Adventures in Systems Biology

Jane Hillston. LFCS, University of Edinburgh

18th July 2005

Joint work with Muffy Calder and Stephen Gilmore

Jane Hillston. LFCS, University of Edinburgh.

Outline

Introduction

Case Study in Systems Biology

Continuous Approximation and Differential Equations

Case Study in Web Services

Jane Hillston. LFCS, University of Edinburgh.

Outline

Introduction

Case Study in Systems Biology

Continuous Approximation and Differential Equations

Case Study in Web Services

Jane Hillston. LFCS, University of Edinburgh.

The PEPA project

► The PEPA project started in Edinburgh in 1991.

Jane Hillston. LFCS, University of Edinburgh.

The PEPA project

- ► The PEPA project started in Edinburgh in 1991.
- It was motivated by problems encountered when carrying out performance analysis of large computer and communication systems, based on numerical analysis of Markov processes.

Jane Hillston. LFCS, University of Edinburgh.

The PEPA project

- The PEPA project started in Edinburgh in 1991.
- It was motivated by problems encountered when carrying out performance analysis of large computer and communication systems, based on numerical analysis of Markov processes.
- Process algebras offered a compositional description technique supported by apparatus for formal reasoning.

The PEPA project

- The PEPA project started in Edinburgh in 1991.
- It was motivated by problems encountered when carrying out performance analysis of large computer and communication systems, based on numerical analysis of Markov processes.
- Process algebras offered a compositional description technique supported by apparatus for formal reasoning.
- Performance Evaluation Process Algebra (PEPA) sought to address these problems by the introduction of a suitable process algebra.

The PEPA project

- The PEPA project started in Edinburgh in 1991.
- It was motivated by problems encountered when carrying out performance analysis of large computer and communication systems, based on numerical analysis of Markov processes.
- Process algebras offered a compositional description technique supported by apparatus for formal reasoning.
- Performance Evaluation Process Algebra (PEPA) sought to address these problems by the introduction of a suitable process algebra.
- The project has sought to investigate and exploit the interplay between the process algebra and the continuous time Markov chain (CTMC).

PEPA for Performance Modelling

Model Construction: Compositionality leads to

- ease of construction
- reusable submodels
- easy to understand models

Jane Hillston. LFCS, University of Edinburgh.

PEPA for Performance Modelling

Model Construction: Compositionality leads to

- ease of construction
- reusable submodels
- easy to understand models

Model Manipulation: Equivalence relations lead to

- term rewriting/state space reduction techniques
- aggregation techniques based on lumpability

PEPA for Performance Modelling

Model Construction: Compositionality leads to

- ease of construction
- reusable submodels
- easy to understand models

Model Manipulation: Equivalence relations lead to

- term rewriting/state space reduction techniques
- aggregation techniques based on lumpability

Model Solution: Formal semantics: lead to

- automatic identification of classes of models susceptible to efficient solution
- use of logics to express performance measures

Example PEPA Case Studies

- Multiprocessor access-contention protocols (Gilmore, Hillston and Ribaudo, Edinburgh and Turin)
- Protocols for fault-tolerant systems (Clark, Gilmore, Hillston and Ribaudo, Edinburgh and Turin)
- Multimedia traffic characteristics (Bowman et al, Kent)
- Database systems (The STEADY group, Heriot-Watt University)
- Software Architectures (Pooley, Bradley and Thomas, Heriot-Watt and Durham)
- Switch behaviour in active networks (Hillston, Kloul and Mokhtari, Edinburgh and Versailles)

Motivation: why use process algebras to model pathways?

 Process algebraic formulation makes interactions/constraints explicit – not the case with classical ordinary differential equation (ODE) models. Structure can also be apparent.

Jane Hillston. LFCS, University of Edinburgh.

Motivation: why use process algebras to model pathways?

- Process algebraic formulation makes interactions/constraints explicit – not the case with classical ordinary differential equation (ODE) models. Structure can also be apparent.
- Equivalence relations allow formal comparison of high-level descriptions.

Motivation: why use process algebras to model pathways?

- Process algebraic formulation makes interactions/constraints explicit – not the case with classical ordinary differential equation (ODE) models. Structure can also be apparent.
- Equivalence relations allow formal comparison of high-level descriptions.
- Access to techniques for reasoning about livelocks, deadlocks

Motivation: why use process algebras to model pathways?

- Process algebraic formulation makes interactions/constraints explicit – not the case with classical ordinary differential equation (ODE) models. Structure can also be apparent.
- Equivalence relations allow formal comparison of high-level descriptions.
- Access to techniques for reasoning about livelocks, deadlocks and when using a stochastic process algebra, performance (transient or steady state) and sophisticated reasoning using stochastic logics such as CSL and probabilistic model checking.

PEPA

Stochastic Process Algebra

Attractive features of process algebras

Jane Hillston. LFCS, University of Edinburgh.

PEPA

Stochastic Process Algebra

Attractive features of process algebras

Compositionality

Jane Hillston. LFCS, University of Edinburgh.

Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition

Jane Hillston. LFCS, University of Edinburgh.

Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition
- Parsimony

Jane Hillston. LFCS, University of Edinburgh.



Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition
- Parsimony

+ Quantification

Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

< □ > < □ > < 三 > < 三 > < 三 > のへ()

Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition
- Parsimony

- + Quantification
 - Actions have duration

Jane Hillston. LFCS, University of Edinburgh.

Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition
- Parsimony

- + Quantification
 - Actions have duration
 - Probabilistic branching

Jane Hillston. LFCS, University of Edinburgh.

Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition
- Parsimony

- + Quantification
 - Actions have duration
 - Probabilistic branching

Performance Evaluation Process Algebra (PEPA)

Jane Hillston. LFCS, University of Edinburgh.

Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition
- Parsimony

- + Quantification
 - Actions have duration
 - Probabilistic branching

Performance Evaluation Process Algebra (PEPA)

Models are constructed from components which engage in activities.



Jane Hillston. LFCS, University of Edinburgh.

Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition
- Parsimony

- + Quantification
 - Actions have duration
 - Probabilistic branching

Performance Evaluation Process Algebra (PEPA)

Models are constructed from components which engage in activities.



Jane Hillston. LFCS, University of Edinburgh.

Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition
- Parsimony

- + Quantification
 - Actions have duration
 - Probabilistic branching

Performance Evaluation Process Algebra (PEPA)

Models are constructed from components which engage in activities.



Jane Hillston. LFCS, University of Edinburgh.

Stochastic Process Algebra

Attractive features of process algebras

- Compositionality
- Formal definition
- Parsimony

- + Quantification
 - Actions have duration
 - Probabilistic branching

Performance Evaluation Process Algebra (PEPA)

Models are constructed from components which engage in activities.



Jane Hillston. LFCS, University of Edinburgh.

The corresponding Continuous Time Markov Chain (CTMC) is derived automatically from the structured operational semantics which define the language:

Jane Hillston. LFCS, University of Edinburgh.



The corresponding Continuous Time Markov Chain (CTMC) is derived automatically from the structured operational semantics which define the language:

PEPA MODEL

Jane Hillston. LFCS, University of Edinburgh.

The corresponding Continuous Time Markov Chain (CTMC) is derived automatically from the structured operational semantics which define the language:

PEPA SOS rules

Jane Hillston. LFCS, University of Edinburgh.

The corresponding Continuous Time Markov Chain (CTMC) is derived automatically from the structured operational semantics which define the language:



Jane Hillston. LFCS, University of Edinburgh.

The corresponding Continuous Time Markov Chain (CTMC) is derived automatically from the structured operational semantics which define the language:



Jane Hillston. LFCS, University of Edinburgh.

The corresponding Continuous Time Markov Chain (CTMC) is derived automatically from the structured operational semantics which define the language:



Note that contrary to classical process algebras, the multiplicity of transitions is important.

The corresponding Continuous Time Markov Chain (CTMC) is derived automatically from the structured operational semantics which define the language:



Note that contrary to classical process algebras, the multiplicity of transitions is important.

The natural equivalence relation in the process algebra strong equivalence or Markovian bisimulation generates a lumpable partition in the underlying CTMC.

The corresponding Continuous Time Markov Chain (CTMC) is derived automatically from the structured operational semantics which define the language:



Note that contrary to classical process algebras, the multiplicity of transitions is important.

The natural equivalence relation in the process algebra strong equivalence or Markovian bisimulation generates a lumpable partition in the underlying CTMC.
PEPA

$$S ::= (\alpha, r).S | S + S | A$$
$$P ::= S | P \bowtie_{L} P | P/L$$

Jane Hillston. LFCS, University of Edinburgh.

PEPA

$$S ::= (\alpha, r).S | S + S | A$$
$$P ::= S | P \bowtie_{L} P | P/L$$

PREFIX: $(\alpha, r).S$ designated first action

Jane Hillston. LFCS, University of Edinburgh.

PEPA

$$S ::= (\alpha, r).S | S + S | A$$

$$P ::= S | P \bowtie_{L} P | P/L$$

PREFIX: CHOICE: $(\alpha, r).S$ designated first action S+S competing components (race policy)

Jane Hillston. LFCS, University of Edinburgh.

PEPA

- $S ::= (\alpha, r).S | S + S | A$ $P ::= S | P \bowtie_{L} P | P/L$
- PREFIX: $(\alpha, r).S$ designated first actionCHOICE:S+Scompeting components
(race policy)CONSTANT: $A \stackrel{def}{=} S$ assigning names

Jane Hillston. LFCS, University of Edinburgh.

PEPA

- $S ::= (\alpha, r).S | S + S | A$ $P ::= S | P \bowtie_{L} P | P/L$
- PREFIX: $(\alpha, r).S$ designated first actionCHOICE:S + Scompeting components
(race policy)CONSTANT: $A \stackrel{def}{=} S$ assigning namesCOOPERATION: $P \bowtie_L P$ $\alpha \notin L$ concurrent activity
(individual actions)
 $\alpha \in L$ cooperative activity

(shared actions)

PEPA

 $S ::= (\alpha, r).S | S + S | A$ $P ::= S | P \bowtie_{L} P | P/L$

PREFIX: CHOICE:	$(\alpha, r).S$ S + S	designated first action competing components (race policy)
CONSTANT: COOPERATION:	$A \stackrel{def}{=} S$ $P \bowtie_L P$	assigning names $\alpha \notin L$ concurrent activity (<i>individual actions</i>) $\alpha \in L$ cooperative activity (<i>shared actions</i>)
HIDING:	P/L	abstraction $\alpha \in L \Rightarrow \alpha \to \tau$

Jane Hillston. LFCS, University of Edinburgh.

Outline

Introduction

Case Study in Systems Biology

Continuous Approximation and Differential Equations

Case Study in Web Services

Jane Hillston. LFCS, University of Edinburgh.

New application domains: biochemical signalling pathways

 Biological advances mean that much more is now known about the components of cells and the interactions between them.

Jane Hillston. LFCS, University of Edinburgh.

New application domains: biochemical signalling pathways

- Biological advances mean that much more is now known about the components of cells and the interactions between them.
- Systems biology aims to develop a better understanding of the processes involved.

Jane Hillston. LFCS, University of Edinburgh.

New application domains: biochemical signalling pathways

- Biological advances mean that much more is now known about the components of cells and the interactions between them.
- Systems biology aims to develop a better understanding of the processes involved.
- Stochastic process algebras have found a new role in developing models for systems biology, allowing biologists to test hypotheses and prioritise experiments.

