Fluid Stochastic Petri Nets: Theory, Applications, and Solution[1]

Graham Horton[4], Vidyadhar G. Kulkarni[5], David M. Nicol[2] and Kishor S. Trivedi[3]

[4]Lehrstuhl für Rechnerstrukturen (IMMD 3),
Universität Erlangen-Nürnberg, Martensstr. 3,
91058 Erlangen, Germany.

[5]Department of Operations Research
University of North Carolina, Chapel Hill, NC 27599, U.S.A.

[2]Department of Computer Science
College of William and Mary, Williamsburg, VA 23187-8795, U.S.A.

[3]Center for Advanced Computing and Communication
Department of Electrical and Computer Engineering
Duke University, Durham, NC 27708-0291, U.S.A.

Abstract

In this paper we introduce a new class of stochastic Petri nets in which one or more places can hold fluid rather than discrete tokens. We define a class of fluid stochastic Petri nets in such a way that the discrete and continuous portions may affect each other. Following this definition we provide equations for their transient and steady-state behavior. We present several examples showing the utility of the construct in communication network modeling and reliability analysis, and discuss important special cases. We then discuss numerical methods for computing the transient behavior of such nets. Finally, some numerical examples are presented.

1 Introduction

One of the difficulties encountered while using Petri nets is that the reachability graph tends to be very large in practical problems. Drawing a parallel with fluid flow approximations in performance analysis of queueing systems, we may define fluid within a Petri-net to approximate token movement. Alternatively, some physical systems have not previously admitted a Petri net modeling approach, as they explicitly contain continuous fluid-like quantities which are controlled with discrete logic. This paper presents a new methodology for modeling such systems.

Stochastic fluid flow models are increasingly used in the performance analysis of communications [3, 10, 13] and manufacturing systems. On the other hand, stochastic Petri nets with discrete places provide a useful framework for specifying and solving performance and reliability models of discrete event dynamic systems [1, 6, 9, 17, 19]. It is natural to extend the stochastic Petri net framework to Fluid Stochastic Petri Nets (FSPNs) by introducing places with continuous tokens and arcs with fluid flow so as to handle stochastic fluid flow systems. This paper extends the model in an earlier paper [18] by allowing the level of fluid in continuous places to affect the enabling of timed transitions and the rates of fluid flow into and out of continuous places. Rules for transition enabling and firing are extended to reflect the notion that flow through a fluid place represents token movement.

An FSPN contains two types of places: discrete places containing a non-negative integer number of tokens, and continuous places containing fluid. Transition firings are determined by both discrete places and continuous places, and fluid flow is permitted through the enabled timed transitions in the Petri net. Associating exponentially distributed or zero firing time with transitions, we can then write the differential equations for the underlying stochastic process. We also provide additional examples of FSPN usage, and discuss numerical issues arising in the solution of the underlying dynamic equations.

The main motivation of this paper is to put the research by Mitra and his colleagues [3, 10, 13] in the context of Petri nets, to make some extensions and to investigate the numerical transient analysis of such stochastic fluid models.

The paper is organized as follows. In Section 2 we develop the fluid model of stochastic Petri nets and in Section 3 we discuss their analysis. Examples and special cases are described in Section 4, numerical solution techniques are described in Section 5, while numerical examples are given in Section 6. The paper concludes in Section 7.

2 The Stochastic Fluid Model

Following the customary notation [8, 14, 16] for defining Petri nets and their extensions, we define a fluid stochastic Petri net (FSPN) as a 8-tuple $(\mathcal{P}, \mathcal{T}, \mathcal{A}, m_0, \mathcal{F}, \mathcal{W}, \mathcal{R}, \mathcal{G})$. $\mathcal{P}$ is a set of places partitioned into a

set of discrete places $\mathcal{P}_d$ and a set of continuous places $\mathcal{P}_c$. The number of fluid places is $F \geq 0$, indexed by $k = 1, 2, \cdots, F$. In a graphical representation, we shall depict continuous places by means of two concentric circles. The set of transitions $\mathcal{T}$ is partitioned into a set of (exponentially distributed firing) timed transitions $\mathcal{T}_E$ and a set of immediate transitions $\mathcal{T}_I$. The set of directed arcs $\mathcal{A}$ is partitioned into two subsets $\mathcal{A}_c$ and $\mathcal{A}_d$. $\mathcal{A}_c$ is a subset of $(\mathcal{P}_c \times \mathcal{T}_E) \cup (\mathcal{T}_E \times \mathcal{P}_c)$ while $\mathcal{A}_d$ is a subset of $(\mathcal{P}_d \times \mathcal{T}) \cup (\mathcal{T} \times \mathcal{P}_d)$. In a graphical representation, arcs in $\mathcal{A}_c$ are drawn as double lines (to suggest a pipe) while those in $\mathcal{A}_d$ are drawn as single lines.

Let $m_d = (\#p_i, i \in \mathcal{P}_d)$ be the vector of the number of tokens in discrete places and let $\vec{x} = (x_k, k \in \mathcal{P}_c)$ be the vector of the fluid levels in continuous places. We will say that m_d is the discrete marking of the net. Let M denote the number of discrete markings, which will be indexed by the symbols i and j, with $i, j = 1, 2, \cdots, M$. The complete state (marking) of a fluid Petri net is described by the vector $m = (\vec{x}, m_d)$ where m_d is the marking of the discrete part of the state and $\vec{x}$ keeps track of the fluid levels in the continuous places. Let $\mathcal{M}$ be the set of all complete markings $(\vec{x}, m_d)$ and $\mathcal{M}_d$ be the set of all discrete markings. The initial marking is $m_0 = (\vec{x}_0, m_{d0})$.

In our formulation an enabled transition in $\mathcal{T}_E$ may drain fluid out of its continuous input places, and may pump fluid into its continuous output places. The rates of flow may be dependent on the complete marking $(\vec{x}, m_d)$. In a general formulation of a stochastic Petri net (embodied, for example, by SPNP [7, 8]) the conditions for enabling can be specified either through explicit arcs or through Boolean functions known as guards. We will allow both of these possibilities. We will continue to use the enabling and firing rules employed in SPNP with the additional possibility of a guard associated with a timed transition being able to base the enabling condition not only on the discrete marking of the net but also on the continuous part. We disallow the enabling of immediate transitions by fluid levels as this would lead to $\vec{x}$-dependent vanishing markings, which cannot be eliminated in a manner analogous to that of GSPNs. Thus the guard function $\mathcal{G}$ is defined for any timed transition in $\mathcal{T}_E$ so that $\mathcal{G} : \mathcal{T} \times \mathcal{M} \to \{0, 1\}$. For a timed transition $\tau \in \mathcal{T}_E$, $\mathcal{G}(\tau, m)$ is a Boolean function that will be evaluated in each marking, and if it evaluates to *true*, the transition may be enabled; otherwise τ is disabled. Upon firing, the transition removes a specified number of tokens from each discrete input place, and deposits a specified number of tokens in each discrete output place. The basic extension we have made from [18] here is that fluid levels in continuous places can change the enabling/disabling of timed transitions in the discrete part of the net. A discrete marking m_d is said to be a vanishing marking if one or more immediate transitions are enabled by it; otherwise it is a tangible marking.

