On Markov Reward Modelling with Second Order FSPNs

Katinka Wolter

Technische Universität Berlin,
Institut für Technische Informatik,
Franklinstr. 28/29
10587 Berlin, Germany.
katinka@cs.tu-berlin.de

Andrea Zisowsky

Technische Universität Berlin,
Fachbereich Mathematik,
Strasse des 17. Juni 136,
10623 Berlin, Germany.
zisowsky@math.tu-berlin.de

May 12, 1999

Abstract

In this paper second order fluid stochastic Petri nets (FSPNs) will be used for modelling reward in a performability model. Two variations of the performability model are presented. In the first one two fluid places are used, one of which represents the profit (reward) obtained by operating the system and the other one the buffer, that is approximated continuously. In the second model only one fluid place is used, which represents the costs (negative reward) arising due to repair. The costs increase continuously while the system is in state of repair (which is a rate reward in the model). Additional costs incur each time the buffer fails (which is an impulse reward in the model). With a numerical solution algorithm the distribution of the reward and its mean are computed.

1 Introduction

Fluid Petri net models have gained much attention over the last years and a number of different classes of hybrid Petri nets have been proposed. For the specification and qualitative analysis the autonomous non-timed and timed Petri nets (ACPNs) [DA93] have been used. Performance and dependability evaluation has been carried out with fluid stochastic Petri nets (FSPNs) [TK93, HKNT98]. The discrete part of an FSPN taken on its own is a generalized stochastic Petri net (GSPN) [ABC+95] and the continuous part alone is equivalent to a fluid queueing system [New71, Kob74]. In FSPNs the continuous model part is represented by fluid places and fluid arcs. The level in the fluid places changes at deterministic rate. The stochastic behaviour of the model is due to the discrete part and the dependencies between discrete and continuous submodel. In this paper a specific class of FSPNs is considered, the second order FSPNs. In second order FSPNs the fluid places are filled and emptied at normally distributed rate. Accordingly, the continuous model part has probabilistic behaviour too. The underlying equations are second order partial differential equations, hence the prefix *second order*.

In [Hor96] FSPNs have been used for modelling accumulated rate reward. Reward structures are added to Markov models, giving Markov reward models, in order to obtain measures of interest from these models. This is a concept widely known for discrete Markov models. Two kinds of rewards can be distinguished, rate reward and impulse reward. Rate reward is the reward obtained while the model is in a certain state. Impulse reward is achieved with the completion of an action, that is with the transition from one model state to the next. Both can be instantaneous transient or steady state rewards, which is the value of the reward at some point in time, or accumulated reward, the reward that has accumulated over a period of time [SM91].

In most situations one is limited to obtaining the mean of the accumulated reward not its distribution. The mean accumulated reward is computed through integration over the state probabilities of the Markov model multiplied with some reward rate for each state. In states where no reward is obtained the reward rate equals zero. In the computation of the mean impulse reward the state transitions rate, that gives the impulse is a factor in the reward rate. It has been shown in [Hor96] that the distribution of accumulated reward is computed by solving a differential equation like those underlying an FSPN.

In the FSPN a fluid place holds the accumulated reward. In each state where rate reward is accumulated, the fluid place is filled at the reward rate. The main advantage of modelling reward with a fluid model is, that the full probability density is obtained automatically and from it the mean can be computed.

In this paper we also consider impulse reward, which has not been done before. With a discrete state transition a non-zero amount of reward is obtained and added to the reward accumulated so far in the fluid place. In the example costs will be modelled.

The second novelty in this paper is that a model with two fluid places will be solved numerically. Simulation for these models has been done earlier [CNT97]. In the example the first fluid place represents a continuously approximated buffer while the second one represents the reward. Hence, the mixed continuous and discrete state space has two continuous dimensions. A numerical method is applied for solving the underlying differential equations.

In the next section the second order FSPN formalism is introduced. In section 3, the main part of the paper, the two models will be discussed. Both are a modified version of the performability model in [Mey82]. For both models the underlying differential equations are given. In the algorithm used for their solution the equations are generated automatically from a (textual) description of the Petri net. The transient evolution of the probability density is given as a result for both examples. With section 4 the paper concludes.

2 The FSPN Formalism

For the definition of second order FSPNs the common notations for GSPNs [ABC$^+$95] and for FSPNs [HKNT98] are taken over and extended as follows.

A second order FSPN is formally defined as a 9-tuple FSPN $= (\mathcal{P}, \mathcal{T}, \mathcal{A}, m_0, g, \lambda, r, b, w)$ where the set of places ($\mathcal{P} = \mathcal{P}_c \cup \mathcal{P}_d$) is divided into the *fluid (continuous)* and the *discrete* places. Discrete places (the elements of $\mathcal{P}_d$) are drawn as single lined circles and hold an integer number of *tokens*, whereas the continuous, or fluid, places (the elements of $\mathcal{P}_d$) are drawn as two concentrical circles and they hold a real-valued amount of *fluid*. The set of transitions $\mathcal{T} = \mathcal{T}_E \cup \mathcal{T}_I$ is composed of the exponentially distributed, and the immediate transitions, respectively. Exponentially distributed transitions are drawn as empty rectangles and the immediate transitions as thin bars.

