

# *The Needham Lecture*

## Tuning Systems: From Composition to Performance

JANE HILLSTON

*School of Informatics, University of Edinburgh, Scotland, UK*  
*Email: Jane.Hillston@ed.ac.uk*

---

**This paper gives a summary of some of the work of the Performance Evaluation Process Algebra (PEPA) project, which was awarded the 2004 Roger Needham Award from the BCS. Centred on the PEPA modelling formalism, the project has sought to balance theory and practice. Theoretical developments have been tested and validated by application to a wide range of problems and such case studies have provided the stimulus for new directions in theory. Both aspects of the work are presented in summary as well as some current and future research topics.**

---

### 1. INTRODUCTION

The Performance Evaluation Process Algebra (PEPA) project was started in Edinburgh in 1991. It was motivated by problems encountered when carrying out performance analysis of large computer and communication systems, based on numerical analysis of Markov processes. Performance analysis seeks to predict the behaviour of a system with respect to dynamic properties such as the number of requests that can be satisfied per unit time (throughput) and response time. Markov processes, whilst offering general applicability, are difficult to construct for large systems, so an intermediate system description language is often used. In the early 1990s the most common of these were queueing networks and stochastic Petri nets (SPNs). Queueing networks, whilst very powerful when applicable, have limited expressiveness and lack formal interpretation. SPN models have formal interpretation but do not have the explicit structure found in queueing networks, which greatly eases model construction.

Process algebras such as calculus of communicating systems (CCS) and communicating sequential processes (CSP) [1, 2] offer a compositional description technique supported by apparatus for formal reasoning. From a performance perspective, however, they lack the timing information essential to derive performance estimates: actions are instantaneous and choices are non-deterministic. PEPA sought to address these problems by the introduction of a suitably quantified process algebra. By associating an exponentially distributed delay with each activity of the process algebra, it is possible to generate a continuous time Markov chain (CTMC) using PEPA.

The PEPA project has sought to investigate and exploit the interplay between process algebra and the underlying CTMC.

The remainder of this paper is structured as follows. In Section 2 a more detailed account of the motivation and

background of the project is given. Some of the theoretical developments of the PEPA project are described in Section 3. It is unwise practice to develop theory without checking it against real applications and some of the case studies which have been undertaken using PEPA are outlined in Section 4. This section also gives an account of the tool support for PEPA and recent efforts to extend the applicability of PEPA modelling. To conclude, some directions for future and ongoing work are presented in Section 5.

### 2. BACKGROUND

In this section we explain in more detail the context in which PEPA was developed, that of performance modelling of computer and communication systems using CTMCs. This is followed by a brief overview of process algebras, and an introduction to the idea of stochastic process algebra (SPA).

#### 2.1. Performance modelling using CTMCs

There are a variety of approaches available for the performance evaluation of computer and communication systems. If the system exists it may be possible to monitor the system directly. However, in general, such an approach is time-consuming, difficult and lacks generality. Therefore, it is often preferable to model the system rather than use such direct experimentation. Indeed when the system is yet to be constructed, modelling is the only option.

Performance models may be analysed by simulation, numerical solution or analytical solution. Simulation models have the advantage of being insensitive to state space size. Unfortunately, such models are time-consuming to analyse and bring the intellectual burden of evaluating the trustworthiness of the results by the calculation of confidence intervals.

In contrast, analytic solution, in which an expression for the performance measure of interest is derived in terms of the input parameters of the model, can be extremely efficient to use. However, constructing such solutions is very much the domain of the expert and typically each system requires a bespoke solution.

The numerical solution of a Markov chain offers a compromise between these two extremes. Some assumptions about the system, particularly with respect to the timing of events, are needed. But the resulting models are relatively straightforward to solve, relying only on simple linear algebra techniques. For moderately sized models the generator matrix of the Markov chain can be stored as a dense matrix, admitting direct solution methods with good numerical accuracy. For larger models sparse matrices are needed, necessitating the use of iterative solution techniques with some loss of numerical precision. The largest models require even more ingenuity in the representation of the matrix using Kronecker methods or binary decision diagram (BDD)-based storage. Here, convergence becomes an issue and these storage schemes inevitably lead to much longer solution times because they are data structures tuned for compactness not speed of access.

**2.2. Process algebra**

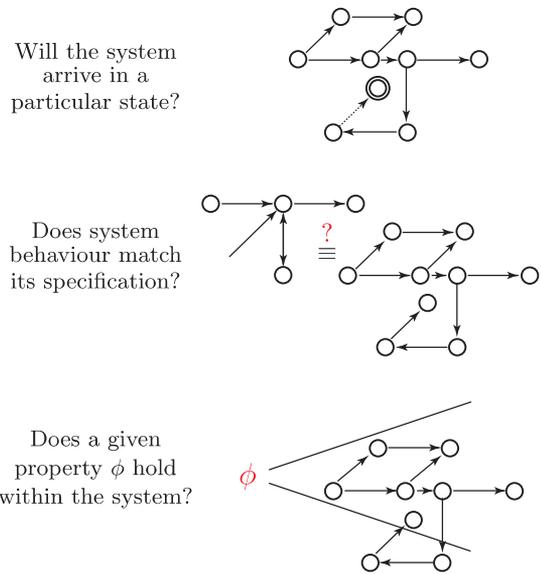
Process algebras are abstract languages used for the specification and design of concurrent systems. The most widely known process algebras are Milner’s CCS [1] and Hoare’s CSP [2]. SPAs take inspiration from both these formalisms.

In the process algebra approach systems are modelled as collections of entities, called agents, which execute atomic actions. These actions are the building blocks of the language and they are used to describe sequential behaviours which may run concurrently, and synchronizations or communications between them.

In CCS two agents communicate when one performs an action, say  $a$ , while the other performs the complementary action  $\bar{a}$ . The resulting communication action has the distinguished label  $\tau$ , which represents an internal action that is invisible to the environment. Agents may proceed with their internal actions simultaneously, but it is important to note that the semantics given to the language imposes an interleaving on such concurrent behaviour.

The communication mechanism in CSP is different as there is no notion of complementary actions. In CSP two agents communicate by simultaneously executing actions with the same label. Since during the communication the joint action remains visible to the environment, it can be reused by other concurrent processes so that more than two processes can be involved in the communication (multi-way synchronization). This is the communication mechanism adopted by most of the SPA languages.

In either case the language is given a small-step structured operational semantics [3], and this is used to generate a labelled transition system. This can be regarded as a derivative tree or graph, in which language terms form the



**FIGURE 1.** Functional analysis of process algebra.

nodes and transitions (actions) are the arcs, which records all the possible evolutions of a language expression or model. This graph can be used for functional verification. For example (Figure 1):

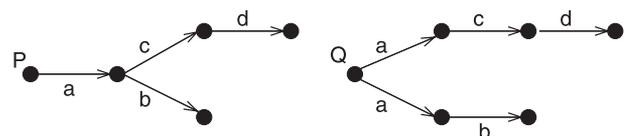
*Reachability analysis* considers whether there is an evolution of the system which will arrive at a particular state or exhibit a particular behaviour, e.g. establishing freedom from deadlock or livelock.

*Specification matching* contrasts a process algebra expression of the expected behaviour of the system, a specification, with a model of the actual implementation. If under an appropriate notion of equivalence the two are equivalent, it can be verified that the implementation will deliver the expected behaviour.

*Model checking.* Many process algebras are equipped with complementary logics which allow properties of the system or its evolution to be expressed.<sup>1</sup> Model checking is a procedure for establishing whether a particular property  $\phi$  will be respected by the system.

Various forms of equivalence relations have been defined for process algebras, based on the structured operational semantics, and they play a fundamental role in model analysis. An important class of relations are based on the notion of bisimulation: two agents are bisimilar if, from the perspective of an observer, each is able to mimic the behaviour of the other.

Consider the agents,  $P$  and  $Q$ , below:



<sup>1</sup>Specifications can be expressed in such a logic as an alternative to a process algebra model.

Although the two are trace equivalent, i.e. they can generate identical sequences of actions, they are not bisimulation equivalent because when  $Q$  has performed an action  $a$  it cannot offer the choice of a  $b$  or a  $c$  in the manner of  $P$ .

### 2.3. Stochastic process algebra

In order to carry out performance analysis of a system, it is essential to record information about the timing characteristics of the system and the relative probabilities of alternative behaviours. Without this quantified information it is not possible to derive quantitative measures such as expected response time or throughput. Therefore, in order to create a process algebra suitable for performance modelling, this quantification was added by associating a random variable with all of the actions of the algebra [4]. In the case of PEPA it is assumed that these random variables are governed by a negative exponential distribution [5]. This is because it is only in that case that it is possible, in general, to associate a CTMC with the process algebra model. Explicit probabilities need not be used to differentiate alternative behaviours—implicit probabilities may be derived from the relative timings of the actions involved.

The benefits of adding quantification can be seen if we reconsider the functional analysis scenarios presented in Figure 1. In each case we can enrich the question asked to take into account timing or probability information. For example, we can now query how long it might take to reach a particular state or behaviour, whether the probability profiles of the specification and the implementation models match, and the probability that a particular property holds.

## 3. THEORY

The theoretical development underpinning PEPA has focused on the interaction between the process algebra and the underlying mathematical structure, the CTMC. The work can be broadly categorized into three areas:

- designing the language
- manipulating models
- solving models and deriving measures.

These will be discussed in the following subsections.

