Verifying Continuous-Time Markov Chains

Lecture 1+2: Discrete-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 20, 2011

Overview

Overview

Probabilities help

- When analysing system performance and dependability
  - to quantify arrivals, waiting times, time between failure, QoS, ...
- When modelling unreliable and unpredictable system behavior
  - to quantify message loss, processor failure
  - to quantify unpredictable delays, express soft deadlines, ...
- When building protocols for networked embedded systems
  - randomized algorithms
- When problems are undecidable deterministically
  - repeated reachability of lossy channel systems, ...

Illustrative example: Security

Security: Crowds protocol **[Reiter & Rubin, 1998]**

- A protocol for anonymous web browsing (variants: mCrowds, BT-Crowds)
- Hide user's communication by random routing within a crowd
  - sender selects a crowd member randomly using a uniform distribution
  - selected router flips a biased coin:
    - with probability $1 - p$: direct delivery to final destination
    - otherwise: select a next router randomly (uniformly)
  - once a routing path has been established, use it until crowd changes
- Rebuild routing paths on crowd changes
- Property: Crowds protocol ensures "probable innocence":
  - probability real sender is discovered $< \frac{1}{2}$ if $N \geqslant \frac{p}{p-\frac{1}{2}} \cdot (c+1)$
  - where N is crowd's size and c is number of corrupt crowd members

Illustrative example: Leader election

Distributed system: Leader election **[Itai & Rodeh, 1990]**

- A round-based protocol in a synchronous ring of $N > 2$ nodes
  - the nodes proceed in a lock-step fashion
  - each slot = 1 message is read + 1 state change + 1 message is sent
  - ⇒ this synchronous computation yields a discrete-time Markov chain
- Each round starts by each node choosing a uniform id $\in \{1, \ldots, K\}$
- Nodes pass their selected id around the ring
- If there is a unique id, the node with the maximum unique id is leader
- If not, start another round and try again ...

Properties of leader election

Almost surely eventually a leader will be elected

$$\mathbb{P}_{=1}(\lozenge \textit{leader elected})$$

With probability at least 0.8, a leader is elected within k steps

$$\mathbb{P}_{\geqslant 0.8}\left(\lozenge^{\leqslant k} \textit{leader elected}\right)$$

Probability to elect a leader within L rounds

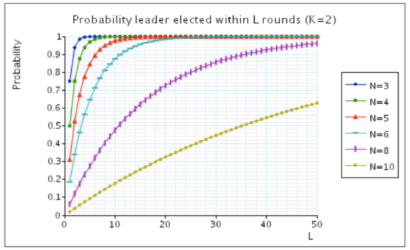

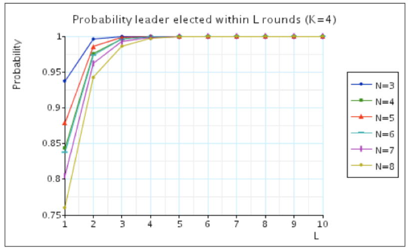

$$\mathbb{P}_{\leqslant q}\left(\lozenge^{\leqslant (N+1)\cdot L} \textit{ leader elected}\right)$$

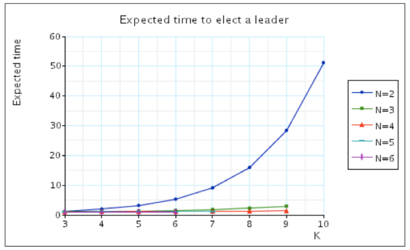

What is probabilistic model checking?

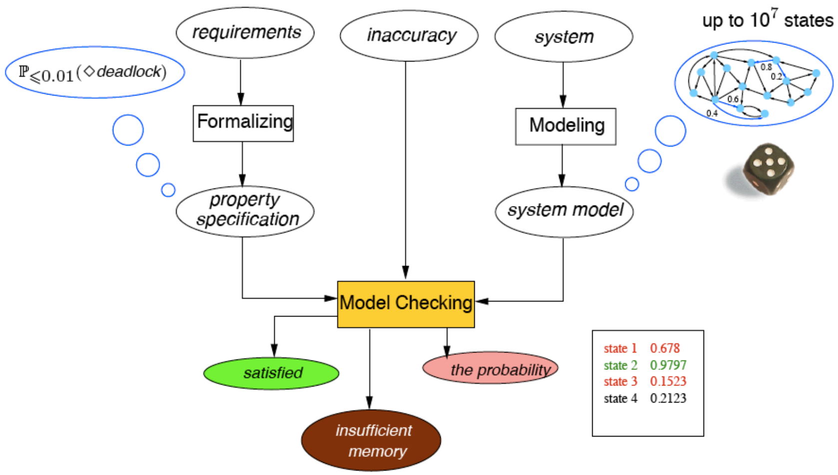

Probabilistic models

	Nondeterminism no	Nondeterminism yes
Discrete time	discrete-time Markov chain (DTMC)	Markov decision process (MDP)
Continuous time	CTMC	CTMDP

Other models: probabilistic variants of (priced) timed automata, or hybrid automata

Probability theory is simple, isn't it?

In no other branch of mathematics
is it so easy to make mistakes
as in probability theory

Henk Tijms, "Understanding Probability" (2004)

Overview

1. Motivation
2. What are discrete-time Markov chains?
3. Reachability probabilities
4. Qualitative reachability and all that
5. Verifying probabilistic CTL
6. Expressiveness of probabilistic CTL
7. Probabilistic bisimulation
8. Verifying ω-regular properties

Geometric distribution

Geometric distribution

Let X be a discrete random variable, natural $k > 0$ and $0 < p \leqslant 1$. The mass function of a *geometric distribution* is given by:

$$Pr\{ X = k \} = (1-p)^{k-1} \cdot p$$

We have $E[X] = \frac{1}{p}$ and $Var[X] = \frac{1-p}{p^2}$ and cdf $Pr\{ X \leqslant k \} = 1 - (1-p)^k$.

Geometric distributions and their cdf's

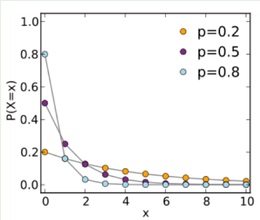

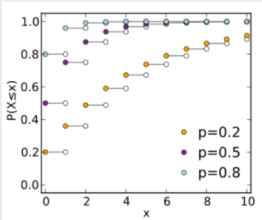

Memoryless property

Theorem

1. For any random variable X with a geometric distribution:

$$Pr\{X = k + m \mid X > m\} \ = \ Pr\{X = k\} \quad \text{for any} \quad m \in T, k \geqslant 1$$

This is called the memoryless property, and X is a memoryless r.v..

2. Any discrete random variable which is memoryless is geometrically distributed.

Joint distribution function

Joint distribution function

The *joint* distribution function of stochastic process $X = \{ X_t \mid t \in T \}$ is given for n, $t_1, \ldots, t_n \in T$ and $d_1, \ldots, d_n$ by:

$$F_X(d_1, \ldots, d_n; t_1, \ldots, t_n) = Pr\{ X(t_1) \leqslant d_1, \ldots, X(t_n) \leqslant d_n \}$$

The shape of F_X depends on the stochastic dependency between $X(t_i)$.

Stochastic independence

Random variables X_i on probability space $\mathcal{P}$ are *independent* if:

$$F_X(d_1, \ldots, d_n; t_1, \ldots, t_n) \ = \ \prod_{i=1}^{n} F_X(d_i; t_i) \ = \ \prod_{i=1}^{n} Pr\{ X(t_i) \leqslant d_i \}.$$

The next state of the stochastic process only depends on the current state, and not on states assumed previously. This is the Markov property.

Markov property

Markov process

A discrete-time stochastic process $\{ X(t) \mid t \in T \}$ over state space $\{ d_0, d_1, \ldots \}$ is a *Markov process* if for any $t_0 < t_1 < \ldots < t_n < t_{n+1}$:

$$Pr\{ X(t_{n+1}) = d_{n+1} \mid X(t_0) = d_0, X(t_1) = d_1, \ldots, X(t_n) = d_n \}$$

$$=$$

$$Pr\{ X(t_{n+1}) = d_{n+1} \mid X(t_n) = d_n \}$$

The distribution of $X(t_{n+1})$, given the values $X(t_0)$ through $X(t_n)$, only depends on the current state $X(t_n)$.

The conditional probability distribution of future states of a Markov process only depends on the current state and not on its further history.

Invariance to time-shifts

Time homogeneity

Markov process $\{ X(t) \mid t \in T \}$ is *time-homogeneous* iff for any $t' < t$:

$$Pr\{ X(t) = d \mid X(t') = d' \} \; = \; Pr\{ X(t - t') = d \mid X(0) = d' \}.$$

A time-homogeneous stochastic process is invariant to time shifts.

Discrete-time Markov chain

A *discrete-time Markov chain* (DTMC) is a time-homogeneous Markov process with discrete parameter T and discrete state space.

Discrete-time Markov chain

Discrete-time Markov chain

A *discrete-time Markov chain* (DTMC) is a time-homogeneous Markov process with discrete parameter T and discrete state space S.

Transition probabilities

The *(one-step) transition probability* from $s \in S$ to $s' \in S$ at epoch $n \in \mathbb{N}$ is given by:

$$p^{(n)}(s, s') \; = \; Pr\{ X_{n+1} = s' \mid X_n = s \} \; = \; Pr\{ X_1 = s' \mid X_0 = s \}$$

where the last equality is due to time-homogeneity.

Since $p^{(n)}(\cdot) = p^{(k)}(\cdot)$, the superscript (n) is omitted, and we write $p(\cdot)$.

Transition probability matrix

Discrete-time Markov chain

A *discrete-time Markov chain* (DTMC) is a time-homogeneous Markov process with discrete parameter T and discrete state space S.

Transition probability matrix

Let $\mathbf{P}$ be a function with $\mathbf{P}(s_i, s_j) = p(s_i, s_j)$. For finite state space S, function $\mathbf{P}$ is called the *transition probability matrix* of the DTMC with state space S.

Properties

1. $\mathbf{P}$ is a (right) *stochastic* matrix, i.e., it is a square matrix, all its elements are in $[0, 1]$, and each row sum equals one.
2. $\mathbf{P}$ has an eigenvalue of one, and all its eigenvalues are at most one.
3. For all $n \in \mathbb{N}$, $\mathbf{P}^n$ is a stochastic matrix.

DTMCs — A transition system perspective

Discrete-time Markov chain

A DTMC $\mathcal{D}$ is a tuple $(S, \mathbf{P}, \iota_{\mathrm{init}}, AP, L)$ with:

- S is a countable nonempty set of states
- $\mathbf{P} : S \times S \rightarrow [0, 1]$, transition probability function s.t. $\sum_{s'} \mathbf{P}(s, s') = 1$
- $\iota_{\mathrm{init}} : S \rightarrow [0, 1]$, the initial distribution with $\sum\limits_{s \in S} \iota_{\mathrm{init}}(s) = 1$
- AP is a set of atomic propositions.
- $L : S \rightarrow 2^{AP}$, the labeling function, assigning to state s, the set $L(s)$ of atomic propositions that are valid in s.

Initial states

- $\iota_{\mathrm{init}}(s)$ is the probability that DTMC $\mathcal{D}$ starts in state s
- the set $\{ s \in S \mid \iota_{\mathrm{init}}(s) > 0 \}$ are the possible initial states.

Simulating a die by a fair coin [Knuth & Yao]

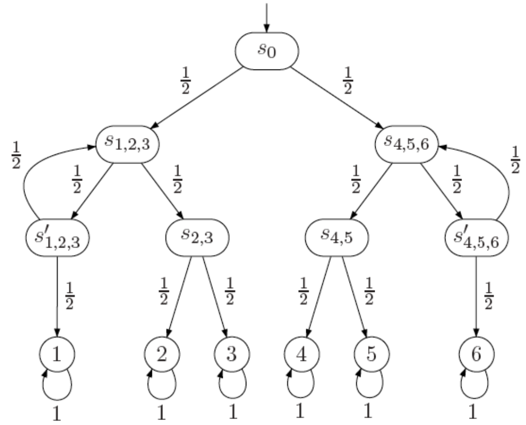

Heads = "go left"; tails = "go right". Does this DTMC adequately model a fair six-sided die?

