IA169 System Verification and Assurance

Verification of Systems with Probabilities

Vojtěch Řehák

Motivation example

Fail-repair system

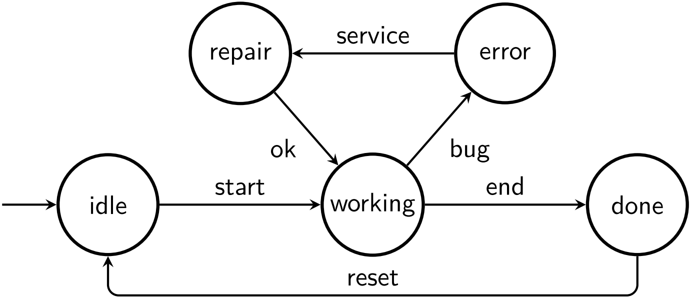

What are the properties of the model?

- $G(\text{working} \implies F \text{ done})$ **NO**
- $G(\text{working} \implies F \text{ error})$ **NO**
- $FG(\text{working} \vee \text{error} \vee \text{repair})$ **NO**

Motivation example

Fail-repair system

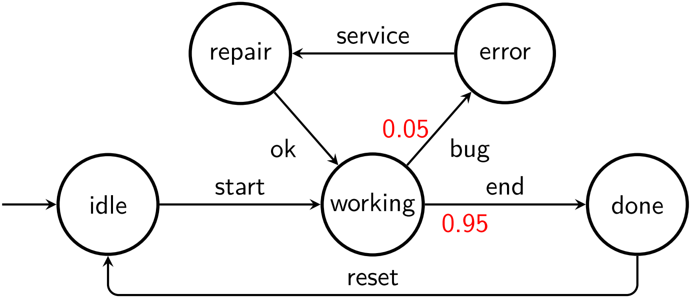

- What is the probability of reaching "done" from "working" with no visit of "error"?
- What is the probability of reaching "done" from "working" with at most one visit of "error"?
- What is the probability of reaching "done" from "working"?

Discrete-time Markov Chains (DTMC)

Probabilistic models

Discrete-time Markov Chains (DTMC)

- Standard model for probabilistic systems.
- State-based model with probabilities on branching.
- Based on the current state, the succeeding state is given by a discrete probability distribution.
- Markov property ("memorylessness") — only the current state determines the successors (the past states are irrelevant).
- Probabilities on outgoing edges sums to 1 for each state.
- Hence, each state has at least one outgoing edge ("no deadlock").

Model of a queue

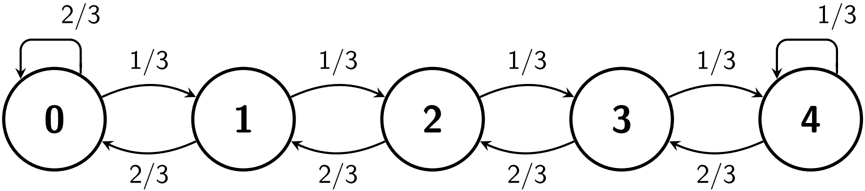

Queue for at most 4 items. In every time tick, a new item comes with probability $1/3$ and an item is consumed with probability $2/3$.

What if a new items comes with probability $p = 1/2$ and an item is consumed with probability $q = 2/3$?

Model of the new queue

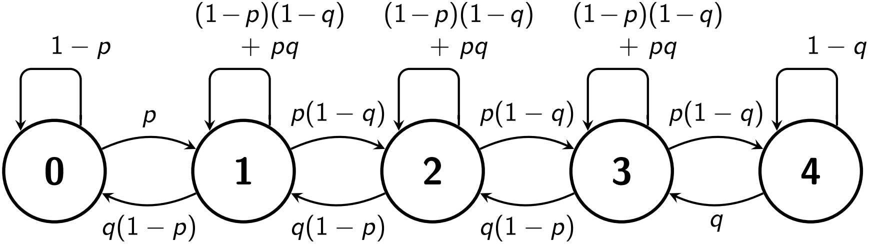

DTMC - formal definition

Discrete-time Markov Chain is given by

- a set of states S,
- an initial state s_0 of S,
- a probability matrix $P : S \times S \to [0, 1]$, and
- an interpretation of atomic propositions $I : S \to AP$.

