Computation of the Distribution of Accumulated Reward with Fluid Stochastic Petri-Nets[1]

Graham Horton[2], Kishor S. Trivedi[3]

[1]Institut für Informatik III,
Universität Erlangen-Nürnberg, Martensstr. 3,
91058 Erlangen, Germany.

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

Abstract

We describe the recently introduced Fluid Stochastic Petri-Nets as a means of computing the distribution of the accumulated rate reward in a GSPN. In practice, it is the expected value of a reward which is computed, a quantity which is dependent solely on the solution of the underlying Markov chain. Until now, the instantaneous reward rates have been a function of the GSPN marking only, and the Markov chain itself was not influenced by the development of the reward value. In this paper it is shown that FSPNs may be used to simulate GSPN reward models and that they allow an important generalization in that both the firing rates of the GSPN and the reward rate may depend on the current reward value. Example models of repairable systems are used to demonstrate the additional capabilities.

1 Introduction

We are interested in this paper in the use of rate rewards applied to Markov chains to model the reliability and/or performance of computer and other systems. In particular we consider the case where the Markov chain is represented by a Generalized Stochastic Petri-Net (GSPN).

Rewards are widely used in conjunction with Markov chain models of systems as a means of specifying measures of interest that depend upon the stochastic behaviour of these systems. In the case where the Markov chain is described by a GSPN, rate rewards are associated with specific markings. GSPN models augmented with rewards will be called GSPN-reward models. Software packages such as SPNP [1] are available which allow the specification and solution of GSPN-reward models[4]. Typically the user is limited to obtaining the expected value of an accumulated reward at a given instant of time, which is described by an ordinary differential equation (ODE) similar to that of the Markov chain itself. Often, however, the distribution of the reward may be of more interest to the user. This distribution is defined by a linear system of hyperbolic partial differential equations (PDE), which is, of course, substantially more challenging to solve numerically. The particular property of rate rewards which we address in this paper is the fact that they are passive quantities, i.e. they influence neither the behaviour of the underlying Markov chain nor the reward rates; indeed, the Markov chain may in principle be solved first independently, and the reward measures then derived from the Markov chain solution. We will show that FSPN models of rate rewards allow a greater degree of modeling power in this regard. We do not consider impulse rewards in this paper.

Fluid Stochastic Petri-Nets (FSPNs) were introduced [4, 3] as a generalization of the well-known GSPNs to allow certain places to contain a non-negative, continuous − as opposed to integer − number of tokens. Semantics for the corresponding arcs and transitions have been defined which regulate the flow of (fluid) tokens in the net. Dependencies between the fluid and discrete parts of the net are also allowed. The mixed continuous/discrete modelling paradigm of FSPNs leads to a correspondingly discrete and continuous-state Markov process whose transient behaviour is described by a linear system of hyperbolic PDEs. These equations are a more general form of the equations of the distribution of accumulated reward. This added generality allows us to conclude that reward models described by FSPNs allow a greater degree of modelling power than has been typically used previously.

In the following section, we briefly review GSPN reward models. In section 3 we describe FSPNs and their governing equations and show using a performability example how they can be used to model rewards. In section 4 we present two small models of systems with failures and repairs that demonstrate the added modelling power. Section 5 concludes the paper.

---

[4]GSPN-reward models accepted by SPNP also allow other extensions such as variable cardinality arcs, enabling functions, etc. For this reason models accepted by SPNP are also called stochastic reward nets.

2 Rate Rewards in GSPNs

Rate rewards are an often-used addition to discrete-state stochastic models which greatly enhance the usefulness of these techniques. With a reward, the user can define the rate of accumulation of a quantity of interest associated with each state of the process. Reward measures of interest are the instantaneous reward rate at any point in time or the expected value and distribution of the accumulated reward. It is the latter measure which concerns us here. One of the most well-known uses of rewards in the reliability modelling field is the so-called *performability* [6]. In a system which provides computational resources and which is subject to failures, reward rates may be defined as the computational power (= computational units processed per unit of time) available in each system state. The performability is then defined as the integral of the total reward rate over a defined time interval.

