

Process algebras for quantitative analysis: Lecture 3 — Case Studies

Jane Hillston

School of Informatics
The University of Edinburgh
Scotland

10th February 2010

Outline

- 1 Recap
- 2 Case Studies
- 3 Upgrading a PC LAN
- 4 Roland the Gunslinger
- 5 Web Service Composition

Outline

- 1 Recap
- 2 Case Studies
- 3 Upgrading a PC LAN
- 4 Roland the Gunslinger
- 5 Web Service Composition

Dynamic behaviour

- The behaviour of a model is dictated by the [semantic rules](#) governing the combinators of the language.

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**.

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.

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 code:

```

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

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

P1 <- P1
  
```

The right sidebar shows a table with the following data:

Utilisation	Throughput	Population
Action	Throughput	
run	0.6666666666666667	0.6666666666666667
start	0.6666666666666667	0.6666666666666667
stop	0.6666666666666667	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 Upgrading a PC LAN
- 4 Roland the Gunslinger
- 5 Web Service Composition

PEPA Case Studies (1)

- Multiprocessor access-contention protocols ([Gilmore, Hillston and Ribaudo](#), [Edinburgh and Turin](#))

PEPA Case Studies (1)

- Multiprocessor access-contention protocols ([Gilmore, Hillston and Ribaudó, Edinburgh and Turin](#))
- Protocols for fault-tolerant systems ([Clark, Gilmore, Hillston and Ribaudó, Edinburgh and Turin](#))

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](#))

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](#))

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](#))

PEPA Case Studies (1)

- Multiprocessor access-contention protocols ([Gilmore, Hillston and Ribaudó, Edinburgh and Turin](#))
- Protocols for fault-tolerant systems ([Clark, Gilmore, Hillston and Ribaudó, 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](#))

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](#))

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](#))

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](#))

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](#))

Grid Scheduling

In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.

Grid Scheduling

In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.




P1

P2

P3

Grid Scheduling

In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.



P1


P2

P3

A **schedule** maps tasks to processors

Grid Scheduling


In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.



In a grid processors are **heterogeneous**...

Grid Scheduling


In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.



In a grid processors are heterogeneous and **dynamic**.

Grid Scheduling


In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.



- Current performance parameters obtained from the monitoring.

Grid Scheduling


In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.



- Current performance parameters obtained from the monitoring.
- Highly **abstract model components** configured to represent different scheduling possibilities.

Grid Scheduling


In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.



- Current performance parameters obtained from the monitoring.
- Highly **abstract model components** configured to represent different scheduling possibilities.
- Fast evaluation and comparison of alternatives.

Grid Scheduling


In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.



- Current performance parameters obtained from the monitoring.
- Highly **abstract model components** configured to represent different scheduling possibilities.
- Fast evaluation and comparison of alternatives.

Grid Scheduling

In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.




The dynamic nature may mean that tasks have to be re-scheduled during the course of the computation.

- Current performance parameters obtained from the monitoring.
- Highly **abstract model components** configured to represent different scheduling possibilities.
- Fast evaluation and comparison of alternatives.

Grid Scheduling

In the **ENHANCE** research project we investigated ways to use performance predictions to improve scheduling decisions in large computational grids.



The dynamic nature may mean that tasks have to be **re-scheduled** during the course of the computation.

- Current performance parameters obtained from the monitoring.
- Highly **abstract model components** configured to represent different scheduling possibilities.
- Fast evaluation and comparison of alternatives.


PEPA Case Studies (3)

- Locks and movable bridges in inland shipping in Belgium ([Knapen](#), [Hasselt](#))




PEPA Case Studies (3)

- Locks and movable bridges in inland shipping in Belgium ([Knapen, Hasselt](#))
- Robotic workcells ([Holton, Gilmore and Hillston, Bradford and Edinburgh](#))




PEPA Case Studies (3)

- Locks and movable bridges in inland shipping in Belgium ([Knapen](#), [Hasselt](#))
- Robotic workcells ([Holton](#), [Gilmore and Hillston](#), [Bradford and Edinburgh](#))
- Cellular telephone networks ([Kloul](#), [Fourneau and Valois](#), [Versailles](#))



PEPA Case Studies (3)

- Locks and movable bridges in inland shipping in Belgium ([Knapen, Hasselt](#))
- Robotic workcells ([Holton, Gilmore and Hillston, Bradford and Edinburgh](#))
- Cellular telephone networks ([Kloul, Fourneau and Valois, Versailles](#))
- Automotive diagnostic expert systems ([Console, Picardi and Ribaud, Turin](#))



Outline

- 1 Recap
- 2 Case Studies
- 3 Upgrading a PC LAN**
- 4 Roland the Gunslinger
- 5 Web Service Composition

Upgrading a PC LAN

Suppose we wish to determine the **mean waiting time** for data packets at a PC connected to a local area network, operating as a token ring.