Craps

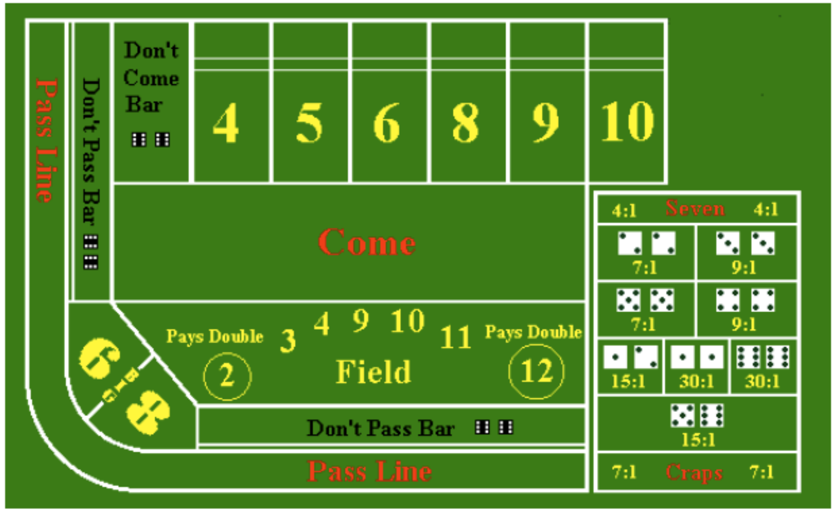

Craps

- Roll two dice and bet
- Come-out roll ("pass line" wager):
  - outcome 7 or 11: win
  - outcome 2, 3, or 12: lose ("craps")
  - any other outcome: roll again (outcome is "point")
- Repeat until 7 or the "point" is thrown:
  - outcome 7: lose ("seven-out")
  - outcome the point: win
  - any other outcome: roll again

A DTMC model of Craps

- Come-out roll:
  - 7 or 11: win
  - 2, 3, or 12: lose
  - else: roll again
- Next roll(s):
  - 7: lose
  - point: win
  - else: roll again

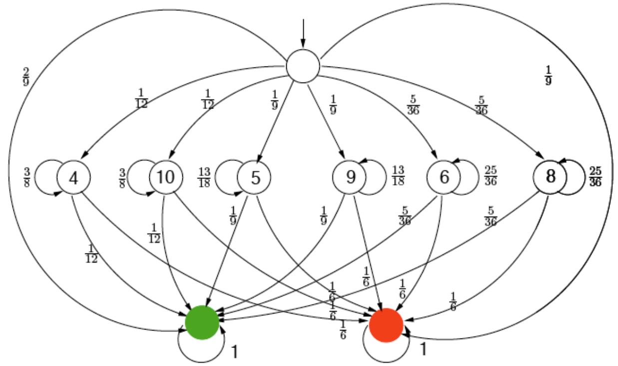

State residence time distribution

Let T_s be the number of epochs of DTMC $\mathcal{D}$ to stay in state s:

$$\begin{aligned} Pr\{\,T_s = 1\,\} &= 1 - \mathbf{P}(s,s) \\ Pr\{\,T_s = 2\,\} &= \mathbf{P}(s,s) \cdot (1 - \mathbf{P}(s,s)) \\ \ldots\ldots & \ldots \ \ldots\ldots \\ Pr\{\,T_s = n\,\} &= \mathbf{P}(s,s)^{n-1} \cdot (1 - \mathbf{P}(s,s)) \end{aligned}$$

So, the state residence times in a DTMC obey a *geometric* distribution.

The expected number of time steps to stay in state s equals $E[T_s] = \frac{1}{1-\mathbf{P}(s,s)}$.
The variance of the residence time distribution is $Var[T_s] = \frac{\mathbf{P}(s,s)}{(1-\mathbf{P}(s,s))^2}$.

Recall that the geometric distribution is the only discrete probability distribution that possesses the memoryless property.

Determining n-step transition probabilities

n-step transition probabilities

The probability to move from s to s' in $n \in \mathbb{N}$ steps is inductively defined:

$$p_{s,s'}(0) \ = \ 1 \quad \text{if } s = s', \quad \text{and 0 otherwise,}$$

$p_{s,s'}(1) = \mathbf{P}(s,s')$, and for $n > 1$ by the Chapman-Kolmogorov equation:

$$p_{s,s'}(n) = \sum_{s''} p_{s,s''}(l) \cdot p_{s'',s'}(n{-}l) \quad \text{for all } 0 < l < n$$

For $l = 1$ and $n > 0$ we obtain: $p_{s,s'}(n) = \sum_{s''} p_{s,s''}(1) \cdot p_{s'',s'}(n{-}1)$

$\mathbf{P}^{(n)} = \mathbf{P}^{(1)} \cdot \mathbf{P}^{(n-1)} = \mathbf{P} \cdot \mathbf{P}^{(n-1)}$ is the n-step transition probability matrix

Repeating this scheme: $\mathbf{P}^{(n)} = \mathbf{P} \cdot \mathbf{P}^{(n-1)} = \ldots = \mathbf{P}^{n-1} \cdot \mathbf{P}^{(1)} = \mathbf{P}^n$.

Transient probability distribution

Transient distribution

$\mathbf{P}^n(s,t)$ equals the probability of being in state t after n steps given that the computation starts in s.

The probability of DTMC $\mathcal{D}$ being in state t after exactly n transitions is:

$$\Theta_n^{\mathcal{D}}(t) \ = \ \sum_{s \in S} \iota_{\text{init}}(s) \cdot \mathbf{P}^n(s,t)$$

$\Theta_n^{\mathcal{D}}(t)$ is called the *transient state probability* at epoch n for state t. The function $\Theta_n^{\mathcal{D}}$ is the *transient state distribution* at epoch n of DTMC $\mathcal{D}$.

When considering $\Theta_n^{\mathcal{D}}$ as vector $(\Theta_n^{\mathcal{D}})_{t \in S}$ we have:

$$\Theta_n^{\mathcal{D}} \ = \ \iota_{\text{init}} \cdot \underbrace{\mathbf{P} \cdot \mathbf{P} \cdot \ldots \cdot \mathbf{P}}_{n \text{ times}} \ = \ \iota_{\text{init}} \cdot \mathbf{P}^n.$$

Overview

Paths in a DTMC

State graph

The *state graph* of DTMC $\mathcal{D}$ is a digraph $G = (V, E)$ with V are the states of $\mathcal{D}$, and $(s, s') \in E$ iff $\mathbf{P}(s, s') > 0$.

Paths

Paths in $\mathcal{D}$ are maximal (i.e., infinite) paths in its state graph. Thus, a path is an infinite sequence of states $s_0 s_1 s_2 \ldots\ldots$ with $\mathbf{P}(s_i, s_{i+1}) > 0$ for all i.

Let $Paths(\mathcal{D})$ denote the set of paths in $\mathcal{D}$, and $Paths^*(\mathcal{D})$ the set of finite prefixes thereof.

Direct successors and predecessors

$Post(s) = \{ s' \in S \mid \mathbf{P}(s, s') > 0 \}$ and $Pre(s) = \{ s' \in S \mid \mathbf{P}(s', s) > 0 \}$ are the set of direct successors and predecessors of s respectively. $Post^*(s)$ and $Pre^*(s)$ are the reflexive and transitive closure of $Post$ and Pre.

Measurable space

Sample space

A *sample space* Ω of a chance experiment is a set of elements that have a 1-to-1 relationship to the possible outcomes of the experiment.

σ-algebra

A *σ-algebra* is a pair $(\Omega, \mathcal{F})$ with $\Omega \neq \varnothing$ and $\mathcal{F} \subseteq 2^\Omega$ a collection of subsets of sample space Ω such that:

1. $\Omega \in \mathcal{F}$
2. $A \in \mathcal{F} \;\Rightarrow\; \Omega - A \in \mathcal{F}$ complement
3. $(\forall i \geqslant 0.\ A_i \in \mathcal{F}) \;\Rightarrow\; \bigcup_{i \geqslant 0} A_i \in \mathcal{F}$ countable union

The elements in $\mathcal{F}$ of a σ-algebra $(\Omega, \mathcal{F})$ are called *events*.
The pair $(\Omega, \mathcal{F})$ is called a *measurable space*.

Let Ω be a set. $\mathcal{F} = \{ \varnothing, \Omega \}$ yields the smallest σ-algebra; $\mathcal{F} = 2^\Omega$ yields the largest one.

Probabilities

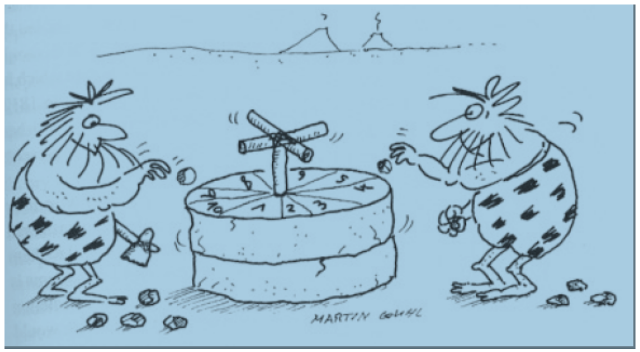

Probability space

Probability space

A *probability space* $\mathcal{P}$ is a structure $(\Omega, \mathcal{F}, Pr)$ with:

- $(\Omega, \mathcal{F})$ is a σ-algebra, and
- $Pr : \mathcal{F} \to [0, 1]$ is a *probability measure*, i.e.:
  1. $Pr(\Omega) = 1$, i.e., Ω is the certain event
  2. $Pr\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} Pr(A_i)$ for any $A_i \in \mathcal{F}$ with $A_i \cap A_j = \varnothing$ for $i \neq j$, where $\{ A_i \}_{i \in I}$ is finite or countably infinite.

The elements in $\mathcal{F}$ of a probability space $(\Omega, \mathcal{F}, Pr)$ are called *measurable* events.

Paths and probabilities

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

Intuition

For a given state s in DTMC $\mathcal{D}$:

- Sample space := set of all infinite paths starting in s
- Events := sets of infinite paths starting in s
- Basic events := cylinder sets
- Cylinder set of finite path $\hat{\pi}$:= set of all infinite continuations of $\hat{\pi}$

Probability measure on DTMCs

Cylinder set

The *cylinder set* of finite path $\hat{\pi} = s_0\, s_1 \ldots s_n \in \mathit{Paths}^*(\mathcal{D})$ is defined by:

$$\mathit{Cyl}(\hat{\pi}) \;=\; \{\, \pi \in \mathit{Paths}(\mathcal{D}) \mid \hat{\pi} \text{ is a prefix of } \pi \,\}$$

The cylinder set spanned by finite path $\hat{\pi}$ thus consists of all infinite paths that have prefix $\hat{\pi}$. Cylinder sets serve as basic events of the smallest σ-algebra on $\mathit{Paths}(\mathcal{D})$.

σ-algebra of a DTMC

The σ-algebra associated with DTMC $\mathcal{D}$ is the smallest σ-algebra that contains all cylinder sets $\mathit{Cyl}(\hat{\pi})$ where $\hat{\pi}$ ranges over all finite path fragments in $\mathcal{D}$.

Probability measure on DTMCs

Cylinder set

The cylinder set of finite path $\hat{\pi} = s_0\, s_1 \ldots s_n \in \mathit{Paths}^*(\mathcal{D})$ is defined by:

$$\mathit{Cyl}(\hat{\pi}) \;=\; \{\, \pi \in \mathit{Paths}(\mathcal{D}) \mid \hat{\pi} \text{ is a prefix of } \pi \,\}$$

Probability measure

Pr is the unique *probability measure* on the σ-algebra on *$\mathit{Paths}(\mathcal{D})$* defined by:

$$\mathit{Pr}(\mathit{Cyl}(s_0 \ldots s_n)) \;=\; \iota_{\text{init}}(s_0) \cdot \mathbf{P}(s_0\, s_1 \ldots s_n)$$

where $\mathbf{P}(s_0\, s_1 \ldots s_n) \;=\; \prod_{0 \leqslant i < n} \mathbf{P}(s_i, s_{i+1})$ for $n > 0$ and $\mathbf{P}(s_0) = 1$.

Some events of interest

Let DTMC $\mathcal{D}$ with (possibly infinite) state space S.

(Simple) reachability

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

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

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

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

Constrained reachability

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

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

More events of interest

Repeated reachability

Repeatedly visit a state in G; formally:

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

Persistence

Eventually reach in a state in G and always stay there; formally:

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

Measurability

Measurability theorem

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

Proof:

To show this, every event will be expressed as allowed operations (complement and/or countable unions) of the events — our cylinder sets!— in the σ-algebra on infinite paths in a DTMC.

Note that $\Box G = \overline{\Diamond \overline{G}}$ and $\Diamond\Box G = \overline{\Box\Diamond\overline{G}}$.

It remains to prove the measurability for the remaining three cases.