Consider the n-state Markov chain with unknown row vector π and generator matrix Q, whose transient behaviour is given by the ordinary differential equation

$$\frac{d\pi}{dt} = \pi Q \quad . \tag{1}$$

Assign to the i-th state of the Markov chain the non-negative reward rate r_i and let $\rho(t)$ be the value of the accumulated reward by time t. Define the diagonal matrix $R = diag(r_1 \ldots r_n)$ and the row vector

$$H(x,t) = (H_1(x,t) \ldots H_n(x,t)) \quad ,$$

where $H_i(x,t)$ is the cumulative distribution function of the accumulated reward, i.e. the probability that at time t, $\rho \le x$ and the Markov chain is in state i. Then the vector H is given by the solution of the linear system of hyperbolic PDEs [7, 8]

$$\frac{\partial H}{\partial t} = -\frac{\partial H}{\partial x} R + HQ \quad . \tag{2}$$

The domain of Equation (2) is $(x,t) \in ((0,\infty),(0,\infty))$. Initial conditions are obtained by stipulating an initial state probability vector for the Markov chain and the cdf of the initial reward value. Boundary conditions for Equation (2) are of both Dirichlet and von Neumann type and are given by

$$H_i(0,t) \;=\; 0 \qquad \text{if } r_i > 0, \; t > 0 \tag{3}$$

$$\lim_{x \to \infty} \frac{\partial H_i}{\partial x} \;=\; 0 \tag{4}$$

for all $1 \le i \le n$. Equation (3) states that the probability of the reward value being equal to zero after a non-zero time has elapsed must be zero for all discrete states with positive reward rates. Condition (4) follows from the fact that the functions H_i are cdfs and must therefore approach asymptotic values at infinity. Note that equations (1) and (2) are related via

$$\pi(t) = H(\infty,t) \quad .$$

Both the matrix of reward rates R and the generator matrix Q are independent of x. This means both that the reward rate is fixed for any state of the Markov chain and that the Markov chain is independent of the value of the reward. Thus we can solve for the Markov chain and the reward distributions independently using (1) and (2) respectively. Since the reward values have no influence on the development of future reward values or the Markov chain, we may think of them as "passive" rewards. Conversely, allowing the dependencies $R = R(x)$ and $Q = Q(x)$ would be equivalent to letting the instantaneous reward rates and the Markov chain itself (i.e. the firing rates of the timed transitions of the GSPN) be influenced by the current accumulated reward value. It is this added flexibility which is possessed by FSPNs and which will be described in the next section.

3   Fluid Stochastic Petri-Nets

Fluid Stochastic Petri-Nets (FSPNs) were introduced [4, 3] as an extension to the well-known GSPNs. In addition to containing all the features of GSPNs, FSPNs allow certain "fluid" places to contain a *continuous* number of tokens, represented by a non-negative, real number. These continuous (fluid) tokens may enter and leave the fluid places via (fluid) arcs, whose cardinality, which may also be a non-negative real quantity, determines the rate of flow of the tokens. Transitions on fluid arcs may be enabled or disabled, permitting or shutting off the flow of fluid tokens accordingly. Fluid tokens flow continuously and non-stochastically, and the stochastic process describing the change in fluid levels over time is a Markov process with zero variance parameter. Fluid tokens may be considered to be analogous to a physical fluid, arcs correspond to pipes (and arc cardinality to pipe cross-sectional area), and fluid places correspond to storage tanks or buckets. However, since the fluid part of the net may be influenced by the discrete GSPN portion, the overall behaviour of the fluid levels is of a probabilistic nature.

The state of the net at any time t is described by its marking vector $M(t)$. In an FSPN with F fluid places and D discrete places $M(t)$ is naturally defined to be

$$M(t) = (X_1(t), \ldots, X_F(t), n_1(t), \ldots, n_D(t))$$

where $X_k(t) \in \mathbb{R}^+ \cup \{0\}$ denotes the fluid level in the k-th fluid place and $n_i(t) \in \mathbb{N} \cup \{0\}$ denotes the number of discrete tokens in the i-th discrete place at time t. We allow the arc cardinalities and the firing rates and the enabling functions of the timed transitions to be dependent on the current state of the FSPN, as defined by the marking $M(t)$.

