

COMPUTING THE INVARIANT LAW OF A FLUID MODEL

by

L.C.G. Rogers & Z. Shi

Queen Mary & Westfield College, University of London

Abstract: In this paper, we discuss a variety of methods for computing the Wiener-Hopf factorization of a finite Markov chain associated to a fluctuating additive functional. The importance of this is that the equilibrium law of a fluid model can be expressed in terms of these Wiener-Hopf factors. The diagonalization methods considered are actually quite efficient, and provide an effective solution to the problem.

Keywords and phrases: Markov chain, Wiener-Hopf factorization, fluid model, invariant law, noisy Wiener-Hopf problem.

AMS 1991 subject classifications: 60J27, 60K15, 60K20, 60K25.

1. Introduction

Fluid approximations to queueing systems have been widely and successfully studied in recent years; as a representative (but by no means exhaustive) sample, see Anick, Mitra & Sondhi [1], Mitra [7], Stern & Elwalid [9], Gaver & Lehoczy [4], [5], Asmussen [2]. The key idea in any fluid approximation is that the service requirements of the customers are typically small, and there are typically many customers in a system at any time, so the total number of customers may be thought of as a continuous variable, which changes in a nearly deterministic way.

The simplest prototype example is clearly explained in Mitra [7]. One has n pipes, each of which can pour fluid into a large tank of capacity a , and in the bottom of the tank there are k taps, each of which can be either open or closed. Each pipe which is pouring fluid into the tank does so at rate θ , and each open tap empties fluid out at rate ρ , so that if ξ_t is the amount of fluid in the tank at time t , Z_t is the number of open taps and Y_t the number of pipes pouring water into the tank, then

$$(1) \quad \frac{d\xi_t}{dt} = \theta Y_t - \rho Z_t$$

at least when $0 < \xi_t < a$. If the tank is full, surplus input overflows and is lost.

One typically assumes that the pair (Y_t, Z_t) is a Markov chain, and one is concerned with the equilibrium distribution of ξ . The great virtue of such models is that the size of the statespace of the chain is nk , whereas if one had sought to model the queueing system discretely, the statespace would have been impossibly large. However, computing the equilibrium distribution is now a problem outside the conventional theory of Markov chains, and is intrinsically quite difficult. It is closely linked to the problem of computing the *Wiener-Hopf factorization of a Markov chain*, a problem first addressed by Barlow, Rogers & Williams [3]. The situation considered there is of a Markov chain X on a finite statespace E , together with an additive functional

$$(2) \quad \phi_t \equiv \int_0^t v(X_s) ds,$$

where $v : E \rightarrow R \setminus \{0\}$ is some function, which takes positive *and* negative values, so that ϕ is a *fluctuating* additive functional. Now one may define time changes

$$(3) \quad \tau_t^\pm \equiv \inf \{u : \pm \phi_u > t\},$$

and the time-changed processes

$$(4) \quad Y_t^\pm \equiv X(\tau_t^\pm).$$

The process Y^+ lives in $E_+ \equiv \{i : v(i) > 0\}$, because at any time τ_t^+ , ϕ must be increasing and therefore X must be in E_+ .

Clearly Y^+ , Y^- are Markov chains; *what are their generators?* It is extraordinarily difficult to answer this question in closed form except in the most elementary examples, but Barlow, Rogers & Williams [3] provide a characterization of the generators G_\pm of Y^\pm , in terms of the matrices

$$(5.i) \quad \Pi_+(i, j) \equiv P^i(X(\tau_0^+) = j, \tau_0^+ < \infty), \quad (i \in E_-, j \in E_+),$$

$$(5.ii) \quad \Pi_-(j, i) \equiv P^j(X(\tau_0^-) = i, \tau_0^- < \infty), \quad (j \in E_+, i \in E_-).$$

To state their results, if S is any finite set, let $\bar{Q}(S)$ be the set of $S \times S$ irreducible Q -matrices, that is, matrices whose off-diagonal entries are non-negative, whose row sums are non-positive, and which generate an irreducible chain. We say that $Q \in \bar{Q}(S)$ is *recurrent* if the row sums of Q are zero, otherwise we say $Q \in \bar{Q}(S)$ is *transient*. Let V be the diagonal matrix $\text{diag}(v(i))$.

