

SPA for quantitative analysis: Lecture 3 — Case Studies and Tools

Jane Hillston

LFCS, School of Informatics
The University of Edinburgh
Scotland

5th March 2013

Outline

- 1 Recap
- 2 Case Studies
- 3 Roland the Gunslinger
- 4 Tools
- 5 Web Service Composition
- 6 Querying models

Outline

- 1 Recap
- 2 Case Studies
- 3 Roland the Gunslinger
- 4 Tools
- 5 Web Service Composition
- 6 Querying models

Dynamic behaviour

- The behaviour of a model is dictated by the **semantic rules** governing the combinators of the language.
- The possible evolutions of a model are captured by applying these rules exhaustively, generating a **labelled transition system**.
- This can be viewed as a graph in which each node is a state of the model (comprised of the local states of each of the components) and the arcs represent the actions which can cause the move from one state to another.
- The language is also equipped with **observational equivalence** which can be used to compare models.

Dynamic behaviour

- The behaviour of a model is dictated by the **semantic rules** governing the combinators of the language.
- The possible evolutions of a model are captured by applying these rules exhaustively, generating a **labelled transition system**.
- This can be viewed as a graph in which each node is a state of the model (comprised of the local states of each of the components) and the arcs represent the actions which can cause the move from one state to another.
- The language is also equipped with **observational equivalence** which can be used to compare models.

Dynamic behaviour

- The behaviour of a model is dictated by the **semantic rules** governing the combinators of the language.
- The possible evolutions of a model are captured by applying these rules exhaustively, generating a **labelled transition system**.
- This can be viewed as a graph in which each node is a state of the model (comprised of the local states of each of the components) and the arcs represent the actions which can cause the move from one state to another.
- The language is also equipped with **observational equivalence** which can be used to compare models.

Dynamic behaviour

- The behaviour of a model is dictated by the **semantic rules** governing the combinators of the language.
- The possible evolutions of a model are captured by applying these rules exhaustively, generating a **labelled transition system**.
- This can be viewed as a graph in which each node is a state of the model (comprised of the local states of each of the components) and the arcs represent the actions which can cause the move from one state to another.
- The language is also equipped with **observational equivalence** which can be used to compare models.

PEPA Eclipse Plug-In input

$$P_1 \stackrel{def}{=} (start, r_1).P_2 \quad P_2 \stackrel{def}{=} (run, r_2).P_3 \quad P_3 \stackrel{def}{=} (stop, r_3).P_1$$
$$P_1 \parallel P_1$$

PEPA Eclipse Plug-In input

$$P_1 \stackrel{\text{def}}{=} (\text{start}, r_1).P_2 \quad P_2 \stackrel{\text{def}}{=} (\text{run}, r_2).P_3 \quad P_3 \stackrel{\text{def}}{=} (\text{stop}, r_3).P_1$$

$$P_1 \parallel P_1$$

State space

- 1 $P_1 \parallel P_1$
- 2 $P_1 \parallel P_2$
- 3 $P_2 \parallel P_1$
- 4 $P_1 \parallel P_3$
- 5 $P_2 \parallel P_2$
- 6 $P_3 \parallel P_1$
- 7 $P_3 \parallel P_2$
- 8 $P_3 \parallel P_2$
- 9 $P_3 \parallel P_3$

PEPA Eclipse Plug-In input

$$P_1 \stackrel{\text{def}}{=} (\text{start}, r_1).P_2 \quad P_2 \stackrel{\text{def}}{=} (\text{run}, r_2).P_3 \quad P_3 \stackrel{\text{def}}{=} (\text{stop}, r_3).P_1$$

$$P_1 \parallel P_1$$

CTMC representation computed by the plug-in

$$\begin{pmatrix} -2r_1 & r_1 & r_1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -r_1 - r_2 & 0 & r_2 & r_1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -r_1 - r_2 & 0 & r_1 & r_2 & 0 & 0 & 0 \\ r_3 & 0 & 0 & -r_1 - r_3 & 0 & 0 & 0 & r_1 & 0 \\ 0 & 0 & 0 & 0 & -2r_2 & 0 & r_2 & r_2 & 0 \\ r_3 & 0 & 0 & 0 & 0 & -r_1 - r_3 & r_1 & 0 & 0 \\ 0 & r_3 & 0 & 0 & 0 & 0 & -r_2 - r_3 & 0 & r_2 \\ 0 & 0 & r_3 & 0 & 0 & 0 & 0 & -r_2 - r_3 & r_2 \\ 0 & 0 & 0 & r_3 & 0 & r_3 & 0 & 0 & -2r_3 \end{pmatrix}$$

The PEPA Eclipse Plug-in processing the model

The screenshot shows the Eclipse IDE with the PEPA plug-in. The main editor displays the following PEPA model:

```

r1 = 1.0; r2 = 1.0; r3 = 1.0;

P1 = (start, r1).P2;
P2 = (run, r2).P3;
P3 = (stop, r3).P1;

P1 ⇔ P1
  
```

The left sidebar shows a project navigator with various PEPA models, including:

- figure 23.csv
- Figure 7.csv
- Figure 9.csv
- finance.pepa
- finance.tex
- ghost.pepa
- HomeBPEL.pepa
- incomplete.pepa
- kdc.pepa
- localDeadlock.pepa
- mobileagent.pepa.m
- mobileagent.pepa
- mobileagent.pepa.cmd
- PC-LAN4.pepa
- PC-LAN6.pepa
- PQ.pepa
- PQ.pepa.filters
- printer.pepa
- proces.generator
- proces.pepa
- proces.pepa.filters
- proces.statepace
- responsetime.pepa
- tiny.cdf.csv
- tiny.gdf.csv
- tiny.pepa**
- tinyFailures.pepa
- WEB1.pepa
- WEB2.pepa
- WEB4.pepa
- worms.pepa

The right sidebar shows a table with Utilisation, Throughput, and Population metrics:

Utilisation	Throughput	Population
Action	Throughput	
run	0.6666666666666667	
start	0.6666666666666667	
stop	0.6666666666666667	

The bottom console shows the state space view with 9 states:

```

9 states
1 P1 P1 0.111111111111111109
2 P2 P1 0.111111111111111111
3 P1 P2 0.111111111111111111
4 P3 P1 0.111111111111111105
5 P2 P2 0.111111111111111111
6 P1 P3 0.111111111111111113
7 P3 P2 0.111111111111111111
8 P2 P3 0.111111111111111112
9 P3 P3 0.111111111111111116
  
```

Outline

- 1 Recap
- 2 Case Studies
- 3 Roland the Gunslinger
- 4 Tools
- 5 Web Service Composition
- 6 Querying models

PEPA Case Studies (1)

- Multiprocessor access-contention protocols ([Gilmore, Hillston and Ribaudo](#), [Edinburgh and Turin](#))
- Protocols for fault-tolerant systems ([Clark, Gilmore, Hillston and Ribaudo](#), [Edinburgh and Turin](#))
- Multimedia traffic characteristics ([Bowman et al](#), [Kent](#))
- QoS protocols for mobile devices ([Wang et al.](#), [EE department](#), [Edinburgh](#))
- Software Architectures ([Pooley, Bradley and Thomas](#), [Heriot-Watt and Durham](#))
- Switch behaviour in active networks ([Hillston, Kloul and Mokhtari](#), [Edinburgh and Versailles](#))

PEPA Case Studies (1)

- Multiprocessor access-contention protocols ([Gilmore, Hillston and Ribaudo](#), [Edinburgh and Turin](#))
- Protocols for fault-tolerant systems ([Clark, Gilmore, Hillston and Ribaudo](#), [Edinburgh and Turin](#))
- Multimedia traffic characteristics ([Bowman et al](#), [Kent](#))
- QoS protocols for mobile devices ([Wang et al.](#), [EE department](#), [Edinburgh](#))
- Software Architectures ([Pooley, Bradley and Thomas](#), [Heriot-Watt and Durham](#))
- Switch behaviour in active networks ([Hillston, Kloul and Mokhtari](#), [Edinburgh and Versailles](#))

PEPA Case Studies (1)

- Multiprocessor access-contention protocols ([Gilmore, Hillston and Ribaudo](#), [Edinburgh and Turin](#))
- Protocols for fault-tolerant systems ([Clark, Gilmore, Hillston and Ribaudo](#), [Edinburgh and Turin](#))
- Multimedia traffic characteristics ([Bowman et al](#), [Kent](#))
- QoS protocols for mobile devices ([Wang et al.](#), [EE department](#), [Edinburgh](#))
- Software Architectures ([Pooley, Bradley and Thomas](#), [Heriot-Watt and Durham](#))
- Switch behaviour in active networks ([Hillston, Kloul and Mokhtari](#), [Edinburgh and Versailles](#))

PEPA Case Studies (1)

- Multiprocessor access-contention protocols ([Gilmore, Hillston and Ribaldo, Edinburgh and Turin](#))
- Protocols for fault-tolerant systems ([Clark, Gilmore, Hillston and Ribaldo, Edinburgh and Turin](#))
- Multimedia traffic characteristics ([Bowman et al, Kent](#))
- QoS protocols for mobile devices ([Wang et al., EE department, Edinburgh](#))
- Software Architectures ([Pooley, Bradley and Thomas, Heriot-Watt and Durham](#))
- Switch behaviour in active networks ([Hillston, Kloul and Mokhtari, Edinburgh and Versailles](#))

PEPA Case Studies (1)

- Multiprocessor access-contention protocols ([Gilmore, Hillston and Ribaudo, Edinburgh and Turin](#))
- Protocols for fault-tolerant systems ([Clark, Gilmore, Hillston and Ribaudo, Edinburgh and Turin](#))
- Multimedia traffic characteristics ([Bowman et al, Kent](#))
- QoS protocols for mobile devices ([Wang et al., EE department, Edinburgh](#))
- Software Architectures ([Pooley, Bradley and Thomas, Heriot-Watt and Durham](#))
- Switch behaviour in active networks ([Hillston, Kloul and Mokhtari, Edinburgh and Versailles](#))

PEPA Case Studies (1)

- Multiprocessor access-contention protocols ([Gilmore, Hillston and Ribaudo, Edinburgh and Turin](#))
- Protocols for fault-tolerant systems ([Clark, Gilmore, Hillston and Ribaudo, Edinburgh and Turin](#))
- Multimedia traffic characteristics ([Bowman et al, Kent](#))
- QoS protocols for mobile devices ([Wang et al., EE department, Edinburgh](#))
- Software Architectures ([Pooley, Bradley and Thomas, Heriot-Watt and Durham](#))
- Switch behaviour in active networks ([Hillston, Kloul and Mokhtari, Edinburgh and Versailles](#))

PEPA Case Studies (2)

- Task scheduling in a Grid-based processing system, ([Benoit *et al.*, Edinburgh](#))
- Probability of airbag deployment ([Clark, Gilmore and Tribastone, Edinburgh](#))
- Crowd Interactions in Smart Environments ([Harrison, Massink and Latella, Newcastle and ISTI](#))
- Spread of a computer virus via a computer network ([Bradley, Gilmore and Hillston, Imperial and Edinburgh](#))
- Disease spread within animal populations ([Benkrine, McCaig, Norman and Shankland, Stirling](#))

PEPA Case Studies (2)

- Task scheduling in a Grid-based processing system, ([Benoit *et al.*, Edinburgh](#))
- Probability of airbag deployment ([Clark, Gilmore and Tribastone, Edinburgh](#))
- Crowd Interactions in Smart Environments ([Harrison, Massink and Latella, Newcastle and ISTI](#))
- Spread of a computer virus via a computer network ([Bradley, Gilmore and Hillston, Imperial and Edinburgh](#))
- Disease spread within animal populations ([Benkrine, McCaig, Norman and Shankland, Stirling](#))

PEPA Case Studies (2)

- Task scheduling in a Grid-based processing system, ([Benoit et al., Edinburgh](#))
- Probability of airbag deployment ([Clark, Gilmore and Tribastone, Edinburgh](#))
- Crowd Interactions in Smart Environments ([Harrison, Massink and Latella, Newcastle and ISTI](#))
- Spread of a computer virus via a computer network ([Bradley, Gilmore and Hillston, Imperial and Edinburgh](#))
- Disease spread within animal populations ([Benkrine, McCaig, Norman and Shankland, Stirling](#))

PEPA Case Studies (2)

- Task scheduling in a Grid-based processing system, ([Benoit *et al.*, Edinburgh](#))
- Probability of airbag deployment ([Clark, Gilmore and Tribastone, Edinburgh](#))
- Crowd Interactions in Smart Environments ([Harrison, Massink and Latella, Newcastle and ISTI](#))
- Spread of a computer virus via a computer network ([Bradley, Gilmore and Hillston, Imperial and Edinburgh](#))
- Disease spread within animal populations ([Benkrine, McCaig, Norman and Shankland, Stirling](#))

PEPA Case Studies (2)

- Task scheduling in a Grid-based processing system, ([Benoit *et al.*, Edinburgh](#))
- Probability of airbag deployment ([Clark, Gilmore and Tribastone, Edinburgh](#))
- Crowd Interactions in Smart Environments ([Harrison, Massink and Latella, Newcastle and ISTI](#))
- Spread of a computer virus via a computer network ([Bradley, Gilmore and Hillston, Imperial and Edinburgh](#))
- Disease spread within animal populations ([Benkrine, McCaig, Norman and Shankland, Stirling](#))

Outline

- 1 Recap
- 2 Case Studies
- 3 Roland the Gunslinger**
- 4 Tools
- 5 Web Service Composition
- 6 Querying models

Roland the Gunslinger

- This sequence of small examples are based around a character called Roland Deschain.
- Roland is a gunslinger and his life consists of wandering around firing his gun.
- We will consider Roland in a number of different scenarios.
- These are not intended to be serious but they serve to
 - illustrate the main features of the language,
 - give you some experience of how models are constructed, and
 - demonstrate a variety of solution techniques.

