

Process algebras for quantitative analysis: Lecture 5 — Collective dynamics

Jane Hillston

School of Informatics
The University of Edinburgh
Scotland

12th February 2010

Outline


- 1 Introduction
 - Collective Dynamics
- 2 Continuous Approximation
- 3 Fluid-Flow Semantics
- 4 Case studies
 - Internet worms
 - Scalable Web Services
- 5 Final remarks

Outline

- 1 Introduction
 - Collective Dynamics
- 2 Continuous Approximation
- 3 Fluid-Flow Semantics
- 4 Case studies
 - Internet worms
 - Scalable Web Services
- 5 Final remarks


Collective Dynamics

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



Collective Dynamics

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



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

Collective Behaviour

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



Process Algebra and Collective Dynamics

Process algebra are well-suited to modelling such systems

Process Algebra and Collective Dynamics

Process algebra are well-suited to modelling such systems

- Developed to represent concurrent behaviour compositionally;

Process Algebra and Collective Dynamics

- Process algebra are well-suited to modelling such systems
- Developed to represent concurrent behaviour compositionally;
 - Capture the interactions between individuals explicitly;

Process Algebra and Collective Dynamics

- Process algebra are well-suited to modelling such systems
- Developed to represent concurrent behaviour compositionally;
 - Capture the interactions between individuals explicitly;
 - Incorporate formal apparatus for reasoning about the behaviour of systems;

Process Algebra and Collective Dynamics

Process algebra are well-suited to modelling such systems

- Developed to represent concurrent behaviour compositionally;
- Capture the interactions between individuals explicitly;
- Incorporate formal apparatus for reasoning about the behaviour of systems;
- Stochastic extensions, such as PEPA, enable quantified behaviour of the dynamics of systems.

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

Performance as an emergent behaviour

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

Performance as an emergent behaviour

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

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

Performance as an emergent behaviour

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

Example Service Level Agreement

90% of requests receive a response within 3 seconds.

Performance as an emergent behaviour

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

Example Service Level Agreement

90% of requests receive a response within 3 seconds.

Qualitative Service Level Agreement

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

Novelty

The novelty of the CODA project has been twofold:

Novelty

The novelty of the CODA project has been twofold:

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

Novelty

The novelty of the CODA project has been twofold:


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

Biochemical signalling pathways

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

Stochastic Process Algebra

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



A simple example: processors and resources

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

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

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


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

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

A simple example: processors and resources

$$\begin{aligned}
 Proc_0 &\stackrel{\text{def}}{=} (task1, r_1).Proc_1 \\
 Proc_1 &\stackrel{\text{def}}{=} (task2, r_2).Proc_0 \\
 Res_0 &\stackrel{\text{def}}{=} (task1, r_3).Res_1 \\
 Res_1 &\stackrel{\text{def}}{=} (reset, r_4).Res_0
 \end{aligned}$$

$$Proc_0 \bowtie_{\{task1\}} Res_0$$




A simple example: processors and resources

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

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

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

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



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

Deriving quantitative data

PEPA models can be analysed for quantified dynamic behaviour in a number of different ways.

Deriving quantitative data

PEPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a [Markov Process \(CTMC\)](#).

Deriving quantitative data

PEPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a [Markov Process \(CTMC\)](#).



Deriving quantitative data

PEPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language also may be used to generate a [stochastic simulation](#).

Deriving quantitative data

PEPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language also may be used to generate a [stochastic simulation](#).



Deriving quantitative data

PEPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a **system of ordinary differential equations (ODEs)**.

Deriving quantitative data

PEPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a **system of ordinary differential equations (ODEs)**.



Deriving quantitative data

PEPA models can be analysed for quantified dynamic behaviour in a number of different ways.

The language may be used to generate a **system of ordinary differential equations (ODEs)**.



Each of these has tool support so that the underlying model is derived automatically according to the predefined rules.

Outline

- 1 Introduction
 - Collective Dynamics
- 2 Continuous Approximation**
- 3 Fluid-Flow Semantics
- 4 Case studies
 - Internet worms
 - Scalable Web Services
- 5 Final remarks

Solving discrete state models

Under the SOS semantics a PEPA model is mapped to a **Continuous Time Markov Chain (CTMC)** with global states determined by the local states of all the participating components.

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

Solving discrete state models

Under the SOS semantics a PEPA model is mapped to a **Continuous Time Markov Chain (CTMC)** with global states determined by the local states of all the participating components.

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

Alternatively they may be studied using **stochastic simulation**. Each run generates a single trajectory through the state space. Many runs are needed in order to obtain average behaviours.

Solving discrete state models

Under the SOS semantics a PEPA model is mapped to a **Continuous Time Markov Chain (CTMC)** with global states determined by the local states of all the participating components.

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

Alternatively they may be studied using **stochastic simulation**. Each run generates a single trajectory through the state space. Many runs are needed in order to obtain average behaviours.

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

Continuous Approximation

The major limitation of the CTMC approach is the state space explosion problem.

Continuous Approximation

The major limitation of the CTMC approach is the **state space explosion** problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

Continuous Approximation

The major limitation of the CTMC approach is the **state space explosion** problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

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

Continuous Approximation

The major limitation of the CTMC approach is the **state space explosion** problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

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



Continuous Approximation

The major limitation of the CTMC approach is the **state space explosion** problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

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




Continuous Approximation

The major limitation of the CTMC approach is the **state space explosion** problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

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



Continuous Approximation

The major limitation of the CTMC approach is the **state space explosion** problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

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

-
-
-
-
-
-
-
-
-

Continuous Approximation

The major limitation of the CTMC approach is the **state space explosion** problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

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




Continuous Approximation

The major limitation of the CTMC approach is the **state space explosion** problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

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



Continuous Approximation

The major limitation of the CTMC approach is the state space explosion problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

Use continuous state variables to approximate the discrete state space.



Continuous Approximation

The major limitation of the CTMC approach is the **state space explosion** problem.

State space explosion becomes an ever more challenging problem as the scale and complexity of modern systems increase.

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



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

New mathematical structures: differential equations

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

New mathematical structures: differential equations

- Use a more abstract state representation rather than the CTMC complete state space.
- Assume that these state variables are subject to continuous rather than discrete change.

