

APPROXIMATE SOLUTION OF A CLASS OF QUEUEING NETWORKS WITH BREAKDOWNS

NIGEL THOMAS*, DAVID THORNLEY[†] and HARF ZATSCHLER[†]

Abstract: In this paper we study a class of open queueing network where servers suffer breakdowns and are subsequently repaired. The network topology is a pipeline with feedback from the final node to the first. Each node consists of a number of queues each with an unreliable server. There are no losses from the queues in this system, however jobs are routed according to the distribution of operational servers at each node in the pipeline. This model is in general intractable, however an iterative technique is presented which combines a number of earlier results to generate an approximation to steady state measures found by simulation.

Keywords: Queueing theory, breakdowns, approximation, decomposition

1. INTRODUCTION

Queueing networks with breakdowns are a class of problem that are of obvious practical interest and have consequently been considered for many years. However, the vast majority of studies that have been made concern only single queue models or solve more general topologies using simulation. A number of papers have addressed the problem of queues in parallel, most notably Mitrani and Wright [Mitrani and Wright, 1994] who analysed a system of nodes in parallel which suffered failures that caused all jobs to be lost, incoming jobs were then routed away from failed nodes, this resulted in an interesting trade off in performance between response time and job loss. Models without loss on failure are not without practical application, particularly in transaction processing and manufacturing. Thomas and Mitrani [Thomas and Mitrani, 1995] started with the same basic model as [Mitrani and Wright, 1994], but changed the nature of the failure so that queues were preserved during repair periods. The same authors also considered an extension to their model [Thomas and Mitrani, 1998] where a pipeline was constructed where each node was a system of parallel queues. It was not possible to solve this model exactly, instead they considered each node in series and compared a simple Poisson approximation with

a Markov-modulated arrival process based on the configuration of operational servers at the previous node in the pipeline.

In this paper we present an extension to the model presented in [Thomas and Mitrani, 1998] to consider the existence of a feedback loop which returns jobs to the start of the pipeline with a given probability. The existence of such a loop means that the approach used previously will no longer be applicable because all the nodes are now dependent of their predecessor, whereas in [Thomas and Mitrani, 1998] the first node had only external arrivals. We employ an iterative approach recently applied to Markovian process algebra [Thomas et al, 2003]. In the context of the queueing systems described here this iterative method is extremely close to that applied recently in [Harrison et al, 2002]. The approach described in [Thomas et al, 2003] requires that all shared actions (in queueing terms this refers to departures from one node which become arrivals at another) are represented in a reduced model to estimate the marginal distribution for each component (node) in turn. The method is repeated until convergence over a particular measure is reached and hence all the marginal distributions are found. In this paper the reduced model is a single queue with Markov-modulated arrivals and convergence is required of the parameters of the arrival process at each node. As in earlier studies these marginal queue size distributions do not in general give rise to a product form solution, but nevertheless can be used to find many performance measures of interest, such as average response time and utilisation.

*Department of Computer Science, University of Durham, UK.
nigel.thomas@durham.ac.uk

[†]Department of Computing, Imperial College London, UK.
{djt,hz3}@doc.ic.ac.uk

2 THE MODEL

Jobs arrive into the system in a Poisson stream with rate λ . There are K nodes in series and in node i there are N_i servers in parallel, each with an associated unbounded queue, to which incoming jobs may be directed. Server j at node i goes through alternating independent operative and inoperative periods, distributed exponentially with means $1/\xi_{i,j}$ and $1/\eta_{i,j}$ respectively. While it is operative, the jobs in its queue receive service of an exponentially distributed duration with mean $1/\mu_{i,j}$, and leave the node upon completion to proceed to the next (if any) node in the pipeline. When a server becomes inoperative (breaks down), the corresponding queue, including the job in service (if any), remains in place. Services that are interrupted in this way are eventually resumed from the point of interruption. On completion of service at the final node a proportion of jobs, $0 < \phi \leq 1$, leave the system and the remainder return to the first node. The system model is illustrated in Figure 1.