The motivation behind FSPNs is both to allow a fluid approximation to a discrete model in the manner, for example, of [2], and to allow the modelling of systems that inherently contain continuous quantities. In the former case, the large discrete state space generated by the accumulation of tokens in a place representing a buffer, for example, may sometimes be more efficiently
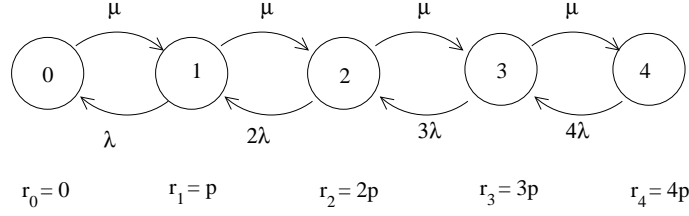
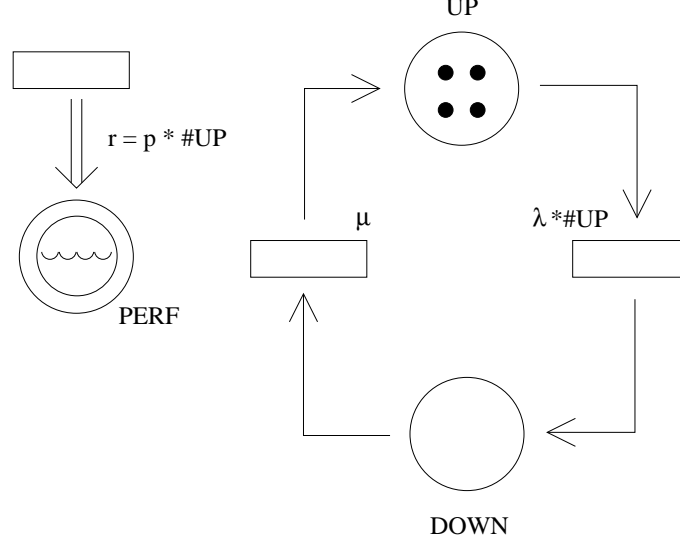Figure 1: Markov chain of performability model

Figure 2: FSPN representation of a simple performability model

represented by a continuous function. In the latter case, we are able to model a much larger class of systems, in particular those with a mixed discrete and continuous state space. FSPNs will also allow the representation of non-Markovian Stochastic Petri-Nets and, as described here, the modeling of the distribution of accumulated reward.

In order to illustrate FSPNs we give a small performability example. Consider a multiprocessor system consisting of four processors, each delivering a performance equal to p. Each processor may fail at rate λ and undergo a repair at rate μ. When both repair and failure times are exponentially distributed the stochastic process describing the multiprocessor system is the five-state birth-death Markov chain of Figure 1. The reward rate representing the instantaneous performance delivered when the Markov chain is in state i is simply $r_i = p * i$, since the index of each state corresponds to the number of functioning processors in the model.

Figure 2 shows an FSPN that represents this performability model. In order to distinguish them from their discrete counterparts, we draw fluid places as two concentric circles and fluid arcs as double arrows. The states of the multiprocessor are represented by the discrete places UP and DOWN, and the reward by the fluid place PERF. The amount of fluid in PERF corresponds

4

to the accumulated reward. Fluid accumulates in the fluid place at the marking-dependent rate r_i. The behaviour of this FSPN corresponds exactly to that of the GSPN reward model.

In the following, we restrict ourselves to FSPNs with one fluid place only. The equations for the case with more than one fluid place may be found in [3]. We denote the fluid level by X, time by t, and index the discrete markings of the FSPN (the states of the Markov chain) using i. We may generate the state-space of the discrete part of the net and eliminate any vanishing markings that may be present, resulting in an n-state Markov chain. $m(t)$ denotes the discrete marking occupied by the net at time t. We may then define the cdf

$$H_i(t, x) = P(X(t) \leq x, m(t) = i)$$

as the probability that at time t, the fluid level will be less than or equal to x, and that the discrete part of the net will be in the marking i. We then set

$$H(t, x) = (H_1(t, x) \ldots H_n(t, x)) \quad .$$

Q denotes the generator matrix of the underlying Markov chain and R the diagonal matrix of fluid flow rates. In the case that fluid flows simultaneously into and/or out of a fluid place via more than one fluid arc, r_i denotes the net rate of fluid inflow or outflow. We allow the rates of the Markov chain to be functions of the fluid levels, giving $Q = Q(x)$ and the fluid flow rates to also be dependent on fluid levels, giving $R = R(x)$. When the latter is the case, we write $r_i(x)$ to indicate a specific net fluid flow rate when the fluid level is equal to x.

