Taming real-time stochastic systems

Nathalie Bertrand – Inria Rennes, France

Formats'19 – Amsterdam

joint work with Patricia Bouyer, Thomas Brihaye and Pierre Carlier

Outline

Introduction

Finite and countable stochastic systems

Real-time stochastic systems
- General stochastic transitions systems
- Taming general stochastic transition systems
- Stochastic timed automata

Conclusion

Motivations

Various applications call for models with real time and probabilities

- clock synchronisation protocols
- root contention protocol
- CSMA: random backoff retransmission time
- molecular reactions
- thermostatically controlled loads
- etc.

Models from the literature combining real time and probabilities

- CTMC
- generalized semi-Markov processes
- stochastic timed automata
- Markov automata
- stochastic differential equations
- continuous-space pure jump Markov processes

Motivations

Various applications call for models with real time and probabilities

- clock synchronisation protocols
- root contention protocol
- CSMA: random backoff retransmission time
- molecular reactions
- thermostatically controlled loads
- etc.

Models from the literature combining real time and probabilities

- CTMC
- generalized semi-Markov processes
- stochastic timed automata
- Markov automata
- stochastic differential equations
- continuous-space pure jump Markov processes

Example

Generalized semi-Markov process for a 2-machine network

- crash events follow exponential distribution
- reboot events follow a uniform distribution

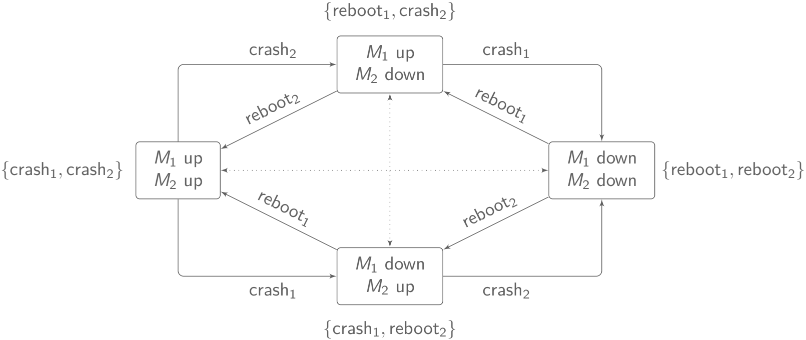

In state (M_1 up, M_2 down)

- delays are sampled for events crash_1 and reboot_2
- race condition determines next state

Example

Generalized semi-Markov process for a 2-machine network

- crash events follow exponential distribution
- reboot events follow a uniform distribution

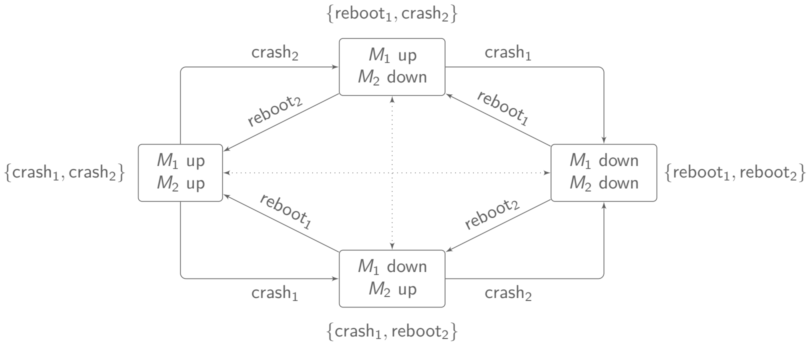

In state (M_1 up, M_2 down)

- delays are sampled for events crash_1 and reboot_2
- race condition determines next state

Real-time stochastic systems

Challenges

- intricate combination of dense time and probabilities
- uncountable state-space
- uncountable branching
- countinuous probability distributions

Model checking objectives

- qualitative: decide if probability of a given property is 1
- quantitative: compute the probability of a given property, or approximate it with arbitrary precision

Real-time stochastic systems

Challenges

- intricate combination of dense time and probabilities
- uncountable state-space
- uncountable branching
- countinuous probability distributions

Model checking objectives

- qualitative: decide if probability of a given property is 1
- quantitative: compute the probability of a given property, or approximate it with arbitrary precision

Outline

Discrete-time Markov chains

Discrete-time Markov chain (DTMC)
$\mathcal{M} = (S, s_0, \delta)$ with $\delta : S \to \mathsf{Dist}(S)$

Examples:

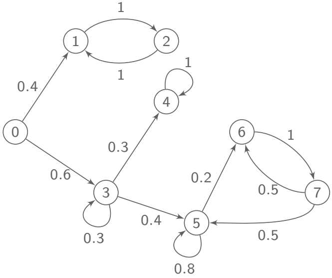

finite Markov chain

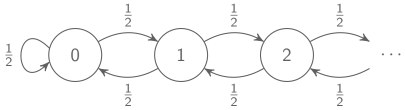

countable Markov chain

Finite DTMC - Quantitative model checking

Aim: Compute probability of reachability property **F**Goal

For state $s \in S$, let $x_s = \mathbb{P}_s(\mathbf{F}\mathrm{Goal})$.

$$x_s = \begin{cases} x_s = 1 & \text{if } s \in \mathrm{Goal} \\ x_s = 0 & \text{if } s \not\models \exists \mathbf{F}\mathrm{Goal} \\ x_s = \sum_{t \in S} \mathbf{P}(s,t)\, x_t & \text{otherwise.} \end{cases}$$

$\rightarrow$ resolution of a system of linear equations

Finite DTMC - Quantitative model checking

Aim: Compute probability of reachability property **F**Goal

For state $s \in S$, let $x_s = \mathbb{P}_s(\mathbf{F}\mathrm{Goal})$.

$$x_s = \begin{cases} x_s = 1 & \text{if } s \in \mathrm{Goal} \\ x_s = 0 & \text{if } s \not\models \exists \mathbf{F}\mathrm{Goal} \\ x_s = \sum_{t \in S} \mathbf{P}(s,t)\ x_t & \text{otherwise.} \end{cases}$$

$\rightarrow$ resolution of a system of linear equations

Countable DTMC - Quantitative model checking

Aim: Compute probability of reachability property $\mathbf{F}\text{Goal}$
$\text{Trap} = \{s \in S \mid s \not\models \exists \mathbf{F}\text{Goal}\}$

Approximation scheme [IN97]

given precision ε

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Goal}) \\ p_n^{\text{no}} = \mathbb{P}_{s_0}(\neg\text{Goal}\ \mathbf{U}_{\leq n}\text{Trap}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$p_1^{\text{yes}} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_1^{no}$

$p_2^{\text{yes}} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_2^{no}$

Does it terminate?

[IN97] P. Iyer and M. Narasimha. Probabilistic lossy channel systems. TAPSOFT'97, LNCS 1214, 667–681.

