$$

- The probability to *take some* outgoing transition from s in $[0, t]$ is:

$$\int_0^t r(s)\cdot e^{-r(s)\cdot x}\,dx \;=\; 1 - e^{-r(s)\cdot t}$$

Enzyme-catalysed substrate conversion

Enzyme - Wikipedia, the free encyclopedia - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://en.wikipedia.org/wiki/Enzyme

reaction, the reaction is *effectively* irreversible. Under these conditions the enzyme will, in fact, only catalyze the reaction in the thermodynamically allowed direction.

stabilizes the transition state, reducing the energy needed to form this species and thus reducing the energy required to form products.

Kinetics

Main article: Enzyme kinetics

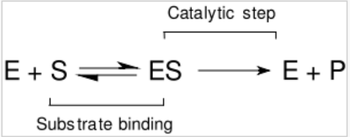

Mechanism for a single substrate enzyme catalyzed reaction. The enzyme (E) binds a substrate (S) and produces a product (P).

Enzyme kinetics is the investigation of how enzymes bind substrates and turn them into products. The rate data used in kinetic analyses are obtained from enzyme assays.

In 1902 Victor Henri [45] proposed a quantitative theory of enzyme kinetics, but his experimental data were not useful because the significance of the hydrogen ion concentration was not yet appreciated. After Peter Lauritz Sørensen had defined the logarithmic pH-scale and introduced the concept of buffering in 1909[46] the German chemist Leonor Michaelis and his Canadian postdoc Maud Leonora Menten repeated Henri's experiments and confirmed his equation which is referred to as Henri-Michaelis-Menten kinetics (sometimes also Michaelis-Menten kinetics).[47] Their work was further developed by G. E. Briggs and J. B. S. Haldane, who derived kinetic equations that are still widely used today.[48]

The major contribution of Henri was to think of enzyme reactions in two stages. In the first, the substrate binds reversibly to the enzyme, forming the enzyme-substrate complex. This is sometimes called the Michaelis complex. The enzyme then catalyzes the chemical step in the reaction and releases the product.

Enzymes can catalyze up to several million reactions per second. For example, the reaction catalyzed by orotidine 5'-phosphate decarboxylase will consume half of its substrate in 78 million years if no enzyme is present. However, when the decarboxylase is added, the same process takes just 25 milliseconds.[49] Enzyme rates depend on solution conditions and substrate concentration. Conditions that denature the protein abolish enzyme activity, such as high temperatures, extremes of pH or high salt concentrations, while raising substrate concentration tends to increase activity. To find the maximum speed of an enzymatic reaction, the substrate concentration is increased until a constant rate of product formation is seen. This is shown in the saturation curve on the right. Saturation happens because, as substrate concentration increases, more and more of the free enzyme is converted into the substrate-bound ES form. At the maximum velocity (V_{max}) of the enzyme, all the enzyme active sites are bound to substrate, and the amount of ES complex is the same as the total amount of enzyme. However, V_{max} is only one kinetic constant of enzymes. The amount of substrate needed to achieve a given rate of reaction is also important. This is given by the Michaelis-Menten constant (K_m), which is the substrate concentration required for an enzyme to reach one-half its maximum velocity. Each enzyme has a characteristic K_m for a given substrate, and this can show how tight the binding of the substrate is to the enzyme. Another useful constant is

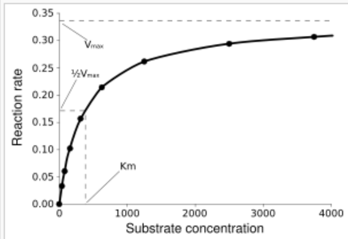

Saturation curve for an enzyme reaction showing the relation between the substrate concentration (S) and rate (*v*).

Done

Stochastic chemical kinetics

- Types of reaction described by stochiometric equations:

$$E + S \underset{k_2}{\overset{k_1}{\rightleftharpoons}} ES \xrightarrow{k_3} E + P$$

- N different types of molecules that randomly collide

  where state $X(t) = (x_1, \ldots, x_N)$ with $x_i = \#$ molecules of sort i

- Reaction probability within infinitesimal interval $[t, t{+}\Delta)$:

$$\alpha_m(\vec{x}) \cdot \Delta \;=\; \Pr\{\text{reaction } m \text{ in } [t, t{+}\Delta) \mid X(t) = \vec{x}\}$$

  where $\alpha_m(\vec{x}) = k_m \cdot \#$ possible combinations of reactant molecules in $\vec{x}$

- Process is a continuous-time Markov chain

Enzyme-catalyzed substrate conversion as a CTMC

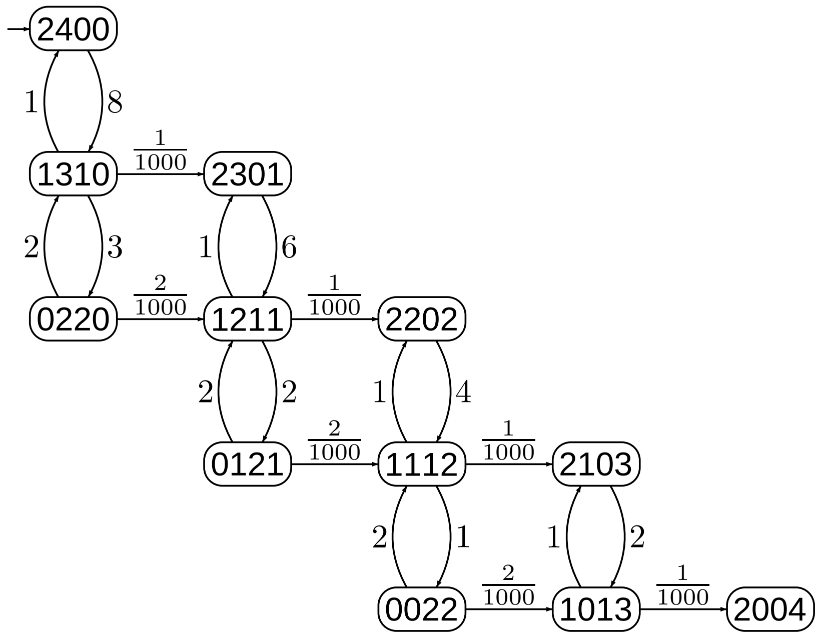

States:	*init*	*goal*
enzymes	2	2
substrates	4	0
complex	0	0
products	0	4

Transitions: $E + S \underset{1}{\overset{1}{\rightleftharpoons}} C \xrightarrow{0.001} E + P$

e.g., $(x_E, x_S, x_C, x_P) \xrightarrow{0.001 \cdot x_C} (x_E + 1, x_S, x_C - 1, x_P + 1)$ for $x_C > 0$

CTMCs are omnipresent!

- Markovian queueing networks (Kleinrock 1975)
- Stochastic Petri nets (Molloy 1977)
- Stochastic activity networks (Meyer & Sanders 1985)
- Stochastic process algebra (Herzog *et al.*, Hillston 1993)
- Probabilistic input/output automata (Smolka *et al.* 1994)
- Calculi for biological systems (Priami *et al.*, Cardelli 2002)

CTMCs are one of the most prominent models in performance analysis

Content of this lecture

- Introduction
  - motivation, DTMCs, PCTL model checking
- Negative exponential distribution
  - definition, usage, properties
- Continuous-time Markov chains
  - definition, semantics, examples

$\Rightarrow$ Performance measures

  - transient and steady-state probabilities, uniformization

Time-abstract evolution of a CTMC

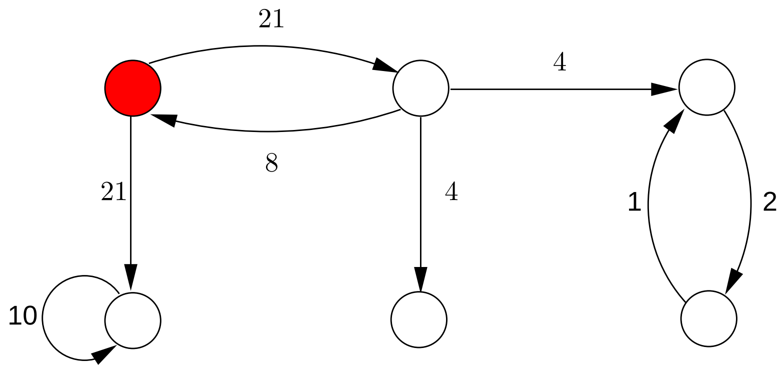

zero-th epoch

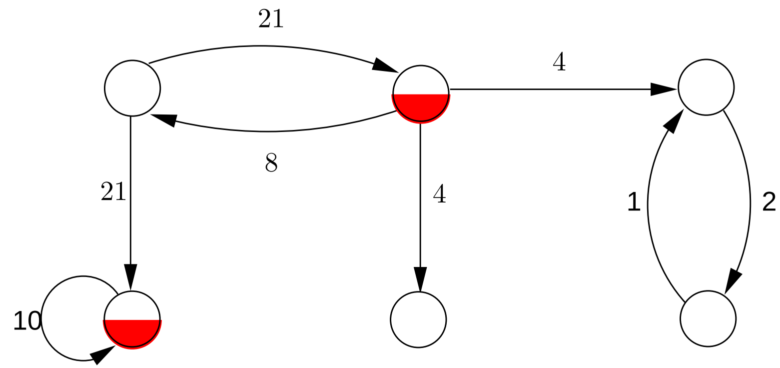

first epoch

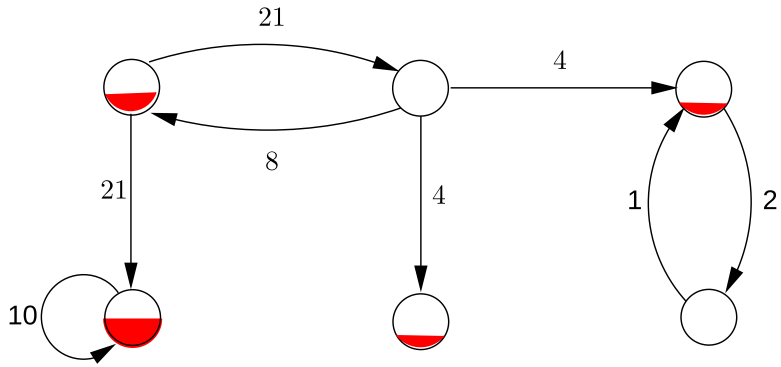

second epoch

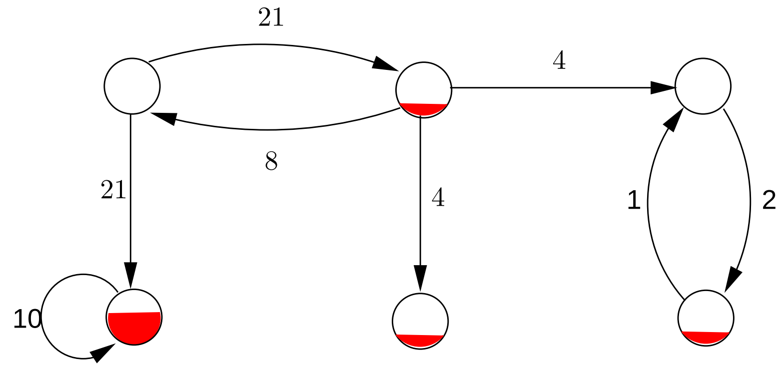

third epoch

On the long run

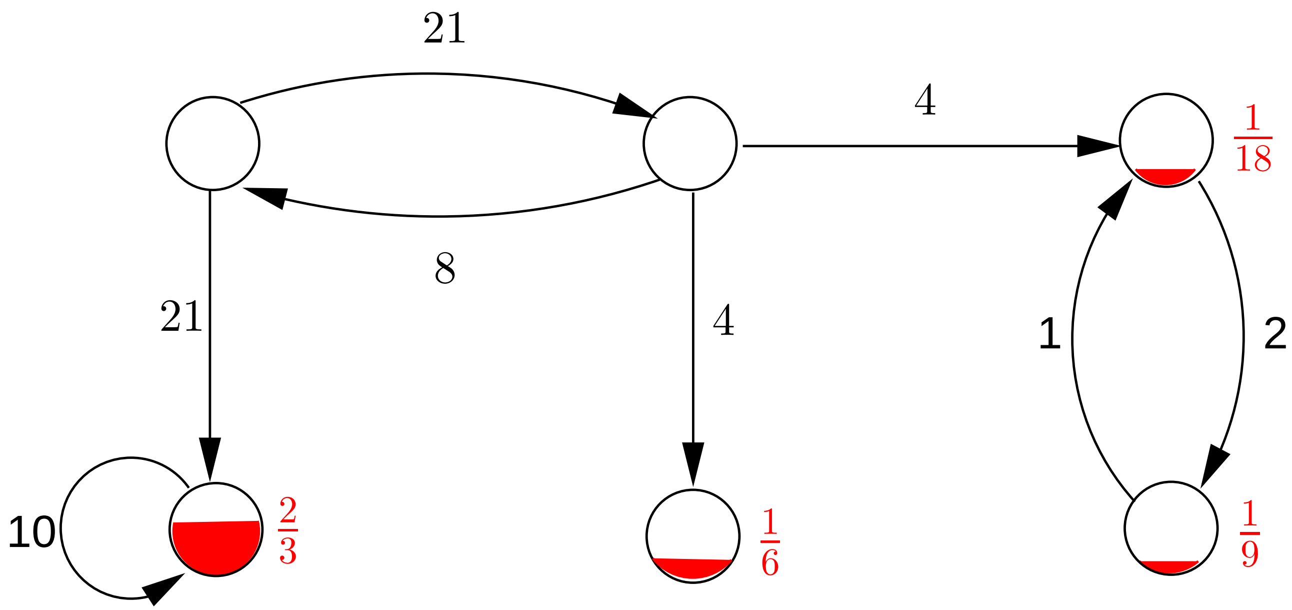

Transient distribution of a CTMC

Let $X(t)$ denote the state of a CTMC at time $t \in \mathbb{R}_{\geqslant 0}$.

Probability to be in state s at time t:

$$
\begin{aligned}
p_s(t) &= \Pr\{\, X(t) = s \,\} \\
&= \sum_{s' \in S} \Pr\{\, X(0) = s' \,\} \cdot \Pr\{\, X(t) = s \mid X(0) = s' \,\}
\end{aligned}
$$

Transient probability vector $\underline{p}(t) = (p_{s_1}(t), \ldots, p_{s_k}(t))$ satisfies:

$$\underline{p}'(t) = \underline{p}(t) \cdot (\mathbf{R} - \mathbf{r}) \qquad \text{given} \qquad \underline{p}(0)$$

where $\mathbf{r}$ is the diagonal matrix of vector $\underline{r}$.

