

Synchronized negative customers in an unreliable server queue

David J Thornley

Department of Computing, Imperial College London
djt@doc.ic.ac.uk

Abstract. Negative customers are an accepted means of removing jobs from systems in queueing theory. A lack of synchronization with causal events can adversely affect results. We introduce a formulation for a queue with unreliable servers in which the arrival of negative customers is synchronized with the processor breakdown. Probabilistic re-sampling on repair can also be modelled, resulting in an interesting transition structure in the Markov chain representing the joint states of queue length and number of active processors. This corresponds to a MAP-like process in which one arrival class is of negative customers. Unlike MAP/MAP or MAP/PH queues, job addition/removal processes are co-modulated. We demonstrate that the introduction of synchronization affects the resulting queue solution when compared to the use of independent negative customer arrivals. We suggest that the synchronised negative customer may be of use in modelling systems in which the processor has local storage, and the job information is moved rather than copied. A synchronised re-sample may correspond to the processor having been suspended, perhaps for cooling rather than replacement, then re-activating.

1 Introduction

Multiprocessor servers with breakdowns and repairs can be modelled using a Markov modulated queue in which the modulation state tracks the number of active processors. Mitrani and Chakka [1] examine the solution for the steady state of such a queue in which the queue length is not altered by breakdown or repair.

One accepted means for removing jobs from queues which have not completed processing is the use of negative customers [3]. A potential drawback to this approach can be a lack of synchronization with motivating events, which affects resulting performance measures. We explore the use of an enriched MAP structure to model queues with this behaviour, and associated re-sampling.

To examine the importance of synchronized loss, we contrast the results from using synchronized loss with those derived from using independent streams of negative customers to approximate the loss rates. We find that the queue length distribution due to synchronized losses cannot be adequately matched, particularly at the full and empty queues, hence mis-calculating loss, blocking and utilization.

We can extend our generalised modulation concept to model re-sampling of a job when a processor recovers. This may correspond to a situation where a deactivated server has local persistent storage which takes jobs from the queue.

The jobs lost (and resampled) may be batched. If these batches are of fixed size, then the appropriate terms in the Kolmogorov balance equations are simply shifted to transition from/to queue lengths differing by the correct amount. The diagonal transition formulation is also compatible with techniques [2] for including geometrically batched processes.

In this paper, we provide Kolmogorov balance equations governing synchronized loss/resample behaviour, outline solution methods, and present some results for a simple example which emphasizes the significance of synchronization. The result is an interesting mathematical structure, with potentially important consequences for the steady state of queues in such systems.

2 Breakdowns and repairs

In a uniform multi-processor server, we have N processors, each with the same processing rate¹ μ . These break down independently as a Poisson point process with rate b . Inactive processors are similarly repaired

¹ Heterogeneous servers are provided by taking a Cartesian product of the ensemble of modulation processes for each server's activity

independently at rate r . The distinct combinations of activity/inactivity are represented as distinct phases in a Markov modulation process. The joint state probabilities of queue length and modulation state form a (semi-)finite 2D lattice strip. Figure 1 shows small excerpts of such a strip for a range of queueing behaviours associated with breakdowns and repairs. The horizontal dimension gives the number of active processors in each column, and the vertical dimension is the queue length, increasing upwards.

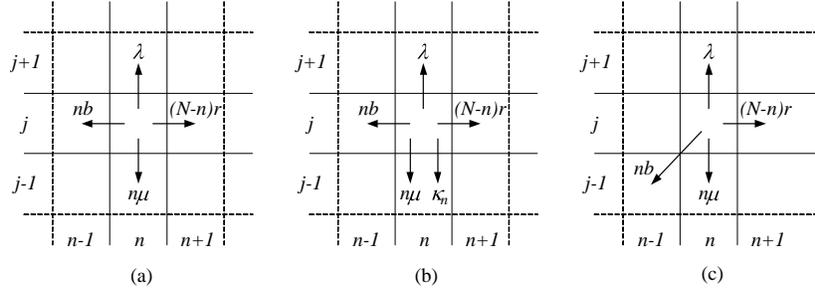


Fig. 1. State transitions in a queue with breakdowns and repairs. (a) is as in [1], (b) has losses due to remove-from-head negative customers, and (c) models synchronized losses.

