Formal Methods in Software Development

Probabilistic Model Checking

Ivano Salvo

Computer Science Department

Sapienza Università di Roma

Lesson **11**, December 14th, 2020

Probabilistic Systems

Real systems are often dependent on phenomena of a stochastic nature. Here, we address **verification of probabilistic systems**.

By contrast, **probabilistic verification** means no complete coverage ("there is no error with a probability of 90%").

* **Randomized algorithms**: several algorithms (distributed) such leader election use tossing coins to break symmetries.

* Modelling **unreliable or unpredictable** behaviours (ex: message loss, system failures): modelling that with nondeterminism can be too coarse. In late stage of model design, probabilistic valuation can take place of nondet.

* **Performance evaluation**: distribution of inputs, messages, etc. are importat to evaluate quantitative aspects such as waiting time, queue length, expected time between failures.

Verifying Probabilistic Systems

We will see:

- **Markov chains** as generalisation of Kripke structures: in this view we will have a "state based" approach to Markov chains;
- A logic for defining probabilistic properties (here probabilities are in the syntax): **PCTL**.

Quantitative properties: "The probability for delivering a message in the next t time units is 98%"

Qualitative properties: A desired event happens almost surely (i.e. with probability 1) or a bad event occurs almost never (i. e. with probability 0): reachability, persistency, repeated reachability.

Lesson 11a:

Markov Chains

Markov Chains: definition

Definition: A (discrete time) **Markov chain** is a tuple $\mathcal{M} = (S, \mathcal{P}, \iota, AP, L)$ where:

S, AP, L as usual are states, atomic propositions and labelling

$\mathcal{P} : S \times S \rightarrow [0,1]$ is the **transition probability function**, such that for all $s \in S$, $\sum_{t \in S} \mathcal{P}(s, t) = 1$

$\iota : S \rightarrow [0,1]$ is the **initial distribution**, such that $\sum_{s \in S} \iota(s) = 1$

$\mathcal{M}$ is finite if S and AP are finite, and the size of $\mathcal{M}$ is:

$$|\mathcal{M}| = |S| + |\{ (s, t) \in S : \mathcal{P}(s, t) > 0\}|$$

(it is the size of the **underlying digraph**)

We will identify $\mathcal{P}$ with the matrix of probability $[\mathcal{P}(s, t)]_{s,t \in S}$ where the row $\mathcal{P}(s, \cdot)$ contains probability to reach successors of s, and the column $\mathcal{P}(\cdot, s)$ contains probability to enter state s from its predecessors.

States such that $\iota(s)>0$ are **initial states** and it is the probability that system evolution starts in s.

Markov Chains: Example

Let us consider an error prone **communication protocol**, that with probability 10% can loose a message. The message is sent until it is eventually delivered.

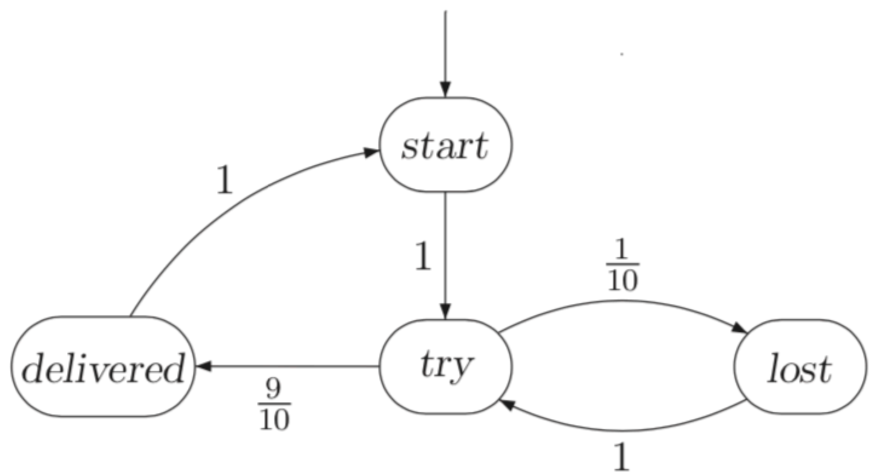

Probability matrix and initial states (start, try, lost, delivered):

$$\mathbf{P} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{10} & \frac{9}{10} \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \qquad \iota_{\text{init}} = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

Markov Chains: Example

Observe that in the underline Kripke structure (**without probabilities**) we can check LTL or CTL properties, like:

$\mathbf{G}\,\mathbf{X}^{100}$ delivered and $\mathbf{EG}\neg$delivered

Both these two properties **does not hold**, even though with very low probability.

In particular, the second **has probability 0!**

Probabilistic model checking allow **quantitative properties** to be checked.

Qualitative properties are a special case, when we ask for an event to have probability 0 or 1.

