Process Algebras and Their Application to Performance Modelling

PROCEEDINGS OF THE THIRD WORKSHOP ON PROCESS ALGEBRA AND PERFORMANCE MODELLING

Edinburgh, Scotland, UK

The impetus behind most of the best known work on building and analysing formal models of a computer system has been to ensure correct and safe behaviour. When the system being modelled contains concurrently active components which cooperate and share work then process algebras are often considered to be an appropriate formal modelling approach. Process algebra models contain both structural and dynamic information and can be rigorously analysed using logical and algebraic methods.

Expectations of a dependable computer system go beyond correct and safe functionality to include satisfaction of timing constraints. The discipline of performance analysis addresses those issues. One of its central concerns is predicting whether a system will respond to its users in a timely manner. Performance analysis is commonly based on stochastic or Markovian models obtained from a high-level representation. These models can be solved by analytical or numerical methods.

Queueing networks have been widely used as performance models for over twenty-five years but they have been overtaken by the advent of large distributed computer systems. It now seems that in order to tackle models of the size and complexity needed for the performance analysis of the next generation of computer systems novel structured techniques will be required for the construction, manipulation and solution of large Markov chains.

Stochastic—often Markovian—process algebras (SPA or MPA) provide expressive high-level languages with enough structure to support the description of the large Markov chains needed for performance analysis of modern computer systems. They are amenable to formal analysis using both the logical and algebraic methods of classical process algebras and the analytical and numerical methods of performance models. All of the charms of more familiar process algebras are still here: the languages are small and defined formally by structured operational semantics; natural equivalences exist to relate processes which an external observer could not distinguish; and complicated processes may be put together as the composition of simpler parts.

If these new approaches bring about an increase in our capability to express more sophisticated designs then model simplification and approximate solution techniques must surely be destined to play increasingly important roles. One of the most promising aspects of the work on stochastic process algebras is the prospect that these techniques, which currently rely on the expertise of experienced performance analysts, may be automated and applied by less experienced users.

The series of workshops on Process Algebra and Performance Modelling was founded to encourage the study of the interplay between functional and performance analysis in a process algebra setting. The first workshop was arranged in May 1993 by Jane Hillston and Faron Moller of the University of Edinburgh. It brought together a diverse group of researchers who shared a common interest in deriving quantitative results from process algebra models which had previously only been used to carry out qualitative analysis of systems. This was a new departure both for formal methods and for performance modelling and the approaches proposed were varied.

Ulrich Herzog and Michael Rettelbach of Universität Erlangen-Nürnberg organised the second workshop in July 1994. The growing maturity of the field was reflected in a greater consensus throughout the material presented at the workshop. In addition to work on stochastic extensions of process algebras themselves there were also papers examining the relationships between the different proposed approaches.

The papers presented here represent the third workshop: once again held in Edinburgh, in June 1995. These papers show that the work on stochastic process algebras is now developing in two clearly identifiable, and equally important, directions. Both are perhaps inspired by the problems of dealing with the state space explosion which occurs in stochastic modelling. The first direction addresses problems of model complexity by offering the modeller a more compact representation of the system. The second direction lies in advancing the search for structured solution techniques for large models.

The first two papers of the proceedings enrich the expressiveness of a stochastic process algebra by incorporating immediate actions which have two important roles within a modelling paradigm.

The first role—representing "management and control" activities whose duration is negligible—is investigated in the paper by Hermanns, Rettelbach and Weiß. In this paper, *Formal Characterisation of Immediate Actions in SPA with Nondeterministic Branching*, previous theoretical results about the stochastic process algebra Markovian TIPP are extended to a new version of the language which includes immediate actions. Moreover it is shown, by the existence of a suitable equivalence relation, that the transitions corresponding to these actions may be safely eliminated before the underlying continuous time Markov chain (CTMC) is generated.

The second role—representing decisions and choices which do not consume system resources—is investigated by Rettelbach in his paper *Probabilistic Branching in Markovian Process Algebras*. Here too an equivalence relation is developed which could be used to automatically reduce the transition system without affecting the integrity of the performance measures.