Extracellular signalling

Extracellular signalling — communication between cells.

- signalling molecules released by one cell migrate to another;
- these molecules enter the cell and instigate a pathway, or series of reactions, which carries the information from the membrane to the nucleus;
- the Ras/Raf-1/MEK/ERK pathway conveys differentiation signals to the nucleus of a cell.



Special relevance to cancer research because when pathways operate abnormally cells divide uncontrollably.

Jane Hillston. LFCS, University of Edinburgh.			
Adventures in Systems Biology	∢ @ ≻		

The ERK signalling pathway



Researchers at the Beatson Institute in Glasgow are investigating the hypothesis that the protein RKIP has a regulatory influence on the ERK signalling pathway.

Jane Hillston. LFCS, University of Edinburgh.



Jane Hillston. LFCS, University of Edinburgh.



 Bi-directional arrows denote both forward and backward reactions;

Jane Hillston. LFCS, University of Edinburgh.



- Bi-directional arrows denote both forward and backward reactions;
- Uni-directional arrows denote reactions which are disassociations.

Jane Hillston. LFCS, University of Edinburgh.



- Bi-directional arrows denote both forward and backward reactions;
- Uni-directional arrows denote reactions which are disassociations.
- Each reagent has a variable concentration, denoted m_i.

Jane Hillston. LFCS, University of Edinburgh.



- Bi-directional arrows denote both forward and backward reactions;
- Uni-directional arrows denote reactions which are disassociations.
- Each reagent has a variable concentration, denoted m_i.
- Each reaction has a corresponding rate constant, e.g. k3, but the rate at which the reaction takes place is the product of this rate constant and the current concentrations of the used substrates.

The ERK signalling pathway



Jane Hillston. LFCS, University of Edinburgh.

PEPA models of the pathway

We have constructed two, complementary, PEPA models of the pathway.

Jane Hillston. LFCS, University of Edinburgh.

PEPA models of the pathway

We have constructed two, complementary, PEPA models of the pathway.

The first focuses on the reagents and the variations in their concentrations caused by the reactions they undertake. Each distinct protein in the pathway is represented as a distinct PEPA component.

PEPA models of the pathway

We have constructed two, complementary, PEPA models of the pathway.

- The first focuses on the reagents and the variations in their concentrations caused by the reactions they undertake. Each distinct protein in the pathway is represented as a distinct PEPA component.
- The second focuses on the sub-pathways within the pathway. Starting from a reagent which exhibits an initial concentration each sub-pathway tracks the valid reactions the substrate may pass through before returning the initial concentration of the initial substrate.

The dynamics or kinetics of the pathway are dependent on the concentrations of the reactants.

Jane Hillston. LFCS, University of Edinburgh.

The dynamics or kinetics of the pathway are dependent on the concentrations of the reactants.

In both our models concentrations are treated in an abstract fashion.

Jane Hillston. LFCS, University of Edinburgh.

The dynamics or kinetics of the pathway are dependent on the concentrations of the reactants.

In both our models concentrations are treated in an abstract fashion.

▶ In the first we distinguish only high and low concentrations.

The dynamics or kinetics of the pathway are dependent on the concentrations of the reactants.

In both our models concentrations are treated in an abstract fashion.

- ► In the first we distinguish only high and low concentrations.
- In the second concentrations are not represented explicitly at all but may be inferred from the state of the appropriate pathways.

PEPA components of the reagent-centric model



Jane Hillston. LFCS, University of Edinburgh.

PEPA components of the reagent-centric model



Each reagent gives rise to a pair of PEPA definitions, one for high concentration and one for low concentration.

Jane Hillston. LFCS, University of Edinburgh.

Configuration of the reagent-centric model



Jane Hillston. LFCS, University of Edinburgh.

PEPA components of the pathway-centric model



Jane Hillston. LFCS, University of Edinburgh.

PEPA components of the pathway-centric model



For each reagent that has an initial concentration we define the sub-pathway generated by the progression of that reagent.

Jane Hillston. LFCS, University of Edinburgh. Adventures in Systems Biology

Configuration of the pathway-centric model

To complete the model we must ensure that the pathways are constrained to interact appropriately, i.e. an association can only occur when all necessary reagents are in a state to make the reaction.



Note that this is much simpler in this case than for the reagent model.

Jane Hillston.	LFCS, I	University	of	Edinburgh
Adventures in	System	s Biology		

Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels:

Jane Hillston. LFCS, University of Edinburgh.

Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.

Jane Hillston. LFCS, University of Edinburgh.

- Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.
- The reagent-centric model can be regarded as a fine-grained view of the system.

- Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.
- The reagent-centric model can be regarded as a fine-grained view of the system.
- The pathway-centric model can be regarded as a more structural, coarse-grained view of the system.

- Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.
- The reagent-centric model can be regarded as a fine-grained view of the system.
- The pathway-centric model can be regarded as a more structural, coarse-grained view of the system.
- Applying the structured operational semantics reveals that they are strongly bisimilar
Commentary on the models

- Neither model currently "correctly" captures the rate of interaction – concentrations are discretized and rates are assumed to be constant within levels: In these examples only using high and low to modify rates in the sense of enabling or disabling activities.
- The reagent-centric model can be regarded as a fine-grained view of the system.
- The pathway-centric model can be regarded as a more structural, coarse-grained view of the system.
- Applying the structured operational semantics reveals that they are strongly bisimilar (in fact, in this case, isomorphic).

The state space



Jane Hillston. LFCS, University of Edinburgh.

