Reachability in Infinite Markov Chains

UMONS Université de Mons

1. Outline

- **Goal:** develop techniques to automatically assess the reliability of complex systems.
- **Problem at hand:** quantify the likelihood that some events happen in **stochastic** and **timed** environments.
- **Plan:** follow a successful approach to understand this problem for countable Markov chains [ABM07] and for general stochastic transition systems [BBBC18] and use it in the setting of *stochastic hybrid systems*.

2. Markov chains

A *Markov chain* is a tuple $\mathcal{M} = (S, \rightarrow, P)$ where

- S is a countable set of states,
- $\rightarrow \subseteq S \times S$ is a transition relation,
- $P : S \times S \rightarrow [0, 1]$ such that for all $s \in S$, $P(s, \cdot)$ is a probability distribution on the transitions from s.

A Markov chain can be used to describe sequences of states in which the probability of each state depends solely on the previous state.

3. Markov chain running example

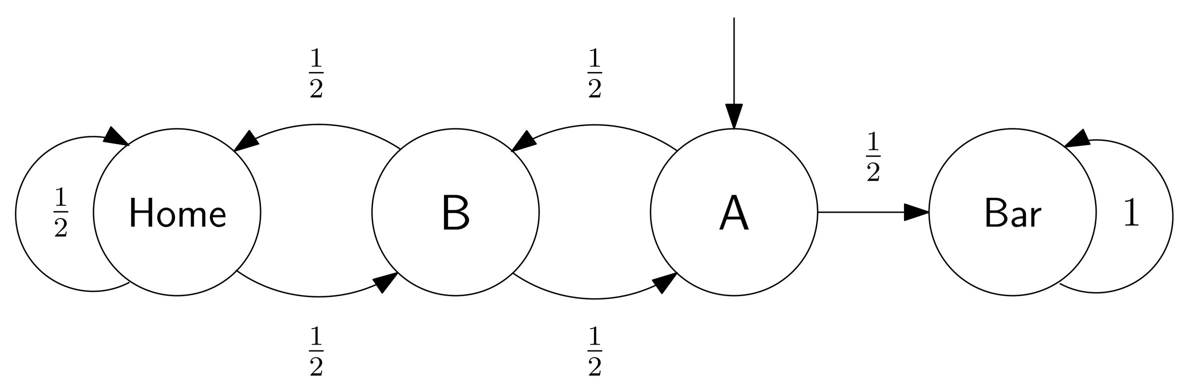

We model a situation in which a man starts in state A and then goes either to B or Bar with probability $\frac{1}{2}$ ($P(\mathsf{A}, \mathsf{B}) = P(\mathsf{A}, \mathsf{Bar}) = \frac{1}{2}$). Once he is at the Bar, he never leaves it. He wants to know how likely he is to go back Home.

4. Runs

- A *run* of $\mathcal{M} = (S, \rightarrow, P)$ is an infinite sequence $s_0 s_1 s_2 \ldots$ of states such that for all $i \geq 0$, $P(s_i, s_{i+1}) > 0$. The set of runs of $\mathcal{M}$ is denoted $\mathsf{Runs}(\mathcal{M})$.
- Given $s_0 \in S$ an initial state, we can define a probability $\mathsf{Prob}^{\mathcal{M}}_{s_0}$ on the runs of $\mathcal{M}$.
- Given a set of runs, we would like to quantify the probability that a run from this set happens.
- In our example, the probability of the set of runs starting with $\mathsf{A} \rightarrow \mathsf{B} \rightarrow \mathsf{Home} \ldots$ is easy to compute:

$$\mathsf{Prob}^{\mathcal{M}}_{\mathsf{A}}(\mathsf{A} \rightarrow \mathsf{B} \rightarrow \mathsf{Home} \ldots) = P(\mathsf{A}, \mathsf{B}) \cdot P(\mathsf{B}, \mathsf{Home}) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}.$$

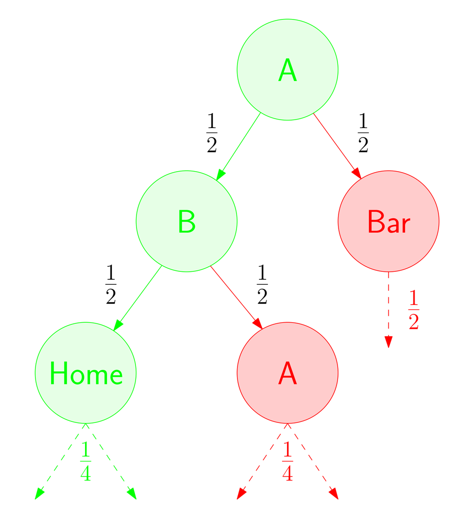

5. Quantitative reachability problem

Let $\mathcal{M} = (S, \rightarrow, P)$ a Markov chain and $F \subseteq S$ a set of states.
The set of runs eventually reaching a state in F is denoted $\Diamond F$.
A standard problem is to compute the probability of ever reaching any state in F from state s_0 (i.e. $\mathsf{Prob}^{\mathcal{M}}_{s_0}(\Diamond F)$). Since runs are infinite and the number of states can be infinite, we would be satisfied if we could calculate a close-enough approximation of this value.

