

Stochastic Process Algebras — From Individuals to Populations

Jane Hillston

Laboratory for Foundations of Computer Science
University of Edinburgh

24th February 2011

Outline



- 1 Introduction
 - Stochastic Process Algebra
- 2 Interpreting SPA for performance modelling
 - Identity and Individuality
 - Collective Dynamics
- 3 Continuous Approximation
 - Numerical illustration
- 4 Example
 - Model
 - Model Evaluation
- 5 Conclusions
 - Alternative interpretations

Outline

- 1 Introduction
 - Stochastic Process Algebra
- 2 Interpreting SPA for performance modelling
 - Identity and Individuality
 - Collective Dynamics
- 3 Continuous Approximation
 - Numerical illustration
- 4 Example
 - Model
 - Model Evaluation
- 5 Conclusions
 - Alternative interpretations



Process Algebra

- Models consist of **agents** which engage in **actions**.





Process Algebra

- Models consist of **agents** which engage in **actions**.





Process Algebra

- Models consist of **agents** which engage in **actions**.





Process Algebra

- Models consist of **agents** which engage in **actions**.





Process Algebra

- Models consist of **agents** which engage in **actions**.



Process Algebra



- Models consist of **agents** which engage in **actions**.



- The structured operational (interleaving) semantics of the language is used to generate a **labelled transition system**.

Process Algebra

- Models consist of **agents** which engage in **actions**.





- The structured operational (interleaving) semantics of the language is used to generate a **labelled transition system**.

Process algebra
model

Process Algebra

- Models consist of **agents** which engage in **actions**.





- The structured operational (interleaving) semantics of the language is used to generate a **labelled transition system**.

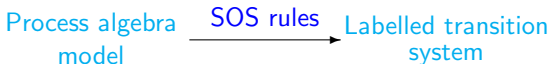


Process Algebra

- Models consist of **agents** which engage in **actions**.





- The structured operational (interleaving) semantics of the language is used to generate a **labelled transition system**.


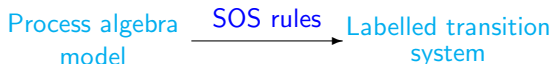


Process Algebra

- Models consist of **agents** which engage in **actions**.



- The structured operational (interleaving) semantics of the language is used to generate a **labelled transition system**.



A simple example: processors and resources

$$Proc_0 \stackrel{def}{=} task1.Proc_1$$

$$Proc_1 \stackrel{def}{=} task2.Proc_0$$

$$Res_0 \stackrel{def}{=} task1.Res_1$$

$$Res_1 \stackrel{def}{=} reset.Res_0$$

$$Proc_0 \parallel_{task1} Res_0$$

A simple example: processors and resources


$$Proc_0 \stackrel{def}{=} task1.Proc_1$$

$$Proc_1 \stackrel{def}{=} task2.Proc_0$$

$$Res_0 \stackrel{def}{=} task1.Res_1$$

$$Res_1 \stackrel{def}{=} reset.Res_0$$

$$Proc_0 \parallel_{task1} Res_0$$



Stochastic process algebras

Process algebras where models are decorated with quantitative information used to generate a stochastic process are [stochastic process algebras \(SPA\)](#).

Stochastic process algebras


Process algebras where models are decorated with quantitative information used to generate a stochastic process are **stochastic process algebras (SPA)**.

This extension was motivated by a desire to bring this **formal** and **compositional** approach to modelling to bear in performance analysis supporting the derivation of measures such as **throughput**, **utilisation** and **response time**.

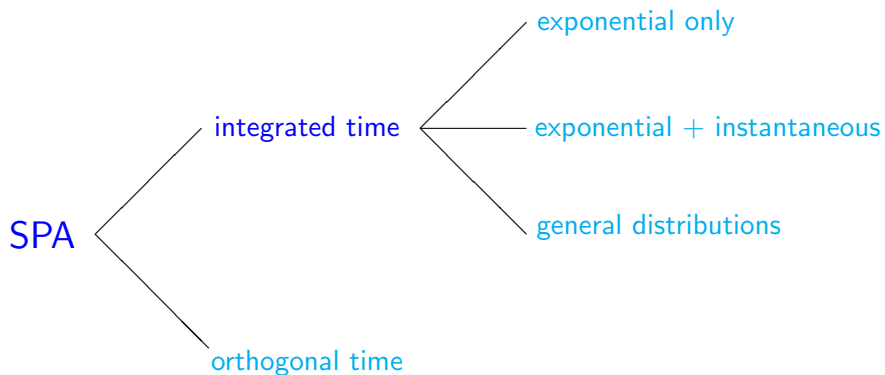
SPA Languages

SPA

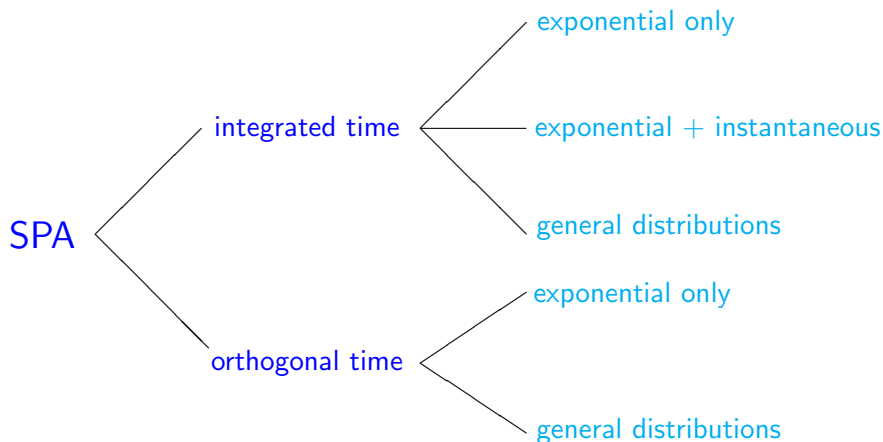
SPA Languages



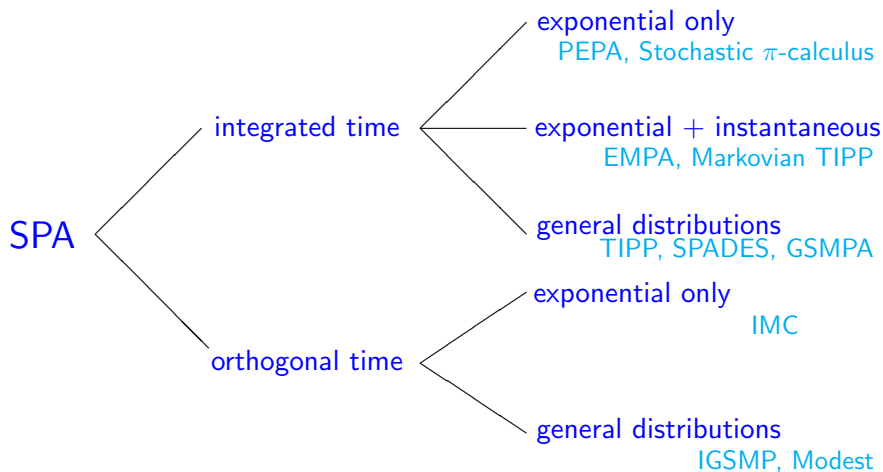
SPA Languages



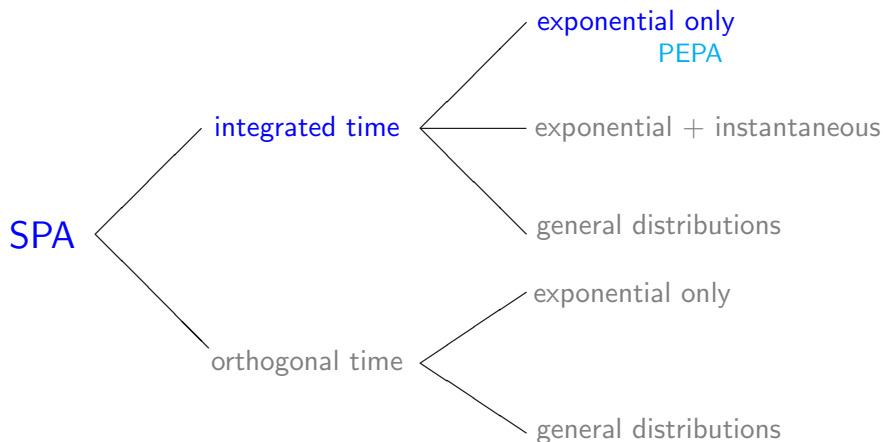
SPA Languages



SPA Languages




SPA Languages




Performance Evaluation Process Algebra

- Models are constructed from **components** which engage in **activities**.



Performance Evaluation Process Algebra


- Models are constructed from **components** which engage in **activities**.



- The language is used to generate a **Continuous Time Markov Chain (CTMC)** for performance modelling.

Performance Evaluation Process Algebra

- Models are constructed from **components** which engage in **activities**.




- The language is used to generate a **Continuous Time Markov Chain (CTMC)** for performance modelling.


PEPA
MODEL

Performance Evaluation Process Algebra

- Models are constructed from **components** which engage in **activities**.




- The language is used to generate a **Continuous Time Markov Chain (CTMC)** for performance modelling.




Performance Evaluation Process Algebra

- Models are constructed from **components** which engage in **activities**.




- The language is used to generate a **Continuous Time Markov Chain (CTMC)** for performance modelling.



Performance Evaluation Process Algebra

- Models are constructed from **components** which engage in **activities**.




- The language is used to generate a **Continuous Time Markov Chain (CTMC)** for performance modelling.



Performance Evaluation Process Algebra

- Models are constructed from **components** which engage in **activities**.



- The language is used to generate a **Continuous Time Markov Chain (CTMC)** for performance modelling.



