Continuous-time Markov chains

Lecture #18 of Probabilistic Models for Concurrency

Joost-Pieter Katoen

Lehrstuhl II: Programmiersprachen u. Softwarevalidierung

E-mail: `katoen@cs.rwth-aachen.de`

July 13, 2005

Overview Lecture #18

⇒ *Continuous-time Markov chains*

- The negative exponential distribution
- What is a continuous-time Markov chain?
- Transient distribution: uniformization
- Limiting and stationary distribution
- Simulation and bisimulation

Probability density

- X is a random variable
  - on a sample space with probability measure $\Pr$
  - assume the set of possible values of X is a continuous interval
- X is *continuously distributed* if there exists a function $f(x)$ such that:

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

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

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

Exponential distribution

Continuous random variable X is *exponential* with parameter $\lambda > 0$ if its density is

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

Cumulative distribution of X:

$$F_X(d) \;=\; \int_0^d \lambda{\cdot}e^{-\lambda\cdot x}\,dx = 1 - e^{-\lambda\cdot d}$$

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

Exponential pdf and cdf

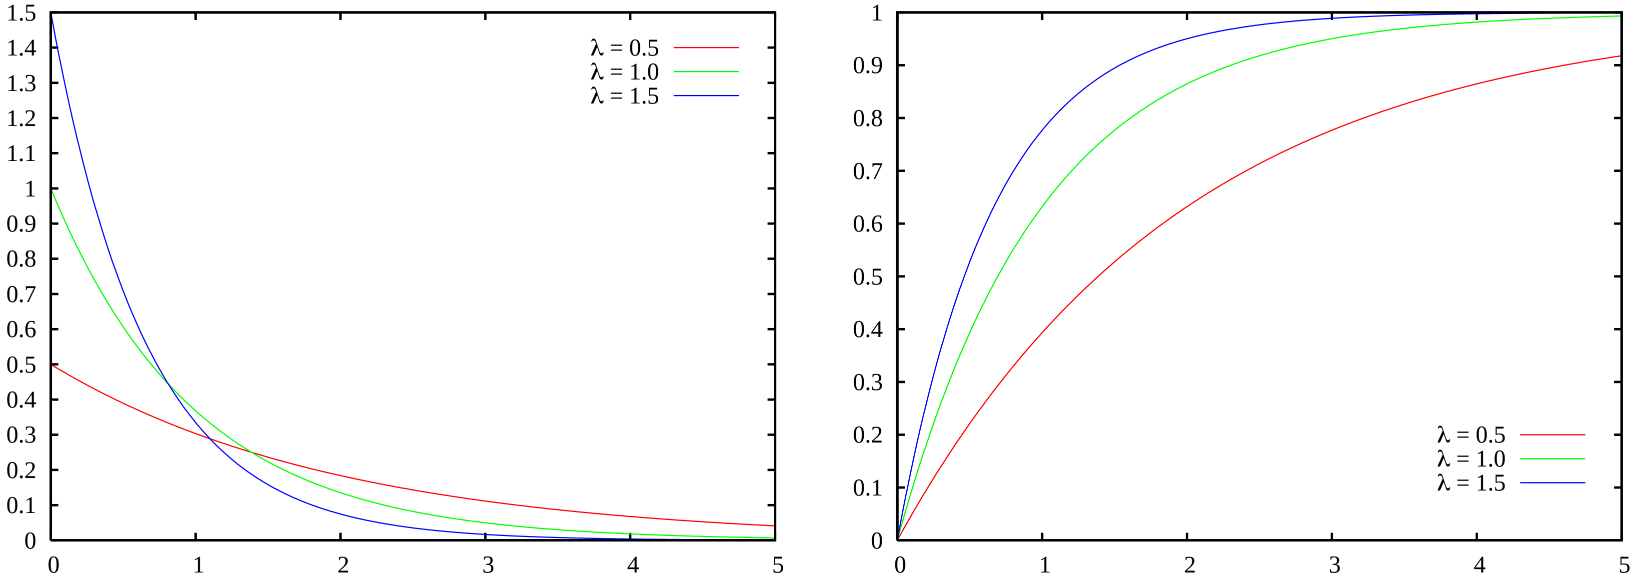

Exponential distributions

- have *nice mathematical* properties (cf. next slide)
- are *adequate* for many real-life phenomena
  - describes the time for a continuous process to change state
  - the time until you have your next car accident (failure rates)
  - the inter-arrival times (i.e., the times between customers entering a shop)
- combinations can *approximate* general distributions arbitrarily closely
- maximal *entropy* probability distribution if just the mean is known

