Formal Methods for System Design

Chapter 6: Model checking probabilistic systems

Mickael Randour

Mathematics Department, UMONS

October 2023

1 Markov chains

2 Reachability and limit behavior

3 PCTL: probabilistic CTL

4 Weighted Markov chains: venturing into the land of quantitative specifications

1 Markov chains

2 Reachability and limit behavior

3 PCTL: probabilistic CTL

4 Weighted Markov chains: venturing into the land of quantitative specifications

Probabilistic systems

Why?

Many real-life systems exhibit *stochastic aspects*. Some examples:

- message loss in communication protocols,
- randomized algorithms (e.g., leader election in distributed systems using coin-tossing to break symmetry),
- quantitative evaluation of system performance (e.g., expected response time).

Probabilities vs. non-determinism

Enriching TSs with actual probabilities instead of simply non-determinism can be useful to analyze more precisely the behavior of a system, on the *quantitative* level.

E.g., some systems may be unable to totally prevent message loss but be able to keep the *probability of this event very small*, which in practice may be sufficient.

Some formal models for probabilistic systems

	Stochastic transitions only	Stochastic & non-deterministic transitions
Discrete time	DT Markov chain (MC)	Markov decision process (MDP)
Continuous time	CTMC	CTMDP

$\Longrightarrow$ **Focus of this chapter.**

But first, who is Markov?

Someone with an awesome mustache!

Yes, but also. . .

Andrey Andreyevich Markov

- Russian mathematician, 1856-1922,
- studied **stochastic processes**.

In 1913, he studied how letters succeed each other in a novel of Alexander Pushkin: he saw that the probability of a letter depends *almost exclusively* on its direct predecessor.

$\Longrightarrow$ Appearance of the **Markov property**.

The models studied here are called “Markov” models because they satisfy this property: they are not all due to Markov.

Markov property

Markov property

A stochastic process satisfies the Markov property if the conditional probability distribution of future states of the process (conditional on *both past and present states*) depends *only upon the present state*, not on the sequence of events that preceded it.

E.g., game of the goose, Brownian motion, Markov chains. . .

Markov chains

An example: simple communication protocol

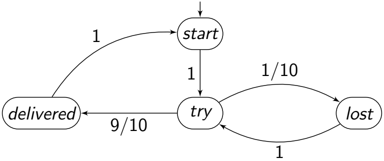

- As in the last chapter, we do not care about *actions*.
- Transitions are marked with probabilities.
  - ▷ Messages are lost with probability 1/10.

Natural questions could be:

- What is the probability that a message is *eventually* delivered?
- Same but *in at most 3 tries*?
- What is the expected (i.e., “average”) number of tries before a message is delivered?

$\Longrightarrow$ **We will see how to answer such questions.**

Markov chains

Formal definition

Definition: (discrete-time) Markov chain (MC)

A (discrete-time) Markov chain (MC) is a tuple $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ where

- S is a countable, nonempty set of *states*;
- $\mathbf{P}\colon S \times S \to [0, 1]$ is the *transition probability function* such that for all $s \in S$, $\sum_{s' \in S} \mathbf{P}(s, s') = 1$;
- $\iota_{\text{init}}\colon S \to [0, 1]$ is the *initial distribution* such that $\sum_{s \in S} \iota_{\text{init}}(s) = 1$;
- AP is the set of *atomic propositions* and $L\colon S \to 2^{AP}$ the *labeling function*.

We mainly consider *finite* MCs.

⚠ **For algorithmic purposes, probabilities supposed rational.**

Markov chains

Related concepts

Classical notions introduced for TSs carry over to MCs:

- *Successors*. State s' is a successor of s iff $\mathbf{P}(s, s') > 0$.
- *Paths*. Same idea.

$\Longrightarrow$ Essentially, one can see an MC as a TS by forgetting the probabilities and applying previously studied techniques.

$\Longrightarrow$ **Next, we focus on techniques specific to MCs.**

$\Longrightarrow$ **This lecture is only an introduction to the rich theory of MCs and related probabilistic models...**

Markov chains

Back to the example

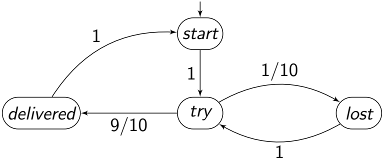

- $S = \{\textit{start}, \textit{try}, \textit{lost}, \textit{delivered}\}$,
- Initial states and transition function seen as matrices:

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

- For $T = \{\textit{lost}, \textit{delivered}\}$,
  $\mathbf{P}(\textit{try}, T) = \begin{pmatrix} 0 & 0 & \frac{1}{10} & \frac{9}{10} \end{pmatrix} \cdot \begin{pmatrix} 0 & 0 & 1 & 1 \end{pmatrix}^{T} = 1.$

Markov chains

Another example: Knuth's die (aka, how to throw a die by tossing a coin)

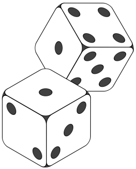

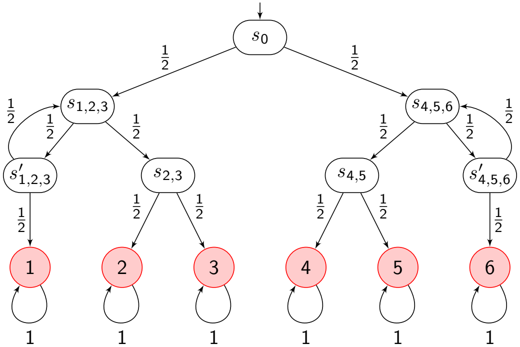

$\implies$ Are you convinced that this MC simulates a **fair** die?

$\implies$ **How can we prove it?**

$\implies$ **Need to properly define a probability measure.**

But let's start with intuition. . .

Markov chains

Another example: Knuth's die (aka, how to throw a die by tossing a coin)

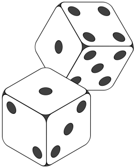

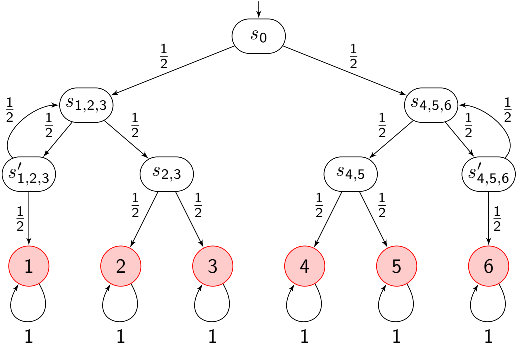

What is the probability to be in s' after n steps, starting from s?

- ▷ $p_{s,s'}(0) = 1$ iff $s' = s$, 0 otherwise. $p_{s,s'}(1) = \mathbf{P}(s, s')$.
- ▷ $p_{s,s'}(n) = \sum_{s'' \in S} p_{s,s''}(m) \cdot p_{s'',s'}(n-m)$ for $n > 1$, $0 < m < n$.
  (Chapman–Kolmogorov equation)

Markov chains

Another example: Knuth's die (aka, how to throw a die by tossing a coin)

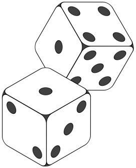

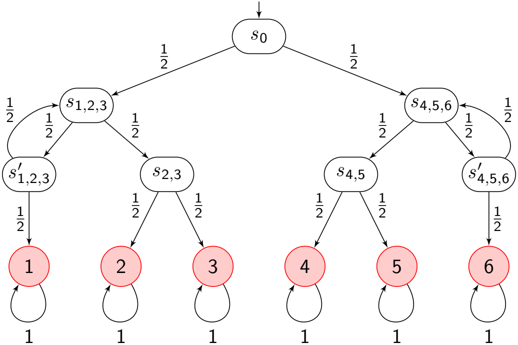

Probability to be in s' from the initial distribution, after n steps?

▷ Now using matrices: $p_{\iota_{\text{init}}, s'}(n) = \sum_{s \in S} \iota_{\text{init}}(s) \cdot \mathbf{P}^n(s, s')$.

↪ Here $\mathbf{P}^n$ is the n-th power of matrix $\mathbf{P}$.

Markov chains

Another example: Knuth's die (aka, how to throw a die by tossing a coin)

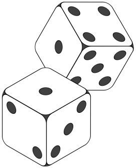

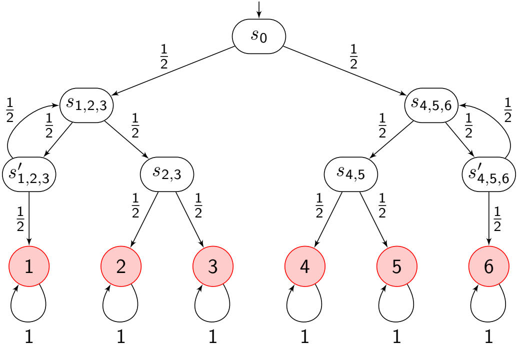

Here,

- ▷ after 1 step, probability $1/2$ to be in either $s_{1,2,3}$ or $s_{4,5,6}$;
- ▷ after 2 steps, $1/4$ for each state of level 3;
- ▷ after 3 steps, $1/8$ for each leaf and for both $s_{1,2,3}$ and $s_{4,5,6}$.

Markov chains

Another example: Knuth's die (aka, how to throw a die by tossing a coin)

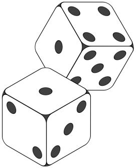

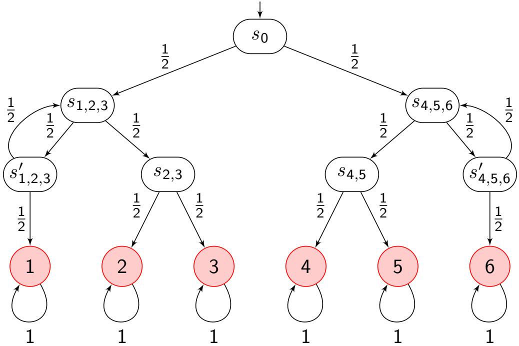

$\implies$ Leaves are **absorbing states**.

Continuing, after 5 steps, $\frac{1}{8} + \frac{1}{8} \cdot \frac{1}{4}$ for each leaf and $\frac{1}{8} \cdot \frac{1}{4}$ for $s_{1,2,3}$ and $s_{4,5,6}$.

$\implies$ **At the limit, we obtain $1/6$ for each leaf.**

Markov chains

Another example: Knuth's die (aka, how to throw a die by tossing a coin)

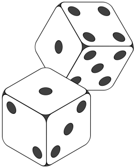

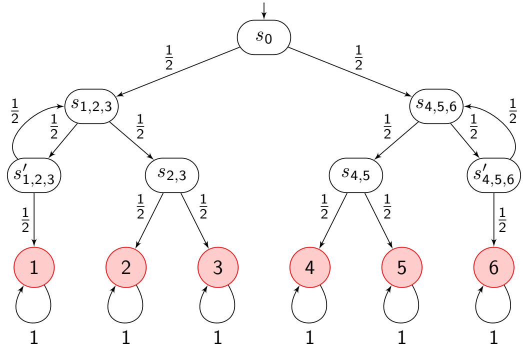

Observe that at any point in time, all outcomes of the die (i.e., leaves of the MC) are equally likely.