The equations for the cumulative probability distributions H are

$$\frac{\partial H}{\partial t} + \frac{\partial}{\partial x}(HR) = HQ - \int_0^x H \frac{\partial Q}{\partial x} \, dx \tag{5}$$

Note that the integral term only appears when Q depends on x and that R may be taken out of the partial derivative when it is independent of x. The boundary conditions of Equation (5) are

$$H_i(t, 0) \;\; = \;\; 0 \;\; \text{if} \;\; r_i(0) > 0 \tag{6}$$

$$\lim_{x \to \infty} \frac{\partial H_i}{\partial x} \;\; = \;\; 0. \tag{7}$$

The boundary condition (6) states that there is zero probability of there being no fluid in the fluid place if there is a net fluid inflow. Condition (7) stems from H_i being a cdf and therefore having an asymptotic value at infinity. The domain of Equation (5) is $(x, t) \in ((0, \infty), (0, \infty))$, although in practice we will only be interested in the finite domain $(0 \ldots tmax, 0 \ldots xmax)$. The initial conditions for Equation (5) are obtained from the cdf of the required initial state of the system.

Comparing Equations (2) and (5) and their boundary conditions, it is easy to see that the former is a special case of the latter obtained by restricting the matrices R and Q to be independent of x.

5

The FSPN of figure 2 corresponds exactly to the reward model and we may solve equation (5) to obtain the distribution of the accumulated reward of the performability model.

We now see that allowing $R = R(x)$, i.e. letting the fluid flow rate be a function of the current fluid level in a place corresponds to a reward rate which is a function of the current reward value. Similarly, allowing $Q = Q(x)$, i.e. rates in the discrete part of the net to depend on the fluid level corresponds to letting the transition rates of the reward model be a function of the current reward level. As far as we know, neither of these possibilities has hitherto been considered. Because they can influence the reward and transition rates, we could think of these as "active" rewards. We feel that this added generality allows new modelling opportunities, which we attempt to illustrate in the next section.

4   Application Examples and Solutions

In this section we present two example GSPN reward models which may be represented by an FSPN but not by standard means. We give in each case the corresponding FSPN and a numerical solution of the model equations. The numerical solutions presented are not intended to be realistic, but merely to illustrate the form of the results. The numerical solutions were obtained by discretizing the $\partial/\partial x$ terms with a first-order upwind scheme and the time derivative by the Forward Euler method.

4.1   TMR system with failures and repairs

Consider a computer system that is embedded in a critical application. In such a case, one might decide to utilize triple modular redundancy (TMR), whereby three identical computers are used and a voter compares the results computed by each. As long as at least two results coincide, these are taken to be correct.

We consider as an example a TMR system which is required to survive a given mission time, i.e. the system should be operational until the mission time is reached, and its status thereafter is unimportant. We allow the possibility of repair of each computer. We consider the system to be operational if at least two of the three computers are operational. In order to save the expense of an unnecessary repair, we allow the option of not repairing a processor if the mission time has almost been reached. More precisely, we define a threshold time τ, after which in the situation of one down and two up processors, the down processor will no longer be repaired. This model is similar to that considered in [5], for which a Markov decision process is defined. There the decision to repair a failed processor may be taken at any time, and an optimization problem is defined whether or not a repair should take place at any point in time.
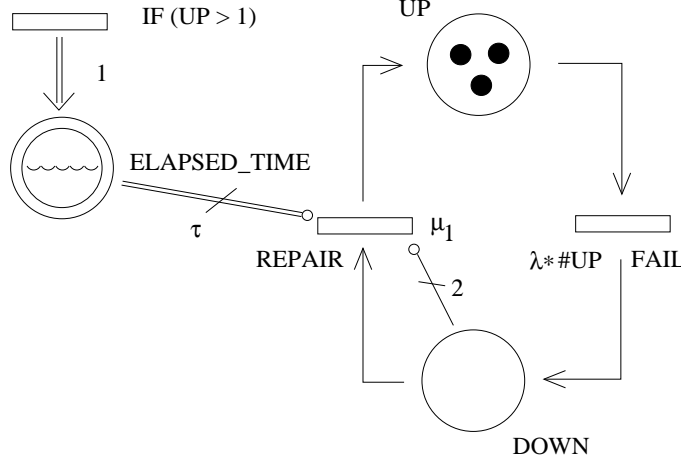
6

Figure 3: FSPN Representation of TMR Model

The Markov chain representing the state of the TMR system consists of four nodes, where the i-th node represents the state in which i processors are operational. There are transitions $i \rightarrow i - 1$ of rate λi, and a transition $2 \rightarrow 3$ representing a repair, which, however, ceases to exist when the system has survived until time $t > \tau$. We thus have the situation that Markov chain representing the state of the TMR system is a function of the reward value. Therefore, this model cannot be represented by a classical reward model. Alternatively, in this simple case, we recognize that this is a non-homogeneous Markov chain, and that we may use a reward represented by an FSPN to model such a problem.

