Geometrically batched networks

Peter Harrison, David Thornley, Harf Zatschler
Department of Computing,
Imperial College London,
UK
{pgh|djt|hz3}@doc.ic.ac.uk

Abstract

We develop approximate solutions for the equilibrium queue length probability distribution of queues in open Markovian networks by considering each queue independently, constructing its arrival process as the join of each contributing queue's departing traffic. Without modulation and non-unit batches, we would only need to consider mean internal traffic rates, modelling each queue as M/M/1 to give an exact result by Jackson's theorem. However, bursty traffic significantly affects steady state queue lengths; for given throughput, mean queue length varies linearly with mean batch size. All batch sizes are geometrically distributed, so each queue is Markovian and has known analytical solution. Our analysis is based on properties of the output processes of these queues, their superposition and splitting, to form the arrival processes at all queues. In general, this leads to a fixed-point problem for the network's equilibrium. The numerical results of our approach are compared with simulation and show promising accuracy.

1. Introduction

Batched processes provide enhanced descriptive power in queueing models; see [9], for example. This can be especially important for modelling the kind of bursty traffic typically observed in present day communication systems [1]. We have an automated formulation and solution mechanism that incorporates batches of geometrically distributed size in Markov modulated multiprocessor queues of finite or infinite capacity. Geometric distributions can be scaled and superimposed to produce a class of monotonic, unbounded convex probability density functions. Networks of such queues can be analysed by considering each queue in isolation, with an arrival process formed by the superposition of the departure processes from all connected queues, together with any external arrivals. A key issue is to approximate accurately this superposition.

More specifically, we are developing a Markov modulated queue in which all jumps in the queue length are geometrically distributed and which can accept a number of independent batched arrival processes. Unbounded transitions in the queue resulting from the use of geometric batch arrivals lead to Kolmogorov equations of 'infinite range'. Our method produces an ensemble of transformed balance equations of minimal, finite range, a large region of which can be represented as a linear homogeneous matrix recurrence relation in the vectors of the modulation state occupation probabilities at each queue length. This localization of the batched equations facilitates solution of unbounded queues via matrix geometric or spectral expansion methods, and is crucial to producing manageable complexity in queues with large, finite capacity.

We find that the ability of this queue to allow independent streams rather than aggregated arrivals [5,10] improves network solution accuracy. In this paper, we outline the formulation of the queue, describe the iterative solution of a network of such queues and examine the effect of different traffic approximation schemes.

2. Queue formulation

The Kolmogorov equations of a queue which includes geometrically batched processes (arrivals and/or service completions) can involve terms corresponding to an infinite range of queue lengths, but their structure allows transformation of the balance equations into a form that includes terms only from a small, local set of queue lengths. The original motivating algebraic concept behind this approach can be found in [3], where particular instances of batched queues are considered.

The transformation procedure has similar goals to Gaussian elimination, which *per se* is inapplicable to unstructured infinite systems. The balance equations are effectively diagonalised, but using an efficient procedure which takes advantage of the structure of their formulation.

Our general formulation solves queues with multiple processors, multiple service streams (multiple streams per processor allow the mixing of batches to construct more complex batch distributions in a Poisson point process), breakdowns and repairs, and both positive and negative arrivals. All the processes are modulated by an independent, finite state Markov chain. Each process may be independently described in every modulation state, and all parameters can be specified (as a constant) within each of a finite number of ranges of queue lengths.

Independent work by Chakka formulates solutions to queues with multiple arrival streams, based on the methods first used in [3]. Although these yield a satisfying theoretical motivation for the underlying algebraic approach, we develop an alternate, algorithmic approach. Problematically, the original methods have difficulty where multiple arrival (or departure streams) in a given modulation state have the same batch size distribution, or if there are different numbers of streams in different modulation states. Our new methods automatically take account of these effects during the construction of the balance equations.

2.1. A Repeating Region

It is well known that over a large range of queue lengths, the probability flux pattern does not change. In this region, the balance equations are of the form:

$$\sum_{i=-u}^{d} \mathbf{v}_{j+i} A_i = \mathbf{0}$$

Where $\mathbf{v}_j$ is the vector of modulation state occupation probabilities at queue length j, and A_i is the matrix of rates into state j with jump size i. We refer to this range of queue lengths, where the coefficients in the (transformed) Kolmogorov equations form a linear homogeneous recurrence relation, as the *repeating region*. The matrix geometric literature often refers to *non*-linear equations – this describes the polynomial in the matrix R used to generate the solution, of the form $\mathbf{v}_j = \mathbf{u}.R^j$.