THEOREM A (Barlow-Rogers-Williams [3]). *A Wiener-Hopf factorization of $V^{-1}Q$ is a quadruple $(Z_+, Q_+; Z_-, Q_-)$, where Z_+ is $E_- \times E_+$, Z_- is $E_+ \times E_-$, $Q_\pm \in \bar{Q}(E_\pm)$ and*

$$(6) \quad V^{-1}Q \begin{pmatrix} I & Z_- \\ Z_+ & I \end{pmatrix} = \begin{pmatrix} I & Z_- \\ Z_+ & I \end{pmatrix} \begin{pmatrix} Q_+ & 0 \\ 0 & -Q_- \end{pmatrix}.$$

- (i) The quadruple $(\Pi_+, G_+; \Pi_-, G_-)$ is always a Wiener-Hopf factorization of $V^{-1}Q$.
- (ii) If $Q \in \bar{Q}(E)$ is transient, then the Wiener-Hopf factorization of $V^{-1}Q$ is unique.

See Rogers [8] for a proof of Theorem A and a discussion of the result. For our current purposes, the main point about Theorem A is that *the invariant distribution of the process ξ can be simply expressed in terms of Π_\pm and G_\pm* ; see §§4-5 of Rogers [8]. As an example, we have that if the buffer has infinite capacity, and if $j \in E_+$, then

$$\lim_{t \rightarrow \infty} P(X_t = j, \xi_t > x) = m_j(\exp(x\hat{G}_-)1)_j,$$

where m is the invariant distribution of the underlying chain X , and \hat{G}_- is the generator of the minus process defined from the reversal of X . The whole story is in [8].

Thus to find the invariant distribution of ξ , we have to compute Π_\pm and G_\pm . There are a variety of ways to do this, which we shall describe, compare, and evaluate in §2. Broadly speaking, there are two *types* of methods: recursive methods, and diagonalization methods. It turns out that the diagonalization methods are numerically far superior, which is comforting, since all the applications so far (for example, [1], [5], [7], [9]) are based on diagonalization procedures.

Next we consider a variant of the fluid model, in which the buffer content process ξ is not simply an additive functional of a Markov chain, but has some (small) Brownian noise superimposed on it. To be more precise, if

$$(7) \quad \phi_t = \int_0^t v(X_s)ds + \epsilon B_t,$$

where B is Brownian motion independent of the chain X , then

$$(8) \quad \xi_t = \phi_t + L_t^0 - L_t^a,$$

where L^0 (respectively, L^a) is an increasing process which grows only when ξ is at 0 (respectively, a), and are such as to constrain ξ to $[0, a]$. As before, we may define time-changes τ^\pm by (3), and new processes Y^\pm by (4). This time, Y^\pm can live anywhere in E , and have generators Γ_\pm . It turns out once again that the invariant law of ξ can be simply expressed in terms of Γ_\pm (see §7 of Rogers [8]), so the problem is to compute Γ_\pm . As Kennedy & Williams [6], Asmussen [2] and Rogers [8] all show, Γ_\pm are characterised as follows.

THEOREM B (Kennedy-Williams; Asmussen; Rogers). *The generator Γ_+ (respectively Γ_-) is the unique $Z \in \bar{Q}(E)$ such that*

$$(9.i) \quad \frac{1}{2}\epsilon^2 Z^2 - VZ + Q = 0$$

(respectively,

$$(9.ii) \quad \frac{1}{2}\epsilon^2 Z^2 + VZ + Q = 0. \quad)$$

In §3 of this paper, we discuss a number of numerical algorithms for computing Γ_{\pm} . Once again, there are recursive methods, and a diagonalization method; and once again, the diagonalization method is far superior.

