MOVEP 2012 Tutorial
Safety, Dependability and Performance Analysis of Extended AADL Models

Part 6: Performability Evaluation

European Space Agency
European Space Research and Technology Centre

RWTH Aachen University
Software Modeling and Verification Group
Thomas Noll

Fondazione Bruno Kessler
Centre for Scientific and Technological Research
Alessandro Cimatti

MOVEP 2012 School; December 7, 2012; Marseille, France

Contents

Outline

Introduction

Error models

AADL error models are finite automata enriched with probabilistic failures (and repairs).

Introduction

Error models

AADL error models are finite automata enriched with probabilistic failures (and repairs).

Two kinds of error models can be distinguished:

- Discrete-time (cf. Marta Kwiatkowska's tutorial)
  - Failures and repairs are modeled by discrete probabilities
  - Instantaneous probabilistic decision to fail (or repair)

  ⇒ Discrete-Time Markov Chains (DTMCs)

Introduction

Error models

AADL error models are finite automata enriched with probabilistic failures (and repairs).

Two kinds of error models can be distinguished:

- Discrete-time (cf. Marta Kwiatkowska's tutorial)
  - Failures and repairs are modeled by discrete probabilities
  - Instantaneous probabilistic decision to fail (or repair)
  - ⇒ Discrete-Time Markov Chains (DTMCs)
- Continuous-time (here)
  - Failures and repairs are modeled by continuous probabilities
  - Occur after a random duration
  - Mostly exponential distributions
  - ⇒ Continuous-Time Markov Chains (CTMCs)

Introduction

Error models

AADL error models are finite automata enriched with probabilistic failures (and repairs).

Two kinds of error models can be distinguished:

- Discrete-time (cf. Marta Kwiatkowska's tutorial)
    - Failures and repairs are modeled by discrete probabilities
    - Instantaneous probabilistic decision to fail (or repair)

    ⇒ Discrete-Time Markov Chains (DTMCs)
- Continuous-time (here)
    - Failures and repairs are modeled by continuous probabilities
    - Occur after a random duration
    - Mostly exponential distributions

    ⇒ Continuous-Time Markov Chains (CTMCs)

As error models are interweaved with non-probabilistic nominal models, in fact decision processes result. We consider deterministic decision processes.

Random Timing

Why Exponential Distributions?

- Are adequate for many real-life phenomena
  - the time until a radioactive particle decays
  - the time between successive car accidents
  - inter-arrival times of jobs, telephone calls in a fixed interval

Why Exponential Distributions?

- Are adequate for many real-life phenomena
  - the time until a radioactive particle decays
  - the time between successive car accidents
  - inter-arrival times of jobs, telephone calls in a fixed interval
- Are the continuous counterpart of the geometric distribution

Why Exponential Distributions?

- Are adequate for many real-life phenomena
  - the time until a radioactive particle decays
  - the time between successive car accidents
  - inter-arrival times of jobs, telephone calls in a fixed interval
- Are the continuous counterpart of the geometric distribution
- Heavily used in physics, performance, and reliability analysis

Why Exponential Distributions?

- Are adequate for many real-life phenomena
  - the time until a radioactive particle decays
  - the time between successive car accidents
  - inter-arrival times of jobs, telephone calls in a fixed interval
- Are the continuous counterpart of the geometric distribution
- Heavily used in physics, performance, and reliability analysis
- Can approximate general distributions arbitrarily closely

Negative Exponential Distributions

Definition (Exponential distribution)

The density of an exponentially distributed random variable Y with rate $\lambda \in \mathbb{R}_{>0}$ is:

$$f_Y(x) = \lambda \cdot e^{-\lambda \cdot x} \text{ for } x > 0 \text{ and } f_Y(x) = 0 \text{ otherwise}$$

Negative Exponential Distributions

Definition (Exponential distribution)

The density of an exponentially distributed random variable Y with rate $\lambda \in \mathbb{R}_{>0}$ is:

$$f_Y(x) = \lambda \cdot e^{-\lambda \cdot x} \text{ for } x > 0 \text{ and } f_Y(x) = 0 \text{ otherwise}$$

The cumulative distribution of Y is:

$$F_Y(d) = \int_0^d \lambda \cdot e^{-\lambda \cdot x} \, dx = 1 - e^{-\lambda \cdot d}.$$

Negative Exponential Distributions

Definition (Exponential distribution)

The density of an exponentially distributed random variable Y with rate $\lambda \in \mathbb{R}_{>0}$ is:

$$f_Y(x) = \lambda \cdot e^{-\lambda \cdot x} \text{ for } x > 0 \text{ and } f_Y(x) = 0 \text{ otherwise}$$

The cumulative distribution of Y is:

$$F_Y(d) = \int_0^d \lambda \cdot e^{-\lambda \cdot x}\, dx = 1 - e^{-\lambda \cdot d}.$$

The rate $\lambda \in \mathbb{R}_{>0}$ uniquely determines an exponential distribution.

Negative Exponential Distributions

Definition (Exponential distribution)

The density of an exponentially distributed random variable Y with rate $\lambda \in \mathbb{R}_{>0}$ is:

$$f_Y(x) = \lambda \cdot e^{-\lambda \cdot x} \text{ for } x > 0 \text{ and } f_Y(x) = 0 \text{ otherwise}$$

The cumulative distribution of Y is:

$$F_Y(d) = \int_0^d \lambda \cdot e^{-\lambda \cdot x} \, dx = 1 - e^{-\lambda \cdot d}.$$

The rate $\lambda \in \mathbb{R}_{>0}$ uniquely determines an exponential distribution.