<i>s</i> ₁	$\begin{split} &(Raf-1_{H}^{*},RKIP_{H},Raf-1^{*}/RKIP_{L},Raf-1^{*}/RKIP/ERK-PP_{L},\\ &ERK_{L},RKIP-P_{L},RKIP-P/RP_{L},RP_{H},MEK_{L},\\ &MEK/Raf-1_{L}^{*},MEK-PP_{H},MEK-PP/ERK_{L},ERK-PP_{H}) \end{split}$	(Pathway ₅₀ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)
-----------------------	---	---

Jane Hillston. LFCS, University of Edinburgh.

<i>s</i> 1	$\begin{split} &(Raf-1_H^*,RKIP_H,Raf-1^*/RKIP_L,Raf-1^*/RKIP/ERK-PP_L,\\ &ERK_L,RKIP-P_L,RKIP-P/RP_L,RP_H,MEK_L,\\ &MEK/Raf-1_L^*,MEK-PP_H,MEK-PP/ERK_L,ERK-PP_H) \end{split}$	(Pathway ₅₀ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)
<i>s</i> 2	$\begin{split} &(Raf-1_{H}^{*},RKIP_{H},Raf-1^{*}/RKIP_{L},Raf-1^{*}/RKIP/ERK-PP_{L},\\ &ERK_{L},RKIP-P_{L}RKIP-P/RP_{L},RP_{H},MEK_{H},\\ &MEK/Raf-1_{L}^{*},MEK-PP_{L},MEK-PP/ERK_{L},ERK-PP_{H}) \end{split}$	(Pathway ₅₁ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)

Jane Hillston. LFCS, University of Edinburgh.

<i>s</i> 1	$\begin{split} &(Raf-1_H^*,RKIP_H,Raf-1^*/RKIP_L,Raf-1^*/RKIP/ERK\text{-}PP_L,\\ &ERK_L,RKIP\text{-}P_L,RKIP\text{-}P/RP_L,RP_H,MEK_L,\\ &MEK/Raf-1_L^*,MEK\text{-}PP_H,MEK\text{-}PP/ERK_L,ERK\text{-}PP_H) \end{split}$	(Pathway ₅₀ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)
s 2	$\begin{split} & (Raf-1_{\mathrm{H}}^*,RKIP_{\mathrm{H}},Raf-1^*/RKIP_{\mathrm{L}},Raf-1^*/RKIP/ERK\operatorname{-PP}_{\mathrm{L}},\\ & ERK_{\mathrm{L}},RKIP\operatorname{-P}_{\mathrm{L}}RKIP\operatorname{-P}/RP_{\mathrm{L}},RP_{\mathrm{H}},MEK_{\mathrm{H}},\\ & MEK/Raf-1_{\mathrm{L}}^*,MEK\operatorname{-PP}_{\mathrm{L}},MEK\operatorname{-PP}/ERK_{\mathrm{L}},ERK\operatorname{-PP}_{\mathrm{H}}) \end{split}$	(Pathway ₅₁ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)
s 3	$\begin{split} & (Raf-1_L^*,RKIP_L,Raf-1^*/RKIP_H,Raf-1^*/RKIP/ERK-PP_L, \\ & ERK_L,RKIP-P_L,RKIP-P/RP_L,RP_H,MEK_L, \\ & MEK/Raf-1_L^*,MEK-PP_H,MEK-PP/ERK_L,ERK-PP_H) \end{split}$	(Pathway ₅₀ , Pathway ₄₁ , Pathway ₃₀ , Pathway ₂₁ , Pathway ₁₀)

Jane Hillston. LFCS, University of Edinburgh.

<i>s</i> 1	$\begin{split} &(Raf-1_H^*,RKIP_H,Raf-1^*/RKIP_L,Raf-1^*/RKIP/ERK\operatorname{-PP}_L,\\ &ERK_L,RKIP\operatorname{-P}_L,RKIP\operatorname{-P}/RP_L,RP_H,MEK_L,\\ &MEK/Raf-1_L^*,MEK\operatorname{-PP}_H,MEK\operatorname{-PP}/ERK_L,ERK\operatorname{-PP}_H) \end{split}$	(Pathway ₅₀ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)
<i>s</i> 2	$\begin{split} & (Raf-1_{\mathrm{H}}^*,RKIP_{\mathrm{H}},Raf-1^*/RKIP_{\mathrm{L}},Raf-1^*/RKIP/ERK\operatorname{-PP}_{\mathrm{L}},\\ & ERK_{\mathrm{L}},RKIP\operatorname{-P}_{\mathrm{L}}RKIP\operatorname{-P}/RP_{\mathrm{L}},RP_{\mathrm{H}},MEK_{\mathrm{H}},\\ & MEK/Raf-1_{\mathrm{L}}^*,MEK\operatorname{-PP}_{\mathrm{L}},MEK\operatorname{-PP}/ERK_{\mathrm{L}},ERK\operatorname{-PP}_{\mathrm{H}}) \end{split}$	(Pathway ₅₁ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)
s 3	$\begin{split} & (Raf-1_L^*,RKIP_L,Raf-1^*/RKIP_H,Raf-1^*/RKIP/ERK-PP_L,\\ & ERK_L,RKIP-P_L,RKIP-P/RP_L,RP_H,MEK_L,\\ & MEK/Raf-1_L^*,MEK-PP_H,MEK-PP/ERK_L,ERK-PP_H) \end{split}$	(Pathway ₅₀ , Pathway ₄₁ , Pathway ₃₀ , Pathway ₂₁ , Pathway ₁₀)
:	:	:

Jane Hillston. LFCS, University of Edinburgh.