Proof for $\Diamond G$

Which event (in our σ-algebra) does $\Diamond G$ formally mean?

the union of all cylinders $Cyl(s_0 \ldots s_n)$ where

$s_0 \ldots s_n$ is a finite path in $\mathcal{D}$ with $s_0, \ldots, s_{n-1} \notin G$ and $s_n \in G$, i.e.,

$$\Diamond G \;=\; \bigcup_{s_0 \ldots s_n \in Paths^*(\mathcal{D}) \cap (S \setminus G)^* G} Cyl(s_0 \ldots s_n)$$

Thus $\Diamond G$ is measurable.

As all cylinder sets are pairwise disjoint, its probability is defined by:

$$\begin{aligned} Pr(\Diamond G) &= \sum_{s_0 \ldots s_n \in Paths^*(\mathcal{D}) \cap (S \setminus G)^* G} Pr\big(Cyl(s_0 \ldots s_n)\big) \\ &= \sum_{s_0 \ldots s_n \in Paths^*(\mathcal{D}) \cap (S \setminus G)^* G} \iota_{\text{init}}(s_0) \cdot \mathbf{P}(s_0 \ldots s_n) \end{aligned}$$

A similar proof strategy applies to the case $\overline{F} \,\mathsf{U}\, G$.

Reachability probabilities: Knuth's die

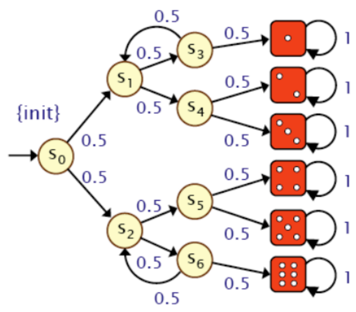

- Consider the event $\Diamond 4$
- Using the previous theorem we obtain:

  $$Pr(\Diamond 4) = \sum_{s_0 \ldots s_n \in (S \setminus 4^*) 4} \mathbf{P}(s_0 \ldots s_n)$$

- This yields: $\mathbf{P}(s_0 s_2 s_5 4) + \mathbf{P}(s_0 s_2 s_6 s_2 s_5 4) + \ldots\ldots$
- Or: $\displaystyle\sum_{k=0}^{\infty} \mathbf{P}(s_0 s_2 (s_6 s_2)^k s_5 4)$
- Or: $\displaystyle\frac{1}{8} \sum_{k=0}^{\infty} \left(\frac{1}{4}\right)^k$
- Geometric series: $\displaystyle\frac{1}{8} \cdot \frac{1}{1-\frac{1}{4}} = \frac{1}{8} \cdot \frac{4}{3} = \frac{1}{6}$

There is however an simpler way to obtain reachability probabilities!

Reachability probabilities in finite DTMCs

Problem statement

Let $\mathcal{D}$ be a DTMC 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 $\mathcal{D}$ 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 $s \in Pre^*(G) \setminus G$:

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

Reachability probabilities: Knuth's die

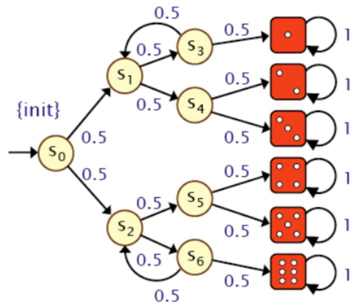

- Consider the event $\Diamond 4$
- Using the previous characterisation we obtain:

$x_1 = x_2 = x_3 = x_5 = x_6 = 0$ and $x_4 = 1$

$x_{s_1} = x_{s_3} = x_{s_4} = 0$

$x_{s_0} = \frac{1}{2} x_{s_1} + \frac{1}{2} x_{s_2}$

$x_{s_2} = \frac{1}{2} x_{s_5} + \frac{1}{2} x_{s_6}$

$x_{s_5} = \frac{1}{2} x_5 + \frac{1}{2} x_4$

$x_{s_6} = \frac{1}{2} x_{s_2} + \frac{1}{2} x_6$

- Gaussian elimination yields:

$x_{s_5} = \frac{1}{2}$, $x_{s_2} = \frac{1}{3}$, $x_{s_6} = \frac{1}{6}$, and $\boxed{x_{s_0} = \frac{1}{6}}$

Linear equation system

Reachability probabilities as linear equation system

- Let $S_? = Pre^*(G) \setminus G$, the states that can reach G by > 0 steps
- $\mathbf{A} = \big(\mathbf{P}(s,t)\big)_{s,t \in S_?}$, the transition probabilities in $S_?$
- $\mathbf{b} = (b_s)_{s \in S_?}$, the probs to reach G in 1 step, i.e., $b_s = \sum_{u \in G} \mathbf{P}(s,u)$

Then: $\mathbf{x} = (x_s)_{s \in S_?}$ with $x_s = Pr(s \models \Diamond G)$ is the unique solution of:

$$\mathbf{x} = \mathbf{A} \cdot \mathbf{x} + \mathbf{b} \quad \text{or} \quad (\mathbf{I} - \mathbf{A}) \cdot \mathbf{x} = \mathbf{b}$$

where $\mathbf{I}$ is the identity matrix of cardinality $|S_?| \times |S_?|$.

Reachability probabilities: Knuth's die

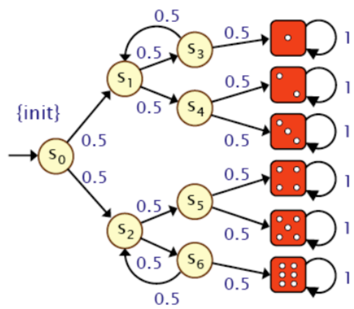

- Consider the event $\Diamond 4$
- $S_? = \{\, s_0, s_2, s_5, s_6 \,\}$

$$\begin{pmatrix} 1 & -\frac{1}{2} & 0 & 0 \\ 0 & 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & 0 & 1 & 0 \\ 0 & -\frac{1}{2} & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} x_{s_0} \\ x_{s_2} \\ x_{s_5} \\ x_{s_6} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ \frac{1}{2} \\ 0 \end{pmatrix}$$

- Gaussian elimination yields:

$x_{s_5} = \frac{1}{2}$, $x_{s_2} = \frac{1}{3}$, $x_{s_6} = \frac{1}{6}$, and $\boxed{x_{s_0} = \frac{1}{6}}$

Constrained reachability probabilities

Problem statement

Let $\mathcal{D}$ be a DTMC with finite state space S, $s \in S$ and $\overline{F}, G \subseteq S$.

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

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

Characterisation of constrained reachability probabilities

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

$$x_s \;=\; \sum_{t \in S \setminus G} \mathbf{P}(s,t) \cdot x_t \;+\; \sum_{u \in G} \mathbf{P}(s,u)$$

Iteratively computing reachability probabilities

Theorem

The vector $\mathbf{x} = \Big(Pr(s \models \overline{F} \,\mathsf{U}\, G)\Big)_{s \in S_?}$ is the *unique* solution of:

$$\mathbf{y} \;=\; \mathbf{A} \cdot \mathbf{y} + \mathbf{b}$$

with $\mathbf{A}$ and $\mathbf{b}$ as defined before.

Furthermore, let:

$$\mathbf{x}^{(0)} = \mathbf{0} \quad \text{and} \quad \mathbf{x}^{(i+1)} \;=\; \mathbf{A} \cdot \mathbf{x}^{(i)} + \mathbf{b} \text{ for } 0 \leqslant i.$$

Then:

1. $\mathbf{x}^{(n)}(s) \;=\; Pr(s \models \overline{F} \,\mathsf{U}^{\leqslant n}\, G)$ for $s \in S_?$
2. $\mathbf{x}^{(0)} \leqslant \mathbf{x}^{(1)} \leqslant \mathbf{x}^{(2)} \leqslant \ldots \leqslant \mathbf{x}$
3. $\mathbf{x} = \lim_{n \to \infty} \mathbf{x}^{(n)}$

where $\overline{F} \,\mathsf{U}^{\leqslant n} G$ contains those paths that reach G via $\overline{F}$ within n steps.

Remark

Iterative algorithms to compute x

There are various algorithms to compute $\mathbf{x} = \lim_{n \to \infty} \mathbf{x}^{(n)}$ where:

$$\mathbf{x}^{(0)} = \mathbf{0} \quad \text{and} \quad \mathbf{x}^{(i+1)} \;=\; \mathbf{A} \cdot \mathbf{x}^{(i)} + \mathbf{b} \text{ for } 0 \leqslant i.$$

The Power method computes vectors $\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \ldots$ and aborts if:

$$\max_{s \in S_?} |\, x_s^{(n+1)} - x_s^{(n)} \,| \;<\; \varepsilon \quad \text{for some small tolerance } \varepsilon$$

This technique guarantees convergence.

Alternative iterative techniques: e.g., Jacobi or Gauss-Seidel, successive overrelaxation (SOR).

Example: Knuth's die

- Let $G = \{\, 1, 2, 3, 4, 5, 6 \,\}$
- Then $Pr(s_0 \models \Diamond G) = 1$
- And $Pr(s_0 \models \Diamond^{\leqslant k} G)$ for $k \in \mathbb{N}$ is given by:

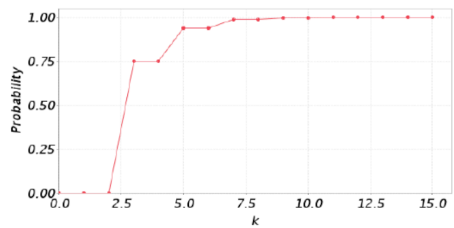

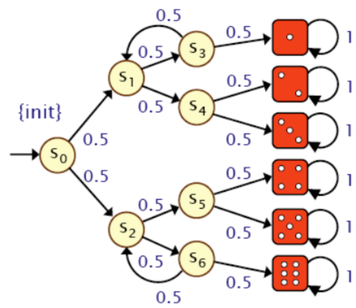

Reachability probability = transient probabilities

Aim

Compute $Pr(\Diamond^{\leqslant n} G)$ in DTMC $\mathcal{D}$. Observe that once a path π reaches G, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Let DTMC $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ and $G \subseteq S$. The DTMC $\mathcal{D}[G] = (S, \mathbf{P}_G, \iota_{\text{init}}, AP, L)$ 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

$$\underbrace{Pr(\Diamond^{\leqslant n} G)}_{\text{reachability in } \mathcal{D}} = \underbrace{Pr(\Diamond^{= n} G)}_{\text{reachability in } \mathcal{D}[G]} = \underbrace{\iota_{\text{init}} \cdot \mathbf{P}_G^n}_{\text{in } \mathcal{D}[G]} = \Theta_n^{\mathcal{D}[G]}$$

Constrained reachability = transient probabilities

Aim

Compute $Pr(\overline{F} \,\mathsf{U}^{\leqslant n}\, G)$ in DTMC $\mathcal{D}$. Observe (as before) that once a path π reaches G via $\overline{F}$, then the remaining behaviour along π is not important. Now also observe that once $s \in F \setminus G$ is reached, 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

$$\underbrace{Pr(\overline{F} \,\mathsf{U}^{\leqslant n}\, G)}_{\text{reachability in } \mathcal{D}} = \underbrace{Pr(\Diamond^{= n} G)}_{\text{reachability in } \mathcal{D}[F \cup G]} = \underbrace{\iota_{\text{init}} \cdot \mathbf{P}_{F \cup G}^n}_{\text{in } \mathcal{D}[F \cup G]} = \Theta_n^{\mathcal{D}[F \cup G]}$$

Overview

Qualitative properties

Quantitative properties

Comparing the probability of an event such as $\Box G$, $\Diamond\Box G$ and $\Box\Diamond G$ with a threshold $\sim p$ with $p \in (0,1)$ and $\sim$ a binary comparison operator $(=, <, \leqslant, \geqslant, >)$ yields a quantitative property.

Example quantitative properties

$Pr(s \models \Diamond\Box G) > \frac{1}{2}$ or $Pr(s \models \Diamond^{\leqslant n} G) \leqslant \frac{\pi}{5}$

Qualitative properties

Comparing the probability of an event such as $\Box G$, $\Diamond\Box G$ and $\Box\Diamond G$ with a threshold > 0 or $= 1$ yields a qualitative property. Any event E with $Pr(E) = 1$ is called almost surely.

Example qualitative properties

$Pr(s \models \Diamond\Box G) > 0$ or $Pr(s \models \Diamond^{\leqslant n} G) = 1$

Verifying qualitative properties

Remark

In the following we will concentrate on almost sure events, i.e., events E with $Pr(E) = 1$. This suffices, as $Pr(E) > 0$ if and only if not $Pr(\overline{E}) = 1$.

Graph notions

Let $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be a (possibly infinite) DTMC.