Figure 1(a) shows the transitions away from state (n, j) in the semi-finite lattice describing the transitions in the Markov modulated queue with single breakdowns and repairs.

$$\mathbf{v}_{j-1}[\Lambda] + \mathbf{v}_j[Q - M - \Lambda] + \mathbf{v}_{j+1}[M] = \mathbf{0}$$

Where the modulator's instantaneous transition matrix Q is of the same form as [1] as given below². Figure 1 b) and c) show the loss of the job in process using respectively independent negative customers and synchronized transition.

The work in [1] allows for all active processors to break down simultaneously at a given rate, which we will call b_0 , and simultaneous repair of all inactive processors, which we will call r_0 . These are distinct from the breakdown and repair rates of individual processors, which are b and r respectively. When queue length does not change on breakdown or repair, this is entirely encapsulated within the transition structure of the modulator.

$$Q = \begin{pmatrix} -\Sigma \dots & Nr & & & r_0 \\ b_0 + b & -\Sigma \dots & (N-1)r & & r_0 \\ b_0 & 2b & -\Sigma \dots & \ddots & \vdots \\ \vdots & & & \ddots & \ddots \\ b_0 & & & & Nb - \Sigma \dots \end{pmatrix}$$

Associated with this modulation structure, the processing rates $\mu_i = (i-1)\hat{\mu}$ in diagonal matrix M give the modulated processing rates, as $\hat{\mu}$ is the rate of an individual processor. Matrix $\Lambda = \lambda I$ the arrival rate in all modulation states.

Losses on breakdown and/or resampling on repair require the addition of a vertical component to the originally horizontal breakdown and repair transitions. The resulting transitions are diagonal, except for breakdowns at the empty queue which are necessarily constrained to being horizontal. The treatment of repair transitions at a full queue, which cannot increase the queue length, deserves attention. Which job is lost depends on the queueing formalism. In a first come, first served queue, it may be more natural to lose the job at the end of the queue, and re-sample the job at the recovered processor.

² In our examples, b_0 and r_0 are set to zero for clarity.

It is also interesting to note that, when diagonal transitions are used, a particular quality of traditional Markov modulated queues is also lost: horizontal transition rates in the lattice are no longer independent of queue length. The “modulator” Q is not of itself a complete instantaneous generator for an ergodic chain over the majority of the queue.

3 Losses modelled by negative customers

To illustrate the effect of the loss being synchronized with the server breakdown, we briefly examine the results of attempting to approximate the behaviour of the queue using *independent* negative customers [3] in a Poisson arrival stream. This is contrasted with the behaviour of the queue with *synchronized* loss of the job in process, which may be considered to be caused by a negative customer removing from the head of the queue, which is *triggered* by the processor breakdown.

In section 7, we show example results of using the same arrival rate of independent negative customers as the processor breakdown rate, and compare this with using a rate chosen to result in a mean queue length equal to the synchronized loss distribution.

To correctly model the loss of a job in progress when the processor breaks down, we use appropriate transitions in the 2D lattice of joint states of queue length and modulation state. Let function $D(A)$ return a diagonal matrix of the row sums of A . The Kolmogorov balance equations for the queue with synchronized losses is as follows:

$$\begin{aligned} \mathbf{v}_j[R - D(R) + B - D(B) - M_j - A] + \mathbf{v}_{j+1}[M_{j+1} + B] &= \mathbf{0}, \text{ for } j = 0 \\ \mathbf{v}_{j-1}[A] + \mathbf{v}_j[R - D(R) - D(B) - M_j - A] \\ &\quad + \mathbf{v}_{j+1}[M_{j+1} + B] = \mathbf{0}, \text{ for } 0 < j < L \\ \mathbf{v}_{j-1}[A] + \mathbf{v}_j[R - D(R) - D(B) - M_j] &= \mathbf{0}, \text{ for } j = L < \infty \end{aligned}$$