The set of arcs $\mathcal{A} = \mathcal{A}_d \cup \mathcal{A}_c \cup \mathcal{A}_j$ is divided into three subsets, the discrete arcs (the elements of $\mathcal{A}_d$), the continuous, or fluid, arcs ($\mathcal{A}_c$) and the jump arcs ($\mathcal{A}_j$). The discrete and the jumps arcs are drawn as single lined arrows, whereas the fluid double-lined arcs are drawn like pipes. The jump arcs only connect transitions and fluid places, and vice versa. Immediate transitions can be connected to a fluid place by a jump arc only if the immediate transition is also connected to a discrete place. In other words, the fluid jumps caused by the firing of an immediate transition can not occur without a discrete model part connected to that same immediate transition. Opposed to this, jumps caused by the firing of a timed transition may stand alone.

The *marking* $\mathbf{m} = (s, \mathbf{z}) \in \mathcal{M}$ consists of a discrete part $s = (\#p_i, i \in \mathcal{P}_d)$, where $\#p_i$ denotes the number of discrete tokens in the i-th discrete place p_i, and a continuous part, a vector representing the fluid level in each fluid place, $\mathbf{z} = (z_k, k \in \mathcal{P}_c) \in \mathcal{Z}$. The initial marking is $\mathbf{m}_0 = (s_0, \mathbf{z}_0)$. The total number of discrete states is S, the set of all discrete states is denoted $\mathcal{S}$.

$\mathrm{r}_{t,p}$, $\mathrm{r}_{p,t} : \mathcal{A}_c \times \mathcal{M} \rightarrow \mathbb{R}^2$ is the flow rate function along the fluid arc connecting the timed transition t and the fluid place p or vice versa. The flow rate $r = (\mu, \sigma^2)$ is a normally distributed random variable specified by expectation μ and variance σ^2. Let Z_k be the fluid level in place $k \in \mathcal{P}_c$. The change rate of the fluid level is again a normally distributed random variable with

$$\frac{dZ_k}{dt} = \sum_{t \text{ enabled in } \mathbf{m}} r_{t,p}(\mathbf{m}) - \sum_{t \text{ enabled in } \mathbf{m}} r_{p,t}(\mathbf{m}).$$

Reflecting barriers assure, that the fluid level in each place stays within its range $[z^{\min}, z^{\max}]$.

The jump height function $\mathrm{b}_{t,p}$, $\mathrm{b}_{p,t} : \mathcal{A}_j \times \mathcal{M} \rightarrow \mathcal{F}$ is a probability density function associated to a jump arc. The size of a fluid jump is a random variable drawn from that distribution.

$g_t : \mathcal{M} \rightarrow \mathbb{B}$ is the guard of transition t that can be a function of the discrete and the continuous state. $\lambda_t : \mathcal{M} \rightarrow \mathbb{R} \cup \{\infty\}$ is a function of both the continuous and the discrete marking. Immediate transitions have firing rate $\lambda_t = \infty$.

The weight function $w_t : \mathcal{S} \rightarrow \mathbb{R}$ is defined for immediate transitions. The firing probability of each of the enabled immediate transitions in a vanishing state is $w_t(s) / \sum_{t_i \text{enabled in } \mathbf{m}} w_{t_i}(s)$.

The enabling rules for the transitions are the same as in discrete Petri net models. The fluid and the jump arcs act only on the fluid places and do not influence the enabling conditions of the transitions.

2.1 The Model Parameters

The dynamics of the CTMC underlying the discrete model part are defined through the generator matrix $\mathbf{Q}$, whereas the continuous dynamics are defined by the fluid change rate $\frac{dZ_k}{dt}$ that describes the change in the fluid level over time. The change in the fluid level is determined by the difference between the inflow and the outflow, each of which is normally distributed. The parameters of the normal distribution are either specified by the user or calculated as the diffusion approximation to the discrete transition at the origin of the inflow arcs and at the destination of the outflow arcs [Wol97].

Let μ_a and σ_a^2 be the expectation and variance of an arbitrary distribution of the discrete input transition to the fluid place P_k. Then the infinitesimal mean and variance of the fluid inflow rate are

$$\mu_\alpha = \frac{1}{\mu_a}, \qquad \sigma_\alpha^2 = \frac{\sigma_a^2}{(\mu_a)^3}. \tag{1}$$

and for the fluid outflow rate μ_δ and σ_δ^2 are analogous. (See [Kle76, New71] for a derivation of this formula).

All the incoming and outgoing streams in one marking add and for the flow rate $R_k(s, \mathbf{z})$ of the fluid place k in every marking $s \in \mathcal{S}$ the following parameters are valid

$$\mu_{k,s} = \sum_{\substack{\text{inflows} \\ \text{to k in s}}} \mu_\alpha - \sum_{\substack{\text{outflows} \\ \text{of k in s}}} \mu_\delta \quad \forall k \in \mathcal{P}_c, s \in \mathcal{S}, \tag{2}$$

$$\sigma_{k,s}^2 = \sum_{\substack{\text{inflows} \\ \text{to k in s}}} \sigma_\alpha^2 + \sum_{\substack{\text{outflows} \\ \text{of k in s}}} \sigma_\delta^2 \quad \forall k \in \mathcal{P}_c, s \in \mathcal{S}. \tag{3}$$

because variances must add. Note, that this simple linear combination only holds, because all the single fluid streams are normally distributed. The expectations and variances for each fluid place are collected into diagonal matrices (if there is none in a discrete state, the entry is equal to zero):

$$\boldsymbol{M}_k = \mathbf{diag}(\mu_{k,1}, \ldots, \mu_{k,S}), \tag{4}$$

$$\mathbf{\Sigma}_k^2 = \mathbf{diag}(\sigma_{k,1}^2, \ldots, \sigma_{k,S}^2). \tag{5}$$

If a hybrid system is modelled $\mu_\alpha, \sigma_\alpha^2, \mu_\delta$ and σ_δ^2 are given by the user and (1) is not needed.

