REGULAR STOCHASTIC PETRI NETS

C. Dutheillet S. Haddad

Laboratoire MASI - CNRS UA 818
Université Pierre & Marie Curie
4 place Jussieu
75252 Paris Cedex 05

ABSTRACT

An extension of regular nets, a class of colored nets, to a stochastic model is proposed. We show that the symmetries in this class of nets make it possible to develop a performance evaluation by constructing only a graph of symbolic markings, which vertices are classes of states, instead of the whole reachability graph. Using algebraic techniques, we prove that all the states in a class have the same probability, and that the coefficients of the linear system describing the lumped Markov process can be calculated directly from the graph of symbolic markings.

Topics : Higher-level net models, stochastic nets.

I - INTRODUCTION

As they are a function of the number and the complexity of the processes to be represented, Petri net models of distributed computer systems must quickly face problems due to the huge size of the reachability graph.

The use of higher level models (e.g., colored nets, Pr/T nets) can make the task of the modeler easier. However, the existing analysis methods often make it necessary to go back to an equivalent ordinary net, so that there is no improvement in the complexity.

In order to face the exponential increase in the number of states, Jensen [Jen 81a] has introduced in colored Petri nets some equivalence relations that take into account the symmetries of the model. He showed that those relations were directly depending on the type of the color functions. He also defined analysis methods associated with the use of simple functions [Jen 81a, Hub 84].

Research has been done to extend the symmetries in colored nets to the stochastic domain [Zen 85, Lin 87, Chi 88]. Some attempts to reuse the symmetry properties have been proposed. One approach [Zen 85] has been to extend the reduced graph developed by Jensen to a stochastic model. However, when using this technique, the partition of states in classes can lead to a non Markovian process. In [Lin 87], the reduced process is actually Markovian, but no method is provided to build it automatically. The approach in [Chi 88] was to build the reachability graph, to group the states in classes according to some symmetry relation, and then to partition these classes in subclasses until the resulting process is Markovian. In fact, as soon as the color functions are general, no existing method avoids the development of the whole reachability graph. Indeed the simplification due to the symmetries can be used only once the reachability graph has been constructed, even sometimes requiring an expensive preliminary analysis.

This paper aims at showing that in the case of colored nets that do not use general functions, the results obtained during the structural analysis can be extended to the performance evaluation. In order to optimize the analysis, our study is based on a particular class of colored nets, the regular nets [Had 87]. In a regular net, the objects in a class have similar behaviors. The color domains of places and transitions are Cartesian products of object classes. The color functions are Cartesian products and linear combinations of two basic functions, one selecting a unique object in a class, the other synchronizing all the objects of the class. For this class of nets, a condensed representation of the reachability graph, the symbolic reachability graph, can be defined. The symmetry properties of this graph are used to simplify the quantitative analysis of the model.

The paper is organized as follows. In the next section, we give the definition and an example of regular net. Section 3 will presents the construction of the graph of symbolic markings. The stochastic model derived from a regular net will be presented in Section 4. In Section 5, we propose an algorithm for computing the state probabilities from the graph of symbolic markings. Our algorithm maps the symbolic reachability graph on a lumped Markov chain and uses the labels of the symbolic arcs to compute the transition rates. In the last section we prove the correctness of the algorithm. The proof organizes into three steps. First we prove that the developed Markov chain always has a solution such that all the ordinary markings within a symbolic marking have the same probability. Then we show that the lumping condition is verified by the Markov chain, and as a consequence, that the linear system to solve can be reduced. Finally we prove that this reduced system is the one obtained with our algorithm. Moreover, as the number of ordinary markings within a symbolic marking is computed by the algorithm, the ordinary probabilities can be derived without additional operations.

II - REGULAR NETS

Even if they do not have the same expression power as general colored nets, regular nets allow one to model a large class of systems. They have been the starting point for developing important theoretical results, such as reductions and computation of linear place-invariants. They have also allowed to formalize the parametrization which leads to a validation of the system that does not depend on the values of some parameters such as the number of sites, or the number of processes.

2.1 Definition :

A regular net RN = <P, T, C, I^-, I^+> is defined by :

P the set of places,

T the set of transitions,

C the set of object classes : $C = \{C_1, \ldots, C_n\}$, with $C_i \cap C_j = \emptyset$,

I^+ and I^- the input and output matrices defined on P x T, which elements I^+, $I^-(p, t)$ are standard color functions of p (defined below).

The color domains C(p) for a place, and C(t) for a transition, are defined as follows : a color domain is made either of the neutral color, or of a Cartesian product of object classes such that all the elements in the product are distinct.

Definition 2.1 : A normalized RN is an RN in which all places and transitions have the same color domain $\boldsymbol{C} = C_1 \times \ldots \times C_n$.

As every RN can be transformed in a normalized RN without modifying its structural behavior, we will limit our study to that specific class of regular nets. However, the case of non-normalized regular nets and the transformation rule are presented in [Dut 89].

Definition 2.2 : Marking of an object.

A marking is a function $m : P \times \boldsymbol{C} \rightarrow \mathbb{N}$, such that m(p, c) is the number of marks of color c in p. The marking of an object c_i can be defined by the function below :

$m : C_i \rightarrow [P \times C_1 \times \ldots \times C_{i-1} \times C_{i+1} \times \ldots \times C_n \rightarrow \mathbb{N}]$

$m(c_i)(p, c_1, \ldots, c_{i-1}, c_{i+1}, \ldots, c_n) = m(p, c)$ where $c = (c_1, \ldots, c_n)$.

Definition 2.3 : Marking of a color.

The marking of a color c can be defined by the function below :

$$m : \boldsymbol{C} \rightarrow [P \rightarrow \mathbb{N}]$$

$$m(c)(p) = m(p, c)$$

As the three definitions are equivalent, we use the same letter for all the marking functions.

The standard color functions of a regular net are defined from two basic functions :

$\mathbf{X_i} : C_i \times \boldsymbol{C} \rightarrow \mathbb{N}$ such that $X_i(c_i', (c_1, \ldots, c_n)) = Id(c_i, c_i')$
$= (\text{If } c_i = c_i' \text{ then } 1 \text{ else } 0)$.

An arc labeled X_i distinguishes exactly one object in the class C_i. The behavior of this object will be independent of the behavior of the other objects of the class when firing the transition.

$\mathbf{S_i} : C_i \times \boldsymbol{C} \rightarrow \mathbb{N}$ such that $S_i(c_i', (c_1, \ldots, c_n)) = 1$.

