

Performance evaluation for global computation

Linda Brodo¹, Pierpaolo Degano², Stephen Gilmore³,
Jane Hillston³, and Corrado Priami⁴

- ¹ Istituto Trentino Cultura—IRST, Via Sommarive 18, 38050 Povo (Trento), Italy.
Email: brodo@itc.it
- ² Dipartimento di Informatica, Università di Pisa F.Buonarroti 2, I-56127 Pisa, Italy.
Email: degano@di.unipi.it
- ³ Laboratory for Foundations of Computer Science, The University of Edinburgh,
Edinburgh EH9 3JZ, Scotland. Email {stg, jeh}@lfc.ed.ac.uk
- ⁴ Università di Trento, Dipartimento di Informatica e Telecomunicazioni, Via
Sommarive 14, 38050 Povo (Trento), Italy. Email: priami@dit.unitn.it

Abstract. Global computing applications co-ordinate distributed computations across widely-dispersed hosts. Such systems present formidable design and implementation challenges to software developers and synchronisation, scheduling and performance problems come to the fore. Complex systems such as these can benefit from the application of high-level performance analysis methods founded on timed process algebras. In this paper we compare the use of two such approaches, the PEPA nets and EOS methods, illustrating our presentation with the example of modelling Web services.

1 Introduction

Our main concern here is comparing existing process algebraic primitives against the needs arising when modelling global applications with a view to determining their run-time performance. Communication and especially mobility are possibly the two main features characterising global computing. There are different approaches to their representation. One widely-studied approach represents mobility implicitly through the communication of links. A name n , representing a communication channel, is passed to an agent that now becomes connected through n to all the agents that know the link n . In this way the topology of the interconnecting network varies while the (distributed) computation goes on. The typical representative of this class is the well-known π -calculus [18].

An alternative approach to representing mobility and communication is taken by the PEPA nets formalism, which combines the process algebra PEPA with a Petri net infrastructure [10]. In this formalism, which can be regarded as a high-level Petri net formalism, *places* are process algebra contexts and *tokens* are process algebra components. Mobility is modelled explicitly by the firing of a transition in the Petri net which has the result of a component moving from one place to another. Communication is restricted to be local and is modelled by the usual process algebra communication between components. We have previously

studied the relationship between PEPA nets and the π -calculus, by translating a subset of the PEPA nets formalism into the stochastic π -calculus [3].

The objective of performance evaluation is to analyse the dynamic behaviour of a system and predict performance *indices* or *measures* such as throughput, utilisation or response time. If a process algebra model is to be used for this purpose certain aspects of the behaviour of the model must be quantified. For example, in classical process algebras alternative behaviours are modelled by a non-deterministic choice. However from such a model no predictions about the likelihood of differing behaviours can be made. Therefore when the objective is performance evaluation, non-deterministic choice is replaced by probabilistic choice. Similarly, in the dynamic behaviour of a system the durations of actions (or equivalently the delays between events) are important and must be incorporated into the model.

Many probabilistic and timed extensions of process algebras have appeared in the literature in the last 15 years, however the most prevalent approach taken in performance evaluation is exemplified by PEPA [12]. In this language all actions have an associated duration which is specified by a random variable, governed by a negative exponential distribution. In PEPA probabilistic choice is not modelled explicitly; when more than one activity is possible it is assumed that the activities race in the sense that each draws from the corresponding distribution function to obtain a duration for this instance of the activity. The activity with the shorter duration sample is the one which will be performed first, thus “winning the race”. In practice, since all durations are governed by a negative exponential distributions, the relative probability of activities in competition can be derived by a simple formula. The stochastic π -calculus [20] adopts similar constructs.

In this way, our process algebra models can be used to generate a Continuous Time Markov Chain (CTMC) which can be solved to obtain a steady state probability distribution from which performance measures can be derived. In recent work we have extended PEPA to allow the durations of activities to be defined via functions rather than explicitly [14, 13]. In the context of global computing this means that the duration of an activity can depend on the state of other components. In PEPA nets, again the target representation for performance analysis is a CTMC, and so both process algebra transitions, and Petri net firings have an associated duration which is negative exponentially distributed. As previously conflicts may be solved by the race policy but it is also possible to assign different priorities to different Petri net transitions, giving some firings priority over others [10].

In the EOS approach [7], the transition labels are enhanced so that they record the application of inference rules. The designer of the application under analysis can define evaluation functions that determine the rates of transitions, by inspecting enhanced labels, as these represent the low level routines performed by the run-time support to execute the transition itself. Since the rates are also affected by the target architecture, its peculiarities will also affect the evaluation functions, the parameters of which are then the enhanced labels and the architectural details [19].

Structure of this paper: The paper is organized as follows. In the next section we describe our running example of Web services. Section 3 recalls the basics of PEPA nets, while Sect. 4 introduces the PEPA nets semantics and shows how performance analysis can be carried out on the running example. Section 5 describes the EOS approach on the π -calculus and shows how it can be used to perform a quantitative analysis of the Web service system. Finally, we draw some conclusions.

2 Example: modelling Web services

Web services provide a technological platform which enables global computation. In a Web services architecture clients and services are loosely coupled and geographically distributed. Services are obtained by discovery from registries and directories. Web service descriptions specify the interfaces and locations of services. Method invocation and transport of data are performed by asynchronous message passing. The computing platform is heterogeneous and architecture-neutral. The implementation platform is also heterogeneous; web service clients may be implemented in a different implementation language from the service application. All of the above qualities typify global computation: distributed computations across heterogenous platforms utilising discovery services to effect remote evaluation.

Another typical quality of global computations is that they take place across administrative domains. In consequence they must coordinate communication and evaluation across different security contexts. Firewalls are used in distributed systems to safeguard systems against attack, preventing unrestricted communication between remote sites. Their presence is a necessity but one which causes problems for some communications protocols. Web services however are accessed by HTTP. The use of the HTTP protocol virtually eliminates the complications caused by firewalls.

Web services achieve global accessibility in practice by adherence to open standards which are widely supported and used. Both communication protocols and data formats are standardised. Web services are globally positioned by giving each a unique Uniform Resource Identifier (URI). Clients and services in a Web services architecture exchange XML-encoded messages using the standard SOAP protocol. The use of XML for data carriage provides an abstraction barrier over the language-dependent in-memory data formats used in application programs. The SOAP protocol provides a high-level transport and may itself be layered over native network protocols such as SMTP or HTTP. A special-purpose language WSDL (Web Services Description Language) exists for describing interface “contracts” between Web service provider and client.

Web services applications incorporate many significant practical advances over previous generations of distributed systems technology. One cost of their considerable advantages is that they are resource-intensive systems. Web services include many layers of encapsulation which would not be needed in traditional binary communication protocols. Service lookup is an overhead, as is

XML-encoding. The XML language itself is a verbose, human-readable encoding format which is engineered for clarity, not for compactness. This has the consequence that XML-encoded method calls are weighty data items which incur significant transmission costs. The use of the HTTP protocol is another overhead. Network reliability, host availability problems and distributed system faults further degrade performance. For these reasons, Web services provide a highly appropriate example for performance modelling techniques such as those presented in this paper.

Process algebras are excellent tools for modelling Web services because they naturally support peer-to-peer architectures. The co-operator/co-operand style of process algebras allows an intuitive encoding of control flow logics such as callbacks. A process algebra which provides direct support for location-awareness is an added benefit. This provides the right conceptual modelling concepts to represent mobile code systems ranging from the asynchronous remote procedure call method provided by Web services to more complex configurations as embodied in the remote evaluation, code-on-demand or mobile agent paradigms.

3 PEPA nets

PEPA nets extend the PEPA [12] stochastic process algebra by connecting individual PEPA models together as the places of a coloured stochastic Petri net. PEPA components travel from place to place as the tokens of the net.