New mathematical structures: differential equations

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

New mathematical structures: differential equations

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

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

Simple example revisited [QEST'05]

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

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

$$Res_0 \stackrel{def}{=} (task1, r_1).Res_1$$

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

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

Simple example revisited [QEST'05]

$$\begin{aligned}
 Proc_0 &\stackrel{\text{def}}{=} (task_1, r_1).Proc_1 \\
 Proc_1 &\stackrel{\text{def}}{=} (task_2, r_2).Proc_0 \\
 Res_0 &\stackrel{\text{def}}{=} (task_1, r_1).Res_1 \\
 Res_1 &\stackrel{\text{def}}{=} (reset, r_4).Res_0
 \end{aligned}$$

$$Proc_0[N_P] \boxtimes_{\{task_1\}} Res_0[N_R]$$

CTMC interpretation

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

Simple example revisited [QEST'05]

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

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

$$Res_0 \stackrel{def}{=} (task1, r_1).Res_1$$

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

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

ODE interpretation

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

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

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

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

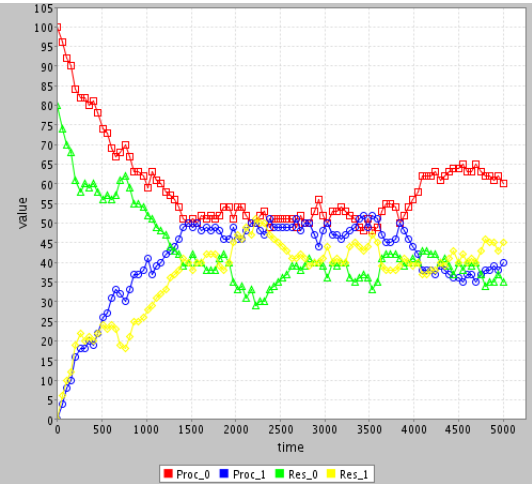
$$\frac{dx_3}{dt} = -r_1 \min(x_1, x_3) + r_4 x_4$$

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


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

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


100 processors and 80 resources (simulation run A)




100 processors and 80 resources (simulation run D)




100 Processors and 80 resources (average of 100 runs)



100 processors and 80 resources (average of 1000 runs)



100 processors and 80 resources (ODE solution)



Outline

- 1 Introduction
 - Collective Dynamics
- 2 Continuous Approximation
- 3 Fluid-Flow Semantics**
- 4 Case studies
 - Internet worms
 - Scalable Web Services
- 5 Final remarks

Deriving a Fluid Approximation of a PEPA model

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

Deriving a Fluid Approximation of a PEPA model

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

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

Deriving a Fluid Approximation of a PEPA model

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

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

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

Fluid Structured Operational Semantics

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

Fluid Structured Operational Semantics

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

- 1 Remove excess components (*Context Reduction*)

Fluid Structured Operational Semantics

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

- ① Remove excess components (*Context Reduction*)
- ② Collect the transitions of the reduced context (*Jump Multiset*)

Fluid Structured Operational Semantics

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

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

Fluid Structured Operational Semantics

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

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

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

Context Reduction

$$\begin{aligned} Proc_0 &\stackrel{def}{=} (task1, r_1).Proc_1 \\ Proc_1 &\stackrel{def}{=} (task2, r_2).Proc_0 \\ Res_0 &\stackrel{def}{=} (task1, r_3).Res_1 \\ Res_1 &\stackrel{def}{=} (reset, r_4).Res_0 \\ System &\stackrel{def}{=} Proc_0[N_P] \bowtie_{\{transfer\}} Res_0[N_R] \end{aligned}$$



$$\mathcal{R}(System) = \{Proc_0, Proc_1\} \bowtie_{\{task1\}} \{Res_0, Res_1\}$$

Context Reduction

$$\begin{aligned} Proc_0 &\stackrel{def}{=} (task1, r_1).Proc_1 \\ Proc_1 &\stackrel{def}{=} (task2, r_2).Proc_0 \\ Res_0 &\stackrel{def}{=} (task1, r_3).Res_1 \\ Res_1 &\stackrel{def}{=} (reset, r_4).Res_0 \\ System &\stackrel{def}{=} Proc_0[N_P] \boxtimes_{\{transfer\}} Res_0[N_R] \end{aligned}$$

$$\Downarrow$$

$$\mathcal{R}(System) = \{Proc_0, Proc_1\} \boxtimes_{\{task1\}} \{Res_0, Res_1\}$$

Population Vector

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

Location Dependency

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

Location Dependency

$$\text{System} \stackrel{\text{def}}{=} Proc_0[N'_C] \underset{\{task1\}}{\boxtimes} Res_0[N_S] \parallel Proc_0[N''_C]$$



$$\{Proc_0, Proc_1\} \underset{\{task1\}}{\boxtimes} \{Res_0, Res_1\} \parallel \{Proc_0, Proc_1\}$$

Population Vector

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

Fluid Structured Operational Semantics by Example

$$\begin{aligned}
 Proc_0 &\stackrel{def}{=} (task1, r_1).Proc_1 \\
 Proc_1 &\stackrel{def}{=} (task2, r_2).Proc_0 \\
 Res_0 &\stackrel{def}{=} (task1, r_3).Res_1 \\
 Res_1 &\stackrel{def}{=} (reset, r_4).Res_0 \\
 System &\stackrel{def}{=} Proc_0[N_P] \begin{array}{c} \boxtimes \\ \{transfer\} \end{array} Res_0[N_R] \\
 &\xi = (\xi_1, \xi_2, \xi_3, \xi_4)
 \end{aligned}$$

$$\frac{Proc_0 \xrightarrow{task1, r_1} Proc_1}{Proc_0 \xrightarrow{task1, r_1 \xi_1} *_ Proc_1}$$

Fluid Structured Operational Semantics by Example

$$\begin{aligned} Proc_0 &\stackrel{def}{=} (task1, r_1).Proc_1 \\ Proc_1 &\stackrel{def}{=} (task2, r_2).Proc_0 \\ Res_0 &\stackrel{def}{=} (task1, r_3).Res_1 \\ Res_1 &\stackrel{def}{=} (reset, r_4).Res_0 \\ System &\stackrel{def}{=} Proc_0[N_P] \underset{\{transfer\}}{\boxtimes} Res_0[N_R] \\ &\xi = (\xi_1, \xi_2, \xi_3, \xi_4) \end{aligned}$$

$$\begin{array}{c} Proc_0 \xrightarrow{task1, r_1} Proc_1 \\ \hline Proc_0 \xrightarrow{task1, r_1 \xi_1} *_ Proc_1 \end{array} \qquad \begin{array}{c} Res_0 \xrightarrow{task1, r_3} Res_1 \\ \hline Res_0 \xrightarrow{task1, r_3 \xi_3} *_ Res_1 \end{array}$$

