Probabilistic Model Checking

Marta Kwiatkowska
Dave Parker

Oxford University Computing Laboratory

ESSLLI'10 Summer School, Copenhagen, August 2010

Course overview

- 5 lectures: Mon–Fri, 11am–12.30pm
  - Introduction
  - 1 – Discrete time Markov chains
  - 2 – Markov decision processes
  - 3 – Continuous-time Markov chains
  - 4 – Probabilistic model checking in practice
  - 5 – Probabilistic timed automata
- Course materials available here:
  - http://www.prismmodelchecker.org/lectures/esslli10/
  - lecture slides, reference list

Probabilistic models

	Fully probabilistic	**Nondeterministic**
Discrete time	Discrete-time Markov chains (DTMCs)	Markov decision processes (MDPs) (probabilistic automata)
Continuous time	Continuous-time Markov chains (CTMCs)	CTMDPs/IMCs
		Probabilistic timed automata (PTAs)

Part 3

Continuous-time Markov chains

Time in DTMCs

- Time in a DTMC (or MDP) 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 exponential distributions
  - suits modelling of: performance/reliability (e.g. of computer networks, manufacturing systems, queueing networks), biological pathways, chemical reactions, …

Overview (Part 3)

- Exponential distribution and its properties

- Continuous-time Markov chains (CTMCs)
  - definition, race conditions, examples
  - paths and probability spaces

- CSL: A temporal logic for CTMCs

- CSL model checking
  - uniformisation, steady-state probabilities

- Extensions: Costs & rewards

Continuous probability distributions

- Defined by:
  - cumulative distribution function

$$F(t) = \Pr(X \le t) = \int_{-\infty}^{t} f(x)\, dx$$

  - where f is the probability density function
  - Pr(X=t) = 0 for all t

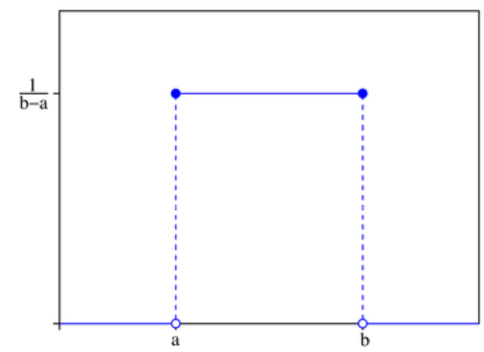

- Example: uniform distribution: U(a,b)

$$f(t) = \begin{cases} 1/_{b-a} & \text{if } a \le t \le b \\ 0 & \text{otherwise} \end{cases}$$

$$F(t) = \begin{cases} 0 & \text{if } t < a \\ {}^{t-a}/_{b-a} & \text{if } a \le t < b \\ 1 & \text{if } t \ge b \end{cases}$$

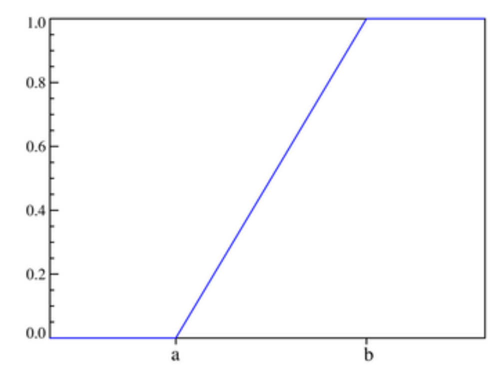

Exponential distribution

- A continuous random variable X is exponential with parameter $\lambda > 0$ if the density function is given by:

$$f(t) = \begin{cases} \lambda \cdot e^{-\lambda \cdot t} & \text{if } t > 0 \\ 0 & \text{otherwise} \end{cases}$$

$\boxed{\lambda = \text{“rate”}}$

- Cumulative distribution function (for $t \geq 0$):

$$F(t) = \Pr(X \leq t) = \int_0^t \lambda \cdot e^{-\lambda \cdot x} dx = [-e^{-\lambda \cdot x}]_0^t = 1 - e^{-\lambda \cdot t}$$

- Other properties:
  - negation: $\Pr(X > t) = e^{-\lambda \cdot t}$
  - mean (expectation): $E[X] = \int_0^\infty x \cdot \lambda \cdot e^{-\lambda \cdot x} dx = \frac{1}{\lambda}$
  - variance: $\mathrm{Var}(X) = 1/\lambda^2$

