Probabilistic/Stochastic Transition Systems (Markov Chains)

Computational Models for Complex Systems

Paolo Milazzo

Dipartimento di Informatica, Università di Pisa
http://pages.di.unipi.it/milazzo
milazzo@di.unipi.it

Laurea Magistrale in Informatica
A.Y. 2019/2020

Introduction

Transition systems describe all the possibile behaviors of a systems

Alternative behaviors are described through non-determinstic choices

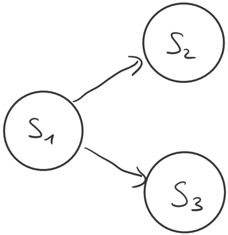

Non-determinism allows choices between alternative behaviors to be modeled without describing the choice criterion

Introduction

Sometimes the choice criterion is known to be

- probabilistic, or
- due to a (stochastic) race between poisson processes (race condition)

This leads to the definition of

- Probabilistic Transition Systems (PTSs)
  aka Discrete Time Markov Chains (DTMCs)
- Stochastic Transition Systems (STSs)
  aka Continuous Time Markov Chains (CTMCs)

See also:

- Dave Parker's Lectures on Probabilistic Model Checking
  (in particular, Lectures 2,3,8,9)
  Available here:
  `https://www.prismmodelchecker.org/lectures/pmc/`

Probability Example

Modeling a 6-sided dice using a fair coin

- algorithm due to Knuth/Yao
- start at 0, toss a coin
- upper branch when H
- lower branch when T
- repeat until value chosen

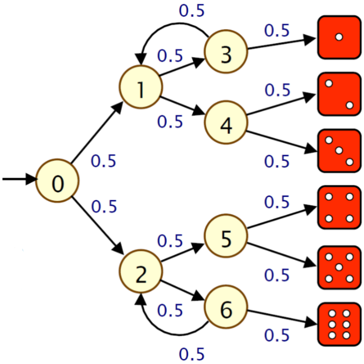

Is this algorithm correct?

- e.g. probability of obtaining a 4?
- Obtain as disjoint union of events
- THH, TTTHH, TTTTTHH, ...
- Probability: $(1/2)^3 + (1/2)^5 + (1/2)^7 + \ldots = 1/6$

Discrete Time Markov Chains (DTMCs)

Let's extend Transition Systems with probabilities...

Definition: Discrete Time Markov Chain (DTMC)

A Discrete Time Markov Chain is a pair (S, P) where

- S is a set of states and
- $P : S \times S \to [0, 1]$ is the probability transition matrix such that, for all $s \in S$ it holds:

$$\sum_{s' \in S} P(s, s') = 1$$

The probability transition matrix can be expressed equivalently as a probabilistic transition relation $\rightarrow \subseteq S \times [0, 1] \times S$ such that $(s, p, s') \in \rightarrow$ (or $s \xrightarrow{p} s'$) if and only if $P(s, s') = p > 0$ (if $p = 0$ the transition is usually omitted).

Discrete Time Markov Chains (DTMCs)

When the set of states is finite, $S = \{s_0, s_1, \ldots, s_n\}$, the probability transition matrix can actually be represented as a square matrix:

$$P = \begin{bmatrix} p_{00} & p_{01} & p_{02} & \cdots & p_{0n} \\ p_{10} & p_{11} & p_{12} & \cdots & p_{1n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ p_{n0} & p_{n1} & p_{n2} & \cdots & p_{nn} \end{bmatrix}$$

where $p_{ij} = P(s_i, s_j)$ and the sum of each row is equal to 1.

A simple DTMC example

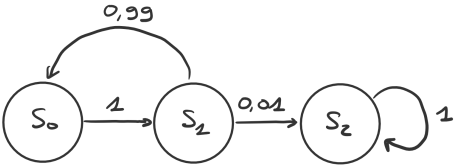

$$S = \{s_0, s_1, s_2\} \qquad P = \begin{bmatrix} 0 & 1 & 0 \\ 0.99 & 0 & 0.01 \\ 0 & 0 & 1 \end{bmatrix}$$

Some notes

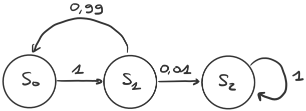

In DTMC we usually have an initial state or (more generally) a probability distribution of initial states, represented as a vector

- $[1, 0, 0]$ means that s_0 is the initial state
- $[0.5, 0.5, 0]$ means that s_0 and s_1 are equally likely to be initial states

The constraint $\sum_{s' \in S} P(s, s') = 1$ implies that

- every state has at least one outgoing transition (otherwise the sum would be 0)
- hence, deadlocks correspond to states with a self-loop

