

A NET LEVEL PERFORMANCE ANALYSIS OF STOCHASTIC PETRI NETS

W. HENDERSON¹, D. LUCIC¹ AND P. G. TAYLOR²

(Received October 1988; revised January 1989)

Abstract

Stochastic Petri Nets are used extensively to find performance measures for communication protocols. This paper illustrates how equilibrium distributions for the markings of a wide class of nets can be found directly without the need to generate a large state space and then resort to equilibrium balance equations.

1. Introduction

A Petri net is a graph model useful for analysing systems which exhibit concurrent and conflicting behaviour. A good introduction to the properties and applications of Petri net theory is given in Peterson [11]. Of particular interest to researchers in communications is the use of Petri nets in the specification and verification of protocols (see Symons [13], Billington, Wheeler and Wilbur-Han [3] and Billington [2]).

By allowing the firing times of transitions to be non-negative random variables a Petri net can be modified to a Stochastic Petri Net (SPN). The purpose of an SPN model is to enable us to discover additional information, such as performance measures, about the protocol under consideration (see Marsan, Balbo and Conti [9], Florin and Natkin [5] and Marsan, Chiola and Fumagalli [10]).

The major problem in finding such performance measures is the need to work with the equilibrium equations based on the reachability graph of the SPN. The size of the reachability graph increases exponentially with both the number of places and the number of tokens in the initial marking and

¹Applied Mathematics Department, University of Adelaide, South Australia 5001.

²Mathematics Department, University of Western Australia, Australia.

© Copyright Australian Mathematical Society 1989, Serial-fee code 0334-2700/89

consequently, even for simple nets, these equations are an extremely difficult medium through which to work.

A similar problem in networks of queues was partially overcome by finding product form solutions for a reasonably wide ranging class of networks. Much of the early work in queueing networks was not generally applicable to SPNs because only single arrivals and single services were permitted, whereas most SPNs have groups of tokens moving through the net at the same time. Some very recent work by Henderson, Pearce, Taylor and Van Dijk [6] and Henderson and Taylor [7] has found that simple equilibrium solutions can still be found for certain networks in which batch arrivals and batch services can occur. In this paper we will show how the ideas used in Henderson and Taylor [7] can be successfully applied to classes of SPNs and in particular how some nets with conflict sets, probabilistic routing and marking dependent firing rates yield a tractable equilibrium solution. A major advantage of adopting this approach is that the analysis is not based on the reachability graph created from the SPN. The technique involves finding the equilibrium distribution for a Markov chain with states comprising the transitions in the SPN. Using this information we can directly find the equilibrium distribution for the markings of the net.

Balbo, Bruell and Chanta [1] have applied the knowledge of product form solutions to the analysis of SPNs, by replacing subnets which possess a product form solution with a single “equivalent server” transition, thereby reducing computational complexity. Since the SPNs presented in this paper possess a closed norm solution it should be possible to extend Balbo et al’s [1] work to cover a wider class of SPNs.

In Section 2 we introduce the necessary notation. Sections 3 and 4 give our main results and examples are given in Section 5.

2. Description of the model

Consider a live and bounded SPN with a finite set $\mathcal{P} = \{1, \dots, P\}$ of places and a finite set \mathcal{T} of transitions. The Markov process representing the SPN has markings $\mathbf{m} \in Z^P$, to denote the state when there are $m(i)$ tokens at place i . When the state is \mathbf{m} , transition $t \in \mathcal{T}$ has a state dependent firing rate given by $q(\mathbf{m}, t)$, and when transition t fires it moves tokens from an input bag $\mathbf{I}(t) \in Z^P$ to one of a set of possible output bags $\mathbf{O}_j(t) \in Z^P$ with probability $p(\mathbf{I}(t), t, \mathbf{O}_j(t))$ where $\sum_j p(\mathbf{I}(t), t, \mathbf{O}_j(t)) = 1$. Thus the marking of the net changes from \mathbf{m} to $\mathbf{m} - \mathbf{I}(t) + \mathbf{O}_j(t)$ with probability $p(\mathbf{I}(t), t, \mathbf{O}_j(t))$.

