Process Algebras for Collective Dynamics (Extended Abstract)

Jane Hillston

Laboratory for Foundations of Computer Science,
The University of Edinburgh, Scotland

Quantitative Analysis

Stochastic process algebras extend classical process algebras such as CCS [1] and CSP [2] with quantified notions of time and probability. Examples include PEPA [3], EMPA [4], MoDeST [5] and IMC [6]. These formalisms retain the compositional structure of classical process algebras and the additional information captured within the model allows analysis to investigate additional properties such as dynamic behaviour and resource usage.

Stochastic process algebras have been successfully applied to quantitative evaluation of systems for over a decade. For example, in the context of performance analysis, PEPA has been used to describe both software and hardware systems and has helped to incorporate early performance prediction into the design process. Moreover, recently there has been considerable interest in using stochastic process algebras for modelling intracellular processes in systems biology.

In all these models the entities in the system under study are represented as components in the process algebra. The structured operational semantics of the language is used to identify all possible behaviours of the system as a labelled transition system. With suitable restrictions on the form of random variables used to govern delays within the model to be negative exponentially distributed this labelled transition system can be interpreted as a continuous time Markov chain (CTMC). This provides access to a wide array of analysis techniques, usually in terms of the evolution of the probability distribution over states of the model over time.

This has the advantage that it is a fine-grained view of the system, allowing the quantitative characteristics of individual entities to be derived. Unfortunately it has the disadvantage that generation and manipulation of the necessary CTMC can be very computationally expensive, or even intractable, due to the well-known *state space explosion* problem. This problem becomes particularly acute in situations where there are large numbers of entities exhibiting similar behaviour interacting within a system. Often in these situations whilst it is important to capture the behaviour of individual entities accurately the dynamics of the system are most fruitfully considered at a population level. Examples include the spread of disease through a population, the behaviour of crowds during emergency evacuation of a building or scalability studies involving a large number of users trying to access a service. In these cases we are interested in the *collective* rather than the *individual* dynamics.

Collective Dynamics

Process algebras have several attractions for modelling for collective dynamics. The behaviour of individuals, and particularly their interactions are important for such systems, and the compositional approach of the process algebra allows the modeller to capture the exact form of interactions and constraints between entities. However, standard approaches to analysis of process algebra models remain focused on the behaviour of individuals and are inherently discrete event-based. As explained above, this leads to the state space explosion problem and makes it difficult to construct models large enough to exhibit the population level effects which we are interested in.

Thus at Edinburgh we have been investigating the use of process algebras for collective dynamics based on alternative semantics for the constructed models, which consider the population rather than the individuals. As observed above, the semantics of individual-oriented stochastic process algebra models generally gives rise to a discrete state space with dynamics captured by a continuous time Markov chain. In the context of collective dyanamics, an alternative mathematical framework based on sets of ordinary differential equations is used. This may be regarded as a fluid approximation of the discrete state model [7] and recent work has shown how this may be accessed directly via a novel symbolic structured operational semantics [8]. This provides a framework in which to establish the relationship between the two alternative forms of representation.

References

1. Milner, R.: A Calculus of Communicating Systems. Number 92 in LNCS. Springer-Verlag (1980)
2. Hoare, C.: Communicating Sequential Processes. Prentice Hall (1985)
3. Hillston, J.: A Compositional Approach to Performance Modelling. Cambridge University Press (1996)
4. Bernardo, M., Gorrieri, R.: A tutorial on EMPA: a theory of concurrent processes with nondeterminism, priorities, probabilities and time. TCS **202** (1998) 1–54
5. D'Argenio, P., Hermanns, H., Katoen, J.P., Klaren, R.: Modest — a modelling and description language for stochastic timed systems. In: Proc. of PAPM-PROBMIV 2001. Volume 2165 of LNCS., Springer-Verlag (2001)
6. Hermanns, H.: Interactive Markov Chains. Springer-Verlag (2002)
7. Hillston, J.: Fluid flow approximation of PEPA models. In: Proc. of the 2nd International Conference on Quantitative Evaluation of Systems. (2005)
8. Tribastone, M., Gilmore, S., Hillston, J.: Scalable differential analysis of process algebra models. IEEE Transactions on Software Engineering (2010) To appear.