A PEPA net differentiates between two types of change of state. We refer to these as *firings* of the net and *transitions* of PEPA components. Each are special cases of PEPA activities. Transitions of PEPA components will typically be used to model small-scale changes of state as components undertake activities. Firings of the net will typically be used to model large-scale changes of state such as context switches, breakdowns and repairs, one thread yielding to another, or a mobile software agent moving from one network host to another.

A firing in a PEPA net causes the transfer of one token from one place to another. The token which is moved is a PEPA component, which causes a change in the subsequent evaluation both in the source (where existing cooperations with other components now can no longer take place) and in the target (where previously disabled cooperations are now enabled by the arrival of an incoming component which can participate in these interactions). Firings have global effect because they involve components at more than one place in the net.

A transition in a PEPA net takes place whenever a transition of a PEPA component can occur (either individually, or in cooperation with another component). Components can only cooperate if they are resident in the same place in the net. The PEPA net formalism does not allow components at different places in the net to cooperate on a shared activity. An analogy is with message-passing distributed systems without shared-memory where software components on the same host can exchange information without incurring a communication overhead but software components on different hosts cannot. Additionally we do not allow a firing to coincide with a transition which is shared, i.e. it is not possible

for two components in one place to cooperate *and* transfer to another place as an atomic action. Thus transitions in a PEPA net have local effect because they involve only components at one place in the net. Maintaining this strict distinction between firings and transitions is essential in order to provide the separation into macro- and micro-step state changes that we are seeking to represent.

Each place has a distinct alphabet for transitions and firings, meaning that the same action type cannot be used for both. Thus there can be no ambiguity between such micro- and macro-scale transitions.

A PEPA net is made up of PEPA *contexts*, one at each place in the net. A context consists of a number of *static* components (possibly zero) and a number of *cells* (at least one). Like a memory location in an imperative program, a cell is a storage area to be filled by a datum of a particular type. In particular in a PEPA net, a cell is a storage area dedicated to storing a PEPA component. The components which fill cells can circulate as the tokens of the net. In contrast, the static components cannot move.

We use the notation $Q[-]$ to denote a context which could be filled by the PEPA component Q or one with the same alphabet. If Q has derivatives Q' and Q'' only and no other component has the same alphabet as Q then there are four possible values for such a context: $Q[-]$, $Q[Q]$, $Q[Q']$ and $Q[Q'']$. $Q[-]$ enables no transitions. $Q[Q]$ enables the same transitions as Q . $Q[Q']$ enables the same transitions as Q' . $Q[Q'']$ enables the same transitions as Q'' . As usual with PEPA components we require that the component has an ergodic definition so that it is always possible to return to a state which one has previously reached. This has as a consequence that if Q' is a derivative of Q then it is also the case that Q is a derivative of Q' , for any Q and Q' .

The introduction of contexts requires an extension to the syntax of PEPA. This extension is presented in Table 1.

For any token component its action type set can be partitioned in distinct subsets corresponding to transitions and firings respectively. For a component Q we will denote these sets by $\mathcal{A}_t(Q)$ and $\mathcal{A}_f(Q)$, where $\mathcal{A}_t(Q)$ is the set of local transitions currently enabled in Q and $\mathcal{A}_f(Q)$ is the set of firings currently enabled for Q . Note that for a firing to be enabled the token must enable the corresponding activity, it must be in a place connected to a net-level transition of the same type and there must be an empty cell at the output place of the transition of the correct token type.

We use capitalised names to denote PEPA components (such as P and Q) and lowercase for PEPA transitions (such as a and b). We use bold capitalised names for PEPA net places (such as \mathbf{P}_1 and \mathbf{P}_2) and bold lowercase for PEPA net firings (such as \mathbf{a} and \mathbf{b}).

3.1 Markings in a PEPA net

The *marking* of a classical Petri net records the number of tokens which are resident at each place in the net. Since the tokens of a classical Petri net are indistinguishable it is sufficient to record their number and one could present the marking of a Petri net with places P_1 , P_2 and P_3 as $(P_1 : 2, P_2 : 1, P_3 : 0)$. If

$N ::= D^+ M$	(net)
(definitions and marking)	
$M ::= (M_{\mathbf{P}}, \dots)$	(marking)
$M_{\mathbf{P}} ::= \mathbf{P}[C, \dots]$	(place marking)
(marking vectors)	
$D ::= I \stackrel{\text{def}}{=} S$	(component defn)
$\mathbf{P}[C] \stackrel{\text{def}}{=} P[C]$	(place defn)
$\mathbf{P}[C, \dots] \stackrel{\text{def}}{=} P[C] \underset{L}{\bowtie} P$	(place defn)
	(identifier declarations)
$S ::= (\alpha, r).S$	(prefix)
$S + S$	(choice)
I	(identifier)
(sequential components)	
$P ::= P \underset{L}{\bowtie} P$	(cooperation)
P/L	(hiding)
$P[C]$	(cell)
I	(identifier)
(concurrent components)	
$C ::= \cdot$	(empty)
S	(full)
	(cell term expressions)

Table 1. The syntax of PEPA extended with contexts

an ordering is imposed on the places of the net a more compact representation of the marking can be used. Place names are omitted and the marking can be written using vector notation thus, $(2, 1, 0)$.

For a PEPA net, we can denote a marking by $(\mathbf{P}_1[Q], \mathbf{P}_2[-], \mathbf{P}_3[-])$ (the token at place \mathbf{P}_1 is in state Q ; the other places have no tokens). In general, a context may have more than one parameter, to be filled by PEPA components of different types. We denote the i th component of a marking M by M_i . For example, $(\mathbf{P}_1[Q], \mathbf{P}_2[-], \mathbf{P}_3[-])_1$ is $\mathbf{P}_1[Q]$.

It is simple to define a function to count the number of tokens in a PEPA net term and this function proves to be useful in practice.

$$\begin{aligned}
\text{tokens}(P) &= 0 \\
\text{tokens}(P[-]) &= 0 \\
\text{tokens}(P[P']) &= 1 \\
\text{tokens}(P \underset{L}{\bowtie} Q) &= \text{tokens}(P) + \text{tokens}(Q) \\
\text{tokens}(P/L) &= \text{tokens}(P)
\end{aligned}$$

3.2 Net-level transitions in a PEPA net

Transitions at the net-level of a PEPA net are labelled in a similar way to the labelled multi-transition system which records the unfolding of the state space of a PEPA model. A labelling function ℓ maps transition names into pairs of names such as (α, r) where it is possible that $\ell(t_i) = \ell(t_j)$ but $t_i \neq t_j$. The first

element of a pair (α, r) specifies an *activity* which must be performed in order for a component to move from the input place of the transition to the output place. The activity type records formally the activity which must be performed if the transition is to fire. The second element is an exponentially-distributed random variable which quantifies the *rate* at which the activity can progress in conjunction with the component which is performing it.

As an example, suppose that Q is a component which is currently at place \mathbf{P}_1 and that it can perform an activity α with rate r_1 to produce the derivative Q' . Further, say that the net has a transition between \mathbf{P}_1 and \mathbf{P}_2 labelled by (α, r_2) . If Q performs activity α in this setting it will be removed from \mathbf{P}_1 (leaving behind an empty cell) and Q' will be deposited into \mathbf{P}_2 (filling an empty cell there).

3.3 Net structure of a PEPA net

The class of nets that we currently use for modelling the net structure of a PEPA net is restricted to *structural state machines*, i.e. nets whose transitions can have only one input place and one output place. This means that we can represent conflicts at the net level, while synchronisations are not allowed. This is consistent with the fact that PEPA components cannot cooperate on a shared activity when they are resident in different places.