It is worth noting that there are a variety of SPN structures which can be transformed to give a net of the above form. Two examples follow.

1. If a set A of transitions enabled together can fire simultaneously as well as independently we can alter the original SPN by adding a new transition to replace the firing time “clock” for the set A . The new transition will have input bag given by $\bigcup_{t \in A} \mathbf{I}(t)$ and output bag $\bigcup_{t \in A} \mathbf{O}(t)$, where $\mathbf{I}(t)$ and $\mathbf{O}(t)$ are input and output bags respectively for transition t . This new SPN is equivalent to the original. The reason for making such a technical change is to keep the notation relatively simple.

2. Without loss of generality we can assume that there is a one to one correspondence between input bags and transitions. Any SPN in which a set of transitions, $\mathcal{B} \subseteq \mathcal{T}$ have a common input bag can be modelled by amalgamating the transitions in \mathcal{B} into a single transition, noting the fact that whenever one of these transitions in enabled then all are enabled. For any state \mathbf{m} in which the transitions in \mathcal{B} are enabled the firing time distribution of the new transition will be negative exponential, with parameter $\sum_{t \in \mathcal{B}} q(\mathbf{m}, t)$. The routing probabilities for the next marking are appropriately weighted by the probabilities, $q(\mathbf{m}, t) [\sum_{s \in \mathcal{B}} q(\mathbf{m}, s)]^{-1}$, $t \in \mathcal{B}$ indicating which transition has fired.

The above procedures can be performed on any SPN before analysis takes place. Thus without loss of generality we can assume henceforth that the SPN is structured so that transitions fire independently of one another and no two transitions have the same input bag.

If t is not enabled in \mathbf{m} then $q(\mathbf{m}, t)$ is equal to zero. Otherwise we assume that $q(\mathbf{m}, t)$ is of the form

$$q(\mathbf{m} + \mathbf{I}(t), t) = \frac{\varphi(\mathbf{m})}{\Phi(\mathbf{m} + \mathbf{I}(t))} \chi(\mathbf{I}(t)), \quad (1)$$

where $\varphi(\cdot)$, $\Phi(\cdot)$ and $\chi(\cdot)$ are arbitrary but given positive functions.

The following examples illustrate the flexibility of this definition.

EXAMPLE 2.1. The standard SPN models assume that transition firing times are marking independent. In such a SPN transition t fires first with rate $q(t)$ (i.e. mean firing time $[q(t)]^{-1}$). This can be modelled using equation (1) by choosing $\varphi(\cdot) = \Phi(\cdot) \equiv 1$ and $\chi(\mathbf{I}(t)) = q(t)$.

EXAMPLE 2.2 The functions $\varphi(\cdot)$ and $\Phi(\cdot)$ have an important role when the nets being modelled have marking dependent transition firing rates. For example, consider a net which has the structure in Figure 1 where C is an arbitrary SPN configuration.

When the marking is \mathbf{m} we wish to choose $\varphi(\cdot)$ and $\Phi(\cdot)$ so that the transitions in C have marking independent firing times and so that the firing rate for t^* is $q(t^*, m(1), m(2))$, i.e. dependent on the number of tokens in $p(1)$

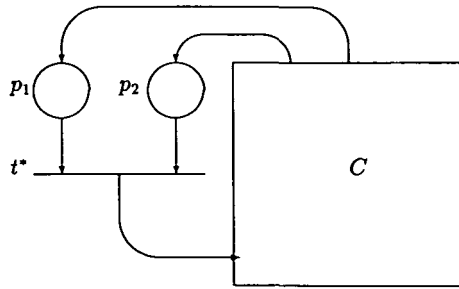


FIGURE 1.

and $p(2)$. Define

$$\chi(\mathbf{I}(t)) = \begin{cases} 1 & \text{if } t = t^* \\ q(t) & \text{otherwise} \end{cases}$$

and

$$\Phi(\mathbf{m}) = \varphi(\mathbf{m}) = \prod_{l=0}^{\min(m(1), m(2))} [q(t^*, m(1) - l, m(2) - l)]^{-1}$$

Substitution into (1) gives

$$q(\mathbf{m}, t) = \begin{cases} q(t^*, m(1), m(2)) & \text{if } t = t^* \\ q(t) & \text{otherwise} \end{cases}$$

as required.