The jump height function $b(.)$ can be an arbitrary probability density. For all discrete states the functions are collected into a diagonal matrix $\mathbf{B} = \mathbf{diag}(b_1(.), \ldots, b_S(.))$.

3 The Models

Two models shall be studied in this paper. Both are based on the performability model of a buffered multiprocessor system presented in [Mey82]. The first one uses two fluid places, one for the buffer content and the other for accumulating reward, that flows in as long as servers and buffer are operational. The second model has only one fluid place, that represents the costs arising from repair of processors and buffer. In the first model processors and buffer may fail but only the processors can be repaired, failure of the buffer leads to an absorbing failed state of the system. Opposed to this in the second model both processors and buffer can be repaired.

3.1 Two-dimensional Fluid Rate Reward Model

The regarded performability model has three processors that may fail independently at rate λ_p. The probability of recovering from a processor fault, the coverage, is c_p. The buffer has L stages that can also fail independently at rate λ_b. If one of the states fails, it causes a crash of the complete system. The

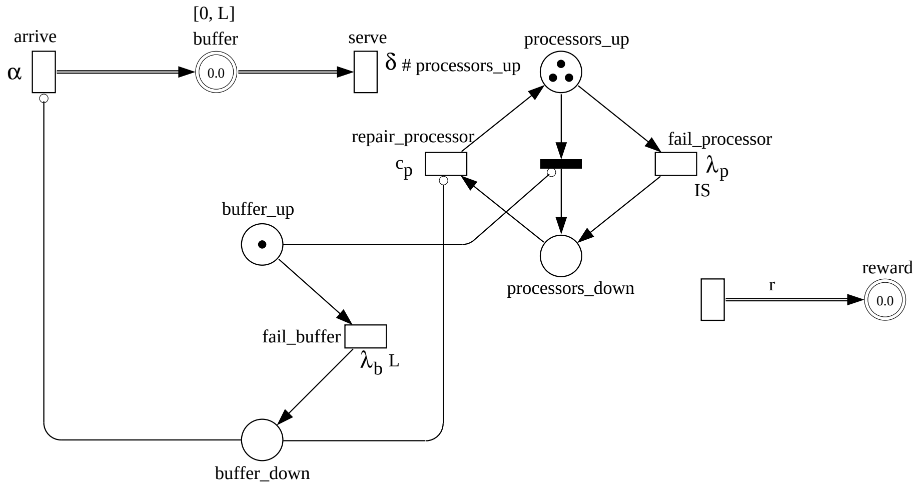

Figure 1: FSPN model of the performability model

failure rate of the whole buffer is then $L\lambda_b$. Jobs arrive to the system at rate α and they are served at rate δ per processor. So the buffer is being filled as long as it is operational and there is an outflow as long as at least one processor is functioning. Reward is accumulated at rate r per operational processor. The reward inflow r_r is a function of the global state of the system and is specified for each discrete state of the model.

Figure 1 shows an FSPN model of the system. The buffer is represented by the one fluid place. Its capacity $L = 4$. When the buffer is down all processors are put into the set of failed processors and the system terminates because none of the transitions is enabled. The second fluid place represents the accumulated reward. The rate function r_r is defined in such a way, that there is an inflow of 0.1 units per operational processor and per time unit only if the buffer is not empty. Hence the fluid flow rate accumulating reward depends on the content of the buffer.

The reachability graph is shown in figure 2. Only the discrete states are labelled, the components for the two fluid levels are omitted. The numbers in the small circles denote the numbering of the discrete states as it is used for the parameter vectors and the results. The failure rate of the buffer has been set to $\lambda_b = 10^{-4}$ per hour and the processor failure rate is $\lambda_p = 10^{-3}$ per hour. The mission time is 80 hours and the coverage $c_p = 0.999$.

For the computation of the fluid change rates (1), (2) and (3) have been used. In all states s where the buffer is operational the expectation μ_b of the fluid change rate of the buffer is computed following (2):

$$\mu_b = \mu_\alpha - \#\text{servers\_up} \cdot \mu_\delta$$

and the variance σ_b^2 equals according to (3)

$$\sigma_b^2 = \sigma_\alpha^2 + \#\text{servers\_up} \cdot \sigma_\delta^2 .$$

For the exponential distribution (1) evaluates to $\mu_\alpha = \alpha$ and $\sigma_\alpha^2 = \alpha$. The same holds for the exponentially