A triple modular redundant system

- 3 processors and a single voter:
  - processors run same program; voter takes a majority vote
  - each component (processor and voter) is failure-prone
  - there is a single repairman for repairing processors and voter

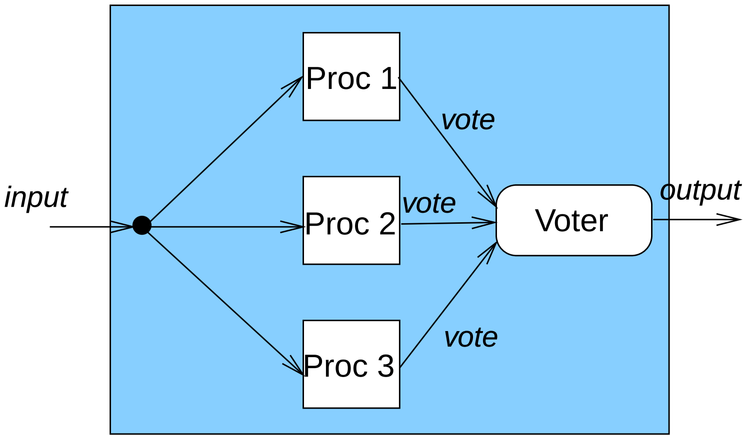

- Modelling assumptions:
  - if voter fails, entire system goes down
  - after voter-repair, system starts "as new"
  - state $=$ ($\#$processors, $\#$voters)

Modelling a TMR system as a CTMC

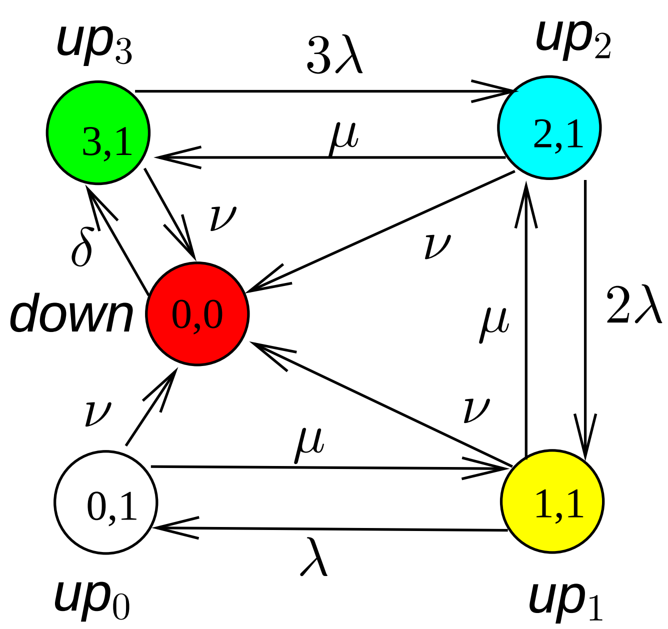

- processor failure rate is λ fph;
  its repair rate is μ rph
- voter failure rate is ν fph;
  its repair rate is δ rph
- rate matrix: e.g., $\mathbf{R}((3,1),(2,1)) = 3\lambda$
- exit rates: e.g., $r((3,1)) = 3\lambda + \nu$
- probability matrix: e.g.,

$$\mathbf{P}((3,1),(2,1)) = \frac{3\lambda}{3\lambda+\nu}$$

Transient probabilities

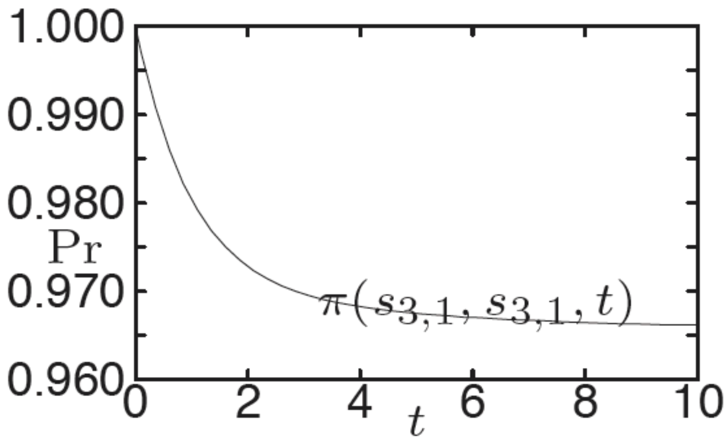

$p_{s_{3,1}}(t)$ for $t \leqslant 10$ hours

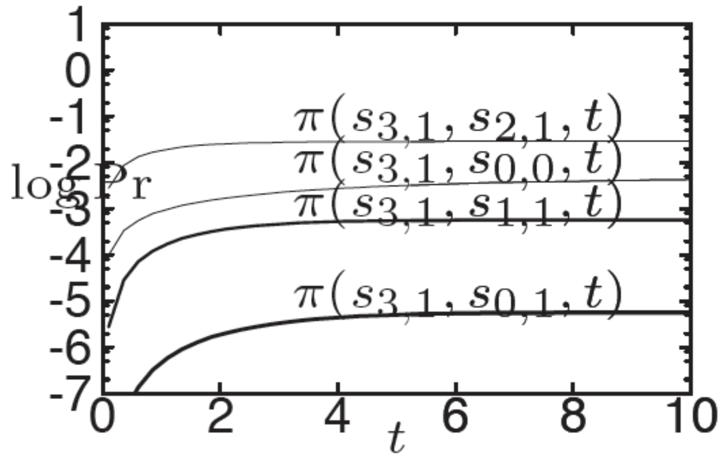

$p(t)$ for $t \leqslant 10$ hours (log-scale)

$\lambda = 0.01$ fph, $\nu = 0.001$ fph

$\mu = 1$ rph and $\delta = 0.2$ rph

(© book by B.R. Haverkort)

Steady-state distribution of a CTMC

For any finite and strongly connected CTMC it holds:

$$p_s = \lim_{t \to \infty} p_s(t) \quad \Leftrightarrow \quad \lim_{t \to \infty} p'_s(t) = 0 \quad \Leftrightarrow \quad \lim_{t \to \infty} p_s(t) \cdot (\mathbf{R} - \mathbf{r}) = 0$$

$\boxed{\begin{array}{c} \text{Steady-state probability vector } \underline{p} = (p_{s_1}, \ldots, p_{s_k}) \text{ satisfies:} \\[1em] \underline{p} \cdot (\mathbf{R} - \mathbf{r}) \; = \; 0 \quad \text{where} \quad \sum_{s \in S} p_s = 1 \end{array}}$

Steady-state distribution

s	$s_{3,1}$	$s_{2,1}$	$s_{1,1}$	$s_{0,1}$	$s_{0,0}$
$p(s)$	$9.655 \cdot 10^{-1}$	$2.893 \cdot 10^{-2}$	$5.781 \cdot 10^{-4}$	$5.775 \cdot 10^{-6}$	$4.975 \cdot 10^{-3}$

The probability of $\geqslant$ two processors and the voter are up

once the CTMC has reached an equilibrium is 0.9655+0.02893 $\approx$ 0.993

$\lambda = 0.01$ fph, $\nu = 0.001$ fph

$\mu = 1$ rph and $\delta = 0.2$ rph

Computing transient probabilities

- Transient probability vector $\underline{p}(t) = (p_{s_1}(t), \ldots, p_{s_k}(t))$ satisfies:

$$\underline{p}'(t) = \underline{p}(t) \cdot (\mathbf{R}{-}\mathbf{r}) \quad \text{given} \quad \underline{p}(0)$$

- Solution using Taylor-Maclaurin expansion:

$$\underline{p}(t) \;=\; \underline{p}(0) \cdot e^{(\mathbf{R}-\mathbf{r}) \cdot t} \;=\; \underline{p}(0) \cdot \sum_{i=0}^{\infty} \frac{((\mathbf{R}{-}\mathbf{r}) \cdot t)^i}{i!}$$

- Main problems: infinite summation + numerical instability due to
  - non-sparsity of $(\mathbf{R}{-}\mathbf{r})^i$ and presence positive and negative entries

Uniform CTMCs

- A CTMC is uniform if $r(s) = r$ for all s for some $r \in \mathbb{R}_{>0}$
- Any CTMC can be changed into a weak bisimilar uniform CTMC
- Let $r \in \mathbb{R}_{>0}$ such that $r \geqslant \mathit{max}_{s \in S}\, r(s)$
  - $\frac{1}{r}$ is at most the shortest mean residence time in CTMC $\mathcal{C}$
- Then $u(r, \mathcal{C}) = (S, \overline{\mathbf{P}}, \overline{r}, L)$ with $\overline{r}(s) = r$ for any s, and:

$$\overline{\mathbf{P}}(s, s') = \frac{r(s)}{r} \cdot \mathbf{P}(s, s') \text{ if } s' \neq s \quad \text{and} \quad \overline{\mathbf{P}}(s, s) = \frac{r(s)}{r} \cdot \mathbf{P}(s, s) + 1 - \frac{r(s)}{r}$$

Uniformization

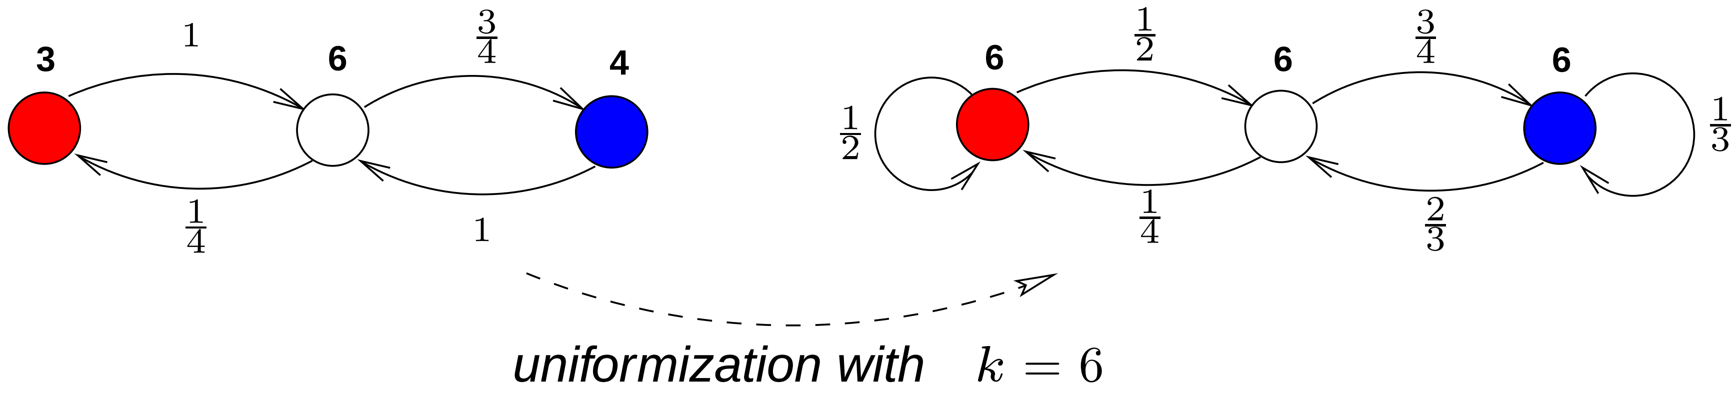

all state transitions in CTMC $u(r, \mathcal{C})$ occur at an average pace of r per time unit

Computing transient probabilities

- Now: $\underline{p}(t) = \underline{p}(0) \cdot e^{r \cdot (\overline{\mathbf{P}} - \mathbf{I})t} = \underline{p}(0) \cdot e^{-rt} \cdot e^{r \cdot t \cdot \overline{\mathbf{P}}} = \sum_{i=0}^{\infty} \underbrace{e^{-r \cdot t} \frac{(r \cdot t)^i}{i!}}_{\text{Poisson prob.}} \cdot \overline{\mathbf{P}}^i$
- Summation can be truncated *a priori* for a given error bound $\varepsilon > 0$:

$$\left\| \sum_{i=0}^{\infty} e^{-rt} \frac{(rt)^i}{i!} \cdot \underline{p}(i) - \sum_{i=0}^{k_\varepsilon} e^{-rt} \frac{(rt)^i}{i!} \cdot \underline{p}(i) \right\| = \left\| \sum_{i=k_\varepsilon+1}^{\infty} e^{-rt} \frac{(rt)^i}{i!} \cdot \underline{p}(i) \right\|$$

- Choose k_ε minimal s.t.: $\sum_{i=k_{\varepsilon+1}}^{\infty} e^{-rt} \frac{(rt)^i}{i!} = 1 - \sum_{i=0}^{k_\varepsilon} e^{-rt} \frac{(rt)^i}{i!} \leqslant \varepsilon$

Transient probabilities: example

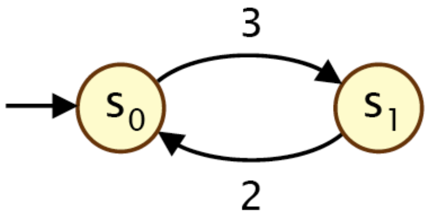

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \underline{r} = \begin{bmatrix} 3 \\ 2 \end{bmatrix} \text{ and } \overline{\mathbf{P}}_3 = \begin{bmatrix} 0 & 1 \\ \frac{2}{3} & \frac{1}{3} \end{bmatrix}$$

Let initial distribution $\underline{p}(0) = (1, 0)$, and time bound $t{=}1$.

Then:

$$\underline{p}(0) \cdot \sum_{i=0}^{\infty} e^{-3} \frac{3^i}{i!} \cdot \overline{\mathbf{P}}^i$$

$$= (1, 0) \cdot e^{-3} \tfrac{1}{0!} \cdot \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + (1, 0) \cdot e^{-3} \tfrac{3}{1!} \cdot \begin{bmatrix} 0 & 1 \\ \frac{2}{3} & \frac{1}{3} \end{bmatrix}$$

$$+ (1, 0) \cdot e^{-3} \tfrac{9}{2!} \cdot \begin{bmatrix} 0 & 1 \\ \frac{2}{3} & \frac{1}{3} \end{bmatrix}^2 + \ldots\ldots$$

$$\approx (0.404043, 0.595957)$$

CTMC paths

- An infinite path σ in a CTMC $\mathcal{C} = (S, \mathbf{P}, r, L)$ is of the form:

$$\sigma \;=\; s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \xrightarrow{t_2} s_3 \ldots\ldots$$

  with s_i is a state in S, $t_i \in \mathbb{R}_{>0}$ is a duration, and $\mathbf{P}(s_i, s_{i+1}) > 0$.

- A Borel space on infinite paths exists (cylinder construction)
  - reachability, timed reachability, and ω-regular properties are measurable

- A path is Zeno if $\sum_i t_i$ is converging

- Theorem: the probability of the set of Zeno paths in any CTMC is 0

Summarizing

- Negative exponential distribution
  - suitable for many practical phenomena
  - nice mathematical properties
- Continuous-time Markov chains
  - Kripke structures with exponential state residence times
  - used in many different fields, e.g., performance, biology, . . .
- Performance measures
  - transient probability vector: where is a CTMC at time t?
  - steady-state probability vector: where is a CTMC on the long run?