Fluid Structured Operational Semantics by Example

$$\begin{aligned}
 Proc_0 &\stackrel{def}{=} (task1, r_1).Proc_1 \\
 Proc_1 &\stackrel{def}{=} (task2, r_2).Proc_0 \\
 Res_0 &\stackrel{def}{=} (task1, r_3).Res_1 \\
 Res_1 &\stackrel{def}{=} (reset, r_4).Res_0 \\
 System &\stackrel{def}{=} Proc_0[N_P] \underset{\{transfer\}}{\boxtimes} Res_0[N_R] \\
 &\xi = (\xi_1, \xi_2, \xi_3, \xi_4)
 \end{aligned}$$

$$\frac{
 \frac{Proc_0 \xrightarrow{task1, r_1} Proc_1}{Proc_0 \xrightarrow{task1, r_1 \xi_1} *_ Proc_1} \quad
 \frac{Res_0 \xrightarrow{task1, r_3} Res_1}{Res_0 \xrightarrow{task1, r_3 \xi_3} *_ Res_1}
 }{
 Proc_0 \underset{\{task1\}}{\boxtimes} Res_0 \xrightarrow{task1, r(\xi)} *_ Proc_1 \underset{\{task1\}}{\boxtimes} Res_1
 }$$

Apparent Rate Calculation


$$\frac{\frac{Proc_0 \xrightarrow{task1, r_1} Proc_1}{Proc_0 \xrightarrow{task1, r_1 \xi_1} *_ Proc_1} \quad \frac{Res_0 \xrightarrow{task1, r_3} Res_1}{Res_0 \xrightarrow{task1, r_3 \xi_3} *_ Res_1}}{Proc_0 \boxtimes_{\{task1\}} Res_0 \xrightarrow{task1, r(\xi)} *_ Proc_1 \boxtimes_{\{task1\}} Res_1}$$

Apparent Rate Calculation

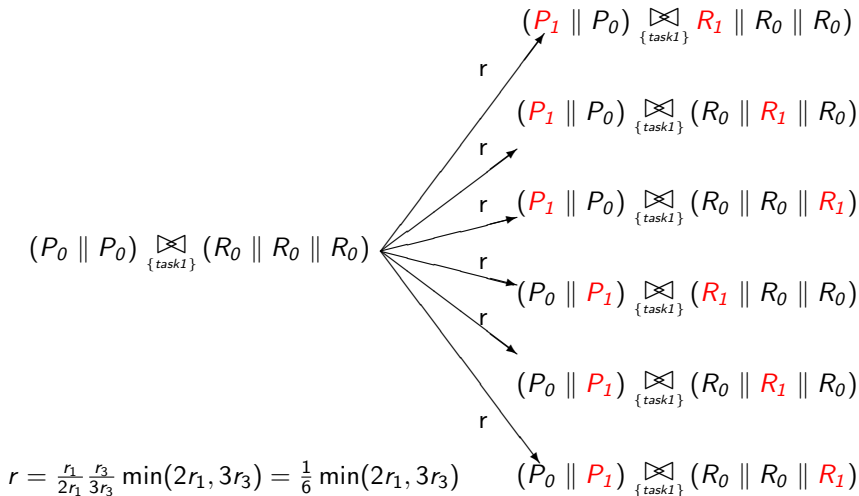
$$\frac{\frac{Proc_0 \xrightarrow{task1, r_1} Proc_1}{Proc_0 \xrightarrow{task1, r_1 \xi_1} *_ Proc_1} \quad \frac{Res_0 \xrightarrow{task1, r_3} Res_1}{Res_0 \xrightarrow{task1, r_3 \xi_3} *_ Res_1}}{Proc_0 \boxtimes_{\{task1\}} Res_0 \xrightarrow{task1, r(\xi)} *_ Proc_1 \boxtimes_{\{task1\}} Res_1}$$

$$\begin{aligned} r(\xi) &= \frac{r_1 \xi_1}{r_{task1}^*(Proc_0, \xi)} \frac{r_3 \xi_4}{r_{task1}^*(Res_0, \xi)} \min(r_{task1}^*(Proc_0, \xi), r_{task1}^*(Res_0, \xi)) \\ &= \frac{r_1 \xi_1}{r_1 \xi_1} \frac{r_3 \xi_4}{r_3 \xi_4} \min(r_1 \xi_1, r_3 \xi_4) \\ &= \min(r_1 \xi_1, r_3 \xi_4) \end{aligned}$$


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



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



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



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

$$(2, 0, 3, 0) \xrightarrow{\min(2r_1, 3r_3)} (1, 1, 2, 1)$$

$$(P_0 \parallel P_0) \bowtie_{\{task1\}} (R_0 \parallel R_0 \parallel R_0)$$

$$(P_1 \parallel P_0) \bowtie_{\{task1\}} (R_1 \parallel R_0 \parallel R_0)$$

$$(P_1 \parallel P_0) \bowtie_{\{task1\}} (R_0 \parallel R_1 \parallel R_0)$$

$$(P_1 \parallel P_0) \bowtie_{\{task1\}} (R_0 \parallel R_0 \parallel R_1)$$

$$(P_0 \parallel P_1) \bowtie_{\{task1\}} (R_1 \parallel R_0 \parallel R_0)$$

$$(P_0 \parallel P_1) \bowtie_{\{task1\}} (R_0 \parallel R_1 \parallel R_0)$$

$$(P_0 \parallel P_1) \bowtie_{\{task1\}} (R_0 \parallel R_0 \parallel R_1)$$

$$r = \frac{r_1}{2r_1} \frac{r_3}{3r_3} \min(2r_1, 3r_3) = \frac{1}{6} \min(2r_1, 3r_3)$$

Jump Multiset

$$Proc_0 \underset{\{task1\}}{\boxtimes} Res_0 \xrightarrow{task1, r(\xi)} * Proc_1 \underset{\{task1\}}{\boxtimes} Res_1$$
$$r(\xi) = \min(r_1 \xi_1, r_3 \xi_3)$$

Jump Multiset

$$Proc_0 \underset{\{task1\}}{\boxtimes} Res_0 \xrightarrow{task1, r(\xi)} * Proc_1 \underset{\{task1\}}{\boxtimes} Res_1$$
$$r(\xi) = \min(r_1 \xi_1, r_3 \xi_3)$$

$$Proc_1 \underset{\{task1\}}{\boxtimes} Res_0 \xrightarrow{task2, \xi_2 r_2} * Proc_0 \underset{\{task1\}}{\boxtimes} Res_0$$

Jump Multiset

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

$$r(\xi) = \min(r_1 \xi_1, r_3 \xi_3)$$

$$Proc_1 \underset{\{task1\}}{\boxtimes} Res_0 \xrightarrow{task2, \xi_2 r_2}_* Proc_0 \underset{\{task1\}}{\boxtimes} Res_0$$

$$Proc_0 \underset{\{task1\}}{\boxtimes} Res_1 \xrightarrow{reset, \xi_4 r_4}_* Proc_0 \underset{\{task1\}}{\boxtimes} Res_0$$