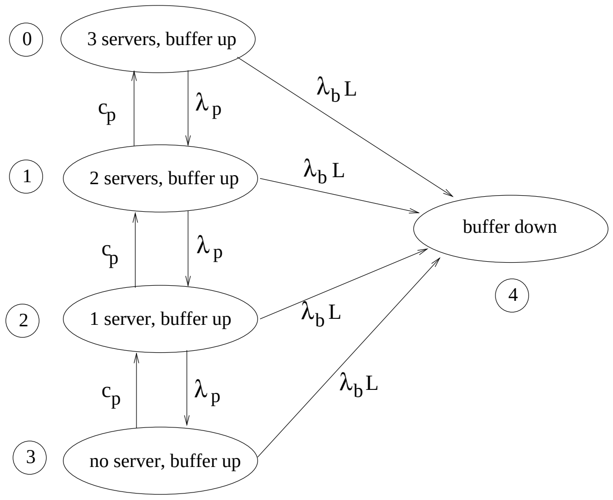

Figure 2: Reachability graph of the FSPN model

distributed service with rate δ. With $\alpha = 0.75$ and $\delta = 0.5$ the fluid flow parameters evaluate to

$$\boldsymbol{M}_b = \mathbf{diag}(-0.75, -0.25, 0.25, 0.75, 0.0) \tag{6}$$

and

$$\boldsymbol{\Sigma}_b^2 = \mathbf{diag}(2.25, 1.75, 1.25, 0.75, 0.1). \tag{7}$$

Reward is accumulated at normally distributed rate with expectation 0.1 and variance 0.1 per operational processor. Hence, the fluid flow parameters for the second fluid place are

$$\boldsymbol{M}_b = \mathbf{diag}(0.3 \cdot r_b, 0.2 \cdot r_b, 0.1 \cdot r_b, 0.0, 0.0) \tag{8}$$

with

$$r_b = \begin{cases} 1.0 & \text{if } z_1 \geq 0.2 \\ 0.0 & \text{else} \end{cases}$$

and

$$\boldsymbol{\Sigma}_b^2 = \mathbf{diag}(0.3, 0.2, 0.1, 0.1, 0.1). \tag{9}$$

The reachability graph gives a graphical representation of the stochastic process underlying this model. Augmenting it with two continuous variables, one for each fluid place, gives a description of the full stochastic process which is formally defined as

$$\{(S(t), Z_1(t), Z_2(t)), t \in \mathbb{R}_0^+\}$$

where $S(t)$ is the discrete marking at time t, and $Z_k(t)$ is a random variable, representing the fluid level in fluid place k at time t.

At time t the transient probability of the model being in discrete state i with fluid level in an infinitesimal environment around (z_1, z_2), is called the *fluid density* and is denoted by $\pi_i(t, z_1, z_2) = \frac{\partial}{\partial z_1} \frac{\partial}{\partial z_1} P(S(t) = i, Z_1(t) \leq z_1, Z_2(t) \leq z_2)$. This density is the component wise solution of the following differential equation in the vector form that describes the models' dynamics.

$$\frac{\partial \boldsymbol{\pi}(t, z_1, z_2)}{\partial t} + \frac{\partial}{\partial z_1}\Big(\boldsymbol{\pi}(t, z_1, z_2)\mathbf{M}_b(z_1)\Big) + \frac{\partial}{\partial z_2}\Big(\boldsymbol{\pi}(t, z_1, z_2)\mathbf{M}_r(z_2)\Big) - \frac{1}{2}\frac{\partial^2}{\partial z_1^2}\Big(\boldsymbol{\pi}(t, z_1, z_2)\boldsymbol{\Sigma}_b^2(z_1)\Big) - \frac{1}{2}\frac{\partial^2}{\partial z_2^2}\Big(\boldsymbol{\pi}(t, z_1, z_2)\boldsymbol{\Sigma}_r^2(z_2)\Big) - \boldsymbol{\pi}(t, z_1, z_2)\mathbf{Q}(z) = 0 \qquad \forall t, \forall z \in \mathring{\Omega}. \quad (10)$$

This is a system of diffusion-convection equations, also called Fokker-Planck equations, that is coupled through the generator matrix of the discrete Markov process $\mathbf{Q}$. The expectations define the drift and the variances the diffusion.