It is usual with coloured Petri nets to associate functions with arcs, offering a generalisation of the usual, basic “functions” offered by arc multiplicities. In PEPA nets the arc functions are implicit. The modification of a token which takes place when it is fired is wholly specified by the action type of the firing, the definition of the token and the semantics. Furthermore, although we allow multiple tokens within net places, only one token can move at each firing. Thus arc multiplicities greater than one are not allowed.

4 Semantics

The PEPA language is formally defined by a small-step operational semantics. In order to describe the firing rule for PEPA nets formally we need a relational operator which is to be used to express the fact that there exists a particular transition in the net superstructure. This operator must have the properties that it identifies the source and target of the transition and that it records the activity which is to be performed in order for a component to cross this transition, moving from the source to the target. We use the notation

$$\mathbf{P}_1 \xrightarrow{(\alpha, r)} \mathbf{P}_2$$

to capture the information that there is a transition connecting place \mathbf{P}_1 to place \mathbf{P}_2 labelled by (α, r) . This relation captures static information about the structure of the net, not dynamic information about its behaviour. We could describe the net structure in a PEPA net using a list of such declarations but the more familiar graphical presentation of a net presents the same information in a more accessible way.

Definition 1. A PEPA net \mathcal{N} is a tuple $\mathcal{N} = (\mathcal{P}, \mathcal{T}, I, O, \ell, \pi, \mathcal{C}, D, M_0)$ such that

- \mathcal{P} is a finite set of places;
- \mathcal{T} is a finite set of net transitions;
- $I : \mathcal{T} \rightarrow \mathcal{P}$ is the input function;
- $O : \mathcal{T} \rightarrow \mathcal{P}$ is the output function;
- $\ell : \mathcal{T} \rightarrow (\mathcal{A}_f, \mathbb{R}^+ \cup \{\top\})$ is the labelling function, which assigns a PEPA activity ((type, rate) pair) to each transition. The rate determines the negative exponential distribution governing the delay associated with the transition;
- $\pi : \mathcal{A}_f \rightarrow \mathbb{N}$ is the priority function which assigns priorities (represented by natural numbers) to firing action types;
- $\mathcal{C} : \mathcal{P} \rightarrow P$ is the place definition function which assigns a PEPA context, containing at least one cell, to each place;
- D is the set of token component definitions;
- M_0 is the initial marking of the net.

The semantic rules for PEPA nets are provided in Table 2. The Cell rule conservatively extends the PEPA semantics to define that a cell which is filled by a component Q has the same transitions as Q itself. A healthiness condition on the rule (also called a *typing judgement*) requires a context such as $Q[-]$ to be filled with a component which has the same alphabet as Q . We write $Q =_a Q'$ to state that Q and Q' have the same alphabet. There are no rules to infer transitions for an empty cell because an empty cell enables no transitions.

The Transition rule states that the net has local transitions which change only a single component in the marking vector. This rule also states that these transitions agree with the transitions which are generated by the PEPA semantics (including the extension for contexts). Recall that the transition and firing alphabets of any place must be distinct.

The Firing rule takes one marking of the net to another marking by performing a PEPA activity and moving a PEPA component from the input place to the output place. This has the effect that two entries in the marking vector change simultaneously.

4.1 The net bisimulation relation

In this section we define a bisimulation relation for PEPA nets called *net bisimulation*. This relation is important both in theory and in practice. In the evolution of the state space of a model by our tool we only store states up to net bisimulation, i.e. we carry out automatic aggregation over equivalent states. This provides a dramatic reduction in the state space of the model under certain conditions.

Our relation is defined in the style of Larsen and Skou [16], based on a conditional transition rate between *markings*, rather than the strong equivalence relation of PEPA which considers the transition rates between components. The *conditional transition rate* from marking M to marking M' via action type α , denoted $q(M, M', \alpha)$, is the sum of the activity rates labelling arcs connecting the

Cell:

$$\frac{Q' \xrightarrow{(\alpha, r)} Q''}{Q[Q'] \xrightarrow{(\alpha, r)} Q[Q'']} \quad (Q =_a Q')$$

Transition:

$$\frac{M_{\mathbf{P}} \xrightarrow{(\alpha, r)} M'_{\mathbf{P}}}{(\dots, M_{\mathbf{P}}, \dots) \xrightarrow{(\alpha, r)} (\dots, M'_{\mathbf{P}}, \dots)} \quad (\alpha \in \mathcal{A}_t)$$

Enabling:

$$\frac{Q \xrightarrow{(\alpha, r_1)} Q' \quad \mathbf{P}_i \xrightarrow{(\alpha, r_2)} \mathbf{P}_j}{(\dots, \mathbf{P}_i[\dots, Q, \dots], \dots, \mathbf{P}_j[\dots, -, \dots], \dots) \xrightarrow{(\alpha, R)}_{\pi(\alpha)} (\dots, \mathbf{P}_i[\dots, -, \dots], \dots, \mathbf{P}_j[\dots, Q', \dots], \dots)} \quad (\alpha \in \mathcal{A}_f)$$

Firing:

$$\frac{M \xrightarrow{(\alpha, r)}_n M' \quad M \xrightarrow{(\beta, s)}_m M''}{M \xrightarrow{(\alpha, r)} M'} \quad (n \geq m)$$

Table 2. Additional semantic rules for PEPA nets

corresponding nodes in the derivation graph which are labelled by the action type α . The *total conditional transition rate* from a marking M to a set of markings E is defined as

$$q[M, E, \alpha] = \sum_{M' \in E} q(M, M', \alpha)$$

Definition 2. An equivalence relation over markings, $\mathcal{R} \subseteq M \times M$, is a net bisimulation if whenever $(M, M') \in \mathcal{R}$ then for all $\alpha \in \mathcal{A}$ and for all equivalence classes $E \in M/\mathcal{R}$,

$$q[M, E, \alpha] = q[M', E, \alpha]$$

4.2 PEPA net model of a Web service

In modelling our Web services example as a PEPA net we first identify three components: *Client*, *WebService* and *SOAPmessage*. We begin with the simplest of these, the *SOAPmessage*. The lifecycle of this component is that it is built using a message composition API, then launched over the network and then read using an XML parser. This leads to another message which is the continuation of the lifetime of this component. This component plays the role of passive data in our application so in its description it leaves unspecified (\top) the rates at which these actions are performed, allowing the cooperating partner in the synchronisation to determine these rates.

$$\begin{aligned} \text{SOAPmessage} &\stackrel{\text{def}}{=} (\text{compose_message}, \top). \\ &\quad (\mathbf{launch}, \top). \\ &\quad (\text{read_message}, \top).\text{SOAPmessage} \end{aligned}$$

A *Client* divides its time between local computation, the details of which we do not model here, and Web services interactions. When the client comes to a phase in its local computation where it realises that it needs to use a Web service it interacts with the discovery service to obtain a specification of the service. It then composes a SOAP message to send to the service. The communication with the remote Web service is asynchronous so the client returns to its local computation, anticipating that a reply will come later. When a message is returned from the service the client will read it and make use of the results in the remainder of its computation.

$$\begin{aligned}
Client &\stackrel{def}{=} (local_computing, r_l).Client \\
&\quad + (discover, r_d).Client_1 \\
Client_1 &\stackrel{def}{=} (compose_message, r_c^C).Client_2 \\
Client_2 &\stackrel{def}{=} (local_computing, r_l).Client_2 \\
&\quad + (read_message, r_r^C).Client
\end{aligned}$$

The lifetime of a Web service is modelled as a simple loop. Web services requests are received and read; these lead to the execution of a Web service and the composition of a message to return the results.

$$\begin{aligned}
WebService &\stackrel{def}{=} (read_message, r_r^S).WebService_2 \\
WebService_2 &\stackrel{def}{=} (transact_service, r_s).WebService_3 \\
WebService_3 &\stackrel{def}{=} (compose_message, r_c^S).WebService
\end{aligned}$$

The places of the net specify that there is a cell (a storage place) for a SOAP message at the client side and at the Web service side. The message synchronises on composition and reading activities.

$$\begin{aligned}
P_1[s] &\stackrel{def}{=} SOAPmessage[s] \underset{L}{\bowtie} WebService \\
P_2[s] &\stackrel{def}{=} SOAPmessage[s] \underset{L}{\bowtie} Client \\
&\text{where } L = \{ compose_message, read_message \}
\end{aligned}$$

The initial marking of the net places a token on the client side, in its initial state: $(P_1[-], P_2[SOAPmessage])$.