EXAMPLE 2.3. More generally, let

$$\chi(\mathbf{I}(t^*)) = 1 \quad \text{and} \quad \Phi(\mathbf{m}) = q(t^*, m(1), m(2))\varphi(\mathbf{m} - \mathbf{e}_1 - \mathbf{e}_2),$$

where \mathbf{e}_i , $i = 1, 2$ is a unit vector with a 1 in the i th position and zeros elsewhere. This again gives the result $q(\mathbf{m}, t^*) = q(t^*, m(1), m(2))$ without placing any restrictions in C .

32. The routing process

The central feature of our analysis is to consider the transitions of the SPN to be themselves states in a Markov chain, which we call the routing process. This is achieved by first considering the input and output bags of the SPN to be the states of a Markov chain and under suitable conditions finding a one to one correspondence between the states of this Markov chain and the transitions of the net.

Although each transition t of our SPN has a unique input bag the probabilistic routing allows for a number of different output bags. Let $\mathcal{O}(t)$ be the

set of output bags of transition t and define

$$\mathcal{R} = \bigcup_{t \in \mathcal{T}} \mathcal{O}(t) \cup \bigcup_{t \in \mathcal{T}} \{\mathbf{I}(t)\}.$$

\mathcal{R} is the set of all vectors which are either output or input bags for the SPN. We assume the net is such that \mathcal{R} is finite.

Define single step transition probabilities on the set \mathcal{R} as

$$\bar{p}(\mathbf{r}, \mathbf{r}') = p(\mathbf{I}(t), t, \mathbf{O}_j(t))$$

whenever there exists a transition t and a j , with $\mathbf{r} = \mathbf{I}(t)$, $\mathbf{r}' = \mathbf{O}_j(t)$. Otherwise let $\bar{p}(\mathbf{r}, \mathbf{r}') = \delta_{\mathbf{r}\mathbf{r}'}$ where $\delta_{\mathbf{r}\mathbf{r}'}$ is the Kronecker delta. Note that the requirement that no vector can be the input bag of two distinct transitions means that these probabilities are well defined.

Now define the set,

$$\mathcal{F} = \{f(\cdot) : f(\mathbf{r}) > 0, \chi(\mathbf{r})f(\mathbf{r}) = \sum_{\mathbf{r}'} \chi(\mathbf{r}')f(\mathbf{r}')\bar{p}(\mathbf{r}', \mathbf{r}), \forall \mathbf{r} \in \mathcal{R}\},$$

where $\chi(\mathbf{r})$ is the function given in (1) when $\mathbf{r} = \mathbf{I}(t)$ for some t and $\chi(\mathbf{r}) = 1$ otherwise. Note that $\chi(\cdot)f(\cdot)$ is an invariant measure for the routing process.

In Section 4 we need \mathcal{F} to be nonempty. The effect of this assumption on the structure of the net is examined in the following results.

LEMMA 1. *For \mathcal{F} to be nonempty all vectors $\mathbf{I}(t), t \in \mathcal{T}$ must be in positive recurrent communicating classes of the routing process.*

PROOF. For t to fire in some marking, $\chi(t) > 0$. Assume there exists a function $f(\cdot) \in \mathcal{F}$, then $\chi(\mathbf{I}(t))f(\mathbf{I}(t))$ is positive. Since the set \mathcal{R} is finite, and $\chi(\cdot)f(\cdot)$ is an invariant measure of the routing process this implies that $\mathbf{I}(t)$ is in a positive recurrent communicating class.