2. The classical Wiener-Hopf problem.

Throughout this section, we shall make the simplifying assumption that

$$v : E \rightarrow \{-1, +1\}.$$

For the computation of Π_+ and Π_- , this loses no generality since we could replace Q by $\Delta^{-1}Q$, where Δ is the diagonal matrix $\text{diag}(|v(i)|)$, and this converts (6) to the standard form where

$$V = \begin{pmatrix} I^{(E_+ \times E_+)} & 0^{(E_+ \times E_-)} \\ 0^{(E_- \times E_+)} & -I^{(E_- \times E_-)} \end{pmatrix}.$$

We shall now similarly partition Q according to the partition $E = E_+ \cup E_-$, and write it as

$$Q = \begin{pmatrix} A^{(E_+ \times E_+)} & B^{(E_+ \times E_-)} \\ C^{(E_- \times E_+)} & D^{(E_- \times E_-)} \end{pmatrix}.$$

So if we look only at the E_+ -columns of (6), we have

$$(10) \quad \begin{pmatrix} A & B \\ -C & -D \end{pmatrix} \begin{pmatrix} I \\ Z_+ \end{pmatrix} = \begin{pmatrix} I \\ Z_+ \end{pmatrix} Q_+,$$

which can be equivalently expressed as

$$(11.i) \quad A + BZ_+ = Q_+,$$

$$(11.ii) \quad -C - DZ_+ = Z_+Q_+.$$

One can easily eliminate Q_+ to obtain the key quadratic equation

$$(12) \quad DZ_+ + Z_+A + C + Z_+BZ_+ = 0.$$

Until further notice, we concentrate on calculating Π_+ , so abbreviate Z_+ to Z .

This innocent-looking equation is the starting point of the recursive methods for calculating Π_+ . Theorem A guarantees that one solution to (12) is given by $Z = \Pi_+$ - but are there any others? Williams [10] proves that Π_+ is in fact the *minimal* solution to (12), but the complete resolution of the uniqueness story has only recently been given (Rogers [8]). In brief, the only time that (10) has a solution (Z_+, Q_+) other than (Π_+, G_+) is if Q_+ is recurrent, G_+ is transient, and in that case Q_+ can be represented in terms of G_+ as

$$Q_+ = G_+ - (G_+1)\mu,$$

where μ is the left eigenvector of G_+ of largest real part. Unfortunately, it is not so easy to express G_+ in terms of Q_+ .

We shall not dwell on possible non-uniqueness, because the methods we examine here all give (in principle) Π_+ . First we shall describe three recursive methods.

(i) *Williams' method* [10]. If we replace C by θC ($0 < \theta < 1$), then we get

$$DZ_+(\theta) + Z_+(\theta)A + \theta C + Z_+(\theta)BZ_+(\theta) = 0,$$

with

$$Z_+(\theta) = \sum_{n \geq 1} \theta^n Y_n.$$

Comparing coefficients of θ^n gives

$$DY_1 + Y_1A = -C,$$

$$DY_{n+1} + Y_{n+1}A = -\sum_{i=1}^n Y_i B Y_{n+1-i}, \quad (n \geq 1).$$

By solving the equation above we will get the solution to (12) with $Z_+ = \sum_{n \geq 1} Y_n$. See Williams [10] for the complete story. This method is rather slow indeed, as apart from the difficulty of solving the type of equation

$$(14) \quad DZ + ZA = -R,$$

(for given R), it is slow to compute $\sum_{i=1}^n Y_i B Y_{n+1-i}$. Moreover, the previous values of Y_i must all be stored to permit the calculation. The following recursive method is an improvement on Williams' method.