Integrated analysis


Qualitative verification can now be complemented by quantitative verification.

Integrated analysis

Qualitative verification can now be complemented by **quantitative** verification.

Reachability analysis

How long will it take
for the system to arrive
in a particular state?




Integrated analysis

Qualitative verification can now be complemented by **quantitative** verification.

Model checking

Does a given property ϕ
hold within the system
with a given probability?




Integrated analysis

Qualitative verification can now be complemented by **quantitative** verification.

Model checking

For a given starting state
how long is it until
a given property ϕ holds?



Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

| | |
|-----------------------|--------------|
| $(\alpha, f).P$ | Prefix |
| $P_1 + P_2$ | Choice |
| $P_1 \boxtimes_L P_2$ | Co-operation |
| P/L | Hiding |
| C | Constant |

Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

| | |
|-----------------------|--------------|
| $(\alpha, f).P$ | Prefix |
| $P_1 + P_2$ | Choice |
| $P_1 \boxtimes_L P_2$ | Co-operation |
| P/L | Hiding |
| C | Constant |

Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

| | |
|---------------------|--------------|
| $(\alpha, f).P$ | Prefix |
| $P_1 + P_2$ | Choice |
| $P_1 \bowtie_L P_2$ | Co-operation |
| P/L | Hiding |
| C | Constant |

Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

| | |
|---------------------|--------------|
| $(\alpha, f).P$ | Prefix |
| $P_1 + P_2$ | Choice |
| $P_1 \bowtie_L P_2$ | Co-operation |
| P/L | Hiding |
| C | Constant |

Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

| | |
|---------------------|--------------|
| $(\alpha, f).P$ | Prefix |
| $P_1 + P_2$ | Choice |
| $P_1 \bowtie_L P_2$ | Co-operation |
| P/L | Hiding |
| C | Constant |

Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

| | |
|-----------------------|--------------|
| $(\alpha, f).P$ | Prefix |
| $P_1 + P_2$ | Choice |
| $P_1 \boxtimes_L P_2$ | Co-operation |
| P/L | Hiding |
| C | Constant |

Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

| | |
|-----------------------|--------------|
| $(\alpha, f).P$ | Prefix |
| $P_1 + P_2$ | Choice |
| $P_1 \boxtimes_L P_2$ | Co-operation |
| P/L | Hiding |
| C | Constant |

$P_1 \parallel P_2$ is a derived form for $P_1 \boxtimes_{\emptyset} P_2$.

Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

| | |
|-----------------------|--------------|
| $(\alpha, f).P$ | Prefix |
| $P_1 + P_2$ | Choice |
| $P_1 \boxtimes_L P_2$ | Co-operation |
| P/L | Hiding |
| C | Constant |

$P_1 \parallel P_2$ is a derived form for $P_1 \boxtimes_{\emptyset} P_2$.

When working with large numbers of entities, we write $P[n]$ to denote an **array** of n copies of P executing in parallel.

Performance Evaluation Process Algebra

PEPA **components** perform **activities** either independently or in **co-operation** with other components.

| | |
|-----------------------|--------------|
| $(\alpha, f).P$ | Prefix |
| $P_1 + P_2$ | Choice |
| $P_1 \boxtimes_L P_2$ | Co-operation |
| P/L | Hiding |
| C | Constant |

$P_1 \parallel P_2$ is a derived form for $P_1 \boxtimes_{\emptyset} P_2$.

When working with large numbers of entities, we write $P[n]$ to denote an **array** of n copies of P executing in parallel.

$$P[5] \equiv (P \parallel P \parallel P \parallel P \parallel P)$$

Rates of interaction: bounded capacity

Stochastic process algebras differ in how they define the rate of synchronised actions. In PEPA cooperation between components gives rise to **shared actions**, the rate of which are governed by the assumption of **bounded capacity**.

Rates of interaction: bounded capacity

Stochastic process algebras differ in how they define the rate of synchronised actions. In PEPA cooperation between components gives rise to **shared actions**, the rate of which are governed by the assumption of **bounded capacity**.

Bounded capacity

No component can be made to carry out an action in cooperation faster than its own defined rate for the action.

Rates of interaction: bounded capacity

Stochastic process algebras differ in how they define the rate of synchronised actions. In PEPA cooperation between components gives rise to **shared actions**, the rate of which are governed by the assumption of **bounded capacity**.

Bounded capacity

No component can be made to carry out an action in cooperation faster than its own defined rate for the action.

Thus **shared actions** proceed at the **minimum of the rates** in the participating components.

Rates of interaction: bounded capacity

Stochastic process algebras differ in how they define the rate of synchronised actions. In PEPA cooperation between components gives rise to **shared actions**, the rate of which are governed by the assumption of **bounded capacity**.

Bounded capacity

No component can be made to carry out an action in cooperation faster than its own defined rate for the action.

Thus **shared actions** proceed at the **minimum of the rates** in the participating components.

In contrast **independent actions** do not constrain each other and if there are multiple copies of a action enabled in independent concurrent components their **rates are summed**.

A simple example: processors and resources

$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0 \boxtimes_{\{task1\}} Res_0$$

A simple example: processors and resources


$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0 \boxtimes_{\{task1\}} Res_0$$



A simple example: processors and resources


$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0 \boxtimes_{\{task1\}} Res_0$$



A simple example: processors and resources


$$Proc_0 \stackrel{\text{def}}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{\text{def}}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{\text{def}}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{\text{def}}{=} (reset, r_4).Res_0$$

$$Proc_0 \boxtimes_{\{task1\}} Res_0$$




$$Q = \begin{pmatrix} -R & R & 0 & 0 \\ 0 & -(r_2 + r_4) & r_4 & r_2 \\ r_2 & 0 & -r_2 & 0 \\ r_4 & 0 & 0 & -r_4 \end{pmatrix}$$

Outline

- 1 Introduction
 - Stochastic Process Algebra
- 2 Interpreting SPA for performance modelling
 - Identity and Individuality
 - Collective Dynamics
- 3 Continuous Approximation
 - Numerical illustration
- 4 Example
 - Model
 - Model Evaluation
- 5 Conclusions
 - Alternative interpretations


Solving discrete state models

Under the SOS semantics a SPA model is mapped to a **CTMC** with global states determined by the local states of all the participating components.




Solving discrete state models

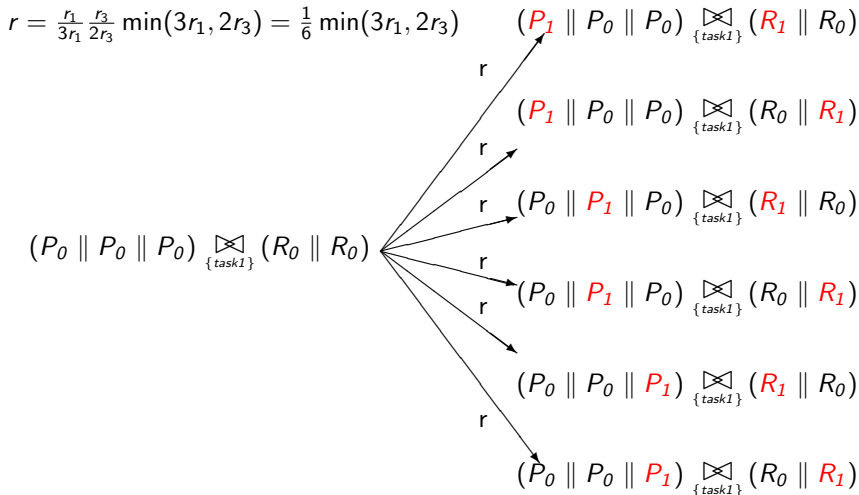
Under the SOS semantics a SPA model is mapped to a **CTMC** with global states determined by the local states of all the participating components.




Modelling at the level of individuals



Modelling at the level of individuals




Solving discrete state models



When the size of the state space is not too large they are amenable to **numerical solution** (linear algebra) to determine a **steady state** or **transient probability distribution**.


Solving discrete state models



When the size of the state space is not too large they are amenable to **numerical solution** (linear algebra) to determine a **steady state** or **transient probability distribution**.

$$Q = \begin{pmatrix} q_{1,1} & q_{1,2} & \cdots & q_{1,N} \\ q_{2,1} & q_{2,2} & \cdots & q_{2,N} \\ \vdots & \vdots & & \vdots \\ q_{N,1} & q_{N,2} & \cdots & q_{N,N} \end{pmatrix}$$

Solving discrete state models



When the size of the state space is not too large they are amenable to **numerical solution** (linear algebra) to determine a **steady state** or **transient probability distribution**.

$$Q = \begin{pmatrix} q_{1,1} & q_{1,2} & \cdots & q_{1,N} \\ q_{2,1} & q_{2,2} & \cdots & q_{2,N} \\ \vdots & \vdots & & \vdots \\ q_{N,1} & q_{N,2} & \cdots & q_{N,N} \end{pmatrix}$$

$$\pi(t) = (\pi_1(t), \pi_2(t), \dots, \pi_N(t))$$

State space explosion

As the number of components, or the complexity of behaviour within components, grows the state space may become so large that it is infeasible to solve the underlying CTMC.

$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0[N_P] \boxtimes_{\{task1\}} Res_0[N_R]$$

State space explosion

As the number of components, or the complexity of behaviour within components, grows the state space may become so large that it is infeasible to solve the underlying CTMC.

