BISS 2013:
Stochastic Process Algebras for Quantitative Analysis Exam

In all cases the deadline is 31st May 2013, and your exam should be mailed, as a pdf document to `jane.hillston@ed.ac.uk`. In the case of OPTION 2 you should also send your `.pepa` file.

You should indicate just one of the options, and please register your intention to take the exam by 30th April 2013.

OPTION 1: Literature review

As mentioned in the lectures there are a number of different stochastic process algebras. The first option is to carry out a study of **one** of the alternatives to PEPA and then write a summary of your chosen formalism and contrast it with PEPA. The suggested alternative languages are:

- **EMPA**: Marco Bernardo, Roberto Gorrieri, *A Tutorial on EMPA: A Theory of Concurrent Processes with Nondeterminism, Priorities, Probabilities and Time*, in Theorical Computer Science, 202(1-2): 1-54 (1998)
- **IMC**: Holger Hermanns, *Interactive Markov Chains: The Quest for Quantified Quality*, Lecture Notes in Computer Science 2428, Springer 2002, isbn 3-540-44261-8
- **SCCP**: Luca Bortolussi, *Stochastic Concurrent Constraint Programming* in Electronic Notes in Theoretical Computer Science, 164(3): 65-80 (2006)
- **Stochastic π-Calculus**: Corrado Priami, *Stochastic pi-Calculus*, in The Computer Journal, 38(7): 578-589 (1995)
- **Bio-PEPA**: Federica Ciocchetta, Jane Hillston, *Bio-PEPA: A framework for the modelling and analysis of biological systems*, Theoretical Computer Science, 410(33-34): 3065-3084 (2009)

In each case I have given the definitive reference for the language but you may find it useful to look at some other papers, to deepen your understanding of the formalism and how it is used.

Your document should include a bibliography and may include small examples to illustrate features of the two languages, particularly their different forms of expression. Your document should be approximately 2500 words (∼5 pages).

OPTION 2: Modelling study

In this option the intention is for you to get some experience of building a PEPA model and conducting some experiments upon it. Below I suggest a specific system which could be interesting to consider but I am open to negotiation if there is another system which you particularly want to model in PEPA. However you **must** get my agreement at least one month before the deadline (ie. 30th April) if you are planning to submit a different model.

The scenario you are asked to consider is the spread of information in an online social network. It is well-known that some topics or messages go viral in such networks, meaning that the majority of nodes are aware of the message, whilst other topics remain localised and do not spread far. Your task is to build a simple PEPA model that can be used to study this phenomenon.

- We assume that the network consists of a large number of subnetworks, each of which may or may not have the message circulating within it at any given time. Users are represented by nodes and nodes can both receive messages from and publish messages onto any of the networks.
- We assume that we are concerned with the spread of a single tagged message (received from the network in the *message* action); all other messages are treated as *data*.
- It is also assumed that each node has a limited capacity so that the arrival of more data may overwrite the message even before it has been read. You may handle this probabilistically, or you may use different states to capture the state of your node (for example, assuming that the capacity of a node is limited to 5 items (message and/or data arrivals)) but this latter approach will lead to a much larger state space.
- If the message persists long enough at a node to be read, then the node may or may not choose to pass the message on, and in that case they may publish onto a network that already has the message or not. This choice is probabilistically determined by the proportion of subnetworks which currently are carrying the message.
- Whilst a subnetwork may receive the tagged message several times from different (or even the same) nodes; we assume that a node that currently holds the message does not receive it again, although it may receive it again after the message has been overwritten by other data.

Thus, using probabilistic representation of capacity, a node may be in three distinct states:

1. Here it does not have the tagged message; it can receive the message from the network (and move to state 2) or it may receive more data and remain in the same state;
2. Here it has received the tagged message but the message has not yet been read; it can receive more data that may not overwrite the message (remain in state 2) or may overwrite it (return to state 1); alternatively the message may be read moving the node to state 3.
3. Here the message has been received and read and the node may choose to propagate the message, by itself publishing the message to a subnetwork; as in state 2 it will continue to receive other data.

You should similarly think about the different states and transitions for a subnetwork. You will then be ready to build a simple PEPA model consisting of a number of nodes and a number of subnetworks and the interactions between them. Think about the starting state of your model.

From this basic model it is up to you to choose what behaviour you are interested in and how to explore it. For example, the scenario above assumes that some nodes start with the message and is suitable for a transient analysis of how the message propagates over time, but the message will always eventually be lost. Instead you could introduce an action to spontaneously create new copies of the message. Similarly I have not described any direct transfer between subnetworks except via nodes, but this could be added to the model if you felt it was appropriate.

Whichever model you develop, you are asked to write a report on your modelling study which includes

1. A description of the system you have modelled, and any assumptions that you have made in constructing your model;

2. A description of your experimental plan, i.e. what you set out to investigate and how you chose to do it. For example, this would explain which aspects of the model you chose to vary and which you kept the same throughout; it would also detail which analysis techniques you used;

3. A report of your results, including some graphs and/or tables;

4. A brief conclusion explaining what you have learned about the system through your modelling study.

OPTION 3: Semantic study

In this option the intention is for you to get some experience directly with working on the semantics and equivalence relations of a stochastic process algebra. PEPA currently has a probabilistic choice that is defined implicitly, in the sense that the relative probability of the alternative derivations is based on the race policy governing the execution of the enabled actions. The goal of this small project is to consider the implications of adding an explicit probabilistic choice operator to the syntax of PEPA. For example writing $P +_p Q$ would assume that a Bernoulli trial is used to decide whether P or Q proceeds, with the probabilities p and $1 - p$ respectively. Your task is to define a variant of PEPA which either includes probabilistic choice as an additional operator or which replaces the existing choice based on the race condition with a probabilistic choice. For your stochastic process algebra you should write a brief report in which you:

1. Define the syntax and give an informal definition;

2. Give a structural operational semantics in the Plotkin-style

3. Define a bisimulation-based equivalence relation for the language

4. Establish whether your bisimulation is a congruence for all the operators of the language.