An arc labeled S_i means that all the objects in the class C_i play a similar part. It corresponds to a synchronization of all the objects of the class if it labels an input arc. It represents a diffusion to all the objects of the class if it labels an output arc.

Those two basic functions can be combined by :

$\mathbf{a_i.S_i + b_i.X_i} : C_i \times \boldsymbol{C} \rightarrow \mathbb{N}$ such that $a_i.S_i + b_i.X_i(c_i', (c_1, \ldots, c_n))$
$= (\text{If } c_i = c_i' \text{ then } (a_i+b_i) \text{ else } a_i)$.

As a consequence, we have $a_i \geq 0$ and $(a_i+b_i) \geq 0$.

Definition 2.4 : A standard color function of a regular net $< a_1.S_1+b_1.X_1, \ldots, a_n.S_n + b_n.X_n>$, denoted by $\Pi_{i=1,n}< a_i.S_i+b_i.X_i>$, is defined as :

$\Pi_{i=1,n}< a_i.S_i+b_i.X_i> : \boldsymbol{C} \times \boldsymbol{C} \rightarrow \mathbb{N}$,

$\Pi_{i=1,n}< a_i.S_i+b_i.X_i> ((c'_1, \ldots, c'_n), (c_1, \ldots, c_n)) = \Pi_{i=1,n} <a_i.S_i + b_i.X_i>(c'_i, (c_1, \ldots, c_n))$

The initial marking of a regular net, denoted by m_0, must be symmetric, i.e., it must verify the following property :

$$\forall p \in P, \forall c, c' \in C(p), m_0(p, c) = m_0(p, c').$$

2.2 Example :

We consider the following net which models the behavior of a distributed database. This model is derived from a model presented in [Jen 81b], and has also been studied in [Had 87]. A site is made of an active and a passive part.The active part of a site can modify a file, whereas the passive part only takes into account the modifications performed by the other sites. When a site modifies a file, its active part sends a message to all the other sites so that they can take the modification into account, whereas it waits for the acknowledgements (transition T1). The passive part of a site that receives a message modifies its copy of the file (transition T3). If there are several messages for a site, the modifications are done one at a time. Once its copy is modified, a site sends back a message to the one that originated the modification (transition T4). Once all the sites have acknowledged the modification, the waiting site resumes its activity and another site can in turn work on the modified file (transition T2).

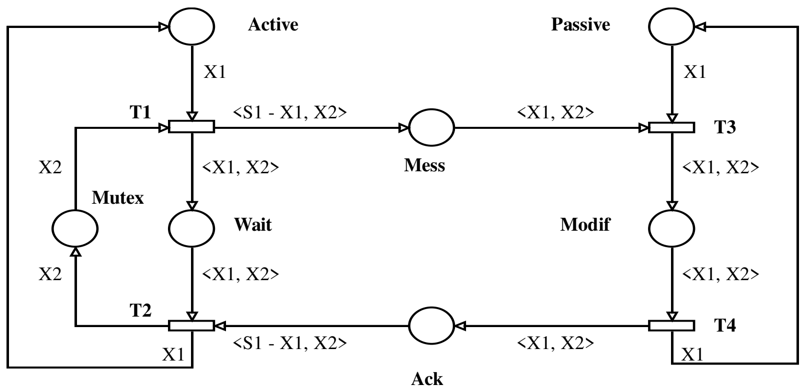

We consider a distributed database with three sites and two files. There are two color classes in the net: the site class SIT = {s1, s2, s3}, and the file class FIL = {f1, f2}. The initial marking of the net is the following : Active = Passive = SIT, and Mutex = FIL.

In the next section, we show that it is possible to develop a reduced reachability graph for regular nets, the symbolic reachbility graph.

III - SYMBOLIC REACHABILITY GRAPH

In this section, we first present some properties of regular nets that are used in the construction of the symbolic reachability graph (SRG). Unlike Jensen's reachability tree [Hub 84] that develops only once equivalent subtrees, the SRG is obtained by grouping states a priori, thus avoiding to develop any reachable state of the unfolded net. The method for constructing the SRG is then developed. We define the representation of a symbolic marking, and we present a symbolic firing rule that can be used directly on the symbolic markings.

3.1 Definitions and properties :

The properties given here result from the definition of the color domains. Also due to the specificity of the color functions, a basic symmetry property of the model is presented.

Definition 3.1 : let s_i a permutation of C_i. A permutation $s = \langle s_1, \ldots, s_n \rangle$ of $(C_1 \times \ldots \times C_n)$ is defined by

$$s(\langle c_1, \ldots, c_n \rangle) = \langle s_1(c_1), \ldots, s_n(c_n) \rangle.$$

The group of the permutation of the type $\langle s_1, \ldots, s_n \rangle$ will be denoted by S.

Definition 3.2 : let m a marking, $s \in S$ a permutation. Then s.m is a marking defined by :

$$\forall\, p \in P, \quad s.m(p, c) = m(p, s(c)).$$

Proposition 3.1 : s.m defines an operation of the group S on the marking set MS, i.e.,

$$\forall\, s, s' \in S, \forall\, m \in MS, \quad (s \circ s').m = s.(s'.m)$$
$$id.m = m.$$

Definition 3.3 : the orbit of m, reg(m), is defined by $reg(m) = \{s.m, s \in S\}$.

The three following corollaries are standard properties of the operation of a group on a set [Lan 77].

Corollary 3.1 : the orbits reg(m) define a partition of MS that induces an equivalence relation **R** :

$$m \mathbf{R} m' \Leftrightarrow reg(m) = reg(m').$$

The equivalence classes of the relation **R** are called symbolic markings, denoted by M.

Corollary 3.2 : Let s a permutation, and M a symbolic marking. Let $f_s : M \to M$, defined by

$$\forall\, m \in M, \quad f_s(m) = s.m.$$

Then f_s is a bijection.

Corollary 3.3 : Let m, m' $\in$ M. $|\{s \in S,\ s.m = m'\}| = |S| / |M|$.

We finally give the basic theorem [Had 87], which is a consequence of the type of the color functions:

Theorem 3.1 : $\forall$ m, m' $\in$ M, $\forall$ t $\in$ T, $\forall$ c $\in$ $(C_1 \text{ x } \ldots \text{ x } C_n)$, $\forall$ s $\in$ S,

$$m[t(c)>m' \Leftrightarrow s.m[t(s(c))>s.m'.$$

3.2 Construction of the SRG :

The construction of the SRG requires that we define the representation of a symbolic marking, and the firing rule that can be applied directly on these symbolic markings. We also present some properties of the SRG. These properties are used to prove that the SRG is relevant for performance evaluation.