Properties

- An exponential distribution possesses the *memory-less property*

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

- Let X and Y be exponential random variables with rate λ and μ
- $\min(X, Y)$ is exponentially distributed with rate $\lambda{+}\mu$
- $\Pr\{X = \min(X, Y)\} = \frac{\lambda}{\lambda+\mu}$
- $\max(X, Y)$ is not an exponential, but a phase-type distribution

exponential distributions are closed under $\min$*, but not under* $\max$

Continuous-time Markov chain

- A *time-homogeneous continuous-time Markov chain* (CTMC) is
  - a Markov process
  - with continuous parameter T and discrete state space $X(t)$
  - which is time-homogeneous

- $p_s(t) = \Pr\{ X(t) = s \}$ probability to be in state s at time instant t

- Probability of being in state s' at time t when in s at step $t' < t$:

$$\begin{aligned} p_{s,s'}(t', t) &= \Pr\{ X(t) = s' \mid X(t') = s \} \\ &= \Pr\{ X(t{-}t') = s' \mid X(0) = s \} \end{aligned}$$

Another perspective

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

- S is a countable set of states
- $\mathbf{R} : S \times S \rightarrow \mathbb{R}_{\geqslant 0}$, a *rate matrix*
  - $\mathbf{R}(s, s') = \lambda$ means that the average speed of going from s to s' is $\frac{1}{\lambda}$
- $E(s) = \sum_{s' \in S} \mathbf{R}(s, s') = \mathbf{R}(s, S)$ is the *exit rate* of state s
  - s is called absorbing whenever $E(s) = 0$

a CTMC is a transition system (unlabeled transitions)
where transitions are equipped with continuous probabilities

Interpretation

- 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:

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

- The probability to take an outgoing transition from s within $[0, t]$ is:

$$1 - e^{-E(s) \cdot t}$$

Time-abstract evolution of an example CTMC

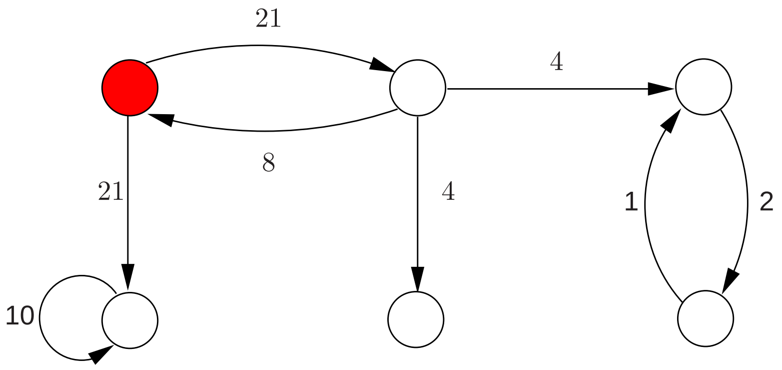

zero-th epoch

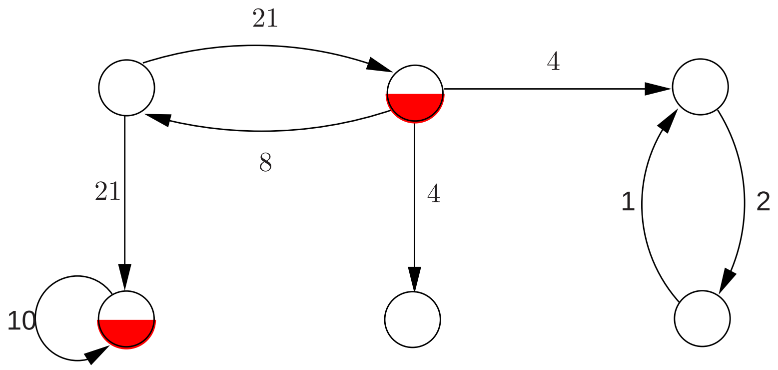

first epoch

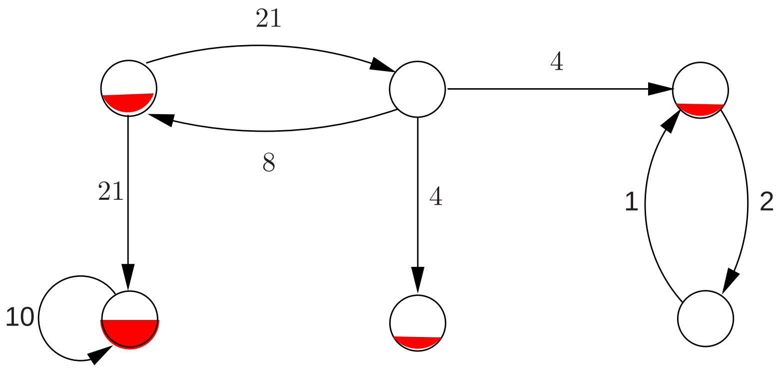

second epoch

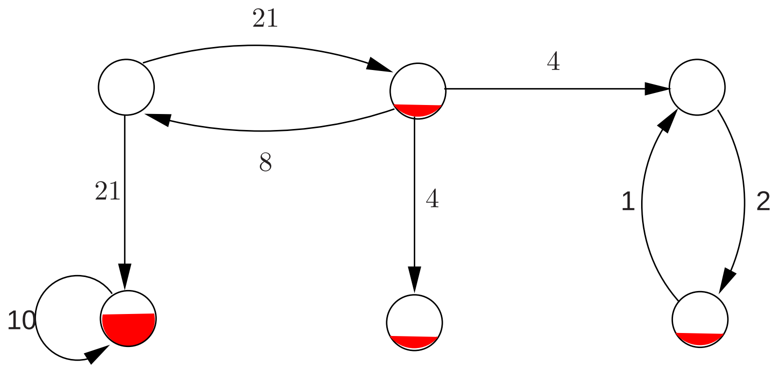

third epoch

On the long run

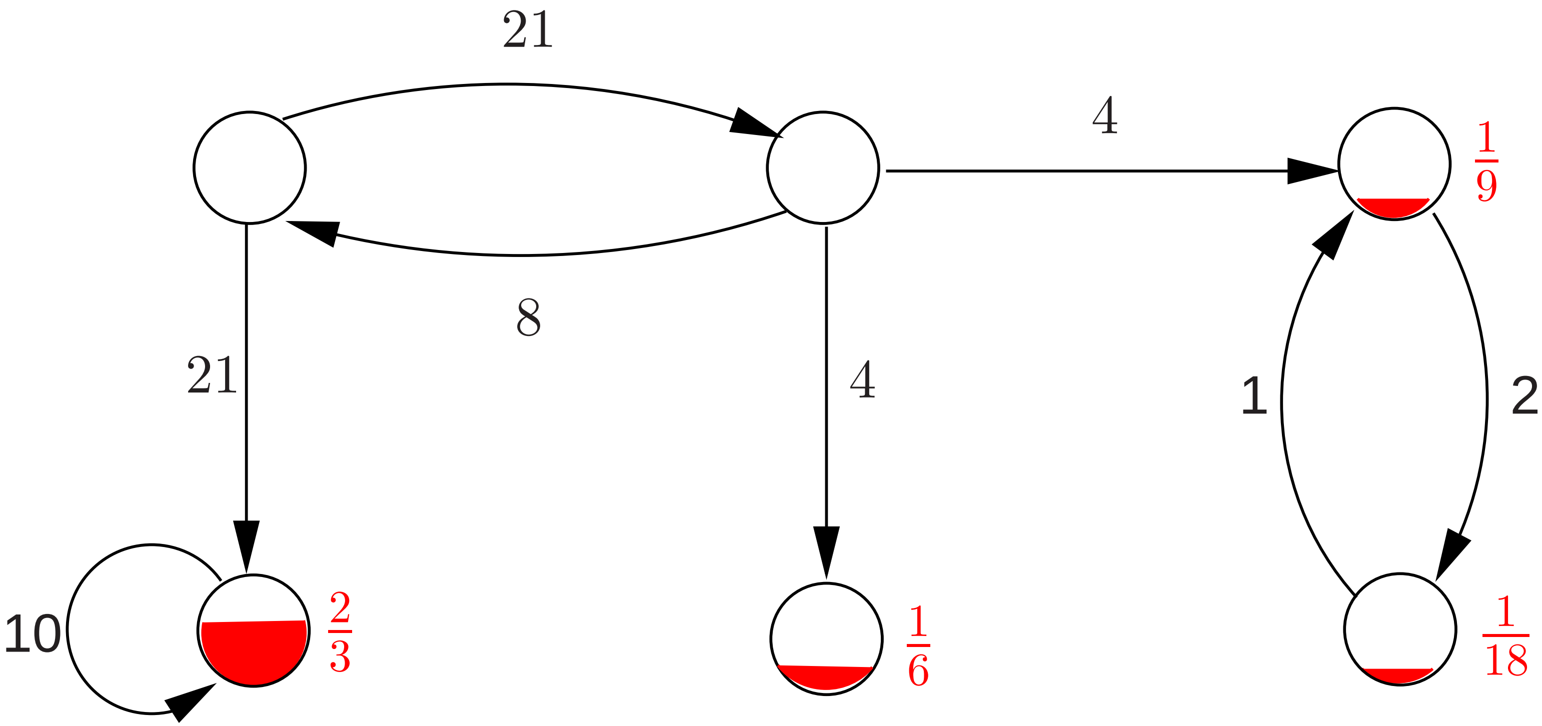

Embedded DTMC

The *embedded* DTMC of the CTMC $(S, \mathbf{R})$ is $(S, \mathbf{P})$ where

$$\mathbf{P}(s, s') = \begin{cases} \frac{\mathbf{R}(s,s')}{E(s)} & \text{if } E(s) > 0 \\ 0 & \text{otherwise} \end{cases}$$