Countable DTMC - Quantitative model checking

Aim: Compute probability of reachability property $\mathbf{F}\mathsf{Goal}$

$\mathsf{Trap} = \{s \in S \mid s \not\models \exists \mathbf{F}\mathsf{Goal}\}$

Approximation scheme [IN97]

given precision ε

$$\begin{cases} p_n^{\mathrm{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\mathsf{Goal}) \\ p_n^{\mathrm{no}} = \mathbb{P}_{s_0}(\neg \mathsf{Goal}\ \mathbf{U}_{\leq n}\mathsf{Trap}) \end{cases}$$

until $p_n^{\mathrm{yes}} + p_n^{\mathrm{no}} \geq 1 - \varepsilon$

$$p_1^{yes} \leq \mathbb{P}(s \models \mathbf{F}\mathsf{Goal}) \leq 1 - p_1^{no}$$

$$p_2^{yes} \leq \mathbb{P}(s \models \mathbf{F}\mathsf{Goal}) \leq 1 - p_2^{no}$$

Does it terminate?

[IN97] P. Iyer and M. Narasimha. Probabilistic lossy channel systems. TAPSOFT'97, LNCS 1214, 667–681.

Countable DTMC - Quantitative model checking

Aim: Compute probability of reachability property **F**Goal
Trap $= \{s \in S \mid s \not\models \exists \mathbf{F}\text{Goal}\}$

Approximation scheme [IN97]

given precision ε

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Goal}) \\ p_n^{\text{no}} = \mathbb{P}_{s_0}(\neg\text{Goal}\ \mathbf{U}_{\leq n}\text{Trap}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$p_1^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_1^{no}$

$p_2^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_2^{no}$

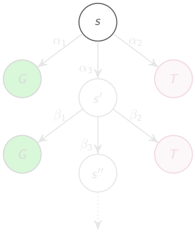

Does it terminate?

[IN97] P. Iyer and M. Narasimha. Probabilistic lossy channel systems. TAPSOFT'97, LNCS 1214, 667–681.

Countable DTMC - Quantitative model checking

Aim: Compute probability of reachability property **F**Goal

Trap $= \{s \in S \mid s \not\models \exists \mathbf{F}\text{Goal}\}$

Approximation scheme [IN97]

given precision ε

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Goal}) \\ p_n^{\text{no}} = \mathbb{P}_{s_0}(\neg\text{Goal } \mathbf{U}_{\leq n}\text{Trap}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$$p_1^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_1^{no}$$

$$p_2^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_2^{no}$$

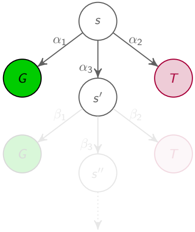

Does it terminate?

[IN97] P. Iyer and M. Narasimha. Probabilistic lossy channel systems. TAPSOFT'97, LNCS 1214, 667–681.

Countable DTMC - Quantitative model checking

Aim: Compute probability of reachability property **F**Goal
Trap $= \{s \in S \mid s \not\models \exists \mathbf{F}\text{Goal}\}$

Approximation scheme [IN97]

given precision ε

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Goal}) \\ p_n^{\text{no}} = \mathbb{P}_{s_0}(\neg\text{Goal}\ \mathbf{U}_{\leq n}\text{Trap}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$$p_1^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_1^{no}$$

$$p_2^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_2^{no}$$

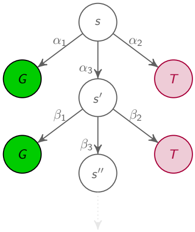

Does it terminate?

[IN97] P. Iyer and M. Narasimha. Probabilistic lossy channel systems. TAPSOFT'97, LNCS 1214, 667–681.

Countable DTMC - Quantitative model checking

Aim: Compute probability of reachability property $\mathbf{F}\text{Goal}$

$\text{Trap} = \{s \in S \mid s \not\models \exists\mathbf{F}\text{Goal}\}$

Approximation scheme [IN97]

given precision ε

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Goal}) \\ p_n^{\text{no}} = \mathbb{P}_{s_0}(\neg\text{Goal}\ \mathbf{U}_{\leq n}\text{Trap}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$$p_1^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_1^{no}$$

$$p_2^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_2^{no}$$

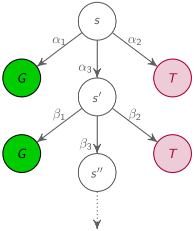

Does it terminate?

[IN97] P. Iyer and M. Narasimha. Probabilistic lossy channel systems. TAPSOFT'97, LNCS 1214, 667–681.

Countable DTMC - Quantitative model checking

Aim: Compute probability of reachability property $\mathbf{F}\text{Goal}$

$\text{Trap} = \{s \in S \mid s \not\models \exists \mathbf{F}\text{Goal}\}$

Approximation scheme [IN97]

given precision ε

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Goal}) \\ p_n^{\text{no}} = \mathbb{P}_{s_0}(\neg\text{Goal}\ \mathbf{U}_{\leq n}\text{Trap}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

$$p_1^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_1^{no}$$

$$p_2^{yes} \leq \mathbb{P}(s \models \mathbf{F}\text{Goal}) \leq 1 - p_2^{no}$$

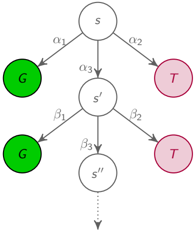

Does it terminate?

[IN97] P. Iyer and M. Narasimha. Probabilistic lossy channel systems. TAPSOFT'97, LNCS 1214, 667–681.

Non-converging example

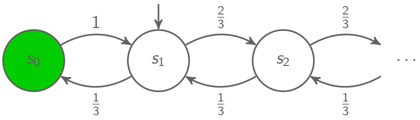

For Goal $= \{s_0\}$

- Trap $= \emptyset$, thus $\forall n,\ p_n^{no} = \mathbb{P}_{s_1}(\mathbf{F}^{\leq n}\text{Trap}) = 0$
- $\mathbb{P}_{s_1}(\mathbf{F}\text{Goal}) = \frac{2}{3}$, thus $\forall n,\ p_n^{yes} \leq \frac{2}{3}$
- $p_n^{yes} + p_n^{no} \leq \frac{2}{3}$

Decisiveness

Decisiveness [ABM07]
$\mathcal{M}$ is *decisive* w.r.t. Goal if $\forall s \in S$, $\mathbb{P}_s(\mathbf{F}\text{Goal} \vee \mathbf{F}\text{Trap}) = 1$

Examples of decisive Markov chains:
finite Markov chains, probabilistic lossy channel systems, probabilistic vector addition systems, noisy Turing machines

Counterexample:

not decisive w.r.t. Goal $= \{s_0\}$
since $\mathbb{P}_{s_1}(\mathbf{F}\text{Goal} \vee \mathbf{F}\text{Trap}) = \frac{2}{3}$

Termination for decisive Markov chains [ABM07]
If $\mathcal{M}$ is decisive w.r.t. Goal, then the approximation scheme is correct and terminates.

