

A Feature Construct for PEPA

Stephen Gilmore and Jane Hillston
Laboratory for Foundations of Computer Science
The University of Edinburgh
{stg, jeh}@dcs.ed.ac.uk

April 30, 1999

Abstract

We show how the PEPA performance modelling language could be extended with a *feature construct* which can be used to describe modifications to PEPA models. We provide this construct with an operational description which conservatively extends the operational semantics of the PEPA language. We apply the feature construct in a small case study.

1 Introduction

PEPA (Performance Evaluation Process Algebra) [1] is a performance modelling notation. It is also a process algebra, a concise mathematical language which is amenable to formal reasoning. PEPA is defined by an unambiguous semantics which makes clear the meaning of all models which are expressed in the language. It has been used to investigate the behaviour and performance of a diversity of distributed and concurrent systems [2, 3, 4, 5, 6, 7].

Constructing performance models of distributed systems is a worthwhile activity. Distributed systems often have designs which are both complex and novel. An ill-considered design decision can lead to an implementation which fails to achieve planned levels of service or has unnecessarily high running costs.

As is the case for other performance modelling notations, PEPA can be applied either prospectively to assess the viability of a candidate design for a yet-to-be-constructed system or retrospectively to provide insight into a fully functional operation. Our novel contribution in this paper is to show how the PEPA language could be extended with syntactic support for the description

of enhancements to models which reflect enhancements to the system under study. We use the term *features* to describe these enhancements to both systems and models, in keeping with the use of this term in the telecommunications industry and in software development. With this extension, the PEPA language can more easily be used throughout the entire lifetime of a complex distributed system, tracking adaptive and corrective maintenance in a formal setting.

Improving the suitability of the PEPA language for modelling complex systems is a useful extension. The difficulty in computer system development stems from the desire to create sophisticated and comprehensive products. These are constructed piece by piece and are subject to many unpredictable revisions over time. This is our motivation for considering a formal means of expressing the addition of new features to an existing system. The fact that a system has interesting performance qualities which are worth investigating via the construction and analysis of system models does not make it impervious to modifications and the addition of new features.

One way to address the problem of modelling the addition of features is to modify the extant system description in such a way that the new feature is incorporated or interwoven into the model description. This might at first seem to be an attractive option since it does not require any modification to the existing modelling language which was used to describe the original system. The modeller can simply think that they are only constructing a model of a more complex system, namely the one which includes the added feature. However, this approach has the disadvantage that it reduces the intellectual leverage which the identification of features gives both to system designers and system builders. Firstly, the valuable documentation function which features provide has not been exploited with the consequence that no formal record of the change history of the system is being created. Secondly, the dependency of the added feature upon components of the existing system will be unclear. Effectively, *all* of the existing system description has been considered to be essential to the description of the new feature. It is not usual that this is the case in practice.

We adopt the principle that feature descriptions require a different form of expression from system descriptions. Particularly, their descriptions should make clear the dependencies of components of the existing system. This places a demand upon our modelling language to provide some distinctive syntactic support for the formal expression of features. Other authors have also argued that the addition of a feature construct to an existing modelling language is the right method by which to make progress in this problem [8, 9].

2 Design of the feature construct

PEPA is a small language with essential, simple combinators. A description appears in Appendix A. Briefly, PEPA *components* perform *timed activities*. Each PEPA model defines a labelled multi-transition system which can be read as a Continuous Time Markov Chain (CTMC) by ignoring the activity names which label the arcs from one state to another.

A feature construct for PEPA must add value without incurring unnecessary loss of simplicity. (The feature construct can itself be seen as a *feature* which is being added to the existing PEPA language. Many of the good practices which are applied when adding functional features to software systems have their analogues here where we are adding a model structuring feature to an existing modelling language.)

Following [8], we consider that a feature construct should describe features formally as self-contained units of functionality. It should be possible to consider features in isolation, without complete knowledge of the system to which they are being added. A feature construct should be general-purpose, allowing a number of different types of features to be added. However, it should not allow undisciplined modifications which would be hard to reason about or understand. This last requirement would rule out of consideration as candidates a number of powerful, but low-level, macro-like operators.

In the particular setting of the PEPA modelling language we want a feature construct which is applicable, general, clear and easy to explain. It must have a formal definition. Two candidates present themselves as being possibly suitable; *re-binding* and *parametric definition*.