Markov Chains: terminology

$Paths(\mathcal{M})$ denotes the set of paths, $Paths_{\mathrm{fin}}(\mathcal{M})$ finite paths. When $\mathcal{M}$ is clear from the context, and s is a state, we can use $Paths(s)$ and $Paths_{\mathrm{fin}}(s)$ to denote paths starting at s.

Direct **successors** of a state s are denoted by $Post(s)$. $Post^*(s)$ is the set of states reachable from s.

Similarly, direct **predecessors** of s are denoted by $Pre(s)$. $Pre^*(s)$ is the set of states backward reachable from s.

These notions are naturally extended to sets.

A state s of a Markov Chain $\mathcal{M}$ is said **absorbing** if $Post^*(s)=\{s\}$, that is $\mathcal{P}(s, s)=1$ and $\mathcal{P}(s, t)=0$ when $s \neq t$.

A taste of probability: σ-algebras

Definition: A ***σ*-algebra** is a pair $(O, \mathcal{E})$ where O is a nonempty set (**outcomes**) and $\mathcal{E} \subseteq \mathcal{P}(O)$ is the set of **events** and it contains the **empty set** and it is **closed under complementation** and **countable unions**. More formally:

- $\varnothing \in \mathcal{E}$,
- If $E \in \mathcal{E}$ then $O \setminus E \in \mathcal{E}$,
- If $E_1, E_2 \ldots \in \mathcal{E}$ then $\cup_{i \geq 1} E_i \in \mathcal{E}$.

Observations: $O \in \mathcal{E}$ as the complement of $\varnothing$. $\mathcal{E}$ is closed under countable intersections, since $\cap_{i\geq1} E_i = O \setminus \cup_{i\geq1} (O \setminus E_i)$.

$\mathcal{P}(O)$ is always a σ-algebra and also $\mathcal{E}=\{\varnothing, O\}$.

Definition: A **probability measure** on $(O, \mathcal{E})$ is a function $Pr : \mathcal{E} \rightarrow [0,1]$ such that $Pr(O)=1$, and for a family of pairwise disjoint sets: $Pr(\cup_{i\geq1} E_i) = \sum_{i\geq1} Pr(E_i)$.

A **probability space** is the triple $(O, \mathcal{E}, Pr)$.

Probability spaces: properties

When O is countable, fixing a function $\mu : O \rightarrow [0,1]$, such that $\sum_{e \in O} \mu(e) = 1$ defines a probability measure on $(O, \mathcal{P}(O))$, defined by $Pr(E)= \sum_{e \in E} \mu(e)$.

Since $E \cup (O \setminus E)=O$ and $E \cap (O \setminus E)= \emptyset$, we have $Pr(O \setminus E)=1- Pr(E)$. In particular, $Pr(\emptyset)=1 - Pr(O)=0$.

Probability measures are **monotonic**: if $E \subseteq F$, then

$$Pr(F) = Pr(E) + Pr(F \setminus E) \geq Pr(E).$$

For each set $P \subseteq \mathcal{P}(O)$, there exists a **smallest σ-algebra $\mathcal{E}_P$** that contains P. $\mathcal{E}_P$ is called the σ-algebra **generated by P**, and P is the **basis**.

Example: Let us consider the experiment of **tossing a fair coin once**. The set O of outcomes is {head, tail}. The singletons {head}, {tail} can be the set of events. The smallest σ-algebra containg such events is $\mathcal{P}$({head, tail}) with:

$$Pr(\emptyset)=0,\ Pr(\{\text{head}\})=Pr(\{\text{tail}\})=1/2, \text{ and } Pr(\{\text{head,tail}\})=1.$$

σ-algebras and Markov chains

Definition: The **cylinder set** of a finite path $\pi = s_0 s_1 \dots s_n$ is $Cyl(\pi) = \{\pi' \mid \pi' = \pi\pi''\}$.

The **σ-algebra $\mathcal{E}_{\mathcal{M}}$ associated with a Markov chain $\mathcal{M}$** is generated by all $Cyl(\pi)$, for any π finite path in $\mathcal{M}$.

$$Pr(Cyl(s_0 s_1 \dots s_n)) = \iota(s_0) \prod_{0 \leq i < n} \mathcal{P}(s_i, s_{i+1})$$

Notation: We will use **LTL-like syntax** to denote events in the probability space $(\mathrm{Path}_{\mathcal{M}}, \mathcal{E}_{\mathcal{M}}, Pr)$.

For example, if $B \subseteq S$, "**F** B" is the set of paths that reach the set B after a finite number of steps.

"**GF** B" is the event of visiting B infinitely often.

Sometimes we will write $\pi \vDash \varphi$ for $\pi \in \varphi$ and we denote with $Pr(s \vDash \varphi)$ the probability of φ to hold in the state s, that is $Pr(\{\pi \in \mathrm{Path}(s) \mid \pi \vDash \varphi\}$.

Reachability problems

As for classical Model Checking, one of the basic problems is **reachability**: here, the problem is to **compute the probability of reaching** a given set of states $B \subseteq S$.