Firing the operational semantics of the example generates the state space depicted in Fig. 1 with the transition system given in Fig. 2. By erasing activity names from the labelled transition system we obtain the CTMC given in Fig. 3.

As a concrete illustration of numerical evaluation take $r_d = r_m = 17.03$, $r_c^C = r_r^S = r_c^S = r_r^C = 3.28$ and $r_s = 1.10$. The value of r_l is immaterial because the self-loops on states which are visible at the process algebra level are not represented at the Markov chain level. In the Markov chain representation we are concerned with balancing flow into a state against flow out of a state, so self-loops have no role.

Denote the infinitesimal generator matrix of the CTMC in Fig 3 by \mathbf{Q} . As usual, we solve $\boldsymbol{\pi}\mathbf{Q} = \mathbf{0}$ subject to $\sum \boldsymbol{\pi} = 1$ giving $(0.025, 0.132, 0.025, 0.132, 0.394, 0.132, 0.025, 0.132)$.

1	$(SOAPmessage[-] \xrightarrow{L} WebService,$ $SOAPmessage[SOAPmessage] \xrightarrow{L} Client)$
2	$(SOAPmessage[-] \xrightarrow{L} WebService,$ $SOAPmessage[SOAPmessage] \xrightarrow{L} Client_1)$
3	$(SOAPmessage[-] \xrightarrow{L} WebService,$ $SOAPmessage[(\mathbf{launch}, \top).(read_message, \top).SOAPmessage] \xrightarrow{L} Client_2)$
4	$(SOAPmessage[(read_message, \top).SOAPmessage] \xrightarrow{L} WebService,$ $SOAPmessage[-] \xrightarrow{L} Client_2)$
5	$(SOAPmessage[SOAPmessage] \xrightarrow{L} WebService_2,$ $SOAPmessage[-] \xrightarrow{L} Client_2)$
6	$(SOAPmessage[SOAPmessage] \xrightarrow{L} WebService_3,$ $SOAPmessage[-] \xrightarrow{L} Client_2)$
7	$(SOAPmessage[(\mathbf{launch}, \top).(read_message, \top).SOAPmessage] \xrightarrow{L} WebService,$ $SOAPmessage[-] \xrightarrow{L} Client_2)$
8	$(SOAPmessage[-] \xrightarrow{L} WebService,$ $SOAPmessage[(read_message, \top).SOAPmessage] \xrightarrow{L} Client_2)$

Fig. 1. Reachable state space of the PEPA nets Web services model shown as the markings of (P_1, P_2)

1	$-(local_computing, r_l) \rightarrow$	1	5	$-(transact_service, r_s) \rightarrow$	6
1	$-(discover, r_d) \rightarrow$	2	5	$-(local_computing, r_l) \rightarrow$	5
2	$-(compose_message, r_c^C) \rightarrow$	3	6	$-(compose_message, r_c^S) \rightarrow$	7
3	$-(local_computing, r_l) \rightarrow$	3	6	$-(local_computing, r_l) \rightarrow$	6
3	$-(\mathbf{launch}, r_m) \rightarrow$	4	7	$-(\mathbf{launch}, r_m) \rightarrow$	8
4	$-(read_message, r_r^S) \rightarrow$	5	7	$-(local_computing, r_l) \rightarrow$	7
4	$-(local_computing, r_l) \rightarrow$	4	8	$-(local_computing, r_l) \rightarrow$	8
			8	$-(read_message, r_r^C) \rightarrow$	1

Fig. 2. The transition system of the Web services example




Fig. 3. The CTMC of the Web services example

4.3 Using logic to specify performance measures

We now explain how to specify performance measures of interest with respect to a PEPA net model by using a probabilistic modal logic. The appropriate logic for PEPA nets is one which can specify performance measures over the places of the net, and has the capability of expressing requirements on tokens in addition to requirements on the transitions and firings of the net.

We introduce the PML_ν logic by means of a two-level grammar which separates the specification of place formulae and token formulae from the specification of transition and firing activities. Behaviour at the transition and firing level is captured by formulae of a sub-logic, PML_μ .

This separation of PML_ν formulae from PML_μ formulae enforces a syntactic restriction on the allowable terms in the logic whereby places cannot refer to the local state at another place. This reflects the global computing idiom that it is impossible to know the global state of the system. This restriction also strongly supports the PEPA nets modelling rule which forbids communication between components at different places in the net, as in distributed systems without shared memory.

We present the sub-logic PML_μ first. The constant true is represented by \mathbf{tt} . Conjunction and negation are denoted as usual. The term ∇_α represents the inability of a process to perform an α action. The diamond operator specifies an activity α , a rate μ , and a succeeding formula which is to be satisfied by all one-step α -derivatives. The accumulated rate of these α activities must be at least μ . We use $\phi, \phi_1, \phi_2, \dots$, to range over PML_μ formulae.

$$\begin{aligned} \phi ::= & \mathbf{tt} \\ & | \neg\phi \\ & | \phi_1 \wedge \phi_2 \\ & | \nabla_\alpha \\ & | \langle \alpha \rangle_\rho \phi \end{aligned}$$

The meaning of the PML_μ connectives is given by reference to the transition relation of the PEPA net semantics. We require an additional simple auxilliary definition:

Definition 1 *Let S be a set of states. $P \xRightarrow{(\alpha, \lambda)} S$ if for all successors $P' \in S$, $P \xrightarrow{\alpha} P'$, and $\sum\{r : P \xrightarrow{(\alpha, r)} P', P' \in S\} = \lambda$.*

Now let P be a model of a PEPA net process. Then

$$\begin{aligned} P & \models_\mu \mathbf{tt} \\ P & \models_\mu \neg\phi \quad \text{iff } P \not\models_\mu \phi \\ P & \models_\mu \phi_1 \wedge \phi_2 \quad \text{iff } P \models_\mu \phi_1 \wedge P \models_\mu \phi_2 \\ P & \models_\mu \nabla_\alpha \quad \text{iff } P \not\xrightarrow{\alpha} \\ P & \models_\mu \langle \alpha \rangle_\rho \phi \quad \text{iff } P \xRightarrow{(\alpha, \lambda)} S \text{ for some } \lambda \geq \rho, \text{ and for all } P' \in S, P' \models_\mu \phi. \end{aligned}$$

It is convenient to introduce a number of derived operators. These add no expressive power to the logic but they shorten the statement of realistic performance measures in PML_μ .

$$\begin{aligned} \mathbf{ff} &\stackrel{\text{def}}{=} \neg \mathbf{tt} \\ [\alpha]_\rho \phi &\stackrel{\text{def}}{=} \neg \langle \alpha \rangle_\rho \neg \phi \\ \Delta_\alpha &\stackrel{\text{def}}{=} \neg \nabla_\alpha \\ \phi_1 \vee \phi_2 &\stackrel{\text{def}}{=} \neg((\neg \phi_1) \wedge (\neg \phi_2)) \end{aligned}$$

The PML_ν logic has as atomic propositions all of the formulae of PML_μ . In addition it has conjunction and negation, place formulae and token formulae. We use $\psi, \psi_1, \psi_2, \dots$, to range over PML_ν formulae.

$$\begin{aligned} \psi ::= & \phi \\ & | \neg \psi \\ & | \psi_1 \wedge \psi_2 \\ & | P_i[\phi] \\ & | \#P_i \sim n \end{aligned}$$

where $\sim = \{=, \neq, <, \leq, >, \geq\}$.