Strongly connected component

- $T \subseteq S$ is *strongly connected* if for any $s, t \in T$, states s and $t \in T$ are mutually reachable via edges in T.
- T is a *strongly connected component* (SCC) of $\mathcal{D}$ if it is strongly connected and no proper superset of T is strongly connected.
- SCC T is a *bottom SCC* (BSCC) if no state outside T is reachable from T, i.e., for any state $s \in T$, $\mathbf{P}(s, T) = \sum_{t \in T} \mathbf{P}(s, t) = 1$.
- Let $BSCC(\mathcal{D})$ denote the set of BSCCs of DTMC $\mathcal{D}$.

Evolution of an example DTMC

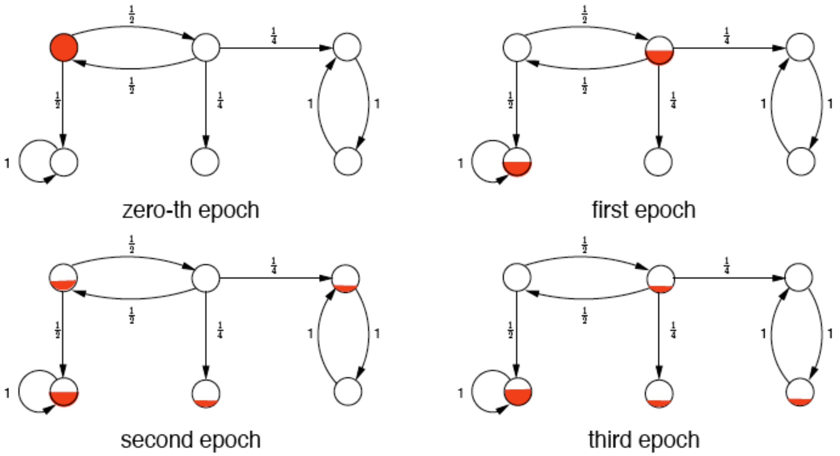

Which states have a probability > 0 when repeating this on the long run?

On the long run

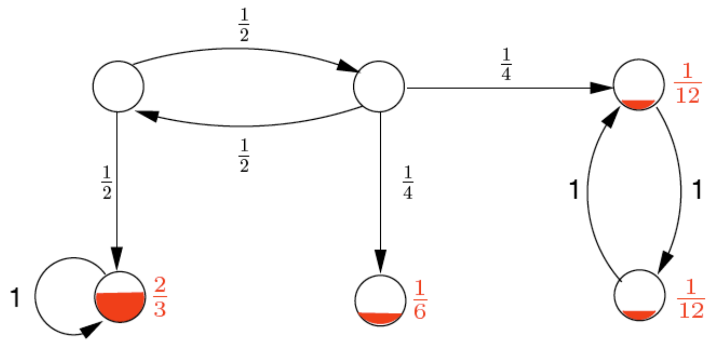

The probability mass on the long run is only left in BSCCs.

Measurability

Lemma

For any state s in (possibly infinite) DTMC $\mathcal{D}$:

$$\{ \pi \in Paths(s) \mid \inf(\pi) \in BSCC(\mathcal{D}) \} \text{ is measurable}$$

where $\inf(\pi)$ is the set of states that are visited infinitely often along π.

Proof:

1. For BSCC T, $\{ \pi \in Paths(s) \mid \inf(\pi) = T \}$ is measurable as:
$$\{ \pi \in Paths(s) \mid \inf(\pi) = T \} = \bigcap_{t \in T} \Box\Diamond t \cap \Diamond\Box T.$$
2. As $BSCC(\mathcal{D})$ is countable, we have:
$$\{ \pi \in Paths(s) \mid \inf(\pi) \in BSCC(\mathcal{D}) \} = \bigcup_{TS \in BSCC(\mathcal{D})} \bigcap_{t \in T} \Box\Diamond t \wedge \Diamond\Box T.$$

Fundamental result

Long-run theorem

For each state s of a finite Markov chain $\mathcal{D}$:

$$Pr_s\{ \pi \in Paths(s) \mid \inf(\pi) \in BSCC(\mathcal{M}) \} = 1.$$

Intuition

Almost surely any finite DTMC eventually reaches a BSCC and visits all its states infinitely often.

Almost sure reachability

Recall: an absorbing state in a DTMC is a state with a self-loop with probability one.

Almost sure reachability theorem

For finite DTMC with state space S, $s \in S$ and $G \subseteq S$ a set of absorbing states:

$$Pr(s \models \Diamond G) = 1 \quad \text{iff} \quad s \in S \setminus Pre^*(S \setminus Pre^*(G)).$$

Note: $S \setminus Pre^*(S \setminus Pre^*(G))$ are states that cannot reach states from which G cannot be reached.

Proof:

Show that both sides of the equivalence are equivalent to $Post^*(t) \cap G \neq \varnothing$ for each state $t \in Post^*(s)$. Rather straightforward.

Computing almost sure reachability properties

Aim:

For finite DTMC $\mathcal{D}$ and $G \subseteq S$, determine $\{ s \in S \mid Pr(s \models \Diamond G) = 1 \}$.

Algorithm

1. Make all states in G absorbing yielding $\mathcal{D}[G]$.
2. Determine $S \setminus Pre^*(S \setminus Pre^*(G))$ by a graph analysis:
   2.1 do a backward search from G in $\mathcal{D}[G]$ to determine $Pre^*(G)$.
   2.2 followed by a backward search from $S \setminus Pre^*(G)$ in $\mathcal{D}[G]$.

This yields a time complexity which is linear in the size of the DTMC $\mathcal{D}$.

Repeated reachability

Almost sure repeated reachability theorem

For finite DTMC with state space S, $G \subseteq S$, and $s \in S$:

$Pr(s \models \Box\Diamond G) = 1$ iff for each BSCC $T \subseteq Post^*(s)$. $T \cap G \neq \varnothing$.

Proof:

Immediate consequence of the long-run theorem.

Almost sure repeated reachability

Almost sure repeated reachability theorem

For finite DTMC with state space S, $G \subseteq S$, and $s \in S$:

$Pr(s \models \Box\Diamond G) = 1$ iff for each BSCC $T \subseteq Post^*(s)$. $T \cap G \neq \varnothing$.

Example:

$B = \{ s_3, s_4, s_5 \}$

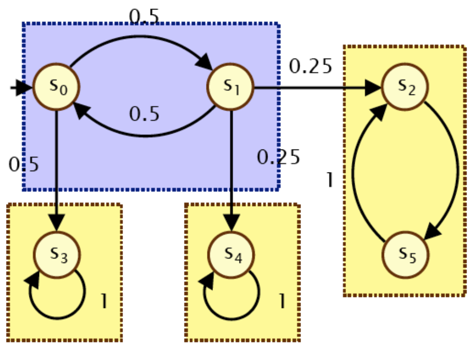

Almost sure persistence

Almost sure persistence theorem

For finite DTMC with state space S, $G \subseteq S$, and $s \in S$:

$Pr(s \models \Diamond\Box G) = 1$ if and only if $T \subseteq G$ for any BSCC $T \subseteq Post^*(s)$

Example:

$\{ s_2, s_3, s_4, s_5 \}$

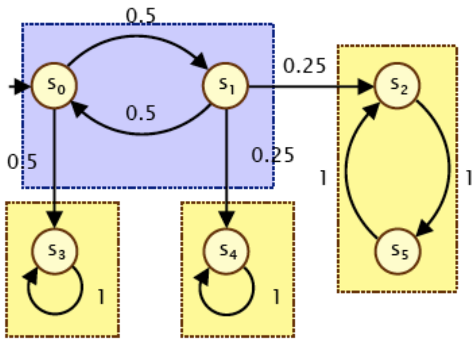

A remark on infinite Markov chains

Graph analysis for infinite DTMCs does not suffice!

Consider the following infinitely countable DTMC, known as random walk:

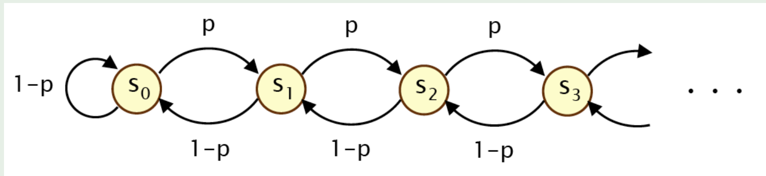

The value of rational probability p does affect qualitative properties:

$$Pr(s \models \Diamond\, s_0) = \begin{cases} 1 & \text{if } p \leqslant \frac{1}{2} \\ < 1 & \text{if } p > \frac{1}{2} \end{cases} \quad \text{and}$$

$$Pr(s \models \Box\Diamond\, s_0) = \begin{cases} 1 \text{ if } p \leqslant \frac{1}{2} \\ 0 \text{ if } p > \frac{1}{2} \end{cases}$$

Quantitative properties

Quantitative repeated reachability theorem

For finite DTMC with state space S, $G \subseteq S$, and $s \in S$:

$$Pr(s \models \Box\Diamond G) \;=\; Pr(s \models \Diamond U)$$

where U is the union of all BSCCs T with $T \cap G \neq \varnothing$.

Quantitative repeated reachability theorem

For finite DTMC with state space S, $G \subseteq S$, and $s \in S$:

$$Pr(s \models \Diamond\Box G) \;=\; Pr(s \models \Diamond U)$$

where U is the union of all BSCCs T with $T \subseteq G$.

Remark

Thus probabilities for $\Box\Diamond G$ and $\Box\Diamond G$ are reduced to reachability probabilities. These can be computed by solving a linear equation system.

Summary

- Executions of a DTMC are strongly fair with respect to all probabilistic choices.
- A finite DTMC almost surely ends up in a BSCC on the long run.
- Almost sure reachability = double backward search.
- Almost sure $\Box\Diamond G$ and $\Diamond\Box G$ properties can be checked by BSCC analysis and reachability.
- Probabilities for $\Box\Diamond G$ and $\Diamond\Box G$ reduce to reachability probabilities.

Take-home message

For finite DTMCs, qualitative properties do only depend on their state graph and not on the transition probabilities! For infinite DTMCs, this does not hold.

Overview

Probabilistic Computation Tree Logic

- PCTL is a language for formally specifying properties over DTMCs.
- It is a branching-time temporal logic based on CTL.
- Formula interpretation is Boolean, i.e., a state satisfies a formula or not.
- 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.

PCTL syntax [Hansson & Jonsson, 1994]

Probabilistic Computation Tree Logic: Syntax

PCTL consists of state- and path-formulas.

- PCTL *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.

- PCTL *path formulae* are formed according to the following grammar:

$$\varphi \ ::= \ \bigcirc\Phi \ \Big| \ \Phi_1 \,\mathsf{U}\, \Phi_2 \ \Big| \ \Phi_1 \,\mathsf{U}^{\leqslant n}\, \Phi_2$$

  where Φ, Φ_1, and Φ_2 are state formulae and $n \in \mathbb{N}$.

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

Probabilistic Computation Tree Logic

- PCTL *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.

- PCTL *path formulae* are formed according to the following grammar:

$$\varphi \ ::= \ \bigcirc\Phi \ \Big| \ \Phi_1 \,\mathsf{U}\, \Phi_2 \ \Big| \ \Phi_1 \,\mathsf{U}^{\leqslant n}\, \Phi_2 \quad \text{where } n \in \mathbb{N}.$$

Intuitive semantics

- $s_0 s_1 s_2 \ldots \models \Phi \,\mathsf{U}^{\leqslant n}\, \Psi$ if Φ holds until Ψ holds within n steps.
- $s \models \mathbb{P}_J(\varphi)$ if probability that paths starting in s fulfill φ lies in J.

Semantics of $\mathbb{P}$-operator

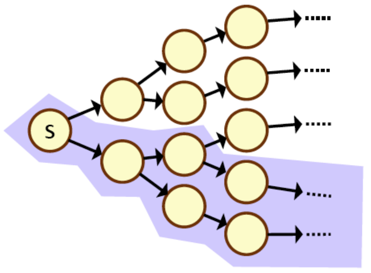

- $s \models \mathbb{P}_J(\varphi)$ if:
  - the probability of all paths starting in s fulfilling φ lies in J.
- Example: $s \models \mathbb{P}_{>\frac{1}{2}}(\Diamond a)$ if
  - the probability to reach an a-labeled state from s exceeds $\frac{1}{2}$.
- Formally:
  - $s \models \mathbb{P}_J(\varphi)$ if and only if $Pr_s\{\, \pi \in Paths(s) \mid \pi \models \varphi \,\} \in J$.