The firing rate function $\mathcal{F}$ is defined for timed transitions $\mathcal{T}_E$ so that $\mathcal{F} : \mathcal{T}_E \times \mathcal{M} \to \mathbb{R}^+$. Thus if a timed transition τ is enabled in (tangible) marking m, it fires with rate $\mathcal{F}(\tau, m)$. Note once again that in [18], these rates were not allowed to depend upon the fluid levels but now we do allow the firing rates to be dependent on fluid levels.

As in [18], the weight function $\mathcal{W}$ is defined for immediate transitions $\mathcal{T}_I$ so that $\mathcal{W} : \mathcal{T}_I \times \mathcal{M}_d \to \mathbb{R}^+$.

Thus if an immediate transition i is enabled in (a vanishing) marking m_d, it fires with probability

$$\frac{\mathcal{W}(i, m_d)}{\displaystyle\sum_{\theta \in \mathcal{T}_I \text{ enabled in } m_d} \mathcal{W}(\theta, m_d)}.$$

Next we describe the evolution of the continuous part of the marking. The flow rate function $\mathcal{R}$ is defined for the arcs connecting a continuous place and a timed transition so that $\mathcal{R} : \mathcal{A}_c \times \mathcal{M} \to \mathrm{I\!R}^+ \cup \{0\}$. Thus when the FSPN marking is $m_t \in \mathcal{M}$ at time t, fluid can leave place $k \in \mathcal{P}_c$ along the arc $(k, \tau) \in \mathcal{A}_c$ at rate $\mathcal{R}((k, \tau), m_t)$ and can enter the continuous place k at rate $\mathcal{R}((\tau, k), m_t)$ along the arc $(\tau, k) \in \mathcal{A}_c$ for each $\tau \in \mathcal{T}_E$ that is enabled in m_t. The instantaneous rate at which fluid builds in a place $k \in \mathcal{P}_c$ at time t, in marking m_t, is then given by

$$r_k(m_t) = \sum_{\tau \in \mathcal{T}_E \text{ enabled in } m_t} \mathcal{R}((\tau, k), m_t) - \sum_{\tau \in \mathcal{T}_E \text{ enabled in } m_t} \mathcal{R}((k, \tau), m_t).$$

We require that for every discrete marking m_d and arc (τ, k), the rate $\mathcal{R}((\tau, k), (\vec{x}, m_d))$ be a "nice" function of $\vec{x}$, e.g., it is piecewise continuous. Observe that since m_t contains continuous levels $\vec{x}$, $r_k(m_t)$ may change as a function of t even if the discrete part of m_t does not change. Once again we have extended the definition in [18] by allowing these rates to be dependent on the fluid levels.

Now let $X_k(t)$ be the fluid level at time t in a continuous place $k \in \mathcal{P}_c$. We assume that there is an upper bound on the fluid content, that is, $X_k(t) \leq B_k$ for all $t \geq 0$. If there is no such upper bound, we set B_k to ∞. Then the sample path of $X_k(t)$ satisfies the differential equation

$$\frac{dX_k(t)}{dt} = \begin{cases} [r_k(m_t)]^+ & \text{if} \quad X_k(t) = 0 \\ [r_k(m_t)]^- & \text{if} \quad X_k(t) = B_k \\ r_k(m_t) & \text{if} \quad 0 < X_k(t) < B_k \text{ and } r_k(m_{t-}) r_k(m_{t+}) \geq 0 \\ 0 & \text{if} \quad 0 < X_k(t) < B_k \text{ and } r_k(m_{t-}) r_k(m_{t+}) < 0 \end{cases} \tag{1}$$

In the case $X_k(t) = 0$ and $r_k(m_t) < 0$, we set the actual rate equal to zero (denoted by $[r_k(m_t)]^+ = \max(r_k(m_t), 0)$) in order to maintain $X_k(t) \geq 0$. In the case that $X_k(t) = B_k$ and $r_k(m_t) > 0$, we set the actual rate equal to zero (denoted by $[r_k(m_t)]^- = \min(r_k(m_t), 0)$) in order to maintain $X_k(t) \leq B_k$. For the explantion of the remaining cases, we refer the reader to [10], Section II. The key observation (for the fourth case) is that a sign change from $+$ to $-$ in $r_k(m_t)$ at m_t will "trap" $X_k(t)$ to be constant. Finally, let $M_d(t)$ be the discrete marking at time t.

In the next section we study the joint process $(X(t), M_d(t))$, where $X(t) = [X_k(t), k \in \mathcal{P}_c]$.

3   Analysis

Recall that $X_k(t)$ is the fluid level in the k^{th} continuous place at time t. The reachability graph corresponding to the discrete part of the net gives rise to a stochastic process that is a Markov process with the state space $\mathcal{M}$. In the special case that the discrete part of the net is not affected by the fluid levels, discrete part of the net gives rise to a CTMC[1]. Let S be the discrete state space and let $\mathbf{Q}(\vec{x}) = [q_{ij}(\vec{x})]$ be the matrix of transition rates derived from the firing rate function $\mathcal{F}$ of section 2. S corresponds to the set of tangible discrete markings in $\mathcal{M}_d$. For every $\vec{x}$ and place $k \in \mathcal{P}_c$ define the diagonal matrix $\mathbf{R}_k(\vec{x}) = diag(r_k(\vec{x}, m_d))$, $m_d \in S)$.

Define the distribution function $H(t, \vec{x}, m_d) = P(X(t) \leq \vec{x}, \text{ and } M_d(t) = m_d)$ and let $\vec{H}(t, \vec{x}) = [H(t, \vec{x}, m_d), m_d \in S]$ be a row vector.