$\text{Path}(\mathbf{F}\ B) = \text{Path}_{fin}(\mathcal{M}) \cap (S \setminus B)^* B$ is the set of path that reach B.

$$Pr(\mathbf{F}\ B) = \sum_{\pi \in \text{Path}(\mathbf{F}\ B)} Cyl(\pi)$$

Example [COMMUNICATION PROTOCOL]: The probability of reaching the state delivered depends on the cylinder of:

$$\pi = \text{start try (lost try)}^n \text{ delivered}$$

from which we derive:

$$Pr(\mathbf{F} \text{ delivered}) = \sum_{n \geq 0} (1/10)^n \, 9/10 = 1$$

Intuition: any message will be eventually delivered. If we put a bound on retransmissions, say 3, we have:

$$Pr(\mathbf{F} \text{ delivered}) = 9/10 + 1/10 * 9/10 + 1/100 * 9/10 = 0.999$$

Computing probabilities

Lex $x_s = Pr(s \vDash \mathbf{F}\, B)$. For s ∈ B, $x_s = 1$. For $s \in S \setminus B$, we have:

reach B in one step

$$x_s = \textstyle\sum_{t \in S \setminus B} P(s, t) \cdot x_t + \sum_{u \in B} P(s, u) \qquad (\ast)$$

This is a sort of "**probabilistic expansion law**". By considering only states in $S'=Pre^*(B)\setminus B$, $(\ast)$ $x = (x_s)_{s \in S'}$ becomes: $x = \mathbf{A}\, x + \mathbf{b}$, where $\mathbf{A}$ is $(\mathcal{P}(s, t))_{s,\, t \in S'}$ and $\mathbf{b}$ is the probability of reaching S' in one step which can be rewritten as $(\mathbf{I} - \mathbf{A})\, x = \mathbf{b}$, where $\mathbf{I}$ is the identity matrix of size $|S'| \times |S'|$.

Example [COMMUNICATION PROTOCOL]: let B = {delivered} and S'={start, try, lost}. We can easily obtain the following equations:

$$x_{\text{start}} = x_{\text{try}} \qquad x_{\text{try}} = 1/10\; x_{\text{lost}} + 9/10 \qquad x_{\text{lost}} = x_{\text{try}}$$

that correspond to the system (**the solution is 1 for all states**):

$$\begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & -\frac{1}{10} \\ 0 & -1 & 1 \end{pmatrix} x = \begin{pmatrix} 0 \\ \frac{9}{10} \\ 0 \end{pmatrix}$$

Algorithm

First compute the set S'. This can be done simply by a backward visit starting from B.

Then generate the matrix **A** and the vector **b** and solve the linear system $(\mathbf{I} - \mathbf{A})x = \mathbf{b}$.

Problem: This system **can have more than one solutions** when **I – A** is **singular**. We are interested is the **least solution** in [0,1].

Solution: apply an **iterative method** (instead of direct methods) for a more general problem **constrained reachability** (property of the form C **U** B).

Iterative constrained reachability

Let B, $C \subseteq S$. We consider the problem of reaching B via a finite path fragment in C, that is C **U** B. For $n \geq 0$, the event $C\, \mathbf{U}^{\leq n}\, B$ is as C **U** B, but it is required that B is reached in **at most n steps**.

We partition S as follows:

- $S \setminus (C\, \mathbf{U}\, B) \subseteq \mathbf{S_0} \subseteq \{ s \in S \mid Pr(s \vDash C\, \mathbf{U}\, B)=0 \}$
- $B \subseteq \mathbf{S_1} \subseteq \{ s \in S \mid Pr(s \vDash C\, \mathbf{U}\, B)=1\}$
- $\mathbf{S_?} = S \setminus (S_0 \cup S_1)$

Theorem: Let $(x_s)_{s \in S?}$ be the least fixed point of the operator $\mathbf{\Upsilon} : [0,1]^n \rightarrow [0,1]^n$ defined by: $\mathbf{\Upsilon}(y) = \mathbf{A}y + \mathbf{b}$, where n is the cardinality of $S_?$, $\mathbf{A}$ is the probability transition restricted on states in $S_?$, and $\mathbf{b}$ is the vector of probability of enter B in one step. Then, if x^0 is $\mathbf{0}$ and $x^{n+1} = \mathbf{\Upsilon}(x^n)$, we have:

- $x_s^{\,n} = Pr(s \vDash C\, \mathbf{U}^{\leq n}\, S_1)$,
- $x_s^{\,0} \leq x_s^{\,1} \leq x_s^{\,2} \leq \ldots$,
- $\boldsymbol{x} = lim_{n \rightarrow \infty}\, x^n$

Iterative Algorithm

The previous theorem suggests an iterative algorithm to compute x_s. $x^0 = \mathbf{0}$ and $x^{n+1} = \boldsymbol{\Upsilon}(x^n)$. Since this sequence converges, we can stop when $|x^{n+1} - x^n| < \varepsilon$, for some small tolerance ε.