Lemma (Variance and expectation)

If Y is exponentially distributed with rate $\lambda \in \mathbb{R}_{>0}$, then

$$\textit{expectation } E[Y] = \frac{1}{\lambda} \quad \textit{and} \quad \textit{variance } X[Y] = \frac{1}{\lambda^2}$$

Exponential PDF and CDF

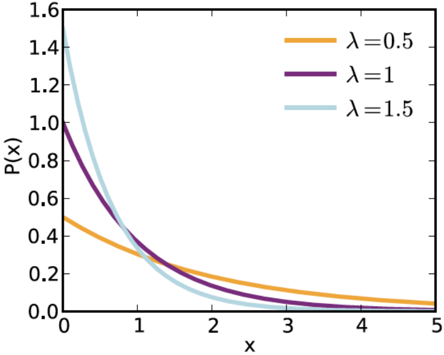

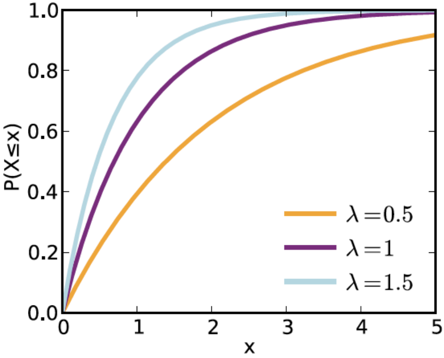

The higher λ, the faster the CDF approaches 1.

Continuous-Time Markov Chains

Definition (Continuous-Time Markov Chain)

A CTMC $\mathcal{C}$ is a tuple $(S, \mathbf{R}, \iota_{\mathrm{init}})$ with:

- S is a countable and non-empty set of states

Continuous-Time Markov Chains

Definition (Continuous-Time Markov Chain)

A CTMC $\mathcal{C}$ is a tuple $(S, \mathbf{R}, \iota_{\text{init}})$ with:

- S is a countable and non-empty set of states
- $\mathbf{R} : S \times S \to \mathbb{R}_{\geq 0}$ is a transition rate function (where $\sum_{s' \in S} \mathbf{R}(s, s') > 0$ for each $s \in S$)

Continuous-Time Markov Chains

Definition (Continuous-Time Markov Chain)

A CTMC $\mathcal{C}$ is a tuple $(S, \mathbf{R}, \iota_{\text{init}})$ with:

- S is a countable and non-empty set of states
- $\mathbf{R} : S \times S \to \mathbb{R}_{\geq 0}$ is a transition rate function (where $\sum_{s' \in S} \mathbf{R}(s, s') > 0$ for each $s \in S$)
- $\iota_{\text{init}} : S \to [0, 1]$ gives the initial distribution with $\sum_{s \in S} \iota_{\text{init}}(s) = 1$

Continuous-Time Markov Chains

Definition (Continuous-Time Markov Chain)

A CTMC $\mathcal{C}$ is a tuple $(S, \mathbf{R}, \iota_{\mathrm{init}})$ with:

- S is a countable and non-empty set of states
- $\mathbf{R} : S \times S \to \mathbb{R}_{\geq 0}$ is a transition rate function (where $\sum_{s' \in S} \mathbf{R}(s, s') > 0$ for each $s \in S$)
- $\iota_{\mathrm{init}} : S \to [0, 1]$ gives the initial distribution with $\sum_{s \in S} \iota_{\mathrm{init}}(s) = 1$

Additional notions

- The possible initial states are given by the set $\{\, s \in S \mid \iota_{\mathrm{init}}(s) > 0 \,\}$.

Continuous-Time Markov Chains

Definition (Continuous-Time Markov Chain)

A CTMC $\mathcal{C}$ is a tuple $(S, \mathbf{R}, \iota_{\text{init}})$ with:

- S is a countable and non-empty set of states
- $\mathbf{R} : S \times S \to \mathbb{R}_{\geq 0}$ is a transition rate function (where $\sum_{s' \in S} \mathbf{R}(s, s') > 0$ for each $s \in S$)
- $\iota_{\text{init}} : S \to [0, 1]$ gives the initial distribution with $\sum_{s \in S} \iota_{\text{init}}(s) = 1$

Additional notions

- The possible initial states are given by the set $\{ s \in S \mid \iota_{\text{init}}(s) > 0 \}$.
- The exit rate of a state, $\mathbf{R} : S \to \mathbb{R}_{>0}$, is determined by

