Making Stochastic Process Algebras Count: Modelling Collective Dynamics

Jane Hillston

Laboratory for Foundations of Computer Science and Centre for Systems Biology at Edinburgh University of Edinburgh

16th November 2010

Outline

1 Introduction

- Stochastic Process Algebra
- Collective Dynamics
- 2 Continuous Approximation
 - State variables
 - Numerical illustration
- 3 Fluid-Flow Semantics
 - Fluid Structured Operational Semantics
- 4 Examples
 - Secure Web Service use
 - Emergency egress

5 Conclusions

Outline

1 Introduction

- Stochastic Process Algebra
- Collective Dynamics
- 2 Continuous Approximation
 - State variables
 - Numerical illustration

3 Fluid-Flow Semantics

Fluid Structured Operational Semantics

4 Examples

- Secure Web Service use
- Emergency egress

5 Conclusions

—Stochastic Process Algebra

Process Algebra





—Stochastic Process Algebra

Process Algebra





—Stochastic Process Algebra

Process Algebra





—Stochastic Process Algebra

Process Algebra





—Stochastic Process Algebra

Process Algebra





—Stochastic Process Algebra

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.

Stochastic Process Algebra

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

-Stochastic Process Algebra

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 SOS rules model

—Stochastic Process Algebra

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 SOS rules Labelled transition system

—Stochastic Process Algebra

Process Algebra



└─Stochastic Process Algebra

A simple example: processors and resources

 $Proc_0 \parallel_{task1} Res_0$

–Stochastic Process Algebra

A simple example: processors and resources



-Stochastic Process Algebra

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 Algebra

Stochastic Process Algebra

Models are constructed from components which engage in activities.



-Stochastic Process Algebra

Stochastic Process Algebra

Models are constructed from components which engage in activities.



-Stochastic Process Algebra

Stochastic 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.

SPA MODEL

-Stochastic Process Algebra

Stochastic 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.

SPA SOS rules

-Stochastic Process Algebra

Stochastic Process Algebra

Models are constructed from components which engage in activities.





-Stochastic Process Algebra

Stochastic Process Algebra

Models are constructed from components which engage in activities.





-Stochastic Process Algebra

Stochastic Process Algebra

Models are constructed from components which engage in activities.





Stochastic Process Algebra

Integrated analysis

Qualitative verification can now be complemented by quantitative verification.

–Stochastic Process Algebra

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?



Stochastic Process Algebra

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?



–Stochastic Process Algebra

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?



└─Stochastic Process Algebra

$(\alpha, f).P$	Prefix
$P_1 + P_2$	Choice
$P_1 \bowtie P_2$	Co-operation
P/L	Hiding
X	Variable

└─Stochastic Process Algebra

$(\alpha, f).P$	Prefix
$P_1 + P_2$	Choice
$P_1 \bowtie P_2$	Co-operation
P/L	Hiding
X	Variable

└─Stochastic Process Algebra

$(\alpha, f).P$	Prefix
$P_1 + P_2$	Choice
$P_1 \bowtie P_2$	Co-operation
P/L	Hiding
X	Variable

└─Stochastic Process Algebra

$(\alpha, f).P$	Prefix
$P_1 + P_2$	Choice
$P_1 \bowtie P_2$	Co-operation
P/L	Hiding
X	Variable

└─Stochastic Process Algebra

$(\alpha, f).P$	Prefix
$P_1 + P_2$	Choice
$P_1 \bowtie P_2$	Co-operation
P/L	Hiding
X	Variable

└─Stochastic Process Algebra

$(\alpha, f).P$	Prefix
$P_1 + P_2$	Choice
$P_1 \bowtie P_2$	Co-operation
P/L	Hiding
X	Variable

–Stochastic Process Algebra

Performance Evaluation Process Algebra (PEPA)

$(\alpha, f).P$	Prefix
$P_1 + P_2$	Choice
$P_1 \bowtie P_2$	Co-operation
P/L	Hiding
Χ	Variable

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

Stochastic Process Algebra

Performance Evaluation Process Algebra (PEPA)

$$\begin{array}{ll} (\alpha, f).P & \operatorname{Prefix} \\ P_1 + P_2 & \operatorname{Choice} \\ P_1 \Join P_2 & \operatorname{Co-operation} \\ P/L & \operatorname{Hiding} \\ X & \operatorname{Variable} \end{array}$$

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

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

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

–Stochastic Process Algebra

PEPA: some recent case studies

 QoS protocols for mobile devices (Wang, Lauriston and Hillston, Edinburgh)
–Stochastic Process Algebra

PEPA: some recent case studies

- QoS protocols for mobile devices (Wang, Lauriston and Hillston, Edinburgh)
- Switch behaviour in active networks (Hillston, Kloul and Mokhtari, Edinburgh and Versailles)

–Stochastic Process Algebra

PEPA: some recent case studies

- QoS protocols for mobile devices (Wang, Lauriston and Hillston, Edinburgh)
- Switch behaviour in active networks (Hillston, Kloul and Mokhtari, Edinburgh and Versailles)
- Task scheduling in a Grid-based processing system (Benoit, Cole, Gilmore and Hillston, Edinburgh)

–Stochastic Process Algebra

PEPA: some recent case studies

- QoS protocols for mobile devices (Wang, Lauriston and Hillston, Edinburgh)
- Switch behaviour in active networks (Hillston, Kloul and Mokhtari, Edinburgh and Versailles)
- Task scheduling in a Grid-based processing system (Benoit, Cole, Gilmore and Hillston, Edinburgh)
- Probability of timely airbag deployment (Clark, Gilmore and Tribastone, Edinburgh)

-Stochastic Process Algebra

Does timeliness matter ...?

There is sometimes a perception in software development that performance does not matter much, or that it is easily fixed later by buying a faster machine.

-Stochastic Process Algebra

Does timeliness matter ...?

There is sometimes a perception in software development that performance does not matter much, or that it is easily fixed later by buying a faster machine.

On the contrary — studies have shown that response time is a key feature in user satisfaction and trust in systems.

-Stochastic Process Algebra

Does timeliness matter ...?

There is sometimes a perception in software development that performance does not matter much, or that it is easily fixed later by buying a faster machine.

On the contrary — studies have shown that response time is a key feature in user satisfaction and trust in systems.

In a recent study by Amazon they artificially delayed page loading times in increments of 100 milliseconds. Even such very small delays were observed to result in substantial and costly drops in revenue.

Gary Linden, Amazon, quoted on http://perspectives.mvdirona.com/2009/10/31/TheCostOfLatency

-Stochastic Process Algebra

Does timeliness matter...?

There is sometimes a perception in software development that performance does not matter much, or that it is easily fixed later by buying a faster machine.

On the contrary — studies have shown that response time is a key feature in user satisfaction and trust in systems.

In a recent study by Amazon they artificially delayed page loading times in increments of 100 milliseconds. Even such very small delays were observed to result in substantial and costly drops in revenue.

Gary Linden, Amazon, quoted on http://perspectives.mvdirona.com/2009/10/31/TheCostOfLatency

AOL, Bing and Google report similar findings.

Collective Dynamic

Collective Dynamics

The behaviour of many systems can be interpreted as the result of the collective behaviour of a large number of interacting entities.



Collective Dynamics

Collective Dynamics

The behaviour of many systems can be interpreted as the result of the collective behaviour of a large number of interacting entities.



For such systems we are often as interested in the population level behaviour as we are in the behaviour of the individual entities.

- Collective Dynamics

Collective Behaviour

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



Collective Dynamics

Collective Behaviour

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



- Collective Dynamics

Collective Behaviour

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



- Collective Dynamics

Collective Behaviour

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



Spread of H1N1 virus in 2009

—Collective Dynamics

Collective Behaviour

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



Love Parade, Germany 2006

Collective Dynamics

Collective Behaviour

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



Map of the Internet 2009

Collective Dynamics

Collective Behaviour

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



Self assessment tax returns 31st January each year

- Collective Dynamics

Process Algebra and Collective Dynamics

- Collective Dynamics

Process Algebra and Collective Dynamics

Process algebra are well-suited to constructing models of such systems:

Developed to represent concurrent behaviour compositionally;

- Collective Dynamics

Process Algebra and Collective Dynamics

- Developed to represent concurrent behaviour compositionally;
- Represent the interactions between individuals explicitly;

- Collective Dynamics

Process Algebra and Collective Dynamics

- Developed to represent concurrent behaviour compositionally;
- Represent the interactions between individuals explicitly;
- Stochastic extensions allow the dynamics of system behaviour to be captured;

- Collective Dynamics

Process Algebra and Collective Dynamics

- Developed to represent concurrent behaviour compositionally;
- Represent the interactions between individuals explicitly;
- Stochastic extensions allow the dynamics of system behaviour to be captured;
- Incorporate formal apparatus for reasoning about the behaviour of systems.

