Verifying Continuous-Time Markov Chains

Lecture 3+4: Continuous-Time Markov Chains

Joost-Pieter Katoen

RWTH Aachen University
Software Modeling and Verification Group

http://www-i2.informatik.rwth-aachen.de/i2/mvps11/

VTSA Summerschool, Liège, Belgium

September 21, 2011

Overview

Overview

Time in discrete-time Markov chains

The advance of time in DTMCs

- Time in a DTMC proceeds in discrete steps
- Two possible interpretations:
  1. accurate model of (discrete) time units
     - e.g., clock ticks in model of an embedded device
  2. time-abstract
     - no information assumed about the time transitions take
- State residence time is geometrically distributed

Continuous-time Markov chains

- dense model of time
- transitions can occur at any (real-valued) time instant
- state residence time is (negative) exponentially distributed

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:

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

  where f satisfies: $f(x) \geqslant 0$ for all x and $\int_{-\infty}^{\infty} f(x)\, dx = 1$
- $F_X(d)$ is the *(cumulative) probability distribution function*
- $f(x)$ is the *probability density function*

Negative exponential distribution

Density of 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 r.v. Y with rate $\lambda \in \mathbb{R}_{>0}$ is:

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

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

Variance and expectation

Let r.v. Y be exponentially distributed with rate $\lambda \in \mathbb{R}_{>0}$. Then:

- Expectation $E[Y] = \int_0^\infty x{\cdot}\lambda{\cdot}e^{-\lambda\cdot x}\, dx \;=\; \frac{1}{\lambda}$
- Variance $Var[Y] = \int_0^\infty (x - E[X])^2 \lambda{\cdot}e^{-\lambda\cdot x}\, dx \;=\; \frac{1}{\lambda^2}$

Exponential pdf and cdf

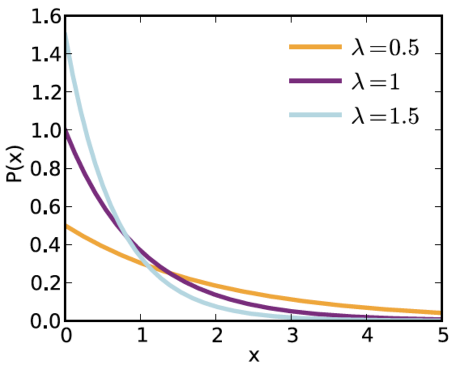

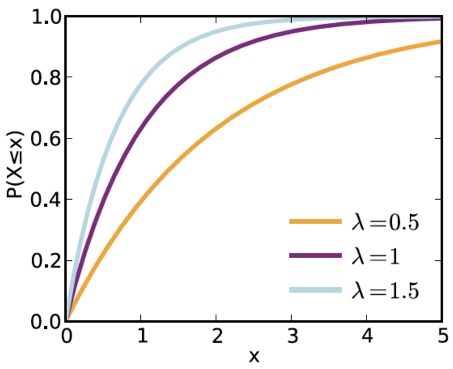

The higher λ, the faster the cdf approaches 1.

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 the *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

Theorem

1. For any exponentially distributed random variable X:

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

2. Any cdf which is memoryless is a negative exponential one.

Proof:

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 contraposition, using the total law of probability.

Closure under minimum

Minimum closure theorem

For independent, exponentially distributed random variables X and Y with rates $\lambda, \mu \in \mathbb{R}_{>0}$, the 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:

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

Closure under minimum

Minimum closure theorem for several exponentially distributed r.v.'s

For independent, exponentially distributed random variables $X_1, X_2, \ldots, X_n$ with rates $\lambda_1, \lambda_2, \ldots, \lambda_n \in \mathbb{R}_{>0}$ the r.v. $\min(X_1, X_2, \ldots, X_n)$ is exponentially distributed with rate $\sum_{0<i\leqslant n} \lambda_i$, i.e.,:

$$Pr\{\min(X_1, X_2, \ldots, X_n) \leqslant t\} \;=\; 1 - e^{-\sum_{0<i\leqslant n} \lambda_i \cdot t} \quad \text{for all } t \in \mathbb{R}_{\geqslant 0}.$$

Proof:

Generalization of the proof for the case of two exponential distributions.

Winning the race with two competitors

The minimum of two exponential distributions

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

The minimum of several exponentially distributed r.v.'s

For independent, exponentially distributed random variables $X_1, X_2, \ldots, X_n$ with rates $\lambda_1, \lambda_2, \ldots, \lambda_n \in \mathbb{R}_{>0}$ it holds:

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

Proof:

Generalization of the proof for the case of two exponential distributions.

Overview

Continuous-time Markov chain

Continuous-time Markov chain

A CTMC is a tuple $(S, \mathbf{P}, r, \iota_{\text{init}}, AP, L)$ where

- $(S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ is a DTMC, and
- $r : S \to \mathbb{R}_{>0}$, the exit-rate function

Interpretation

- residence time in state s is exponentially distributed with rate $r(s)$.
- phrased alternatively, the average residence time of state s is $\frac{1}{r(s)}$.
- thus, the higher the rate $r(s)$, the shorter the average residence time in s.

Example

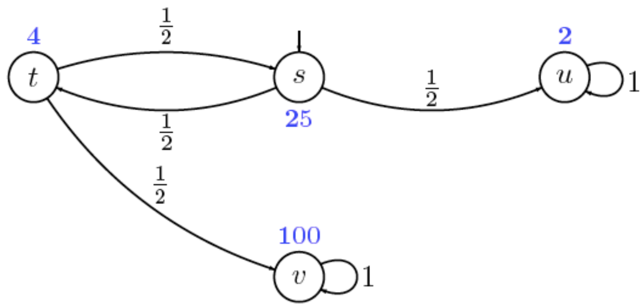

$r(s) = 25,\ r(t) = 4,\ r(u) = 2$ and $r(v) = 100$

Example: a classical perspective

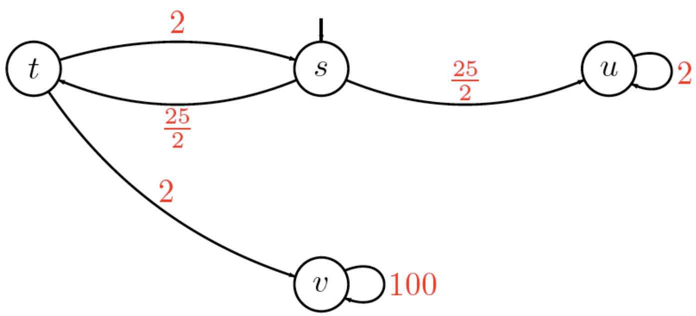

$r(s) = 25,\ r(t) = 4,\ r(u) = 2$ and $r(v) = 100$

The transition rate $\mathbf{R}(s, s') = \mathbf{P}(s, s') \cdot r(s)$

We use $(S, \mathbf{P}, r, \iota_{\text{init}}, AP, L)$ and $(S, \mathbf{R}, \iota_{\text{init}}, AP, L)$ interchangeably.

CTMC semantics by example

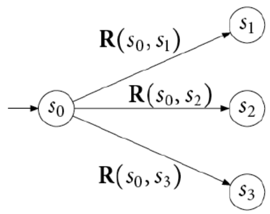

CTMC semantics

- Transition $s \to s'$:= 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

Enabledness

The probability that transition $s \to s'$ is *enabled* in $[0, t]$ is $1 - e^{-\mathbf{R}(s,s')\cdot t}$.

State-to-state timed transition probability

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

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

Residence time distribution

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

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

CTMC semantics

State-to-state timed transition probability

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

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

CTMC semantics

Residence time distribution

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

Kinetics

[edit]

Main article: Enzyme kinetics

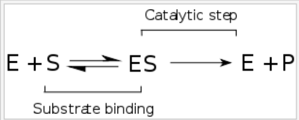

Mechanism for a single substrate enzyme catalyzed reaction. The enzyme (E) binds a substrate (S) and produces a product (P).

Enzyme kinetics is the investigation of how enzymes bind substrates and turn them into products. The rate data used in kinetic analyses are commonly obtained from enzyme assays, where since the 90s, the dynamics of many enzymes are studied on the level of individual molecules.

In 1902 Victor Henri[57] 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[58] 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 (termed also Michaelis-Menten kinetics).[59] Their work was further developed by G. E. Briggs and J. B. S. Haldane, who derived kinetic equations that are still widely considered today a starting point in solving enzymatic activity.[60]

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. Note that the simple Michaelis Menten mechanism for the enzymatic activity is considered today a basic idea, where many examples show that the enzymatic activity involves structural dynamics. This is incorporated in the enzymatic mechanism while introducing several Michaelis Menten pathways that are connected with fluctuating rates [44][45][46]. Nevertheless, there is a mathematical relation connecting the behavior obtained from the basic Michaelis Menten mechanism (that was indeed proved correct in many experiments) with the generalized Michaelis Menten mechanisms involving dynamics and activity; [61] this means that the measured activity of enzymes on the level of many enzymes may be explained with the simple Michaelis-Menten equation, yet, the actual activity of enzymes is richer and involves structural dynamics.

Source: wikipedia (June 2011)

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 $[t, t{+}\Delta)$:

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

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

- This process is a continuous-time Markov chain.

Enzyme-catalyzed substrate conversion as a CTMC

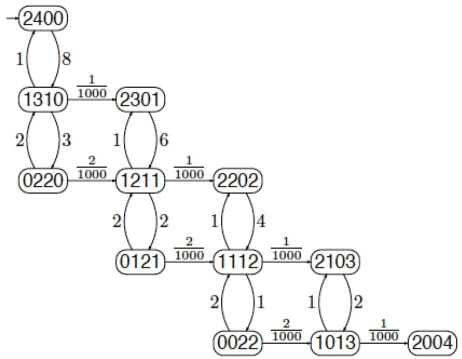

States:	*init*	*goal*
enzymes	2	2
substrates	4	0
complex	0	0
products	0	4

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

Summary

Main points

- Exponential distributions are closed under minimum.
- The probability to win a race amongst several exponential distributions only depends on their rates.
- A CTMC is a DTMC where state residence times are exponentially distributed.
- CTMC semantics distinguishes between enabledness and taking a transition.
- CTMCs are frequently used as semantical model for high-level formalisms.