Matrix M_j gives the processing rates $\mu_{m,j} = \min(j\mu, \mu_m)$ at queue length j , where $\mu_m = (m-1)\mu$ for a queue with N homogeneous processors. The matrix B provides the transitions in the lattice due to breakdowns (unspecified elements are zero):

$$B = \begin{pmatrix} 0 & & & & 0 \\ b & \ddots & & & \\ & \ddots & 0 & & \\ & & (N-1)b & 0 & \\ 0 & & & Nb & 0 \end{pmatrix}$$

The R matrix describes the repair behaviour, and is detailed below.

For an independent negative customer arrival, the transition is solely in the queue length dimension (vertical). We introduce a component in what we call the modulation dimension (horizontal) to represent the change in number of active processors.

As is normal practice with Markov modulated queues, the parameters of the queue’s behaviour are represented in matrix form, in order that we may associate the rates and batch size distribution parameters with the correct state occupation probabilities within vectors of the probabilities at a given queue length.

To achieve a diagonal transition, we effectively use B as an *off-diagonal* negative customer rate matrix. This contrasts with the negative customer rate matrix used in the independent approximation, which is diagonal, thus representing purely vertical transitions in the queue lattice strip, with the horizontal transitions due to breakdown being encapsulated in the *independent* modulation structure.

As we have mentioned, when we have synchronized loss (and repair later), the modulator matrix Q is no longer itself a generator of an ergodic chain except at the empty queue. For example, consider a system with one unreliable processor. Considering the balance equations given above, note that the part of the expression corresponding to the modulator Q in more traditional Markov modulated queues is as follows:

$$\begin{aligned} Q &= R - D(R) + B - D(B), \text{ for } j = 0 \\ &R - D(R) - D(B), \text{ for } 0 < j \leq L \end{aligned}$$

In a two state example, we have:

$$\begin{aligned}
 B &= \begin{pmatrix} 0 & 0 \\ b & 0 \end{pmatrix}, R = \begin{pmatrix} 0 & r \\ 0 & 0 \end{pmatrix} \\
 Q &= \begin{pmatrix} -r & r \\ b & -b \end{pmatrix}, \text{ for } j = 0 \\
 &= \begin{pmatrix} -r & r \\ 0 & -b \end{pmatrix}, \text{ for } 0 < j \leq L
 \end{aligned}$$

This Q matrix at levels $0 < j \leq L$ does not provide a horizontal transition back to modulation state 1 from 2. This is instead provided by the breakdown rates given in B . At the empty queue, no jobs are lost on breakdown of a server, so this transition becomes horizontal, and is incorporated into the Q matrix.

If processors are allowed to fail while unoccupied, this requires differentiation between broken processors which were occupied on breakdown, and those which were not. This would use three-states per processor, and require a large state space for the system as a whole comprising the joint state of all processors. Instead, we propose that a processor is unlikely to fail when unoccupied, and this leads to a modified breakdown rate in state (n, j) of $\min Nb, jb$, as only occupied processors break down.

4 Synchronized re-sampling on processor repair

We now consider allowing the job which was taken out of the system when the processor broke down to be re-sampled when the processor recovers.