- Collective Dynamics

Process Algebra and Collective Dynamics

Process algebra are well-suited to constructing models of such systems:

- Developed to represent concurrent behaviour compositionally;
- Represent the interactions between individuals explicitly;
- Stochastic extensions allow the dynamics of system behaviour to be captured;
- Incorporate formal apparatus for reasoning about the behaviour of systems.

Unfortunately, whilst theoretically possible, the analysis of such systems through the standard process algebra approaches — explicitly building the state space — is not generally feasible.

Collective Dynamics

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.



Collective Dynamics

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.



— Collective Dynamics

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.



— Collective Dynamics

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.



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

— Collective Dynamics

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.



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

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

Collective Dynamics

Solving discrete state models

Alternatively 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.



Collective Dynamics



As the size of the state space becomes large it becomes infeasible to carry out numerical solution and extremely time-consuming to conduct stochastic simulation.

Collective Dynamics

The CODA project

In the CODA project we have been developing stochastic process algebras and associated theory, tailored to the construction and evaluation of the collective dynamics of large systems of interacting entities.

Collective Dynamics

The CODA project

In the CODA project we have been developing stochastic process algebras and associated theory, tailored to the construction and evaluation of the collective dynamics of large systems of interacting entities.

One approach to this is to keep the discrete state representation in the model and to evaluate it algorithmically rather than analytically, i.e. carry out a discrete event simulation of the model to explore its possible behaviours.

Collective Dynamics

The CODA project

In the CODA project we have been developing stochastic process algebras and associated theory, tailored to the construction and evaluation of the collective dynamics of large systems of interacting entities.

One approach to this is to keep the discrete state representation in the model and to evaluate it algorithmically rather than analytically, i.e. carry out a discrete event simulation of the model to explore its possible behaviours.

Another approach is to make a shift to population statistics.

- Collective Dynamics

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.

- Collective Dynamics

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.

- Collective Dynamics

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.

- Collective Dynamics

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.
Collective Dynamics

Example: 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.

Collective Dynamics

Example: 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.

Collective Dynamics

Example: 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".

Collective Dynamics

Novelty

The novelty in this approach is twofold:

- Collective Dynamics

Novelty

The novelty in this approach is twofold:

 Linking process algebra and continuous mathematical models for dynamic evaluation represents a paradigm shift in how such systems are studied.

- Collective Dynamics

Novelty

The novelty in this approach is twofold:

- Linking process algebra and continuous mathematical models for dynamic evaluation represents a paradigm shift in how such systems are studied.
- The prospect of formally-based quantified evaluation of dynamic behaviour could have significant impact in application domains such as:

- Collective Dynamics

Novelty

The novelty in this approach is twofold:

- Linking process algebra and continuous mathematical models for dynamic evaluation represents a paradigm shift in how such systems are studied.
- The prospect of formally-based quantified evaluation of dynamic behaviour could have significant impact in application domains such as:

Large scale software systems

Issues of scalability are important for user satisfaction and resource efficiency but such issues are difficult to investigate using discrete state models.

- Collective Dynamics

Novelty

The novelty in this approach is twofold:

- Linking process algebra and continuous mathematical models for dynamic evaluation represents a paradigm shift in how such systems are studied.
- The prospect of formally-based quantified evaluation of dynamic behaviour could have significant impact in application domains such as:

Biochemical signalling pathways

Understanding these pathways has the potential to improve the quality of life through enhanced drug treatment and better drug design.

- Collective Dynamics

Novelty

The novelty in this approach is twofold:

- Linking process algebra and continuous mathematical models for dynamic evaluation represents a paradigm shift in how such systems are studied.
- The prospect of formally-based quantified evaluation of dynamic behaviour could have significant impact in application domains such as:

Epidemiological systems

Improved modelling of these systems could lead to improved disease prevention and treatment in nature and better security in computer systems.

- Collective Dynamics

Novelty

The novelty in this approach is twofold:

- Linking process algebra and continuous mathematical models for dynamic evaluation represents a paradigm shift in how such systems are studied.
- The prospect of formally-based quantified evaluation of dynamic behaviour could have significant impact in application domains such as:

Crowd dynamics

Technology enhancement is creating new possibilities for directing crowd movements in buildings and urban spaces, for example for emergency egress, which are not yet well-understood.

Outline

Introduction

Stochastic Process Algebra

Collective Dynamics

2 Continuous Approximation

- State variables
- Numerical illustration

3 Fluid-Flow Semantics

Fluid Structured Operational Semantics

4 Examples

- Secure Web Service use
- Emergency egress

5 Conclusions

Alternative Representations



Alternative Representations



Alternative Representations



- Continuous Approximation
 - -State variables

- Continuous Approximation
 - -State variables

Use continuous state variables to approximate the discrete state space.

0 0 0

- Continuous Approximation
 - -State variables



- Continuous Approximation
 - -State variables

Use continuous state variables to approximate the discrete state space.

0 0 0 0 0

- Continuous Approximation
 - -State variables



- Continuous Approximation
 - -State variables

|--|

- Continuous Approximation
 - -State variables



- Continuous Approximation
 - -State variables

Use continuous state variables to approximate the discrete state space.

0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0

- Continuous Approximation
 - -State variables

Use continuous state variables to approximate the discrete state space.

- Continuous Approximation
 - -State variables

Use continuous state variables to approximate the discrete state space.

- Continuous Approximation
 - -State variables

Use continuous state variables to approximate the discrete state space.

- Continuous Approximation
 - -State variables

Use continuous state variables to approximate the discrete state space.

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

-State variables

New mathematical structures: differential equations

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

.

- Continuous Approximation
 - -State variables

- **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.

- Continuous Approximation
 - -State variables

- **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.

- Continuous Approximation
 - -State variables

- **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.
- Instead the trajectory of the ODEs estimates the expected behaviour of the CTMC.

- Continuous Approximation
 - -State variables

- 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.
- Instead the trajectory of the ODEs estimates the expected behaviour of the CTMC.

Appropriate for models in which there are large numbers of components of the same type, i.e. models of populations and situations of collective dynamics.

-State variables

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.

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

-State variables

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 of a counting variable is

-State variables

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 of a counting variable is
decrease by 1 if the component participates in the action

-State variables

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 of 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

-State variables

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 of 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.
Continuous Approximation

-State variables

Models suitable for counting abstraction

 In Bio-PEPA components are parameterised with a counting variable and the definition of an action records the impact that an action has on the counting variable.