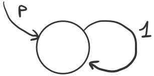

Coins and dice example as a DTMC

Let's reformulate Knuth/Yao's algorithm as a DTMC:

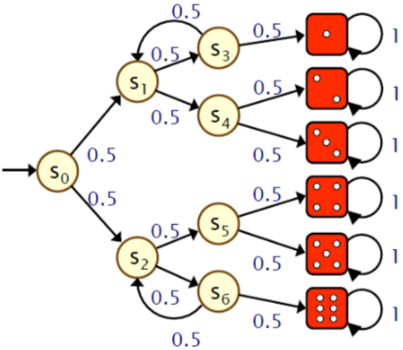

$$S = \{s_0, s_1, \ldots, s_6, 1, 2, \ldots, 6\}$$

$$s_{init} = s_0$$

$$P = \begin{bmatrix} 0 & 0.5 & 0.5 & \ldots \\ 0 & 0 & 0 & \ldots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Paths and their probabilities

A path of a DTMC is the analogous of a (maximal) trace for a Transition System

Definition: Path

A path π of a DTMC (S, P) with initial state s_0, is a (possibly infinite) sequence of states $\pi = s_0, s_1, s_2, \ldots$ such that for each s_{i+1} with $i \in \mathbb{N}$ in π it holds $P(s_i, s_{i+1}) > 0$.

The probability of a path is simply the product of the probabilities of its transitions:

$$Prob(s_0, s_1, s_2, \ldots, s_n) = \prod_{i=0}^{n-1} P(s_i, s_i + 1)$$

$$Prob(s_0, s_1, s_2, \ldots) = \prod_{i \in \mathbb{N}} P(s_i, s_i + 1)$$

Probabilistic reachability

In a DTMC it is possible to compute the probability that the system will reach a given state

- Reachability = property expressing whether a given state can be reached (there exists a path leading to it)
- Probabilistic reachability = probability of reaching a given state (probabilities of all the paths leading to it)

Paths are independent events: their probabilities can be summed!

Definition: Probabilistic Reachability

The probability of reaching state s of a DTMC $(S, \rightarrow)$ from the initial state s_0, is the sum of the probabilities of all paths leading to it.

$$ProbReach(s_0, s) = \sum_{\pi \in Reach(s_0,s)} Prob(\pi)$$

where $Reach(s_0, s)$ is the (possibly infinite) set of paths reaching s.

Probabilistic reachability: example

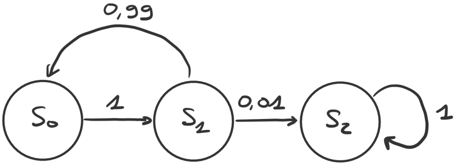

$$
\begin{aligned}
ProbReach(s_0, s_2) &= 1 \cdot 0.01 \\
&+ 1 \cdot 0.99 \cdot 1 \cdot 0.01 \\
&+ (1 \cdot 0.99)^2 \cdot 1 \cdot 0.01 \\
&\quad \vdots \\
&+ (1 \cdot 0.99)^n \cdot 1 \cdot 0.01 \\
&\quad \vdots \\
&= 1
\end{aligned}
$$

Probabilistic reachability: example

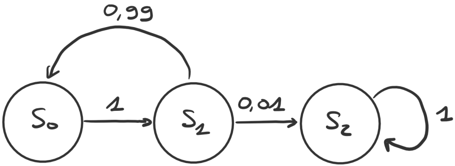

In this example, the infinite sum can be avoided by observing that the only path not leading to s_2 is the infinite path $\pi_{01} = s_0, s_1, s_0, s_1, s_0, \ldots$

So, $ProbReach(s_0, s_2) = 1 - Prob(\pi_{01})$

But π_{01} is a single infinite path with a loop containing a transition with a probability strictly smaller than 1

- $P(\pi_{01}) = (0.99 \cdot 1)^{\infty} = 0$
- $ProbReach(s_0, s_2) = 1 - Prob(\pi_{01}) = 1$

Probabilistic reachability: example

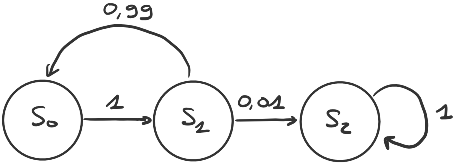

Another (more general) way to avoid the infinite summation, is by reformulating *ProbReach* in terms of a linear system of equations

The idea:

- the probability of reaching s_2 from s_2 is 1
- the probability of reaching s_2 from s_1 is 0.01 plus the probability of reaching s_0 in one step, and then of reaching s_2 from there
- the probability of reaching s_2 from s_0 is the probability of reaching s_1 in one step, and then of reaching s_2 from there