$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0[N_P] \boxtimes_{\{task1\}} Res_0[N_R]$$

CTMC interpretation

| Processors (N_P) | Resources (N_R) | States ($2^{N_P+N_R}$) |
|----------------------|---------------------|--------------------------|
| 1 | 1 | 4 |
| 2 | 1 | 8 |
| 2 | 2 | 16 |
| 3 | 2 | 32 |
| 3 | 3 | 64 |
| 4 | 3 | 128 |
| 4 | 4 | 256 |
| 5 | 4 | 512 |
| 5 | 5 | 1024 |
| 6 | 5 | 2048 |
| 6 | 6 | 4096 |
| 7 | 6 | 8192 |
| 7 | 7 | 16384 |
| 8 | 7 | 32768 |
| 8 | 8 | 65536 |
| 9 | 8 | 131072 |
| 9 | 9 | 262144 |
| 10 | 9 | 524288 |
| 10 | 10 | 1048576 |

Achieving aggregation

- If we sacrifice looking at the **identity** of each component we can often achieve substantial state space reduction by **aggregation**.

Achieving aggregation

- If we sacrifice looking at the **identity** of each component we can often achieve substantial state space reduction by **aggregation**.
- This is supported by a shift in how we view the state of a model, based on a **counting abstraction**.

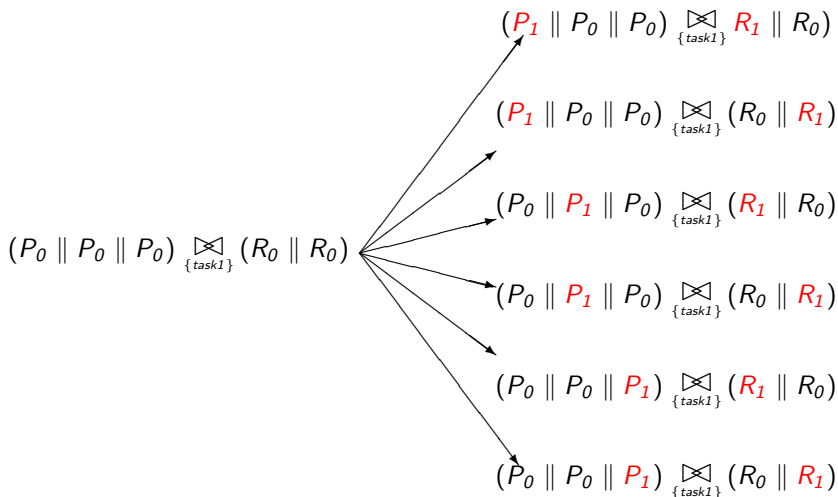
Achieving aggregation

- If we sacrifice looking at the **identity** of each component we can often achieve substantial state space reduction by **aggregation**.
- This is supported by a shift in how we view the state of a model, based on a **counting abstraction**.
- The **syntactic** nature of PEPA (and other SPAs) makes models easily understood by humans, but not so convenient for computers to directly apply these tools and approaches.

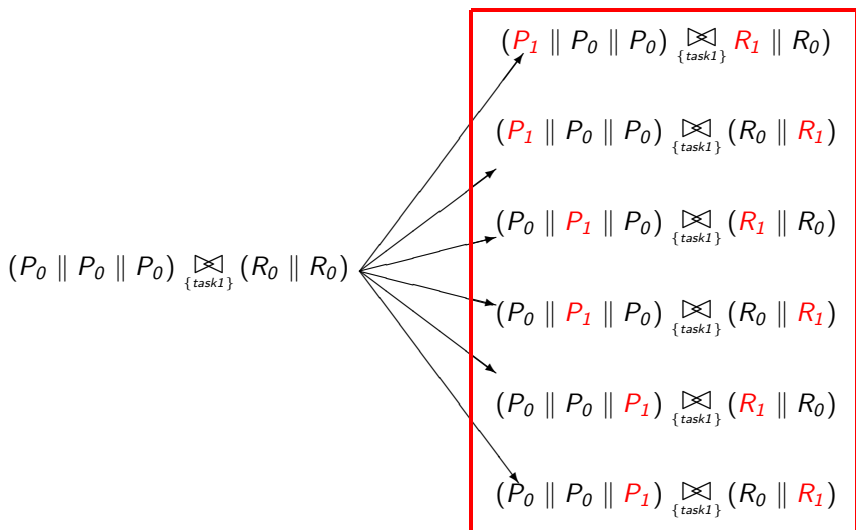
Achieving aggregation

- If we sacrifice looking at the **identity** of each component we can often achieve substantial state space reduction by **aggregation**.
- This is supported by a shift in how we view the state of a model, based on a **counting abstraction**.
- The **syntactic** nature of PEPA (and other SPAs) makes models easily understood by humans, but not so convenient for computers to directly apply these tools and approaches.
- By shifting to a **numerical state representation** we can more readily exploit results such as aggregation and access to alternative mathematical interpretations (i.e. **fluid approximation**).

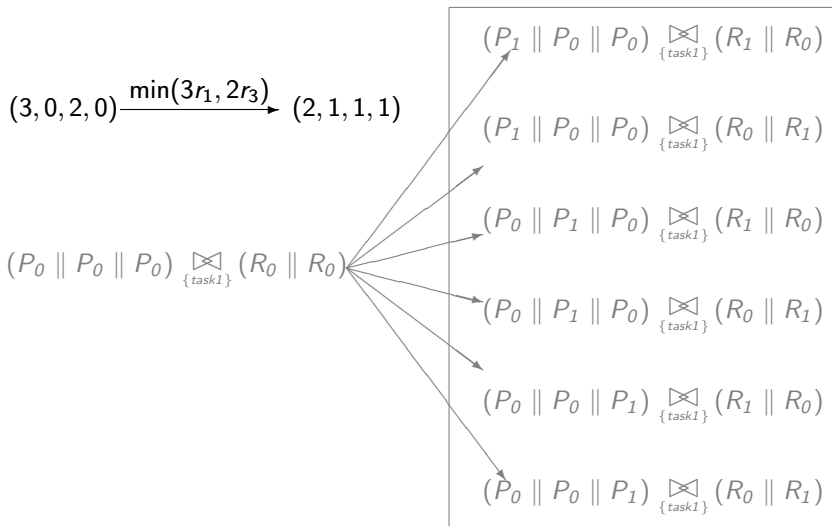
Counting abstraction to generate the *Lumped* CTMC



Counting abstraction to generate the *Lumped* CTMC



Counting abstraction to generate the *Lumped* CTMC



Using this result in practice

There are well-known algorithms such as [Paige and Tarjan](#) for finding the maximal partition of a graph according to some equivalence.

Using this result in practice

There are well-known algorithms such as [Paige and Tarjan](#) for finding the maximal partition of a graph according to some equivalence.

However in practice we would much rather construct the aggregated state space directly.

Using this result in practice

There are well-known algorithms such as [Paige and Tarjan](#) for finding the maximal partition of a graph according to some equivalence.

However in practice we would much rather construct the aggregated state space directly.

The first approach to this used [canonical forms](#) but still worked syntactically to identify states. [Gilmore, Hillston and Ribaud, IEEE TSE 2001].

Using this result in practice

There are well-known algorithms such as [Paige and Tarjan](#) for finding the maximal partition of a graph according to some equivalence.


However in practice we would much rather construct the aggregated state space directly.

The first approach to this used [canonical forms](#) but still worked syntactically to identify states. [Gilmore, Hillston and Ribaud, IEEE TSE 2001].

A more recent approaches shift to a [counting abstraction](#) and a [numerical representation of states and transitions](#).

Solving discrete state models

Even with aggregation models may become too large to solve the underlying CTMC. As an alternative they may be studied using **stochastic simulation**. Each run generates a single trajectory through the state space. Many runs are needed in order to obtain average behaviours.



Discrete Event Simulation


The numerical solution of the CTMC is seen as being the **exact** result.

Discrete Event Simulation


The numerical solution of the CTMC is seen as being the **exact** result.

Analysing a model via **discrete event simulation** can also produce useful results, often with some measure of **confidence** in the results. This is achieved by repeatedly taking random walks through the state space and observing the results, which can be very **computationally intensive**.


100 processors and 80 resources (simulation run A)




100 processors and 80 resources (simulation run B)




100 processors and 80 resources (simulation run C)



100 processors and 80 resources (simulation run D)



100 processors and 80 resources (average of 10 runs)



Collective dynamics

For some SPA models we can make considerable gains in efficiency when solving the model if we take a **collective dynamics** view of the system.

Collective dynamics

For some SPA models we can make considerable gains in efficiency when solving the model if we take a **collective dynamics** view of the system.

Collective dynamics considers the behaviour of populations of similar entities which can interactive with each other in seemingly simple ways to produce phenomena at the population level.

Collective dynamics

For some SPA models we can make considerable gains in efficiency when solving the model if we take a **collective dynamics** view of the system.

Collective dynamics considers the behaviour of populations of similar entities which can interactive with each other in seemingly simple ways to produce phenomena at the population level.

In this case we lose the **identity** of components and even **individuality**, but for many models this is an approximation we are willing to make for the efficiency, or even tractability, of the models.

Collective Behaviour

In the natural world there are many instances of collective behaviour and its consequences:



Collective Behaviour

In the natural world there are many instances of collective behaviour and its consequences:



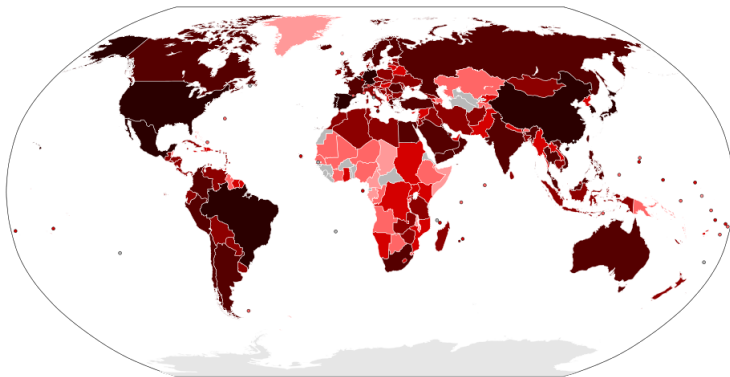
Collective Behaviour

In the natural world there are many instances of collective behaviour and its consequences:



Collective Dynamics

This is also true in the man-made and engineered world:



Spread of H1N1 virus in 2009

Collective Dynamics

This is also true in the man-made and engineered world:



Love Parade, Germany 2006

Collective Dynamics

This is also true in the man-made and engineered world:

The screenshot shows a web browser window titled "HMRC: Login". The address bar contains the URL "https://online.hmrc.gov.uk/login?GAREASONCODE=-1&GARESOURCE=" and the search bar contains "Inland Revenue Tax Returns". The browser's toolbar shows various search engines like Apple, Yahoo!, Google Maps, YouTube, Wikipedia, and News (1075). The page header features the HM Revenue & Customs logo and the text "Online Services" with links for "HMRC home", "Contact us", and "Help".