- Continuous Approximation
 - -Numerical illustration

Simple example revisited

$$\begin{array}{lll} Proc_{0} & \stackrel{def}{=} & (task1, r_{1}).Proc_{1} \\ Proc_{1} & \stackrel{def}{=} & (task2, r_{2}).Proc_{0} \\ Res_{0} & \stackrel{def}{=} & (task1, r_{3}).Res_{1} \\ Res_{1} & \stackrel{def}{=} & (reset, r_{4}).Res_{0} \end{array}$$

 $Proc_0[N_P] \bigotimes_{_{\{task1\}}} Res_0[N_R]$

Continuous Approximation

Simple example revisited

$$\begin{array}{rcl} Proc_{0} & \stackrel{def}{=} & (task1, r_{1}).Proc_{1} \\ Proc_{1} & \stackrel{def}{=} & (task2, r_{2}).Proc_{0} \\ Res_{0} & \stackrel{def}{=} & (task1, r_{3}).Res_{1} \\ Res_{1} & \stackrel{def}{=} & (reset, r_{4}).Res_{0} \end{array}$$

 $Proc_0[N_P] \bigotimes_{\text{{task1}}} Res_0[N_R]$

СТІ	MC interpret	tation
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

- Continuous Approximation
 - -Numerical illustration

Simple example revisited

$$\begin{array}{lll} Proc_{0} & \stackrel{def}{=} & (task1, r_{1}).Proc_{1} \\ Proc_{1} & \stackrel{def}{=} & (task2, r_{2}).Proc_{0} \\ \end{array}$$

 $Res_0 \stackrel{\text{def}}{=} (task1, r_3).Res_1$ $Res_1 \stackrel{\text{def}}{=} (reset, r_4).Res_0$

 $Proc_0[N_P] \underset{\{task1\}}{\bowtie} Res_0[N_R]$

- *task*1 decreases *Proc*₀ and *Res*₀
- *task*1 increases *Proc*₁ and *Res*₁
- task2 decreases Proc1
- task2 increases Proc0
- reset decreases Res1
- reset increases Res₀

- Continuous Approximation
 - Numerical illustration

Simple example revisited

$$\begin{array}{lll} Proc_{0} & \stackrel{def}{=} & (task1, r_{1}).Proc_{1} \\ Proc_{1} & \stackrel{def}{=} & (task2, r_{2}).Proc_{0} \\ \end{array}$$

- $Res_0 \stackrel{\text{def}}{=} (task1, r_3).Res_1$
- $Res_1 \stackrel{def}{=} (reset, r_4).Res_0$

 $Proc_0[N_P] \underset{{task1}}{\boxtimes} Res_0[N_R]$

$$\frac{dx_1}{dt} = -\min(r_1 x_1, r_3 x_3) + r_2 x_2 x_1 = \text{no. of } Proc_1$$

- *task*1 decreases *Proc*₀
- task1 is performed by Proc₀ and Res₀
- task2 increases Proc₀
- *task*2 is performed by *Proc*₁

Continuous Approximation

Numerical illustration

Simple example revisited

$$Proc_0[N_P] \bigotimes_{\text{{task1}}} Res_0[N_R]$$

ODE interpretation $\frac{\mathrm{d}x_1}{\mathrm{d}t} = -\min(r_1 \, x_1, r_3 \, x_3) + r_2 \, x_2$ $x_1 = \text{no. of } Proc_1$ $\frac{\mathrm{d}x_2}{\mathrm{d}t} = \min(r_1 x_1, r_3 x_3) - r_2 x_2$ $x_2 = no.$ of *Proc*₂ $\frac{\mathrm{d}x_3}{\mathrm{d}t} = -\min(r_1\,x_1,r_3\,x_3) + r_4\,x_4$ $x_3 = \text{no. of } Res_0$ $\frac{dx_4}{dt} = \min(r_1 x_1, r_3 x_3) - r_4 x_4$ $x_4 = no.$ of Res₁

- Continuous Approximation
 - Numerical illustration

100 processors and 80 resources (simulation run A)



- Continuous Approximation
 - Numerical illustration

100 processors and 80 resources (simulation run B)



- Continuous Approximation
 - Numerical illustration

100 processors and 80 resources (simulation run C)



- Continuous Approximation
 - Numerical illustration

100 processors and 80 resources (simulation run D)



Continuous Approximation

Numerical illustration

100 processors and 80 resources (average of 10 runs)



Continuous Approximation

Numerical illustration

100 Processors and 80 resources (average of 100 runs)



- Continuous Approximation
 - Numerical illustration

100 processors and 80 resources (average of 1000 runs)



- Continuous Approximation
 - Numerical illustration

100 processors and 80 resources (ODE solution)



Outline

1 Introduction

- Stochastic Process Algebra
- Collective Dynamics
- 2 Continuous Approximation
 - State variables
 - Numerical illustration

3 Fluid-Flow Semantics

- Fluid Structured Operational Semantics
- 4 Examples
 - Secure Web Service use
 - Emergency egress

5 Conclusions

Fluid Structured Operational Semantics

Deriving a Fluid Approximation of a PEPA model

The aim is to represent the CTMC implicitly (avoiding state space explosion), and to generate the set of ODEs which are the fluid limit of that CTMC.

Fluid Structured Operational Semantics

Deriving a Fluid Approximation of a PEPA model

The aim is to represent the CTMC implicitly (avoiding state space explosion), and to generate the set of ODEs which are the fluid limit of that CTMC.

The exisiting (CTMC) SOS semantics is not suitable for this purpose because it constructs the state space of the CTMC explicitly.

Fluid Structured Operational Semantics

Deriving a Fluid Approximation of a PEPA model

The aim is to represent the CTMC implicitly (avoiding state space explosion), and to generate the set of ODEs which are the fluid limit of that CTMC.

The exisiting (CTMC) SOS semantics is not suitable for this purpose because it constructs the state space of the CTMC explicitly.



-Fluid Structured Operational Semantics

Deriving a Fluid Approximation of a PEPA model

The aim is to represent the CTMC implicitly (avoiding state space explosion), and to generate the set of ODEs which are the fluid limit of that CTMC.

The exisiting (CTMC) SOS semantics is not suitable for this purpose because it constructs the state space of the CTMC explicitly.

Nevertheless we are able to define a structured operational semantics which defines the possible transitions of an abitrary abstract state and from this derive the ODEs.

-Fluid Structured Operational Semantics

Deriving a Fluid Approximation of a PEPA model

The aim is to represent the CTMC implicitly (avoiding state space explosion), and to generate the set of ODEs which are the fluid limit of that CTMC.

The exisiting (CTMC) SOS semantics is not suitable for this purpose because it constructs the state space of the CTMC explicitly.

Nevertheless we are able to define a structured operational semantics which defines the possible transitions of an abitrary abstract state and from this derive the ODEs.



-Fluid Structured Operational Semantics

Fluid Structured Operational Semantics

In order to get to the implicit representation of the CTMC we need to:

-Fluid Structured Operational Semantics

Fluid Structured Operational Semantics

In order to get to the implicit representation of the CTMC we need to:

1 Make the counting abstraction (*Context Reduction*)

-Fluid Structured Operational Semantics

Fluid Structured Operational Semantics

In order to get to the implicit representation of the CTMC we need to:

- **1** Make the counting abstraction (*Context Reduction*)
- 2 Collect the transitions of the reduced context (Jump Multiset)

-Fluid Structured Operational Semantics

Fluid Structured Operational Semantics

In order to get to the implicit representation of the CTMC we need to:

- **1** Make the counting abstraction (*Context Reduction*)
- 2 Collect the transitions of the reduced context (Jump Multiset)
- **3** Calculate the rate of the transitions in terms of an arbitrary state of the CTMC.

-Fluid Structured Operational Semantics

Fluid Structured Operational Semantics

In order to get to the implicit representation of the CTMC we need to:

- **1** Make the counting abstraction (*Context Reduction*)
- 2 Collect the transitions of the reduced context (Jump Multiset)
- **3** Calculate the rate of the transitions in terms of an arbitrary state of the CTMC.

Once this is done we can extract the vector field $F_{\mathcal{M}}(x)$ from the jump multiset.

Fluid Structured Operational Semantics

Context Reduction

$$\begin{array}{l} Proc_{0} \stackrel{def}{=} (task1, r_{1}).Proc_{1} \\ Proc_{1} \stackrel{def}{=} (task2, r_{2}).Proc_{0} \\ Res_{0} \stackrel{def}{=} (task1, r_{3}).Res_{1} \\ Res_{1} \stackrel{def}{=} (reset, r_{4}).Res_{0} \\ System \stackrel{def}{=} Proc_{0}[N_{P}] \underset{\{transfer\}}{\bowtie} Res_{0}[N_{R}] \\ \psi \\ \mathcal{R}(System) = \{Proc_{0}, Proc_{1}\} \underset{\{task1\}}{\bowtie} \{Res_{0}, Res_{1}\} \end{array}$$

-Fluid Structured Operational Semantics

Context Reduction

$$\begin{array}{rcl} Proc_{0} & \stackrel{def}{=} & (task1, r_{1}).Proc_{1} \\ Proc_{1} & \stackrel{def}{=} & (task2, r_{2}).Proc_{0} \\ Res_{0} & \stackrel{def}{=} & (task1, r_{3}).Res_{1} \\ Res_{1} & \stackrel{def}{=} & (reset, r_{4}).Res_{0} \\ System & \stackrel{def}{=} & Proc_{0}[N_{P}] \underset{\{transfer\}}{\boxtimes} Res_{0}[N_{R}] \\ & \downarrow \\ \mathcal{R}(System) = \{Proc_{0}, Proc_{1}\} \underset{\{task1\}}{\boxtimes} \{Res_{0}, Res_{1}\} \end{array}$$