3. Network Solution

We solve networks of geometrically batched queues iteratively using approximated departure traffic at each queue. The aggregate arrival process at each queue is then the superposition of the approximated departure processes of the feeding queues.

We initialise the solution with a coarse approximation based on the traffic equations of the network. For a network of M/M/c queues, this traffic would yield the correct solution for every queue, by Jackson's theorem. The presence of batches, modulation, finite queues and negative customers renders this procedure approximate. We therefore refine this solution in an iterative process involving three stages: re-approximating the departure link traffic of each queue, solving each individual queue based on that new traffic, then assessing convergence criteria and repeating if they are not met.

There is a wide range of methods for approximating link traffic, as has been widely explored in the literature. In [5], the traffic arriving at a queue is generally a single geometrically batched stream which has been calculated to approximate the *aggregate* behaviour of all the contributing departure processes. Our queue formulation allows us to treat the departure processes (which have been approximated) as individual streams, arriving independently at target queues.

3.1. Departure Process Approximation.

Our queues employ geometrically batched processing, which leads to bursty departures, and hence a requirement for batched link traffic. The distribution of batch sizes in the departure process of a queue with batched processing is not exactly geometrically distributed. The batch size distribution from a queue with length j is necessarily bounded above by j. If we approximate this with a geometric distribution which matches the mean batch size of the departures, this over-estimates the contribution of large batch sizes. The effect on subsequent queues is to cause an over-estimate of mean queue length. We find that the mean queue length resulting from geometrically batched arrivals increases linearly with the mean batch size for fixed throughput.

To improve the traffic approximation, we superpose the batched traffic with unbatched traffic (of unit batch size) to form an *enriched* approximation, as unbatched traffic tends to cause an under-estimate of mean queue length. The (per job) throughputs of these two approximating departure sub-streams sum to the throughput determined for individual jobs at the queue. There are a number of regimes that could be followed to make use of the two streams, and of course we could choose more streams. In this initial investigation, we seek to verify that the model of the network can be improved by superposition of departure traffic to form arrival processes.

When using single batched Poisson streams to approximate the departure processes, we compute the mean batch size, d_i say, of departures from each queue i. Each batch size distribution parameter is then set to give the same mean d_i, and each rate parameter to give the correct throughput, λ_i say. The enriched departure approximation at each queue is a superposition of a batched Poisson stream with mean batch size d_i and rate $\beta\lambda_i/d_i$ and an independent unbatched Poisson stream with rate $(1-\beta)\lambda_i$. This yields a different (smaller) mean batch size. The weights are chosen heuristically, for example in terms of node utilisation.

The arrivals to any given queue in a network are formed by the assumed independent superposition of the feeding queues' approximated departure processes. This superposition is simply achieved by presenting each link's traffic as an independent arrival stream to the queue. Each traffic component can be independently modulated, so the modulation structure of the target queue becomes the Cartesian product of those of the feeding departure process approximations, and any modulation of the queue's processing rate.

3.2. Solution of Individual Queues

Once the departure process approximations have been constructed at the beginning of an iterative step, we can solve the component queues based on the new arrival processes constructed as described earlier.

We formulate the equations using localizing transformations, which yield a finite recurrence relation for the repeating region (which is infinite in an unbounded queue), and explicit equations that impose the necessary boundary conditions at the empty and full (if finite capacity) queue. The solution to such a homogeneous linear matrix recurrence relation is formed as an inner product of a number of geometric terms and a set of basis vectors. These series can be represented explicitly, as in spectral expansion techniques [11], or implicitly using optimized matrix geometric methods [2,4,12].

Spectral expansion provides for an arbitrary number of components to the solution – in fact the exact number required. Matrix geometric methods provide a square M by M matrix, where M is the number of modulation states. This can have the correct number of degrees of freedom (free variables bound by the transformed equations), more than sufficient, or too few to give a solution. In queues using geometric batches and multiple arrival streams, we may find that the number d of eigenvalues is greater than M^2. To represent the solution using matrix geometric methods, we therefore would require $n > 1$ matrices R_i ($1 \le i \le n$).

$$\mathbf{v}_j = \sum_{i=1}^{n} \mathbf{u}_i R^j \text{, in matrix geometric terms}$$
$$= \sum_{i=1}^{d} \alpha_i \boldsymbol{\psi}_i \lambda_i^j \text{, in spectral expansion terms}$$

where the α_i are constants and (λ_i, ψ_i) are the eigenvalue-eigenvector pairs of a certain characteristic equation corresponding to the transformed Kolmogorov equations for the repeating region.