$$\mathbf{R}(s) := \sum_{s' \in S} \mathbf{R}(s, s').$$

CTMC Semantics by Example

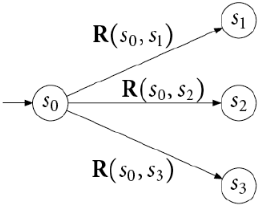

CTMC Semantics by Example

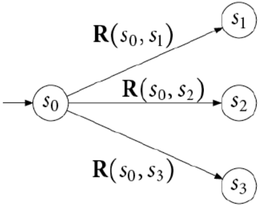

Example

CTMC semantics

- Transition $s \rightarrow s'$:= r.v. $X_{s,s'}$ with rate $\mathbf{R}(s, s')$

CTMC Semantics by Example

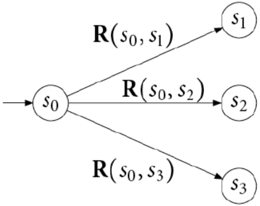

Example

CTMC semantics

- Transition $s \to s' :=$ r.v. $X_{s,s'}$ with rate $\mathbf{R}(s, s')$
- Probability to go from state s_0 to, say, state s_2 is:

$$Pr\{X_{s_0,s_2} \leq X_{s_0,s_1} \cap X_{s_0,s_2} \leq X_{s_0,s_3}\}$$

CTMC Semantics by Example

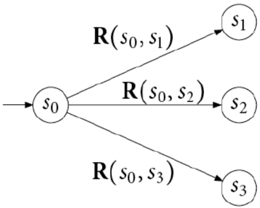

Example

CTMC semantics

- Transition $s \to s'$:= r.v. $X_{s,s'}$ with rate $\mathbf{R}(s, s')$
- Probability to go from state s_0 to, say, state s_2 is:

$$Pr\{X_{s_0,s_2} \leq X_{s_0,s_1} \cap X_{s_0,s_2} \leq X_{s_0,s_3}\}$$

$$= \frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0, s_1) + \mathbf{R}(s_0, s_2) + \mathbf{R}(s_0, s_3)}$$

CTMC Semantics by Example

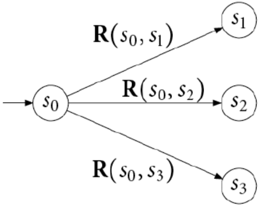

Example

CTMC semantics

- Transition $s \rightarrow s' :=$ r.v. $X_{s,s'}$ with rate $\mathbf{R}(s, s')$
- Probability to go from state s_0 to, say, state s_2 is:

$$Pr\{X_{s_0,s_2} \leq X_{s_0,s_1} \cap X_{s_0,s_2} \leq X_{s_0,s_3}\}$$

$$= \frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0, s_1) + \mathbf{R}(s_0, s_2) + \mathbf{R}(s_0, s_3)}$$

$$= \frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0)}$$

CTMC Semantics by Example

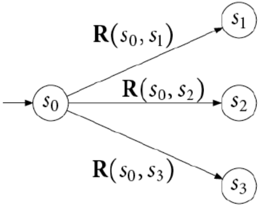

Example

CTMC semantics

- Transition $s \rightarrow s'$:= r.v. $X_{s,s'}$ with rate $\mathbf{R}(s, s')$
- Probability to go from state s_0 to, say, state s_2 is:

$$Pr\{X_{s_0,s_2} \leq X_{s_0,s_1} \cap X_{s_0,s_2} \leq X_{s_0,s_3}\}$$

$$= \frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0, s_1) + \mathbf{R}(s_0, s_2) + \mathbf{R}(s_0, s_3)}$$

$$= \frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0)}$$

- Probability of staying at most time t in s_0 is:

$$Pr\{\min\{X_{s_0,s_1}, X_{s_0,s_2}, X_{s_0,s_3}\} \leq t\}$$

CTMC Semantics by Example

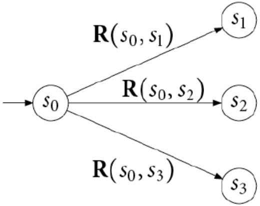

Example

CTMC semantics

- Transition $s \to s' :=$ r.v. $X_{s,s'}$ with rate $\mathbf{R}(s, s')$
- Probability to go from state s_0 to, say, state s_2 is:

$$
\begin{aligned}
& Pr\{X_{s_0,s_2} \leq X_{s_0,s_1} \cap X_{s_0,s_2} \leq X_{s_0,s_3}\} \\
= & \frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0, s_1) + \mathbf{R}(s_0, s_2) + \mathbf{R}(s_0, s_3)} \\
= & \frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0)}
\end{aligned}
$$

- Probability of staying at most time t in s_0 is:

$$
\begin{aligned}
& Pr\{\min\{X_{s_0,s_1}, X_{s_0,s_2}, X_{s_0,s_3}\} \leq t\} \\
= & 1 - e^{-(\mathbf{R}(s_0,s_1)+\mathbf{R}(s_0,s_2)+\mathbf{R}(s_0,s_3))\cdot t}
\end{aligned}
$$