2.1 Re-binding component definitions

PEPA is a *compositional* description language so in principle new features could be installed by re-binding the definitions of key components. Such an extension to the PEPA language would meet most of our criteria for a feature construct. Re-binding definitions is a general-purpose concept and it can certainly be described both formally and with clarity. Unfortunately, it fails to meet our key criteria of applicability because it cannot describe the most general case of making unforeseen extensions to an existing system. The use of re-binding as a feature construct is only applicable in the cases where the designers of the system have previously loaded the system with re-programmable hooks. It is usual to describe such systems as *feature-ready* because they have been designed in the anticipation of the addition of new features in particular ways. Extensible software systems such as Web browsers with a “plug-in” capability are an example of feature-ready systems.

2.2 Parametric component definitions

Our winning candidate for a feature construct for PEPA is the use of parameterised components. The parameters capture the dependency of the feature on the existing system. By defining our new feature in terms of the existing system, monitored for essential behaviour, we allow for system reconfiguration and the introduction of new components.

3 Semantics of the feature construct

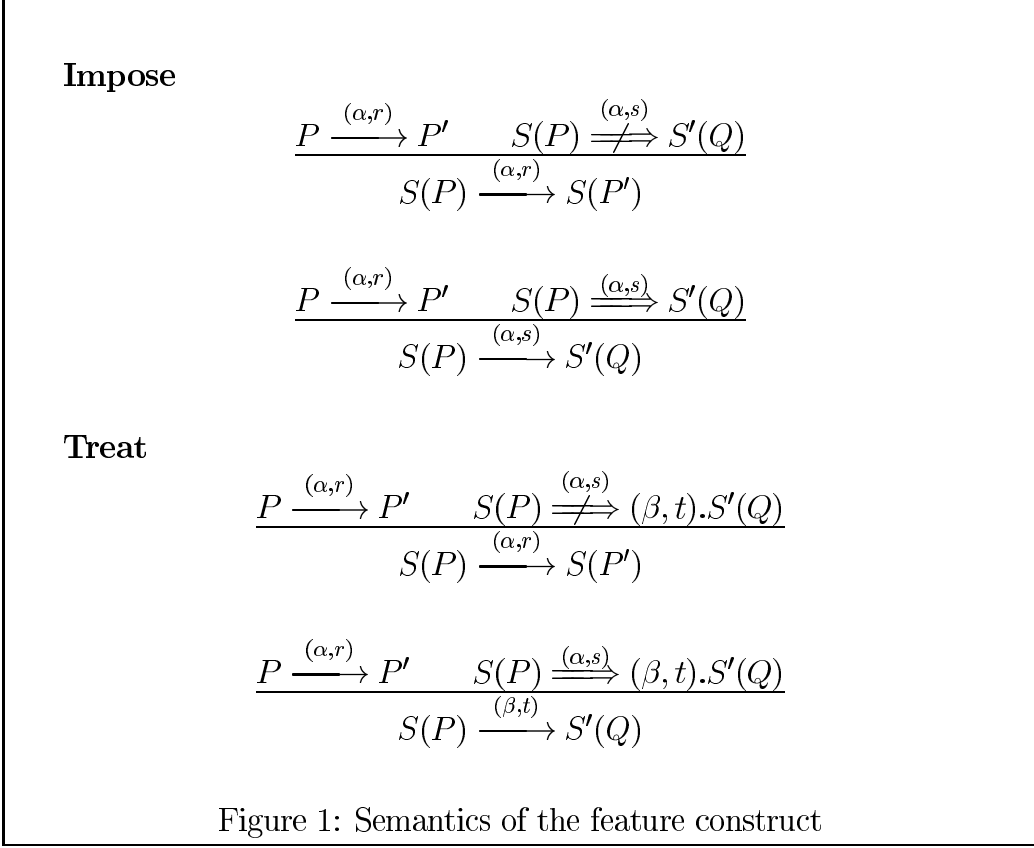
Simply put, we allow components to be parameterised by other components. The parameterisation is first-order, that is we do not allow parameterised components to be passed as parameters. More formally, parameters range over PEPA expressions extended with formal parameter identifiers in addition to the identifiers used for PEPA constants. To preclude syntactic ambiguity, we use the convention that if a constant appears in a formal parameter expression then it denotes the component bound to that constant name and it is not a re-use of that identifier with another meaning in the body of the parameterised component definition. Thus it is not possible to make a hole in the scope of a component definition by re-using the identifier of that component as the identifier of a formal parameter.