### 3.1. Designing the language

In the early 1990s several SPAs, motivated by the desire to add quantification to process algebra models and make them suitable for performance modelling, appeared in the literature. These included TIPP [6] from the University of Erlangen, EMPA [7] from the University of Bologna and SPADE [8] from Imperial College, in addition to PEPA. PEPA was the first language to be developed with the intention of generating a CTMC which could be solved numerically, but versions of TIPP and EMPA from around the same time were similarly Markovian based.

As previously mentioned, PEPA extends classical process algebra by associating a random variable, representing duration, with every action. These random variables are assumed to be negative exponentially distributed and this

leads to a clear relationship between the process algebra model and a CTMC. Via this underlying CTMC performance measures can be extracted from the model.

PEPA models are described as interactions of components. Each component can perform a set of actions: an action  $a \in \mathcal{Act}$  is described by a pair  $(\alpha, r)$ , where  $\alpha \in \mathcal{A}$  is the type of the action and  $r \in \mathbb{R}^+$  is the parameter of a negative exponential distribution governing its duration. Whenever a process  $P$  can perform an action, an instance of the probability distribution is sampled: the resulting number specifies how long it will take to complete the action in this instance.

A small but powerful set of combinators is used to build up complex behaviour from simpler behaviour. The combinators are familiar from classical process algebra: prefix, choice, parallel composition (cooperation) and abstraction (hiding). Each of the combinators is informally introduced below in terms of a very simple model of a web-based information system, and the formal operational semantics are shown in Figure 2.

*Prefix* ( $\cdot$ ): A component may have purely sequential behaviour, repeatedly undertaking one activity after another and eventually returning to the beginning of its behaviour. A simple example is a web server, which allows one data transfer at a time. Each browser requiring web pages will need to acquire access to the server and only when the transfer is complete will the server be released and available again for acquisition.

$$Server \stackrel{def}{=} (get, \top).(downld, \mu).(rel, \top).Server$$

In some cases, as here, the rate of an action is outside the control of this component. Such actions are carried out jointly with another component, with this component playing a passive role. For example, the server is passive with respect to the *get* action and this is recorded by the distinguished symbol  $\top$  (called ‘top’).

*Choice* ( $+$ ): A choice between two possible behaviours is represented as the sum of the possibilities. For example, if we consider a browser in the information system, displaying the current data may have two possible outcomes: demand for access to data available in the local cache (with probability  $p_1$ ) or demand for access to data stored at the remote server (with probability  $p_2 = 1 - p_1$ ). These alternatives are represented as shown below:

$$Browser \stackrel{def}{=} (display, p_1\lambda).(cache, m).Browser \\ + (display, p_2\lambda).(get, g).(downld, \top).(rel, r).Browser$$

A race condition governs the behaviour of simultaneously enabled actions and the continuous nature of the probability distributions ensures that the actions cannot occur simultaneously. Thus a sum will behave as either one summand or the other. When an action has more than one possible outcome, e.g. the display action in the browser, it is represented by a choice of separate actions, one for each possible outcome. The rates of these actions are chosen to reflect their relative probabilities.

<p><b>Prefix</b></p> $\frac{}{(\alpha, r).E \xrightarrow{(\alpha, r)} E}$ <p><b>Choice</b></p> $\frac{E \xrightarrow{(\alpha, r)} E'}{E + F \xrightarrow{(\alpha, r)} E'}$ $\frac{F \xrightarrow{(\alpha, r)} F'}{E + F \xrightarrow{(\alpha, r)} F'}$	<p><b>Cooperation</b></p> $\frac{E \xrightarrow{(\alpha, r)} E'}{E \boxtimes_L F \xrightarrow{(\alpha, r)} E' \boxtimes_L F} \quad (\alpha \notin L)$ $\frac{F \xrightarrow{(\alpha, r)} F'}{E \boxtimes_L F \xrightarrow{(\alpha, r)} E \boxtimes_L F'} \quad (\alpha \notin L)$ $\frac{E \xrightarrow{(\alpha, r_1)} E' \quad F \xrightarrow{(\alpha, r_2)} F'}{E \boxtimes_L F \xrightarrow{(\alpha, R)} E' \boxtimes_L F'} \quad (\alpha \in L)$ <p>where <math>R = \frac{r_1}{r_\alpha(E)} \frac{r_2}{r_\alpha(F)} \min(r_\alpha(E), r_\alpha(F))</math></p>	<p><b>Hiding</b></p> $\frac{E \xrightarrow{(\alpha, r)} E'}{E/L \xrightarrow{(\alpha, r)} E'/L} \quad (\alpha \notin L)$ $\frac{E \xrightarrow{(\alpha, r)} E'}{E/L \xrightarrow{(\tau, r)} E'/L} \quad (\alpha \in L)$ <p><b>Constant</b></p> $\frac{E \xrightarrow{(\alpha, r)} E'}{A \xrightarrow{(\alpha, r)} E'} \quad (A \stackrel{def}{=} E)$
--	---	--

FIGURE 2. PEPA structured operational semantics.

**Cooperation** ( $\boxtimes$ ): We have already anticipated that the browser and the server in the example will be working together within the same system. This will require them to cooperate when the browser needs to download data which is not available locally. In contrast, the local activities of the browser can be carried out independently of the server. Cooperation over given actions is reflected in the cooperation by the cooperation set  $L = \{get, download, rel\}$  in this case. Actions in this set require the simultaneous involvement of both components. The resulting action, a shared action, will have the same type as the two contributory actions and a rate reflecting the rate of the action in the slowest participating component. This is discussed in more detail in the following subsection.

If, for simplicity, we assume that the system consists of just two browsers, the system is represented as the cooperation of the browsers and the server as follows:

$$WEB \stackrel{def}{=} (Browser \parallel Browser) \boxtimes_L Server$$

$$L = \{get, download, rel\}$$

The combinator  $\parallel$  is a degenerate form of the cooperation combinator, formed when two components behave completely independently, without any cooperation between them (i.e.  $L = \emptyset$ ), as in the case of the two browsers.

**Abstraction** ( $/$ ): It is often convenient to hide some actions, making them private to the component or components involved. The duration of the actions is unaffected, but their type becomes hidden, appearing instead as the unknown type  $\tau$ . Components cannot synchronize on  $\tau$ . For example, as we further develop the model of the information system we may wish to hide the access of a browser to its local cache. This might lead to a new representation of the browser:

$$Browser' \stackrel{def}{=} Browser / \{cache\}$$

and a corresponding new representation of the system:

$$WEB' \stackrel{def}{=} (Browser' \parallel Browser') \boxtimes_L Server$$

$$L = \{get, download, rel\}$$

The use of the hiding combinator in this way has two implications. First, it ensures that no components added to the model at a later stage can interact, or interfere, with this action of the browser. Second, private actions are deemed to have no contribution to the performance measures being calculated and this might subsequently suggest simplifications to the model.

Throughout the simple example above we have used constants (names) such as *Server* to associate identifiers with behaviours. Using recursive definitions we are able to describe components with infinite behaviours.

### 3.1.1. Cooperation

Communication or parallel composition is the essence of compositionality in process algebras. It gives structure to models, indicating which actions can be undertaken concurrently and which cannot. In PEPA there is no concept of conjugate or complementary actions, as there is CCS, because the aim was to capture something more general than a communication. Thus, there need not be strict input and output roles assigned to the participants. Instead, a multi-way synchronization framework is adopted as in CSP. This means that components or agents jointly perform actions of the same type, when the parallel composition dictates it. Note that in PEPA the cooperation combinator is in fact a family of combinators, since its meaning varies according to the contents of the cooperation set  $L$ .

Additional consideration is needed since the actions which are to be performed jointly may each have been assigned rates (durations) in their respective components. The issue of what it means for two timed actions to synchronize is a vexed one and the various SPAs have adopted a variety of solutions to this problem. This issue is discussed in [9] and in detail in Bradley's thesis [10].

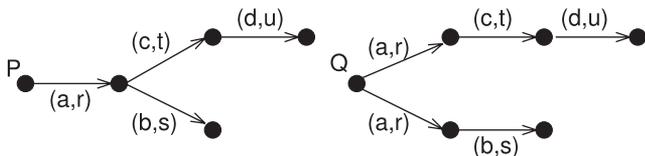
In PEPA it is assumed that each component has bounded capacity to carry out activities of any particular type, determined by the apparent rate. For a component  $P$  and action type  $\alpha$ , the apparent rate of  $\alpha$  in  $P$ , denoted  $r_\alpha(P)$ , is the sum of the rates of each action  $\alpha$  enabled in  $P$ . This corresponds to the rate at which  $P$  appears to an external observer to carry out an action  $\alpha$ , due to the superposition principle of the negative exponential distribution. The definition of cooperation in PEPA is based on the assumption that a component cannot be made to exceed its bounded capacity, meaning that the apparent rate of the shared action will be the minimum of the apparent rates of the components involved.

### 3.1.2. Semantics and equivalence relations

The semantic rules of PEPA generate a labelled transition system, just as in the case of classical process algebra. However, there are some significant differences introduced by the inclusion of quantified information. In particular, it is important to note that the semantics gives rise to a multi-transition system, i.e. it is not sufficient to record the existence of a transition or arc between two nodes. The multiplicity of the transition is also important. This is because the apparent rate of a term which has two copies of the same arc, generated by two instances of the same action, will differ from that of a term with only one instance.

Once a derivation graph has been generated for a particular model, this forms the basis of the underlying CTMC on which performance analysis will be carried out. A state of the CTMC is associated with each node of the graph, and the transition rate between states is simply the sum of the rates of actions labelling arcs between the corresponding nodes. Thus, during its evolution each syntactic term of the PEPA model corresponds to a state of the CTMC. It is established in [5] that this generates a unique Markov process.