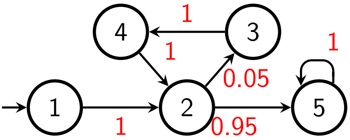

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0.05 & 0 & 0.95 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Back to our questions

Fail-Repair System

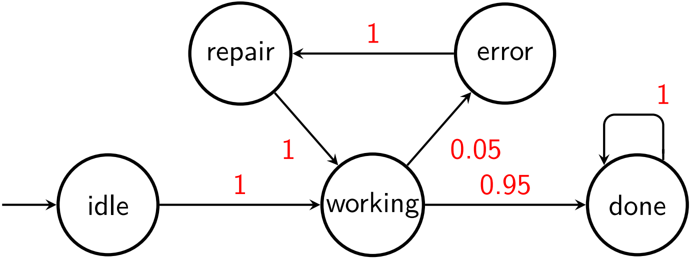

- What is the probability of reaching “done” from “working” with no visit of “error”?
- What is the probability of reaching “done” from “working” with at most one visit of “error”?
- What is the probability of reaching “done” from “working”?

Markov chain analysis

Transient analysis

- distribution after k-steps
- reaching/hitting probability
- hitting time

Long run analysis

- probability of infinite hitting
- stationary (invariant) distribution
- mean inter visit time
- long run limit distribution

Property Specification

Recall some non-probabilistic specification languages:

LTL formulae

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid X\,\varphi \mid \varphi\,U\,\varphi$$

CTL formulae

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid EX\,\varphi \mid E[\varphi\,U\,\varphi] \mid EG\,\varphi$$

Syntax of CTL*

state formula	$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid E\,\psi$
path formula	$\psi ::= \varphi \mid \neg\psi \mid \psi \vee \psi \mid X\,\psi \mid \psi\,U\,\psi$

Property specification languages

We need to quantify probability that a certain behaviour will occur.

Probabilistic Computation Tree Logic (PCTL)

Syntax of PCTL

state formula	$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid P_{\bowtie b}\psi$
path formula	$\psi ::= X\,\varphi \mid \varphi\, U\, \varphi \mid \varphi\, U^{\leq k}\, \varphi$

where

- $b \in [0,1]$ is a probability bound,
- $\bowtie \in \{\leq, <, \geq, >\}$, and
- $k \in \mathbf{N}$ is a bound on the number of steps.

A PCTL formula is always a state formula.

$\alpha\, U^{\leq k}\, \beta$ is a bounded until saying that α holds until β within k steps.
For $k = 3$ it is equivalent to $\beta \vee (\alpha \wedge X\,\beta) \vee (\alpha \wedge X\,(\beta \vee \alpha \wedge X\,\beta))$.

Some tools also supports $P_{=?}\psi$ asking for the probability that the specified behaviour will occur.

PCTL examples

We can also use derived operators like G, F, $\wedge$, $\Rightarrow$, etc.

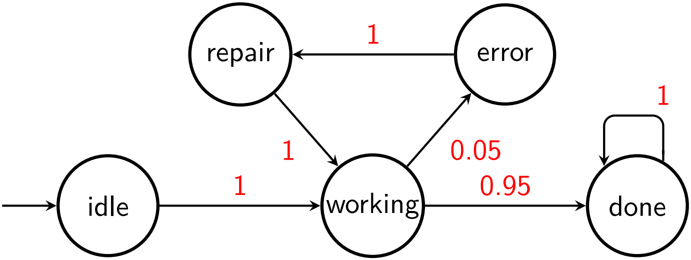

Probabilistic reachability $P_{\geq 1}(\, F\, done\,)$

- probability of reaching the state *done* is equal to 1

Probabilistic bounded reachability $P_{>0.99}(\, F^{\leq 6}\, done\,)$

- probability of reaching the state *done* in at most 6 steps is > 0.99

Probabilistic until $P_{<0.96}(\,(\neg error)\, U\,(done)\,)$

- probability of reaching *done* with no visit of *error* is less than 0.96

Qualitative vs. quantitative properties

Qualitative PCTL properties

- $P_{\bowtie b}\,\psi$ where b is either 0 or 1

Quantitative PCTL properties

- $P_{\bowtie b}\,\psi$ where b is in $(0, 1)$

In DTMC where zero probability edges are erased, it holds that

- $P_{>0}(\,X\,\varphi)$ is equivalent to $EX\,\varphi$
  - there is a next state satisfying φ
- $P_{\geq 1}(\,X\,\varphi)$ is equivalent to $AX\,\varphi$
  - the next states satisfy φ
- $P_{>0}(\,F\,\varphi)$ is equivalent to $EF\,\varphi$
  - there exists a finite path to a state satisfying φ

but

- $P_{\geq 1}(\,F\,\varphi)$ is **not** equivalent to $AF\,\varphi$

There is no CTL formula equivalent to $P_{\geq 1}(\,F\,\varphi)$,
and no PCTL formula equivalent to $AF\,\varphi$.