To begin with, we assume that the rate functions $R_k(\vec{x})$ are differentiable functions of $\vec{x}$ and transition rates $\mathbf{Q}(\vec{x})$ are piecewise right continuous functions of $\vec{x}$. We also assume that the capacities of the continuous places are infinite. Under these assumptions, it can be shown that $(X(t), M_d(t))$ has a density $h(t, \vec{x}, m_d)$ for all $\vec{x}$ with non-negative components and $m_d \in S$ and has a probability mass $c(t, \vec{x}, m_d)$ if $\vec{x}$ has at least one component equal to 0, where

$$c(t, \vec{x}, m_d) = \frac{P(X_k(t) = 0 \text{ if } x_k = 0, \ X_k(t) \in (x_k, x_k + dx_k) \text{ if } x_k > 0 \ \forall k)}{\prod_{k:x_k>0} dx_k}$$

Let $\vec{h}(t, \vec{x})$ denote the corresponding row vector of $h(t, \vec{x}, m_d)$.

The next theorem gives the coupled system of partial differential equations satisfied by $\vec{h}$ and c. These equations describe the transient behavior of the FSPN.

Theorem 1:

The equations satified by $\vec{h}$ are

$$\frac{\partial \vec{h}}{\partial t} + \sum_k \frac{\partial(\vec{h} \mathbf{R}_k(\vec{x}))}{\partial x_k} = \vec{h} \mathbf{Q}(\vec{x}) \tag{2}$$

$$c(t, \vec{x}, m_d) = 0$$
$$\text{if, for any } k, \qquad x_k = 0 \text{ and } r_k(\vec{x}, m_d) > 0 \tag{3}$$

$$\frac{\partial}{\partial t} c(t, \vec{x}, m_d) + \sum_{k:x_k=0} h(t, \vec{x}, m_d) r_k(\vec{x}, m_d) +$$
$$\sum_{k:x_k>0} \frac{\partial}{\partial x_k} h(t, \vec{x}, m_d) r_k(\vec{x}, m_d) = \sum_{i \in S} c(t, \vec{x}, i) q_{i, m_d}(\vec{x})$$
$$\text{if} \qquad x_k = 0 \ \Rightarrow \ r_k(\vec{x}, m_d) < 0 \quad \forall k \tag{4}$$

4

<u>Proof:</u>

A rigorous proof can be obtained by using the same technique as in [12]. Here we provide an intuitive proof for the $\vec{h}$-equation.

Assume for simplicity only one fluid place and $r_1(x, i) > 0$. Consider a portion of the function $h(t, x, i)$ at one instant in time in the vicinity of the location $x = x_l$. We consider a cell surrounding this point whose left and right boundaries are located at x_- and x_+ respectively. Let the discrete values $h(t, x_l, i)$ and $q_{ij}(x_l)$ represent the mean values of h and q_{ij} respectively within the cell. This situation is illustrated in Figure 1.


Figure 1: Derivation of FSPN equations

Probability mass must be conserved. We may therefore derive a balance equation for the change in probability mass inside the cell:

$$\left[\begin{array}{c} change\ in\ probability \\ mass\ in\ cell \end{array} \right] = \left[\begin{array}{c} total\ mass \\ entering\ cell \end{array} \right] - \left[\begin{array}{c} total\ mass \\ leaving\ cell \end{array} \right].$$

Probability mass enters the cell at the location x_- at the rate $r(x_-, i)$ and leaves at x_+ at the rate $r(x_+, i)$. In addition the cell gains probability from the corresponding cell of each other discrete marking m_j, $j \neq i$ at the rate $q_{ji}(x_l)$ and loses probability to them at a total rate of $-q_{ii}(x_l)$. The balance equation for the cell probability mass becomes

$$\Delta x \frac{dh(t, x_l, i)}{dt} = h(t, x_-, i) r_1(x_-, i) - h(t, x_+, i) r_1(x_+, i) + \Delta x (\sum_{j \neq i} q_{ji}(x_l) h(t, x_l, j) + q_{ii}(x_l) h(t, x_l, i)). \quad (5)$$

The equation is identical for the case $r_1(x, i) < 0$. Dividing the vector form of equation (5) by Δx and taking the limit as $\Delta x \to 0$ we obtain

$$\frac{\partial \vec{h}}{\partial t} + \frac{\partial (\vec{h} \mathbf{R}_1(\vec{x}))}{\partial x} = \vec{h} \mathbf{Q}(\vec{x}).$$

The argument generalizes easily to the case of more than one fluid place, yielding equation (2). The cell is in general a hypercube of sidelength Δx.

5

The boundary conditions follow from the fact that $h(t, \vec{x}, i)$ is a probability density function.

$\square$

Theorem 2:

The equations for the cumulative probability distributions $\vec{H}$ are

$$\frac{\partial \vec{H}}{\partial t} + \sum_k \frac{\partial}{\partial x_k}(\vec{H}\mathbf{R}_k(\vec{x})) = \vec{H}\mathbf{Q}(\vec{x}) - \int_0^{x_F} \ldots \int_0^{x_1} \vec{H}\frac{\partial \mathbf{Q}(\vec{x})}{\partial x_k} \, dx_1 \ldots dx_F \tag{6}$$

with the boundary conditions

$$H(t, \vec{x}, i) = 0 \quad \text{at} \quad x_k = 0, \quad \text{if} \quad r_k(\vec{x}, i) > 0 \tag{7}$$

$$\lim_{x_k \to \infty} \frac{\partial \vec{H}}{\partial x_k} = 0. \tag{8}$$

Proof:

Using the abbreviations

$$\frac{\partial \vec{H}}{\partial \vec{x}} = \frac{\partial^{(F)} \vec{H}}{\partial x_1 \ldots \partial x_F}$$

$$\int_{\vec{0}}^{\vec{x}} d\vec{x} = \int_0^{x_F} \ldots \int_0^{x_1} dx_1 \ldots dx_F$$

and noting that

$$H(t, \vec{x}, i) = c(t, \vec{x}, i) + \int_{\vec{0}}^{\vec{x}} h(t, \vec{x}, i) d\vec{x}$$

we substitute into Equation (2) and integrate with respect to $\vec{x}$:

$$\int_{\vec{0}}^{\vec{x}} \frac{\partial}{\partial t}\left(\frac{\partial \vec{H}}{\partial \vec{x}}\right) d\vec{x} + \int_{\vec{0}}^{\vec{x}} \sum_k \frac{\partial}{\partial x_k}\left(\frac{\partial \vec{H}}{\partial \vec{x}}\mathbf{R}_k(\vec{x})\right) d\vec{x} = \int_{\vec{0}}^{\vec{x}} \frac{\partial \vec{H}}{\partial \vec{x}}\mathbf{Q}(\vec{x}) \, d\vec{x}$$

which yields (6). Equation (4) is implicitly contained in Equation (6), along with the boundary condition (7).