Exponential distribution – Examples

Probability distribution function

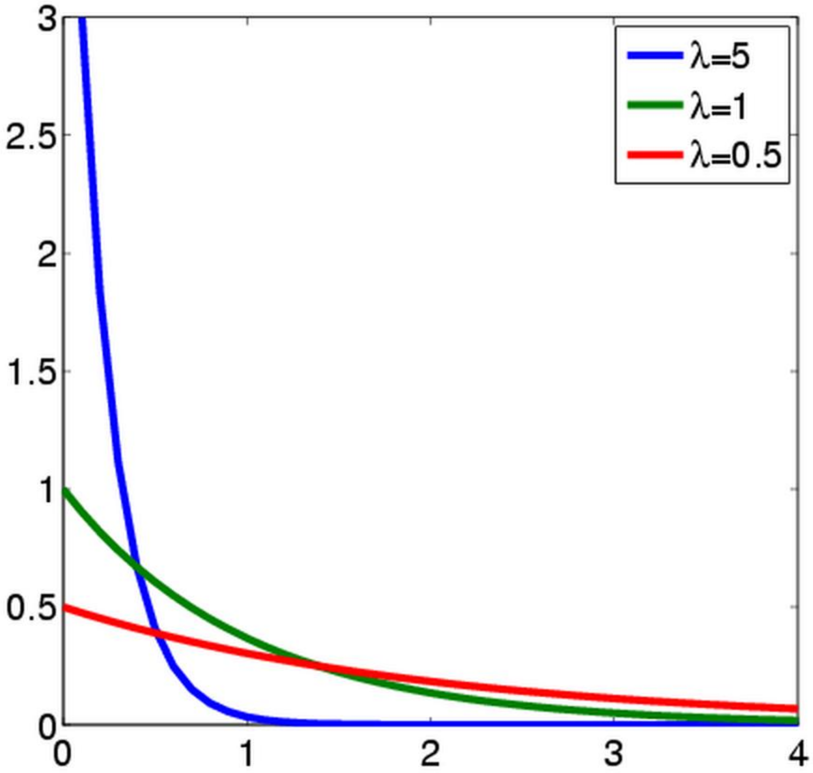

Cumulative distribution function

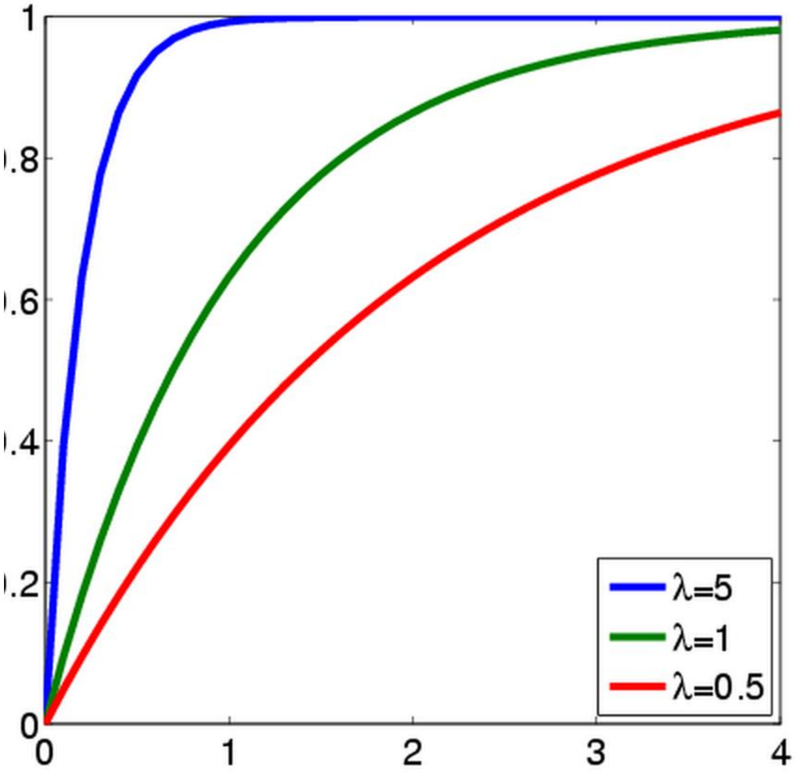

- The more λ increases, the faster the c.d.f. approaches 1

Exponential distribution

- Adequate for modelling many real-life phenomena
  - failures
    - e.g. time before machine component fails
  - inter-arrival times
    - e.g. time before next call arrives to a call centre
  - biological systems
    - e.g. times for reactions between proteins to occur

- Maximal entropy if just the mean is known
  - i.e. best approximation when only mean is known

- Can approximate general distributions arbitrarily closely
  - phase-type distributions

Exponential distribution – Properties

- Two useful properties of the exponential distribution:

- The exponential distribution is memoryless:
  - $\Pr(X>t_1+t_2 \mid X>t_1) = \Pr(X>t_2)$
  - it is the only memoryless continuous distribution
  - the discrete-time equivalent is the geometric distribution

- The minimum of two independent exponential distributions is an exponential distribution (parameter is sum)
  - $X_1 \sim \text{Exponential}(\lambda_1)$, $X_2 \sim \text{Exponential}(\lambda_2)$
  - $Y = \min(X_1,X_2) \sim \text{Exponential}(\lambda_1+\lambda_2)$
  - generalises to minimum of n distributions

Overview (Part 3)

- Exponential distribution and its properties
- Continuous-time Markov chains (CTMCs)
  - definition, race conditions, examples
  - paths and probability spaces
- CSL: A temporal logic for CTMCs
- CSL model checking
  - uniformisation, steady-state probabilities
- Extensions: Costs & rewards

Continuous-time Markov chains

- Continuous-time Markov chains (CTMCs)
  - labelled transition systems augmented with rates
  - continuous time delays, exponentially distributed

- Formally, a CTMC C is a tuple $(S, s_{init}, \mathbf{R}, L)$ where:
  - S is a finite set of states (“state space”)
  - $s_{init} \in S$ is the initial state
  - $\mathbf{R} : S \times S \rightarrow \mathbb{R}_{\geq 0}$ is the transition rate matrix
  - $L : S \rightarrow 2^{AP}$ is a labelling with atomic propositions

- Transition rate matrix assigns rates to each pair of states
  - used as a parameter to the exponential distribution
  - transition between s and s’ when $\mathbf{R}(s,s')>0$
  - probability triggered before t time units: $1 - e^{-\mathbf{R}(s,s')\cdot t}$

Simple CTMC example

- Modelling a queue of jobs
  - initially the queue is empty
  - jobs arrive with rate 3/2 (i.e. mean inter-arrival time is 2/3)
  - jobs are served with rate 3 (i.e. mean service time is 1/3)
  - maximum size of the queue is 3
  - state space: $S = \{s_i\}_{i=0..3}$ where s_i indicates i jobs in queue

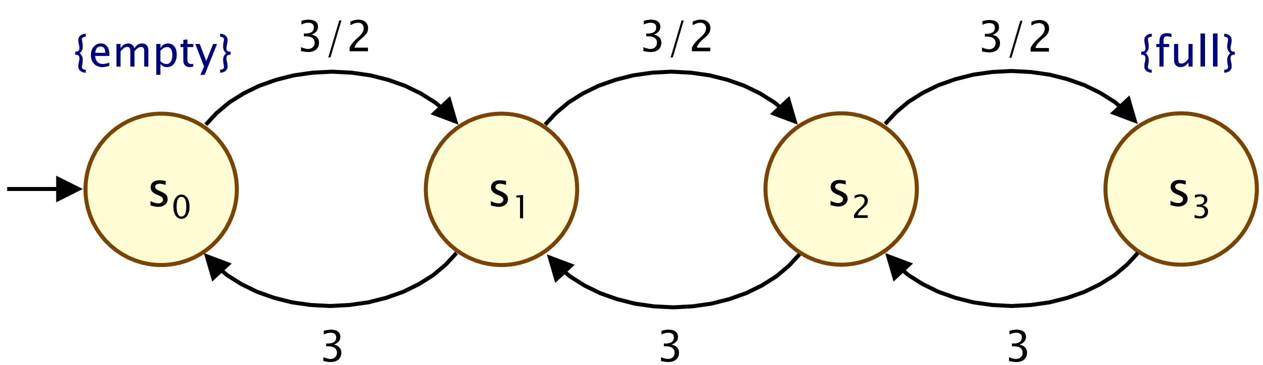

Race conditions

- What happens when there exists multiple s' with **R**(s,s')>0?
  - race condition: first transition triggered determines next state
  - two questions:
  - 1. How long is spent in s before a transition occurs?
  - 2. Which transition is eventually taken?

- 1. Time spent in a state before a transition
  - minimum of exponential distributions
  - exponential with parameter given by summation:

$$E(s) = \sum_{s' \in S} R(s, s')$$

  - probability of leaving a state s within [0,t] is $1 - e^{-E(s) \cdot t}$
  - E(s) is the exit rate of state s
  - s is called absorbing if E(s)=0 (no outgoing transitions)

Race conditions…

- 2. Which transition is taken from state s?
  - the choice is independent of the time at which it occurs
  - e.g. if $X_1 \sim \text{Exponential}(\lambda_1)$, $X_2 \sim \text{Exponential}(\lambda_2)$
  - then the probability that $X_1 < X_2$ is $\lambda_1/(\lambda_1+\lambda_2)$
  - more generally, the probability is given by…

- The embedded DTMC: $emb(C)=(S,s_{init},\mathbf{P}^{emb(C)},L)$
  - state space, initial state and labelling as the CTMC
  - for any $s,s' \in S$

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

- Probability that next state from s is s' given by $\mathbf{P}^{emb(C)}(s,s')$

Two interpretations of a CTMC

- Consider a (non-absorbing) state $s \in S$ with multiple outgoing transitions, i.e. multiple $s' \in S$ with $\mathbf{R}(s,s')>0$

- 1. Race condition
  - each transition triggered after exponentially distributed delay
    - probability triggered before t time units: $1 - e^{-\mathbf{R}(s,s')\cdot t}$
  - first transition triggered determines the next state