[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr: Decisive Markov Chains. LMCS 3(4) (2007)

Decisiveness

Decisiveness [ABM07]
$\mathcal{M}$ is *decisive* w.r.t. Goal if $\forall s \in S$, $\mathbb{P}_s(\mathbf{F}\text{Goal} \vee \mathbf{F}\text{Trap}) = 1$

Examples of decisive Markov chains:
finite Markov chains, probabilistic lossy channel systems, probabilistic vector addition systems, noisy Turing machines

Counterexample:

not decisive w.r.t. Goal $= \{s_0\}$
since $\mathbb{P}_{s_1}(\mathbf{F}\text{Goal} \vee \mathbf{F}\text{Trap}) = \frac{2}{3}$

Termination for decisive Markov chains [ABM07]
If $\mathcal{M}$ is decisive w.r.t. Goal, then the approximation scheme is correct and terminates.

[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr: Decisive Markov Chains. LMCS 3(4) (2007)

Decisiveness

Decisiveness [ABM07]
$\mathcal{M}$ is *decisive* w.r.t. Goal if $\forall s \in S,\ \mathbb{P}_s(\mathbf{F}\mathsf{Goal} \vee \mathbf{F}\mathsf{Trap}) = 1$

Examples of decisive Markov chains:
finite Markov chains, probabilistic lossy channel systems, probabilistic vector addition systems, noisy Turing machines

Counterexample:

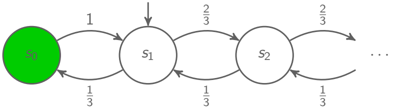

not decisive w.r.t. Goal $= \{s_0\}$
since $\mathbb{P}_{s_1}(\mathbf{F}\mathsf{Goal} \vee \mathbf{F}\mathsf{Trap}) = \frac{2}{3}$

Termination for decisive Markov chains [ABM07]
If $\mathcal{M}$ is decisive w.r.t. Goal, then the approximation scheme is correct and terminates.

[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr: Decisive Markov Chains. LMCS 3(4) (2007)

Decisiveness

Decisiveness [ABM07]

$\mathcal{M}$ is *decisive* w.r.t. Goal if $\forall s \in S,\ \mathbb{P}_s(\mathbf{F}\text{Goal} \vee \mathbf{F}\text{Trap}) = 1$

Examples of decisive Markov chains:
finite Markov chains, probabilistic lossy channel systems, probabilistic vector addition systems, noisy Turing machines

Counterexample:

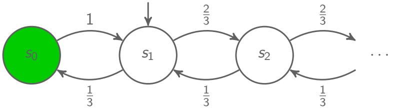

not decisive w.r.t. Goal $= \{s_0\}$
since $\mathbb{P}_{s_1}(\mathbf{F}\text{Goal} \vee \mathbf{F}\text{Trap}) = \frac{2}{3}$

Termination for decisive Markov chains [ABM07]

If $\mathcal{M}$ is decisive w.r.t. Goal, then the approximation scheme is correct and terminates.

[ABM07] P.A. Abdulla, N. Ben Henda, R. Mayr: Decisive Markov Chains. LMCS 3(4) (2007)

Beyond reachability - repeated reachability

Aim: Compute probability of repeated reachability property **GF**Goal

variant of approximation scheme using coTrap $= \{s \in S \mid s \not\models \exists \mathbf{F}\mathsf{Trap}\}$

Approximation scheme

given precision ε

$$\begin{cases} q_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\mathsf{coTrap}) \\ q_n^{\text{no}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\mathsf{Trap}) \end{cases}$$

until $q_n^{\text{yes}} + q_n^{\text{no}} \geq 1 - \varepsilon$

$$q_n^{\text{yes}} \leq \mathbb{P}_{s_0}(\mathbf{GF}\mathsf{Goal}) \leq 1 - q_n^{\text{no}}$$

Termination for decisive Markov chains

If $\mathcal{M}$ is decisive w.r.t. Goal **and** Trap, the approximation scheme is correct and terminates.

Beyond reachability - repeated reachability

Aim: Compute probability of repeated reachability property $\mathbf{GF}\mathsf{Goal}$

variant of approximation scheme using $\mathsf{coTrap} = \{s \in S \mid s \not\models \exists\mathbf{F}\mathsf{Trap}\}$

Approximation scheme

given precision ε

$$\begin{cases} q_n^{\mathsf{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\mathsf{coTrap}) \\ q_n^{\mathsf{no}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\mathsf{Trap}) \end{cases}$$

until $q_n^{\mathsf{yes}} + q_n^{\mathsf{no}} \geq 1 - \varepsilon$

$$q_n^{\mathsf{yes}} \leq \mathbb{P}_{s_0}(\mathbf{GF}\mathsf{Goal}) \leq 1 - q_n^{\mathsf{no}}$$

Termination for decisive Markov chains

If $\mathcal{M}$ is decisive w.r.t. Goal **and** Trap, the approximation scheme is correct and terminates.

Beyond reachability - repeated reachability

Aim: Compute probability of repeated reachability property $\mathbf{GF}\text{Goal}$

variant of approximation scheme using $\text{coTrap} = \{s \in S \mid s \not\models \exists \mathbf{F}\text{Trap}\}$

Approximation scheme

given precision ε

$$\begin{cases} q_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{coTrap}) \\ q_n^{\text{no}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Trap}) \end{cases}$$

until $q_n^{\text{yes}} + q_n^{\text{no}} \geq 1 - \varepsilon$

$$q_n^{\text{yes}} \leq \mathbb{P}_{s_0}(\mathbf{GF}\text{Goal}) \leq 1 - q_n^{\text{no}}$$

Termination for decisive Markov chains

If $\mathcal{M}$ is decisive w.r.t. Goal **and** Trap, the approximation scheme is correct and terminates.

Beyond reachability - ω-regular properties

Aim: Compute probability of Muller property $\mathsf{Inf} \in \mathcal{F}$

Attractor
Attr is an attractor for $\mathcal{M}$ if $\forall s \in S,\ \mathbb{P}_s(\mathbf{F}\mathsf{Attr}) = 1$

$\mathcal{M}$ admits a finite attractor Attr $\implies$ $\mathcal{M}$ is decisive w.r.t. any Goal

- From Attr build Graph(Attr) and compute its BSCCs.
- Identify BSCC that are good regarding the Muller condition.

C_1 is good iff $\exists F \in \mathcal{F}$

$\forall q\ (C_1 \to^\star q) \Rightarrow (q \in F)$

$\forall q\ (q \in F) \Rightarrow (q \to^\star C_1)$

- $\mathbb{P}_s(\mathsf{Inf} \in \mathcal{F}) = \sum_{C \text{ Good BSCC}} \mathbb{P}_s(\mathbf{F}C)$.
- Use approximation scheme to compute $\mathbb{P}_s(\mathbf{F}C)$

[ABRS05] P. A. Abdulla, **NB**, A. Rabinovich, P. Schnoebelen: Verification of probabilistic systems with faulty communication. I&C, 2005.