Roland the Gunslinger

- This sequence of small examples are based around a character called Roland Deschain.
- Roland is a gunslinger and his life consists of wandering around firing his gun.
- We will consider Roland in a number of different scenarios.
- These are not intended to be serious but they serve to
 - illustrate the main features of the language,
 - give you some experience of how models are constructed, and
 - demonstrate a variety of solution techniques.

Roland the Gunslinger

- This sequence of small examples are based around a character called Roland Deschain.
- Roland is a gunslinger and his life consists of wandering around firing his gun.
- We will consider Roland in a number of different scenarios.
- These are not intended to be serious but they serve to
 - illustrate the main features of the language,
 - give you some experience of how models are constructed, and
 - demonstrate a variety of solution techniques.

Roland the Gunslinger

- This sequence of small examples are based around a character called Roland Deschain.
- Roland is a gunslinger and his life consists of wandering around firing his gun.
- We will consider Roland in a number of different scenarios.
- These are not intended to be serious but they serve to
 - illustrate the main features of the language,
 - give you some experience of how models are constructed, and
 - demonstrate a variety of solution techniques.

Roland alone

In the first scenario we consider Roland alone, with the single activity of firing his gun which is a six-shooter. When his gun is empty Roland will reload the gun and then continue shooting.

$$Roland_6 \stackrel{def}{=} (fire, r_{fire}).Roland_5$$

$$Roland_5 \stackrel{def}{=} (fire, r_{fire}).Roland_4$$

$$Roland_4 \stackrel{def}{=} (fire, r_{fire}).Roland_3$$

$$Roland_3 \stackrel{def}{=} (fire, r_{fire}).Roland_2$$

$$Roland_2 \stackrel{def}{=} (fire, r_{fire}).Roland_1$$

$$Roland_1 \stackrel{def}{=} (fire, r_{fire}).Roland_{empty}$$

$$Roland_{empty} \stackrel{def}{=} (reload, r_{reload}).Roland_6$$

Roland with two guns

All self-respecting gun-slingers have one gun in each hand. If we suppose that Roland has two guns then he should be allowed to fire either gun independently. A simplistic model of this has two instances of *Roland* in parallel:

$$Roland_6 \parallel Roland_6$$

But this does not capture the fact that Roland needs both hands in order to reload either gun. The simplest solution is to assume that Roland only reloads both guns when both are empty.

$$Roland_6 \begin{array}{c} \boxtimes \\ \{reload\} \end{array} Roland_6$$

From now on we restrict Roland to his shotgun, which has two shots and requires both hands for firing.

Roland with two guns

All self-respecting gun-slingers have one gun in each hand. If we suppose that Roland has two guns then he should be allowed to fire either gun independently. A simplistic model of this has two instances of *Roland* in parallel:

$$Roland_6 \parallel Roland_6$$

But this does not capture the fact that Roland needs both hands in order to reload either gun. The simplest solution is to assume that Roland only reloads both guns when both are empty.

$$Roland_6 \begin{array}{c} \boxtimes \\ \{reload\} \end{array} Roland_6$$

From now on we restrict Roland to his shotgun, which has two shots and requires both hands for firing.

Roland with two guns

All self-respecting gun-slingers have one gun in each hand. If we suppose that Roland has two guns then he should be allowed to fire either gun independently. A simplistic model of this has two instances of *Roland* in parallel:

$$Roland_6 \parallel Roland_6$$

But this does not capture the fact that Roland needs both hands in order to reload either gun. The simplest solution is to assume that Roland only reloads both guns when both are empty.

$$Roland_6 \begin{array}{c} \boxtimes \\ \{reload\} \end{array} Roland_6$$

From now on we restrict Roland to his shotgun, which has two shots and requires both hands for firing.

Roland meets an Enemy

- Upon his travels Roland encounters some enemies and when he does so he must fight them.
- Roland is the wildest gunslinger in the west so we assume that no enemy has the skill to seriously harm Roland.
- Each time Roland fires he might miss or hit his target.
- But with nothing to stop him he will keep firing until he successfully hits (and kills) the enemy.
- We assume that some sense of cowboy honour prevents any enemy attacking Roland if he is already involved in a gun fight.

The model

$$Roland_{idle} \stackrel{def}{=} (attack, r_{attack}).Roland_2$$
$$Roland_2 \stackrel{def}{=} (hit, r_{hit}).(reload, r_{reload}).Roland_{idle} \\ + (miss, r_{miss}).Roland_1$$
$$Roland_1 \stackrel{def}{=} (hit, r_{hit}).(reload, r_{reload}).Roland_{idle} \\ + (miss, r_{miss}).Roland_{empty}$$
$$Roland_{empty} \stackrel{def}{=} (reload, r_{reload}).Roland_2$$

Parameter settings for the *Roland₂* model

<i>parameter</i>	<i>value</i>	<i>explanation</i>
r_{fire}	1.0	Roland can fire the gun once per-second
$p_{hit-success}$	0.8	Roland has an 80% success rate
r_{hit}	0.8	$r_{fire} \times p_{hit-success}$
r_{miss}	0.2	$r_{fire} \times (1 - p_{hit-success})$
r_{reload}	0.3	It takes Roland about 3 seconds to reload
r_{attack}	0.01	Roland is attacked once every 100 seconds

Steady state analysis

We can calculate the probability that at arbitrary time Roland is involved in a battle.

This can be based on the steady state probability that *Roland* is in any of the states in which a battle is on-going, i.e. *Roland*₂, *Roland*₁ and *Roland*_{empty}.

```
State Measure 'roland peaceful'  
mean          9.5490716180e-01  
State Measure 'roland in battle'  
mean          0.0450928382e-01
```

> 95% chance that Roland is not currently involved in a gun battle.

Steady state analysis

We can calculate the probability that at arbitrary time Roland is involved in a battle.

This can be based on the steady state probability that *Roland* is in any of the states in which a battle is on-going, i.e. *Roland*₂, *Roland*₁ and *Roland*_{empty}.

```
State Measure 'roland peaceful'  
mean          9.5490716180e-01  
State Measure 'roland in battle'  
mean          0.0450928382e-01
```

> 95% chance that Roland is not currently involved in a gun battle.

Steady state analysis

We can calculate the probability that at arbitrary time Roland is involved in a battle.

This can be based on the steady state probability that *Roland* is in any of the states in which a battle is on-going, i.e. *Roland*₂, *Roland*₁ and *Roland*_{empty}.

```
State Measure 'roland peaceful'  
mean          9.5490716180e-01  
State Measure 'roland in battle'  
mean          0.0450928382e-01
```

> 95% chance that Roland is not currently involved in a gun battle.

Steady state analysis

We can calculate the probability that at arbitrary time Roland is involved in a battle.

This can be based on the steady state probability that *Roland* is in any of the states in which a battle is on-going, i.e. *Roland*₂, *Roland*₁ and *Roland*_{empty}.

```
State Measure 'roland peaceful'  
mean          9.5490716180e-01  
State Measure 'roland in battle'  
mean          0.0450928382e-01
```

> 95% chance that Roland is not currently involved in a gun battle.

Passage-Time Analysis

Passage-time analysis allows us to calculate measures such as the probability that Roland has killed his enemy at a given time after he is attacked.


This would involve calculating the probability that the model performs a *hit* action within the given time after performing an *attack* action.

Passage-Time Analysis

Passage-time analysis allows us to calculate measures such as the probability that Roland has killed his enemy at a given time after he is attacked.


This would involve calculating the probability that the model performs a *hit* action within the given time after performing an *attack* action.

Passage-Time Analysis results



The probability that Roland will successfully perform a *hit* action a given time after an *attack*. Gun battles typically last about 5 seconds and occurs about every 100 seconds. The probability that Roland has performed a *hit* action five seconds after an *attack* action is $\approx 90\%$

Passage-Time Analysis results



The probability that Roland has performed a *miss* action a given time after an *attack* action. These probabilities are much lower because Roland's probability of success are high and if he successfully kills an enemy we must wait for the next attack in order to have a chance of seeing a *miss*.

Cooperation

- In the previous model Roland's enemies were represented only implicitly.
- We now consider a model in which the enemies appear explicitly and allow them to fight back.
- However for now we still assume that they are rather ineffectual and so they never seriously injure Roland.
- This model can be used to calculate properties such as the likelihood that an enemy will manage to fire one shot before they are killed by Roland.

Cooperation

- In the previous model Roland's enemies were represented only implicitly.
- We now consider a model in which the enemies appear explicitly and allow them to fight back.
- However for now we still assume that they are rather ineffectual and so they never seriously injure Roland.
- This model can be used to calculate properties such as the likelihood that an enemy will manage to fire one shot before they are killed by Roland.

Cooperation

- In the previous model Roland's enemies were represented only implicitly.
- We now consider a model in which the enemies appear explicitly and allow them to fight back.
- However for now we still assume that they are rather ineffectual and so they never seriously injure Roland.
- This model can be used to calculate properties such as the likelihood that an enemy will manage to fire one shot before they are killed by Roland.

Cooperation

- In the previous model Roland's enemies were represented only implicitly.
- We now consider a model in which the enemies appear explicitly and allow them to fight back.
- However for now we still assume that they are rather ineffectual and so they never seriously injure Roland.
- This model can be used to calculate properties such as the likelihood that an enemy will manage to fire one shot before they are killed by Roland.

Revised Model

$$Roland_{idle} \stackrel{def}{=} (attack, \top).Roland_2$$

$$Roland_2 \stackrel{def}{=} (hit, r_{hit}).(reload, r_{reload}).Roland_{idle} \\ + (miss, r_{miss}).Roland_1$$

$$Roland_1 \stackrel{def}{=} (hit, r_{hit}).(reload, r_{reload}).Roland_{idle} \\ + (miss, r_{miss}).Roland_{empty}$$

$$Roland_{empty} \stackrel{def}{=} (reload, r_{reload}).Roland_2$$

$$Enemies_{idle} \stackrel{def}{=} (attack, r_{attack}).Enemies_{attack}$$

$$Enemies_{attack} \stackrel{def}{=} (fire, r_{e-miss}).Enemies_{attack} \\ + (hit, \top).Enemies_{idle}$$

$$Roland_{idle} \quad \boxtimes_{\{hit, attack\}} \quad Enemies_{idle}$$

Additional parameters

<i>parameter</i>	<i>value</i>	<i>explanation</i>
r_{attack}	0.01	Roland is attacked once every 100 seconds
r_{e-miss}	0.3	Enemies can fire only once every 3 seconds

Levels of abstraction

- Notice that in this model the behaviour of the enemy has been simplified.
- There is no running out of bullets or reloading.
- This model can be thought of as an approximation to a more complicated component similar to the one which models Roland.
- Here the rate at which the enemy fires encompasses all of the actions, including the reloading of an empty gun.
- We may choose to model a component in such an abstract way when the focus of our modelling is really elsewhere in the model.

Levels of abstraction

- Notice that in this model the behaviour of the enemy has been simplified.
- There is no running out of bullets or reloading.
- This model can be thought of as an approximation to a more complicated component similar to the one which models Roland.
- Here the rate at which the enemy fires encompasses all of the actions, including the reloading of an empty gun.
- We may choose to model a component in such an abstract way when the focus of our modelling is really elsewhere in the model.

Levels of abstraction

- Notice that in this model the behaviour of the enemy has been simplified.
- There is no running out of bullets or reloading.
- This model can be thought of as an approximation to a more complicated component similar to the one which models Roland.
- Here the rate at which the enemy fires encompasses all of the actions, including the reloading of an empty gun.
- We may choose to model a component in such an abstract way when the focus of our modelling is really elsewhere in the model.