Model Checking
Continuous-Time Markov Chains

Joost-Pieter Katoen

Software Modeling and Verification Group

RWTH Aachen University

associated to University of Twente, Formal Methods and Tools

Lecture at MOVEP Summerschool, July 1, 2010

Content of this lecture

- Continuous Stochastic Logic
  - syntax, semantics, examples
- CSL model checking
  - basic algorithms and complexity
- Bisimulation
  - definition, minimization algorithm, examples
- Priced continuous-time Markov chains
  - motivation, definition, some properties

Content of this lecture

⇒ Continuous Stochastic Logic

- syntax, semantics, examples

- CSL model checking
  - basic algorithms and complexity
- Bisimulation
  - definition, minimization algorithm, examples
- Priced continuous-time Markov chains
  - motivation, definition, some properties

Continuous-time Markov chain

A *continuous-time Markov chain* (CTMC) is a tuple $(S, \mathbf{P}, r, L)$ where:

- S is a countable (today: finite) set of *states*
- $\mathbf{P} : S \times S \rightarrow [0, 1]$, a *stochastic matrix*
  - $\mathbf{P}(s, s')$ is one-step probability of going from state s to state s'
  - s is called *absorbing* iff $\mathbf{P}(s, s) = 1$
- $r : S \rightarrow \mathbb{R}_{>0}$, the *exit-rate function*
  - $r(s)$ is the rate of exponential distribution of residence time in state s

CTMC paths

- An infinite path σ in a CTMC $\mathcal{C} = (S, \mathbf{P}, r, L)$ is of the form:

$$\sigma \;=\; s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \xrightarrow{t_2} s_3 \ldots\ldots$$

  with s_i is a state in S, $t_i \in \mathbb{R}_{>0}$ is a duration, and $\mathbf{P}(s_i, s_{i+1}) > 0$.

- A Borel space on infinite paths exists (cylinder construction)
  - reachability, timed reachability, and ω-regular properties are measurable

- Let *Paths*(s) denote the set of infinite path starting in state s

Reachability probabilities

- Let $\mathcal{C} = (S, \mathbf{P}, r, L)$ be a finite CTMC and $G \subseteq S$ a set of states
- Let $\Diamond G$ be the set of infinite paths in $\mathcal{C}$ reaching a state in G
- Question: what is the probability of $\Diamond G$ when starting from s?
  - what is the probability mass of all infinite paths from s that eventually hit G?
- As state residence times are not relevant for $\Diamond G$, this is simple

Probabilistic reachability

- $\Pr(s, \Diamond G)$ is the least solution of the set of linear equations:

$$\Pr(s, \Diamond G) = \begin{cases} 1 & \text{if } s \in G \\ \sum_{s' \in S} \mathbf{P}(s, s') \cdot \Pr(s', \Diamond G) & \text{otherwise} \end{cases}$$

- Unique solution by pre-computing *Sat*$(\forall \Diamond G)$ and *Sat*$(\exists \Diamond G)$
  - this is a standard graph analysis (as in CTL model checking)

- This is the same as in the first lecture this morning

Continuous stochastic logic (CSL)

- CSL equips the until-operator with a time interval:
  - let interval $I \subseteq \mathbb{R}_{\geqslant 0}$ with rational bounds, e.g., $I = [0, 17]$
  - $\Phi \, \mathsf{U}^I \Psi$ asserts that a Ψ-state can be reached via Φ-states

    ... while reaching the Ψ-state at some time $t \in I$
- CSL contains a probabilistic operator $\mathbb{P}$ with arguments
  - a path formula, e.g., *good* $\mathsf{U}^{[0,12]}$*bad*, and
  - a probability interval $J \subseteq [0, 1]$ with rational bounds, e.g., $J = [0, \frac{1}{2}]$
- CSL contains a long-run operator $\mathbb{L}$ with arguments
  - a state formula, e.g., $a \wedge b$ or $\mathbb{P}_{=1}(\Diamond \Phi)$, and
  - a probability interval $J \subseteq [0, 1]$ with rational bounds

The branching-time logic CSL

- For $a \in AP$, $J \subseteq [0,1]$ and $I \subseteq \mathbb{R}_{\geqslant 0}$ intervals with rational bounds:

$$\boxed{\begin{array}{c} \Phi \;::=\; a \;\big|\; \neg\Phi \;\big|\; \Phi \wedge \Phi \;\big|\; \mathbb{L}_J(\Phi) \;\big|\; \mathbb{P}_J(\varphi) \\[2em] \varphi \;::=\; \Phi \,\mathsf{U}\, \Phi \;\big|\; \Phi \,\mathsf{U}^I\, \Phi \end{array}}$$

- $s_0 t_0 s_1 t_1 s_2 \ldots \models \Phi \,\mathsf{U}^I\, \Psi$ if Ψ is reached at $t \in I$ and prior to t, Φ holds
- $s \models \mathbb{P}_J(\varphi)$ if the probability of the set of φ-paths starting in s lies in J
- $s \models \mathbb{L}_J(\Phi)$ if starting from s, the probability of being in Φ on the long run lies in J

Derived operators

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

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

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

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

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

Timed reachability formulas

- In $\geqslant$ 92% of the cases, a goal state is legally reached within 3.1 sec:

$$\mathbb{P}_{\geqslant 0.92}\left(\textit{legal}\ \mathsf{U}^{\leqslant 3.1}\ \textit{goal}\right)$$

- Almost surely stay in a legal state for at least 10 sec:

$$\mathbb{P}_{=1}\left(\Box^{\leqslant 10}\ \textit{legal}\right)$$

- Combining these two constraints:

$$\mathbb{P}_{\geqslant 0.92}\left(\textit{legal}\ \mathsf{U}^{\leqslant 3.1}\ \mathbb{P}_{=1}\left(\Box^{\leqslant 10}\ \textit{legal}\right)\right)$$

Long-run formulas

- The long-run probability of being in a safe state is at most 0.00001:

$$\mathbb{L}_{\leqslant 10^{-5}}\,(\textit{safe})$$

- On the long run, with at least "five nine" likelihood almost surely a goal state can be reached within one sec.:

$$\mathbb{L}_{\geqslant 0.99999}\,\big(\mathbb{P}_{=1}(\Diamond^{\leqslant 1}\textit{goal})\big)$$

- The probability to reach a state that in the long run guarantees more than five-nine safety exceeds $\frac{1}{2}$:

$$\mathbb{P}_{>0.5}\,(\Diamond\,\mathbb{L}_{>0.99999}(\textit{safe}))$$

CSL semantics

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

$s \models a$	iff $a \in L(s)$
$s \models \neg \Phi$	iff not $(s \models \Phi)$
$s \models \Phi \wedge \Psi$	iff $(s \models \Phi)$ and $(s \models \Psi)$
$s \models \mathbb{L}_J(\Phi)$	iff $\lim_{t \to \infty} \Pr\{ \sigma \in \textit{Paths}(s) \mid \sigma@t \models \Phi \} \in J$
$s \models \mathbb{P}_J(\varphi)$	iff $\Pr\{ \sigma \in \textit{Paths}(s) \mid \sigma \models \varphi \} \in J$
$\sigma \models \Phi \, \mathsf{U}^I \, \Psi$	iff $\exists t \in I. \; ((\forall t' \in [0, t). \, \sigma@t' \models \Phi) \wedge \sigma@t \models \Psi)$

where $\sigma@t$ is the state along σ that is occupied at time t

Content of this lecture

- Continuous Stochastic Logic
  - syntax, semantics, examples

⇒ CSL model checking
  - basic algorithms and complexity

- Bisimulation
  - definition, minimization algorithm, examples

- Priced continuous-time Markov chains
  - motivation, definition, some properties

CSL model checking

- Let $\mathcal{C}$ be a finite CTMC and Φ a CSL formula.
- Problem: determine the states in $\mathcal{C}$ satisfying Φ
- Determine *Sat*(Φ) by a recursive descent over parse tree of Φ
- For the propositional fragment ($\neg, \wedge, a$): do as for CTL
- How to check formulas of the form $\mathbb{P}_J(\varphi)$?
  - φ is an until-formula: do as for PCTL, i.e., linear equation system
  - φ is a time-bounded until-formula: integral equation system
- How to check formulas of the form $\mathbb{L}_J(\Psi)$?
  - graph analysis + solving linear equation system(s)

Model-checking the long-run operator

- For a strongly-connected CTMC:

$$s \in \mathit{Sat}(\mathbb{L}_J(\Phi)) \quad \text{iff} \quad \sum_{s' \in \mathit{Sat}(\Phi)} p(s') \in J$$

$\Longrightarrow$ this boils down to a standard steady-state analysis

- For an arbitrary CTMC:
  - determine the *bottom* strongly-connected components (BSCCs)
  - for BSCC B determine the steady-state probability of a Φ-state
  - compute the probability to reach BSCC B from state s

$$s \in \mathit{Sat}(\mathbb{L}_J(\Phi)) \quad \text{iff} \quad \sum_{B} \left(\Pr\{\, s \models \Diamond B \,\} \cdot \sum_{s' \in B \cap \mathit{Sat}(\Phi)} p^{B}(s') \right) \in J$$

Verifying long-run properties: an example

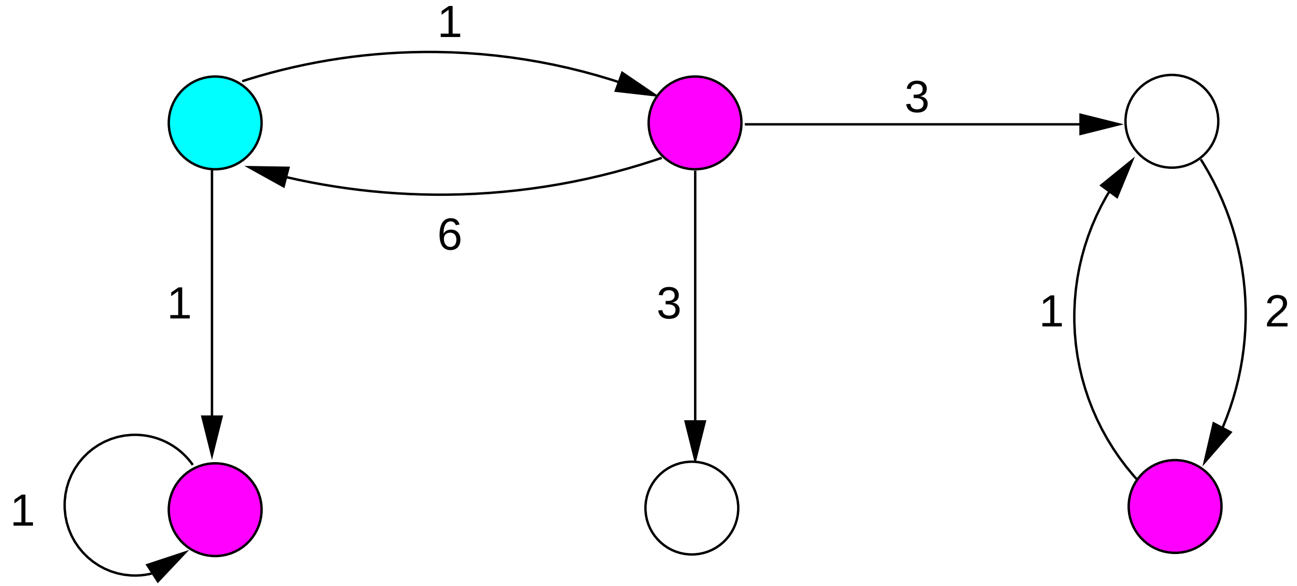

determine the bottom strongly-connected components

Verifying long-run properties: an example

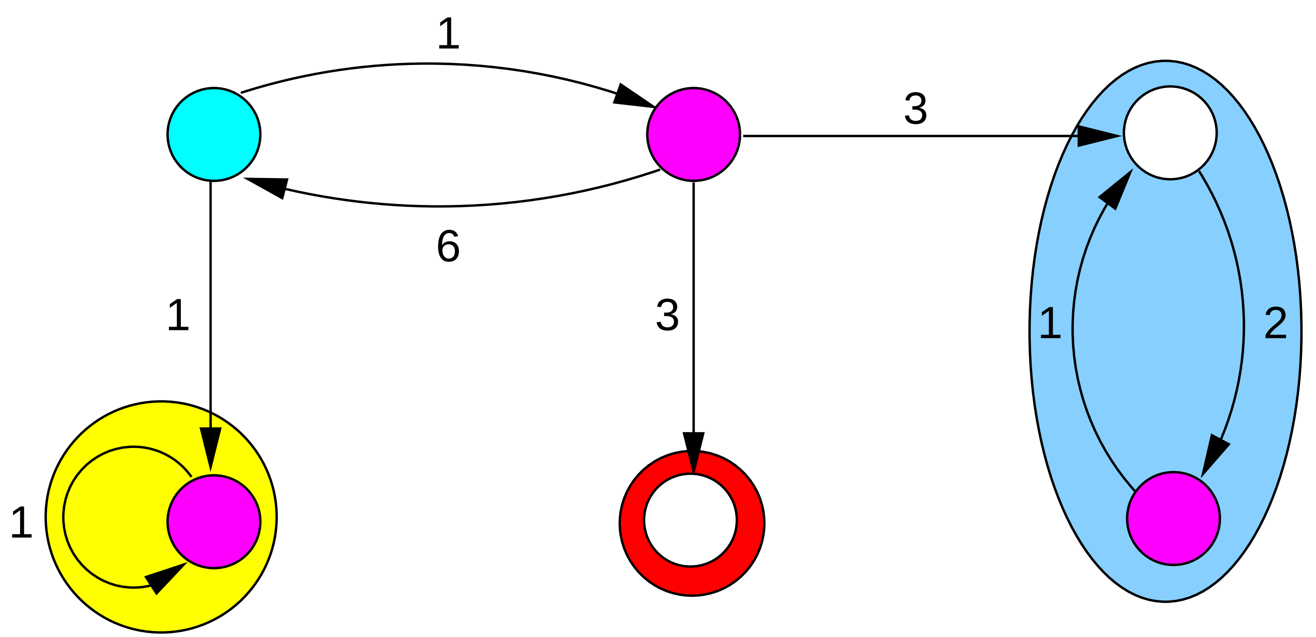

$$s \models \mathbb{L}_{>\frac{3}{4}}(\textit{magenta}) \quad \text{iff} \quad \Pr\{s \models \Diamond at_{yellow}\} \cdot p^{yellow}(\textit{magenta}) + \Pr\{s \models \Diamond at_{blue}\} \cdot p^{blue}(\textit{magenta}) > \frac{3}{4}$$