We use spectral expansion in the examples provided here, as we do not yet have an efficient means for finding a set of independent generator matrices R_i. We speculate that using the previous iteration's solution for an ensemble of independent R_i matrices at each queue as a starting point in the new iteration will improve efficiency.

The queues in our network take multiple independent arrival streams, leading to a matrix recurrence relation of degree $C+A$, where C is the number of independent batched service completion streams, and A is the number of independent batched arrival streams. With M modulation states, we therefore find up to $M(C+A)$ eigenvalues with independent eigenvectors associated with the characteristic equation. This number of eigenvalues is reduced by one for each instance of an arrival or departure process having a zero rate in a given modulation state, or a pair of processes having the same batch size distribution parameter. For example, with a single unreliable processor, in the modulation state corresponding to the inactive processor, there is no service completion stream, so there is one less eigenvalue than the maximum.

4. Results

We have taken a small class of network topologies and examined the effect different approximation strategies have on the accuracy of the steady state solution of the network. We compare our approximated results with simulation of the network in a standard fashion. The approximation method has a marked effect on queue length distributions, and we have found that a single target measure, mean queue length, sheds light on this effect.

The range of approximation methods are referred to in the graphs as follows: "simulated" comes from simulation, "Poisson" comes from using a single (unbatched) Poisson stream for departures, "aggregated" refers to the use of a single geometrically batched stream for arrivals, aggregating the behaviour of all contributing queues, "batched" refers to our use of independent batched streams from each contributing queue, and "weighted" refers to the use of enriched departure streams, in this case Poisson superposed with batched in appropriate proportion. In the networks considered, λ is the external arrival rate to server S, ρ_i is the load at queue C_i and ϕ_i is the service completion batch size distribution parameter at queue C_i. All of our analytical solutions were implemented in Mathematica® [13].

4.1. Simple Feed Forward

A feed forward network (figure 4, when n=1 and p_1=0) does not require iteration for its solution. In the following examples, λ=1.

Figures 1 and 2 show the results, in terms of the error in the mean queue length at C_1, of solving the network using Poisson, batched and weighted approximations. The batched approximation tends to overestimate the mean queue length, and the Poisson traffic leads to a serious underestimation. A combination of these two approximations can be used to correct these deficiencies, as in the weighted result. Our initial approach has been to share throughput between the batched and unbatched streams in a ratio ρ:(1-ρ), where ρ is the utilization of the queue in the current iteration. We chose this, as a relatively idle queue tends to produce small batches, and the busy queue, larger ones. The results are promising, and could be improved with a more sophisticated ratio.

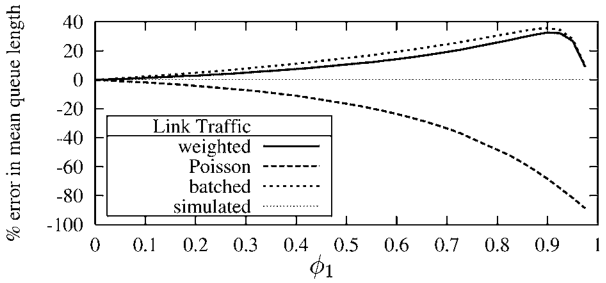

Figure 1. Percentage error in mean queue length of C_1 against ϕ_1 with ϕ_2=0.5.

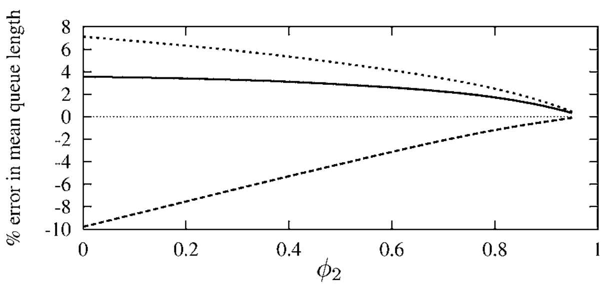

Figure 2. Percentage error in mean queue length of C_1 against ϕ_2 with ϕ_1=0.3.