Upgrading a PC LAN

Suppose we wish to determine the **mean waiting time** for data packets at a PC connected to a local area network, operating as a token ring.

The transmission medium supports no more than **one transmission at any given time**. To resolve conflicts, a token is passed round the network from one node to another in round robin order.

Upgrading a PC LAN

Suppose we wish to determine the **mean waiting time** for data packets at a PC connected to a local area network, operating as a token ring.

The transmission medium supports no more than **one transmission at any given time**. To resolve conflicts, a token is passed round the network from one node to another in round robin order.

A node has control of the medium, i.e. it can transmit, only whilst it holds the token. In a PC LAN every PC corresponds to a node on the network. Other nodes on the network might be peripheral devices such as printers or faxes but for the purposes of this study we make no distinction and assume that all nodes are PCs.

Upgrading a PC LAN

There are currently four PCs (or similar devices) connected to the LAN in a small office, but the company has recently recruited two new employees, each of whom will have a PC. Our task is to find out how the delay experienced by data packets at each PC will be affected if another two PCs are added.

Modelling Assumptions

Each PC can only store **one data packet waiting for transmission at a time**, so at each visit of the token there is either one packet waiting or no packet waiting. The average rate at which each PC generates data packets for transmission is known to be λ .

Modelling Assumptions

Each PC can only store **one data packet waiting for transmission at a time**, so at each visit of the token there is either one packet waiting or no packet waiting. The average rate at which each PC generates data packets for transmission is known to be λ .


We also know the mean duration, d , of a data packet transmission, and the mean time, m , taken for the token to pass from one PC to the next.

Modelling Assumptions

Each PC can only store **one data packet waiting for transmission at a time**, so at each visit of the token there is either one packet waiting or no packet waiting. The average rate at which each PC generates data packets for transmission is known to be λ .

We also know the mean duration, d , of a data packet transmission, and the mean time, m , taken for the token to pass from one PC to the next.

It is assumed that if another data packet is generated, whilst the PC is transmitting, this second data packet must wait for the next visit of the token before it can be transmitted. In other words, each PC can **transmit at most one data packet per visit** of the token.



Modelling the system: choosing components

The first stage in developing a model of the system in PEPA is to determine the components of the system and the actions which they can undertake.

Modelling the system: choosing components

The first stage in developing a model of the system in PEPA is to determine the components of the system and the actions which they can undertake.

It seems clear that one type of component should be used to represent the PCs. Since we are not told that they differ in any way we may assume that the four/six PC may be represented by four/six components with essentially the same behaviour. However since the order in which the token visits nodes is important we do need to distinguish the PCs.

Modelling the system: choosing components

The first stage in developing a model of the system in PEPA is to determine the components of the system and the actions which they can undertake.

It seems clear that one type of component should be used to represent the PCs. Since we are not told that they differ in any way we may assume that the four/six PC may be represented by four/six components with essentially the same behaviour. However since the order in which the token visits nodes is important we do need to distinguish the PCs.

We also need to represent the medium in some way. In fact for the token ring it is only necessary to represent the token as this is sufficient to determine whether a PC can engage in its desired actions or not.

Modelling the system: choosing activities

The description of the PC is very simple in this case. It only has two activities which it can undertake:

- generate a data packet;
- transmit a data packet.

Moreover we are told that it can only hold one data packet at a time and so these activities must be undertaken sequentially.

Modelling the system: choosing activities

The description of the PC is very simple in this case. It only has two activities which it can undertake:

- generate a data packet;
- transmit a data packet.

Moreover we are told that it can only hold one data packet at a time and so these activities must be undertaken sequentially.

This suggests the following PEPA component for the i th PC:

$$\begin{aligned} PC_{i0} &\stackrel{\text{def}}{=} (arrive, \lambda).PC_{i1} \\ PC_{i1} &\stackrel{\text{def}}{=} (serve_i, \mu).PC_{i0} \end{aligned}$$

Modelling the system: choosing activities

The description of the PC is very simple in this case. It only has two activities which it can undertake:

- generate a data packet;
- transmit a data packet.

Moreover we are told that it can only hold one data packet at a time and so these activities must be undertaken sequentially.

This suggests the following PEPA component for the i th PC:

$$PC_{i0} \stackrel{def}{=} (arrive, \lambda).PC_{i1}$$
$$PC_{i1} \stackrel{def}{=} (serve_i, \mu).PC_{i0}$$

However we will see that this needs some refinement.

Modelling the system: choosing activities

For the token we can think of its current state being characterised by its current position. Thus, if there are N PCs in the network the states of the token correspond to the values $\{1, 2, \dots, N\}$.