Beyond reachability - ω-regular properties

Aim: Compute probability of Muller property $\mathsf{Inf} \in \mathcal{F}$

Attractor

Attr is an attractor for $\mathcal{M}$ if $\forall s \in S,\ \mathbb{P}_s(\mathbf{F}\mathsf{Attr}) = 1$

$\mathcal{M}$ admits a finite attractor Attr $\implies$ $\mathcal{M}$ is decisive w.r.t. any Goal

- ▷ From Attr build Graph(Attr) and compute its BSCCs.
- ▷ Identify BSCC that are good regarding the Muller condition.

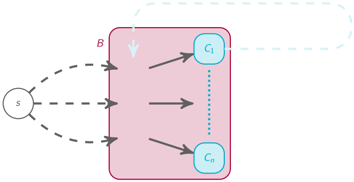

C_1 is good iff $\exists F \in \mathcal{F}$

$\forall q\ (C_1 \to^\star q) \Rightarrow (q \in F)$

$\forall q\ (q \in F) \Rightarrow (q \to^\star C_1)$

- ▷ $\mathbb{P}_s(\mathsf{Inf} \in \mathcal{F}) = \sum_{C \text{ Good BSCC}} \mathbb{P}_s(\mathbf{F}C)$.
- ▷ Use approximation scheme to compute $\mathbb{P}_s(\mathbf{F}C)$

[ABRS05] P. A. Abdulla, **NB**, A. Rabinovich, P. Schnoebelen: Verification of probabilistic systems with faulty communication. I&C, 2005.

Beyond reachability - ω-regular properties

Aim: Compute probability of Muller property $\mathsf{Inf} \in \mathcal{F}$

Attractor

Attr is an attractor for $\mathcal{M}$ if $\forall s \in S,\ \mathbb{P}_s(\mathbf{F}\mathsf{Attr}) = 1$

$\mathcal{M}$ admits a finite attractor Attr $\implies$ $\mathcal{M}$ is decisive w.r.t. any Goal

- ▷ From Attr build Graph(Attr) and compute its BSCCs.
- ▷ Identify BSCC that are good regarding the Muller condition.

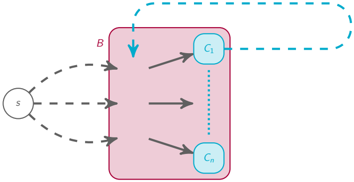

C_1 is good iff $\exists F \in \mathcal{F}$

$\forall q\ (C_1 \to^\star q) \Rightarrow (q \in F)$

$\forall q\ (q \in F) \Rightarrow (q \to^\star C_1)$

- ▷ $\mathbb{P}_s(\mathsf{Inf} \in \mathcal{F}) = \sum_{C \text{ Good BSCC}} \mathbb{P}_s(\mathbf{F}C)$.
- ▷ Use approximation scheme to compute $\mathbb{P}_s(\mathbf{F}C)$

[ABRS05] P. A. Abdulla, **NB**, A. Rabinovich, P. Schnoebelen: Verification of probabilistic systems with faulty communication. I&C, 2005.

Beyond reachability - ω-regular properties

Aim: Compute probability of Muller property $\mathsf{Inf} \in \mathcal{F}$

Attractor

Attr is an attractor for $\mathcal{M}$ if $\forall s \in S,\ \mathbb{P}_s(\mathbf{F}\mathsf{Attr}) = 1$

$\mathcal{M}$ admits a finite attractor Attr $\implies$ $\mathcal{M}$ is decisive w.r.t. any Goal

- From Attr build Graph(Attr) and compute its BSCCs.
- Identify BSCC that are good regarding the Muller condition.

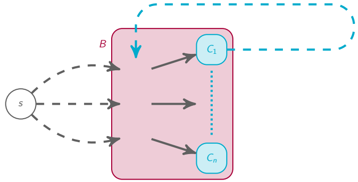

C_1 is good iff $\exists F \in \mathcal{F}$

$\forall q\ (C_1 \to^\star q) \Rightarrow (q \in F)$

$\forall q\ (q \in F) \Rightarrow (q \to^\star C_1)$

- $\mathbb{P}_s(\mathsf{Inf} \in \mathcal{F}) = \sum_{C \text{ Good BSCC}} \mathbb{P}_s(\mathbf{F}C)$.
- Use approximation scheme to compute $\mathbb{P}_s(\mathbf{F}C)$

[ABRS05] P. A. Abdulla, **NB**, A. Rabinovich, P. Schnoebelen: Verification of probabilistic systems with faulty communication. I&C, 2005.

Beyond reachability - ω-regular properties

Aim: Compute probability of Muller property $\mathsf{Inf} \in \mathcal{F}$

Attractor

Attr is an attractor for $\mathcal{M}$ if $\forall s \in S,\ \mathbb{P}_s(\mathbf{F}\mathsf{Attr}) = 1$

$\mathcal{M}$ admits a finite attractor Attr $\implies$ $\mathcal{M}$ is decisive w.r.t. any Goal

- From Attr build Graph(Attr) and compute its BSCCs.
- Identify BSCC that are good regarding the Muller condition.

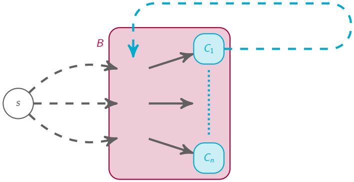

C_1 is good iff $\exists F \in \mathcal{F}$

$$\forall q\ (C_1 \to^\star q) \Rightarrow (q \in F)$$

$$\forall q\ (q \in F) \Rightarrow (q \to^\star C_1)$$

- $\mathbb{P}_s(\mathsf{Inf} \in \mathcal{F}) = \sum_{C \text{ Good BSCC}} \mathbb{P}_s(\mathbf{F}C)$.
- Use approximation scheme to compute $\mathbb{P}_s(\mathbf{F}C)$

[ABRS05] P. A. Abdulla, **NB**, A. Rabinovich, P. Schnoebelen: Verification of probabilistic systems with faulty communication. I&C, 2005.

Outline

Introduction

Finite and countable stochastic systems

Real-time stochastic systems
- General stochastic transitions systems
- Taming general stochastic transition systems
- Stochastic timed automata

Conclusion

Outline

Stochastic transitions systems

Stochastic transition systems (STS)
$\mathcal{T} = (S, \Sigma, \kappa)$ with (S, Σ) a measurable space and $\kappa : S \times \Sigma \to [0, 1]$ a Markov kernel such that $\forall s \in S,\ \kappa(s, \cdot) \in \mathsf{Dist}(S)$

Examples of STS

- countable Markov chains: $\Sigma = 2^S$
- continuous time Markov chains (CTMC)
- stochastic timed automata
- generalised semi-Markov processes
- etc.

