Exploring correctness and accuracy of solutions to matrix polynomial equations in queues

David Thornley and Harf Zatschler
Department of Computing, Imperial College London
180 Queen's Gate, South Kensington, London SW7 2RH
djt & hz3@doc.ic.ac.uk

Abstract

Spectral expansion and matrix analytic methods are important solution mechanisms for matrix polynomial equations. These equations are encountered in the steady-state analysis of Markov chains with a semi-finite or finite two dimensional lattice of states, which describe a significant class of finite and infinite queues. We prove that the limited size of the eigenspectrum of the matrix geometric representation used in matrix analytic solution mechanisms confines its applicability to systems with a number of eigenvalues less than or equal to the dimension of the matrices used to form the solution. As well as proving this limitation, we relate our experience of a practical queue with generalized exponential traffic whose steady state cannot be represented using one or two rate matrices. We also provide an explanation for the numerical issues creating difficulty in finding matrix geometric solutions for finite queues. While we have not found a solution to these numerical issues, we do outline the steps required to enable complete matrix geometric solutions with larger eigenspectra, but which may not be efficient. On the other hand, we identify a case where care must be taken when using spectral expansion. Essentially, the eigensystem of a finite queue degenerates at saturation. We therefore formulate an enhanced spectral expansion method using generalized eigenvectors, which we prove gives a complete solution, even at saturation. We conclude that the state of the art requires the use of efficient matrix analytic methods where applicable, but correct solution in the general case is currently only guaranteed using generalized spectral expansion. We suggest that use of matrix analytic tools directed toward efficiency for solving a given queueing system should be preceded by an analysis of the eigensystem of the solution through spectral expansion, whether algebraic or numerical, to verify that the solutions produced by such tools are correct.

1 Introduction

Queues form the basis of many performabilty models, with a Markov modulation process instantaneously selecting arrival and departure descriptions from a finite set. The steady state joint distribution of queue length and modulation state occupation is at the core of many analyses. In this paper, we give an example of a form of queue for which the solution given by current favoured matrix analytic (MA) or geometric (MG) tools would be presumed correct, but which can in fact be incorrect. We perform an analysis which explains that the class of queues for which the matrix geometric or analytic solution is correct, while large, is strictly limited to those whose eigenspectrum is less than or equal to the dimension of the generator matrix (R) or auxiliary matrix (G) used in the matrix geometric and analytic approaches [3]. We discovered this shortcoming through analysis of spectral expansion (SE) methods.

The problem lies in the fact that the eigenspectrum of the solution for the steady state of a queue can be larger than matrix geometric and analytic methods can represent. In an infinite queue where the missing eigenvalues are all on or outside the unit disc, these would be ascribed a zero coefficient as they cannot be normalized. However, if the missing eigenvalues are less than one, or the queue is finite, the solution found may be incorrect. We have an example queue based on geometrically batched traffic for which both of these conditions can arise without unusual construction of the paramterization, leading to circumstances in which a generator matrix is inadequate. We prove that the matrix generator representation for the solution of queues can be incorrect by proving that the solution must embody the same eigenspectrum as a generalized spectral expansion (GSE) solution (see section 4) – which we prove to be complete – and providing an example queue for which this spectrum cannot be embodied by the generator matrix.

Our example queue uses generalized exponential traffic, and we explain how the matrix generated solution fails in section 5.5. We propose a potential mechanism for extend-

ing the matrix analytic approach in section 5.6 which could provide a correct solution algebraically, but which is subject to numerical issues briefly examined in section 6.2.

After establishing the background of a requirement for the solution to matrix recurrence relations in performability problems in section 2, we provide an exposition of the spectral expansion solution in section 3, and explain the degeneracy of the solution at saturation, which requires GSE, given in section 4. We then consider the MG and MA solutions in section 5 and relate them to that of GSE in section 5.1, then describe an extension the MG solution in section 5.2, analyse the size of the eigenspectrum of the solution in section 5.3, then show that the MG and MA solutions can be inadequate in section 5.5. We propose a means for enlarging the class of finite queues soluble using MG and MA by folding the system in section 5.6, but demonstrate in section 6 with an analysis of the accuracy of the numerical solution that this may be problematic. The methods for finding solutions to these problems are not considered, as we are principally concerned with correctness.

2 Background

The system performance analysis problem is commonly cast as a probabilistic state-transition system. In particular, a common tool is the Markov modulated queue, whose state space of queue length by modulation state takes the form of a semi-finite 2D lattice strip of states. The modulation state is selected from a finite set, which might represent the state of activity of a processor, or type of traffic currently arriving. The queue length may be finite or infinite. We consider systems in which transitions between these joint modulation and queue length states occur as Poisson point processes, so the solution to their steady state occupation probabilities is found by using Kolmogorov probability flux balance equations. The literature is replete with methods for approximating or reducing systems with unbounded queue length jumps or non-Poisson transitions to systems in which the transitions are bounded to a finite range. This leads to the major part of the queue complying with a simple linear homogeneous recurrence relation in a vector-valued state.

State transitions in such problems exhibit a "repeating region" within which neighbouring probability vectors $\overrightarrow{\mathbf{v}}_j$ follow a matrix recurrence relation.

$$\sum_{i=0}^{M} \overrightarrow{\mathbf{v}}_{j+i} \underline{\mathbf{A}}_i = \overrightarrow{\mathbf{0}} \tag{1}$$

The order M of this recurrence relation is determined by batch transitions within the queue. Its order is given by the sum of the maximal batch size for downward transitions and the maximal one for upward transitions in the queue, *e.g.*when arrival and departure batches are of unit size, its order is 2. Queues with geometrically distributed customer batches exhibit arbitrarily large batch sizes, so the matrix recurrence relation could be of infinite order. To enable the solution of such problems using spectral expansion or matrix geometric methods, we use an algorithm introduced in [13] that transforms the problem into an equivalent whose repeating region contains only finite batch sizes. The transformed equations are called *localized balance equations*. In the transformed system, we find that $M = n^{up} + n^{down}$, where n^{up} is the number of distinct, geometrically batched upward transitions (customer arrivals) and n^{down} similarly the number of distinct, geometrically batched downward transitions (service completions or customer removal via negative customers).

Two related methods applicable to solve for repeating regions in queues are derivatives of the matrix geometric method and the spectral expansion method. The methods differ in the representation of the relationship between adjacent levels.

Spectral expansion, resurrected and advocated by Mitrani and Chakka in [8] (after it was dismissed as improper in 1981 by Neuts [10]), on the other hand calculates repeating region SOPs via linear combinations of appropriately scaled eigenvectors inherent to the recurrence relation.

The questions of relative stability and efficiency of the two methods is vexed, as it pivots on the development of new computational routines. For example, in 1981 the MG method was based on an iterative Simple Substitution (SS) [10] method which took many iterations to converge, especially for heavily-loaded systems.

A major improvement for MG methods was introduced in 1993 by the Logarithmic Reduction (LR) algorithm [7] of Latouche and Ramaswami. Each iteration is computationally more expensive, but many fewer iteration steps are required, even for systems whose load approaches unity. Subsequent work by Bini and Meuni [1] improved on this further. However, even more recently, the numerical stability of Ramaswami's original approach have proven important, as discussed in a recent enlightening treatment of the efficient and powerful ETAQA method for calculating moments of queue length distributions fo M/G/1-type processes [6].