Modelling the system: choosing activities

For the token we can think of its current state being characterised by its current position. Thus, if there are N PCs in the network the states of the token correspond to the values $\{1, 2, \dots, N\}$.

When it is at the i th PC then the token may

- transmit a token if there is one to transmit and then walk on;
or
- walk on at once if there is no token waiting;

Modelling the system: choosing activities

For the token we can think of its current state being characterised by its current position. Thus, if there are N PCs in the network the states of the token correspond to the values $\{1, 2, \dots, N\}$.

When it is at the i th PC then the token may

- transmit a token if there is one to transmit and then walk on;
- or
- walk on at once if there is no token waiting;

$$Token_i \stackrel{def}{=} (walk_{i+1}, \omega).Token_{i+1} + (serve_i, \mu).(walk_{i+1}, \omega).Token_{i+1}$$

Refining the components

Using the current definition the token will resolve the choice between transmitting a token and walking on **probabilistically** (i.e. through the race condition).

Refining the components

Using the current definition the token will resolve the choice between transmitting a token and walking on **probabilistically** (i.e. through the race condition).

In order to ensure that the choice is made dependent on the state of PC being visited, we add a **walkon** action to the PC when it is empty, and impose a cooperation between the PC and the Token for both **walkon** and **serve**.

Refining the components

Using the current definition the token will resolve the choice between transmitting a token and walking on **probabilistically** (i.e. through the race condition).

In order to ensure that the choice is made dependent on the state of PC being visited, we add a **walkon** action to the PC when it is empty, and impose a cooperation between the PC and the Token for both **walkon** and **serve**.

$$PC_{i0} \stackrel{def}{=} (arrive, \lambda).PC_{i1} + (walkon_2, \omega).PC_{i0}$$
$$PC_{i1} \stackrel{def}{=} (serve_i, \mu).PC_{i0}$$

Complete model: four PC case

$$PC_{10} \stackrel{def}{=} (arrive, \lambda).PC_{11} + (walkon_2, \omega).PC_{10}$$

$$PC_{11} \stackrel{def}{=} (serve_1, \mu).PC_{10}$$

$$PC_{20} \stackrel{def}{=} (arrive, \lambda).PC_{21} + (walkon_3, \omega).PC_{20}$$

$$PC_{21} \stackrel{def}{=} (serve_2, \mu).PC_{20}$$

$$PC_{30} \stackrel{def}{=} (arrive, \lambda).PC_{31} + (walkon_4, \omega).PC_{30}$$

$$PC_{31} \stackrel{def}{=} (serve_3, \mu).PC_{30}$$

$$PC_{40} \stackrel{def}{=} (arrive, \lambda).PC_{41} + (walkon_1, \omega).PC_{40}$$

$$PC_{41} \stackrel{def}{=} (serve_4, \mu).PC_{40}$$

$$Token_1 \stackrel{def}{=} (walk_{on_2}, \omega).Token_2 + (serve_1, \mu).(walk_2, \omega).Token_2$$

$$Token_2 \stackrel{def}{=} (walk_{on_3}, \omega).Token_3 + (serve_2, \mu).(walk_3, \omega).Token_3$$

$$Token_3 \stackrel{def}{=} (walk_{on_4}, \omega).Token_4 + (serve_3, \mu).(walk_4, \omega).Token_4$$

$$Token_4 \stackrel{def}{=} (walk_{on_1}, \omega).Token_1 + (serve_4, \mu).(walk_1, \omega).Token_1$$

$$LAN \stackrel{def}{=} (PC_{10} \parallel PC_{20} \parallel PC_{30} \parallel PC_{40}) \boxtimes_L Token_1$$

where $L = \{walk_{on_1}, walk_{on_2}, walk_{on_3}, walk_{on_4},$
 $serve_1, serve_2, serve_3, serve_4\}$.

Here we have arbitrarily chosen a starting state in which all the PCs are empty and the Token is at PC1.

State space size

N	2	3	4	5	6	8
$ S $	16	48	128	320	768	4096

State space size

N	2	3	4	5	6	8
$ S $	16	48	128	320	768	4096

N	10	20	30
$ S $	2048	4.194304×10^7	6.442450×10^{10}

Performance results

In order to calculate performance measures from the model we must first assign values to the parameters.

The following values were assigned to parameters of the model in both cases (the unit of time is assumed to be milliseconds),

$$\lambda = 0.01 \quad d = 10 \quad \mu = 0.1 \quad m = 1.0 \quad \omega = 1.0$$

Performance results

In order to calculate performance measures from the model we must first assign values to the parameters.