CTMC Semantics by Example

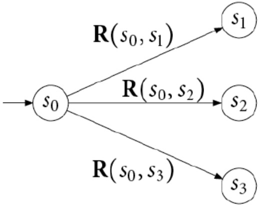

Example

CTMC semantics

- Transition $s \to s' :=$ r.v. $X_{s,s'}$ with rate $\mathbf{R}(s, s')$
- Probability to go from state s_0 to, say, state s_2 is:

$$Pr\{X_{s_0,s_2} \leq X_{s_0,s_1} \cap X_{s_0,s_2} \leq X_{s_0,s_3}\}$$

$$= \frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0, s_1) + \mathbf{R}(s_0, s_2) + \mathbf{R}(s_0, s_3)}$$

$$= \frac{\mathbf{R}(s_0, s_2)}{\mathbf{R}(s_0)}$$

- Probability of staying at most time t in s_0 is:

$$Pr\{\min\{X_{s_0,s_1}, X_{s_0,s_2}, X_{s_0,s_3}\} \leq t\}$$

$$= 1 - e^{-(\mathbf{R}(s_0,s_1)+\mathbf{R}(s_0,s_2)+\mathbf{R}(s_0,s_3))\cdot t}$$

$$= 1 - e^{-\mathbf{R}(s_0)\cdot t}$$

CTMC Semantics

Residence time distribution

The probability to take some outgoing transition from s in $[0, t]$ is:

$$\int_0^t \mathbf{R}(s)\cdot e^{-\mathbf{R}(s)\cdot x}\,dx \;=\; 1 - e^{-\mathbf{R}(s)\cdot t}$$

CTMC Semantics

Residence time distribution

The probability to take some outgoing transition from s in $[0, t]$ is:

$$\int_0^t \mathbf{R}(s) \cdot e^{-\mathbf{R}(s) \cdot x} \, dx \; = \; 1 - e^{-\mathbf{R}(s) \cdot t}$$

State-to-state timed transition probability

The probability to move from s to s' in $[0, t]$ is:

$$\frac{\mathbf{R}(s, s')}{\mathbf{R}(s)} \cdot \left(1 - e^{-\mathbf{R}(s) \cdot t}\right).$$

CTMCs are Omnipresent!

- Markovian queueing networks (Kleinrock 1975)
- Stochastic Petri nets (Molloy 1977)
- Stochastic activity networks (Meyer & Sanders 1985)
- Stochastic process algebra (Herzog *et al.*, Hillston 1993)
- Probabilistic input/output automata (Smolka *et al.* 1994)
- Calculi for biological systems (Priami *et al.*, Cardelli 2002)

CTMCs are Omnipresent!

- Markovian queueing networks (Kleinrock 1975)
- Stochastic Petri nets (Molloy 1977)
- Stochastic activity networks (Meyer & Sanders 1985)
- Stochastic process algebra (Herzog *et al.*, Hillston 1993)
- Probabilistic input/output automata (Smolka *et al.* 1994)
- Calculi for biological systems (Priami *et al.*, Cardelli 2002)

CTMCs are one of the most prominent models in performance analysis!

Paths in a CTMC

Definition (Timed paths)

(Timed) paths in a CTMC $\mathcal{C}$ are maximal (i.e., infinite) paths of alternating states and time instants:

$$\pi = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots$$

such that $s_i \in S$ and $t_i \in \mathbb{R}_{>0}$.

Here each t_i is the amount of time spent in state s_i.

Paths in a CTMC

Definition (Timed paths)

(Timed) paths in a CTMC $\mathcal{C}$ are maximal (i.e., infinite) paths of alternating states and time instants:

$$\pi \;=\; s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots$$

such that $s_i \in S$ and $t_i \in \mathbb{R}_{>0}$.

Here each t_i is the amount of time spent in state s_i.

Notations

- $Paths(s)$: set of paths starting in $s \in S$
- $Paths(\mathcal{C})$: set of paths starting in some initial state of $\mathcal{C}$
- $\pi[i] := s_i$: $(i{+}1)$-st state along timed path π
- $\pi@t$: state occupied in π at time $t \in \mathbb{R}_{\geq 0}$, i.e. $\pi@t := \pi[i]$ where i is the smallest index such that $\sum_{j=0}^{i} t_j > t$

Outline

Timed Reachability Events

Let CTMC $\mathcal{C}$ with (possibly infinite) state space S.

(Simple) timed reachability

Eventually reach a state in $G \subseteq S$ in the interval I:

Timed Reachability Events

Let CTMC $\mathcal{C}$ with (possibly infinite) state space S.

(Simple) timed reachability

Eventually reach a state in $G \subseteq S$ in the interval I:

$$\Diamond^{I}\, G \;=\; \{\, \pi \in Paths(\mathcal{C}) \mid \exists t \in I.\, \pi@t \in G \,\}$$

Timed Reachability Events

Let CTMC $\mathcal{C}$ with (possibly infinite) state space S.

(Simple) timed reachability

Eventually reach a state in $G \subseteq S$ in the interval I:

$$\Diamond^{I} G \;=\; \{\, \pi \in Paths(\mathcal{C}) \mid \exists t \in I.\, \pi@t \in G \,\}$$

Invariance, i.e., always stay in state in G in the interval I:

Timed Reachability Events

Let CTMC $\mathcal{C}$ with (possibly infinite) state space S.

(Simple) timed reachability

Eventually reach a state in $G \subseteq S$ in the interval I:

$$\Diamond^{I}\, G \;=\; \{\, \pi \in Paths(\mathcal{C}) \mid \exists t \in I.\, \pi@t \in G \,\}$$

Invariance, i.e., always stay in state in G in the interval I:

$$\Box^{I}\, G \;=\; \{\, \pi \in Paths(\mathcal{C}) \mid \forall t \in I.\, \pi@t \in G \,\} \;=\; \overline{\Diamond^{I}\, \overline{G}}.$$

Timed Reachability Events

Let CTMC $\mathcal{C}$ with (possibly infinite) state space S.