The SE method for steady state analysis of queues [8], proposed in 1995, depends heavily on mature eigenvalue and eigenvector packages [14, 12]. We cannot predict the formulation of a method that will improve efficiency by a similar degree to that seen between simple substitution and logarithmic/cyclical reduction for matrix geometric methods. Despite this, matrix analytic and geometric solutions can be wrong, and this has not previously been recognized, and an efficient but incorrect solution may be undesirable depending on the character of the error.

3 Spectral expansion solution

The spectral expansion solution to the linear homogeneous matrix equation (1) is the sum of geometric series provided by the eigenvalues ξ_i of the characteristic matrix

$$\underline{\mathbf{Q}}(\xi) = \sum_{k=0}^{M} \xi^k \underline{\mathbf{A}}_k \tag{2}$$

found by setting $\det \underline{\mathbf{Q}}(\xi) = 0$. Each of the eigenvalues ξ_i is then used to determine its corresponding eigenvector $\overrightarrow{\boldsymbol{\psi}}_i$ from the set of linear equations $\overrightarrow{\boldsymbol{\psi}}_i \underline{\mathbf{Q}}(\xi_i) = 0$. Spectral expansion uses a solution representation of the following form:

$$\overrightarrow{\mathbf{v}}_j = \sum_{i=1}^{n^*} \alpha_i \lambda_i^j \overrightarrow{\boldsymbol{\psi}}_i \qquad \epsilon^b \leq j \leq \epsilon^t \tag{3}$$

Where $\overrightarrow{\mathbf{v}}_j$ is the vector of state occupation probabilities for a given queue length. Let there be n^* eigenvalue (λ_i) / eigenvector ($\overrightarrow{\boldsymbol{\psi}}_i$) pairs. Each of these defines a basis function component, and by summing the ensemble with each component scaled appropriately, we can satisfy the boundary and normalization conditions.

When the queue is unbounded, any eigenvalues of magnitude greater than or equal to 1 (*i.e.* lying on or outside the unit disk in the Argand diagram) must take zero coefficients, as their infinite sum does not converge, and hence the solutions cannot be normalized. Eigenvectors with zero eigenvalue can also not contribute to the solution except at the boundary where they form the kernel (raised to the power zero).

3.1 Number of eigenvalues

The number of eigenvalues n^* pertaining to this system is dictated by the polynomial order of the determinant of the characteristic matrix in ξ as well as the size of the matrices involved. Thus, before considering zero eigenvalues and those situated outside the unit disc, we have $n^* = MN$

Zero eigenvalues occur when one or more of the $\underline{\mathbf{A}}_i$ matrices in expression (2) are singular. For each increment that the rank of $\underline{\mathbf{A}}_i$ differs from its dimension M, exactly one additional zero eigenvalue is introduced.

Normally, every one of the n^* eigenvalue/eigenvector pairs that have not been dropped corresponds to one of the free variables α_i that the SE method uses to constrain the repeating region. An exception to this occurs when an eigenvalue is complex. The characteristic equation derived from (2) has only real coefficients so that any complex eigenvalues appear in complex conjugate pairs. The geometric series of each of these traces a spiral in the Argand diagram. Since the solution must be real, the imaginary parts must cancel out. As a consequence, the corresponding α coefficients must also be mutually complex conjugate, reducing the degrees of freedom by one for every complex conjugate eigenvalue pair encountered.

3.2 Limitations of the SE method

The practical deployment of the spectral expansion method has not changed significantly since early treatments, for example by Mitrani in 1995 [8]. The main desirable improvements would be in terms of efficiency, as solutions are generally accurate. The only fundamental issue that needs to be addressed occurs in the vicinity of saturation. A queue is said to be saturated when customers within the repeating region arrive with the exact overall rate at which they depart (due to service and negative customers).

When a queue approaches this saturation, a non-unit eigenvalue, smaller in magnitude than 1, approaches the absolute value 1. There is always another unit eigenvalue, for example in a Markov modulated queue it is the one associated with eigenvector $\overrightarrow{\pi}$, the equilibrium solution vector of the modulation chain.

The proximity of these two eigenvalues at 1 introduces instability in the solution and eventually, at the exact point of saturation (which of course is only relevant for a finite capacity queue), the eigenvalue 1 occurs with multiplicity two, *i.e.* is replicated. However, there exists only one eigenvector for this eigenvalue. To see why, we consider that if the unit eigenvalue were to have an eigenspace of dimention more than one, it would imply that the modulator has more than one distinct steady state in a Markov modulated queue. In [8], it is stated that in practice there is a full set of linearly independent eigenvectors. We also find this, except in the case described above.

4 Generalized Spectral Expansion

In the steady state solution, one eigenvector disappears at the exact saturation point. This discontinuity in the number of solution eigenvectors causes conditioning problems in the linear system used to evaluate the free variables imposing the boundary conditions.

We extend the analysis of the characteristic matrix $\underline{\mathbf{Q}}(\xi)$ using generalized eigenvectors which removes the discontinuity and allows us to prove that our generalized spectral solution is complete. Running through this analysis also allows us to identify the size of the eigensystem of the problem. Our work with multiple streams of geometric batches [13] revealed these shortcomings, and we now briefly outline the stages of analysis.

Consider a recurrence relation of order M.

$$\sum_{k=0}^{M} \overrightarrow{\mathbf{v}}_{j+k} \underline{\mathbf{A}}_k = \overrightarrow{\mathbf{0}} \qquad \epsilon^b \leq j \leq \epsilon^t \tag{4}$$

where ϵ^b and ϵ^t are the bottom and top levels of the queue at which this relation holds, *i.e.* the bottom and top of the repeating region.

If $\underline{\mathbf{A}}_M$ is non-singular, we can re-write this as

$$\overrightarrow{\mathbf{v}}_{j+M} = -\sum_{k=0}^{M-1} \overrightarrow{\mathbf{v}}_{j+k} \underline{\mathbf{A}}_k \underline{\mathbf{A}}_M^{-1}$$

or, equivalently

$$\overrightarrow{\mathbf{v}}_{j+M} = [\overrightarrow{\mathbf{v}}_j, \ldots, \overrightarrow{\mathbf{v}}_{j+M-1}] \underline{\mathbf{W}} \tag{5}$$

where

$$\underline{\mathbf{W}} = \begin{pmatrix} -\underline{\mathbf{A}}_0 \underline{\mathbf{A}}_M^{-1} \\ \vdots \\ -\underline{\mathbf{A}}_{M-1} \underline{\mathbf{A}}_M^{-1} \end{pmatrix}$$

When $\underline{\mathbf{A}}_M$ is singular, but $\underline{\mathbf{A}}_0$ is not, we can similarly write

$$\overrightarrow{\mathbf{v}}_j = -\sum_{k=1}^{M} \overrightarrow{\mathbf{v}}_{j+k} \underline{\mathbf{A}}_k \underline{\mathbf{A}}_0^{-1}$$

and continue in an analogous manner. When both $\underline{\mathbf{A}}_M$ and $\underline{\mathbf{A}}_0$ are singular, we can use a transformation of the variable $\xi \mapsto \frac{1+\xi}{1-\xi}$ in (2). Such a transformation has already been mentioned in [8]. The resulting matrix polynomial has a modified matrix $\underline{\mathbf{A}'}_0$ that is always non-singular, so the above methods can be used for the solution process. The eigenvalues found are transformed back via $\xi \mapsto \frac{\xi-1}{\xi+1}$.