Population Vector

$$\xi = (\xi_1, \xi_2, \xi_3, \xi_4)$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Location Dependency

 $\textit{System} \stackrel{\tiny def}{=} \textit{Proc}_0[N'_C] \underset{\tiny \{\textit{task1}\}}{\bowtie} \textit{Res}_0[N_S] \parallel \textit{Proc}_0[N''_C]$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Location Dependency

 $System \stackrel{\text{def}}{=} Proc_0[N'_C] \underset{\{task1\}}{\bowtie} Res_0[N_S] \parallel Proc_0[N''_C]$ \downarrow $\{Proc_0, Proc_1\} \underset{\{task1\}}{\bowtie} \{Res_0, Res_1\} \parallel \{Proc_0, Proc_1\}$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Location Dependency

$$System \stackrel{det}{=} Proc_0[N'_C] \underset{\{task1\}}{\boxtimes} Res_0[N_S] \parallel Proc_0[N''_C]$$
$$\Downarrow$$
$$\{Proc_0, Proc_1\} \underset{\{task1\}}{\boxtimes} \{Res_0, Res_1\} \parallel \{Proc_0, Proc_1\}$$

Population Vector

$$\xi = (\xi_1, \xi_2, \xi_3, \xi_4, \xi_5, \xi_6)$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Fluid-Flow Semantics

Fluid Structured Operational Semantics

$$\frac{\underset{Proc_{0}}{\xrightarrow{task1,r_{1}}} Proc_{1}}{\underset{Proc_{0}}{\xrightarrow{task1,r_{1}\xi_{1}}} * Proc_{1}}$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

$$\frac{\Pr{oc_0} \xrightarrow{task1, r_1} \Pr{oc_1}}{\Pr{oc_0} \xrightarrow{task1, r_1 \xi_1} * \Pr{oc_1}} \qquad \frac{\operatorname{Res_0} \xrightarrow{task1, r_3} \operatorname{Res_1}}{\operatorname{Res_0} \xrightarrow{task1, r_3 \xi_3} * \operatorname{Res_1}}$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

$$\frac{\frac{Proc_{0} \xrightarrow{task1, r_{1}} Proc_{1}}{Proc_{0} \xrightarrow{task1, r_{1}\xi_{1}} Proc_{1}} \xrightarrow{Res_{0} \xrightarrow{task1, r_{3}} Res_{1}}{Res_{0} \xrightarrow{task1, r_{3}\xi_{3}} Res_{1}}}{\frac{Proc_{0} \bigotimes_{_{\{task1\}}} Res_{0} \xrightarrow{task1, r(\xi)}}{Proc_{1} \bigotimes_{_{\{task1\}}} Res_{1}}}$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Apparent Rate Calculation



Fluid-Flow Semantics

Fluid Structured Operational Semantics

Apparent Rate Calculation

$$\frac{\frac{Proc_{0} \xrightarrow{task1, r_{1}} Proc_{1}}{exsk1, r_{1}\xi_{1}}}{Proc_{0} \xrightarrow{task1, r_{1}\xi_{1}} Proc_{1}} \xrightarrow{Res_{0} \xrightarrow{task1, r_{3}\xi_{3}} Res_{1}}{Res_{0} \xrightarrow{task1, r_{3}\xi_{3}} Res_{1}}$$

$$r(\xi) = \min\left(r_1\xi_1, r_3\xi_4\right)$$
Fluid Structured Operational Semantics

$f(\xi, I, \alpha)$ as the Generator Matrix of the Lumped CTMC

 $(\underset{\bigstar}{\overset{P_1}{\models}} \parallel P_0 \parallel P_0) \bigotimes_{\scriptscriptstyle \{task1\}} {\overset{R_1}{\models}} \parallel R_0)$ $(P_1 \parallel P_0 \parallel P_0) \underset{\text{{\tiny {\{ task1 \}}}}{\boxtimes}}{\boxtimes} (R_0 \parallel R_1)$ $(P_0 \parallel P_1 \parallel P_0) \bigotimes_{\substack{\{\text{task}\}}} (R_1 \parallel R_0)$ $(P_0 \parallel P_0 \parallel P_0) \underset{\text{{\scriptsize {fask1}}}}{\boxtimes} (R_0 \parallel R_0)$ $(P_0 \parallel P_1 \parallel P_0) \bigotimes_{\text{{task1}}} (R_0 \parallel R_1)$ $(P_0 \parallel P_0 \parallel P_1) \bigotimes_{\text{{task1}}} (R_1 \parallel R_0)$ $(P_0 \parallel P_0 \parallel P_1) \underset{\text{{task1}}}{\bowtie} (R_0 \parallel R_1)$

Fluid Structured Operational Semantics

$f(\xi, I, \alpha)$ as the Generator Matrix of the Lumped CTMC



Fluid-Flow Semantics

Fluid Structured Operational Semantics

$f(\xi, I, \alpha)$ as the Generator Matrix of the Lumped CTMC

$$(3,0,2,0) \xrightarrow{\min(3r_{1},2r_{3})} (2,1,1,1)$$

$$(P_{1} \parallel P_{0} \parallel P_{0}) \underset{\{task1\}}{\boxtimes} (R_{1} \parallel R_{0})$$

$$(P_{1} \parallel P_{0} \parallel P_{0}) \underset{\{task1\}}{\boxtimes} (R_{0} \parallel R_{1})$$

$$(P_{0} \parallel P_{0} \parallel P_{0}) \underset{\{task1\}}{\boxtimes} (R_{0} \parallel R_{0})$$

$$(P_{0} \parallel P_{0} \parallel P_{0}) \underset{\{task1\}}{\boxtimes} (R_{0} \parallel R_{0})$$

$$(P_{0} \parallel P_{0} \parallel P_{0}) \underset{\{task1\}}{\boxtimes} (R_{0} \parallel R_{0})$$

$$(P_{0} \parallel P_{0} \parallel P_{0}) \underset{\{task1\}}{\boxtimes} (R_{0} \parallel R_{0})$$

$$(P_{0} \parallel P_{0} \parallel P_{0}) \underset{\{task1\}}{\boxtimes} (R_{0} \parallel R_{0})$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Jump Multiset

$$\frac{Proc_0}{{}_{\{taskI\}}} \underset{\{taskI\}}{\boxtimes} \frac{Res_0}{\underset{\{taskI,r(\xi)\}}{\longrightarrow}} \frac{Proc_1}{Proc_1} \underset{\{taskI\}}{\boxtimes} \frac{Res_1}{Res_1}$$
$$r(\xi) = \min(r_1\xi_1, r_3\xi_3)$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Jump Multiset

$$\frac{Proc_0}{{}_{\{taskI\}}} \underset{\{taskI\}}{\overset{[taskI]}{\longrightarrow}} \underset{Res_0}{\overset{[taskI]}{\longrightarrow}} \underset{\{taskI\}}{\overset{[taskI]}{\longrightarrow}} \underset{Res_1}{\overset{[taskI]}{\longrightarrow}} \underset{r(\xi) = \min(r_1\xi_1, r_3\xi_3)}{\overset{[taskI]}{\longrightarrow}}$$

$$\frac{Proc_{1}}{{}_{\{task1\}}}Res_{0} \xrightarrow{task2,\xi_{2}r_{2}} * \frac{Proc_{0}}{{}_{\{task1\}}}Res_{0}$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Jump Multiset

$$\frac{\operatorname{Proc}_{0}}{\operatorname{task}_{1}} \underset{\operatorname{task}_{1}}{\bowtie} \operatorname{Res}_{0} \xrightarrow{\operatorname{task}_{1}, r(\xi)} \operatorname{Proc}_{1} \underset{\operatorname{task}_{1}}{\bowtie} \operatorname{Res}_{1}}{\operatorname{r}(\xi) = \min\left(r_{1}\xi_{1}, r_{3}\xi_{3}\right)}$$

$$\frac{Proc_{1}}{{}_{\{task1\}}}Res_{0} \xrightarrow{task2,\xi_{2}r_{2}} * \frac{Proc_{0}}{{}_{\{task1\}}}Res_{0}$$

$$Proc_{0} \bigotimes_{\substack{\{taskI\}}} Res_{1} \xrightarrow{reset, \xi_{4}r_{4}} * Proc_{0} \bigotimes_{\substack{\{taskI\}}} Res_{0}$$

Fluid-Flow Semantics

-Fluid Structured Operational Semantics

Equivalent Transitions

Some transitions may give the same information:

$$\begin{array}{c|c} Proc_{0} & \underset{\{task1\}}{\boxtimes} Res_{1} & \xrightarrow{reset, \xi_{4}r_{4}} & Proc_{0} & \underset{\{task1\}}{\boxtimes} Res_{0} \\ Proc_{1} & \underset{\{task1\}}{\boxtimes} Res_{1} & \xrightarrow{reset, \xi_{4}r_{4}} & Proc_{1} & \underset{\{task1\}}{\boxtimes} Res_{0} \end{array}$$

i.e., Res_1 may perform an action independently from the rest of the system.