Levels of abstraction

- Notice that in this model the behaviour of the enemy has been simplified.
- There is no running out of bullets or reloading.
- This model can be thought of as an approximation to a more complicated component similar to the one which models Roland.
- Here the rate at which the enemy fires encompasses all of the actions, including the reloading of an empty gun.
- We may choose to model a component in such an abstract way when the focus of our modelling is really elsewhere in the model.

Levels of abstraction

- Notice that in this model the behaviour of the enemy has been simplified.
- There is no running out of bullets or reloading.
- This model can be thought of as an approximation to a more complicated component similar to the one which models Roland.
- Here the rate at which the enemy fires encompasses all of the actions, including the reloading of an empty gun.
- We may choose to model a component in such an abstract way when the focus of our modelling is really elsewhere in the model.

Model Validation

It is also sometimes useful to carry out a validation of the model by calculating a metric which we believe we already know the value of.

For example in this model we could make such a sanity check by calculating the probability that the model is in a state in which Roland is idle but the enemies are not, or vice versa.

This should never occur and hence the probability should be zero.

Model Validation

It is also sometimes useful to carry out a validation of the model by calculating a metric which we believe we already know the value of.

For example in this model we could make such a sanity check by calculating the probability that the model is in a state in which Roland is idle but the enemies are not, or vice versa.

This should never occur and hence the probability should be zero.

Model Validation

It is also sometimes useful to carry out a validation of the model by calculating a metric which we believe we already know the value of.

For example in this model we could make such a sanity check by calculating the probability that the model is in a state in which Roland is idle but the enemies are not, or vice versa.

This should never occur and hence the probability should be zero.

Sensitivity Analysis

- **Sensitivity analysis** studies how much influence particular parameter values, such as activity rates, have on performance metrics calculated for the system as a whole.
- A single activity in a PEPA model may have a significant impact on the dynamics of the model, or, conversely, may exert very little influence.
- Sensitivity analysis is performed by solving the model many times while varying the rates slightly.
- For this model we chose to vary three of the rates involved and measured the passage time between an *attack* and a *hit* activity, for each combination of rates.

Sensitivity Analysis

- **Sensitivity analysis** studies how much influence particular parameter values, such as activity rates, have on performance metrics calculated for the system as a whole.
- A single activity in a PEPA model may have a significant impact on the dynamics of the model, or, conversely, may exert very little influence.
- Sensitivity analysis is performed by solving the model many times while varying the rates slightly.
- For this model we chose to vary three of the rates involved and measured the passage time between an *attack* and a *hit* activity, for each combination of rates.

Sensitivity Analysis


- **Sensitivity analysis** studies how much influence particular parameter values, such as activity rates, have on performance metrics calculated for the system as a whole.
- A single activity in a PEPA model may have a significant impact on the dynamics of the model, or, conversely, may exert very little influence.
- Sensitivity analysis is performed by solving the model many times while varying the rates slightly.
- For this model we chose to vary three of the rates involved and measured the passage time between an *attack* and a *hit* activity, for each combination of rates.

Sensitivity Analysis

- **Sensitivity analysis** studies how much influence particular parameter values, such as activity rates, have on performance metrics calculated for the system as a whole.
- A single activity in a PEPA model may have a significant impact on the dynamics of the model, or, conversely, may exert very little influence.
- Sensitivity analysis is performed by solving the model many times while varying the rates slightly.
- For this model we chose to vary three of the rates involved and measured the passage time between an *attack* and a *hit* activity, for each combination of rates.

Sensitivity Analysis: Results


Sensitivity of cumulative distribution function to hitSuccess



The effect of varying the $p_{hit-success}$ parameter.

Sensitivity Analysis: Results


Sensitivity of cumulative distribution function to fireRate



The effect of varying the r_{fire} parameter.

Sensitivity Analysis: Results

Sensitivity of cumulative distribution function to reloadRate



The effect of varying the r_{reload} parameter.

Accurate Enemies

- We now allow the enemies of Roland to actually hit him. This means that Roland may die. It is important to note that this has the consequence that the model will always deadlock. The underlying Markov process is no longer ergodic.
- We assume that the enemies can only hit Roland once every 50 seconds. This rate approximates the rate of a more detailed model in which we would assign a process to the enemies which is much like that of the process which describes Roland.
- The only new parameter is $r_{e\text{-hit}}$ which is assigned a value 0.02 to reflect this assumption.

Accurate Enemies

- We now allow the enemies of Roland to actually hit him. This means that Roland may die. It is important to note that this has the consequence that the model will always deadlock. The underlying Markov process is no longer ergodic.
- We assume that the enemies can only hit Roland once every 50 seconds. This rate approximates the rate of a more detailed model in which we would assign a process to the enemies which is much like that of the process which describes Roland.
- The only new parameter is $r_{e\text{-hit}}$ which is assigned a value 0.02 to reflect this assumption.

Accurate Enemies

- We now allow the enemies of Roland to actually hit him. This means that Roland may die. It is important to note that this has the consequence that the model will always deadlock. The underlying Markov process is no longer ergodic.
- We assume that the enemies can only hit Roland once every 50 seconds. This rate approximates the rate of a more detailed model in which we would assign a process to the enemies which is much like that of the process which describes Roland.
- The only new parameter is $r_{e\text{-hit}}$ which is assigned a value 0.02 to reflect this assumption.

New Roland

$$Roland_{idle} \stackrel{def}{=} (attack, \top).Roland_2$$

$$Roland_2 \stackrel{def}{=} (hit, r_{hit}).(reload, r_{reload}).Roland_{idle}$$

$$+ (miss, r_{miss}).Roland_1$$

$$+ (e-hit, \top).Roland_{dead}$$

$$Roland_1 \stackrel{def}{=} (hit, r_{hit}).(reload, r_{reload}).Roland_{idle}$$

$$+ (miss, r_{miss}).Roland_{empty}$$

$$+ (e-hit, \top).Roland_{dead}$$

$$Roland_{empty} \stackrel{def}{=} (reload, r_{reload}).(reload, r_{reload}).Roland_2$$

$$+ (e-hit, \top).Roland_{dead}$$

$$Roland_{dead} \stackrel{def}{=} Stop$$

New Enemy

$$\begin{array}{l}
 \textit{Enemies}_{idle} \\
 \textit{Enemies}_{attack} \\
 \\
 \textit{Roland}_{idle}
 \end{array}
 \begin{array}{c}
 \begin{array}{c} \underline{\underline{def}} \\ \underline{\underline{def}} \\ + \end{array} \\
 \\
 \begin{array}{c} \boxtimes \\ \{hit, attack, e-hit\} \end{array}
 \end{array}
 \begin{array}{l}
 (\textit{attack}, r_{\textit{attack}}).\textit{Enemies}_{attack} \\
 (\textit{e-hit}, r_{\textit{e-hit}}).\textit{Enemies}_{idle} \\
 (\textit{hit}, \top).\textit{Enemies}_{idle} \\
 \\
 \textit{Enemies}_{idle}
 \end{array}$$

Model Analysis

Steady-State Analysis Since there is an infinite supply of enemies eventually Roland will always die and the model will deadlock.

Transient Analysis Transient analysis on this model can be used to calculate the probability that Roland is dead after a given amount of time. As time increases this should tend towards probability 1.

Passage-Time Analysis Passage-time analysis could be used to calculate the probability of a given event happening at a given time after another given event, e.g. from an attack on Roland until he dies or wins the gun fight.

Model Analysis

Steady-State Analysis Since there is an infinite supply of enemies eventually Roland will always die and the model will deadlock.

Transient Analysis Transient analysis on this model can be used to calculate the probability that Roland is dead after a given amount of time. As time increases this should tend towards probability 1.

Passage-Time Analysis Passage-time analysis could be used to calculate the probability of a given event happening at a given time after another given event, e.g. from an attack on Roland until he dies or wins the gun fight.

Model Analysis

Steady-State Analysis Since there is an infinite supply of enemies eventually Roland will always die and the model will deadlock.

Transient Analysis Transient analysis on this model can be used to calculate the probability that Roland is dead after a given amount of time. As time increases this should tend towards probability 1.

Passage-Time Analysis Passage-time analysis could be used to calculate the probability of a given event happening at a given time after another given event, e.g. from an attack on Roland until he dies or wins the gun fight.

Roland makes a friend

In the next revision of the model we introduce an accomplice who is befriended by Roland and who, when Roland is attacked, fights alongside him.

In this scenario cooperation is used to synchronise between components of the model such that they observe events which they neither directly cause nor are directly affected by.

Whenever either Roland or the accomplice kills the enemy the other must witness this action, so as to stop firing at a dead opponent (it would be a waste of ammunition!).

Roland makes a friend

In the next revision of the model we introduce an accomplice who is befriended by Roland and who, when Roland is attacked, fights alongside him.

In this scenario cooperation is used to synchronise between components of the model such that they observe events which they neither directly cause nor are directly affected by.

Whenever either Roland or the accomplice kills the enemy the other must witness this action, so as to stop firing at a dead opponent (it would be a waste of ammunition!).

Roland makes a friend

In the next revision of the model we introduce an accomplice who is befriended by Roland and who, when Roland is attacked, fights alongside him.

In this scenario cooperation is used to synchronise between components of the model such that they observe events which they neither directly cause nor are directly affected by.

Whenever either Roland or the accomplice kills the enemy the other must witness this action, so as to stop firing at a dead opponent (it would be a waste of ammunition!).

A new component for Roland

$$\begin{aligned}
 \text{Roland}_{idle} &\stackrel{\text{def}}{=} (\text{attack}, \top). \text{Roland}_2 \\
 &+ (\text{befriend}, r_{\text{befriend}}). \text{Roland}_{idle}
 \end{aligned}$$

$$\begin{aligned}
 \text{Roland}_2 &\stackrel{\text{def}}{=} (\text{hit}, r_{\text{hit}}). \text{Roland}_{hit} + (\text{miss}, r_{\text{miss}}). \text{Roland}_1 \\
 &+ (\text{a-hit}, \top). \text{Roland}_{idle}
 \end{aligned}$$

$$\begin{aligned}
 \text{Roland}_1 &\stackrel{\text{def}}{=} (\text{hit}, r_{\text{hit}}). \text{Roland}_{hit} \\
 &+ (\text{miss}, r_{\text{miss}}). \text{Roland}_{empty} \\
 &+ (\text{a-hit}, \top). (\text{reload}, r_{\text{reload}}). \text{Roland}_{idle}
 \end{aligned}$$

$$\text{Roland}_{hit} \stackrel{\text{def}}{=} (\text{enemy-die}, \top). (\text{reload}, r_{\text{reload}}). \text{Roland}_{idle}$$

$$\begin{aligned}
 \text{Roland}_{empty} &\stackrel{\text{def}}{=} (\text{reload}, r_{\text{reload}}). \text{Roland}_2 \\
 &+ (\text{a-hit}, \top). (\text{reload}, r_{\text{reload}}). \text{Roland}_{idle}
 \end{aligned}$$

Synchronising Roland and the Accomplice

- When there is an accomplice, he and Roland fight together against the enemy.
- This involves some cooperation between them.
- However we do not want to leave Roland vulnerable when there is no accomplice present because some of his actions become blocked.
- To prevent this we introduce a dummy component representing the **absence** of an accomplice.
- In this state the accomplice component will **passively participate** in any attack which Roland makes.

Synchronising Roland and the Accomplice

- When there is an accomplice, he and Roland fight together against the enemy.
- This involves some cooperation between them.
- However we do not want to leave Roland vulnerable when there is no accomplice present because some of his actions become blocked.
- To prevent this we introduce a dummy component representing the **absence** of an accomplice.
- In this state the accomplice component will **passively participate** in any attack which Roland makes.

Synchronising Roland and the Accomplice

- When there is an accomplice, he and Roland fight together against the enemy.
- This involves some cooperation between them.
- However we do not want to leave Roland vulnerable when there is no accomplice present because some of his actions become blocked.
- To prevent this we introduce a dummy component representing the **absence** of an accomplice.
- In this state the accomplice component will **passively participate** in any attack which Roland makes.

Synchronising Roland and the Accomplice

- When there is an accomplice, he and Roland fight together against the enemy.
- This involves some cooperation between them.
- However we do not want to leave Roland vulnerable when there is no accomplice present because some of his actions become blocked.
- To prevent this we introduce a dummy component representing the **absence** of an accomplice.
- In this state the accomplice component will **passively participate** in any attack which Roland makes.