The role of the reward in representing this situation is obvious: the instantaneous reward rates are defined by

$$r_i = \begin{cases} 1 & i = 2, 3 \\ 0 & i = 0, 1 \end{cases} ,$$

in other words, the reward value is simply a clock that measures the elapsed time that the system has survived.

The corresponding FSPN representation of the reward model is very simple: the net contains one fluid and two discrete places (ELAPSED_TIME, UP, DOWN), with the initial marking (0.0, 3, 0). Figure 3 shows the FSPN, where the time to failure of each processor is exponentially distributed with rate parameter λ and the time to repair is exponentially distributed with rate parameter μ_1. However, repair may be shut-off altogether as described earlier. Thus the overall repair rate is denoted by μ, where

$$\mu = \begin{cases} \mu_1 & \text{if (ELAPSED\_TIME} < \tau \text{ and DOWN} < 2) \\ 0 & \text{otherwise} \end{cases}$$

Note the use of the fluid inhibitor arc to prevent repairs after the mission time has been reached.
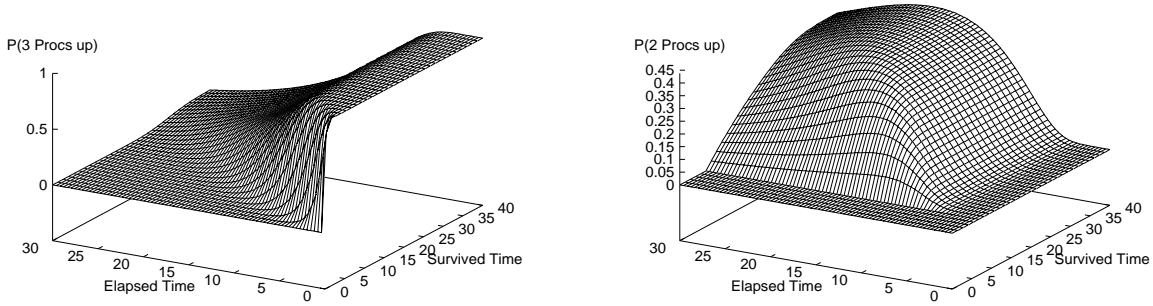
Figure 4: Numerical Solution of TMR Model

Figure 4 contains a portion of the results of a numerical solution of this model. Shown is the probability that the survived time is less than or equal to the elapsed time and that the number of operational processors is three (left picture) or two (right picture). The results for zero and one operational processors are omitted. Note that, in order to obtain best presentation of the curves, the diagrams have been rotated so that elapsed time goes from right to left (0 to 30 years) and survived time (0 to 40 years) into the page. The parameters used are $\mu_1 = 365$, (i.e. a repair takes one day on average), $\tau = 7$ years, and $\lambda = 12$ (failure rate is once per month). Note that these parameters are hardly realistic for a mission time of 10 years, say. They merely serve to illustrate the shape of the curves.

4.2   Cooling System

We consider as a second example a case in which the model itself contains a continuous quantity which is modeled by a reward. Here, physical considerations require that the reward rate be a function of the current reward value. In addition, we allow the Markov chain's behaviour to also depend on the current reward value.

Consider the cooling system of a power station, in which a liquid coolant is heated by an energy source and is cooled by refrigeration. The refrigeration system is subject to failures and repairs. Safety considerations demand that the coolant temperature T not exceed a critical value τ_2. Parameters are such that the refrigeration unit, when operational, can reduce the temperature of the coolant at a greater rate than that at which it is increased by the energy source. Repairs are assumed to be carried out at either a standard rate μ_1, or, if the coolant temperature exceeds a threshold temperature τ_1, at an increased rate μ_2:

$$\mu(T) = \begin{cases} \mu_1 & \text{if } T < \tau_1 \\ \mu_2 & \text{otherwise} \end{cases} \quad .$$
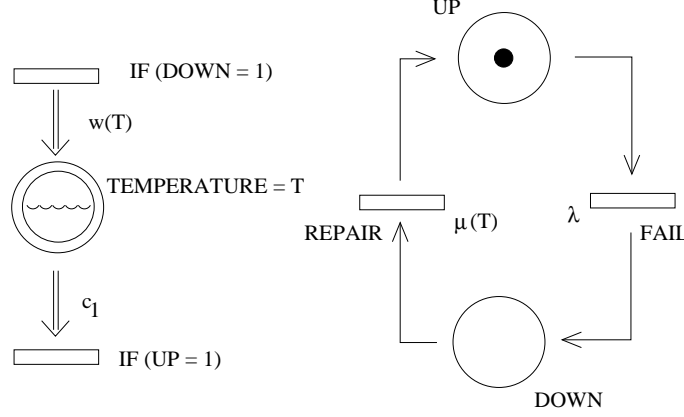
8