Verifying long-run properties: an example

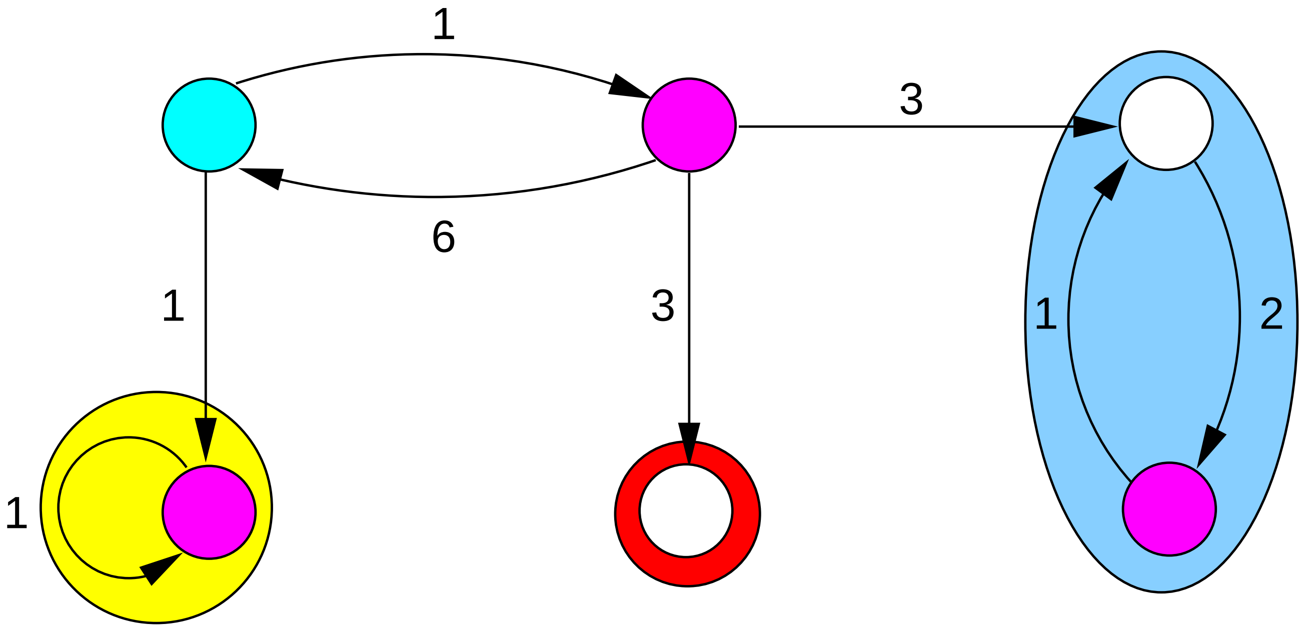

$$s \models \mathbb{L}_{>\frac{3}{4}}(\textit{magenta}) \quad \text{iff} \quad \Pr\{s \models \Diamond at_{yellow}\} \cdot \underbrace{p^{yellow}(\textit{magenta})}_{=1}$$

$$+ \Pr\{s \models \Diamond at_{blue}\} \cdot \underbrace{p^{blue}(\textit{magenta})}_{=\frac{2}{3}} > \frac{3}{4}$$

Verifying long-run properties: an example

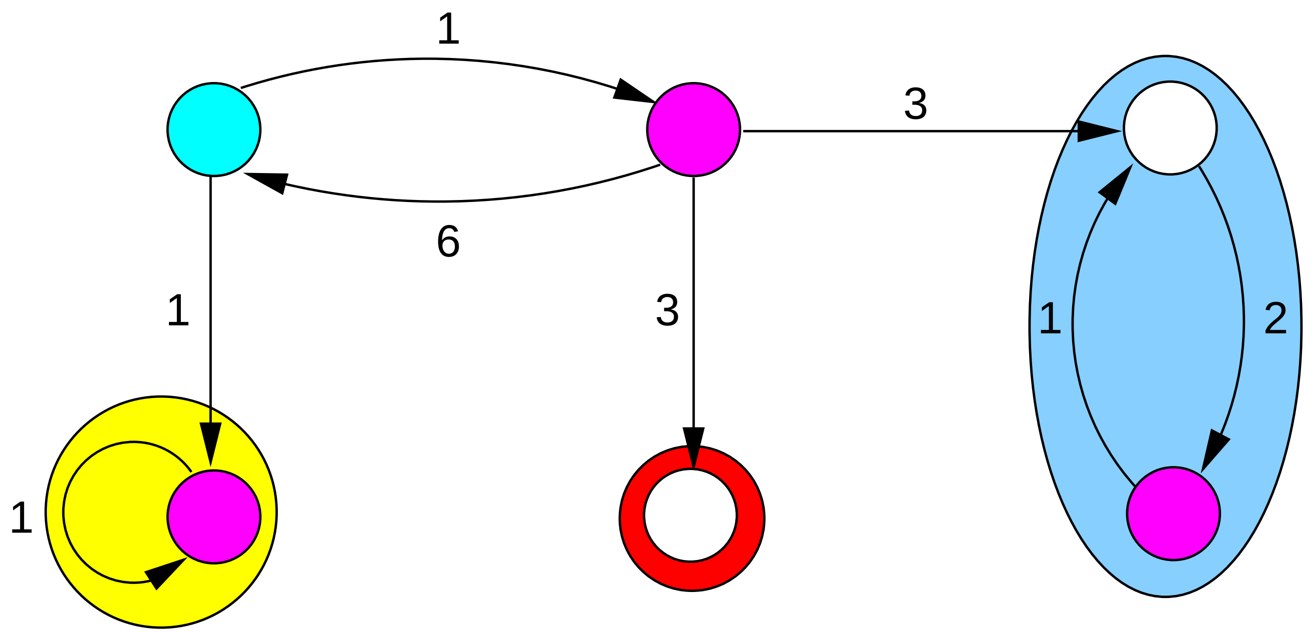

$s \models \mathbb{L}_{>\frac{3}{4}}(\textit{magenta})$ iff $\Pr\{s \models \Diamond at_{yellow}\} + \frac{2}{3}\Pr\{s \models \Diamond at_{blue}\} > \frac{3}{4}$

Verifying long-run properties: an example

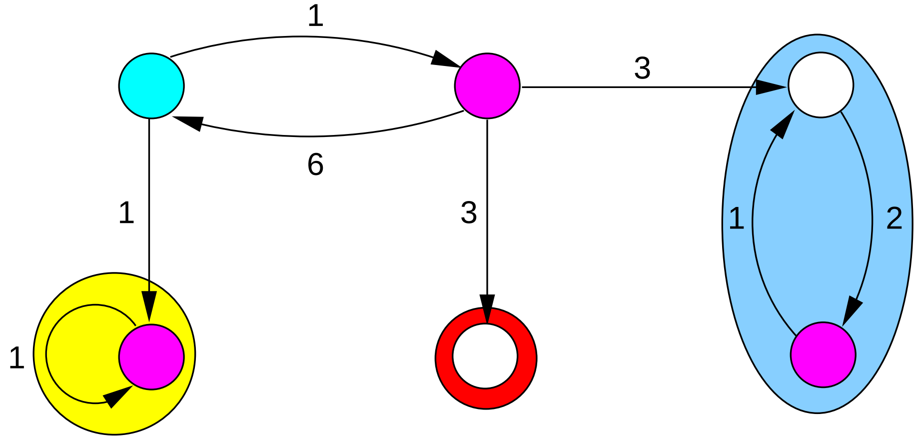

$$s \models \mathbb{L}_{>\frac{3}{4}}(\textit{magenta}) \quad \text{iff} \quad \Pr\{s \models \Diamond at_{yellow}\} + \tfrac{2}{3}\Pr\{s \models \Diamond at_{blue}\} > \tfrac{3}{4}$$

$$\begin{aligned}
\Pr\{s \models \Diamond at_{yellow}\} &= \tfrac{1}{2} + \tfrac{1}{2}\Pr\{s' \models \Diamond at_{yellow}\} \\
\Pr\{s' \models \Diamond at_{yellow}\} &= \tfrac{1}{2}\Pr\{s \models \Diamond at_{yellow}\}
\end{aligned}$$

$$\Rightarrow \Pr\{s \models \Diamond at_{yellow}\} = \tfrac{1}{2}\sum_{k=0}^{\infty}\left(\tfrac{1}{4}\right)^k = \tfrac{2}{3}$$

Verifying long-run properties: an example

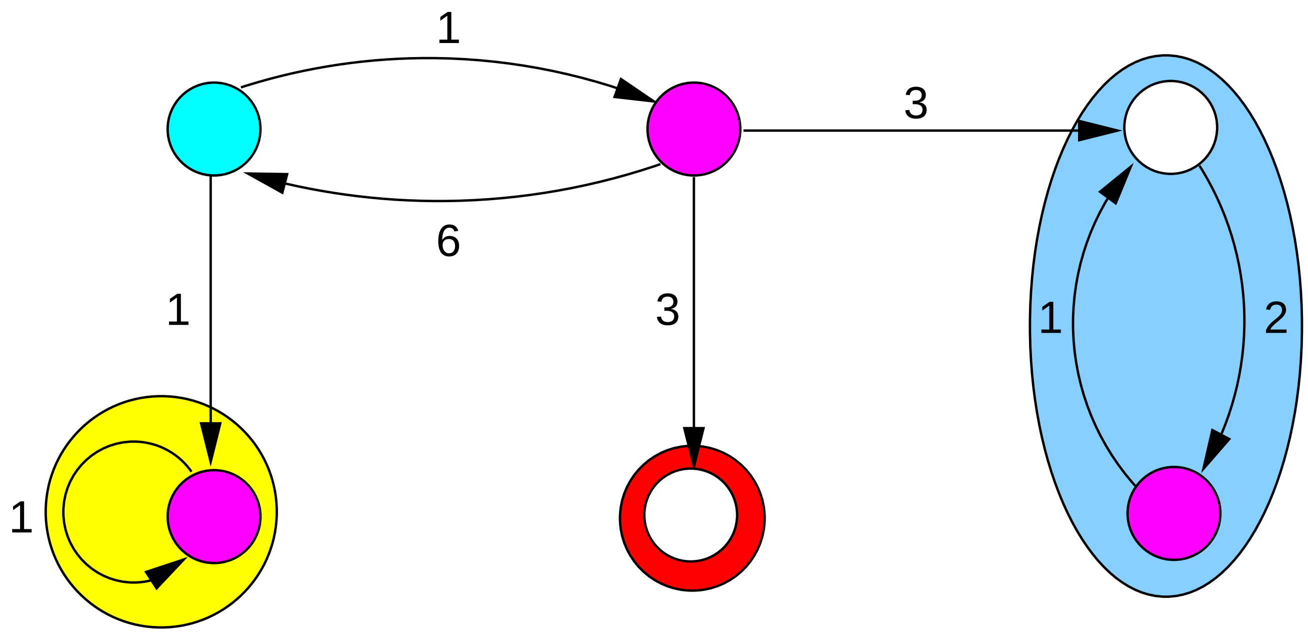

$s \models \mathbb{L}_{>\frac{3}{4}}(\textit{magenta})$ iff $\underbrace{\Pr\{s \models \Diamond at_{yellow}\}}_{\frac{2}{3}} + \frac{2}{3}\underbrace{\Pr\{s \models \Diamond at_{blue}\}}_{\frac{1}{6}} > \frac{3}{4}$

Verifying long-run properties: an example

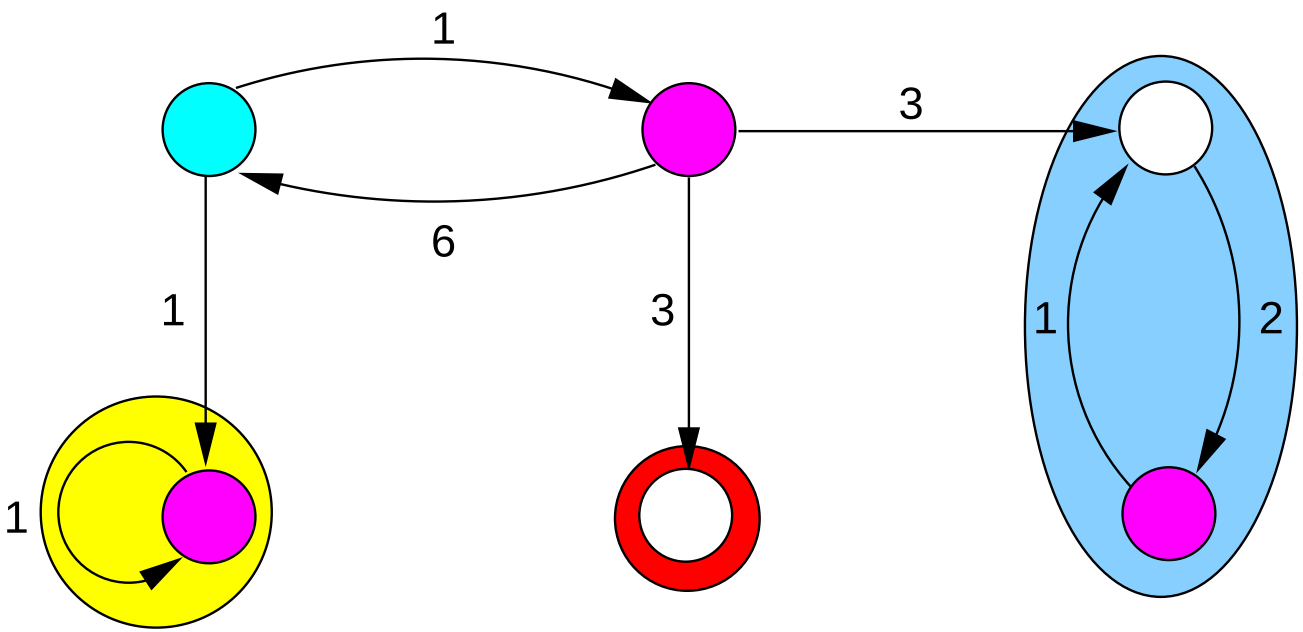

$$s \models \mathbb{L}_{>\frac{3}{4}}(\textit{magenta}) \quad \text{iff} \qquad \frac{2}{3} + \frac{2}{3}\cdot\frac{1}{6} > \frac{3}{4}$$

Verifying long-run properties: an example

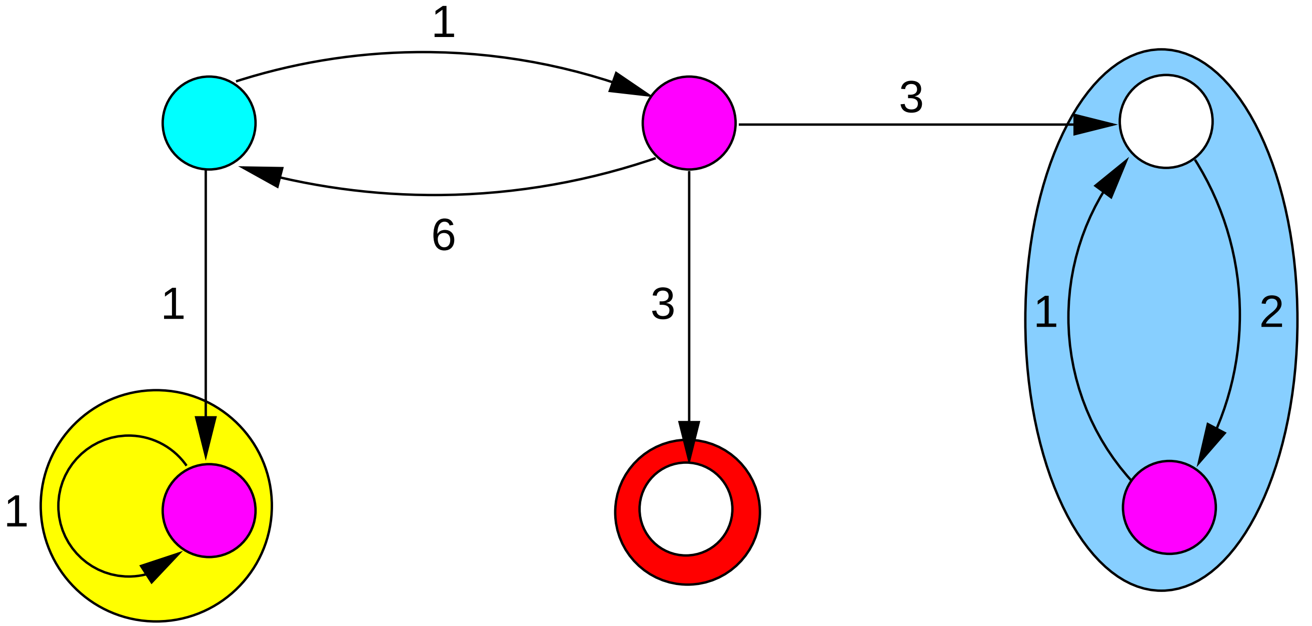

Thus: $s \models \mathbb{L}_{>\frac{3}{4}}(\textit{magenta})$ as $\underbrace{\frac{2}{3} + \frac{2}{3} \cdot \frac{1}{6}}_{\frac{7}{9}} > \frac{3}{4}$