The following values were assigned to parameters of the model in both cases (the unit of time is assumed to be milliseconds),

$$\lambda = 0.01 \quad d = 10 \quad \mu = 0.1 \quad m = 1.0 \quad \omega = 1.0$$

We wish to calculate the average waiting time of data packets at any PC in the network. Since all the PCs are statistically identical, we can arbitrarily choose one as representative—we select PC_1 .

Derivation of performance measures

Roughly speaking, there are three ways in which performance measures can be derived from the steady state distribution of a Markov process.

These different methods can be thought of as corresponding to different types of measure:

- [state-based measures](#), e.g. utilisation;
- [rate-based measures](#), e.g. throughput;
- other measures which fall outside the above categories, e.g. response time.

State-based measures

State-based measures are those which clearly correspond to the probability that the model is in a state, or a subset of states, which satisfy some condition.

State-based measures

State-based measures are those which clearly correspond to the probability that the model is in a state, or a subset of states, which satisfy some condition.

For example, **utilisation** will correspond to those states where a resource is in use.

State-based measures

State-based measures are those which clearly correspond to the probability that the model is in a state, or a subset of states, which satisfy some condition.

For example, **utilisation** will correspond to those states where a resource is in use.

Thus in order to calculate utilisation we sum the steady state probabilities of being in any of the states where the resource is in use.

Rate-based measures

Rate-based measures are those which correspond to the predicted rate at which some event occurs.

Rate-based measures

Rate-based measures are those which correspond to the predicted rate at which some event occurs.

This will be the product of the rate of the event, and the probability that the event is enabled, i.e. the probability of being in one of the states from which the event can occur.

Rate-based measures

Rate-based measures are those which correspond to the predicted rate at which some event occurs.

This will be the product of the rate of the event, and the probability that the event is enabled, i.e. the probability of being in one of the states from which the event can occur.

Thus to calculate the **throughput** of the transmission we consider the probability of being a state where transmission can occur and multiply it by μ the transmission rate.

Calculating waiting time

Waiting time, which is the **residence time** of a data packet at the interface minus the **transmission time** cannot be derived directly from the steady state distribution since it is neither a state-based nor a rate-based performance measure.

Calculating waiting time

Waiting time, which is the **residence time** of a data packet at the interface minus the **transmission time** cannot be derived directly from the steady state distribution since it is neither a state-based nor a rate-based performance measure.

However, if we know the average number of data packets at the interface, and the average throughput, we can apply **Little's law** to calculate residence time

$$\mathcal{R} = \frac{\mathcal{N}}{X}$$

Calculating waiting time

Once we have the residence time the waiting time can be calculated as:

$$W = \frac{N}{X} - 1/\mu$$

Calculating waiting time

Once we have the residence time the waiting time can be calculated as:

$$W = \frac{\mathcal{N}}{X} - 1/\mu$$

There will be a data packet waiting in the interface whenever the first PC is in state PC_{11} . So the **average number of data packets**, \mathcal{N} , is the total probability of being in one of these states.

Calculating waiting time

Once we have the residence time the waiting time can be calculated as:

$$W = \frac{\mathcal{N}}{X} - 1/\mu$$

There will be a data packet waiting in the interface whenever the first PC is in state PC_{11} . So the **average number of data packets**, \mathcal{N} , is the total probability of being in one of these states.

Similarly, a data transmission can occur whenever there is a data packet waiting for transmission and the token is at PC1. So the **average throughput** X will be the rate at which transmission occurs, μ , multiplied by the total probability of being in one of these states.

Predicted waiting time

Based on these expressions and the calculated steady state distributions we find the following values:
for 4 PCs

$$\text{average waiting time, } \mathcal{W} = \frac{0.1333}{0.008666} - \frac{1}{0.1} = 15.3862 - 10 = 5.3862$$

Predicted waiting time

Based on these expressions and the calculated steady state distributions we find the following values:

for 4 PCs

$$\text{average waiting time, } \mathcal{W} = \frac{0.1333}{0.008666} - \frac{1}{0.1} = 15.3862 - 10 = 5.3862$$

for 6 PCs

$$\text{average waiting time, } \mathcal{W} = \frac{0.1719}{0.00828} - \frac{1}{0.1} = 20.7678 - 10 = 10.7678$$

Predicted waiting time

Based on these expressions and the calculated steady state distributions we find the following values:

for 4 PCs

$$\text{average waiting time, } \mathcal{W} = \frac{0.1333}{0.008666} - \frac{1}{0.1} = 15.3862 - 10 = 5.3862$$

for 6 PCs

$$\text{average waiting time, } \mathcal{W} = \frac{0.1719}{0.00828} - \frac{1}{0.1} = 20.7678 - 10 = 10.7678$$

Thus the average waiting time for data packets will almost double when two more PCs are added to the network.