Overview

Transient distribution of a CTMC

Transient state probability

Let $X(t)$ denote the state of a CTMC at time $t \in \mathbb{R}_{\geqslant 0}$. The probability to be in state s at time t is defined by:

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

Theorem: transient distribution as linear differential equation

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

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

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

Transient distribution theorem

Theorem: transient distribution as linear differential equation

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

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

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

Computing transient probabilities

The 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 standard knowledge yields: $\underline{p}(t) = \underline{p}(0) \cdot e^{(\mathbf{R}-\mathbf{r}) \cdot t}$.

Computing a matrix exponential

First attempt: use Taylor-Maclaurin expansion. This yields

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

But: numerical instability due to fill-in of $(\mathbf{R}-\mathbf{r})^i$ in presence of positive and negative entries in the matrix $\mathbf{R}-\mathbf{r}$.

Uniformization

Let CTMC $\mathcal{C} = (S, \mathbf{P}, r, \iota_{\text{init}}, AP, L)$ with S finite.

Uniform CTMC

CTMC $\mathcal{C}$ is uniform if $r(s) = r$ for all $s \in S$ for some $r \in \mathbb{R}_{>0}$.

Uniformization [Gross and Miller, 1984]

Let $r \in \mathbb{R}_{>0}$ such that $r \geqslant max_{s \in S}\ r(s)$. Then $unif(r, \mathcal{C})$ is the tuple $(S, \overline{\mathbf{P}}, \overline{r}, \iota_{\text{init}}, AP, L)$ with $\overline{r}(s) = r$ for all $s \in S$, and:

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

It follows that $\overline{\mathbf{P}}$ is a stochastic matrix and $unif(r, \mathcal{C})$ is a CTMC.

Uniformization: example

Uniformization

Let $r \in \mathbb{R}_{>0}$ such that $r \geqslant max_{s \in S}\ r(s)$. Then $unif(r, \mathcal{C}) = (S, \overline{\mathbf{P}}, \overline{r}, \iota_{\text{init}}, AP, L)$ with $\overline{r}(s) = r$ for all $s \in S$, and:

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

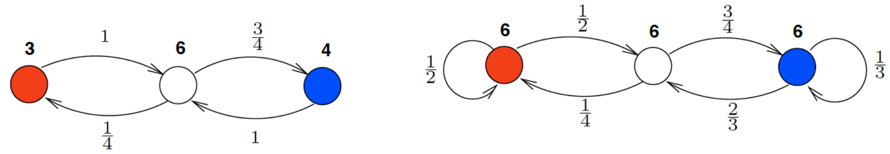

CTMC $\mathcal{C}$ and its uniformized counterpart $unif(6, \mathcal{C})$

Uniformization: intuition

Uniformization

Let $r \in \mathbb{R}_{>0}$ such that $r \geqslant max_{s \in S}\ r(s)$. Then $unif(r, \mathcal{C}) = (S, \overline{\mathbf{P}}, \overline{r}, \iota_{\text{init}}, AP, L)$ with $\overline{r}(s) = r$ for all $s \in S$, and:

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

Intuition

- Fix all exit rates to (at least) the maximal exit rate r occurring in CTMC $\mathcal{C}$.
- Thus, $\frac{1}{r}$ is the shortest mean residence time in the CTMC $\mathcal{C}$.
- Then normalize the residence time of all states with respect to r as follows:
  1. replace an average residence time $\frac{1}{r(s)}$ by a shorter (or equal) one, $\frac{1}{r}$
  2. decrease the transition probabilities by a factor $\frac{r(s)}{r}$, and
  3. increase the self-loop probability by a factor $\frac{r-r(s)}{r}$

  That is, slow down state s whenever $r(s) < r$.

Strong bisimulation on DTMCs

Probabilistic bisimulation [Larsen & Skou, 1989]

Let $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be a DTMC and $R \subseteq S \times S$ an equivalence. Then: R is a *probabilistic bisimulation* on S if for any $(s, t) \in R$:

1. $L(s) = L(t)$, and
2. $\mathbf{P}(s, C) = \mathbf{P}(t, C)$ for all equivalence classes $C \in S/R$

where $\mathbf{P}(s, C) = \sum_{s' \in C} \mathbf{P}(s, s')$.

For states in R, the probability of moving by a single transition to some equivalence class is equal.

Probabilistic bisimilarity

Let $\mathcal{D}$ be a DTMC and s, t states in $\mathcal{D}$. Then: s is *probabilistically bisimilar* to t, denoted $s \sim_p t$, if there exists a probabilistic bisimulation R with $(s, t) \in R$.

Strong bisimulation on CTMCs

Probabilistic bisimulation [Buchholz, 1994]

Let $\mathcal{C} = (S, \mathbf{P}, r, \iota_{\text{init}}, AP, L)$ be a CTMC and $R \subseteq S \times S$ an equivalence. Then: R is a *probabilistic bisimulation* on S if for any $(s, t) \in R$:

1. $L(s) = L(t)$, and
2. $r(s) = r(t)$, and
3. $\mathbf{P}(s, C) = \mathbf{P}(t, C)$ for all equivalence classes $C \in S/R$

The last two conditions amount to $\mathbf{R}(s, C) = \mathbf{R}(t, C)$ for all equivalence classes $C \in S/R$.

Probabilistic bisimilarity

Let $\mathcal{C}$ be a CTMC and s, t states in $\mathcal{C}$. Then: s is *probabilistically bisimilar* to t, denoted $s \sim_m t$, if there exists a probabilistic bisimulation R with $(s, t) \in R$.

Weak bisimulation on DTMCs

Weak probabilistic bisimulation [Baier & Hermanns, 1996]

Let $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be a DTMC and $R \subseteq S \times S$ an equivalence. Then: R is a *weak probabilistic bisimulation* on S if for any $(s, t) \in R$:

1. $L(s) = L(t)$, and
2. if $\mathbf{P}(s, [s]_R) < 1$ and $\mathbf{P}(t, [t]_R) < 1$, then:

$$\frac{\mathbf{P}(s, C)}{1 - \mathbf{P}(s, [s]_R)} = \frac{\mathbf{P}(t, C)}{1 - \mathbf{P}(t, [t]_R)} \quad \text{for all} C \in S/R, C \neq [s]_R = [t]_R.$$

3. s can reach a state outside $[s]_R$ iff t can reach a state outside $[t]_R$.

For states in R, the conditional probability of moving by a single transition to another equivalence class is equal. In addition, either all states in an equivalence class C almost surely stay there, or have an option to escape from C.

Weak bisimulation on DTMCs

Weak probabilistic bisimulation [Baier & Hermanns, 1996]

Let $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be a DTMC and $R \subseteq S \times S$ an equivalence. Then: R is a *weak probabilistic bisimulation* on S if for any $(s, t) \in R$:

1. $L(s) = L(t)$, and
2. if $\mathbf{P}(s, [s]_R) < 1$ and $\mathbf{P}(t, [t]_R) < 1$, then:

$$\frac{\mathbf{P}(s, C)}{1 - \mathbf{P}(s, [s]_R)} = \frac{\mathbf{P}(t, C)}{1 - \mathbf{P}(t, [t]_R)} \quad \text{for all} C \in S/R, C \neq [s]_R = [t]_R.$$

3. s can reach a state outside $[s]_R$ iff t can reach a state outside $[t]_R$.

Probabilistic weak bisimilarity

Let $\mathcal{D}$ be a DTMC and s, t states in $\mathcal{D}$. Then: s is *probabilistically weak bisimilar* to t, denoted $s \approx_p t$, if there exists a probabilistic weak bisimulation R with $(s, t) \in R$.

Weak bisimulation on DTMC: example

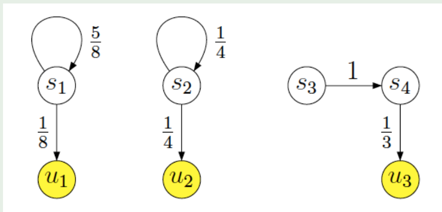

The equivalence relation R with $S/R = \{ \{s_1, s_2, s_3, s_4\}, \{u_1, u_2, u_3\} \}$ is a weak bisimulation. This can be seen as follows. For $C = \{ u_1, u_2, u_3 \}$ and s_1, s_2, s_4 with $\mathbf{P}(s_i, [s_i]_R) < 1$ we have:

$$\frac{\mathbf{P}(s_1, C)}{1 - \mathbf{P}(s_1, [s_1])} = \frac{1/8}{1-5/8} = \frac{1/4}{1-1/4} = \frac{\mathbf{P}(s_2, C)}{1 - \mathbf{P}(s_2, [s_2])} = \frac{1/3}{1} = \frac{\mathbf{P}(s_4, C)}{1 - \mathbf{P}(s_4, [s_4])}.$$

Note that $\mathbf{P}(s_3, [s_3]_R) = 1$. Since s_3 can reach a state outside $[s_3]$ as s_1, s_2 and s_4, it follows that $s_1 \approx_p s_2 \approx_p s_3 \approx_p s_4$.

Reachability condition

Remark

Consider the following DTMC:

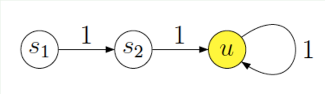

It is not difficult to establish $s_1 \approx s_2$. Note: $\mathbf{P}(s_1, [s_1]) = 1$, but $\mathbf{P}(s_2, [s_2]_R) < 1$. Both s_1 and s_2 can reach a state outside $[s_1]_R = [s_2]_R$. The reachability condition is essential to establish $s_1 \approx s_2$ and cannot be dropped: otherwise s_1 and s_2 would be weakly bisimilar to an equally labelled absorbing state.

Weak bisimulation on CTMCs

Weak probabilistic bisimulation **[Bravetti, 2002]**

Let $\mathcal{C} = (S, \mathbf{P}, r, \iota_{\text{init}}, AP, L)$ be a CTMC and $R \subseteq S \times S$ an equivalence. Then: R is a *weak probabilistic bisimulation* on S if for any $(s, t) \in R$:

1. $L(s) = L(t)$, and
2. $\mathbf{R}(s, C) = \mathbf{R}(t, C)$ for all $C \in S/R$ with $C \neq [s]_R = [t]_R$

Weak probabilistic bisimilarity

Let $\mathcal{C}$ be a CTMC and s, t states in $\mathcal{C}$. Then: s is *weak probabilistically bisimilar* to t, denoted $s \approx_m t$, if there exists a weak probabilistic bisimulation R with $(s, t) \in R$.