The main content area is titled "Welcome to Online Services" and is divided into two sections: "Existing users" and "New user".

Existing users

Please enter your User ID and password, then click the 'Login' button below.

Please note: Fields are not case sensitive.

User ID:

Password:

[Digital Certificate user](#)

[Lost User ID?](#)

[Lost password?](#)

[Lost or expired Activation PIN?](#)

[If you have lost both your User ID and password please contact the HM Revenue & Customs \(HMRC\) Online Services Helpdesk.](#)

New user

To register for online services please click the 'Register' button below.

[Digital Certificate user](#)

[Frequently Asked Questions \(FAQs\)](#)

[Computer requirements](#)

[View a demo of HMRC's services](#)

[Registration and Enrolment process](#)

Self assessment tax returns 31st January each year

Process Algebra and Collective Dynamics

Some large process algebra models can be considered to exhibit collective dynamics

Process Algebra and Collective Dynamics

Some large process algebra models can be considered to exhibit collective dynamics

- Each component type captures the behaviour of one type of individual;

Process Algebra and Collective Dynamics

Some large process algebra models can be considered to exhibit collective dynamics

- Each component type captures the behaviour of one type of individual;
- The compositional structure of the model makes explicit interaction between component types;

Process Algebra and Collective Dynamics

Some large process algebra models can be considered to exhibit collective dynamics

- Each component type captures the behaviour of one type of individual;
- The compositional structure of the model makes explicit interaction between component types;
- When there are many instances of the individual component types these may be regarded as a population;

Process Algebra and Collective Dynamics

Some large process algebra models can be considered to exhibit collective dynamics

- Each component type captures the behaviour of one type of individual;
- The compositional structure of the model makes explicit interaction between component types;
- When there are many instances of the individual component types these may be regarded as a population;
- Through the interactions of these populations group or complex behaviours may emerge at the population level.

Population statistics: emergent behaviour

A shift in perspective allows us to model the interactions between individual components but then only consider the system as a whole as an interaction of populations.

Population statistics: emergent behaviour

A shift in perspective allows us to model the interactions between individual components but then only consider the system as a whole as an interaction of populations.

This allows us to model much larger systems than previously possible but in making the shift we are no longer able to collect any information about individuals in the system.

Population statistics: emergent behaviour

A shift in perspective allows us to model the interactions between individual components but then only consider the system as a whole as an interaction of populations.

This allows us to model much larger systems than previously possible but in making the shift we are no longer able to collect any information about individuals in the system.

To characterise the behaviour of a population we **count** the number of individuals within the population that are exhibiting certain behaviours rather than tracking individuals directly.

Population statistics: emergent behaviour

A shift in perspective allows us to model the interactions between individual components but then only consider the system as a whole as an interaction of populations.

This allows us to model much larger systems than previously possible but in making the shift we are no longer able to collect any information about individuals in the system.

To characterise the behaviour of a population we **count** the number of individuals within the population that are exhibiting certain behaviours rather than tracking individuals directly.

Furthermore we make a **continuous approximation** of how the counts vary over time.

Performance as an emergent behaviour

In this framework we must think about the performance of a system from the collective point of view. Service providers often want to do this in any case. For example making contracts in terms of [service level agreements](#).

Performance as an emergent behaviour

In this framework we must think about the performance of a system from the collective point of view. Service providers often want to do this in any case. For example making contracts in terms of [service level agreements](#).

Example Service Level Agreement

90% of requests receive a response within 3 seconds.

Performance as an emergent behaviour

In this framework we must think about the performance of a system from the collective point of view. Service providers often want to do this in any case. For example making contracts in terms of [service level agreements](#).

Example Service Level Agreement

90% of requests receive a response within 3 seconds.

Qualitative Service Level Agreement

Less than 1% of the responses received within 3 seconds will read “System is overloaded, try again later”.

Outline

- 1 Introduction
 - Stochastic Process Algebra
- 2 Interpreting SPA for performance modelling
 - Identity and Individuality
 - Collective Dynamics
- 3 Continuous Approximation
 - Numerical illustration
- 4 Example
 - Model
 - Model Evaluation
- 5 Conclusions
 - Alternative interpretations

Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.

Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.



Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.




Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.



Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.




Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.



Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.



Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.



Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.



Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.



Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.



Continuous Approximation

Use **continuous state variables** to approximate the discrete state space.



Use **ordinary differential equations** to represent the evolution of those variables over time.

New mathematical structures: differential equations

- 1 Use a **counting abstraction** rather than the CTMC complete state space.

New mathematical structures: differential equations

- 1 Use a **counting abstraction** rather than the CTMC complete state space.
- 2 Assume that these state variables are subject to **continuous** rather than **discrete** change.

New mathematical structures: differential equations

- 1 Use a **counting abstraction** rather than the CTMC complete state space.
- 2 Assume that these state variables are subject to **continuous** rather than **discrete** change.
- 3 No longer aim to calculate the probability distribution over the entire state space of the model.

New mathematical structures: differential equations

- 1 Use a **counting abstraction** rather than the CTMC complete state space.
- 2 Assume that these state variables are subject to **continuous** rather than **discrete** change.
- 3 No longer aim to calculate the probability distribution over the entire state space of the model.
- 4 Instead the trajectory of the ODEs estimates the **expected** behaviour of the CTMC.

Models suitable for counting abstraction

- In the **PEPA** language multiple instances of components are represented explicitly — we write $P[n]$ to denote an **array** of n copies of P executing in parallel.

$$P[5] \equiv (P \parallel P \parallel P \parallel P \parallel P)$$

Models suitable for counting abstraction

- In the **PEPA** language multiple instances of components are represented explicitly — we write $P[n]$ to denote an **array** of n copies of P executing in parallel.

$$P[5] \equiv (P \parallel P \parallel P \parallel P \parallel P)$$

- The impact of an action on a counting variable is

Models suitable for counting abstraction

- In the **PEPA** language multiple instances of components are represented explicitly — we write $P[n]$ to denote an **array** of n copies of P executing in parallel.

$$P[5] \equiv (P \parallel P \parallel P \parallel P \parallel P)$$

- The impact of an action on a counting variable is
 - decrease by 1 if the component participates in the action

Models suitable for counting abstraction

- In the **PEPA** language multiple instances of components are represented explicitly — we write $P[n]$ to denote an **array** of n copies of P executing in parallel.

$$P[5] \equiv (P \parallel P \parallel P \parallel P \parallel P)$$

- The impact of an action on a counting variable is
 - decrease by 1 if the component participates in the action
 - increase by 1 if the component is the result of the action

Models suitable for counting abstraction

- In the **PEPA** language multiple instances of components are represented explicitly — we write $P[n]$ to denote an **array** of n copies of P executing in parallel.

$$P[5] \equiv (P \parallel P \parallel P \parallel P \parallel P)$$

- The impact of an action on a counting variable is
 - decrease by 1 if the component participates in the action
 - increase by 1 if the component is the result of the action
 - zero if the component is not involved in the action.

Simple example revisited

$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0[N_P] \boxtimes_{\{task1\}} Res_0[N_R]$$

Simple example revisited

$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0[N_P] \boxtimes_{\{task1\}} Res_0[N_R]$$

CTMC interpretation

| Processors (N_P) | Resources (N_R) | States ($2^{N_P+N_R}$) |
|----------------------|---------------------|--------------------------|
| 1 | 1 | 4 |
| 2 | 1 | 8 |
| 2 | 2 | 16 |
| 3 | 2 | 32 |
| 3 | 3 | 64 |
| 4 | 3 | 128 |
| 4 | 4 | 256 |
| 5 | 4 | 512 |
| 5 | 5 | 1024 |
| 6 | 5 | 2048 |
| 6 | 6 | 4096 |
| 7 | 6 | 8192 |
| 7 | 7 | 16384 |
| 8 | 7 | 32768 |
| 8 | 8 | 65536 |
| 9 | 8 | 131072 |
| 9 | 9 | 262144 |
| 10 | 9 | 524288 |
| 10 | 10 | 1048576 |

Simple example revisited

$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0[N_P] \boxtimes_{\{task1\}} Res_0[N_R]$$

- *task1* decreases $Proc_0$ and Res_0
- *task1* increases $Proc_1$ and Res_1
- *task2* decreases $Proc_1$
- *task2* increases $Proc_0$
- *reset* decreases Res_1
- *reset* increases Res_0

Simple example revisited

$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0[N_P] \boxtimes_{\{task1\}} Res_0[N_R]$$

$$\frac{dx_1}{dt} = -\min(r_1 x_1, r_3 x_3) + r_2 x_2$$

$x_1 = \text{no. of } Proc_1$

- *task1* decreases $Proc_0$
- *task1* is performed by $Proc_0$ and Res_0
- *task2* increases $Proc_0$
- *task2* is performed by $Proc_1$