3.2.1 Optimal representation of a symbolic marking :

The basic principle for representing a symbolic marking consists in grouping in a subclass all the objects of a class that have the same marking. Therefore, permuting two objects within a subclass will not modify the marking. The identity of the objects in a subclass is then forgotten and only the number of objects is taken into account for each subclass. In this goal, each object class C_i is partitioned in a number of subclasses, $C_i = \{C_{i,1}, \ldots, C_{i,si}\}$, such that all the objects in a subclass have the same marking. The cardinalities of each subclass, which values are in $\mathbb{N}^+$, verify $\sum_{j=1,si} |C_{i,j}| = |C_i|$. The marking of each place is then similar to an ordinary marking where the subclasses are considered as objects. Thus we can define the marking of a subclass, the marking of a product of subclasses and the permutation on subclasses in the same way as they are defined for objects in Definitions 2.2, 2.3 and 3.1.
Moreover, the grouping must be maximal, i.e., two different subclasses must have different markings. The representation of symbolic markings as defined above is unique within a permutation $<s_1, \ldots, s_n>$ of the set of subclasses. However, it is possible to define and to calculate a canonical representation for each symbolic marking by an adequate ordering based on the marking of the products of subclasses [Had 87].

Notice that the decomposition in subclasses is local to a symbolic marking. Thus, the subclass $C_{i,j}$ appearing in a symbolic marking M and the same subclass $C_{i,j}$ appearing in a symbolic marking M' may not have related meanings.

3.2.2 Example :

The following reachable ordinary marking, in which the site s2 has sent a message to the two other sites s1 and s3,

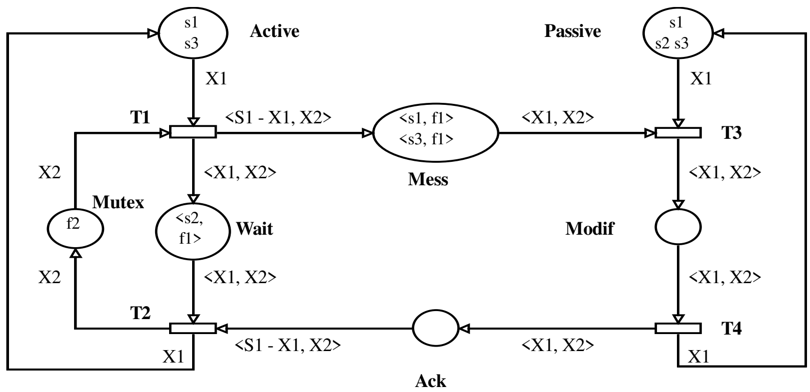

is represented by the symbolic marking :

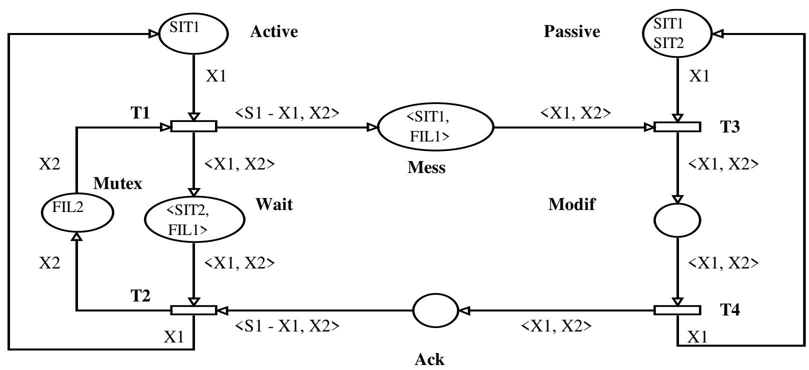

where $|SIT1| = 2$, $|SIT2| = 1$, $|FIL1| = 1$, and $|FIL2| = 1$.
In fact, this symbolic marking represents any ordinary reachable marking in which one site has sent a message to both others.

3.2.3 Symbolic firing rule :

In order to build the SRG, we first define a symbolic firing rule on the symbolic markings which must be sound i.e., an ordinary marking enables a colored transition if and only if its symbolic marking enables an equivalent symbolic firing and the ordinary marking obtained by the firing belongs to the symbolic marking obtained by the symbolic firing.

The first effect of the symbolic firing will be to split each instantiated subclass in two subclasses, one with the object instantiated in the underlying firing and the other with the remaining objects of the subclass. Thus we formally define this splitting.

Notation : In M, we will denote the partition of the class C_i in $\{C_{i,1}, \ldots, C_{i,si}\}$ by $\mathbb{C}_i = \{C_{i,1}, \ldots, C_{i,si}\}$. If confusion may arise, $|C_{i,j}|_M$ will denote the cardinality of $C_{i,j}$ in M.

Definition 3.4 : Let M a symbolic marking. Then $M[C_{1,u1}, \ldots, C_{n,un}]$ is a symbolic marking defined by :

- If $|C_{i,ui}| > 1$ then

  $\mathbb{C}_i = \{C_{i,1}, \ldots, C_{i,si}, C_{i,si+1}\}$ with

  $|C_{i,si+1}| = |C_{i,ui}|_{\mathbb{M}} - 1$, $|C_{i,ui}| = 1$, $|C_{i,j}| = |C_{i,j}|_M$ for any $j \neq u_i$ and $j \neq s_i+1$

  Else the partition of C_i is unchanged.

- The marking of the old subclasses is unchanged and the marking of the new subclass $C_{i,si+1}$ is the same as the one of $C_{i,ui}$.

Notice that in $M(C_{1,u1}, \ldots, C_{n,un})$ the grouping is not always maximal and that even if the grouping is maximal the representation is not always canonical. But it does not matter since this symbolic marking is just an intermediate marking and it will not appear in the SRG.

The instantiation of a transition in a symbolic firing will be made by choosing a subclass per class instead of an object per class in an ordinary firing. Thus we must define the value of the colored functions for subclasses. This definition is the same as the one for the objects. In the case where an instantiated subclass contains more than one object, the symbolic firing should be enabled for the object instantiated in the underlying firing and for the other objects of the subclass which are not instantiated. Thus the definition should be different but since we apply our definitions on split markings, this case never appears.