This is captured by the procedure used for the construction of the generator function $f(\xi, I, \alpha)$

Fluid-Flow Semantics

Construction of $f(\xi, I, \alpha)$

$$Proc_0 \bigotimes_{\{task1\}} Res_1 \xrightarrow{reset, \xi_4 r_4} Proc_0 \bigotimes_{\{task1\}} Res_0$$

Fluid-Flow Semantics

Construction of
$$f(\xi, I, \alpha)$$

$$Proc_0 \bigotimes_{\{task1\}} Res_1 \xrightarrow{reset, \xi_4r_4} Proc_0 \bigotimes_{\{task1\}} Res_0$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Construction of
$$f(\xi, I, \alpha)$$

$$Proc_{0} \underset{\{task1\}}{\boxtimes} Res_{1} \xrightarrow{reset, \xi_{4}r_{4}} * Proc_{0} \underset{\{task1\}}{\boxtimes} Res_{0}$$

- Take *I* = (0, 0, 0, 0)
- Add -1 to all elements of / corresponding to the indices of the components in the lhs of the transition

$$I = (-1, 0, 0, -1)$$

Add +1 to all elements of / corresponding to the indices of the components in the rhs of the transition

$$I = (-1 + 1, 0, +1, -1) = (0, 0, +1, -1)$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Construction of
$$f(\xi, I, \alpha)$$

$$Proc_0 \underset{\{task1\}}{\bowtie} Res_1 \xrightarrow{reset, \xi_4r_4} Proc_0 \underset{\{task1\}}{\bowtie} Res_0$$

- Take *I* = (0, 0, 0, 0)
- Add -1 to all elements of / corresponding to the indices of the components in the lhs of the transition

$$l = (-1, 0, 0, -1)$$

Add +1 to all elements of / corresponding to the indices of the components in the rhs of the transition

$$I = (-1 + 1, 0, +1, -1) = (0, 0, +1, -1)$$

$$f(\xi, (0, 0, +1, -1), reset) = \xi_4 r_4$$

Fluid-Flow Semantics

Construction of $f(\xi, I, \alpha)$

Fluid-Flow Semantics

Construction of
$$f(\xi, I, \alpha)$$

$$Proc_0 \underset{\{task1\}}{\bowtie} Res_0 \xrightarrow{task1, r(\xi)} * Proc_1 \underset{\{task1\}}{\bowtie} Res_1$$

$$f(\xi, (-1, +1, -1, +1), task1) = r(\xi)$$

Fluid-Flow Semantics

Construction of
$$f(\xi, I, \alpha)$$

$$\begin{array}{c|c} \operatorname{Proc}_{0} & \operatornamewithlimits{\Join}_{\{task1\}} \operatorname{Res}_{0} & \xrightarrow{task1, r(\xi)} * & \operatorname{Proc}_{1} & \operatornamewithlimits{\Join}_{\{task1\}} \operatorname{Res}_{1} \\ \operatorname{Proc}_{1} & \operatornamewithlimits{\bowtie}_{\{task1\}} \operatorname{Res}_{0} & \xrightarrow{task2, \xi_{2}r_{2}'} * & \operatorname{Proc}_{0} & \operatornamewithlimits{\bowtie}_{\{task1\}} \operatorname{Res}_{0} \end{array}$$

$$\begin{array}{rcl} f(\xi,(-1,+1,-1,+1),\mathit{task1}) &=& r(\xi) \\ f(\xi,(+1,-1,0,0),\mathit{task2}) &=& \xi_2 r_2 \end{array}$$

Fluid-Flow Semantics

Construction of
$$f(\xi, I, \alpha)$$

$$\begin{array}{c|c} \operatorname{Proc}_{0} & [\bowtie]_{\{task1\}} \operatorname{Res}_{0} & \xrightarrow{task1, r(\xi)} * & \operatorname{Proc}_{1} & [\bowtie]_{\{task1\}} \operatorname{Res}_{1} \\ \operatorname{Proc}_{1} & [\bowtie]_{\{task1\}} \operatorname{Res}_{0} & \xrightarrow{task2, \xi_{2}r'_{2}} * & \operatorname{Proc}_{0} & [\bowtie]_{\{task1\}} \operatorname{Res}_{0} \\ \operatorname{Proc}_{0} & [\bowtie]_{\{task1\}} \operatorname{Res}_{1} & \xrightarrow{\operatorname{reset}, \xi_{4}r_{4}} * & \operatorname{Proc}_{0} & [\bowtie]_{\{task1\}} \operatorname{Res}_{0} \end{array}$$

$$\begin{array}{rcl} f(\xi,(-1,+1,-1,+1),\mathit{task1}) &=& r(\xi) \\ f(\xi,(+1,-1,0,0),\mathit{task2}) &=& \xi_2 r_2 \\ f(\xi,(0,0,+1,-1),\mathit{reset}) &=& \xi_4 r_4 \end{array}$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Capturing behaviour in the Generator Function

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Capturing behaviour in the Generator Function

Numerical Vector Form

$$\xi = (\xi_1, \xi_2, \xi_3, \xi_4) \in \mathbb{N}^4, \quad \xi_1 + \xi_2 = N_P \text{ and } \xi_3 + \xi_4 = N_R$$

-Fluid-Flow Semantics

-Fluid Structured Operational Semantics

Capturing behaviour in the Generator Function

Numerical Vector Form

$$\xi = (\xi_1, \xi_2, \xi_3, \xi_4) \in \mathbb{N}^4, \quad \xi_1 + \xi_2 = N_P \text{ and } \xi_3 + \xi_4 = N_R$$

Generator Function

$$\begin{array}{rcl} f(\xi,(-1,1,-1,1),\mathit{task1}) &=& \min(r_1\xi_1,r_3\xi_3) \\ f(\xi,l,\alpha): & f(\xi,(1,-1,0,0),\mathit{task2}) &=& r_2\xi_2 \\ & f(\xi,(0,0,1,-1),\mathit{reset}) &=& r_4\xi_4 \end{array}$$

Fluid-Flow Semantics

Fluid Structured Operational Semantics

Extraction of the ODE from f

Generator Function

$$\begin{array}{lll} f(\xi,(-1,1,-1,1),\textit{task1}) &=& \min(r_1\xi_1,r_3\xi_3) \\ f(\xi,(1,-1,0,0),\textit{task2}) &=& r_2\xi_2 \\ f(\xi,(0,0,1,-1),\textit{reset}) &=& r_4\xi_4 \end{array}$$

Differential Equation

$$\begin{aligned} \frac{dx}{dt} &= F_{\mathcal{M}}(x) = \sum_{l \in \mathbb{Z}^d} l \sum_{\alpha \in \mathcal{A}} f(x, l, \alpha) \\ &= (-1, 1, -1, 1) \min(r_1 x_1, r_3 x_3) + (1, -1, 0, 0) r_2 x_2 \\ &+ (0, 0, 1, -1) r_4 x_4 \end{aligned}$$

Fluid-Flow Semantics

—Fluid Structured Operational Semantics

Extraction of the ODE from f

Generator Function

$$\begin{array}{lll} f(\xi,(-1,1,-1,1),\textit{task1}) &=& \min(r_1\xi_1,r_3\xi_3) \\ f(\xi,(1,-1,0,0),\textit{task2}) &=& r_2\xi_2 \\ f(\xi,(0,0,1,-1),\textit{reset}) &=& r_4\xi_4 \end{array}$$

Differential Equation

$$\frac{dx_1}{dt} = -\min(r_1x_1, r_3x_3) + r_2x_2$$
$$\frac{dx_2}{dt} = \min(r_1x_1, r_3x_3) - r_2x_2$$
$$\frac{dx_3}{dt} = -\min(r_1x_1, r_3x_3) + r_4x_4$$
$$\frac{dx_4}{dt} = \min(r_1x_1, r_3x_3) - r_4x_4$$

Fluid Structured Operational Semantics

Consistency results

The vector field \(\mathcal{F}(x)\) is Lipschitz continuous i.e. all the rate functions governing transitions in the process algebra satisfy local continuity conditions.

