

iFM & ABZ 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
Joost-Pieter Katoen & Thomas Noll



Fondazione Bruno Kessler
Centre for Scientific and Technological Research
Marco Bozzano & Alessandro Cimatti

iFM & ABZ 2012; June 18, 2012; Pisa, Italy

- 1 Introduction
- 2 Verifying Discrete-Time Markov Chains
- 3 Verifying Continuous-Time Markov Chains
- 4 Tool Demo

- 1 Introduction
- 2 Verifying Discrete-Time Markov Chains
- 3 Verifying Continuous-Time Markov Chains
- 4 Tool Demo

Error models

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

Two kinds of error models can be distinguished:

- Discrete-time

- Failures and repairs are modeled by discrete probabilities
- Instantaneous probabilistic decision to fail (or repair)

⇒ discrete-time Markov chains (DTMCs)

- Continuous-time

- Failures and repairs are modeled by continuous probabilities
- Mostly exponential distributions
- Failures and repairs occur after a random duration

⇒ 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.

- 1 Introduction
- 2 Verifying Discrete-Time Markov Chains
- 3 Verifying Continuous-Time Markov Chains
- 4 Tool Demo

Let's start easy

Discrete-time Markov chain

A DTMC \mathcal{D} is a tuple $(S, \mathbf{P}, \ell_{\text{init}})$ with:

- S is a countable nonempty set of states
- $\mathbf{P} : S \times S \rightarrow [0, 1]$, transition probability function s.t. $\sum_{s'} \mathbf{P}(s, s') = 1$
- $\ell_{\text{init}} : S \rightarrow [0, 1]$, the initial distribution with $\sum_{s \in S} \ell_{\text{init}}(s) = 1$


Initial states

- the set $\{s \in S \mid \ell_{\text{init}}(s) > 0\}$ are the possible initial states.

Paths

Paths through a DTMC are infinite sequence of states.

Simulating a die by a fair coin [Knuth & Yao]



Heads = “go left”; tails = “go right”. Does this DTMC adequately model a fair six-sided die?

Some events of interest

(Simple) reachability

Eventually reach a state in $G \subseteq S$. Formally:

$$\Diamond G = \{ \pi \in Paths(\mathcal{D}) \mid \exists i \in \mathbb{N}. \pi[i] \in G \}$$

Invariance, i.e., always stay in state in G :

$$\Box G = \{ \pi \in Paths(\mathcal{D}) \mid \forall i \in \mathbb{N}. \pi[i] \in G \} = \overline{\overline{\Diamond \bar{G}}}.$$

Constrained reachability

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

$$\bar{F} \cup G = \{ \pi \in Paths(\mathcal{D}) \mid \exists i \in \mathbb{N}. \pi[i] \in G \wedge \forall j < i. \pi[j] \notin F \}$$

In a similar way, $\Box \Diamond G$ and $\Diamond \Box G$ are defined.

Reachability probabilities in finite DTMCs

Problem statement

Let \mathcal{D} be a DTMC with finite state space S , $s \in S$ and $G \subseteq S$.


Aim: determine $Pr(s \models \Diamond G) = Pr_s\{\pi \in Paths(s) \mid \pi \in \Diamond G\}$.

Characterisation of reachability probabilities

- Let variable $x_s = Pr(s \models \Diamond G)$ for any state s
 - if G is not reachable from s , then $x_s = 0$
 - if $s \in G$ then $x_s = 1$
- For any state $s \in Pre^*(G) \setminus G$:

$$x_s = \underbrace{\sum_{t \in S \setminus G} P(s, t) \cdot x_t}_{\text{reach } G \text{ via } t \in S \setminus G} + \underbrace{\sum_{u \in G} P(s, u)}_{\text{reach } G \text{ in one step}}$$