A useful lemma

Let $\mathcal{C}$ be a CTMC and R an equivalence relation on S with $(s, t) \in R$. Then: the following two statements are equivalent:

1. If $\mathbf{P}(s, [s]_R) < 1$ and $\mathbf{P}(t, [t]_R) < 1$ then for all $C \in S/R$, $C \neq [s]_R = [t]_R$:

$$\frac{\mathbf{P}(s, C)}{1 - \mathbf{P}(s, [s]_R)} = \frac{\mathbf{P}(t, C)}{1 - \mathbf{P}(t, [t]_R)} \quad \text{and} \quad \mathbf{R}(s, S \setminus [s]_R) = \mathbf{R}(t, S \setminus [t]_R)$$

2. $\mathbf{R}(s, C) = \mathbf{R}(t, C)$ for all $C \in S/R$ with $C \neq [s]_R = [t]_R$.

Proof:

Left as an exercise.

Weak bisimulation on CTMCs: example

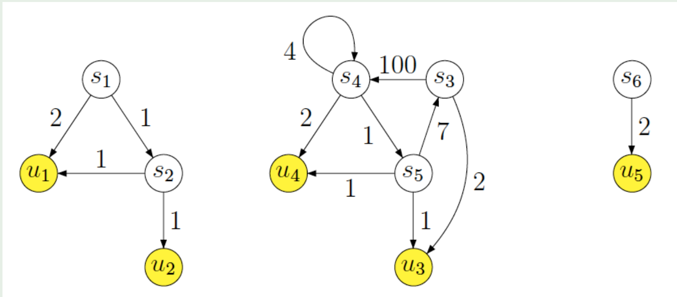

Equivalence relation R with $S/R = \{ \{s_1, s_2, s_3, s_4, s_5, s_6\}, \{u_1, u_2, u_3, u_4, u_5\} \}$ is a weak bisimulation on the CTMC depicted above. This can be seen as follows. For $C = \{ u_1, u_2, u_3, u_4, u_5 \}$, we have that all s-states enter C with rate 2. The rates between the s-states are not relevant.

Properties (without proof)

Strong and weak bisimulation in uniform CTMCs

For all uniform CTMCs $\mathcal{C}$ and states s, u in $\mathcal{C}$, we have:

$$s \sim_m u \quad \text{iff} \quad s \approx_m u \quad \text{iff} \quad s \sim_p u.$$

For any CTMC $\mathcal{C}$, we have: $\mathcal{C} \approx_m \mathit{unif}(r, \mathcal{C})$ with $r \geqslant \max_{s \in S} r(s)$.

Preservation of transient probabilities

For all CTMCs $\mathcal{C}$ with states s, u in $\mathcal{C}$ and $t \in \mathbb{R}_{\geqslant 0}$, we have:

$$s \approx_m u \quad \text{implies} \quad \underline{p}(t) = \underline{p}(t)$$

where $\underline{p}(0) = \mathbf{1}_s$ and $\underline{p}(0) = \mathbf{1}_u$ where $\mathbf{1}_s$ is the characteristic function for state s, i.e., $\mathbf{1}_s(s') = 1$ iff $s = s'$.

Computing transient probabilities

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

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

Standard knowledge yields: $\underline{p}(t) \;=\; \underline{p}(0) \cdot e^{(\mathbf{R} - \mathbf{r}) \cdot t}$.

As uniformization preserves transient probabilities, we replace $\mathbf{R} - \mathbf{r}$ by its variant for the uniformized CTMC, i.e., $\overline{\mathbf{R}} - \overline{\mathbf{r}}$. We have:

$$\overline{\mathbf{R}}(s, s') = \overline{\mathbf{P}}(s, s') \cdot \overline{r}(s) = \overline{\mathbf{P}}(s, s') \cdot r \quad \text{and} \quad \overline{\mathbf{r}} = \mathbf{I} \cdot r.$$

Thus:

$$\underline{p}(0) \cdot e^{(\overline{\mathbf{R}} - \overline{\mathbf{r}}) \cdot t} \;=\; \underline{p}(0) \cdot e^{(\overline{\mathbf{P}} \cdot r - \mathbf{I} \cdot r) \cdot t} \;=\; \underline{p}(0) \cdot e^{(\overline{\mathbf{P}} - \mathbf{I}) \cdot r \cdot t} \;=\; \underline{p}(0) \cdot e^{-rt} \cdot e^{r \cdot t \cdot \overline{\mathbf{P}}}.$$

Computing transient probabilities

$$\underline{p}(t) \;=\; \underline{p}(0) \cdot e^{(\overline{\mathbf{R}} - \overline{\mathbf{r}}) \cdot t} \;=\; \underline{p}(0) \cdot e^{(\overline{\mathbf{P}} \cdot r - \mathbf{I} \cdot r) \cdot t} \;=\; \underline{p}(0) \cdot e^{(\overline{\mathbf{P}} - \mathbf{I}) \cdot r \cdot t} \;=\; \underline{p}(0) \cdot e^{-rt} \cdot e^{r \cdot t \cdot \overline{\mathbf{P}}}.$$

Computing a matrix exponential

Exploit Taylor-Maclaurin expansion. This yields:

$$\underline{p}(0) \cdot e^{-rt} \cdot e^{r \cdot t \cdot \overline{\mathbf{P}}} \;=\; \underline{p}(0) \cdot e^{-rt} \cdot \sum_{i=0}^{\infty} \frac{(r \cdot t)^i}{i!} \cdot \overline{\mathbf{P}}^i \;=\; \underline{p}(0) \cdot \sum_{i=0}^{\infty} \underbrace{e^{-r \cdot t} \frac{(r \cdot t)^i}{i!}}_{\text{Poisson prob.}} \cdot \overline{\mathbf{P}}^i$$

As $\overline{\mathbf{P}}$ is a stochastic matrix, computing the matrix exponential $\overline{\mathbf{P}}^i$ is numerically stable.

Intermezzo: Poisson distribution

Poisson distribution

The Poisson distribution is a discrete probability distribution that expresses the probability of a given number i of events occurring in a fixed interval of time $[0, t]$ if these events occur with a known average rate r and independently of the time since the last event. Formally, the pdf is:

$$f(i; r \cdot t) \;=\; e^{-r \cdot t} \frac{(r \cdot t)^i}{i!}$$

where r is the mean of the Poisson distribution.

Remark

The Poisson distribution can be derived as a limiting case to the binomial distribution as the number of trials goes to infinity and the expected number of successes remains fixed.

Transient probabilities: example

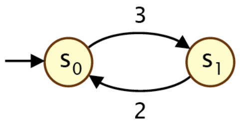

$$\mathbf{P} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \underline{r} = \begin{bmatrix} 3 \\ 2 \end{bmatrix} \text{ and } \overline{\mathbf{P}}_3 = \begin{bmatrix} 0 & 1 \\ \frac{2}{3} & \frac{1}{3} \end{bmatrix}$$

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

$$\begin{aligned} \underline{p}(1) &= \underline{p}(0) \cdot \sum_{i=0}^{\infty} e^{-3} \frac{3^i}{i!} \cdot \overline{\mathbf{P}}^i \\ &= (1,0) \cdot e^{-3} \tfrac{1}{0!} \cdot \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} + (1,0) \cdot e^{-3} \tfrac{3}{1!} \cdot \begin{bmatrix} 0 & 1 \\ \frac{2}{3} & \frac{1}{3} \end{bmatrix} \\ &\quad + (1,0) \cdot e^{-3} \tfrac{9}{2!} \cdot \begin{bmatrix} 0 & 1 \\ \frac{2}{3} & \frac{1}{3} \end{bmatrix}^2 + \ldots\ldots \\ &\approx (0.404043, 0.595957) \end{aligned}$$

Truncating the infinite sum

Computing transient probabilities

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

- Summation can be truncated *a priori* for a given error bound $\varepsilon > 0$.
- The *error* that is introduced by truncating at summand k_ε is:

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

- Strategy: choose k_ε minimal such that:

$$\sum_{i=k_{\varepsilon+1}}^{\infty} e^{-rt} \frac{(rt)^i}{i!} = \sum_{i=0}^{\infty} e^{-rt} \frac{(rt)^i}{i!} - \sum_{i=0}^{k_\varepsilon} e^{-rt} \frac{(rt)^i}{i!} = 1 - \sum_{i=0}^{k_\varepsilon} e^{-rt} \frac{(rt)^i}{i!} \leqslant \varepsilon$$

Summary

Main points

- Bisimilar states are equally labelled and their cumulative rate to any equivalence class coincides.
- Weak bisimilar states have equal conditional probabilities to move to some equivalence class, and can either both leave their class or both can't.
- Uniformization normalizes the exit rates of all states in a CTMC.
- Uniformization transforms a CTMC into a weak bisimilar one.
- Transient distribution are obtained by solving a system of linear differential equations.
- These equations can be solved conveniently on the uniformized CTMC.

Overview

1. Negative exponential distributions
2. What are continuous-time Markov chains?
3. Transient distribution
4. Timed reachability probabilities
5. Verifying continuous stochastic CTL
6. Verifying linear real-time properties

Paths in a CTMC

Timed paths

Paths in CTMC $\mathcal{C}$ are maximal (i.e., infinite) paths of alternating states and time instants:

$$\pi = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots$$

such that $s_i \in S$ and $t_i \in \mathbb{R}_{>0}$. Let $Paths(\mathcal{C})$ be the set of paths in $\mathcal{C}$ and $Paths^*(\mathcal{C})$ the set of finite prefixes thereof.

Time instant t_i is the amount of time spent in state s_i.

Notations

- Let $\pi[i] := s_i$ denote the $(i{+}1)$-st state along the timed path π.
- Let $\pi\langle i \rangle := t_i$ the time spent in state s_i.
- Let $\pi@t$ be the state occupied in π at time $t \in \mathbb{R}_{\geqslant 0}$, i.e. $\pi@t := \pi[i]$ where i is the smallest index such that $\sum_{j=0}^{i} \pi\langle j \rangle > t$.