Remark: Sets S_0 and S_1 are **not uniquely** identified. For example, $S_0 = S \setminus (C\ \mathbf{U}\ B)$ and $S_1 = B$ suffices. However, the largest S_0 and S_1, the faster is the convergence (smaller matrices, etc.). A reasonable choice is:

$$S_0 = \{\, s \in S \mid Pr(s \vDash C\ \mathbf{U}\ B)=0 \,\} \text{ and } S_1 = \{\, s \in S \mid Pr(s \vDash C\ \mathbf{U}\ B)=1 \,\}$$

Bounded Until Properties. Taking $S_0 = S \setminus C\ \mathbf{U}\ B$ and $S_1 = B$ and $S_? = C \setminus B$ we have that $x^n(s) = \Pr(s \vDash C\ \mathbf{U}^{\leq n}\ B)$.

Remark: The n^{th} power of $\mathbf{A}$ contains probabilities to reach a state in exactly n steps. More precisely, $\mathbf{A}^n(s, t)$ is the sum of probabilities of all paths of the form $s=s_0 s_1 \ldots s_n=t$.

In other words: $\mathbf{A}^n(s, t)=\Pr(s \vDash S\ \mathbf{U}^{=n}\ t)$

Lesson 11b:

Qualitative properties

Qualitative properties

Qualitative properties require some event to happen with probability **1** or, dually, check if some event occurs with probability **0**.

Most of qualitative properties can be established just looking at the underlying digraph, because in a ***finite*** Markov chain ***almost surely*** paths **eventually enter in a Bottom Strongly Connected Component** (BSCC).

Persistence Properties. The event **GF** *B* is measurable. This event can be written as a countable intersections of countable unions of cylinder sets (prove this equality is an easy exercise):

$$\mathbf{GF}\ B = \bigcap_{n \geq 0} \bigcup_{m \geq n} Cyl(\text{“}m{+}1^{th} \text{ state is in B”})$$

Persistence properties of the form **FG** *B* are measurable as the complement of **GF** *B*. As a matter of fact, **FG** $B = S \setminus (\mathbf{GF}\ S \setminus B)$.

Probabilistic Choice & Fairness

In a Markov chain, if a state t is visited infinitely often, then almost surely all finite path fragments starting in t will be taken infinitely often.

Here "almost surely" has to be read as conditional probability: an event E **holds almost surely under another event** D, if $Pr(D) = Pr(E \cap D)$.

Theorem: Let $\mathcal{M}$ be a finite Markov Chain, and $s, t \in S$. Then:

$$Pr(s \vDash \mathbf{GF}\ t) = Pr(\bigwedge_{\pi \in \mathrm{PathFin}(t)} \mathbf{GF}\ \pi)$$

The above theorem implies that **each transition** (t, t') such that $\mathcal{P}(t, t') > 0$ will be taken almost surely if t is visited infinitely often. In this sense, **probabilistic choice is strongly fair**.

Theorem: Let $\mathcal{M}$ be a MC, and $s \in S$. Then:

$$Pr(\{\pi \in \mathrm{Path}(s) \mid inf(\pi) \in \mathrm{BSCC}(\mathcal{M})\} = 1$$

In every MC, almost surely, a path ends in a BSCC of $\mathcal{M}$.

Almost sure reachability

The problem of **almost sure reachability** amounts to determine the set of states that reach a given set of goal states B almost surely.

Theorem: Let $\mathcal{M}$ be a finite MC, $s \in S$, and $B \subseteq S$. Then the following statements are equivalent:

1. $Pr(s \vDash \mathbf{F}\ B) = 1$
2. $Post^*(t) \cap B \neq \varnothing$ for each $t \in Post^*(s)$
3. $s \in \mathrm{S} \setminus Pre^*(\mathrm{S} \setminus Pre^*(B))$

This theorem gives a purely graph-theoretic characterisation (**condition** (**3**)) of almost-sure reachability. Observe that from s such that $Pr(s \vDash \mathbf{F}\ B) = 1$ **we cannot go outside** $Pre^*(B)$.

Algorithm: Build the MC $\mathcal{M}_B$ where all states in B are made absorbing. Then use two backward reachability on $\mathcal{M}_B$ to compute the set of states $\mathrm{S} \setminus Pre^*(\mathrm{S} \setminus Pre^*(B))$ [the first from B and the second from $\mathrm{S} \setminus Pre^*(B)$]

Qualit. constrained reachability

The problem of **qualitative constrained reachability** amounts to determine the sets of states S_0 and S_1 such that:
$S_0 = \{ s \in S \mid Pr(s \vDash C\,\mathbf{U}\,B)=0 \}$ and $S_1 = \{ s \in S \mid Pr(s \vDash C\,\mathbf{U}\,B)=1\}$.