COROLLARY 1. *For \mathcal{F} to be nonempty,*

- (a) *all $\mathbf{r} \in \mathcal{R}$ must be the input bag for some transition t ,*
- (b) *all $\mathbf{r} \in \mathcal{R}$ must be an output bag for some transition t .*

PROOF. (a) Assume \mathcal{F} is nonempty and that \mathbf{r} is not an input bag for any transition t . Since state \mathbf{r} is not the input bag to some transition then $\bar{p}(\mathbf{r}, \mathbf{r}') = \delta_{\mathbf{r}\mathbf{r}'}$ and \mathbf{r} is an absorbing state in the routing process. Hence \mathbf{r} resides in its own communicating class. Also, $\mathbf{r} \in \mathcal{O}(t)$ for some transition t means that there exists an \mathbf{r}' such that $\bar{p}(\mathbf{r}', \mathbf{r}) > 0$, implying that \mathbf{r}' is not in a positive recurrent communicating class. This contradicts Lemma 1.

(b) If \mathbf{r} is not the output bag for some transition t then there exists no \mathbf{r}' with $\bar{p}(\mathbf{r}', \mathbf{r}) > 0$. Thus \mathbf{r} must be a transient state. This contradicts Lemma 1 in conjunction with Corollary 1(a).

COROLLARY 2. *If \mathcal{F} is nonempty there exists a one to one correspondence between distinct elements of \mathcal{R} and elements of \mathcal{T} .*

PROOF. It follows from Corollary 1(a) that each vector $\mathbf{r} \in \mathcal{R}$ is the input bag for some transition $t \in \mathcal{T}$. By the assumption in Section 2 each input bag has a unique transition. The result follows.

REMARK. Corollary 2 states that for nets with nonempty \mathcal{F} each output bag of a transition is the input bag of another transition. That is each $\mathbf{r} \in \mathcal{R}$ can be uniquely identified with a transition. Consequently we can define the routing process to be a Markov chain on the set of transitions \mathcal{T} , rather than on the set of input and output bags \mathcal{R} . Without loss of generality we will retain the same notation by assuming that the one step transition probabilities for the routing process on \mathcal{T} are $\bar{p}(t, s) = \bar{p}(\mathbf{I}(t), \mathbf{I}(s))$, $s, t \in \mathcal{T}$. \mathcal{F} can also be redefined as

$$\mathcal{F} = \{f(\cdot) : f(t) > 0, \chi(t)f(t) = \sum_{s \in \mathcal{T}} \chi(s)f(s)\bar{p}(s, t) \forall t \in \mathcal{T}\}$$

and $\chi(\cdot)f(\cdot)$ is an invariant measure for the routing process. As the routing process consists of only positive recurrent communicating classes it may be assumed that $\chi(t)f(t) > 0 \forall t \in \mathcal{T}$.

4. Closed form solution

THEOREM 1. *Assume that there exists a function $f(\cdot) \in \mathcal{F}$ and a function $\{g(\cdot) : Z^P \rightarrow R\}$ which have the property that, for all $t \in \mathcal{T}$ and $\mathbf{m} + \mathbf{I}(t)$ in the reachability graph,*

$$\frac{g(\mathbf{m} + \mathbf{I}(t))}{g(\mathbf{m} + \mathbf{I}(s))} = \frac{f(t)}{f(s)} \tag{2}$$

whenever $\bar{p}(t, s) > 0$.

Then the equilibrium distribution of the SPN is given by

$$\pi(\mathbf{m}) = K\Phi(\mathbf{m})g(\mathbf{m}) \tag{3}$$

where $\Phi(\cdot)$ is given in (1) and K is a normalising constant.

PROOF. Define the reversed time transition probabilities for the routing process as

$$\bar{p}^R(s, t) = \frac{\chi(t)f(t)\bar{p}(t, s)}{\chi(s)f(s)} \tag{4}$$

and postulate an equilibrium distribution for the SPN given by (3) subject to (2).

Note that any marking change of the SPN begins with a marking $\mathbf{m} + \mathbf{I}(t)$ with $t \in \mathcal{T}$ enabled. Transition t fires producing an output bag which, because \mathcal{F} is nonempty, is the input bag of another transition, s say. The resultant marking is $\mathbf{m} + \mathbf{I}(s)$, and transition s is enabled. In reversed time the markings and enabled transitions are unchanged but the reverse operation occurs.