Derived operators

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

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

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

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

Correctness of Knuth's die

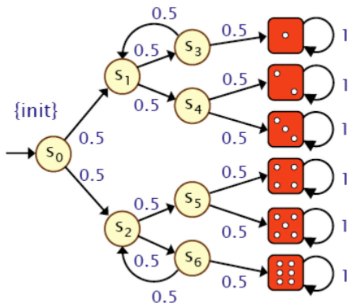

Correctness of Knuth's die

$\mathbb{P}_{=\frac{1}{6}}(\Diamond 1) \ \wedge \ \mathbb{P}_{=\frac{1}{6}}(\Diamond 2) \ \wedge \ \mathbb{P}_{=\frac{1}{6}}(\Diamond 3) \ \wedge \ \mathbb{P}_{=\frac{1}{6}}(\Diamond 4) \ \wedge \ \mathbb{P}_{=\frac{1}{6}}(\Diamond 5) \ \wedge \ \mathbb{P}_{=\frac{1}{6}}(\Diamond 6)$

Measurability

PCTL measurability

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

Proof (sketch):

Three cases:

1. $\bigcirc \Phi$:
   - cylinder sets constructed from paths of length one.
2. $\Phi \, \mathsf{U}^{\leqslant n} \, \Psi$:
   - (finite number of) cylinder sets from paths of length at most n.
3. $\Phi \, \mathsf{U} \, \Psi$:
   - countable union of paths satisfying $\Phi \, \mathsf{U}^{\leqslant n} \, \Psi$ for all $n \geqslant 0$.

PCTL model checking

PCTL model checking problem

Input: a finite DTMC $\mathcal{D} = (S, \mathbf{P}, \iota_{\mathrm{init}}, AP, L)$, state $s \in S$, and PCTL 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: $s \models \mathbb{P}_J(\bigcirc\, \Phi)$ if and only if $Pr(s \models \bigcirc\, \Phi) \in J$.

Lemma

$Pr(s \models \bigcirc\, \Phi) \;=\; \sum_{s' \in Sat(\Phi)} \mathbf{P}(s, s')$.

Algorithm

Considering the above equation for all states simultaneously yields:

$$\big(Pr(s \models \bigcirc\, \Phi)\big)_{s \in S} \;=\; \mathbf{P} \cdot \mathbf{b}_\Phi$$

with $\mathbf{b}_\Phi$ the characteristic vector of $Sat(\Phi)$, i.e., $b_\Phi(s) = 1$ iff $s \in Sat(\Phi)$.

Checking the next-step operator reduces to a single matrix-vector multiplication.

Example

Consider DTMC:

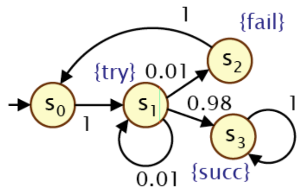

and PCTL-formula:

$$\mathbb{P}_{\geqslant 0.9}\,(\bigcirc\,(\neg try \vee succ))$$

1. $Sat(\neg try \vee succ) \;=\; (S \setminus Sat(try)) \cup Sat(succ) \;=\; \{ s_0, s_2, s_3 \}$
2. We know: $\big(Pr(s \models \bigcirc\, \Phi)\big)_{s \in S} \;=\; \mathbf{P} \cdot \mathbf{b}_\Phi$ where $\Phi = \neg try \vee succ$
3. Applying that to this example yields:

$$\big(Pr(s \models \bigcirc\, \Phi)\big)_{s \in S} \;=\; \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0.01 & 0.01 & 0.98 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 \\ 0 \\ 1 \\ 1 \end{pmatrix} \;=\; \begin{pmatrix} 0 \\ 0.99 \\ 1 \\ 1 \end{pmatrix}$$

4. Thus: $Sat(\mathbb{P}_{\geqslant 0.9}(\bigcirc\,(\neg try \vee succ)) \;=\; \{ s_1, s_2, s_3 \}$.

Time complexity

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

Time complexity of PCTL model checking

For finite DTMC $\mathcal{D}$ and PCTL state-formula Φ, the PCTL model-checking problem can be solved in time

$$\mathcal{O}(\, poly(size(\mathcal{D})) \cdot n_{\max} \cdot |\Phi| \,)$$

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

Time complexity

Time complexity of PCTL model checking

For finite DTMC $\mathcal{D}$ and PCTL state-formula Φ, the PCTL model-checking problem can be solved in time

$$\mathcal{O}(\, poly(size(\mathcal{D})) \cdot n_{\max} \cdot |\Phi| \,).$$

Proof (sketch)

1. For each node in the parse tree, a model-checking is performed; this yields a linear complexity in $|\Phi|$.
2. The worst-case operator is (unbounded) until.
   2.1 Determining $S_{=0}$ and $S_{=1}$ can be done in linear time.
   2.2 Direct methods to solve linear equation systems are in $\Theta(|S_?|^3)$.
3. Strictly speaking, $\mathsf{U}^{\leqslant n}$ could be more expensive for large n.
   But it remains polynomial, and n is small in practice.

Some practical verification times

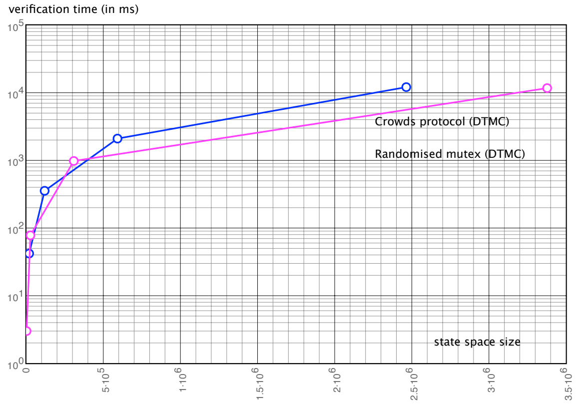

- command-line tool MRMC ran on a Pentium 4, 2.66 GHz, 1 GB RAM laptop.
- PCTL formula $\mathbb{P}_{\leqslant p}(\Diamond obs)$ where *obs* holds when the sender's id is detected.

Summary

- PCTL is a variant of CTL with operator $\mathbb{P}_J(\varphi)$.
- Sets of paths fulfilling PCTL path-formula φ are measurable.
- PCTL model checking is performed by a recursive descent over Φ.
- The next operator amounts to a single matrix-vector multiplication.
- The bounded-until operator $\mathsf{U}^{\leqslant n}$ amounts to n matrix-vector multiplications.
- The until-operator amounts to solving a linear equation system.
- The worst-case time complexity is polynomial in the size of the DTMC and linear in the size of the formula.

Overview

Qualitative PCTL

Qualitative PCTL

State formulae in the *qualitative fragment* of PCTL (over *AP*):

$$\Phi ::= \mathrm{true} \;\Big|\; a \;\Big|\; \Phi_1 \wedge \Phi_2 \;\Big|\; \neg\Phi \;\Big|\; \mathbb{P}_{>0}(\varphi) \;\Big|\; \mathbb{P}_{=1}(\varphi)$$

where $a \in AP$, and φ is a path formula formed according to the grammar:

$$\varphi ::= \bigcirc\,\Phi \;\Big|\; \Phi_1 \,\mathsf{U}\, \Phi_2.$$

Remark

The probability bounds $= 0$ and < 1 can be derived:

$$\mathbb{P}_{=0}(\varphi) \equiv \neg\mathbb{P}_{>0}(\varphi) \quad \text{and} \quad \mathbb{P}_{<1}(\varphi) \equiv \neg\mathbb{P}_{=1}(\varphi)$$

So, in qualitative PCTL, there is no bounded until, and only > 0, $= 0$, > 1 and $= 1$ thresholds.

Qualitative PCTL

Qualitative PCTL

State formulae in the *qualitative fragment* of PCTL (over AP):

$$\Phi ::= \text{true} \;\big|\; a \;\big|\; \Phi_1 \wedge \Phi_2 \;\big|\; \neg\Phi \;\big|\; \mathbb{P}_{>0}(\varphi) \;\big|\; \mathbb{P}_{=1}(\varphi)$$

where $a \in AP$, and φ is a path formula formed according to the grammar:

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

Examples

$\mathbb{P}_{=1}(\Diamond\mathbb{P}_{>0}(\bigcirc\, a))$ and $\mathbb{P}_{<1}(\mathbb{P}_{>0}(\Diamond a)\,\mathsf{U}\, b)$ are qualitative PCTL formulas.

CTL versus qualitative PCTL

Equivalence of PCTL and CTL Formulae

The PCTL formula Φ is *equivalent* to the CTL formula Ψ, denoted $\Phi \equiv \Psi$, if $Sat(\Phi) = Sat(\Psi)$ for each DTMC $\mathcal{D}$.

Example

The simplest such cases are path formulae involving the next-step operator:

$$\begin{aligned}\mathbb{P}_{=1}(\bigcirc\, a) &\equiv \forall\bigcirc\, a\\ \mathbb{P}_{>0}(\bigcirc\, a) &\equiv \exists\bigcirc\, a\end{aligned}$$

And for $\exists\Diamond$ and $\forall\Box$ we have:

$$\begin{aligned}\mathbb{P}_{>0}(\Diamond a) &\equiv \exists\Diamond a\\ \mathbb{P}_{=1}(\Box a) &\equiv \forall\Box a.\end{aligned}$$

CTL versus qualitative PCTL

(1) $\mathbb{P}_{>0}(\Diamond a) \equiv \exists\Diamond a$ and (2) $\mathbb{P}_{=1}(\Box a) \equiv \forall\Box a$.

Proof:

(1) Consider the first statement.

- $\Rightarrow$ Assume $s \models \mathbb{P}_{>0}(\Diamond a)$. By the PCTL semantics, $Pr(s \models \Diamond a) > 0$. Thus, $\{\, \pi \in Paths(s) \mid \pi \models \Diamond a \,\} \neq \varnothing$, and hence, $s \models \exists\Diamond a$.
- $\Leftarrow$ Assume $s \models \exists\Diamond a$, i.e., there is a finite path $\hat{\pi} = s_0\, s_1 \ldots s_n$ with $s_0 = s$ and $s_n \models a$. It follows that all paths in the cylinder set $Cyl(\hat{\pi})$ fulfill $\Diamond a$. Thus:

$$Pr(s \models \Diamond a) \;\geqslant\; Pr_s(Cyl(s_0\, s_1 \ldots s_n)) \;=\; \mathbf{P}(s_0 s_1 \ldots s_n) \;>0.$$

So, $s \models \mathbb{P}_{>0}(\Diamond a)$.

(2) The second statement follows by duality.

CTL versus qualitative PCTL

(1) $\mathbb{P}_{>0}(\Diamond a) \equiv \exists\Diamond a$ and (2) $\mathbb{P}_{=1}(\Box a) \equiv \forall\Box a$.

(3) $\mathbb{P}_{>0}(\Box a) \not\equiv \exists\Box a$ and (4) $\mathbb{P}_{=1}(\Diamond a) \not\equiv \forall\Diamond a$.

Example

Consider the second statement (4). Let s be a state in a (possibly infinite) DTMC. Then: $s \models \forall\Diamond a$ implies $s \models \mathbb{P}_{=1}(\Diamond a)$. The reverse direction, however, does not hold. Consider the example DTMC:

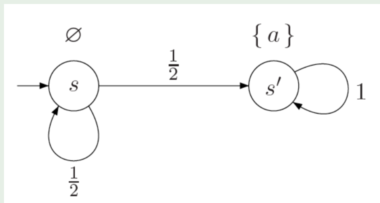

$s \models \mathbb{P}_{=1}(\Diamond a)$ as the probability of path s^ω is zero. However, the path s^ω is possible and violates $\Diamond a$. Thus, $s \not\models \forall\Diamond a$.

Statement (3) follows by duality.

Almost-sure-reachability not in CTL

Almost-sure-reachability not in CTL

1. There is no CTL formula that is equivalent to $\mathbb{P}_{=1}(\Diamond a)$.
2. There is no CTL formula that is equivalent to $\mathbb{P}_{>0}(\Box a)$.

Proof:

We provide the proof of 1.; 2. follows by duality: $\mathbb{P}_{=1}(\Diamond a) \equiv \neg\mathbb{P}_{>0}(\Box\neg a)$. By contraposition. Assume $\Phi \equiv \mathbb{P}_{=1}(\Diamond a)$. Consider the infinite DTMC $\mathcal{D}_p$:

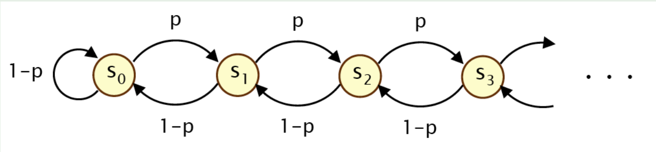

The value of p does affect reachability: $Pr(s \models \Diamond\, s_0) = \begin{cases} 1 & \text{if } p \leqslant \frac{1}{2} \\ < 1 & \text{if } p > \frac{1}{2} \end{cases}$

Almost-sure-reachability not in CTL

There is no CTL formula that is equivalent to $\mathbb{P}_{=1}(\Diamond a)$.