Paths and probabilities

To reason quantitatively about the behavior of a CTMC, we need to define a probability space over its paths.

Intuition

For a given state s in CTMC $\mathcal{C}$:

- Sample space := set of all interval-timed paths $s_0\, I_0 \ldots I_{k-1}\, s_k$ with $s = s_0$
- Events := sets of interval-timed paths starting in s
- Basic events := cylinder sets
- Cylinder set of finite interval-timed paths := set of all infinite timed paths with a prefix in the finite interval-timed path

Probability measure on DTMCs

Cylinder set

Let $s_0, \ldots, s_k \in S$ with $\mathbf{P}(s_i, s_{i+1}) > 0$ for $0 \leqslant i < k$ and $I_0, \ldots, I_{k-1}$ non-empty intervals in $\mathbb{R}_{\geqslant 0}$. The *cylinder set* of $s_0\, I_0\, s_1\, I_1 \ldots I_{k-1}\, s_k$ is defined by:

$$Cyl(s_0, I_0, \ldots, I_{k-1}, s_k) = \{ \pi \in Paths(\mathcal{C}) \mid \forall 0 \leqslant i \leqslant k.\, \pi[i] = s_i \text{ and } i < k \Rightarrow \pi\langle i \rangle \in I_i \}$$

The cylinder set spanned by $s_0, I_0, \ldots, I_{k-1}, s_k$ thus consists of all infinite timed paths that have a prefix $\hat{\pi}$ that lies in $s_0, I_0, \ldots, I_{k-1}, s_k$. Cylinder sets serve as basic events of the smallest σ-algebra on $Paths(\mathcal{C})$.

σ-algebra of a CTMC

The σ-algebra associated with CTMC $\mathcal{C}$ is the smallest σ-algebra $\mathcal{F}(Paths(s_0))$ that contains all cylinder sets $Cyl(s_0, I_0, \ldots, I_{k-1}, s_k)$ where $s_0 \ldots s_k$ is a path in the state graph of $\mathcal{C}$ (starting in s_0) and $I_0, \ldots, I_{k-1}$ range over all sequences of non-empty intervals in $\mathbb{R}_{\geqslant 0}$.

Probability measure on CTMCs

Cylinder set

The *cylinder set* $Cyl(s_0, I_0, \ldots, I_{k-1}, s_k)$ of $s_0\, I_0 \ldots I_{k-1}\, s_k$ is defined by:

$$\{ \pi \in Paths(\mathcal{C}) \mid \forall 0 \leqslant i \leqslant k.\, \pi[i] = s_i \text{ and } i < k \Rightarrow \pi\langle i \rangle \in I_i \}$$

Probability measure

Pr is the unique *probability measure* on the σ-algebra $\mathcal{F}(Paths(s_0))$ defined by induction on k as follows: $Pr(Cyl(s_0)) = \iota_{\text{init}}(s_0)$ and for $k > 0$:

$$Pr(Cyl(s_0, I_0, \ldots, I_{k-1}, s_k)) = Pr(Cyl(s_0, I_0, \ldots, I_{k-2}, s_{k-1})) \cdot \int_{I_{k-1}} \mathbf{R}(s_{k-1}, s_k) \cdot e^{-r(s_{k-1})\tau}\, d\tau.$$

Solving the integral

$$Pr(Cyl(s_0, I_0, \ldots, I_{k-2}, s_{k-1})) \cdot \mathbf{P}(s_{k-1}, s_k) \cdot \left(e^{-r(s_k) \cdot \inf I_{k-1}} - e^{-r(s_k) \cdot \sup I_{k-1}} \right).$$

Zeno theorem

Zeno path

Path $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \xrightarrow{t_2} s_3 \ldots\ldots$ is called Zeno [1] if $\sum_i t_i$ converges.

Intuition

In case $\sum_i t_i$ does not diverge, the timed path represents an"unrealistic" computation where infinitely many transitions are taken in a finite amount of time. Example:

$$s_0 \xrightarrow{1} s_1 \xrightarrow{\frac{1}{2}} s_2 \xrightarrow{\frac{1}{4}} s_3 \ldots s_i \xrightarrow{\frac{1}{2^i}} s_{i+1} \ldots$$

In real-time systems, such executions are typically excluded from the analysis.

Thanks to the following theorem, Zeno paths do not harm for CTMCs.

Zeno theorem

For all states s in any CTMC, $Pr\{\, \pi \in Paths(s) \mid \pi \text{ is Zeno} \,\} = 0$.

[1] Zeno of Elea (490–430 BC), philosopher, famed for his paradoxes.

Proof of Zeno theorem

Zeno theorem

For all states s in any CTMC, $Pr\{\, \pi \in Paths(s) \mid \pi \text{ is Zeno} \,\} = 0$.

Reachability events

Let CTMC $\mathcal{C}$ with (possibly infinite) state space S.

(Simple) reachability

Eventually reach a state in $G \subseteq S$. Formally:

$$\Diamond G \;=\; \{\, \pi \in Paths(\mathcal{C}) \mid \exists i \in \mathbb{N}.\, \pi[i] \in G \,\}$$

Invariance, i.e., always stay in state in G:

$$\Box G \;=\; \{\, \pi \in Paths(\mathcal{C}) \mid \forall i \in \mathbb{N}.\, \pi[i] \in G \,\} \;=\; \overline{\Diamond \overline{G}}.$$

Constrained reachability

Or "reach-avoid" properties where states in $F \subseteq S$ are forbidden:

$$\overline{F} \,\mathsf{U}\, G \;=\; \{\, \pi \in Paths(\mathcal{C}) \mid \exists i \in \mathbb{N}.\, \pi[i] \in G \wedge \forall j < i.\, \pi[j] \notin F \,\}$$

Measurability

Measurability theorem

Events $\Diamond G$, $\Box G$, $\overline{F} \,\mathsf{U}\, G$, $\Box\Diamond G$ and $\Diamond\Box G$ are measurable on any CTMC.

Proof:

Left as an exercise.

Reachability probabilities in finite CTMCs

Problem statement

Let $\mathcal{C}$ be a CTMC with finite state space S, $s \in S$ and $G \subseteq S$.

Aim: determine $Pr(s \models \Diamond G) = Pr_s(\Diamond G) = Pr_s\{ \pi \in Paths(s) \mid \pi \models \Diamond G \}$

where Pr_s is the probability measure in $\mathcal{C}$ with single initial state s.

Characterisation of reachability probabilities

- Let variable $x_s = Pr(s \models \Diamond G)$ for any state s
  - if G is not reachable from s, then $x_s = 0$
  - if $s \in G$ then $x_s = 1$
- For any state $s \in Pre^*(G) \setminus G$:

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

Verifying CTMCs

Verifying untimed properties

So, computing reachability probabilities is exactly the same as for DTMCs. The same holds for constrained reachability, persistence and repeated reachability. In fact, all PCTL and LTL formulas can be checked on the embedded DTMC $(S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ using the techniques described before in these lecture slides.

Justification:

As the above temporal logic formulas or events do not refer to elapsed time, it is not surprising that they can be checked on the embedded DTMC.

Timed reachability events

Let CTMC $\mathcal{C}$ with (possibly infinite) state space S.

(Simple) timed reachability

Eventually reach a state in $G \subseteq S$ in the interval I. Formally:

$$\Diamond^I G = \{ \pi \in Paths(\mathcal{C}) \mid \exists t \in I.\, \pi@t \in G \}$$

Invariance, i.e., always stay in state in G in the interval I:

$$\Box^I G = \{ \pi \in Paths(\mathcal{C}) \mid \forall t \in I.\, \pi@t \in G \} = \overline{\Diamond^I \overline{G}}.$$

Constrained timed reachability

Or "reach-avoid" properties where states in $F \subseteq S$ are forbidden:

$$\overline{F}\, \mathrm{U}^I G = \{ \pi \in Paths(\mathcal{C}) \mid \exists t \in I.\, \pi@t \in G \wedge \forall d < t.\, \pi@d \notin F \}$$

Measurability

Measurability theorem

Events $\Diamond^I G$, $\Box^I G$, and $\overline{F}\, \mathrm{U}^I G$ are measurable on any CTMC.

Proof:

Left as an exercise.

Timed reachability probabilities in finite CTMCs

Problem statement

Let $\mathcal{C}$ be a CTMC with finite state space S, $s \in S$, $t \in \mathbb{R}_{\geqslant 0}$ and $G \subseteq S$.

Aim: $Pr(s \models \Diamond^{\leqslant t} G) = Pr_s(\Diamond^{\leqslant t} G) = Pr_s\{\, \pi \in Paths(s) \mid \pi \models \Diamond^{\leqslant t} G \,\}$

where Pr_s is the probability measure in $\mathcal{C}$ with single initial state s.

Characterisation of timed reachability probabilities

- Let function $x_s(t) = Pr(s \models \Diamond^{\leqslant t} G)$ for any state s
  - if G is not reachable from s, then $x_s(t) = 0$ for all t
  - if $s \in G$ then $x_s(t) = 1$ for all t
- For any state $s \in Pre^*(G) \setminus G$:

$$x_s(t) = \int_0^t \sum_{s' \in S} \underbrace{\mathbf{R}(s,s') \cdot e^{-r(s) \cdot x}}_{\substack{\text{probability to move to} \\ \text{state } s' \text{ at time } x}} \cdot \underbrace{x_{s'}(t-x)}_{\substack{\text{prob. to fulfill} \\ \Diamond^{\leqslant t-x} G \text{ from } s'}} \, dx$$

Reachability

Reachability probabilities in finite DTMCs and CTMCs

Can be obtained by solving a system of linear equations for which many efficient techniques exists.

Timed reachability probabilities in finite CTMCs

Can be obtained by solving a system of Volterra integral equations. This is in general a non-trivial issue, inefficient, and has several pitfalls such as numerical stability.

Solution

Reduce the problem of computing $Pr(s \models \Diamond^{\leqslant t} G)$ to an alternative problem for which well-known efficient techniques exist: computing transient probabilities (see previous lecture).

