A Precedence PEPA Model for Performance and Reliability Analysis

Jean-Michel Fourneau and Leïla Kloul

PRiSM, Université de Versailles Saint-Quentin,
45 Av. des Etats Unis, 78000 Versailles, France
{jmf, kle}@prism.uvsq.fr

Abstract. We propose new techniques to simplify the computation of the cycle times and the absorption times for a large class of PEPA models. These techniques allow us to simplify the model description to reduce the number of states of the underlying Markov chain. The simplification processes are associated with stochastic comparisons of random variables. Thus the simplified models are stochastic bounds for the original ones.

1 Introduction

In the recent years, several researchers have investigated ways to solve steady-state distributions for Stochastic Process Algebra models with exponential duration of activities such as PEPA models [10]. The tensor based representation [11] allows us to build large state spaces in a very efficient manner. However solving the steady-state distribution remains a difficult problem even if the bisimulation technique allows us to reduce the state space. Recently the process algebra formalism has also been used to solve transient problems [7], still under the Markovian assumption.

Here, we advocate a completely different approach which is not totally related to this Markovian assumption. First, we want to compute the distribution of the cycle time (if the model is well defined) or the distribution of the absorption time (if the model has an absorbing state) instead of the steady-state distribution. The cycle time is the delay between two successive visits to a specific state while the absorption time is the time until absorption. Cycle time is closely related to the throughput of the system while the distribution of the absorption time allows us to define the reliability of a system. By taking the average of these distributions, one can obtain the mean throughput and the average population with Little's formula or the mean time to failure. These quantities are in general significant for models based on customer's point of view rather than server's states.

We propose a two-level hierarchical approach. At the higher level, we consider a precedence PEPA model. Each component of the precedence model is a sub-model isolated from the other components. Because of the exponential duration of the activities in a PEPA component, these sub-models can be associated with continuous time Phase type distributions.

A. Horváth and M. Telek (Eds.): EPEW 2006, LNCS 4054, pp. 1–15, 2006.

Computing absorption time distribution is usually done by uniformization and analysis of transient discrete-time Markov chains. This technique requires a large number of vector-matrix multiplications. The matrix size is the number of states in the Markov chain. So it is important to find techniques which can be used to reduce this number of states. Cycle times computation are not necessarily based on Markovian assumption, even if exponential delays of individual activities may lead to the usual Markovian numerical analysis. For a class of decision-free Petri nets, cycle times are defined by recurrence relations [3]. Furthermore these relations are linear but on the max-plus semigroup. Such structures have been studied extensively in the context of random variables (see for instance Baccelli et al [1]). For more general systems, the computation of the cycle times is a complex problem. The stochastic comparison appears to be a promising technique to cope with this complexity.

If we need to compute the cycle time of a PEPA model which is too complex to analyse numerically, we design automatically a new model such that its cycle time is a bound for the exact one. This bound is stochastic: we do not compare reals but distribution functions. Thus stochastic bounds are far more accurate than worst-case analysis. If the new model has a reduced state space, we may then use numerical methods (or even analytical results) to efficiently solve the problem. Note that bounding some performance measures is often sufficient as quite often we only need to verify the requirements in terms of threshold. Stochastic bounds may also be applied to Markov chains (see [8] for a survey of the various techniques involved and [12] for an example of delays due to a Fair Queueing discipline).

Here we propose high level techniques which transform a PEPA model into simpler PEPA model. These techniques are based on stochastic bounds. They allow us to divide the problem into sub-problems or to replace a complete PEPA sub-model by a single activity. Here we just give some theoretical results, we will present in a sequel paper the algorithms we need and some numerical results.

The rest of the paper is organised as follows. In Section 2, we present some concepts of stochastic comparison while Section 3 gives a simple introduction to PEPA, the SPA we consider. Section 4 is devoted to the precedence PEPA model. Section 5 contains the main results of the paper. Finally in Section 6, we conclude our work with some remarks and future work.

2 A Simple Introduction to Stochastic Comparison

We restrict ourselves to finite Continuous Time Markov Chains (CTMC). Stoyan [14] defined the strong stochastic ordering (“st” ordering for short) by the set of non-decreasing functions. Bounds on the distribution imply bounds on these functions as well. Important performance measures such as average population, loss rates or tail probabilities are non decreasing functions. The second part of the definition for discrete random variables is much more convenient for an algebraic formulation and an algorithmic setting.

