(日) (同) (三) (三)

2

Stochastic Modelling Part I: Foundations

Stephen Gilmore LFCS, University of Edinburgh

GLOBAN Summerschool Lygby, Denmark

25th August 2006

Stephen Gilmore. LFCS, University of Edinburgh.

Global Computing

What distinguishes global computing from local computing?

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations ・ロト・御 ト・ 画・ ・ 画・ うらの

イロト イヨト イヨト イヨト

-큰

Global Computing: Some differences

Physical distance

Network latency

Stephen Gilmore. LFCS, University of Edinburgh.

Global Computing: Some differences

Physical distance

- Network latency
- Partial failures
 - Server may be down
 - Routers may be down

Stephen Gilmore. LFCS, University of Edinburgh.

Global Computing: Some differences

Physical distance

- Network latency
- Partial failures
 - Server may be down
 - Routers may be down
- Scale
 - Workload characterisation

Stephen Gilmore. LFCS, University of Edinburgh.

Global Computing: Some differences

Physical distance

- Network latency
- Partial failures
 - Server may be down
 - Routers may be down
- Scale
 - Workload characterisation
- Resource sharing
 - Network contention
 - CPU load

· < /⊒ > < ∃ > <

- Physical distance need to represent time
 - Network latency
- Partial failures
 - Server may be down
 - Routers may be down
- Scale
 - Workload characterisation
- Resource sharing
 - Network contention
 - CPU load

(日) (同) (三) (三)

- Physical distance need to represent time
 - Network latency
- Partial failures randomness and probability
 - Server may be down
 - Routers may be down
- Scale
 - Workload characterisation
- Resource sharing
 - Network contention
 - CPU load

(日) (同) (三) (三)

- Physical distance need to represent time
 - Network latency
- Partial failures randomness and probability
 - Server may be down
 - Routers may be down
- Scale need to quantify population sizes
 - Workload characterisation
- Resource sharing
 - Network contention
 - CPU load

- Physical distance need to represent time
 - Network latency
- Partial failures randomness and probability
 - Server may be down
 - Routers may be down
- Scale need to quantify population sizes
 - Workload characterisation
- Resource sharing need to express percentages
 - Network contention
 - CPU load

Modelling Global Computing: The challenges

Time What representation of time will we use?

Randomness What kind of random number distributions will we use?

- Probability How can we have probabilities in the model without uncertainty in the results?
 - Scale How can we escape the state-space explosion problem?

Percentages What can it mean to have a fraction of a process?

Stephen Gilmore. LFCS, University of Edinburgh.



Quality of Service issues

Can the server maintain reasonable response times?

2

Stephen Gilmore. LFCS, University of Edinburgh.



Scalability issues

How many times do we have to replicate this service to support all of the subscribers?

Stephen Gilmore. LFCS, University of Edinburgh.



Scalability issues

Will the server withstand a distributed denial of service attack?

2

Stephen Gilmore. LFCS, University of Edinburgh.



Service-level agreements

What percentage of downloads do complete within the time we advertised?

Stephen Gilmore. LFCS, University of Edinburgh.

Outline

Introduction

Interplay: Process Algebra and Markov Process

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

Outline

Introduction

Interplay: Process Algebra and Markov Process

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations ・ロト ・回ト ・ ヨト ・

2

Performance Modelling using CTMC



Stephen Gilmore. LFCS, University of Edinburgh.



A negative exponentially distributed duration is associated with each transition.

イロト イヨト イヨト イヨト

2

Performance Modelling using CTMC



these parameters form the entries of the infinitesimal generator matrix Q

Stephen Gilmore. LFCS, University of Edinburgh.



Image: A math a math

In steady state the probability flux out of a state is balanced by the flux in.

Stephen Gilmore. LFCS, University of Edinburgh.



"Global balance equations" captured by $\pi Q = 0$ solved by linear algebra

Stephen Gilmore. LFCS, University of Edinburgh.



▲ロ ▶ ▲ 圖 ▶ ▲ 圖 ▶ ▲ 圖 ▶ ▲ 国 ▶ ▲ 国 ▶

Stephen Gilmore. LFCS, University of Edinburgh.



Stephen Gilmore. LFCS, University of Edinburgh.

<ロト </p>

Performance Modelling using CTMC



Stephen Gilmore. LFCS, University of Edinburgh.

Outline

Introduction

Interplay: Process Algebra and Markov Process

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

Stochastic process algebras

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

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

・ロト ・ 日 ト ・ ヨ ト ・ ヨ ト

4

SPA Languages

SPA

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-큰

SPA Languages



Stephen Gilmore. LFCS, University of Edinburgh.

• • • • • • • • • • • •

2

SPA Languages



Stephen Gilmore. LFCS, University of Edinburgh.

- * 伊 * * 注 * *

SPA Languages



Stephen Gilmore. LFCS, University of Edinburgh.

SPA Languages



Stephen Gilmore. LFCS, University of Edinburgh.

SPA Languages



Stephen Gilmore. LFCS, University of Edinburgh.

SPA Languages



Stephen Gilmore. LFCS, University of Edinburgh.

SPA Languages



Stephen Gilmore. LFCS, University of Edinburgh.

SPA Languages



Stephen Gilmore. LFCS, University of Edinburgh.
SPA Languages



Stephen Gilmore. LFCS, University of Edinburgh.

イロン イ理と イヨン ・ ヨン・

-2

The Importance of Being Exponential



Stephen Gilmore. LFCS, University of Edinburgh.

The Importance of Being Exponential



The memoryless property of the negative exponential distribution means that residual times do not need to be recorded.

(日) (同) (三) (三)

The exponential distribution and the expansion law

We retain the expansion law of classical process algebra:

$$\begin{aligned} (\alpha, r).Stop \parallel (\beta, s).Stop = \\ (\alpha, r).(\beta, s).(Stop \parallel Stop) + (\beta, s).(\alpha, r).(Stop \parallel Stop) \end{aligned}$$

only if the negative exponential distribution is used.

< ロ > < 同 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ >

3

Performance Evaluation Process Algebra

 Models are constructed from components which engage in activities.



Stephen Gilmore. LFCS, University of Edinburgh.

< ロ > < 同 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ >

3

Performance Evaluation Process Algebra

 Models are constructed from components which engage in activities.



Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-큰

Performance Evaluation Process Algebra

 Models are constructed from components which engage in activities.



Stephen Gilmore. LFCS, University of Edinburgh.

< ロ > < 同 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ >

3

Performance Evaluation Process Algebra

 Models are constructed from components which engage in activities.



Stephen Gilmore. LFCS, University of Edinburgh.

Performance Evaluation Process Algebra

 Models are constructed from components which engage in activities.



The language is used to generate a CTMC for performance modelling.

(日) (同) (三) (三)

3

Performance Evaluation Process Algebra

 Models are constructed from components which engage in activities.



The language is used to generate a CTMC for performance modelling.

PEPA MODEL

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

3

Performance Evaluation Process Algebra

Models are constructed from components which engage in activities.



The language is used to generate a CTMC for performance modelling.

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

3

Performance Evaluation Process Algebra

 Models are constructed from components which engage in activities.



The language is used to generate a CTMC for performance modelling.



Stephen Gilmore. LFCS, University of Edinburgh.

Performance Evaluation Process Algebra

Models are constructed from components which engage in activities.



The language is used to generate a CTMC for performance modelling.



Stephen Gilmore. LFCS, University of Edinburgh.

Performance Evaluation Process Algebra

Models are constructed from components which engage in activities.



The language is used to generate a CTMC for performance modelling.

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-2

PEPA

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

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

Stephen Gilmore. 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

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations ・ロト ・ 日 ト ・ ヨ ト ・ ヨ ト

-

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

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-2

PEPA

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

PREFIX:	$(\alpha, r).S$	designated first action
CHOICE:	S + S	competing components
CONSTANT:	$A \stackrel{{}_{\scriptscriptstyle def}}{=} S$	assigning names

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-큰

PEPA

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

- PREFIX: $(\alpha, r).S$ defCHOICE:S + ScoCONSTANT: $A \stackrel{\text{def}}{=} S$ asCOOPERATION: $P \bowtie_L P$ α
- r).Sdesignated first actionScompeting components $\stackrel{f}{=} S$ assigning names $\stackrel{f}{=} P$ $\alpha \notin L$ individual actions
 - $\alpha \in \mathbf{L}$ shared actions

Stephen Gilmore. LFCS, University of Edinburgh.

▲□▶ ▲圖▶ ▲圖▶ ▲圖▶

æ

PEPA

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

PREFIX:	$(\alpha, r).S$	designated first action
CHOICE:	S + S	competing components
CONSTANT:	$A \stackrel{{}_{\scriptscriptstyle def}}{=} S$	assigning names
COOPERATION:	$P \bowtie_{L} P$	$\alpha \notin \mathbf{L}$ individual actions
	-	$lpha \in \mathcal{L}$ shared actions
HIDING:	P/L	abstraction $\alpha \in \mathbf{L} \Rightarrow \alpha \to \tau$

Stephen Gilmore. LFCS, University of Edinburgh.

· < /⊒ > < ∃ > <

Interplay between process algebra and Markov process

The theoretical development underpinning PEPA has focused on the interplay between the process algebra and the underlying mathematical structure, the Markov process.

Interplay between process algebra and Markov process

- The theoretical development underpinning PEPA has focused on the interplay between the process algebra and the underlying mathematical structure, the Markov process.
- From the process algebra side the Markov chain had a profound influence on the design of the language and in particular on the interactions between components.

Interplay between process algebra and Markov process

- The theoretical development underpinning PEPA has focused on the interplay between the process algebra and the underlying mathematical structure, the Markov process.
- From the process algebra side the Markov chain had a profound influence on the design of the language and in particular on the interactions between components.
- From the Markov chain perspective the process algebra structure has been exploited to find aspects of independence even between interacting components.

Example: Browsers, server and download

$$\textit{Server} ~~ \stackrel{\textit{\tiny def}}{=} ~~ (\textit{get}, \top).(\textit{download}, \mu).(\textit{rel}, \top).\textit{Server}$$

WEB
$$\stackrel{\text{def}}{=} (Browser \parallel Browser) \bowtie_{l} Server$$

where $L = \{get, download, rel\}$

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-큰

Integrated analysis

Qualitative verification can now be complemented by quantitative verification.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

Integrated analysis: Reachability analysis

Reachability analysis

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



Stephen Gilmore. LFCS, University of Edinburgh.

Integrated analysis: Specification matching

Specification matching

With what probability

does system behaviour match its specification?



A 🖓 h

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

Integrated analysis: Specification matching

Specification matching

Does the "frequency profile" of the system match that of the specification?



Image: A matrix A

Stephen Gilmore. LFCS, University of Edinburgh.

Integrated analysis: Model checking

Model checking

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



Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

Integrated analysis: Model checking

Model checking

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



Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

イロト イヨト イヨト イヨト

-큰

Parallel Composition

 Parellel composition is the basis of the compositionality in a process algebra

Stephen Gilmore. LFCS, University of Edinburgh.

Parallel Composition

Parellel composition is the basis of the compositionality in a process algebra — it defines which components interact and how.

(日) (周) (三) (三)

Parallel Composition

- Parellel composition is the basis of the compositionality in a process algebra — it defines which components interact and how.
- In classical process algebra is it often associated with communication.

Parallel Composition

- Parellel composition is the basis of the compositionality in a process algebra — it defines which components interact and how.
- In classical process algebra is it often associated with communication.
- When the activities of the process algebra have a duration the definition of parallel composition becomes more complex.

Who Synchronises...?

Even within classical process algebras there is variation in the interpretation of parallel composition:

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

▲ @ ▶ ▲ ∃ ▶ ▲

Who Synchronises...?

Even within classical process algebras there is variation in the interpretation of parallel composition:

CCS-style

- Actions are partitioned into input and output pairs.
- Communication or synchronisation takes places between conjugate pairs.
- The resulting action has silent type \(\tau\).
Even within classical process algebras there is variation in the interpretation of parallel composition:

CCS-style

- Actions are partitioned into input and output pairs.
- Communication or synchronisation takes places between conjugate pairs.
- The resulting action has silent type \(\tau\).