Definition 3.5 : Let M a symbolic marking. Then :
$<a_i.S_i + b_i.X_i>$ is a function from $\mathbb{C}_i \times \prod_{j=1,n} \mathbb{C}_j \to \mathbb{N}$ and
$<a_i.S_i + b_i.X_i> (C_{i,vi} , (C_{1,u1}, \ldots, C_{n,un})) =$ If $u_i \neq v_i$ then a_i else (a_i+b_i)

Definition 3.6 : Let M a symbolic marking. Then :
$\prod_{j=1,n} <a_i.S_i + b_iX_i>$ is a function from $\prod_{i=1,n} \mathbb{C}_i \times \prod_{i=1,n} \mathbb{C}_i$ ----> $\mathbb{N}$ and
$\prod_{j=1,n} <a_i.S_i + b_iX_i> ((C_{1,v1}, \ldots, C_{n,vn}) , (C_{1,u1}, \ldots, C_{n,un})) =$
$\prod_{j=1,n} (<a_i.S_i + b_iX_i> (C_{i,vi} , (C_{1,u1}, \ldots, C_{n,un})))$

Let M_j a symbolic marking, t a transition, and $(C_{1,u1}, \ldots, C_{n,un})$ a tuple of subclasses such that $C_{i,ui} \in C_i$. $C_{i,ui}$ is the distinguished subclass of C_i for the firing of t.

Definition 3.7 : t is enabled from M for $(C_{1,u1}, \ldots, C_{n,un})$ iff :
$\forall\, p \in P$,
$M[C_{1,u1}, \ldots, C_{n,un}] (p, C_{1,v1}, \ldots, C_{n,vn}) \geq I^-(p, t) ((C_{1,v1}, \ldots, C_{n,vn}), (C_{1,u1}, \ldots, C_{n,un}))$

The symbolic marking M' obtained by firing $t(C_{1,u1}, \ldots, C_{n,un})$ is calculated with the three following steps :

Step 1 : We apply the incidence functions on $M[C_{1,u1}, \ldots, C_{n,un}]$ giving a new symbolic marking M_1
$\forall\, p \in P$ (we denote $I = I^- - I^+$),

$$M_1(p, C_{1,v1}, \ldots, C_{n,vn}) = M[C_{1,u1}, \ldots, C_{n,un}] (p, C_{1,v1}, \ldots, C_{n,vn}) + I(p, t) ((C_{1,v1}, \ldots, C_{n,vn}), (C_{1,u1}, \ldots, C_{n,un}))$$

Step 2 : as the grouping of states may not be maximal in M_1, it consists in grouping all the subclasses that have same markings giving a new marking M_2. In fact, only the splitted subclasses may be equivalent to previously existing ones.

Step 3 : calculation for M_2 of the canonical representative marking M'.

3.2.4 Example :
We apply the technique to our example. From the initial marking, we show the possible transition firings in the reachability graph (RG), and we represent the same step in the SRG. We will consider the marking of the different places in the following order :

(Active) (Passive) (Mess) (Mutex) (Wait) (Modif) (Ack).

The possible transition firings in the RG are the following ones :
T1(s1,f1),T1(s1,f2),T1(s2,f1),T1(s2,f2),T1(s3,f1),T1(s3,f2)

The SRG corresponding to the same step is the following one : At the beginning, places Active and Passive are marked with a subclass SIT1 of SIT, which is equal to SIT, and Mutex is marked with FIL1 which is a subclass of FIL, FIL1 = FIL. The firing of transition T1 for any couple of objects in SIT1 x FIL1 leads to a unique symbolic marking. This marking is obtained by splitting SIT1 in two subclasses SIT1 and SIT2, and FIL1 in two subclasses FIL1 and FIL2. The splitting is necessary because the couple of objects for which T1 has been fired now have a marking different from that of the other objects of their former subclasses.

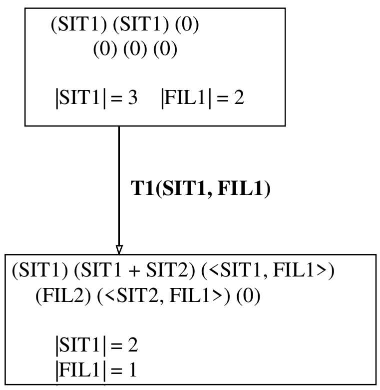

The complete SRG in case of two sites and two files would have the following structure. Each node of the graph can be interpreted without knowing the identity of the file or the site that marks any place.

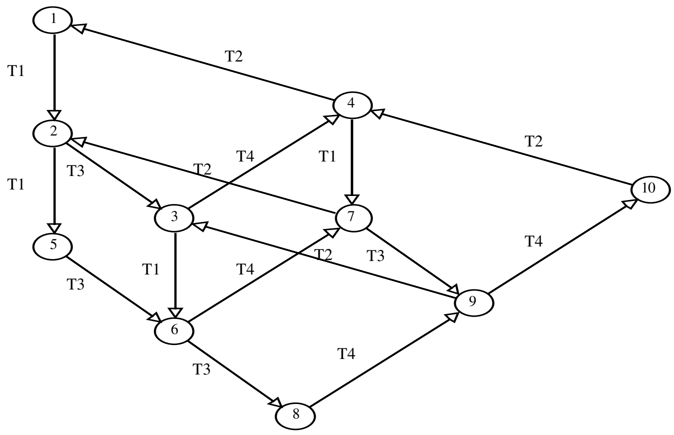

1 : no file being modified 2 : 1 file being modified, a message has been sent 3 : 1 file being modified, the second site is performing the modification 4 : 1 file being modified, the second site has sent an acknowledgement 5 : 2 files being modified, 2 messages have been sent 6 : 2 files being modified, one site is performing the modification 7 : 2 files being modified, one site has sent an acknowledgement 8 : 2 files being modified, both sites are performing a modification 9 : 2 files being modified, one site has sent an acknowledgement, the other is performing a modification 10 : 2 files being modified, 2 acknowledgements have been sent

3.2.5 Construction :

The algorithm for constructing the SRG is different from the construction of the ordinary RG only by the firing rule and the labels of the arcs that are made of the transitions and the tuples of firing subclasses, whereas in the RG we have the transition and the tuple of firing objects. Notice that the initial symbolic marking only contains the initial ordinary marking because of the symmetry of the initial marking.

3.2.6 Some properties of the SRG :

Many properties have been proved on the SRG, such as quasi-liveness, and the possibility of finding a home state. However, we will only present some properties that are useful for the derivation of a stochastic model. All the proofs of the propositions given here are in [Had 87].