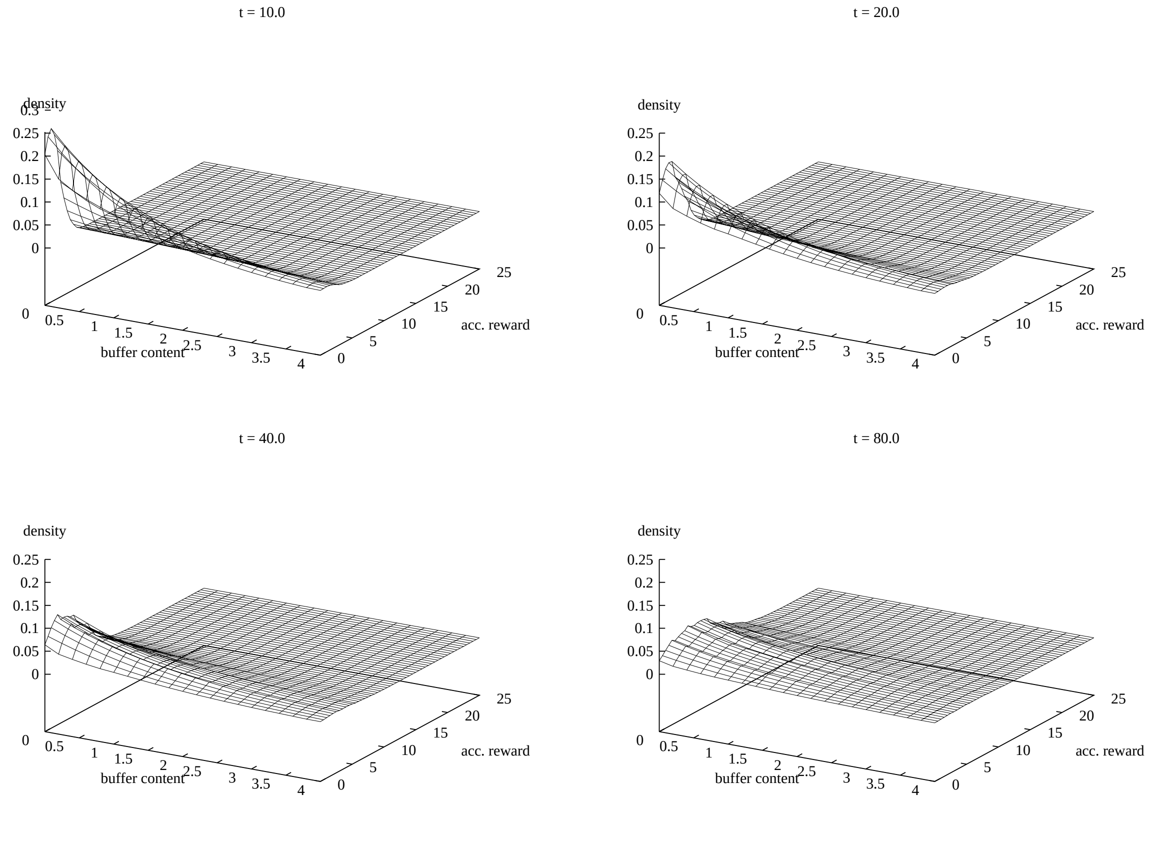

Figure 3: Probability densities in selected discrete states

An initial as well as a boundary condition on either end of the regarded domain must be imposed. The initial condition is defined by the models' configuration at time zero, whereas the boundary conditions are determined by the fact that the integral over a density always evaluates to one.

Let $\boldsymbol{\pi}_0$ be the vector of initial discrete state probabilities and let $\mathbf{z}_0$ be the vector of initial fluid levels, then the inital conditions are

$$\boldsymbol{\pi}(0, z_1, z_2) = \boldsymbol{\pi}_0 \cdot \delta(z_1 - z_{1,0})\delta(z_2 - z_{2,0}), \tag{11}$$

where $\delta(z_k - z_{k,0})$ is the Dirac impulse at $z_k = z_{k,0}$.

The boundary condition is equivalent to a reflecting barrier at the origin and at the upper bound [CM65]

$$\sum_{k=1}^{2} \left[\frac{1}{2} \frac{\partial}{\partial z_k} \boldsymbol{\pi}(t, z_1, z_2) \boldsymbol{\Sigma}_k^2(z_1, z_2) - \boldsymbol{\pi}(t, z_1, z_2) \mathbf{M}_k(z_1, z_2) \right]_{z=\{z_1^{\min}, z_1^{\max}\} \times \{z_2^{\min}, z_2^{\max}\}} = 0. \tag{12}$$

Equation (10) subject to (11) and (12) has been discretized and solved numerically [Zis98]. A theta method has been applied, that is a flexible discretization having the fully implicit, explicit and the Crank-Nicolson scheme as special cases. For the convection term an upwind scheme has been used, the coupling term is always discretized explicitly. Both examples presented here have been solved with the fully implicit discretization.