- 2. Separate delay/transition
  - remain in s for delay exponentially distributed with rate E(s)
    - i.e. probability of taking an outgoing transition from s within [0,t] is given by $1-e^{-E(s)\cdot t}$
  - probability that next state is s' is given by $\mathbf{P}^{emb(C)}(s,s')$
    - i.e. $\mathbf{R}(s,s')/E(s) = \mathbf{R}(s,s') \ / \ \Sigma_{s'\in S} \mathbf{R}(s,s')$

Continuous-time Markov chains

- Infinitesimal generator matrix

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

- Alternative definition: a CTMC is:
  - a family of random variables $\{ X(t) \mid t \in \mathbb{R}_{\geq 0} \}$
  - $X(t)$ are observations made at time instant t
  - i.e. $X(t)$ is the state of the system at time instant t
  - which satisfies…

- Memoryless (Markov property)

$$P[X(t_k)=s_k \mid X(t_{k-1})=s_{k-1}, \ldots, X(t_0)=s_0] = P[X(t_k)=s_k \mid X(t_{k-1})=s_{k-1}]$$

Simple CTMC example…

$C = (S, s_{init}, R, L)$

$S = \{s_0, s_1, s_2, s_3\}$

$s_{init} = s_0$

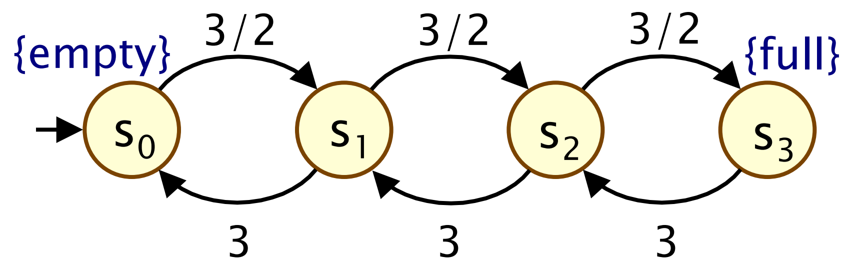

$AP = \{empty, full\}$

$L(s_0)=\{empty\}$, $L(s_1)=L(s_2)=\varnothing$ and $L(s_3)=\{full\}$

$$R = \begin{bmatrix} 0 & 3/2 & 0 & 0 \\ 3 & 0 & 3/2 & 0 \\ 0 & 3 & 0 & 3/2 \\ 0 & 0 & 3 & 0 \end{bmatrix} \quad P^{emb(C)} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 2/3 & 0 & 1/3 & 0 \\ 0 & 2/3 & 0 & 1/3 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad Q = \begin{bmatrix} -3/2 & 3/2 & 0 & 0 \\ 3 & -9/2 & 3/2 & 0 \\ 0 & 3 & -9/2 & 3/2 \\ 0 & 0 & 3 & -3 \end{bmatrix}$$

transition rate matrix

embedded DTMC

infinitesimal generator matrix

Example 2

- 3 machines, each can fail independently
  - failure rate λ, i.e. mean-time to failure (MTTF) = $1/\lambda$
  - modelled as exponential distributions
- One repair unit
  - repairs a single machine at rate μ (also exponential)
- State space:
  - $S = \{s_i\}_{i=0..3}$ where s_i indicates i machines operational

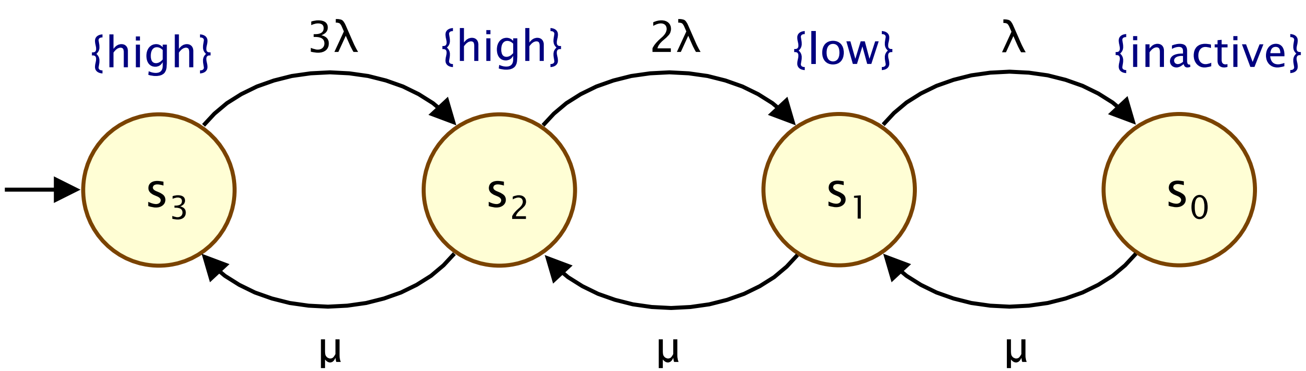

Example 3

- Chemical reaction system: two species A and B
- Two reactions:

$$A + B \underset{k_2}{\overset{k_1}{\rightleftarrows}} AB \qquad A \xrightarrow{k_3}$$

  - reversible reaction under which species A and B bind to form AB (forwards rate = $|A| \cdot |B| \cdot k_1$, backwards rate = $|AB| \cdot k_2$)
  - degradation of A (rate $|A| \cdot k_3$)
  - $|X|$ denotes number of molecules of species X

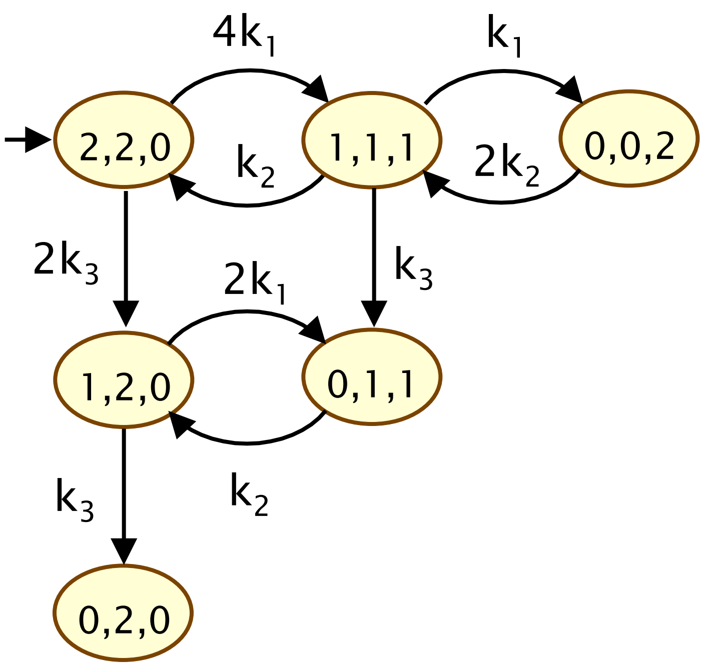

- CTMC with state space
  - $(|A|,|B|,|AB|)$
  - initially $(2,2,0)$

Paths of a CTMC

- An infinite path ω is a sequence $s_0 t_0 s_1 t_1 s_2 t_2 \ldots$ such that
  - $R(s_i,s_{i+1}) > 0$ and $t_i \in \mathbb{R}_{>0}$ for all $i \in \mathbb{N}$
  - amount of time spent in the jth state: $time(\omega,j)=t_j$
  - state occupied at time t: $\omega@t=s_j$

    where j smallest index such that $\Sigma_{i\leq j}\, t_j \geq t$
- A finite path is a sequence $s_0 t_0 s_1 t_1 s_2 t_2 \ldots t_{k-1} s_k$ such that
  - $R(s_i,s_{i+1}) > 0$ and $t_i \in \mathbb{R}_{>0}$ for all $i<k$
  - s_k is absorbing ($R(s,s') = 0$ for all $s' \in S$)
  - amount of time spent in the ith state only defined for $j\leq k$:

    $time(\omega,j)=t_j$ if $j<k$ and $time(\omega,j)=\infty$ if $j=k$
  - state occupied at time t: if $t\leq\Sigma_{i\leq k}\, t_j$ then $\omega@t$ as above
    otherwise $t>\Sigma_{i\leq k}\, t_j$ then $\omega@t=s_k$

- Path(s) denotes all infinite and finite paths starting in s

Recall: Probability spaces

- A σ-algebra (or σ-field) on Ω is a family Σ of subsets of Ω closed under complementation and countable union, i.e.:
  - if $A \in \Sigma$, the complement $\Omega \setminus A$ is in Σ
  - if $A_i \in \Sigma$ for $i \in \mathbb{N}$, the union $\cup_i A_i$ is in Σ
  - the empty set $\varnothing$ is in Σ
- Elements of Σ are called measurable sets or events
- Theorem: For any family F of subsets of Ω, there exists a unique smallest σ-algebra on Ω containing F
- Probability space (Ω, Σ, Pr)
  - Ω is the sample space
  - Σ is the set of events: σ-algebra on Ω
  - $Pr : \Sigma \to [0,1]$ is the probability measure:

    $Pr(\Omega) = 1$ and $Pr(\cup_i A_i) = \Sigma_i \, Pr(A_i)$ for countable disjoint A_i

Probability space

- Sample space: Path(s) (set of all paths from a state s)
- Events: sets of infinite paths
- Basic events: cylinders
  - cylinders = sets of paths with common finite prefix
  - include time intervals in cylinders

- Cylinder is a sequence $s_0,I_0,s_1,I_1,\ldots,I_{n-1},s_n$
  - $s_0,s_1,s_2,\ldots,s_n$ sequence of states where $\mathbf{R}(s_i,s_{i+1})>0$ for $i<n$
  - $I_0,I_1,I_2,\ldots,I_{n-1}$ sequence of of nonempty intervals of $\mathbb{R}_{\geq 0}$

- $\mathrm{Cyl}(s_0,I_0,s_1,I_1,\ldots,I_{n-1},s_n)$ set of (infinite and finite paths):
  - $\omega(i)=s_i$ for all $i \leq n$ and $\mathrm{time}(\omega,i) \in I_i$ for all $i < n$

Probability space

- Define measure over cylinders by induction
- $Pr_s(Cyl(s))=1$
- $Pr_s(Cyl(s,I,s_1,I_1,\ldots,I_{n-1},s_n,I',s'))$ equals:

$$Pr_s(Cyl(s,I,s_1,I_1,\ldots,I_{n-1},s_n)) \cdot P^{emb(C)}(s_n,s') \cdot \left(e^{-E(s_n)\cdot \inf I'} - e^{-E(s_n)\cdot \sup I'}\right)$$

probability transition from s_n to s' (defined using embedded DTMC)

probability time spent in state s_n is within the interval I'

Probability space

- Probability space (Path(s), $\Sigma_{Path(s)}$, Pr_s) [BHHK03]
- Sample space Ω = Path(s) (infinite and finite paths)
- Event set $\Sigma_{Path(s)}$
  - least σ-algebra on Path(s) containing all cylinders sets $Cyl(s_0,I_0,\ldots,I_{n-1},s_n)$ where:
    - $s_0,\ldots,s_n$ ranges over all state sequences with $\mathbf{R}(s_i,s_{i+1})>0$ for all i
    - $I_0,\ldots,I_{n-1}$ ranges over all sequences of non-empty intervals in $\mathbb{R}_{\geq 0}$ (where intervals are bounded by rationals)
- Probability measure Pr_s
  - Pr_s extends uniquely from probability defined over cylinders

Probability space – Example

- Probability of leaving the initial state s_0 and moving to state s_1 within the first 2 time units of operation?

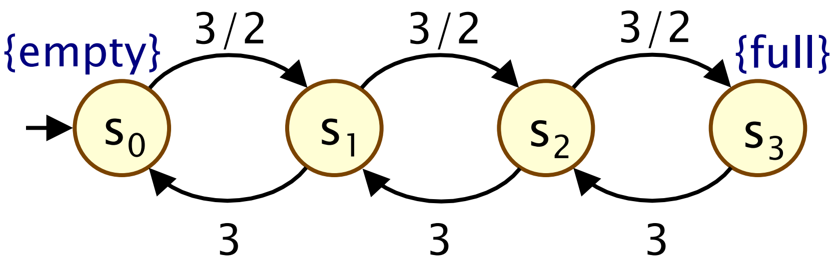

- Cylinder $Cyl(s_0,(0,2],s_1)$
- $Pr_{s0}(Cyl(s_0,(0,2],s_1))$

$$= Pr_{s0}(Cyl(s_0)) \cdot \mathbf{P}^{emb(C)}(s_0,s_1) \cdot (e^{-E(s0)\cdot 0} - e^{-E(s0)\cdot 2})$$

$$= 1 \cdot 1 \cdot (e^{-3/2\cdot 0} - e^{-3/2\cdot 2})$$

$$= 1 - e^{-3}$$

$$\approx 0.95021$$

Transient and steady-state behaviour

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

Overview (Part 3)

- Exponential distribution and its properties
- Continuous-time Markov chains (CTMCs)
  - definition, race conditions, examples
  - paths and probability spaces
- CSL: A temporal logic for CTMCs
- CSL model checking
  - uniformisation, steady-state probabilities
- Extensions: Costs & rewards

CSL

- Temporal logic for describing properties of CTMCs
  - CSL = Continuous Stochastic Logic [ASSB00,BHHK03]
  - extension of (non-probabilistic) temporal logic CTL
  - transient, steady-state and path-based properties
- Key additions:
  - probabilistic operator P (like PCTL)
  - steady state operator S
- Example: down → $P_{>0.75}$ [¬fail $U^{\leq[1,2.5]}$ 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 $\mathbb{R}_{\geq 0}$, $p \in [0,1]$, and $\sim \in \{<,>,\leq,\geq\}$
  - unbounded until U is a special case: $\phi_1\ U\ \phi_2 \equiv \phi_1\ U^{[0,\infty)}\ \phi_2$

- Quantitative properties: $P_{=?}\,[\ \psi\ ]$ and $S_{=?}\,[\ \phi\ ]$
  - where P/S is the outermost operator

CSL semantics for CTMCs

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

  - $s \models a \quad \Leftrightarrow \quad a \in L(s)$
  - $s \models \Phi_1 \wedge \Phi_2 \quad \Leftrightarrow \quad s \models \Phi_1$ and $s \models \Phi_2$
  - $s \models \neg\Phi \quad \Leftrightarrow \quad s \models \Phi$ is false
  - $s \models P_{\sim p}\ [\psi] \quad \Leftrightarrow \quad \mathrm{Prob}(s, \psi) \sim p$
  - $s \models S_{\sim p}\ [\Phi] \quad \Leftrightarrow \quad \Sigma_{s' \models \Phi}\ \underline{\pi}_s(s') \sim p$