PEPA has been equipped with a number of equivalence relations which have been shown to be useful for a variety of purposes [5]. The most significant is strong equivalence, sometimes termed Markovian bisimulation. Just as with the bisimulation for classical process algebra, the central notion here is that each of the pair of components should be able to mimic the behaviour of the other from the perspective of an external observer. This observer is now assumed to have the ability to time the behaviour over many repetitions and thus deduce information about the apparent rates of actions. This means that for components to be strongly equivalent they must have the same apparent rate for all action types. Therefore, if we consider again the two processes,  $P$  and  $Q$ , considered earlier, now enhanced with activity rates,



we can immediately deduce that they are not bisimilar because the apparent rates with respect to  $a$  in the initial

states are not the same. Note that this is a bisimulation in the same style as the bisimulation defined by Larsen and Skou for a probabilistic variant of CCS [11].

In 1960, Kemeny and Snell established that if we partition the state space of a CTMC and then form a new stochastic process in which each partition forms a state and the transition rate between states is the superposed transition rate of all transitions in one partition to the other, this stochastic process will satisfy the Markov property if and only if the partition has a property called lumpability [12]. An important property of strong equivalence in PEPA is that it induces a lumpable partition on the underlying CTMC. This forms the basis of an exact model reduction technique termed aggregation, which is discussed in the next subsection.

## 3.2. Manipulating models

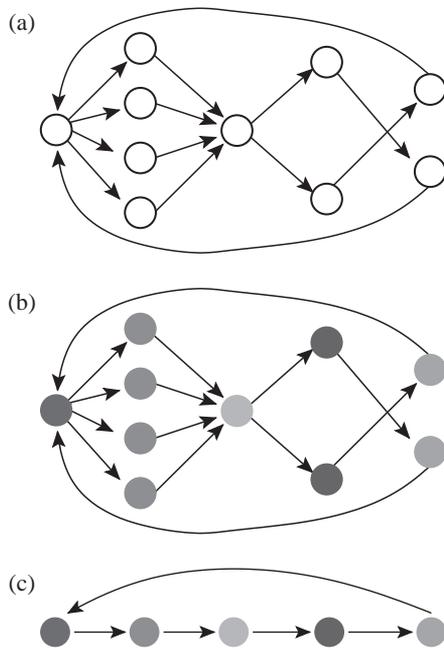
PEPA, like all state-based modelling techniques, suffers from problems of state space explosion. The compositionality of the process algebra can greatly aid model construction, but it can readily result in a model which is too large to be solved directly. As explained earlier, the study of techniques which can reduce the state space of models, or otherwise make them more amenable to solution, has been an active area of research in performance modelling for over 20 years. Two such techniques are model simplification and model aggregation, and the application of these techniques in the process algebra setting has been investigated. The challenge for PEPA has been to define such model manipulation techniques, in the context of the process algebra, in such a way that they can subsequently be applied automatically, based on the formally defined equivalence relations.

### 3.2.1. Model simplification

In model simplification the objective is to replace one model by another. Thus, an equivalence relation is used to establish behavioural or observational equivalence between models. The aim is to find a replacement model which is more desirable from a solution point of view, e.g. smaller state space, special class of the model, and so on. Once the desirable model has replaced the original, the underlying CTMC is generated as usual by associating one state with each node in the labelled transition system generated by the semantics. Equivalence relations that have been used in this way are weak isomorphism [5, 13] and strong bisimulation [5] in PEPA.

### 3.2.2. Model aggregation

In model aggregation the objective is to take a more abstract view of the system, and thus regard the model at a coarser granularity. Here, an equivalence relation is used to establish behavioural or observational equivalence between states within a model. In effect, this results in an alternative mapping from the labelled transition system, given by the semantics of the model, to the underlying CTMC. The equivalence relation is used to partition the nodes of the labelled transition system into equivalence classes. Then,



**FIGURE 3.** State-state equivalence for aggregation (a) original state space; (b) applying an equivalence relation to form a partition; (c) aggregate; reducing each partition to a single state.

instead of the usual one-to-one correspondence between nodes and states, one state in the CTMC is associated with each equivalence class of nodes (Figure 3) [14].

Establishing the link between strong equivalence and lumpability in the underlying Markov process forms the basis for an automatic procedure to reduce the size of models. Moreover, since the equivalence relation is a congruence the reduction can be applied to isolated components of the model, meaning that the state space of the complete model need never be constructed [14]. An algorithm to apply this technique on the fly, at the process algebra level, during state space generation has been developed and implemented for PEPA models [15].

### 3.3. Solving models

Despite the successes that have been gained in techniques for model manipulation, it still remains the case that in many instances the matrix characterizing the underlying CTMC and the corresponding steady state probability vector are simply too large to be readily stored on standard computing equipment. A variety of approaches to this problem for numerical solutions of CTMCs have appeared in the literature, including using disk-based storage [16], Kronecker and BDD-based representation of these entities [17, 18] and decomposed solutions of various forms [19].

Much of the work on PEPA has studied the use of decomposed solutions. It is clear that there is great advantage to be gained if the compositional structure of a PEPA model can be used during model solution, i.e. if the CTMCs corresponding to the components can be solved separately and their solutions combined to obtain a solution, exact or

approximate, of the whole CTMC. This work is discussed in this subsection.

In many cases the techniques which are applied are well known at the CTMC level. The advantage of characterizing the corresponding class of PEPA models is that by ‘lifting’ the definition from the stochastic process level to a formally defined high-level modelling paradigm we can facilitate the automatic detection of these structures when they occur, thus avoiding the construction of the original CTMC.

#### 3.3.1. Product form models

One class of CTMCs which are susceptible to an efficient solution technique are those which exhibit a product form equilibrium distribution. Consider a CTMC  $X(t)$ , whose state space  $\mathcal{S}$  is of the form  $\mathcal{S} \subseteq S_1 \times S_2$ , i.e. each state  $\mathbf{s} = s_1 \times s_2$  contains two pieces of information capturing different aspects of the current state. In general, these aspects may be dependent in many ways. When the process  $X(t)$  exhibits a product form solution, i.e. the steady state probability of an arbitrary state  $\mathbf{s}$ ,  $\pi(\mathbf{s})$ , can be expressed as  $\pi_1(s_1) \times \pi_2(s_2)$ , it indicates that these different aspects of the state description are independent with respect to the steady state.

Product form distributions have been widely used in the analysis of queueing networks and, owing to their efficient solution, have contributed to the popularity of queueing networks in performance analysis. For example, Jackson networks [20] and their generalization BCMP networks [21], have been widely employed. In these cases the underlying CTMC is known to have a reversible or quasi-reversible structure.

Work on finding PEPA models which give rise to product form solutions has drawn on the previous work on queueing networks. Essentially, this can be seen as an investigation of when components interact and yet remain statistically independent at steady state. It is clear that when a PEPA model consists of completely independent components, i.e.  $P \parallel Q$ , the steady state distribution will have product form:

$$\pi(P \parallel Q) = \pi_P(P) \times \pi_Q(Q)$$

where  $\pi_P$  and  $\pi_Q$  are the steady state distributions over the local states of  $P$  and  $Q$  respectively. However, few real systems consist of components which are independent in this way. The challenge has been to find circumstances in which components  $P$  and  $Q$  which interact, still exhibit statistical independence. A number of classes of such models have been identified.

*Reversible models.* Informally, a reversible Markov process is one which behaves identically when we observe it with time reversed as when we observe it with time flowing forward. More formally, an irreducible, stationary CTMC  $X(t)$  is reversible if it satisfies the detailed balance equations:

$$\pi(j)q(j, k) = \pi(k)q(k, j) \quad (1)$$

where  $q(j, k)$  is the instantaneous transition rate from state  $j$  to state  $k$  and  $\pi(\cdot)$  is the steady state probability distribution.

An initial study of the structure of the state space of SPA models giving rise to reversible CTMCs was presented by Bhabuta *et al.* in [22]. In [23], Hillston and Thomas, identify syntactic conditions which a PEPA model must satisfy in order for the underlying process to be reversible. The problem is tackled in two stages. First, a basic class of sequential components which give rise to reversible structures are identified. Then, assuming that a known class of reversible PEPA components exist, the authors investigate under what circumstances the conditions for reversibility will be preserved if reversible components are composed using the combinators of the PEPA language.

Fundamental to the basic class of reversible sequential components is the notion of a reverse pair. A pair of action types  $(\alpha, -\alpha)$  forms a reverse pair if, in any state, any  $\alpha$  transition leads to a state in which a  $-\alpha$  transition leads back to the original state. This ability to ‘undo’ any transition in the subsequent transition seems to be fundamental to reversibility. It is clear to see that this is a necessary condition for Equation (1) to be satisfied.

*Quasi-reversible models.* Like reversibility, quasi-reversibility originates in queueing theory. Formally, a stationary CTMC  $X(t)$  is quasi-reversible if for all times  $t_0$  the state  $X(t_0)$  is independent of

- (i) the input process after  $t_0$  and
- (ii) the output process before  $t_0$ .

Rather than the detailed balance equations which characterized reversibility, a quasi-reversible process satisfies partial balance equations:

$$\pi(i) \sum_{j \in S'} q(i, j) = \sum_{j \in S'} \pi(j) q(j, i) \quad (2)$$

for all states  $i$  and a corresponding subset of states  $S'$ . More details of the definition of quasi-reversibility can be found in [24].