Probabilistic reachability: example

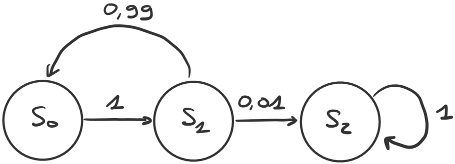

Another (more general) way to avoid the infinite summation, is by reformulating *ProbReach* in terms of a linear system of equations

This leads to a mutually recursive reformulation of *ProbReach*:

- $ProbReach(s_2, s_2) = 1$
- $ProbReach(s_1, s_2) = 0.01 \cdot ProbReach(s_2, s_2) + 0.99 \cdot ProbReach(s_0, s_2)$
- $ProbReach(s_0, s_2) = 1 \cdot ProbReach(s_1, s_2)$

Probabilistic reachability: example

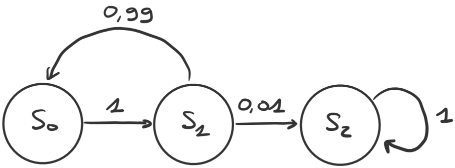

Another (more general) way to avoid the infinite summation, is by reformulating *ProbReach* in terms of a linear system of equations

Let's denote $ProbReach(s, s_2)$ as x_s to obtain:

$$\begin{cases} x_{s_2} = 1 \\ x_{s_1} = 0.01 x_{s_2} + 0.99 x_{s_0} \\ x_{s_0} = x_{s_1} \end{cases}$$

From which we obtain easily $x_{s_0} = 1$

Probabilistic reachability: coins and dice example

Let's compute the probability of rolling a 6

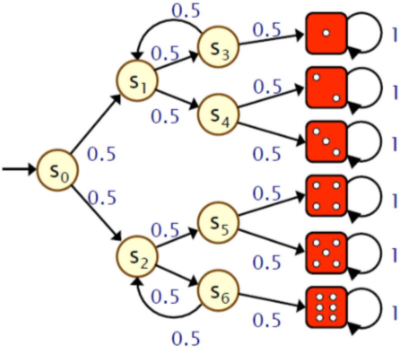

$$
\begin{cases}
x_6 = 1 \\
x_{s_6} = \frac{1}{2} x_{s_2} + \frac{1}{2} x_6 \\
x_{s_2} = \frac{1}{2} x_{s_6} + \frac{1}{2} x_{s_5} \\
x_{s_5} = 0 \\
x_{s_0} = \frac{1}{2} x_{s_2} + \frac{1}{2} x_{s_1} \\
x_{s_1} = 0
\end{cases}
$$

Probabilistic reachability: coins and dice example

Let's compute the probability of rolling a 6

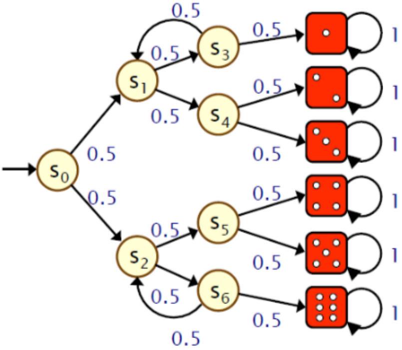

$$\begin{cases} x_{s_6} = \frac{1}{2} x_{s_2} + \frac{1}{2} \\ x_{s_2} = \frac{1}{2} x_{s_6} \\ x_{s_0} = \frac{1}{2} x_{s_2} \end{cases}$$

Probabilistic reachability: coins and dice example

Let's compute the probability of rolling a 6

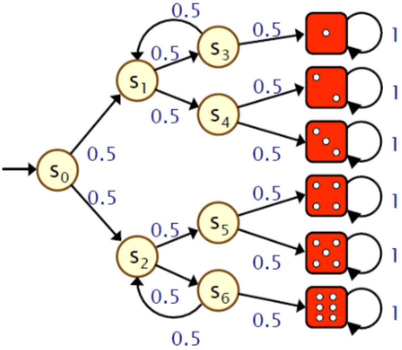

$$\begin{cases} x_{s_6} = \frac{2}{3} \\ x_{s_2} = \frac{1}{3} \\ x_{s_0} = \frac{1}{6} \end{cases}$$

Computing probabilistic reachability

We have seen that computing probabilistic reachability $ProbReach(s, s')$ amounts to solving a system of linear equations in order to obatain $x_{s'}$:

$$\forall S_i \in S \qquad X_{s_i} = \begin{cases} 1 & \text{if } s_i = s' \\ \sum_{s_j \in S} P(s_j, s_i) X_{s_i} & \text{otherwise} \end{cases}$$

where P is the probability transition matrix of the DTMC

This can be done by applying iterative computational algebra methods

Continuous Time Markov Chains (CTMCs)

This time let's extend Transition Systems with stochastic rates...