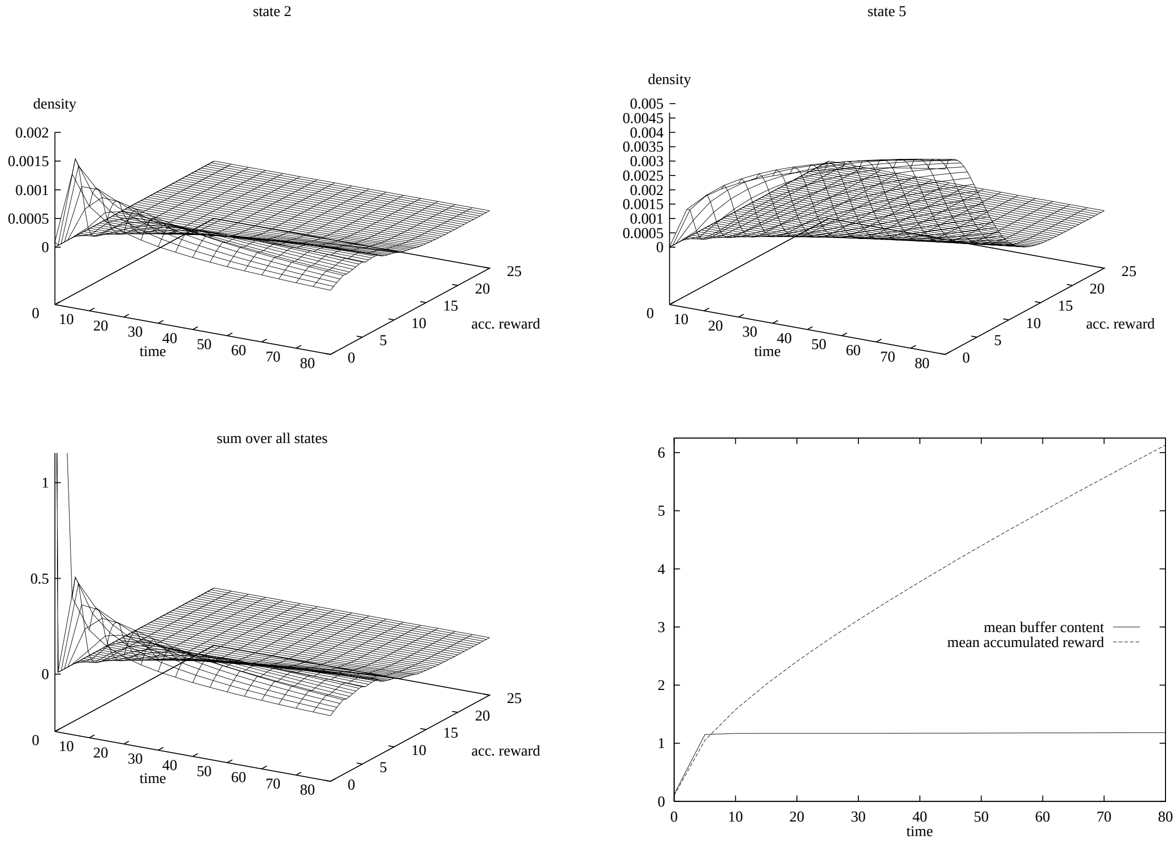

Figure 4: Transient probability density in selected discrete states

In each each discrete state a two-dimensional density is obtained, that evolves over time.[1] The sum

[1] Strictly speaking this is a density only after normalization. Nonetheless, we will call the fractions of a density also density.

over all discrete states is a two-dimensional probability density at each time point. The change of all functions can be observed like a film. Figure 3 shows the probability density in state one at selected time points. The density of the accumulated reward is obtained by integrating at each time point over the buffer content, the second continuous dimension. The result has one dimension less, so it can be plotted over time. Figure 4 shows the transient evolution of the density of the accumulated reward in state one (with one failed processor), in state 4 (where the buffer is failed) and for the sum over all discrete states. The latter is the density of the accumulated reward, independent of what discrete state the model is in. The lower right plot in figure 4 shows the transient evolution of the means of both continuous variables.

3.2 Rate and Impulse Reward Model

The second example regarded in this paper is another variation of the performability model. The buffer is modelled with a discrete place and there is only one fluid place. This fluid place represents the costs that arise from repair of the processors as well as the buffer. Each failure of the buffer costs an additional price once, whereas the repair costs per unit of time. The failure and repair rate of the processors and arrival and service rate are the same as in the first model. For obtaining illustrative results the failure rate of the buffer has been set to $\lambda_b = 0.01$. Repair of the buffer takes on the average 6 hours ($1/r_b = 6$ hours). The costs that are incurred by the repair of processors and buffer are normally distributed with mean 0.1 and variance 0.05 per processor, and mean 0.5 and variance 0.1 for the repair of the buffer and maintenance of the whole system. Each failure of the buffer costs additional deterministic amount of 0.5 units.

The buffer has $L = 2$ stages. The failure rate of the whole buffer is $L\lambda_b$ and the repair rate of the whole system after a buffer failure is r_b. Upon failure of the buffer data is saved, so that the system can continue processing after repair. With the repair of the buffer the failed processors are repaired too.

Figure 5 shows an FSPN model of the system. The costs are represented by the fluid place. A jump arc connects transition *fail_buffer* with the place *costs*. When the buffer is down all the failed processors are put back into the set of operational processors. As long as repair of the buffer takes all the processors are in maintenance and repair, and the transition representing the buffer failure is disabled as is the arrival of jobs. No service takes place when the buffer is empty. The main difference between this model and the first one is that the fluid place is not only filled through the fluid arcs but also through a jump arc.

The reachability graph of this model has 15 states (it is not depicted), three blocks of the reachability graph of the two-dimensional model, one for each number of packets in the buffer. The fluid flow parameters are the same in all three blocks and they evaluate to $\mathbf{M}_b = \mathbf{diag}(\mu_{r_1} + \mu_{r_2}) = \mathbf{diag}(\boldsymbol{\mu}, \boldsymbol{\mu}, \boldsymbol{\mu})$ with $\boldsymbol{\mu} = (0.0, 0.1, 0.2, 0.3, 0.5)$ and $\mathbf{\Sigma}^2{}_b = \mathbf{diag}(\sigma^2_{r_1} + \sigma^2_{r_2}) = \mathbf{diag}(\boldsymbol{\sigma}, \boldsymbol{\sigma}, \boldsymbol{\sigma})$ with $\boldsymbol{\sigma} = (0.05, 0.05, 0.1, 0.15, 0.1)$.

The jump height matrix is $\mathbf{B} = \mathbf{diag}(\mathbf{b}, \mathbf{b}, \mathbf{b})$ where $\mathbf{b} = (det(0.5), det(0.5), det(0.5), det(0.5), 0.0)$ and finally the jump rates are $\mathbf{\Lambda} = \mathbf{diag}(\boldsymbol{\lambda}, \boldsymbol{\lambda}, \boldsymbol{\lambda})$ where $\boldsymbol{\lambda} = (0.02, 0.02, 0.02, 0.02, 0.0)$.