S_0 corresponds to the set of states satisfying $\neg\mathbf{E}\,(C\,\mathbf{U}\,B)$ and can be computed by a backward reachability from B.

As for S_1, we reduce the problem to an almost sure reachablity in a slightly modified Markov chain $\mathcal{M}'$. We make absorbing all states in B and in $S \setminus (C\,\mathbf{U}\,B)$.

- $\Pr_{\mathcal{M}}(s \vDash C\,\mathbf{U}\,B) = \Pr_{\mathcal{M}'}(s \vDash \mathbf{F}\,B)$ for all $s \in C \setminus B$
- $\Pr_{\mathcal{M}}(s \vDash C\,\mathbf{U}\,B) = \Pr_{\mathcal{M}'}(s \vDash \mathbf{F}\,B) = 1$ for all $s \in B$
- $\Pr_{\mathcal{M}}(s \vDash C\,\mathbf{U}\,B) = \Pr_{\mathcal{M}'}(s \vDash \mathbf{F}\,B) = 0$ for all $s \in S \setminus (C \setminus B)$

This give a polynomial algorithm (the transformation from $\mathcal{M}$ to $\mathcal{M}'$ is clearly linear in the size of $\mathcal{M}$).

Qualitative repeated reachability

Corollary: Let $\mathcal{M}$ be a finite MC, $s \in S$, and $B \subseteq S$. Then the following are equivalent:

- $Pr(s \vDash \mathbf{GF}\ B) = 1$
- $T \cap B \neq \varnothing$ for each BSCC T reachable from s.
- $s \vDash \mathbf{AG\ EF}\ B$.

Corollary: Let $\mathcal{M}$ be a finite MC, $s \in S$, and $B \subseteq S$ and let V be the union of all BSCC T of $\mathcal{M}$ such that $T \cap B \neq \varnothing$. Then:

$$Pr(s \vDash \mathbf{GF}\ B) = Pr(s \vDash \mathbf{F}\ V)$$

Lesson 12a:

Probabilistic CTL

Probabilistic CTL

Probabilistic CTL (PCTL) extends the syntax of CTL with a **probabilistic operator** $\mathbf{P}_{[a,b]}(\varphi)$ whose intended semantics is that the probability of φ is in the interval $[a, b]$ $(0 \leq a \leq b \leq 1)$.

In PCTL we can define **quantitative properties** to be checked in a Markov chain.

The **interpretation** of formula is **boolean**. $\mathbf{P}_{[a,b]}(\varphi)$ is the probabilistic counterpart of the path quantifiers **E** and **A**.

Example: In the communication protocol, the PCTL formula:

$$\mathbf{P}_{=1}(\mathbf{F}\ \text{delivered}) \wedge \mathbf{P}_{=1}(\mathbf{G}\ (\text{try} \rightarrow \mathbf{P}_{\geq 0.99}\ (\mathbf{F}^{\leq 3}\ \text{delivered})))$$

asserts that **almost surely** some message will be delivered and that almost surely, for any attempt to send a message, with probability at least 99% the message will be sent within 3 steps.

Probabilistic CTL: Syntax

State formulae:

$$\psi ::= \mathsf{true} \mid a \mid \psi_1 \wedge \psi_2 \mid \neg\psi \mid \mathbf{P}_J(\varphi)$$

where $a \in AP$, φ is a path formula, and $J \subseteq [0,1]$ is an interval with rational bounds.

Path formulae:

$$\varphi ::= \mathbf{X}\,\psi \mid \psi_1\,\mathbf{U}\,\psi_2 \mid \psi_1\,\mathbf{U}^{\leq n}\,\psi_2$$

where ψ_1, ψ_2 are state formulae and n is a natural number.

As in CTL, temporal operators **X** and **U** are required to be preceded by **P**. Intervals can be abbreviated: $\mathbf{P}_{\leq 0.5}$ means $\mathbf{P}_{[0,0.5]}$, $\mathbf{P}_{=1}$ means $\mathbf{P}_{[1,1]}$, and $\mathbf{P}_{>0}$ means $\mathbf{P}_{]0,1]}$ etc.

Semantics is similar to CTL. Step bounded until $\psi_1\,\mathbf{U}^{\leq n}\,\psi_2$ requires that ψ_2 holds after at most n steps.

Probabilistic CTL: semantics

The semantics is the same as that of CTL, except for $\mathbf{P}_J(\varphi)$ and bounded until. We have:

$$s \vDash \mathbf{P}_J(\varphi) \text{ iff } Pr(s \vDash \varphi) \in J$$

$$\pi \vDash \psi_1 \, \mathbf{U}^{\leq n} \, \psi_2 \text{ iff } \exists 0 \leq j \leq n.\ \pi_j \vDash \psi_2 \wedge (\forall 0 \leq k < j.\ \pi_k \vDash \psi_1)$$

Formally, we need to check whether events specified by PCTL path formulae are measurable.