(ii) *An improved recursive method*. This second method generates a sequence Z_n of non-negative matrices by the recursion

$$(15) \quad Z_0 = 0, \quad DZ_{n+1} + Z_{n+1}A = -C - Z_n B Z_n.$$

The following representation

$$(16) \quad Z_{n+1} = \int_0^\infty e^{tD} (C + Z_n B Z_n) e^{tA} dt$$

makes it clear that the Z_n increase to some solution K of (12), and likewise it is easy to see that if Z were any non-negative solution to (12), then for all n , $Z_n(i, j) \leq Z(i, j)$ and hence $K(i, j) \leq Z(i, j)$, for all i, j . Since Π_+ is the minimal solution to (12), $K = \Pi_+$.

How does one go about solving the equation (15)? The main difficulty is that solving the equation (14) with R known, is not a trivial matter. One approach might be to observe that the solution to (16) can be expressed as

$$Z = \int_0^\infty e^{tD} R e^{tA} dt,$$

but this does not lend itself to rapid computation. Another observation is that if D and A were *diagonal* matrices, then the computation of Z is trivial. We made a program which firstly reduced A, D to diagonal form, and then used the recursion (14) to compute the Z_n . The main conclusions about the method were

- It is stable, and converges monotonically to the true value;
- It is quite slow;
- It requires a pair of diagonalizations, so can be no real improvement over the diagonalization method discussed below.

The last point deserves amplification. Diagonalizing an $N \times N$ matrix is an $O(N^3)$ operation. If we now have to diagonalize two square matrices of order about $N/2$, the best we can hope for is to reduce the computation time by a factor of about 4.

(iii) *Kurtz' method*. This method was shown to us by Tom Kurtz. Firstly write $D = D_\Delta + D_+$, $A = A_\Delta + A_+$, where D_Δ (resp. A_Δ) is the diagonal part of D (resp. A). Now consider the recursion

$$(17) \quad Z_0 = 0, \quad D_\Delta Z_{n+1} + Z_{n+1} A_\Delta = -C - Z_n B Z_n - D_+ Z_n - Z_n A_+.$$

It is trivial to verify inductively that the Z_n are now all non-negative ($Z_n(i, j) \geq 0, \forall i, j, n$), and are increasing ($Z_{n+1}(i, j) \geq Z_n(i, j), \forall i, j, n$). Moreover, if Z is any non-negative solution to (12), then equally easily $Z \geq Z_n$ for each n . Hence the monotone limit of the Z_n will be the smallest non-negative solution to (12), namely Π_+ .

The Kurtz method has one great advantage over the previous recursive method (ii); solving (17) for Z_{n+1} is now trivial, because D_Δ and A_Δ are diagonal. One small disadvantage is that there is now no probabilistic interpretation of Z_n , so estimating how far Z_n is from Π_+ may not be so easy.

We carried out the Kurtz method on a number of examples where the off-diagonal entries of Q were independent negative exponential random variables, so that essentially Q had no structure. The main points about the method were:

- It is stable, and converges monotonically to the solution;
- It is slow, but not as slow as the previous recursive method (ii).

It is worth pointing out that in any recursive method which involves many matrix multiplications, one can not expect a fast algorithm, since multiplying an $n \times k$ matrix by a $k \times m$ matrix has $O(nmk)$ operations so that each *step* of the recursion is $O(n_+ n_- (n_+ + n_-))$, where n_\pm is the size of E_\pm . Thus we could not expect the overall order of a recursive method to be better than the diagonalization method described below (which is $O((n_+ + n_-)^3)$), although it might be better if (say) $n_+ = O(1)$.

We also tried to accelerate the convergence of these recursive methods, by looking at Z_{n-1}, Z_n, Z_{n+1} , and trying to estimate from these what the limit would be, then jumping to this estimated matrix and continuing the algorithm from there. Although this certainly got more quickly to the correct answer, it did not *always* get to the correct answer! We discarded this and the other recursive methods in favour of the diagonalization method.