Definition: Continuous Time Markov Chain (CTMC)

A Continuous Time Markov Chain is a pair (S, R) where

- S is a set of states and
- $R : S \times S \to \mathbb{R}^{\geq 0}$ is the transition rate matrix

The transition rate matrix can be expressed equivalently as a stochastic transition relation $\to \subseteq S \times \mathbb{R}^{\geq 0} \times S$ such that $(s, r, s') \in \to$ (or $s \xrightarrow{r} s'$) if and only if $R(s, s') = r > 0$ (if $r = 0$ the transition is usually omitted).

Race conditions

What happens when there exist multiple s' with $R(s, s') > 0$?

- race condition: the "fastest" transition determines the next state of the system
- Two questions:
  - 1) How long is spent in s before a transition occurs?
  - 2) Which transition is eventually taken?

1) Time spent in a state before a transition

- minimum of exponential distributions
- exponential with parameter given by the summation:

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

- E(s) is called exit rate of state s

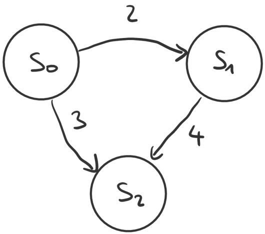

Race conditions

2) Which transition is taken from state s?

- the choice is independent from the time at which it occurs
- the probability is proportional to the rate of each transition

More generally, the probability of the next transition to occur is given by the embedded DTMC of the CTMC...

Embedded DTMC of a CTMC

The embedded DTMC of a CTMC describes the state changes of the CTMC by ignoring time

It is obtained by normalizing the transition rates of the CTMC with respect to the exit rate of each state

Definition: Embedded DTMC

Given a CTMC (S, R), its embedded DTMC is the DTMC (S, P) where, for any $s, s' \in S$

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

Embedded DTMC of a CTMC

An example:

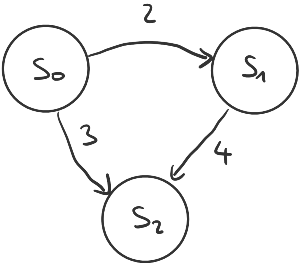

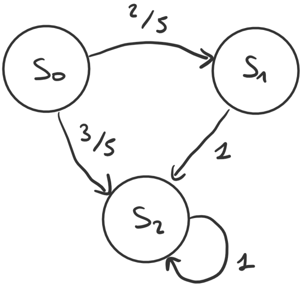

Time Independent Probabilistic Reachability on CTMCs

Given a CTMC, what is the probability of reaching a state s at any time?

- It corresponds to probabilistic reachability of the same state in the embedded DTMC
- So, it can be computed by applying computational algebra methods based on the transition probability matrix P of the embedded DTMC

Transient (Time-dependent) Probabilistic Reachability on CTMCs

Given a CTMC, what is the probability of the system to be in a state s at a given time?

This can be answered by introducing uniformised DTMCs

Uniformisation

Given a CTMC, its uniformised DTMC is obtained by choosing a uniformisation rate q bigger or equal to all the rates of the CTMC

- each rate r of the CTMC is transformed into probability r/q
- self-loops are added where necessary

Example (with $q = 10$):

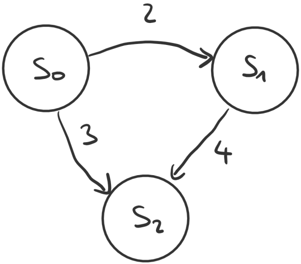

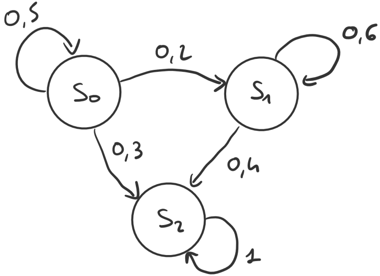

Uniformisation

Some notes:

- a transition in the uniformised DTMC describes a step with duration $1/q$
- q should be chosen big enough to assume that at most one transition can occur during a $1/q$ time interval

Transient probabilistic reachability of a CTMC can now be computed as probabilistic reachability in the uniformised DTMC, by taking the length of the paths in the DTMC into account.

- again, can be computed by performing matrix computations on the transition probability matrix of the DTMC

Application to chemical reactions

Now, we could apply reachability analysis to chemical reactions...

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