The operational definition of the feature construct presented in Figure 1 is built structurally from the transition relation on PEPA components (written \longrightarrow) and the transitions of parameterised components (written \Longrightarrow). The definition of the transition relation for PEPA components is in Figure 6 in Appendix A. The definition of the transition relation for parameterised components is in Figure 2. It depends on the definition given in Figure 3 of the judgement relation for contexts (written \vdash).

3.1 Impose and treat

A useful separation of concerns in describing feature integration is the distinction between *imposing* new behaviour on the system and *treating* the existing behaviour in a new way. In a state-based modelling approach such as ours this divides into a concern in the former case that the system should not reach certain states and in the later case that the system should not perform certain activities.

In both cases there are two possibilities. Either the transition which the existing system would perform has a matching transition in the transitions of the parameterised component or it does not (written $\not\Rightarrow$). In the cases where there is a matching transition the new state of the system is dictated



by the parameterised component as Q replaces P' . The effect of the Treat rule is to replace an existing activity α by a new activity β .

It is easy to see that the expressiveness of the Treat rule subsumes that of the Impose rule. When the metavariables α and β denote the same PEPA activity and the metavariables s and t denote the same PEPA rate then the Treat rule expresses the same adaptation of the existing system behaviour as the Impose rule. It is also easy to see that the Impose rule by itself is not sufficient because it is not possible to prevent the existing system model from performing its *first* activity. This is a limitation because we wish to be able to impose new behaviour on *any* state of the existing system, even the initial state. The Treat rule does not suffer from this limitation.

Since the Treat rule subsumes the Impose rule, and is applicable in more situations, why then do we keep the Impose rule at all? The reason is that in practice it gives the most convenient form of expression to new features. The more general Treat rule would need to be used almost always to simulate the Impose rule, so an additional rule seems to be worthwhile.

Overriding

$$\frac{}{C \vdash (\alpha, r).S(E) \xrightarrow{(\alpha, r)} S(E)}$$
$$\frac{}{C \vdash (\alpha, r).(\beta, t).S(E) \xrightarrow{(\alpha, r)} (\beta, t).S(E)}$$

Option

$$\frac{C \vdash E_1 \xrightarrow{(\alpha, r)} S(E)}{C \vdash E_1 + E_2 \xrightarrow{(\alpha, r)} S(E)} \quad \frac{C \vdash E_2 \xrightarrow{(\alpha, r)} S(E)}{C \vdash E_1 + E_2 \xrightarrow{(\alpha, r)} S(E)}$$
$$\frac{C \vdash E_1 \xrightarrow{(\alpha, r)} (\beta, t).S(E)}{C \vdash E_1 + E_2 \xrightarrow{(\alpha, r)} (\beta, t).S(E)} \quad \frac{C \vdash E_2 \xrightarrow{(\alpha, r)} (\beta, t).S(E)}{C \vdash E_1 + E_2 \xrightarrow{(\alpha, r)} (\beta, t).S(E)}$$

Application

$$\frac{C \vdash E' \xrightarrow{(\alpha, r)} S'(F)}{S(P) \xrightarrow{(\alpha, r)} S'(Q)} \quad (S(E) \stackrel{\text{def}}{=} E' \text{ and } C \vdash E = P, F = Q)$$
$$\frac{C \vdash E' \xrightarrow{(\alpha, r)} (\beta, t).S'(F)}{S(P) \xrightarrow{(\alpha, r)} (\beta, t).S'(Q)} \quad (S(E) \stackrel{\text{def}}{=} E' \text{ and } C \vdash E = P, F = Q)$$

Figure 2: Transition relation for parameterised definitions

3.2 Parameterised definitions

Parameterised components are used to express the intervention of the new feature on the existing system. We use the dot notation to express overriding state and activities and the subsequent evolution to a new parameterised system. We use the plus notation in parameterised components to express the option of monitoring a range of activities.