4.2. A Tandem Queue

Consider a tandem queue with feedback, as obtained using the structure in figure 4, setting n=1, p_1>0, and λ=1. This is the simplest network that requires iteration for its solution.

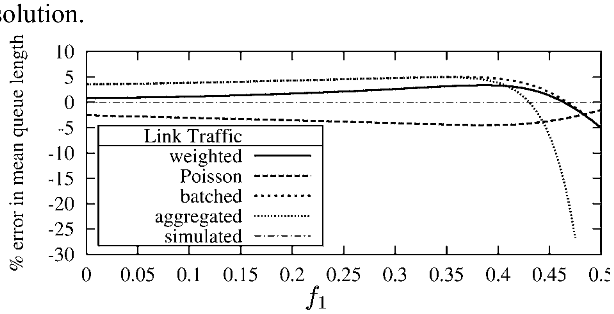

Figure 3. Error in mean queue length of C_1 for all strategies when varying feedback fraction f_1.

The Poisson underestimation continues up to a point where that approximation does not cause sufficient losses in the finite queue, where the real traffic often occurs in large batch sizes. Note that our use of independent input streams does not fare significantly better than the aggregation approximation in this case, as there are few streams. In this example, C_1 is saturated at f_1=0.5, which means that the distribution of departure batch sizes is closer to geometric, and the use of independent batched streams becomes more accurate.

4.3 Multiple Clients Feeding Back

We now take a single server node with external Poisson arrivals, and split its output between a number of heterogeneous clients which feed back a proportion of their output to the server. We use six client-nodes with independent processing characteristics. The server receives six independent arrival processes, as any independent Poisson streams aggregate to a single stream with the summed rate, and one client is unbatched. For the example results, we vary only f_6. The remaining parameters ϕ_1 through ϕ_6 are 0.0, 0.1, 0.3, 0.4, 0.5, 0.2 and f_1 through f_5 are all 0.4. The processor rate for each client is chosen so that its capacity is 1. The server has rate 4, batch parameter 0.3, and λ=3.

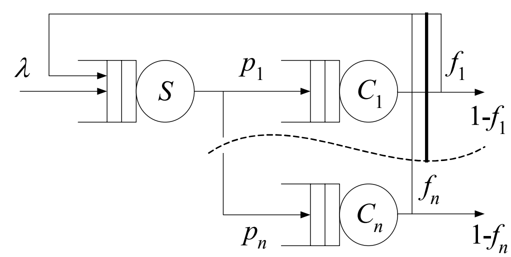

Figure 4. A network comprising a server S taking external Posson arrivals rate λ, and a proportion f_i fed back from n clients C_i, $1 \leq i \leq n$.

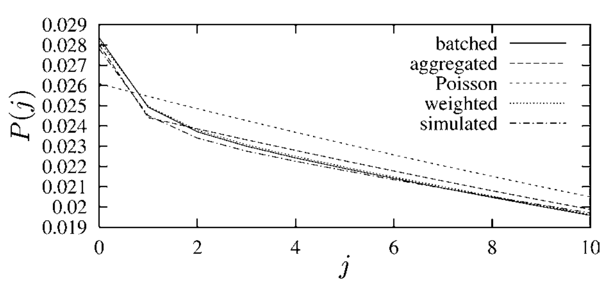

Figure 5. Queue length distribution at server for f_6=0.75.

Figure 5 shows part of the queue length distribution at the server. The arrival process is quite complex, yet all three batched approximations fare similarly, although enriched does best, and aggregated does least well. However, if we look at the mean queue length error at the server, aggregation looks most accurate. In contrast, the mean queue length error at C_6 is best approximated using enriched traffic, especially at higher feedback probabilities.

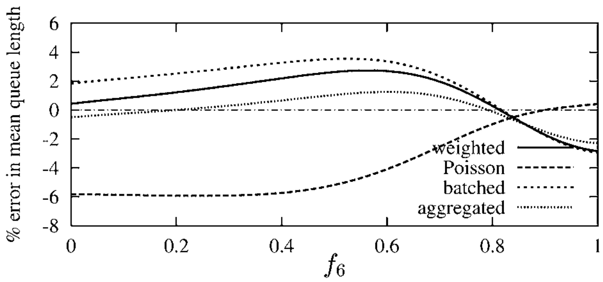

Figure 6. Mean queue length error for S, varying f_6.