Even within classical process algebras there is variation in the interpretation of parallel composition:

CCS-style

- Actions are partitioned into input and output pairs.
- Communication or synchronisation takes places between conjugate pairs.
- The resulting action has silent type τ.

Even within classical process algebras there is variation in the interpretation of parallel composition:

CCS-style

- Actions are partitioned into input and output pairs.
- Communication or synchronisation takes places between conjugate pairs.
- The resulting action has silent type τ.

CSP-style

- No distinction between input and output actions.
- Communication or synchronisation takes place on the basis of shared names.
- The resulting action has the same name.

Even within classical process algebras there is variation in the interpretation of parallel composition:

CCS-style

- Actions are partitioned into input and output pairs.
- Communication or synchronisation takes places between conjugate pairs.
- The resulting action has silent type τ.

CSP-style

- No distinction between input and output actions.
- Communication or synchronisation takes place on the basis of shared names.
- The resulting action has the same name.

Even within classical process algebras there is variation in the interpretation of parallel composition:

CCS-style

- Actions are partitioned into input and output pairs.
- Communication or synchronisation takes places between conjugate pairs.
- The resulting action has silent type τ.

CSP-style

- No distinction between input and output actions.
- Communication or synchronisation takes place on the basis of shared names.
- The resulting action has the same name.

Even within classical process algebras there is variation in the interpretation of parallel composition:

CCS-style

- Actions are partitioned into input and output pairs.
- Communication or synchronisation takes places between conjugate pairs.
- The resulting action has silent type τ.

CSP-style

- No distinction between input and output actions.
- Communication or synchronisation takes place on the basis of shared names.
- The resulting action has the same name.

Most stochastic process algebras adopt CSP-style synchronisation.

-2

Timed Synchronisation

The issue of what it means for two timed activities to synchronise is a vexed one....

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

・ロト ・ 日 ト ・ ヨ ト ・ ヨ ト

-큰

Timed Synchronisation



Stephen Gilmore. LFCS, University of Edinburgh.

・ロト ・聞ト ・ヨト ・ヨト

-큰

Timed Synchronisation



Barrier Synchronisation

Stephen Gilmore. LFCS, University of Edinburgh.

・ロト ・ 日 ト ・ ヨ ト ・ ヨ ト

4

Timed Synchronisation



Stephen Gilmore. LFCS, University of Edinburgh.

Image: A math a math

2

-

Timed Synchronisation



s is no longer exponentially distributed

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

-큰

Timed Synchronisation



algebraic considerations limit choices

Stephen Gilmore. LFCS, University of Edinburgh.

・ロト ・聞ト ・ヨト ・ヨト

-큰

Timed Synchronisation



TIPP: new rate is product of individual rates

Stephen Gilmore. LFCS, University of Edinburgh.

▲□▶ ▲圖▶ ▲圖▶ ▲圖▶

4

Timed Synchronisation



Stephen Gilmore. LFCS, University of Edinburgh.

* ロ > * 個 > * 注 > * 注 >

-큰

Timed Synchronisation



EMPA: one participant is passive

Stephen Gilmore. LFCS, University of Edinburgh.

4

Timed Synchronisation



Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

2

Timed Synchronisation



bounded capacity: new rate is the minimum of the rates

Stephen Gilmore. LFCS, University of Edinburgh.

Timed Synchronisation



Stephen Gilmore. LFCS, University of Edinburgh.

Cooperation in PEPA

In PEPA each component has a bounded capacity to carry out activities of any particular type, determined by the apparent rate for that type.

- 4 @ > - 4 @ > - 4 @ >

Cooperation in PEPA

- In PEPA each component has a bounded capacity to carry out activities of any particular type, determined by the apparent rate for that type.
- Synchronisation, or cooperation cannot make a component exceed its bounded capacity.

Cooperation in PEPA

- In PEPA each component has a bounded capacity to carry out activities of any particular type, determined by the apparent rate for that type.
- Synchronisation, or cooperation cannot make a component exceed its bounded capacity.
- Thus the apparent rate of a cooperation is the minimum of the apparent rates of the co-operands.

(日) (同) (三) (三)

2

Equivalence Relations

In process algebra equivalence relations are defined based on the notion of observability.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

-큰

Equivalence Relations



Stephen Gilmore. LFCS, University of Edinburgh.

-큰

Equivalence Relations



Stephen Gilmore. LFCS, University of Edinburgh.

-큰

Equivalence Relations



Stephen Gilmore. LFCS, University of Edinburgh.

Equivalence Relations

In PEPA observation is assumed to include the ability to record timing information over a number of runs.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

-큰

Equivalence Relations



Stephen Gilmore. LFCS, University of Edinburgh.

-큰

Equivalence Relations



Stephen Gilmore. LFCS, University of Edinburgh.

-큰

Equivalence Relations



Stephen Gilmore. LFCS, University of Edinburgh.

-큰

Equivalence Relations



Stephen Gilmore. LFCS, University of Edinburgh.

Equivalence Relations

The resulting equivalence relation is a **bisimulation** in the style of Larsen and Skou, and coincides with the Markov process notion of **lumpability**.

Aggregation and lumpability

Model aggregation: use a state-state equivalence to establish a partition of the state space of a model, and replace each set of states by one macro-state, i.e. take a different stochastic representation of the same model.

- **(())) (())**

Aggregation and lumpability

Model aggregation: use a state-state equivalence to establish a partition of the state space of a model, and replace each set of states by one macro-state, i.e. take a different stochastic representation of the same model. A lumpable partition is the only partition of a Markov process which preserves the Markov property.

Aggregation and lumpability



▲□▶ ▲□▶ ▲目▶ ▲目▶ 三回 - のへで

Stephen Gilmore. LFCS, University of Edinburgh.

Aggregation and lumpability



Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part I: Foundations

* ロ > * 個 > * 注 > * 注 >

2

Aggregation and lumpability



Stephen Gilmore. LFCS, University of Edinburgh.
(日) (同) (三) (三)

3

Stochastic Modelling Part II: Markovian methods

Stephen Gilmore LFCS, University of Edinburgh

GLOBAN Summerschool Lygby, Denmark

25th August 2006

Stephen Gilmore. LFCS, University of Edinburgh.

Outline

A modelling language

A semantics for the modelling language

Tools for the modelling language

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part II: Markovian methods ▲口 > ▲母 > ▲目 > ▲目 > ▲目 > ▲日 > ④ < ⊙

Outline

A modelling language

A semantics for the modelling language

Tools for the modelling language

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part II: Markovian methods ◆ロ > ◆母 > ◆臣 > ◆臣 > 「臣 」のへで

Performance Evaluation Process Algebra

- PEPA (Performance Evaluation Process Algebra) is a high-level modelling language for distributed systems. It can be used to develop models of existing systems (abstraction) or designs for proposed ones (specification).
- PEPA can capture performance information in a process algebra setting. It is a stochastic process algebra.
- The definitive reference for PEPA is A Compositional Approach to Performance Modelling, Jane Hillston, Cambridge University Press, 1996.

Strengths of stochastic process algebras

SPAs have strengths in the areas of semantic definition, inherent compositionality and the existence of important equivalence relations (including bisimulation). This relation provides the basis for aggregation of PEPA models.

Terminology

The components in a PEPA model engage, cooperatively or individually, in activities. Each activity has an action type which corresponds to the actions of the system being modelled.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part II: Markovian methods ・ロト ・聞 ・ ・目 ・ ・ 目・ ・ の へ ()・

Terminology

The components in a PEPA model engage, cooperatively or individually, in activities. Each activity has an action type which corresponds to the actions of the system being modelled.

To represent unimportant or unknown actions there is a distinguished action type, τ .

▲ @ ▶ ▲ @ ▶ ▲

Quantitative aspects

Every activity in PEPA has an associated activity rate which may be any positive real number, or the distinguished symbol " \top ", meaning unspecified, read as 'top'.

Quantitative aspects

Every activity in PEPA has an associated activity rate which may be any positive real number, or the distinguished symbol " \top ", meaning unspecified, read as 'top'.

Components and activities are primitives. PEPA also provides a small set of combinators.

4

PEPA syntax

S	::=	(α, r).S	(prefix)
		$S_1 + S_2$	(choice)
		X	(variable)
С	::=	$C_1 \stackrel{\bowtie}{\underset{L}{\bowtie}} C_2$	(cooperation)
		C / L	(hiding)
		5	(sequential)

Stephen Gilmore. LFCS, University of Edinburgh.

PEPA: informal semantics (sequential sublanguage)

$(\alpha, r).S$

The activity (α, r) takes time Δt (drawn from the exponential distribution with parameter r).

$S_1 + S_2$

In this choice either S_1 or S_2 will complete an activity first. The other is discarded.

(日) (同) (三) (三)

PEPA: informal semantics (combinators)

$\begin{array}{ccc} C_1 & \bigvee_L & C_2 \\ & & \text{All activities of } C_1 \text{ and } C_2 \text{ with types in } L \text{ are shared: others remain individual.} \\ & & \text{NOTATION: write } C_1 \parallel C_2 \text{ if } L \text{ is empty.} \end{array}$

C / LActivities of C with types in L are hidden (τ type activities) to be thought of as internal delays.

イロト イヨト イヨト イヨト

-큰

Example: M/M/1/N/N queue

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

æ

Example: M/M/1/N/N queue



 $Queue_i \equiv Arrival_i \bigotimes_{\{serve\}} Server$

Stephen Gilmore. LFCS, University of Edinburgh.

-큰

Expansion Law

$$P \bowtie_{L} Q =$$

Stephen Gilmore. LFCS, University of Edinburgh.

4

Expansion Law

- - -

$$P \bowtie_{L} Q =$$

$$\sum \{ (\alpha, r) \cdot (P' \bowtie_{L} Q) \colon P \xrightarrow{(\alpha, r)} P'; \alpha \notin L \} +$$

Stephen Gilmore. LFCS, University of Edinburgh.

-큰

Expansion Law

$$P \bowtie_{L} Q =$$

$$\sum \{ (\alpha, r) . (P' \bowtie_{L} Q) : P \xrightarrow{(\alpha, r)} P'; \alpha \notin L \} +$$

$$\sum \{ (\alpha, r) . (P \bowtie_{L} Q') : Q \xrightarrow{(\alpha, r)} Q'; \alpha \notin L \} +$$

Stephen Gilmore. LFCS, University of Edinburgh.

<ロ> (日) (日) (日) (日) (日)

-큰

Expansion Law

$$P \bowtie_{L} Q =$$

$$\sum \{ (\alpha, r) \cdot (P' \bowtie_{L} Q) : P \xrightarrow{(\alpha, r)} P'; \alpha \notin L \} +$$

$$\sum \{ (\alpha, r) \cdot (P \bowtie_{L} Q') : Q \xrightarrow{(\alpha, r)} Q'; \alpha \notin L \} +$$

$$\sum \{ (\alpha, r) \cdot (P' \bowtie_{L} Q') :$$

$$P \xrightarrow{(\alpha, r_{1})} P'; Q \xrightarrow{(\alpha, r_{2})} Q'; \alpha \in L \}$$

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

Synchronisation

What should be the impact of synchronisation on rate? There are many possibilities.

 Restrict synchronisations to have one active partner and one passive partner.

Synchronisation

What should be the impact of synchronisation on rate? There are many possibilities.

- Restrict synchronisations to have one active partner and one passive partner.
- Choose a function which satisfies a small number of algebraic properties.

Synchronisation

What should be the impact of synchronisation on rate? There are many possibilities.

- Restrict synchronisations to have one active partner and one passive partner.
- Choose a function which satisfies a small number of algebraic properties.
- Have the rate limited by the slowest participant in terms of apparent rate. This is the approach adopted by PEPA.

Bounded capacity

Within the cooperation framework, PEPA assumes bounded capacity: that is, a component cannot be made to perform an activity faster by cooperation, so the rate of a shared activity is the minimum of the apparent rates of the activity in the cooperating components.