The meaning of PML_ν formulae (\models_ν) is defined in terms of the meaning of PML_μ formulae (\models_μ) and the token counting function for PEPA nets. Let M be a marking of a PEPA net. Then,

$$\begin{aligned} M \models_\nu \phi & \quad \text{iff } M \models_\mu \phi \\ M \models_\nu \neg \psi & \quad \text{iff } M \not\models_\nu \psi \\ M \models_\nu \psi_1 \wedge \psi_2 & \quad \text{iff } M \models_\nu \psi_1 \wedge M \models_\nu \psi_2 \\ M \models_\nu P_i[\phi] & \quad \text{iff } M_i \models_\mu \phi \\ M \models_\nu \#P_i \sim n & \quad \text{iff } \text{tokens}(M_i) \sim n. \end{aligned}$$

4.4 Selecting states of the web services model

Performance measures characterising the long-run behaviour of the system are calculated from the computation of the probability of being in selected subsets of the states of the system.

We now use PML_ν to characterise some of the states of the Web services PEPA net model, illustrating its use as a specification language for performance measures.

The first value which we might wish to quantify is the *next-read probability*. This is the probability that one of the components of the model can read a message as its next action. We related this formula to the concrete subset of states of the Web services model as shown below:

$$\|\Delta_{\text{read_message}}\| = \{4, 8\}$$

A slightly more specialised quantity is the *server next-read probability*. This is the probability that the Web service component can read a message as its

next action. Again we relate a PML_ν formula to a subset of the state space, in this case this just turns out to be just a single state.

$$\|P_1[\Delta_{read_message}]\| = \{4\}$$

As a final example we can specify the *blocking probability*. This describes the cases where a Web services request message is being processed at the server side and the client is delayed awaiting the reply, performing local computation only. For the present simple example, there are many ways to express this property some of which would also be applicable in a more complex, multi-threaded version of the model. The most direct expressions seem to come from stating the number of tokens at one of the places.

$$\|\#P_1 = 1\| = \|\#P_2 = 0\| = \{4, 5, 6, 7\}$$

5 Enhanced operational semantics

In this section we introduce enhanced operational semantics (EOS for short) [6]. EOS is built upon operational semantics by enriching labels of transitions with the (partial) encodings of their proofs. As a test-bed we use here the π -calculus [18, 17].

Definition 3. *Let \mathcal{N} be a countable infinite set of names which is ranged over by a, b, \dots, x, y, \dots with $\tau \notin \mathcal{N}$. We also assume a set \mathcal{A} of agent identifiers ranged over by A, A_1, \dots . Processes (denoted by $P, Q, R, \dots \in \mathcal{P}$) are built from names according to the syntax*

$$P ::= \mathbf{0} \mid \pi.P \mid (\nu x)P \mid [x = y]P \mid P|P \mid P + P \mid A(y_1, \dots, y_n)$$

where π may be either $x(y)$ for input, or $\bar{x}y$ for output (where x is the subject, singled out by a function sbj and y the object, singled out by a function obj) or τ for silent moves. The order of precedence among the operators is the order (from left to right) listed above. Hereafter, the trailing $\mathbf{0}$ will be omitted.

We briefly recall the intuitive semantics. The process null $\mathbf{0}$ can perform no actions. The prefix π is the first atomic action that the process $\pi.P$ can perform. The input $x(y)$ binds the occurrences of the variable y in the prefixed process P . Roughly, a name will be received on the channel x and it will substitute the free occurrences of the placeholder y in P . The output prefix $\bar{x}z$ sends the name z along the channel x without binding z . In $(\nu x)P$, the restriction operator creates a new (unique) name x whose scope is P . In $[x = y]P$, the matching operator tests if the names x and y are equal. If so, the process P is executed, otherwise the execution of the process stops. The operator $|$ defines the parallel composition of processes. In the composition $P_1 | P_2$ the two processes act independently and they may communicate if they share a common channel name. The summation operator defines the non deterministic choice: $P_1 + P_2$ behaves either as P_1 or as P_2 . For each agent identifier A there is a unique defining equation of the form

$A(\tilde{y}) \stackrel{\text{def}}{=} P$, where \tilde{y} is a list of distinct parameters which are the free names of the process P . Each occurrence of an agent identifier $A(\tilde{z})$ will be replaced by the process P , substituting the list of formal parameters \tilde{y} by the list of actual parameters \tilde{z} .

To enrich the labels of transitions we introduce tags used to record the rules applied in their derivation and we call the new label *proof terms*.

Definition 4 (proof terms). Let $\vartheta \in \{\|_0, \|_1, +_0, +_1\}^*$. Then the set Θ of proof terms (with metavariable θ) is defined by the following syntax

$$\theta ::= \vartheta\mu \mid \vartheta(\|_0\vartheta_0\mu_0, \|_1\vartheta_1\mu_1)$$

with $\mu_0 = x(z)$ iff μ_1 is either $\bar{x}z$ or $\bar{x}(z)$, or vice versa.

Function $\ell : \Theta \rightarrow \text{Act}$ is defined as


$$\ell(\theta\mu) = \mu; \quad \ell(\vartheta(\|_0\vartheta_0\mu_0, \|_1\vartheta_1\mu_1)) = \tau.$$

We only consider tags to record the occurrences of the parallel and summation operators. Instead, in [21] there are tags also for the other π -calculus operators: restriction (νx) , constant definition (\tilde{y}) , and matching $=_m$. The enhanced operational semantics is defined by the operational rules in Tab. 4 and by the minimal congruence, induced by α -congruence and by the rules in Tab. 3. Note that the $|$ and $+$ operators are no longer commutative and associative.

$$\begin{array}{ll} (\nu x)(\nu x')T \equiv (\nu x')(\nu x)P & A(\tilde{y}) \equiv P, \text{ if } A(\tilde{y}) \stackrel{\text{def}}{=} P \\ (\nu x)(T_0|T_1) \equiv ((\nu x)T_0)|T_1, \text{ if } x \notin \text{fn}(T_1) & (\nu x)\mathbf{0} \equiv \mathbf{0} \end{array}$$

Table 3. Structural congruence for the π -calculus.

An interpretation of the proof terms As already noted, the structural congruence in Tab. 3 allows the parallel operator to be neither commutative nor associative. Moreover, we also discarded the congruence rule $P|\mathbf{0} \equiv P$. These conditions allow us to interpret the sequence of parallel tags in the proof terms as abstract addresses of each subprocess. For example, consider the process $((P_1 | P_2) + Q) | ((R_1 + R_2) | T)$. Its syntax tree is as follows:



$Act : \mu.P \xrightarrow{\mu} P$	$Ide : \frac{P\{\tilde{y}/\tilde{x}\} \xrightarrow{\theta} P'}{Q(\tilde{y}) \xrightarrow{\theta} P'}, Q(\tilde{x}) \stackrel{def}{=} P$
$Sum_0 : \frac{P \xrightarrow{\theta} P'}{P + Q \xrightarrow{+0\theta} P'}$	$Par_0 : \frac{P \xrightarrow{\theta} P'}{P Q \xrightarrow{\parallel_0\theta} P' Q}, bn(\mu) \cap fn(Q) = \emptyset$
$Sum_1 : \frac{Q \xrightarrow{\theta} Q'}{P + Q \xrightarrow{+1\theta} Q'}$	$Par_1 : \frac{Q \xrightarrow{\theta} Q'}{P Q \xrightarrow{\parallel_1\theta} P' Q'}, bn(\mu) \cap fn(Q) = \emptyset$
$Com_0 : \frac{P \xrightarrow{\bar{x}y} P', Q \xrightarrow{x(z)} Q'}{P Q \xrightarrow{\langle \parallel_0\bar{x}y, \parallel_1x(z) \rangle} P' Q'\{y/z\}}$	$Open : \frac{P \xrightarrow{\bar{x}y} P'}{(\nu y)P \xrightarrow{\bar{x}(y)} P'}, y \neq x$
$Com_1 : \frac{Q \xrightarrow{x(z)} Q', P \xrightarrow{\bar{x}y} P'}{Q P \xrightarrow{\langle \parallel_0x(z), \parallel_1\bar{x}y \rangle} Q'\{y/z\} P'}$	$Res : \frac{P \xrightarrow{\theta} P'}{(\nu x)P \xrightarrow{\theta} (\nu x)P'}, x \notin n(\theta)$
$Close_0 : \frac{P \xrightarrow{\bar{x}(y)} P', Q \xrightarrow{x(y)} Q'}{P Q \xrightarrow{\langle \parallel_0\bar{x}(y), \parallel_1x(y) \rangle} (\nu y)(P' Q')}$	
$Close_1 : \frac{Q \xrightarrow{\bar{x}(y)} Q', P \xrightarrow{x(y)} P'}{Q P \xrightarrow{\langle \parallel_0x(y), \parallel_1\bar{x}(y) \rangle} (\nu y)(Q' P')}$	

Table 4. Late transition system of π -calculus.

Each subprocess is uniquely identified by the sequence of parallel tags, *e.g.* the process P_2 has $\|_0\|_1$ as abstract address. We introduce the function *add* which extracts the sequence of parallel tags (abstract address) from the proof terms, *e.g.* $add(\|_0 +_0 \|_1) = \|_0\|_1$.

5.1 Stochastic semantics

We interpret the proof terms by means of the special function $\$$ (which will be defined later on and will be called the *cost function*). This function associates a *rate* with each proof term. An aleatory variable X_θ , which expresses the time duration of the θ action, is associated with each proof term θ . Moreover, the values that X_θ can assume are regulated by an exponential function $f_\theta(x) = \lambda e^{-\lambda x}$ for such a parameter λ . We will use the property of exponential distributions that the mean of an exponential variable $E[X_\theta]$ is the inverse of its distribution parameter: $\lambda = \frac{1}{E[X_\theta]}$. We will calculate a reasonable empirical mean time duration for each action and its inverse will be assumed to be the rate associated with the relative proof term. Also the probability distribution functions enjoy the *memoryless property*. Roughly, any time an action becomes enabled, it restarts its elapsing time as it would the first time that it is enabled.

Cost Functions A performance prediction may be useful at specification time both when the future system implementation is known and when the specification is not connected to any particular implementation. In practice, in the first case performance results enrich the specification by giving a taste of potential technical improvements of the implementation of the real system. In the second case many implementations can be imagined for the same specification and this could give hints in choosing the more adequate implementation. In our view there are two logical phases: the first one requires that the system functionalities that must be satisfied are described using a specification language; the second phase some high-level quantitative values are associated with the processes or actions of the specification. These two tasks can be executed at the same time or separately. To ensure the generality of our approach, we assume to know only two common measures of distributed systems: the throughput¹ of the communication channels and the size of the messages exchanged.

Intuitively, the rate r to be associated with an action π depends on the duration of the execution of the action itself and on the time spent in performing some run time support operations of the target machine. The proof terms record these low-level operations of the run time support. Thus, for example, an action π fired after a choice costs more than the same action occurring deterministically. We model this performance degradation by introducing a scaling factor for r in correspondence with any operation of the routine implementing the action, as we will discuss later.

¹ The number of bits, characters, or blocks passing through a data communication channel. Throughput may vary greatly from its theoretical maximum. Throughput is expressed in data units per period of time; *e.g.* as blocks per second.

We need two auxiliary functions: *th* and *size*. The first one associates a throughput with a triple $(\vartheta_0, \vartheta_1, name)$, where ϑ_0 and ϑ_1 are the abstract addresses of the subprocesses that are communicating, and *name* is the channel name they use to interact. The function *size* associates a byte size with the couple $(\vartheta, name)$, where *name* is the sent data and ϑ is the abstract address from where the name has been sent. We use also the function *min* which returns the minimum value between its two arguments. The definition of the cost function follows:

$$\begin{aligned} \$(\vartheta\mu) &= \frac{size(obj(\mu))}{th(add(\vartheta), \epsilon, sbj(\mu))} \times \$_o(\vartheta) \\ \$(\vartheta(\vartheta_0\mu_0, \vartheta_1\mu_1)) &= \frac{size(obj(\mu_i))}{th(add(\vartheta\vartheta_0), add(\vartheta\vartheta_1), sbj(\mu_i))} \times min(\$_o(\vartheta\vartheta_0), \$_o(\vartheta\vartheta_1)) \end{aligned}$$

where the function $\$_o$ associates a real number, in the interval $(0, 1]$, with the occurrences of the summation operator:

$$\begin{aligned} \$_o(+_i\vartheta) &= \$_o(+_i) \times \$_o(\vartheta) \\ \$_o(\|_i \cdots \|_j +_k\vartheta) &= \$_o(\|_i \cdots \|_j +_k) \times \$_o(\|_i \cdots \|_j \vartheta) \\ \$_o(\|_i \cdots \|_j) &= 1 \\ &where\ i, j, k \in \{0, 1\}. \end{aligned}$$

The value returned by $\$_o$ represents a *slowing factor* due to the time spent on the execution of a nondeterministic choice, taking care also of the abstract address of the $+$ operator. This is done by recursively keeping all the parallel tags which are in the θ label (see the second equation of $\$_o$ definition). The definition of $\$_o$ heavily depends on the knowledge of the technical implementation of the system. We associate a *a slowing factor* with the $+$ operator, as we included the tags $+_0$ and $+_1$ for the summation in Def. 4. Nevertheless, it is possible to extend the definition of the function $\$_o$ including other tags for the other operators in those cases where it is known how to associate a rate with the execution of the operators. We summarize this execution in three simple steps: (1) introduce a tag for the operator in the proof term definition; (2) drop the structural congruence rules which involve the operator; (3) give a more structured definition for the function $\$_o$ associating a *slowing factor* with the operator, depending also on the abstract addresses used. The above definition is an extension of the cost function definitions in [2, 19, 7], deriving from some laboratory experiments which we made with student collaborations.

5.2 π -calculus model of a Web service

We adopt the model of the Web service presented in Sect. 4.2 considering four components: *Client*, *WebService*, *SOAPmsg* and *Discover*. Some comments about the notation; we shall write \bar{x} to represent an output action that does not send any message and $x()$ to represent an input action which does not receive any information, respectively. We then write $x.a$ for a place-holder that will be replaced with the value a . Moreover we do not write the list of the free names

associated with a constant definition, for instance we shall write *Client* without any parameter. We only write one parameter in the definition of $Client_2(rdMsg)$ to force the idea that it is a restricted name.

The *Discover* process interacts with the *Client* process by accepting a request on the public channel *dis* and replying by sending a private name to the client performing the request.

$$Discover \stackrel{def}{=} dis(x_askDes).(\nu des)\overline{x_askDes} des.Discover$$

The *SOAPmsg* process uses the public channel *client* to accept a description of the service, *x_des*, from any of the possible clients seeking to use a web service. The client will send also the name, *x_rdMsg*, over which he wants to receive the answer. To guarantee that the two sequential messages are sent by the same client, the *SOAPmsg* process waits for the second message on the name received as the first name, *x_des*. This mechanism encodes security features because cryptographic messages usually contain a secret key for the answer. Note that this mechanism is widely used in this example.