To enable the use of standard matrix operations and measures that require square matrices, we add a set of $M-1$ vector equations of the form $\overrightarrow{\mathbf{v}}_{j+k} = \overrightarrow{\mathbf{v}}_{j+k}$ to the one in (5) that do not change the solution set of the original matrix equation as follows.

$$[\overrightarrow{\mathbf{v}}_{j+1}, \overrightarrow{\mathbf{v}}_{j+2}, \ldots, \overrightarrow{\mathbf{v}}_{j+M}] = [\overrightarrow{\mathbf{v}}_j, \overrightarrow{\mathbf{v}}_{j+1}, \ldots, \overrightarrow{\mathbf{v}}_{j+M-1}] \underline{\mathbf{F}}$$

where

$$\underline{\mathbf{F}} = \begin{bmatrix} \underline{\mathbf{0}} & \underline{\mathbf{0}} & \underline{\mathbf{0}} & \cdots & \underline{\mathbf{0}} & -\underline{\mathbf{A}}_0 \underline{\mathbf{A}}_M^{-1} \\ \underline{\mathbf{I}} & \underline{\mathbf{0}} & \underline{\mathbf{0}} & \cdots & \underline{\mathbf{0}} & -\underline{\mathbf{A}}_1 \underline{\mathbf{A}}_M^{-1} \\ \underline{\mathbf{0}} & \underline{\mathbf{I}} & \underline{\mathbf{0}} & \cdots & \underline{\mathbf{0}} & -\underline{\mathbf{A}}_2 \underline{\mathbf{A}}_M^{-1} \\ \vdots & \vdots & \vdots & & & \vdots \\ \underline{\mathbf{0}} & \underline{\mathbf{0}} & \underline{\mathbf{0}} & \cdots & \underline{\mathbf{I}} & -\underline{\mathbf{A}}_{M-1} \underline{\mathbf{A}}_M^{-1} \end{bmatrix} \tag{6}$$

Define a new, compound, $1 \times NM$ solution vector

$$\overrightarrow{\mathbf{w}}_j = [\overrightarrow{\mathbf{v}}_j, \overrightarrow{\mathbf{v}}_{j+1}, \ldots, \overrightarrow{\mathbf{v}}_{j+M-1}]$$

This gives

$$\overrightarrow{\mathbf{w}}_{j+1} = \overrightarrow{\mathbf{w}}_j \underline{\mathbf{F}} \Leftrightarrow \overrightarrow{\mathbf{w}}_{j+i} = \overrightarrow{\mathbf{w}}_j \underline{\mathbf{F}}^i \quad \epsilon^b \leq j \leq j+i \leq \epsilon^t \tag{7}$$

An analysis of this system using the Jordan form and generalized eigenvectors appears in [4]. Here, we reproduce the expression for the generalized spectral expansion of the solution to the queueing system. This is required to represent the solution to the queue when an eigenvalue is repeated, as occurs when the queue saturates exactly. The standard spectral expansion solution is a special case of this.

We express the solution using generalized eigenvectors as follows, this time for $j > 0$:

$$\overrightarrow{\mathbf{v}}_{\epsilon^b+j} = \sum_{k=1}^{r} \left[\sum_{l=1}^{s_k} \alpha_{k,l} \left[\sum_{m=0}^{l-1} \binom{j}{m} \lambda_k^{j-m} \overrightarrow{\boldsymbol{\zeta}}_{k,l-m} \right] \right] \tag{8}$$

If all Jordan blocks $1 \ldots r$ of $\underline{\mathbf{F}}$ are of size $s_k = 1$, we obtain the familiar spectral expansion representation

$$\overrightarrow{\mathbf{v}}_{\epsilon^b+j} = \sum_{k=1}^{NM} \alpha_{k,1} \lambda_k^j \overrightarrow{\boldsymbol{\zeta}}_{k,1}$$

In practice the need for generalized eigenvectors in addition to simple eigenvectors has, so far, only been found to arise in a few special cases. As mentioned previously, one of these is given by a Markov modulated finite queue at "saturation" , *i.e.* where, within the repeating region, the mean customer arrival rate exactly equals the mean customer departure rate. In this case, a generalized eigenvector with eigenvalue 1 exists. This generalized eigenvector, conjointly with the eigenvector $\overrightarrow{\pi}$ given by the steady-state of the modulator $\overrightarrow{\pi} Q = \overrightarrow{\mathbf{0}}$, that is present in the solution to every queue, form a Jordan block of size 2.

5 Matrix Geometric solution

Unlike spectral expansion, which represents the probabilities within the repeating region ($\epsilon^b \leq j \leq \epsilon^t$) using eigenvalues and vectors, matrix geometric methods [10] utilize a non-singular matrix $\underline{\mathbf{R}}$ to describe changes in probability between two adjacent levels *i.e.*

$$\overrightarrow{\mathbf{v}}_{j+1} = \overrightarrow{\mathbf{v}}_j \underline{\mathbf{R}} \Leftrightarrow \overrightarrow{\mathbf{v}}_{j+i} = \overrightarrow{\mathbf{v}}_j \underline{\mathbf{R}}^i \quad \epsilon^b \leq j \leq j+i \leq \epsilon^t \tag{9}$$

or, by anchoring the expression at the bottom of the repeating region,

$$\overrightarrow{\mathbf{v}}_j = \overrightarrow{\mathbf{f}} \underline{\mathbf{R}}^{j-\epsilon^b} \qquad \epsilon^b \leq j \leq \epsilon^t \tag{10}$$

Here, $\overrightarrow{\mathbf{f}}$ takes the place of the α's in SE to be constrained by the boundary conditions imposed by the rest of the model.

The power we raise the matrix $\underline{\mathbf{R}}$ to could be offset by any integer as any change here can be absorbed in a corresponding change in the value of $\overrightarrow{\mathbf{f}}$.

By substituting (10) into the matrix recurrence (1) we see that a solution using this representation can only exist if $\underline{\mathbf{R}}$ satisfies the matrix polynomial equation

$$\sum_{k=0}^{M} \underline{\mathbf{R}}^k \underline{\mathbf{A}}_k = \underline{\mathbf{0}} \tag{11}$$

Efficient methods of solving for $\underline{\mathbf{R}}$ are presented in [7] and [9].

5.1 Correspondence in eigenmodes

Let $\hat{\underline{\mathbf{R}}}$ be a non-singular matrix[1] found using one of the above methods that satisfies the matrix polynomial equation (11). Having full rank N, the matrix $\hat{\underline{\mathbf{R}}}$ has N non-zero eigenvalues ζ_n and corresponding eigenvectors $\overrightarrow{\boldsymbol{\chi}}_n$.