Apparent rate

The total capacity of a component P to carry out activities of type α is termed the apparent rate of α in P, denoted $r_{\alpha}(P)$.

Apparent rate

The total capacity of a component P to carry out activities of type α is termed the apparent rate of α in P, denoted $r_{\alpha}(P)$.

It is used heavily when calculating the pairwise cooperation rate: when cooperating with another component, the bounded capacity principle ensures that the overall rate of cooperation does not exceed either of the consistuent apparent rates.

Apparent rate: definition

To summarise the original ruleset from [Hillston 96], the apparent rate function can be defined as:

$$r_{\alpha}(P) = \sum_{\substack{P \xrightarrow{(\alpha,\lambda_i)}}} \lambda_i \tag{1}$$

where $\lambda_i \in \mathbb{R}^+ \cup \{n \top \mid n \in \mathbb{Q}, n > 0\}$, $n \top$ is shorthand for $n \times \top$ and \top represents the passive action rate that inherits the rate of the coaction from the cooperating component.

(日) (同) (三) (三)

æ

Properties of \top (the "unspecified" symbol)

\top requires the following arithmetic rules:

$$m\top < n\top \quad : \quad \text{for } m < n \text{ and } m, n \in \mathbb{Q}$$
$$r < n\top \quad : \quad \text{for all } r \in \mathbb{R}, n \in \mathbb{Q}$$
$$m\top + n\top = (m+n)\top \quad : \quad m, n \in \mathbb{Q}$$
$$\frac{m\top}{n\top} = \frac{m}{n} \quad : \quad m, n \in \mathbb{Q}$$

Stephen Gilmore. LFCS, University of Edinburgh.

< ロ > < 同 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ >

Properties of \top (the "unspecified" symbol)

 \top requires the following arithmetic rules:

$$\begin{array}{rcl} m\top < n\top & : & \text{for } m < n \text{ and } m, n \in \mathbb{Q} \\ & r < n\top & : & \text{for all } r \in \mathbb{R}, n \in \mathbb{Q} \\ m\top + n\top = (m+n)\top & : & m, n \in \mathbb{Q} \\ & & \frac{m\top}{n\top} = \frac{m}{n} & : & m, n \in \mathbb{Q} \end{array}$$

Note that $(r + n\top)$ is undefined for all $r \in \mathbb{R}$ in PEPA therefore disallowing components which enable both active and passive actions in the same action type at the same time, e.g. $(\alpha, \lambda).P + (\alpha, \top).P'$.

<ロト </p>

2

Outline

A modelling language

A semantics for the modelling language

Tools for the modelling language

To apply probability theory to the process under study, we view it as a random experiment.

- To apply probability theory to the process under study, we view it as a random experiment.
- The sample space of a random experiment is the set of all individual outcomes of the experiment.

- To apply probability theory to the process under study, we view it as a random experiment.
- The sample space of a random experiment is the set of all individual outcomes of the experiment.
- These individual outcomes are also called sample points or elementary events.

- To apply probability theory to the process under study, we view it as a random experiment.
- The sample space of a random experiment is the set of all individual outcomes of the experiment.
- These individual outcomes are also called sample points or elementary events.
- An event is a subset of a sample space.

Random variables

We are interested in the dynamics of a system as events happen over time. A function which associates a (real-valued) number with the outcome of an experiment is known as a random variable. Formally, a random variable X is a real-valued function defined on a sample space Ω .

Measurable functions

If X is a random variable, and x is a real number, we write $X \leq x$ for the event

$$\{ \omega : \omega \in \Omega \text{ and } X(\omega) \leq x \}$$

(日) (同) (三) (三)

-큰

Measurable functions

If X is a random variable, and x is a real number, we write $X \le x$ for the event

$$\{ \omega : \omega \in \Omega \text{ and } X(\omega) \leq x \}$$

and we write X = x for the event

$$\{\,\omega:\omega\in\Omega \text{ and } X(\omega)=x\,\}$$

Measurable functions

If X is a random variable, and x is a real number, we write $X \le x$ for the event

$$\{\,\omega:\omega\in\Omega\,\, ext{and}\,\,X(\omega)\leq x\,\}$$

and we write X = x for the event

$$\{\,\omega:\omega\in\Omega \text{ and } X(\omega)=x\,\}$$

Another property required of a random variable is that the set $X \le x$ is an event for each real x. This is necessary so that probability calculations can be made. A function having this property is said to be a measurable function or measurable in the Borel sense.
イロト イヨト イヨト イヨト

-큰

Distribution function

For each random variable X we define its distribution function F for each real x by

$F(x) = \Pr[X \le x]$

Distribution function

For each random variable X we define its distribution function F for each real x by

$$F(x) = \Pr[X \le x]$$

We associate another function $p(\cdot)$, called the probability mass function of X (pmf), for each real x:

$$p(x) = \Pr[X = x]$$

Stephen Gilmore. LFCS, University of Edinburgh.

Continuous random variables

A random variable X is continuous if p(x) = 0 for all real x.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part II: Markovian methods ▲ロト▲聞と▲臣と▲臣と 臣 のべの

Continuous random variables

A random variable X is continuous if p(x) = 0 for all real x.

(If X is a *continuous* random variable, then X can assume infinitely many values, and so it is reasonable that the probability of its assuming any *specific* value we choose beforehand is zero.)

Continuous random variables

A random variable X is continuous if p(x) = 0 for all real x.

(If X is a *continuous* random variable, then X can assume infinitely many values, and so it is reasonable that the probability of its assuming any *specific* value we choose beforehand is zero.)

The distribution function for a continuous random variable is a continuous function in the usual sense.

Exponential random variables, distribution function

The random variable X is said to be an *exponential random* variable with parameter λ ($\lambda > 0$) or to have an exponential distribution with parameter λ if it has the distribution function

$$F(x) = \begin{cases} 1 - e^{-\lambda x} & \text{for } x > 0 \\ 0 & \text{for } x \le 0 \end{cases}$$

Stephen Gilmore. LFCS, University of Edinburgh.

Exponential random variables, distribution function

The random variable X is said to be an *exponential random* variable with parameter λ ($\lambda > 0$) or to have an exponential distribution with parameter λ if it has the distribution function

$$F(x) = \left\{ egin{array}{cc} 1-e^{-\lambda x} & ext{for } x>0 \ 0 & ext{for } x\leq 0 \end{array}
ight.$$

Some authors call this distribution the negative exponential distribution.

(日) (同) (三) (三)

3

Exponential random variables, density function

The density function f = dF/dx is given by

$$f(x) = \begin{cases} \lambda e^{-\lambda x} & \text{if } x > 0\\ 0 & \text{if } x \le 0 \end{cases}$$

Stephen Gilmore. LFCS, University of Edinburgh.

< ロ > < 同 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ >

3

Mean, or expected value

If X is a continuous random variable with density function $f(\cdot)$, we define the mean or expected value of X, $\mu = E[X]$ by

$$\mu = E[X] = \int_{-\infty}^{\infty} xf(x)dx$$

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

Mean, or expected value, of the exponential distribution

Suppose X has an exponential distribution with parameter $\lambda > 0$. Then

$$\mu = E[X] = \int_{-\infty}^{\infty} x \lambda e^{-\lambda x} dx$$

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

Mean, or expected value, of the exponential distribution

Suppose X has an exponential distribution with parameter $\lambda > 0$. Then

$$\mu = E[X] = \int_{-\infty}^{\infty} x \lambda e^{-\lambda x} dx = \frac{1}{\lambda}$$

Stephen Gilmore. LFCS, University of Edinburgh.

Exponential inter-event time distribution

The time interval between successive events can also be deduced. Let F(t) be the distribution function of T, the time between events. Consider Pr(T > t) = 1 - F(t):

Pr(T > t) = Pr(No events in an interval of length t)

Exponential inter-event time distribution

The time interval between successive events can also be deduced. Let F(t) be the distribution function of T, the time between events. Consider Pr(T > t) = 1 - F(t):

Pr(T > t) = Pr(No events in an interval of length t)= 1 - F(t)

Stephen Gilmore. LFCS, University of Edinburgh.

Exponential inter-event time distribution

The time interval between successive events can also be deduced. Let F(t) be the distribution function of T, the time between events. Consider Pr(T > t) = 1 - F(t):

Pr(T > t) = Pr(No events in an interval of length t)= 1 - F(t) $= 1 - (1 - e^{-\lambda t})$

Stephen Gilmore. LFCS, University of Edinburgh.

Exponential inter-event time distribution

The time interval between successive events can also be deduced. Let F(t) be the distribution function of T, the time between events. Consider Pr(T > t) = 1 - F(t):

Pr(T > t) = Pr(No events in an interval of length t)= 1 - F(t) $= 1 - (1 - e^{-\lambda t})$ $= e^{-\lambda t}$

Stephen Gilmore. LFCS, University of Edinburgh.

The memoryless property of the exponential distribution is so called because the time to the next event is independent of when the last event occurred.

Suppose that the last event was at time 0. What is the probability that the next event will be after t + s, given that time t has elapsed since the last event, and no events have occurred?

Suppose that the last event was at time 0. What is the probability that the next event will be after t + s, given that time t has elapsed since the last event, and no events have occurred?

$$\Pr(T > t + s \mid T > t) = \frac{\Pr(T > t + s \text{ and } T > t)}{\Pr(T > t)}$$

Suppose that the last event was at time 0. What is the probability that the next event will be after t + s, given that time t has elapsed since the last event, and no events have occurred?

$$Pr(T > t + s | T > t) = \frac{Pr(T > t + s \text{ and } T > t)}{Pr(T > t)}$$
$$= \frac{e^{-\lambda(t+s)}}{e^{-\lambda t}}$$

Suppose that the last event was at time 0. What is the probability that the next event will be after t + s, given that time t has elapsed since the last event, and no events have occurred?

$$Pr(T > t + s | T > t) = \frac{Pr(T > t + s \text{ and } T > t)}{Pr(T > t)}$$
$$= \frac{e^{-\lambda(t+s)}}{e^{-\lambda t}}$$
$$= e^{-\lambda s}$$

Stephen Gilmore. LFCS, University of Edinburgh.

Suppose that the last event was at time 0. What is the probability that the next event will be after t + s, given that time t has elapsed since the last event, and no events have occurred?

$$Pr(T > t + s | T > t) = \frac{Pr(T > t + s \text{ and } T > t)}{Pr(T > t)}$$
$$= \frac{e^{-\lambda(t+s)}}{e^{-\lambda t}}$$
$$= e^{-\lambda s}$$

This value is independent of t (and so the time already spent has not been remembered).

PEPA activities and rates

When enabled an activity, $a = (\alpha, \lambda)$, will delay for a period determined by its associated distribution function, i.e. the probability that the activity *a* happens within a period of time of length *t* is $F_a(t) = 1 - e^{-\lambda t}$.

PEPA activities and rates

We can think of this as the activity setting a timer whenever it becomes enabled. The time allocated to the timer is determined by the rate of the activity. If several activities are enabled at the same time each will have its own associated timer. When the first timer finishes that activity takes place—the activity is said to complete or succeed. This means that the activity is considered to "happen": an external observer will witness the event of an activity of type α . An activity may be preempted, or aborted, if another one completes first.

PEPA and Markov processes

In a PEPA model if we define the stochastic process X(t), such that $X(t) = C_i$ indicates that the system behaves as component C_i at time t, then X(t) is a Markov process which can be characterised by a matrix, \boldsymbol{Q} .

PEPA and Markov processes

In a PEPA model if we define the stochastic process X(t), such that $X(t) = C_i$ indicates that the system behaves as component C_i at time t, then X(t) is a Markov process which can be characterised by a matrix, \boldsymbol{Q} .

A stationary or equilibrium probability distribution, $\pi(\cdot)$, exists for every time-homogeneous irreducible Markov process whose states are all positive-recurrent.

PEPA and Markov processes

In a PEPA model if we define the stochastic process X(t), such that $X(t) = C_i$ indicates that the system behaves as component C_i at time t, then X(t) is a Markov process which can be characterised by a matrix, Q.