APPROXIMATE QUANTITATIVE REACHABILITY

Inputs

- A Markov chain $\mathcal{M} = (S, \rightarrow, P)$,
- An initial state s_0,
- A set of states $F \subseteq S$,
- A rational $\epsilon > 0$.

Output A rational θ such that $\theta \leq \mathsf{Prob}^{\mathcal{M}}_{s_0}(\Diamond F) \leq \theta + \epsilon$.

In our previous example, let us assume that our goal is to reach Home ($F = \{\mathsf{Home}\}$). We notice that there is a positive probability to reach Home from A and B but not from Bar. How could we approximate the probability of reaching Home?

6. Approximation scheme [IN97]

For $F \subseteq S$, the *avoid-set* $\widetilde{F} = \{s \in S \mid \mathsf{Prob}^{\mathcal{M}}_s(\Diamond F) = 0\}$ is the set of states from which F is non-reachable.
For any $n \geq 0$, we can compute the probability of reaching F and $\widetilde{F}$ from an initial state s in less than n steps:

$$\begin{cases} p^{Yes}_n &= \mathsf{Prob}^{\mathcal{M}}_s(\Diamond_{\leq n} F), \\ p^{No}_n &= \mathsf{Prob}^{\mathcal{M}}_s(\neg F \ \mathbf{U}_{\leq n} \ \widetilde{F}). \end{cases}$$

To do so, we *unfold* the Markov chain from the initial state. We notice that every time we reach a state in F or $\widetilde{F}$, we can stop the unfolding. In our example, if $F = \{\mathsf{Home}\}$ and $\widetilde{F} = \{\mathsf{Bar}\}$,

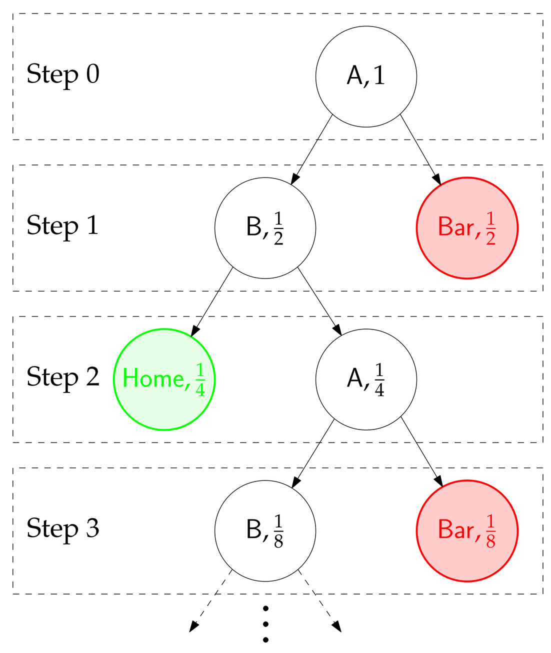

$\rightsquigarrow p^{Yes}_0 = 0,\ p^{No}_0 = 0,$

$\rightsquigarrow p^{Yes}_1 = 0,\ p^{No}_1 = \frac{1}{2},$

$\rightsquigarrow p^{Yes}_2 = \frac{1}{4},\ p^{No}_2 = \frac{1}{2},$

$\rightsquigarrow p^{Yes}_3 = \frac{1}{4},\ p^{No}_3 = \frac{1}{2} + \frac{1}{8} = \frac{5}{8} \ldots$

- For all $n \geq 0$, $p^{Yes}_n \leq \mathsf{Prob}^{\mathcal{M}}_s(\Diamond F) \leq 1 - p^{No}_n$.
- Moreover, $(p^{Yes}_n)_n$ and $(p^{No}_n)_n$ are both non-decreasing sequences.
- We stop the algorithm when $(1 - p^{No}_n) - p^{Yes}_n \leq \epsilon$ for a fixed $\epsilon > 0$.

This algorithm works well on this example but unfortunately, it does not always terminate.