Timed reachability probabilities = transient probabilities

Aim

Compute $Pr(s \models \Diamond^{\leqslant t} G)$ in CTMC $\mathcal{C}$. Observe that once a path π reaches G within t time, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Let CTMC $\mathcal{C} = (S, \mathbf{P}, r, \iota_{\text{init}}, AP, L)$ and $G \subseteq S$. The CTMC $\mathcal{C}[G] = (S, \mathbf{P}_G, r, \iota_{\text{init}}, AP, L)$ with $\mathbf{P}_G(s,t) = \mathbf{P}(s,t)$ if $s \notin G$ and $\mathbf{P}_G(s,s) = 1$ if $s \in G$.

All outgoing transitions of $s \in G$ are replaced by a single self-loop at s.

Lemma

$$\underbrace{Pr(s \models \Diamond^{\leqslant t} G)}_{\text{timed reachability in } \mathcal{C}} = \underbrace{Pr(s \models \Diamond^{= t} G)}_{\text{timed reachability in } \mathcal{C}[G]} = \underbrace{\underline{p}(t) \text{ with } \underline{p}(0) = \mathbf{1}_s}_{\text{transient prob. in } \mathcal{C}[G]}.$$

Constrained timed reachability probabilities

Problem statement

Let $\mathcal{C}$ be a CTMC with finite state space S, $s \in S$, $t \in \mathbb{R}_{\geqslant 0}$ and $G, F \subseteq S$.

Aim: $Pr(s \models \overline{F} \,\mathsf{U}^{\leqslant t} G) = Pr_s(\overline{F} \,\mathsf{U}^{\leqslant t} G) = Pr_s\{\, \pi \in Paths(s) \mid \pi \models \overline{F} \,\mathsf{U}^{\leqslant t} G \,\}$.

Characterisation of timed reachability probabilities

- Let function $x_s(t) = Pr(s \models \overline{F} \,\mathsf{U}^{\leqslant t} G)$ for any state s
  - if G is not reachable from s via $\overline{F}$, then $x_s(t) = 0$ for all t
  - if $s \in G$ then $x_s(t) = 1$ for all t
- For any state $s \in Pre^*(G) \setminus (F \cup G)$:

$$x_s(t) = \int_0^t \sum_{s' \in S} \underbrace{\mathbf{R}(s,s') \cdot e^{-r(s) \cdot x}}_{\substack{\text{probability to move to} \\ \text{state } s' \text{ at time } x}} \cdot \underbrace{x_{s'}(t-x)}_{\substack{\text{prob. to fulfill} \\ \overline{F} \,\mathsf{U}^{\leqslant t-x} G \text{ from } s'}} \, dx$$

Constrained timed reachability = transient probabilities

Aim

Compute $Pr(s \models \overline{F} \, \mathrm{U}^{\leqslant t} \, G)$ in CTMC $\mathcal{C}$. Observe (as before) that once a path π reaches G within time t via $\overline{F}$, then the remaining behaviour along π is not important. Now also observe that once $s \in F \setminus G$ is reached within time t, then the remaining behaviour along π is not important. This suggests to make all states in G <u>and</u> $F \setminus G$ absorbing.

Lemma

$$\underbrace{Pr(s \models \overline{F} \, \mathrm{U}^{\leqslant t} \, G)}_{\text{timed reachability in } \mathcal{C}} \;=\; \underbrace{Pr(s \models \Diamond^{=t} G)}_{\text{timed reachability in } \mathcal{C}[F \cup G]} \;=\; \underbrace{\underline{p}(t) \text{ with } \underline{p}(0) = \mathbf{1}_s}_{\text{transient prob. in } \mathcal{C}[F \cup G]} \;.$$

Strong and weak bisimulation

Bisimulation preserves timed reachability events

Let $\mathcal{C}$ be a CTMC with state space S, $s, u \in S$, $t \in \mathbb{R}_{\geqslant 0}$ and $G, F \subseteq S$. Then:

1. $s \sim_m u$ implies $Pr(s \models \overline{F} \, \mathrm{U}^{\leqslant t} \, G) \;=\; Pr(u \models \overline{F} \, \mathrm{U}^{\leqslant t} \, G)$
2. $s \approx_m u$ implies $Pr(s \models \overline{F} \, \mathrm{U}^{\leqslant t} \, G) \;=\; Pr(u \models \overline{F} \, \mathrm{U}^{\leqslant t} \, G)$

provided F and G are closed under $\sim_m$ and $\approx_m$, respectively.

Proof:

Left as an exercise.

Summary

Main points

- Cylinder sets in a CTMC are paths that share interval-timed path prefixes.
- Reachability, persistence and repeated reachability can be checked as on DTMCs.
- Timed reachability probabilities can be characterised as Volterra integral equation system.
- Computing timed reachability probabilities can be reduced to transient probabilities.
- Weak and strong bisimulation preserves timed reachability probabilities.

Overview

Continuous Stochastic Logic

- CSL is a language for formally specifying properties over CTMCs.
- It is a branching-time temporal logic based on CTL.
- Formula interpretation is Boolean, i.e., a state satisfies a formula or not.
- Like in PCTL, the main operator is $\mathbb{P}_J(\varphi)$
  - where φ constrains the set of paths and J is a threshold on the probability.
  - it is the probabilistic counterpart of $\exists$ and $\forall$ path-quantifiers in CTL.
- The new features are a timed version of the next and until-operator.
  - $\bigcirc^I \Phi$ asserts that a transition to a Φ-state can be made at time $t \in I$.
  - $\Phi \, \mathsf{U}^I \Psi$ asserts that a Ψ-state can be reached via Φ-states at time $t \in I$.

CSL syntax [Baier, Katoen & Hermanns, 1999]

Continuous Stochastic Logic: Syntax

CSL consists of state- and path-formulas.

- CSL *state formulas* over the set AP obey the grammar:

$$\Phi \ ::= \ \text{true} \ \big| \ a \ \big| \ \Phi_1 \wedge \Phi_2 \ \big| \ \neg\Phi \ \big| \ \mathbb{P}_J(\varphi)$$

  where $a \in AP$, φ is a path formula and $J \subseteq [0,1]$, $J \neq \varnothing$ is a non-empty interval.
- CSL *path formulae* are formed according to the following grammar:

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

  where Φ, Φ_1, and Φ_2 are state formulae and $I \subseteq \mathbb{R}_{\geqslant 0}$ an interval.

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

Continuous Stochastic Logic

- CSL *state formulas* over the set AP obey the grammar:

$$\Phi \ ::= \ \text{true} \ \big| \ a \ \big| \ \Phi_1 \wedge \Phi_2 \ \big| \ \neg\Phi \ \big| \ \mathbb{P}_J(\varphi)$$

  where $a \in AP$, φ is a path formula and $J \subseteq [0,1]$, $J \neq \varnothing$.
- CSL *path formulae* are formed according to the following grammar:

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

  where Φ, Φ_1, and Φ_2 are state formulae and $I \subseteq \mathbb{R}_{\geqslant 0}$ an interval.

Intuitive semantics

- $s_0 t_0 s_1 t_1 \ldots \models \Phi \, \mathsf{U}^I \, \Psi$ if Ψ is reached at $t \in I$ and prior to t, Φ holds.
- $s \models \mathbb{P}_J(\varphi)$ if probability that paths starting in s fulfill φ lies in J.

Derived operators

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

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

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

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

Paths in a CTMC

Timed paths

Paths in CTMC $\mathcal{C}$ are maximal (i.e., infinite) paths of alternating states and time instants:

$$\pi = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots$$

such that $s_i \in S$ and $t_i \in \mathbb{R}_{>0}$. Let $Paths(\mathcal{C})$ be the set of paths in $\mathcal{C}$ and $Paths^*(\mathcal{C})$ the set of finite prefixes thereof.

Notations

- Let $\pi[i] := s_i$ denote the $(i+1)$-st state along the timed path π.
- Let $\pi\langle i \rangle := t_i$ the time spent in state s_i.
- Let $\pi@t$ be the state occupied in π at time $t \in \mathbb{R}_{\geqslant 0}$, i.e. $\pi@t := \pi[i]$ where i is the smallest index such that $\sum_{j=0}^{i} \pi\langle j \rangle > t$.

Example properties

- Transient probabilities to be in *goal* state at time point 4:

$$\mathbb{P}_{\geqslant 0.92}\left(\Diamond^{=4}\, goal\right)$$

- With probability $\geqslant 0.92$, a goal state is reached legally:

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

- ... in maximally 137 time units: $\mathbb{P}_{\geqslant 0.92}\left(\neg\, illegal \;\mathsf{U}^{\leqslant 137}\; goal\right)$
- ... once there, remain there almost surely for the next 31 time units:

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

CSL semantics (1)

Notation

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

Satisfaction relation for state formulas

The satisfaction relation $\models$ is defined for CSL state formulas by:

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

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

This is as for PCTL, except that Pr is the probability measures on cylinder sets of timed paths in CTMC $\mathcal{C}$.

CSL semantics (2)

Satisfaction relation for path formulas

Let $\pi = s_0\, t_0\, s_1\, t_1\, s_2 \ldots$ be an infinite path in CTMC $\mathcal{C}$.

The satisfaction relation $\models$ is defined for state formulas by:

$$\begin{array}{lll} \pi \models \bigcirc^I \Phi & \text{iff} & s_1 \models \Phi \wedge t_0 \in I \\ \pi \models \Phi \,\mathsf{U}^I\, \Psi & \text{iff} & \exists t \in I.\ ((\forall t' \in [0,t).\ \pi@t' \models \Phi) \wedge \pi@t \models \Psi) \end{array}$$

Standard next- and until-operators

- $X\Phi \equiv \bigcirc^I \Phi$ with $I = \mathbb{R}_{\geqslant 0}$.
- $\Phi \,\mathsf{U}\, \Psi \equiv \Phi \,\mathsf{U}^I\, \Psi$ with $I = \mathbb{R}_{\geqslant 0}$.

Measurability

CSL measurability

For any CSL path formula φ and state s of CTMC $\mathcal{C}$,
the set $\{ \pi \in Paths(s) \mid \pi \models \varphi \}$ is measurable.

Proof:

Rather straightforward; left as an exercise.

CSL model checking

CSL model checking problem

Input: a finite CTMC $\mathcal{C} = (S, \mathbf{P}, r, \iota_{\text{init}}, AP, L)$, state $s \in S$, and CSL state formula Φ

Output: yes, if $s \models \Phi$; no, otherwise.