A stationary or equilibrium probability distribution, $\pi(\cdot)$, exists for every time-homogeneous irreducible Markov process whose states are all positive-recurrent.

This distribution is found by solving the global balance equation

$$\pi oldsymbol{Q} = oldsymbol{0}$$

subject to the normalisation condition

$$\sum \pi(C_i) = 1.$$

Stephen Gilmore. LFCS, University of Edinburgh.

PEPA and time

All PEPA models are time-homogeneous since all activities are time-homogeneous: the rate and type of activities enabled by a component are independent of time.

PEPA and irreducibility and positive-recurrence

The other conditions, irreducibility and positive-recurrent states, are easily expressed in terms of the derivation graph of the PEPA model. We only consider PEPA models with a finite number of states so if the model is irreducible then all states must be positive-recurrent i.e. the derivation graph is strongly connected.



In terms of the PEPA model this means that all behaviours of the system must be recurrent; in particular, for every choice, whichever path is chosen it must eventually return to the point where the choice can be made again, possibly with a different outcome.

Continuous-Time Markov Chains (CTMCs)

A Markov process with discrete state space and discrete index set is called a Markov chain. The future behaviour of a Markov chain depends only on its current state, and not on how that state was reached. This is the Markov, or memoryless, property.

$$Pr(X(t_{n+1}) = x_{n+1} | X(t_n) = x_n, \dots, X(t_0) = x_0)$$

= $Pr(X(t_{n+1}) = x_{n+1} | X(t_n) = x_n)$

Stephen Gilmore. LFCS, University of Edinburgh.

Structured Operational Semantics

PEPA is defined using a Plotkin-style structured operational semantics (a "small step" semantics).

<ロ> (日) (日) (日) (日) (日)

-큰

Structured Operational Semantics

PEPA is defined using a Plotkin-style structured operational semantics (a "small step" semantics).

Prefix

$$(\alpha, r).E \xrightarrow{(\alpha, r)} E$$

Stephen Gilmore. LFCS, University of Edinburgh.

< ロ > < 同 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ >

-큰

Structured Operational Semantics

PEPA is defined using a Plotkin-style structured operational semantics (a "small step" semantics).

Prefix

$$(\alpha, r).E \xrightarrow{(\alpha, r)} E$$

Choice

$$\frac{E \xrightarrow{(\alpha,r)} E'}{E + F \xrightarrow{(\alpha,r)} E'}$$

$$\frac{F \xrightarrow{(\alpha,r)} F'}{E + F \xrightarrow{(\alpha,r)} F'}$$

Stephen Gilmore. LFCS, University of Edinburgh.

< ロ > < 同 > < 回 > < 回 > < 回 > < 回

Structured Operational Semantics: Cooperation ($\alpha \notin L$)

Cooperation

$$\frac{E \xrightarrow{(\alpha,r)} E'}{E \bowtie_{L} F \xrightarrow{(\alpha,r)} E' \bowtie_{L} F} (\alpha \notin L)$$

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

3

Structured Operational Semantics: Cooperation ($\alpha \notin L$)

Cooperation

$$\frac{E \xrightarrow{(\alpha,r)} E'}{E \bowtie_{L} F \xrightarrow{(\alpha,r)} E' \bowtie_{L} F} (\alpha \notin L)$$

$$\frac{F \xrightarrow{(\alpha,r)} F'}{E \bowtie_{L} F \xrightarrow{(\alpha,r)} E \bowtie_{L} F'} (\alpha \notin L)$$

Stephen Gilmore. LFCS, University of Edinburgh.

• • • • • • • • • • • •

3

Structured Operational Semantics: Cooperation ($\alpha \in L$)

Cooperation
$$\frac{E \xrightarrow{(\alpha, r_1)} E' F \xrightarrow{(\alpha, r_2)} F'}{E \bowtie_{L} F \xrightarrow{(\alpha, R)} E' \bowtie_{L} F'} (\alpha \in L)$$

Stephen Gilmore. LFCS, University of Edinburgh.
Image: A math a math

3

Structured Operational Semantics: Cooperation ($\alpha \in L$)

Cooperation
$$\frac{E \xrightarrow{(\alpha,r_1)} E' F \xrightarrow{(\alpha,r_2)} F'}{E \bowtie_{L} F \xrightarrow{(\alpha,R)} E' \bowtie_{L} F'} (\alpha \in L)$$

where
$$R = \frac{r_1}{r_{\alpha}(E)} \frac{r_2}{r_{\alpha}(F)} min(r_{\alpha}(E), r_{\alpha}(F))$$

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-큰

Structured Operational Semantics: Hiding

Hiding

$$\frac{E \xrightarrow{(\alpha,r)} E'}{E/L \xrightarrow{(\alpha,r)} E'/L} (\alpha \notin L)$$

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-큰

Structured Operational Semantics: Hiding

Hiding

$$\frac{E \xrightarrow{(\alpha,r)} E'}{E/L \xrightarrow{(\alpha,r)} E'/L} (\alpha \notin L)$$

$$\frac{E \xrightarrow{(\alpha,r)} E'}{E/L \xrightarrow{(\tau,r)} E'/L} (\alpha \in L)$$

Stephen Gilmore. LFCS, University of Edinburgh.

< ロ > < 同 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < 回 > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ >

-큰

Structured Operational Semantics: Constants

Constant

$$\frac{E \xrightarrow{(\alpha,r)} E'}{A \xrightarrow{(\alpha,r)} E'} (A \stackrel{def}{=} E)$$

Stephen Gilmore. LFCS, University of Edinburgh.

· < /⊒ > < ∃ > <

Properties of the definition (1)

PEPA has no "nil" (a deadlocked process). This is because the PEPA language is intended for modelling non-stop processes (such as Web servers, operating systems, or manufacturing processes) rather than for modelling terminating processes (a compilation, a sorting routine, and so forth).

A (10) × (10) × (10)

Roll your own!

When we are interested in transient behaviour we use the deadlocked process *Stop* to signal a component which performs no further actions.

$$Stop \stackrel{\text{def}}{=} \left(\left((a, r).Stop \right) \underset{\{a, b\}}{\bowtie} \left((b, r).Stop \right) \right) / \{a, b\}$$

Stephen Gilmore. LFCS, University of Edinburgh.

Properties of the definition (2)

Cooperation in PEPA is multi-way. Two, three, four or more partners may cooperate, and they all need to synchronise for the activity to happen.

Properties of the definition (2)

Cooperation in PEPA is multi-way. Two, three, four or more partners may cooperate, and they all need to synchronise for the activity to happen.

This comes from the fact that synchronisation has the form $a, a \rightarrow a$ (as in CSP) instead of $a, \bar{a} \rightarrow \tau$ (as in CCS and the π -calculus).

Properties of the definition (2)

Cooperation in PEPA is multi-way. Two, three, four or more partners may cooperate, and they all need to synchronise for the activity to happen.

This comes from the fact that synchronisation has the form $a, a \rightarrow a$ (as in CSP) instead of $a, \overline{a} \rightarrow \tau$ (as in CCS and the π -calculus).

This is used to have "witnesses" to events (known as stochastic probes).

Properties of the definition (3)

Because of its mapping onto a CTMC, PEPA has an interleaving semantics.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part II: Markovian methods ▲ロト ▲圖 ▶ ▲ 国 ▶ ▲ 国 ▶ ▲ 国 ▶ ▲ 国 ▶

- **(())) (())) ())**

Properties of the definition (3)

Because of its mapping onto a CTMC, PEPA has an interleaving semantics.

Other modelling formalisms based on CTMCs are also based on an interleaving semantics (e.g. Generalised Stochastic Petri nets).

Model solution and analysis

A continuous time Markov chain (CTMC) is generated from a PEPA model via its structured operational semantics.

Model solution and analysis

A continuous time Markov chain (CTMC) is generated from a PEPA model via its structured operational semantics.

Linear algebra is used to solve the model in terms of equilibrium behaviour.

Model solution and analysis

A continuous time Markov chain (CTMC) is generated from a PEPA model via its structured operational semantics.

Linear algebra is used to solve the model in terms of equilibrium behaviour.

The resulting probability distribution is seldom the ultimate goal of performance analysis; a modeller derives performance measures from this distribution via a reward structure.

A logical foundation for the specification language

The expression, and testing for satisfaction of equilibrium properties, can be seen to be closely related to the specification, and model checking of a formula expressed in Larsen and Skou's probabilistic modal logic (PML). We give a modified interpretation of such formulae suitable for reasoning about PEPA's continuous time models.

A logical foundation for the specification language

The expression, and testing for satisfaction of equilibrium properties, can be seen to be closely related to the specification, and model checking of a formula expressed in Larsen and Skou's probabilistic modal logic (PML). We give a modified interpretation of such formulae suitable for reasoning about PEPA's continuous time models.

We exploit the operators of modal logic to be more discriminating about which states contribute to the reward measure. In particular, we can select a state based on model behaviour which is not immediately local to the state.

<ロ> (日) (日) (日) (日) (日)

-큰

Larsen and Skou's PML



Stephen Gilmore. LFCS, University of Edinburgh.

Relation to PEPA

Defn.
$$P \xrightarrow{(\alpha,\nu)} S$$
 if for all $P' \in S$, $P \xrightarrow{\alpha} P'$ and $\sum \{r \mid P \xrightarrow{(\alpha,r)} P', P' \in S\} = \nu$.

メロト メポト メヨト メヨト

-큰

Relation to PEPA

Defn.
$$P \xrightarrow{(\alpha,\nu)} S$$
 if for all $P' \in S$, $P \xrightarrow{\alpha} P'$ and $\sum \{r \mid P \xrightarrow{(\alpha,r)} P', P' \in S\} = \nu$.

Let P be a model of a PEPA process.

$$P \models tt$$

メロト メポト メヨト メヨト

-큰

Relation to PEPA

Defn.
$$P \xrightarrow{(\alpha,\nu)} S$$
 if for all $P' \in S$, $P \xrightarrow{\alpha} P'$ and $\sum \{r \mid P \xrightarrow{(\alpha,r)} P', P' \in S\} = \nu$.

Let P be a model of a PEPA process.

$$P \models \mathsf{tt}$$
$$P \models \neg F \text{ if } P \not\models F$$

Stephen Gilmore. LFCS, University of Edinburgh.

<ロ> (日) (日) (日) (日) (日)

-큰

Relation to PEPA

Defn.
$$P \xrightarrow{(\alpha,\nu)} S$$
 if for all $P' \in S$, $P \xrightarrow{\alpha} P'$ and $\sum \{r \mid P \xrightarrow{(\alpha,r)} P', P' \in S\} = \nu$.

Let P be a model of a PEPA process.

$$P \models \mathsf{tt}$$

$$P \models \neg F \text{ if } P \not\models F$$

$$P \models F_1 \land F_2 \text{ if } P \models F_1 \text{ and } P \models F_2$$

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-큰

Relation to PEPA

Defn.
$$P \xrightarrow{(\alpha,\nu)} S$$
 if for all $P' \in S$, $P \xrightarrow{\alpha} P'$ and $\sum \{r \mid P \xrightarrow{(\alpha,r)} P', P' \in S\} = \nu$.

Let P be a model of a PEPA process.

$$P \models \mathsf{tt}$$

$$P \models \neg F \text{ if } P \not\models F$$

$$P \models F_1 \land F_2 \text{ if } P \models F_1 \text{ and } P \models F_2$$

$$P \models \nabla_{\alpha} \text{ if } P \xrightarrow{\alpha}$$

Stephen Gilmore. LFCS, University of Edinburgh.

Relation to PEPA

Defn.
$$P \xrightarrow{(\alpha,\nu)} S$$
 if for all $P' \in S$, $P \xrightarrow{\alpha} P'$ and $\sum \{r \mid P \xrightarrow{(\alpha,r)} P', P' \in S\} = \nu$.