Stochastic timed automata

Stochastic timed automata (STA):
timed automata with random delays

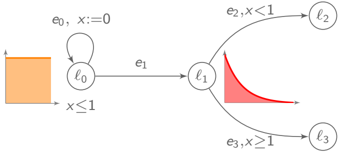

distrib. on possible delays
uniform distribution in ℓ_0
exponential distribution in ℓ_1

Semantics: from state s

1. pick a delay τ following the distribution
2. choose at random an edge enabled at $s + \tau$

STA remedy artefacts of standard timed automata such as

- arbitrary precision
- time-convergence

Markov models with uncountable state-space; real-time behaviour

[BBBM+14] **NB**, P. Bouyer, T. Brihaye, Q. Menet *et al.* Stochastic timed automata. LMCS, 2014.

Stochastic timed automata

Stochastic timed automata (STA):
timed automata with random delays

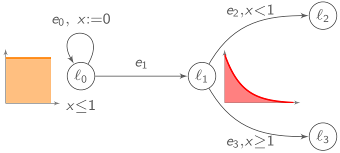

distrib. on possible delays
uniform distribution in ℓ_0
exponential distribution in ℓ_1

Semantics: from state s

1. pick a delay τ following the distribution
2. choose at random an edge enabled at $s + \tau$

STA remedy artefacts of standard timed automata such as

- arbitrary precision
- time-convergence

Markov models with uncountable state-space; real-time behaviour

[BBBM+14] **NB**, P. Bouyer, T. Brihaye, Q. Menet *et al.* Stochastic timed automata. LMCS, 2014.

Stochastic timed automata

Stochastic timed automata (STA):
timed automata with random delays

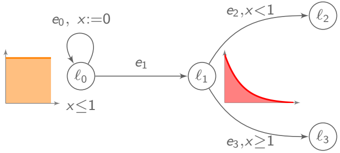

distrib. on possible delays
uniform distribution in ℓ_0
exponential distribution in ℓ_1

Semantics: from state s

1. pick a delay τ following the distribution
2. choose at random an edge enabled at $s + \tau$

STA remedy artefacts of standard timed automata such as

- arbitrary precision
- time-convergence

Markov models with uncountable state-space; real-time behaviour

[BBBM+14] **NB**, P. Bouyer, T. Brihaye, Q. Menet *et al.* Stochastic timed automata. LMCS, 2014.

Stochastic timed automata

Stochastic timed automata (STA):
timed automata with random delays

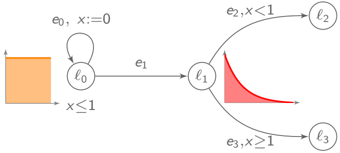

distrib. on possible delays
uniform distribution in ℓ_0
exponential distribution in ℓ_1

Semantics: from state s

1. pick a delay τ following the distribution
2. choose at random an edge enabled at $s + \tau$

STA remedy artefacts of standard timed automata such as

- arbitrary precision
- time-convergence

Markov models with uncountable state-space; real-time behaviour

[BBBM+14] **NB**, P. Bouyer, T. Brihaye, Q. Menet *et al.* Stochastic timed automata. LMCS, 2014.

Outline

Extending decisiveness beyond countable DTMC

Difficulty: $s \models \exists \mathbf{F}\,\mathsf{Goal} \not\Rightarrow \mathbb{P}_s(\mathbf{F}\,\mathsf{Goal}) > 0$

Example on STA:

- $(\ell_0, 0) \models \exists \mathbf{F}(\ell_3, \cdot)$
- $\mathbb{P}_{(\ell_0,0)}(\mathbf{F}(\ell_3, \cdot)) = 0$

$\rightarrow$ trap must be redefined: $\mathsf{Trap} = \{s \in S \mid \mathbb{P}_s(\mathbf{F}\,\mathsf{Goal}) = 0\}$

Decisiveness

$\mathcal{T}$ is *decisive* w.r.t. Goal if $\forall \mu \in \mathsf{Dist}(S)$, $\mathbb{P}_\mu(\mathbf{F}\,\mathsf{Goal} \vee \mathbf{F}\,\mathsf{Trap}) = 1$

How to check decisiveness?

How to perform qualitative/quantitative model checking of decisive STS?

Through the lense of an **abstraction** function!

Extending decisiveness beyond countable DTMC

Difficulty: $s \models \exists \mathbf{F}\mathsf{Goal} \not\Rightarrow \mathbb{P}_s(\mathbf{F}\mathsf{Goal}) > 0$

Example on STA:

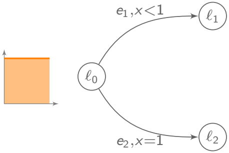

- $(\ell_0, 0) \models \exists \mathbf{F}(\ell_3, \cdot)$
- $\mathbb{P}_{(\ell_0,0)}(\mathbf{F}(\ell_3, \cdot)) = 0$

$\rightarrow$ trap must be redefined: $\mathsf{Trap} = \{s \in S \mid \mathbb{P}_s(\mathbf{F}\mathsf{Goal}) = 0\}$

Decisiveness

$\mathcal{T}$ is *decisive* w.r.t. Goal if $\forall \mu \in \mathsf{Dist}(S),\ \mathbb{P}_\mu(\mathbf{F}\mathsf{Goal} \vee \mathbf{F}\mathsf{Trap}) = 1$

How to check decisiveness?

How to perform qualitative/quantitative model checking of decisive STS?

Through the lense of an **abstraction** function!

Extending decisiveness beyond countable DTMC

Difficulty: $s \models \exists \mathbf{F}\text{Goal} \not\Rightarrow \mathbb{P}_s(\mathbf{F}\text{Goal}) > 0$

Example on STA:

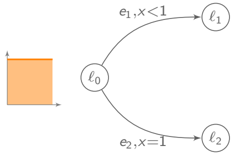

- $(\ell_0, 0) \models \exists \mathbf{F}(\ell_3, \cdot)$
- $\mathbb{P}_{(\ell_0,0)}(\mathbf{F}(\ell_3, \cdot)) = 0$

$\rightarrow$ trap must be redefined: $\text{Trap} = \{s \in S \mid \mathbb{P}_s(\mathbf{F}\text{Goal}) = 0\}$

Decisiveness
$\mathcal{T}$ is *decisive* w.r.t. Goal if $\forall \mu \in \text{Dist}(S),\ \mathbb{P}_\mu(\mathbf{F}\text{Goal} \vee \mathbf{F}\text{Trap}) = 1$

How to check decisiveness?
How to perform qualitative/quantitative model checking of decisive STS?
Through the lense of an **abstraction** function!