Equivalent Transitions

Some transitions may give the same information:

$$\begin{array}{ccc}
 Proc_0 \quad \boxtimes_{\{task1\}} \quad Res_1 & \xrightarrow{reset, \xi_4 r_4} * & Proc_0 \quad \boxtimes_{\{task1\}} \quad Res_0 \\
 Proc_1 \quad \boxtimes_{\{task1\}} \quad Res_1 & \xrightarrow{reset, \xi_4 r_4} * & Proc_1 \quad \boxtimes_{\{task1\}} \quad Res_0
 \end{array}$$

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

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

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

$$Proc_0 \begin{array}{c} \boxtimes \\ \{task1\} \end{array} Res_1 \xrightarrow{reset, \xi_4 r_4} *_ Proc_0 \begin{array}{c} \boxtimes \\ \{task1\} \end{array} Res_0$$

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

$$Proc_0 \begin{array}{c} \boxtimes \\ \{task1\} \end{array} Res_1 \xrightarrow{reset, \xi_4 r_4} *_ Proc_0 \begin{array}{c} \boxtimes \\ \{task1\} \end{array} Res_0$$

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

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

$$Proc_0 \begin{array}{c} \boxtimes \\ \{task1\} \end{array} Res_1 \xrightarrow{reset, \xi_4 r_4} *_ Proc_0 \begin{array}{c} \boxtimes \\ \{task1\} \end{array} Res_0$$

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

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

- Add $+1$ to all elements of l corresponding to the indices of the components in the rhs of the transition

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

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

$$Proc_0 \begin{array}{c} \boxtimes \\ \{task1\} \end{array} Res_1 \xrightarrow{reset, \xi_4 r_4} *_ Proc_0 \begin{array}{c} \boxtimes \\ \{task1\} \end{array} Res_0$$

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

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

- Add $+1$ to all elements of l corresponding to the indices of the components in the rhs of the transition

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

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

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

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

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

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

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

$$\begin{array}{ccc}
 Proc_0 \boxtimes_{\{task1\}} Res_0 & \xrightarrow{task1, r(\xi)}_* & Proc_1 \boxtimes_{\{task1\}} Res_1 \\
 Proc_1 \boxtimes_{\{task1\}} Res_0 & \xrightarrow{task2, \xi_2 r'_2}_* & Proc_0 \boxtimes_{\{task1\}} Res_0
 \end{array}$$

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

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

$$\begin{array}{ccc}
 Proc_0 \boxtimes_{\{task1\}} Res_0 & \xrightarrow{task1, r(\xi)}_* & Proc_1 \boxtimes_{\{task1\}} Res_1 \\
 Proc_1 \boxtimes_{\{task1\}} Res_0 & \xrightarrow{task2, \xi_2 r'_2}_* & Proc_0 \boxtimes_{\{task1\}} Res_0 \\
 Proc_0 \boxtimes_{\{task1\}} Res_1 & \xrightarrow{reset, \xi_4 r_4}_* & Proc_0 \boxtimes_{\{task1\}} Res_0
 \end{array}$$

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

$$f(\xi, (+1, -1, 0, 0), task2) = \xi_2 r_2$$

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

Capturing behaviour in the Generator Function

$$\begin{aligned} Proc_0 &\stackrel{\text{def}}{=} (task1, r_1).Proc_1 \\ Proc_1 &\stackrel{\text{def}}{=} (task2, r_2).Proc_0 \\ Res_0 &\stackrel{\text{def}}{=} (task1, r_3).Res_1 \\ Res_1 &\stackrel{\text{def}}{=} (reset, r_4).Res_0 \\ System &\stackrel{\text{def}}{=} Proc_0[N_P] \bowtie_{\{transfer\}} Res_0[N_R] \end{aligned}$$

Capturing behaviour in the Generator Function

$$\begin{aligned} Proc_0 &\stackrel{\text{def}}{=} (task1, r_1).Proc_1 \\ Proc_1 &\stackrel{\text{def}}{=} (task2, r_2).Proc_0 \\ Res_0 &\stackrel{\text{def}}{=} (task1, r_3).Res_1 \\ Res_1 &\stackrel{\text{def}}{=} (reset, r_4).Res_0 \\ System &\stackrel{\text{def}}{=} Proc_0[N_P] \underset{\{transfer\}}{\boxtimes} Res_0[N_R] \end{aligned}$$

Numerical Vector Form

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

Capturing behaviour in the Generator Function

$$\begin{aligned}
 Proc_0 &\stackrel{def}{=} (task1, r_1).Proc_1 \\
 Proc_1 &\stackrel{def}{=} (task2, r_2).Proc_0 \\
 Res_0 &\stackrel{def}{=} (task1, r_3).Res_1 \\
 Res_1 &\stackrel{def}{=} (reset, r_4).Res_0 \\
 System &\stackrel{def}{=} Proc_0[N_P] \bowtie_{\{transfer\}} Res_0[N_R]
 \end{aligned}$$

Numerical Vector Form

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

Generator Function

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

Extraction of the ODE from f

Generator Function

$$f(\xi, (-1, 1, -1, 1), task1) = \min(r_1\xi_1, r_3\xi_3)$$

$$f(\xi, (1, -1, 0, 0), task2) = r_2\xi_2$$

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

Differential Equation

$$\frac{dx}{dt} = F_{\mathcal{M}}(x) = \sum_{l \in \mathbb{Z}^d} l \sum_{\alpha \in \mathcal{A}} f(x, l, \alpha)$$

$$= (-1, 1, -1, 1) \min(r_1x_1, r_3x_3) + (1, -1, 0, 0)r_2x_2 \\ + (0, 0, 1, -1)r_4x_4$$

Extraction of the ODE from f

Generator Function

$$f(\xi, (-1, 1, -1, 1), \text{task1}) = \min(r_1\xi_1, r_3\xi_3)$$

$$f(\xi, (1, -1, 0, 0), \text{task2}) = r_2\xi_2$$

$$f(\xi, (0, 0, 1, -1), \text{reset}) = r_4\xi_4$$

Differential Equation

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

$$\frac{dx_2}{dt} = \min(r_1x_1, r_3x_3) - r_2x_2$$

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

$$\frac{dx_4}{dt} = \min(r_1x_1, r_3x_3) - r_4x_4$$

Outline

- 1 Introduction
 - Collective Dynamics
- 2 Continuous Approximation
- 3 Fluid-Flow Semantics
- 4 Case studies**
 - Internet worms
 - Scalable Web Services
- 5 Final remarks

Internet worms: Background

- Internet worms are malicious programs that exploit operating system security weaknesses to propagate themselves.