Let P be a model of a PEPA process.

$$P \models \text{tt}$$

$$P \models \neg F \text{ if } P \not\models F$$

$$P \models F_1 \land F_2 \text{ if } P \models F_1 \text{ and } P \models F_2$$

$$P \models \nabla_{\alpha} \text{ if } P \xrightarrow{\alpha}$$

$$P \models \langle \alpha \rangle_{\mu} F \text{ if } P \xrightarrow{(\alpha,\nu)} S \text{ for some } \nu \ge \mu,$$
and for all $P' \in S, P' \models F$

Stephen Gilmore. LFCS, University of Edinburgh.

Modal characterisation of strong equivalence

Let P be a model of a PEPA process. Then

```
P \cong Q iff for all F, P \models F iff Q \models F
```

That is to say that two PEPA processes are strongly equivalent (in particular, their underlying Markov chains are lumpably equivalent) if and only if they both satisfy, in the setting where rates are quantified, the same set of PML formulae.

Case study: active badges

We have used the PEPA modelling language and its accompanying specification language to analyse the configuration of a location tracking system based on active badges. Active badges transmit unique infra-red signals which are detected by networked sensors. These report locations back to a central database.

Case study: active badges

The badges are battery-powered and the tradeoff in the system is between the conservation of battery power and the accuracy of the information harvested from the sensors. When transmissions from badges collide, the badges sleep for a randomly determined time before retrying.

Active badges: the PEPA model

The PEPA model of this system tracks the progress of one badge-wearer around three connected corridors (numbered 14, 15 and 16). The activities which are performed in the system include the badge registering with a sensor (at rate r), the person moving to another corridor (at rate m) and a sensor reporting back to the central database (at rate s).

イロト イヨト イヨト イヨト

-큰

Active badges: the PEPA model

Person

$$P_{14} \stackrel{def}{=} (reg_{14}, r).P_{14} + (move_{15}, m).P_{15}$$

$$P_{15} \stackrel{def}{=} (reg_{15}, r).P_{15} + (move_{14}, m).P_{14} + (move_{16}, m).P_{16}$$

$$P_{16} \stackrel{def}{=} (reg_{16}, r).P_{16} + (move_{15}, m).P_{15}$$

イロト イヨト イヨト イヨト

-큰

Active badges: the PEPA model

Person

$$\begin{array}{rcl}
P_{14} & \stackrel{def}{=} & (reg_{14}, r).P_{14} + (move_{15}, m).P_{15} \\
P_{15} & \stackrel{def}{=} & (reg_{15}, r).P_{15} + (move_{14}, m).P_{14} + (move_{16}, m).P_{16} \\
P_{16} & \stackrel{def}{=} & (reg_{16}, r).P_{16} + (move_{15}, m).P_{15}
\end{array}$$

Sensor

$$\begin{array}{rcl} S_{14} & \stackrel{def}{=} & (reg_{14}, \top).(rep_{14}, s).S_{14} \\ S_{15} & \stackrel{def}{=} & (reg_{15}, \top).(rep_{15}, s).S_{15} \\ S_{16} & \stackrel{def}{=} & (reg_{16}, \top).(rep_{16}, s).S_{16} \end{array}$$

Stephen Gilmore. LFCS, University of Edinburgh.

Active badges: the PEPA model

Database

DB_{14}	def =	$(rep_{14}, \top).DB_{14} + (rep_{15}, \top).DB_{15} + (rep_{16}, \top).DB_{16}$
DB_{15}	def =	$(rep_{14}, \top).DB_{14} + (rep_{15}, \top).DB_{15} + (rep_{16}, \top).DB_{16}$
DB_{16}	$\stackrel{\tiny def}{=}$	$(rep_{14}, \top).DB_{14} + (rep_{15}, \top).DB_{15} + (rep_{16}, \top).DB_{16}$

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part II: Markovian methods ▲ロト ▲掃 ▶ ▲ 臣 ▶ ▲ 臣 ▶ ─ 臣 ─ のへで

イロト イヨト イヨト イヨト

-큰

Active badges: the PEPA model

Database

$$\begin{array}{rcl} DB_{14} & \stackrel{def}{=} & (rep_{14}, \top).DB_{14} + (rep_{15}, \top).DB_{15} + (rep_{16}, \top).DB_{16} \\ DB_{15} & \stackrel{def}{=} & (rep_{14}, \top).DB_{14} + (rep_{15}, \top).DB_{15} + (rep_{16}, \top).DB_{16} \\ DB_{16} & \stackrel{def}{=} & (rep_{14}, \top).DB_{14} + (rep_{15}, \top).DB_{15} + (rep_{16}, \top).DB_{16} \end{array}$$

System

$$P_{14} \bowtie_{l} (S_{14} \parallel S_{15} \parallel S_{16}) \bowtie_{M} DB_{14}$$

where
$$L = \{ reg_{14}, reg_{15}, reg_{16} \}$$

 $M = \{ rep_{14}, rep_{15}, rep_{16} \}$

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

2

Probability that the database holds inaccurate information



Stephen Gilmore. LFCS, University of Edinburgh.

Outline

A modelling language

A semantics for the modelling language

Tools for the modelling language

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part II: Markovian methods

▲ロト ▲圖 ト ▲ 臣 ト ▲ 臣 ト → 臣 → のへで

A high-level language for performance modelling

PEPA is a high-level language for performance modelling. PEPA models describe stochastic (in fact, Markovian) processes.

. ↓ □ ▶ ↓ □ ▶ ↓

A high-level language for performance modelling

PEPA is a high-level language for performance modelling. PEPA models describe stochastic (in fact, Markovian) processes.

Without a high-level modelling language, the modeller would be forced to work with unstructured matrix representations of stochastic processes.

A high-level language for performance modelling

PEPA is a high-level language for performance modelling. PEPA models describe stochastic (in fact, Markovian) processes.

Without a high-level modelling language, the modeller would be forced to work with unstructured matrix representations of stochastic processes.

Process algebras are useful because they allow the definition of equivalence relations between model components and these relations may be used in model simplification.
The PEPA Workbench

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

The PEPA Workbench: functionality

The workbench will report errors in the model function:

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

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

The PEPA Workbench: functionality

The workbench will report errors in the model function:

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

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

The workbench provides a simple pattern language for selecting states from the stationary distribution.

イロト イポト イヨト イヨト

-

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

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part II: Markovian methods

-큰

<ロト <回ト < 回ト

PEPA Workbench input $P_1 \stackrel{\text{def}}{=} (start, r_1).P_2 \qquad P_2 \stackrel{\text{def}}{=} (run, r_2).P_3 \qquad P_3 \stackrel{\text{def}}{=} (stop, r_3).P_1$ $P_1 \parallel P_1$

PEPA Workbench output



Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part II: Markovian methods

Stochastic Modelling Part III: Scalability

Stephen Gilmore LFCS, University of Edinburgh

GLOBAN Summerschool Lygby, Denmark

25th August 2006

• • • • • • • • • • • •

3

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

2

Outline

Numerical evaluation via CTMCs and ODEs

Case study in Web Services

Commentary and comparison

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

Performance Evaluation Process Algebra

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

Stephen Gilmore. LFCS, University of Edinburgh.

Performance Evaluation Process Algebra

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

The rate at which an activity is performed is quantified by some component in each co-operation. The symbol \top indicates that the rate value is quantified elsewhere (not in this component).

Stephen Gilmore. LFCS, University of Edinburgh.

Performance Evaluation Process Algebra

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

The rate at which an activity is performed is quantified by some component in each co-operation. The symbol \top indicates that the rate value is quantified elsewhere (not in this component).

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

PEPA: informal semantics (sequential sublanguage)

$(\alpha, r).S$

The activity (α, r) takes time Δt (drawn from the exponential distribution with parameter r).

$S_1 + S_2$

In this choice either S_1 or S_2 will complete an activity first. The other is discarded.

(日) (同) (三) (三)

PEPA: informal semantics (combinators)

$\begin{array}{ccc} C_1 & \bigsqcup_{L} & C_2 \\ & & \text{All activities of } C_1 \text{ and } C_2 \text{ with types in } L \text{ are shared: others remain individual.} \\ & & \text{NOTATION: write } C_1 \parallel C_2 \text{ if } L \text{ is empty.} \end{array}$

C / LActivities of C with types in L are hidden (τ type activities) to be thought of as internal delays.

<ロ> (日) (日) (日) (日) (日)

-큰

Derived forms and additional syntax

$P_1 \parallel P_2$ is a derived form for $P_1 \bowtie_{0} P_2$.

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

3

Derived forms and additional syntax

 $P_1 \parallel P_2$ is a derived form for $P_1 \bowtie_{0} P_2$.

Because we are interested in transient behaviour we use the deadlocked process *Stop*.

Stephen Gilmore. LFCS, University of Edinburgh.

Derived forms and additional syntax

$P_1 \parallel P_2$ is a derived form for $P_1 \bowtie_{0} P_2$.

Because we are interested in transient behaviour we use the deadlocked process *Stop*.

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

(日) (同) (三) (三)

Derived forms and additional syntax

 $P_1 \parallel P_2$ is a derived form for $P_1 \bowtie P_2$.

Because we are interested in transient behaviour we use the deadlocked process *Stop*.

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

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

Stephen Gilmore. LFCS, University of Edinburgh.

Background: Deterministic processes

A process is a deterministic process if knowledge of its values up to and including time t allows us to unambiguously predict its value at any infinitesimally later time t + dt.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part III: Scalability ・ロト・御ト・当ト・当ト・ 油・ ろんの

Background: ODEs are memoryless deterministic processes

A set of ordinary differential equations defines a memoryless deterministic process.

$$\begin{aligned} \mathbf{X}(t+dt) &= \mathbf{X}(t) + f(\mathbf{X}(t), t) dt \\ \frac{d\mathbf{X}}{dt} &= f(\mathbf{X}, t) \end{aligned}$$

Stephen Gilmore. LFCS, University of Edinburgh.

Background: Stochastic processes

A process is a stochastic process if knowledge of its values up to and including time t allows us to probabilistically predict its value at any infinitesimally later time t + dt.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part III: Scalability ・ロト・西ト・西ト・西ト・日・ シックの

Background: Stochastic processes

A process is a stochastic process if knowledge of its values up to and including time t allows us to probabilistically predict its value at any infinitesimally later time t + dt.

Stochastic processes subsume deterministic processes.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part III: Scalability

Outline

Numerical evaluation via CTMCs and ODEs

Case study in Web Services

Commentary and comparison

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

æ

Modelling with quantified process algebras

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1) \end{array}$$

Stephen Gilmore. LFCS, University of Edinburgh.

Modelling with quantified process algebras

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, \textit{r}).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, \textit{r}).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, \textit{r}).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

This example defines a system with nine reachable states:

1. $P_1 \parallel P_1$	4 . $P_2 \parallel P_1$	7. $P_3 \parallel P_1$
2. $P_1 \parallel P_2$	5. $P_2 \parallel P_2$	8. <i>P</i> ₃ <i>P</i> ₂
3. <i>P</i> ₁ ∥ <i>P</i> ₃	6. <i>P</i> ₂ <i>P</i> ₃	9 . <i>P</i> ₃ ∥ <i>P</i> ₃

The transitions between states have quantified duration r which can be evaluated against a CTMC or ODE interpretation.