An alternative approach to handling large models is represented by the work presented in the next four papers. Here the emphasis is on finding efficient solution techniques for large models which exploit their structure. Useful structure can sometimes be readily identified in the process algebra presentation of the model.

The paper by Haverkort, *In Search of Probability Mass: Probabilistic Evaluation of High-Level Specified Markov Models*, describes an iterative algorithm which avoids constructing a large part of the entire state space for the model while still giving results which are good approximations to the exact solution.

In *A Simple Time Scale Decomposition Technique for Stochastic Process Algebras* Hillston and Mertsiotakis describe how to recognise models for which a good approximate solution can be obtained when only some actions are considered at a time. An algorithm to implement the proposed technique is presented and a prototype implementation is described.

By far the most efficient class of solutions for Markov processes are the so-called *product form* solutions which take advantage of apparent independence between submodels. When submodels behave as if independent they can be solved in isolation thereby greatly reducing the size and complexity of the problem. The component structure and clear definition of submodel interactions present in SPA models has encouraged researchers to seek to identify when such models exhibit product form solution. In these proceedings we include two papers on this topic, showing two distinct approaches to the problems.

Sereno's approach, presented in *Towards a Product Form Solution for Stochastic Process Algebras*, is derived from earlier work on Stochastic Petri Nets (SPN). In contrast, Harrison and Hillston's work, *Exploiting Quasi-reversible Structure in Markovian Process Algebra Models* extends previous work on queueing networks. An interesting area for future work would be to compare these two distinct approaches within the common framework of stochastic process algebras.

It would be an unwise practice to develop the theory of stochastic process algebras without simultaneously assessing their practical utility. For this reason we are very pleased to include a paper by Holton which describes the construction and solution of a model of a real manufacturing process: *A PEPA Specification of an Industrial Production Cell*.

It is also very important for the process algebra approach to stochastic modelling to learn from existing, successful approaches such as queueing networks and generalised stochastic Petri nets. The benefits of understanding the relationship between these formalisms are apparent in the work on efficient solution techniques: as we noted, the proposed algorithms were inspired by earlier work on queueing networks and stochastic Petri nets. Deep insight into the relationship between algebras and nets can be obtained by comparing the semantics of the formalisms—as in the case of the papers by Bernardo *et al* and Ribaudo.

In *A Distributed Semantics for EMPA Based on Stochastic Contextual Nets* Bernardo *et al.* develop a semantic model of Extended Markovian Process Algebra (EMPA) based on an extended class of generalised stochastic Petri nets, *contextual nets*. The appeal of this semantics is its simplicity, with only one rule capturing the behaviour of all operators at the process algebra level.

The paper *On the Aggregation Techniques in Stochastic Petri Nets and Stochastic Process Algebras* makes use of a net-based semantics to compare approaches to state space aggregation which are available in SPA and Stochastic Well-formed Nets (SWN), a class of coloured SPN.

Different paradigms exhibit distinct strengths and weaknesses and a better understanding of the relationships between them can have mutual benefit as characteristics and techniques are imported from one to the other. In addition to the work enhancing SPA with techniques inspired by other formalisms, process algebra ideas have been incorporated into existing performance modelling languages. In *Compositional Construction of SWN models*, Rojas investigates ways in which SWN can be enriched with some of the structural and reasoning apparatus of SPA. Her work promises a fresh approach to the construction and analysis of Petri net based models.

This issue concludes with two papers which explore intriguing next possible directions for formal stochastic process modelling. Brinksma *et al.* investigates the potential of using a true concurrency approach in *A*

Stochastic Causality-Based Process Algebra. In *Stochastic π-Calculus* Priami considers the flexibility which process mobility might bring to performance modelling whilst staying within the Markovian framework.

The editors of this special issue wish to thank all of the participants who attended the third workshop on Process Algebras for Performance Modelling. The next in the series will be held in Turin in July 1996. Our task as editors this time has been made much easier by an excellent body of referees who discharged their duties both thoroughly and quickly. Our thanks also go to all the authors for meeting the tight deadlines which we set without compromising on the rigour or clarity of their papers.

Stephen Gilmore and Jane Hillston
Department of Computer Science
University of Edinburgh
Email: {stg,jeh}@dcs.ed.ac.uk