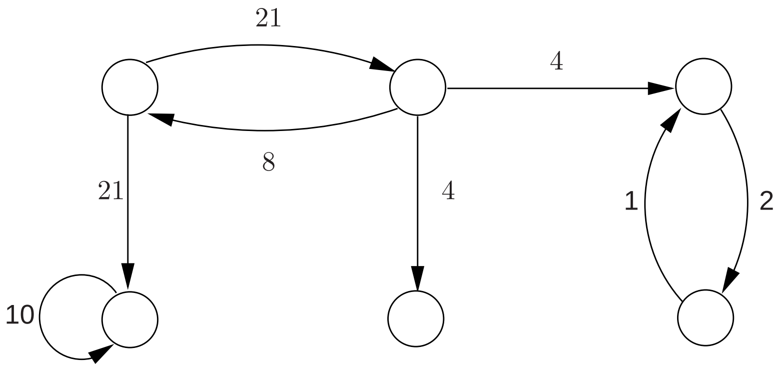

a CTMC

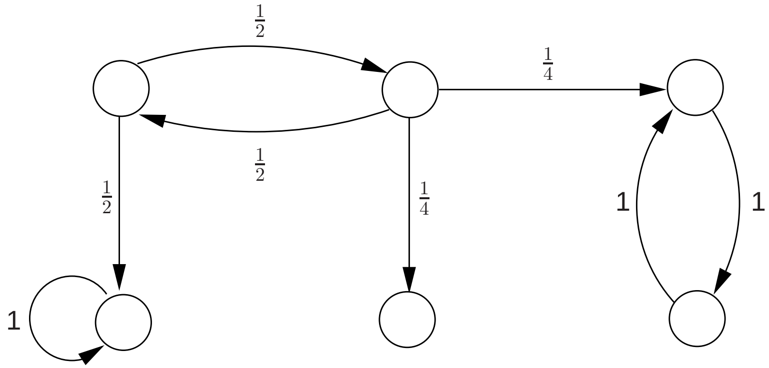

iits embedded DTMC

Transient distribution of a CTMC

Probability to be in state s at time t:

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

Using $\underline{p}(t) = (p_{s_0}(t), p_{s_1}(t), \ldots, p_{s_k}(t))$ we obtain in matrix form:

$$\boxed{\underline{p}'(t) = \underline{p}(t) \cdot \mathbf{Q} \quad \text{given} \quad \underline{p}(0)}$$

where $\mathbf{Q} = \mathbf{R} - \textit{diag}(E)$ is the infinitesimal generator matrix