(Simple) timed reachability

Eventually reach a state in $G \subseteq S$ in the interval I:

$$\Diamond^{I} G \;=\; \{\, \pi \in Paths(\mathcal{C}) \mid \exists t \in I.\, \pi@t \in G \,\}$$

Invariance, i.e., always stay in state in G in the interval I:

$$\Box^{I} G \;=\; \{\, \pi \in Paths(\mathcal{C}) \mid \forall t \in I.\, \pi@t \in G \,\} \;=\; \overline{\Diamond^{I} \overline{G}}.$$

Constrained timed reachability

Or "reach-avoid" properties where states in $F \subseteq S$ are forbidden:

Timed Reachability Events

Let CTMC $\mathcal{C}$ with (possibly infinite) state space S.

(Simple) timed reachability

Eventually reach a state in $G \subseteq S$ in the interval I:

$$\Diamond^I G = \{ \pi \in Paths(\mathcal{C}) \mid \exists t \in I.\, \pi@t \in G \}$$

Invariance, i.e., always stay in state in G in the interval I:

$$\Box^I G = \{ \pi \in Paths(\mathcal{C}) \mid \forall t \in I.\, \pi@t \in G \} = \overline{\Diamond^I \overline{G}}.$$

Constrained timed reachability

Or "reach-avoid" properties where states in $F \subseteq S$ are forbidden:

$$\overline{F}\, \mathrm{U}^I G = \{ \pi \in Paths(\mathcal{C}) \mid \exists t \in I.\, \pi@t \in G \wedge \forall d < t.\, \pi@d \notin F \}$$

Measurability

Theorem (Measurability theorem)

Events $\Diamond^{I}\, G$, $\Box^{I}\, G$, *and* $\overline{F}\, \mathrm{U}^{I}\, G$ *are measurable on any CTMC.*

Timed Reachability Probabilities in Finite CTMCs

Problem statement

Let $\mathcal{C}$ be a CTMC with finite state space S, $s \in S$, $t \in \mathbb{R}_{\geq 0}$ and $G \subseteq S$.

Timed Reachability Probabilities in Finite CTMCs

Problem statement

Let $\mathcal{C}$ be a CTMC with finite state space S, $s \in S$, $t \in \mathbb{R}_{\geq 0}$ and $G \subseteq S$.

Aim: compute $Pr(s \models \lozenge^{\leq t}\, G) \;=\; Pr_s\{\, \pi \in Paths(s) \mid \pi \models \lozenge^{\leq t}\, G \,\}$

where Pr_s is the probability measure in CTMC $\mathcal{C}$ with single initial state s.

Timed Reachability Probabilities in Finite CTMCs

Problem statement

Let $\mathcal{C}$ be a CTMC with finite state space S, $s \in S$, $t \in \mathbb{R}_{\geq 0}$ and $G \subseteq S$.

Aim: compute $Pr(s \models \Diamond^{\leq t} G) = Pr_s\{ \pi \in Paths(s) \mid \pi \models \Diamond^{\leq t} G \}$

where Pr_s is the probability measure in CTMC $\mathcal{C}$ with single initial state s.

Characterisation of timed reachability probabilities

- Let function $x_s(t) = Pr(s \models \Diamond^{\leq t} G)$ for any state s
  - if G is not reachable from s, then $x_s(t) = 0$ for all t

Timed Reachability Probabilities in Finite CTMCs

Problem statement

Let $\mathcal{C}$ be a CTMC with finite state space S, $s \in S$, $t \in \mathbb{R}_{\geq 0}$ and $G \subseteq S$.

Aim: compute $Pr(s \models \lozenge^{\leq t} G) = Pr_s\{ \pi \in Paths(s) \mid \pi \models \lozenge^{\leq t} G \}$

where Pr_s is the probability measure in CTMC $\mathcal{C}$ with single initial state s.

Characterisation of timed reachability probabilities

- Let function $x_s(t) = Pr(s \models \lozenge^{\leq t} G)$ for any state s
  - if G is not reachable from s, then $x_s(t) = 0$ for all t
  - if $s \in G$ then $x_s(t) = 1$ for all t

Timed Reachability Probabilities in Finite CTMCs

Problem statement

Let $\mathcal{C}$ be a CTMC with finite state space S, $s \in S$, $t \in \mathbb{R}_{\geq 0}$ and $G \subseteq S$.

Aim: compute $Pr(s \models \lozenge^{\leq t} G) = Pr_s\{\pi \in Paths(s) \mid \pi \models \lozenge^{\leq t} G\}$

where Pr_s is the probability measure in CTMC $\mathcal{C}$ with single initial state s.

Characterisation of timed reachability probabilities

- Let function $x_s(t) = Pr(s \models \lozenge^{\leq t} G)$ for any state s
  - if G is not reachable from s, then $x_s(t) = 0$ for all t
  - if $s \in G$ then $x_s(t) = 1$ for all t
- For any state $s \in Pre^*(G) \setminus G$:

$$x_s(t) = \int_0^t \sum_{s' \in S} \underbrace{\mathbf{R}(s,s') \cdot e^{-\mathbf{R}(s) \cdot x}}_{\text{probability to move to state } s' \text{ at time } x} \cdot \underbrace{x_{s'}(t-x)}_{\text{probability to fulfill } \lozenge^{\leq t-x} G \text{ from } s'} \, dx$$

Solving Reachability Problems

Reachability probabilities in finite CTMCs

Solve a system of linear equations (using some efficient techniques).