- 4 同 1 4 回 1 4 回 1

Analysis based on Continuous-time Markov Chains

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using transient analysis we can evaluate the probability of each state with respect to time. For t = 0:

1. 1.0000	4. 0.0000	7. 0.0000
2. 0.0000	5. 0.0000	8. 0.0000
3. 0.0000	6. 0.0000	9. 0.0000

Analysis based on Continuous-time Markov Chains

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, \textit{r}).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, \textit{r}).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, \textit{r}).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using transient analysis we can evaluate the probability of each state with respect to time. For t = 1:

1. 0.1642	4. 0.1567	7. 0.0842
2. 0.1567	5. 0.1496	8. 0.0804
3. 0.0842	<u>6</u> . 0.0804	9. 0.0432

Analysis based on Continuous-time Markov Chains

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using transient analysis we can evaluate the probability of each state with respect to time. For t = 2:

1.	0.1056	4.	0.1159	7.	0.1034
2.	0.1159	5.	0.1272	8.	0.1135
3.	0.1034	6.	0.1135	9.	0.1012

Analysis based on Continuous-time Markov Chains

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using transient analysis we can evaluate the probability of each state with respect to time. For t = 3:

1.	0.1082	4.	0.1106	7.	0.1100
2.	0.1106	5.	0.1132	8.	0.1125
3.	0.1100	6.	0.1125	9.	0.1119

Analysis based on Continuous-time Markov Chains

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using transient analysis we can evaluate the probability of each state with respect to time. For t = 4:

1.	0.1106	4.	0.1108	7.	0.1111
2.	0.1108	5.	0.1110	8.	0.1113
3.	0.1111	6.	0.1113	9.	0.1116

Analysis based on Continuous-time Markov Chains

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, \textit{r}).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, \textit{r}).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, \textit{r}).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using transient analysis we can evaluate the probability of each state with respect to time. For t = 5:

1.	0.1111	4.	0.1110	7.	0.1111
2.	0.1110	5.	0.1110	8.	0.1111
3.	0.1111	6.	0.1111	9.	0.1111

Analysis based on Continuous-time Markov Chains

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using transient analysis we can evaluate the probability of each state with respect to time. For t = 6:

1. 0.	1111	4.	0.1111	7.	0.1111
2. 0.	1111	5.	0.1110	8.	0.1111
3. 0.	1111	6.	0.1111	9.	0.1111

Analysis based on Continuous-time Markov Chains

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using transient analysis we can evaluate the probability of each state with respect to time. For t = 7:

1. 0.1	1111	4.	0.1111	7.	0.1111
2. 0.	1111	5.	0.1111	8.	0.1111
3. 0.1	1111	6.	0.1111	9.	0.1111

Analysis based on Ordinary Differential Equations

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using the ordinary differential equation semantics we can compute the expected number of each type of component.

For
$$t = 0$$
: P_1 2.0000
 P_2 0.0000
 P_3 0.0000

Analysis based on Ordinary Differential Equations

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using the ordinary differential equation semantics we can compute the expected number of each type of component.

For
$$t = 1$$
: P_1 0.8121
 P_2 0.7734
 P_3 0.4144

Stephen Gilmore. LFCS, University of Edinburgh.

Analysis based on Ordinary Differential Equations

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using the ordinary differential equation semantics we can compute the expected number of each type of component.

For
$$t = 2$$
: P_1 0.6490
 P_2 0.7051
 P_3 0.6457

Analysis based on Ordinary Differential Equations

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, \textit{r}).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, \textit{r}).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, \textit{r}).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using the ordinary differential equation semantics we can compute the expected number of each type of component.

For
$$t = 3$$
: P_1 0.6587
 P_2 0.6719
 P_3 0.6692

Stephen Gilmore. LFCS, University of Edinburgh.

Analysis based on Ordinary Differential Equations

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using the ordinary differential equation semantics we can compute the expected number of each type of component.

For
$$t = 4$$
: P_1 0.6648
 P_2 0.6665
 P_3 0.6685

Stephen Gilmore. LFCS, University of Edinburgh.
Analysis based on Ordinary Differential Equations

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, r).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, r).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, r).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using the ordinary differential equation semantics we can compute the expected number of each type of component.

For
$$t = 5$$
: P_1 0.6666
 P_2 0.6663
 P_3 0.6669

Analysis based on Ordinary Differential Equations

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\scriptscriptstyle def}{=} (\textit{start}, \textit{r}).P_2 & P_2 \stackrel{\scriptscriptstyle def}{=} (\textit{run}, \textit{r}).P_3 & P_3 \stackrel{\scriptscriptstyle def}{=} (\textit{stop}, \textit{r}).P_1 \\ \textit{System} \stackrel{\scriptscriptstyle def}{=} (P_1 \parallel P_1) \end{array}$$

Using the ordinary differential equation semantics we can compute the expected number of each type of component.

For
$$t = 6$$
: P_1 0.6666
 P_2 0.6666
 P_3 0.6666

Analysis based on Ordinary Differential Equations

Tiny example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1) \end{array}$$

Using the ordinary differential equation semantics we can compute the expected number of each type of component.

For
$$t = 7$$
: P_1 0.6666
 P_2 0.6666
 P_3 0.6666

Analysis based on Ordinary Differential Equations

Slightly larger example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1 \parallel P_1) \end{array}$$

A slightly larger example with a third copy of the process also initiated in state P_1 .

For
$$t = 0$$
: P_1 3.0000
 P_2 0.0000
 P_3 0.0000

Stephen Gilmore. LFCS, University of Edinburgh.

Analysis based on Ordinary Differential Equations

Slightly larger example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1 \parallel P_1) \end{array}$$

A slightly larger example with a third copy of the process also initiated in state P_1 .

For
$$t = 1$$
: P_1 1.1782
 P_2 1.1628
 P_3 0.6590

Stephen Gilmore. LFCS, University of Edinburgh.

Analysis based on Ordinary Differential Equations

Slightly larger example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1 \parallel P_1) \end{array}$$

A slightly larger example with a third copy of the process also initiated in state P_1 .

For
$$t = 2$$
: P_1 0.9766
 P_2 1.0754
 P_3 0.9479

Stephen Gilmore. LFCS, University of Edinburgh.

Analysis based on Ordinary Differential Equations

Slightly larger example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1 \parallel P_1) \end{array}$$

A slightly larger example with a third copy of the process also initiated in state P_1 .

For
$$t = 3$$
: P_1 0.9838
 P_2 1.0142
 P_3 1.0020

Analysis based on Ordinary Differential Equations

Slightly larger example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1 \parallel P_1) \end{array}$$

A slightly larger example with a third copy of the process also initiated in state P_1 .

For
$$t = 4$$
: P_1 0.9981
 P_2 0.9995
 P_3 1.0023

Stephen Gilmore. LFCS, University of Edinburgh.

Analysis based on Ordinary Differential Equations

Slightly larger example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1 \parallel P_1) \end{array}$$

A slightly larger example with a third copy of the process also initiated in state P_1 .

For
$$t = 5$$
: P_1 1.0001
 P_2 0.9996
 P_3 1.0003

Analysis based on Ordinary Differential Equations

Slightly larger example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1 \parallel P_1) \end{array}$$

A slightly larger example with a third copy of the process also initiated in state P_1 .

For
$$t = 6$$
: P_1 1.0001
 P_2 0.9999
 P_3 1.0000

Analysis based on Ordinary Differential Equations

Slightly larger example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1 \parallel P_1) \end{array}$$

A slightly larger example with a third copy of the process also initiated in state P_1 .

For
$$t = 7$$
: P_1 1.0000
 P_2 0.9999
 P_3 0.9999

Stephen Gilmore. LFCS, University of Edinburgh.

Analysis based on Ordinary Differential Equations

Slightly larger example

$$\begin{array}{ll} P_1 \stackrel{\text{\tiny def}}{=} (start, r).P_2 & P_2 \stackrel{\text{\tiny def}}{=} (run, r).P_3 & P_3 \stackrel{\text{\tiny def}}{=} (stop, r).P_1 \\ System \stackrel{\text{\tiny def}}{=} (P_1 \parallel P_1 \parallel P_1) \end{array}$$

A slightly larger example with a third copy of the process also initiated in state P_1 .

For
$$t = 8$$
: P_1 1.0000
 P_2 1.0000
 P_3 1.0000

What just happened?

An ODE specifies how the value of some continuous variable varies over continuous time. For example, the temperature in a container may be modelled by an ODE describing how the temperature will change dependent on the current temperature and pressure. The pressure can be similarly modelled and the equations together form a system of ODEs describing the state of the container.

What just happened?

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 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 ODE.

Isn't this just the Chapman-Kolmogorov equations?

It is possible to perform transient analysis of a continuous-time Markov chain by solving the Chapman-Kolmogorov differential equations:

$$\frac{d\pi(t)}{dt} = \pi(t)Q$$

[Stewart, 1994]

Stephen Gilmore. LFCS, University of Edinburgh.

Isn't this just the Chapman-Kolmogorov equations?

It is possible to perform transient analysis of a continuous-time Markov chain by solving the Chapman-Kolmogorov differential equations:

$$\frac{d\pi(t)}{dt} = \pi(t)Q$$

[Stewart, 1994]

That's not what we're doing. We go directly to ODEs.

Stephen Gilmore. LFCS, University of Edinburgh.

What's the value proposition?

The bottleneck for Markovian modelling of systems is the size of the solution vector, which is bounded by the product of the state-space sizes of the processes which are composed in parallel ("state-space explosion").

Stephen Gilmore. LFCS, University of Edinburgh.

What's the value proposition?

- The bottleneck for Markovian modelling of systems is the size of the solution vector, which is bounded by the product of the state-space sizes of the processes which are composed in parallel ("state-space explosion").
- The size of the solution vector for the system of ODEs may be exponentially smaller.

・ロト ・聞ト ・ヨト ・ヨト

2

Outline

Numerical evaluation via CTMCs and ODEs

Case study in Web Services

Commentary and comparison

Stephen Gilmore. 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.



- 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.

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.

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.

PEPA model: Second party clients



- 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.

(日) (同) (三) (三)

-큰

PEPA model: Second party clients



Stephen Gilmore. LFCS, University of Edinburgh.

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.

PEPA model: Brokers



 $\stackrel{def}{=}$ $(request_{h}, \top)$. Broker_{dec_input} Broker_{idle} $\stackrel{def}{=}$ Broker_{dec_input} (*decrypt_{sp}*, *r_{b_dec_sp}*).Broker_{enc_input} $\stackrel{def}{=}$ Broker_{enc_input} (encrypt_{ws}, r_{b_enc_ws}).Broker_{sending} $\stackrel{def}{=}$ Broker_{sending} (request_{ws}, r_{b_rea}).Broker_{waiting} def = Broker_{waiting} $(response_{ws}, \top)$. Broker_{dec_resp} def = Broker_{dec_resp} (decrypt_{ws}, r_{b dec ws}).Broker_{enc_resp} def = Broker_{enc_resp} $(encrypt_{sp}, r_{b_{enc_{sp}}})$. Broker_{replying} $\stackrel{def}{=}$ Broker_{replying} (response_b, r_{b_resp}).Broker_{idle}

Stephen Gilmore. LFCS, University of Edinburgh.

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.

(日) (同) (日) (日)

-큰

PEPA model: First party clients



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

Stephen Gilmore. LFCS, University of Edinburgh.

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.

4

PEPA model: Web service



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}	$\stackrel{def}{=}$	$(execute_{ws}, r_{ws_exec}).WS_{securing}$
WS _{securing}	$\stackrel{def}{=}$	$(encryptResp_{ws}, r_{ws_enc_b}).WS_{responding}$
WS _{responding}	def =	$(response_{ws}, r_{ws_resp_b}).WS_{idle}$
WS _{method}	$\stackrel{def}{=}$	(execute _{ws} , r _{ws_exec}).WS _{returning}
WS _{returning}	def =	(result _{ws} , r _{ws_res}).WS _{idle}

Stephen Gilmore. LFCS, University of Edinburgh.

PEPA model: Web service



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

Stephen Gilmore. LFCS, University of Edinburgh.

PEPA model: Web service



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

Stephen Gilmore. LFCS, University of Edinburgh.

4

PEPA model: Web service



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}	$\stackrel{def}{=}$	$(execute_{ws}, r_{ws_exec}).WS_{securing}$
WS _{securing}	$\stackrel{def}{=}$	$(encryptResp_{ws}, r_{ws_enc_b}).WS_{responding}$
WS _{responding}	def =	$(response_{ws}, r_{ws_resp_b}).WS_{idle}$
WS _{method}	$\stackrel{\rm def}{=}$	(execute _{ws} , r _{ws_exec}).WS _{returning}
WS _{returning}	def =	(result _{ws} , r _{ws_res}).WS _{idle}

Stephen Gilmore. LFCS, University of Edinburgh.
(日) (同) (日) (日)

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
$$\mathcal{K} = \{ request_b, response_b \}$$

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

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

Stephen Gilmore. 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.