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
  - $\mathrm{Prob}(s, \psi) = \mathrm{Pr}_s \{\omega \in \mathrm{Path}_s \mid \omega \vDash \psi \}$

if ω(0) is absorbing, ω(1) not defined

- Semantics of path formulae:
  - for a path ω of the CTMC:
  - $\omega \vDash X\,\phi \quad \Leftrightarrow \quad \omega(1)$ 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 in the interval I where ϕ_2 is true and ϕ_1 is true at all preceding time instants

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

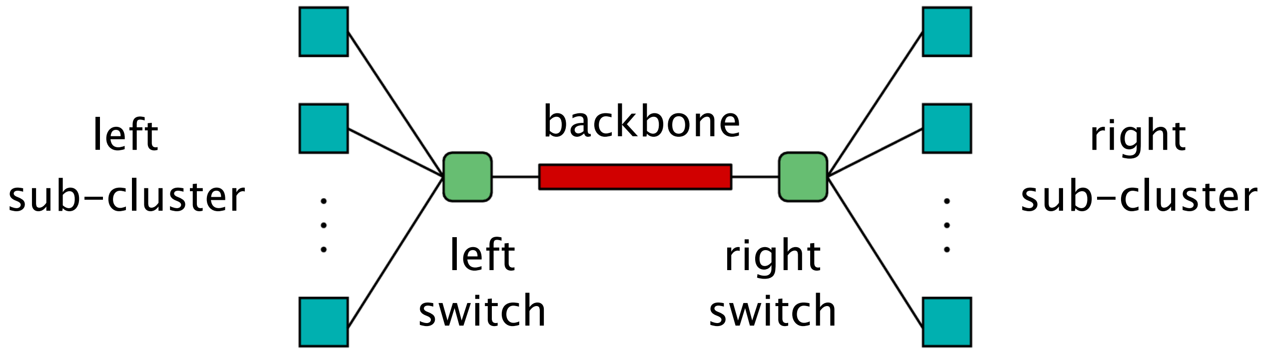

  - minimum QoS: at least ¾ of the workstations operational and connected via switches
  - premium QoS: all workstations operational and connected via switches

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
- $P_{<0.05}$ [$F^{[0,10]}$ $\neg$minimum]
  - the probability that the QoS drops below minimum within 10 hours is less than 0.05
- $\neg$minimum $\rightarrow$ $P_{<0.1}$ [$F^{[0,2]}$ $\neg$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

Overview (Part 3)

- Exponential distribution and its properties
- Continuous-time Markov chains (CTMCs)
  - definition, race conditions, examples
  - paths and probability spaces
- CSL: A temporal logic for CTMCs
- CSL model checking
  - uniformisation, steady-state probabilities
- Extensions: Costs & rewards

CSL model checking

- Model checking a CSL formula ϕ on a CTMC
  - basic algorithm proceeds by induction on parse tree of ϕ
  - non-probabilistic operators (true, a, $\neg$, $\wedge$) identical to PCTL

- Main task: computing probabilities for $P_{\sim p}$ [·] and $S_{\sim p}$ [·]

- Untimed properties can be verified on the embedded DTMC
  - properties of the form: $P_{\sim p}$ [X ϕ] or $P_{\sim p}$ [ϕ_1 U ϕ_2]
  - use algorithms for checking PCTL against DTMCs

- Which leaves…
  - time-bounded until operator: $P_{\sim p}$ [ϕ U^I ϕ]
  - steady-state operator: $S_{\sim p}$ [ϕ]

Model checking – Time-bounded until

- Compute Prob(s, $\phi_1 \, U^I \, \phi_2$) for all states where I is an arbitrary interval of the non-negative real numbers

- Note:
  - Prob(s, $\phi_1 \, U^I \, \phi_2$) = Prob(s, $\phi_1 \, U^{cl(I)} \, \phi_2$) where cl(I) denotes the closure of the interval I
  - Prob(s, $\phi_1 \, U^{[0,\infty)} \, \phi_2$) = $\mathrm{Prob}^{emb(C)}$(s, ϕ_1 U ϕ_2) where emb(C) is the embedded DTMC

- Therefore, 3 remaining cases to consider:
  - I = [0,t] for some $t \in \mathbb{R}_{\geq 0}$ (described in this lecture)
  - I = [t,t'] for some $t \leq t' \in \mathbb{R}_{\geq 0}$ or I = $[t,\infty)$ for some $t \in \mathbb{R}_{\geq 0}$

- Two methods: 1. Integral equations; 2. Uniformisation

Time-bounded until (integral equations)

- Computing the probabilities reduces to determining the least solution of the following set of integral equations:

- $\text{Prob}(s, \phi_1\ U^{[0,t]}\ \phi_2)$ equals
  - 1 if $s \in \text{Sat}(\phi_2)$,
  - 0 if $s \in \text{Sat}(\neg\phi_1 \wedge \neg\phi_2)$
  - and otherwise equals

probability of moving from s to s' at time x

probability, in state s', of satisfying until before t-x time units elapse

$$\int_0^t \sum_{s' \in S} \left(P^{emb(C)}(s,s') \cdot E(s) \cdot e^{-E(s) \cdot x} \right) \cdot \text{Prob}(s', \phi_1\ U^{[0,t-x]}\ \phi_2)\ dx$$

- One possibility: solve these integrals numerically
  - e.g. trapezoidal, Simpson and Romberg integration
  - expensive, possible problems with numerical stability

Time-bounded until (uniformisation)

- Reduction to transient analysis…
  - on a modified CTMC C'

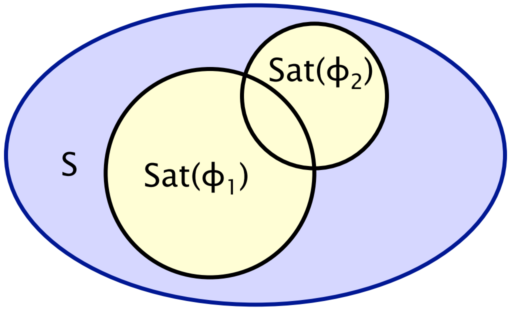

- Make all ϕ_2 states absorbing
  - in such a state $\phi_1\ U^{[0,x]}\ \phi_2$ holds with probability 1

- Make all $\neg\phi_1 \wedge \neg\phi_2$ states absorbing
  - in such a state $\phi_1\ U^{[0,x]}\ \phi_2$ holds with probability 0

- Formally: modified CTMC C' = $C[\phi_2][\neg\phi_1 \wedge \neg\phi_2]$
  - where for CTMC $C=(S,s_{init},\mathbf{R},L)$, let $C[\theta]=(S,s_{init},\mathbf{R}[\theta],L)$ where $\mathbf{R}[\theta](s,s')=\mathbf{R}(s,s')$ if $s \notin Sat(\theta)$ and 0 otherwise

Time-bounded until (uniformisation)

- Problem then reduces to calculating transient probabilities in the modified CTMC C' :

$$\mathrm{Prob}(s, \phi_1\ U^{[0,t]}\ \phi_2) \ = \ \sum_{s' \in Sat(\phi_2)} \underline{\pi}^{C'}_{s,t}(s')$$

$\underline{\pi}_{s,t}{}^{C'}(s')$:
transient probability in C':
starting in state s,
the probability of being
in state s' at time t