Time-bounded reachability

- $s \models \mathbb{P}_J\left(\Phi \, \mathsf{U}^I \, \Psi\right)$ if and only if $\Pr\{s \models \Phi \, \mathsf{U}^I \, \Psi\} \in J$
- For $I = [0, t]$, $\Pr\{s \models \Phi \, \mathsf{U}^{\leqslant t} \Psi\}$ is the least solution of:
  - 1 if $s \in$ *Sat*(Ψ)
  - if $s \in$ *Sat*(Φ) $-$ *Sat*(Ψ):

$$\int_0^t \sum_{s' \in S} \underbrace{\mathbf{R}(s, s') \cdot e^{-r(s) \cdot x}}_{\substack{\text{probability to move to} \\ \text{state } s' \text{ at time } x}} \cdot \underbrace{\Pr\{s' \models \Phi \, \mathsf{U}^{\leqslant t-x} \, \Psi\}}_{\substack{\text{probability to fulfill } \Phi \, \mathsf{U} \, \Psi \\ \text{before time } t-x \text{ from } s'}} \, dx$$

  - 0 otherwise

Reduction to transient analysis

- For an arbitrary CTMC $\mathcal{C}$ and property $\varphi = \Phi \, \mathsf{U}^{\leqslant t} \, \Psi$ we have:
  - φ is fulfilled once a Ψ-state is reached before t along a Φ-path
  - φ is violated once a $\neg \, (\Phi \, \vee \, \Psi)$-state is visited before t
- This suggests to transform the CTMC $\mathcal{C}$ as follows:
  - make all Ψ-states and all $\neg \, (\Phi \, \vee \, \Psi)$-states absorbing
- Theorem: $\underbrace{s \models \mathbb{P}_J(\Phi \, \mathsf{U}^{\leqslant t} \, \Psi)}_{\text{in } \mathcal{C}}$ iff $\underbrace{s \models \mathbb{P}_J(\Diamond^{=t} \, \Psi)}_{\text{in } \mathcal{C}'}$
- Then it follows: $s \models_{\mathcal{C}'} \mathbb{P}_J(\Diamond^{=t} \, \Psi)$ iff $\underbrace{\sum_{s' \models \Psi} p_{s'}(t)}_{\text{transient probs in } \mathcal{C}'} \in J$

Example: TMR with $\mathbb{P}_J((\textit{green} \vee \textit{blue})\, \mathsf{U}^{[0,3]}\, \textit{red})$

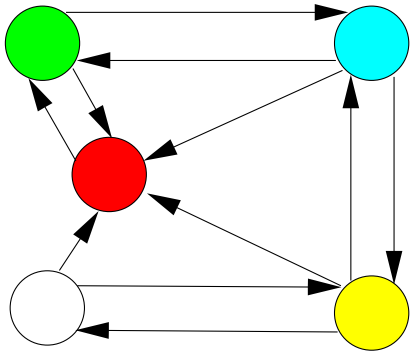

transformation

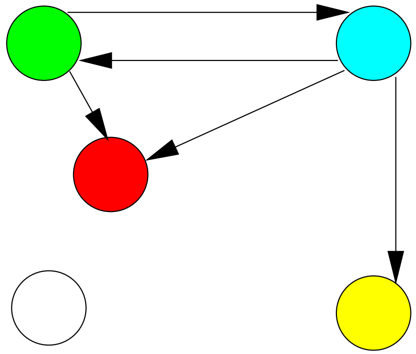

uniformisation

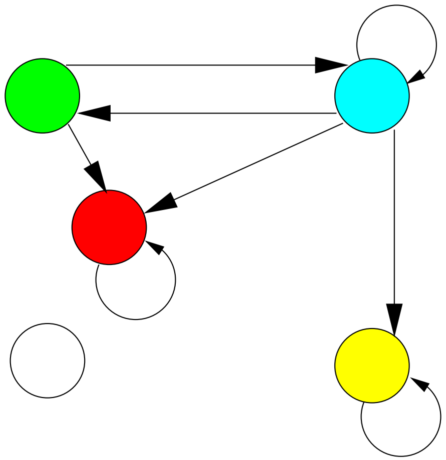

recursive computation

like PCTL

bounded until

Interval-bounded reachability

- For any path σ that fulfills $\Phi \, \mathsf{U}^{[t,t']} \, \Psi$ with $0 < t \leqslant t'$:
  - Φ holds continuously up to time t, and
  - the suffix of σ starting at time t fulfills $\Phi \, \mathsf{U}^{[0,t'-t]} \, \Psi$

- Approach: divide the problem into two:

$$\underbrace{\sum_{s' \models \Phi} p^{\mathcal{C}'}(s, s', t)}_{\text{check } \Box^{[0,t]} \, \Phi} \cdot \underbrace{\sum_{s'' \models \Psi} p^{\mathcal{C}''}(s', s'', t'-t)}_{\substack{\text{check } \Phi \, \mathsf{U}^{[0,t'-t]} \, \Psi \\ \text{with starting distribution } \underline{p}^{\mathcal{C}'}(t)}}$$

  - where CTMC $\mathcal{C}'$ equals $\mathcal{C}$ with all Φ-states absorbing
  - and CTMC $\mathcal{C}''$ equals $\mathcal{C}$ with all Ψ and $\neg(\Phi \vee \Psi)$-states absorbing

Verification times

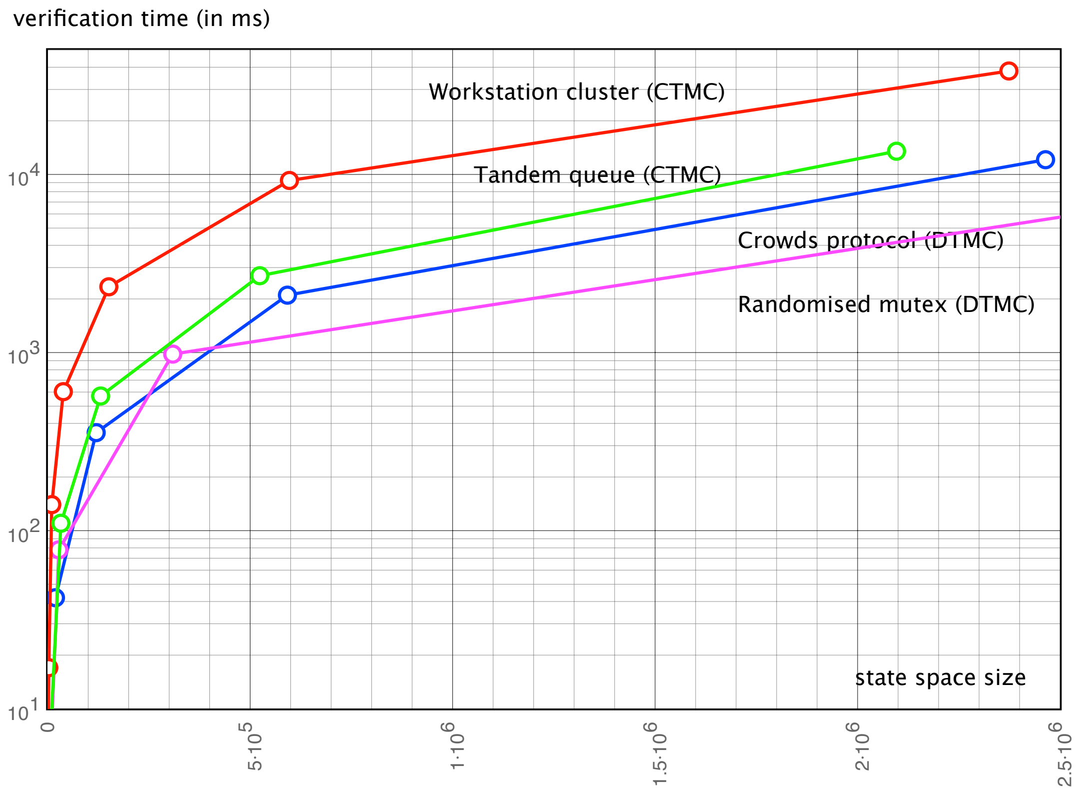

command-line tool MRMC ran on a Pentium 4, 2.66 GHz, 1 GB RAM laptop

Reachability probabilities

	Nondeterminism no	Nondeterminism yes
Reachability	linear equation system DTMC	linear programming MDP
Timed reachability	transient analysis CTMC	discretisation + linear programming CTMDP

Summary of CSL model checking

- Recursive descent over the parse tree of Φ
- Long-run operator: graph analysis + linear system(s) of equations
- Time-bounded until: CTMC transformation and uniformization
- Worst case time-complexity: $\mathcal{O}(|\Phi|\cdot(|\,\mathbf{R}\,|\cdot r\cdot t_{max} + |\,S\,|^{2.81}))$

  with $|\Phi|$ the length of Φ, uniformization rate r, t_{max} the largest time bound in Φ
- Tools:

  PRISM (symbolic), MRMC (explicit state), YMER (simulation), VESTA (simulation), . . .

Content of this lecture

- Continuous Stochastic Logic
  - syntax, semantics, examples
- CSL model checking
  - basic algorithms and complexity

$\Rightarrow$ Bisimulation
  - definition, minimization algorithm, examples
- Priced continuous-time Markov chains
  - motivation, definition, some properties

Probabilistic bisimulation

- Traditional LTL/CTL model checking: (Fisler & Vardi, 1998)
  - significant reductions in state space (upto logarithmic)
  - cost of bisimulation minimisation significantly exceeds model checking time
- Pros:
  - fully automated and efficient abstraction technique
  - enables compositional minimization
- Our interest:

does bisimulation minimization as pre-computation step
of probabilistic model checking pay off?

Probabilistic bisimulation

- Let $\mathcal{C} = (S, \mathbf{P}, r, L)$ be a CTMC and R an equivalence relation on S

- R is a probabilistic bisimulation on S if for any $(s, s') \in R$ it holds:

  1. $L(s) = L(s')$
  2. $r(s) = r(s')$
  3. $\mathbf{P}(s, C) = \mathbf{P}(s', C)$ for all $C \in S/R$, where $\mathbf{P}(s, C) = \sum_{u \in C} \mathbf{P}(s, u)$

  Note that the last two conditions together equal $\mathbf{R}(s, C) \; = \; \mathbf{R}(s', C)$.

- States s and s' are bisimilar, denoted $s \sim s'$, if:

  $\exists$ a probabilistic bisimulation R on S with $(s, s') \in R$

Example

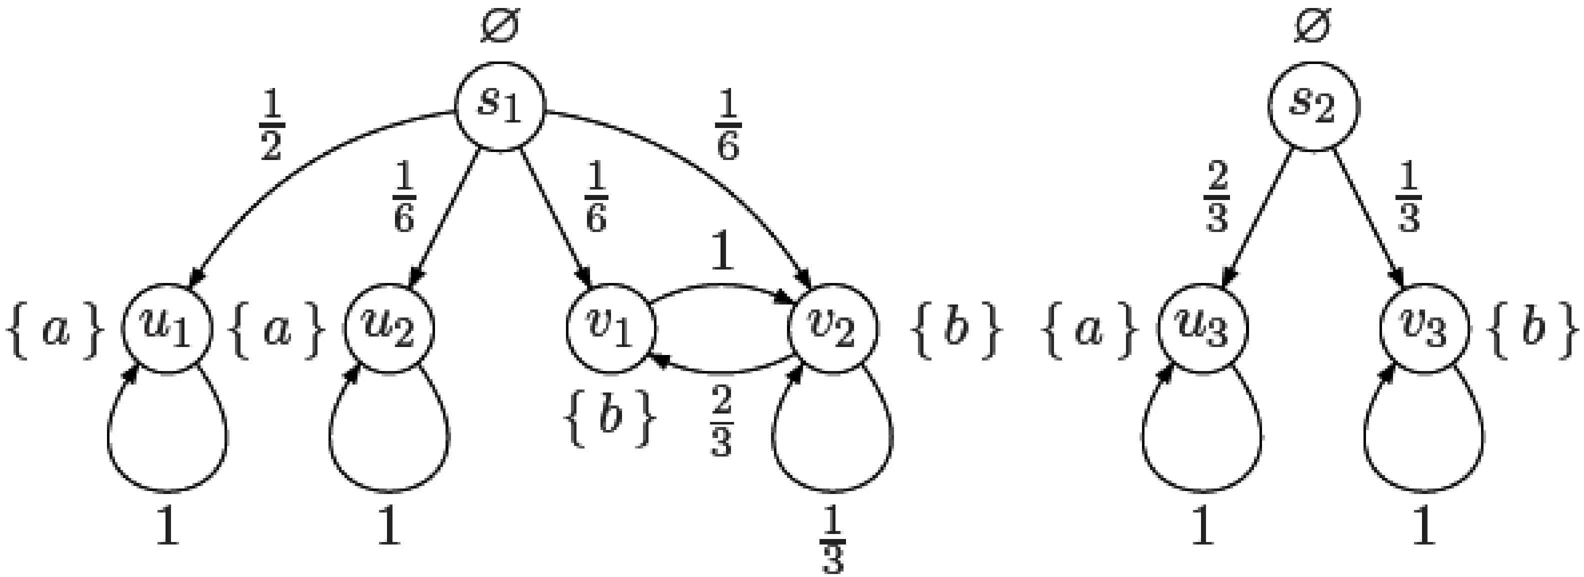

for simplicity, all states have the same exit rate (= uniform CTMC)