Theorem: For each PCTL path formula φ and state s of a Markov chain, $\text{Path}(s, \varphi) = \{\pi \in \text{Path}(s) \mid \pi \vDash \varphi\}$ is measurable.

Proof: Induction on φ. If $\varphi \equiv \mathbf{X}\, \varphi'$, then Path(s, φ) is the union of Path(t, φ'), such that $t \vDash \varphi'$. If $\varphi \equiv \psi_1 \, \mathbf{U}^{\leq n} \, \psi_2$, then Path(s, φ) is the union of all cylinder sets $\text{Cyl}(s_0 s_1 \ldots s_k)$, where $k \leq n$, $s_k \vDash \psi_2$ and $s_i \vDash \psi_1$ for $0 \leq i < k$. If $\varphi \equiv \psi_1 \, \mathbf{U} \, \psi_2$, then Path(s, φ) can be written as $\bigcup_{n \geq 0} \{\pi \in \text{Path}(s) \mid \pi \vDash \psi_1 \, \mathbf{U}^{\leq n} \, \psi_2\}$. $\square$

Probabilistic CTL: equivalences

As usual, other operators, such as **F** and **R** as well as other boolean connectives can be derived using duality.
For example: $\mathbf{F}^{\leq n}\, \psi \equiv \mathsf{true}\ \mathbf{U}^{\leq n}\, \psi$.

We have that $\mathbf{P}_{<p}(\varphi) \equiv \mathbf{P}_{>p}(\neg\varphi)$ and $\mathbf{P}_{]a,b]}(\varphi) \equiv \neg\mathbf{P}_{\leq a}(\varphi) \wedge \mathbf{P}_{>b}(\varphi)$.

Be careful to the duality between lower and upper bounds! Therefore we could limit to consider only upper-bounds and one between $\mathbf{P}_{=1}$ and $\mathbf{P}_{=0}$ for qualitative properties.

If an event E holds with probability **at most** p, then the complementary event E holds with probability **at least** 1-p.

For example:

$$\mathbf{P}_{\leq p}(\mathbf{G}\ \varphi) \equiv \mathbf{P}_{\geq 1-p}(\mathbf{F}\ \neg\varphi) \text{ and } \mathbf{P}_{]p,q]}(\mathbf{G}^{\leq n}\ \varphi) \equiv \mathbf{P}_{[1-q,\ 1-p[}(\mathbf{F}^{\leq n}\ \neg\varphi).$$

PCTL: proving equivalences

Let us consider the equivalence:

$$\mathbf{P}_{>0}(\mathbf{X}\ \mathbf{P}_{>0}(\mathbf{F}\ \psi)) \equiv \mathbf{P}_{>0}(\mathbf{F}\ \mathbf{P}_{>0}(\mathbf{X}\ \psi))$$

($\Rightarrow$) Let s be such that $s \models \mathbf{P}_{>0}(\mathbf{X}\ \mathbf{P}_{>0}(\mathbf{F}\ \psi))$, then there exists t, such that $\mathcal{P}(s, t)>0$ and $t \models \mathbf{P}_{>0}(\mathbf{F}\ \psi)$ and therefore there exists a finite path $t_0 t_1 \ldots t_k$ where $t=t_0$ and $t_k \models \psi$. Therefore $t_{k-1} \models \mathbf{X}\ \psi$. Since $s\ t_0 t_1 \ldots t_{k-1}$ is a path fragment starting in s with positive probability, we have $s \models \mathbf{P}_{>0}(\mathbf{F}\ \mathbf{P}_{>0}(\mathbf{X}\ \psi))$.

($\Leftarrow$) Conversely, if $s \models \mathbf{P}_{>0}(\mathbf{F}\ \mathbf{P}_{>0}(\mathbf{X}\ \psi))$ then there exists a path fragment $s_0 s_1 \ldots s_k$ with $s=s_0$ and $s_k \models \mathbf{P}_{>0}(\mathbf{X}\ \psi)$, but this means that s_k has a successor t such that $t \models \psi$. This means that the path fragment $s_1 \ldots s_k t$ is a witness for $s_1 \models \mathbf{P}_{>0}(\mathbf{F}\ \psi)$ and hence $s \models \mathbf{P}_{>0}(\mathbf{X}\ \mathbf{P}_{>0}(\mathbf{F}\ \psi))$.

PCTL model checking

The problem is to verify in a Markov chain if $s \vDash \varphi$, where φ is a PCTL formula. As for CTL, **the idea is to compute set of states Sat(ψ) for all subformulae ψ** of φ. For propositional subformulae, the problem is essentially the same as in CTL, so the interesting case is to determine $\mathrm{Sat}(\mathbf{P}_J\, \psi) = \{s \in S \mid Pr(s \vDash \psi) \in J\}$.