We can uniquely express the vector $\overrightarrow{\mathbf{v}}_{\epsilon^b}$ as a linear combination of these eigenvectors,

$$\overrightarrow{\mathbf{v}}_{\epsilon^b} = \sum_{n=1}^{N} \gamma_n \overrightarrow{\boldsymbol{\chi}}_n \tag{12}$$

Then from equation (10) we have

$$\overrightarrow{\mathbf{v}}_j = \sum_{n=1}^{N} \gamma_n \zeta_n^{j-\epsilon^b} \overrightarrow{\boldsymbol{\chi}}_n \qquad \epsilon^b \le j \le \epsilon^t \tag{13}$$

We now compare this expression to the one derived for spectral expansion

$$\overrightarrow{\mathbf{v}}_j = \sum_{i=1}^{n^*} \alpha_i \xi_i^j \overrightarrow{\boldsymbol{\psi}}_i \qquad \epsilon^b \le j \le \epsilon^t \tag{14}$$

If the Matrix Geometric method gives a correct solution, it must be equal to that derived when using spectral expansion[2], *i.e.*

$$\sum_{n=1}^{N} \gamma_n \zeta_n^{j-\epsilon^b} \overrightarrow{\boldsymbol{\chi}}_n = \sum_{i=1}^{n^*} \alpha_i \xi_i^j \overrightarrow{\boldsymbol{\psi}}_i \qquad \epsilon^b \le j \le \epsilon^t \tag{15}$$

Proposition 1. *Any non-zero eigenvalue ζ_n of $\hat{\underline{\mathbf{R}}}$ (found using matrix geometric methods) is equal to an eigenvalue ξ_m found during spectral expansion and the sets of non-zero eigenvalues $\zeta = \{\zeta_n\}$ and $\xi = \{\xi_m\}$ satisfy $\zeta \subseteq \xi$.*

[1] Here we expand the case of an R with no Jordan blocks of size greater than one for brevity. In other cases, we use equation 8 and the resulting comparison of coefficients is equivalent but more involved

[2] Or indeed generalized expansion for more effort

Proof. Assume that ζ is not a subset of ξ. Therefore there exists a non-zero eigenvalue $\zeta_m \in \zeta$ such that $\zeta_m \notin \xi$. The equality (15) holds for any choice of the free variables γ_n, so that we chose $\gamma_n = \delta_{n,m}$ and try to determine the values of the free variables, α_i that the spectral expansion method uses to represent this particular solution.

$$\sum_{n=1}^{N} \delta_{n,m} \zeta_n^{j-\epsilon^b} \overrightarrow{\boldsymbol{\chi}}_n = \sum_{i=1}^{n^*} \alpha_i \xi_i^j \overrightarrow{\boldsymbol{\psi}}_i \qquad \epsilon^b \le j \le \epsilon^t \tag{16}$$

$$\zeta_m^{j-\epsilon^b} \overrightarrow{\boldsymbol{\chi}}_m = \sum_{i=1}^{n^*} \alpha_i \xi_i^j \overrightarrow{\boldsymbol{\psi}}_i \qquad \epsilon^b \le j \le \epsilon^t$$

$$\overrightarrow{\boldsymbol{\chi}}_m = \sum_{i=1}^{n^*} \alpha_i \frac{\xi_i^j}{\zeta_m^{j-\epsilon^b}} \overrightarrow{\boldsymbol{\psi}}_i \qquad \epsilon^b \le j \le \epsilon^t$$

$$\overrightarrow{\boldsymbol{\chi}}_m = \sum_{i=1}^{n^*} \frac{\alpha_i}{\zeta_m^{-\epsilon^b}} \left(\frac{\xi_i}{\zeta_m}\right)^j \overrightarrow{\boldsymbol{\psi}}_i \qquad \epsilon^b \le j \le \epsilon^t$$

The LHS of this expression is constant, whereas the right hand side varies with j, except in the following two cases.

Case 1. $\alpha_i = 0$ for all i. From this it follows that an eigenvector $\overrightarrow{\boldsymbol{\chi}}_m = \overrightarrow{\mathbf{0}}$. This contradicts the assumption that the matrix $\underline{\mathbf{R}}$ found using matrix geometric methods has full rank.

Case 2. Either $\frac{\xi_i}{\zeta_m} = 1$ or $\alpha_i = 0$ for all $1 \le i \le n^*$. It follows that for at least one i, we have $\xi_i = \zeta_m$. This contradicts the assumption that $\zeta_m \notin \xi$. □

Proposition 2. *Let $\overrightarrow{\boldsymbol{\chi}}$ be the set of non-zero eigenvectors of the matrix geometric solution matrix $\underline{\mathbf{R}}$, and similarly $\overrightarrow{\boldsymbol{\psi}}$ the set of eigenvectors found using spectral expansion. The set of (non-zero) eigen*vectors *satisfy the property that any member $\overrightarrow{\boldsymbol{\chi}}_m \in \overrightarrow{\boldsymbol{\chi}}$ lies within a subspace[3] spanned by members $\overrightarrow{\boldsymbol{\psi}}_n \in \overrightarrow{\boldsymbol{\psi}}$, with the same associated eigenvalue.*

Proof. Again, we chose $\gamma_n = \delta_{n,m}$ in the expression (15) and hence equation (16). Since there is no j-dependence on the RHS, $\alpha_i = 0$ whenever $\xi_i \neq \zeta_m$. From the above, we know that there is at least one i such that $\xi_i = \zeta_m$. Let there be a total of $k \ge 1$ equal eigenvalues. After re-labelling all terms for which we did not set the coefficient $\alpha_i = 0$, we have

$$\overrightarrow{\boldsymbol{\chi}}_m = \sum_{i=1}^{k} \frac{\alpha_i}{\zeta_m^{-\epsilon^b}} \left(\frac{\xi_i}{\zeta_m}\right)^j \overrightarrow{\boldsymbol{\psi}}_i \qquad \epsilon^b \le j \le \epsilon^t$$

i.e. $$\overrightarrow{\boldsymbol{\chi}}_m = \sum_{i=1}^{k} \frac{\alpha_i}{\zeta_m^{-\epsilon^b}} 1^j \overrightarrow{\boldsymbol{\psi}}_i \qquad \epsilon^b \le j \le \epsilon^t$$

[3] Usually, this subspace only has dimension 1, and hence the eigenvectors are pairwise parallel

It follows that $\overrightarrow{\boldsymbol{\chi}}_m$ is a vector that lies within the subspace spanned by $\left\{\overrightarrow{\boldsymbol{\psi}}_i\right\}_{1\leq i\leq k}$, where the spanning vectors are all eigenvectors of the SE method with the same eigenvalue, ζ_m. □

5.2 Backward series

A matrix solution to the forward series gives N of the eigenvalues and eigenvectors contained within the complete solution. We can use the backward series to gain access to up to N more.

$$\sum_{k=0}^{M} \overrightarrow{\mathbf{v}}_{j+k}\underline{\mathbf{A}}_{M-k} = \underline{\mathbf{0}} \tag{17}$$

This matrix recurrence has the same structure as the original forward series given in (4) involving $\underline{\mathbf{R}}$ and we can use the same algorithms to find a numerical approximation $\underline{\hat{\mathbf{B}}}$ to $\underline{\mathbf{B}}$ with $\overrightarrow{\mathbf{v}}_{j+1} = \overrightarrow{\mathbf{v}}_j\underline{\mathbf{B}}$, a solution to

$$\sum_{k=0}^{M} \underline{\mathbf{B}}^k\underline{\mathbf{A}}_{M-k} = \underline{\mathbf{0}} \tag{18}$$

When using the forward and backward series, the repeating region is represented as follows

$$\overrightarrow{\mathbf{v}}_j = \overrightarrow{\mathbf{v}}_{\epsilon^b}\underline{\mathbf{R}}^{j-\epsilon^b} + \overrightarrow{\mathbf{v}}_{\epsilon^t}\underline{\mathbf{B}}^{\epsilon^t-j}$$

This representation is only well-defined when the model in question has a finite state space. For infinite queues, for example, the backward series cannot be used.

5.3 Size of the eigenspectrum

Depending on the number of eigenvalues $n^* = |\xi|$ found by spectral expansion in relation to those of the matrix geometric method, which is necessarily fixed at $2N$ (or N), we can make three observations.