Simple example revisited

$$Proc_0 \stackrel{def}{=} (task1, r_1).Proc_1$$

$$Proc_1 \stackrel{def}{=} (task2, r_2).Proc_0$$

$$Res_0 \stackrel{def}{=} (task1, r_3).Res_1$$

$$Res_1 \stackrel{def}{=} (reset, r_4).Res_0$$

$$Proc_0[N_P] \boxtimes_{\{task1\}} Res_0[N_R]$$

ODE interpretation

$$\frac{dx_1}{dt} = -\min(r_1 x_1, r_3 x_3) + r_2 x_2$$

$x_1 = \text{no. of } Proc_1$

$$\frac{dx_2}{dt} = \min(r_1 x_1, r_3 x_3) - r_2 x_2$$

$x_2 = \text{no. of } Proc_2$


$$\frac{dx_3}{dt} = -\min(r_1 x_1, r_3 x_3) + r_4 x_4$$

$x_3 = \text{no. of } Res_0$


$$\frac{dx_4}{dt} = \min(r_1 x_1, r_3 x_3) - r_4 x_4$$

$x_4 = \text{no. of } Res_1$


100 processors and 80 resources (simulation run A)




100 processors and 80 resources (simulation run B)




100 processors and 80 resources (simulation run C)




100 processors and 80 resources (simulation run D)




100 processors and 80 resources (average of 10 runs)




100 Processors and 80 resources (average of 100 runs)



100 processors and 80 resources (average of 1000 runs)



100 processors and 80 resources (ODE solution)



Outline

- 1 Introduction
 - Stochastic Process Algebra
- 2 Interpreting SPA for performance modelling
 - Identity and Individuality
 - Collective Dynamics
- 3 Continuous Approximation
 - Numerical illustration
- 4 **Example**
 - **Model**
 - **Model Evaluation**
- 5 Conclusions
 - Alternative interpretations

Case study: A Virtual University



Location, Time, and Size



Case study: e-University Course Selection

The e-University was one of the case studies considered in the recently completed SENSORIA project, an EU project focussed on service-oriented computing.

Case study: e-University Course Selection

The e-University was one of the case studies considered in the recently completed SENSORIA project, an EU project focussed on service-oriented computing.

In the [Course Selection](#) scenario students obtain information about the courses available at their education establishment and may enrol in those for which specific requirements are satisfied.

Case study: e-University Course Selection

The e-University was one of the case studies considered in the recently completed SENSORIA project, an EU project focussed on service-oriented computing.

In the [Course Selection](#) scenario students obtain information about the courses available at their education establishment and may enrol in those for which specific requirements are satisfied.

The model will not consider other services which may be deployed in an actual application (e.g. authentication services) because their impact on performance is assumed to be negligible.

Case study: e-University Course Selection

The e-University was one of the case studies considered in the recently completed SENSORIA project, an EU project focussed on service-oriented computing.

In the [Course Selection](#) scenario students obtain information about the courses available at their education establishment and may enrol in those for which specific requirements are satisfied.

The model will not consider other services which may be deployed in an actual application (e.g. authentication services) because their impact on performance is assumed to be negligible.

We assume a constant population of students e.g. the application is only accessible after the university's matriculation process is complete.

The model: services

Access to the system is through the [University Portal](#). There are four services in this model:

[Course Browsing](#) allows the user to navigate through the University's course offerings;

[Course Selection](#) allows the user to submit a tentative course plan which will be validated against the University's requirements and the student's curriculum;

[Student Confirmation](#) will enforce the student to check relevant personal details;

[Course Registration](#) will confirm the student's selection.

The model: services

Access to the system is through the [University Portal](#). There are four services in this model:

[Course Browsing](#) allows the user to navigate through the University's course offerings;


[Course Selection](#) allows the user to submit a tentative course plan which will be validated against the University's requirements and the student's curriculum;

[Student Confirmation](#) will enforce the student to check relevant personal details;

[Course Registration](#) will confirm the student's selection.

These components make use of a [Database](#) service, which in turn maintains an event log through a [Logger](#) service.

The model: deployment diagram



Solid lines show request/response pairings whilst dashed lines show deployment on processors.

The model: multi-threading

- The [University Portal](#) instantiates a pool of threads.
- Each thread deals with one request from one student for one of the services offered.
- During processing the thread is locked in the sense that it cannot be acquired by further incoming requests.
- When the request is fulfilled the thread clears its state and becomes available again.

The model: multi-threading

- The [University Portal](#) instantiates a pool of threads.
- Each thread deals with one request from one student for one of the services offered.
- During processing the thread is locked in the sense that it cannot be acquired by further incoming requests.
- When the request is fulfilled the thread clears its state and becomes available again.

Analogous multi-threaded behaviour will be given to [Database](#) and [Logger](#).

The model: multi-threading

- The [University Portal](#) instantiates a pool of threads.
- Each thread deals with one request from one student for one of the services offered.
- During processing the thread is locked in the sense that it cannot be acquired by further incoming requests.
- When the request is fulfilled the thread clears its state and becomes available again.

Analogous multi-threaded behaviour will be given to [Database](#) and [Logger](#).

If all threads are busy requests are queued. There will also be contention with many threads running on the same processor.

The model: modelling processors and services

$$\begin{aligned} \text{Processor}_1 &\stackrel{\text{def}}{=} (\text{acq}, r_{\text{acq}}).\text{Processor}_2 \\ \text{Processor}_2 &\stackrel{\text{def}}{=} (\text{type}_1, r_1).\text{Processor}_1 \\ &+ (\text{type}_2, r_2).\text{Processor}_1 \\ &+ \dots \\ &+ (\text{type}_n, r_n).\text{Processor}_1 \end{aligned}$$

The model: modelling processors and services

$$Processor_1 \stackrel{def}{=} (acq, r_{acq}).Processor_2$$

$$Processor_2 \stackrel{def}{=} (type_1, r_1).Processor_1$$

$$+ (type_2, r_2).Processor_1$$

$$+ \dots$$

$$+ (type_n, r_n).Processor_1$$

$$A \stackrel{def}{=} (req_{A,B}, r_{reqA}).(\text{reply}_{A,B}, r_{repA}).A'$$

$$B \stackrel{def}{=} (req_{A,B}, r_{reqB}).(\text{execute}, r).(\text{reply}_{A,B}, r_{repB}).B'$$

The model: modelling processors and services

$$\begin{aligned}
 \text{Processor}_1 &\stackrel{\text{def}}{=} (\text{acq}, r_{\text{acq}}).\text{Processor}_2 \\
 \text{Processor}_2 &\stackrel{\text{def}}{=} (\text{type}_1, r_1).\text{Processor}_1 \\
 &\quad + (\text{type}_2, r_2).\text{Processor}_1 \\
 &\quad + \dots \\
 &\quad + (\text{type}_n, r_n).\text{Processor}_1
 \end{aligned}$$

$$A \stackrel{\text{def}}{=} (\text{req}_{A,B}, r_{\text{req}A}).(\text{reply}_{A,B}, r_{\text{rep}A}).A'$$

$$B \stackrel{\text{def}}{=} (\text{req}_{A,B}, r_{\text{req}B}).(\text{execute}, r).(\text{reply}_{A,B}, r_{\text{rep}B}).B'$$

The communication between A and B will be expressed by means of the cooperation operator $A \bowtie_L B$, $L = \{\text{req}_{A,B}, \text{reply}_{A,B}\}$.

The model: University Portal

A single thread of execution for the application layer **University Portal** is implemented as a sequential component which initially accepts requests for any of the services provided:

$$\begin{aligned} \textit{Portal} &\stackrel{\textit{def}}{=} (\textit{req}_{\textit{student},\textit{browse}}, \nu). \textit{Browse} \\ &+ (\textit{req}_{\textit{student},\textit{select}}, \nu). \textit{Select} \\ &+ (\textit{req}_{\textit{student},\textit{confirm}}, \nu). \textit{Confirm} \\ &+ (\textit{req}_{\textit{student},\textit{register}}, \nu). \textit{Register} \end{aligned}$$

The rate ν is used throughout this model in all the request/reply activities.

The model: Course Browsing (example service)

The action type acq_{ps} is used to obtain exclusive access to processor PS .

$$Browse \stackrel{def}{=} (acq_{ps}, \nu). Cache$$

$$Cache \stackrel{def}{=} (cache, 0.95r_{cache}). Internal \\ + (cache, 0.05r_{cache}). External$$

$$Internal \stackrel{def}{=} (acq_{ps}, \nu). (internal, r_{int}). BrowseRep$$

$$External \stackrel{def}{=} (req_{external, read}, \nu). (reply_{external, read}, \nu). \\ (acq_{ps}, \nu). (external, r_{ext}). BrowseRep$$

$$BrowseRep \stackrel{def}{=} (reply_{student, browse}, \nu). Portal$$

The model: Course Browsing (example service)

The action type acq_{ps} is used to obtain exclusive access to processor PS .

$$Browse \stackrel{def}{=} (acq_{ps}, \nu).Cache$$

$$Cache \stackrel{def}{=} (cache, 0.95r_{cache}).Internal \\ + (cache, 0.05r_{cache}).External$$

$$Internal \stackrel{def}{=} (acq_{ps}, \nu).(internal, r_{int}).BrowseRep$$

$$External \stackrel{def}{=} (req_{external, read}, \nu).(reply_{external, read}, \nu). \\ (acq_{ps}, \nu).(external, r_{ext}).BrowseRep$$

$$BrowseRep \stackrel{def}{=} (reply_{student, browse}, \nu).Portal$$

The other services, **Course Selection**, **Student Confirmation** and **Course Registration** are modelled analogously.

The model: the database

The **Database** offers two functions: **reading** and **writing**.