Internet worms: Background

- Internet worms are malicious programs that exploit operating system security weaknesses to propagate themselves.
- While the security flaws go unpatched, the worm spreads epidemic-like and uses large amounts of available bandwidth.

Internet worms: Background

- Internet worms are malicious programs that exploit operating system security weaknesses to propagate themselves.
- While the security flaws go unpatched, the worm spreads epidemic-like and uses large amounts of available bandwidth.
- Far more destructive is the worms' effect on the Internet routing infrastructure, as the worms tend to overload the connecting routers with nonexistent IP lookups.

Internet worms: Background

- Internet worms are malicious programs that exploit operating system security weaknesses to propagate themselves.
- While the security flaws go unpatched, the worm spreads epidemic-like and uses large amounts of available bandwidth.
- Far more destructive is the worms' effect on the Internet routing infrastructure, as the worms tend to overload the connecting routers with nonexistent IP lookups.
- Worms like Nimbda, Slammer, Code Red, Sasser and Code Red 2 have caused the Internet to become unusable for many hours at a time until security patches could be applied and routers fixed.

Internet worms: Background

- Internet worms are malicious programs that exploit operating system security weaknesses to propagate themselves.
- While the security flaws go unpatched, the worm spreads epidemic-like and uses large amounts of available bandwidth.
- Far more destructive is the worms' effect on the Internet routing infrastructure, as the worms tend to overload the connecting routers with nonexistent IP lookups.
- Worms like Nimbda, Slammer, Code Red, Sasser and Code Red 2 have caused the Internet to become unusable for many hours at a time until security patches could be applied and routers fixed.
- The estimated cost of computer worms and related activities is about \$50 billion a year.

An Internet-scale Problem

- We wish to study the emergent behaviour of Internet worms as they spread to thousands and then hundreds-of-thousands of hosts.

An Internet-scale Problem

- We wish to study the emergent behaviour of Internet worms as they spread to thousands and then hundreds-of-thousands of hosts.
- Explicit state-based methods for calculating steady-state, transient or passage-time measures are limited to state-spaces of the order of 10^9 .

An Internet-scale Problem

- We wish to study the emergent behaviour of Internet worms as they spread to thousands and then hundreds-of-thousands of hosts.
- Explicit state-based methods for calculating steady-state, transient or passage-time measures are limited to state-spaces of the order of 10^9 .
- By transforming our stochastic process algebra model into a set of ODEs, we can obtain a plot of model behaviour against time for models with global state spaces in excess of 10^{10000} states.

Susceptible-Infective-Removed (SIR) model

- We apply a version of an SIR model of infection to various computer worm attack models.

Susceptible-Infective-Removed (SIR) model

- We apply a version of an SIR model of infection to various computer worm attack models.
- An SIR model explicitly represents the total number of susceptible, infective and removed hosts in a system and is more commonly used to model disease epidemics.

Susceptible-Infective-Removed (SIR) model

- We apply a version of an SIR model of infection to various computer worm attack models.
- An SIR model explicitly represents the total number of susceptible, infective and removed hosts in a system and is more commonly used to model disease epidemics.

$$\frac{ds(t)}{dt} = -\beta s(t) i(t)$$

Susceptible-Infective-Removed (SIR) model

- We apply a version of an SIR model of infection to various computer worm attack models.
- An SIR model explicitly represents the total number of **susceptible**, **infective** and **removed** hosts in a system and is more commonly used to model disease epidemics.

$$\frac{ds(t)}{dt} = -\beta s(t) i(t)$$

$$\frac{di(t)}{dt} = \beta s(t) i(t) - \gamma i(t)$$

Susceptible-Infective-Removed (SIR) model

- We apply a version of an SIR model of infection to various computer worm attack models.
- An SIR model explicitly represents the total number of susceptible, infective and removed hosts in a system and is more commonly used to model disease epidemics.

$$\frac{ds(t)}{dt} = -\beta s(t) i(t)$$

$$\frac{di(t)}{dt} = \beta s(t) i(t) - \gamma i(t)$$

$$\frac{dr(t)}{dt} = \gamma i(t)$$

Susceptible-Infective-Removed over a network

- This is our most basic infection model and is used to verify that we get recognisable qualitative results.

Susceptible-Infective-Removed over a network

- This is our most basic infection model and is used to verify that we get recognisable qualitative results.
- Initially, there are N susceptible computers and one infected computer.

Susceptible-Infective-Removed over a network

- This is our most basic infection model and is used to verify that we get recognisable qualitative results.
- Initially, there are N susceptible computers and one infected computer.
- As the system evolves more susceptible computers become infected from the growing infective population.

Susceptible-Infective-Removed over a network

- This is our most basic infection model and is used to verify that we get recognisable qualitative results.
- Initially, there are N susceptible computers and one infected computer.
- As the system evolves more susceptible computers become infected from the growing infective population.
- An infected computer can be patched so that it is no longer infected or susceptible to infection.

Susceptible-Infective-Removed over a network

- This is our most basic infection model and is used to verify that we get recognisable qualitative results.
- Initially, there are N susceptible computers and one infected computer.
- As the system evolves more susceptible computers become infected from the growing infective population.
- An infected computer can be patched so that it is no longer infected or susceptible to infection.
- This state is termed **removed** and is an absorbing state for that component in the system.

Susceptible-Infective-Removed over a network

- The capacity of the network is dictated by the parameter M , the number of concurrent, independent connections that the network can sustain.

Susceptible-Infective-Removed over a network

- The capacity of the network is dictated by the parameter M , the number of concurrent, independent connections that the network can sustain.
- Additionally, an attempted network connection can fail or timeout as indicated by the *fail* action.

Susceptible-Infective-Removed over a network

- The capacity of the network is dictated by the parameter M , the number of concurrent, independent connections that the network can sustain.
- Additionally, an attempted network connection can fail or timeout as indicated by the *fail* action.
- This might be due to network contention or the lack of availability of a susceptible machine to infect.

Susceptible-Infective-Removed over a network

- The capacity of the network is dictated by the parameter M , the number of concurrent, independent connections that the network can sustain.
- Additionally, an attempted network connection can fail or timeout as indicated by the *fail* action.
- This might be due to network contention or the lack of availability of a susceptible machine to infect.
- As large scale worm infections tend not to waste time determining whether a given host is already infected or not, we assume that a certain number of infections will attempt to reinfect hosts; in this instance, the host is unaffected.