Proposition 3.2 : The reachability property is equivalent for the ordinary and the symbolic markings :

$$m \in RG \Leftrightarrow reg(m) \in SRG.$$

Proposition 3.3 : There is an exact relation between the arcs of the RG and the SRG.
Let M = reg(m) and M' two symbolic markings of the SRG. Let $A_{m,M'}$ the set of arcs going out of m to any m' ∈ M', $A_{M,M'}$ the set of symbolic arcs leading from M to M'. Then there is an application mapping $A_{m,M'}$ on $A_{M,M'}$ such that the reciprocal image of a symbolic arc labelled by $t(C_{1,u1}, \ldots, C_{n,un})$ is a set of arcs labelled by some $t(c_1, \ldots, c_n)$.
The cardinality of this set is $\Pi_{j=1}^{n} | C_{j,uj} |$.

IV - REGULAR STOCHASTIC PETRI NETS

Stochastic Petri nets (SPN's) [Mol 81, Flo 85] are a tool well adapted to the modeling and the performance evaluation of distributed computer systems. In an SPN, a firing delay is associated with each transition. Firing delays are instances of random variables that have a negative exponential probability distribution. The probability that two timed transitions sample exactly the same delay time is zero, so that priority zero transitions are assumed to fire one at a time. The selection of the transition instance to fire among the set of the enabled ones follows a "race" policy [Ajm 85] (the transition that has drawn the least delay is the one that fires). Under these assumptions, the reachability graph is isomorphic to a Markov chain [Mol 81].

Regular Stochastic Petri Nets :

The association of a timing semantics to colored nets, and the extension of some notions such as conflict or confusion, have been clearly defined in [Chi 88]. However, it has also been shown that it was almost impossible to avoid the construction of the reachability graph of the unfolded net when trying to analyze the model. This is the reason why we decided to build our stochastic model only on a subclass of colored nets

As regular nets are a class of colored nets, the rules defined in [Chi 88] apply to Regular Stochastic Petri Nets (RSPN's). A RSPN is then defined as a couple (RN, λ), where RN defines a Regular Net, and λ is a weight function $\lambda : T \times \boldsymbol{C} \rightarrow \mathbb{R}^+$ associating a positive real number with each transition. The value of $\lambda_{t(c)}$ is the mean delay between the moment when (t, c) is enabled and the moment when it actually fires.

<u>Restrictions on the model</u> :

We will neither consider the case of immediate transitions, nor marking-dependent weights, but these extensions have been developed in [Dut 89]. We will further impose that all the color instances of a transition have the same firing rate, so that λ can be redefined as a function $T \rightarrow \mathbb{R}^+$. As all the objects in a class behave the same way, it is not a heavy restriction to believe that they have the same timing constraints.

V - COMPUTATION OF THE STATE PROBABILITIES

All the performance measures of a specific system can be obtained from the computation of the steady-state probabilities. However, complex systems can generate large reachability graphs, thus making the probability computation impossible. In this section, we present an algorithm for computing the steady state probabilities without developping the whole reachability graph. This algorithm is based on a stochastic aggregation of states that is directly derived from the symbolic reachability graph. We first recall the algorithm used in a general case, then we present an algorithm that can be used in the case of regular nets.

5. 1 Usual algorithm

The probability vector P, where p(i) is the steady probability of state i, is usually obtained from the following algorithm :

(1) Unfolding of the net.

(2) Construction of the reachability graph of the unfolded net (the unicity of P is ensured iff there is only one absorbing strongly connected component in the graph), and valuation of the arcs with the rate of the associated transition.

(3) Computation of the square matrix A defined by

the dimension of A is the number N of reachable markings,

if $i \neq j$, $a_{i,j} = \sum_{t(c) \text{ s.t. the firing of } t(c) \text{ leads from } i \text{ to } j} \lambda_{t(c)}$,

$a_{i,i} = - \sum_{i \neq j} a_{i,j}$

(4) Resolution of the system

$$P.A = 0 \text{ and } \| P \| = \sum_{i=1,..,N} p(i) = 1$$

This algorithm is very general and does not use the particular structure of the Markov chain derived from a regular net. We now present an algorithm which is valid only in the case of regular nets.

5.2 Improved algorithm

The algorithm presented in this section uses some properties of regular nets in order to reduce the complexity of the resolution. In the sequel, a symbolic marking of the SRG will be denoted M_I, I = 1, …, N.

(1) Construction of the SRG.

The number of ordinary markings within a symbolic marking is computed during this step with the following formula :

$|M| = 1 / K \,.\, \prod_{i=1,..N} (|C_i|! / (|C_{i1}|! \times \ldots \times |C_{i,ni}|!))$

where n_i is the number of subclasses of C_i in the symbolic marking M, and K is the cardinality of the set of the permutations that leave M unchanged when applied to the subclasses of object classes :

$K = | \{<s_1, \ldots, s_n> \in S, <s_1, \ldots, s_n>.M = M\} |$.

(2) Computation of the square matrix A^* defined by :

the dimension of A^* is the number N' of reachable symbolic markings,

$$\text{if } I \neq J,\ a_{I,J} = \sum_{t(D_1,\ldots,D_n) \in A_{M,M'}} \lambda_t \,.\, |D_1| \,.\, \ldots\, . \,|D_n|$$

$$a_{I,I} = - \sum_{I \neq J} a_{I,J}$$

(3) Resolution of the system

$$Q^*.A^* = 0 \text{ and } \| Q^* \| = \sum_{I=1,..,N'} q^*(M_I) = 1$$

(4) Computation of the state probability vector Q defined by $q(i) = q^*(M_I) / | M_I |$ for any marking i belonging to the symbolic marking M_I.

The fourth step is optional and depends on the will of the user to get either the probabilities of the classes of states, or the probabilities of the individual states.

N' is generally much less than N, and as the whole reachability graph is never built, the algorithm provides a real improvement in the place necessary for storing the data. Moreover, the improvement is still greater when looking at the complexity of the computation, as the exact resolution of the matricial system has a complexity $O(n^3)$. As for approximate methods, the complexity depends both on the number of states and the number of arcs, our algorithm is still an improvement. It is therefore possible to obtain exact results for large reachability graphs. The results obtained are also often more significant for the user, because the symbolic markings can be easily interpreted. In the next section, we will prove the correctness of the improved algorithm by showing that vectors P and Q are equal.

VI - PROOF OF THE ALGORITHM

In this section, we detail the different steps that make the transformation of the algorithm correct. The process associated with the unfolded net is supposed to be ergodic. We first give a sketch of the method we will use to prove the correctness of our improved algorithm. The proof organizes into three steps.