The model: the database

The **Database** offers two functions: **reading** and **writing**.

$$\begin{aligned}
 \text{Database} &\stackrel{\text{def}}{=} (\text{req}_{\text{external}, \text{read}}, \nu). \text{Read} \\
 &\quad + (\text{req}_{\text{register}, \text{write}}, \nu). \text{Write} \\
 \text{Read} &\stackrel{\text{def}}{=} (\text{acq}_{pd}, \nu). (\text{read}, r_{\text{read}}). \text{ReadReply} \\
 \text{ReadReply} &\stackrel{\text{def}}{=} (\text{reply}_{\text{external}, \text{read}}, \nu). \text{Database} \\
 \text{Write} &\stackrel{\text{def}}{=} (\text{acq}_{pd}, \nu). (\text{write}, r_{\text{write}}). \text{LogWrite} \\
 \text{LogWrite} &\stackrel{\text{def}}{=} (\text{req}_{\text{database}, \text{log}}, \nu). \\
 &\quad (\text{reply}_{\text{database}, \text{log}}, \nu). \text{WriteReply} \\
 \text{WriteReply} &\stackrel{\text{def}}{=} (\text{reply}_{\text{register}, \text{write}}, \nu). \text{Database}
 \end{aligned}$$

The model: the database

The **Database** offers two functions: **reading** and **writing**.

$$\begin{aligned}
 \text{Database} &\stackrel{\text{def}}{=} (\text{req}_{\text{external}, \text{read}}, \nu). \text{Read} \\
 &\quad + (\text{req}_{\text{register}, \text{write}}, \nu). \text{Write} \\
 \text{Read} &\stackrel{\text{def}}{=} (\text{acq}_{pd}, \nu). (\text{read}, r_{\text{read}}). \text{ReadReply} \\
 \text{ReadReply} &\stackrel{\text{def}}{=} (\text{reply}_{\text{external}, \text{read}}, \nu). \text{Database} \\
 \text{Write} &\stackrel{\text{def}}{=} (\text{acq}_{pd}, \nu). (\text{write}, r_{\text{write}}). \text{LogWrite} \\
 \text{LogWrite} &\stackrel{\text{def}}{=} (\text{req}_{\text{database}, \text{log}}, \nu). \\
 &\quad (\text{reply}_{\text{database}, \text{log}}, \nu). \text{WriteReply} \\
 \text{WriteReply} &\stackrel{\text{def}}{=} (\text{reply}_{\text{register}, \text{write}}, \nu). \text{Database}
 \end{aligned}$$

Reading is a purely local computation.

Writing additionally uses the **Logger** service.

The model: the Logger

The **Logger** service accepts requests from **Student Confirmation** and **Database** respectively. It is deployed on the same processor as **Database**, (processor PD) so one thread execution may be modelled as:

$$\begin{aligned}
 \text{Logger} &\stackrel{\text{def}}{=} (\text{req}_{\text{confirm}, \text{log}}, \nu). \text{LogConfirm} \\
 &\quad + (\text{req}_{\text{database}, \text{log}}, \nu). \text{LogDatabase} \\
 \text{LogConfirm} &\stackrel{\text{def}}{=} (\text{acq}_{pd}, \nu). \\
 &\quad (\text{log}_{\text{conf}}, r_{lgc}). \text{ReplyConfirm} \\
 \text{ReplyConfirm} &\stackrel{\text{def}}{=} (\text{reply}_{\text{confirm}, \text{log}}, \nu). \text{Logger} \\
 \text{LogDatabase} &\stackrel{\text{def}}{=} (\text{acq}_{pd}, \nu). \\
 &\quad (\text{log}_{db}, r_{lgd}). \text{ReplyDatabase} \\
 \text{ReplyDatabase} &\stackrel{\text{def}}{=} (\text{reply}_{\text{database}, \text{log}}, \nu). \text{Logger}
 \end{aligned}$$

The model: The processor PD

$$\begin{aligned} PD_1 &\stackrel{def}{=} (acq_{pd}, \nu).PD_2 \\ PD_2 &\stackrel{def}{=} (read, r_{read}).PD_1 + (write, r_{write}).PD_1 \\ &\quad + (log_{conf}, r_{lgc}).PD_1 + (log_{db}, r_{lgd}).PD_1 \end{aligned}$$

The model: Student Workload

A student is modelled as a sequential component which interacts with the university portal and accesses all of the services available. The behaviour is cyclic and the student interposes some think time between successive requests.

The model: Student Workload

A student is modelled as a sequential component which interacts with the university portal and accesses all of the services available. The behaviour is cyclic and the student interposes some think time between successive requests.

$$StdThink \stackrel{def}{=} (think, r_{think}).StdBrowse$$

$$StdBrowse \stackrel{def}{=} (req_{student,browse}, \nu).(reply_{student,browse}, \nu).StdSelect$$

$$StdSelect \stackrel{def}{=} (req_{student,select}, \nu).(reply_{student,select}, \nu).StdConfirm$$

$$StdConfirm \stackrel{def}{=} (req_{student,confirm}, \nu).(reply_{student,confirm}, \nu).StdRegister$$

$$StdRegister \stackrel{def}{=} (req_{student,register}, \nu).(reply_{student,register}, \nu).StdThink$$

The model: System Equation

The system equation captures the **multiplicity** of threads and processors.

$$StdThink[N_S] \bowtie_* \left((Portal[N_P] \bowtie_{M_1} ValUni[N_P] \bowtie_{M_1} ValCur[N_P]) \right. \\ \left. \bowtie_{M_2} Database[N_D] \bowtie_{M_3} Logger[N_L] \right) \bowtie_* (PS_1[N_{PS}] \bowtie_{\emptyset} PD_1[N_{PD}])$$

where

$$M_1 = \{fork, join\}$$

$$M_2 = \{req_{external,read}, reply_{external,read}, req_{register,write}, reply_{register,write}\}$$

$$M_3 = \{req_{confirm,log}, reply_{confirm,log}, req_{database,log}, reply_{database,log}\}$$

The separate validating threads **ValUni** and **ValCur** inherit the multiplicity levels of the thread **Portal** which spawns them.

Qualitative Analysis

The state space size for this model grows very rapidly — for example, adding one portal thread can result in an increase in state space size by a factor of ten.

| N_S | N_P | N_D | N_L | N_{PS} | N_{PD} | <i>Size</i> |
|-------|-------|-------|-------|----------|----------|-------------|
| 1 | any | any | any | 1 | any | 48 |
| 1 | any | any | any | ≥ 2 | any | 49 |
| 2 | 1 | 1 | 1 | 1 | 1 | 230 |
| 3 | 1 | 1 | 1 | 1 | 1 | 680 |
| 3 | 2 | 2 | 2 | 2 | 2 | 5540 |
| 10 | 2 | 2 | 2 | 2 | 2 | 512116 |
| 10 | 3 | 2 | 2 | 2 | 2 | 5075026 |

Qualitative Analysis

The state space size for this model grows very rapidly — for example, adding one portal thread can result in an increase in state space size by a factor of ten.

Analysis based on the [explicit representation](#) of the state space is nevertheless a valuable tool for [validating the correctness](#) of the model.

Qualitative Analysis

The state space size for this model grows very rapidly — for example, adding one portal thread can result in an increase in state space size by a factor of ten.

Analysis based on the [explicit representation](#) of the state space is nevertheless a valuable tool for [validating the correctness](#) of the model.

When there is only one student the state space is fairly small regardless of the multiplicity levels of threads and processors because at most only one of them will be used.

Qualitative Analysis

The state space size for this model grows very rapidly — for example, adding one portal thread can result in an increase in state space size by a factor of ten.

Analysis based on the [explicit representation](#) of the state space is nevertheless a valuable tool for [validating the correctness](#) of the model.

When there is only one student the state space is fairly small regardless of the multiplicity levels of threads and processors because at most only one of them will be used.

Another form of qualitative analysis can be based on visual inspection of the reachability graph, which can be walked through to generate possible trajectories of the system.

Model Verification

We may verify that the model respects policies of **exclusive access** to threads and processors by direct inspection of the state space.

| <i>Student</i> | <i>Portal</i> | <i>Portal</i> | <i>PS</i> |
|---|-----------------------------------|-----------------------------------|-----------|
| Action type cache | | | |
| $(reply_{student}, browse, \nu).StdSelect$ | Cache | Portal | PS_2 |
| $(reply_{student}, browse, \nu).StdSelect$ | Portal | Cache | PS_2 |
| Action type prepare | | | |
| $(reply_{student}, select, \nu).StdConfirm$ | Portal | $(prepare, r_{prep}).ForkPrepare$ | PS_2 |
| $(reply_{student}, select, \nu).StdConfirm$ | $(prepare, r_{prep}).ForkPrepare$ | Portal | PS_2 |
| Action type confirm | | | |
| $(reply_{student}, confirm, \nu).StdRegister$ | Portal | $(confirm, r_{con}).LogStudent$ | PS_2 |
| $(reply_{student}, confirm, \nu).StdRegister$ | $(confirm, r_{con}).LogStudent$ | Portal | PS_2 |
| Action type register | | | |
| $(reply_{student}, register, \nu).StdThink$ | $(register, r_{reg}).Store$ | Portal | PS_2 |
| $(reply_{student}, register, \nu).StdThink$ | Portal | $(register, r_{reg}).Store$ | PS_2 |

This fragment of the state space when there are two threads for the portal and one instance of all other components.