Synchronising Roland and the Accomplice

- When there is an accomplice, he and Roland fight together against the enemy.
- This involves some cooperation between them.
- However we do not want to leave Roland vulnerable when there is no accomplice present because some of his actions become blocked.
- To prevent this we introduce a dummy component representing the **absence** of an accomplice.
- In this state the accomplice component will **passively participate** in any attack which Roland makes.

Synchronising Roland and the Accomplice

- When there is an accomplice, he and Roland fight together against the enemy.
- This involves some cooperation between them.
- However we do not want to leave Roland vulnerable when there is no accomplice present because some of his actions become blocked.
- To prevent this we introduce a dummy component representing the **absence** of an accomplice.
- In this state the accomplice component will **passively participate** in any attack which Roland makes.

$$\begin{aligned}
 Acmpl_{abs} &\stackrel{def}{=} (befriend, r_{befriend}).Acmpl_{idle} \\
 &\quad + (hit, \top).Acmpl_{abs} \\
 &\quad + (attack, \top).Acmpl_{abs}
 \end{aligned}$$

Component for the Accomplice

$$\begin{aligned}
 \mathit{Acmpl}_{idle} &\stackrel{\text{def}}{=} (\mathit{attack}, \top). \mathit{Acmpl}_2 \\
 \mathit{Acmpl}_2 &\stackrel{\text{def}}{=} (\mathit{a-hit}, r_{\mathit{a-hit}}). \mathit{Acmpl}_{hit} + (\mathit{hit}, \top). \mathit{Acmpl}_{idle} \\
 &\quad + (\mathit{miss}, r_{\mathit{miss}}). \mathit{Acmpl}_1 \\
 &\quad + (\mathit{enemy-hit}, \top). \mathit{Acmpl}_{abs} \\
 \mathit{Acmpl}_1 &\stackrel{\text{def}}{=} (\mathit{a-hit}, r_{\mathit{a-hit}}). \mathit{Acmpl}_{hit} \\
 &\quad + (\mathit{hit}, \top). (\mathit{reload}, r_{\mathit{a-reload}}). \mathit{Acmpl}_{idle} \\
 &\quad + (\mathit{miss}, r_{\mathit{miss}}). \mathit{Acmpl}_{empty} \\
 &\quad + (\mathit{enemy-hit}, \top). \mathit{Acmpl}_{abs} \\
 \mathit{Acmpl}_{hit} &\stackrel{\text{def}}{=} (\mathit{enemy-die}, \top). (\mathit{reload}, r_{\mathit{a-reload}}). \mathit{Acmpl}_{idle} \\
 \mathit{Acmpl}_{empty} &\stackrel{\text{def}}{=} (\mathit{reload}, r_{\mathit{a-reload}}). \mathit{Acmpl}_2 \\
 &\quad + (\mathit{enemy-hit}, \top). \mathit{Acmpl}_{abs} \\
 &\quad + (\mathit{hit}, \top). (\mathit{reload}, r_{\mathit{a-reload}}). \mathit{Acmpl}_{idle}
 \end{aligned}$$

Parameter Settings for the Accomplice

<i>parameter</i>	<i>value</i>	<i>explanation</i>
$r_{befriend}$	0.001	Roland befriends a stranger once every 1000 seconds
r_{a-fire}	1.0	the accomplice can also fire once per second
$p_{a-hit-success}$	0.6	the accomplice has a 60 percent accuracy
r_{a-hit}	0.6	$r_{fire} \times p_{hit-success}$
r_{a-miss}	0.4	$r_{fire} \times (1.0 - p_{hit-success})$
$r_{a-reload}$	0.25	it takes the accomplice 4 seconds to reload

Component for the Enemy

The component representing the enemy is similar to before.

$$\begin{aligned}
 Enemies_{idle} &\stackrel{def}{=} (attack, r_{attack}).Enemies_{attack} \\
 Enemies_{attack} &\stackrel{def}{=} (enemy-hit, r_{e-hit}).Enemies_{attack} \\
 &\quad + (enemy-die, r_{die}).Enemies_{idle}
 \end{aligned}$$

The system equation is as follows:

$$(Roland_2 \boxtimes_{\{attack, hit, a-hit, befriend\}} Acmpl_{abs}) \boxtimes_{\{attack, enemy-die, enemy-hit\}} Enemies_{idle}$$

Model Analysis

Steady-State Analysis

- As before we can determine the probability that Roland is involved in a gun battle at an arbitrary time.
- We could also determine the likelihood that Roland has an accomplice at an arbitrary time.
- Since Roland cannot perform a befriending action while currently involved in a battle, the probability that Roland is in such a battle clearly affects the probability that he is alone in his quest.
- So, for example, if Roland's success rate is reduced then gun battles will take longer to resolve, hence Roland will be involved in a gun battle more often, and therefore he will befriend fewer accomplices.

Model Analysis

Steady-State Analysis

- As before we can determine the probability that Roland is involved in a gun battle at an arbitrary time.
- We could also determine the likelihood that Roland has an accomplice at an arbitrary time.
- Since Roland cannot perform a befriending action while currently involved in a battle, the probability that Roland is in such a battle clearly affects the probability that he is alone in his quest.
- So, for example, if Roland's success rate is reduced then gun battles will take longer to resolve, hence Roland will be involved in a gun battle more often, and therefore he will befriend fewer accomplices.

Model Analysis

Steady-State Analysis

- As before we can determine the probability that Roland is involved in a gun battle at an arbitrary time.
- We could also determine the likelihood that Roland has an accomplice at an arbitrary time.
- Since Roland cannot perform a befriending action while currently involved in a battle, the probability that Roland is in such a battle clearly affects the probability that he is alone in his quest.
- So, for example, if Roland's success rate is reduced then gun battles will take longer to resolve, hence Roland will be involved in a gun battle more often, and therefore he will befriend fewer accomplices.

Model Analysis

Steady-State Analysis

- As before we can determine the probability that Roland is involved in a gun battle at an arbitrary time.
- We could also determine the likelihood that Roland has an accomplice at an arbitrary time.
- Since Roland cannot perform a befriending action while currently involved in a battle, the probability that Roland is in such a battle clearly affects the probability that he is alone in his quest.
- So, for example, if Roland's success rate is reduced then gun battles will take longer to resolve, hence Roland will be involved in a gun battle more often, and therefore he will befriend fewer accomplices.

Model Analysis

Transient Analysis

An example transient analysis would be to determine the expected time after Roland has set off before he meets his first accomplice.

Model Analysis

Passage-Time Analysis

- An example analysis would be to calculate the passage-time from an *attack* action until the death of the enemy or of the accomplice.
- Since all gun battles now end in the enemy being killed stopping the analysis there would give us the expected duration of any one gun battle.
- There is also the possibility to start the analysis from the *befriend* action and stop it with the death of the accomplice.

Model Analysis

Passage-Time Analysis

- An example analysis would be to calculate the passage-time from an *attack* action until the death of the enemy or of the accomplice.
- Since all gun battles now end in the enemy being killed stopping the analysis there would give us the expected duration of any one gun battle.
- There is also the possibility to start the analysis from the *befriend* action and stop it with the death of the accomplice.

Model Analysis

Passage-Time Analysis

- An example analysis would be to calculate the passage-time from an *attack* action until the death of the enemy or of the accomplice.
- Since all gun battles now end in the enemy being killed stopping the analysis there would give us the expected duration of any one gun battle.
- There is also the possibility to start the analysis from the *befriend* action and stop it with the death of the accomplice.

Hiding

- Currently there is nothing in the model to stop an enemy from disrupting the interaction between Roland and his accomplice, e.g. by performing a *befriend* action.
- One way to avoid this is to 'hide' those actions only Roland and the accomplice should cooperate on.
- To do this for our model we can simply change the system equation:

$$((Roland_2 \bowtie_{L_1} Acmpl)/L_1) \bowtie_{L_2} Enemies_{idle}$$

where $L_1 = \{hit, a-hit, befriend\}$ and
 $L_2 = \{attack, enemy-die, enemy-hit\}$.

Hiding

- Currently there is nothing in the model to stop an enemy from disrupting the interaction between Roland and his accomplice, e.g. by performing a *befriend* action.
- One way to avoid this is to 'hide' those actions only Roland and the accomplice should cooperate on.
- To do this for our model we can simply change the system equation:

$$((Roland_2 \bowtie_{L_1} Acmpl) / L_1) \bowtie_{L_2} Enemies_{idle}$$

where $L_1 = \{hit, a-hit, befriend\}$ and
 $L_2 = \{attack, enemy-die, enemy-hit\}$.

Hiding

- Currently there is nothing in the model to stop an enemy from disrupting the interaction between Roland and his accomplice, e.g. by performing a *befriend* action.
- One way to avoid this is to 'hide' those actions only Roland and the accomplice should cooperate on.
- To do this for our model we can simply change the system equation:

$$((Roland_2 \bowtie_{L_1} Acmpl)/L_1) \bowtie_{L_2} Enemies_{idle}$$

where $L_1 = \{hit, a-hit, befriend\}$ and
 $L_2 = \{attack, enemy-die, enemy-hit\}$.

Outline

- 1 Recap
- 2 Case Studies
- 3 Roland the Gunslinger
- 4 Tools**
- 5 Web Service Composition
- 6 Querying models

The PEPA Eclipse Plug-in

Calculating by hand the transitions of a PEPA model and subsequently expressing these in a form which was suitable for solution was a tedious task prone to errors. The PEPA Eclipse Plug-in relieves the modeller of this work.

The PEPA Eclipse Plug-in: functionality

The plug-in will report errors in the model function:

- deadlock,
- absorbing states,
- static synchronisation mismatch (cooperations which do not involve active participants).

The plug-in also generates the transition graph of the model, computes the number of states, formulates the Markov process matrix Q and communicates the matrix to a solver.

The plug-in provides a simple pattern language for selecting states from the stationary distribution.

The PEPA Eclipse Plug-in: functionality

The plug-in will report errors in the model function:

- deadlock,
- absorbing states,
- static synchronisation mismatch (cooperations which do not involve active participants).

The plug-in also generates the transition graph of the model, computes the number of states, formulates the Markov process matrix Q and communicates the matrix to a solver.

The plug-in provides a simple pattern language for selecting states from the stationary distribution.

The PEPA Eclipse Plug-in: functionality

The plug-in will report errors in the model function:

- deadlock,
- absorbing states,
- static synchronisation mismatch (cooperations which do not involve active participants).

The plug-in also generates the transition graph of the model, computes the number of states, formulates the Markov process matrix Q and communicates the matrix to a solver.

The plug-in provides a simple pattern language for selecting states from the stationary distribution.

The PEPA Eclipse Plug-in processing the model

The screenshot displays the Eclipse IDE with the PEPA plug-in. The main editor shows the following PEPA model:

```

r1 = 1.0; r2 = 1.0; r3 = 1.0;

P1 = (start, r1).P2;
P2 = (run, r2).P3;
P3 = (stop, r3).P1;

P1 <- P1
  
```

The left sidebar (Navigator) lists several files, including:

- Figure 23.csv
- Figure 7.csv
- Figure 9.csv
- finance.pepa
- finance.tex
- ghost.pepa
- HomeBPEL.pepa
- incomplete.pepa
- kdc.pepa
- localDeadlock.pepa
- mobileagent_pepa.m
- mobileagent.pepa
- mobileagent.pepa.cmd
- PC-LAN4.pepa
- PC-LAN6.pepa
- PQ.pepa
- PQ.pepa.filters
- printer.pepa
- procces.generator
- procces.pepa
- procces.pepa.filters
- procces.stateSpace
- responsetime.pepa
- tiny-cdf.csv
- tiny-pdf.csv
- tiny.pepa**
- tinyFailures.pepa
- WEB1.pepa
- WEB2.pepa
- WEB4.pepa
- worms.pepa

The right sidebar (Outline) shows a table of performance metrics:

Utilisation	Throughput	Population
run	0.6666666666666667	
start	0.6666666666666667	
stop	0.6666666666666667	

The bottom panel (Problems) shows the state space view with 9 states:

State	Process	Label	Value
1	P1	P1	0.111111111111111109
2	P2	P1	0.111111111111111111
3	P1	P2	0.111111111111111111
4	P3	P1	0.111111111111111105
5	P2	P2	0.111111111111111111
6	P1	P3	0.111111111111111113
7	P3	P2	0.111111111111111111
8	P2	P3	0.111111111111111112
9	P3	P3	0.111111111111111116

The PEPA website

`http://www.dcs.ed.ac.uk/pepa`

From the website the PEPA Eclipse Plug-in and some other tools are available for download.

There is also information about people involved in the PEPA project, projects undertaken and a collection of published papers.