$\underline{p}(t)$ is the transient-state probability vector at time t

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

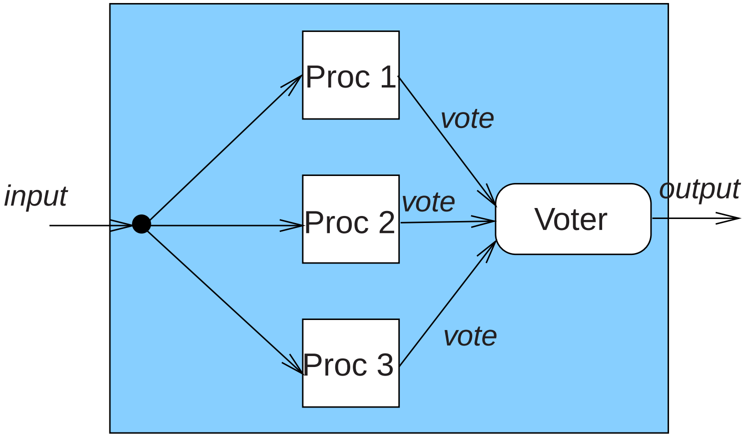

- Modelling assumptions:
  - if voter fails, entire system goes down
  - after voter-repair, system starts “as new”
  - state $=$ (#processors, #voters)

Modelling a TMR system as a CTMC

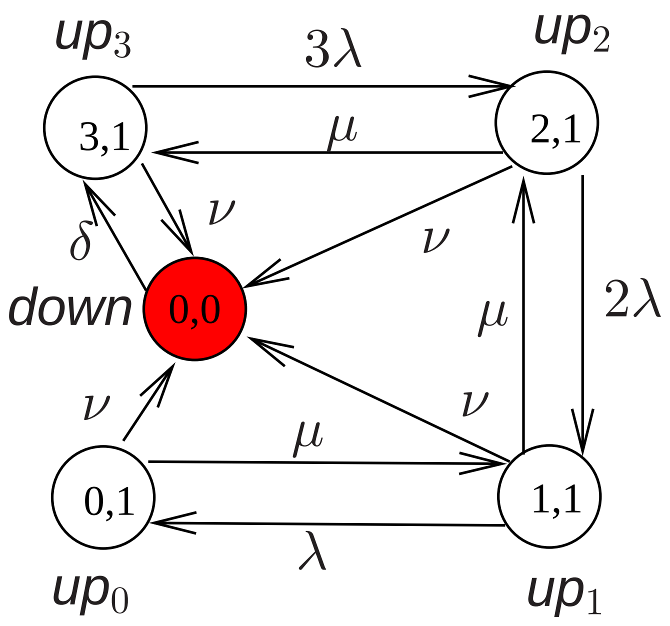

- processor failure rate is λ fph;
  its repair rate is μ rph
- voter failure rate is ν fph;
  its repair rate is δ rph
- rate matrix: e.g., $\mathbf{R}((3,1),(2,1)) = 3\lambda$
- exit rates: e.g., $E(3,1) = 3\lambda{+}\nu$
- probability matrix: e.g.,