$$\begin{array}{c}
\frac{}{C \vdash P = P} \quad \frac{}{C, I = P \vdash I = P} \quad \frac{C \vdash E = P \quad C \vdash L}{C \vdash E = P, L} \\
\\
\frac{C \vdash E_1 = P_1 \quad C \vdash E_2 = P_2}{C \vdash E_1 \underset{L}{\bowtie} E_2 = P_1 \underset{L}{\bowtie} P_2} \quad \frac{C \vdash E = P}{C \vdash E/L = P/L}
\end{array}$$

Figure 3: The judgement relation for definition contexts

3.3 Contexts

We introduce definition contexts to explain the binding of actual component parameters to formal parameter identifiers. Recalling the syntactic restriction on formal parameter identifiers from Section 3 on page 4 we note that contexts cannot contain re-definitions of identifiers of constants. The notation $C, I = P$ therefore denotes a context C extended by the binding of an identifier I to a component P . Such an expression can be used to judge that the identifiers I and P both denote the same component. The notation $E = P, L$ is used to denote a list of equalities beginning with one between E and P and continuing with those in L .

4 Example

We present here an example which serves only to help explain the use of the feature construct, not to act as a compelling defense of the use of formal notations in performance modelling. We consider first the very simple model of a transmitter which transmits at rate t to a receiver which receives data at rate r . This transmission is conducted through the medium of a network which passively cooperates with the transmitter and the receiver. This is the meaning of the \top symbol used in the occurrences of the transmit and receive activities specified in the description of the network, that it is passive with respect to these activities.

The feature which we add to this simple system is a new component which monitors the network bandwidth which the transmitter has been able to obtain and signals whenever this drops below a critical threshold. In this circumstance the transmitter halves its transmission rate (for example, in a multimedia application by sampling an analogue input signal at half of the previous rate or in another application by applying data compression).