The boundary condition (7) follows from the observation that the fluid level in place k cannot remain at 0 for a positive amount of time in state m_d if the net rate $r_k(\vec{x}, m_d) > 0$. The boundary condition (8) comes from the observation that H represents a probability distribution with respect to each x_k, and must therefore approach an asymptotic value x_k tends to infinity (or reaches its maximal value).

$\square$

This method of deriving partial differential equations that represent conservation laws is well established in computational fluid dynamics, where it is known as the "Finite Volume" technique [15].

The assumptions on $\mathbf{R}_k(\vec{x})$ can be relaxed and we can allow $\mathbf{R}_k(\vec{x})$ to be piecewise differentiable and fluid places to have finite capacities. This may introduce non-zero probability masses at the inter-region boundaries and will need to be explicitly accounted for.

The domain of Equations (2) and (6) is

$$0 < \quad x_k \quad < B_k$$
$$0 < \quad t \quad < \infty,$$

although in practice we will only be interested in the finite domain $0 < t < tmax$, $0 < x_k < \min\{xmax, B_k\}$ (where $xmax$ is finite when $B_k = \infty$).

The initial conditions for Equation (6) are

$$H(0, \vec{x}, i) \quad = \quad 1 \quad \text{if} \quad i = m_{d0} \text{ and } \vec{x} \geq \vec{x}_0$$
$$H(0, \vec{x}, i) \quad = \quad 0 \quad \text{otherwise.}$$

The initial conditions for Equation (2) are

$$c(0, \vec{x}, i) = \delta(m_0)$$

where δ is the delta function.

Now suppose the following limits exist

$$\vec{F}(\vec{x}) = \lim_{t \to \infty} \vec{H}(t, \vec{x}).$$

Then from Theorem 2, we see that the steady-state distribution $\vec{F}(\vec{x})$ obeys the following system of differential equations

$$\sum_k \frac{\partial}{\partial x_k} \left(\vec{F}(\vec{x}) \mathbf{R}_k(\vec{x}) \right) = \vec{F}(\vec{x}) \mathbf{Q}(\vec{x}) - \int_{\vec{0}}^{\vec{x}} \vec{H} \frac{\partial \mathbf{Q}(\vec{x})}{\partial x_k} \, d\vec{x}, \tag{9}$$

with the normalization condition $\lim_{\vec{x} \to \infty} \vec{F}(\vec{x})\vec{e} = 1$, where $\vec{e}$ is a column vector of all 1's.

We note that the steady-state distribution $\vec{F}$ exists when

$$\lim_{x_k \to \infty} \sum_i \pi_i(\vec{x}) r_k(\vec{x}, i) < 0 \quad k = 1 \dots F$$

where $\pi(\vec{x})$ is the solution of

$$\pi(\vec{x}) \mathbf{Q}(\vec{x}) \quad = \quad 0$$
$$\sum \pi_i(\vec{x}) = 1$$

7

3.1 FSPN with a Single Continuous Place

In the special case of a single continuous place, Equation (9) reduces to:

$$\frac{d}{dx}(\vec{F}(x)\mathbf{R}(x)) = \vec{F}(x)\mathbf{Q}(x) - \int_0^x \vec{F}\frac{\partial \mathbf{Q}(x)}{\partial x}\ dx. \tag{10}$$

In the following subsections, we consider three special cases of the Equation (10).

3.1.1 Constant Case

In the special case that $\mathbf{R}(x)$ and $\mathbf{Q}(x)$ are both independent of x, following [3], solution of such an equation is of the form $\vec{F}(x) = \vec{h}e^{\lambda x}$ where $\vec{h}$ is a row vector and λ is a scalar. Substituting in (10) we have

$$\vec{h}(\lambda\mathbf{R} - \mathbf{Q}) = 0. \tag{11}$$

If a non-zero $\vec{h}$ is to satisfy the above equation, we must have $\det(\lambda\mathbf{R} - \mathbf{Q}) = 0$. The number of solutions of $\det(\lambda\mathbf{R} - \mathbf{Q}) = 0$ equals the number of non-zero diagonal elements of $\mathbf{R}(x)$. Let these solutions be denoted by $\lambda_1, \lambda_2 \cdots, \lambda_k$. Let $\vec{h}_i$ be the solution to $\vec{h}_i(\lambda_i\mathbf{R} - \mathbf{Q}) = 0$. Then the general solution to (10) is given by

$$\vec{F}(x) = \sum_{i=1}^k a_i\vec{h}_ie^{\lambda_i x} \tag{12}$$

where the scalars a_i need to be determined from the boundary conditions and the boundedness of $\vec{F}(x)$. It is known that the number of λ_i's with positive real part equals the number of negative diagonal entries of $\mathbf{R}$. The coefficient a_i corresponding to an eigenvalue λ_i with $Re(\lambda_i) > 0$ must be zero in order to maintain boundedness of $\vec{F}(x)$. The remaining coefficients a_i are uniquely determined by the boundary condition $F(0, m) = 0$ if $r(m) > 0$.

Now if $\mathbf{R}(x)$ and $\mathbf{Q}(x)$ are both piecewise constant functions of x, we can apply the above procedure for each different segment and piece the individual solutions together [10]. In the general case, we can use numerical solution methods for linear odes that are available. We refer here to explicit methods such as RKF-45 or implicit methods such as implicit Runge-Kutta [4] or TR-BDF2 [5] and so on.

4 Examples

Next we examine a number of examples to illustrate the modeling power of FSPNs.

Figure 2: FSPN Model of Statistically Multiplexed Network Switch

4.1 Statistical Multiplexing of a Network Switch

For our first example, we recast the problem studied in [10] as a FSPN. We have K sources of ATM cells. Each source emits both high priority and low priority cells; the arrival rate of each depends on the state of a two-state Markov chain, e.g., if the source's Markov chain is in state $s \in [1, 2]$, then the rate at which high priority cells is produced is $\lambda_{hi}^{(s)}$, and the rate at which low priority cells is produced is $\lambda_{low}^{(s)}$. All cells are delivered to a single switch, with buffer capacity B. All low priority cells are discarded when the buffer level is greater than some $B_t < B$.