$\Longrightarrow$ **Proper simulation of a fair die with a fair coin.**

Markov chains

Another example: Knuth's die (aka, how to throw a die by tossing a coin)

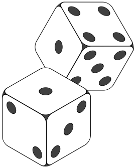

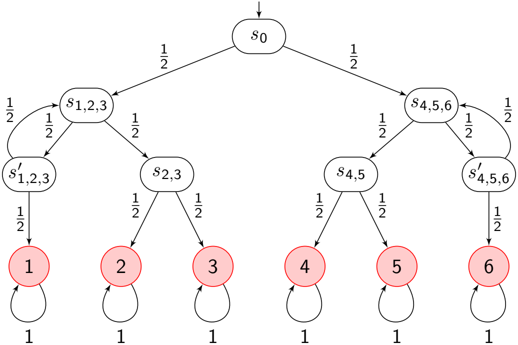

Technically possible to visit $s_{1,2,3}$ infinitely often (hence never reaching a leaf) but **probability of such an event is null**.
$\Longrightarrow$ Upcoming concepts of *bottom strongly connected components (BSCCs)* (here, the leaves) and *transient states* (here, everything else).

Markov chains

Back to lossy communication again

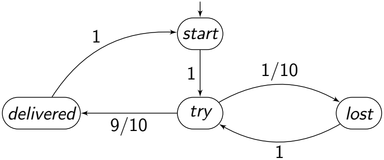

Here, also, it seems that the *probability that a message is eventually delivered is one*, while the path $\pi = \mathit{start} \cdot (\mathit{try} \cdot \mathit{lost})^{\omega}$ is a perfectly valid path in the underlying TS.

$\Longrightarrow$ **Let's discuss how one can define a proper notion of probability on MCs.**

Probability measure on MCs

Defining a probability space

Goal

To reason about the behavior of MCs, we need to define a *probability space* over (sets of) paths.

⚠ Doing this formally requires measure theory and notions such as σ-algebrae.

$\Longrightarrow$ **Here, we only sketch the main steps.**

$\Longrightarrow$ **For a formal presentation, see the book.**

Probability measure on MCs

Intuition

- What are the possible *outcomes* of an MC?
  - ▷ All (infinite) paths in $Paths(\mathcal{M})$ (defined as for TSs).
- What are the *events* we want to characterize?
  - ▷ Subsets of $Paths(\mathcal{M})$. E.g., given a target set T, what is the probability of the event $\{\pi \in Paths(\mathcal{M}) \mid \pi \models \Diamond T\}$, often written as $\Diamond T$?

$\Longrightarrow$ To define properly those events and be able to put a probability measure on them, we rely on **cylinder sets**.

Probability measure on MCs

Cylinder sets

Definition: cylinder set of a finite path

The *cylinder set* of $\widehat{\pi} = s_0 \dots s_n \in \mathit{Paths}_{\mathit{fin}}(\mathcal{M})$ is defined as

$$\mathit{Cyl}(\widehat{\pi}) = \{\pi \in \mathit{Paths}(\mathcal{M}) \mid \widehat{\pi} \text{ is a prefix of } \pi\}.$$

It is the set of all infinite continuations of $\widehat{\pi}$.

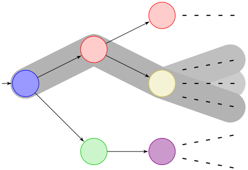

Seeing an MC through its infinite tree unfolding, one can picture cylinder sets as the combination of a finite branch + the corresponding subtree. E.g., here in grey, is the cylinder set of the finite path 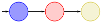.

Probability measure on MCs

Probability space

Probability space of an MC

The set of *events* of the probability space for an MC $\mathcal{M}$ contains *all cylinder sets* $Cyl(\widehat{\pi})$ where $\widehat{\pi}$ ranges over all finite paths in $Paths_{fin}(\mathcal{M})$.

Now, what is the probability of a cylinder set?

Probability measure on MCs

Probability of cylinder sets

Definition: cylinder set of a finite path

The *cylinder set* of $\widehat{\pi} = s_0 \dots s_n \in \mathit{Paths}_{\mathit{fin}}(\mathcal{M})$ is defined as

$$\mathit{Cyl}(\widehat{\pi}) = \{\pi \in \mathit{Paths}(\mathcal{M}) \mid \widehat{\pi} \text{ is a prefix of } \pi\}.$$

It is the set of all infinite continuations of $\widehat{\pi}$.

Probability measure

There exists a unique *probability measure* $\mathbb{P}^{\mathcal{M}}$ defined by

$$\mathbb{P}^{\mathcal{M}}(\mathit{Cyl}(s_0 \dots s_n)) = \iota_{\text{init}}(s_0) \cdot \mathbf{P}(s_0 \dots s_n)$$

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

$\Longrightarrow$ **Essentially the probability of prefix $s_0 \dots s_n$.**

Probability measure on MCs

Measurable events (1/2)

Measurability

To be able to define the probability of an event, this event must be measurable.

Good news

Cylinder sets are measurable, and any event defined using *complement and/or countable unions* of cylinder sets are also measurable.

Examples

Events such as $\Diamond T$, $\Box T$, $C \mathsf{U} T$, $\Diamond\Box T$ and $\Box\Diamond T$ are measurable.

$\Longrightarrow$ **See next slide.**

Probability measure on MCs

Measurable events (2/2)

Take the case $\Diamond T$. This event can be expressed as the countable union of *all cylinders* $Cyl(s_0 \ldots s_n)$ where $s_0, \ldots, s_{n-1} \notin T$ and $s_n \in T$:

$$\Diamond T = \bigcup_{s_0 \ldots s_n \in Paths_{fin}(\mathcal{M}) \cap (S \setminus T)^* T} Cyl(s_0 \ldots s_n).$$

Hence it is measurable. Since all cylinders are pairwise disjoint, its probability (we drop $\mathcal{M}$ when the context is clear) is given by

$$\begin{aligned}\mathbb{P}(\Diamond T) &= \sum_{s_0 \ldots s_n \in Paths_{fin}(\mathcal{M}) \cap (S \setminus T)^* T} \mathbb{P}(Cyl(s_0 \ldots s_n)) \\ &= \sum_{s_0 \ldots s_n \in Paths_{fin}(\mathcal{M}) \cap (S \setminus T)^* T} \iota_{\text{init}}(s_0) \cdot \mathbf{P}(s_0 \ldots s_n)\end{aligned}$$

Probability measure on MCs

Back to Knuth's die

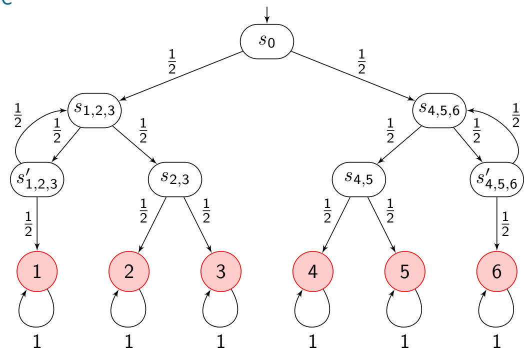

Using this approach, we can formalize the probability of $\Diamond 2$.

$$\mathbb{P}(\Diamond 2) = \sum_{s_0 \ldots s_n \in (S \setminus 2)^* 2} \mathbf{P}(s_0 \ldots s_n)$$

$$= \mathbf{P}(s_0 s_{1,2,3} s_{2,3} 2) + \mathbf{P}(s_0 s_{1,2,3} s'_{1,2,3} s_{1,2,3} s_{2,3} 2) + \ldots$$

Probability measure on MCs

Back to Knuth's die

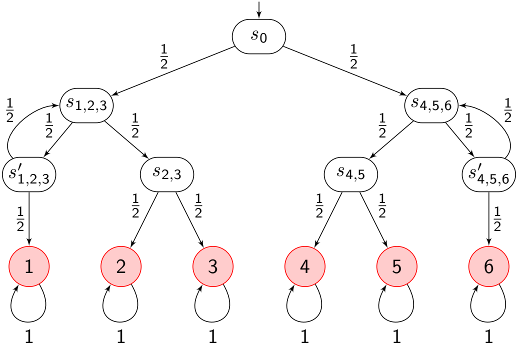

Thus $\mathbb{P}(\Diamond 2) = \sum_{i=0}^{\infty} \mathbf{P}(s_0 s_{1,2,3}(s'_{1,2,3}s_{1,2,3})^i s_{2,3}2) = \sum_{i=0}^{\infty} \frac{1}{8} \cdot \left(\frac{1}{4}\right)^i$.

This is a geometric series: $\mathbb{P}(\Diamond 2) = \frac{1}{8} \cdot \frac{1}{1-\frac{1}{4}} = \frac{1}{6}$.

$\Longrightarrow$ **Applying the same process to all leaves we get that the die is indeed fair.**

Probability measure on MCs

Back to Knuth's die

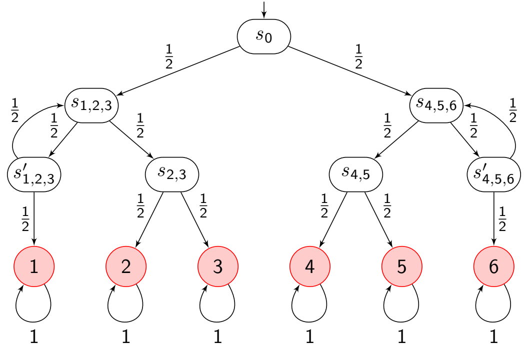

Thus $\mathbb{P}(\diamondsuit 2) = \sum_{i=0}^{\infty} \mathbf{P}(s_0 s_{1,2,3} (s'_{1,2,3} s_{1,2,3})^i s_{2,3} 2) = \sum_{i=0}^{\infty} \frac{1}{8} \cdot \left(\frac{1}{4}\right)^i$.

This is a geometric series: $\mathbb{P}(\diamondsuit 2) = \frac{1}{8} \cdot \frac{1}{1-\frac{1}{4}} = \frac{1}{6}$.

$\implies$ **We will see easier ways to compute reachability probabilities in the next section.**

1 Markov chains

2 Reachability and limit behavior

3 PCTL: probabilistic CTL

4 Weighted Markov chains: venturing into the land of quantitative specifications

Reachability

Via linear equations

Goal: given an MC $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$, $T \subseteq S$ and $s \in S$, we want to compute $\mathbb{P}_s(\Diamond T) = \mathbb{P}_s(\{\pi \in \mathit{Paths}(s) \mid \pi \models \Diamond T\})$, where $\mathbb{P}_s$ denotes the probability measure in $\mathcal{M}$ with single initial state s.