In [25], a PEPA characterization of this class is presented. As in the work on reversibility, the approach is to first find simple instances of PEPA processes which give rise to quasi-reversible structures in their underlying CTMC. Then, conditions are established under which these components can be composed whilst maintaining the quasi-reversible property. Again the notion of a reverse pair is important and strong restrictions are placed on the interactions between components: each must be a flow cooperation. This means that the ‘positive’ half of a reverse pair in one component is carried out in cooperation with the ‘negative’ half of a reverse pair in another.

*Routing process models.* Sereno’s work, reported in [26], derives product form criteria for PEPA models based on earlier work on product form criteria for SPNs [27]. The SPN results rely on defining a Markov chain whose states correspond to the transitions of the SPN, the so-called routing chain. This chain exists only when severe restrictions are placed on the forms of synchronization and resource contention which can be represented in the net.

Sereno’s approach to PEPA is completely analogous to the earlier work on SPNs—he defines a Markov chain in which the states correspond to the actions of the SPA model. This is called the routing process. Sereno shows that if the state space of the routing process can be partitioned into equivalence classes of enabling actions (roughly speaking, one action enables another if the post-set of one is the pre-set of the other; we take the transitive closure of that relation), then a product form solution exists. Moreover, the partition forms the basis for the decomposition.

*Boucherie resource contention models.* Other classes of models have been considered in which the interaction between components is indirect, i.e. the components themselves are composed in parallel (without cooperation) in the model definition, but they compete over cooperation with a third component. Boucherie characterized a class of Markov chains which fit this framework. According to his definition, CTMCs which are otherwise independent compete for exclusive access to shared resources, causing blocking while a resource is held [28].

In [29] Hillston and Thomas characterize this class of CTMC in PEPA. As in the underlying CTMCs, the PEPA models consist of non-interacting components which give rise to the constituent processes of the underlying CTMC. These components compete, via synchronization with resource components. A PEPA component is termed a resource if it is never free to act independently. The general form of these process algebra terms and the resulting product form is, schematically:

$$\pi((P \parallel Q) \bowtie_L R) = B \times \pi_P(P \bowtie_L R) \times \pi_Q(Q \bowtie_L R)$$

where the component  $R$  represents the resource,  $\pi_P$  and  $\pi_Q$  are the steady state distributions over the derivatives of  $P \bowtie_L R$  and  $Q \bowtie_L R$  respectively, and  $B$  is the normalizing constant. The decomposition is formed by considering each of the model terms ( $P$  and  $Q$  in this case) acting in cooperation with the resource ( $R$ ) in isolation. Although presented here informally, these conditions are defined as formal syntactical conditions which can be checked on the model specification.

*Queueing discipline models.* In his PhD thesis [13], Clark defined a new combinator  $Q_{A,\xi}$  for PEPA which forces sequential components within its scope to observe a first-come-first-served (FCFS) discipline with respect to action types within the set  $A$ . Moreover, the rates of activities of those types are no longer controlled by the individual components but by the vector  $\xi$ . This is a derived combinator, meaning that any expression involving the combinator can be re-expressed using the existing PEPA combinators. In particular, for a set of components,  $S_1, \dots, S_n$ ,

$$Q_A(S_1, \dots, S_n) \equiv (S_1 \parallel \dots \parallel S_n) \bowtie_{M_\xi} R_\xi$$

for suitably chosen  $M_\xi$  and  $R_\xi$ .

This class of models is shown to be insensitive and therefore to have a product form solution so that the steady state probability of the complete model can be written as an expression involving the steady state probabilities of the individual models solved in isolation. In [30] it is established that this class of models is related to BCMP queueing networks, capturing infinite server and FCFS stations from the user's perspective.

Unlike the other classes which have been discussed above, characterization does not necessitate the definition of syntactic rules which may be used to check whether any model instance belongs to the class or not. Instead, the use of the derived combinator means that models can be constructed with a guaranteed product form solution.

*Quasi-separable models.* As with reversibility and quasi-reversibility, the notion of quasi-separability is one which has been developed in relation to queueing networks, in particular queueing networks in which breakdowns occur [31]. It is assumed that the CTMC is comprised of a number of components and that there are two pertinent pieces of information for each component. A representation of the whole process can then be formed as a pair of vectors, each vector capturing one piece of information for each component. For a process to be quasi-separable it must be possible to analyse the behaviour of a component, say component  $i$ , given the  $i$ th element of the first vector and all elements of the second, or vice versa. This allows the complete process to be reduced to a number of submodels, each of which contains all the information about exactly one component.

For such processes it is not possible to find the exact solution of the steady state distribution as a product of the local steady state distributions. Nevertheless, decomposed solutions can lead to exact results for the local steady state distributions and many performance measures.

In [32], Thomas and Gilmore present a characterization of PEPA models which are quasi-separable. It is assumed in this characterization that the information which must be included in each decomposed submodel is not distributed between the components but maintained by a single scheduler component. There are several conditions for the way in which this component may interact with the other components of the model. Furthermore, the individual components have no direct interaction between them—they must be in parallel composition with no synchronization, i.e. each of these components interacts only with the scheduler. The model is decomposed into a set of models, each comprising of a single component considered with the scheduler, in isolation.

Recent work on product form PEPA models has taken a slightly different form. In Harrison's work on the Reversed Compound Agent Theorem (RCAT) the process algebra has been used to establish a framework in which the relationships between different classes of product form CTMCs can be compared [33]. Within this framework Harrison has been able to demonstrate that the product forms which arise in Jackson networks [20] and G-networks [34] are based on the

same fundamental mechanisms: this becomes apparent when they are represented in PEPA.

### 3.3.2. Aggregated decomposed solutions

Product form models have the benefit of yielding the exact solution of the complete model. A variety of other techniques have been developed for decomposed solutions of CTMCs which impose less stringent conditions on the candidate models but which yield only approximate results. In many cases these are forms of aggregated decomposed solutions. Owing to the richer interactions between submodels it is not sufficient to only consider the submodels in isolation when forming a solution to the whole model. In addition, the decomposed solution involves a stochastic representation of the interactions between the submodels, the aggregated model. Some work has been done on considering PEPA, and other SPA, models in this framework.

*Time scale decomposition.* The work on time scale decomposition in SPA is based on the notion of near completely decomposable CTMCs [35] and inspired by previous work on time scale decomposition of SPN models [36]. A characterization of a near completely decomposable CTMC at the matrix level is that the matrix is block structured with elements in the diagonal blocks being at least an order of magnitude larger than elements in the off-diagonal blocks. This implies that the model is made up of subsystems whose internal interactions are much more frequent than the interactions between subsystems. As a consequence it can be assumed that the subsystems reach an internal equilibrium between external interactions.

The initial classification of SPA models susceptible to time scale decomposition [37], relied on a classification of the sequential components within a model into fast or slow; this in turn was based on a classification of all actions relative to some threshold rate. A component is considered to be fast if it enables fast or passive actions; a component is considered to be slow if it enables only slow actions. Only models comprised of fast and slow components were considered. Submodels were formed by allowing evolution only via fast actions. The aggregated model was formed by allowing evolution between such subsets of states via slow actions [37].

Later work by Mertsiotakis [38], tackles the problem of hybrid components—sequential components which cannot be classified as either fast or slow since they enable both fast and slow actions.

Related work on other SPAs includes:

- Mertsiotakis and Silva's work on decomposition of a class of SPA models, termed decision free processes [38, 39], based on an earlier work on throughput approximation in a class of SPNs called *marked graphs* [40].
- Bohnenkamp and Haverkort's work on near-independence [41], exploiting the notion of near-independence introduced by Ciardo and Trivedi in [42] in the context of SPNs.
- This was later expanded by the same authors in [43], which aims to reformulate the underlying Markov process of an SPA model as a set of semi-Markov

processes. These are then solved via their embedded Markov chains and evaluations of the time distributions between synchronizations between the SPA components.

#### 4. APPLICATIONS

Developing models for real applications has always been part of the PEPA project. This allows us to demonstrate to ourselves and others that the theory we have developed is useful. It is also a valuable source of inspiration for new theories and future directions.

In this section we provide a brief overview of some of the case studies which have been undertaken by ourselves and others using PEPA, followed by an account of some of the tools that support modelling with PEPA and then more detailed accounts of two recent projects which seek to make PEPA modelling more accessible to practitioners.

##### 4.1. Case studies

As originally intended, PEPA has been applied to study the performance characteristics of a number of computer and communication systems. Initial examples focussed on well-known standard performance evaluation abstractions such as multi-server, multi-queue systems [5] and various queueing systems [44]. However, over time more realistic case studies emerged, both from the PEPA group and from others. For example, in [45] the performance impact of fault-tolerant protocols within a distributed system framework is evaluated. In [46] Bowman *et al.* develop a model of multimedia traffic characteristics and use it to derive the quality of service measures such as jitter, throughput and latency. In an investigation of the ways in which to ease the development of parallel database systems, the STEADY group at the Heriot-Watt University proposed the use of performance estimators. PEPA was used to verify the output of the performance estimators for a number of particular configurations and therefore improve confidence in the approach [47].

In recent work a group at the PRISM Laboratory of the University of Versailles are working on a novel, active, rule-based approach to active networks (networks in which intermediate nodes supplement routing of data with some computation) [48]. A PEPA model was used to study the impact of the 'active' traffic on the non-active cross-traffic in terms of loss rate and latency within an active switch [49]. Furthermore, the models were validated against simulation models of the same system and showed very good agreement [50].

In addition, the formalism has been applied to a number of other problems which are beyond the usual arena of computer performance evaluation.

*Inland shipping.* Luk Knapen of Hasselt applied PEPA to study traffic flow within the inland shipping network of Belgium, focussing in particular on the locks and movable bridges.