Postulate that the marking change rate for the SPN in reversed time is

$$q^R(\mathbf{m} + \mathbf{I}(s), \mathbf{m} + \mathbf{I}(t)) = \frac{\varphi(\mathbf{m})\chi(s)}{\Phi(\mathbf{m} + \mathbf{I}(s))} \bar{p}^R(s, t) \tag{5}$$

From Kelly [8], the validity of

$$\pi(\mathbf{m} + \mathbf{I}(t))q(\mathbf{m} + \mathbf{I}(t), t)\bar{p}(t, s) = \pi(\mathbf{m} + \mathbf{I}(s))q^R(\mathbf{m} + \mathbf{I}(s), \mathbf{m} + \mathbf{I}(t)) \tag{6}$$

and

$$\sum_{t \in \mathcal{T}} q(\mathbf{m}, t) = \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{T}} q^R(\mathbf{m}, \mathbf{m} + \mathbf{I}(s) - \mathbf{I}(t)) \tag{7}$$

is sufficient to establish that $\pi(\mathbf{m})$ as given by (3) subject to (2) is the equilibrium distribution of the SPN and that the postulate of (5) is correct.

This can be easily verified using (1), (2), (3), (4) and (5).

Conventionally the phrase “product form” has been associated with the equilibrium solution of Jackson networks which involves a product over the nodes of the network. In Example 5.1, following, it is shown how Jackson networks and their solution can be derived as a special case of the above results.

Theorem 1 also gives a product form solution. The product is between two distinct terms. The first of these is $\Phi(\cdot)$, a function related to the transition firing rates and to marking dependent properties of the SPN. The second, $g(\cdot)$, can be evaluated by analysing the network skeleton provided by the transitions and the routing process. The effect is to reduce the problem of finding a stationary distribution for the number of tokens in each place of a SPN from that of solving balance equations for a Markov process on the markings, which grows exponentially in size with an increase in the number of tokens, to solving balance equations for a Markov chain on the transitions, the size of which does not change with an increase in the number of tokens. Example 5.2 following, illustrates this effect.

5. Examples

EXAMPLE 5.1. Jackson networks. Consider a SPN as defined in Section 2 with the additional property that each input and output bag of any transition contains just a single arc. Let $\chi(t) = 1, \forall t \in \mathcal{T}$ and $\Phi(\mathbf{m}) = \varphi(\mathbf{m}) = \prod_{t \in \mathcal{T}} \prod_{l=1}^{m(t)} (q(t, l))^{-1}$. With these parameter values and assumptions on the

net structure, $f(\cdot) \in \mathcal{F}$ satisfies

$$f(t) = \sum_{s \in \mathcal{S}} f(s)\bar{p}(s, t) \quad \forall t \in \mathcal{T} \tag{8}$$

and $q(\mathbf{m}, t) = q(m(t), t)$. A function $g(\cdot)$ given by $g(\mathbf{m}) = \prod_{t \in \mathcal{T}} f(t)^{m(t)}$ satisfies (2) for all \mathbf{m} and the equilibrium solution for the SPN is

$$\pi(\mathbf{m}) = K \prod_{t \in \mathcal{T}} f(t)^{m(t)} \left[\prod_{l=1}^{m(t)} q(l, t) \right]^{-1}. \tag{9}$$

With the above choices the SPN under consideration becomes a Jackson network of queues. Each place has its own transition. Tokens are served one at a time, through transition t with service rate $q(m(t), t)$ when $m(t)$ tokens are available. These tokens are then independently routed according to the probability distribution $\bar{p}(t, s)$, $t, s \in \mathcal{T}$. Equations (8) are the traffic equations for the Jackson network and the standard solution is given by (9).