(iv) *The diagonalization method.* Let m denote the invariant distribution of Q . We shall assume for the moment that $m(E_+) \neq m(E_-)$, a condition that will hold in all but a few freak situations. Just to fix our ideas, suppose that $m(E_+) > m(E_-)$. In this case, $\phi_t \rightarrow +\infty$ a.s., and the matrix Π_- is strictly substochastic, whereas the matrix Π_+ has row-sums equal to 1. This means that the matrix

$$S = \begin{pmatrix} I & \Pi_- \\ \Pi_+ & I \end{pmatrix}$$

is invertible, and so from the Wiener-Hopf factorization (6), the matrix $V^{-1}Q$ is similar to the block diagonal matrix

$$\begin{pmatrix} G_+ & 0 \\ 0 & -G_- \end{pmatrix}$$

If now we take the eigenvalues $\lambda_1, \dots, \lambda_N$ of $V^{-1}Q$ (assumed for the moment distinct) with corresponding eigenvectors f_1, \dots, f_N , and supposing that the first $n_+ = |E_+|$ of the eigenvalues are of non-positive real part, we have, for each j ,

$$M_j(t) \equiv f_j(X_t) e^{-\lambda_j \phi_t} \quad \text{is a martingale,}$$

and

$$\text{for } j = 1, \dots, n_+, \quad M_j \text{ is bounded on } [0, \tau_0^+].$$

Since $\tau_0^+ < \infty$ a.s. ($\phi_t \rightarrow \infty$ a.s.) we may use the optional sampling theorem to conclude

$$(18) \quad f_j(i) = E^i[f_j(Y_0^+)], \quad (j = 1, \dots, n_+, i \in E_-).$$

(This argument is familiar from many earlier works in this area - see, for example, Barlow, Rogers & Williams [3]).

Thus we have for each $j = 1, \dots, n_+$

$$f_j(i) = (\Pi_+ f_j^+)(i), \quad (i \in E_-),$$

where f_j^+ is the restriction to E_+ of f_j . So if

$$F = (f_1, \dots, f_{n_+}), F_+ = (f_1^+, \dots, f_{n_+}^+),$$

we have in matrix notation

$$(19) \quad \begin{pmatrix} I \\ \Pi_+ \end{pmatrix} F_+ = F.$$

If we know the diagonalization of $V^{-1}Q$, we know F ; and F_+ is invertible - so from (19) we can calculate Π_+ by inverting F_+ ! The diagonalization method therefore firstly diagonalizes $V^{-1}Q$, then computes F from this, and finally computes Π_+ .

Remarks. The assumptions made in the above discussion can all be lifted quite easily. If the eigenvalues are *not* assumed distinct, we take a basis f_1, \dots, f_N of Jordan vectors ($(V^{-1}Q - \lambda)^r f = 0$ for some λ , for some $r \geq 1$), and the identity (18) still holds good for all those Jordan vectors whose eigenvalues have non-positive real part - see Barlow, Rogers & Williams [3] and Rogers [8] for further discussion. In practice, a matrix whose Jordan form is *not* diagonal is a freak. The assumption that $m(E_+) > m(E_-)$ is also superfluous; if $m(E_+) < m(E_-)$ then there are N_+ eigenvalues whose real part is strictly negative and exactly as we argued for (18), one now finds

$$f_j(i) = E^i[f_j(Y_0^+) : \tau_0^+ < \infty], \quad (j = 1, \dots, n_+, i \in E_-).$$

The balanced case $m(E_+) = m(E_-)$ needs a bit more care, but is not too difficult. Rogers [8] proves that the Jordan decomposition of $V^{-1}Q$ now contains a Jordan block $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$. There will now be exactly $n_+ - 1$ eigenvalues of $V^{-1}Q$ whose real part is strictly negative, f_1, \dots, f_{n_+-1} , say, and by setting $f_{n_+} \equiv 1$, the statement (18) remains valid, and from that we deduce Π_+ as before.