*Robotic workcells.* Robert Holton of the University of Bradford used PEPA models to analyse the performance

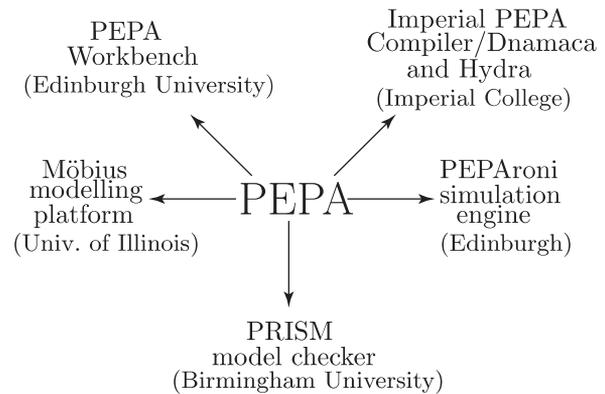


FIGURE 4. Tool support for PEPA modelling.

and functional correctness of a robotic workcell designed for an automated manufacturing system [51, 52].

*Cellular telephone networks.* A team from the PRISM Laboratory at the University of Versailles considered the problem of dimensioning in a cellular telephone network. They used a PEPA model to study the impact of allocating bandwidth resources between micro and macro-cell levels on call blocking and dropping [53], by taking advantage of automatic aggregation [15].

*Automotive diagnostic expert systems.* Console *et al.* of the University of Turin constructed a PEPA model of an automatic diagnostic system to be deployed in a car. A large number of sensors were placed around the car and some number could trigger an alarm. The role of the PEPA model was to provide probabilistic reasoning to resolve the likely cause of the alarm based on previous observations of the timing and frequency of individual faults [54].

##### 4.2. Tool support

Case studies of the size and complexity described in the previous subsection are only possible if the modelling process has adequate support. Fortunately, the PEPA formalism has been supported in a number of different tools offering a variety of different analysis techniques. Such support has been a strong factor in encouraging others to use the formalism.

The major tools which support PEPA are summarized in Figure 4 and briefly described in the subsections below.

###### 4.2.1. The PEPA Workbench

Initially developed in 1993, the PEPA Workbench has undergone several revisions, but has nevertheless maintained the same core functionality. It provides a parser which can apply the operational semantics to derive the derivation graph capturing all the possible evolutions of the model. It can render this as the infinitesimal generator matrix of a CTMC in formats suitable for both internal numerical solvers (biconjugate gradient algorithm or successive over-relaxation) or external numerical computing platforms such as Maple and Matlab. In addition, it includes facilities to

automatically derive some performance measures such as throughput, and a one-step debugger which can show the evolution of a model one activity at a time [55, 56].

The tool exists in two versions: the original version written in the functional programming language Standard ML, and a later edition written in Java [57]. The simple high-level tool was used to gain experience and insights. This was subsequently extended to a better engineered, more portable version. The ML edition of the PEPA Workbench is still used in this way, as a test bed for extensions of the PEPA language, such as PEPA nets [58], and for new algorithms [15]. In contrast, some less research-oriented extensions, such as the inclusion of transient solution facilities [59], are found only in the Java edition.

Generating the CTMC underlying a PEPA model and finding its steady state probability vector is rarely, if ever, the final objective of PEPA modelling. Formal tool support for querying performance models is an area which has received little attention until recently, despite its practical importance. Whilst some effort has been applied in this direction for PEPA models, it remains an area in which there is much scope for future work.

At the most basic level the modeller wishes to construct a reward structure over the state space of the CTMC, to be used in conjunction with the steady state probability vector to derive performance measures. For steady state measures the reward structure is a vector recording a ‘reward’ for each state, although for many states the reward value will be zero. Thus, the problem becomes one of identifying the appropriate set of states to attach a non-zero reward to. Clearly, when the CTMC arises from a stochastic process algebra model we prefer to characterize the state at the process algebra level. PEPA analysis tools have been developed which tackle this problem in two distinct ways.

*The PEPA State Finder.* The PEPA State Finder is used with the ML edition of the PEPA Workbench. It identifies subsets of states using regular expression pattern matching, applied to the table of PEPA expressions which make up the states of the model. Recall that there is a one-to-one correspondence between the syntactic forms of the PEPA process as it evolves and the states of the CTMC. The Workbench maintains a table to record this correspondence and using regular expression pattern matching the PEPA State Finder is able to extract the states of interest. For example, it is possible to use an expression such as  $* | (next, r) . *$  to return the state numbers of all the states in which the second component enables a  $(next, r)$  activity. This could then be used to construct a reward structure suitable for calculating the throughput of  $next$  in the second component: the value  $r$  is placed in the reward vector at each position corresponding to a (numerical) state found by the PEPA State Finder.

*PML<sub>μ</sub>.* A more sophisticated means of specifying rewards is described in Clark’s PhD thesis [13], and developed around the stochastic logic PML<sub>μ</sub>. Inspired by the probabilistic modal logic of Laren and Skou, PML [11], PML<sub>μ</sub> is able to differentiate PEPA terms which perform the same activities

but at different rates. The key to this is a modification to the Hennessy–Milner logic in which the diamond operator becomes decorated with a rate. The semantics of an expression in the logic is a subset of states, and thus logical expressions may be used, in conjunction with a value, to specify a reward structure. Clark extended the ML edition of the PEPA Workbench to include support for PML<sub>μ</sub> and associated reward structures [13].

#### 4.2.2. *The Imperial PEPA Compiler*

The recently developed Imperial PEPA Compiler (IPC) incorporates an alternative parser for PEPA models [60], thus providing a bridge to alternative analysis tools developed at Imperial College by Knottenbelt and his group [61, 62].

The IPC tool translates an input PEPA model into the Petri net notation provided by Dnamaca [61]; its support for the PEPA language is comprehensive. Apparent rates are supported, as are anonymous components. These are two advantages over the PEPA-to-PRISM compiler, and a richer class of PEPA models can be analysed by IPC/Dnamaca as a result.

The steady state probability distribution represents the behaviour of the system at equilibrium, where the influence of the initial state of the system is no longer measurable. Some performance measures of interest cannot be derived from the results of steady state analysis. Examples of performance measures in the class of non-equilibrium measurements include mean time to failure analysis, as computed in the evaluation of dependable systems. Other examples include the probabilistic quality-of-service guarantees which underpin most commercial service level agreements (SLAs): e.g. the probability that a 10-node cluster should be able to process 3000 database transactions in <6 s should be > 0.915; or a train service should not run > 10 min late >20% of the time.

More generally, such measures necessitate the computation of passage-time quantiles which detail the probability of passing through the system evolution from a start state to an end state (or set of starting states to a set of end states). The computation of such measures depends on the aggregate time behaviour across a whole system of complex interactions. The computation of passage-time quantiles depends on transient analysis of the CTMC, which is more expensive than steady state analysis at both run-time and memory consumption.

Via IPC, the unique solution capabilities of Dnamaca become available and because of this it is possible to efficiently perform passage-time analysis over PEPA models [60, 62]. Start and end points are specified using the concept of stochastic probes developed by Argent-Katwala, Bradley and Dingle [63]. Stochastic probes are themselves PEPA components which have been generated from regular expression-based inputs.

#### 4.2.3. *PEPAroni simulation engine*

Simulation has proved to be a useful alternative to numerical analysis of the underlying CTMC in two cases. First, if the size of the model is prohibitively large for numerical analysis

simulation can be used, although issues of run length can arise. Second, in the context of extending the expressiveness of PEPA with general distributions [13] numerical solution is no longer exact in most cases. Principally motivated by this second case, Clark implemented a simulator for PEPA models, called the PEPAroni simulator.

#### 4.2.4. The PRISM model checker

PRISM is a probabilistic model checker developed by Kwiatkowska's group at the University of Birmingham. It supports discrete time Markov chains and Markov decision processes as well as CTMCs. The standard input to PRISM is a model described in a simple reactive modules language. PEPA was integrated into the tool via a compiler which translates PEPA models into this language. There was also some work required to extend the model checker to support PEPA's combinators (cooperation and hiding).

Integration into PRISM enables model checking of the CTMC underlying a PEPA model against properties expressed in continuous stochastic logic [64]. It also provides access to the efficient numerical solutions of PRISM based on MTBDDs [18] and sparse matrix representation. PRISM has been applied successfully to a number of PEPA (and PEPA net) case studies [65, 66].

#### 4.2.5. The Möbius modelling platform

The Möbius modelling framework [67] was developed at the University of Illinois Urbana-Champaign. It is both a multi-formalism and multi-paradigm modelling tool, i.e. it aims to offer the user a choice of model description techniques and solution methods. Moreover, it is designed to allow a model to be composed of submodels which may be expressed in different formalisms. It has a broad spectrum of users in North America. Integrating PEPA into Möbius offered opportunities to present SPA to users who were previously unfamiliar with the formalism, and to explore the possibilities of interaction between modelling formalisms [68].

### 4.3. New application areas

In this section we describe two recent research projects which have sought, in different ways, to extend the applicability of PEPA. In both cases the use of PEPA becomes much more transparent to the users, who interact with the tools through other high-level system descriptions. Whilst the use of PEPA has been more widespread than perhaps anticipated, it remains the case that the majority of users are academic. Our objective is not to increase the user community of PEPA per se but to encourage the use of sound performance analysis as the basis for design and deployment decisions. Thus, in these latest developments the use of PEPA is in some ways hidden from the user.