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

Transient probabilities

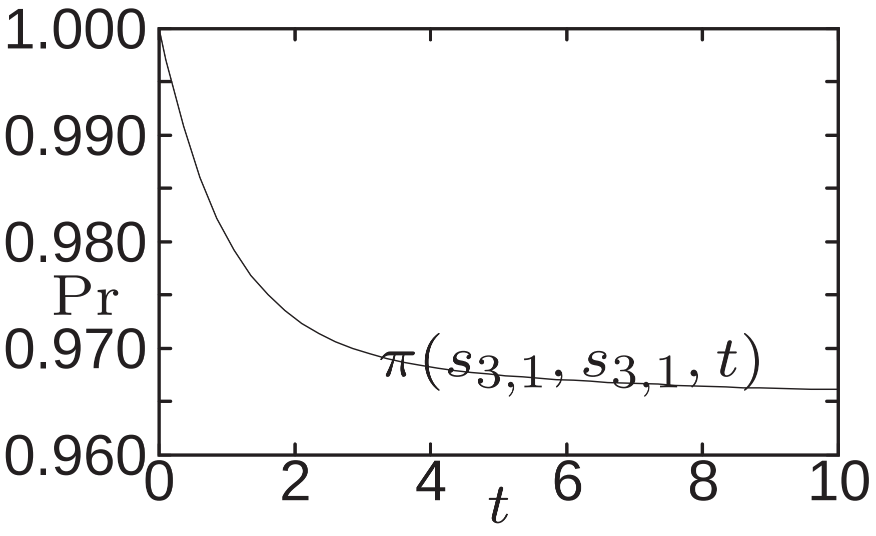

$p_{s_{3,1}}(t)$ for first 10 hours

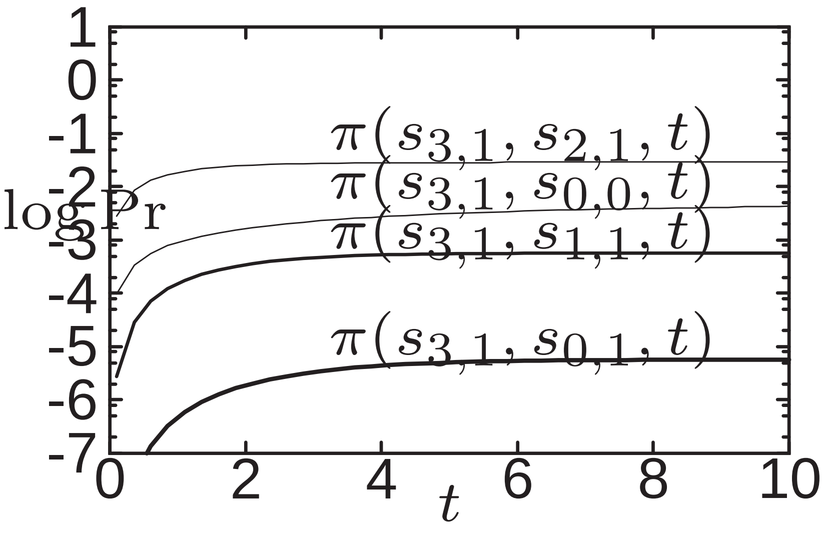

$p(t)$ for first 10 hours (logscale)

$\lambda = 0.01$ failures per hour (fph), $\nu = 0.001$ fph

$\mu = 1$ repairs per hour (rph) and $\delta = 0.2$ rph

© B.R. Haverkort)

Steady-state distribution of a CTMC

Assuming a stationary distribution exists (e.g., finite and irreducible):

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

Using $\underline{p} = (p_{s_0}, p_{s_1}, \ldots, p_{s_k})$ we obtain in matrix form:

$$\boxed{\underline{p}\cdot\mathbf{Q} \;=\; 0 \quad \text{where} \quad \textstyle\sum_{s\in S}\, p_s = 1}$$

$\underline{p}$ is the steady-state probability vector
and is obtained by solving a system of linear equations

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 is 0.994

$\lambda = 0.01$ failures per hour (fph), $\nu = 0.001$ fph
$\mu = 1$ repairs per hour (rph) and $\delta = 0.2$ rph

Computing transient probabilities

- Solution to $\underline{p}'(t) = \underline{p}(t) \cdot \mathbf{Q}$ is: $\underline{p}(t) = \underline{p}(0) \cdot e^{\mathbf{Q}t} = \underline{p}(0) \cdot \sum_{i=0}^{\infty} \frac{(\mathbf{Q}t)^i}{i!} \quad (*)$
- Main problems: infinite summation + numerical instability due to
  - $\mathbf{Q}^i$ becomes non-sparse with positive and negative entries
- Solution: transform CTMC $(S, \mathbf{R})$ into DTMC $(S, \mathbf{U})$ with
  - $\mathbf{U} := \mathbf{I} + \frac{\mathbf{Q}}{q}$ with $q \geqslant \max_i \{ E(s_i) \}$

