

Modeling and Verification of Probabilistic Systems

Joost-Pieter Katoen

Lehrstuhl für Informatik 2
Software Modeling and Verification Group

<http://moves.rwth-aachen.de/teaching/ss-14/movep14/>

June 26, 2014

Continuous-time Markov chain

Continuous-time Markov chain

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

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

Let $\mathbf{R}(s, s') = \mathbf{P}(s, s') \cdot r(s)$ be the transition rate of transition (s, s')

Interpretation

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

Overview

- 1 Recall: continuous-time Markov chains
- 2 Probability measure on CTMC paths
- 3 Reachability probabilities
 - Untimed reachability
 - Timed reachability
 - Reduction to transient analysis
 - Bisimulation and timed reachability
- 4 Summary

CTMC semantics

Enabledness

The probability that transition $s \rightarrow s'$ 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:

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

Residence time distribution

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

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

Paths in a CTMC

Timed paths

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

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

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 $t \in \mathbb{R}_{\geq 0}$, i.e. $\pi @ t := \pi[i]$ where i is the smallest index such that $\sum_{j=0}^i \pi\langle j \rangle > t$.

Overview

- 1 Recall: continuous-time Markov chains
- 2 Probability measure on CTMC paths
- 3 Reachability probabilities
 - Untimed reachability
 - Timed reachability
 - Reduction to transient analysis
 - Bisimulation and timed reachability
- 4 Summary

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 $s_0 l_0 \dots l_{k-1} s_k$ 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 CTMCs

Cylinder set

Let $s_0, \dots, s_k \in S$ with $\mathbf{P}(s_i, s_{i+1}) > 0$ for $0 \leq i < k$ and l_0, \dots, l_{k-1} non-empty intervals in $\mathbb{R}_{\geq 0}$ with rational bounds. The *cylinder set* of $s_0 l_0 s_1 l_1 \dots l_{k-1} s_k$ is defined by:

$$Cyl(s_0, l_0, \dots, l_{k-1}, s_k) = \left\{ \pi \in Paths(\mathcal{C}) \mid \forall 0 \leq i \leq k. \pi[i] = s_i \text{ and } i < k \Rightarrow \pi\langle i \rangle \in l_i \right\}$$

The cylinder set spanned by $s_0, l_0, \dots, l_{k-1}, s_k$ thus consists of all infinite timed paths that have a prefix $\hat{\pi}$ that lies in $s_0, l_0, \dots, l_{k-1}, s_k$. 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 $Cyl(s_0, l_0, \dots, l_{k-1}, s_k)$ where $s_0 \dots s_k$ is a path in the state graph of \mathcal{C} (starting in s_0) and l_0, \dots, l_{k-1} range over all sequences of non-empty intervals in $\mathbb{R}_{\geq 0}$.

Probability measure on CTMCs

Cylinder set

The *cylinder set* $Cyl(s_0, l_0, \dots, l_{k-1}, s_k)$ of $s_0 l_0 \dots l_{k-1} s_k$ is defined by:

$$\{ \pi \in Paths(\mathcal{C}) \mid \forall 0 \leq i \leq k. \pi[i] = s_i \text{ and } i < k \Rightarrow \pi(i) \in l_i \}$$

Probability measure

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

$$Pr(Cyl(s_0, l_0, \dots, l_{k-1}, s_k)) = Pr(Cyl(s_0, l_0, \dots, l_{k-2}, s_{k-1})) \cdot \int_{l_{k-1}} \mathbf{R}(s_{k-1}, s_k) \cdot e^{-r(s_{k-1})\tau} d\tau.$$

Solving the integral

$$Pr(Cyl(s_0, l_0, \dots, l_{k-2}, s_{k-1})) \cdot \mathbf{P}(s_{k-1}, s_k) \cdot (e^{-r(s_{k-1}) \cdot \inf l_{k-1}} - e^{-r(s_{k-1}) \cdot \sup l_{k-1}}).$$

Proof of Zeno theorem

Zeno theorem

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

Proof:

On the blackboard.

Zeno theorem

Zeno path

Path $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \xrightarrow{t_2} s_3 \dots$ 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:

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

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, $Pr\{\pi \in Paths(s) \mid \pi \text{ is Zeno}\} = 0$.

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

Overview

- 1 Recall: continuous-time Markov chains
- 2 Probability measure on CTMC paths
- 3 Reachability probabilities
 - Untimed reachability
 - Timed reachability
 - Reduction to transient analysis
 - Bisimulation and timed reachability
- 4 Summary

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 :

$$\square G = \{ \pi \in Paths(\mathcal{C}) \mid \forall i \in \mathbb{N}. \pi[i] \in G \} = \overline{\overline{\diamond G}}$$

Constrained reachability

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

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

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 $s \in Pre^*(G) \setminus G$:

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

Measurability

Measurability theorem

Events $\diamond G$, $\square G$, $\overline{F} U G$, $\square \diamond G$ and $\diamond \square G$ are measurable on any CTMC.