The numerical computations were done using the NAG library routine F01AAF to diagonalize $V^{-1}Q$. This routine was fast and accurate. In the following table we compare the time taken (on a SUN SPARC station 2) to calculate Π_+ by the Kurtz method and by the diagonalization method. The matrices Q were generated randomly by the method described above. The columns are labelled by (n_+, n_-) , where $n_{\pm} = |E_{\pm}|$ as before.

Table 1

	(10,15)	(48,50)	(100,100)
time for diagonalization method	0.49s	9.82s	60.88s
time for Kurtz method	4.76s	1942.31s	50533.0*s
number of iterations for Kurtz method	64	460	1378

*The method here had not yet converged, though the mean value of the difference $Z_{n+1}(i, j) - Z_n(i, j)$ for $i, j = 1, \dots, 100$ was less than 10^{-8} .

Observe that the diagonalization method is overwhelmingly superior!

3. Noisy Wiener-Hopf factorization. We now consider the noisy Wiener-Hopf problem; we have to solve the equation (9.i-ii) for the unknown Q -matrices. Firstly we consider two recursive methods.

(i) *Kennedy & Williams' method* [6]. If we write $Q = \Delta + R$, where Δ is the diagonal part of Q , we introduce a killing with probability $1 - \theta$ at each jump of the chain by replacing Q by $Q(\theta) \equiv \Delta + \theta R$. If now the solution to (9.i) with $Q(\theta)$ in place of Q is represented as

$$Z(\theta) = \sum_{n \geq 0} \theta^n \Gamma_n,$$

then by comparing coefficients of θ^n for each n in (9.i) we arrive at

$$(20.i) \quad \frac{1}{2}\epsilon^2\Gamma_0^2 - V\Gamma_0 + \Delta = 0,$$

$$(20.ii) \quad \frac{1}{2}\epsilon^2(\Gamma_0\Gamma_1 + \Gamma_1\Gamma_0) - V\Gamma_1 + R = 0,$$

$$(20.iii) \quad \frac{1}{2}\epsilon^2 \sum_{r=0}^n \Gamma_r\Gamma_{n-r} - V\Gamma_n = 0, \quad (n \geq 2).$$

Now (20.i) can be solved explicitly, since Γ_0 is a diagonal matrix of non-positive entries, and then trivially from (20.ii) we can calculate Γ_1 (since Γ_0 is diagonal); and finally, recasting (20.iii) in the form

$$(V - \frac{1}{2}\epsilon^2\Gamma_0)\Gamma_n - \frac{1}{2}\epsilon^2\Gamma_n\Gamma_0 = \frac{1}{2}\epsilon^2 \sum_{r=1}^{n-1} \Gamma_r\Gamma_{n-r},$$

it is a simple matter to calculate Γ_n from Γ_r , $r < n$. The solution we seek is $Z(1) = \sum \Gamma_n$.

However, this method, as in the case of Williams' method for the classical Wiener-Hopf problem, is computationally most undesirable. For one thing, the calculation of the sum-of-products $\sum_{r=1}^{n-1} \Gamma_r\Gamma_{n-r}$ is slow, and can not be done recursively; it has to be recomputed at each stage. Moreover, the previous values of Γ_r must all be stored to permit this calculation! We reject this method in favour of a somewhat better recursive method.

(ii) *An improved recursive method.* Instead of trying to calculate Γ_+ to satisfying (9.i), we attempt to calculate

$$K = \Gamma_+ + \alpha I,$$

where α is going to be chosen big enough. In terms of K , the equation to be satisfied is easily derived from (9.i):

$$(21) \quad \frac{1}{2}\epsilon^2 K^2 - (V + \epsilon^2\alpha)K + (\alpha V + \frac{1}{2}\epsilon^2\alpha^2 I + Q) = 0.$$

Now if we take $\alpha > 0$ sufficiently large, all diagonal entries of the diagonal matrix $D \equiv V + \epsilon^2\alpha$ will be positive, and all entries of $C \equiv \alpha V + \frac{1}{2}\epsilon^2\alpha^2 I + Q$ will be non-negative. Thus the recursive scheme

$$(22) \quad K_0 = 0, \quad DK_{n+1} = \frac{1}{2}\epsilon^2 K_n^2 + C$$

will generate an increasing sequence K_n , which, as before, is easily seen to be bounded above by $\Gamma_+ + \alpha I$, and which converges, to the unique solution $\Gamma_+ + \alpha I$.