Proof:

We have: $Pr(s \models \Diamond\, s_0) = \begin{cases} 1 & \text{if } p \leqslant \frac{1}{2} \\ < 1 & \text{if } p > \frac{1}{2} \end{cases}$

Thus, in $\mathcal{D}_{\frac{1}{4}}$ we have $s \models \mathbb{P}_{=1}(\Diamond s_0)$ for all states s, while in $\mathcal{D}_{\frac{3}{4}}$, e.g., $s_1 \not\models \mathbb{P}_{=1}(\Diamond s_0)$. Hence: $s_1 \in Sat_{\mathcal{D}_{\frac{1}{4}}}(\mathbb{P}_{=1}(\Diamond s_0))$ but $s_1 \notin Sat_{\mathcal{D}_{\frac{3}{4}}}(\mathbb{P}_{=1}(\Diamond s_0))$.
For CTL-formula Φ —by assumption $\Phi \equiv \mathbb{P}_{=1}(\Diamond s_0)$— we have:

$$Sat_{\mathcal{D}_{\frac{1}{4}}}(\Phi) \;=\; Sat_{\mathcal{D}_{\frac{3}{4}}}(\Phi).$$

Hence, state s_1 either fulfills the CTL formula Φ in both DTMCs or in none of them. This, however, contradicts $\Phi \equiv \mathbb{P}_{=1}(\Diamond s_0)$.

$\forall\Diamond$ is not expressible in qualitative PCTL

1. There is no qualitative PCTL formula that is equivalent to $\forall\Diamond a$.
2. There is no qualitative PCTL formula that is equivalent to $\exists\Box a$.

Fair CTL

Fair paths

In fair CTL, path formulas are interpreted over fair infinite paths, i.e., paths π that satisfy

$$fair \;=\; \bigwedge_{s \in S} \; \bigwedge_{t \,\in\, Post(s)} (\Box\Diamond s \rightarrow \Box\Diamond t).$$

A path π such that $\pi \models fair$ is called fair. Let $Paths_{fair}(s)$ be the set of fair paths starting in s.

Fair CTL semantics

The fair semantics of CTL is defined by the satisfaction $\models_{fair}$ which is defined as $\models$ for the CTL semantics, except that:

$$\begin{array}{lll} s \models_{fair} \exists\varphi & \text{iff} & \text{there exists } \pi \in Paths_{fair}(s).\, \pi \models_{fair} \varphi \\ s \models_{fair} \forall\varphi & \text{iff} & \text{for all } \pi \in Paths_{fair}(s).\, \pi \models_{fair} \varphi. \end{array}$$

Fairness theorem

Qualitative PCTL versus fair CTL theorem

Let s be an arbitrary state in a finite DTMC. Then:

$$\begin{array}{lll} s \models \mathbb{P}_{=1}(\Diamond a) & \text{iff} & s \models_{fair} \forall \Diamond a \\ s \models \mathbb{P}_{>0}(\Box a) & \text{iff} & s \models_{fair} \exists \Box a \\ s \models \mathbb{P}_{=1}(a \,\mathsf{U}\, b) & \text{iff} & s \models_{fair} \forall (a \,\mathsf{U}\, b) \\ s \models \mathbb{P}_{>0}(a \,\mathsf{U}\, b) & \text{iff} & s \models_{fair} \exists (a \,\mathsf{U}\, b) \end{array}$$

Comparable expressiveness

Qualitative PCTL and fair CTL are equally expressive.

Almost sure repeated reachability

Almost sure repeated reachability is PCTL-definable

For finite DTMC $\mathcal{D}$, state $s \in S$ and $G \subseteq S$:

$$s \models \mathbb{P}_{=1}(\Box \mathbb{P}_{=1}(\Diamond G)) \quad \text{iff} \quad Pr_s\{\pi \in Paths(s) \mid \pi \models \Box \Diamond G\} = 1.$$

We abbreviate $\mathbb{P}_{=1}(\Box \mathbb{P}_{=1}(\Diamond G))$ by $\mathbb{P}_{=1}(\Box \Diamond G)$.

Remark:

For CTL, universal repeated reachability properties can be formalized by the combination of the modalities $\forall \Box$ and $\forall \Diamond$:

$$s \models \forall \Box \forall \Diamond G \quad \text{iff} \quad \pi \models \Box \Diamond G \text{ for all } \pi \in Paths(s).$$

Repeated reachability probabilities

Repeated reachability probabilities are PCTL-definable

For finite DTMC $\mathcal{D}$, state $s \in S$, $G \subseteq S$ and interval $J \subseteq [0, 1]$ we have:

$$s \models \underbrace{\mathbb{P}_J(\Diamond \mathbb{P}_{=1}(\Box \mathbb{P}_{=1}(\Diamond G)))}_{=\mathbb{P}_J(\Box \Diamond G)} \quad \text{if and only if} \quad Pr(s \models \Box \Diamond G) \in J.$$

Remark:

By the above theorem, $\mathbb{P}_{>0}(\Box \Diamond G)$ is PCTL definable. Note that $\exists \Box \Diamond G$ is not CTL-definable (but definable in a combination of CTL and LTL, called CTL*).

Almost sure persistence

Almost sure persistence is PCTL-definable

For finite DTMC $\mathcal{D}$, state $s \in S$ and $G \subseteq S$:

$$s \models \mathbb{P}_{=1}(\Diamond \mathbb{P}_{=1}(\Box G)) \quad \text{iff} \quad Pr_s\{\pi \in Paths(s) \mid \pi \models \Diamond \Box G\} = 1.$$

We abbreviate $\mathbb{P}_{=1}(\Diamond \mathbb{P}_{=1}(\Box G))$ by $\mathbb{P}_{=1}(\Diamond \Box G)$.

Remark:

Note that $\forall \Diamond \Box G$ is not CTL-definable. $\Diamond \Box G$ is a well-known example formula in LTL that cannot be expressed in CTL. But by the above theorem it can be expressed in PCTL.

Persistence probabilities

Persistence probabilities are PCTL-definable

For finite DTMC $\mathcal{D}$, state $s \in S$, $G \subseteq S$ and interval $J \subseteq [0,1]$ we have:

$$s \models \underbrace{\mathbb{P}_J(\Diamond \mathbb{P}_{=1}(\Box G))}_{=\mathbb{P}_J(\Diamond\Box G)} \quad \text{if and only if} \quad Pr(s \models \Diamond\Box G) \in J.$$

Proof:

Left as an exercise. Hint: use the long run theorem.

Summary

- Qualitative PCTL only allow the probability bounds > 0 and $= 1$.
- There is no CTL formula that is equivalent to $\mathbb{P}_{=1}(\Diamond a)$.
- There is no PCTL formula that is equivalent to $\forall\Box a$.
- These results do not apply to finite DTMCs.
- $\mathbb{P}_{=1}(\Diamond a)$ and $\forall\Box a$ are equivalent under fairness.
- Repeated reachability probabilities are PCTL definable.

Take-home messages

Qualitative PCTL and CTL have incomparable expressiveness. Qualitative and fair CTL are equally expressive. Repeated reachability and persistence probabilities are PCTL definable. Their qualitative counterparts are not expressible in CTL.

Overview

Probabilistic bisimulation: intuition

Intuition

- Strong bisimulation is used to compare labeled transition systems.
- Strongly bisimilar states exhibit the same step-wise behaviour.
- Our aim: adapt bisimulation to discrete-time Markov chains.
- This yields a probabilistic variant of strong bisimulation.

- When do two DTMC states exhibit the same step-wise behaviour?
- Key: if their transition probability for each equivalence class coincides.

Probabilistic bisimulation

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 $(s, t) \in R$.

Probabilistic bisimulation

Probabilistic bisimulation

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 $C \in S/R$.

Remarks

As opposed to bisimulation on states in transition systems, any probabilistic bisimulation is an equivalence.

Example

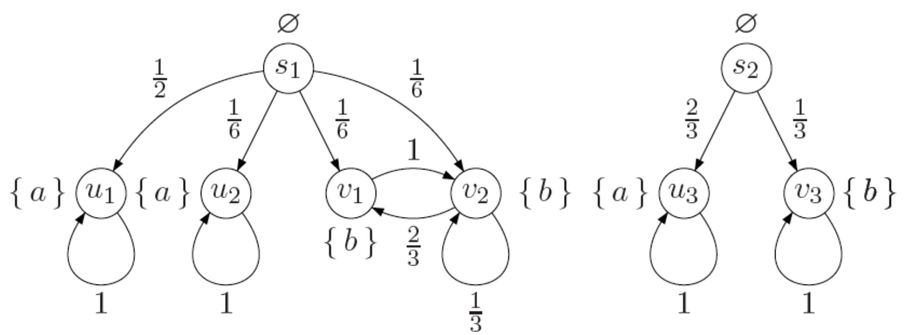

Quotient under $\sim_p$

Quotient DTM under $\sim_p$

For $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ and probabilistic bisimulation $\sim_p \subseteq S \times S$ let

$$\mathcal{D}/\sim_p \;=\; (S', \mathbf{P}', \iota'_{\text{init}}, AP, L'), \quad \text{the } \textit{quotient} \text{ of } \mathcal{D} \text{ under } \sim_p$$

where

- $S' = S/\sim_p = \{ [s]_{\sim_p} \mid s \in S \}$ with $[s]_{\sim_p} = \{ s' \in S \mid s \sim_p s' \}$
- $\mathbf{P}'([s]_{\sim_p}, [s']_{\sim_p}) = \mathbf{P}(s, [s']_{\sim_p})$
- $\iota'_{\text{init}}([s]_{\sim_p}) = \sum_{s' \in [s]_{\sim_p}} \iota_{\text{init}}(s)$
- $L'([s]_{\sim_p}) = L(s)$.

Remarks

The transition probability from $[s]_{\sim_p}$ to $[t]_{\sim_p}$ equals $\mathbf{P}(s, [t]_{\sim_p})$. This is well-defined as $\mathbf{P}(s, C) = \mathbf{P}(s', C)$ for all $s \sim_p s'$ and all bisimulation equivalence classes C.

Craps

- Come-out roll:
  - 7 or 11: win
  - 2, 3, or 12: lose
  - else: roll again
- Next roll(s):
  - 7: lose
  - point: win
  - else: roll again

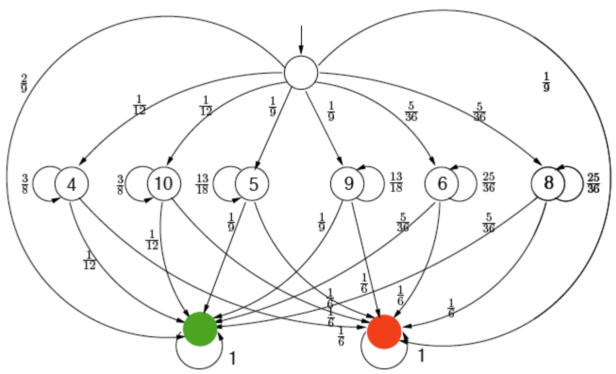

Quotient DTMC of Craps under $\sim_p$

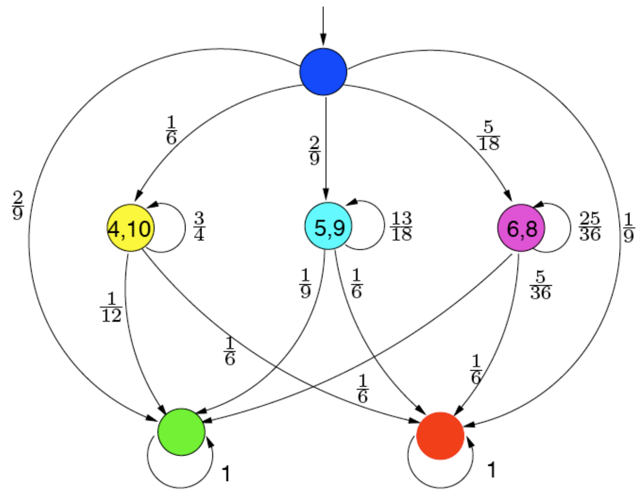

Preservation of PCTL-formulas

Bisimulation preserves PCTL

Let $\mathcal{D}$ be a DTMC and s, t states in $\mathcal{D}$. Then:

$$s \sim_p t \quad \text{if and only if} \quad s \text{ and } t \text{ are PCTL-equivalent.}$$

Remarks

$s \sim_p t$ implies that

1. transient probabilities, reachability probabilities,
2. repeated reachability, persistence probabilities
3. all qualitative PCTL formulas

for s and t are equal.