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

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

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

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.
- 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.
- 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

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part III: Scalability

Numerical evaluation via CTMCs and ODEs

Commentary and comparison

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

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part III: Scalability

Second party clients	Brokers	Web service instance	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
					1		1			

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part III: Scalability

Numerical evaluation via CTMCs and ODEs

ŝ

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

Ш

Running times from analyses (in seconds)

11

Ш

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part III: Scalability

Numerical evaluation via CTMCs and ODEs

Ш

Ш

Ш

Ш

▲ロト▲聞と▲臣と▲臣と「臣」 釣べ⊙

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

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part III: Scalability

三 のへの

イロト イヨト イヨト イヨト

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

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part III: Scalability

三 のへで

イロト イヨト イヨト イヨト

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

Stephen Gilmore. 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

Stephen Gilmore. 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.
- ► 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



Stephen Gilmore. LFCS, University of Edinburgh.

Brokers



Stephen Gilmore. LFCS, University of Edinburgh.

First party clients



Stephen Gilmore. LFCS, University of Edinburgh.

Web service



Stephen Gilmore. LFCS, University of Edinburgh.

Outline

Numerical evaluation via CTMCs and ODEs

Case study in Web Services

Commentary and comparison

・ロト ・母 ト ・ヨト ・ヨト ・ヨー うへの

Stephen Gilmore. LFCS, University of Edinburgh.

· < /⊒ > < ∃ > <

Commentary and comparison

Previous performance modelling with PEPA used continuous-time Markov chains (CTMCs). These admit steady-state and transient analysis (by solving the CTMC).

Stephen Gilmore. LFCS, University of Edinburgh.

- Previous performance modelling with PEPA used continuous-time Markov chains (CTMCs). These admit steady-state and transient analysis (by solving the CTMC).
- Steady-state is cheaper but less informative. Transient is more informative but more expensive.

Stephen Gilmore. LFCS, University of Edinburgh.

- Previous performance modelling with PEPA used continuous-time Markov chains (CTMCs). These admit steady-state and transient analysis (by solving the CTMC).
- Steady-state is cheaper but less informative. Transient is more informative but more expensive.
- Major drawback: state-space explosion. Generating the state-space is slow. Solving the CTMC is slow.

- Previous performance modelling with PEPA used continuous-time Markov chains (CTMCs). These admit steady-state and transient analysis (by solving the CTMC).
- Steady-state is cheaper but less informative. Transient is more informative but more expensive.
- Major drawback: state-space explosion. Generating the state-space is slow. Solving the CTMC is slow.
- In practice effective only to systems of size 10⁶ states, even when using clever storage representations.

(日) (同) (三) (三)

Commentary and comparison

 Mapping PEPA to ODEs admits *course-of-values* analysis by solving the ODE (akin to transient analysis).

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

Commentary and comparison

- Mapping PEPA to ODEs admits *course-of-values* analysis by solving the ODE (akin to transient analysis).
- Major benefit: avoids state-space generation entirely.

Stephen Gilmore. LFCS, University of Edinburgh.

- Mapping PEPA to ODEs admits *course-of-values* analysis by solving the ODE (akin to transient analysis).
- ► Major benefit: avoids state-space generation entirely.
- Major benefit: ODE solving is effective in practice, leaning towards suitability for interactive experimentation. Good for modellers, gaining more insights into the system behaviour.

- Mapping PEPA to ODEs admits *course-of-values* analysis by solving the ODE (akin to transient analysis).
- ► Major benefit: avoids state-space generation entirely.
- Major benefit: ODE solving is effective in practice, leaning towards suitability for interactive experimentation. Good for modellers, gaining more insights into the system behaviour.
- Effective for systems of size 10^{10^6} states and beyond.

· < /⊒ > < ∃ > <

Discussion: process algebras and ODEs

Models in the PEPA stochastic process algebra are concise, and in direct style they generate a system of ODEs the number of which is linear in the number of distinct component types in the PEPA model.

Stephen Gilmore. LFCS, University of Edinburgh.

. ↓ □ ▶ ↓ □ ▶ ↓

Discussion: process algebras and ODEs

- Models in the PEPA stochastic process algebra are concise, and in direct style they generate a system of ODEs the number of which is linear in the number of distinct component types in the PEPA model.
- Thus there is no hidden cost in the use of the high-level language but there are many advantages.

Stephen Gilmore. LFCS, University of Edinburgh.

Discussion: process algebras and ODEs

- Models in the PEPA stochastic process algebra are concise, and in direct style they generate a system of ODEs the number of which is linear in the number of distinct component types in the PEPA model.
- Thus there is no hidden cost in the use of the high-level language but there are many advantages.
 - PEPA models can be checked for freedom from deadlock, satisfaction of logical properties, or compared using relations such as bisimulation.

Discussion: process algebras and ODEs

- Models in the PEPA stochastic process algebra are concise, and in direct style they generate a system of ODEs the number of which is linear in the number of distinct component types in the PEPA model.
- Thus there is no hidden cost in the use of the high-level language but there are many advantages.
 - PEPA models can be checked for freedom from deadlock, satisfaction of logical properties, or compared using relations such as bisimulation.
 - As a compositional modelling language PEPA components can be re-used in other models, promoting best practice.

Stephen Gilmore. LFCS, University of Edinburgh.

False and true concurrency

In the process algebra world, algebras with an interleaving semantics are termed false concurrency. PEPA [Hillston 1994] was the first timed process algebra to have an interleaving semantics allowing it to generate a CTMC. The interleaving semantics gives rise to the state-space explosion problem.

Process algebras without an interleaving semantics are termed true concurrency process algebras. The search for a true concurrency timed process algebra has been a ten-year open problem.

PEPA [Hillston 2005] is the first timed process algebra to have a true concurrency semantics via the mapping to ODEs. The true concurrency semantics avoids the state-space explosion problem.

イロト イヨト イヨト イヨト

2

Stochastic Modelling Part IV: Variance

Stephen Gilmore LFCS, University of Edinburgh

GLOBAN Summerschool Lygby, Denmark

25th August 2006

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part IV: Variance

(日) (同) (三) (三)

Review: Problems of scale

Global computing systems are typically ubiquitous, may exhibit mobility, and are distributed. There are many places where computations can happen, and many replications of services.

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic Modelling. Part IV: Variance

Review: Problems of scale

- Global computing systems are typically ubiquitous, may exhibit mobility, and are distributed. There are many places where computations can happen, and many replications of services.
- In many types of stochastic modelling one runs into problems of scale.

Review: Problems of scale

- Global computing systems are typically ubiquitous, may exhibit mobility, and are distributed. There are many places where computations can happen, and many replications of services.
- In many types of stochastic modelling one runs into problems of scale.
- Continuous-Time Markov Chain

Review: Problems of scale

- Global computing systems are typically ubiquitous, may exhibit mobility, and are distributed. There are many places where computations can happen, and many replications of services.
- In many types of stochastic modelling one runs into problems of scale.
- Continuous-Time Markov Chain
 - 6 copies of a process with 10 states \Rightarrow 10⁶ states
Review: Problems of scale

- Global computing systems are typically ubiquitous, may exhibit mobility, and are distributed. There are many places where computations can happen, and many replications of services.
- In many types of stochastic modelling one runs into problems of scale.
- Continuous-Time Markov Chain
 - 6 copies of a process with 10 states $\Rightarrow 10^6$ states
 - The state-space explosion problem

Stephen Gilmore. LFCS, University of Edinburgh.

Review: Markovian methods and differential equations

 Markovian methods allow us to apply a thorough stochastic treatment in system modelling, giving an exact representation of a stochastic process.

▲ @ ▶ ▲ ∃ ▶ ▲

Review: Markovian methods and differential equations

- Markovian methods allow us to apply a thorough stochastic treatment in system modelling, giving an exact representation of a stochastic process.
- Differential equation methods allow us to scale this analysis to large problem sizes but rest on the assumption of continuity for component populations.

Modelling challenges: multi-scale systems

What about systems in which some components are present in large numbers and others are not?

Stephen Gilmore. LFCS, University of Edinburgh.

Modelling challenges: multi-scale systems

- What about systems in which some components are present in large numbers and others are not?
- ► The system will exhibit a high degree of variance.

Modelling challenges: multi-scale systems

- What about systems in which some components are present in large numbers and others are not?
- ► The system will exhibit a high degree of variance.
- Simulation allows us to investigate this.

Outline

Comparing stochastic simulation and ODEs

Stochastic simulation algorithms

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part IV: Variance ・ロト・御・・田・・田・・田・ シック

Outline

Comparing stochastic simulation and ODEs

Stochastic simulation algorithms

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part IV: Variance ▲ロト ▲御 ト ▲ 臣 ト ▲ 臣 ト 二臣 … の Q ()?

A simple example: processors and resources

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

$Proc_0[P] \bigotimes_{_{\{task1\}}} Res_0[R]$

Stephen Gilmore. LFCS, University of Edinburgh.

A simple example: processors and resources

$$\begin{array}{l} Proc_{0} \stackrel{def}{=} (task1, \top).Proc_{1} \\ Proc_{1} \stackrel{def}{=} (task2, r_{2}).Proc_{0} \\ Res_{0} \stackrel{def}{=} (task1, r_{1}).Res_{1} \\ Poc \stackrel{def}{=} (reset c) Poc \end{array}$$

$$Res_1 \cong (reset, s).Res_0$$

$$Proc_0[P] \bigotimes_{\{task1\}} Res_0[R]$$

Processors (P)	Resources (R)	States (2^{P+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

CTMC interpretation

Stephen Gilmore. LFCS, University of Edinburgh.

イロト イヨト イヨト イヨト

-큰

A simple example: processors and resources

$$Proc_{0} \stackrel{def}{=} (task1, \top).Proc_{1}$$

$$Proc_{1} \stackrel{def}{=} (task2, r_{2}).Proc_{0}$$

$$Res_{0} \stackrel{def}{=} (task1, r_{1}).Res_{1}$$

$$Res_{1} \stackrel{def}{=} (reset, s).Res_{0}$$

$$Proc_{0}[P] \underset{\{task1\}}{\bowtie} Res_{0}[R]$$

$$\frac{dRes_{1}}{dt} = r_{1} \min(Proc_{0}, Res_{0})$$

$$-r_{2} Proc_{1}$$

$$\frac{dRes_{0}}{dt} = -r_{1} \min(Proc_{0}, Res_{0})$$

$$+s Res_{1}$$

$$\frac{dRes_{1}}{dt} = r_{1} \min(Proc_{0}, Res_{0})$$

$$-s Res_{1}$$

Stephen Gilmore. LFCS, University of Edinburgh.

-

-큰

Processors and resources (SSA run A)



Stephen Gilmore. LFCS, University of Edinburgh.

э

Processors and resources (SSA run B)



Stephen Gilmore. LFCS, University of Edinburgh.

-큰

-

Processors and resources (SSA run C)



Stephen Gilmore. LFCS, University of Edinburgh.

-큰

-

Processors and resources (SSA run D)



Stephen Gilmore. LFCS, University of Edinburgh.

표 관 문

Processors and resources (average of 10 runs)



Stephen Gilmore. LFCS, University of Edinburgh.

Processors and resources (average of 100 runs)



Stephen Gilmore. LFCS, University of Edinburgh.

Processors and resources (average of 1000 runs)



Stephen Gilmore. LFCS, University of Edinburgh.

Processors and resources (average of 10000 runs)



Stephen Gilmore. LFCS, University of Edinburgh.

Processors and resources (ODE solution)



Stephen Gilmore. LFCS, University of Edinburgh.

From realisations to ensembles

As we repeatedly sample from the underlying random number distributions the average of the samples tends to the mean.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part IV: Variance

(日) (同) (三) (三)

2

Processors and resources: scaling out

What effect does increasing the number of copies have?

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part IV: Variance

э

Processors and resources (single SSA run, 100/80)



Stephen Gilmore. LFCS, University of Edinburgh.