The Kolmogorov balance equations a queue with synchronized loss on breakdown, and synchronized re-sampling of the job on repair is as follows:

$$\begin{aligned}
 \mathbf{v}_0[R - D(R) + B - D(B) - M_0 - \Lambda] + \mathbf{v}_1[M_1 + B] &= \mathbf{0}, \text{ for } j = 0 \\
 \mathbf{v}_{j-1}[\Lambda + R] + \mathbf{v}_j[-D(R) - D(B) - M_j - \Lambda] \\
 &\quad + \mathbf{v}_{j+1}[M_{j+1} + B] = \mathbf{0}, \text{ for } 0 < j < L \\
 \mathbf{v}_{L-1}[\Lambda + R] + \mathbf{v}_L[-D(R) - D(B) - M_L] &= \mathbf{0}, \text{ for } L < \infty
 \end{aligned}$$

If the queue is full, and a processor is inactive, the transition in our representation is horizontal, but care must be taken with the semantics of job re-sampling. It may be appropriate to discard the job at the end of the queue and replace it with one which had been ‘‘stored’’ with the inactive processor, but within the processor, effectively shifting all jobs in the queue back by one position. The meaning of the waiting room location notionally taken by the inactive processor in this case is not obvious. It may be better to disallow arrivals to the queue at values of $j \geq L - c$: this remains an unresolved issue which is to be analysed against practical interpretations of the synchronised activity, and may vary from application to application.

In many circumstances, it may not be worth dealing with these diagonal transitions, as the behaviour is closely approximated by the original formulation used in [1]. However, if we consider a more general formulation, in which the job in a broken or deactivated processor may not be re-sampled, the transition structure is somewhat richer.

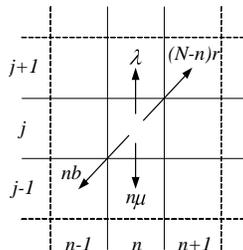


Fig. 2. State transitions in a queue with deterministic loss and resampling on breakdown/repair.

5 Probabilistic resampling

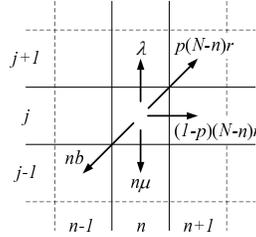


Fig. 3. State transitions in the main body of a queue with synchronized loss on breakdown and probabilistic re-sampling synchronized with repair.

Some circumstances may motivate the probabilistic discarding of jobs allocated to an inactive processor. For example, if we consider a media stream in which data may become obsolete, it may be that we choose not to re-sample it. This corresponds to a general transition structure as shown in figure 3 modelling probabilistic re-sampling, in which p is the probability of re-sampling the job, which is be worked into the Kolmogorov balance equations as follows:

$$\begin{aligned}
 \mathbf{v}_0[(1-p)R - D(R) + B - D(B) - M_0 - A] + \mathbf{v}_1[M_1 + B] &= \mathbf{0} \\
 \mathbf{v}_{j-1}[A + pR] + \mathbf{v}_j[(1-p)R - D(R) - D(B) - M_j - A] \\
 &\quad + \mathbf{v}_{j+1}[M_{j+1} + B] = \mathbf{0}, \text{ for } 0 < j < L \\
 \mathbf{v}_{L-1}[A + R] + \mathbf{v}_j[R - D(R) - D(B) - M_L] &= \mathbf{0}, \text{ for } L < \infty
 \end{aligned}$$

When $p = 1$, this is identical to the system in the previous section. This transition structure is a Markov arrival process (or MAP) with two classes of customer; one positive, one negative. Transitions in the modulation structure representing a repair are probabilistically accompanied by a positive customer arrival. Modulation transitions representing a breakdown are accompanied by the arrival of a negative customer. The independent arrivals, which cause strictly vertical transitions in the lattice are superposed and unaffected by this transition structure. The processing completion transitions, which also cause strictly vertical transitions, are at rates selected by the modulation state. Independent external arrivals may be worked into the system using standard methods, either using Kronecker products to construct a joint state space for arrivals and the intrinsic queue behaviour, or treating arrivals as G-type, which is often found to improve efficiency when the joint modulation state space is large.