Characterization of reachability probabilities. Let $x_s = \mathbb{P}_s(\Diamond T)$ for all $s \in S$.

- ▷ If T cannot be reached from s, then $x_s = 0$ (cf. underlying graph).
- ▷ If $s \in T$, then $x_s = 1$.
- ▷ For any $s \in \mathit{Pre}^*(T) \setminus T$:

$$x_s = \underbrace{\sum_{s' \in S \setminus T} \mathbf{P}(s, s') \cdot x_{s'}}_{\text{reach } T \text{ via } s' \in S \setminus T} + \underbrace{\sum_{s'' \in T} \mathbf{P}(s, s'')}_{\text{reach } T \text{ in one step}} .$$

Reachability

Back to Knuth's die again

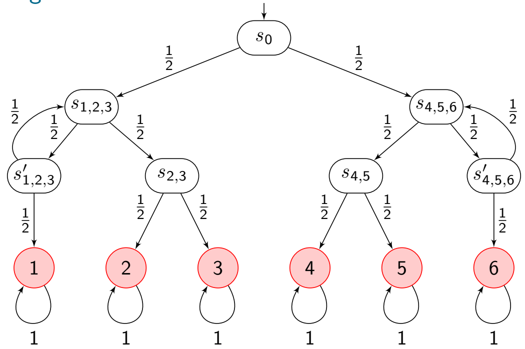

Computing $\mathbb{P}_{s_0}(\Diamond 2)$ via linear equations instead of infinite series?

- ▷ $x_2 = 1$ and $x_1 = x_3 = x_4 = x_5 = x_6 = 0$.
- ▷ $x_{s_{4,5}} = x_{s'_{4,5,6}} = x_{s_{4,5,6}} = 0$ and $x_{s_{2,3}} = \frac{1}{2}$.
- ▷ $x_{s_{1,2,3}} = \frac{1}{2} x_{s'_{1,2,3}} + \frac{1}{2} x_{s_{2,3}}$ and $x_{s'_{1,2,3}} = \frac{1}{2} x_{s_{1,2,3}}$.

Reachability

Back to Knuth's die again

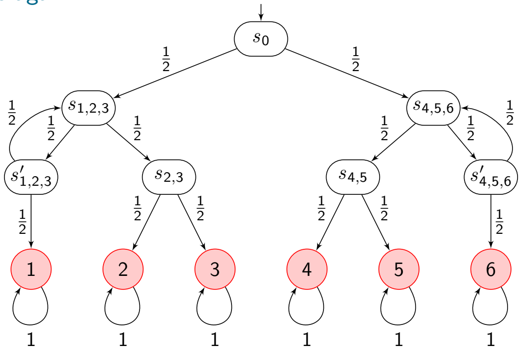

Solving $x_{s_{1,2,3}} = \frac{1}{2}x_{s'_{1,2,3}} + \frac{1}{2}x_{s_{2,3}}$ and $x_{s'_{1,2,3}} = \frac{1}{2}x_{s_{1,2,3}}$ yields:

- ▷ $x_{s_{1,2,3}} = \frac{1}{3}$ and $x_{s'_{1,2,3}} = \frac{1}{6}$.
- ▷ Finally, $x_{s_0} = \frac{1}{2}x_{s_{1,2,3}} = \frac{1}{6}$.

$\Longrightarrow$ **We obtain the correct result in a simpler way.**

Constrained reachability

Going further

We can generalize this approach, and formulate it using matrices, to deal with events of the type $C \,\mathsf{U}\, T$.

Theorem

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be a finite MC with $C, T \subseteq S$. Let

- $S_{=0} = Sat(\neg\exists(C \,\mathsf{U}\, T))$ (i.e., states for which no path exists),
- $T \subseteq S_{=1} \subseteq \{s \in S \mid \mathbb{P}(s \models C \,\mathsf{U}\, T) = 1\}$ (i.e., states for which we know the probability to be one),
- $S_{?} = S \setminus (S_{=0} \cup S_{=1})$.

Then, vector $(\mathbb{P}(s \models C \,\mathsf{U}\, T))_{s \in S_?}$ is the unique solution of the equation system $\mathsf{x} = \mathsf{A}\mathsf{x} + \mathsf{b}$ where $\mathsf{A} = (\mathbf{P}(s, s'))_{s,s' \in S_?}$ and $\mathsf{b} = (\mathbf{P}(s, S_{=1}))_{s \in S_?}$.

$\Longrightarrow$ **Essentially the same ideas as before, but let's work it out on a blackboard example.**

Constrained reachability

Example: summary

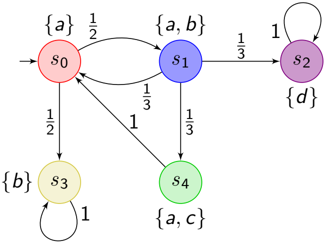

- $AP = \{a, b, c, d\}$.
- $\mathbb{P}^{\mathcal{M}}(\underbrace{\neg c}_{C} \,\mathsf{U}\, \underbrace{d}_{T})$?
- $S_{=0} = \{s_3, s_4\}$, $S_{=1} = \{s_2\}$.

Equation system:

$$\begin{pmatrix} x_{s_0} \\ x_{s_1} \end{pmatrix} = \begin{pmatrix} 0 & \frac{1}{2} \\ \frac{1}{3} & 0 \end{pmatrix} \cdot \begin{pmatrix} x_{s_0} \\ x_{s_1} \end{pmatrix} + \begin{pmatrix} 0 \\ \frac{1}{3} \end{pmatrix}.$$

Solution: $x_{s_0} = \frac{1}{5}$ and $x_{s_1} = \frac{2}{5}$.

$$\implies \mathbb{P}^{\mathcal{M}}(\neg c \,\mathsf{U}\, d) = \frac{1}{5}.$$

Constrained reachability

Deriving other events

Observe that being able to compute the probability of event $C \mathsf{U} T$ also permits to consider other classical events:

- $\Diamond T = S \mathsf{U} T$,
- $\Box T = \overline{\Diamond \overline{T}}$ (complement),
  - ▷ Hence $\mathbb{P}^{\mathcal{M}}(\Box T) = 1 - \mathbb{P}^{\mathcal{M}}(\Diamond \overline{T})$.
- We will come back to $\Diamond\Box T$ and $\Box\Diamond T$ when considering *limit behavior* of MCs and BSCCs.

Constrained reachability

Iterative approach via least fixed point computation

Theorem

For $S_{=0} = Sat(\neg\exists(C \,\mathsf{U}\, T))$, $S_{=1} = T$ and $S_? = S \setminus (S_{=0} \cup S_{=1})$, the vector $\mathrm{x} = (\mathbb{P}(s \models C \,\mathsf{U}\, T))_{s \in S_?}$ is the *least fixed point* of the operator $\Upsilon\colon [0,1]^{S_?} \to [0,1]^{S_?}$ given by

$$\Upsilon(\mathrm{y}) = \mathrm{A} \cdot \mathrm{y} + \mathrm{b}.$$

Furthermore, if $\mathrm{x}^{(0)} = 0$ is the vector consisting of zeros only, and $\mathrm{x}^{(n+1)} = \Upsilon(\mathrm{x}^{(n)})$ for $n \geq 0$, then

- $\mathrm{x}^{(n)} = (x_s^{(n)})_{s \in S_?}$ where $x_s^{(n)} = \mathbb{P}(s \models C \,\mathsf{U}^{\leq n}\, T)$ for $s \in S_?$,
- $\mathrm{x}^{(0)} \leq \mathrm{x}^{(1)} \leq \ldots \leq \mathrm{x}$, and
- $\mathrm{x} = \lim_{n \to \infty} \mathrm{x}^{(n)}$.

$\Longrightarrow$ **This also gives a way to compute the reachability probability in at most n steps.**

Constrained reachability

Iterative approach: example for the lossy communication protocol

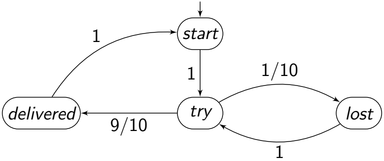

Recall those two questions:

- What is the probability that a message is *eventually* delivered?
  - ▷ $\mathbb{P}^{\mathcal{M}}(\diamondsuit \mathit{delivered}) = 1$.
- Same but *in at most 3 tries*?
  - ▷ $\mathbb{P}^{\mathcal{M}}(\diamondsuit^{\leq \text{ 3 tries}} \mathit{delivered}) = 999/1000$.

$\Longrightarrow$ **Blackboard computation.**

Limit behavior of MCs

Intuition

Recall the two examples studied before.

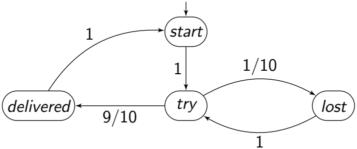

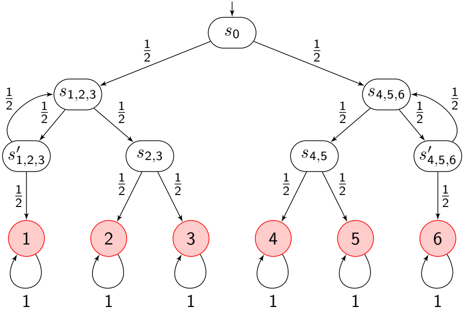

In the left MC, looping on *lost* forever has probability zero: hence all states will be visited infinitely often with probability one.

In the right MC, with probability one we reach one of the absorbing leaves and the other states are never seen again.

Limit behavior of MCs

Intuition

Recall the two examples studied before.

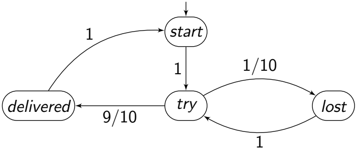

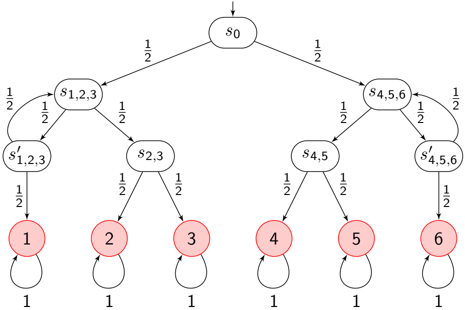

Each leaf in the right MC, as well as the whole left MC are **bottom strongly connected components**: intuitively, it is impossible to leave and all states are visited infinitely often with probability one.

Every other state of the right MC is visited finitely often with probability one: they are called **transient states**.

Limit behavior of MCs

Bottom strongly connected components (BSCCs)

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be an MC and $T \subseteq S$.