• If $n^* < 2N$, there are fewer eigenmodes present in the recurrence relation than both matrices contain. There are $2N - n^*$ eigenvalues between $\underline{\hat{\mathbf{R}}}$ and $\underline{\hat{\mathbf{B}}}^{-1}$ that are not constrained. We either use zero-eigenvalues (paired with some linearly independent eigenvectors) or use a number of eigenvector/eigenvalue pairs twice, once for $\underline{\hat{\mathbf{R}}}$ and another time for $\underline{\hat{\mathbf{B}}}^{-1}$. The introduction of zero-eigenmodes will make one of $\underline{\mathbf{R}}$ and $\underline{\mathbf{B}}$ singular, whereas using the same eigenmode twice yields two non-singular matrices. While normally singular matrices are to be avoided, it turns out that in practice the former method is numerically more stable for the subsequent solution process in which $\underline{\hat{\mathbf{R}}}$ and $\underline{\hat{\mathbf{B}}}^{-1}$ will be used. When using the formula $\overrightarrow{\mathbf{v}}_j = \overrightarrow{\mathbf{v}}_{\epsilon^b}\underline{\mathbf{R}}^{j-\epsilon^b} + \overrightarrow{\mathbf{v}}_{\epsilon^t}\underline{\mathbf{B}}^{\epsilon^t-j}$ to calculate, we raise both matrices to possibly large powers. This process is unproblematic for a singular matrix because its zero-subspace maps vectors to zero, thereby not allowing any errors to grow when raised to a power.

• When $n^* = 2N$, we have a direct correspondence in the solutions given by the spectral expansion and matrix geometric methods, with the only difference lying in the representation of the solution. We can seamlessly move from the matrix geometric method to spectral expansion by finding the eigenvalues of $\underline{\hat{\mathbf{R}}}$ and $\underline{\hat{\mathbf{B}}}$, and back by constructing $\underline{\hat{\mathbf{R}}} = \underline{\boldsymbol{\Psi}}^{-1}\underline{\boldsymbol{\Xi}}\,\underline{\boldsymbol{\Psi}}$ (similarly for $\underline{\hat{\mathbf{B}}}$), where the i-th column of $\underline{\boldsymbol{\Psi}}$ is $\overrightarrow{\boldsymbol{\psi}}_i$ and $\underline{\boldsymbol{\Xi}}$ is a diagonal matrix with i-th element ξ_i. Then, *e.g.*

$$\overrightarrow{\boldsymbol{\psi}}_i\underline{\hat{\mathbf{R}}} = \overrightarrow{\boldsymbol{\psi}}_i\underline{\boldsymbol{\Psi}}^{-1}\underline{\boldsymbol{\Xi}}\,\underline{\boldsymbol{\Psi}} = \overrightarrow{\mathbf{e}}_i\underline{\boldsymbol{\Xi}}\,\underline{\boldsymbol{\Psi}} = \xi_i\overrightarrow{\mathbf{e}}_i\underline{\boldsymbol{\Psi}} = \xi_i\overrightarrow{\boldsymbol{\psi}}_i$$

• In the case $n^* > 2N$, there are more eigenmodes present in the recurrence relation than can be represented by two matrices $\underline{\hat{\mathbf{R}}}$ and $\underline{\hat{\mathbf{B}}}$ of dimension N. In terms of the spectral expansion representation of a solution derived using the matrix geometric method, this forces the coefficients α_i to be zero whenever an eigenmode i is not represented by either the forward or backward matrix. Depending on the parameters of the model being solved, this can invalidate a solution obtained through the matrix geometric solution. If the queue is infinite, then the MG method will produce a solution, but this solution could be incorrect. If the queue is finite, however, the proposed solution will fail to meet the boundary conditions at the full and empty queue, so will be more obviously incorrect. Note that this failure to solve is distinct from that resulting from numerical inaccuracy, which we examine in section 6.

5.4 Modified MG method using multiple generator matrices

Should just one pair of matrices ($\underline{\hat{\mathbf{R}}}, \underline{\hat{\mathbf{B}}}$) be insufficient in accurately representing all present eigenmodes in the system, multiple solution matrices $\underline{\mathbf{R}}_k$ $(1 \leq k \leq k^F)$ and $\underline{\mathbf{B}}_k$ $(1 \leq k \leq k^B)$ of (11) and (18) have to be sought. The augmented probability representation is

$$\overrightarrow{\mathbf{v}}_j = \sum_{k=1}^{k^F} \overrightarrow{\mathbf{v}}_{k,\epsilon^b}\underline{\mathbf{R}}_k^{j-\epsilon^b} + \sum_{k=1}^{k^B} \overrightarrow{\mathbf{v}}_{k,\epsilon^t}\underline{\mathbf{B}}_k^{\epsilon^t-j} \tag{19}$$

The introduction of multiple solution matrices necessitates the presence of additional boundary conditions given by $\overrightarrow{\mathbf{v}}_{k,\epsilon^b}$ and $\overrightarrow{\mathbf{v}}_{k,\epsilon^t}$. This is only to be expected as the Spectral Expansion method also allocated one boundary condition per arrival stream in the modulated queues that we have covered, for example [13] and the Matrix Geometric method is simply an alternative representation of a solution within a repeating region.

As before it is the case that the eigenvalues and eigenvectors of matrix geometric solution matrices $\underline{\mathbf{R}}_k$ and $\underline{\mathbf{B}}_k^{-1}$ are subsets of ξ and $\overrightarrow{\psi}$ respectively. To minimize the number $k^F + k^B$ of necessary matrices and with it the computational complexity of the modified matrix geometric representation, care should be taken to choose solution matrices such that the sets of their eigenvalues (and eigenvectors) are pairwise disjoint.

A potential difficulty arises from the fixed size N of all matrices involved. If the number of necessary eigenmodes n^* is not divisible by N, one of the solution-matrices will have insufficient eigenvalues (and eigenvectors) to be fully specified. To remedy this situation, we pad the representation with linearly independent zero-eigenvalue eigenvectors.

So far no algorithm is known that does not involve prior knowledge of all eigenvalues and eigenvectors.

5.5 Limitations of the MG method

Our experimentation with spectral expansion has revealed important characteristics of a class of queues which cannot be solved using matrix geometric methods in their current form. The problem lies in the spectrum of the solution. Matrices as used in MG have a fixed number of eigenvalues, equal to the number N of modulation states. Currently known algorithms can provide either one or two of these matrices, depending on whether the queue is infinite or finite.

Infinite queues do not require the full set of eigenvalues, in particular only a number n^* of those which have magnitude strictly less than one. In our geometrically batched queues, the value of n^* has been found to be directly related to the number of distinct geometrically batched arrival processes, n^{up}, over all modulation states and for an infinite queue that is not saturated. In practice we find, disregarding degenerate[4] arrival streams, $n^* = n^{up}N$.

The MG method only provides one matrix, $\underline{\mathbf{R}}$, to represent the solution, and for $n^{up} > N$ it is therefore not able to accurately solve the problem.

When queues have a finite buffer, contributions from eigenvalues outside the unit disc can also be normalized and need to be included in the solution process. As in the infinite case, the number of these depends on n^{up}, the number of distinct arrival processes. In addition to these, n^{down}, the number of distinct service and killing processes need to be considered and we have $n^* = (n^{up} + n^{down})N$ for the finite case. By way of forward and backward series, explained in section 5.2, the MG method can represent a total of $2N$ eigenvalues, so that an exact solution for a queue where $n^{up} + n^{down} > 2N$ is not possible in general.