Outline

- 1 Recap
- 2 Case Studies
- 3 Upgrading a PC LAN
- 4 Roland the Gunslinger**
- 5 Web Service Composition

Roland the Gunslinger

- This sequence of small examples are based around a character called Roland Deschain.

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.

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.

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

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,

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

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 way to model this is to have two instances of *Roland* in parallel:

$$Roland_6 \parallel Roland_6$$

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 way to model this is to have two instances of *Roland* in parallel:

$$Roland_6 \parallel Roland_6$$

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

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

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 way to model this is to have two instances of *Roland* in parallel:

$$Roland_6 \parallel Roland_6$$

However, this model does not capture the fact that Roland needs both hands in order to reload either gun. The simplest way to fix this 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 bullets 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.

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 $Roland_{empty}$.

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 $Roland_{empty}$.

Or we can calculate the probability that *Roland* is in the state $Roland_{idle}$ and subtract it from 1.

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 $Roland_{empty}$.

Or we can calculate the probability that *Roland* is in the state $Roland_{idle}$ and subtract it from 1.

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

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 $Roland_{empty}$.

Or we can calculate the probability that *Roland* is in the state $Roland_{idle}$ and subtract it from 1.

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

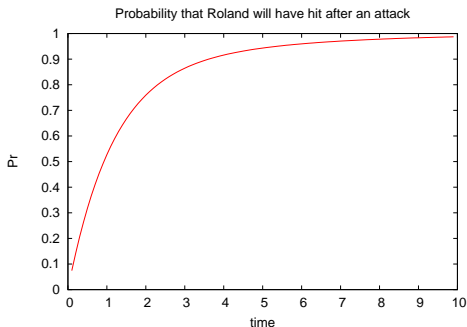
Passage-time analysis allows us to calculate measures such as the probability that Roland has killed his enemy at a given time after he is attacked.

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 five seconds and one occurs about once 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.

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.

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.

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

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

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

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

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

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

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

$$Roland_2 \quad \begin{array}{c} \boxtimes \\ \{hitattack\} \end{array} \quad Enemies_{idle}$$

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.

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.

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.

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.

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.

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.

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.

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


- **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.

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.

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.

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

$$\begin{array}{l}
 \textit{Enemies}_{idle} \\
 \textit{Enemies}_{attack} \\
 + \\
 \\
 \textit{Roland}_2
 \end{array}
 \begin{array}{l}
 \stackrel{def}{=} \\
 \stackrel{def}{=} \\
 \\
 \\
 \boxtimes \\
 \{hit, attack, e-hit\}
 \end{array}
 \begin{array}{l}
 (\textit{attack}, r_{attack}).\textit{Enemies}_{attack} \\
 (e\textit{-hit}, r_{e\textit{-hit}}).\textit{Enemies}_{idle} \\
 (\textit{hit}, \top).\textit{Enemies}_{idle} \\
 \\
 \textit{Enemies}_{idle}
 \end{array}$$

Model Analysis

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

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.

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.

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.

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

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

$$+ (befriend, r_{befriend}).Roland_{idle}$$

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

$$+ (a-hit, \top).Roland_{idle}$$

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

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

$$+ (a-hit, \top).(reload, r_{reload}).Roland_{idle}$$

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

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

$$+ (a-hit, \top).(reload, r_{reload}).Roland_{idle}$$

Synchronising Roland and the Accomplice

- When there is an accomplice, he and Roland fight together against the enemy.

Synchronising Roland and the Accomplice

- When there is an accomplice, he and Roland fight together against the enemy.
- This involves some cooperation between them.

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.

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.

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

$$Acmpl_{idle} \stackrel{def}{=} (attack, \top).Acmpl_2$$

$$Acmpl_2 \stackrel{def}{=} (a-hit, r_{a-hit}).Acmpl_{hit} + (hit, \top).Acmpl_{idle} \\ + (miss, r_{miss}).Acmpl_1 \\ + (enemy-hit, \top).Acmpl_{abs}$$

$$Acmpl_1 \stackrel{def}{=} (a-hit, r_{a-hit}).Acmpl_{hit} \\ + (hit, \top).(reload, r_{a-reload}).Acmpl_{idle} \\ + (miss, r_{miss}).Acmpl_{empty} \\ + (enemy-hit, \top).Acmpl_{abs}$$

$$Acmpl_{hit} \stackrel{def}{=} (enemy-die, \top).(reload, r_{a-reload}).Acmpl_{idle}$$

$$Acmpl_{empty} \stackrel{def}{=} (reload, r_{a-reload}).Acmpl_2 \\ + (enemy-hit, \top).Acmpl_{abs} \\ + (hit, \top).(reload, r_{a-reload}).Acmpl_{idle}$$

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.