Basic algorithm

In order to check whether $s \models \Phi$ do:

1. Compute the satisfaction set $Sat(\Phi) = \{ s \in S \mid s \models \Phi \}$.
2. This is done recursively by a bottom-up traversal of Φ's parse tree.
   - The nodes of the parse tree represent the subformulae of Φ.
   - For each node, i.e., for each subformula Ψ of Φ, determine $Sat(\Psi)$.
   - Determine $Sat(\Psi)$ as function of the satisfaction sets of its children:

     e.g., $Sat(\Psi_1 \wedge \Psi_2) = Sat(\Psi_1) \cap Sat(\Psi_2)$ and $Sat(\neg\Psi) = S \setminus Sat(\Psi)$.
3. Check whether state s belongs to $Sat(\Phi)$.

Core model checking algorithm

Probabilistic operator $\mathbb{P}$

In order to determine whether $s \in Sat(\mathbb{P}_J(\varphi))$, the probability $Pr(s \models \varphi)$ for the event specified by φ needs to be established. Then

$$Sat(\mathbb{P}_J(\varphi)) \;=\; \{s \in S \mid Pr(s \models \varphi) \in J\}.$$

Let us consider the computation of $Pr(s \models \varphi)$ for all possible φ.

The next-step operator

Recall that: $s \models \mathbb{P}_J(\bigcirc^I \Phi)$ if and only if $Pr(s \models \bigcirc^I \Phi) \in J$.

Lemma

$$Pr(s \models \bigcirc^I \Phi) \;=\; \underbrace{\left(e^{-r(s)\cdot \inf I} - e^{-r(s)\cdot \sup I}\right)}_{\text{probability to leave } s \text{ in interval } I} \cdot \sum_{s' \in Sat(\Phi)} \mathbf{P}(s, s').$$

Algorithm

Considering the above equation for all states simultaneously yields:

$$\left(Pr(s \models \bigcirc \Phi)\right)_{s \in S} \;=\; \mathbf{b}_I^T \cdot \mathbf{P}$$

with $\mathbf{b}_I$ is defined by $b_I(s) = e^{-r(s)\cdot \inf I} - e^{-r(s)\cdot \sup I}$ if $s \in Sat(\Phi)$ and 0 otherwise, and $\mathbf{b}_I^T$ is the transposed variant of $\mathbf{b}_I$.

Time-bounded until (1)

Recall that: $s \models \mathbb{P}_J(\Phi \, \mathsf{U}^{\leqslant t} \, \Psi)$ if and only if $Pr(s \models \Phi \, \mathsf{U}^{\leqslant t} \, \Psi) \in J$.

Lemma

Let $S_{=1} = Sat(\Psi)$, $S_{=0} = S \setminus (Sat(\Phi) \cup Sat(\Psi))$, and $S_? = S \setminus (S_{=0} \cup S_{=1})$. Then:

$$Pr(s \models \Phi \, \mathsf{U}^{\leqslant t} \, \Psi) = \begin{cases} 1 & \text{if } s \in S_{=1} \\ 0 & \text{if } s \in S_{=0} \\ \int_0^t \sum_{s' \in S} \mathbf{R}(s,s') \cdot e^{-r(s) \cdot x} \cdot Pr(s' \models \Phi \, \mathsf{U}^{\leqslant t-x} \, \Psi) \, dx & \text{otherwise} \end{cases}$$

This is a slight generalisation of the Volterra integral equation system for timed reachability.

Time-bounded until (2)

Let $S_{=1} = Sat(\Psi)$, $S_{=0} = S \setminus (Sat(\Phi) \cup Sat(\Psi))$, and $S_? = S \setminus (S_{=0} \cup S_{=1})$. Then:

$$Pr(s \models \Phi \, \mathsf{U}^{\leqslant t} \, \Psi) = \begin{cases} 1 & \text{if } s \in S_{=1} \\ 0 & \text{if } s \in S_{=0} \\ \int_0^t \sum_{s' \in S} \mathbf{R}(s,s') \cdot e^{-r(s) \cdot x} \cdot Pr(s' \models \Phi \, \mathsf{U}^{\leqslant t-x} \, \Psi) \, dx & \text{otherwise} \end{cases}$$

Recall that

$$\underbrace{Pr(s \models \overline{F} \, \mathsf{U}^{\leqslant t} \, G)}_{\text{timed reachability in } \mathcal{C}} = \underbrace{Pr(s \models \Diamond^{=t} G)}_{\text{in } \mathcal{C}[F \cup G]} = \underbrace{\underline{p}(t) \text{ with } \underline{p}(0) = \mathbf{1}_s}_{\text{transient prob. in } \mathcal{C}[F \cup G]} .$$

Phrased using CSL state formulas

$$\underbrace{Pr(s \models \Phi \, \mathsf{U}^{\leqslant t} \, \Psi)}_{\text{timed reachability in } \mathcal{C}} = \underbrace{Pr(s \models \Diamond^{=t} \Psi)}_{\text{in } \mathcal{C}[Sat(\neg\Phi) \cup Sat(\Psi)]} = \underbrace{\underline{p}(t) \text{ with } \underline{p}(0) = \mathbf{1}_s}_{\mathcal{C}[Sat(\neg\Phi) \cup Sat(\Psi)]} .$$

Time-bounded until (3)

Algorithm for checking $Pr(s \models \Phi \, \mathsf{U}^{\leqslant t} \, \Psi) \in J$

1. If $t = \infty$, then use approach for until (as in PCTL): solve a system of linear equations.
2. Determine recursively $Sat(\Phi)$ and $Sat(\Psi)$.
3. Make all states in $S \setminus Sat(\Phi)$ and $Sat(\Psi)$ absorbing.
4. Uniformize the resulting CTMC with respect to its maximal rate.
5. Determine the transient probability at time t using s as initial distribution.
6. Return yes if transient probability of all Ψ-states lies in J, and no otherwise.

Time-bounded until (4)

Possible optimizations

1. Make all states in $S \setminus Sat(\exists(\Phi \, \mathsf{U} \, \Psi))$ absorbing.
2. Make all states in $Sat(\forall(\Phi \, \mathsf{U} \, \Psi))$ absorbing.
3. Replace the labels of all states in $S \setminus Sat(\exists(\Phi\Psi))$ by unique label zero.
4. Replace the labels of all states in $Sat(\forall(\Phi \, \mathsf{U} \, \Psi))$ by unique label one.
5. Perform bisimulation minimization on all states.

The last step collapses all states in $S \setminus Sat(\exists(\Phi \, \mathsf{U} \, \Psi))$ into a single state, and does the same with all states in $Sat(\forall(\Phi \, \mathsf{U} \, \Psi))$.

Preservation of CSL-formulas

Bisimulation and CSL-equivalence coincide

Let $\mathcal{C}$ be a finitely branching CTMC and s, t states in $\mathcal{C}$. Then:

$$s \sim_m t \quad \text{if and only if} \quad s \text{ and } t \text{ are CSL-equivalent.}$$

Remarks

If for CSL-formula Φ we have $s \models \Phi$ but $t \not\models \Phi$, then it follows $s \not\sim_m t$. A single CSL-formula suffices!

Preservation of CSL-formulas

Weak bisimulation and CSL-without-next-equivalence coincide

Let $\mathcal{C}$ be a finitely branching CTMC and s, t states in $\mathcal{C}$. Then:

$$s \approx_m t \quad \text{if and only if} \quad s \text{ and } t \text{ are CSL-without-next-equivalent.}$$

Here. CSL-without-next is the fragment of CSL where the next-operator $\bigcirc$ does not occur.

Remarks

If for CSL-without-next-formula Φ we have $s \models \Phi$ but $t \not\models \Phi$, then it follows $s \not\approx_m t$.

Uniformization and CSL

Uniformization and CSL

For any finite CTMC $\mathcal{C}$ with state space S, $r \geqslant \max\{ r(s) \mid s \in S \}$ and Φ a CSL-without-next-formula:

$$Sat^{\mathcal{C}}(\Phi) \;=\; Sat^{\mathcal{C}'}(\Phi) \quad \text{where } \mathcal{C}' = unif(r, \mathcal{C}).$$

Uniformization and CSL

For any uniformized CTMC: CSL-equivalence coincides with CSL-without-next-equivalence.

Time complexity

Let $|\Phi|$ be the size of Φ, i.e., the number of logical and temporal operators in Φ.

Time complexity of CSL model checking

For finite CTMC $\mathcal{C}$ and CSL state-formula Φ, the CSL model-checking problem can be solved in time

$$\mathcal{O}\big(\, poly(size(\mathcal{C})) \cdot t_{\max} \cdot |\Phi| \,\big)$$

where $t_{\max} \;=\; \max\{ t \mid \Psi_1 \mathsf{U}^{\leqslant t} \Psi_2 \text{ occurs in } \Phi \}$ with and $t_{\max} = 1$ if Φ does not contain a time-bounded until-operator.

Some practical verification times

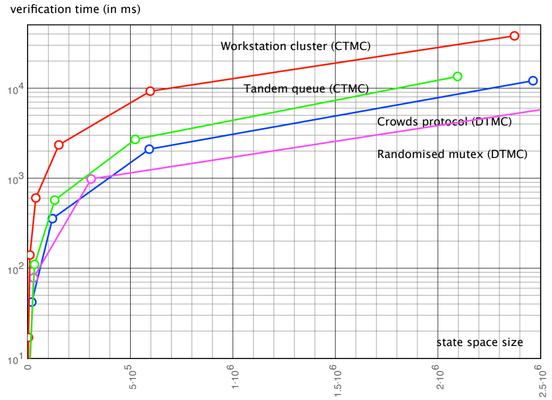

- command-line tool MRMC ran on a Pentium 4, 2.66 GHz, 1 GB RAM laptop.
- CSL formulas are time-bounded until-formulas.

Summary

- CSL is a variant of PCTL with timed next and timed until.
- Sets of paths fulfilling CSL path-formula φ are measurable.
- CSL model checking is performed by a recursive descent over Φ.
- The timed next operator amounts to a single vector-matrix multiplication.
- The time-bounded until-operator $\mathsf{U}^{\leqslant t}$ is solved by uniformization.
- The worst-case time complexity is polynomial in the size of the CTMC and linear in the size of the formula.