Reachability probabilities: Knuth's die



- Consider the event $\Diamond 4$
- Using the previous characterisation we obtain:

$$x_1 = x_2 = x_3 = x_5 = x_6 = 0 \text{ and } x_4 = 1$$

$$x_{s_1} = x_{s_3} = x_{s_4} = 0$$

$$x_{s_0} = \frac{1}{2}x_{s_1} + \frac{1}{2}x_{s_2}$$

$$x_{s_2} = \frac{1}{2}x_{s_5} + \frac{1}{2}x_{s_6}$$

$$x_{s_5} = \frac{1}{2}x_5 + \frac{1}{2}x_4$$

$$x_{s_6} = \frac{1}{2}x_{s_2} + \frac{1}{2}x_6$$

- Gaussian elimination yields:

$$x_{s_5} = \frac{1}{2}, x_{s_2} = \frac{1}{3}, x_{s_6} = \frac{1}{6}, \text{ and } x_{s_0} = \frac{1}{6}$$

Linear equation system

Reachability probabilities as linear equation system

- Let $S_? = Pre^*(G) \setminus G$, the states that can reach G by > 0 steps
- $\mathbf{A} = (\mathbf{P}(s, t))_{s, t \in S_?}$, the transition probabilities in $S_?$
- $\mathbf{b} = (b_s)_{s \in S_?}$, the probs to reach G in 1 step, i.e., $b_s = \sum_{u \in G} \mathbf{P}(s, u)$

Then: $\mathbf{x} = (x_s)_{s \in S_?}$ with $x_s = Pr(s \models \Diamond G)$ is the **unique** solution of:

$$\mathbf{x} = \mathbf{A} \cdot \mathbf{x} + \mathbf{b} \quad \text{or} \quad (\mathbf{I} - \mathbf{A}) \cdot \mathbf{x} = \mathbf{b}$$

where \mathbf{I} is the identity matrix of cardinality $|S_?| \times |S_?|$.

Repeated reachability and persistence

Long-run theorem

Almost surely any finite DTMC eventually reaches a BSCC and visits all its states infinitely often.

Repeated reachability = Reachability

For finite DTMC with state space S , $G \subseteq S$, and $s \in S$:

$$Pr(s \models \Box \Diamond G) = Pr(s \models \Diamond U)$$

where U is the union of all BSCCs T with $T \cap G \neq \emptyset$.

Persistency = Reachability

For finite DTMC with state space S , $G \subseteq S$, and $s \in S$:

$$Pr(s \models \Diamond \Box G) = Pr(s \models \Diamond U)$$

where U is the union of all BSCCs T with $T \subseteq G$.

Probabilistic bisimulation: intuition

Intuition

- Strong bisimulation is used to **compare** labeled transition systems.
 - Strongly bisimilar states exhibit the same step-wise behaviour.
 - Our aim: adapt bisimulation to discrete-time Markov chains.
 - This yields a probabilistic variant of strong bisimulation.
-
- When do two DTMC states exhibit the same step-wise behaviour?
 - **Key: if their transition probability for each equivalence class coincides.**

Probabilistic bisimulation

Probabilistic bisimulation

[Larsen & Skou, 1989]

Let $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}})$ be a DTMC and $R \subseteq S \times S$ an equivalence. R is a *probabilistic bisimulation* on S if for any $(s, t) \in R$:

$$\mathbf{P}(s, C) = \mathbf{P}(t, C) \text{ for all equivalence classes } C \in S/R.$$

where $\mathbf{P}(s, C) = \sum_{s' \in C} \mathbf{P}(s, s')$.

For states in R , the probability of moving by a single transition to some equivalence class is equal.

Probabilistic bisimilarity

Let \mathcal{D} be a DTMC and s, t states in \mathcal{D} . Then: s is *probabilistically bisimilar* to t , denoted $s \sim_p t$, if there *exists* a probabilistic bisimulation R with $(s, t) \in R$.

Probabilistic bisimulation

Probabilistic bisimulation


Let $\mathcal{D} = (S, \mathbf{P}, \ell_{\text{init}})$ be a DTMC and $R \subseteq S \times S$ an equivalence. Then: R is a *probabilistic bisimulation* on S if for any $(s, t) \in R$:

$$\mathbf{P}(s, C) = \mathbf{P}(t, C) \text{ for all equivalence classes } C \in S/R.$$

Remarks

As opposed to bisimulation on states in transition systems, **any** probabilistic bisimulation is an equivalence.