Definition 1. *Let X and Y be random variables taking values on a totally ordered space. Then X is said to be less than Y in the strong stochastic sense, that is, $X <_{st} Y$ iff $E[f(X)] \leq E[f(Y)]$ for all non decreasing functions f whenever the expectations exist.*

If X and Y take values on the finite state space $\{1, 2, \ldots, n\}$ with p and q as probability distribution vectors, then X is said to be less than Y in the strong stochastic sense, that is, $X <_{st} Y$ iff $\sum_{j=k}^{n} p_j \leq \sum_{j=k}^{n} q_j$ for $k = 1, 2, \ldots, n$.

Example 1. *Let $a = (0.1, 0.3, 0.4, 0.2)$ and $b = (0.1, 0.1, 0.5, 0.3)$. We have $a <_{st} b$ as:*

$$\left[\begin{array}{ll} 0.2 & \leq 0.3 \\ 0.2 + 0.4 & \leq 0.3 + 0.5 \\ 0.2 + 0.4 + 0.3 & \leq 0.3 + 0.5 + 0.1 \end{array}\right.$$

Sufficient conditions for comparison for CTMC are known for a long time [14]. The stochastic comparison of CTMC implies that their steady-state and transient distributions are also ordered.

Theorem 1 (Stoyan [14], page 193). *Let us consider two CTMC $Z1$ and $Z2$ on the same state space whose transition rate matrix are respectively $Q1$ and $Q2$. If*

1. *$Z1_0 <_{st} Z2_0$*
2. *$\sum_{k \geq l} Q1(i,k) \leq \sum_{k \geq l} Q2(j,k)$ for all $i \leq j$ and for all l which satisfy $l \leq i$ or $l \geq j$.*

then $Z1 <_{st} Z2$.

It may be important to compare Phase type random variables with exponential ones because it allows building a smaller Markov chain. Let us first define a family of random variables well known in reliability modelling [4].

Definition 2 (New Better than Used in Expectation). *Let X_t be the residual time of X, given that $X > t$. X is said to be NBUE if $E(X_t) \leq E(X)$ for all t.*

For instance, Erlang, uniform and constant random variables are NBUE. This family leads to another stochastic ordering: the increasing convex ordering which is used to compare random variables with exponentials.

Definition 3. *Let X and Y be two random variables on the same space ϵ, X is smaller in increasing convex order than Y, if and only if $E(f(X)) \leq E(f(Y))$ for all convex and non decreasing functions f on ϵ, provided that the expectations exist. The relation is denoted by $X <_{icx} Y$.*

Property 1 ([14]). *If X is NBUE of mean m, then X is smaller in increasing convex ordering than an exponentially distributed random variable of mean m.*

The icx ordering also provides a very intuitive lower bound.

Property 2 ([14]). *For any arbitrary positive random variable X, $E(X) <_{icx} X$.*

We also have two very simple properties which will be used to derive bounds at the higher level of a model from bounds obtained at the lower level.

Property 3. *The Max and Plus operators are convex and non decreasing functions.*

Property 4. *Let X and Y be two r.v. such that $X <_{icx} Y$, then for all convex and non decreasing function f, we have $f(X) <_{icx} f(Y)$.*

Finally, we can compare the absorption time of Markov chains [5] as stated in the following property.

Property 5. *Let $Z1$ and $Z2$ be two homogeneous Markov chains with an absorbing state n and let $T_a(Z1)$ and $T_a(Z2)$ denote absorption times for the two chains. If $Z1 <_{st} Z2$ or $Z1 <_{icx} Z2$ then $T_a(Z2) <_{st} T_a(Z1)$.*

Note that the "st" comparison of absorption times is now on random variables T_a defined on the time instants, not on the states.

3 PEPA

In PEPA, a system is viewed as a set of *components* which carry out *activities*. Each activity is characterised by an *action type* and a duration which is exponentially distributed. Thus each activity is defined by a couple (α, r) where α is the action type and r is the *activity rate*. Because of the exponential distribution of the activity duration, the underlying Markov process of a PEPA model is a continuous time Markov process.

PEPA formalism provides a set of combinators which allows expressions to be built, defining the behaviour of components, via the activities they engage in. Below, we present informally the combinators we are interested in and which are necessary to our model. For more details about the formalism, see [10].

Constant: noted $S \overset{def}{=} P$, it allows us to assign names to components. To component S, we assign the behaviour of component P.

Prefix: noted $(\alpha, r).P$, this combinator is the basic mechanism by which the behaviours of components are constructed. The component carries out activity (α, r) and subsequently behaves as component P.

Choice: noted $P_1 + P_2$, this combinator represents competition between components. The system may behave either as component P_1 or as P_2. All current activities of the components are enabled. The first activity to complete, determined by the race condition, distinguishes one of these components, the other is discarded.

Cooperation: noted $P_1 \underset{L}{\bowtie} P_2$, it allows the synchronisation of components P_1 and P_2 over the activities in the cooperation set L. Components may proceed independently with activities whose types do not belong to this set. A particular case of the cooperation is when $L = \emptyset$. In this case, components proceed with activities independently and are noted $P_1 || P_2$.

