Lecture 8
Continuous-time Markov chains

Alessandro Abate

Department of Computer Science
University of Oxford

Overview

- Transient probabilities
  - uniformisation
- Steady-state probabilities
- CSL: Continuous Stochastic Logic
  - syntax
  - semantics
  - examples

Recall CTMC notions

- Continuous-time Markov chain: $C = (S, s_{init}, \mathbf{R}, L)$
  - $\mathbf{R} : S \times S \to \mathbb{R}_{\geq 0}$ is the transition rate matrix
  - rates interpreted as parameters of exponential distributions

- Embedded DTMC: $emb(C) = (S, s_{init}, \mathbf{P}^{emb(C)}, L)$

$$\mathbf{P}^{emb(C)}(s,s') = \begin{cases} \mathbf{R}(s,s')/E(s) & \text{if } E(s) > 0 \\ 1 & \text{if } E(s) = 0 \text{ and } s = s' \\ 0 & \text{otherwise} \end{cases}$$

- Infinitesimal generator matrix

$$\mathbf{Q}(s,s') = \begin{cases} \mathbf{R}(s,s') & s \neq s' \\ -\sum_{s \neq s'} \mathbf{R}(s,s') & \text{otherwise} \end{cases}$$

Transient and steady-state behaviour

- Transient behaviour
  - state of the model at a particular time instant
  - $\underline{\pi}^{C}_{s,t}(s')$ is the probability, having started in state s, of being in state s' at time t in CTMC C
  - $\underline{\pi}^{C}_{s,t}(s') = Pr_s\{\ \omega \in Path^C(s) \mid \omega@t=s'\ \}$

- Steady-state behaviour
  - state of the model in the long-run
  - $\underline{\pi}^{C}_{s}(s')$ is probability, having started in state s, of being in state s' in the long run
  - $\underline{\pi}^{C}_{s}(s') = \lim_{t\to\infty} \underline{\pi}^{C}_{s,t}(s')$
  - intuitively: long-run percentage of time spent in each state

Computing transient probabilities

- Consider simple example and compare to case for DTMCs
- What is the probability of being in state s_0 at time t?
- DTMC vs CTMC:

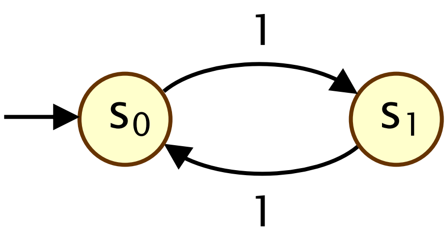

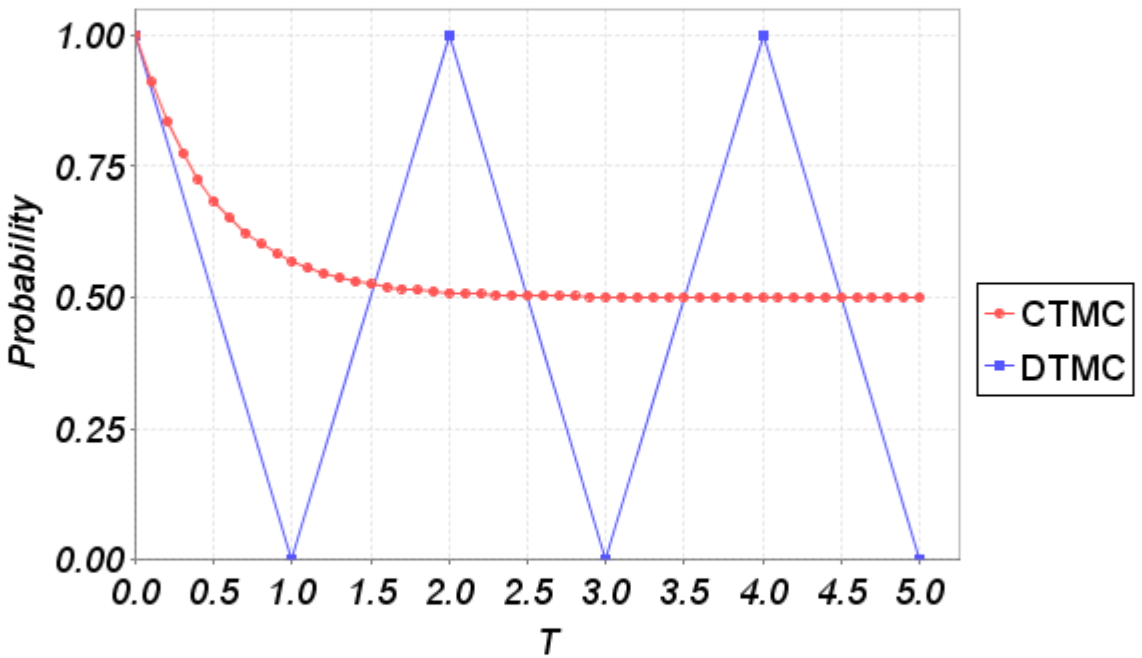

Computing transient probabilities

- Π_t – matrix of transient probabilities at time t
  - $\Pi_t(s,s') = \underline{\pi}_{s,t}(s')$
- Π_t solution of the differential equation: $\Pi_t' = \Pi_t \cdot Q$
  - where **Q** is the infinitesimal generator matrix
- Can be expressed as a matrix exponential and therefore evaluated as a power series

$$\Pi_t = e^{Q \cdot t} = \sum_{i=0}^{\infty} (Q \cdot t)^i / i!$$

- computation potentially unstable numerically
- probabilities instead computed using uniformisation

Uniformisation

- We build the uniformised DTMC unif(C) of CTMC C
- If C $=(S,s_{init},\mathbf{R},L)$, then unif(C) $= (S,s_{init},\mathbf{P}^{unif(C)},L)$
  - set of states, initial state and labelling the same as C
  - $\mathbf{P}^{unif(C)} = \mathbf{I} + \mathbf{Q}/q$
  - $\mathbf{I}$ is the $|S|\times|S|$ identity matrix
  - $q \geq \max \{ E(s) \mid s \in S \}$ is the uniformisation rate

- Each time step (epoch) of uniformised DTMC corresponds to one exponentially distributed delay with rate q
  - if $E(s)=q$ transitions the same as embedded DTMC (residence time has the same distribution as one epoch)
  - if $E(s)<q$ add self loop with probability $1-E(s)/q$ (residence time longer than $1/q$ so one epoch may not be 'long enough')

Uniformisation – Example

- CTMC C:

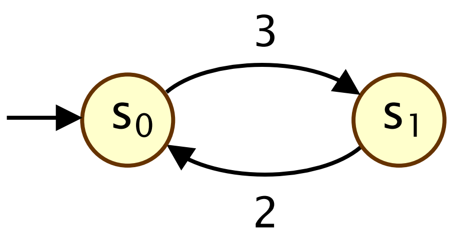

$$\mathbf{R} = \begin{bmatrix} 0 & 3 \\ 2 & 0 \end{bmatrix} \qquad \mathbf{Q} = \begin{bmatrix} -3 & 3 \\ 2 & -2 \end{bmatrix}$$

- Uniformised DTMC unif(C)
  - let uniformisation rate q = $\max_s$ { E(s) } = 3

$$\mathbf{P}^{\mathrm{unif(C)}} = \mathbf{I} + \mathbf{Q}/q = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \begin{bmatrix} -1 & 1 \\ 2/3 & -2/3 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 2/3 & 1/3 \end{bmatrix}$$

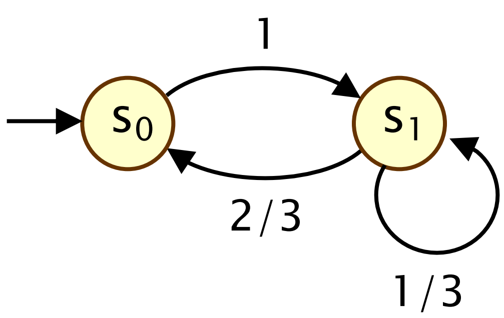

Uniformisation

- Using the uniformised DTMC the transient probabilities can be expressed by:

$$\Pi_t = e^{Q \cdot t}$$

Uniformisation

- Using the uniformised DTMC the transient probabilities can be expressed by:

$$
\begin{aligned}
\Pi_t &= e^{Q \cdot t} = e^{q \cdot (\mathbf{P}^{unif(C)} - \mathbf{I}) \cdot t} = e^{(q \cdot t) \cdot \mathbf{P}^{unif(C)}} \cdot e^{-q \cdot t} \\
&= e^{-q \cdot t} \cdot \left(\sum_{i=0}^{\infty} \frac{(q \cdot t)^i}{i!} \cdot \left(\mathbf{P}^{unif(C)} \right)^i \right) \\
&= \sum_{i=0}^{\infty} \left(e^{-q \cdot t} \cdot \frac{(q \cdot t)^i}{i!} \right) \cdot \left(\mathbf{P}^{unif(C)} \right)^i \\
&= \sum_{i=0}^{\infty} \gamma_{q \cdot t, i} \cdot \left(\mathbf{P}^{unif(C)} \right)^i
\end{aligned}
$$

ith Poisson probability with parameter $q \cdot t$

$\mathbf{P}^{unif(C)}$ is stochastic (all entries in [0,1] & rows sum to 1); therefore computations with **P** are more numerically stable than **Q**

Uniformisation

$$\Pi_t = \sum_{i=0}^{\infty} \gamma_{q \cdot t,i} \cdot \left(\mathbf{P}^{unif(C)} \right)^i$$

- $(\mathbf{P}^{unif(C)})^i$ is probability of jumping between each pair of states in i steps
- $\gamma_{q \cdot t,i}$ is the ith Poisson probability with parameter $q \cdot t$
  - the probability of i steps occurring in time t, given each has delay exponentially distributed with rate q
- Can truncate the (infinite) summation using techniques of Fox and Glynn [FG88], which allow efficient computation of the Poisson probabilities, and provide error bounds

Uniformisation

- Computing $\underline{\pi}_{s,t}$ for a fixed state s and time t
  - can be computed efficiently using matrix-vector operations
  - pre-multiply the matrix Π_t by the initial distribution
  - in this case: $\underline{\pi}_{s,0}(s')$ equals 1 if s=s' and 0 otherwise

$$\underline{\pi}_{s,t} = \underline{\pi}_{s,0} \cdot \Pi_t = \underline{\pi}_{s,0} \cdot \sum_{i=0}^{\infty} \gamma_{q \cdot t,i} \cdot \left(\mathbf{P}^{unif(C)} \right)^i$$

$$= \sum_{i=0}^{\infty} \gamma_{q \cdot t,i} \cdot \underline{\pi}_{s,0} \cdot \left(\mathbf{P}^{unif(C)} \right)^i$$

  - compute iteratively to avoid the computation of matrix powers

$$\left(\underline{\pi}_{s,t} \cdot \mathbf{P}^{unif(C)} \right)^{i+1} = \left(\underline{\pi}_{s,t} \cdot \mathbf{P}^{unif(C)} \right)^{i} \cdot \mathbf{P}^{unif(C)}$$

Uniformisation – Example

- CTMC C, uniformised DTMC for q=3

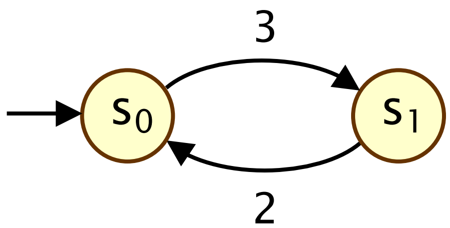

$$\mathbf{R} = \begin{bmatrix} 0 & 3 \\ 2 & 0 \end{bmatrix} \quad \mathbf{Q} = \begin{bmatrix} -3 & 3 \\ 2 & -2 \end{bmatrix} \quad \mathbf{P}^{unif(C)} = \begin{bmatrix} 0 & 1 \\ 2/3 & 1/3 \end{bmatrix}$$