#### 4.3.1. The DEGAS project

In the CEC-funded DEGAS project<sup>2</sup> we have been investigating ways to make performance modelling using

<sup>2</sup>Design Environments for Global ApplicationS project IST-2001-32072 funded by the FET Proactive Initiative on Global Computing.

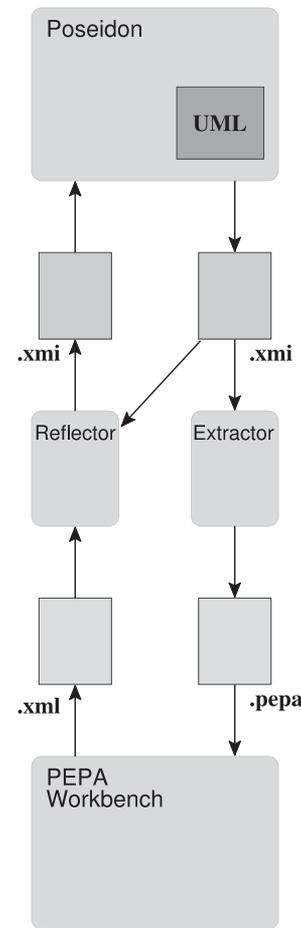


FIGURE 5. Architecture of the DEGAS analysis environment.

PEPA more accessible to software designers who may be unfamiliar with process algebra. We have sought to take advantage of the popularity of the unified modelling language (UML) by providing a way to derive performance measures from a suitably annotated UML model. A schematic view of this process is shown in Figure 5.

The key functionality is provided by a pair of software modules, the extractor and the reflector [69]. These form a bridge between the UML modelling environment and the PEPA tools. Annotations are added to the UML model according to a predetermined stereotype. The UML is then saved in the usual way in .xmi format. The extractor produces a corresponding .pepa format which can be loaded into the PEPA Workbench. This allows a steady state probability distribution corresponding to the states of the PEPA model to be derived. However, this is still inaccessible to the UML modeller—it is essential that results are reported in terms which make sense to the software designer, i.e. in terms of the original UML model. This functionality is provided by the reflector module which aggregates the steady state probability distribution data to produce suitable annotations to the UML model. This has required a modification to the PEPA Workbench so that results can be written out in XML format. The reflector then combines these with the original .xmi file.

This scheme is not specific to UML models. For example, an extractor has recently been developed for models developed in BPEL 4WS, a web service composition language [70].

#### 4.3.2. *The ENHANCE project*

The EPSRC-funded ENHANCE project<sup>3</sup> seeks to assist in the development of efficient applications in computational grid environments. The thrust of the project is two-fold.

- The use of high-level programming schemes such as Cole's skeletons facilitate correct high-level implementation of systems with complex coordination patterns. These patterns (the 'skeletons') are abstracted into a reusable, efficient library of coordination combinators. The use of this library elevates the programming model and eliminates low-level parallel programming errors such as sends with unmatched receives.
- Predefined performance models (PEPA components) corresponding to these skeletons, as well as the grid infrastructure, can be appropriately parameterized to represent any given configuration. Thus, alternative configurations can be evaluated prior to scheduling (and re-scheduling) in order to select the schedule with the best predicted performance.

A prototype tool, AMoGeT, supporting this framework has been developed [71]. This tool uses the Network Weather Service to query the available grid in order to obtain a short-term forecast of compute load on these machines. It assimilates this information with a PEPA model composed of the present program and grid infrastructure. In fact, a set of PEPA models are constructed, one for each candidate schedule. Solving these models within AMoGeT derives performance information which can be used to place processes on processors in order to reduce the makespan of the application overall.

## 5. FUTURE WORK

As in the past the plan is to continue the development of both new theory and new applications of PEPA in tandem. In this section we provide brief overviews of the current work in both areas.

### 5.1. New theory

#### 5.1.1. *PEPA nets*

Over the last decade mobility has had a major impact on the way we design, implement and manage many computer systems. Mobility may be manifest in the form of devices which change location and spontaneously connect/disconnect, or in the form of executable code which is moved around the network for a variety of reasons. In either case the effect is that the context in which computation is taking place is dynamically changing, and these changes

will affect the performance of the system. The PEPA nets formalism has been designed to capture information about mobility and so allow performance models of such systems to be readily and naturally developed [72].

A PEPA net is an SPN with coloured tokens. The tokens represent mobile objects with state and behaviour, where we use the term mobile loosely to characterize objects which may find themselves in different contexts during execution. The tokens are described using PEPA.

The use of SPNs for performance models is well established [73] and coloured variants, e.g. stochastic well-formed nets (SWNs) [74], have also been developed. However, the use of colours in PEPA nets offers something quite distinct—the possibility of differentiating between two types of change of state within a system. Unlike SWNs where tokens remain indistinguishable within their colour classes, tokens within PEPA nets are autonomous components. Firings of the net will typically be used to model macro-step (or global) changes of state, whereas transitions within the PEPA tokens are typically used to model micro-step (or local) changes of state as components undertake activities. Thus, modelling with PEPA nets uses both Petri nets and process algebra together as a single, structured performance modelling formalism. Moreover, we have demonstrated that PEPA nets offer some expressivity which is not directly offered by either PEPA or Petri nets [72]. Their modelling capabilities have been demonstrated in a number of case studies including MobileIP [75], a decentralized peer-to-peer emergency medical application [76] and a distributed multi-user role-playing game [77].

#### 5.1.2. *New mathematical models*

As detailed earlier, a major challenge for all discrete state-based performance modelling formalisms (i.e. queueing networks, SPNs and SPAs) is the problem of state space explosion. Much of the research effort of the performance analysis community for the last 20 years has been devoted to tackling this problem, using techniques such as those discussed in Section 3.3. However, recent work in the PEPA group has been studying an alternative approach. In this approach we move away from the assumption of discrete state space and instead consider continuous approximations. Where there are sufficient numbers of components, early indications are that this is a viable alternative to steady state Markovian analysis [78] for some problems.

When the state variables are assumed to be continuous, and activity rates are taken to be constant, the evolution of the system can be described by a set of ordinary differential equations. When this is generalized to allow activity rates to be governed by probability distributions rather than being deterministic the evolution of the system can be described by a set of random differential equations [79]. A further generalization, introducing more uncertainty, is offered by stochastic differential equations [80]. The investigation of the use of these alternative mathematical models is the subject of ongoing research.

<sup>3</sup>Enhancing the Performance Predictability of Grid Applications with Patterns and Process Algebras, EPSRC grant number GR/S21717/01.

## 5.2. New application domains

As already mentioned, the elegance of the PEPA modelling language, the available tool support and the access to quantified analysis have attracted researchers from other fields to develop PEPA models. This has led to the development of models in some novel areas. We conclude by outlining two of these areas below.

*Biochemical signalling pathways.* We have begun exploratory work in the area of biochemical signalling pathways using the ERK signalling pathway as an example [81]. A PEPA model has been developed to represent the pathway and was subsequently analysed using both standard Markovian analysis [82] and the continuous state space approximation technique described above [83].

In the longer term it is anticipated that new description languages and solution techniques will be needed for this new domain of application.

- Process algebras model a system by focusing on the activities and agents that undertake them. Whilst this view of a system comprised of individuals who interact to produce a collective behaviour has strong resonances within the biological community, the forms of activity and interaction which may be witnessed are more complex. For example, the interaction which captures inhibition without blocking is not readily represented in existing process algebras.
- Issues of scalability of analysis mechanisms have been a major concern within the work on performance evaluation, but the systems considered in the biological domain take this problem even further. In many instances existing techniques are not going to adequately cope, suggesting that new analysis techniques, or at least new mappings between the process algebra and analysis techniques, are needed.

*Security and timing attacks.* Many security analyses can be undertaken without information about the timing of message exchanges and therefore the system description does not need to record timing information. However, timing attacks are a form of security breach in which the timing, rather than just the ordering, of events is crucial.

Timing attacks can be mounted by timing the exchange of messages during a session between two parties. The parts of the communication which do not require user interaction have sufficient repeatability such that the timings of these interactions can be used to derive information. (Any user interaction usually introduces sufficient random delay to mask any information which could have been obtained from the timing information.) Where a secure session between two parties can be monitored by an eavesdropper, then information about the time parts of the interaction can be recorded. In the case where there is a difference in the time taken by different types of interactions (e.g. between a successful interaction and one that fails), then information can be leaked from the communication via the time recorded. One approach to this problem is to introduce delays into the

faster interactions in order to mask the difference between fast and slow interactions.

In [84] we have investigated the use of PEPA to analyse modified protocols to ensure that the introduced delays do indeed have the desired effect of masking the timing differences between interactions.

## ACKNOWLEDGEMENTS

Many people have contributed to the PEPA project in a variety of ways. In particular, the author would like to thank Ashok Argent-Katwala, Anne Benoit, Jeremy Bradley, Linda Brodo, Muffy Calder, Catherine Canevet, Graham Clark, Nick Dingle, Amani El-Rayes, Stephen Gilmore, Zully Grant-Duff, Valentin Haenel, Peter Harrison, Robert Holton, Jon Hunter, Leïla Kloul, Marta Kwiatkowska, Vassilis Mertsiotakis, Amdjed Mokhtari, Gethin Norman, Dave Parker, Corrado Priami, Matthew Prowse, Marina Ribaud, Matteo Sereno, Fotis Stathopoulos, Joanna Tomasik, Nigel Thomas and Feng Wan.