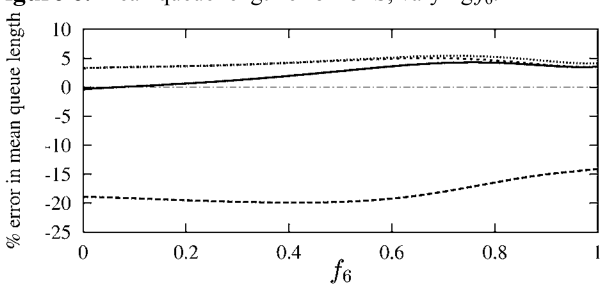

Figure 7. Mean queue length error for C_6, varying f_6.

5. Further work

We will investigate Markov modulated behaviour, to express unreliable processors, interrupted traffic, and more complex traffic behaviour (e.g. finite approximations to long tailed distributions), e.g. the internet.

We can approximate instantaneous batches using IPPs. A simple IPP has states *on*, with Poisson arrivals rate λ, and *off*, with no arrival stream, and transitions from *off* to *on* at rate a, and back at rate b. An arrival occurs on a transition from *off* to *on*, which we implement using an unusual 'diagonal transition' in the 2D Markov chain, so the mean number of arrivals in one visit to the active state is $(b+\lambda)/b$. As b increases to infinity (with constant $(b+\lambda)/b$) the process converges to a Poisson point arrival process with geometrically distributed batch sizes, rate a, and batch size distribution parameter $\lambda/(b+\lambda)$.

This structure increases the number of modulation states, which gives a single matrix R for the matrix geometric method more degrees of freedom, allowing solution for more arrival streams using a single matrix. Small inaccuracies are introduced because the batches are not instantaneous: the effect is to introduce over-sized batches in some circumstances. We are looking at the use of negative customers [6,7] to counteract this.

We also intend to address the question of response time distributions. This has been solved in [8] for a Markov modulated M/M/1 queue with batches, negative customers and infinite capacity. Our iterative method could therefore be applied immediately to the sojourn time distribution at a single such queue in a network, but the accuracy of the approximation would be more questionable, this metric being notoriously less robust.

6. Conclusions

Geometrically batched processes allow us to express burstiness in a network, and we expect to produce efficient and accurate means for solving for the steady state of such networks. Our superposition of modified batched streams has enabled matching of network behaviour to within closer bounds than previously achieved by aggregation methods.

7. Acknowledgements

This work is funded by EPSRC grant number GR/N16068.

References

[1] M. Bhabuta and P.G. Harrison. Analysis of ATM traffic on the London MAN, In *Proceedings of the 4th International Conference on Performance Modelling and Evaluation of ATM Networks*, Ilkely, Chapman and Hall, 1997.

[2] D.A. Bini, G. Latouche, B. Meini. Solving matrix polynomial equations arising in queueing problems *Linear Algebra and its Applications* **340**, 1 Jan 2002, pp 225-244

[3] 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), 2001, pp. 881-919,

[4] W. Fisher and K.S. Meier-Hellstern. The (Markov Modulated Poisson Process) MMPP Cookbook, *Performance Evaluation*, **18**, 22-July 1996, pp.149-171

[5] R.J. Fretwell and D.D. Kouvatsos. Correlated Traffic Modelling and Batch Renewal Markov Modulated Processes, In *Proc. 4th IFIP Workshop on Perf. Modelling and Evaluation of ATM Networks*, Chapman & Hall, 1997, pp. 20-44.

[6] E. Gelenbe. Product form queueing networks with negative and positive customers, *Journal of Applied Probability*, **28**, 1991, pp. 656-663.

[7] E. Gelenbe. G-Networks with signals and batch removal, *Probability in the Engineering and Informational Sciences*, 7, 1993, pp. 335-342.

[8] P.G. Harrison. The MM CPP/GE/c/L G-queue: sojourn time distribution, *Queueing Systems: Theory and Applications*, **41**(3), 2002, pp 271-298.

[9] P.G. Harrison and N.M. Patel, Performance Modelling of Communication Networks and Computer Architectures, Addison-Wesley, 1993.

[10] D.D. Kouvatsos and S.G. Denazis, ``Entropy Maximised Queueing Networks with Blocking and Multiple Job Classes", *Performance Evaluation* **17**, 1993, pp.189-205

[11] 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,** 1995, pp. 241-260.

[12] M. Neuts. Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach, Dover Publications, 1995.

[13] Stephen Wolfram, The Mathematica Book, 4th ed., Wolfram Media/Cambridge University Press, 1999.