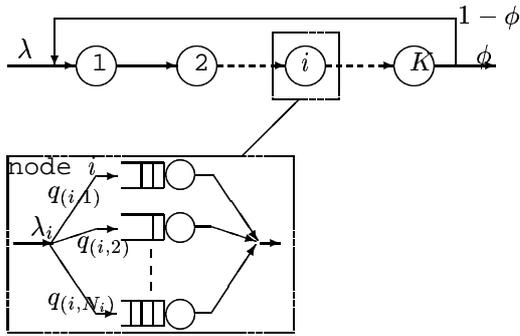


Figure 1. A single source to a pipeline of K nodes, split between the queues in each node

The external arrival rate is given in Figure 1 as λ_i , and since no jobs are lost the overall arrival rate at all nodes will be the same, namely, λ/ϕ . However, since the arrivals at node i depend on the departures from node $i - 1$ then the arrival stream will, in general, cease to be Poisson (the only case where this is not true is for node 1 if $\phi = 1$). The *system configuration* at any moment is specified by the subset, σ , of servers that are currently operative (that subset may be empty, or it may be the set of all servers): $\sigma \subset \Omega_N$, where $\Omega_N = \{(1, 1), (1, 2), \dots, (1, N_1), (2, 1), \dots, (K, N_K)\}$, where the pair $\{i, j\}$ represents server j at node i . There are of course 2^N possible system configurations, where $N = \sum_{i=1}^K N_i$. In general it is more convenient to consider the subset σ_i whose elements are those servers at node i which are operative. The set of

all servers at node i is denoted by Ω_{N_i} . Clearly $\sigma_i \subset \Omega_{N_i} \subset \Omega_N$ and $\sigma_i \subset \sigma$. The steady-state marginal probability, p_{σ_i} , of configuration σ_i at node i is given by

$$p_{\sigma_i} = \prod_{j \in \sigma_i} \frac{\eta_{i,j}}{\xi_{i,j} + \eta_{i,j}} \prod_{j \in \bar{\sigma}_i} \frac{\xi_{i,j}}{\xi_{i,j} + \eta_{i,j}}, \quad \sigma_i \subset \Omega_{N_i},$$

And the steady-state marginal probability, p_σ , of configuration σ is given by

$$p_\sigma = \prod_{i,j \in \sigma} \frac{\eta_{i,j}}{\xi_{i,j} + \eta_{i,j}} \prod_{i,j \in \bar{\sigma}} \frac{\xi_{i,j}}{\xi_{i,j} + \eta_{i,j}}, \quad \sigma \subset \Omega_N,$$

where $\bar{\sigma}_i$ is the complement of σ_i with respect to Ω_{N_i} , $\bar{\sigma}$ is the complement of σ with respect to Ω_N and an empty product is by definition equal to 1. These expressions follow from the fact that servers break down and are repaired independently of each other.

If, at the time of arrival at node i , a new job finds the node in configuration σ_i , then it is directed to the queue at server j with probability $q_{i,j}(\sigma_i)$. These decisions are independent of each other, of past history, of the sizes of the various queues and of the state of any other node in the pipeline. Thus, a routing policy at node i is defined by specifying 2^{N_i} vectors,

$$\mathbf{q}_i(\sigma_i) = [q_{i,1}(\sigma_i), q_{i,2}(\sigma_i), \dots, q_{i,N_i}(\sigma_i)], \quad \sigma_i \subset \Omega_{N_i},$$

such that for every σ_i ,

$$\sum_{j=1}^{N_i} q_{i,j}(\sigma_i) = 1$$

There are clearly many strategies that can be employed using this system and a number have been studied previously. Intuitively, it seems better not to send jobs to queues where the server is inoperative, unless that is unavoidable. This suggests the following strategy: If the subset of operative servers at node i in the current system configuration is σ_i , and that subset is non-empty, send jobs to queue j only if $j \in \sigma_i$, with probability proportional to q_j :

$$q_j(\sigma_i) = \frac{q_j}{\sum_{l \in \sigma} q_l}, \quad j \in \sigma.$$

If σ is empty (i.e. all servers are broken), send jobs to queue j with probability q_j ($j = 1, 2, \dots, N_i$). Note that this strategy does not take account of the states of servers at other nodes in the system. However the existence of other nodes may have an effect on the optimal routing vector for a given strategy, for instances

in spreading jobs when all preceding servers are operative, but directing jobs only to fast servers when few preceding servers are operative.