Susceptible-Infective-Removed over a network

$$S \stackrel{\text{def}}{=} (\text{infect}S, \top).I$$

$$I \stackrel{\text{def}}{=} (\text{infect}I, \beta).I + (\text{infect}S, \top).I + (\text{patch}, \gamma).R$$

$$R \stackrel{\text{def}}{=} \text{Stop}$$

Susceptible-Infective-Removed over a network

$$S \stackrel{\text{def}}{=} (\text{infect}S, \top).I$$

$$I \stackrel{\text{def}}{=} (\text{infect}I, \beta).I + (\text{infect}S, \top).I + (\text{patch}, \gamma).R$$

$$R \stackrel{\text{def}}{=} \text{Stop}$$

$$\text{Net} \stackrel{\text{def}}{=} (\text{infect}I, \top).\text{Net}'$$

$$\text{Net}' \stackrel{\text{def}}{=} (\text{infect}S, \beta).\text{Net} + (\text{fail}, \delta).\text{Net}$$

Susceptible-Infective-Removed over a network

$$S \stackrel{\text{def}}{=} (\text{infect}S, \top).I$$

$$I \stackrel{\text{def}}{=} (\text{infect}I, \beta).I + (\text{infect}S, \top).I + (\text{patch}, \gamma).R$$

$$R \stackrel{\text{def}}{=} \text{Stop}$$

$$\text{Net} \stackrel{\text{def}}{=} (\text{infect}I, \top).\text{Net}'$$


$$\text{Net}' \stackrel{\text{def}}{=} (\text{infect}S, \beta).\text{Net} + (\text{fail}, \delta).\text{Net}$$

$$\text{Sys} \stackrel{\text{def}}{=} (S[N] \parallel I) \bowtie_L \text{Net}[M]$$

$$\text{where } L = \{ \text{infect}I, \text{infect}S \}$$

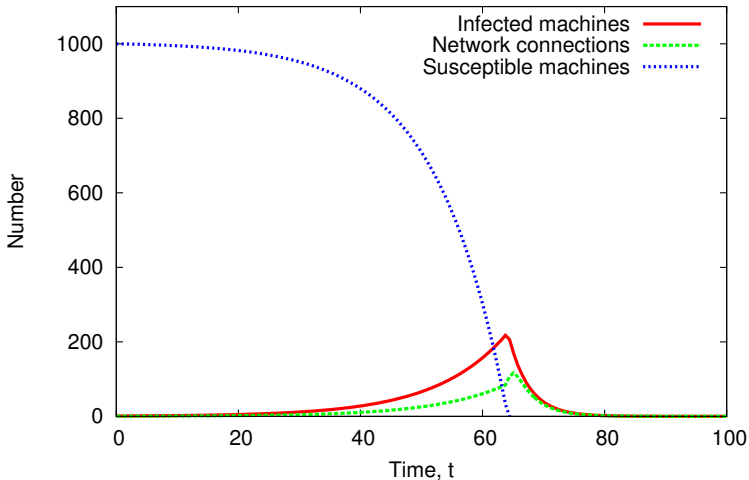
Patch rate $\gamma = 0.1$. Connection failure rate $\delta = 0.5$

Worm infection dynamics for gamma=0.1, delta=0.5




Patch rate $\gamma = 0.3$. Connection failure rate $\delta = 0.5$

Worm infection dynamics for gamma=0.3



Increasing machine patch rate γ from 0.1 to 0.3

Infected machines for different values of gamma



Susceptible-Infective-Removed-Reinfection (SIRR) model

- As with the SIR model, we constrain infection to occur over a limited network resource, constrained by the number of independent network connections in the system, M .

Susceptible-Infective-Removed-Reinfection (SIRR) model

- As with the SIR model, we constrain infection to occur over a limited network resource, constrained by the number of independent network connections in the system, M .
- A small modification in the process model of infection allows for removed computers to become susceptible again after a delay.

Susceptible-Infective-Removed-Reinfection (SIRR) model

- As with the SIR model, we constrain infection to occur over a limited network resource, constrained by the number of independent network connections in the system, M .
- A small modification in the process model of infection allows for removed computers to become susceptible again after a delay.
- We use this to model a faulty or incomplete security upgrade or the mistaken removal of security patches which had previously defended the machine against attack.

Susceptible-Infective-Removed-Reinfection (SIRR) model

$$S \stackrel{\text{def}}{=} (\text{infect}S, \top).I$$

$$I \stackrel{\text{def}}{=} (\text{infect}I, \beta).I + (\text{infect}S, \top).I + (\text{patch}, \gamma).R$$

$$R \stackrel{\text{def}}{=} (\text{unsecure}, \mu).S$$

Susceptible-Infective-Removed-Reinfection (SIRR) model

$$S \stackrel{\text{def}}{=} (\text{infect}S, \top).I$$

$$I \stackrel{\text{def}}{=} (\text{infect}I, \beta).I + (\text{infect}S, \top).I + (\text{patch}, \gamma).R$$

$$R \stackrel{\text{def}}{=} (\text{unsecure}, \mu).S$$

$$\text{Net} \stackrel{\text{def}}{=} (\text{infect}I, \top).\text{Net}'$$

$$\text{Net}' \stackrel{\text{def}}{=} (\text{infect}S, \beta).\text{Net} + (\text{fail}, \delta).\text{Net}$$

Susceptible-Infective-Removed-Reinfection (SIRR) model

$$S \stackrel{\text{def}}{=} (\text{infect}S, \top).I$$

$$I \stackrel{\text{def}}{=} (\text{infect}I, \beta).I + (\text{infect}S, \top).I + (\text{patch}, \gamma).R$$

$$R \stackrel{\text{def}}{=} (\text{unsecure}, \mu).S$$

$$\text{Net} \stackrel{\text{def}}{=} (\text{infect}I, \top).\text{Net}'$$


$$\text{Net}' \stackrel{\text{def}}{=} (\text{infect}S, \beta).\text{Net} + (\text{fail}, \delta).\text{Net}$$

$$\text{Sys} \stackrel{\text{def}}{=} (S[1000] \parallel I) \boxtimes_L \text{Net}[M]$$

$$\text{where } L = \{\text{infect}I, \text{infect}S\}$$


Unsecured SIR model (200 network channels)

Worm infection dynamics for $N=200$




Unsecured SIR model (50 network channels)

Worm infection dynamics for N=50



Unsecured SIR model (20 network channels)

Worm infection dynamics for N=20



Conclusions

- The scale of the effects of Internet worms defeats attempts to model their behaviour in very close detail, and thus impedes the analysis which has the potential to bring understanding of their function and distribution.

Conclusions

- The scale of the effects of Internet worms defeats attempts to model their behaviour in very close detail, and thus impedes the analysis which has the potential to bring understanding of their function and distribution.
- Process algebra modelling allows the details of interactions to be recorded on the individual level but then abstracted away into appropriate population-based representations.