## TOOL AVAILABILITY

The PEPA tools are mostly available in open-source form under the GNU Public License. The PEPA Workbench and related tools are available from the PEPA web site at [www.dcs.ed.ac.uk/pepa](http://www.dcs.ed.ac.uk/pepa). The Möbius multi-paradigm modelling framework is available from the University of Illinois at [www.crhc.uiuc.edu/PERFORM/mobius-software.html](http://www.crhc.uiuc.edu/PERFORM/mobius-software.html). PRISM is available from [www.cs.bham.ac.uk/~dyp/prism/](http://www.cs.bham.ac.uk/~dyp/prism/). The IPC is available from [www.doc.ic.ac.uk/ipc/](http://www.doc.ic.ac.uk/ipc/).

## REFERENCES

- [1] Milner, R. (1989) *Communication and Concurrency*. Prentice-Hall.
- [2] Hoare, C. A. R. (1985) *Communicating Sequential Processes*. Prentice-Hall.
- [3] Plotkin, G. D. (1981) *A structural approach to operational semantics*. Technical Report DAIMI FN-19, University of Aarhus.
- [4] Herzog, U. (1990) *Formal description, time and performance analysis: a framework*. Technical Report 15/90, IMMD VII, Friedrich-Alexander-Universität, Erlangen-Nürnberg, Germany.
- [5] Hillston, J. (1996) *A Compositional Approach to Performance Modelling*. Cambridge University Press.
- [6] Götz, N., Herzog, U. and Rettelbach, M. (1992) *TIPP—a language for timed processes and performance evaluation*. Technical Report 4/92, IMMD7, University of Erlangen-Nürnberg, Germany.
- [7] Bernardo, M. and Gorrieri, R. (1998) A tutorial on EMPA: a theory of concurrent processes with nondet erminism, priorities, probabilities and time. *TCS*, **202**, 1–54.
- [8] Strulo, B. and Harrison, P.G. (2000) Spades—a process algebra for discrete event simulation. *J. Logic Comput.*, **10**(1), 3–42.
- [9] Hillston, J. (1994) The nature of synchronisation. In Herzog, U. and Rettelbach, M. (eds) *Proc. 2nd Int. Workshop on*

- Process Algebras and Performance Modelling*, Erlangen, July, pp. 51–70.
- [10] Bradley, J. (1999) Towards Reliable Modelling with Stochastic Process Algebras. PhD thesis, Department of Computer Science, University of Bristol.
- [11] Larsen, K. and Skou, A. (1991) Bisimulation through probabilistic testing. *Inform. Comput.*, **94**(1), 1–28.
- [12] Kemeny, J. G. and Snell, J. L. (1960) *Finite Markov Chains*. Van Nostrand.
- [13] Clark, G. (2000) Techniques for the Construction and Analysis of Algebraic Performance Models. PhD thesis, The University of Edinburgh.
- [14] Hillston, J. (1995) Compositional markovian modelling using a process algebra. In Stewart, W. J. (ed) *Numerical Solution of Markov Chains*. Kluwer.
- [15] Gilmore, S., Hillston, J. and Ribaud, M. (2001) An efficient algorithm for aggregating PEPA models. *IEEE Trans. Software Eng.*, **27**(5), 449–464.
- [16] Deavours, D. D. and Sanders, W. H. (1998) An efficient disk-based tool for solving large Markov models. *Perform. Evaluation*, **33**, 67–84.
- [17] Hillston, J. and Kloul, L. (2001) An efficient Kronecker representation for PEPA models. In de Alfaro, L. and Gilmore, S. (eds) *Proc. First Joint PAPM-PROBMIV Workshop*, Aachen, Germany, September 12–14, *LNCS 2165*, Springer-Verlag.
- [18] Hermanns, H., Meyer-Kayser, J. and Siegle, M. (1999) Multi-terminal binary decision diagrams to represent and analyse continuous time Markov chains. In *Proc. 3rd Int. Workshop on the Numerical Solution of Markov Chains*, Zaragoza, Spain, September 6–10, pp. 188–207.
- [19] Hillston, J. (2001) Exploiting structure in solution: decomposing composed models, *FMPA Lecture Notes*. Springer-Verlag.
- [20] Jackson, J. R. (1963) Jobshop-like queueing systems. *Manage. Sci.*, **10**, 131–142.
- [21] Baskett, F., Chandy, K. M., Muntz, R. R. and Palacios, F. G. (1975) Open, closed and mixed networks of queues with different classes of customers. *J. ACM*, **22**(2), 248–260.
- [22] Bhabuta, M., Harrison, P. G. and Kanani, K. (1995) Detecting reversibility in Markovian Process Algebra. In *Performance Engineering of Computer and Telecommunications Systems*, Liverpool John Moores University, Springer-Verlag.
- [23] Hillston, J. and Thomas, N. (1998) A syntactical analysis of reversible PEPA models. In *Proc. 6th Process Algebra and Performance Modelling Workshop*, Nice, France, September 11–12. University of Verona.
- [24] Kelly, F. (1979) *Reversibility and Stochastic Processes*. Wiley.
- [25] Harrison, P. and Hillston, J. (1995) Exploiting quasi-reversible structures in Markovian process algebra models. *Comput. J.*, **38**(6), 510–520.
- [26] Sereno, M. (1995) Towards a product form solution of stochastic process algebras. *Comput. J.*, **38**(6), 622–632.
- [27] Henderson, W. and Taylor, P. G. (1991) Embedded processes in stochastic petri nets. *IEEE Trans. Softw. Eng.*, **17**(2), 108–116.
- [28] Boucherie, R. J. (1994) A Characterisation of Independence for Competing Markov Chains with Applications to Stochastic Petri Nets. *IEEE Trans. Softw. Eng.*, **20**(7), 536–544.
- [29] Hillston, J. and Thomas, N. (1999) Product form solution for a class of PEPA models. *Perform. Evaluation*, **35**(3–4), 171–192.
- [30] Clark, G. and Hillston, J. (2002) Product form solution for an insensitive stochastic process algebra structure. *Perform. Evaluation*, **50**(2–3), 129–151.
- [31] Mitrani, I. and Wright, P. E. (1994) Routing in the presence of breakdowns. *Perform. Evaluation*, **20**, 151–164.
- [32] Thomas, N. and Gilmore, S. (1998) Applying quasi-separability to markovian process algebra. In *Proc. 6th Process Algebra and Performance Modelling Workshop*, Nice, France, September 11–12. University of Verona.
- [33] Harrison, P. G. (2003) Turning back time in Markovian process algebra. *TCS*, **290**, 1947–1986.
- [34] Gelenbe, E. (1991) Queueing networks with negative and positive customers. *J. Appl. Probab.*, **28**, 656–663.
- [35] Courtois, P. J. (1977) *Decomposability: Queueing and Computer System Applications*. ACM Series. Academic Press, New York.
- [36] Blakemore, A. and Tripathi, S. (1993) Automated time scale decomposition of SPNs. In *Proc. of 5th International Workshop on Petri Nets and Performance Models (PNPM '93)*, Toulouse, October 20–22.
- [37] Hillston, J. and Mertsiotakis, V. (1995) A simple time scale decomposition technique for stochastic process algebras. *Comput. J.*, **38**(7), 566–577.
- [38] Mertsiotakis, V. (1998) Approximate Analysis Methods for Stochastic Process Algebras. PhD thesis, Universität Erlangen–Nürnberg, Martensstraße 3, 91058 Erlangen.
- [39] Mertsiotakis, V. and Silva, M. (1997) A throughput approximation algorithm for decision free processes. In Ribaud, M. (ed.) *Proc. 7th Int. Workshop on Petri Nets and Performance Models*, St Malo, France, June 3–6.
- [40] Jungnitz, H. (1992) Approximation Methods for Stochastic Petri Nets. PhD thesis, Rensselaer Polytechnic Institute.
- [41] Bohnenkamp, H. and Haverkort, B. (1997) Decomposition methods for the solution of stochastic process algebra models: a proposal. In Brinksma, E. and Nymeyer, A. (eds) *Proc. 5th Process Algebra and Performance Modelling Workshop*, Enschede, The Netherlands, June 26–27.
- [42] Ciardo, G. and Trivedi, K. S. (1992) A decomposition approach for stochastic petri net models. *Perfor. Evaluation*, **18**, 37–59.
- [43] Bohnenkamp, H. and Haverkort, B. (1998) Semi-Numerical Solution of Stochastic Process Algebra Models. In Priami, C. (ed.) *Proc. 6th Process Algebra and Performance Modelling Workshop*, Nice, France, September 11–12.
- [44] Thomas, N. and Hillston, J. (1997) *Using Markovian process algebra to specify interactions in queueing systems*. Technical Report ECS-LFCS-97-373, LFCS, The University of Edinburgh.
- [45] Clark, G., Gilmore, S., Hillston, J. and Ribaud, M. (2000) Exploiting modal logic to express performance measures. In *Proc. 11th Int. Conf. Computer Performance Evaluation: Modelling Techniques and Tools*, Schaumburg, Illinois, USA, March 27–31, *LNCS 1786*, pp. 211–227. Springer-Verlag.
- [46] Bowman, H., Bryans, J. and Derrick, J. (1998) Analysis of a multimedia stream using stochastic process algebra. In Priami, C. (ed.) *6th Int. Workshop on Process Algebras and Performance Modelling*, Nice, France, September 11–12, pp. 51–69.