Example



Quotient under \sim_p

Quotient DTM under \sim_p

For $\mathcal{D} = (S, \mathbf{P}, \iota_{\text{init}})$ and probabilistic bisimulation $\sim_p \subseteq S \times S$ let

$$\mathcal{D}/\sim_p = (S', \mathbf{P}', \iota'_{\text{init}}), \quad \text{the quotient of } \mathcal{D} \text{ under } \sim_p$$

where

- $S' = S/\sim_p = \{[s]_{\sim_p} \mid s \in S\}$ with $[s]_{\sim_p} = \{s' \in S \mid s \sim_p s'\}$
- $\mathbf{P}'([s]_{\sim_p}, [s']_{\sim_p}) = \mathbf{P}(s, [s']_{\sim_p})$
- $\iota'_{\text{init}}([s]_{\sim_p}) = \sum_{s' \in [s]_{\sim_p}} \iota_{\text{init}}(s')$

Remarks


The transition probability from $[s]_{\sim_p}$ to $[t]_{\sim_p}$ equals $\mathbf{P}(s, [t]_{\sim_p})$. This is well-defined as $\mathbf{P}(s, C) = \mathbf{P}(s', C)$ for all $s \sim_p s'$ and all \sim_p equivalence classes C .

- Come-out roll:


- 7 or 11: win
- 2, 3, or 12: lose
- else: roll again

- Next roll(s):

- 7: lose
- point: win
- else: roll again



Bisimulation quotient DTMC of Craps



Bisimulation preserves reachability probabilities

Let \mathcal{D} be a DTMC and s, t states in \mathcal{D} . Then:

$$s \sim_p t \text{ implies } Pr(s \models \Diamond G) = Pr(t \models \Diamond G)$$

for every \sim_p -closed set of states $G \subseteq S$.

Remarks

$s \sim_p t$ implies that (repeated) reachability probabilities, and persistence probabilities for s and t coincide.

In fact, \sim_p coincides with probabilistic CTL equivalence.

The coarsest bisimulation quotient can be computed in $\mathcal{O}(m \cdot \log n)$

- 1 Introduction
- 2 Verifying Discrete-Time Markov Chains
- 3 Verifying Continuous-Time Markov Chains**
- 4 Tool Demo

Random timing



Negative exponential distribution

Density of exponential distribution

The density of an *exponentially distributed* r.v. Y with *rate* $\lambda \in \mathbb{R}_{>0}$ is:

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

The cumulative distribution of r.v. Y with rate $\lambda \in \mathbb{R}_{>0}$ is:

$$F_Y(d) = \int_0^d \lambda \cdot e^{-\lambda \cdot x} dx = [-e^{-\lambda \cdot x}]_0^d = 1 - e^{-\lambda \cdot d}.$$



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

Variance and expectation

Let r.v. Y be exponentially distributed with rate $\lambda \in \mathbb{R}_{>0}$. Then:

$$\text{Expectation } E[Y] = \frac{1}{\lambda} \text{ and variance } X[Y] = \frac{1}{\lambda^2}$$

Exponential pdf and cdf




The higher λ , the faster the cdf approaches 1.

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
- Yield a *maximal entropy* if only the mean is known

Continuous-time Markov chains

A CTMC is a DTMC with an *exit rate* function $r : S \rightarrow \mathbb{R}_{>0}$ where $r(s)$ is the rate of an exponential distribution.




$$r(s) = 25, r(t) = 4, r(u) = 2 \text{ and } r(v) = 100$$

Example: a classical perspective


A CTMC is a DTMC with an *exit rate* function $r : S \rightarrow \mathbb{R}_{>0}$ where $r(s)$ is the rate of an exponential distribution.

A CTMC is a DTMC where transition probability function \mathbf{P} is replaced by a *transition rate* function \mathbf{R} . We have $\mathbf{R}(s, s') = \mathbf{P}(s, s') \cdot r(s)$.



$$r(s) = 25, r(t) = 4, r(u) = 2 \text{ and } r(v) = 100$$

CTMC semantics by example



CTMC semantics

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

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

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

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

State-to-state timed transition probability

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

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

Residence time distribution

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

$$\int_0^t r(s) \cdot e^{-r(s) \cdot x} dx = 1 - e^{-r(s) \cdot t}$$

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

Timed paths

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

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

such that $s_j \in S$ and $t_j \in \mathbb{R}_{>0}$.