<i>s</i> 1	$\begin{split} &(Raf-1_H^*,RKIP_H,Raf-1^*/RKIP_L,Raf-1^*/RKIP/ERK\text{-}PP_L,\\ &ERK_L,RKIP\text{-}P_L,RKIP\text{-}P/RP_L,RP_H,MEK_L,\\ &MEK/Raf-1_L^*,MEK\text{-}PP_H,MEK\text{-}PP/ERK_L,ERK\text{-}PP_H) \end{split}$	(Pathway ₅₀ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)
<i>s</i> 2	$\begin{split} &(Raf-1_H^*,RKIP_H,Raf-1^*/RKIP_L,Raf-1^*/RKIP/ERK-PP_L,\\ &ERK_L,RKIP-P_LRKIP-P/RP_L,RP_H,MEK_H,\\ &MEK/Raf-1_L^*,MEK-PP_L,MEK-PP/ERK_L,ERK-PP_H) \end{split}$	(Pathway ₅₁ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)
<i>s</i> 3	$\begin{split} &(Raf-1_L^*,RKIP_L,Raf-1^*/RKIP_H,Raf-1^*/RKIP/ERK-PP_L,\\ &ERK_L,RKIP-P_L,RKIP-P/RP_L,RP_H,MEK_L,\\ &MEK/Raf-1_L^*,MEK-PP_H,MEK-PP/ERK_L,ERK-PP_H) \end{split}$	(Pathway ₅₀ , Pathway ₄₁ , Pathway ₃₀ , Pathway ₂₁ , Pathway ₁₀)
:		
\$28	$\begin{split} & (Raf-1_{\mathrm{L}}^*,RKIP_{\mathrm{L}},Raf-1^*/RKIP_{\mathrm{L}},Raf-1^*/RKIP/ERK-PP_{\mathrm{L}},\\ & ERK_{\mathrm{L}},RKIP-P_{\mathrm{H}},RKIP-P/RP_{\mathrm{L}},RP_{\mathrm{H}},MEK_{\mathrm{L}},\\ & MEK/Raf-1_{\mathrm{H}}^*,MEK-PP_{\mathrm{L}},MEK-PP/ERK_{\mathrm{L}},ERK-PP_{\mathrm{H}}) \end{split}$	(Pathway ₅₀ , Pathway ₄₀ , Pathway ₃₀ , Pathway ₂₀ , Pathway ₁₀)

Jane Hillston. LFCS, University of Edinburgh.

Quantified analysis – *k8product*

Approximating a variation in the initial concentration of RKIP by varying the rate constant k1, we can assess the impact on the production of ERK-PP.



Jane Hillston. LFCS, University of Edinburgh.

Quantified analysis - k14 product

Similarly we can assess the impact on the production of MEK-PP.



Jane Hillston. LFCS, University of Edinburgh.

Outline

Introduction

Case Study in Systems Biology

Continuous Approximation and Differential Equations

Case Study in Web Services

Jane Hillston. LFCS, University of Edinburgh.

Deriving differential equation: overview

The ODEs are the familiar mathematical model for the biochemists.

Jane Hillston. LFCS, University of Edinburgh.

Deriving differential equation: overview

- The ODEs are the familiar mathematical model for the biochemists.
- There should be one equation for each reagent/concentration, indicating how the concentration varies over time.

Jane Hillston. LFCS, University of Edinburgh.

Deriving differential equation: overview

- The ODEs are the familiar mathematical model for the biochemists.
- There should be one equation for each reagent/concentration, indicating how the concentration varies over time.
- Standard solution tools are available for solution of this equations, which are known and trusted by the biologists.

Deriving differential equation: overview

- The ODEs are the familiar mathematical model for the biochemists.
- There should be one equation for each reagent/concentration, indicating how the concentration varies over time.
- Standard solution tools are available for solution of this equations, which are known and trusted by the biologists.
- From the reagent-centric PEPA model we can see the influence of the reactions on each concentration – this is the basis of the derivation via a directed bipartite graph termed the activity graph.

Deriving differential equations: activity graph



Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

▲□▶ ▲□▶ ▲三▶ ▲三▶ 三 のへで

Deriving differential equations: activity matrix

	k1	k2	k3	k4	k5	k6	k7	k8	k9	k10	k11	k12	k13	k14	k15
Raf-1*	-1	+1	0	0	$^{+1}$	0	0	0	0	0	0	-1	$^{+1}$	$^{+1}$	0
RKIP	-1	$^{+1}$	0	0	0	0	0	0	0	0	$^{+1}$	0	0	0	0
Raf-1*/RKIP	$^{+1}$	-1	-1	$^{+1}$	0	0	0	0	0	0	0	0	0	0	0
Raf-1*/RKIP/ERK-PP	0	0	+1	-1	-1	0	0	0	0	0	0	0	0	0	0
ERK	0	0	0	0	$^{+1}$	-1	+1	0	0	0	0	0	0	0	0
RKIP-P	0	0	0	0	$^{+1}$	0	0	0	-1	$^{+1}$	0	0	0	0	0
MEK-PP	0	0	0	0	0	-1	+1	$^{+1}$	0	0	0	0	0	$^{+1}$	-1
MEK-PP/ERK	0	0	0	0	0	+1	-1	-1	0	0	0	0	0	0	0
ERK-PP	0	0	-1	$^{+1}$	0	0	0	$^{+1}$	0	0	0	0	0	0	0
RP	0	0	0	0	0	0	0	0	-1	$^{+1}$	$^{+1}$	0	0	0	0
RKIP-P/RP	0	0	0	0	0	0	0	0	$^{+1}$	-1	-1	0	0	0	0
MEK	0	0	0	0	0	0	0	0	0	0	0	-1	$^{+1}$	0	$^{+1}$
MEK/Raf-1*	0	0	0	0	0	0	0	0	0	0	0	$^{+1}$	-1	-1	0

Each row corresponds to a single reagent; the entries in a row indicate whether an activity (column) increases the concentration (+1), decreases it (-1) or has no impact (0). Each column corresponds to a single reaction; the negative entries indicate those substrates which are *used* in the reaction.