Now, the *SOAPmsg* is ready to send to the *WebService* the client's request on the public channel *web*. Then it will receive the answer that will be sent to the *Client* along the channel *x_rdMsg*.

$$SOAPmsg \stackrel{def}{=} client(x_des).x_des(x_rdMsg).\overline{web} x_des. \\ x_des(x_service)\overline{x_rdMsg} x_service.SOAPmsg$$

The *Client* process interleaves the activity of looking for web-services with various other activities that we generically express with input/output action on the channel *localComp*.

When the *Client* realizes that he needs to use a web service, he sends the request *askDes* to a discovery service, by means of a public channel *dis*. The *Client* sends the description, *x_des*, received by the discovery service, to the *SOAPmsg* using the channel *client* and he also sends the private name *rdMsg* over which he will asynchronously wait for the answer.

$$Client \stackrel{def}{=} localComp().Client \\ + \\ ((\nu askDes)\overline{dis} askDes.askDes(x_des). \\ (\nu rdMsg)\overline{client} x_des.\overline{x_des}(rdMsg).Client_2(rdMsg) \\)$$

$$Client_2(rdMsg) \stackrel{def}{=} localComp().Client_2(rdMsg) \\ + \\ rdMsg(x_service).Client$$

Finally, the *WebService* process interacts only with the *SOAPmsg* receiving the request on the public channel *web* and replying with the answer *service* on

the channel x_des .

$$WebService \stackrel{def}{=} web(x_des).(\nu service)\overline{x_des} service.WebService$$

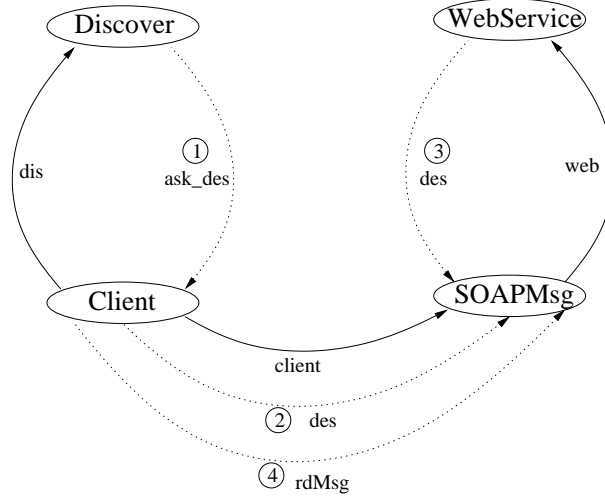


Fig. 4. The changes in the channel communication topology.

In this example the mobility is implemented by means of the channel-passing mechanism. In Fig. 4 we graphically show the changes in the channel communication topology. The arcs between the three processes are the communication channels: the one with dotted lines are temporary channels and the numbers indicate the time when the channel is used.

The complete system is given by putting in parallel the above three processes:

$$P \stackrel{def}{=} (Client \mid Discover) \mid (SOAPmsg \mid WebService)$$

The transition system system generated by applying the enhanced operational semantics is illustrated in Fig. 5.

The system is quite simple, but our modeling allows us to add in the system other *Client*, *SOAPmsg*, *WebService* and *Discover* as we want without changing the code of each process. Each session is identified by a unique message *des* and each sequence of messages in the session is connected by the mechanism of channel passing.

The table with the definitions of processes is in Fig. 6.

In order to obtain the rates for each proof term, we apply the cost function $\$$ with the measures in Tab. 5². For simplicity we always assume $\$_o(\vartheta) = 1$, for all labels.

² The values of the throughput are slight modifications of the real values in [15]

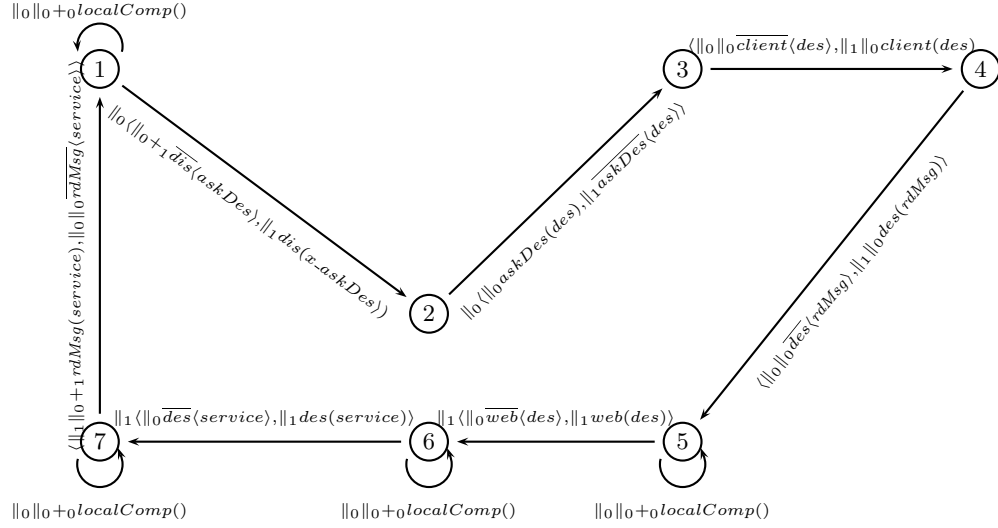


Fig. 5. Transition system of the process P .

1	$(Client \mid Discover) \mid (SOAPmsg \mid Webservice)$
2	$((\nu des)askDes(des).Discover \mid askDes(x_des).(\nu rdMsg)client \ x_des.x_des\langle rdMsg \rangle.Client_2(rdMsg)) \mid (SOAPmsg \mid Webservice)$
3	$(Discover \mid (\nu rdMsg)client \ des.x_des\langle rdMsg \rangle.Client_2(rdMsg)) \mid (SOAPmsg \mid Webservice)$
4	$(Discover \mid des\langle rdMsg \rangle.Client_2(rdMsg)) \mid (des(x_rdMsg).\overline{web} \ des.des(x_service).x_rdMsg \ x_service.SOAPmsg \mid Webservice)$
5	$(Discover \mid Client_2(rdMsg)) \mid (\overline{web} \ des.des(x_service).\overline{rdMsg} \ x_service.SOAPmsg \mid Webservice)$
6	$(Discover \mid Client_2(rdMsg)) \mid (des(x_service).\overline{rdMsg} \ x_service.SOAPmsg \mid (\nu service)\overline{des} \ service.Webservice)$
7	$(Discover \mid Client_2(rdMsg)) \mid (rdMsg \ service.SOAPmsg \mid Webservice)$

Fig. 6. Processes of the states of P transition system.

<i>th</i> : <i>Kbytes/secs.</i>		<i>size</i> : <i>Kbytes</i>	
$(\ _0\ _0, \ _0\ _1, des)$	871.58	$(\ _0\ _0, askDes)$	50
$(\ _0\ _0, \ _0\ _1, askDes)$	453.8	$(\ _0\ _1, des)$	200
$(\ _0\ _0, \ _1\ _0, client)$	974.53	$(\ _0\ _0, des)$	250
$(\ _0\ _0, \ _1\ _0, des)$	328.14	$(\ _0\ _0, rdMsg)$	100
$(\ _1\ _0, \ _1\ _1, des)$	254.6	$(\ _1\ _0, des)$	300
$(\ _1\ _0, \ _1\ _1, web)$	554.6	$(\ _1\ _0, service)$	500
$(\ _1\ _0, \ _0\ _0, rdMsg)$	749.65	$(\ _0\ _0, service)$	600

Table 5. Functions *th* and *size* applied to the process *P*.