EXAMPLE 5.2. A SPN with conflict sets. Consider the SPN depicted in Figure 2. In this example the routing process on the transitions consists of three communicating classes, $C_1 = \{t_1, t_3, t_7\}$, $C_2 = \{t_4, t_5\}$, $C_3 = \{t_2, t_6\}$ with routing probabilities given by, $\bar{p}(t_1, t_3) = \bar{p}(t_3, t_7) = \bar{p}(t_7, t_1) = 1$, $\bar{p}(t_4, t_5) = \bar{p}(t_5, t_4) = 1$ and $\bar{p}(t_2, t_6) = \bar{p}(t_6, t_2) = 1$.

Assume that $\chi(t_i) = q(i)$, $1 \leq i \leq 7$. Consequently

$$\chi(t_i)f(t_i) = q(i)f(t_i) = \begin{cases} \alpha_1 & \text{if } i = 1, 3, 7 \\ \alpha_2 & \text{if } i = 4, 5 \\ \alpha_3 & \text{if } i = 2, 6 \end{cases}$$

where $\alpha_1, \alpha_2, \alpha_3$ are arbitrary and for this example can be set to unity.

Hence $f(t_i) = [q(i)]^{-1}$.

A function $g(\mathbf{m})$ which takes the form

$$g(\mathbf{m}) = \left[\frac{q(2)}{q(6)} \right]^{m(6)} \left[\frac{q(4)}{q(5)} \right]^{m(8)} \left[\frac{q(1)}{q(7)} \right]^{m(7)} \left[\frac{q(1)}{q(3)} \right]^{\min\{m(4), m(5)+m(8)\}} \tag{10}$$

satisfies (2) for all pairs of immediately reachable markings.

As an example, for markings $\mathbf{m} + \mathbf{I}(t_1)$ and $\mathbf{m} + \mathbf{I}(t_3)$, the left hand side of (2) is

$$\frac{g(\mathbf{m} + \mathbf{I}(t_1))}{g(\mathbf{m} + \mathbf{I}(t_3))} = \frac{1}{\frac{q(1)}{q(3)}} = \frac{q(3)}{q(1)} = \frac{f(t_1)}{f(t_3)} \tag{11}$$

as required.

The equilibrium distribution for this SPN is

$$\pi(\mathbf{m}) = K\Phi(\mathbf{m})g(\mathbf{m}) \tag{12}$$

with $g(\mathbf{m})$ given by (10).

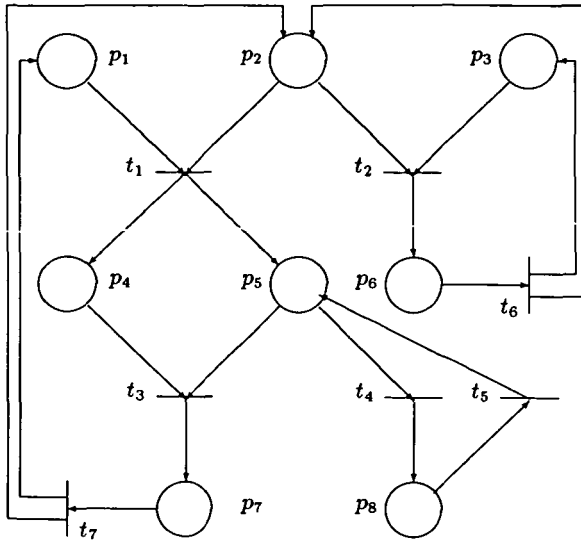


FIGURE 2.

1. Note that neither $\varphi(\cdot)$ nor $\Phi(\cdot)$ have yet been specified but that the solution has been found. As shown in Section 2 choosing $\varphi(\cdot)$ and $\Phi(\cdot)$ appropriately results in meaningful marking dependent firing rates.