(1) All the markings within a symbolic marking have the same probability.

(2) We prove that the Markovian lumping condition is fulfilled, and the linear system corresponding to the ordinary matricial equation can be reduced to a system with fewer variables, but which is no longer probabilistic (the solution is not a probability vector). However, the reduced system can be transformed into a probabilistic system by a simple change of variables.

(3) This probabilistic system is exactly the one computed by our algorithm.

Notation : According to the context, the index of a symbolic marking will be denoted with a capital or a lower case letter. An ordinary marking of the RG will be denoted (i, j) where M_I is the associated symbolic marking, and j is an order number within the symbolic marking. Let s S a permutation. Then s.(i, j) will be denoted (i, s.j), where s.j is the order number of s.(i, j) within M_I.

6.1 Equiprobability of the markings :

In this section, we show that as the linear system associated to the Markov chain has a unique solution, the solution is a probability vector such that all the markings in a symbolic marking have the same probability.

We will denote by Eq(i, k) the equation of the system Eq defined by P.A = 0, corresponding to the ordinary marking number k of the symbolic marking number i:

Eq(i,k) : $\Sigma_{J=1,..N} \Sigma_{(j,q) \in M_I} p_{(j, q)} \cdot a_{(j, q)(i, k)} = 0$

Proposition 6.1 : The transition rate between a state (j, q) and a state (i, k) is the same as the transition rate between (j, s.q) and (i, s.k).

$$\forall (j, q), \forall (i, k), \forall s \in S, \quad a_{(j, q)(i, k)} = a_{(j, s.q)(i, s.k)}.$$

Proof : Two cases must be considered.

case 1 : $(j, q) \neq (i, k)$: this is a straightforward consequence of the basic Theorem 3.1.

case 2 : $(j, q) = (i, k)$: $a_{(i, s.k)(i, s.k)}$ is written as

$$a_{(i,\ s.k)(i,\ s.k)} = -\left[\sum_{\substack{J=1 \\ J\neq I}}^{N} \sum_{(j,\ q)\in M_J} a_{(i,\ s.k)(j,\ q)} + \sum_{\substack{(i,\ q)\in M_I \\ q\neq s.k}} a_{(i,\ s.k)(i,\ q)}\right]$$

From Corollary 3.2, we know that a permutation s defines a bijection on a symbolic marking M_I. Therefore, it is equivalent to sum on (i, q) or (i, s.q).

$$a_{(i,\ s.k)(i,\ s.k)} = -\left[\sum_{\substack{J=1 \\ J\neq I}}^{N} \sum_{(j,\ q)\in M_J} a_{(i,\ s.k)(j,\ s.q)} + \sum_{\substack{(i,\ q)\in M_I \\ q\neq k}} a_{(i,\ s.k)(i,\ s.q)}\right]$$

Applying case 1, we obtain :

$$a_{(i,\ s.k)(i,\ s.k)} = -\left[\sum_{\substack{J=1 \\ J\neq I}}^{N} \sum_{(j,\ q)\in M_J} a_{(i,\ k)(j,\ q)} + \sum_{\substack{(i,\ q)\in M_I \\ q\neq k}} a_{(i,\ k)(i,\ q)}\right] = a_{(i,\ k)(i,\ k)} \quad \underline{\Delta\Delta\Delta}$$

<u>Proposition 6.2</u> : $\{(p_{(i,\ s.k)})_{(i,\ k)}\}$ is a solution of the same equation as $\{(p_{(i,\ q)})_{(i,\ k)}\}$.

$$\forall\ (i,\ k),\quad \sum_{J=1}^{N} \sum_{(j,\ q)\in M_J} p_{(j,\ s.q)} \cdot a_{(j,\ q)(i,\ k)} = 0.$$

<u>Proof</u> : we write Eq(i, s.k) :

$$\sum_{J=1}^{N} \sum_{(j,\ q)\in M_J} p_{(j,\ q)} \cdot a_{(j,\ q)(i,\ s.k)} = 0.$$

and applying Corollary 3.2, we change the index (j, q) by (j, s.q)

$$\sum_{J=1}^{N} \sum_{(j,\ q)\in M_J} p_{(j,\ s.q)} \cdot a_{(j,\ s.q)(i,\ s.k)} = 0.$$

Applying Proposition 6.1, we get

$$\sum_{J=1}^{N} \sum_{(j,\ q)\in M_J} p_{(j,s.q)} \cdot a_{(j,\ q)(i,\ k)} = 0.$$

As we assume that the system is ergodic, the solution is unique. So $p_{(i,\ k)} = p_{(i,\ s.k)}$, and all the markings in a symbolic marking have the same probability.

However, we are going to show that in case the system has multiple solutions, there is always one such that all the markings in a symbolic marking have the same probability, and this solution is computed by our algorithm. If the reader is only interested in ergodic systems, he can go directly to Section 6.2.

Proposition 6.3 : $\left\{ \left(\frac{1}{|M_I|} \cdot \sum_{(i,\ k') \in M_I} p_{(i,\ k')} \right)_{(i,\ k)} \right\}$ is a solution of Eq, i.e.,

Proof : we apply Proposition 6.2, and we sum on all the possible permutations.

$$\sum_{s \in S} \sum_{J=1}^{N} \sum_{(j,\ q) \in M_J} p_{(j, s.q)} \cdot a_{(j,\ q)(i,\ k)} = 0.$$

$$\sum_{J=1}^{N} \sum_{(j,\ q) \in M_J} \left(\sum_{s \in S} p_{(j,\ s.q)} \right) \cdot a_{(j,\ q)(i,\ k)} = 0.$$

The application of Corollary 3.3 gives :

$$\sum_{J=1}^{N} \sum_{(j,\ q) \in M_J} \left(\frac{|S|}{|M_J|} \cdot \sum_{(j,\ q') \in M_J} p_{(j,\ q')} \right) \cdot a_{(j,\ q)(i,\ k)} = 0.$$

Dividing by |S| we get :

$$\sum_{J=1}^{N} \sum_{(j,\ q) \in M_J} \left(\frac{1}{|M_J|} \cdot \sum_{(j,\ q') \in M_J} p_{(j,\ q')} \right) \cdot a_{(j,\ q)(i,\ k)} = 0.$$

6.2 Reduction of the linear system :

We prove here that the steady-state probability of being in a symbolic marking can be calculated from a system with a reduced number of variables.

<u>Proposition 6.4</u> : The transition rate out of a symbolic marking M_J to an ordinary marking in M_I has the same value for every marking in M_I. Conversely, the transition rate out of an ordinary marking in M_I to a symbolic marking M_J has the same value for every marking in M_I.