Model Verification

We may verify that the model respects policies of **exclusive access** to threads and processors by direct inspection of the state space.

| <i>Student</i> | <i>Portal</i> | <i>Portal</i> | <i>PS</i> |
|---|-----------------------------------|-----------------------------------|-----------|
| Action type cache | | | |
| $(reply_{student}, browse, \nu).StdSelect$ | Cache | Portal | PS_2 |
| $(reply_{student}, browse, \nu).StdSelect$ | Portal | Cache | PS_2 |
| Action type prepare | | | |
| $(reply_{student}, select, \nu).StdConfirm$ | Portal | $(prepare, r_{prep}).ForkPrepare$ | PS_2 |
| $(reply_{student}, select, \nu).StdConfirm$ | $(prepare, r_{prep}).ForkPrepare$ | Portal | PS_2 |
| Action type confirm | | | |
| $(reply_{student}, confirm, \nu).StdRegister$ | Portal | $(confirm, r_{con}).LogStudent$ | PS_2 |
| $(reply_{student}, confirm, \nu).StdRegister$ | $(confirm, r_{con}).LogStudent$ | Portal | PS_2 |
| Action type register | | | |
| $(reply_{student}, register, \nu).StdThink$ | $(register, r_{reg}).Store$ | Portal | PS_2 |
| $(reply_{student}, register, \nu).StdThink$ | Portal | $(register, r_{reg}).Store$ | PS_2 |

This fragment of the state space when there are two threads for the portal and one instance of all other components.

A necessary condition for **correctness** is that if one thread is engaged in an activity the other is idle.

Model Verification

In general we do not wish to conduct such verification by hand but it may be readily undertaken using automatic tools such as the probabilistic model checker, [PRISM](#).

Model Verification

In general we do not wish to conduct such verification by hand but it may be readily undertaken using automatic tools such as the probabilistic model checker, [PRISM](#).

Indeed PRISM supports PEPA as one of its input languages.

Markovian Analysis: Performance Bounds

Since exact Markovian analysis relies on an explicit enumeration of the state space it is limited to relatively small-scale systems.

Markovian Analysis: Performance Bounds

Since exact Markovian analysis relies on an explicit enumeration of the state space it is limited to relatively small-scale systems.

But performance analysis of even these small scale systems can offer valuable insight into the behaviour of the system — for the e-University system it may be used to derive performance bound estimates.

Markovian Analysis: Performance Bounds

Since exact Markovian analysis relies on an explicit enumeration of the state space it is limited to relatively small-scale systems.

But performance analysis of even these small scale systems can offer valuable insight into the behaviour of the system — for the e-University system it may be used to derive performance bound estimates.

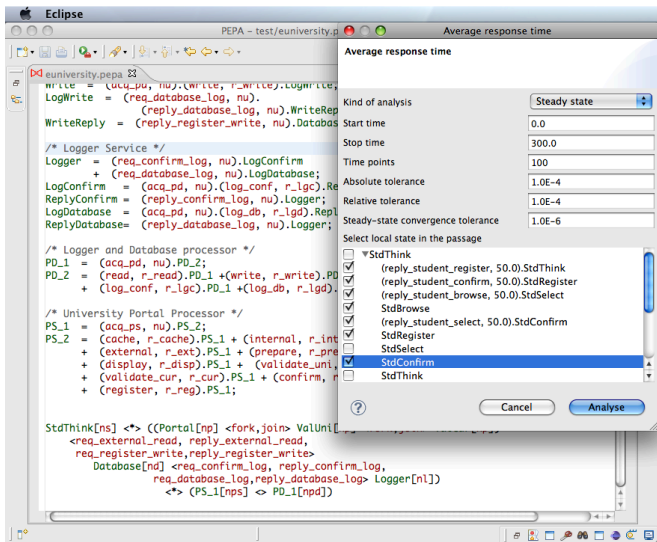
If we consider the system with only one student, the performance in terms of response time perceived by the user will be **optimal** because there is **no contention for resources** — the response time for larger populations of students can only get worse.

Average response times

With only one student the state space is kept manageable and bounds can be computed quickly and accurately.

| <i>System Configuration</i> | | | | | | | | |
|-----------------------------|-------|-------|----------|----------|---------|---------|---------|---------|
| N_P | N_D | N_L | N_{PD} | N_{PS} | $N_S=1$ | $N_S=2$ | $N_S=3$ | $N_S=4$ |
| 1 | 1 | 1 | 1 | 1 | 3.195 | 3.694 | 4.390 | 5.357 |
| 1 | 1 | 1 | 1 | 2 | 3.064 | 3.522 | 4.155 | 5.032 |
| 3 | 3 | 3 | 3 | 3 | 3.064 | 3.065 | 3.066 | 3.074 |

Calculating average response time



Scalability and Optimisation

To study the system under **realistically sized user workloads** many more instances of the components need to be considered.

Scalability and Optimisation

To study the system under **realistically sized user workloads** many more instances of the components need to be considered.

This gives rise to a model whose state space would far exceed the size that can be handled via explicit state space techniques such as Markovian analysis.

Scalability and Optimisation

To study the system under **realistically sized user workloads** many more instances of the components need to be considered.

This gives rise to a model whose state space would far exceed the size that can be handled via explicit state space techniques such as Markovian analysis.

Fortunately, as multiple instances of components are involved we can consider this as a system exhibiting **collective behaviour** and resort to fluid approximation.

Scalability and Optimisation

To study the system under **realistically sized user workloads** many more instances of the components need to be considered.

This gives rise to a model whose state space would far exceed the size that can be handled via explicit state space techniques such as Markovian analysis.

Fortunately, as multiple instances of components are involved we can consider this as a system exhibiting **collective behaviour** and resort to fluid approximation.

The model gives rise to a set of **63 coupled ODEs** which can be solved, for example using a fifth-order Runge-Kutta numerical integrator.

Scalability and Optimisation

To study the system under **realistically sized user workloads** many more instances of the components need to be considered.

This gives rise to a model whose state space would far exceed the size that can be handled via explicit state space techniques such as Markovian analysis.

Fortunately, as multiple instances of components are involved we can consider this as a system exhibiting **collective behaviour** and resort to fluid approximation.

The model gives rise to a set of **63 coupled ODEs** which can be solved, for example using a fifth-order Runge-Kutta numerical integrator.

Of course an alternative would be to use **stochastic simulation**.

Scalability analysis

$$N_P = N_D = N_L = 80, N_{PD} = 40, N_{PS} = 40.$$

| N_S | <i>ODE Runtime</i> | | <i>CTMC Runtime</i> | | <i>Error</i> |
|-------|--------------------|-------|---------------------|--------|--------------|
| 1 | 2.975 | 2.7 s | 3.091 | 436 s | 3.74% |
| 300 | 2.975 | 2.7 s | 3.105 | 2656 s | 4.17% |
| 325 | 2.975 | 2.6 s | 3.329 | 3017 s | 10.62% |
| 350 | 3.686 | 2.9 s | 3.863 | 6505 s | 4.57% |
| 400 | 5.999 | 7.3 s | 5.993 | 4465 s | 0.10% |
| 500 | 10.623 | 6.6 s | 10.534 | 3845 s | 0.84% |
| 600 | 15.248 | 6.6 s | 15.233 | 2985 s | 0.10% |

Stochastic simulation was conducted using the method of batch means, terminated when the 95% confidence interval was $\leq 1\%$ of the average.

Optimisation

In addition to the workload parameter N_S there are 20 other parameters in the model: 15 rate parameters and 5 concurrency levels for threads and processors.

Optimisation

In addition to the workload parameter N_S there are 20 other parameters in the model: 15 rate parameters and 5 concurrency levels for threads and processors.

A typical scenario may be that the rate parameters are hardware-dependent and cannot be revised, and the workload is fixed, determined by the student population.

Optimisation

In addition to the workload parameter N_S there are 20 other parameters in the model: 15 rate parameters and 5 concurrency levels for threads and processors.

A typical scenario may be that the rate parameters are hardware-dependent and cannot be revised, and the workload is fixed, determined by the student population.

The modeller's role is then to determine an optimal configuration for the concurrency levels of the system with respect to user-perceived performance, e.g. average response time.

Optimisation

In addition to the workload parameter N_S there are 20 other parameters in the model: 15 rate parameters and 5 concurrency levels for threads and processors.

A typical scenario may be that the rate parameters are hardware-dependent and cannot be revised, and the workload is fixed, determined by the student population.

The modeller's role is then to determine an optimal configuration for the concurrency levels of the system with respect to user-perceived performance, e.g. average response time.

Let us assume a workload population of 350 students and assume that the response time previously found for this workload is acceptable.

Optimisation

The table shows the response times calculated with different system configurations of similar size.

| <i>Conf.</i> | N_P | N_D | N_L | N_{PS} | N_{PD} | <i>Response time</i> |
|--------------|-------|-------|-------|----------|----------|----------------------|
| <i>A</i> | 80 | 80 | 80 | 40 | 40 | 3.686 |
| <i>B</i> | 70 | 70 | 70 | 40 | 40 | 3.686 |
| <i>C</i> | 60 | 60 | 60 | 40 | 40 | 4.506 |
| <i>D</i> | 70 | 70 | 70 | 35 | 35 | 5.998 |
| <i>E</i> | 70 | 50 | 50 | 40 | 40 | 3.686 |
| <i>F</i> | 70 | 20 | 20 | 40 | 40 | 3.686 |
| <i>G</i> | 70 | 20 | 15 | 40 | 40 | 4.278 |
| <i>H</i> | 70 | 15 | 20 | 40 | 40 | 5.024 |

Optimisation

The table shows the response times calculated with different system configurations of similar size.

| <i>Conf.</i> | N_P | N_D | N_L | N_{PS} | N_{PD} | <i>Response time</i> |
|--------------|-------|-------|-------|----------|----------|----------------------|
| <i>A</i> | 80 | 80 | 80 | 40 | 40 | 3.686 |
| <i>B</i> | 70 | 70 | 70 | 40 | 40 | 3.686 |
| <i>C</i> | 60 | 60 | 60 | 40 | 40 | 4.506 |
| <i>D</i> | 70 | 70 | 70 | 35 | 35 | 5.998 |
| <i>E</i> | 70 | 50 | 50 | 40 | 40 | 3.686 |
| <i>F</i> | 70 | 20 | 20 | 40 | 40 | 3.686 |
| <i>G</i> | 70 | 20 | 15 | 40 | 40 | 4.278 |
| <i>H</i> | 70 | 15 | 20 | 40 | 40 | 5.024 |