Quotient Markov chain

For $\mathcal{C} = (S, \mathbf{R}, L)$ and probabilistic bisimulation $\sim \subseteq S \times S$ let

$$\mathcal{C}/\sim = (S', \mathbf{R}', L'), \qquad \text{the quotient of } \mathcal{C} \text{ under } \sim$$

where

- $S' = S/\sim = \{ [s]_\sim \mid s \in S \}$ with $[s]_\sim = \{ s' \in S \mid s \sim s' \}$
- $\mathbf{R}' : S' \times S' \to [0, 1]$ is defined such that for each $s \in S$ and $C \in S$:

$$\mathbf{R}'([s]_\sim, C) = \mathbf{R}(s, C)$$

- $L'([s]_\sim) = L(s)$

it follows that $\mathcal{C} \sim \mathcal{C}/\sim$

Modelling a TMR system as a CTMC

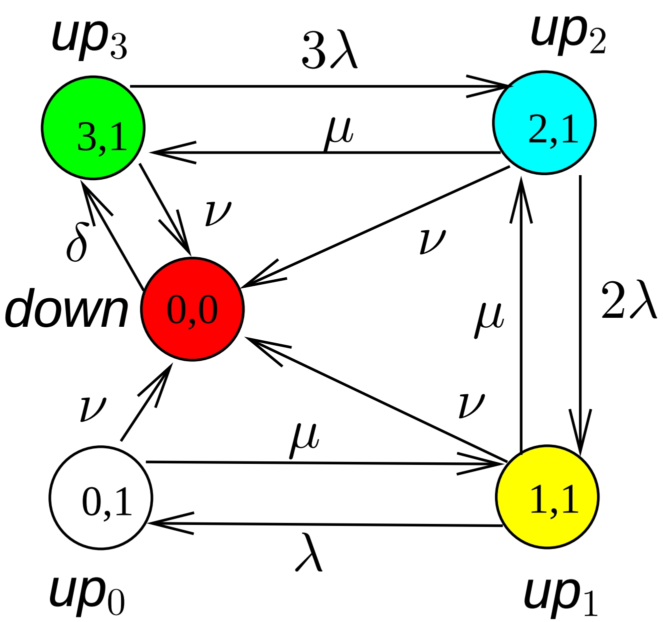

- processor failure rate is λ fph;
  its repair rate is μ rph
- voter failure rate is ν fph;
  its repair rate is δ rph
- rate matrix: e.g., $\mathbf{R}((3,1),(2,1)) = 3\lambda$
- exit rates: e.g., $r((3,1)) = 3\lambda+\nu$
- probability matrix: e.g.,

$$\mathbf{P}((3,1),(2,1)) = \frac{3\lambda}{3\lambda+\nu}$$

A bisimilar TMR model

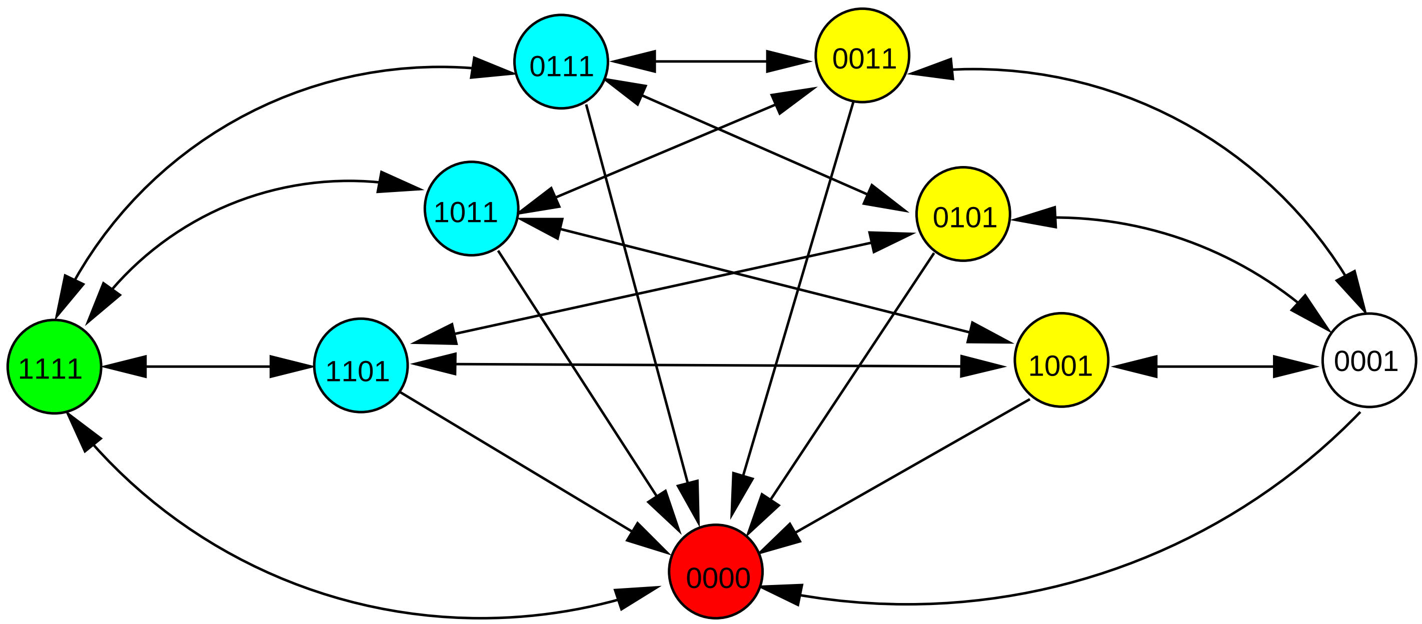

$$\mathbf{R}'([s]_{\sim_m}, C) = \mathbf{R}(s, C) = \sum_{s' \in C} \mathbf{R}(s, s')$$

Preservation of state probabilities

- Let $\mathcal{C} = (S, \mathbf{R}, L)$ be a CTMC with initial distribution $\underline{p}(0)$
- For any $C \in S_0/\sim$ we have:

$$\underline{p'}_C(t) \;=\; \sum_{s \in C} \underline{p}_s(t) \qquad \text{for any } t \geqslant 0$$

- If the steady-state distribution exists, then it follows:

$$\underline{p'}_C \;=\; \lim_{t \to \infty} \underline{p'}_C(t) \;=\; \lim_{t \to \infty} \sum_{s \in C} \underline{p}_s(t) \;=\; \sum_{s \in C} \underline{p}_s$$

Logical characterization

For any finite CTMC with states s and s':

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

The quotient under the coarsest bisimulation can be obtained by partition-refinement in time-complexity $\mathcal{O}(|\mathbf{R}| \cdot \log |S|)$

Craps

- Roll two dice and bet on outcome
- Come-out roll (“pass line” wager):
  - outcome 7 or 11: win
  - outcome 2, 3, and 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