6 Solving for the steady state

We may use a range of solution techniques to use the balance equations to provide the steady state of the queue. The two most common for a queue with a homogeneous region are matrix geometric methods and spectral expansion, described and compared in *e.g.* [7]. Since then, there have been a number of advances in matrix geometric iterative solution methodology to improve the range of problems for which it is more efficient. In summary, we find that the matrix geometric method is more efficient in the majority of circumstances, identified by low to medium utilization. Spectral expansion could be used in situations where the number of non-zero eigenvalues in the system is greater than the number of modulation states, in particular arising in some circumstances where arrivals are batched.

Some performance measure calculations are made simpler by access to the steady state of the queue in spectral form, and in this circumstance, it may be more efficient to use spectral expansion directly.

A repeating region The vector state occupation probabilities in a region of a queue described by a homogeneous matrix geometric series can be represented by a matrix geometric series. We will label the shortest queue length in the repeating region k in the following. For example, in the case of the single processor queue with synchronized losses on breakdown, the region of the joint state modulation queue length probability lattice lies between queue lengths 1 and $L - 1$, where L is the maximum queue length. When we do not consider resampling, we have:

$$\mathbf{v}_{j-1} \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix} + \mathbf{v}_j \begin{pmatrix} -r-\lambda & r \\ 0 & -b-\lambda-\mu \end{pmatrix} + \mathbf{v}_{j+1} \begin{pmatrix} 0 & 0 \\ b & \mu \end{pmatrix} = \mathbf{0}$$

At the empty queue, we have no downward transitions due to processing, and the loss component of a breakdown transition is absent, and since the first matrix in the series is not singular, the vector of state occupation probabilities is not part of the repeating region.

The solution to the linear homogeneous matrix equation given above is the sum of geometric series provided by the eigenvalues ξ_i of the characteristic equation $Q(\xi) = \sum_{k=0}^{n^u+n^d} \mathbf{v}_{j-n^u} Q_i$ found by setting $\det |Q(\xi)| = 0$, projected onto the corresponding left eigenvectors ψ_i from $\psi_i Q(\xi_i) = 0$.

Each eigenvalue/eigenvector pair defines a basis function component, and by summing the ensemble with each component scaled appropriately, we can satisfy the boundary conditions defined by the balance equations of the levels neighbouring the repeating region, and the normalization constraint.

At queue lengths j falling under the repeating region balance equation, we may use explicit eigensystem representation, or summarize this as a matrix geometric term, as follows:

$$\begin{aligned} \text{Spectral expansion: } \mathbf{v}_j &= \sum_{i=1}^{\epsilon^n} \alpha_i \xi_i^{j-k} \psi_i \\ \text{Matrix geometric: } \mathbf{v}_j &= \mathbf{v}_k A^{j-k} \end{aligned}$$

The coefficients α_i in the spectral expansion representation are free variables to be constrained by the boundary conditions imposed by the rest of the queue (the processor filling region, and in the case of a finite queue, the full queue region.) The elements of vector \mathbf{v}_k are free variables, and matrix A (constructed either directly from the eigensystem, or computed iteratively for higher efficiency) provides the progression between vectors on the series for successive j .

The boundary conditions are imposed by the inclusion of balance equations which include both explicit \mathbf{v} vectors outside the repeating region, and vectors defined by the eigensystem, or powers of A .

When the queue is infinite, any eigenvalues of magnitude greater than or equal to 1 (*i.e.* lying on or outside the unit disk in the Argand plane) must take zero coefficients, as their infinite sum does not converge, and hence cannot be normalized. If there are fewer eigenvalues less than one than modulation states, then iterative matrix geometric methods will generate a singular matrix A , and k should be chosen such that $j - k$ is greater than or (ideally) equal to zero at the lower bound of the repeating region.

7 Results

We present a number of simple examples of a single unreliable processor queue. For speed of execution, we employed an iterative matrix geometric solution method [4] to solve for the steady state. We use the terms *active phase* and *inactive phase* to refer to the modulation states representing the corresponding state of the processor.