In a cooperation, the rate of a shared activity is defined as the rate of the slowest component. For a component P_1 and an action type α, the working capacity is termed the *apparent rate* of α in P_1. It is the sum of the rates of the α type activities enabled in P_1. The apparent rate of α in a cooperation between P_1 and P_2 over α will be the minimum of the apparent rate of α in P_1 and the apparent rate of α in P_2.

The rate of an activity may be unspecified for a component and is noted $\top$. Such a component is said to be *passive* with respect to this action type and the rate of this shared activity is defined by the other component in cooperation.

In PEPA, when a component C carries out an activity (α, r) and subsequently behaves as component C', this one is said to be a *derivative* of C. From any PEPA component C, the *derivative set*, denoted $ds(C)$, is the set of derivatives (behaviours) which can evolve from the component. This set is defined recursively.

The evolution of a PEPA model is governed by the Structured Operational Semantics (SOS) rules of the language [10]. These rules define the admissible transitions or state changes associated with each combinator.

Necessary (but not sufficient) conditions for the ergodicity of the Markov process in terms of the structure of the PEPA model have been identified and can be readily checked [10]. These conditions imply that the model must be a *cyclic* PEPA component. The model should be constructed as a cooperation of *sequential* components, i.e. components constructed using only prefix, choice and constants. This leads to formally define the syntax of PEPA expressions in terms of *model components* P and *sequential components* S:

$$P ::= A \mid P \underset{L}{\bowtie} P \mid P/L \qquad\qquad S ::= (\alpha, r).S \mid S + S \mid A_s$$

where A denotes a constant which is either a model or a sequential component and A_s denotes a constant which is a sequential component. Thus the compositional structure of PEPA models is at the level of the cooperating components; such models are considered as *well-defined*.

4 The Precedence PEPA Model

We consider that a system is represented by a set of components which have the same general behaviour as they wake up, proceed with their activities and then make other components wake up. The components are assumed to be initially asleep (off) and cannot proceed with the execution of their activities unless they are woken up. We assume a precedence relation between the enabling of the components in the set as the results of some components can be used as an input by other components. The components are labelled to allow a representation of this precedence relation. We assume the following properties for the set of labels:

1. the set is totally ordered,
2. the set has a unique minimal element which is denoted by $Comp_0$ for convenience,
3. and the set has a unique maximal element which is denoted by $Comp_n$.

We assume that $Comp_0$ constitutes the starting component of the system and $Comp_n$ the last one to be enabled. When $Comp_n$ completes, the system is assumed to have the same behaviour, restarting from the beginning, i.e. $Comp_0$ (Figure 1). Furthermore, we assume that the precedence relation between the components is a Directed Acyclic Graph (DAG) modified by this return arc from $Comp_n$ to $Comp_0$.

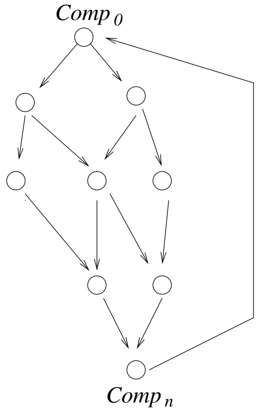

Fig. 1. The precedence relation between the components

Our system specifications allow us to consider two kinds of analysis, performance analysis and reliability analysis. The former exploits the presence of the return arc from $Comp_n$ to $Comp_0$ to compute performance measures such as the cycle times. The latter is only possible if we have in our system an absorbing state, that is the precedence relation between the components is a real DAG. Moreover, $Comp_n$ must contain an absorbing state.

4.1 Formal Description of the System

To represent the precedence relation characterising our system, we define two families of sets P_i and S_i. P_i is the set of components which must complete their activities before $Comp_i$ is woken up and S_i represents the set of components which are enabled when $Comp_i$ has completed its activities. Note that the two families of sets have to be consistent.

We describe the system using $n + 1$ components. Each component $Comp_k$, $k = 0 \dots n$, is woken up thanks to activity $wake\_up_{jk}$ where j is a predecessor of $Comp_k$, that is $j \in P_k$. Once awake, the component can then proceed with its own activities $\alpha_{k,l}$, $l = 1..m_k$, where m_k is the number of activities of $Comp_k$. Note that these activities $(\alpha_{k,l})$ are all individual activities and once $Comp_k$ has finished executing them, it will wake up the components which are in its set of successors S_k.