How the transient probabilities are computed?

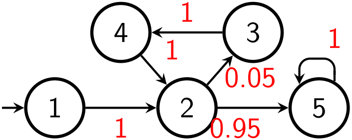

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0.05 & 0 & 0.95 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Probability in the k-th state when starting in 1

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \end{bmatrix} \times P = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \end{bmatrix} \times P^2 = \begin{bmatrix} 0 & 0 & 0.05 & 0 & 0.95 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \end{bmatrix} \times P^3 = \begin{bmatrix} 0 & 0 & 0 & 0.05 & 0.95 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \end{bmatrix} \times P^4 = \begin{bmatrix} 0 & 0.05 & 0 & 0 & 0.95 \end{bmatrix}$$

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \end{bmatrix} \times P^5 = \begin{bmatrix} 0 & 0 & 0.0025 & 0 & 0.9975 \end{bmatrix}$$

How the transient probabilities are computed?

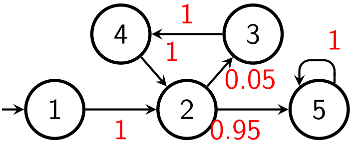

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0.05 & 0 & 0.95 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Probability of being in 5 or 2 in the k-th state

$$P \times \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \end{bmatrix}^T = \begin{bmatrix} 1 & 0.95 & 0 & 1 & 1 \end{bmatrix}^T$$

$$P^2 \times \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \end{bmatrix}^T = \begin{bmatrix} 0.95 & 0.95 & 1 & 0.95 & 1 \end{bmatrix}^T$$

$$P^3 \times \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \end{bmatrix}^T = \begin{bmatrix} 0.95 & 1 & 0.95 & 0.95 & 1 \end{bmatrix}^T$$

$$P^4 \times \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \end{bmatrix}^T = \begin{bmatrix} 1 & 0.9975 & 0.95 & 1 & 1 \end{bmatrix}^T$$

$$P^5 \times \begin{bmatrix} 0 & 1 & 0 & 0 & 1 \end{bmatrix}^T = \begin{bmatrix} 0.9975 & 0.9975 & 1 & 0.9975 & 1 \end{bmatrix}^T$$

Unbounded reachability - optional slide

Unbounded reachability

Let $p(s, A)$ be the probability of reaching a state in A from s.

It holds that:

- $p(s, A) = 1$ for $s \in A$
- $p(s, A) = \sum_{s' \in succ(s)} P(s, s') * p(s', A)$ for $s \notin A$

where $succ(s)$ is a set of successors of s and $P(s, s')$ is the probability on the edge from s to s'.

Theorem

- The minimal non-negative solution of the above equations equals to the probability of unbounded reachability.

Long Run Analysis

Long run analysis

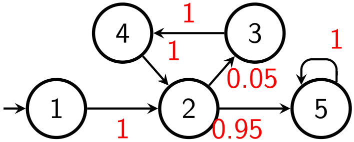

Recall that we reach the state 5(*done*) with probability 1.

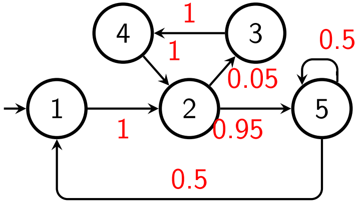

What are the states visited infinitely often with probability 1?

Transient and recurrent states

Definitions

- A state of DMTC is called **transient** iff there is a positive probability that the system will not return back to this state.
- A state s of DMTC is called **recurrent** iff, starting from s, the system eventually returns back to the s with probability 1.

Theorem

- Every transient state is visited finitely many times with probability 1.
- Each recurrent state is with probability 1 either **not visited** or **visited infinitely many times**.[1]

[1] This holds only in DTMC models with finitely many states.

Transient vs. recurrent states

Which states are transient? Which states are recurrent?

Decompose the graph representation onto strongly connected components.

