

GLOBALAN Stochastic Modelling Exercise

Stephen Gilmore
Laboratory for Foundations of Computer Science
The University of Edinburgh

stg@inf.ed.ac.uk

1 About this exercise

This exercise is on applying stochastic modelling to a global computing system. You are at liberty to choose the particular system which you model but it should have at least some of the typical features of a global computing system discussed in the course: physically distributed, large-scale, and resource sharing, with partial failure modes.

Below are examples and non-examples of Global Computing systems.

Global Computing	
Examples	Non-examples
A replicated Web service used by users from many locations	A printer which sometimes jams
A compute cluster such as a Beowulf or Condor pool	A single user device which rarely fails (e.g. an iPod)

The steps which you need to take are these:

1. consider the system which you wish to model and identify the components of interest and the activities which they undertake either individually or in co-operation;
2. describe these components behaviourally, using finite-state representations;
3. express your model as a high-level description in the PEPA stochastic process algebra;
4. count, measure or estimate quantities of interest in the model (e.g. rates, probabilities, and population sizes), adding these to the model;
5. process your model with the PEPA Workbench or the P.Dizzy differential equation and stochastic simulation solver or the PRISM probabilistic model checker;
6. compute performance results from your model.

2 Submission description

Present your submission as a short paper (10 pages) containing:

- an abstract;
- an introduction containing an informal, but careful, explanation of the problem which you are modelling;
- the PEPA model of your system, explaining how this relates to the informal description;

- a section on parameter estimation, explaining how numerical values have been measured or estimated;
- the results obtained from the modelling tools;
- your conclusions.

3 Take home message from the course

The key to success in stochastic modelling is *abstraction*. It is essential to leave out as much detail as possible. Start with a small, but unrealistic, model and try to gradually incorporate detail to move towards a larger, more realistic model.

4 Deadline, submission instructions

Send your submission by email to `stg@inf.ed.ac.uk` (with CC to `globan`) by **November 1st**. Submissions will be assessed by **December 24th**.

5 Resources

As an introduction to PEPA you have the slides on stochastic modelling from the GLOBAN summerschool and papers on the CD. Key papers on the CD are

- *Tuning systems: from composition to performance*, Jane Hillston, Computer Journal, 2005.
— An introduction to PEPA and stochastic modelling with emphasis on applications.
- *Process algebras for quantitative analysis*, Jane Hillston, LICS, 2005.
— An introduction to PEPA and its associated theory.
- *Fluid-flow approximation for PEPA models*, Jane Hillston, QEST, 2005.
— The differential equation semantics for PEPA.

Many other papers and example PEPA models are available on-line at the PEPA Web site <http://www.dcs.ed.ac.uk/pepa>

Software tools such as the PEPA Workbench and the P.Dizzy differential equation solver and stochastic simulator are available from the Web address <http://www.dcs.ed.ac.uk/pepa/LATEST>

These tools are freely available for download. They are implemented in the Java programming language and require Java 1.5 or higher to execute. They will not work with earlier versions of the Java programming language. Test your version of Java using `java --version`. Additionally the P.Dizzy modelling tool requires the Java build tool Ant to be available. This is built-in to Java environments such as NetBeans so using this will be enough.

The PRISM probabilistic model checker is available for download from <http://www.cs.bham.ac.uk/~dxdp/prism>