Jane Hillston. LFCS, University of Edinburgh.

Deriving differential equations: activity matrix

	k1	k2	k3	<i>k</i> 4	<i>k</i> 5	<i>k</i> 6		conc.
Raf-1*	-1	+1	0	0	+1	0		m_1
RKIP	-1	+1	0	0	0	0		<i>m</i> ₂
Raf-1*/RKIP	+1	-1	-1	+1	0	0		(m ₃)
Raf-1*/RKIP/ERK-PP	0	0	+1	-1	-1	0		<i>m</i> ₄
ERK	0	0	0	0	+1	-1		<i>m</i> 5
RKIP-P	0	0	0	0	+1	0		<i>m</i> ₆
MEK-PP	0	0	0	0	0	-1		<i>m</i> 7
MEK-PP/ERK	0	0	0	0	0	+1		<i>m</i> 8
ERK-PP	0	0	-1	+1	0	0		<i>m</i> 9
:	:	:	:	:	:	:	·	

$$\frac{dm_3(t)}{dt} = k_1 m_1(t)m_2(t) - k_2 m_3(t) - k_3 m_3(t)m_9(t) + k_4 m_4(t)$$

Jane Hillston. LFCS, University of Edinburgh.

Differential equations

Using this approach, we can validate our PEPA model against the system of ODEs indepedently derived by the biochemists:

$$\begin{array}{rcl} \frac{dm_1(t)}{dt} &=& -k_1m_1(t)m_2(t) + k_2m_3(t) + k_5m_4(t) - k_{12}m_1(t)m_{12}(t) \\ && +k_{13}m_{13}(t) + k_{14}m_{13}(t) \\ \hline \frac{dm_2(t)}{dt} &=& -k_1m_1(t)m_2(t) + k_2m_3(t) + k_{11}m_{11}(t) \\ \hline \frac{dm_3(t)}{dt} &=& k_1m_1(t)m_2(t) - k_2m_3(t) - k_3m_3(t)m_9(t) + k_4m_4(t) \\ &\vdots &\vdots \\ \hline \frac{dm_{13}(t)}{dt} &=& k_{12}m_1(t)m_12(t) - k_{13}m_{13}(t) - k_{14}m_{13}(t) \end{array}$$

Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

1

There are several advantages to be gained by introducing a process algebra model as an intermediary to the derivation of the ODEs.

Jane Hillston. LFCS, University of Edinburgh.

There are several advantages to be gained by introducing a process algebra model as an intermediary to the derivation of the ODEs.

The ODEs can be automatically generated from the descriptive process algebra model, thus reducing human error.

There are several advantages to be gained by introducing a process algebra model as an intermediary to the derivation of the ODEs.

- The ODEs can be automatically generated from the descriptive process algebra model, thus reducing human error.
- The formality of the process algebra model and its underlying semantics allow us to derive properties of the model, such as freedom from deadlock, before numerical analysis is carried out.

There are several advantages to be gained by introducing a process algebra model as an intermediary to the derivation of the ODEs.

- The ODEs can be automatically generated from the descriptive process algebra model, thus reducing human error.
- The formality of the process algebra model and its underlying semantics allow us to derive properties of the model, such as freedom from deadlock, before numerical analysis is carried out.
- The algebraic formulation of the model makes clear the interactions between the biochemical entities, or substrates. The style of modelling is descriptive, close to informal graphical representations that biochemists already use.

For a generation, performance modellers have seen their choices as being:

Closed form analytical models;

Jane Hillston. LFCS, University of Edinburgh.

For a generation, performance modellers have seen their choices as being:

- Closed form analytical models;
- Simulations; or

Jane Hillston. LFCS, University of Edinburgh.

For a generation, performance modellers have seen their choices as being:

- Closed form analytical models;
- Simulations; or
- Numerical solution of continuous time Markov chains (CTMC)

For a generation, performance modellers have seen their choices as being:

- Closed form analytical models;
- Simulations; or
- Numerical solution of continuous time Markov chains (CTMC)

The major limitations of the CTMC approach are the state space explosion problem and the reliance on exponential distributions.

New mathematical structures: differential equations

Use a more abstract state representation rather than the CTMC complete state space.

Jane Hillston. LFCS, University of Edinburgh.

New mathematical structures: differential equations

- Use a more abstract state representation rather than the CTMC complete state space.
- No longer aim to calculate the probability distribution over the entire state space of the model.

New mathematical structures: differential equations

- Use a more abstract state representation rather than the CTMC complete state space.
- No longer aim to calculate the probability distribution over the entire state space of the model.
- Assume that these state variables are subject to continuous rather than discrete change.

New mathematical structures: differential equations

- Use a more abstract state representation rather than the CTMC complete state space.
- No longer aim to calculate the probability distribution over the entire state space of the model.
- Assume that these state variables are subject to continuous rather than discrete change.

Only appropriate for some models, but results are promising in those cases.

- In a PEPA model the state at any current time is the local derivative or state of each component of the model.
- When we have large numbers of repeated components it can make sense to represent the state of the system as the count of the current number of each possible local derivative or component type.
- We can approximate the behaviour of the model by treating the number of each component type as a continuous variable, and the state of the model as a whole as the set of such variables.
- The evolution of each such variable can then be described by an ordinary differential equation (assuming rates are deterministic).

- In a PEPA model the state at any current time is the local derivative or state of each component of the model.
- When we have large numbers of repeated components it can make sense to represent the state of the system as the count of the current number of each possible local derivative or component type.
- We can approximate the behaviour of the model by treating the number of each component type as a continuous variable, and the state of the model as a whole as the set of such variables.
- The evolution of each such variable can then be described by an ordinary differential equation (assuming rates are deterministic).