Fluid Structured Operational Semantics

Consistency results

- The vector field \(\mathcal{F}(x)\) is Lipschitz continuous i.e. all the rate functions governing transitions in the process algebra satisfy local continuity conditions.
- The generated ODEs are the fluid limit of the family of CTMCs generated by f(ξ, l, α): this family forms a sequence as the initial populations are scaled by a variable n.

Fluid Structured Operational Semantics

Consistency results

- The vector field \(\mathcal{F}(x)\) is Lipschitz continuous i.e. all the rate functions governing transitions in the process algebra satisfy local continuity conditions.
- The generated ODEs are the fluid limit of the family of CTMCs generated by f(ξ, l, α): this family forms a sequence as the initial populations are scaled by a variable n.
- We can prove this using Kurtz's theorem: Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes, T.G. Kurtz, J. Appl. Prob. (1970).

Fluid Structured Operational Semantics

Consistency results

- The vector field \(\mathcal{F}(x)\) is Lipschitz continuous i.e. all the rate functions governing transitions in the process algebra satisfy local continuity conditions.
- The generated ODEs are the fluid limit of the family of CTMCs generated by f(ξ, l, α): this family forms a sequence as the initial populations are scaled by a variable n.
- We can prove this using Kurtz's theorem: Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes, T.G. Kurtz, J. Appl. Prob. (1970).
- Moreover Lipschitz continuity of the vector field guarantees existence and uniqueness of the solution to the initial value problem.

Outline

I Introduction

- Stochastic Process Algebra
- Collective Dynamics
- 2 Continuous Approximation
 - State variables
 - Numerical illustration

3 Fluid-Flow Semantics

Fluid Structured Operational Semantics

4 Examples

- Secure Web Service use
- Emergency egress

5 Conclusions

-Secure Web Service use

Example: Secure Web Service use



The example which we consider is a Web service which has two types of clients:

-Secure Web Service use

Example: Secure Web Service use



- The example which we consider is a Web service which has two types of clients:
 - first party application clients which access the web service across a secure intranet, and

Secure Web Service use

Example: Secure Web Service use



- The example which we consider is a Web service which has two types of clients:
 - first party application clients which access the web service across a secure intranet, and
 - second party browser clients which access the Web service across the Internet.

Secure Web Service use

Example: Secure Web Service use



- The example which we consider is a Web service which has two types of clients:
 - first party application clients which access the web service across a secure intranet, and
 - second party browser clients which access the Web service across the Internet.
- Second party clients route their service requests via trusted brokers.

-Secure Web Service use

Scalability and replication



 To ensure scalability the Web service is replicated across multiple hosts.

-Secure Web Service us

Scalability and replication



- To ensure scalability the Web service is replicated across multiple hosts.
- Multiple brokers are available.

└─ Secure Web Service use

Scalability and replication



- To ensure scalability the Web service is replicated across multiple hosts.
- Multiple brokers are available.
- There are numerous first party clients behind the firewall using the service via remote method invocations across the secure intranet.

└─Secure Web Service use

Scalability and replication



- To ensure scalability the Web service is replicated across multiple hosts.
- Multiple brokers are available.
- There are numerous first party clients behind the firewall using the service via remote method invocations across the secure intranet.
- There are numerous second party clients outside the firewall.

-Secure Web Service use

Security and use of encryption



 Second party clients need to use encryption to ensure authenticity and confidentiality. First party clients do not.

Secure Web Service use

Security and use of encryption



- Second party clients need to use encryption to ensure authenticity and confidentiality. First party clients do not.
- Brokers add decryption and encryption steps to build end-to-end security from point-to-point security.

Secure Web Service use

Security and use of encryption



- Second party clients need to use encryption to ensure authenticity and confidentiality. First party clients do not.
- Brokers add decryption and encryption steps to build end-to-end security from point-to-point security.
 - When processing a request from a second party client brokers decrypt the request before re-encrypting it for the Web service.
- Examples

Secure Web Service use

Security and use of encryption



- Second party clients need to use encryption to ensure authenticity and confidentiality. First party clients do not.
- Brokers add decryption and encryption steps to build end-to-end security from point-to-point security.
 - When processing a request from a second party client brokers decrypt the request before re-encrypting it for the Web service.
 - When the response to a request is returned to the broker it decrypts the response before re-encrypting it for the client.

-Secure Web Service use

PEPA model: Second party clients



A second party client composes service requests, encrypts these and sends them to its broker.

-Secure Web Service use



- A second party client composes service requests, encrypts these and sends them to its broker.
- It then waits for a response from the broker.

└─ Secure Web Service use



- A second party client composes service requests, encrypts these and sends them to its broker.
- It then waits for a response from the broker.
- The rate at which the first three activities happen is under the control of the client.

Secure Web Service use



- A second party client composes service requests, encrypts these and sends them to its broker.
- It then waits for a response from the broker.
- The rate at which the first three activities happen is under the control of the client.
- The rate at which responses are produced is determined by the interaction of the broker and the service endpoint.

└─ Secure Web Service use



└─ Secure Web Service use

PEPA model: Brokers



• The broker is inactive until it receives a request.

-Secure Web Service use

PEPA model: Brokers



- The broker is inactive until it receives a request.
- It then decrypts the request before re-encrypting it for the Web service to ensure end-to-end security.

Secure Web Service use

PEPA model: Brokers



- The broker is inactive until it receives a request.
- It then decrypts the request before re-encrypting it for the Web service to ensure end-to-end security.
- It forwards the request to the Web service and then waits for a response.

└─Secure Web Service use

PEPA model: Brokers



- The broker is inactive until it receives a request.
- It then decrypts the request before re-encrypting it for the Web service to ensure end-to-end security.
- It forwards the request to the Web service and then waits for a response.
- The corresponding decryption and re-encrytion are performed before returning the response to the client.

-Secure Web Service use

PEPA model: Brokers



Broker_{idle} Broker_{dec_input} Broker_{enc_input} Broker_{sending} Broker_{waiting} Broker_{dec_resp} Broker_{enc_resp} Broker_{replying}

 $\stackrel{def}{=}$ $(request_{h}, \top)$. Broker_{dec_input} $\stackrel{def}{=}$ (decrypt_{sp}, r_{b_dec_sp}).Broker_{enc_input} $\stackrel{def}{=}$ (encrypt_{ws}, r_{h enc ws}).Broker_{sending} def = (request_{ws}, r_{b_req}).Broker_{waiting} $\stackrel{def}{=}$ $(response_{ws}, \top)$. Broker_{dec resp} $\stackrel{def}{=}$ (decrypt_{ws}, r_{b_dec_ws}).Broker_{enc_resp} $\stackrel{def}{=}$ (encrypt_{sp}, r_{b_enc_sp}).Broker_{replying} $\stackrel{def}{=}$ (response_b, r_{b_resp}).Broker_{idle}

-Secure Web Service use

PEPA model: First party clients



The lifetime of a first party client mirrors that of a second party client except that encryption need not be used when all of the communication is conducted across a secure intranet. - Examples

Secure Web Service use

PEPA model: First party clients



- The lifetime of a first party client mirrors that of a second party client except that encryption need not be used when all of the communication is conducted across a secure intranet.
- Also the service may be invoked by a remote method invocation to the host machine instead of via HTTP.