The FSPN for this problem is illustrated in Figure 2. The number of tokens in place p_i reflects the number of sources in state i, $i = 1, 2$. The rate at which sources change from state 1 to 2 (alt., 2 to 1) is $\#p_1\lambda_1$ (alt., $\#p_2\lambda_2$). Letting r_i^{high} and r_i^{low} denote the rate at which high and low priority cells are generated by a source in state i, the aggregate arrival rate of fluid to the buffer from sources in state i is a function of the discrete marking :

$$f_i(\#p_1, \#p_2) = \#p_i(r_i^{high} + r_i^{low}).$$

The overall arrival rate as a function of the discrete marking is

$$\gamma(\#p_1, \#p_2) = f_1(\#p_1, \#p_2) + f_2(\#p_1, \#p_2).$$

The fluid level in the place *loss* reflects the total cell loss since time 0 (like the cell buffer, it is initially empty). The rate $c(x)$ at which cells are successfully switched out of the buffer with level x is $c(x) = c$ if $x > 0$, and $c(x) = 0$ if $x = 0$. Of more interest is the function $L(\#p_1, \#p_2, x)$ describing the rate of cell loss

Figure 3: FSPN for the Machine Breakdown Model

(both into and out of the transition so labeled):

$$L(\#p_1, \#p_2, x) = \begin{cases} [\gamma(\#p_1, \#p_2) - c]^+ & \text{if } 0 \le x \le B_t \\ \#p_1 r_1^{low} + \#p_2 r_2^{low} + [\#p_1 r_1^{high} + \#p_2 r_2^{high} - c]^+ & \text{if } B_t < x < B \\ \#p_1 r_1^{low} + \#p_2 r_2^{low} + [\#p_1 r_1^{high} + \#p_2 r_2^{high} - c]^+ - & \\ \quad [c - \#p_1 r_1^{high} + \#p_2 r_2^{high}]^+ & \text{if } x = B_t \end{cases}$$

When $x \le B_t$, low priority cells are not discarded automatically, and any loss is the difference between the aggregate arrival rate and the service rate. Low priority cells are dropped automatically when $B_t < x \le B$, and further loss may be due to an inability of the server to keep up with the aggregate high priority arrival rate.

4.2   Machine Breakdown Model

In the previous example, fluid levels have no effect on the behavior of the discrete portion of the FSPN. The next example shows how fluid levels may affect the firing rate of discrete transitions. Consider a system with N statistically identical and independent components, each with a failure rate that depends on the overall system load [11]; the repair rate is μ. Work arrives to the system at a constant rate r and is completed at rate d per functioning machine. Work here is considered to be a non-negative real quantity. We model this system as an FSPN shown in Figure 3.

The model has two discrete places (p_u and p_d), one continuous place and three timed transitions (t_f, t_r and t_{always}). The number of tokens in p_u models the number of functioning machines (with the initial value N) while the number of tokens in p_d is the number of failed machines undergoing repair. The firing of the transition t_f represents a machine failure; given positive failure rates λ_1 and λ_2 and positive scaling factor α, the load-dependent firing rate of this transition with fluid level x is taken to be

$$f(\#p_u, x) = \#p_u(\lambda_1 + \lambda_2(1.0 - e^{-\alpha x})).$$
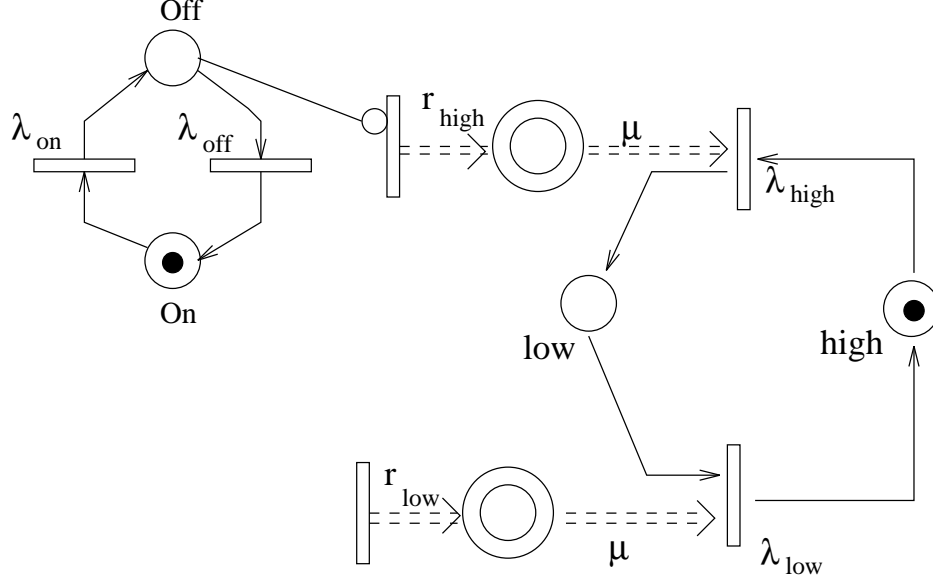
10

Figure 4: FSPN model of an alternating switch

Repair of a machine is modeled by the firing of t_r; its firing rate is μ times the number of tokens in p_d. The transition t_{always} is always enabled and it continuously pumps fluid at rate r into the continuous place. The rate of firing of this transition can be chosen to be any positive value. Whenever transition t_f is enabled, it drains fluid from the continuous place at rate d times the number of tokens in place p_u.

4.3   Alternating Switch

In the next example we model a communcation switch with two arrival streams. One stream is bursty, with a high arrival rate r_{high} when active. It is described as an alternating process, that is on for an exponential period of time with rate λ_{on}, and off for an exponential period of time with rate λ_{off}. The other stream is slow, but constant (r_{low}). The switch services workload from either stream at rate μ.

The switch is designed to allocate exponentially distributed time-slots to each stream, with rates λ_{high} and λ_{low}. The time slots alternate—the fast stream is given a slot, then the slow stream, and so on. The FSPN modeling this switch is illustrated in Figure 4.

5   Numerical Transient Analysis

5.1   Discretization

We choose to solve the equations for the probability distribution (6) rather than the density equations (2), since the latter would require the numerical treatment of a delta function used to describe the initial condition. We describe here the case $\mathbf{R}(\vec{x}) = \mathrm{const}$, $\mathbf{Q}(\vec{x}) = \mathrm{const}$.