2. The initial marking has not been used to find the solution given by (10), therefore any choice of the initial marking would result in the SPN having the same equilibrium solution. The solution given by (10) has normalising constant K , as its only unknown. Once given the initial marking, the reachability graph comprising the set of markings \mathcal{M} may be found. By summing (1) over the markings in \mathcal{M} and noting that $\sum_{m \in \mathcal{M}} \pi(m) = 1$, the normalising constant can be calculated by

$$K = \sum_{m \in \mathcal{M}} [\Phi(m)g(m)]^{-1}.$$

EXAMPLE 5.3. Consider the SPN illustrated in Figure 3. As the input bags of the transitions t'_1 and t''_1 are identical we follow the pre-analysis procedure outlined in Section 2. We amalgamate these two transitions and introduce probabilistic routing to create a SPN of the form given in Figure 4.

If the SPN of Figure 3 has mean transition firing times, $[q(1')]^{-1}$, $[q(1'')]^{-1}$, $[q(2)]^{-1}$, $[q(3)]^{-1}$, $[q(4)]^{-1}$, for transitions $t'_1, t''_1, t_2, t_3, t_4$ respectively, then the SPN of Figure 4 has mean firing times $[q(1)]^{-1}$, $[q(2)]^{-1}$, $[q(3)]^{-1}$, $[q(4)]^{-1}$, with $q(1) = q(1') + q(1'')$. The probabilistic routing when t_1 fires produces output bags $(0, 1, 2, 0, 0)$ (solid arcs) and $(0, 0, 1, 1, 0)$ (dashed arcs) with probabilities, $a = q(1')/q(1)$ and $b = q(1'')/q(1)$ respectively.

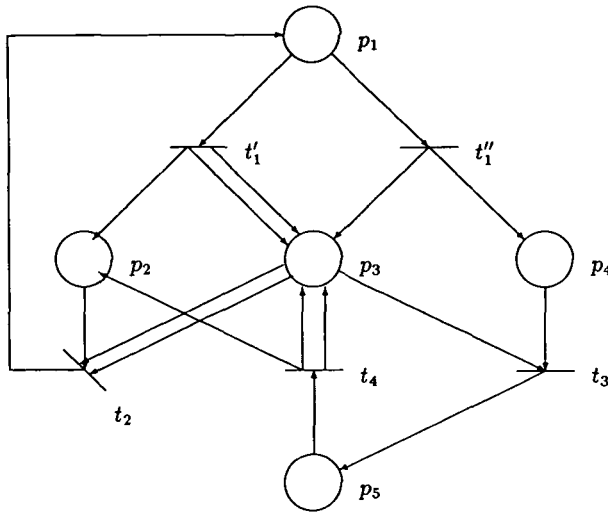


FIGURE 3.

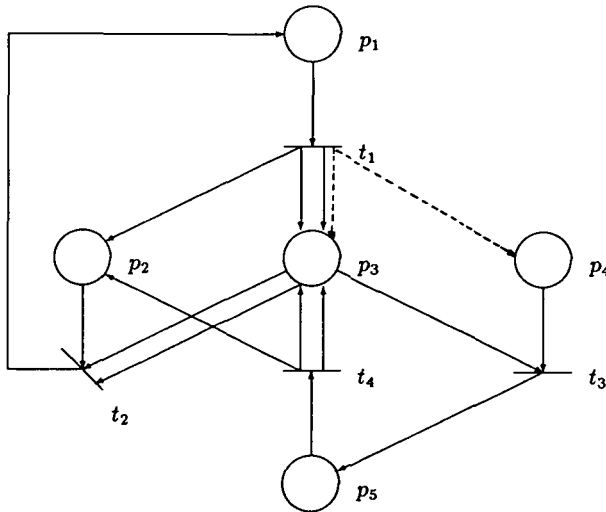


FIGURE 4.

The routing process has one step probabilities $\bar{p}(t_1, t_2) = a$, $\bar{p}(t_1, t_3) = b$, $\bar{p}(t_2, t_1) = \bar{p}(t_3, t_4) = \bar{p}(t_4, t_2) = 1$ and invariant measure, $(1, 1, b, b)$.