Extending decisiveness beyond countable DTMC

Difficulty: $s \models \exists \mathbf{F}\text{Goal} \not\Rightarrow \mathbb{P}_s(\mathbf{F}\text{Goal}) > 0$

Example on STA:

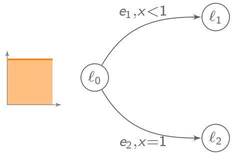

- $(\ell_0, 0) \models \exists \mathbf{F}(\ell_3, \cdot)$
- $\mathbb{P}_{(\ell_0,0)}(\mathbf{F}(\ell_3, \cdot)) = 0$

$\rightarrow$ trap must be redefined: $\text{Trap} = \{s \in S \mid \mathbb{P}_s(\mathbf{F}\text{Goal}) = 0\}$

Decisiveness

$\mathcal{T}$ is *decisive* w.r.t. Goal if $\forall \mu \in \text{Dist}(S),\ \mathbb{P}_\mu(\mathbf{F}\text{Goal} \vee \mathbf{F}\text{Trap}) = 1$

How to check decisiveness?

How to perform qualitative/quantitative model checking of decisive STS?

Through the lense of an **abstraction** function!

Extending decisiveness beyond countable DTMC

Difficulty: $s \models \exists \mathbf{F}\mathsf{Goal} \not\Rightarrow \mathbb{P}_s(\mathbf{F}\mathsf{Goal}) > 0$

Example on STA:

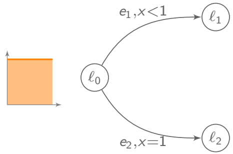

- $(\ell_0, 0) \models \exists \mathbf{F}(\ell_3, \cdot)$
- $\mathbb{P}_{(\ell_0,0)}(\mathbf{F}(\ell_3, \cdot)) = 0$

$\rightarrow$ trap must be redefined: $\mathsf{Trap} = \{s \in S \mid \mathbb{P}_s(\mathbf{F}\mathsf{Goal}) = 0\}$

Decisiveness

$\mathcal{T}$ is *decisive* w.r.t. Goal if $\forall \mu \in \mathsf{Dist}(S),\ \mathbb{P}_\mu(\mathbf{F}\mathsf{Goal} \vee \mathbf{F}\mathsf{Trap}) = 1$

How to check decisiveness?
How to perform qualitative/quantitative model checking of decisive STS?

Through the lense of an **abstraction** function!

Abstractions

Abstraction

For two STS $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2)$, and $\alpha : (S_1, \Sigma_1) \to (S_2, \Sigma_2)$ a measurable function

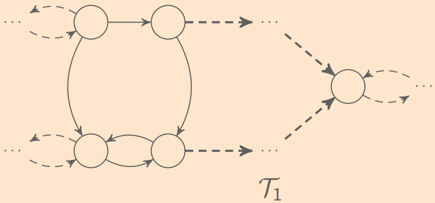

Abstractions

Abstraction

For two STS $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2)$, and $\alpha : (S_1, \Sigma_1) \to (S_2, \Sigma_2)$ a measurable function

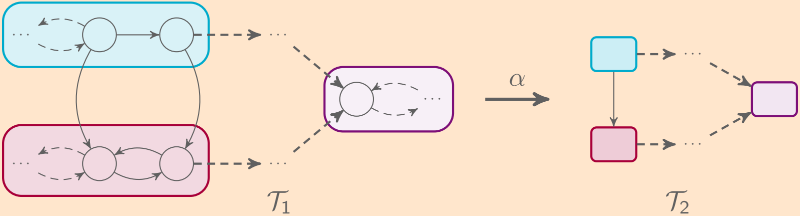

Abstractions

Abstraction

For two STS $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2)$, and $\alpha : (S_1, \Sigma_1) \to (S_2, \Sigma_2)$ a measurable function

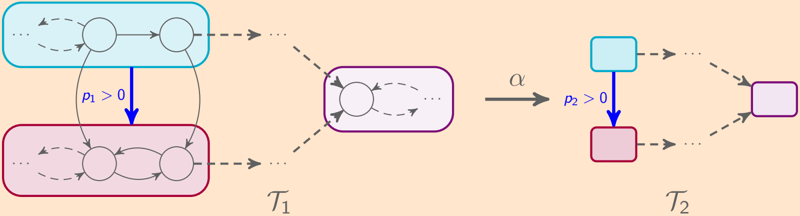

- $\mathcal{T}_2$ is an α-abstraction of $\mathcal{T}_1$ whenever $p_1 > 0 \iff p_2 > 0$.
- $\mathcal{T}_2$ is a sound α-abstraction of $\mathcal{T}_1$ whenever for each $B \in \Sigma_2$:

$$\mathbb{P}^{\mathcal{T}_2}(\mathbf{F}B) = 1 \quad \Rightarrow \quad \mathbb{P}^{\mathcal{T}_1}(\mathbf{F}\alpha^{-1}(B)) = 1$$

Abstractions

Abstraction

For two STS $\mathcal{T}_1 = (S_1, \Sigma_1, \kappa_1)$ and $\mathcal{T}_2 = (S_2, \Sigma_2, \kappa_2)$, and $\alpha : (S_1, \Sigma_1) \to (S_2, \Sigma_2)$ a measurable function

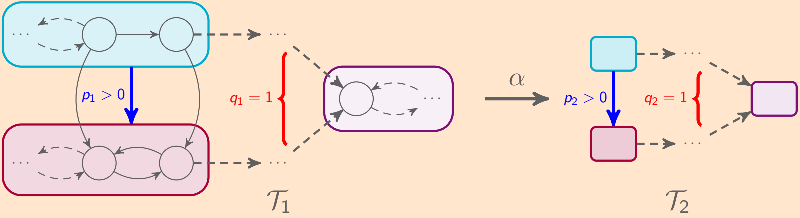

- $\mathcal{T}_2$ is an α-abstraction of $\mathcal{T}_1$ whenever $p_1 > 0 \iff p_2 > 0$.
- $\mathcal{T}_2$ is a sound α-abstraction of $\mathcal{T}_1$ whenever for each $B \in \Sigma_2$:

$$\mathbb{P}^{\mathcal{T}_2}(\mathbf{F}B) = 1 \quad \Rightarrow \quad \mathbb{P}^{\mathcal{T}_1}(\mathbf{F}\alpha^{-1}(B)) = 1$$

Abstractions, decisiveness and attractors

Example

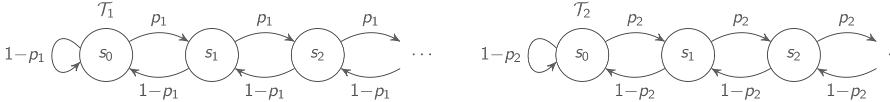

$$\alpha(s_i) = s_i, \quad \text{for all } i \in \mathbb{N}.$$

- $\mathcal{T}_2$ is an α-abstraction of $\mathcal{T}_1$ as soon as $p_1, p_2 \in (0,1)$.
- $\mathcal{T}_2$ is a sound α-abstraction of $\mathcal{T}_1$ iff $(p_1 > \frac{1}{2} \iff p_2 > \frac{1}{2})$.