The stochastic process under consideration is a Markov process similar to the one in the first example with only one continuous component. The fluid density in state i is $\pi_i(t, z) = d/dz\ P(S(t) = i, Z(t) \leq z)$

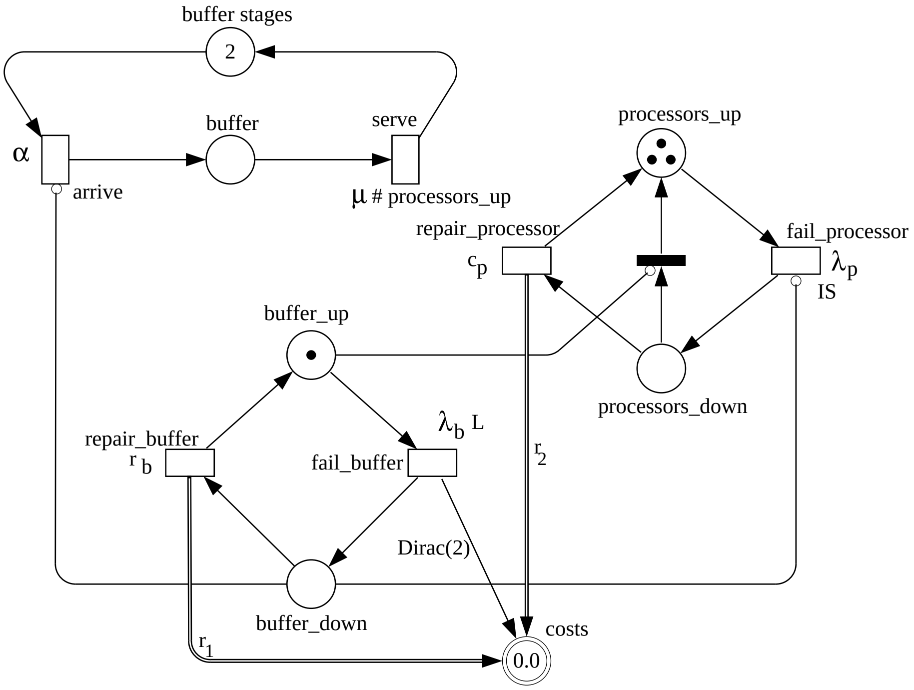

Figure 5: FSPN model of the second performability model

and the equation describing the models dynamics is

$$\frac{\partial \boldsymbol{\pi}(t,z)}{\partial t} + \frac{\partial}{\partial z}\Big(\boldsymbol{\pi}(t,z)\mathbf{M}(z)\Big) - \frac{1}{2}\frac{\partial^2}{\partial z^2}\Big(\boldsymbol{\pi}(t,z)\boldsymbol{\Sigma}^2(z)\Big) - \boldsymbol{\pi}(t,z)\Big(\mathbf{Q}(z) - \boldsymbol{\Lambda}(z)\Big) - \int_{z_{\min}}^{z} \boldsymbol{\pi}(t,y)\mathbf{B}(z-y)\,dy\,\boldsymbol{\Lambda}(z) = 0 \qquad \forall t, \forall z \in \mathring{\Omega}. \quad (13)$$

The derivatives are those of the first example reduced to one dimension. The integrals are equivalent to those in the Takács integro-differential equation [Tak55], they are the convolution of the jump height distribution with the fluid density function, but the Takács equation is a first order differential equation.

The solutions of this equation are plotted in figure 6 for the discrete state (from top to bottom, left to right) 'buffer empty, two processors up, buffer up', 'buffer empty, buffer down' and for the sum over all discrete states. The behaviour is dominated by the fluid flow, so the waves due to the jumps are not visible. The mean accumulated reward (lower right plot) increases almost constantly at rate 0.0875. The rate is so low, because the model is with probability 0.8 (this measure is not plotted) in states where all processors and the buffer are operational.

4 Conclusions

Two models have been presented in order to demonstrate the modelling power of second order FSPNs in Markov reward modelling. Both are modifications of a well-known performability model. It has been