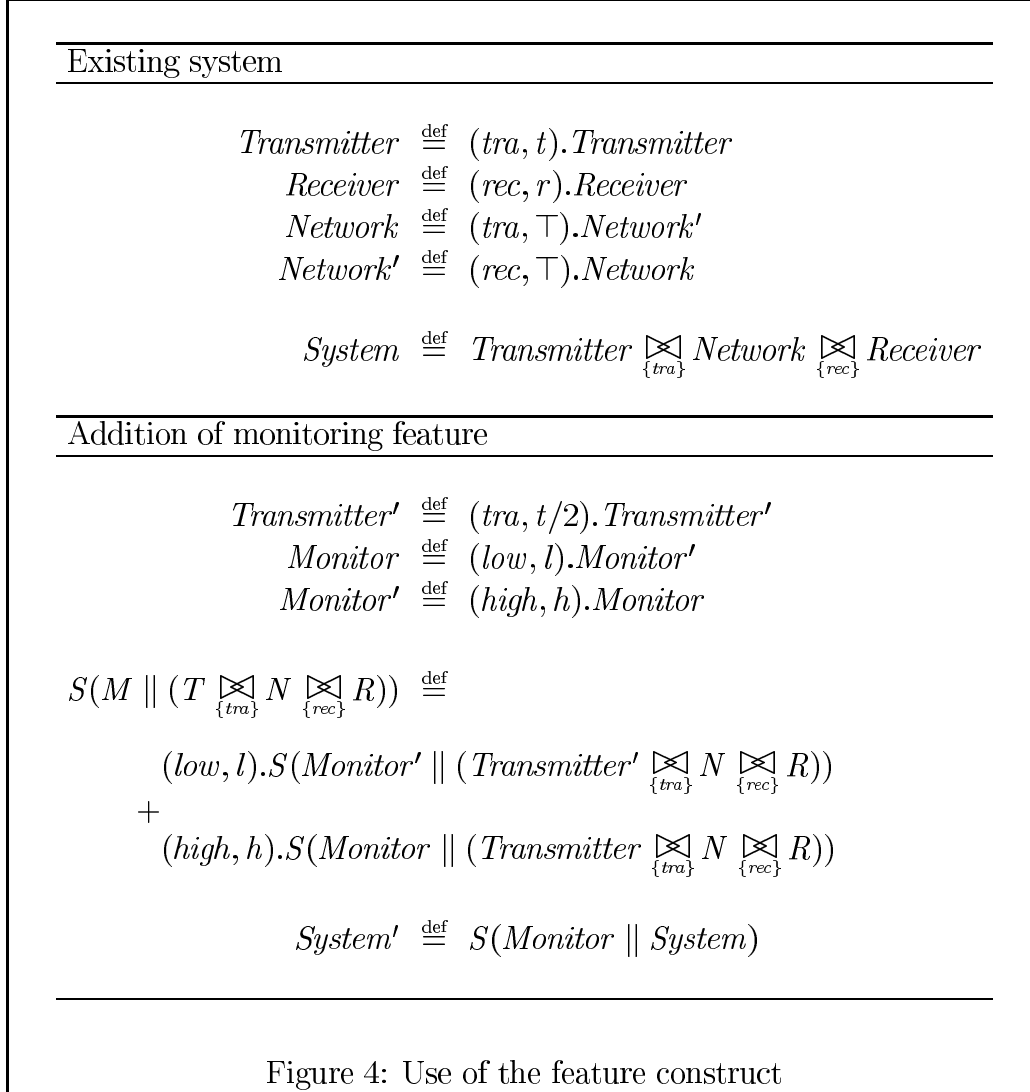


Figure 4: Use of the feature construct

It is possible that the congestion on the network will subsequently reduce. In this case the monitor will signal that it is suitable for the transmitter to resume transmitting at the higher rate. The details of how the monitor measures the consumption of the available bandwidth are abstracted away in the model and the signals to switch between transmitting at the low rate and the high rate are simply modelled by stochastic events with parameters l and h respectively. From a performance modelling point of view, the effect of adding the new feature is to turn a simple Markovian model into a Markov modulated process. The new feature ensures that the correct *Transmitter* component is used whenever the *Monitor* component changes state.

The fact that there was no planning for the subsequent addition of a monitoring feature in the original model is reflected in the fact that the activities which are of interest to the feature $S()$ are the *low* and *high* activities which are themselves newly added to the system via the *Monitor* component.

5 Future work

We have presented a method of describing additional features of a system separately from the description of the system itself. As features are progressively added, the complexity of a system inevitably grows. One avenue of future work is the investigation of the estimation of the additional complexity which is brought to the system by the addition of a single feature. Concretely, we could ask for a formula which computes how the size of the state space of the extended system will increase as a function of the state spaces of the systems which are used as actual parameters of parameterised components.

Conclusions

If performance modelling notations and tools are to realise the valuable contribution which they promise for the development of reliable and efficient complex software systems they must provide support not only for the initial design of systems but also for their correction, revision, and subsequent extension. We have shown how a new construct could be added to the PEPA performance modelling language: *parameterised components*. The new construct satisfies many of the desired goals of a feature construct and in addition promotes *structured whole-lifecycle performance modelling* of complex software systems. We have given the construct an operational semantics which builds upon the existing semantics of PEPA. We have applied the construct in a small case study.

Acknowledgements

The authors are grateful to Mark Ryan for suggestions which helped to improve the structure of the paper and the presentation of the feature construct in particular. Stephen Gilmore is supported by the ‘Distributed Commit Protocols’ grant from the EPSRC and by Esprit Working group FIREworks. Jane Hillston is supported by the EPSRC ‘COMPA’ grant.

References

- [1] J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
- [2] D.R.W. Holton. A PEPA specification of an industrial production cell. In S. Gilmore and J. Hillston, editors, *Proceedings of the Third International Workshop on Process Algebras and Performance Modelling*, pages 542–551. Special Issue of *The Computer Journal*, 38(7), December 1995.
- [3] S. Gilmore, J. Hillston, D.R.W. Holton, and M. Rettelbach. Specifications in Stochastic Process Algebra for a Robot Control Problem. *International Journal of Production Research*, 34(4):1065–1080, 1996.
- [4] A. El-Rayes, M. Kwiatkowska, and S. Minton. Analysing performance of lift systems in PEPA. In R. Pooley and J. Hillston, editors, *Proceedings of the Twelfth UK Performance Engineering Workshop*, pages 83–100, Department of Computer Science, The University of Edinburgh, September 1996.
- [5] E. W. Dempster, N. T. Tomov, J. Lü, C. S. Pua, M. H. Williams, A. Burger, H. Taylor, and P. Broughton. Verifying a performance estimator for parallel DBMSs. In *Proceedings of EuroPar (EuroPar'98)*, September 1998.
- [6] H. Bowman, J. Bryans, and J. Derrick. Analysis of a multimedia stream using stochastic process algebra. In C. Priami, editor, *Sixth International Workshop on Process Algebras and Performance Modelling*, pages 51–69, Nice, September 1998.
- [7] L. Kloul, J.M. Fourneau, and F. Valois. Performance modelling of hierarchical cellular networks using PEPA. In J. Hillston, editor, *Proceedings of the Seventh International Workshop on Process Algebras and Performance Modelling*, Zaragoza, Spain, September 1999.
- [8] M. C. Plath and M. D. Ryan. Plug and play features. In W. Bouma, editor, *Feature Interactions in Telecommunications Systems V*. IOS Press, 1998.
- [9] M. D. Ryan. Feature-oriented programming: A case study using the SMV language. Technical report, School of Computer Science, The University of Birmingham, UK, September 1997.