Conclusions

- The scale of the effects of Internet worms defeats attempts to model their behaviour in very close detail, and thus impedes the analysis which has the potential to bring understanding of their function and distribution.
- Process algebra modelling allows the details of interactions to be recorded on the individual level but then abstracted away into appropriate population-based representations.
- The scale of problems which can be modelled in this way vastly exceeds those which are founded on explicit state representations.

Conclusions

- The scale of the effects of Internet worms defeats attempts to model their behaviour in very close detail, and thus impedes the analysis which has the potential to bring understanding of their function and distribution.
- Process algebra modelling allows the details of interactions to be recorded on the individual level but then abstracted away into appropriate population-based representations.
- The scale of problems which can be modelled in this way vastly exceeds those which are founded on explicit state representations.
- We believe the modelling methods exemplified here to be generally useful for analysing the behaviour of populations of interacting processes with complex dynamics.

Virtual University Scenario

- A **Virtual University** is a federation of *real* universities, each contributing courses and degrees.

Virtual University Scenario

- A *Virtual University* is a federation of *real* universities, each contributing courses and degrees.
- Sharing of knowledge is promoted by providing students with a wider selection of subjects.

Virtual University Scenario

- A **Virtual University** is a federation of *real* universities, each contributing courses and degrees.
- Sharing of knowledge is promoted by providing students with a wider selection of subjects.
- Services are replicated across the physical sites.

Virtual University Scenario

- A **Virtual University** is a federation of *real* universities, each contributing courses and degrees.
- Sharing of knowledge is promoted by providing students with a wider selection of subjects.
- Services are replicated across the physical sites.
- By agreement in the university, students may connect to any site to download content and use services, not just the one which is geographically closest.

Case Study: A Virtual University



Location, Time, and Size




Replicating Web Services

Two viable approaches to cope with increasing user demand:

Replicating Web Services

Two viable approaches to cope with increasing user demand:


- use a service broker for routing



Replicating Web Services

Two viable approaches to cope with increasing user demand:

- use a service broker for routing




- decentralised routing

Replicating Web Services


Two viable approaches to cope with increasing user demand:

- use a service broker for routing




- decentralised routing

Decentralised Routing




- 1 A client contacts a university site to download content.

Decentralised Routing



- 1 A client contacts a university site to download content.
- 2 The site either serves the request or forwards it to another site.

Decentralised Routing



- 1 A client contacts a university site to download content.
- 2 The site either serves the request or forwards it to another site.
- 3 The decision is made in accord with the local service policy.

Model in PEPA

Clients

$$\begin{aligned}
 \text{Client}_i & \stackrel{\text{def}}{=} (\text{connect}_1, c_{1,i}).(\text{download}_1, d_{1,i}).\text{Idle}_i \\
 & + (\text{connect}_2, c_{2,i}).(\text{download}_2, d_{2,i}).\text{Idle}_i \\
 & \dots \\
 & + (\text{connect}_m, c_{m,i}).(\text{download}_m, d_{m,i}).\text{Idle}_i \\
 & + (\text{overload}, \top).\text{Client}_i \\
 \text{Idle}_i & \stackrel{\text{def}}{=} (\text{idle}, r_{\text{idle},i}).\text{Client}_i
 \end{aligned}$$

$$(1 \leq i \leq k)$$

Model in PEPA

Clients

$$\begin{aligned}
 \text{Client}_i & \stackrel{\text{def}}{=} (\text{connect}_1, c_{1,i}).(\text{download}_1, d_{1,i}).\text{Idle}_i \\
 & + (\text{connect}_2, c_{2,i}).(\text{download}_2, d_{2,i}).\text{Idle}_i \\
 & \dots \\
 & + (\text{connect}_m, c_{m,i}).(\text{download}_m, d_{m,i}).\text{Idle}_i \\
 & + (\text{overload}, \top).\text{Client}_i \\
 \text{Idle}_i & \stackrel{\text{def}}{=} (\text{idle}, r_{\text{idle},i}).\text{Client}_i
 \end{aligned}$$

$$(1 \leq i \leq k)$$

Model in PEPA

Clients

$$\begin{aligned}
 \text{Client}_i & \stackrel{\text{def}}{=} (\text{connect}_1, c_{1,i}).(\text{download}_1, d_{1,i}).\text{Idle}_i \\
 & + (\text{connect}_2, c_{2,i}).(\text{download}_2, d_{2,i}).\text{Idle}_i \\
 & \dots \\
 & + (\text{connect}_m, c_{m,i}).(\text{download}_m, d_{m,i}).\text{Idle}_i \\
 & + (\text{overload}, \top).\text{Client}_i \\
 \text{Idle}_i & \stackrel{\text{def}}{=} (\text{idle}, r_{\text{idle},i}).\text{Client}_i
 \end{aligned}$$

$$(1 \leq i \leq k)$$

Model in PEPA

Content mirrors

$$\begin{aligned}
 \text{Mirror}_j &\stackrel{\text{def}}{=} (\text{connect}_j, f_j(s)).\text{MirrorUploading}_j \\
 \text{MirrorUploading}_j &\stackrel{\text{def}}{=} (\text{download}_j, \top).\text{Mirror}_j
 \end{aligned}$$

$$(1 \leq j \leq m)$$

Model in PEPA

Content mirrors

$$\begin{aligned}
 \text{Mirror}_j &\stackrel{\text{def}}{=} (\text{connect}_j, f_j(s)).\text{MirrorUploading}_j \\
 \text{MirrorUploading}_j &\stackrel{\text{def}}{=} (\text{download}_j, \top).\text{Mirror}_j \\
 &\quad (1 \leq j \leq m)
 \end{aligned}$$

Service policies as functional rates in PEPA

The Bologna policy

Serve all requests while load is less than 75%. If more, and the loads at UNIFI, UPISA, LMU and UEDIN are at least 60%, 60%, 40% and 20% then serve the request if load is less than 95%.