- To compute for all states s:

$$\underline{\mathrm{Prob}}(\phi_1\ U^{[0,t]}\ \phi_2) = \Pi^{C'}_t \cdot \underline{\phi_2}$$

  - where $\underline{\phi_2}$ is a 0-1 vector characterising ϕ_2
  - and $\Pi^{C'}_t$ is the matrix of all transient probabilities in C'

Computing transient probabilities

- Π_t – matrix of transient probabilities
  - $\Pi_t(s,s')=\underline{\pi}_{s,t}(s')$

- Π_t solution of the differential equation: $\Pi_t' = \Pi_t \cdot Q$
  - Q 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
  - probabilities instead computed using uniformisation

Uniformisation

- Uniformised DTMC unif(C) of CTMC C $=(S,s_{init},\mathbf{R},L)$:
  - $\text{unif(C)} = (S,s_{init},\mathbf{P}^{\text{unif(C)}},L)$
  - set of states, initial state and labelling the same as C
  - $\mathbf{P}^{\text{unif(C)}} = \mathbf{I} + \mathbf{Q}/q$
  - $\mathbf{I}$ 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:

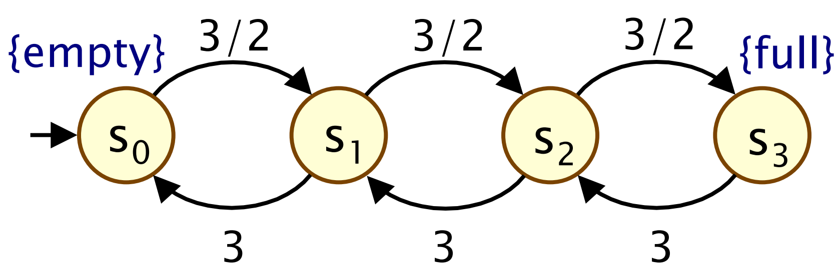

$$\mathbf{R} = \begin{bmatrix} 0 & 3/2 & 0 & 0 \\ 3 & 0 & 3/2 & 0 \\ 0 & 3 & 0 & 3/2 \\ 0 & 0 & 3 & 0 \end{bmatrix}$$

- Uniformised DTMC unif(C)
  - let uniformisation rate $q = \max_s \{ E(s) \} = 4.5$
  - $\mathbf{P}^{\mathrm{unif(C)}} = \mathbf{I} + \mathbf{Q}/q$

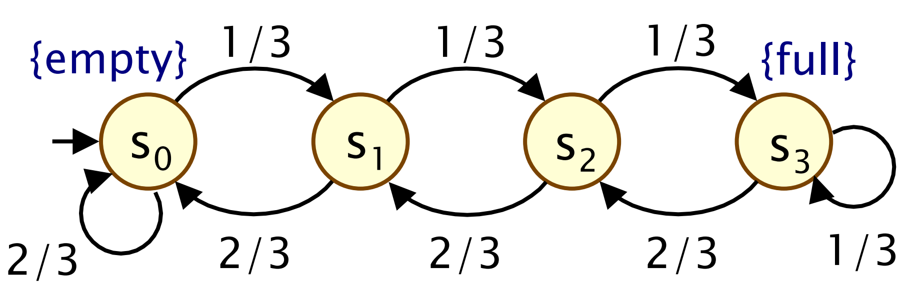

$$\mathbf{P}^{\mathrm{unif(C)}} = \begin{bmatrix} 2/3 & 1/3 & 0 & 0 \\ 2/3 & 0 & 1/3 & 0 \\ 0 & 2/3 & 0 & 1/3 \\ 0 & 0 & 2/3 & 1/3 \end{bmatrix}$$

Uniformisation

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

$$\Pi_t = e^{Q\cdot t} = e^{q\cdot(P^{unif(C)} - I)\cdot t} = e^{(q\cdot t)\cdot P^{unif(C)}} \cdot e^{-q\cdot t}$$

$$= e^{-q\cdot t} \cdot \left(\sum_{i=0}^{\infty} \frac{(q\cdot t)^i}{i!} \cdot \left(P^{unif(C)}\right)^i \right)$$

$$= \sum_{i=0}^{\infty} \left(e^{-q\cdot t} \cdot \frac{(q\cdot t)^i}{i!} \right) \left(P^{unif(C)}\right)^i$$

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

ith Poisson probability with parameter q·t

$P^{unif(C)}$ stochastic (all entries in [0,1] & rows sum to 1), therefore computations with **P** 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 the techniques of Fox and Glynn [FG88], which allow efficient computation of the Poisson probabilities

Time-bounded until (uniformisation)

- Recall that for model checking, we require:

$$\underline{\text{Prob}}(\phi_1\ \mathrm{U}^{[0,t]}\ \phi_2) = \Pi_t^{C'} \cdot \underline{\phi_2}$$

- So, using uniformisation:

$$\underline{\text{Prob}}(\phi_1\ \mathrm{U}^{[0,t]}\ \phi_2) = \sum_{i=0}^{\infty} \left(\gamma_{q \cdot t,i} \cdot \left(\mathbf{P}^{\mathrm{unif}(C')} \right)^i \cdot \underline{\phi_2} \right)$$

- This can be computed efficiently using matrix-vector multiplication (avoiding matrix powers):

$$\left(\mathbf{P}^{\mathrm{unif}(C')} \right)^0 \cdot \underline{\phi_2} = \underline{\phi_2}$$

$$\left(\mathbf{P}^{\mathrm{unif}(C')} \right)^{i+1} \cdot \underline{\phi_2} = \mathbf{P}^{\mathrm{unif}(C')} \cdot \left(\left(\mathbf{P}^{\mathrm{unif}(C')} \right)^i \cdot \underline{\phi_2} \right)$$

Time-bounded until – Example

- $P_{>0.65}$ [$F^{[0,7.5]}$ full] $\equiv$ $P_{>0.65}$ [true $U^{[0,7.5]}$ full]
  - "probability of the queue becoming full within 7.5 time units"
- State s_3 satisfies full and no states satisfy ¬true
  - in C[full][¬true ∧¬ full] only state s_3 made absorbing

$$\begin{bmatrix} 2/3 & 1/3 & 0 & 0 \\ 2/3 & 0 & 1/3 & 0 \\ 0 & 2/3 & 0 & 1/3 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

matrix of unif(C[full][¬true ∧¬full]) with uniformisation rate $\max_{s \in S} E(s)$ =4.5

s_3 made absorbing

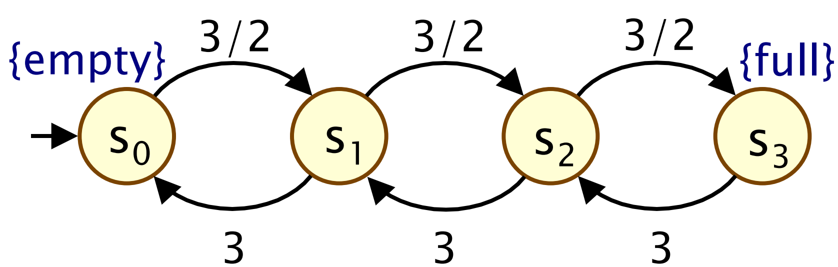

Time-bounded until – Example

- Computing the summation of matrix-vector multiplications

$$\underline{\text{Prob}}(\phi_1\ U^{[0,t]}\ \phi_2) = \sum_{i=0}^{\infty}\left(\gamma_{q\cdot t,i} \cdot \left(\mathbf{P}^{\text{unif}(C')} \right)^i \cdot \underline{\phi_2} \right)$$