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

PCTL* syntax

Probabilistic Computation Tree Logic: Syntax

PCTL* consists of state- and path-formulas.

- PCTL* *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.

- PCTL* *path formulae* are formed according to the following grammar:

$$\varphi ::= \Phi \;\big|\; \neg\varphi \;\big|\; \varphi_1 \wedge \varphi_2 \;\big|\; \bigcirc\varphi \;\big|\; \varphi_1 \,\mathsf{U}\, \varphi_2$$

  where Φ is a state formula and φ, φ_1, and φ_2 are path formulae.

Bounded until in PCTL*

Bounded until

Bounded until can be defined using the other operators:

$$\varphi_1 \, \mathsf{U}^{\leqslant n} \varphi_2 \;=\; \bigvee_{0 \leqslant i \leqslant n} \psi_i \quad \text{where } \psi_0 = \varphi_2 \text{ and } \psi_{i+1} = \varphi_1 \wedge \bigcirc \psi_i \text{ for } i \geqslant 0.$$

Examples in PCTL* but not in PCTL

$\mathbb{P}_{>\frac{1}{4}}(\bigcirc a \, \mathsf{U} \bigcirc b)$ and $\mathbb{P}_{=1}(\mathbb{P}_{>\frac{1}{2}}(\Box\Diamond a) \vee \mathbb{P}_{\leqslant\frac{1}{3}}(\Diamond\Box b))$.

Preservation of PCTL*-formulas

Bisimulation preserves PCTL*

Let $\mathcal{D}$ be a DTMC and s, t states in $\mathcal{D}$. Then:

$$s \sim_p t \quad \text{if and only if} \quad s \text{ and } t \text{ are PCTL}^*\text{-equivalent.}$$

Remarks

1. Bisimulation thus preserves not only all PCTL but also all PCTL* formulas.
2. By the last two results it follows that PCTL- and PCTL*-equivalence coincide. Thus any two states that satisfy the same PCTL formulas, satisfy the same PCTL* formulas.

PCTL⁻ syntax

Simple Probabilistic Computation Tree Logic: Syntax

PCTL⁻ only consists of state-formulas. These formulas over the set AP obey the grammar:

$$\Phi \;::=\; a \;\Big|\; \Phi_1 \wedge \Phi_2 \;\Big|\; \mathbb{P}_{\leqslant p}(\bigcirc \Phi)$$

where $a \in AP$ and p is a probability in $[0, 1]$.

Remarks

This is a truly simple logic. It does not contain the until-operator. Negation is not present and cannot be expressed. Only upper bounds on probabilities.

The next theorem shows that PCTL-, PCTL*- and PCTL⁻-equivalence coincide.

Preservation of PCTL

PCTL/PCTL* and Bisimulation Equivalence

Let $\mathcal{D}$ be a DTMC and s_1, s_2 states in $\mathcal{D}$. Then, the following statements are equivalent:

(a) $s_1 \sim_p s_2$.
(b) s_1 and s_2 are PCTL*-equivalent, i.e., fulfill the same PCTL* formulas
(c) s_1 and s_2 are PCTL-equivalent, i.e., fulfill the same PCTL formulas
(d) s_1 and s_2 are PCTL⁻-equivalent, i.e., fulfill the same PCTL⁻ formulas

Proof:

1. (a) $\Longrightarrow$ (b): by structural induction on PCTL* formulas.
2. (b) $\Longrightarrow$ (c): trivial as PCTL is a sublogic of PCTL*.
3. (c) $\Longrightarrow$ (d): trivial as PCTL− is a sublogic of PCTL.
4. (d) $\Longrightarrow$ (a): involved. First finite DTMCs, then for arbitrary DTMCs.

IEEE 802.11 group communication protocol

	original DTMC			quotient DTMC		red. factor	
OD	states	transitions	ver. time	blocks	total time	states	time
4	1125	5369	122	71	13	**15.9**	**9.00**
12	37349	236313	7180	1821	642	**20.5**	**11.2**
20	231525	1590329	50133	10627	5431	**21.8**	**9.2**
28	804837	5750873	195086	35961	24716	**22.4**	**7.9**
36	2076773	15187833	5103900	91391	77694	**22.7**	**6.6**
40	3101445	22871849	7725041	135752	127489	**22.9**	**6.1**

Summary

- Bisimilar states have equal transition probabilities to all equivalence classes.
- $\sim_p$ is the coarsest probabilistic bisimulation.
- In a quotient DTMC all states are equivalence classes under $\sim_p$.
- Bisimulation, i.e., $\sim_p$, and PCTL-equivalence coincide.
- PCTL, PCTL* and PCTL$^-$-equivalence coincide.
- To show $s \not\sim_p t$, show $s \models \Phi$ and $t \not\models \Phi$ for $\Phi \in$ PCTL$^-$.
- Bisimulation may yield up to exponential savings in state space.

Take-home message

Probabilistic bisimulation coincides with a notion from the sixties, named (ordinary) lumpability.

Overview

Paths and traces

Paths

A *path* in DTMC $\mathcal{D}$ is an infinite sequence of states $s_0 s_1 s_2 \ldots\ldots$ with $\mathbf{P}(s_i, s_{i+1}) > 0$ for all i.

Let $Paths(\mathcal{D})$ denote the set of paths in $\mathcal{D}$, and $Paths^*(\mathcal{D})$ the set of finite prefixes thereof.

Trace

The *trace* of path $\pi = s_0\, s_1\, s_2 \ldots$ is $trace(\pi) = L(s_0)\, L(s_1)\, L(s_2) \ldots$. The trace of finite path $\widehat{\pi} = s_0\, s_1 \ldots s_n$ is $trace(\widehat{\pi}) = L(s_0)\, L(s_1) \ldots L(s_n)$.

The set of traces of a set Π of paths: $trace(\Pi) = \{\, trace(\pi) \mid \pi \in \Pi \,\}$.

LT properties

Linear-time property

A *linear-time property* (LT property) over AP is a subset of $(2^{AP})^\omega$. An LT-property is thus a set of infinite traces over 2^{AP}.

Intuition

An LT-property gives the admissible behaviours of the DTMC at hand.

Probability of LT properties

The *probability* for DTMC $\mathcal{D}$ to exhibit a trace in P (over AP) is:

$$Pr^{\mathcal{D}}(P) \; = \; Pr^{\mathcal{D}}\{\, \pi \in Paths(\mathcal{D}) \mid trace(\pi) \in P \,\}.$$

For state s in $\mathcal{D}$, let $Pr(s \models P) \; = \; Pr_s\{\, \pi \in Paths(s) \mid trace(\pi) \in P \,\}$.

We will later identify a rich set P of LT-properties—those that include all LTL formulas—for which $\{\, \pi \in Paths(\mathcal{D}) \mid trace(\pi) \in P \,\}$ is measurable.

Safety properties

Safety property

LT property P_{safe} over AP is a *safety property* if for all $\sigma \in (2^{AP})^\omega \setminus P_{safe}$ there exists a finite prefix $\widehat{\sigma}$ of σ such that:

$$P_{safe} \cap \underbrace{\left\{ \sigma' \in (2^{AP})^\omega \mid \widehat{\sigma} \text{ is a prefix of } \sigma' \right\}}_{\text{all possible extensions of } \widehat{\sigma}} \; = \; \varnothing.$$

Any such finite word $\widehat{\sigma}$ is called a *bad prefix* for P_{safe}.

Regular safety property

A safety property is *regular* if its set of bad prefixes constitutes a regular language (over the alphabet 2^{AP}). Thus, the bad prefixes of a regular safety property can be represented by a finite-state automaton.

Probability of a regular safety property

Let $\mathcal{A} = (Q, 2^{AP}, \delta, q_0, F)$ be a deterministic finite-state automaton (DFA) for the bad prefixes of regular safety property P_{safe}:

$$P_{safe} \; = \; \{\, A_0\, A_1\, A_2 \ldots \in (2^{AP})^\omega \mid \forall n \geqslant 0.\, A_0\, A_1 \ldots A_n \notin \mathcal{L}(\mathcal{A}) \,\}.$$

Assume δ to be total, i.e., $\delta(q, A)$ is defined for each $A \subseteq AP$ and each state $q \in Q$. Furthermore, let $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be a finite DTMC. Our interest is to compute the probability

$$Pr^{\mathcal{D}}(P_{safe}) \; = \; 1 - \sum_{s \in S} \iota_{\text{init}}(s) \cdot Pr(s \models \mathcal{A}) \quad \text{where}$$

$$Pr(s \models \mathcal{A}) \; = \; Pr^{\mathcal{D}}_s\{\, \pi \in Paths(s) \mid trace(\pi) \notin P_{safe} \,\}.$$

These probabilities can be obtained by considering a product of DTMC $\mathcal{D}$ with DFA $\mathcal{A}$.

Product Markov chain

Product Markov chain

Let $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be a DTMC and $\mathcal{A} = (Q, 2^{AP}, \delta, q_0, F)$ be a DFA. The *product* $\mathcal{D} \otimes \mathcal{A}$ is the DTMC:

$$\mathcal{D} \otimes \mathcal{A} \; = \; (S \times Q, \mathbf{P}', \iota'_{\text{init}}, \{\, accept \,\}, L')$$

where $L'(\langle s, q \rangle) = \{\, accept \,\}$ if $q \in F$ and $L'(\langle s, q \rangle) = \varnothing$ otherwise, and

$$\iota'_{\text{init}}(\langle s, q \rangle) \; = \; \begin{cases} \iota_{\text{init}}(s) & \text{if } q = \delta(q_0, L(s)) \\ 0 & \text{otherwise.} \end{cases}$$

The transition probabilities in $\mathcal{D} \otimes \mathcal{A}$ are given by:

$$\mathbf{P}'(\langle s, q \rangle, \langle s', q' \rangle) \; = \; \begin{cases} \mathbf{P}(s, s') & \text{if } q' = \delta(q, L(s')) \\ 0 & \text{otherwise.} \end{cases}$$

Product Markov chain

Remarks

- For each path $\pi = s_0\, s_1\, s_2 \ldots$ in DTMC $\mathcal{D}$ there exists a unique run $q_0\, q_1\, q_2 \ldots$ in DFA $\mathcal{A}$ for $trace(\pi) = L(s_0)\, L(s_1)\, L(s_2) \ldots$ and $\pi^+ = \langle s_0, q_1\rangle\, \langle s_1, q_2\rangle\, \langle s_2, q_3\rangle \ldots$ is a path in $\mathcal{D} \otimes \mathcal{A}$.
- The DFA $\mathcal{A}$ does not affect the probabilities, i.e., for each measurable set Π of paths in $\mathcal{D}$ and state s:

$$Pr_s^{\mathcal{D}}(\Pi) = Pr^{\mathcal{D}\otimes\mathcal{A}}_{\langle s, \delta(q_0, L(s))\rangle} \underbrace{\{ \pi^+ \mid \pi \in \Pi \}}_{\Pi^+}$$

- For $\Pi = \{ \pi \in Paths^{\mathcal{D}}(s) \mid trace(\pi) \notin P_{safe} \}$, the set Π^+ is given by:

$$\Pi^+ = \{ \pi^+ \in Paths^{\mathcal{D}\otimes\mathcal{A}}(\langle s, \delta(q_0, L(s))\rangle) \mid \pi^+ \models \Diamond accept \}.$$

Quantitative analysis of regular safety properties

Theorem for analysing regular safety properties

Let P_{safe} be a regular safety property, $\mathcal{A}$ a DFA for the set of bad prefixes of P_{safe}, $\mathcal{D}$ a DTMC, and s a state in $\mathcal{D}$. Then:

$$\begin{aligned} Pr^{\mathcal{D}}(s \models P_{safe}) &= Pr^{\mathcal{D}\otimes\mathcal{A}}(\langle s, q_s\rangle \not\models \Diamond accept) \\ &= 1 - Pr^{\mathcal{D}\otimes\mathcal{A}}(\langle s, q_s\rangle \models \Diamond accept) \end{aligned}$$

where $q_s = \delta(q_0, L(s))$.

Remarks

1. For finite DTMCs, $Pr^{\mathcal{D}}(s \models P_{safe})$ can thus be computed by determining reachability probabilities of *accept* states in $\mathcal{D} \otimes \mathcal{A}$. This amounts to solving a linear equation system.
2. For qualitative regular safety properties, i.e., $Pr^{\mathcal{D}}(s \models P_{safe}) > 0$ and $Pr^{\mathcal{D}}(s \models P_{safe}) = 1$, a graph analysis of $\mathcal{D} \otimes \mathcal{A}$ suffices.