Service policies as functional rates in PEPA

The Bologna policy

Serve all requests while load is less than 75%. If more, and the loads at UNIFI, UPISA, LMU and UEDIN are at least 60%, 60%, 40% and 20% then serve the request if load is less than 95%.

$$f_{\text{UNIBO}} = \begin{cases} \top & \text{if } \text{MirrorUploading}_{\text{UNIBO}} < 75 \\ \top & \text{if } \text{MirrorUploading}_{\text{UNIBO}} < 95, \\ & \text{MirrorUploading}_{\text{UNIFI}} \geq 60, \\ & \text{MirrorUploading}_{\text{UPISA}} \geq 60, \\ & \text{MirrorUploading}_{\text{LMU}} \geq 40, \\ & \text{MirrorUploading}_{\text{UEDIN}} \geq 20 \\ 0 & \text{otherwise} \end{cases}$$

Model in PEPA

Dealing with overload

$$\textit{Overload} \stackrel{\text{def}}{=} (\textit{overload}, o(s)).\textit{Overload}$$

$$o(s) = \begin{cases} \top & f_i(s) = 0, \quad 1 \leq i \leq m \\ 0 & \text{otherwise} \end{cases}$$

Model in PEPA

Dealing with overload

$$\textit{Overload} \stackrel{\text{def}}{=} (\textit{overload}, o(s)).\textit{Overload}$$


$$o(s) = \begin{cases} \top & f_i(s) = 0, \quad 1 \leq i \leq m \\ 0 & \text{otherwise} \end{cases}$$

The system as a whole with client and mirror site populations

$$(\textit{Client}_1[p_1] \parallel \textit{Client}_2[p_2] \parallel \dots \parallel \textit{Client}_k[p_k]) \\ \boxtimes_L (\textit{Mirror}_1[q_1] \parallel \textit{Mirror}_2[q_2] \parallel \dots \parallel \textit{Mirror}_m[q_m])$$


Numerical Results

$$r_{idle} = 0.001$$




Numerical Results

$$r_{idle} = 0.01$$




Numerical Results

$$r_{idle} = 0.02$$




Numerical Results

$$r_{idle} = 0.03$$




Numerical Results

$$r_{idle} = 0.04$$




Numerical Results

$$r_{idle} = 0.05$$



Numerical Results

$$r_{idle} = 0.06$$



Comments

- We have a modelling approach which captures both functional and non-functional properties of large-scale systems.

Comments

- We have a modelling approach which captures both functional and non-functional properties of large-scale systems.
- PEPA gives us insights into the performance of the system.

Comments

- We have a modelling approach which captures both functional and non-functional properties of large-scale systems.
- PEPA gives us insights into the performance of the system.
 - We applied the continuous-space semantics of PEPA and were able to see service policies at work.

Comments

- We have a modelling approach which captures both functional and non-functional properties of large-scale systems.
- PEPA gives us insights into the performance of the system.
 - We applied the continuous-space semantics of PEPA and were able to see service policies at work.
 - Analysis carried out on a system of 17 ODEs as opposed to an underlying CTMC of over 270 million states.

Comments

- We have a modelling approach which captures both functional and non-functional properties of large-scale systems.
- PEPA gives us insights into the performance of the system.
 - We applied the continuous-space semantics of PEPA and were able to see service policies at work.
 - Analysis carried out on a system of 17 ODEs as opposed to an underlying CTMC of over 270 million states.
- This framework is general. Elements that may be changed are

Comments

- We have a modelling approach which captures both functional and non-functional properties of large-scale systems.
- PEPA gives us insights into the performance of the system.
 - We applied the continuous-space semantics of PEPA and were able to see service policies at work.
 - Analysis carried out on a system of 17 ODEs as opposed to an underlying CTMC of over 270 million states.
- This framework is general. Elements that may be changed are
 - Classes of clients

Comments

- We have a modelling approach which captures both functional and non-functional properties of large-scale systems.
- PEPA gives us insights into the performance of the system.
 - We applied the continuous-space semantics of PEPA and were able to see service policies at work.
 - Analysis carried out on a system of 17 ODEs as opposed to an underlying CTMC of over 270 million states.
- This framework is general. Elements that may be changed are
 - Classes of clients
 - Deployment of service providers

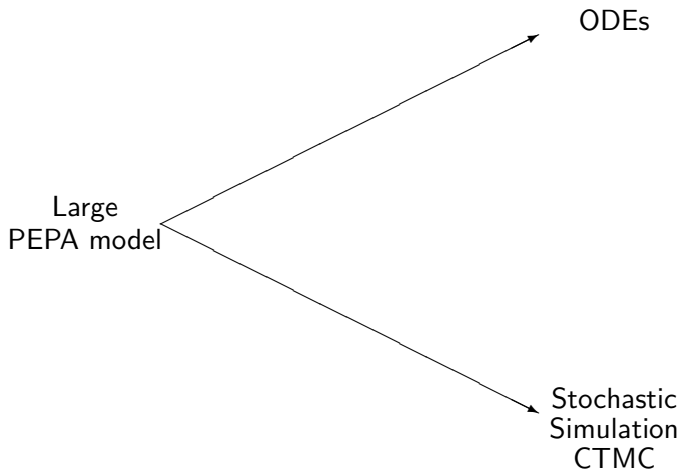
Comments

- We have a modelling approach which captures both functional and non-functional properties of large-scale systems.
- PEPA gives us insights into the performance of the system.
 - We applied the continuous-space semantics of PEPA and were able to see service policies at work.
 - Analysis carried out on a system of 17 ODEs as opposed to an underlying CTMC of over 270 million states.
- This framework is general. Elements that may be changed are
 - Classes of clients
 - Deployment of service providers
 - Different load balancing policies

Outline

- 1 Introduction
 - Collective Dynamics
- 2 Continuous Approximation
- 3 Fluid-Flow Semantics
- 4 Case studies
 - Internet worms
 - Scalable Web Services
- 5 Final remarks**

Alternative Representations



Consistency results

Consistency results

- The vector field $\mathcal{F}(x)$ is Lipschitz continuous

Consistency results

- The vector field $\mathcal{F}(x)$ is Lipschitz continuous
- The generated ODEs are the fluid limit of the family of CTMCs generated by $f(\xi, l, \alpha)$

Consistency results

- The vector field $\mathcal{F}(x)$ is Lipschitz continuous
- The generated ODEs are the fluid limit of the family of CTMCs generated by $f(\xi, l, \alpha)$
- We can prove this using Kurtz's theorem:
Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes, T.G. Kurtz, J. Appl. Prob. (1970).

Conclusions

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

Conclusions

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

Conclusions

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

Thanks!

Acknowledgements: collaborators

Thanks to many co-authors and collaborators: **Jeremy Bradley**, Muffy Calder, Federica Ciocchetta, Allan Clark, Jie Ding, Adam Duguid, Vashti Galpin, Nil Geisweiller, **Stephen Gilmore**, Marco Stenico, **Mirco Tribastone**, and others.

Acknowledgements: funding

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

Thanks!

Acknowledgements: collaborators

Thanks to many co-authors and collaborators: **Jeremy Bradley**, Muffy Calder, Federica Ciocchetta, Allan Clark, Jie Ding, Adam Duguid, Vashti Galpin, Nil Geisweiller, **Stephen Gilmore**, Marco Stenico, **Mirco Tribastone**, and others.

Acknowledgements: funding

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

More information:

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