The behaviour of the components of the system are modelled using the following equations:

$$Comp_0 \overset{def}{=} (start, w_1).(\alpha_{0,1}, r_{0,1})...(\alpha_{0,m_0}, r_{0,m_0}). \prod_{i \in S_0} (wake\_up_{0i}, s).Comp_0$$

$$Comp_i \overset{def}{=} (wake\_up_{0i}, \top).(\alpha_{i,1}, r_{i,1})...(\alpha_{i,m_i}, r_{i,m_i}). \prod_{k \in S_i} (wake\_up_{ik}, s).Comp_i \qquad \forall i \in S_0$$

$$Comp_k \overset{def}{=} \prod_{j \in P_k} (wake\_up_{jk}, \top).(\alpha_{0,k}, r_{0,k})...(\alpha_{k,m_k}, r_{k,m_k}). \prod_{j \in S_k} (wake\_up_{kj}, s).Comp_k \qquad \forall k \in S_i$$

$$Comp_n \overset{def}{=} \prod_{j \in P_n} (wake\_up_{jn}, \top).(\alpha_{n,1}, r_{n,1})...(\alpha_{n,m_n}, r_{n,m_n}).(end, w_2).Comp_n$$

where the notation of the form $\prod_{k \in A_i} (\beta_{ik}, r)$ refers to $(\beta_{ii_1}, r).(\beta_{ii_2}, r). \cdots . (\beta_{i_{|A_i|}}, r)$.

The use of $\prod_{k \in S_i} (wake\_up_{ik}, s)$ allows us to model the case where $Comp_i$ wakes up all the components in its successors set one by one. Whereas the use of $\prod_{j \in P_n} (wake\_up_{jn}, \top)$ like in $Comp_n$ models the case where a component has to wait for several predecessors to complete their activities before proceeding with its own activities.

Additionally, we consider another component *Clock*, which allows starting, and restarting the system only once $Comp_n$ has completed its activities. This additional component has to synchronise with $Comp_0$ on activity *start* then on activity *end* with $Comp_n$.

$$\begin{aligned} Clock &\overset{def}{=} (start, \top).Clock_0 \\ Clock_0 &\overset{def}{=} (end, \top).Clock \end{aligned}$$

The behaviour of the complete system is modelled as the interaction of its components as follows:

$$\begin{aligned} System \overset{def}{=} Clock \underset{\{start,end\}}{\bowtie} (\ldots(Comp_0 \underset{\{wake\_up_{0i}/i \in S_0\}}{\bowtie} (\ldots||Comp_i||\ldots)_{i \in S_0}) \\ \underset{\{wake\_up_{ik}/k \in S_i\}}{\bowtie} (\ldots||Comp_k||\ldots)\ldots)_{k \in S_i} \ldots(\ldots||Comp_j||\ldots)_{j \in P_n} \\ \underset{\{wake\_up_{jn}/j \in P_n\}}{\bowtie} Comp_n)\ldots) \end{aligned}$$

4.2 Reliability Analysis Using the PEPA Model

Component *Clock* is only necessary in the case where a performance analysis is targeted as it allows modelling the return arc of the precedence relation between the components. Reciprocally, whenever reliability analysis is the objective, component *Clock* is not only unnecessary, but has to be removed from the model. As all its activities have an unspecified rate ($\top$), its removal from the model has no impact on the remaining components.

Moreover, to ensure that the underlying Markov chain of the model has an absorbing state, we need the last component $Comp_n$ to not return to its initial state once activity *end* has been completed. Therefore, we have to redefine $Comp_n$ as follows:

$$Comp_n \overset{def}{=} \prod_{j \in P_n} (wake\_up_{jn}, \top).(\alpha_{n,1}, r_{n,1})...(\alpha_{n,m_n}, r_{n,m_n}).(end, w_2).Comp_n^*$$

In the first definition of $Comp_n$, the first and the last derivatives (states) were the same, that is $Comp_n$. In the new definition, the first derivative is still $Comp_n$, but the last one is different. It is denoted by $Comp_n^*$ and models the absorbing state of the system. □

As explained above, finding the absorption times and the cycle times for our class of systems are connected problems. The former assumes that the durations of the activities are independent, which is the case in our model. Thus the delays from successive beginnings of the first component form a renewal process. In the case of the latter, once the last component has completed its activities, the first component is woken up.

In the following we only consider the absorption times. The computation of the cycle times can be easily deduced from the results developed for the absorption times.

5 Reliability Analysis: Computing the Absorption Times

The PEPA model is a two-level hierarchy model, the component level and the model level. Therefore the computation of the bounds on the absorption times rely on two different classes of techniques according to the hierarchy level considered. However, all these techniques are based on the recurrence equations we can obtain at the higher level of our hierarchy.

In the following, once we show how to obtain the recurrence equations, we first propose techniques which can be applied on the PEPA sub-models (components). Then we show how we can modify the precedence relation between the PEPA components to derive simpler models.