-큰

Processors and resources (single SSA run, 1,000/800)



Stephen Gilmore. LFCS, University of Edinburgh.

Processors and resources (single SSA run, 10,000/8,000)



Stephen Gilmore. LFCS, University of Edinburgh.

-2

Processors and resources (single SSA run, 100,000/80,000)



Stephen Gilmore. LFCS, University of Edinburgh.

Systems of many components

Each specific run is individually in closer and closer agreement with the differential equation approach as the number of components in the system increases.

Outline

Comparing stochastic simulation and ODEs

Stochastic simulation algorithms

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part IV: <u>Variance</u> ▲ロト ▲課 ト ▲語 ト ▲語 ト 「語」 のへで

Stochastic simulation: Propensity functions

As explicitly derived by Gillespie, the stochastic model uses a form often termed the propensity function that gives the probability a_{μ} of activity μ occurring in time interval (t, t + dt).

$$a_{\mu} \mathrm{d}t = h_{\mu} c_{\mu} \mathrm{d}t$$

where the *M* possible activities are given an arbitrary index μ $(1 \le \mu \le M)$, h_{μ} denotes the number of possible combinations of components involved in μ , and c_{μ} is a stochastic rate constant.

Stochastic: Grand probability function

The stochastic formulation proceeds by considering the grand probability function $Pr(\mathbf{X}; t) \equiv probability$ that there will be at time t, X_i copies of component S_i , where $\mathbf{X} \equiv (X_1, X_2, \dots, X_N)$ is a vector of populations.

Stephen Gilmore. LFCS, University of Edinburgh. Stochastic Modelling. Part IV: Variance

Stochastic: Grand probability function

The stochastic formulation proceeds by considering the grand probability function $Pr(\mathbf{X}; t) \equiv probability$ that there will be at time t, X_i copies of component S_i , where $\mathbf{X} \equiv (X_1, X_2, \dots, X_N)$ is a vector of populations.

Evidently, knowledge of this function provides a complete understanding of the probability distribution of all possible states at all times.

▲ @ ▶ ▲ ∃ ▶ ▲

Stochastic: Infinitesimal time interval

By considering a discrete infinitesimal time interval (t, t + dt) in which either no activities occur or exactly one does we see that there exist only M + 1 distinct configurations at time t that can lead to the state **X** at time t + dt.

Stochastic: Infinitesimal time interval

By considering a discrete infinitesimal time interval (t, t + dt) in which either no activities occur or exactly one does we see that there exist only M + 1 distinct configurations at time t that can lead to the state **X** at time t + dt.

$$Pr(\mathbf{X}; t + dt)$$

= Pr(**X**; t) Pr(no state change over dt)

Stephen Gilmore. LFCS, University of Edinburgh.

▲ @ ▶ ▲ ∃ ▶ ▲

Stochastic: Infinitesimal time interval

By considering a discrete infinitesimal time interval (t, t + dt) in which either no activities occur or exactly one does we see that there exist only M + 1 distinct configurations at time t that can lead to the state **X** at time t + dt.

$$\begin{aligned} \mathsf{Pr}(\mathbf{X}; t + \mathsf{d}t) \\ &= \mathsf{Pr}(\mathbf{X}; t) \, \mathsf{Pr}(\mathsf{no \ state \ change \ over \ d}t) \\ &+ \sum_{\mu=1}^{M} \mathsf{Pr}(\mathbf{X} - \mathbf{v}_{\mu}; t) \, \mathsf{Pr}(\mathsf{state \ change \ to \ }\mathbf{X} \ \mathsf{over \ d}t) \end{aligned}$$

Stephen Gilmore. LFCS, University of Edinburgh.

Stochastic: Infinitesimal time interval

By considering a discrete infinitesimal time interval (t, t + dt) in which either no activities occur or exactly one does we see that there exist only M + 1 distinct configurations at time t that can lead to the state **X** at time t + dt.

$$\begin{aligned} &\mathsf{Pr}(\mathbf{X}; t + \mathsf{d}t) \\ &= \mathsf{Pr}(\mathbf{X}; t) \, \mathsf{Pr}(\mathsf{no \ state \ change \ over \ d}t) \\ &+ \sum_{\mu=1}^{M} \mathsf{Pr}(\mathbf{X} - \mathbf{v}_{\mu}; t) \, \mathsf{Pr}(\mathsf{state \ change \ to \ }\mathbf{X} \ \mathsf{over \ d}t) \end{aligned}$$

where \mathbf{v}_{μ} is a vector defining the result of reaction μ on state vector \mathbf{X} , i.e. $\mathbf{X} \to \mathbf{X} + \mathbf{v}_{\mu}$ after an occurrence of reaction μ .

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (三) (三)

2

Stochastic: State change probabilities

Pr(no state change over dt)

$$1-\sum_{\mu=1}^{M}a_{\mu}(\mathbf{X})\mathsf{d}t$$

Stephen Gilmore. LFCS, University of Edinburgh.
Stochastic: State change probabilities

Pr(no state change over dt)

$$1-\sum_{\mu=1}^{M}a_{\mu}(\mathbf{X})\mathsf{d}t$$

Pr(state change to X over dt)

$$\sum_{\mu=1}^{M} \mathsf{Pr}(\mathbf{X} - \mathbf{v}_{\mu}; t) a_{\mu}(\mathbf{X} - \mathbf{v}_{\mu}) \mathsf{d}t$$

Stephen Gilmore. LFCS, University of Edinburgh.

Learning from other disciplines

Computing is not the only discipline where problems of scale and multi-scale occur. A recent trend has been to borrow simulation methods from computational biology and computational chemistry.

Stochastic simulation algorithms

Gillespie's Stochastic Simulation Algorithm (SSA) is essentially an exact procedure for numerically simulating the time evolution of a well-stirred chemically reacting system by taking proper account of the randomness inherent in such a system.

Stochastic simulation algorithms

Gillespie's Stochastic Simulation Algorithm (SSA) is essentially an exact procedure for numerically simulating the time evolution of a well-stirred chemically reacting system by taking proper account of the randomness inherent in such a system.

It is rigorously based on the same microphysical premise that underlies the chemical master equation and gives a more realistic representation of a system's evolution than the deterministic reaction rate equation (RRE) represented mathematically by ODEs.

(日) (同) (三) (三)

3

Gillespie's exact SSA (1977)

The algorithm takes time steps of variable length, based on the rate constants and population size of each chemical species.

Stephen Gilmore. LFCS, University of Edinburgh.

・ロト ・ 同ト ・ ヨト ・ ヨ

Gillespie's exact SSA (1977)

- The algorithm takes time steps of variable length, based on the rate constants and population size of each chemical species.
- The probability of one reaction occurring relative to another is dictated by their relative propensity functions.

Gillespie's exact SSA (1977)

- The algorithm takes time steps of variable length, based on the rate constants and population size of each chemical species.
- The probability of one reaction occurring relative to another is dictated by their relative propensity functions.
- One random variable is then used to choose which reaction will occur, and another random variable determines how long the step will last.

Gillespie's exact SSA (1977)

- The algorithm takes time steps of variable length, based on the rate constants and population size of each chemical species.
- The probability of one reaction occurring relative to another is dictated by their relative propensity functions.
- One random variable is then used to choose which reaction will occur, and another random variable determines how long the step will last.
- The chemical populations are altered according to the definition of the reaction and the process is repeated.

Gillespie's SSA is a Monte Carlo Markov Chain simulation

The SSA is a Monte Carlo type method. With the SSA one may approximate any variable of interest by generating many trajectories and observing the statistics of the values of the variable. Since many trajectories are needed to obtain a reasonable approximation, the efficiency of the SSA is of critical importance.

Computational cost of Gillespie's exact algorithm

The cost of this detailed stochastic simulation algorithm is the likely large amounts of computing time. The key issue is that the time step for the next reaction can be very small indeed if we are to guarantee that only one reaction can take place in a given time interval.

Computational cost of Gillespie's exact algorithm

The cost of this detailed stochastic simulation algorithm is the likely large amounts of computing time. The key issue is that the time step for the next reaction can be very small indeed if we are to guarantee that only one reaction can take place in a given time interval.

Increasing the molecular population or number of reaction mechanisms necessarily requires a corresponding decrease in the time interval. The SSA can be very computationally inefficient especially when there are large numbers of molecules or the propensity functions are large.

(日) (周) (三) (三)

Gillespie's tau-leap method (2001)

Gillespie proposed two new methods, namely the τ -leap method and the midpoint τ -leap method in order to improve the efficiency of the SSA while maintaining acceptable losses in accuracy.

Daniel T. Gillespie.

Approximate accelerated stochastic simulation of chemically reacting systems.

J. Comp. Phys., 115(4):1716–1733, 2001.

Gillespie's tau-leap method (2001)

Gillespie proposed two new methods, namely the τ -leap method and the midpoint τ -leap method in order to improve the efficiency of the SSA while maintaining acceptable losses in accuracy.

Daniel T. Gillespie.

Approximate accelerated stochastic simulation of chemically reacting systems.

J. Comp. Phys., 115(4):1716–1733, 2001.

The key idea here is to take a larger time step and allow for more reactions to take place in that step, but under the proviso that the propensity functions do not change too much in that interval. By means of a Poisson approximation, the tau-leaping method can "leap over" many reactions.

Stephen Gilmore. LFCS, University of Edinburgh.

Gillespie's tau-leap method (drawback)

The use of approximation in Poisson methods leads to the possibility of negative molecular numbers being predicted — something with no physical explanation.

Gillespie's Modified Poisson tau-leap methods (2005)

Gillespie's modified Poisson tau-leaping method introduces a second control parameter whose value dials the procedure from the original Poisson tau-leaping method at one extreme to the exact SSA at the other.

Stephen Gilmore. LFCS, University of Edinburgh.

(日) (同) (日) (日)

Gillespie's Modified Poisson tau-leap methods (2005)

Gillespie's modified Poisson tau-leaping method introduces a second control parameter whose value dials the procedure from the original Poisson tau-leaping method at one extreme to the exact SSA at the other.

Any reaction channel with a positive propensity function which is within n_c firings of exhausting its reactants is termed a *critical* reaction.

Y. Cao, D. Gillespie, and L. Petzold.

Avoiding negative populations in explicit tau leaping.

J. Chem. Phys, 123(054104), 2005.

<ロト </p>

Gillespie's Modified Poisson tau-leap methods (2006)

The modified algorithm chooses τ in such a way that no more than *one* firing of *all* the critical reactions can occur during the leap. The probability of producing a negative population is reduced to nearly zero.

Stephen Gilmore. LFCS, University of Edinburgh.

Gillespie's Modified Poisson tau-leap methods (2006)

The modified algorithm chooses τ in such a way that no more than *one* firing of *all* the critical reactions can occur during the leap. The probability of producing a negative population is reduced to nearly zero.

If a negative population *does* occur the leap can simply be rejected and repeated with τ reduced by half, or the entire simulation can be abandoned and repeated for larger n_c .

Y. Cao, D. Gillespie, and L. Petzold.

Efficient stepsize selection for the tau-leaping simulation method.

J. Chem. Phys, 2006, 124(4):044109.

イロト イヨト イヨト イヨト

2

Family of stochastic simulation algorithms

FASTEST, BEST	
Discrete, exact	Continuous, approximate
	Modified Poisson $ au$ leap (2005)
	au leap (2001)
Logarithmic Direct Method (2006)	
Sorting Direct Method (2005)	
Optimised Direct Method (2004)	
Next Reaction Method (2000)	
Direct Method (1977)	
First Reaction Method (1977)	
SLOWEST, WORST	

Stephen Gilmore. LFCS, University of Edinburgh.

Conclusions

- Stochastic methods can be used to investigate the temporal and behavioural nature of global computing systems.
- Numerical solution of Markov chains gives a thorough stochastic treatment of models of modest size.
- Differential equation methods scale well to increasing population sizes allowing global computing systems to be addressed.
- Stochastic simulation algorithms make a bridge between the two approaches: there has been a recent explosion of interest in the subject with many new variants of Gillespie's algorithm.