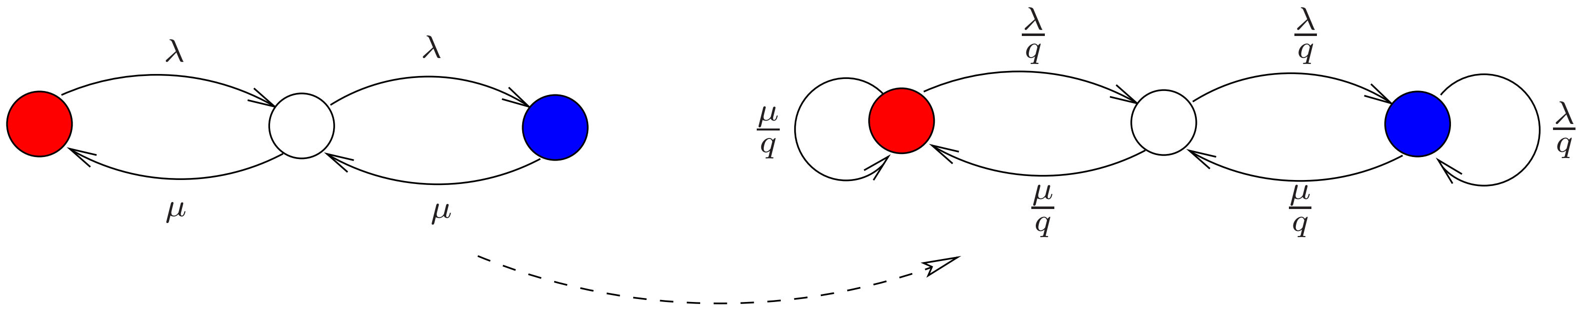

uniformisation with $q = \lambda + \mu$

Computing transient probabilities

- Now (*): $\underline{p}(t) = \underline{p}(0) \cdot e^{q(\mathbf{U}-\mathbf{I})t} = \underline{p}(0) \cdot e^{-qt} \cdot e^{qt\mathbf{U}} = \sum_{i=0}^{\infty} \underbrace{e^{-qt} \frac{(qt)^i}{i!}}_{\text{Poisson prob.}} \underline{p}(i)$

- Summation can be truncated *a priori* for a given error bound ε:

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

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

$\Rightarrow$ Transient analysis of a CTMC $\approx$ transient analysis of a DTMC

Markovian bisimulation

- Let $\mathcal{C} = (S, \mathbf{R})$ be a CTMC and R an equivalence relation on S
- R is a *Markovian bisimulation* on S if for any $(s, s') \in R$:

$$\underbrace{\mathbf{P}(s, \cdot) \ \equiv_R \ \mathbf{P}(s', \cdot) \quad \text{and} \quad E(s) = E(s')}_{\mathbf{R}(s,\cdot) \ \equiv_R \ \mathbf{R}(s',\cdot)}$$

where $\equiv_R$ denotes the lifting of R on *Distr*(S) defined by:

$$\mu \ \equiv_R \ \mu' \quad \text{iff} \quad \mu(C) = \mu'(C) \qquad \text{for all} \qquad C \in S/R$$

Quotient transition system

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

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

where

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

$$\mathbf{R}'(C, C') \ = \sum_{s \in C} \mathbf{R}(s, C')$$

Preservation of state probabilities

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

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

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

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

Markovian simulation

- Let $\mathcal{C} = (S, \mathbf{R})$ be a CTMC and R a binary relation on S
- R is a *Markovian simulation* on S if for all $(s, s') \in R$:

$$\mathbf{P}(s, \cdot) \sqsubseteq_R \mathbf{P}(s', \cdot) \quad \text{and} \quad E(s) \leqslant E(s')$$