Solving Reachability Problems

Reachability probabilities in finite CTMCs

Solve a system of linear equations (using some efficient techniques).

Timed reachability probabilities in finite CTMCs

Solve a system of Volterra integral equations.

Solving Reachability Problems

Reachability probabilities in finite CTMCs

Solve a system of linear equations (using some efficient techniques).

Timed reachability probabilities in finite CTMCs

Solve a system of Volterra integral equations.

This is in general non-trivial, inefficient, and has several pitfalls such as numerical stability.

Solving Reachability Problems

Reachability probabilities in finite CTMCs

Solve a system of linear equations (using some efficient techniques).

Timed reachability probabilities in finite CTMCs

Solve a system of Volterra integral equations.

This is in general non-trivial, inefficient, and has several pitfalls such as numerical stability.

Solution

Reduce the problem of computing $Pr(s \models \Diamond^{\leq t} G)$ to an alternative problem for which well-known efficient techniques exist: computing transient probabilities.

Timed Reachability Probabilities = Transient Probabilities

Aim

Compute $Pr(s \models \Diamond^{\leq t} G)$ in CTMC $\mathcal{C}$.

Timed Reachability Probabilities = Transient Probabilities

Aim

Compute $Pr(s \models \Diamond^{\leq t} G)$ in CTMC $\mathcal{C}$. Observe that once a path π reaches G within time t, then the remaining behaviour along π is not important.

Timed Reachability Probabilities = Transient Probabilities

Aim

Compute $Pr(s \models \Diamond^{\leq t} G)$ in CTMC $\mathcal{C}$. Observe that once a path π reaches G within time t, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Timed Reachability Probabilities = Transient Probabilities

Aim

Compute $Pr(s \models \Diamond^{\leq t} G)$ in CTMC $\mathcal{C}$. Observe that once a path π reaches G within time t, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Let CTMC $\mathcal{C} = (S, \mathbf{R}, \iota_{\text{init}})$ and $G \subseteq S$. Then CTMC $\mathcal{C}[G] := (S, \mathbf{R}_G, \iota_{\text{init}})$ with

$$\mathbf{R}_G(s,t) := \begin{cases} \mathbf{R}(s,t) & \text{if } s \notin G \\ \mathbf{R}(s) & \text{if } s \in G, t = s \\ 0 & \text{if } s \in G, t \neq s \end{cases}$$

Aim

Compute $Pr(s \models \Diamond^{\leq t} G)$ in CTMC $\mathcal{C}$. Observe that once a path π reaches G within time t, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Let CTMC $\mathcal{C} = (S, \mathbf{R}, \iota_{\text{init}})$ and $G \subseteq S$. Then CTMC $\mathcal{C}[G] := (S, \mathbf{R}_G, \iota_{\text{init}})$ with

$$\mathbf{R}_G(s,t) := \begin{cases} \mathbf{R}(s,t) & \text{if } s \notin G \\ \mathbf{R}(s) & \text{if } s \in G, t = s \\ 0 & \text{if } s \in G, t \neq s \end{cases}$$

(thus: all outgoing transitions of $s \in G$ are replaced by a single self-loop at s)

Timed Reachability Probabilities = Transient Probabilities

Aim

Compute $Pr(s \models \Diamond^{\leq t} G)$ in CTMC $\mathcal{C}$. Observe that once a path π reaches G within time t, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Let CTMC $\mathcal{C} = (S, \mathbf{R}, \iota_{\text{init}})$ and $G \subseteq S$. Then CTMC $\mathcal{C}[G] := (S, \mathbf{R}_G, \iota_{\text{init}})$ with

$$\mathbf{R}_G(s,t) := \begin{cases} \mathbf{R}(s,t) & \text{if } s \notin G \\ \mathbf{R}(s) & \text{if } s \in G, t = s \\ 0 & \text{if } s \in G, t \neq s \end{cases}$$

(thus: all outgoing transitions of $s \in G$ are replaced by a single self-loop at s)

Lemma

$$\underbrace{Pr(s \models \Diamond^{\leq t} G)}_{\textit{timed reachability in } \mathcal{C}} \quad =$$

Aim

Compute $Pr(s \models \Diamond^{\leq t} G)$ in CTMC $\mathcal{C}$. Observe that once a path π reaches G within time t, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Let CTMC $\mathcal{C} = (S, \mathbf{R}, \iota_{\text{init}})$ and $G \subseteq S$. Then CTMC $\mathcal{C}[G] := (S, \mathbf{R}_G, \iota_{\text{init}})$ with

$$\mathbf{R}_G(s,t) := \begin{cases} \mathbf{R}(s,t) & \text{if } s \notin G \\ \mathbf{R}(s) & \text{if } s \in G, t = s \\ 0 & \text{if } s \in G, t \neq s \end{cases}$$

(thus: all outgoing transitions of $s \in G$ are replaced by a single self-loop at s)

Lemma

$$\underbrace{Pr(s \models \Diamond^{\leq t} G)}_{\textit{timed reachability in } \mathcal{C}} = \underbrace{Pr(s \models \Diamond^{= t} G)}_{\textit{timed reachability in } \mathcal{C}[G]} =$$