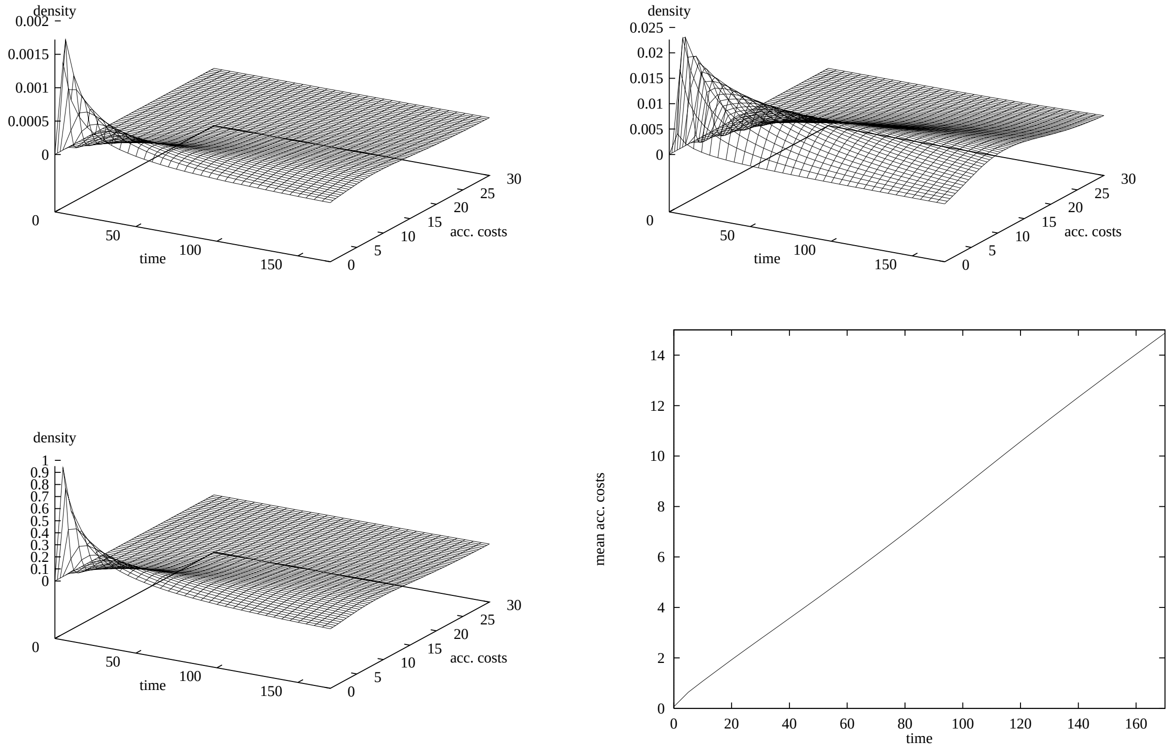

Figure 6: Transient probability density in selected discrete states

shown that not only accumulated rate reward can be modelled with an FSPN, but also accumulated impulse reward.

Rate reward is accumulated in an FSPN through the inflow along a fluid arc to a fluid place while a transition is enabled, whereas impulse reward is added upon the firing of a transition as indicated by a jump arc. Jump arcs are labelled with a probability distribution. An amount sampled from that distribution is added to the content of the fluid place at the destination of the jump arc upon the firing of the transition at its origin. The example in this paper uses the Dirac impulse as probability distribution to model deterministic jump height.

The underlying equations are differential equations that also include convolution integrals. A numerical discretization method has been used to solve the equations. The numerical algorithm is also able to handle models with two fluid places. So it was possible in the first example to use one fluid place for the accumulated reward and the other one for a continuously approximated buffer. Further work will be on good numerical algorithms that are more run-time and memory efficient, so that larger models can be analyzed and on the convenient computation of different reward measures from the fluid densities.

References

[ABC+95] M. Ajmone Marsan, G. Balbo, G. Chiola, S. Donatelli, and G. Francheschinis. *Modelling with Generalized Stochastic Petri Nets.* John Wiley & Sons, 1995.

[CM65] D.R. Cox and H.D. Miller. *The Theory of Stochastic Processes.* Chapman and Hall, 1965.

[CNT97] G. Ciardo, D. Nicol, and K.S. Trivedi. Discrete-event Simulation of Fluid Stochastic Petri Nets. In *Proc. Seventh International Workshop on Petri Nets and Performance Models - PNPM'97*, pages 217–225, Saint Malo, France, June 3–6 1997. IEEE-CS Press.

[DA93] R. David and H. Alla. Autonomous and Timed Continuous Petri nets. In G. Rozenberg, editor, *Advances in Petri Nets 1993*, volume 674 of *Lecture Notes in Computer Science*, pages 71 – 90. Springer-Verlag, 1993.

[HKNT98] G. Horton, V. G. Kulkarni, D. M. Nicol, and K. S. Trivedi. Fluid Stochastic Petri Nets: Theory, Applications and Solution. *European Journal of Operations Research*, 105(1):184–201, February 1998. Also published as ICASE report no. 96-5.

[Hor96] G. Horton. Computation of the Distribution of Accumulated Reward with Fluid Stochastic Petri nets. In *Proc. 2nd IEEE Int. Computer Performance and Dependability Symposium (IPDS)*, pages 90–95, Urbana-Champaign, Illinois, USA, 4-6 September 1996. IEEE-CS Press.

[Kle76] L. Kleinrock. *Queueing Systems Vol II: Computer Applications.* John Wiley, 1976.

[Kob74] H. Kobayashi. Application of the Diffusion Approximation to Queueing Networks II: Nonequilibrium Distributions and Applications to Computer Modeling. *Journal of the ACM*, 21(3):459–469, July 1974.

[Mey82] J. F. Meyer. Closed-Form Solutions of Performability. *IEEE Transactions on Computers*, C-31(7):648–657, July 1982.

[New71] G. F. Newell. *Applications of Queueing Theory.* Chapman and Hall, Ltd. (London), 1971.

[SM91] W. H. Sanders and J. F. Meyer. A Unified Approach for Specifying Measures of Performance, Dependability, and Performability. In A. Avizienis and J. Laprie, editors, *Dependable Computing for Critical Applications*, volume 4 of *Dependable Computing and Fault-Tolerant Systems*, pages 215–237. Springer, 1991.

[Tak55] L. Takács. Investigation of waiting time problems by reduction to Markov processes. *Acta Math. Acad. Sci. Hungar.*, 6:101–129, 1955.

[TK93] K. S. Trivedi and V. G. Kulkarni. FSPNs: Fluid Stochastic Petri Nets. In *Proc. 14th Int. Conf. on the Application and Theory of Petri Nets*, pages 24–31, Chicago, 1993.

[Wol97] K. Wolter. Second Order Fluid Stochastic Petri Nets: an Extension of GSPNs for Approximate and Continuous modelling. In *Proc. 1st World Congress on Systems Simulation (WCSS'97)*, pages 328–332, Singapore, Sept. 1–3 1997.

[Zis98] A. Zisowsky. Entwurf und Implementierung eines Verfahrens für die transiente Analyse fluider stochastischer Petri-Netze. Master's thesis, TU Berlin, August 1998.