As for the operator **X**, it suffices to multiply the matrix $\mathcal{P}$ by the characteristic vector of Sat(ψ):

$$Pr(s \vDash \mathbf{X}\, \psi) = \sum_{s' \in \mathrm{Sat}(\psi)} \mathcal{P}(s, s')$$

If we have formulae of the form $\psi_1\, \mathbf{U}^{\leq n}\, \psi_2$ or $\psi_1\, \mathbf{U}\, \psi_2$, we can just use technique we have seen for constrained reachability, where C=Sat(ψ_1) and B=Sat(ψ_2).

As for the bounded operator $\mathbf{U}^{\leq n}$ we have to stop after n iterations.

PCTL model checking

Theorem: Let $\mathcal{M}$ be a finite MC and φ be a PCTL formula. The model checking problem $\mathcal{M} \models \varphi$ can be solved in time $\mathcal{O}(poly(size(\mathcal{M})) \cdot n_{\max} \cdot |\varphi|)$ where $n_{\max}$ is the maximum step bound that appears in formulae of the form $\psi_1\ \mathbf{U}^{\leq n}\ \psi_2$.

For efficiency reasons, **qualitative properties** such as $\mathbf{P}_{=1}(\psi_1\ \mathbf{U}\ \psi_2)$ or $\mathbf{P}_{>0}(\psi_1\ \mathbf{U}\ \psi_2)$ are solved by using graph-based algorithms [this avoids solving systems of linear equations].

A **counterexample** or **witness** in PCTL is a set of path fragments that show the refutation or satisfaction of a formula.

Counterexamples and witnesses

Example: If $s \nvDash \mathbf{P}_{\leq p}(\mathbf{F}\ \psi)$, then $Pr(s \vDash \mathbf{F}\ \psi)>p$. A proof is a set Π of finite path fragments such that for all $\pi \in \Pi$, $\pi = s_0 s_1 \ldots s_k$, $s_k \vDash \psi$ and for $i<k$, $s_i \nvDash \psi$ and $\sum_{\pi \in \Pi} Pr(\pi)>p$.

If $s \nvDash \mathbf{P}_{\geq p}(\mathbf{F}\ \psi)$, is obtained by a set Π of path that refute $\mathbf{F}\ \psi$. These paths have the shape $\pi = s_0 s_1 \ldots s_k$, for $i \leq k$, $s_i \nvDash \psi$ s_i, and s_k belongs to a BSCC C of $\mathcal{M}$ such that $\mathrm{C} \cap \mathrm{Sat}(\psi) = \varnothing$. Moreover we must have that $\sum_{\pi \in \Pi} Pr(\pi)>1-p$. The cylinder sets $Cyl(\pi)$ satisfies $\mathbf{G}\neg\psi$ paths.

To compute $Pr(s \vDash \mathbf{G}\neg\psi)$ it is necessary to consider paths that reach a BSCC T of $\mathcal{M}$ such that $\mathrm{C} \cap \mathrm{Sat}(\psi) = \varnothing$ through $\neg\psi$ states: we can collect all such paths (increasing k) until the probability is greater than $1-p$.

PCTL model checking: Example

Let us consider the MC below. Let us assume that we are checking the property $\mathbf{P}_{\leq 1/2}(\mathbf{F}\ b)$ and that s_0 is the initial state.

$\mathcal{M} \not\models \mathbf{P}_{\leq 1/2}(\mathbf{F}\ b)$ is witnessed by three paths:

$$\{s_0 s_1 t_1,\ s_0 s_1 s_2 t_1\ ,\ s_0 s_2 t_1\}$$

whose probability is 0.2+0.2+0.15=0.55>0.5=1/2.

Observe that the counterexample is not unique. There are other paths such as $s_0 s_1 s_2 t_2$ and $s_0 s_2 t_2$.

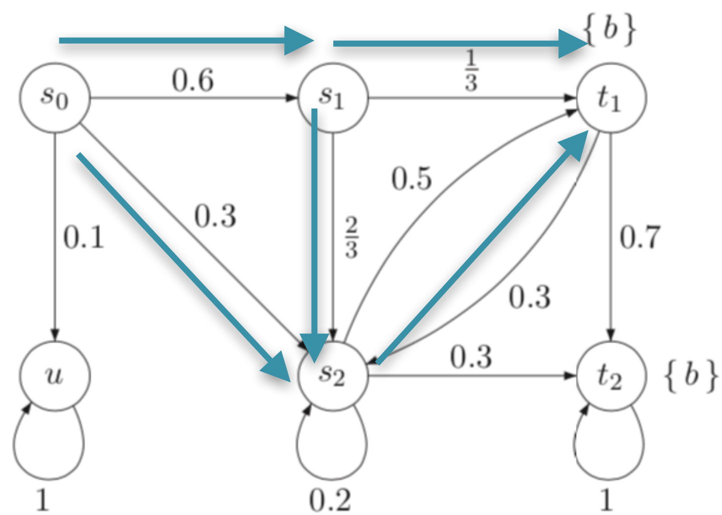

Qualitative fragment of PCTL

The goal here is to compare the expressive power of PCTL wrt CTL. **It is evident that quantitative properties cannot be expressed in CTL**. But what about **qualitative properties**?