- Initial distribution: $\underline{\pi}_{s0,0} = [\ 1, 0\ ]$
- Transient probabilities for time t = 1:

$$\underline{\pi}_{s0,1} = \sum_{i=0}^{\infty} \gamma_{q \cdot t,i} \cdot \underline{\pi}_{s0,0} \cdot \left(\mathbf{P}^{unif(C)} \right)^i$$

$$= \gamma_{3,0} \cdot [1,0] \cdot \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \gamma_{3,1} \cdot [1,0] \cdot \begin{bmatrix} 0 & 1 \\ 2/3 & 1/3 \end{bmatrix} + \gamma_{3,2} \cdot [1,0] \cdot \begin{bmatrix} 0 & 1 \\ 2/3 & 1/3 \end{bmatrix}^2 + \ldots$$

$$\approx [\ 0.404043,\ 0.595957\ ]$$

Steady-state probabilities

- Limit $\underline{\pi}^{C}_{s}(s') = \lim_{t\to\infty} \underline{\pi}^{C}_{s,t}(s')$
  - exists for all finite CTMCs (see next slide)

- As for DTMCs, need to consider the underlying graph structure of the Markov chain:
  - reachability (between pairs) of states
  - bottom strongly connected components (BSCCs)
  - one special case to consider: absorbing states are BSCCs
  - note: can do this equivalently on embedded DTMC

- CTMC is irreducible if all its states belong to a single BSCC; otherwise reducible

Periodicity

- Unlike for DTMCs, do not need to consider periodicity

- DTMC/CTMC:

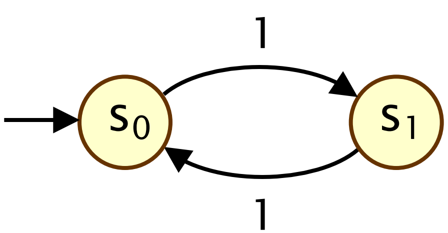

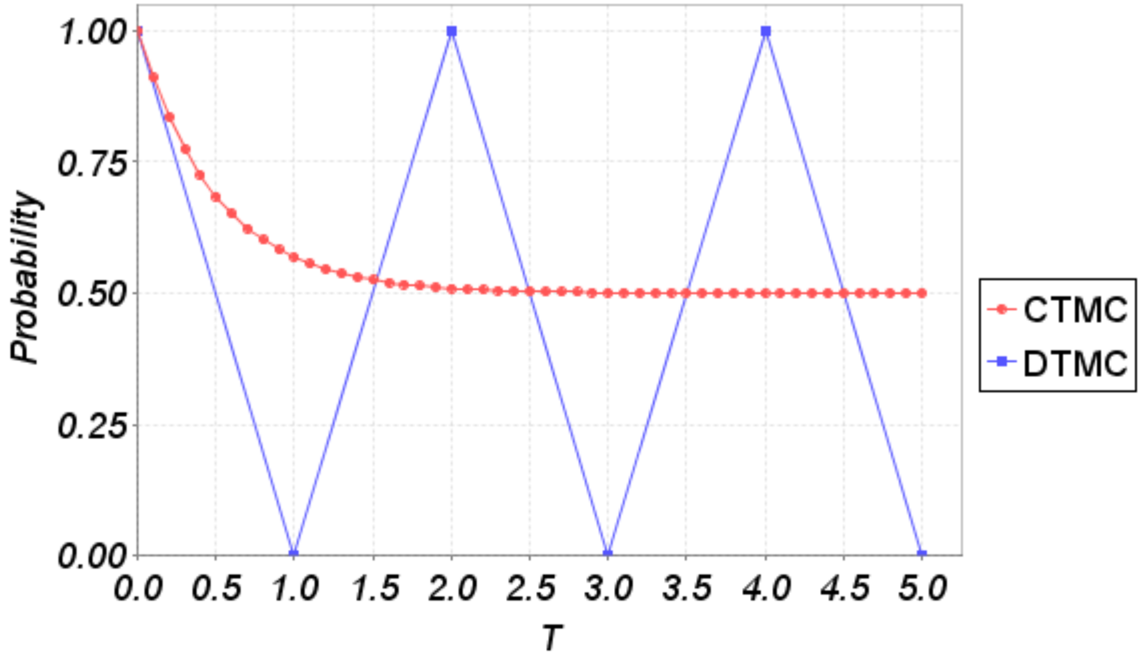

Irreducible CTMCs

- For an irreducible CTMC:
  - the steady-state probabilities are independent of the starting state: denote these steady-state probabilities by $\underline{\pi}^C(s')$