Proof:

Consider $\diamond G$. $\diamond G$ is the union of all cylinders $Cyl(s_0, [0, \infty), \dots, [0, \infty), s_n)$ where $s_0, \dots, s_{n-1} \notin G$ and $s_n \in G$. As the set of state sequences $s_0 \dots s_n$ is countable, $\diamond G$ is a countable union of cylinders. Thus $\diamond G$ is measurable. The proof for $\square \diamond G$ goes along similar lines, using the proof principle for DTMCs.

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 $(S, \mathbf{P}, l_{\text{init}}, AP, L)$ 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 :

$$\square^I G = \{ \pi \in Paths(\mathcal{C}) \mid \forall t \in I. \pi @ t \in G \} = \overline{\diamond^I \bar{G}}.$$

Constrained timed reachability

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

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

Timed reachability probabilities in finite CTMCs

Problem statement

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

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

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

Characterisation of timed reachability probabilities

- ▶ Let function $x_s(t) = Pr(s \models \diamond^{\leq t} G)$ 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 $s \in Pre^*(G) \setminus G$:

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

Measurability

Measurability theorem

Events $\diamond^I G$, $\square^I G$, and $\bar{F}U^I G$ are measurable on any CTMC.

Proof (sketch):

Consider $\diamond^I G$ where $I = [0, t]$. $\diamond^{\leq t} G$ is the union of $Cyl(s_0, I_0, \dots, I_{n-1}, s_n)$ with $s_0, \dots, s_{n-1} \notin G$, $s_n \in G$, and $\sup(I_0) + \dots + \sup(I_{n-1}) \leq t$. The set of state sequences $s_0 \dots s_n$ is countable and the set of rational bounded intervals I_0, \dots, I_{n-1} is countable. Thus $\diamond^{\leq t} G$ is a countable union of cylinders, and thus is measurable. The proof for the remaining case $\bar{F}U^I G$ is similar and left as an exercise.

Timed reachability probabilities: example

On the blackboard.

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

Example

Timed reachability probabilities = transient probabilities

Aim

Compute $Pr(s \models \diamond^{\leq t} G)$ 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 $\mathcal{C} = (S, \mathbf{P}, r, \iota_{init}, AP, L)$ and $G \subseteq S$. The CTMC $\mathcal{C}[G] = (S, \mathbf{P}_G, r, \iota_{init}, AP, L)$ 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

$$\underbrace{Pr(s \models \diamond^{\leq t} G)}_{\text{timed reachability in } \mathcal{C}} = \underbrace{Pr(s \models \diamond^{=t} G)}_{\text{timed reachability in } \mathcal{C}[G]} = \underbrace{\sum_{s' \in G} p_{s'}(t)}_{\text{transient prob. in } \mathcal{C}[G]} \text{ with } \underline{p}(0) = \mathbf{1}_s$$

Constrained timed reachability probabilities

Problem statement

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

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

Characterisation of timed reachability probabilities

- ▶ Let function $x_s(t) = Pr(s \models \overline{F} U^{\leq t} G)$ 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 $s \in Pre^*(G) \setminus (F \cup G)$:

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

Constrained timed reachability = transient probabilities

Aim

Compute $Pr(s \models \bar{F} U^{\leq t} G)$ in CTMC \mathcal{C} . Observe (as before) that once a path π reaches G within time t via \bar{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 and $F \setminus G$ absorbing.

Lemma

$$\underbrace{Pr(s \models \bar{F} U^{\leq t} G)}_{\text{timed reachability in } \mathcal{C}} = \underbrace{Pr(s \models \diamond^{=t} G)}_{\text{timed reachability in } \mathcal{C}[F \cup G]} = \underbrace{\sum_{s' \in G} p_{s'}(t)}_{\text{transient prob. in } \mathcal{C}[F \cup G]} \text{ with } \underline{p}(0) = \mathbf{1}_s.$$

Example

Strong and weak bisimulation

Bisimulation preserves timed reachability events

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

1. $s \sim_m u$ implies $Pr(s \models \bar{F} U^{\leq t} G) = Pr(u \models \bar{F} U^{\leq t} G)$
2. $s \approx_m u$ implies $Pr(s \models \bar{F} U^{\leq t} G) = Pr(u \models \bar{F} U^{\leq t} G)$

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

Proof:

Left as an exercise.

Overview

- 1 Recall: continuous-time Markov chains
- 2 Probability measure on CTMC paths
- 3 Reachability probabilities
 - Untimed reachability
 - Timed reachability
 - Reduction to transient analysis
 - Bisimulation and timed reachability
- 4 Summary

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 preserve timed reachability probabilities.