- s' simulates s, denoted $s \sqsubseteq_m s'$, if there exists

  a Markovian simulation R on S such that $(s, s') \in R$

Example

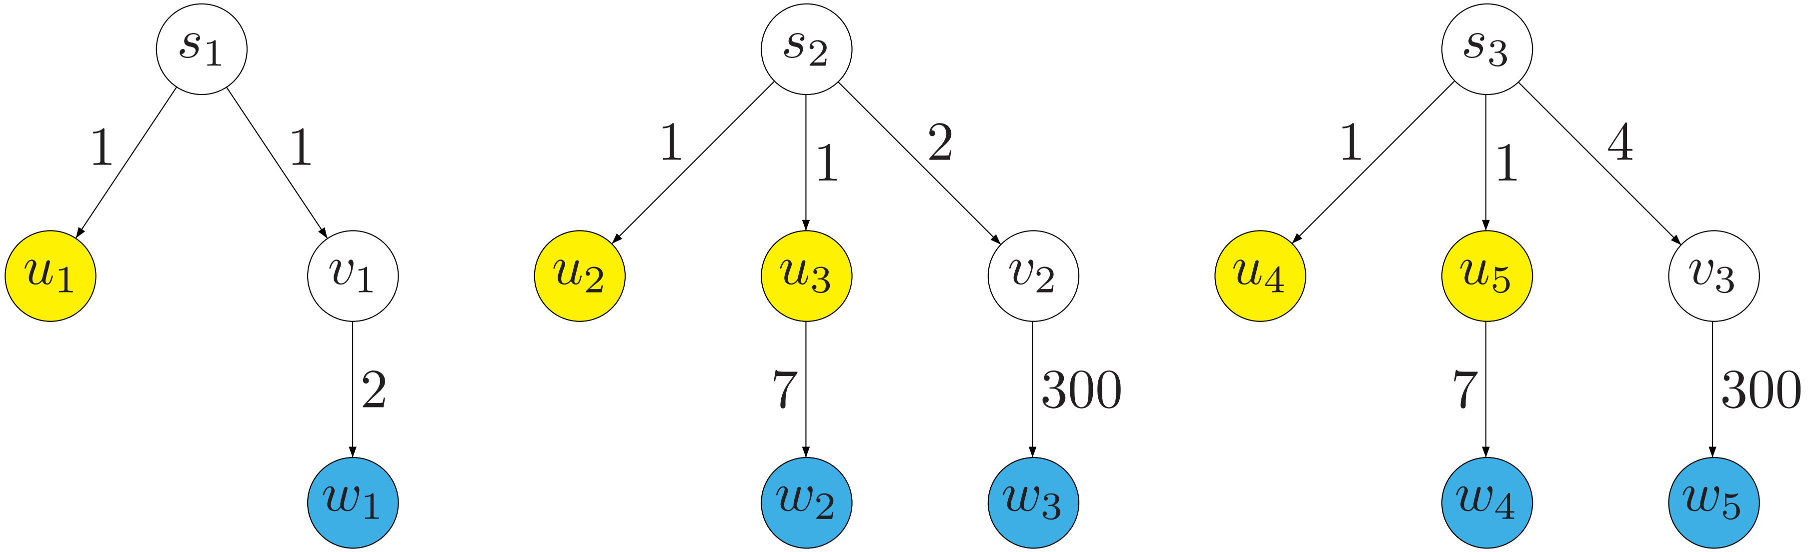

$s_1 \sqsubseteq_m s_2$ but $s_2 \not\sqsubseteq_m s_3$

Some properties

1. $\sqsubseteq_m$ is a pre-order
2. $s \sim_m s'$ implies $s \sqsubseteq_m s'$
3. Markovian simulation equivalence $\simeq_m$ agrees with $\sim_m$
4. $s \sim_m s'$ implies $s \sim_p s'$ in the embedded DTMC
5. $s \sqsubseteq_m s'$ implies $s \sqsubseteq_p s'$ in the embedded DTMC
6. if $E(s) = E$ for any s then $\sqsubseteq_m$ agrees with $\simeq_m$

Relating (bi)simulation on CTMCs and DTMCs

$$\begin{tikzcd}
\sim_m \arrow[r, shift left] \arrow[d, shift left] & \sim_d \arrow[l, dashed, shift left] \arrow[d, shift left] \\
\simeq_m \arrow[u, shift left] \arrow[r, shift left] \arrow[d, shift left] & \simeq_d \arrow[l, dashed, shift left] \arrow[u, shift left] \arrow[d, shift left] \\
\sqsubseteq_m \arrow[u, dashed, shift left] \arrow[r, shift left] & \sqsubseteq_d \arrow[l, dashed, shift left] \arrow[u, dashed, shift left]
\end{tikzcd}$$

$R \longrightarrow R'$ means that R is coarser than R'

Probabilistic timed reachability

For any $C \subseteq S$ such that $C = C{\uparrow}_{\sqsubseteq_m}$:

$$s \sqsubseteq_m s' \quad \Rightarrow \quad \underbrace{\Pr\left\{s \overset{\leqslant t}{\rightsquigarrow} C\right\}}_{p(s,t,C)} \leqslant \Pr\left\{s' \overset{\leqslant t}{\rightsquigarrow} C\right\} \qquad \text{for any } t \geqslant 0$$

where $p(s,t,C) = \lim_{n \to \infty} p(s,t,n,C)$ with:

$$p(s,t,n,C) = \begin{cases} 1 & \text{if } s \in C \\ \displaystyle\int_0^t \sum_{s' \in S} \mathbf{R}(s,s') \cdot e^{-E(s) \cdot x} \cdot p(s', t-x, n-1, C)\, dx & \text{if } s \notin C \text{ and } n > 0 \\ 0 & \text{otherwise} \end{cases}$$

generalization possible by forbidding paths visiting $B \subseteq S$ with $B = B{\uparrow}_{\sqsubseteq_m}$