- T is *strongly connected* if for any $s, s' \in T$, there is a path *via edges in T* from s to s'.
- T is a *strongly connected component (SCC)* of $\mathcal{M}$ if T is strongly connected and no proper superset of T is strongly connected.
- T is a *bottom SCC (BSCC)* of $\mathcal{M}$ if T is an SCC and no state outside T can be reached, i.e., for any $s \in T$, $\mathbf{P}(s, T) = 1$.

$\Longrightarrow$ **Once in a BSCC, we never leave it, and we visit all states infinitely often with probability one.**

Intuition. Anytime we see a state, positive probability to visit any other state in the future thanks to strong connectivity. Since we never leave the BSCC, this possibility appears repeatedly and the probability that we never visit a given state again is zero.

Limit behavior of MCs

BSCCs: examples

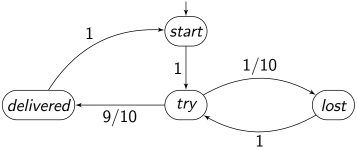

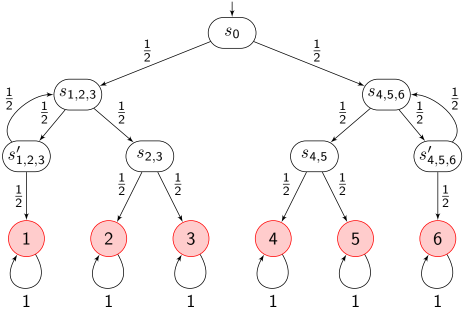

In the left MC, $\{try, lost\}$ is strongly connected but not an SCC because S is a proper superset and is an SCC. Furthermore, S is a BSCC.

In the right MC, $\{s'_{1,2,3}, s_{1,2,3}\}$ and $\{s'_{4,5,6}, s_{4,5,6}\}$ are SCCs but not BSCCs (because of the probability leaks). All leaves are BSCCs.

Limit behavior of MCs

Fundamental theorem

Theorem

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ and $s \in S$. Then,

$$\mathbb{P}_s(\{\pi \in Paths(s) \mid \text{inf}(\pi) \text{ is a BSCC of } \mathcal{M}\}) = 1.$$

Recall that inf(π) is the set of states visited infinitely often along π.

$\Longrightarrow$ **We end up in a BSCC with probability one.**

Important consequence: if we are interested in the long-run behavior of the MC (e.g., prefix-independent properties like $\Box\Diamond T$), then it suffices to check which BSCCs are reached with positive probability and what happens in them.

Limit behavior of MCs

Application to classical events

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be an MC, $s \in S$ and $T \subseteq S$.

- **Infinitely often.** Repeated reachability can be reduced to reachability of good BSCCs:

$$\mathbb{P}_s(\Box\Diamond T) = \mathbb{P}_s(\Diamond U)$$

  where U is the union of all BSCCs B in $\mathcal{M}$ such that $B \cap T \neq \emptyset$.

- **Persistence.** Same idea:

$$\mathbb{P}_s(\Diamond\Box T) = \mathbb{P}_s(\Diamond U)$$

  where U is the union of all BSCCs B in $\mathcal{M}$ such that $B \subseteq T$.

$\Longrightarrow$ **Blackboard example for $\Box\Diamond T$.**

Limit behavior of MCs

Example: summary

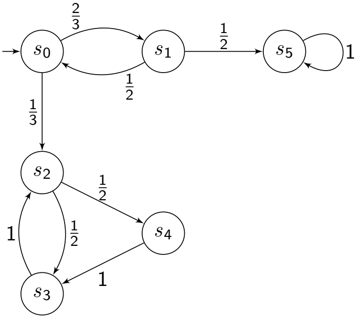

- $\mathbb{P}^{\mathcal{M}}(\Box\Diamond T)$ for $T = \{s_1, s_4\}$?
- BSCCs:
  - ▷ $B_1 = \{s_2, s_3, s_4\}$ (good, $B_1 \cap T \neq \emptyset$),
  - ▷ $B_2 = \{s_5\}$ (bad, $B_2 \cap T = \emptyset$).
- Hence, $\mathbb{P}^{\mathcal{M}}(\Box\Diamond T) = \mathbb{P}^{\mathcal{M}}(\Diamond s_2)$.

Applying the same approach as before, we have:

- $S_{=0} = \{s_5\}$, $S_{=1} = \{s_2, s_3, s_4\}$ and $S_? = \{s_0, s_1\}$.
- Solving $\mathsf{x} = \mathsf{A}\mathsf{x} + \mathsf{b}$ yields $\mathsf{x}_{s_0} = \frac{1}{2}$ hence $\mathbb{P}^{\mathcal{M}}(\Box\Diamond T) = \frac{1}{2}$.

Limit behavior of MCs

Steady-state distribution of a BSCC

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be an MC such that S is a BSCC. E.g., the lossy communication protocol.

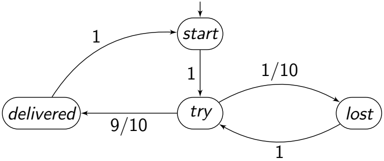

We can compute its **steady-state (or stationary) distribution**: the expected portion of time spent in each state in the long-run.

Steady-state distribution

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be an MC such that S is a BSCC. Then, there exists a unique stochastic vector v (i.e., $\mathrm{v} \in [0,1]^S$, $\sum_i v_i = 1$) such that $\mathrm{v}\mathbf{P} = \mathrm{v}$. This vector is the steady-state distribution.

Limit behavior of MCs

Steady-state distribution: example

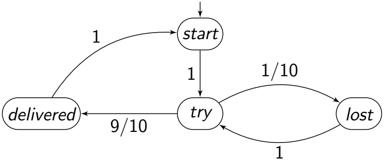

Consider the order $\{start, try, lost, delivered\}$. We are looking for a probabilistic vector v such that:

$$\begin{pmatrix} v_s & v_t & v_l & v_d \end{pmatrix} \cdot \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & \frac{1}{10} & \frac{9}{10} \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} v_s & v_t & v_l & v_d \end{pmatrix}.$$

Using $v_s + v_t + v_l + v_d = 1$, we obtain $\mathrm{v} = \begin{pmatrix} \frac{9}{29} & \frac{10}{29} & \frac{1}{29} & \frac{9}{29} \end{pmatrix}$.

Limit behavior of MCs

Steady-state distribution: an unusual application

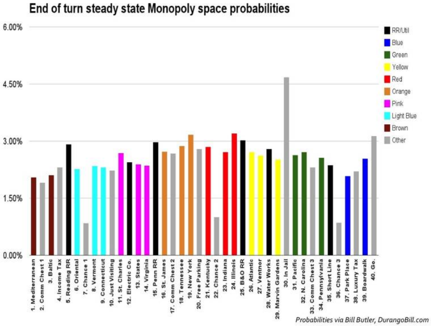
End of turn steady state Monopoly space probabilities

Probabilities via Bill Butler, DurangoBill.com

Under mild hypotheses, the Monopoly boardgame can be seen as a Markov chain consisting of a unique BSCC.

$\Longrightarrow$ **Studies have shown which squares are the most commonly visited.**

Limit behavior of MCs

Steady-state distribution: an unusual application

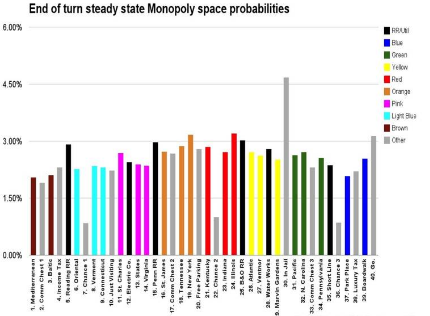

Probabilities via Bill Butler, DurangoBill.com

- ▷ After jail, Illinois Avenue is the most visited square with more than 3% of the total time (whereas a fair board would have all squares at 2.5%).
- ▷ Most cost-efficient squares: orange squares.

1. Markov chains
2. Reachability and limit behavior
3. PCTL: probabilistic CTL
4. Weighted Markov chains: venturing into the land of quantitative specifications

What is probabilistic CTL?

- PCTL is a *branching-time temporal logic* to express *properties of states* in an MC.
- Essentially, a CTL-like logic for probabilistic systems.

Main difference

CTL

Paths quantified using $\forall$ and $\exists$.

PCTL

Paths probability quantified using $\mathcal{P}_J(\phi)$ where $J \subseteq [0, 1]$ and ϕ is a path formula.

$\Longrightarrow$ Intuitively, $s \models \mathcal{P}_J(\phi)$ iff $\mathbb{P}_s(\{\pi \in Paths(s) \mid \pi \models \phi\}) \in J$.

$\Longrightarrow$ PCTL additionally includes the bounded until $\mathsf{U}^{\leq n}$ introduced before.

PCTL syntax

Core syntax

PCTL syntax

Given the set of atomic propositions AP, PCTL *state formulae* are formed according to the following grammar:

$$\Phi ::= \mathsf{true} \mid a \mid \Phi \wedge \Psi \mid \neg \Phi \mid \mathcal{P}_J(\phi)$$

where $a \in AP$, $J \subseteq [0, 1]$ is an interval with rational bounds, and ϕ is a path formula. PCTL *path formulae* are formed according to the following grammar:

$$\phi ::= \bigcirc \Phi \mid \Phi \,\mathsf{U}\, \Psi \mid \Phi \,\mathsf{U}^{\leq n} \Psi$$

where Φ and Ψ are state formulae and $n \in \mathbb{N}$.

$\Longrightarrow$ As for quantifiers in CTL, the syntax of PCTL enforces the presence of the probability operator $\mathcal{P}_J$ before every temporal operator.

PCTL syntax

Core syntax

PCTL syntax

Given the set of atomic propositions AP, PCTL *state formulae* are formed according to the following grammar:

$$\Phi ::= \mathsf{true} \mid a \mid \Phi \wedge \Psi \mid \neg\Phi \mid \mathcal{P}_J(\phi)$$

where $a \in AP$, $J \subseteq [0, 1]$ is an interval with rational bounds, and ϕ is a path formula. PCTL *path formulae* are formed according to the following grammar:

$$\phi ::= \bigcirc\, \Phi \mid \Phi \,\mathsf{U}\, \Psi \mid \Phi \,\mathsf{U}^{\leq n} \Psi$$

where Φ and Ψ are state formulae and $n \in \mathbb{N}$.

⚠ **Notations:** in the book, notations $\mathbb{P}$ for probability and $\mathcal{P}_J$ for the PCTL operator are replaced by Pr and $\mathbb{P}_J$ respectively.