A Summary of the PEPA language

The PEPA language provides a small set of combinators. These allow language terms to be constructed defining the behaviour of components, via the activities they undertake and the interactions between them. The syntax may be formally introduced by means of the grammar shown in Figure 5.

$S ::=$	(sequential components)
$(\alpha, r).S$	(prefix)
$S + S$	(choice)
C_S	(constant)
 $P ::=$	 (model components)
$P \underset{L}{\bowtie} P$	(cooperation)
P/L	(hiding)
C	(constant)

Figure 5: The syntax of PEPA

In the grammar S denotes a *sequential component* and P denotes a *model component* which executes in parallel. C stands for a constant which denotes either a sequential or a model component, as defined by a defining equation. C when subscripted with an S stands for constants which denote sequential components. The component combinators, together with their names and interpretations, are presented informally below.

Prefix: The basic mechanism for describing the behaviour of a system is to give a component a designated first action using the prefix combinator, denoted by a full stop. For example, the component $(\alpha, r).S$ carries out activity (α, r) , which has action type α and an exponentially distributed duration with parameter r , and it subsequently behaves as S . Sequences of actions can be combined to build up a life cycle for a component.

Choice: The life cycle of a sequential component may be more complex than any behaviour which can be expressed using the prefix combinator alone. The choice combinator captures the possibility of competition

between different possible activities. The component $P + Q$ represents a system which may behave either as P or as Q . The activities of both P and Q are enabled. The first activity to complete distinguishes one of them: the other is discarded. The system will behave as the derivative resulting from the evolution of the chosen component.

Constant: It is convenient to be able to assign names to patterns of behaviour associated with components. Constants are components whose meaning is given by a defining equation.

Hiding: The possibility to abstract away some aspects of a component's behaviour is provided by the hiding operator, denoted by the division sign in P/L . Here, the set L of visible action types identifies those activities which are to be considered internal or private to the component. These activities are not visible to an external observer, nor are they accessible to other components for cooperation. Once an activity is hidden it only appears as the unknown type τ ; the rate of the activity, however, remains unaffected.

Cooperation: Most systems are comprised of several components which interact. In PEPA direct interaction, or *cooperation*, between components is represented by the butterfly combinator. The set which is used as the subscript to the cooperation symbol determines those activities on which the *cooperands* are forced to synchronise. Thus the cooperation combinator is in fact an indexed family of combinators, one for each possible *cooperation set* L (we write $P \parallel Q$ as an abbreviation for $P \bowtie_L Q$ when L is empty). When cooperation is not imposed, namely for action types not in L , the components proceed independently and concurrently with their enabled activities. However if a component enables an activity whose action type is in the cooperation set it will not be able to proceed with that activity until the other component also enables an activity of that type. The two components then proceed together to complete the *shared activity*. The rate of the shared activity may be altered to reflect the work carried out by both components to complete the activity.

In some cases, when an activity is known to be carried out in cooperation with another component, a component may be *passive* with respect to that activity. This means that the rate of the activity is left unspecified and is determined upon cooperation, by the rate of the activity in the other component. All passive actions must be synchronised in the final model.

Model components capture the structure of the system in terms of its *static* components. The dynamic behaviour of the system is represented by the evolution of these components, either individually or in cooperation. The form of this evolution is governed by a set of formal rules which give an operational semantics of PEPA terms. The semantic rules, in the structured operational style of Plotkin, are presented in Figure 6 without further comment; the interested reader is referred to [1] for more details. The rules are read as follows: if the transition(s) above the inference line can be inferred, then we can infer the transition below the line. The notation r_α^E which is used in the third cooperation rule denotes the apparent rate of α in E .