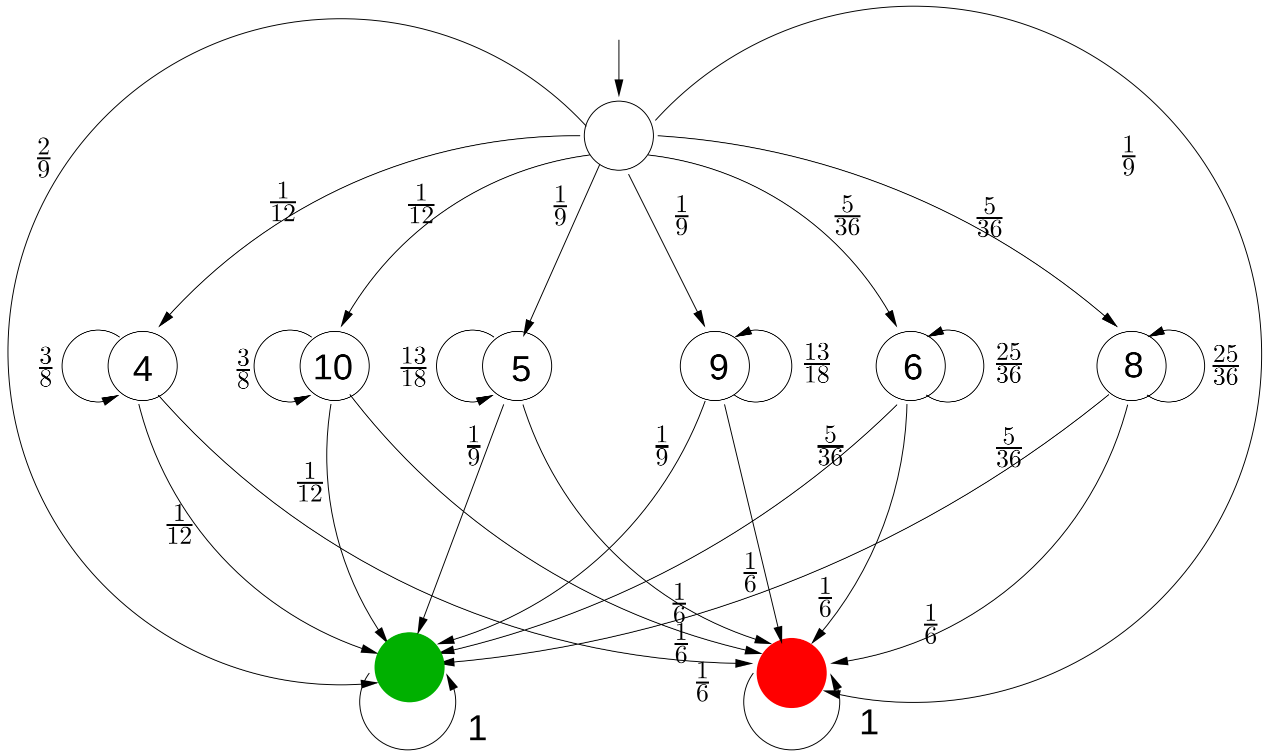

Minimizing Craps

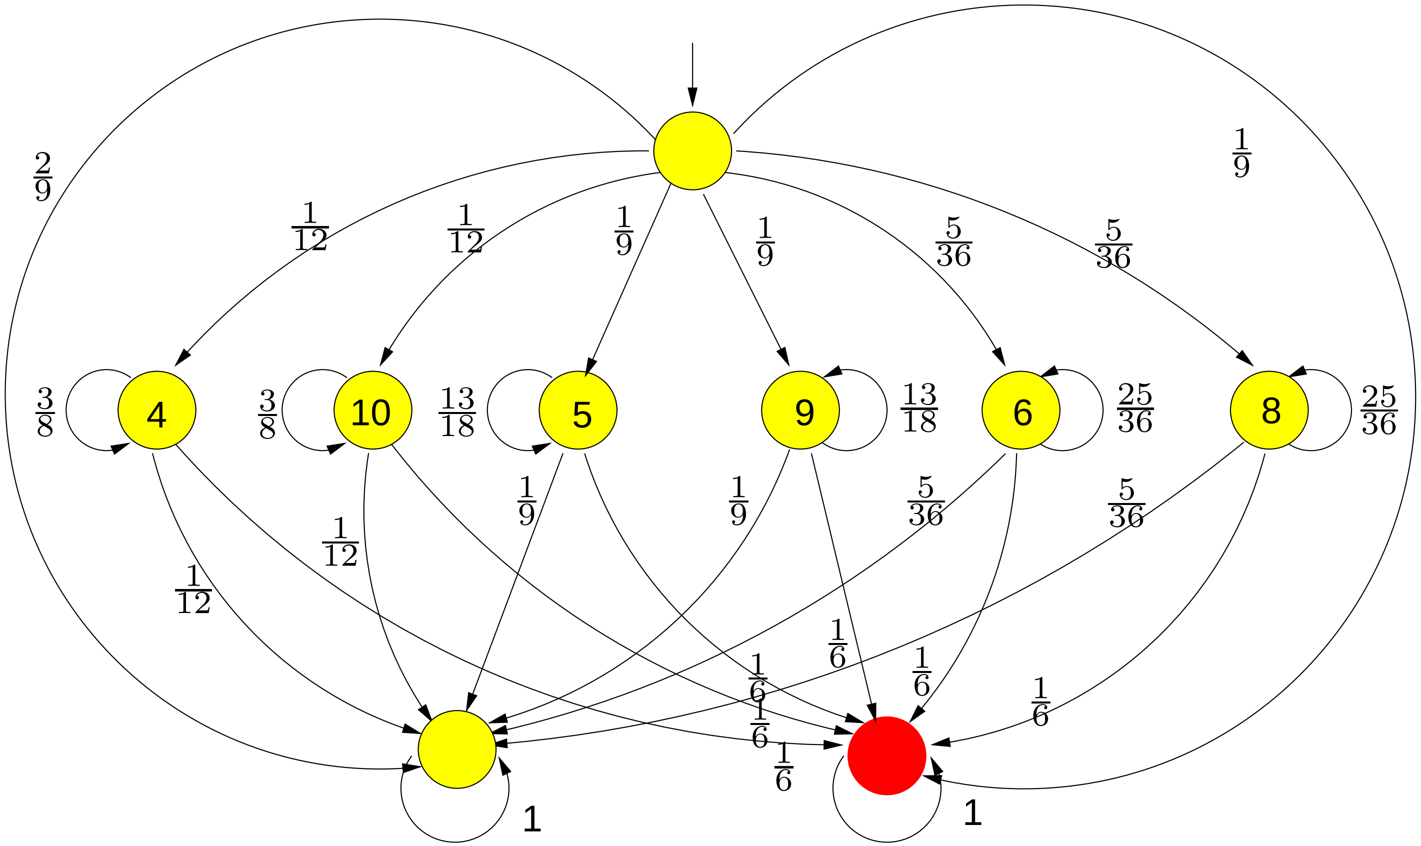

initial partitioning for the atomic propositions $AP = \{ \mathit{loss} \}$

A first refinement

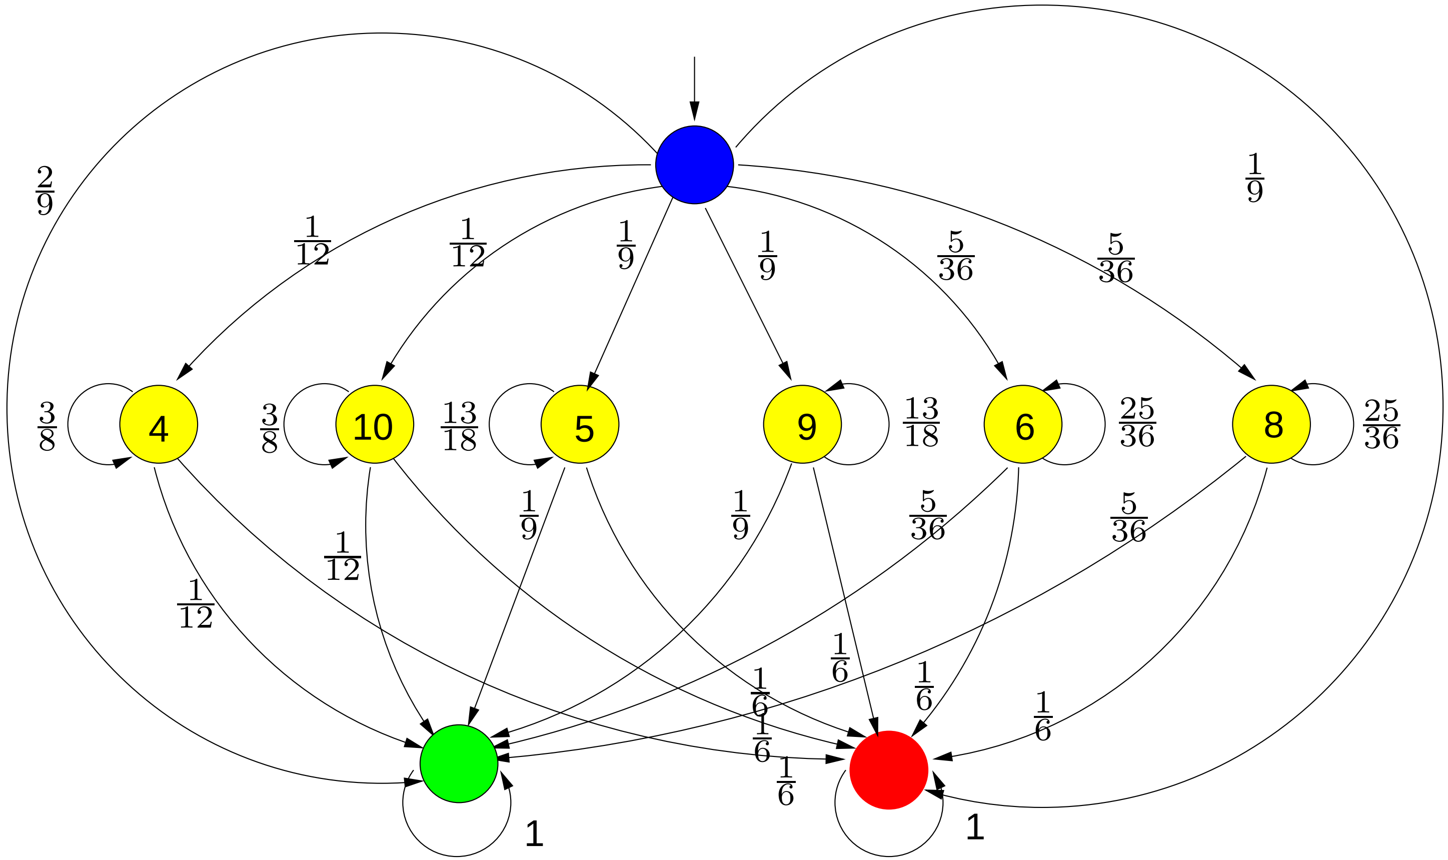

refine (“split”) with respect to the set of red states

A second refinement

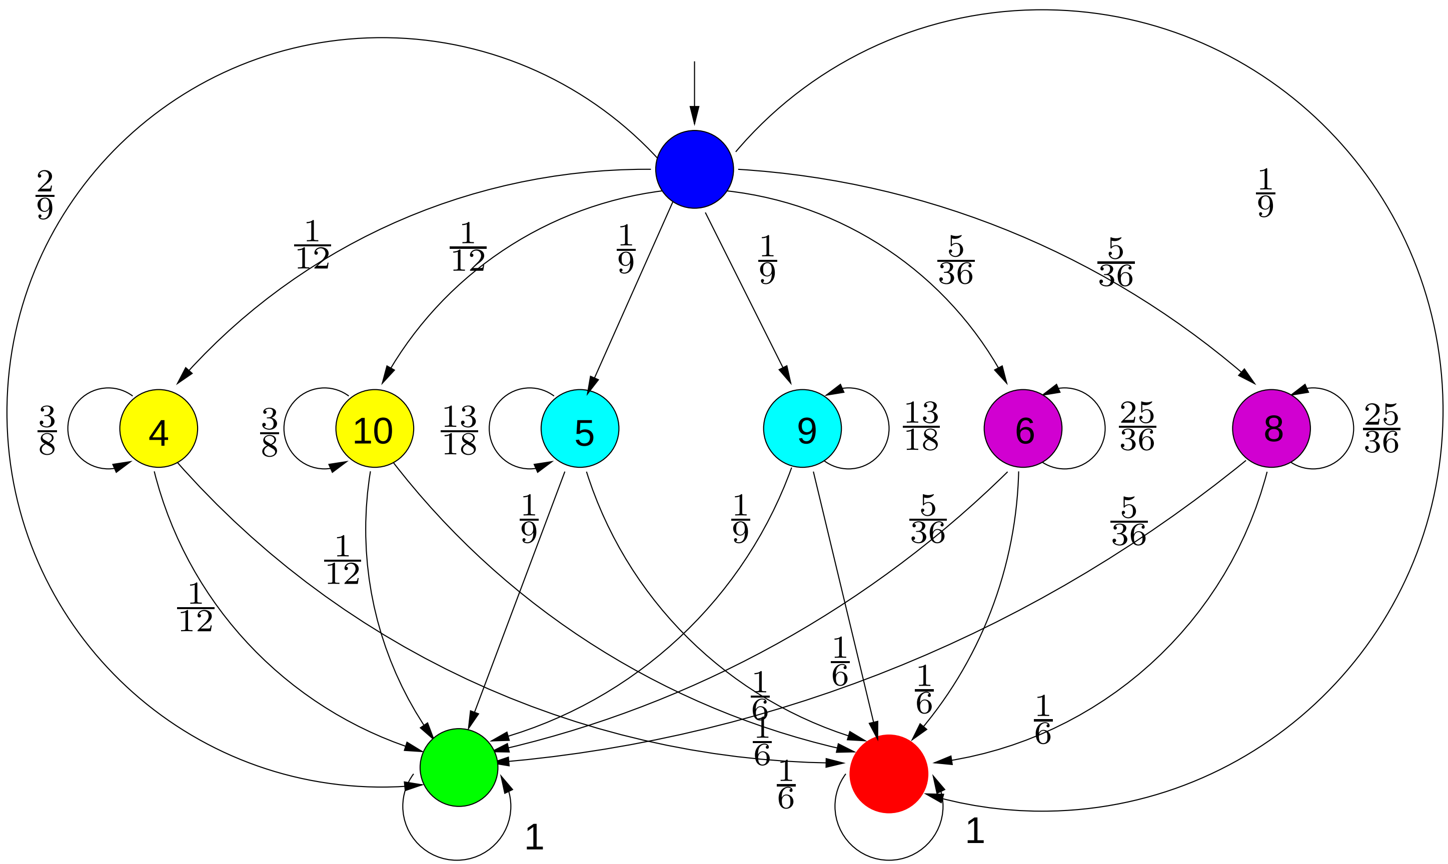

refine (“split”) with respect to the set of green states

Quotient DTMC

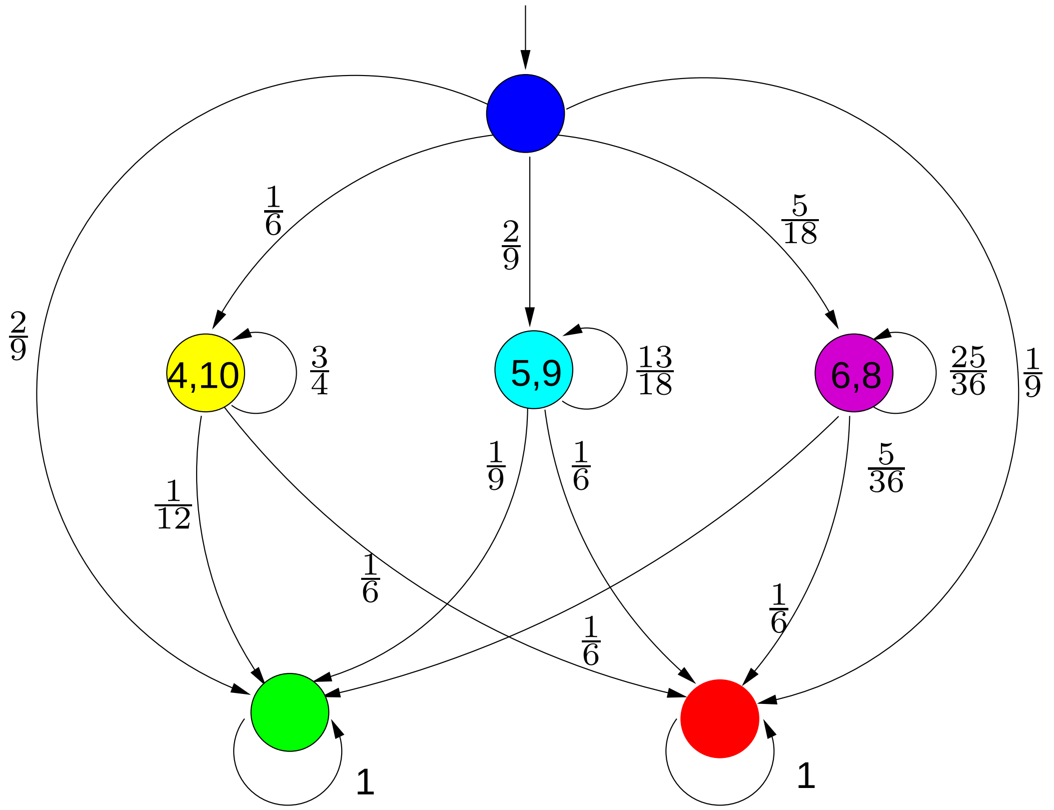

IEEE 802.11 group communication protocol

	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**

all verification times concern timed reachability properties

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

Overview

	strong bisimulation $\sim$	weak bisimulation $\approx$	strong simulation $\sqsubseteq$	weak simulation $\precapprox$
logical preservation	CSL	$\text{CSL}_{\setminus \bigcirc}$	safeCSL	$\text{safeCSL}_{\setminus \bigcirc}$
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	$\mathcal{O}(m \log n)$	$\mathcal{O}(n^3)$	–	–

Content of this lecture

- Continuous Stochastic Logic
  - syntax, semantics, examples
- CSL model checking
  - basic algorithms and complexity
- Bisimulation
  - definition, minimization algorithm, examples

$\Rightarrow$ Priced continuous-time Markov chains
  - motivation, definition, some properties

Power consumption in mobile ad-hoc networks

- Single battery-powered mobile phone with ad-hoc traffic
- Two types of traffic: ad-hoc traffic and ordinary calls
  - offer transmission capabilities for data transfer between third parties (altruism)
  - normal call traffic
- Prices are used to model power consumption
  - in *doze* mode (20 mA), calls can neither be made nor received
  - active calls are assumed to consume 200 mA
  - ad-hoc traffic and call handling takes 120 mA; idle mode costs 50 mA
  - total battery capacity is 750 mAh; price equals one mA

A priced stochastic Petri net model

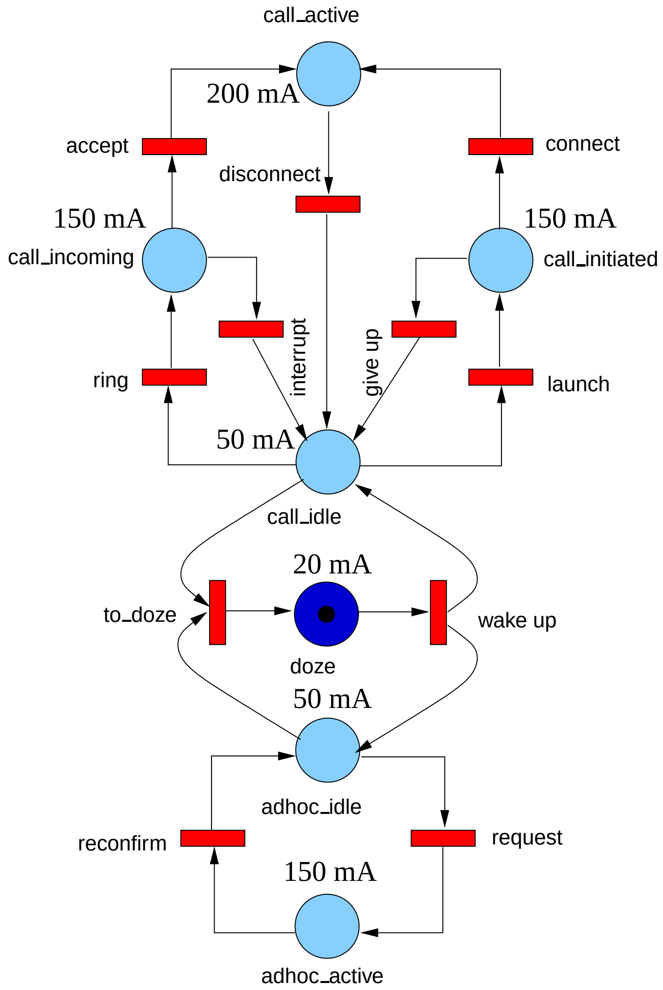

transition	mean time (in min)	rate (per h)
accept	20	180
connect	10	360
disconnect	4	15
doze	5	12
give up	1	60
interrupt	1	60
launch	80	0.75
reconfirm	4	15
request	10	6
ring	80	0.75
wake up	16	3.75

Required properties

- The probability to receive a call within 24 hours exceeds 0.23
- The probability to receive a call while having consumed at most 80% power exceeds 0.99
- The probability to launch a call before consuming at most 80% power within 24 hours – while using the phone only for ad-hoc transfer beforehand – exceeds 0.78

Priced continuous-time Markov chains

A CMRM is a triple $(S, \mathbf{R}, L, \rho)$ where:

- S is a set of states, $\mathbf{R}$ a rate matrix and L a labelling (as before)
- $\rho : S \to \mathbb{R}_{\geqslant 0}$ is a price function

Interpretation:

- Staying t time units in state s costs $\rho(s) \cdot t$

Cumulating price

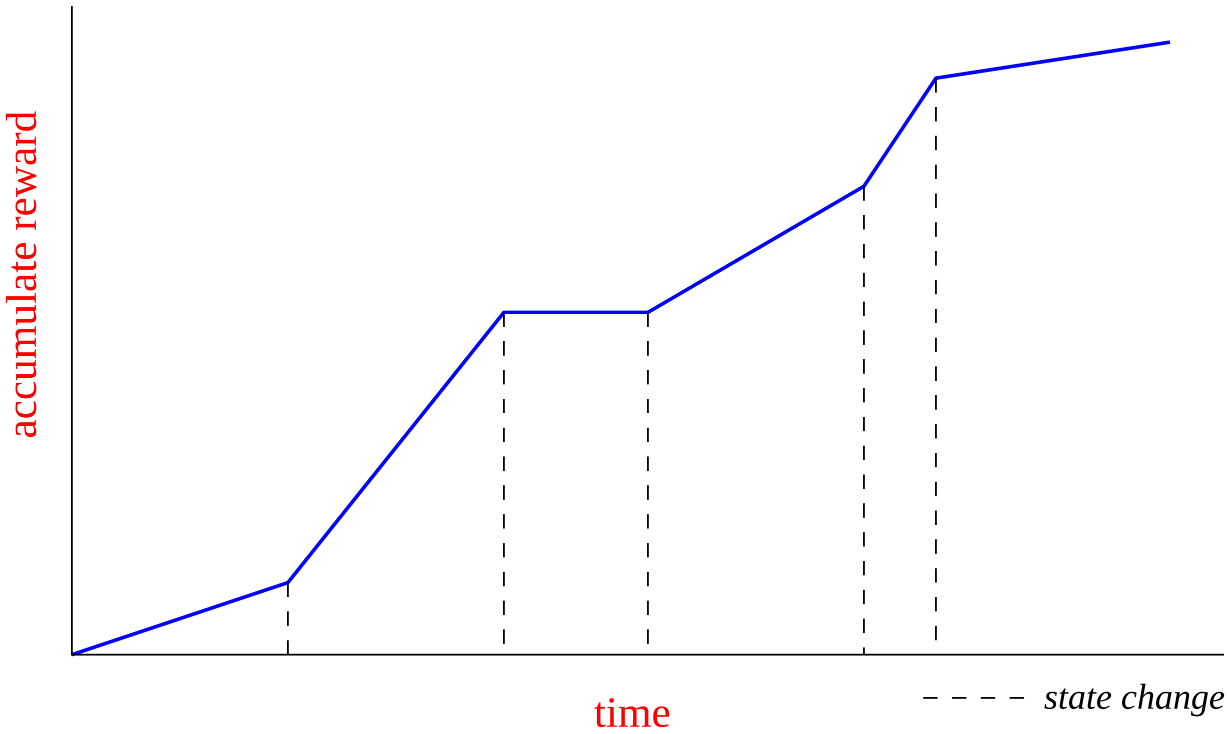

Time- and cost-bounded reachability

- In $\geqslant$ 92% of the cases, a goal state is reached with *cost at most 62*:

$$\mathcal{P}_{\geqslant 0.92}\left(\neg\, \textit{illegal} \;\; \mathsf{U}_{\leqslant 62} \;\; \textit{goal}\right)$$

- . . . . . . within 133.4 time units: $\mathcal{P}_{\geqslant 0.92}\left(\neg\, \textit{illegal} \;\; \mathsf{U}^{\leqslant 133.4}_{\leqslant 62} \;\; \textit{goal}\right)$

- Possible to put constraints on:
  - the *likelihood* with which certain behaviours occur,
  - the *time frame* in which certain events should happen, and
  - the *prices* (or: rewards) that are allowed to be made.