PCTL syntax

Derived operators

As usual, other operators can be derived from this core syntax:

- Boolean connectives ($\vee$, $\rightarrow$, etc) are derived in the usual way,
- $\Diamond\Phi \equiv \mathsf{true}\,\mathsf{U}\,\Phi$, $\Diamond^{\leq n}\Phi \equiv \mathsf{true}\,\mathsf{U}^{\leq n}\Phi$,
- the "always" is obtained using the duality of the events: e.g., $\mathcal{P}_{\leq p}(\Box\Phi) = \mathcal{P}_{\geq 1-p}(\Diamond\neg\Phi)$.

Operators W and R can be obtained similarly.

PCTL: examples

Knuth's die

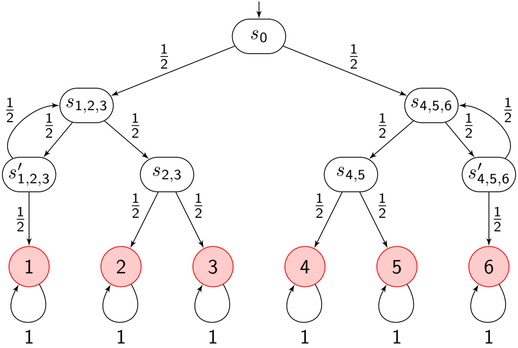

We express that all numbers should have probability 1/6 in PCTL:

$$\Phi = \bigwedge_{1 \leq i \leq 6} \mathcal{P}_{=\frac{1}{6}}(\Diamond i).$$

This PCTL formula should hold in s_0, and we proved that it does.

PCTL: examples

Lossy communication protocol

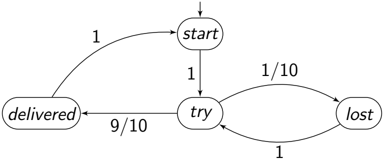

The PCTL formula

$$\Phi = \mathcal{P}_{=1}(\Diamond \mathit{delivered}) \land \mathcal{P}_{=1}\Big(\Box\big(\mathit{try} \rightarrow \mathcal{P}_{\geq 0.99}(\Diamond^{\leq 3} \mathit{delivered})\big)\Big)$$

expresses that

- with probability one, at least one message will be delivered (first conjunct),
- with probability one, every attempt to send a message results in the message being delivered within 3 steps with probability 0.99 (second conjunct).

PCTL semantics

For state formulae

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be an MC, $a \in AP$, $s \in S$, Φ and Ψ be PCTL state formulae and ϕ be a PCTL path formula.

Satisfaction for state formulae

$s \models \Phi$ iff formula Φ holds in state s.

$$
\begin{array}{lll}
s \models \text{true} & & \\
s \models a & \text{iff} & a \in L(s) \\
s \models \Phi \wedge \Psi & \text{iff} & s \models \Phi \text{ and } s \models \Psi \\
s \models \neg\Phi & \text{iff} & s \not\models \Phi \\
s \models \mathcal{P}_J(\phi) & \text{iff} & \mathbb{P}(s \models \phi) \in J
\end{array}
$$

where $\mathbb{P}(s \models \phi) = \mathbb{P}_s(\{\pi \in Paths(s) \mid \pi \models \phi\})$.

PCTL semantics

For path formulae

Let $\pi = s_0 s_1 s_2 \ldots$.

Satisfaction for path formulae

$\pi \models \phi$ iff path π satisfies ϕ.

$$\begin{array}{lll} \pi \models \bigcirc\Phi & \text{iff} & s_1 \models \Phi \\ \pi \models \Phi \,\mathsf{U}\, \Psi & \text{iff} & \exists j \geq 0,\ s_j \models \Psi \text{ and } \forall 0 \leq i < j,\ s_i \models \Phi \\ \pi \models \Phi \,\mathsf{U}^{\leq n} \Psi & \text{iff} & \exists 0 \leq j \leq n,\ s_j \models \Psi \text{ and } \forall 0 \leq i < j,\ s_i \models \Phi \end{array}$$

PCTL semantics

For Markov chains (1/2)

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$ be an MC and Φ a PCTL state formula over AP.

Definition: satisfaction set

The **satisfaction set** $Sat_{\mathcal{M}}(\Phi)$ (or briefly, $Sat(\Phi)$) for formula Φ is

$$Sat(\Phi) = \{s \in S \mid s \models \Phi\}.$$

The classical formulation of the PCTL model checking problem is to check whether a given state s belongs to $Sat(\Phi)$ or not.

$\Longrightarrow$ **What about satisfaction for an MC?**

PCTL semantics

For Markov chains (2/2)

Remark: any MC $\mathcal{M}$ with $|Supp(\iota_{\text{init}})| > 1$ can be equivalently presented as an MC $\mathcal{M}'$ with one additional state s_{init} such that $\iota'_{\text{init}}(s_{\text{init}}) = 1$ and $\mathbf{P}'(s_{\text{init}}, s) = \iota_{\text{init}}(s)$ for any state s of $\mathcal{M}$.

Let Φ be a PCTL formula and Φ' be the same formula where bounded until properties and nexts are shifted by one step (because of the additional initial transition). We easily define *satisfaction of PCTL formula Φ for the MC $\mathcal{M}$* as

$$\mathcal{M} \models \Phi \iff \mathcal{M}' \models \Phi' \iff s_{\text{init}} \models \Phi'.$$

PCTL semantics

The $\mathcal{P}$ operator

We have seen that $s \models \mathcal{P}_J(\phi)$ iff $\mathbb{P}_s(\{\pi \in Paths(s) \mid \pi \models \phi\}) \in J$.

Potential problem?

Recall it only makes sense if the considered event is *measurable*.

$\Longrightarrow$ **Are all sets defined by PCTL path formulae measurable?**

$\Longrightarrow$ **Fortunately, yes. It can be proved that they are using an approach similar to what we did for $\lozenge T$.**

$\Longrightarrow$ **See the book.**

PCTL model checking

Decision problem

Definition: PCTL model checking problem

Given an MC $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L)$, a state $s \in S$ and a PCTL state formula Φ, decide if $s \models \Phi$ or not.

Sketch of the algorithm

- Same skeleton as for CTL: *recursive computation of $Sat(\Phi)$ via bottom-up traversal of the parse tree of Φ*.
- What is new: **how to deal with subformulae $\Psi = \mathcal{P}_J(\phi)$?**
  - $Sat(\mathcal{P}_J(\phi)) = \{s \in S \mid \mathbb{P}(s \models \phi) \in J\}$.
  - Hence **we need to compute $\mathbb{P}(s \models \phi)$ for $s \in S$**.

$\implies$ **If we learn how to do this, we are done: we already know the rest of the algorithm.**

PCTL model checking

Computing $\mathbb{P}(s \models \phi)$ (1/2)

We have three possible path formulae to consider: $\phi = \bigcirc \Phi$, $\phi = \Phi \,\mathsf{U}\, \Psi$ and $\phi = \Phi \,\mathsf{U}^{\leq n} \Psi$. All other ones can be derived from the core syntax.

1. Let $\phi = \bigcirc \Phi$. Then we simply have:

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

by definition of the transition probability function $\mathbf{P}$ in $\mathcal{M}$.

$\Longrightarrow$ **Easily achieved by a single matrix-vector multiplication (see slide 10).**

PCTL model checking

Computing $\mathbb{P}(s \models \phi)$ (2/2)

2 Let $\phi = \Phi \, \mathsf{U} \, \Psi$. Then:

$$\mathbb{P}(s \models \Phi \, \mathsf{U} \, \Psi) = \mathbb{P}(s \models C \, \mathsf{U} \, T)$$

for $C = \mathit{Sat}(\Phi)$ and $T = \mathit{Sat}(\Psi)$.

$\Longrightarrow$ **We saw how to compute this through a linear equation system (which can be done in polynomial time).**

3 Let $\phi = \Phi \, \mathsf{U}^{\leq n} \Psi$. Then:

$$\mathbb{P}(s \models \Phi \, \mathsf{U}^{\leq n} \Psi) = \mathbb{P}(s \models C \, \mathsf{U}^{\leq n} T)$$

for $C = \mathit{Sat}(\Phi)$ and $T = \mathit{Sat}(\Psi)$.

$\Longrightarrow$ **We saw how to compute this via the iterative approach: it requires $\mathcal{O}(n)$ matrix-vector multiplications.**

PCTL model checking

Complexity

Complexity of the PCTL model checking algorithm

The time complexity for an MC $\mathcal{M}$ and a PCTL formula Φ is $\mathcal{O}(\mathit{poly}(|\mathcal{M}|) \cdot n_{\max} \cdot |\Phi|)$, where $n_{\max}$ is the maximal step bound appearing in a subformula of Φ or $n_{\max} = 1$ if Φ contains no bounded until operator.

$\Longrightarrow$ **Polynomial (⚠$n_{\max}$) time, as for CTL model checking.**

Remark: qualitative PCTL

For *qualitative* PCTL properties (i.e., $\mathcal{P}_{=1}$ or $\mathcal{P}_{>0}$), more efficient algorithms exist: **graph-based techniques suffice** (as the actual values of the probabilities do not matter).

PCTL vs. CTL

Recall that CTL gives us quantifiers $\forall$ and $\exists$ whereas PCTL gives us operator $\mathcal{P}_J$.

$\Longrightarrow$ **Can we compare their expressiveness?**

E.g., is $s \models \mathcal{P}_{=1}(\phi) \iff s \models \forall\phi$? Is $s \models \mathcal{P}_{>0}(\phi) \iff s \models \exists\phi$? For any path formula ϕ? For some of them?

PCTL vs. CTL

Example

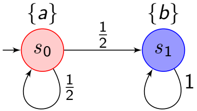

Here, we have that:

- $s_0 \models \mathcal{P}_{=1}(\Diamond b)$ but $s_0 \not\models \forall\Diamond b$,
- $s_0 \models \exists\Box a$ but $s_0 \not\models \mathcal{P}_{>0}(\Box a)$.

Remark: sure vs. almost-sure properties

We often say that a property satisfied for all paths is **sure** whereas a property satisfied with probability one is **almost-sure**.

PCTL vs. CTL

In full generality

Non-exhaustive list of relations:

$$s \models \mathcal{P}_{=1}(\Diamond\Phi) \;\underset{\Longleftarrow}{\overset{\not\Longrightarrow}{}}\; s \models \forall\Diamond\Phi$$

$$s \models \mathcal{P}_{>0}(\Diamond\Phi) \iff s \models \exists\Diamond\Phi$$

$$s \models \mathcal{P}_{=1}(\bigcirc\,\Phi) \iff s \models \forall\bigcirc\,\Phi$$

$$s \models \mathcal{P}_{>0}(\bigcirc\,\Phi) \iff s \models \exists\bigcirc\,\Phi$$

$$s \models \mathcal{P}_{=1}(\Box\Phi) \iff s \models \forall\Box\Phi$$

$$s \models \mathcal{P}_{>0}(\Box\Phi) \;\underset{\not\Longleftarrow}{\overset{\Longrightarrow}{}}\; s \models \exists\Box\Phi$$

Expressiveness

PCTL and CTL are incomparable.