Optimisation

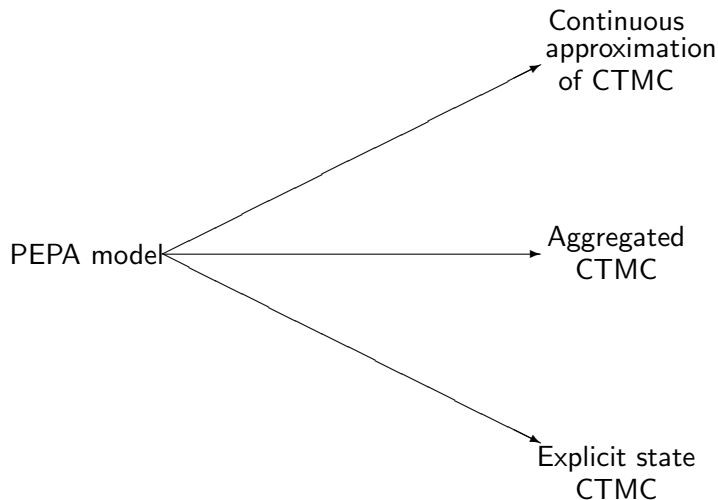
The table shows the response times calculated with different system configurations of similar size.

| <i>Conf.</i> | N_P | N_D | N_L | N_{PS} | N_{PD} | <i>Response time</i> |
|--------------|-------|-------|-------|----------|----------|----------------------|
| <i>A</i> | 80 | 80 | 80 | 40 | 40 | 3.686 |
| <i>B</i> | 70 | 70 | 70 | 40 | 40 | 3.686 |
| <i>C</i> | 60 | 60 | 60 | 40 | 40 | 4.506 |
| <i>D</i> | 70 | 70 | 70 | 35 | 35 | 5.998 |
| <i>E</i> | 70 | 50 | 50 | 40 | 40 | 3.686 |
| <i>F</i> | 70 | 20 | 20 | 40 | 40 | 3.686 |
| <i>G</i> | 70 | 20 | 15 | 40 | 40 | 4.278 |
| <i>H</i> | 70 | 15 | 20 | 40 | 40 | 5.024 |

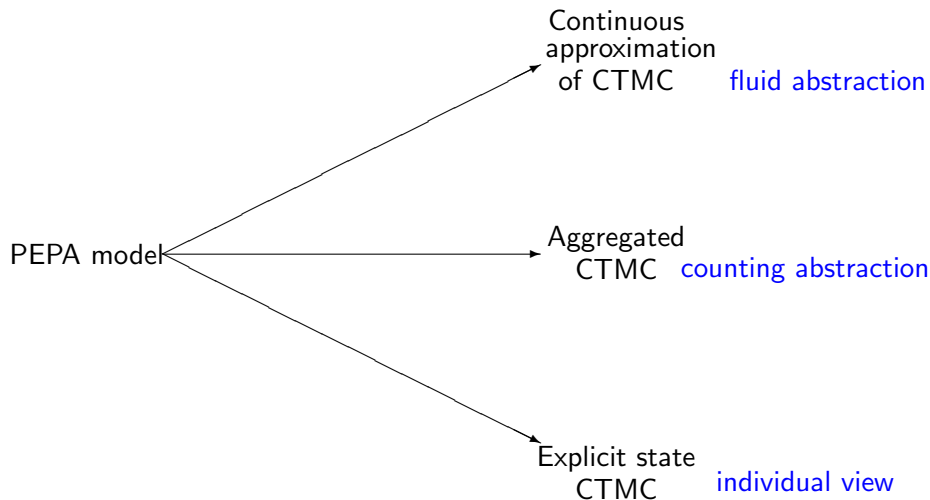
Outline

- 1 Introduction
 - Stochastic Process Algebra
- 2 Interpreting SPA for performance modelling
 - Identity and Individuality
 - Collective Dynamics
- 3 Continuous Approximation
 - Numerical illustration
- 4 Example
 - Model
 - Model Evaluation
- 5 Conclusions
 - Alternative interpretations

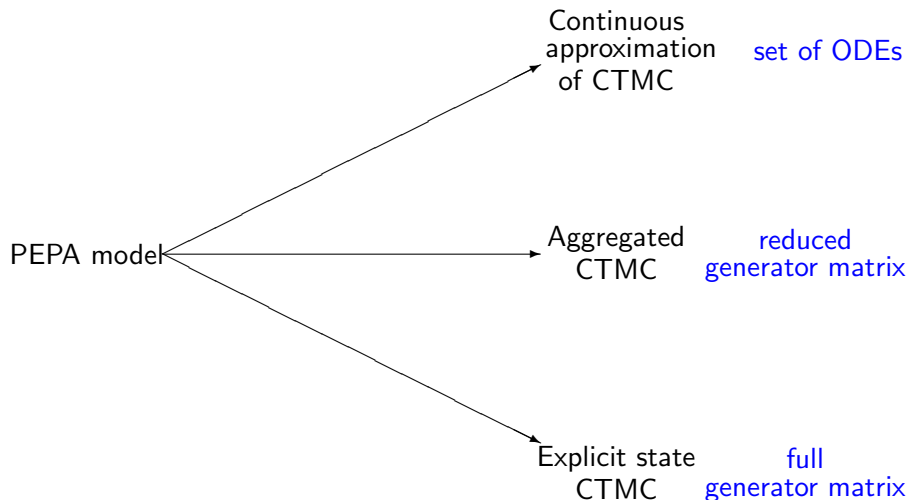
Scalable Representations



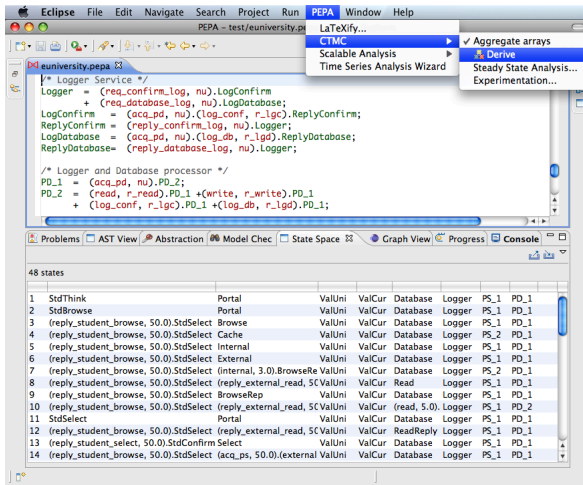
Scalable Representations



Scalable Representations



Eclipse Plug-in for PEPA



Robust tool support is essential to make develop techniques practical.

Other applications

PEPA, and the associated analysis techniques, were originally developed with the objective of studying computer systems.

However, it has also be adopted by modelling a wide-range of other types of system:

- **Locks and movable bridges** in inland shipping in Belgium (Knapen, Hasselt)
- **Automotive on-board diagnostics** expert systems (Console, Picardi and Ribaud)
- **Biological cell signalling pathways** (Calder, Duguid, Gilmore and Hillston)
- **Crowd dynamics** in informatic environments (Harrison, Latella and Massink)

Conclusions

- Many interesting and important systems can be regarded as examples of **collective dynamics** and **emergent behaviour**.

Conclusions

- Many interesting and important systems can be regarded as examples of **collective dynamics** and **emergent behaviour**.
- Process algebras, such as PEPA, are well-suited to modelling the behaviour of such systems in terms of the individuals and their interactions.

Conclusions

- Many interesting and important systems can be regarded as examples of **collective dynamics** and **emergent behaviour**.
- Process algebras, such as PEPA, are well-suited to modelling the behaviour of such systems in terms of the individuals and their interactions.
- **Continuous approximation** allows a rigorous mathematical analysis of the average behaviour of such systems.

Conclusions

- Many interesting and important systems can be regarded as examples of **collective dynamics** and **emergent behaviour**.
- Process algebras, such as PEPA, are well-suited to modelling the behaviour of such systems in terms of the individuals and their interactions.
- **Continuous approximation** allows a rigorous mathematical analysis of the average behaviour of such systems.
- This alternative view of systems has opened up many and exciting new research directions.

Thanks!

Thanks!

Acknowledgements: collaborators

Thanks to many co-authors and collaborators: Jeremy Bradley, Luca Bortolussi, Allan Clark, Graham Clark, Adam Duguid, Vashti Galpin, Nil Gesweiller, **Stephen Gilmore**, Marco Stenico, Nigel Thomas, **Mirco Tribastone**, and others.

Thanks!

Acknowledgements: collaborators

Thanks to many co-authors and collaborators: Jeremy Bradley, Luca Bortolussi, Allan Clark, Graham Clark, Adam Duguid, Vashti Galpin, Nil Gesweiller, **Stephen Gilmore**, Marco Stenico, Nigel Thomas, **Mirco Tribastone**, and others.

Acknowledgements: funding

The PEPA project has received funding from the EPSRC, BBSRC, the Royal Society and the CEC IST-FET programme.

Thanks!

Acknowledgements: collaborators

Thanks to many co-authors and collaborators: Jeremy Bradley, Luca Bortolussi, Allan Clark, Graham Clark, Adam Duguid, Vashti Galpin, Nil Gesweiller, **Stephen Gilmore**, Marco Stenico, Nigel Thomas, **Mirco Tribastone**, and others.

Acknowledgements: funding

The PEPA project has received funding from the EPSRC, BBSRC, the Royal Society and the CEC IST-FET programme.

More information:

<http://www.dcs.ed.ac.uk/pepa>