Transferring decisiveness and attractors

If $\mathcal{T}_2$ is a sound α-abstraction of $\mathcal{T}_1$,

$$\mathcal{T}_2 \text{ decisive w.r.t. Goal} \implies \mathcal{T}_1 \text{ decisive w.r.t. } \alpha^{-1}(\text{Goal})$$

$$\text{Attr} \in \Sigma_2 \text{ attractor for } \mathcal{T}_2 \implies \alpha^{-1}(\text{Attr}) \text{ attractor for } \mathcal{T}_1$$

Abstractions, decisiveness and attractors

Example

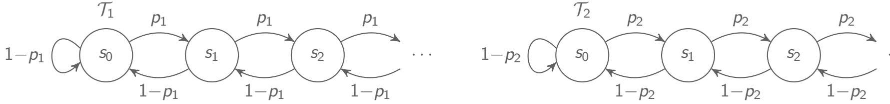

$$\alpha(s_i) = s_i, \quad \text{for all } i \in \mathbb{N}.$$

- $\mathcal{T}_2$ is an α-abstraction of $\mathcal{T}_1$ as soon as $p_1,\ p_2 \in (0,1)$.
- $\mathcal{T}_2$ is a sound α-abstraction of $\mathcal{T}_1$ iff $(p_1 > \frac{1}{2} \iff p_2 > \frac{1}{2})$.

Transferring decisiveness and attractors

If $\mathcal{T}_2$ is a sound α-abstraction of $\mathcal{T}_1$,

$$\mathcal{T}_2 \text{ decisive w.r.t. Goal} \implies \mathcal{T}_1 \text{ decisive w.r.t. } \alpha^{-1}(\text{Goal})$$

$$\text{Attr} \in \Sigma_2 \text{ attractor for } \mathcal{T}_2 \implies \alpha^{-1}(\text{Attr}) \text{ attractor for } \mathcal{T}_1$$

Abstractions, decisiveness and attractors

Example

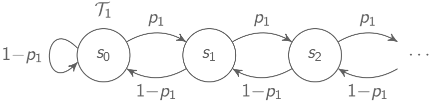

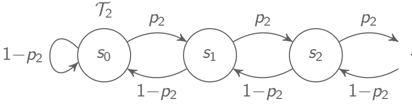

$$\alpha(s_i) = s_i, \quad \text{for all } i \in \mathbb{N}.$$

- $\mathcal{T}_2$ is an α-abstraction of $\mathcal{T}_1$ as soon as $p_1, p_2 \in (0,1)$.
- $\mathcal{T}_2$ is a sound α-abstraction of $\mathcal{T}_1$ iff $(p_1 > \frac{1}{2} \iff p_2 > \frac{1}{2})$.

Transferring decisiveness and attractors

If $\mathcal{T}_2$ is a sound α-abstraction of $\mathcal{T}_1$,

$$\mathcal{T}_2 \text{ decisive w.r.t. Goal} \implies \mathcal{T}_1 \text{ decisive w.r.t. } \alpha^{-1}(\text{Goal})$$

$$\text{Attr} \in \Sigma_2 \text{ attractor for } \mathcal{T}_2 \implies \alpha^{-1}(\text{Attr}) \text{ attractor for } \mathcal{T}_1$$

Reachability model checking for decisive STS

Approximation scheme
given precision ε

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Goal}) \\ p_n^{\text{no}} = \mathbb{P}_{s_0}(\neg\text{Goal}\ \mathbf{U}_{\leq n}\text{Trap}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

Quantitative reachability analysis
Let $\mathcal{T}$ be a general STS. If $\mathcal{T}$ is decisive w.r.t. Goal, then the p_n^{yes} and $(1 - p_n^{\text{no}})$ both converge to $\mathbb{P}(\mathbf{F}\text{Goal})$.

Applicability: the approximation scheme is effective if

- Trap can computed
- one can evaluate the values p_n^{yes} and $(1 - p_n^{\text{no}})$;
  i.e. one can compute (or approximate!) probabilities of cylinders of the form $\text{Cyl}(S \cdots S\text{Goal})$ and $\text{Cyl}(\neg\text{Goal} \cdots \neg\text{Goal}\text{Trap})$

Reachability model checking for decisive STS

Approximation scheme
given precision ε

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Goal}) \\ p_n^{\text{no}} = \mathbb{P}_{s_0}(\neg\text{Goal }\mathbf{U}_{\leq n}\text{Trap}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

Quantitative reachability analysis
Let $\mathcal{T}$ be a general STS. If $\mathcal{T}$ is decisive w.r.t. Goal, then the p_n^{yes} and $(1-p_n^{\text{no}})$ both converge to $\mathbb{P}(\mathbf{F}\text{Goal})$.

Applicability: the approximation scheme is effective if

- Trap can computed
- one can evaluate the values p_n^{yes} and $(1-p_n^{\text{no}})$;
  i.e. one can compute (or approximate!) probabilities of cylinders of the form $\text{Cyl}(S \cdots S\text{Goal})$ and $\text{Cyl}(\neg\text{Goal} \cdots \neg\text{Goal}\text{Trap})$

Reachability model checking for decisive STS

Approximation scheme
given precision ε

$$\begin{cases} p_n^{\text{yes}} = \mathbb{P}_{s_0}(\mathbf{F}_{\leq n}\text{Goal}) \\ p_n^{\text{no}} = \mathbb{P}_{s_0}(\neg\text{Goal}\ \mathbf{U}_{\leq n}\text{Trap}) \end{cases}$$

until $p_n^{\text{yes}} + p_n^{\text{no}} \geq 1 - \varepsilon$

Quantitative reachability analysis
Let $\mathcal{T}$ be a general STS. If $\mathcal{T}$ is decisive w.r.t. Goal, then the p_n^{yes} and $(1-p_n^{\text{no}})$ both converge to $\mathbb{P}(\mathbf{F}\text{Goal})$.

Applicability: the approximation scheme is effective if

- Trap can computed
- one can evaluate the values p_n^{yes} and $(1-p_n^{\text{no}})$;
  i.e. one can compute (or approximate!) probabilities of cylinders of the form $\text{Cyl}(S \cdots S\text{Goal})$ and $\text{Cyl}(\neg\text{Goal} \cdots \neg\text{Goal}\text{Trap})$

Model checking ω-regular properties

Framework:

- $\mathcal{T}_1$ general STS
- $\mathcal{T}_2$ countable Markov chain with finite attractor
- $\mathcal{T}_2$ sound α-abstraction of $\mathcal{T}_1$

Model checking Muller properties

- almost-sure model checking of Muller property in $\mathcal{T}_1$ reduces to almost-sure model checking of *reachability property in* $\mathcal{T}_2$;
- computation of the probability of Muller property on $\mathcal{T}_1$ reduces to computation of a *reachability probability in* $\mathcal{T}_1$.