What can we define in PCTL?

Two examples

Repeated reachability ("infinitely often"):

$$s \models \underbrace{\mathcal{P}_J(\Diamond\mathcal{P}_{=1}(\Box\mathcal{P}_{=1}(\Diamond a)))}_{\mathcal{P}_J(\Box\Diamond a)} \iff \mathbb{P}(s \models \Box\Diamond a) \in J.$$

$\Longrightarrow$ *The formula essentially states that we have probability within J to reach a BSCC B such that $B \cap Sat(a) \neq \emptyset$.*

Persistence:

$$s \models \mathcal{P}_J(\Diamond\mathcal{P}_{=1}(\Box a)) \iff \mathbb{P}(s \models \Diamond\Box a) \in J.$$

$\Longrightarrow$ *The formula essentially states that we have probability within J to reach a BSCC B such that $B \subseteq Sat(a)$.*

Understanding a PCTL formula: example

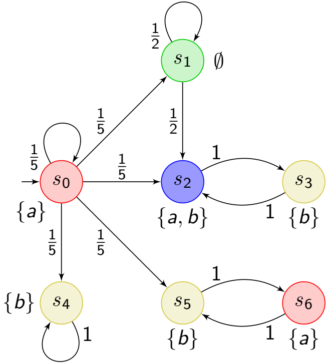

Seems too complex?

- ▷ Reach BSCC B s.t. $B \cap \mathit{Sat}(a) \neq \emptyset \ldots$
- ▷ and $B \subseteq \mathit{Sat}(b) \ldots$

$\Longrightarrow$ Only $\{s_2, s_3\}$ is fine.

- ▷ in at most 3 steps...
- ▷ following a path in $\mathit{Sat}(a) \ldots$

$\Longrightarrow$ Visiting s_1 is not allowed.

- ▷ with probability $\geq 2/9$.

Consider checking the PCTL formula Φ for s_0: $\boxed{\Longrightarrow \textbf{Yes}, s_0 \models \Phi.}$

$$\Phi = \mathcal{P}_{\geq \frac{2}{9}}\Big(a \,\mathsf{U}^{\leq 3} \big(\mathcal{P}_{=1}(\Box(\mathcal{P}_{=1}(\Diamond a))) \wedge \mathcal{P}_{=1}(\Box b)\big)\Big).$$

Thus $\Phi \equiv \mathcal{P}_{\geq \frac{2}{9}}(s_0 \,\mathsf{U}^{\leq 3} s_2)$. $\mathbb{P}_{s_0}(s_0 \,\mathsf{U}^{\leq 3} s_2) = \frac{1}{5} + \frac{1}{5^2} + \frac{1}{5^3} = \frac{31}{125} > \frac{2}{9}$.

Beyond PCTL

For classical TSs, we saw that several logics exist beyond CTL, including LTL and CTL*.

For MCs also, several other formalisms exist.

- Probabilities of **linear-time properties** can be computed using an approach similar to LTL model checking:
  1. Represent the complement LT property through an automaton $\mathcal{A}$ (here a deterministic Rabin automaton).
  2. Compute the product MC $\mathcal{M} \otimes \mathcal{A}$.
  3. Check a reachability/persistence property on the product.
- The logic **PCTL* extends PCTL in the same way as CTL* extends CTL**: by allowing LTL formulae as path formulae.
  - ▷ Just as CTL/CTL* properties are preserved by bisimulation, PCTL/PCTL* properties are preserved by *probabilistic bisimulation*, the adaptation of the notion to MCs.

1 Markov chains

2 Reachability and limit behavior

3 PCTL: probabilistic CTL

4 Weighted Markov chains: venturing into the land of quantitative specifications

Quantitative specifications

Usefulness

As discussed in Ch. 1, in practical applications, it is often necessary to consider the **performance of a system**. E.g.,

- reaching a target state *using a minimal amount of energy*,
- *minimizing the average response time* of a request-response system.

$\Longrightarrow$ To reason about such quantities, we need to **enrich the classical models of TSs and MCs with weights** representing quantitative changes (e.g., time taken by an action, consumed energy).

$\Longrightarrow$ **We need specific techniques for each type of quantitative property we want to model.**

Quantitative specifications

A quick glance

The theory of quantitative specifications is huge. We only illustrate two particular cases in the context of MCs:

1. **Shortest path** (or cost-bounded reachability).
   - ▷ Each transition has a cost and we want to consider the *cost-to-target* (i.e., sum of the costs up to reaching the target).
2. **Mean-payoff** (or long-run average).
   - ▷ Each transition has a reward and we want to consider the *average reward per transition in the long-run*.

For both settings, we consider two problems:

1. Computing the **expected value** of the quantitative property for an MC.
2. Computing the **probability** to obtain a value within a given interval.

Weighted Markov chain

Definition: weighted Markov chain (WMC)

A weighted Markov chain (WMC) is a tuple $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L, w)$ where

- S, $\mathbf{P}$, ι_{init}, AP and L are defined as for traditional MCs,
- $w \colon S \times S \to \mathbb{Z}$ is a (partial) *weight function* assigning an integer weight to each transition (s, s') such that $\mathbf{P}(s, s') > 0$.

Illustration: weights appear besides probabilities on transitions.

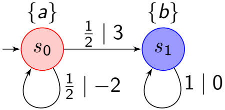

Remark

In the book weights are on *states*. Both formalisms are equivalent.

Shortest path

The setting

Idea: generalization of the graph problem to MCs to model probabilistic aspects of real-life systems, e.g., traffic, accidents. . .

Restriction

We consider only *non-negative weights*, i.e., $w \colon S \times S \to \mathbb{N}$.

Example: lossy communication protocol.

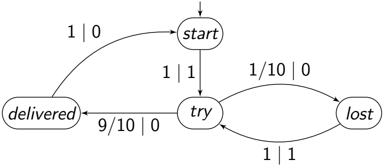

$\Longrightarrow$ We put 1 on transitions entering *try* as we want to reason on the number of tries needed before reaching *delivered*.

Shortest path

Cost-to-target

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L, w)$ be a WMC and $T \subseteq S$ be the set to reach. We introduce the **truncated sum** payoff function that assigns the *cumulative cost to target* to paths of the MC.

Definition: truncated sum

The truncated sum up to T is a function $\mathsf{TS}^T : \mathit{Paths}(\mathcal{M}) \rightarrow \mathbb{N} \cup \{\infty\}$ whose values are given by

$$\mathsf{TS}^T(\pi) = \begin{cases} \sum_{i=0}^{n-1} w(s_i, s_{i+1}) \text{ if } (\forall\, 0 \leq i < n,\ s_i \notin T) \wedge s_n \in T \\ \infty \text{ if } \pi \not\models \Diamond T \end{cases}$$

where $\pi = s_0 s_1 \ldots \in \mathit{Paths}(\mathcal{M})$.

Shortest path

Cost-to-target: example

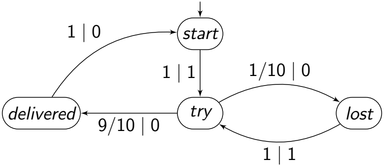

For $T = \{\textit{delivered}\}$, we have:

- $\mathsf{TS}^T((\textit{start} \cdot \textit{try} \cdot \textit{delivered})^\omega) = 1 + 0 = 1$,
- $\mathsf{TS}^T((\textit{start} \cdot \textit{try} \cdot \textit{lost} \cdot \textit{try} \cdot \textit{delivered})^\omega) = 1 + 0 + 1 + 0 = 2$,
- $\mathsf{TS}^T(\textit{start} \cdot (\textit{try} \cdot \textit{lost})^\omega) = \infty$ because T is never reached.

$\Longrightarrow$ **First interesting question: what is the expected cost-to-target, i.e., the average number of tries before a message is delivered?**

Shortest path

Expected cost-to-target

Expected cost-to-target

For $s \in S$ and $T \subseteq S$, the *expected cost-to-target* $\mathbb{E}_s(\mathsf{TS}^T)$ is obtained as follows:

- if $\mathbb{P}_s(\Diamond T) < 1$, then $\mathbb{E}_s(\mathsf{TS}^T) = \infty$;
- otherwise,

$$\mathbb{E}_s(\mathsf{TS}^T) = \sum_{r=0}^{\infty} r \cdot \mathbb{P}_s(\{\pi \in Paths(s) \mid \mathsf{TS}^T(\pi) = r\}).$$

$\Longrightarrow$ *Coincides with the intuition of "average cost-to-target".*

The second equality can be equivalently written as

$$\mathbb{E}_s(\mathsf{TS}^T) = \textstyle\sum_{\widehat{\pi}} \mathbf{P}(\widehat{\pi}) \cdot \mathsf{TS}^T(\widehat{\pi})$$

for $\widehat{\pi} \in \{s_0 \dots s_n \in Paths_{fin}(s) \mid (\forall\, 0 \leq i < n,\ s_i \notin T) \wedge s_n \in T\}$.

Shortest path

Expected cost-to-target: illustration

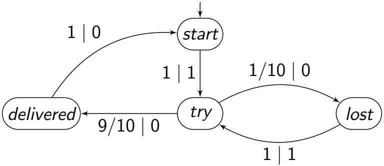

Applying the definition for $T = \{\textit{delivered}\}$, we obtain:

$$\mathbb{E}_s(\mathsf{TS}^T) = \frac{9}{10}\cdot 1 + \frac{9}{100}\cdot 2 + \frac{9}{1000}\cdot 3 + \frac{9}{10000}\cdot 4 + \ldots$$
$$= \frac{9}{10}\cdot\sum_{r=1}^{\infty} r\cdot\left(\frac{1}{10}\right)^{r-1} = \frac{9}{10}\cdot\frac{1}{(1-\frac{1}{10})^2} = \frac{9}{10}\cdot\left(\frac{10}{9}\right)^2 = \frac{10}{9}.$$

$\Longrightarrow$ **On average, the message is delivered after $10/9$ tries.**

Shortest path

Expected cost-to-target: simpler approach

Based on the technique used for constrained reachability, we can also use a *linear equation system*.