Outline

- 1 Recap
- 2 Case Studies
- 3 Roland the Gunslinger
- 4 Tools
- 5 Web Service Composition**
- 6 Querying models

Web Service Composition: Introduction

We consider an example of a business application which is composed from a number of offered web services.

A user accesses the application via an SMS message requesting directions to the nearest facility (post-office, restaurant, bank etc.) and receives a response as an MMS message containing a map.

Since the application involves a users' current location there is an access control issue since it must be ensured that the web service consumer has the requisite authority to execute the web service it requests.

Moreover the service provider imposes a restriction that only one request may be handled for each SMS message received.

Web Service Composition: Introduction

We consider an example of a business application which is composed from a number of offered web services.

A user accesses the application via an SMS message requesting directions to the nearest facility (post-office, restaurant, bank etc.) and receives a response as an MMS message containing a map.

Since the application involves a users' current location there is an access control issue since it must be ensured that the web service consumer has the requisite authority to execute the web service it requests.

Moreover the service provider imposes a restriction that only one request may be handled for each SMS message received.

Web Service Composition: Introduction

We consider an example of a business application which is composed from a number of offered web services.

A user accesses the application via an SMS message requesting directions to the nearest facility (post-office, restaurant, bank etc.) and receives a response as an MMS message containing a map.

Since the application involves a users' current location there is an access control issue since it must be ensured that the web service consumer has the requisite authority to execute the web service it requests.

Moreover the service provider imposes a restriction that only one request may be handled for each SMS message received.

Web Service Composition: Introduction

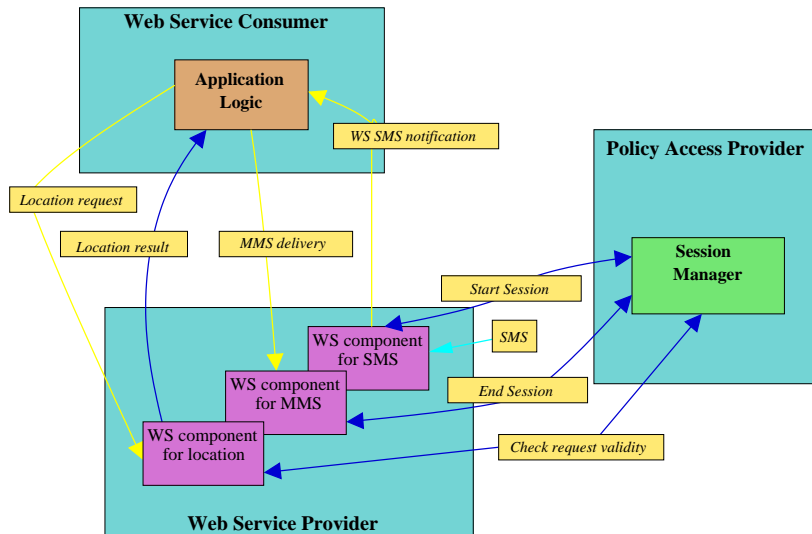
We consider an example of a business application which is composed from a number of offered web services.

A user accesses the application via an SMS message requesting directions to the nearest facility (post-office, restaurant, bank etc.) and receives a response as an MMS message containing a map.

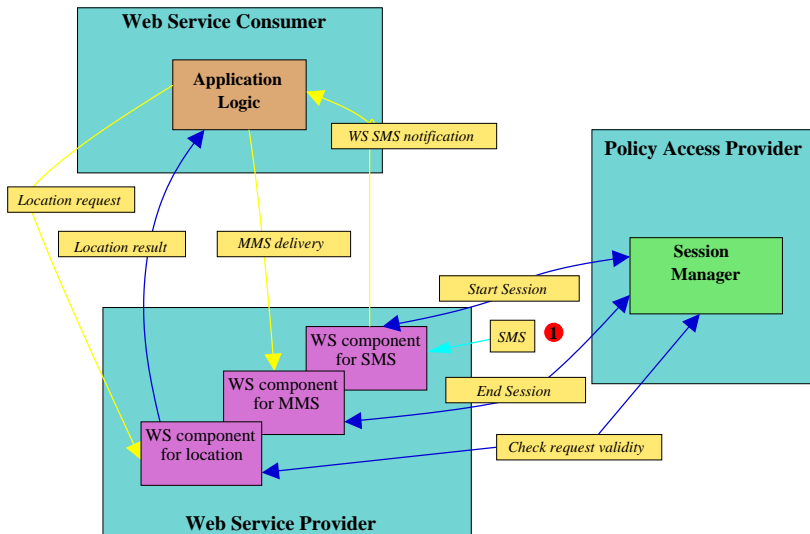
Since the application involves a users' current location there is an access control issue since it must be ensured that the web service consumer has the requisite authority to execute the web service it requests.

Moreover the service provider imposes a restriction that only one request may be handled for each SMS message received.