Overview

1. Negative exponential distributions
2. What are continuous-time Markov chains?
3. Transient distribution
4. Timed reachability probabilities
5. Verifying continuous stochastic CTL
6. Verifying linear real-time properties

Robot navigation

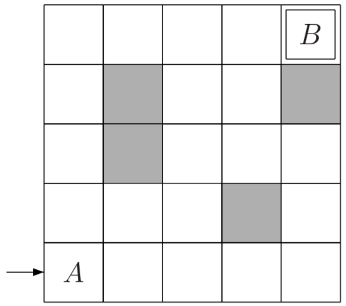

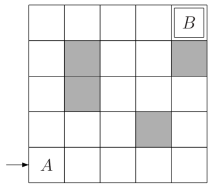

The robot randomly moves through the cells, and resides in a cell for an exponentially distributed amount of time.

Robot navigation: property

Property:

What is the probability to reach B from A within 10 time units while residing in any dangerous zone for at most 2 time units?

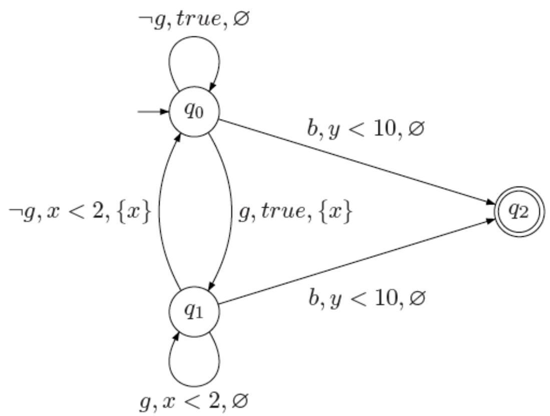

Deterministic timed automata

A Deterministic Timed Automaton (DTA) $\mathcal{A}$ is a tuple $(\Sigma, X, Q, q_0, F, \rightarrow)$:

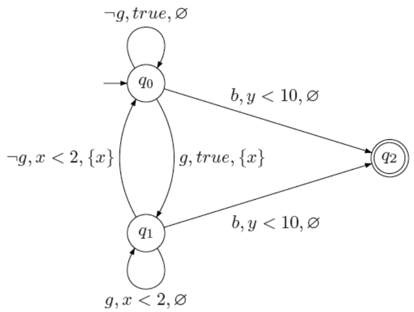

- Σ - *alphabet*
- X - finite set of *clocks*
- Q - finite set of *locations*
- $q_0 \in Q$ - *initial* location
- $F \subseteq Q$ - *accept* locations
- $\rightarrow \in Q \times \Sigma \times \mathcal{C}(X) \times 2^X \times Q$ - *transition relation*;

Determinism: $q \xrightarrow{a,g,X} q'$ and $q \xrightarrow{a,g',X'} q''$ implies $g \cap g' = \varnothing$

Model checking Markov chains

	branching time	linear time	
discrete-time (DTMC D)	PCTL	LTL	
	linear equations [HJ94] (⋆)	automata-based [V85,CSS03] (⋆⋆)	tableau-based [CY95]
	PTIME	PSPACE-C	

Model checking Markov chains

	branching time		linear time	
discrete-time (DTMC D)	PCTL		LTL	
	linear equations [HJ94] (⋆)		automata-based [V85,CSS03] (⋆⋆)	tableau-based [CY95]
	PTIME		PSPACE-C	
continuous-time (CTMC C)	untimed PCTL		untimed LTL	
	emb(C) (⋆)		*emb*(C) (⋆⋆)	
	PTIME		PSPACE-C	

Model checking Markov chains

	branching time		linear time	
discrete-time (DTMC D)	PCTL		LTL	
	linear equations [HJ94] (⋆)		automata-based [V85,CSS03] (⋆⋆)	tableau-based [CY95]
	PTIME		PSPACE-C	
continuous-time (CTMC C)	untimed PCTL	real-time CSL	untimed LTL	
	$emb(C)$ (⋆)	integral equations [BHHK03]	$emb(C)$ (⋆⋆)	
	PTIME	PTIME	PSPACE-C	

Model checking Markov chains

	branching time		linear time	
discrete-time (DTMC D)	PCTL		LTL	
	linear equations [HJ94] (⋆)		automata-based [V85,CSS03] (⋆⋆)	tableau-based [CY95]
	PTIME		PSPACE-C	
continuous-time (CTMC C)	untimed PCTL	real-time CSL	untimed LTL	real-time DTA
	$emb(C)$ (⋆)	integral equations [BHHK03]	$emb(C)$ (⋆⋆)	integral equations of second type (PDPs)
	PTIME	PTIME	PSPACE-C	PSPACE-C

What are we interested in?

Problem statement:

Given model CTMC C and specification DTA A, determine the fraction of runs in C that satisfy A:

$$Pr(C \models A) \;:=\; Pr^{C}\{\text{Paths in } C \text{ accepted by } A\}$$

Theoretical facts

Well-definedness

For any CTMC C and DTA A, the set $\{$Paths in C accepted by $A\}$ is measurable.

Characterizing the probability of $C \models A$

$Pr(C \models A)$ equals the reachability probability of accepting paths in $C \otimes A$.

Characterizing the probability of $C \models A$ under finite acceptance

$Pr(C \models A)$ equals the reachability probability of accepting paths in $C \otimes RG(A)$.

Characterizing the probability of $C \models A$ under Muller acceptance

$Pr(C \models A)$ equals the reachability probability of accepting terminal strongly connected components in $C \otimes RG(A)$.

Product construction

CTMC C with state space S | DTA A with state space Q

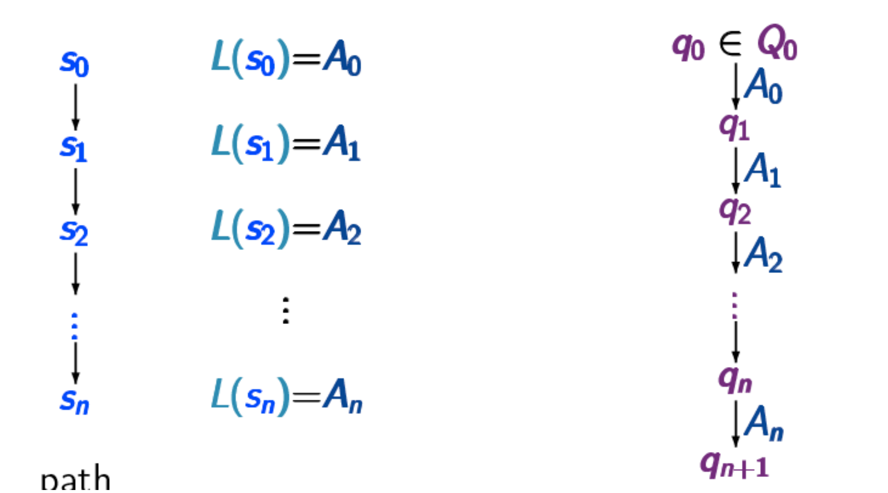

Product construction $\otimes$

CTMC C with state space S | DTA A with state space Q

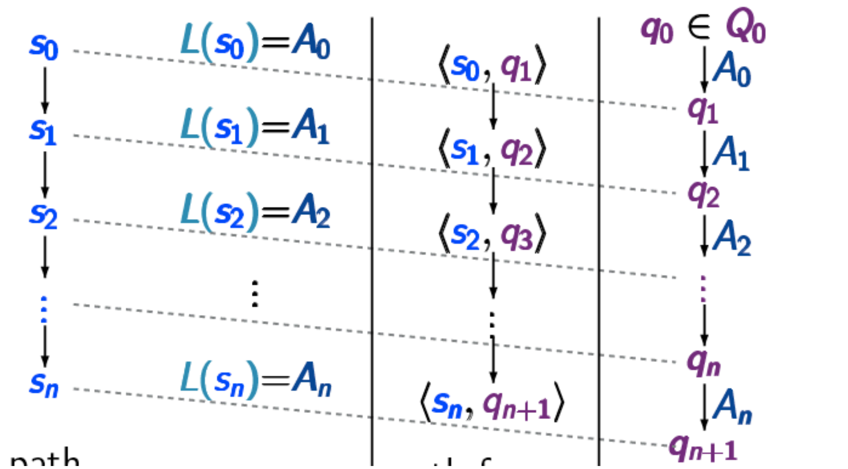

product $C \otimes A$

Product construction: example

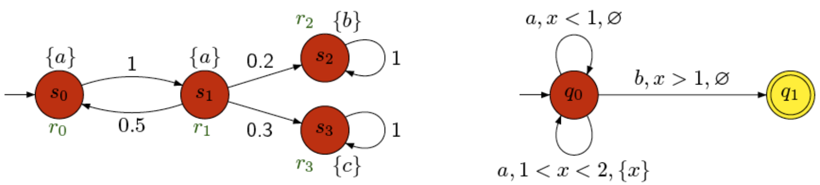

An example CTMC C (left) and DTA A (right)

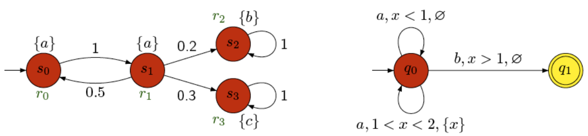

An example CTMC C (left up) and DTA A (right up) and $C \otimes RG(A)$ (below)

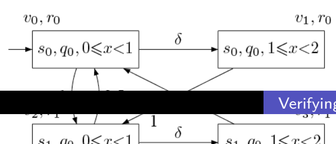

One-clock DTA: partitioning $C \otimes RG(A)$

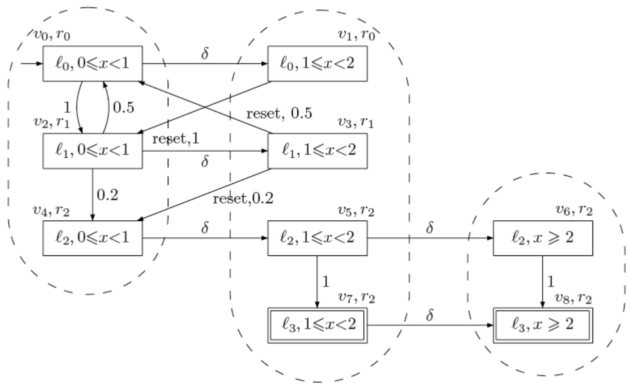

- constants $c_0 < \ldots < c_m$ in A yields $m+1$ subgraphs.
- subgraph i captures behaviour of C and A in $[c_i, c_{i+1})$.
- any subgraph is a CTMC, resets lead to subgraph 0, delays to $i+1$.
- a subgraph with its resets yields an "augmented" CTMC.