The bounds on the absorption times are obtained from the recurrence equations which can be established on instants of transition. Let (t_i) (resp. (b_i)) be the completion time (resp. the wake up time) of component $Comp_i$. The main results come from the type of equations connecting instants t_i to other instants t_j if $Comp_j$ is a predecessor of $Comp_i$ in the precedence model.

Let d_i be the service time of $Comp_i$, that is the time required for $Comp_i$ to proceed with all its activities $\alpha_{i,l}$, $l = 1..m_i$. As $Comp_i$ is a PEPA component where all the activities have exponential durations, the total duration of $Comp_i$ has a continuous Phase type distribution (PH in the following).

Clearly for all i, we have $t_i = b_i + d_i$. Now it is important to note that $Comp_i$ wakes up as soon as all the components in its predecessors set P_i have completed all their activities. Thus

$$b_i = max_{j \in P_i}(t_j)$$

After substitution, we get:

$$t_i = d_i + max_{j \in P_i}(t_j) \tag{5.1}$$

Thus we obtain a linear equation on vector (t_i) using two operators: the addition and the maximum. Such linear equations have been extensively studied as they allow new types of analytical or numerical methods which are not based on exponential delays or embedded Markov chains. In this paper we assume that activities have exponential durations, but as the random variables d_i model a PEPA sub-model $(Comp_i)$ duration, d_i has a PH distribution.

5.1 Bounds Due to Service-Time of Activities

Using equation 5.1 and properties 3 and 4 we obtain the first comparison results if all the random variables d_i are New Better than Used in Expectation (NBUE). Indeed, the NBUE property implies the relation between a single random variable and an exponential one with the same mean. Equation 5.1 and property 3 show that t_i is defined using two increasing and convex operators. Property 4 states that the relation holds for the absorption time.

Property 6. *Consider a precedence PEPA model such that the PH distribution associated with* $Comp_i$*, a component of the model, is NBUE. The absorption time is upper bounded in the increasing convex sense by the absorption time of the same model where* $Comp_i$ *is replaced by a single activity with rate* $E(d_i)$*.*

Similarly, we have a lower bound provided by constant random variable with the same mean using property 1.

Property 7. *Consider a precedence PEPA model with arbitrary random variables. The absorption time is lower bounded in the increasing convex sense by the absorption time of the same model where the PH distribution associated with* $Comp_i$*, a component of the model, is replaced by a constant with the same mean.*

Note that, in this case, the resulting model is not a usual PEPA model anymore as we have a component with a non exponential duration. One can also obtain a lower bound of the completion time by a very simple argument on the duration of any component. This is stated in the following property.

Property 8. *For all positive random variables* X*, we have zero* $\leq_{st} X$ *where zero is considered as the constant r.v. with mean* 0*.*

Finally,

Property 9. *Consider a precedence PEPA model. The absorption time is lower bounded in the strong stochastic sense by the absorption time of the same model where* $Comp_i$*, a component of the model, has been removed and where we have added arcs* (C_1, C_2) *in the precedence model for all components* C_1 *and* C_2 *such that arcs* $(C_1, Comp_i)$ *and* $(Comp_i, C_2)$ *were in the initial precedence model.*

So the main question remaining now is whether a PH distribution is NBUE or not. To the best of our knowledge such a problem has never been studied before.

5.1.1 Phase Type and NBUE Distributions

A Phase type distribution is the absorption time of a transient Markov chain on state space $1..N$. It is defined by the initial distribution (say σ) and the transition rate matrix Q. Let Y be this chain and X the absorption time of Y knowing σ. Without loss of generality we assume that there exists only one absorbing state which is the last one (i.e N). Thus

$$Q = \left[\begin{array}{c|c} \quad T \quad & t \\ \hline 0 & 0 \end{array}\right]$$

with $t = -Te$, e being a column of 1. Without loss of generality we assume that the initial distribution σ is $(1, 0, \ldots, 0)$. Indeed, a general distribution can be considered if we add an extra state at the beginning.

Let X_t be the residual time before absorption, given that $X \geq t$. Remember that the distribution of X is NBUE iff $E(X_t) \leq E(X)$ for all t.

At time t, chain Y is in state j with probability $Pr(Y_t = j|Y_0 = 1)$. Let μ_k be the rate of activity k. The expectation of the remaining time before absorption in Y can be computed using the mean number of passages in any state of Y before being absorbed. Of course these quantities depend on the initial state of the chain. Let $a_{i,j}$ be the average number of visits to state i when the initial state of chain Y is j. Clearly, we have:

$$E(X) = \sum_{i=1}^{N-1} \frac{a_{i,1}}{\mu_i}$$

Similarly because of the memoryless property, the remaining time after t is obtained by conditioning on the state reached at time t as follows:

$$E(X_t) = \sum_{j=1}^{N-1} Pr(Y_t = j|Y_0 = 1) \sum_{i=1}^{N-1} \frac{a_{i,j}}{\mu_i}$$

Now, we must compare E and $E(X_t)$ to check if a distribution is NBUE. First we obtain a very simple result which is quite useful.