The system state at time t is specified by the pair $[I(t), \mathbf{J}(t)]$, where $I(t)$ indicates the current configuration (the configurations can be numbered, so that $I(t)$ is an integer in the range $0, 1, \dots, 2^N - 1$), and $\mathbf{J}(t)$ is an integer vector whose k 'th element, $J_k(t)$, is the number of jobs in queue k ($k = 1, 2, \dots, N$). The integer k is used here instead of the pair i, j for simplicity, the relationship between k and i, j is a simple 1 to 1 mapping such that

$$j + \sum_{x=1}^{i-1} N_x = k$$

Under the assumptions that have been made, $X = \{[I(t), \mathbf{J}(t)], t \geq 0\}$ is an irreducible Markov process. The condition for ergodicity of X is that, for every queue i, j , the overall arrival rate is lower than the overall service capacity:

$$\sum_{\forall \sigma_i} \lambda_i p_{\sigma_i} q_{i,j}(\sigma_i) < \mu_{i,j} \frac{\eta_{i,j}}{\xi_{i,j} + \eta_{i,j}}$$

$, i = 1, 2, \dots, K, j = 1, 2, \dots, N_i.$

When the routing probabilities at each node depend on the system configuration, the process X is not separable (i.e., it does not have a product-form solution). Consequently, the problem of determining its equilibrium distribution is intractable in general. On the other hand, the quantities of principal interest are expressed in terms of averages only; they are the steady-state mean queue sizes, L_k , and the overall average response time, W , given by

$$W = \frac{1}{\lambda} \sum_{i=1}^K \sum_{j=1}^{N_i} L_{i,j}.$$

To determine those performance measures, it is not necessary to know the joint distribution of all queue sizes; the marginal distributions of the N queues in isolation are sufficient. Unfortunately, the isolated queue processes, $\{J_k(t), t \geq 0\}$ ($k = 1, 2, \dots, N$), are not Markov. As mentioned earlier the arrival stream at any node i is not Poisson since it depends on the activity of all the previous nodes (*ad infinitum* given $\phi < 1$), this makes an exact solution of the marginal queue size distributions almost as intractable a problem as solving the joint distribution of all queue sizes. However, it is possible to obtain good approximate solutions for the marginal queue size distributions by assuming the

arrival stream at node i to be Markov-modulated Poisson. However, unlike the pipeline model presented in [Thomas and Mitrani, 1998], there is no explicit start point in this model where the arrival process at a node is known, therefore an iterative solution is employed and a further approximation is used to start the process. This iterative process and the formation of the Markov-modulated arrival processes are discussed in Sections 3 and 4.

Consider the stochastic processes $Y_{i,j}$,

$$Y_{i,j} = \{[I^*(t), J_{i,j}(t)], t \geq 0\}$$

$, i = 1, 2, \dots, K, j = 1, 2, \dots, N_i,$

which model the joint behaviour of the configuration and the size of an individual queue i, j , where $I^*(t)$ indicates the current mode of the Markov-modulated arrival process (MMPP). The state space of $Y_{i,j}$ is infinite in one dimension only, which simplifies the solution considerably and makes it tractable for reasonably large values of the number of modes in the MMPP, I_{max} . The important observation here is that, with the assumption of a Markov-modulated Poisson process, $Y_{i,j}$ is an irreducible Markov process, for every i, j . This is because the arrivals into, and departures from queue i, j during a small interval $(t, t + \Delta t)$ depend only on the approximated system configuration and the size of queue i, j at time t , and not on the sizes of the other queues. As mentioned earlier, without the approximation of the arrival stream to a Markov-modulated Poisson process, this statement would not be true, since a job only arrives at node $i + 1$ after successfully completing service at node i , therefore making the queue size at any node dependent on all previous nodes of service. It is then necessary to find the equilibrium distribution of $Y_{i,j}$:

$$p_{i,j}(x, y) = \lim_{t \rightarrow \infty} P[I^*(t) = x, J_{i,j}(t) = y]$$

$, x = 0, 1, \dots, I_{max} - 1, y = 0, 1, \dots$

Given the probabilities $p_{i,j}(x, y)$, the average size of queue i, j is obtained from

$$L_{i,j} = \sum_{y=1}^{\infty} y \sum_{x=0}^{I_{max}-1} p_{i,j}(x, y)$$

There are three established approaches to solving systems of this kind, matrix geometric methods [Neuts, 1981], solution by generating functions and spectral expansion [Mitrani and Chakka, 1995]. We have employed spectral expansion due primarily to familiarity with this technique and this choice is somewhat arbitrary. The use of spectral expansion has some issues

regarding stability with respect to deriving eigenvalues, although the method is well known, elegant and efficient. Since it appears in detail elsewhere we do not present the application of the spectral method here and the interested reader is directed to that earlier work [Mitrani and Chakka, 1995].