- [47] Dempster, E. W., Tomov, N. T., Lü, J., Pua, C. S., Williams, M. H., Burger, A., Taylor, H. and Broughton, P. (1998) Verifying a performance estimator for parallel DBMSs. In *Proc. EuroPar (EuroPar'98)*, Southampton, September 2–4.
- [48] Bouzeghoub, M., Kloul, L. and Mokhtari, A. (2002) *A new active network framework based on active rules*. Technical Report 2002/21, PRiSM, Université de Versailles.
- [49] Hillston, J., Kloul, L. and Mokhtari, A. (2003) Active nodes performance analysis using PEPA. In *Proc. 19th UK Performance Engineering Workshop*, University of Warwick, July 9–10, pp. 244–256.
- [50] Hillston, J., Kloul, L. and Mokhtari, A. (2004) Towards a feasible active networking scenario. *Telecommun. Sys.*, **27**(2–4), 413–438.
- [51] Gilmore, S., Hillston, J., Holton, D. R. W. and Rettelbach, M. (1996) Specifications in Stochastic Process Algebra for a Robot Control Problem. *Int. J. Prod. Res.*, **34**(4), 1065–1080.
- [52] Holton, D. R. W. (1995) A PEPA specification of an industrial production cell. *Comput. J.*, **38**(7), 542–551.
- [53] Forneau, J. M., Kloul, L. and Valois, F. (2002) Performance modelling of hierarchical cellular networks using PEPA. *Perform. Evaluation*, **50**(2–3), 83–99.
- [54] Console, L., Picardi, C. and Ribaud, M. (2000) Diagnosis and diagnosability analysis using process algebras. In *Proc. 11th Int. Workshop on Principles of Diagnosis (DX00)*, Morelia, Mexico, June 8–11.
- [55] Gilmore, S. and Hillston, J. (1994) The PEPA workbench: a tool to support a process algebra-based approach to performance modelling. In *Proc. 7th Int. Conf. on Modelling Techniques and Tools for Computer Performance Evaluation*, Vienna, Austria, May 3–6, *LNCS 794*, pp. 353–368. Springer-Verlag.
- [56] Clark, G., Gilmore, S. and Hillston, J. (1999) The PEPA performance modelling tools. In Hillston, J. (ed.) *Proc. 7th Workshop on Process Algebra and Performance Modelling*, Zaragoza, Spain, September 6–10. University of Zaragoza Press.
- [57] Hunter, J. (1999) Re-evaluation of the PEPA Workbench. Master's thesis, School of Computer Science, The University of Edinburgh.
- [58] Gilmore, S., Hillston, J. and Ribaud, M. (2002) PEPA nets: a structured performance modelling formalism. In *Proc. 12th Int. Conf. on Modelling Tools and Techniques for Computer and Communication System Performance Evaluation*, London, UK, April 14–17, *LNCS 2324*, pp. 111–130. Springer-Verlag.
- [59] Wan, F. (2000) Interface Engineering and Transient Analysis for the PEPA Workbench. Master's thesis, School of Computer Science, The University of Edinburgh, UK.
- [60] Bradley, J. T., Dingle, N. J., Gilmore, S. T. and Knottenbelt, W. J. (2003) Derivation of passage-time densities in PEPA models using IPC: the imperial PEPA compiler. In *Proc. 11th IEEE/ACM Int Symp. on Modeling, Analysis and Simulation of Computer and Telecommunications Systems*, Orlando, FL, October 12–15, pp. 344–351. IEEE Computer Society Press.
- [61] Knottenbelt, W. J. (1996) Generalised Markovian analysis of timed transition systems. Master's thesis, University of Cape Town, South Africa.
- [62] Bradley, J. T., Dingle, N. J., Gilmore, S. T. and Knottenbelt, W. J. (2003) Extracting passage times from PEPA models with the HYDRA tool: a case study. In Jarvis, S. (ed.) (2003) *Proc. of 19th UK Performance Engineering Workshop*, University of Warwick, July 9–10, pp. 79–90.
- [63] Argent-Katwala, A., Bradley, J. T. and Dingle, N. J. (2004) Expressing performance requirements using regular expressions to specify stochastic probes over process algebra models. In *Proc. 4th Int. Workshop on Software and Performance*, Redwood Shores, CA, January 14–16, pp. 49–58. ACM Press.
- [64] Kwiatkowska, M., Norman, G. and Parker, D. (2002) PRISM: Probabilistic symbolic model checker. In *Proc. of 12th Int. Conf. on Modelling Tools and Techniques for Computer and Communication System Performance Evaluation*, London, UK, April 14–17, *LNCS 2324*, pp. 200–204. Springer-Verlag.
- [65] Gilmore, S. and Kloul, L. (2003) A unified tool for performance modelling and prediction. In *Proc. 22nd Int Conf. on Computer Safety, Reliability and Security (SAFECOMP 2003)*, Potsdam, Germany, September 21–24, *LNCS 2788*, pp. 179–192. Springer-Verlag.
- [66] Gilmore, S., Hillston, J., Kloul, L. and Ribaud, M. (2004). Software performance modelling using PEPA nets. In *Proc. 4th Int. Workshop on Software and Performance*, Redwood Shores, CA, January 14–16, pp. 13–24. ACM Press.
- [67] Clark, G., Courtney, T., Daly, D., Deavours, D., Derisavi, S., Doyle, J. M., Sanders, W. H. and Webster, P. (2001) The Möbius modeling tool. In *Proc. 9th Int. Workshop on Petri Nets and Performance Models*, Aachen, Germany, September 11–14, pp. 241–250.
- [68] Clark, G. and Sanders, W. H. (2001) Implementing a stochastic process algebra within the Möbius modeling framework. In de Alfaro, L. and Gilmore, S. (eds) *Proc. First Joint PAPM-PROBMIV Workshop*, Aachen, Germany, September 12–14, *LNCS 2165*, pp. 200–215. Springer-Verlag.
- [69] Canevet, C., Gilmore, S., Hillston, J., Prowse, M. and Stevens, P. (2003) Performance modelling with UML and stochastic process algebras. *IEE Proc. Comput. Dig. Tech.*, **150**(2), 107–120.
- [70] Mitchell, B. and Hillston, J. (2004) Analysing web service composition with PEPA. In *Proc. 3rd Workshop on Process Algebras and Stochastically Timed Activities*, Edinburgh, UK June 10–11, pp. 33–44.
- [71] Benoit, A., Cole, M., Gilmore, S. and Hillston, J. (2005) Realistic performance evaluation of skeleton-based grid applications using the Network Weather Service. *Comput. J.* (to appear).
- [72] Gilmore, S., Hillston, J., Ribaud, M. and Kloul, L. (2003) PEPA nets: A structured performance modelling formalism. *Perform. Evaluation*, **54**(2), 79–104.
- [73] Ajmone Marsan, M., Balbo, G., Conte, G., Donatelli, S. and Franceschinis, G. (1995) *Modelling with Generalized Stochastic Petri Nets*. John Wiley.
- [74] Chiola, G., Duthellet, C., Franceschinis, G. and Haddad, S. (1993) Stochastic well-formed colored nets and symmetric modeling applications. *IEEE Trans. on Comput.*, **42**(11), 1343–1360.
- [75] Gilmore, S., Hillston, J. and Kloul, L. (2004). PEPA nets. In Calzarossa, M. C. and Gelenbe, E. (eds) *Performance Tools and Applications to Networked Systems: Revised Tutorial Lectures*, *LNCS 2965*, pp. 311–335. Springer-Verlag.

- [76] Gilmore, S., Haenel, V., Hillston, J. and Kloul, L. (2004). PEPA nets in practice: modelling a decentralised peer-to-peer emergency medical application. In Núñez, M. *et al.* (eds) *Applying Formal Methods: Testing, Performance, and M/E-Commerce, Proc. EPEW 2004*, Toledo, Spain, September 30–October 2, *LNCS 3236*, pp. 262–277. Springer-Verlag.
- [77] Gilmore, S., Kloul, L. and Piazza, D. (2004) Modelling role-playing games using PEPA nets. In *Proc. 19th Int. Symp. on Computer and Information Sciences (ISCIS 2004)*, Kemer-Antalya, Turkey, October 27–29, *LNCS 3280*, pp. 523–532. Springer-Verlag.
- [78] Benoit, A., Cole, M., Gilmore, S. and Hillston, J. (2005) Enhancing the effective utilisation of grid clusters by exploiting on-line performability analysis. In *Proc. 2nd Grid Performability Workshop*, 2005. to appear.
- [79] Soong, T. T. (1973) *Random Differential Equations in Science and Engineering*. Academic Press.
- [80] Oksendal, B. K. (2003) *Stochastic Differential Equations*. Springer.
- [81] Cho, K.-H., Shin, S.-Y., Kim, H.-W., Wolkenhauer, O., McFerran, B. and Kolch, W. (2003) Mathematical modeling of the influence of RKIP on the ERK signaling pathway. In Priami, C. (ed.) *Computational Methods in Systems Biology (CSMB'03)*, Rovereto, Italy, February 24–26, *LNCS 2602*, pp. 127–141. Springer-Verlag.
- [82] Calder, M., Gilmore, S. and Hillston, J. (2004) Modelling the influence of RKIP on the ERK signaling pathway using the stochastic process algebra PEPA. In *Proc. BioConcur'04*, London, England, August 30.
- [83] Calder, M., Gilmore, S. and Hillston, J. (2005) Automatically deriving ODEs from process algebra models of signalling pathways. To appear in *CMSB'05*, Edinburgh, Scotland, April 3–5.
- [84] Buchholtz, M., Gilmore, S., Hillston, J. and Nielson, F. (2004) Securing statically-verified communications protocols against timing attacks. In *Proc. 1st Int. Workshop on Practical Applications of Stochastic Modelling*, Royal Society, London, September 4, pp. 61–79.