We perform a semi-discretization of (6) in the coordinate directions x_k, using a first-order upwind method. In the nomenclature of section 3, the upwind scheme approximates values at the boundaries of the cell by those at the neighbouring grid points in the upstream direction, relative to the flow direction r. We choose a finite domain $0 < x_k < xmax$ and use G equidistant grid points. The grid spacing (mesh size) Δx is therefore equal to $xmax/(G-1)$. Note that in order that the boundary condition (8) be approximately satisfied, $xmax$ must be chosen to be sufficiently large. We use the notation $r_{k,l_1,\ldots,l_F,i} = r_k(l_1\Delta x,\ldots,l_F\Delta x,i)$, $q_{i,j,l_1,\ldots,l_F} = q_{ij}(l_1\Delta x,\ldots,l_F\Delta x)$ and $\tilde{H}_{i,l_1,\ldots,l_F}(t) = H(t,l_1\Delta x,\ldots,l_F\Delta x,i)$. The upwind discretization is given by

$$\frac{\partial}{\partial x_k} H(t,\ldots,x_k = l_k\Delta x,\ldots,i) \approx \left\{ \begin{array}{ll} \frac{1}{\Delta x}\left(\tilde{H}_{i,\ldots,l_k,\ldots}(t) - \tilde{H}_{i,\ldots,(l-1)_k,\ldots}(t)\right) & \text{if} \quad r_{k,\ldots,l_k,\ldots,i} \geq 0 \\ \frac{1}{\Delta x}\left(\tilde{H}_{i,\ldots,(l+1)_k,\ldots}(t) - \tilde{H}_{i,\ldots,l_k,\ldots}(t)\right) & \text{if} \quad r_{k,\ldots,l_k,\ldots,i} < 0 \end{array} \right. . \tag{13}$$

We obtain from the semi-discretization the linear system of ordinary differential equations

$$\frac{d\tilde{H}}{dt} = \tilde{H}\tilde{\mathbf{Q}} \tag{14}$$

where the unknown vector

$$\tilde{H} = \left(\tilde{H}_{1,0_1,\ldots,0_F},\ \tilde{H}_{1,0_1,\ldots,1_F},\ \ldots,\ \tilde{H}_{M,G-1_1,\ldots,G-1_F}\right)$$

is obtained by a lexicographic ordering of the unknowns by fluid place and by the discrete marking.

The matrix $\tilde{\mathbf{Q}}$ is given by

$$\tilde{\mathbf{Q}} = \bar{\mathbf{Q}} + \mathbf{W}$$

where $\bar{\mathbf{Q}}$ represents the CTMC of the discrete part of the net, and $\mathbf{W}$ the discretization of the space derivatives multiplied by the flow rates.

The matrix $\bar{\mathbf{Q}}$ is given by

$$\bar{\mathbf{Q}} = \left[\begin{array}{ccc} Q_{11} & \ldots & Q_{1M} \\ \vdots & \ddots & \vdots \\ Q_{M1} & \ldots & Q_{MM} \end{array} \right]$$

12

where

$$Q_{ij} = \begin{bmatrix} q_{ij,0_1,\ldots,0_F} & & \\ & \ddots & \\ & & q_{ij,(G-1)_1,\ldots,(G-1)_F} \end{bmatrix}.$$

Note that the matrix $\bar{\mathbf{Q}}$ is a CTMC matrix. The block-diagonal matrix $\mathbf{W}$ has the form

$$\mathbf{W} = -\frac{1}{\Delta x} \begin{bmatrix} W_1 & & 0 \\ & \ddots & \\ 0 & & W_M \end{bmatrix}$$

where $W_i \in \mathbb{R}^{G^F \times G^F}$ represents the discretization of the second term of equation (6) for the i-th discrete marking using the upwind method (13). Each W_i is a sparse, banded matrix in which each column is either zero or has a positive main diagonal coefficient and non-positive off-diagonal coefficients containing the rate values $r_{i,\ldots}$. In addition, the main diagonal coefficient is the negative sum of the other entries in its column. Each W_i, and thus also $\mathbf{W}$, therefore represents the transpose of a CTMC.

The linear system of ordinary differential equations (14) thus defined may be integrated by any standard method. In the numerical examples in section 6 we will use the Forward Euler scheme

$$\tilde{H}(t + \Delta t) \approx \tilde{H}(t) + \Delta t \tilde{H}(t) \tilde{Q}.$$

Note that for this scheme the discretization mesh sizes must satisfy

$$\max_{k,l_1,\ldots,l_F,i} r_{k,\ldots,i} \Delta t < \Delta x \tag{15}$$

in order that the integration be stable.

5.2   Numerical Integration in a Transformed Domain

In order to avoid the difficulty of solving equation (6) in an arbitrarily large and *a priori* unknown domain, we can perform the coordinate transformation

$$y_k = 1 - e^{-x_k} \tag{16}$$

which maps the infinite interval $x_k = [0, \infty]$ to the finite interval $y_k = [0, 1]$. Equation (6) then becomes

$$\frac{\partial \vec{H}}{\partial t} + \sum_k \frac{\partial \vec{H}}{\partial y_k} \mathbf{R}_k(\vec{y})(1 - y_k) = \vec{H} \mathbf{Q}(\vec{y}) - \int_0^y \vec{H} \frac{\partial \mathbf{Q}}{\partial \vec{y}} dy \tag{17}$$

for $\vec{H}(t, \vec{y})$. The boundary and initial conditions remain unchanged.

This minor modification to the equations makes their solution significantly easier, since the decision where to place $xmax$ must no longer be made, and the danger of an inappropriate choice is avoided. The

13

coordinate transformation has the additional advantage of compressing the infinite interval of the large x values, where in many cases the solution shows virtually no structure, into a small space. Furthermore, we can obtain numerical values at $x = \infty$, which correspond to the probabilities for the discrete markings, whereas in the untransformed case these must be approximated by values obtained at $x = xmax$.

5.3   Problem Size

Recall that G denotes the number of gridpoints in each dimension x_i, F the number of fluid places, and M the number of discrete markings. The number of time-steps to be integrated is T.

The computational complexity for the solution is $\mathcal{O}(TMG^F)$ floating point operations, since for each of T timesteps we must increment each solution value in an F-dimensional grid of sidelength G for each of M discrete markings using $O(1)$ operations. Note that for an explicit integration method such as Forward Euler, because the condition (15) must be satisfied, an increase in G must ultimately be accompanied by a proportionate increase in T.

The storage requirements of the algorithm are at least $4MG^F$ bytes since for each of M discrete markings we must store an F-dimensional grid of floating point numbers with sidelength G. Solutions at successive time-steps can be overwritten.