Outline

Tameable stochastic timed automata

Natural abstraction: Markov chain build on region automaton

STA with an attractor

- single-clock STA

$$\text{Attr} = \{(\ell, 0)\} \cup \{(\ell, r) \mid \forall (\ell, r) \to^* (\ell', r'),\ r' = r\}$$

- reactive STA, *i.e.* STA where from every state all delays are possible

$$\text{Attr} = \{(\ell, r) \mid \forall x,\ x = 0 \text{ or } x > M \text{ in } r\}$$

Model checking STA

- we recover all known decidability/approximability results...
- ... and extend them! (*e.g.* Muller properties)

Tameable stochastic timed automata

Natural abstraction: Markov chain build on region automaton

STA with an attractor

- single-clock STA

$$\text{Attr} = \{(\ell, 0)\} \cup \{(\ell, r) \mid \forall (\ell, r) \to^* (\ell', r'),\ r' = r\}$$

- reactive STA, *i.e.* STA where from every state all delays are possible

$$\text{Attr} = \{(\ell, r) \mid \forall x,\ x = 0 \text{ or } x > M \text{ in } r\}$$

Model checking STA

- we recover all known decidability/approximability results...
- ... and extend them! (*e.g.* Muller properties)

Beyond tameable STA

Finding sound abstractions that are decisive is non trivial...

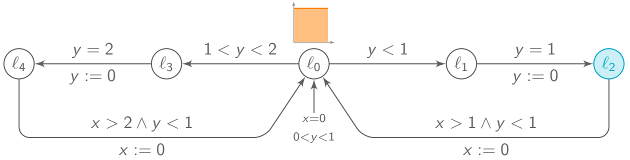

Converging phenomenon:
on entering ℓ_0 the value of clock y increases $\rightarrow \mathbb{P}(\mathbf{F}\ell_2) < 1$

First attempt

$\rightarrow \mathbb{P}_q^{\mathcal{T}_2}(\mathbf{F}\ell_2) = 1$. **not sound!**

Second attempt

$\rightarrow \mathbb{P}_q^{\mathcal{T}_2}(\mathbf{F}\ell_2) < 1$. **not decisive!**

Beyond tameable STA

Finding sound abstractions that are decisive is non trivial...

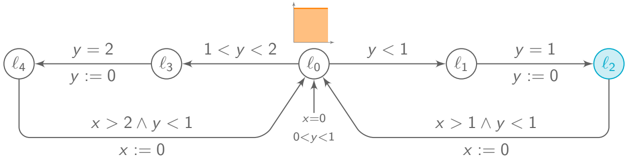

Converging phenomenon:
on entering ℓ_0 the value of clock y increases $\quad \rightarrow \mathbb{P}(\mathbf{F}\ell_2) < 1$

First attempt

$\rightarrow \mathbb{P}_q^{\mathcal{T}_2}(\mathbf{F}\ell_2) = 1$. **not sound!**

Second attempt

$\rightarrow \mathbb{P}_q^{\mathcal{T}_2}(\mathbf{F}\ell_2) < 1$. **not decisive!**

Beyond tameable STA

Finding sound abstractions that are decisive is non trivial...

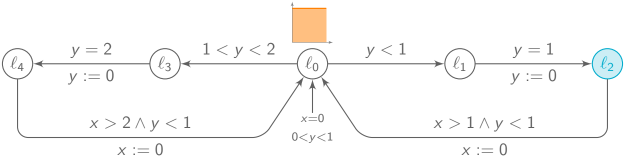

Converging phenomenon:
on entering ℓ_0 the value of clock y increases $\rightarrow \mathbb{P}(\mathbf{F}\ell_2) < 1$

First attempt

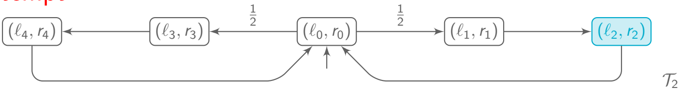

$\rightarrow \mathbb{P}_q^{\mathcal{T}_2}(\mathbf{F}\ell_2) = 1$. **not sound!**

Second attempt

$\rightarrow \mathbb{P}_q^{\mathcal{T}_2}(\mathbf{F}\ell_2) < 1$. **not decisive!**

Beyond tameable STA

Finding sound abstractions that are decisive is non trivial...

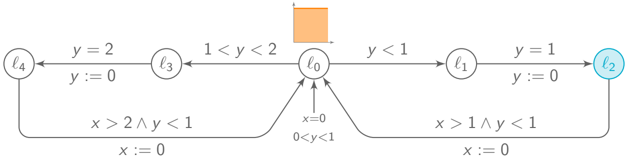

Converging phenomenon:
on entering ℓ_0 the value of clock y increases $\rightarrow \mathbb{P}(\mathbf{F}\ell_2) < 1$

First attempt

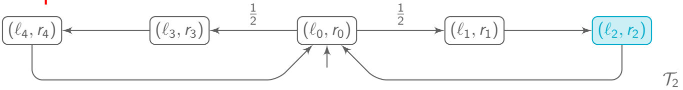

$\rightarrow$ $\mathbb{P}^{\mathcal{T}_2}_q(\mathbf{F}\ell_2) = 1$. **not sound!**

Second attempt

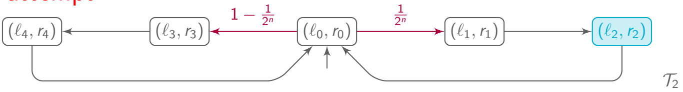

$\rightarrow$ $\mathbb{P}^{\mathcal{T}_2}_q(\mathbf{F}\ell_2) < 1$. **not decisive!**

Outline

Contributions

- decisiveness and attractor notions for general stochastic systems
- generic approach to analysing countinuous stochastic systems
  - algorithms for qualitative model checking
  - approximation schemes for quantitative model checking
- application to subclasses of real-time systems
  - stochastic timed automata
  - generalised semi-Markov processes
  - stochastic time Petri nets
- recovering and extending known results from the literature

$\rightarrow$ more results and all technical details in the article

NB, Patricia Bouyer, Thomas Brihaye and Pierre Carlier
When are Stochastic Transition Systems Tameable?
JLAMP 2018

Future work

- applicability to other classes of systems
  - candidate: timed lossy channel systems
- convergence speed of the approximation schemes
- beyond purely stochastic systems
  - decisiveness of Markov decision processes; already for countable MDP!
- compositionality of the approach

Dank U!

Future work

- applicability to other classes of systems
  - candidate: timed lossy channel systems
- convergence speed of the approximation schemes
- beyond purely stochastic systems
  - decisiveness of Markov decision processes; already for countable MDP!
- compositionality of the approach

Dank U!