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

The system equation is as follows:

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

Model Analysis

Steady-State Analysis

- As before we can determine the probability that Roland is involved in a gun battle at an arbitrary time.

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.

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.

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.

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.

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.

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.

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 Upgrading a PC LAN
- 4 Roland the Gunslinger
- 5 Web Service Composition**

Web Service Composition: Introduction

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

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.

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.

Web Service Composition: Introduction

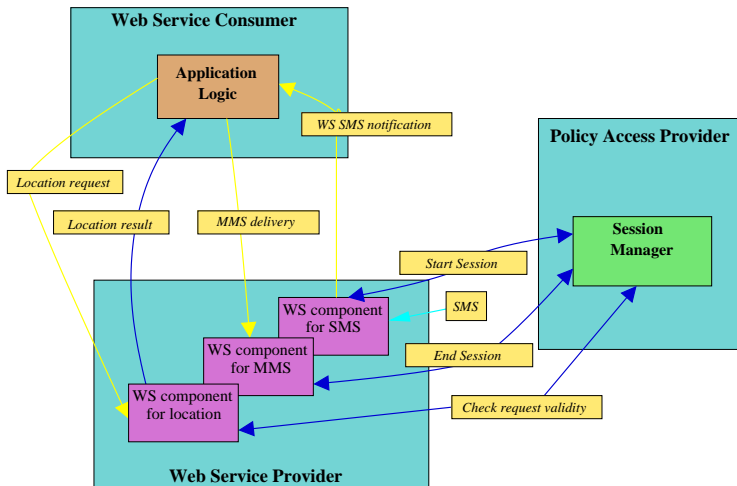
We consider an example of a business application which is composed from a number of offered web services.

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.

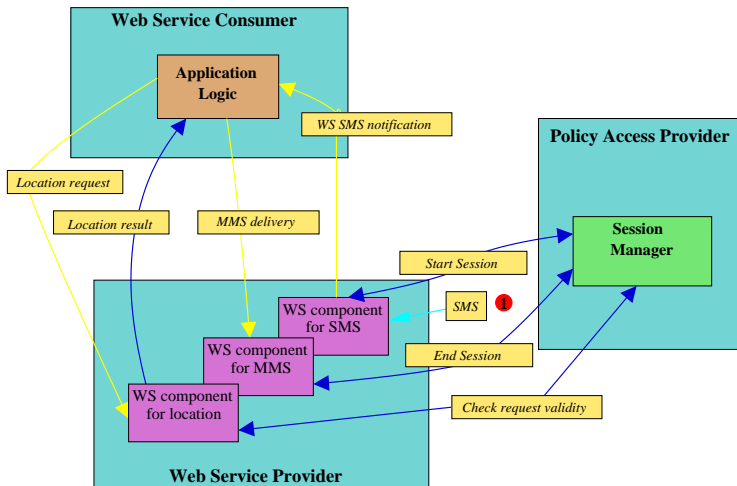
Since the application involves a users' current location there is an access control issue since it must be ensured that the web service consumer has the requisite authority to execute the web service it requests.