We have the following B and R matrices in the Kolmogorov balance equations described earlier:

$$B = \begin{pmatrix} 0 & 0 \\ b & 0 \end{pmatrix}, R = \begin{pmatrix} 0 & r \\ 0 & 0 \end{pmatrix}$$

We use a short finite queue as it highlights the effect on the peak in queue length probability at the full queue due to the zero processing rate in the inactive phase. We see that this peak is de-emphasised when the independent negative customer approximation is used.

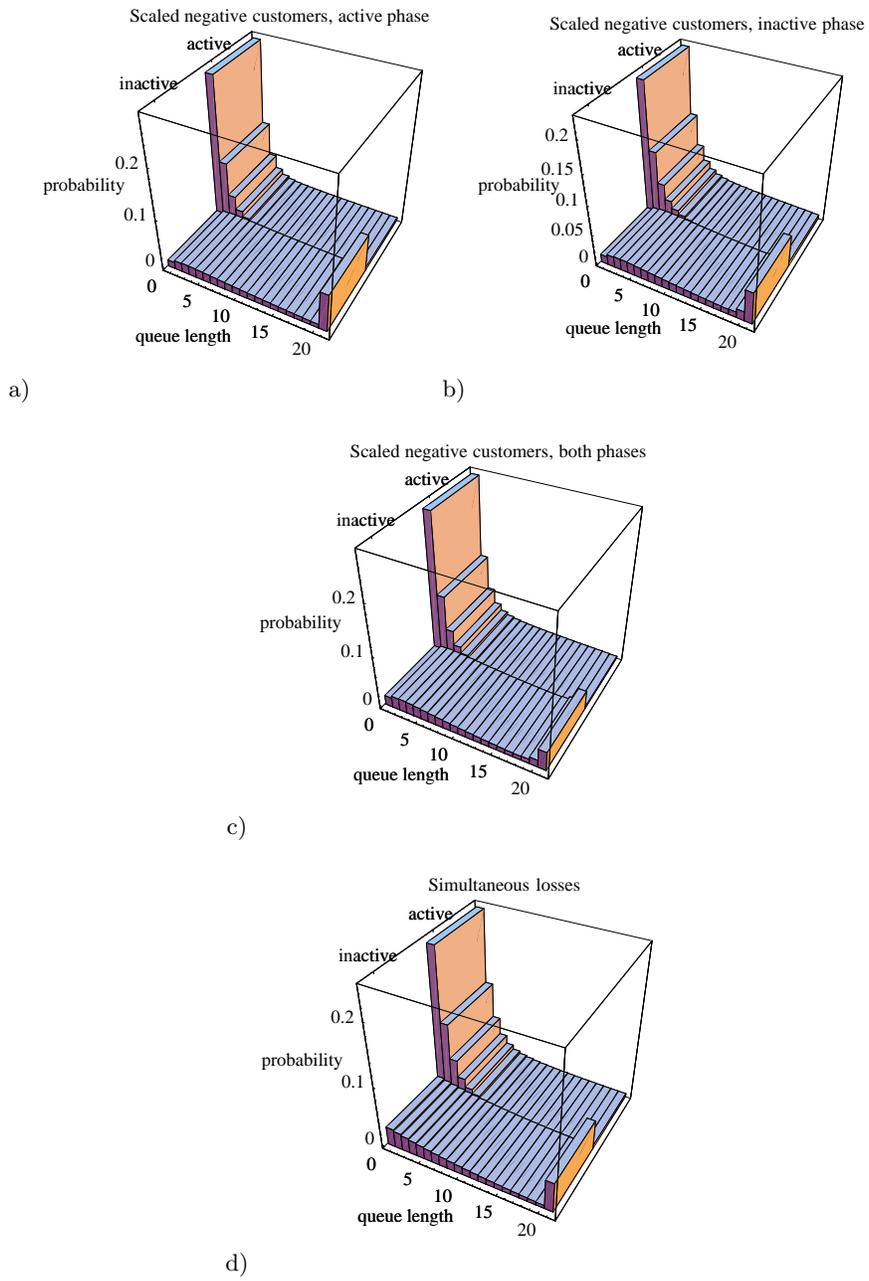


Fig. 4. Joint modulation state queue length probability plots for the synchronized loss case, and the three example approximations.