  – yields $\underline{\text{Prob}}(\ F^{[0,7.5]}\ \text{full}\ ) \approx [\ 0.6482,\ 0.6823,\ 0.7811,\ 1\ ]$

- $P_{>0.65}[\ F^{[0,7.5]}\ \text{full}\ ]$ satisfied in states s_1, s_2 and s_3

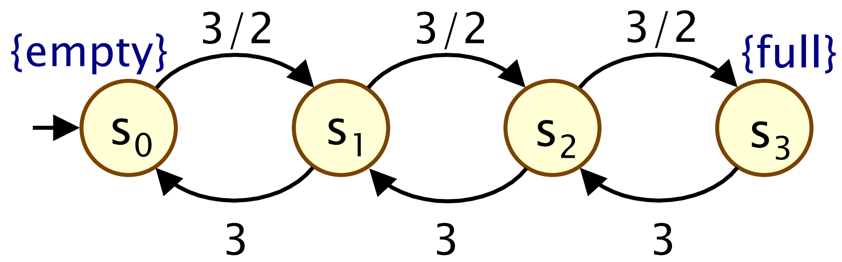

Model Checking – Steady-state

- A state s satisfies the formula $S_{\sim p}[\phi]$ if $\Sigma_{s' \models \phi}\ \underline{\pi}^{C}_{s}(s') \sim p$
  - $\underline{\pi}^{C}_{s}(s')$ is the probability, having started in state s, of being in state s' in the long run
  - thus model checking reduces to computing and then summing steady-state probabilities for the CTMC

- Steady-state probabilities: $\underline{\pi}^{C}_{s}(s') = \lim_{t\rightarrow\infty} \underline{\pi}^{C}_{s,t}(s')$
  - limit exists for all finite CTMCs
  - need to consider underlying graph structure of CTMC
  - i.e. its bottom strongly connected components (BSCCs)
  - irreducible CTMC (comprises one BSCC)
    - solution of one linear equation system
  - reducible CTMC (multiple BSCCs)
    - solve for each BSCC, combine results

Irreducible CTMCs

- For an irreducible CTMC:
  - the steady-state probabilities are independent of the starting state: denote the 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 **Q** is the infinitesimal generator matrix of C

- Solved by standard means:
  - direct methods, such as Gaussian elimination
  - iterative methods, such as Jacobi and Gauss-Seidel

Balance equations

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

balance the rate of leaving and entering a state

normalisation

For all $s \in S$:

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

$$\Leftrightarrow$$

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

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

Steady-state – Example

- Model check $S_{<0.1}$[full] on CTMC:

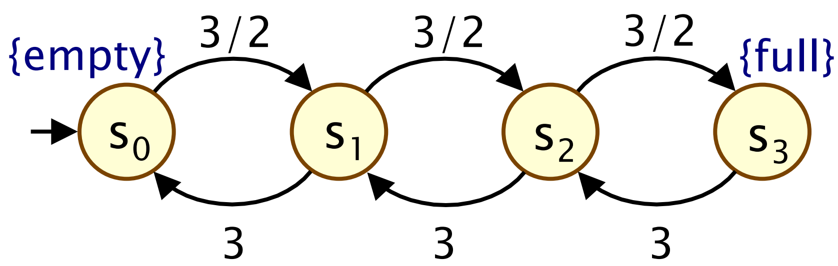

$$\mathbf{Q} = \begin{bmatrix} -3/2 & 3/2 & 0 & 0 \\ 3 & -9/2 & 3/2 & 0 \\ 0 & 3 & -9/2 & 3/2 \\ 0 & 0 & 3 & -3 \end{bmatrix}$$

- CTMC is irreducible (comprises a single BSCC)
  - steady state probabilities independent of starting state

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

Steady-state – Example

- Model check $S_{<0.1}$[full] on CTMC:

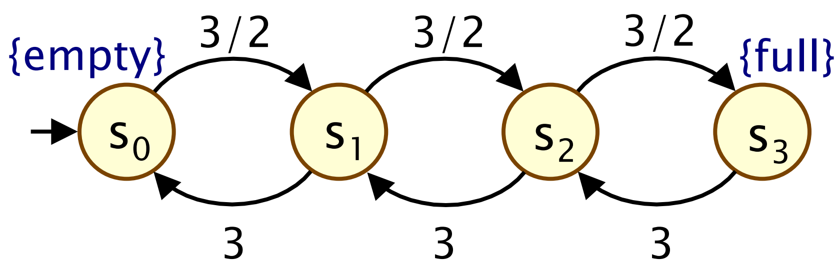

- Solve:

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

  - solution: $\underline{\pi}$ = [8/15, 4/15, 2/15, 1/15]
  - $\Sigma_{s' \vDash Sat(full)}\ \underline{\pi}(s') = 1/15 < 0.1$
  - so all states satisfy $S_{<0.1}$[full]

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 $Prob^{emb(C)}(s, F\ T)$ of reaching each T from s

- Then:

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

CSL model checking complexity

- For CSL model checking of a CTMC, complexity is:
  - linear in $|\Phi|$ and polynomial in $|S|$
  - linear in $q \cdot t_{max}$ (t_{max} is maximum finite bound in intervals)
- Unbounded until ($P_{\sim p}[\Phi_1 \ U^{[0,\infty)} \ \Phi_2]$) and steady-state ($S_{\sim p}[\Phi]$)
  - require solution of linear equation system of size $|S|$
  - can be solved with Gaussian elimination: cubic in $|S|$
  - precomputation algorithms (max $|S|$ steps)
- Time-bounded until ($P_{\sim p}[\Phi_1 \ U^I \ \Phi_2]$)
  - at most two iterative sequences of matrix-vector products
  - operation is quadratic in the size of the matrix, i.e. $|S|$
  - total number of iterations bounded by Fox and Glynn
  - the bound is linear in the size of $q \cdot t$ (q uniformisation rate)

Overview (Part 3)

- Exponential distribution and its properties
- Continuous-time Markov chains (CTMCs)
  - definition, race conditions, examples
  - paths and probability spaces
- CSL: A temporal logic for CTMCs
- CSL model checking
  - uniformisation, steady-state probabilities
- Extensions: Costs & rewards

Rewards (or costs)

- Like DTMCs, we can augment CTMCs with rewards
  - real-valued quantities assigned to states and/or transitions
  - can be interpreted in two ways: instantaneous/cumulative
  - properties considered here: expected value of rewards
  - formal property specifications in an extension of CSL

- For a CTMC $(S,s_{init},\mathbf{R},L)$, a reward structure is a pair $(\underline{\rho},\iota)$
  - $\underline{\rho} : S \rightarrow \mathbb{R}_{\geq 0}$ is a vector of state rewards
  - $\iota : S \times S \rightarrow \mathbb{R}_{\geq 0}$ is a matrix of transition rewards

- For cumulative reward-based properties of CTMCs
  - state rewards interpreted as rate at which reward gained
  - if the CTMC remains in state s for $t \in \mathbb{R}_{>0}$ time units, a reward of $t \cdot \underline{\rho}(s)$ is acquired