Linear system for expected cost-to-target

Let $S_{=1} = \{s \in S \mid \mathbb{P}_s(\Diamond T) = 1\}$. Values $x_s = \mathbb{E}_s(\mathsf{TS}^T)$ form the unique solution to the following system:

$$x_s = \begin{cases} 0 & \text{if } s \in T \\ \sum_{s' \in Post(s)} \mathbf{P}(s, s') \cdot (w(s, s') + x_{s'}) & \text{if } s \in S_{=1} \setminus T \\ \infty & \text{otherwise.} \end{cases}$$

$\Longrightarrow$ *The total expected cost in a state can be split up into the cost of the next transition + the expected total cost from the next state, both subject to the probability distribution over successors.*

Shortest path

Expected cost-to-target: revisited illustration

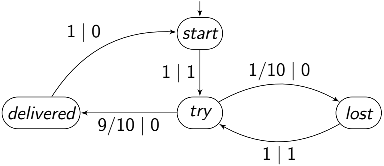

For $T = \{\mathit{delivered}\}$, with the linear system approach, we have:

$$\begin{cases} x_s &= 1 + x_t \\ x_t &= \frac{1}{10} \cdot x_l + \frac{9}{10} \cdot x_d \\ x_l &= 1 + x_t \\ x_d &= 0 \end{cases} \implies \begin{cases} x_s &= \frac{10}{9} \\ x_t &= \frac{1}{9} \\ x_l &= \frac{10}{9} \\ x_d &= 0 \end{cases}$$

$\implies$ **We obtain $\mathbb{E}_s(\mathsf{TS}^T) = 10/9$ as expected.**

Shortest path

Expected cost-to-target: complexity

Complexity

Given a WMC $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L, w)$, $s \in S$ and $T \subseteq S$, computing the expected cost-to-target $\mathbb{E}_s(\mathsf{TS}^T)$ takes polynomial time in $|\mathcal{M}|$.

Shortest path

Cost-bounded reachability probability

Different problem: fix a bound $b \in \mathbb{N}$ and compute the **probability** to reach T with cost $\leq b$.

Cost-bounded reachability (CBR) probability

For $s \in S$, $T \subseteq S$, the *CBR probability for bound $b \in \mathbb{N}$* is $\mathbb{P}_s(\mathsf{TS}^T \leq b) = \mathbb{P}_s(\{\pi \in Paths(s) \mid \mathsf{TS}^T(\pi) \leq b\})$.

$\Longrightarrow$ Several formulations of the solution exist. In the next slide, we present one based on a **reduction to computing a simple reachability probability on a unfolded MC**.

Key idea

We are only interested in paths π reaching T with $\mathsf{TS}^T(\pi) \leq b$ $\Longrightarrow$ *anything that happens once the cumulative cost is $> b$ is useless* (recall that weights are non-negative).

Shortest path

Cost-bounded reachability probability: reduction to reachability

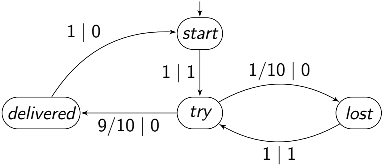

To compute $\mathbb{P}_s(\mathsf{TS}^T \leq b)$ for $T = \{\mathit{delivered}\}$ and $b = 2$, we *unfold this MC up to the bound*, integrating the current sum in the new states, and we stop a branch as soon as (i) T is reached, or (ii) the sum exceeds b.

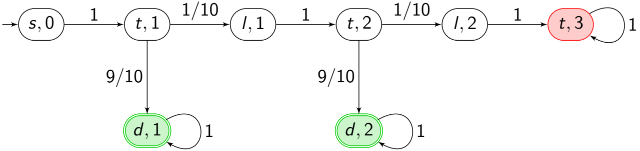

Shortest path

Cost-bounded reachability probability: reduction to reachability

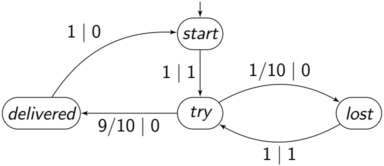

Let $\mathcal{M}$ be the original WMC, and $\mathcal{M}_b$ the unfolded unweighted MC. We have a relation between paths π in $\mathcal{M}$ and π' in $\mathcal{M}_b$ and

$$\mathsf{TS}^T(\pi) \leq b \iff \pi' \models \Diamond T' \text{ where } T' = T \times \{0, 1, \ldots, b\}.$$

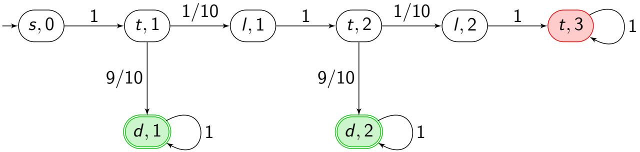

Shortest path

Cost-bounded reachability probability: reduction to reachability

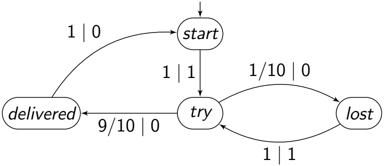

Hence $\mathbb{P}_s(\mathsf{TS}^T \leq b) = \mathbb{P}_{(s,0)}(\Diamond T')$, which we can compute (e.g., using the classical linear equation system in $\mathcal{M}_b$) to obtain $\mathbb{P}_s(\mathsf{TS}^T \leq b) = 9/10 + 9/100 = 99/100$ as naturally expected.

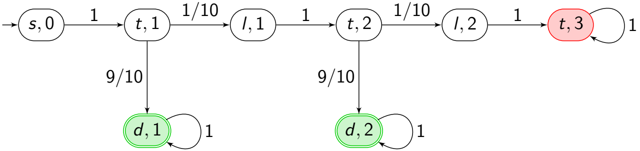

Shortest path

Cost-bounded reachability probability: complexity

Complexity of the algorithm

Given a WMC $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L, w)$, $s \in S$, $T \subseteq S$ and $b \in \mathbb{N}$, computing the CBR probability $\mathbb{P}_s(\mathsf{TS}^T \leq b)$ takes polynomial time in $|\mathcal{M}_b|$, hence pseudo-polynomial time in $|\mathcal{M}|$.

$\Longrightarrow$ **With regard to the binary encoding of the problem, the time needed can be exponential!**

$\Longrightarrow$ **The exponential blow-up cannot be avoided!**

Hardness

The decision problem associated to the CBR probability, i.e., deciding whether $\mathbb{P}_s(\mathsf{TS}^T \leq b)$ exceeds a given probability or not, is in PSPACE and PosSLP-hard [HK15], which is higher than NP-hard.

Shortest path

Additional remarks

- Computing the expected cost-to-target is easier than computing the CBR probability: P vs. PSPACE-easy and NP-hard.
- Both quantities can be used in a quantitative extension of PCTL called Probabilistic Reward CTL (PRCTL).

Mean-payoff

The setting

Idea: quantifying the average reward/cost per transition in the long run, e.g., energy consumption per action, response time...

Unrestricted weights

We accept *both positive and negative weights*, i.e., $w\colon S \times S \to \mathbb{Z}$.

Example: we want to characterize the average energy consumption per transition in the long-run.

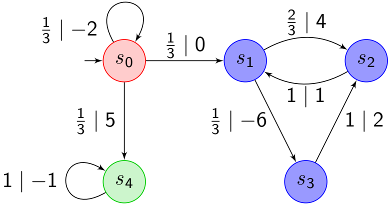

Mean-payoff

Definition of the payoff

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L, w)$ be a WMC. The **mean-payoff** function assigns the *long-run average weight* to paths of the WMC.

Definition: mean-payoff

The mean-payoff is a function $\mathsf{MP}\colon \mathit{Paths}(\mathcal{M}) \to \mathbb{R}$ whose values are given by

$$\mathsf{MP}(\pi) = \liminf_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} w(s_i, s_{i+1})$$

where $\pi = s_0 s_1 \ldots \in \mathit{Paths}(\mathcal{M})$.

Mean-payoff

Example

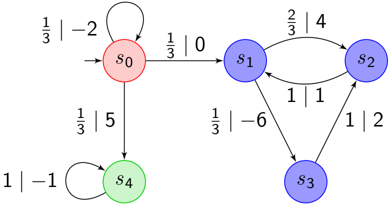

We have:

- $\mathsf{MP}((s_0)^\omega) = \liminf_{n\to\infty} \dfrac{1}{n} \cdot (-2n) = -2.$
- $\mathsf{MP}((s_0)^3(s_4)^\omega) = \liminf_{n\to\infty} \dfrac{2\cdot(-2) + 5 + (n-3)\cdot(-1)}{n} = -1.$

  $\Longrightarrow$ **Mean-payoff is prefix-independent:** $\mathsf{MP}(\pi) = \mathsf{MP}(\pi')$ **for any suffix** π' **of** π**.**
- $\mathsf{MP}(s_0(s_1 s_2)^\omega) = \liminf_{n\to\infty} \dfrac{\frac{n}{2}\cdot 4 + \frac{n}{2}\cdot 1}{n+1} = 2.5.$

  $\Longrightarrow$ **Average of the cycle.**

Mean-payoff

BSCCs

As for the shortest path, we want to consider both the *expected mean-payoff* and the *probability of achieving a given bound*.

$\Longrightarrow$ **We first consider BSCCs where an important result links both quantities.**

Theorem

Let $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L, w)$ be a WMC such that S is a BSCC. Then, there exists a value $\nu \in \mathbb{Q}$ such that for all $s \in S$,

1. $\mathbb{E}_s(\mathsf{MP}) = \nu$, and
2. $\mathbb{P}_s(\mathsf{MP} = \nu) = 1$.

$\Longrightarrow$ **Key result: in a BSCC, the expected mean-payoff is the same in all states and it is achieved almost-surely. It follows from definition of BSCCs and prefix independence of the mean-payoff.**

Mean-payoff

Computing the expected mean-payoff in BSCCs

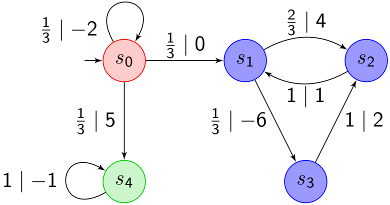

For BSCC $B_1 = \{s_4\}$, we trivially have that $\mathbb{E}_{B_1}(\mathsf{MP}) = -1$. What about $B_2 = \{s_1, s_2, s_3\}$?

Intuitively, we are interested in the "average behavior" of the BSCC in the long-run... which is described by its *steady-state distribution*.

Mean-payoff

Computing the expected mean-payoff in BSCCs

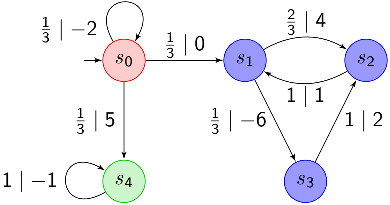

For BSCC $B_1 = \{s_4\}$, we trivially have that $\mathbb{E}_{B_1}(\mathsf{MP}) = -1$. What about $B_2 = \{s_1, s_2, s_3\}$?