One-clock DTA: partitioning $C \otimes RG(A)$

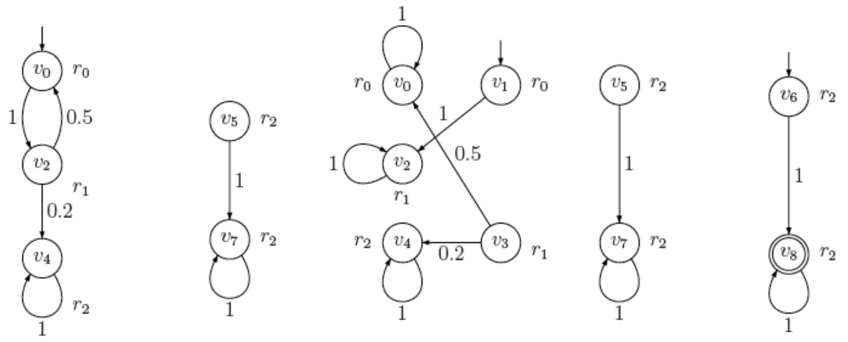

(a) $\mathcal{C}_0$ (b) $\mathcal{C}_1$ (c) $\mathcal{C}_1^a$ (d) $\mathcal{C}_2$

One-clock DTA: characterizing $Pr(C \models A)$

Theorem

For CTMC C with initial distribution α, 1-clock DTA A we have that:

$$Pr(C \models A) \;=\; \alpha \cdot \mathbf{u}$$

where $\mathbf{u}$ is the solution of the linear equation system $\mathbf{x} \cdot \mathbf{M} = \mathbf{f}$, with

$$\mathbf{M} = \left(\begin{array}{c|c} \mathbf{I}_{n_0} - \mathbf{B}_{m-1} & \mathbf{A}_{m-1} \\ \hline \hat{\mathbf{P}}^a_m & \mathbf{I}_{n_m} - \mathbf{P}_m \end{array} \right)$$

and $\mathbf{f}$ is the characterizing vector of the final states in subgraph m, and $\mathbf{A}$ and $\mathbf{B}$ are obtained from transient probabilities in all subgraphs.

One-clock DTA: algorithm

Algorithm 1 Verifying a CTMC against a 1-clock DTA

Require: a CTMC $\mathcal{C}$ with initial distr. α, a 1-clock DTA $\mathcal{A}$ with constants $c_0, \ldots, c_m$
Ensure: $\Pr(\mathcal{C} \models \mathcal{A})$

$\mathcal{G}(\mathcal{A})$:= `buildRegionGraph`($\mathcal{A}$);
Product := `buildProduct`($\mathcal{C}, \mathcal{G}(\mathcal{A})$); $\{\mathcal{C} \otimes \mathcal{G}(\mathcal{A})\}$
subGraphs $\{\mathcal{G}_i\}_{0 \leqslant i \leqslant m}$:= `partitionProduct`(*Product*);
for each subGraph $\mathcal{G}_i$ **do**
  $\mathcal{C}_i$:= `buildAugmentedCTMC`($\mathcal{G}_i$); {build augmented CTMC cf. Definition 6}
end for
$\{\mathcal{C}'_i\}_{0 \leqslant i \leqslant m}$:= `lumpGroupCTMCs`($\{\mathcal{C}_i\}_{0 \leqslant i \leqslant m}$); {lump a group of CTMCs, see Alg. 2}
for each CTMC $\mathcal{C}'_i$ **do** *TransProb*$_i$:= `computeTransientProb`($\mathcal{C}'_i, \Delta c_i$); **end for**
linearEqSystem := `buildLinearSystem`($\{TransProb_i\}_{0 \leqslant i \leqslant m}$); {cf. Theorem 2}
probVector := `solveLinearSystem`(*linearEqSystem*);
return $\alpha \cdot$ *probVector*;

Reachability in (our) PDPs

- For single-clock DTA, reachability probabilities in (our) PDPs are characterized by the least solution of a linear equation system, whose coefficients are solutions of some ordinary differential equations (ODEs).
- For these coefficients either an analytical solution (for small state space) can be obtained or an arbitrarily closely approximated solution can be determined efficiently.
- In multi-clock DTA, reachability probabilities in (our) PDPs are characterized as the least solution of a Volterra integral equation system of the second type.
- This solution can be approximated by solving a system of partial differential equations (PDEs).

Robot navigation revisited

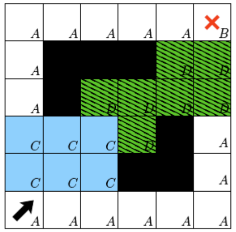

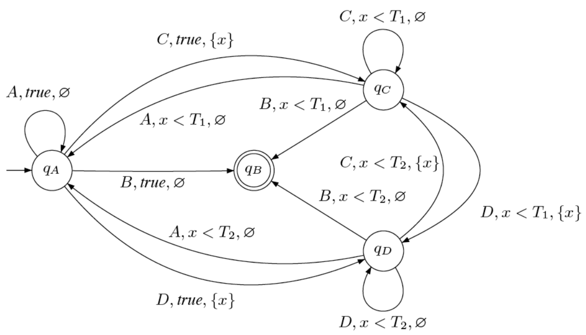

Black squares are walls. The residence time in consecutive C-cells $< T_1$.
The residence time in consecutive D-cells $< T_2$.

Verification results

	#CTMC	No lumping			With lumping			
N	states	# ⊗ states	time(s)	%transient	#blocks	time(s)	%transient	%lumping
10	100	148	0.09	59%	78	0.09	43%	32%
20	400	702	6.7	18%	380	7.1	14%	7%
30	900	1248	32	17%	619	26	14%	6%
40	1600	2672	119	13%	1296	93	10%	5%
50	2500	4174	135	17%	2015	138	12%	7%
60	3600	4232	309	16%	1525	261	12%	7%
70	4900	8661	904	12%	4212	1130	7%	3%
80	6400	9529	1753	12%	4339	1429	14%	4%
90	8100	9812	2433	8%	2613	1922	6%	5%

Product construction and solving the linear equation system is most time-consuming

Systems biology: immune-receptor signaling

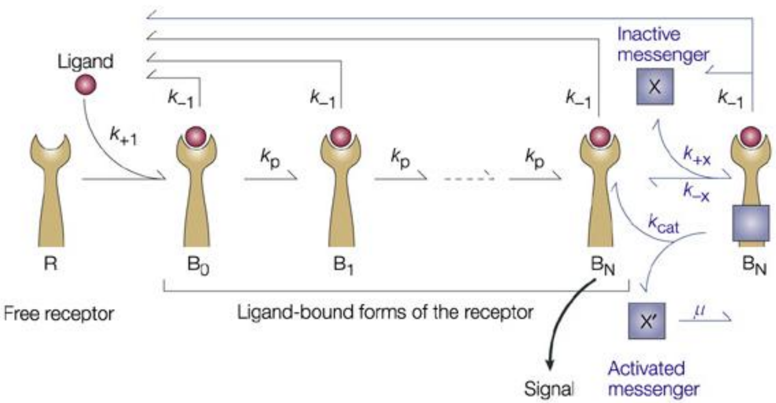

[Goldstein et. al., Nat. Reviews Immunology, 2004]

Systems biology: immune-receptor signaling

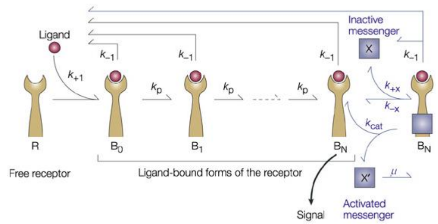

- M ligands can react with a receptor R with rate k_{+1} yielding a ligand-receptor LR
- LR undergoes a sequence of N modifications with a constant rate k_p yielding $B_1, \ldots, B_N$
- LR B_N can link with an inactive messenger with rate k_{+x} yielding a ligand-receptor-messenger (LRM).
- The LRM decomposes into an active messenger with rate k_{cat}

Verification results

M	#CTMC states	No lumping # ⊗ states	No lumping time(s)	With lumping #blocks	With lumping time(s)	%transient	%lumping
1	18	31	0	13	0	0%	0%
2	150	203	0.06	56	0.05	58%	39%
3	774	837	1.36	187	0.84	64%	30%
4	3024	2731	17.29	512	9.19	73%	24%
5	9756	7579	152.54	1213	73.4	76%	21%
6	27312	18643	1547.45	2579	457.35	78%	20%
7	68496	41743	11426.46	5038	3185.6	85%	14%
8	157299	86656	23356.5	9200	11950.8	81%	18%
9	336049	169024	71079.15	15906	38637.28	76%	22%
10	675817	312882	205552.36	26256	116314.41	71%	26%

In the case of no lumping, 99% of time is spent on transient analysis

Multi-multi-core model checking

N	4 Cores time(s)	4 Cores speedup	20 Cores time(s)	20 Cores speedup
3	0.45	3.03	0.42	3.22
4	5.3	3.26	3.44	5.02
5	44.73	3.41	15.87	9.61
6	620.16	2.50	160.58	9.64
7	4142.19	2.76	949.32	12.04
8	8168.62	2.86	1722.63	13.56
9	23865.17	2.98	5457.01	13.03
10	70623.46	2.91	16699.22	12.31

Parallelization of the transient analysis only; not the lumping.

Summary

Take-home messages

- Checking CTMCs against deterministic timed automata (DTA).
- Efficient numerical algorithm for one-clock DTA:
  - using standard means: region construction, graph analysis, transient analysis, linear equation systems.
  - three orders of magnitude faster than alternative approaches.
  - natural support for parallelization and bisimulation minimization.
- Discretization approach for multiple-clock DTA with error bounds.