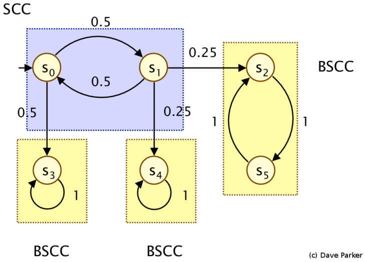

Theorem [1]

- A state is **recurrent** if and only if it is in a **bottom strongly connected component**. All other states are **transient**.

[1] This holds only in DTMC models with finitely many states.

Irreducible Markov Chain

For the sake of infinite behaviour, we will concentrate on bottom strongly connected components only.

Definition

- A Markov chain is said to be **irreducible** if every state can be reached from every other state in a finite number of steps.

Theorem

- A Markov chain is **irreducible** if and only if its graph representation is a single strongly connected component.

Corollary

- All states of a finite irreducible Markov chain are recurrent.

Stationary (Invariant) Distribution

Definition

- Let P be the transition matrix of a DTMC and $\vec{\lambda}$ be a probability distribution on its states. If

$$\vec{\lambda}P = \vec{\lambda},$$

then $\vec{\lambda}$ is a **stationary** (or **steady-state** or **invariant** or **equilibrium**) **distribution** of the DTMC.

Question:
How many stationary distributions can a Markov chain have?
Can it be more than one?
Can it be none?

Answer: It can be more that one. For example, in the Drunkard's walk

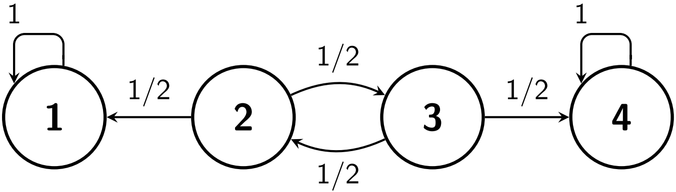

both $(1, 0, 0, 0)$ and $(0, 0, 0, 1)$ are stationary distributions.

But, this is not an irreducible Markov chain.

Stationary Distributions

Theorem

- In every finite irreducible DTMC there is a unique invariant distribution.

Q: Can it be none?

Theorem

- For each finite DTMC, there is an invariant distribution.

Q: How can we compute the invariant distribution of a finite irreducible Markov chain?

Stationary Distribution & Cut-sets

Again, we can construct a set of equations that express the result.

Theorem

- Let P be a transition matrix of a finite irreducible DTMC and $\vec{\pi} = (\pi_1, \pi_2, \ldots, \pi_n)$ be a stationary distribution corresponding to P. For any state i of the DTMC, we have

$$\sum_{j \neq i} \pi_j P_{j,i} = \sum_{j \neq i} \pi_i P_{i,j}.$$

Mean Portion of Visited States and Inter Visit Time

Theorem

- Let us have a finite irreducible DTMC and the unique stationary distribution $\vec{\pi}$. It holds that

  $\pi_i = lim_{n\to\infty} E($ # of visits of state i during the first n steps$)/n$.

- Let us have a finite irreducible DTMC and the unique stationary distribution $\vec{\pi}$. It holds that

  $$\pi_i = 1/m_i$$

  where m_i is the mean inter visit time of state i.

For example:

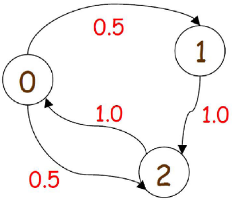

aperiodic

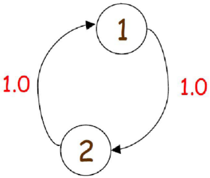

periodic

Definition

- A state s is **periodic** if there exists an integer $\Delta > 1$ such that length of every path from s to s is divisible by Δ.
- A Markov chain is **periodic** if any state in the chain is periodic.
- A state or chain that is not periodic is **aperiodic**.

Theorem

- Let us have a finite aperiodic irreducible DTMC and the unique stationary distribution $\vec{\pi}$. It holds that

$$\vec{\pi} = lim_{n\to\infty}\vec{\lambda}P^n$$

where $\vec{\lambda}$ is an arbitrary distribution on states.

Q: What this is good for?

DTMC Extensions - Communication and Nondeterminism

Last remark on some DTMC extensions.

Modules and synchronisation

- modules
- actions
- rewards

Decision extension

- Markov Decision Processes (MDP)
- **Pmin** and **Pmax** properties