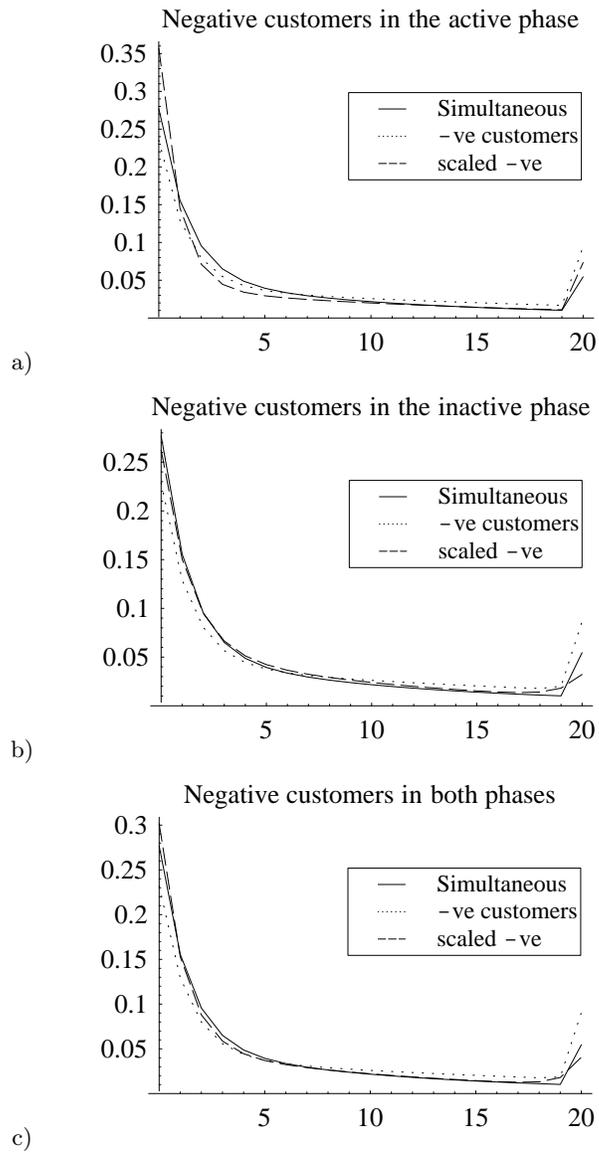


Fig. 5. Comparing queue length distributions from the best-estimate using negative customers in the *inactive*, *active* and both phases with the synchronized loss characteristic.

First, we calculate the queue length distribution of the queue when synchronized losses are enabled. This is the reference behaviour, and the behaviour resulting from use of independent negative customers is plotted in comparison with this.

We then use independent negative customers to approximate this behaviour. We first use a negative customer rate equal to the breakdown rate in the active phase, and show the distribution for this. We then multiply the negative customer rate by an adjustment factor to match the mean queue length of the reference behaviour. We show these distributions superposed for comparison.

Negative customers at the breakdown rate are then used in the inactive phase, and similar comparisons made, followed by an examination of the results of negative customers present in both phases at equal rates.

7.1 Observations on the effect of synchronization

Figure 4 shows the joint queue length activity state characteristics of the single unreliable processor queue with losses due to a) independent Poisson arrivals of negative customers (PAN) in the active state, b) PAN in the inactive state, c) PAN at equal rates in both activity states and d) with synchronized loss on breakdown (SLB). Note that the PAN characteristics all show a dip in the joint probability of very short queue lengths in the inactive state, which is absent in the SLB characteristic.