Timed Reachability Probabilities = Transient Probabilities

Aim

Compute $Pr(s \models \Diamond^{\leq t} G)$ in CTMC $\mathcal{C}$. Observe that once a path π reaches G within time t, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Let CTMC $\mathcal{C} = (S, \mathbf{R}, \iota_{\text{init}})$ and $G \subseteq S$. Then CTMC $\mathcal{C}[G] := (S, \mathbf{R}_G, \iota_{\text{init}})$ with

$$\mathbf{R}_G(s,t) := \begin{cases} \mathbf{R}(s,t) & \text{if } s \notin G \\ \mathbf{R}(s) & \text{if } s \in G, t = s \\ 0 & \text{if } s \in G, t \neq s \end{cases}$$

(thus: all outgoing transitions of $s \in G$ are replaced by a single self-loop at s)

Lemma

$$\underbrace{Pr(s \models \Diamond^{\leq t} G)}_{\textit{timed reachability in } \mathcal{C}} = \underbrace{Pr(s \models \Diamond^{= t} G)}_{\textit{timed reachability in } \mathcal{C}[G]} = \underbrace{\underline{p}(t) \textit{ with } \underline{p}(0) = \mathbf{1}_s}_{\textit{transient prob. in } \mathcal{C}[G]} .$$

Outline

Approach

Approach in the COMPASS toolset

1. Weave the nominal behaviour and error model (model extension)
2. The semantics yields an continuous-time decision process

Approach

Approach in the COMPASS toolset

1. Weave the nominal behaviour and error model (model extension)
2. The semantics yields an continuous-time decision process
3. Apply (BDD-based) bisimulation minimisation to this process (not shown here)
4. Mostly this yields a CTMC
5. Verify it using the techniques explained before

Approach

Approach in the COMPASS toolset

1. Weave the nominal behaviour and error model (model extension)
2. The semantics yields an continuous-time decision process
3. Apply (BDD-based) bisimulation minimisation to this process (not shown here)
4. Mostly this yields a CTMC
5. Verify it using the techniques explained before
6. For timed reachability, cover the entire range from 0 to t

Current work is on directly analysing the stochastic decision process

Example: Sensor-Filter Data Acquisition System

- models a data acquisition system
- the value is read by a sensor, filtered by a filter, and returned as output
- two redundant sensors `sensor1` and `sensor2`
- two redundant filters `filter1` and `filter2`
- a central `Monitor` detects anomalies in the output of either the sensors or the filters, and issues a system reconfiguration (`switchS` or resp. `switchF`) whenever needed

Acquisition system

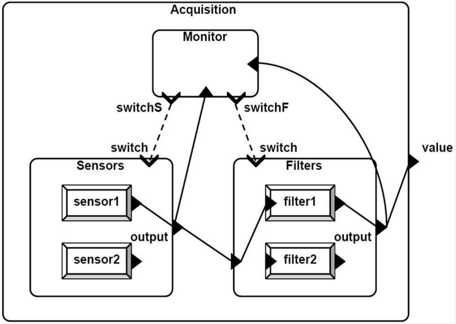

Modeling Sensor and Filter Errors

Sensor error model:

- two faulty states: Drifted and Dead
- Poisson distribution

Filter error model:

- two faulty states: Degrade and Dead
- Poisson distribution

Sensor error model

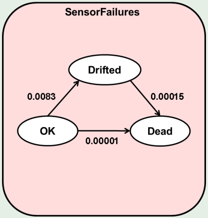

Filter error model

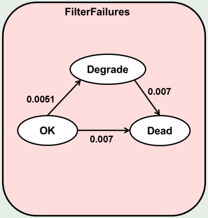

Sensor Error Model in AADL

```
error model SensorFailures
  features
    OK: initial state;
    Drifted: error state;
    Dead: error state;
end SensorFailures;

error model implementation SensorFailures.Impl
  events
    drift: error event occurrence poisson 0.083;
    die: error event occurrence poisson 0.00001;
    dieByDrift: error event
      occurrence poisson 0.00015;
  transitions
    OK -[ die ]-> Dead;
    OK -[ drift ]-> Drifted;
    Drifted -[ dieByDrift ]-> Dead;
end SensorFailures.Impl;
```

Sensor error model

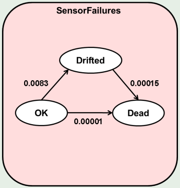

Defining Fault Injections

- in state `Dead`, the output of the sensor is stuck at 15
- in state `Dead`, the output of the filter is stuck at 0

Fault injections

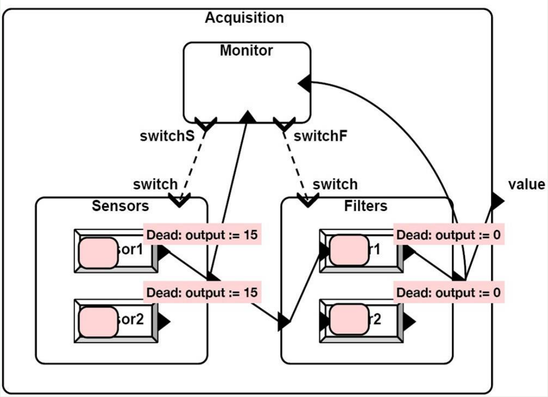

Properties of Interest

Some properties of interest

- A filter or a sensor fails
- A sensor fails
  - `sensor1` fails
  - `sensor2` fails
- Filters fail twice
- Monitor reacts to filter failures
- Sensors or filters die within 76 hours
- `sensor2` fails before `filter2` within 512 hours