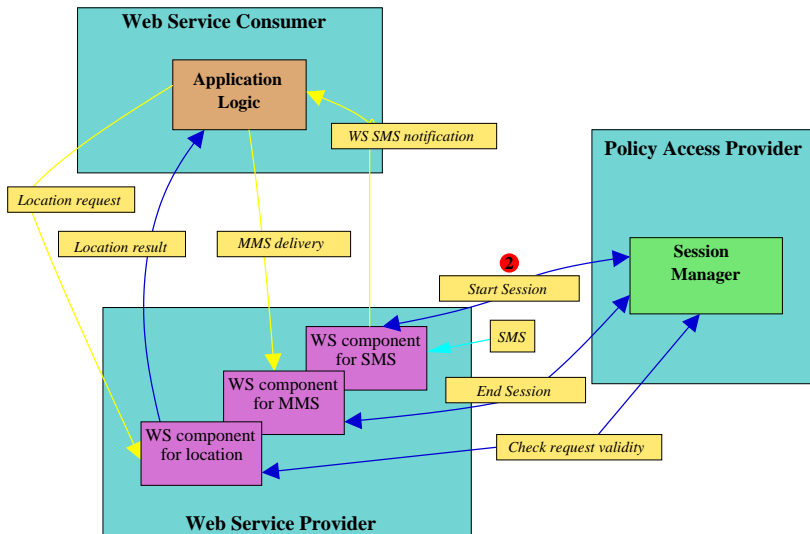
Schematic view



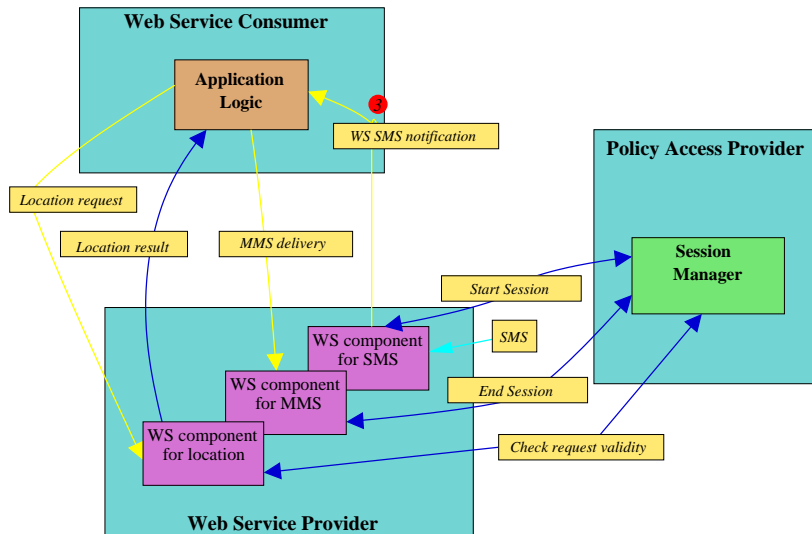
Schematic view



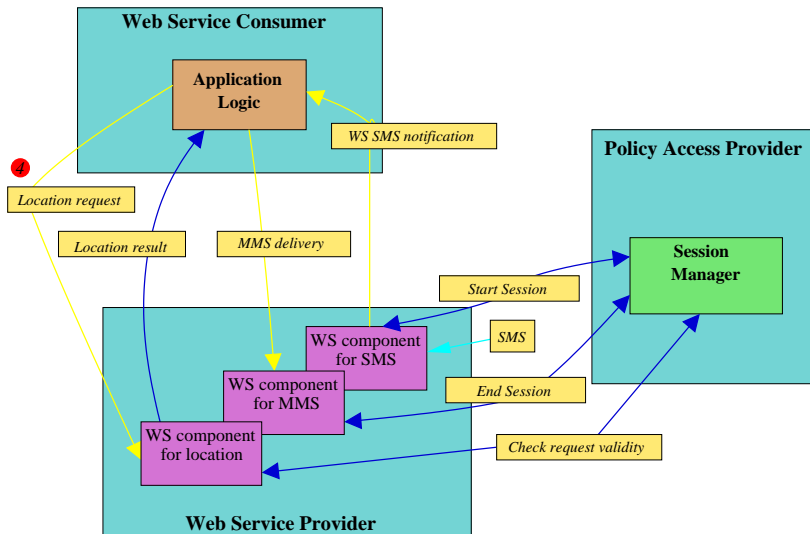
Schematic view



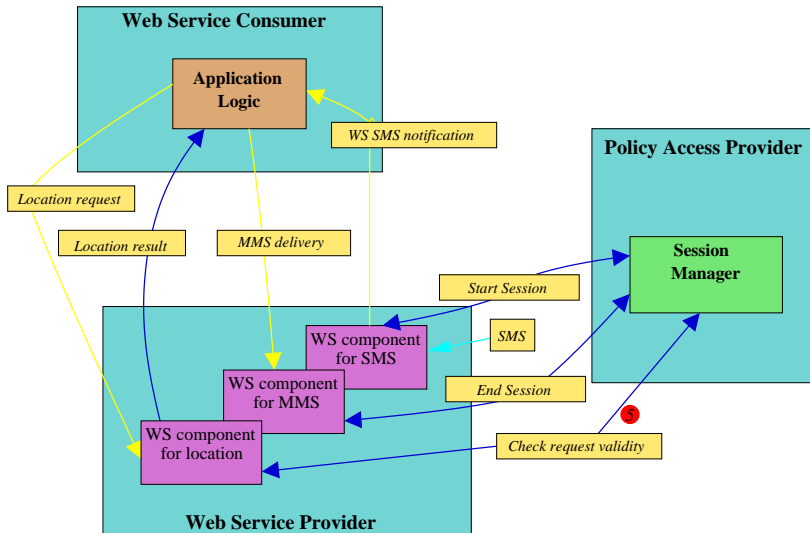
Schematic view



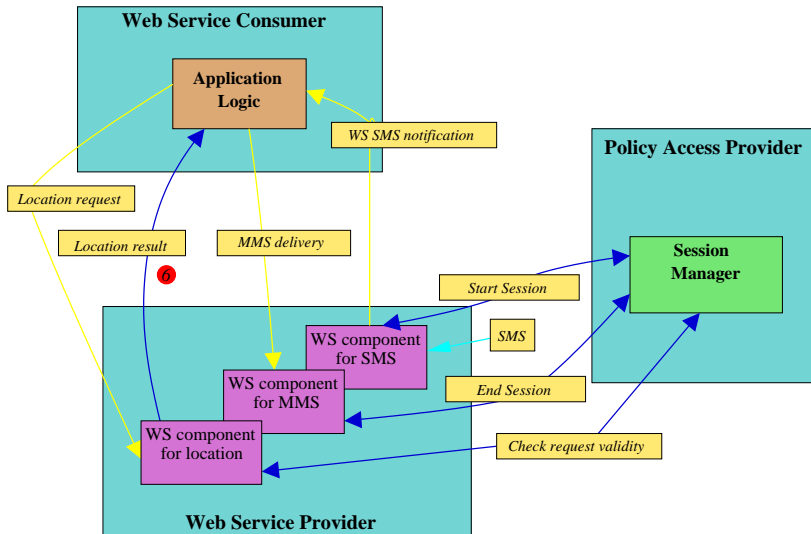
Schematic view



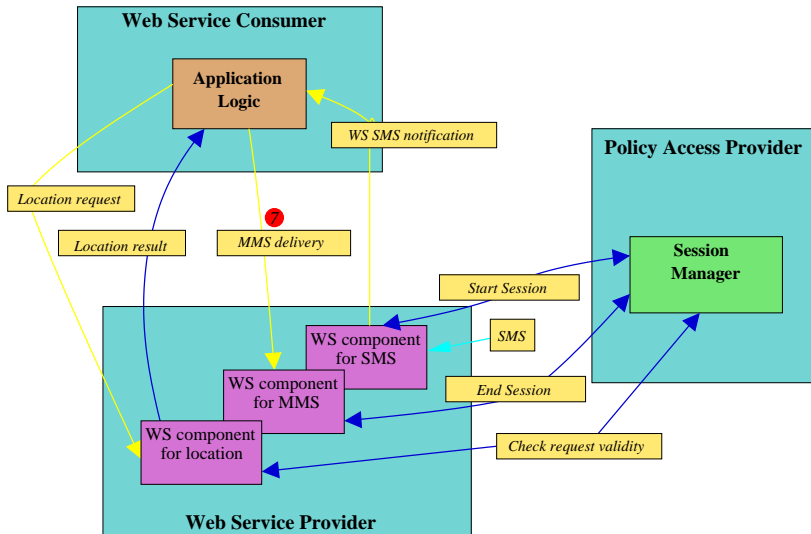
Schematic view



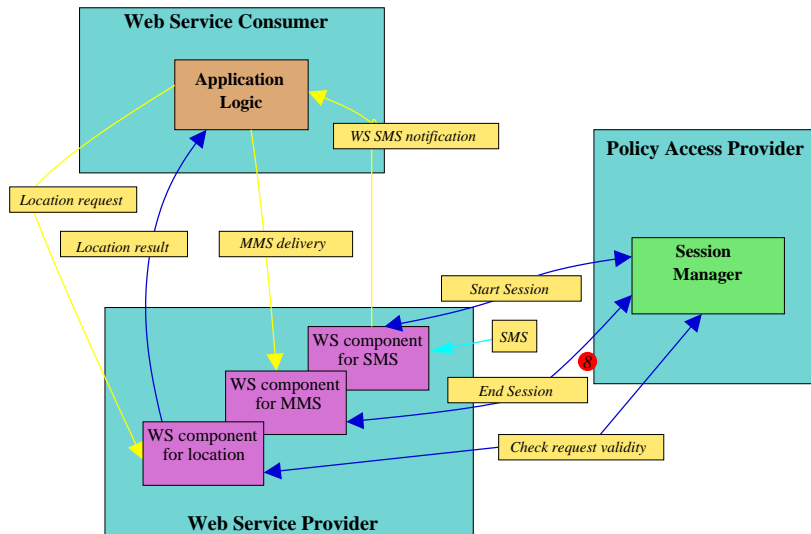
Schematic view



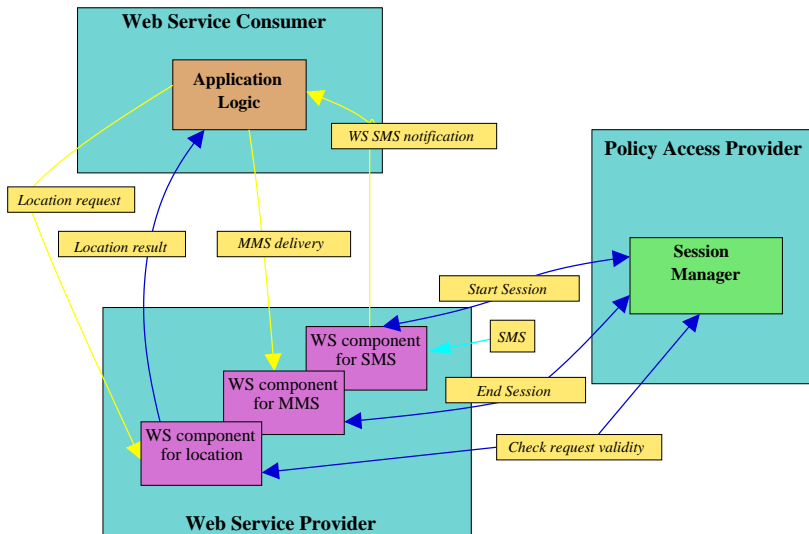
Schematic view



Schematic view



Schematic view



The PEPA model

The PEPA model of the system consists of four components:

- The user;
- The web service provider;
- The web service consumer, and
- The policy access provider.

The Web Service Provider consists of three distinct elements but the web service consumer is associated with a session which accesses each element in sequence.

Concurrency is introduced into the model by allowing multiple sessions rather than by representing the constituent web services separately.

The PEPA model

The PEPA model of the system consists of four components:

- The user;
- The web service provider;
- The web service consumer, and
- The policy access provider.

The Web Service Provider consists of three distinct elements but the web service consumer is associated with a session which accesses each element in sequence.

Concurrency is introduced into the model by allowing multiple sessions rather than by representing the constituent web services separately.

The PEPA model

The PEPA model of the system consists of four components:

- The user;
- The web service provider;
- The web service consumer, and
- The policy access provider.

The Web Service Provider consists of three distinct elements but the web service consumer is associated with a session which accesses each element in sequence.

Concurrency is introduced into the model by allowing multiple sessions rather than by representing the constituent web services separately.

The PEPA model

The PEPA model of the system consists of four components:

- The user;
- The web service provider;
- The web service consumer, and
- The policy access provider.

The Web Service Provider consists of three distinct elements but the web service consumer is associated with a session which accesses each element in sequence.

Concurrency is introduced into the model by allowing multiple sessions rather than by representing the constituent web services separately.

Component *Customer*

The customer's behaviour is simply modelled with two local states.

$$\begin{aligned} \textit{Customer} &\stackrel{\textit{def}}{=} (\textit{getSMS}, r_1).\textit{Customer}_1 \\ \textit{Customer}_1 &\stackrel{\textit{def}}{=} (\textit{getMap}, \top).\textit{Customer} \\ &\quad + (\textit{get404}, \top).\textit{Customer} \end{aligned}$$

We associate the user-perceived system performance with the throughput of the *getMap* action which can be calculated directly from the steady state probability distribution of the underlying Markov chain.

Component *Customer*

The customer's behaviour is simply modelled with two local states.

$$\begin{aligned} \textit{Customer} &\stackrel{\textit{def}}{=} (\textit{getSMS}, r_1).\textit{Customer}_1 \\ \textit{Customer}_1 &\stackrel{\textit{def}}{=} (\textit{getMap}, \top).\textit{Customer} \\ &\quad + (\textit{get404}, \top).\textit{Customer} \end{aligned}$$

We associate the user-perceived system performance with the throughput of the *getMap* action which can be calculated directly from the steady state probability distribution of the underlying Markov chain.

Component *Customer*

The customer's behaviour is simply modelled with two local states.

$$\begin{aligned} Customer &\stackrel{def}{=} (getSMS, r_1).Customer_1 \\ Customer_1 &\stackrel{def}{=} (getMap, \top).Customer \\ &\quad + (get404, \top).Customer \end{aligned}$$

We associate the user-perceived system performance with the throughput of the *getMap* action which can be calculated directly from the steady state probability distribution of the underlying Markov chain.

Component *WSConsumer*

Once a session has been started, it initiates a request for the user's current location and waits for a response.

For valid requests, location is returned and used to compute the appropriate map, which is then sent via an MMS message, using the web service.

$$\begin{aligned} WSConsumer &\stackrel{def}{=} (notify, \top).WSConsumer_2 \\ WSConsumer_2 &\stackrel{def}{=} (locReq, r_4).WSConsumer_3 \\ WSConsumer_3 &\stackrel{def}{=} (locRes, \top).WSConsumer_4 \\ &\quad + (locErr, \top).WSConsumer \\ WSConsumer_4 &\stackrel{def}{=} (compute, r_7).WSConsumer_5 \\ WSConsumer_5 &\stackrel{def}{=} (sendMMS, r_9).WSConsumer \end{aligned}$$

Component *WSConsumer*

Once a session has been started, it initiates a request for the user's current location and waits for a response.

For valid requests, location is returned and used to compute the appropriate map, which is then sent via an MMS message, using the web service.

$$\begin{aligned} WSConsumer &\stackrel{def}{=} (notify, \top).WSConsumer_2 \\ WSConsumer_2 &\stackrel{def}{=} (locReq, r_4).WSConsumer_3 \\ WSConsumer_3 &\stackrel{def}{=} (locRes, \top).WSConsumer_4 \\ &\quad + (locErr, \top).WSConsumer \\ WSConsumer_4 &\stackrel{def}{=} (compute, r_7).WSConsumer_5 \\ WSConsumer_5 &\stackrel{def}{=} (sendMMS, r_9).WSConsumer \end{aligned}$$

Component *WSConsumer*

Once a session has been started, it initiates a request for the user's current location and waits for a response.

For valid requests, location is returned and used to compute the appropriate map, which is then sent via an MMS message, using the web service.

$$\begin{aligned}
 WSConsumer &\stackrel{def}{=} (notify, \top).WSConsumer_2 \\
 WSConsumer_2 &\stackrel{def}{=} (locReq, r_4).WSConsumer_3 \\
 WSConsumer_3 &\stackrel{def}{=} (locRes, \top).WSConsumer_4 \\
 &\quad + (locErr, \top).WSConsumer \\
 WSConsumer_4 &\stackrel{def}{=} (compute, r_7).WSConsumer_5 \\
 WSConsumer_5 &\stackrel{def}{=} (sendMMS, r_9).WSConsumer
 \end{aligned}$$

Component *WSProvider*

The use of sessions restricts a user's access to the services of the Web Service Provider to be sequential.

We assume that there is a distinct instance of the component *WSProvider* for each distinct session.

The *checkValid* action is represented twice, to capture the two possible distinct outcomes of the action.

- If the check is successful the location must be returned to the Web Service Consumer in the form of a map (*getMap*).
- If the check revealed an invalid request (*locErr*) then an error must be returned to the Web Service Consumer (*get404*) and the session terminated (*stopSession*).

Component *WSProvider*

The use of sessions restricts a user's access to the services of the Web Service Provider to be sequential.

We assume that there is a distinct instance of the component *WSProvider* for each distinct session.

The *checkValid* action is represented twice, to capture the two possible distinct outcomes of the action.

- If the check is successful the location must be returned to the Web Service Consumer in the form of a map (*getMap*).
- If the check revealed an invalid request (*locErr*) then an error must be returned to the Web Service Consumer (*get404*) and the session terminated (*stopSession*).

Component *WSProvider*

The use of sessions restricts a user's access to the services of the Web Service Provider to be sequential.

We assume that there is a distinct instance of the component *WSProvider* for each distinct session.

The *checkValid* action is represented twice, to capture the two possible distinct outcomes of the action.

- If the check is successful the location must be returned to the Web Service Consumer in the form of a map (*getMap*).
- If the check revealed an invalid request (*locErr*) then an error must be returned to the Web Service Consumer (*get404*) and the session terminated (*stopSession*).

Component *WSProvider*

The use of sessions restricts a user's access to the services of the Web Service Provider to be sequential.

We assume that there is a distinct instance of the component *WSProvider* for each distinct session.

The *checkValid* action is represented twice, to capture the two possible distinct outcomes of the action.

- If the check is successful the location must be returned to the Web Service Consumer in the form of a map (*getMap*).
- If the check revealed an invalid request (*locErr*) then an error must be returned to the Web Service Consumer (*get404*) and the session terminated (*stopSession*).

Component *WSProvider*

The use of sessions restricts a user's access to the services of the Web Service Provider to be sequential.

We assume that there is a distinct instance of the component *WSProvider* for each distinct session.

The *checkValid* action is represented twice, to capture the two possible distinct outcomes of the action.

- If the check is successful the location must be returned to the Web Service Consumer in the form of a map (*getMap*).
- If the check revealed an invalid request (*locErr*) then an error must be returned to the Web Service Consumer (*get404*) and the session terminated (*stopSession*).