Property 10. *If, for all state i, we have for all state j, $a_{i,j} \leq a_{i,1}$ then the PH distribution is NBUE.*

Proof: If $a_{i,j} \leq a_{i,1}$ for all j, then any convex sum of $a_{i,j}$ is smaller than $a_{i,1}$. And $\sum_{j=1}^{N-1} Pr(Y_t = j|Y_0 = 1)a_{i,j}$ is such a convex sum. Finally we get $E(X_t) \leq E(X)$.
Property 10 allows us to derive the following one:

Property 11. *The hypoexponential distribution is NBUE.*

Remember that the hypoexponential distribution is a generalisation of the Erlang distribution where the exponential stages do not have the same rate. For an hypoexponential distribution we get $a_{i,1} = 1$ as we visit every stage exactly once. If we begin at stage j, the number of visits is 1 or 0 depending if the stage to visit is after j or before j. Thus $a_{i,j} \leq a_{i,1}$ for all i, j.

The hypoexponential distribution is easy to detect from a PEPA specification of a component. It is a set of successive individual activities without any choice operator.

Theorem 2. *If a PEPA component C is constructed using only the prefix operator, the rates of successive individual activities of C are the rates of the stages of an hypoexponential distribution and the completion time of C is NBUE. These individual activities of C can therefore be aggregated into a single individual activity with the same mean.*

Proof: Consider a PEPA component which consists of a sequence of individual activities in which the only operator used is the prefix. As each activity α in the sequence is exponentially distributed with rate r_α, these rates constitute the rates of the stages of an hypoexponential distribution. As the hypoexponential distribution is NBUE, according to Property 11, the completion time of the sequence of activities of the component is NBUE. Consequently, this sequence of activities can be aggregated and replaced by a single individual activity with the same mean. □

According to Theorem 2, we can replace a PEPA component with successive individual activities by a component with a single individual activity. In this context, the stochastic comparison allows a drastic reduction of the complexity. Moreover, it allows a new type of aggregation which is not exact, but which provides proved bounds.

Thus, in our precedence PEPA model, we can aggregate the sequence of individual activities $(\alpha_{k,1}, r_{k,1}).(\alpha_{k,2}, r_{k,2}) \ldots (\alpha_{k,m_k}, r_{k,m_k})$ of component $Comp_k$ into a single activity (α_k, r_k) where $r_k = (\frac{1}{r_{k,1}} + \frac{1}{r_{k,2}} + \ldots + \frac{1}{r_{k,m_k}})^{-1}$.

More generally we get the following characterisation:

Property 12. *Consider an arbitrary PH distribution. If for all state i and j we have*

$$\sum_{i=1}^{N-1} \frac{a_{i,j}}{\mu_i} \leq \sum_{i=1}^{N-1} \frac{a_{i,1}}{\mu_i}$$

then the PH distribution is NBUE.

Proof: Again $E(X_t)$ is a convex sum with coefficients $Pr(Y_t = j|Y_0 = 1)$ of the first quantities in the relation. Thus if the set of inequalities is satisfied for all j and i, we get $E(X_t) \leq E(X)$.

Let us now consider acyclic PH distributions. Assume that the states of Y are ordered according to the natural ordering associated to this directed acyclic graph.

Such a family of distributions have been shown to be very efficient when we have to fit a general distribution [6]. This family is quite large and it contains Coxian distribution. For an acyclic PH distribution using this numbering assumption, we clearly have $a_{i,j} = 0$ if $j < i$. As the graph of Y does not contain any directed cycle, any state will be visited zero or once. The expected number of visits is also the probability of visit. It is quite simple to compute $a_{i,1}$ and $a_{i,j}$ from the transition probability matrix embedded in matrix T.

Once we have computed $a_{i,j}$ for all i and j we can check the sufficient relations stated in Property 12. In general Coxian distribution are not NBUE but Property 12 gives a very simple way to check it. Note that Property 12 also applies when the PH is not acyclic. However, the complexity of computing $a_{i,j}$ is now much higher.

Finally one can bound an arbitrary acyclic PH distribution by an hypoexponential distribution.