The main conclusions about this method are:

- It is quite simple, and stable, converging monotonically to the true value;
- It is quite slow, because of the many matrix operations.

Notice that we have considerable freedom in our choice of the parameter α . We tried some examples using this method by choosing different α . It turned out that the larger the value of α is, the longer it takes to compute.

Once again, though, it is the diagonalization method which wins.

(iii) *The diagonalization method.* If we consider the matrix

$$(23) \quad M \equiv \begin{pmatrix} 2\epsilon^{-2}V & I \\ -2\epsilon^{-2}Q & 0 \end{pmatrix},$$

then if we take an eigenvector of M with eigenvalue λ of non-positive real part, the eigenvalue equation

$$(24) \quad M \begin{pmatrix} f \\ g \end{pmatrix} \equiv \begin{pmatrix} 2\epsilon^{-2}V & I \\ -2\epsilon^{-2}Q & 0 \end{pmatrix} \begin{pmatrix} f \\ g \end{pmatrix} = \lambda \begin{pmatrix} f \\ g \end{pmatrix}$$

implies that

$$\frac{1}{2}\epsilon^2\lambda^2 f - \lambda Vf + Qf = 0,$$

and hence

$$f(X_t)e^{-\lambda\phi_t} \text{ is a martingale.}$$

If $\text{Re}(\lambda) < 0$, we use the optimal sampling theorem as before to see that

$$Ef(Y_t^+) = e^{\lambda t} f,$$

and hence that

$$(25) \quad \Gamma_+ f = \lambda f.$$

Assuming for the moment that $m(E_+) > m(E_-)$ and that all eigenvalues of M are distinct, then Γ_- is a transient Q-matrix, so has N eigenvalues in the open left half plane. An exactly similar argument to that above shows that if $-\lambda$ is an eigenvalue of M with strictly positive real part, then λ is an eigenvalue of Γ_- , and conversely. Thus in this case, M has exactly N eigenvalues with strictly positive real part, exactly $N - 1$ eigenvalues with strictly negative real part, together with the zero eigenvalue. In particular, then, if we had diagonalized M , we would know which eigenvectors/values belong to Γ_+ ; we could then compute Γ_+ from the eigenvectors and eigenvalues, and likewise Γ_- .

To treat the delicate case $m(E_+) = m(E_-)$, we observe that M has a one-dimensional kernel, the left kernel being spanned by $(0, m)$, and that

$$(m, y)M = (0, m).$$

Here, m is the invariant law of Q , and y is such that

$$mV = yQ.$$

(Such a y exists, because $mv1 = 0$.) Thus the Jordan form of M contains a Jordan block

$$\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

(It is easy to see that the Jordan block of the zero eigenvalue must be 2×2 .) We conclude that in this case the Jordan form of M has exactly $N - 1$ eigenvalues in the open left half plane, $N - 1$ in the open right half plane (assuming the eigenvalues are distinct).

Finally, we present some results of the computations. We took the three Q-matrices we used in Section 2, and added noise $\epsilon = 10^{-1}$ in the first instance, and then $\epsilon = 10^{-3}$ in the second. The interest in taking ϵ small is of course that this should approximate the classical Wiener-Hopf situation studied in Section 2. Table 2 displays the times (in seconds) taken by the diagonalization method.