In a similar manner to that adopted in Example 5.2, let

$$g(\mathbf{m}) = \left[b \frac{q(2)}{q(4)} \right]^{m(5)} \left[b \frac{q(2)}{q(3)} \right]^{m(4)} \left[\frac{q(1)}{q(2)} \right]^\delta$$

where $\delta = \min(m(2) + m(5) + m(4), [(m(3) + 2m(5) + m(4))/2])$ and $[(m(3) + 2m(5) + m(4))/2]$ denotes the integer part of $(m(3) + 2m(5) + m(4))/2$.

Again it is simple to show that $g(\mathbf{m})$ satisfies (2) for all pairs of immediately reachable markings and the equilibrium distribution is again given by $\pi(\mathbf{m}) = K\Phi(\mathbf{m})g(\mathbf{m})$.

This example involves transitions with common input bags, but as in Example 5.2, there is no need to specify the form of the functions $\Phi(\cdot)$ and $\varphi(\cdot)$, nor is there a need to give the initial marking. Again the initial marking specifies the reachability graph and therefore influences the solution only through the normalising constant.

6. Conclusion

This paper has considered the relevance to SPNs of recent work on batch movement around networks. We have introduced some basic ideas on finding equilibrium distributions based on a net analysis rather than the underlying state space analysis. Since many common protocols produce reachability graphs with 10^6 or 10^7 markings, a performance analysis based on the solution of the global balance equations is unmanageable.

Consequently approaches such as the one presented in this paper become crucial if a performance analysis is required for any complex protocol. Calculating the normalising constant may also be difficult in our method. However since we have shown the SPN has a type of product form solution, there is scope to apply methods similar to the convolution algorithms and mean value analysis of queueing networks (see Buzen [4] and Reiser [12]).

Acknowledgement

The authors would like to thank the Teletraffic Research Centre at the University of Adelaide and Telecom Australia for supporting the research presented in this paper.

References

- [1] G. Balbo, S. Bruell and S. Ghanta, "Combining queueing network and generalized stochastic Petri net models for the analysis of some software blocking phenomena", *IEEE SE-12*, No. 4 (1986) 561–576.
- [2] J. Billington, "Protocol Engineering and Nets", *Proc. Workshop on Appl. and Theo. of Petri Nets, Zaragoza (1987)* 137–156.

- [3] J. Billington, G. Wheeler and M. C. Wilbur-Han, "Protean: a high level Petri net tool for the specification and verification of communication protocols", *IEEE Trans. Soft. Eng. SE-14* (1988).
- [4] J. P. Buzen, "Computational algorithm for closed networks with exponential servers", *Comm. Assoc. Comput. Mach.*, **16** (1973) 527–531.
- [5] G. Florin and S. Natkin, "Evaluation based upon stochastic Petri nets of the maximum throughput of a full duplex protocol", in *Appl. and Theo. of Petri Nets*, (eds: C. Girault and W. Reisig), *Informatik-Fachberichte*, **52** (1985) 280–288.
- [6] W. Henderson, C. E. M. Pearce, P. G. Taylor and N. M. van Dijk, "Closed queueing networks and batch services", submitted 1988.
- [7] W. Henderson and P. Taylor, "Open networks of queues with batch arrivals and batch services", submitted 1988.
- [8] F. Kelly, *Reversibility and stochastic networks*, (Wiley, London, 1979).
- [9] M. A. Marson, G. Balbo and G. Conti, "A class of generalised stochastic Petri nets for the performance evaluation of multiprocessor systems", *ACM Trans. Comp. Syst.* **2** (1984) 93–122.
- [10] M. A. Marsan, G. Chiola and A. Fumagalli, "An accurate performance model of CSMA/CD Bus LAN", *Euro. Workshop on Appl. and Theo. of Petri Nets* **7** (1986).
- [11] J. L. Peterson, *Petri net theory and the modeling of systems* (Prentice-Hall, 1981).
- [12] M. Reiser, "Mean value analysis and convolutional method for queue dependent servers in closed queueing networks", *Perf. Eval.* **1** (1981) 7–18.
- [13] F. J. W. Symons, "Introduction to numerical Petri nets, a general graphical model of concurrent processing systems", *Aust. Tele. Res.* **14** (1980).