Figure 5 shows the queue length probabilities resulting from the independent negative customer approximation methods plotted against the reference characteristic with synchronized losses. The solid line is the SLB characteristic, the dotted line shows the result of using a negative customer rate equal to the breakdown rate, and the dashed line shows the characteristic after the negative customer rate has been set to a value which results in the same mean queue length as the reference behaviour.

The factors by which the negative customer arrival rate was multiplied to achieve the mean queue length matches shown (to two significant figures) are a) 17 for the active phase, b) 7.7 for the inactive phase and c) 5.0 for both phases (each being b before adjustment). These values were found by a bisection search. Figure 5a) shows that the utilisation is significantly under-estimated when using negative arrivals in the active phase, and that the full queue probability is over-estimated. If negative customers are used in the inactive state, the match is improved. Figure 5b) shows that when negative customers are used only in the inactive phase, the utilisation can be closely matched, but the full queue probability is underestimated. In figure 5c) when negative customers are used in both phases, a closer balance is struck: the utilisation is still well matched, and the full queue probability match is improved.

There may be scope for improving the match achieved by approximation using negative customers by careful choice of different rates in the two phases, but it will not be possible to avoid the dip in the queue length probability near the empty queue in the inactive phase, which is “topped up” from the active phase when synchronized losses are enabled. In figure 4, compare the steady rise toward the empty queue in the inactive phase with synchronized losses (top of the figure) with the dip in the characteristics using negative customers.

8 Further work

The identification of appropriate values for p , the probability of re-sampling a job on repair is an open question. We expect it will be appropriate to solve for this iteratively over solutions for the queue as a whole when attempting to effect, for example, a job obsolescence deadline. Identification of a class of problems for which p is of a functional form still maintaining the presence of a repeating region – whether exact or approximate – enabling the use of matrix geometric methods will be of interest.

9 Conclusions

The introduction of synchronized negative customers by re-modelling the transition pattern in the 2D lattice of states of a Markov modulated queue enhances the potential for accurate modelling of the effect of losses in an unreliable server. Synchronized re-sampling complements this behaviour to complete the picture in terms of transition patterns. These are achieved by using an enriched MAP which can decrease as well as

increase queue length within a single modulation structure. Common methods for independent modulated or G-type arrival streams arriving at modulated queues should be susceptible to inclusion within this regime with relatively minor modification.

References

1. I. Mitrani and R. Chakka. Spectral expansion solution for a class of Markov models: Application and comparison with the matrix-geometric method, *Performance Evaluation* **23** pp. 241-260, 1995.
2. P.G.Harrison, D.J. Thornley and H. Zatschler. Geometrically batched networks, in Proceedings ISICIS'02, 17th International Symposium on Computer and Information Sciences, CRC Press 2002
3. E. Gelenbe. Product form queueing networks with negative and positive customers, *Journal of Applied Probability* **28**, pp. 656-663, 1991.
4. Bini DA, Latouche G, Meini B. Solving matrix polynomial equations arising in queueing problems *Linear Algebra and its Applications* 340: 225-244 Jan 1 2002
5. R. Chakka and P.G. Harrison. A Markov modulated multi-server queue with negative customers - The MM CPP/GE/c/L G-queue. *Acta Informatica* **37**(11-12), pp. 881-919, 2001
6. Stephen Wolfram, *The Mathematica Book*, 4th ed., (Wolfram Media/Cambridge University Press, 1999)
7. Boudewijn Haverkort, Raymond Marie, Gerardo Rubino, Kishor Trivedi, *Performability Modelling Techniques and Tools*, John Wiley and Sons, 2001.
8. P.G. Harrison. The MM CPP/GE/c/L G-queue: sojourn time distribution, *Queueing Systems: Theory and Applications*, 41 (3): 271-298 JUL 2002