- These probabilities can be computed as
  - the unique solution of the linear equation system:

$$\underline{\pi}^C \cdot \mathbf{Q} = \underline{0} \quad \text{and} \quad \sum_{s \in S} \underline{\pi}^C(s) = 1$$

where $\mathbf{Q}$ is the infinitesimal generator matrix of C

- Solved by standard means (cf. Lec. 5):
  - direct methods, such as Gaussian elimination
  - iterative methods, such as Jacobi and Gauss-Seidel

Balance equations

$$\underline{\pi}^C \cdot \mathbf{Q} = \underline{0} \quad \text{and} \quad \sum_{s \in S} \underline{\pi}^C(s) = 1$$

balance the rate of leaving and entering a state

normalisation

For all $s \in S$:

$$\underline{\pi}^C(s) \cdot \left(-\Sigma_{s' \neq s} \mathbf{R}(s,s')\right) + \Sigma_{s' \neq s} \underline{\pi}^C(s') \cdot \mathbf{R}(s',s) = 0$$

$$\Leftrightarrow$$

$$\underline{\pi}^C(s) \cdot \Sigma_{s' \neq s} \mathbf{R}(s,s') = \Sigma_{s' \neq s} \underline{\pi}^C(s') \cdot \mathbf{R}(s',s)$$

Corresponds to: $\underline{\pi}^C \cdot \mathbf{P} = \underline{\pi}^C$ where **P** is matrix for embedded DTMC

Steady-state – Example

- Solve: $\underline{\pi} \cdot \mathbf{Q} = 0$ and $\Sigma\, \underline{\pi}(s) = 1$

$$\mathbf{Q} = \begin{bmatrix} -3/2 & 3/2 & 0 & 0 \\ 3 & -9/2 & 3/2 & 0 \\ 0 & 3 & -9/2 & 3/2 \\ 0 & 0 & 3 & -3 \end{bmatrix}$$

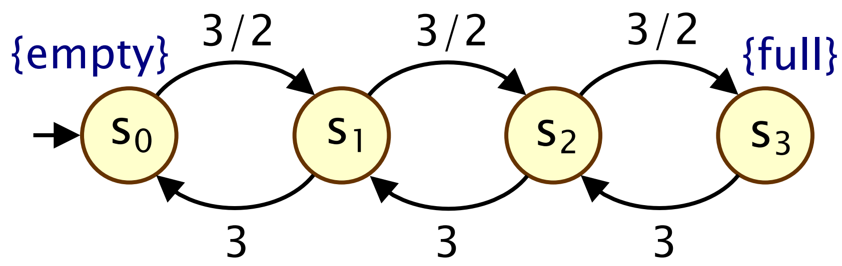

$$\begin{aligned}
-3/2 \cdot \underline{\pi}(s_0) + 3 \cdot \underline{\pi}(s_1) &= 0 \\
3/2 \cdot \underline{\pi}(s_0) - 9/2 \cdot \underline{\pi}(s_1) + 3 \cdot \underline{\pi}(s_2) &= 0 \\
3/2 \cdot \underline{\pi}(s_1) - 9/2 \cdot \underline{\pi}(s_2) + 3 \cdot \underline{\pi}(s_3) &= 0 \\
3/2 \cdot \underline{\pi}(s_2) - 3 \cdot \underline{\pi}(s_3) &= 0 \\
\underline{\pi}(s_0) + \underline{\pi}(s_1) + \underline{\pi}(s_2) + \underline{\pi}(s_3) &= 1
\end{aligned}$$

$$\underline{\pi} = [\, 8/15,\ 4/15,\ 2/15,\ 1/15 \,]$$
$$\sim [\, 0.533,\ 0.267,\ 0.133,\ 0.067 \,]$$

Reducible CTMCs

- For a reducible CTMC:
  - the steady-state probabilities $\underline{\pi}^C(s')$ depend on start state s
- Find all BSCCs of CTMC, denoted bscc(C)
- Compute:
  - steady-state probabilities $\underline{\pi}^T$ of sub-CTMC for each BSCC T
  - probability $ProbReach^{emb(C)}(s, T)$ of reaching each T from s

    (intuitive computation w/ emb(C) shall become clearer in Lec 10)
- Then:

$$\underline{\pi}_s^C(s') = \begin{cases} ProbReach^{emb(C)}(s,T) \cdot \underline{\pi}^T(s') & \text{if } s' \in T \text{ for some } T \in bscc(C) \\ 0 & \text{otherwise} \end{cases}$$