Time instant t_j is the amount of time spent in state s_j .

Notations

- Let $\pi[i] := s_i$ denote the $(i+1)$ -st state along the timed path π .
- Let $\pi@t$ be the 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 $\sum_{j=0}^i t_j > t$.

Zeno theorem

Zeno path

Path $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \xrightarrow{t_2} s_3 \dots$ is called **Zeno**^a if $\sum_i t_i$ converges.

^aZeno of Elea (490-430 BC), philosopher, famed for his paradoxes.

Example

$$s_0 \xrightarrow{1} s_1 \xrightarrow{\frac{1}{2}} s_2 \xrightarrow{\frac{1}{4}} s_3 \dots s_i \xrightarrow{\frac{1}{2^i}} s_{i+1} \dots$$

In timed automata, such executions are typically excluded from the analysis.

Zeno theorem

For all states s in any CTMC, $Pr\{\pi \in Paths(s) \mid \pi \text{ is Zeno}\} = 0$.

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 . Formally:

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

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:

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

Measurability theorem

Events $\diamond' G$, $\square' G$, and $\overline{F} U' G$ 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 $G \subseteq S$.

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

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
 - if $s \in G$ then $x_s(t) = 1$ for all t
- For any state $s \in Pre^*(G) \setminus G$:

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

Reachability probabilities in finite DTMCs and 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} . Observe that once a path π reaches G within t time, then the remaining behaviour along π is not important. This suggests to make all states in G absorbing.

Let CTMC $\mathcal{C} = (S, \mathbf{P}, r, \iota_{\text{init}})$ and $G \subseteq S$. The CTMC $\mathcal{C}[G] = (S, \mathbf{P}_G, r, \iota_{\text{init}})$ with $\mathbf{P}_G(s, t) = \mathbf{P}(s, t)$ if $s \notin G$ and $\mathbf{P}_G(s, s) = 1$ if $s \in G$.

All outgoing transitions of $s \in G$ are replaced by a single self-loop at s .

Lemma

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

Transient distribution theorem

Theorem: transient distribution as ordinary differential equation

The **transient** probability vector $\underline{p}(t) = (p_{s_1}(t), \dots, p_{s_k}(t))$ satisfies:

$$\underline{p}'(t) = \underline{p}(t) \cdot (\mathbf{R} - \mathbf{r}) \quad \text{given} \quad \underline{p}(0)$$

where \mathbf{r} is the diagonal matrix of vector \underline{r} .

Solution technique:

Transform the CTMC (again), and then truncate a Taylor-MacLaurin expansion. This yields a **polynomial-time approximation** algorithm.

Probabilistic bisimulation

[Buchholz, 1994]

Let $\mathcal{C} = (S, \mathbf{P}, r, \iota_{\text{init}})$ be a CTMC and $R \subseteq S \times S$ an *equivalence*. Then: R is a *probabilistic bisimulation* on S if for any $(s, t) \in R$:

- ① $r(s) = r(t)$, and
- ② $\mathbf{P}(s, C) = \mathbf{P}(t, C)$ for all equivalence classes $C \in S/R$

The last two conditions amount to $\mathbf{R}(s, C) = \mathbf{R}(t, C)$ for all equivalence classes $C \in S/R$.

Probabilistic bisimilarity

Let \mathcal{C} be a CTMC and s, t states in \mathcal{C} . Then: s is *probabilistically bisimilar* to t , denoted $s \sim_m t$, if there *exists* a probabilistic bisimulation R with $(s, t) \in R$.

Bisimulation preserves timed reachability probabilities

Let \mathcal{C} be a CTMC and s, t states in \mathcal{C} . Then:

$$s \sim_m t \text{ implies } Pr(s \models \diamond^{\leq t} G) = Pr(t \models \diamond^{\leq t} G)$$

for every \sim_p -closed set of states $G \subseteq S$ and any t .

Remarks

$s \sim_p t$ implies that (repeated) reachability probabilities, and persistence probabilities for s and t coincide.

In fact, \sim_p coincides with probabilistic CTL equivalence.