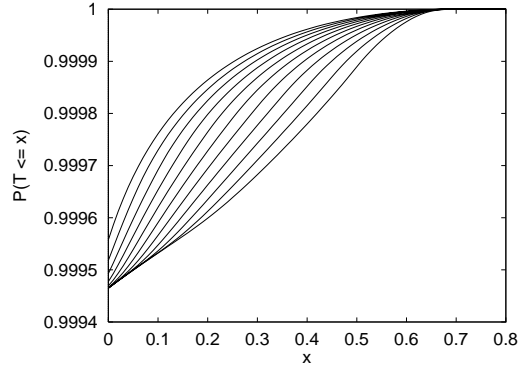Figure 5: FSPN Representation of Cooling System Model

Figure 6: Steady-State Solution $P(T \leq x)$ of Cooling Model. $\tau_1 = 0$, (0.05), 0.5

We model the coolant temperature – the continuous quantity – by a reward, and thus by a fluid place in the FSPN. The (constant) temperature of the energy source is τ_0 and the temperature of the unrefrigerated coolant rises at a rate proportional to the temperature gradient: $w(T) = dT/dt = \tau_0 - T$. The refrigeration unit cools at the rate

$$ c = \begin{cases} c_1 & \text{if } (\text{UP} = 1) \\ 0 & \text{otherwise} \end{cases} . $$

Note that we have extended the reward model to allow a negative reward rate when $c > w$.

We are interested in computing the cumulative probability distribution of the coolant temperature at various points in time. This will allow us to determine, for example, the probability that the critical temperature will be exceeded or that the emergency repair must be initiated. Figure 6 shows a numerically computed steady-state solution, i.e. the probability that, at steady-state, the coolant temperature T is less than or equal to x. The parameter values used are $c_1 = 0.8$, $\lambda = \frac{1}{3650}$, $\mu_1 = 0.2$, $\mu_2 = 2.0$, $\tau_0 = 0.75$. Results for various threshold values τ_1 from 0.0 (lowest curve) to 0.5 (uppermost curve) in steps of 0.05 are shown.

9

5   Conclusions

We have shown how the recently introduced Fluid Stochastic Petri-Nets may be used to represent GSPN rate reward models. The transient solution of the FSPN corresponds to the distribution of the accumulated reward. Both are described by a system of linear hyperbolic PDEs, whereby the FSPN equations generalize the reward equations. By allowing transition rates of the discrete part of the FSPN to be dependent on a fluid level, we can simulate an "active reward", i.e. one whose value can influence the behaviour of the Markov chain. By allowing the rate of fluid flow to have the same dependency, we can model a reward-dependent instantaneous reward rate. This added functionality allows the construction and solution of a larger class of reward models.

References

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

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

[3] G. Horton, V. Kulkarni, D. Nicol, K. Trivedi: Fluid Stochastic Petri Nets: Theory, Applications, and Solution, *ICASE Technical Report* 96-5. Submitted for publication to *European Journal of Operations Research*.

[4] V. Kulkarni, K. Trivedi: "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.

[5] H. de Meer, K. Trivedi, M. Dal Cin: Guarded repair of Dependable Systems, *Theoretical Computer Science* 128 (1994).

[6] J. F. Meyer: On evaluating performability of degradable computing systems, *IEEE Trans. Comput.*, Vol. C-29, No.8, Aug. 1980.

[7] K. Pattipati, Y. Li, H. Blom: A Unified Framework for the Performability Evaluation of Fault-Tolerant Computer Systems. *IEEE Trans. Comp.*, Vol. 42 No. 3, March 1993.

[8] A. Reibman, R. Smith, and K. S. Trivedi. Markov and Markov reward model transient analysis: An overview of numerical approaches. *European J. Oper. Res.*, 40:257–267, 1989.