3 APPROXIMATION USING AN MMPP

In the previous section it was stated that the arrivals at node i could be approximated by a Markov-modulated arrival process. In the study of the simpler pipeline model [Thomas and Mitrani, 1998] a comparison was made between a simple Poisson approximation and an MMPP where the modes correspond to the operational state at the preceding node, i.e. σ_{i-1} , thus the MMPP used at node i will have $2^{N_{i-1}}$ modes. In each mode the arrival rate is calculated as the sum of the departure streams in that operational state:

$$\sum_{j=1}^{N_i} \mu_{i,j} (p_{\sigma_i} - p_{i,j,0}(\sigma_i))$$

If this model was specified using Markovian process algebra and the technique described in [Thomas et al, 2003] applied then there would in fact be $4^{N_{i-1}}$ modes in the MMPP. This is because the separation of *actions* in process algebra gives rise to separate modes not only when each server is operational or not, but also whether its queue is empty or not. Hence the mode in the MMPP is described by the superset of the pairs (o_i, e_i) , where $o_i \in \{0, 1\}$ indicates whether the server i is operational or not and $e_i \in \{0, 1\}$ indicates whether the queue at i is empty or not. The arrival rate in each mode is the sum of service rates for each server that is both working and has a non-empty queue. The transitions between modes in this case are somewhat more complex. Clearly the transitions arising from a change in operational state are the same as previously, and the transition rate from empty to non-empty is simply the average arrival rate into that queue in that operational state. The transition from non-empty to empty is calculated as the service rate multiplied by the probability that there is exactly one job in the queue given that it is non-empty. Clearly this requires knowing the probabilities $p_{i-1,j}(\sigma_{i-1}, 0)$ and $p_{i-1,j}(\sigma_{i-1}, 1)$ for all j and σ_{i-1} , except $\sigma_i = \emptyset$.¹ The number of modes in the MMPP clearly has implications for the amount of work required to solve the model. Therefore a move from $2^{N_{i-1}}$ to $4^{N_{i-1}}$ is clearly not

¹Clearly the predecessor of node 1 is node K , hence we interpret $1 - 1$ as K in this instance.

desirable when N_{i-1} is large except if there is a significant increase in accuracy of the approximation.

4 ITERATIVE SOLUTION

If the probability of leaving the system after the last node is not certain, $\phi < 1$, then there is no node for which the input process is entirely known. To address this problem an iterative approach is adopted (from [Thomas et al, 2003]) to tackle this problem as follows.

1. Calculate the rate of an equivalent Poisson arrival process at node 1.
2. Solve to find the approximate marginal queue size probabilities at node 1.
3. Use the calculated queue empty probabilities to generate a Markov-modulated arrival process at the next node.
4. Use this to calculate the approximate marginal queue size probabilities at the next node.
5. Repeat steps 3 and 4 until convergence criterion is satisfied (or abandon).

The convergence criterion employed is that the same MMPP is calculated twice in succession (to some number of decimal places) for any given node. The equivalent Poisson stream at node 1 is easily shown to have rate $\frac{1}{\phi} \lambda$ in steady state since there is no job loss.

5 NUMERICAL RESULTS

Figures 2 and 3 show illustrate that the MMPP approximation based on the operative state at the preceding node is generally very successful at predicting average response time except when the periods of operation and inoperation were very long (Figure 3) and at high load (Figure 2). One of the reasons for this inaccuracy is that during long inoperative periods the number of jobs at the preceding node will become much larger (all the servers may be broken or sufficient such that the remainder are saturated) and so on repair there will be a period of continuous service before steady state behaviour can eventually be resumed. The more complex MMPP approximation including states where the preceding queues are empty are similar in accuracy to the MMPP case. The advantage of these two methods of approximation is that they appear to offer an upper

and lower bound respectively, but only in the absence of feedback.