Computing $\mathbb{E}_{B_2}(\mathsf{MP})$:

▷ Compute the *steady-state distribution* v s.t. $\mathsf{v}\mathbf{P} = \mathsf{v}$.

$$\begin{cases} \mathsf{v}_{s_1} = \mathsf{v}_{s_2} \\ \mathsf{v}_{s_2} = \frac{2}{3}\mathsf{v}_{s_1} + \mathsf{v}_{s_3} \\ \mathsf{v}_{s_3} = \frac{1}{3}\mathsf{v}_{s_1} \\ \mathsf{v}_{s_1} + \mathsf{v}_{s_2} + \mathsf{v}_{s_3} = 1 \end{cases} \implies \begin{cases} \mathsf{v}_{s_1} = \frac{3}{7} \\ \mathsf{v}_{s_2} = \frac{3}{7} \\ \mathsf{v}_{s_3} = \frac{1}{7} \end{cases}$$

Mean-payoff

Computing the expected mean-payoff in BSCCs

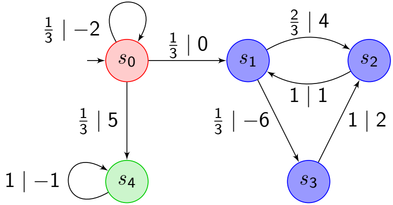

For BSCC $B_1 = \{s_4\}$, we trivially have that $\mathbb{E}_{B_1}(\mathsf{MP}) = -1$. What about $B_2 = \{s_1, s_2, s_3\}$?

Computing $\mathbb{E}_{B_2}(\mathsf{MP})$:

- ▷ *Steady-state distribution* $\mathsf{v} = (\frac{3}{7} \quad \frac{3}{7} \quad \frac{1}{7})$.
- ▷ Compute the *one-step expected reward column-vector* e.

$$\begin{cases} \mathsf{e}_{s_1} = \frac{2}{3} \cdot 4 + \frac{1}{3} \cdot (-6) \\ \mathsf{e}_{s_2} = 1 \\ \mathsf{e}_{s_3} = 2 \end{cases} \implies \begin{cases} \mathsf{e}_{s_1} = \frac{2}{3} \\ \mathsf{e}_{s_2} = 1 \\ \mathsf{e}_{s_3} = 2 \end{cases}$$

Mean-payoff

Computing the expected mean-payoff in BSCCs

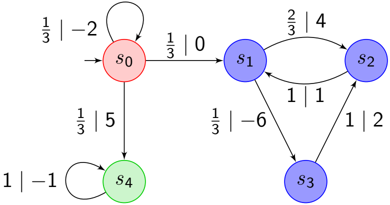

For BSCC $B_1 = \{s_4\}$, we trivially have that $\mathbb{E}_{B_1}(\mathsf{MP}) = -1$. What about $B_2 = \{s_1, s_2, s_3\}$?

Computing $\mathbb{E}_{B_2}(\mathsf{MP})$:

- *Steady-state distribution* $\mathsf{v} = (\frac{3}{7} \quad \frac{3}{7} \quad \frac{1}{7})$.
- *One-step expected reward column-vector* $\mathsf{e} = (\frac{2}{3} \quad 1 \quad 2)^T$.
- Finally, compute $\mathbb{E}_{B_2}(\mathsf{MP}) = \mathsf{v} \cdot \mathsf{e}$.

$$\mathbb{E}_{B_2}(\mathsf{MP}) = \frac{3}{7} \cdot \frac{2}{3} + \frac{3}{7} \cdot 1 + \frac{1}{7} \cdot 2 = 1.$$

Mean-payoff

Computing the expected mean-payoff in BSCCs

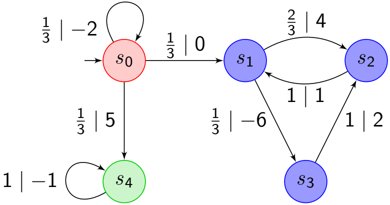

For BSCC $B_1 = \{s_4\}$, we trivially have that $\mathbb{E}_{B_1}(\mathsf{MP}) = -1$. What about $B_2 = \{s_1, s_2, s_3\}$?

Computing $\mathbb{E}_{B_2}(\mathsf{MP})$:

- ▷ *Steady-state distribution* $\mathsf{v} = (\frac{3}{7} \quad \frac{3}{7} \quad \frac{1}{7})$.
- ▷ *One-step expected reward column-vector* $\mathsf{e} = (\frac{2}{3} \quad 1 \quad 2)^T$.
- ▷ $\mathbb{E}_{B_2}(\mathsf{MP}) = \mathsf{v} \cdot \mathsf{e} = 1$.

$\Longrightarrow$ **We can do this for all BSCCs of any WMC.**

$\Longrightarrow$ **And by the last theorem, we also get that for all s in BSCC B, $\mathbb{P}_s(\mathsf{MP} = \mathbb{E}_B(\mathsf{MP})) = 1$.**

Mean-payoff

Computing the expected mean-payoff in BSCCs: complexity

Complexity

Given a WMC $\mathcal{M} = (S, \mathbf{P}, \iota_{\mathrm{init}}, AP, L, w)$ with BSCCs $B_1, \ldots, B_k$, the following properties hold:

- $1 \leq k \leq |S|$ (as BSCCs are disjoint by definition),
- computing the expected mean-payoff values $\mathbb{E}_{B_1}(\mathsf{MP}), \ldots, \mathbb{E}_{B_k}(\mathsf{MP})$ takes polynomial time in $|\mathcal{M}|$.

Mean-payoff

Dealing with general WMCs: expected mean-payoff

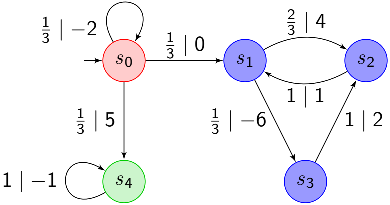

We know that $\mathbb{E}_{B_1}(\mathsf{MP}) = -1$ and $\mathbb{E}_{B_2}(\mathsf{MP}) = 1$ for $B_1 = \{s_4\}$ and $B_2 = \{s_1, s_2, s_3\}$.

$\Longrightarrow$ Can we compute $\mathbb{E}_{s_0}(\mathsf{MP})$?

Since the mean-payoff is *prefix-independent*, we only care about the long-run behavior and the long-run behavior almost-surely only happens in. . . BSCCs.

$\Longrightarrow$ **The global expected mean-payoff is simply the weighted average between all reachable BSCCs.**

Mean-payoff

Dealing with general WMCs: expected mean-payoff

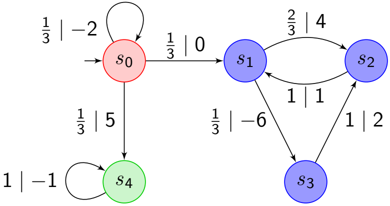

We know that $\mathbb{E}_{B_1}(\mathsf{MP}) = -1$ and $\mathbb{E}_{B_2}(\mathsf{MP}) = 1$ for $B_1 = \{s_4\}$ and $B_2 = \{s_1, s_2, s_3\}$.

$\Longrightarrow$ Can we compute $\mathbb{E}_{s_0}(\mathsf{MP})$?

Hence,

$$\begin{aligned}\mathbb{E}_{s_0}(\mathsf{MP}) &= \mathbb{P}_{s_0}(\Diamond B_1) \cdot \mathbb{E}_{B_1}(\mathsf{MP}) + \mathbb{P}_{s_0}(\Diamond B_2) \cdot \mathbb{E}_{B_2}(\mathsf{MP}) \\ &= \frac{1}{2} \cdot (-1) + \frac{1}{2} \cdot 1 = 0.\end{aligned}$$

$\Longrightarrow$ **The expected mean-payoff is zero for this WMC.**

Mean-payoff

Dealing with general WMCs: probability of achieving a given mean-payoff

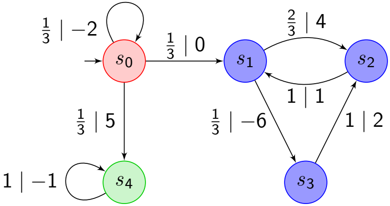

We know that $\mathbb{E}_{B_1}(\mathsf{MP}) = -1$ and $\mathbb{E}_{B_2}(\mathsf{MP}) = 1$ for $B_1 = \{s_4\}$ and $B_2 = \{s_1, s_2, s_3\}$.

$\Longrightarrow$ Can we compute the probability $\mathbb{P}_{s_0}(\mathsf{MP} \geq 0)$?

Using the same arguments, it suffices to compute

$$\mathbb{P}_{s_0}(\mathsf{MP} \geq 0) = \sum_{B_i \text{ s.t. } \mathbb{E}_{B_i}(\mathsf{MP}) \geq 0} \mathbb{P}_{s_0}(\Diamond B_i).$$

$\Longrightarrow$ **The probability of reaching a BSCC with an adequate expected mean-payoff.**

Mean-payoff

Dealing with general WMCs: probability of achieving a given mean-payoff

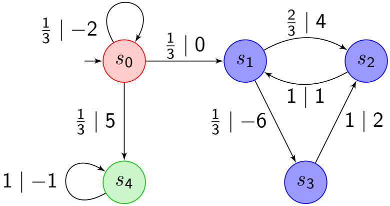

We know that $\mathbb{E}_{B_1}(\mathsf{MP}) = -1$ and $\mathbb{E}_{B_2}(\mathsf{MP}) = 1$ for $B_1 = \{s_4\}$ and $B_2 = \{s_1, s_2, s_3\}$.

$\implies$ Can we compute the probability $\mathbb{P}_{s_0}(\mathsf{MP} \geq 0)$?

Hence,

$$\mathbb{P}_{s_0}(\mathsf{MP} \geq 0) = \mathbb{P}_{s_0}(\Diamond B_2) = \frac{1}{2}.$$

$\implies$ **Mean-payoff ≥ 0 is obtained with probability $\frac{1}{2}$.**

Mean-payoff

Dealing with general WMCs: complexity

For both problems, we need to compute

1. the expected mean-payoff of BSCCs,
   $\hookrightarrow$ Takes polynomial time.
2. reachability probabilities toward BSCCs.
   $\hookrightarrow$ Takes polynomial time.

Complexity

Given a WMC $\mathcal{M} = (S, \mathbf{P}, \iota_{\text{init}}, AP, L, w)$, both computing its expected mean-payoff and computing the probability of paths with a mean-payoff greater than a given bound $b \in \mathbb{Q}$ requires polynomial time in $|\mathcal{M}|$.

Remark: those quantities can also be formalized in PRCTL.

Complexity wrap-up

	Shortest path	*Mean-payoff*
Expected value	P	P
Probability	PSPACE-easy/NP-hard	P

References I

C. Baier and J.-P. Katoen.
Principles of model checking.
MIT Press, 2008.

Jerzy Filar and Koos Vrieze.
Competitive Markov decision processes.
Springer, 1997.

Christoph Haase and Stefan Kiefer.
The odds of staying on budget.
In *Proc. of ICALP*, LNCS 9135, pages 234–246. Springer, 2015.