Component *WSProvider*

$$\begin{aligned} \text{WSProvider} &\stackrel{\text{def}}{=} (\text{getSMS}, \top). \text{WSProvider}_2 \\ \text{WSProvider}_2 &\stackrel{\text{def}}{=} (\text{startSession}, r_2). \text{WSProvider}_3 \\ \text{WSProvider}_3 &\stackrel{\text{def}}{=} (\text{notify}, r_3). \text{WSProvider}_4 \\ \text{WSProvider}_4 &\stackrel{\text{def}}{=} (\text{locReq}, \top). \text{WSProvider}_5 \\ \text{WSProvider}_5 &\stackrel{\text{def}}{=} (\text{checkValid}, 99 \cdot \top). \text{WSProvider}_6 \\ &+ (\text{checkValid}, \top). \text{WSProvider}_{10} \end{aligned}$$

Component *WSProvider* cont.

$$WSProvider_6 \stackrel{def}{=} (locRes, r_6).WSProvider_7$$
$$WSProvider_7 \stackrel{def}{=} (sendMMS, \top).WSProvider_8$$
$$WSProvider_8 \stackrel{def}{=} (getMap, r_8).WSProvider_9$$
$$WSProvider_9 \stackrel{def}{=} (stopSession, r_2).WSProvider$$
$$WSProvider_{10} \stackrel{def}{=} (locErr, r_6).WSProvider_{11}$$
$$WSProvider_{11} \stackrel{def}{=} (get404, r_8).WSProvider_9$$

Component *PAP*Provider

We consider a stateless implementation of the policy access provider.

$$\begin{aligned} \mathit{PAP}Provider &\stackrel{\text{def}}{=} (\mathit{startSession}, \top). \mathit{PAP}Provider \\ &+ (\mathit{checkValid}, r_5). \mathit{PAP}Provider \\ &+ (\mathit{stopSession}, \top). \mathit{PAP}Provider \end{aligned}$$

Model Component *WSComp*

The complete system is composed of some number of instances of the components interacting on their shared activities:

$$\begin{aligned}
 WSComp \stackrel{def}{=} & \left((Customer[N_C] \underset{L_1}{\bowtie} WSPProvider[N_{WSP}]) \right. \\
 & \quad \underset{L_2}{\bowtie} WSCConsumer[N_{WSC}] \\
 & \quad \quad \underset{L_3}{\bowtie} PAPProvider[N_{PAP}]
 \end{aligned}$$

where the cooperation sets are

$$L_1 = \{getSMS, getMap, get404\}$$

$$L_2 = \{notify, locReq, locRes, locErr, sendMMS\}$$

$$L_3 = \{startSession, checkValid, stopSession\}$$

Parameter Values

<i>param.</i>	<i>value</i>	<i>explanation</i>
r_1	0.0010	rate customers request maps
r_2	0.5	rate session can be started
r_3	0.1	notification exchange between consumer and provider
r_4	0.1	rate requests for location can be satisfied
r_5	0.05	rate the provider can check the validity of the request
r_6	0.1	rate location information can be returned to consumer
r_7	0.05	rate maps can be generated
r_8	0.02	rate MMS messages can be sent from provider to customer
r_9	$10.0 * r_8$	rate MMS messages can be sent via the Web Service

Steady State Analysis for System Tuning

- Suppose that we want to design the system in such a way that it can handle 30 independent customers.
- Some parameters such as the network delays may be constrained by the available technology.
- However, there are a number of degrees of freedom which let us vary, for example, the number of threads of control of the components of the system.
- The aim of the analysis is to deliver a satisfactory service in a cost-effective way.
- The simplest example of a cost function may be a linearly dependency on the number of copies of a component or the rate at which an activity is performed.

Steady State Analysis for System Tuning

- Suppose that we want to design the system in such a way that it can handle 30 independent customers.
- Some parameters such as the network delays may be constrained by the available technology.
- However, there are a number of degrees of freedom which let us vary, for example, the number of threads of control of the components of the system.
- The aim of the analysis is to deliver a satisfactory service in a cost-effective way.
- The simplest example of a cost function may be a linearly dependency on the number of copies of a component or the rate at which an activity is performed.

Steady State Analysis for System Tuning

- Suppose that we want to design the system in such a way that it can handle 30 independent customers.
- Some parameters such as the network delays may be constrained by the available technology.
- However, there are a number of degrees of freedom which let us vary, for example, the number of threads of control of the components of the system.
- The aim of the analysis is to deliver a satisfactory service in a cost-effective way.
- The simplest example of a cost function may be a linearly dependency on the number of copies of a component or the rate at which an activity is performed.


Steady State Analysis for System Tuning

- Suppose that we want to design the system in such a way that it can handle 30 independent customers.
- Some parameters such as the network delays may be constrained by the available technology.
- However, there are a number of degrees of freedom which let us vary, for example, the number of threads of control of the components of the system.
- The aim of the analysis is to deliver a satisfactory service in a cost-effective way.
- The simplest example of a cost function may be a linearly dependency on the number of copies of a component or the rate at which an activity is performed.

Steady State Analysis for System Tuning

- Suppose that we want to design the system in such a way that it can handle 30 independent customers.
- Some parameters such as the network delays may be constrained by the available technology.
- However, there are a number of degrees of freedom which let us vary, for example, the number of threads of control of the components of the system.
- The aim of the analysis is to deliver a satisfactory service in a cost-effective way.
- The simplest example of a cost function may be a linearly dependency on the number of copies of a component or the rate at which an activity is performed.

Throughput of the *getMap* action

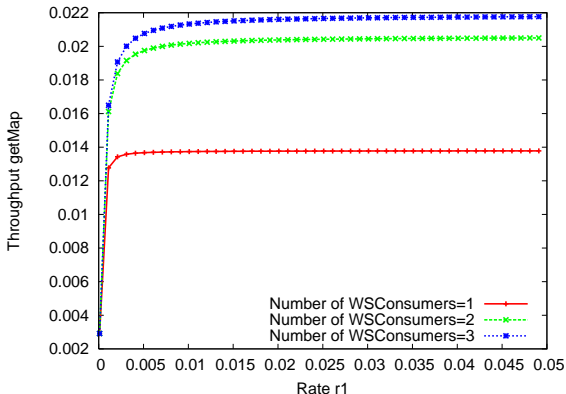


As the number of customers varies between 1 and 30 for various numbers of copies of the *WSPProvider* component.

Throughput of the *getMap* action

- Under heavy load increasing the number of providers initially leads to a sharp increase in the throughput. However the gain deteriorates so that the system with four copies is just 8.7% faster than the system with three.
- In the following we settle on three copies of *WSProvider*.

Throughput of *getMap* action



As the request arrival rate (r_1) varies for differing numbers of *WSConsumer*.


Throughput of *getMap* action

- Every line starts to plateau at approximately $r_1 = 0.010$ following an initial sharp increase. This suggests that the user is the bottle next in the system when the arrival rate is lower. Conversely, at high rates the system becomes congested.
- Whilst having two copies of *WSConsumer*, corresponding to two operating threads of control, improves performance significantly, the subsequent increase with three copies is less pronounced.
- So we set the number of copies of *WSConsumer* to 2.

Optimising the number of copies of *PAP*Provider

- Here we are particularly interested in the overall impact of the rate at which the validity check is performed.
- Slower rates may mean more computationally expensive validation.
- Faster rates may involve less accuracy and lower security of the system.

Throughput of *getMap* action



As the validity check rate (r_5) varies for differing numbers of *PAPProvider*.

Throughput of *getMap* action

- A sharp increase followed by a constant levelling off suggests that optimal rate values lie on the left of the plateau, as faster rates do not improve the system considerably.
- As for the optimal number of copies of *PAPProvider*, deploying two copies rather than one dramatically increases the quality of service of the overall system.
- With a similar approach as previously discussed, the modeller may want to consider the trade-off between the cost of adding a third copy and the throughput increase.

An alternative design for *PAProvider*

- The original design of *PAProvider* is **stateless**.
- Any of its services can be called at any point, the correctness of the system being guaranteed by implementation-specific constraints such as session identifiers being uniquely assigned to the clients and passed as parameters of the method calls.
- Alternatively we may consider a **stateful** implementation, modelled as a sequential component with three local states.
- This implementation has the consequence that there can never be more than N_{PAP} *WSProvider* which have started a session with a *PAProvider*

An alternative design for *PAP*Provider


- The original design of *PAP*Provider is **stateless**.
- Any of its services can be called at any point, the correctness of the system being guaranteed by implementation-specific constraints such as session identifiers being uniquely assigned to the clients and passed as parameters of the method calls.
- Alternatively we may consider a **stateful** implementation, modelled as a sequential component with three local states.
- This implementation has the consequence that there can never be more than N_{PAP} *WSP*Provider which have started a session with a *PAP*Provider

Component *PAProvider* — Stateful Version

It maintains a thread for each session and carries out the validity check on behalf of the Web Service Provider.

$$\begin{aligned} PAProvider &\stackrel{def}{=} (startSession, \top).PAProvider_2 \\ PAProvider_2 &\stackrel{def}{=} (checkValid, r_5).PAProvider_3 \\ PAProvider_3 &\stackrel{def}{=} (stopSession, \top).PAProvider \end{aligned}$$

Throughput of *getMap* action



As the validity check rate (r_5) varies for differing numbers of *PAPProvider* (stateful version).

Throughput of *getMap* action

- In this case the incremental gain in adding more copies has become more marked.
- However, the modeller may want to prefer the original version, as three copies of the stateful provider deliver about as much as the throughput of only one copy of the stateless implementation.

Acknowledgements

The models of Roland the Gunslinger are due to Allan Clark and the model of the web service composition is joint work with Mirco Tribastone.

Outline

- 1 Recap
- 2 Case Studies
- 3 Roland the Gunslinger
- 4 Tools
- 5 Web Service Composition
- 6 Querying models**

Querying models

So far we have focussed on the construction of the model and demonstrated the use of some particular examples to derive quantitative measures.

PEPA is complemented by a couple of formal approaches to query models.

- stochastic model checking based on CSL formulae; and
- (eXtended) stochastic probes within the PEPA model.

Querying models

So far we have focussed on the construction of the model and demonstrated the use of some particular examples to derive quantitative measures.

PEPA is complemented by a couple of formal approaches to query models.

- stochastic model checking based on CSL formulae; and
- (eXtended) stochastic probes within the PEPA model.

Model checking

Model checking requires two inputs:

- a description of the system, usually given in some high-level modelling formalism such as a process algebra description, or a Petri net;
- a specification of one or more desired properties of the system, normally using temporal logics such as CTL (Computational Tree Logic) or LTL (Linear-time Temporal Logic).

Model checking

Model checking requires two inputs:

- a description of the system, usually given in some high-level modelling formalism such as a process algebra description, or a Petri net;
- a specification of one or more desired properties of the system, normally using temporal logics such as CTL (Computational Tree Logic) or LTL (Linear-time Temporal Logic).

Model checking

From the high-level description the model checker constructs a **labelled transition system** which captures all possible behaviours of the system.

The **model checking algorithms** then automatically verify whether or not each property is satisfied in the system.

Model checking

From the high-level description the model checker constructs a **labelled transition system** which captures all possible behaviours of the system.

The **model checking algorithms** then automatically verify whether or not each property is satisfied in the system.

Stochastic model checking

In **stochastic model checking** it is assumed that the labelled transition system is a **Continuous Time Markov Chain (CTMC)**.

This makes stochastic process algebras suitable high-level language for stochastic model checking.

The logic is also enhanced to query not just **logical behaviour** (whether some property is satisfied or not) but also **quantified behaviour** (e.g. the probability that a property is satisfied at a particular time).

Stochastic model checking

In [stochastic model checking](#) it is assumed that the labelled transition system is a [Continuous Time Markov Chain \(CTMC\)](#).

This makes stochastic process algebras suitable high-level language for stochastic model checking.

The logic is also enhanced to query not just [logical behaviour](#) (whether some property is satisfied or not) but also [quantified behaviour](#) (e.g. the probability that a property is satisfied at a particular time).

Stochastic model checking

In [stochastic model checking](#) it is assumed that the labelled transition system is a [Continuous Time Markov Chain \(CTMC\)](#).

This makes stochastic process algebras suitable high-level language for stochastic model checking.

The logic is also enhanced to query not just [logical behaviour](#) (whether some property is satisfied or not) but also [quantified behaviour](#) (e.g. the probability that a property is satisfied at a particular time).

Model checking

There are two broad approaches to model checking:

- **Explicit state model checking** (exhaustive exploration for all possible states/executions): exact results obtained via numerical computation.
- **Statistical model-checking** (discrete event simulation and sampling over multiple runs): approximate results.

PRISM model and model checking

- Probabilistic model checking in PRISM is based on a CTMC and the logic CSL.
- Formally the mapping from PEPA is based on the structured operational semantics, generating the underlying CTMC in the usual way.
- In practice PEPA is an input language for PRISM with a direct mapping between PEPA components and the interacting, reactive modules of the PRISM input language.
- Note, however, that this places a restriction to have synchronisations in which only one participant is active as PRISM cannot handle the apparent rate based calculations of cooperation in PEPA.

PRISM model and model checking

- Probabilistic model checking in PRISM is based on a CTMC and the logic CSL.
- Formally the mapping from PEPA is based on the structured operational semantics, generating the underlying CTMC in the usual way.
- In practice PEPA is an input language for PRISM with a direct mapping between PEPA components and the interacting, reactive modules of the PRISM input language.
- Note, however, that this places a restriction to have synchronisations in which only one participant is active as PRISM cannot handle the apparent rate based calculations of cooperation in PEPA.

PRISM model and model checking

- Probabilistic model checking in PRISM is based on a CTMC and the logic CSL.
- Formally the mapping from PEPA is based on the structured operational semantics, generating the underlying CTMC in the usual way.
- In practice PEPA is an input language for PRISM with a direct mapping between PEPA components and the interacting, reactive modules of the PRISM input language.
- Note, however, that this places a restriction to have synchronisations in which only one participant is active as PRISM cannot handle the apparent rate based calculations of cooperation in PEPA.