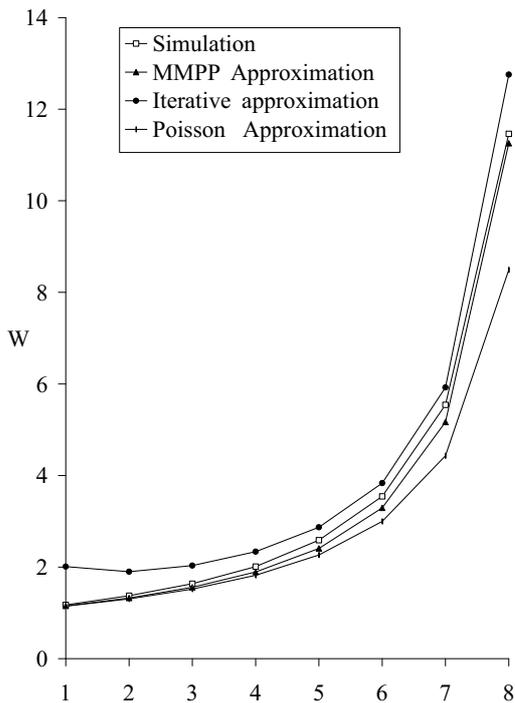


Figure 2. Average response time at node 2 against arrival rate; $\mu_i = 10, \eta_i = 0.1, \xi_i = 0.01, \phi = 1$

The divergence between a simple Poisson approximation and the MMPP approximation becomes more exaggerated when the feedback probability ϕ is decreased. Under these conditions the traffic at each node becomes increasingly less Poisson and the accuracy of both approximations diminishes dramatically. Unlike the pipeline case these approximations give an over estimate of the average response time in the presence of feedback. This is due to the tightly coupled nature of this model. The key point to observe is when the bulk of the arrivals occur into a queue, not how bursty they are. If node 1 breaks down then all arrivals into node 2 are blocked and so (after a number of services at node 2) is the feedback. This means that during most of the breakdown period the arrivals into node 1 are just the external arrivals. The same happens during breakdowns at node 2, although obviously here the external jobs have to pass through node 1 before reaching node 2. Thus in both cases relatively few jobs arrive at a node when it is broken so the queue sizes don't grow too much. Contrast this to the Poisson approximation; here jobs are assumed to arrive at a node regardless of its state, a constant rate of λ/ϕ . Hence in the feedback case the

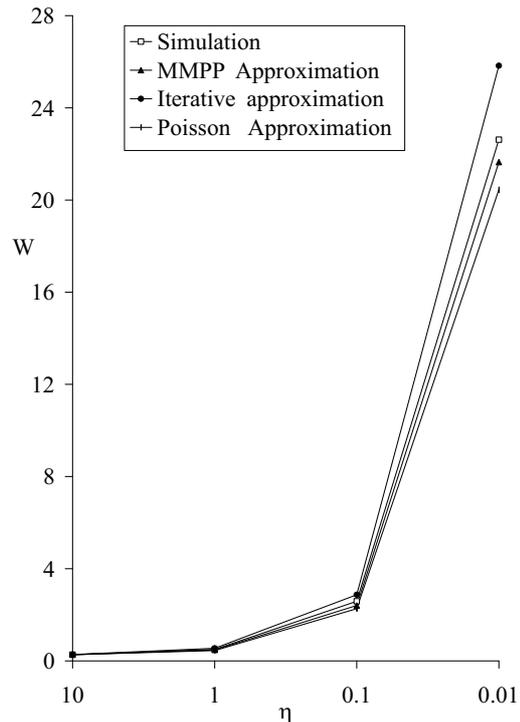


Figure 3. Average response time at node 2 against repair rate; $\mu_i = 10, \lambda = 5, \xi_i = \eta_i/10, \phi = 1$

Poisson approximation ceases to be a lower bound. The more complex approximation captures some of this behaviour, but it's still fairly crude.

Figures 4 and 5 show the behaviour of the same two node model where $\phi = 0.5$. In Figure 4 the average system response time is shown when the external arrival rate is varied. Note that the effective rate of arrivals is in fact twice the rate given on average, but the arrivals are made much more bursty by the occurrence of breakdowns. At low load the iterative approach clearly gives the best approximation, but this becomes less accurate as load increases. At even higher load, nearer saturation, the Poisson approximation becomes more accurate. This is due to the fact that the queue lengths are generally so large that the queues rarely empty during a breakdown at the other node, and so the "steady" arrivals of the Poisson approximation are more realistic.