- In a PEPA model the state at any current time is the local derivative or state of each component of the model.
- When we have large numbers of repeated components it can make sense to represent the state of the system as the count of the current number of each possible local derivative or component type.
- We can approximate the behaviour of the model by treating the number of each component type as a continuous variable, and the state of the model as a whole as the set of such variables.
- The evolution of each such variable can then be described by an ordinary differential equation (assuming rates are deterministic).

- In a PEPA model the state at any current time is the local derivative or state of each component of the model.
- When we have large numbers of repeated components it can make sense to represent the state of the system as the count of the current number of each possible local derivative or component type.
- We can approximate the behaviour of the model by treating the number of each component type as a continuous variable, and the state of the model as a whole as the set of such variables.
- The evolution of each such variable can then be described by an ordinary differential equation (assuming rates are deterministic).

- The PEPA definitions of the component specify the activities which can increase or decrease the number of components exhibited in the current state.
- The cooperations show when the number of instances of another component will have an influence on the evolution of this component.

- The PEPA definitions of the component specify the activities which can increase or decrease the number of components exhibited in the current state.
- The cooperations show when the number of instances of another component will have an influence on the evolution of this component.
Differential equations from arbitrary PEPA models

- The PEPA definitions of the component specify the activities which can increase or decrease the number of components exhibited in the current state.
- The cooperations show when the number of instances of another component will have an influence on the evolution of this component.

Derivation of the system of ODES representing the PEPA model then proceeds via an activity matrix in much the same way as for the high/low biochemical pathway model.

Outline

Introduction

Case Study in Systems Biology

Continuous Approximation and Differential Equations

Case Study in Web Services

Jane Hillston. LFCS, University of Edinburgh.



- 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.



- 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.



- 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.

Jane Hillston. LFCS, University of Edinburgh.



- 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.



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

Jane Hillston. LFCS, University of Edinburgh.



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

Jane Hillston. LFCS, University of Edinburgh.



- 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.



- 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.



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

Jane Hillston. LFCS, University of Edinburgh.



- 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.



- 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.



- 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.

Jane Hillston. LFCS, University of Edinburgh.



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

Jane Hillston. LFCS, University of Edinburgh.



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

Jane Hillston. LFCS, University of Edinburgh.



- 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.



- 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.



Jane Hillston. LFCS, University of Edinburgh.



The broker is inactive until it receives a request.

Jane Hillston. LFCS, University of Edinburgh.



- ► 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.



- ► 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 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.

Jane Hillston. LFCS, University of Edinburgh.



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

 $\stackrel{def}{=}$

 $\stackrel{def}{=}$

 $\stackrel{def}{=}$

 $\stackrel{def}{=}$

 $\stackrel{def}{=}$

 $\stackrel{def}{=}$

def =

 $\stackrel{def}{=}$

 $(request_{b}, \top).Broker_{dec_input}$ $(decrypt_{sp}, r_{b_dec_sp}).Broker_{enc_input}$ $(encrypt_{ws}, r_{b_enc_ws}).Broker_{sending}$ $(request_{ws}, r_{b_req}).Broker_{waiting}$ $(response_{ws}, \top).Broker_{dec_resp}$ $(decrypt_{ws}, r_{b_dec_ws}).Broker_{replying}$ $(response_{b}, r_{b_enc_sp}).Broker_{idle}$

Jane Hillston. LFCS, University of Edinburgh.



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.



- 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.



- 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.



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

Jane Hillston. LFCS, University of Edinburgh.



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.

Jane Hillston. LFCS, University of Edinburgh.



- 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.



- 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.



WS _{idle}	def =	$(request_{ws}, \top)$. $WS_{decoding}$
	+	$(invoke_{ws}, \top).WS_{method}$
WS _{decoding}	def =	$(decryptReq_{ws}, r_{ws_dec_b}).WS_{execution}$
WS _{execution}	def =	$(execute_{ws}, r_{ws_exec}).WS_{securing}$
WS _{securing}	def =	(encryptResp _{ws} , r _{ws_enc_b}).WS _{responding}
VS _{responding}	def =	(response _{ws} , r _{ws_resp_b}).WS _{idle}
WS _{method}	def =	$(execute_{ws}, r_{ws_exec}).WS_{returning}$
WS _{returning}	def =	(result _{ws} , r _{ws_res}).WS _{idle}

Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

l



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

Jane Hillston. LFCS, University of Edinburgh.



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

 $\begin{array}{l} \stackrel{\text{def}}{=} & (request_{ws}, \top).WS_{decoding} \\ + & (invoke_{ws}, \top).WS_{method} \\ \stackrel{\text{def}}{=} & (decryptReq_{ws}, r_{ws_dec_b}).WS_{execution} \\ \stackrel{\text{def}}{=} & (execute_{ws}, r_{ws_exec}).WS_{securing} \\ \stackrel{\text{def}}{=} & (encryptResp_{ws}, r_{ws_enc_b}).WS_{responding} \\ \stackrel{\text{def}}{=} & (response_{ws}, r_{ws_resp_b}).WS_{idle} \\ \stackrel{\text{def}}{=} & (execute_{ws}, r_{ws_exec}).WS_{returning} \\ \stackrel{\text{def}}{=} & (result_{ws}, r_{ws_res}).WS_{idle} \\ \end{array}$

Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

<ロト < @ ト < 注 > < 注 > 注 の < @</p>



WS _{idle}	def =	$(request_{ws}, \top)$. $WS_{decoding}$
	+	$(invoke_{ws}, \top).WS_{method}$
WS _{decoding}	def =	$(decryptReq_{ws}, r_{ws_dec_b}).WS_{execution}$
WS _{execution}	def =	$(execute_{ws}, r_{ws_exec}).WS_{securing}$
WS _{securing}	def =	(encryptResp _{ws} , r _{ws_enc_b}).WS _{responding}
VS _{responding}	def =	(response _{ws} , r _{ws_resp_b}).WS _{idle}
WS _{method}	def =	$(execute_{ws}, r_{ws_exec}).WS_{returning}$
WS _{returning}	def =	(result _{ws} , r _{ws_res}).WS _{idle}

Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

l

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.

$$System \stackrel{\text{def}}{=} (SPC_{idle} \bigotimes_{\mathcal{K}} Broker_{idle}) \bigotimes_{\mathcal{L}} (WS_{idle} \bigotimes_{\mathcal{M}} FPC_{idle})$$
where
$$\mathcal{K} = \{ request_b, response_b \}$$

$$\mathcal{L} = \{ request_{ws}, response_{ws} \}$$

$$\mathcal{M} = \{ invoke_{ws}, result_{ws} \}$$

Jane Hillston. LFCS, University of Edinburgh.
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.

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

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.

Cost of analysis

Performance models admit many different types of analysis. Some have lower evaluation cost, but are less informative, such as steady-state analysis. Others have higher evaluation cost, but are more informative, such as transient analysis.

Cost of analysis

- Performance models admit many different types of analysis. Some have lower evaluation cost, but are less informative, such as steady-state analysis. Others have higher evaluation cost, but are more informative, such as transient analysis.
- We compare ODE-based evaluation against other techniques which could be used to analyse the model.

Cost of analysis

- Performance models admit many different types of analysis. Some have lower evaluation cost, but are less informative, such as steady-state analysis. Others have higher evaluation cost, but are more informative, such as transient analysis.
- We compare ODE-based evaluation against other techniques which could be used to analyse the model.
- We compare against steady-state and transient analysis 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.

Comparison of analysis types

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

Jane Hillston. LFCS, University of Edinburgh.

Comparison of analysis types

- We report only a single run of the transient analysis and simulation. In practice, due to the stochastic nature of the analyses, these 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.

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

Jane Hillston. LFCS, University of Edinburgh.

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

Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●

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

Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

▲ロト ▲母 ▶ ▲目 ▶ ▲目 ▶ ▲ ● ◆ ○ ◆ ○ ◆

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

Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

▲□▶ ▲□▶ ▲□▶ ▲□▶ ▲□ ● ● ●

	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
I	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
I	100	100	100	100	-	_	-	-	_	2.78	2.78

Jane Hillston. LFCS, University of Edinburgh.

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

Jane Hillston. LFCS, University of Edinburgh.

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

Jane Hillston. LFCS, University of Edinburgh.

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

Jane Hillston. LFCS, University of Edinburgh.

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.

Jane Hillston. LFCS, University of Edinburgh.

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.

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.
- The graphs show fluctuations in the numbers of components with respect to time from t = 0 to t = 100 for estimated values of rates for the activities of the system. We can observe an initial flurry of activity until the system stabilises into its steady-state equilibrium at time (around) t = 50.

Second party clients



Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

<ロト < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □

Brokers



Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

<ロ> <0</p>

First party clients



Jane Hillston. LFCS, University of Edinburgh.

Web service



Jane Hillston. LFCS, University of Edinburgh.

Adventures in Systems Biology

▲□▶ ▲圖▶ ▲≧▶ ▲≧▶ _ 差 _ 釣�()~

Comparison with Continuous-Time Markov Chain solution

For a system without blocking (1000 instances of all components) we compared the steady-state values of state variables with values derived from the the steady-state probabilities computed by the PEPA Workbench for the model with one instance of each component.

Comparison with Continuous-Time Markov Chain solution

- For a system without blocking (1000 instances of all components) we compared the steady-state values of state variables with values derived from the the steady-state probabilities computed by the PEPA Workbench for the model with one instance of each component.
 - The steady state probabilities of the PEPA Workbench were scaled by 1000.

Comparison with Continuous-Time Markov Chain solution

- For a system without blocking (1000 instances of all components) we compared the steady-state values of state variables with values derived from the the steady-state probabilities computed by the PEPA Workbench for the model with one instance of each component.
 - The steady state probabilities of the PEPA Workbench were scaled by 1000.
- We found good agreement with the results obtained by ODE solution after the system stabilises into its steady-state equilibrium.

Outline

Introduction

Case Study in Systems Biology

Continuous Approximation and Differential Equations

Case Study in Web Services

Jane Hillston. LFCS, University of Edinburgh.

The excursion into systems biology has been very worthwhile:

Jane Hillston. LFCS, University of Edinburgh.

The excursion into systems biology has been very worthwhile:

the reaction of the biochemists from the Beatson Institute have been very positive:

Jane Hillston. LFCS, University of Edinburgh.

The excursion into systems biology has been very worthwhile:

- the reaction of the biochemists from the Beatson Institute have been very positive:
 - the presentation of the system in two contrasting styles, which can be proved to be equivalent;

Jane Hillston. LFCS, University of Edinburgh.

The excursion into systems biology has been very worthwhile:

- the reaction of the biochemists from the Beatson Institute have been very positive:
 - the presentation of the system in two contrasting styles, which can be proved to be equivalent;
 - the quantified reasoning at the Markov chain level even with crude approximation of concentrations – giving estimations of transient properties such as throughput over time, and first passage time distributions.

The excursion into systems biology has been very worthwhile:

- the reaction of the biochemists from the Beatson Institute have been very positive:
 - the presentation of the system in two contrasting styles, which can be proved to be equivalent;
 - the quantified reasoning at the Markov chain level even with crude approximation of concentrations – giving estimations of transient properties such as throughput over time, and first passage time distributions.
- the derivation of differential equations appears to offer an interesting alternative to existing modelling approaches to performance evaluation of large scale models.