The coarsest bisimulation quotient can be computed in $\mathcal{O}(m \cdot \log n)$

- 1 Introduction
- 2 Verifying Discrete-Time Markov Chains
- 3 Verifying Continuous-Time Markov Chains
- 4 Tool Demo

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
- 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

Demo Example: Sensor-Filter Acquisition System

Redundant Sensor-Filter Example: Nominal Model

- models a value 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 either the output of the sensor or the filter, and issues a system reconfiguration (**switchS** or **switchF**) whenever needed


Acquisition System



Sensor Error model:

- two faulty states: **Drifted** and **Dead**
- poisson distribution


Sensor Error Model



Filter Error model:

- two faulty states: **Degrade** and **Dead**
- poisson distribution

Filter Error Model




Sensor: SLIM Error Model

```
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



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

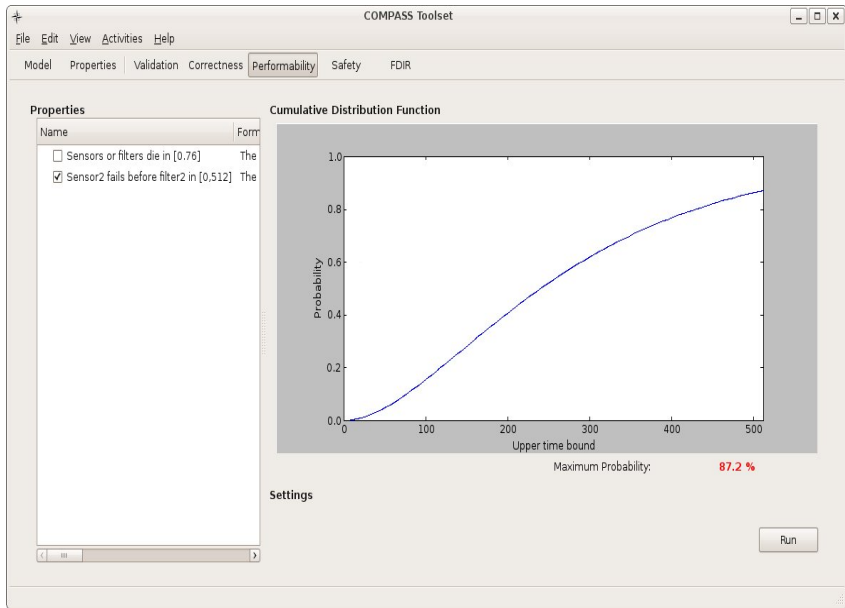


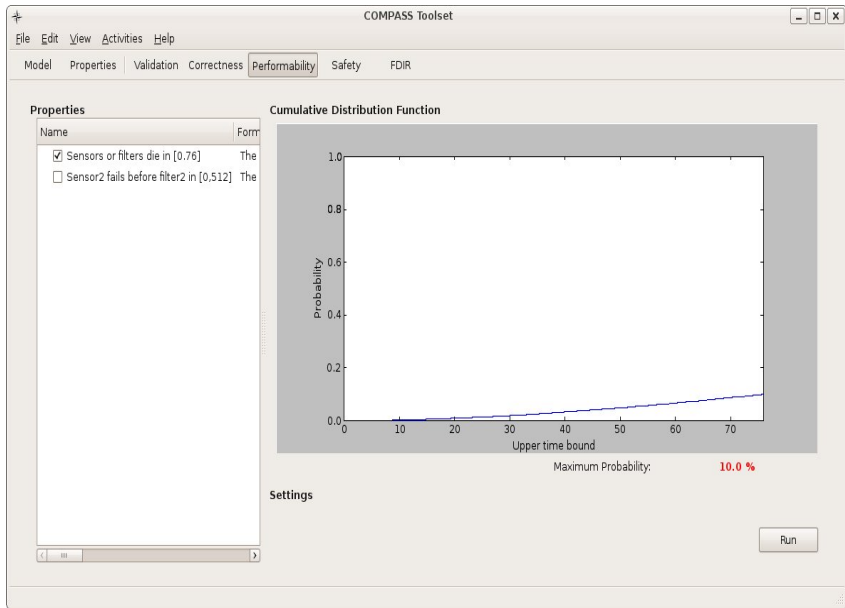
Some properties of interest

- 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

Example: Sensor filter example

Recapitulate the sensor filter example with error model.





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](#))