The resulting Markov chain is:

$$\begin{array}{c}
 1 \xrightarrow{17.03\dots} 2 \\
 5 \xrightarrow{0.84\dots} 6
 \end{array}
 \left| \begin{array}{c}
 2 \xrightarrow{2.26\dots} 3 \\
 6 \xrightarrow{1.10\dots} 7
 \end{array} \right|
 \begin{array}{c}
 3 \xrightarrow{3.89\dots} 4 \\
 7 \xrightarrow{1.24\dots} 1
 \end{array}
 \left| \begin{array}{c}
 4 \xrightarrow{3.28\dots} 5
 \end{array}
 \right.$$

and the stationary distribution is $(0.57, 0.076, 0.131, 0.110, 0.028, 0.037, 0.041)$. Adopting the rewards technique [5] we can analyze in more detail the probability results. For example, if we are interested in the probability of using the data name *des*, independently of its location usage, we consider the following reward array $(0, 0, 1, 0, 1, 1, 0)$. Thus we obtain that the probability that the system is using the name *des* is 0.196.

6 Conclusions

If they were to be viewed purely formally as high-level description languages for specifying continuous-time Markov chains, then PEPA nets and the Stochastic π -calculus would be considered equally expressive. That is to say, for a given CTMC *C*, it is possible to construct a high-level model in either formalism such that the underlying CTMC derived from the model is isomorphic to *C*. This is a fundamental agreement in expressive power, but it is a rather weak one, similar to the agreement that all programming languages are Turing complete. In this paper and in related work [4] we have sought to understand the connections between these formalisms more thoroughly.

The modelling paradigms supported by PEPA nets and the Stochastic π -calculus EOS approach have a common root in using interleaving models of concurrent systems to first describe and then analyse the temporal behaviour of global and mobile code applications. However, there are many opportunities in such an enterprise to exercise creativity in the expression of concepts such as process mobility and performance metrics over models of mobile code systems. The differences between the PEPA nets approach and the EOS approach highlight points where different design choices were made.

Inside the behavioural description of a system the modeller needs to represent sequential execution and causal ordering of events. Over this aspect of the

behavioural modelling there is close agreement between PEPA nets and EOS. However, process algebras also need to represent the concurrent composition of sequential behaviours and concepts such as synchronisation, parameterisation, naming and scoping. In stochastically timed process algebras particularly there are many ways to design and justify the synchronisation operator for processes [11, 1] and different design decisions are naturally taken in the PEPA nets and EOS approaches.

Adjacent to this, and perhaps of greater importance, is the use of the process algebra machinery in defining the meaning of terms in the language and legitimising their analysis. The differences between PEPA nets and the EOS approach are most pronounced here. The EOS approach encodes the rules which are used to produce the derivatives of a process as proof terms in their derivations. This information is implicit in a PEPA net one-step derivation although a proof of any derivation could be obtained by revisiting the operational semantics of the language or by using an EOS semantics for PEPA nets [8]. The proof terms play a central role in the performance analysis process for EOS. The evaluation cost function is defined over the proof terms of the language and hence built into the language at the same level as the operational semantics. The cost function and the operational semantics interoperate, with structural congruence rules for operators being disabled by their use in the definition of the evaluation cost function.

In contrast for PEPA nets, performance measures over a model are defined outside the operational semantics for the language, and this separation is highlighted by the use of a separate logical language, PML_{ν} , for the expression of these measures. This separation means that the interpretation of the language constructs is unchanged across models and so tools supporting the language can perform optimisations such as quotienting by PEPA's bisimulation equivalence [9]. This operation is performed by rewriting the terms denoting process derivatives to amalgamate syntactically distinct terms which represent processes which no external observer could distinguish. This has the effect of reducing the state space of the system and therefore reducing the numerical computation effort which is needed to find the steady-state probability distribution for a given assignment of values to the symbolic rates of the model.

Despite these differences in methodology the present paper illustrates that the two modelling approaches can be used effectively in modelling real-world global computing applications and complement each other well in practical use. Both of the modelling methods used here are continuing to develop both in theory and in practical application. When, as in the present paper, we can compare modelling idioms in use we have the opportunity to see how to import analysis methods and techniques from one formalism to the other, to the benefit of both.

Acknowledgements: The authors are supported by the DEGAS (Design Environments for Global ApplicationS) IST-2001-32072 project funded by the FET Proactive Initiative on Global Computing.

References

1. J. T. Bradley and N. Davies. Reliable performance modelling with approximate synchronisations. In *Proceedings of the 7th International Workshop on Process Algebra and Performance Modelling (PAPM'99)*, pages 99–118, Prensas Universitarias de Zaragoza, September 1999.
2. L. Brodo, P. Degano, and C. Priami. A tool for quantitative analysis of calculus processes. In *ICALP Satellite Workshops*, pages 535–550, 2000.
3. L. Brodo, S. Gilmore, J. Hillston, and C. Priami. A stochastic π -calculus semantics for PEPA nets. In *Proceedings of the workshop on Process Algebras and Stochastically Timed Activities (PASTA)*, pages 1–17, Edinburgh, Scotland, June 2002.
4. L. Brodo, S. Gilmore, J. Hillston, and C. Priami. Mapping coloured stochastic Petri nets to stochastic process algebras. In P. Kemper, editor, *Proceedings of the ICALP Workshop on Stochastic Petri nets*, June 2003. To appear.
5. G. Clark and J. Hillston. Towards automatic derivation of performance measures from PEPA models. In *Proceedings of UKPEW*, 1996.
6. P. Degano and C. Priami. Non-interleaving semantics for mobile processes. *TCS: Theoretical Computer Science*, 216, 1999.
7. P. Degano and C. Priami. Enhanced operational semantics: a tool for describing and analyzing concurrent systems. *ACM Computing Surveys*, 33(2):135–176, 2001.
8. S. Gilmore and J. Hillston. An enhanced operational semantics for PEPA nets. DEGAS project internal document, May 2002.
9. S. Gilmore, J. Hillston, and M. Ribaud. An efficient algorithm for aggregating PEPA models. *IEEE Transactions on Software Engineering*, 27(5):449–464, May 2001.
10. S. Gilmore, J. Hillston, and M. Ribaud. PEPA nets: A structured performance modelling formalism. In T. Field, P.G. Harrison, J. Bradley, and U. Harder, editors, *Proceedings of the 12th International Conference on Modelling Tools and Techniques for Computer and Communication System Performance Evaluation*, number 2324 in Lecture Notes in Computer Science, pages 111–130, London, UK, April 2002. Springer-Verlag.
11. J. Hillston. The nature of synchronisation. In U. Herzog and M. Rettelbach, editors, *Proceedings of the Second International Workshop on Process Algebras and Performance Modelling*, pages 51–70, Erlangen, November 1994.
12. J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
13. J. Hillston and L. Kloul. Formal techniques for performance analysis: blending SAN and PEPA. Submitted for publication, 2002.
14. J. Hillston and L. Kloul. From SAN to PEPA: A technology transfer. In *Proceedings of the workshop on Process Algebras and Stochastically Timed Activities (PASTA)*, pages 56–76, Edinburgh, Scotland, June 2002.
15. Information Networks Division Hewlett-Packard Company. *Netperf: A Network Performance Benchmark*, February 1996. Revision 2.1.
16. K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, September 1991.
17. R. Milner. *Communicating and Mobile Systems: The π Calculus*. Cambridge University Press, Cambridge, England, 1999.
18. R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes. *Information and Computation*, 100(1):1–77, September 1992.

19. C. Nottegar, C. Priami, and P. Degano. Performance evaluation of mobile processes via abstract machines. *IEEE Transactions on Software Engineering*, 27(10):867–889, 2001.
20. C. Priami. Stochastic π -calculus. In S. Gilmore and J. Hillston, editors, *Proceedings of the Third International Workshop on Process Algebras and Performance Modelling*, pages 578–589. Special Issue of *The Computer Journal*, 38(7), December 1995.
21. C. Priami. Language-based performance prediction for distributed and mobile systems. *INFCTRL: Information and Computation (formerly Information and Control)*, 175(2):119–145, 2002.