$\forall$ J, $\forall$ (i, k), $\forall$ s,

$$\sum_{(j,\, q) \in M_J} a_{(j,\, q)(i,\, k)} = \sum_{(j,\, q) \in M_J} a_{(j,\, q)(i,\, s.k)}$$

$$\sum_{(j,\, q) \in M_J} a_{(i,\, k)(j,\, q)} = \sum_{(j,\, q) \in M_J} a_{(i,\, s.k)(j,\, q)}$$

Note that the second equation of Proposition 2 is equivalent to the strong Markovian lumping condition [Kem 60], which ensures that the aggregation of states will preserve the Markovian property of the process.

<u>Proof</u> : The result is obtained by changing the indexes of the sums and applying Proposition 6.1.

ΔΔΔ

Let :

$$\overline{P}_I = \frac{1}{|M_I|} \cdot \sum_{(i,k) \in M_I} p_{(i,k)}$$

be the probability of any marking in M_I for the equiprobable solution, and :

$$\hat{p}_I = |M_I| \cdot \overline{p}_I$$

the probability of being in the symbolic marking M_I. We define also the following equation :

$$\overline{Eq}(i,\, k) \ : \ \sum_{J=1}^{N} \sum_{(j,\, q) \in M_J} \overline{p}_J \cdot a_{(j,\, q)(i,\, k)} = 0.$$

Proposition 6.5 : The system $\overline{Eq}$ is such that the equations associated with all the markings in a symbolic marking are the same, i.e.,

$$\forall\ (i, k),\quad \overline{Eq}(i, k) = \overline{Eq}(i, s.k)$$

Proof : $\overline{Eq}(i, k)$: $\sum_{J=1}^{N} \sum_{(j, q)\in M_J} \overline{p}_J \cdot a_{(j, q)(i, k)} = 0.$

$$\sum_{J=1}^{N} \overline{p}_J \cdot \sum_{(j, q)\in M_J} a_{(j, q)(i,k)} = 0.$$

Applying Proposition 6.4, $\sum_{J=1}^{N} \overline{p}_J \cdot \sum_{(j, q)\in M_J} a_{(j, q)(i, s.k)} = 0.$

As a consequence, we can consider only one equation by symbolic marking. The system we have to solve is then :

$$\begin{cases} \sum_{J=1}^{N} \overline{p}_J \cdot \sum_{(j, q)\in M_J} a_{(j, q)(i, k)} = 0 \\ \sum_{I=1}^{N} |M_I| \cdot \overline{p}_I = 1. \end{cases}$$

Let $\overline{a}_{I, J} = |M_J| \cdot \sum_{(i, k)\in M_I} a_{(i, k)(j, q)}$

Proposition 6.6 : $\sum_{J=1}^{N} \overline{a}_{I, J} = 0.$

Proof : $\sum_{J=1}^{N} \overline{a}_{I, J} = \sum_{j=1}^{N} |M_J| \cdot \sum_{(i, k)\in M_I} a_{(i, k)(j, q)}$

Applying Proposition 6.4,

$$= \sum_{J=1}^{N} \sum_{(j, q)\in M_J} \sum_{(i, k)\in M_I} a_{(i, k)(j, q)}$$

$$= \sum_{(i,\ k)\in M_I} \sum_{J=1}^{N} \sum_{(j,\ q)\in M_J} a_{(i,\ k)(j,\ q)} = 0$$

Multiplying the first equation by $|M_I|$, we can transform our system in

$$\begin{cases} \sum_{J=1}^{N} \overline{p}_J \,.\, \overline{a}_{J,\ I} = 0 \\ \sum_{I=1}^{N} |M_I| \,.\, \overline{p}_I = 1 . \end{cases}$$

Introducing the new notation $\hat{a}_{I,\ J} = \dfrac{\overline{a}_{I,\ J}}{|M_I|}$, we obtain the final system :

$$\begin{cases} \sum_{J=1}^{N} \hat{p}_J \,.\, \hat{a}_{J,\ I} = 0 \\ \sum_{I=1}^{N} \hat{p}_I = 1. \end{cases}$$

It can be verified that the system we have to solve is a stochastic system, and the usual techniques for calculating the steady state probabilities can then be used. The same system would have been obtained using Markovian lumping techniques. And in that case too, it would have been necessary to develop an additional demonstration to prove that all the markings within a symbolic marking have the same probability.

Applied to our example, the above technique allows us to transform a system with 487 variables for the ordinary markings in a system with 46 variables for the symbolic markings. The gain considerably increases with the cardinalities of the object classes. Moreover, the probabilities of the symbolic markings are often more significant for the modeler than the probabilities of the ordinary markings, making it all the more useful to compute directly the values for the lumped states.

In the next part, we will show that the coefficients $\hat{a}_{I,\ J}$ can be derived directly from the SRG,

and that it is possible to calculate the number of ordinary markings in a symbolic marking, thus allowing us to derive the ordinary probabilities from the probabilities of the symbolic markings.

6.3 Computation of the coefficients of the reduced linear system :

The properties we are going to use to compute the coefficients of the linear system are directly linked to Proposition 3.3.

Let φ the function mapping $A_{m,M'}$, the set of arcs leading from an ordinary marking m to any marking in the symbolic marking M', on $A_{M,M'}$, the set of symbolic arcs leading from M, with $m \in M$, to M'. Then we have the following properties :

(1) φ is surjective,

(2) the same transition labels an arc and its image by φ :

$$\varphi[t'(c_1, \ldots, c_n)] = t(D_1, \ldots, D_n) \Rightarrow t' = t,$$

(3) $| \varphi^{-1} [t(D_1, \ldots, D_n)] | = \Pi_{i=1}^{n} | D_i |$

We denote λ_t the rate associated with a transition t, t(A) the transition labeling the symbolic arc A, and D_i^A the subclass of C_i instantiating t(A). So, if t(a) is the transition labeling an ordinary arc a, then t(a) = t(A) for any $a \in \varphi^{-1}(A)$.