Theorem 3. *Let X be an arbitrary acyclic PH distribution associated with transition rate matrix Q. Let Z be the hypoexponential associated with transition matrix R. Assume now that the states of the chains are ordered according to the DAG. If*

$$R(i, i+1) = \sum_{j \geq i+1} Q(i, j) \quad \textit{and} \quad R(i, j) = 0 \ \forall j \neq i, i+1$$

then $X <_{st} Z$.

Proof: As X and Z are PH distributions, they are also absorption times of CTMC. Theorem 1 states that the comparison of CTMC can be easily checked. Property 5 shows that the comparison of CTMC implies the comparison of absorption times. So it is sufficient here to state that the chains associated with the distributions satisfy both conditions of Theorem 1.

- The first condition is trivial as the initial distribution is the same.
- Remember that the states of the chains are ordered according to the DAG. Thus matrix Q is upper triangular. The lower triangle of Q and R clearly satisfy the constraints of Theorem 1. Finally one can easily check the upper triangle part of the relation as $R(i, i+1) = \sum_{j \geq i+1} Q(i, j)$.

Thus one can transform any acyclic PH distribution into an hypoexponential one. The PEPA sub-models (components) are transformed as well. The numbering of activities defines the sequential ordering of the activities. The definition of matrix R gives the activity rates in this transformed PEPA sub-model.

Let us now turn to other techniques based on the precedence relation at the higher level, the model level.

5.2 Changing the Precedence Model to Obtain Bounds

Such transformations of the model are strongly related to the rules proposed by Bacelli and Liu [2] for queueing networks with synchronisations and by Vincent and Pekergin for tasks graph [13]. Even if the problems are not the same, they all share this property of linear evolution equation with max-plus operators (see also [3] for Petri nets). The main transformations which have been proposed consist of the addition or the deletion of a node, an edge, a place or a transition. However these modifications of the graph do not always help for the resolution of the model.

To compute an upper bound for our system, we propose to add a new component in the precedence model. The main idea is to make the model separable. Then we divide the model into two sub-models which are analysed in isolation.

Assume that the precedence model has n components. Let us assume that the components of the model are ordered according to a topological ordering consistent with the precedence relation: if there exists a directed edge from $Comp_i$ to $Comp_j$ then $i < j$. We first add a new component (say $Comp_{n+1}$). Then we modify the directed edges of the precedence model. Let m be an arbitrary integer between 1 and n. We add directed edges in the precedence model from any component $Comp_i$ in $1..m$ to $n+1$ and from component $Comp_{n+1}$ to any component in $m+1..n$. Such a component is denoted as a *star*.

Note that now the model is not correctly ordered: the star component ($Comp_{n+1}$) does not have a correct index according to the precedence relation. However the new model is still a precedence model.

Let us now prove that this transformation provides an upper bound. First we reorder the state according to the new precedence model. The component we have added receives number $m+1$. Without loss of generality we assume that the components between 1 and m keep the same numbers they had before the insertion while the components numbers previously between $m+1$ and n increase by 1.

Considering equation 5.1, we just derive the new sets P_i as a function of the sets before the insertion. We have:

$$\begin{cases} P_i & \leftarrow P_i & \forall i = 1..m \\ P_{m+1} & \leftarrow \{1..m\} & \\ P_i & \leftarrow P_{i-1} \bigcup \{m+1\} & \forall i = m+2..n \end{cases}$$

Let us denote by t'_i the new values of the completion time. Clearly we have $t'_i = t_i$ for all $i \leq m$. As sets P_i are now larger, we also have: $t'_i \geq t_{i-1} + d_{m+1}$.

Theorem 4. *Let m be an arbitrary integer in 2..n-1, the absorption time is upper bounded in the strong stochastic sense by the absorption time of the same model with a star component added with label $m+1$.*

Adding a *star* has also an effect on the resolution algorithm. The model is now separable into two sub-models containing $Comp_1$ to $Comp_m$ for the first one and $Comp_{m+1}$ to $Comp_{n+1}$ for the second one. Indeed, to be absorbed in $Comp_{n+1}$, one must visit first $Comp_{m+1}$ and then travel from $Comp_{m+1}$ to $Comp_{n+1}$. We compute the time to be absorbed in $Comp_{m+1}$ in a sub-model and then we compute the time to be absorbed in $Comp_{n+1}$ knowing that the initial component is $Comp_{m+1}$ at time 0. These two random variables are independent