In addition to this limitation, there is the practical obstacle that *e.g.* the LR matrix geometric method [7] is only applied to third order matrix recurrence relations, leading to the solution of a quadratic matrix polynomial. Matrix analytic techniques seek the solution to a higher order problem, but still represent the solution using a single matrix, or potentially with a pair of matrices (one forward and backward as we suggest in section 5.2).

5.6 Queue Folding

The twin problem of too few eigenvalues and too high order of recurrence relation can be overcome in some circumstances by *folding* the state-space of the queue. This procedure, already mentioned for fixed batch size problems in [5], is equivalent to *reblocking* techniques [2], but equally applicable to queues with geometrically distributed batches, once an equivalent set of localized balance equations is found [13]. In reblocking, the state transition matrix for the underlying Markov chain is repartitioned into larger blocks (matrices) than would arise naturally from the problem so be solved. For example, when solving modulated queues with N modulation states, the natural block size is N and consequently $\underline{\mathbf{R}}$ and $\underline{\mathbf{B}}$ are $N \times N$. With reblocking they are *e.g.* $2N$ and $2N \times 2N$, respectively.

6 Accuracy

For simplicity of reproduction of results, the finite examples with two modulation states we used have odd maximum queue lengths to fit exactly into the folded structure. Solution to queues with even maximum queue lengths requires the inclusion of notional dummy states in the lattice with zero occupation probability, and is not additionally enlightening.

We use Mathematica® [15] to support our analysis work, and results in this paper were calculated using a precision of 10^{-20}. We used the built-in function `LUDecomposition` to solve the linear systems we described as matrix equations. Eigenvalues are found using the Mathematica function `Solve` applied to the polynomial determinant equation, and the eigenvectors by using `NullSpace` applied to the transpose of the characteristic matrix with the appropriate eigenvalue substituted.

For the finite waiting room example, we solved for the probability of queue states 1) directly using untransformed balance equations, 2) directly using localised balance equations, 3) directly using a folded system, 4) by spectral expansion of the localised equations, 5) by spectral expansion of the folded localised equations and 6) by matrix geometric methods applied to the folded localised equations. In every

[4] These are either unbatched arrival processes (*i.e.* batch parameter 0), or an arrival process whose batch parameter is equal to that of another arrival process, and hence can be combined into one.

case, the results for the equilibrium occupation probabilities of the joint modulation and queue length states solved using methods 1 through 5 were identical to within the numerical precision used. We therefore compare only spectral expansion and (unfolded) matrix geometric methods explicitly.

L	j	SE
49	0	$(0.14265372, 0.07371261)$
	49	$(2.91845e-5, 1.11857e-5)$
15	0	$(0.1509138, 0.7798059)$
	15	$(1.0946906e-2, 0.4447833e-2)$

L	j	MG
49	0	$(0.1426\mathbf{4900}, 0.07371\mathbf{004})$
	49	$(2.9\mathbf{5770}e-5, 1.\mathbf{2028}e-5)$
15	0	$(0.150\mathbf{5570}, 0.77\mathbf{74897})$
	15	$(1.\mathbf{113158}e-2, 0.\mathbf{54306}e-2)$

In solving the finite queues, we found that the matrix geometric methods returned results which differed from spectral expansion to a significant degree, of the order of $10^{-7} \ldots 10^{-5}$, *i.e.* not within the minimal precision. To illustrate this with meaningful values, we show the probability of the queue being empty and full for each modulation state, thus indicating the utilization and blocking probability. The values are given for spectral expansion (SE) which agree with the direct solution of the explicit Markov chain, and matrix geometric methods (MG), which do not.

6.1 The effect of raising λ to a power

The solution to the system using spectral expansion is

$$\overrightarrow{\mathbf{v}}_{j+\epsilon^b} = \sum_k \alpha_k \lambda_k^j \overrightarrow{\psi}_k$$

but we use estimates $\hat{\lambda}_k$ and $\hat{\overrightarrow{\psi}}$, which are perturbed by the numerical representation. Using the vector $\overrightarrow{\epsilon}$ to represent the maximum error in eigenvectors and ϵ_k for the error in the eigenvalues, we have an approximation $\hat{\overrightarrow{\mathbf{v}}}_j$ to the state occupation vectors:

$$\hat{\overrightarrow{\mathbf{v}}}_{j+\epsilon^b} = \sum_k \alpha_k (\lambda_k + \varepsilon_k)^j (\overrightarrow{\psi}_k + \overrightarrow{\varepsilon})$$

This gives an error varying with j arising from

$$\hat{\lambda}^j - \lambda^j = (\lambda + \varepsilon)^j - \lambda^j = j\varepsilon\lambda^{j-1} + o(\varepsilon^2) \qquad (20)$$

The term $j\varepsilon\lambda^{j-1}$ decreases with j when $j/(j+1) > \lambda$, which requires $\lambda < 1$ and j sufficiently large. If $\lambda > 1$, then the error always increases with j. The implications of this for infinite queues, in which eigenvalues on or outside the unit disc are excluded, is that errors are not amplified. In finite queues, however, in which all eigenvalues are included, errors in the larger eigenvalues will be amplified in the balance equations at the full queue. Note, however, that errors in the eigenvectors are not amplified. Thus, satisfaction of the boundary conditions at the full queue are not compromised, as these cover a small range of queue lengths differing by a small power of the eigenvalues.

6.2 The effect of raising $\underline{\mathbf{R}}$ and $\underline{\mathbf{B}}$ to powers

The matrices providing the solution in the matrix geometric method encapsulate (between them) the entire eigensystem of the solution. Either or both of these may have eigenvalues greater than 1. Since the matrix $\underline{\mathbf{R}}$ provides components of the solution to both finite and infinite queues, and the approximation and error behaviour of $\hat{\underline{\mathbf{B}}}$ exhibits the same form, we examine $\underline{\mathbf{R}}$ and its approximation in detail.

The numerical estimate $\hat{\underline{\mathbf{R}}} = \underline{\mathbf{R}} + \underline{\mathbf{D}}$, where $\underline{\mathbf{R}}$ is the "perfect" solution, and $\underline{\mathbf{D}}$ is a matrix of errors due principally to rounding in the numerical representation. Examining the role of the term $\underline{\mathbf{D}}$ in the solution, we find a term which does not simplify in the same way as in expression (20), as the matrices are not commutative: expanding the result of raising the numerical estimate $\hat{\underline{\mathbf{R}}}$ to the power j yields the following:

$$\begin{aligned}
&\hat{\underline{\mathbf{R}}}^j = (\underline{\mathbf{R}} + \underline{\mathbf{D}})^j \\
&= \underline{\mathbf{R}}^j + (\underline{\mathbf{D}}\underline{\mathbf{R}}^{j-1} + \underline{\mathbf{R}}\underline{\mathbf{D}}\underline{\mathbf{R}}^{j-2} + \cdots \underline{\mathbf{R}}^{j-1}\underline{\mathbf{D}}) + o(\underline{\mathbf{D}}^2) \\
&= \underline{\mathbf{R}}^j + (\underline{\mathbf{\Omega}}_0 + \underline{\mathbf{\Omega}}_1 + \cdots + \underline{\mathbf{\Omega}}_{j-1}) + o(\underline{\mathbf{D}}^2) \\
&= \underline{\mathbf{R}}^j + \underline{\mathbf{\Omega}} + o(\underline{\mathbf{D}}^2)
\end{aligned}$$

where $\underline{\mathbf{\Omega}}_i = \underline{\mathbf{R}}^i \underline{\mathbf{D}} \underline{\mathbf{R}}^{j-1-i}$ and $\underline{\mathbf{\Omega}} = \sum_i \underline{\mathbf{\Omega}}_i$. To obtain an error estimate, we consider the effect that the perturbation $\underline{\mathbf{D}}$ has on the eigenvectors of $\hat{\underline{\mathbf{R}}}$. Let matrix $\underline{\mathbf{R}}$ have eigenvectors $\overrightarrow{\psi}_i$ associated with eigenvalues λ_i labelled such that $\lambda_i \leq \lambda_{i+1}$.