Thus, as in classical process algebra, the semantics of each term in PEPA is given via a labelled *multi-transition* system—the multiplicities of arcs are significant. In the transition system a state corresponds to each syntactic term of the language, or *derivative*, and an arc represents the activity which causes one derivative to evolve into another. The complete set of reachable states is termed the *derivative set* of a model and these form the nodes of the *derivation graph* formed by applying the semantic rules exhaustively.

The timing aspects of components’ behaviour are not represented in the states of the derivation graph, but on each arc as the parameter of the negative exponential distribution governing the duration of the corresponding activity. The interpretation is as follows: when enabled an activity $a = (\alpha, r)$ will delay for a period sampled from the negative exponential distribution with parameter r . If several activities are enabled concurrently, either in competition or independently, we assume that a *race condition* exists between them. Thus the activity whose delay before completion is the least will be the one to succeed. The evolution of the model will determine whether the other activities have been *aborted* or simply *interrupted* by the state change. In either case the memoryless property of the negative exponential distribution eliminates the need to record the previous execution time.

When two components carry out an activity in cooperation the rate of the shared activity will reflect the working capacity of the slower component. We assume that each component has a capacity for performing an activity type α , which cannot be enhanced by working in cooperation (it still must carry out its own work), unless the component is passive with respect to that activity type. For a component P and an action type α , this capacity is termed the *apparent rate* of α in P . It is the sum of the rates of the α type activities enabled in P . The apparent rate of α in a cooperation between P and Q over α will be the minimum of the apparent rate of α in P and the apparent rate of α in Q .

The derivation graph is the basis of the underlying Continuous Time Markov Chain (CTMC) which is used to derive performance measures from

Prefix

$$\frac{}{(\alpha, r).E \xrightarrow{(\alpha, r)} E}$$

Cooperation

$$\frac{E \xrightarrow{(\alpha, r)} E'}{E \bowtie_L F \xrightarrow{(\alpha, r)} E' \bowtie_L F} \quad (\alpha \notin L) \quad \frac{F \xrightarrow{(\alpha, r)} F'}{E \bowtie_L F \xrightarrow{(\alpha, r)} E \bowtie_L F'} \quad (\alpha \notin L)$$

$$\frac{E \xrightarrow{(\alpha, r_1)} E' \quad F \xrightarrow{(\alpha, r_2)} F'}{E \bowtie_L F \xrightarrow{(\alpha, R)} E' \bowtie_L F'} \quad (\alpha \in L) \quad \text{where } R = \frac{r_1}{r_\alpha^E} \frac{r_2}{r_\alpha^F} \min(r_\alpha^E, r_\alpha^F)$$

Choice

$$\frac{E \xrightarrow{(\alpha, r)} E'}{E + F \xrightarrow{(\alpha, r)} E'} \quad \frac{F \xrightarrow{(\alpha, r)} F'}{E + F \xrightarrow{(\alpha, r)} F'}$$

Hiding

$$\frac{E \xrightarrow{(\alpha, r)} E'}{E/L \xrightarrow{(\alpha, r)} E'/L} \quad (\alpha \notin L) \quad \frac{E \xrightarrow{(\alpha, r)} E'}{E/L \xrightarrow{(\tau, r)} E'/L} \quad (\alpha \in L)$$

Constant

$$\frac{E \xrightarrow{(\alpha, r)} E'}{A \xrightarrow{(\alpha, r)} E'} \quad (A \stackrel{\text{def}}{=} E)$$

Figure 6: The operational semantics of PEPA

a PEPA model. The graph is systematically reduced to a form where it can be treated as the state transition diagram of the underlying CTMC. Each derivative is then a state in the CTMC. The *transition rate* between two derivatives P and Q in the derivation graph is the rate at which the system changes from behaving as component P to behaving as Q . It is denoted by $q(P, Q)$ and is the sum of the activity rates labelling arcs connecting node P to node Q . In order for the CTMC to be *ergodic* its derivation graph must be strongly connected. Some necessary conditions for ergodicity, at the syntactic level of a PEPA model, have been defined [1]. These syntactic conditions are imposed by the grammar introduced earlier.

A.1 Availability of the modelling tools

The PEPA modelling tools, together with user documentation and papers and example PEPA models are available from the PEPA Web page at the address <http://www.dcs.ed.ac.uk/pepa>.