7. Counterexample

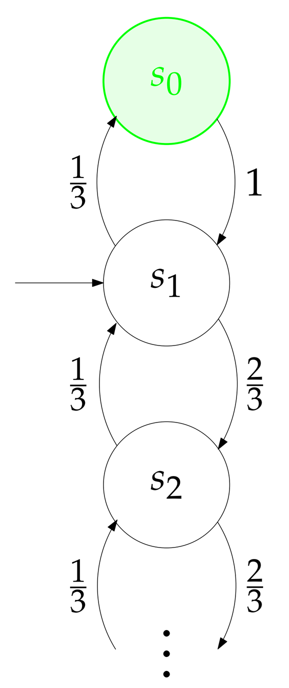

- Infinite number of states (random walk on the positive integers).
- We start in s_1, $F = \{s_0\} \implies \widetilde{F} = \emptyset$. Therefore, for $n \geq 0$, $p^{No}_n = 0$.
- We can compute via other means that $\mathsf{Prob}(\Diamond F) = \frac{2}{3}$, so for $n \geq 0$, $p^{Yes}_n \leq \frac{2}{3}$.

$\implies (1 - p^{No}_n) - p^{Yes}_n \geq \frac{1}{3}$ for any n.

$\implies$ if $0 < \epsilon < \frac{1}{3}$, the algorithm does not terminate.

8. When does it terminate? $\rightsquigarrow$ Decisiveness

Definition ([ABM07]). *A Markov chain $\mathcal{M}$ is decisive w.r.t. $F \subseteq S$ if for any initial state $s \in S$,*

$$\mathsf{Prob}^{\mathcal{M}}_s(\Diamond F \vee \Diamond \widetilde{F}) = 1.$$

Theorem ([ABM07]). *If $\mathcal{M}$ is decisive w.r.t. F, then the approximation scheme to compute $\mathsf{Prob}^{\mathcal{M}}(\Diamond F)$ is correct and terminates.*

Many classes of stochastic systems turn out to be decisive:

- finite Markov chains,
- Markov chains with a finite attractor and globally coarse ones [ABM07],
- reactive/single-clock stochastic timed automata [Car17].

9. Our goal...

...is to prove that *stochastic o-minimal hybrid systems* verify some decisiveness assumption.

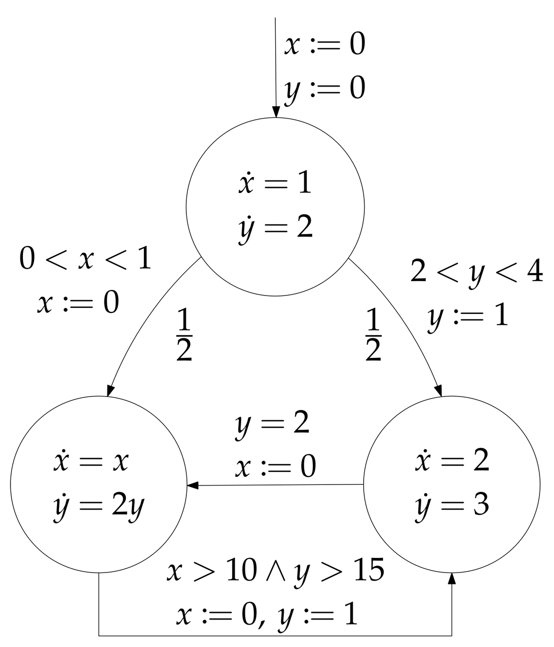

This model consists of

- finitely many discrete states,
- finitely many continuous variables,
- guards and resets on each edge,
- continuous distributions on time delays,
- discrete distributions on edges.

The set of states is thus uncountable ($S \times \mathbb{R}^n$, where n is the number of continuous variables). Stochastic o-minimal hybrid systems have two interesting properties making decisiveness possible:

- every variable has to be reset at each edge (*strong reset*);
- existence of a *finite time-abstract bisimulation*.

References

[ABM07] Parosh Aziz Abdulla, Noomene Ben Henda, and Richard Mayr. Decisive Markov chains. *Logical Methods in Computer Science*, 3(4), 2007.

[BBBC18] Nathalie Bertrand, Patricia Bouyer, Thomas Brihaye, and Pierre Carlier. When are stochastic transition systems tameable? *J. Log. Algebr. Meth. Program.*, 99:41–96, 2018.

[Car17] Pierre Carlier. *Verification of Stochastic Timed Automata. (Vérification des automates temporisés et stochastiques)*. PhD thesis, University of Paris-Saclay, France, 2017.

[IN97] S. Purushothaman Iyer and Murali Narasimha. Probabilistic lossy channel systems. In Michel Bidoit and Max Dauchet, editors, *TAPSOFT'97, Proceedings*, volume 1214 of *Lecture Notes in Computer Science*, pages 667–681. Springer, 1997.

Faculté des Sciences

Université de Mons

Pierre Vandenhove
Ongoing work with P. Bouyer, T. Brihaye, M. Randour, C. Rivière

Service de Mathématiques effectives
Département de Mathématiques