Secure Web Service use

PEPA model: First party clients



- The lifetime of a first party client mirrors that of a second party client except that encryption need not be used when all of the communication is conducted across a secure intranet.
- Also the service may be invoked by a remote method invocation to the host machine instead of via HTTP.
- Thus the first party client experiences the Web service as a blocking remote method invocation.

└─ Secure Web Service use

PEPA model: First party clients



$$\begin{array}{lll} FPC_{idle} & \stackrel{\text{def}}{=} & (compose_{fp}, r_{fp_cmp}).FPC_{calling} \\ FPC_{calling} & \stackrel{\text{def}}{=} & (invoke_{ws}, r_{fp_inv}).FPC_{blocked} \\ FPC_{blocked} & \stackrel{\text{def}}{=} & (result_{ws}, \top).FPC_{idle} \end{array}$$

-Secure Web Service use

PEPA model: Web service



There are two ways in which the service is executed, leading to a choice in the process algebra model taking the service process into one or other of its two modes of execution.

Secure Web Service use

PEPA model: Web service



- There are two ways in which the service is executed, leading to a choice in the process algebra model taking the service process into one or other of its two modes of execution.
- In either case, the duration of the execution of the service itself is unchanged.

Secure Web Service use

PEPA model: Web service



- There are two ways in which the service is executed, leading to a choice in the process algebra model taking the service process into one or other of its two modes of execution.
- In either case, the duration of the execution of the service itself is unchanged.
- The difference is only in whether encryption is needed and whether the result is delivered via HTTP or not.

-Secure Web Service use

PEPA model: Web service



 $\begin{array}{ccc} WS_{idle} & \stackrel{\hookrightarrow}{=} \\ & & + \\ WS_{decoding} & \stackrel{def}{=} \\ WS_{execution} & \stackrel{def}{=} \\ WS_{securing} & \stackrel{def}{=} \\ WS_{responding} & \stackrel{def}{=} \\ WS_{method} & \stackrel{def}{=} \\ WS_{returning} & \stackrel{def}{=} \end{array}$

-Secure Web Service use

PEPA model: Web service



WS_{idle} WS_{decoding} WS_{execution} WS_{securing} WS_{responding} WS_{method} WS_{returning}

-Secure Web Service use

PEPA model: Web service



WS_{idle} WS_{decoding} WS_{execution} WS_{securing} WS_{responding} WS_{method} WS_{returning}

-Secure Web Service use

PEPA model: Web service



 $\begin{array}{ccc} WS_{idle} & \stackrel{\hookrightarrow}{=} \\ & & + \\ WS_{decoding} & \stackrel{def}{=} \\ WS_{execution} & \stackrel{def}{=} \\ WS_{securing} & \stackrel{def}{=} \\ WS_{responding} & \stackrel{def}{=} \\ WS_{method} & \stackrel{def}{=} \\ WS_{returning} & \stackrel{def}{=} \end{array}$

-Secure Web Service use

PEPA model: System composition

In the initial state of the system model we represent each of the four component types being initially in their idle state.

$$\begin{aligned} & \textit{System} \stackrel{\text{\tiny def}}{=} (\textit{SPC}_{\textit{idle}} \Join \textit{Broker}_{\textit{idle}}) \bowtie_{\mathcal{L}} (\textit{WS}_{\textit{idle}} \Join \textit{FPC}_{\textit{idle}}) \\ & \text{where} \quad \mathcal{K} = \{\textit{request}_b, \textit{response}_b\} \\ & \mathcal{L} = \{\textit{request}_{ws}, \textit{response}_{ws}\} \\ & \mathcal{M} = \{\textit{invoke}_{ws}, \textit{result}_{ws}\} \end{aligned}$$

PEPA model: System composition

In the initial state of the system model we represent each of the four component types being initially in their idle state.

$$\begin{aligned} System \stackrel{\text{def}}{=} (SPC_{idle} \bigotimes_{\mathcal{K}} Broker_{idle}) & \underset{\mathcal{L}}{\boxtimes} (WS_{idle} \bigotimes_{\mathcal{M}} FPC_{idle}) \\ \text{where} \quad \mathcal{K} = \{ request_b, response_b \} \\ \mathcal{L} = \{ request_{ws}, response_{ws} \} \\ \mathcal{M} = \{ invoke_{ws}, result_{ws} \} \end{aligned}$$

This model represents the smallest possible instance of the system, where there is one instance of each component type. We evaluate the system as the number of clients, brokers, and copies of the service increase.

-Secure Web Service use



 We compare fluid approximation, ODE-based, evaluation against other techniques which could be used to analyse the model.

Secure Web Service use

Cost of analysis

- We compare fluid approximation, ODE-based, evaluation against other techniques which could be used to analyse the model.
- We compare against steady-state and transient analysis based on an explicit state representation, as implemented by the PRISM probabilistic model-checker (which provides PEPA as one of its input languages). We also compare against Monte Carlo Markov Chain simulation.

-Secure Web Service use

Comparison of analysis types

We report only a single run of the simulation analysis. In practice, due to the stochastic nature of the analysis, this would need to be re-run multiple times to produce results comparable to the ODE-based analysis. - Examples

Secure Web Service use

Comparison of analysis types

- We report only a single run of the simulation analysis. In practice, due to the stochastic nature of the analysis, this would need to be re-run multiple times to produce results comparable to the ODE-based analysis.
- Moreover, note that the number of ODEs is constant regardless of the number of components in the system, whilst the state space grows dramatically.

Making stochastic process algebras count — Jane Hillston

└─Secure Web Service use

└─Secure Web Service use

Second party clients	Brokers	Web service instances	First party clients	Number of states in the full state-space	Number of states in the aggregated state-space	Sparse matrix steady-state	Matrix/MTBDD steady-state	Transient solution for time $t = 100$	MCMC simulation one run to $t = 100$	ODE solution
1	1	1	1	48	48	1.04	1.10	1.01	2.47	2.81
2	2	2	2	6,304	860	2.15	2.26	2.31	2.45	2.81

└─Secure Web Service use

Second party clients	Brokers	Web service instances	First party clients	Number of states in the full state-space	Number of states in the aggregated state-space	Sparse matrix steady-state	Matrix/MTBDD steady-state	Transient solution for time $t = 100$	MCMC simulation one run to $t = 100$	ODE solution
1	1	1	1	48	48	1.04	1.10	1.01	2.47	2.81
2	2	2	2	6,304	860	2.15	2.26	2.31	2.45	2.81
3	3	3	3	1,130,496	161,296	172.48	255.48	588.80	2.48	2.83

└─Secure Web Service use

Second party clients	Brokers	Web service instances	First party clients	Number of states in the full state-space	Number of states in the aggregated state-space	Sparse matrix steady-state	Matrix/MTBDD steady-state	Transient solution for time $t=100$	MCMC simulation one run to $t = 100$	ODE solution
1	1	1	1	48	48	1.04	1.10	1.01	2.47	2.81
2	2	2	2	6,304	860	2.15	2.26	2.31	2.45	2.81
3	3	3	3	1,130,496	161,296	172.48	255.48	588.80	2.48	2.83
4	4	4	4	>234M	_	_	_	_	2.44	2.85

└─Secure Web Service use

Second party clients	Brokers	Web service instances	First party clients	Number of states in the full state-space	Number of states in the aggregated state-space	Sparse matrix steady-state	Matrix/MTBDD steady-state	Transient solution for time $t = 100$	MCMC simulation one run to $t = 100$	ODE solution
1	1	1	1	48	48	1.04	1.10	1.01	2.47	2.81
2	2	2	2	6,304	860	2.15	2.26	2.31	2.45	2.81
3	3	3	3	1,130,496	161,296	172.48	255.48	588.80	2.48	2.83
4	4	4	4	>234M	_	_	_	_	2.44	2.85
100	100	100	100	_	_	_	_	_	2.78	2.78

└─Secure Web Service use

Second party clients	Brokers	Web service instances	First party clients	Number of states in the full state-space	Number of states in the aggregated state-space	Sparse matrix steady-state	Matrix/MTBDD steady-state	Transient solution for time $t = 100$	MCMC simulation one run to $t = 100$	ODE solution
1	1	1	1	48	48	1.04	1.10	1.01	2.47	2.81
2	2	2	2	6,304	860	2.15	2.26	2.31	2.45	2.81
3	3	3	3	1,130,496	161,296	172.48	255.48	588.80	2.48	2.83
4	4	4	4	>234M	_	-	_	_	2.44	2.85
100	100	100	100	_	_	-	_	_	2.78	2.78
1000	100	500	1000	-		-	-	_	3.72	2.77