CSL

- Temporal logic for describing properties of CTMCs
  - CSL = Continuous Stochastic Logic [ASSB00,BHHK03]
  - extension of (non-probabilistic) temporal logic CTL
- Key additions:
  - probabilistic operator P (like PCTL)
  - steady state operator S
  - temporal operators over dense time intervals
- Example: $\mathrm{down} \rightarrow P_{>0.75}\,[\, \neg\mathrm{fail}\ U^{[1,2.5]}\ \mathrm{up}\,]$
  - when a shutdown occurs, the probability of a system recovery being completed between 1 and 2.5 hours without further failure is greater than 0.75
- Example: $S_{<0.1}$[insufficient_routers]
  - in the long run, the chance that an inadequate number of routers are operational is less than 0.1

CSL syntax

- CSL syntax:

ψ is true with probability ~p

  - $\phi ::= \text{true} \mid a \mid \phi \wedge \phi \mid \neg\phi \mid P_{\sim p}\,[\psi] \mid S_{\sim p}\,[\phi]$ (state formulae)

  - $\psi ::= X\,\phi \quad \mid \quad \phi\; U^{I}\; \phi$ (path formulae)

"next"

"time bounded until"

in the "long run" ϕ is true with probability ~p

  - where a is an atomic proposition, I interval of $\mathbb{R}_{\geq 0}$ and $p \in [0,1]$, $\sim \in \{<,>,\leq,\geq\}$

- A CSL formula is always a state formula
  - path formulae only occur inside the P operator

CSL semantics for CTMCs

- CSL formulae interpreted over states of a CTMC
  - $s \vDash \phi$ denotes ϕ is "true in state s" or "satisfied in state s"
- Semantics of state formulae:
  - for a state s of the CTMC $(S, s_{init}, \mathbf{R}, L)$:

  - $s \vDash a \quad \Leftrightarrow \quad a \in L(s)$
  - $s \vDash \phi_1 \wedge \phi_2 \quad \Leftrightarrow \quad s \vDash \phi_1$ and $s \vDash \phi_2$
  - $s \vDash \neg\phi \quad \Leftrightarrow \quad s \vDash \phi$ is false
  - $s \vDash P_{\sim p}[\psi] \quad \Leftrightarrow \quad Prob(s, \psi) \sim p$
  - $s \vDash S_{\sim p}[\phi] \quad \Leftrightarrow \quad \Sigma_{s' \vDash \phi} \underline{\pi}_s(s') \sim p$

Probability of, starting in state s, satisfying the path formula ψ

Probability of, starting in state s, being in state s' in the long run

CSL semantics for CTMCs

- Prob(s, ψ) is the probability, starting in state s, of satisfying the path formula ψ
  - $\text{Prob}(s, \psi) = \text{Pr}_s \{\omega \in \text{Path}_s \mid \omega \vDash \psi \}$

if ω(0) is absorbing ω(1) not defined

- Semantics of path formulae:
  - for a path ω of the CTMC:
  - $\omega \vDash X\ \phi \quad \Leftrightarrow \quad$ ω(1) is defined and $\omega(1) \vDash \phi$
  - $\omega \vDash \phi_1\ U^I\ \phi_2 \quad \Leftrightarrow \quad \exists t \in I.\ (\ \omega@t \vDash \phi_2 \wedge \forall t'<t.\ \omega@t' \vDash \phi_1)$

there exists a time instant within the interval I where ϕ_2 is true, and ϕ_1 is true at all preceding time instants (also before interval I)

More on CSL

- Basic logical derivations:
  - false, $\phi_1 \vee \phi_2$, $\phi_1 \rightarrow \phi_2$
- Normal (unbounded) until is a special case
  - $\phi_1 \ U \ \phi_2 \equiv \phi_1 \ U^{[0,\infty)} \ \phi_2$
- Derived path formulae:
  - $F \ \phi \equiv \text{true} \ U \ \phi$, $F^I \ \phi \equiv \text{true} \ U^I \ \phi$
  - $G \ \phi \equiv \neg(F \ \neg\phi)$, $G^I \ \phi \equiv \neg(F^I \ \neg\phi)$
- Negate probabilities: …
  - e.g. $\neg P_{>p} [\ \psi\ ] \equiv P_{\leq p} [\ \psi\ ]$, $\neg S_{\geq p} [\ \phi\ ] \equiv S_{<p} [\ \phi\ ]$
- Quantitative properties
  - of the form $P_{=?} [\ \psi\ ]$ and $S_{=?} [\ \phi\ ]$
  - where P/S is the *outermost* operator
  - fit for experiments, patterns, trends, …