Checking time- and cost-bounded reachability

- $s \models \mathbb{P}_L(\Phi \, \mathsf{U}^I_J \, \Psi)$ if and only if $\Pr\{s \models \Phi \, \mathsf{U}^I_J \, \Psi\} \in L$

- For $I = [0, t]$ and $J = [0, r]$, $\Pr\{s \models \Phi \, \mathsf{U}^{\leqslant t}_{\leqslant r} \Psi\}$ is the least solution of:

  - 1 if $s \models \Psi$
  - if $s \models \Phi$ and $s \not\models \Psi$:

$$\int_{K(s)} \sum_{s' \in S} \mathbf{R}(s, s') \cdot e^{-r(s) \cdot x} \cdot \Pr\{s' \models \Phi \, \mathsf{U}^{\leqslant t - x}_{\leqslant r - \rho(s) \cdot x} \, \Psi\} \, dx$$

    where $K(s) = \{ \, x \in I \mid \rho(s) \cdot x \in J \, \}$ is subset of I whose price lies in J

  - 0 otherwise

Duality: model transformation

- Key concept: exploit duality of time advancing and price increase

- The dual of an MRM $\mathcal{C}$ with $\rho(s) > 0$ into MRM $\mathcal{C}^*$:

$$\mathbf{R}^*(s, s') = \frac{\mathbf{R}(s, s')}{\rho(s)} \quad \text{and} \quad \rho^*(s) = \frac{1}{\rho(s)}$$

  state space S and the state-labelling L in $\mathcal{C}$ are unaffected

- So, accelerate state s if $\rho(s) < 1$ and slow it down if $\rho(s) > 1$

Duality theorem

- Transform any state-formula by swapping price and time bounds:

$$\left(\Phi \, \mathsf{U}_{J}^{I} \, \Psi\right) * \;=\; \Phi^{*} \, \mathsf{U}_{I}^{J} \, \Psi^{*}$$

- Duality theorem: $\underbrace{s \models \mathbb{P}_L \left(\Phi \, \mathsf{U}_{J}^{I} \, \Psi\right)}_{\text{in } \mathcal{C}}$ iff $\underbrace{s \models \mathbb{P}_L \left(\Phi^{*} \, \mathsf{U}_{I}^{J} \, \Psi^{*}\right)}_{\text{in } \mathcal{C}^{*}}$

$\Rightarrow$ $\boxed{\text{Verifying } \mathsf{U}_{J} \text{ (in } \mathcal{C}\text{) is identical to model-checking } \mathsf{U}^{J} \text{ (in } \mathcal{C}^{*}\text{)}}$

Proof sketch

$$\Pr_{\mathcal{C}^*}(s \models \Diamond^{\leqslant c}_{\leqslant t} G)$$

$=$ (* for $s \notin G$ *)

$$\int_{K^*} \sum_{s' \in S} \mathbf{R}^*(s, s') \cdot e^{-r^*(s) \cdot x} \cdot \Pr_{\mathcal{C}^*} \left(s' \models \Diamond^{\leqslant c \ominus x}_{\leqslant t \ominus \rho^*(s) \cdot x} G \right) \, dx$$

$=$ (* substituting $y = \frac{x}{\rho(s)}$ *)

$$\int_{K} \sum_{s' \in S} \mathbf{R}(s, s') \cdot e^{-r(s) \cdot y} \cdot \Pr_{\mathcal{C}^*} \left(s' \models \Diamond^{\leqslant c \ominus \rho(s) \cdot y}_{\leqslant t \ominus y} G \right) \, dy$$

$=$ (* $\mathcal{C}$ and $\mathcal{C}^*$ have same digraph, equation system has unique solution *)

$$\int_{K} \sum_{s' \in S} \mathbf{R}(s, s') \cdot e^{-r(s) \cdot y} \cdot \Pr_{\mathcal{C}} \left(s' \models \Diamond^{\leqslant c \ominus \rho(s) \cdot y}_{\leqslant t \ominus y} G \right) \, dy$$

$=$ (* $s \notin G$ *)

$$\Pr_{\mathcal{C}^*} \left(s \models \Diamond^{\leqslant t}_{\leqslant c} G \right)$$

Reduction to transient rate probabilities

Consider the formula $\Phi \, \mathsf{U}^{\leqslant t}_{\leqslant c} \, \Psi$ on MRM $\mathcal{C}$

- Approach: *transform* the MRM $\mathcal{C}$ as follows
  - make all Ψ-states and all $\neg\,(\Phi \,\vee\, \Psi)$-states absorbing
  - equip all these absorbing states with price 0
- Theorem: $\underbrace{s \models \mathbb{P}_J(\Phi \, \mathsf{U}^{\leqslant t}_{\leqslant c} \, \Psi)}_{\text{in MRM } \mathcal{C}}$ iff $\underbrace{s \models \mathbb{P}_J(\Diamond^{=t}_{\leqslant c} \, \Psi)}_{\text{in MRM } \mathcal{C}'}$
- This amounts to compute the transient rate distribution in $\mathcal{C}'$

$\Rightarrow$ Algorithms to compute this measure are not widespread!

A discretization approach

- *Discretise* both time and accumulated price as (small) d
  - probability of > 1 transition in d time-units is negligible (Tijms & Veldman 2000)

- $\Pr(s \models \Diamond^{[t,t]}_{\leqslant c} \Psi) \;\approx\; \sum_{s' \models \Psi} \sum_{k=1}^{c/d} F^{t/d}(s', k) \cdot d$

- Initialization: $F^1(s, k) = 1/d$ if $(s, k) = (s_0, \underline{\rho}(s_0))$, and 0 otherwise

- $F^{j+1}(s, k) = \underbrace{F^j(s, k{-}\rho(s)) \cdot (1{-}r(s) \cdot d)}_{\text{be in state } s \text{ at epoch } j} + \sum_{s' \in S} \underbrace{F^j(s', k{-}\rho(s')) \cdot \mathbf{R}(s', s) \cdot d}_{\text{be in } s' \text{ at epoch } j}$

- Time complexity: $\mathcal{O}(|S|^3 \cdot t^2 \cdot d^{-2})$ (for all states)

Discretization

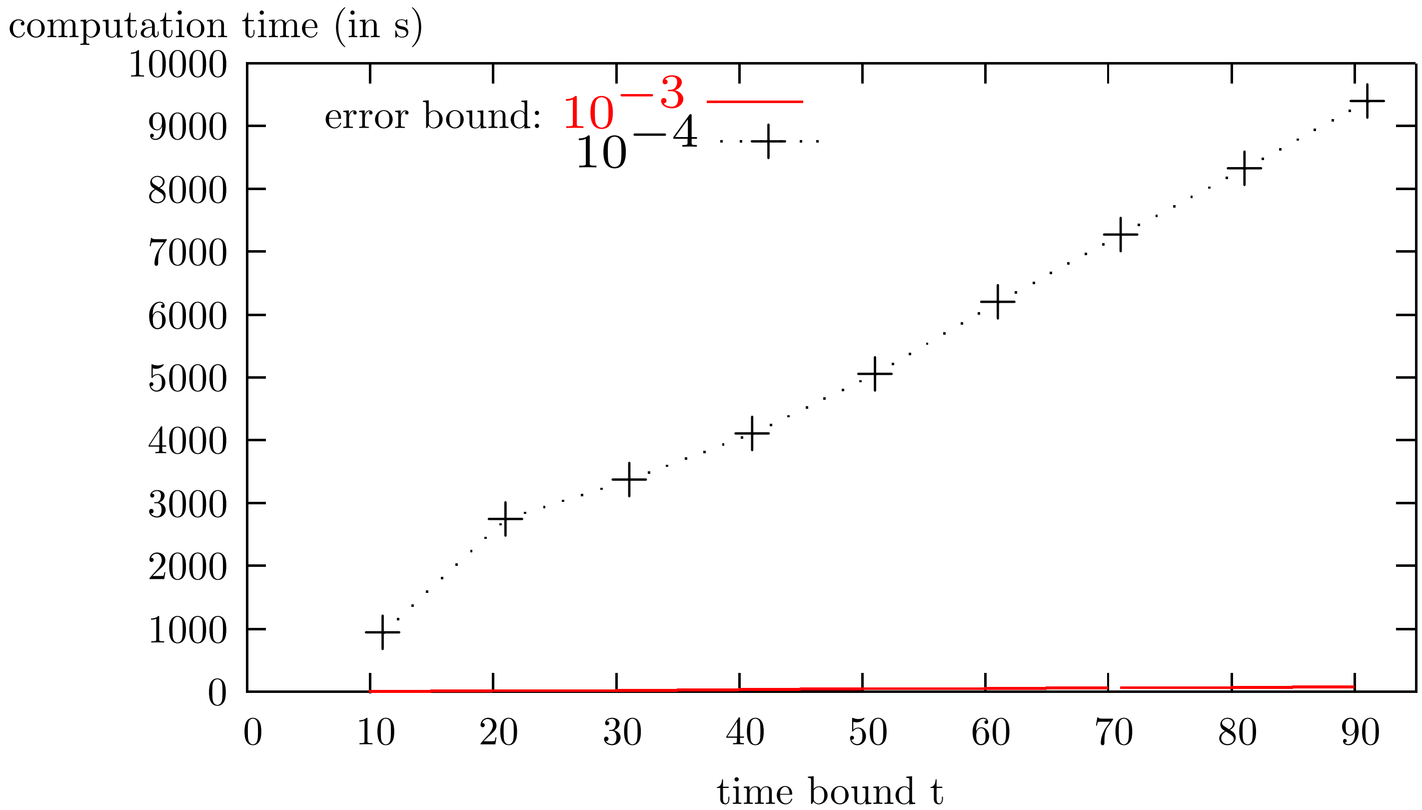

about 300 states; error bound not known

Discretization

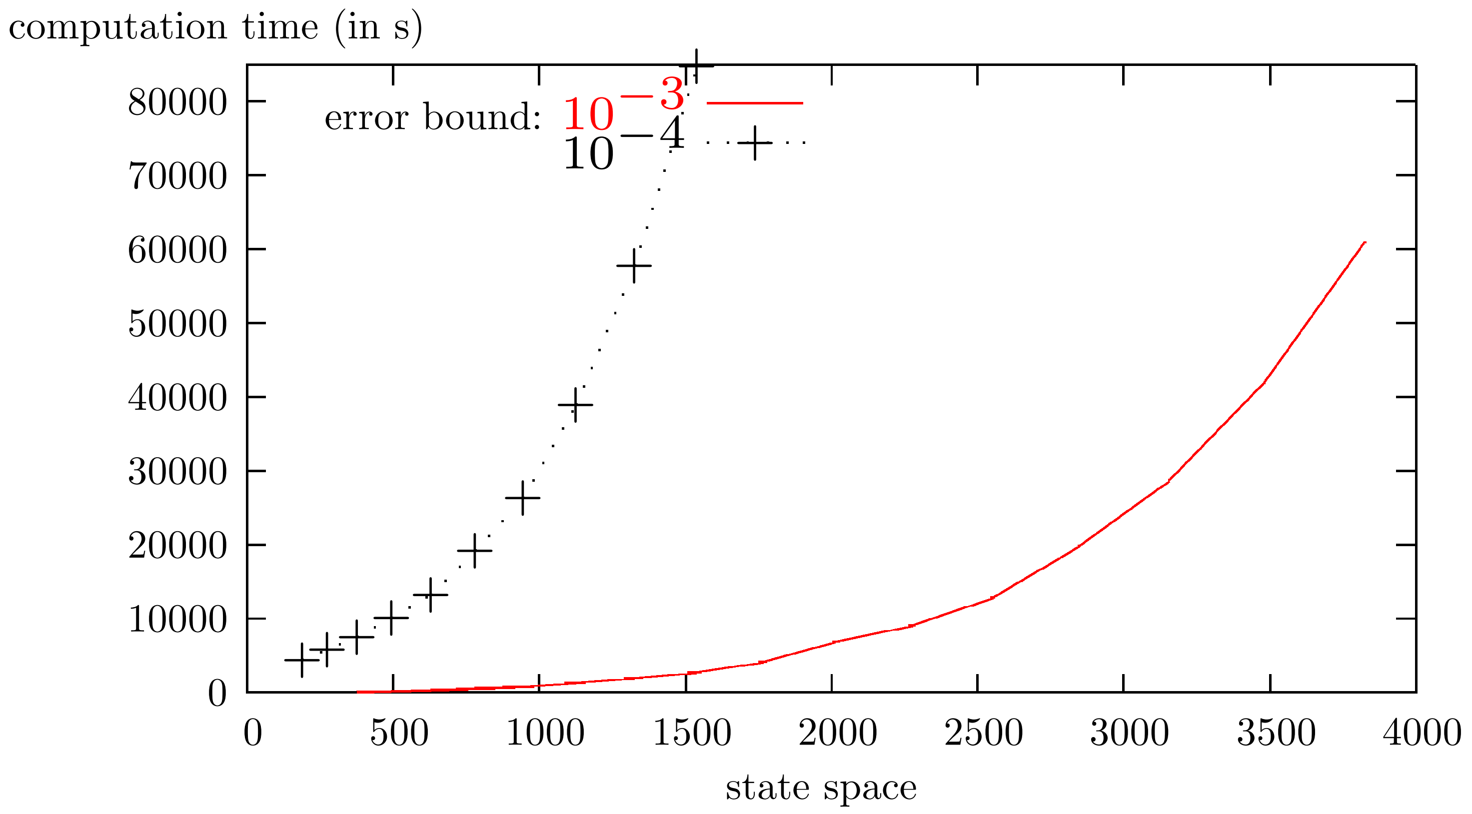

Perspectives

- Linear real-time specifications (MTL, timed automata)
- Aggressive abstraction techniques
- Counterexample generation
- Continuous-time Markov decision processes
- Parametric model checking
- Infinite-state model checking
- . . . . . .

CTMC model checking

- . . . . . . is a mature automated technique
- . . . . . . has a broad range of applications
- . . . . . . is supported by powerful software tools
- . . . . . . extendible to prices
- . . . . . . supported by aggressive abstraction

more information: `www.mrmc-tool.org`

- CTMC model checking
  - CSL: [Baier, Haverkort, Hermanns & Katoen, IEEE Trans. Softw. Eng., 2003]
  - linear timed specifications: [Chen, Han, Katoen & Mereacre, LICS 2009]
- Bisimulation minimization
  - [Derisavi, Hermanns & Sanders, IPL 2005], [Valmari & Franceschinis, TACAS 2010]
  - [Katoen, Kemna, Zapreev & Jansen, TACAS 2007]
- Priced continuous-time Markov chain model checking
  - [Baier, Haverkort, Hermanns & Katoen, ICALP 2000]
  - [Baier, Cloth, Haverkort, Hermanns & Katoen, DSN 2005/FMSD 2010]
- CTMC abstraction
  - 3-valued abstraction: [Katoen, Klink, Leucker & Wolf, CONCUR 2008]
  - compositional abstraction: [Katoen, Klink and Neuhäusser, FORMATS 2009]