6 CONCLUSIONS

The method proposed here seeks to extend the pipeline model to include a feedback which is incorporated into

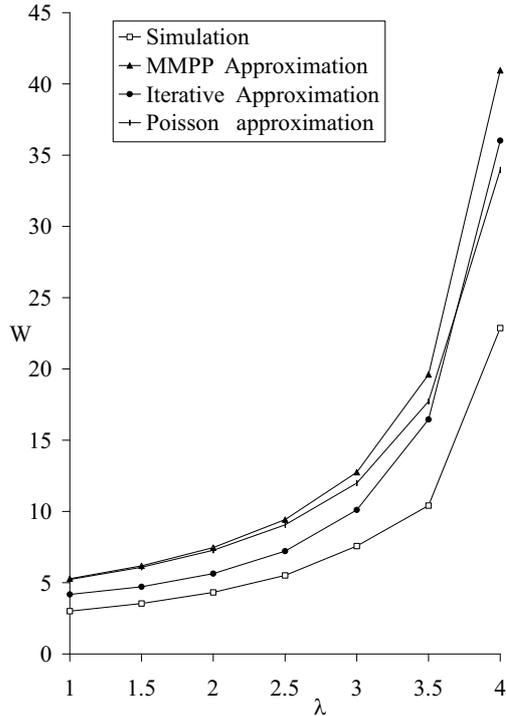


Figure 4. Average system response time against arrival rate; $\mu_i = 10, \eta_i = 0.1, \xi_i = 0.01, \phi = 0.5$

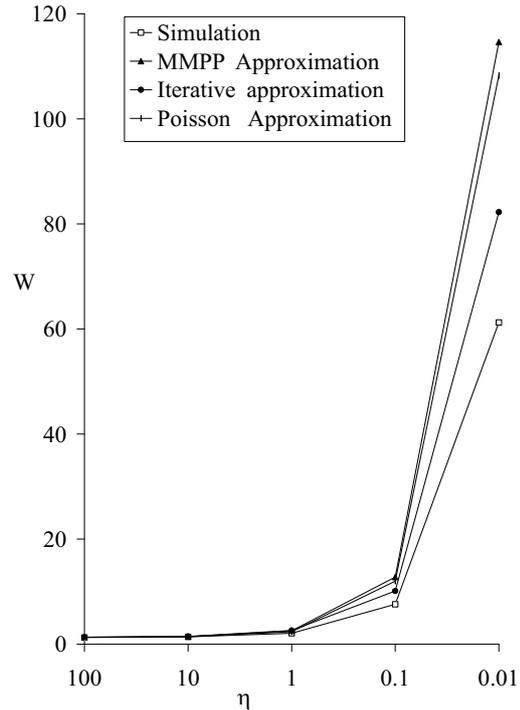


Figure 5. Average system response time against repair rate; $\mu_i = 10, \lambda = 3, \xi_i = \eta_i/10, \phi = 0.5$

the solution method using an iterative approach. The class of models that it is possible to consider using this method includes more general network structures, although it is evaluated here with a simple loop. Simple approximations work well under most conditions but become increasingly less accurate as load and inoperative periods increase. These approximations form a lower bound to the exact solution when there is no feedback, but become overly pessimistic when feedback exists. The more complex approximation with iterative solution method performs somewhat better under these conditions, but there is still considerable room for improvement, particularly at high load. Incorporating a burst of services following repair may increase the accuracy of this method, this remains an area of continuing investigation.

REFERENCES

Harrison P. Thornley D. and Zatschler H. 2002, Geometrically Batched Networks, in: *Proceedings of the Seventeenth International Symposium On Computer and Information Sciences*.

Mitrani I. and Chakka R. 1995, Spectral Expansion Solution for a Class of Markov Models: Application and Comparison with the Matrix-Geometric Method, *Performance Evaluation*, 23, pp. 241-260.

Mitrani I. and Wright P. 1994, Routing in the Presence of Breakdowns, *Performance Evaluation* 20, pp. 151-164.

Neuts M. 1981, *Matrix geometric solutions in stochastic models*, John Hopkins Univ Press.

Thomas N. and Mitrani I. 1995, Routing among different servers, in: F. Baccelli, A. Jean-Marie and I. Mitrani, eds., *Quantitative Methods in Parallel Systems*, Springer-Verlag.

Thomas N. and Mitrani I. 1998, Approximate solution of a pipeline with server vacations, in: *Proceedings of 12th European Simulation Multiconference*.

Thomas N. Bradley J. and Thornley D. 2003, An approximate solution of PEPA models using component substitution, *IEE Proceedings Computers and Digital Techniques*.