State formulae:

$$\psi \ ::= \ \mathsf{true} \ | \ a \ | \ \psi_1 \wedge \psi_2 \ | \ \neg\psi \ | \ \mathbf{P}_{>0}(\varphi) \ | \ \mathbf{P}_{=1}(\varphi)$$

where $a \in AP$, φ is a path formula.

Path formulae:

$$\varphi \ ::= \ \mathbf{X}\,\psi \ | \ \psi_1 \,\mathbf{U}\, \psi_2$$

where ψ_1, ψ_2 are state formulae.

Observations: $\mathbf{P}_{=0}(\varphi)=\neg\mathbf{P}_{>0}(\varphi)$ and $\mathbf{P}_{<1}(\varphi)=\neg\mathbf{P}_{=1}(\varphi)$.

Definition: The PCTL formula φ is **equivalent** to the CTL formula ψ, notation $\varphi\equiv\psi$ iff $\mathrm{Sat}(\varphi)=\mathrm{Sat}(\psi)$ **for all MC $\mathcal{M}$**.

"Trivial" Equivalences

It is well-known that `almost surely' differs from **A**, because of some path with zero probability. In the MC below, we have $s \vDash \mathbf{P}_{=1}(\mathbf{F}\ a)$ but $s \nvDash \mathbf{A}\ \mathbf{F}\ a$. The converse always holds.

For certain formulae, $\mathbf{P}_{=1}$ corresponds to **A** and $\mathbf{P}_{>0}$ corresponds to **E**. For example: $s \vDash \mathbf{P}_{=1}(\mathbf{X}\ \varphi) \Leftrightarrow s \vDash \mathbf{A}\ \mathbf{X}\ \varphi$ and $s \vDash \mathbf{P}_{>0}(\mathbf{X}\ \varphi) \Leftrightarrow s \vDash \mathbf{EX}\ \varphi$.

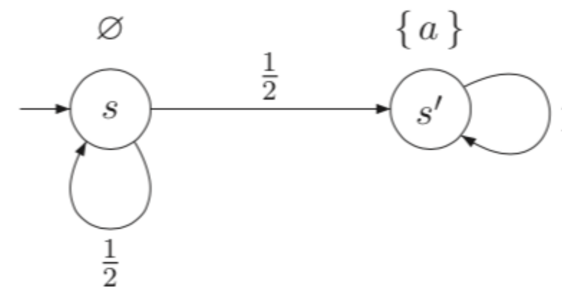

We have also: $s \vDash \mathbf{P}_{>0}(\mathbf{F}\ \varphi) \Leftrightarrow s \vDash \mathbf{EF}\ \varphi$ and $s \vDash \mathbf{P}_{=1}(\mathbf{G}\ \varphi) \Leftrightarrow s \vDash \mathbf{AG}\ \varphi$

We show how to prove this statements:

Assuming $s \vDash \mathbf{P}_{>0}(\mathbf{F}\ \varphi)$, we have $Pr(s \vDash \mathbf{F}\ \varphi)>0$ that implies that there exists a finite path fragment whose last state satisfies φ. But this path fragment is a witness of $s \vDash \mathbf{E}\ \mathbf{F}\ \varphi$ in CTL.

Conversely, assuming $s \vDash \mathbf{E}\ \mathbf{F}\ \varphi$ we have that there exist a finite path fragment and its cylinder satisfies $s \vDash \mathbf{P}_{>0}(\mathbf{F}\ \varphi)$.

The other statement follows by duality.

PCTL and fairness

As we have seen, often a 0-probability loop makes the difference between a PCTL property $\mathbf{P}_{=1}(\varphi)$ and a CTL $\mathbf{A}\ \varphi$.

Let us define the following strong fairness constraints:

$$\text{sfair} = \bigwedge_{s \in S} \bigwedge_{t \in \text{post}(s)} \mathbf{GF}\ s \rightarrow \mathbf{GF}\ t$$

Then we have the following equivalences:

$s \models \mathbf{P}_{=1}(\varphi\ \mathbf{U}\ \psi) \Leftrightarrow s \models_{\text{sfair}} \mathbf{A}\ (\varphi\ \mathbf{U}\ \psi)$ and $s \models \mathbf{P}_{>0}(\mathbf{G}\ \varphi) \Leftrightarrow s \models_{\text{sfair}} \mathbf{E}\ \mathbf{G}\ \varphi$

Therefore, qualitative **PCTL** is a **sort of CTL plus strong fairness**.

Lesson 11

That's all Folks…

…Questions?