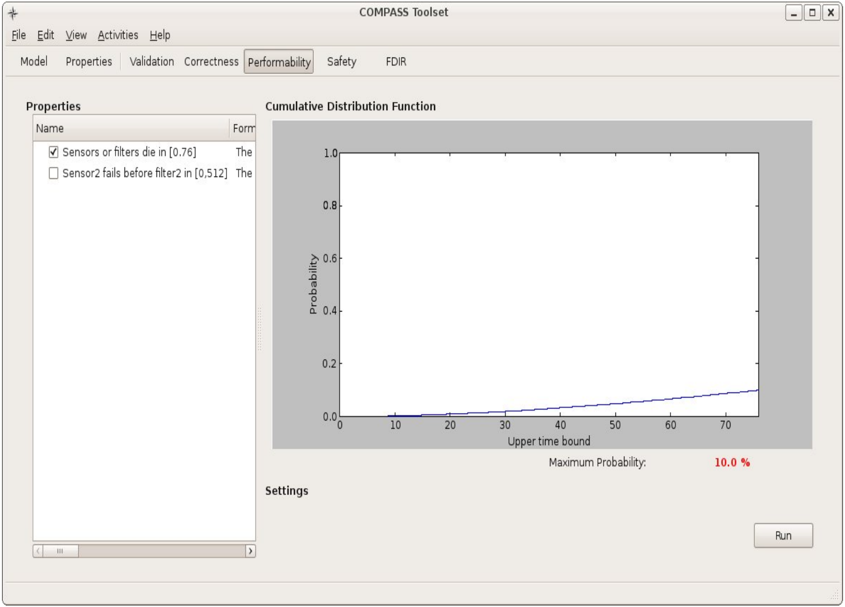
COMPASS Toolset
File Edit View Activities Help
Model Properties Validation Correctness Performability Safety FDIR
Properties
Name
Form
Sensors or filters die in [0.76]
The
Sensor2 fails before filter2 in [0,512]
The
Cumulative Distribution Function
Probability
Upper time bound
1.0
0.8
0.6
0.4
0.2
0.0
0
10
20
30
40
50
60
70
Maximum Probability:
10.0 %
Settings
Run

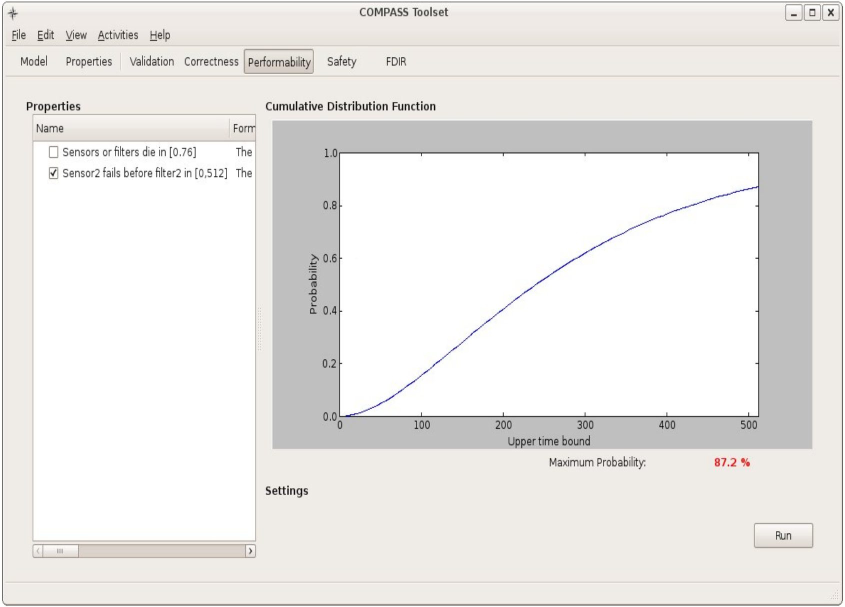
COMPASS Toolset
File Edit View Activities Help
Model Properties Validation Correctness Performability Safety FDIR
Properties
Name
Form
Sensors or filters die in [0.76]
The
Sensor2 fails before filter2 in [0,512]
The
Cumulative Distribution Function
Probability
0.0
0.2
0.4
0.6
0.8
1.0
0
100
200
300
400
500
Upper time bound
Maximum Probability:
87.2 %
Settings
Run

Outline

Further Information

- Probabilistic model checking (Baier et. al, CACM 2011)
  (Kwiatkowska et. al, SFM 2011)
  (Baier & Katoen, Principles of Model Checking)
- CTMC model checking (Baier et. al, IEEE TSE 2003)
- Probabilistic bisimulation (Larsen & Skou, Inf. Comp 1989)
  (Kemeny & Snell, 1960)
  (Buchholz, Appl. Prob. 1994)
- Bisimulation minimisation (Derisavi et. al, IPL 2005)
  (Valmari & Franceschinis, TACAS 2010)
- Stochastic decision processes (Guck et. al, NFM 2012)