Table 2

(E_+ , E_-)	$\epsilon = 10^{-1}$	$\epsilon = 10^{-3}$
(10,15)	1.29	1.17
(48,50)	43.71	39.24
(100,100)	323.76	291.22

We also tried the recursive method (ii) above on an example of $(|E_+|, |E_-|) = (3, 100)$, with $\epsilon = 10^{-3}$, and α the minimal value such that all diagonal entries of $V + \epsilon^2\alpha$ is positive and all entries of $\alpha V + \epsilon^2\alpha^2 I/2 + Q$ is non-negative, and this was very slow - it took 3500s. Observe that the calculations in Table 2 took a lot longer than the classical Wiener-Hopf calculations of Section 2, of the order of 4.5 times as long. This is in fact better than expected, in that the heart of the calculation is diagonalization of a matrix, which is an $O(N^3)$ operation; thus when we diagonalize a matrix of double the order, in the noisy Wiener-Hopf, we would expect an eight fold increase in computing time.

We have thus calculated the matrices Π_+ in §2 and Γ_+ in §3. When ϵ is very small, and if we write

$$\Gamma_+ = \begin{pmatrix} \Gamma_{++}^{(E_+ \times E_+)} & \Gamma_{+-}^{(E_+ \times E_-)} \\ \Gamma_{-+}^{(E_- \times E_+)} & \Gamma_{--}^{(E_- \times E_-)} \end{pmatrix},$$

we would expect that the Γ_+ chain viewed in E_+ to look very much like the chain with generator $G_+ = A + B\Pi_+$. Therefore, for example, $-(\Gamma_{--})^{-1}\Gamma_{-+}$ should be like Π_+ . Indeed, we computed the matrix $H \equiv \Pi_+ + (\Gamma_{--})^{-1}\Gamma_{-+}$ (which should be approaching zero when ϵ tends to zero) for our example where $(|E_+|, |E_-|) = (100, 100)$ and $\epsilon = 10^{-3}$; it turned out that the sum of squares of all the entries in H is of order 10^{-11} . Our intuition is therefore justified.

REFERENCES

- [1] Anick, D., Mitra, D. & Sondhi, M.M.: Stochastic theory of a data-handling system with multiple sources. *Bell System Tech. J.* 61 (1982) 1871-1894.
- [2] Asmussen, S.: Markov-modulated reflected Brownian motion and fluid flow. (1992) *Preprint*.
- [3] Barlow, M.T., Rogers, L.C.G. & Williams, D.: Wiener-Hopf factorization for matrices. *Sém. Probab. XIV*, Lect. Notes Math. vol.784 (Springer 1980) 324-331.
- [4] Gaver, D.P. & Lehoczky, J.P.: Performance evaluation of voice/data queueing systems. In: *Applied Probability - Computer Science: The Interface* (Eds.: R.L. Disney & T.J. Ott), vol.I (1982) 329-346.
- [5] Gaver, D.P. & Lehoczky, J.P.: Channels that cooperatively service a data stream and voice messages. *IEEE Trans. Comm.* 30 (1982) 1153-1161.
- [6] Kennedy, J. & Williams, D.: Probabilistic factorization of a quadratic matrix polynomial. *Math. Proc. Camb. Phil. Soc.* 107 (1990) 591-600.
- [7] Mitra, D.: Stochastic theory of a fluid model of producers and customers coupled by a buffer. *Adv. Appl. Probab.* 20 (1988) 646-676.
- [8] Rogers, L.C.G.: Fluid models in queueing theory and Wiener-Hopf factorization of Markov chains. (1992) *Preprint*.
- [9] Stern, T.E. & Elwalid, A.I.: Analysis of separable Markov-modulated rate models for information-handling systems. *Adv. Appl. Probab.* 23 (1991) 105-139.
- [10] Williams, D.: A “potential-theoretic” note on the quadratic Wiener-Hopf equation for Q-matrices. *Sém. Probab. XVI*, Lect. Notes Math. vol.920 (Springer 1982) 91-94.

L.C.G. Rogers & Z. Shi
School of Mathematical Sciences
Queen Mary & Westfield College
Mile End Road
London E1 4NS