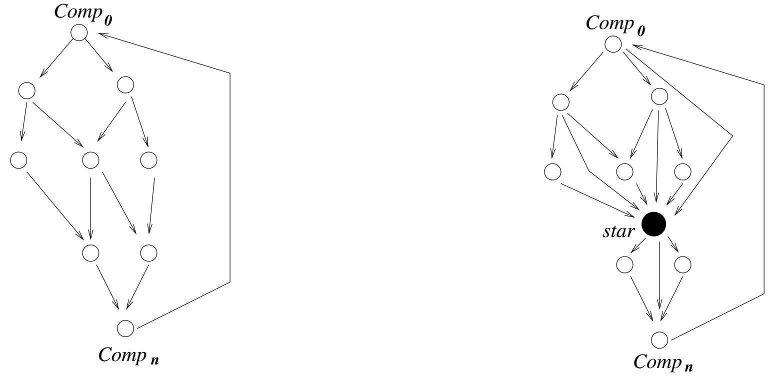

Fig. 2. Adding a star

and the distribution of the global absorption time is the convolution of the two distributions obtained from the sub-models. If we assume that the components set is equally divided by the insertion of the *star*, the stochastic comparison allows a drastic reduction of the complexity of the analysis.

6 Conclusion

The approach we have presented here constitutes a first step towards a new hierarchical resolution of hierarchical models. Indeed, we must improve our resolution techniques which are now far away from our modelling skills. Stochastic comparison is a very efficient approach to simplify models and obtain bounds. We can apply bounds on the transition times like in this paper or on the states. Both approaches rely on a monotonicity property which is implicit on transition instants associated with a precedence model. Precedence PEPA models are quite general but a natural extension to this work will be the generalisation of this type of method for an even larger set of models. It must be clear that this approach requires hierarchical models where the high level exhibits some monotonicity property which must be consistent with the comparison we made for sub-models absorption times. These are the key properties of the approach. However they are limited neither to precedence PEPA models nor to (max,+) semi-ring (again see [1]). For instance, PEPA nets, a new hierarchical modelling technique [9], are based on a high level model which is mainly a Finite State Machine (a simplified Petri Net with a limited interconnection between places and transitions) and on sub-models associated with the places of the net. The sub-models are PEPA models. It is worthy to remark that when the vertex cut of the directed cycles of the FSM has size one then we can derive the same linear equations on (max,+) semi-ring, from the PEPA net model, as the ones we have obtained here. Indeed when we remove this directed arc, the graph of the FSM becomes a DAG and it shows the relation with a precedence PEPA model. Thus all the techniques presented here can be applied to this new modelling technique as well.

References

1. F. Baccelli, G. Cohen, G.J. Olsder, J.P. Quadrat, “Synchronization and Linearity, an algebra for Discrete Event Systems”, Wiley & Sons, 1992.
2. F. Baccelli, Z. Liu, “On the executions of parallel programs on multiprocessor systems - a queueing theory approach”, JACM, V 37, N2, pp:373–414, 1990.
3. F. Baccelli, Z. Liu, “Comparison properties of stochastic decision free Petri nets”, Inria Report 1433, 1991.
4. R. E. Barlow, F. Proschan, “Statistical Theory of Reliability and Life Testing”, 1981.
5. M. Ben Mamoun, A. Busic, J.M. Fourneau, N. Pekergin, “Increasing Convex Monotone Markov Chains: Theory, Algorithm and Applications”, Submitted.
6. A. Bobbio, A. Horvath, M. Scarpa, M. Telek, “Acyclic Discrete Phase type distributions; properties and a parameter estimation algorithm”, Performance evaluation Vol.54, pp:1–32, 2003.
7. J. T. Bradley, N. Dingle, S.T. Gilmore, W.J. Knottenbelt “Extracting Passage Time form PEPA models with the Hydra tool: a case study”, UKPEW 2003.
8. J.M. Fourneau, N. Pekergin. “An algorithmic approach to stochastic bounds”, LNCS 2459, Performance evaluation of complex systems: Techniques and Tools, pp:64–88.
9. S. Gilmore and J. Hillston and L. Kloul and M. Ribaudo, “PEPA nets: A structured performance modelling formalism”, Performance Evaluation, vol.54, pp:79–104, Elsevier Science, 2003.
10. J. Hillston, *A compositional Approach to Performance Modelling*, PhD Thesis, The University of Edinburgh, 1994.
11. J. Hillston, L. Kloul, “An Efficient Kronecker Representation for PEPA Models”, LNCS, N.2165, pp:120-135, Springer Verlag, Proceedings of the Joint International Workshop, PAPM-PROBMIV 2001, Aachen, Germany, September 2001.
12. N. Pekergin, “Stochastic delay bounds on fair queueing algorithms”, Proceedings of INFOCOM’99, pp:1212–1220, New York, 1999.
13. N. Pekergin, J.M. Vincent, “Stochastic Bounds on Execution Times of Parallel Programs”, IEEE Transactions on Software Engineering, Vol. 17, No. 17, October 1991.
14. D. Stoyan, “Comparison Methods for Queues and Other Stochastic Models”, Wiley & Sons, New York, 1983.