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

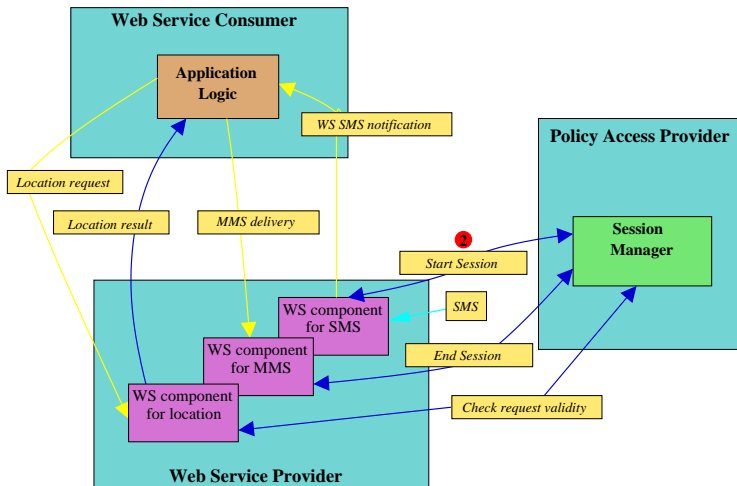
Schematic view



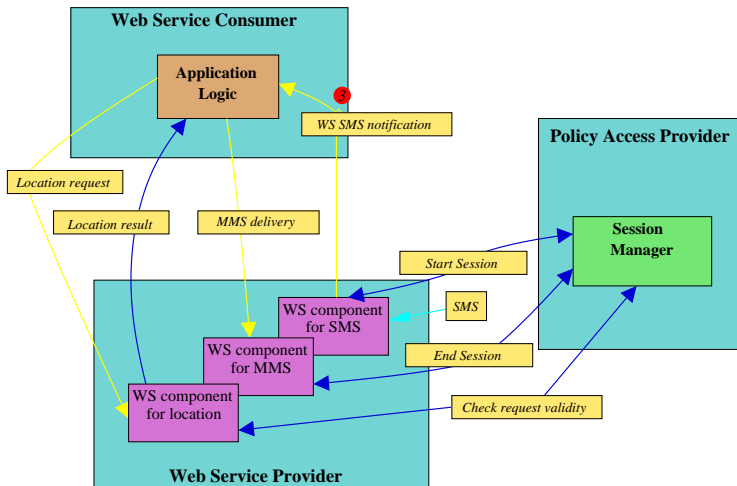
Schematic view



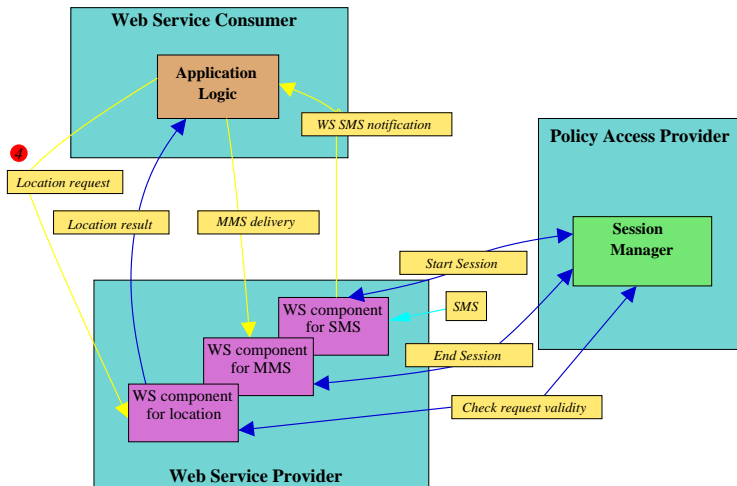
Schematic view



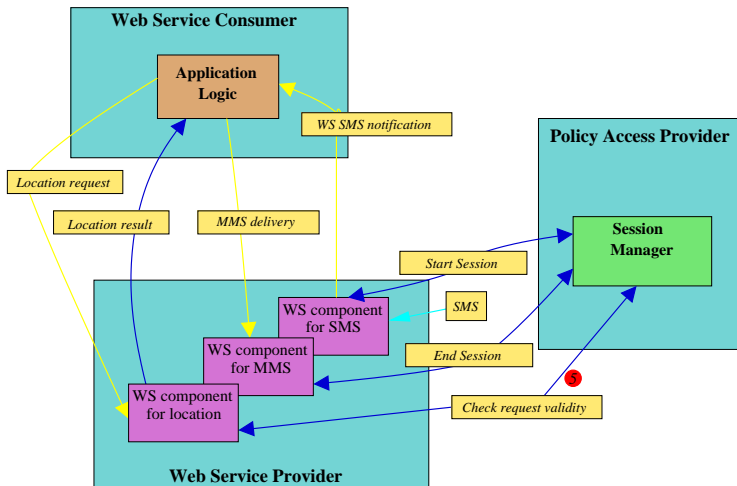
Schematic view



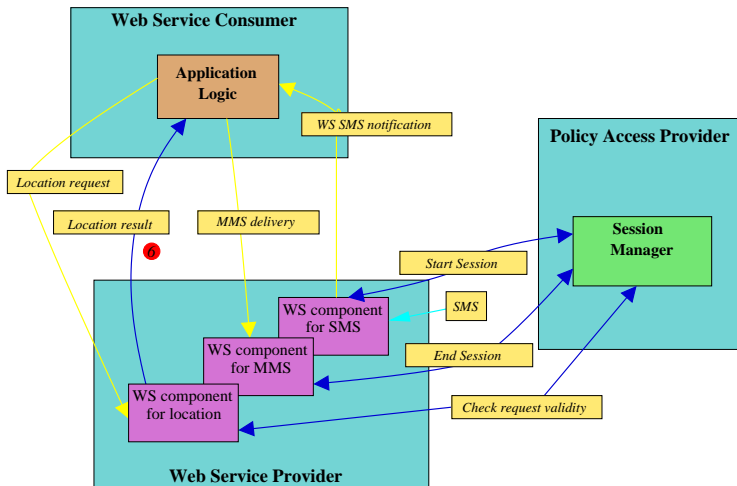
Schematic view



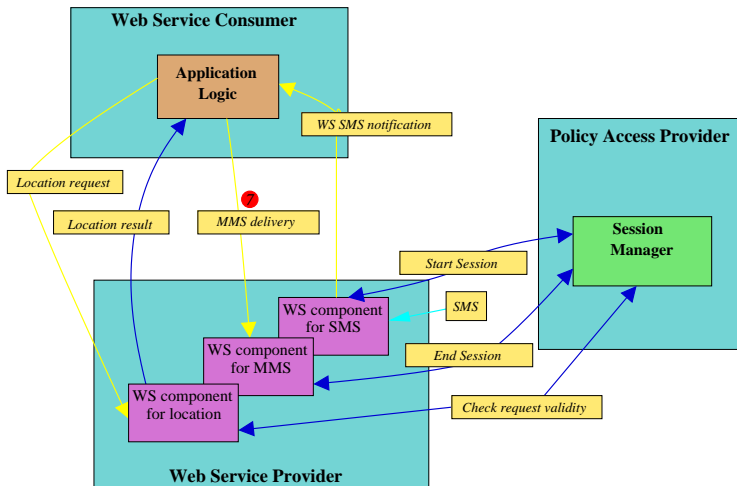
Schematic view



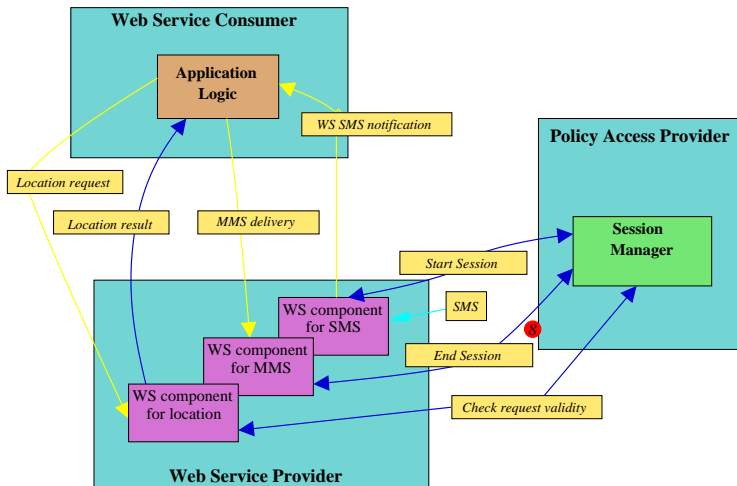
Schematic view




Schematic view



Schematic view



Schematic view



The PEPA model

The PEPA model of the system consists of four components:

The 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 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.

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.

Component *Customer*

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

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

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 *WSConsumer*

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

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.

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.

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.

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.

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*).

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 *WSP*Provider

$$\begin{aligned} WSP_{Provider} &\stackrel{def}{=} (getSMS, \top).WSP_{Provider_2} \\ WSP_{Provider_2} &\stackrel{def}{=} (startSession, r_2).WSP_{Provider_3} \\ WSP_{Provider_3} &\stackrel{def}{=} (notify, r_3).WSP_{Provider_4} \\ WSP_{Provider_4} &\stackrel{def}{=} (locReq, \top).WSP_{Provider_5} \\ WSP_{Provider_5} &\stackrel{def}{=} (checkValid, 99 \cdot \top).WSP_{Provider_6} \\ &+ (checkValid, \top).WSP_{Provider_{10}} \end{aligned}$$

Component *WSPProvider* cont.

$$WSPProvider_6 \stackrel{def}{=} (locRes, r_6).WSPProvider_7$$
$$WSPProvider_7 \stackrel{def}{=} (sendMMS, \top).WSPProvider_8$$
$$WSPProvider_8 \stackrel{def}{=} (getMap, r_8).WSPProvider_9$$
$$WSPProvider_9 \stackrel{def}{=} (stopSession, r_2).WSPProvider$$
$$WSPProvider_{10} \stackrel{def}{=} (locErr, r_6).WSPProvider_{11}$$
$$WSPProvider_{11} \stackrel{def}{=} (get404, r_8).WSPProvider_9$$

Component *PAP*Provider

We consider a stateless implementation of the policy access provider.

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

Model Component $WSComp$

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

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

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.

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.

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.


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.

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 *WSPProvider*.

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.

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 *PAP*Provider — Stateful Version

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

$$\begin{aligned} PAPProvider &\stackrel{def}{=} (startSession, \top).PAPProvider_2 \\ PAPProvider_2 &\stackrel{def}{=} (checkValid, r_5).PAPProvider_3 \\ PAPProvider_3 &\stackrel{def}{=} (stopSession, \top).PAPProvider \end{aligned}$$

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.

Thank you

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.