Since the sidelength of the grids G may typically be of the order of 50 or more when a simple discretization is used, we see that this would seriously limit the size of the FSPNs that can be solved. Future work must therefore include strategies for reducing the amount of memory needed to represent the function H.

5.4   Time Integration by Randomization

Randomization is a numerical method widely used for the solution of systems of ODEs of the form (14). It has the advantages of high numerical stability and low roundoff error, *a priori* specification of absolute error tolerance requirements, and is in addition often found to be faster than numerical integration schemes. The method has superior roundoff error behavior when the matrix $\mathbf{P} = (I + \frac{1}{\lambda}\tilde{\mathbf{Q}})$ has all non-negative entries, where $\lambda \geq \max |\tilde{\mathbf{Q}}_{ii}|$. This is, for example, the case for a CTMC.

As the following Theorem shows, the semi-discretized FSPN equation (14) also has this property, indicating that randomization may be the method of choice for computing transient solutions.

Theorem 3:

For the matrix $\mathbf{P} = (I + \frac{1}{\lambda}\tilde{\mathbf{Q}})$ where $\tilde{\mathbf{Q}}$ is the matrix of equation (14) holds
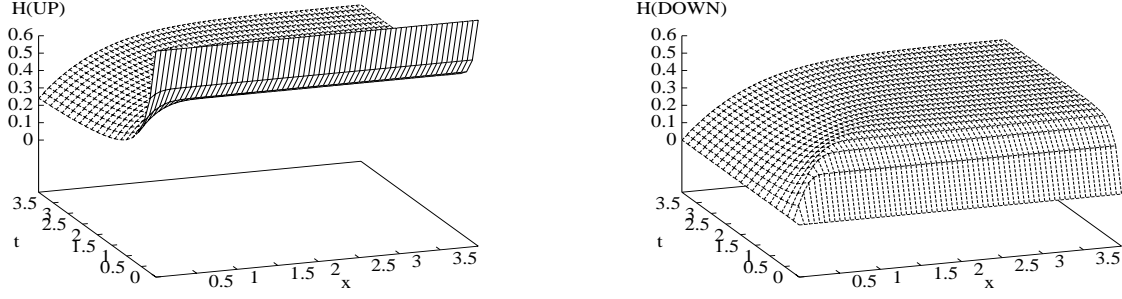
$$\mathbf{P}_{ij} \geq 0 \quad \forall i, j$$

Figure 5: Numerical Solution of Breakdown Model

Proof:

It is sufficient to show

$$\tilde{\mathbf{Q}}_{ii}, \mathbf{W}_{ii} \leq 0, \quad \text{and} \quad \tilde{\mathbf{Q}}_{ij}, \mathbf{W}_{ij} \geq 0 \quad i \neq j \tag{18}$$

$$-\tilde{\mathbf{Q}}_{ii} \geq \tilde{\mathbf{Q}}_{ij}, \ \tilde{\mathbf{Q}}_{ji} \quad i \neq j \tag{19}$$

We have

$$\bar{\mathbf{Q}}_{ii}, \mathbf{W}_{ii} \leq 0, \qquad \bar{\mathbf{Q}}_{ij}, \mathbf{W}_{ij} \geq 0 \quad i \neq j$$

$$\bar{\mathbf{Q}}_{ii} = -\sum_{i \neq j} \mathbf{Q}_{ij}, \qquad \bar{\mathbf{W}}_{ii} = -\sum_{i \neq j} \mathbf{W}_{ij}$$

since $\bar{\mathbf{Q}}$ is a CTMC and $\mathbf{W}$ represents an upwind discretization. Equations (18) and (19) follow directly from $\tilde{\mathbf{Q}} = \bar{\mathbf{Q}} + \mathbf{W}$.

6 Numerical Examples

6.1 Machine Breakdown Model

For our first illustration of the behaviour of an FSPN we choose the model of section 4.2. We consider the case of one processor only, and parameter values of $\lambda_1 = 2$, $\lambda_2 = 0$, $\mu = 3$, $r = 1$ and $d = 2$. In this case, both $\mathbf{R}$ and $\mathbf{Q}$ are independent of $\vec{x}$. We solve the equations for the distribution (6) in the range $0 \leq t \leq 4$, $0 \leq x \leq 4$. We discretized with stepsizes $\Delta x = 1/64$ and $\Delta t = 1.0e - 4$. The numerical results for this problem are shown in Figure 5.

For this simple case, it can be shown via Laplace transforms that the steady state solution is given by
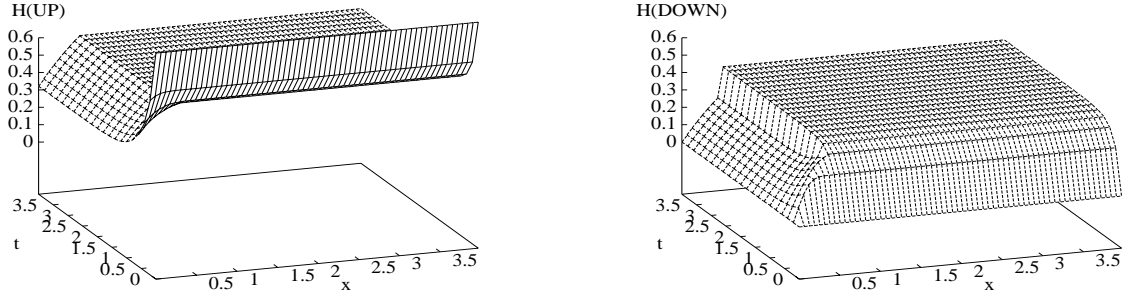
$$H(x, UP) = 0.4(1 - e^{-x})$$

15

Figure 6: Numerical Solution of Modified Breakdown Model

$$H(x, DOWN) \quad = \quad 0.2(3 - 2e^{-x})$$

The transient solution values obtained at $t = 4$ agree closely with these results.

Now we modify the example to allow dependency of the matrices $\mathbf{R}$ and $\mathbf{Q}$ on the fluid level X. First we extend the model to contain a backup server which is only used when the primary server is down and the load level reaches a value of 0.5. We thus have

$$d = \begin{cases} 2 & m_d = UP \\ 0 & m_d = DOWN, \ x < 0.5 \\ 2 & m_d = DOWN, \ x \geq 0.5 \end{cases}$$

The dependency of $\mathbf{Q}$ on x is obtained by setting

$$\lambda = \lambda_1 + \lambda_2(1 - e^{-\alpha x})$$

choosing $\alpha = 1$, $\lambda_1 = 1$ and $\lambda_2 = 2$. The results of this computation are depicted in Figure 6. Note the discontinuity and change of slope at $x = 0.5$, when the backup server is started.