Reward structures – Examples

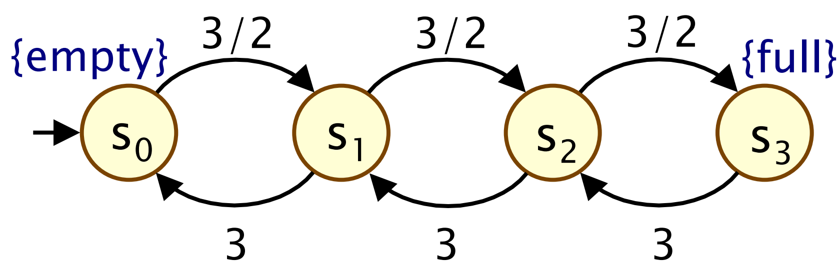

instantaneous

- Example: “size of message queue”
  - $\underline{\rho}(s_i)=i$ and $\iota(s_i,s_j)=0 \ \forall i,j$

cumulative

- Example: “time for which queue is not full”
  - $\underline{\rho}(s_i)=1$ for $i<3$, $\underline{\rho}(s_3)=0$ and $\iota(s_i,s_j)=0 \ \forall i,j$

Reward structures – Examples

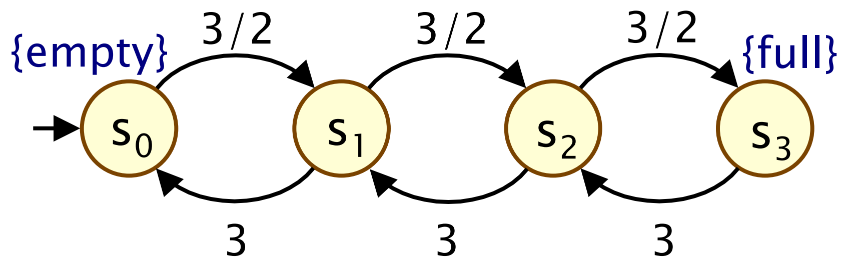

cumulative

- Example: “number of requests served”

$$\rho = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad \text{and} \quad \iota = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

CSL and rewards

- PRISM extends CSL to incorporate reward-based properties
  - adds R operator like the one added to PCTL

expected reward is ~r

  - ϕ ::= … | $R_{\sim r}$ [$I^{=t}$] | $R_{\sim r}$ [$C^{\leq t}$] | $R_{\sim r}$ [F ϕ] | $R_{\sim r}$ [S]

“instantaneous” “cumulative” “reachability” “steady-state”

  - where $r,t \in \mathbb{R}_{\geq 0}$, $\sim \in \{<,>,\leq,\geq\}$

- $R_{\sim r}$ [·] means “the expected value of · satisfies ~r”

Types of reward formulae

- Instantaneous: $R_{\sim r}$ [$I^{=t}$]
  - the expected value of the reward at time-instant t is ~r
  - "the expected queue size after 6.7 seconds is at most 2"
- Cumulative: $R_{\sim r}$ [$C^{\leq t}$]
  - the expected reward cumulated up to time-instant t is ~r
  - "the expected requests served within the first 4.5 seconds of operation is less than 10"
- Reachability: $R_{\sim r}$ [F ϕ]
  - the expected reward cumulated before reaching ϕ is ~r
  - "the expected requests served before the queue becomes full"
- Steady-state $R_{\sim r}$ [S]
  - the long-run average expected reward is ~r
  - "expected long-run queue size is at least 1.2"

Reward properties in PRISM

- Quantitative form:
  - e.g. $R_{=?}$ [$C^{\leq t}$]
  - what is the expected reward cumulated up to time-instant t?
- Add labels to R operator to distinguish between multiple reward structures defined on the same CTMC
  - e.g. $R_{\{num\_req\}=?}$ [$C^{\leq 4.5}$]
  - "the expected number of requests served within the first 4.5 seconds of operation"
  - e.g. $R_{\{pow\}=?}$ [$C^{\leq 4.5}$]
  - "the expected power consumption within the first 4.5 seconds of operation"

Reward formula semantics

- Formal semantics of the four reward operators:
  - $s \models R_{\sim r}[\, I^{=t} \,]$ $\Leftrightarrow$ $\mathrm{Exp}(s, X_{I=t}) \sim r$
  - $s \models R_{\sim r}[\, C^{\leq t} \,]$ $\Leftrightarrow$ $\mathrm{Exp}(s, X_{C\leq t}) \sim r$
  - $s \models R_{\sim r}[\, F\ \Phi \,]$ $\Leftrightarrow$ $\mathrm{Exp}(s, X_{F\Phi}) \sim r$
  - $s \models R_{\sim r}[\, S \,]$ $\Leftrightarrow$ $\lim_{t\rightarrow\infty}(\, 1/t \cdot \mathrm{Exp}(s, X_{C\leq t})\,) \sim r$

- where:
  - Exp(s, X) denotes the expectation of the random variable $X : \mathrm{Path}(s) \rightarrow \mathbb{R}_{\geq 0}$ with respect to the probability measure Pr_s

Reward formula semantics

- Definition of random variables:
  - path $\omega = s_0 t_0 s_1 t_1 s_2 \ldots$

state of ω at time t

$$X_{I=k}(\omega) = \underline{\rho}(\omega @ t)$$

time spent in state s_i

time spent in state s_{j_t} before t time units have elapsed

$$X_{C \le t}(\omega) = \sum_{i=0}^{j_t - 1} \big(t_i \cdot \underline{\rho}(s_i) + \iota(s_i, s_{i+1}) \big) + \left(t - \sum_{i=0}^{j_t - 1} t_i \right) \cdot \underline{\rho}(s_{j_t})$$

$$X_{F\phi}(\omega) = \begin{cases} 0 & \text{if } s_0 \in Sat(\phi) \\ \infty & \text{if } s_i \notin Sat(\phi) \text{ for all } i \geq 0 \\ \sum_{i=0}^{k_\phi - 1} t_i \cdot \underline{\rho}(s_i) + \iota(s_i, s_{i+1}) & \text{otherwise} \end{cases}$$

  - where $j_t = \min\{ j \mid \Sigma_{i \le j}\, t_i \geq t \}$ and $k_\phi = \min\{ i \mid s_i \vDash \phi \}$

Model checking reward formulae

- Instantaneous: $R_{\sim r}$ [$I^{=t}$]
  - reduces to transient analysis (state of the CTMC at time t)
  - use uniformisation
- Cumulative: $R_{\sim r}$ [$C^{\leq t}$]
  - extends approach for time-bounded until
  - based on uniformisation
- Reachability: $R_{\sim r}$ [F ϕ]
  - can be computed on the embedded DTMC
  - reduces to solving a system of linear equations
- Steady-state: $R_{\sim r}$ [S]
  - similar to steady state formulae $S_{\sim r}$ [ϕ]
  - graph based analysis (compute BSCCs)
  - solve systems of linear equations (compute steady state probabilities of each BSCC)

Summary

- Exponential distribution
  - suitable for modelling failures, waiting times, reactions, …
  - nice mathematical properties
- Continuous-time Markov chains
  - transition delays modelled as exponential distributions
  - probability space over paths
- CSL: Continuous Stochastic Logic
  - extension of PCTL for properties of CTMCs
- CSL model checking
  - extension of PCTL model checking for DTMCs
  - uniformisation: efficient iterative method for transient prob.s

- Tomorrow: Probabilistic model checking in practice
  - PRISM, tool demo, counterexamples, bisimulation