└─Secure Web Service use

Second party clients	Brokers	Web service instances	First party clients	Number of states in the full state-space	Number of states in the aggregated state-space	Sparse matrix steady-state	Matrix/MTBDD steady-state	Transient solution for time $t=100$	MCMC simulation one run to $t = 100$	ODE solution
1	1	1	1	48	48	1.04	1.10	1.01	2.47	2.81
2	2	2	2	6,304	860	2.15	2.26	2.31	2.45	2.81
3	3	3	3	1,130,496	161,296	172.48	255.48	588.80	2.48	2.83
4	4	4	4	>234M	_	-	-	-	2.44	2.85
100	100	100	100	-	_	-	-	-	2.78	2.78
1000	100	500	1000	_	_	_	_	_	3.72	2.77
1000	1000	1000	1000	-	-	-	-	-	5.44	2.77

-Secure Web Service use

Time series analysis via ODEs

We now consider the results from our solution of the PEPA Web Service model as a system of ODEs with the number of clients of both kinds, brokers, and web service instances all 1000.
Secure Web Service use

Time series analysis via ODEs

- We now consider the results from our solution of the PEPA Web Service model as a system of ODEs with the number of clients of both kinds, brokers, and web service instances all 1000.
- The results as presented from our ODE integrator are time-series plots of the number of each type of component behaviour as a function of time.

└─ Secure Web Service use

Time series analysis via ODEs

- We now consider the results from our solution of the PEPA Web Service model as a system of ODEs with the number of clients of both kinds, brokers, and web service instances all 1000.
- The results as presented from our ODE integrator are time-series plots of the number of each type of component behaviour as a function of time.
- We can observe an initial flurry of activity until the system stabilises into its steady-state equilibrium at time (around) t = 50.

└─Secure Web Service use

Second party clients



└─Secure Web Service use

Brokers



└─Secure Web Service use

First party clients



└─Secure Web Service use

Web service



Emergency egress

Designing for human crowd dynamics

- Widespread take up of mobile and communicating computational devices is making ubiquitous systems a reality and creating new ways for us to interact with our environment.
- One application is to provide routing information to help people navigate through unfamiliar locations.
- In these cases the dynamic behaviour of the system as a whole is important to ensure the satisfaction of the users.
- Emergency egress can be regarded as a particular case, when the location may be familiar but circumstances may alter the usual topology and make efficient movement particularly important.

Emergency egress

Example scenario

RA 211		18w	18e			SE 13
LW 25	HA 133					LE 16
SW 22	RB 92	16w	_	RC 98	18e	·

The layout of the building is described in terms of the arrangement of the rooms, hallways, landing and stairs. Each has a capacity and may have an initial occupancy.

Emergency egress

Example scenario

RA 211		18w	18e			SE 13
LW 25	HA 133					LE 16
SW 22	RB 92	16w		RC 98	18e	·

The layout of the building is described in terms of the arrangement of the rooms, hallways, landing and stairs. Each has a capacity and may have an initial occupancy.

Bio-PEPA components describe the behaviours of individuals, but also rooms and information dissemination.

Emergency egress

Model specification

```
// BUILDING LAYOUT (COMPARTMENTS)
location ra : size = normal room, type= compartment;
location d1 ra ha : size = normal door, type= compartment;
. . .
// PARAMETERS SET UP
                                //3 + (60/7):
to rad1 = 6;
from d1 = door exit rate;
occupancy d1 = D1 ra e@d1 ra ha + D1 ra w@d1 ra ha + ... :
full d1 = H(capacity d1 - occupancy d1);
switch d1 = open d1*full d1;
// AGENT DYNAMICS (FUNCTIONAL RATES)
// From ra to ha through dl
kineticLawOf ra e in d1 ra ha : fMA (to ra d1 * switch d1 * ra e in safe);
kineticLawOf ha e out dl ra ha : fMA (from dl * open dl * safeDl ra e * allowance ha);
kineticLawOf ra_w_in_d1_ra_ha : ...
kineticLawOf ha w out d1 ra ha : ...
// ... and back
kineticLawOf ha e in dl ra ha : ...
// AGENT DEFINITIONS (SEQUENTIAL PROCESSES)
RA = (ra e in d1 ra ha, 1) \iff RA e e ra + \dots
      (ra e out dl ra ha, 1) >> RA e@ra + ...
RA w = ...
HAe =
```

```
. . .
```

- Examples
 - Emergency egress



One stochastic simulation run

- Examples
 - Emergency egress



10 stochastic simulation runs

- Examples
 - Emergency egress



500 stochastic simulation runs

- Examples
 - Emergency egress



ODE numerical simulation

Emergency egress

Example results: rerouting through mediation



Room occupancy over time without rerouting capability

Emergency egress

Example results: rerouting through mediation



Room occupancy over time with rerouting capability

Outline

Introduction

- Stochastic Process Algebra
- Collective Dynamics
- 2 Continuous Approximation
 - State variables
 - Numerical illustration

3 Fluid-Flow Semantics

Fluid Structured Operational Semantics

4 Examples

- Secure Web Service use
- Emergency egress

Conclusions

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

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

- Many interesting and important systems can be regarded as examples of collective dynamics and emergent behaviour.
- Process algebras, such as PEPA and Bio-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.

- Many interesting and important systems can be regarded as examples of collective dynamics and emergent behaviour.
- Process algebras, such as PEPA and Bio-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.

On-going work

Time series plots counting the populations of components over time tell us a great deal about the dynamics of the system but are not necessary the information we require.

On-going work

- Time series plots counting the populations of components over time tell us a great deal about the dynamics of the system but are not necessary the information we require.
- Recent work has establish how performance measures such as throughput, and average response time can be derived from the ODE solutions.

On-going work

- Time series plots counting the populations of components over time tell us a great deal about the dynamics of the system but are not necessary the information we require.
- Recent work has establish how performance measures such as throughput, and average response time can be derived from the ODE solutions.
- On-going work is investigating the use of probes to query the model by adding components to the model whose sole purpose is to gather statistics.

On-going work

- Time series plots counting the populations of components over time tell us a great deal about the dynamics of the system but are not necessary the information we require.
- Recent work has establish how performance measures such as throughput, and average response time can be derived from the ODE solutions.
- On-going work is investigating the use of probes to query the model by adding components to the model whose sole purpose is to gather statistics.
- Using this technique we can now derive measures such as cumulative response time from the fluid approximation of the model.

Thanks!

Thanks!

Acknowledgements: collaborators

Thanks to many co-authors and collaborators: Andrea Bracciali, Jeremy Bradley, Luca Bortolussi, Federica Ciocchetta, **Allan Clark**, Jie Ding, Adam Duguid, Vashti Galpin, **Stephen Gilmore**, Diego Latella, Mieke Massink, **Mirco Tribastone**, and others.

Thanks!

Acknowledgements: collaborators

Thanks to many co-authors and collaborators: Andrea Bracciali, Jeremy Bradley, Luca Bortolussi, Federica Ciocchetta, **Allan Clark**, Jie Ding, Adam Duguid, Vashti Galpin, **Stephen Gilmore**, Diego Latella, Mieke Massink, **Mirco Tribastone**, and others.

Acknowledgements: funding

Thanks to EPRSC for the Process Algebra for Collective Dynamics grant and the CEC IST-FET programme for the SENSORIA project which have supported this work.

Thanks!

Acknowledgements: collaborators

Thanks to many co-authors and collaborators: Andrea Bracciali, Jeremy Bradley, Luca Bortolussi, Federica Ciocchetta, **Allan Clark**, Jie Ding, Adam Duguid, Vashti Galpin, **Stephen Gilmore**, Diego Latella, Mieke Massink, **Mirco Tribastone**, and others.

Acknowledgements: funding

Thanks to EPRSC for the Process Algebra for Collective Dynamics grant and the CEC IST-FET programme for the SENSORIA project which have supported this work.

More information:

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

Alternative Representations



Alternative Representations