ω-regular languages

Infinite repetition of languages

Let Σ be a finite alphabet. For language $\mathcal{L} \subseteq \Sigma^*$, let $\mathcal{L}^\omega$ be the set of words in $\Sigma^* \cup \Sigma^\omega$ that arise from the infinite concatenation of (arbitrary) words in Σ, i.e.,

$$\mathcal{L}^\omega = \{ w_1 w_2 w_3 \ldots \mid w_i \in \mathcal{L}, i \geqslant 1 \}.$$

The result is an ω-language, i.e., $\mathcal{L} \subseteq \Sigma^*$, provided that $\mathcal{L} \subseteq \Sigma^+$, i.e., $\varepsilon \notin \mathcal{L}$.

ω-regular expression

An *ω-regular expression* G over the Σ has the form: $\mathrm{G} = \mathrm{E}_1.\mathrm{F}_1^\omega + \ldots + \mathrm{E}_n.\mathrm{F}_n^\omega$ where $n \geqslant 1$ and $\mathrm{E}_1, \ldots, \mathrm{E}_n, \mathrm{F}_1, \ldots, \mathrm{F}_n$ are regular expressions over Σ such that $\varepsilon \notin \mathcal{L}(\mathrm{F}_i)$, for all $1 \leqslant i \leqslant n$.

The *semantics* of G is defined by $\mathcal{L}_\omega(\mathrm{G}) = \mathcal{L}(\mathrm{E}_1).\mathcal{L}(\mathrm{F}_1)^\omega \cup \ldots \cup \mathcal{L}(\mathrm{E}_n).\mathcal{L}(\mathrm{F}_n)^\omega$ where $\mathcal{L}(\mathrm{E}) \subseteq \Sigma^*$ denotes the language (of finite words) induced by the regular expression E.

ω-regular expressions

ω-regular expression

An *ω-regular expression* G over the Σ has the form: $\mathrm{G} = \mathrm{E}_1.\mathrm{F}_1^\omega + \ldots + \mathrm{E}_n.\mathrm{F}_n^\omega$ where $n \geqslant 1$ and $\mathrm{E}_1, \ldots, \mathrm{E}_n, \mathrm{F}_1, \ldots, \mathrm{F}_n$ are regular expressions over Σ such that $\varepsilon \notin \mathcal{L}(\mathrm{F}_i)$, for all $1 \leqslant i \leqslant n$.

The semantics of G is defined by $\mathcal{L}_\omega(\mathrm{G}) = \mathcal{L}(\mathrm{E}_1).\mathcal{L}(\mathrm{F}_1)^\omega \cup \ldots \cup \mathcal{L}(\mathrm{E}_n).\mathcal{L}(\mathrm{F}_n)^\omega$ where $\mathcal{L}(\mathrm{E}) \subseteq \Sigma^*$ denotes the language (of finite words) induced by the regular expression E.

Example

Examples for ω-regular expressions over the alphabet $\Sigma = \{ A, B, C \}$ are

$$(A + B)^* A(AAB + C)^\omega \quad \text{or} \quad A(B + C)^* A^\omega + B(A + C)^\omega.$$

ω-regular properties

ω-regular property

LT property P over AP is called *ω-regular* if $P = \mathcal{L}_\omega(G)$ for some ω-regular expression G over the alphabet 2^{AP}.

Example

Let $AP = \{ a, b \}$. Then some ω-regular properties over AP are:

- always a, i.e., $(\{ a \} + \{ a, b \})^\omega$.
- eventuallty a, i.e., $(\varnothing + \{ b \})^*.(\{ a \} + \{ a, b \}).(2^{AP})^\omega$.
- infinitely often a, i.e., $((\varnothing + \{ b \})^*.(\{ a \} + \{ a, b \}))^\omega$.
- from some moment on, always a, i.e., $(2^{AP})^*.(\{ a \} + \{ a, b \})^\omega$.

Deterministic Rabin automata

Deterministic Rabin automaton

A *deterministic Rabin automaton* (DRA) $\mathcal{A} = (Q, \Sigma, \delta, q_0, \mathcal{F})$ with

- Q, $q_0 \in Q_0$, Σ is an alphabet, and $\delta : Q \times \Sigma \to Q$ as before
- $\mathcal{F} = \{ (L_i, K_i) \mid 0 < i \leqslant k \}$ with $L_i, K_i \subseteq Q$, is a set of *accept pairs*

A *run* for $\sigma = A_0 A_1 A_2 \ldots \in \Sigma^\omega$ denotes an infinite sequence $q_0\, q_1\, q_2 \ldots$ of states in $\mathcal{A}$ such that $q_0 \in Q_0$ and $q_i \xrightarrow{A_i} q_{i+1}$ for $i \geqslant 0$.

Run $q_0\, q_1\, q_2 \ldots$ is *accepting* if for some pair (L_i, K_i), the states in L_i are visited finitely often and the states in K_i infinitely often. That is, an accepting run should satisfy

$$\bigvee_{0<i\leqslant k} (\Diamond\Box\neg L_i \wedge \Box\Diamond K_i).$$

Deterministic Rabin automata

DRA and ω-regular languages

The class of languages accepted by DRAs agrees with the class of ω-regular languages.

Thus, the language of any DRA $\mathcal{A}$ is ω-regular. Vice versa, for any ω-regular language $\mathcal{L}$, a DRA $\mathcal{A}$ exists such that $\mathcal{L}_\omega(\mathcal{A}) = \mathcal{L}$.

The proof of this theorem is outside the scope of this lecture.

Verifying DRA properties

Product of a Markov chain and a DRA

The product of DTMC $\mathcal{D}$ and DRA $\mathcal{A}$ is defined as the product of a Markov chain and a DFA, except that the labeling is defined differently.

Let the acceptance condition of $\mathcal{A}$ is $\mathcal{F} = \{ (L_1, K_1), \ldots, (L_k, K_k) \}$. Then the sets L_i, K_i serve as atomic propositions in $\mathcal{D} \otimes \mathcal{A}$. The labeling function L' in $\mathcal{D} \otimes \mathcal{A}$ is the obvious one: if $H \in \{ L_1, \ldots, L_k, K_1, \ldots, K_k \}$, then $H \in L'(\langle s, q \rangle)$ if and only if $q \in H$.

Accepting BSCC

A BSCC T in $\mathcal{D} \otimes \mathcal{A}$ is *accepting* if and only if there exists some index $i \in \{ 1, \ldots, k \}$ such that:

$$T \cap (S \times L_i) = \varnothing \quad \text{and} \quad T \cap (S \times K_i) \neq \varnothing.$$

Thus, once such an accepting BSCC T is reached in $\mathcal{D} \otimes \mathcal{A}$, the acceptance criterion for the DRA $\mathcal{A}$ is fulfilled almost surely.

Verifying DRA objectives

Verifying DRA objectives theorem

Let $\mathcal{D}$ be a finite DTMC, s a state in $\mathcal{D}$, $\mathcal{A}$ a DRA, and let U be the union of all accepting BSCCs in $\mathcal{D} \otimes \mathcal{A}$. Then:

$$Pr^{\mathcal{D}}(s \models \mathcal{A}) \;=\; Pr^{\mathcal{D}\otimes\mathcal{A}}(\langle s, q_s\rangle \models \Diamond U) \quad \text{where} q_s = \delta(q_0, L(s)).$$

Thus: $Pr^{\mathcal{D}}(\mathcal{A}) \;=\; \sum_{s \in S} \iota_{\text{init}}(s) \cdot Pr^{\mathcal{D}\otimes\mathcal{A}}(\langle s, \delta(q_0, L(s))\rangle \models \Diamond U)$. The computation of probabilities for satisfying ω-regular properties boils down to computing the reachability probabilities for certain BSCCs in $\mathcal{D} \otimes \mathcal{A}$. Again, a graph analysis and solving systems of linear equations suffice. The time complexity is polynomial in the size of $\mathcal{D}$ and $\mathcal{A}$.

Measurability

Measurability theorem for ω-regular properties [Vardi 1985]

For any DTMC $\mathcal{D}$ and ω-regular LT property P, the set

$$\{\, \pi \in Paths(\mathcal{D}) \mid trace(\pi) \in P \,\}$$

is measurable.

Proof (sketch)

Represent P by a DRA $\mathcal{A}$ with accept sets $\{\,(L_1, K_1), \ldots, (L_k, K_k)\,\}$. Let $\varphi_i = \Diamond\Box\neg L_i \wedge \Box\Diamond K_i$ and Π_i the set of paths satisfying φ_i. Then $\Pi = \Pi_1 \cup \ldots \cup \Pi_k$. In addition, $\Pi_i \;=\; \Pi_i^{\Diamond\Box} \cap \Pi_i^{\Box\Diamond}$ where $\Pi_i^{\Diamond\Box}$ is the set of paths π in $\mathcal{D}$ such that $\pi^+ \models \Diamond\Box\neg L_i$, and $\Pi_i^{\Box\Diamond}$ is the set of paths π in $\mathcal{D}$ such that $\pi^+ \models \Box\Diamond K_i$. It remains to show that $\Pi_i^{\Diamond\Box}$ and $\Pi_i^{\Box\Diamond}$ are measurable. This goes along the same lines as proving that $\Diamond\Box\, G$ and $\Box\Diamond\, G$ are measurable.

Linear temporal logic

Linear Temporal Logic: Syntax [Pnueli 1977]

LTL formulas over the set AP obey the grammar:

$$\varphi \;::=\; a \;\Big|\; \neg\varphi \;\Big|\; \varphi_1 \wedge \varphi_2 \;\Big|\; \bigcirc\, \varphi \;\Big|\; \varphi_1 \,\mathsf{U}\, \varphi_2$$

where $a \in AP$ and φ, φ_1, and φ_2 are LTL formulas.

LTL semantics

LTL semantics

The LT-property induced by LTL formula φ over AP is:

$$Words(\varphi) \;=\; \left\{\sigma \in \left(2^{AP}\right)^{\omega} \mid \sigma \models \varphi\right\}, \text{where } \models \text{ is the smallest relation s.t.:}$$

$$\begin{array}{lllll}
\sigma & \models & \text{true} & & \\
\sigma & \models & a & \text{iff} & a \in A_0 \quad (\text{i.e., } A_0 \models a) \\
\sigma & \models & \varphi_1 \wedge \varphi_2 & \text{iff} & \sigma \models \varphi_1 \text{ and } \sigma \models \varphi_2 \\
\sigma & \models & \neg\,\varphi & \text{iff} & \sigma \not\models \varphi \\
\sigma & \models & \bigcirc\,\varphi & \text{iff} & \sigma^1 = A_1 A_2 A_3 \ldots \models \varphi \\
\sigma & \models & \varphi_1 \,\mathsf{U}\, \varphi_2 & \text{iff} & \exists j \geqslant 0.\ \sigma^j \models \varphi_2 \;\text{ and }\; \sigma^i \models \varphi_1,\ 0 \leqslant i < j
\end{array}$$

for $\sigma = A_0 A_1 A_2 \ldots$ we have $\sigma^i = A_i A_{i+1} A_{i+2} \ldots$ is the suffix of σ from index i on.

Some facts about LTL

LTL is ω-regular

For any LTL formula φ, the set $Words(\varphi)$ is an ω-regular language.

LTL are DRA-definable

For any LTL formula φ, there exists a DRA $\mathcal{A}$ such that $\mathcal{L}_\omega = Words(\varphi)$ where the number of states in $\mathcal{A}$ lies in $2^{2^{|\varphi|}}$.

Verifying a DTMC against LTL formulas

Complexity of LTL model checking **[Vardi 1985]**

The qualitative model-checking problem for finite DTMCs against LTL formula φ is PSPACE-complete, i.e., verifying whether $Pr(s \models \varphi) > 0$ or $Pr(s \models \varphi) = 1$ is PSPACE-complete.

Recall that the LTL model-checking problem for finite transition systems is also PSPACE-complete.

Summary

Summary

- Verifying a DTMC $\mathcal{D}$ against a DFA $\mathcal{A}$, i.e., determining $Pr(\mathcal{D} \models \mathcal{A})$, amounts to computing reachability probabilities of accept states in $\mathcal{D} \otimes \mathcal{A}$.
- For DBA objectives, the probability of infinitely often visiting an accept state in $\mathcal{D} \otimes \mathcal{A}$.
- DBA are strictly less powerful than ω-regular languages.
- Deterministic Rabin automata are as expressive as ω-regular languages.
- Verifying DTMC $\mathcal{D}$ agains DRA $\mathcal{A}$ amounts to computing reachability probabilities of accepting BSCCs in $\mathcal{D} \otimes \mathcal{A}$.

Take-home message

Model checking a DTMC against various automata models reduces to computing reachability probabilities in a product.