PRISM model and model checking

- Probabilistic model checking in PRISM is based on a CTMC and the logic CSL.
- Formally the mapping from PEPA is based on the structured operational semantics, generating the underlying CTMC in the usual way.
- In practice PEPA is an input language for PRISM with a direct mapping between PEPA components and the interacting, reactive modules of the PRISM input language.
- Note, however, that this places a restriction to have synchronisations in which only one participant is active as PRISM cannot handle the apparent rate based calculations of cooperation in PEPA.

The CSL logic

The syntax of CSL is as follows:

$$\begin{aligned} \phi ::= & \text{true} \mid a \mid \neg\phi \mid \phi \wedge \phi \mid \\ & \mathbf{P}_{\sim p}[\phi \mathbf{U}^I \phi] \mid \mathbf{S}_{\sim p}[\phi] \mid \\ & \mathbf{R}_{\sim r}[I^=t] \mid \mathbf{R}_{\sim r}[C^{\leq t}] \mid \mathbf{R}_{\sim r}[\mathbf{F} \phi] \mid \mathbf{R}_{\sim r}[\mathbf{S}] \end{aligned}$$

where a is an **atomic proposition**, $\sim \in \{<, \leq, \geq, >\}$, $p \in [0, 1]$, I is an interval of $\mathbb{R}^{\geq 0}$ and $r, t \in \mathbb{R}^{\geq 0}$.

P and **S** are **probabilistic operators** which include a **probabilistic bound** $\sim p$.

R is a **reward operator** with a **reward bound** $\sim r$.

Probabilistic operators

A formula $\mathbf{P}_{\sim p}[\phi \mathbf{U}' \phi]$ is true in a state s if the probability of the formula $(\phi \mathbf{U}' \phi)$ being satisfied from state s meets the bound $\sim p$.

A formula of type $\phi_1 \mathbf{U}' \phi_2$ is an **until** formula.

It is true of a path σ through the state space if, for some time instant $t \in I$, at time t in the path σ the CSL subformula ϕ_2 is true and the subformula ϕ_1 is true at all preceding time instants.

A formula $\mathbf{S}_{\sim p}[\phi]$ is true in state s if the probability that the formula ϕ being satisfied in a steady state reached from state s meets the bound $\sim p$.

Probabilistic operators

A formula $\mathbf{P}_{\sim p}[\phi \mathbf{U}' \phi]$ is true in a state s if the probability of the formula $(\phi \mathbf{U}' \phi)$ being satisfied from state s meets the bound $\sim p$.

A formula of type $\phi_1 \mathbf{U}' \phi_2$ is an **until** formula.

It is true of a path σ through the state space if, for some time instant $t \in I$, at time t in the path σ the CSL subformula ϕ_2 is true and the subformula ϕ_1 is true at all preceding time instants.

A formula $\mathbf{S}_{\sim p}[\phi]$ is true in state s if the probability that the formula ϕ being satisfied in a steady state reached from state s meets the bound $\sim p$.

Probabilistic operators

A formula $\mathbf{P}_{\sim p}[\phi \mathbf{U}' \phi]$ is true in a state s if the probability of the formula $(\phi \mathbf{U}' \phi)$ being satisfied from state s meets the bound $\sim p$.

A formula of type $\phi_1 \mathbf{U}' \phi_2$ is an **until** formula.

It is true of a path σ through the state space if, for some time instant $t \in I$, at time t in the path σ the CSL subformula ϕ_2 is true and the subformula ϕ_1 is true at all preceding time instants.

A formula $\mathbf{S}_{\sim p}[\phi]$ is true in state s if the probability that the formula ϕ being satisfied in a steady state reached from state s meets the bound $\sim p$.

The CSL Reward operator

The CSL **reward operator** \mathbf{R} is used to express properties concerning the expected value of rewards.

$\mathbf{R}_{\sim r}[I=t]$ asserts that the expected value of the state reward at time instant t meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[C \leq t]$ refers to the expected reward accumulated up until t .

$\mathbf{R}_{\sim r}[\mathbf{F} \phi]$ asserts that the expected reward accumulated before a state satisfying ϕ is reached meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[\mathbf{S}]$ asserts that the long-run/steady state expected reward meets the bound $\sim r$.

The CSL Reward operator

The CSL **reward operator** \mathbf{R} is used to express properties concerning the expected value of rewards.

$\mathbf{R}_{\sim r}[I=t]$ asserts that the expected value of the state reward at time instant t meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[C \leq t]$ refers to the expected reward accumulated up until t .

$\mathbf{R}_{\sim r}[\mathbf{F} \phi]$ asserts that the expected reward accumulated before a state satisfying ϕ is reached meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[\mathbf{S}]$ asserts that the long-run/steady state expected reward meets the bound $\sim r$.

The CSL Reward operator

The CSL **reward operator** \mathbf{R} is used to express properties concerning the expected value of rewards.

$\mathbf{R}_{\sim r}[I=t]$ asserts that the expected value of the state reward at time instant t meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[C^{\leq t}]$ refers to the expected reward accumulated up until t .

$\mathbf{R}_{\sim r}[\mathbf{F} \phi]$ asserts that the expected reward accumulated before a state satisfying ϕ is reached meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[\mathbf{S}]$ asserts that the long-run/steady state expected reward meets the bound $\sim r$.

The CSL Reward operator

The CSL **reward operator** \mathbf{R} is used to express properties concerning the expected value of rewards.

$\mathbf{R}_{\sim r}[I=t]$ asserts that the expected value of the state reward at time instant t meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[C \leq t]$ refers to the expected reward accumulated up until t .

$\mathbf{R}_{\sim r}[\mathbf{F} \phi]$ asserts that the expected reward accumulated before a state satisfying ϕ is reached meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[\mathbf{S}]$ asserts that the long-run/steady state expected reward meets the bound $\sim r$.

The CSL Reward operator

The CSL **reward operator** \mathbf{R} is used to express properties concerning the expected value of rewards.

$\mathbf{R}_{\sim r}[I=t]$ asserts that the expected value of the state reward at time instant t meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[C \leq t]$ refers to the expected reward accumulated up until t .

$\mathbf{R}_{\sim r}[\mathbf{F} \phi]$ asserts that the expected reward accumulated before a state satisfying ϕ is reached meets the bound $\sim r$.

$\mathbf{R}_{\sim r}[\mathbf{S}]$ asserts that the long-run/steady state expected reward meets the bound $\sim r$.

Computation in PRISM

The underlying computation in PRISM for explicit state model checking involves a combination of:

graph-theoretical algorithms, for conventional temporal logic model checking and **qualitative** probabilistic model checking;

numerical computation, for **quantitative** probabilistic model checking, i.e. calculation of probabilities and reward values.

Statistical model checking

The basic idea of statistical model checking is to **simulate** the system for finitely many runs, and use **statistics** to infer whether the samples provide **evidence** for the satisfaction or violation of the property of interest.

Advantages of statistical model checking

- Much larger models can be handled since the state space does not need to be constructed and stored all at once.
- Since the approach is based on **observations** and **samples** it can be applied to any system which is **executable** — the underlying stochastic process does not need to be a CTMC.
- Since many independent samples are required it is susceptible to **coarse-grained parallelization**.

Advantages of statistical model checking

- Much larger models can be handled since the state space does not need to be constructed and stored all at once.
- Since the approach is based on **observations** and **samples** it can be applied to any system which is **executable** — the underlying stochastic process does not need to be a CTMC.
- Since many independent samples are required it is susceptible to **coarse-grained parallelization**.

Advantages of statistical model checking

- Much larger models can be handled since the state space does not need to be constructed and stored all at once.
- Since the approach is based on **observations** and **samples** it can be applied to any system which is **executable** — the underlying stochastic process does not need to be a CTMC.
- Since many independent samples are required it is susceptible to **coarse-grained parallelization**.

Advantages of statistical model checking

- Much larger models can be handled since the state space does not need to be constructed and stored all at once.
- Since the approach is based on **observations** and **samples** it can be applied to any system which is **executable** — the underlying stochastic process does not need to be a CTMC.
- Since many independent samples are required it is susceptible to **coarse-grained parallelization**.

These advantages are off-set by the disadvantage that it is an **approximation** compared with the exact, explicit state approach.

General framework

Consider a CTMC \mathcal{X} and a property ϕ .

An **execution** or **run** of \mathcal{X} is a, possibly infinite, sequence of states in \mathcal{X} .

We wish to decide $\mathbf{P}_{\sim p}[\phi]$, i.e. whether \mathcal{X} satisfies ϕ with probability satisfying the bound $\sim p$.

The result of each execution is taken to be the result of a Bernoulli trial, 0 or 1, according to whether ϕ is satisfied or not.

Let q be the probability that ϕ is satisfied, then we seek to establish if $q \sim p$.

General framework

Consider a CTMC \mathcal{X} and a property ϕ .

An **execution** or **run** of \mathcal{X} is a, possibly infinite, sequence of states in \mathcal{X} .

We wish to decide $\mathbf{P}_{\sim p}[\phi]$, i.e. whether \mathcal{X} satisfies ϕ with probability satisfying the bound $\sim p$.

The result of each execution is taken to be the result of a Bernoulli trial, 0 or 1, according to whether ϕ is satisfied or not.

Let q be the probability that ϕ is satisfied, then we seek to establish if $q \sim p$.

General framework

Consider a CTMC \mathcal{X} and a property ϕ .

An **execution** or **run** of \mathcal{X} is a, possibly infinite, sequence of states in \mathcal{X} .

We wish to decide $\mathbf{P}_{\sim p}[\phi]$, i.e. whether \mathcal{X} satisfies ϕ with probability satisfying the bound $\sim p$.

The result of each execution is taken to be the result of a Bernoulli trial, 0 or 1, according to whether ϕ is satisfied or not.

Let q be the probability that ϕ is satisfied, then we seek to establish if $q \sim p$.

General framework

Consider a CTMC \mathcal{X} and a property ϕ .

An **execution** or **run** of \mathcal{X} is a, possibly infinite, sequence of states in \mathcal{X} .

We wish to decide $\mathbf{P}_{\sim p}[\phi]$, i.e. whether \mathcal{X} satisfies ϕ with probability satisfying the bound $\sim p$.

The result of each execution is taken to be the result of a Bernoulli trial, 0 or 1, according to whether ϕ is satisfied or not.

Let q be the probability that ϕ is satisfied, then we seek to establish if $q \sim p$.

General framework

Consider a CTMC \mathcal{X} and a property ϕ .

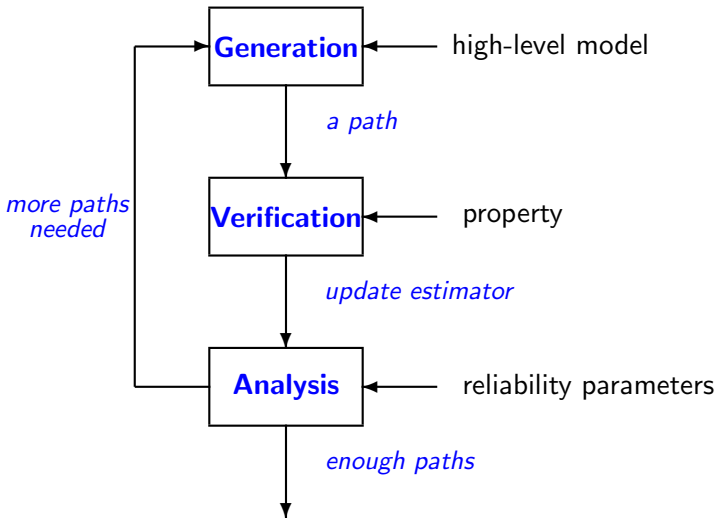
An **execution** or **run** of \mathcal{X} is a, possibly infinite, sequence of states in \mathcal{X} .

We wish to decide $\mathbf{P}_{\sim p}[\phi]$, i.e. whether \mathcal{X} satisfies ϕ with probability satisfying the bound $\sim p$.

The result of each execution is taken to be the result of a Bernoulli trial, 0 or 1, according to whether ϕ is satisfied or not.

Let q be the probability that ϕ is satisfied, then we seek to establish if $q \sim p$.

Schematic for statistical model checking



References

- A. Clark, S. Gilmore, J. Hillston, and M. Tribastone, chapter *Stochastic Process Algebras*, in Formal Methods for Performance Evaluation: 7th Intl. School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2007, LNCS 4486, pages 132-179. Springer-Verlag, Bertinoro, Italy, May-June 2007.
- <http://www.dcs.ed.ac.uk/pepa/>
- <http://www.prismmodelchecker.org/>