CSL example – Workstation cluster

- Case study: Cluster of workstations [HHK00]
  - two sub-clusters (N workstations in each cluster)
  - star topology with a central switch
  - components can break down, single repair unit

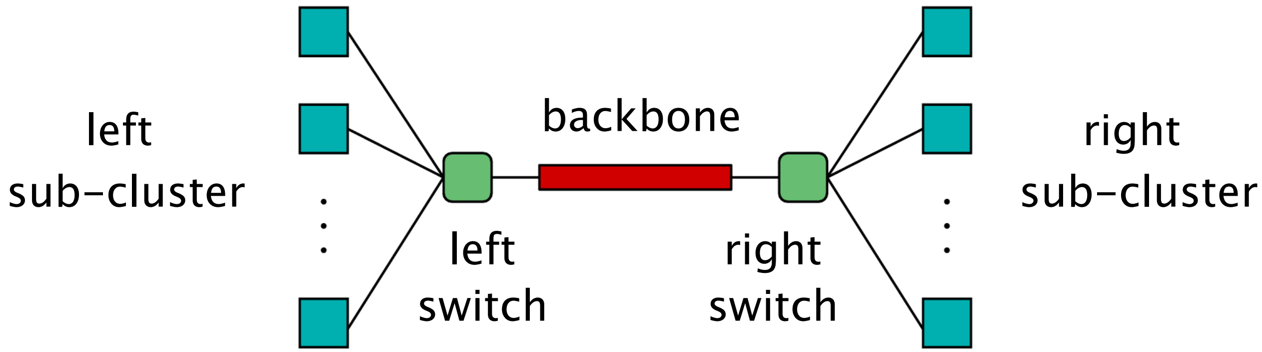

  - minimum QoS: at least ¾ of the workstations operational and connected via switches
  - premium QoS: all workstations operational and connected via switches

CSL example – Workstation cluster

- $S_{=?}$ [minimum]
  - the probability in the long run of having minimum QoS
- $P_{=?}$ [$F^{[t,t]}$ minimum]
  - the (transient) probability at time instant t of minimum QoS
- $P_{<0.05}$ [$F^{[0,10]}$ ¬minimum]
  - the probability that the QoS drops below minimum within 10 hours is less than 0.05
- ¬minimum → $P_{<0.1}$ [$F^{[0,2]}$ ¬minimum]
  - when facing insufficient QoS, the chance of facing the same problem after 2 hours is less than 0.1

CSL example – Workstation cluster

- $\text{minimum} \rightarrow P_{>0.8}\,[\ \text{minimum}\ U^{[0,t]}\ \text{premium}\ ]$
  - the probability of going from minimum to premium QoS within t hours without violating minimum QoS is at least 0.8
- $P_{=?}\,[\ \neg\text{minimum}\ U^{[t,\infty)}\ \text{minimum}\ ]$
  - the chance it takes more than t time units to recover from insufficient QoS
- $\neg\text{r\_switch\_up} \rightarrow P_{<0.1}\,[\neg\text{r\_switch\_up}\ U\ \neg\text{l\_switch\_up}\ ]$
  - if the right switch has failed, the probability of the left switch failing before it is repaired is less than 0.1
- $P_{=?}\,[\ F^{[2,\infty)}\ S_{>0.9}[\ \text{minimum}\ ]\ ]$
  - the probability of it taking more than 2 hours to get to a state from which the long-run probability of minimum QoS is >0.9

Summing up…

- Transient probabilities (time instant t)
  - computation with uniformisation: efficient iterative method
- Steady-state (long-run) probabilities
  - like DTMCs
  - requires graph analysis
  - irreducible case: solve linear equation system
  - reducible case: steady-state for sub-CTMCs + reachability
- CSL: Continuous Stochastic Logic
  - extension of PCTL for properties of CTMCs