As $\underline{\mathbf{D}}$ is a perturbation matrix with small entries distributed about zero, we have that (in general) for any $1 \leq i \leq N$,

$$\overrightarrow{\psi}_i \underline{\mathbf{D}} = \sum_{k=1}^{N} \beta_{i,k} \overrightarrow{\psi}_k$$

where $\beta_{i,k}$ are also small.

Consider the h^{th} error term, $\underline{\mathbf{\Omega}}_h = \underline{\mathbf{R}}^h \underline{\mathbf{D}} \underline{\mathbf{R}}^{j-1-h}$, when eigenvector $\overrightarrow{\psi}_i$ is multiplied by it

$$\begin{aligned}
\overrightarrow{\psi}_i \underline{\mathbf{\Omega}}_h &= \overrightarrow{\psi}_i \underline{\mathbf{R}}^h \underline{\mathbf{D}} \underline{\mathbf{R}}^{j-1-h} \\
&= \lambda_i^h \overrightarrow{\psi}_i \underline{\mathbf{D}} \underline{\mathbf{R}}^{j-1-h} \\
&= \lambda_i^h \left(\textstyle\sum_{k=1}^{N} \beta_{i,k} \overrightarrow{\psi}_k \right) \underline{\mathbf{R}}^{j-1-h} \\
&= \lambda_i^h \textstyle\sum_{k=1}^{N} \beta_{i,k} \lambda_k^{j-1-h} \overrightarrow{\psi}_k
\end{aligned}$$

to evaluate the total error in evaluating $\overrightarrow{\psi}_i\hat{\underline{\mathbf{R}}}^j$, we sum up all the individual errors

$$\begin{aligned}
\overrightarrow{\psi}_i\underline{\mathbf{\Omega}} &= \overrightarrow{\psi}_i(\underline{\mathbf{D}}\underline{\mathbf{R}}^{j-1} + \underline{\mathbf{R}}\underline{\mathbf{D}}\underline{\mathbf{R}}^{j-2} + \cdots + \underline{\mathbf{R}}^{j-1}\underline{\mathbf{D}}) \\
&= \textstyle\sum_{h=0}^{j-1} \overrightarrow{\psi}_i\underline{\mathbf{\Omega}}_h \\
&= \textstyle\sum_{h=0}^{j-1} \left(\lambda_i^h \sum_{k=1}^{N} \beta_{i,k}\lambda_k^{j-1-h}\overrightarrow{\psi}_k\right) \\
&= \textstyle\sum_{k=1}^{N} \left(\beta_{i,k}\overrightarrow{\psi}_k \sum_{h=0}^{j-1} \lambda_i^h\lambda_k^{j-1-h}\right) \\
&= \textstyle\sum_{k=1}^{N} \beta_{i,k}\overrightarrow{\psi}_k \frac{\lambda_i^j-\lambda_k^j}{\lambda_i-\lambda_k}
\end{aligned}$$

The behaviour for large j is prescribed by the largest eigenvalue, λ_N, as

$$\overrightarrow{\psi}_i\underline{\mathbf{\Omega}} \doteq \begin{cases} \beta_{i,N}\overrightarrow{\psi}_N\lambda_N^{j-1} & \text{if } i \neq N \\ \sum_{k=1}^{N-1}\beta_{i,k}\overrightarrow{\psi}_k\lambda_N^{j-1} & \text{if } i = N \end{cases}$$

We can now quantify the asymptotic effect of numerical inaccuracy in $\hat{\underline{\mathbf{R}}}$ on the eigenvectors of $\underline{\mathbf{R}}$.

$$\begin{aligned}
\overrightarrow{\psi}_i\hat{\underline{\mathbf{R}}}^j &= \overrightarrow{\psi}_i(\underline{\mathbf{R}}^j + \underline{\mathbf{\Omega}} + o(\underline{\mathbf{D}}^2)) \\
&\doteq \lambda_i^j\overrightarrow{\psi}_i + \begin{cases} \beta_{i,N}\overrightarrow{\psi}_N\lambda_N^{j-1} & \text{if } i \neq N \\ \sum_{k=1}^{N-1}\beta_{i,k}\overrightarrow{\psi}_k\lambda_N^{j-1} & \text{if } i = N \end{cases} \\
&\doteq \begin{cases} \beta_{i,N}\overrightarrow{\psi}_N\lambda_N^{j-1} & \text{if } i \neq N \\ \lambda_N^j\overrightarrow{\psi}_N & \text{if } i = N \end{cases}
\end{aligned} \tag{21}$$

which means that every eigenvector $\overrightarrow{\psi}_i$ asymptotically approaches the eigenvector $\overrightarrow{\psi}_N$ when multiplied by large powers of $\hat{\underline{\mathbf{R}}}$.

This behaviour is mirrored in the use of $\hat{\underline{\mathbf{B}}}$ in providing the additional eigensystem components required for use in finite queues.

This tendency for any vector, when multiplied by high powers of the matrix geometric matrix $\hat{\underline{\mathbf{R}}}$, to approach the eigenvector associated with the largest eigenvalue is not surprising, as this phenomenon gives us the well-known power method for finding the largest eigenvalue of a matrix.

The above error estimate shows that the numerical solution for large j becomes dominated by the eigenvector with largest eigenvalue. Within the solution process for infinite systems all eigenvalues – including λ_N – are contained strictly within the unit disc due to the requirement of normalization. Consequently, for all i, the error estimate $\overrightarrow{\psi}_i\underline{\mathbf{\Omega}}$ dominates the probability solution vector for large j, but also vanishes, allowing an accurate solution to be achieved using the matrix geometric method.

In large, finite systems, the largest eigenvalue λ_N is outside the unit disc and the error term $\overrightarrow{\psi}_i\underline{\mathbf{\Omega}}$ ceases to vanish so that, with increasing queue length, numerical errors are amplified. The particular value of $\mathbf{f}\underline{\mathbf{R}}^{j-\epsilon^b}$ (from $\overrightarrow{\mathbf{v}}_j = \overrightarrow{\mathbf{v}}_{\epsilon^b}\underline{\mathbf{R}}^{j-\epsilon^b} + \overrightarrow{\mathbf{v}}_{\epsilon^t}\underline{\mathbf{B}}^{\epsilon^t-j}$) may not be particularly inaccurate, because in the perfect solution, the largest eigenvalue will dominate. However, the contributions from the other components to the solution, in particular those arising from the smaller eigenvalues, quickly lose precision. As a consequence the boundary condition equations that specify $\overrightarrow{\mathbf{f}}$ and $\overrightarrow{\mathbf{b}}$, linking the top and bottom of the queue, can only be evaluated in an inaccurate manner. Subsequent solution yields similarly inaccurate values for $\overrightarrow{\mathbf{f}}$ and $\overrightarrow{\mathbf{b}}$, which cause inaccuracies throughout the queue.