Figure 7 shows the results for the unmodified breakdown model, using the equation in transformed coordinates (17). Here integration until $t = 24$ has been performed. Note that the solution obtained for $H(DOWN)$ at $t = 24$ appears as a straight line through the origin with gradient 0.4, corresponding to the analytic solution of Equation (20).

6.2   ATM Switch Model

We now consider the multiplexed network switch model considered in section 4.1. We set the number of sources K to one, parameters $B1 = 3000$ and $B = 6000$. High priority packets arrive at the rate 6×10^3 and low priority packets at a rate of 4×10^3 per second when the source is in state 1, and at rates 4×10^3 and

Figure 7: Probability distributions for breakdown model in transformed coordinates. Left: Server up; Right: Server down.


Figure 8: Probability distributions for ATM switch model. Left: Source state 1; Right: Source state 2.

2×10^3 per second respectively when the source is in state 2. The switch can process both types of packet at a rate of 5×10^3 per second. The exponentially distributed rates of change between high and low priority packet generation are $\lambda_1 = 0.5$ and $\lambda_2 = 0.5$. The results for the number of packets in the buffer are shown in Figure 8. The left and right results at $t = 10$ qualitatively match those of Elwalid and Mitra [10], Figure 2, right and center, respectively. An appropriate choice of parameters also yields a result, not illustrated here, which is similar to [10], Figure 2, left.

7 Conclusion

We have defined a new class of stochastic Petri nets by introducing places with continuous tokens and arcs with fluid flows. This new class of fluid stochastic Petri nets (FSPNs) should be useful in modeling stochastic fluid flow systems, and may also be useful in modeling processes that control physical systems. Our model formulation permits the discrete and continuous parts to affect each other, endowing FSPNs with the ability

to both control the fluid flow, and have the discrete control decisions be affected by observed fluid flow. We have provided formal definition of FSPNs and developed the rules for their dynamic evolution. We have derived coupled systems of partial differential equations for the transient and the steady-state behavior of FSPNs. Spectral representation of the FSPNs with a single continuous place can be adapted from the literature on stochastic fluid flow models. We have presented a number of examples illustrating the modeling power of FSPNs, have considered issues arising in the numerical solution of the dynamical equations, and have provided numerically solved examples.

References

[1] M. Ajmone Marsan, G. Balbo, and G. Conte, *Performance Models of Multiprocessor Systems*, MIT Press, Cambridge, MA, 1986.

[2] M. Ajmone Marsan and G. Chiola, "On Petri nets with deterministic and exponentially distributed firing times," *Lecture Notes in Computer Science*, Vol. 266, pp. 132–145. Springer-Verlag, 1987.

[3] D. Anick, D. Mitra and M. Sondhi, "Stochastic Theory of Data-Handling Systems," *The Bell System Technical Journal,* Vol. 61, No. 8, pp. 1871-1894, Oct. 1982.

[4] O. Axelsson, A class of A-stable methods. *BIT* Vol. 9, pp. 185–199, 1969.

[5] R. E. Bank, W. M. Coughran, W. Fichtner, E. H. Grosse, D. J. Rose, and R. K. Smith, Transient simulation of Silicon devices and circuits. *IEEE Transactions on Computer-Aided Design* Vol. 4, pp. 436–451, 1985.

[6] C. G. Cassandras, *Discrete Event Systems: Modeling and Performance Analysis*, Aksen Associates, Holmwood, IL, 1993.

[7] G. Ciardo, J. K. Muppala, and K. S. Trivedi, "SPNP: Stochastic Petri Net Package," *International Conference on Petri Nets and Performance Models*, Kyoto, Japan, pp. 142–150, 1989.

[8] G. Ciardo, A. Blakemore, P. F. J. Chimento, J. K. Muppala, and K. S. Trivedi, "Automated generation and analysis of Markov reward models using Stochastic Reward Nets," in C. Meyer and R. J. Plemmons, editors, *Linear Algebra, Markov Chains, and Queueing Models*, Vol. 48 of *IMA Volumes in Mathematics and its Applications*, Springer-Verlag, 1992.

[9] G. Ciardo, J. K. Muppala, and K. S. Trivedi, "Analyzing concurrent and fault-tolerant software using stochastic Petri nets," *J. Par. and Distr. Comp.*, Vol. 15, No. 3, pp. 255–269, July 1992.

[10] A. I. Elwalid and D. Mitra, "Statistical Multiplexing with Loss Priorities in Rate-Based Congestion Control of High-Speed Networks," *IEEE Transaction on Communications*, Vol. 42, No. 11, pp. 2989-3002, November 1994.

[11] R. K. Iyer, D. J. Rosetti and M. C. Hsueh, "Measurement and Modeling of Computer Reliability as Affected by System Activity," *ACM Transactions on Computer Systems*, Vol. 4, pp. 214 - 237, August 1986.

[12] R. J. Karandikar and V. G. Kulkarni, "Second-order fluid flow models: Reflected Brownian motion in a random environment," *Operations Research*, Vol. 43, No. 1, pp. 77–88, 1995.

[13] D. Mitra, "Stochastic Theory of Fluid Models of Multiple Failure-Susceptible Producers and Consumers Coupled by a Buffer," *Advances in Applied Probability,* Vol. 20, pp. 646-676, 1988.

[14] T. Murata, "Petri Nets: Properties, analysis and applications," *Proceedings of the IEEE*, Vol. 77, No. 4, pp. 541–579, April 1989.

[15] S. V. Patankar, *Numerical Heat Transfer and Fluid Flow*, McGraw-Hill, New York, 1980.

[16] J. L. Peterson, *Petri Net Theory and the Modeling of Systems*, Prentice-Hall, Englewood Cliffs, NJ, USA, 1981.

[17] T. Robertazzi, *Computer Networks and Systems: Queueing Theory and Performance Evaluation*, Springer-Verlag, 1990.

[18] K. S. Trivedi, V. G. Kulkarni, "FSPNs: Fluid Stochastic Petri Nets," *Lecture Notes in Computer Science*, Vol 691, M. Ajmone Marsan (ed.), *Proc. 14th International Conference on Applications and Theory of Petri Nets*, Springer-Verlag, Heidelberg, pp. 24-31, 1993.

[19] N. Viswanadham and Y. Narahari, *Performance Modeling of Automated Manufacturing Systems*, Prentice-Hall, Englewood Cliffs, NJ, 1992.