Then the coefficients of the reduced linear system can be directly computed with the following formula :

<u>Proposition 6.7</u> : $$\hat{a}_{I,J} = \sum_{A \in \mathcal{A}_{M_I,M_J}} \lambda_{t(A)} \cdot |D_1^A| \cdot \ldots \cdot |D_n^A|$$

<u>Proof</u> : From the definition of $\hat{a}_{I,J}$, we have for any $(i, k) \in M_I$:

$$\hat{a}_{I,J} = \sum_{(j,q) \in M_J} a_{(i,k)(j,q)} = \sum_{a \text{ leading from } (i,k) \text{ to } (j,q) \in M_J} \lambda_{t(a)}$$

The set of arcs leading from (i, k) to $(j, q) \in M_J$ can be partitioned according to their image by φ. Thus we get :

$$\hat{a}_{I,J} = \sum_{A \in \mathcal{A}_{M_I,M_J}} \sum_{a \in \varphi^{-1}(A)} \lambda_{t(a)}$$

As φ preserves the transition names, this can be also written

$$\hat{a}_{I,J} = \sum_{A \in \mathcal{A}_{M_I,M_J}} \sum_{a \in \varphi^{-1}(A)} \lambda_{t(A)}$$

which is equal to

$$\hat{a}_{I,\,J} = \sum_{A \in \mathcal{A}_{M_I, M_J}} \left|\varphi^{-1}(A)\right| \cdot \lambda_{t(A)}$$

Applying Proposition 3.3, we obtain

$$\hat{a}_{I,\,J} = \sum_{A \in \mathcal{A}_{M_I, M_J}} \prod_{i=1}^{n} \left|D_i^A\right| \cdot \lambda_{t(A)} \qquad \Delta\Delta\Delta.$$

Those values can therefore be calculated directly from the SRG, by giving to an arc a weight depending on the cardinalities of the subclasses of its label. The values of the coefficients for I = J are derived of the nullity of the sum.

VII - CONCLUSION

The symmetry properties of regular nets have been used to prove many structural results. In this paper, we have shown that the symmetries can be used also in the case of a quantitative analysis.
We have proved that all the states in a symbolic marking have the same probability, and therefore, that the probabilities of all the ordinary markings can be derived from the resolution of a system which size depends only on the number of symbolic markings. We have given an algorithm for calculating the coefficients of the reduced system directly from the graph of symbolic markings.

The advantages of our method are twofold . On the one hand, the user can choose if he wants the probabilities of the ordinary markings, or if he is only interested in the probabilities of the symbolic markings which are often more significant. On the other hand, the reduction of the system will bring a dramatic improvement in the memory space and the CPU time required to solve large models, and will increase the class of models that can be analytically solved.

Our research directions will be to develop a software tool that will automatically construct the graph of symbolic markings. This tool could be later interfaced with powerful softwares for stochastic Petri nets, such as GreatSPN [Chi 87], or RDPS [Flo 86]. The introduction of immediate transitions and marking-dependent weights has been presented in [Dut 89]. We now intend to analyze less restricted classes of nets that still have symmetry properties. These extended nets include successor functions [Had 88], a non-symmetric initial marking, or color domains with several occurrences of the same object class.

References

[Ajm 85] M. Ajmone Marsan, G. Balbo, A. Bobbio, G. Chiola, G. Conte, and A. Cumani, "On Petri Nets with Stochastic Timing", in proc. International Workshop on Timed Petri Nets, pp 80-87, IEEE-CS Press, Torino, Italy (July 1985).

[Ajm 87] M. Ajmone Marsan, G. Balbo, G. Chiola, and G. Conte, "Generalized Stochastic Petri Nets Revisited : Random Switches and Priorities", in proc. International Workshop on Petri Nets and Performance Models, pp 44-53, IEEE-CS Press, Madison, WI, USA (August 1987).

[Chi 87] G. Chiola, "A Graphical Petri Net Tool for Performance Analysis", 3rd International Workshop on Modeling Techniques and Performance Evaluation, AFCET, Paris (March 1987).

[Chi 88] G. Chiola, G. Bruno, and T. Demaria, "Introducing a Color Formalism into Generalized Stochastic Petri Nets", in proc. 9th European Workshop on Application and Theory of Petri Nets, pp 202-215, Venezia, Italy (June 1988).

[Dut 89] C. Dutheillet, S. Haddad, “Aggregation of States in Colored Stochastic Petri Nets. Application to a Multiprocessor Architecture”, in proc. PNPM 89, pp 40-49, IEEE-CS Press, Kyoto, Japan (December 1989).

[Flo 85] G. Florin, S. Natkin, "Les Réseaux de Petri Stochastiques", AFCET TSI, vol. 4, no. 1, pp 143-160 (January 1985).

[Flo 86] G. Florin, "RDPS, a Software Package for the Validation and the Evaluation of Dependable Computer Systems", 3rd Proc. IFAC SAFECOMP Workshop, Sarlat, France (October 1986).

[Had 87] S. Haddad, "Une Catégorie Régulière de Réseaux de Petri de Haut Niveau : Définition, Propriétés et Réductions", Thèse de Doctorat, RR87/197, Laboratoire MASI, Université Paris VI, Paris, France (October 1987).

[Had 88] S. Haddad and J.M. Couvreur, "Towards a General and Powerful Computation of Flows for Parametrized Coloured Nets", in proc. 9th European Workshop on Application and Theory of Petri Nets, pp 202-215, Venezia, Italy (June 1988).

[Hub 84] P. Huber, A.M. Jensen, L.O. Jepsen, and K. Jensen, "Towards Reachability Trees for High Level Petri Nets", 5th European Workshop on Application and Theory of Petri Nets, Aarhus, Denmark (1984).

[Jen 81a] K. Jensen, "Coloured Petri Nets and the Invariant Method", Theoretical Computer Science 14, pp 317-336 (1981).

[Jen 81b] K. Jensen, “How to Find Invariants for Coloured Petri Nets”, Mathematical Foundations of Computer Science, LNCS 118, Springer - Verlag (1981).

[Kem 60] J. G. Kemeny and J. L. Snell, "Finite Markov Chains", Van Nostrand, Princeton, NJ (1960).

[Lan 77] S. Lang, “Algebra”, Addison - Wesley (1977).

[Lin 87] Chuang Lin and D. Marinescu,
"On Stochastic High Level Petri Nets", in proc. International Workshop on Petri Nets and Performance Models, pp 44-53, IEEE-CS Press, Madison, WI, USA (August 1987).

[Mol 81] M. K. Molloy,
"On the Integration of Delay and Throughput Measures in Distributed in Distributed Processing Models", Ph.D. Dissertation, University of California, Los Angeles, CA, USA (September 1981).

[Zen 85] A. Zenie,
"Colored Stochastic Petri Nets", in proc. International Workshop on Timed Petri Nets, pp 262-271, IEEE-CS Press, Torino, Italy (July 1985).