6.3 Numerical example of $\mathbf{f}.\hat{R}^j$

We have presented an asymptotic error estimate for calculations performed in solving a finite queue using the matrix geometric method. The limit $j \to \infty$ is never achieved for any given problem so we show the behaviour of the eigenvectors in our example to illustrate how quickly errors can accumulate. To this end we examine the angle $\alpha(\overrightarrow{\psi}_i\hat{\underline{\mathbf{R}}}^j, \overrightarrow{\psi}_4) = \cos^{-1}\frac{\overrightarrow{\psi}_i\underline{\mathbf{R}}^j.\overrightarrow{\psi}_4}{\left\|\overrightarrow{\psi}_i\underline{\mathbf{R}}^j\right\|_2\left\|\overrightarrow{\psi}_4\right\|_2}$ ($\overrightarrow{\psi}_4$ being the dominant eigenvector and $\left\|\overrightarrow{\mathbf{x}}\right\|_2$ the Euclidean norm of a vector $\overrightarrow{\mathbf{x}}$) for various j when working with double precision conforming with IEEE floating point standards. In

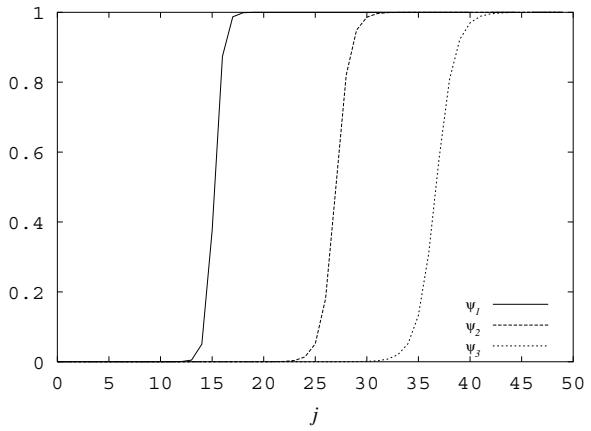

Figure 1. Cosine of angle to dominant eigenvector with power j of matrix F.

figure 6.3, we show the convergence of each of the non-dominant eigenvectors toward the *direction* [5] of the dominant eigenvector applied to successive powers of $\hat{\underline{\mathbf{R}}}$. The vertical axis between 0 and 1 indicates the progression in terms of the fraction of the angle (given by its cosine) between the true vector and the dominant vector traversed.

We might wish to reduce this effect by considering alternative schemes for constructing the value of $\mathbf{f}.\hat{\underline{\mathbf{R}}}^j$. Figure 6.3 shows the result of successive post-multiplication by $\underline{\mathbf{R}}$ of an intermediate result, seeded at $\mathbf{f}$. If instead, we attempt to calculate $\underline{\mathbf{R}}^j$ more efficiently by use of intermediate, previously calculated squares of the matrix, *e.g.* we

[5] The relative magnitude of the vectors was not considered

use $\underline{\hat{\mathbf{R}}}^4 = \underline{\hat{\mathbf{R}}}^2\underline{\hat{\mathbf{R}}}^2$ instead of $\underline{\hat{\mathbf{R}}}^4 = \underline{\hat{\mathbf{R}}}\underline{\hat{\mathbf{R}}}\underline{\hat{\mathbf{R}}}\underline{\hat{\mathbf{R}}}$, the trajectory of the estimate of $\mathbf{f}.\underline{\hat{\mathbf{R}}}^j$ becomes erratic after the initial departure from it's correct direction, but gives overall smaller errors. No matter how accurately the matrix $\underline{\hat{\mathbf{R}}}^j$ is calculated, the lower error bound (21) holds.

7 Conclusion

Beginning with an analogy of the simple scalar solution to Fibonnacci's recurrence, MG has become ever more sophisticated and efficient while the more recent SE has focused on a more general approach. In this paper we have extended SE to a provably complete solution method using generalized eigenvectors. This provided a framework in which to compare traditional MG and SE. We have identified models in which MG is inadequate in its basic form, but also indicated how it may be possible to expand the class of solutions. Standard SE also has its limitations, but these are less stringent. We also considered numerical issues arising from high powers of matrices to help explain stability issues in using MG solutions.

This paper therefore proves the existence of a class of queues for which the matrix geometric representation and the matrix analytic approach fails. It seems unlikely that this example will be unique, so further exploration of the eigenspectrum of queue solutions is motivated.

Until the class of queueing systems which does not comply with a fixed-dimension matrix generated solution is fully characterized, we suggest that the use of mature and efficient matrix geometric or analytic tools for solving a given queueing system should be preceded by an analysis of the eigensystem of the solution through spectral expansion, whether algebraic or numerical, to verify that the solutions produced by such tools are correct.

Acknowledgement: The authors would like to thank Peter Harrison for his guidance and participation in maintaining mathematical rigour.

References

[1] D. Bini and B. Meini. Improved cyclic reduction for solving queueing problems. *Numerical Algorithms*, 15:57–74, 1997.

[2] D. Bini and B. Meini. Solving block banded block toeplitz systems with structured blocks: algorithms and applications. In *Structured Matrices: Recent Developments in Theory and Computation*, New York, 2001. Nova Science Publisher Inc.

[3] A. Riska and E. Smirni. M/G/1-type Markov Processes: A tutorial In *Performance Evaluation of Complex Systems:Techniques and Tools*, Performance 2002 Tutorial Lectures, M. Calzarossa and S. Tucci (eds), Springer-Verlag, Lecture Notes in Computer Science (2459), Rome, Italy, September 2002

[4] H. Zatschler *Steady-State and Response Time Analysis of Modulated Queues and Networks with Batches* PhD thesis, University of London, August 2004, http://aesop.doc.ic.ac.uk/pubs/mmcpp-agen/

[5] B. Haverkort and A. Ost. Steady-state analysis of infinite stochastic petri nets: comparing the spectral expansion and the matrix-geometric method. In *Proceedings of 7th International Workshop on Petri Nets and Performance Models*, 1997.

[6] A. Stathopoulos, A, Riska,Z. Hua, and E. Smirni Bridging ETAQA and Ramaswami's formula for the solution of M/G/1-type processes In Performance Evaluation Volume 62, Issues 1-4 , October 2005, Pages 331-348

[7] G. Latouche and V. Ramaswami. A logarithmic reduction algorithm for quasi-birth-death processes. *Journal of Applied Probability*, 30:650–674, 1993.

[8] I. Mitrani and R. Chakka. Spectral expansion solution for a class of markov models: Application and comparison with the matrix-geometric method. *Performance Evaluation*, 23:241–260, 1995.

[9] V. Naoumov, U. Krieger, and D. Wagner. Analysis of a multi-server delay-loss system with a general markovian arrival process. In *Proceedings of First International Conference on Matrix-Analytic Methods*, July 1997.

[10] M. Neuts. *Matrix Geometric Solutions in Stochastic Models*. Johns Hopkins University Press, Baltimore, 1981.

[11] M. Pease. *Mathematics in Science and Engineering, vol 16*. Academic Press Inc., 1965.

[12] B. Smith and al. *Matrix Eigensystem Routines – EISPACK Guide*. Springer-Verlag, New York, 1976.

[13] D. Thornley, H. Zatschler, and P. Harrison. An automated formulation of queues with multiple geometric batch processes. In *Proceedings of HETNETS'03*, July 2003.

[14] J. Wilkinson and C. C. Reinsch. *vol. II of Handbook for Automatic Computation*. Springer-Verlag, New York, 1971.

[15] S. Wolfram. *The Mathematica Book 4th ed.* Cambridge University Press, 1999.